Impact of Human Mobility on Opportunistic Forwarding Algorithms

Augustin Chaintreau, Pan Hui, Jon Crowcroft, *Fellow*, *IEEE*,
Christophe Diot, Richard Gass, and James Scott

Abstract—We study data transfer opportunities between wireless devices carried by humans. We observe that the distribution of the intercontact time (the time gap separating two contacts between the same pair of devices) may be well approximated by a power law over the range [10 minutes; 1 day]. This observation is confirmed using eight distinct experimental data sets. It is at odds with the exponential decay implied by the most commonly used mobility models. In this paper, we study how this newly uncovered characteristic of human mobility impacts one class of forwarding algorithms previously proposed. We use a simplified model based on the renewal theory to study how the parameters of the distribution impact the performance in terms of the delivery delay of these algorithms. We make recommendations for the design of well-founded opportunistic forwarding algorithms in the context of human-carried devices.

Index Terms—Computer systems organization, communication/networking and information technology, mobile computing, algorithm/protocol design and analysis, mobile environments, mathematics of computing, probability and statistics.

1 Introduction

The increasing popularity of devices equipped with wireless network interfaces (such as cell phones or PDAs) offers new communication services opportunities. Such mobile devices can transfer data in two ways—by transmitting over a wireless (or wired) network interface and by taking advantage of the user's mobility. They form a P*o*cket Switched Network [1]. Communication services that rely on this type of data transfer will strongly depend on human mobility characteristics and on how often such transfer opportunities arise. Therefore, they will require networking protocols that are different from those used on the Internet. Since two (or more) ends of the communication might not be connected simultaneously, it is impossible to maintain routes or to access centralized services such as the DNS.

In order to better understand the constraints of opportunistic data transfer, we analyze eight distinct data sets collected in networks with mobile devices. Three data sets come from experiments we conducted ourselves. We define the intercontact time as the time between two transfer opportunities for the same devices. We observe in the eight traces that the intercontact time tail distribution is slowly varying over a large range. Inside this range, the intercontact time distribution can be compared to a power law. We study the impact of those large intercontact times on the actual performance and theoretical limits of a general class of opportunistic forwarding algorithms that we call "oblivious forwarding algorithms." Algorithms in this class do not use the identities of the devices that are met, nor the recent history of contacts, nor the time of day in order to make forwarding decisions. Instead, forwarding decisions are based on statically defined forwarding rules that bound the number of data replicas or the number of hops.

Based on our experimental observations, we develop a simplified model of opportunistic contact between human-carried wireless devices. Our model makes several independence assumptions which are common in the literature of mobile ad hoc routing. We do not claim that this model captures the performance of different forwarding algorithms accurately. Rather, it serves our purpose, which is to demonstrate how the tail of intercontact times influences the performance of oblivious forwarding algorithms and how these should be configured to offer reasonable guarantees.

Experimental results are presented in Section 2. In Section 3, we model contact opportunities based on our observations and we analyze the delay that wireless devices would experience using a class of forwarding algorithms previously studied in the literature. Section 4 is dedicated to related work. The paper concludes with a brief summary of contributions and a presentation of future work, including a discussion of the implications of our assumptions.

2 Experimental Results

2.1 Data Sets

In order to conduct informed design of opportunistic forwarding algorithms, it is important to analyze the frequency and duration of contacts between human-carried communicating devices. Ideally, an experiment would cover a large user base over a large time period, as well as include data on connection opportunities encountered 24 hours a day.

- *A. Chaintreau and C. Diot are with Thomson Research, 46 quai A. Le Gallo, 92648 Boulogne, France. E-mail: {augustin.chaintreau, christophe.diot}@thomson.net.*
- *P. Hui and J. Crowcroft are with the University of Cambridge, 15 J J Thomson Av., Cambridge, CB3 0FD. E-mail: {pan.hui, jon.crowcroft}@cl.cam.ac.uk.*
- *J. Scott is with Microsoft Research, 7 J J Thomson Av., CB3 0FB, UK. E-mail: james.scott@microsoft.com.*
- *R. Gass is with Intel Research, 4720 Forbes Avenue, Pittsburgh, PA 15213. E-mail: richard.gass@intel.com.*

Manuscript received 1 June 2006; revised 20 June 2006; accepted 16 Mar. 2007; published online 28 Mar. 2007.
For information on obtaining reprints of this article, please send e-mail to: tmc@computer.org, and reference IEEECS Log Number TMCSI-0151-0606.
Digital Object Identifier no. 10.1109/TMC.2007.1060.

1536-1233/07/$25.00 © 2007 IEEE Published by the IEEE CS, CASS, ComSoc, IES, & SPS

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

TABLE 1
Comparison of Data Collected in the Eight Experiments

User Population	Cambridge	Hong Kong	Infocom	Toronto	UCSD	Dartmouth	MIT BT	MIT GSM
Device	iMote	iMote	iMote	PDA	PDA	Laptop/PDA	Cell Phone	Cell Phone
Network type	Bluetooth	Bluetooth	Bluetooth	Bluetooth	WiFi	WiFi	Bluetooth	GSM
Contact type	direct	direct	direct	direct	AP-based	AP-based	direct	AP-based
Duration (days)	3	5	4	16	77	114	246	246
Granularity (seconds)	120	120	120	120	120	300	300	10
Devices participating	12	37	41	23	273	6,648	100	100
# of internal contacts	4,229	280	22,459	2,802	195,364	4,058,284	54,667	572,190
# of internal contacts/pair/day	10	0.042	3.4	0.35	0.034	0.00080	0.022	0.23
Recorded external devices	148	1,840	197	N/A	N/A	N/A	N/A	N/A
# of external contacts	2,441	6,881	5,723	N/A	N/A	N/A	N/A	N/A

We examined two types of data sets. First, we use publicly available traces measuring connectivity between clients and access points (APs) in several wireless networks (using WiFi or GSM technology); contacts between the clients were deduced from the traces following an assumption that we discuss below. Second, we collected our own traces of direct contacts recorded using small portable wireless radio devices (iMotes) that were distributed to different groups of people. We found a few other traces of direct contacts and we have included them for comparison with ours. In total, there are eight data sets and each of their characteristics are summarized in Table 1.

2.1.1 AP-Based Data Sets

The University of California, San Diego (UCSD) [2] and Dartmouth University [3] traces make use of WiFi networking, with the former including client-based logs of the visibility of access points (APs), while the latter includes SNMP logs from the access points. The durations of the different logs traces are three and four months, respectively. Since we required data about device-to-device transmission opportunities, the raw data sets were unsuitable for our experiment and required preprocessing. For both data sets, we made the assumption that mobile devices seeing the same AP would also be able to communicate directly (in ad hoc mode). Consequently, a list of transmission opportunities was deduced for each pair of devices, which corresponds to the time intervals for which they share at least one AP.

The traces from the Reality Mining project [4] at the Massachusetts Institute of Technology (MIT) Media Lab include records of visible GSM cell towers, collected by 100 cell phones distributed to students and faculty on the campus during 9 months. We have assumed, as above, that two devices are in contact whenever they are connected with the same cell tower.

Unfortunately, the assumptions we have made for all these data sets introduce inaccuracies. On the one hand, it is overly optimistic since two devices attached to the same (WiFi or GSM) base station may still be out of range of each other. On the other hand, the data might omit connection opportunities, such as when two devices pass each other at a place where there is no instrumented access point. Another potential issue with these data sets is that the devices are not necessarily colocated with their owners at all times. Despite these inaccuracies, these traces are a valuable source of data spanning many months and including thousands of devices. In addition, considering two devices connected to the same base station as being potentially in contact is not altogether unreasonable. These devices may indeed be able to communicate locally through the base station.

2.1.2 Direct Contact Data Sets

In order to complement the previous traces, we did our own experiment using Intel iMotes, which are embedded devices similar to Crossbow motes,[1] except that they communicate via Bluetooth. We programmed the iMotes to log contact data every 120 s for all visible Bluetooth devices (including iMotes as well as other Bluetooth devices such as cell phones). Each contact is represented by a tuple (MAC address, start time, and end time). The experimental settings are described in detail in [1]; an anonymized version of our data is now available to other research groups in the CRAWDAD[2] archive.

We include in this paper the results from three iMote-based experiments. We first obtained data from 12 doctoral students and faculty comprising a research group at the University of Cambridge. The second experiment included a group of 37 participants in Hong Kong selected in such a way that they do not belong to the same work or social group and, in particular, that none of them knows each other. The third experiment was conducted during the IEEE INFOCOM 2005 conference in Miami, where iMotes were carried by 41 attendees for 4 days. The contacts collected by iMotes are classified into two groups: the sighting of another iMote is classified as an "internal" contact, while the sighting of other types of Bluetooth devices is called an "external" contact. The external contacts are numerous and they provide a measure of the wireless networking opportunities present at that time. Internal contacts, on the other hand, represent the data transfer opportunities among participants if they were all equipped with devices which are always-on and always-carried.

In addition to our own experiment, we found two data sets with direct contacts and included them for comparison: A research group from the University of Toronto has collected direct contact traces using 23 Bluetooth-enabled PDAs distributed to a group of students. These devices performed a Bluetooth inquiry every 120 seconds and this data was logged. This methodology does not require devices to be in range of any AP in order to collect contacts, but it does require that the PDAs are carried by the

1. See www.xbow.com.
2. See crawdad.cs.dartmouth.edu.

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

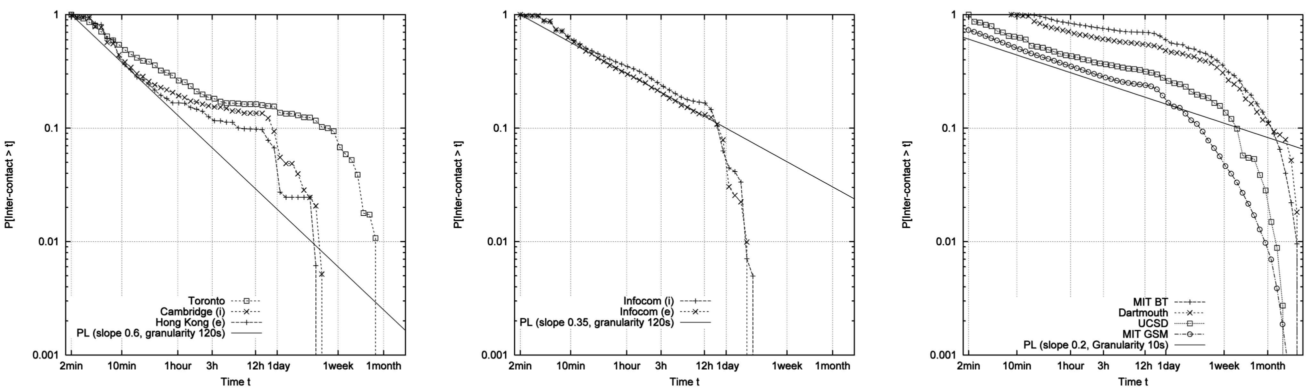

(a) (b) (c)

Fig. 1. Aggregated distribution of the intercontact time in eight data set experiments: (a) iMote-based experiments at Cambridge and Hong Kong, and the Toronto experiment, (b) iMote-based experiment at INFOCOM, and (c) data collected at UCSD, Dartmouth, and MIT.

participants and that the participants keep them charged. The data set we use comes from an experiment that lasted 16 days. The traces from the Reality Mining project [4] include direct Bluetooth sightings, recorded every 300 seconds by each participant's cell phone.

2.2 Definitions

We are interested in how the characteristics of transfer opportunities impact data forwarding decisions. In this paper, we focus on how often such opportunities occur, but not in their duration. We decided not to analyze how much data can be transported during a transfer opportunity because this strongly depends on the wireless technology used. Later in our analysis (see Section 3), we will assume that all contacts last a single time slot and we will address two extreme cases corresponding to a lower and upper bounds of the amount of data that could be transferred in each connection opportunity.

We define the *intercontact time* as the time elapsed between two successive contact periods for a given pair of devices. Intercontact time characterizes the frequency with which data can be transferred between networked devices; it has rarely been studied in the literature. Two remarks must be made with regard to this definition: First, the intercontact time is computed once at the end of each contact period, at the time interval between the end of this contact and the beginning of the next contact with the same device.[3] An alternative option would be to compute the *remaining intercontact time* seen at any time t: for each pair of devices, it is the time it takes after t before these devices meet again (a formal definition is given in Section 3). Intercontact time and remaining intercontact time have different distributions, which are related, for a renewal process, via a classical result known as the waiting time paradox (see [5, p. 147]). A similar relation holds for stationary processes in the theory of Palm Calculus (see [6, p. 15]). We choose to study the first definition of "intercontact time seen at the end of a contact period" as the second gives too much weight to large intercontact times. In other words, the definition we have chosen is the most conservative one in the presence of large values.

3. Note that we did not include the time between the beginning of the experiment and the first contact for a pair, nor the time between the last contact of a pair and the end of the experiment.

Second, the intercontact time distribution is influenced by the experiment's duration and its granularity (i.e., the time elapsed between two successive scannings for the same device). Intercontact times that last more than the duration of the experiment cannot be observed, and intercontact times close to the duration are less likely to be observed. In a similar way, we cannot observe the intercontact times that last less than the granularity of measurement (which ranges from two to five minutes for different experiments).

Another measure of the frequency of transfer opportunities that could be considered is the *inter-any-contact time*, i.e., for a given device, the time elapsed between two successive contacts with any other device. This measure is very much dependent on the density of wireless devices during the experiment as it characterizes time that devices spend without meeting any other device. This measure was studied for most of these data sets in [1]. We do not present further results here.

2.3 Intercontact Time Characterization

We study the empirical distribution of the intercontact times obtained for all experiments shown in Fig. 1 and Fig. 2.

For all plots, an empirical distribution of the intercontact times was first computed separately for each pair of devices that met at least twice. It is hard to study the characteristics of the distributions for all pairs individually because there are many such distributions and some of them may only include a few observed values. This is why we follow a two-step approach: First, we present the distribution obtained when all pairs' distributions are combined, each with an equal weight, in a distribution that we call the *aggregated distribution*. Second, we use a parametric model motivated by this first part and estimate the parameter of the *individual distribution* for each pair.

2.3.1 Aggregated Distribution

Fig. 1 presents the aggregated distribution for different data sets. All plots show the complementary cumulative distribution function using a log-log scale.

For iMote experiments, "(i)" indicates that the data set shown is obtained using internal contacts only, while "(e)" indicates that the data set shown includes only external contacts. For the first two iMote experiments (labeled

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

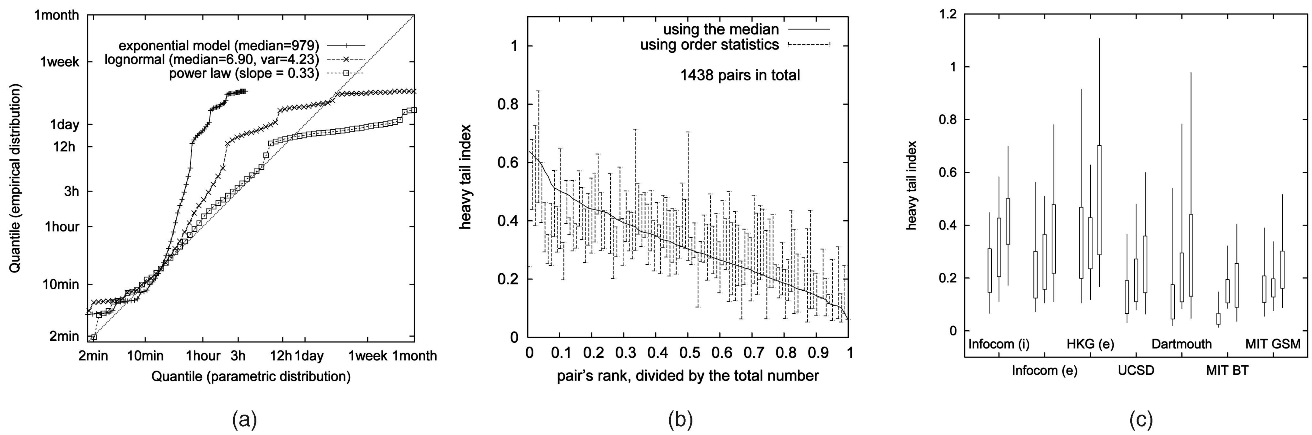

Fig. 2. iMote-based experiment at INFOCOM: (a) Quantile-quantile plot of a comparison between the aggregated distribution of the intercontact time and three parametric models, (b) estimation of the heavy tail index of the power law applied separately for each pair, and (c) summary of the results obtained in all data sets.

Cambridge and Hong Kong), we present only one case here (corresponding, respectively, to internal and external contacts). They are shown in Fig. 1a, which also includes the distribution obtained among pairs of experimental devices in the trace from the University of Toronto. Distributions belonging to the iMote-based experiment at INFOCOM are shown in Fig. 1b, where distributions associated with internal and external contacts have been plotted separately for comparison. Fig. 1c presents the distribution of intercontact time computed using traces from other experiments than ours.

Let us first note that, although intercontact times are short in most cases, the occurrence of large intercontact times is far from negligible: In the three iMote-based experiments, 17 to 30 percent of intercontact times are greater than one hour, and 3 to 7 percent of all intercontact times are greater than one day. In the Toronto data sets, 14 percent of intercontact times last more than a day and 8 percent more than a week. These large intercontact times are even more frequent in the traces collected at UCSD, Dartmouth, and MIT, the most extreme case being the MIT trace using Bluetooth sightings, where up to 60 percent of the intercontact times observed are above 1 day. The variation between data sets is significant. It can be expected given the diversity of communication technologies and population studied, as well as the impact of experimental conditions (granularity, duration). But, there are also common properties that we now discuss in more detail.

We now concentrate on the region between 10 minutes and 1 day. In this region, all data sets exhibit the same characteristics: The CCDF is slowly varying and it is lower bounded by the CCDF of a power law distribution, which may, in some cases, be a close approximation. This contradicts the exponential decay of the tail, which characterizes the most common mobility models found in the literature (see Section 4), and we prove in the next section that this can have a significant impact on the performance of opportunistic networking algorithms.

To justify the above claim, we studied the quantile-quantile plot comparison between the empirical distribution found and three parametric models (exponential, log-normal, and power law). An example is shown in Fig. 2a for the distribution based on internal contacts during the INFOCOM experiment. All parametric models have been set to take the same median value as the empirical distribution. We also normalize the power law to fit the granularity $t = 120$ seconds and the log-normal distribution such that the logarithm of both the empirical variable and the model have the same variance. Not surprisingly, we observe that the three models deviate significantly from the empirical findings for values above one day. As expected, the exponential distribution is very far from the empirical one and the quantile for the log-normal distribution deviates from the empirical case by a nonnegligible factor. The power law distribution, in contrast, remains close to the empirical one for values of up to 18 hours, and it seems to be the most appropriate model to apply. In other data sets, the power law may sometimes not match the empirical findings, as well as in this example, but among these three models, it is always the one closest to the empirical distribution. For values above one day, we expect models with additional parameters (e.g., following a Weibull distribution) to improve the match with the empirical distribution, but that is beyond the scope of this paper.

The most notable difference we observe between data sets is that the fit with a power law is better for the data sets that contain the largest number of points, such as in Fig. 1b and Fig. 1c. We also observe that the slope of the power law that is a lower bound on the range [10 minutes; 1 day] is different between data sets: This is 0.6 for the iMote experiments at Cambridge and Hong Kong, as well as for the Toronto data set, 0.35 for the iMote-based experiment at INFOCOM, and 0.2 for the traces collected at UCSD, Dartmouth, and MIT. In all cases, it is below 1. The value of this slope, which is also called the "heavy tail index," is critical for the performance of opportunistic forwarding algorithms (see the analysis in Section 3), and we discuss it further below.

Fig. 1b shows that the distribution is almost unchanged if one considers internal or external contacts. The same observation was made for other iMote experiments [1], except for the experiment conducted in Hong Kong, where, as expected, very few internal contacts were logged. Some variations of the heavy tail index have been observed depending on the time of the day [1].

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

2.3.2 Individual Distributions for Each Pair

So far, we have studied the aggregated distribution where all pairs have been combined together, and we found that it can be approximated by a power law for values up to 1 day. In this section, we assume that this claim can be made individually for all pairs, although the parameter of this power law, also called the heavy tail index, may be different among them. This approach allows us to study the heterogeneity between pairs via a single parameter; some of these results also measure the accuracy of the above assumption for each pair.

Estimator for the Heavy Tail Index. Let us consider a pair of nodes. The sample of the intercontact times observed for this pair will be denoted by $X_1, \ldots, X_n$, its order statistics by $X_{(1)} \leq \ldots \leq X_{(n)}$, and its median value by $\mathtt{m}$. All times will be given in seconds. If we assume that this sample follows a power law with granularity 120 s and heavy tail index α, we have $\mathbb{P}[X \geq x] = (x/120)^{-\alpha}$, such that an estimator of α based on the sample's median $\mathtt{m}$ is given by

$$\breve{\alpha} = \frac{\ln(2)}{\ln(\mathtt{m}) - \ln(120)}.$$

More generally, one can consider all order statistics $X_{(i)}$ that fit in the range [10 minutes; 1 day] and estimate α based on each of them. This creates a collection of estimators for the value of α as follows:

$$\left\{ \frac{\ln(n) - \ln(n-i)}{\ln(X_{(i)}) - \ln(120)} \;\middle|\; 600 \leq X_{(i)} \leq 86{,}400,\ i < n \right\}.$$

We denote by $\underline{\alpha}$ and $\overline{\alpha}$, respectively, the minimum and maximum value in this set above. This is equivalent to plotting the empirical CCDF for this sample in a log-log scale and bounding this CCDF from above and below by two straight lines that go through probability 1 at time value 120 s. The slopes of these lines would be equal, respectively, to $-\underline{\alpha}$ and $-\overline{\alpha}$. In contrast to $\breve{\alpha}$, these two estimators are not centered around the value of α, and they do not converge to this value when the sample becomes large. They rather serve the purpose of a heuristic analysis; they characterize some bounds that are verified by each pair. Note also that, intuitively, the difference $\overline{\alpha} - \underline{\alpha}$ indicates how the conditional distribution of the sample in this range differs from a pure power law.

In Fig. 2b, we plot the values of $\breve{\alpha}$ and the interval $[\underline{\alpha}, \overline{\alpha}]$ for all pairs of iMotes during the experiment conducted at INFOCOM. One can expect that the heavy tail index takes different values among pairs, as some participants are more likely to meet often than others. We initially ranked all pairs according to their value for $\breve{\alpha}$ in the decreasing order. Although we have computed these values for all pairs, we only draw the interval $[\underline{\alpha}; \overline{\alpha}]$ for 100 pairs chosen arbitrarily according to their rank (one every 14) in order to keep the figure readable. As shown in Fig. 2b, estimations of α for different pairs may indeed vary between 0.05 and 1. Between these two extreme values, which are very rarely observed, estimates for almost all pairs lie between 0.1 and 0.7 depending on the estimator. Note that all estimates of α are smaller than 1; the only exceptions are the upper estimate $\overline{\alpha}$ for three pairs (i.e., less than 0.2 percent of pairs in this case). The median-based estimate lies in [0.2 ; 0.4] for half of the pairs and the lower estimates (respectively, the upper estimate) lies in [0.14 ; 0.32] (respectively, [0.32 ; 0.5]) again for half of the pairs.

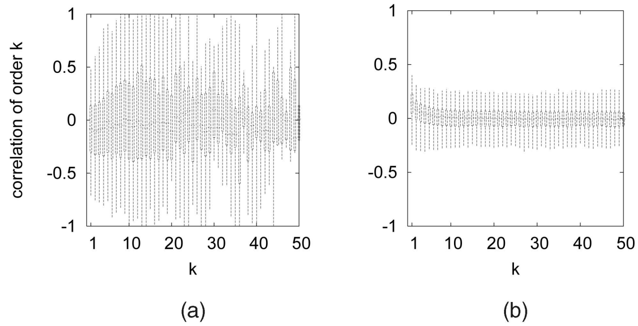

Fig. 3. Correlation coefficients for the sequence of intercontact times: (a) for all pairs of iMotes in the INFOCOM data set and (b) for all pairs of devices in the MIT GSM data set.

These results have three major implications: First, the heterogeneity among pairs implies different possible values for α, which are centered around the value already observed when studying the aggregate distribution (i.e., 0.33). Second, the difference between the median estimator and the heuristic bounds we defined above remains within 0.25 except in a few cases. Last, the upper estimate $\overline{\alpha}$ almost never goes above 1, which establishes that the intercontact time distribution for each pair is lower bounded in this range by a power law with a heavy tail index smaller than 1.

The same results have been obtained for other data sets, and they are summarized in Fig. 2c. For each data set indicated, we show the distribution of values obtained among pairs for the three estimators defined above. Each estimator stands for one box-plot: From left to right, $\underline{\alpha}, \breve{\alpha}, \overline{\alpha}$; the thick part indicates the values found in 50 percent of the pairs and the thin part contains the region where 90 percent of the pairs are found.

In the Hong Kong and Dartmouth data sets, where contacts are sparser, intercontact time samples for each pair contains fewer values. As a consequence, the difference between estimators can grow significantly. We even observe that $\overline{\alpha}$ goes slightly beyond 1 for 10 percent of the pairs in Hong Kong data set, although this could be an artifact of our conservative estimate.

Correlation. We study the autocorrelation coefficients to see how the value of the intercontact time may depend on the previous values for the same pair. The results are shown in Fig. 3 for all order k up to 50. Since the intercontact time distribution usually has no finite variance, we computed the correlation coefficient on the values of the logarithm of the intercontact times. Note that, since a correlation coefficient was computed for each pair, we present for all orders k the average value we observed among all pairs, as well as the interval containing 50 percent and 90 percent of the centered values (respectively, in the thick box and the thin bar).

In the INFOCOM data set, the variation of the coefficient among pairs is quite important, although most pairs remain reasonably noncorrelated (the thick box always remains less than 0.30 away from zero). Overall, we observe a slightly negative correlation over all pairs on average, which reduces as k grows. Correlation coefficients are smaller when the data set is large (as seen, for example, in the MIT

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

GSM trace shown here, as well as for all other long traces). This tends to indicate that these coefficient pairs would be closer to zero if the iMote experiment could be done with a longer duration, and that the sample of intercontact times collected for each pair was bigger.

Based on the above results, we assume in the next section that the intercontact time distribution follows a power law for each pair. To simplify the analysis, we assume, in addition, that the coefficient is the same for all pairs, that the sequence of intercontact times is i.i.d. (i.e., correlation coefficients are null), and that they are independent between pairs. This simplification allows us to characterize the performance of forwarding algorithms quite generally. Some of the results we present can be extended to stationary ergodic sequences or correlation between pairs, but that is left for future work.

3 Forwarding with Power Law-Based Opportunities

We now analyze the impact of our findings on the performance of a class of forwarding algorithms. We first define our abstract model of the opportunistic behavior of mobile users based on our experimental observations.

3.1 Assumptions and Forwarding Algorithms

3.1.1 Contact Process Model

We consider a slotted time $t = 0, 1, \ldots$. For a given pair of devices (d, d'), let us introduce its contact process $(U_t^{(d,d')})_{t\geq 0}$ defined by

$$U_t^{(d,d')} = \begin{cases} 1 & \text{if } d \text{ and } d' \text{ are in contact during slot } t, \\ 0 & \text{otherwise.} \end{cases}$$

For the pair (d, d'), we consider the sequence of the time slots $T_0^{(d,d')} < T_1^{(d,d')} < \ldots < T_k^{(d,d')} < \ldots$ that describes all the values of $t \in \mathbb{N}$ such that $U_t^{(d,d')} = 1$.

We do not include in this model the contact time representing the duration of each contact, assuming that each contact starts and ends during the same time slot. This is justified here by the fact that we are interested in a model accounting for consequences of large values of the intercontact time. It was observed (see [1]) that the contact time distribution may also be approximated by a power law, but over a range that is much smaller than the range for the intercontact time.

Under this condition, the time $\tau_k^{(d,d')} = T_{k+1}^{(d,d')} - T_k^{(d,d')}$ for any d, d' and $k \geq 0$ is the intercontact time after the kth contact of this pair. We suppose in our model that it has the same law as X, which follows a power law with heavy tail index $\alpha > 0$:

$$P[X \geq t] = t^{-\alpha} \text{ for all } t = 1, 2, \ldots. \tag{1}$$

Note that X is not bounded but is finite almost surely. It may easily be seen that X has a finite mean if and only if $\alpha > 1$.

In addition, we assume that the contact process $(U_t^{(d,d')})_{t\geq 0}$ of each node pair (d, d') is a renewal process and that contact processes associated with different pairs are independent. In other words, the intercontact times in the sequence $(\tau_k^{(d,d')})_{k\geq 0}$ are i.i.d., for all (d, d'), and sequences belonging to different pairs are independent.

We come back to these assumptions later in Section 5. Note that these assumptions are shared explicitly or implicitly by most of the analyses of currently proposed mobility models. This is because it is typically very difficult to analyze models where dependence may arise between different devices or between successive events occurring with one or more devices.

Even if we do not explicitly model the contact time (each contact lasts one time slot), we need to take into consideration the fact that a contact may last long enough to transmit a significant amount of data. We then introduce two situations:

- the *short contact case*, where only one data unit can be sent between the two devices during each contact, and
- the *long contact case*, where we assume that all queues in the two devices can be completely emptied during each contact.

These two cases represent lower and upper bounds for the evaluation of bandwidth. The number of data units transmitted in a contact (whether short or long) is defined as a data bundle.[4] The long and the short case differ from a queuing standpoint. In the long contact case, as soon as a data unit has arrived in a node, it can be sent to all other nodes that are met. In the short contact case, only one data unit is sent at once and, therefore, data can accumulate in the memory of a relaying device.

Note that our model does not take into account explicit geographical locations or movement of devices; rather, it directly describes the processes of contacts between devices. The results of this section extend to any mobility model which creates independent contact processes for all pairs of devices that follow this same law.

For any pair of devices (d, d'), let us introduce the **remaining intercontact time** observed at time slot t: It is an integer denoted by $R_t^{(d,d')}$ and defined as

$$R_t^{(d,d')} = \min\left\{ t' - t \;\middle|\; t' \geq t \text{ and } U_{t'}^{(d,d')} = 1 \right\}.$$

As the contacts for each pair are supposed to follow a renewal process, $R_t^{(d,d')}$ is a homogeneous Markov Chain. As shown in Appendix A, it is recurrent and ergodic if and only if $\alpha > 1$.

3.1.2 Forwarding Algorithms

We are interested in a general class of forwarding algorithms, which all rely on other devices to act as relays, carrying data between a source device and a destination device that might not be contemporaneously connected. These relay devices are chosen purely based on contact opportunism and not using any stored information that describes the current state of the network. The only information used in forwarding is the identity of the destination so that a device knows when it meets the destination for a bundle. We call such algorithms "oblivious" as they could be in reality quite complex and, as we will see, very efficient in some cases.

4. In DTN standards, a *bundle* usually denotes a large object with a collection of data units.

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

The following two algorithms provide bounds for the class of algorithm described above:

- Wait-and-forward: The source waits until its next direct contact with the destination to communicate.
- Flooding: A device forwards all its received data to any device which it encounters, keeping a copy for itself.

The first algorithm uses minimal resources but can incur very long delays and does not take full advantage of the ad hoc network capacity. The second algorithm, initially proposed in [7], delivers data with the minimum possible latency but does not scale well in terms of bandwidth, storage, and battery usage. In between these two extreme algorithms, there is a whole range of algorithms that differ in the number of relays used to maximize the chance of reaching the destination with a delay as small as possible while avoiding flooding. The most important reasons not to flood are to minimize memory requirements and related power consumption in relay devices and to delete the backlog of previously relayed messages that are still waiting to be delivered and could be outdated. A number of strategies, based on time-outs, buffer management, limit on the number of hops, and/or duplicate copies, have been proposed [7], [8], [9] to minimize replication and backlog.

3.2 Analysis of the Two-Hop Relaying Algorithm

Having described the class of "oblivious" algorithms we are considering in this work, we now introduce the two-hop relaying algorithm [10] and evaluate its performance for the model of power law intercontact times that we have described. Results are generalized to all oblivious algorithms in the following section.

3.2.1 Description

The two-hop relaying algorithm was introduced by Grossglauser and Tse in [10]. This forwarding algorithm operates as follows: When a source has a bundle to send to a destination, it forwards it once to the first device that it meets. This first device is either a relay device or the destination itself. If it is the destination, the bundle is delivered in one hop; otherwise, the device acts as a relay and stores the bundle in a queue corresponding to this destination. Bundles from this queue will be delivered when the relay device meets the destination. Bundles for the same destination are delivered by a relay device in a first-come-first-served order. As queuing may occur in the devices that act as relays, in the short contact case, the forwarding process of bundles sent by the source to a relay needs to be of lower intensity than the bundles sent by this relay to the destination. This is the case in the implementation proposed in [10] and we make the same assumption below.

We choose this algorithm to start our study of the impact of power law intercontact times on opportunistic forwarding for the following three reasons:

- In the short contact case, this algorithm was shown [10] to maximize the capacity of dense mobile ad hoc networks under the condition that devices' locations are i.i.d., distributed uniformly in a bounded region.
- The mobility process of the devices is an important parameter. Grossglauser and Tse [10] assumed that each node moves at each time slot to an i.i.d. position, which implies that the intercontact time is geometrically distributed. This result holds more generally for any Markovian evolution (such as a random walk) defined on a finite domain. It has also been shown when devices move according to the random way-point mobility model (see the analysis of Section 3 in [11]).
- Grossglauser and Tse [10] and Sharma and Mazumdar [11] have shown that data experiences a finite expected delay under these conditions.

3.2.2 Analysis

We consider N mobile devices which transmit data according to the two-hop relaying algorithm described above. Instead of the mobility model used in [10], we assume that contacts between devices follow the model that we have introduced in the beginning of this section.

To ensure stability in the relay's queuing mechanism, we assume that the source s is not saturated: Bundles are created at s during a sequence of time slots, denoted by $(t_k^{(s)})_{k\in\mathbb{Z}}$. The same assumption is made for the long contact case, although stability of the queue occupancy is not an issue in this context as the queue is emptied after each contact with the destination.

We have the following result, which is a consequence from the regenerative theorem (or Smith's formula):

Theorem 1. *For a pair of source-destination devices (s, d), let $t_k^{(s)}$ be the time when the kth bundle is created at s to be sent to d and let $t_k^{(d)}$ be the time when it is delivered to d. Letting $D_k = t_k^{(d)} - t_k^{(s)}$, we have, starting from any initial condition:*

1. *If $\alpha < 2$, $\lim_{k\to\infty} \mathbb{E}[D_k] = +\infty$.*
2. *If $\alpha > 2$ and we assume that all contacts are long, $\lim_{k\to\infty} \mathbb{E}[D_k] = \bar{D} < +\infty$ and we have*

$$\bar{R} \leq \bar{D} \leq 2\bar{R}, \text{ where } \bar{R} = \frac{1}{2} + \frac{\mathbb{E}[X^2]}{2 \cdot \mathbb{E}[X]}.$$

3. *If $\alpha > 2$ and we assume that all contacts are short, when each source sends data to a unique and distinct destination with rate $\lambda < \frac{N-1}{2\mathbb{E}[X]}$, then the delay of a bundle has finite expectation.*

Proof. We study first the case of long contacts, where any amount of information may be exchanged when a contact occurs between two devices.

We analyzed here a single source-destination pair. The two-hop relaying strategy uses multiple routes to transport bundles belonging to this pair because any other contacted device may act as a relay. This bundle is transmitted to the first relay that is met by s after time $t_k^{(s)}$. Letting r_k be this relay, we have $r_k = \mathrm{argmin}_{r'\neq s} R^{(s,r')}_{t_k^{(s)}}$, and this transmission occurs at time $t_k^{(r)} = t_k^{(s)} + \min_{r'\neq s} R^{(s,r')}_{t_k^{(s)}}$. The bundle is then delivered to destination d at time ;$t_k^{(d)} = t_k^{(r)} + R^{(r_k,d)}_{t_k^{(r)}}$. We can rewrite

$$D_k = t_k^{(d)} - t_k^{(s)} = \min_{r\neq s} R^{(s,r)}_{t_k^{(s)}} + R^{(r_k,d)}_{t_k^{(r)}}. \tag{2}$$

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

Let us first establish the positive result from Theorem 1.2 that the two-hop relaying strategy achieves a delay with finite mean if $\alpha > 2$.

Proving Theorem 1.2. In this case, $\mathbb{E}[X^2]$ is finite and

$$\mathbb{E}\left[\sum_{t=T_0^{(d,d')}}^{T_1^{(d,d')}-1} R_t^{(d,d')}\right] = \mathbb{E}[X(X+1)/2] < \infty,$$

for any pair (d, d') of devices. By Smith's formula (see (5) in the Appendix), we have

$$\lim_{t\to\infty}\mathbb{E}\left[R_t^{(d,d')}\right] = \frac{\mathbb{E}[X^2] + \mathbb{E}[X]}{2\mathbb{E}[X]}.$$

The process $(\min_{r\neq s} R_t^{(s,r)})_{t\geq 0}$ is taken as a minimum of a finite number of independent processes, corresponding to pairs $\{(s,r) \,|\, r \neq s\}$, which all have the same law.

Hence,

$$\lim_{t\to\infty}\mathbb{E}\left[\min_{r\neq s} R_t^{(s,r)}\right] \leq \frac{\mathbb{E}[X^2] + \mathbb{E}[X]}{2\mathbb{E}[X]}.$$

Lemma 2 can then be applied to this process, with $(t_k^{(s)})_{k\geq 0}$, which is independent from it; this proves

$$\lim_{k\to\infty}\mathbb{E}\left[\min_{r\neq s} R_{t_k^{(s)}}^{(s,r)}\right] \leq \frac{\mathbb{E}[X^2] + \mathbb{E}[X]}{2\mathbb{E}[X]}.$$

If we consider the collection of random variables $((R_t^{(r,d)})_{t\geq 0})_{r\neq s}$, the condition in Lemma 2.1 is met. As $(t_k^{(r)})_{k\geq 0}$ and $(r_k)_{k\geq 0}$ depend only on $(t_k^{(s)})_{k\geq 0}$ and contacts processes belonging to other pairs than $\{(r,d) \,|\, r \neq s\}$, they are independent from the collection above, and we have

$$\lim_{k\to\infty}\mathbb{E}\left[R_{t_k^{(r)}}^{(r_k,d)}\right] = \left(\frac{1 + \mathbb{E}[X^2]}{2\mathbb{E}[X]}\right).$$

Using (2), we have

$$R_{t_k^{(r)}}^{(r_k,d)} \leq D_k = \min_{r\neq s} R_{t_k^{(s)}}^{(s,r)} + R_{t_k^{(r)}}^{(r_k,d)};$$

hence,

$$\frac{1}{2}\left(1 + \frac{\mathbb{E}[X^2]}{\mathbb{E}[X]}\right) \leq \lim_{k\to\infty}\mathbb{E}[D_k] \leq \left(1 + \frac{\mathbb{E}[X^2]}{\mathbb{E}[X]}\right).$$

Note that this result holds if the law of X is replaced by any law that admits a finite second moment.

Proving Theorem 1.1 for $1 < \alpha < 2$. As $\alpha > 1$, Smith's Formula (5) holds in this case for any function f verifying the integrability condition.

Let r denote any device different from s. For convenience, let us denote $X_1 = T_1^{(r,d)} - T_0^{(r,d)}$. Then, we have, for any A that may be chosen arbitrarily large,

$$\frac{A(A+1)}{2}\mathbb{I}_{\{X_1\geq A\}} \leq \sum_{t=T_0^{(r,d)}}^{T_1^{(r,d)}-1} \min(R_t^{(r,d)}, A) \leq A \cdot X_1.$$

These variables are positive; they all have a finite expectation by comparison with the right term. This proves the integrability condition required in (5) for the function $f(x) = \min(x, A)$; hence, we obtain

$$\lim_{t\to\infty}\mathbb{E}\left[\min(R_t^{(r,d)}, A)\right] \geq \frac{\frac{A(A+1)}{2}\mathbb{P}[X_1 \geq A]}{\mathbb{E}[X_1]} \geq \frac{A^2 \cdot A^{-\alpha}}{2.\mathbb{E}[X_1]}.$$

As this inequality holds for arbitrary large A and $\alpha < 2$, we have $\lim_{t\to\infty}\mathbb{E}[R_t^{(r,d)}] = +\infty$. The collection of processes $((R^{(r,d)})_{t\geq 0})_{r\neq s}$ verifies the condition in Lemma 2.2. As $(t_k^{(r)})_{k\geq 0}$ and $(r_k)_{k\geq 0}$ are independent of this collection, we can therefore deduce that

$$\lim_{k\to\infty}\mathbb{E}\left[R_{t_k^{(r)}}^{(r_k,d)}\right] = +\infty;\ \text{hence,}\ \lim_{k\to\infty}\mathbb{E}[D_k] = +\infty.$$

Proving Theorem 1.1 for $\alpha \leq 1$. In this case, for any device r, the Markov chain defining $(R_t^{(r,d)})_{t\geq 1}$ is recurrent null, so that Orey's theorem (see [5, p. 131]) implies

$$\lim_{t\to\infty}\mathbb{P}\left[R_t^{(r,d)} = i\right] = 0 \text{ for all } i.$$

In particular, for any A,

$$\lim_{t\to\infty}\mathbb{P}\left[R_t^{(r,d)} < A\right] = 0 \text{ and } \lim_{t\to\infty}\mathbb{P}\left[R_t^{(r,d)} \geq A\right] = 1.$$

We have $\mathbb{E}[R_t^{(r,d)}] \geq A \cdot \mathbb{P}[R_t^{(r,d)} \geq A]$. As a consequence, and because the result holds for any arbitrary A, we have $\lim_{t\to\infty}\mathbb{E}[R_t^{(r,d)}] = +\infty$. This holds for any device r. Another application of Lemma 2 with Condition 2 allows us to prove $\lim_{k\to\infty}\mathbb{E}[R_{t_k^{(r)}}^{(r_k,d)}] = +\infty$.

The short-contact case. The result in Theorem 1.1 follows from the long-contact case, as the delay in the short contact case is always larger. The proof of Theorem 1.3 is a little more complex, but follows from classical results on Palm Calculus in discrete time and random walks. It may be found in Appendix B. □

To summarize, we have identified two regions where the behavior of the two-hop relaying algorithm would differ under the power law intercontact time assumption: When α is greater than 2, the algorithm converges to a finite expected delay as in the case of an exponential decay. When α is smaller than 2, the two-hop forwarding algorithm does not converge to a finite expected delay, as the delay that can be expected, starting from any initial condition, grows without bound with time. This remains true even for the long contact case, where data exchange is unlimited during contacts, and queuing in relay devices therefore has no impact on the delay experienced. In other words, the region $\alpha > 2$ may be thought of as the *stability region* of the two-hop relaying algorithm.

3.3 Generalization

In this section, we characterize the stability region (defined as the values of α for which an algorithm achieves a bounded delay) for the general class of oblivious algorithms. We conduct the following proofs in the long contact case only. We further assume, when $\alpha > 1$ and, therefore, that a steady state exists, that the system has reached its

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

stationary behavior; otherwise, when $\alpha \leq 1$, we start from any initial condition.

We generalize the two-hop relaying algorithm as follows: Instead of sending a single copy of a given data unit to a unique relay, the source will send m copies of each data unit: one to each of the first m relays that it meets. As we have assumed that the contact processes belonging to these relays are independent, the source may thereby reduce the total transmission delay by increasing its probability to pick a relay with a small delay to the destination among the m relays to which it has forwarded the message. This observation is made rigorous in the following lemma:

Lemma 1. *Let $(R_t^{(d_1,d'_1)})_{t\geq 0},\ldots,(R_t^{(d_m,d'_m)})_{t\geq 0}$ be the remaining intercontact times for m different pairs of devices $(d_i,d'_i)_{1\leq i\leq m}$. We suppose that they have reached their steady state.*

If we suppose that $m>1$ and $1+\frac{1}{m}<\alpha<2$, then

$$\mathbb{E}\Big[R_t^{(d_1,d'_1)}\Big]=\ldots=\mathbb{E}\Big[R_t^{(d_m,d'_m)}\Big]=+\infty \text{ and}$$
$$\mathbb{E}\Big[\min\Big(R_t^{(d_1,d'_1)},\ldots,R_t^{(d_m,d'_m)}\Big)\Big]<\infty.$$

Proof. As $\alpha>1$, Lemma 3.2 holds: A unique stationary distribution exists for the product chain

$$R_t^{(d_1,d'_1)},\ldots,R_t^{(d_m,d'_m)}$$

given as the product of the stationary distribution for each component. Hence,

$$\mathbb{P}\Big[\min\Big(R_t^{(d_1,d'_1)},\ldots,R_t^{(d_m,d'_m)}\Big)>i\Big]=\Big(\mathbb{P}\Big[R_t^{(d_1,d'_1)}>i\Big]\Big)^m \leq\left(\frac{1}{c_1(\alpha-1)}\right)^m(i+1)^{-m\cdot(\alpha-1)}.$$

The expectation of the minimum is therefore finite as soon as $-m\cdot(\alpha-1)<-1$ or, equivalently, $\alpha>1+\frac{1}{m}$. □

This result shows that, for α smaller than 2, the expected time to meet the destination is infinite. However, the expected time for the destination to meet a group of m devices may have a finite expected value, provided that $\alpha>1$ and that m is large enough. This observation is the key component in the next result, which proves that using a two-hop relaying strategy with m relays is sufficient to extend the stability region to any value of $\alpha>1$. This theorem also proves that the case $\alpha<1$, which is observed in most data sets, is of quite a different nature, as even unlimited flooding does not achieve a bounded delay. We comment on this difference further in Section 5.

Theorem 2. *Let us consider a source destination pair (s,d) and $t_k^{(s)},t_k^{(d)},D_k$ defined as in Theorem 1. We assume that all contacts are long.*

1. *If $\alpha>2$, there exists a forwarding algorithm using only one copy of the data, with a finite expected delay, such that, starting from any initial condition, $\lim_{k\to\infty}\mathbb{E}[D_k]=\bar{D}<+\infty$.*
2. *If $1<\alpha<2$, $m\in\mathbb{N}$ is chosen such that $\alpha>1+\frac{1}{m}$, and the network contains at least $N\geq 2m$ devices, there exists an algorithm using m relay devices such that, in steady state, $\mathbb{E}[D_k]=\bar{D}<+\infty$.*
3. *If $\alpha\leq 1$, for a network containing a finite number of devices and any forwarding algorithm, including flooding, we have, starting from any initial condition, $\lim_{k\to\infty}\mathbb{E}[D_k]=+\infty$.*

Proof.

Proving Theorem 2.1. This proof is just a reminder of the result of Theorem 1. The two-hop relaying algorithm may be chosen and it achieves a finite expected delay.

Proving Theorem 2.2. Let us assume that $\alpha>1+1/m$ and $N\geq 2m$, where $m\in\mathbb{N}$. The forwarding algorithm that we consider in this case is a two-hop relaying algorithm using m different relays.

Step 1. A bundle is created at time t in the source (denoted as device s). It is first transmitted to the m first devices that are met. We estimate first the time when each of these m relays are all contacted and have received the bundle. Let us consider the collection of remaining inter-contact time with all the other devices $(R_t^{(s,r)})_{r\neq s}$. This collection contains $N-1$ variables. If we consider a version of this collection, sorted for each time t, in the increasing order, the time to contact m different devices at time t is the $m\mathrm{th}$ value of this sorted sequence. Corollary 2, which is a simple variation of Lemma 1 shown in Appendix C, tells that this variable is of finite expected value if $\alpha>1+1/(N-1-m+1)$. This last assumption is automatically verified as $N-1-m+1=N-m\geq m$ by assumption.

Step 2. At time t', a copy of the bundle is present in each of the m relays that we denote $r_1,\ldots,r_m$. We now consider the vector $(R_{t'}^{(r_1,d)},\ldots,R_{t'}^{(r_m,d)})$ which describes the times needed for each of these relays to get in contact with the destination. The time length elapsed until the packet is delivered to the destination is taken as the minimum of this values. An application of Lemma 1 tells us that this time has a finite expected value.

As a consequence, the overall delay, from the time of creation of the bundle in the source to the delivery at the destination, is the sum of two variables with finite expectations. It is hence of finite expected value.

Proving Theorem 2.3. Let us consider, in this case, for a source s and any other device r in the network, the remaining time $R_t^{(s,r)}$ at time t until the next contact. As $\alpha<1$, all of these sequences of random variables are irreducible null recurrent Markov chains. By Orey's theorem ([5, p. 131]), we then have that $\lim\mathbb{P}[R_t^{(s,r)}=i]=0$ for all i when t tends to infinity. In particular, for any arbitrary large A, we have $\lim\mathbb{P}[R_t^{(d,d')}\leq A]=0$, so that

$$\mathbb{P}\left[\min_{r\neq s}R_t^{(s,r)}\geq A\right]=\mathbb{P}\left[\bigcap_{r\neq s}\Big\{R_t^{(s,r)}\geq A\Big\}\right]\to_{t\to\infty}1.$$

Consequently, $\mathbb{E}[\min_{r\neq s}R_t^{(s,r)}]$ diverges for large t.

As a consequence, starting from any initial condition, the time for a source to reach any other device is of infinite expectation as times increase. No forwarding algorithm, no matter how redundant, can then transport a packet within a finite expected delay using only opportunistic contacts between devices. □

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

Note. By comparison, the result in Theorem 2.3 applies to any case that includes short contacts as well as long contacts. Generally, a network containing N devices admits forwarding algorithms that achieve a bounded expected delay for any $\alpha > 1 + \frac{1}{\lfloor N/2 \rfloor}$. One example of those is flooding (that may use up to $N-2$ relays); that is not the only one, as a forwarding algorithm using only $\lfloor N/2 \rfloor$ relays is sufficient.

3.4 Summary

At this stage, we have established the following results for the class of oblivious forwarding algorithms defined in Section 3.1 in the long contact case:

- For $\alpha > 2$, any algorithm from the class we considered achieves a delay with finite mean.
- If $1 < \alpha < 2$, the two-hop relaying algorithm, introduced in [10], is not stable in the sense that the delay incurred has an infinite expectation. It is, however, still possible to design an oblivious forwarding algorithm that achieves a delay with finite mean. This requires that m duplicate copies of the data are produced and forwarded, where m must be greater than $\frac{1}{\alpha-1}$ and the network must contain at least $\frac{2}{\alpha-1}$ devices.
- If $\alpha < 1$, none of these algorithms, including flooding, can achieve a transmission delay with a finite expectation.

In other words, we have characterized the performance of all these algorithms in the face of extreme conditions (i.e., heavy tailed intercontact times). The last case, where $\alpha < 1$, corresponds to the most extreme situation, and the result we provide in this case seems at first unsatisfactory: None of the algorithms we have introduced can guarantee a finite expected delay. To make the matter worse, this case, where $\alpha < 1$, seems to be typical of the intercontact time distribution in the [10 minutes; 1 day] range for all the scenarios we have previously studied empirically. This implies overall that the expected delay for all the scenarios we have discussed before should be at least of the order of one day. Note that this was shown for any forwarding algorithms used and, even when queuing delay in relay devices are neglected. In fact, this is a negative result, and we come back to interpret it and discuss its implications in Section 5.

4 Related Work

Our opportunistic communication model is related to both Delay-Tolerant Networking and Mobile Ad Hoc Networking.[5] Research works on MANET, DTN, and, more recently, Pocket Switched Networks [12] confirm the importance of the problem we address, as several propositions were made to use mobile devices as relays for data transport. Such an approach has been used to enable communication where no contemporaneous path may be found [7], to efficiently gather information in a network of low power sensors [13], [14], [15], and to improve the spatial reuse of dense MANET [10], [11]. All those works have proved that the mobility model used has a strong impact on the performance of the algorithms they use.

5. www.dtnrg.org and www.ietf.org/html.charters/manet-charter.html.

We did not find any previous work studying the characteristics of intercontact time for users of portable wireless devices. However, we have identified related work in the area of modeling and forwarding algorithms.

A common property of the most common mobility models is that the tail of the intercontact time distribution decays exponentially. In other words, for these models, the intercontact time is light tailed: This is the case for i.i.d. location of devices in a bounded region (as assumed in [10]) or, more generally, for any random walk defined on a finite region using a comparison argument. It is also the case for the popular random way-point model as demonstrated in the Section 3 of [11]. This fact is commonly used in the analysis of mobile networks to estimate the delay of a scheme. It was shown recently that devices moving according to a Brownian motion in a bounded region exhibit heavy tailed intercontact time with a finite variance (corresponding in our analysis to the case $\alpha > 2$) (see [16] and the associated technical report).

The most relevant work is the algorithm proposed by Grossglauser and Tse in [10], further analyzed by Sharma and Mazumdar in [11]. The two-hop relay forwarding algorithm was initially introduced to study how the mobility of devices impacts the capacity of the network. Our work starts from very different assumptions. Most notably, we do not model the bandwidth limitation due to interference, as we focus only on the delay induced by mobility. However, some of the results that we show could be used to characterize the delay obtained in such contexts.

5 Conclusion

We have analyzed several network scenarios for opportunistic data transfer among mobile devices carried by humans using eight experimental data sets. For all data sets, we observe that the intercontact time between two devices can be approximated by a power law in the [10 minutes; 1 day] range. We prove in a simple model the following major results: Power law conditions may be addressed with finite expected delay by "oblivious forwarding algorithms" as long as the heavy tail index of the power law is greater than 1. When the heavy tail index is smaller than 1, the expected delay cannot be bounded for any forwarding algorithm of that type, even when one ignores the queuing occurring in each relay device. We have measured a heavy tail index smaller than 1 in all data sets. As a consequence, the expected delay is at least of the order of one day.

These observations bring new practical recommendations to evaluate the performance of forwarding algorithms. Most of the mobility models commonly used today are characterized by a light tailed intercontact time distribution for any pair of nodes. This property has been used in the past to estimate the delay in these networks. Our empirical findings of intercontact time distributions, in contrast, are well-approximated by a power law for values up to one day. Some mobility models can, in theory, be modified to account for this last property; this may be a future research direction. Another complementary direction, which is chosen in this paper, is to directly model opportunities between devices instead of their geographical locations. This approach has the advantage that it can be directly compared with a growing sets of real-life connectivity

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

traces, now publicly available. We believe that this is a practical solution, at least for some of the issues to be addressed in opportunistic networking.

More generally, our results are dealing with the feasibility of forwarding in opportunistic networks and their consequence requires further attention. At least three different directions may be followed:

- First, it might be that reasoning with an expected value of delay is not suitable, since the possible occurrence of a long delay is unavoidable whenever a forwarding algorithm is used. Applications for such networks should therefore be designed to cope with this aspect of opportunistic communication.
- Second, note that we did not model the general case where contact processes for a pair of nodes are heterogeneous or contain significant correlation. It is still possible that a finite expected delay exists in a more complex model that accurately reproduces the statistical properties of our data sets. This direction is appealing but it requires the removal of one of the modeling assumptions that we have made and which are common for most of the results currently known in this area. It also necessitates the design of a forwarding algorithm that differentiates between nodes; some schemes of that type have been only recently proposed [17], [18], [19].
- Third, one can investigate how to add connection opportunities in a mobile network using special devices or partial infrastructure that could, in some cases, be already available.

We are currently working on performing more human mobility experiments using different type of devices, and diverse sociological groups, in order to follow the directions we mention above. One of our long term goals is to study the properties of the actual traffic created by users in an opportunistic data network.

Appendix A

Preliminary Results

A.1 Independent Composition and Limit Expectation

Lemma 2. *Let $((F_t^{(i)})_{t\in\mathbb{N}})_{i\in I}$ be a finite collection of sequences of real valued random variables verifying $\lim_{t\to\infty}\mathbb{E}[F_t^{(i)}]=l$, where $l\in\mathbb{R}\cup\{+\infty\}$ and*

1. $\forall i,t,\mathbb{E}[F_t^{(i)}]\in\mathbb{R}$ *and* $l\in\mathbb{R}$ *or*
2. $\forall i,t,\mathbb{E}[F_t^{(i)}]\in\mathbb{R}\cup\{+\infty\}$ *and* $l=+\infty$.

Let $(t_k)_{k\in\mathbb{N}}$ and $(i_k)_{k\in\mathbb{N}}$ be two $\mathbb{N}$ valued processes, independent from F, such that $\lim_{k\to\infty} t_k=+\infty\ a.s.$

We then have $\lim_{k\to\infty}\mathbb{E}[F_{t_k}^{(i_k)}]=l$.

Proof. Let us first develop the following expectation:

$$\begin{aligned}\mathbb{E}\left[F_{t_k}^{(i_k)}\right] &= \sum_{i\in I}\sum_{t\geq 0}\sum_{j\geq 0} j\,\mathbb{P}\left[i_k=i,\ t_k=t,\ F_t^{(i)}=j\right]\\ &= \sum_{i\in I}\sum_{t\geq 0}\sum_{j\geq 0} j\,\mathbb{P}\left[i_k=i\right]\mathbb{P}\left[t_k=t\right]\mathbb{P}\left[F_t^{(i)}=j\right]\\ &= \sum_{i\in I}\sum_{t\geq 0}\mathbb{P}\left[i_k=i\right]\mathbb{P}\left[t_k=t\right]\mathbb{E}\left[F_t^{(i)}\right].\end{aligned}\tag{3}$$

If we suppose Lemma 2.1, we have $l<+\infty$ and

$$\forall\varepsilon>0,\exists T\ s.t.\ \left(t>T\Longrightarrow\left|\mathbb{E}\left[F_t^{(i)}\right]-l\right|<\frac{\varepsilon}{2}\right).$$

If $M=\sup_{i\in I,t\leq T}\left|\mathbb{E}\left[F_t^{(i)}\right]-l\right|$, there exists K such that

$$k>K\Longrightarrow\mathbb{P}\left[t_k\geq T\right]\geq 1-\frac{\varepsilon}{2\cdot M}$$

and, hence, $\left|\mathbb{E}\left[F_{t_k}^{(i_k)}\right]-l\right|$ can be bounded from above by

$$\begin{aligned}&\sum_{i\in I}\sum_{t\geq 0}\mathbb{P}[i_k=i]\,\mathbb{P}[t_k=t]\left|\mathbb{E}\left[F_t^{(i)}\right]-l\right|\\ &\leq\sum_{i\in I}\mathbb{P}[i_k=i]\left(M\cdot\sum_{t\leq T}\mathbb{P}[t_k=t]\right.\\ &\qquad\left.+\sum_{t>T}\mathbb{P}[t_k=t]\left|\mathbb{E}\left[F_t^{(i)}\right]-l\right|\right)\\ &\leq\sum_{i\in I}\mathbb{P}[i_k=i](\varepsilon/2+\varepsilon/2)\leq\varepsilon.\end{aligned}$$

If we suppose Lemma 2.2, we have $l=+\infty$ and

$$\forall A>0,\exists T,(t>T\Longrightarrow\mathbb{E}\left[F_t^{(i)}\right]\geq 2\cdot(A+1)).$$

Let $M'=\sup_{i\in I,t\leq T}\max\left(-\mathbb{E}\left[F_t^{(i)}\right],0\right)$. $\exists K$ such that

$$k>K\Longrightarrow\mathbb{P}[t_k\geq T]\geq\max(1/2,1-1/M'),$$

and

$$\begin{aligned}\mathbb{E}\left[F_{t_k}^{(i_k)}\right]&=\sum_{i\in I}\sum_{t\geq 0}\mathbb{P}[i_k=i]\mathbb{P}[t_k=t]\mathbb{E}\left[F_t^{(i)}\right]\\ &\geq\sum_{i\in I}\mathbb{P}[i_k=i]\left(-M\cdot\sum_{t\leq T}\mathbb{P}[t_k=t]\right.\\ &\qquad\left.+\sum_{t>T}\mathbb{P}[t_k=t]\mathbb{E}\left[F_t^{(i)}\right]\right)\\ &\geq\sum_{i\in I}\mathbb{P}[i_k=i]\left(-1+\frac{1}{2}(2(A+1))\right)\geq A.\end{aligned}$$

A.2 Remaining Intercontact

Because the contact process $(U_t^{(d,d')})_{t\geq 0}$ is a renewal process, the sequence $(R_t^{(d,d')})_{t\geq 0}$ of integers is a Homogeneous Markov Chain in $\mathbb{N}$ such that

$$\begin{cases}R_{t+1}^{(d,d')}=R_t^{(d,d')}-1 & \text{if } R_t^{(d,d')}>0,\\ R_{t+1}^{(d,d')}=i-1\text{ with prob. }\mathbb{P}[X=i] & \text{if } R_t^{(d,d')}=0.\end{cases}\tag{4}$$

This Markov Chain is clearly irreducible and aperiodic as $\mathbb{P}[X=1>0]$ and it is recurrent as X is almost surely finite. The following lemma characterizes its properties, which depend on the value of α based on classical results from the theory of Markov chains.

Lemma 3. *For any devices d,d',e,e' such that $(d,d')\neq(e,e')$, we have*

1. *If $\alpha>1$, $(R_t^{(d,d')})_{t\geq 0}$ is ergodic.*
2. *If $\alpha>1$, the chain $(R_t^{(d,d')},R_t^{(e,e')})_{t\geq 0}$ is ergodic and admits the following stationary distribution:*

$$\pi(i,j)=\frac{(i+1)^{-\alpha}(j+1)^{-\alpha}}{(c_1)^2},\ where\ c_1=\sum_{i'\geq 0}(i'+1)^{-\alpha}$$

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

such that we have in steady state

$$\frac{(i+2)^{-(\alpha-1)}}{c_1(\alpha-1)} \le \mathbb{P}\left[R_t^{(d,d')} > i\right] \le \frac{(i+1)^{-(\alpha-1)}}{c_1(\alpha-1)}.$$

3. *If* $\alpha \le 1$, $(R_t^{(d,d')})_{t\ge 0}$ *is recurrent null.*

Proof. Let us introduce $\mathtt{ret}_0$, the time for $R(d,d')$ to return in the state 0. From the structure of the Markov chain (4), starting from state 0, we can easily deduce that $\mathbb{E}_0[\mathtt{ret}_0] = \mathbb{E}[X]$. If $\alpha > 1$, we have $\mathbb{E}[X] < +\infty$, proving Lemma 3.1, and if $\alpha \le 1$, we have $\mathbb{E}[X] = +\infty$, proving Lemma 3.3.

By Lemma 3.1, we know that the Markov chain $R(d,d')$ is recurrent positive; hence, it admits a stationary distribution. It is easy to check, from its regenerative structure, that it is given by $\pi(i) = c_1(i+1)^{-\alpha}$, where $c_1 = 1/\sum_{i\ge 0}(i+1)^{-\alpha}$.

The same result holds for $R(e,e')$. As these two Markov Chains are independent, one can then check easily that the product Markov chain $(R(d,d'), R(e,e'))$, which is irreducible and aperiodic, admits a stationary distribution given by the product of the measure. It is hence ergodic.

In steady state, we have

$$\mathbb{P}\left[R_t^{(d,d')} > i\right] = \sum_{j>i} \pi(i) = \frac{1}{c_1}\sum_{j>i}(j+1)^{-\alpha}.$$

As the function $x \mapsto (x+1)^{-\alpha}$ is nonincreasing, we have

$$\int_{i+1}^{\infty}(x+1)^{-\alpha}dx \le \sum_{j>i}(j+1)^{-\alpha} \le \int_{i}^{\infty}(x+1)^{-\alpha}dx,$$

thus

$$\frac{(i+2)^{-(\alpha-1)}}{\alpha-1} \le \sum_{j>i}(j+1)^{-\alpha} \le \frac{(i+2)^{-(\alpha-1)}}{\alpha-1},$$

which completes the proof for Lemma 3.2. □

Smith's formula for $\alpha > 1$**.** For any devices d and d', the process $(R_t^{(d,d')})_{t\ge 0}$ is regenerative with respect to the delayed renewal sequence $(T_k^{(d,d')})_{k\ge 0}$. If we assume $\alpha > 1$, we have $\mathbb{E}[X] < +\infty$; hence, the interevent of the sequence $(T_k^{(d,d')})_{k\ge 0}$ admits a finite mean. We know, in this case (see [5, p. 148]), that

$$\lim_{t\to\infty} \mathbb{E}\left[f\left(R_t^{(d,d')}\right)\right] = \frac{\mathbb{E}\left[\sum_{t=T_0^{(d,d')}}^{T_1^{(d,d')}-1} f\left(R_t^{(d,d')}\right)\right]}{\mathbb{E}\left[T_1^{(d,d')} - T_0^{(d,d')}\right]}$$
$$\text{for any } f \text{ verifying } \mathbb{E}\left[\sum_{t=T_0^{(d,d')}}^{T_1^{(d,d')}-1}\left|f\left(R_t^{(d,d')}\right)\right|\right] < \infty. \tag{5}$$

Appendix B
Queuing with a Process of Service Instant

In constrast with classical queuing systems, the nodes of a mobile network only serve bundles from a given queue when they are in contact with the corresponding destination. In this section, we extend some well-known results on queues to handle this constraint.

Let us consider a queue receiving customers according to a point process $\mathbf{a} = \{a_k \mid k \in \mathbb{Z}\}$, that may be served only at some service instant, which follow a process $s = \{s_m \mid m \in \mathbb{Z}\}$. We make the following assumptions:

- A customer arriving at time t joins the queue and can be served starting from $t+1$.
- At each service instant, one customer from the queue is served, except if the queue is empty.

Hence, this system behaves as if the time slot were divided in two parts: In the first half of the time slot, a customer from the queue is served if the slot is a service instant; in the second half, new customers join the queue. Let us introduce $Q(t)$, the number of customers present after the first half of time slot t is completed. The process Q follows the recursion

$$Q(t) = \max(0, Q(t-1) + N_{\mathbf{a}}(t-1) - N_{\mathbf{s}}(t)), \tag{6}$$

where $N_{\mathbf{a}}$ (respectively, $N_{\mathbf{s}}$) denotes the counting measure associated with the point process $\mathbf{a}$ (respectively, $\mathbf{s}$).

B.1 Stationarity, Little's Law

Results are shown in the θ_stationary ergodic framework (see [6]). We assume here that θ is a measurable mapping $\Omega \to \Omega$, which preserves the probability measure (i.e., $\mathbb{P} \circ \theta = \mathbb{P}$) and is ergodic (all θ invariant events have probability 0 or 1).

A point process is called *stationary* with respect to θ if its counting measure verifies $N(\theta(\omega), C) = N(\omega, C+t)$, where $C \subset \mathbb{Z}$ and $\omega \in \Omega$. We define its *intensity* as $\mathbb{E}[N(0)]$.

The next result follows closely the proof of the stability regime for a single server queue (see [6, pp. 83-87]). It shows that, under a simple stability condition, the system admits a steady state that is stationary in a strong sense (compatible with the shift θ). The expected delay of a customer through this system is then given by a generalized Little Formula.

Notation. Following the usual convention of Palm calculus (here, in discrete time), we denote by $\mathbb{P}_{\mathbf{a}}^0$ the probability measure $\mathbb{P}$ under the condition that point process $\mathbf{a}$ has a point in $t = 0$. We number customer k with the convention that customer $k = 0$ denotes the last customer that arrived strictly before 1. We denote by V_k the sojourn time of customer k.

Lemma 4. *If* $\mathbf{a}, \mathbf{s}$ *are two stationary point processes with respect to* θ, *with respective intensities* λ, μ *such that* $\lambda < \mu$,

1. *There exists an initial condition,* $\tilde{Q} < \infty$ *a.s., such that the queue process verifies* $\tilde{Q}(t) = \tilde{Q} \circ \theta^t$.
2. *In this stationary regime,* $\mathbb{E}_{\mathbf{a}}^0[\tilde{V}_k] = 1 + \frac{1}{\lambda}\mathbb{E}[\tilde{Q}]$.
3. *If the queue starts empty,* $\limsup_{k\to\infty} \mathbb{E}[V_k] \le 1 + \frac{1}{\lambda}\mathbb{E}[\tilde{Q}]$.

Proof. We define the sequence of variables indexed by T

$$\tilde{Q}^{[T]} = \max_{-T \le -t \le 0}(N_{\mathbf{a}}(-t,\ldots,-1) - N_{\mathbf{s}}(-t+1,\ldots,0)).$$

Clearly, this sequence is positive, nondecreasing, and verifies

$$\tilde{Q}^{[T+1]} \circ \theta = \max\left(0, \tilde{Q}^{[T]} + N_{\mathbf{a}}(0) - N_{\mathbf{s}}(1)\right). \tag{7}$$

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

It then admits an a.s. limit, denoted by $\tilde{Q}$, verifying

$$\tilde{Q} \circ \theta = \max\big(0, \tilde{Q} + N_{\mathbf{a}}(0) - N_{\mathbf{s}}(1)\big).$$

This limit may take infinite values. Note that, since $\{\tilde{Q} = \infty\}$ is θ_invariant and θ is ergodic, it then has probability 1 or 0. In other words, either this limit is a.s. infinite or it is a.s. finite.

We can rewrite (7) as

$$\tilde{Q}^{[T+1]} \circ \theta = \tilde{Q}^{[T]} - \min\Big(\tilde{Q}^{[T]}, N_{\mathbf{s}}(1) - N_{\mathbf{a}}(0)\Big),$$

such that

$$\begin{aligned}\mathbb{E}\Big[\min\Big(\tilde{Q}^{[T]}, N_{\mathbf{s}}(1) - N_{\mathbf{a}}(0)\Big)\Big] &= \mathbb{E}\Big[\tilde{Q}^{[T]} - \tilde{Q}^{[T+1]} \circ \theta\Big] \\ &= \mathbb{E}\Big[\tilde{Q}^{[T]} - \tilde{Q}^{[T+1]}\Big] \leq 0.\end{aligned}$$

By monotone convergence, we deduce

$$\mathbb{E}[\min\big(\tilde{Q}, N_{\mathbf{s}}(1) - N_{\mathbf{a}}(0)\big)] \leq 0.$$

Assuming that $\tilde{Q}$ is a.s. infinite, the minimum above is then always given by the second term, which implies that $\mathbb{E}[N_{\mathbf{s}}(1) - N_{\mathbf{a}}(0)] = \mu - \lambda \leq 0$. By the converse induction,

$$\mu > \lambda \Longrightarrow \tilde{Q} < \infty \text{ a.s.},$$

which proves Lemma 4.1.

Lemma 4.2 is an application of the Campbell-Mecke egality:

$$\begin{aligned}\mathbb{E}[\tilde{Q}] &= \mathbb{E}\left[\sum_{k\in\mathbb{Z}} \mathbb{I}_{\{a_k \leq -1\}} \mathbb{I}_{\{a_k + V_k \geq 1\}}\right] \\ &= \lambda \sum_{v>0} \sum_{k\in\mathbb{Z}} \mathbb{I}_{\{k \leq -1\}} \mathbb{I}_{\{k+v \geq 1\}} \mathbb{P}^0_{\mathbf{a}}[V_0 = v] \\ &= \lambda \sum_{v>0} (v-1) \mathbb{P}^0_{\mathbf{a}}[V_0 = v] = \lambda(\mathbb{E}(0)_{\mathbf{a}}[V_0] - 1).\end{aligned}$$

We have a.s. $V_k \leq \tilde{V}_k$, which proves Lemma 4.3. □

B.2 Expected Queue Length

Lemma 5. *Assume a and s are two renewal point processes*

- *with intensities $\lambda < \mu$,*
- *such that interevent distribution F_a has a finite mean, and*
- *such that the interevent distribution F_s has a finite variance. Then,*

$$\mathbb{E}\left[\max_{t>0}(N_{\mathbf{a}}(1,\ldots,t) - N_{\mathbf{s}}(1,\ldots,t))\right] < \infty.$$

Proof. We recall the classical result on random walks (see [20, p. 270]): For $(Z_k)_k$ i.i.d., $\mathbb{E}[Z_k] < 0, \mathbb{E}[(Z_k^+)^2] < \infty$, we have

$$\mathbb{E}\left[\max_{k\geq 0}(Z_1 + \ldots + Z_k)\right] < \infty. \tag{8}$$

Let us prove first, for any $\nu > \lambda$,

$$\mathbb{E}\left[\max_{t>0}(N_{\mathbf{a}}(1,\ldots,t) - t\nu)\right] < \infty.$$

Let us denote by $S_1, S_2, \cdots$ the sequence of points of the process $\mathbf{a}$ that belongs to $\{0,1,2,\ldots\}$. They may be seen as the result of a random walk $S_n = X_1 + \ldots + X_n$, where variables $(X_k)_k$ are i.i.d. and follow the interevent distribution. The above expectation may be rewritten

$$\mathbb{E}\left[\max_{n>0}(n - \nu \cdot S_n)\right] = \nu \cdot \mathbb{E}\left[\max_{n>0}(Y_1 + \ldots + Y_n)\right],$$

where $Y_k = \frac{1}{\nu} - X_k$. Note that $\mathbb{E}[(Y_k)^2] \leq \nu^2 < \infty$ and $\mathbb{E}[Y_k] < 0$, proving by (8) that the above expectation is finite.

Next, we prove that, for any $\nu < \mu$,

$$\begin{aligned}\mathbb{E}\left[\max_{t>0}(t\cdot\nu - N_{\mathbf{s}}(1,\ldots,t))\right] &= \mathbb{E}\left[\max_{n>0}\big(S'_n \cdot \nu - n\big)\right] \\ &= \nu \cdot \mathbb{E}\left[\max_{n>0}(Z_1 + \ldots + Z_n)\right] < \infty\end{aligned}$$

as $Z_k = X'_k - \frac{1}{\nu}$, $\mathbb{E}[Z_k] < 0$, and $\mathbb{E}\Big[(Z_k^+)^2\Big] \leq \mathbb{E}\Big[(X'_k)^2\Big] < \infty$.

To conclude, we choose ν such as $\lambda < \nu < \mu$ and we have

$$\begin{aligned}&\mathbb{E}\left[\max_{t>0}(N_{\mathbf{a}}(1,\ldots,t) - N_{\mathbf{s}}(1,\ldots,t))\right] \\ &= \mathbb{E}\left[\max_{t>0}(N_{\mathbf{a}}(1,\ldots,1) - t\cdot\nu + t\cdot\nu - N_{\mathbf{s}}(1,\ldots,t))\right] \\ &\leq \mathbb{E}\left[\max_{t>0}(N_{\mathbf{a}}(1,\ldots,t) - t\cdot\nu) + \max_{t>0}(t\cdot\nu - N_{\mathbf{s}}(1,\ldots,t))\right].\end{aligned}$$

□

Corollary 1. *If $\mathbf{a}$ and $\mathbf{s}$ satisfy the conditions of Lemma 5,*

$$\mathbb{E}[\tilde{Q}] < \infty \textit{ and } \mathbb{E}^0_{\mathbf{a}}[\tilde{V}_k] < \infty.$$

Proof. According to the proof of Lemma 4,

$$\tilde{Q} = \max_{t\geq 0}(N_{\mathbf{a}}(-t,\ldots,-1) - N_{\mathbf{s}}(-t+1,\ldots,0)).$$

The result is then following the above lemma. □

Appendix C
Proof of Theorem 1 in the Short-Contact Case

Let us summarize the results from the above sections: In a queue with arrival $\mathbf{a}$ and service instant $\mathbf{s}$, customers experienced a finite expected delay if 1) $\mathbf{a}$ and $\mathbf{s}$ are renewal processes, 2) the stability condition is verified, and 3) interservice-instant distribution has a finite variance. Note that Conditions 2 and 3 are necessary. In the following, we present a scheme ensuring that all queues implemented in the mobile nodes verify Conditions 1, 2, and 3. It may be improved at the cost of an additional effort to weaken Assumption 1.

Each source devices s maintain a set of $N-1$ *source queues* corresponding to each other device. We assume that bundles are created in each of these queues according to a renewal process with intensity $\lambda < \frac{1-p}{2\cdot\mathbb{E}[X]}$ with $p > 0$. When another

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

device d is met during an odd time slot, a bundle from the queue associated with d is served, if this queue is not empty. The device d may be the destination for s, but, otherwise, the bundle is entering a *relay queue* (see below). For technical reasons, we also assume that, with a small probability p, taken independently, an *independent blocking* occurs and no bundle at all is sent by the source during this contact.

All devices (including all sources) maintain, in addition, $N-1$ *relay queues*, each one corresponding to a given destination. When a bundle is received during an odd time-slot (as described above), it is entering the *relay queue* corresponding to its destination. If another device d is met during an even time slot, a bundle for destination d is sent, unless the corresponding queue is empty.

Let us prove that bundles experience finite expected delay in each of these queues:

- Each *source queue* receives and serves bundles according to stationary processes that satisfy Conditions 1, 2, and 3.
- A *relay queue* satisfies Conditions 2 and 3; unfortunately, the arrival process in this queue is not a renewal process. Nevertheless, the same result holds by a comparison. All arrival times of a bundle in this *relay queue* are included in a *quasisaturated* renewal process (that includes all meeting times with the source corresponding to the destination of the queue, without independent blocking). Note that the expected delay in a *relay queue* is never larger than the expected delay in the same queue with a *quasisaturated* arrival process. One can check easily that this last case verifies Conditions 1, 2, and 3, proving that the expected delay is finite in both cases.

We deduce that all sources can transmit to their destination at a rate smaller than $\frac{(N-1)\cdot(1-p)}{2\mathbb{E}[X]}$ such that bundles experienced a finite expected delay. As p may be chosen arbitrarily, the same result holds for any rate smaller than $\frac{N-1}{2\cdot\mathbb{E}[X]}$.

C.1 Proof of Corollary of Lemma 1

For any real numbers $(x_1,\dots,x_m)$ and $i \le m$, let us denote by $\mathtt{ord}(i,(x_1,\dots,x_m))$ the ith element of the sequence after it is reordered in the increasing order. In particular, $\mathtt{ord}(1,(x_1,\dots,x_m)) = \min(x_1,\dots,x_m)$. We have:

Corollary 2. *Let* $\left(R_t^{(d_1,d'_1)}\right)_{t\ge0},\dots,\left(R_t^{(d_m,d'_m)}\right)_{t\ge0}$ *be the remaining intercontact times for* m *different pairs of devices* $(d_i,d'_i)_{1\le i\le m}$. *We suppose that* $\alpha > 1+\frac{1}{m-j+1}$, *then*

$$\mathbb{E}\Big[\mathtt{ord}\Big(j,\Big(R_t^{(d_1,d'_1)},\dots,R_t^{(d_m,d'_m)}\Big)\Big)\Big] < \infty.$$

Proof. Let $M_j = \mathtt{ord}\Big(j, R_t^{(d_1,d'_1)},\dots,R_t^{(d_m,d'_m)}\Big)$.

$$\mathbb{P}[M_j > n] = \mathbb{P}\Big[\#\Big\{i \;\Big|\; R_t^{(d_i,d'_i)} > n\Big\} \ge m-j+1\Big].$$

This is the probability that at least $m-j+1$ events occur on a collection of m variables. Note that all these events are independent, and each of them occurs with the same probability $p \le \frac{(n+1)^{-(\alpha-1)}}{c_1(\alpha-1)}$. As a consequence, the above probability may be rewritten as

$$\sum_{k=m-j+1}^{m}\binom{k}{m}p^k(1-p)^{m-k} \le p^{m-j+1}\sum_{k=m-j+1}^{m}\binom{k}{m}.$$

This proves that, for $c_2 = \frac{1}{c_1(\alpha-1)}\sum_{k=m-j+1}^{m}\binom{k}{m}$,

$$\mathbb{P}[M_j > n] \le c_2(n+1)^{-(m-j+1)(\alpha-1)},$$

which implies that $\mathbb{E}[M_j] < \infty$ as soon as

$$\alpha > 1 + \frac{1}{m-j+1}.$$

□

ACKNOWLEDGMENTS

The authors would like to gratefully acknowledge Dave Kotz and Tristan Henderson at Dartmouth College, Geoff Voelker at the University of California, San Diego, Eyal de Lara and Jing Su at the University of Toronto, and the Reality Mining project at the Massachusetts Institute of Technology for providing their data. They would also like to thank Vincent Hummel and Ralph Kling at Intel for their support, Mark Crovella for his insightful comments on their findings, Laurent Massoulié for his advice on some of the proofs, and the participants of the three experiments they conducted. This work was partially funded by the Information Society Technologies program of the European Commission in the framework of the FET-SAC HAGGLE project (027918).

REFERENCES

[1] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket Switched Networks and the Consequences of Human Mobility in Conference Environments," *Proc. ACM SIGCOMM First Workshop Delay Tolerant Networking and Related Topics,* 2005.

[2] M. McNett and G.M. Voelker, "Access and Mobility of Wireless PDA Users," technical report, Dept. of Computer Science and Eng., Univ. of California, San Diego, 2004.

[3] T. Henderson, D. Kotz, and I. Abyzov, "The Changing Usage of a Mature Campus-Wide Wireless Network," *Proc. MobiCom '04,* pp. 187-201, 2004.

[4] N. Eagle and A. Pentland, "Reality Mining: Sensing Complex Social Systems," *J. Personal and Ubiquitous Computing,* 2005.

[5] P. Bremaud, *Markov Chains, Gibbs Field, Monte Carlo Simulation and Queues,* first ed. Springer-Verlag, 1999.

[6] F. Baccelli and P. Bremaud, *Elements of Queuing Theory,* second ed. Springer-Verlag, 2003.

[7] A. Vahdat and D. Becker, "Epidemic Routing for Partially-Connected Ad Hoc Networks," technical report, Univ. of California, San Diego, 2000.

[8] X. Chen and A. Murphy, "Enabling Disconnected Transitive Communication in Mobile Ad Hoc Networks," *Proc. Workshop Principles of Mobile Computing,* pp. 21-27, 2001.

[9] J.A. Davis, A.H. Fagg, and B.N. Levine, "Wearable Computers as Packet Transport Mechanisms in Highly-Partitioned Ad-Hoc Networks," *Proc. Fifth IEEE Int'l Symp. Wearable Computers (ISWC '01),* p. 141, 2001.

[10] M. Grossglauser and D. Tse, "Mobility Increases the Capacity of Ad Hoc Wireless Networks," *IEEE/ACM Trans. Networking,* vol. 10, no. 4, pp. 477-486, 2002.

[11] G. Sharma and R.R. Mazumdar, "Delay and Capacity Trade-Off in Wireless Ad Hoc Networks with Random Way-Point Mobility," preprint, School of Electrical and Computer Eng., Purdue Univ., 2005, http://min.ecn.purdue.edu/linx/paper/it05.pdf.

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.

[12] J. Scott, P. Hui, J. Crowcroft, and C. Diot, "Haggle: A Networking Architecture Designed around Mobile Users," *Proc. Third IFIP Wireless on Demand Network Systems Conf.,* 2006.
[13] S. Jain, R.C. Shah, G. Borriello, W. Brunette, and S. Roy, "Exploiting Mobility for Energy Efficient Data Collection in Sensor Networks," *Proc. IEEE Symp. Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks,* 2004.
[14] 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," *Proc. 10th Int'l Conf. Architectural Support for Programming Languages and Operating Systems (ASPLOS-X),* pp. 96-107, 2002.
[15] T. Small and Z.J. Haas, "The Shared Wireless Infostation Model: A New Ad Hoc Networking Paradigm (or Where There Is a Whale, There Is a Way)," *Proc. MobiHoc '03,* pp. 233-244, 2003.
[16] X. Lin, G. Sharma, R.R. Mazumdar, and N.B. Shroff, "Degenerate Delay-Capacity Tradeoffs in Ad-Hoc Networks with Brownian Mobility," *IEEE/ACM Trans. Networking,* vol. 14, pp. 2777-2784, 2006.
[17] H. Dubois-Ferriere, M. Grossglauser, and M. Vetterli, "Age Matters: Efficient Route Discovery in Mobile Ad Hoc Networks Using Encounter Ages," *Proc. MobiHoc '03,* pp. 257-266, 2003.
[18] A. Lindgren, A. Doria, and O. Schelén, "Probabilistic Routing in Intermittently Connected Networks," *SIGMOBILE Mobile Computing Comm. Rev.,* vol. 7, no. 3, pp. 19-20, 2003.
[19] J. Leguay, T. Friedman, and V. Conan, "Evaluating Mobility Pattern Space Routing for DTNs," *Proc. INFOCOM,* 2006.
[20] S. Asmussen, *Applied Probability and Queues,* second ed. Springer-Verlag, 2003.

Augustin Chaintreau joined the Thomson Paris Research Lab in 2006, soon after receiving the PhD degree in computer science and mathematics from Ecole Normale Supérieure de Paris. He worked previously at INRIA under the supervision of Francois Baccelli, and in collaboration with Alcatel Bell and IBM research on the nature and the consequences of interaction between TCP flows. He visited Intel Research Cambridge for one year until 2005. His current research interests are in designing and analyzing algorithms for opportunistic, wireless, and overlay networks.

Pan Hui received the bachelor's and Mphil degrees from the Electricical and Electronic Enginnering Department at the University of Hong Kong in 2002 and 2004, respectively. He is currently a PhD student in the System Research Group, Computer Laboratory, University of Cambridge. He began work on the PhD degree in 2004 and is expected to graduate in fall 2007. See his Web page at http://www.cl.cam.ac.uk/~ph315/.

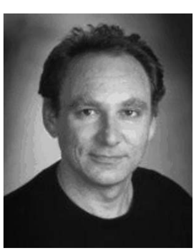

Jon Crowcroft is the Marconi Professor of Networked Systems in the Computer Laboratory of the University of Cambridge. Prior to that, he was professor of networked systems in the Computer Science Department at University College London. He is a fellow of the ACM, the British Computer Society, the IEE, the Royal Academy of Engineering, and the IEEE. He was a member of the IAB and was the general chair for ACM SIGCOMM from 1995 to 1999. He is on the editorial team for COMNET,and on the program committees for ACM SIGCOMM and IEEE INFOCOM. He has published five books, the latest of which is *Linux TCP/IP Implementation* (Wiley, 2001). Currently, he is the Principle Investigator for the CMI-funded Communications Research Network, which is a 3M pound government/industry/academia multidisciplinary collaboration to automate the successful exploitation of disruptive communications technologies.

Christophe Diot received the PhD degree in computer science from Institut National Polytechnique de Grenoble in 1991. He was with INRIA Sophia-Antipolis from October 1993 to September 1998, Sprint (Burlingame, California) from October 1998 to April 2003, and Intel Research (Cambridge, United Kingdom) from May 2003 to September 2005. He joined Thomson in October 2005 to start and manage the Paris Research Lab (http://parislab.thomson.net). Dr. Diot's future research activities will focus on communication services and platforms. He is a fellow of the ACM.

Richard Gass is a research engineer at Intel Research in the Pittsburgh lablet. His interests are in opportunistic communication, wireless networks, and vehicular ad hoc networks. He is also enrolled as a PhD student at the University of Pierre and Marie Curie in Paris.

James Scott joined Microsoft Research in Cambridge, United Kingdom, in January 2007, and previously worked at Intel Research in Cambridge, United Kingdom, for four years. He received the PhD degree from the University of Cambridge under Professor Andy Hopper. His research interests are broad and include mobile networking and ubiquitous computing.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.

Authorized licensed use limited to: University of Illinois. Downloaded on January 29, 2010 at 17:43 from IEEE Xplore. Restrictions apply.