



Pocket Switched Networks:
Real-world mobility
and its consequences for
opportunistic forwarding

Augustin Chaintreau, Pan Hui, Jon Crowcroft,
Christophe Diot, Richard Gass, James Scott

February 2005

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<http://www.cl.cam.ac.uk/>

© 2005 Augustin Chaintreau, Pan Hui, Jon Crowcroft,
Christophe Diot, Richard Gass, James Scott

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

<http://www.cl.cam.ac.uk/TechReports/>

ISSN 1476-2986

Pocket Switched Networks: Real-world mobility and its consequences for opportunistic forwarding

Augustin Chaintreau[†], Pan Hui^{*}, Jon Crowcroft^{*},
Christophe Diot[†], Richard Gass[†] and James Scott[†]

[†] Intel Research Cambridge

{augustin.chaintreau, christophe.diot, richard.gass, james.w.scott}@intel.com

^{*} University of Cambridge

{pan.hui, jon.crowcroft}@cl.cam.ac.uk

Abstract

Opportunistic networks make use of human mobility and local forwarding in order to distribute data. Information can be stored and passed, taking advantage of the device mobility¹, or forwarded over a wireless link when an appropriate contact is met². Such networks fall into the fields of mobile ad-hoc networking and delay-tolerant networking. In order to evaluate forwarding algorithms for these networks, accurate data is needed on the intermittency of connections.

In this paper, the inter-contact time between two transmission opportunities is observed empirically using four distinct sets of data, two having been specifically collected for this work, and two provided by other research groups.

We discover that the distribution of inter-contact time follows an approximate power law over a large time range in all data sets. This observation is at odds with the exponential decay expected by many currently used mobility models. We demonstrate that opportunistic transmission schemes designed around these current models have poor performance under approximate power-law conditions, but could be significantly improved by using limited redundant transmissions.

1 Introduction

With the explosive deployment of mobile wireless devices in recent years, the research community has focused on the networking issues these represent. Such devices can transfer data in two ways - firstly by transmitting it over a wireless (or wired) network interface, and secondly by being carried from location to location by their user. The use of device mobility for data transmission requires fundamentally different networking protocols than those such as TCP/IP, which rely on contemporaneous connectivity between their endpoints.

¹carried, for example, in one's trouser *pocket*!

²for example, as one moves into a *pocket* of wireless connectivity!

We use the term opportunistic networking to refer to data exchanges based on the connection “opportunities” that arise whenever mobile devices happen to come into wireless range due to the mobility of their users. Such situations are prevalent in many regions of the world where broadband access infrastructure has limited coverage (as well as cost and application constraints) ; Thus, users have “islands of connectivity” (e.g. home or work), but are also likely to sporadically be in range of other users while in between. It is also worth noting that access infrastructure is vulnerable to natural disasters or other failures, and that in such exceptional situations, opportunistic networking may be the only feasible way to carry important data.

Opportunistic networking is a sub-class of both Delay-Tolerant Networking (DTN) [1] and Mobile Ad-Hoc Networking (MANET) [2]. In this paper we are concerned specifically with the part of the design space where nodes are often out of contact with one another.

This paper makes three key contributions.

1. We present the design and results of an experiment which provides a picture of some typical human mobility patterns in the context of opportunistic connections.
2. We analyse both the data gathered by ourselves as well as two other large data sets made available from other mobility experiments and show that an approximate power law holds over an extended range of values for the inter-contact times of nodes in all four experiments.
3. We provide a proof that one major class of previously proposed opportunistic forwarding algorithms will not perform well in these conditions and propose measures which may increase the performance of such algorithms.

We next discuss the related work and then present these contributions in the order above. The paper concludes with a brief summary of contributions and presentation of future work.

2 Related work

We have investigated three areas for related work:

Measurement : A number of research groups have conducted studies into mobility in the context of networking. Many of these are aimed at analysing and informing the design of infrastructure-based networks, but their data and results are also relevant for opportunistic networks. This category includes Balazinska and Castro’s study [3] as well as the data gathered at UCSD [4] and Dartmouth [5] which we analyse in this paper.

Modeling : Much of the work in DTN and MANETs concerns the modeling of mobility or location [6, 7]. The goal of the models has typically been to drive the evaluation of routing schemes which assume that the majority of nodes are connected most of the time. This is not likely to be relevant since we are in the part of the design space where nodes are often disconnected. Indeed, the purpose of this paper is to model the distribution of these disconnection times and its impact on forwarding decisions, as noted next.

A common property of many mobility models found in the literature is that the inter-contact distribution decays exponentially over time.

One of the simplest examples, introduced in [8], in the case where nodes locations are i.i.d. with a uniform distribution in a bounded region, the success of communication between each sender-receiver pair has a fixed probability $p > 0$ at each time slot t .

In this situation, if we consider a particular sender-receiver pair, the remaining time to the next contact is distributed exponentially following the distribution X :

$$\mathbb{P}[X > t] = p^t \tag{1}$$

This property is also true for the popular random waypoint model, see [9]. In this article, a brownian motion model is analysed as well. The authors claim that the inter-contact time is in this case stochastically bounded by an exponential random variable.

Forwarding : Su et al.[10] used traces of human mobility patterns gathered with PDAs to evaluate the feasibility of opportunistic networking. Some sensor networks act in an opportunistic fashion, such as Zebranet [11], which uses opportunistic connections between zebra-mounted nodes to transfer sensor data and thus collect statistics about zebra populations. Similar research exists with whales [12]. However, these experiments are largely pragmatic proofs-of-concept.

The most relevant work when trying to find forwarding algorithms for networks that are frequently disconnected, is the algorithm proposed by Grossglauser and Tse in [8], further analysed in [9]. The principal motivation of that work was somewhat different from ours: it was to find the available increase in capacity of the multi-hop radio network as a function of the mobility and the number of nodes. In the process of exploiting mobility and trading it off against transmission, the authors created an opportunistic forwarding algorithm.

Given the nature of intermittent connectivity, it is likely that successful forwarding algorithms are based purely on locally learned information. The regime in some senses is even more resource starved than a MANET where one eschews proactive routing. Thus we need to measure what we can statistically learn locally in a variety of likely scenarios, and then use those measurements to drive the evaluation of appropriate forwarding algorithms.

As envisaged in the introduction, one can imagine that such a regime might operate only in parts of a network, where in other parts, connectivity is maintained. We will discuss this in the conclusions under further work.

3 Gathering data on transfer opportunities

In order to conduct informed design of forwarding policies and algorithms for opportunistic networks, it is important to gather data on the frequency and duration of contact between humans (and the devices they carry). However, this is not easy to gather – ideally, a data set would cover a large user base over a large time period, as well as include data on connection opportunities encountered twenty-four hours a day. This presents many practicality issues ; dealing with deployment of mobile devices to a large user population, the battery life of the devices, and minimising the inconvenience to users of carrying the devices so that they are willing to do so at all times.

3.1 Data from the research community

We first examined the data made available to the community by people who have performed previous measurement exercises. Two data sets emerged, namely from UCSD [4] and Dartmouth [5]. Both 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.

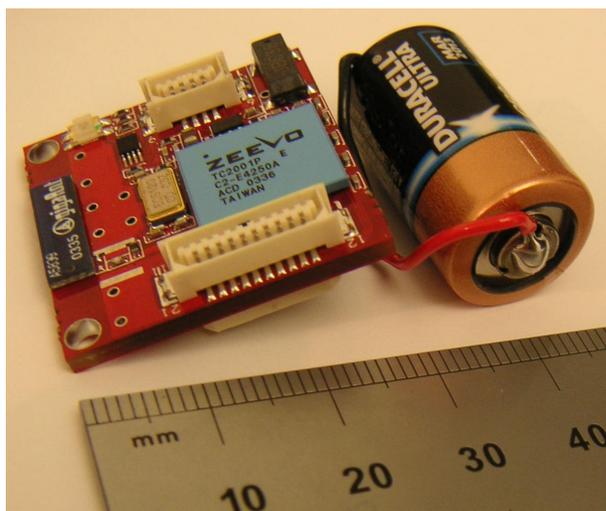


Figure 1: An iMote connected to a CR2 battery

The durations of the logs 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 pre-processing. For both data sets, we made the assumption that mobile devices in sight of the same AP would also be able to communicate directly (in ad-hoc mode), and created a list of transmission opportunities by determining, for each pair of nodes, the set of time regions for which they shared at least one AP.

Unfortunately, this assumption introduces inaccuracies. On one hand, it is overly optimistic, since two devices attached to the same access point may still be out of range of each other. On the other hand, the data might omit connection opportunities, since two devices may pass each other at a place where there is no instrumented access point, and this contact would not be logged. Furthermore, it is hard to ensure that the devices are in fact co-located with their owner at all times. Despite these inaccuracies, the WiFi traces are a valuable source of data, since they span many months and include thousands of nodes. In addition, considering two devices connected to the same AP are potentially in contact is not altogether unreasonable, as these devices could indeed communicate through the AP, without using end-to-end connectivity.

3.2 iMote experiments

In response to the limitations of the previous traces for our work, we set up our own experiments making use of the iMote platform made by Intel Research. iMotes (see Figure 1) are derived from the Berkeley Mote³, with the current version based around the Zeevo TC2001P system-on-a-chip providing an ARM7 processor and Bluetooth support. Along with a 950mAh CR2 battery, each iMote was enclosed in packaging designed to be convenient for test subjects to continually carry. Two types of packaging were made available : some iMotes were made into keyfobs while others were enclosed in small boxes (see Figure 2). Subjects were asked to pick the form factor which allowed them to conveniently keep the iMote with them at all times, with most simply attaching the iMote to their keys.

We programmed the iMotes to perform an inquiry for five seconds every two minutes. While

³see <http://www.xbow.com/>

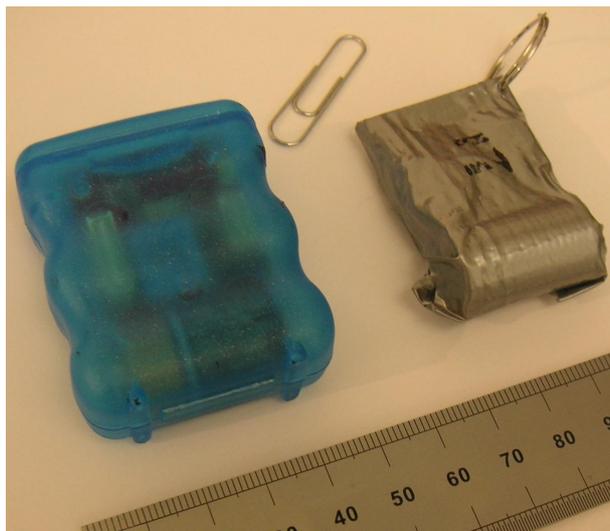


Figure 2: Packaged iMotes in boxed and keyfob form factors.

the Bluetooth specification indicates ten seconds should be used, our experience with the iMote hardware is that the vast majority of nearby nodes are seen in the first five seconds or not at all. Performing an inquiry is power-intensive and it is important to minimise time spent in this mode. The iMotes spend 120 seconds (plus or minus twelve seconds with a uniform random distribution) in a sleep mode where they are able to respond to inquiries. The reason for introducing randomness is that iMotes cannot respond to inquiry while themselves performing inquiry. Without randomness, iMotes might synchronise in such a way that they would never see each other. The two minute inquiry period was chosen to provide an estimated lifetime of one week for the iMotes with CR2 batteries. During the experiment, each iMote used flash memory to log “contact” data for all visible Bluetooth nodes (including iMotes as well as other Bluetooth devices), with each contact being represented by a tuple (MAC address, start time, end time). A twenty-four hour pilot deployment was performed in order to iron out software bugs and refine the deployment methodology and packaging mechanisms.

Two iMote experiments (“iMote A” and “iMote B”) were conducted. Experiment iMote A included seventeen researchers and interns working at Intel Research Cambridge, while iMote B involved eighteen doctoral students and faculty comprising a research group at the University of Cambridge Computer Lab. Unfortunately, real world factors contributed to the malfunction of some of the iMotes. As a result, the two experiments resulted in data from nine and twelve iMotes.

The iMotes were collected after their battery had expired, and the flash memory was read. A number of post-processing steps were undertaken on the data. A time basis consistent for all iMotes was reconstructed by (1) using the known start time of the experiment, and by using a program to find “mutual” sightings (i.e. where two iMotes see each other) of long duration, and (2) determining the clock offset between the two iMotes. The synchronisation was checked manually. iMotes contacts were classified as “internal” with other iMotes and “external” with other types of device. External contacts are a valuable source of data. iMotes are deployed to a small set of advanced users. The external contacts are numerous and include anyone who has an active Bluetooth device in the vicinity of the iMote users, thereby measuring the actual Bluetooth deployment. The internal contacts, on the other hand, represent the contacts that

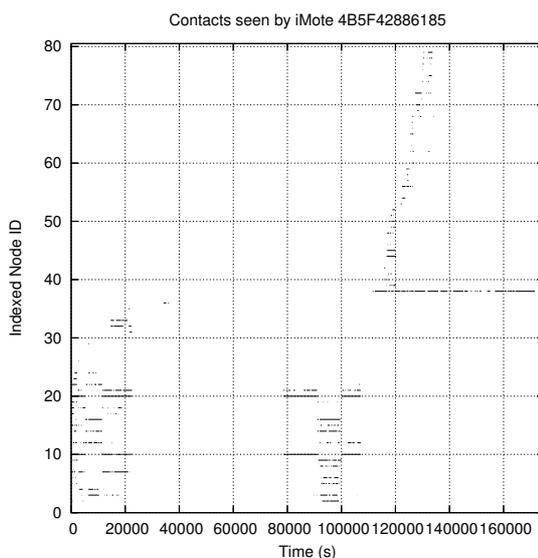


Figure 3: Contact history of one iMote.

each of our participants would have, if they were equipped with devices (such as smart phones) which are always-on and always-carried, which supported opportunistic networking.

A typical log is illustrated in Figure 3. On the Y axis, the various contacts have been indexed, with all iMotes having low index (i.e. 1-18), while external contacts have higher indexes. This allows us to quickly view the data.

An anonymised version of our data will be made available to other research groups on demand.

3.3 Comparison of data sets

A summary of the features of the four data sets is shown in Table 1. These data sets give us a chance to study four different user populations, using two different wireless technologies. The iMote experiments have the advantages that the users are more likely to carry the small-form-factor iMote at all times, and that logging takes place wherever the user is and not just when the users are near APs. The WiFi-based experiments have larger user populations and durations, and include all contacts occurring at the instrumented locations. Thus, the data sets are complementary in many ways.

User Population	Intel	Cambridge	UCSD	Dartmouth
Device	iMote	iMote	PDA	Laptop/PDA
Network type	Bluetooth	Bluetooth	WiFi	WiFi
Duration (days)	3	5	77	114
Granularity (seconds)	120	120	20	300
Nodes participating	141	238	261	6648
Number of contacts	3,984	8,856	175,105	4,058,284

Table 1: Comparison of data sets used

4 Characterization of contact opportunities

This section reports our observations on the four mobility data sets described above. The consequences of these observations on opportunistic networking forwarding algorithms will be described in the following section.

4.1 Definitions

We are interested in the characteristics of connection opportunities, i.e. how many and when do they occur, how often and how long. We choose to characterize these opportunities in term of contacts. To do so, we define two parameters represented on Figure 4 :

- The *contact duration* is the time interval for which two network devices can communicate when they come into range. The number of such contacts and the distribution of contact durations is an important factor in determining the capacity of opportunistic networks. It gives insight on how much data can be transferred at each opportunity.
- The *inter-contact time* is the time interval between two contacts. This parameter strongly affects the feasibility of opportunistic networking, and has rarely been studied in the literature. The nature of the distribution will affect the choice of suitable forwarding algorithms to be used to maximize the successful transmission of messages in a bounded delay.

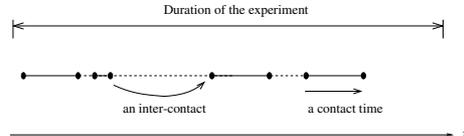


Figure 4: Contact and Inter-contact times for a pair of nodes.

In this work, we focus on the analysis of the distribution of inter-contact time. For completeness, we also briefly discuss the contact time distribution. Two remarks must be made at this point. First, we computed the distributions of these two variables per event (i.e. considering each one of the value taken by these periods of time). We did not compute the distribution of the variable as seen at a random instant in time by the pair. These two distributions are in correspondence by a classical result from renewal theory (see [13]).

In addition to that, our results are influenced by the duration and granularity of the experiments. For event lengths approaching the duration of the experiment there is an artificially lower likelihood of observation, and events lasting longer than the experiment cannot be observed. Similarly, for short event lengths, the data is affected by the granularity of measurement (which ranges from twenty seconds to five minutes in our data sets).

4.2 Inter-contact time

We study inter-contact time first as it is the parameter that has the most significant impact on the feasibility of opportunistic networking. Inter-contact time affects the frequency with

which packets can be transferred between networked devices. In Figure 5 we plot inter-contact time distributions for all four data sets (i.e. iMote A, iMote B, UCSD, and Dartmouth). The significant region is the middle of the graph, with the leftmost and rightmost parts showing artifacts due to the granularity and duration of the experiments as described above. In this region, all four experiments show the same behavior, an *approximate power law*, as evidenced by the straightness of the curve. A power law is characterised by its coefficient reflecting the slope of the line on log-log graphs — this coefficient is important for the analysis presented in Section 5. For the iMote experiments, the coefficient is 0.5 for the range [100s; one day], with a slightly convex distribution (flatter than that of a power law). The Dartmouth data exhibits an approximate power law with a coefficient of 0.33 on the range [100s; 1 week]. The UCSD distribution coefficient is also 0.33, but over a more limited range [100s; two days].

The approximate power law shape means that the inter-contact distribution is heavy-tailed over this range - i.e. the tail distribution function decreases slowly. This is contrary to the exponential decay of many mobility models put forward in the literature. As a result, opportunistic networking algorithms which have been designed around exponential models must be re-evaluated in the light of our observations (see next section). In the Bluetooth traces, 15% of inter-contact durations are greater than one hour, and 5% are greater than one day. In the WiFi experiments, the large experiment duration allows us to analyse large inter-contact durations. In the Dartmouth trace, we find that they are far from negligible: 20% are more than a day, 10% are more than a week. In the UCSD trace, 15% are more than a day, and 4% are more than one week.

While the WiFi experiments have longer durations, longer inter-contact times may be affected by the more limited mobility of laptops or PDAs as their users may not carry them all the time. As one might expect, the iMote experiments show lower inter-contact times (as illustrated by the lower coefficient of the approximate power law). This is an encouraging sign for the field of opportunistic networking - with the always-carried, always-on nature of devices such as smart phones, more connection opportunities will be found. What is striking is that the same overall pattern, the approximate power law, seems to apply to both Bluetooth and WiFi data despite the differences in experimental methodology.

Internal vs. external iMote contacts

While the above iMote plots used both internal and external contacts, it is also instructive to look at each of these classes in isolation. These are shown in Figure 6. We observe that the internal contact plot is almost indistinguishable from the plot containing all contacts (Figure 5), while the external contact plot has a very similar distribution, again showing an approximate power law. This indicates that our experimental methodology using iMotes is suitable for studying the behavior of the deployed Bluetooth user base.

4.3 Impact of time of day on opportunities

Since human movement patterns have daily cycles, it is possible that the results above arise as an artifact of averaging over entire days, and that during given parts of the day different behaviour is seen. In order to address this possibility, we split the day into 3 hour bins. For each inter-contact gap in the experiments, we add the duration to the bin in which the gap starts. Results are shown in Figure 7.

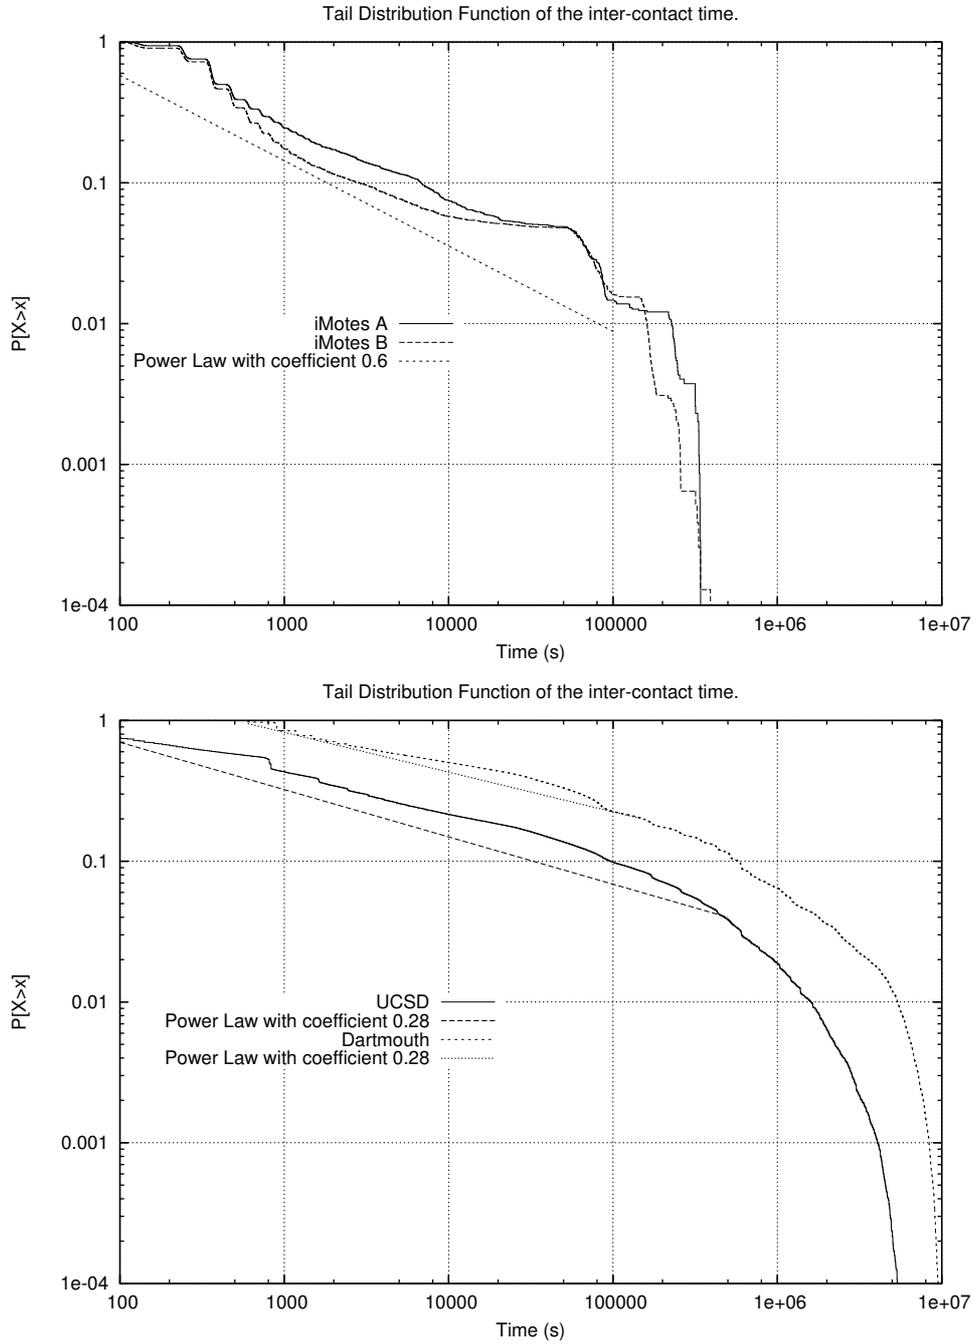


Figure 5: Distribution of inter-contact time for our four experimental data sets ; our Bluetooth-based experiments on the top, WiFi-based experiments on the bottom.

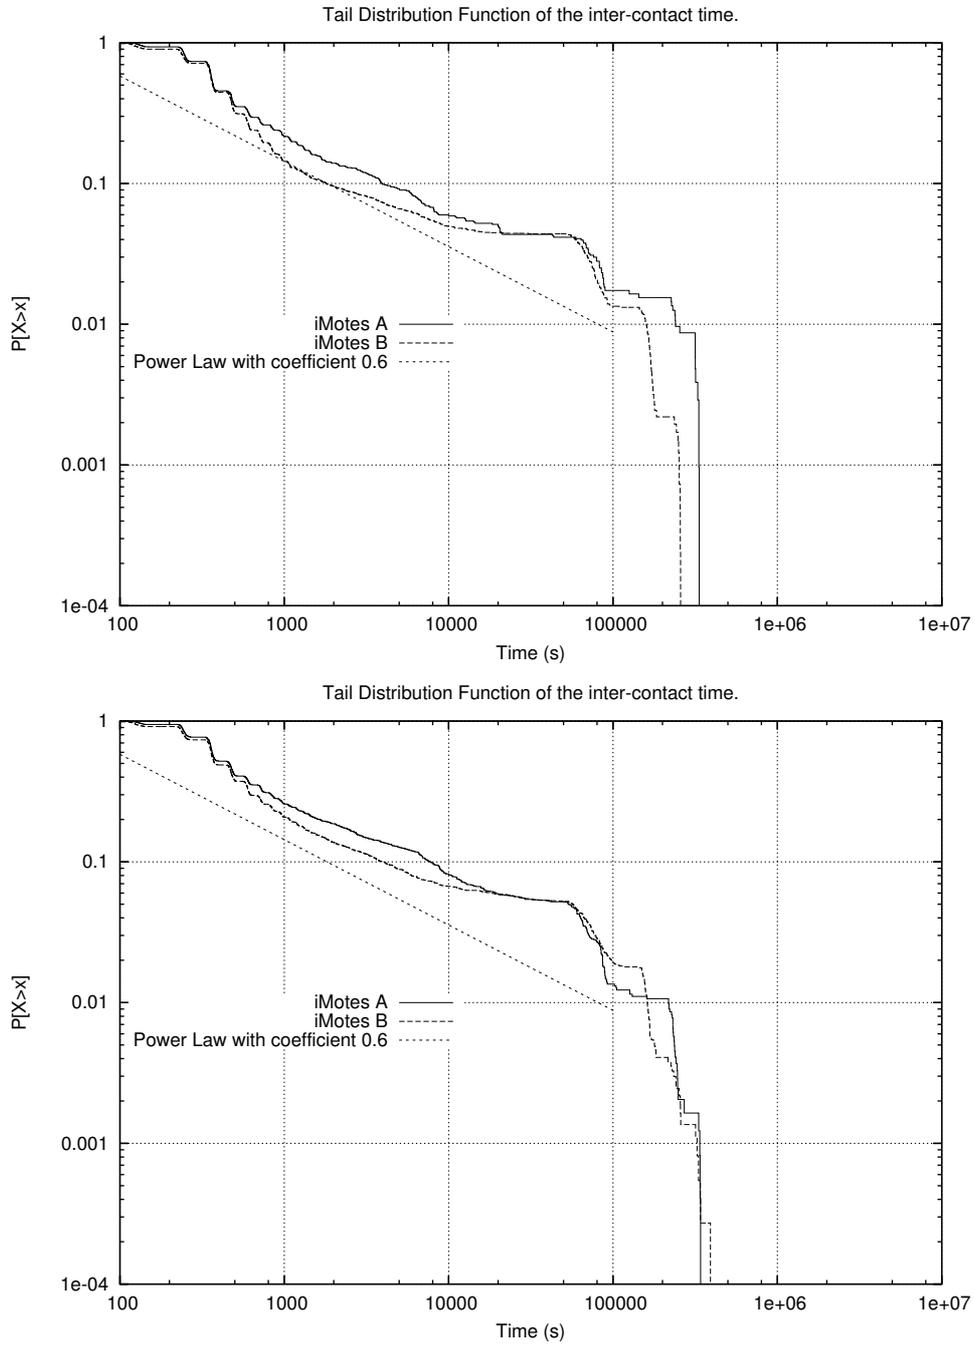


Figure 6: Distribution of inter-contact times between iMotes and other iMotes (top), and between iMotes and other Bluetooth devices (bottom).

It can be observed on Figure 7 (top) that the t.d.f. is becoming flat up to value of the order of a day for 15% or 10% of intercontact, at period 15:00-18:00 and 18:00-21:00. This is explained by the proximity of a period of inactivity, immediately after, for evening and night, that increases the proportion of large values of inter-contact times starting in this period.

As Figure 7 shows, the inter-contact time exhibits an approximate power law in each of the time bins, with the only difference being the coefficient of the power law, which is easily explained by the human diurnal cycle.

4.4 Contact duration

The contact duration is one of the factors affecting the amount of data that can be transferred between nodes when they come into range. Other relevant factors include relative distance patterns, the speed of movement, the discovery latency, and wireless network congestion. While the capacity of an opportunistic contact is obviously an important topic for opportunistic networking, it is not the focus of this paper. Nevertheless, for completeness we discuss some aspects of contact durations.

Figure 8 shows the distributions of contact durations in all four data sets. The maximum contact duration is around 3 hours in the iMote experiments, 3 days in the Dartmouth data set, and around half a day in the UCSD data. The latter two can be explained by the experimental methodologies; the UCSD PDAs were only able to last 8 hours before recharging, while the Dartmouth data may reflect laptops left in offices. 50% of the contact last less than 6 minutes for iMotes experiments, when only 20% of the contact last less than this in UCSD traces, and around 30% for traces from Dartmouth. The most important result is that contact time distributions in these cases seem to belong to the same category that inter-contact time. They are close to a power law on a range of value. But the range of value is much more narrow, and the coefficient of the power law is higher (around 1.5 for the Bluetooth-based datasets, and 1.8 to 2 for the WiFi datasets).

5 Networking with power law-based opportunities

In this section, we study the impact of heavy tailed inter-contact times on the actual performance and theoretical limits of a general class of opportunistic forwarding algorithms. These algorithms can be characterized as “stateless” in that they do not maintain history data or attempt to predict performance in the future. Instead, each node takes advantage of contact opportunities to forward as many packets as possible while in contact with another node. We start by defining formally the class of algorithms we study.

5.1 Opportunistic forwarding algorithms

We are interested in a general class of opportunistic forwarding algorithms that relies on intermediate nodes to carry data between a source and a destination that might not be contemporaneously connected. Intermediate nodes are chosen purely based on contact opportunism and not using any stored routing information.

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

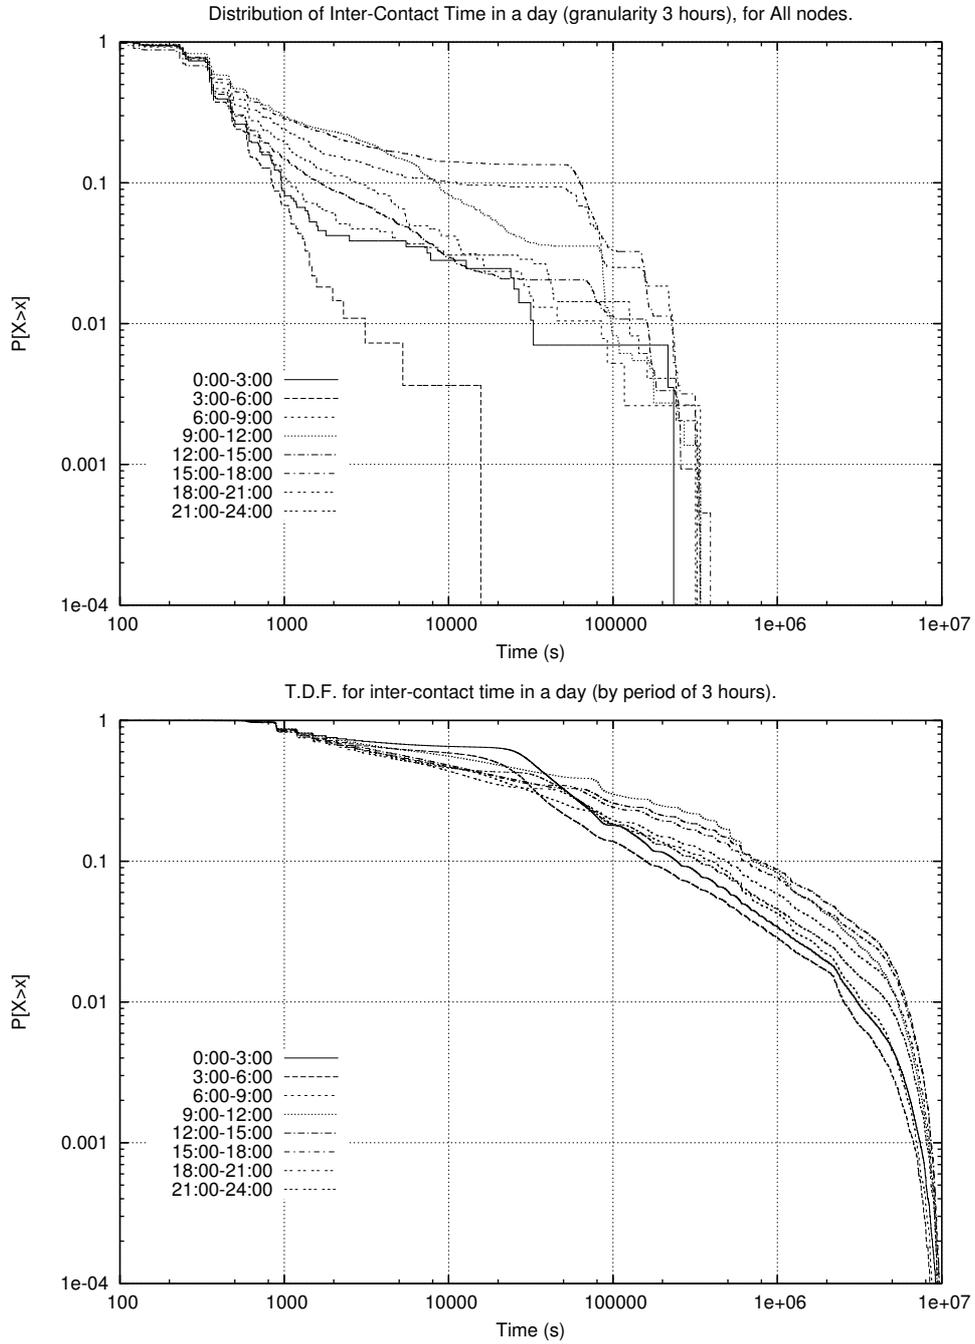


Figure 7: Distribution of inter-contact times depending on the time of the day, for Bluetooth-based experiment data sets (top) and WiFi-based experiment (bottom) data sets.

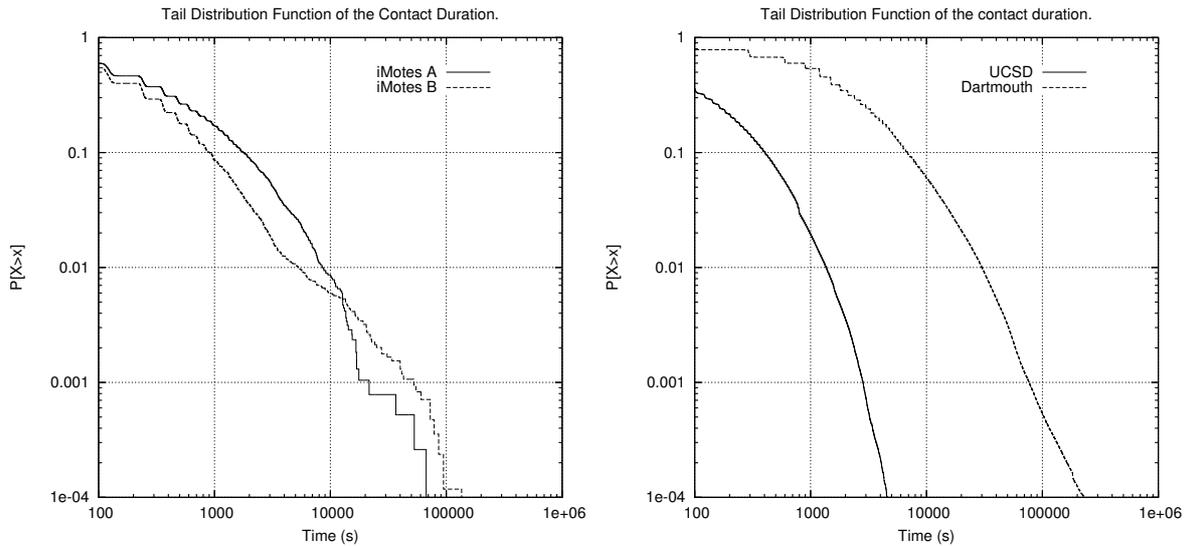


Figure 8: Distribution of contact duration for Bluetooth (left) and WiFi (right) data sets.

- wait-and-forward : The source simply waits until its next contact with the destination to communicate.
- flooding : a node forwards all its packets to any node which it encounters, keeping copies 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 delivers packets with the minimum possible latency, but does not scale well in terms of bandwidth, storage, and battery usage.

In between these two extreme cases, there is a whole set of middle-ground algorithms that use various numbers of intermediate nodes, contacts, and packet duplicates. Starting from the most conservative algorithms (wait-and-forward), we analyse incrementally the capacity of our class of algorithms, in the face of power law behaviour of the inter-contact delay distribution.

We first study the two-hop relaying algorithm introduced by Grossglauser and Tse in [8]. The two-hop relaying scheme operates as follows. Time is divided in a sequence of even and odd slots. During every even time slot, packets are sent from sources to relays, with the first available contact chosen as a relay. The only exception is if the destination is the first node to come into contact with the source, as the packet is transmitted directly. Otherwise, the relay will keep the packet in memory and the source does not send this packet again. Relays deliver to the destination only; going through a second relay is not permitted.

Each relay maintains one queue per destination. We assume that buffers are unlimited, and that packets in each queue are transmitted to the destination in a first-come first-served fashion, at the next time that the relay encounters the destination, in an odd time slot. As queuing is used in the intermediate nodes, the forwarding process of packets sent by the source to a relay needs to be of lower intensity than the packets sent by this relay to the destination. This is the case in the implementation proposed in [8] and we make the same assumption below.

This algorithm is a good candidate to start our study of the impact of power law inter-contact times on opportunistic forwarding for the following three reasons :

- The algorithm was shown to maximize the capacity of dense ad-hoc networks, under the condition that nodes are i.i.d. uniformly in a bounded region.
- This result depends strongly on the mobility process of nodes. Authors of [8] assumed an exponential decay of the inter-contact time. The same result has been proven for packets following Brownian motion or random waypoint mobility model [9].
- [8] and [9] have shown that that the packets experience a finite expected delay under these conditions.

5.2 Analysis of the two-hop relaying algorithm

In this section we first construct a communication model based on the following assumption : Inter-contact times between any pair of nodes follow an unbounded random variable with distribution X according to a power-law with coefficient $\alpha > 0$. With a discrete slotted time t ,

$$P[X \geq t] = (\max(1, t))^{-\alpha} \text{ for all } t = 1, 2, \dots \quad (2)$$

Later in this section, we will discuss the more general case where the inter-contact time is bounded, and where its distribution is subject to a lower bound by a power law over a restricted interval.

5.2.1 The unbounded power law case

We consider N nodes which transmit packets according to the two-hop relaying algorithm described above. Instead of the mobility model used in [8], we introduce a communication model where a node's location is ignored. In this model the times at which two nodes can exchange data is modeled as a renewal point process of events. The time between two events is distributed according to the random variable X already described.

Our goal is to evaluate the transmission delay associated to a packet. We therefore assume that all contacts last one time slot and that inter-contact periods last a variable number of slots. This simplification is supported by our experimental results, in which all data sets exhibited contact times several orders of magnitude less than inter-contact times (see Section 4).

To simulate the fact that a contact may last longer than one slot, we introduce two cases :

- the short contact case : where only one packet can be sent between nodes for each time slot where they are in contact.
- the long contact case : where two nodes in contact can exchange an arbitrary amount of data.

These two cases represent lower and upper bounds on the amount of data that can be exchanged during a contact period.

We define the set of packets transmitted in a single contact as a data bundle. In the long contact case, a data bundle contains an arbitrary number of packets. In the short contact case, a data bundle is comprised of a single packet. Bundles are treated differently by the relay in the two contact cases. For long contacts, a relay can empty its entire queue every time a destination is in contact. In the short contact case, packets are sent one at a time whenever the destination is in contact. Therefore, multiple contacts may be needed to empty the queue.

In the short contact case, since we need to ensure stability in the relay's queuing mechanism, we assume that bundles are created at the source at times following a point process with finite intensity λ (which we also call the demand intensity). The same assumption can be made for the long contact case although stability is not an issue in this context.

We introduce for each packet/bundle k , the delay D_k it experiences from its creation time to the time where it is received by the destination. We also introduce $M(t)$, the number of packets/bundles received by the destination at time t .

A consequence of the waiting time paradox is :

Theorem 1 *For a source with a positive demand intensity λ , transmitting information using the two-hop relaying algorithm, we have :*

- *For $\alpha < 2$, the sequence of delays experienced by packets/bundles diverges with expectation $\lim \mathbb{E}[D_k] = +\infty$, and if, in addition, $\alpha > 1$ in the Cesaro sense :*

$$\lim_{k \rightarrow \infty} \frac{D_1 + D_2 + \dots + D_k}{k} = \infty \text{ a.s.}$$

- *For the long contact case and $\alpha > 2$, a stationary regime exists where the delay of a bundle transmitted from a source to a destination has a finite expected value. The sequence of delay converges almost surely in the Cesaro sense.*
- *For the short contact case and $\alpha > 3$, if $\lambda < \frac{N-1}{\mathbb{E}[X]}$, a stationary regime exists where the delay of a bundle transmitted from a source to a destination has a finite expected value.*

Proof : The first half of the proof deals with the long contact case. The result for short contacts can then be deduced by increasing stochastic comparison [14], for any positive traffic demand λ .

long contact, for $1 < \alpha < 2$: We consider a source communicating with a relay at times $(S_k)_{k \geq 1}$; the times when the relay can communicate with the destination will be denoted by $(R_k)_{k \geq 1}$.

$$\begin{cases} R_0 = 0 \\ R_k = R_{k-1} + X_k \quad \text{for } k = 1, 2, \dots \end{cases},$$

where $(X_k)_{k \geq 1}$ are i.i.d. according to X . We have the same form for the sequence $(S_k)_{k \geq 1}$, with the inter-contact sequence $(Y_k)_{k \geq 1}$ independent according to X . We introduce for any time t the remaining time F_t (resp. G_t) until the next source-relay contact (resp. relay-destination). Note that F_t is also the remaining delay that will be experienced by a bundle that arrives in the relay at time t . It is a well-known result that both are Markov chains such that :

$$\begin{cases} F_{t+1} = F_t - 1 & \text{if } F_t \geq 1 \\ F_{t+1} = j \text{ with prob. } \mathbb{P}[X = j + 1] & \text{if } F_t = 0 \end{cases},$$

This Markov chain is irreducible, aperiodic, recurrent (as X is a.s. finite) and positive (as $\mathbb{E}[T_0] = \mathbb{E}[X] < +\infty$). From the regenerative form of invariant measure one can see that its stationary distribution π is given by $\frac{(i+1)^{-\alpha}}{\sum_{i \geq 0} (i+1)^{-\alpha}}$.

The key argument is that the product chain $(F_t, G_t)_{t \geq 1}$ is irreducible and admits a stationary measure (given by the product of π), hence it is ergodic and we have almost surely :

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T f(F_t, G_t) = \sum_{i \geq 0, j \geq 0} \pi(i, j) f(i, j) \quad (3)$$

for any function f integrable with the distribution π .

In particular the LHS for $f^{(A)} : (i, j) \mapsto \min(i, A) \mathbb{I}_{\{j=0\}}$ is a.s. lower bounded by $\text{Cst} \sum_{i=0}^A \frac{i}{(i+1)^{-\alpha}}$, which diverges for $A \rightarrow \infty$. Hence we can prove by comparison for all A :

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T F_t \mathbb{I}_{\{G_t=0\}} = +\infty$$

such that by taking $T = S_k$ we have the Cesaro divergence :

$$\lim \frac{F_{S_1} + \dots + F_{S_k}}{S_k} = \lim \frac{D_1 + \dots + D_k}{k \cdot \mathbb{E}[X]} = +\infty$$

To prove that the expectation diverges, we have for any $A \geq 1$:

$$\frac{A(A+1)}{2} \mathbb{I}_{\{X_1 \geq A\}} \leq \sum_{t=0}^{X_1-1} \min(F_t, A) \leq A \cdot X_1,$$

where all these variables are positive with a finite expectation. Then, as $(R_k)_{k \geq 1}$ is a regenerative sequence for $(F_t)_{t \geq 1}$, the regenerative theorem tells us :

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}[F_t] &\geq \lim_{t \rightarrow \infty} \mathbb{E}[\min(F_t, A)] \\ &\geq \frac{\frac{A(A+1)}{2} \mathbb{P}[X_1 \geq A]}{\mathbb{E}[X_1]} \geq \frac{A^2 \cdot A^{-\alpha}}{2 \cdot \mathbb{E}[X_1]} \end{aligned}$$

$$\text{hence } \frac{\mathbb{E}[M(t)]}{t} \rightarrow 0 \text{ for } t \rightarrow \infty . \quad (4)$$

using $M(S_k + D_k)/S_k + D_k = k/S_k + D_k$.

long contact case for $0 < \alpha < 1$: We first prove that the expected delay diverges. In this case, as $\mathbb{E}[X] = \infty$, the Markov chain defining $(F_t)_{t \geq 1}$ is recurrent null, so that Orey's theorem [13] tells us :

$$\lim_{t \rightarrow \infty} \mathbb{P}[F_t = i] = 0 \text{ for all } i$$

In particular, for any A ,

$$\lim_{t \rightarrow \infty} \mathbb{P}[F_t < A] = 0 \text{ hence } \mathbb{P}[F_t \geq A] \rightarrow_{t \rightarrow \infty} 1$$

This implies : $\mathbb{E}[F_t] \geq \mathbb{E}[F_t \mathbb{I}_{\{F_t \geq A\}}] \geq A$ which proves that $\lim_{k \rightarrow \infty} \mathbb{E}[D_k] = \lim_{k \rightarrow \infty} \mathbb{E}[F_{S_k}] = \lim_{t \rightarrow \infty} \mathbb{E}[F_t] = +\infty$.

long contact case for $\alpha > 2$: For such values of the parameter α , the product Markov chain (F_t, G_t) is positive recurrent (see above). In addition, the function $f : (i, j) \mapsto i \mathbb{I}_{\{j=0\}}$ is integrable w.r.t. the stationary distribution π as defined below, hence Equation (3) proves :

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T F_t \mathbb{I}_{\{G_t=0\}} = \sum_{i \geq 1} i \cdot \pi(i) \cdot \pi(0) < +\infty \text{ a.s.}$$

This proves an almost certain convergence in the Cesaro sense for the sequence of delay D_k .

The regenerative theorem, applied to the process F_t proves that :

$$\lim_{t \rightarrow \infty} \mathbb{E}[F_t] = \frac{\sum_{i=1}^{X_1} F_t}{\mathbb{E}[X]} = \frac{\mathbb{E}[X^2] + \mathbb{E}[X]}{2\mathbb{E}[X]} < +\infty$$

this proves the convergence in expectation of the delay, as $D_k = F_{S_k}$.

short contact case for $\alpha > 3$: We will use the same notation as above and consider one queue in a given relay. We assume the point process S_k and R_k to be in a steady state. What is now different is that the delay of packet k is not given by F_{S_k} , as existing remaining packets may block the new one for some time before it is transmitted to the destination. Another difference is that bundles are created in the source according to a point process with intensity λ . We will first prove that packets arriving in a relay through a point process with intensity $\frac{\lambda}{N-1}$ remains for a time with finite expectation.

We introduce $W(t)$ the remaining load in the relays for this destination as seen immediately after time t . When a packet arrived and the relay contains $m \geq 0$ packets to be sent $W(t)$ is given by

$$\inf \{R_{k+m}, k \geq 1 \text{ such that } R_k > t\} - t$$

For this case, a method similar to the one used in [9] can be used. When a packet enters the relay, it adds an additional service request to the value of $W(t)$. There is two cases to distinguish: either a packet enters a queue that is empty, in which case the additional service request is of the same distribution as F_t , or it adds to a number of packets waiting in line already. In the latter case, the additional service request is with law X . As can be seen from the expression of F_t , its distribution is always greater, for the stochastic increasing order, than the one of X . Hence the load in this queuing system is stochastically smaller than the load in a single server queue with clients all requesting independent service with law F_t .

As we have assumed $\alpha > 3$, the second moment of the law of F_t is finite (as can be seen from the expression of π), it is then possible to use Kingman's bound [15] to show that under the stability condition (verified as $\frac{\lambda}{N-1} < \frac{1}{\mathbb{E}[X]}$), the expected load in the corresponding single server queue is finite. By comparison, the expected load in the original queueing system is finite. Note that $W(t)$ is also the time at which a packet arriving at time t leaves the queue, this proves that the sojourn time of a packet in the relay is of finite expected value. The proof is exactly the same to show that the delay from packets created at the source with intensity λ stay a delay with a finite expectation in the source before being sent to the relay. The departure from the source to a relay can be shown to be a point process with intensity $\frac{\lambda}{N-1}$, which complete the proof.

We note here that the same result is likely to hold under the same stability condition for the more general case $\alpha > 2$. One way to look at this problem is through the framework of queues in a random environment (see [14]). This case is the subject of ongoing research. \square

Corollary 1 *As a consequence from the proof :*

$$\text{if } \alpha < 2, \lim_{t \rightarrow \infty} \frac{\mathbb{E}[M(t)]}{t} = 0$$

If $\alpha > 2$ (or $\alpha > 3$ for the short contact case),

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}[M(t)]}{t} = \min\left(\frac{N-1}{\mathbb{E}[X]}, \lambda\right) > 0$$

To summarize, we have identified two regions where the behavior of the two-hop forwarding algorithm would differ under the power law inter-contact time assumption : For a value of α that is greater than 2 (or 3 if contacts are short) the algorithm converges, as in the case of an exponential decay, to a finite expected delay. For smaller value, it will exhibit a very different performance, as delay will grow unboundedly with time, and throughput will vanish in expectation.

5.2.2 Restricted power law

We have shown above that the two-hop relaying algorithm stabilizes, under the assumption (2) of unbounded inter-contact time following a power law, if and only if the coefficient of the power law is greater than 2. In particular, for smaller values of the coefficient α , which is what we observe in all four experimental datasets, the empirical throughput may vanish with time, and the average delay of successive packets sent will grow to infinity.

In practice, the consequence of this result is that the delay observed by a source in this forwarding scheme will almost surely diverge out of the range of inter-contact times for which a power law applies with coefficient smaller than 2.

In other words, the probability that a packet is blocked in one of the very long inter-contact times becomes larger with time. In our empirical data sets, the expected delay would not be close to moderate values. It would be dominated instead by the extreme values found in the power-law region. For the Bluetooth experiments, this is on the order of days, while for WiFi this is in weeks.

This remains true even in the optimistic cases of unlimited bulk exchange between nodes, with no queuing expected in the relay, and the best possible use of contact opportunities between the chosen relay and the destination.

On the other hand, the situation is not so pessimistic if a power law applies with a stronger coefficient (greater than 3). In this case, the two-hop relaying algorithm converges for a stable demand to the stationary regime of a queuing system, with finite expected delay. A positive throughput is guaranteed to the source, that does not necessarily depend on the size of the range where this power law applies.

5.3 Generalization to stateless algorithm

We have seen that the performance of the two-hop relaying algorithm strongly depends on the value of the parameter α which controls the distribution of large inter-contact times. In this section we prove that the same results apply more generally to any stateless forwarding algorithm. Depending on the value of α , the delay might have bounded expectation for all forwarding algorithms, some, or for none of them. This is important since different mobility conditions lead to different values of α ; for example, Section 4.3 showed the variation at different times of the day.

We show that, within the class of stateless forwarding algorithms, there is a subclass characterized by multiple transmissions to different relay nodes that outperforms the two-hop algorithm with regard to stability in the long contact case.

Instead of sending a single copy of a message to a relay, the source will send each packet to the first m relays opportunistically met. Provided that the contacts of these relays are independent, the source may reduce the transmission delay by increasing its probability to pick a relay

with a small delay to the destination among the N relays to which it has forwarded the message. We study this subclass of algorithm for $N = 2$.

Lemma 1 *Let $(F_t)_{t \geq 1}$ and $(F'_t)_{t \geq 1}$ the independent remaining time of a renewal process, with inter-event following the distribution of Equation (2) for $\frac{3}{2} < \alpha < 2$.*

$$\begin{aligned} \text{Then } \mathbb{E}[F_t] &= \mathbb{E}[F'_t] = +\infty \\ \text{and } \mathbb{E}[\min(F_t, F'_t)] &< \infty. \end{aligned}$$

Proof : The stationary distribution exists as $\alpha > 1$, it extends naturally to the product chain, where it is given by

$$\begin{aligned} \pi(i, j) &= \text{Cst} \frac{1}{(i+1)^{-\alpha}(j+1)^{-\alpha}} \text{ such that} \\ \mathbb{E}[\min(F_t, F'_t)] &= \text{Cst} \sum_{i,j} \frac{\min(i, j)}{(i+1)^{-\alpha}(j+1)^{-\alpha}} \end{aligned}$$

This series is convergent if $\alpha > \frac{3}{2}$, as :

$$\begin{aligned} \sum_{i,j} \frac{\min(i,j)}{i^\alpha j^\alpha} &= 2 \cdot \sum_{i,j} \frac{\min(i,j)}{i^\alpha j^\alpha} \mathbb{I}_{\{i \leq j\}} \\ &= 2 \cdot \sum_j \frac{1}{j^\alpha} \sum_{i=1}^j i^{1-\alpha} \\ &\leq 2 \cdot \sum_j \frac{1}{j^\alpha} (\text{Cst} + j^{2-\alpha}) \\ &\leq 2 \left(\sum_j \frac{1}{j^\alpha} \text{Cst} + \sum_j j^{2-2\alpha} \right). \end{aligned}$$

□

This result shows that, for certain values of the parameter α , the expected time to get in contact with a given node is infinite, but the delay before contacting one node in a given set may have a finite expected value. More precisely, for $2 > \alpha > \frac{3}{2}$, any set containing at least two nodes has this property.

The same result can be extended for an arbitrary number of relays m :

Lemma 2 *Let $(F_t^{(1)})_{t \geq 1}, \dots, (F_t^{(m)})_{t \geq 1}$ be independent remaining times of a renewal process, with inter-event distributed according to Equation (2), for $\frac{m+1}{m} < \alpha < 2$.*

$$\begin{aligned} \text{Then } \mathbb{E}[F_t^{(1)}] &= \dots = \mathbb{E}[F_t^{(m)}] = +\infty \\ \text{and } \mathbb{E}[\min(F_t^{(1)}, \dots, F_t^{(m)})] &< \infty. \end{aligned}$$

Proof : We can prove using the same argument than previous lemma :

$$\mathbb{E}[\min(F_t^{(1)}, \dots, F_t^{(m)})] = \text{Cst} \sum_{i_1, \dots, i_m} \frac{\min(i_1, \dots, i_m)}{(i_1+1)^{-\alpha} \dots (i_m+1)^{-\alpha}},$$

such that in particular it is finite if we have $f(m, \alpha, 1) < \infty$ for

$$f(m, \alpha, \beta) = \sum_{i_1, \dots, i_m} \frac{(\min(i_1, \dots, i_m))^\beta}{(i_1)^{-\alpha} \dots (i_m)^{-\alpha}}.$$

We will prove that if $(m + 1)/m < \alpha$, and $\beta \leq 1$, then $f(m, \alpha, \beta) < \infty$.

For $m = 1$, this is true as for $\alpha > 2$ and $\beta \leq 1$, $f(1, \alpha, \beta) = \sum_i i^{\beta-\alpha} < +\infty$.

For $m = 2$, the result is an easy generalization from the previous lemma, and we have more generally :

$$\begin{aligned}
f(m, \alpha, \beta) &\leq m. \sum_{i_2, \dots, i_m} \left(\sum_{i_1=1}^{\min(i_2, \dots, i_m)} (i_1)^{\beta-\alpha} \right) \frac{1}{(i_2)^{-\alpha} \dots (i_m)^{-\alpha}} \\
&\leq m. \sum_{i_2, \dots, i_m} \frac{\text{Cst} + \text{Cst}(\min(i_2, \dots, i_m))^{\beta+1-\alpha}}{(i_2)^{-\alpha} \dots (i_m)^{-\alpha}} \\
&\leq \text{Cst} + \text{Cst} f(m-1, \alpha, \beta+1-\alpha) \\
&\leq \dots \\
&\leq \text{Cst} + \text{Cst} \underbrace{f(1, \alpha, \beta + (m-1)(1-\alpha))}_{=\sum_i i^{\beta+(m-1)(1-\alpha)-\alpha}}.
\end{aligned}$$

The result follow as

$$\beta + (m-1)(1-\alpha) - \alpha \leq -1 + \underbrace{m+1 - m\alpha}_{<0}.$$

□

This observation is the key component in the next result, which proves that such a method can be used with two relays to build a stable algorithm with $\alpha > \frac{3}{2}$, in the long contact case. More generally we can construct a method using m duplicate copies that is stable under $\alpha > \frac{m+1}{m}$. We also need to make assumption on the size of the network.

The next result also sets out a theoretical bound on the application of stateless forwarding algorithms to transport packets through contact opportunities with power law properties. As we will see, for inter-contact times distributed with $\alpha < 1$, no algorithm based on opportunistic contacts, including unlimited flooding, can support stable (i.e. bounded-delay) communication.

Theorem 2 *In an opportunistic network, for the long contact case, if we assume that the contacts for each pair (r, s) of nodes follow an independent stationary renewal process with the distribution of Equation (2) with parameter $\alpha_{r,s} \geq \alpha$.*

- if $\alpha > 2$, there is a forwarding protocol using only one copy of the packet, with a finite expected delay.
- if $\alpha > \frac{3}{2}$, and we suppose that the network contains at least $N = 4$ nodes, there is a forwarding protocol using 2 duplicate copies, such that the expected delay is finite.
- More generally, if $\alpha > 1$, and if we suppose that the network contains at least $N > \frac{2}{\alpha-1}$ nodes there exists an algorithm using m duplicate copies with $m > \frac{1}{\alpha-1}$, that is achieving a finite expected delay.
- if $\alpha_{r,s} \leq 1$ for all pairs (r,s) , and we suppose a finite network with N nodes, no forwarding protocol (with any number of copies, including flooding) can achieve a finite expected delay.

Proof : First result : $\alpha > 2$: This is only a reminder for long contacts of the result shown in the previous section, the algorithm chosen is the two-hop relaying strategy.

Second result : $1 < \alpha < 2$: We assume the same coefficient $\alpha_{r,s}$ for every pair of nodes. The result can be deduced in the general case by an increasing stochastic comparison. Let's treat the case $\alpha > 3/2$. The forwarding algorithm that we will consider is a two-hop relaying using two relays.

STEP 1 : The packet k is transmitted by the source (denoted as node 1) to the first two intermediate nodes that are met. We first estimate the time for which each of these two relays will be both contacted and will receive the packet.

For a given node r , the inter-contact time of the pair $(1, r)$ follows a power law distribution $P(X \geq u) = (\max(1, u))^{-\alpha}$. In particular, the remaining time $F_t^{(1,r)}$ until the next contact is an irreducible and positive recurrent Markov Chain, with stationary distribution $\pi(i) = \frac{(i+1)^{-\alpha}}{\sum_{i \geq 0} (i+1)^{-\alpha}}$ (see the proof of Theorem 1). In particular, as there are $N - 1$ nodes from which the source can choose two relays from, they will be chosen as the two minimum coordinates of the random vector $(F_t^{(1,2)}, \dots, F_t^{(1,N)})$. Let us denote them by $(\check{F}_t^{(1,r_1)}, \check{F}_t^{(1,r_2)})$, sorted increasingly. As there are at least $N - 1 \geq 3$ nodes, $\check{F}_t^{(1,r_2)}$ is the result of the minimum of two independent Markov Chain with law F_t , hence it has a finite expectation by the lemma.

STEP 2 : At time $t + F_t^{(1,r_2)}$, a copy of the packet is present in each of this two relays. The remaining time for each of them to get in contact with the destination follows the same distribution that $F_t^{(1,r)}$. In particular the minimum of these two independent variables has a finite expectation. The total latency since the packet is created in t is then the sum of two variables with finite expectation.

GENERALIZATION TO ANY m : Let us assume $\alpha > 1$ and $m > 1/(\alpha - 1)$. The same proof can be applied more generally to a forwarding algorithm using m duplicate copies.

A source forwards a copy of the packet to each of the first m relays that it meets. We assume that the network contains at least $N \geq 2m$ nodes, hence the successive waiting time to meet the corresponding relays is of finite expected value : The first relay met is chosen among $N - 1$, and even the m -th relay met by the source is chosen as the first relay met of the $N - 1 - (m - 1) \geq m$ remaining.

When m relays receive the packet, the time-to-delivery is the minimum of the times before the destination meet each of the set m , which has been shown to have a finite expected value.

Third result : $\alpha < 1$: Let us consider in this case, for all pairs (r, s) of nodes in the network, the remaining time $F_t^{(s,r)}$ a time t until the next contact. They are all irreducible null recurrent Markov chains. By Orey's theorem, we then have that $\lim P(F_t^{(s,r)} = i) = 0$ for all i when t tends to infinity.

In particular $\lim P(F_t^{(s,r)} \leq M) = 0$, so that $\lim \bigcap_{(r,s)} P(F_t^{(s,r)} \geq M) = 1$.

Hence $\lim P(\min_{(r,s)} F_t^{(s,r)} \geq M) = 1$ and $\mathbb{E} \left[\min_{(r,s)} F_t^{(s,r)} \right]$ diverges for large t .

As a consequence, when two nodes are in contact in such a network, the expectation of the time until any of them is in contact with any other nodes is diverging with time. No forwarding algorithm, no matter how redundant, can then transport a packet within a finite expected delay. \square

The second result, which demonstrates the improvement introduced by two redundant copies of the same packet, relies on the set of chosen relays being independently chosen with respect to their contact patterns. Although it seems to be an assumption that is accepted in the related literature on opportunistic communications, it may be the most difficult to achieve in a practical

manner. Extension of this result to nodes with lightly correlated contact patterns, and heuristics to choose a reasonable set of relays in practical cases, are appealing future directions for research.

At this stage, we have established the following results for the class of stateless forwarding algorithm defined in 5.1, in the long contact case :

- If $\alpha < 1$, none of these algorithms can achieve a transmission delay with a finite expectation.
- If $1 < \alpha < 2$, the algorithm introduced by [8] is not stable as the delay has an infinite expectation. It is however possible to build a forwarding algorithm that achieves a bounded delay, using a number of duplicate copies of the packet. The number of necessary copies m used must be greater than $\frac{1}{1-\alpha}$ and the network must contain at least $\frac{2}{1-\alpha}$ nodes.
- For $\alpha > 2$ any algorithm from our class converges.

Note that by comparison the flooding algorithm can achieve a forwarding with bounded delay in all the cases $\alpha > 1$, if the network is large enough. Using such an aggressive forwarding scheme is however never necessary to have a bounded delay ; for any value of $\alpha > 1$ there exists an algorithm using only a fixed number of copies m that does not depend on network's size.

All the divergence results apply to the short contact case as well, as it is more constrained than the long contact case. It is not clear at this stage if redundant copy can improve the domain of stability of forwarding algorithm in the short contact case. This case is left to future investigation.

Therefore, in the cases of both our Bluetooth and WiFi datasets, there is no stateless opportunistic algorithm that can guarantee a transmission delay with finite expectation, since $\alpha < 1$. This poses a challenge for designers of opportunistic communication systems, which could be tackled by using hybrid opportunistic and infrastructure-based communications, or by creating stateful forwarding algorithms, which learn from network history, e.g. locating well-connected clusters of nodes for which a power law of inter-contact times may have a tractable coefficient.

The conclusion from this section is that forwarding algorithms designed for non-end-to-end connectivity can still take advantage of connection opportunities in the context of inter-contact times which follow a power law. However, certain conditions must be met to ensure stable communications. The strongest heavy tailed properties of inter-contact times, especially the ones that were exhibited by our Bluetooth experiments ($\alpha < 0.6$) and by WiFi traces ($\alpha < 0.33$), result in diverging delay for any stateless forwarding algorithm.

6 Summary, conclusion and future work

We study an ad-hoc network scenario, called opportunistic networking, where inherent mobility and the occasional connection with other devices are used to transfer data.

We establish a first major result, which is that in four different and independent data sets, the distribution of the “inter-contact time” between nodes in an opportunistic networking environment follows an *approximate power law* over a large range, with power law coefficient less than one. This result is not consistent with the exponential decay predicted by all existing node mobility models used to date in ad-hoc networking.

We show that a class of stateless forwarding algorithms, that have been proved to deliver packets with a bounded delay in the case of exponential decay inter-contact times, have indeed an infinite expected delay when mobility follows approximate power law inter-contact with coefficient under 1. We prove that using multiple intermediate relays is sufficient for these algorithms to converge when the power law coefficient is located between 1 and 2. Above two, these algorithms converge naturally.

The implications of our work for the research community are as follows:

1. Current mobility models (e.g. random waypoint, uniformly distributed locations) do not have the characteristics observed in our human mobility experiments. New mobility models are therefore required in order to facilitate evaluation of potential opportunistic data transmission schemes.
2. Little work has been done in the area of informed design of opportunistic forwarding algorithms — this remains an area ripe for study. Suitable directions for work might involve the sharing of recent contact information between nodes, leading to a more careful selection of potential relay nodes which are likely to have a short path to the destination, while also being independently moving as compared to other chosen relay nodes.

We plan to continue our research in a number of directions. Firstly, we wish to characterise the contacts patterns between nodes as well as the inter-contact patterns, looking at how contact duration, node speed, relative distance, discovery latency, and network congestion affect the capacity of an opportunistic contact for Bluetooth and WiFi networks. Secondly, we wish to explore tractable forwarding algorithms for opportunistic networks, taking into account the lessons above of avoiding long-delay paths and of sending redundant copies over independently-moving relays. Finally, we wish to examine a hybrid of opportunistic and infrastructure based networking, and the delay and bandwidth characteristics present in such conditions.

7 Acknowledgements

We would like to gratefully acknowledge Tristan Henderson at Dartmouth College and Geoff Voelker at UCSD for providing their data, Vincent Hummel and Ralph Kling at Intel for their extensive support with the iMote platform, Mark Crovella for his insightful comments on our findings, and the participants of our mobility study at Intel Research and Cambridge University.

References

- [1] R. Patra S. Jain, K. Fall. Routing in a delay tolerant network. In *Proceedings of ACM SIGCOMM*, 2004.
- [2] E. Royer and C-K Toh. A review of current routing protocols for ad-hoc mobile wireless networks. *IEEE Personal Communications Magazine*, pages 46–55, 1999.
- [3] M. Balazinska and P. Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. In *1st International Conference on Mobile Systems, Applications, and Services (MobiSys)*, San Francisco, CA, May 2003.

- [4] M. McNett and G. M. Voelker. Access and mobility of wireless pda users. Technical report, Computer Science and Engineering, UC San Diego, 2004.
- [5] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of ACM Mobicom*, 2004.
- [6] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. *Wireless Communications and Mobile Computing: Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, 2(5):483–502, 2002.
- [7] A. Jardosh, E. M. Belding-Royer, K. C. Almeroth, and S. Suri. Towards realistic mobility models for mobile ad hoc networks. *MobiCom, San Diego, CA*, 2003.
- [8] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking*, 10(4):477–486, 2002.
- [9] G. Sharma and R. R. Mazumdar. Scaling laws for capacity and delay in wireless ad hoc networks with random mobility. In *Proceedings of ICC 2004*. IEEE, 2004.
- [10] J. Su, A. Chin, A. Popivanova, A. Goel, and E. de Lara. User mobility for opportunistic ad-hoc networking. In *Proceedings of WMCSA 2004*, 2004.
- [11] 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 zebrantet. In *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 96–107, 2002.
- [12] 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). In *Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 233–244, 2003.
- [13] P. Bremaud. *Markov Chains, Gibbs Field, Monte Carlo Simulation and Queues*. Springer-Verlag, 1999.
- [14] F. Baccelli and P. Bremaud. *Elements of Queuing Theory*. Springer-Verlag, second edition, 2003.
- [15] J. F. C. Kingman. Inequalities in the theory of queues. *J. Roy. Statistical Society, Series B*, 32:102–110, 1970.