Impact of Human Mobility on Opportunistic Forwarding Algorithms

Augustin Chaintreau‡, Pan Hui*, Jon Crowcroft*,
Christophe Diot‡, Richard Gass†, and James Scott†,

‡Thomson Research
46 quai A. Le Gallo
92648 Boulogne FRANCE
augustin.chaintreau@thomson.net
christophe.diot@thomson.net

*University of Cambridge
15 JJ Thomson Avenue
Cambridge, CB3 0FD, UK
pan.hui@cl.cam.ac.uk, jon.crowcroft@cl.cam.ac.uk

† Intel Research
richard.gass@intel.com, jamesscott@acm.org

***Abstract*— This paper studies data transfer opportunities between wireless devices carried by humans. We observe that the distribution of the inter-contact 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.**

I. 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 user's mobility. They form a Pocket 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 in 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 inter-contact time as the time between two transfer opportunities, for the same devices. We observe in the eight traces that the inter-contact time tail distribution is slowly varying over a large range. Inside this range the inter-contact time distribution can be compared to a power law. We study the impact of those large inter-contact times on the actual performance and theoretical limits of a general class of opportunistic forwarding algorithms that we call "naive forwarding algorithms". Algorithms in this class do not use the identities of the devices that are met, nor the recent history of contacts, or the time of the 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. It is based on 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 inter-contact times influences the performance of naive forwarding algorithms, and how these should be configured to offer reasonable guarantees.

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

II. Experimental Analysis

A. 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 carrying 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 twenty-four hours a day.

We examined two types of data sets. First, we use public traces measuring connectivity between clients and access points 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, the characteristics of each of them are summarized in Table I.

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

TABLE I

COMPARISON OF DATA COLLECTED IN THE EIGHT EXPERIMENTS.

1) AP-based data sets: UCSD [2] and Dartmouth [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 pre-processing. 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 MIT Media Lab include records of visible GSM cell towers, collected by 100 cellphones 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 assumption we have made for all these data sets introduces inaccuracies. On 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 co-located 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) 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 120s 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, end time). The experimental settings are described in detail in [1]; an anonymous version of our data is now available to other research groups on the CRAWDAD[2] archive.

We include in this paper the results from three iMote-based experiments. The first obtained data from twelve doctoral students and faculty comprising a research group at the University of Cambridge. The second experiment included a group of thirty seven 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: The University of Toronto have collected direct contact traces using 23 Bluetooth-enabled PDAs distributed to a group of students. These devices performed a Bluetooth inquiry each 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 participants, and 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.

B. 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 on 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 III), we will assume that all contacts last a single time slot and we will address two extreme cases

[1] See www.xbow.com

[2] See crawdad.cs.dartmouth.edu

corresponding to a lower and upper bounds of the amount of data that could be transferred in each connection opportunity.

We define the *inter-contact time* as the time elapsed between two successive contact periods for a given pair of devices. Inter-contact 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 inter-contact time is computed once at the end of each contact period, as 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 inter-contact time* seen at any time t, for each pair of devices: it is the time it takes after t, before a given pair of devices meet again (a formal definition is given in Section III). Inter-contact time and remaining inter-contact time have different distributions, which are related, for a renewal process, via a classical result known as the waiting time paradox (see p.147 in [5]). A similar relation holds for stationary processes, in the theory of Palm Calculus (see p.15 in [6]). We choose to study the first definition of "inter-contact time seen at the end of a contact period", as the second gives too much weight to large values of inter-contact times. In other words the definition we have chosen is the most conservative one in the presence of large values.

Second, the inter-contact time distribution is influenced by the experiment's duration and its granularity (i.e. the time elapsed between two successive scanning for the same device). Inter-contact times that last more than the duration of the experiment cannot be observed, and inter-contact times close to the duration are less likely to be observed. In a similar way, we cannot observe the inter-contact times that last less than the granularity of the measurement (which ranges from two to five minutes among 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, due to a lack of space.

C. Inter-contact time characterization

We study the empirical distribution of the inter-contact times obtained for all experiments shown in Figures 1 and 2.

For all plots, an empirical distribution of the inter-contact 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 that 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.

1) Aggregated distribution: Figure 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 two first iMote experiments (labeled Cambridge and Hong Kong) we present only one case here (corresponding respectively to internal and external contacts). They are shown in Figure 1 (left), 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 Figure 1 (middle), where distributions associated with internal and external contacts have been plotted separately for comparison. Figure 1 (right) presents the distribution of inter-contact computed using traces from other experiments than ours.

Let us first note that, although inter-contact are short in most cases, the occurrence of large inter-contacts is far from negligible: in the three iMote based experiment, 17 to 30% of inter-contact times are greater than one hour, and 3 to 7% of all inter-contacts are greater than one day. In the Toronto data sets, 14% of inter-contacts last more than a day, and 8% more than a week. These large inter-contacts are even more present in the traces collected in UCSD, Dartmouth and MIT, the most extreme case being the MIT trace using Bluetooth sightings, where up to 60% of the inter-contacts observed are above one 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 condition (granularity, duration). But they also present common properties that we now discuss in more details.

We now concentrate on the region between 10 minutes and one day. In this region, all data sets exhibit the same characteristics: the CCDF is slowly varying, it is lower bounded by the CCDF of a power law distribution, that may in some cases be a good approximation. This contradicts the exponential decay of the tail which characterizes the most common mobility models found in the literature (see Section IV), and we prove in the next section that it 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 Figure 2 (left) 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.

[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.

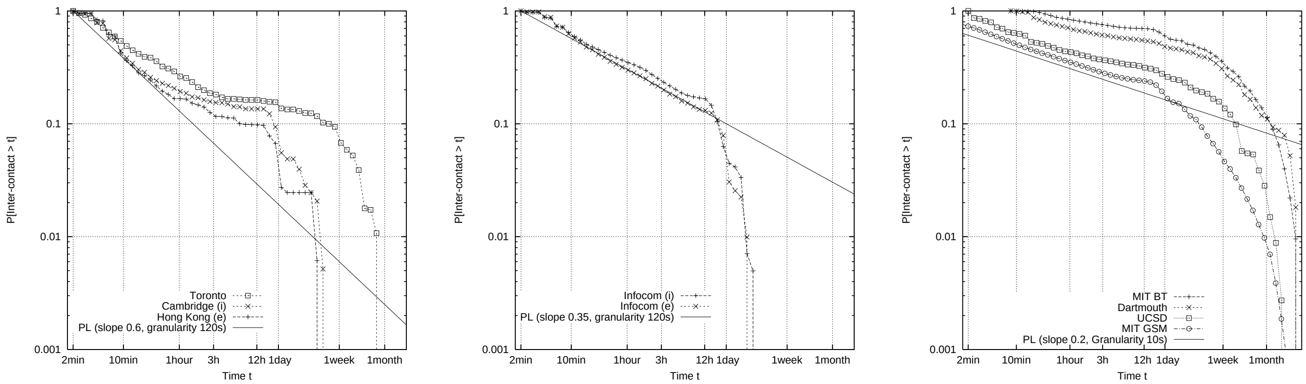

Fig. 1. Aggregated distribution of the inter-contact time in eight data sets experiments: iMote-based experiments at Cambridge and Hong Kong, and Toronto experiment (left), iMote-based experiment at INFOCOM (middle), data collected at UCSD, Dartmouth and MIT (right).

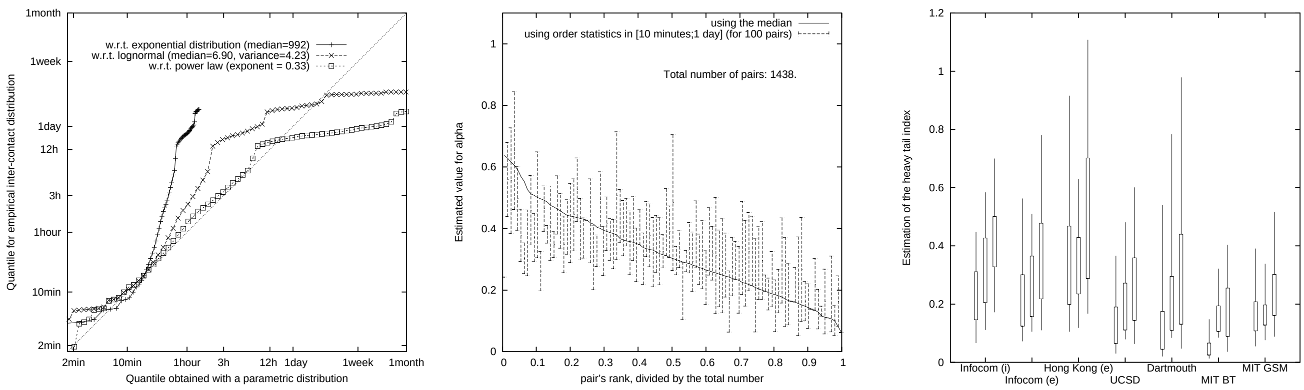

Fig. 2. iMote-based experiment at Infocom: Quantile-quantile plot of comparison between the aggregated distribution of the inter-contact time and three parametric models (left), estimation of the power law coefficient applied separately for each pair (middle), summary of results obtained in all data sets (right).

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, the quantile for the log-normal distribution deviates from the empirical case by a non negligible factor. The power law distribution, by opposition, remains close to the empirical one for values 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 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 Figure 1 (middle) and (right). We also observe that the coefficient 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 Toronto data sets, 0.35 for the iMote based experiment at Infocom, and 0.2 for traces collected in UCSD, Dartmouth and MIT. In all cases, it is below 1. The value of this coefficient, which is also called the "heavy tail index", is critical for the performance of opportunistic forwarding algorithms (see the analysis in Section III), and we discuss it further below.

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

2) Individual distribution 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 inter-contacts 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 120s 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 is given by:

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

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. It creates a collection of estimators for the value of α, as follows:

$$\left\{ \frac{\ln(n) - \ln(n-i)}{\ln(X_{(i)}) - \ln(120)} \;\middle|\; 600 \leq X_{(i)} \leq 86400,\ i < n \right\}.$$

We denote by $\underline{\alpha}$ and $\overline{\alpha}$ respectively the minimum and maximum value in this set above. It is equivalent to plot the empirical CCDF for this sample in a log-log scale, and bound this CCDF from above and below by two straight lines that go through probability 1 at time value 120s. These slopes would be equal respectively to $-\underline{\alpha}$ and $-\overline{\alpha}$. By opposition to $\check{\alpha}$, 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 Figure 2 (middle), we plot the values of $\check{\alpha}$ and the interval $[\underline{\alpha}, \overline{\alpha}]$ for all pairs of iMotes during the experiment conducted at Infocom. One can expect that the coefficient 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 $\check{\alpha}$, 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 Figure 2 (middle) 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 lies 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% of pairs in this case). The median based estimate lies in [0.2 ; 0.4] for half of the pairs, the lower estimates (resp. the upper estimate) lies in [0.14 ; 0.32] (resp. [0.32 ; 0.5]) again for half of the pairs.

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 inter-contact distribution for each pair is lower bounded in this range by a power law with a coefficient smaller than 1.

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

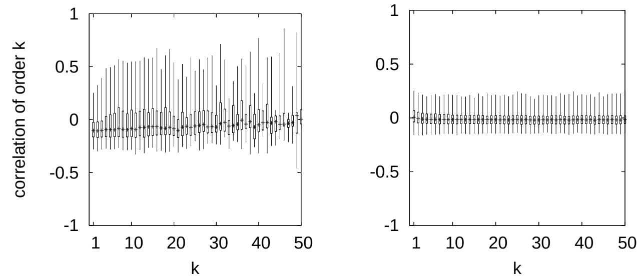

Fig. 3. Correlation coefficients for the sequence of inter-contact times: for all pairs of iMotes in the Infocom data set (left), for all pairs of devices in MIT GSM data set (right).

In the Hong Kong and Dartmouth data sets, where contacts are sparser, inter-contact 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% of the pairs in Hong Kong data set, although it is probably an artefact of our conservative estimate.

Correlation: in Figure 3 we show the auto-correlation coefficient for all order k up to 50. A correlation coefficient was computed for each pair, we present for all k the average value we observed among all pairs, as well as the interval containing 50% and 90% of the centered values (respectively, in the thick box and the thin bar).

In the Infocom data set, we observe a slightly negative correlation on average over all pairs, which reduces as k goes large. The variation of the coefficient among pairs is quite important, although most pairs remain reasonably non-correlated (the thick box remains always less than 0.15 away from zero). Correlation coefficients are smaller when the data set is large (as seen for example in the MIT GSM trace shown here, as well as for all other long traces). This tends to indicate that these coefficients for all pairs would be closer to zero if the iMote experiment could be done with a longer duration, and that the sample of inter-contacts collected for each pair was bigger.

In the next section, we are assuming that all inter-contacts in a sequence are independent (i.e. correlation coefficients are all taken as zero). 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, but that is left for future work.

III. Forwarding with Power Law–Based Opportunities

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

A. Assumptions and Forwarding Algorithms

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 inter-contact 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 of inter-contact 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 inter-contact time after the k_th contact of this pair. We suppose in our model that it has the same law as X, which follows a power law with coefficient $\alpha > 0$:

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

Note that X is not bounded but is almost surely finite. 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 inter-contact 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 V. Note that these assumptions are verified, or implicitly assumed, in most of the analyzes 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.
- 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 a lower and an upper bound 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 the relay device.

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

Note that our model does not take into account explicit geographical location 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 inter-contact 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 $\alpha > 1$.

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 "naive", although they could be in reality quite complex and, as we will see, very efficient in some cases.

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 reason not to flood is to minimize memory requirements and related power consumption in relay devices, and to delete the backlog of previously relayed message 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.

B. Analysis of the two-hop relaying algorithm

Having described the class of "naive" 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 inter-contacts that we have described. Results are generalized to the general class of naive algorithms in the following section.

1) Description: The two-hop relaying algorithm was introduced by Grossglauser and Tse in [10]. This forwarding algorithm operates as follow: when a source has a bundle to send to a destination, it forwards it once to the first devices 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 inter-contact 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. The authors of [10] assumed that each node moves at each time slot to an i.i.d. position, which implies that the inter-contact 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]).
- [10] and [11] have shown that data experiences a finite expected delay under these conditions.

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, that 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 k_th bundle is created at s to be sent to d, and let $t_k^{(d)}$ be the time when it is delivered to d. Let $D_k = t_k^{(d)} - t_k^{(s)}$, we have, starting from any initial condition:*

(i) *If $\alpha < 2$, $\lim_{k\to\infty}\mathbb{E}[D_k] = +\infty$.*

(ii) *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*

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

(iii) *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; that is 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 $t_k^{(s)}$. Let r_k be this relay, we have $r_k = \text{argmin}_{r'\neq s} R^{(s,r')}_{t_k^{(s)}}$; 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:

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

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

Proving (ii) **:** In this case, $\mathbb{E}\left[X^2\right]$ 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 $(\min_{r\neq s} R_t^{(s,r)})_{t\ge 0}$ 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.

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

Lemma 2 can then be applied to this process, with $(t_k^{(s)})_{k\ge 0}$ which is independent from it; this proves $\lim_{k\to\infty}\mathbb{E}\left[\min_{r\neq s} R^{(s,r)}_{t_k^{(s)}}\right] \le \frac{\mathbb{E}\left[X^2\right]+\mathbb{E}\left[X\right]}{2\mathbb{E}\left[X\right]}$.

If we consider the collection of random variables $((R_t^{(r,d)})_{t\ge 0})_{r\neq s}$, the condition (i) of Lemma 2 is met. As $(t_k^{(r)})_{k\ge 0}$ and $(r_k)_{k\ge 0}$ depend only on $(t_k^{(s)})_{k\ge 0}$ and contacts processes belonging to other pairs than $\{(r,d)\,|\,r\neq s\}$, they are independent from the collection above, and we have $\lim_{k\to\infty}\mathbb{E}\left[R^{(r_k,d)}_{t_k^{(r)}}\right] = \left(\frac{1+\mathbb{E}[X^2]}{2\mathbb{E}[X]}\right)$. Using (2), we have

$$R^{(r_k,d)}_{t_k^{(r)}} \le \quad D_k = \min_{r\neq s} R^{(s,r)}_{t_k^{(s)}} + R^{(r_k,d)}_{t_k^{(r)}},\ \text{ hence}$$

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

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

Proving (i)**, 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 $X_1 = T_1^{(r,d)} - T_0^{(r,d)}$, we have for any A, that may be chosen arbitrary 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 A arbitrary large, and $\alpha < 2$, we have: $\lim_{t\to\infty} \mathbb{E}\left[R_t^{(r,d)}\right] = +\infty$. The collection of processes $((R^{(r,d)})_{t\geq 0})_{r\neq s}$ verifies condition (b) of Lemma 2. As $(t_k^{(r)})_{k\geq 0}$ and $(r_k)_{k\geq 0}$ are independent of this collection, we can therefore deduce that

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

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

The short-contact case: The result (i) follows from the long-contact case, as the delay in the short contact case is always larger. The proof of (iii) 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 inter-contact 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 has therefore no impact on the delay experienced. In other words, the region $\alpha > 2$ may be thought as the *stability region* of the two-hop relaying algorithm.

C. 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 naive 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 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 remaining inter-contact 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.*

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

$$\textit{then } \mathbb{E}\left[R_t^{(d_1,d_1')}\right] = \ldots = \mathbb{E}\left[R_t^{(d_m,d_m')}\right] = +\infty$$
$$\textit{and } \mathbb{E}\left[\min(R_t^{(d_1,d_1')}, \ldots, R_t^{(d_m,d_m')})\right] < \infty.$$

Proof: As $\alpha > 1$, Lemma 3 (ii) 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,

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

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 a quite different nature, as even unlimited flooding does not achieve a bounded delay. We comment on this difference further in Section V.

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.*

(i) *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.* $\lim_{k\to\infty} \mathbb{E}\left[D_k\right] = \bar{D} < +\infty$.

(ii) *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:* $\mathbb{E}[D_k] = \bar{D} < +\infty$*.*

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

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

Consequently, $\mathbb{E}\left[\min_{r\neq s} R_t^{(s,r)}\right]$ 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 increases. No forwarding algorithm, no matter how redundant, can then transport a packet within a finite expected delay, using only opportunistic contact between devices. ■

Note: By comparison, the result (iii) 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.

D. Summary

At this stage, we have established the following results for the class of so-called naive forwarding algorithm defined in III-A, 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 by [10], is not stable in the sense that the delay incurred has an infinite expectation. It is however still possible to build a naive 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 inter-contact 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 inter-contact distribution in the [10 minutes; 1 day] range for all the scenarios we have previously studied empirically. This overall implies 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 that is a negative result, and we come back to interpret it and discuss its implications in Section V.

IV. 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 was considered to enable communication where no contemporaneous path may be found [7], to gather efficiently

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

information in a network of low power sensors [13], [14], [15], or 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 propose.

We did not find any previous work studying the characteristics of inter-contact 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 inter-contact distribution decays exponentially. In other words, for these models, the inter-contact 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]. It was shown recently that devices moving according to a Brownian motion in a bounded region, exhibit heavy tailed inter-contact 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 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.

V. 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 inter-contact time between two devices can be approximated by a power law on the [10 minutes; 1 day] range. We prove in a simple model the following major results: power law condition may be addressed in terms of delay by "naive forwarding algorithms" as long as the heavy tail index of the power law is greater than 1. When, by opposition, the heavy tail index is smaller than 1, the expected delay cannot be bounded for any forwarding algorithm of that type, even when one neglects 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 inter-contact distribution for any pair of nodes. That seems at odds with the empirical evidence of inter-contact distribution, for values up to 1 day, which is well approximated by a heavy tail distribution. Some of these 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 geographical locations. This approach has the advantage that it can be directly compared with a growing sets of real-life connectivity 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 deal with the feasibility of forwarding in opportunistic networks and their consequence requires more attention. At least three different directions may be followed.

- First, it might be that reasoning with expected value of delay is not suitable, since the possible occurrence of a long delay is unavoidable, whatever 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 contains significant correlation. It is still possible that a finite expected delay exist in a more complex model that reproduces accurately the statistical properties of our data sets. This direction is appealing but it requires to remove one of the modeling assumptions that we have made and that are instrumental for most of the results currently known in this area. It also necessitates to design a forwarding algorithm that differentiate 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. It looks promising but the impact of this partial infrastructure should be carefully studied.

We currently work 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 goal is to study the properties of the actual traffic created by users in an opportunistic data network.

VI. Acknowledgments

We would like to gratefully acknowledge Dave Kotz and Tristan Henderson at Dartmouth College, Geoff Voelker at University of California San Diego, Eyal de Lara and Jing Su at University of Toronto, and the Reality Mining project at MIT, for providing their data. We would like to thank also Vincent Hummel and Ralph Kling at Intel for their support, Mark Crovella for his insightful comments on our findings, Laurent Massoulié for his advice on some of the proofs, and the participants of the three experiments we conducted.

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," in *Proceedings of ACM SIGCOMM first workshop on delay tolerant networking and related topics*, 2005.

[2] M. McNett and G. M. Voelker, "Access and mobility of wireless PDA users," Computer Science and Engineering, UCSD, Tech. Rep., 2004.

[3] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," in *MobiCom '04: Proceedings of the 10th annual international conference on Mobile computing and networking*, 2004, pp. 187–201.

[4] N. Eagle and A. Pentland, "Reality mining: Sensing complex social systems," *Journal of Personal and Ubiquitous Computing*, 2005.

[5] P. Bremaud, *Markov Chains, Gibbs Field, Monte Carlo Simulation and Queues*, 1st ed. Springer-Verlag, 1999.

[6] F. Baccelli and P. Bremaud, *Elements of Queuing Theory*, 2nd ed. Springer-Verlag, 2003.

[7] A. Vahdat and D. Becker, "Epidemic routing for partially-connected ad hoc networks," UCSD, Tech. Rep., 2000.

[8] X. Chen and A. Murphy, "Enabling disconnected transitive communication in mobile ad hoc networks," in *Proceedings of the Workshop on Principles of Mobile Computing*, 2001, pp. 21–27.

[9] J. A. Davis, A. H. Fagg, and B. N. Levine, "Wearable computers as packet transport mechanisms in highly-partitioned ad-hoc networks," in *ISWC '01: Proceedings of the 5th IEEE International Symposium on Wearable Computers*, 2001, p. 141.

[10] M. Grossglauser and D. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Transactions on 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," 2005, preprint, School of ECE, Purdue University. [Online]. Available: http://min.ecn.purdue.edu/ linx/paper/it05.pdf

[12] J. Scott, P. Hui, J. Crowcroft, and C. Diot, "Haggle: a networking architecture designed around mobile users," in *Proceedings of The Third IFIP WONS Conference*, 2006.

[13] S. Jain, R. C. Shah, G. Borriello, W. Brunette, and S. Roy, "Exploiting mobility for energy efficient data collection in sensor networks," in *Proceedings of IEEE WiOpt*, 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," in *ASPLOS-X: Proceedings of the 10th international conference on Architectural support for programming languages and operating systems*, 2002, pp. 96–107.

[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)," in *MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, 2003, pp. 233–244.

[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. Netw.*, vol. 14, no. SI, 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," in *MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, 2003, pp. 257–266.

[18] A. Lindgren, A. Doria, and O. Schelén, "Probabilistic routing in intermittently connected networks," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 7, no. 3, pp. 19–20, 2003.

[19] J. Leguay, T. Friedman, and V. Conan, "Evaluating mobility pattern space routing for DTNs," in *Proceedings of IEEE Infocom*, 2006.

[20] S. Asmussen, *Applied Probability and Queues*, 2nd ed. Springer-Verlag, 2003.

APPENDIX

A. Preliminary Results

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, $\lim_{t\to\infty} \mathbb{E}\left[F_t^{(i)}\right] = l$, where $l \in \mathbb{R}\cup\{+\infty\}$, and*

$$(a)\ \forall i,t,\ \mathbb{E}\left[F_t^{(i)}\right] \in \mathbb{R},\quad \textit{and}\quad l\in\mathbb{R},$$

$$\textit{or } (b)\ \forall i,t,\ \mathbb{E}\left[F_t^{(i)}\right] \in \mathbb{R}\cup\{+\infty\}\quad \textit{and}\quad l=+\infty\ .$$

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. .

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

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 (a), we have $l<+\infty$ and

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

Let $M = \sup_{i\in I, t\leq T}\left|\mathbb{E}\left[F_t^{(i)}\right]-l\right|$, there exists K s.t. $k>K \implies \mathbb{P}\left[t_k\geq T\right] \geq 1-\dfrac{\varepsilon}{2\cdot M}$ and hence

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

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

Let us now suppose (b), we have $l=+\infty$ and

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

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

$$k>K \implies \mathbb{P}\left[t_k\geq T\right] \geq \max\left(1/2, 1-1/M'\right),\quad\text{and}$$

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

■

2) Remaining inter-contact: 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 an Homogeneous Markov Chain in $\mathbb{N}$ such that:

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

This Markov Chain is clearly irreducible and aperiodic as $\mathbb{P}\left[X = 1\right] > 0$, 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*

(i) *If $\alpha > 1$, $(R_t^{(d,d')})_{t\geq 0}$ is ergodic.*

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

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

such that, we have in steady state

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

(iii) *If $\alpha \leq 1$, $(R_t^{(d,d')})_{t\geq 0}$ 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 $\mathbb{E}_0[\mathtt{ret}_0] = \mathbb{E}\left[X\right]$. If $\alpha > 1$, we have $\mathbb{E}\left[X\right] < +\infty$, proving (i), and if $\alpha \leq 1$, we have $\mathbb{E}\left[X\right] = +\infty$, proving (iii).

By (i), 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 $c_1 = 1/\sum_{i\geq 0}(i+1)^{-\alpha}$.

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 non-increasing, we have:

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

$$\text{thus} \quad \frac{(i+2)^{-(\alpha-1)}}{\alpha-1} \leq \sum_{j>i}(j+1)^{-\alpha} \leq \frac{(i+2)^{-(\alpha-1)}}{\alpha-1}$$

which completes the proof for (ii). ■

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

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

B. Queuing with a Process of Service Instant

By opposition to classical queueing 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 $\mathbf{s} = \{\ s_m \ \mid m \in \mathbb{Z}\ \}$. We make the following assumptions:

- A customers 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 queues 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 the time-slot t is completed. The process Q follows the recursion:

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

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

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 $\mathbb{E}\left[N(0)\right]$.

The next result follows closely the proof of the stability regime for a single server queue (see [6] p.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* $\lambda < \mu$*,*

(i) *There exists an initial condition,* $\tilde{Q} < \infty$ *a.s. , such that the queue process verifies:* $\tilde{Q}(t) = \tilde{Q} \circ \theta^t$*.*

(ii) *In this stationary regime:* $\mathbb{E}^0_{\mathbf{a}}\left[\tilde{V}_k\right] = 1 + \frac{1}{\lambda}\mathbb{E}\left[\tilde{Q}\right]$

(iii) *If the queue starts empty,* $\limsup_{k\to\infty} \mathbb{E}\left[V_k\right] \leq 1 + \dfrac{1}{\lambda}\mathbb{E}\left[\tilde{Q}\right]$

Proof: We define the sequence of variables indexed by T

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

Clearly this sequence is positive, non-decreasing, and verifies

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

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

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

This limit may take infinite values. Note that $\{\tilde{Q} = \infty\}$ 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:

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

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

By monotone convergence, we deduce

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

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

$$\mu > \lambda \implies \tilde{Q} < \infty \ \text{ a.s. } ; \ \text{ which proves } (i).$$

(ii) is an application of the Campbell-Mecke egality:

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

We have a.s. $V_k \leq \tilde{V}_k$, which proves (iii). ■

2) Expected queue length:

Lemma 5 *Assume* $\mathbf{a}$ *and* $\mathbf{s}$ *are two renewal point processes,*
- *with intensities* $\lambda < \mu$*,*
- *such that inter-event distribution* $F_{\mathbf{a}}$ *has a finite mean,*
- *and the inter-event distribution* $F_{\mathbf{s}}$ *has a finite variance,*

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

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

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

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

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

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 inter-event distribution. The above expectation may be rewritten

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

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

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

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

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

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

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

■

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

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

Proof: According to the proof of Lemma 4,

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

The result is then following the above lemma. ■

3) Proof of Theorem 1 in the short-contact case: Let us summarize results from the above subsections: 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) inter-service-instant distribution has a finite variance. Note that conditions (2) and (3) are necessary. In the following, we present a scheme insuring that all queues implemented in the mobile nodes verify conditions (1),(2) and (3). It may be improved at the cost of an additionnal 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 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 reason we also assume that with a small probability p, taken independently, an *independent blocking* occur, 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 (1), (2) and (3).
- A *relay queue* satisfy (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 *quasi-saturated* renewal process (that include 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 always not larger than the expected delay in the same queue with a *quasi-saturated* arrival process. One check easily that this last case verify all 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. Proof of Corollary of Lemma 1

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

Corollary 2 *Let* $(R_t^{(d_1,d'_1)})_{t\ge 0},\ldots,(R_t^{(d_m,d'_m))})_{t\ge 0}$ *be the remaining inter-contact times for* m *different pairs of devices* $(d_i,d'_i)_{1\le i\le m}$. *We suppose that* $\alpha > 1+\frac{1}{m-j+1}$,

$$\text{then } \mathbb{E}\left[\texttt{ord}\left(j,(R_t^{(d_1,d'_1)},\ldots,R_t^{(d_m,d'_m)})\right)\right] < \infty.$$

Proof: Let $M_j = \texttt{ord}\left(j, R_t^{(d_1,d'_1)},\ldots,R_t^{(d_m,d'_m)}\right)$

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

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

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

which implies that $\mathbb{E}\left[M_j\right] < \infty$ as soon as $\alpha > 1+\frac{1}{m-j+1}$.