

The Diameter of Opportunistic Mobile Networks

Augustin Chaintreau Abderrahmen Mtibaa Laurent Massoulié Christophe Diot

Thomson, Paris, France. email: firstname.lastname@thomson.net

ABSTRACT

Portable devices have more data storage and increasing communication capabilities everyday. In addition to classic infrastructure based communication, these devices can exploit human mobility and opportunistic contacts to communicate. We analyze the characteristics of such opportunistic forwarding paths. We establish that opportunistic mobile networks in general are characterized by a small diameter, a destination device is reachable using only a small number of relays under tight delay constraint. This property is first demonstrated analytically on a family of mobile networks which follow a random graph process. We then establish a similar result empirically with four data sets capturing human mobility, using a new methodology to efficiently compute all the paths that impact the diameter of an opportunistic mobile networks. We complete our analysis of network diameter by studying the impact of intensity of contact rate and contact duration. This work is, to our knowledge, the first validation that the so called “small world” phenomenon applies very generally to opportunistic networking between mobile nodes.

1. INTRODUCTION

The proliferation of powerful portable devices has created a new environment for networking. As opposed to conventional communication that relies on infrastructure, these devices can use hop-by-hop opportunistic data forwarding between each other. In this environment, a device should decide whether or not to transfer a message at the time it meets another one. How to select the next hop towards the destination in a way to minimize delay and maximize success rate is so far unknown. It is in general difficult to design opportunistic

forwarding algorithms, as their performance depends extensively on the characteristic of the mobility present in the network [2]. We claim that it is critical to study first the properties of the paths made available between nodes by opportunistic contacts and mobility.

Among topological properties that directly impact forwarding, it is essential to characterize the diameter of an opportunistic network. The diameter bounds the number of hops needed to construct a path over time between two nodes. Opportunistic forwarding seems in general significantly harder for networks with large diameters. Note that as opposed to paths in a static graph, the paths in opportunistic mobile networks have both time-position and hop-number; we wish to characterize both of them since each has a great impact on the feasibility of forwarding. We define formally the diameter for any opportunistic mobile network as follows: the number of hops needed to achieve a high proportion (e.g. 99%) of the success rate of flooding, under any time constraint. This work does not aim at characterizing one forwarding algorithm, but instead it contributes to a better understanding of the performance of all algorithms with regard to hops and delays.

This paper makes the following contributions.

- We prove that the diameter for a family of opportunistic networks, described by a process of random graphs, increases slowly with the network size. A phase transition characterizes when paths that are short both in terms of delay and hop-number may be found. This result is analogous to the short diameter observed among vertexes of a random graph. We also prove that the hop-number of the delay-optimal path varies little with the contact rate, especially when the network is sparse (i.e. when the contact rate is low) (§3).
- Since the definition of the diameter requires to know at all time the delay-optimal path between two nodes, we propose an efficient algorithm that computes these paths exhaustively (§4).
- We apply our technique to four mobility traces. The results validate the implication of our analy-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CoNEXT'07, December 10-13, 2007, New York, NY, U.S.A.
Copyright 2007 ACM 978-1-59593-770-4/ 07/ 0012 ...\$5.00.

sis for measurement of human mobility in diverse environments; we generally observe diameter between 4 and 6 hops (§5). Lastly, we investigate empirically the impact of different characteristics of opportunistic contacts (duration, intensity of contacts) on the network diameter (§6).

To the best of our knowledge, this work presents for the first time both simple analytical results and empirical evidence validating that the so called “small world” phenomenon is relevant for delay-efficient opportunistic networking.

2. RELATED WORK

Opportunistic mobile networks can be seen as a class of delay-tolerant networks. As opposed to other works in this area, opportunistic forwarding do not make the assumption that details or summary about future device mobility are known (e.g. [7]) nor that mobility may be partly controlled to serve the network’s need (e.g. [19]). Grossglauser and Tse pioneered theoretical analysis of opportunistic forwarding; they established in [5] that mobility increases the capacity of a network, when devices are densely deployed and follow a regular mobility process. Most of the forwarding algorithms proposed since that time (see e.g. [2] and references therein) includes for each packet a time-out and a maximum number of hops, to avoid consuming too much resource. These parameters should depend on the properties of paths available thanks to the mobility, although this aspect remains little known today. In [16], the authors advocate to use opportunistic connections and social network properties to improve data dissemination in mobile networks. Our paper justifies that the small world phenomenon can be beneficial in such context.

Previous works that characterize the impact of mobility on opportunistic forwarding have focused on the distribution of the time between two successive contacts for the same pair, also called the inter-contact time [2],[9, 18, 17]. [18] identifies modes in the contact process among buses, created by periodic schedule. In [17], using a data set extracted from students lectures schedule, the authors studies minimal delay and a “hop distance” separately; the latter is computed using a static graph extracted from the mobility. None of these previous works has studied the delay and hop-number properties of paths available over time in a general context. Pioneer experimental works collecting mobility traces did not consider multi-hop paths properties [6, 13, 4].

The characterization of paths length in a random graph is not a new research topic [8], but to the best of our knowledge, none of these works considered a graph that evolves with time. Up to now, the characterization of dynamic networks has considered contemporaneous paths in a graph with an increasing set of vertexes [12]. One essential new feature in our model is that the path

itself is not drawn at a given point in time, but should follow a sequence of steps in a chronological way. These structures, known as temporal networks, or sometimes called “evolving graphs”, have been studied from an algorithmic standpoint [15, 10, 1]. None of these works analyzed large random structures with such property.

3. RANDOM TEMPORAL NETWORKS

We model opportunistic mobile networks as a temporal network, i.e. a graph with a static set of nodes, and a set of edges that may change with time. In this paper we discuss the properties of sequences of edges that verify a chronological property. These sequence are generally called paths. As a consequence we need to study jointly how paths behave with respect to time and with respect to the sequence of nodes they follow.

We analyze in this section a family of simple temporal networks, *random temporal networks*, which capture dynamic edges using a sequence of uniform random graphs. In other words, we assume that during each time slot, a contact between two nodes may occur with a fixed probability, independently of other nodes and other time slots. A couple of variants of this model are presented to include the impact of continuous time, and different assumptions on the latency/bandwidth found locally. One of our main findings is a condition for the existence of paths with constraints on both delay and hop-number.

Before starting this study, note that random temporal networks follow simplifying assumptions about mobility that are not usually met in practice. This point is discussed further in §3.4. We present a more general model of temporal network in §4. It is used to study properties of paths found in empirical traces, where these assumptions need not hold.

3.1 Definitions & Assumptions

Consider a network made of N nodes. We are studying the properties of large graphs, assuming that the contact rate (i.e. the average number of contacts made by a node in a unit of time) remains a constant λ .

For any functions f, g of N , we write $f(N) = \tilde{\Theta}(g(N))$ if there exist four positive constants c, C and α, β such that, for N sufficiently large, we have

$$c(\ln N)^{-\alpha} \cdot g(N) \leq f(N) \leq C(\ln N)^{\beta} \cdot g(N).$$

We use the following elementary consequence:

PROPOSITION 1. *If $f(N) = \tilde{\Theta}(N^a)$ with $a \in \mathbb{R}$, then $a < 0 \implies \lim_{\infty} f = 0$, $a > 0 \implies \lim_{\infty} f = \infty$,*

3.1.1 Discrete-time model

This may be seen as a generalization of the uniform random graph introduced by Erdős and Rényi. A classical uniform random graph is a graph $G = (V, E)$ such that the set of edges E is a random variable verifying:

- $\mathbb{P}[(u, v) \in E] = p$ for any pair of nodes (u, v) .
- The events $\{(u, v) \in E\}$, for all pairs (u, v) , are mutually independent.

Note that since we assume an average contact rate λ , we have in this case $p = \frac{\lambda}{N-1}$.

We define a random temporal network as a collection of graphs $\{G_t = (V, E_t) \mid t \in \mathbb{N}\}$ such that

- $\mathbb{P}[(u, v) \in E_t] = p = \frac{\lambda}{N-1}$ for any pair (u, v) .
- The events $\{(u, v) \in E_t\}$, for all pairs (u, v) and all time t , are mutually independent.

The pairs (u, v) may be directed or undirected without changing any of the definitions above.

3.1.2 Continuous-time model

In the discrete-time model above, for any pair of nodes (u, v) the set of indices t such that edge (u, v) is in E_t is a sequence of integers in \mathbb{N} separated by geometric random variables. A natural generalization of the model to a continuous time setting is then to assume that, for any pairs of nodes (u, v) , the times of contact are separated by exponential random variables. In other words, this process of time instants constitutes a Poisson process. All the results of this paper have been also demonstrated in continuous time (see [3]).

3.1.3 Paths in long contact case/short contact case

A path from u to v in a temporal network is a sequence $u = u_0 \rightsquigarrow^{t_1} u_1 \rightsquigarrow \dots \rightsquigarrow^{t_k} u_k = v$ such that:

- (i) $(u_{i-1}, u_i) \in E_{t_i}$ for all $i = 1, \dots, k$,
- (ii) $t_{i+1} \geq t_i$ for all $i = 1, \dots, k$.

We assume that all contacts have a fixed duration that is either one time slot (in the discrete time model) or negligible (in the continuous time model). To model the impact of a limited local bandwidth, or local latency, on the properties of paths, we introduce two cases.

We define the *long contact case* as the model where any number of edges may be used in a single time slot, as allowed by the definition of a path found above. In other words we assume that a single time slot is sufficiently long to exchange across several contacts.

On the contrary, we define the *short contact case* as the model where only one contact may be used in a single time slot. In other words, we require that all paths verify, in addition to the conditions above:

- (ii') $t_{i+1} \geq t_i + 1$ for all $i = 1, \dots, k$.

3.2 Phase transition

In this section, we first prove a result on the expected number of paths between two nodes, when delay and hops are constrained. This result is then used to describe a phase transition for the appearance of paths in a random temporal network.

3.2.1 Expected number of paths with constraints

We describe now our main analytical results, which characterize the expected number of paths with constraints on both delay and hop-number. A source u and a destination v are fixed in advance and without loss of generality we assume that the packet is ready to be sent at the source at time $t = 0$.

Since we are interested in large networks, with a constant contact rate per node, we let N go to infinity and limit the time and the number of hops allowed in the paths by a slowly increasing function of N .

LEMMA 1. *Let us denote the maximum time t_N and number of hops k_N of a path allowed in the network. We assume that they are given as a function of N by:*

$$\begin{cases} t_N = \lfloor \tau \cdot \ln(N) \rfloor, \\ k_N = \lfloor \gamma \cdot t_N \rfloor = \lfloor \gamma \cdot \tau \cdot \ln(N) \rfloor, \end{cases} \quad (1)$$

where τ and γ are two positive constants.

Let us denote by Π_N the number of paths from u to v under the above constraints. Then, as N grows large

- for short contacts, $\mathbb{E}[\Pi_N] = \tilde{\Theta}\left(N^{-1+\tau(\gamma \ln(\lambda)+h(\gamma))}\right)$
where $h : x \in [0; 1] \mapsto -x \ln(x) - (1-x) \ln(1-x)$.
- for long contacts, $\mathbb{E}[\Pi_N] = \tilde{\Theta}\left(N^{-1+\tau(\gamma \ln(\lambda)+g(\gamma))}\right)$
where $g : x \in [0; 1] \mapsto (1+x) \ln(1+x) - x \ln(x)$.

PROOF. This proof follows from the memoryless property of geometric and exponential distributions. We first estimate the probability of success of a path whose nodes are fixed in advance. The expectation is then this probability multiplied by the number of possible combinations. Stirling formula is used to complete the argument. The complete proof can be found in [3]. \square

COROLLARY 1. *Under the assumption above we have, in the short contact case*

$$\begin{cases} \mathbb{E}[\Pi_N] \rightarrow 0 & \text{if } 1/\tau > \gamma \ln(\lambda) + h(\gamma) \\ \mathbb{E}[\Pi_N] \rightarrow \infty & \text{if } 1/\tau < \gamma \ln(\lambda) + h(\gamma). \end{cases}$$

The first result implies that, when $1/\tau > \gamma \ln(\lambda) + h(\gamma)$:

$$\mathbb{P}[\text{There exists a path with constraints (1)}] \rightarrow 0$$

All these results holds for the long contact case when replacing the function h by g .

PROOF. The two first results follows directly from Lemma 1 and Proposition 1. The last result is a consequence from the Markov inequality, which may be written here $\mathbb{P}[\Pi_N \geq 1] \leq \mathbb{E}[\Pi_N]$. \square

3.2.2 Phase transition in the short contact case

The results from the previous section prove that, depending on the values of the two constant numbers τ and γ , as well as the contact rate λ , one of the two following statements holds: either there almost surely does

not exist a path satisfying the logarithmic bound (1); or the number of paths that satisfy these conditions grows on average to infinity with N . Moreover, one can tell which statement holds simply by comparing the values of $1/\tau$ and $\gamma \ln(\lambda) + h(\gamma)$.




Figure 1: Phase transition (short contact case)

This result might be interpreted as follows. Note that the short contact case implies $\gamma \leq 1$. The function $\gamma \in [0; 1] \mapsto \gamma \ln(\lambda) + h(\gamma)$ admits a maximum, that is given by $M = \ln(\lambda + 1)$ and this maximum is attained when $\gamma = \frac{\lambda}{1+\lambda}$. This is illustrated on Figure 1 where the value of this function of γ was plotted for three different values of λ . We can deduce the following dichotomy:

- If $\tau < 1/M = 1/\ln(1 + \lambda)$, then $1/\tau$ is always larger than $\gamma \ln(\lambda) + h(\gamma)$ for any value of γ . As a consequence, almost surely there does not exist a path with delay less than $\tau \ln(N)$ when N is large.
- If $\tau > 1/\ln(1+\lambda)$, then the super-critical condition from Corollary 1 is verified for $\gamma \in [\gamma_1; \gamma_2]$, an interval which contains $\frac{\lambda}{1+\lambda}$. As a consequence, the average number of paths with delay $\tau \ln(N)$ and $\gamma \tau \ln(N)$ hops is unbounded for large N ,

The delay-optimal path corresponds to the critical value of τ for which a path is likely to be found. From a heuristic standpoint, we then expect the delay-optimal path to have delay $t \approx \frac{\ln(N)}{\ln(1+\lambda)}$. When τ approaches this critical value, the interval of possible values for γ becomes small, and centered around $\frac{\lambda}{1+\lambda}$. Hence, we expect that the delay optimal path has hop-number given by $k \approx \frac{\lambda \ln(N)}{(1+\lambda) \cdot \ln(1+\lambda)}$. As an example, when $\lambda = 0.5$, we expect a delay growing with N as $t \approx 2.47 \cdot \ln(N)$ and we expect a number of hops $k \approx 1.64 \cdot \ln(N)$.

3.2.3 Phase transition in the long contact case

In the long contact case, the expected number of paths with delay and hop constraints, for large N , becomes large if $1/\tau < \gamma \ln(\lambda) + g(\gamma)$. In contrast with the short contact case, the properties of this function of γ change with the value of λ (see Figure 2). As a result, we look at the cases $\lambda < 1$ and $\lambda > 1$ separately.




Figure 2: Phase transition (long contact case)

- When $\lambda < 1$, the result is very similar to the short contact case. The function of γ admits a maximum $M = -\ln(1 - \lambda)$, attained when $\gamma = \frac{\lambda}{1-\lambda}$.

Following the same heuristic as the short contact case, we expect the delay-optimal path to have delay $t \approx \frac{\ln(N)}{-\ln(1-\lambda)}$ and $k \approx \frac{\lambda \ln(N)}{-(1-\lambda) \cdot \ln(1-\lambda)}$ hops. As an example, when $\lambda = 0.5$, we expect a delay $t \approx 1.69 \cdot \ln(N)$ and the same number of hops.

- When $\lambda > 1$, we notice a difference with the short contact case, since the function of γ is increasing and unbounded. In that case, for any τ , even an arbitrary small constant, some paths exist with a delay less than $\tau \ln(N)$.

To compute the hop-numbers of these paths, we proceed as follows. Since $1/\tau$ is large, γ should then be sufficiently large to satisfy the condition of the corollary. The function of γ is then close to its asymptote $\gamma \mapsto 1 + \gamma \ln(\lambda)$. We deduce that the smallest value of γ verifying the condition is $\frac{1-\tau}{\tau \ln(\lambda)}$, hence $k = \gamma \tau \ln(N)$ given by $k \approx \ln(N)/\ln(\lambda)$.

This last regime should not come as a surprise: in a static random graph, there exists a phase transition when λ approaches 1. In particular, when λ is greater than 1, there almost surely exists a unique connected component with a large size (see Theorem 5.4 p.109 in [8]). In our model, since the long contact case allows

one to use any number of hops during the same time slot, this property implies that there does exist paths for any arbitrarily small τ . In other words, although the network is still changing with time, it is essentially “almost-simultaneously connected”.

3.3 Impact of the contact rate

We have seen above that the value of the contact rate λ impacts qualitatively the phase transition in the long contact case, as λ becomes larger than 1. More generally, when λ is less than 1, it has a direct impact on the delay of the paths present in a random temporal network. We note in particular that as λ gets small, the minimum value of τ for which a path may exist (respectively $1/\ln(1+\lambda)$ and $-1/\ln(1-\lambda)$ for the short and the long contact case) becomes large.

In contrast with delay, the hop-number of a path varies little as λ changes. In particular, when λ becomes small, the hop-number for the delay-optimal path in both short and long contact cases no longer depends on λ and it converges to $\ln(N)$.




Figure 3: Hop-number of the delay-optimal path seen as a function of the contact rate

One may think of the following explanation for this result: As λ decreases, the network is essentially completely disconnected except occasionally for one or several disjoint pairs of nodes. Therefore, if we rescale time by merging periodically a fixed number of time slots together, the network appears to be very similar (a few disjoint pairs of nodes) except for an appropriate scaling of the contact rate. As a consequence, decreasing λ linearly increases the time needed to create a path, but is unlikely to impact the actual hop-number of this path. The analysis above makes this intuition rigorous.

We summarize our results in Figure 3. The y -axis represents the estimate of the hop-number for the delay-

optimal path, normalized by $\ln(N)$. We see that in both dense and sparse regimes, short and long contacts are in good agreement. They differ only near the critical value $\lambda = 1$, where the long contact case has a singularity. This singularity is not likely to occur in practice as bandwidth limits the hops in a single time slot.

3.4 Discussions

We have studied the properties of random temporal networks, as they become large. We have shown that a source destination pair admits on average many forwarding paths using a small number of hops and over a short number of time-slots (i.e. both grow as the logarithm of the network size). Note that the delay of the optimal paths depends heavily on the contact rate between the nodes. On the other hand, the hop-number of these paths seems rather insensitive to changes in the contact rate λ , with the only exception of the long-contact case around the value $\lambda = 1$.

Inspired by these results, one may compute the diameter of a network using the paths with an optimal delay, and deduce that temporal networks generally have a small diameter, for almost any contact rate. However, our model is based on several important simplifications that may impact in general the delay and hop-number of paths in opportunistic networks:

- **Size:** The results above assume that the network is large enough to apply equivalence class results.
- **Homogeneity:** We have assumed that nodes contact others uniformly at random. In practice this is not true as people tend to come close to each other according to their habits and the communities of interest that they share.
- **Inter-contact time statistics:** Since we assume that contacts between nodes follow Bernoulli or Poisson Processes, the distribution of time between two contacts of a pair is light tailed. Previous experiments have shown that this assumption holds only at the timescale of days and weeks [2, 9].
It is nevertheless possible to extend all of the results we have obtained so far to contacts described by a renewal process with general inter-contact time distribution with finite variance. We expect this to have a major impact on the delay of a path, but a relatively small impact on hop-number.
- **Stationarity:** This model does not include periodic diurnal cycles in the variation of the contact rate that are typically found in human mobility. The network may change from a highly mobile and dense subset of contacts into a sparse and slowly varying subset of contacts. Again, we expect this effect to impact the delay of paths in temporal network, but not much their hop-number.

In the rest of this article, after having formally defined the diameter of any temporal network, we estimate its value for mobility traces where in general none of these assumptions holds. Our goal is to demonstrate that the fundamental insight brought up by this simple model translate in practice into similar qualitative trends.

4. NETWORK DIAMETER DEFINITION AND MEASUREMENT METHOD

In this section, we first formalize our definition of network diameter. We then describe an algorithm to study efficiently and exhaustively the properties of delay optimal paths in opportunistic networks. This algorithm is instrumental to study the diameter found in large traces, as done in the next section. It is one of the contributions of this work.

4.1 Definition of diameter

The diameter of a network is an upper bound on the number of intermediate hops needed to find at least one path between any two nodes. What is essentially new in a temporal network is that one has to specify whether this path should also satisfy a condition on its time characteristics. Inspired by the results from the previous section, we define the diameter such that we require this path to be almost always optimal.

For $k \in \mathbb{N} \cup \{\infty\}$, and $t \geq 0$ let $\Pi(t, k)$ be 1 if there exists a path that uses at most k hops and succeeds to deliver a packet within t seconds, let $\Pi(t, k)$ be 0 otherwise. This variable depends on the source, destination, and starting time of the packet. For $\varepsilon > 0$, we define the $(1-\varepsilon)$ -diameter as the minimum integer k such that

$$\forall t \geq 0, \mathbb{P} [\Pi(t, k) = 1] \geq (1 - \varepsilon) \cdot \mathbb{P} [\Pi(t, \infty) = 1] .$$

This expression can be interpreted in two ways: in a stochastic model of a stationary homogeneous network, like the one we studied in §3, the probability is chosen according to the distribution of the random variables $\Pi(t, k)$ seen at any time between any two nodes. In a mobility data set, we consider the empirical probability combining uniformly observations for all sources, destinations and starting times. In both case, this expression states that, for any delay-constraint, it is almost as likely to find a successful path within k hops than it is with any more hops.

One may think of this definition as an example of “competitive analysis” since the success ratio (i.e., the probability to find a path within t seconds and at most k hops) is not described in absolute term, but instead is compared with an optimal strategy. It helps the definition to adapt to variable environments. However, it requires one to know beforehand the performance of the delay-optimal paths at all time, which is why we formally describe in the rest of this section how they can be extracted from traces.

4.2 Paths in temporal networks

Each data set may be seen as a temporal network. More precisely we represent it as a graph where edges are all labeled with a time interval, and there may be multiple edges between two nodes. A vertex represents a device. An edge from device u to device v , with label $[t^{\text{beg}}; t^{\text{end}}]$, represents a contact, where u sees v during this time interval. The set of edges of this graph therefore includes all the contacts recorded by each device.

Paths associated with a sequence of contacts: We intend to characterize and compute in an efficient way *all* the sequences of contacts that are available to transport a message in the network. Note that we allow simultaneous contacts to be used, as in the long contact case defined in §3.1. It is possible to include a positive transmission delay in all these definitions, we expect that the diameter will be smaller in that case.

A sequence $(e_i = (u_{i-1}, u_i, [t_i^{\text{beg}}; t_i^{\text{end}}]))_{i=1, \dots, n}$ of contacts is *valid* if it can be associated with a time respecting path from u_0 to u_n . In other words, it is valid if there exists a non-decreasing sequence of times $t_1 \leq t_2 \leq \dots \leq t_n$ such that $t_i^{\text{beg}} \leq t_i \leq t_i^{\text{end}}$ for all i . An equivalent condition is given by:

$$\forall i = 1, \dots, n, t_i^{\text{end}} \geq \max_{j < i} \left\{ t_j^{\text{beg}} \right\} . \quad (2)$$

The time-respecting path associated with a sequence of contacts (e_1, \dots, e_n) is not unique, but we can characterize all of them as follows. Let us formally define the *last departure* of this sequence as $\text{LD}(\mathbf{e}) = \min_i \{ t_i^{\text{end}} \}$; and the *earliest arrival* as $\text{EA}(\mathbf{e}) = \max_i \{ t_i^{\text{beg}} \}$.

From the definition of a time respecting path, we have

- (i) All paths associated with this sequence of contacts verify $t_1 \leq \text{LD}$ and $t_n \geq \text{EA}$.

This property shows that the last departure is in fact the maximum possible starting time of a path using this sequence of contacts. Similarly the earliest arrival denotes the minimum possible ending time for a path using this sequence. These two optimums are attained, as one can immediately check the following.

- (ii) If $\text{LD} \leq \text{EA}$, there is a path with $t_1 = \text{LD}$, $t_n = \text{EA}$.
- (iii) If $\text{EA} \leq \text{LD}$, there is a path with $t_1 = t_2 = \dots = t_n = t$ for all $t \in [\text{EA}; \text{LD}]$.

Concatenation: Concatenating two sequences of contacts that both verify Eq. (2) does not necessarily create a compound sequence that verifies Eq. (2) (see [3] for a counterexample). We can characterize exactly when concatenation between two sequences is possible:

- (iv) Two sequences $(\mathbf{e}), (\mathbf{e}')$ of contacts such that $u_n = u'_0$ and that both verify Eq. (2) can be concatenated into a sequence of contacts $\mathbf{e}' \circ \mathbf{e}$ satisfying Eq. (2) if and only if $\text{EA}(\mathbf{e}') \leq \text{LD}(\mathbf{e})$.

When the condition above is verified, we can deduce the values LD, EA associated with the concatenated sequence as follows: $EA(e' \circ e) = \max(EA(e), EA(e'))$, and $LD(e' \circ e) = \min(LD(e), LD(e'))$ (see examples in Figure 4). Note that $EA = t^{\text{beg}} \leq t^{\text{end}} = LD$ for a sequence made with a single contact, but sequences with multiple contacts, like Figure 4 (a), might not verify $EA \leq LD$.




Figure 4: Two examples of concatenation

4.3 Delay-optimal paths

So far we have been describing a method to characterize when a sequence of contacts supports a time-respecting path, and when we can concatenate them. However, computing all of them in general is very costly. In this section, we formally define paths that are optimal in terms of delay, and therefore are susceptible to impact the diameter of the network.

Delivery function: As a consequence of (ii) and (iii), for a message created at u_0 at time t , if $t \leq LD$ then there exists a path associated with the sequence of contacts e , that transports this message and delivers it to u_n at time $\max(t, EA)$. Otherwise, when $t > LD$, no path based on these contacts exists to transport the message. The optimal delivery time of a message created at time t , on a path using this sequence of contacts,

$$\text{del}(t) = \begin{cases} \max(t, EA) & \text{if } t \leq LD, \\ \infty & \text{else.} \end{cases}$$

Similarly the optimal delivery time for any paths that use one of the sequences of contacts e_1, \dots, e_n is

$$\text{del}(t) = \min\{\max(t, EA_k), 1 \leq k \leq n \text{ s.t. } t \leq LD_k\} \quad (3)$$

where the minimum of an empty set is equal to ∞ .

Optimal paths: We say that a time respecting path, leaving device u_0 at time t_{dep} , arriving in device u_n at time t_{arr} , is strictly dominated in case there exists another path from u_0 to u_n with starting and ending times $t'_{\text{dep}}, t'_{\text{arr}}$ such that ($t'_{\text{dep}} \geq t_{\text{dep}}$ and $t'_{\text{arr}} \leq t_{\text{arr}}$), and if at least one of these inequalities is strict. A path is said *optimal* if no other path strictly dominates it.

According to (ii) and (iii) above, among the paths associated with a sequence of contacts with values (LD, EA) the optimal ones are the following: if $LD \leq EA$, this is the path starting at time LD and arriving at time EA . Otherwise, when $LD > EA$, all paths that start and arrive at the message generation time $t \in [EA; LD]$ are optimal.




Figure 5: Example of a delivery function, and the corresponding pairs of values $(LD_i, EA_i)_{i=1,2,3,4}$.

An example of delivery function is shown in Figure 5. Note that the value of the delivery function (y-axis) may be infinite. Pairs (LD_1, EA_1) to (LD_3, EA_3) satisfy $EA \leq LD$, they may correspond to direct source-destination contacts, or sequence of contacts that all intersect at some time; the fourth pair verifies $LD_4 < EA_4$, hence it does not correspond to a contemporaneous connectivity. The message needs to leave the source before LD_4 , and remains for sometime in an intermediate device before being delivered later at time EA_4 .

4.4 Efficient computation of optimal paths

We construct the set of optimal paths, and delivery function for all source-destination pairs, as an induction on the set of contacts in the traces. We represent the delivery function for a given source-destination pair by a list of pairs of values (LD, EA) . The key element in the computation is that only a subset of these pairs is needed to characterize the function del . This subset corresponds to the number of discontinuities of the delivery function, and the number of optimal paths that can be constructed with different contact sequences.

We use the following observation: We assume that the values $(LD_k, EA_k)_{k=1, \dots, n}$, used to compute the delivery function as in (3), are increasing in their first coordinate. Then, as $k = n, n-1, \dots$, we note that the k -th pair can always be removed, leaving the function del unchanged, unless this pair verifies:

$$EA_k = \min\{EA_l \mid l \geq k\}. \quad (4)$$

In other words, a list such that all pairs verify this condition describes all optimal paths, and the function del , using a minimum amount of information.

As a new contact is added to the graph, new sequences of contacts can be constructed thanks to the concatenation rule (fact (iv) shown above). This creates a new set of values (LD, EA) to include in the list

of different source-destination pairs. This inclusion can be done so that only the values corresponding to an optimal path are kept, following condition (4).

We show that our method can also be used to identify all paths that are optimal inside certain classes, for instance the class of paths with at most k hops. This can be done by computing all the optimal paths associated with sequences of at most k contacts, starting with $k = 1$, and using concatenation with edges on the right to deduce the next step.

Compared with previous generalized Dijkstra’s algorithm [1, 7], this algorithm computes directly representation of paths for all starting times. That is essential to have an exhaustive search for paths, and estimate the diameter of a network at any time-scale. We found algorithm UW2 in [15] to be the closest to ours. We have introduced here an original specification through a concise representation of optimal paths which makes it feasible to analyze long traces with hundred thousands of contacts. Recently we have found that another algorithm has been developed independently to study minimum delay in DTN [18]. It works as follows: a packet is created for any beginning and end of contacts; a discrete event simulator is used to simulate flooding; the results are then merged using linear extrapolation.

5. EMPIRICAL RESULTS

In this section we present first our data sets. Then we analyze the characteristics of optimal paths in opportunistic networks using the methodology described in the previous section.

5.1 Mobility data sets

We use four experimental data sets. Three were collected by the Haggle Project [2]. They include two experiments conducted during conferences, *Infocom05*, *Infocom06*, and one experiment conducted in *Hong-Kong*. The fourth data set was collected by the MIT Reality Mining Project [4]. Table 1 summarizes important characteristics of these data sets.

In the Haggle experiments, people were asked to carry an experimental device (i.e., an iMote [2]) with them at all times. These devices log all contacts between experimental devices (i.e., called here internal contacts) using a periodic scanning every t seconds, where t is called granularity. In addition, they log contacts with other external Bluetooth devices (i.e., external contacts) that they meet opportunistically (e.g., cell phones, PDAs).

In *Infocom05*, and *Infocom06* data sets, the experimental devices were distributed to students attending the conference. The largest experiment is *Infocom06* with 78 participants. By default we are presenting here results for internal contacts only; results with internal and external contacts are very similar.

In *Hong-Kong*, people carrying the experimental de-

vices were chosen carefully in a Hong Kong bar to avoid social relationships between them. These people returned the iMote at the same bar a week later. As a consequence, there are only few internal contacts, which is why we are presenting here results with both internal and external contacts.

Note that our data sets may not include all opportunistic encounters between participants of the experiment, because of the time between two scans, hardware limitations, software parameters, and interference [14]. For the same reasons, it is also possible that some contacts appear shorter than they are. In addition, we do not have access to the contact history of the external devices and, as consequence, we miss some of the direct contacts between them as well. We examine the impact of this “sampling” effect in §6.

The *Reality Mining* data set includes records from Bluetooth contacts and GSM base stations for a group of cell-phones distributed to 100 MIT students during 9 months. We only show results from the Bluetooth data set. We also made the same observations on the GSM data set, as well as other publicly available data sets, including traces from campus WLAN in Dartmouth [6] and UCSD [13]. These results can be found in [3].

5.2 Preliminary observations

Figure 6 shows the next time the device is in range of another device (z-axis) as a function of time (x-axis), for six representative participants (y-axis) chosen in three data sets. Hence, the diagonal on this plot represents a period of uninterrupted contact, while the steps correspond to period with no contact at all.




Figure 6: Time of the next contact with any other device, as seen by six participants in Hong Kong, Reality Mining, and Infocom05.

These results confirm that the data sets we have used exhibit variable contact characteristics over time. The nodes in *Hong Kong* and *Reality Mining* exhibit low

Experimental data set	<i>Infocom05</i>	<i>Infocom06</i>	<i>Hong-Kong</i>	<i>Reality Mining</i> BT
Duration (days)	3	4	5	246
Granularity (seconds)	120	120	120	300
Number of experimental Devices	41	78	37	100
Number of internal Contacts	22,459	182,951	560	54,667
Rate of contact (only internal contacts)	1.10	1.70	8.6×10^{-3}	15×10^{-3}
Number of external Devices	223	4,649	831	N/A
Number of external Contacts	1,173	11,630	2,507	N/A
Rate of contact (incl. external contacts)	1.39	2.53	98×10^{-3}	N/A

Table 1: Characteristics of the four experimental data sets

contact rate and they go through periods of complete disconnection that might sometimes last during more than one day (e.g. node 1). We also notice some periods of high contact rate, where the node is always in reach of one or several devices. These periods are usually rare and short in *Hong Kong* and *Reality Mining*. In contrast, nodes in *Infocom05* are almost always in a high contact period, except at night.




Figure 7: Distribution of contact duration

Figure 7 plots distributions of contact duration for the four data sets. This figure shows that contact duration vary a lot in all traces (from a couple of minutes up to several hours). Above 75% of contacts (143,502 contacts) in *Infocom06* are only one slot long (i.e. 2 minutes). This can be partly explained by the sampling effect mentioned in §5.1. However we still find around 0.4% (765 contacts) that are longer than one hour. This has two important consequences: First, a path computation technique representing each contact as an interval of time, rather than a collection of discrete instantaneous contacts, should scale more easily. Second, it is an issue whether the contacts that are short are actually useful to carry data. Removing those short contacts may actually impact the diameter of the network. We discuss these points further in §6.2.

5.3 Properties of Delay-optimal Paths

Following the methodology described earlier, we compute the sequence of optimal paths found between any source and destination in the network, varying the maximum number of hops in a path between 1 and 4, and infinity. Figure 8 shows, for a given source-destination pair in *Hong Kong*, the delivery function for different maximum number of hops (i.e. intermediate relays).




Figure 8: Example of a delivery function, for a given source-destination pair during *Hong-Kong*, for paths with various number of hops.

In the example shown here, there is no path from the source to the destination when using paths with less than 3 hops. When 4 hops can be used, the number of optimal paths increases to 5. We see also that there is no optimal path with more than 4 hops; the delivery function is identical when the maximum hop is set to 4 or infinity. [3] contains similar results on other traces.

5.3.1 Distribution of delay

From the sequence of delay-optimal paths we deduce the delay obtained by the optimal path at all time. We combine all the observations of a trace uniformly among all sources, destinations, and for every starting time (in




Figure 9: CDF function for the optimal transmission delay, observed for all source-destination pairs.

seconds). We present this aggregated sample of observations via its empirical CDF in Figure 9; the maximum hop-number varies between 1, 4 or 6, and infinity. Note that the value of the CDF for a given time t is equal to the probability to successfully find a path within time t , when sources, destinations and message generation time are chosen at random. If no path exists, we include an infinite value in the distribution. Under each figure, we indicate the value of the diameter at confidence level 99%. We present the distribution on a [2 minutes, 1 week] time period, as paths with larger delay are not likely to be of any use (about this point, and for the diameter seen at every timescale, see [3]).

We notice first, that for all time-scale and all data sets, the difference of the CDF between 4 to 6 hops and unlimited hops is very small (as found by the value of the diameter). Second, we do observe some difference across data sets. The *Infocom05* data set is by far the best connected: a direct contact to the destination may be found within 1 day with probability 65%, whereas this is the case in less than 3% for both other data sets. In addition, the relative improvement introduced by using paths with several hops is not the same depending on the time-scale and the data sets. This improvement seems almost negligible in *Hong-Kong* for small timescales (less than an hour), whereas in *Infocom05* this is where it is the largest. For large timescales (more than 6 hours), we have the exact opposite. We conjecture that it is related to the contact rate, or contact intensity, between the participating nodes (high in conference, as may be seen in Table 1, low in other data sets) and explore this further in §6.1. Note that we have observed variation inside the same data set as well, for instance when studying the CDF of the minimum delay during day time only (see [3]). The result confirms the correlation between multi-hop delay improvement at small time-scale and high contact rate.

These results validate the small value of the diameter

compared to network size, for all the environments we have studied, as in random temporal networks. We now wish to study further the impact of the contact rate.

6. MOBILITY CHARACTERISTICS AND NETWORK DIAMETER

The previous section assumes that all the contacts found in a trace can be used to exchange information between two nodes. In practice, some contacts may not be available for forwarding, because of a collision or if these contacts are too short. In this section we apply a “contact removal” technique to a mobility trace: Each contact is either kept or removed according to a given rule fixed in advance. In a second step the diameter and the delay for this network are measured, using the same methodology as before. One could also add or move contacts in a trace, but we choose here not to do that because it seems harder to interpret results where some contacts are created artificially. We use in this section the second day of *Infocom06*, to start from a data set that contains a large number of contacts.

6.1 Contact rate

We vary the rate of contacts in a network by removing each contact independently with the same fixed probability p . Figure 10 plots the empirical CDF of the minimal delay, or delay obtained with optimal paths, for *Infocom06*, before and after probabilistic contact removal with $p = 0.9$ and $p = 0.99$. The results shown are the average of 5 independent experiences.

As expected, removing contacts deteriorates the delay performance, especially for small time-scale. The probability to reach a destination within 10 minutes drops from 35% to 0.2% when 99% of the contacts are removed, while the probability of success within in 6 hours decreases from 90% to 15%. On the other hand random contact removal does not seem to impact the




Figure 10: Empirical CDF of minimum delay when contacts are removed randomly (*Infocom06*)

diameter of the network, which remains under 5 hops. Moreover, as contacts are removed the improvement introduced by using several intermediate hops becomes less important at small time-scale, and remains important at large time-scale. This confirms the intuition that improvement at small time scales is related to high contact rate values (as seen in §5.3.1).

6.2 Contact duration

We now remove each contact if it lasts less than t seconds, where t is a fixed threshold. Figure 11 presents the empirical CDF of the minimal delay after we remove contacts that last less than 2, 10 and 30 minutes (respectively 75%, 92% and 99% of contacts removed). Contacts that are less than 2 minutes are those with a device seen during a single scan. When these contacts are removed, the success probability is halved at any time scale, all the results we have obtained remain.

Interestingly the diameter changes when we only keep contacts that last more than 10 minutes. First note that the probability to find a path with delay less than 10 minutes remains above 5%, while it was 2% when we had removed 90% of contacts randomly. In other words, keeping only the longest contacts maintain more available paths within a small delay. However this comes at the cost of an increased diameter.

One explanation for this phenomenon is that contacts are with different types: long contacts with less mobile nodes, or familiar people, and short contact where people are met that are belonging to any other part of the group. This result suggests that opportunistic schemes have to take advantage of short contacts (less than 10 minutes), not only because there are many, but also because those may help to keep the diameter small.

As a summary, Figure 12 plots the diameter (with confidence level 99%) for each delay separately. As suggested before, we observe that the diameter decreases with delay for high contact rate. In contrast, the diam-

eter increases with delay when the contact rate is low. In between, we find an intermediate regime where the diameter might be larger for a narrow range of time. We conjecture this is because the network remains connected but lacks shortcuts between far-away nodes.

7. CONCLUSION

This work establishes the existence of the so-called “small-world” phenomenon in opportunistic mobile networks. From a theoretical perspective, we have proved that the diameter grows slowly with the network size in a simple random case. We have also analyzed multiple human contact traces and observed that the network diameter generally varies between 3 and 6 hops, for networks containing from 40 up to a hundred nodes. This result holds for sparse and dense networks, and the diameter varies only a little when contacts are removed. This result has important impacts on how to design forwarding algorithms in opportunistic networks. In particular, it indicates that messages can be discarded after a few number of hops without occurring more than a marginal performance cost.

We now describe some important future extensions for this work. First, note that Corollary 1 proves that the expected number of paths becomes large under supercritical condition, but it does not prove that a path exists almost surely. Proving this last result is more difficult, and beyond the scope of this paper as the classical concentration method known as “second moment” (e.g. p.54 in [8]) does not apply. Second, this paper proves that short paths generally exist between any two nodes, but it does not indicate whether these paths can be found efficiently by a distributed algorithm using local information in the nodes. This topic has been addressed on static graphs with a more complex model [11], but the extension of these results to temporal networks remains an open problem. Third, the traces we used are unlikely to follow the simple prop-




Figure 11: Empirical CDF of minimum delay when short contacts are removed (*Infocom06*)




Figure 12: Diameter as a function of delay

erties of random temporal networks. Extending these results to study the impact of memory and heterogeneity in contact processes on the diameter of the network is an important research direction to pursue.

8. ACKNOWLEDGMENT

We would like to gratefully acknowledge Roch Guerin, Don Towsley and George Varghese for their insightful comments on early version of this article.

9. REFERENCES

- [1] B. Bui-Xuan, A. Ferreira, and A. Jarry. Computing shortest, fastest, and foremost journeys in dynamic networks. *International Journal of Foundations of Computer Science*, 14(2):267–285, 2003.
- [2] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, J. Scott, and R. Gass. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Trans. Mob. Comp.*, 6(6):606–620, 2007.
- [3] A. Chaintreau, A. Mtibaa, L. Massoulié, and C. Diot. The diameter of opportunistic mobile networks. Technical Report CR-PRL-2007-07-0001, Thomson, 2007. www.thlab.net/tr/CR-PRL-2007-07-0001.pdf.
- [4] N. Eagle and A. Pentland. Reality mining: Sensing complex social systems. *Journal of Personal and Ubiquitous Computing*, 10(4):255–268, 2005.
- [5] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Trans. on Net.*, 10(4):477–486, 2002.
- [6] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of ACM MobiCom*, 2004.
- [7] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Proceedings of ACM SIGCOMM*, 2004.
- [8] S. Janson, T. Luczak, and A. Ruciński. *Random Graphs*. John Wiley & Sons, 2000.
- [9] T. Karagiannis, J.-Y. Le Boudec, and M. Vojnovic. Power law and exponential decay of inter contact times between mobile devices. In *Proceedings of ACM MobiCom*, 2007.
- [10] D. Kempe, J. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. In *Proceedings of ACM STOC*, 2000.
- [11] J. Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proceedings of ACM STOC*, 2000.
- [12] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proceeding of ACM SIGKDD*, 2005.
- [13] M. McNett and G.M. Voelker. Access and mobility of wireless PDA users. *SIGMOBILE Mob. Comput. Commun. Rev.*, 7(4):55–57, 2003.
- [14] E. Nordström, C. Diot, R. Gass, and P. Gunningberg. Experiences from measuring human mobility using bluetooth inquiring devices. In *Proceedings of Workshop ACM MobiEval*, 2007.
- [15] A. Orda and R. Rom. Shortest-path and minimum-delay algorithms in networks with time-dependent edge-length. *J. ACM*, 37(3):607–625, 1990.
- [16] M. Papadopoulou and H. Schulzrinne. Seven degrees of separation in mobile ad hoc networks. In *Proceedings of IEEE GLOBECOM*, 2000.
- [17] V. Srinivasan, M. Motani, and W. Tsang Ooi. Analysis and implications of student contact patterns derived from campus schedules. In *Proceedings of ACM MobiCom*, 2006.
- [18] X. Zhang, J. Kurose, B. N. Levine, D. Towsley, and H. Zhang. Study of a bus-based disruption tolerant network: Mobility modeling and impact on routing. In *Proceedings of ACM MobiCom*, 2007.
- [19] W. Zhao, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proceedings of ACM MobiHoc*, 2004.