

Density-Aware Routing in Highly Dynamic DTNs: The RollerNet Case

Pierre-Ugo Tournoux, *Student Member, IEEE*, Jérémie Leguay,
Farid Benbadis, John Whitbeck, *Student Member, IEEE*,
Vania Conan, and Marcelo Dias de Amorim, *Member, IEEE*

Abstract—We analyze the dynamics of a mobility data set collected in a pipelined disruption-tolerant network (DTN), a particular class of intermittently-connected wireless networks characterized by a 1-D topology. First, we collected and investigated traces of contact times among thousands of participants of a rollerblading tour in Paris. The data set shows extreme dynamics in the mobility pattern of a large number of nodes. Most strikingly, fluctuations in the motion of the rollerbladers cause a typical *accordion* phenomenon—the topology expands and shrinks with time, thus influencing connection times and opportunities between participants. Second, we show through an analytical model that the accordion phenomenon, through the variation of the average node degree, has a major impact on the performance of epidemic dissemination. Finally, we test epidemic dissemination and other existing forwarding schemes on our traces, and conclude that routing should adapt to the varying, though predictable, nature of the network. To this end, we propose DA-SW (Density-Aware Spray-and-Wait), a measurement-oriented variant of the spray-and-wait algorithm that tunes, in a dynamic fashion, the number of a message copies to be disseminated in the network. The particularity of DA-SW is that it relies on a set of abaci that represents the three phases of the accordion phenomenon: aggregation, expansion, and stabilization. We show that DA-SW leads to performance results that are close to the best case (obtained with an oracle).

Index Terms—Delay-tolerant networking, mobility traces, opportunistic networking, routing.

1 INTRODUCTION

IN disruption-tolerant networks, node mobility has a major impact on the performance of communication protocols. It is, therefore, fundamental to reach a fine-grained understanding of how nodes move. Existing synthetic mobility models (such as the omnipresent random waypoint) are unable to characterize realistic scenarios [2]. For this reason, many initiatives have provided the research community with real traces; examples are human mobility in conferences [3] and on university campuses [4]. The analysis of such traces has allowed better characterization of node mobility and its impact on the dynamics of the network topology. The main observations are based on the concepts of contact and intercontact times, which generally follow power law distributions.

The RollerNet experiment studies a class of DTNs that follows a pipelined shape. It gathers groups of up to 15,000 people that go rollerblading Sunday afternoon in Paris. During three hours, rollerbladers travel about

20 miles, covering a large portion of the city, as shown in Fig. 1. They are guided by staff members and assisted by public safety forces. Participants pause for about 30 minutes after the first half of the course; some of them leave, others join. Providing these participants with ad hoc communication capabilities would be highly appreciated. Applications include announcements of itinerary hazards, participant search, and gaming, to cite a few. Because of the extreme mobility and intermittent connectivity (as we will see in Section 3), the rollerblading tour is a typical scenario of disruption-tolerant network (DTN) [3], [5], [6]. Such networks suffer from frequent connectivity disruptions, which means that there is no guarantee that an end-to-end path exists between a given pair of nodes at a given time.

The RollerNet experiment adopts the classic approach of logging contacts between participants of the roller tour using Bluetooth devices. Our experiment, however, introduces a new dimension to the data. Besides the contacts between the measurement nodes (iMotes from Intel), we also store contacts with cell phones that have their Bluetooth device on. Although directed, the resulting connectivity graph exhibits a number of properties that had not previously been observed.

One observation is of particular interest: the *accordion phenomenon*. The accordion phenomenon, also known as string instability or the slinky effect, can occur in pipelined sets of interconnected systems [7]. Broadly speaking, string stability means that if the initial states of all the interconnected systems are bounded then those states are uniformly bounded for all time. When such stability cannot be guaranteed, a small variation in the state of one system can greatly impact the states of the other systems. This phenomenon has been widely studied in the context of

- P.-U. Tournoux is with LAAS/CNRS, ISAE, 10 Avenue Edouard Belin—BP 54032 31055 Toulouse Cedex 4, France, and also with the UPS/INSA/INP/ISAE/LAAS, Universités de Toulouse, F-31077 Toulouse, France. E-mail: pierre-ugo.tournoux@isae.fr.
- J. Leguay, F. Benbadis, and V. Conan are with Thales Communications, 160 Boulevard de Valmy—BP 82, 92704 Colombes Cedex, France. E-mail: {jeremie.leguay, farid.benbadis, vania.conan}@fr.thalesgroup.com.
- J. Whitbeck is with Thales Communications and LIP6/CNRS—UPMC Sorbonne Universités, France. E-mail: john.whitbeck@lip6.fr.
- M.D. de Amorim is with the LIP6/CNRS—UPMC Sorbonne Universités, 4, place Jussieu, 75005 Paris, France. E-mail: marcelo.amorim@lip6.fr.

Manuscript received 24 Aug. 2009; revised 15 July 2010; accepted 25 Oct. 2010; published online 21 Dec. 2010.

For information on obtaining reprints of this article, please send e-mail to: tmc@computer.org, and reference IEEECS Log Number TMC-2009-08-0348. Digital Object Identifier no. 10.1109/TMC.2010.247.

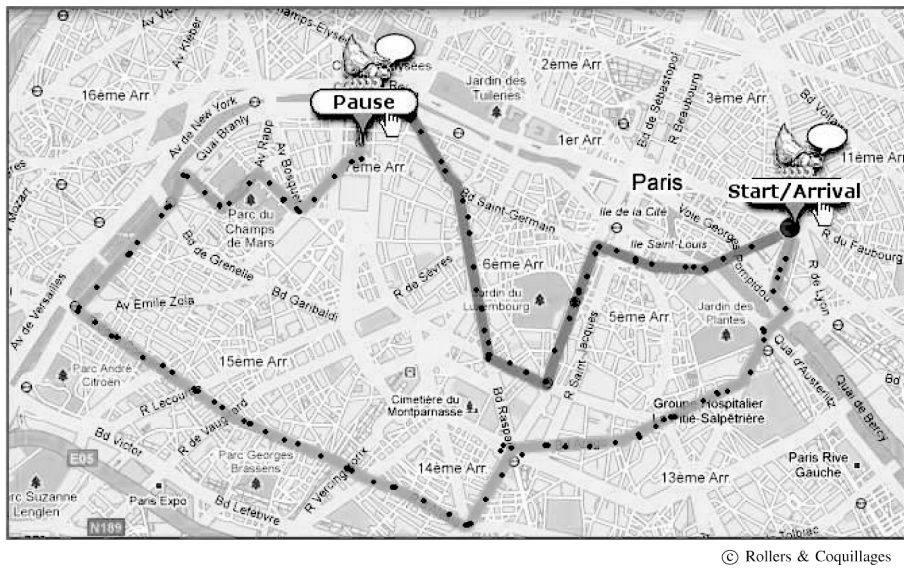


Fig. 1. This map shows an example of a rollerblading trip a Sunday afternoon in Paris. Note that the course is split into two halves with a 30-minute pause in between.

distances between vehicles on a road, as it creates moving traffic jams. If one car slows down slightly, due to the other drivers' bounded reaction time, it can create a traffic jam among the following cars. The rollerbladers' movements exhibit a very similar behavior. Indeed, not only participants have a delayed reaction to others' movements, but leaders have to adapt their speed to the slope, red lights, and crossings. The resulting accordion phenomenon translates into alternating phases of compression and expansion of the rollerblading crowd. As expected, metrics such as the average node degree, the graph diameter, and the number of connected components vary accordingly.

We focus on the impact of the accordion phenomenon on the delivery delay and traffic overhead. To this end, we evaluate two major routing strategies on our data set, namely epidemic and spray-and-wait [8]. We study a generalization of the estimation of the epidemic delay proposed by Spyropoulos et al. to the class of d -regular connectivity graphs [9]. This shows that the expected delay in epidemic routing depends on the average node degree of the graph. Because of the accordion phenomenon, the average node degree and thus the expected epidemic delay fluctuate during the rollerblading tour. The epidemic approach serves as a benchmark for our analysis; it provides both a lower-bound for delay and an upper-bound for overhead.

For the spray-and-wait strategy, both delay and overhead are directly related to the number of copies disseminated in the spray phase. Indeed, in the absence of any time constraints, spraying a small number of copies may suffice. However, if a target value for the average delay is set, then one must disseminate a larger number of copies at the cost of a greater overhead. Unfortunately, due to the accordion phenomenon, the delay varies considerably for a given number of copies. This results in an often very poor delay-overhead tradeoff.

To counter this problem, we propose DA-SW (Density-Aware Spray-and-Wait), an adaptive variant of the spray-and-wait algorithm that controls the communication overhead while keeping the delay within the expected bounds. The idea behind DA-SW is to exploit the accordion

phenomenon and tune the dissemination effort to respond adequately to the density variations. More specifically, the source sends more copies of a bundle when the topology is sparse and fewer copies when the topology becomes denser. Our results reveal that DA-SW completely counters the accordion phenomenon and dynamically adapts to the current conditions of the network when using a density oracle.

In short, the contributions of this paper are:

1. *New traces.* We present a new data set obtained through real measurements in a rollerblading tour. Because the characteristics of our scenario are inherently different from existing traces, we identify a number of properties that have not yet been observed in other works.
2. *Epidemic in d -regular graphs.* We propose a model for the evaluation of the expected delay of epidemic routing in d -regular connectivity graphs.
3. *An adaptive routing.* We propose and evaluate a new routing strategy adapted to the mobility observed in our experiment.

This paper extends prior work [1] by presenting significant additional contributions on:

1. *Contact traces visualization.* Based on dynamic graph visualization techniques, we show how the network topology evolves according to mobility patterns.
2. *Positioning.* To estimate a node's position in the crowd, we rely on two techniques based, respectively, on transmission delay [1] and graph-drawing inference [10].
3. *Selection of routing paradigms.* As the rollerblading tour plots a high level of interactions between participants and a dense network topology compared to prior DTN data sets (see Section 6), one may question the use of the DTN rather than the MANET [11] paradigm. We show that MANET routing may fail in such dense but highly dynamic networks but that it can well complement DTN routing.

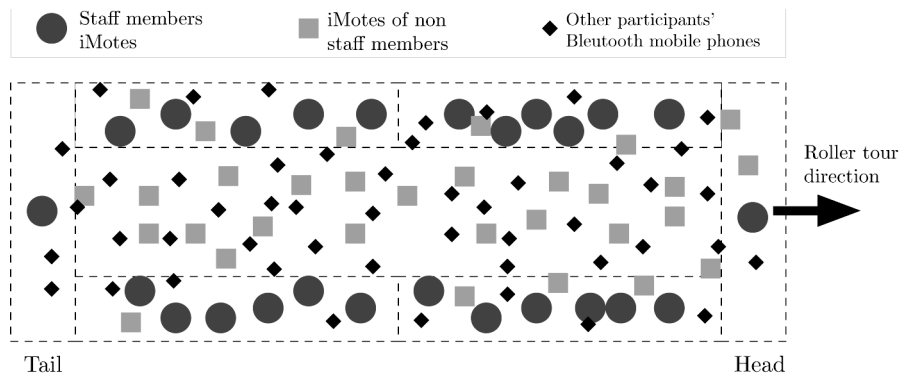


Fig. 2. Organization of the RollerNet experiment.

4. *Adaptive DA-SW based on multiple abaci.* In prior work [1], we showed that DA-SW mitigates the accordion phenomenon while suffering from delayed reactions to the evolution of the topology. To solve this issue, we propose to better anticipate changes in the topology dynamic by tuning the dissemination effort regarding the current evolution: aggregation, expansion, stabilization. In order to fine-tune the dissemination intensity, we propose to rely on four abaci, which are computed for the different network phases.
5. *Capacity limitation.* The literature on intermittently-connected mobile networks frequently assumes that contacts provide enough capacity for two nodes to exchange whatever the amount of data they have. In this paper, instead, we consider that links have limited capacity. This introduces supplementary constraints in the performance of the system. We show, however, that our approach still outperforms other solutions under such conditions.
6. *Locality.* Instead of relying on the average density of the whole network, nodes consider their local connectivity degree to estimate the number of replica to send. We observe through extensive analysis that this has little effect on the overall performance of the system, as our proposal leads to results that are close to the best case given by an oracle.

The remainder of this paper is organized as follows: In Section 2, we describe in detail the conduct of the experiment. Section 3 analyzes data collected during the tour and describes the observed accordion phenomenon. A model of the impact of topology density on the epidemic routing is presented in Section 4, while the presentation of DA-SW, our adaptive routing algorithm, is described in Section 5. Related work and previous analysis of real traces are presented in Section 6. Finally, we present our concluding remarks and identify directions for future work in Section 7.

2 THE ROLLERNET EXPERIMENT

We capture mobility of participants by deploying contact loggers in the RollerNet experiment [12]. We adopt the traditional approach that consists in deploying contact loggers (Intel iMotes in our case) on a number of participants. Every time a logger discovers a neighboring

device, it stores the time and identity of the latter.¹ Nevertheless, contrary to experiments related in the literature, we also log contacts with other participating Bluetooth devices, like cell phones or PDAs.

Out of the collected data, we characterize the interactions between people over the duration of the rollerblading tour. Such information is helpful in the design of new forms of applications in the domains of content delivery, location services, and gaming. Such applications should become available in a near future on mobile phones and will take advantage of the phones' ability to communicate directly with other phones in their vicinity, without passing through the infrastructure.

2.1 Deployment

The data set we present in this paper has been collected on 20 August 2006. According to organizers and police information, about 2,500 people participated in the rollerblading tour (rain showers just before the tour resulted in a number of participants below the average). The total duration of the tour was about 3 hours, composed of two sessions of 80 minutes, interspersed with a break of 20 minutes.

We deployed contact loggers on 62 volunteers of three types: friends of ours, members of rollerblading associations, and staff operators. In addition to the loggers, we asked other participants to activate Bluetooth on their cell phones. The iMotes also use Bluetooth technology and log periodically (every 15 seconds) the encounters they have with other devices (iMotes or cell phones).²

Fig. 2 presents a schematic view of the deployment organization. During a tour, staff members are organized into six groups. Over the duration of an outing, these groups have an almost static position relative to each other; there are two groups on each side of the crowd, one group at the front, and another one at the end (groups are represented by dashed lines in the figure). In order to distribute coverage uniformly among the crowd, 25 iMotes were entrusted among these six groups. Deploying in this manner, however, does not guarantee that the topology will be connected at all times (indeed, we are interested in studying the disruptions that emerge from changing network conditions for devices of limited range). Furthermore,

1. The iMotes are small sensor platforms with an ARM7 processor, some on board storage, and Bluetooth capability.

2. We detected more than 1,000 cell phones; some were participants of the rollerblading tour, others probably nearby pedestrians.

we made sure that among these 25 iMotes, 2 were given respectively to one staff member always staying at the tail and one staff member always staying at the front. The remaining iMotes were distributed as follows: 11 iMotes to friends and 26 iMotes to members of roller skating associations. These latter 37 devices do not have assigned positions.

2.2 Sampling Issues

All the iMotes ran the same software. They performed Bluetooth discovery for 5 seconds every 15 seconds. As using Bluetooth, two nodes scanning simultaneously are unable to see each other. In order to avoid synchronization of two iMotes around the same cycle clock, we introduced a backoff time of ± 5 seconds around the start time. Due to Bluetooth's short range, contacts are logged when people get close to each other (typically within ten meters).

Because there is a small risk that the Bluetooth inquiry may occasionally miss a contact even though it is present, we made the decision that if a contact is seen at a given inquiry I_i , but not at the subsequent inquiry I_{i+1} , we will still assume that the recorded contact was never broken if we observe it again at the following inquiry I_{i+2} . Conversely, if a contact is first seen at inquiry I_i , but not at inquiries I_{i+1} and I_{i+2} , we consider that this contact duration was equivalent to 1 second.

As this mechanism is implemented directly on the iMotes, and as we record only the beginning and ending of contacts, we drastically reduced memory consumption. Although we could not observe any intercontact time of less than two intervals, this assumption was also made in previous experiments, described in Section 6.

2.3 Data Processing

We also performed a number of post-treatments on the raw traces. First, as no hardware resets have been experienced during the experiment, we had just performed a simple time synchronization using the starting times of each iMotes. Second, as Bluetooth suffers from MAC address corruption, we filtered the data set and retained only the 1,112 MAC addresses of devices that appeared at least twice.³ Third, we anonymized the traces for privacy reasons. And finally, we associated with each iMotes a descriptor containing information about the individual carrying the node (whether she or he is a staff member or a regular participant, for instance).

3 DATA ANALYSIS AND THE ACCORDION PHENOMENON

In this section, we analyze the dynamics of the rollerblading tour. In order to avoid any biases from the traces obtained during the break phase and from the changes occurred between the rollerblading phases of the tour, we analyze the first 5,000 seconds of our data set (before the break). During this period, we recorded 39,305 contacts between the iMotes and 50,193 contacts between the iMotes and 638 external Bluetooth devices (other 413 external devices were detected during the break and the second part of the tour).

3. Some observed MAC addresses are curiously similar to iMotes' MAC addresses. Most of the time, only few bytes are different from a known iMote, e.g., 4B5F42887128 instead of 4E5F42887128.

We first focus on the *accordion phenomenon* and how it affects the speed at which information can be disseminated. For the sake of completeness, we also present results from commonly used analysis which describe pairwise interactions between nodes in contact-oriented data sets.

3.1 Accordion Phenomenon

The accordion phenomenon is due to the alternation of two mobility patterns: pause phases, when skaters wait for safety forces to block the next sequence of crossroads, and rolling phases. To understand the dynamics of network connectivity and characterize accordion effects in the tour, we use the three following metrics:

- **Average node degree.** We compute the average node degree in a discretized fashion. For each node, the node degree is the number of contacts the node had within an interval of length δ seconds. We then study the evolution of this metric across successive time intervals of δ seconds.
- **Connected components.** As the tour expands and shrinks with time, the topology disconnects. To better understand this dynamics, we measure, on graphs obtained every δ seconds, the number of connected components and the size of the largest (giant) connected one. For this measure, we consider only connections between iMote devices, and do not take into consideration isolated nodes.
- **Average delay.** This metric provides, every α seconds, the average delay required for a bundle to travel from any node i to any node j given the contact opportunities that we measured. In this case also, only iMote devices are considered; other devices do not play the role of relays.

Fig. 3 presents results for $\delta = 15$ and $\alpha = 30$. This figure presents five plots (from top to bottom) showing the evolution of:

1. the average node degree for iMote devices considering contacts with external devices,
2. the average node degree considering only contacts between iMotes,
3. the size of the giant connected component,
4. the number of connected components, and
5. the average transmission delay between all pairs of iMotes.

We observe that all these measures follow *cyclic oscillations being in phase or in opposite-phase with each other*. The average node degree is low when participants are rolling (nodes are sparsely distributed in space) and high when the tour shrinks and stops (participants tend to concentrate). It varies between 17.1 and 50.2 when considering contacts only with mobile phones, and between 2.9 and 7.8 when considering contacts between iMotes only. Due to the high number of external devices and the fact that part of them left the tour after $t = 3,000$ seconds (due to a little rain shower), oscillations appear fuzzier way when considering external devices.

Furthermore, the number of connected components and the size of the giant connected component oscillate, respectively, in the ranges 1-6 and 19-62. Note that the maximum size of the giant connected component is necessarily 62, as we have only considered iMotes.

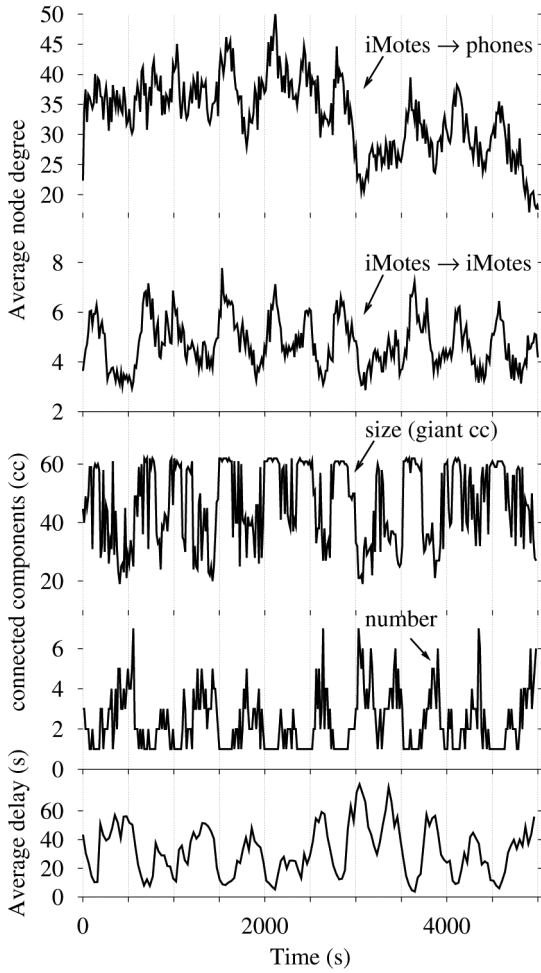


Fig. 3. Dynamics of the roller tour's topology. From bottom to top, we have: average delay, number of connected components, size of the largest (giant) connected component, average node degree when only iMotes are considered, and average node degree of iMotes with external Bluetooth devices. Those figures are valuable for $\delta = 15$ and $\alpha = 30$.

Although it is intuitive that disconnections occur more frequently when the topology is sparse, it is interesting to underline that a giant connected component dominates.

Finally, oscillations of the average delay are in opposite-phase with the average node degree, and in phase with the number of connected components. They vary between 3.7 seconds and 78.5 seconds. This result highlights two facts. First, data can be disseminated in average between any two iMotes in 32.4 seconds; our initial feeling was that this value would be much higher. This result is encouraging as it allows a large number of applications to operate under reasonable time constraints. Second, data dissemination is affected by oscillations of network connectivity, being slower when the local density of node is low or when the topology becomes disconnected.

All these synchronized oscillations characterize the so-called accordion phenomenon.

This accordion phenomenon frequently appears in road traffic on highways. This emergent phenomenon, also known as *traffic waves* or *mobile traffic jams*, is due to traveling disturbances in the distribution of cars. However, rather than coming from propagating fluctuations in relative speeds of participants, the oscillations observed in the roller tour are mainly imposed by the alternation of two




Fig. 4. Minimum delay and number of relays (when minimizing delay) from tail to head.

mobility patterns: pause phases and rolling phases. To better understand the physics of the accordion phenomenon in road traffic, we invite readers to review Kerner's work [13] on the *three phase traffic theory*, which explains how traffic conditions can alternate from free flow or congestion to moving jams.

We now look at the effects of the accordion phenomenon on an end-to-end perspective considering the two specific iMotes deployed at the head and tail of the pipeline. Fig. 4 presents the evolution of the minimum delay and the number of relays (when minimizing the delay), required for data to flow from the tail to the head. This figure shows oscillations which are similar and in phase with previous observations. 17.5 relays and 52.8 seconds are required on average. We can also see that the delay roughly varies with the number of relays.

As in a typical accordion phenomenon, acceleration and deceleration at the front propagates to the back of the tour with certain latency. This is observable in the contact traces. Fig. 5 plots the average node degree of the first third of the crowd (i.e., the head group) and the last third of the crowd (i.e., the tail group). Since nodes are not equipped with any positioning system, we used two nontrivial relative positioning techniques to place each node relatively to the head and tail nodes. The first is based on the *minimum delay* required to send data to head iMote. The longer the delay, the further they are from the head. The second is based on a *plausible* mobility inference technique, which reconstructs a realistic node mobility which could have produced the measured contact trace. This is explained in more detail in the next section.

Fig. 5 shows results for both techniques. Taken individually, each technique is debatable and leads to slightly different results due to small differences in the head and tail group compositions. However, both support the same interpretation. Fig. 5 exhibits also two phenomena. First, a spike in the average node degree in the tail group is almost always preceded by a spike in the head group. This dephasing is characteristic of an accordion phenomenon. Second, the average node degree in the tail group sometimes remains high while that of the head group drops or fluctuates. This phenomenon, particularly visible around 1,000 and 2,500 seconds, is due to the fact that good skaters are often positioned at the front so that beginners stay behind.

The dephasing of the average node degree is examined in more details in Fig. 6. The arrows emphasize the various phases (D for *deceleration*, A for *acceleration*, C_h for constant speed at high density, and C_l for constant speed at low




Fig. 5. Average node degree for head and tail groups. Two position estimation methods are used here. (a) Using inferred mobility. (b) Using message delay to head node.

density) for both the head and tail groups. Each group goes through cyclical sequences of $D \rightarrow C_h \rightarrow A \rightarrow C_l$ phases. However, the phases in the head group always precede those in the tail group. For example, when the head group's first D phase overlaps with the tail group's previous low average degree C_l phase. This happens because the tour is slowing down, and the head participants start slowing before the tail ones. Later, around 1,700 seconds, the head nodes start accelerating again, leaving the tail nodes behind. In fact, the tail nodes remain in their constant speed phase (C_h) until about 1,800 seconds when they finally start accelerating to catch up with the head nodes. Both groups begin a deceleration phase (D) after 2,050 seconds, thereby starting a new cycle. These observations highlight the dynamic effect of the accordion phenomenon on the contact trace.

3.2 Evolution of Connectivity Patterns: Some Visual Examples

In order to guide our intuition in analyzing the extremely dynamic connectivity in RollerNet, we created a real-time animation of the connectivity graph [14]. This is an application of the plausible mobility heuristic which aims at inferring node movement based only on knowledge of past and future contacts [10]. Recovering the original node movement is clearly impossible as too much information is lost when simply keeping the contact trace. However, it is possible to infer a *plausible* movement, i.e., one that could have produced the initial contact trace.

A complete description of the heuristic is beyond the scope of this paper, but the basic idea is to gradually push

nodes closer to each other up to the moment when the contact appears and push them away from each other in the absence of any contact. Furthermore, nodes all have a maximum speed (10 m/s in this paper) and past contacts restrain two nodes from moving too quickly away from each other. Other constraints, such as obstacles or borders may be added.

The plausible mobility approach works particularly well with the RollerNet trace, because of the surprising amount of spatial information hidden in it. Indeed, we know that nodes 27 and 38 remained, respectively, at the front and the back of the rollerblading tour. Furthermore, since the rollerbladers go down Paris streets, we can approximate the shape of the crowd as a 50-meter wide rectangle.⁴ The iMotes also have a fairly short range (less than 30 meter). Combining all this information with the plausible mobility inference techniques described above, it is possible to obtain an animation of the 62 iMotes movement (Fig. 7).

The various snapshots in Fig. 7 are representative of the various phases the rollerblading tour goes through (see Fig. 6). They alternate between D (Fig. 7a), C_h (Fig. 7b), A (Fig. 7c), and C_l (Fig. 7d), before beginning a new cycle with a new deceleration phase (Fig. 7e). As expected, the diameter of connected components progressively grows during the acceleration phases, eventually causing them to split into several smaller connected components.

The movement thus inferred from the contact information is likely correlated to the actual movement, though we could not test this hypothesis in the absence of a corresponding GPS trace. However, the resulting animation provides great insights into the dynamics of the RollerNet trace.

3.3 Connectivity Space: Is RollerNet a MANET or a DTN Case?

For routing, we chose to use the DTN rather than the MANET paradigm for following two main reasons:

Disconnection. While end-to-end connectivity may be present at times (Figs. 7b, 7d), the network topology is typically composed of several disjoint components (Figs. 7a, 7c, 7e). As noted before in Fig. 3, end-to-end connectivity is correlated with high average node degree. However, dramatic network splits can occur even under high average node degree, as evidenced in Fig. 7c. The DTN paradigm adapts

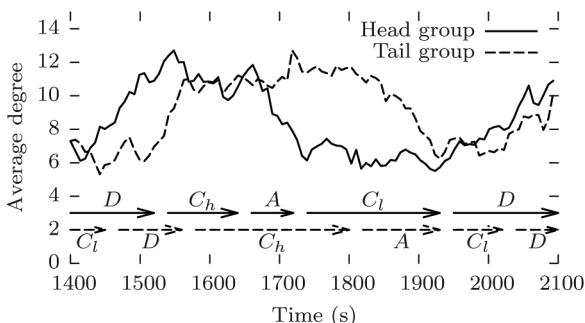


Fig. 6. Zoom on period between 1,400 s and 2,100 s on Fig. 5. A stands for acceleration period, C_h for constant speed at high density, C_l for roughly constant speed at low density, and D for deceleration.

4. This is probably slightly overestimated. Most avenues on the path of this rollerblading tour range from a dozen meters wide to the 70-meter wide Champs-Élysées.

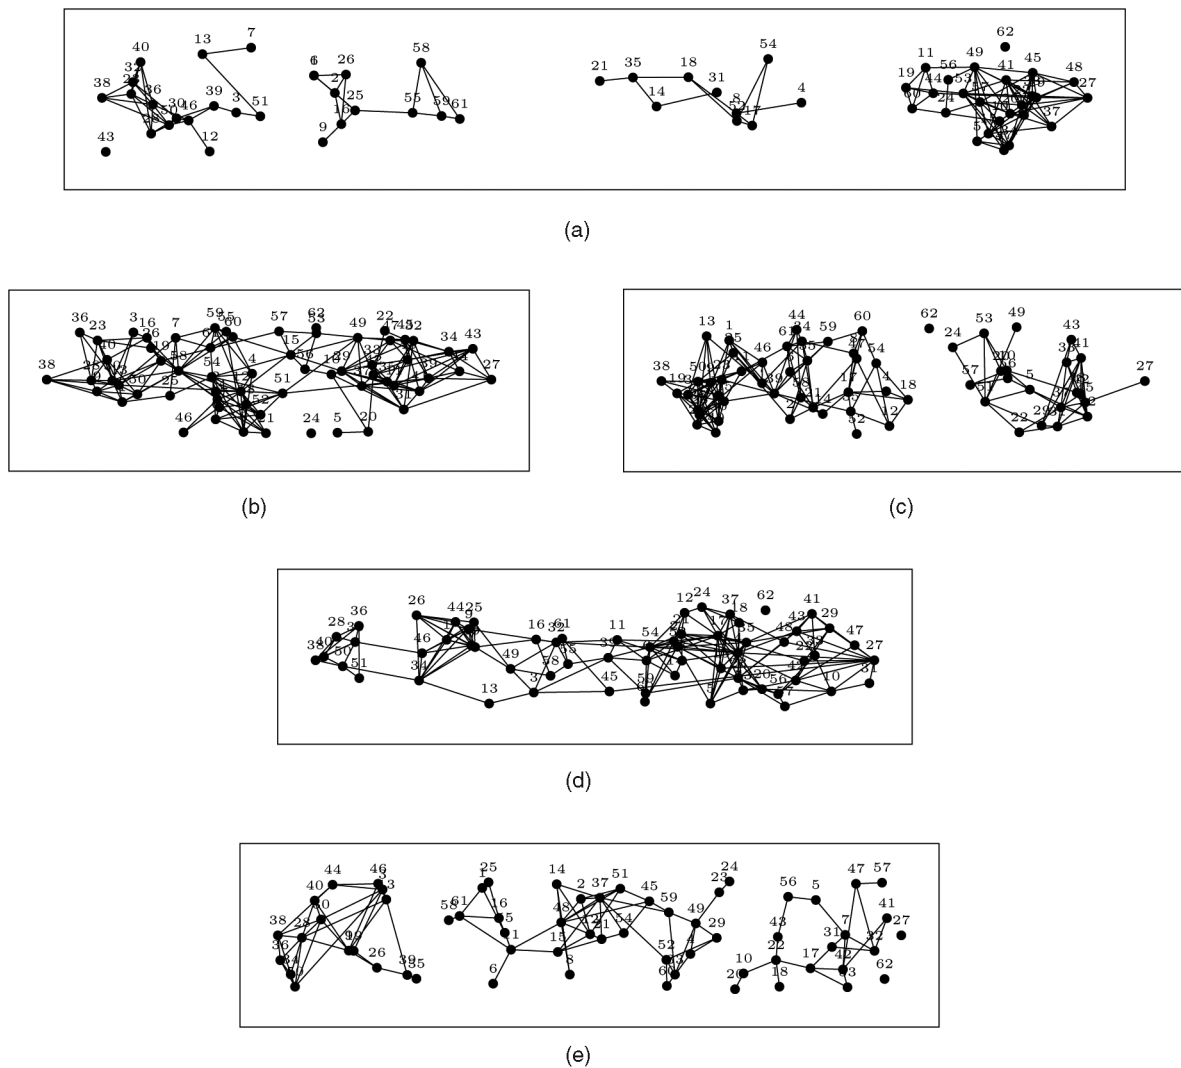


Fig. 7. Snapshots of the connectivity trace at different times. Each snapshot corresponds to one of the head group phases of Fig. 6. Nodes 27 and 38 are, respectively, at the front and at the back of the rollerblading tour. (a) 1,450 s (D). (b) 1,600 s (C_h). (c) 1,700 s (A). (d) 1,900 s (C_l). (e) 2,050 s (D).

seamlessly to both the connected and disconnected phases, whereas MANET will fail whenever a network split occurs.

Speed of topology change. MANET may also fail even during the connected phases simply because the topology evolves too quickly for a traditional MANET routing protocol to keep up. Indeed the sheer instability of the links (on average 11 links go up or down every second) would result in a deluge of topology updates in the proactive case and *route error* and *new route request* messages in the reactive case. Furthermore example, in the case of a standard proactive protocol such as OLSR [15] with its recommended parameters, there is an 8-9 second delay between the moment a link physically goes down and when that information is propagated through the network. Given the speed of topology change, each OLSR router will, therefore, have a very incorrect view of the current network topology. In the case of a reactive protocol such as AODV [16], by the time a route has been established (and packets can begin to be send), there is a high probability that it will already have been invalidated, thereby stopping the packet flow and causing a reflooding of the network to repair the route. DTN protocols, on the other hand, handle high mobility naturally, by focusing on per encounter events.

More generally, routing paradigms for mobile wireless networks are diversely adapted to the parameters of node density and node mobility. In Fig. 8, which maps the different routing approaches on the bi-dimensional mobile wireless network space, traditional MANET routing appears in the top left corner. When the density of nodes diminishes, end-to-end connectivity can disappear. In such sparse networks, nodes have very few, if any, neighbors within their transmission ranges. The topology eventually




Fig. 8. Mobility versus density: When different paradigms apply.




Fig. 9. Pairwise average contact and intercontact time distributions.

splits into several noncommunicating connected components. This is typically the realm of delay-tolerant networking which one can further subdivide in two [17]—the Assisted DTNs (A-DTN), in case of low mobility of nodes, or Unassisted DTNs (U-DTN) where mobility is high. The latter corresponds to traditional DTN scenarios. RollerNet is, to the best of our knowledge, the first connectivity trace to be both very dense and highly mobile.

3.4 Pairwise Interactions

We then look at pairwise interactions between nodes considering the commonly used contact times and intercontact times. For any pair of nodes i and j , their interactions over time plot instants at which they are in contact, called *contact times*, and instants at which they are not in contact, called *intercontact times*. Considering the two statistical processes which generate these two kinds of time intervals helps understanding pairwise interactions of nodes.

We compare results for the RollerNet data set with those of the popular data set introduced by Chaintreau et al. [3] in order to see how much they differ. In this other experiment, iMotes were also used to acquire proximity contacts that occurred between participants of the INFOCOM 2005 conference. People intending the conference were asked to carry one of these sensors in their pocket all the time. They collected data from 41 iMotes over three days. The devices performed Bluetooth inquiry scans every 2 minutes. For each pair of nodes (i, j) , we considered, as for RollerNet, that i and j were in contact if either one saw the other. We now use INFOCOM to refer to this data set.

Fig. 9 presents pairwise average contact and intercontact time distributions. In both RollerNet and INFOCOM, pairwise interactions are heterogeneous. We observe that intercontact times are larger than contact times being in average respectively 250 seconds and 4 seconds for RollerNet and 10,870 seconds and 100 seconds for INFOCOM.⁵ Finally, the order of magnitude of the two processes is lower in RollerNet than in INFOCOM.

In order to better characterize pairwise processes, we have tested for whether the distribution of intercontact times between any two nodes can be modeled either by an exponential, a log-normal, or a Pareto distribution. For this purpose, we use the Cramer-Smirnov-Von-Mises [18] statistical hypothesis test. Recall that such a statistical test can only *reject* or *fail to reject* a given hypothesis. So, when the hypothetical distribution is rejected by the test, we are certain that the distribution computed over the data does not match. On the other hand, when the test fails to reject the hypothesis, we only assume that the distribution matches

5. We recall that contacts observed during only one inquiry in the RollerNet data set are considered to last one second.

TABLE 1
Intercontact Time Distribution, Fitting Results

	Infocom	RollerNet
Number of pairs tested	737	1,394
Exponential	3.3 %	49.6 %
Pareto	5.0 %	6.7 %
Log-normal	99.3 %	99.8 %
None	0.7 %	0.2 %

the data to a confidence level $1 - \alpha$. We used a relatively high level of confidence ($\alpha = 0.01$) and also visually cross-checked the goodness of fits. For each pair of nodes (i, j) having at least 9 contacts, we have compared the cumulative distribution I_N^{ij} of the N intercontact times observed and the hypothesized cumulative distribution functions (CDF). Table 1 presents, for each data set, the proportion of pairs for which the distribution of intercontact times fits exponential, Pareto, and log-normal distributions. We also show the proportion of pairs that were rejected for all three hypothetical distributions. Note that, for a given node pair, several distributions may fit the intercontact distribution.

These intercontact times are found to follow a log-normal distribution. One notable observation is that log-normal tends to fit better than exponential or Pareto for all three data sets. The main reason is that the log-normal distribution offers a more versatile model to capture the variability in intercontact patterns across the different pairs of nodes. We also observe that a larger number of pairs fit an exponential distribution in RollerNet when compared to INFOCOM.

4 IMPACT OF DENSITY ON EPIDEMIC ROUTING

Because of the accordion phenomenon, node density and, consequently, node degree vary with time. We have seen that higher degrees offer more favorable conditions for data transfers, improved delivery rates, and shorter delivery times. In this section, we investigate more precisely the relationship between average degree and delivery delay.

Let us consider a DTN consisting of m nodes. We characterize the network density of contact opportunities by the average degree d of the nodes. Every node can only ever meet an average of d other nodes and it will never meet the $m - d - 1$ remaining ones. We model contact time instants between neighboring nodes as independent Poisson Processes with parameter λ . This implies intercontact times with exponential distribution, which matches half of the cases in the RollerNet data set (see Table 1). The neighbors of all nodes are chosen with uniform probability and produce a random d -regular neighborhood graph.

We are interested in computing E_{epid}^d , the average delay for transferring a bundle between any two nodes using epidemic routing. For $d \geq 3$, according to the properties of random regular graphs [19], the neighborhood graph is always *almost surely* d -connected, meaning that E_{epid}^d is well

defined. A random walk on a grid leads to a connected network (i.e., $d = m - 1$) with exponentially distributed intercontact times [9]. In this case, we have that $E_{epid}^{m-1} = \frac{1}{\lambda} \frac{H_{m-1}}{m-1}$, where $H_n = \sum_{p=1}^n \frac{1}{p}$ is the n th harmonic number. In the RollerNet data set, degrees vary between 3 and 10 for a total of 62 nodes, so we would like to study the delay of epidemic routing for $3 \leq d < 61$ (i.e., $3 \leq d < m - 1$).

The epidemic delay is computed by gradually coloring all nodes, starting from the source. After each step p , p nodes have received the bundle, and any of them can meet only $s(p)$ other nodes (their aggregate neighbors). This first encounter takes an average of $\frac{1}{p s(p) \lambda}$. At each step, one has the same probability $\frac{1}{m-1}$ of encountering the target destination. Adding all delay contributions for each of the $m - 1$ steps in the coloring procedure, we obtain the following:

$$E_{epid}^d = \frac{1}{\lambda(m-1)} \sum_{p=1}^{m-1} \frac{m-p}{p s(p)}. \quad (1)$$

In the random d -regular graph, one can compute $s(p)$ recursively, starting from $s(1) = d$ (the first node has an average of d neighbors). Suppose that the coloring process has reached step p ; let us call $N(p)$ the set of colored nodes and $S(p)$ their disjoint neighbors (the neighbors of all nodes in $N(p)$ minus the nodes themselves). The number of elements in $N(p)$ and $S(p)$ are, respectively, p and $s(p)$.

Let x be the node in $N(p)$ that is the first to meet node $y \in S(p)$. To generate the two sets at step $p + 1$, y is added to $N(p)$ and removed from $S(p)$. Let us now look at the neighbors of x ; x has at least one neighbor in $S(p)$, y . Each of the $(m - 2)$ remaining nodes (all the nodes except x and y) have the same probability of being any of the x 's $d - 1$ remaining neighbors. Thus, the average number of those $d - 1$ nodes falling in $N(p)$ is $\frac{(d-1)(p-1)}{m-2}$. Similarly, the average number of nodes falling in $S(p)$ is $\frac{(d-1)(s(p)-1)}{m-2}$.

To sum up, $S(p + 1)$ is obtained from $S(p)$ by removing y , adding the $d - 1$ neighbors of y , and removing the neighbors of y which fall in $N(p)$ and in $S(p)$, which translates into

$$s(p + 1) = s(p) - 1 + d - 1 - \frac{p-1}{n_d} - \frac{s(p)-1}{n_d}, \quad (2)$$

where $n_d = \frac{m-2}{d-1}$. Solving the recursive equation above leads to the following formula for $s(p)$, $1 \leq p \leq m - 1$:

$$s(p) = m - p - \left(1 - \frac{1}{n_d}\right)^{(p-1)} (m - 1 - d). \quad (3)$$

We can now derive the following result on the behavior of E_{epid}^d for $3 \leq d \leq m - 1$:

- For $d \ll m$, we have $E_{epid}^d \approx \frac{K(d)}{\lambda}$, where $K(d) = \frac{2}{d-2} \sum_{p=1}^{\infty} \frac{1}{p(p+d-2)}$. This means that for nodes with a limited number of neighbors, the epidemic delay is not dependent on the size of the network. The formula is obtained from a parabolic approximation of (2). In this case, the limiting factor is the slow time it takes for the epidemic dissemination to meet all the remaining nodes.




Fig. 10. Delay versus node degree for epidemic ($\lambda = 1$).

- For d approaching $m - 1$, we obtain $E_{epid}^d \approx \frac{1}{\lambda} \left(\frac{1}{d} + \frac{H_{m-1}-1}{m-1}\right)$. In this case, the contaminated set may see all the remaining nodes very quickly (after only one step in the approximation) and the limiting factor is the time to see each node one after the other.

Fig. 10 plots the epidemic delay as a function of node degree d for a network of 62 nodes and average intercontact times $\frac{1}{\lambda} = 1$ (normalized). The two extreme behaviors are shown. For the smallest degrees, the delay drops hyperbolically; for $d \geq 4$, we have $K(d) \approx \frac{2}{d}$. For $d \geq 20$, the decrease in delay gets smoother and approaches $\frac{H_{61}}{61} \approx 0.077$.

RollerNet falls in the first category with average degrees (3 to 10) much smaller than the number of iMotes (62). The limiting factor for epidemic routing is the degree, whatever the total size of the network. We observe that node density has a direct impact on the delay experienced with epidemic routing. The formula given in 1 also shows a sharp hyperbolic drop in the delay with increased node density. One can anticipate that practical routing strategies will behave in a similar way. Monitoring the density variations, especially for small node degrees, is one of the keys for performing efficient routing in DTNs.

This result comforts our strategy of designing a routing scheme in which a node dynamically selects the number of bundle copies in function of its connectivity degree. This is the subject of the following section.

5 MITIGATING THE ACCORDION EFFECT

As shown in the previous section, delay varies according to node density when using epidemic routing. Because density in RollerNet follows the accordion phenomenon, the same behavior is observed in terms of delay. We, therefore, need an adaptive strategy that mitigates the accordion effects by controlling the number of copies disseminated in the network as a function of the density. We first investigate the benefits of such a strategy with O-SW, an Oracle-based representation of spray-and-wait routing strategy [8]. O-SW represents the best result one can obtain with the different versions of the SW strategy, i.e., at each time instant we pick the best result that could be obtained by selecting the right number of copies in SW (see, Section 5-B). The results reveal that it is possible to control the average delay while reducing the overhead if the right number of copies to be sent is known. We then propose DA-SW (Density-Aware Spray-and-Wait), a distributed version of O-SW. DA-SW relies on an estimation of node degree to infer, at each instant, the proper number of copies to disseminate.

5.1 Spray-and-Wait Routing

To control flooding in DTNs, Spyropoulos et al. propose Spray-and-Wait (SW), a routing strategy in which the source distributes a number of bundle copies to relay nodes that wait to meet the destination [8]. SW faces a natural tradeoff. The number of bundle replicas is related to the probability for at least one of the relay nodes to meet the destination. This has a clear impact on the delay; nevertheless, the more copies are sent, the higher the overhead.

Among the existing methods to disseminate copies, we choose in this work the *binary spraying*, which is optimal when node movements are independent and identically distributed [8]. In binary spraying, when a node A carrying n copies of a bundle meets node B (that does not have any copy of the same bundle), it transfers $n/2$ copies to B and keeps the remaining $n/2$. A node stops sharing its copies either when it meets the destination or when it has one copy remaining. We call SW(n) the binary spray-and-wait algorithm where sources create n copies of the bundles.

5.2 Oracle-Based Spray-and-Wait

Our first proposal is an oracle-based adaptive algorithm that we call O-SW (Oracle-based Spray-and-Wait). The idea here is to show how a SW-based routing protocol which has, at each instant, a global knowledge of the current and future states of the network would behave. For this purpose, we run the SW(n) algorithm with the following values of n : {10, 16, 20, 25, 32, 40, 50, 64, 80, 100, 128}. At each time, we consider the number of copies that leads to the best result as the choice of O-SW. We show in this way that we can guarantee some expected delay if we know the right number of copies to disseminate.

To evaluate routing performance, we replayed in The One [20], a commonly used DTN simulator, the RollerNet traces (the same number of nodes, 62, and the same contacts). As we are interested in investigating maximum routing capacities, we first make simple assumptions regarding lower layers; in particular, we consider infinite bandwidth between nodes, contention-free access to the medium, and infinite buffers in the nodes. In order to capture short-term variations of network connectivity, we chose to generate bundles every 30 seconds between all pairs of iMotes. We use the following metrics to evaluate routing performance:

- **Bundle delay:** the time needed for the first copy of a bundle to reach its destination. Note that since our bundle delivery ratio is always 1, our average bundle delay calculations are not biased by discarding the delay of undelivered bundles.
- **Overhead:** the total number of transmissions needed for bundle delivery. In other words, it quantifies the overhead or the effort spent by the dissemination scheme till a bundle is delivered.

We show in Fig. 11 the results in terms of delay and overhead obtained with O-SW and four instantiations of the SW(n) algorithm ($n = \{10, 16, 32, 64\}$). We first observe that, as expected, the higher the number of copies in SW(n), the lower the delay. Note that it is possible, with O-SW, to guarantee an average delay by switching from one variant of SW(n) to another (in Fig. 11, the target delay is 70 seconds). In opposition to SW(n) protocols, O-SW leads to an overhead which follows the oscillations of node degree previously seen in Fig. 3. Its overhead varies between 9.8 (that of SW(10))




Fig. 11. Delay and overhead over time for SW(n) and O-SW.

and 25.2 (that of SW(128)). Using only 11 different values for n , O-SW achieves an average delay of 70.8 seconds with a standard deviation of 3.7 seconds. This result motivates the use of adaptive routing strategies which control the dissemination effort with regards to node density.

5.3 Density-Aware Spray-and-Wait

In the previous section, we showed that it is worth using an adaptive algorithm to control the delay/overhead. Nevertheless, obtaining the best case (i.e., with the oracle) is not attainable. To this end, we propose DA-SW (Density-Aware Spray-and-Wait), a distributed version of O-SW. The goal of DA-SW is to choose the right number n of copies of a bundle, in a dynamic fashion, based solely on local information of nodes, in order to achieve constant average delay.

Thus, each node using DA-SW possesses a number of abaci, which play the role of lookup table. For a given couple of node degree and targeted average delay, an abacus provides the right number of copies to be sent. The abaci DA-SW uses are computed from previous measurement experiments (for RollerNet, this is done once and reused in future tours, as we expect to have similar behaviors).

To build an abacus, we run a certain number of simulations with different versions of SW. Each version of SW sends a fixed number of copies of each bundle, which varies from 2 to 64. For each version (each number of copies), we set the average node degree as a parameter, which varies from 5 to 12. Then, for each couple (*number of copies, average node degree*), we measure the average delay for a message to travel from source to destination. The results are shown in Fig. 12, where curves have been derived from least square fitted simulation results. We can also observe in this figure the expected lower bound, obtained by using epidemic routing.

We can see that the ranking in delay obtained with the different versions of SW is preserved and that delay decreases faster with node degree when using more and




Fig. 12. Average delay for SW(n) and Epidemic versus average node degree.

more copies. DA-SW interpolates the right value of n , the number of copies to be sent, from this abacus using, as input, the observed current average node degree and the targeted average delay. For instance, if the current node degree is 10 and the requested transmission delay of the bundle is set to 120 seconds, the node will spread eight copies of this bundle.

Our goal is to get *as close as possible to the results obtained with O-SW*. Fig. 13 shows simulation results for DA-SW when targeting a delay of 70 seconds. DA-SW achieves an average delay of 68.1 seconds with a standard deviation of 13.9 seconds, whereas O-SW achieves an average delay of 70.8 seconds with a standard deviation of 3.7 seconds. Both DA-SW and O-SW require equivalent overheads, respectively, 19.9 and 19.1 (in number of copies). These results show that DA-SW is able to mitigate the connectivity variations caused by accordion effects to reach a given average delay.

5.4 Phase Transition Impact

As we observed in Fig. 13, DA-SW suffers from delay variations compared to O-SW due to a lack of anticipation of aggregation and disaggregation phases (respectively when node density goes up and down). Indeed, when looking at evolutions of overheads for DA-SW and O-SW, we see that they plot similar patterns, but that of DA-SW is behind schedule compared to that of O-SW.

As a consequence, node degree should not be interpreted in the same way depending of the state of the roller tour. A bundle sent when the skaters are accelerating experiences a decreasing node density on its way to destination. Similarly,

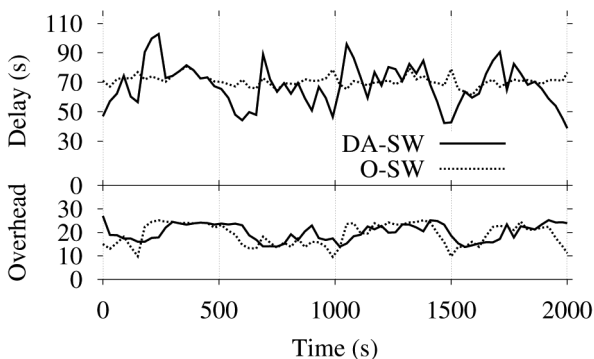


Fig. 13. Average delay and overhead over time for DA-SW and O-SW.

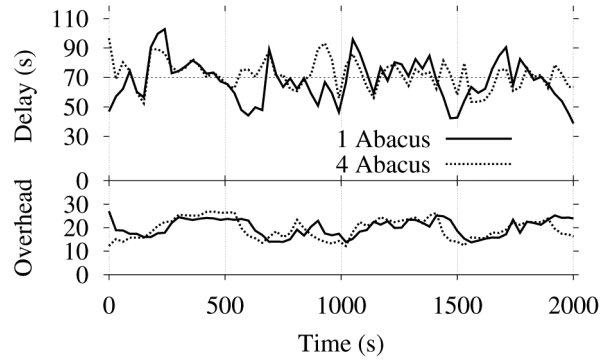


Fig. 14. Average delay and overhead over time for DA-SW using one and four abacus.

observing a constant node degree at source node should not have the same meaning whether the roller tour is stopped or not. To verify that discriminating these phases would lead to higher performance for DA-SW, we performed simulations using an a priori knowledge of the timing of phases. In reality, as mobility patterns follow cyclic variations, maintaining a history of observed node degrees would allow to discriminate these phases.

Fig. 14 compares results for the two variants of DA-SW using one and four abacus.⁶ We clearly see that the four abacus variant corrects the time shift that we observed in Fig. 13 in terms of number of copies selected between O-SW and DA-SW with one abacus. This adaptive behavior leads to a significant improvement of DA-SW accuracy especially during transition phases. It results in a smoother delay for DA-SW which now achieves an average of 71.16 seconds with a gain of 35 percent in standard deviation being now 10.3 seconds while the average overhead remains the same. It confirms that the knowledge of the current trend in the evolution of the network connectivity allows to predict more accurately the delay experienced by the bundles in the network.

5.5 Resource Limitations

The ability of a DTN to transport data does not only depend on network connectivity. It is also driven by the availability of key resources such as storage and bandwidth. This section investigates the robustness of DA-SW with regards to limitations of bandwidth. Nodes are assumed to share with their neighbors a Bluetooth bandwidth of 250KB/s. We do not consider here limitations on storage as in RollerNet the storage capacity of devices is typically several orders of magnitudes higher than the Bluetooth throughput.

We consider first the case where nodes know the average connectivity degree of the network. Fig. 15 shows the average delay over time for different load demands using DA-SW. In these simulations, 1 message of different sizes (1, 10, 20, and 30 KB) is issued at a random node every second. We observe that the average delay is, respectively, 64.5, 65.4, 68.6, and 72.3 seconds. DA-SW achieves in average a delay close to its target of 70 seconds, although it increases with the message size when the participants are rolling (i.e., when the contact opportunities are less

6. Referring to the acceleration, deceleration and constant phases with high and low node density.




Fig. 15. Average delay over time for DA-SW for different message sizes (1 message is issued at a random node every second).

frequent). In terms of throughput, DA-SW supports the input loads considered even though the average delay starts to overcome the targeted delay for message sizes higher than 30 KB. Without Bluetooth contact sampling issues, more frequent and longer contacts would have enhanced the DTN capacity and thus allowed higher message sizes. In any case, we advocate that DA-SW's performance is sufficient for RollerNet applications (we expect that short messaging will be one of the main applications to exchange practical information, such as relative position and route information).

We now consider a local estimation of node degree, simply using the total number of contacts that nodes see during the last 30 seconds. The results are shown in Fig. 16. We observe that with a message size of 20 KB and the same traffic pattern used to obtain Fig. 15, the average delay is similar. The average bundle delay is also preserved, with an average of 68.6 seconds compared to 68.0 seconds. However, the standard variation increases of 5 percent due to estimation errors variation, both positively and negatively.

6 RELATED WORK

There have been a number of efforts to collect mobility data from DTN scenarios in the latest years. The Reality Mining experiment conducted at MIT has captured proximity, location, and activity information from 100 subjects over an academic year [21]. Each participant had an application running on her/his mobile phone to record proximity with others through periodic Bluetooth scans and location using information provided by the phone on the cellular network. The UMass DieselNet project investigated the connectivity among 40 buses in western Massachusetts [6], [22]. In addition to RollerNet, other experiments with iMotes have been conducted by the Huggle project [3], [23], which explores networking possibilities for mobile users using peer-to-peer connectivity between them in addition to existing infrastructures.

For the sake of completeness, we show the similarities and differences between the RollerNet study and the previous ones. We summarize in Table 2 the main parameters and measurement results from all the experiments, extending the information provided by Chaintreau et al. [3] (*INFOCOM*) and Leguay et al. [23] (*Cambridge*). The experiments *Intel* and *Cam-U* were performed in corporate and research lab settings, with the participants being researchers and graduate students. The *INFOCOM* experiment was conducted during *INFOCOM* 2005 and the *Cambridge* experiment investigated the feasibility of a




Fig. 16. Average delay over time for DA-SW with local and global average degree estimations. Message size is equal to 20 KB.

city-wide content distribution architecture composed of short range wireless access points. We can see from Table 2 that the RollerNet data set is essentially different from existing ones. In particular,

1. it brings a new dimension to the analysis by offering a much larger number of contacts;
2. it leads to enhanced accuracy with a higher frequency of neighbor discovery beacons;
3. it uses more loggers than previous examples; and
4. it reveals that nodes see each other much more frequently than in other scenarios.

Additional works have been done to gather data that can be used, after some processing, as DTN-like data. For instance, Henderson et al. [24] have deployed one of the most extensive trace collection efforts to gather information about its Wi-Fi access network at Dartmouth College. These data have been used to characterize the mobility of users [25] or to evaluate DTN routing protocols [26]. Similar Wi-Fi based data have been used to analyze mobility such as that of ETH Zürich [27].

Concerning the routing strategy itself, there is a large number of them in the literature. PROPHET relies on a probabilistic metric calculated using history of encounters and transitivity [28]. Burns et al. use both information of contacts between nodes and visits at specific locations to determine routes [29]. The movement patterns are structured and each node learns the probability that another node can successfully deliver a message to the destination. Burgess et al. propose the protocol MaxProp in the context of a real DTN deployment on 40 buses [22]. This protocol uses meeting probabilities to find paths in association to complementary mechanisms for improving performance in terms of delivery ratio and latency, such as buffer management and transmission scheduling. Leguay et al. define a high-dimensional Euclidean space, called MobySpace, constructed upon nodes mobility patterns [26]. The specific MobySpace protocol is based on the nodes visit frequency to each possible location. Bin Tariq et al. propose Optimized Waypoints, a distributed algorithm to determine ferry routes to achieve certain properties of end-to-end connectivity, such as delay and message loss among the nodes in the network [30].

Our work definitely contributes to this domain by revealing new findings in a pipelined DTN and bringing a new data set that is far from being fully analyzed.

TABLE 2
Comparison with Previous iMote Experiments

	Intel	Cam-U	Infocom	Cambridge	RollerNet
Duration (days)	3	5	3	10	0.12
δ (mins)	2	2	2	10	0.25
Devices	8	12	41	36	62
Internal contacts	1,091	4,229	22,459	8,545	58,850
Average #					
Contacts/pair/day	6.5	6.4	4.6	1.5	280.0
External devices	92	159	197	3,586	1,051
External contacts	1,173	2,507	5,791	10,469	73,661

7 CONCLUSION AND OPEN ISSUES

The main contributions of this paper are threefold. First, we present a new data set obtained through real measurements in a rollerblading tour which falls into the class of pipelined DTNs. It exhibits a high density of contacts and, more importantly, a clear accordion behavior. To the best of our knowledge, this is the first time an experiment with DTNs reveals such a phenomenon. The RollerNet data set is available to the research community through the CRAW-DAD public archive [31]; it can be used either by the DTN community or by anyone interested in the general problem related to the accordion effect. Second, we theoretically motivate that monitoring node density variations, especially for the small node degrees, is key to performing efficient routing in DTNs, as it has a direct impact in delay on epidemic routing. Finally, we propose DA-SW (Density-Aware Spray-and-Wait), an adaptive variant of the spray-and-wait algorithm that controls the communication overhead while keeping the delay within the expected bounds. DA-SW relies on the local node degree and a set of abaci to define how many copies of a message need to be spread in the network in order to reach a certain delivery delay.

Future work along these lines includes investigating predictors for the average degree so that DA-SW would better anticipate changes in mobility patterns. To this end, we intend to perform other experiments of the same kind using other radio technologies such as Wi-Fi or Zigbee (which would improve contact sampling accuracy). Other interesting perspectives are to study and use social communities in the roller tour. Hui et al. [32] have shown that it is possible to detect social grouping in a decentralized fashion and that it can be used to make efficient forwarding decisions. Despite the bandwidth limitations we have considered in this paper, further work also remains to be done on other resource constraints in terms of node buffers and energy consumption. Balasubramanian

et al. have shown that routing strategies should fully integrate those parameters when making their forwarding decisions [33].

ACKNOWLEDGMENTS

The authors thank James Scott and Pan Hui, formerly from Intel Research, Cambridge, for the iMotes that they lent to them and for all the useful discussions. They also thank Timur Friedman and Bruno Dalouche from UPMC Sorbonne Universités for their support and technical help while preparing the deployment. Finally, they are grateful to Philippe Moulié and all the staff members from Rollers & Coquillages [34]. An earlier version of this paper appeared at INFOCOM 2009 [1]. This version is more detailed and contains significantly more results.

REFERENCES

- [1] P.-U. Tournoux, J. Leguay, F. Benbadis, V. Conan, M.D. de Amorim, and J. Whitbeck, "The Accordion Phenomenon: Analysis, Characterization, and Impact on DTN Routing," *Proc. IEEE INFOCOM*, pp. 1116-1124, 2009.
- [2] D. Johnson and D. Maltz, *Dynamic Source Routing in Ad Hoc Wireless Networks*. Kluwer Academic, vol. 353, 1996.
- [3] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms," *Proc. IEEE INFOCOM*, 2006.
- [4] M. Kim, D. Kotz, and S. Kim, "Extracting a Mobility Model from Real User Traces," *Proc. IEEE INFOCOM*, 2006.
- [5] K. Fall, "A Delay-Tolerant Network Architecture for Challenged Internets," *Proc. ACM SIGCOMM*, 2003.
- [6] UMass DieselNet, "A Bus-Based Disruption Tolerant Network," <http://prisms.cs.umass.edu/dome/umassdieselnet>, 2009.
- [7] X. Liu and A. Goldsmith, "Effects of Communication Delay on String Stability in Vehicle Platoons," *Proc. IEEE Intelligent Transportation Systems Conf. (ITSC '01)*, 2001.
- [8] T. Spyropoulos, K. Psounis, and C. Raghavendra, "Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks," *Proc. ACM SIGCOMM Workshop Delay-Tolerant Networking (WDTN '05)*, 2005.

- [9] T. Spyropoulos, K. Psounis, and C.S. Raghavendra, "Efficient Routing in Intermittently Connected Mobile Networks: The Single-Copy Case," *IEEE/ACM Trans. Networking*, vol. 16, no. 1, pp. 63-76, Feb. 2008.
- [10] J. Whitbeck, M. Amorim, and V. Conan, "Plausible Mobility: Inferring Movement from Contacts," *Proc. ACM Int'l Workshop Mobile Opportunistic Networking (MobiOpp '10)*, 2010.
- [11] S. Corson, *Mobile Ad Hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations*, IETF RFC 2501, pp. 1-13, Jan. 1999.
- [12] "RollerNet: Analysis and Use of Mobility in Rollerblade Tours," <http://www-rp.lip6.fr/rollernet>, 2011.
- [13] B.S. Kerner, *The Physics of Traffic, Empirical Freeway Pattern Features, Engineering Application, and Theory*. Springer, 2004.
- [14] J. Whitbeck, "Plausible Mobility," <http://plausible.lip6.fr>, 2009.
- [15] T. Clausen and P. Jacquet, *Optimized Link State Routing Protocol (OLSR)*, IETF RFC 3626, Oct. 2003.
- [16] C. Perkins, E. Belding-Royer, and S. Das, *Ad Hoc On-Demand Distance Vector (AODV) Routing*, IETF RFC 3561, July 2003.
- [17] V. Borrel, M.H. Ammar, and E.W. Zegura, "Understanding the Wireless and Mobile Network Space: A Routing-Centered Classification," *Proc. Second ACM Workshop Challenged Networks*, 2007.
- [18] W. Eadie, *Statistical Methods in Experimental Physics*. Elsevier Science Ltd., 1971.
- [19] B. Bollobás, *Random Graphs*, second ed. Cambridge Univ., 2001.
- [20] A. Keränen, J. Ott, and T. Kärkkäinen, "The ONE Simulator for DTN Protocol Evaluation," *Proc. Second Int'l Conf. Simulation Tools and Techniques (SIMUTools '09)*, 2009.
- [21] N. Eagle and A. Pentland, "Reality Mining: Sensing Complex Social Systems," *Personal and Ubiquitous Computing*, vol. 10, pp. 255-268, 2005.
- [22] J. Burgess, B. Gallagher, D. Jensen, and B.N. Levine, "MaxProp: Routing for Vehicle-Based Disruption Tolerant Networking," *Proc. IEEE INFOCOM*, 2006.
- [23] J. Leguay, A. Lindgren, J. Scott, T. Friedman, and J. Crowcroft, "Opportunistic Content Distribution in an Urban Setting," *Proc. ACM Workshop Challenged Networks*, 2006.
- [24] T. Henderson, D. Kotz, and I. Abyzov, "The Changing Usage of a Mature Campus-Wide Wireless Network," *Proc. ACM MobiCom*, 2004.
- [25] M. Kim and D. Kotz, "Classifying the Mobility of Users and the Popularity of Access Points," *Proc. Int'l Workshop Location- and Context-Awareness (LoCA '05)*, vol. 3479, pp. 198-209, May 2005.
- [26] J. Leguay, T. Friedman, and V. Conan, "Evaluating Mobility Pattern Space Routing for DTNs," *Proc. IEEE INFOCOM*, 2006.
- [27] C. Tudeuce and T. Gross, "A Mobility Model Based on WLAN Traces and Its Validation," *Proc. IEEE INFOCOM*, 2005.
- [28] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic Routing in Intermittently Connected Networks," *ACM SIGMOBILE Mobile Computing and Comm. Rev.*, vol. 7, no. 3, July 2003.
- [29] B. Burns, O. Brock, and B.N. Levine, "MV Routing and Capacity Building in Disruption Tolerant Networks," *Proc. IEEE INFOCOM*, 2005.
- [30] M.M.B. Tariq, M. Ammar, and E. Zegura, "Message Ferry Route Design for Sparse Ad Hoc Networks with Mobile Nodes," *Proc. ACM MobiHoc*, 2006.
- [31] J. Leguay and F. Benbadis, "CRAWDAD Data Set UPMC/RollerNet (v. 2009-02-02)," <http://crawdada.cs.dartmouth.edu/upmc/rollernet>, 2011.
- [32] P. Hui, J. Crowcroft, and E. Yoneki, "BUBBLE Rap: Social-Based Forwarding in Delay Tolerant Networks," *Proc. ACM MobiHoc*, 2008.
- [33] A. Balasubramanian, B.N. Levine, and A. Venkataramani, "DTN Routing as a Resource Allocation Problem," *Proc. ACM SIGCOMM*, 2007.
- [34] "Rollers & Coquillages," <http://www.rollers-coquillages.org>, 2011.



Pierre-Ugo Tournoux received the MSc degree from Pierre & Marie Curie University, Paris, France, in 2008, and the PhD degree from the University of Toulouse, France. He is currently a post doctoral fellow at the Institut Supérieur de l'Aéronautique et de l'Espace, University of Toulouse, France. His research interests include routing, reliability mechanisms, and transport in challenged networks. He is a student member of the IEEE.



Jérémie Leguay is a researcher at Thales Communications, Colombes, France. He received the engineering degree in 2003 from the Ecole Française d'Électronique et d'Informatique (EFREI), Villejuif, France, and the MSc degree in computer science from Linköping University in Sweden in 2004. From 2004 to 2007, he was a PhD candidate at the Computer Science Laboratory (LIP6) of Pierre & Marie Curie University, Paris, France. His current research interests are in network protocols and communication services in the contexts mobile networks, sensor networks, and large-scale distributed systems.



Farid Benbadis received the PhD degree from the University Pierre et Marie Curie, Paris, France, in 2007. He spent one year as a postdoctoral fellow at Orange Labs, France Telecom R&D, France, where he worked on content distribution mechanisms. He is currently a research engineer at Thales Communications, Colombes, France. His research interests include network protocol in wired and wireless environments. He recently developed works on spontaneous topologies and ad hoc routing protocols. Currently, he is working on wireless ad hoc/mesh networks.



John Whitbeck received the MSc degree from Paris-Diderot University, Paris, France, and the engineering degree from Ecole Polytechnique, Paris, France, in 2008. He is currently working toward the PhD degree at Universités Pierre Pierre et Marie Curie (UPMC) Sorbonne Universités, Paris, France, under a CIFRE grant with Thales Communications. His research focuses delay/disruption tolerant networks from both practical and theoretical perspectives. He is a student member of the IEEE.



uted system design.

Vania Conan received the engineering degree in 1990, and PhD degree in computer science, in 1996, from Ecole des Mines de Paris, France, and a masters' degree in artificial intelligence from Universitèrre Pierre et Marie Curie, Paris, in 1991. He is presently a senior research engineer at Thales Communications, Colombes, France, and manager of collaborative research projects. His current research topics include ad hoc, delay tolerant networking, and distrib-



Marcelo Dias de Amorim is a CNRS permanent researcher at the computer science laboratory (LIP6) of Universitèrre Pierre et Marie Curie (UPMC) Sorbonne Universités, France. His research interests focus on the design and evaluation of dynamic networks as well as service-oriented architectures. He is a member of the IEEE. More information can be found at <http://www-npa.lip6.fr/~amorim>.