Impact of Communication Infrastructure on Forwarding in Pocket Switched Networks

Anders Lindgren
dugdale@sm.luth.se
Luleå University of Technology

Christophe Diot
christophe.diot@thomson.net
Thomson Research

James Scott
james.w.scott@intel.com
Intel Research Cambridge

ABSTRACT

Recently, it has been established on multiple experimental data sets that human contact processes exhibit heavy-tailed inter-event distributions. This characteristic makes it difficult to transport data with a finite transfer time in a network of mobile devices, relying on opportunistic contacts only. Using various experimental data sets, we analyze how different types of communication infrastructure impact the feasibility of data transfers among mobile devices.

The first striking result is that the heavy tailed nature of the contact processes persists after infrastructure is introduced. We establish experimentally that infrastructure improves significantly multiple opportunistic contact properties, relevant to opportunistic forwarding algorithms. We discuss how infrastructure can be used to design simpler and more efficient (in terms of delay and number of hops) opportunistic forwarding algorithms. In addition to this, for the first time in a study like this, the communication pattern of nodes is taken into account in the analysis. We also show that node pairs that have a real-life history of communication have contact properties that are better for opportunistic message forwarding to each other than what other node pairs have.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design

General Terms

Experimentation, Design, Measurement, Performance, Theory

Keywords

Delay Tolerant Networking, Mobility analysis

1. INTRODUCTION

With the increasing penetration and ubiquity of devices such as PDAs and cellular phones that contain multiple communication interfaces as well as huge memory space, the classic communication architectures look more and more like anomalies, or like dinosaurs that cannot adapt to the very mobile and nomadic nature of today's communication services. Many communication opportunities arise where the traditional assumptions of small delays and continuous connectivity no longer hold. In such scenarios, it is possible to make use of the mobility of nodes to distribute messages throughout the network to intermediate nodes that may keep the message until they meet the destination, or get "closer" to the destination. This communication paradigm in intermittently connected networks has been called *Pocket Switched Networks* [3].

Recently, some communication architectures [2, 10] and protocols [7] have been proposed that could be applied to the situation we describe above. Unfortunately most of them could not be tested on a real-life case; but their evaluations had been conducted with simulations using simple mobility models. We take another approach in this paper; based on real traces, we study the nature of the networks and the human mobility patterns that make them useful to carry new communication services. Chaintreau et al. took an experimental approach in [3] where they analyzed traces from several wireless networks, and characterized the human mobility patterns arising from them. They have shown that the inter-contact time distribution seen between two mobile devices exhibits a heavy tail on the range [5min;1day], similar to a power-law with coefficient less than 1. Chaintreau et al. establish mathematically that this distribution is so heavy tailed that any forwarding algorithm that would only rely on multiple copies of the messages, or on waiting to see the destination, would rarely be successful in delivering the data to its destination.

In this paper, we study how this negative statement might be weakened when some level of infrastructure is deployed. It is highly likely that a network will not consist solely of mobile nodes, but that it will also contain a certain degree of infrastructure, though this might not have complete coverage of the network region. The existence of infrastructure also seems to be a intuitive way to shorten the tail of the distribution by adding some predictability and/or periodicity in contact characteristics.

If wireless access points are connected to each other (through the Internet or some other network), these access points can provide shortcuts between mobile nodes that are on physically separate locations, as long as they are both in contact with an access point. If the access points are equipped with persistent storage, mobile nodes will not even have to be in contact with the access point at the same time to be able to communicate. In Sect. 2.2, we discuss these different kinds of infrastructure in more detail. We then analyze the impact of various types of infrastructure on mobile device contact characteristics in Sect. 3. First, we analyze the metrics previously identified by [3], and also introduce and discuss new metrics which has an impact on forwarding efficiency. The analysis is done using real mobility data sets, using different wireless technologies such as Bluetooth, WiFi and GSM.

All previous studies have also made the assumption that all node pairs are equally likely to communicate. Based on additional information given with the traces, we focus in this paper on the node pairs that are likely to communication. It turns out to have a significant impact on the performance of the network.

To address this problem, we analyze multiple real-world mobile traffic traces. Some of these traces rely on infrastructure, and some do not. We find that:

- Infrastructure does not remove the power-law property of the inter-contact time distribution.
- Infrastructure does significantly increase the number of contact opportunities, and lower inter-contact times.
- Infrastructure and multi-hop neighbourhood networking both independently give orders of magnitude improvement in delay. The even better news is that when they are combined their effects are added.
- Previous analysis has been too pessimistic in considering inter-contact times between all node pairs. We show that when only considering node pairs that actually have communicated, the contact opportunities are much better.

This work is original in three ways. To our knowledge, this is the first systematic analysis of mobile device contacts properties on experimental data, encompassing multiple months of experimentation and multiple networking technologies. Second, we show how different types of infrastructure impact the communication capabilities of mobile devices. Finally, it is the first work in this area, in which the communication patterns of nodes have been considered.

2. METHODOLOGY

We now define what we call "infrastructure" and describe the data sets used for this work. The experiments have been performed under quite different circumstances, which gives us great confidence in our results.

2.1 Experimental Data Sets

To perform our analysis, we use two publicly available data sets that log human mobility with various wireless technologies. In this section, the data sets are described, and the conditions in which they were gathered are explained.

UCSD Data Set

This data set was collected at University of California, San Diego (UCSD) using 300 PDAs that were used by students on the campus [8]. These PDAs were programmed to log all reachable IEEE 802.11 access points. During the experiment, the PDAs were in contact with 524 different access points. As this data set only logs contacts between mobile nodes and access points (and not contacts between two mobile nodes) we made some additional assumption for the infrastructure-less analysis to be possible. The method we use to transform this data set will be described in the Sect. 2.3.

MIT Data Set

The Reality Mining project [5] deployed 100 smart phones to students and staff at MIT. Over a period of 9 months, these phones were running software that logged contacts with other Bluetooth enabled devices by doing Bluetooth device discovery every five minutes, as well as logging information about the cellular tower they are associated with (a total of 31545 different cellular towers were logged). This leaves us with two types of contacts collected with two different wireless technologies (in the following section identified as MIT-cell for the contact data based on the log of cellular towers and MIT-bt for the contact data based on the log of Bluetooth discoveries).

2.2 Infrastructure in PSN

In previous work, the assumption has usually been made that networks rely either on infrastructure (or at least global connectivity), or mobile nodes. In practice, it is very likely that a network will consist of a number of mobile nodes that can connect both to each other and to infrastructure such as WiFi hot spots when available. We consider two different kinds of fixed infrastructure, both of which we assume consist of wireless access points that mobile nodes can communicate with in addition to other mobile nodes.

The first form of infrastructure we analyze is commonly known as hot spots. These hot spots are connected to each other through a private network (i.e. ISP provided hot spots) or through the Internet (open home networks). We assume that the hot spot connection is of sufficient capacity so that any transfer limitations lie in the wireless channel. hot spots allow communication to take place between all mobile nodes connected to any hot spot at a given point in time. Thus, two mobile nodes do not need to be at the same location to be able to communicate with each other, but the network of access points create shortcuts that allow mobile nodes to communicate with nodes at other locations as long as they are both in contact with *any* access point.

We also consider an infrastructure where access points are equipped with persistent storage. These access points can be connected or not. Because of this storage capability, it is possible to allow a message to be passed between two non-contemporaneously connected mobile nodes. Note, however, that such multi-hop communication opportunities are unidirectional, as only the node with the earlier contact with an access point can send a message to the node with the later contact, and not vice versa (unless, of course, the first node has one more contact with an access point at a later instant).

Mobile infrastructures also exist, consisting of, for example, a data transport node attached to a public transit bus or some other scheduled vehicle. In this work, we do not address mobile infrastructure, but focus on the more commonly available fixed infrastructure.

In the rest of this paper, we will systematically compare three different scenarios: no infrastructure, hot spot infrastructure and hot spot with memory. We will not study the case where access points have storage capabilities but are not connected due to some computational complexities involved in that analysis. We do however hope to be able to study that in future works.

2.3 Dealing with Infrastructure in Data Sets

Only the MIT-bt data set contains real device to device contacts. The other to data sets (UCSD and MIT-cell) are infrastructure based. In order to analyze properties of contacts between mobile nodes in those traces, we made the following assumption: two mobile nodes are considered to have a contact when they are simultaneously connected to the same access point. This methodology was described first in [3]. While this assumption can be discussed (most notably because node contacts that do not happen close to an access point will not be logged), it is realistic and still provides valuable data. In addition, when compared to other mobility traces collected using mobile devices only, the traces processed with our methodology exhibit very similar characteristics.

We applied this assumption to UCSD and MIT-cell. The MIT-cell trace with no infrastructure can be compared to the MIT-bt one. Despite technology differences, we see Sect. 3 that this methodology does not introduces considerable bias in the properties of the data set.

Note that both UCSD and MIT-cell are interesting from an infrastructure standpoint. In MIT-cell, the cellular coverage is very ubiquitous and the cell towers network provides high coverage. Therefore, if all cell towers are used as PSN infrastructure, all devices in reach of a cell tower can see most of the other devices and the data set would be of somewhat limited interest for this study. On the other hand, the UCSD infrastructure is more dispersed and connectivity is only provided in some campus hot spots. Therefore, in all infrastructure scenarios, we will vary the amount of infrastructure available to mobile devices by randomly selecting a subset of the infrastructure in order to investigate the impact of different levels of infrastructure.

2.4 Communicating Nodes in Data Sets

In previous studies, all analysis have been done assuming that all node pairs are equally likely to communicate. For example, when studying the inter-contact time distributions, the distribution of the inter-contacts between all node pairs have been considered. This is however overly pessimistic. If two nodes never attempt to communicate, the inter-contact time between the nodes (or any other metrics relating to that node pair) is not relevant to the performance of the system. Thus, it is questionable whether that data should be included in the analysis. This has however been the only possible way to analyze the data as no other information has been available.

In addition to logging contacts between nodes, the MIT data set also logs all communication (phone calls made or received, and text messages exchanged) that users do with their phones. We use this data to extract a subset of the node pairs that have had some form of communication during the experiment. The fact that they have communicated with each other indicate that there exist some sort of social relationship between them that also makes it more likely that some data transfer between them would occur (i.e., you will most likely not send a message to someone with whom you have no existing relationship). We use the contacts between these node pairs as an additional data set to investigate the impact of only studying node pairs that actually communicate. We denote these versions of the MIT data sets as MIT-cell-comm and MIT-bt-comm. Out of the 4560 possible combinations of node pairs, 115 node pairs are included in MIT-bt-comm.

3. ANALYSIS

This section consists of three parts. First, we analyze the heavy tailed nature of the inter-contact time distribution, and how it is affected by the introduction of infrastructure. Having observed very limited changes, we investigate other parameters. The second subsection is dedicated to other contact and inter-contact properties, such as number of occurrences and duration. Then, we study the impact of infrastructure on the network connectivity and on the minimum delay needed to transfer packets between 2 nodes (making the assumption that we have a forwarding algorithm that can find the optimal path). As some of the metrics investigated in this paper deal only with direct contacts between pairs of nodes, adding persistent storage in the infrastructure has no impact. As a consequence, in such cases, we only show the performance using hot spots without persistent storage in the graphs, which can be seen as the type of infrastructure is not specified in those figures.

3.1 Inter-contact Distribution Analysis

The distribution of inter-contact times between mobile nodes has been shown to exhibit power-law characteristics [3] on numerous data sets. The plots in Fig. 1 show the complementary cumulative distribution function (CCDF) of inter-contact times for the UCSD and MIT data sets, both without infrastructure, and also with varying degrees of infrastructure of connected APs. For the MIT data sets, the distribution without infrastructure is shown both for the MIT-bt data as well as for the MIT-cell data. The assumptions described in Sect. 2.3 were applied to the UCSD and MIT-cell data to generate device to device contacts for the analysis. To reduce the number of curves, this graph only shows the case when all the access points in the UCSD data set are active as infrastructure, as the difference is very small. For clarity, the distribution is shown using a log-log scale; for reference, we included an exact power-law distribution that is linear on this scale.

The MIT data from the Reality Mining project has not been previously analyzed in this way. Thus, it is interesting to see in Fig. 1(b) and Fig. 1(c) that it exhibits similar properties as the UCSD data in that the distribution shows a clear power-law characteristic up until around 12 hours. As mentioned above, in addition to looking at the MIT-bt data where Bluetooth contacts were used, we also used the assumptions from Sect. 2.3 on the MIT-cell data to create contacts between devices. In Fig. 1(b) and Fig. 1(c), the inter-contact distribution is shown both for the Bluetooth contacts as well as the contacts created using this assumption (compare the curves "0% infrastructure, MIT-bt", and "No infrastructure, MIT-cell"). It is worth noting that the difference in the distribution is rather small between using the different methods for determining the node pair contacts, despite the very different assumptions that needed to be done for the two methods. The fact that a difference do exist between the two different distributions is likely to be due to the very different properties of the communication technologies used in the two data sets. The much larger

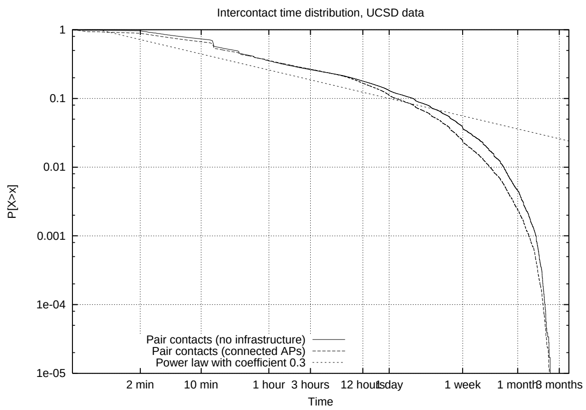

(a) UCSD data.

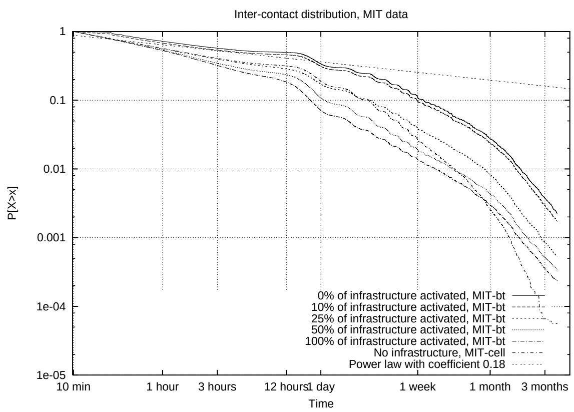

(b) MIT data.

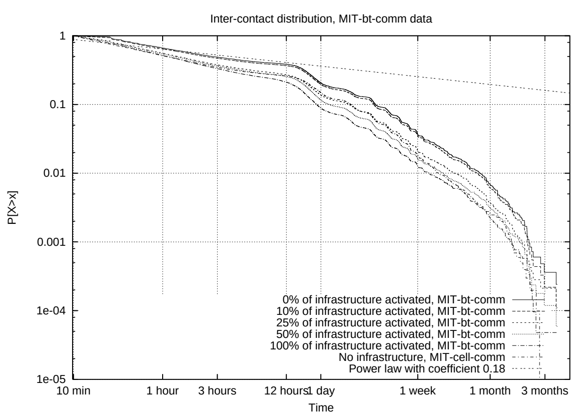

(c) MIT-comm data.

Figure 1: CCDFs of distribution of inter-contact times.

range of cellular towers compared to the range of Bluetooth will inherently add a bias such as more frequent contacts and shorter inter-contact times which can be seen in the distribution. What is of real interest is however the fact that the general shape and the overall properties of both distributions are very similar despite the very different technologies used and assumptions made. This gives us added confidence in that using simultaneous access point sightings is a good way of estimating node-to-node contacts. As the Bluetooth contacts are a more accurate measure of contacts between nodes, only those are considered for the rest of the paper. In all results where some infrastructure is present, the Bluetooth data (MIT-bt) is used for contacts between devices, and the MIT-cell contacts are used for contacts with infrastructure.

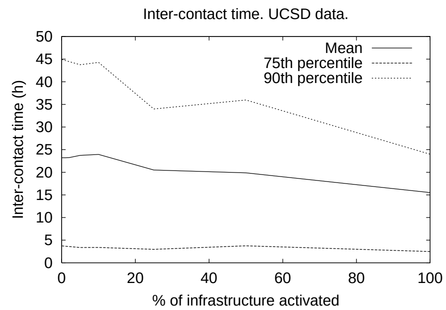

(a) UCSD data.

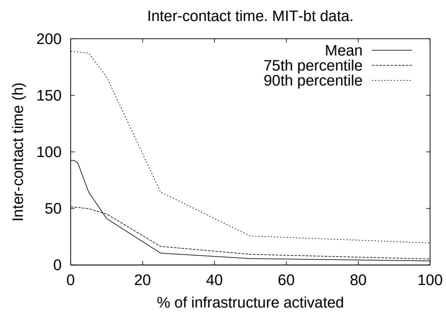

(b) MIT-bt data.

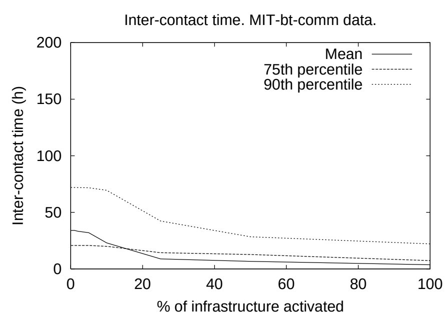

(c) MIT-bt-comm data.

Figure 2: Mean and 75th and 90th percentiles of inter-contact times.

The main motivation for introducing infrastructure into the network was the hope that it would change the properties of the inter-contact distribution. However, as observed in Fig. 1, the power-law property of the distribution is still there when the infrastructure is introduced both in the UCSD and MIT data sets. The introduction of the infrastructure do cause a visible change in the distributions, and it is notable that infrastructure reduces inter-contact times. This can be seen by the shift of the curves in the graphs as more and more infrastructure is introduced. In the UCSD case, the distributions are practically identical for inter-contacts times of 12 hours or less, but after that, the infrastructure make some impact in shortening the inter-contact times. This change is much more evident for the MIT-bt data, which should be due to the nature of the infrastructures. In the MIT-bt data, infrastructure is much more ubiquitous than for the UCSD data, and due to the assumptions used, contacts in the UCSD data only occur in the presence of infrastructure, as opposed to the the MIT-bt data that contain other contacts as well. Another significant difference between the two data sets is that inter-contacts tend to be shorter in the UCSD data set than in the MIT data set. For example, when no infrastructure is in use, approximately 75% of the UCSD inter-contact times are shorter than 12 hours, while only 50% of the MIT inter-contacts are shorter than 12 hours.

Comparing Fig. 1(b) and Fig. 1(c), we can see that while we still have the same general properties, the inter-contact distribution (especially when there is no infrastructure available) indicate shorter inter-contact times for MIT-comm than for the general MIT data. Around 10% of the inter-contacts in the general MIT data set were longer than a week, while the same is true for less than 5% of the MIT-comm inter-contacts. It is encouraging to see that for the nodes that are more likely to desire communication, the inter-contact times are shorter than for other nodes. This indicate that previous results have indeed been too pessimistic when studying the inter-contact times between all node pairs.

Even though a slight shift in the distribution is caused by the introduction of infrastructure, the slope of the distribution still remains close to the original one in the power-law region of the plot. This may seem surprising at first glance, but it is actually quite apparent why the change in the distribution is not larger. The infrastructure is likely to add many contacts between nodes, and thus also create many new short inter-contact times as new contacts are inserted. This will obviously also affect the very long inter-contact times, but here, the effect on the shape of the first part of the distribution will not be that large. If a new contact occurs in the middle of one very long inter-contact time, two new inter-contact times will be created that still are very long compared to the many very short times. Thus, the heavy-tailed properties of the distribution remains, but the tail will be somewhat shorter. We can for example see this effect for the MIT-bt data where, as infrastructure is added, the percentage of inter-contact times that are longer than one week goes from around 10% to less than 2%.

As seen above, the shape of the distributions in Fig. 1 remains mostly the same when infrastructure is deployed, and the differences that do exist seem small at first. The implications of these differences are however significant. As the amount of infrastructure increases, there is a significant drop in the mean inter-contact time. This, as well as a big drop in the 75th and 90th percentiles of the inter-contact times is shown in Fig. 2. For the MIT data, the improvement is very significant with a reduction of up to around 80 hours, or 85% of the mean inter-contact time, even with quite limited amount of infrastructure, and even larger improvements as more infrastructure is added. For the UCSD data, the decrease in mean inter-contact times is not as large as for the MIT data, but it still constitutes an 8 hour reduction of the inter-contactimes. The fact that the data have been collected in different experiments for very different infrastructure deployments explains why the MIT datasets exhibits the strongest impact when an infrastructure is gradually deployed. The percentile values are also influential on network performance as they mean that even if we have a very heavy-tail distribution, where there might exist some very long delays, we will still be able to get a communication opportunity between two nodes in a reasonable amount of time in 75% or 90% of the cases. For these percentile values, we can see even larger improvements, with reductions of up to 170 hours.

Once again, Fig. 2(c) verify that the inter-contact times between nodes that have a desire to communicate is very much lower than in the general case, especially when no or little infrastructure is in use (with more infrastructure in use, social relations between nodes become less important as more contacts can be provided by the infrastructure, and thus the improvement is not as large).

In the investigation of the impact an introduction of infrastructure would have on the inter-contact time distribution of our data sets, it was determined that the basic properties of the distribution did not change significantly. It was however, possible to see that the mean inter-contact time was decreased by a large amount as more infrastructure was added. Even when only some infrastructure was added, mean inter-contact times were reduced by up to five times. Thus, even though the inter-contact time distribution kept its power-law properties, the existence of infrastructure in the network should still improve forwarding performance in Pocket Switched Networks.

3.2 Network Properties

The inter-contact time is important as it shows what kind of delays that there can be in the system. On the other hand, the number of contacts between nodes, and even more importantly, the total duration of contacts between nodes gives an indication on how much data can be transferred in the network, and thus of the capacity of the network.

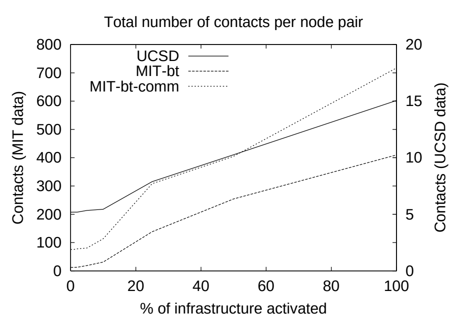

Figure 3: Number of contacts per node pair.

In Fig. 3, we show the number of contacts between each node pair, and Fig. 4 shows the total contact time between each node pair over the duration of the trace (please note that the scales on the y axis for the UCSD and MIT curves are different). As more infrastructure nodes are activated, both the number of contacts between the nodes, as well as the total duration in which they are in contact (translating into the amount of data they can exchange) increase. This supports our belief that the addition of infrastructure increases both the total capacity of the network as well as the

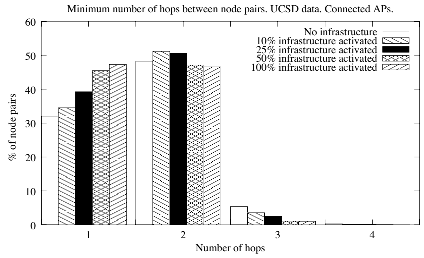

(a) UCSD data. Connected APs.

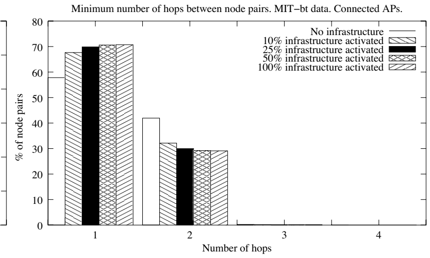

(b) MIT-bt data. Connected APs.

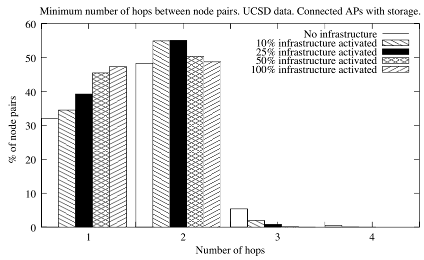

(c) UCSD data. Connected APs with storage.

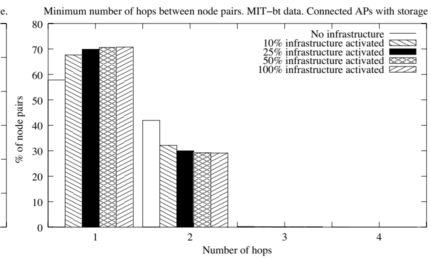

(d) MIT-bt data. Connected APs with storage.

Figure 5: Minimum number of hops between node pairs.

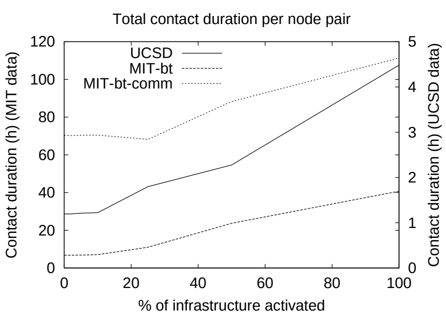

Figure 4: Total contact time per node pair.

number of communication opportunities.

Not very surprisingly, the graphs also show that node pairs that have communicated have 2-5 times as many contacts and up to seven times as long contact time than those that have not communicated. This is to be expected as the existence of communication is a sign of some form of social relationship. In everyday life, people normally spend more time together with people they have relationships with than they do with strangers. Thus, it is also natural that these nodes have more contacts.

3.2.1 Minimum number of hops

Unless all nodes meet very frequently, making direct forwarding feasible, it is believed that message delivery in a PSN will have to rely on multi-hop forwarding. This kind of forwarding opportunistically makes use of other intermediate nodes that it has contacts with to forward a message to its destination. When introducing multi-hop forwarding mechanisms, the overhead of forwarding a message several times is also added. Therefore, to know which tradeoffs to make in the design of the forwarding algorithm, it is of interest to study the minimum number of hops between any given node pair.

Figure 5 shows histograms of how many node pairs that can reach each other with a certain maximum number of hops. Since the communication can be asymmetric, each node pair is counted twice, once in each direction, as the hop count might be different (and communication might even only be possible in one direction). The MIT nodes are already very well connected (almost everybody has the possibility to communicate with everybody else even without infrastructure), but we still see that the addition of infrastructure reduces the number of hops required between many of the node pairs. In the MIT-bt-comm data set, almost all node pairs have a minimum hop count of one, so therefore no plot for that data set is included here.

For the UCSD data, the connectivity is not as complete as for MIT-bt. Therefore, we can see that the introduction of infrastructure significantly increases the number of node pairs that have the possibility to communicate. Furthermore, it also reduces the number of hops required for nodes that already had communication possibilities.

These results show that infrastructure can both enable communication between node pairs that previously had no possible communication path, and also reduce the number of hops between two nodes, effectively reducing the necessary resource requirements of the system. It is very interesting to note that almost all nodes can reach each other in only two hops, and if a forwarding algorithm were to use a maximum of four hops, it would be almost guaranteed to reach all nodes. It is however important to remember that these figures only consider the shortest possible path in terms of hop counts, so there might exist faster (shorter delay) paths for getting a message between two nodes that involve more hops than the minimum. This is discussed in more detail in Sect. 3.3.

3.3 Delay analysis

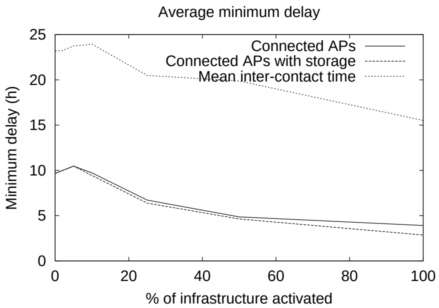

(a) UCSD data.

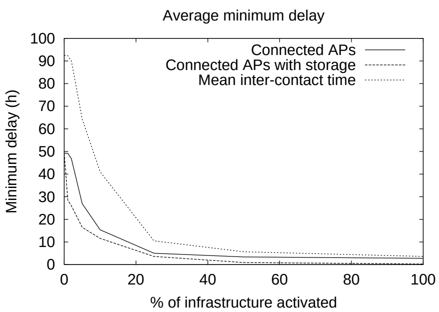

(b) MIT-bt data.

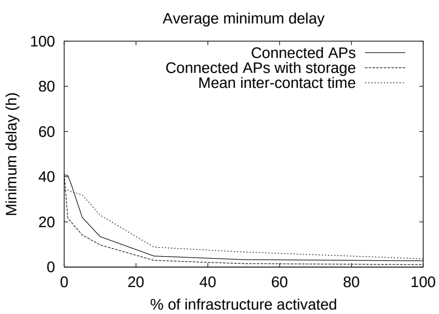

(c) MIT-bt-comm data.

Figure 6: Minimum expected delay between node pairs.

Applications that will run in a P*o*cket Switched Network have to be designed to be delay tolerant, but it is still vital for the perceived user experience to keep the delay as low as possible. The inter-contact times indicate what kind of delays to expect when only relying on direct communication between source and destination. Using a good forwarding algorithm and multiple hops through the network, it might be possible to achieve better performance. In this section, we calculate the minimum delay for delivering a message between node pairs at any given point in time, assuming intermediate nodes can be used and that an optimal forwarding algorithm is in place, a network characteristic not previously studied. Figure 6 shows the average minimum delay between node pairs for varying degrees of infrastructure. We also plot the mean inter-contact time in the same graph for comparison to see how large the improvement is when multiple hops are allowed. The average delay for messages decreases between 50% and 90% as the amount of infrastructure increases. As expected, we see that the impact of the infrastructure is even larger when the infrastructure nodes are equipped with persistent storage as this creates more communication opportunities.

Comparing the curves for the multi-hops delays with the inter-contact times, we see that the multi-hop delays are significantly lower than the inter-contact times. This means that there tend to exist multi-hop paths in the network that are more efficient than direct communication in terms of delay. Such a path exist whenever there are a series of intermediate nodes that are able to deliver a message from a source to its destination before a contact between the source and destination occurs. In many cases, the use of a multi-hop path cut the expected delivery time in half, or reduce it even further. A consequence of this is that there is a need for good and efficient forwarding algorithms that are able to make use of these communication opportunities effectively. While this is true regardless of whether or not any infrastructure is present, the largest gain appears when the combination of multi-hop communication and some infrastructure is used. As an example, we can see that for the MIT-bt data, the average minimum delay can be reduced from the more than 90 hours with only direct communication to around 5 hours when multi-hop communication is allowed and 25% of the infrastructure is active. This is a very significant reduction, and more importantly, it is likely to make the difference between the network being useable or not for many applications. While a delay of around 5 hours would be tolerable to many users, very long delays of above 90 hours are less likely to be acceptable. These curves also show the benefits of having persistent storage present in the infrastructure. At all times, this added storage makes lower delays through the network possible to achieve.

It is also interesting to note in Fig. 6(c) that the improvements caused by using multi-hop paths in the MIT-bt-comm data set was much smaller than for the other two data sets. Since the nodes in that data set have a social relationship, they meet quite frequently. Thus their inter-contact times are short enough that it is not very common for a multi-hop path to exist that can offer a significantly lower delay.

4. RELATED WORK

Research about Delay Tolerant Networks (DTN) and intermittently connected networks, which Pocket Switched Networks is a special case of, has grown tremendously lately. Most work in the area has however been mostly geared towards presenting new communication architectures or protocols, and less on more fundamental analysis of underlying issues, though some such research exist.

Chaintreau et al. [3] observed a power-law property in the distribution of inter-contact times in a number of experimental data sets. The authors mathematically proved that this property will cause certain simple stateless forwarding algorithms to have unbounded worst case delays. This highlights the heavy-tailed nature of the inter-contact times stemming from human mobility. This is very different from for example the exponential inter-contact times generated by the popular random way-point model used in many simulation studies. Other work has shown that user behavior such as association times, amount of data transer and inter-session time in infrastructure wireless networks also exhibit heavy tail properties [1, 4, 6].

Previous work has also introduced a number of different kinds of infrastructure, both fixed and mobile, to support the protocol or communication system presented in that work. The DAKNet [9] system of consists fixed Internet kiosks in villages in developing countries in conjunction with data collection units on buses or motorcycles that perform all data transfer. In the Shared Wireless Infostation Model (SWIM) [11], a network of intermittently connected nodes is used to gather oceanographic data from sensors attached to whales through an infrastructure of buoys to which the whales can upload their data. In [12], a system is proposed where a mobile infrastructure of *message ferries* is in use, which are dedicated nodes that use their mobility to improve system performance. These infrastructures has however all been specialized for a particular case, and no more general analysis of the impact of infrastructure has been made. Thus, our work is significantly different from all of the previous work. Instead of proposing the use of one particular type of infrastructure to solve a specific task, we identify the different possible types of infrastructure likely to be present in a Pocket Switched Network, and analyze what the possible impact of introducing that into the network can be.

5. CONCLUSIONS

We have shown experimentally that communication infrastructure does not modify the power-law shape of the contact process identified by Chaintreau et al. in the context of the mobility of wireless devices. However, we show that the introduction of some infrastructure, makes is much easier to forward messages opportunistically as it (1) increases the number of opportunities and (2) reduces significantly the high percentile of transmission delay. Adding some storage capabilities to communication infrastructure increases even more the number of opportunities to reach a given destination. Our results also show that allowing multi-hop communication will give orders of magnitude improvement in delivery delay.

This is also the first time in a study like this that the communication pattern of nodes is taken into account in the analysis. We are able to show that node pairs that have a real-life history of communication have contact properties that are better for opportunistic message forwarding to each other than what other node pairs have. This also means that previous analysis have been too pessimistic in considering inter-contact times between all node pairs.

However, the power-law nature of the inter-contact time keeps constraining the type of forwarding algorithms that can be designed. Given paths exist, and that 90% of them have acceptable delays (for most applications, below 5 hours or so), the challenge now is to design forwarding algorithm that can find these path, avoiding messages to be caught in an extremely long inter-contact between two intermediate nodes. We believe that we can use the contact history of nodes to decide whether or not an intermediate node can forward a message toward its destination.

Acknowledgements

The authors would like to thank Augustin Chaintreau for his comments and feedback on the work.

6. REFERENCES

[1] M. Balazinska and P. Castro. Characterizing mobility and network usage in a corporate wireless local-area network. In *Proceedings of the First International Conference on Mobile Systems, Applications, and Services (MobiSys 2003)*, 2003.

[2] V. G. Cerf, S. C. Burleigh, A. J. Hooke, L. Torgerson, R. C. Durst, K. L. Scott, K. Fall, and H. S. Weiss. Delay-tolerant network architecture. Internet Draft draft-irtf-dtnrg-arch-02.txt (work in progress), March 2003.

[3] A. Chaintreau, P. Hui, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket switched networks: Real-world mobility and its consequences for opportunistic forwarding. Technical Report UCAM-CL-TR-617c, University of Cambridge, Computer Lab, February 2005.

[4] F. Chinchilla, M. Lindsey, and M. Papadopouli. Analysis of wireless information locality and association patterns in a campus. In *Proceedings of IEEE International Conference on Computer Communications (INFOCOM'04)*, 2004.

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

[6] R. Jain, A. Shivaprasad, D. Lelescu, and X. He. Towards a model of user mobility and registration patterns. *Mobile Computing and Communications Review*, 8(4):59–62, 2004.

[7] A. Lindgren, A. Doria, and O. Schelén. Probabilistic routing in intermittently connected networks. In *Proceedings of The First International Workshop on Service Assurance with Partial and Intermittent Resources (SAPIR 2004)*, August 2004.

[8] M. McNett and G. M. Voelker. Access and mobility of wireless pda users. *Mobile Computing Communications Review*, 9(2):40–55, April 2005.

[9] A. Pentland, R. Fletcher, and A. A. Hasson. A road to universal broadband connectivity. In *Proceedings of the 2nd International Conference on Open Collaborative Design for Sustainable Innovation (dyd 02), Bangalore, India*, Dec 2002.

[10] J. Scott, P. Hui, J. Crowcroft, and C. Diot. Haggle: A networking architecture designed around mobile users. In *To apper in Proceedings of the Third Annual Conference on Wireless On demand Network Systems and Services (WONS 2006)*, January 2006.

[11] T. Small and Z. Haas. The shared wireless infostation model - a new ad hoc networking paradigm (or where there is a whale, there is a way). In *Proceedings of The Fourth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2003)*, pages 233–244, June 2003.

[12] W. Zhao, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC 2004)*, May 24–26 2004.