Understanding Urban Interactions from Bluetooth Phone Contact Traces

Anirudh Natarajan, Mehul Motani and Vikram Srinivasan

Electrical & Computer Engineering, National University of Singapore

Abstract. The increasing sophistication of mobile devices has enabled several mobile social software applications, which are based on opportunistic exchange of data amongst devices in proximity of each other. Examples include Delay Tolerant Networking (DTN) and PeopleNet. In this context, understanding user interactions is essential to designing algorithms which are efficient and enhance the user experience. In our experiment, users were handed Bluetooth enabled phones and asked to carry them all the time to log information about other devices in their proximity. Data was logged over several months, with over 350,000 contacts logged and over 10,000 unique devices discovered in this period. [1] This paper analyzes this data by charactering the distributions of metrics such as contact time and inter-pair-contact time, and introducing several other important metrics useful for understanding user interactions. We find that most metrics follow a power law, except for inter-pair-contact time. We also look for patterns in user interactions, with the hope that these can be exploited for better algorithm design.

1 Introduction

The increase in the capabilities of mobile communication devices has led to a plethora of interesting new applications. One category of these applications are Delay Tolerant Networking (DTN) applications such as Haggle [3], PeopleNet [8] and Serendipity [4]. The performance of these DTN application critically depends on patterns of user interactions. Previous work such as [3], [4] and [6] have attempted to study and characterize these interactions. However these studies were limited in scope and scale. Our experiment was done in the urban setting of Singapore where mobile phone penetration rates are exremely high. Reports from the Infocomm Development Authority of Singapore (IDA) [10], quote it to be 95.5%, more than a year ago. While we cannot exacty estimate the Bluetooth phone penetration rate, from our data we were able to find one new Bluetooth device roughly every 10-15 minutes.

Select people, whom we call probes, were given Bluetooth mobile phones and asked to carry the phones all the time to log information about other devices in their proximity. The phones handed out would perform Bluetooth device discoveries every 30 seconds and log the Bluetooth name of the external device

[1] The dataset will be made available to the research community

logged, the MAC address of the device and the time of the rendezvous. Data was logged over four months, with over $350,000$ logs and over $10,000$ unique devices discovered in this period. Compared to other studies employing a similar approach, our study is of a much larger scale. We discuss these differences in detail in Section 2. Section 3 lays out the details regarding how we went about collecting the data and the challenges we faced.

In Section 4 we characterize user interactions via the usual distributions of metrics of contact time, inter-pair-contact time (between specific pairs of devices) and inter-contact time (between any two devices). We also introduce new metrics, that shed light on the way people cluster based on the notion of a meeting. A meeting is an event that occurs when a user is in contact with at least one other device for some minimum duration.

Finally, in Section 5, we look at patterns in user interactions, with the premise that predictability can lead to the design of better opportunistic algorithms for DTN-like aplications. We looked at the time series of the number of contacts over different time scales and also looked for the occurence of common contacts over these time scales.

2 Related Work

There have been two kinds of trace based studies in the recent past. The first collects and analyzes traces collected from WiFi device association patterns with access points [1, 2, 5, 7, 11]. In the second approach, which is adopted by our study, users were handed Bluetooth based devices and interactions encountered by these users were studied.

One of the largest and first studies stems from the Serendipity [4] project at MIT. In their study, 100 users at MIT were handed Nokia 6600 phones and asked to carry them around for around 9 months. Apart from logging usage patterns of the phone, such as how often different applications on the phone were used, they also studied interactions with other test users by scanning for Bluetooth devices every 5 minutes. The University of Toronto study [6], and the Haggle studies [3], gave out between 8 and 41 Bluetooth enabled devices over a course of a few days and analyzed contact and inter contact times.

Table 1 shows the difference between our studies and the others. It is clear that the size and scope of our data is much larger to the ones that have been obtained in the other studies. We discovered more than $10,000$ unique devices and had over $350,000$ logs. This is orders of magnitude larger than data collected by any of the other studies. The granularity at which we chose to do our Bluetooth discovery is at an interval of 30 seconds. Other studies use a device discovery period of at least 120 seconds. Note that one could not go much lower in granularity as the Bluetooth discovery process takes roughly 10-20 seconds to complete. Moreover, we collected data over 4 months (with 296 man days of logging) which is larger than the Haggle and Toronto studies. Although [4] had a much larger number of volunteers, who logged data over a much larger duration, we notice that the number of unique devices discovered is much smaller. We be-

	Our Study	Haggle/ Intel	Haggle/ Cambridge	Haggle/ Infocom	Toronto	Serendipity
Device	Phone	iMote	iMote	iMote	PDA	Phone
Devices participating	12 (3 static, 9 mobile)	8	12	41	23	100
Duration	4 Months	3 days	5 days	3 days	16 days	9 months
Location	City-wide	Intel Campus	Campus	Conference	Campus	City-wide
Granularity	30 secs.	120 secs.	120 secs.	120 secs.	120 secs.	300 secs.
Logs	362,599	2,264	6,736	28,250	2,802	NA
Unique devices discovered	10,673	92	159	197	N/A	2798

Table 1. A comparison of the different studies of peer-to-peer contact pattern traces between Bluetooth devices in terms of scope, duration, and amount of data collected.

lieve that this is due to the fact that at the time the study was done, Bluetooth penetration rates were much lower. Moreover, device discovery was performed only every 5 minutes. When we looked at their data we found it difficult to mine statistics such as contact duration and inter-contact times accurately.

Moreover, the Haggle study etc., analyses only contact time and inter contact time distributions. However, with the constant increase in the number of applications based on opportunistically exploiting the proximity of devices it is clear we need to analyze several additional parameters.

For example in PeopleNet [8], the goal is not to transfer information from a certain source to destination. PeopleNet is a large distributed geographic database. Queries hop around from one mobile device to another in search of information that resides on these devices. In such a scenario, it is clear that the performance of the system is determined by how people aggregate rather than precise contact patterns between pairs of devices. Another example, stems from the notion of using sensors embedded in mobile phones to gather data from the environment. In [9], the authors propose an algorithm to aggregate data from different mobile devices which exploits the aggregation patterns of people. In these contexts, it is clear that one needs to understand aggregation dynamics. In our analysis of the data, we have designed metrics to understand these aggregation patterns.

3 Methodology

To allow us to get a wide variety of data we chose 12 probes [2]. Of these 3 were static and 9 were mobile. The static devices were customized, line powered, Bluetooth access points running on embedded Linux and these were placed in

[2] Due to certain logistical constraints we could not have more than 12 probes

three of the busiest lecture theaters on National University of Singapore campus. The 9 mobile probes were chosen to get as diverse a sampling of various social behavior patterns. 5 students on campus, 2 faculty members and 2 students who lived off campus carried mobile phones with the software that logged the Bluetooth device discoveries. After collecting the data we did realize that our choices did give us a varied set of behaviors. As expected, the 2 students living off campus logged the most contacts, logging around 170 distinct devices for every man day logged. Interestingly, the static probes discovered the least number of distinct devices per day. The maximum was 13.2 distinct devices per day. This clearly highlights the importance of mobility to increasing the potential for opportunistic data relay algorithms.

The reason phones were chosen instead of iMotes was that phones are personal devices that people already have a reason to carry around. This meant that users would remember to recharge the phones and always carry it with them over long durations (months). Further, mobile phones have more than 6MB of memory whereas iMotes have only 64KB.

Having narrowed down the choice, we picked Nokia 6600 and Panasonic X800 phones as they were the most reliable. In particular, HP PDA's and Sony Ercisson's consistently logged fewer devices than the former two devices under identical conditions.

The phones and the static devices conducted Bluetooth device discoveries every 30 seconds and logged the MAC addresses, the date and the time when the device was found. The static devices were programmed to upload their data to a central MySQL server once every day. The mobile probes had to transfer their data by activating a program on their computers that would then automatically transfer the data from the PC to the central server.

The main challenge faced in collecting the data was the finite battery life. Due to Bluetooth device discovery being an energy consuming process, phones would run out of power and the logging would stop. Often phones needed to be recharged every day in order to log continuously. Another source of error was human error. Despite our persistent attempts to remind the probes to keep the logging program switched on at all times, they had a tendency to switch it on in crowded areas which skewed the data. The logging program would also crash from time to time. This error could occur a few minutes or a few days after the logging program was switched on. Despite our best efforts we were unable to avoid this error which seems to have originated from the OS of the phone. On some of the phones when the program crashed an audible beep was made which reminded volunteers to turn on the program.

Due to the format in which the data was logged we were unable to ascertain the exact times for the occurence of these errors. However, we estimate from our data that on average the mobile probes were not logging for 24.5% of the time. From interviews with our probes, these outages seem to have been random and uniformly distributed over time. While we did miss potential contacts, our logs clearly mark the beginning and ending of any period when logging was performed. During these periods all potential contacts were recorded. In this paper we make inferences over these periods and hence the inferences about contact patterns are valid. Further, due to the random nature of the outages and the long interval over which we recorded data, the patterns that we look for in Sec. 5 are not significantly affected.

4 Metrics of Interest

The following definitions are crucial to understanding the metrics presented in this section. We define two devices to be in *contact* if they are in Bluetooth range of each other. We define a device to be in a *meeting*, if it is in *contact* with *at least* one other device for more than τ_c seconds.

Contact Time: This is the duration for which two bluetooth enabled devices are in contact. Contact time is a useful metric that spans all the applications that we have in mind. If contact times are longer on an average, it implies that more data can be exchanged/forwarded during each contact. For a DTN like application, this will affect the system throughput.

As has been established in earlier studies, the contact time distribution in Fig. 1(a) follows a power law relationship. 80% of the contacts are short in duration lasting less than 9 minutes. Contacts never last longer than 3 hours. The mean slope of the distribution is 0.84642, with a variance of 0.02 across the different mobile users. In other words, contact times are independent of user behavior.

Inter-pair-contact time: This is the time duration between two successive contacts of a specific pair of devices. If a pair of devices have been in contact only once, then the inter-pair-contact time is infinity. We look at inter-pair-contact times, conditioned on the event that this pair of devices has been in contact at least twice. This metric is of particular interest in DTN applications. For example in [3], they analyze the impact of the power law characteristic of inter-pair-contact time distributions on the performance of DTN.

The loglog plot for the inter-pair-contact time is shown in Fig. 1(b). Although the loglog plot looks linear over small time scales (within a day), it does not appear to follow a power law over large time scales. Note that in this study, there are no artifacts due to either granularity or time scale of the study, which were present in the Haggle and Toronto Bluetooth based studies. The best fit line to this curve has a slope of 0.414. We found that 80% of inter pair contacts occur within 2 hours. In this case, the slopes for individual users was quite different, which reflects individuality in user behavior. In [3], the authors found the inter-pair contact time to follow the power law with a coefficient of 0.6. We suspect the differences are due to the different environments encountered by the probes.

The two metrics above were investigated in previous DTN related studies. They focus on contacts between two specific people. However, not all applications can benefit from this information. An application like PeopleNet might be interested in just meeting any other users and swapping queries with them. It would benefit from knowing how and at what frequency people come together. The meeting metrics described below allow us to understand how people aggregate in groups.

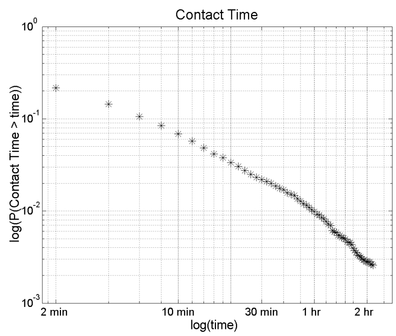

(a) Contact Time

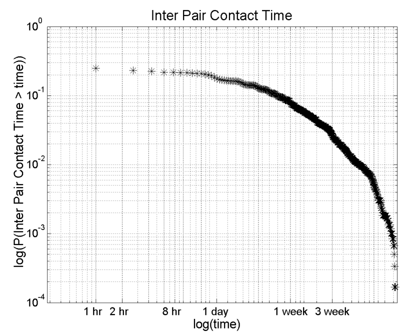

(b) Inter-Pair-Contact Time

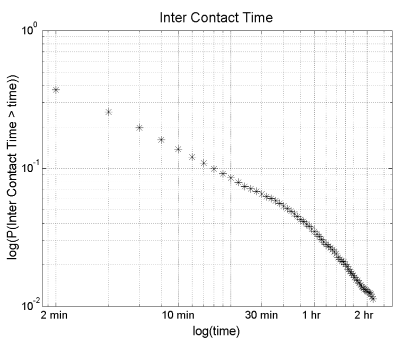

(c) Inter-Contact Time

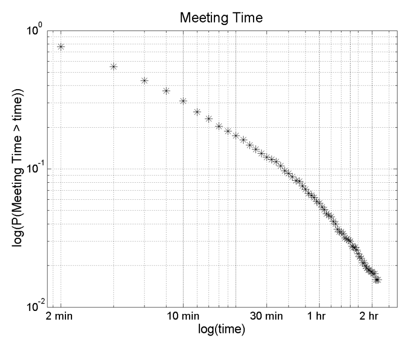

(d) Meeting Time

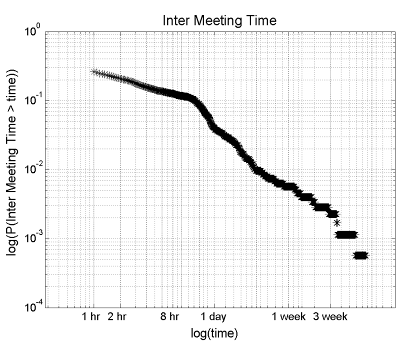

(e) Inter-Meeting Time

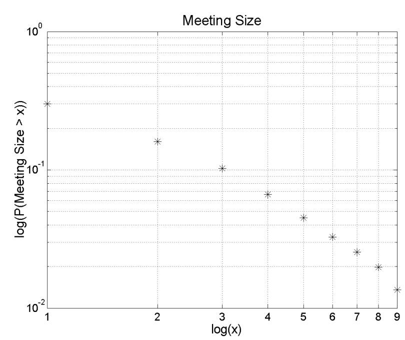

(f) Meeting Size

Fig. 1. LogLog Plots of the Distribution

Inter-contact time: This is the time between two successive contacts. For example if a device A is discovered at 1PM followed by a device B at $1:08$ PM. Then, the inter-contact time is 8 minutes. This metric allows us to make inferences about the penetration of Bluetooth devices. It also determines the frequency with which one has opportunity to exchange data with devices in the proximity. This is key to the performance of PeopleNet and Serendipity like applications. In a PeopleNet like application, this metric will determine the time to generate a match between a pair of matching queries.

From Fig. 1(c), we see that the inter-contact time does follow a power law distribution. Inter-contact times never exceed a day in duration. The slope is 0.55. 80% of inter contact times occur within 40 minutes. However, if we constrain ourselves to consider only the work day, i.e. between 8AM and 8PM most of the inter-contact times were between 4 minutes to 13 minutes for various users.

Meeting Time: This is the duration for which a device is contiguously in a meeting state. We will illustrate this with an example. Suppose node N comes in contact with device A at 1PM and stays in contact for more than τ_c seconds, then N enters the meeting state at 1PM. Assume that N is in contact with A for one hour and B comes in contact at $1:05$ PM and stays in contact with node N until $2:05$ PM. At $2:05PM$, N is in contact with no other devices. Then we say that the meeting time is 65 minutes. Meeting times allow us to understand how people aggregate in groups. This is particularly relevant to PeopleNet like applications where this data could be used to estimate how long the application will have to swap queries with other users.

For distributions related to meetings, we computed them for different thresholds τ_c ranging from 30 seconds to 120 seconds. In the following plots, $\tau_c = 30$ seconds. The conclusions are very similar for the other threshold times also. We see from Fig. 1(d) that the distributions for the meeting times is well approximated by a power law. The slope for the meeting time is 0.76 and 80% of meetings last less than 30 minutes.

Inter-Meeting Time: This is the time interval between successive meetings. Again this is well approximated by a power law with a slope of 0.39. We note that there is a similarity between the distribution of inter-pair-contact times and inter meeting times.

Meeting Size: This is the total number of devices which are in contact over the duration of a meeting. This metric too follows a power law. The maximum meeting size we have discovered is 124. Interestingly, we noted that one can make predictions regarding the meeting time by knowing the meeting size. We found that if the meeting size was 1 the average meeting time was 13 minutes, whereas for meetings of size 2 and greater, the meetings lasted 17 minutes on average. Another interesting observation is that while 80% of the meetings are of size 1, 80% of all the contacts made are made in meetings of size 2 and above. This implies that aggregation centres, have the most potential to be exploited by the DTN applications.

Average Instantaneous Meeting Size: This average is computed for each meeting, by weighting each device in the meeting by the fraction of time for which that device was in contact with the probe. This helps us to understand aggregation dynamics. For any meeting, if the average is approximately the same as the meeting size, then the environment is fairly static. On the other hand, if the average is much less than the meeting size, this implies that the environment is dynamic with devices coming in and going out (e.g., in a mall).

We found that for most meetings, irrespective of the meeting size, the average instantaneous meeting size was close to 1. This implies that users are often discovered in dynamic environments such as malls and coffee shops and are constantly moving in and out of contact. Based on this we can make the conclusion, that whenever we make contact with a new device we must try and exchange information immediately, as the device might quickly leave the meeting.

5 Looking for Patterns

With the belief that predictability in contact patterns can aid algorithm design we looked for patterns in the following two ways.

5.1 Time Series

We analyzed individual user contact patterns and plotted time series at different time scales and looked for patterns in the number of distinct devices that the probes saw on varying time scales. When looking at a PeopleNet like application, it would be helpful to know when exactly a probe is likely to find herself in a crowded area which increases the opportunity to find matches for queries.

From Fig. 2(a), which captures the number of distinct devices seen by a certain probe on a particular day of the week, we see a clear diurnal pattern. There are no devices discovered between 9PM and 9AM for this probe. We observed similar pattern for most users. Typically, we found that no devices are discovered between midnight and 8 AM in the morning. We see see a fairly large variance in the number of devices discovered at any given hour.

Next we look at the number of contacts made by a user on each day of the week. This is shown in Fig. 2(b). We see that there is a large variance in the number of users seen on a given day making predictability hard. There also does not seem to be any clear difference between week days and weekends.

The time series data that we have looked at focuses on patterns in the number of contacts made. Consider, however, DTN applications where information needs to be passed from a particular source to a particular destination. An application such as this might not be so concerned with how many contacts will be made at different times, rather with whom those contacts might be and when they are likely to be made. To address this issue of estimating the chance of meeting specific people repeatedly we introduce the notion of commonality.

5.2 Commonality

For any given time scale, the commonality is computed by looking at all subsequent pairs of times. Assume we are considering times t_i and t_{i+1} at a particular

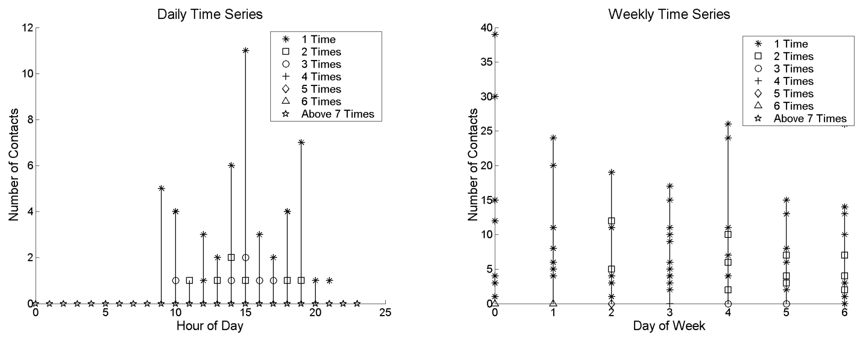

(a) (b)

Fig. 2. Fig. 2(a) shows the number of contacts in each hour across a day on a particular day of the week. Fig. 2(b) shows the number of contacts made on every day of the week. 0 represents Sunday, 1 stands for Monday and so on. The different symbols are used to refer to the number of times the particular number of contacts found. For example, in Fig. 2(b), 0 devices were found 6 times on a monday and 30 devices were found once on a sunday.

time scale. Let A_i be the set of users seen at time t_i. Then the commonality for this pair of times is given by $C_i = \frac{|A_i \cap A_{i+1}|}{|A_i \cup A_{i+1}|}$. We then average over all i.

We first looked at the fraction of identical devices seen across subsequent hours on the same day (HSD). We found that the commonality ranged between 0.1 to 0.2. When we increased the time scale to subsequent days (SD) we found it had lesser commonality. For thoroughness, we also considered a daily time scale to look at commonality between contacts at specific times on subsequent days (e.g. 10 AM everyday) (STSD), specific days of the week (DW) (e.g. all devices discovered on Thursdays) and at specific times on a weekly time scale (TDW) (e.g. devices discovered at 10 AM on Thursdays). We summarise the maximum commonality seen at all these time scales in table 2. The data indicates that the more time that has elapsed since the last time contact was made with a particular device, the lesser the chances of meeting that device again.

Time Scale	Maximum Commonality
HSD	0.2
SD	0.14
STSD	0.07
DW	0.04
TDW	0.03

Table 2. Maximum commonality across different time scales. Refer paper for the meanings of acronyms. As time scale increases, there is reducing commonality between devices discovered

6 Discussion and Conclusions

In summary, we have performed an extensive data collection and analysis experiment with Bluetooth phones, getting a true sampling of user interactions. We believe that the data significantly builds on the existing data from other studies. In analyzing the data, apart from the usual metrics such as contact time and inter-contact times, we proposed several new metrics which help to understand specific behavior such as how users cluster in groups. We must note however, that the metrics are affected by the penetration of Bluetooth devices in the environment. We also looked for patterns in the interactions for the probes. In terms of the number of contacts seen there exists a certain amount of predictability. When we looked for correlations between the devices seen at different time scales, we found very little correlation.

In our future work an important area that needs attention is making the logging process more reliable. We have succeeded in modifying the program and making it far more stable. Another step we plan is to periodically text probes to keep their logging programs switched on.

References

1. P. Bahl A. Balachandran, G. Voelker and P. Rangan. Characterizing user behavior and network performance in a public wireless lan. In *In Proceedings of ACM SIGEMTRICS 2002*, Marina Del Ray, USA, June 2002.
2. T. Gross C. Tuduce. A mobility model based on wlan traces and its validation. In *Proceedings of Infocom 2005*, Miami, U.S.A, March 2005.
3. A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Pocket switched networks: Real world mobility and its consequences for opportunistic forwarding. In *Technical Report Number 617*, February 2005.
4. N. Eagle and A. Pentland. Social serendipity: Mobilizing social software. *IEEE Pervasive Computing*, 4(2):28–34, 2005.
5. T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide network. In *Proceedings of MobiCom 2004*, pages 22–31, September 2004.
6. J. C. Cai J. Su, A. Chin and E. DeLara. User mobility for opportunistic ad-hoc networking. In *In Proceedings of IEEE WMCSA 2004*, English Lake District, UK, December 2004.
7. M. McNett and G. M. Voelker. Access and mobility of wireless PDA users. *Mobile Computing and Communications Review*, 9(2):40–55, 2005.
8. M. Motani, V. Srinivasan, and P. Nuggenhalli. PeopleNet: Engineering a wireless virtual social network. In *Proceedings of MobiCom 2005*, Cologne, Germany, August 2005.
9. M. Motani, V. Srinivasan, and O. Wei Tsang. Analysis and implications of student contact patterns derived from campus schedules. In *Proceedings of MobiCom 2006*, California, U.S.A., September 2006.
10. InfoComm Development Authority of Singapore. Telecom services statistics for may 2005. http://www.ida.gov.sg/idaweb/media/.
11. D. Tang and M. Baker. Analysis of a metropolitan-area wireless network. In *In Proceedings of ACM MOBICOM 1999*, Boston, USA, 1999.