

User Mobility for Opportunistic Ad-Hoc Networking

Jing Su, Alvin Chin, Anna Popivanova, Ashvin Goel[†], Eyal de Lara
Department of Computer Science

[†]Department of Electrical and Computer Engineering
University of Toronto

Abstract

As mobile devices become increasingly pervasive and commonly equipped with short-range radio capabilities, we observe that it might be possible to build a network based only on pair-wise contact of users. By using user mobility as a network transport mechanism, devices can intelligently route latency-insensitive packets using power-efficient short-range radio. Such a network could provide communication capability where no network infrastructure exists, or extend the reach of established infrastructure. To collect user mobility data, we ran two user studies by giving instrumented PDA devices to groups of students to carry for several weeks. We evaluate our work by providing empirical data that suggests that it is possible to make intelligent routing decisions based on only pair-wise contact, without previous knowledge of the mobility model or location information.

1. Introduction

Mobile devices continue to increase in popularity and pervasiveness. With each new generation, the capabilities of these devices continue to grow. In particular, devices are now more commonly equipped with short-range radio communication capabilities. These short-range radios, such as Bluetooth, are ideal for mobile devices because of the relatively light power requirements.

As the number of radio-equipped mobile devices increases, we observe that it might be possible to build a network based on the pair-wise contact of users and their devices. Thus instead of extending radio signal coverage to connect nodes and form links, we utilize user mobility to transport packets between nodes. Mobile ad-hoc networks (MANETs) formed by these devices could be used for sending and routing latency-insensitive packets. Potentially, such a network could give mobile devices rich communication capabilities for longer periods of time, using low-power radios as a supplement to heavier-weight established infrastructure. Not only could such a network provide a

supplemental lower cost medium to devices, but it could also extend established infrastructure. For example, this could effectively extend the reach of wireless hot-spots without resorting to cellular networks.

We observe that people usually do not move randomly [10], and in fact have organized movement patterns; furthermore, this organized movement can possibly result in regular and predictable meeting patterns between individuals. For example, Bob takes the subway every weekday morning at 8:30am. While waiting for the subway, he often stands next to or walks by the same group of strangers. The patterns can also be less serendipitous, for example Bob might have regular and predictable meetings with co-workers and friends.

If mobility and contact patterns can be exploited, then user mobility can be used for packet transport, allowing devices to make better packet routing and radio operation decisions. Though epidemic propagation [2] guarantees the shortest latency, a smarter routing and replication algorithm could significantly improve memory and power efficiency with minimal performance impact.

In this paper, we present an empirical study to determine whether pair-wise contact between users equipped with short-range wireless mobile devices can form an ad-hoc network, and whether the contact information can be used to potentially make better routing decisions. We performed our study by enlisting volunteer students to carry instrumented Bluetooth-enabled PDA devices for several weeks. Wi-Fi (802.11) was not chosen because of its heavier power consumption.

Our study makes three main contributions. First, we show that networks can be formed using user mobility, and that there is potentially exploitable regularity in the patterns. This could allow for routing and radio operation decisions that would be more efficient than epidemic propagation. Second, we share our experiences in conducting a user mobility data collection experiment. Unlike previous studies, our work does not rely on base stations, location tracking, predefined mobility patterns, or active user input. Furthermore, we instrument our experiment using consumer electronic

devices, potentially allowing any radio-equipped device to participate in this network. Finally, we identify limitations in power management mechanisms in currently available mobile devices, particularly in radio management.

The rest of this paper is organized as follows. Section 2 provides a description of the experimental design, implementation, and deployment. Section 3 describes our analysis and experimental results. Section 4 reflects on our experience gathering traces of pair-wise contacts between users. Section 5 examines related work, and finally Section 6 concludes our work.

2. Experiment and User Study

We investigate whether real user mobility and opportunistic pair-wise interactions between users can be exploited to provide data communication.

Deploying data collection devices to real users requires careful design considerations. The following section describes the design requirements of the experiment in terms of data collection and user impact. We then describe the implementation concerns and decisions with respect to the design requirements, followed by the deployment of the experiment in two separate user studies.

2.1. Design Requirements

The experiment's primary objective is to collect trace data of pair-wise contact over the duration of a working day. For the purposes of this experiment, we do not strive to transfer real data, detect connection quality, measure bandwidth, or track user¹ location.

To address the issue of real user mobility, we need to provide users with an instrumented device to carry. The instrumentation must satisfy three requirements: there should be some motivation for the user to carry the device as often as possible, the data collection should work independent of the user's activities, and the device should last at least an eight-hour work-day.

We provided users with a device that provides useful functionality to encourage them to carry the device. The instrumentation software runs invisibly in the background with minimal impact on the usability of the devices. Though we could have used specialized devices designed for this experiment, we felt that using commodity devices helps highlight our motivation of networking consumer mobile devices in interesting ways.

Our aim is to detect opportunistic pair-wise contact, even when users might not be aware of it. Contact could take place while at a meeting, waiting at an elevator, or even walking by another participant. Users may not be aware

of who may or may not be a participant, and they may not be using their devices during that moment of contact. Nevertheless, it is desirable to record such contact since it presents a communication opportunity. It is highly likely that most of the time these devices will be carried in pockets or bags. Therefore we opted to use radio, as opposed to infrared, which is omni-directional and does not require line-of-sight.

Power management is an important issue with mobile devices. Inadequate power management can render the device unusable and prevent it from gathering data. Since many mobile devices rely on disk-less storage, an extended power outage can result in lost user and experiment data. Because users carry these devices every day, they might not have time to charge the devices until they get home. Requiring users to recharge the devices mid-day would be disruptive to their daily routine and increases the likeliness of the device being forgotten or left behind. To cover a working day, we estimated that the devices should last at least eight hours, including standard usage as well as background radio operation and data gathering.

It should be noted that security and privacy are not issues as far as the experiment is concerned. Devices do not track or share user information, and the mapping of devices to users is kept confidential. The data used for analysis is anonymized before use. At this time, we also do not consider the security and privacy concerns of actually implementing a functioning network using this method. This work is primarily concerned with determining whether such a network is feasible.

2.2. Implementation

We chose to use PDA (personal digital assistant) devices, in particular Palm Tungsten T handhelds (herein referred to as Palm devices) running the Palm Operating System (PalmOS). Because sufficient battery life is a major concern, the PocketPC platform was not a viable option. Similarly, due to power concerns, we chose to use Bluetooth radio instead of Wi-Fi, though Wi-Fi is currently more commonly available. Wi-Fi can consume between 10 to 50 times more power than Bluetooth in low-usage modes² [6]. The Tungsten T devices have a built-in Bluetooth radio, which is slightly more power efficient than using an add-on card. They also can be updated with any number of available third-party applications, which helped increase its appeal to the users. To gather data, we developed a custom Palm application [19, 20] to run in the background and periodically use the radio to search for other users. Because PalmOS is a single-threaded event-driven platform, we use a self-setting alarm timer to grab background processing time. When the timer is triggered, we asynchronously operate the radio to listen for nearby devices and announce our presence. The application

1 We refer to participants of the user studies as "users".

2 On average, 90% of time in sleep mode, and 10% in RX and TX.

then sets another timer and sleeps. For most applications, this technique produces no observable hindrance to the user experience.

The frequency at which the devices announce and listen on their radios affects battery longevity. However, because we aim to capture serendipitous contact, longer sleep times may result in the device missing brief contacts. We made a best-effort attempt to have the protocol capture what we call the “walk-by”. Assuming a 10-meter radius of the Bluetooth antenna, and an average walking speed of 2 m/s, there is a 10-second window of opportunity to detect a user walking past a stationary user.

After several implementation iterations, we developed a minimal protocol to stretch the battery life to the target range. At the start of each user study, all Palm devices are NTP (Network Time Protocol) time synchronized so that radio usage can be minimized and the odds of successful communication increased [12]. The Bluetooth radio on the Palm devices are half-duplex, which required a scheme for allowing each device to announce their presence as well as listen for other nearby devices. At an established time epoch, all devices power their radio simultaneously. Each device will then listen on their radios for a random 1 to 3 seconds. Immediately after the random listen interval, the device will broadcast its presence for 3 seconds, followed by another random 1 to 3 seconds of listening. The device will then place its Bluetooth radio in a low power non-listening sleep state, and wait until the next epoch to repeat the cycle.

The randomized sleep intervals provide a crude medium access mechanism, while minimizing the amount of time that the Bluetooth radio must be powered. Though it’s possible that a pair of Palm devices that meet might pick the same random intervals and not find each other, we expect it to be uncommon.

We experimentally found that under normal user activity, a 16-second period for this protocol achieves approximately 8 to 10 hours of battery life. Though this period means we fall short of catching the “walk-by” window, increasing the frequency resulted in unacceptable sacrifice in battery life. Conversely, decreasing the frequency resulted in the devices failing to capture brief serendipitous contact.

Technical issues on platform limitations and power conservation, which led to using such a conservative and carefully managed communication protocol, are discussed in Section 4.1.

2.3. User Study

We conducted two separate user studies for our experiment. Each study involved approximately 20 students in total from two separate classes in different departments: Computer Science (CS) and Electrical and Computer Engineering (ECE). We made sure that participants were not enrolled

in both classes simultaneously, as we wish to examine how much interactivity there is between students in those classes.

We acknowledge the inherent limitations in the size and selection of our user pool. Twenty participants is not a large number, considering how the students can be anywhere on or off campus. At the very least, they are expected to meet once per week during class times, predisposing them to an a priori pattern. Despite these limitations, these initial user studies helped examine some interesting questions regarding the feasibility of using user mobility for packet transport. Do the users meet more often than just during class time? Is there a bias in which intermediate node is most successful at delivering packets to a particular destination? Is there robustness in the network or is packet transport reliant on a few nodes? Can this trace data be used to begin formulating better routing decisions that result in more efficient network usage with minimal latency impacts compared to epidemic?

Our first user study involved only graduate students during the autumn of 2003 and lasted for two-and-a-half weeks. Nine students were in a CS graduate course, eight students were in a graduate ECE course, and one student was unrelated to either of those two courses. The second user study involved only undergraduate students during the spring of 2004 and lasted for eight weeks. Ten students were in an undergraduate CS class and ten in an undergraduate ECE class.

In addition, we deployed three Palm m125 handheld devices³ which we hid throughout the computer science building. These devices are not base stations and play no special role in the experiment or network. One can think of them as users with very little mobility. The m125 devices were included in the study to help examine the following questions. If we assumed the m125 devices were base stations or stationary people, how often would users pass by one? Do they play a critical role in distributing packets through the network? For the first user study, the m125 devices were hidden near locations frequented by graduate students. For the second user study, the devices were hidden near undergraduate labs.

3. Experimental Results

Our analysis focuses on examining the connectivity, abstract capacity, and pattern of contact for the data collected in the user studies. In the following sections we will discuss the contact information seen from the collected data, followed by the methods used for determining other characteristics of the network.

3 For both user studies, 04, 05, and 09 are m125 devices.

node	adjacent to
01	03 07 11 16 17
02	03 04 05 07 09 11 14 15 16 17 20 23
03	01 02 05 06 07 11 14 15 16 17 20
04	02 07 11 14 15 16 17
05	02 03 07 11 14 15 17 20 23
06	03 07 08 10 11 12 13 14 15 16 17 18 19
07	01 02 03 04 05 06 08 09 10 11 12 13 14 15 16 17 18 20 23
08	06 07 09 10 12 13 15 18 19
09	02 07 08 11 14 15 17 20 23
10	06 07 08 11 12 13 14 18 19
11	01 02 03 04 05 06 07 09 10 12 13 14 15 16 17 18 19 20 23
12	06 07 08 10 11 13 18 19
13	06 07 08 10 11 12 18 19
14	02 03 04 05 06 07 09 10 11 15 16 17 20 23
15	02 03 04 05 06 07 08 09 11 14 16 17 20 23
16	01 02 03 04 06 07 11 14 15 17 20
17	01 02 03 04 05 06 07 09 11 14 15 16
18	06 07 08 10 11 12 13 19
19	06 08 10 11 12 13 18
20	02 03 05 07 09 11 14 15 16 23
23	02 05 07 09 11 14 15 20

Table 1. User Study 1, Adjacency

node	orig	no class	no m125s	no class, no m125	only m125
01	20	19	17	16	0
02	20	19	17	16	12
03	20	19	17	16	1
04	20	19	0	0	12
05	20	19	0	0	12
06	20	19	17	16	0
07	19	17	16	14	11
08	20	19	17	16	13
09	20	19	0	0	13
10	20	19	17	16	0
11	20	19	17	16	12
12	20	19	17	16	0
13	20	19	17	16	0
14	20	19	17	16	13
15	20	19	17	16	13
16	20	19	17	16	12
17	20	19	17	16	12
18	20	19	17	16	0
19	20	0	17	0	0
20	19	17	16	14	10
23	19	17	16	14	12

Table 2. User Study 1, Reachability

3.1. Data and Definitions

The collected data contains lists of tuples of the form $(timestamp, node, node)$, where nodes are Palm device IDs. This list is created by merging the data logs from all of the devices, and then sorting them into timestamp order. Because the logs are simply a record of detected pair-wise contact, there is no inherent ordering in the two node fields. The fields only indicate a pair that detected one another during that instance of time.

The first user study collected 64,160 tuples and the second user study collected 14,796. In Section 3.5, we discuss some of the reasons for the smaller trace size from the second user study. Note that the size of the log is not necessarily indicative of the amount of diversity in user contact. The logs contain an entry for every pair-wise contact detected after every periodic search cycle of 16 seconds. Thus two individuals sitting together for extended periods of time can generate many log records, but this communication pattern does not result in a very dynamic network. Conversely, a group of people may make frequent but short visits to each other, resulting in a short but diverse trace.

3.2. Connectivity

In this section, we examine the direct contact and reachability of the nodes. Reachability refers to a node’s ability to send a packet, via some path of intermediary nodes, to a selected destination node. The path traversed by the packet must obey the chronological ordering of node contacts found in the trace data.

Table 1 shows the direct contact (adjacency) for each of the nodes in user study 1. For any given node, the table lists the set of all nodes directly contacted over the lifetime of the data trace. The table for the second user study is not shown due to space limitations. Overall, the second user study had adjacency patterns very similar to those shown in Table 1—some nodes with very low adjacency and others very high.

node	orig	no class	no m125s	no class, no m125	only m125
01	21	21	18	18	0
02	19	19	17	16	7
03	19	19	17	17	4
04	19	19	0	0	9
05	20	20	0	0	9
06	21	21	18	18	0
07	21	21	18	17	0
08	21	20	18	16	0
09	20	20	0	0	9
10	19	19	17	17	8
11	21	21	18	17	0
12	21	19	18	17	7
13	21	19	18	16	8
14	19	19	17	17	8
15	21	20	18	16	0
16	21	21	18	17	0
17	20	20	18	18	0
18	17	16	7	6	0
19	21	21	18	17	0
20	20	20	18	18	9
21	21	21	18	17	0
22	21	21	18	17	9
23	21	21	18	17	0

Table 3. User Study 2, Reachability

Tables 2 and 3 show the number of other nodes reachable from any given node via direct contact or through intermediate nodes. We calculated the reachability by exhaustive and complete search over the lifetime of the trace. The column *orig* shows the reachability given the original trace data, with no filtering. In *no class*, we remove all traces which take place 15-minutes before, during, and 15-minutes after class times for each of the nodes. The column *no m125* shows the reachability with all entries involving the three static nodes (Palm m125) removed. The combined effect of removing class times and m125 nodes is shown in column *no class, m125*. Finally, *only m125* shows the reachability for an infrastructure setup, where nodes only communicated with the m125 nodes.

The *no class* and *no m125* columns show that the connectivity is not reliant upon class time or the three hidden nodes. We also examined the effects of independently removing each of the nodes, in turn, from the traces, and find that there is no significant drop in reachability.

These results suggest that nodes have significant contact and reachability between one another, and that the network

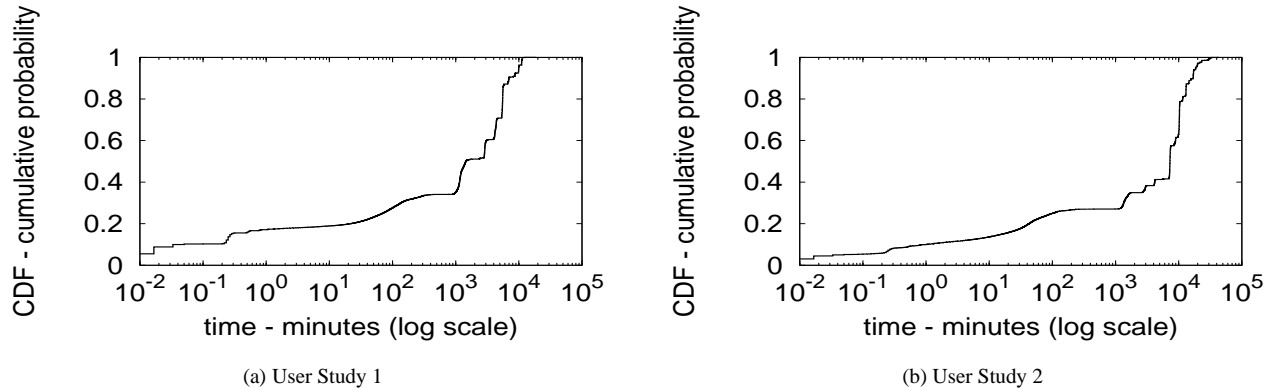


Figure 1. CDF of multi-hop packet delivery latency

does not necessarily depend on specific “hub” nodes for connectivity. Instead, there are multiple redundant paths available for reaching any particular node. Moreover, this suggests there is a measurable amount of interaction between the devices, which may be sufficient for networking.

Column *only m125* of Tables 2 and 3, compared to column *orig*, shows the potential for increased reachability if pair-wise communication is utilized. This suggests that there can be potential networking gains if pair-wise communication is utilized to extend the reach of wireless hot-spots.

Surprisingly, the *only m125* column for the second user study produced much lower reachability numbers than expected. This might suggest that not as many undergraduates utilized the labs as anticipated, or that there is a significant impact from the power loss problems discussed in Section 4.

3.3. Capacity

In this section we examine the potential capacity of the network, focusing on the first user study. The results of the second user study will be discussed in Sections 3.5 and 4.

3.3.1. Epidemic Propagation To examine the potential capacity of the network, we developed a software suite to take our data traces as input, and simulate epidemic propagation across the traces. Thus the *simulator* output provides us with an abstract view of a potential network and its topology. The simulator assumes that all nodes have infinite amounts of memory and have infinite instantaneous bandwidth when radio contact is made.

As the simulator runs over the trace data, it keeps track of who each node last saw. Whenever a node meets another node that is different from the one it last saw, an *event* is triggered for the nodes in question. Conceptually the event simply represents a point in the trace where there is an interesting change in the network and interesting packet flow can take place.

The simulator uses these events to trigger packet generation at each of the nodes. Each participating node in the event generates a single broadcast packet. The participants then exchange packets until their queues are synchronized (i.e. they all have copies of the same packets).

While events are occasions of new direct contact, we define *multi-hop meetings* as occasions of new contact between pairs of source-destination nodes via more than one hop. During an event between a particular pair of destination and intermediary nodes, we say that a multi-hop meeting occurs between a particular source (not part of the event) and the destination if the destination receives at least one new packet from the source. Recall that in epidemic propagation, nodes never accept duplicates. Thus the new packets from the source represent a new contact occasion. Because the intermediary node can potentially be the facilitator for multiple source nodes to the destination, there can be multiple multi-hop meetings for a given event. However, there is only one multi-hop meeting per source immaterial of the number of new packets from that source.

Our primary interest in this investigation is the feasibility of forming a MANET using user mobility. In particular we focus our efforts on exploring the characteristics of multi-hop routes. However, it is quite possible that a group of nodes might be close enough together that their radios have sufficient coverage to form a connected ad-hoc network. For example, a group of three users working on a group project, shuffling amongst each other around a lab area, might trigger many simulated events and packet transfers. However, this scenario is not of interest to our study since it does not exploit real user mobility.

To remove these interactions from the simulator output, we ignore any packets delivered (from any node to any other node) with an end-to-end latency of less than two minutes. Figures 1(a) and 1(b) show the cumulative packet delivery latency for all multi-hop packets. Between the 1 and 10 minute interval, the rate of successful packet deliveries is

from	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	23
01	0	5	3	3	8	3	2	2	4	2	4	2	2	6	6	4	14	2	1	4	2
02	5	0	10	3	12	3	7	4	9	3	14	3	2	10	18	6	5	3	1	16	5
03	3	10	0	3	12	3	6	3	6	3	10	3	2	10	11	1	10	3	1	7	3
04	2	4	2	0	4	2	2	2	6	2	2	2	2	2	4	2	2	2	1	3	2
05	7	10	10	3	0	4	10	3	11	3	14	3	2	11	10	10	3	3	1	4	5
06	4	5	5	3	5	0	2	4	5	9	6	2	3	6	4	4	10	6	4	2	1
07	0	9	7	1	9	0	0	0	6	0	5	0	0	17	6	8	6	0	0	11	3
08	4	5	4	2	4	9	2	0	3	14	4	2	12	4	3	3	4	6	2	2	1
09	6	11	10	4	9	3	6	4	0	3	7	3	2	11	8	7	9	3	1	4	4
10	3	4	3	2	4	8	2	7	3	0	3	3	6	3	3	3	4	4	4	2	1
11	3	11	9	2	12	5	9	5	6	1	0	3	1	11	10	6	6	1	1	7	8
12	3	3	3	2	3	11	1	2	3	<u>16</u>	3	0	14	<u>3</u>	3	3	3	4	3	2	1
13	4	4	4	2	4	6	2	6	4	<u>7</u>	3	5	0	4	4	3	4	5	3	2	1
14	6	10	14	5	15	3	9	4	11	3	9	3	2	0	9	10	17	3	1	8	1
15	6	<u>23</u>	11	2	9	2	4	2	10	3	11	3	2	22	0	8	16	3	1	10	4
16	5	10	1	2	10	2	5	3	5	3	7	3	2	8	9	0	11	3	1	8	3
17	8	7	7	3	4	4	7	3	6	3	10	3	2	9	12	9	0	3	1	6	3
18	3	3	3	2	3	9	2	<u>25</u>	3	14	3	5	13	3	<u>3</u>	3	3	0	4	2	1
19	3	3	3	2	3	1	2	<u>2</u>	3	2	2	1	1	3	3	3	3	2	0	2	1
20	2	6	5	1	2	1	6	1	5	1	4	1	1	5	6	7	6	1	0	0	3
23	2	5	4	1	5	1	5	1	2	1	13	1	1	2	6	3	3	1	0	4	0

Table 4. User Study 1, Number of multi-hop meetings between node pairs




Figure 2. User Study 1, CDF of multi-hop packet arrivals for selected pairs of nodes

relatively flat. We chose the 2-minute cut-off as an educated guess. This threshold hopefully removes packets that are delivered without using significant user mobility, while keeping packets that are delivered quickly via user mobility.

For the remainder of this paper, unless otherwise noted, we discuss only multi-hop packets. We focus on multi-hop packets to study the feasibility of using mobility and multiple hand-off to provide effective networking. Overall, the packets delivered using a single-hop accounted for less than 18% of the total number delivered. This means that over 82% of packets were delivered faster using some multi-hop path.

3.3.2. Simulation Results Table 4 shows the number of multi-hop meetings over which packets sent from a given source were received at a given destination. Though the number of multi-hop meetings is small, recall that our network

is very sparse. We highlight six pairs of nodes (underlined) to examine in closer detail. These pairs are chosen because they have the highest multi-hop meeting activity spread evenly (roughly speaking) over the lifetime of the trace. In our very sparse network, it is understandable to have many pessimistic cases. We hypothesize that with a denser network of devices, communication opportunities will increase and provide better numbers.

We choose three pairs of nodes, $12 \rightarrow 10$, $15 \rightarrow 02$, and $18 \rightarrow 08$, that could have met directly, but have packets that are delivered in less time using user mobility and multi-hop paths. We also choose three pairs of nodes, $08 \rightarrow 14$, $12 \rightarrow 14$, and $18 \rightarrow 15$, that could only deliver packets using multi-hop paths. That is, these pairs never meet directly.

Figure 2 shows cumulative distribution frequency (CDF) plots for the six selected pairs of nodes. The packets that are plotted are the first copy of each packet to arrive.




Figure 3. User Study 1, Hop count proportion

The three node pairs in Figure 2(a) are pairs that have adjacency but are able to deliver some packets faster using a multi-hop path. In particular, between the node pair $18 \rightarrow 08$, user mobility is able to deliver over 50% of the packets within less than one hour, despite our very sparse network. Figure 2(b) shows three pairs that have no adjacency – that is, the nodes of each pair never meet directly in the traces.

Finally, we show the distribution of hop counts for these selected pairs of nodes in Figure 3. These plots show the diverse range of hop counts that exist for traversing between pairs of nodes. Figure 3(a) shows the contribution of direct communication relative to other multi-hop communication paths. This graph includes all packets, single and multi-hop, from all nodes in the trace. Figure 3(b) shows the hop count distribution for node pairs which have adjacency, but found faster communication using multi-hop. The pairs shown in Figure 3(c) have no adjacency, and therefore can only communicate via multi-hop.

3.4. First-Hand-Off Nodes

As an initial step toward finding patterns that would allow us to make better routing and replication decisions, we examine the proportion of first arrival packets delivered from a given source to destination, based on the first intermediary node the source handed the packet to. Note that a destination node may receive more than one copy of any given packet due to replication in the network. Here we only consider the first arrival of any given packet, ignoring subsequent duplicates. We call the first intermediary node the *first-hand-off node*. If a large proportion of the successful packet deliveries are done by a small number of first-hand-off nodes, then the source node might be able to achieve high rates of successful delivery while minimizing the number of replicas it sends into the network.

In Table 5, we see the proportions of packets handled by each respective first-hand-off node. The *node* column lists all of the nodes that the source met directly, which contributed in delivering a packet to the destination. The second column shows the proportion of all packets successfully delivered

from source to destination, via the respective first-hand-off node. Note that a packet might take several hops to reach the destination, but here we are focusing on the effects of the first hop.

Table 5 illustrates possible routing and replication optimizations the source node can make. For example, based on the bias shown in Table 5(a), node 15 could replicate to nodes 14 and 11, in hopes of successfully reaching node 02.

Table 6 shows a similar set of node pairs and their first-hand-off node proportions, except none of these node pairs make direct contact with one another in the traces at all. Thus in the previous table we filter out traces of direct contact, but in this table the selected node pairs never directly meet (i.e. no adjacency).

We found that most node pairs that directly meet have a strong first-hand-off bias. Most node pairs that do not directly meet have a more even distribution of proportions, as seen in Tables 6(a) and 6(d). For nodes with an even distribution of proportions, the source does not have a clear first choice for deciding which other node to prefer. However, subsequent hand-offs can have a stronger bias for the next hand-off. For example, in Table 6(a), node 08 has a 41% proportion using node 10 as a second-hand-off⁴, and node 11 has a 100% proportion using node 05 as a second-hand-off. A more sophisticated routing protocol which includes some form of feedback would allow the original source node to make better decisions.

3.5. Discussion

Due to the sparse nature of our network from the limited number of participants and potentially large ground area covered, it can be expected that most nodes will require long periods of time to communicate.

However, as network density increases, we expect the number of nodes able to find quick multi-hop paths to destination nodes will increase, and the need for high-latency paths will decrease.

⁴ Not shown in tables due to space limitation.

node	%
14	62.9
11	14.3
07	10.8
06	6.0
23	3.9
16	0.6
05	0.6
17	0.5
03	0.3
20	0.1
09	0.1

(a) 15 to 02

node	%
18	50.7
08	23.5
11	11.9
13	8.5
06	5.1
19	0.2
07	0.0

(b) 12 to 10

node	%
12	38.4
10	33.6
11	10.9
06	10.7
13	6.2
19	0.2

(c) 18 to 08

node	%
02	45.4
06	18.4
14	16.9
17	9.6
07	8.1
20	0.6
11	0.6
15	0.2

(d) 03 to 16

Table 5. User Study 1, Packet delivery proportion by first-hand-off node

node	%
08	29.1
11	27.5
12	20.6
10	16.2
06	3.4
13	3.0
19	0.2
07	0.0

(a) 18 to 15

node	%
15	54.9
10	27.8
18	9.3
12	7.6
13	0.2
06	0.1
09	0.1
07	0.1

(b) 08 to 14

node	%
08	46.5
18	16.6
11	14.5
10	13.1
13	6.4
06	2.7
19	0.1
07	0.0

(c) 12 to 14

node	%
18	21.2
13	16.6
08	15.7
14	15.0
11	15.0
12	12.8
06	3.7
07	0.1

(d) 10 to 15

Table 6. User Study 1, Packet delivery proportion by first-hand-off node (no adjacency)

As discussed in Section 3.2, the number of direct contacts that a node makes is not necessarily an indicator of its ability to reach other nodes, or be an intermediary message carrier. If the number of participants in the network increased, we expect the magnitude of adjacency for nodes to decrease relative to the total number of nodes. However, the data trends suggest that nodes will continue to maintain high reachability.

4. Experiences

In hindsight, we find that our original estimate of an 8 to 10 hour work-day is insufficient for our user-base. After the first user study with graduate students, we believed our estimate worked well. However, the second user study proved to require even more working battery life. In post-experiment interviews, we found that graduate students kept chargers at their office, and would regularly recharge the devices while at their desk. Thus most graduate students did not fully exercise the eight-hour battery life.

Most undergraduate students cannot recharge their devices mid-day. Indeed, from the onset of the second user study, a significant number of the users could not finish their work-day without draining their devices. Though we established a strict regimen of collecting data on a weekly basis from the students, they often suffered catastrophic data loss from battery exhaustion, losing several days worth of data.

Furthermore, we also found that graduate students were far more conservative with the Palm devices. Few used more than the basic features, and most only carried the devices diligently without much usage. After the first experiment, many participants mentioned that they understood the experimental nature of the software and objective, and treated the device with delicate care.

In contrast, the undergraduate students used the devices quite liberally. Within two weeks of the second user study, we found that most of the participants had downloaded significant numbers of third-party software to use on the Palm devices, including numerous games. Clearly the usage patterns of the undergraduates were significantly more demanding than we had anticipated.

4.1. Technical Limitations

The most fundamental and important implementation detail of the experiment is power management on the PalmOS platform. Though there are numerous other oddities of the platform that we work around, power management proves to be the most critical.

Previous research has shown that power consumption on mobile devices is dominated by the radio and display [3]. The fundamental problem with the PalmOS API is that it is designed to be user activity driven, and does not provide interfaces for managing specific resources used in background tasks. Furthermore the API only provides a limited set of common operations available to applications, and reserves full control over resource management. In order for our software to use the Bluetooth radio, it must wake and power the whole Palm device, including the display.

This poses a significant challenge for us when attempting to periodically use the Bluetooth radio. Our software cannot initialize and utilize the radio without first asking the PalmOS to power the rest of the device, including the display. Thus even when the device is in a user's bag, periodically using the Bluetooth radio to search for other devices also means paying for the activation of the display for the duration of the radio communication. Since the Tungsten T devices provide backlit displays, the power cost is quite significant. We searched for a method to disable the back-lighting while performing the periodic radio communication. Unfortunately, the PalmOS API only provides a method for toggling the back-lighting, but no method for querying the current setting.

Again, the API assumes the user will call the toggle through the application until the desired setting is reached.

The Palm Tungsten T devices keep a reserve of power in the batteries, in case the power levels run too low. Should this happen, the devices refuse to power on, and utilizes the remaining power to maintain memory state. At the start of every communication cycle, our software checks the battery status of the device, and disables further radio usage if the power level runs too low. Though we have the detector set to trigger well before the critical low-battery level, we still experienced numerous total power failures in the second user study. In these cases, we suspect that the reserve amount was insufficient to last the many hours between power failure and when the students finally get home to recharge.

Finally, due to extreme power limitations, we use a clock-synchronized radio protocol for data gathering. To ensure that the clocks do not drift too far, we visit each Palm device at least once per week with a NTP-synchronized laptop to re-synchronize the time.

5. Related Work

Previous work has looked into epidemic algorithms for data propagation [2] using node mobility. Mobility is used to increase the chance of nodes finding intermediary carriers and help relieve medium contention [4, 5, 7, 14, 22, 23]. Epidemic propagation places no bounds on power, storage space, or time; the data can be propagated through an environment by using mobile devices as carriers for storing and forwarding data [5]. In these networks, nodes carry data until it meets a base station where the data is offloaded for analysis. Examples of such mobile ad-hoc networks include ZebraNet [13, 18] and SWIM [22] which have been created and physically deployed in real environments, using zebras and whales for nodes, respectively. Zhao et al. [23] use mobility for data delivery in MANETS, similar to [21]. However, these systems use a tiered system for data collection. Furthermore Zhao et al. have control over the mobility of the nodes. In our work, we try to determine patterns of node mobility, without any control over their movement. Though our system could be used for data gathering, our objective is to build towards a supplemental networking platform.

Most studies of MANETs, whether for data retrieval, distribution, or networking, use simulated movement and theoretical mobility models [8, 9, 16, 17]. Our work is unique in that we use trace data involving real user contact to analyze our work.

A more realistic approach is to obtain traces from a real environment and use these traces as a model for simulation. Jetcheva et al. [11] used a fleet of city buses as mobile nodes to obtain mobility trace data. They then simulated potential latency and routing characteristics, assuming various radio coverage models, using the collected data. Our work is

unique in that we make no assumptions about radio coverage or mobility model. We study the usage of real user mobility (and only mobility) for the formation of networks. Our collection of data traces is novel in that we collect contact information of people, and we have no predetermined model of mobility.

Kotz et al. [15] provide an extensive study of large wireless network environments. Their work provides supplemental research on wireless activity and metrics that our study does not address. However, their work focuses on traces of peers (using Wi-Fi) interacting with infrastructure. Our work expands on their efforts by focusing on detecting user mobility and peer-to-peer contact patterns. Instead of studying infrastructure, we aim to study user interaction patterns, to potentially form a supplemental networking platform based on user mobility and contact.

Finally, our experiences support the need for better hinting mechanisms and OS awareness in power management for mobile devices as described in [1].

6. Conclusion

We present an experimental study to test the feasibility of using user mobility and opportunistic pair-wise contact to form an ad-hoc network. Using commodity mobile devices, we instrumented two user studies for experimentally collecting trace data of user contact. Our approach is unique in that we do not have a predetermined model of user mobility, and strive to provide a networking model based only on pair-wise contact.

The results of the experiment are promising, showing that user mobility can potentially be used to form a network. Using this trace data, we simulate an idealized network using epidemic propagation, and observe that nodes exhibit signs of regularity and affinity of contact. Furthermore, in many cases, success of message delivery from any source to destination is not evenly distributed amongst the intermediary nodes. Thus source nodes can potentially use this information for better routing decisions.

We also describe our experiences developing and deploying instrumented devices to real users. Our experiences show that power management is still an area with much room for improvement, especially for background ambient operation. Many power conservation research projects focus on energy management for devices in use. But power management methods and device peripheral design will need to take into account how devices might be used in two different modes: active use and background operation.

As future work, we plan to instrument another user study, with improved device battery longevity, and collect longer traces. We believe a more focused user group, for example nurses in a hospital or elder care facility. From the messaging characteristics, we plan to develop a routing protocol, and

compare its effectiveness and efficiency with protocols such as directed diffusion [9]. As an initial step, we will implement the first-hand-off-node heuristic to examine its impact on replication and latency. Finally, we plan to apply machine learning techniques to the contact patterns between nodes. We hope to discover regularity models in the node contact patterns, and use this information to provide better routing decisions, estimate network quality and confidence levels, and better radio power management methods.

6.1. Acknowledgments

We would like to thank our reviewers, our shepherd Yih-Chun Hu, and all the anonymous participants of the user studies.

References

- [1] M. Anand, E. B. Nightingale, and J. Flinn. Ghosts in the machine: interfaces for better power management. In *Proceedings of the 2nd international conference on Mobile systems, applications, and services*, pages 23–35. ACM Press, 2004.
- [2] N. Bailey. *The mathematical theory of infectious diseases and its applications* (second edition). Hafner Press, 1975.
- [3] K. Barr and K. Asanovic. Energy aware lossless data compression. In *Proceedings of the First International Conference on Mobile Systems, Applications, and Services (MobiSys 2003)*, San Francisco, California, May 2003.
- [4] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. *Wireless Communications and Mobile Computing: Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, 2(5):483–502, 2002.
- [5] J. A. Davis, A. H. Fagg, and B. N. Levine. Wearable computers as packet transport mechanisms in highly-partitioned ad-hoc networks. In *Proceedings of the Fifth International Symposium on Wearable Computers*, pages 141–149, Zurich, Switzerland, Oct. 2001.
- [6] F. Erina and P. Francesco. Bluetooth and ieee 802.11 wireless protocols: a survey and a comparison. CMPSCI Tech. Rep. 2003-TR-36, ISTI - CNR (Institute of the National Research Council), Moruzzi 1, 56124 Pisa, Italy, 2003.
- [7] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In *Proceedings of INFOCOM*, pages 1360–1369, Anchorage, Alaska, Apr. 2001.
- [8] K. M. Hanna, B. N. Levine, and R. Manmatha. Mobile distributed information retrieval for highly-partitioned networks. In *Proceedings of 11th IEEE International Conference on Network Protocols*, pages 38–49, Atlanta, Georgia, Nov. 2003.
- [9] C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva. Directed diffusion for wireless sensor networking. *IEEE/ACM Transactions on Networking*, 11(1), Feb. 2003.
- [10] A. Jardosh, E. M. Belding-Royer, K. C. Almeroth, and S. Suri. Towards realistic mobility models for mobile ad hoc networks. In *MobiCom, San Diego, CA, 2003*.
- [11] J. G. Jetcheva, Y.-C. Hu, S. PalChaudhuri, A. K. Saha, and D. B. Johnson. Design and evaluation of a metropolitan area multitier wireless ad hoc network architecture. In *Proceedings of the Fifth IEEE Workshop on Mobile Computing Systems and Applications (WMCSA 2003)*, pages 32–43, Pittsburgh, Pennsylvania, Oct. 2003.
- [12] C. E. Jones, K. M. Sivalingam, P. Agrawal, and J. C. Chen. A survey of energy efficient network protocols for wireless networks. *Wirel. Netw.*, 7(4):343–358, 2001.
- [13] P. Juang, H. Oki, Y. Wang, M. Martonosi, L.-S. Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebrant. In *Proceedings of the Tenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)*, San Jose, California, Oct. 2002.
- [14] A. Kansal, A. A. Somasundara, D. D. Jea, M. B. Srivastava, and D. Estrin. Intelligent fluid infrastructure for embedded networks. In *Proceedings of ACM MobiSys 2004*, pages 111–123, Boston, Massachusetts, June 2004.
- [15] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. In *Proceedings of the 8th annual international conference on Mobile computing and networking*, pages 107–118. ACM Press, 2002.
- [16] A. Lagar Cavilla, G. S. Baron, T. E. Hart, L. Litty, and E. de Lara. Simplified simulation models for indoor manet evaluation are not robust. In *Proceedings of the First IEEE Conference on Sensor and Ad-Hoc Communications and Networks*, Oct. 2004.
- [17] Z. Li, B. Li, D. Xu, and X. Zhou. ifbw: Middleware-assisted rendezvous-based information access for mobile ad hoc applications. In *Proceedings of the First International Conference on Mobile Systems, Applications and Services (MobiSys 2003)*, pages 71–84, San Francisco, CA, May 2003.
- [18] T. Liu, C. M. Sadler, P. Zhang, and M. Martonosi. Implementing software on resource-constrained mobile sensors: Experiences with impala and zebrant. In *Proceedings of the Second International Conference on Mobile Systems, Applications and Services (MobiSys 2004)*, pages 256–269, Boston, Massachusetts, June 2004.
- [19] Pilrc: Pilot resource compiler. Sourceforge.net, <http://sourceforge.net/projects/pilrc>, 2004.
- [20] gcc prc-tools. Sourceforge.net, <http://prc-tools.sourceforge.net>, 2003.
- [21] R. C. Shah, S. Roy, S. Jain, and W. Brunette. Datamules: Modelling a three tiered architecture for sparse sensor networks. In *First IEEE International Workshop on Sensor Network Protocols and Applications (SNPA)*, May 2003.
- [22] T. Small and Z. J. Haas. The shared wireless infostation model - a new ad hoc networking paradigm (or where there is a whale, there is a way). In *Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 233–244, Annapolis, Maryland, 2003.
- [23] 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*, pages 187–198. ACM Press, 2004.