Measurements from an 802.11b Mobile Ad Hoc Network

Vincent Lenders, Jörg Wagner and Martin May
Computer Engineering and Networks Laboratory
ETH Zurich
8092 Zurich, Switzerland
{lenders, may}@tik.ee.ethz.ch, jwagner@ee.ethz.ch

Abstract

This paper analyzes the characteristics of a multi-hop 802.11b mobile ad hoc network. We present data gathered from a mobile network of 20 devices carried by test users over 5 days in an indoor environment. The data is analyzed with regard to (i) the number of reachable devices, (ii) the node degree, (iii) the average path length, (iv) the link lifetime, (v) and the route lifetime. Despite the relatively high node density and low node mobility in our setup, we observe frequent network partitioning and considerably high path lengths (as large as 7 hops). However, the usability of these long paths is questionable as their lifetime is short. We believe that our measurements are representative for typical indoor environments and that the results can and should be used for evaluating networking protocols as well as to validate existing or to derive new mobility models.

1. Introduction

Mobile ad hoc networks consist of autonomous devices communicating over radio without any support from a fixed infrastructure. Common scenarios for ad hoc networks include dynamic communication networks for emergency and rescue operations, disaster relief efforts, and similar tasks, where typically no communication infrastructure is present prior to the deployment of the ad hoc network.

Our visionary application of an ad hoc network is to network human-carried devices without any fixed infrastructure. Such a network would allow users to exchange messages, share files, or content of any type while moving around.

The open research question we address in this paper is how a relatively dense IEEE 802.11 network of human-carried devices would appear. In particular, we are interested in the impact of human mobility on the variation of the network topology.

We tackle this question in this paper by observing the network topology over time from a twenty-nodes ad hoc network we deployed in an office environment. The network consists of commodity Personal Digital Assistants (PDAs) equipped with IEEE 802.11b wireless interfaces which are turned into ad hoc mode. Each PDA is carried by a test user over five consecutive days. All users are working on the same floor at a research lab located in the ETH Zurich campus.

On each PDA, a monitoring tool determines every second the devices that are in direct transmission range, and stores this information on a FLASH card. From these periodic messages, we then reconstruct snapshots of the topology for various moments in time. In particular, we analyze (i) the number of devices each node can reach, (ii) the node degree, (iii) the average path length between node pairs, (iv) the link lifetime, as well as (v) the route lifetime. The first three metrics (i)-(iii) provide useful insights to understand network partitioning and the network size. The link and route lifetime reveal the network dynamics in our experiments, and specifically the impact of human mobility.

The rest of the paper is organized as follows. In the next section, we survey related work. In Section 3, we describe the experimental setup. In Section 4, we analyze the network characteristics and we conclude in Section 5.

2. Related Work

Many studies of multi-hop ad hoc networks have been conducted with simulations including synthetic mobility models. However, it is unclear if the networks in simulations are realistic. In the following, we review related works analyzing networks with real user mobility.

APE [7] was proposed as a testbed for reproducible mobility measurements. The testbed was used to capture the effect of mobility on the network characteristics with up to 37 nodes. Their measurements mainly differentiate from ours in the way the test users move. In their scenarios, people move with strict choreographic instructions for small time periods of maximum 400 seconds, whereas in our measurements, test users move without instructions and the measurements last for multiple days. Furthermore, they use the signal strengths between the nodes to capture the effect of mobility whereas in our analysis, we use the number of lost packets between two nodes.

Measurements with real user mobility were also conducted with PDAs [10] and Intel iMotes [4] using Bluetooth. In [10], twenty PDAs were distributed to students at a campus. In [4], the authors distributed iMotes to conference attenders. In both experiments, each device periodically scanned the neighborhood with Bluetooth to detect devices in direct proximity. Compared to our experiment, the users were much more widespread and multi-hop end-to-end connectivity between nodes was less frequent. Therefore, the authors analyze their traces for delay-tolerant or opportunistic networking scenarios. In a delay-tolerant network, the time until a data chunk is delivered to a destination is the dominant metric to look at, compared to the duration of an end-to-end path as examined in this work. Furthermore, due to the limitations of the current Bluetooth implementation, it is not possible to exchange period beacons or to scan the neighborhood at each node with a frequency of 1 s^{-1} as we do in our experiment.

User mobility was studied based on 802.11 WLAN access points associations of users with their laptops at a campus in [6, 11]. However, these traces reflect a different kind of mobility compared to the traces used in our work. Users with laptops are typically not moving while they are associated with an access point; generally, they move only when the laptops are switched off. The result from these traces characterize the users' association behavior with access points and cannot be easily adopted to study the network characteristics of a mobile ad hoc network. In [9], PDAs with 802.11 WLAN interfaces were used instead of laptops. The authors observe a higher degree of mobility than with laptops. However, since the PDAs are used in infrastructure mode (we operate the PDAs in ad hoc mode in our experiment), the effect of mobile users outside the range of an access point is not captured.

A testbed with mobile cars was setup at CMU [8]. This testbed is different from ours. A smaller number of nodes are used, the measurements are outside, and cars have different motion characteristics than people.

3. Testbed with Real Human Mobility

In this section, we present the testbed we use for the analysis and how we determine the topology from the packet traces.

3.1. Experimental Setup

For our measurements, we used a network of 20 identical HP iPAQs hx2400 running Microsoft Windows Mobile Edition 2003. The iPAQs were connected using their integrated IEEE 802.11b WLAN module (a Samsung SWL-2750C chipset operating at 2.4 GHz) turned into ad hoc mode.

During summer 2005, we conducted several experiments where the described iPAQs were distributed to twenty test users during five consecutive working days (from 10am to 5pm when all users were working). The test users were researchers, staff members, and students at the Swiss Federal Institute of Technology (ETH) in Zurich, all working on the same floor (the floor plan and the working place of each test user is shown in Figure 1). The test users were instructed to carry the PDA with them throughout the day and to recharge the battery whenever necessary (the autonomy of the battery unfortunately was less than a working day, approximately 5 hours, when WLAN was turned on). A majority of the test users were researchers and spent most of the time at their desks. The users became mobile mainly due to lunch and coffee breaks, for going to the rest room, picking up printouts in the hallway, or meeting each other for discussions. Some test users occasionally left the building or even the campus for some time during the experiment.

Since the experiment took place in a working environment, other devices such as WLAN access points and laptops operating at close channels are possible causes of interference. To mitigate the effect of interference with our testbed, we disabled all access points on the floor operating at neighboring channels for the period of the experiments.

3.2. Network Connectivity Model

In our experiment, the users did not use the PDAs to actually communicate with each other. Instead, we ran an autonomous monitoring tool on each iPAQ to detect other devices in transmission range. The tool periodically sent IP broadcast packets. Note that 802.11b broadcast packets are not acknowledged at the MAC Layer and transmitted at a fixed bit rate of 1 Mbit/s. The devices in direct transmission range which received a broadcast packet, stored the arrival time of the packet, the identity (the IP address) of the sender, and the packet's sequence number on an external Compact Flash card. The devices were synchronized every morning before the measurement period. Due to the short observation periods (1 day monitoring before resynchronization) and the "long" inter-packet delay (1 second), we did not encounter any inaccuracies from clock drifts.

We determine connectivity between nodes based on the history of received/lost packets. That is, we construct a *con-*

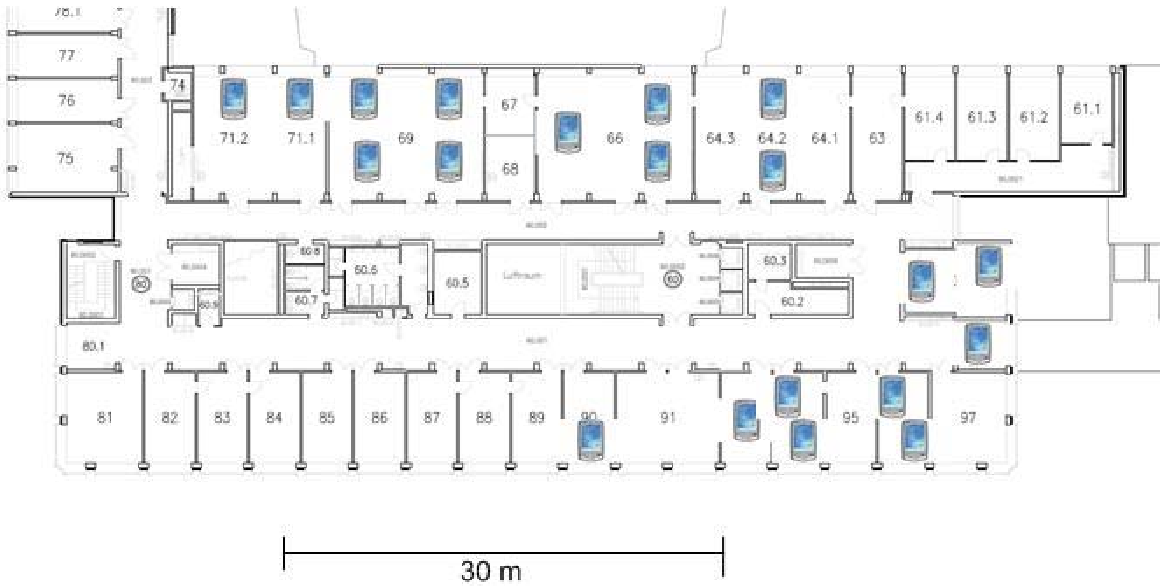

Figure 1. Floor plan of testbed location. The working place of each test user is marked with a PDA.

nectivity graph for each broadcast time interval. In a connectivity graph, there exists a link from node a to node b if b receives enough packets from a within a time window. Note that we calculate the amount of received packets for all individual node pairs in both directions and thus, unidirectional links might occur. For the analysis in this paper, we define a link when at least 50% of the packets are received within the time window.

We also thought of using the received signal strength indication (RSSI) from the MAC layer to measure the link quality and assign links in the connectivity graph. Unfortunately, the RSSI information on the used devices was not reliable enough when the wireless module was operated in ad hoc mode.

A problem we observed was BSSID partitioning (also observed in [2]). This occurs after the convergence of network partitions. There, even if all devices use the same SSID and the same channel, the devices do not see each other because of different BSSIDs. Eventually, all devices will agree on the same BSSID, however this convergence might turn out to be slow. As we did not have access to the network driver, which would have been necessary to solve the problem at its root, we implemented a tool which reset the wireless interface of a device if multiple partitions coexisted for longer time periods. This workaround helped us to mitigate the long convergence time, but the problem remained and could be a small source of bias in the analysis.

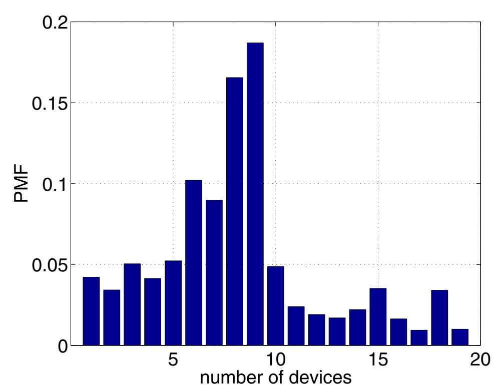

Figure 2. Reachable Devices

4. Analysis

In this section, we analyze the traces from our measurements. The numbers we present next are extracted from an experiment with a duration of five consecutive days. The settings used in this experiment were: 1 second for the broadcast time interval at the devices; 7 seconds for the time window size to reconstruct the connectivity graphs. In other experiments conducted at our lab, we used different parameter settings, however, the obtained results were very similar to those presented in this paper.

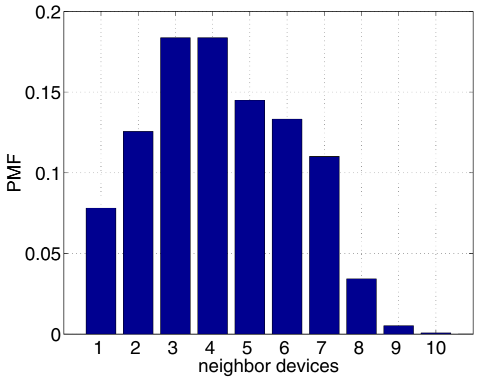

Figure 3. Node Degree

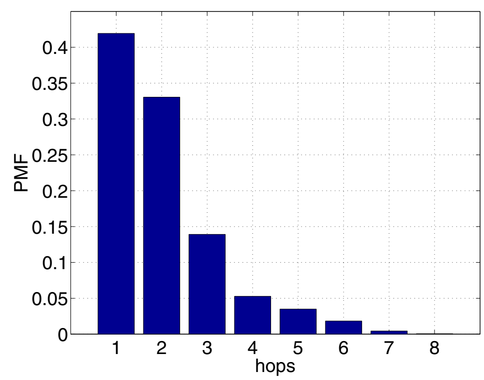

Figure 4. Shortest Path Length

4.1. Reachable Devices

The first metric we look at is the number of devices a node can reach in the network. In Figure 2, we plot the probability mass function (PMF) of this metric averaged over all devices. The values represent the number of devices reachable directly or via multiple hops from the connectivity graphs. If a node cannot reach another node, the network must be partitioned. The distribution has a peak at 9 nodes. Thus, the probability is high that each node can reach half of the nodes. Notice that at certain moments, all 20 iPAQs were connected into one single network.

Against our expectations, the relatively high device density was not sufficient to establish a fully connected network all the time. The wireless signal attenuation due to walls and objects in the rooms was too high and caused a considerably high number of partitions.

4.2. Node Degree

Another metric we consider is the node degree. We define the node degree as the average number of devices that are direct neighbors in the connectivity graph. The probability mass function of this metric averaged over all devices is plotted in Figure 3. The peak of the distribution is at four neighbors; the mean is at 4.25. Since the distribution is close to zero for 10 neighbors, we conclude that at no point in time during the measurement period, the test users were all in transmission range. At most, there were 9 nodes in direct transmission range.

4.3. Path Length

We next analyze the path lengths from the connectivity graphs. For this purpose, we calculated the shortest paths from any two node pairs over time. The PMF is plotted in Figure 4. As expected, most nodes are close to each other and, thus have a distance of one hop (with a probability of approximately 0.43). Surprisingly for a network size of 20 nodes which were relatively close to each other, there were some quite long shortest-paths consisting of 5, 6, and even 7 hops. Notice however, that these long paths are routes in the connectivity graphs which were obtained from the periodic beacon messages. There is no guarantee that data packets were actually sent over those routes.

4.4. Link Residual Lifetime

The previous three metrics were not directly indicating the degree of dynamics in the network. To capture this effect, we first look at the lifetime of links in the network. There are two statistical view points to look at link lifetimes: (i) the *total lifetime* of a link describes the time interval between the moment the link appeared until it breaks; and (ii) the *residual lifetime* represents the time interval between a sample moment after the creation until the link breaks. Generally, it is not important whether the total or the residual lifetime is used since the distribution of the total lifetime can be converted into the distribution for the residual lifetime, and vice versa, with the law of total probability. From an application or user perspective, it is more interesting to look at the residual lifetime since communication starts at arbitrary moments and not necessarily when a new link becomes available. In the following, we therefore analyze the residual lifetime.

In Figure 5, we plot the probability mass function and the corresponding cumulative distribution function (CDF) for the link residual lifetime. We observe that many links break after a small amount of time. After 100 seconds, 20% of the links are unavailable. After 500 seconds, already 55% of the links are not available anymore. Interestingly, for long time intervals, we find that a significant amount of links are still available. After 3500 seconds, approximately 3% of the links are still available.

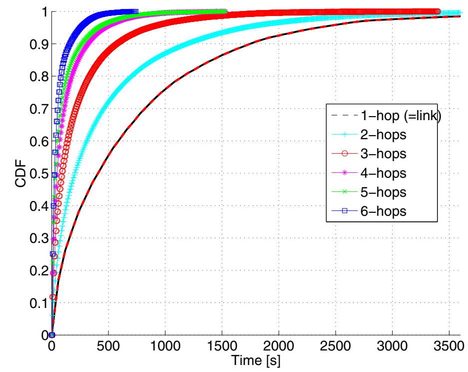

Figure 6. CDFs of Route Residual Lifetime for Different Route Lengths

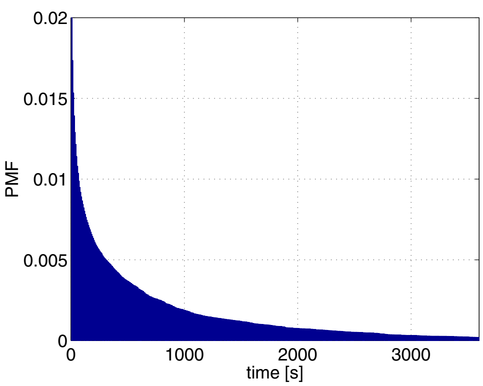

(a) PMF

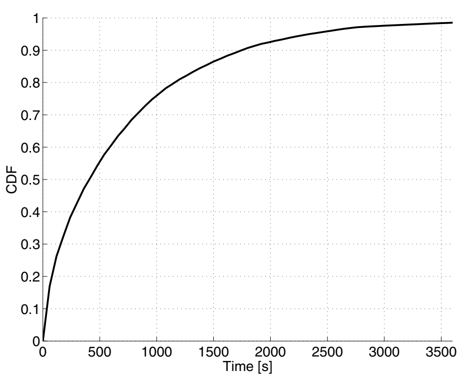

(b) CDF

Figure 5. Link Residual Lifetime

We explain this trend as follows. The links that break quickly are mostly unstable links which break due to transmission errors caused by fading, interference, collisions, or noise. Links with a lifetime larger than a few hundred seconds, break either due to transmission errors or user mobility. However, links which have been up for longer periods are very stable and break most frequently due to user mobility. Some links thus had a lifetime of more than 3500 seconds, i.e., the users did not move within this period, for example, when the device had to be recharged.

4.5. Route Residual Lifetime

The CDF of the residual route lifetime is plotted in Figure 6 individually for routes of different lengths. The distributions are obtained by counting the remaining lifetime of the shortest route between all node pairs form the network connectivity graphs. Note that since the nodes are mobile, it is possible that, while monitoring the lifetime of a route between two nodes, a shorter alternative route between these nodes becomes available. However, we always count the remaining lifetime of the initially computed shortest route. The dashed curve in the plot is the CDF for 1-hop routes (i.e., a link) and is identical to the link lifetime CDF in Figure 5. For larger routes, the expected lifetime significantly decreases. For 2-hops routes, the probability that a route lasts longer than 3500 seconds is almost equal to zero. And for 6-hops routes, the probability is almost zero after 500 seconds. Furthermore, 90 % of the 3-hops routes are already broken after approximately 500 seconds.

While in our experimental setup the users were moving only occasionally, the lifetime of routes happened to be short. We conclude that in an office environment as the one used in this experiment, multi-hop communication is a challenge for applications that require persistent end-to-end connections for longer time periods.

5. Conclusions

We studied the network characteristics of an 802.11b ad hoc network with 20 mobile users in a typical office environment. The main findings are:

- We observed that, despite the relatively high node density in our experiment, network partitioning occurred very frequently.
- The link lifetime distribution shows that many unstable links break after a short link duration, but few very stable links result in long lifetimes.
- We further observed quite large path lengths of up to 6 and even 7 hops in the network. However, the lifetime of such long paths is very short. 90 % of the 6-hops paths are broken after 200 seconds.

We believe that our measurements are representative for typical indoor or office environments. Our results could and should therefore be used to evaluate networking protocols. For example, the presented link lifetime distribution could be used to adapt existing mobility models, or, if necessary, to derive different new mobility models. As future work, we plan to compare our empirical results with frequently used mobility models as for example the random waypoint [5] or the random reference point group mobility model [3] to verify if they result in similar link and route lifetime distributions.

6 Acknowledgements

We would like to thank all the participants of the experiments who carried the PDAs. Vincent Lenders is partly funded by the Swiss National Science Foundation (SNF) under grant 200021-103578.

References

[1] IETF Mobile Ad-hoc Networking (manet) Working Group. http://www.ietf.org/html.charters/manet-charter.html, 2004.

[2] D. Aguayo, J. Bricket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *Proceedings of ACM SIGCOMM'04*, Portland, Oregon, USA, October 2004.

[3] X. Hong, M. Gerla, G. Pei, and C.-C. Chiang. A Group Mobility Model for Ad Hoc Wireless Networks. In *Proceedings of ACM/IEEE MSWiM'99*, Seattle, WA, August 1999.

[4] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket Switched Networks and Human Mobility in Conference Environments. In *Proceedings of the ACM SIGCOMM 2005 Workshop on Delay-Tolerant Networking (WDTN)*, Philadelphia, Pennsylvania, USA, August 2005.

[5] D. B. Johnson, D. A. Maltz, Y.-C. Hu, and J. G. Jetcheva. The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR). February 2002.

[6] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. In *Proceedings of the Conference on Mobile Computing and Networking (MOBICOM)*, Altanta, Georgia, USA, September 2002.

[7] H. Lundgren, D. Lundberg, J. Nielsen, E. Nordström, and C. Tschudin. A Large-scale Testbed for Reproducible Ad hoc Protocol Evaluations. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, Orlando, Florida, USA, March 2002.

[8] D. A. Maltz, J. Broch, and D. B. Johnson. Quantitative Lessons From a Full-Scale Multi-Hop Wireless Ad Hoc Network Testbed. In *Proceedings of the IEEE Wireless Communications and Networking Conference*, Chicago, USA, September 2000.

[9] M. McNett and G. M. Voelker. Access and Mobility of Wireless PDA Users. In *ACM SIGMOBILE Mobile Computing and Communications Review*, April 2005.

[10] J. Su, A. Chin, A. Popivanova, A. Goel, and E. de Lara. User Mobility for Opportunistic Ad Hoc Networking. In *Proceedings of the EEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, English Lake District, UK, December 2004.

[11] C. Tuduce and T. Gross. A Mobility Model based on WLAN Traces and its Validation. In *Proceedings of INFOCOM*, Miami, Florida, USA, March 2005.