

Analyzing the Impact of Mobility in Ad Hoc Networks

Vincent Lenders, Jörg Wagner, and Martin May
Computer Engineering and Networks Laboratory
ETH Zurich

8092 Zurich, Switzerland

lenders@tik.ee.ethz.ch, jwagner@ee.ethz.ch, may@tik.ee.ethz.ch

ABSTRACT

The impact of mobility on the link and route lifetimes in ad hoc networks is of major importance for the design of efficient MAC and network layer protocols. Up to now, no real-life measurements were used to study the effect of node mobility on link and route lifetime distributions. In this paper, we present data gathered from a real network of 20 test users and analyze it with regard to link and route lifetime distributions. Besides link breakage due to node mobility, links might also break due to diverse sources of interference or packet collisions. We develop a statistical framework to distinguish between the mobility and interference or collision errors. With this framework, we are able to determine and analyze the lifetime distributions for both error types separately. We use this framework together with our measurements to validate two commonly used stochastic mobility models including the random waypoint and the random reference group mobility model. The results show that the distributions of the two stochastic mobility models match very closely the empirical link lifetime distribution.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols; C.4 [Performance of Systems]

General Terms

Design, Experimentation, Measurement

Keywords

MANET, Human Mobility, PDA, Link Lifetime, Route Lifetime

1. INTRODUCTION

The availability of small-size, portable devices and inexpensive wireless communication technology has pushed the development of mobile and pervasive computing. Specifically, Mobile Ad hoc NETWORKS (MANET)[1] have received a lot of attention in the research community. As a result, multiple protocols and applications for MANETs have been developed. Since not so many MANETs

have been deployed yet, most of the research in this area is based on simulations. Among the multitude of simulation parameters (traffic and communication pattern, mobility model, propagation model, etc.), the mobility model has a major impact on the link and route lifetime distribution, and therefore also on the protocol and application performance. Hence, for an efficient design of MAC and network protocols, one has to understand the impact of the node mobility on the link and route lifetimes.

Until now, only very few evaluations have been published that are based on measured data of real human mobility and behavior. One obvious reason is the lack of deployed MANETs available to conduct experiments. A further reason is the inherent difficulty to build experimental testbeds with human mobility as real test users need to be involved and coordinated in the experiments.

In this paper, we study and analyze the impact of human mobility on the link and route lifetime distributions in a real MANET. Throughout summer 2005, we repeatedly conducted experiments in a typical office environment. We distributed twenty commodity Personal Digital Assistants (PDAs) with IEEE 802.11b wireless interfaces to different test users. The involved test users were researchers and students working on the same floor in a building and are thus in wireless transmission range most of the time. For data capturing, each PDA runs a set of monitoring tools to determine and trace connectivity to other PDAs.

We use the collected data from this test network to analyze the impact of human mobility on the link and route lifetimes. However, the lifetime of links and routes is not only determined by user mobility but also by different sources of failures like packet collisions, or interference from any sender emitting at the same frequency band. In the absence of motion or location detection sensors, it is not directly possible to distinguish with IEEE 802.11b hardware between link failures due to human mobility or due to a collision/interference failure. To solve this problem, we develop a statistical framework which allows for analyzing the link and route lifetimes while distinguishing between the two causes of failures.

We find that failures due to human mobility and interference or collision errors result in very different link and route lifetime distributions. If we presume that there occur only interruptions due to mobility, 90% of all links would still be available after five minutes. In a scenario where interruptions occur only due to interference/collision failures, only 67% would remain after five minutes. With a long-term perspective, we observe a different trend. After 55 minutes, only 3% of the links would still be available with mobility interruptions only, whereas 35% would remain with collision or interference interruptions only.

Finally, we compare two commonly used mobility models with our measurements and our framework. In particular, we use the random waypoint [10] and the random reference point group mo-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

REALMAN'06, May 26, 2006, Florence, Italy.

Copyright 2006 ACM 1-59593-360-3/06/0005 ...\$5.00.

bility model [7]. Comparing with these two stochastic mobility models, we obtain a very similar link lifetime distribution as those measured in our testbed.

To summarize, the contributions of this paper are the following:

- We present the link and route lifetime distributions of an 802.11b ad hoc network consisting of 20 devices carried by human beings.
- We present a statistical framework to analyze link and route lifetime distributions with mobility. We use this framework to distinguish between collision or interference errors and errors due to mobility, and determine the distribution of each error type.
- We compare the empirical link lifetime with the lifetime determined from simulations with the random waypoint mobility model and the random reference point group mobility model, and show that the simulated distributions are very close to the empirical distribution.

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 and show the overall link and route lifetime distributions. In Section 4, we present our statistical framework for the route and link lifetimes and use it on the empirical data in Section 5. We compare the empirical lifetimes with simulations in Section 6 and conclude in Section 7.

2. RELATED WORK

In this section, we discuss related work. We distinguish between approaches that use synthetic mobility models and approaches which use network traces to study the impact of mobility in wireless networks.

2.1 Mobility Models

Numerous stochastic mobility models, such as the random waypoint or the random reference group mobility model, have been proposed for evaluating protocols in ad hoc networks. A survey of the most frequently used stochastic models is presented in [4]. Previous works on the impact of mobility on the link and route lifetime have been conducted mainly with these models. One of the earliest analysis was done by Mc Donald and Znati in [13]. They used a slightly modified random waypoint mobility model and derived expressions for the probability of link and path availability. In [6], link lifetimes were studied from simulations with the random waypoint mobility model, the Gauss-Markov scenario, and the Manhattan grid scenario. In [16], the link and path duration is studied for four different mobility models including the random waypoint, the reference point group mobility, the freeway, and the Manhattan mobility model. These studies are helpful to understand how the lifetime of links and routes are affected by the model dependant mobility patterns. However, they do not discuss or validate whether the used stochastic mobility patterns are realistic.

More recently, more complex mobility models have been proposed which do no longer model node mobility as a random process. For example, the predetermined time tables of public buses in a city was used by Jetcheva et al. in [9] to predict the location of each bus at specific points in time. But, such a model can only be applied to mobile entities which have a deterministic motion path. Another approach proposed in [15] is based on strategies. The motion of each node depends on a strategy that the node uses to achieve a given goal. While being more realistic than stochastic

motion, it remains unclear how close this model reflects real life behavior. An interesting approach was proposed in [5]. The authors use a multi-player game engine (Quake II) to study player movement in a virtual world. However, it has not been shown how close the movement in the virtual space corresponds to typical movement in a real world.

2.2 Real User Mobility

User mobility was studied based on 802.11 WLAN access point associations of users with their laptops at a campus in [11, 18]. 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 user's association behavior with access points and cannot be easily adopted to study the lifetime of links and paths in a mobile ad hoc network. In [14], 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.

Measurements with real user mobility were conducted with PDAs [17] and Intel iMotes [8] using Bluetooth. In [17], 20 PDAs were distributed to students at a campus. In [8], the authors distributed iMotes to conference attenders. In both experiments, each device periodically scanned the neighborhood with Bluetooth to detect devices in direct proximity. These experiments resemble more our measurement scenario compared to the setup used with access points. However, in the two mentioned experiments, 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 a frequency of 1 s^{-1} as we do in our experiment.

To the best of our knowledge, our work is the first to study the effect of user mobility on the link and path duration with real network traces from an ad hoc network.

3. TESTBED WITH REAL HUMAN MOBILITY

In this section, we first describe the experimental setup. Then, we explain how we reconstructed the connectivity graphs from our data traces, and finally, we plot the overall link and route lifetime distributions.

Note that a preliminary study of the testbed was described in [12]. Please refer to this paper for general information on the test network such as for example the average node degree, the average path length, or number of reachable devices averaged over time.

3.1 Experimental Setup

We established a testbed network at ETH Zurich. The network consists of 20 identical HP iPAQs hx2400 running Microsoft Windows Mobile Edition 2003. The iPAQs are connected using their integrated IEEE 802.11b WLAN module (a Samsung SWL-2750C chipset) turned into ad hoc mode.

During summer 2005, we conducted several experiments where the described iPAQs were distributed to 20 test users during five

consecutive working days (from 10am to 5pm). The test users were researchers, staff members, and students of a networking research lab, all working on the same floor (see [12] for a detailed map of the setup). 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 five 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 also left the building or even the campus for a certain amount of time during the experiment.

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 sequence number of the packet on an external Compact Flash card. The devices were synchronized every morning before the measurement period. Due to the short observation periods (one day monitoring before re-synchronization) and the "long" inter-packet delay (one 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 *connectivity 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.

There are several reasons that a node might not receive packets from another node. Two nodes might possibly be too far apart to be able to receive packets from each other. But also simultaneous transmissions from a third node might cause packet collisions. In fact, we observed in our testbed that when many users were in direct transmission range, a significant amount of packets was lost due to collisions. Furthermore, collocated devices emitting at the same or neighboring frequency bands might cause interference. Since our experiments took place in a lab environment, all sorts of other WLAN and Bluetooth devices not part of the testbed were on during the experiments.

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 co-existed

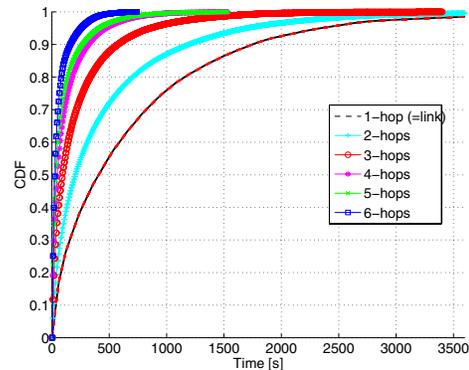


Figure 1: CDFs of Empirical Overall Link and Route Residual Lifetime

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.

3.3 Overall Link and Route Residual Lifetimes

There are two statistical view points to look at link and route lifetimes: (i) the *total lifetime* of a link or route describes the time interval between the moment the link (or route) appeared until it breaks; and (ii) the *residual lifetime* represents the time interval between a sample moment after the creation until the link or path 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 of the residual lifetime, and vice versa. 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 route becomes available. In the remaining part of the paper, without stating it explicitly, we always consider the residual lifetime.

We next analyze the overall residual lifetime of the links and routes in the test network. The overall residual link lifetime is the duration of a link between two nodes. The overall residual lifetime is affected by both, errors due to mobility and errors due to collisions or interference. In Figure 1, we plot the cumulative distribution function (CDF) for the links (denoted with 1-hop). 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 is still available. For example after 3500 seconds, approximately 3% of the links are still available.

The CDF for the overall residual route lifetime is also plotted in Figure 1 individually for routes of different lengths. The distributions are obtained by counting the remaining lifetime of the shortest route between all node pairs in the network connectivity graph. Note that since the nodes are mobile, it is possible that, while monitoring the lifetime of a route, a shorter alternative route becomes available. However, we always count the remaining lifetime of the initially computed shortest route. For larger routes, the expected lifetime significantly decreases. For two-hops routes, the probability that a route lasts longer than 3500 seconds is almost equal to zero. And for two-hops routes, the probability is almost zero after 500 seconds.

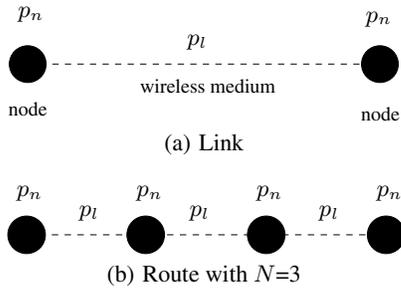


Figure 2: Failure Model

4. GENERAL FRAMEWORK FOR RESIDUAL LIFETIMES

In this section, we present our statistical framework to model the lifetime of links and routes. Our main goal is to determine the impact of mobility. For this purpose, we separate the possible failure reasons into two classes: failures which are due to node mobility and failures which are independent of node mobility like collisions or interference. With these two failure probabilities, we first derive the lifetime CDFs for the case where all nodes are mobile, and then for the case where some nodes are mobile and some nodes are fixed. The framework is then used in the next section to analyze the data from our experiment.

4.1 Basic Framework with all Mobile Nodes

In our basic model, all users are mobile and behaving in the same way (the case including nodes behaving differently is handled in the next subsection). Furthermore, all devices are reliable and functioning over time. Our goal is to determine the reason for link breaks. Therefore, we differentiate between link transmission failures caused by movement of one of the linked nodes, and other failures caused by interference from other devices. We define the corresponding failure probabilities as follows:

- $p_n(t) = P[T < t]$: The probability that the link between two nodes breaks due to movement of one of the two nodes.
- $p_l(t) = P[T < t]$: The probability that the link between two nodes breaks due to disturbances from other devices. They might be packet collisions from other nodes competing for the wireless channel, or general interference from other sources sending at the same frequency. In the rest of this paper, we always refer to interference to describe this type of failures.

With these two independent failure probabilities, we are now able to formulate the overall lifetime probability of a link. A link consists of two nodes connected over a wireless transmission medium (see Figure 2(a)). For a link to last longer than time interval t , none of the nodes may move away during that time and no interference may occur. Therefore, the probability that a link lasts longer than time t is

$$P[T_{\text{link}} > t] = (1 - p_l(t))(1 - p_n(t))^2, \quad (1)$$

and the overall link lifetime CDF is

$$\begin{aligned} P[T_{\text{link}} < t] &= 1 - P[T_{\text{link}} > t] \\ &= 1 - (1 - p_l(t))(1 - p_n(t))^2. \end{aligned} \quad (2)$$

The link lifetime in the absence of interference is then obtained by setting $p_l(t) = 0$

$$P[T_{\text{link}} < t | p_l(t) = 0] = 1 - (1 - p_n(t))^2 = 2p_n(t) - p_n^2(t), \quad (3)$$

and for the link lifetime CDF in the absence of node mobility, we set $p_n(t) = 0$

$$P[T_{\text{link}} < t | p_n(t) = 0] = p_l(t). \quad (4)$$

The same reasoning leads us to the conditional lifetime CDF of routes, meaning the lifetime CDF of routes in the absence of mobility or interference failures. A route consists of a set of nodes connected with links (see Figure 2(b)). If N is the number of links in a route, $N + 1$ is the number of nodes which are part of the route (N is 3 in the figure). For a route to last longer than the time interval t , none of the $N + 1$ nodes may move away and none of the N wireless connections between the nodes may be disturbed with interference during that time. Thus, the probability that a route lasts longer than time t is

$$P[T_{\text{route}} > t] = (1 - p_l(t))^N (1 - p_n(t))^{N+1}, \quad (5)$$

and the route lifetime CDF is

$$P[T_{\text{route}} < t] = 1 - (1 - p_l(t))^N (1 - p_n(t))^{N+1}. \quad (6)$$

The conditional route lifetime CDF in the absence of interference is then

$$P[T_{\text{route}} < t | p_l(t) = 0] = 1 - (1 - p_n(t))^{N+1}, \quad (7)$$

and the conditional route lifetime CDF in the absence of node mobility

$$P[T_{\text{route}} < t | p_n(t) = 0] = 1 - (1 - p_l(t))^N. \quad (8)$$

4.2 Extension with Fixed Nodes

The model introduced above is based on the assumption that all the nodes behave in the same way (that all nodes are mobile). However, if some nodes are static, this assumption is no longer true and the derived lifetimes are no longer correct. In the following, we relax from this strict assumption and extend our framework for cases where nodes behave differently, namely part of the nodes are mobile, and the remaining nodes are static.

For space reasons, we only derive the overall route lifetime, and refer to the technical report version of the paper for the conditional route lifetimes and link lifetime distributions. These distributions can easily be derived in the same manner as shown in the previous subsection.

Let M be the number of mobile nodes out of $N + 1$ nodes in a route. Then, according to the law of total probability, the overall probability that a route does not break during the time interval t is

$$P[T_{\text{route}} < t] = \sum_{m=0}^{N+1} P[T_{\text{route}} < t | M = m] \cdot P[M = m], \quad (9)$$

where $P[M = m]$ is the probability that a route contains m mobile nodes. According to Equation (6), the conditional route lifetime probability is given by

$$P[T_{\text{route}} < t | M = m] = 1 - (1 - p_l(t))^N (1 - p_n(t))^m. \quad (10)$$

To obtain the probability that a particular route contains m mobile nodes, we assume that the nodes are spatially distributed in such a way that the probability of the next hop in a route to be fixed or mobile is equal to the number of remaining fixed/mobile nodes

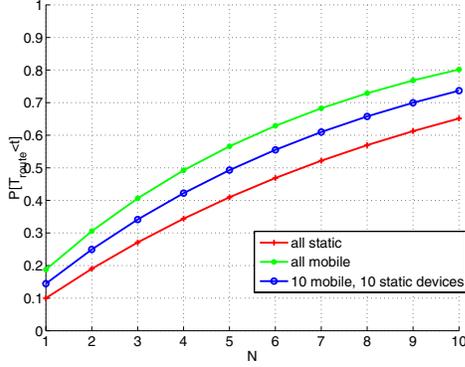


Figure 3: Probability that the residual lifetime of a route is smaller than time t (see Equation (12)) vs. the number of hops for static (red), mobile (green) and partially mobile (blue) ad hoc networks. $p_l = 0.1$ and $p_n = 0.05$ are used in the plot.

divided by the total number of remaining nodes. This assumption implies that the probability that a route contains m mobile nodes is hypergeometrically distributed. Thus, if \mathcal{M} is the set of mobile nodes and \mathcal{F} is the set of fixed nodes in the network, then the probability that a route contains m mobile nodes is

$$P[M = m] = \frac{\binom{|\mathcal{M}|}{m} \cdot \binom{|\mathcal{F}|}{N+1-m}}{\binom{|\mathcal{M}|+|\mathcal{F}|}{N+1}}, \quad (11)$$

where $|\mathcal{M}|$ and $|\mathcal{F}|$ denote the number of nodes in each set. Substituting Equation (10) and (11) into (9) finally yields

$$P[T_{\text{route}} < t] = \sum_{m=0}^{N+1} (1 - (1 - p_l(t))^N \cdot (1 - p_n(t))^m) \cdot \frac{\binom{|\mathcal{M}|}{m} \cdot \binom{|\mathcal{F}|}{N+1-m}}{\binom{|\mathcal{M}|+|\mathcal{F}|}{N+1}} \quad (12)$$

To illustrate the effect of the model extension for fixed nodes, we plot Equation 12 in Figure 3 for typical values of fixed nodes in a network of 20 nodes. The curves show the probability that a route breaks within time t depending on the route lengths. For the plot, we assumed the failure probabilities to be $p_n = 0.05$ and $p_l = 0.1$. The bullet curve shows the distribution when all nodes are mobile ($|\mathcal{F}| = 0$, $|\mathcal{M}| = 20$), the crossed curve when all nodes are fixed ($|\mathcal{F}| = 20$, $|\mathcal{M}| = 0$), and the curve with circles when 10 nodes are fixed and 10 nodes are mobile ($|\mathcal{F}| = 10$, $|\mathcal{M}| = 10$). Notice the larger discrepancy of the green and blue curve to the red curve for larger routes.

5. ANALYSIS OF TESTBED TRACES

In this section, we use our framework to analyze the cause of link and route failures in the PDA test network. In particular, we are interested in how often link failures were due to moving nodes. Note that as such, it is not possible to answer this question by just looking at the packet traces. When we observe that a node no longer receives broadcast packets from another node for a certain amount of time, we cannot directly tell if the node moved away or collisions/interference was the cause.

The method to separate the two causes of failures works as follows. We first estimate the failure probabilities p_n and p_l from the given overall route lifetime CDF (see Figure 1). The estimation

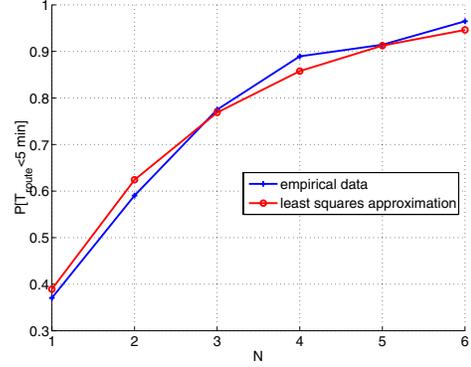


Figure 4: Probability that the lifetime of a route is shorter than 5 minutes versus the route length. Empirical and least square approximation of Equation (12) on empirical data are plotted. Resulting error probabilities are $p_l = 0.33$ and $p_n = 0.04$.

$t[\text{min}]$	p_l	p_n
0	0	0
5	0.33	0.04
10	0.40	0.18
15	0.44	0.30
20	0.46	0.40
25	0.50	0.47
30	0.54	0.54
35	0.55	0.63
40	0.57	0.69
45	0.60	0.74
50	0.61	0.79
55	0.65	0.82

Table 1: Estimated failure probabilities in test network. p_l is the probability that a link breaks within time t due to interference or collisions. p_n is the probability that a node moves within time t in such a way that an existing link breaks.

method is based on the fact that the failure probabilities are used in a different power order in the CDF for the overall route lifetime (see Equation (6) and (12)). Then, we insert the two failure probabilities in the equations for the conditional lifetime distributions. Finally we fit an analytical distribution on the resulting conditional CDF.

5.1 Estimation of Failure Probabilities

To estimate the failure probabilities p_n and p_l , we first plot the overall route lifetime probability for constant time values against the number of hops per route on the horizontal axis. Then, we make a least square approximation of Equation (12) on each curve. An example for $t = 5$ min is shown in Figure 4. The curve with crosses represents the empirical distribution. This distribution stands for the probability that a route of N links will last less than five minutes. The curve with circles is the least square approximation of Equation (12) on the curve with crosses. Each approximation per time value yields an estimation of the corresponding failure probabilities. The failure probabilities for $t = 5$ min are $p_l(5 \text{ min}) = 0.33$ and $p_n(5 \text{ min}) = 0.04$.

By repeating this procedure for different values of t , we obtain the failure probability distributions for $p_n(t)$ and $p_l(t)$. The values for $t = 0, 5, \dots, 55$ min are listed in Table 1. The two distributions

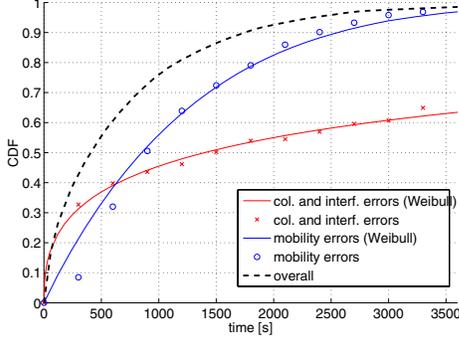


Figure 5: CDF of conditional residual link lifetime

show an interesting trend. For small time values (a few minutes), the probability p_n that a node moves in such a way that a link breaks is significantly smaller than the probability p_l that a link breaks due to interference. For $t = 30$ min, the probability for both failures is equal and for large time values, the failure probability that a link breaks due to node mobility dominates. We find this trend very intuitive. At places with many possible sources of interference, a link will likely break quickly. In contrast, users tend to stay at the same location for quite some time but then move away with a high probability.

Note that for the estimation of the failure probabilities, we assumed that five nodes are fixed ($|\mathcal{F}| = 5$), and the remaining fifteen nodes are mobile ($|\mathcal{M}| = 15$). These values are based on our personal observation on how often people forgot their PDAs at their desk, and how often the PDAs had to be recharged. We also used different values of fixed and mobile nodes to see the sensitivity of these two parameters. We observed that small variations in the parameters do not have an impact on the trends we are presenting.

To get the link lifetime CDF with only mobility failures and with only interference failures, we first insert the estimated failure probabilities p_n and p_l in the analytical equations derived in previous section (Equation (3) and (4)). The result is plotted in Figure 5. The circles represent the values with only mobility failures and the crosses represent the values with only interference failures. For comparison, we also plot the overall link lifetime CDF (including both sort of failures and represented with the dashed curve). When comparing the CDF with only mobility failures to the CDF with only interference errors from the figure, we clearly distinguish two separate regions. For small time values, most links break due to interference or collisions errors. At around $t = 600$ seconds, both CDFs intersect and the probability that a link breaks is equal for both sort of failures. For larger time values, the mobility failures dominate the link failures. Another interesting trend is that in contrast to the CDF with mobility failures which is very close to 1 for $t = 3500s$, the CDF with only interference errors reaches approximately 0.63. That means that when there is no mobility, links break quickly or remain stable for a long period.

We also investigated how well an analytical distribution fits the measurements. Since we are considering lifetime distributions, the Weibull distribution is a natural choice as it is used in reliability theory to model the lifetime of objects. The density function of the Weibull distribution is given by

$$f(t) = \alpha \cdot \beta \cdot t^{\beta-1} \cdot e^{-\alpha t^\beta}, t \geq 0. \quad (13)$$

The Weibull distribution is a generalization of well known distributions such as the exponential distribution or the Raleigh distri-

bution. However, the Weibull distribution models the total lifetime of objects whereas in our paper, we consider the residual lifetime. We therefore transform the Weibull density function for residual lifetimes. We achieve this by using the law of total probability:

$$f_r(t) = \int_{-\infty}^{\infty} f_r(t|\tau) \cdot f_t(\tau) \cdot d\tau, \quad (14)$$

where $f_t(\tau)$ denotes the probability density function (PDF) of the total lifetime. Given the total lifetime τ of a link, the residual lifetime is uniformly distributed in the interval $[0, \tau]$. Therefore, we can write

$$f_r(t) = \begin{cases} \int_0^\infty \frac{1}{\tau} \cdot I\{0 < t < \tau\} \cdot \alpha \cdot \beta \cdot \tau^{\beta-1} \cdot e^{-\alpha \tau^\beta} \cdot d\tau, & t \geq 0, \\ 0, & t < 0. \end{cases} \quad (15)$$

$$= \begin{cases} \int_t^\infty \alpha \cdot \beta \cdot \tau^{\beta-2} \cdot e^{-\alpha \tau^\beta} \cdot d\tau, & t \geq 0, \\ 0, & t < 0. \end{cases} \quad (16)$$

The according CDF is found by integration

$$F_r(t) = \begin{cases} \int_0^t \int_\eta^\infty \alpha \cdot \beta \cdot \tau^{\beta-2} \cdot e^{-\alpha \tau^\beta} \cdot d\tau \cdot d\eta, & t \geq 0, \\ 0, & t < 0. \end{cases} \quad (17)$$

Note that none of the appearing integrals can be solved analytically.

The parameters of the original Weibull distribution which best fits our data and which we used in Figure 5 are $\alpha = 1.35 \cdot 10^{-6}$, $\beta = 1.725$ for the CDF with mobility failures only and $\alpha = 0.021$, $\beta = 0.43$ for the CDF with only interference failures. As we expected, the good fit of the distribution with the empirical values shows that the Weibull distribution models well the conditional link absolute lifetimes in our experiment.

5.2 Conditional Route Lifetime CDF

We plot the Weibull approximations of the CDF for the conditional route lifetimes in Figure 6. In Figure 6(a), we see the CDFs for 2-hops routes, in Figure 6(b) for 3-hops routes and in Figure 6(c) for 6-hops routes. The intersection point of the CDFs for routes of different lengths remains approximately at the same time value around $t = 600$ seconds. However, as the route length increases, the probability that a route remains stable for that long significantly decreases. For 6-hops routes, almost all routes are broken before $t = 600s$. Thus, we conclude that for large routes, the effect of collision and interference errors dominates the effect of mobility failures.

6. COMPARISON WITH MOBILITY MODELS

Many protocols for mobile ad hoc networks have been evaluated with simulations using synthetic mobility models. However, few results were published that verify if these models really reflect the mobility behavior of real users. In this section, we describe our results that are going in that direction. Specifically, we compare the link lifetime from our experiment with the lifetime distribution we obtain from simulations of two popular mobility models: the random waypoint model [10] and the random reference point group mobility model [7]. Note that this direct comparison is only possible because we managed in the previous section to isolate the cause of link failures into mobility failures and failures due to collisions or interference.

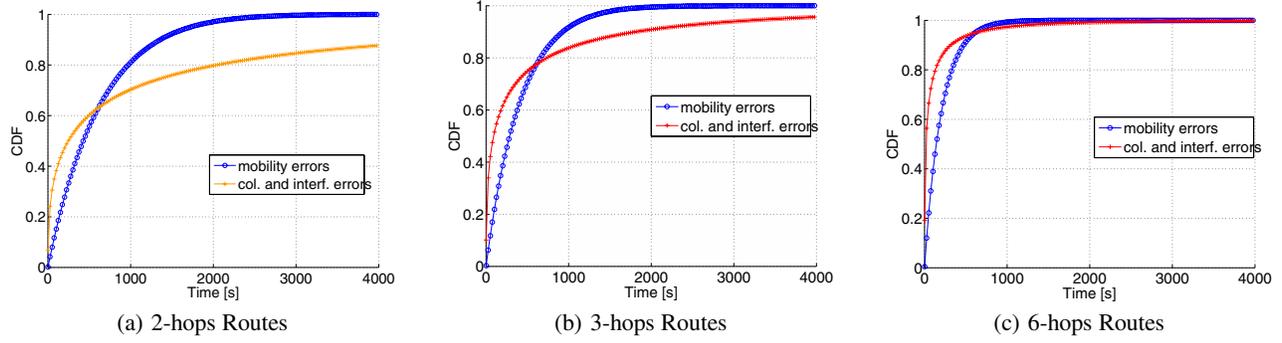


Figure 6: Conditional Route Residual Lifetime

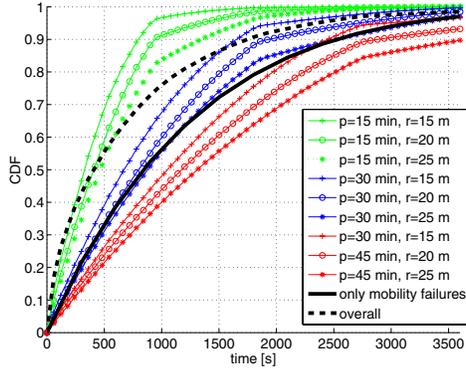


Figure 7: Link Residual Lifetime with Random Waypoint Mobility Model

6.1 Simulation Parameters

We set the parameters of the simulation to be as close as possible to the experiment settings. Thus, the number of nodes is set to twenty, whereas 15 nodes are mobile and five nodes are fixed. The simulation plane is set to a rectangle of 100x20 meters which roughly corresponds to the effective size of the floor where the users were working. The node speed is set to 1m/s corresponding to the walking speed of humans. The wireless range r of each device varies between 15 and 25 meters and we assume that there exists a bidirectional link between two devices if their geometric distance is smaller than the wireless range.

As mobility models, we use the random waypoint and the random reference point group mobility model as described in [10] and [7]. Since the behavior of the node mobility changes after an initial startup phase [3], we remove the beginning of each simulation run and consider the simulations only after we observe convergence. Both, the random waypoint and the random reference point group mobility model, have the notion of pause time to specify the time a node has to wait after it arrives at its destination point. We vary the pause time p between 15 and 45 minutes based on our observation during the experiment.

6.2 Results

We first present the simulation results with the random waypoint mobility model. In Figure 7, we plot the link lifetime CDFs obtained with simulations for different wireless ranges and different pause times. To compare with the empirical CDF, we have to use

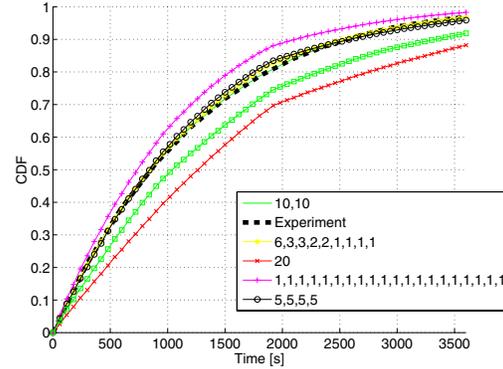


Figure 8: Link Residual Lifetime with Random Reference Point Group Mobility Model

the conditional lifetime CDF with only mobility errors which is the continuous curve in the plot. For completeness, we also plot the overall lifetime CDF (dashed curve) with both sort of failures. The main issue we observe is that the shape of the empirical CDF with only mobility failures fits very well to the CDFs obtained with simulations. For a pause time of $p = 30$ minutes and a wireless range of $r = 25$ meters (which are close to our observations of the test users), the empirical CDF is almost identical to the simulated CDF. Notice that the overall empirical CDF is not matching any simulated CDFs.

We show the results of the simulations with the random reference point group mobility model in Figure 8. We used different numbers of groups and different group sizes. We show the resulting CDF for one group of 20 users (20), two groups of ten users (10,10), four groups of five users (5,5,5,5), nine groups of different sizes (6,3,3,2,2,1,1,1,1), and 20 groups of one user (1,...,1). The empirical CDF is the dashed curve. The CDFs for four groups of five users and nine groups of different size are almost identical to the empirical CDF. The remaining CDFs have a similar shape but depending on the number of groups, links tend to break earlier or later than in our experiment.

We did not plot the route residual lifetime CDFs from the simulations due to space limitations, but they also fit for both models very well to the empirical CDF.

We conclude that both models, the random waypoint and the random reference point group, model user mobility in such a way that the resulting link and route residual lifetime distributions are the

same as in our experiment. For applications and network protocols which are affected by these metrics, both synthetic models are thus suitable.

7. CONCLUSIONS

The goal of this paper was to analyze the impact of human mobility on the link and route lifetime of mobile ad hoc networks. For this purpose, we analyzed the data gathered from a real ad hoc network of 20 PDAs connected via 802.11b wireless interfaces. In order to separate mobility and collisions/interference as causes of link failures, we developed a general framework to model the link and route lifetimes. With the help of this framework, we were able to study the impact of both sort of failures separately. We found that interruptions due to human mobility and collisions/interference have a completely different impact on the lifetime of links and routes. We can distinguish two separate regions: for small lifetimes, most link breaks are due to interference/collision errors while for large lifetimes, most failures occur due to mobility. Such a result is crucial for the design of new MAC and link layer protocols. The more likely the error is due to mobility, the less efforts should be wasted to re-establish the connection. Finally, we compared the empirical link lifetime with those obtained by statistical mobility models. The results show that the distribution of the random waypoint and the random reference point group mobility models are very close the empirical distribution.

Although the framework for the residual lifetimes was developed for this analysis, it is not specific to our testbed. The framework is generic and can be applied to different type of data sets. As part of our future work, we plan to use the acquired insights for the design of new MAC and network layer protocols.

8. 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.

9. 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] C. Bettstetter, G. Resta, and P. Santi. The Node Distribution of the Random Waypoint Mobility Model for Wireless Ad Hoc Networks. *IEEE Transactions on Mobile Computing*, 2(3):257–269, 2003.
- [4] T. Camp, J. Boleng, and V. Davies. A Survey of Mobility Models for Ad Hoc Network Research. In *Wireless Communications and Mobile Computing (WCMC): Special issue on mobile ad hoc networking: research, trends and applications*, volume 2, pages 483–502, 2002.
- [5] F. Fitzek, L. Badia, M. Zorzi, G. Schulte, P. Seeling, and T. Henderson. Mobility and Stability Evaluation in Wireless MultiHop Networks Using MultiPlayer Games. In *ACM NETGAMES '03*, Redwood City, CA, USA, May 2003.
- [6] M. Gerharz, C. de Waal, M. Frank, and P. Martini. Link Stability in Mobile Wireless Ad Hoc Networks. In *Proceedings of the IEEE Conference on Local Computer Networks (LCN)*, Tampa, Florida, USA, November 2002.
- [7] 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.
- [8] 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.
- [9] 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 IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, pages 32–43, Pittsburgh, Pennsylvania, USA, October 2003.
- [10] 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.
- [11] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. In *Proceedings of the Conference on Mobile Computing and Networking (MOBICOM)*, Atlanta, Georgia, USA, September 2002.
- [12] V. Lenders, J. Wagner, and M. May. Measurements from an 802.11b Mobile Ad Hoc Network. In *Proceedings of the IEEE WoWMoM workshop on advanced experimental activities on wireless networks and systems (EXPONWIRELESS)*, Niagara-Falls/Bufalo, New York, USA, June 2006.
- [13] A. B. McDonald and T. Znati. A Path Availability Model for Wireless Ad Hoc Networks. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, New Orleans, USA, September 1999.
- [14] M. McNett and G. M. Voelker. Access and Mobility of Wireless PDA Users. In *ACM SIGMOBILE Mobile Computing and Communications Review*, April 2005.
- [15] L. Petrak, O. Landsiedel, and K. Wehrle. Framework for Evaluation of Networked Mobile Games. In *NetGames '05*, Hawthorne, New York, USA, October 2005.
- [16] N. Sadagopan, F. Bai, B. Krishnamachari, and A. Helmy. PATHS: Analysis of PATH Duration Statistics and their Impact on Reactive MANET Routing Protocols. In *Proceedings of the Fourth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, Annapolis, Maryland, USA, June 2003.
- [17] J. Su, A. Chin, A. Popivanova, A. Goel, and E. de Lara. User Mobility for Opportunistic Ad Hoc Networking. In *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, English Lake District, UK, December 2004.
- [18] C. Tudeuce and T. Gross. A Mobility Model based on WLAN Traces and its Validation. In *Proceedings of INFOCOM*, Miami, Florida, USA, March 2005.