


AN EMPIRICAL STUDY OF THE IMPACT OF MOBILITY ON LINK FAILURES IN AN 802.11 AD HOC NETWORK

VINCENT LENDERS, JÖRG WAGNER, SIMON HEIMLICHER, MARTIN MAY,
AND BERNHARD PLATTNER, ETH ZÜRICH



A key point in coupling research and real-life applications is to understand how real-world conditions impact practical networking aspects. To gain more realistic insights, the authors deploy an indoor IEEE 802.11 mobile ad hoc network comprising 20 PDAs carried by volunteers for one week.

ABSTRACT

A great deal of research has been done during the past few years in the area of wireless self-organizing networks. Generally, this research has been supported by either simulation or theoretical analysis, both relying on strong assumptions. However, a key point in coupling research and real-life applications is to understand how real-world conditions impact practical networking aspects. To gain more realistic insights, we deploy an indoor IEEE 802.11 mobile ad hoc network comprising 20 PDAs carried by volunteers for one week. In a subsequent analysis, we explore the impact of mobility and interference on the observed network behavior. A major finding of our analysis is that mobility is the most dominant cause of link failure for links with a long lifetime, whereas other causes (unrelated to mobility) are responsible for the breakage of links with short lifetimes. This inherent property could be used by network protocols in self-organizing networks to optimize link or route repair decisions depending on the age of a link at the time it fails.

Introduction

Mobile ad hoc networks consist of mobile devices communicating over wireless links without any support from a fixed infrastructure. Originally envisioned for a vast set of applications, such as disaster recovery or tactical communication, mobile ad hoc networks also provide attractive opportunities to connect mobile users in urban areas at a low cost, based on portable devices like PDAs, mobile phones, media players, and so on. We expect that these kinds of ad hoc networks will enable a multitude of new useful applications in the near future. For example, users could exchange music files, photos, or podcasts directly when they are in proximity without the necessity to detour through an intermediate fixed infrastructure.

Although many have acknowledged the potential benefits of mobile ad hoc networks of

personal devices, there have been surprisingly few reports on real deployments and analyses so far [1, 2]. As a consequence, many systems and protocols for mobile ad hoc networks have been designed and evaluated based on a set of strong assumptions. For example, many ad hoc routing protocols assume bidirectional links. Another assumption that thoroughly impacts protocol design is that the network remains connected over time even while nodes move. However, both these assumptions have been shown to be inadequate in real-world environments (e.g., [3]). Designing systems and applications that eventually will work in the field first requires a profound understanding of the underlying characteristics and behavior of real networks.

In this article, we take a step in this direction and present an analysis of an 802.11b wireless ad hoc network with real user mobility. The ad hoc network comprises 20 PDAs that are carried by volunteers working on the same floor in an office building for one week. Our analysis explores common network characteristics and how link failures are affected by mobility. In addition, we compare how different prevailing mobility models affect the network characteristics as an attempt to characterize the mobility in our real-life experiment.

IEEE 802.11 MOBILE AD HOC NETWORK IN AN OFFICE ENVIRONMENT

The traces for our analysis were collected from an experimental ad hoc network we deployed at ETH Zurich. The network consists of 20 identical HP iPAQs connecting via IEEE 802.11b in ad hoc mode. To determine connectivity in the network, each device periodically sends an IP broadcast packet every 0.5 seconds. The devices in direct transmission range that receive such a broadcast packet store the arrival time of the packet, the identity (the IP address) of the sender, and a sender-specific sequence number of the packet on an external compact flash card

for later offline analysis. These broadcast packets are the only traffic we inject into the network.

Twenty volunteer test users carried the devices during five consecutive working days (from 10 a.m. to 5 p.m.). The test users were researchers, staff members, and students of a networking research lab, all working on the same floor having a rectangular shape of a size of 100 meters by 30 meters. The test users were instructed to carry the PDA with them throughout the day and to recharge the battery whenever necessary. 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, going to the rest room, picking up printouts in the hallway, or meeting each other for discussions. Only a few test users occasionally left the building or the campus for a certain period during the experiment. Therefore the network is relatively dense.

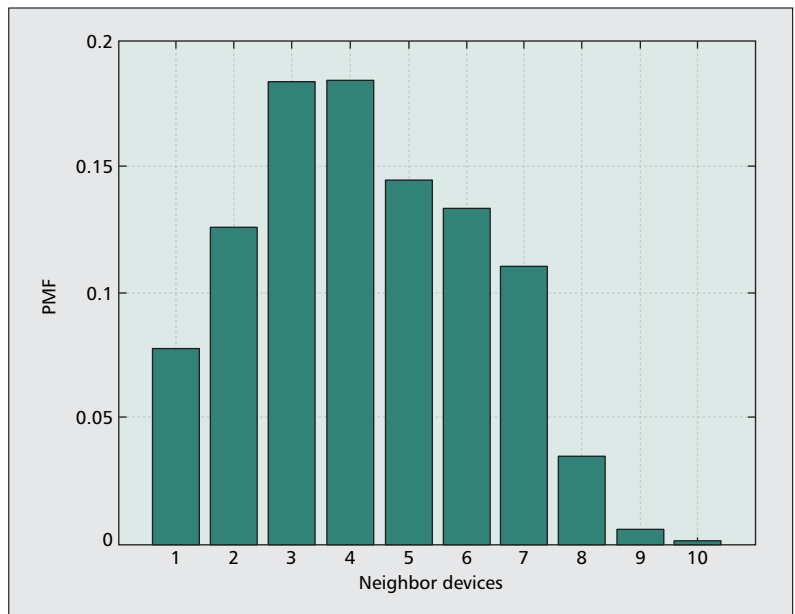
Because our experiments took place in a productive environment, other WLAN and Bluetooth devices that were not part of our testbed but were emitting at neighboring frequency bands could interfere with transmissions from our network. However, due to the high penetration of WLAN in buildings today, such external sources of interference are realistic in office environments.

NETWORK ANALYSIS

Our analysis is done offline based on the connectivity we observe from the periodic broadcast packets in the traces. To facilitate the analysis, we calculate a *connectivity graph* for each broadcast time interval. In a connectivity graph of time t , there exists a link from node a to node b if b receives at least 50 percent of all packets from a within a time window $[t - \tau, t]$ where τ is a short time period (seven seconds in this article). Based on these connectivity graphs, we determine the following metrics:

- *Node degree*: The node degree of a node is the number of outgoing links it has in the connectivity graph. This metric indicates whether users tend to form large groups or stay alone.
- *Path length*: The path length is the length in hops of the shortest path between two nodes.
- *Link and route lifetime*: The lifetimes capture the level of dynamics in the network. There are two statistical view points to look at lifetimes:
 - The *total lifetime* of a link describes the time interval between the moment the link appeared until it breaks.
 - The *residual lifetime* represents the time interval between a sample moment after the creation until the link or path breaks.

Generally, it is not relevant whether the total or the residual lifetime is used, as 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 because communication starts at arbitrary moments and not necessarily when a new route becomes available. For the remainder of the article, we always refer to the residual lifetime



■ Figure 1. Node degree.

unless stated otherwise. The route lifetime is defined by counting the remaining lifetime of the shortest route between two nodes. Note that because the nodes are mobile, it is possible that while assessing the lifetime of a route, a shorter alternative route becomes available. However, we always count the remaining lifetime of the initially computed shortest route.

Our analysis of these metrics is based on their average over 35 hours (five days of seven hours).

NODE DEGREE


The probability mass function (PMF) of the node degree averaged over all devices is plotted in Fig. 1. The peak of the distribution is at four neighbors; the mean is at 4.25. Because the distribution is close to zero for ten 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 ten nodes in direct transmission range.

PATH LENGTH

The PMF of the shortest path length averaged over all node pairs during the experiment is plotted in Fig. 2. As expected, most nodes are close to each other and have thus a distance of one hop (with a probability of approximately 0.43). Surprisingly, for a network size of 20 nodes that were relatively close to each other, there were some long shortest-paths, consisting of five, six, and even seven hops. Note, however, that these long paths are routes in the connectivity graphs that were obtained from the periodic beacon messages. This does not mean that data packets were actually sent over those multihop routes.

LINK AND ROUTE LIFETIME

The cumulative distribution function (CDF) of the route lifetime is plotted individually in Fig. 3 for routes of different lengths. The distribution shows that many links tend to break after a short time. Consider, for example, the link lifetime



■ Figure 2. Shortest path length.

distribution (one hop). After 100 seconds, 20 percent of the links are unavailable; after 500 seconds, 55 percent of the links are not available anymore. Interestingly, for long time intervals, we find that a significant number of links are still available. For example, after 3500 seconds, the distribution shows that approximately 3 percent of the links are still available.

We can see that for larger routes, the lifetime significantly decreases. For two-hop routes, the probability that a route lasts longer than 3500 seconds is almost equal to zero. And for six-hop routes, the probability is almost zero after 500 seconds.

IMPACT OF MOBILITY ON LINK AND ROUTE FAILURES

Our next goal is to understand more deeply how user mobility impacts the link and route lifetimes shown in the previous section. The issue is that these lifetimes are not only affected by the node mobility, but also by other sources of failures like packet collisions and interference. Unfortunately, with standard 802.11b hardware, it is not possible to directly determine the cause among these two fundamentally different sorts of failures without the help of motion detection or location sensors. For this reason, we have developed a model that allows us to distinguish the causes of failures based on the statistical properties of the collected traces.

The basic idea of the model is to separate the link failures into two classes, based upon the cause of the problem:

- Failures that are due to node mobility
- Failures like collisions or interference, which are independent of node mobility

We assign to each of these causes a separate failure probability, which allows us to deduce two CDFs for the lifetime:

- One conditioned that mobility is the only cause of failures

- The other conditioned that all failures are due to other reasons

We provide a brief sketch of the model before we employ it for analyzing the cause of link and route failures in the collected traces. The details of the model are given in [4]. Let us consider the following probabilities of failure within residual time T (starting from an arbitrary time instance of observation):

- $p_n(T) = P[T_{\text{link}} < T]$: The probability that the residual lifetime of a link between two nodes, T_{link} , is less than T , given the link breaks due to node movement.
- $p_l(T) = P[T_{\text{link}} < T]$: The probability that the residual lifetime of a link between two nodes, T_{link} , is less than T given the link breaks due to disturbances from other devices (packet collisions or general interference from other sources transmitting at the same frequency). Subsequently, we refer to interference to describe this type of failure.

We implicitly assume that all nodes in the network behave identically and independently of each other. Moreover, we make the assumption that the two failure events themselves are statistically independent of each other. Finally, we assume that no other than these two failure events occurred.

Our goal is to estimate both the $p_n(T)$ and $p_l(T)$ of the CDF through the empirical distributions of the residual lifetime of N -hop routes (not conditioned on the failure source) for different N . These empirical distributions are immediately accessible in our traces.

To extrapolate the $p_n(T)$ and $p_l(T)$ of the CDF, we exploit a distinct difference in the way these distributions enter the probability of a route of N hops being interrupted due to *any* reason within less than time T . This difference is understood by realizing that a route failure caused by mobility can either be due to movement of the source node, movement of any relay node, or movement of the destination node. Route failures caused by interference, in contrast, are solely due to issues arising at receiving nodes, that is, not due to the source node. In essence, for a route of N hops, the probability that the route is interrupted due to any reason in less than time T from the observation time instance is given by $p^{(N)}(T) = 1 - (1 - p_n(T))^{N+1} (1 - p_l(T))^N$. Here, the powers $N + 1$ and N account for the fact that there are $N + 1$ statistically independent chances of a failure due to node mobility, but only N chances for a failure due to interference.

We obtain estimates of $p_n(T)$ and $p_l(T)$ by matching the $p^{(N)}(T)$ of the CDFs, for $N = 1, \dots, 6$ with the respective empirical CDFs obtained from the traces. More specifically, for arbitrary T , we obtain $p_n(T)$ and $p_l(T)$ through a least-squares approximation of the empirical CDFs through $p^{(N)}(T)$.

We finally approximate the estimated $p_n(T)$ and $p_l(T)$ of the CDFs by a distribution derived from a Weibull distribution. The Weibull distribution is commonly used for modeling the absolute (as opposed to residual) lifetime of objects. Because the distribution of the absolute lifetime is fully determined by the distribution of the residual lifetime (and vice versa), we obtain the

corresponding residual lifetime distribution through a straightforward change of variables in the probability distribution function (PDF). Figure 4 shows the Weibull distributions optimally approximating $p_n T(T)$ and $p_l(T)$ in a least-squares sense upon the change of variables, that is, our estimated distributions for the absolute link lifetimes conditioned on failures due to mobility and interference, respectively.

The key observation in Fig. 4 is the following: depending on the time horizon, one type of failure is more likely than the other. Accordingly, we can infer the reason why a link breaks depending on how long it has existed. Assume that we divide the lifetime axis into three regions (I, II, and III in Fig. 4). When a link breaks after a short time (region I), it is most likely that the link broke due to interference. If a link breaks in region II, we cannot reliably determine the reason of the failure because the probability is approximately equal for both sorts of failures. However, if a link breaks in region III (for example, after 1500 seconds), it is more likely that the link has broken due to node mobility. The information why a link break occurred could be used by a routing protocol to decide, for example, if a link should be removed from the routing table or not. In particular, if a link breaks due to interference or packet collisions, it might reappear in the near future and should thus not be removed from the routing table. Other applications of this information in the medium access control (MAC) or transport layers are also imaginable.

COMPARISON OF TRACES WITH SIMULATIONS

To obtain more insight into the characteristics of the mobility of our subjects, we compare the empirical node degree and link lifetime distributions with results from three popular stochastic mobility models. Although none of these models is designed to match the mobility of people in an office environment, the comparison allows us to see which of the diverse models matches best our empirical data.

SIMULATION OF MOBILITY MODELS


We employ a simple simulation model that often is used in existing literature: nodes are connected to all other nodes that are within transmission range of their radio device. Time is slotted into intervals of one second. The remaining simulation parameters match our experiment, that is, duration is 35 hours, simulation area is $100 \text{ m} \times 30 \text{ m}$, node speed is 1 m/s , and the transmission range is 20 m .

We use three widely used mobility models that span a wide spectrum of mobility patterns:

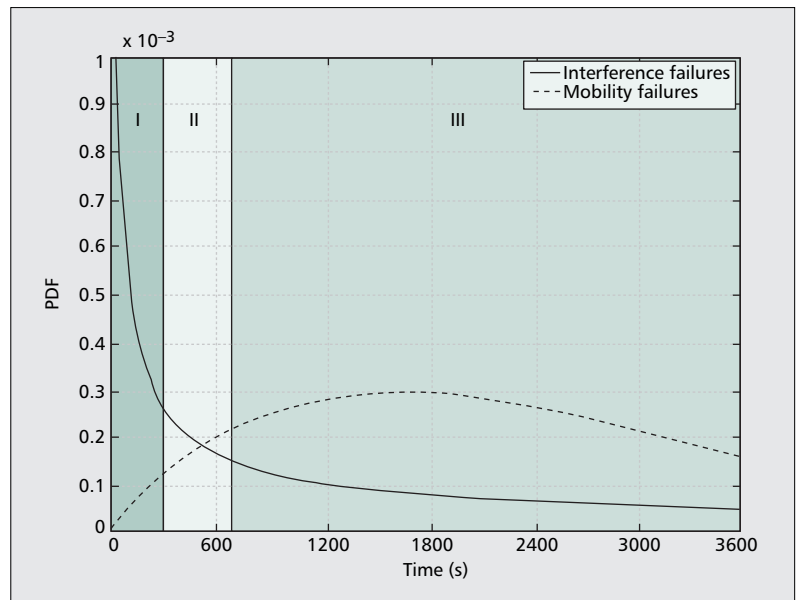
- The random waypoint [5]
- The random walk [6]
- The Manhattan mobility model [6]

In all three models, nodes begin their journey from points uniformly distributed in the area and then proceed differently as follows:

- In the *random waypoint* model, nodes move toward the next so-called waypoint, a point that is chosen from a uniform distribution



■ Figure 3. CDFs of link and route residual lifetime.



■ Figure 4. Conditional total link lifetime PDF with only interference failures.

upon arrival at the previous waypoint. This results in nodes moving in the same direction for a rather long distance of 36.5 m on average.

- In the *random walk* model, nodes move one *movement unit* in a direction chosen randomly between 0° and 360° and then pick another direction, similar to Brownian motion. In contrast to the random waypoint model, the random walk model has nodes moving in zigzaggy lines or roam about a limited area for extended periods.
- In the *Manhattan* model, node movement is constrained to a Cartesian grid. Nodes move along a grid line until they encounter an intersection, where they move ahead with probability 0.5 or turn right or left with probability 0.25, respectively. Thus, in contrast to the

other two models, the Manhattan model allows only a finite number of locations in any time slot. The movement of nodes is similar to random walk, but the probability to keep moving ahead is larger than turning, and the probability to reverse is zero.

The latter two models have a parameter that controls how local the mobility is. We use the two values 2m and 5m for the movement and grid unit, respectively, to study the impact of these parameters. In the random waypoint model, nodes pause for a certain time before changing direction. Because the people in our measurements also pause at their table or at

lunch for considerable periods of time, we adapt the random walk and the Manhattan model accordingly and add a pause time as follows: nodes draw a distance from the distance distribution of the random waypoint model and then pause after this distance. The pause time is distributed exponentially, and its mean is set to match the mean total-link lifetime; on condition that there are only mobility failures (Fig. 4).

NODE DEGREE DISTRIBUTION

Figure 5 shows the PMF of the node degree for the different models and our empirical data. The three mobility models produce very similar distributions, albeit with different maxima. The random walk model with a 2-m movement unit (denoted by move = 2m) and the Manhattan model with a 5-m grid unit (grid = 5m) have more weight of the distribution around the maximum value.

The degree distribution of our empirical data exhibits quite a different shape, particularly at degree values of five and seven. This indicates that the node distribution in our experiment was non-uniform, as both Manhattan and random walk have uniform node distributions, and the node degree solely depends on the node position in our simulations.

LINK LIFETIME DISTRIBUTION

To compare the mobility patterns of our empirical data with the simulated models, we use the residual-link lifetime distribution; conditioned there are only mobility failures ($P[T_{\text{link}} < t | p_i(T) = 0]$). The results are shown as a logarithmic plot of the complementary cumulative distribution function (CCDF) in Fig. 6. As the link lifetime is affected only by the mobility characteristics of the models, this plot allows us to draw conclusions about the kind of mobility the people in our experiments performed. Although the slope of the three models differs considerably, all simulated data approximate an exponential distribution. In the particular scenario we consider, the random-waypoint model seems to provide the best match for the empirical data. To improve the fit of a random walk and the Manhattan model, we would have to increase the step and grid size, respectively, resulting in movements with longer straight lines.

CONCLUSIONS

We have analyzed the characteristics of an indoor 802.11b ad hoc network with real user mobility. Our traces include measurements of more than 35 hours trace time. We first analyzed the network topology. A deeper analysis of the link and route lifetimes showed that mobility and interference, the two dominating failure causes, have a completely different impact. On the one hand, when links with short lifetimes fail, it is usually due to interference or packet collisions. On the other hand, when links with a long lifetime fail, it is usually because of node mobility. For instance, this information could be used by routing protocols to optimize route repair decisions. Finally, we compared the impact of the real user mobility to the random

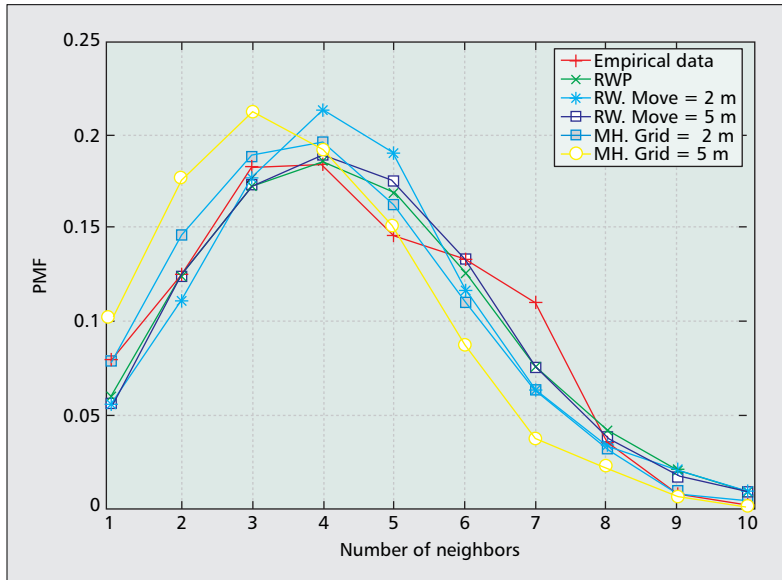


Figure 5. Node degree distribution of the mobility models compared to the empirical result. The model names are abbreviated as follows: RWP: random waypoint; RW: random walk; MH: Manhattan. The parameters of the RW and MH models are: Move: movement unit; Grid: grid unit. The mean node degree of the empirical data and the three models is approximately 4.25.




Figure 6. Comparison of the conditional empirical link residual lifetime distribution with simulations of the mobility models.

waypoint, the random walk, and the Manhattan mobility model. The comparison of these three popular mobility models has shown that the random waypoint mobility model exhibits a very similar node degree, as well as link lifetime distribution as the corresponding empirical distributions obtained from our measurements, whereas the other two models have considerably longer link lifetimes with our choice of parameters.

REFERENCES

- [1] A. Chaintreau *et al.*, "Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms," *Proc. IEEE INFOCOM*, Barcelona, Spain, Apr. 2006.
- [2] J. Su *et al.*, "User Mobility for Opportunistic Ad Hoc Networking," *Proc. IEEE Wksp. Mobile Comp. Sys. and Apps.*, English Lake District, UK, Dec. 2004.
- [3] D. Kotz, C. Newport, and C. Elliott, "The Mistaken Axioms of Wireless-Network Research," Dartmouth College Comp. Sci. tech. rep. TR2003-467, July 2003.
- [4] V. Lenders, J. Wagner, and M. May, "Analyzing the Impact of Mobility in Ad Hoc Networks," *ACM/Sigmobile REALMAN Wksp.*, Florence, Italy, May 2006.
- [5] D. Johnson and D. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks," *Mobile Computing*, T. Imelinsky and H. Korth, Eds, vol. 353, Kluwer Academic, 1996, pp. 153–81.
- [6] T. Camp, J. Boleng, and V. Davies, "A Survey of Mobility Models for Ad Hoc Network Research," *Wireless Commun. and Mobile Comp.*, Special Issue on Mobile Ad Hoc Networking: Research, Trends, and Applications, vol. 2, 2002, pp. 483–502.

BIOGRAPHIES

VINCENT LENDERS [M] (lenders@tik.ee.ethz.ch) holds an M.Sc. (2001) and a Ph.D., both in electrical engineering, from the Swiss Federal Institute of Technology in Zurich (ETH Zurich), Switzerland. After his Ph.D., he was appointed as a doctoral research fellow at Princeton University, New Jersey. Since 2008 he has been working as a research fellow at armasuisse, Switzerland. His main research interests include mobile ad hoc networking and security. He is a member of ACM and the Internet Society. He has also served on the program committees of multiple networking conferences.

JÖRG WAGNER (jwagner@nari.ee.ethz.ch) received a Master's degree in information technology and electrical engineering from ETH Zurich in October 2006. He is currently working toward his Ph.D. at the same university. His research interests lie in the broad fields of physical layer design and

applied information theory, as well as wireless and ad hoc networking.

SIMON HEIMLICH [M] (heimlicher@tik.ee.ethz.ch) received a Master's degree in information technology and electrical engineering from ETH Zurich in 2005 and is currently pursuing his Ph.D. degree at the same institution. His research interests are modeling and analysis of data transport in partially connected networks. He is a member of the ACM.

MARTIN MAY [M] (may@tik.ee.ethz.ch) received a Master's degree in computer science from the University of Mannheim in 1996. In 1999 he received his Ph.D. degree at INRIA Sophia Antipolis from the University of Nice, France. He did most of his thesis work on Internet QoS mechanisms at INRIA, but was also a technical staff member of Lucent Bell-Labs Research, Holmdel, New Jersey, and Sprintlabs, Burlingame, Massachusetts. Until the beginning of 2000, he continued his research as a post-doctoral member of research staff at Sprintlabs, Burlingame. In 2000 he founded a startup company in France where he worked in the field of content networking. After selling the company in 2004, he went back to academia where he joined ETH Zurich as a senior research associate. His research interests are in future Internet architectures and network security. Since 2006 he has coordinated a large EU-funded project on Future Internet Technologies with the goal to develop new networking paradigms and node architectures for a future Internet, the ANA project. In 2008 he joined Thomson Research Laboratory, Paris, France, where he leads research activities in mobile networking and future Internet technologies. He is a member of ACM and the Internet Society. He is an editor of *ACM Sigcomm Computer Communication Review*, has chaired multiple workshops and conferences, and has also served on technical program committees for many networking conferences.

BERNHARD PLATTNER [M] (plattner@tik.ee.ethz.ch) is a professor of computer engineering at ETH Zurich, where he leads the Communication Systems Group. He has been the principal investigator or co-P.I. of numerous national and international projects in the area of computer networking. His current research interests are in self-organizing networks, mobile ad hoc networks, and practical aspects of information security. He has also directed research on active networks, starting as early as 1996, and multimedia applications for high-speed networks. In 1996–1998 he served as head of the Faculty of Electrical Engineering at ETH Zurich. In 2005–2007 he was vice-rector of ETH for Bachelor/Master studies. He is a member of the ACM and the Internet Society. He served as the program or general chair of various international conferences, such as ACM SIGCOMM '91, INET '94, and IWAN '02, and has served on the program committees of other major conferences, such as IEEE INFOCOM.

A major finding of our analysis is that mobility is the most dominant cause of link failure for links with a long lifetime, whereas other causes (unrelated to mobility) are responsible for the breakage of links with short lifetimes.