Link Stability in Mobile Wireless Ad Hoc Networks

Michael Gerharz, Christian de Waal, Matthias Frank, Peter Martini
Institute of Computer Science IV, University of Bonn, Germany
{*gerharz,dewaal,matthew,martini*}*@cs.uni-bonn.de*

Abstract

In this paper, we develop adaptive metrics to identify stable links in a mobile wireless networking environment based on the analysis of link durations in several different mobility scenarios. Our metrics only rely on online statistical evaluation of observed link durations. Neither do they require information on signal strength, radio conditions, or spacing of the mobile devices, nor do they depend on the availability of additional hardware such as GPS receivers or a synchronisation of the devices. We demonstrate the ability of the metrics to select stable links with a high probability in a wide range of scenarios.

1. Introduction

The number of applications available for wireless communications is growing rapidly: mobile telephony is ubiquitous nowadays, wireless hotspots are spreading everywhere, and also ad hoc networking is growing mature these days. A key characteristic of these scenarios is the dynamic behaviour of the involved communication partners. Communication protocols will have to deal with a frequently changing network topology. However, many applications require stable connections to guarantee a certain degree of QoS.

In access networks, access point handovers may disrupt the data transfer. In addition, service contexts may need to be transferred to the new access points, introducing additional overhead and delays to the connection.

In ad hoc networks, mobile services enable peer-to-peer connections for voice or data traffic. Using stable links is crucial for establishing stable paths between connection peers. Rerouting is especially costly in these networks without infrastructure, since it usually results in (at least partly) flooding the network.

The *stability* of a link is given by its probability to persist for a certain time span, which is not necessarily linked with its probability to reach a very high age.

Little work has been published so far on this topic. The related concept of signal stability, well known from cellular networks, has been used to find the right time for a handover. Variations in the received signal strength may hint to the movement pattern of the connection peers and thus allow an estimation of a probable connection loss. However, received signal strength is largely dependent on actual radio conditions. Due to fading effects those measurements are subject to large fluctuations.

Another method of estimating the distance and the relative speed of mobile nodes is the use of GPS receivers. But generally, using GPS is infeasible under many conditions. It is e.g. unavailable in indoor environments and it is not suitable for small devices due to its high power consumption.

Using the current age of a link to predict the residual lifetime is a relatively unexplored approach, and fundamental research on the relation of a link's age to its residual lifetime is still missing. This raises the question if a simple and reliable method for stable link selection may be based upon a statistical analysis of link durations.

The rest of this paper is structured as follows. In the following section we provide an overview of existing research on the topic of link stability in mobile networks. In section 3, we present a characterisation of the mobility models that we used for our evaluations. Section 4 summarises the results on link durations in several different scenarios. In section 5, we explain how the observations from section 4 can be used to choose more stable links even in highly dynamic scenarios. After a few thoughts on practical implementation issues in section 6, we summarise our results and outline perspectives for further work in section 7.

2. Related Work

One of the earliest works in the context of link stability is the development of *Associativity Based Routing* (ABR) [11]. The idea behind this ad hoc routing protocol is to prefer *stable* links over *transient* links. A link is considered to be stable if it exists for a time of at least $A_{\text{thresh}} = 2r_{\text{tx}}/v$, where r_{tx} is the transmission range and v denotes the rela-

This work was supported in part by the German Federal Ministry of Education and Research (BMBF) as part of the IPonAir project (http://www.IPonAir.de/).

tive speed of two devices. It is left open how to determine the relative speed v among the mobiles which in turn determines A_{thresh}. The significance of this value will become clear in later sections. ABR measures the lifetime of a link using hello messages which are periodically broadcast.

The motivation behind this approach has been found in assuming an implicit grouping for links that reach a certain age. After a time of A_{thresh}, it is assumed that the nodes move with a similar speed in a similar direction and thus are likely to stay together for a relatively long period of time. However, this assumption is justified only intuitively for dynamic scenarios.

Signal Stability Adaptive Routing (SSA) [5] follows a similar approach. It distinguishes *strongly connected* from *weakly connected* links where a link is considered to be strongly connected, if it has been active for a certain predefined amount of time. However, this concept is considered only as a supplement to SSA's signal strength based approach and has been found to perform poorly in [5].

Also based on signal strength measurements is the *Routelifetime Assessment Based Routing* (RABR) ([2]). It tries to predict the time when the received signal strength falls below a critical threshold using a measured value of average change in received signal strength.

Another prediction method for link durations is presented in [6]. The method is based on distance measurements between mobile devices. A refinement in [7] takes possible changes in speed or direction of motions into account. The distance between connection peers may be acquired with the help of GPS receivers or signal strength measurements. Apart from the shortcomings of these two methods, the problem with this approach is that the distance of a receiver is only a very vague hint on link availability. In realistic environments the coverage area of a radio transmission hardly ever has a circular shape and is subject to strong fluctuations.

A further approach based on the availability of GPS measurements has been suggested in [10]. The *Flow Oriented Routing Protocol* (FORP) follows a similar approach of calculating a link's residual lifetime from a mobile's own speed and the speed and distance of the connected party. However, this method strongly depends on the assumption of a free-space propagation model and on having GPS equipment available for distance measurements and time synchronisation. These requirements can hardly be presumed in a realistic environment.

A theoretical link availability prediction method is presented in [9]. However, this method is specifically based on the Random Walk Model (which is very similar to the Random Waypoint Model described in section 3) and it is not clear how well the results apply to real world scenarios. Furthermore, the method actually estimates the availability of a radio link at a certain point in time, not its stability.

3. Mobility Models & Simulation Setup

This section provides an overview of the mobility models and the simulation setup which we used to evaluate link stability in mobile wireless networks.

3.1. Mobility Models

One of the best known mobility models is the *Random Waypoint* model (cf. e.g. [3]) in its numerous variants. The simplicity of this model is probably the reason for its widespread use: At the beginning of the simulation, the nodes are randomly placed on the simulation area. Each node picks a random destination on the area and moves there with a random speed. Having reached the destination, the node pauses for a random time. All random choices mentioned are uniformly distributed in predefined ranges.

Variations of this model ([4]) include choosing directions rather than destinations or preferring certain destinations, so called *attraction points*.

Several studies analyse the characteristics of the Random Waypoint model (e.g. [4]). First of all, node distribution is non-uniform in this model. Rather, a clustering of nodes in the centre of the modeled area can be observed. While this could be prevented by choosing directions rather than destinations, another property is common to both approaches which is that a node may change its movement direction and speed abruptly. Obviously, this results in a rather jagged movement pattern which some might deem unnatural.

A plethora of more sophisticated mobility models have been proposed in the literature (cf. [3] for a survey). Most of them concentrate on particular selected aspects of mobility.

Thus, although one single model may not be able to provide a complete picture of link stability in real world, analysing different mobility models, all of which model very different scenarios, should provide insight to how well stable links may be identified in a wide range of scenarios. Put in more general words, the adaptiveness of a protocol can best be judged by applying it to a wide range of different scenarios.

In the following, we describe the models that we used in our simulations in addition to the Random Waypoint model.

Movements in the *Gauss-Markov* model [8] are smoother than the sudden stops and sharp turns which arise from the Random Waypoint model. In this model, a mobile has a certain speed and direction, which are updated in discrete time intervals. The new values are randomly chosen from normal distributions with the old values as means.

Since the motion vectors are changed regularly, there is a stronger erratic component present than in the Random Waypoint model, where mobiles moving into each other's transmission range are likely to keep their directions and velocities until they part again.

With the *Manhattan Grid* mobility model [1], a grid is placed on the simulation area, modelling streets or, on a smaller scale, paths in an exhibition hall or the like. The mobiles move along the paths, change their speed with a certain probability in intervals of certain distances, and at crossings, there is a certain probability for the mobiles to turn.

As the distances between the grid points grow larger against the transmission range of the devices, an implicit grouping is provoked in the sense that mobiles which move on the same path are connected for a certain period of time.

We have extended this model for our simulations to support pause times of mobiles and defined a reasonable non-zero minimum speed for devices in motion. The pause times were implemented to allow a better comparison to the other models, which also have this feature. In the Gauss-Markov model, mobiles may choose a speed of zero. While [1] states that a minimum speed of zero also applies to the Manhattan Grid model, this cannot be consistently integrated with this models idea to change the behaviour of mobiles in distance intervals (in contrast, the Gauss-Markov model uses time intervals). Another problem arising with the use of distance intervals is that extremely low speeds are kept up for long time spans.

3.2. Simulation Setup

The speed of the mobile nodes was assumed to be walking speed, i.e. up to 1.5m/s in most of the cases, however some simulations were done with higher maximum speeds of up to 20m/s to show that the results are also valid for increased mobility.

The transmission range of the radio devices was chosen as 50m with a free space propagation model. A link was considered to be established when two nodes reached each other's transmission radius and considered broken when their distance exceeded the transmission radius.

Although this assumption is based on ideal radio conditions, it may be expected that our results will principally be similar when using more realistic propagation models. This is justified in section 6.

Transmission range and speed of the mobile nodes are related parameters, since scaling the speed up has the same effect as scaling the transmission range down, as long as the area dimensions are adjusted accordingly.

The number of nodes was chosen to allow for both a connected network and a sufficiently large sample of link durations. For all models, 25 nodes per 100m x 100m simulation turned out to be a reasonable choice. It should be noted that this parameter is not critical when the view is limited to link lifetimes. A lower node density results in lower probability for the network being connected which does not seem useful for ad hoc networking scenarios. A higher node density would enlargen the number of measured link durations which would be even more favourable for the presented method.

Therefore, 225 nodes were simulated in a 300m x 300m area for all scenarios presented in this paper.

Further parameter settings of the scenarios presented throughout this paper are the following, if not stated otherwise: In Gauss-Markov scenarios, the movement vectors were updated every 10s, and the standard deviations used for speed and angle were 0.1m/s and $\pi/8$, respectively. In Manhattan Grid scenarios, a new speed was chosen every 10m with a probability of 0.2, and the turn probability at a crossing was set to 0.5. When using pause times, the mobiles decided to pause with a probability of 0.05 every 10m. A block size of 100m x 100m leads to scenarios containing 3 x 3 blocks (16 grid points).

The simulation time has to be chosen carefully in order not to influence the outcome of the measurements. First of all, the initial part of the simulations may not be representative for the link duration distribution for several reasons. Some of the mobility models show a different distribution of mobile nodes once they have reached a stable state (e.g. Random Waypoint Model). Therefore, the evaluation of the simulation should start after the models have reached a stable state. Furthermore, the initial part must be long enough as to allow for old links. It does not make sense to start the evaluation of the simulation at a time when there is no decent chance to see links having reached a reasonably old age. The same is true for the end of the simulation. Here, the set of samples is not representative, because links which are still active at the end of the simulation have to be omitted from the evaluation, since their residual lifetime is not known.

4. Link Characteristics in Different Mobility Models

In section 2 we described how ABR and SSA favour older links over younger ones. Reasons for this approach are given only intuitively, but are not justified analytically. To study this heuristic, we analysed link lifetimes in scenarios generated with the mobility models described above with a number of different parameter settings. The simulations show that the key to stable link selection is more complex than just selecting the oldest link.

4.1. Link Durations

In order to get a first impression of the distribution of link durations, we counted the observed link durations from a simulation run and plotted them into a histogram. Figures 1, 2, and 3 present some examples for several mobility models. The parameters used for these scenarios are the

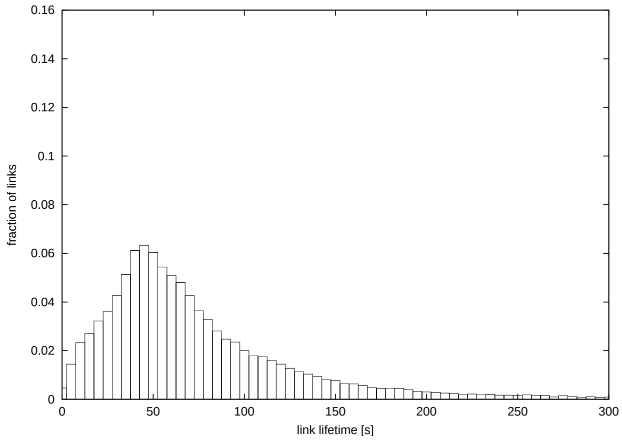

Figure 1. Distribution of link durations in a Gauss-Markov scenario

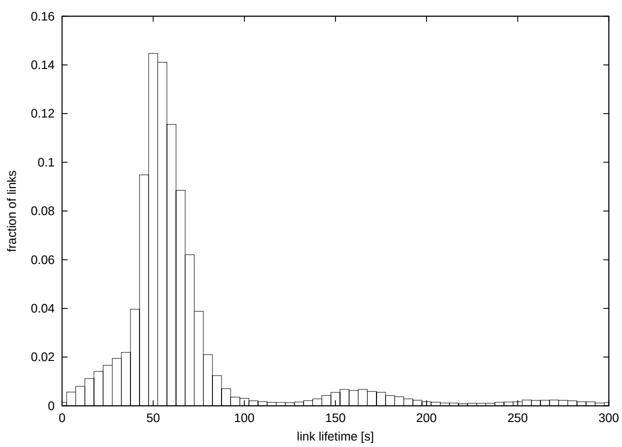

Figure 2. Distribution of link durations in a Manhattan Grid scenario

standard parameters described above; the Manhattan Grid and Random Waypoint scenarios are without pause times. About $80,000$ links were recorded in the Gauss-Markov and Manhattan Grid scenarios, and about $150,000$ in the Random Waypoint scenario, pointing out the latter's high dynamics.

All link duration histograms have in common a peak roughly at the transit time of two mobiles crossing each other's transmission range on a (nearly) direct path, i.e. at the familiar A_{thresh} from the ABR algorithm. This suggests that a more appropriate choice for A_{thresh} should at least add an offset to this threshold to incorporate the fact that a large number of links is likely to disappear immediately when chosen at that age, not exactly what can be called stable.

Older links occur relatively seldom, with less than 20%

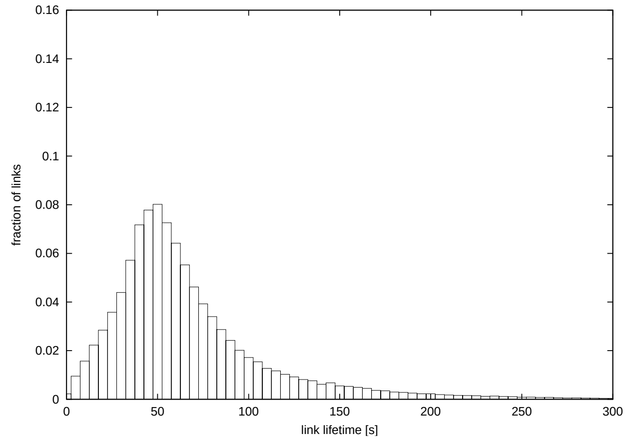

Figure 3. Distribution of link durations in a Random Waypoint scenario

reaching 100s in many scenarios. Other characteristics are rather specific to the mobility models and will be elaborated in the following.

Depending on the parameters for the simulation setup, the Random Waypoint and Gauss-Markov scenarios showed rather similar characteristics. First, we describe how the variation of these parameters influences the outcome of the simulations.

In the Random Waypoint scenarios, a larger simulation area permits the choice of more distant destinations. If two mobiles choose similar movement patterns to reach far away destinations, this will result in a long link duration.

In Gauss-Markov scenarios, a similar effect may be observed when the update intervals are increased. The longer these intervals get, the more seldom the short links occur. This indicates that for short update intervals, there is a high probability that two mobiles moving towards each other will choose different movement directions already in their next update and hence separate again. As an overall result we observe that the less frequently these updates occur, the more do the Gauss-Markov scenarios resemble the Random Waypoint scenarios.

Changing the speed variation in a Gauss-Markov scenario has a strong influence on the link durations. With higher values for this parameter, short links occur much more often.

In contrast, increasing the angle variation did not make links more short-lived. This is due to the fact that with higher variations in the movement vectors' angles, the mobiles tend to erratically move around in locally bounded areas. As an illustration, one can think of the motion patterns becoming more and more similar to those of a Brownian motion, where all angles between 0 and 2π are equally likely.

In order to detect stable links in Gauss-Markov scenarios it is necessary to identify those nodes that have come so close (relatively to the degree of change) that their chance of leaving each other's transmission range due to a different decision is low.

In Manhattan Grid scenarios, the grid structure, and especially the block size, determines the distribution of link durations. In principle there are two cases how two nodes may meet: They may come from opposite directions, cross each other's transmission range, and separate again. This results in the link durations corresponding to the first peak in figure 2 and is by far the most common case. This peak is more distinct than with other models, because two nodes almost always move on a direct path when they meet. Alternatively, two nodes may move in the same direction when they meet. From that moment on, they will stay connected as long as they take the same decisions on the following crossings. At every crossing, it will become clear if they stay together for another block crossing time. This is visible in the consecutive peaks in figure 2. Of course, these effects get blurred when pause times are introduced to the model and when speed variations are higher. The clue to finding stable links would be to identify those links that connect to nodes moving in the same direction.

Now, considering the strong variations in the link durations, is it possible to select the best of several available links in terms of residual lifetime without knowing the topology of the area and the movement patterns of the mobile nodes? This leads to the question of how residual link lifetimes are distributed.

4.2. Residual Link Lifetime

Mathematically, the density d_a of residual link lifetimes for links of age a may be calculated from the density d and the distribution D of the link lifetimes as

$$d_a(t) = \frac{d(t+a)}{1-D(a)} = \frac{d(t+a)}{\int_a^\infty d(t)dt}. \qquad (1)$$

Intuitively, this means shifting the graph of the density function of the link durations left by a and scaling the y-axis such that the area between the graph and the x-axis is 1 again. Obviously, d_0 equals d.

The mean value E_a of the residual lifetime of a link of age a is given by

$$E_a = \int_0^\infty t d_a(t)dt = \frac{\int_a^\infty t d(t)dt}{\int_a^\infty d(t)dt} - a = \frac{E - D(a)}{1-D(a)} - a, \quad (2)$$

where E is the mean value of the link durations.

Figures 4, 5 and 6 show the averages and three quantiles of the distributions approximated by recording link lifetimes in different simulation runs. The intervals around

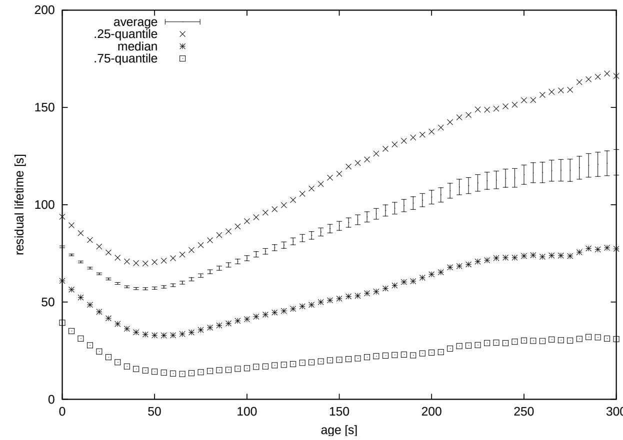

Figure 4. Residual lifetime in a Gauss-Markov scenario (cf. fig. 1)

the average values give the 0.95 confidence intervals for the mean value.

Obviously, the distribution of link durations lacks the memoryless property. Instead, the residual lifetime of a link is ruled by its current age.

A common phenomenon to all scenarios considered is a minimum at the modal value of the link duration distribution. Independently of the model used, average residual link lifetimes for young links drop by a constant factor with each second the links grow older (in equation 1, the denominator decreases only slightly due to the small number of links going down at the beginning). In an intuitive sense this means that not many assumptions may be made concerning the movement of two mobiles that just came into each other's transmission range. Thus, the expected residual lifetime is dominated by the fact that for every single link, its remaining time decreases with increasing age.

As already pointed out in section 4.1, these observations suggest that ABRs threshold value for stable links, A_{thresh}, should be chosen slightly higher than the average transit time of two nodes. Take a look e.g. at the .75-quantile in figure 5(a) where a Manhattan Grid scenario is presented. If a link is between 50s and 75s old, there is a 75% probability for it to disappear almost immediately. Shifting A_{thresh} up would improve stability significantly.

As the age of a link grows beyond the density's modal value, its probability to disappear in the next moment diminishes – the expected residual lifetime starts to rise again (the denominator in equation 1 becomes small). This is depicted by the assertion that if links reach a certain age, it becomes likely that the movement vectors of the corresponding nodes are similar and that they will stay connected for an even longer period of time, in the Manhattan Grid model e.g., because they took the same decision at the last cross-

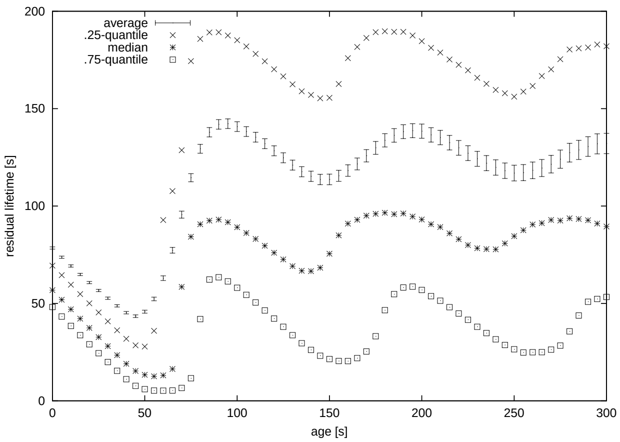

(a) continuous movement (cf. fig. 2)

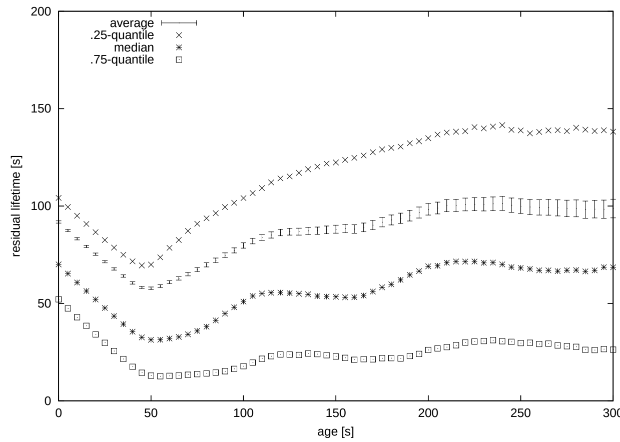

(b) pause approx. every 200m for max. 150s

Figure 5. Residual lifetime in a Manhattan Grid scenario

ing.

We want to emphasise that the rise in average residual lifetime in some of the scenarios (e.g. figure 4) is mainly due to only a small fraction of links with extremely high lifetimes. This becomes visible when comparing the average residual lifetime with the plotted quantiles. The medians do not rise as sharply as the corresponding averages and run well below the averages. Thus, the high average value must be caused by relatively few links with exceptionally high residual lifetimes. The .75-quantiles, i.e. the minimum residual lifetimes of a majority of all links, point this up even more, because their increase is much slower than the average lifetime. As a consequence, old links will not be stabler than young links in many cases, although they have a higher expected residual lifetimes.

As a consequence, using the average value as link selection criterion should yield the highest link lifetime *on average*, but this should not be misunderstood as *stability*. Imagine e.g. a user requesting a mobile service. He is not interested in how long a link *might* live. If he needs a reliable link for a rather short period of time, a link with a certain chance to grow extremely old is not helpful if this chance is small. Analogously, a multi-hop path breaks when just one link goes down, and therefore it does not make a path long-lived if it includes just a small fraction of links that grow very old. In both examples, risk should be minimised by selecting links having good chances of being available for a certain period of time, albeit on the cost of a smaller average link lifetime.

We will now go into some details on what causes the partly very different shape of the curves for the various scenarios.

The residual lifetimes in Manhattan Grid scenarios have an oscillating shape for old links. This corresponds to the successive peaks in the distribution of link durations(fig. 2). In section 4.1 it has been pointed out that these iterations are due to common decisions on grid crossings. If two nodes that share a link decide to take the same direction at a crossing this usually means they will stay together for the whole grid length. To illustrate this, figure 5(a) shows a particularly extreme example for this behaviour. This example was simulated without pause times and a quite low speed standard deviation (0.2m/s). The effect fades with higher standard deviations in speed and the introduction of pause times (cf. figure 5(b) for an example with a maximal pause time of 150s).

The impact of the pause time is particularly evident in the Random Waypoint scenario of figure 6. With a pause time of 300s, the oldest links have a residual lifetime just as small a the local minimum of links with ages of around 50s.

Generally, the average residual lifetime of links in Random Waypoint scenarios does not rise as drastically as in the other models when using the same simulation area. Nevertheless, with larger simulation areas the mobiles get a higher chance of moving together without one of them changing its direction or stopping abruptly, thus provoking a higher gain for old connections.

Why do older links have a much higher residual lifetime in Gauss-Markov scenarios compared to Random Waypoint scenarios? In the Random Waypoint model, the new movement vectors are independent of the previous ones. Thus, a similar movement in the past is not an indicator for a similar

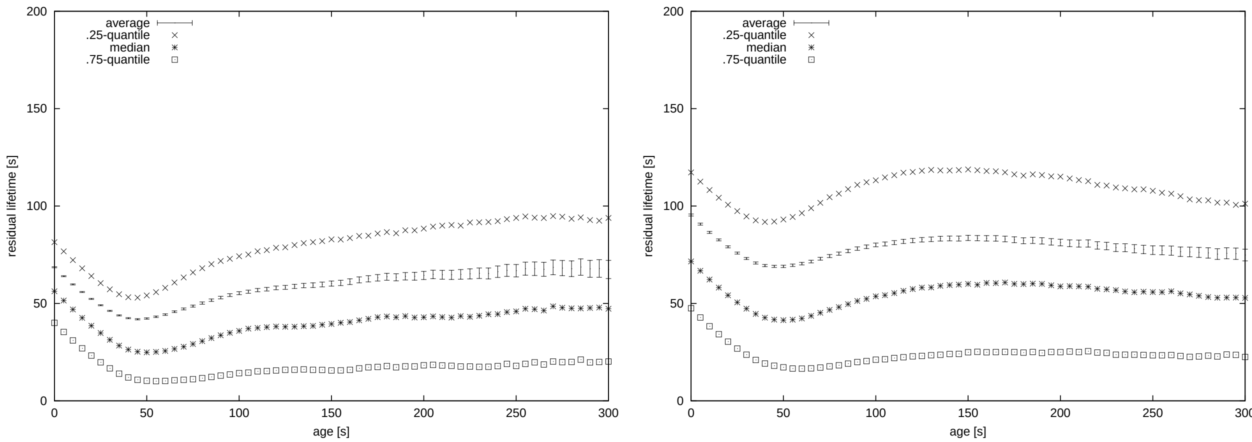

(a) continuous movement (cf. fig. 3)

(b) max. 300s pause time

Figure 6. The influence of the pause time in Random Waypoint scenarios

movement in the future. In contrast to this, in the Gauss-Markov model, the new movement vectors are determined based on the old ones. Once two mobile nodes chose similar parameters, they have the same seed values for their future decisions.

Furthermore, a non-zero relative speed and angle between two mobiles' movement vectors determines when the link between them disappears, given the movement vectors are not updated before. Since these updates are frequent in Gauss-Markov scenarios, there is a decent chance for differences between the movements of two mobiles to be averaged out in the long run. In contrast, this is very unlikely in Random Waypoint scenarios.

For these reasons, Gauss-Markov scenarios tend to show a steeper rise in residual lifetimes (cf. fig. 4). Of course, this effect fades when the update intervals grow large or the standard deviations for the speed changes rise.

5. Estimating Link Stability

The considerations in section 4 lead to the derivation of two practical methods which enable a mobile device to estimate residual link lifetimes solely based on observations of link lifetimes in the past. It should be emphasised that the goal of these considerations is not to predict the exact residual lifetime of a link. This is not important when deciding which of several links are stable, meaning they are most likely of all to stay available for some time.

As a side note, predicting the lifetime of a link seems to be a very difficult task when considering the high variations we see in section 4.1.

5.1. Metrics

To enable mobile devices to make smart decisions on the basis of the observations from section 4, they must save the link durations they observed. In practice, the link lifetimes would be counted in an array d using a certain granularity, which for the sake of simplicity we assume to be unit of time in the following. This array has a finite number of $N+1$ elements. Thus, links with a lifetime in $[a-0.5; a+0.5[$ would be counted as having lived a units of time where $a \in \{0, \ldots, N\}$. In the following formulas, the terms a and $a+s$ are to be discretised to this domain.

The straightforward way for a mobile device to estimate the average residual lifetime of a link of age a is to calculate the following formula, following equation 2:

$$R(a) = \frac{\sum_{t=a}^{N} t\,d[t]}{\sum_{t=a}^{N} d[t]} - a \tag{3}$$

This would imply a complexity linear in the number of array elements for the evaluation of a link and a constant complexity to update the information about link durations. Alternatively, arrays managing values for the numerator and the denominator in the formula above could be used to evaluate a link's quality with constant complexity and update the link duration information with complexity linear in the number of array elements.

A user requesting a stable link for a well-known time span s might want to calculate the probability for the link to stay available for this time as

$$P_s(a) = \frac{\sum_{t=a+s}^{N} d[t]}{\sum_{t=a}^{N} d[t]}. \tag{4}$$

However, exact specifation of s cannot be expected to be known in advance, so that instead, the device could calculate some α-quantile as

$$Q_\alpha(a) = \max\{s \mid P_s(a) \geq \alpha\}. \qquad (5)$$

Here, the same complexity considerations apply as above.

As a summary, we have identified two link stability metrics: When choosing the link with the highest average residual lifetime of a link based on equation 3, it may be expected to reach high residual link lifetimes on average over many choices. In contrast, choosing the link with the highest α-quantile for some α based on equation 5 hopefully minimises the risk that the link's residual lifetime is very short. In our experiments, we have used the .75-quantile. Finding the optimal quantile is subject to further research.

We compare these two metrics to the following reference metrics: Randomly choosing a link and not considering stability at all is the approach of most protocols so far, and choosing the oldest link is an approach motivated by some previous publications on link stability ([11], [5]).

5.2. Results

The results obtained for the metrics presented above were compared in several mobility scenarios with a duration of three hours. A node's decisions were based on the statistical data of the observations from the past 30 minutes, and the link duration array with a granularity of 5s contained 100 elements. An aging of the knowledge base is essential in reality, especially in scenarios that show temporally or locally different characteristics.

For each node, the selections according to all of the above criteria were recorded for 120 randomly chosen points in time, restricted to the second hour of the simulations for the reasons explained in 3.2.

For each choice i, let c_i denote the residual lifetime of the link chosen and r_i the average residual lifetime of all available links. With these values, we calculated the *average gain* as

$$g = \frac{\sum_i c_i}{\sum_i r_i} - 1.$$

At first we take a look at the average gain of the different methods. Figure 7 shows a representative overview comparing the Random Waypoint scenario without pause times (cf. figures 3 and 6(a)), the Manhattan Grid scenario with pause times (cf. figure 5(b)) and two Gauss-Markov scenarios, one of which has an increased standard deviation for the speed with 0.5m/s (for the other scenario, cf. figures 1 and 4).

The different metrics show very different behaviour across the various mobility models. The oldest link metric performs best in scenarios where similar decisions of

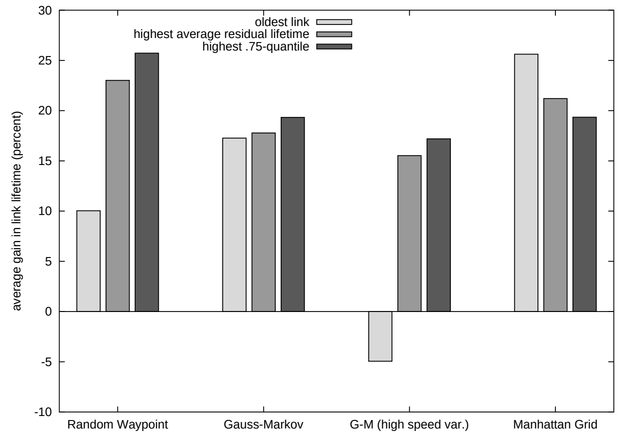

Figure 7. Average gain in link lifetime for different link metrics

two nodes have the greatest impact on the residual lifetime. This is e.g. the case in Manhattan Grid scenarios (as visible in figure 7), where taking the same turn usually results in staying together for the whole lane. Another example would be a Random Waypoint scenario on a large area where nodes may choose very distant destinations, so that similar movements of two nodes result in a long connection time. However, if there are frequent changes in the movement patterns of the nodes, the oldest link metric might perform even worse than randomly choosing a link, as in the Gauss-Markov scenario with the high variation in speed.

The average based metric always outperforms the random metric. The results are relatively stable across a diverse set of scenarios we have evaluated. The gain compared to a random selection was consistently above 15%. Thus, the estimation of the average residual lifetime seems to adapt to a wider range of mobility patterns. However, in certain scenarios, the average gain is significantly smaller than with the oldest link metric. The reason for this is the increased probability for the latter approach to choose links that have an exceptionally long residual lifetime. In other words, the average based metric sometimes prefers young over old links where it should not.

This may be due to the lack of statistical data on very old links. Since these links occur relatively seldom, estimation based metrics maybe favour a young link over an older link only because the statistics indicate a short lifetime of the very few (and thus unrepresentative) samples including values that are decreased to the maximum value that can be stored in the finite array.

The performance of the quantile based metric is comparable to that of the average based metric.

A different picture is obtained by studying the probability of a link failure during the serving of a request. In

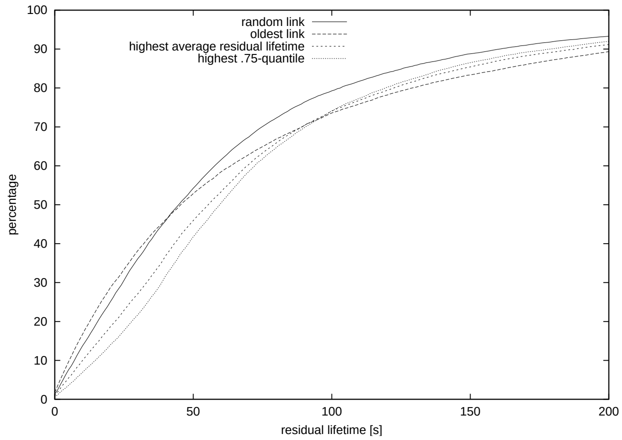

Figure 8. Distributions of residual link lifetimes of links chosen according to different metrics in the Gauss-Markov scenario

figure 8, the distributions of residual link lifetimes of links selected according to the different metrics are depicted for the Gauss-Markov scenario with the lower speed variation. While the different strategies show similar average gain values for this scenario, our proposed metrics show significantly better results in avoiding short-lived links with a residual lifetime of less than about 90s. Also note that choosing the oldest link yields even more links with a residual lifetime below about 40s than randomly choosing a link.

These conclusions become obvious in figure 9, where we have plotted a comparison of the different link rating metrics for various mobility models and for different applications which require a link for 10s and 120s respectively. (Please notice the different scaling of the diagrams.)

Especially if a link is required for a relatively short amount of time, the most promising approach seems to be the quantile based metric. Apparently, being more conservative here in order to a have a good chance that most of the links reach a required age proves to be more robust. The average based metric also performs well and consistently provides better results than randomly choosing a link. Depending on the scenario parameters, the probability to choose an unfavourable link may be reduced to about a third when compared to the random choice strategy.

At first sight, it may be surprising that using the oldest link metric often is even worse than randomly choosing a link, given its high average gain. The reason is the inability of this metric to judge when the oldest of the available link is a particularly bad choice. It has been pointed out before that the oldest link metric has a relatively high risk of selecting particularly short lived links.

For long lasting links, the situation is different. The probability of choosing a link that reaches at least the age of 120s is biggest for the oldest link metric. However, this probability is only one out of five anyhow, and the performance of the other metrics (including the random choice) is close to that of the oldest link metric.

All in all, the benefits of the estimation based metrics are clearly visible in most of the scenarios. The metrics are able to adapt to a wide range of scenarios and minimise the risk of selecting particularly short living links. This advantage increases even more when the link selection criteria are applied to the path selection process of a routing protocol. Since a path becomes unavailable when just one link goes down, the failure probabilities rise exponentially in the number of hops.

6. Practical Considerations

How mobile devices obtain knowledge about link lifetimes, depends on the link layer protocol. Bluetooth e.g. requires an explicit connection setup before communication between devices is possible, whereas with IEEE 802.11 (Wireless LAN), this is not the case and devices are able to receive packets from all other devices, given they use the same channel.

If an explicit connection setup is not provided, beacons must be broadcast regularly to exchange neighbor information. However, a device should not yet consider a link to exist after the reception of just one beacon. Bear in mind that with more diffuse radio conditions links may oscillate between up and down due to fading effects. To eliminate these oscillations, a stabilisation period should be introduced which in effect means a reduction of the transmission range. This should not affect the principal results presented in this paper.

A device should have knowledge about other devices in its neighbourhood for several reasons. Some transmit power control algorithms try to optimise between network connectivity and spatial re-use by adjusting the number of neighbours. Many routing protocols need link state information, and quality of service support is hardly imaginable without it. For efficiency reasons, there will most probably be a generic beaconing layer for connectionless link layer protocols.

7. Conclusions & Further Work

In this paper, we have examined the concept of identifying stable links in mobile wireless networks based solely on the mobile devices observations of link ages. We have shown that across several different mobility models, the expected residual lifetime of a link varies with its age. Based on these observations, two link stability metrics were derived and compared to a metric which favours the oldest

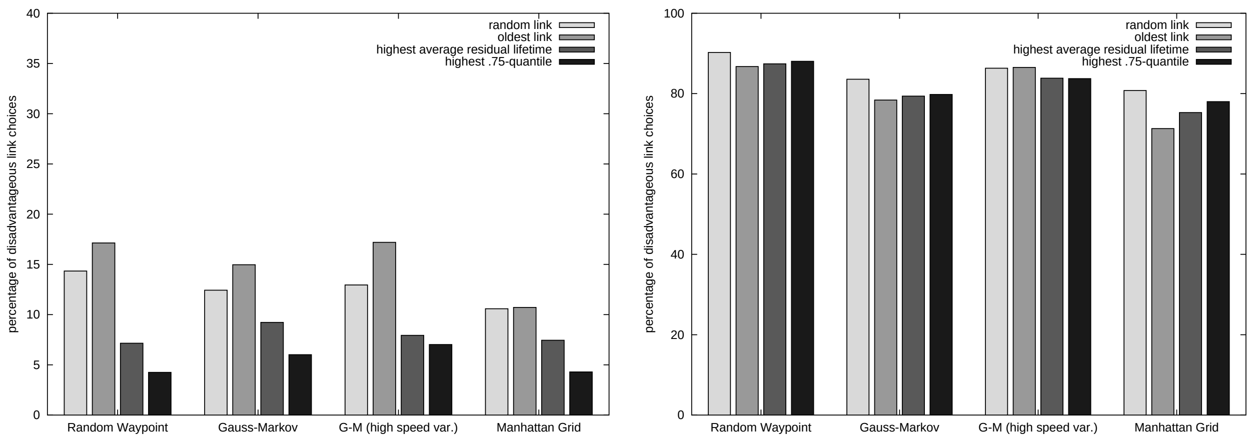

(a) Fraction of chosen links with a residual lifetime of less than 10s

(b) Fraction of chosen links with a residual lifetime of less than 120s

Figure 9. Probability of a link failure during a request

link as well as to a random choice of links. Always choosing the oldest link sometimes works very well, but often bears the risk of getting particularly bad choices. Choosing the link based on its estimated average residual lifetime outperforms the random choice in all scenarios. The conservative approach of using the .75-quantile instead of the average value outperformed the other metrics in most of the scenarios. Especially if a link is required only for a short period of time, the chance to choose a wrong link may be lowered to less than a third when compared to randomly choosing a link. The two stability metrics were shown to be adaptive to a wide range of different scenarios due to their online analysis of observed link durations.

Further work will concentrate on refining the metrics. Additionally, it will be interesting to see whether the performance of the metrics may be enhanced by taking into account other factors such as signal strength. Finally, extending the link stability metric to a path rating metric seems promising for the use in mobile ad hoc networks.

References

[1] Selection procedures for the choice of radio transmission technologies of the UMTS (TS30.03 v3.2.0). TS 30.03, 3GPP, April 1998.

[2] S. Agarwal, A. Ahija, J. P. Singh, and R. Shorey. Route-lifetime Assessment Based Routing (RABR) Protocol for Mobile Ad-hoc Networks. In *Proc. IEEE International Conference on Communications 2000 (ICC'00)*, volume 3, pages 1697–1701.

[3] C. Bettstetter. Smooth is Better than Sharp: A Random Mobility Model for Simulation of Wireless Networks. In *Proc. 4th ACM International Workshop on Modeling, Analysis, and Simulation of Wireless and Mobile Systems (MSWiM'01)*, pages 19–27, Rome, Italy, July 2001.

[4] C. Bettstetter and C. Wagner. The Spatial Node Distribution of the Random Waypoint Mobility Model. In *Proc. 1st German Workshop on Mobile Ad-Hoc Networks (WMAN'02)*, pages 41–58, Ulm, Germany, March 2002.

[5] R. Dube, C. D. Rais, K.-Y. Wang, and S. K. Tripathi. Signal Stability based Adaptive Routing (SSA) for Ad-Hoc Mobile Networks. *IEEE Personal Communication*, February 1997.

[6] D. He, J. Shengming, and R. Jianqiang. A Link Availability Prediction Model for Wireless Ad Hoc Networks. In *Proc. of International Workshop on Wireless Networks and Mobile Computing, Taipei, Taiwan*, April 2000.

[7] S. Jiang, D. He, and J. Rao. A Prediction-based Link Availability Estimation for Mobile Ad Hoc Networks. In *Proc. IEEE INFOCOM*, pages 1745–1752. IEEE, April 2001.

[8] B. Liang and Z. Haas. Predictive distance-based mobility management for PCS networks. In *Proc. IEEE INFOCOM*. IEEE, March 1999.

[9] A. B. McDonald and T. F. Znati. A Path Availability Model for Wireless Ad-Hoc Networks. In *Proc. IEEE WCNC*, pages 35–40. IEEE, September 1999.

[10] W. Su, S. Lee, and M. Gerla. Mobility Prediction and Routing in Ad Hoc Wireless Networks. *International Journal of Network Management, Wiley & Sons*, 11:3–30, 2001.

[11] C.-K. Toh. Associativity Based Routing For Ad Hoc Mobile Networks. *Wireless Personal Communications Journal, Special Issue on Mobile Networking and Computing Systems*, 4(2):103–139, March 1997.