Strategies for Finding Stable Paths in Mobile Wireless Ad Hoc Networks

Michael Gerharz, Christian de Waal, Peter Martini
Institute of Computer Science IV, University of Bonn
D-53117 Bonn, Germany
{*gerharz,dewaal,martini*}*@cs.uni-bonn.de*

Paul James
Nokia Research Center
D-44807 Bochum, Germany
paul.james@nokia.com

Abstract

In this paper, we introduce statistical methods to estimate the stability of paths in a mobile wireless ad hoc environment. Identifying stable paths helps to reduce control traffic and the number of connection interruptions. By means of simulation, we analyse the stability of paths chosen according to a variety of strategies, including those used by the well-known routing protocols AODV and DSR, under a variety of different mobility patterns. This offers new insights into the relation between a path's stability and other characteristics and shows that our statistical metrics are able to identify stable paths in a wide range of scenarios.

1. Introduction

In the past years, a vast number of routing protocols for mobile wireless ad hoc networks has been introduced. One reason for the considerable effort in this area is that routing in such networks is a very challenging task, since network topology changes occur frequently. Each time this happens, an established route may break. If this is the case, data packets must be re-routed quickly, which potentially involves a high amount of control traffic and yet may not be able to avoid interruptions perceived by the users of realtime applications. Therefore, the re-routing of existing connections should be carried out as seldom as possible; in other words, routes should be established along "stable", i.e. durable paths.

In the design of most ad hoc routing protocols, this issue has not been addressed. E.g., DSR simply tries to establish any shortest path in order to save bandwidth. AODV establishes the path along which the RREQ was propagated with the lowest delay, which is likely to be a path with generally low delay and hop count.

Only a few protocols aim at establishing stable routes (cf. section 2), and these either depend on specific hardware features or on unverified assumptions. In this paper, we provide some fundamental work on the issue of path stability.

In section 2, we present previous work in the area of path stability. In section 3, we give a brief introduction to our approach for link stability estimation, which has previously been presented in [4]. In section 4, we present several approaches to estimate the stability of paths. Some of these are simple, intuitive strategies used in related work, others utilise the link stability metrics presented in section 3. In section 5, we present the results of simulations comparing the different approaches to conventional routing strategies that try to find shortest paths (e.g. DSR) or close-to-shortest paths (e.g. AODV). Finally, section 6 draws conclusions and outlines subjects for further work.

2. Related Work

The idea behind *Associativity Based Routing* (ABR) [11] is to prefer *stable* links over *transient* links. A link is considered 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 relative 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 motivation behind this approach is the assumption that after a connection time of A_{thresh}, the corresponding nodes are likely to be moving with a similar speed and direction and thus to stay together for a relatively long period of time. Similar stability concepts have been used in a plethora of ad hoc routing protocols, e.g. [9]. However, in [4], it is shown that the assumption that old links are stabler than younger ones is not generally applicable.

Signal Stability Adaptive Routing (SSA) [2] 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. Paths are established exclusively along stable links. In contrast to ABR, the motivation behind the stability threshold in SSA is merely that weak links are likely to suffer from signal strength fluctuation and

---

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

should thus be avoided. In this sense, SSA is an optimisation independent of the stability approaches discussed in this paper.

Also based on signal strength measurements is the *Routelifetime Assessment Based Routing* (RABR) ([1]). 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. However, their approach relies on the movement patterns being rather steady. Furthermore, their stability estimation disregards the effects of path loss as well as the possibly strong fluctuations in signal strength caused by small scale fading effects.

Different prediction heuristics for a link's residual lifetime are introduced in in [10] as well as in [5], which is refined in [6]. Both approaches rely on the availability of GPS receivers or equivalent equipment to acquire distance and velocity information of neighbouring nodes. Apart from the disadvantages associated with these methods, e.g. inavailability in indoor environments or strong battery demands, the problem with this approach is that the distance of a receiver is only a very vague hint on link availability in realistic scenarios.

A theoretical path availability prediction method is presented in [8]. However, this method is specifically based on the Random Walk Model 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 until that point.

[7] points out that shortest paths are often rather instable, an observation that we can confirm with this paper. As the reason they identify the so-called *edge effect*: Shortest paths are likely to consist of long-distance links. However, a long-distance link is likely to break already as the result of rather small positional changes of the corresponding nodes.

3. Link Stability

A number of stability approaches is based on the assumption that old links are more stable than younger links. However, this is not valid in general, in fact the opposite might be true as well. Nevertheless, a non-trivial dependency of a link's residual lifetime on its current age may be observed in many scenarios. As an example, figure 1 compares the average residual lifetime of a Random Waypoint and a Manhattan Grid scenario (cf. sec. 5.1) each with 500 nodes roaming on a square area of 10^6 m^2 with a transmission range of 140m. A detailed analysis of residual link lifetimes in different mobility scenarios may be found in [4].

This dependency can be utilised to select a stable link among several alternatives with high probability. Several possible metrics have been introduced in [4] and are described in section 3.2 together with their strengths and weaknesses. Since some of these metrics rely on statistical data of link lifetimes, section 3.1 addresses how nodes attain this information.

3.1. Collecting Statistical Data On Link Lifetimes

Devices may count observed link lifetimes in an array d of length $N+1$ using a certain granularity, which for the sake of simplicity we assume to be unit of time in the following. Thus, lifetimes in $[a-0.5;a+0.5[$ would be counted as lifetimes of a time units, $a \in \{0,\ldots,N\}$. Terms describing link ages are to be discretised to this domain.

Important parameters are the array size (i.e. the number of array "buckets"), the granularity, and the observation time span (i.e. the time span after which recorded information is considered outdated and purged).

The array size is mainly limited by the available memory. For a fixed number of buckets, the granularity determines the maximal lifetime that may still be recorded. This means that the granularity should not be chosen too fine in order to ensure meaningful statistics for old links. However, with decreasing granularity, the estimations based on the collected link lifetime data also lose accuracy.

The observation time span must be set to satisfy two conflicting goals: It must be long enough to allow for the collection of a representative sample of link lifetimes (note that the amount of data collected does not only depend on the observation time span, but also on the density of the network). On the other hand, a too long observation span bears the risk of using outdated information.

3.2. Link Stability Metrics

This section provides a brief overview of the link stability metrics introduced in [4] and their underlying assumptions.

Select the oldest link. Several scenarios show a large expected residual lifetime for old links, e.g. Random Waypoint scenarios with long pause times or Gauss-Markov scenarios with large angle standard deviations. Thus, in these scenarios it is a straightforward and easy to implement strategy to select the oldest link. However, this strategy is only applicable in scenarios with a high probability of achieving a sufficient number of long lived links. Otherwise, the oldest link might easily be one whose expected residual lifetime is in the local minimum e.g. observed in figure 1(a). Furthermore, some scenarios even show shorter expected residual lifetimes for old links, e.g. some Gauss-Markov scenarios with large speed standard deviation.

Select the youngest link. Almost all scenarios show a descent of residual lifetimes in the early ages of a link as ob-

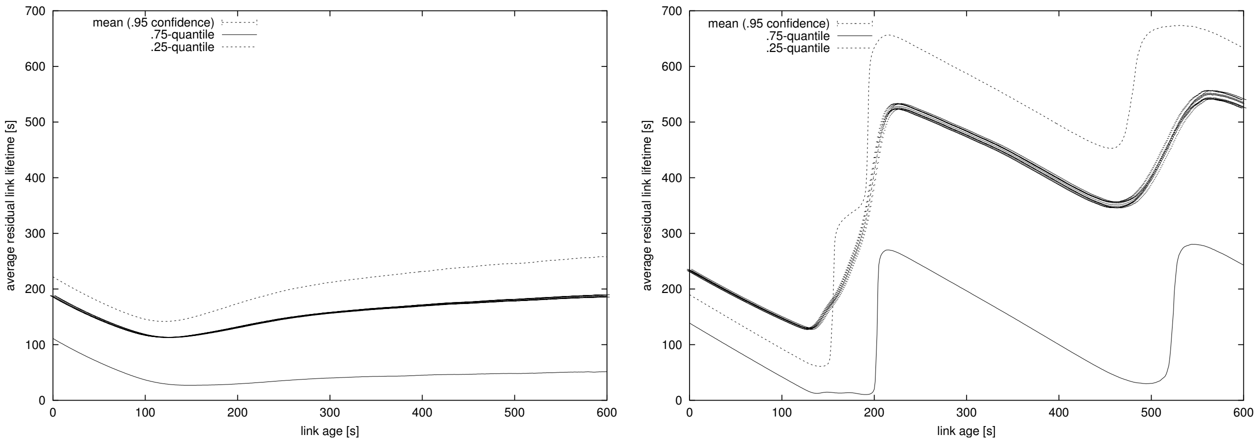

(a) Random Waypoint

(b) Manhattan Grid, 3x3 blocks

Figure 1. Dependency of a link's residual lifetime on its current age

served in figure 1. In scenarios where only a small portion of the links grows beyond this first period of a link's lifetime, it is straightforward to select the youngest link. Note that this is the case in many commonly used mobility patterns. This strategy does not perform well in rather static scenarios that show a large fraction of old links with a large expected residual lifetime, as e.g. Random Waypoint scenarios with long pause times.

Select the link with maximum expected residual lifetime. The expected residual lifetime of a link may be calculated from the collected statistical data as:

$$R(a) = \frac{\sum_{t=a}^{N} t d[t]}{\sum_{t=a}^{N} d[t]} - a \tag{1}$$

This metric promises to be more adaptive than the previously presented metrics. However, figure 1 shows that oftentimes the large average value results from only a small fraction of links having a very large residual lifetime whereas usually a large fraction of the links breaks much earlier (note e.g. figure 1(b) at 500s). This implies that the probability of choosing a rather short-lived link is relatively high. This property is undesirable for a wide range of applications including realtime and interactive applications where the number of interruptions should be at a minimum.

Select the link with maximum "persistence probability". This metric directly aims at minimising the number of interruptions during a certain time span. If the connection time is known for an application, it is possible to select a link based on its probability to persist for this amount of time:

$$P_s(a) = \frac{\sum_{t=a+s}^{N} d[t]}{\sum_{t=a}^{N} d[t]}. \tag{2}$$

Usually, this connection time is not known in advance. It might e.g. be too long (the failure probability is high in general) or too short (the failure probability is very low in general) to be a sensible parameter choice. Consequently, the challenge is to choose this persistence time appropriately to achieve good performance in a wide range of scenarios and applications.

Select the link with the lowest failure probability. If a connection time as required for the persistence probability metric is not applicable, it is still possible to choose a link based on a more conservative approach than the expected residual lifetime by choosing the link with the largest residual lifetime α-quantile:

$$Q_\alpha(a) = \max\{s \mid P_s(a) \geq \alpha\}. \tag{3}$$

If α is chosen close to 1, this metric provides a fairly reliable estimation of the minimum time this link is likely to persist.

4. Path Stability

In principle, the link stability metrics introduced in section 3 may be extended to a path stability metric in what may appear to be arbitrary ways. However, an advantageous characteristic of a path stability metric is its cumulative computability, meaning that the cost of an n-hop path can be calculated from the cost of an $(n-1)$-hop subpath and the cost of the *n*th hop. The cumulative computability allows a path stability metric to be integrated into a distance vector routing protocol.

This section introduces several possibilities to define a path stability metric. Basically, these metrics split into two categories. The first is based on the quantitative estimation of a link's stability using some specific metric, and the second is based on the qualitative estimation of link stabilities, i.e. the categorisation into stable and instable links.

Minimise the Number of Instable Links. If links can be classified as *stable* and *instable* according to the applied link stability metric, an intuitive way of rating the stability of a path is to minimise the number of instable links along a path. If two paths with the same number of instable links exist, the shorter one should be selected. Metrics that support the classification of links into stable and instable categories are to rate a link stable if it is either younger or older than a predefined threshold.

It should be noted that this is a different approach than that of ABR, which defines the "association stability" of a path as the fraction of stable links along this path (links are rated stable if they are older than a certain threshold). This has the disadvantage that the absolute number of instable links is not considered. Therefore, a 5-hop path with 2 instable links will be rated more stable than a 2-hop path with 1 instable link.

Maximise the Expected Residual Lifetime. A path's expected residual lifetime could only be estimated by an instance that knows the age of every link along the path, i.e. it is not "cumulative". To make things worse, it would either have to be assumed that the distribution of link lifetimes is the same for all stations and all locations within the network, or the instance calculating the expected residual lifetime would have to know the distributions of all other devices along the path. Either way, this approach is impractical and would probably not work too well anyway for the same reason already given in the context of link stability.

Maximise the Persistence Probability. The persistence probability of a path may be estimated by multiplying the persistence probabilities of all links along the path. Although theoretically, the residual lifetimes of the links are certainly not independent of each other, this metric performs quite well because the dependency is rather marginal: The fact that a link is available for a certain time span is not of great significance even for a neighbouring link.

As with the persistence of links, the main question for this approach is how the predefined target time should be chosen. If it is too long, the benefit that may be observed for the target time might be outweighed by the loss for shorter times (cf. figure 2). If it is too short, many links may be estimated to have a persistence probability of 1, and this approach converges towards a shortest path metric. However, if the application requirements suggest such a short connection time, avoiding the few links with a smaller estimated persistence probability could be worth the effort.

Maximise a Residual Lifetime Quantile. The minimum residual lifetime a path will reach with a given probability can be calculated from the path's persistence probabilities for different time spans in the same way as described for links in section 3.2. In a strict sense, this is not a "cumulative" calculation, although the calculation of each single persistence probability is. In contrast to the calculation of the expected residual lifetime, the amount of information that needs to be transmitted is bounded, but one value would have to be transmitted for each array bucket, making this approach impractical for a larger array length.

Avoid Instable Links. Since it usually is a single link break that invalidates a route, another possible approach is to define the stability of a path to be the stability of the most instable link along this path. The stability of a link may e.g. be rated by an α-quantile or the average of the residual lifetimes estimated for the link's current age, or by its estimated persistence probability. To keep the hop count within sensible bounds, a shortest path should be preferred from all paths with equal stabilities.

5. Simulation Results

We have evaluated the stability concepts introduced in section 4 by extensive simulations for a wide range of mobility patterns using BonnMotion ([3]) which simulates the different path selection strategies assuming ideal radio conditions. All simulation results are presented together with the .95 confidence interval.

With high probability, link durations in a realistic environment will not conform to any of the scenarios we have simulated. However, as long as there is no measured data on link lifetimes in mobile ad hoc networks available, we are restricted to the use of data resulting from simulations. For this reason, we take great care to analyse scenarios with strongly varying characteristics, in order to capture different aspects of real-life scenarios. Also, the characteristics of real-life scenarios will surely vary between different scenarios, and also between different points in time within the same scenario. Therefore, it is important that stability estimation approaches are adaptive to different conditions.

5.1. Simulation Setup

The simulations were performed on the following mobility models. A thorough treatment of these models may be

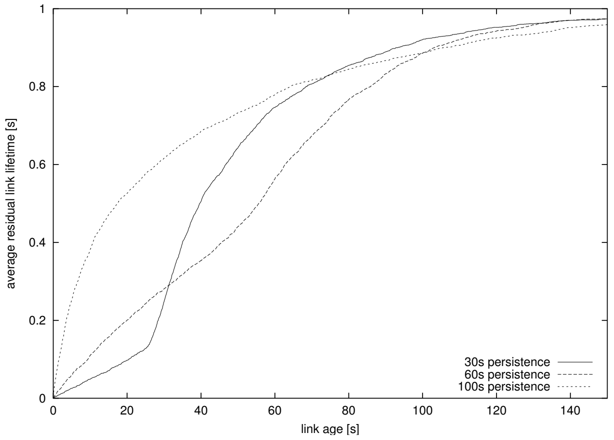

Figure 2. Distribution of route lifetimes for different persistence probabilities

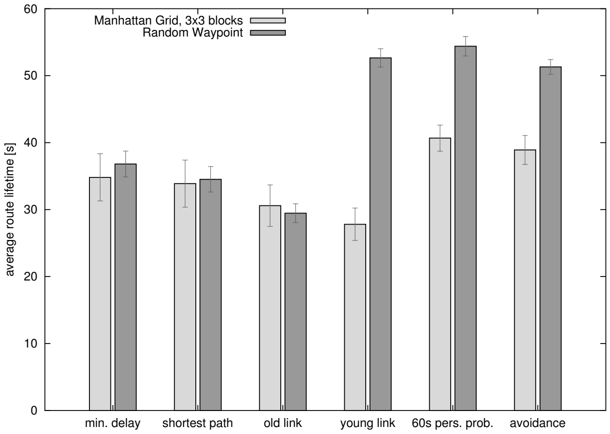

Figure 3. Average route lifetimes with different discovery strategies

found in [4].

Random Waypoint scenarios were simulated with 200, 500, and 800 nodes on an area of 10^6 square metres with transmission ranges of 140m and 200m. The nodes velocities were chosen from the interval 0.5m/s-1.5m/s. Maximal pause times ranged from no pause at all to 1800 seconds.

Gauss-Markov scenarios were simulated with 500 nodes on an area of 10^6 square metres with a transmission range of 140m. The update interval was switched between 10s and 50s. The standard deviation for the velocity was 0.1m/s with a maximal value of 1.5 m/s. The standard deviation for the angle was $\pi/4$ and $\pi/8$.

Manhattan Grid scenarios were simulated with 500 nodes on an area of 10^6 square metres. The update distance was set to 10m with an update probability of 0.2. The standard deviation for the velocity was set to 0.1m/s with a mean value of 1m/s and a minimum value of 0.5m/s. The turn probability was 0.5.

For the evaluation of the stability approaches we used the following representative sample of stability metrics:

1. *Minimise the number of too old links.* A link is deemed stable if its age is below 30 seconds.

2. *Minimise the number of too young links.* A link is deemed stable if its age is above 150 seconds.

3. *Maximise the persistence probability* for 60 seconds. Figure 2 shows a comparison of different values for the persistence duration in a Random Waypoint scenario with 500 nodes on a 200mx5000m area. Obviously, this metric is adaptable to different applications. For our simulations, we chose 60s as the target time for the persistence metrics. Many TCP connections as e.g. an email transfer or the download of a web page require only a few seconds, but e.g. VoIP sessions can be expected to have a duration in the order of magnitude of 60s.

4. *Avoid instable links* with the residual lifetime .95-quantile as the stability criterion.

As a reference, we used two commonly found path selection schemes, namely shortest path which is used by such routing protocols as DSR and many proactive protocols, and minimum delay such as found in AODV. The latter is simulated by assigning random delays to all links in the network and selecting the path with the shortest sum of link delays.

5.2. Different Views on Path Stability

5.2.1. Distribution of Route Lifetimes

Figure 3 compares the average lifetime of routes discovered with different strategies for a Manhattan Grid scenario with 3x3 blocks and a Random Waypoint scenario with 500 nodes and no pause time as described in section 5.1. Statistics on the link lifetimes of these scenarios are displayed in figure 1(b).

Whereas the average route lifetime may be improved by some of the stability approaches in this Random Waypoint scenario, the same is not true for this Manhattan Grid scenario. However, it has been pointed out previously that for many applications more important than the average route lifetime is a route's probability of being available for a certain time span.

To analyse this, figure 4 provides the route lifetime distributions corresponding to the average route lifetimes of the Manhattan Grid scenario shown in figure 3. Although the average route lifetimes hardly differ in this scenario, the risk of selecting routes that last shorter than 100s is clearly reduced with the two statistical metrics. Whereas the average lifetime of paths chosen according to the avoidance

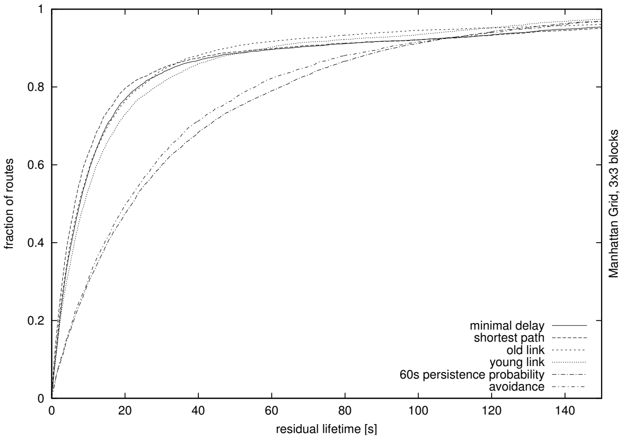

Figure 4. Distribution of route lifetimes with different discovery strategies

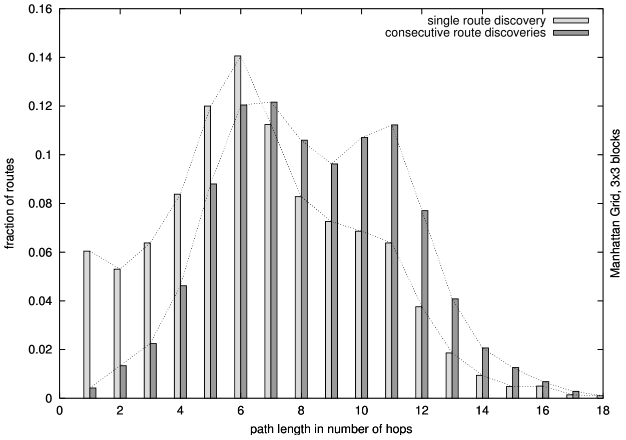

Figure 5. Distribution of path lengths for single and consecutive route discoveries

metric e.g. is very similar to the average path lifetime resulting from the minimum delay metric, the median route lifetime is more than doubled. For longer connections, this deficit is compensated. However, the route failure probability in these cases is already well above 90% rendering this advantage insignificant for most scenarios.

5.2.2. The Impact of Route Re-Discoveries

The importance of avoiding short-lived links becomes especially clear when taking a look at the number of route discoveries that have to be performed during a connection. This metric has a twofold significance: Reducing the number of route discoveries means reducing the routing overhead in the network, thereby increasing the network's capacity. This is especially important with routing protocols that employ flooding for the route discovery process. On the other hand, it is very important for realtime applications as e.g. voice conversations that the number of route discoveries is kept at a minimum because a route discovery often results in a disruption of the conversation.

The distributions of route lifetimes as measured in figure 4 provide only a vague hint on the number of route discoveries necessary for a certain connection duration. The actual number is larger than the value that may be estimated from this random sample of route lifetime measurements, because consecutive route discoveries are not independent random samples. Especially when a discovered route has a short lifetime, the topology of the network and in particular the distance between source and destination will be roughly the same for the successive route discovery. This means that the lifetime of the new route is likely to be in the same order of magnitude as that of the previous path.

This is confirmed by the fact that consecutively discovered routes exhibit an increased average path length when compared to independently discovered routes, visual in figure 5 which shows the distribution of path lengths for the minimum delay metric. In combination with the observation that the average lifetimes of paths of equal length remain constant in both experiments, this results in a shorter overall average route lifetime.

The more likely it is that routes with short lifetimes are established, the stronger this effect is. Consequently, it can be expected that the impact of a stability approach is increased in situations where multiple consecutive route discoveries are required. This can be observed in figure 6 which shows the number of route discoveries during a 180s connection for the different path selection strategies (in relation to the minimum delay strategy).

5.2.3. Path Length

In general, the average lifetime of a path decreases with its length. The reason for this is that longer paths have a higher probability of including a short lived link. Specifically, for the shortest path metric the probability of encountering the edge effect increases. The minimum delay metric tends towards short paths, but by circumventing the shortest path in many cases, it may avoid the selection of the worst links, i.e. those that cause the edge effect to occur. For the stability metrics, the probability of making a wrong selection increases with every additional hop. The avoidance metric e.g. has an error probability of $1 - \alpha$ for each link.

Although stable paths as identified by the different stability metrics are usually not shortest paths, the length of a route should not rise unnecessarily. Longer routes require

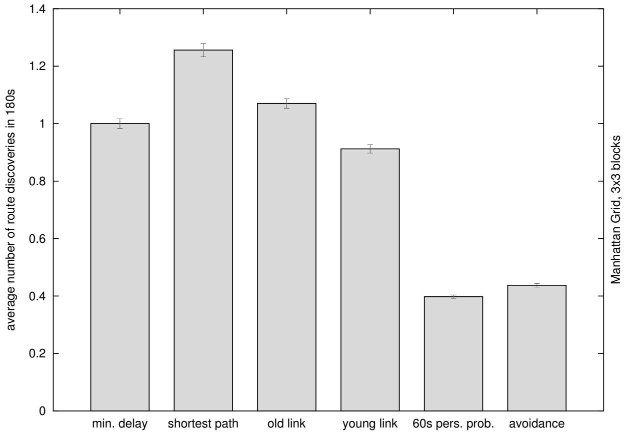

Figure 6. Route discoveries with different stability metrics (normalised to min. delay)

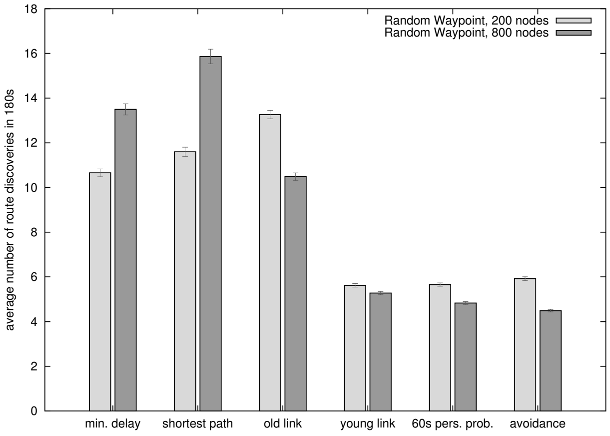

Figure 7. Influence of the node density on path stability

more packet transmissions and thereby reduce network capacity. This effect has to be traded off against the benefit of a possibly reduced routing overhead due to fewer route failures.

On average, the hop count of routes discovered with the avoidance metric was nearly twice the hop count of a shortest path in many scenarios because this metric is absolutely insensitive towards the path length. In contrast, the length of a stable path discovered with the persistence metric remains within sensible bounds in most scenarios because with every additional hop having a persistence probability of less than 1, the path's stability estimation is reduced. Both, young and old link metric, limit path length by selecting the shortest path with a minimum number of instable links.

5.3. The Influence of Mobility Model Parameters

This section analyses some details of the specific mobility scenarios and their influence on the performance of the different path stability metrics. First, the influence of general parameters such as the node density, transmission range and shape of the area will be described. Afterwards, specific properties of the different mobility models will be analysed.

Node Density. The statistical approaches benefit from an increased node density for two reasons (cf. figure 7). It has already been pointed out that the accuracy of the link lifetime statistics will be improved by an increased node density, because more samples may be recorded. Furthermore, a higher node density also leads to a higher network connectivity, which in turn has the effect of offering more alternate paths. It is not only this number of alternatives that counts, but with an increasing network connectivity, it also becomes less probable that there are specific bottleneck links that are common to all paths, which significantly limits the benefit of stability approaches over a random choice.

Both, shortest path and minimum delay metric suffer from increased node density due to the increased probability that a shortest path contains a link that is subject to the edge effect. Both the young and the old link metric benefit from increased node density in this scenario, because the probability of having links available that match the stability criterion increases.

Transmission Range. Node density and transmission range are closely related parameters. In fact, increasing the node density has the same effect as increasing the transmission range, if at the same time the velocity of the nodes is adjusted. In other words, increasing transmission ranges may be regarded as increasing node density and decreasing node velocity.

Consequently, higher transmission ranges lead to longer link lifetimes and thus to fewer rediscoveries in general (cf. figure 8). Therefore, the parameter choices for the statistical metrics might have to be adjusted (cf. section 3.1).

Shape of the Area. The shape of the area has an influence on several parameters. First of all, the distance between source and destination is likely to be higher for rectangular areas compared to quadratically shaped areas.

As already noted previously, a path becomes less stable with each additional link because it might be particularly short-living. For this reason, we see that the average number of route discoveries during a 180s connection rises dramatically for the minimum delay and the shortest path metric (cf. figure 9). The shortest path metric's property to force

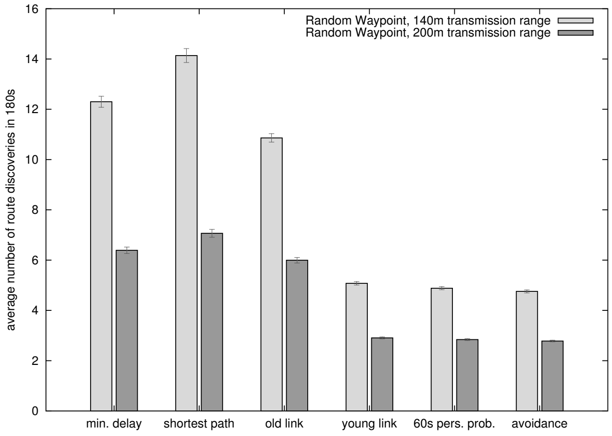

Figure 8. Influence of the transmission range on path stability

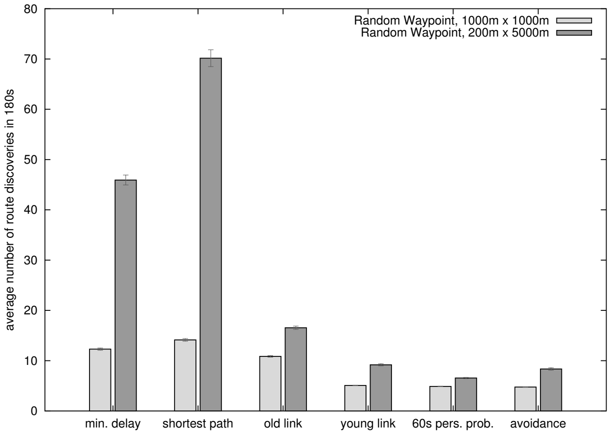

Figure 9. Influence of the area's shape on path stability

the selection of short-living long-distance links with an even higher probability lead to an increase of more than 50% in the number of route discoveries in comparison to the minimum delay metric (in comparison to a 15% overhead for the quadratic area).

The increase in the average number of route discoveries for the different stability metrics is rather limited, because their ability to avoid particularly bad links limits the impact of longer paths. Furthermore, the identification of stable links is more evident in rectangular ares due to the following reason: The more an area resembles a straight line, the more focused are the movements of the nodes, such that in extreme cases, nodes merely move in the same or in opposite directions. Thus, when it has become clear that two nodes move in the same direction (the corresponding link's age has passed the maximal lifetime of two nodes moving in the opposite direction), it is likely that they will stay together quite long.

Although in contrast, the young link metric is likely to establish links between nodes moving in opposite directions, the results in figure 9 suggest that it is able to avoid short-living links quite well. The reason is that even if the nodes move in opposite directions, it is quite unlikely that this is a short-living link since it has just recently been established. However, the young link metric is not able to catch up with the statistical approaches in rectangular areas where one dimension is significantly smaller than the transmission range, because it does not select the often superior old links which are likely to result from nodes moving in the same direction. This can also be observed in Manhattan Grid scenarios, cf. e.g. figure 6.

Pause Times. Pause times introduce more statics to a scenario leading to an increased fraction of old links and thus to a higher average lifetime of links. This benefits the old link metric directly and it outperforms all other metrics in the scenario shown in figure 10, which compares a Random Waypoint scenarios using 30min maximal pause time with one using no pause at all.

The young link metric cannot profit much from increased pause times, because this metric prefers links that have just come up and as such result from moving nodes.

The effect of the statistical metrics is limited by the observation time span which was 500s in our simulations. This explains why these metrics do not profit from increased pause times in these simulations (cf. section 3.1).

Gauss-Markov Scenarios. An increased update interval results in longer residual lifetimes for old links as well as a pronounced minimum expected residual lifetime, i.e. the extremal values are more distinct for larger update intervals. Because less updates are performed on the nodes' movement vectors during a certain time span, the probability decreases that two similar vectors diverge from each other within this time. This explains the longer residual lifetimes for old links. On the other hand, the probability also decreases that contrary movement vectors are assimilated before a link between two nodes moving in opposite directions breaks. Therefore, a higher number of links goes down about after the time it takes a node to traverse another node's coverage area with the typical relative speed, which explains the distinguished minimum residual lifetime. In general, Gauss-Markov scenarios with large update intervals resemble Random Waypoint scenarios (cf. [4]).

These two contrary effects lead to an only marginal ef-

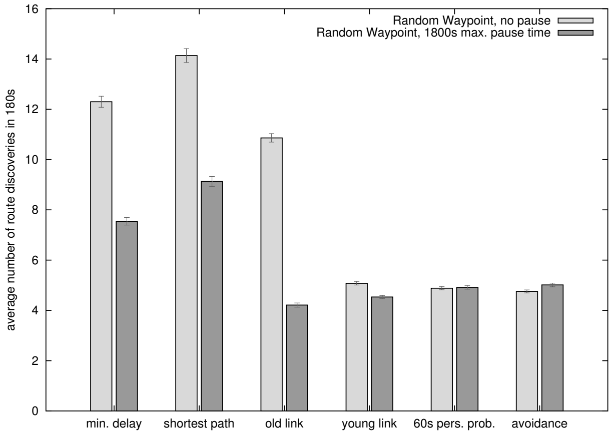

Figure 10. Influence of pause times on path stability

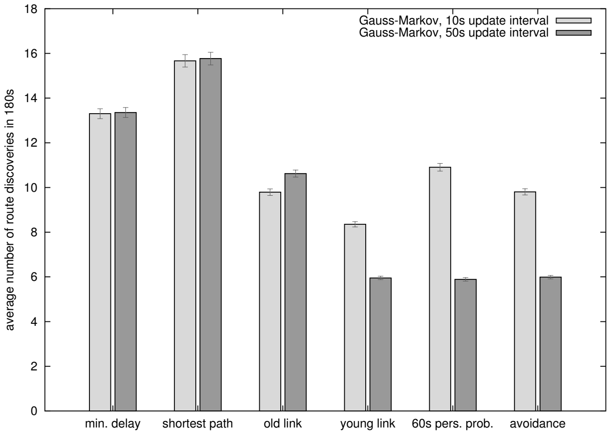

Figure 11. Influence of Gauss-Markov update interval on path stability

fect on the performance of the two reference metrics. On the other hand, the stability approaches benefit from the greater differences in residual lifetime. The old link metric cannot profit from increased update intervals due to the decreased number of old links.

An increased standard deviation of the velocity increases the disconnection probability of two nodes moving in the same direction, due to their increased non-zero relative speed. This decreases the residual lifetime of old links which has the greatest impact on the old link metric. But also the other stability approaches suffer from the aligned residual lifetimes such that the stability improvement in this scenario is rather small.

An increased standard deviation of the angle leads to bigger changes in the movement direction after every update. This more erratic movement makes the nodes rather stationary in the sense that the distances covered over a certain time span become smaller on average. Since this way it takes the nodes longer to cover certain distances, the expected residual lifetime of a short-distance link is higher. In effect, the residual lifetime of old links increases with increasing angle standard deviation, resulting in a benefit for the old link metric which will be largest for a standard deviation of π. On the other hand, the erratic movements result in a large fraction of just established links vanishing in the next update period. Thus, the young link metric may fall significantly short of the reference metrics for large standard deviations. The performance of the reference metrics is only marginally affected by the angle's standard deviation, because the increased residual lifetime of old links is compensated by a large fraction of links breaking early. The statistical approaches adapt to the changing conditions by preferring the older links.

Manhattan Grid Scenarios. In [4], it is outlined that in Manhattan Grid scenarios, the lifetime of a link depends to a large extent on the direction that the corresponding nodes choose at crossings. As a consequence, the size of the blocks has a major impact on link stability. Larger blocks result in a long connection period for nodes that take the same turn at a crossing (as long as the velocity standard deviation and the pause probability is small). This is directly reflected in the expected residual lifetime of links (cf. figure 1(b)). Due to the distinct lifetime expectations for links of different age, the statistical approaches are able to improve stability significantly compared to the reference metrics, independent of the number of blocks. The good performance of the young link metric for smaller block sizes is due to its ability to avoid the particularly bad choices. However, missing the promising of the old links limits this effect for larger block sizes. The old link metric is generally not able to identify the minima particularly well, especially for rather long paths. The particularly bad performance of the old link metric for 3x3 blocks is due to the large probability for links of age 150s to 200s of breaking almost immediately.

6. Conclusions & Further Work

In all scenarios that we evaluated, selecting a path based on its estimated probability to persist for a minimum amount of time proved to reduce the number of route discoveries significantly. In extreme cases, we could observe a reduction of the number of route discoveries to less than 15% compared to a minimum delay path and less than 10% compared to the shortest path.

In many scenarios, the number of route discoveries can also be reduced significantly by using a simple categorisa-

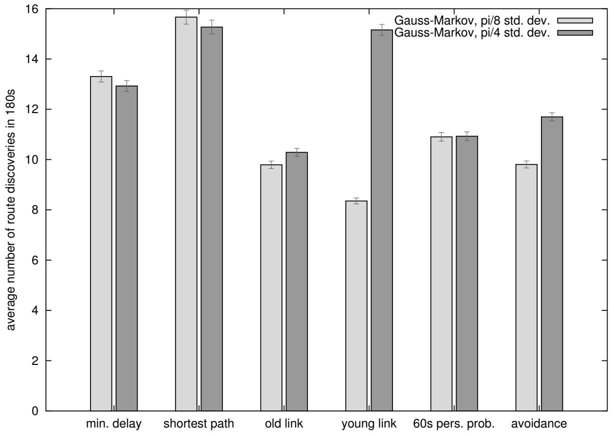

Figure 12. Influence of Gauss-Markov angle standard deviation on path stability

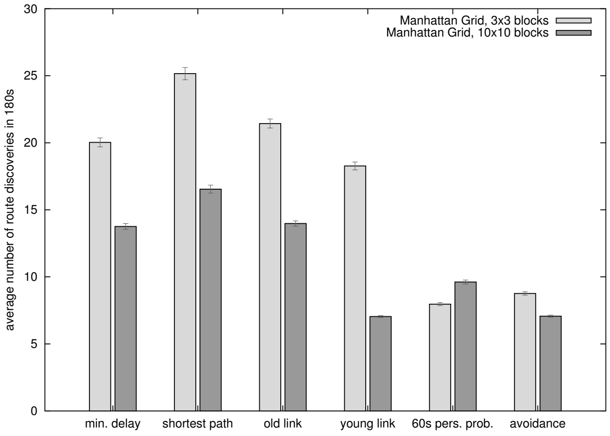

Figure 13. Influence of the Manhattan Grid block size on path stability

tion of links into stable and instable links, where the stability criterion is an age either above or below a certain threshold value. However, a good choice of this criterion depends on the scenario characteristics. Generally, it can be stated that preferring old links is beneficial in rather static scenarios, whereas preferring young links is beneficial in rather dynamic scenarios. Choosing the wrong criterion can lead to a severe increase in route failures even in comparison to randomly choosing links. In contrast, the statistical approaches are highly adaptive to a vast range of different scenarios.

Also, we have confirmed the previously observed instability of shortest paths. Especially for connections over longer distances, it is advantageous not to use a shortest path.

Further work is necessary to practically apply path stability metrics: Certainly, it is not practical to manually set the various parameters that control the stability metrics (and how observed link lifetimes are stored for a statistical approach). Therefore, heuristics on how to adapt to network characteristics automatically should be explored. Also, possibilities of integrating a stability metric into existing ad hoc routing protocols must be developed and evaluated.

Furthermore, the observation that stable paths tend to be longer than randomly chosen paths opens up an important question: How does the capacity gain resulting from the lower number of route discoveries trade off against the reduced capacity resulting from the increased path length?

References

[1] S. Agarwal, A. Ahuja, 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)*, pages 1697–1701.

[2] 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.

[3] M. Gerharz and C. de Waal. Bonnmotion: A mobility scenario generation and analysis tool. http://web.cs.uni-bonn.de/IV/BonnMotion/.

[4] M. Gerharz, C. de Waal, M. Frank, and P. Martini. Link stability in mobile wireless ad hoc networks. In *Proceedings of the 27th Annual IEEE Conference on Local Computer Networks (LCN'02)*, pages 30–39, Tampa, FL, November 2002.

[5] 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.

[6] 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.

[7] G. Lim, K. Shin, S. Lee, H. Yoon, and J. S. Ma. Link stability and route lifetime in ad-hhoc wireless networks. In *International Workshop on Ad Hoc Networking (IWAHN'02)*, August 2002.

[8] 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.

[9] R. Sivakumar, P. Sinha, and V. Bhargavan. CEDAR: a core-extraction distributed ad hoc routing algorithm. *IEEE Journal on Selected Areas in Communication*, 17(8), August 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.