

The stability of paths in a dynamic network

Kuipers, FA; Van Mieghem, PFA

DOI

<https://doi.org/10.1145/1095921.1095936>

Publication date

2005

Document Version

Accepted author manuscript

Published in

Proceedings of the 1st international conference on emerging networking experiments and technologies - CoNext'05

Citation (APA)

Kuipers, FA., & Van Mieghem, PFA. (2005). The stability of paths in a dynamic network. In s.n. (Ed.), *Proceedings of the 1st international conference on emerging networking experiments and technologies - CoNext'05* (pp. 105-114). Association for Computing Machinery (ACM).
<https://doi.org/10.1145/1095921.1095936>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

The Stability of Paths in a Dynamic Network

Fernando Kuipers
Delft University of Technology
P.O. Box 5031
Delft, The Netherlands
F.A.Kuipers@ewi.tudelft.nl

Huijuan Wang
Delft University of Technology
P.O. Box 5031
Delft, The Netherlands
bestwhj@gmail.com

Piet Van Mieghem
Delft University of Technology
P.O. Box 5031
Delft, The Netherlands
P.VanMieghem@ewi.tudelft.nl

ABSTRACT

Dynamic networks appear in several contexts: QoS routing faces the difficult problem of accurately and efficiently maintaining, distributing and updating network state information, and in wireless ad hoc networking, signal strength fluctuations complicate the choice of stable paths. In this paper we will focus on the stability of paths in a network with dynamically changing link weights. The level of path stability has a direct relation to the number of updates that are necessary to maintain an accurate view of the network state. If a small change in the network state does not affect the shortest path, then such a change need not be distributed throughout the network. We evaluate path stability by adding noise and observing the change in paths.

Categories and Subject Descriptors

C.2.2 [Computer Systems Organization]: Computer Communication Networks—*Network Protocols*

General Terms

Theory, Performance

Keywords

Network dynamics, Link-state update policy, shortest path, link weight perturbation, Quality of Service

1. INTRODUCTION

The need for a Quality of Service (QoS) aware Internet has been acknowledged by the research and business community for many years. Consider the simple example of telephony over Internet (VoIP). Two parties can communicate if the packetized voice is well within 200 ms mouth-to-ear delay. In addition, the communication requires a certain bound on the number of packets that are lost. Hence, as QoS measures we have time and loss. In fact, we may require also additional criteria such as bandwidth, monetary cost, and

delay variation (jitter). The exact QoS routing algorithm SAMCRA [17] can handle an arbitrarily large amount m of QoS measures (such as time, distance, cost, jitter, etc.) and even returns the optimal path in the m -dimensional space subject to a constraints vector. The latter means in the VoIP example that the path must guarantee an end-to-end delay within, 200 ms. Thus, for each QoS measure a QoS constraint can be included.

Despite the existence of QoS algorithms like SAMCRA [11], the Internet still lacks a (widely) working QoS architecture. Assuming the Internet infrastructure has enough resources to be able to provide QoS, then the problem is that of accurately and efficiently maintaining, distributing and updating the dynamic QoS link weights. Monitoring any change along the Internet is simply not possible and even not desirable, because not all changes are important. Further there is a topology range of interest: not all details of the entire global Internet are needed to determine a path P_{AB} from source A to destination B . If we first look at the time-scale in a network topology as illustrated in Figure 1, we distinguish between changes that occur (1) infrequently and (2) frequently. The first kind reflects topology changes due to failures and the joining/leaving of nodes. In the current Internet, only this kind of topology changes is considered. Its dynamics are relatively well understood. The key point is that the time between two “first kind” topology changes is long compared to the time needed to flood this information across the whole network. Thus, the topology databases on which routing relies, converge rapidly (with respect to the frequency of updates) to the new situation.

The second type of rapidly varying changes are typically related to the consumption of resources or to the traffic flowing through the network. The link weight coupling to state information seriously complicates the dynamics of flooding because the flooding convergence time T can be longer than the change rate Δ of some measure (such as available bandwidth).

Figure 1 illustrates how the bandwidth BW on a link may change as a function of time. In contrast to the first kind changes where $T \ll \Delta$, in the second kind changes, T can be of the same order as Δ . Apart from this, the second type changes necessitate the definition of a significant change that will trigger the process of flooding. In the first kind, every change was significant enough to start the flooding. The second kind significant change may be influenced by the flooding convergence time T and is, generally, strongly related to the traffic load in (a part of) the network. An optimal update strategy for the second type changes is highly desirable

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

CoNEXT'05, October 24–27, 2005, Toulouse, France.
Copyright 2005 ACM 1-59593-097-X/05/0010 ...\$5.00.

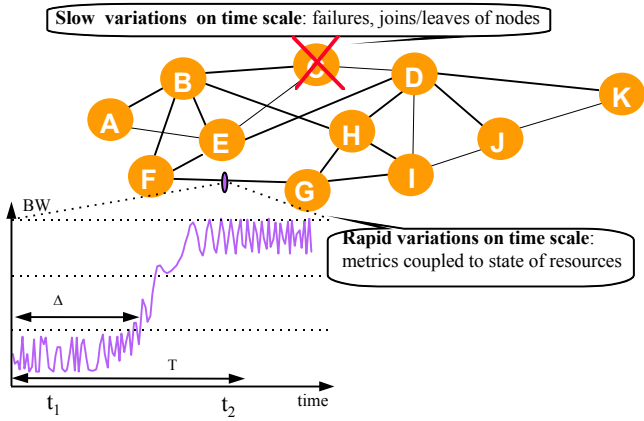


Figure 1: Network topology changes on different time scales. BW stands for bandwidth.

in future multimedia networks that are characterized by the broad variability in traffic profiles and QoS requirements.

The Internet is shown to possess a power-law like degree distribution [5], while ad hoc networks may vary from lattice structures to random graphs [8]. Moreover, the link weights in these ad hoc networks depend on a fluctuating signal strength. Since paths strongly depend on both link weight structure and graph properties, the network dynamics will depend on these factors, even to the extent that some control strategies successful in a certain class of graphs may not work properly in other graphs. We will examine different classes of graphs.

Throughout this paper we use the following notation to represent a network: $G(N, L)$ denotes a network topology, where $\{N\}$ is the set of N nodes and $\{L\}$ is the set of L links.

The goal of this paper is to evaluate the influence of dynamic link weights on the stability of paths. Our philosophy is that one must first understand the underlying concepts of dynamic networks, before developing link-state update protocols. The research presented in this paper will enhance our knowledge of how to properly choose the link-state update triggering thresholds or hold-down timers and may provide a first indication of what level of inaccuracy in the link weights is acceptable. The rest of the paper is organized as follows: in Section 2, we will provide an overview of related work. Section 3 provides our research framework. Section 4 evaluates path stability under constant noise, and Section 5 under uniformly distributed noise. In Section 6 path stability is bounded in case noise is only added to the path. In Section 7 we end with conclusions.

2. RELATED WORK

The current Internet only considers infrequent topology changes and consequently the study of these changes dates back to the early days of the ARPANET. One of the topics studied is that of end-to-end Internet path stability. Paxson [14] defined two types of stability, namely “prevalence,” meaning the likelihood that a particular route is encountered and “persistence,” the likelihood that a route remains unchanged over a long period of time. Based on measurements he found that *Internet paths are heavily dominated*

by a single prevalent route, but the time over which routes persist show wide variation, ranging from seconds to days.

In the context of QoS routing or ad hoc networking, the frequent changes in the network resources become a decisive factor and many problems emerge, e.g.: (a) how to predict the traffic load, (b) when to update the network with new information, (c) how to update the network, and (d) how to cope with inaccurate network state information. Below we briefly review the literature related to these problems.

2.1 Traffic prediction

Anjali *et al.* [1] proposed an algorithm to estimate the available bandwidth of a link in MPLS networks. They used a linear prediction model that is solved through Wiener-Hopf equations. Sang and Li [15] assessed the predictability of traffic by considering how far into the future a traffic rate process can be predicted with bounded error and what the minimum prediction error is over a specified prediction time interval. They used two models, namely the auto-regressive moving average and the Markov-modulated Poisson process and concluded that the applicability of traffic prediction is limited by the deteriorating prediction accuracy with increasing prediction interval. Jain and Dovrolis [10] targeted the end-to-end available bandwidth and stated that *the variability of the available bandwidth increases significantly as the utilization of the “low capacity” link increases, which makes a lightly loaded network have a more predictable and smooth throughput.* You and Chandra [21] and Basu *et al.* [3] analyzed Internet data measured at a campus and modeled this data using auto-regressive processes. Papagiannaki *et al.* [13] studied the evolution of IP backbone traffic at the larger time scale of hours and introduced a methodology to predict when and where link additions/upgrades have to take place in an IP backbone. They used mathematical tools to process historical information and extracted trends in traffic evolution at different time scales.

2.2 Network update triggering

Frequently updating the network, through the dissemination of link-state advertisements (LSA), can cause a significant overhead. Different link-state update policies have been proposed, which are reviewed in [12], [16], [2]. The link-state update policies can be classified as either periodic based (LSA at fixed intervals) or trigger-based (LSA at a certain event) and may use either a hold-down timer or the moving-average principle [12] to reduce the number of LSA.

2.3 Network update distribution

The current Internet disseminates its network state through the entire autonomous domain by using broadcast (flooding). In broadcast every router replicates the network state information onto all of its outgoing links. This method is too costly when the frequency of updates is expected to be high. To reduce the overhead in broadcasting, Garcia and Spohn [6] proposed the adaptive link-state protocol (ALP). A router in ALP disseminates link-state updates incrementally to its neighbors for only those links along paths (trees) used to reach destinations. Huang and McKinley [9] also proposed a tree-based protocol (T-LSR) that only constructs a single tree shared by every router for the dissemination of LSA and combines it with broadcast to make the protocol robust against node/link failures.

2.4 Inaccurate network state

The dynamics of link weights, prohibits an always accurate and up-to-date view of the network resources. The level of accuracy in state information depends on the choice of update strategy and can seriously impact the effectiveness (in terms of blocking) of path selection algorithms. A discussion of routing under inaccurate state information can be found in [7].

3. RESEARCH FRAMEWORK

In this section we provide a framework for the remaining part of this paper. The goal is to evaluate the stability of a path in a dynamic environment in a mathematical and simulative way. Our simulations consist of generating minimally 10^4 graphs (from a particular class of graphs) with link weights according to a specific distribution. We assign one weight w_l per link $l \in \{L\}$. We can consider this graph to be a snapshot in time of our dynamically changing network. In this graph, we compute the shortest path P between a source A and a destination B . Next, we perturb all the link weights in the graph by adding “link weight noise” $\tilde{w}_{l,\alpha}$ with strength α , such that the resulting link weight equals $w_l + \tilde{w}_{l,\alpha}, \forall l \in \{L\}$. This new graph represents a snapshot of the network at a later point in time and the noise represents the impact of the arrival/departure of flows over time on the resources. The level of noise is related to the period of time and the size and arrival rate of the flows. We recompute the shortest path P' between A and B in the perturbed graph and compare this path with the previously retrieved path P . We store the difference in path structure, i.e. how many links are different Δl , and the difference in path weight Δw . Three classes of graphs are investigated, namely the class of random graphs $G_p(N)$ [4], with link density p independent of N , the class of square two-dimensional lattices, and the class of Internet-like power-law graphs [5], with exponent $\tau = 2.4$ in the nodal degree distribution $\Pr[d = k] \sim k^{-\tau}$. The source A and destination B are chosen randomly.

By varying α ($\alpha > 0$, corresponding to the level of noise or perturbation), we are able to evaluate the perturbation threshold that causes P and P' to differ. The expected difference in weight between P and P' gives an indication of the size of the link-state update thresholds that should be used. When adding negative noise, a link weight may become negative. Negative or zero link weights are not considered realistic link weights. To assure positive link weights we truncate negative or zero link weights at a very small value ($\varepsilon = 10^{-5}$) near zero¹.

4. ADDING CONSTANT NOISE

The simplest scenario for adding noise is adding a constant noise to all links in the graph $G(N, L)$. More generally, consider a uniform distribution $U(b, c)$ to which we add a constant noise α . This results in the perturbed graph $G'(N, L)$ with link weights uniformly distributed in the range $(b + \alpha, c + \alpha]$. A scaling of this distribution, does not affect its properties and we can therefore use the distribution $U(a, 1)$, with $a = \frac{b+\alpha}{c+\alpha}$ representing a relative measure for noise.

The value for a has a large influence, as shown in Figure 2. We have considered (10^6) complete graphs with $N = 25$ nodes and simulated with $a = 0$ up to 0.5 in steps of 0.05. Since the source and destination are always connected by the

¹In our mathematical analysis we use 0.

direct link in the complete graph, the weight of the shortest path can never exceed 1.




Figure 2: Probability density function of the weight of the shortest path in the complete graph $G_{1.0}(25)$ for different a in the link weight distribution $U(a, 1)$.

When $a = \frac{1}{2}$, the link weights are uniformly distributed within $(\frac{1}{2}, 1]$ and the shortest path must be the direct link, because the weight of any path with $h > 1$ hops obeys $w(P_{h>1}) = \sum_{j=1}^h w(n_j \rightarrow n_{j+1}) \geq \sum_{j=1}^2 w(n_j \rightarrow n_{j+1}) > 1 \geq w(P_{h=1})$.

In fact, in a complete graph with uniformly distributed link weights specified by $U(a, 1)$ with $0 < a < \frac{1}{2}$, if the weight of the shortest path lies within $(a, 2a]$, the shortest path can only be the direct link. The probability for the direct link to be distributed within $(a, 2a]$ is $\frac{a}{1-a}$ and the corresponding probability density is

$$f_w(x) = \frac{\frac{a}{1-a}}{a} = \frac{1}{1-a}, \quad x \in (a, 2a]$$

which explains the uniform part in Figure 2.

When $\frac{1}{3} \leq a < \frac{1}{2}$, a similar reasoning can be applied to show that paths with more than two hops can never become the shortest path. When the weight of the direct link lies in $(2a, 1]$, paths with one hop and two hops compete to be the shortest path. In different classes of graphs, a similar analysis² can be made, as long as we can compute the probability that an h -hop path exists. Hence, given that the link weights are uniformly distributed in the range $(b, c]$, it is possible to predict the influence of adding constant noise.

5. ADDING UNIFORM NOISE

In this section we will evaluate the influence of adding uniformly distributed noise. We have assigned the link weights as follows: $w_l = U(0, 1)$ and $\tilde{w}_{l,\alpha} = \alpha U(-0.5, 0.5), \forall l \in \{L\}$. Initially all link weights were assigned a value chosen from a uniform distribution in the range $(0, 1]$. In the second scenario we added uniformly distributed noise in the range $\alpha[-0.5, 0.5]$. We have also simulated with a Gaussian distribution $N(\mu = 0, \sigma = \alpha)$ for the link weights/noise: the results were similar to those of the uniform distribution.

²For more mathematical and simulation results we refer to [20].

5.1 Mathematical analysis

We will provide upper bounds on the difference in path weights $\Delta w = w(P') - w(P)$, between the perturbed path P' and the unperturbed path P in any class of graphs.

By construction,

$$w(P) = \min_{P \subset \{P_{AB}\}} \left[\sum_{l \in P} w_l \right]$$

$$w(P'; \alpha) = \min_{P' \subset \{P_{AB}\}} \left[\sum_{l' \in P'} (w_{l'} + \tilde{w}_{l'; \alpha}) \right]$$

where w_l are the unperturbed link weights and where $\tilde{w}_{l'; \alpha}$ is the perturbation with strength $\alpha \geq 0$. Clearly, $w(P'; 0) = w(P)$ and the maximum possible perturbed weight is bounded by

$$w(P'; \alpha) \leq w(P) + \frac{\alpha}{2} h_P$$

where h_P is the hopcount of the shortest non-perturbed path. The other extreme in case of truncation is $w(P'; \alpha) = 0$, which occurs if there is a path from A to B with all link weights zero. Hence, denoting $\Delta w = w(P'; \alpha) - w(P)$,

$$-w(P) \leq \Delta w \leq \frac{\alpha}{2} h_P$$

In practice, the relevant range for α is limited to $w_{l'} + \tilde{w}_{l'; \alpha} \geq 0$, where truncation is not necessary. The probability that a perturbed link weight is smaller than zero is found from (3) in the Appendix, with $h = 1$ and $z = 0$ as

$$\Pr[w_{l'} + \tilde{w}_{l'; \alpha} \leq 0] = \begin{cases} \frac{\alpha}{8}, & \alpha \leq 2 \\ \frac{1}{2} - \frac{1}{2\alpha}, & \alpha > 2 \end{cases} \quad (1)$$

The probability never exceeds 50%.

Roughly, the probability to have a zero weight path is bounded from below by,

$$\Pr[w(P'; \alpha) = 0] \geq \prod_{e \in P'} \Pr[w_{l'} + \tilde{w}_{l'; \alpha} \leq 0]$$

$$\approx \begin{cases} \left(\frac{\alpha}{8}\right)^{\mathbb{E}[h]}, & \alpha \leq 2 \\ \left(\frac{1}{2} - \frac{1}{2\alpha}\right)^{\mathbb{E}[h]}, & \alpha > 2 \end{cases}$$

which is only significant for large α and a small expected hopcount $\mathbb{E}[h]$.

5.2 Simulations for Δw

We will first present our results for the difference in path weights $\Delta w = w(P'; \alpha) - w(P)$ for different classes of graphs. For each simulation 10^4 connected graphs were created. We have simulated with the random graphs $G_p(N)$ with link density $p = 0.2$ and $p = 0.01$, the power-law graphs with $\tau = 2.4$, and the square two-dimensional lattices. All graphs consisted of $N = 1000$ nodes, except for the square two-dimensional lattices which contained $N = 1024$ nodes.

5.2.1 The random graph

In Figure 3 we present $\mathbb{E}[\Delta w]$ as a function of the perturbation strength α .

For the class of random graphs with $p = 0.2$, $\mathbb{E}[\Delta w]$ decreases already for very small values of α , although this decrease is only small. For large α , $\mathbb{E}[\Delta w]$ saturates at $\mathbb{E}[\Delta w] \approx -\mathbb{E}[w(P)]$. For the class of random graphs with $p = 0.01$, $\mathbb{E}[\Delta w]$ starts decreasing at larger α than with $p =$




Figure 3: The expected difference in length ($\mathbb{E}[\Delta w] = \mathbb{E}[w(P') - w(P)]$) between P and P' for the class of random graphs ($N = 1000$), as a function of the perturbation strength α .

0.2 and the decrease is steeper. However, $\mathbb{E}[\Delta w]$ again saturates at $\mathbb{E}[\Delta w] \approx -\mathbb{E}[w(P)]$. This implies that $\mathbb{E}[w(P'; \alpha)] \approx 0$ for large α , irrespective of p , as was expected from our mathematical analysis.

In $G_p(N)$, a typical length is $\mathbb{E}[w(P)] \sim \frac{\ln N}{Np}$ [19]. The ratio of the two link densities considered here ($p = 0.2$ and $p = 0.01$), $\frac{0.2}{0.01} = 20$ almost precisely equals the ratio in $\mathbb{E}[\Delta w]$ at saturation, which equals $\frac{\mathbb{E}[\Delta w, p=0.01]}{\mathbb{E}[\Delta w, p=0.2]} = \frac{-0.730}{-0.037} = 19.5$.

The expected hopcount of the shortest path P in $G_p(N)$ with p fixed and uniformly (or exponentially) distributed link weights scales as $O(\log N)$ and the number of paths between source and destination is expected to be large, for N large. If α is high, there is a high probability that several links have truncated weights. Especially for $G_p(N)$ when p is fixed, this results in a high probability that the shortest path only consists of such zero-weight links. This behavior is verified in Figure 3. For smaller link density p , the number of (truncated) links is smaller and the expected hopcount larger. Hence a stronger perturbation is required before the saturation state $\mathbb{E}[\Delta w] \approx -\mathbb{E}[w(P)] \approx -\frac{\ln N}{Np}$ is reached. In fact, since saturation means that the source and destination nodes are connected by truncated links, we can consider a superconducting subgraph, which only contains the truncated links of the original graph. The link density of this subgraph is $p \Pr[w_{l'} + \tilde{w}_{l'; \alpha} \leq 0]$, where $\Pr[w_{l'} + \tilde{w}_{l'; \alpha} \leq 0]$ is given by (1). If α is large enough so that the subgraph is connected, then $\mathbb{E}[\Delta w]$ will saturate. The critical link density equals $p_c \approx \frac{\ln N}{N}$ and hence $p \Pr[w_{l'} + \tilde{w}_{l'; \alpha} \leq 0] \approx \frac{\ln N}{N}$ or $\alpha \approx \frac{8 \ln N}{pN}$ (we have used $\Pr[w_{l'} + \tilde{w}_{l'; \alpha} \leq 0] = \frac{\alpha}{8}$ in (1), because in the simulated cases saturation already occurs for $\alpha \leq 2$), as verified in Figure 3. The minimal α for saturation scales with $\frac{1}{p}$.

5.2.2 The square two-dimensional lattice

For the class of two-dimensional lattices, the expected minimum hopcount [18] equals $\mathbb{E}[h_{\min}] \approx \frac{2}{3} \sqrt{N}$ for large N . Although the expected hopcount of the shortest path is




Figure 4: The expected difference in length ($\mathbb{E}[\Delta w] = \mathbb{E}[w(P') - w(P)]$) between P and P' for the class of two-dimensional lattices ($N = 1024$), as a function of the perturbation strength α .

bounded from below by the minimum hopcount, we still expect a scaling of the hopcount as $O(\sqrt{N})$. For small α there is first a slight decrease ($\mathbb{E}[\Delta w] < 0$), after which a linear increase is manifested. Initially, on average 50% of the link weights are decreased and the remaining link weights increased. In this case, the expected weight of a randomly chosen path remains the same as in the unperturbed graph. However, the shortest path is likely to benefit more from the reduced link weights (by locally circumventing increased link weights), than that it is hampered by the increased link weights. This explains the small decrease for small α . For larger α , some decreased weights are truncated and cannot decrease further than $\varepsilon = 10^{-5}$, while increased weights have no upper bound. If α increases, the increased weights will have an increased influence. For large α approximately 50% of the in total $2(N - \sqrt{N})$ links will be truncated at $\varepsilon = 10^{-5}$. Hence, enough possibilities exist to circumvent a single high link weight. However, due to the relatively high expected hopcount not all high link weights may be circumvented. Since these high link weights increase with α , this explains the absence of saturation in Figure 4. The slope of the linear increase equals 0.054, which can be explained by considering a scaled version of the perturbed graph in case $\alpha \rightarrow \infty$. For large α , approximately 50% of the link weights in the perturbed graph are truncated and the remaining link weights are approximately³ uniformly distributed in the range $(0, \frac{\alpha}{2}]$. Dividing the link weights by α does not alter the shortest path and hence we choose for $\alpha \rightarrow \infty$ a scaled perturbed graph, with 50% of the link weights 0 and 50% of the link weights uniformly distributed in the range $(0, 0.5]$. Hence, the expected weight of the shortest path in the perturbed graph is approximated by $\mathbb{E}[P'] \approx \alpha \mathbb{E}[P'']$, where P'' equals the shortest path in the scaled ($\alpha \rightarrow \infty$) perturbed graph. Since both the shortest path P in the unperturbed graph and P'' are independent of α , we can approximate the slope in Figure 4 by $\mathbb{E}[P'']$. Simulations indicate that $\mathbb{E}[P''] = 0.078$ for $N = 1024$. The small devia-

³The actual distribution can be found via (3) in the Appendix.

tion from the slope in Figure 4, stems from the fact that we have assumed a uniform distribution for the non-truncated links in the (scaled) perturbed graph, which only holds for $\alpha \rightarrow \infty$.

5.2.3 The power-law graph

For the class of power-law graphs, we expect that the hopcount of the shortest path scales as $O(\log N)$, but that it is larger than the expected hopcount in $G_p(N)$, which also scales as $O(\log N)$. Contrary to the class of random graphs, we expect that the number of paths in the power-law graphs between a source and destination node is relatively small. In this case, there are less possibilities to circumvent the links with high weights, which is manifested for large α where $\mathbb{E}[\Delta w]$ increases linearly with α . This linear increase is best observed for smaller N in Figure 5.




Figure 5: The expected difference in length ($\mathbb{E}[\Delta w] = \mathbb{E}[w(P') - w(P)]$) between P and P' for the class of power-law graphs ($N = 100$), as a function of the perturbation strength α . The link weights were chosen according to $w_l = U(0, 1)$ and $\tilde{w}_{l,\alpha} = N(0, \alpha)$, $\forall l \in \{L\}$.

The slope for $\tau = 2.0$ is smaller than for $\tau = 2.4$, which is expected because a smaller τ leads to denser graphs and hence more possibilities to circumvent high link weights. The decrease for small α , as observed for the class of lattices, is also visible here.

We have also simulated with link weights with a granularity of 10, meaning that the link weights could only take one out of 10 values uniformly distributed in the range $[0, 1]$. In this case the paths P and P' remained the same ($\mathbb{E}[\Delta w] = 0$ and $\text{var}[\Delta w] = 0$) up to $\alpha = 2$, which confirms that *choosing a coarser granularity improves stability*.

5.3 Simulations for Δl

To better evaluate the difference between P and P' , we have also stored the number of different links Δl . More formally, Δl is the sum of the non-overlapping links of P and P' and therefore (between the same source and destination nodes) Δl cannot be 1 or 2. Figures 6-9 display the results for Δl in three different classes of graphs. We have simulated with $\alpha = 0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1.0, 2.0, 5.0$, but to avoid an overload of curves, we have plotted only four values of α .

5.3.1 The random graph




Figure 6: Probability that the number of different links Δl equals k in the class of random graphs $G_{0.2}(1000)$ for different values of α .




Figure 7: The expected difference in links $\mathbb{E}[\Delta l]$ as a function of α in the class of random graphs $G_{0.2}(1000)$.

In Figures 6 and 7 we can observe that for $G_{0.2}(1000)$, a small perturbation α can already trigger a large Δl . An interesting observation is that $\mathbb{E}[\Delta l]$ and $\text{var}[\Delta l]$ increase up to $\alpha = 0.05$, after which they decrease and stabilize for $\alpha > 0.1$. Note that $\alpha = 0.1$ is precisely the value at which $\mathbb{E}[\Delta w]$ for $p = 0.2$ in Figure 3 stabilized. In Figure 7, we have also plotted the expected hopcount $\mathbb{E}[h]$ of the shortest path in the perturbed graph and its variance $\text{var}[h]$. The mean $\mathbb{E}[h]$ and variance $\text{var}[h]$ are similar in shape to $\mathbb{E}[\Delta l]$ and $\text{var}[\Delta l]$. This is expected since the shortest path in the unperturbed graph is independent of α . In Figure 7, $\mathbb{E}[h]$ for $\alpha = 5$ equals 1.92, which corresponds to the expected hopcount of the shortest path in the random graph $G_{p'}(N)$ with constant link weights ($\mathbb{E}[h_{\min}] \simeq 2 - p'$ and $\text{var}[h_{\min}] \simeq p'(1 - p')$, [18]), as follows: if the number of truncated links in the

perturbed random graph $G_p(N)$ with a constant (i.e., truncated) link weight ε scales as $O(N^b)$, $b \geq 1$, then the shortest path in that graph behaves as if all links were constant. In that case the average number of truncated links per node > 1 . We can consider such a graph as “superconducting.” The number of truncated links equals $\frac{L\alpha}{8}$, where $L = \frac{pN(N-1)}{2}$ for $G_p(N)$, and scales as $O(\alpha N^2)$ for α large enough and p independent of N . We expect that

$$\mathbb{E}[h_{\min}] \simeq 2 - p \cdot \Pr[w_{l'} + \tilde{w}_{l';\alpha} \leq 0]$$

and with $p' = p \cdot \Pr[w_{l'} + \tilde{w}_{l';\alpha} \leq 0]$

$$\text{var}[h_{\min}] \simeq p'(1 - p')$$

which is verified in the simulations. For instance, for $\alpha = 5$ we have $\mathbb{E}[h_{\min}] \simeq 2 - 0.2(\frac{1}{2} - \frac{1}{10}) = 1.92$, which perfectly matches the simulation results. Only if $\alpha \lesssim \frac{8 \ln N}{pN}$ our scaling rule does not apply and we can consider the graph to be in “normal” or “transition” regime. The transition from normal to superconducting is observed in Figure 7.

5.3.2 The square two-dimensional lattice




Figure 8: Probability that the number of different links Δl equals k in the class of square two-dimensional lattices ($N = 1024$) for different values of α .

Due to the regular structure of the two-dimensional lattices, Δl can never be odd. We have removed those odd values from Figure 8. The high probabilities that $\Delta l = 0$ in Figure 8 suggest that the two-dimensional lattices are more stable than the random graphs. There is a higher probability that the paths P and P' are identical. Compared to the random graphs, the tail of Δl in the two-dimensional lattices is longer, due to the higher expected hopcount.

5.3.3 The power-law graph

The results for the power-law graphs are displayed in Figure 9. For small values of α the paths can be considered stable, but for high α we see a long tail and an irregular “distribution” function. These irregularities in the distribution suggest regularity or combinatorial confinement (in part of) the topology. The long tail suggests that sometimes




Figure 9: Probability that the number of different links Δl equals k in the class of power-law graphs ($\tau = 2.4$ and $N = 1000$) for different values of α .

the hopcount of the shortest path in a power-law graph can be very large.

6. BOUNDS ON PATH STABILITY

The simulations in the previous sections consisted of perturbing all links in the network, because over time, traffic is sent between different source-destination pairs and hence a large part of the network is affected. We have also examined how an increase in the use of a single shortest path can decrease its attractiveness as a shortest path. In practice, traffic between a source A and destination B is directed to the shortest path between A and B , and the other paths between A and B are not utilized. For instance, if a single flow is allocated on the shortest path, its resources will decrease, and its weight (e.g., delay) will increase. Depending on the size of the flow, the weight of the path may increase in such a way that it no longer is the shortest path anymore. In this section, we will evaluate the minimal total amount of noise that needs to be added to a shortest path in order for it to change. This amount of noise can be bounded. A lower bound is obtained from the difference in weight between the shortest path P_1 and the second shortest path P_2 . The minimal total amount of noise α (which is the sum of the noise on the links constituting the shortest path) must obey $\alpha \geq w(P_2) - w(P_1)$. Since P_1 and P_2 may share some links, adding this lower-bound noise may not change the shortest path. Only if P_1 and P_2 are link disjoint, α equals $w(P_2) - w(P_1)$. By computing the shortest path (P_l) that is link disjoint with P_1 , we can find an upper bound $\alpha \leq w(P_l) - w(P_1)$.

We have simulated with the class of random graphs $G_{0.2}(N)$, complete graphs, and lattices, all with independent uniformly distributed link weights in the range $(0, 1]$. The results for $G_{0.2}(N)$ are plotted in Figures 10-13. Figures 10 and 11 indicate that the difference in weight decreases when N increases, because $\mathbb{E}[w(P)] \approx \frac{\ln N}{N^p}$ [19]. The probability density function for $w(P_2) - w(P_1)$ in Figure 10 seems to follow an exponential distribution, unlike the probability density function for $w(P_l) - w(P_1)$ in Figure 11. However,




Figure 10: The probability density function (pdf) of the difference in weight between the shortest and second shortest path, for different network sizes, in the class of random graphs.




Figure 11: Probability density function of the difference in weight between the shortest path and the second shortest path that is link-disjoint with the first, in the class of random graphs.

for large N , they both tend to zero as observed in Figure 12. The larger and denser a network becomes, the more paths exist, and those paths will be more correlated (i.e., have a higher overlap) and their difference in weight will decrease. However, also $\mathbb{E}[w(P_1)]$ decreases (as $\frac{\ln N}{N^p}$). We have therefore computed the relative difference in weight $\frac{w(P_2) - w(P_1)}{w(P_1)}$ in Figure 13.

Figure 13 shows that the difference in weight decreases a little faster with N than the weight of the shortest path. The results for the complete graphs are similar to those for $G_{0.2}(N)$, especially for N large. The probability density function for the difference $w(P_2) - w(P_1)$ in the class of lattices is displayed in Figure 14.

Although the probability density function for $w(P_2) - w(P_1)$ in the class of lattices deviates a little from an exponential distribution, it provided a good approximation for the class of random graphs and hence might provide a good starting point for modelling the difference in weight




Figure 12: The expected difference in weight between the shortest path and the second shortest (link-disjoint) path, as function of the number of nodes N , in the class of random graphs ($p = 0.2$) and complete graphs ($p = 1$).




Figure 13: Probability density function for the relative difference $\frac{w(P_2) - w(P_1)}{w(P_1)}$, in the class of random graphs.

in any class of graphs. The probability density function for $w(P_1) - w(P_1)$ is displayed in Figure 15. The probability density function tends to a Gaussian distribution with increasing mean. To illustrate this, we have plotted in Figure 16 the expected difference in weight in the class of lattices.

As can be observed from Figure 16, the difference in weights $w(P_2) - w(P_1)$ decreases with N . However, the difference in weights $w(P_1) - w(P_1)$ increases, which makes the difference in weights between the link-disjoint paths a less suitable upper bound in sparse networks that have a large expected hopcount. However, since $w(P_2) - w(P_1)$ gives us a conservative estimate, this is the more important measure to predict the stability of a path in worst case.

Finally, in order to test the accuracy of our lower (and upper) bounds, we have simulated the minimal total amount of noise that needs to be added to a shortest path in order for it to change. We have added a constant noise α to the links on the shortest path and evaluated the minimal α that was required to change the shortest path. We have simulated with the class of random graphs $G_p(N)$, with $p = 0.2$




Figure 14: The probability density function (pdf) of the difference in weight between the shortest and second shortest path, for different network sizes in the class of lattices.




Figure 15: Probability density function of the difference in weight between the shortest path and the second shortest path that is link-disjoint with the first, in the class of lattices.

and exponentially distributed link weights with mean 1. N was chosen in the range [25, 400]. All simulations consisted of 10^4 iterations. In first instance, the shortest path was computed between a (randomly chosen) source and destination node. This path is referred to as SP_{old} . In the next step we added a constant noise to the links in SP_{old} , and gradually increased this noise until SP_{old} was replaced by a new shortest path SP_{new} , for which $w(SP_{new}) < w(SP_{old})$. The minimal noise α (per link), for which this holds was stored. Since $\mathbb{E}[w(SP)] \approx \frac{\ln N}{Np}$, we have plotted $Np\mathbb{E}[\alpha]$, $Np\mathbb{E}[w(SP_{old})]$, and $Np\mathbb{E}[w(SP_{new})]$ on a log-scale in Figure 17.

Indeed, as expected, $Np\mathbb{E}[w(SP_{old})]$ forms a linear line on the log-scale, and hence nicely obeys $\mathbb{E}[w(SP_{old})] \approx \frac{\ln N}{Np}$. Also SP_{new} approximates the $\frac{\ln N}{Np}$ law, which is probably due to the small expected hopcount ($\mathbb{E}[h(SP_{old})] \approx \ln N$) of SP_{old} and the relatively small values of α , which marginally contribute to the overall link weight distribution. In Figure 17, we have also plotted the lower bound




Figure 16: The expected difference in weight between the shortest path and the second shortest (link-disjoint) path, as function of the number of nodes N , in the class of lattices.

scaled per link $Np \left(\frac{w(P_2)}{h(P_1)} - \frac{w(P_1)}{h(P_1)} \right)$ and the upper bound per link $Np \left(\frac{w(P_2)}{h(P_1)} - \frac{w(P_1)}{h(P_1)} \right)$. The actual minimal noise $\mathbb{E}[\alpha]$ is nicely captured in between these two bounds. Moreover, these bounds improve with increasing N . Because the lower bound $w(P_2) - w(P_1)$ is most sensitive, it gives us a practical (worst-case) estimate of the stability of a path in the network.

7. CONCLUSIONS

In this paper we have taken a new and conceptual approach to lay the foundation for the development of link-state update policies. We have combined a mathematical analysis with simulations on thousands of graphs from different classes to be able to draw confident conclusions.

Generally, a small perturbation does not lead to a large difference in path weight. If the difference in path weight is negligible, then it is not necessary to update the network with superfluous information.

In practice, the link weights will have a finite granularity instead of the real values used for the simulations. We expect that a larger granularity will improve the stability of paths and consequently the predictability of network state.

Our results for the difference in links Δl displayed more volatility as a function of the perturbation strength α than the difference in weights Δw . Fortunately, the latter parameter is a more important measure to set a threshold for updating the network state.

Finally, through a mathematical and simulative analysis, we have indicated that the minimal level of perturbation that is required to make a shortest path obsolete, can be bounded. Moreover, this level rapidly decreases to zero for $N \rightarrow \infty$.

Acknowledgments: We would like to thank Eguzki Astiz Lezaun for providing many of the simulation results in Section 5.

8. REFERENCES




Figure 17: Expected weight of the minimal constant noise α that changes the shortest path and the expected weights of the shortest path SP_{old} before the noise and the shortest path SP_{new} after noise.

- [1] T. Anjali, C. Scoglio, J.C. de Oliveira, L.C. Chen, I.F. Akyildiz, J.A. Smith, G. Uhl and A. Sciuto, "A New Path Selection Algorithm for MPLS Networks Based on Available Bandwidth Estimation," Proc. of QoS/ICQT 2002, LNCS 2511, Zurich, Switzerland, pp. 205-214, October 16-18, 2002.
- [2] G. Apostolopoulos, R. Guerin and S. Kamat, "Quality of Service Routing: A Performance Perspective," Proc. of SIGCOMM'98, Vancouver, BC, Canada, pp. 17-28, September 1998.
- [3] S. Basu, A. Mukherjee and S. Klivansky, "Time Series Models for Internet Traffic," Proc. of IEEE INFOCOM'96, pp. 611-620, 1996.
- [4] B. Bollobas, *Random Graphs*, Cambridge University Press, second edition, 2001.
- [5] M. Faloutsos, P. Faloutsos and C. Faloutsos, "On power-law relationships of the Internet topology," Proc. of ACM SIGCOMM'99, Cambridge, Massachusetts, pp. 251-262, 1999.
- [6] J.J. Garcia-Luna-Aceves and M. Spohn, "Scalable Link-State Internet Routing," Proc. of IEEE International Conference on Network Protocols (ICNP), Austin, Texas, USA, October 14-16, 1998.
- [7] R. Guerin and A. Orda, "QoS routing in networks with inaccurate information: theory and algorithms," IEEE/ACM Transactions on Networking, vol. 7, no. 3, pp. 350-364, June 1999.
- [8] R. Hekmat and P. Van Mieghem, "Degree Distribution and Hopcount in Wireless Ad-hoc Networks," Proc. of IEEE ICON'03, Sydney, Australia, pp. 603-609, Sept. 28 - Oct. 3, 2003.
- [9] Y. Huang and P.K. McKinley, "Tree-based link-state routing in the presence of routing information corruption," Computer Communications, vol. 26, pp. 691-699, 2003.
- [10] M. Jain and C. Dovrolis, "End-to-End Available Bandwidth: Measurement Methodology, Dynamics,

and Relation with TCP Throughput,” Proc. of SIGCOMM’02, Pittsburgh, Pennsylvania, USA, August 19-23, 2002.

- [11] F.A. Kuipers, T. Korkmaz, M. Krunz and P. Van Mieghem, “An Overview of Constraint-Based Path Selection Algorithms for QoS Routing,” IEEE Communications Magazine, vol. 40, no. 12, December 2002.
- [12] B. Lekovic and P. Van Mieghem, “Link State Update Policies for Quality of Service Routing,” Proc. of 8th IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT2001), Delft, The Netherlands, October 18, 2001.
- [13] K. Papagiannaki, N. Taft, Z-L Zhang and C. Diot, “Long-Term Forecasting of Internet Backbone Traffic: Observations and Initial Models,” Proc. of IEEE INFOCOM 2003, San Francisco, USA, March 30 - April 3, 2003.
- [14] V. Paxson, “End-to-End Routing Behavior in the Internet,” IEEE/ACM Transactions on Networking, vol. 5, no. 5, pp. 601-615, October 1997.
- [15] A. Sang and S-Q Li, “A predictability analysis of network traffic,” Computer Networks, vol. 39, pp. 329-345, 2002.
- [16] A. Shaikh, J. Rexford and K.G. Shin, “Evaluating the Impact of Stale Link State on Quality-of-Service Routing,” IEEE/ACM Transactions on Networking, vol. 9, no. 2, April 2001.
- [17] P. Van Mieghem and F.A. Kuipers, “Concepts of Exact Quality of Service Algorithms,” IEEE/ACM Transaction on Networking, vol. 12, no. 5, pp. 851-864, 2004.
- [18] P. Van Mieghem, G. Hooghiemstra and R. W. van der Hofstad, “Scaling Law for the Hopcount,” Delft University of Technology, report2000125, 2000, <http://www.nas.ewi.tudelft.nl>.
- [19] P. Van Mieghem, *Performance Analysis of Communications Systems and Networks*, Cambridge University Press, 2005.
- [20] Huijuan Wang, *Analysis of the Shortest Path Problem: Link Weight Structure, Observability and K Shortest Paths*, M.Sc. thesis, Delft University of Technology, March 2005.
- [21] C. You and K. Chandra, “Time Series Models for Internet Data Traffic,” Proc. of 24th conference on local computer networks, pp. 164-171, October 1999.

APPENDIX

A. APPROXIMATE CALCULUS

In this Appendix we provide an approximate calculus of the length of a path as function of the hopcount in a class of graphs with link weight distribution $w_l + \tilde{w}_{l,\alpha}$, where $w_l = U(0, 1)$ and $\tilde{w}_{l,\alpha} = \alpha U(-0.5, 0.5)$, $\forall l \in \{L\}$. The distribution function $F_w^{h*}(z)$ is the probability that a sum of h independent random variables each with cumulative distribution function F_w is at most z and is given by the h -fold convolution:

$$F_w^{h*}(z) = \int_0^z F_w^{(h-1)*}(z-y) f_w(y) dy, \quad h \geq 2,$$

and where $F_w^{h*} = F_w$.

For $F_x(z) = z 1_{0 \leq z \leq 1}$ corresponding to $U(0, 1)$, and using $(x+y)^n = \sum_{j=0}^n \binom{n}{j} x^{n-j} y^j$ and the inverse Laplace transform $\frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} \frac{e^{sa}}{s^{n+1}} ds = \frac{a^n}{n!} 1_{\text{Re}(a)>0}$, we find

$$F_x^{h*}(z) = \sum_{j=0}^{\lfloor z \rfloor} \frac{(-1)^j}{j!(h-j)!} (z-j)^h, \quad 0 \leq z \leq h \quad (2)$$

while the perturbed situation corresponding to $F_w(z) = U(0, 1) + \alpha U(-0.5, 0.5)$ is

$$F_w^{h*}(z) = \frac{1}{\alpha^h} \sum_{j=0}^h \frac{h!(-1)^j}{j!(h-j)!} \sum_{k=0}^h \frac{h!(-1)^k}{k!(h-k)!} \cdot \frac{(z-j + \frac{\alpha h}{2} - \alpha k)^{2h}}{(2h)!} 1_{z-j+\frac{\alpha h}{2}-\alpha k > 0} \quad (3)$$

Since the link weights are independent and also w_l and $\tilde{w}_{l,\alpha}$ are independent, the expected weight of an h -hop path P in the perturbed graph equals $\mathbb{E}[w(P)] = \sum_{i=1}^h \mathbb{E}[w_l + \tilde{w}_{l,\alpha}] = \frac{h}{2}$ and the variance $\text{var}[w(P)] = \sum_{i=1}^h \text{var}[w_l + \tilde{w}_{l,\alpha}] = \frac{h(1+\alpha^2)}{12}$. By the central limit theorem, provided h is large enough, we can approximate the distribution (3) with the Gaussian distribution $\Phi\left(\frac{z-\mu}{\sigma}\right)$, with $\mu = \mathbb{E}[w(P)] = \frac{h}{2}$ and $\sigma = \sqrt{\text{var}[w(P)]} = \sqrt{\frac{h(1+\alpha^2)}{12}}$.