

Routing Stability in Static Wireless Mesh Networks

Krishna Ramachandran, Irfan Sheriff, Elizabeth Belding, Kevin Almeroth

University of California, Santa Barbara

Abstract. Considerable research has focused on the design of routing protocols for wireless mesh networks. Yet, little is understood about the stability of routes in such networks. This understanding is important in the design of wireless routing protocols, and in network planning and management. In this paper, we present results from our measurement-based characterization of routing stability in two network deployments, the UCSB MeshNet and the MIT Roofnet. To conduct these case studies, we use detailed link quality information collected over several days from each of these networks. Using this information, we investigate routing stability in terms of route-level characteristics, such as prevalence, persistence and flapping. Our key findings are the following: wireless routes are weakly dominated by a single route; dominant routes are extremely short-lived due to excessive route flapping; and simple stabilization techniques, such as hysteresis thresholds, can provide a significant improvement in route persistence.

1 Introduction

Applications, such as ‘last-mile’ Internet delivery, public safety, and distributed sensing, are driving the deployment of large-scale multi-hop wireless networks, also known as *mesh networks*. Although wireless routers in such networks are typically stationary, routes in these networks are expected to be unstable. One reason is that wireless links vary widely in their qualities because of multi-path fading effects, external interference and weather conditions. Link quality fluctuations can lead to variations in the quality of mesh routes, which can result in route fluctuations. This type of instability is unique to wireless networks.

Current routing protocols are not intelligent enough to consider routing stability during the selection of routes. A majority of the routing protocols [6] [14] ignore the fact that a route initially discovered has become sub-optimal over time. Route rediscovery is typically triggered by only route breaks and route timeouts. This approach can be detrimental to network performance.

Other routing protocols [2][7] periodically re-evaluate the quality of a route. The evaluation periodicity depends on the rate at which routing protocol control messages are exchanged. This approach fails to adapt to route quality variations that occur at smaller time-scales. However, by always picking the best route available, the resulting routing instability can lead to routing pathologies, such as packet reordering [3], which can severely degrade network performance.

We require a routing protocol that provides the best tradeoff between performance adaptability and routing stability. A detailed investigation of routing stability can help us design such a routing protocol.

Another reason such an analysis is important is because routing stability impacts mesh network management. As an example, channel management schemes [15, 16] in

multi-radio mesh networks assign channels to frequency diversify routes in the mesh. If routes are expected to change, the mesh radios should also be re-assigned channels in order to ensure optimal network performance.

An understanding of routing stability can also help in network planning, such as router placement and radio configuration. For example, stability analysis may suggest that routes to certain regions in the coverage area fluctuate frequently. The reason could be either poor placement of routers or radio misconfiguration.

Although considerable research has focused on the design of routing protocols and routing metrics for wireless mesh networks, there exists no formal study of routing stability in such networks. This paper presents the first measurement-based characterization of routing stability in static wireless mesh networks. We perform our study by answering questions such as: (1) Is there a clear choice of an optimal route between a source-destination pair? (2) If not, how long do such routes persist before a route change (flap) occurs? (3) What benefit does a route flap provide? and (4) What measures can help reduce route flaps?

In order to perform our measurement-based characterization of routing stability, we analyze link-quality information collected over a period of 2-3 days from two mesh network deployments, the UCSB MeshNet*, and the MIT Roofnet**. The MeshNet is a 20-node multi-radio 802.11a/b network deployed indoors on five floors of a typical office building on the UCSB campus. The MIT Roofnet is a 22-node outdoor network spread over four square kilometers in Cambridge, MA.

Clearly, routing stability analysis is influenced by the routing protocol. In order to investigate routing stability independent of any particular routing protocol, we compute high-throughput routes between all pairs of nodes assuming global knowledge of the collected link qualities. Routes are computed greedily, on a per-minute basis in our analysis, using the Dijkstra algorithm with the Weighted Cumulative Expected Transmission Time (WCETT) [7] as the path selection metric. We use WCETT because it has been shown to discover high throughput paths [7]. We compute routes greedily because we want to establish an upper bound on route capacities deliverable by a mesh network. Using the maximum capacities, we seek to understand the tradeoffs with respect to route instability.

The major findings from our study are as follows:

- Mesh routes are weakly dominated by a single route. The median prevalence of the dominant routes on the MeshNet and Roofnet are 65% and 57% respectively.
- Dominant routes are short-lived because of an excessive number of route flaps, most of which last only one minute.
- In a large number of cases, a route flap provides marginal improvement in throughput. 50% of the route flaps on the MeshNet, and 27% on the Roofnet, provide less than a 10% throughput improvement.
- Avoidance of routes that either last only one minute or provide only 10% throughput improvement increases the lifetime of the dominant route up to five-fold on the MeshNet and up to four-fold on the Roofnet.

* <http://moment.cs.ucsb.edu/meshnet>

** <http://pdos.csail.mit.edu/roofnet>

Although the above findings are specific to the two networks we have analyzed, we believe that the trends observed are generally applicable. Some of the findings discussed in this paper are well-known. A major contribution of this paper is a quantitative characterization of the extent of instability.

2 Related Work

Many studies have analyzed routing stability for wireline networks. Paxson reported on routing loops, routing stability, and routing symmetry by analyzing route information collected using *traceroute* [17]. Paxson found that Internet paths are typically dominated by a single route, and that a majority of Internet routes persist for either days or weeks. Labovitz et al. investigated Internet routing stability by analyzing BGP routing messages collected at key vantage points in the Internet [13]. Govindan et al. studied the growth of the Internet from 1994 to 1995 and found that route availability had degraded with the Internet's growth [9]. More recently, considerable attention has been given to routing pathologies because of BGP configuration faults [8, 18].

In the domain of wireless networks, various routing protocols [2, 6, 14] have been proposed for multi-hop wireless networks. Although the discovery of routes has been extensively studied by these efforts, to the best of our knowledge, there exists no formal study of routing stability in such networks. Studies have investigated connectivity between source-destination pairs in mobile ad hoc networks in terms of the lifetime of routes [1]. However, in such networks, node mobility influences the route lifetime. Our focus is on static mesh networks where mobility has little bearing on routing stability. Instead, the stability is influenced by the network topology and variations in link quality.

3 Methodology

Our analysis of routing stability is based on link quality information collected from the UCSB MeshNet and the MIT Roofnet. We start this section by briefly describing the two deployments. We then discuss the technique used to collect link quality information, following which we present the route computation engine that uses the link qualities to compute routes. We end this section with a discussion of some shortcomings in our methodology.

3.1 Network Deployments

The UCSB MeshNet is a multi-radio 802.11a/b network consisting of 20 PC-nodes deployed indoors on five floors of a typical office building in the UCSB campus. Each node is equipped with two types of PCMCIA radios: a Winstrom Atheros-chipset 802.11a radio and a Senao Prism2-chipset 802.11b radio. Each type of radio operates on a band-specific common channel. For rate adaptation, the 802.11b and 802.11a radios use auto-rate feedback [10] and SampleRate [2] respectively. There are 802.11b access points deployed in the building, which operate on various 802.11b channels. There is no external interference in the 802.11a band.

The MIT Roofnet consists of 22-nodes spread over four square kilometers in Cambridge, MA. Each node is a PC equipped with a Prism2-chipset 802.11b radio and an omni-directional antenna that is either roof-mounted or projecting out of a window. All radios operate on the same 802.11b channel. The Roofnet nodes experience interference from other, non-Roofnet access points.

3.2 Link Quality Estimation

Link quality is measured using the Expected Transmission Time (ETT) metric [7], which estimates the total time to transmit a packet on a link. The ETT is calculated from a link's loss rate and its data rate. ETT is given by the equation: $[(packet\ size)/(d_1 * d_2 * bw)]$, where d_1 and d_2 are the link's delivery ratios in the forward and reverse directions, and bw is the average of the link data rate reported by the two end nodes on the link. *packet size* is assumed to be 1500 bytes.

In the case of the MeshNet, the link quality information was collected on three different days. The loss rate was calculated by having each node issue a broadcast probe of size 524 bytes every second on each of its radios. Each node records the number of probes received from each of its neighbors in a 10 second window. The ratio of the number of packets received to the number of packets sent (10) yields a link's delivery ratio. The link data rate is measured using packet pair probing [11]. Every 10 seconds, each node issues packet-pair unicast probes of size 134 bytes and 1134 bytes on each of its radios. The difference in transmission time of the packet pair, as measured by a neighbor, is piggybacked on packet pairs issued by that neighbor. Every 10 seconds, each node reports each of its link's delivery ratio and data rate to a central repository.

In the case of the Roofnet, link delivery ratios are available*** on a per-minute basis for each 802.11b data rate. Since bandwidth information is not available for ETT computation, we set the link's ETT to be the ETT at the lowest data rate. In order to compute link delivery ratios, every 3 seconds, each Roofnet node broadcasts a 1500 byte probe at each of the 802.11b data rates, and a 134 byte probe at 1 Mbps. The 1500 byte probe is used to estimate the delivery probability of a large data packet at each of 802.11b data rates, whereas the 134 byte probe is used to estimate the delivery probability of a 802.11b acknowledgment. We use link delivery ratios on the 12th and 13th of May 2004 in our analysis.

3.3 Route Computation

We compute routes between all source-destination pairs for each minute recorded in our two data sets using an implementation of the Dijkstra's shortest-path algorithm. The quality of a route is computed using the Weighted Cumulative Expected Transmission Time (WCETT) metric [7]. The WCETT of a route is an estimate of the time a packet will take to traverse that route. The estimate is computed by taking into account the data rates, reliabilities, and channel assignments of all links on the path. We set WCETT's channel diversification parameter to 0.5. This setting gives equal weight to a path's channel diversification and its packet delivery rate [7]. In the case of the Roofnet, all

*** <http://pdos.csail.mit.edu/roofnet>

radios operated on a common channel. Hence, channel diversification did not play a role in the route computation for the Roofnet. A total of 6,345 and 11,470 unique routes were observed for the MeshNet and the Roofnet, respectively.

3.4 Shortcomings

Some noteworthy shortcomings in our analysis methodology are worth considering. First, we do not explicitly account for the impact of network load and external networks on the link quality measurements. In the case of the UCSB MeshNet, there was no data traffic on the mesh during the collection period. We are unable to say for a fact that this was the case with the MIT Roofnet because the Roofnet was operational during the link quality monitoring. Both networks experienced interference on the 802.11b band. We believe that the outcome of our analysis does not change per se. However, with our current methodology, we are unable to quantify the extent of the impact of these factors on our results. We plan to address this shortcoming in our future work.

A second consideration is the relationship between routing stability and time-of-day patterns. Routing behavior is expected to be more stable during off-peak hours when external interference and the load on the network are typically low. Our current analysis does not differentiate routing behavior based on time-of-day patterns. We plan to investigate this effect in our future work.

Finally, the configuration of a radio, such as its transmission power, receive sensitivity, and carrier sense threshold, is likely to influence routing stability. A majority of current radios and their drivers do not permit fine-grained control of configuration settings. As a result, an empirical-based analysis of the impact of radio configuration on routing stability is challenging. Software-defined radios are likely to help address this limitation.

4 Stability Analysis

We use three stability metrics in our analysis. First, *prevalence* is the probability of observing a given route [17]. Second, *persistence* represents the duration for which a route lasts before a route change occurs [17]. Third, *route flap* refers to a change in route.

4.1 Route Prevalence and Persistence

For a given source-destination pair, we analyze its routing prevalence in terms of its dominant route. The dominant route is the route observed the most number of times. In order to compute p_d , the prevalence of the dominant route, we note n_p , the total number of times any route was available between the given pair as is observed in the set of routes computed using the technique described in Section 3.3; and k_p , the number of times the dominant route was observed in the same route set. The prevalence p_d is then given as $p_d = k_p/n_p$.

Figure 1 shows the cumulative distribution of the prevalence of the dominant route for all source-destination pairs in the MeshNet and Roofnet. We observe that the dominant routes in both networks have a wide distribution of prevalence values. The median




Fig. 1. Prevalence of the dominant route for all source-destination pairs.




Fig. 2. Persistence of the dominant routes between all source-destination pairs.




Fig. 3. Number of unique routes for all source-destination pairs.

prevalence on the MeshNet and Roofnet are 65% and 57%, respectively. This observation suggests that *routes in static mesh networks are weakly dominated by a single route*.

We next analyze the persistence of the dominant routes. In order to calculate the persistence of the dominant route, we record all the durations observed for each dominant route. The persistence of a dominant route is then computed as the average of all its recorded durations.

Figure 2 plots the cumulative distribution of the persistence values in minutes for the dominant routes. For better clarity, only persistence values in the range of 1-1200 minutes are depicted on the x-axis. We observe that the dominant routes for both networks have a wide distribution of persistence values. The median persistence value for the MeshNet is 9.6 minutes, and the corresponding value for the Roofnet is 3.2 minutes. This result suggests that *routes in static mesh networks are short-lived*.

Note that, in general, the prevalence and persistence of the dominant route in the MeshNet are higher than in the Roofnet. To investigate the reason, we examined the number of unique routes computed between all pairs of nodes in the two networks. Figure 3 shows the cumulative distribution of the number of unique routes for all source-destination pairs. For the median node pair, the MeshNet offers 7 unique routes while the Roofnet offers as many as 17 unique routes. In general, the number of unique routes available between node pairs in the Roofnet is much higher than in the MeshNet. Therefore, there exists a higher probability for a Roofnet node-pair to choose a route other than the dominant route, compared to a MeshNet node-pair. This reason could explain the lower prevalence and persistence values in the Roofnet compared to the MeshNet.

One plausible explanation for the higher number of available routes in the Roofnet lies in the difference in the design of the two networks. The Roofnet is an outdoor 802.11b network, whereas the MeshNet is an indoor 802.11a/b network. In spite of being a dual-radio mesh, we observed that the majority of routes in the MeshNet consisted of 802.11a links. This majority occurs because 802.11a offers significantly higher data rates as compared to 802.11b. Now, 802.11b has a greater range than 802.11a. 802.11a range is further limited in a non-line-of-sight indoor environment as is the case in the MeshNet. Consequently, the Roofnet nodes are better connected with one another than nodes in the Meshnet. This reason could explain why the number of routes available in the Roofnet is much higher than in the MeshNet.

A worthwhile consideration following from the above reasoning is the impact network planning has on routing stability. In the specific case of the MeshNet, network connectivity likely contributed to higher persistence and prevalence values compared to the Roofnet. As another case in point, Camp et al. found that node placement in their Houston urban mesh deployment influenced routing performance [4].

Our analysis of persistence and prevalence indicates that routes in wireless mesh networks are inherently unstable. As a result, one would expect route flaps to occur frequently in a mesh network. The next section investigates the utility of the route flaps by investigating the throughput improvement they offer, and their lifetimes.

4.2 Route Flapping

The methodology to analyze the impact of route flaps is as follows. Every route change between a source-destination pair from one instance of time to the next is recorded as a route flap. For each route flap, we noted the length of time, in minutes, the flap persists before the next flap is observed. Also, for each route flap, we computed the percentage throughput improvement offered by the new route over the old route. Assuming a 1500 byte packet, the throughput of a route can be computed by taking the ratio of packet size to the route's WCETT value.

Figures 4 and 5 plots the percentage throughput improvement offered by a route flap on the y-axis against the lifetime of the flap on the x-axis. Each point corresponds to a route flap. For better clarity, only flap lifetimes in the range 1 through 50 are depicted on the x-axis. Several observations can be made from this figure.

First, the figure shows a high concentration of short-lived route flaps. The long-lived flaps are smaller in number and likely correspond to the dominant routes. Figure 6 plots all the route flaps shown in Figures 4 and 5 as a cumulative distribution of their flap lifetimes. For both networks, over 60% of the route flaps last only a minute; 90% of the route flaps last less than five minutes. The high number of short-lived route flaps contribute to the instability of routing in the two networks, as is observed in our analysis in Section 4.1.

Second, even though a high concentration of short-lived route flaps exists, the throughput improvement offered by these flaps varies widely. For example, in both networks, the one minute route flaps offer throughput improvements as little as 0.001% and as high as 100,000%. The implication of our findings is that *opportunistic throughput maximization through route flaps can lead to significant instability in a mesh network. However, many short-lived routes do provide significant gains in throughput. This suggests*




Fig. 4. Throughput benefit of route flaps in MeshNet.




Fig. 5. Throughput benefit of route flaps in Roofnet.




Fig. 6. Flap lifetimes as a fraction of total routes.




Fig. 7. Percentage throughput improvement as a fraction of total routes.

a routing protocol that provides good stability may have to compromise on throughput gains.

A third observation is that a large number of route flaps provide only a marginal improvement in throughput. Figure 7 plots all the route flaps shown in Figures 4 and 5 as a cumulative distribution of the percentage throughput improvement they provide. 50% of the route flaps in the MeshNet and 27% of the route flaps in the Roofnet provide less than 10% throughput improvement. These route flaps vary in duration from 1 minute to 50 minutes. The implication of this result is that *a routing protocol that always flaps routes will likely achieve only minimal gains in a large number of instances.*

4.3 Can Routing Stability be Improved?

The previous observations suggest that route flapping can be dampened by selectively choosing an alternate route between a source-destination pair. For example, a routing protocol may choose to switch to an alternate route only when the route offers more than 10% throughput improvement over what is currently used. In the specific case of the UCSB MeshNet, such a dampening threshold has the potential to eliminate more than 50% of all route flaps. Another likely dampening metric could be to switch to an alternate route only when the alternative is consistently better than the current route




Fig. 8. Route stability with dampening in MeshNet.




Fig. 9. Route stability with dampening in Roofnet.

for a specified amount of time. For example, this period could be two minutes. In the specific case of the UCSB MeshNet, such a dampening strategy has the potential to eliminate more than 60% of all route flapping.

To investigate the routing stability improvements that can result by applying such dampening techniques, we use two dampening metrics. The first metric is a 10% throughput improvement threshold, i.e., an alternate route is chosen only if it provides better than 10% throughput improvement. The second dampening metric is an alternate route persistence value of two minutes, i.e., the alternate route is available for at least 2 minutes.

Figures 8 and 9 plots the results from our application of the dampening techniques. The graphs depict the persistence values of the dominant routes against the fraction of all dominant routes. In the case of the MeshNet, if we consider the median dominant route, the one minute dampening metric yields a 5-fold increase in persistence. The 10% threshold yields a 4.5-fold increase in persistence. In the case of the Roofnet, the 10% threshold yields a 4-fold increase in persistence whereas the one minute threshold yields a 3-fold increase.

The above results indicate that by *using low thresholds during route selection in a mesh network, the persistence of the dominant routes can be significantly increased, therefore leading to increased stability*. An increase in the persistence will reduce routing pathologies, such as packet reordering [3], but may lower end-to-end throughput. As future work, we plan to investigate the trade-offs between stability and throughput in more detail.

5 Conclusion

We present a measurement-based characterization of routing stability in two static wireless mesh networks. This is a first step towards understanding long term behavior of routes in mesh networks. Some next steps for our continued analysis include: the impact of traffic load and external interference, the correlation between daily and weekly patterns, and the impact of physical layer properties such as transmission power and receiver sensitivity. We believe that the insights gained from this paper can stimulate

more research in understanding mesh routing behavior, which in turn can help us design better routing protocols and network management tools.

References

1. S. Agarwal, A. Ahuja, J. Singh, and R. Shorey. Route-lifetime Assessment Based Routing Protocol for Mobile Ad-hoc Networks. In *IEEE ICC*, New Orleans, LA, June 2000.
2. J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and Evaluation of an Unplanned 802.11b Mesh Network. In *ACM MobiCom*, Cologne, Germany, August 2005.
3. E. Blanton and M. Allman. On Making TCP More Robust to Packet Reordering. *ACM Computer Communication Review*, 32(1):20–30, 2002.
4. J. Camp, J. Robinson, C. Steger, and E. Knightly. Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network. In *ACM/USENIX Mobisys*, Uppsala, Sweden, June 2006.
5. A. Khanna, J. Zinky. The Revised ARPANET Routing Metric. In *ACM SIGCOMM*, Austin, TX, September 1989.
6. T. Clausen and P. Jacquet. Optimized Link State Routing Protocol. Internet Engineering Task Force, RFC 3626, October 2003.
7. R. Draves, J. Padhye, and B. Zill. Routing in Multi-radio, Multi-hop Wireless Mesh Networks. In *ACM MobiCom*, Philadelphia, PA, September 2004.
8. N. Feamster and H. Balakrishnan. Detecting BGP Configuration Faults with Static Analysis. In *USENIX Networked Systems Design and Implementation*, Boston, MA, May 2005.
9. R. Govindan and A. Reddy. An Analysis of Internet Inter-Domain Topology and Route Stability. In *IEEE Infocom*, Washington, DC, 1997.
10. A. Kamerman and L. Monteban. WaveLAN 2: A High-performance Wireless LAN for the Unlicensed Band. In *Bell Labs Technical Journal*, Summer 1997.
11. S. Keshav. A Control-Theoretic Approach to Flow Control. In *ACM Sigcomm*, Zurich, Switzerland, September 1991.
12. C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed Internet Routing Convergence. In *ACM Sigcomm*, Stockholm, Sweden, August 2000.
13. C. Labovitz, G. Malan, and F. Jahanian. Internet Routing Instability. *IEEE Transactions on Networking*, 6(5):515–528, 1998.
14. C. Perkins, E. Belding-Royer, and S. Das. Ad Hoc On-Demand Distance Vector Routing. Internet Engineering Task Force (IETF), RFC 3561, July 2003.
15. K. Ramachandran, E. Belding-Royer, K. Almeroth, and M. Buddhikot. Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks. In *IEEE Infocom*, Barcelona, Spain, April 2006.
16. A. Raniwala and T. Chiueh. Architecture and Algorithms for an IEEE 802.11-based Multi-Channel Wireless Mesh Network. In *IEEE Infocom*, Miami, FL, March 2005.
17. V. Paxson. End-to-end Routing Behavior in the Internet. In *ACM Sigcomm*, Palo Alto, CA, August 1996.
18. K. Varadhan, R. Govindan, and D. Estrin. Persistent Route Oscillations in Inter-domain Routing. *Computer Networks*, 32(1):1–16, January 2000.