

Insights into the Routing Stability of a Multi-hop Wireless Testbed (Invited Paper)

Mehdi Bezahaf¹, Luigi Iannone², Marcelo Dias de Amorim¹, and Serge Fdida¹

¹ LIP6/CNRS – UPMC Univ Paris 06, France
`{bezahaf, amorim, sf}@npa.lip6.fr`

² Technische Universität Berlin – Deutsche Telekom Laboratories, Germany
`luigi@net.t-labs.tu-berlin.de`

Abstract. By nature, links in multi-hop wireless networks have an unpredictable behavior, which directly affects the stability of routes. In this paper, we investigate the stability of the network by addressing some interesting questions related to the presence of both dominant and sub-dominant routes between nodes, in a real deployment. We focus on the persistence of the dominant route and the first four sub-dominant routes. The persistence is computed as the percentage of time that a given route is used. We note that source-destination pairs mostly use the dominant route and two sub-dominant routes to communicate, but with a low persistence. We also investigate the number of hops crossed by these routes and their impact on the stability. It turns out that the larger the number of hops, the larger the number of sub-dominant routes. However, when exceeding four hops, the notion of dominance fades.

Keywords: wireless mesh networks, multi-hop wireless testbed, routing stability, link stability.

1 Introduction

Wireless multi-hop networks consist in a set of wireless nodes, which may or not be mobile, communicating through wireless links. In such networks, intermediate nodes relay packets in order to reach the final destination (potentially a gateway to the Internet). The ordered list of wireless nodes crossed by these packets is referred to as a *path* or *route*.

A myriad of routing protocols have been proposed to build paths that are robust enough to overcome the natural limitation of wireless routes [6,7,11]. Indeed, due to the specificities of wireless links (e.g., interference, collisions, and fading), paths have an intrinsic unstable behavior when compared to routes in wired networks. Despite the countless routing protocols proposed so far, relatively little research work has been done on how to take into account the notion of route stability in route selection algorithms. Some routing protocols such as SSA, and ABR [4,13,8] use models based on signal strength or pilot signal to estimate the link stability. Paul et al. proposed the “affinity” parameter, which

consists in predicting links' lifetime [9]. This parameter is based on the strength and the stability of relationships of a node and its neighbors. Inspired by this idea, Agarwal et al. proposed a routing algorithm that incorporates the link lifetime prediction on the basis of affinity appraisal [2]. Authors also considered the path length while choosing an optimal route for a TCP source. The only study that tries to shed some light on the stability behavior of wireless paths is the one of Ramachandran et al. [12]. They collected, through wireless beacons, information about link quality and then computed, on a per-minute basis and offline, all possible routes between sources and destinations using an implementation of the Dijkstra shortest-path algorithm.

In this paper, we propose to go deeper into the analysis of routing stability by basing our analysis on routing events in a real testbed. We focus on a multi-hop wireless mesh network composed of static backbone nodes [3]. We address the following questions: What is the proportion of time during which the dominant route is active? Is this dominant route persistent? Are there any sub-dominant routes? If yes, what about their dominance? How long does each one of them last? How does the number of hops affect this dominance? By providing answers to these questions, we come up with the following contributions:

- We propose a measurement-based study of routing stability in a real multi-hop wireless mesh network testbed deployment, running a homemade python-based implementation of the DSDV (Destination-Sequenced Distance Vector) routing protocol.
- We investigate routing stability in terms of dominance and persistence of dominant and sub-dominant routes, also evaluating the behavior of routes when the number of hops increases.
- We observe that source-destination pairs mostly use a dominant route and two different sub-dominant routes to communicate between them, but with a low persistence. Moreover, we also observe that with a larger number of hops, a larger number of sub-dominant routes appears. However, when exceeding four hops, the notion of dominance disappear.

The remainder of the paper is organized as follows. In Section 2, we describe our testbed deployment, introduce the measurement setup, and define the scenario of our experiments. In Section 3, we investigate the dominance of the first most used route between each source-destination pair. Then, in Section 4, we investigate the persistence of the rest of the most used routes. In Section 5, we evaluate the impact of hops on routes dominance. We conclude the paper in Section 6 by summarizing our main results.

2 Experimentation Setup

Our analysis of routing stability is based on routing information collected from a real testbed deployment (namely MeshDVnet). We start this section by briefly describing this testbed. We discuss then the topology of the network and the experimentation scenario. We finally explain how we extract route information after collecting all routing tables in a central database.

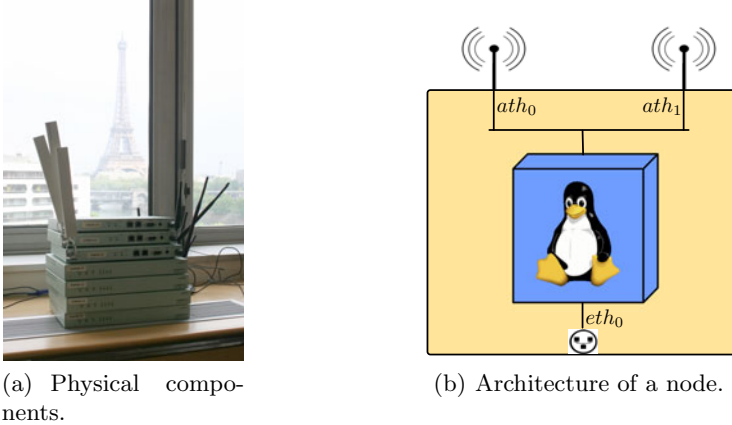


Fig. 1. MeshDVnetBox: our wireless mesh network testbed box

2.1 The MeshDVnet Testbed

Based on IPv6, MeshDVnet is a wireless mesh network testbed deployed at the computer science laboratory of the Université Pierre et Marie Curie [5]. This testbed offers traditional wireless LAN access connectivity to clients, while forwarding all generated traffic between clients and toward the gateway through the wireless routers in a multi-hop fashion. This testbed shares the wireless medium with the lab’s wireless network, the other companies’ networks in the same building, as well as neighborhood building’s networks.

MeshDVnet is logically decoupled into two subnetworks: one formed by the set of routers, which constitutes the *backbone*, and the other formed by the set of *clients*. Each router consists of a Linux kernel v2.6.19 (Slackware distribution) running on Soekris boards net4521 (AMD ElanSC520 133 Mhz) and net5501 (AMD Geode LX processor 500 Mhz) with two wireless cards (IEEE 802.11 a/b/g) using Madwifi driver v0.9.4 [1] (cf., Fig. 1). The first wireless card (ath_0 in Fig. 1(b)) is configured in ad hoc mode for backbone communications, while the second wireless card (ath_1 in Fig. 1(b)) is configured in access point (AP) to give access to clients. Furthermore, an Ethernet interface (eth_0 in Fig. 1(b)) is used for maintenance, supervision, and data collection (but not for traffic forwarding).

MeshDVnet backbone is built with thirteen routers placed around our lab’s building (in a single floor). We have an associated supervision web page that refreshes every two minutes to monitor if routers are up and running. The web page shows also the links between each router and their signal quality. It is publicly available (IPv6 only) at <http://www.infradio-jussieu.lip6.fr/supervision/supervision-mesh-kennedy.html>.

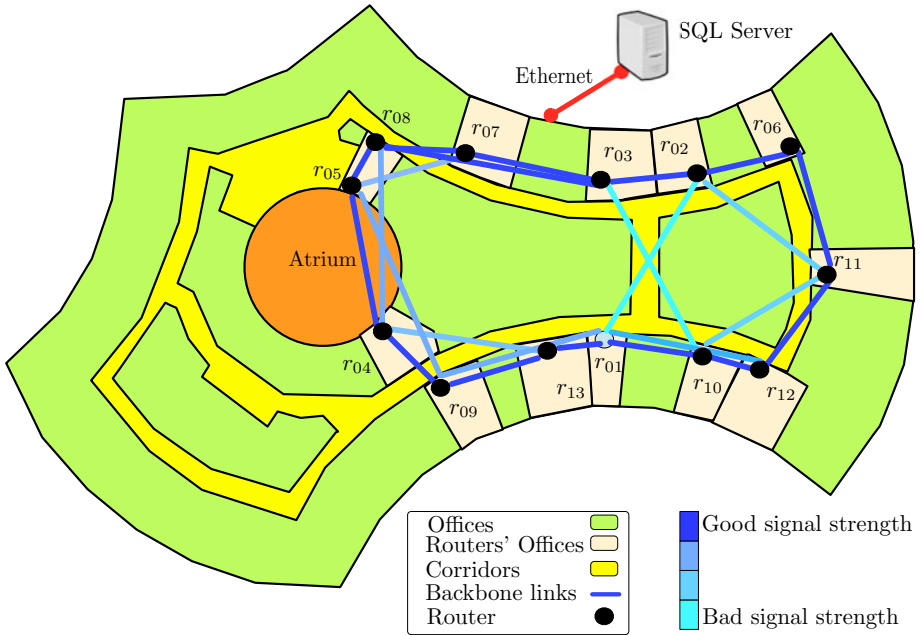


Fig. 2. Routing information collection infrastructure: thirteen routers communicating through wireless interface, and sending their routing information through the Ethernet interface to the central postgresSQL server

2.2 Measurement Setup

For the experimentation, we deployed thirteen routers in our offices as shown in Fig. 2. In order to have a realistic study with a real routing protocol and real routing message exchanges, we use the well-known proactive routing protocol DSDV (Destination-Sequenced Distance Vector) to build and maintain paths between routers [10]. Furthermore, in order to study the effect of the number of hops on the stability of links, we use hop-count as the metric for route selection. We chose DSDV as routing protocol in this first campaign of measurements for the simple reason that it is the protocol readily available on the MeshDVnet testbed.

It is out of the scope of this paper to compare or evaluate the performance of ad hoc routing protocols. Rather, we focus on illustrating the stability of wireless links using a real routing protocol through a measurement-based study. Clearly, the choice of the routing protocol and the routing metric can have, in some way, side effects on the stability of links; however, there is no escape to this situation, since this is part of any real measurement campaign. To the best of our knowledge, this is the first study on stability of wireless links based on measurement and taking into account a real routing protocol. As discussed in Section 6, an analysis of the impact of different routing protocols on the stability of wireless links is left for future work.

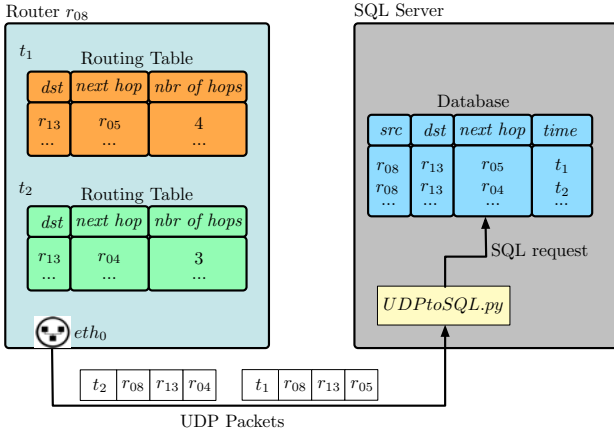


Fig. 3. Database update: router r_{08} detects at t_1 time a new route to reach router r_{13} . It sends a UDP packet to the SQL server to add this information to the database. At t_2 time it shifts to a new route, and updates the database with the new entry.

2.3 Routing Information Collection

Besides the thirteen routers, we deployed a central postgresSQL server to collect all information that we need for our experimentation (cf., Fig. 2). This information is collected through the Ethernet interface of each router (to avoid biasing the wireless medium with interfering traffic). We revisited the DSDV implementation in order to provide the routers with the capability to send their routing table to the database each time a routing table change occurs. With this infrastructure, we are able to centralize the routing information as seen by the routers itself, without any offline computation. The only post-processing we perform consists in the computation of the various statistics that we present in the second part of the paper.

During of the experimentation, which lasted a total of four consecutive days, at each change in a routing table the concerned router notifies in real time this change to the database by sending a UDP packet through its Ethernet interface.

In the postgresSQL server, we wrote a script that receives these UDP packets and translate them to a SQL insert (cf., Fig. 3). In this way, we obtain in real-time a database with all changes between all possible source-destination pairs during the experimentation. Finally, to retrieve all routing history, and to compute dominant route and sub-dominant routes, we execute a python script that iteratively queries the database.

3 Dominant Route

We define the *dominant route* for source-destination pair as the most used route in terms of cumulative duration. The other routes are referred to as *sub-dominant routes*.




Fig. 4. Percentage of time the dominant route was used between each source-destination pairs

One of the first questions that we address in this paper is: *what is the proportion of time of the dominant route?* In order to answer this question, we first extract from our database all possible routes between any pair of nodes in order to compute the percentage of time the dominant route is active. The result is shown in Fig. 4. For each router (x -axis) there are 12 bars (one for each other router in the topology). Note that this figure aims at showing a rough view of route activity; the details are given in the following. The main observation is that most of the source-destination pairs use the same route more than 40% of the time. Furthermore, there are some source-destination pairs that use the same route during all the experimentation (100% of time).

3.1 Route Persistence

A natural consequence of the previous observations concerning Fig. 4 is the question: *are the dominant routes used once for a long time or are they used several times for short periods?* The answer to this question is obtained by computing the *persistence* of the dominant route for each pair.

The persistence is computed as the average of the contiguous durations of the dominant route for each pair of nodes. In other words, it is the average time a dominant route is used before a different route is selected. The results are shown in Fig. 5 (recall that we have 13 nodes, which leads to $13 \times 12 = 156$ different dominant route IDs). We observe that 90 pairs of routers use in




Fig. 5. Persistence of the dominant route between each source-destination pair with a zoom between 0 and 5 minutes. Routes' IDs are sorted relatively to the persistence of the dominant route.

average their dominant route less than one contiguous minute (with very low confidence interval), which represents 57.7% of the cases, while 20% of them use their dominant route during one to five minutes. For the remaining 22.3% of the pairs, the duration ranges from five up to ninety-five minutes, with larger confidence intervals.

An important point is the evaluation the distribution of the route persistence. We selected four sample points corresponding to the persistence of dominant route of four source-destination pairs (respectively $r_{02} \rightarrow r_{06}$, $r_{12} \rightarrow r_{10}$, $r_{02} \rightarrow r_{10}$, $r_{13} \rightarrow r_{03}$), where each sample point represents a specific region of the curve in Fig. 5. In Fig. 6, we show the distribution of these four points. For each plot in this figure, the x -axis corresponds to the x^{th} time the dominant route is used. For instance, in the plot of the first sample, 70 in the x -axis corresponds to the 70th time the dominant route between r_{02} and r_{06} was used, and its persistence is given by the value in the y -axis value.

Note that most of the points are concentrated around the mean value, which indicates that the mean persistence is a reliable value. However, in samples 1 and 2, persistence of the dominant route varies from 1.04 minutes to 191.4 minutes, and from 0.31 minutes to 140.07 minutes, respectively the MIN and the MAX values, which causes the confidence interval (CI) to increase its value. We can also observe that in samples 3 and 4 the dominant route oscillates more than in samples 1 and 2 (cf., NP values in Fig. 6), but varies less in samples 3 and 4 than in samples 1 and 2 (cf., MAX/MIN values in Fig. 6).

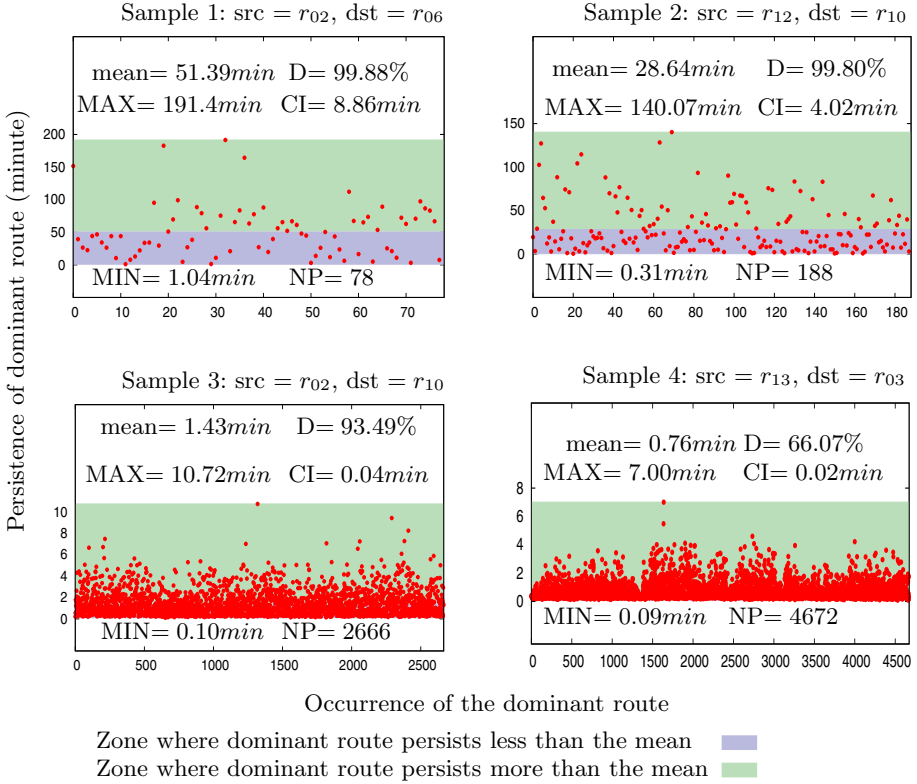


Fig. 6. Distribution of persistence of the dominant route between four source-destination pairs ($r_{02} \rightarrow r_{06}$, $r_{12} \rightarrow r_{10}$, $r_{02} \rightarrow r_{10}$, $r_{13} \rightarrow r_{03}$). Mean is the limit between the two zones. In the graphs, D represents the dominance ratio, CI is the confidence interval, and NP is the total number point of dominant route’s oscillation.

3.2 Route Oscillation

We evaluate now how much the dominant routes oscillate. We define the oscillation of a route as the number of times during the experimentation that this route appears between the same source-destination pair.

Fig. 7 shows the results sorted by their dominance. We observe that 90 dominant routes, which represents 57.7% of all dominant routes, oscillate in average 5,000 times (in four days), where their dominance varies between 14% and 85%. We also notice that a route can be dominant up to 57.7% even with a high oscillation; however, the higher the dominance, the higher the stability of the route (less oscillation).

Fig. 8 shows which kind of dominant route is the most persistent according to its dominance. The obtained result confirms the previous one. In fact, we see that




Fig. 7. Dominance vs. oscillation




Fig. 8. Dominance vs. persistence

routes with a high dominance level are the most persistent, which corresponds to the previous result (i.e., dominant routes with a high dominance ratio are more stable).

4 Sub-dominant Routes

The results presented in Section 3 clearly show that, during most of the time, source-destination pairs use the dominant route to communicate. *What happens in the rest of the time?* In order to know if there exist sub-dominant routes, we apply the previous analysis to the remaining routes.

We start our investigation with the dominance ratio of the four most prevalent sub-dominant routes. We sort all source-destination pairs according to the dominance ratio of their dominant route. Fig. 9 shows the percentage of activity of the main dominant route and the correspondent four main sub-dominant routes. As expected, the larger is the ratio of the dominant route, the lower the number sub-dominant routes used during the remaining time.

Fig. 10 shows the persistence of the sub-dominant routes. For the sake of clarity, we consider only the first 100 pairs for which the persistence of their dominant route is less than 2 minutes. Exceeding this value, the persistence of sub-dominant routes is negligible. We observe that the persistence of the first and second sub-dominant routes is equivalent to the persistence of the main dominant route, when the latter lasts less than half a minute. We also observe an interesting phenomenon. Some first and second sub-dominant routes are more




Fig. 9. Dominance of the five most used routes for each source-destination pair




Fig. 10. Persistence of the five most used routes for each source-destination pair

persistent than the dominant route in some cases, when the persistence of this latter is quite small. In fact, a dominant route that oscillates a lot can be less persistent than a less dominant route with long rare apparition. These cases correspond to source-destination pairs that use some rare routes to communicate (e.g., the corridor in the middle of the topology shown in Fig. 2).

5 Impact of the Number of Hops

In the following, we investigate the influence of the number of hops on the stability of routes. We consider alternative routes, persistence between pairs, and the possible use of sub-dominant routes.

We define physical hops between a source and a destination as the geographical distance between them in term of hops (network vision). Given that the routers in our testbed are static, the number of physical hops between a source and a destination is static as well. We also define logical hops as the number of hops used by the routing protocol. In fact, a destination can be in the range of the source (one physical hop), but given the links' perturbation, the source can choose another route with more logical hops.

To better understand the difference between physical and logical hops let us have a look to Fig. 11. For instance, the pair (r_{01}, r_{12}) in Fig. 11(a) is two physical hops pair, even if r_{01} uses the three logical hops route $(r_{01} \rightarrow r_{13} \rightarrow r_{10} \rightarrow r_{12})$ to reach r_{12} . Given the specific topology of our scenario, each router has two neighbors at X physical hops. For instance, the router r_{01} in Fig. 11(b) has $[r_{13}, r_{10}]$ as neighbors at one hop, $[r_{09}, r_{12}]$ at two hops, and $[r_{04}, r_{11}]$ at



Fig. 11. Physical and logical hops

three hops. We obtain than 26 pairs of routers at each hop level (13 routers \times 2 neighbors at each hop).

Using the notion of physical hops, we classify all source-destination pairs in six classes. Each class corresponds to the number of hops between the source and the destination (a class is called X -hop class if the destination is localized at X physical hops from the source). Note that in the rest of the paper, hops will refer to physical hops.

5.1 Available Routes

We define the number of available routes as the number of different routes used between a source-destination pair during the experiments.

Fig. 12 shows the evaluation of this number between routers at different level of hops. We observe that more we increase the number of hops, more the number of distinct routes increases, which is obvious. However, we notice that for each




Fig. 12. Evolution of number of distinct routes with the number of hops. Each curve is sorted according to the number of distinct routes.

hop, the number of distinct route changes depending on the source router. For example, for the six hops distance pairs, we observe some pairs with 90 ~ 100 distinct routes where other pairs with 240 ~ 270 distinct routes. Another interesting result is the one hop distance pairs. In fact, we expected to obtain a linear line with only one distinct route between each pair, but we observe that some pairs have used more than 30 routes. This result means that pairs use logical multi-hops even if they are only one physical hop away. In fact, if the medium at one physical hop (direct link) experiences bad link quality, routing protocol selects another possible route with more logical hops.

5.2 Persistence

In Section 3.1, we studied the persistence of dominant route in the whole network from a macroscopic point of view. We investigate now the persistence of the routes with regard to the number of hops separating the source-destination pair.

It is intuitive that shorter routes are more persistent. Fig. 13 shows the persistence of the dominant route for each pair of nodes at different distances. We observe that persistence of one-hop distance pair in the worst case is more than 4 minutes and can grow to more than 100 minutes. The persistence of four, five, and six hop distance pairs is practically the same and less than one minute. Dominant route persists more than one minute in 50% of cases for the three hops distance pairs to reach the limit of 10 minutes. At two hops distance pairs, dominant route are more persistent with a maximum of 42 minutes.




Fig. 13. Persistence of dominant route between each pair of source-destination at different distances

5.3 Sub-dominant Routes

Finally, we show the relationship between the number of sub-dominant routes and the number of hops between nodes. As shown in the previous section, we classify source-destination pairs according to their hop distance. Fig. 14 shows the percentage of dominance of the dominant and first four sub-dominant routes. We observe that the larger the number of hops, the smaller the prevalence of the dominant route.

In Section 4, we investigated the dominance of the main sub-dominant routes. We saw that some source-destination pairs use only the dominant route, while other pairs also select a few sub-dominant routes. More specifically, what we can observe from Fig. 14 is that nodes separated by one and two hops mainly use

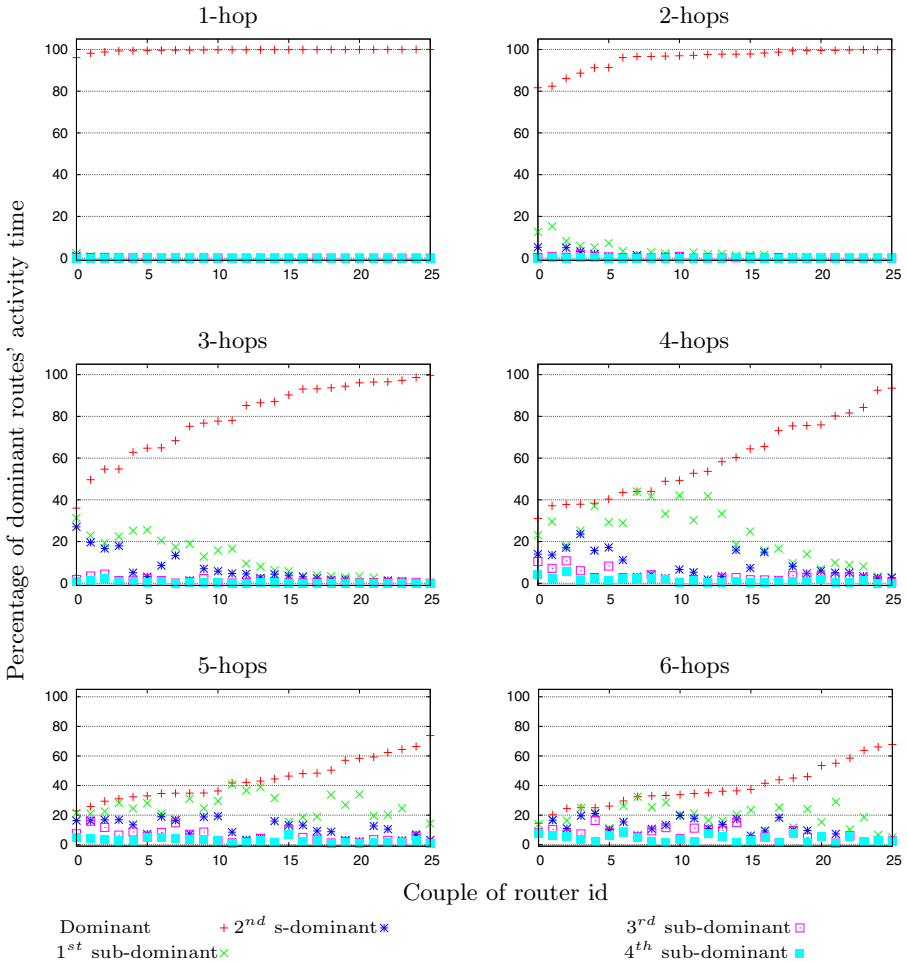


Fig. 14. Percentage of dominance of the five most used routes between pairs of routers. Pairs are classified according to their hop distance.

only the dominant route, whereas nodes separated by three and four hops also use two sub-dominant routes. Exceeding four hops, the dominant route starts losing its dominance and pairs tend to use different routes without notion of dominance.

6 Conclusion

In this paper, we investigated the stability of routes in a real wireless mesh network deployment. For our study, we used DSDV as routing protocol with hop count metric. Our findings can be summarized in three main observations.

Firstly, some pairs of routers use their dominant route with a persistence of almost 100% of the experimentation time, while other pairs show a high degree of instability. Secondly, some pairs alternate in almost perfect balance between their dominant route and their first two different sub-dominant routes. Thirdly, the larger the number of hops separating nodes, the larger the number of selected routes, losing the notion of dominant route.

The above observations help in understanding and quantify the stability behavior of links in multi-hop wireless environment. Our results could be useful in new routing protocols design and implementation, taking into account the number of hops and the dominance of routes. As a next step, we plan to study how much the stability of the links can be altered when we change the routing protocol and/or the metric.

References

1. Madwifi project – multiband atheros driver for wireless fidelity
2. Agarwal, S., Ahuja, A., Singh, J.P., Shorey, R.: Route-lifetime assessment based routing (rabr) protocol for mobile ad-hoc networks. In: IEEE International Conference on Communications, New Orleans, LA, USA (June 2000)
3. Akyildiz, I.F., Wang, X., Wang, W.: Wireless mesh networks: a survey. *Computer Networks* 47(4), 445–487 (2005)
4. Dube, R., Rais, C.D., Yeh Wang, K., Tripathi, S.K.: Signal stability based adaptive routing (ssa) for ad-hoc mobile networks. *IEEE Personal Communications* 4(1), 36–45 (1997)
5. Iannone, L., Fdida, S.: Meshdvt: A distance vector mobility-tolerant routing protocol for wireless mesh networks. In: IEEE ICPS Workshop on Multi-hop Ad hoc Networks: From Theory to Reality, Santorini, Greece (July 2005)
6. Jacquet, P., Mhlehthaler, P., Clausen, T., Laouti, A., Qayyum, A., Viennot, L.: Optimized link state routing protocol for ad hoc networks. In: IEEE International Multitopic Conference, Lahore, Pakistan (December 2001)
7. Johnson, D.B., Maltz, D.A.: Dynamic source routing in ad hoc wireless networks. In: Imielinski, T., Korth, H. (eds.) *Mobile Computing*, vol. 353, pp. 153–181 (1996)
8. Lim, G., Shin, K., Lee, S., Yoon, H., Ma, J.S.: Link stability and route lifetime in ad-hoc wireless networks. In: International Conference on Parallel Processing Workshops, Vancouver, BC, Canada (August 2002)

9. Paul, K., Bandyopadhyay, S., Mukherjee, A., Saha, D.: Communication aware mobile hosts in ad-hoc wireless network. In: IEEE International Conference on Personal Wireless Communications, Jaipur, India (February 1999)
10. Perkins, C., Bhagwat, P.: Highly dynamic destination-sequenced distance-vector routing (dsv) for mobile computers. In: ACM Sigcomm, London, UK (September 1994)
11. Perkins, C.E., Royer, E.M.: Ad-hoc on-demand distance vector routing. In: IEEE Wrokshop Mobile Computing systems and Applications, New Orleans, LA, USA (February 1999)
12. Ramachandran, K., Sheriff, I., Belding, E., Almeroth, K.: Routing stability in static wireless mesh networks. In: Passive and Active Network Measurement, Louvain-la-neuve, Belgium (April 2007)
13. Toh, C.-K.: Associativity-based routing for ad-hoc mobile networks. *Wireless Personal Communications* 4(2), 103–139 (1997)