Performance Study of the Better Approach to Mobile Adhoc Networking (B.A.T.M.A.N.) Protocol in the Context of Asymmetric Links

Alexander Klein and Lothar Braun and Fabian Oehlmann
Institute of Computer Science
Chair for Network Architectures and Services
Technical University Munich, Germany
Email: {klein,braun,oehlmann}@net.in.tum.de

***Abstract*—The performance of a mesh network is mainly affected by the utilized routing protocol and its capability to detect topology changes and to select stable and reliable routes. In real-world scenarios, a high fraction of links have asymmetric characteristics which has to be concerned by routing protocols in order to find optimized routes. The latest version of the Better Approach to Mobile Adhoc Networking protocol (B.A.T.M.A.N.) introduces its own routing metric, addressing the problem of integrating asymmetric links efficiently in the network topology. In this work, we compare version III and IV of the B.A.T.M.A.N protocol in terms of delivery ratio, average end-to-end hop count and generated overhead in two scenarios with a high fraction of asymmetric links to assess the quality of the newly introduced mechanisms on the performance of the protocol.**

***Keywords*-mesh; routing; asymmetric; evaluation; wireless;**

I. Introduction

Mesh networking is a well investigated field of research. Yet, the technological progress has so far not lead to an adoption of the underlying principle of relaying information over multiple hops in today's wireless networks. There are a lot of challenges, such as mobility, limited power, and unstable connections, that mesh networks have to cope with. A large number of different routing protocols have been developed to solve the problem of efficient packet forwarding under these demanding circumstances. The B.A.T.M.A.N. protocol [1] is a recent, ongoing development among them. One of its distinctive features is its inherent capability of handling asymmetric wireless links.

Studies have shown that a significant fraction of unreliable links can be found in wireless testbeds, which exhibit asymmetric characteristics [2].This phenomenon occurs particularly in low-power setups or in mobile networks with heterogeneous devices. Protocol developers are aware of such problems and therefore include mechanisms into their protocols, which aim at improving the protocol performance when asymmetric links are observed.

Our work aims at comparing version III and IV of the B.A.T.M.A.N protocol with respect to delivery ratio, average end-to-end hop count and generated overhead in scenarios with a high fraction of asymmetric links. We assess the quality of the newly introduced mechanisms, by closely examining two strongly differing wireless scenarios with asymmetric links. Additionally, a detailed perspective on their functioning is given. Our results show that B.A.T.M.A.N.'s first simplistic approach of finding reliable routes is tainted with a significant drawback in asymmetric network topologies. The coping strategy introduced in version IV introduces a routing metric that is highly capable of detecting paths with minimum packet loss even in the presence of asymmetric links.

This work is organized as follows. Section II discusses previous performance studies of B.A.T.M.A.N III and IV, and highlights where our work extends and differs from previous evaluations. A detailed description of their asymmetric link treatment is given in Section III. The results of our evaluation are presented in Section IV-A. Finally, our work is concluded in Section V.

II. Related Work

Many studies compare routing B.A.T.M.A.N against other routing protocols, using different implementations and versions. Moreover, they use different performance metrics for their evaluation, with throughput being the most popular one. Due to these differences in both protocol and evaluation metric, the presented papers are hardly comparable to each other. We put these studies into two groups: The first one evaluates version III while the other focuses on version IV of the B.A.T.M.A.N. protocol. To the best of our knowledge, no one has compared the performance of both versions in the context of asymmetric links, which is the focus of our work.

A. Version III Performance Studies

Ikeda et al. [4] performed several testbed evaluations of version III of the B.A.T.M.A.N. protocol. In their work, the effect of the routing protocol on throughput with TCP and UDP traffic is measured in a small setup. Varying characteristics are observed compared to the Optimized Link State Routing protocol (OLSR) [5]. The performance of Adhoc On Demand Distance Vector (AODV) protocol [7] and B.A.T.M.A.N. in mobile scenarios is discussed in [8].

978-1-4673-1239-4/12/$31.00 ©2012 IEEE

The authors conclude that B.A.T.M.A.N. shows a better performance compared to AODV with increasing mobility. Barolli et al. [6] did an experimental comparison in a static scenario in which OLSR with the ETX link quality extension and a differing window size was evaluated against B.A.T.M.A.N. in terms of throughput. In [4] B.A.T.M.A.N. is compared against OLSR with ETX as link quality estimation in a 7×7 grid testbed. B.A.T.M.A.N. achieves a higher throughput while maintaining more stable links than OLSR.

B. Version IV Performance Studies

In [12], the *batmand* routing daemon is tested against current implementations of OLSR and the Babel protocol. An indoor office mesh setup is used to demonstrate that Babel and B.A.T.M.A.N. show better overall results than OLSR. B.A.T.M.A.N. observes slightly better results concerning stability and packet delivery, while Babel offers a quicker reaction to topology changes. The same authors provide an evaluation of the Hybrid Wireless Mesh Protocol (HWMP) in [15]. Its performance levels are below those of OLSR and B.A.T.M.A.N..

In [16], the layer 2 and layer 3 implementation of B.A.T.M.A.N. is evaluated in an office testbed. The resulting performance is compared to that of OLSR and Babel. In this study, Babel considerably outperforms OLSR and B.A.T.M.A.N., while the latter are showing more or less equal performance levels.

III. Asymmetric Link Treatment

Several studies have shown that links in a wireless network can observe losses and are often asymmetric in their transmission quality depending on the direction [3].With regard to the underlying technology, there are many reasons for a link to be asymmetric: The surrounding environment can have great impact since radio wave propagation is affected by obstacles and different height levels of the devices. In addition, devices might suffer from different noise levels in areas with non-homogeneous node density. Technological reasons for asymmetric links can be found in varying power levels of transmissions of some devices, different hardware or battery levels. Moreover, antenna directions of individual devices can play an important role for radio wave propagation.

OLSRv2, Babel and B.A.T.M.A.N.'s evaluate paths between two devices using similar principles: The link quality is measured, using a protocol-specific metric, which is then propagated over the network. While every node in OLSRv2 performs the shortest path calculation on its own, Babel and B.A.T.M.A.N. version IV follow the distance vector paradigm: every node only knows about costs towards its direct neighbors. Contrary to this, B.A.T.M.A.N. version III does not perform any local link measurement at all, thus not having any special notion of dealing with asymmetric links. The rest of this section it will therefore focus on B.A.T.M.A.N. version IV, and does not consider version III.

The utilized algorithms of the protocol are able of finding the shortest path in a directed graph. The usage of asymmetric links strongly depends on the cost metric which determines the weights of the edges in the graph. Equal costs for every link lead to the utilization of the path with the shortest hop count. Such cost functions tend to use links with longer physical distance, which minimize the hops on a path and often has negative impact on the transmission reliability. Thus, OLSRv2 and Babel use the *expected transmission count* (ETX) [17] in their default implementations. Although the metric has symmetric characteristics, it takes asymmetry into consideration because it generates the same value for both directions. Links with asymmetric characteristics are merely punished and therefore less likely used for packet forwarding.

The *expected number of transmissions over forward links* (ETF) metric, as proposed in [2], similarly counts the expected number of transmissions in the desired direction. Therefore, the metric does not punish asymmetric links, but rather exploits those, which provide a good quality in at least one direction. The authors concluded that this metric outperforms ETX not only in networks without hop-to-hop acknowledgments, but also in those making use of them. This relies on the fact that acknowledgments are small in size and sent right after reception, when the channel is still likely to be idle.

B.A.T.M.A.N.'s TQ metric is somehow a mixture of hop count, ETX and ETF: Like ETF, it basically measures the probability of a successful transmission towards the next hop. Yet, it tries to avoid asymmetric links by applying the asymmetric penalty to punish links with a bad quality in the opposing direction akin to ETX. Additionally, a high hop penalty obfuscates the resulting TQ values, increasing the tendency to use the route with the shortest hop count.

A. Protocol Description

This section describes the protocol versions III and IV, as they have been implemented for simulation in the OPNET Modeler.

1) Version Comparison: The mechanisms of both protocol versions can be divided into three main tasks: Assertion of the usability of local links, flooding the network with originator messages, and providing a metric for estimating the quality of the path that is used when routing packets through the network. Both versions hereby work with a single packet format of a fixed size in order to provide a scalable solution. Similar to distance-vector routing, both versions of the protocol do not exchange and maintain global knowledge of the network topology.

The assessment of link usability in both versions depends on a bidirectional link check. Version III does this basically by assuring that OGMs are still traveling in both directions.

Version IV's approach is based on building statistics on local link quality. The number of lost echoes, causing a link to be considered unidirectional, is in version III explicitly preconfigured by a parameter. In version IV, this number is defined by the local window size.

Flooding of messages within the network is controlled by certain rules that specify when to rebroadcast an Originator Message (OGM). While both versions rebroadcast messages that are originated by and received from direct neighbors at all times, to maintain the local link considerations, they behave differently with OGMs of distant nodes. Version III employs a basic principle of forwarding non-duplicates, with the strict policy of only repeating messages that have been received via the currently best ranking neighbor. The path quality metric otherwise might become unreliable and may even cause routing loops. On version IV this problem is not inflicted since the path quality metric is changed to represent the best possible value. Version IV therefore floods the network more reliably, but on the other hand causes more overhead.

The path quality metric in version III is represented by the number of messages that reach a node via its direct neighbors. In contrast, version IV's OGMs propagate path qualities for every direct neighbor to determine their ranking. The major difference of both approaches lies in different assumptions regarding the conditions of the paths. Version III assumes that the path by which most packets reach a node, is the optimal path for forwarding. Version IV evaluates path qualities by merging information from every hop of the path.

B. Evolution of the Protocol

The evolution of version I to III is described in [9], while version IV is explained in [10]. Version V design considerations are provided on the open mesh website [11] by the protocol developers.

- Generation I served mainly as an experimental implementation for testing purposes. Link qualities are measured with respect to the number of OGMs, which have been received via one neighbor. It does not yet verify whether the links that build the network topology are unidirectional or bidirectional.

- Generation II introduces a classification of links in unidirectional and bidirectional as well as a mechanism to prevent nodes from propagating links which are not classified as bidirectional.

- Generation III introduces a strict policy of rebroadcasting OGMs, i.e. they are only rebroadcasted if they have been received via the best ranking neighbor. Furthermore, the detection of bidirectional links has been revised, so that it is time-independent. In addition, the algorithm supports multiple interfaces per node and allows modification of the TTL value for OGMs in order to limit the protocol overhead.

- Generation IV implies fundamental changes to the neighbor ranking mechanism. While former generations based their routing decision on the number of received OGMs, generation IV introduces the transmit quality metric. The metric reflects the probability of a successful transmission of a packet on a certain link. Moreover, packet aggregation may be applied to minimize the delay and optimize the efficiency.

- Generation V is planned to have two major modifications. First, the task of measuring link qualities between direct neighbors will no longer be done by using OGMs. A new packet format is introduced which will be used instead of OGMs for this purpose. However, OGMs are still transmitted to disseminate routing information. Secondly, a reactive mechanism is planned to detect link failures in the network. This is intended to increase the responsiveness to topology changes.

IV. Evaluation

A. Simulation Environment

For our evaluation of the different B.A.T.M.A.N. protocol versions, we employ a simulation environment which allows us to perform experiments with both versions, using the exact same properties for asymmetric links. In our experiments, we vary the fraction of asymmetric links that can be observed in the testbed, allowing us the evaluate the protocol mechanisms under different situations.

Our simulation scenarios have been implemented in OPNET and employ a filter module that artificially simulates the presence of asymmetric links. For each simulation run, a filter matrix is generated that describes link conditions between every node pair in both directions. Upon reception of a packet, the filter randomly eliminates packets, with respect to the successful transmission probabilities given in the filter matrix. Together with properties of transmitter modules, receiver modules, and antenna positions this abstracts the physical layer for our simulation. The OPNET simulator makes use of a free space propagation model in its calculations for the physical layer. Our filter module was built on top of this model, which results in the probability of a successful transmission that is below the one specified in the filter matrix. Furthermore, every filtered packet is still perceived by the MAC module as one that occupies the transmission medium. The data rate of the physical layer was set to 1 Mbit/s.

B. Grid Scenario

A grid scenario is used to simulate a random lossy asymmetric link setup to test the protocols' ability of finding reliable paths under highly asymmetric conditions.

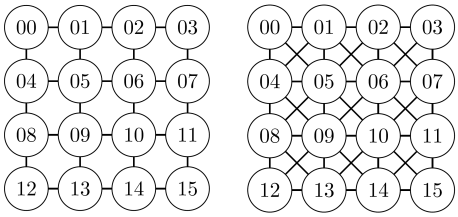

(a) Grid - Topology A (b) Grid - Topology B

Figure 1. Grid Scenario

1) Scenario Description: In total, 49 nodes have been placed in a static 7×7 grid. The distance between the nodes is varied in order to create network topologies with different node densities. In scenario A, the node distance was chosen to be 75 meters, whereas in scenario B the distance was chosen to be 60 meters. A transmission range of 100 meters for each node leads to two different topologies. As depicted in Figure 1, where we show the topologies for 16 participating nodes, the varying node density causes the middle nodes in case A to have 4 neighbors. In case B, the middle nodes have 8 neighbors each. Therefore, these topologies result in 84 links for topology A and 228 for topology B.

The purpose of this setup is to simulate the performance of the protocol while increasing the fraction of asymmetric links within the topology. The simulated grid setup results in a pairwise average hop count between any two nodes of 3.29 and 4.66, respectively. This setup does render the routing decision rather challenging since multiple paths with different quality in terms of reliability exists between any two nodes.

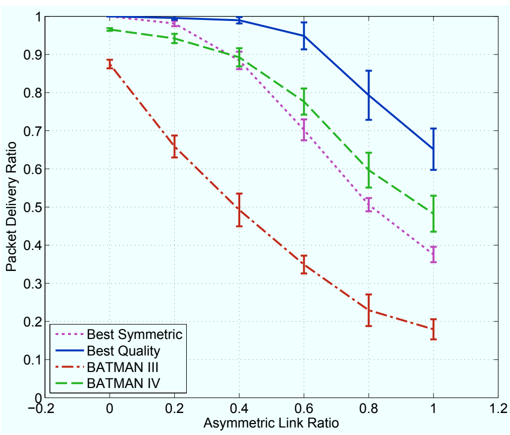

Figure 2. Grid, Topology A, Reliability

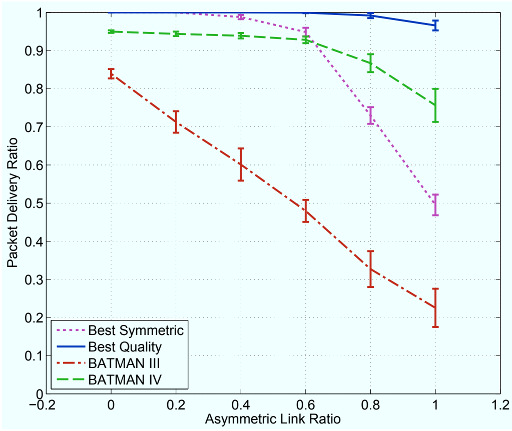

Figure 3. Grid, Topology B, Reliability

The simulation is carried out with a duration of 1100 seconds. The routing mechanism is started after 30 seconds and traffic is sent after 100 seconds. Therefore, the routing decision is made upon fully informed state of the links when the first data packets are to be sent. All nodes generate data packets with a constant packet size of 1024 bits. The packets are sent in constant time intervals of 1 second with small jitter of 1ms to avoid synchronized medium access which would result in collisions. Each node selects a random destination at the beginning of the simulation.

To simulate a level of asymmetry within the topology, originator messages and data packets are being filtered. Hence, at the beginning of the simulation, links are randomly chosen to be asymmetric with a link quality between 0.2 and 0.8 in one direction and 1 in the other. The asymmetric ratio determines the percentage of all present links on which the filter is inflicted, and does not change throughout the simulation run.

C. Results

Figures 2 and 3 show the resulting reliability of the simulation with an increasing asymmetric link ratio from 0 to 1 in steps of 0.2. Another curve shows the best reliability that is theoretically possible to achieve if the protocol would have perfect knowledge of the network. This curve is obtained by calculating the shortest path from every node to its traffic destination and averaging this value for every node and run. The inverse path quality was applied as link costs metric and path costs result in a multiplicative combination of link costs. While the amount of asymmetric links remains the same in the compared simulation runs, the best possible average path costs differ due to the random link qualities in every simulation run.

Additionally, another curve is drawn which we call best symmetric. This curve describes the expected reliability which would be allowed by the filter module if paths are chosen in the following way: the underlying graph that is

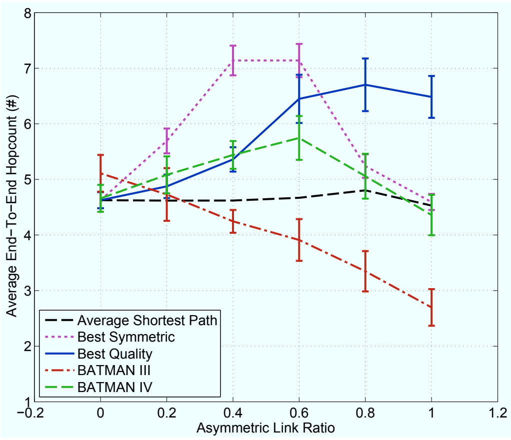

Figure 4. Grid, Topology A, Average End-To-End Hop Count

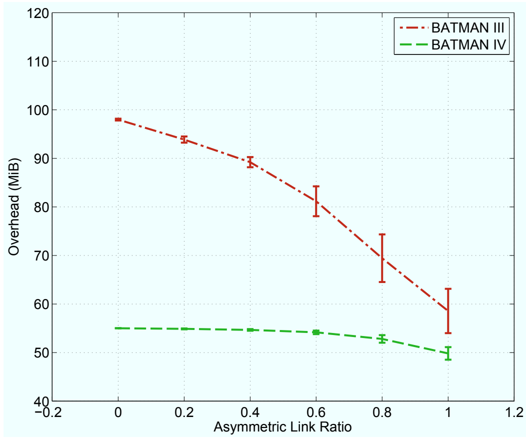

Figure 5. Grid, Topology A, Overhead

used for shortest path calculation is turned into an undirected graph, and link costs are calculated using the ETX metric for both directions. The objective of this curve is to have a comparison to a corresponding symmetric routing metric, which punishes the use of asymmetric links. In the following, our findings for the different topologies are discussed in detail.

1) Topology A: The best expected reliability shows a clear decreasing tendency with an increasing number of asymmetric links. Low asymmetry ratios yield the best possible paths and the packet delivery ratio remains at a value near 1. This is expected since their is a high chance that alternative routes are available allowing avoidance of lossy links. As the asymmetric links ratio increases, the chance of finding alternative routes to minimize packet loss decreases. This results in lower delivery ratios which is reflected by the slope of the graphs.

If we compare the best expected curve to the protocols performances curves, we can see differences in how version III and IV handle asymmetric links. The curve of version III decreases almost linearly with the increasing number of asymmetric links which is symptomatic for its tendency to route in the wrong direction with an asymmetric link topology. This tendency can also be found when looking at the average end-to-end hop count of the received packets in Figure 4. The hop count decreases, which points to the fact that particularly packets that routed over long paths (as in long physical distances) are lost. As longer paths accumulate lower reliabilities, the ones chosen apparently prove to be the lossier ones. Figure 4 also shows that version III does not detect the shortest paths available in the topology, without link loss present. This accounts to the circumstance, that OGMs do not spread as successfully in the center of the setup, where the node density is higher than at the border. Therefore it has a slight tendency to prefer longer paths along the border over the shortest path.

Version IV's curve on the other hand orients itself much on the best expected curve. This shows the protocol indeed tends to choose the path showing the highest quality. Yet the resulting values are lower for the hop penalty applied to the metric may lead to a path with a shorter hop count to be favored over one with a better quality. In addition, due to the asymmetric penalty, links with a high difference of link qualities in each direction may be avoided, even though they're part of the best path. Figure 4 shows that the protocol chooses the shortest possible path without asymmetry in the topology. Then the average hops increase, since the asymmetric links can at first be avoided by evading the asymmetric links with the usage of a longer alternative. For higher ratios the value decreases again, since the growing number of asymmetric links entails a decreasing number of better alternative paths.

Noticeably, version IV's routing metric outperforms the best symmetric reliability in Figure 2 already with an asymmetric link ratio around 0.4. Since the best symmetric curve is an idealized calculation, not taking into account the packet loss caused by the free space propagation model of the simulation, a routing protocol applying such a metric would achieve worse results. This indicates that there can be achieved a high benefit in terms of reliability by choosing an asymmetric metric in general.

The generated overhead is shown Figure 5. For version III it decreases dramatically with a decreasing link quality while version IV maintains an almost constant level. This results from version III's strict packet drop policy. In an ideal unpartitioned network without lossy links, both protocols rebroadcast every non-duplicate once. With lossy links, version III just rebroadcasts OGMs incoming over the best ranking neighbor or if the non-duplicate was received over a path with an equal length. This policy is necessary to prevent routing loops, similar to the feasibility conditions in the Babel protocol. Version IV on the other hand basically rebroadcasts any non-duplicate. There remains a high probability that any OGM reaches every node at least once, even with bad link conditions. Therefore, the amount of overhead remains at a constant level.

2) Topology B: The higher node density in the second topology example leads to a much higher number of links offered by the setup. Both protocols clearly profit from this circumstance such that they have more routing options available.

As can be seen in Figure 3, version III still has a linear decrease with a growing number of asymmetric links. The received packet count metric results in paths, with a high probability of a successful reception from a certain node. In the other direction, these paths have a linear increasing probability of being afflicted with a link loss. This mirrors in the resulting curve.

Version IV maintains an almost constant packet delivery ratio up to an asymmetric link ratio of 0.6. In comparison to the precious simulation it manages to find high reliable paths more easily. While the best expected average reliability remains close to 1, even with every link being asymmetric, version IV's reliability becomes lower. Since every link is then penalized with the asymmetric penalty, the better routing decision is harder to make. Particularly on longer paths, this has a severe impact on the resulting global TQ value since the best path may not be distinguished from a shorter path with less asymmetry.

Compared to the values of the best symmetric curve, version IV's benefit of its link metric exhibits itself only with a high ratio of asymmetric links, then being equally significant as in topology A. Thus, this approach is particularly promising with a lower number of alternative paths.

Figure 6 draws the same picture as Figure 4 in Section IV-C1. Except for both curves in the beginning staying closer to the average shortest path hop count. This stems from the higher probability of finding a high reliable short

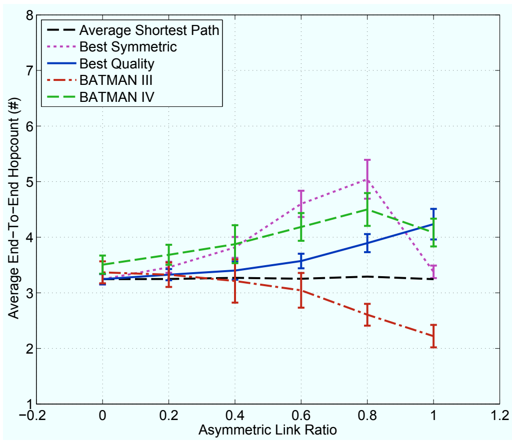

Figure 6. Grid, Topology B, Average End-To-End Hop Count

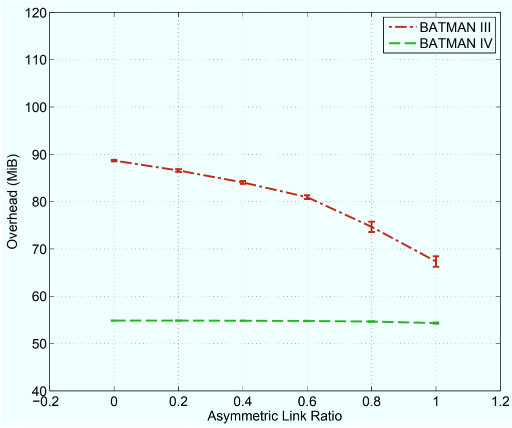

Figure 7. Grid, Topology B, Overhead

path. Noticeable is that, while there are many paths present with a reliability of 1, version IV does not always decide for the shortest one among them. Due to the low chosen hop penalty, equally reliable paths with a low hop difference are not reliably distinguishable.

The overhead generated by the protocols in scenarios A and B is shown in Figure 5 and Figure 7, respectively. The higher number of available links does not affect the behavior significantly. Version III's overhead shows a lower overall level owing to the fact that the average shortest path between a node pair has less hops in this topology. The rebroadcasting mechanism does not make use of longer alternative paths.

V. Conclusion

The B.A.T.M.A.N. protocol was initially created for a mesh environment where participating nodes are normally attached to a wall and connected to a power supply system. Thus, the majority of links was assumed to be bidirectional. The experience from real world deployments has shown that it is not sufficient to detect and avoid asymmetric links.

The core principle of both protocol versions relies in the usage of the sliding window approach for link detection. In version III, a small sliding window yields a fast detection of available paths. Thus, it is very well suited for scenarios with frequent topology changes. However, it is incapable of exploiting asymmetric links which makes it a less good solution for heterogeneous networks where nodes tend to have different characteristics. Our results have shown that version IV is able to estimate the performance of local links in both directions whereas the accuracy depends on the window size. In our future work, we will evaluate version V of B.A.T.M.A.N. protocol which is currently under development. Version V will include additional mechanisms which aim at providing further performance improvements in the context of asymmetric links.

REFERENCES

[1] A. Neumann, C. Aichele, M. Lindner, and S. Wunderlich, *Better Approach To Mobile Ad-hoc Networking (B.A.T.M.A.N.)*, Internet-Draft, pages 1-24, 2008. Network Working Group.

[2] L. Sang, A. Arora, and H. Zhang, *On link asymmetry and one-way estimation in wireless sensor networks*, ACM Transactions on Sensor Networks, 6(2):1-25, March 2010.

[3] D. Kotz, C. Newport, and C. Elliott, *The mistaken axioms of wireless-network research*, Technical Report TR2003-467, Department of Computer Science, Dartmouth College, 2003.

[4] M. Ikeda, G. Marco, L. Barolli, and M. Takizawa, *A BAT in the lab: Experimental results of new link state routing protocol*, In Proceedings of the 22nd International Conference on Advanced Information Networking and Applications, AINA '08, pages 295-302, 2008.

[5] T. Clausen and P. Jacquet, *Optimized Link State Routing protocol (OLSR)*, Technical Report RFC 3626, Network Working Group, October 2003.

[6] L. Barolli, M. Ikeda, G. Marco, A. Durresi, and F. Xhafa, *Performance analysis of OLSR and BATMAN protocols considering link quality parameter*, In Proceedings of the 2009 International Conference on Advanced Information Networking and Applications, AINA '09, pages 307-314, Washington, 2009.

[7] C. Perkins, E. Royer, and S. Das, *Ad hoc on-demand distance vector (AODV) routing*, RFC 3626, pages 1-75, July 2003.

[8] E. Kulla, M. Ikeda, L. Barolli, and R. Miho, *Impact of source and destination movement on MANET performance considering BATMAN and AODV protocols*, In Proceedings of the 2010 International Conference on Broadband, Wireless Computing, Communication and Applications, BWCCA '10, pages 94-101, Washington, 2010.

[9] A. Neumann, C. Aichele, and M. Lindner, *B.A.T.M.A.N. status report*, June 2007, http://downloads.open-mesh.org/batman/papers/batman-status.pdf, viewed April 2012.

[10] S. Wunderlich, M. Lindner, and W. Tsai, *B.A.T.M.A.N.-advanced documentation draft*, May 2009, http://git.open-mesh.org/?p=batman-adv-doc.git, viewed April 2012.

[11] C. Aichele, L. Lssing, M. Lindner, S. Wunderlich, and S. Eckelmann, *Open mesh website*, April 2012, http://www.open-mesh.org/, viewed April 2012.

[12] M. Abolhasan, B. Hagelstein, and J.C.P. Wang, *Real-world performance of current proactive multi-hop mesh protocols*, In Proceedings of the 15th Asia-Pacific conference on Communications, APCC '09, pages 42-45, 2009.

[13] T. Clausen, C. Dearlove, and P. Jacquet, *The Optimized Link State Routing protocol version 2*, http://tools.ietf.org/html/draft-ietf-manetolsrv2-12, July 2011.

[14] D. Johnson, N. Ntlatlapa, and C. Aichele, *Simple pragmatic approach to mesh routing using BATMAN*, In Proceedings of the 2nd IFIP International symposium on Wireless Communications and Information Technology in developing countries, WCITD '08, Pretoria, October 2008.

[15] J.C.P. Wang, B. Hagelstein, and M. Abolhasan, *Experimental evaluation of IEEE 802.11s path selection protocols in a mesh testbed*, In Proceedings of the 4th International Conference on Signal Processing and Communication Systems, ICSPCS '10, pages 1-3, 2010.

[16] D. Murray, M. Dixon, and T. Koziniec, *An experimental comparison of routing protocols in multi hop ad hoc networks*, In Proceedings of the Australasian Telecommunication Networks and Applications Conference, ATNAC '10, pages 159-164, 2010.

[17] D. Couto, D. Aguayo, J. Bicket, and R. Morris, *A high-throughput path metric for multi-hop wireless routing*, In Proceedings of the 9th annual international conference on Mobile computing and networking, MobiCom '03, pages 134-146, New York, 2003. ACM.