

A Comparison of OLSR and OSPF-MDR for Large-Scale Airborne Mobile Ad-Hoc Networks

Greg Kuperman, Leonid Veytser, Bow-Nan Cheng, Scott Moore, and Aradhana Narula-Tam
MIT Lincoln Laboratory
Lexington, MA, USA 02420
{gkuperman, veytser, bcheng, scott.moore, arad}@ll.mit.edu

ABSTRACT

In this paper, we compare the proactive MANET routing schemes of OLSR and OSPF-MDR via high-fidelity simulation, and consider their suitability for large-scale airborne networks. A successful MANET routing scheme must be bandwidth efficient and robust to frequent topology changes. To assess the two protocols, we simulate them in networks with up to 400 mobile nodes, under a variety of network densities. We evaluate them on the basis of the amount of routing overhead generated, the rate of successful packet delivery, and the time it takes until all of the routing tables converge. We find that OLSR requires up to an order magnitude higher router overhead than OSPF-MDR, while providing only a marginal benefit in packet delivery success rates. The largest difference between the two protocols is the time it takes for their routing tables to converge in the presence of packet loss. OLSR has consistent convergence times for networks of all sizes, while the convergence time of OSPF-MDR increases with network size.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*; C.2.2 [Computer-Communication Networks]: Network Protocols—*Routing protocols*

Keywords

OLSR; OSPF-MDR; MANET

1. INTRODUCTION

As the proliferation of advanced, mobile computing devices continues, there is an urgent need to be able to interconnect these systems with a reliable, high-speed network

This work is sponsored by the Assistant Secretary of Defense Research and Engineering (ASD R&E) via Air Force contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

(c) 2014 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the United States Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

AIRBORNE'14, August 11, 2014, Philadelphia, PA, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM Copyright 2014 ACM 978-1-4503-2985-9/14/08...\$15.00.

<http://dx.doi.org/10.1145/2636582.2637347>.

that will be capable of supporting a large number of nodes without any fixed infrastructure. This type of network architecture is commonly referred to as a mobile ad-hoc network (MANET), where wireless nodes must be capable of dynamically forming multi-hop routes between one another in the presence of a frequently-changing network topology. A prime example of such a need is the U.S. Department of Defense's desire for a network-centric military [1], where hundreds, if not thousands, of nodes will be interconnected via a dynamic and robust network. A large part of this network will operate in the airborne domain, which will include high-capacity aerial backbones, tactical edge-networks, and swarms of UAVs. While there has been significant research into enabling routing and connectivity for MANETs, there is still a lack of understanding regarding the behavior of many of these proposed solutions for larger-scale networks. In this paper, we evaluate two proactive MANET routing schemes, Optimized Link State Routing (OLSR) [2] and Open Shortest Path First with MANET Extensions (OSPF-MDR) [3], via high-fidelity simulation, and consider their suitability for large-scale networks, with a particular focus on the airborne domain.

Some of the key limitations of wireless mobile ad-hoc networks that differentiate them from static wired networks include lower link capacity due to noise and interference, and unstable links due to mobile nodes. A successful routing scheme for MANETs must address these limitations by being bandwidth efficient, and being robust to frequent topology changes. In the past two-decades, there have been many proposed protocols for MANET routing, with these approaches being either "proactive" or "reactive" [4]. In proactive routing, each node maintains an updated route to every other node, which is achieved by a periodic flood of link-state information. In reactive routing, a node will "discover" a route to another node only when it has data destined for that node. Proactive schemes have been previously shown to have lower latency and higher data delivery rates than their reactive counterparts [5]; hence, we limit our evaluation to proactive routing schemes.

OLSR has been compared to numerous other proactive routing schemes, and is typically considered to have superior performance [6, 7]. OSPF-MDR is a MANET extension to the OSPF routing protocol, which is widely used in wired networks. OSPF-MDR was developed more recently, and has key components that differentiate it from OLSR, including a novel mechanism for dissemination of control traffic using connected dominating sets (CDS). There has been limited comparison of the two routing schemes to one

another. In [8] and [9], the authors of both papers run tests on 40 node networks, and conclude that OLSR outperforms OSPF-MDR. To our knowledge, no studies have gone beyond networks of 40 nodes. As our simulations show, the performance of the routing protocols changes significantly as the size and density of the network varies. To better understand how these routing protocols behave, it is important that we test them under differing network sizes and densities.

In this paper, we evaluate and compare OLSR and OSPF-MDR for airborne networks that contain up to 400 nodes under a variety of node densities. We focus on airborne networks, where nodes move at faster speeds and have higher mobility than nodes in a ground network. Because of these characteristics, fixed infrastructure networks are infeasible for the airborne domain, and MANET is the most promising solution. This evaluation of the two routing protocols is done via a high-fidelity simulation using OPNET [10]. In Section 2, an overview of OLSR and OSPF-MDR is provided. In Section 3, our simulation results comparing the two routing schemes are presented and discussed. In Section 4, we conclude and offer suggestions for a next step forward in designing appropriate MANET routing schemes for large-scale networks.

2. OVERVIEW OF OLSR AND OSPF-MDR

In this section, an overview of OLSR [2] and OSPF-MDR [3] is provided. In proactive routing, the network topology is periodically disseminated across the network such that each node maintains a route to every other node. The advantage of this approach is that it reduces the delay in finding a path (there should be almost no delay since each node proactively calculates routes), but this comes at the expense of increased routing overhead used to maintain the routing tables.

In OLSR, there are two main message types: Hello and topology control (TC). Each node periodically sends a hello message to all of its direct (one-hop) neighbors. In a hello message, a node lists all of its one-hop neighbors; consequently, a node receiving a hello message is able to learn about all of the nodes in its two-hop neighborhood. Topology control messages contain a node's view of the link states for the entire network. These messages are flooded throughout the network, allowing nodes to proactively form end-to-end paths to any other node. Hello and TC messages are sent periodically, with a default periodicity of 2 and 5 seconds, respectively.

Having each node retransmit every TC message is bandwidth inefficient, and would quickly overwhelm a wireless network if the number of nodes grows too large. In OLSR, a node selects a subset of its neighbors to relay the TC messages, which reduces the overall bandwidth used to disseminate topology information. The nodes selected to retransmit a TC message are known as multipoint relays (MPR). Each node independently selects its own set of MPR nodes according to an algorithm given in [11], with the MPR nodes being chosen such that each two-hop neighbor of a node can be reached via an MPR.

OSPF-MDR has three main message types: Hello, link state advertisement (LSA), and database description (DD). Similar to OLSR, nodes learn about their two-hop neighborhood via hello messages, with hellos being sent every 2 seconds. A key difference from OLSR is that OSPF-MDR uses incremental hellos: only changes from the last hello are

reported, otherwise an empty hello is sent. The LSA message contains a node's view of the network topology, and is similar to OLSR's TC message. To reduce the number of topology messages transmitted, LSAs are only generated if a topology change is detected; unlike the hello message, empty LSAs are not sent. In order to not send too many LSAs in a highly dynamic environment, there is a minimum interval between LSA transmissions, with the default value being 5 seconds.

To be more bandwidth efficient, OSPF-MDR also chooses a subset of a node's neighbors to relay topology information. OSPF-MDR introduces a new algorithm that has nodes coordinate between themselves such that they select a set backbone nodes that form a connected dominating set¹ (CDS). This is in contrast to OLSR's MPR algorithm, where each node independently selects its own set of MPR nodes. The nodes that are selected to be part of the CDS are called MANET designated routers (MDR). Furthermore, OLSR requires that every two-hop neighbor of a node must be able to be reached through an MPR node; OSPF-MDR has no such requirement for MDR nodes. OSPF-MDR also has the option to minimize bandwidth usage by sending LSAs for the entire topology or just part of it. At a minimum, LSAs can contain only the MDR nodes, which will cause all routes to traverse the set of backbone nodes. The database description message is transmitted whenever an MDR node forms a new adjacency with another node, and is used to synchronize the two nodes' databases of link states.

3. SIMULATION OF OLSR AND OSPF-MDR

In this section, we evaluate and compare OLSR and OSPF-MDR using a high-fidelity simulation. The goal is to understand how these routing protocols behave in large-scale MANETs, under a variety of network sizes and densities. We focus our work on the airborne domain, where the high speed and mobility of nodes necessitates a multi-hop MANET solution in order to interconnect all of the aircraft. For a routing protocol to be successful in a MANET environment, it must be bandwidth efficient and resilient to frequent topology changes. To gauge bandwidth efficiency, we measure the amount of routing overhead generated by each of the protocols. For resiliency, we measure two things: packet delivery success rates, which captures how many packets encounter stale or non-existent routes, and convergence time, which measures how long it takes for a node to find a complete set of routes after some change in the network. In Section 3.1, we present the parameters of our simulation, as well as the simulation scenario. In Section 3.2, we present the results of the simulation, and discuss their causes and implications.

3.1 Simulation Parameters and Scenario

For our simulation, we used the OPNET network modeler [10]. All configurable parameters for the two routing protocols are set to their default values as described in their respective IETF RFC documents. For OSPF-MDR, the amount of topology information disseminated in the LSA messages is set to be the minimum required to calculate shortest paths between nodes ($LSAFullness = 1$). Since OLSR also finds shortest paths between nodes, using this LSA dissemination setting for OSPF-MDR allows for a fair comparison between the two protocols. We only wish to

¹In a connected dominating set, every node is either part of the CDS or is adjacent to a node in the CDS.




Figure 1: Average routing overhead per node

measure and compare the performance of the routing algorithms; hence, we want to minimize any cross-layer interaction between the routing and lower layers. For the datalink/physical layer, we choose the OPNET “Smart MAC”, which is configurable with respect to data rates, transmission distances, and packet error rates. For all scenarios, the data rate for any link is set to 24 Mbps, and the transmission distance is set to 300 nautical miles (nmi), which is approximately 560 km. If a node is within another node’s transmission range, then a link exists between those two nodes; otherwise, those nodes cannot communicate.

Two scenarios are run. For the first scenario, we wish to measure the routing overhead and packet delivery ratio when the nodes are mobile. The network has 200 nodes, which are initially distributed in a uniform fashion across three differently sized square regions: 1200 nmi², 1500 nmi², and 1800 nmi². The smaller region yields a higher node density, while the larger region has a lower density. For mobility, the random waypoint model is used, with all nodes moving at the same speed. The simulations are run with nodes moving at five different speeds, ranging from Mach 0.82 to Mach 4.1, where Mach 1 = 0.18 nmi/sec. Data packets are sent in an “all-to-all” fashion, where each node sends a 16 byte UDP packet to every other node every 2 seconds. The packet delivery ratio is the number of UDP packets successfully received at their intended destinations compared to the number of packets transmitted. The routing overhead is all other traffic in the network. The simulation models 60 minutes of network runtime.

In the second scenario, we measure the convergence times of the routing protocols. Specifically, we measure the time it takes from network initialization until each node has a routing entry for every other node. If the network topology is constantly changing, it would be difficult to isolate this metric; hence, the scenario is run without mobility. Instead of mobility, we add a 10% probability of packet error to emulate losses that might be experienced in a wireless environment. Networks from 50 to 400 nodes are simulated, where the nodes are arranged in a grid pattern across a square region of 1800 nmi². Since we are only interested in the convergence time, no data packets are sent.

3.2 Simulation Results

For clarity and exposition, we break the discussion of the simulation results into three components: Routing overhead, packet delivery ratio, and convergence time.

3.2.1 Routing Overhead

In Figure 1, the average amount of routing traffic generated per node is plotted for OLSR and OSPF-MDR for

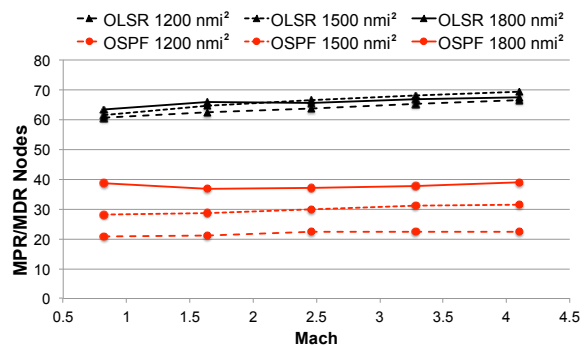


Figure 2: Percent of nodes selected as MPR/MDR

the three different sized networks. Three observations from these results are:

1. OLSR requires significantly higher routing overhead than OSPF-MDR. At Mach 0.82, OLSR generates almost an entire order of magnitude more control traffic than OSPF-MDR does. The gap narrows a bit as the speed increases, but even at Mach 4.1, OLSR creates almost three times as much routing overhead.
2. As network density decreases (area size increases), the amount of routing overhead increases. Additionally, this increase is more pronounced for OSPF-MDR than it is for OLSR.
3. As the speed increases, OSPF-MDR has a much more significant increase in the amount control traffic generated than OLSR does. In fact, OSPF-MDR almost triples going from Mach 0.82 to Mach 4.1, while OLSR only has about a 20% increase.

For the first observation, the cause of the large disparity in router overhead between the two protocols can be explained by examining the differences in how each algorithm chooses the relay nodes that disseminate topology data. OLSR uses the MPR algorithm, as defined in [11], while OSPF-MDR uses its own MDR algorithm, as discussed in Section 2. For MPR selection, each node independently chooses a subset of its neighbors such that all of its two-hop neighbors can be reached via an MPR node. For MDR selection, nodes coordinate between themselves to select a subset of nodes that form a connected dominating set (CDS) to relay topology data. In OLSR, since nodes choose their own set of MPRs independently of one another, there is no attempt to have MPR nodes be common amongst neighboring nodes. This causes redundancy in the number of nodes sending topology data throughout the network, which causes many more topology messages to be sent than otherwise might have been required. In fact, the minimum number of nodes needed to retransmit topology information throughout the network would be the minimum CDS, which is a CDS with the smallest number of nodes forming a backbone. OSPF-MDR uses coordination to select MDR nodes that form a CDS in a distributed fashion.

In more dense networks, nodes are closer to one another (i.e. paths are shorter). Hence, the number of nodes needed to relay topology data will be fewer than a less dense network. For example, in a long line network, every node except the two end nodes must be a relay node. Alternatively, if all nodes were within transmission range of one another (a complete graph), then only one node needs to be designated a relay node (assuming at least one node must be a relay).

In Figure 2, the percentage of nodes that are designated as MPR/MDR are plotted with respect to speed and density. As expected, since OSPF-MDR coordinates between nodes to choose a CDS for MDR nodes, the number of MDR nodes decreases as the network density becomes higher (smaller network area). Since OLSR has nodes independently choose their own set of MPR nodes, the number of MPR nodes remains relatively constant, regardless of network density. In fact, because of this lack of coordination, approximately $\frac{2}{3}$ of nodes become MPRs.

Another factor in allowing OSPF-MDR to generate less router overhead is its use of differential hellos and topology updates. Typically, only the changes are reported, with “full” messages being transmitted more infrequently. This is in contrast to OLSR, which sends a full hello and topology update each period.

The second observation, that overhead increases with decreased network density, particularly for OSPF-MDR, can be explained by seeing how the number of MDR nodes change with respect to density. As density decreases, more MDR nodes are required to form a CDS, which in turn will cause a greater number of topology messages to be transmitted across the network. The reason that OLSR also increases as density increases is because in denser networks, nodes have more neighbors, and the hello message will then be larger in size. The number of hello messages stays constant (one hello transmitted per node), but those hellos are now larger. We note that this also contributes to increased routing overhead in OSPF-MDR.

We now discuss the third observation: Increased speed causes a significantly larger increase in routing overhead for OSPF-MDR than it does for OLSR. This can be explained by the Database Distribution (DD) message that OSPF-MDR has, for which OLSR does not have an equivalent. As described in Section 2, when a new adjacency is formed between an MDR node and some other node, the entire database of topology information is exchanged between the two nodes. When mobility increases, nodes are forming adjacencies at much higher rates than when mobility is low. Consequently, there is an increase in database distribution messages being transmitted.

3.2.2 Packet Delivery Ratio

While our simulations show OSPF-MDR generating lower routing overhead than OLSR, that does not necessarily mean that OSPF-MDR outperforms OLSR. The other key metric that must be evaluated is how successful a routing scheme is at delivering packets to their intended destination. To evaluate how well the two protocols deliver packets, each node generates a set of 16 byte UDP packets every two seconds that are destined to every other node in the network. In a MANET, mobility causes frequent changes in topology, and routing entries can quickly become stale. If a routing table has an entry for a next hop to a node that is no longer within range, a packet will still be transmitted, and consequently will not arrive at the destination. Additionally, if a packet arrives at some node, and there is no longer a routing entry at that node for the intended destination of the packet, the router will drop the packet.

In Figure 3, the percentage of successfully received packets is plotted for both OLSR and OSPF-MDR for the different sized networks. A few observations from these results are:

1. As the speed of the nodes increases, the packet delivery




Figure 3: Successful packet delivery ratio

ratio decreases. At lower speeds, the success rate is up to 99% for both OLSR and OSPF-MDR, regardless of density. At the fastest speeds, the success rates drop to approximately 90% for the two protocols.

2. More dense networks have higher packet delivery ratios than less dense ones. As noted in the previous observation, at lower speeds, density has little impact. But at higher speeds, density does have a large impact on the delivery success rates. At Mach 4.1, the difference in the packet delivery ratio is 5% for both protocols.
3. OLSR has higher delivery rates than OSPF-MDR, but not by a significant margin. The largest gap is 2%, and occurs when nodes are traveling at the fastest speed, operating in the least dense network.

The first observation is as anticipated. As the speed increases, nodes are more likely to move apart from one another. If a packet arrives before the departure of a node is discovered, that node may be used as a next hop, causing the packet to be lost. If a node has moved out of range, a new route may not have been discovered yet, causing the packet to be dropped at the router. The second observation is also not surprising. As network density increases, the path between two nodes becomes shorter. Consequently, with a shorter path, a packet has fewer hops to traverse before it reaches its destination, giving it fewer opportunities to be lost along the way.

The third observation, that OLSR has higher delivery rates than OSPF-MDR, can be explained by considering the discussion from the previous section regarding router overhead. OSPF-MDR tries to minimize the number of MDR nodes, which relay topology information throughout the network. If a node changes position and loses its connection with an MDR node, it may take a few exchanges of messages before it forms an adjacency with some new MDR. This will cause a delay before that node begins receiving topology information again. Another possibility is that the CDS that the MDR nodes form becomes disconnected from one another. In that event, a new CDS will have to be established before topology information can be properly disseminated throughout the network. These events are less likely in OLSR, because of the high number of MPR nodes, which gives a higher redundancy of relay nodes. The probability that a node moves away from all of its MPR nodes, or that a set of MPR nodes becomes entirely disconnected from another set, is lower than a similar event happening with MDR nodes. Nonetheless, OLSR and OSPF-MDR both degrade similarly with respect to speed and density, and only differ by 2% in the most extreme case.




Figure 4: Convergence time of OLSR and OSPF-MDR

3.2.3 Convergence Time

The final metric we consider is the amount of time it takes for all nodes to find routes to all other nodes during network initialization. We label this as the “convergence” time of the routing protocol. To isolate how long a routing protocol takes to find a stable solution, we eliminate mobility from the simulation. Instead of mobility, we add a 10% probability of packet error to emulate losses that might be experienced in a wireless environment. The number of nodes are varied from 50 to 400, and the nodes are arranged in a grid across a square region of 1800 nmi².

The results for convergence time are plotted in Figure 4. Without packet errors, OLSR and OSPF-MDR both have relatively constant convergence times of 10 and 15 seconds, respectively. As expected, the convergence time of OSPF-MDR is higher than that of OLSR. In OLSR, the relay nodes (MPRs) are chosen without coordination between any of the nodes. In OSPF-MDR, coordination is required to select relay nodes such that they form a CDS. While this coordination allows the routing overhead to be significantly reduced, the extra messaging causes longer convergence times.

With packet errors, OLSR still has a relatively constant convergence time of 10 seconds for networks of all sizes. Unlike OLSR, the convergence time for OSPF-MDR increases significantly for larger networks. This high convergence time is due to very few “bad” nodes; typically only one or two nodes require more than 15 seconds to form full routing tables. OLSR is robust to failure because topology messages are flooded across a much larger set of nodes, which allows it to tolerate loss of control packets. OSPF-MDR requires more coordination and has fewer relay nodes, which causes it to be potentially more vulnerable to loss of control packets. We have not yet identified the exact reason as to why these few nodes take so long to converge, but we conjecture that it is due to a specific set of packets being lost, coupled with a particular ordering of these loss events. Such events are low probability, and hence not seen in smaller networks that have fewer control messages. We conjecture that this issue has not surfaced previously because experiments have not yet been conducted on such large mobile networks (400 nodes). Further investigation is underway into better understanding the problem, and devising a solution to allow those few nodes to converge more quickly.

4. CONCLUSION

In this paper, we compared and evaluated OLSR and OSPF-MDR for large-scale airborne MANETs. We ran a high-fidelity simulation of the two protocols using OPNET, using networks up to 400 nodes, with a variety of densities. In particular, we considered the following metrics in our evaluation: routing overhead, packet delivery ratio, and

convergence time. For routing overhead, OSPF-MDR generated far less control traffic than OLSR; under lower mobility and denser networks, OSPF-MDR created an order of magnitude less overhead. This is due to the inherent differences between how the two protocols select nodes to relay topology information throughout the network: OLSR picks many more than OSPF-MDR does. Since OLSR picks more relay nodes, it has redundancy in control traffic; hence, it becomes slightly more robust than OSPF-MDR against topology changes due to mobility. For convergence time, we found that OLSR had a fairly constant convergence time of about 10 seconds with respect to network size and packet failure rate. OSPF-MDR performance was similar without packet losses, but packet losses resulted in significantly longer convergence times for larger networks due to certain loss events that prevented one or two nodes from converging. Further research is being conducted to understand and identify the cause of these complex events that result in long convergence times. Overall, OSPF-MDR seems to be a promising solution for the airborne domain. However, to scale to larger networks the convergence issue must be resolved. Further investigation is recommended.

5. REFERENCES

- [1] C. Wilson, “Network centric operations: Background and oversight issues for congress.” Defense Technical Information Center (DTIC) Document, 2007.
- [2] T. Clausen and P. Jacquet, “Optimized link state routing (olsr),” *IETF RFC 3626*, 2003.
- [3] R. Ogier and P. Spagnolo, “Mobile ad hoc network (manet) extension of ospf using connected dominating set (cds) flooding,” *IETF RFC 5614*, August, 2009.
- [4] E. M. Royer and C.-K. Toh, “A review of current routing protocols for ad hoc mobile wireless networks,” *Personal Communications, IEEE*, vol. 6, no. 2, pp. 46–55, 1999.
- [5] C. Mbarushimana and A. Shahrabi, “Comparative study of reactive and proactive routing protocols performance in mobile ad hoc networks,” in *Advanced Information Networking and Applications Workshops, 2007*, vol. 2. IEEE, 2007, pp. 679–684.
- [6] S. Ade and P. Tijare, “Performance comparison of aodv, dsdv, olsr and dsr routing protocols in mobile ad hoc networks,” *International Journal of Information Technology and Knowledge Management*, vol. 2, no. 2, pp. 545–548, 2010.
- [7] A. K. BR, L. C. Reddy, and P. S. Hiremath, “Performance comparison of wireless mobile ad-hoc network routing protocols,” *IJCSNS*, vol. 8, no. 6, p. 337, 2008.
- [8] J. Fang, T. Goff, and G. Pei, “Comparison studies of ospf-mdr, olsr and composite routing,” in *MILCOM 2010*. IEEE, 2010, pp. 989–994.
- [9] B.-N. Cheng and S. Moore, “A comparison of manet routing protocols on airborne tactical networks,” in *MILCOM 2012*. IEEE, 2012, pp. 1–6.
- [10] O. M. Documentation, “Opnet technologies,” *Inc.[Internet] http://www.opnet.com*, 2003.
- [11] A. Qayyum, L. Viennot, and A. Laouiti, “Multipoint relaying for flooding broadcast messages in mobile wireless networks,” in *Hawaii International Conference on System Sciences, 2002. HICSS.*, pp. 3866–3875.