A WIRELESS INTERFACE TYPE FOR OSPF

Thomas R. Henderson, Phillip A. Spagnolo, and Jae H. Kim
Boeing Phantom Works
P.O. Box 3707, MC 3W-51, Seattle, WA 98124-2207
{thomas.r.henderson, phillip.a.spagnolo, jae.h.kim}@boeing.com

ABSTRACT

We report on a new wireless interface type for the OSPF routing protocol that greatly improves performance in multihop wireless networks. In this paper, we describe how to adapt OSPF to adequately handle wireless networks in a scalable manner. Our simulation-based study first illustrates that the operation of legacy OSPF routers in mobile ad hoc network environments can generate prohibitive amounts of overhead for networks as small as 20 nodes. Next, we study the possibility of constraining the number of OSPF adjacencies formed in the wireless network, and find that, while scalability can be improved in this manner, the performance does not approach that of the OLSR protocol designed specifically for mobile networks. Consequently, we describe how OSPF can be extended to incorporate key design concepts from OLSR, such that the scaling performance of OLSR can be approached while maintaining backward compatibility with legacy OSPF routers. Our "wireless" interface type achieves bandwidth efficiency within a factor of two to three of OLSR, and allows wireless networks to be incorporated into larger heterogeneous networks. Finally, we describe some strategies and open issues regarding the composition of large heterogeneous networks including wireless subnets.

Keywords: OSPFv2; OLSR; wireless networks; ad hoc networks; MANET.

INTRODUCTION

The Open Shortest Path First, version 2 (OSPFv2) routing protocol [1] is one of the most widely used interior routing protocols in the Internet, and a version for IPv6 (OSPFv3) has also been standardized. OSPF is a "link state" routing protocol that disseminates complete topology information to all nodes in the network via a reliable flooding algorithm, and then uses the Dijkstra shortest path algorithm to compute paths. OSPF supports heterogeneous network topologies and classifies a router's interfaces from a set of five possible types: point-to-point, broadcast, non-broadcast multiple access (NBMA), Point-to-Multipoint, and virtual. The two interface types that are most applicable to wireless networks, broadcast and Point-to-Multipoint types, each have limitations. Broadcast interfaces and networks, while being the most bandwidth efficient, rely on election of a "designated router" that has special responsibilities for the broadcast network. In a wireless network without full mesh connectivity, the designated router election protocol may not converge, leading to inoperative routing. In modern OSPFv2 implementations, Point-to-Multipoint interfaces can be configured to capitalize on the underlying broadcast nature of the wireless channel by sending certain protocol messages as multicasts, but scalability is a limitation since each router must maintain an adjacency with every other router ($O(N^2)$).

BACKGROUND

The Internet Engineering Task Force (IETF) has been considering a class of routing protocols optimized for mobile ad hoc networks (MANET). One such protocol, Optimized Link State Routing (OLSR) [2], resembles OSPF in several aspects; two key differences are that OLSR uses an unreliable, periodic flooding algorithm optimized for bandwidth conservation, and OLSR is presently specified only for broadcast-based multihop wireless networks with equal-cost link metrics. While approaches such as OLSR and other MANET protocols are promising, they are not yet full-fledged routing protocols with well-developed transit network interfaces and behaviors.

Few contributions in the literature are aimed at optimizing OSPF operation for wireless networks. The paper by Wollman and Barsoum [3] advocates adjustment of certain default timer values, and heavier reliance on route aggregation, to solve some OSPF performance problems. Baker [4] has suggested OSPF modifications for MANET, in the context of OSPFv3 for IPv6, that allows for roaming between areas, specifies a new MANET interface type, allows for more complicated link metrics, improves scalability by limiting the number of neighbor relationships formed, and considers IPv4 and IPv6 integration issues. BBN has developed a proprietary "Radio OSPF" that aims to leverage side information available at a wireless router to eliminate the sending of OSPF Hello neighbor discovery messages and modify designated router-related flooding [5].

Our goal is to adapt OSPFv2 to be able to adequately handle wireless networks in a scalable manner. We envision that a modification of an existing, commercially proven routing protocol to support wireless operation will provide a faster path to commercial availability of such a capability. Our work makes the following primary contributions:

i) we demonstrate and explain the shortcomings of current OSPFv2 implementations when operating in a mobile, multihop wireless environment;

ii) we evaluate the efficacy of a heuristic that modifies the Point-to-Multipoint interface type to limit the number of adjacencies formed, thereby improving performance, but demonstrate that such a modification still does not approach the performance of native ad hoc routing protocols such as OLSR;

iii) we describe the implementation and performance of a wireless interface type that utilizes key optimization mechanisms used in the ad hoc OLSR routing protocol, and show that it can achieve bandwidth efficiency within a factor of two to three of OLSR while still offering full backward compatibility with heterogeneous OSPF-based networks.

METHODOLOGY

We performed network simulations using the QualNet simulator. Earlier in 2002, we validated basic OSPFv2 performance against the OSPFv2 model in the OPNET simulator. We are currently developing a Linux-based implementation of the proposed wireless interface type, for comparison with simulation results.

We extended the QualNet OSPFv2 implementation as follows: i) we added a multicast-capable Point-to-Multipoint interface type, ii) we added and used layer-2 notification capability (in which a failure to receive an acknowledgement from a link layer unicast resulted in immediate notification to the routing protocol, rather than letting it wait for missing Hello packets), and iii) we added detailed event tracing capabilities to supplement the built-in statistics reporting. We used a "flooding" timer (uniform random variate between [0,100] ms) to suppress immediate response to Link State Request packets, a small amount of broadcast jitter (uniform random variate between [0,1] ms) for various protocol events such as Hello transmissions, and a delay for coalescing ACKs (between [0,100] ms). To avoid synchronization effects, we brought each router up at a random time in the first 30 minutes of a simulation, and then discarded the first 30 minutes of simulation data. Each simulation trial was conducted for 60 minutes total. In conducting simulation trials, different random seeds were used for similar configurations and the results averaged; our data illustrates error bars that represent one sample standard deviation from the mean. Finally, we also developed for QualNet a full implementation of draft version 7 of the OLSR protocol, including an identical layer-2 notification, but without a "full link state" option.

We used the standard 2.4 GHz 802.11b channel and MAC models in QualNet, configured for Ricean fading based on omnidirectional antennas. Our main network scenario of interest was the mobile ad hoc network (MANET) scenario. The scenarios are based on a number of mobile nodes moving in a 500 meter by 500 meter square grid. Each node was a router and contained an 802.11b radio operating at a carrier rate of 1 Mbps. Nodes moved according to a random waypoint model; each node selected a random destination in the grid and selected a random speed between 0 and 10 meter/s, and upon reaching the destination, paused for 30 seconds before moving again. A low amount of user data traffic was configured to measure the delivery performance of the routing; node N periodically sent packets to node N+1, at a rate of 1.6 Kb/s.

The two key parameters varied in these experiments were radio range and the number of nodes. The range was varied from 60 meters up to 500 meters. The effect of reducing the radio range was to move from a fully meshed, yet mobile, topology to one that increasingly reduced the number of edges in the connectivity graph, and correspondingly increased the rate of change in the underlying graph topology. An increase in the number of nodes corresponded to an increase in the node density, which on the one hand increased the degree of the graph when the power was low, but on the other hand increased the RF interference.

BASIC OSPF PERFOMANCE

The OSPFv2 Point-to-Multipoint interface type was originally intended for non-broadcast networks, such as Frame Relay networks, in which not all routers can communicate directly. The type has been generalized by router vendors to also support multicast-capable networks in which all routers may not be able to communicate directly at all times. While multicast-capable Point-to-Multipoint mode allows routers to more efficiently send Hello and Link State Update messages, it still requires each router to set up a pairwise adjacency with each other router. In a dynamic topology, when many routers are involved, the bandwidth required to synchronize databases and perform reliable flooding of updates can severely limit scalability.

Figure 1 illustrates the performance of mobile ad hoc networks of various sizes when each router's wireless interface is configured in multicast-capable Point-to-Multipoint mode. The overhead in Figure 1a counts overhead at the IP level, and is expected to be low at 500m (almost everyone can talk to everyone else all of the time), increase as the radio range is reduced (neighbor relationships become more dynamic, but the network as a whole remains mostly connected), and then decrease again for short-range radios (topology becomes disjoint). As expected, when the range is short, the delivery ratio (Figure 1b) can be improved by adding more nodes to the topology. However, for a network size of 24 nodes or larger, the network is not able to initially converge even for very low rate of topology change in the network (500m radio range), due to excessive routing protocol overhead. As the number of nodes grows above roughly 20, the bandwidth consumed by OSPF becomes prohibitively large (greater than 300 Kb/s for some mobility scenarios), leading to poor data delivery ratios (Figure 1b) and average data packet delays in the tens of seconds or greater. In summary, for a dynamic wireless topology, the simulation data suggests that OSPFv2 may not be able to support more than about 20 nodes on a 1 Mb/s 802.11 channel, with overhead scaling faster than linear with the number of nodes.

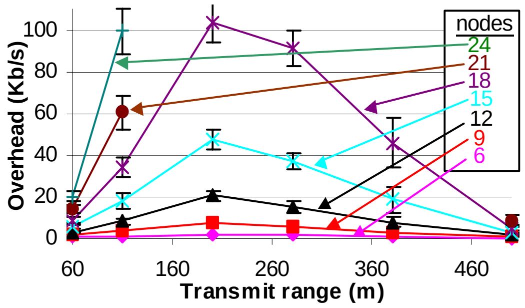

(a) IP overhead vs. mobile transmit range

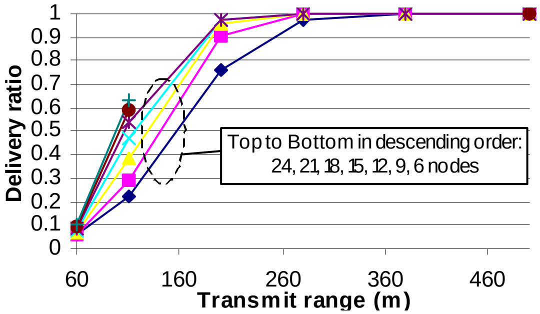

(b) Packet delivery ratio vs. mobile transmit range

Figure 1. OSPFv2 Point-to-Multipoint performance in simulated network for 6-24 nodes. The routing topology did not consistently converge for radio ranges greater than 110m when the network size was 21 nodes or greater, leading to prohibitively high congestion (such points are not plotted).

OLSR PERFORMACE COMPARISON

In contrast to OSPF, OLSR was designed specifically for MANET operation. Figure 2 demonstrates the performance of OLSR for MANETs with 12, 24, 36, 48, and 60 nodes. The overhead generated in OLSR is approximately an order of magnitude less than that of OSPF with Point-to-Multipoint for the corresponding number of nodes, and the rate of overhead growth with increasing numbers of nodes is slower than for OSPF. Additionally, the delivery ratio in OLSR is about the same as OSPF for low number of nodes. We observed that the delivery ratio for more disconnected topologies (radio ranges less than 260m) was lower for OLSR than for OSPF. We conjecture that this is because OSPF, through its reliable flooding mechanism, is better able to disseminate timely routing information than is OLSR, which relies on unreliable flooding and the dissemination of partial topology information.

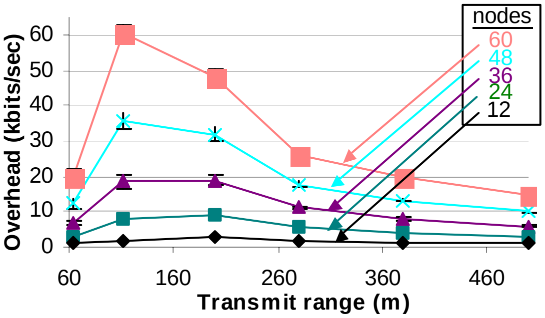

(a) IP overhead vs. mobile transmit range

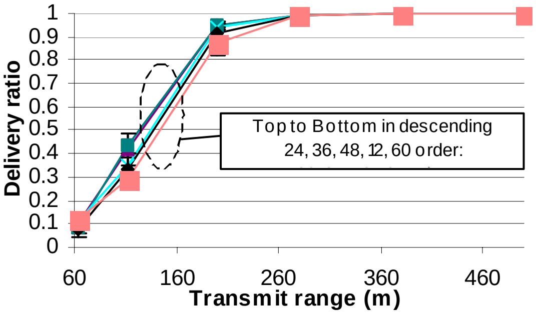

(b) Packet delivery ratio vs. mobile transmit range

Figure 2. OLSR performance in simulated network for 12, 24, 36, 48, and 60 nodes.

LIMITING OSPF ADJACENCIES

The OLSR results suggest room for improvement in OSPF scaling performance. One problem of the Point-to-Multipoint interface is that it requires a pairwise adjacency between each set of communicating routers. Baker suggests in [4] that OSPF scaling could potentially be improved by limiting the number of neighbors with which OSPF routers maintain a relationship. However, while [4] suggests possibly limiting the number of neighbors reported in Hello messages, we instead investigated techniques to limit the number of adjacencies that reach "Full" state. In our modification to Baker's proposal, routers continue to identify and list all routers from which a Hello message has been recently received, but the number of neighbors with which a router attempts to perform a database exchange is constrained to some pre-configured limit. Despite the constraint, a router is permitted to advertise, in its router Link State Advertisement (LSA), a link to any other one-hop router provided that the neighbor state is 2-Way or greater; this allows single hop routes to be used between routers even if they are not fully adjacent. In effect, LSAs are still fully disseminated through the network, but only via a subset of the possible router neighbor relationships.

A distributed algorithm that results in an optimal set of adjacencies given a constraint on the maximum number of adjacencies per node is a difficult problem, one that we speculate is NP-complete. We instead implemented the following heuristic. Assume the number of nodes is N, and the adjacency limit is L (an even number, greater than two) and identical for each node. Subdivide the nodes into groups of size N/(L/2). Allow nodes to form adjacencies with other members of their respective groups (which yields up to L/2 possible adjacencies). Otherwise, if nodes from different groups encounter one another, allow each node to attempt to form an adjacency so long as its total number of adjacencies is less than L (note that in this latter case, the nodes may disagree on whether they desire to form an adjacency). The heuristic should yield roughly L adjacencies in a well-connected network, with a possible maximum of min(3L/2,N) adjacencies (in the rare case of a node encountering L non-group members before encountering an additional L/2 group members).

The intuition behind the above heuristic is that, without changes to message formats, we want to make sure that the router is able to mutually agree with certain other routers that the pair should form an adjacency (one can envision pathological cases in which each router selects a different set of neighbors with which to form adjacencies, with the result being that no adjacencies are formed), but we want routers to also form some adjacencies with other random routers so that the sets of adjacent routers are not disjoint and LSAs are fully circulated.

Figure 3 illustrates the performance of this approach when applied to a 24 node network, again configured as Point-to-Multipoint but with the adjacency limitations. Recall from Figure 1 that standard OSPF was not able to operate successfully with 24 nodes in this mobile network. The results in Figure 3 show reduced overhead without sacrificing routing performance when the adjacency limit is 8 or fewer; when the limit is 12, the overhead is still too much for the 1 Mb/s 802.11 network. Although this method of operation improves scaling performance, it is still substantially worse than what can be obtained with OLSR (Figure 2).

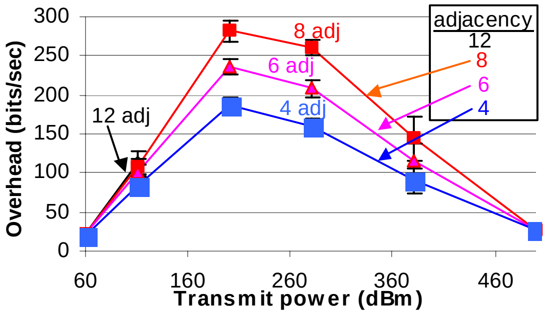

(a) IP overhead vs. mobile transmit range

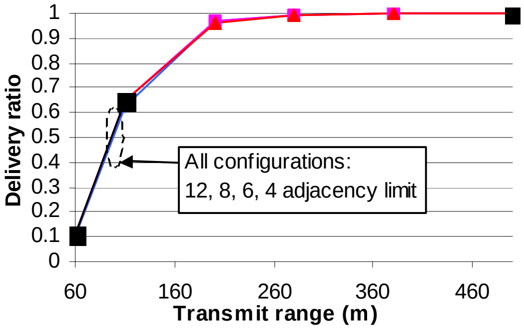

(b) Packet delivery ratio vs. mobile transmit range

Figure 3. OSPFv2 Point-to-Multipoint performance in simulated network for 24 nodes, with OSPFv2 adjacency limits of 12, 8, 6, and 4 nodes.

How can the scaling performance be improved to approach that of OLSR? One possibility is to use OLSR in mobile networks, and to define rules whereby routing information is exchanged between OSPF and OLSR. However, OLSR is not presently specified to handle detailed routing information such as variable cost link metrics and summarized prefixes that are found in heterogeneous IP networks. Another possibility is to operate OLSR as a "subnet" (layer 2) routing protocol and operate OSPF in designated router mode at layer 3. This approach may be viable but would likely require coordination between the layers to avoid redundant neighbor discovery operations, for example. A final possible solution is to attempt to incorporate, within the existing OSPF protocol, the aspects of OLSR that allow it to scale well in mobile networks. In the remainder of this paper, we describe the results of such an integration that allows OSPF to operate like OLSR in wireless networks, yet remains fully compatible with the rest of OSPF. The key is to define a new wireless interface type that replaces the reliable topology dissemination of OSPFv2 with something similar to OLSR's unreliable flooding mechanism.

WIRELESS INTERFACE TYPE

The two main contributors to the difference in overhead between OSPF and OLSR are OSPF's database synchronization and its reliable flooding mechanisms. These OSPF mechanisms are designed to achieve fast network-wide convergence. An alternative, adopted by OLSR, is to make flooding periodic and unreliable, and to optimize the flooding topology so that fewer broadcasts are needed to reach all nodes. A further step is to remove database synchronization procedures, allowing nodes to gradually learn of the topology as they listen to periodic broadcasts. These changes loosen the guarantees on convergence, and hence trade some routing optimality for lower bandwidth overhead. Within such a framework, the periodicity of the topology distribution becomes a key tuning knob, as more frequent flooding should improve routing performance, at the cost of bandwidth overhead.

One weakness of OLSR (as it is presently specified, not an inherent weakness of link state protocols) is its present applicability to heterogeneous networks. For example, OLSR is not presently specified to support variable cost link metrics, and does not specify rules for operating as a transit network. However, by taking core mechanisms of OLSR and defining them as an extension to OSPF, one can approach OLSR performance without losing the other features of OSPF useful for heterogeneous networks, such as support for broadcast networks, demand circuits, and other subnet technologies like Frame Relay and ATM. We have implemented and studied such an extension, which we call a "wireless" interface type for OSPF.

Within the wireless subnet, our wireless interface for OSPF functions as follows:

- Each node distributes a router LSA to all other nodes in the subnet. This router LSA is flooded unreliably via a new message type: the Link State Flood (LSF) message, similar to an OSPF Link State Update (LSU), but with the addition of a flooding sequence number. The LSF is multicast on the wireless interface every LSF_INTERVAL seconds, and also (optionally, not in our simulations) whenever there is a change to the LSA.
- Each node maintains a LSF duplicate table, and discards those LSFs that have already been processed. The monotonically-increasing LSF sequence number is used to detect duplicates.
- Each node selects a "multipoint relay" (MPR) set as described in [2]. MPRs essentially form an overlay lattice on the topology, and are intended to reduce unnecessary flooding transmissions. Only those nodes that have been selected in the MPR set of an LSF-originating node have responsibility to reflood any received LSFs from that node.
- A new "wireless" Hello message is used to implement the MPR selection algorithm as described in [2].
- In the wireless subnet, nodes do not attempt to enter into an OSPF neighbor state grater than 2-Way; if a neighbor is in state 2-Way, packets may be forwarded to it. This avoids database synchronization in the wireless subnet. Nodes will list neighbors in their LSAs if the neighbor state is 2-Way.

Figure 4 shows the overhead and delivery ratio when the wireless interface type is used in the MANET network with 12, 24, 36, and 48 nodes. The LSF_INTERVAL used is the OLSR default of 10 seconds. The routing overhead is significantly lower than that seen with the Point-to-Multipoint interface in Figure 1 and the Point-to-Multipoint interface with limited adjacencies in Figure 3. We also observe that such a decrease in the overhead does not force a tradeoff on data delivery ratio. The delivery ratios are only slightly lower than those of legacy OSPF when the radio range is low.

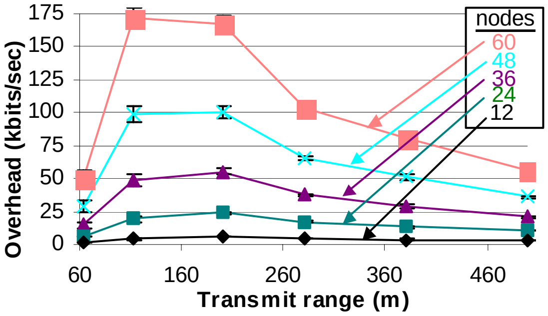

(a) IP overhead vs. mobile transmit range

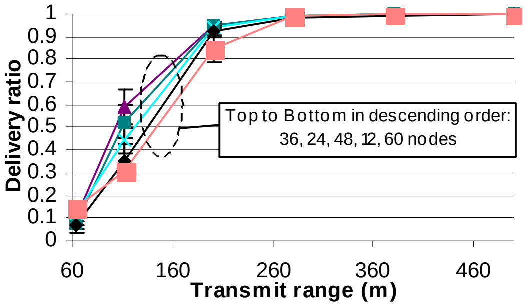

(b) Packet delivery ratio vs. mobile transmit range

Figure 4. Wireless OSPF performance in simulated network.

When compared to OLSR, the wireless interface provides slightly better levels in the delivery ratio, but the overhead is approximately two to three times that found in OLSR. OSPF with the wireless interface generates more overhead because each router in the wireless network must generate its own router LSA. In OLSR, the topology control (TC) is the comparable message to the router LSA, and only nodes selected as MPRs generate these messages. Therefore, if half the nodes are selected as MPRs then OLSR will generate approximately half as much topology control overhead as OSPF. Furthermore, the TC only informs nodes of a subset of the links in the topology (the MPR links), and the generation of TC messages requires the MPR selection algorithm to first converge. We conjecture that the reason that the wireless OSPF delivery ratio outperforms that of OLSR is because of the more thorough dissemination of topology information in OSPF, since each router originates a complete description of its local topology. If our wireless interface were to adopt the partial link state TC mechanism of OLSR, we predict that performance would more closely approximate that of OLSR; however, such a change would represent more of a departure from legacy OSPF database data structures.

Finally, we investigated network behavior as we scaled two parameters: the number of nodes in the flat routing topology, and the LSF_INTERVAL. Using an LSF_INTERVAL of 10 seconds, even though OLSR uses about a third less traffic in large networks than wireless OSPF, both protocols tended to be able to support only up to about 70 nodes in our simulation (on a 1 Mb/s rate channel). This was because of the background unicast data load that scaled linearly with the number of nodes. We highlight this result because it demonstrates that scalability of large networks is not solely based on the routing protocol overhead. However, routing protocol overhead is still important to reduce because the multicast packets can collide with the unicast MAC transmissions. Secondly, we found that the data delivery ratios were relatively insensitive to extending the LSF_INTERVAL to 30 seconds. By extending the LSF_INTERVAL, topology updates can be coalesced at the expense of making cached routing information more stale. Different mobility models may yield different results.

HETEROGENEOUS NETWORKS

The principal benefit of using a modified OSPF rather than OLSR in a multihop wireless network can be seen when considering heterogeneous networks. In this section, we describe our design strategy for integrating the wireless network with the rest of the OSPF network, and illustrate the performance benefits in comparison with a full OLSR approach. An alternative is to deploy separate OSPF and OLSR routing processes, with the OLSR process handling the MANET subnet only, but this approach also has the drawback of requiring the management of route redistribution between routing protocols.

OSPF nodes may have multiple wireless interfaces on different subnets, and may have interfaces with traditional link types (such as point-to-point). In our study, we call such routers "hybrid" routers. Hybrid routers have two additional responsibilities: they need to correctly disseminate information about the wireless subnet to other OSPF subnets, and they need to disseminate LSAs learned from outside the wireless subnet to within the wireless subnet. Since we use standard LSAs, wireless topology information can be flooded outward as usual, allowing backward compatibility with legacy OSPF routers; the wireless subnet appears to the outside as a traditional Point-to-Multipoint subnet. Hybrid routers also distribute any LSAs learned from outside in an LSF message.

In large networks, it is possible that the number of LSAs that need to be distributed into a wireless subnet may be large, so large that the overhead required to periodically flood them exceeds the overhead required to operate OSPFv2 in traditional Point-to-Multipoint fashion. There are a number of potential strategies to alleviate this:

- if the wireless subnet is not used for transit by the larger OSPFv2 network, hybrid routers could suppress the sending of all outside LSAs and instead send a default route. In this case, traditional techniques for designing the wireless network as a "stub" area may be used.
- outside LSAs may be flooded at a lower rate than the internal LSAs, thereby trading off convergence time for bandwidth efficiency.
- hybrid routers may develop additional techniques to suppress redundant LSAs that may be flooded into a wireless network from multiple hybrid routers.

We have developed a soft-state mechanism that implements the last bullet above. Hybrid routers maintain another list, an "outside LSA" list, in which they hold copies of outside LSAs received on their wireless interface. These LSAs are maintained for some interval longer than the LSF_INTERVAL (in our simulations, we used an interval of twice the LSF_INTERVAL). If a hybrid router is considering the flooding of an outside LSA into the wireless subnet, it first scans its outside LSA list to determine if the LSA has been seen on the wireless subnet recently; if so, it suppresses the transmission. This generally leads to only one router disseminating each external LSA.

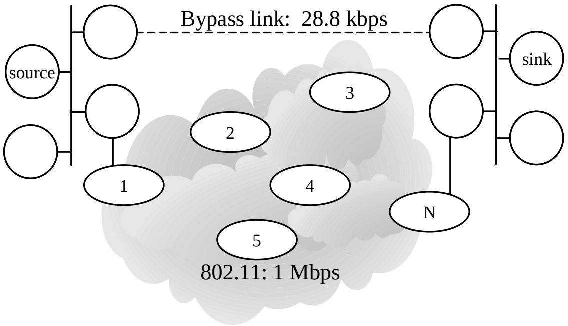

Figure 5. Heterogenous network with transit MANET.

Figure 5 depicts a heterogeneous network composed of two Ethernets linked by a MANET serving as a transit network between two end points. In parallel, a low bandwidth wired link connects two nodes to provide a backup path. The low bandwidth link has a data rate of 28.8 Kb/s while all other links have a rate of 1Mbps. The configuration of the MANET was unmodified from the above simulations. The simulation was run for two configurations: i) all routers are OLSR routers, and ii) all routers are OSPFv2 routers, with wireless interface types as applicable. In the latter case, the nodes numbered 1 and N are hybrid nodes. The hybrid nodes send out these LSAs at different times based on the interface type, and implement the soft-state mechanism for redundant LSA suppression described above. The simulations consisted of periodically sending UDP datagrams at 56kb/s from the source to the sink. In OSPF, all links were assigned a cost of 1 except for the bypass link, which was assigned a cost of 10. No costs were assigned in OLSR because routes are created based only on hop count. The total network overhead, the delivery ratio of the data, and the end-to-end delay were calculated when 12, 24, and 36 nodes were used in the MANET.

As shown in Figure 6, the overhead generated by OLSR is again about half that of OSPF. However, OSPF far exceeds the performance of OLSR in delivery ratio and end-to-end delay. This result is easily predicted since the routing algorithm in OSPF chooses to use the higher bandwidth path through the MANET while OLSR's algorithm chooses the shortest hop count (low bandwidth). Therefore, OSPF is able to send data at 56 Kb/s while OLSR can only send at around 28.8 Kb/s. Using OLSR forces buffers to overflow at the input to the low bandwidth link and incurs large packet delays. Conversely, OSPF is able to use the high bandwidth MANET, and it can use the backup link whenever no route can be found across the MANET. Assigning variable link metrics is one of the key advantages of modifying a mature protocol such as OSPF to function efficiently in a MANET, and is critical for load balancing in many practical network scenarios including, for example, Naval afloat networks. However, we also note that OLSR could be extended to support variable link metrics.

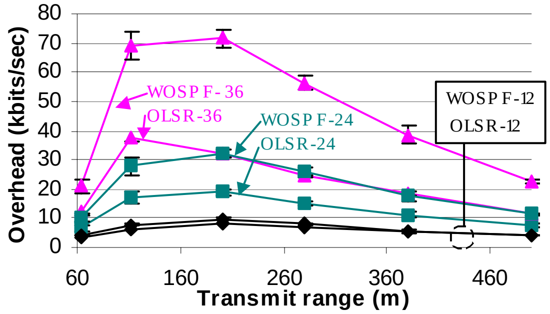

(a) IP overhead vs. mobile transmit range

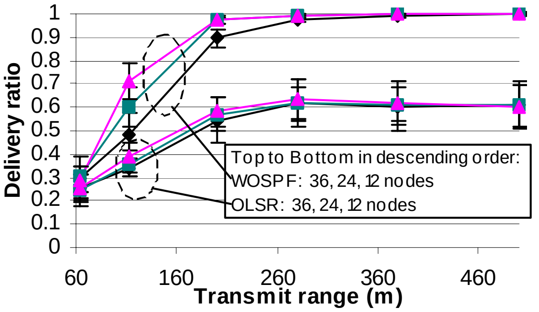

(b) Packet delivery ratio vs. mobile transmit range

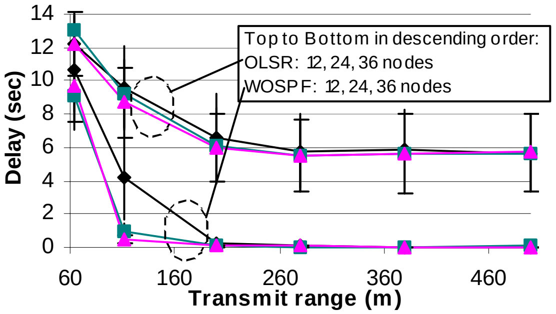

(c) End-to-end delay vs. mobile transmit range

Figure 6. Comparison of OLSR and Wireless OSPFv2 (WOSPF) performance in topology shown in Figure 5.

SUMMARY

We have described a wireless interface type for OSPF that incorporates the bandwidth-efficient mechanisms found in the OLSR MANET routing protocol. We found that, by running OSPFv2 in Point-to-Multipoint mode, our simulation network could not scale beyond sizes of roughly 20 nodes. Our modified OSPF routers approach the bandwidth efficiency and scaling properties of OLSR (using less than half of the bandwidth required by OSPF), while remaining fully compatible with legacy OSPF routers. We have also evaluated the possibility of using the OSPF Point-to-Multipoint interface type for wireless networks, including a technique to limit the number of adjacencies formed between neighboring routers, and found that the scaling of such an approach is inferior to OLSR's mechanisms. We believe that this development approach of incorporating OLSR features into OSPF has a faster transition path, and greater opportunity for interoperability and implementation reuse, than the alternate approach of developing OLSR to become a full-featured routing protocol.

The following topics require further study:

- Baker [4] has proposed additional mechanisms to aid in mobile networks, such as the extension of OSPFv3 mechanisms to support area mobility, and the creation of composite routing metrics that take into account additional factors relevant to wireless networks, such as willingness to route traffic and performance of the link. It may also be possible to include so-called "Hazy-Sighted" link state approaches [6] to improve scalability further.
- An alternative to running a MANET protocol as a layer-3 routing protocol is to run the MANET protocol as a layer-2 protocol, and to configure the layer-3 interface as a broadcast interface. Although such a configuration may lead to (possibly redundant) neighbor discovery operations at layer-2 and layer-3, a net benefit may be gained by reducing the LSA distribution overhead for large networks.
- The scaling performance of our approach is likely dominated by the requirements to disseminate "outside" LSAs into the wireless subnet. Further study is required to determine the operational limits of this approach for large networks, and whether additional mechanisms to improve scalability of outside LSA dissemination can be devised.

ACKNOWLEDGMENTS

This work was funded in part by ONR Contract N00014-02-C-0029 (Program Manager: Dr. Jerry Ferguson). We also thank Fred Baker, Phillipe Jacquet, and Joe Macker for helpful discussions on portions of this work.

REFERENCES

[1] Moy, J., "OSPF Version 2," RFC 2328, April 1998.

[2] T. Clausen, P. Jacquet, A. Laouiti, P Minet, P. Muhlethaler, A. Qayyum. "The Optimized Link State Routing Protocol," *Internet-Draft: draft-ietf-manet-olsr-07.txt (work in progress)*, November 2002.

[3] Wollman, W. and Y. Barsoum, "Overview of Open Shortest Path First, Version 2 (OSPF v2) Routing in the Tactical Environment," *Proceedings of 1995 IEEE MILCOM Conference*, vol. 3, pp. 925-930, October 1995.

[4] Baker, F., "An outsider's view of MANET," *Internet-Draft: draft-baker-manet-review-01.txt (expired)*, March 2002.

[5] Zavgren, J., Presentation at Panel Session on Ad-Hoc Networks, IEEE Milcom Conference, October 1997. Meeting notes available online at: "http://wnl.ece.cornell.edu/milcom_panel.html".

[6] Santivanez, C. et al. (2001), "Making Link-State Routing Scale for Ad Hoc Networks," Proceedings of the 2001 ACM International Symposium on Mobile Ad Hoc Networking and Computing, pp. 22-32, September 2001.