COMPARISON OF PROPOSED OSPF MANET EXTENSIONS

Phillip A. Spagnolo
and Thomas R. Henderson
The Boeing Company
P.O. Box 3707, MC 7L-49
Seattle, WA 98124-2207

Abstract— **The Open Shortest Path First (OSPF) routing protocol performs inefficiently when operated over certain types of mobile ad hoc networks (MANETs), such as those formed by using IEEE 802.11 ad hoc radios. In 2003, the Internet Engineering Task Force (IETF) OSPF working group solicited proposals to extend OSPFv3 for IPv6 to operate efficiently in MANET environments. During design team consideration, two proposals were developed and discussed: Overlapping Relays with Smart Peering (OR/SP) and MANET Designated Routers (MDRs). The two proposals both reduce OSPF overhead using similar ideas, but there are a few key differences. In this paper, we compare the design and operation of these two proposals, and use a simulation-based study to isolate several performance characteristics. Using the results of this comparison, we explain why we consider MDRs to be more suitable than OR/SP for OSPF in MANET environments.**

I. INTRODUCTION

Mobile Ad Hoc Networks (MANETs) are a class of computer networks characterized by the use of wireless communications and the presence of dynamically changing connectivity [1]. MANETs can operate in the absence of fixed infrastructure, but often are deployed as extensions to more stable network topologies.

One MANET use case of interest is the provision of Internet connectivity to mobile users beyond the radio range of a fixed network access point. There are a few possibilities for configuring a MANET in this case. One option is to configure the layer-2 (link or subnet) protocol to perform routing and bridging operations, such that the layer-3 (IP) protocol perceives the MANET to be a single multi-access link. This option has the benefit of shielding the rest of the network from changes due to mobility events, but can cause a loss of visibility of the underlying physical topology as observed by higher layer protocols, and requires support of MANET protocols at layer-2. Another option is to explicitly handle the MANET at layer-3, using routing and autoconfiguration protocols designed to handle the dynamic, multihop, partial mesh connectivity provided by the underlying radios. This option is useful if the underlying layer-2 protocol does not support routing or bridging[1] or if layer-3 is being used to interconnect many heterogeneous layer-2 technologies (e.g., multiple radio networks).

There are a number of possibilities for routing protocols used in layer-3 MANETs. The Internet Engineering Task Force (IETF) has specified a number of experimental MANET routing protocols such as OLSR [2]. These MANET routing protocols have focused on optimizing MANET performance and have not yet been extended to handle more heterogeneous networking environments such as found in larger enterprises. If a MANET is used in a transit networking scenario (i.e., other networks use the MANET as an intermediate network), routing information must be redistributed between the MANET routing protocol and other routing protocols. Redistribution is typically statically configured and may be lossy if one protocol is not capable of fully carrying another's data. The alternative solution for such deployments is to try to reuse an enterprise routing protocol such as Open Shortest Path First (OSPF) [3], [4], but previous experience with OSPF has shown that the protocol does not have suitable mechanisms for operating with high performance over broadcast-based, wireless multihop networks [5]–[7].

A natural question is whether an existing enterprise routing protocol could be extended to operate effectively over MANETs. In 2004, the OSPF Working Group at the Internet Engineering Task Force (IETF) considered a problem statement describing issues with integrating MANET extensions for OSPF [8] and chartered a design team to study alternatives for improving OSPF performance in such networks. During the course of the design team activities, the team and later the working group have actively considered two proposals:

- a proposal known as **Overlapping Relays**, contributed by a team at Cisco Systems [9], and extended with a proposal called **Smart Peering** [10], and
- a proposal known as **MANET Designated Routers**, contributed by Richard Ogier [11].

We have developed an experimental implementation of each proposal, with contributed code from both Cisco Systems and Richard Ogier, and conducted comparative simulation studies. Our implementation is an extension of the `quagga`[2] OSPFv3 routing daemon. To support both experimental and simulation studies, we have developed a wrapper for the `quagga` routing daemon for the Georgia Tech Network Simulator [12]; a packet-level, discrete-event simulator that supports detailed per-packet tracing, and simplified 802.11 or TDMA models.

We have previously published a technical report [13] that contains simulation results and discussion of the two proposals. This paper provides a summary and further results that compare more recent implementations of each approach.

II. OSPF AND MANETS

OSPF is a routing protocol that uses distributed computation and network map dissemination to compute shortest path

[1]e.g., most current IEEE 802.11a/b/g devices operating in IBSS mode

[2]http://www.quagga.net

routes through a network. Each router publishes a local map of its current one-hop neighborhood and floods this information throughout the network, generating a new local map whenever there is locally a topology change. By collecting current copies of the local maps as advertised by each router into a local database, a given router can assemble a global network map, and can run a shortest-path computation algorithm to determine a shortest path to each destination. If all of the databases in the network are synchronized, and each router executes the same path computation algorithm, then each router will generate a consistent set of forwarding tables that do not lead to forwarding loops.

Much of the complexity in OSPF derives from the careful protocol design to drive the network to a synchronized state. There are a few processes that collectively work on this. Any local map (or Link State Advertisement, LSA) is initially flooded throughout the network, to try to achieve the quickest dissemination of the information. However, such packet floods do not guarantee delivery to all nodes, so OSPF requires that each router retransmit its copy of the LSA until explicitly or implicitly acknowledged by all of its neighboring routers. Further, when two routers discover each other, they create what is known as an "adjacency," using a protocol that ensures that their routing databases become synchronized. Retransmissions occur only along these adjacencies, which are bidirectional and pairwise between neighboring nodes.

A. Limitations of OSPF in MANET environments

There are two basic options for adapting OSPF to a layer-3 MANET routing protocol: i) adapt the broadcast-based OSPF interface type to allow for the partial mesh of connectivity that can result in a MANET, or ii) adapt the Point-to-MultiPoint OSPF interface type to operate more efficiently in a MANET. The first approach requires a modification of the designated router election protocol to handle partial mesh connectivity, and is not further discussed herein. We focus instead on the second approach.

The OSPF Point-to-MultiPoint interface allows for multiple point-to-point OSPF links to be aggregated on one physical interface. OSPF Point-to-MultiPoint can operate correctly over MANETs if the protocol is allowed to forward (reflood) LSAs out of the same interface from which they have been received, but the protocol operates inefficiently especially in dense network scenarios. One easy modification is to allow acknowledgments for LSAs to be multicast to all one-hop neighbors, but more is needed. Although there are a variety of optimizations to the Hello protocol and database description procedures that are possible for MANET, our previous results [14] indicated that flooding overhead was the primary contributor to OSPF overhead growth in MANETs. Therefore, the OSPF MANET design team at the IETF has focused on improving OSPF's flooding performance as a first step towards extending OSPF to work better over MANETs.

B. OSPF overhead reduction

Most of the OSPF overhead observed in layer-3 MANETs can be traced to the flooding and retransmission of LSAs.

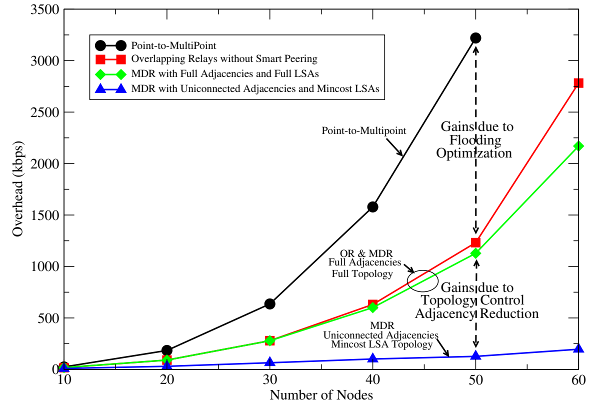

Fig. 1. Overhead reduction due to flooding reduction, topology control, and adjacency control (mobile simulations described in [13])

Three basic strategies can be combined to reduce this overhead:

- **flooding reduction:** Reduce the number of transmissions required to broadcast the LSA;
- **adjacency reduction:** Reduce the number of adjacency relationships maintained by the protocol, without degrading synchronization of the network; and
- **topology reduction:** Reduce the number of available links (neighbors) advertised by routers without reducing the data plane forwarding performance.

Our previous results have indicated that these are complementary objectives and that, when combined, significant overhead savings can be realized without compromising performance. Figure 1 summarizes some simulation results described in more detail in [13], for a 50-node MANET running OSPF modifications over 802.11b radios. The data suggests that improved flooding performance can be complemented by the use of both adjacency and topology reduction. In this paper, we explore this result in more detail by contrasting the design and performance of Overlapping Relays and MANET Designated Routers.

C. Overview of Overlapping Relays

Overlapping Relays (ORs) is a proposal from Cisco Systems to adapt the optimized flooding technique developed for OLSR [2] to function in the reliable, acknowledged flooding framework of OSPF. The basic proposal known as Overlapping Relays is defined in [9], and extended with a proposal called **Smart Peering (SP)** [10]. We abbreviate this combined approach as OR/SP herein. OR/SP includes support for a differential Hello mechanism and the use of LSA floods to serve as acknowledgments to upstream nodes; herein, we focus on the following three components: i) **flooding reduction**, ii) **adjacency reduction**, and iii) **topology reduction**.

The main *flooding reduction* is adapted from OLSR's approach of advertising 2-hop neighborhood (in this case, advertised in a router-LSA) to all 1-hop neighboring routers, thereby allowing each router to independently select a subset of its 1-hop neighbors to serve as ORs, or multi-point relays (MPRs), for LSAs that it originates or refloods. In terms of graph theory, the set of ORs forms a neighbor-selected, source-dependent connected dominating set (CDS) of the 2-hop neighborhood. The heuristic used for OR selection can be varied to allow for certain nodes to be preferentially selected on the basis of OR set stability, link quality, or redundancy. Note that OR relationships are not necessarily symmetrical.

The main flooding reduction is augmented by "non-active" Overlapping Relays. Although, in OSPF, LSAs sent via the flooding mechanism are retransmitted (along OSPF adjacencies) if the downstream node fails to acknowledge the flooded message in a timely manner, non-active ORs allow neighboring nodes that are not ORs to more proactively retransmit an LSA if they detect that the recipient has not yet received it. All 1-hop neighbors that are not selected as ORs can serve as non-active ORs. To avoid redundant retransmissions, non-active ORs must therefore activate before the expiration of the sender's retransmission timer, and must coordinate among themselves to avoid too many duplicate transmissions. Non-active ORs are an optimization and are not required for synchronization.

Adjacency reduction was added to the original OR proposal by an extension known as Smart Peering (SP). A goal of adjacency reduction is to reduce the number of OSPF adjacencies required while still maintaining a high probability of network connectivity via the set of adjacencies, thereby preserving database synchronization. SP works by allowing a router to consider whether to bring up a new adjacency based on reachability among the set of existing adjacencies. If a candidate neighbor is reachable via a path of adjacencies, a direct adjacency to that neighbor can be suppressed. A shortest path tree (SPT) computation via the existing LSA database can be used to determine such reachability. Various heuristics can be applied to determine the local density of adjacencies that a router will strive to maintain, as well as how close (in hops) a neighbor must be reachable via existing adjacencies (i.e., to reduce the diameter or "stretch" of the control plane).

Adjacencies suppressed in this manner are called "unsynchronized" adjacencies, although the term may be a misnomer because the neighbors are synchronized (just not directly), and there is not an adjacency in the OSPF sense (in particular, retransmissions do not occur along these neighbor relationships). Ogier first introduced a similar concept called "routable neighbor" prior to the SP proposal [11].

The OR/SP proposal specifies adjacency reduction as stated above; however, there is another method that has been independently suggested. Given that ORs exist among a subset of the one-hop neighbors of a node, one may instead elect to form adjacencies among only ORs and their selectors (and a few special nodes needed to guarantee connectivity) [15]. However, since OR and selector relationships are not necessarily symmetric, it has been suggested [16] that the number of adjacencies and rate of adjacency formation would exceed that of the SP approach.

Topology reduction can occur with adjacency reduction if the router elects to only include adjacent neighbors in its router LSAs. However, this may severely constrain the available links in the forwarding table, causing data plane forwarding stretch. To mitigate this, a router is allowed to advertise unsynchronized adjacencies as links in its router LSAs, although the technique for selecting such neighbors is unspecified. Readers may note, however, that this introduces a circular dependency in the above adjacency reduction process; if the LSAs are used to determine reachability via adjacencies, but the LSAs contain also unsynchronized adjacencies, the router may not have enough information to determine control plane reachability. Solving this requires that these unsynchronized adjacencies be specially marked in LSAs such that they can be excluded from the SPT computation described above but used in the data plane SPT. The draft defines a "U"-bit in the LSA for this purpose; the U-bit is ignored by legacy routers.

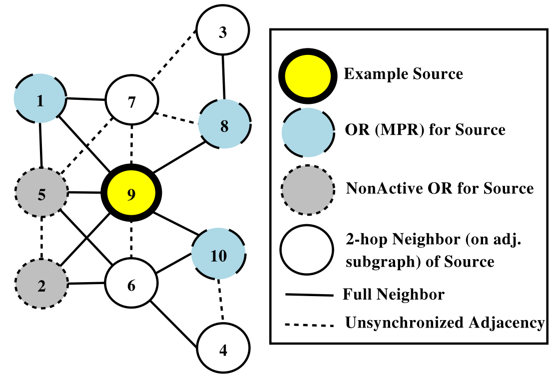

Fig. 2. Overview of Overlapping Relays

In summary, OR/SP uses ORs (multipoint relays) augmented by non-active Overlapping Relays to improve flooding reduction, uses a second SPT computation and modified LSAs to compute a reduced control plane (adjacency) topology while still maintaining connectivity, and allows for advertisement of non-adjacent neighbors (although the algorithm for selection of which non-adjacent neighbors to advertise is unspecified). Figure 2 provides an overview of OR operation in which the node numbering corresponds to OSPF Router ID. Assume that Node 9 is the flooding source under consideration. The SP algorithm will select adjacencies depending on the order in which nodes come up; in this example, assume that the solid lines are the resulting adjacencies. First, note that a node may be two or more hops away from a physical neighbor on the adjacency subgraph (e.g., nodes 9 and 7). Node 9 selects 1, 8, and 10 as its ORs so that the 2-hop neighborhood on the adjacency subgraph is covered. Node 9 may also advertise nodes 7 and 6 as unsynchronized adjacencies.

D. Overview of MANET Designated Routers

In many respects, the OSPF MANET Designated Router (MDR) proposal [11] is similar to OR/SP. Both approaches use a connected dominating set to reduce the number of initial floods of a new LSA through the network, both provide a strategy for adjacency reduction, and both allow for the advertisement of links to non-adjacent neighbors in the router LSAs. We focus herein on identifying the differences in the following key components: i) **flooding reduction**, ii) **adjacency reduction**, and iii) **topology reduction**.

For *flooding reduction*, in contrast to the MPR approach used by OR/SP, MDR-capable routers run a distributed al-

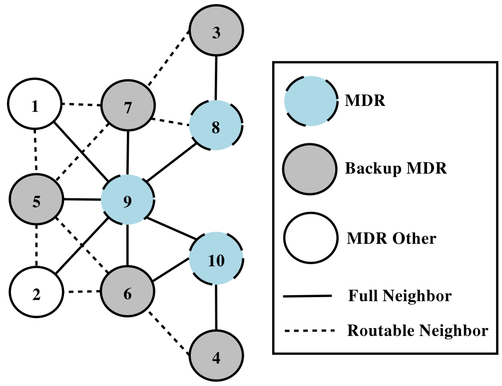

Fig. 3. Overview of MANET Designated Routers

gorithm that determines whether the router is included in a single, shared, connected dominating set (CDS) that is used for forwarding all LSA floods, regardless of the previous hop of the message. That is, it is a self-selected, source-independent CDS. The draft provides an algorithm that can be run, in a distributed fashion, to select a set of MDRs and backup MDRs (BMDRs) that form the backbone. Conceptually, this approach generalizes the notion of OSPF's Designated Router (DR) and Backup DR used on the broadcast interface, although no network-LSA is originated (router-LSAs as in a Point-to-MultiPoint mode are originated). MDRs and BMDRs are responsible for reflooding all LSAs; MDRs provide the initial flood, and (analogous to non-active Overlapping Relays) BMDRs may back up the MDR flooding by retransmitting before the normal OSPF adjacency retransmission occurs. The set of links between MDRs forms a CDS, and between the set of MDRs and BMDRs, a biconnected CDS. Hello messages are used to convey the 2-hop neighbor information necessary for MDR/BMDR election.

Adjacency reduction is accomplished by limiting adjacencies to the MDR/BMDR backbone, and allowing all other routers to form at least two adjacencies with the backbone, similar to how the OSPF broadcast interface works. This ensures that the control plane is biconnected and that any router is either on the backbone or is one-hop away from it, although two neighboring routers may be multiple hops away from each other in the control plane.

The set of adjacencies described above is a much sparser representation of the underlying physical network, for dense topologies. Advertised *topology reduction* could again just be limited to advertising adjacencies, but such an approach would incur significant routing stretch. Ogier introduced the concept of a "routable neighbor", which was adapted by Cisco Systems in the form of the "unsynchronized adjacency" described in the OR/SP approach previously. In Ogier's approach, a router advertises all of its adjacencies (thereby guaranteeing reachability in the data plane) and a subset of its other (routable) neighbors. A number of algorithms are specified for selecting these routable neighbors to advertise, including algorithms that guarantee that the resulting partial-topology LSAs provide min-hop or min-cost paths in the data plane.

In summary, MDRs use a source-independent CDS augmented by backup MDRs to improve flooding reduction, uses this same flooding backbone as the framework for reducing the control plane (adjacency) topology while still maintaining connectivity, and allows for advertisement of non-adjacent neighbors while specifying several algorithms for their selection. Figure 3 provides an overview of MDR operation. Since MDRs are source-independent (in contrast to ORs), an example source is not shown as in Figure 2. For equal comparison with ORs, the figure shows MDRs' optional uniconnected adjacency backbone. The node numbers correspond to Router-IDs in OSPF. Note that the formation of adjacencies is time dependent, so the adjacencies will vary based on when neighbor relationships are formed.

III. Comparison

As mentioned above, both the OR/SP and MDR proposals bear a lot of similarity, with respect to their goals and high-level designs. What are the main differences between them, and are they significant? Using reference research implementations of both proposals, we have conducted selected simulations of both approaches in a basic MANET setting. We configured a number of variants of the proposal, looking to isolate performance characteristics of different aspects of the respective designs. We have studied the core performance (routing protocol overhead and data delivery ratio) as well as additional statistics such as the size and stability of the adjacency subgraph, the stability of the forwarding table, and stretch factors (control and data plane). In this section, we summarize some key findings. Space precludes a full treatment of these results, but we have made our simulation code and scripts available to the interested reader.[3]

A. Methodology

We briefly introduce the simulation methodology, which is described more thoroughly in [13]. We developed extensions to the `quagga` open source routing implementation of OSPFv3. Further, we have instrumented the Georgia Tech Network Simulator (*GTNetS*) discrete-event network simulator to allow our `quagga` extension to run also in simulation. This extension has been developed with the active participation of both OSPF MANET draft authors (Richard Ogier and staff from Cisco Systems) and basic validation has been performed by studying trace files of small sample networks instrumented to exercise portions of the code.

Below, we summarize results from running the OSPF MANET extensions over a simulated, single-channel 802.11b (IBSS) network, in which all nodes used a single broadcast-capable radio interface. Nodes were configured to roam according to a random waypoint mobility model; specifically, each node was configured to move to a random point on a square grid at a fixed velocity chosen from a uniform distribution between zero and `velocity_max`, pause for a fixed amount of time, and then choose a new random location and move again.

In the below, we consider simulations of the following configurations:

[3]http://hipserver.mct.phantomworks.org/ietf/ospf

- **MDR.full.full** MDR backbone using full adjacencies (i.e., disabling the adjacency reduction component of MDR) and reporting full topology LSAs. This configuration is not defined in [11];
- **MDR.uni.full** MDR backbone using unconnected adjacencies and reporting full topology LSAs. This variant uses the minimal set of required adjacencies, and it isolates the effect of reducing adjacencies alone;
- **MDR.uni.min** MDR backbone using unconnected adjacencies and reporting minimal (only adjacent neighbors) LSAs. This variant provides the lowest amount of advertised topology in the OSPF MDR framework;
- **OR.full.full** ORs with full adjacencies (i.e., no SP) and full topology LSAs *(compare with MDR.full.full because only flooding is optmized)*;
- **OR.sp.full** ORs with SP-formed adjacencies and full topology LSAs *(compare with MDR.uni.full because both form minimal adjacency connectivity)*; and
- **OR.sp.min** ORs with SP-formed adjacencies and minimal topology LSAs (only the adjacencies) *(compare with MDR.uni.min because minimal adjacencies are formed and minimal LSAs are created)*.

The above set of configurations allow us to isolate the effects of different aspects of the OSPF MANET proposals, and to permit fairer comparisons between the two approaches. However, MDRs are not limited to these configurations. Two other possible configurations are listed here:

- **MDR.bi.mincost** MDR backbone using biconnected adjacencies for redundancy in the flooding plane and partial-topology LSAs providing min-cost paths (a recommended configuration for OSPF MDR by [11]);
- **MDR.uni.mincost** MDR backbone using unconnected adjacencies and partial-topology LSAs providing min-cost paths;

Our results below are excerpted from a set of simulation runs that varied the following parameters:

- **number of nodes.** We varied the number of nodes between the values of 20, 50, and 70.
- **density.** In a given configuration, we varied the neighbor density primarily by varying each node's radio range. This yielded densities of roughly 1 to 40 neighbors/nodes (0.06 to 0.8 neighbors/total number of nodes).
- **topology change rate.** We primarily looked at two configurations. For a given number of nodes and density, the change rate is most directly influenced by the mobility parameters (velocity, pause time) and radio model (communication gray zone, controlled by "alpha" parameter). Our two mobility configurations yielded average 2-Way neighbor lifetimes of roughly 50 seconds ("high" mobility) and 90 seconds (lower mobility).

In evaluating a routing protocol's performance, there are a number of statistics that can be considered, and some have more relevance in certain environments than others. For the purposes of this study, we have focused on the parameters described in the sections below, which we believe provide a sufficient first-order assessment of how well the protocol is performing. Space precludes including numerical results for all cases, so we have tried to include the most interesting ones for one scenario (50 nodes, high mobility) and just describe the others.

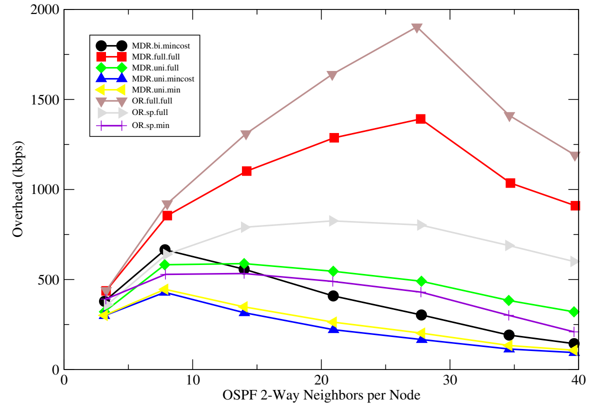

Fig. 4. Overhead as a function of physical neighbor density (50 nodes, high mobility)

B. Core performance

1) Overhead: In many MANETs, the wireless transmission capacity is bandwidth and/or interference limited, so reducing the number of transmissions devoted to control traffic is critical. Here, we consider just the raw overhead generated, not the "total overhead" as defined by Santivanez [17], since we separately consider route stretch below.

Figure 4 contrasts the overhead yielded by the various approaches, as a function of neighbor density for a mobile, 50 node network. We observed that full adjacency and full topology variants generated the most overhead, followed by full topology variants of configurations with thinned adjacencies. The best performance was obtained by the MDR variants we tested with partial topology LSAs. MDRs consistently performed better than ORs when comparing similar configurations of topology and adjacency reduction. The improvement in MDRs was primarily due to a large number of refloods by non-active ORs.

2) Delivery Ratio: This statistic (measure of the fraction of data packets delivered in the network) conveys how well the protocol finds good routes for user data. The statistics are somewhat dependent upon the traffic matrix; we have assumed low amounts of uniformly distributed traffic herein. Reduction of OSPF overhead is practically limited by the need to provide working routes in the network. As we have shown in previous work [13], in general, OSPF MANET modifications do not significantly impact the user data delivery ratio. Similarly, our simulation runs herein do not show an appreciable difference between the performance of the approaches and that of normal OSPF, except that for high mobility scenarios, we found that the OR/SP with minimal LSAs performed a bit worse than the rest, because of the forwarding stretch we discuss below.

C. Adjacency statistics

1) Adjacencies per node: We measured how many adjacencies, on average, each node formed, as a function of the number of physical neighbors. Figure 5 shows that for

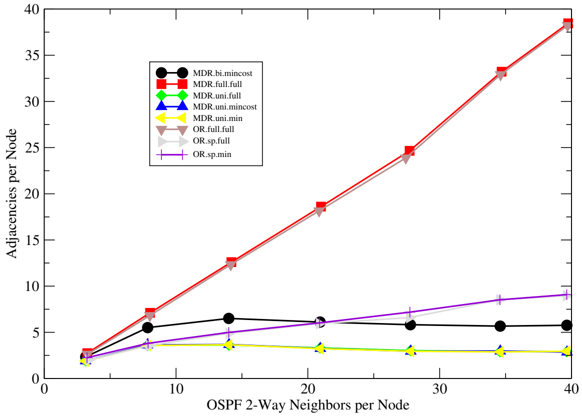

Fig. 5. Adjacencies per node as a function of physical neighbor density (50 nodes, high mobility)

full adjacency variants of OSPF MANET, the number of adjacencies rises almost linearly with the number of physical neighbors, as expected. MDR variants are able to control the number of adjacencies very well as the number of physical neighbors increases; the effect of uni- vs. biconnected adjacencies is clearly seen in this plot. OR/SP variants exhibit low numbers of adjacencies as well, but the curves depart from the comparable MDR unconnected case as the neighbor density grows; we believe that this is because topology changes in OR/SP network may result in uncoordinated additions of new adjacencies by several nodes when there is a detected need to repair the connectivity of the adjacency subgraph, and this probability increases with increased neighbor density.

2) Stability of adjacency subgraph: One question that has arisen with respect to OR/SP is whether the adjacency subgraph is stable or changes more frequently than that of MDRs. Our results suggest that there is not much difference in this measure between the two approaches; the rate of change is stable as the neighbor density increases, and the comparable curves for OR/SP and MDR (uniconnected) overlap.

D. Stability of forwarding tables

1) Rate of LSA generation: The rate of LSA generation is highly correlated to the overhead observed in the MANET, but there is an important distinction in OSPF. Since we are motivated to use OSPF MANET for integrating MANET subnetworks with a larger OSPF network, we may in many cases be concerned about the total overhead generated and flooded outside of the MANET subnet. Although it may be possible to apply OSPF summarization techniques (area hierarchy) to minimize this impact, the question of whether such additional approaches are required or not must be considered. This parameter is also an indirect measure of the stability of the forwarding tables.

We found that, in our high mobility scenarios, the LSA generation rate (per node) bumped up against the OSPF architectural constant `MinLSInterval` of five seconds. This means that the MANET LSA generation was saturated, and higher mobility would tend to result in worse convergence but perhaps not much more overhead. Using OSPF MANET techniques, we were able to reduce the frequency with which routers advertised new LSAs. All of the MANET approaches were able to rise above the five second limit, and all exhibited better performance as the physical neighbor density grew. The best performing (most stable) variant that we tested was MDR uniconnected adjacencies with min-cost LSAs.

2) Stability of unicast routes: We measured the number of route changes per node per second. The worst performing variants were OR/SP and MDR using full adjacencies and full topologies, as expected (more advertised topology leads to more possible changes). Of the reduced adjacency variants, which were more stable by at least a factor of two over the full adjacency variants, those variants also with reduced topology LSAs consistently performed the best.

E. Stretch factors

1) Control plane (flooding) stretch: Flooding stretch considers how many hops are necessary to disseminate routing (LSA) updates. While the flooding path of LSA updates through a network is affected by various factors, approaches that limit the topology of the flooding subgraph (as both OSPF MANET approaches do, for efficiency purposes) run the risk of lengthening the flooding path of some updates.

We initially found OR/SP to perform much more poorly than MDR variants, for the following reason. SP controls the formation of adjacencies in the network, and MPRs must be elected from adjacent neighbors. SP therefore introduced a sparser subgraph for MPR election, and consequently, one of the properties of MPR flooding (flooding along min-hop paths) was eliminated and flooding paths actually could become several hops longer than minimum. We were able to fix this by allowing nodes to "overhear" flooding transmissions of nodes that are not MPR selectors, leading to flooding stretch results comparable to those of legacy OSPF and OSPF MDR. All of our simulation results use this modified version of OR/SP.

2) Data plane (forwarding) stretch: Beyond data delivery ratio, the ability of the routing protocol to find efficient (low-cost) routes is important. Techniques that limit the routing topology (as both OSPF MANET approaches can perform, for overhead scalability improvement) run the risk of degrading the path lengths in the data plane. Figure 6 illustrates one measure of forwarding stretch; the ratio of the sum of data packet forward and receive events to the number of receive events. A lower value of this statistic is better (reduced stretch).

As Figure 6 shows, the data plane route stretch is comparable for all approaches except one: the OR/SP case with minimal LSAs. This is because SP does not control the diameter of the underlying adjacency subgraph, and advertised links. In contrast, MDR with minimal LSAs can achieve near-optimal forwarding stretch since the CDS is at most one-hop worse than the shortest path. This resulted in the poorer data delivery ratio for SP and minimal LSAs also mentioned above. We hypothesize that partial topology techniques such as Ogier's min-cost LSAs could be added to SP, to improve this performance metric. In the low density scenarios, the variations between the approaches are not due to longer routes,

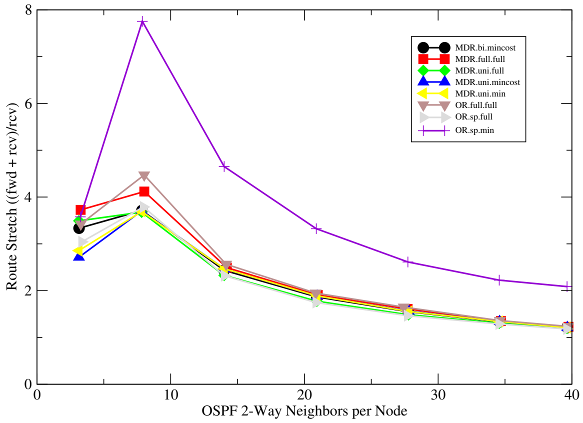

Fig. 6. Data plane forwarding stretch as a function of physical neighbor density (50 nodes, high mobility)

but they are due to some scenarios generating more events where packets are sent and then forwarded but not received.

IV. Conclusions

Our recent simulation results suggest that, to first order, we are able to achieve similar performance with variants of both the OR/SP and MDR proposals in many instances. We believe that each can be tuned to approximate the performance of the other, within reasonable bounds; however, we consistently observed that MDR produces less overhead than a comparably configured OR/SP implementation. The decision between the two approaches should also rest on issues such as implementation complexity and architectural fit with OSPF. Here, there seem to be a few key differences:

1) MDR provides an integrated approach to the problem of flooding efficiency, topology reduction, and adjacency reduction. OR/SP, however, has a separate solution for flooding reduction and adjacency reduction, and probably needs to add an approach for topology reduction to perform adequately. Even though OR/SP would likely need to combine disparate pieces to come closer to MDR performance, one could argue that modularity is not a bad feature of the design, unless it introduces unnecessary complexity.
2) SP has no framework for controlling the graph structure of the adjacency subgraph. As mentioned above, MDR ensures that every node lies on the adjacency backbone or is one-hop away. The SP adjacency subgraph is more arbitrary, depending on the order of events such as neighbor discovery, and leads to considerable forwarding stretch unless additional techniques are used. Furthermore, the loss of connectivity of the adjacency subgraph may trigger several uncoordinated attempts to repair the connectivity from nodes many hops away, leading to more adjacencies than necessary. MDR is less prone to this since nodes can make consistent decisions and know their roles in the topology.
3) MDR seems to provide more flexibility in tuning both the redundancy of adjacencies in the control plane and data plane topology advertised. It is possible that these features (such as biconnected SP) could be added to OR/SP, but it is not clear how much more the complexity might rise or the performance might change.

Work is ongoing within the OSPF working group at IETF to converge on an experimental standard for MANET extensions.

Acknowledgments

The authors would like to especially thank Richard Ogier for many detailed discussions and insights and for providing and testing a substantial part of the MDR implementation, Stan Ratliff and Cisco Systems for providing SP implementation code and participating in the validation and discussion of those results, and Acee Lindem, Abhay Roy, Philippe Jacquet, and Emmanuel Baccelli for technical discussions. Gary Pei contributed modifications of the *GTNetS* 802.11 model. Finally, we thank Fred Templin, Gary Pei, and Ian Chakeres who commented on earlier draft versions. This work was performed under Boeing IR&D Mobile Networking Technology (program manager Jae Kim), who provided helpful comments. This work derives from earlier research funded by ONR under contract N00014-04-C-0003, who approved this manuscript for public release (distribution unlimited).

References

[1] S. Corson and J. Macker, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations," RFC 2501 (Informational), Jan. 1999.
[2] T. Clausen and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)," RFC 3626 (Experimental), Oct. 2003.
[3] J. Moy, "OSPF Version 2," RFC 2328 (Standard), Apr. 1998.
[4] R. Coltun, D. Ferguson, and J. Moy, "OSPF for IPv6," Request for Comments 2740, IETF, December 1999.
[5] W. Wollman and Y. Barsoum, "Overview of Open Shortest Path First, Version 2 (OSPFv2) Routing in the Tactical Environment," in *Proceedings – IEEE Military Communications Conference MILCOM*. Nov. 1995, vol. 3, pp. 925 – 930, IEEE.
[6] P. Sass, "Communications Networks for the Force XXI Digitized Battlefield," *Mob. Netw. Appl.*, vol. 4, no. 3, pp. 139–155, 1999.
[7] T. Henderson, P. Spagnolo, and J. Kim, "A Wireless Interface Type for OSPF," in *Proceedings – IEEE Military Communications Conference MILCOM*. Oct. 2003, vol. 2, pp. 1256 – 1261, IEEE.
[8] F. Baker, "Problem Statement for OSPF Extensions for Mobile Ad Hoc Routing," Internet-Draft (work in progress) draft-baker-manet-ospf-problem-statement-00, IETF, August 2003.
[9] M. Chandra, "Extensions to OSPF to Support Mobile Ad Hoc Networking," Internet-Draft (work in progress) draft-chandra-ospf-manet-ext-03, IETF, April 2005.
[10] A. Roy, "Adjacency Reduction in OSPF using SPT Reachability," Internet-Draft (work in progress) draft-roy-ospf-smart-peering-01, IETF, November 2005.
[11] R. Ogier and P. Spagnolo, "MANET Extension of OSPF using CDS Flooding," Internet-Draft (work in progress) draft-ogier-manet-ospf-extension-07, IETF, March 2006.
[12] George F. Riley, "Large-scale Network Simulations with GTNetS," in *Winter Simulation Conference*, David M. Ferrin and Douglas J. Morrice, Eds. 2003, pp. 676–684, ACM.
[13] T. Henderson, P. Spagnolo, and G. Pei, "Evaluation of OSPF MANET Extensions," Technical Report D950-10897-1, Boeing Phantom Works http://hipserver.mct.phantomworks.org/ietf/ospf, July 2005.
[14] P. Spagnolo T. Henderson and G. Pei, "Design Considerations for a Wireless OSPF Interface," Internet-Draft (work in progress) draft-spagnolo-manet-ospf-design-00, IETF, April 2004.
[15] E. Baccelli, T. Clausen, and P. Jacquet, "OSPF MPR Extension for Ad Hoc Networks," Internet-Draft (work in progress) draft-baccelli-ospf-mpr-ext-01, IETF, March 2006.
[16] R. Ogier, "Advantages of OSPF-MDR," Internet-Draft (work in progress) draft-ogier-ospf-mdr-position-00, IETF, February 2006.
[17] C. Santivanez, R. Ramanathan, and I. Stavrakakis, "Making Link-State Routing Scale for Ad Hoc Networks," in *Proceedings – ACM MobiHOC Conference*, Oct. 2001.