Analysis of OSPF MANET IETF-64 Scenarios

Boeing Technical Report D950-10915-1

Phil Spagnolo and Tom Henderson
Boeing Phantom Works
Seattle, WA 98124

March 3, 2006

phillip.a.spagnolo@boeing.com, thomas.r.henderson@boeing.com,

Abstract

The IETF OSPF WG is evaluating two proposals for OSPF MANET extensions. At IETF-64, Stan Ratliff from Cisco Systems presented data that showed Overlapping Relays with Smart Peering yielded lower overhead than MANET Designated Routers in sparse, highly mobile networks. Since that meeting, we have worked with Cisco to try to replicate and explain the reported results. The IETF-64 results have since been retracted by Cisco. This memo goes into further detail to explain the discrepancies and to present more current results for the scenarios of interest.

1. Introduction

This report outlines our attempt to validate OSPF-MANET results presented by Stan Ratliff from Cisco Systems at the 64th IETF meeting (IETF-64), [1]. The results presented showed that Overlapping Relays (OR)s [2], with Smart Peering (SP) [3], had substantially less overhead than MANET Designated Routers (MDR)s [4], in highly mobile sparse networks. At the meeting and on the OSPF-MANET list, the validity of these results was questioned, but due to the short notice by Cisco before releasing the results at the meeting, there had been no independent confirmation of the results. As a result of this further work, we and the original author came to the conclusion that the previously reported results were not correct and should be retracted [5]. The technical report begins with a presentation of the validation procedure and then results are presented followed by some further analysis.

2. Validation Procedure

The first step in validation was to port Cisco's Smart Peering (SP) code to Boeing's GTNetS code base [6]. Stan Ratliff was very helpful in providing the code and in working through the debugging process.

After the port was completed, we changed certain simulation parameters from the defaults to match those presented at IETF-64. Specifically, the velocity was set to 16 m/s, the pause time was set to 1 sec, the Link State Acknowledgment (LSAck) caching time was set to 100 seconds, and the maximum radio range[1] was varied from 40 to 300 meters in 20 meter increments.

Initial results, did not match those presented at IETF-64, so we began to debug the code. During the debugging process we found the following problems in the OR and SP code, and we explain how they were resolved.

1. OR bug fix: The first hop 2-Way neighbors were being used in the selection of MPRs. This problem did not appear until SP was developed because before that a 2-way link became FULL very quick, but with SP 2-way links may never become FULL. The OR (MPR) selection algorithm now only use Full neighbors for OR selection.
2. Section 4.4 of the SP draft, "Overlapping Relay [OR] election impact," was previously not implemented. It has now been implemented with the method in which flooding is not performed over non-adjacent links. Therefore, the MPR selection algorithm only uses adjacent links for one- and two-hop neighbors.

[1] As explained in [6], the true radio range is probabilistic based on a parameter that controls the dropoff in packet reception probability as a function of distance. The parameter described here sets the maximum range beyond which the probability of reception is zero.

3. SP bug fix: Router-LSAs were not being originated as often as necessary because routability of unsynchronized adjacencies was being checked against the unsynchronized SPF table. The synchronized (routing on adjacencies only) SPF table is now used to determine when a router-LSA should be originated.
4. SP bug fix: The synchronized SPF table was being built using links in its own router-LSA that were unsynchronized. The SPF calculation was changed to use only synchronized links when the synchronized SPF table is calculated.
5. SP bug fix: Each time an LSA was installed with `ospf6_spf_install()` the neighbors were checked for routability and adjacencies where updated. However, it was found that this method was not working because the SPF was not being run prior to checking for routability. Therefore, the routability check and the adjacency update were moved to instead occur after each SPF calculation.

In addition to these changes Cisco suggested that the following optimization be made for SP.

6. If a new router-LSA is originated due to a neighbor down event, then the router-LSA is not flooded until after either any necessary adjacencies are formed or the minLSInterval has passed. This corrects the possible problem that a neighbor down event causes an LSA flood, but then the OSPF router has to wait MinLSInterval before announcing any new adjacencies (formed from previously unsynchronized neighbors).

After these changes, we again performed the simulations, but again we did not obtain results consistent with those reported in [1]. During discussions with Stan Ratliff, we found that in addition to changing the velocity, pause time, and LSAck cache times, Cisco had also changed the AckInterval (0.5 sec), PushbackInterval (7 sec), and RxmtInterval (9 sec). We note here that a PushbackInterval of 7 sec does not conform to Section 3.4.30 of the OR draft, [2], (PushbackInterval + propagation delay < RxmtInterval/2). However, Cisco has found that this value yielded better results. In our implementation, a PushbackInterval of 7 seconds means that the range of Pushback expiration is 7 to 14 seconds; in effect, this change results in the Overlapping Relays aspect of the protocol being significantly deemphasized, with most recovery from flooding loss occurring due to retransmissions instead.

3. Results

We now present some new results obtained with the code changes described in Section 2. In the GTNetS release[2], we ran these simulations with scripts `comp_density_or.pl` and `comp_density_mdr.pl`. We list key simulation parameters in Table 1. We tried to maintain similar timer values in both proposals, so that the proposals could be evenly compared. The PushbackInterval and BackupWaitInterval are the only dissimilar parameters because they are used differently in the respective proposals. Note that we chose a different RxmtInterval, PushbackInterval, and BackupWaitInterval than in the Cisco simulations, for similar comparisons. However, we have data (not shown) that indicates the same trends as in the original Cisco parameters. We also note that the PushbackInterval is also not strictly in compliance with the OR draft because the range extends from 3.5 to 7 seconds which is not less then half of the RxmtInterval (7).

[2] http://hipserver.mct.phantomworks.org/ietf/ospf/releases/060220_release/gtnets_060220.tgz

Parameter	Value	Range (if applicable)
Grid length	500 m^2	N/A
Radio Range	R	R = 40,60,80,…,300
Number of Nodes	N	N = 20,30,40,50
pause_time	1 sec	N/A
velocity	16 m/sec	Uniform(0,16)
HelloInterval	2 sec	N/A
DeadInterval	6 sec	N/A
RxmtInterval	7 sec	6.9 to 7 sec
MinLSInterval	5 sec	N/A
MinLSArrival	1 sec	N/A
AckInterval	0.5 sec	0 to 0.5 sec
PushbackInterval	3.5 sec	3.5 to 7 sec
BackupWaitInterval	0.5 sec	0.5 to 1 sec

Table 1: Simulation parameters used

We ran a multicast-capable Point-to-Multipoint (PTMP) interface type as a baseline case, and three different OSPF-MANET variants:

- OR/SP with Full LSAs
- MDR with Uni-connected adjacencies and Full LSAs
- MDR with Bi-connected adjacencies and Full LSAs

The term "full LSAs" refers to the practice of including all adjacent and routable neighbors as Type-1 links in the router LSAs. Routable neighbors are those non-adjacent neighbors that have reached at least state 2-Way and for which a synchronized path (via adjacencies) exists or has recently existed. We note that the preferred configuration of MDRs by Ogier is to use topology-reduced LSAs (such as min-cost LSAs), but we use Full LSAs here for a more equal comparison with OR/SP using Full LSAs. Strictly speaking, the most similar configuration and comparison of the two approaches is between OR/SP and uni-MDR curves; the bi-MDR curves use a bi-connected MDR backbone for redundancy. The results in each case are for a single simulation run of 1800 seconds, with an initial startup period of 1800 seconds that is discarded from our data.

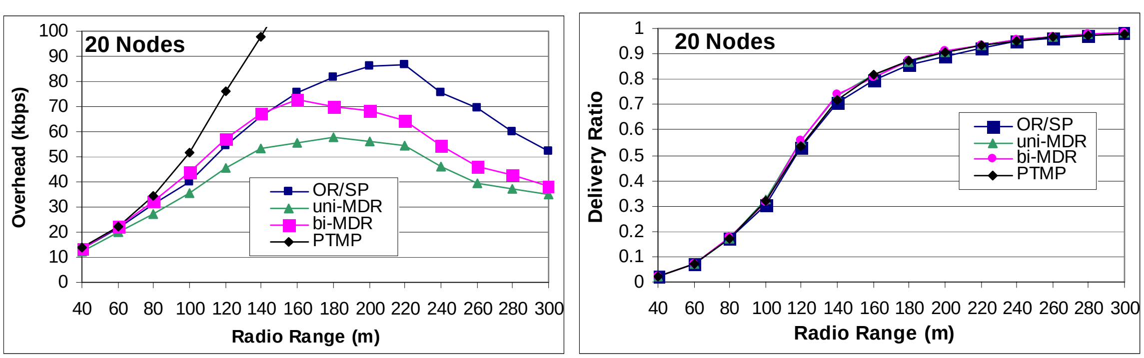

Figure 1: OSPF-MANET comparison for 20 nodes

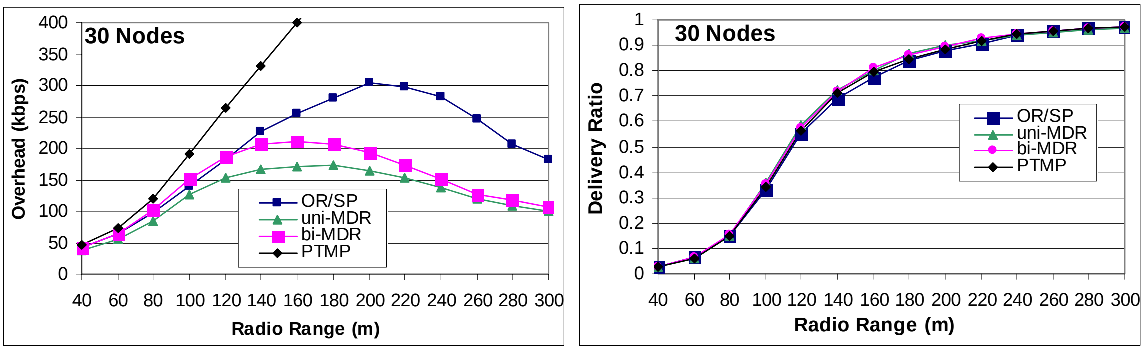

Figure 2: OSPF-MANET comparison for 30 nodes

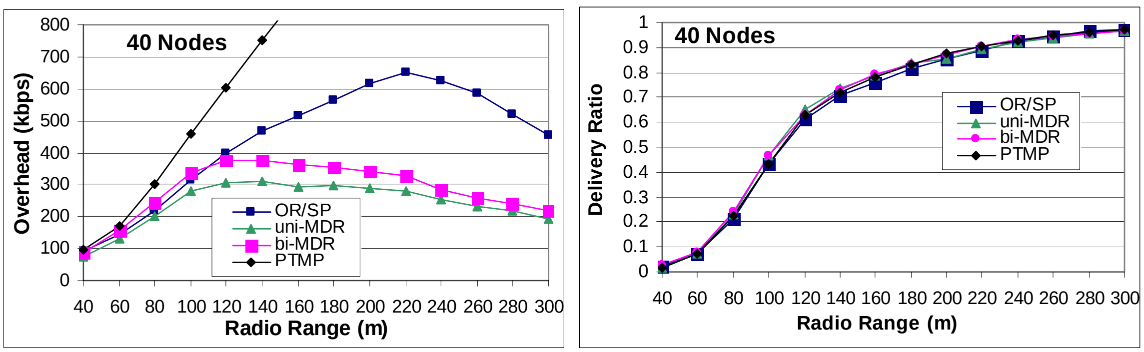

Figure 3: OSPF-MANET comparison for 40 nodes

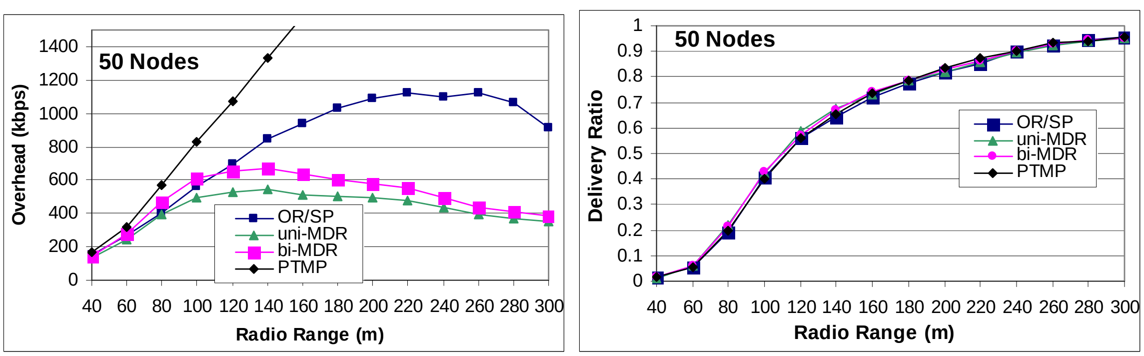

Figure 4: OSPF-MANET comparison for 50 nodes

The results in Figures 1-4 suggest a nearly opposite conclusion from the results presented at IETF-64 [1]. MDRs with both uni- and bi-connected adjacencies and full LSAs have significantly lower overhead than OR/SP with full LSAs as the network scales in density and number of nodes. In some cases OR/SP generates twice as much overhead. However, the IETF-64 results showed OR/SP networks with significantly less overhead. In the cases where, OR/SP and MDR perform about the same, the performance advantage of OSPF-MANET over PTMP is not very significant. In most cases, an OSPF-MANET interface type would probably not be necessary; moreover, the data plane delivery ratio is so low as to question whether the network is even usable in any configuration.

The delivery ratios are nearly the same in all cases, while in Cisco's IETF results [1], MDRs had slightly higher delivery ratio. The change in results can be mostly attributed to the various bugs fixed in the code and the differences in timer values.

4. Further Analysis

In this section, we further analyze the data summarized in Figures 1-4 to better understand what underlying factors contribute to the performance differences.

4.1 Additional Informative Data

We now present further data (Figures 5-8) to highlight the differences between the operation of OR/SP and MDR. In each case, we plot the average number of 2-Way neighbors per node, adjacencies per node, seconds per router-LSA, and router-LSA hopcount/install.

The average number of 2-Way neighbors per node is a measure of the network density as observed at the OSPF protocol level. Holding other parameters constant, this network density can be varied by varying the radio range in our simulations. As shown in Figure 5, the average number of 2-Way neighbors per node is independent of whether OR/SP or MDR is used. This result confirms that changing the topology formed by adjacencies does not impact the 2-way neighbors that are formed.

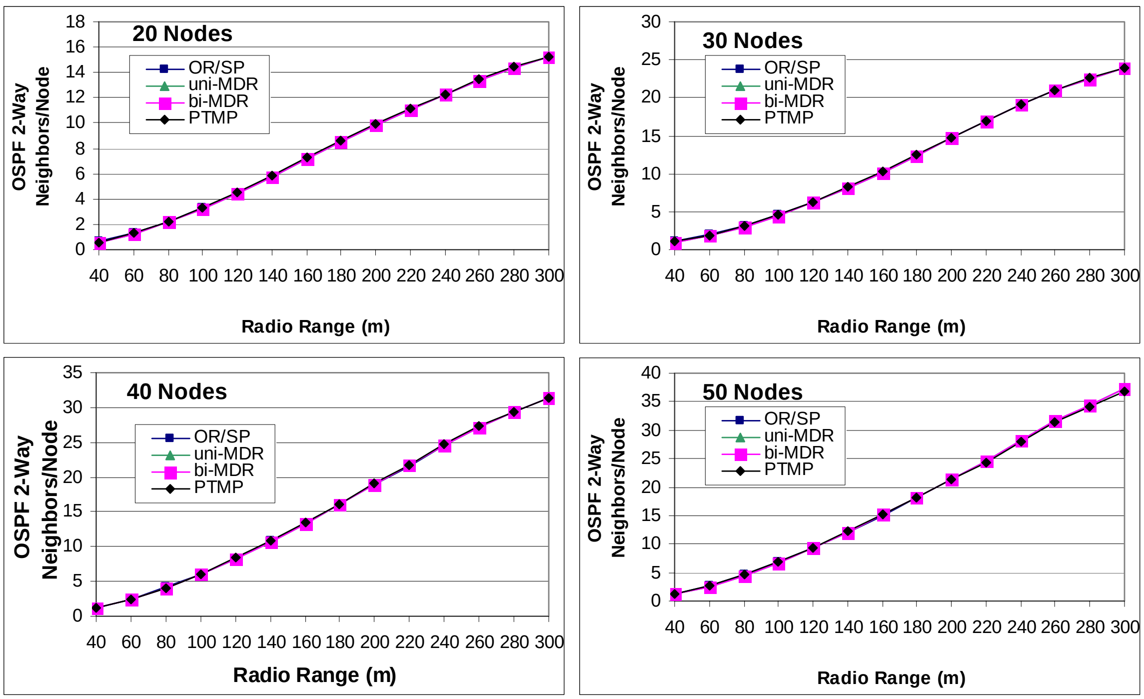

Figure 5: OSPF 2-Way neighbors per node as a function of network density

BOEING®

The plots of the average number of adjacencies per node are much more instructive (Figure 6). We see that MDRs (both uni- and bi-connected) maintain about the same number of adjacencies independent of the number of nodes and radio range. This is not true if the number of neighbors is low enough to limit the number of adjacencies that can be formed. In contrast, the number of adjacencies in OR/SP grows as the number of nodes or the radio range increases. These results are in line with the design of MDRs and OR/SPs. MDRs are designed to become adjacent with a fixed number of 2-Way neighbors based on being uni- or bi-connected. OR/SPs form adjacencies based on routability of its neighbors. This latter method is therefore very dependent on the order in which neighbors are found or lost, and is subject to instances in which routers independently try to correct the loss of routability by bringing up different adjacencies at nearly the same time.

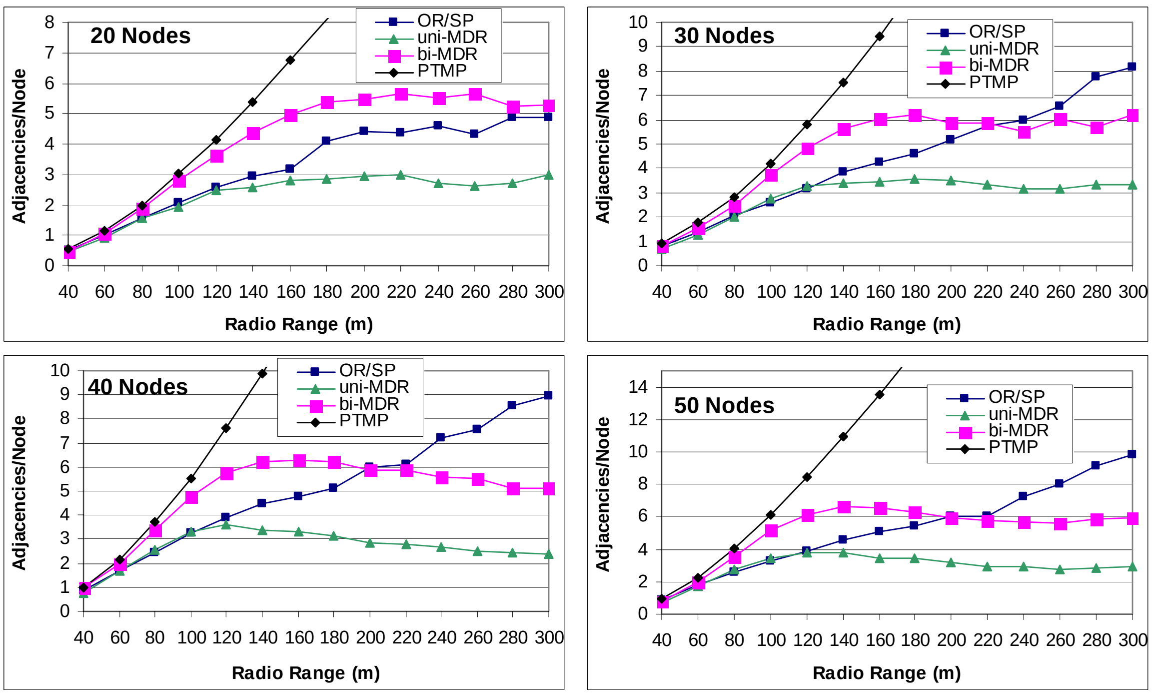

Figure 6: OSPF adjacencies per node as a function of network density

Plots of the number of seconds per router-LSA origination (Figure 7) show the same trend for both approaches. A router-LSA is originated whenever a link is added or subtracted to the LSA. This overall process is further regulated by the architectural constant MinLSInterval, which prevents LSAs from being originated more frequently than every 5 seconds, and has the effect of aggregating changes when changes are more frequent than every 5 seconds. In a dense network (in which the radio range covers most neighbors), the links do not change very often with mobility, and in a sparse network, the number of links to change is fewer. In between, there are many links with frequent changes, resulting in the inverse bell-shaped curves. The plots indicate that OR/SPs originate router-LSAs at a slightly higher rate than MDRs, but that there is not a huge difference between the approaches. The increased frequency may be due to choosing which links to become adjacent with based on routability. Adjacency changes at an OR/SP router can be triggered by a topology change many hops away. We also note that these simulation scenarios use mobility and radio models that cause topology change at a rate that approaches the limit (based on MinLSInterval) allowed by OSPF.

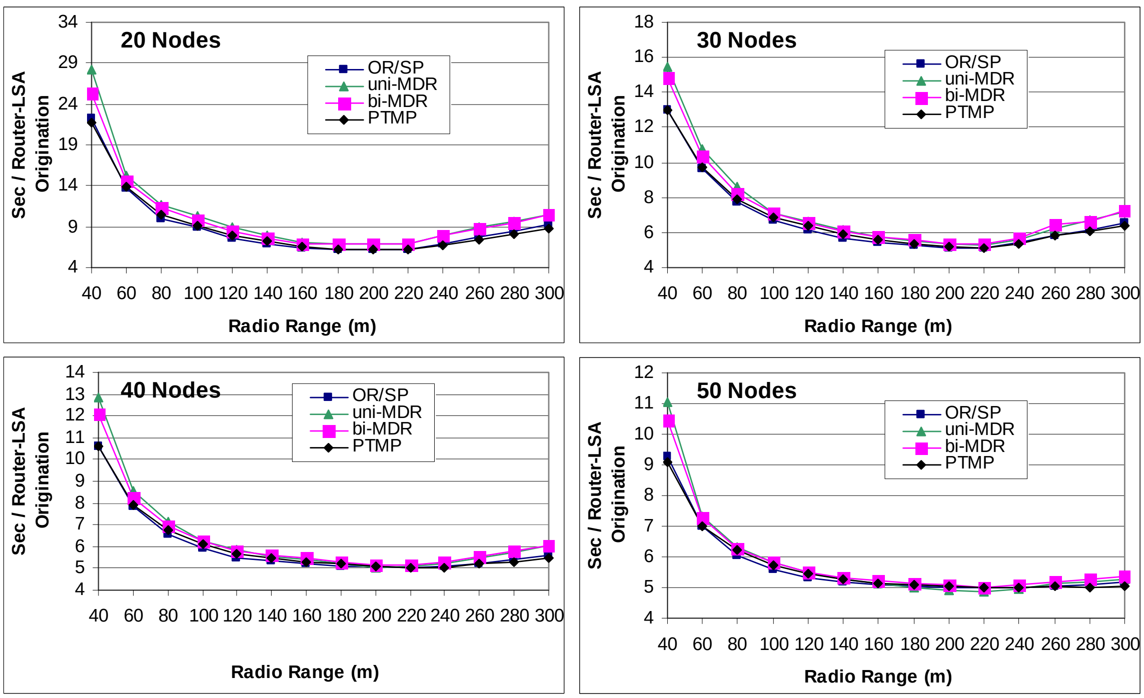

Figure 7: Frequency of router-LSA updates, as function of network density

The data in Figure 7 suggests that LSAs are being originated at roughly the same frequency regardless of the different flooding approaches. Therefore, the difference in overhead between the approaches must be due to the flooding efficiency and/or flooding path lengths, since the size of the LSAs (a function of the 2-Way neighbors in Figure 5) is roughly the same. We instrumented the simulations further to try to measure the difference between the approaches due to flooding path lengths. We created a new statistic, labeled "Router-LSA hopcount/install," to roughly measure the average number of flooding hops required before an LSA is installed in a router's database. We calculate this statistic by keeping two running counts: the number of installed LSAs and the total number of hops taken before installation in a database. Every time an LSA is installed we increment the counter for the number of installs by one, and increment the counter for the number of hops taken to be installed by the LSA hopcount. The LSA hopcount is maintained in a new field in the LSA itself (maintained for simulations only), and is incremented by one every time an LSA is reflooded. At the end of the simulation, we divided the total hopcount by the total number of installs to yield the statistics shown in Figure 8. This metric is imperfect in that it does not account for LSAs that are never installed successfully, but we believe that it is accurate enough to provide insight into the flooding stretch introduced by these proposals.

The plots in Figure 8 clearly show that OR/SP requires many more hops for an LSA to reach its final destination. In some cases, it took twice as many hops for an LSA to be installed when using OR/SP flooding. This is counter-intuitive, since ORs are based on MPRs and hence should be able to flood along min-hop paths. The reason for the flooding stretch in OR/SP is discussed in the next section.

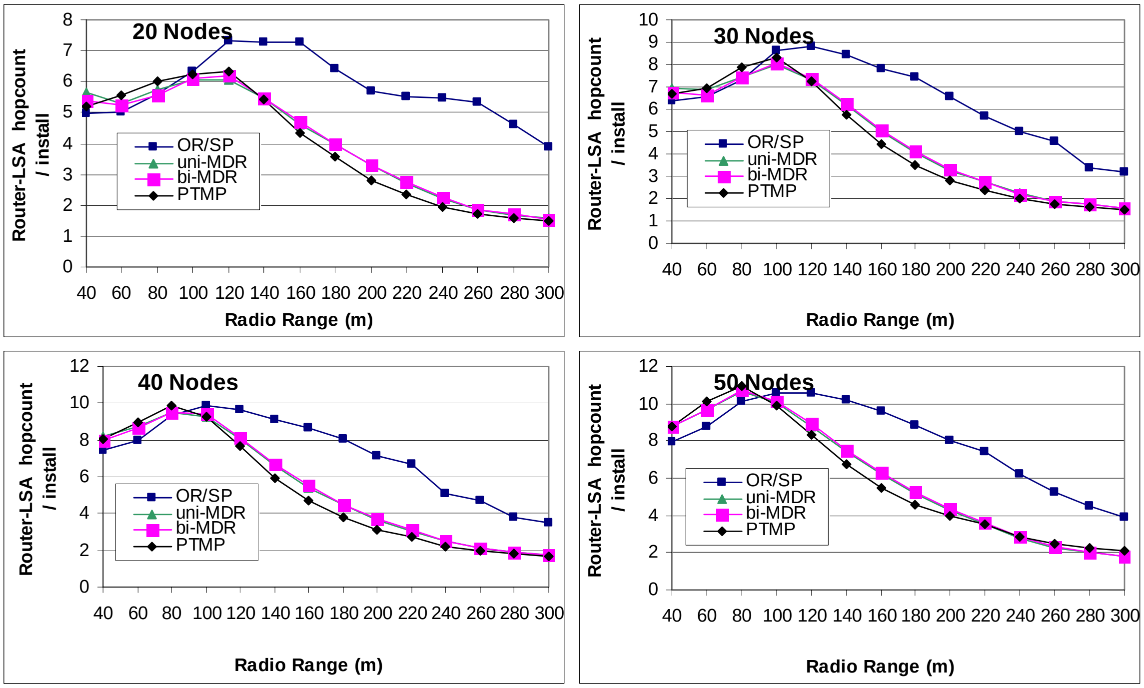

Figure 8: One measure of the flooding stretch incurred by different approaches.

4.2 Database Synchronization

One nagging question that has been unanswered since the first release of our initial OR/SP extensions was why OR/SPs had so many more LSAs out of synchronization, compared with similar PTMP and MDR scenarios, when database exchange was performed. To clarify, upon forming new adjacencies, OR/SP router databases are typically further out-of-sync than MDR routers in the same scenario. During this analysis of MDRs and OR/SPs, we found that the number LSAs out of synchronization was lower in MDRs because MDRs allow the processing of LSUs from routers in state 2-Way or higher. This is allowed in [4], Section 8, bullet 1. In practical terms, the MDR design allows for LSAs to be flooded over unsynchronized links; a commonly used term in wireless literature for this is "overhearing." MDR routers are able to overhear the transmission of neighboring non-adjacent routers and, if new LSAs are discovered, the MDR routers can then reflood them. This design choice in MDRs also influences the Router-LSA hopcount/install statistic, and can explain the results in Figure 8.

We next explored whether the same technique could be used in OR/SPs. The answer seems to be that overhearing can be used in OR/SPs, but it cannot be exploited to the same extent. The problem with accepting LSAs over non-adjacent links in OR/SPs is that these neighbors cannot reflood the LSA, since they are not MPRs for the transmitting source. A neighboring router that is below state Full will never be chosen as an MPR, so it can accept the LSA but it cannot reflood it. In contrast, MDRs are not chosen by a neighbor, so an MDR can accept and flood the LSA regardless of the neighbor state of the neighbor from which it was received. We modified the OR/SP code to allow this overhearing (defined in [4] Section 8); Figure 9 shows a 50 node case of OR/SPs with the ability to receive and process, but not necessarily reflood, LSAs from neighbors above state 2-Way.

With the new OR/SP modification, we found that the overhead improvement due to MDRs is no longer as significant as shown above in Figures 1-4. This new data is plotted in Figure 9, and one can see that the performance of the modified OR/SP now resembles that of uni-connected MDR for most statistics. However, the most notable differences remaining between MDRs and OR/SPs in these simulation scenarios (use of full advertised topology) are the number of adjacencies/node and the overhead at higher densities.

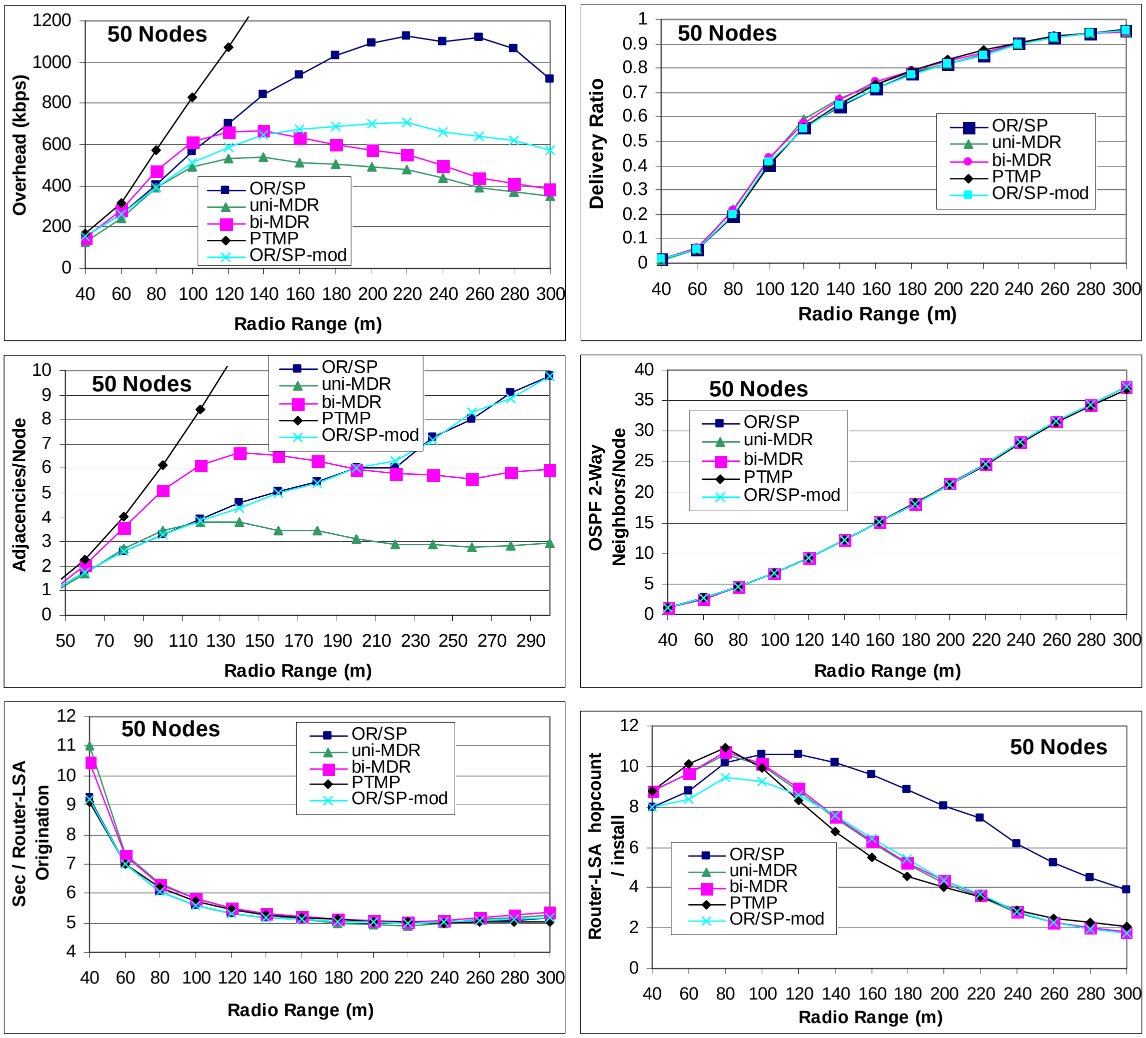

Figure 9: Results for modified OR/SP approach.

The data in Figure 9 raises two questions: Why does the higher number of adjacencies for OR/SP not result in a corresponding increase in overhead (which flattens out as radio range, and hence density, increases), and why is the overhead still generally higher in OR/SP as the density increases.

1. **Why does the overhead not scale with adjacencies in the modified OR/SP?** A larger number of adjacencies can lead to increased overhead due to increased frequency of database synchronization. However, the overall rate of LSA generation, or the size of LSAs themselves, should not scale with the total number of adjacencies since we are considering full topology LSAs (from an LSA topology perspective, both proposals behave as if there were full OSPF adjacencies). We examined the raw statistics concerning database exchange, and found that the database exchange-related traffic is not significant in these scenarios when compared to the flooding overhead. Therefore, the fact that modified OR/SP has more adjacencies than the corresponding MDR case does not contribute significantly to the overhead performance gap between them.
2. **Why more overhead with modified OR/SP?** We found that the additional overhead is mainly due to the number of MPRs that decide to flood an LSA versus the number of MDRs. We do not know whether this performance gap could be further closed by further OR/SP optimizations or timer value adjustments.

5. Summary

This report has documented our attempt to validate and extend the simulation performance results presented by Cisco Systems at IETF-64. We were not able to validate the results presented at IETF-64, so we worked with the authors of the Smart Peering simulation software extension to debug it and adjust timer values appropriately. In section 3, we presented new results for similar scenarios that yielded the opposite conclusion as presented by Cisco; the overhead performance of Overlapping Relays with Smart Peering (OR/SP) was worse than that of a similar MANET Designated Routers (MDR) configuration. In section 4, we have gone further to try to explain the difference between the two approaches, by comparing plots of physical neighbor density, OSPF adjacency density, the frequency of router-LSA origination, and measures of flooding stretch. We found in these scenarios that the flooding overhead in many cases was capped by the OSPF MinLSInterval parameter (5 seconds), that the OR/SP solution did not control the density of OSPF adjacencies as well as did MDRs, and that the flooding stretch by OR/SP was worse than that of MDRs. We explored a simple modification to OR/SP to allow it to overhear neighboring LSA transmissions rather than waiting to receive them on the regular MPR-based flooding path; this modification narrowed the performance gap between OR/SP and MDRs, but did not eliminate it. We do not know whether further tuning of OR/SP parameters and timers might bring it closer to the performance of similarly configured MDRs. We also note that this study has been limited to the case of using full-topology LSAs (including neighbors for which adjacencies have not formed); we did not examine the possible additional scalability that may be realized by using partial topology LSAs [4].

6. Acknowledgments

The authors thank Richard Ogier and Stan Ratliff for working on improvements to the OSPF MANET code base, and also to Richard Ogier, Acee Lindem, Stan Ratliff, and Abhay Roy for clarifying aspects of the respective protocol designs.

7. References

[1] S. Ratliff. “Comparison of OR and MDR Overhead and Delivery Ratio”, Cisco Systems. 8 Nov. 2005 <http://www3.ietf.org/proceedings/05nov/slides/ospf-4/sld1.htm>.

[2] M. Chandra. "Extensions to OSPF to Support Mobile Ad Hoc Networking", draft-chandra-ospf-manet-ext-03.txt (expired), April 2005.

[3] A. Roy. “Adjacency Reduction in OSPF using SPT Reachability”, draft-roy-ospf-smart-peering-01.txt (work in progress), November 2005.

[4] R. Ogier and P. Spagnolo. "MANET Extension of OSPF using CDS Flooding", draft-ogier-manet-ospf-extension-06.txt (work in progress), December 2005.

[5] S. Ratliff. “FW: [Ospf-manet] new software snapshot”, Cisco Systems. 8 Feb. 2006 <http://www1.ietf.org/mail-archive/web/ospf-manet/current/msg00114.html>.

[6] T. Henderson, P. Spagnolo, and G. Pei, “Evaluation of OSPF MANET Extensions,” Boeing Technical Report D950-10897-1, July 21, 2005. Available at http://hipserver.mct.phantomworks.org/ietf/ospf.