

A STUDY OF LINK STATE FLOODING OPTIMIZATIONS FOR SCALABLE WIRELESS NETWORKS

Joseph P. Macker
Justin W. Dean
Information Technology Division
Naval Research Laboratory

ABSTRACT

Several methods have recently been proposed to improve link state protocol efficiency within wireless ad hoc networks. One class of approaches targets the reduction of link state control overhead that contributes to network-wide broadcast storms, therefore enhancing protocol scalability. We have implemented and studied two such methods and demonstrate and discuss their relative performance characteristics. We present early simulation analysis over a number of scalability factors, including average nodal density and network diameter (maximum path length). Based upon initial results, we demonstrate that the two approaches to link state overhead reduction are somewhat complementary and that they provide additional benefit when applied together in many topologies studied. While both approaches are valid scalable flooding techniques, we discuss further the relative merits and potential disadvantages of each technique.

BACKGROUND

There has been recent interest in improving the scalability and operation of link state routing mechanisms in wireless ad hoc networks [CM99]. While applying a hierarchy can reduce protocol overhead through state aggregation, several methods recently proposed reduce classical link state overhead requirements without hierarchies. Classical flooding techniques often used for distributing link state amongst the network nodes contribute significant protocol overhead and traffic congestion as the network scales, especially when applied in wireless networks [SRS02]. We define classical flooding to be the case in which each node throughout the network rebroadcasts the data packet once and only once.

In recent years, more routing control flooding improvements for mobile ad hoc networks (manet) have been proposed and there has been some evaluation of their relative merits [HXG02]. We have chosen to model and analyze some detailed scalability aspects of two recent innovative approaches that address problem of link state flooding in different ways.

Figures 1 and 2 illustrate two such methods used to reduce control flooding overhead for link state algorithms. The




Figure 1: Flooding Reduction via Optimized Relaying

first design method (illustrated in Figure 1) provides a more efficient approach to classical flooding by reducing retransmissions required to reach all nodes. A second design approach (illustrated in Figure 2) applies the assumption that by updating localized regions more frequently than distant regions overhead can be reduced while maintaining effective routing. The applicability assumption here, which may be true for many practical network scenarios, is that differential topology changes between short time epochs are mainly localized.

As shown in Figure 1, there has been several design methods proposed for wireless link state routing that reduce required network retransmissions by dynamically




Figure 2: Flooding Reduction via Time-Space Algorithm

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE A Study of Link State Flooding Optimizations for Scalable Wireless Networks				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Information Technology Division, 4555 Overlook Avenue, SW, Washington, DC, 20375				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

determining a subset of forwarding nodes. Variants of such forms include cluster-based forwarding and source-specific selective forwarding. While there are many variants of this forwarding subset concept and we based our initial modeling and analysis upon the multi-point relay (MPR) design concept from Optimized Link State Routing (OLSR) work [QVL00]. One difference in our design from previous OLSR based work is that we implemented a full link state algorithm using MPRs in addition to the specified partial link state approach. The MPR forwarding method promises to provide flooding overhead reductions for non-sparse network neighborhoods and thus improves overall network scalability improvements for protocols or applications requiring a flooding service.

Figure 3 illustrates a set of example LS time-distance functions that can be applied to the link state flooding process in both time and space. As partially illustrated in Figure 3, several methods have been proposed for more efficient time-space link state dissemination. Two of the more well-known approaches include fisheye state routing (FSR) [PGC00], similar to near sighted LS of Figure 3, and more recently the concept of fuzzy-sighted link state (FSLs) [SRS02]. In both of these cases, link state information is updated less frequently as network distance increases between participating nodes, but the methods by which this is accomplished are different.




Figure 3: Classes of Time-Distance LS Functions

We feel the FSLs is a reasonably general approach in this design space and therefore we have chosen to model and analyze it here as our second method of link state flooding reduction. Both the MPR and FSLs methods promise to provide protocol overhead reduction as a network scales, but relative improvements are anticipated to rely on different network scenario factors and scaling metrics. We expect the MPR approach to provide improvements relative to the density of a network or regions within the network and the FSLs type of approach to provide more overhead reduction in networks of increasing diameter. If

the network is both dense and increasing in diameter we would expect both approaches to provide additional benefit in combination. The goal of this work was to initially examine these hypotheses in terms of several structured simulation studies.

SIMULATION MODELS AND APPROACH

In order to analyze the relative performance and scalability aspects of the various LS flooding optimizations discussed, we created a set of working simulation models. We began by implementing a basic wireless link state routing protocol within the ns2 network simulator [NS2]. For comparison purposes, we also implemented a more classical flooding algorithm to provide a performance baseline of LS overhead reduction. This classical flooding model was excited by a version of discretized LS as shown in Figure 3, so our baseline model should produce less overhead than standard LS. Each node produced link state updates (LSUs) at defined intervals and each node rebroadcasted received LSUs once (duplicate LSUs were detected and dropped).

Study Conditions and Assumptions

All LS models and extensions that we developed and analyzed are based upon full link state routing protocol versions. For example within the OLSR specification, the concept of MPR flooding is not only used to reduce redundant retransmissions but also reduces the amount of link state information required within an update (i.e., only MPRs produce LSU messages and include only the list of MPR selector sets not the full neighbor set)—the result is a partial link state routing algorithm. To extend beyond this, we have developed a MPR flooding variant that optionally provides full link state information. This provides results independent of any additional message reductions that result from other protocol enhancements, such as partial link state algorithms. We, therefore, feel the overhead results we are collecting are conservative.

All LS models studied also use hellos for discovering two hop neighborhoods. Multi-point relaying needs two hop information for its flooding algorithm. Though fuzzy sighted LS (by itself) does not need two-hop neighbor information, a mechanism for discovering local changes is needed to assure correct information for flooding, and hellos are a simple way of achieving this. There are methods for reducing the size of hellos, but for the purposes of this study we use a simple hello process which allows two-hop neighbor information collection and maintenance for consistency.

We chose a number of scalability factors for use in our study, including average nodal density and network diameter (maximum path length). We define nodal density to be the expected number of neighbors any given node

has. N total nodes are placed in a random uniform fashion within a W by W unit square. Each node is defined to have a radius of awareness, defined as R units. Our definition of density is not directly correlated to the ratio of number of nodes per unit area when R is not small compared to W . This is important to note because R is not small compared to W for some of our smaller diameter tests.

To obtain the density for a given simulation we calculated an area of awareness for a randomly placed node in the W by W grid. To do this we used a weighted average for three different regions.

Region 1: node is placed towards the center so that no sides of the simulation square intersect the awareness circle

$$\begin{aligned} area_1 &= \pi r^2 \\ weight_1 &= (w - 2r)^2 / w^2 \end{aligned}$$

Region 2: node is placed such that one side of the simulation area intersects the awareness circle

$$\begin{aligned} area_2 &= (-2/3 + \pi)r^2 \\ weight_2 &= 4(w - 2r)r / w^2 \end{aligned}$$

Region 3: node is placed in the corner so that two sides of the simulation area intersect the awareness circle

$$\begin{aligned} area_3 &= (-5/12 + 3/4\pi)r^2 \\ weight_3 &= 4r^2 / w^2 \end{aligned}$$

Expected maximum network diameter was simply defined by the “maximum” minimum path length in a given scenario given by the diagonal path achievable from opposite corners of the square coverage area. This estimation assumes a well connected network. The maximum expected network diameter formula is simply:

$$Diam_{max} = w\sqrt{2}/r$$

While the simulation models we used are fully capable of supporting motion, our initial analysis presented here is based upon randomly generated static topology scenarios as the network scales in various dimensions.

Classical Flooding (Discretized Link State)

As our baseline flooding algorithm we used discretized LS flooding. At every flooding interval every node broadcasts link state message. When a node receives a new link state message from a neighbor it updates its link state information and rebroadcasts the link state message once and only once. In this way messages are disseminated throughout the network and each node has a picture of the entire network LS topology.

Multi-Point Relaying

Multi-point relays use a local subset of forwarding nodes to reduce the number of forwarded packets. We used the heuristic election algorithm defined for the OLSR specification to select forwarding nodes or MPRs. A two-hop neighborhood is discovered by the hello mechanism. Neighbors that have paths to unique two-hop neighbors are selected as forwarders first. Further election is covered by the heuristic algorithm defined in [QVL00] until all two-hop neighbors are covered by at least one designated forwarder.

Nodes broadcast link state information at every flooding interval same as our classical flooding model. The difference is that link state messages are only forwarded by neighbors which were selected as MPRs. The reduction in flooding overhead comes from the smaller number of forwarding nodes for the same amount of coverage.

Fuzzy-sighted Link State

Once again, the premise behind FLSLS routing is that local topology changes affect routing decisions more frequently than far away changes, and because of this, far away neighbors do not need to be updated as often about local topology changes. Nodes wanting to forward packets to a distant node do not necessarily need to know the entire current path just the general direction in which to send the packet. As the packet moves along, closer to its destination, forwarding nodes have newer information about where the destination node is located. In this manner packets are delivered without having the newest information as long as the relevant topology information has been refreshed effectively. There are a large set of possible refreshing equations in time-space that one can apply to FLSLS. These equations may even be dynamic dependent upon network topology conditions perceived.

¹ Integrated estimation of actual area, we empirically verified estimated results of nodal density and diameter in random network trials.





Figure 4: Overhead for 3-hop and 8-hop network diameter tests

We applied a time-distance LS equation defined below to set the time to live (TTL) field of link state packets.

Fuzzy-sighted link state equation:

$$i \in N, j \in N$$

$$ttl = \begin{cases} i \in 2^N \Rightarrow h(i) = i \\ i \notin 2^N \Rightarrow h(i) = h(i - 2^j) | 2^j < i < 2^{j+1} \end{cases}$$

The value of the TTL field limits how far flooding will travel in terms of hop distance, however flooding cannot go further than the maximum physical hop count of the network no matter the value of the TTL field. TTL values that are greater than or equal to the maximum hop count of the network assure full flooding given a connected network and no dropped packets.

This formula should be well suited for evenly dense networks with movement slow enough for LS intervals to capture local movement effectively. LS update distances will be able to keep up with motion assuming the interval is small enough. We chose this formula for what we believe to be its robustness in multiple different scenarios. The scope of this paper does not include analysis of robustness of the FLS protocol approach in general or that of other time-space equations that could be used.

Fuzzy-sighted Multi-Point Relaying

Just as the name implies this method uses both multi-point relaying and fuzzy-sighted link state methods. Multi-point relays are selected and used in the same manor as regular MPR flooding except flooded messages are only sent as far as the fuzzy-sighted formula dictates. Since both methods address different problem domains they should work together and create a flooding mechanism that produces less overhead then either one of them separately.

SIMULATION RESULTS

Scenario Description

Simulation scenarios in ns2 were based on 60 second trials with no movement and no data traffic. Nodes were placed randomly in a square area for each successive trial. Both random node placement and random jitter added to the flooding of control messages affected the overall number of flooding packets seen. Therefore, multiple runs were executed (up to 10) for each scenario to minimize the effect of random node placement and to bring average results closer to the average density and maximum hop diameters generated analytically. Link state flooding intervals were set to 5 seconds and a random transmission jitter of 2.5 seconds was centered on the 5 second intervals.

We wanted to see the overhead growth of the different flooding algorithms as the network grew in both diameter and density. To see how density alone effects growth we varied the density from an average of 4 to 20 1-hop neighbors by varying the number of nodes in a given area with a fixed maximum hop diameter. To see the affect the network diameter had on growth of overhead we then varied the diameter from 3 to 12 keeping the density fixed. Overall this resulted in 54 different scenarios. For each scenario we generated 5 random placement files which contained node position information. We then simulated the different flooding algorithms on each of the random position files.

From the simulation trace files we were able to count the number of flooded control packets which were sent throughout the network. We counted both originating and forwarded messages. The results are presented in the next section.

Simulation Results

We found that both FSLs and MPR flooding performed better than simple discretized flooding in all scenarios in terms of amount of overhead traffic generated. We also found that the combined FSLs/MPR method produced less flooding overhead than either method alone in all scenarios. Figure 4 shows the results, on a log scale, for a set of small diameter test runs and a set of larger diameter scenario runs. Reductions in overhead caused by MPR flooding methods were more significant than FSLs methods in smaller diameter test runs, less than a maximum of 5 hops. As demonstrated in the small diameter scenario the number of messages sent by the MPR method is always less than the FSLs method. FSLs starts producing less overhead than the MPR style flooding at all densities in scenarios with 8 or larger maximum hop count. Again in figure 3, you can see that FSLs produces less overhead than MPR flooding for every density when the network diameter was set to eight hops.




Figure 5: Overhead for 8 node density simulations

In the 5 to 8 hop diameter test range, reductions in overhead were observed to be significant and of similar magnitude for both methods. Which method produced less overhead was related not only to the diameter parameter but to density as well. Figure 5 shows the growth of overhead as the network diameter increases for tests with a fixed average neighborhood density of eight. At this density FSLs starts producing less overhead than MPR flooding at the crossover point of about six hops. We produced a set of results for each simulated density and noticed that all crossovers occurred at about the same six hop diameter, though there was a slightly increasing trend. The crossover point occurred at greater diameters as density increased. We expected this as MPR flooding should become more efficient as the localized density of the network increases.




Figure 6: Overhead efficiency for 6-hop network diameter

Figure 6 shows the reduction in overhead, when compared to discretized flooding, caused by the different flooding algorithms with increasing density over a fixed network diameter of 6 hops. The downward slope of the MPR curve demonstrates the increased efficiency of MPR flooding as density increases. FSLs on the other hand exhibits roughly the same reduction ratio across all densities given that the diameter remains fixed. The reduction caused by fuzzy-sighted multi-point relaying is about equal to the product of the FSLs and MPR percentage reductions. This relation supports the theory that the overhead reduction methods are orthogonal.

Discussion of Scalability Trends

From these results one can see which flooding algorithm should produce less overhead for a given uniform random scenario. For most small to medium sized networks the reductions afforded by MPR flooding are quite substantial and should outperform reductions caused by FSLs methods alone. MPR flooding overhead grows faster than FSLs dissemination as the network grows in diameter. For large diameter networks MPRs alone may not be sufficient for reducing overhead.

FUTURE CONSIDERATIONS AND WORK

Mobility Issues and Robustness Concerns

It is important to reiterate that the simulation results presented in this paper only take into account packet overhead of different flooding mechanisms. A connection between overall link-state routing performance and the overhead reduction caused by the different flooding algorithms is not explored in this paper. Pro-active algorithms based upon dominating neighbor subsets for flooding, MPR flooding, have been implemented in real protocols and the reduction in overhead of these protocols, caused by more efficient flooding, seems to outweigh the loss in robustness that comes from more redundant

Dynamics and non-homogenous topologies

Power control and Topology Adaptation

CONCLUSION

While we judged both approaches to be valid scalable flooding techniques, we also briefly discussed the relative design assumptions and characteristics of each method. Further work is planned in the studying overall effectiveness and application for such techniques in more mobile and dynamic scenarios.