Simplified Multicast Forwarding in Mobile Ad hoc Networks

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

ABSTRACT

We implemented a working IP multicast forwarding prototype for use in mobile ad hoc networks (MANETs) based upon flooding mechanisms. We present the design of a working experimental prototype and some initial performance results using the NRL mobile network emulation system and various optional flooding approaches within the design framework. In addition, we present supplemental analytical examination of several implemented flooding algorithms for MANET environments and discuss related performance tradeoffs. We conclude by presenting further technical considerations and future work issues.

BACKGROUND AND MOTIVATION

There is growing interest in designing generic user data flooding or simplified multicasting forwarding services for use within mobile ad hoc networks (MANETs)[QVL00, CLOV02, WC02, NTCS99.] Also, previous work developing MANET unicast routing protocols often involves the application of a packet flooding service within the protocol control plane to support basic functionality (e.g., router discovery, proactive link state dissemination). Some novel technical work has been done to optimize the flooding process for specific routing protocol designs and requirements [SRS02, HXG02, QVL00, PGC00, OTL04].

A basic packet forwarding service that reaches all destinations participating within a MANET environment can be a useful generic routing mechanism for an application layer. While the design requirements for such a forwarding mechanism are similar to those often needed in the control plane by many MANET unicast routing protocol layers, it is desirable to provide a more generic forwarding function for use by other applications. There are a number of application areas that could take advantage of a simple, broadcast-type delivery service within a MANET routing region (e.g., multimedia streaming, peer-to-peer middleware multicasting, MANET auto-configuration, and discovery services). In this paper, we discuss the implementation of an initial working prototype and analyze some aspects of such a packet forwarding capability and we model it as a simplified multicast routing service for scoped MANET application.

The simplest design often conceived and adapted for MANET packet flooding is a classical flooding (CF) algorithm. In CF, each participating forwarder node is required to rebroadcast a packet when heard for the first time. This approach is extremely simple and generally only requires duplicate packet detection and a basic forwarding mechanism. However, it is well known that using CF results in a significant number of redundant transmissions often referred to as the *broadcast storm problem* [NTSC99]. In wireless MANET environments, reducing unnecessary channel contention significantly improves network performance. Therefore, reducing the number of required relay nodes is a heightened design goal for this environment. Unfortunately, reducing the number of relay nodes in a MANET environment may also decrease the robustness of overall packet delivery. There exists an interdependent design tradeoff between relay efficiency and delivery robustness that is scenario and system dependent and should be examined carefully.

At a theoretic level, work in the area of minimizing packet forwarders, or relay node sets, is often related to basic graph theory problems. In graph theory, a *dominating set* (DS) for a graph is a set of vertices whose neighbors, along with themselves, constitute all the vertices in the graph. A connected DS (CDS) is a DS forming a connected graph. A minimum CDS (MCDS) is a set such that the number of vertices is the minimum required to form a CDS. Finding a small dominating set is one of the most fundamental problems of traditional graph theory and is in theory often related to the problem of optimizing flooding algorithms in MANETs. Finding an MCDS in a given graph is known to be NP-hard [GJ79.] Beyond these basic static graph theoretic issues, MANET protocol designs require more distributed and dynamic operation. To better explain the design motivations, we formulated three basic competing design characteristics of an effective MANET flooding algorithm solution:

- A resultant cover set that is small compared to the total number of nodes as the network scales in size and density.
- A robust approach somewhat resilient to network mobility and link dynamics.
- A cover set election/maintenance mechanism that is lightweight, distributed, and adaptive in nature.

RELATED WORK AND OVERVIEW

Previous novel work on MANET flooding has been by done by others. In [WC02], taxonomy of flooding algorithms for use in MANET environments was presented and the work examined performance issues related to various approaches. Other important previous work has developed distributed mechanisms that select and maintain reduced relay node sets. As we have

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **2004**	2. REPORT TYPE	3. DATES COVERED **00-00-2004 to 00-00-2004**

4. TITLE AND SUBTITLE **Simplified Multicast Forwarding in Mobile Ad Hoc 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
The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**		**7**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

already mentioned the design tradeoffs are further complicated by wireless contention, topological classes, and the robustness of packet delivery with mobility. In addition, the development of an actual protocol implementation for IP multicast forwarding based upon these flooding algorithms raises additional design tradeoffs and issues. This includes maintenance of protocol state, duplicate packet detection mechanisms, and any associated protocol signaling required by a mechanism.

Our work here concentrates on the design and implementation of an initial working prototype for IP MANET multicast forwarding based upon improved flooding mechanisms. The initial design adapts a variety of flooding algorithms, but additional algorithms can be easily added to the framework for prototype experimentation. In the next section of this paper, we describe the working implementation and the related design issues. Following the implementation section, we analyze the initial working prototype within the NRL mobile network environment [CMW03] and discuss several example performance results. In a later section, we supplement implementation and emulation observations with more scaled examination from several basic analytical models. We conclude by discussing future work and summarizing our work to date.

IP DESIGN APPROACH

In addition to researching various algorithms and approaches for MANET flooding, our goal was to provide a working prototype compatible with existing and emerging IP network protocol frameworks. The implementation approach taken enables the protocol to work as a simplified MANET multicast routing mechanism within a MANET routing area. At present the IP packets are not encapsulated in any additional or specialized IP header, so multicast routing is performed on native IP multicast application packets. One important routing design difference between MANET interfaces and many wired network interfaces is that forwarding out the same interface a packet arrived on is a normal allowed operation. It is important to note that this operation is often disallowed in wired multicast routing designs. Because of this feature, a fairly common requirement in MANET packet flooding is some form of duplicate packet detection. This generally requires some form of packet sequence identification.

Our simplified MANET multicast routing implementation is composed of three parts: A sequence id generator and marker to be used when and if necessary, a duplicate detection module, and a basic multicast packet forwarding module. The sequence generator is responsible for marking each packet with a monotonically increasing unique identification number when existing IP kernel methods are not sufficient or are not predictable. The duplicate detection mechanism is used to remove and detect duplicate packets from both entering the interface forwarding process and from being delivered to upper layer applications. The forwarding module performs basic multicast IP packet forwarding out a particular MANET interface as appropriate.

Sequence Generation and Marking Module

As we mentioned detecting duplicate packets is important in a MANET packet flooding process. In our implementation, IP multicast packets from a particular source are assumed to be marked with a temporally unique identification number in the IPv4 header using the ID field [RFC791]. Unfortunately, in present operating system networking kernels this identification number (ID) for the IP header is not always generated or applied in a consistent manner. As an example, the present LINUX network kernel implementation does not generate a meaningful ID field entry for UDP packets. In order to build a working implementation without encapsulating packets, we built a sequence generation and marking module that can maintain and add a monotonically increasing IP ID field for source-specific multicast packets. When needed on a local system or at a specific gateway this process will also recalculate and replace a proper IP header checksum for the formulated header.

Although we have demonstrated its use in this prototype code, the adoption of the IPv4 ID field for widespread packet duplication detection has some disadvantages. The main disadvantage is that the use and interpretation of the field is known to be non-consistent across operating systems. As an alternative, the use of an encapsulated header or header extension in future implementations may provide more flexibility and consistency across implementations. We leave these design alternatives to be further defined and discussed in future work. A basic sequencing and marking design similar to the one we have formulated can be easily adapted to work with future approaches or can be bypassed when not needed.

Duplicate Detection Module

When designated multicast packets are received by the device driver, they are intercepted and examined for duplicate detection. In the present implementation, source-specific information and the sequence ID field is checked against a cached duplicate history. When a packet is not identified as received or transmitted previously, it will be sent to the forwarding module and to the upper layer. When a packet is marked as previously received, it will be silently dropped. Proper operation of the duplicate detection module relies on the fact that a working and predictable mechanism for packet identification generation and marking is in place.

Multicast Forwarding Module

The multicast forwarding module is flexible in its design and presently supports different flooding design optimizations. The current experimental mechanisms are: classic flooding (CF), source-specific multi-point relay (S-MPR) flooding, and non-source multi-point relay (NS-MPR).

The most basic mechanism implemented is the CF approach. In CF, each node transmits a locally generated or newly received packet exactly once. The duplicate detection technique mentioned in the previous section avoids any duplicate packet retransmissions.

The S-MPR flooding mechanism is based upon the well-known MPR technique and allows only *locally elected* MPRs to retransmit packets that are received from upstream selector nodes. The present algorithm leverages a local 2-hop MPR election mechanism, we build a reduced relay node set for application layer multicast data flooding. Symmetric 2-hop neighbor knowledge can be collected via single HELLO exchanges or through some lower layer mechanism if available. It is well-known that source-specific MPRs compose a connected dominating set and using S-MPR significantly reduces redundant retransmission of packets, especially in dense network neighborhoods [JLMV02]. An implementation disadvantage of S-MPR is previous hop identification is required to perform a proper forwarding match. This previous hop filtering requirement adds some additional state and complexity to the design, but it is functional in our present prototype.

We were further interested in exploring flooding techniques that do not require previous hop information during the forwarding decision process and we initially implemented the NS-MPR forwarding mode to examine one possibility. The NS-MPR mechanism combines all source-specific elected MPRs into a common relay node set. In this case, during the active forwarding process previous hop information is not required and only knowledge that a node is an MPR for at least on neighbor is used. A significant finding of our initial evaluation is that the NS-MPR technique, using the present election algorithm, does not scale well as compared with the S-MPR approach (see analytical section). In other words, a combined resultant relay set is not significantly reduced as compared to a source-specific relay set. Other optimization algorithms to form common relay sets not requiring previous hop knowledge are known and we are presently investigating alternative methods.

PROTOTYPE EMULATION RESULTS

We performed initial examination of our prototype MANET multicast implementation using the NRL mobile network emulation software [CMW03]. Our initial testbed configuration consisted of 10 mobile network nodes in an emulated mobile wireless environment running 802.11b ad hoc wireless interfaces at 2 Mbps link rates. For our initial experiments, one node is the source of the multicast and the other 9 nodes are simultaneous multicast receivers.

We expected the transmission overhead efficiency of various flooding methods to become a more significant factor, in terms of packet delivery effectiveness, as the network traffic load increases. Dependent upon the flooding approach used, the effectiveness of overhead reduction can be a function of the actual topology of the network. To first examine and illustrate some performance trends related to topology, we developed several fixed-topology tests within our emulator. The various classes of fixed network topologies we examined are shown in Figure 1, with the x labeled node representing the multicast data source. The various examples represent various local neighbor and edge node densities. While ten emulation nodes is rather limiting in the number of scenarios that could be investigated, we achieved interesting results nonetheless.

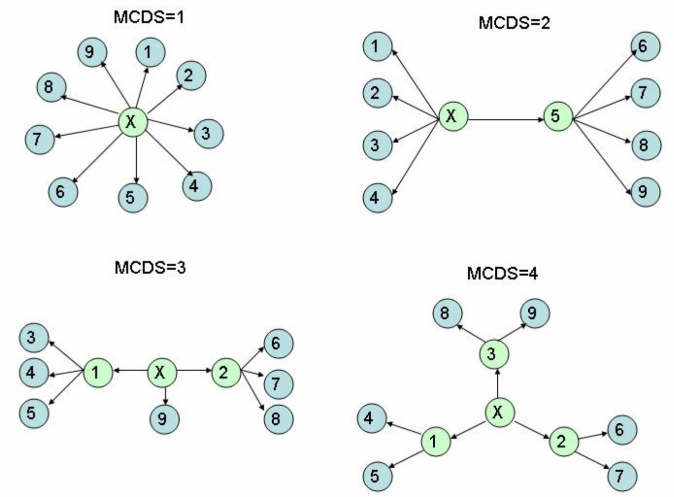

Figure 1: Fixed Topology Examples

In the fixed topology tests, we were interested in examining significant levels of congestion and network contention when using various flooding algorithms to perform MANET multicasting. This is also more typical of the expected operating conditions. In these sets of tests, we increase the offered multicast source traffic from 10-800 kbps incrementally every 10 seconds. In Figure 2, we observe that the supportable traffic transmission rate is saturating around 1300 kbps due to the physical limitations of the 802.11b operating at 2 Mbps. Notice, as expected, that the slope measuring total overhead traffic increases as the number of nodes in the relay set increases. CF flooding is represented as the steepest slope in the picture and is the same for all possible topologies (all nodes transmit a packet once). The other curves represent the effectiveness of the S-MPR mechanism operating over various topologies. The improved relay set efficiency keeps this mechanism from saturating the MANET and allows a higher achievable source rate. The decreasing slope is indicative of the network topology's increase in density allowing a smaller ratio of relay nodes to be used.

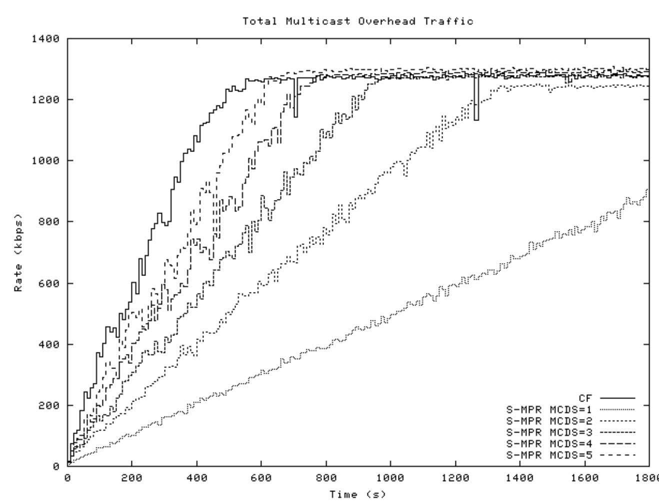

Figure 2 : Total Multicast Overhead in Various Topologies

In Figure 3, we present the data from the fixed topology tests of Figure 2 from a different operational performance perspective. Here we graph the total received multicast data traffic at all nodes as a function of time, topology, and representative flooding algorithm. CF flooding is shown as a solid line and demonstrates the worst maximum achievable performance irregardless of topology. As anticipated, S-MPR mechanism demonstrates increased efficiency as areas of network neighborhood density increase. As density increases, more efficiency in multicast forwarding is possible leading to higher sustainable maximum source rates for multicast application data (e.g., streaming video)

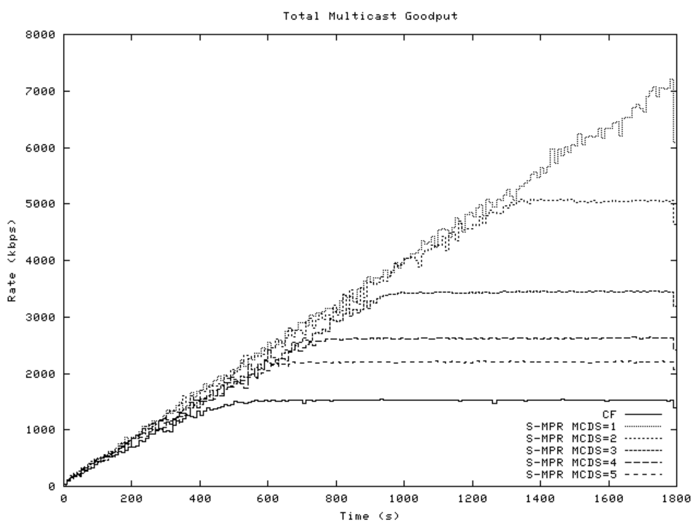

Figure 3: Total Multicast Goodput Across Receiver Set

The previous basic topological test examples demonstrate that as the required relay set increases, the maximal attainable throughput decreases. In a subsequent set of tests, we examined the multicast forwarding performance of the various approaches in a mobile network. The motion model used was a variation of the well-known random way point model. A single node was chosen to source multicast traffic and the offered traffic model was 10, 50, and 100 kbps CBR at 10 minute intervals for a complete running test time of 30 minutes. We measured the overall packets received at all multicast nodes and the total amount of flooding overhead equal to the initial transmission and all forwarded retransmissions.

In Figure 4, we observe that the total number of received multicast packets is similar with all three flooding methods used in our prototype. These results also generally demonstrate that the multicast forwarding method is working for all three methods even with node mobility occurring. CF occasionally shows a slight robustness gain at various times at the expense of redundant transmissions. By examining overhead requirements in Figure 5, we notice a significant difference in the number of overhead packets required throughout the experiment for the different methods. Even with such a small network experiment, this clearly shows the advantage of MPR-assisted flooding and performance gains would be similar for other techniques that reduce the required size of a relay node set.

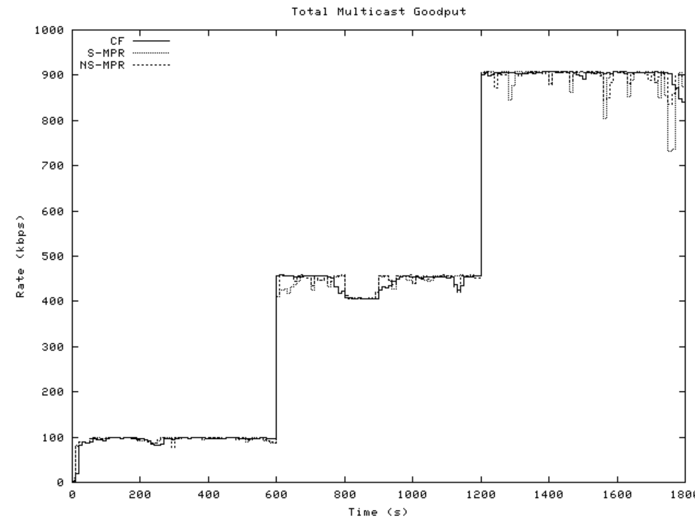

Figure 4: Total Multicast Goodput with Mobility

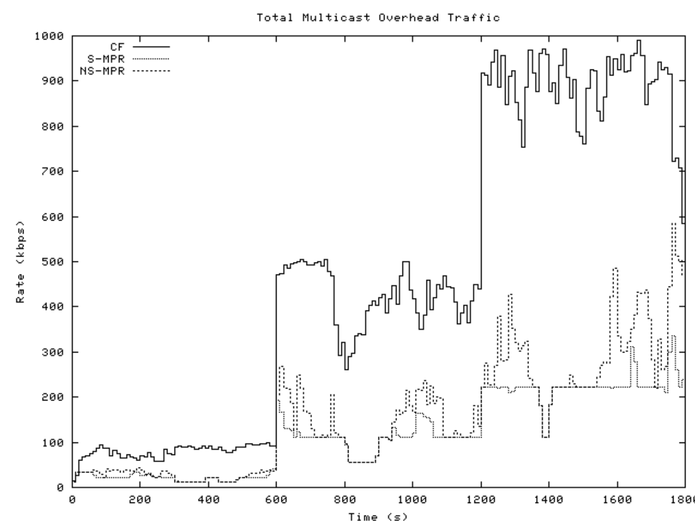

Figure 5: Total Multicast Overhead with Mobility

NETWORK APPLICATION TESTS

In addition to the quantitative test data measurements presented, we successfully conducted a number of functional application experiments using multicast video and VoIP traffic. We have demonstrated the operation of the video conferencing tool (VIC) [VIC01] to perform multicast streaming video in a MANET environment using the forwarding prototype described here. We started these tests by putting 10 nodes in an emulated MANET testbed. Once again each node moved using a random-way point motion model. A source node or set of source nodes provided live video feeds using the simplified MANET multicast forwarding mechanism to all MANET receivers. We increased the video quality rate at the source node and examined when the picture began to degrade at the MANET multicast receivers. As can be directly derived from the previous quantitative results, the more efficient multicast algorithms supported higher effective data rates and associated video quality. This experiment also functionally demonstrated that the MANET multicast flooding engine working with existing multicast capable operating systems and multi-media applications. Additional experiments were performed operating various Voice over IP (VoIP) applications to test functionality and these tests demonstrated similar observed results.

FURTHER ANALYTICAL ANALYSIS

With the mobile network emulator and the working MANET multicast forwarder, we were able to get some interesting preliminary results even with a limited test scenario. We showed that our implementation supports simplified native IP multicast forwarding within a MANET area and that the system can adopt more optimized algorithms (e.g., S-MPR) to decrease the relay node set and support higher sustainable goodput. With the initial emulation experiments we were limited by hardware to examining only small network configurations. To examine additional performance trends as the number of participating network nodes increases we chose a basic analytical approach. We were especially interested in looking at some basic clustering and combined MPR algorithms.

BASIC ANALYTICAL MODELS

We implemented three different analytical flooding methods in our network analyzer: NS-MPR flooding, S-MPR flooding, and a centralized clustering algorithm that creates a single CDS relay node set. The clustering algorithm was used to compare against the efficiency and scalability trends of the S-MPR approach. The analysis examined the expected relay set size (denoted as forwarders along the y-axis) as the number of nodes in the network increases. In this analysis we do not take into account lost packets or overall throughput, only expected flooding overhead in terms of the number relay nodes required to a flood a single application packet. Other important metrics like robustness, forwarding delay, and control overhead needed to support the various methods are not represented as they were in the emulator.

We decided to look at an analysis of both random graph networks and random unit graph networks performance and developed a similar analysis approach taken in [JLMV02]. The random graph networks are defined by (N, p), where N is the number of nodes and p is the probability of a link between any two nodes. We restricted our random unit graph networks to be square with this restriction (N,L) characterizes where N is the number of nodes and L is the length of the side of the grid. There is a link between nodes if the distance between the two nodes is less than one unit. In order to avoid undesirable anomalies in our data from unconnected networks, we generate our test networks at random and then perform an initial check for connectivity. If a network is not connected it is discarded and another network is generated. Each flooding method is run on a successfully connected graph and the resulting number of forwarders to flood data is recorded. Each data point in our graph represents the average number of forwarders a flooding method selected for one hundred randomly generated networks.

ANALYTICAL MODEL RESULTS

In order to compare the analytical model results with the basic working prototype results, we first examined a similar small network of ten nodes. Figure 6 represents ten node networks with increasing connectivity levels. The results show how the various flooding methods reduce redundant forwarding as average connectivity changes. It is shown that S-MPR outperforms both simple cluster-based and NS-MPR flooding over all connectivity probabilities though the corresponding growth trends are relatively similar. This graph supports the emulator findings that fewer messages will be sent using S-MPR rather than NS-MPR flooding, but within this size network NS-MPR seems relatively efficient.

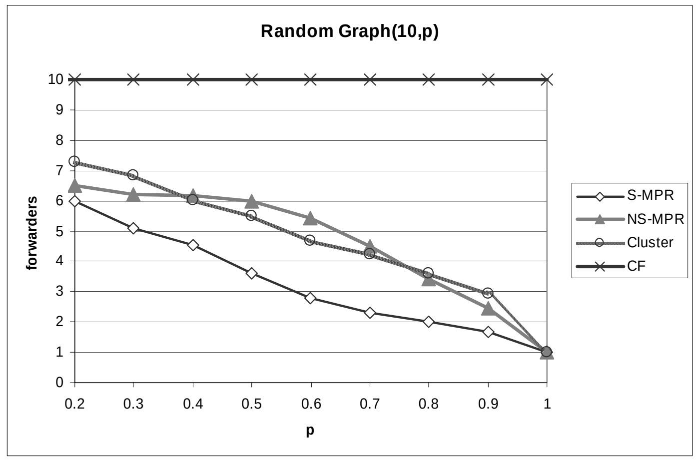

Figure 6: Expected Relay Set Size for 10 node network

Figure 7 shows growth trends as the number of nodes increase in a fixed square area with a diameter of three units using the random unit disc graph model. As shown in the last graph, at ten nodes the overhead increase though apparent is not that significant. However, as the number of nodes in the network increases NS-MPR flooding becomes increasing less efficient at reducing the number of flooding messages. Also, this graph clearly demonstrates the similarity in efficiency trends between S-MPR and a centralized-clustering scheme.

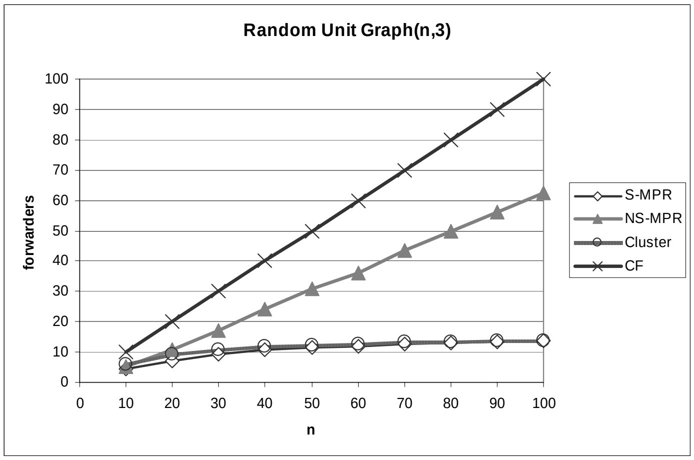

Figure 7: Expected Relay Set Size for Larger Networks

Next we take a look at larger networks in the random unit disc model with density being variable. We can see by the results in Figure 8 that as the network becomes sparser the NS-MPR hybrid relay node mechanism rapidly becomes less efficient than both the basic clustering and S-MPR mechanisms. This result gives us less confidence in the NS-MPR mechanism as an efficient mechanism in larger MANET networks.

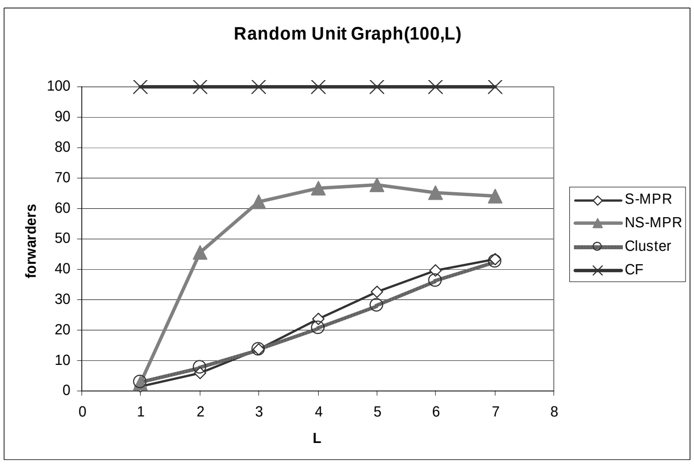

Figure 8: Expected Relay Set Size in RUG with 100 nodes

FURTHER WORK

We have identified a number of areas of ongoing and future work from our present studies. We have demonstrated a working prototype of a simplified MANET multicast routing mechanism based upon improved flooding algorithms. The framework is relatively flexible but the study of additional potential forwarding strategies and design tradeoffs is ongoing. More examination of the packet delivery robustness of various techniques under more complicated mobility scenarios is desirable. We have also pointed out the present duplicate detection mechanism is potentially problematic due to inconsistencies and limitations in IP ID field usage. We plan to investigate additional approaches including possible specialized encapsulation and/or header extension methods. A header extension method is likely appropriate for implementation in IPv6 and we plan to prototype and experiment with such a capability.

CONCLUSION

We have presented an implementation overview of a working protocol prototype that can be used to support simplified multicast IP routing in MANET environments. The simplified model is based upon the adaptation and use of efficient flooding algorithms for mobile wireless networks. In addition to the prototype design, we functionally demonstrated the working code in the NRL mobile network emulator and presented some initial performance results. This framework can be used to support a wide variety of applications desiring multicast routing support within a scoped mobile routing area.

Even for small MANET networks, we demonstrated the potential for flooding optimizations to significantly improve maximum sustainable goodput rates to a group of multicast receivers. We also presented some additional analytical models to compare with our emulation results this helped to validate the present results and to further evaluate performance issues as the network scales. We conclude that, unlike S-MPR, the NS-MPR technique implemented to eliminate previous hop routing dependencies in the forwarding decision does not achieve high efficiency under scaled network conditions. We also presented further technical considerations and future work issues.

While additional research areas remain, we demonstrated that the approach described can be used to provide a simplified multicast routing capability for MANET networks at the IP layer. In addition, we demonstrated a working prototype and the importance of reducing the size of the relay set to improve the maximum achievable goodput within the network.

REFERENCES

[CM99] S. Corson and J. Macker, “Mobile Ad Hoc Networking (MANET): Routing Protocol Performance issues and Evaluation Considerations,” RFC2501, JAN 1999.

[SRS02] C.Santiváñez, R. Ramanathan, I. Stavrakakis, “Making Link-state Routing Scale for Ad hoc Networks,” Proceedings of ACM Mobihoc 2001, Long Beach, CA, Oct 2001.

[HXG02] Xiaoyan Hong, Kaixin Xu, and Mario Gerla, “Scalable Routing Protocols for Mobile Ad Hoc Networks,” IEEE Network Magazine, July-Aug, 2002.

[QVL00] A.Qayyum, L. Viennot,and A.Laouiti, “Multipoint Relaying: An Efficient Technique for flooding in Mobile Wireless Networks,” Research Report RR-3898, INRIA, February 2000.

[MPC00] J. Macker, V. Park, and M.S. Corson, “Mobile and Wireless Internet Services: Putting the Pieces Together,” IEEE Communications Magazine, June 2001.

[PGC00] G. Pei, M. Gerla, and T.-W. Chen, “Fisheye State Routing: A Routing Scheme for Ad Hoc Wireless Networks,” Proceedings of ICC 2000, New Orleans, LA, Jun. 2000.

[P01] Charles Perkins, editor, Ad Hoc Networking, Addison-Wesley. 2001.

[PM00] V. Park, and J. Macker, “Protocol Considerations for Distributed Sensor Networks,” NRL Memorandum Report, NRL/MR/5523—00-8521, Dec 18, 2000.

[CLOV02] T. Clausen, N. Larsen, T. Olesen, and L. Viennot, “Investigating Data Broadcast Performance in Mobile Ad Hoc Networks,” the 5th International Symposium on Wireless Personal Multimedia Communications (WPMC), Oct 2002.

[WC02] Brad Williams, Tracy Camp, “Comparison of Broadcasting Techniques for Mobile Ad hoc Networks,” Proceedings of ACM Mobihoc 2002, Lausanne, Switzerland, June 2002.

[NTCS99] S. Ni, Y. Tseng, Y. Chen, and J. Sheu. The Broadcast Storm Problem in Mobile Ad hoc Networks. Proc. Of Mobicon 99, pgs. 151-162, Aug 1999.

[JLMV02] P. Jacquet, A. Laouiti, P. Minet, L. Viennot, “Performance of multipoint relaying in ad hoc mobile routing protocols,” Networking, May 2002.

[CMW03] W. Chao, J. Macker, J Weston, “NRL Mobile Network Emulator,” Milcom, Oct 2003.

[GJ79] M.R. Garey and D.S. Johnson, “Computers and Intractability: A Guide to the Theory of NP-Completeness.” Freeman and Company, 1979.

[OTL04] R. Ogier, F. Templin, and M. Lewis. Topology Dissemination Based on Reverse-Path Forwarding (TBRPF), IETF RFC 3684, Feb 2004.

[VIC01] http://www-nrg.ee.lbl.gov/vic/.