Evaluation of Distributed Cover Set Algorithms in Mobile Ad hoc Network for Simplified Multicast Forwarding

Joseph Macker
macker@itd.nrl.navy.mil
Justin Dean
jdean@itd.nrl.navy.mil
Ian Downard
iandownard@ieee.org
Brian Adamson
adamson@itd.nrl.navy.mil

PROTEAN Research Group
Information Technology Division
US Naval Research Lab
Washington, DC USA 20375-5000
http://cs.itd.nrl.navy.mil

This paper presents design issues and performance analysis of optimized cover set algorithms supporting Simplified Multicast Forwarding (SMF) of data plane traffic within mobile ad hoc network (MANET) environments. SMF is presently being developed within the Internet Engineering Task Force (IETF) as an experimental specification to provide simplified multicast data dissemination among multi-hop, wireless nodes within peer MANET routing neighborhoods. The SMF protocol allows for a variety of cover set reduction techniques to optimize the simplified data flooding and relaying process amongst routing peers. A variety of distributed algorithms for forming connected dominating sets (CDS) are being considered as mechanisms to reduce the cover set. This paper presents modeling and analysis work of various CDS relay set algorithms and demonstrates working code within an SMF implementation. We provide a brief problem background, discuss models and scenarios, compare various algorithms, and then summarize observations as well as discuss future work. A main purpose of the paper is to begin examining the robustness of the algorithms to mobility and increasing multicast traffic load. We examine the results against a classical flooding baseline for comparison. We observe similar efficiency and robustness performance for several forwarding algorithms of interest.

I. Introduction

Design, specification, and implementation progress has been made in realizing more efficient ways of flooding control messages within mobile ad hoc network (MANET) wireless routing protocol designs. For example, within MANET routing protocols, such as Internet Engineering Task Force (IETF) experimental Request for Comments (RFCs) 3626 [1] and 3684 [2], the flooding process of delivering control packets to all peer nodes is improved by the dynamic election of a reduced relay set. Until recently, the application of these concepts and algorithms has typically been applied to control plane message flooding. However, there is growing interest in flooding data plane traffic efficiently within Internet Protocol (IP) based MANET systems. For small-to-moderate size MANET networks this can provide an alternative localized multicast routing approach.

The development of an experimental specification for Simplified Multicast Forwarding (SMF) is at present an active working group item within the IETF MANET Working Group [3][4][5]. Additionally, early experiments have been done with voice over IP (VoIP), video streaming, chat, and distributed sensor applications using SMF. Two main objectives of SMF are:

1. Develop a specification framework for simple IP multicast packet forwarding on MANET interface types, including duplicate packet detection mechanisms.

2. Apply known efficient flooding or relay set mechanisms to SMF for further reducing contention and congestion in wireless multi-hop scenarios

A near term goal of SMF is to develop a flexible framework supporting multiple relay set mechanisms for ongoing experimentation. This paper presents an

analysis of early SMF experimentation applying various optimized relay set algorithms. The results presented here are based on simulation and complement related work in analytical and emulation studies previously reported [6]. We focus our investigation on connected dominated set (CDS) algorithm performance in terms of non-duplicate traffic delivery ratios under a variety of mobility and loading scenarios. Future extensions to SMF may support receiver pruning and other group specific optimizations. The study here is limited to examining the SMF core approach of delivering IP multicast data to all participating peer SMF routers. The main purpose is to provide a better understanding of the robustness and efficiency of CDS algorithms being specified under both mobile and increased traffic loading conditions.

A dominating set (DS) in a network 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. Our analysis within this paper targets the examination of distributed SMF CDS operation and robustness under both mobility and increased multicast traffic loading. Analysis of both mobility and traffic loading conditions is important to provide a better understanding of anticipated deployment performance. Previous analytical studies provide some understanding of the properties of distributed CDS algorithm cover set efficiency, yet analyses under mobile and dynamic loading conditions are lacking. The overall effectiveness of a CDS algorithm is more than a matter of relay set minimization but is also related to algorithm robustness to network wide loading and topological dynamics. Localized network neighbor discovery mechanisms represent another important design consideration when examining this area of MANET system design robustness. While further work could examine alternate neighbor discovery methods, we use a common approach for all CDS algorithms studied to provide a normalized comparison. We discuss more details of the testing conditions prior to the analysis section.

II. Background

The SMF baseline design limits the scope to best effort multicast forwarding and its applicability is intended to be constrained within a peer MANET routing area. SMF provides a forwarding process compatible with different neighborhood discovery protocols and relay set selection algorithms. Multicast receiver group-based enhancements (e.g., cover set pruning) are possible and intended to be supportable but are not addressed in this study. As a basic flooding capability, it also supports Classical Flooding (CF). In CF, every node forwards each unique packet it receives once. In CF operation, there is no dependency on any relay set algorithm or neighborhood topology information. However, it is hypothesized that more efficient flooding techniques will be preferred due to expected gains in network efficiency and reductions in wireless congestion and contention as traffic load increases. When examining more optimized variants of flooding, we are interested in improvement trends and therefore we use CF as a baseline for comparison throughout our analysis. We are particularly interested in performance comparisons as topology dynamics and loading is varied. Even in small networks such performance tradeoffs are not well understood.

Local network link or neighborhood information is generally used in the dynamic determination of a reduced set of forwarding nodes. In many cases, it is expected that local discovery functions may be provided by mechanisms running concurrently with SMF (e.g., lower layer, MANET unicast protocol, MANET Neighborhood Discovery). Such neighborhood discovery design variants are supported by SMF and will likely affect efficiency and robustness of cover set election. This study does not concentrate on examining neighborhood discovery performance. To conduct CDS comparison independent of discovery processes, we fix the neighborhood discovery process used throughout the study. In particular, we used a common IP layer neighborhood discovery mechanism that provides and maintains local two hop information for all CDS algorithms. This mechanism is based upon an implementation of the HELLO protocol specified within Optimized Link State Routing (OLSR) [1] and provides 2-hop bi-directional neighbor detection and maintenance. The SMF duplicate packet detection (DPD) scheme used is also fixed throughout the study and is based upon packet sequence markings.

Previous work on MANET flooding and reduced relay set mechanisms has been done and our work leverages related accomplishments [6]. In [7], a proposed taxonomy of flooding algorithms for use in MANET environments was presented and the work examined performance issues related to various approaches. Our work does not consider some of the algorithm classes presented in that paper. Rather, we focus on the use of classical flooding and distributed CDS approaches which guarantee full graph connectivity under ideal conditions. The design tradeoffs are further complicated by traffic loading, topologi-

cal classes, and the robustness of packet delivery and set election under mobile scenarios. The actual protocol implementation for IP multicast forwarding based upon these flooding algorithms raises additional design trade offs and issues. This includes maintenance of protocol state, duplicate packet detection mechanisms, packet processing requirements, expected traffic patterns (e.g., one-to-many vs. many-to-many), and requirements for robust performance. The associated protocol signaling and maintenance required by any optimized mechanism must be traded off against the need for these techniques to operate under congested and mobile conditions. Many previous studies focusing purely on flooding for routing control have not examined the robustness of such techniques when significant traffic loading occurs along with mobility. It is important to better understand CDS traffic loading and delivery performance in the context of SMF because the applicability of multicast application routing is broader than routing control plane traffic as mentioned in the introduction.

While in some network and traffic scenarios CF can be a highly robust delivery technique, it is well known that within dense networks, CF results in a significant number of redundant transmissions often referred to as the broadcast storm problem [8]. Reducing unnecessary channel contention within a MANET can significantly improve network performance, especially as traffic load increases. Therefore, minimizing the number of required relay nodes is a heightened design goal for multi-hop wireless networks. Unfortunately, reducing the number of relay nodes in a MANET environment may also decrease the robustness of packet delivery in a mobile topology. There exists an interdependent design trade off between relay efficiency and delivery robustness that is scenario and system dependent and should be considered carefully. Our analysis is focused on contributing to a better understanding of performance and early implementation experience within this design space.

The NRL SMF reference prototype implementation is used throughout the studies presented here. The implementation uses the PROTEAN research group protocol prototyping library (Protolib) C++ toolkit [10] to facilitate a working cross platform implementation in MacOS, Linux, BSD, and Windows as well as simulation systems such as NS-2 and OPNET. The NRL SMF code used in this implementation is also publicly available [11]. The Protolib MANET kernel is an extension to Protolib for NS-2 that provides packet routing and forwarding functions for use within the NS-2 framework. It was originally designed specifically for porting different unicast MANET routing implementations into NS-2, but has evolved to support SMF capability. This code provides the following functions:

- Protocol independent unicast routing table
- Basic NS-2 packet forwarding for unicast or multicast packets
- SMF duplicate packet detection
- Basic network interface multiplexing between application agents
- Transmission of packets to the NS-2 interface queue

As mentioned earlier, we use a common MANET neighborhood discovery process to facilitate CDS election within this study. Our OLSR implementation code [12] has been extended to produce CDS information beyond the standard OLSR Multipoint Relay (MPR) election process described in the OLSR specification and an interface mechanism is used between OLSR and the SMF module to dynamically share this information. This enables a common neighbor link discovery mechanism and provides for several CDS variants. In our testing we have also disabled related unicast routing control plane traffic (e.g., topology control (TC) packet flooding) and we adapt the OLSR implementation solely as a CDS election and maintenance mechanism.

III. Overview of CDS Algorithms Studied

A minimum connected dominating set (MCDS) is a connected dominating set with the smallest cardinality of any CDS in a graph. Calculating the MCDS is NP-hard [13] yet distributed algorithms have been developed that attempt to construct CDS that approximate MCDS characteristics. Here we provide an overview of the algorithms that we have analyzed, implemented, and used in this study. We do not claim this is comprehensive, but it represents classes of several algorithms presently being specified and experimented with for Internet protocol integration and consideration. The following five CDS algorithms are examined within this study:

- Classical Flooding (CF)
- Source-specific Multipoint Relay (S-MPR)
- Non-source-specific MPR (NS-MPR)

- Essential CDS (E-CDS)
- MPR-based CDS (MPR-CDS)

Pseudocode for each of these algorithms is shown in Figure 1. We have implemented and used "Classical Flooding" as a baseline for comparison. This technique represents a basic approach to data flooding with duplicate packet detection. CF is the only algorithm in this study which does not require some form of neighbor discovery. The most significant challenge in implementing this protocol is performing duplicate packet detection, a requirement for all SMF algorithms included in this study.

The S-MPR forwarding algorithm is specified for use in the OLSR routing protocol [1]. Two-hop neighborhood discovery is performed using HELLO message signaling. SMF nodes then select and signal a subset of their bi-connected one-hop neighbors as multi-point relay (MPR) nodes. This subset provides flooding coverage to all two-hop neighbors. An S-MPR node forwards if and only if

- it receives a unique multicast packet from any of its bi-connected neighbors
- the neighbor from which the packet was received has selected the node as an MPR.

Thus, S-MPR requires duplicate detection, MPR selection signaling, as well as packet previous-hop knowledge. S-MPR is the only algorithm included in this study which requires previous-hop knowledge to determine if the packet should be relayed. Since the S-MPR algorithm is previous-hop dependent, different relay sets may result for different packet sources. Multiple source-specific relay sets may be beneficial for overall system wide performance. For example, a distributed network load may help sustain network life for a power-constrained system. However, the per-packet, previous-hop information requirement of S-MPR may be burdensome for some network deployments. S-MPR can guarantee minimal hop paths between all nodes in the network, while maintaining an efficient CDS [14]. S-MPR is the only algorithm included in this study which exhibits both of these properties.

The NS-MPR algorithm is a slight variant on S-MPR. By using the simple technique of combining source-specific MPRs one CDS is formed. In this case, a node forwards every unique packet if and only if

- it has been selected as an MPR by any other node regardless of the previous hop of the packet.

This approach removes the requirement of packet previous-hop knowledge needed for S-MPR while maintaining minimal hop path forwarding. However, the authors realize NS-MPR has poor scaling properties [6] and it is presented for comparison purposes only.

The E-CDS implementation used in this study is based on the essential connected dominating set (E-CDS) algorithm described in a proposal for MANET extensions to OSPF using CDS flooding [15]. This study does not examine the redundant relay set functionality included in the OSPF extensions or the extra reliability which may be afforded by this extension. E-CDS performance and behavior also closely resembles that of restricted knowledge Rule K described in [9], implementation however was based on the proposed OSPF extensions. The E-CDS algorithm produces a common shared set of relay nodes for all nodes in the network similar to NS-MPR.

Nodes select themselves as relays using neighborhood router priority information. Priority values need not be unique and can be a combination of values such as power level and address values. Although uniqueness is not a requirement, duplicate values can result in a greater number of forwarders. For nodes to correctly assign themselves as relays, priority values need to be learned within a two-hop neighborhood. Differing router priority selection methods, while using E-CDS, can result in differing CDSs for a given network.

In this study a node's router priority was based on the number of one-hop neighbors and then the node's address. This metric required additional information to be shared within the two-hop neighborhood to convey nodal density. Using nodal density resulted in a smaller average number of forwarders for the networks scenarios examined. Using alternative metrics as the priority value can result in different E-CDS performance, we do not examine this issue further here. The signaling processes were included in HELLO messages without increasing the number of messages transmitted. After neighborhood discovery and signaling is performed E-CDS nodes select themselves as relays if and only if

- the node's router priority is greater than all its two-hop neighbors,
- or there does not exist a path from the highest priority neighbor to all other one and two hop neighbors using only nodes with greater priorities as relays.

Once a node has selected itself as a relay all unique multicast packets are rebroadcast. E-CDS, unlike S-MPR, does not guarantee minimal hop paths for end-to-end connections. From previous studies, E-CDS scaling closely resembles S-MPR scaling [6]. Because E-CDS uses a shared CDS, there may be higher traffic concentration within the network forwarding paths compared to source based approaches.

The final algorithm studied is the MPR-based Connected Dominating Set (MPR-CDS) algorithm described in [16]. MPR-CDS reduces the number of forwarding nodes to a more efficient subset of MPRs than the simple NS-MPR approach previously described. MPR-CDS requires that nodes know unique network addresses (or any other unique ordering identifier) for each node within their two-hop neighborhood. After neighborhood discovery, a node using MPR-CDS will forward all unique packets if and only if

- the node's identifier is higher than all its one-hop neighbors,
- or, it has been selected as an MPR by the node that has the highest identifier in its one-hop neighborhood.

Like E-CDS, this approach results in a common relay set, and does not guarantee minimal hop paths. It also has no requirement for previous hop knowledge similar to other shared CDS algorithms. MPR-CDS has similar scaling properties to both E-CDS and S-MPR [5].

IV. The Simulation Model in NS-2

In this section we layout the simulation parameters used in our SMF analysis. Performance was measured over a range of data rates as a ratio of observed goodput versus maximum goodput(e.g., all non-duplicate packets received) aggregated across all multicast destinations. For this study, all mobile receivers act as multicast receivers. Aggregate goodput is defined as the number of non-duplicate bits per unit of time received at all destinations (or all group nodes, in the case of multicast) [17]. We use aggregate goodput as a primary metric for measuring the effectiveness of SMF protocols because it expresses aggregate end-to-end data loss. When measured over a range of increasing traffic loads, it is one measure of end-to-end performance. Again, this early study focused on examining aggregate deliverable data statistics under both mobility and increased traffic loading. Future studies would be needed to examine other statistical metrics of interest.

Definitions:

n_0: The node performing the computation
N_1: The set of 1−hop neighbors of n_0
N_2: The set of 2−hop neighbors reachable by n_0 excluding n_0 and all $n \in N_1$
$N_2(y)$: The subset of N_2 neighbors reachable by node y, where $y \in N_1$
$N_1(z)$: The subset of N_1 nodes which have direct connections to z, where $z \in N_2$
$MPRs$: The subset of N_1 which have been selected to forward packets from n_0
$MPR\text{–}Selectors$: The subset of N_1 for whom n_0 has been selected to forward packets
$RtrPri$: an expression of router priority. For example, use $|N_1|$ as a router's priority then break ties with its address.

Procedure for node n_0 to select node n as an MPR:

1. Add n to the $MPRs$ set.
2. Remove node n from N_1.
3. For each y in $N_2(n)$, remove y from N_2.
4. Calculate $N_1(z)$ for all nodes $z \in N_2$.
5. Calculate $N_2(y)$ for all nodes $y \in N_1$.

S-MPR selection algorithm:

1. Calculate $N_1(z)$ for all nodes $z \in N_2$.
2. Calculate $N_2(y)$ for all nodes $y \in N_1$.
3. For each $z \in N_2$ where $|N_1(z)| \equiv 1$, select the node in $N_1(z)$ as an MPR.
4. While $N_2 \neq \emptyset$ select the node y, with the largest $|N_2(y)|$, as MPR.
5. The node shares its $MPRs$, and nodes in the MPR set add the sharing node to their $MPR\text{–}Selectors$ set.
6. Nodes forward all unique multicast packets which are first received from a node in their $MPR\text{–}Selectors$ set.

NS-MPR selection algorithm:

1. Perform stesp 1−5 of S–MPR operation.
2. If N_0's $MPR\text{–}Selectors$ set is nonempty, then N_0 selects itself as a forwarder.

MPR-CDS selection algorithm:

1. Perform steps 1−5 of S–MPR operation.
2. If n_0 has the highest address in $N_1(n_0)$, and n_0's $MPR\text{–}Selectors$ set is nonempty, then n_0 selects itself as a forwarder.
3. If the largest address in $N_1(n_0)$ is in n_0's $MPR\text{–}Selectors$ set, then n_0 selects itself as a forwarder.

E-CDS selection algorithm:

1. If n_0 has a larger value of $RtrPri$ than all nodes in $N_1 \cup N_2$, then n_0 selects itself as an MPR.
2. Otherwise, let R_{max} be the node in N_1 that has the largest $RtrPri$. If there does not exist a path from R_{max} to every other node in $N_1 \cup N_2$ using only nodes that have $RtrPri$ larger than n_0's, then n_0 selects itself as a forwarder.

Figure 1: Pseudocode for SMF Algorithms.

The following simulation parameters are common to all the scenarios presented in this section:

- 25 nodes move in a random walk within a 710 meter square plane for 10 minutes of simulated time. Random walk was chosen over random waypoint so that nodal density would remain relatively constant throughout each simulation [18][19][1]. Ten minutes of simulated time was observed to be sufficient to obtain a stable average measure of goodput statistics in our scenarios. Node directions are reset once per second with a random angle distributed evenly between -18° and 18° relative to the current vector. Node speeds are adjusted once per second within a random range distributed evenly between $\pm 0.5m/s$. Speeds which exceed the maximum or minimum threshold ($0m/s$) specified in the scenario are rounded to the highest or lowest speed, respectively. Motion vectors are reflected off the boundaries of the plane. Initial node placement and motion vectors are randomly generated with even distribution within the running scenario's parameters of network size and node speeds.

- A two-ray ground reflection scheme is used to model radio-propagation as an ideal circle whose 250 meter radius is an absolute limit on signal range.

- The medium access model is the default 802.11 module within the NS-2 environment. The MAC/radio model may significantly effect end-to-end traffic loading results for any multicast mechanism so any future experiments should be reexamined with that in mind.

- OLSR unicast routing is enabled, and topology control messages are sent at 2.5 ± 1.875 second intervals, and OLSR Hello messages are sent at 0.5 ± 0.25 second intervals.

- Source nodes begin multicasting data starting with random uniform distribution after approximately a 6 second scenario startup time. A 5 second startup window was observed to be sufficient to converge the initial localized neighborhood discovery process given the parameters used.

- Multicast data is sent in 256 byte UDP datagrams with a time-to-live of 255 hops. Source nodes send these datagrams at constant frequencies which approximate a data rate between 1kbps and 201kbps.

[1] Previous work [6] has established that network density affects the efficiency of relay sets. In our study, we maintained density as a controlled variable.

There is one additional caveat to how our mobility scenarios were generated. The only cases where it was feasible to guarantee a contiguous network was where the maximum node speed was specified as $1m/s$. In the other scenarios where maximum node speed was specified as $2m/s$, $4m/s$, $8m/s$, $12m/s$, $16m/s$, or $20m/s$, fragments could not be avoided without changing other controlled variables such as network density or randomness. While network fragmentations in these scenarios were small in size and overall duration, they did decrease the goodput achievable by SMF in a way that is independent from the CDS algorithms. Although this artificially lowered aggregate goodput results for datapoints in Figures 3 and 4, the relative comparisons between protocols are valid when connectivity is biased equally between scenarios.

Table 1 expresses the nodal connectivity of our mobility scenarios as an average of the same random walks used to acquire datapoints for Figures 3 and 4. Ideally, SMF protocols would only have been analyzed in scenarios where all 25 nodes remain connected in a single network partition for the duration of the simulation (as was achieved in the $1m/s$ category). However, for cases where that was not possible, the average nodal connectivity was only slightly offset from this ideal.

Mobility Scenario	Average Partitions	Average Network Nodal connectivity
$1m/s$	1	25
$2m/s$	1.13	24.57
$4m/s$	1.22	24.1
$8m/s$	1.17	24.34
$12m/s$	1.15	24.47
$16m/s$	1.17	24.36
$20m/s$	1.18	24.36

Table 1: This table expresses how often each mobility scenario is contained within a single partition of 25 nodes randomly moving for a duration of 10 minutes.

There are several reasons we feel it is important to include average nodal connectivity information along with our results. One reason is we observed increasing mobile network fragmentation events under certain scenario conditions. For example, observe how the network nodal connectivity average of the $4m/s$ scenarios was lower than the others. This helped indicate

that the $4m/s$ scenario experienced more fragmentation than other scenarios and this contributed to the effect in Figure 4, where SMF aggregate goodput was analyzed under increasing mobility as traffic loading remained constant. We hypothesize from these results that SMF strategies using reduced relay sets tend to benefit more (than strategies of nonoptimized relay sets) from the decreased contention that results during network fragmentation under heavy traffic loading. The relationship between Table 1 and Figure 4 also suggests that average network nodal connectivity may be a limiting factor for aggregate goodput in MANETs with heavily loaded SMF traffic. More thorough investigation of this matter is left for future work.

V. NS-2 SMF Simulation Results

SMF and various CDS algorithms were tested in several scenarios of increasingly congested MANETs. Congestion was introduced by increasing the multicast traffic load, as well as increasing the quantity of traffic sources. Figures 2(a) and 2(b) show what percentage of an offered load is delivered by the classical flooding, NS-MPR, S-MPR, E-CDS, and MPR-CDS protocols in relatively slow changing networks where nodes move with a maximum speed of $1m/s$. These results were produced by multicast streams sent from 1 and 4 source nodes, respectively. The number of sources remained unchanged for the duration of each scenario, and each scenario was configured to run through 10 minutes of simulated time. Aggregate goodput is shown relative to increasing load, as expressed by the data rates shown on the x-axis. Each data point represents goodput averaged over 10 simulation runs with identical traffic models but randomized mobility. The errorbars illustrate the deviation of goodput observed over those runs. For fairness, each protocol was simulated with the same set of 10 random mobility scenarios. Observe how the optimized CDS algorithms deliver the highest aggregate goodput as traffic load increases. There is however a slight delivery ratio degradation for such algorithms in lightly loaded cases. This indicates the benefits of additional nodal forwarding redundancy to improve reliability can be quickly outweighed by increased traffic loading.

The next set of tests performed examined algorithm robustness to node mobility under light traffic loading. Figure 3 summarizes these results and demonstrates how the CDS algorithms perform in lightly loaded mobile networks. The mobility patterns used to capture the data points plotted in this figure have maximum node speeds shown on the x-axis. The average node speeds are approximately half the maximum speeds, and the errorbars represent the variance of goodput observed over 10 repeated scenarios with the same random walk mobility patterns as used in Figures 2(a) and 2(b). The traffic in these scenarios was generated by a single source offering a 1kbps stream flooded to every node in the 25 node scenario. Observe how increasing node mobility does not visibly reduce the achievable aggregate goodput of classical flooding nor NS-MPR. This is likely due to the more redundant nature of their forwarding algorithm at the expense of additional overhead. Figure 1 shows this feature to be potentially detrimental as traffic loading increases.

Overall, our results suggest that different optimized

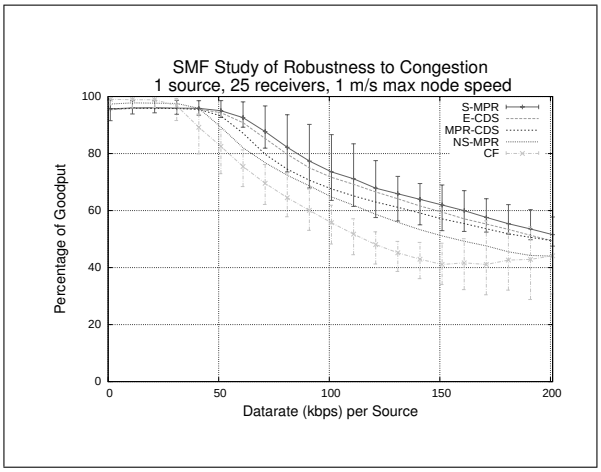

(a) 1 source

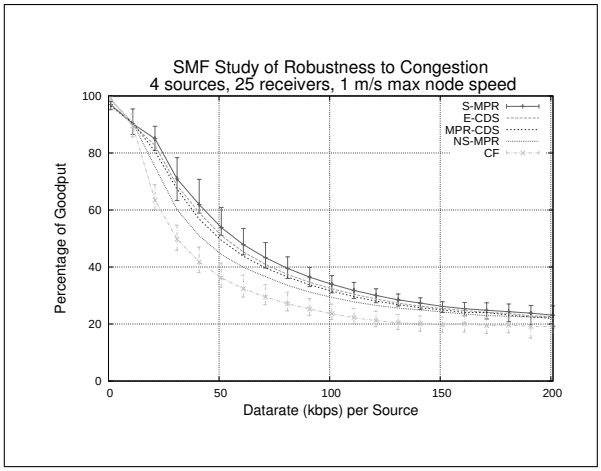

(b) 4 sources

Figure 2: Congestion affects on goodput performance of multicast streams delivered by various SMF schemes in relatively static networks.

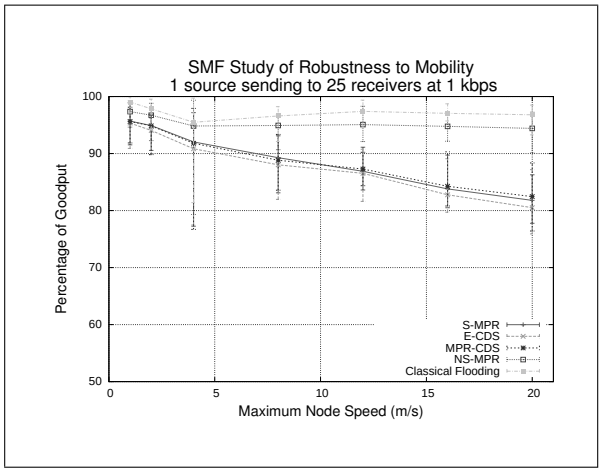

Figure 3: Comparison of SMF protocols performance with different patterns of node mobility in uncongested networks.

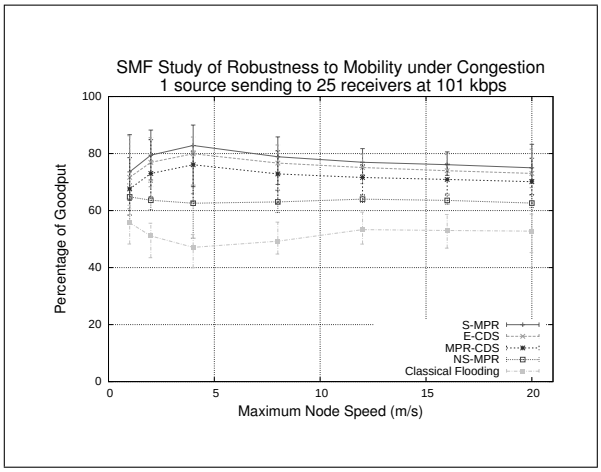

Figure 4: Figure illustrating that SMF protocols are not significantly impacted by node mobility in highly congested networks.

CDS forwarding policies will not significantly differ in their overall performance under varying mobility (at least at the scale we tested). Classical and redundant flooding has a delivery ratio advantage only in the lightly loaded traffic case. Thus, we would expect the impact of traffic loading to dominate goodput more than node mobility and link failures. This expectation is confirmed by Figure 4, which shows that increasing node mobility does not significantly reduce the achievable goodput in highly loaded conditions. Figure 4 illustrates goodput for the same mobility scenarios plotted in Figure 3, but this time multicast traffic is generated by a single source node offering a 101kbps stream via SMF to every node in the 25 node MANET. This datarate exceeded the physical channel's bandwidth capacity, thereby causing increased congestion and contention in the network. Observe how increasing node mobility in this network does not visibly reduce the achievable aggregate goodput.

VI. Future Work

Optimizing SMF CDS algorithm packet delivery performance is somewhat scenario dependent and design decisions could benefit from further study. Here we have presented early analysis of various implemented CDS algorithms undergoing both mobility and increased traffic loading conditions. Future work is planned to better understand the individual flow statistics versus the aggregate performance. Understanding this will allow a better determination of anticipated application performance requirements and the potential for next generation reliable multicast transport protocols to be applied end-to-end. Another interesting aspect of SMF cover set mechanisms which should be studied is the tradeoff between more redundant forwarding sets and robustness. As shown in Figure 1, the value of redundancy in improving aggregate packet delivery is often detrimental as traffic loading increasing. Such approaches have been examined for routing control redundancy where loading can be more tightly managed, but their usefulness in multicast forwarding scenarios with potentially higher traffic loads is less clear. Results would also be useful for control plane traffic of a more bursty nature within on demand protocols.

Analyzing the costs and benefits of supporting multicast group-specific delivery optimization on top of protocols like SMF is another attractive topic for future research. Sustaining robust and efficient multicast routing in MANETs while nodes arbitrarily join and leave groups generally adds additional signaling and state within the mobile network. Design trade offs exist between robustness, overhead, and efficiency when considering such enhancements [20]. These issues cannot be resolved without more scenario specific studies including scale of interest, traffic distribution, and frequency of membership changes. While group management can help prevent unnecessary packet forwarding (and thereby reduce protocol overhead), maintaining those groups becomes more difficult as the underlying network topology becomes more dynamic. Extensions have been made to Ad hoc On-demand Distance Vector (AODV) and OLSR (called MAODV [21] and MOLSR [22]) which include group management for multicast routing in MANETs. Other protocols including the On-Demand

Multicast Routing Protocol (ODMRP) [23] were also designed for highly dynamic, energy constrained, and bandwidth limited MANETs. More examination of specific operational and traffic scenarios versus protocol behavior could be made to better understand future trade offs in MANET specific multicast deployments.

VII. Conclusions

We implemented five different CDS algorithms for SMF, including classical flooding, NS-MPR, S-MPR, E-CDS, and MPR-CDS. We have presented some rationale for the examination of CDS algorithms within a working SMF implementation. We explained our initial working code and its NS2 implementation which is publicly available online at [11]. We presented and discussed initial performance evaluations of SMF CDS algorithms in networks of varying load and mobility. We summarize our present conclusions as follows:

- S-MPR, E-CDS, and MPR-CDS are effective approaches for SMF over the range of congested and dynamic networks studied in this paper.
- S-MPR, E-CDS, MPR-CDS, NS-MPR, and CF all experience reasonable robustness to changes in topology caused by network mobility alone. This is encouraging regarding the potential application of more optimized algorithms like E-CDS and S-MPR for multicasting application data.
- Under increasing traffic loads and mobility the optimized algorithms (S-MPR, MPR-CDS, E-CDS) are clearly more effective in average non-duplicate packet ratio delivery than simpler but more redundant approaches like CF.

In addition to the analysis results we have also demonstrated that SMF can easily support multiple forwarding strategies and CDS protocol types within a working implementation. In summary, the most encouraging result is that shared cover set approaches such as E-CDS and MPR-CDS demonstrate similar robustness and effectiveness within the scenarios studied as compared to source specific S-MPR. This indicates that either approach can likely be used with effectiveness. In fact, SMF has been used in both cases when working in concert with OLSR unicast routing (S-MPR) and a Boeing Quagga [24] implementation of MANET-OSPF (E-CDS variant). Future work is planned to examine additional performance and scenario specific issues.

Acknowledgment

The authors would like to thank the MANET SMF design team members and Richard Ogier for technical discussions and interactions related to this work. Part of the work described here was funded by the Office of Naval Research (ONR).

References

[1] Clausen, T. and P. Jacquet, "Optimized Link State Routing Protocol", IETF RFC 3626, 2003.

[2] Ogier, R., Templin, F., and M. Lewis, "Topology Dissemination Based on Reverse-Path Forwarding", 2003.

[3] J. Macker. "Simplified Multicast (SMF) Progress/Issues",
Proceedings of the 62nd IETF, March 2005.
`http://www3.ietf.org/proceedings/05mar/slides/manet-3.pdf`

[4] C. Perkins. "SMURF - Simplified Manet Multicast Routing/Forwarding", Proceedings of the 63rd IETF, Aug 2005.
`http://www3.ietf.org/proceedings/05aug/slides/manet-8.pdf`

[5] B. Adamson, J, Dean. "Simplified Multicast Forwarding (SMF) Update",
Proceedings of the 64th IETF, November 2005.
`http://www3.ietf.org/proceedings/05nov/slides/manet-6.pdf`

[6] J. Macker, J. Dean, W. Chao, "Simplified Multicast Forwarding in Mobile Ad hoc Networks," IEEE MILCOM 2004 Proceedings, November 2004.

[7] Williams, B. and T. Camp, "Comparison of Broadcasting Techniques for Mobile Ad hoc Networks," Proceedings of ACM Mobihoc 2002.

[8] 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.

[9] Wu, Dai: "An extended localized algorithm for connected dominating set formation in ad hoc wireless networks," IEEE trans. on parallel and distributed system, vol. 15, issue 10, October 2004.

[10] The Protean Protocol Prototyping Library, http://cs.itd.nrl.navy.mil/work/protolib/

[11] The NRL Simplified Multicast Forwarding (SMF) Engine http://cs.itd.nrl.navy.mil/work/smf/

[12] NRL's OLSR implementation for NS-2. http://cs.itd.nrl.navy.mil/work/olsr/

[13] Garey, M. and D. Johnson, "Computers and Intractability: A Guide to the Theory of NP-Completeness.", Freeman and Company, 1979.

[14] Jacquet, P., Laouiti, V., Minet, P., and L. Viennot, "Performance of multipoint relaying in ad hoc mobile routing protocols", Networking, 2002.

[15] Ogier, R., "MANET Extension of OSPF Using CDS Flooding", Proceedings of the 62nd IETF, March 2005. http://www3.ietf.org/proceedings/05mar/slides/ospf-5/sld1.htm

[16] C. Adjih, P. Jacquet, L. Viennot. "Computing Connected Dominating Sets with Multipoint Relays", INRIA RR-4597. August 2004.

[17] D. Newman. Benchmarking Terminology for Firewall Performance. IETF RFC-2647, August 1999.

[18] Christian Bettstetter, Giovanni Resta, Paolo Santi. "The Node Distribution of the Random Waypoint Mobility Model for Wireless Ad Hoc Networks," IEEE Transactions on Mobile Computing, vol. 02, no. 3, pp. 257-269, July-September, 2003.

[19] Jungkeun Yoon, Mingyan Liu, and Brian Noble. "Random waypoint considered harmful," In Proceedings of IEEE INFOCOM'03, vol. 2, pp. 1312–1321, March 2003.

[20] Thomas Kunz, Ed Cheng. "On-Demand Multicasting in Ad-Hoc Networks: Comparing AODV and ODMRP," 22nd IEEE International Conference on Distributed Computing Systems, 2002.

[21] E.M. Royer and C.E. Perkins, "Multicast operation of the ad-hoc on-demand distance vector routing protocol," Proceedings of the 5th annual ACM/IEEE international conference on Mobile Computing and Networking (MobiCom), 1999, pp. 207-218.

[22] A. Laouiti, P. Jacquet, P. Minet, L. Viennot, T. Clausen, C. Adjih, "Multicast Optimized Link State Routing," INRIA research report RR-4721, February 2003.

[23] S.-J. Lee, M. Gerla, C.-C Chiang, "On Demand Multicast Routing Protocol", In Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC), pages 1298-1304, New Orleans, LA, September 1999.

[24] The Boeing OSPF-MANET Extensions to Quagga. http://hipserver.mct.phantomworks.org/ietf/ospf/

Biographies

Joseph P. Macker is a senior communication systems and network research scientist within the Information Technology Division at the Naval Research Laboratory in Washington, D.C. He presently leads the Protocol Engineering and Advanced Networking (Protean) Group that is investigating adaptive networking solutions for both mobile wireless and wired networking architectures. He holds an M.S.E.E. from George Washington University in Communications Theory and a B.S.E.E. from the University of Maryland, College Park. His primary research interests have been in the following areas: self-organizing network protocol and architecture design, multicast technology and data reliability, mobile wireless networking, cooperative software agents in ad hoc environments, multimedia networking, and adaptive sensor networking.

Justin Dean has been a network and communication systems research scientist within the Information Technology Division at the Naval Research Laboratory (NRL) in Washington, D.C since 1999. He holds an Bachelor's of Science in Computer Engineering from the University of Maryland, College Park. Mr. Dean's research interests include algorithm performance of Mobile Ad-hoc Networking (MANET) protocols, extension development for standard networking protocols, development of analytical models, and exploration of network visualization techniques. He is an active participant in the Internet Engineering Task Force (IETF) specifically within the MANET working group.

Ian Downard graduated from the University of Missouri-Rolla with an M.S. in Computer Engineering in 2002. After serving 6 years as a research sci-

entist within the Information Technology Division at the Naval Research Laboratory, he transitioned to the private sector where he currently works as a software engineer for OPNET Technologies. His research interests involve developing tools for analyzing resource discovery and autoconfiguration techniques within MANET environments. Mr. Downard is a member of IEEE and the Sigma Xi scientific research society.

Brian Adamson has been involved in research in radio communications and data networking at the Naval Research Laboratory (NRL) since 1984. He holds an M.S.E.E. from George Washington University in Communications Theory and a B.S.E.E. from West Virginia University. Mr. Adamson's current research interests include data and multimedia network transport, group communications, dynamic routing for wireless networks, and peer-to-peer networking. He also actively participates in the Internet Engineering Task Force (IETF) in the areas of Mobile Ad-hoc Networking (MANET), IPv6, reliable multicast transport, and is currently serving as a chair of the IETF Reliable Multicast Transport (RMT) Working Group.