Elastic Multicast: Design Extensions and Experimentation Results

Brian Adamson, Joseph P. Macker, Jeffery W. Weston
Information Technology Division, Naval Research Laboratory, Washington DC
Email: brian.adamson@nrl.navy.mil, joseph.macker@nrl.navy.mil, jeffery.weston@nrl.navy.mil

***Abstract*—This paper describes a new implementation of the Elastic Multicast (EM) protocol including new design enhancements for improved dynamic operation. The paper also presents additional performance data collected from emulation-based mobile network experiments. EM is a low complexity extension to Simplified Multicast Forwarding (SMF) that adds group-specific dynamic pruning of the SMF-based multicast forwarding mesh for higher rate traffic flows. It therefore reduces overhead by pruning the SMF relay sets in areas where no receivers exist.**

Our experimental emulation results show that, under a variety of mobility conditions and multicast group distribution patterns, EM maintains SMF-like data delivery robustness while significantly reducing overhead. We also demonstrate that a new design feature that provides preemptive ACK messages for active receiver groups leads to lower loss under mobility and sparse receiver groups. Based upon the results, we consider this feature critical to be included in any future EM design. We also present experimental results examining the performance of EM with classical flooding (CF) and connected dominating set (CDS) relay modes. We show, for the experiments examined, that CF provides some reduced loss with minimal additional overhead when used with EM. We also discuss future work and ongoing issues.

I. Introduction

Maintaining tree-based structures to support multicast routing in a mobile ad hoc network (MANET) is often not as effective as it is in wired networks due to the reduced reliability and increased dynamics of mobile wireless mesh network topologies [1], [2]. Therefore, some recent MANET multicast designs have been based upon redundant, mesh-based flooding mechanisms. Such mechanisms can reduce loss and delay and minimize state maintenance in highly dynamic cases, but come at a price of increased network-wide congestion for groups with scoped memberhip. One example is Simplified Multicast Forwarding (SMF) as desribed in IETF RFC 6621 specification [3]. SMF is characterized by flooding based upon duplicate packet detection that provides multicast forwarding resilient to dynamic topology changes in wireless communication environments. SMF can support multicast applications, management, and control and often results in lower message delivery delay and loss due to its resilient and simplified forwarding approach. By supporting options for connected dominating set (CDS) based relaying in addition to classic flooding (CF), the basic SMF specification provides a tradeoff space ranging from high redundancy, high overhead flooding to efficient flooding with reduced forwarding overhead. SMF has been previously analyzed in simulation and emulation research and has seen use in past and recent mobile wireless demonstration systems as well [4], [5]. With group-specific extensions to SMF, significant overhead savings may be achievable when and where dense mesh network topologies stabilize for periods of time or where group data exchange may be scoped within localized regions. This work discusses enhancements and experiments with such a protocol extension called Elastic Multicast (EM) routing.

II. Past Related Work

The initial design concepts for EM routing and initial experimental results were presented in [6] and later outlined in an Internet Draft specification [7]. From [6], the basic goal of EM is to provide an IP multicast service model that dynamically uses redundant forwarding in portions of the network affected by higher rates of topology change, while constructing and maintaining more specific group forwarding for higher rate flows where and when the network is relatively stable. Such protocol extensions can reduce the forwarding overhead of higher data rate traffic flows (e.g., collaborative sensing, tracking) while maintaining some of the resilient adaptivity of SMF-based redundant forwarding under dynamic conditions. In [6] a number of initial basic emulation results were presented using an earlier prototype of EM. In some basic scenarios, it was shown that overhead could be reduced versus using traditional SMF capability as the only multicast forwarding engine. Basic tradeoffs between data loss and overhead efficiency were demonstrated in [6], although more extensive experimentation with different multicast group types is of interest due to the complex tradeoffs that occur between overhead, redundancy, congestion, and contention. In addition, past experimentation work with SMF and CDS forwarding under a variety of mobility and congestion conditions was shown in [5]. This work showed that while classical flooding decreases overall loss statistics in mobile networks with light loads, more optimized forwarding increased throughput benefits as the network became more congested. So while redundant network transmissions increase the resiliency of delivery under dynamics for light traffic loading, the resultant loss due to both self-induced congestion and potential wireless channel contention is interdependent and nontrivial.

III. New Contributions

EM has three main benefits in mobile wireless network deployments. First, it attempts to build and maintain reduced forwarding meshes for high data rates flows resulting in less contention and congestion issues that can lead to higher overall loss and latency. Second, it provides group-specific traffic containment for group exchanges within localized clusters or network regions in contrast to using flooding variants which continuously impact the overhead in wider parts of the network. Third, by providing group-specific extensions to SMF, EM functionality aids in designing more effective and dynamic multicast gateway approaches and management between different autonomous network systems supporting multicast traffic exchanges.

In this work, we will present several areas of new contribution beyond past the EM work in [6]:

1) Addition of dynamic group-specific join/leave support
2) Multiple interface operation
3) Enhancements for dynamic, flow-specific mesh construction and maintenance (e.g., token bucket shaping)
4) Preemptive acknowledgement mechanism to improve late joining performance and reduce loss
5) Experimentation results from a complex scenario involving multiple groups and motion types
6) Demonstration of hybrid EM use, using mixed low and high rate multicast traffic flows

IV. Design

The *nrlsmf* routing prototype implementation was extended to include EM functionality making it straightforward for the EM prototype to take advantage of SMF design for lower rate traffic providing duplicate detection and self-organizing CDS flooding optimizations useful in MANET type operations. The *nrlsmf* implementation also benefits from several new enhancements added to make a more complete multicast routing system design as we will discuss. This new implementation is used within the emulation experiments described in this paper.

A. Overall Design

The goal the EM design is to provide an adaptive, group-specific refinement to SMF. EM leverages the duplicate packet detection and optional CDS flooding optimizations of SMF and adds simple control signaling to opportunistically reduce the SMF flooding mesh for group-specific traffic flows. The following terminology is used in discussing the EM design:

- *Group*: a routeable IP multicast address
- *Node*: an active network node running SMF or EM
- *Flow*: network traffic flow identified by the tuple
  - {source:group:(class):(protocol)}
- *Subscriber*: a node that subscribes to (or joins) a group
- *Active forwarder*: a participating multicast forwarding node
- *EM-ACK*: a control packet sent to upstream forwarders, expressing interest in a flow
- *EM-ADV*: optional control packet flooded to downstream forwarders, providing a list of active flow descriptions

In EM, active SMF relays limit flooding of multicast flows by default with a low-rate token bucket discipline. Thus, limited traffic gets broadcasted to the entire MANET network area as in SMF (optional CDS flooding optimization

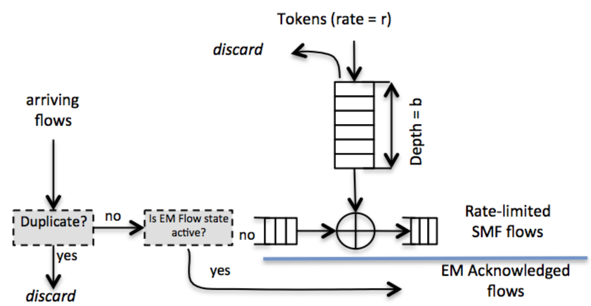

Fig. 1: SMF/EM Flow Rate Limiting

is supported) while higher rate traffic flows are throttled by the token bucket. This constraint is relaxed upon receipt of an acknowledgment message (EM-ACK) from a neighboring downstream router. Figure 1 illustrates this concept with the default token bucket limiting for unacknowledged flows and unlimited forwarding of non-duplicative packets for flows that have been acknowledged. An EM-ACK is sent to an upstream relay upon receipt of non-duplicative multicast packets for a flow when either the receiving router is a member of the group or it is actively receiving EM-ACK message for the given flow from other downstream routers. Since non-duplicative packets are required, only those upstream router(s) delivering the earliest arriving copies of IP multicast packets for a flow are acknowledged. Limited forwarding is reinstated for a flow after a threshold number of packets are received or a designated timeout occurs without any acknowledgement from downstream routers. This simple protocol serves to activate only the SMF relays needed to serve members of the multicast group for a given source-to-group flow specification. Under ideal circumstances (e.g., no wireless packet loss), a minimal set of relays remain activated. Under very high dynamics or degraded packet loss conditions, a downstream router may receive non-duplicative packets from alternating multiple upstream relays, and some redundant forwarders will be at least temporarily activated. However, that also allows for more resilient delivery of the forwarded flow under such conditions.

The low-rate, token bucket limited forwarding of flows allows potential receivers to discover flows with modest overhead and provide immediate delivery of low rate user data. Different forwarders (as a subset of overall SMF relays) will be activated by EM-ACK messages as flows for joined groups are dynamically sensed. Conversely, the count/timeout mechanism deactivates such EM relays when their forwarding service is no longer required by downstream routers due to topology changes or group membership departures. These routing changes are somewhat localized and are able to respond to dynamics quickly since no significant network state needs to be exchanged upon the routers. The *nrlsmf* EM implementation uses a threshold count of received packets without an EM-ACK reception event as the dominant criteria for relay deactivation with a secondary timeout that helps avoid unnecessarily prolonged forwarding of moderate flows or flows with irregular traffic patterns. These packet count triggers and timeouts may be configurable on a per-flow basis (e.g., by traffic class) in a deployed implementation.

The EM-ADV message mentioned above is a proposed message containing a list of active flow descriptions that upstream routers could send as a surrogate for the default low-rate forwarding of user data packets. A single EM-ADV packet could represent many flows. This would allow the protocol to scale for deployments where large numbers of different flows may be present. It is envisioned that hybrid operation could be supported where low rate flooding of user packets is the default behavior for some flow types while the surrogate EM-ADV messages are sent for other flow types. Also, *a priori* or critical flows with broad membership could always be flooded using SMF. The EM framework allows this flexibility.

The steady-state behavior for an active high rate flow is dependent upon the regular generation of EM-ACK messages and the detection of this signaling by upstream relays to keep unlimited-rate forwarding active for a flow (i.e., relaxing the rate limiting of a flow). The *nrlsmf* implementation also governs the maximum rate at which EM-ACK messages are generated for a given flow. So, for high rate flows, the quantity of control messages will remain at some nominal rate. The timeouts and count thresholds are structured to provide a "make-before-break" behavior as mobility results in topology changes. A more detailed protocol specification is planned to more precisely describe the relationship of the EM timers and counters. The token bucket parameters and these timeouts and counts could also be exposed as configurable items for different traffic classes and/or other flow descriptors to meet application utility requirements and network management needs.

B. Group Joins/Leaves and Related Design Issues

Flooding of IP datagrams, as in the case of basic SMF operation, is easier to implement than dynamic, group-specific IP multicasting. There is no need for the receiving kernel or the forwarding nodes to keep track of specific group management joins/leaves for multicast addresses. To support dynamic, group-specific routing it is necessary to be able to learn what group addresses are of interest. The IETF has developed protocols to locally signal group membership changes. The Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) protocols are specified in [8] and are essentially similar protocols, with IGMP being used for IPv4 and MLD used for IPv6. These protocols support both any-source multicast (ASM) and single-source multicast (SSM) semantics for group membership. The EM design extensions and the *nrlsmf* implementation are compatible with both ASM and SSM group membership operations. The current implementation monitors interfaces configured for EM operation for IGMP or MLD messages to dynamically learn group memberships for the local node. A future update to *nrlsmf* will support transmission of IGMP and MLD query messages to learn membership information for hosts attached to other interfaces (e.g., Ethernet).

C. Preemptive Acknowledgement Features

In the earlier designs of EM, an EM-ACK is generated only in response to a received packet for a newly detected flow of interest (e.g., triggered by local node membership). A feature of the updated *nrlsmf*-based implementation is to also actively keep state (within configured memory limits) for any detected flows including the recent, per-flow upstream forwarders so that an EM-ACK may be preemptively (i.e. immediately) sent upstream upon reception of an EM-ACK from a downstream router or upon a local group membership addition. This technique reduces the time needed to activate unlimited, high-rate forwarding of a flow since the token bucket limit prior to activation imposes delay of discovery of upstream relays for a specific given flow, which may decrease loss under dynamic conditions (e.g. mobility).

D. Elastic Multicast Multiple Interface Capability

Another new capability within the present *nrlsmf* EM implementation is support for multiple interface forwarding. Operating the EM protocol over a flooding domain that includes multiple interfaces per-node (i.e., multiple, possibly heterogeneous wireless subnetwork connections) helps mitigate an issue with SMF and the relay set selection algorithms currently used. With existing SMF algorithms, high rate traffic can be unnecessarily flooded over multiple wireless interfaces in a redundant fashion. EM will only activate high-rate forwarding on interfaces upon which EM-ACK messages are received. This can greatly reduce unnecessary, redundant forwarding by routers with multiple wireless interfaces operating in a common routing architecture. The *nrlsmf* implementation enforces this by keeping token bucket and EM forwarding state on a per-interface basis for each interface in an SMF interface group (set of interfaces in a common flooding domain).

E. Multicast Gateway Support

The EM-ADV and EM-ACK signaling messages can be used to support multicast routing gateways to other domains. A gateway node could advertise its presence by injecting EM-ADV messages with a "wildcard" flow description, identifying itself as a potential source for multicast flows. This gateway advertisement would be periodically flooded within the EM SMF domain and nodes would respond using the EM-ACK message to indicate their group membership interests. Since the EM-ADV message would be disseminated and processed on a hop-by-hop basis within the EM/SMF flooding domain, hop count and/or other path metric information could be included and accumulated to allow nodes to possibly selectively acknowledge the lowest cost gateway when multiple gateways are present. This process would allow gateways to learn the group memberships within the multi-hop MANET multicast routing domain. Similarly, a gateway in conjunction with other exterior multicast routing protocol operation, would use EM-ACK messages to receive and relay multicast packets for flows sourced within the local EM domain. This concept is being further developed and planned for implementation and experimentation within *nrlsmf*.

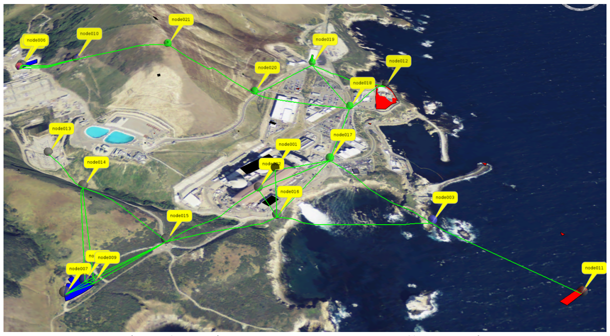

Fig. 2: Disaster Response Scenario

V. Experimental Results and Discussion

A. Fictional Emergency Response Scenario

To provide a rich ad hoc wireless network scenario to exercise the use of EM beyond the contrived topology examples in [6], we chose to use a fictional emergency response event at a nuclear power plant. The idea behind the vignette is loosely based upon an occurrence like the 2011 Fukushima Daiichi disaster and involves the use mobile wireless network technology to provide an infrastructure-free communication capability to aid rescuers and emergency personnel. The scenario involves different mobile units and devices including: medical response teams, repair crews, equipment trucks, airborne units, and wireless-capable sensor devices. While fictitious, the emergency response scenario requires command units, devices, and mobile units to use an ad hoc wireless capability to orchestrate emergency response and coordination. Our basic 21-node scenario is emulated using the Common Open Research Emulator (CORE), the Naval Research Laboratory (NRL) Network Mobility Framework (NMF), and the Extendable Mobile Ad hoc Network Emulator (EMANE) components [9]–[11]. The scenario involves both preplanned and causal motion elements. Non-causal motion involves security or patrol vehicles which are activated to orbit a critical inner area perimeter or can be statically placed in distributed locations. The scenario can also be run with causal-induced motion that is event-driven by communication messages received throughout the scenario. An example is a tasking order causing repair teams or medical rescue teams to move to particular destination locations to perform further tasking. A snapshot of the scenario topology and node names in shown in Figure 2. Within the scenario, we can vary the types and degrees of motion to test different aspects of EM performance.

B. Group-specific multicast traffic

In order to test the combination of SMF and EM operations with more heterogeneous traffic distributions, a variety of traffic flows are established within our experimental scenario using the MGEN [12] traffic tool to provide different combinations of sources and receivers of multicast flows. MGEN provides the necessary dynamic traffic scripting and group join capabilities needed to instantiate group-specific network exchanges. The following is a brief description of the multicast flows and receiver groups that were established within the experimental scenario used. The group communication flows generated range from small, localized clustered memberships to one flow which involves all receivers at the other extreme. In the present testing, all sources generated approximately 8 kbps traffic flows (10 packets per second of size 100 bytes).

- *Group 1*: Sensor: n1 to all 20 nodes (n2-21)
- *Group 2*: Repair: n18 to 6 site repair nodes (n4-9)
- *Group 3*: Rescue: n19 to 2 rescue vehicles (n11-12)
- *Group 4*: Surface patrol (n20) to helo (n11) coordination
- *Group 5*: Command: n21 to 3 patrolling nodes (n18-20)

C. Elastic Multicast Data Collection and Metrics

Several comparative metrics were examined within our EM experiments. Network overhead induced by EM or SMF operation is an important metric and is measured via distributed raw network traffic capture traces from each node in the experiment. Induced network overhead of a particular routing experiment is measured as the transmission events resulting from forwarding sourced multicast flows of interest and also includes the overhead of additional control messages required such as the EM-ACK process. From the distributed raw traffic logs we can determine the amount of forwarded packets within the network. Average goodput and loss statistics are calculated directly using the end-to-end traffic tool MGEN. Average goodput is calculated simply as the amount of user data received within some time window. The average loss is measured using each receivers MGEN log and then aggregating for the group.

D. Motion

Since different types of topology motion will effect EM in different ways we break our experiments down in different scenario motion categories to examine the effects of different types of topology dynamics on temporal and overall performance. The motion categories for the emergency response scenario are the following:

- Static Operations (no motion)
- Looping unit motion only
- Causal motion of units only
- All: Causal + looping motion

E. Summary Results: Goodput, Loss, Overhead

Our main metrics of interest with EM operations is to measure sustained goodput to subscribed receivers while also measuring resulting overhead and loss across a set of flows with different group and distribution characteristics. A main design goal is to sustain efficiency and delivery assurance under different types of multicast distribution and dynamics.

Figures 3,4,5 plot goodput, loss, and overhead respectively for multiple experiments and include multiple motion and group distribution cases as separate data points within each experiment. The title of each graph represent the type of protocol experiment that was run as follows:

- EM(PA) - Elastic Multicast with preemptive ACKs
- EM(NP) - Elastic Multicast without preemptive ACKs
- SMF - SMF flooding only

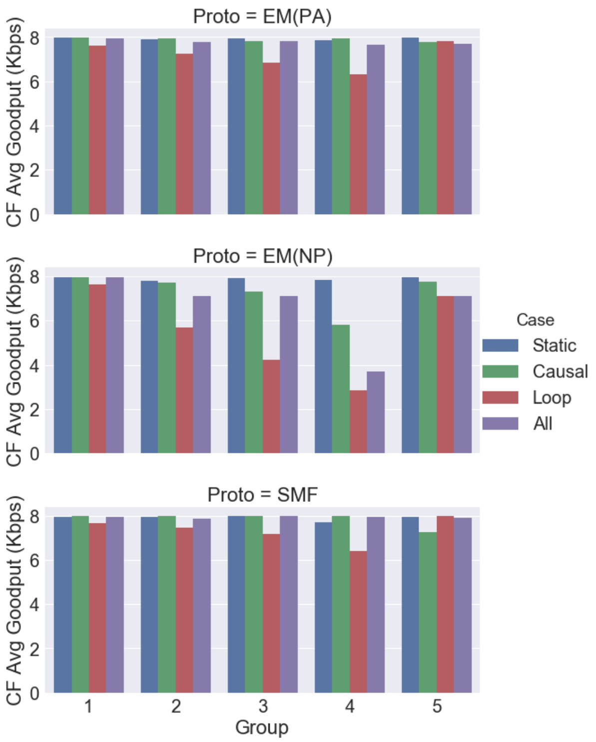

Fig. 3: EM-CF and SMF-CF Goodput

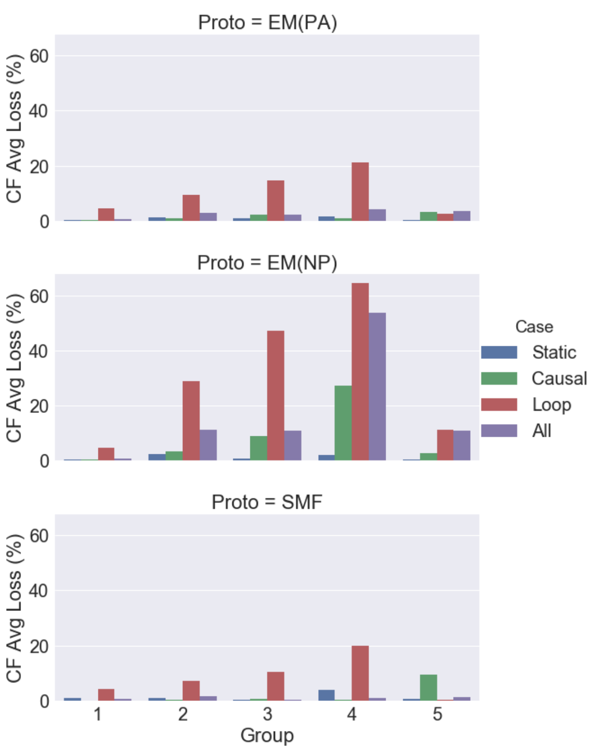

Fig. 4: EM-CF and SMF-CF Recv Loss

The five groups of bar graphs represent the different multicast distribution groups within the experiment and the four bar graphs within each group represent the four types of mobility as indicated in the legend.

From Fig 3 we can see that EM(PA) sustains reasonable goodput performance throughout all tests and is comparable

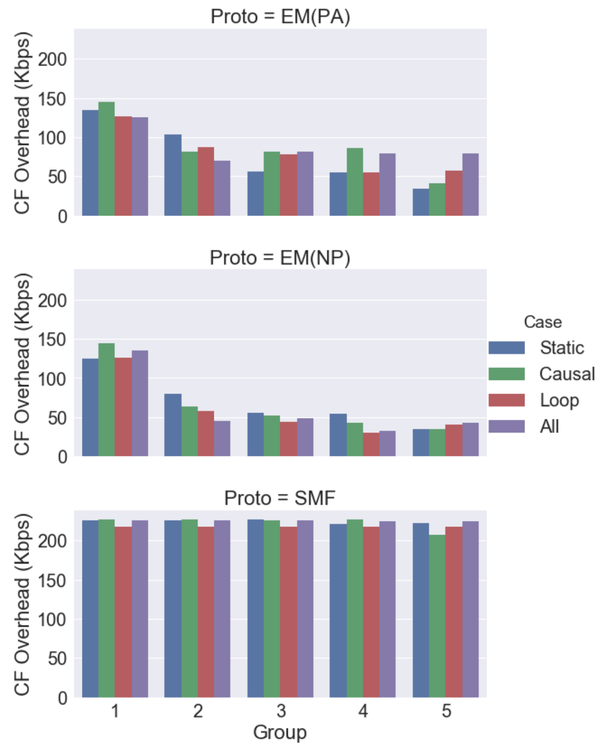

Fig. 5: EM-CF and SMF-CF Overhead

with SMF-only delivery. We contrast this to EM(NP) that shows some reduction in goodput under various motion trials, especially in Group 4 with the ”All” motion category. We can examine Fig 4 to see the corresponding average packet loss increase in the those cases. Since the total overhead metric Fig 5 to deliver the packets within Fig 3 includes both forwarded and ACK packets, we see this decrease for EM(NP) and most of this reduction can be explained by the increased packet loss leading to a lower number of forwarding events and therefore lower average overhead. From these figures, we observe that for these scenarios, the overall EM(PA) provides similar goodput and loss as SMF-CF for all multicast groups and motion types examined. EM(NP) can suffer from increased loss as seen in some of the results. EM has definitely reduced the overhead required to deliver data and this varies across traffic distribution types as is expected. For Group 1, the ”All” receiver case, we see a 25-30% reduction and for the sparser groups this is more significant up to a 75% or so reduction in overhead traffic. So these new results help confirm the fact that EM reduces overhead and can provide dynamic data delivery similar to the more redundant forwarding of SMF. Additionally, the benefit of the new preemptive ACKing mechanism is also seen in some of the performance data.

F. Elastic Multicast and Underlying SMF Relay Set Selection Algorithms

We ran a significant number of cases with variants of CDS algorithms for flooding and we present here a comparison of CF versus using Essential Connected Dominating Set (ECDS) [3] relays. Due to space limitations, we cannot present summary data for all of the SMF relay set selection algorithm

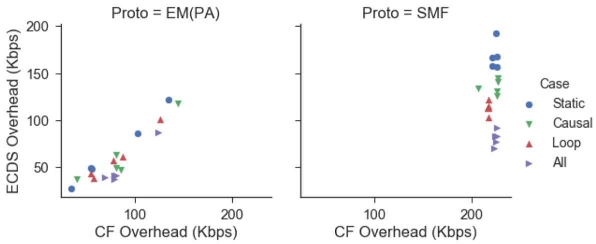

Fig. 6: CF vs. ECDS Forwarding Overhead

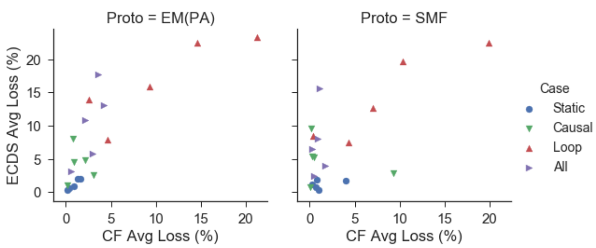

Fig. 7: CF vs. ECDS Recv Loss

variations tests that were run. Figures 6, 7 respectively present scattergrams of CF vs. ECDS overhead and loss statistics for each flow group across the 4 different motion cases and for EM on the left and plain SMF forwarding on the right. We see from Figure 6 that forwarding overhead is significantly reduced by using ECDS across all flows for the case of SMF-only delivery. But when the EM protocol extensions are activated, the resultant overhead becomes more comparable between cases. In Figure 7 we see for some cases that loss is comparable but that ECDS does worse in several cases. This provides some early indication that using CF flooding underneath EM may improve delivery resilience. We conjecture that this is due to the fact that EM has additional paths available for detecting and ACKing flows during mobility with CF but reduces overhead significantly for the more substantial higher rate flows. Addition studies may provide more insight on the tradeoffs between congestion, contention, and adaptivity.

VI. Further Planned Work and Issues

This paper described several advances in the design and implementation of the EM protocol. However, further work is planned to conduct more parameterized examination of concepts such as different relay set selection metrics and algorithms. Support for the EM-ADV control message described and inter-domain multicast gateways will be developed for further experimentation. Additionally, more experiments with heterogenous, multi-interface systems will be conducted.

VII. Summary

A full EM implementation supporting dynamic group join/leave, flexible flow classification, and multi-interface support was developed using the existing *nrlsmf* prototype implementation. Low-delay, preemptive acknowledgement generation was implemented to help reduce flow activation delay and loss under dynamic conditions. Experiments were conducted using existing wireless network emulation components with an emergency disaster response scenario that included different variations of complex motion and multicast traffic distribution.

Summary results showed that EM with preemptive ACKing provided goodput performance comparable to SMF CF mode while significantly reducing the overhead required within the network across different distribution groups. With increasing and differing mobility, there is less overhead savings as expected due to the inability to optimize some of the mesh forwarding, but across all cases examined, EM with preemptive ACKs still delivered traffic with reasonable success. We also showed that without preemptive ACKing, EM experiences significant additional loss. We observed that EM loss also increases with the use of underlying SMF-ECDS flooding versus SMF-CF flooding. These results are early but we have shown and measured the effectiveness of EM operations within a more complex mobile network scenario.

References

[1] JJ Garcia-Luna-Aceves and Ewerton L Madruga. "the core-assisted mesh protocol". *IEEE Journal on Selected Areas in Communications*, 17(8):1380–1394, 1999.

[2] Prasant Mohapatra, Chao Gui, and Jian Li. Group Communications in Mobile Ad hoc Networks. *IEEE Computer*, 37(2):52–59, 2004.

[3] J.P. Macker. Simplified Multicast Forwarding, (RFC 6621). *Internet Engineering Task Force (IETF) Experimental Internet Standard*, 2012.

[4] Claudiu Danilov, Thomas R Henderson, Thomas Goff, Jae H Kim, Joseph Macker, Jeff Weston, Natasha Neogi, Andres Ortiz, and Daniel Uhlig. Experiment and Field Demonstration of an 802.11-based Ground-UAV Mobile Ad hoc Network. In *IEEE Military Communications (MILCOM) Conference, 2009.*, pages 1–7. IEEE, 2009.

[5] Joseph Macker, Ian Downard, Justin Dean, and Brian Adamson. Evaluation of Distributed Cover Set Algorithms in Mobile Ad hoc Network for Simplified Multicast Forwarding. *ACM SIGMOBILE Mobile Computing and Communications Review*, 11(3):1–11, 2007.

[6] C. Danilov, T. R. Henderson, O. Brewer, J. H. Kim, J. Macker, and B. Adamson. Elastic Multicast for Tactical Communications. In *IEEE Military Communications (MILCOM) Conference*, pages 1–6. IEEE, Oct 2012.

[7] B Adamson, C Danilov, and J Macker. Elastic Multicast Routing Protocol. *IETF, Internet Draft Specification,https://tools.ietf.org/html/draft-adamson-elasticmcast-00*, 2013.

[8] B Fenner, H He, B Haberman, and H Sandick. Internet Group Management Protocol (IGMP)/Multicast Listener Discovery (MLD)-based Multicast Forwarding, RFC 4605. *Internet Engineering Task Force (IETF) Proposed Internet Standard*, 2006.

[9] Jeff Ahrenholz, Claudiu Danilov, Thomas R. Henderson, and Jae H. Kim. CORE: A real-time network emulator. In *IEEE Military Communications Conference*, pages 1–7, 2008.

[10] Natalie Ivanic, Brian Rivera, and Brian Adamson. Mobile Ad hoc Network Emulation Environment. In *Military Communications Conference, 2009. MILCOM 2009. IEEE*, pages 1–6. IEEE, 2009.

[11] Jeff Ahrenholz, Tom Goff, and Brian Adamson. Integration of the CORE and EMANE Network Emulators. In *IEEE Military Communications (MILCOM) Conference*, pages 1870–1875. IEEE, 2011.

[12] Brian Adamson. The Multi-Generator (MGEN) User's and Reference Guide 5.0. *https://downloads.pf.itd.nrl.navy.mil/docs/mgen/mgen.html*, March 2017.