Elastic Multicast

Brian Adamson – NRL
Claudiu Danilov – Boeing
Joe Macker – NRL

draft-adamson-elasticmcast-00

7 November 2013
IETF 88 - Vancouver

Outline

- Background and Main Concept
- Protocol Description
- Some Experimental Results
- Implementation Status
- Protocol Design Considerations
- Document Status

Background - SMF

Simplified Multicast Forwarding (SMF)

- MANET Multicast (broadcast) protocol.
- Packets are disseminated to the entire MANET.
  - Duplicates are eliminated (no loops).
  - Optimized through use of Connected Dominating Set (CDS).
- No knowledge of group membership.
- Very efficient in small, highly dynamic mobile networks.
- Not intended to scale to large networks.

Main Concept

- Dynamic pruning of broadcasting nodes in SMF.
  - *Keep the portion of SMF mesh active only where needed*
- Converges to single path multicast trees in stable networks
- Expands to full network broadcast in highly dynamic networks
- Maintains no topology or global membership information.
- Applied to higher bandwidth traffic
- Lower bandwidth traffic still broadcasted to the entire MANET
  - *Proposed provision for surrogate message to "advertise" active flows instead of flooding user traffic for some flows if desired*

Definitions

- **Group:** An IP Multicast address.
- **Node:** A wireless network node running SMF.
- **Flow:** Traffic being characterized by a source address, destination group, and a set of optional parameters such as protocol, traffic class, port, etc.
- **Subscriber:** A node that subscribes to (joins) that group.
- **Active forwarder for a flow:** An SMF node that does not limit packet broadcasting for that flow.
- **EM-ACK message for a flow:** A unicast control packet being sent to an upstream (previous hop) forwarder, indicating interest in a flow.
- **EM-ADV message:** Optional message with a list of active flows that can be disseminated as surrogate for lower rate flooding

Elastic Multicast Protocol (1)

- Each node limits the forwarding rate through a separate token bucket created for each flow they observe.
  - DPD and SMF relay sets (CDS) are also applied.
  - Packets are forwarded via local broadcast.
  - Lower traffic flows get floded to the entire MANET, as with SMF.
  - Higher bandwidth flows are throttled, and get most of their packets dropped. All packets that exceed the token bucket rate for their flow are dropped. Some packets still make it to the entire MANET.
  - Initial empty token bucket can allow for “start up” of some multicast protocols
- Subscribers send periodic EM-ACK messages to an upstream forwarder only if/when they receive new packetsfor groups that are subscribed to from that node.
  - DPD determines new packets.
  - No link is established to the upstream nodes.
  - EM-ACK messages continue to be sent periodically to an upstream node, as long as new subscribed packets continue to be received from that node.

Elastic Multicast Protocol (2)

- Upon receiving an EM-ACK message for a flow, a node becomes an active forwarder for that flow, for a limited interval.
  - A new EM-ACK message received refreshes the interval for that flow.
- Active forwarders send periodic EM-ACK messages to upstream forwarders only if/when they receive new packets that match flows for which they are active forwarders.
  - EM-ACK messages continue to be sent periodically to an upstream forwarder, as long as new packets for their flows continue to be received from that node.

Dynamic Behavior

- **Static network**
  - Multiple trees (backbones) of high speed forwarders are established from sources to multicast destinations.
  - Trees are maintained through periodic Join messages.

- **Dynamic network**
  - More nodes become high speed forwarders due to topology changes
  - The set/number of active forwarders depends of the network dynamics in different parts of the network.
  - Potentially all SMF relays can become active forwarders.
  - When/where mobility slows down, some of the active forwarder timers expire, and they start throttling down the flows they forward.
  - Eventually, if the network becomes static, only a small number of active forwarders will form trees/backbones derived from the full SMF mesh

Elastic Multicast Protocol Messages

- EM-ACK
  - router-id : *identifier of previous hop forwarder*
  - group-addr: *group destination address of flow*
  - source-addr: *source address of flow*
  - protocol, traffic class, ports (*optional descriptors*)

- EM-ADV – *optional list of active flows, each with:*
  - group-addr: *group destination address of flow*
  - source-addr: *source address of flow*
  - protocol, traffic class, ports (*optional descriptors*)
  - dpd-id: *DPD identifier of packet triggering flow advertisement*
  - tagger-id: *originator of flow advertisement*

Current Implementation Approach

Architecture

- Two separate processes
- Packets intercepted and forwarded by SMF process.
- Forwarding decisions done in the Routing Daemon.
- The Routing Daemon processes handles subscriptions and sends EM-ACK messages

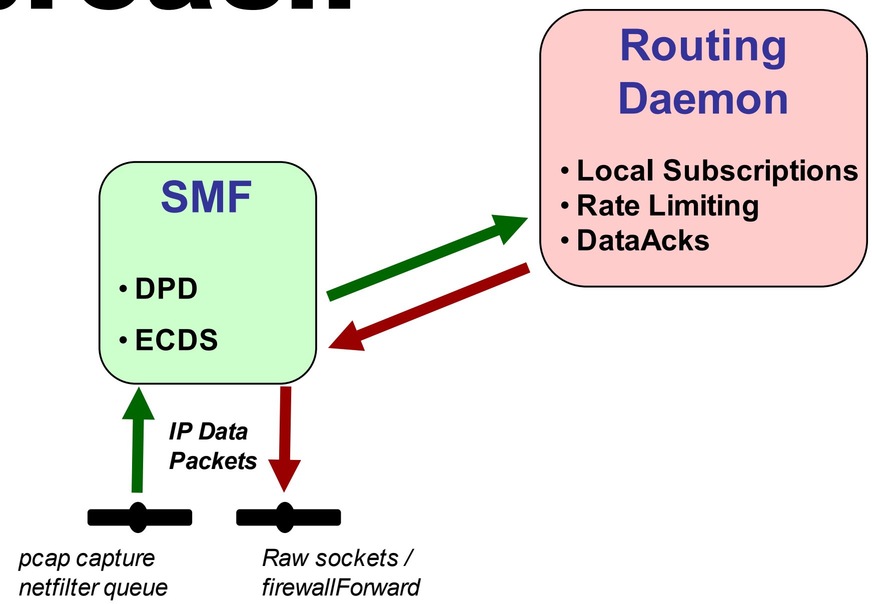

Static Grid Topology

- 20 nodes, static topology
- Single sender at node S
- SMF (Classic Flooding) compared with EM.
- Measured traffic at each node.

- **Single receiver D1 (EM1)**
  - Forwarding nodes: 7, 13, 19
- **Two receivers D1 and D2 (EM2)**
  - Forwarding nodes: 7, 12, 13, 19

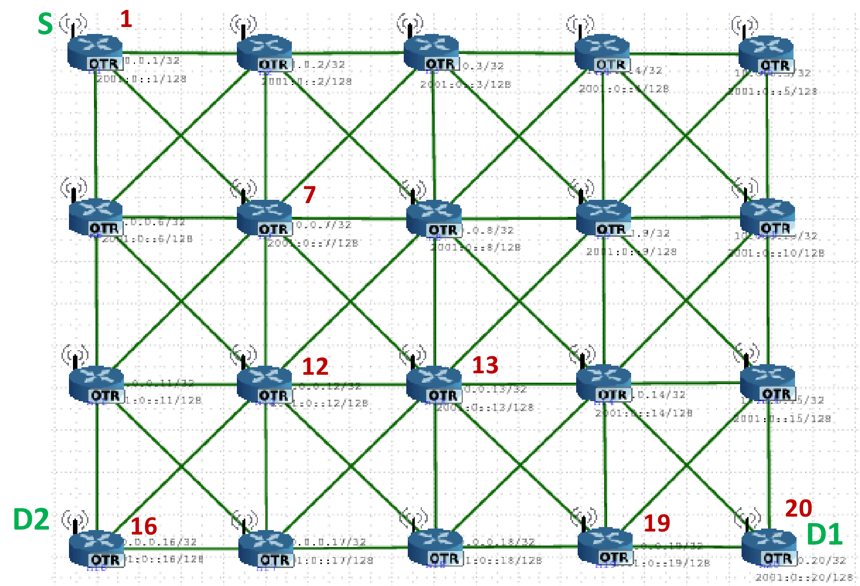

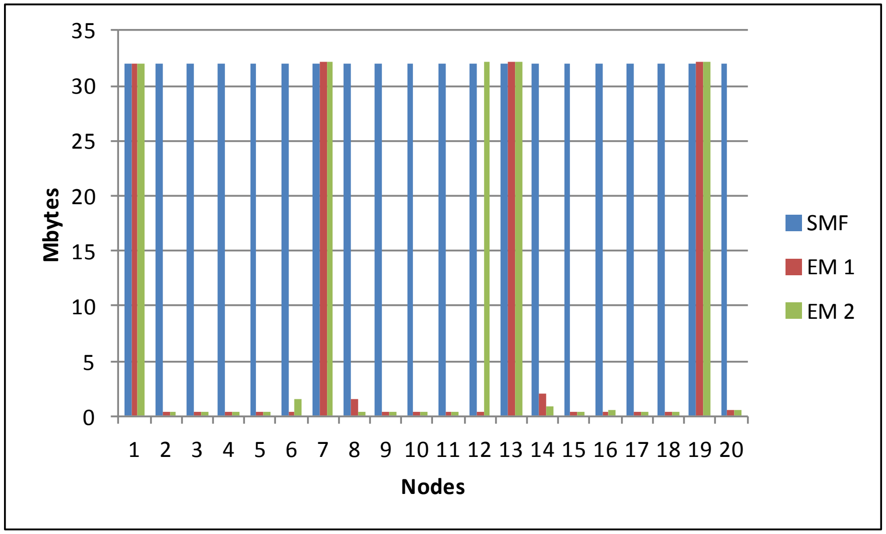

Gauntlet Scenario

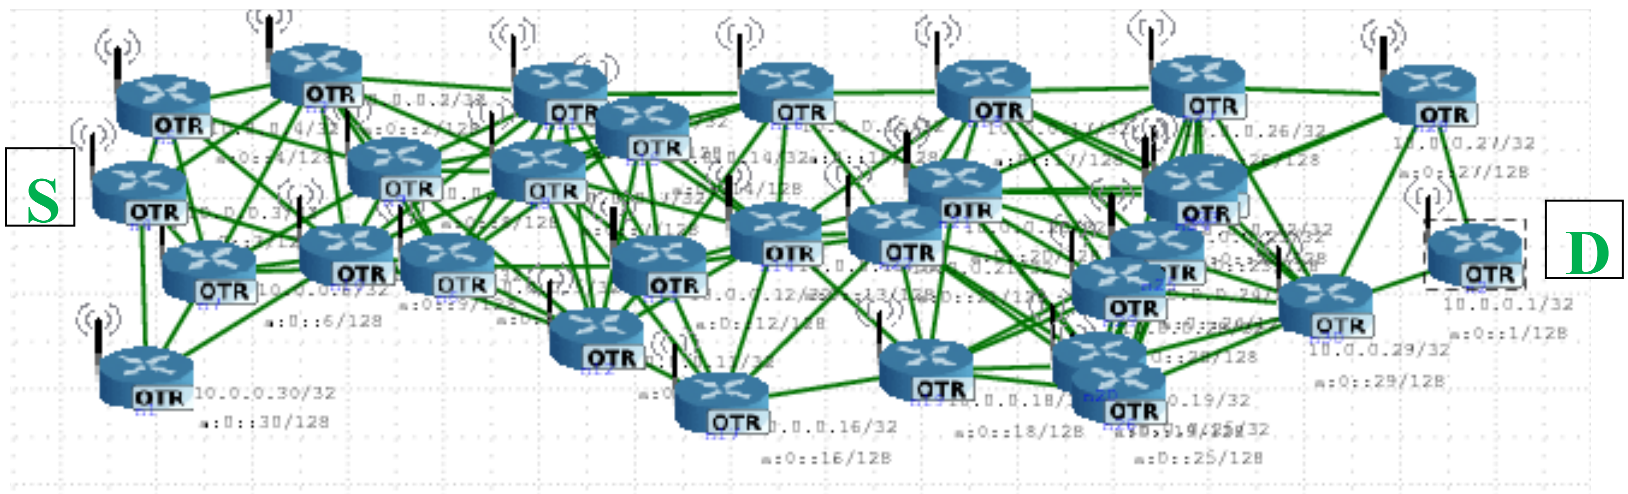

- 30 nodes, random mobility.
- 200 x 800 grid; wireless range 200
- Two end nodes are static (S, D).
- Nodes in the middle of the network move faster.
- Node mobility slows down gradually towards the ends
- Multicast UDP streams sent from S to D for 300 seconds.

Gauntlet Topology

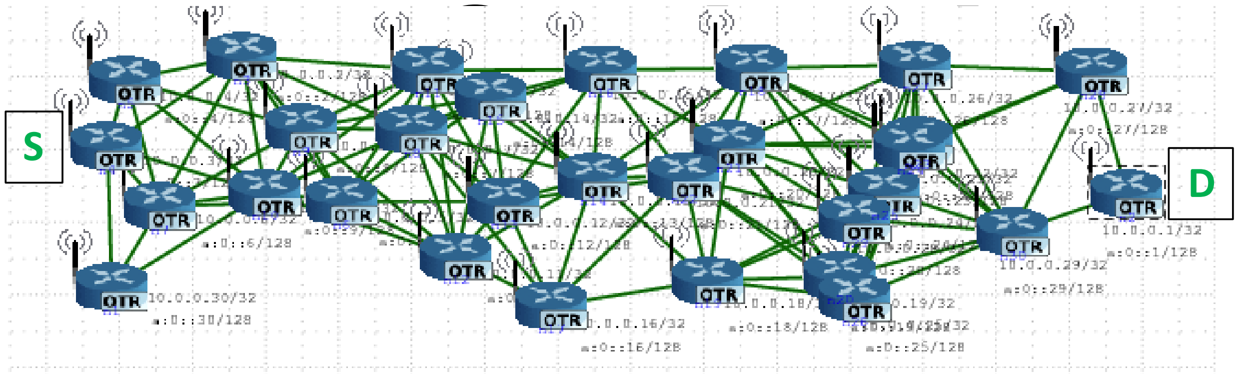

Static

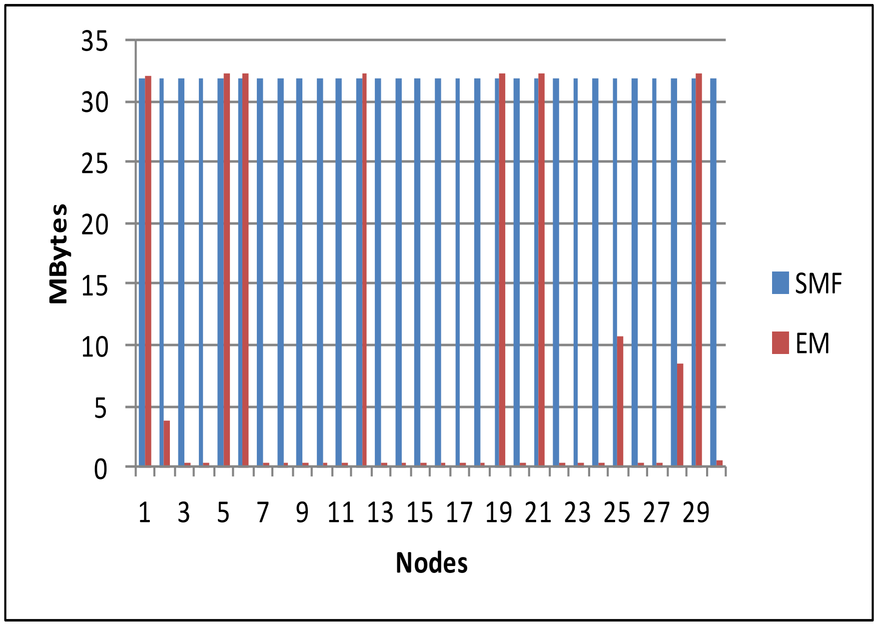

Mobile

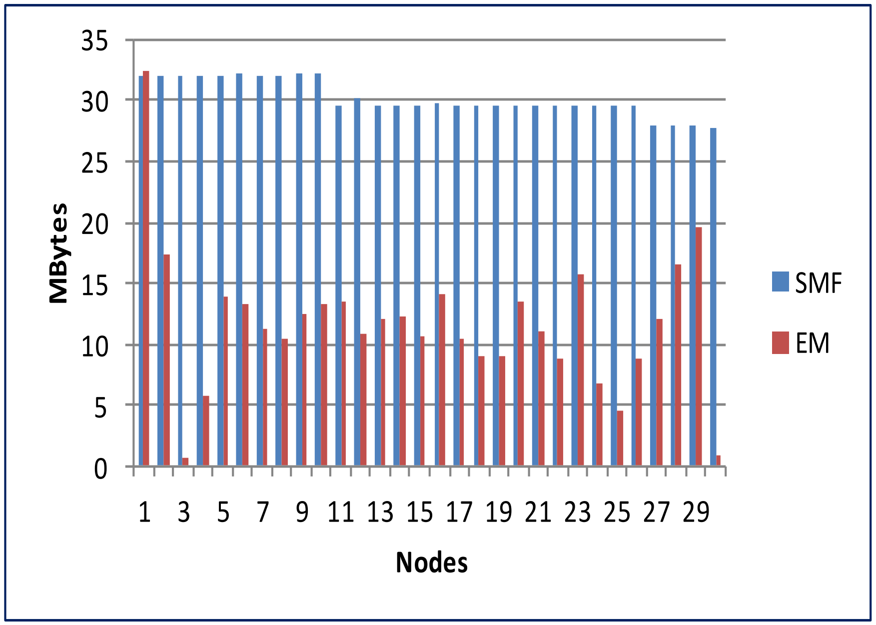

Random Waypoint – Delivery Ratio

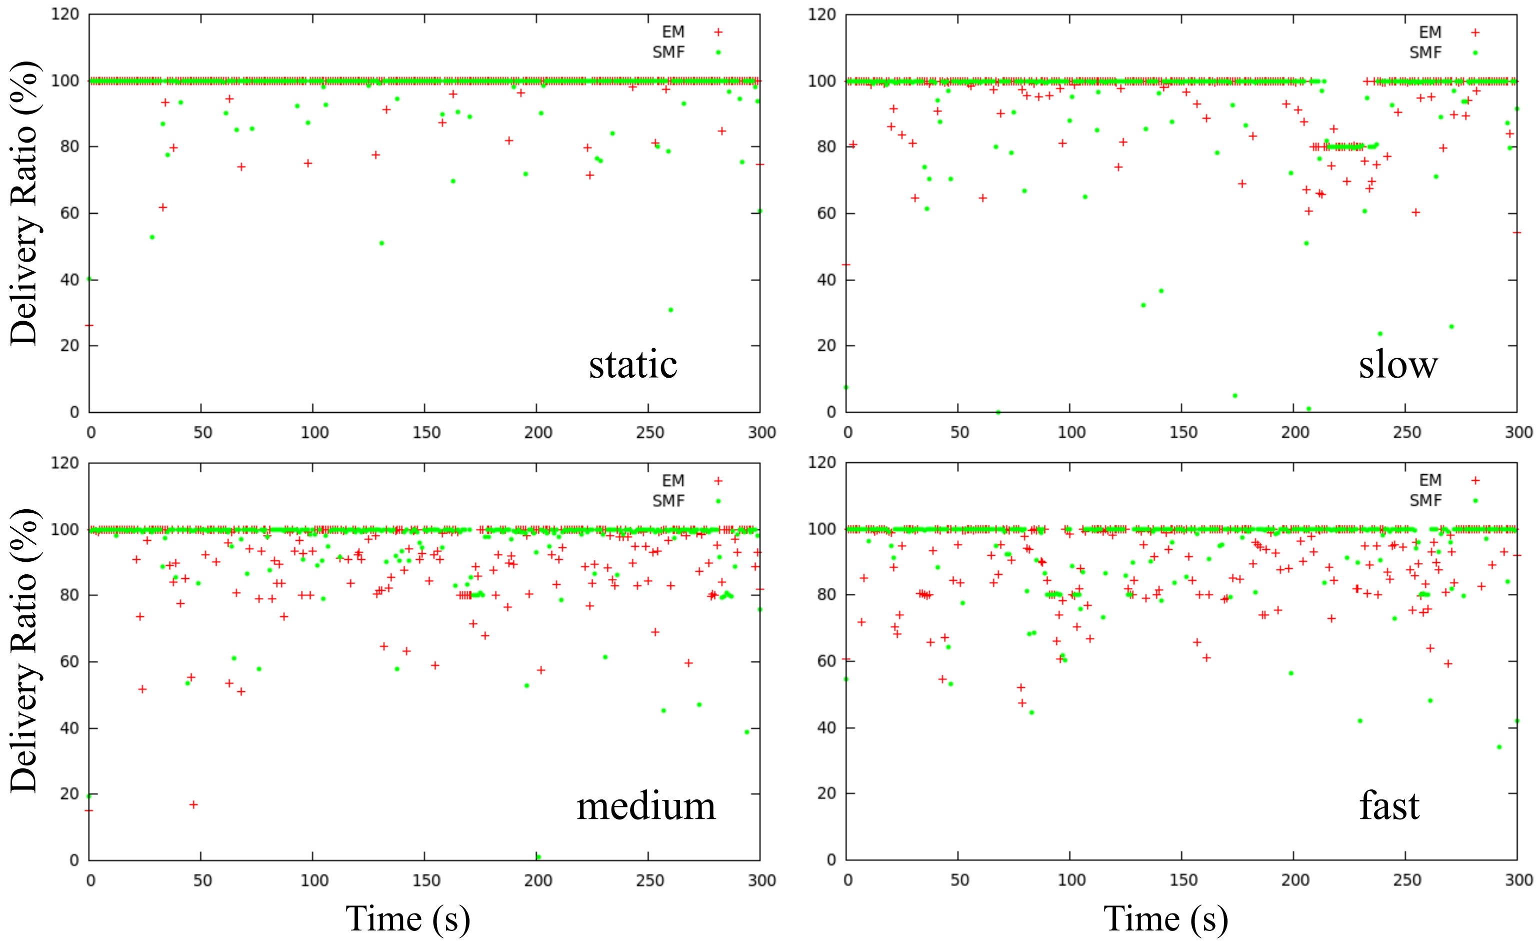

40 mobile nodes; 5 subscribers; one sender; four different scenarios

Random Waypoint – Overhead

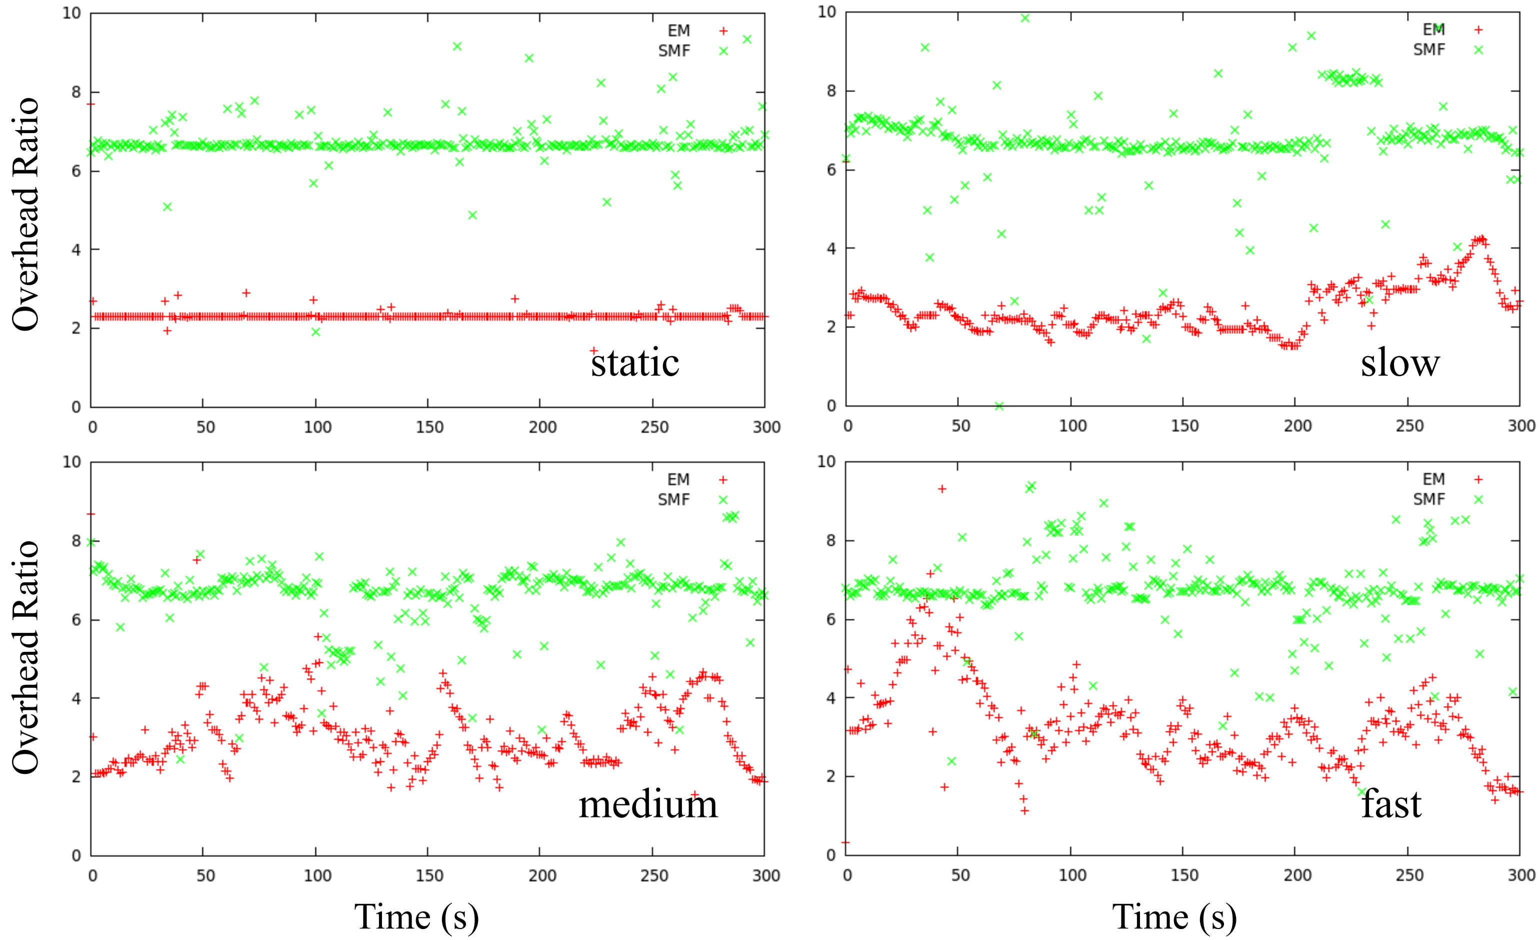

40 mobile nodes; 5 subscribers; one sender; four different scenarios

Implementation Status

- EM-ACK implemented as UDP message format
  - Will be updated to PacketBB format
  - EM-ADV, dynamic group membership not yet implemented
- *nrlsmf* and *arouted* processes both in packet forwarding path
  - Plan to define forwarding information base (FIB) structure for SMF that supports this and potentially other protocol needs for better control and forwarding plane separation
  - Will refactor *nrlsmf*, *arouted*, and *nrlnhdp* implementations to support updated design

Protocol Design Considerations

- IGMP and MLD
  - ASM and SSM support possible
  - "Local host" association paradigms
    - Collocated processes joining groups
      - OS-specific IGMP/MLD join/leave interception
    - Directly attached hosts on non-MANET interfaces
      - Standard IGMP/MLD querier operation sufficient
    - Neighboring non-router MANET hosts
      - Modified IGMP/MLD querier operation needed
- Border Gateways
  - EM-ADV mechanism useful for gateway purposes
  - Concept for "wildcard" EM-ADV message to collect routing area membership and/or pre-plumb set of active forwarders for joined groups
- Proactive Elastic Multicast
  - Nodes joining groups as sources could generated EM-ADV messages
- SMF Relay Set Selection Algorithm considerations
  - ECDS and CF accommodated by EM-ACK generation as described
  - S-MPR and other relay set algorithms may need modified behavior
    - E.g., only send EM-ACK to S-MPR "selector" previous hop forwarders

Document Status

- `draft-adamson-elasticmcast-00`
- It's a rough draft.
- PacketBB (RFC 5444) message formats not yet defined
  - Elastic Multicast could "piggy-back" with other MANET protocol signaling (e.g. NHDP, etc)
- Protocol timeouts not explicitly described
- Details on various protocol design considerations need to be completed

Elastic Multicast – Summary

- Dynamic pruning of SMF relays for specific group memberships
- Converges to single path multicast trees in stable networks
- Expands to full network flood in highly dynamic networks
- Maintains no topology or global membership information.
- Relay set reduction applied to higher bandwidth traffic
- Lower bandwidth traffic still flooded to the entire MANET
  - Or surrogate flow advertisements