Reliable Adaptive Lightweight Multicast Protocol

Ken Tang, Scalable Network Technologies

Katia Obraczka, UC Santa Cruz

Sung-Ju Lee, Hewlett-Packard Laboratories

Mario Gerla, UCLA

Overview

- Ad hoc network introduction
- QualNet network simulator
- Reliable multicast in ad hoc networks
  - Scalable Reliable Multicast (SRM) case study
  - Reliable Adaptive Lightweight Multicast (RALM) protocol
- Conclusion

Reliable Multicast in Ad Hoc Networks

- Challenges in MANETs
  - Node mobility
  - Hidden terminals make MANET sensitive to network load and congestion
- Our goal: design a multicast transport protocol that achieves both *reliability* and *congestion control*

Case Study of the Scalable Reliable Multicast (SRM) Protocol

- Representative of “wired” reliable multicast protocols
  - Negative acknowledgements (NACKs)
  - Multicasting of NACKs
  - Nack’ed packets are retransmitted
  - NACK suppression
  - Local recovery

Scalable Reliable Multicast (SRM)

Representative of “wired” reliable multicast protocols

- Receivers use *repair request* messages to request retransmission of lost data
- *Repair requests* are generated until the lost data is recovered
- Any multicast group member that has the requested data may answer by sending a *repair* message.
- NACKs and data retransmissions are multicast to the entire group
- Suppresses *repair request* and *repair* messages

Snippet of SRM Performance

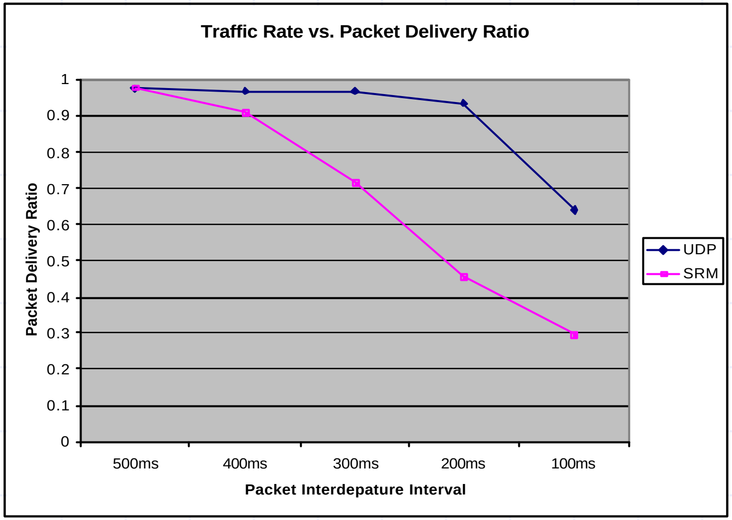

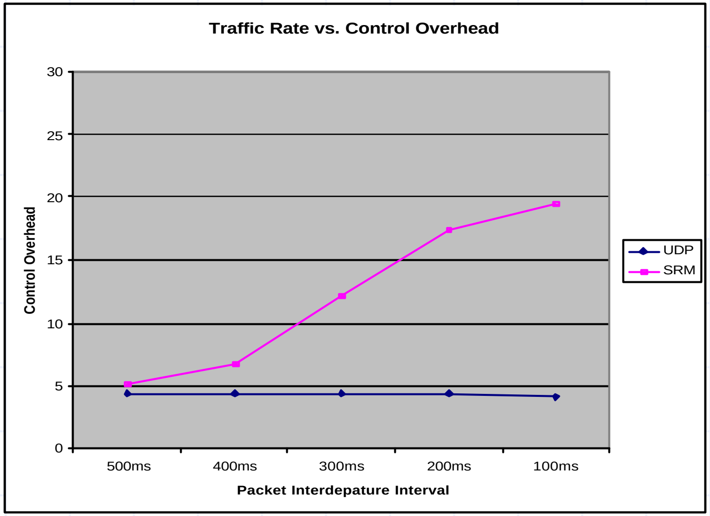

- 50 nodes in 1500m x 1500m area
- 5 sources and 10 receivers
- Traffic rate varies from 2 packets per second to 10 packets per second
- SRM degrades as traffic rate increases
  - Retransmissions increase packet loss (since sources maintain sending rate) which further triggers more retransmissions (as evident by control overhead graph) which leads to even more packet loss
  - Packet loss caused by increased load in the first place. Retransmission

Lessons Learned

- Confirmed that ad hoc networks are extremely sensitive to network load
- Reliability cannot be achieved by retransmission requests alone
  - SRM under-performed plain UDP
- Strong indication that some form of congestion control in conjunction with retransmissions is also needed to accompany reliability

Lessons Learned (cont'd)

- Losses may not be correlated: downstream nodes may still receive packets even if upstream nodes do not, especially considering mobility
- Packet loss may be due to wireless medium error rather than simply congestion

Reliable Adaptive Lightweight Multicast (RALM) Highlights

- Rate-based transmission
- Transmit at "**native rate**" of application as long as no congestion/loss is detected
- When congestion/loss (via NACKs) is detected, fall back to **send-and-wait**
- In send-wait mode control congestion and perform loss recovery
- Reliability achieved with congestion control AND retransmissions

RALM Finite State Diagram

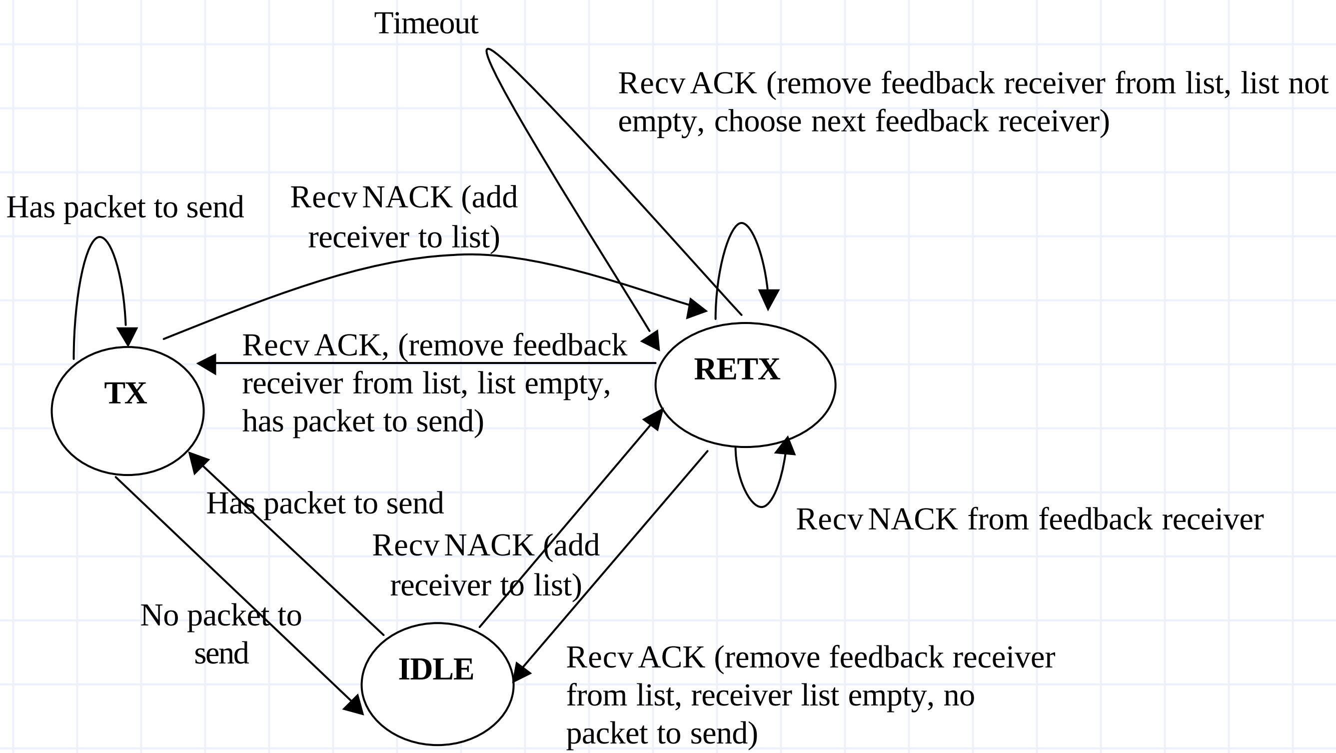

RALM Example

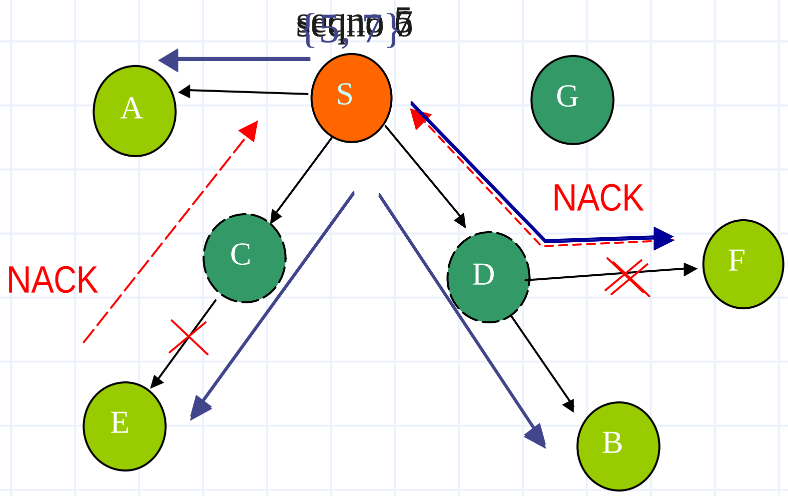

- Node E and node F detect loss
  - Node E detects loss of packet with seqno 5
  - Node F detects loss of packets with seqno 5 and 6
  - All receivers receive seqno 7
  - Both E and F unicast NACK to node 1
- Node E and node F are now recorded in Receiver List for round-robin send-and-wait loss recovery

RALM Example (cont'd)

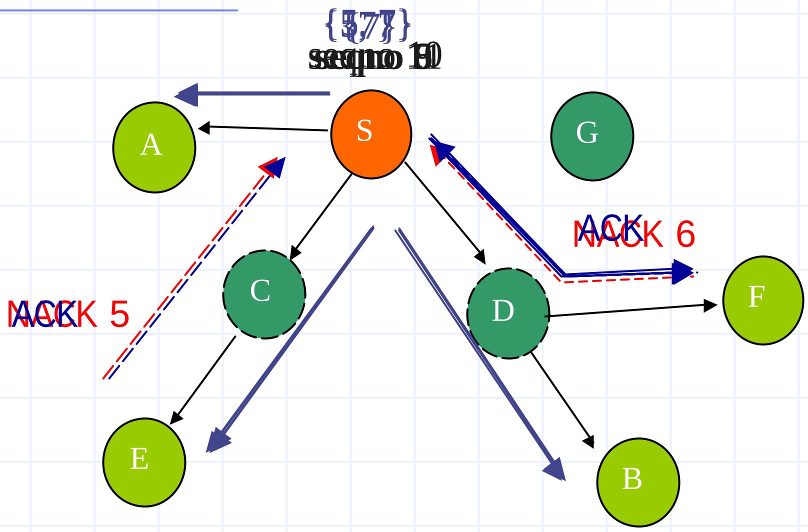

- Node S selects node E as the feedback receiver to reliably transmit seqno 8
  - Only node E may respond
- Node S then selects node F to reliably transmit seqno 9
  - Only node F may respond
- Since there are no more receivers in Receiver List, revert to multicasting at the application sending rate

Feedback Receiver

- Only a single (feedback) receiver acknowledges data
  - Feedback receiver list: list of nodes that have sent NACKs
  - The source specifies the feedback receiver in the multicast data
  - Feedback receiver is rotated in round robin order
  - Unicast ACK or NACK to the source
  - If feedback receiver fails to respond to source after specified number of times, receiver is skipped until the next round

Loss Recovery

- When the feedback receiver detects loss packets, it unicasts a NACK to the source for retransmission
  - Lost packets are requested *one at a time* until it has all the up-to-date packets
  - It slows down the source transmission when congestion is detected
- The source *multicasts* both new and retransmitted packets
  - Other nodes who may have lost those packets will receive the retransmission
- The feedback receiver unicasts ACK to the source once it receives all the packets
  - The source chooses a new feedback receiver from the Receiver List
  - Repeats this process until the list is empty

Simulation Environment

- QualNet for network simulation
- Compare UDP, SRM and RALM on top of ODMRP/AODV/IEEE802.11DCF in various scenarios
  - UDP: no congestion control or error control
  - SRM: error control without congestion control
- 50 nodes in 1500*m* by 1500*m* area
- Channel capacity: 2 Mb/s
- Propagation range: 375 meters
- Two-ray ground reflection model
  - Free space path loss for near sight
  - Plane earth path loss for far sight
- Random waypoint mobility model
- Constant bit rate "application-driven" traffic

Simulation Environment (Cont'd)

- Metrics
  - Packet delivery ratio: Effectiveness and reliability
  - Control overhead
    - The total number of data and control packets sent by routing and transport layer protocols : the number of data packets received by the receivers
    - Efficiency
  - End-to-end latency: Timeliness

Traffic Rate Experiment

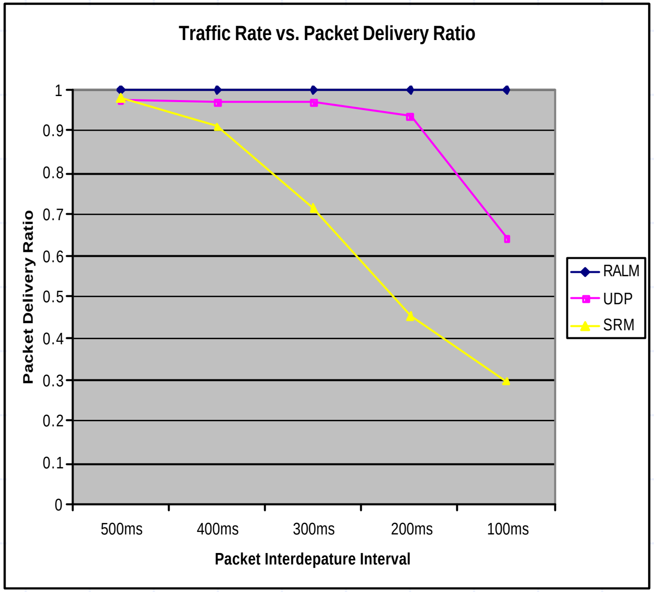

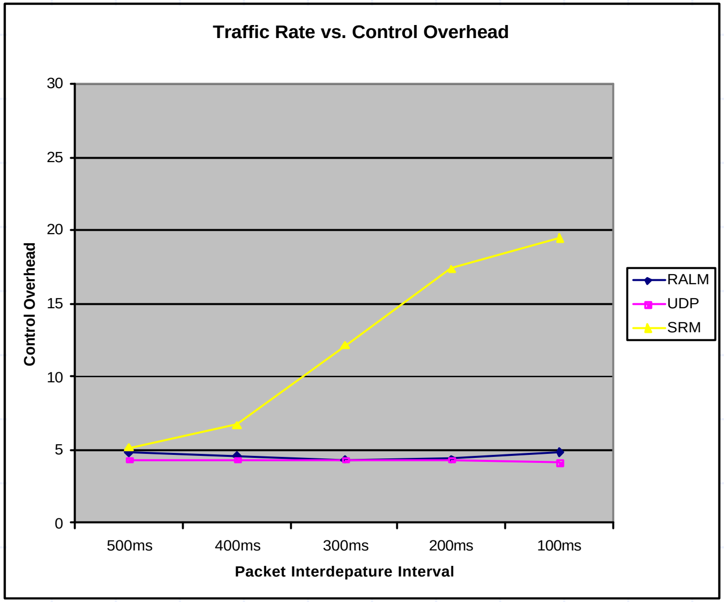

- No mobility
- 5 sources and 10 receivers
- Vary inter-departure rate from 500ms (2 packets per second) to 100ms (10 packets per second)
- RALM: 100% reliability, low control overhead and delay

Mobility Experiments

- 5 sources and 10 receivers
- 2 packets per second
- Random waypoint from 0 m/s to 50 m/s with pause time of 0 s
- UDP outperforms SRM
- 100% data delivery with RALM

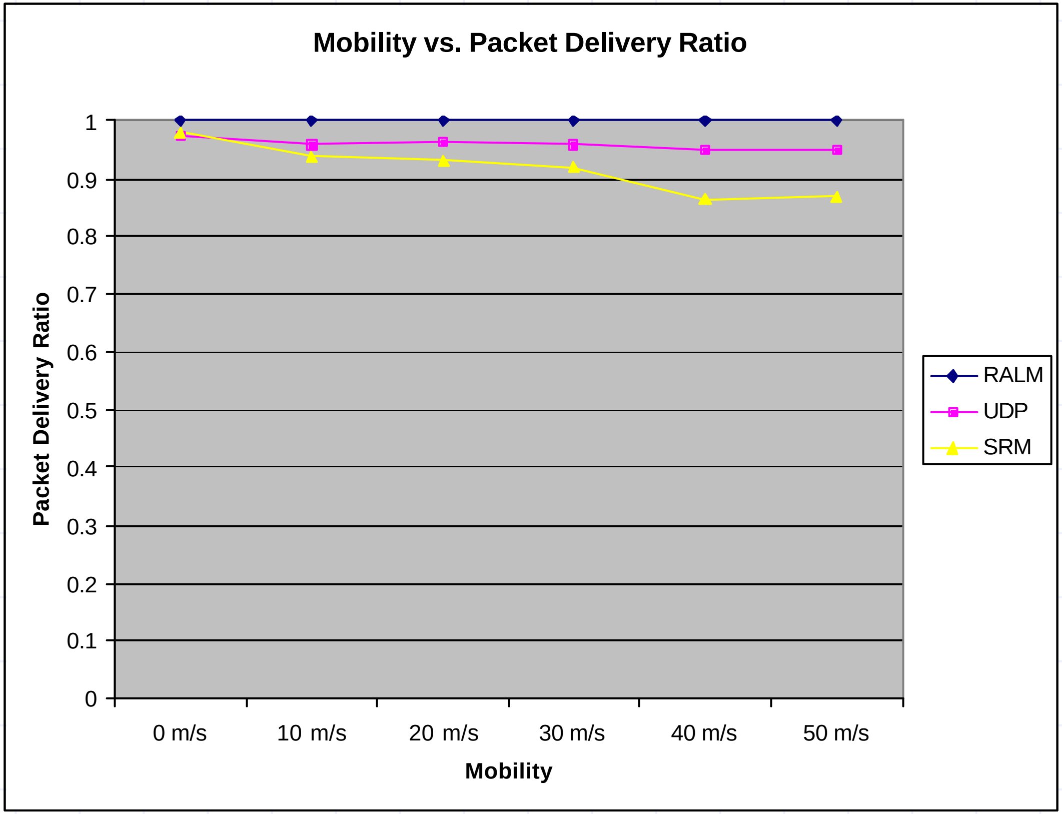

RALM vs. Multiple Unicast TCP Experiments

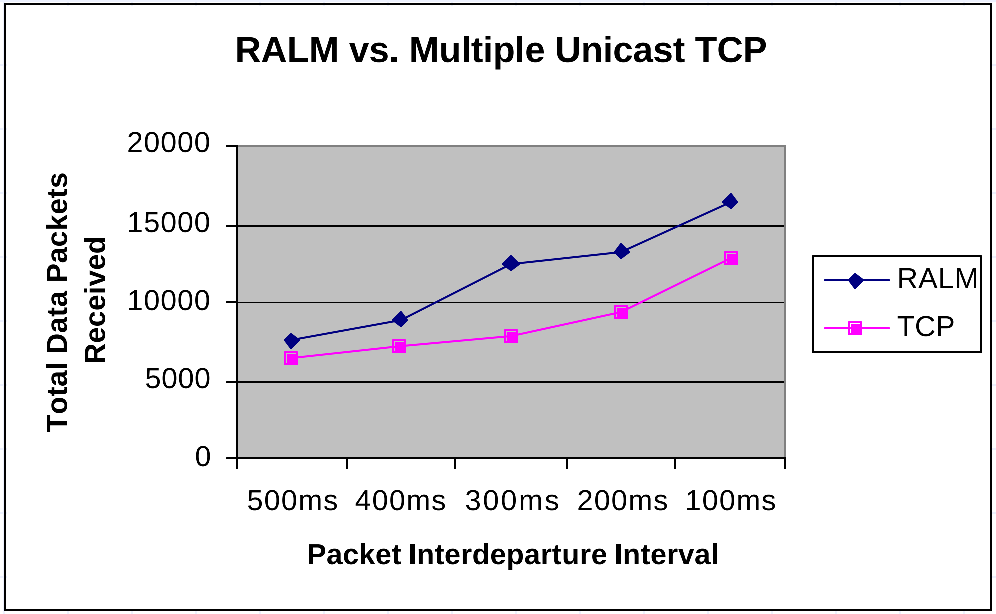

- Same as traffic rate experiment
- Compare RALM to multiple unicast TCP streams
- On average, 25% more packets delivered than TCP
- RALM performance differential grows with increase in receiver set

Conclusion

- Traditional wired network approach to reliable multicast does not work well in ad hoc networks
  - Mobility
  - Hidden-terminal problems
  - Contention-based MAC protocols
- Must take into account also congestion control, not simply error control (i.e., SRM)
- RALM utilizes congestion control in conjunction with reliable delivery to achieve reliability

Ongoing Work

- Discriminate loss from mobility and congestion
- Simulate on top of MAODV
- Compare performance against other ad hoc reliable transport multicast protocols (e.g., anonymous gossip)
- Look at congestion control and reliability at various layers