CSE 222
Graduate Networking

Winter 2001

Lecture 17: Resource Reservation

Geoffrey M. Voelker

Network Quality of Service

- Problem: How do you provide quality of service in the Internet?
  - Bandwidth, delay, “rate”
- Even harder: In multipoint-multipoint communications?
  - Heterogeneous network performance
- Datagram networks
  - Maximize network utilization, support multipoint, robust
  - But, only best-effort (no guarantees)
- Circuit switched networks
  - Guarantees
  - Inefficient for bursty traffic, not robust, no multipoint

The Best of Both

- How can we provide the best of both worlds?
  - Integrated Services Packet Network (ISPN)
  - Essentially, add service guarantees to datagram networks
- To do this, we need:
  - Flow specs defining resource requirements (network, receiver)
  - Routing protocol supporting (1) QoS and (2) multicast
  - Reservation protocol to create and maintain reservations
  - Admission control to prevent over-utilization
  - Packet transmission scheduling algorithm

Reservation Protocol

- The paper focuses on the reservation protocol for two systems
  - Internet Stream Protocol, ST-II (IPv5?)
  - Resource Reservation Protocol, RSVP
- Overview
  - Describe reservation protocols
  - Compare static performance
  - Compare dynamic performance
- In retrospect, should have read the RSVP paper

ST-II

- Service model
  - Simplex data stream multicast from source to receivers
    - » ST-II implements the multicast tree
  - Connect message floods net to establish paths to receivers
    - » Includes flow specification to do resource reservation
    - » Tree supports only one flow spec
    - » Independent Streams model
  - Establish multicast state, reserve network resources
  - Receivers reply with Accept or Refuse
    - » Can adjust flow specification in reply
  - Sender synchronously waits for replies from all receivers
    - » Can adapt to lower flow spec or reject receiver

ST-II (2)

- Group membership
  - New receivers communicate out-of-band with source
  - Receivers reply with Accept or Refuse
  - Sender can adapt or reject a receiver requiring a different flow
  - Receivers leave with a Refuse or Disconnect from source
- Reliability
  - Control messages reliably transmitted hop-by-hop
  - Hello messages used for heart-beat between neighbors

RSVP

- Service model
  - Also simplex distribution tree
    - » RSVP relies upon IP multicast to create and maintain tree
  - Source sends Path message to multicast group address
    - » Includes flow spec for receivers
    - » Notifies intermediate RSVP agents of new tree
  - Receivers listen to multicast group to receive Path messages
    - » Receiver-initiated
    - » Respond with Reservation requests back to source
    - » RSVP agents reserve resources as request goes back to source
    - » Terminates when it reaches a branch that can support the flow

RSVP (2)

- Group membership
  - Handled by underlying IP multicast
- Reliability
  - Soft state built from period messages
  - Path refreshes adapt to changes in multicast distribution tree
    - » From source, adapts to topology changes
  - Reservation refreshes maintain resource reservations
    - » From receivers, used for new reservations

RSVP Reservations

- RSVP reservations have two components
  - Resource allocation
    - » How much to reserve (e.g., 128Kb)
  - Packet filter
    - » What packets get to use the resources (e.g., which sources)
- Packet filters
  - Wildcard – packets from any source can use reservations
  - Fixed filter – packets from a single source only
    - » Changing source requires new setup and admission control
  - Dynamic filter – dynamically choose a set of sources
    - » Without requiring redoing setup and admission control

Analyses

- Static analysis – network resource requirements
  - Application requirements
  - Heterogeneous groups
  - Channel selection
- Dynamic analysis – protocol overhead for adaptation
  - Reliability
  - Group membership
- Or, what does RSVP have that ST-II does not?

Static Analyses

- Application requirements
  - Audio conferencing – at most a few speakers at a time
  - Only need to reserve a small number of audio channels
- Problem: N participants
  - Potentially N sources
  - ST-II needs to allocate a channel for all sources
    - » Resource requirements scale with number of members
    - » Limits maximum size of group
  - RSVP can use a Wildcard reservation for maximum number of simultaneous sources
    - » Resource requirements scale with number of simultaneous sources

Heterogeneous Groups

- What happens when receivers have different resource capabilities?
  - LAN receivers vs. modem receivers
- ST-II
  - Must allocate maximum requested resources along all links
  - Or, limited to lowest common denominator
- RSVP
  - Reserves minimum requirements for all downstream receivers
- Experiment
  - Complex network with random receivers and sources
  - RSVP 27.7% less than ST-II

Channel Selection

- What about selecting from a dynamic set of sources?
  - Channel selection
- Options
  - ST-II: Allocate channels for every source
    - » Too many simultaneous sources to allocate resources
  - RSVP Dynamic Filter: Receiver reserves enough resources for max simultaneous connections
    - » Filtering done in network at RSVP agent
    - » Which sources can be changed dynamically
  - RSVP Fixed Filter: Sources and reservations for them are fixed
    - » Chosen Source model
    - » Changing sources requires entirely new tree

Dynamic Analysis

- Network dynamics
  - What is the overhead of adapting to topology changes?
- ST-II
  - Hello messages among neighbors
  - Scales with number of agents, independent of reservations
- RSVP
  - Period Path and Reservation refresh messages
  - Merging used to ensure that only one reservation message is propogated over link
  - Scales with number of reservations

Group Membership Dynamics

- What is the overhead of group membership changes?
- Control messages
  - ST-II
    - » Connect and Accept exchanged between source and receiver
    - » Overhead on links proportional to # of downstream receivers
    - » More processing overhead closer to source (hot spots)
  - RSVP
    - » One message on link in both directions (Path, Reservation)
    - » Path is multicast, Reservation only has to reach closest branch
- Latency
  - ST-II: One round-trip time
  - RSVP: Depends on timers (e.g., Path refresh rate)

Discussion

- Some common design themes
  - Soft state
  - Receiver-driven
- ST-II maximum group size
  - I'm still unclear how they max group size of floor(bottleneck bandwidth/single stream resource request) + 1 -> why the +1 part?
- Equal comparison
  - Not surprisingly ST-II lost on all of the benchmarks. This is like making performance comparison between Voodoo I (1993) and GeForce3 (2001) graphics cards.

For Next Time...

- Send in an eval for only one of Stefan's papers
  - Both are interesting and well-written