15-441: Computer Networking

Lecture 23: QoS and Mobile/Wireless Networking

Overview

- RSVP
- Differentiated services
- Internet mobility
- TCP Over Noisy Links

Components of Integrated Services

1. Type of commitment
   What does the network promise?
2. Packet scheduling
   How does the network meet promises?
3. Service interface
   How does the application describe what it wants?
4. **Establishing the guarantee**
   How is the promise communicated
   How is admission of new applications controlled?

Service Interfaces

- Guaranteed Traffic
  - Host specifies rate to network
  - Why not bucket size b?
    - If delay not good, ask for higher rate
- Predicted Traffic
  - Specifies (r, b) token bucket parameters
  - Specifies delay D and loss rate L
  - Network assigns priority class
  - Policing at edges to drop or tag packets
    - Needed to provide isolation – why is this not done for guaranteed traffic?
      - WFQ provides this for guaranteed traffic

Resource Reservation Protocol (RSVP)

- Carries resource requests all the way through the network
- Main goal: establish "state" in each of the routers so they "know" how they should treat flows.
  - State = packet classifier parameters, bandwidth reservation, ..
- At each hop consults admission control and sets up reservation. Informs requester if failure

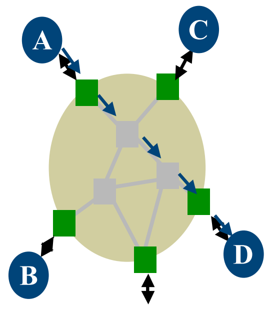

RSVP Motivation

- Resource reservation mechanism for multi-point applications
  - E.g., video or voice conference
  - Heterogeneous receivers
  - Changing membership
- Use network efficiently
  - Minimize reserved bandwidth
  - Share reservations between receivers
  - Limit control overhead (scaling).
  - Adapt to routing changes

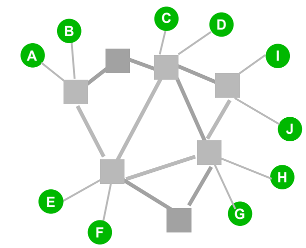

PATH Messages

- PATH messages carry sender's Tspec
  - Token bucket parameters
- Routers note the direction PATH messages arrived and set up *reverse path* to sender
- Receivers send RESV messages that follow reverse path and setup reservations
- If reservation cannot be made, user gets an error

RESV Messages

- Forwarded via reverse path of PATH
- Queuing delay and bandwidth requirements
- Source traffic characteristics (from PATH)
- Filter specification
  - Which transmissions can use the reserved resources
- Router performs admission control and reserves resources
  - If request rejected, send error message

Path and Reservation Messages

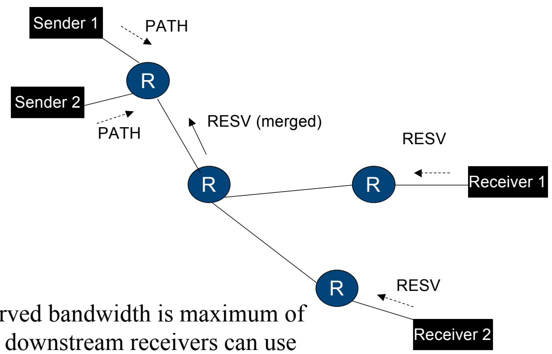

Reserved bandwidth is maximum of what downstream receivers can use

Soft State

- Periodic PATH and RESV msgs refresh established reservation state
  - Path messages may follow new routes
  - Old information times out
- Properties
  - Adapts to changes routes and sources
  - Recovers from failures
  - Cleans up state after receivers drop out

Overview

- RSVP
- Differentiated services
- Internet mobility
- TCP Over Noisy Links

Differentiated Services: Motivation and Design

- Edge routers do fine grain enforcement
  - Typically slower links at edge
  - E.g. mail sorting in post offices
  - Label packets with a type field
    - Uses IP TOS bits
    - E.g. a priority stamp
- Core routers process packets based on packet marking and defined per hop behavior
- More scalable than IntServ
  - No per flow state or signaling

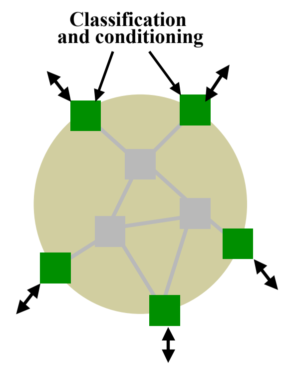

Expedited Forwarding PHB

User sends within profile & network commits to delivery with requested profile
- Strong guarantee
- Possible service: providing a virtual wire
- Admitted based on peak rate

- Rate limiting of EF packets at edges only, using token bucket to shape transmission

- Simple forwarding: classify packet in one of two queues, use priority
  - EF packets are forwarded with minimal delay and loss (up to the capacity of the router)

Expedited Forwarding Traffic Flow

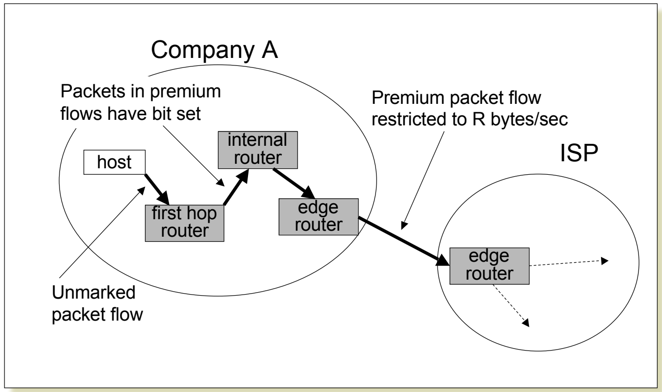

Assured Forwarding PHB

- AF defines 4 classes
  - Strong assurance for traffic within profile & allow source to exceed profile
    - Implement services that differ relative to each other (e.g., gold service, silver service…)
  - Admission based on expected capacity usage profiles
  - Within each class, there are three drop priorities
    - Traffic unlikely to be dropped if user maintains profile

- User and network agree to some traffic profile
  - Edges mark packets up to allowed rate as "in-profile" or high priority
  - Other packets are marked with one of 2 lower "out-of-profile" priorities
  - A congested router drops lower priority packets first
    - Implemented using clever queue management (RED with In/Out bit)

Edge Router Input Functionality

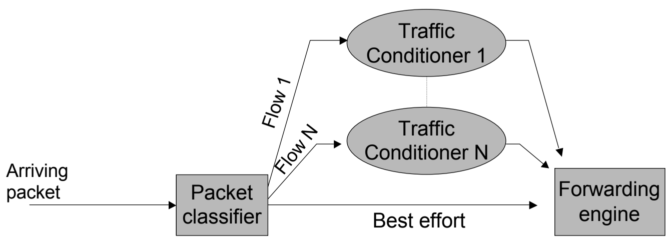

classify packets based on packet header

Traffic Conditioning

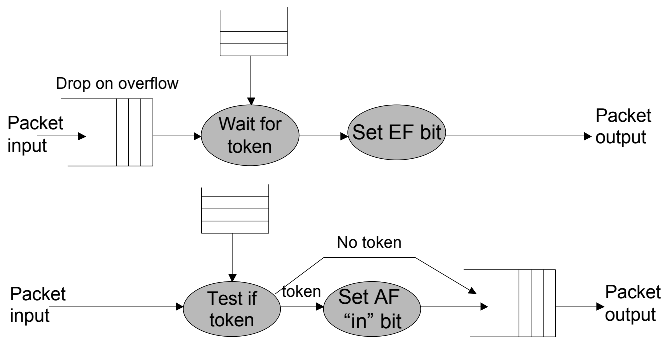

Router Output Processing

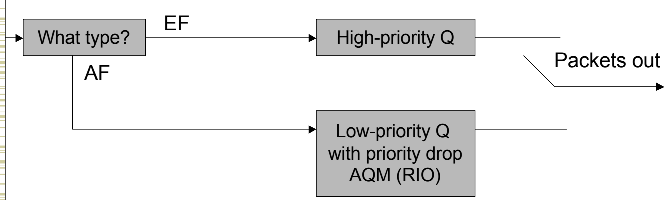

Edge Router Policing

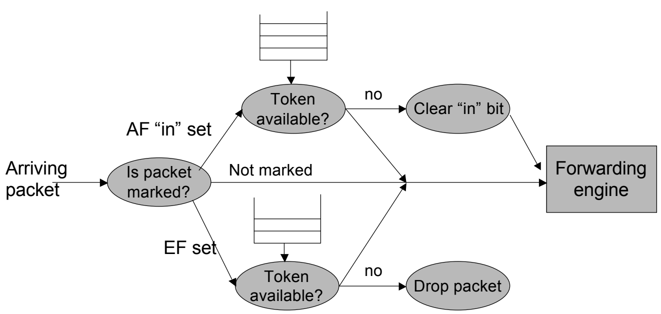

Comparison

	Best-Effort	Diffserv	Intserv
Service	• Connectivity • No isolation • No guarantees	• Per aggregation isolation • Per aggregation guarantee	• Per flow isolation • Per flow guarantee
Service Scope	• End-to-end	• Domain	• End-to-end
Complexity	• No set-up	• Long term setup	• Per flow setup
Scalability	• Highly scalable • (nodes maintain only routing state)	• Scalable (edge routers maintains per aggregate state; core routers per class state)	• Not scalable (each router maintains per flow state)

Overview

- RSVP
- Differentiated services
- Internet mobility
- TCP Over Noisy Links

Wireless Challenges

- Force us to rethink many assumptions
- Need to share airwaves rather than wire
  - Don't know what hosts are involved
  - Host may not be using same link technology
- Mobility
- Other characteristics of wireless
  - Noisy → lots of losses
  - Slow
  - Interaction of multiple transmitters at receiver
    - Collisions, capture, interference
  - Multipath interference

Routing to Mobile Nodes

- Obvious solution: have mobile nodes advertise route to mobile address/32
  - Should work!!!
- Why is this bad?
  - Consider forwarding tables on backbone routers
    - Would have an entry for each mobile host
    - Not very scalable
- What are some possible solutions?

How to Handle Mobile Nodes? (Addressing)

- Dynamic Host Configuration (DHCP)
  - Host gets new IP address in new locations
  - Problems
    - Host does not have constant name/address → how do others contact host
    - What happens to active transport connections?

How to Handle Mobile Nodes? (Naming)

- Naming
  - Use DHCP and update name-address mapping whenever host changes address
  - Fixes contact problem but not broken transport connections

How to Handle Mobile Nodes? (Transport)

- TCP currently uses 4 tuple to describe connection
  - <Src Addr, Src port, Dst addr, Dst port>
- Modify TCP to allow peer’s address to be changed during connection
- Security issues
  - Can someone easily hijack connection?
- Difficult deployment → both ends must support mobility

How to Handle Mobile Nodes? (Link Layer)

- Link layer mobility
  - Learning bridges can handle mobility → this is how it is handled at CMU
  - Encapsulated PPP (PPTP) → Have mobile host act like he is connected to original LAN
    - Works for IP AND other network protocols

How to Handle Mobile Nodes? (Routing)

- Allow mobile node to keep same address and name
- How do we deliver IP packets when the endpoint moves?
  - Can’t just have nodes advertise route to their address
- What about packets from the mobile host?
  - Routing not a problem
  - What source address on packet? → this can cause problems
- Key design considerations
  - Scale
  - Incremental deployment

Basic Solution to Mobile Routing

- Same as other problems in computer science
  - Add a level of indirection
- Keep some part of the network informed about current location
  - Need technique to route packets through this location (interception)
- Need to forward packets from this location to mobile host (delivery)

Interception

- Somewhere along normal forwarding path
  - At source
  - Any router along path
  - Router to home network
  - Machine on home network (masquerading as mobile host)
- Clever tricks to force packet to particular destination
  - "Mobile subnet" – assign mobiles a special address range and have special node advertise route

Delivery

- Need to get packet to mobile's current location
- Tunnels
  - Tunnel endpoint = current location
  - Tunnel contents = original packets
- Source routing
  - Loose source route through mobile current location

Mobile IP (RFC 2290)

- Interception
  - Typically home agent – a host on home network
- Delivery
  - Typically IP-in-IP tunneling
  - Endpoint – either temporary mobile address or foreign agent
- Terminology
  - Mobile host (MH), correspondent host (CH), home agent (HA), foreign agent (FA)
  - Care-of-address, home address

Mobile IP (MH at Home)

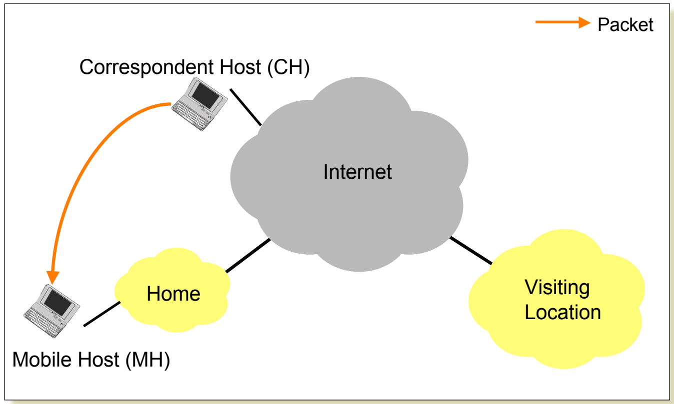

Mobile IP (MH Moving)

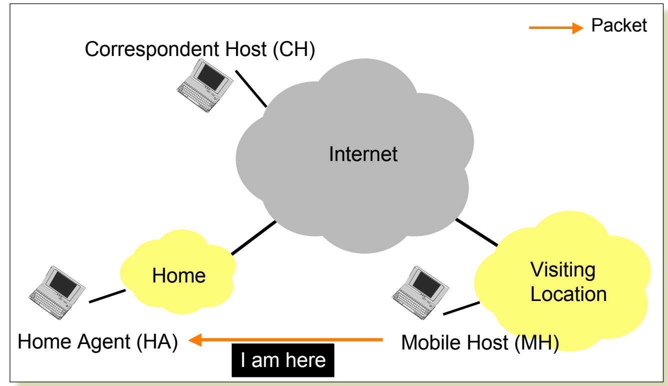

Mobile IP (MH Away – FA)

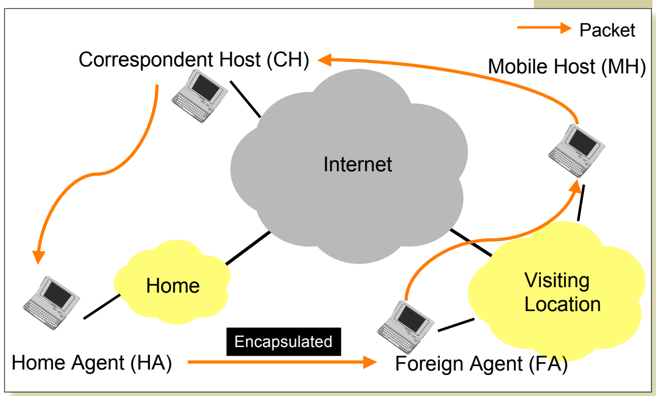

Mobile IP (MH Away - Collocated)

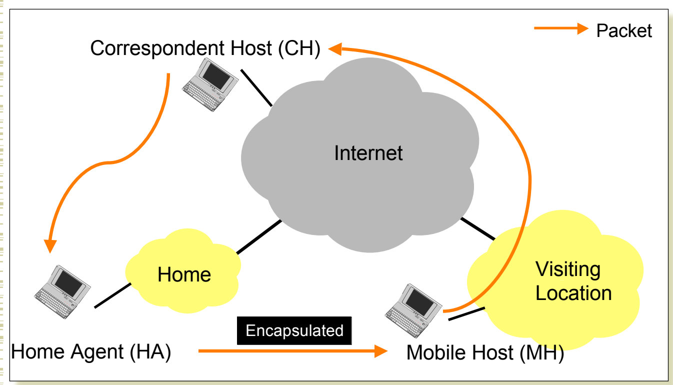

Other Mobile IP Issues

- Route optimality
  - Resulting paths can be sub-optimal
  - Can be improved with route optimization
    - Unsolicited binding cache update to sender
- Authentication
  - Registration messages
  - Binding cache updates
- Must send updates across network
  - Handoffs can be slow
- Problems with basic solution
  - Triangle routing
  - Reverse path check for security

EXTRA SLIDES

The rest of the slides are FYI

RSVP Goals

- Used on connectionless networks
  - Should not replicate routing functionality
  - Should co-exist with route changes
- Support for multicast
  - Different receivers have different capabilities and want different QOS
  - Changes in group membership should not be expensive
  - Reservations should be aggregate – I.e. each receiver in group should not have to reserve
  - Should be able to switch allocated resource to different senders
- Modular design – should be generic “signaling” protocol
- Result
  - Receiver-oriented
  - Soft-state

RSVP Service Model

- Make reservations for simplex data streams
- Receiver decides whether to make reservation
- Control msgs in IP datagrams (proto #46)
- PATH/RESV sent periodically to refresh soft state
- One pass:
  - Failed requests return error messages - receiver must try again
  - No e2e ack for success

RSVP State in Switches

- Have to keep sink tree information.
  - no such thing as inverse multicast routing
- Also have to keep information on sources if filters are used.
  - selected in path message
  - used in aggregation and propagating propagating information to switches
- Also used in limiting protocol overhead.
  - switches do not propagate periodic reservation and path messages
  - they periodically regenerate copies that summarize the information they have
- Raises concerns about scalability.

Receiver Initiated Reservations

- Receiver initiates reservation by sending a reservation over the sink tree.
  - Assumes multicast tree has been set up previously
    - also uses receiver-initiated mechanism
  - Uses existing routing protocol, but routers have to store the sink tree (reverse path from forwarding path)
- Properties.
  - Scales well: can have parallel independent connect and disconnect actions - single shared resource required
  - Supports receiver heterogeneity: reservation specifies receiver requirements and capabilities

DiffServ

- Analogy:
  - Airline service, first class, coach, various restrictions on coach as a function of payment
- Best-effort expected to make up bulk of traffic, but revenue from first class important to economic base (will pay for more plentiful bandwidth overall)
- <u>Not as motivated by real-time</u>! Motivated by economics and assurances
  - Supports QoS for flow aggregates.
  - Architecture does not preclude more fine grain guarantees

Per-hop Behaviors (PHBs)

- Define behavior of individual routers rather than end-to-end services – there may be many more services than behaviors
- Multiple behaviors – need more than one bit in the header
- Six bits from IP TOS field are taken for Diffserv code points (DSCP)

Red with In or Out (RIO)

- Similar to RED, but with two separate probability curves
- Has two classes, “In” and “Out” (of profile)
- “Out” class has lower Min_{thresh}, so packets are dropped from this class first
  - Based on queue length of all packets
- As avg queue length increases, “in” packets are also dropped
  - Based on queue length of only “in” packets

RIO Drop Probabilities