Express: Holbrook and Cheriton [Holbrook99a]

CSci551: Computer Networks
SP2006 Thursday Section
John Heidemann

Key ideas

- routing protocols
  - EXPRESS Grp Mgt Protocol (EGMP)
  - ways to tell how many people are in the group
  - protocols to cope with only one sender (ex. session relay)
- service model
  - concept of the channel (mcast group)
  - channel has only *one* sender
- using simple encryption to secure groups
- deployment problems with IP multicast
  - billing issues: IP multicast bills on what you send
  - solution: need to be able to count members, need to protect channels (encryption), sender needs to control content

IP Multicast Problems

- need billing mechanism
  - need to know number of subscribers
- need access control
  - need to limit who can send and subscribe
  - ISPs concerned about mcast
- IPv4 mcast addresses too limited
- current protocols too complex (particularly PIM)

⇒ *single source* multicast

Express vs. Multicast Problems

- need billing mechanism
  - count how big group is
  - have single source (know who to bill)
- need access control
  - encryption keys per channel (but a bit weak because rtrs must know keys)
  - only allow sender to send to group
- IPv4 mcast addresses too limited
  - identify groups by sender address and additional group number

Express Approach

- all addresses are source specific (S,E)
  - 2^{24} channels per source, (2^{32} sources)
- access control
  - only source can send
  - channels optionally protected by "key" (really just a secret)
- sub-cast support (encapsulate pkt to any router on the tree [if you know who they are])
- best-effort counting service

Express Components

- ECMP: Express Count Mgt Protocol
  - like IGMP, but also adds *count* support
  - counts used to determine receivers or for other things like voting
    - not clear how general
- session relays
  - service at source that can relay data on to tree (similar to PIM tunneling)

Why Single Source?

- easier to count?
  - not really
- good match for applications
  - yes, definitely for some
- simpler routing protocol
  - true… definitely much simplier than PIM
- simplifies access control and billing
  - addresses what they see as the main problem with IP mcast

Observations

- Simpler? yes
- Enough to justify mcast to ISPs? not clear

Another Alternative: Application-level Multicast

- if the ISPs won't give us multicast, we'll take it :-)
- just do it all at the app
- results in some duplicated data on links
- and app doesn't have direct access to unicast routing
- but can work… (ex. Yoid project at ISI); Narada paper

Application-level Multicast Example

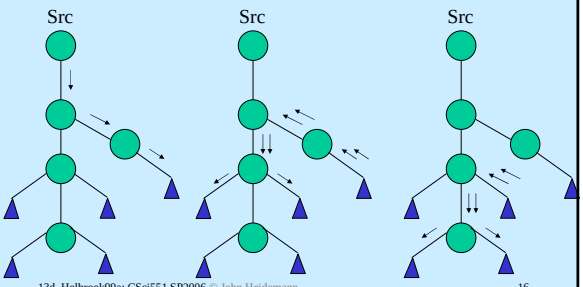

App-level Multicast

- Simplest approach:
  - send data to central site that forwards
- Better approaches:
  - try to balance load on any one link
  - try to topologically cluster relays

Other questions/observations?

- xxx