RSVP Aggregation

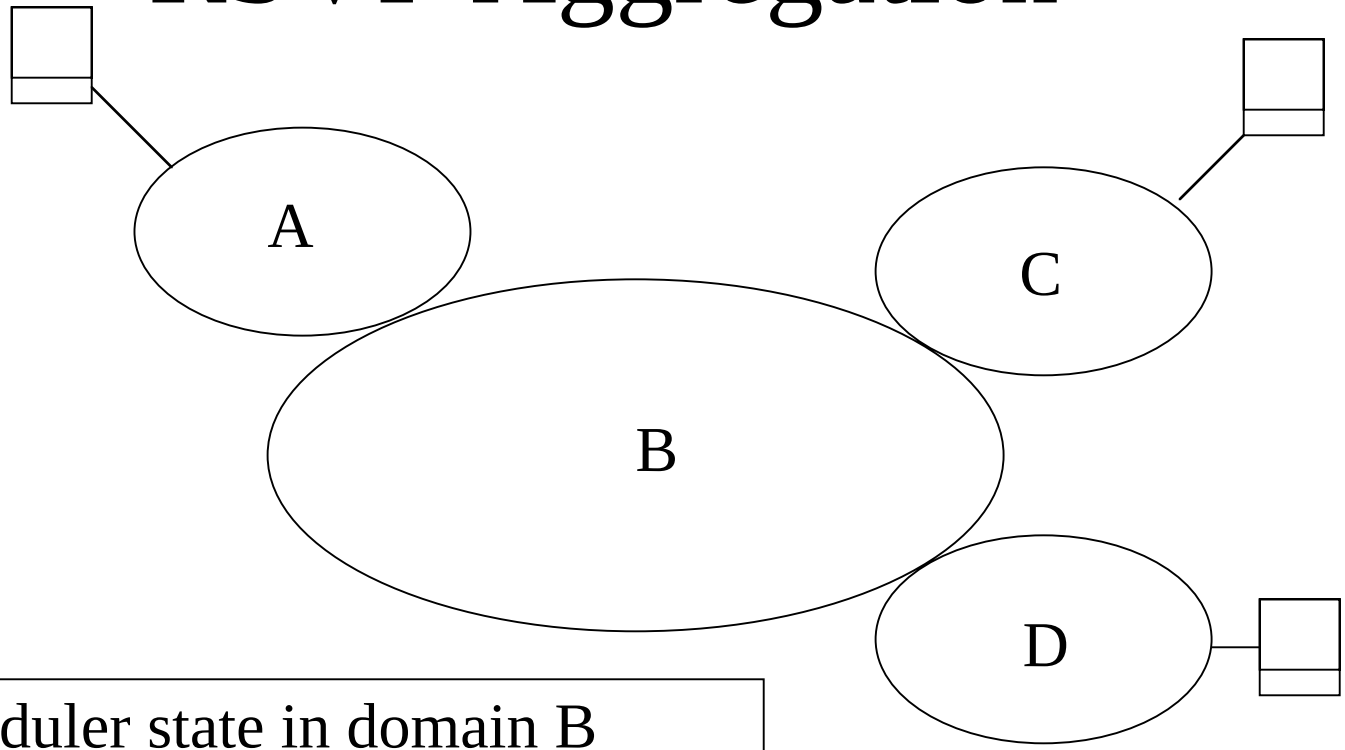

Limit scheduler state in domain B

Limit RSVP control state in domain B

Draft-guerin-aggreg-rsvp-00.txt
Draft-berson-rsvp-aggregation-00.ps

Requirements (in aggregated domains)

- No (limited) awareness of individual RSVP flows
  - Scheduling state
  - RSVP state
- Must satisfy QoS requirements of individual flows
  - Delivered service
  - Path characterization (ADSPEC)
- Must maintain isolation of flows
  - Nonconformant traffic of one flow must not zap others
- Must not limit ability to support individual flow reservations in other domains
  - Particularly, must be able to un-aggregate

Draft-guerin...

- Document space
  - Survey - range of possibilities
  - Assumes unicast
  - Covers full spectrum of issues
    - RSVP state management
    - Aggregate scheduling requirements (briefly)
    - Admission control
    - Path characterization
  - All services

Tunneling Approach

- Point-to-point RSVP tunnels between ingress and egress routers of the aggregated domain
- End-to-end QoS data and RSVP messages tunnel through domain

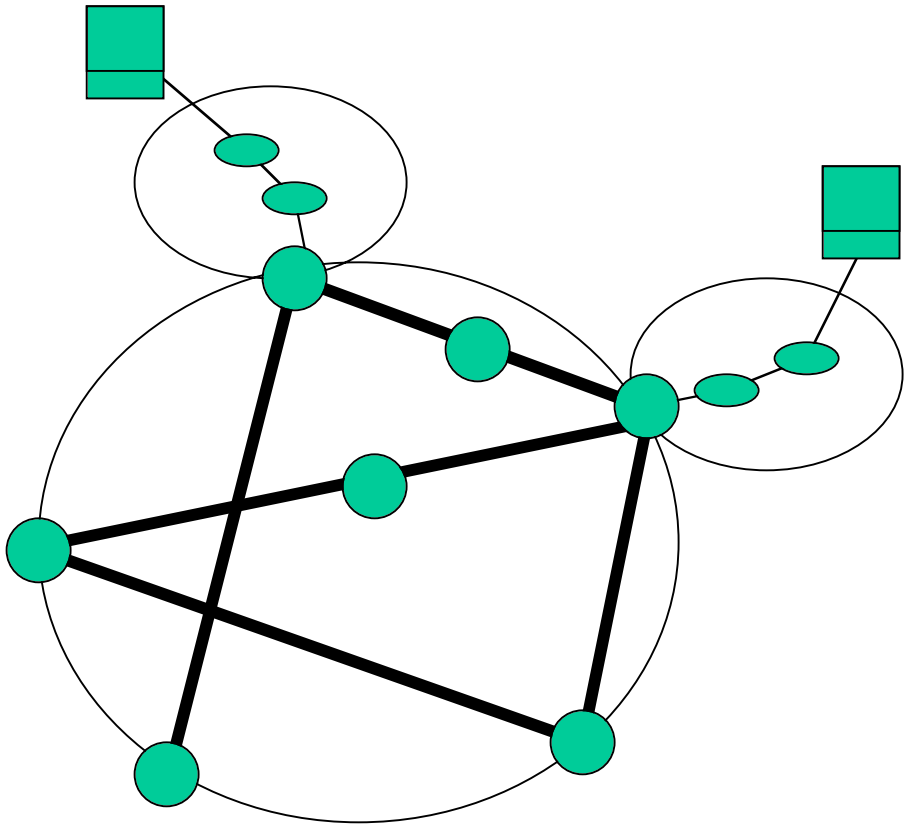

Tunneling

- Ingress, egress routers do RSVP for aggregate pipe resources
- Egress routers adjust reservations for flow arrival/departure and routing changes
- Individual PATH adspecs updated by egress based on cached aggregate PATH adspec

- Drawbacks
  - Difficulty maintaining isolation of individual flows
    - If using same service, must send nonconformant traffic outside the tunnel
  - Encap/decap/data overhead

Tag-based aggregation

- Approach
  - Scheduling class in core of net selected by packet tag set at ingress
  - Class selected based on reservation request
  - Individual RSVP messages sent through domain transparently
  - Individual PATH messages (adspec) updated at egress

- Issues
  - Definition of appropriate scheduling behavior for aggregated flows
  - Transparent location of flow egress/ingress points
  - (Dynamic) allocation of backbone resources
  - Computation of adspec information

Class-based RSVP approach

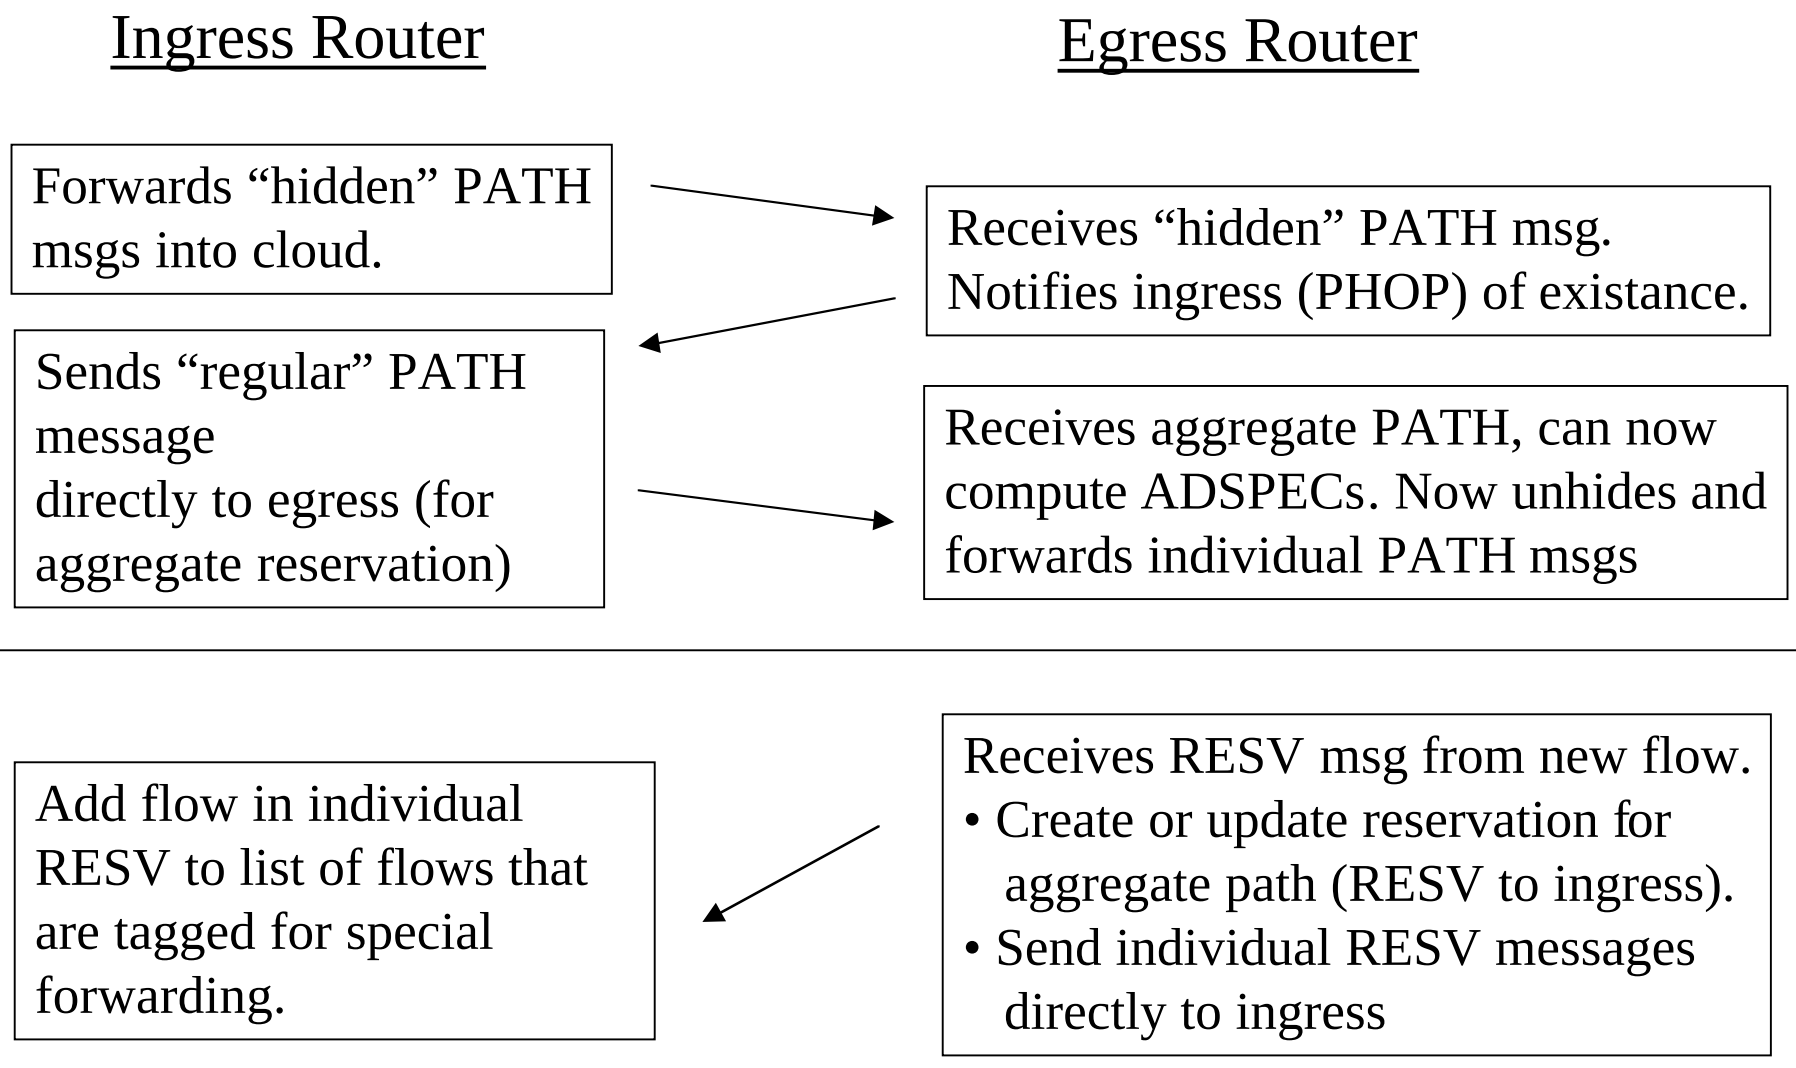

Draft-berson...

- Document space
  - Unicast and multicast
  - Controlled-load like services
    - Assumes measurement-based admission control
    - Does not explicitly discuss adspec/characterization issues
  - Primary focus on RSVP / signalling

Berson - unicast

- Flows scheduled as aggregated traffic classes in the region
- Retains hop-by-hop per-path admission control
- Packets are tagged with class by ingress router
- Interior routers ignore normal RSVP messages (RSVP “off”)
- Arrival of RESV at region ingress triggers hop-by-hop admission control along data path (Adreq)
- Collected result returned to ingress router (Adrep) – starts tagging packets if reservation accepted

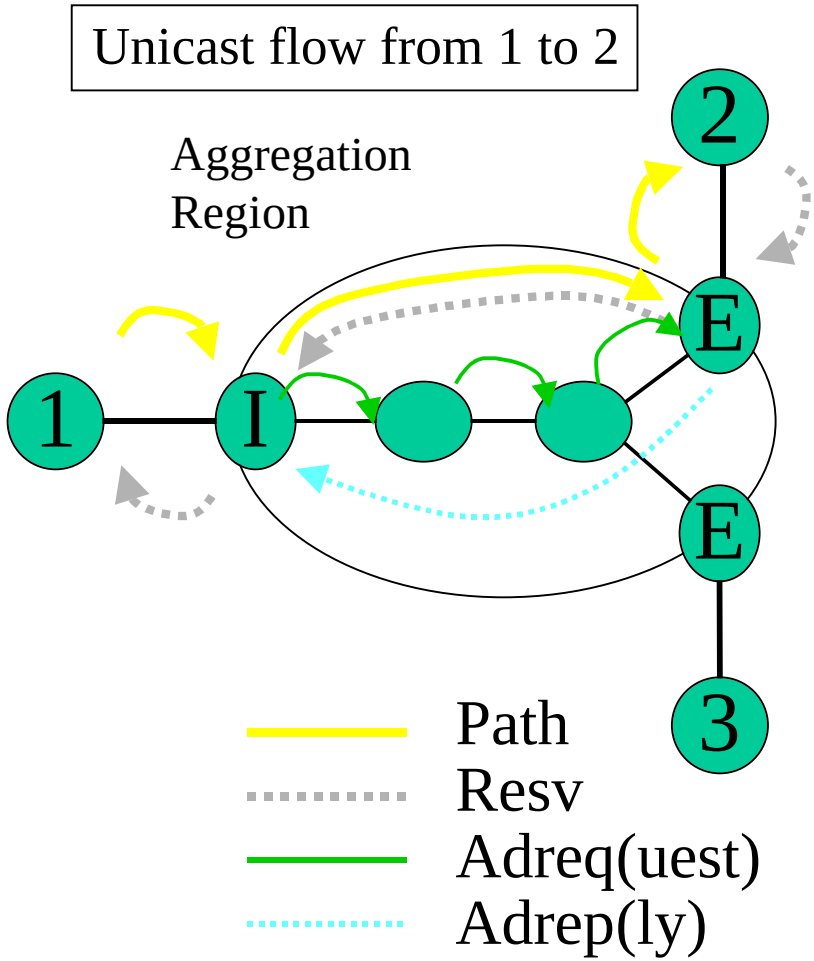

Berson - multicast

- Multicast flows require some state within the region
  - Heterogeneous reservations
  - Qos and best-effort branches
- Creation of split point “wakes up” RSVP session
  - Next PATH message creates RSVP state. Router installs “retagger” to remark packets leaving reserved path to new QoS
- Admission control done over homogeneous reservation segments
  - Newly awake router becomes endpoint for adreq/adrep msgs

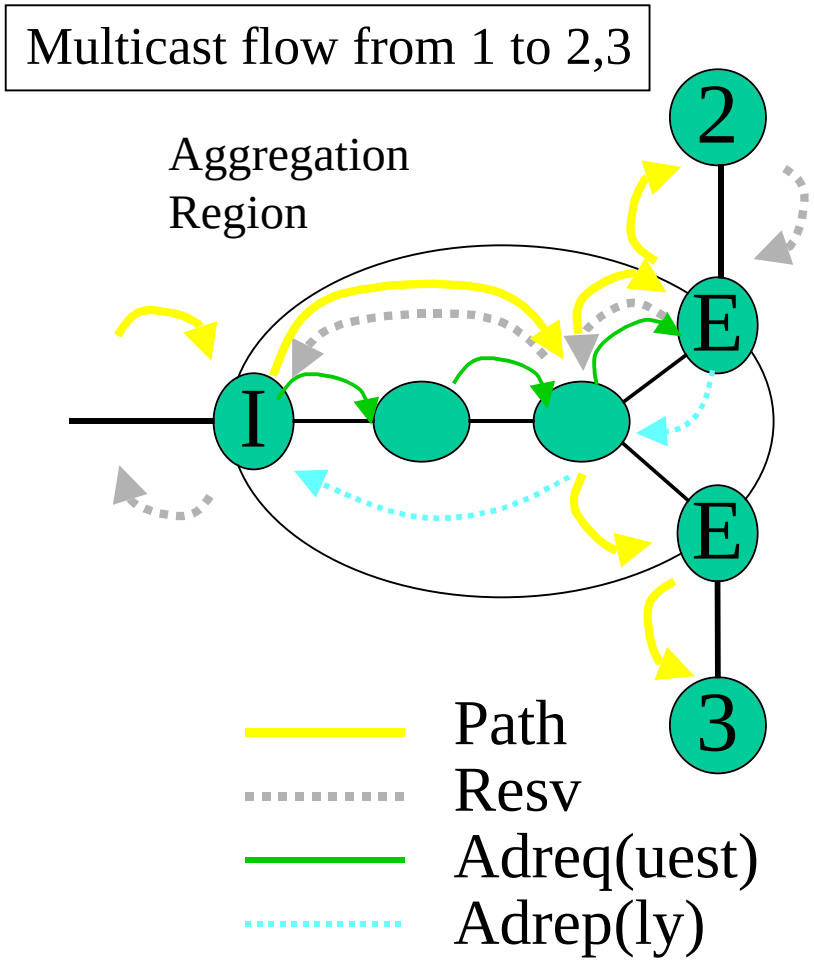

Comments

- Both of these drafts assume per-path dynamic management of aggregate resources
- Both suggest small changes to RSVP
- Neither creates a clear and shining line between “outer” RSVP and “inner” mechanism
  - But both are close
  - Some efficiency gained by lack of complete separation
- Consider these ideas as tools in a toolkit
  - Look for a clean functional interface
  - That can support these and other techniques