

Layering the Inter-Domain Layer

or Putting Programmable Routers to Work

Vytautas Valancius and Nick Feamster
College of Computing
Georgia Institute of Technology

Imagine:

A joint 'live' course is arranged between Princeton University, US and Royal Institute of Technology, Sweden



←→
Bidirectional Video stream: ~7Mbps
Negligible jitter, minimum delay



Problem:

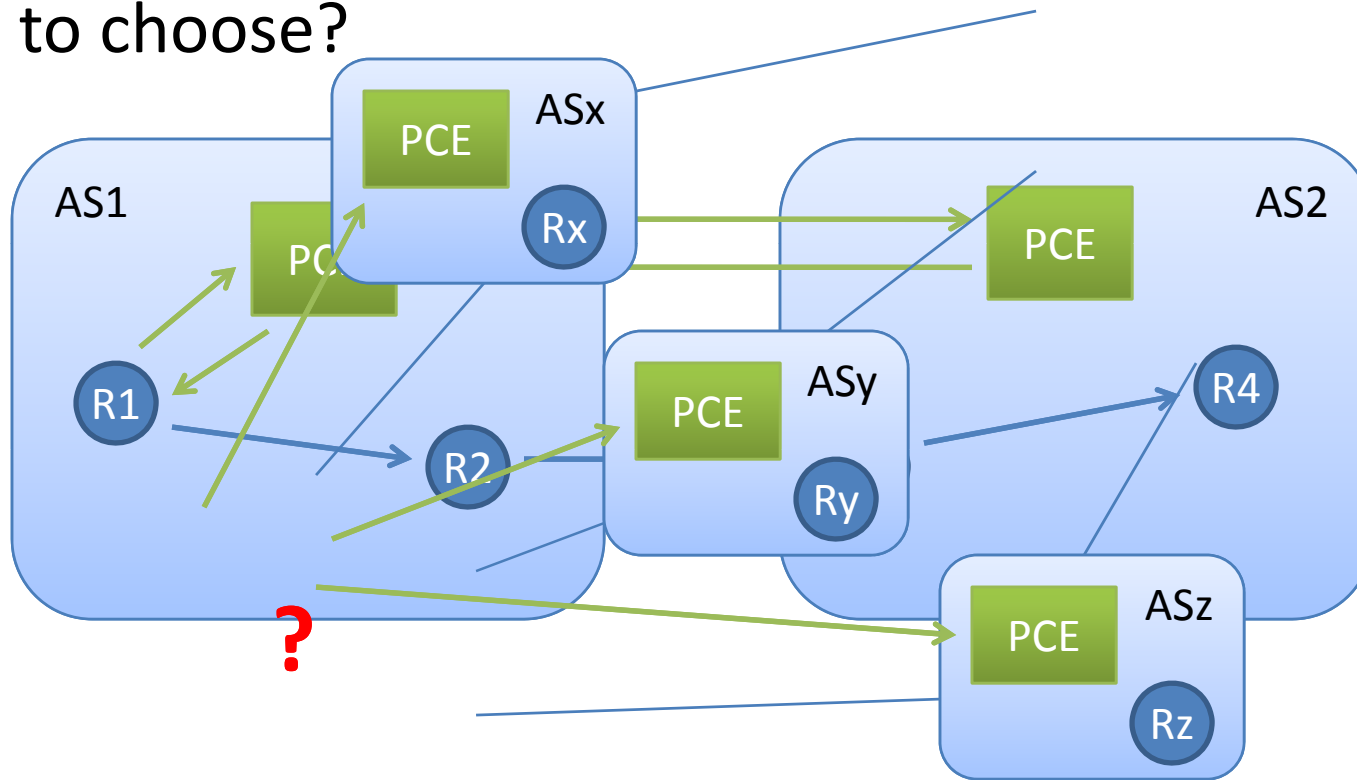
How to allocate inter-domain path meeting certain demands?

State of The Art Today

- Digital circuit or TE MPLS path
 - Must be within one domain (otherwise expect provisioning delays)
- TE path and Path Computation Element (PCE) architecture
 - PCE provides a distributed way to calculate the exact route between domains without exposing topology

Inter-Domain Resource Reservation

Which inter-domain path to choose?



PCE – Path Computation Element

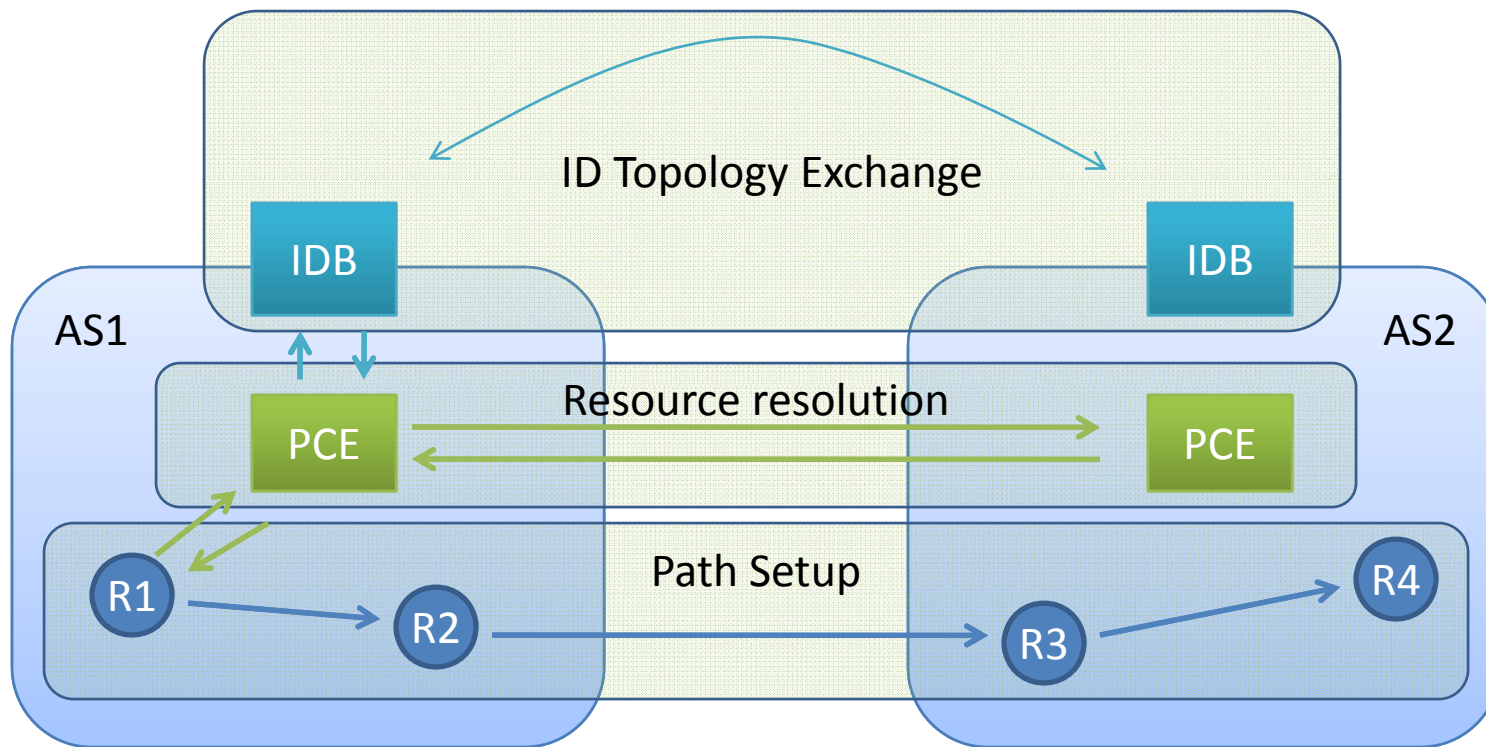
→ Path Request and Reply

→ Path Setup

Position

- For most inter-domain services, it is possible and beneficial to separate control plane into three distinct layers:
 - Inter-domain topology (?)
 - Resource resolution (i.e. PCE)
 - Service setup and monitoring (i.e. RSVP)
- It is possible to replace ‘?’ with a scalable protocol:
 - Only abstract inter-domain information is needed
 - No need to be up-to-date or support policies (irregularities are handled by resource resolution layer)
 - No need to react to datapath failures (service setup and monitoring takes care of that)

How This Would Work?



Research Questions

- What kind of architecture inter-domain topology service should employ?
 - Centralized vs hierarchical vs distributed
 - Domain-state protocol with flooding?
- What kind of information should be exposed?
- How often information is updated?

Thank You!

Backup

- H. Wang et. al. 'A Hardware-Accelerated Implementation of the RSVP-TE Signaling Protocol' – 400'000 path setups/sec