Internet Exchanges and BGP Route Servers

Limelight Networks
Elisa Jasinska <elisa@llnw.com>

Interconnection

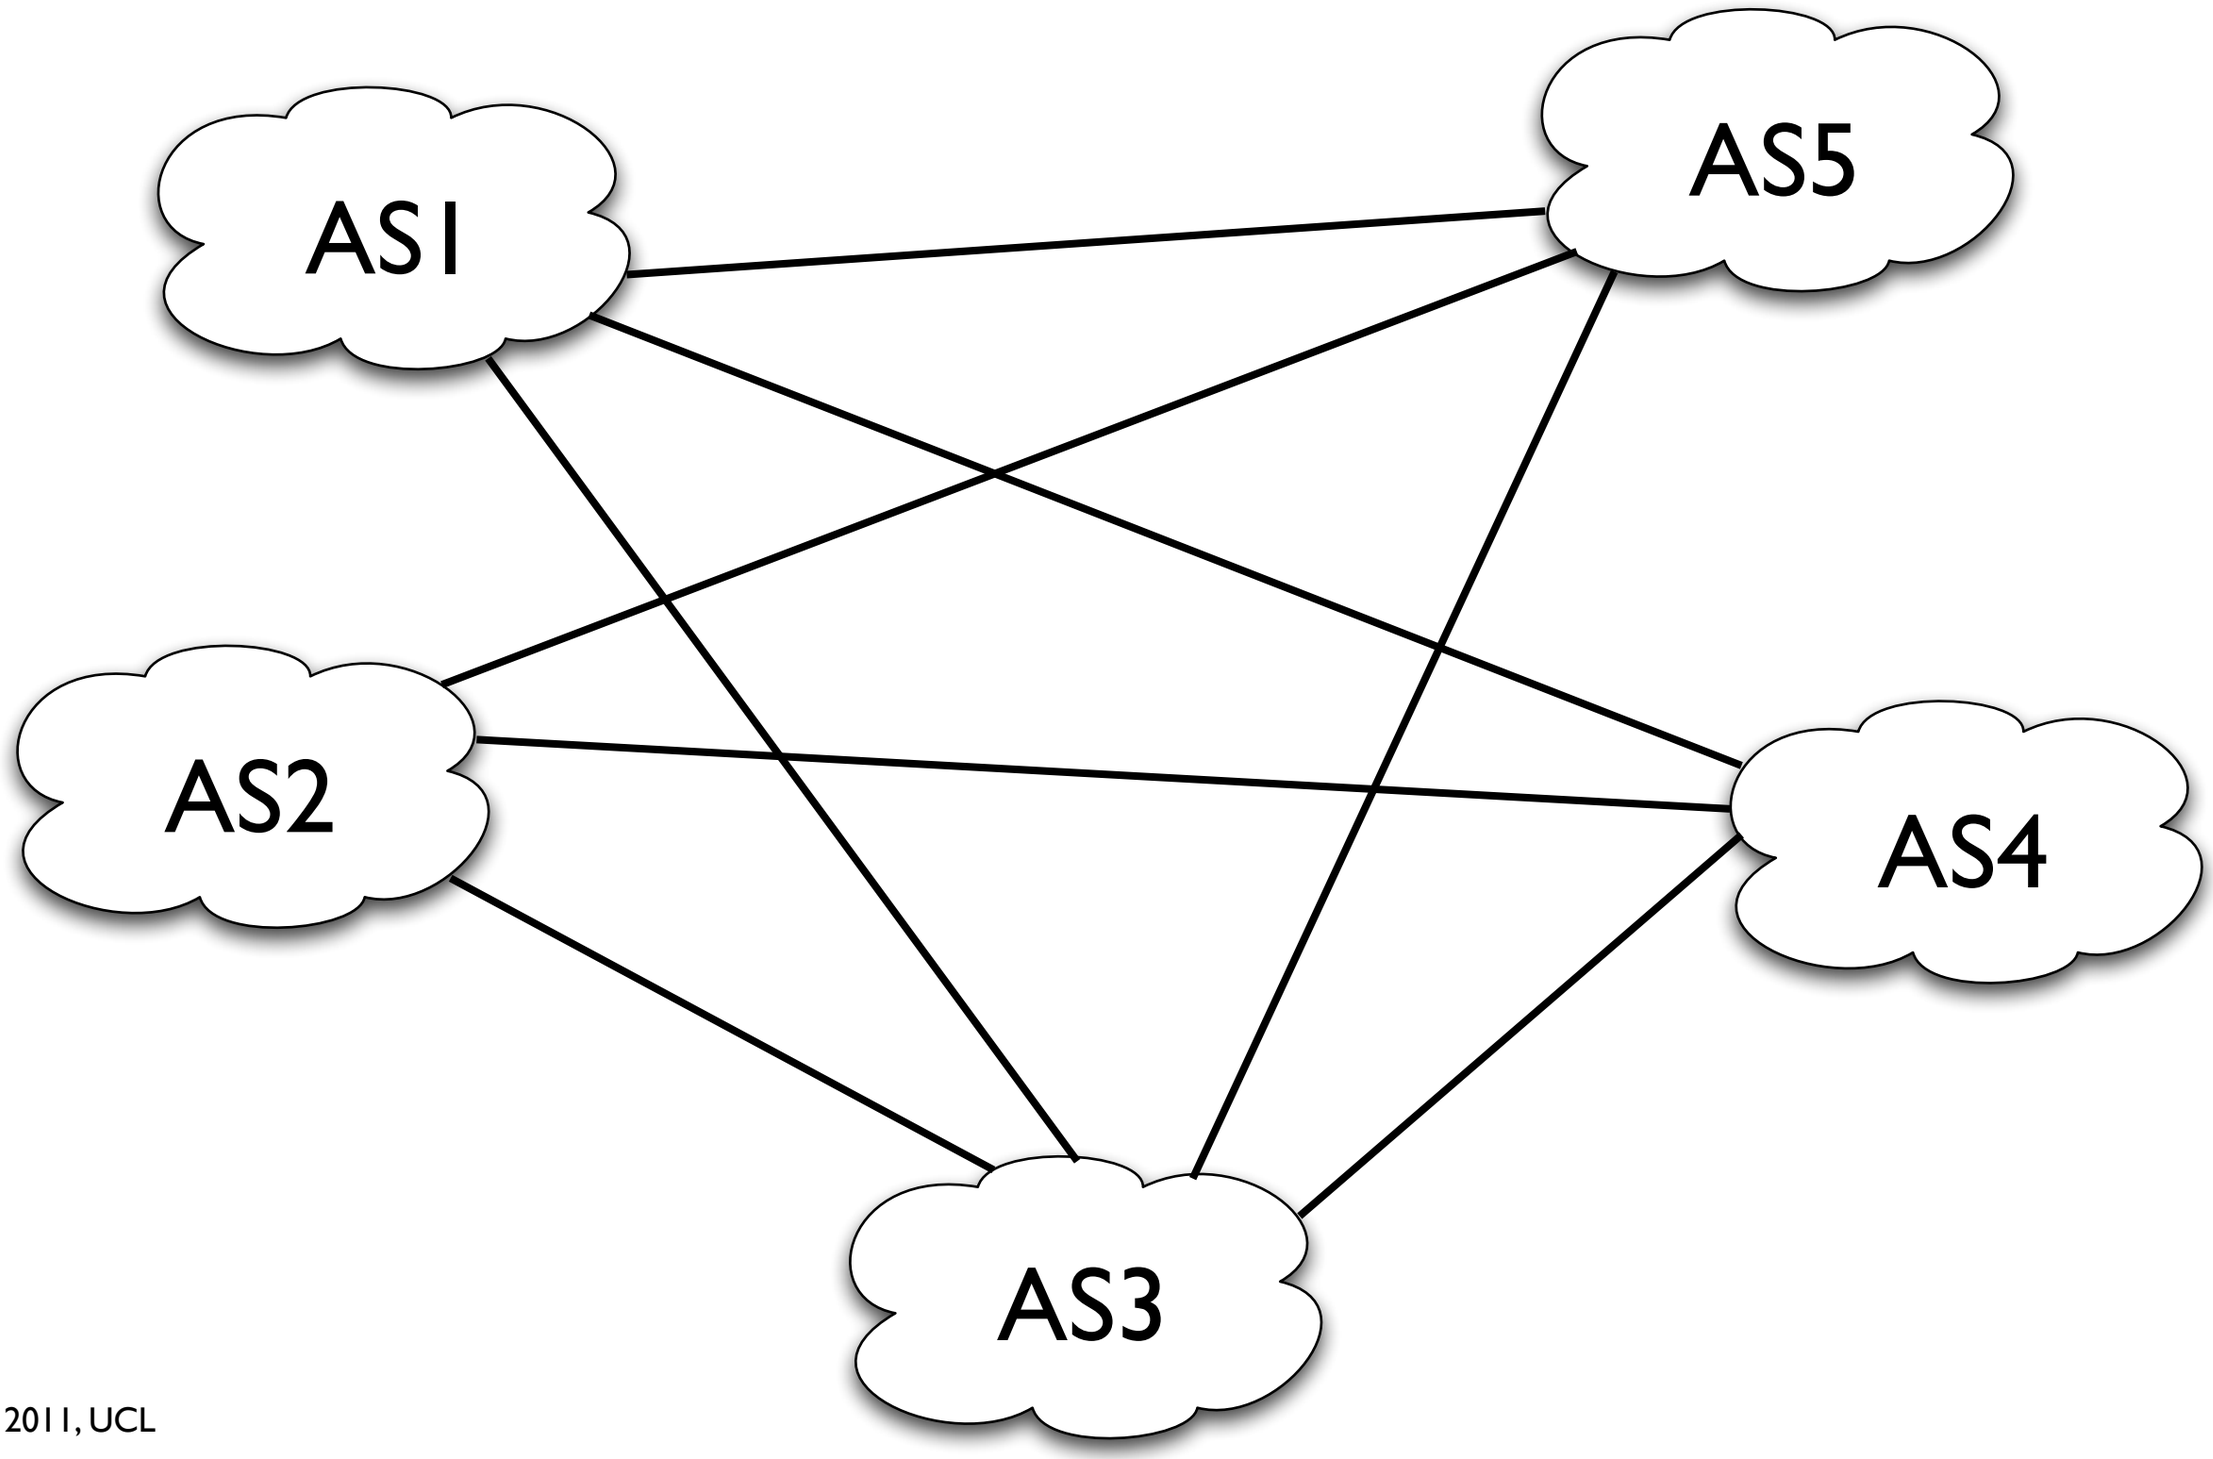

Internet Exchange

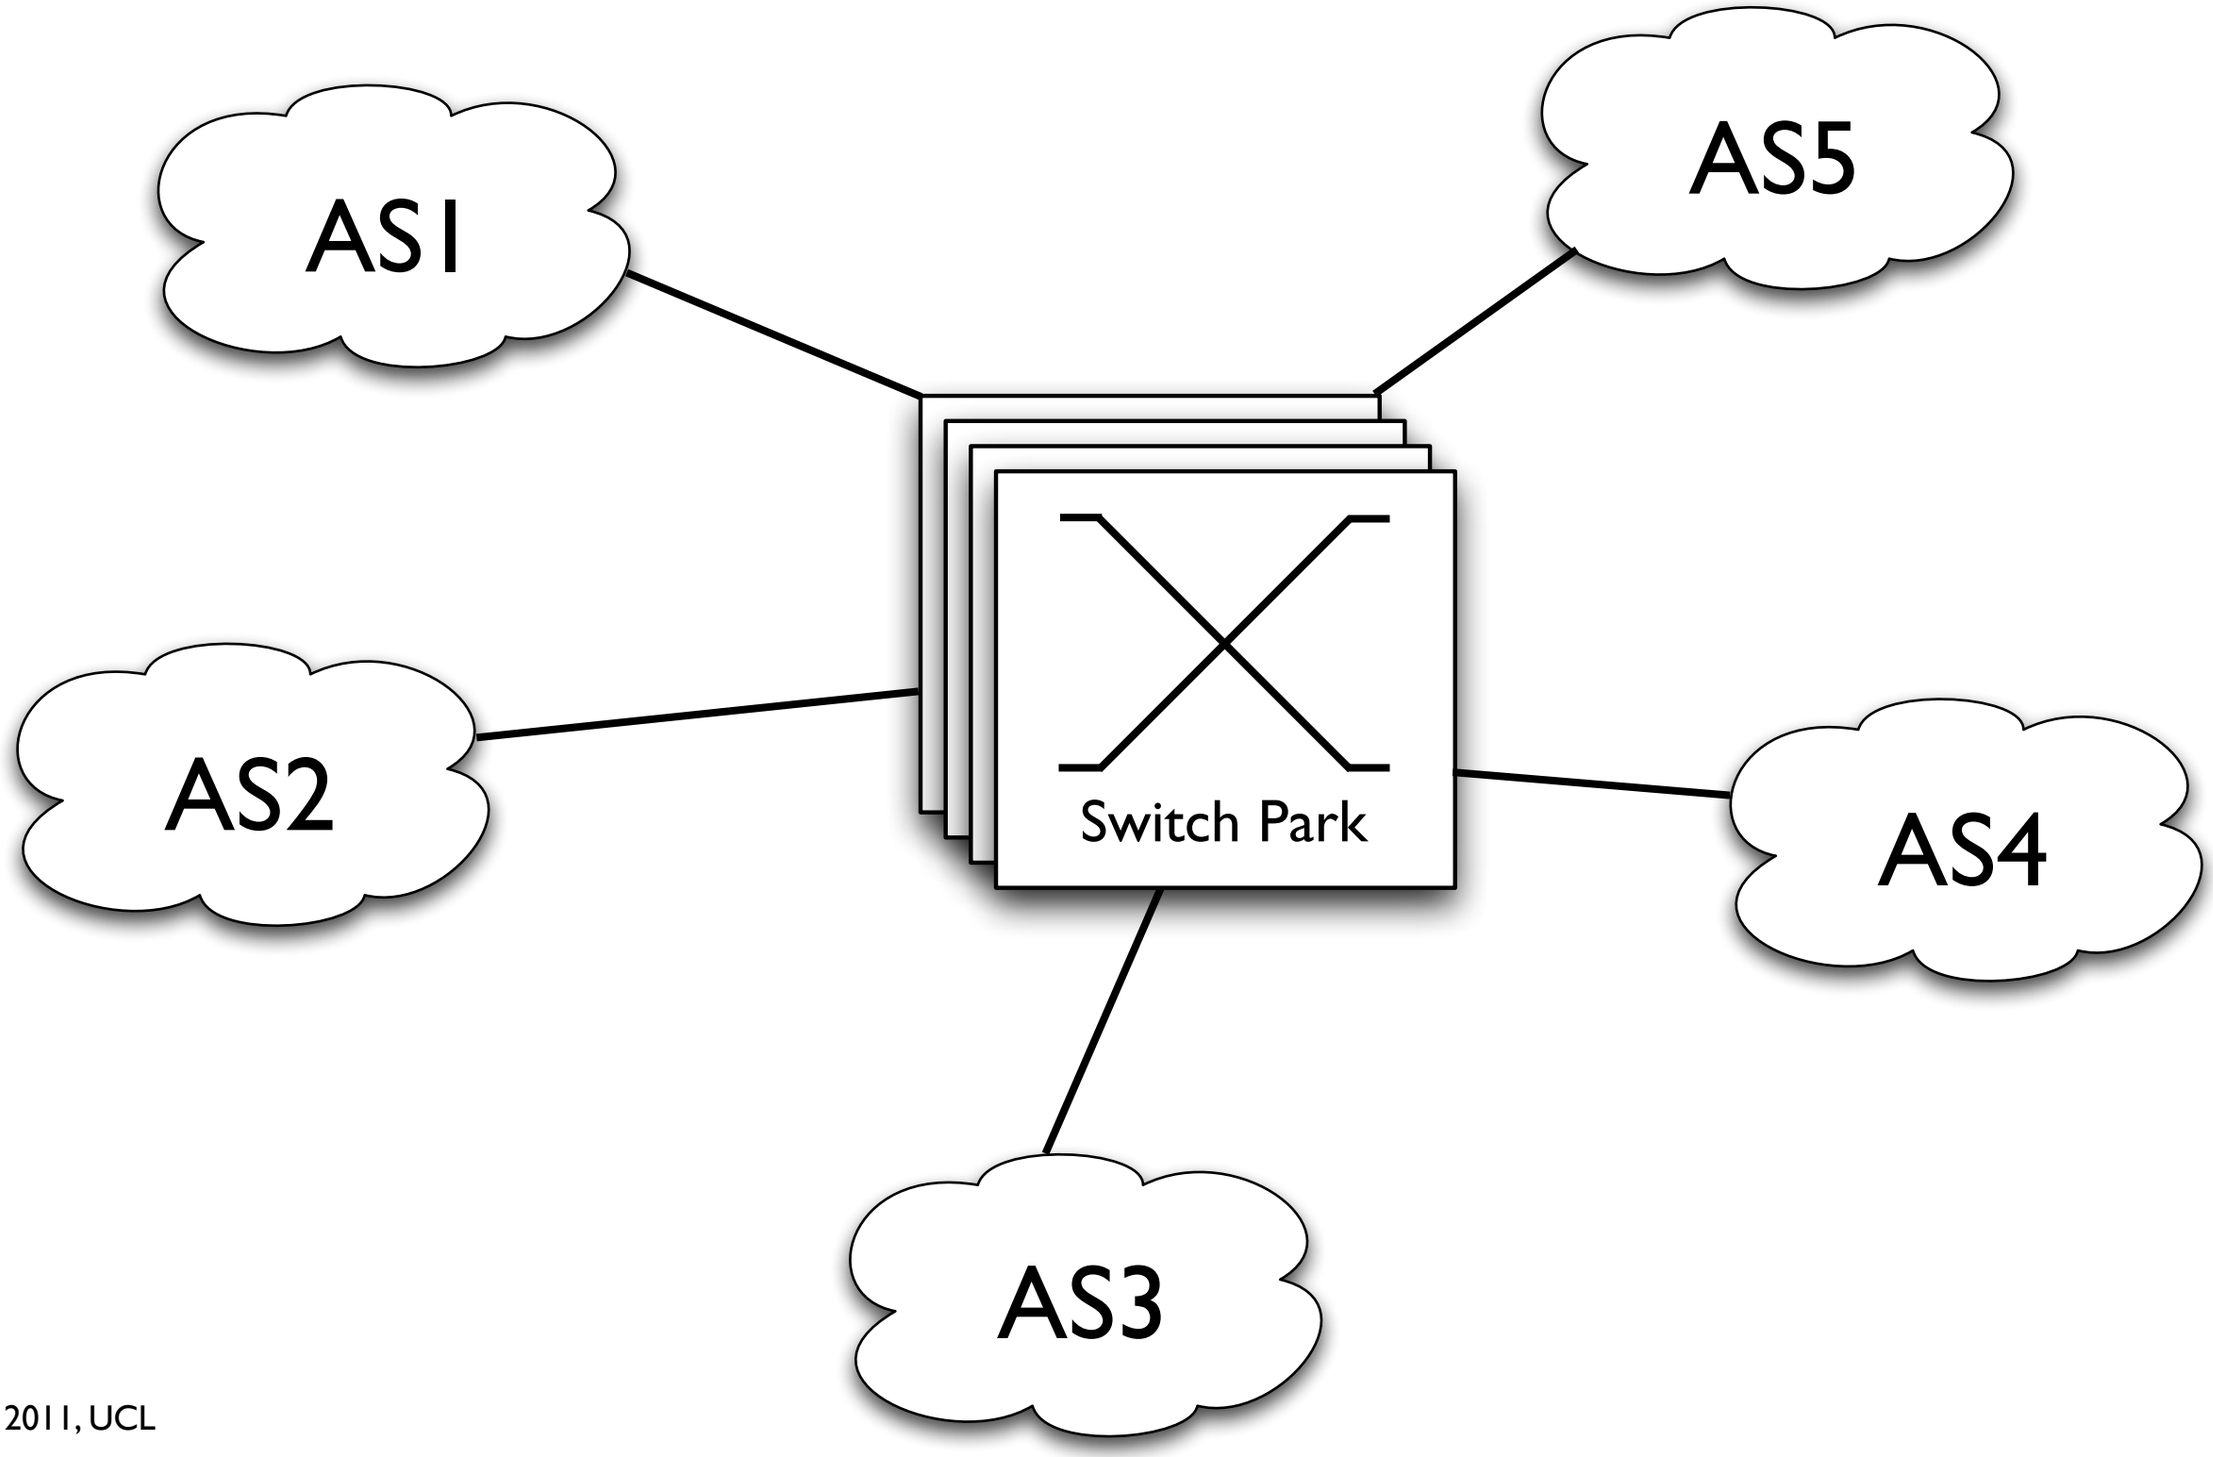

Internet Exchanges

- Layer-2 platform for peers to interconnect
- 121 IXPs in Europe
- Examples: AMS-IX in Amsterdam, LINX in London, DE-CIX in Frankfurt, Espanix in Madrid, Equinix in various locations, ...
- Different structure in different parts of the world:
  - Commercial vs. non-for-profit
  - Independent vs. tied into other services

AMS-IX

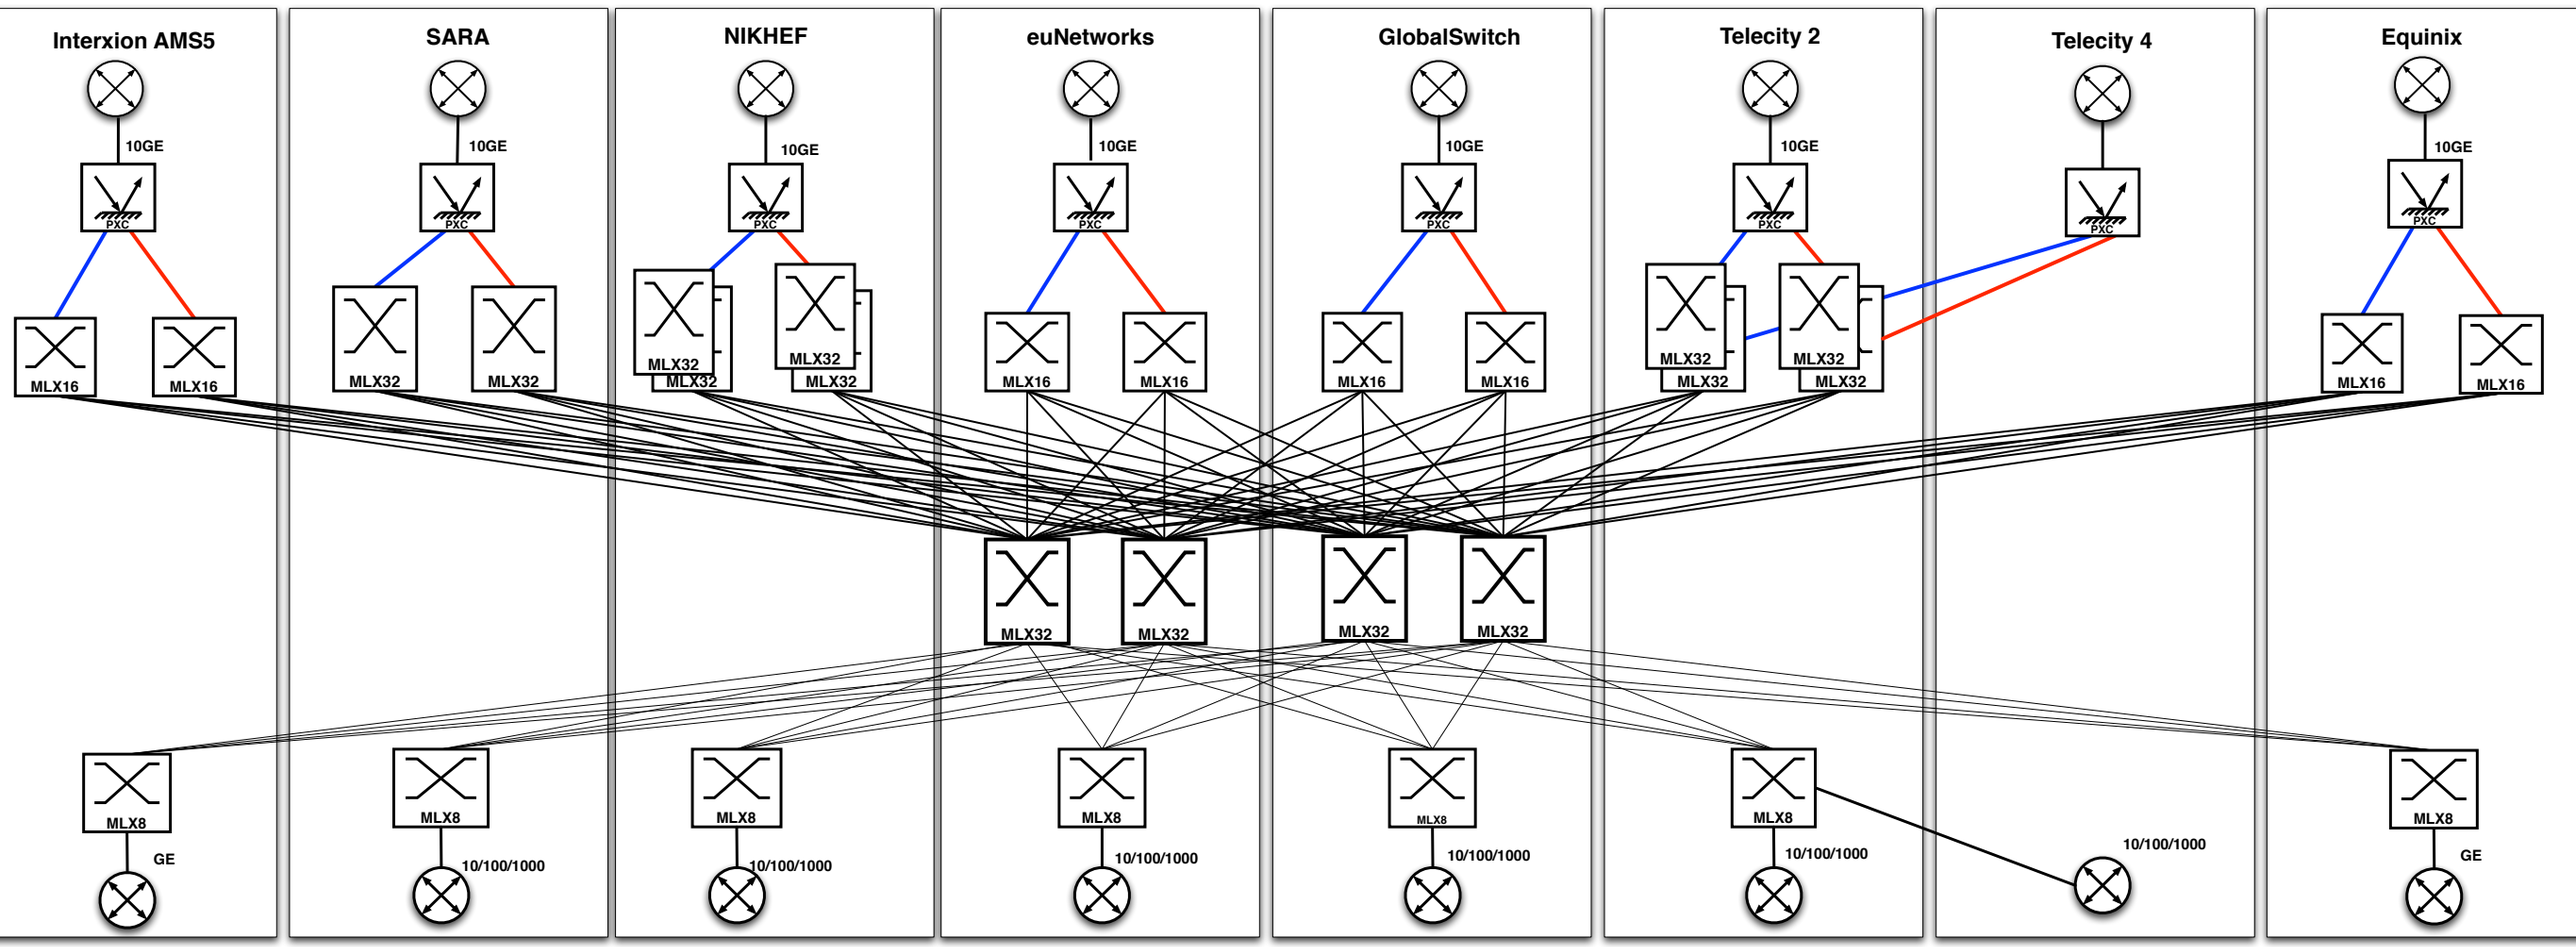

Jul 7, 2011, UCL

eBGP - Full Mesh

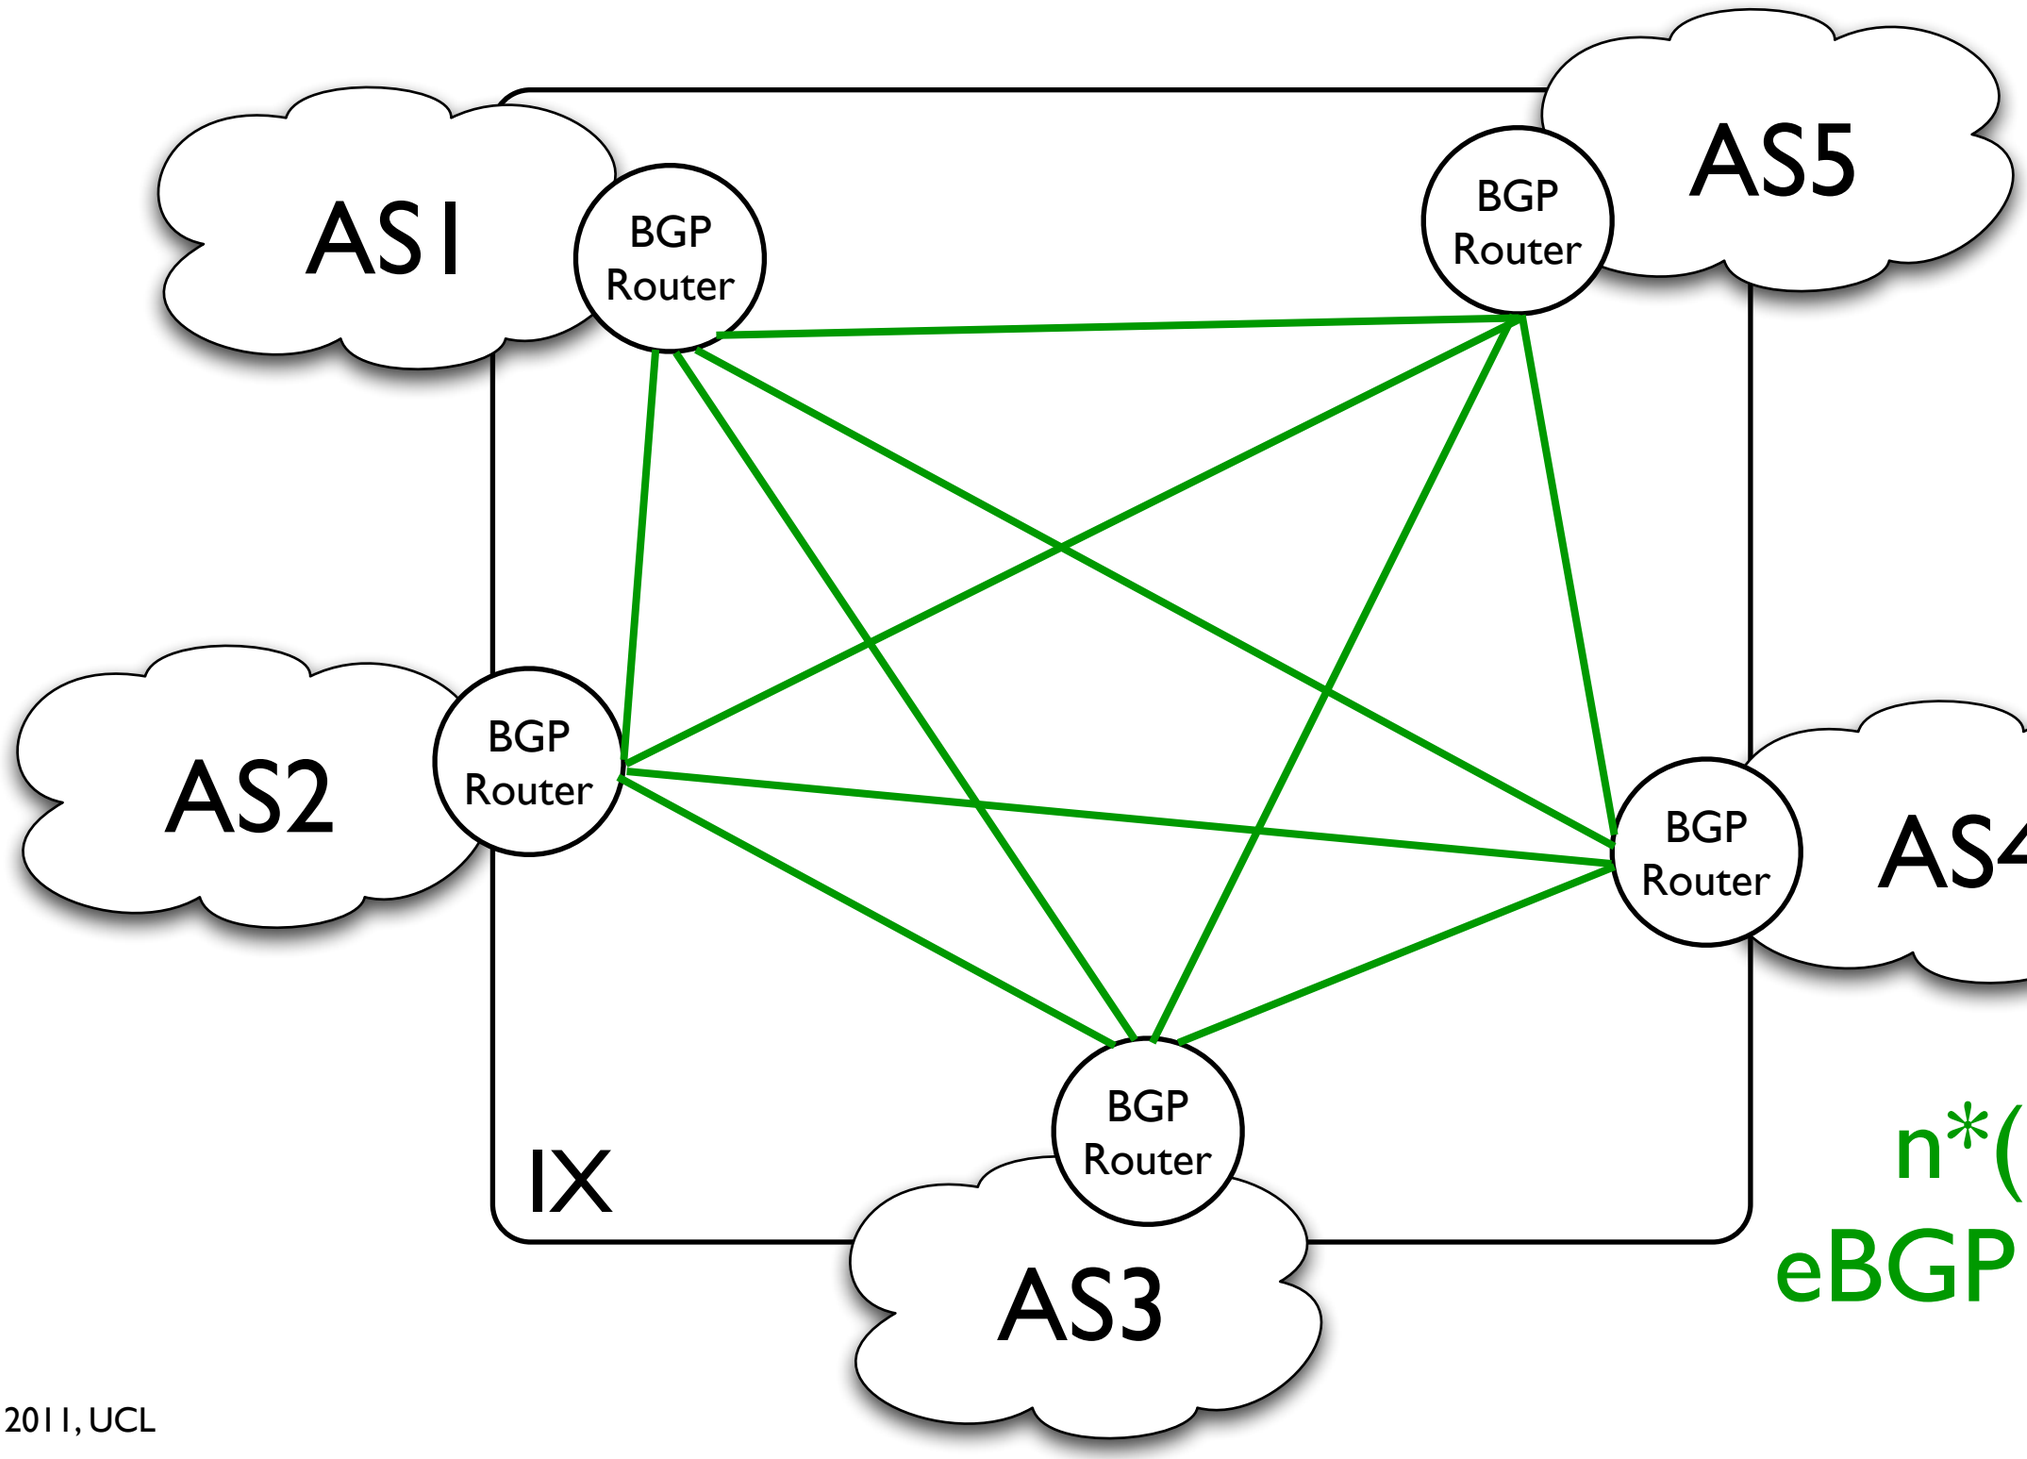

n*(n-1)/2
eBGP Sessions

eBGP - Full Mesh

- Administrative overhead
- Initial resources when connecting to IX

n*(n-1)/2
eBGP Sessions

eBGP - Full Mesh with RS

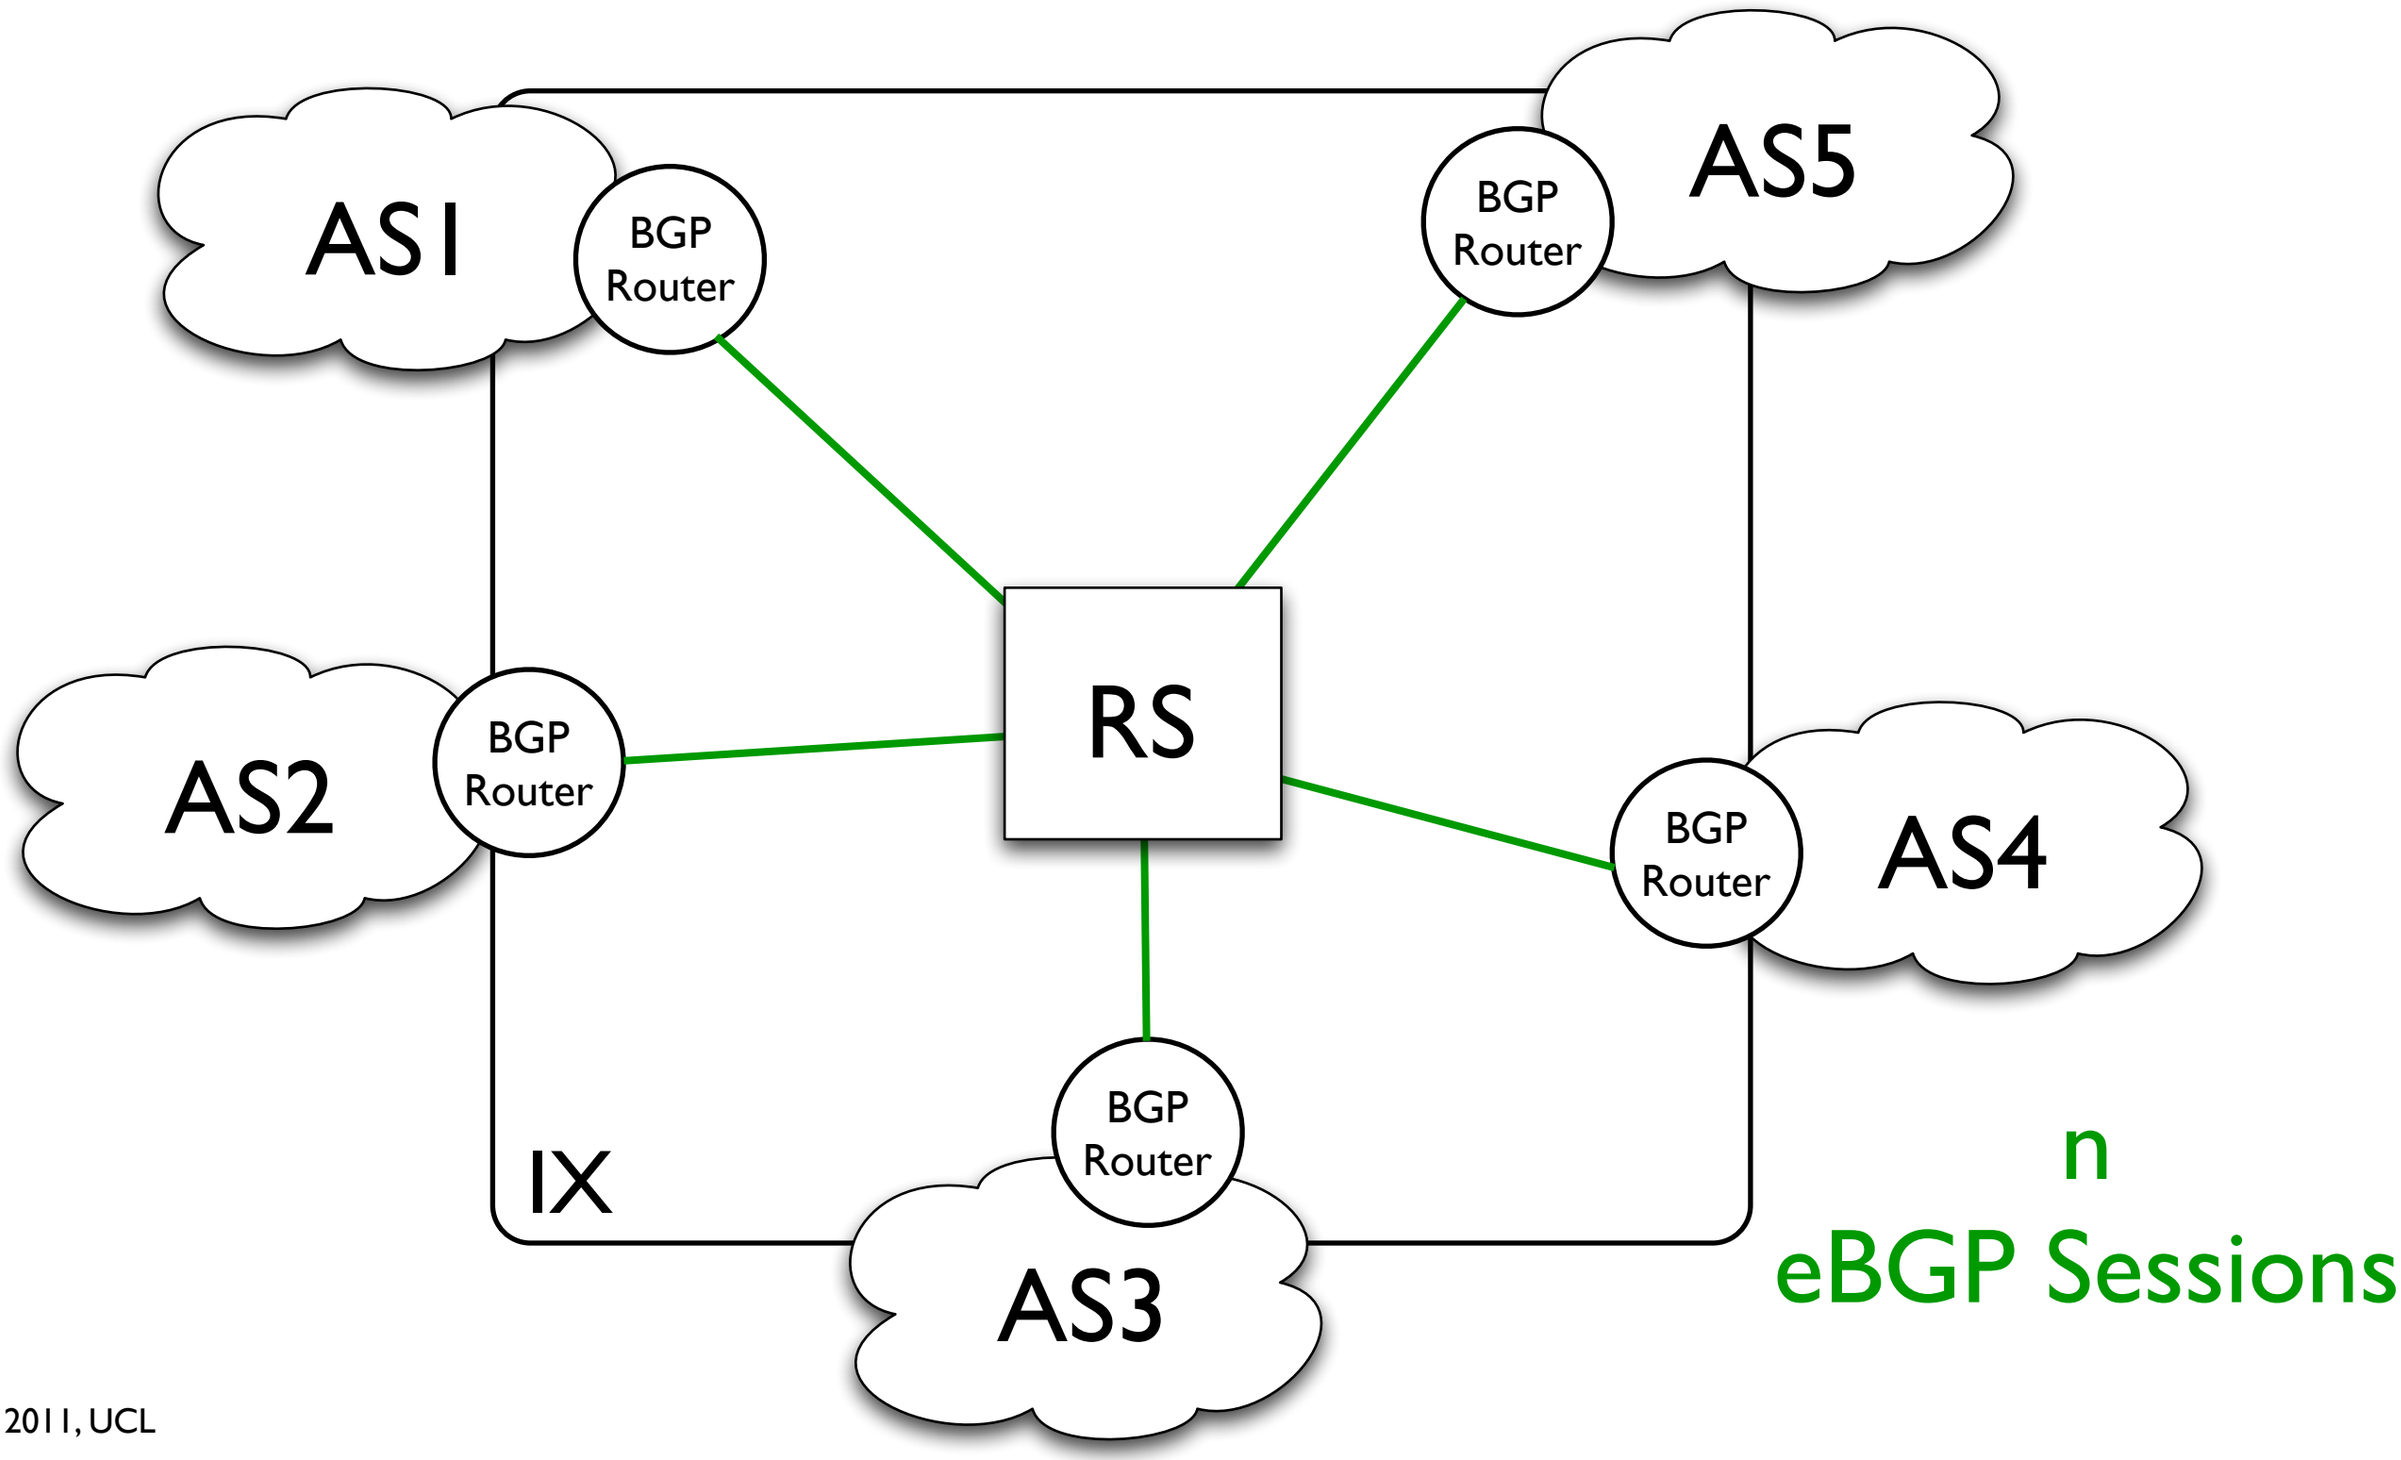

Route Server

- Receive routing information from all client routers
- Compute best paths
- Distribute routing information among all client routers

Route Server

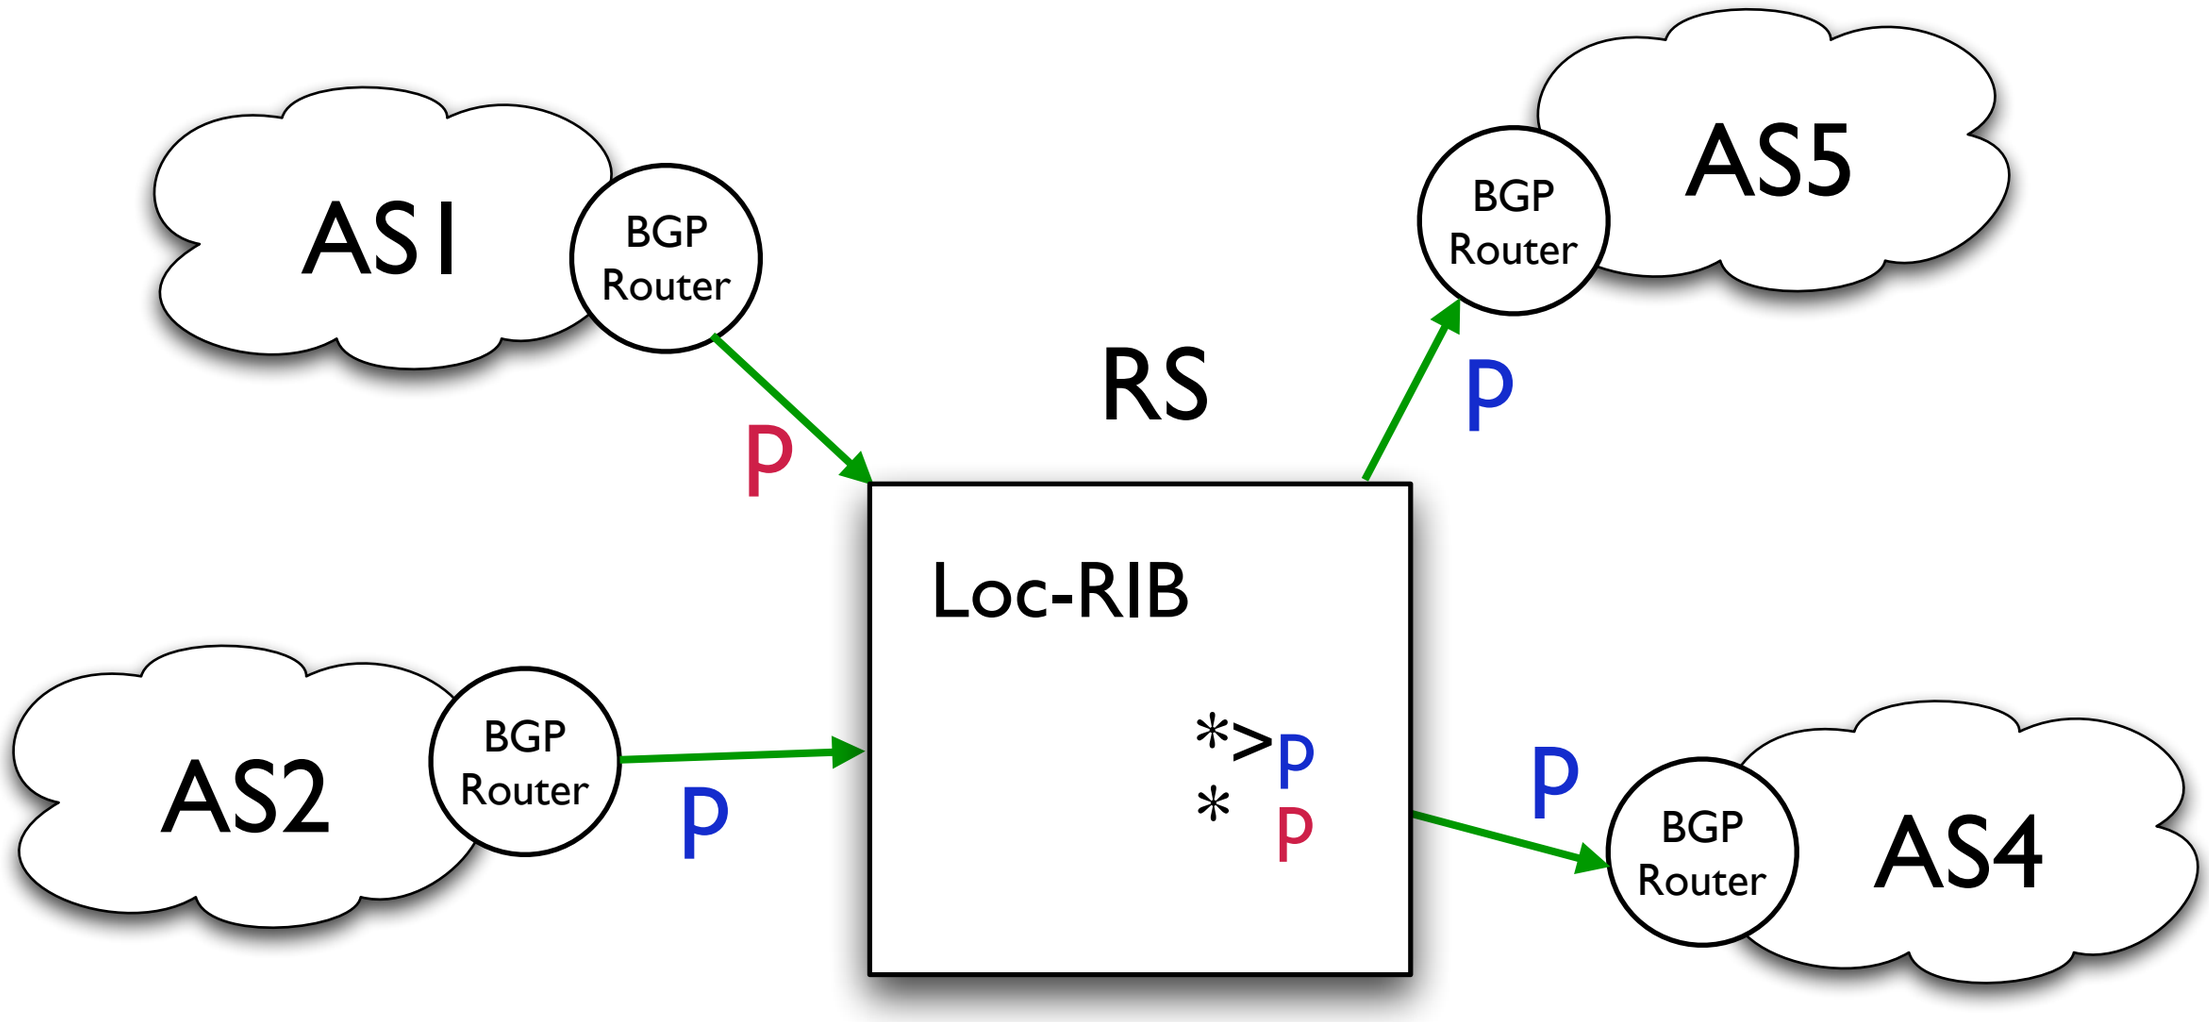

Route Server Implementations

- Open source implementations:
  - Quagga, OpenBGPD, BIRD
  - Implementation differences, no document to describe what a route server is
  - More implementations to come:
    - Cisco IOS, JunOS
  - draft-jasinska-ix-bgp-route-server-02

Attribute Transparency

- BGP RS implementation:
  - MUST NOT change NEXT_HOP
  - SHOULD NOT insert ASN into AS_PATH
  - SHOULD NOT modify, process or remove MULTI_EXIT_DISC
  - SHOULD NOT modify, process or remove Communities

eBGP - Partial Mesh

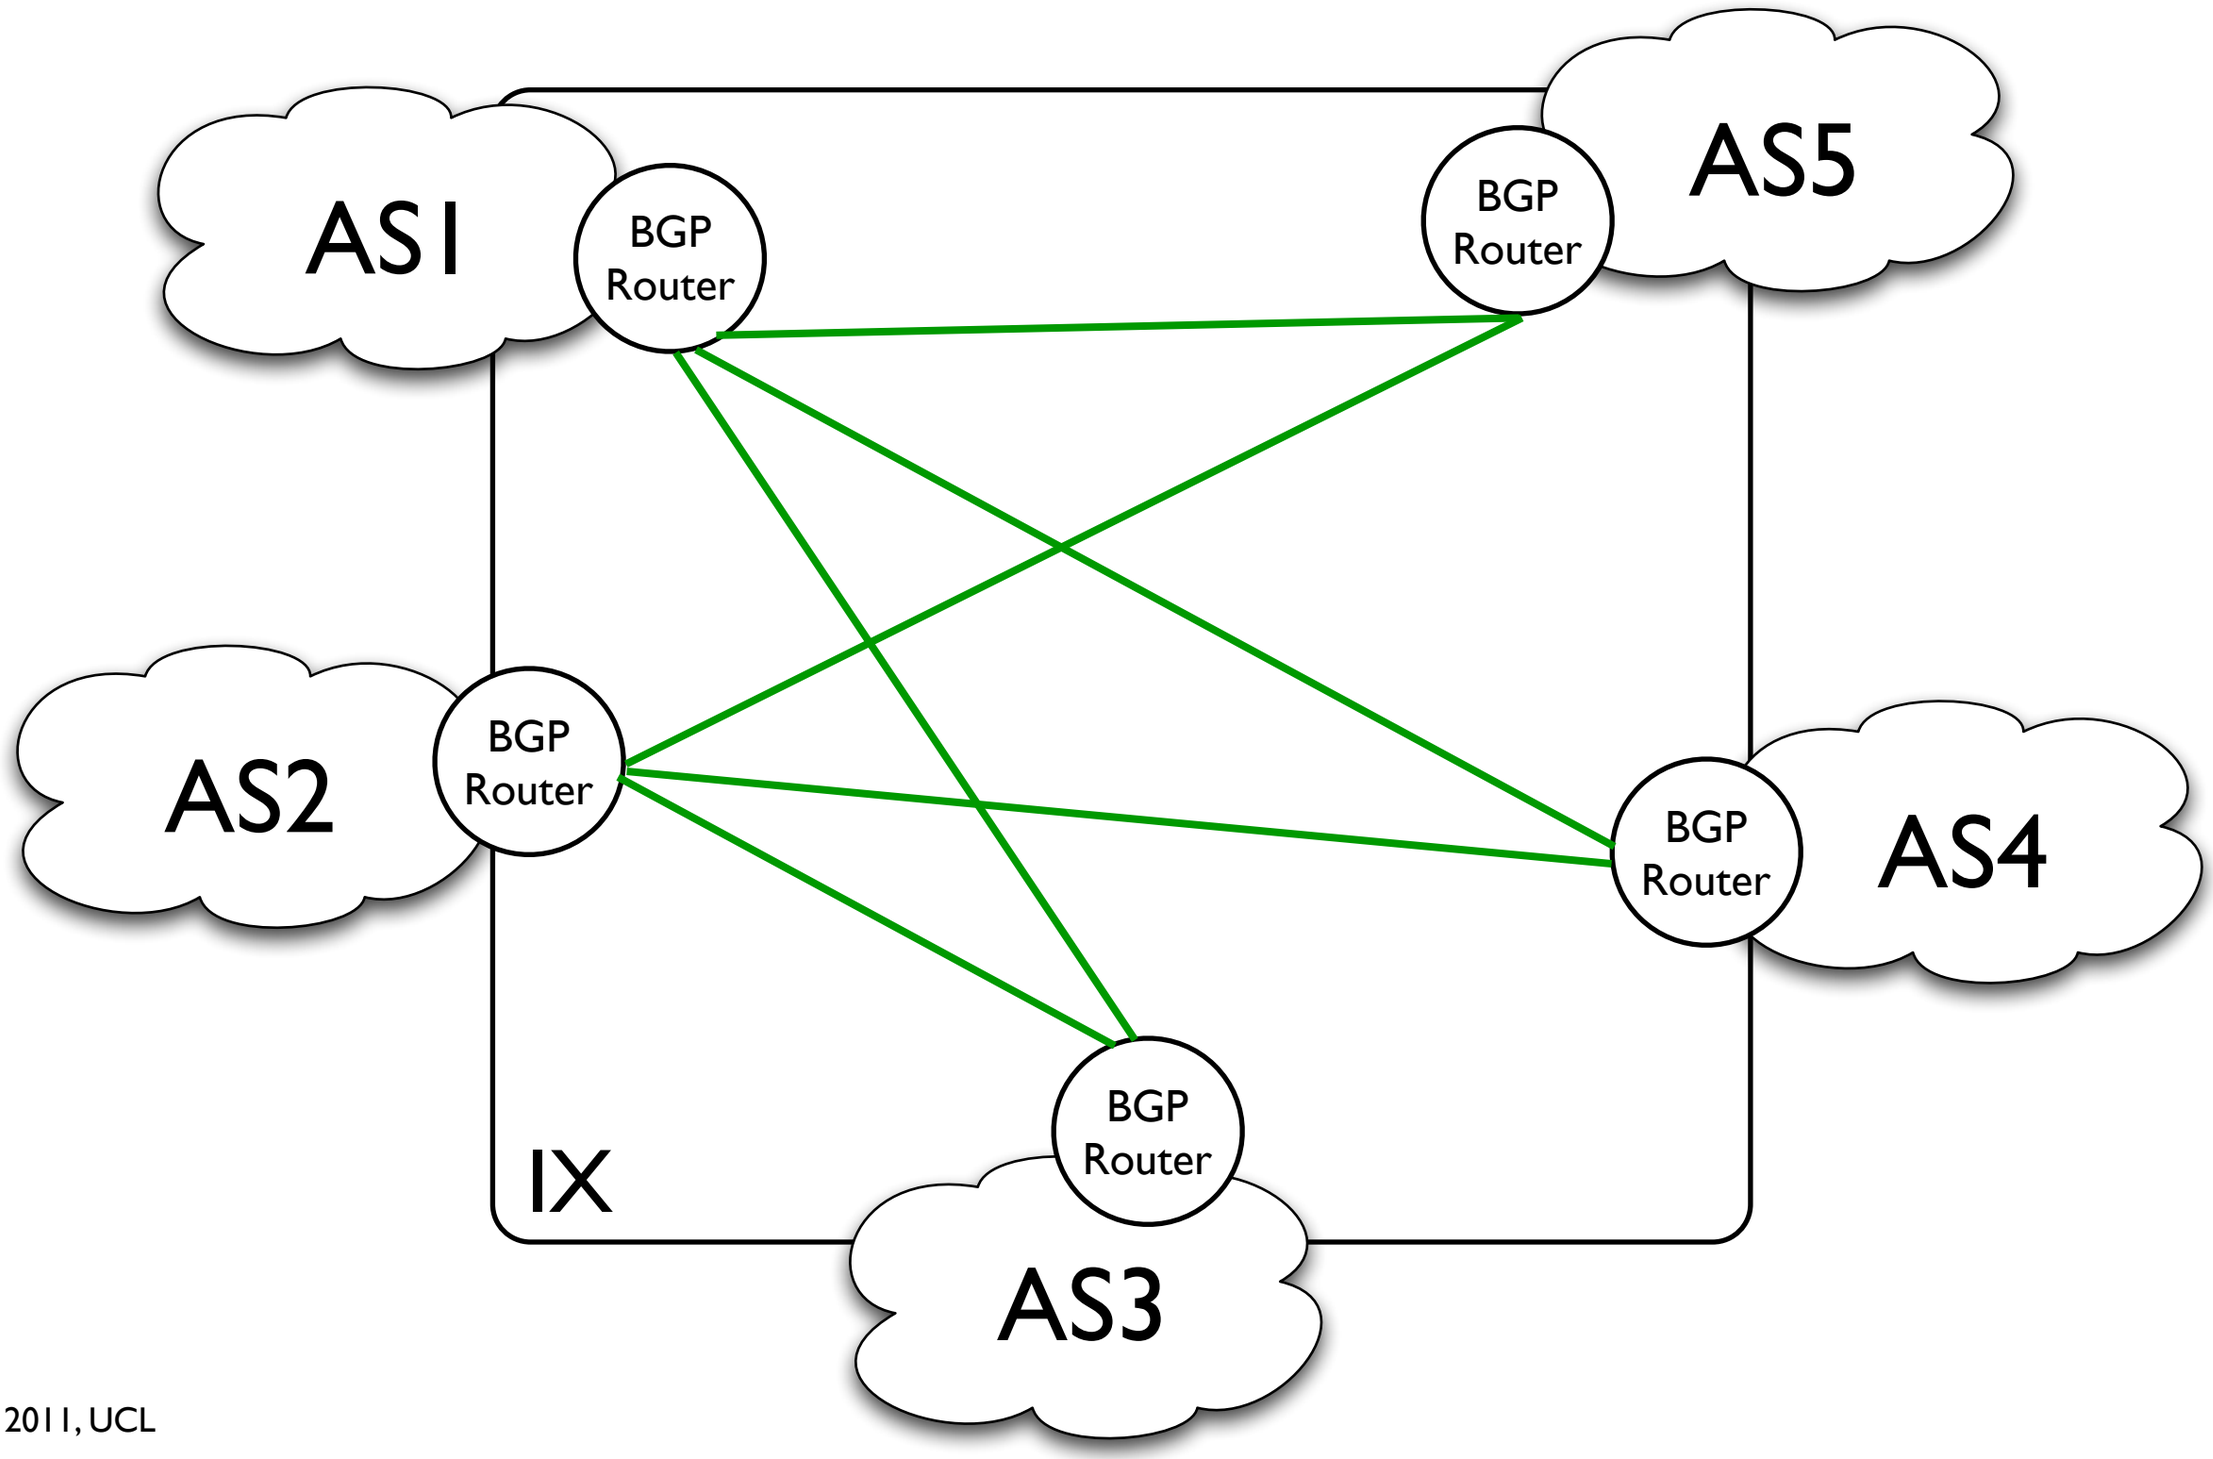

eBGP - Partial Mesh with RS

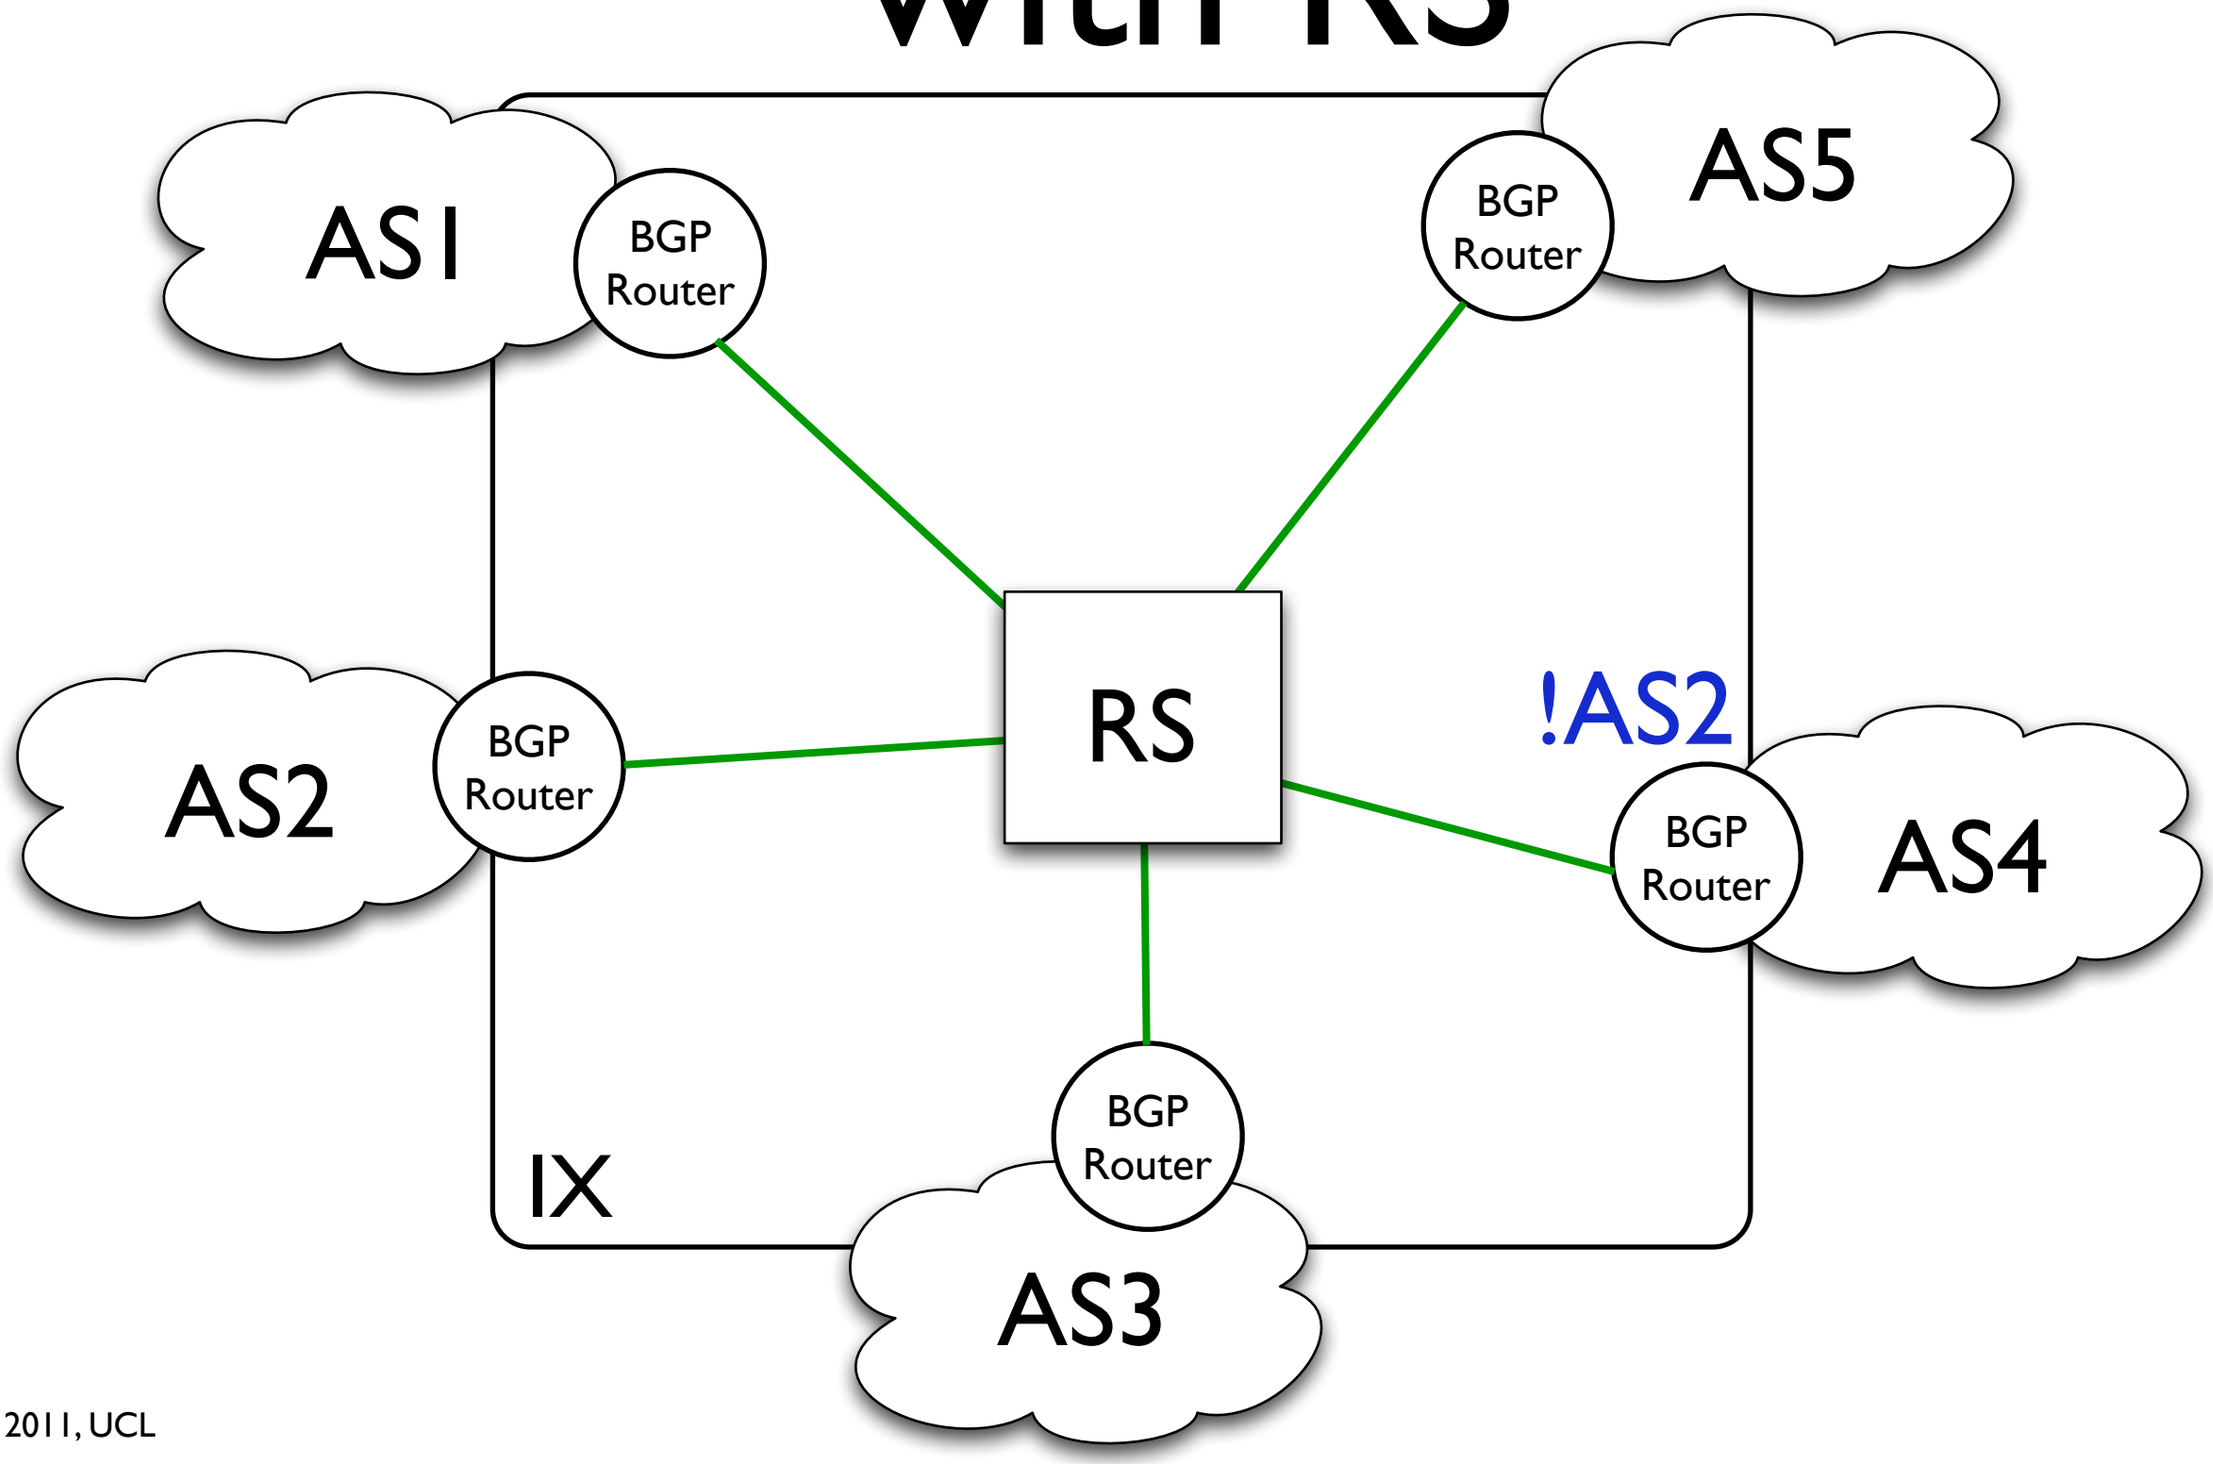

Prefix Hiding

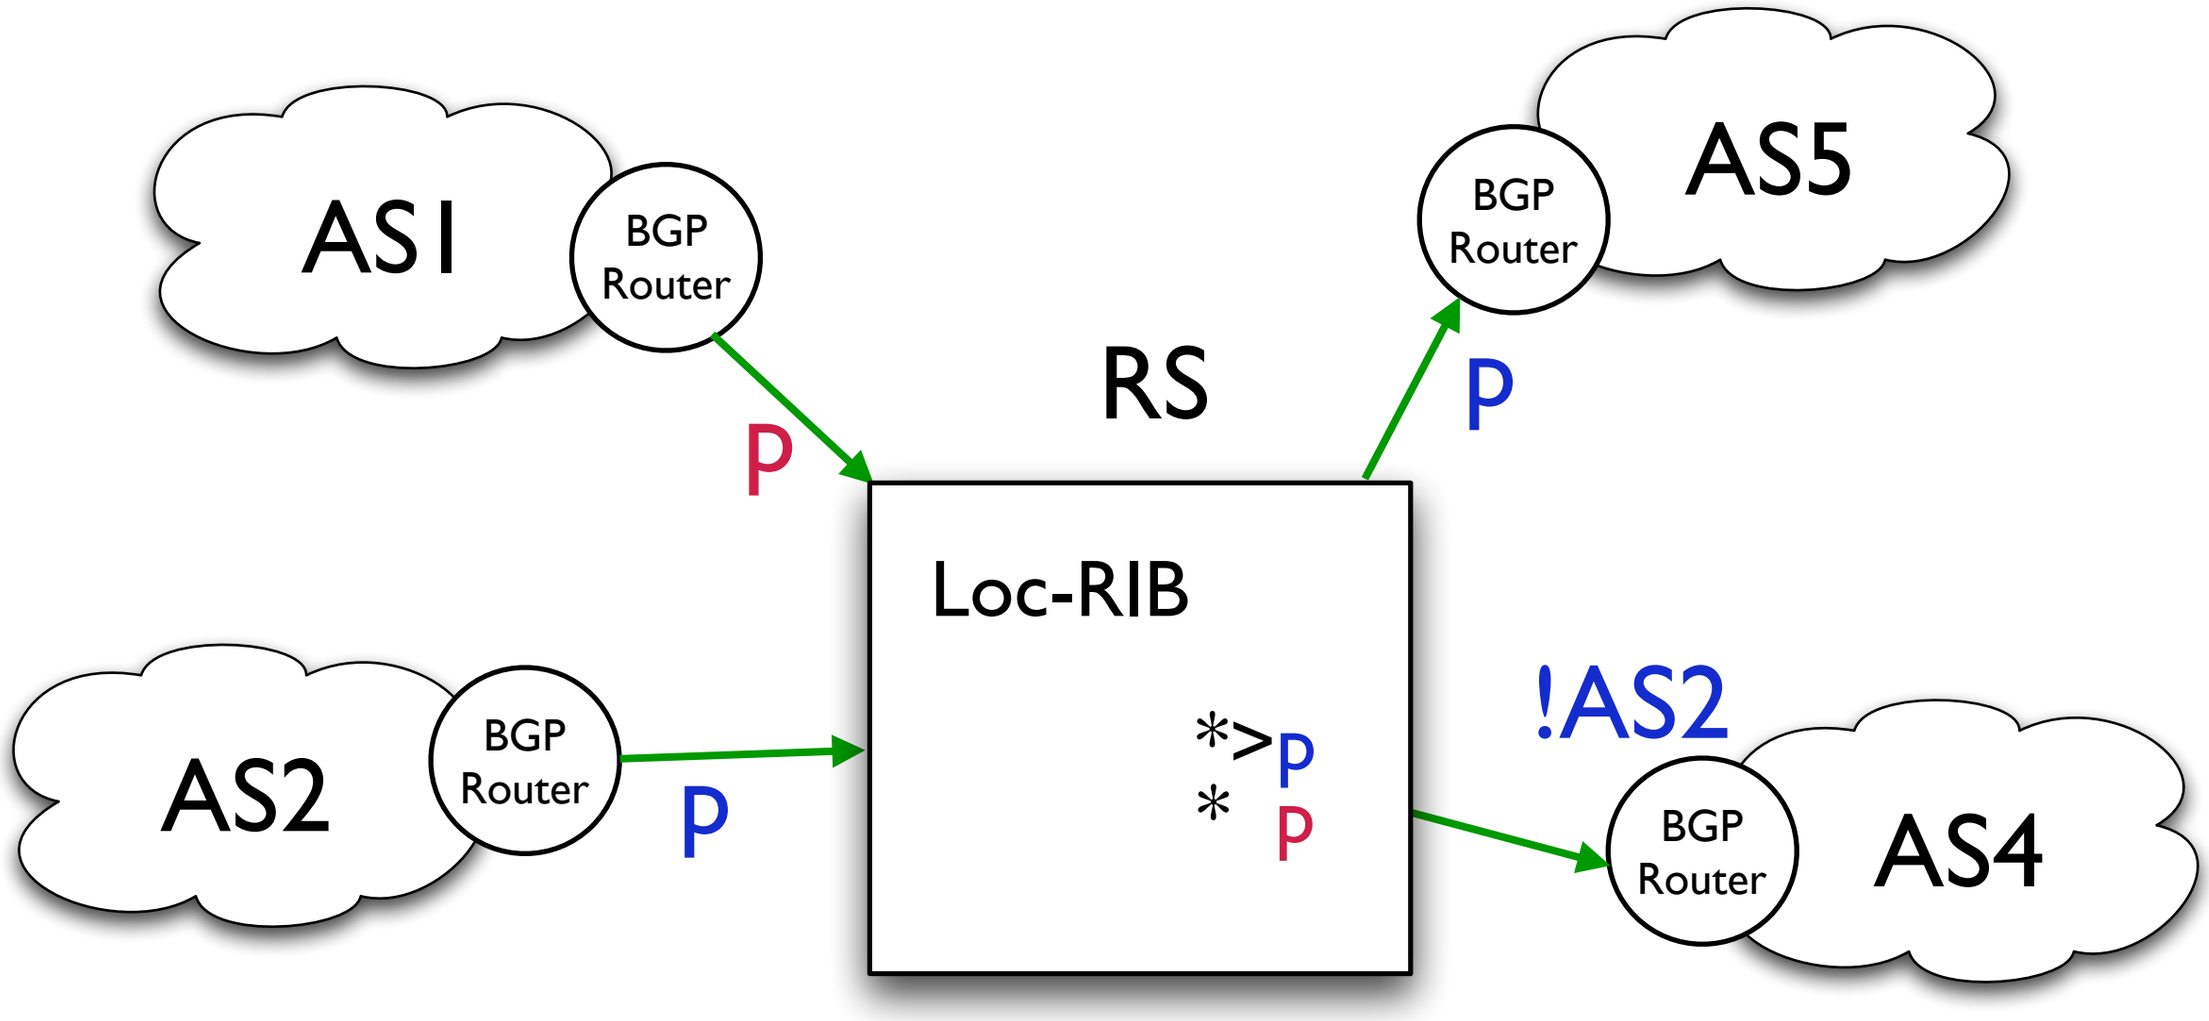

Multiple RIBs

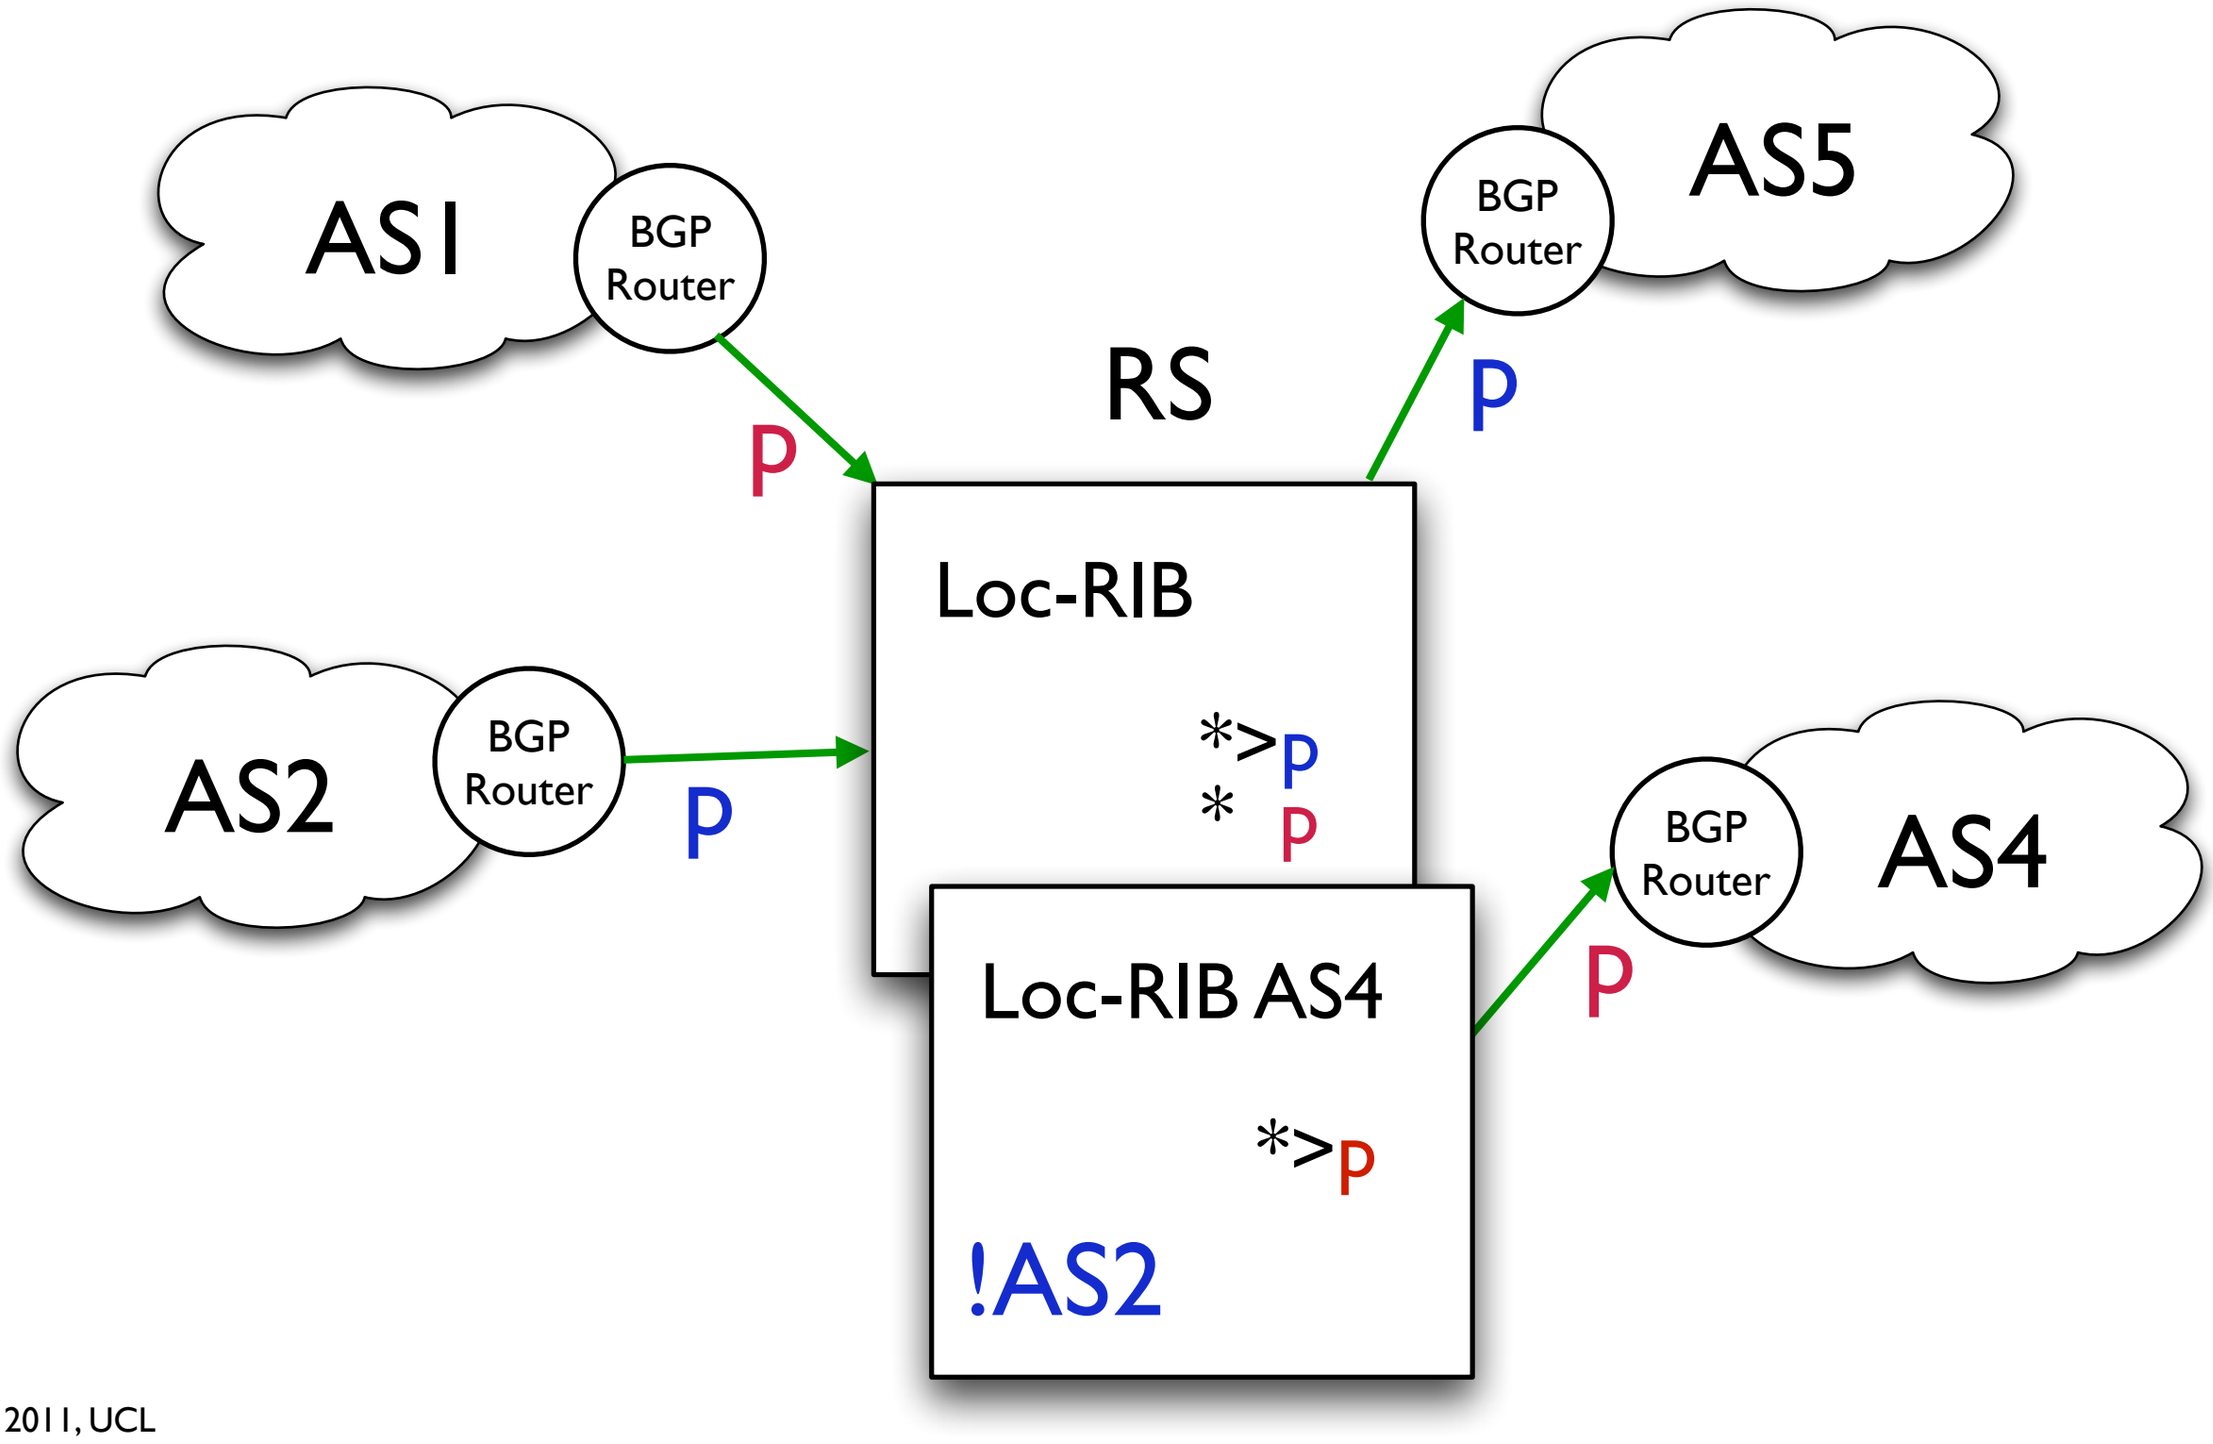

Add-paths

- Another approach to provide policy routing on a route server
- Add-paths on eBGP
- Second-best path sent to route server client in case policy applies (or third, or fourth, etc.)

Thanks!