SCION: Control Plane Overview

Adrian Perrig
Network Security Group, ETH Zürich

ETH zürich

SCION

SCION Control Plane Overview

- Control plane: How to find and distribute end-to-end paths
  [Chapter 2.1, Chapter 7]
  - Path exploration
  - Path registration
  - Path lookup
- Security and reliability aspects
- Service anycast
- SCION control message protocol (SCMP)

Reminder: SCION Isolation Domain (ISD)

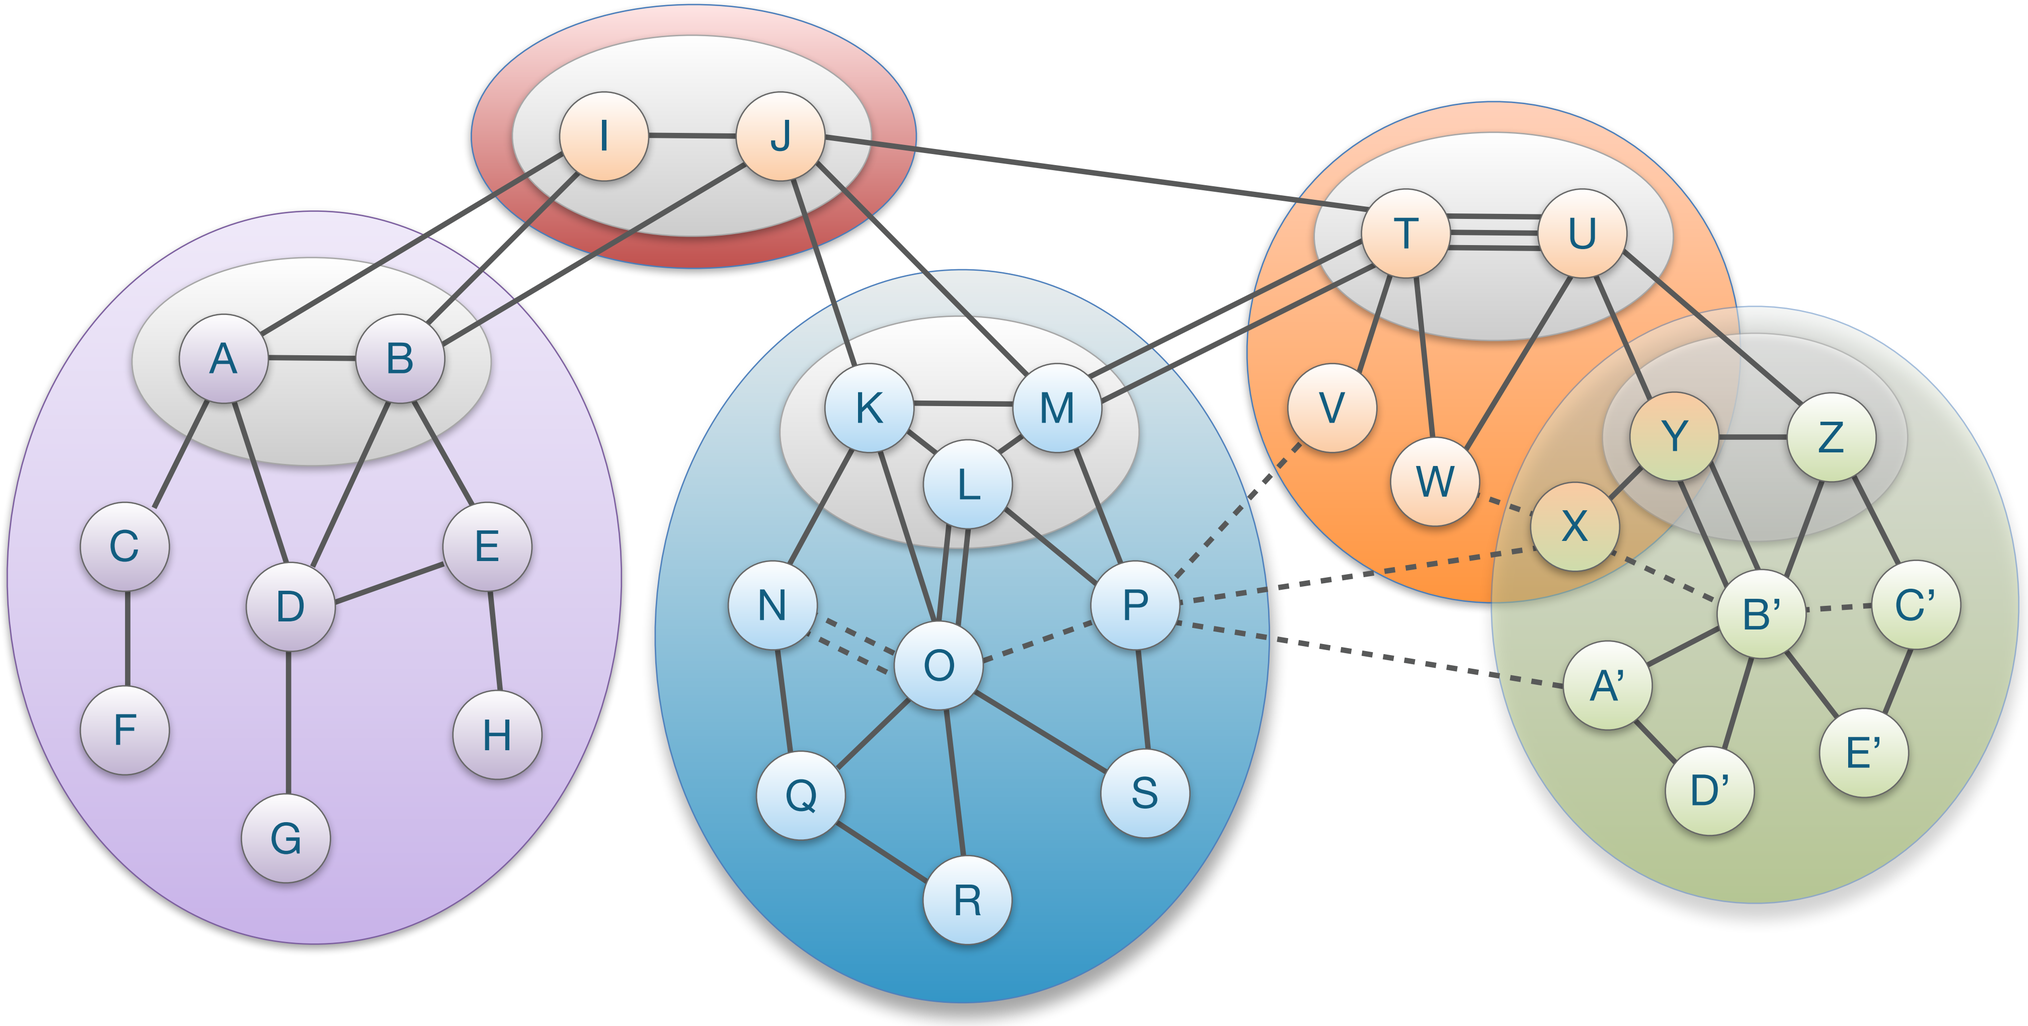

ETH zürich

SCION

Intra-ISD Path Exploration: Beaconing

- Core ASes K, L, M initiate Path-segment Construction Beacons (PCBs), or “beacons”
- PCBs traverse ISD as a policy-constrained multi-path flood
- Each AS receives multiple PCBs representing path segments to a core AS
- Each PCB can be used as an up-path segment to communicate with core AS

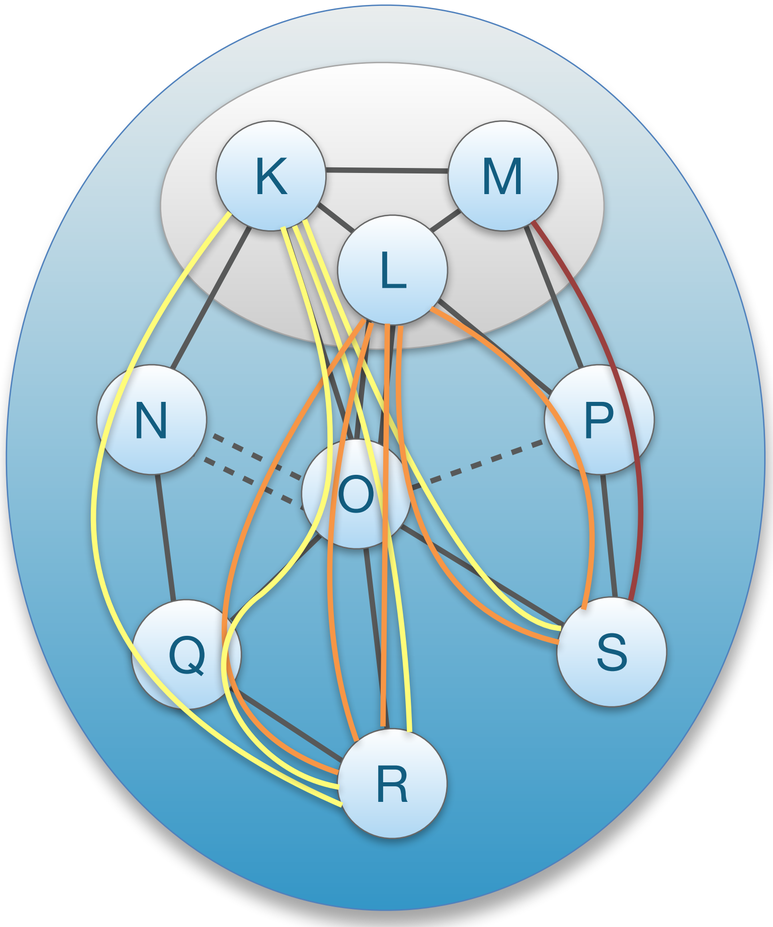

Beaconing in More Detail

- Each AS deploys one or multiple beacon servers
- PCBs are sent via a SCION service anycast packet
- SCION border routers receive PCB and select one beacon server to forward it to
- Beacon servers coordinate to re-send PCBs periodically to downstream ASes
  - Currently every 5 seconds, PCBs are selected and forwarded

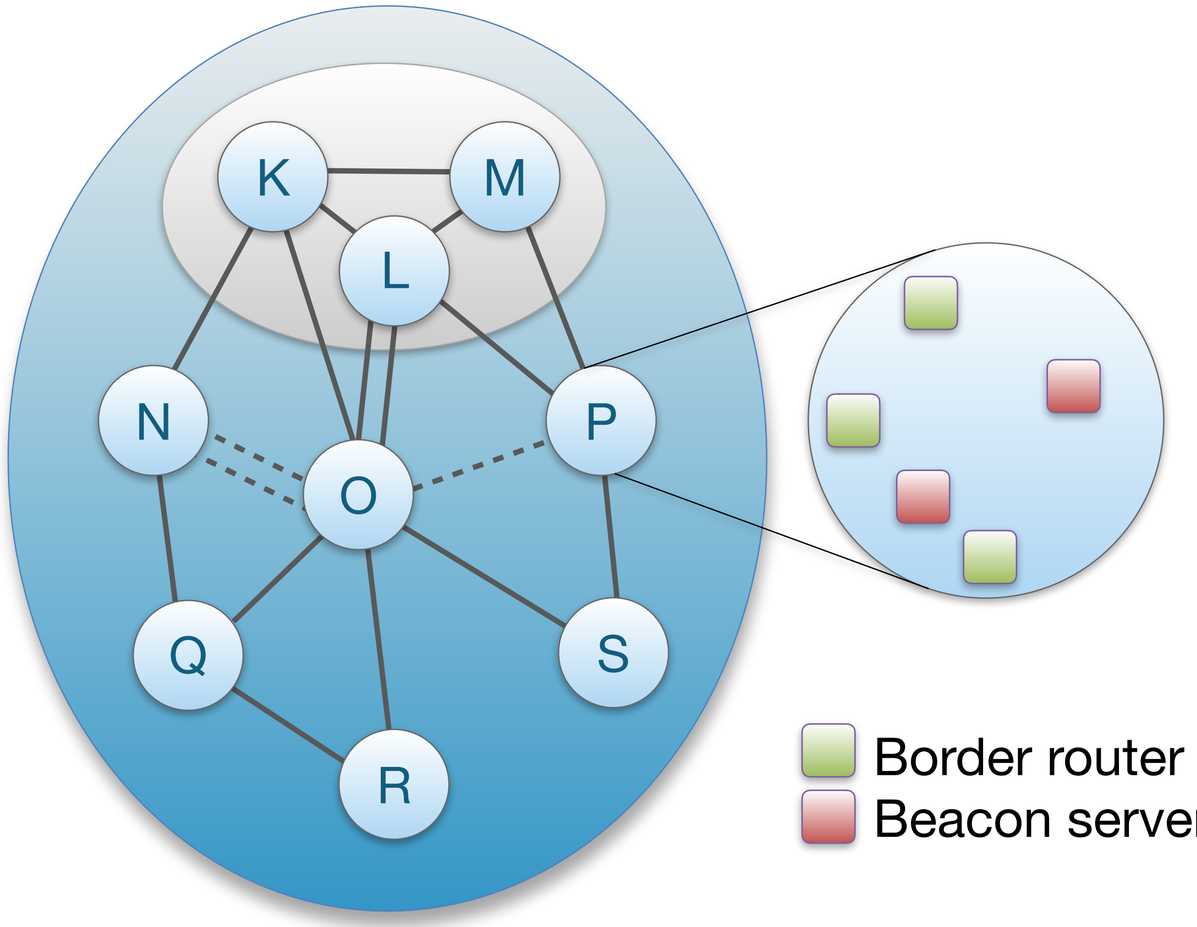

PCB Contents

- A PCB contains an info field with:
  - PCB creation time
- Each AS on path adds:
  - AS name
  - Hop field for data-plane forwarding
    - Link identifiers
    - Expiration time
    - Message Authentication Code (MAC)
  - AS signature

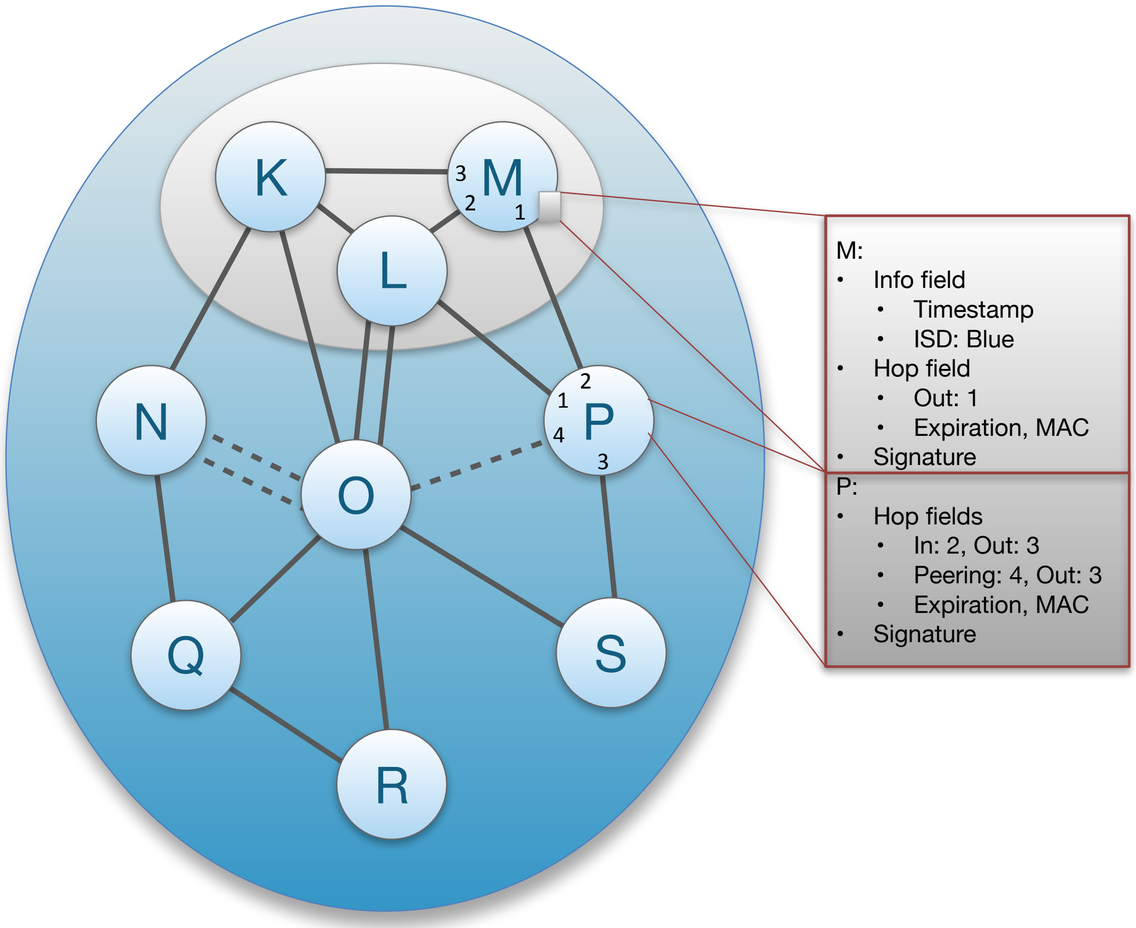

Path Server Infrastructure

- Path servers offer lookup service:
  - ISD, AS → down-path segments, core-path segments
  - Local up-path segment request → up-path segments to core ASes
- Core ASes operate core path server infrastructure
  - Consistent, replicated store of down-path segments and core-path segments
- Each non-core AS runs local path servers
  - Serves up-path segments to local clients
  - Resolves and caches response of remote AS lookups

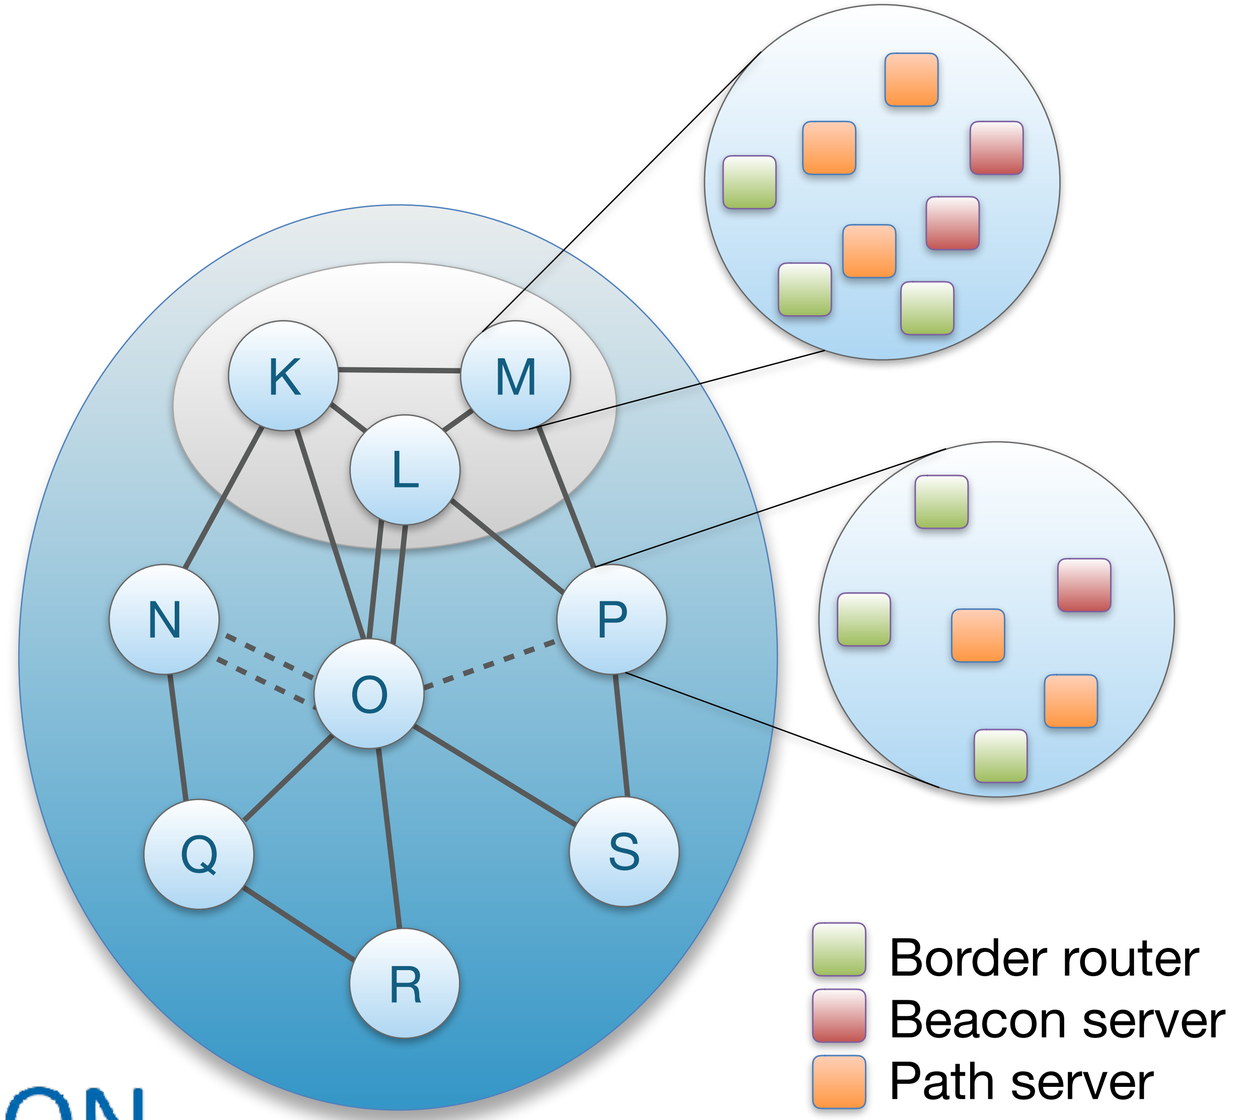

Down-Path Segment Registration

- Each AS' beacon servers select path segments that they wants to announce as down-path segments for others to use to communicate with AS
- Beacon servers upload the selected down-path segments to path servers in core ASes

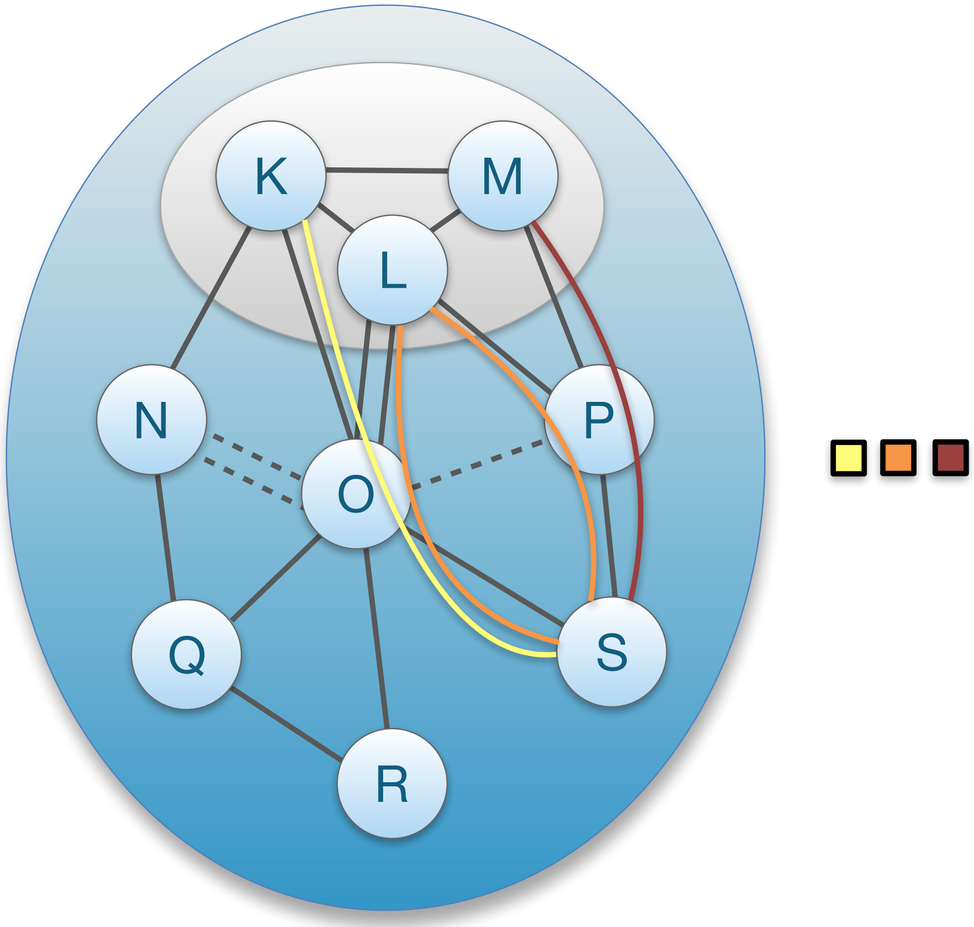

Up-Path Segment Registration

- Each AS' beacon servers select path segments that they wants to announce as up-path segments for local hosts to communicate with other AS
- Beacon servers send the selected up-path segments to local path servers

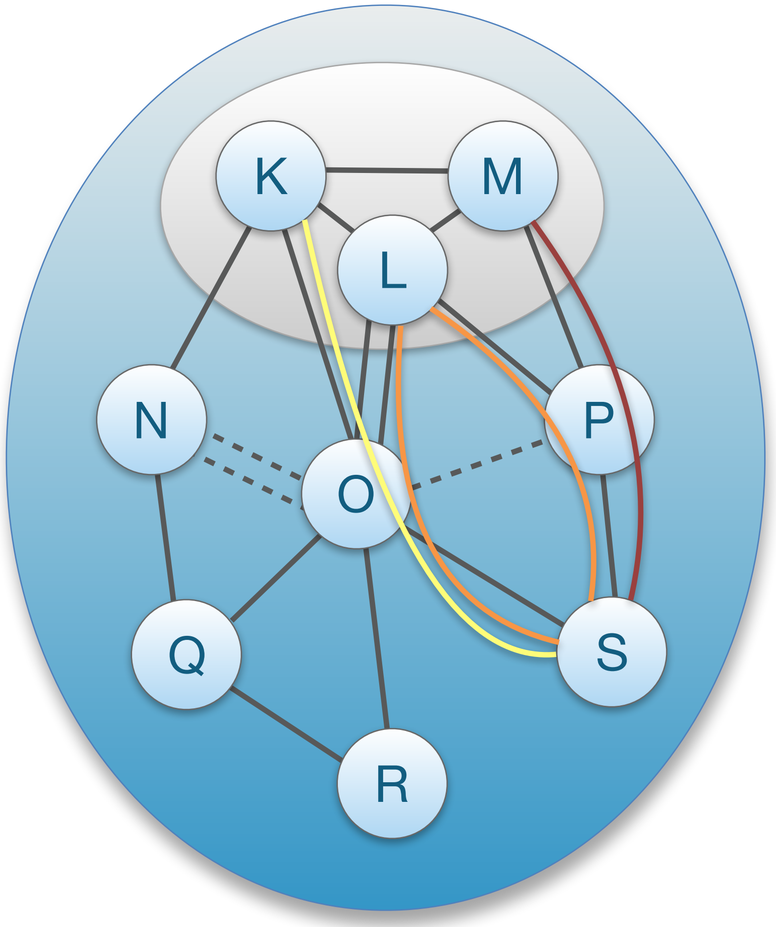

Core Beaconing for Inter-ISD Path Exploration

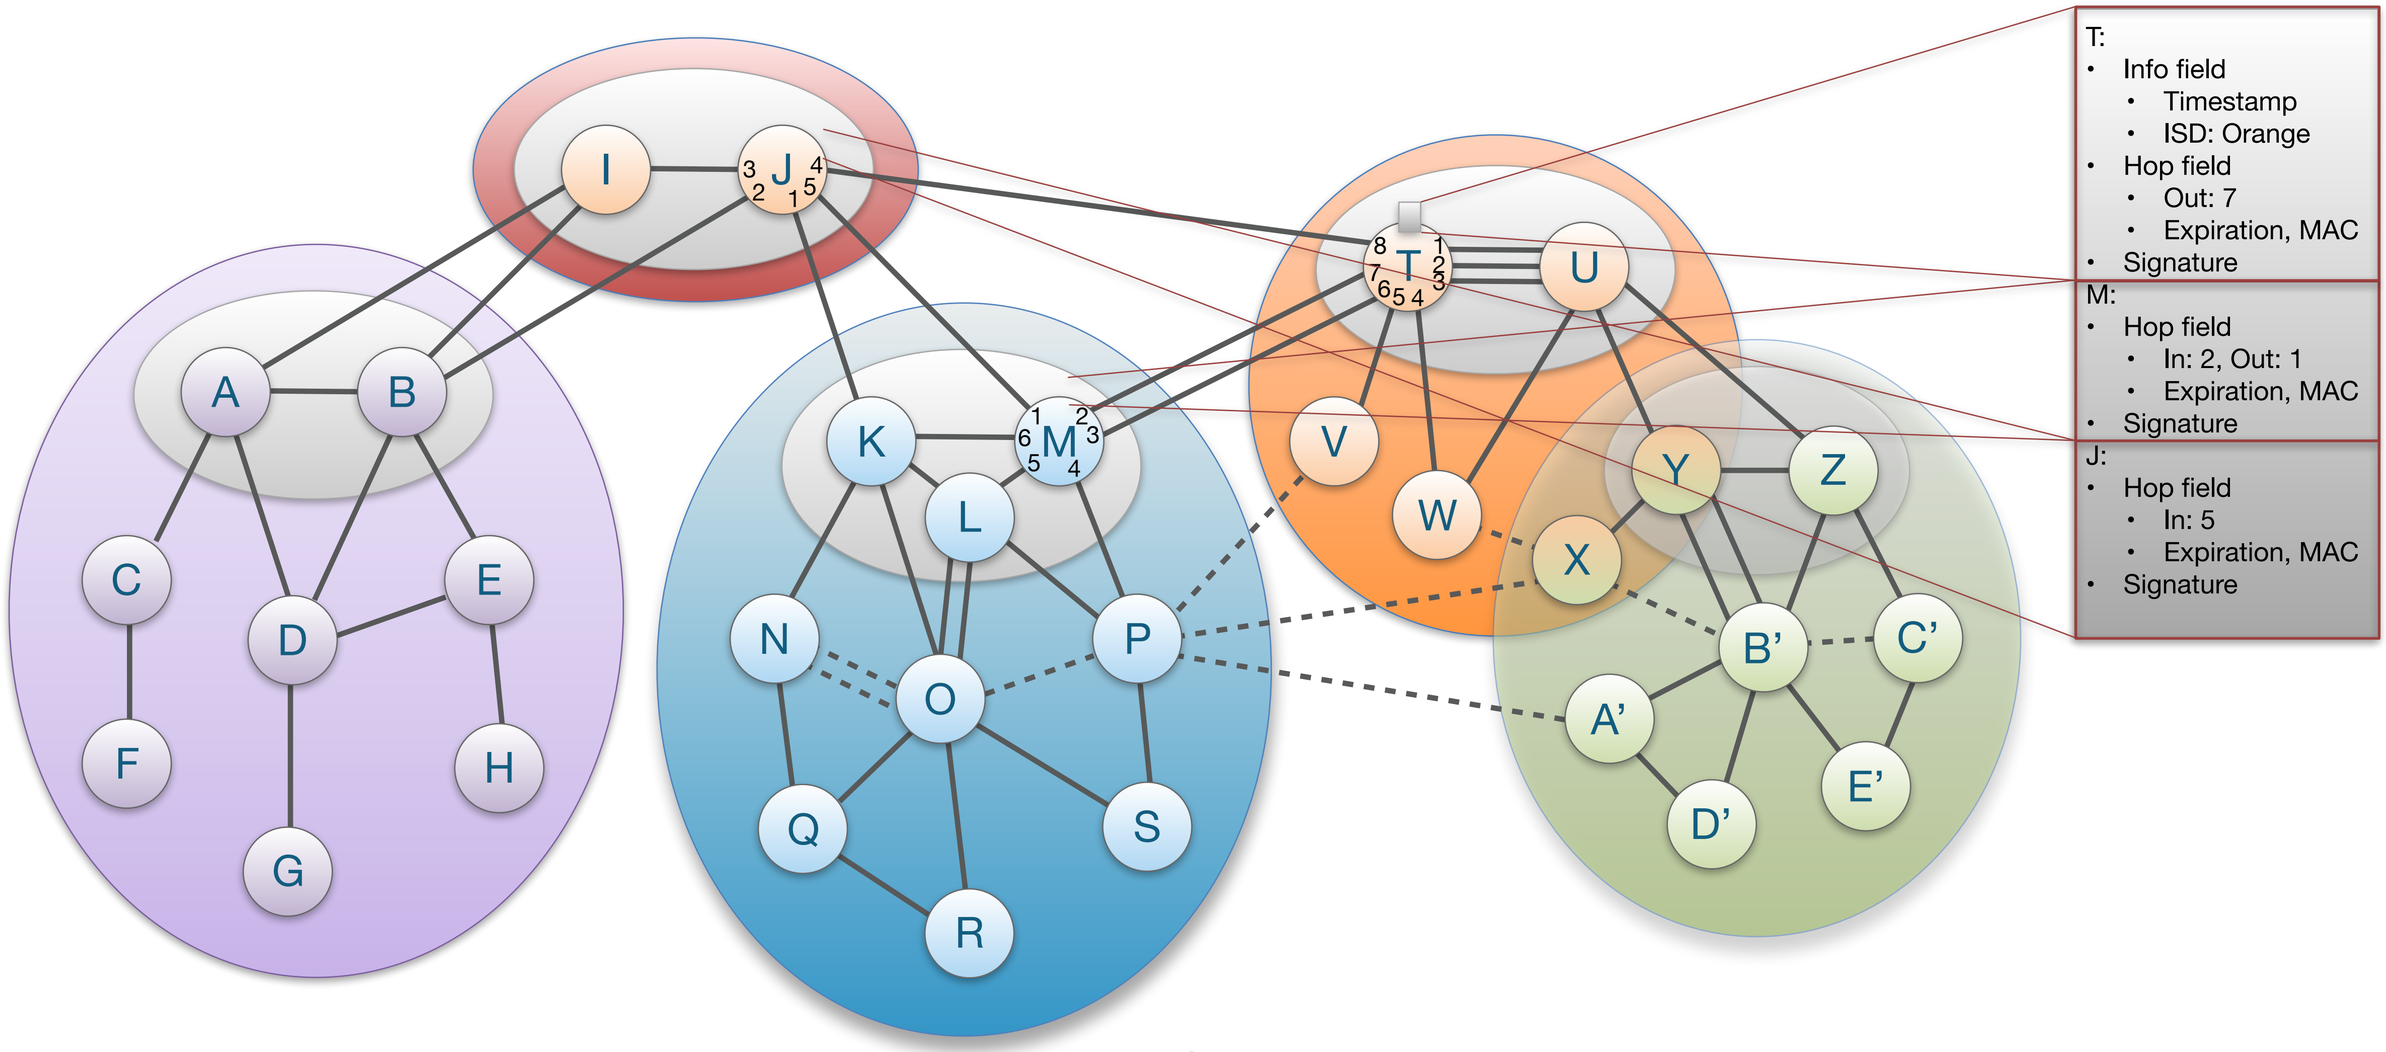

ETH zürich

SCION

Inter-ISD Path Exploration: Sample Core Paths from AS T

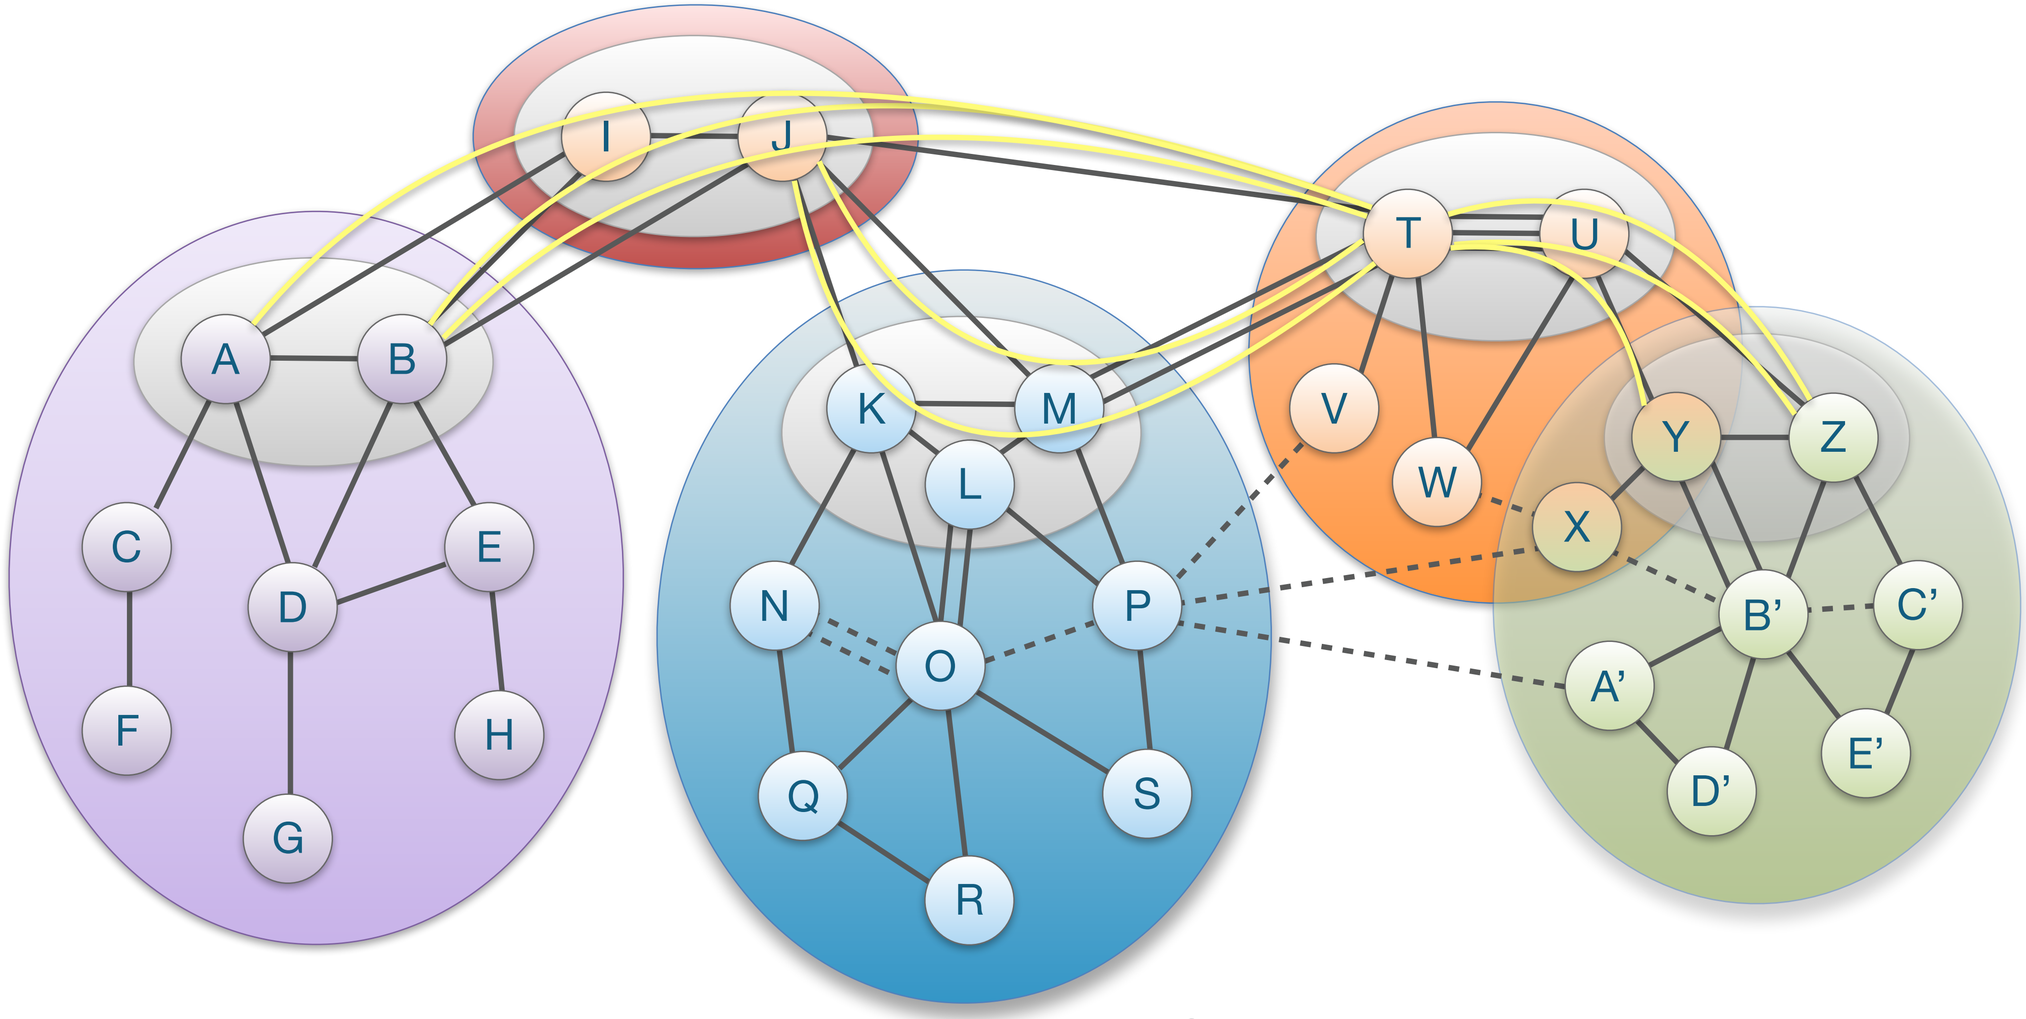

ETH zürich

SCION

Path Lookup: Local ISD

- Client requests path segments to <ISD, AS> from local path server
- If down-path segments are not locally cached, local path server send request to core path server
- Local path server replies
  - Up-path segments to local ISD core ASes
  - Down-path segments to <ISD, AS>
  - Core-path segments as needed to connect up-path and down-path segments

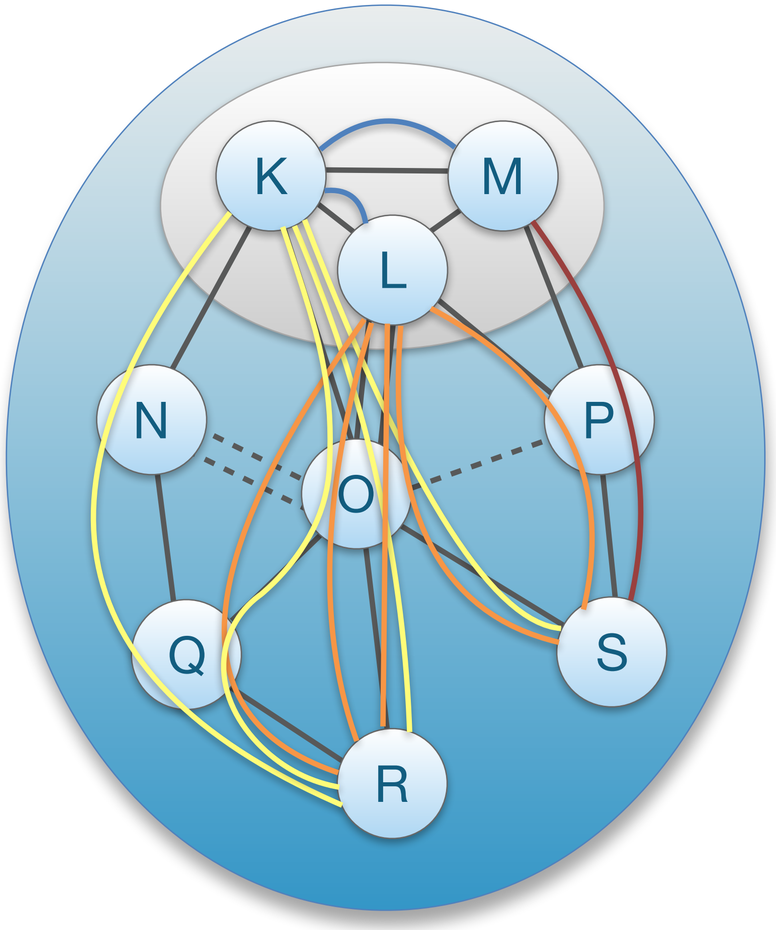

Path Lookup: Remote ISD

- Host contacts local path server requesting <ISD, AS>
- If path segments are not cached, local path server will contact core path server
- If core path server does not have path segments cached, it will contact remote core path server
- Finally, host receives up-, core-, and down-segments

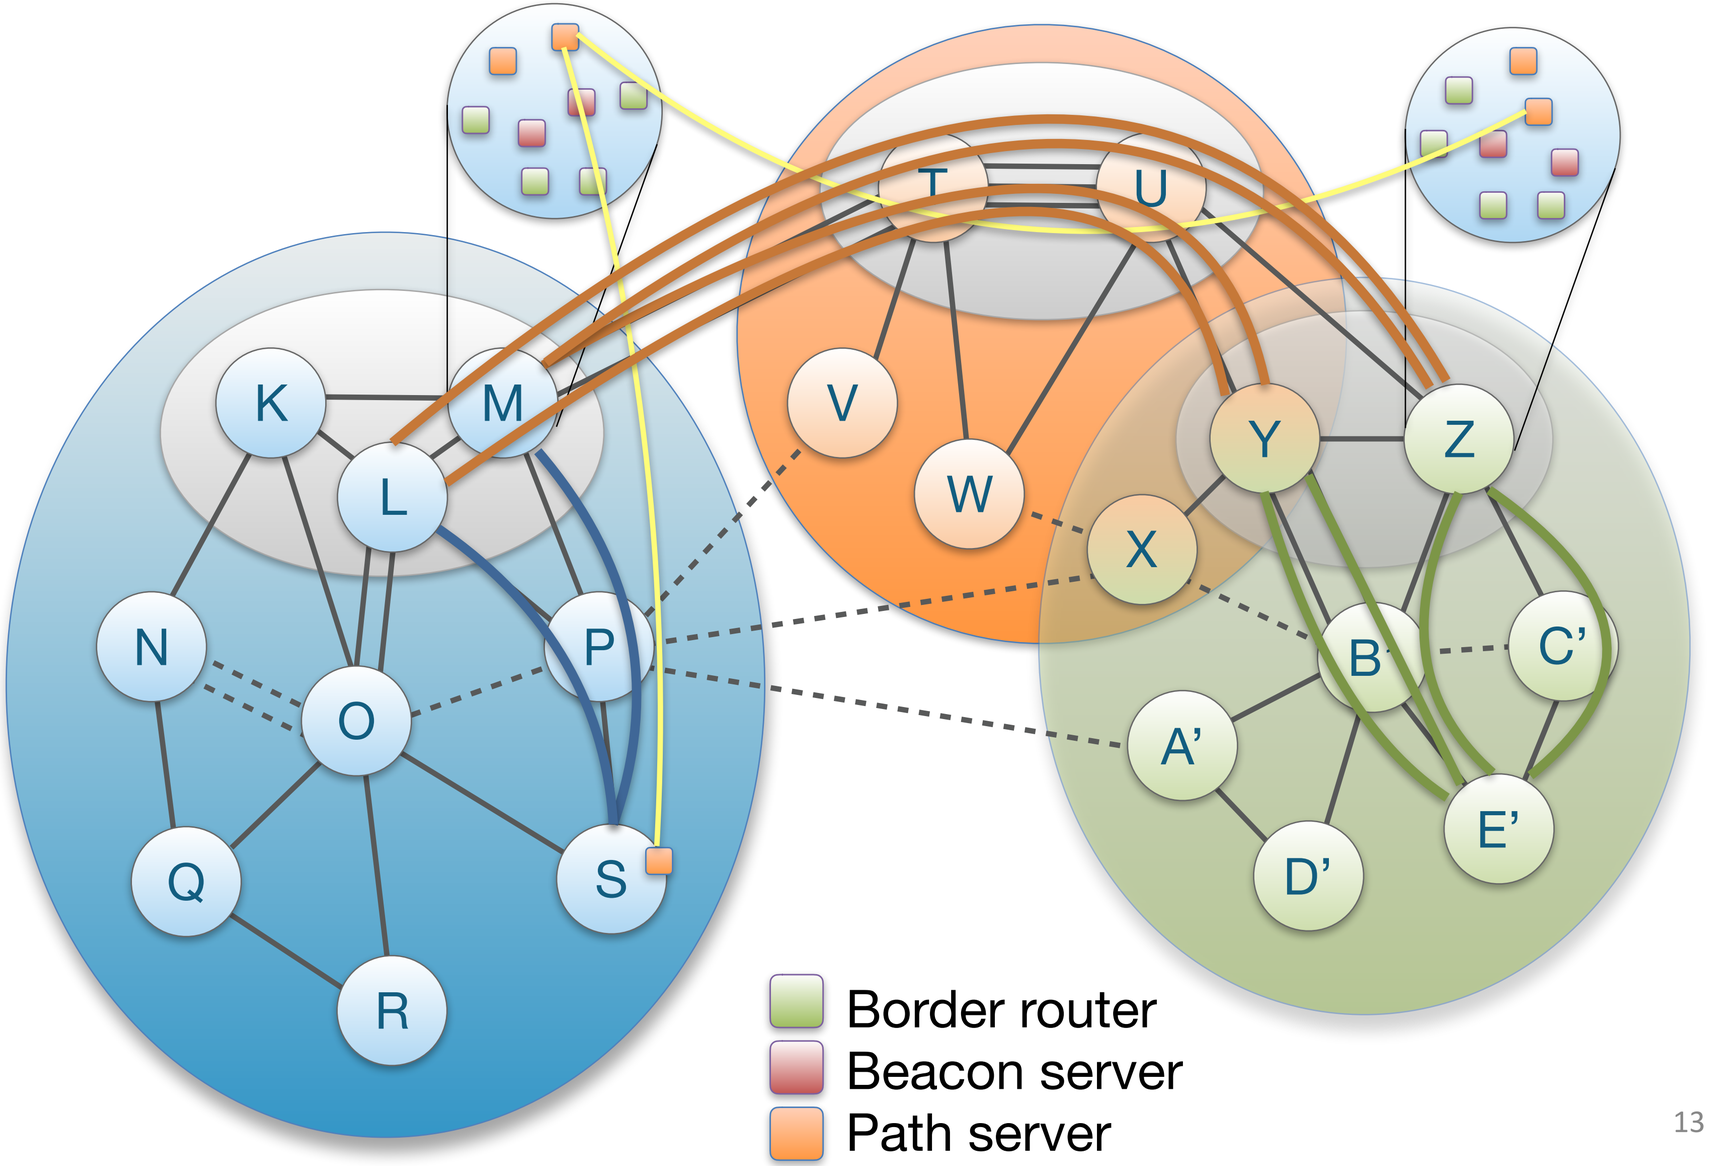

How to Secure PCB Dissemination

- Assumptions
  - Each AS has certificate: $\{AS, K_{AS}, expiration\}_{KcoreAS}$
    - Talks on SCION PKI and control-plane PKI provide more detail on how this works
  - Beacon servers know relevant AS certificates
- Each PCB is signed by core AS that issues it
- Each AS that resends PCB signs updated PCB
- Note: data-plane information (hop fields) are protected with efficient Message Authentication Code

Failed Interface Detection

- Border routers send periodic keep-alive message to neighboring border routers
- Received keep-alive messages are disseminated to all internal beacon server instances
- After a threshold number of keep-alive messages are missing, link is declared inactive

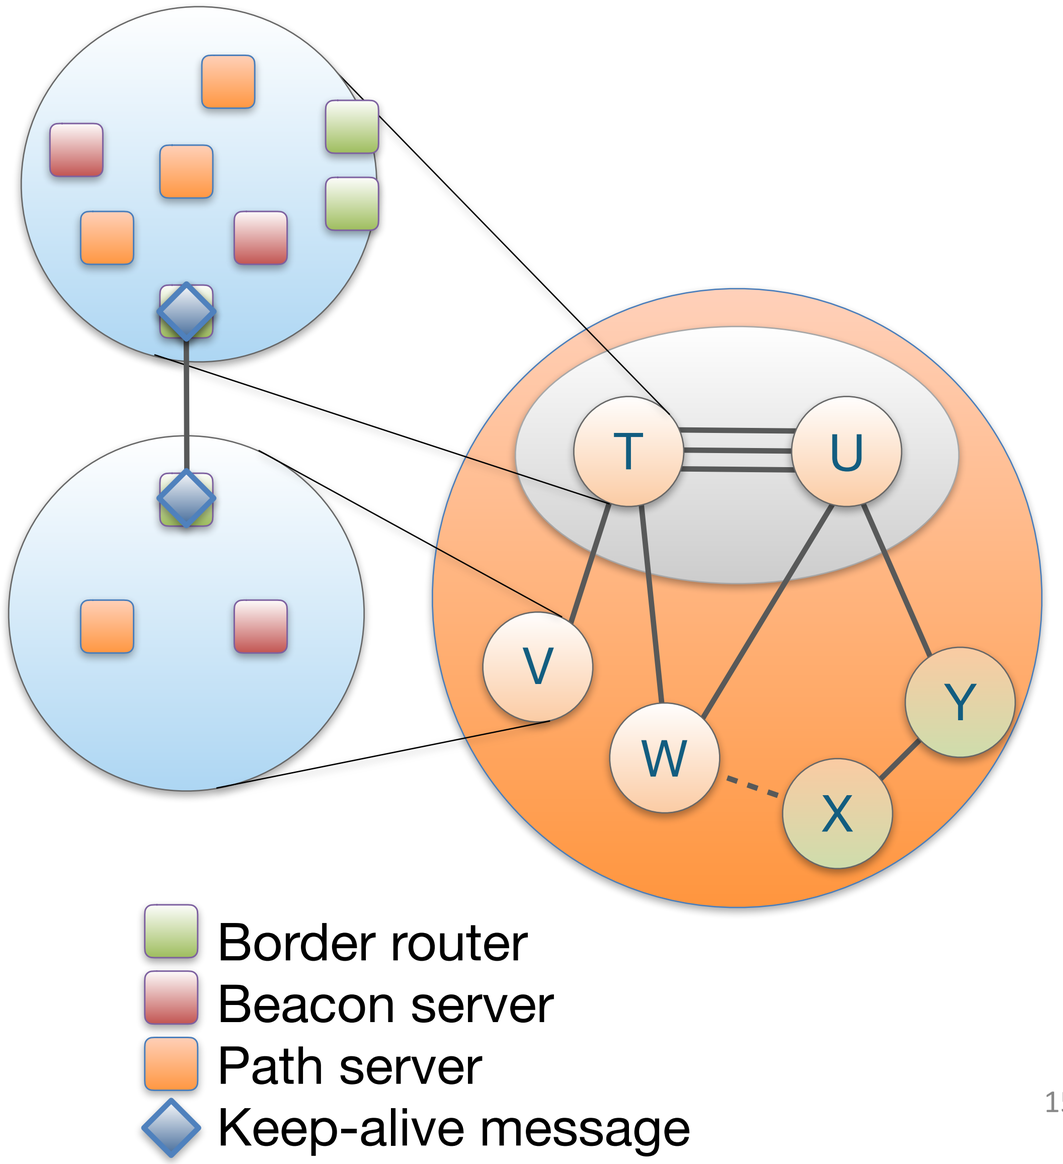

Secure Path Revocation

- Each AS adds a Revocation Token (RT) to the PCB
  - RTs enable efficient authentication of link revocation messages from corresponding AS
- When packet reaches a border router that cannot forward the packet, router sends a link revocation message back to host
- Host re-distributes revocation message to path and beacon servers, to remove path segments containing broken links
- Section 7.3 in SCION book describes this process in detail

Service Anycast

- To support service-based communication, SCION offers service anycast
  - Service address type used as a packet's destination address
- An up-path segment can be included, and a service anycast extension can indicate in which ASes the request should be considered
- Border routers determine if the packet should be sent to a server instance in the AS

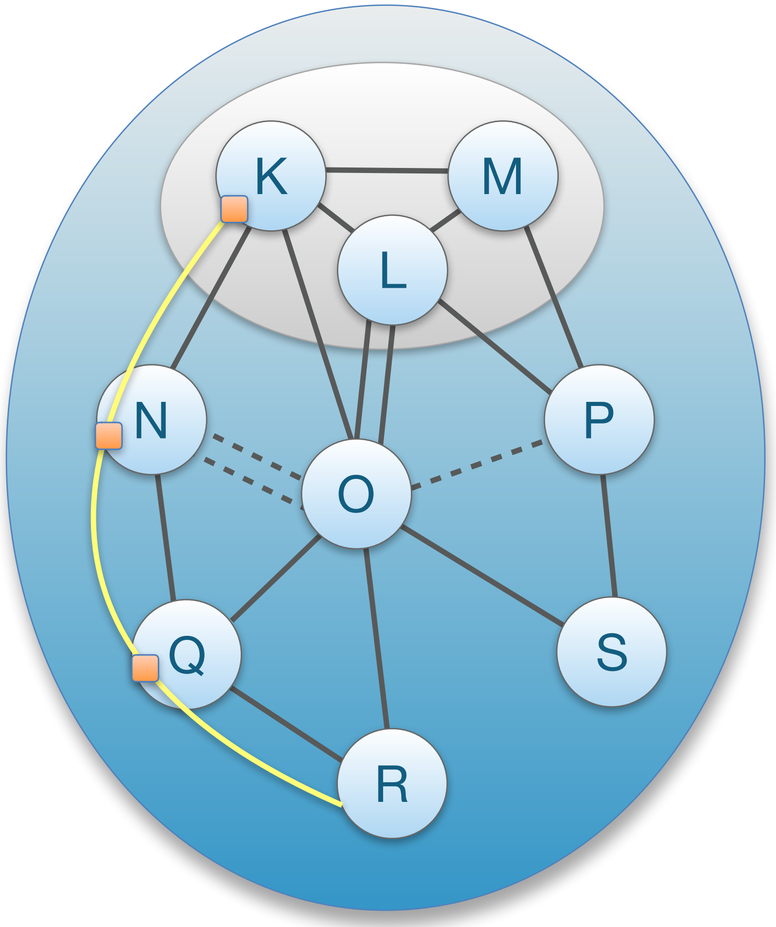

Failure Resilience and Service Discovery

- For reliability, control-plane infrastructure services rely on a consistency service with the following properties
  - Leader election
  - Group membership list
  - Distributed consistent database
- Currently, we are using Apache Zookeeper for this purpose
- Discovery service provides list of active server instances
  - Combination of information from consistency service and static configurations

Failure Resilience: Beacon Service

- All AS beacon server instances connect to consistency service and appear as group members
  - Leader election algorithm determines master beacon server
- PCBs are disseminated with a SCION service address as the destination address
  - SCION border router will select one running beacon server instance to deliver PCB to
  - Receiving beacon server instance re-distributes PCB to all other instances via the consistency service's distributed database
- Master beacon server disseminates PCBs and registers up-path segments at local path server, and down-path segments at core path servers

Failure Resilience: Path Service

- All AS path server instances connect to consistency service and appear as group members
  - Leader election algorithm determines master path server in a core AS
  - No leader election in non-core AS
- Path replication within core AS
  - To handle high load, down-path segment registrations are not disseminated by consistency service
  - Instead, non-master path servers fetch down-path segments from master path server and push registered down-path segments to master path server
  - Down-path segment registrations are also sent to a path server of each core AS
- Path replication within non-core AS
  - Non-core path servers use consistency service for up-path segment replication

SCION Control Message Protocol (SCMP)

- SCMP is analogous to ICMP in the current Internet and provides:
  - Network diagnostic: SCION equivalents of ping or traceroute
  - Error messages: signal problems with packet processing or inform end hosts about network-layer problems
- SCMP is the first secure control message protocol we are aware of
  - Asymmetric authentication (AS certificates) or symmetric authentication (DRKey) are supported

For More Information …

- … please see our web page: www.scion-architecture.net
- Chapter 7 of our book “SCION: A secure Internet Architecture”
  - Available from Springer this Summer 2017
  - PDF available on our web site