Dissemination of Paths in Path-Aware Networks

Christos Pappas
Network Security Group, ETH Zurich

IETF, November 16, 2017

ETH zürich

PANRG Motivation

How does path-awareness extend to the edge?

PANRG Motivation

How does path-awareness extend to the edge?

- endpoint discovery of paths
  - scalability
  - failure-free paths / recovery

PANRG Motivation

How does path-awareness extend to the edge?

- endpoint discovery of paths
  - scalability
  - failure-free paths / recovery

- explicit discovery of path properties by endpoints
  - dynamic properties (e.g., congestion, utilization)
  - static properties(e.g., capacity, disjointness)

PANRG Motivation

How does path-awareness extend to the edge?

- endpoint discovery of paths
  - scalability
  - failure-free paths / recovery

- explicit discovery of path properties by endpoints
  - dynamic properties (e.g., congestion, utilization)
  - static properties(e.g., capacity, disjointness)

- explicit endpoint selection of paths
  - path representation

PANRG Motivation

How does path-awareness extend to the edge?

This talk: Path Dissemination

✤ endpoint discovery of paths
- scalability
- failure-free paths / recovery

✤ explicit discovery of path properties by endpoints
- dynamic properties (e.g., congestion, utilization)
- static properties(e.g., capacity, disjointness)

✤ explicit endpoint selection of paths
- path representation

Path Dissemination

💡 Path control to endpoints = ⬆️path information + ⬆️path choice

Path Dissemination

💡 Path control to endpoints = ⬆path information + ⬆path choice

S1: Path Construction

- topology exploration
- path creation

} network or network & endpoint

S2: Path Selection

- choosing one or more constructed paths

S3: Path Representation (encoding)

- how path choice is expressed to the network

Path Dissemination

Path control to endpoints = ↑path information + ↑path choice

S1: Path Construction

- topology exploration
- path creation

} network or network & endpoint

S2: Path Selection

- choosing one or more constructed paths

S3: Path Representation (encoding)

- how path choice is expressed to the network

Path Dissemination Design Choices

- **The SCION Internet Architecture**
  Barrera et al., in Communications of the ACM, 2017
- NIRA: A New Inter-Domain Routing Architecture
  Yang et al., in ACM ToN, 2007
- Pathlet Routing
  Godfrey et al., in SIGCOMM, 2009

Path Dissemination in SCION - Overview

Path Dissemination in SCION - Overview

S1a: Topology Exploration: Beaconing

S1b: Path-Segment Construction: Creation of up/down-segments

} S1: Path Creation

✣ Registration of Path Segments

Path Dissemination in SCION - Overview

Network
S1a: Topology Exploration: Beaconing S1b: Path-Segment Construction: Creation of up/down-segments ✣ Registration of Path Segments

} S1: Path Creation

Path Dissemination in SCION - Overview

Network

S1a: Topology Exploration: Beaconing

S1b: Path-Segment Construction: Creation of up/down-segments

✣ Registration of Path Segments

} S1: Path Creation

Endpoints

✣ Lookup of Path Segments

S2: Path Combination and Selection

S3: Path Representation

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

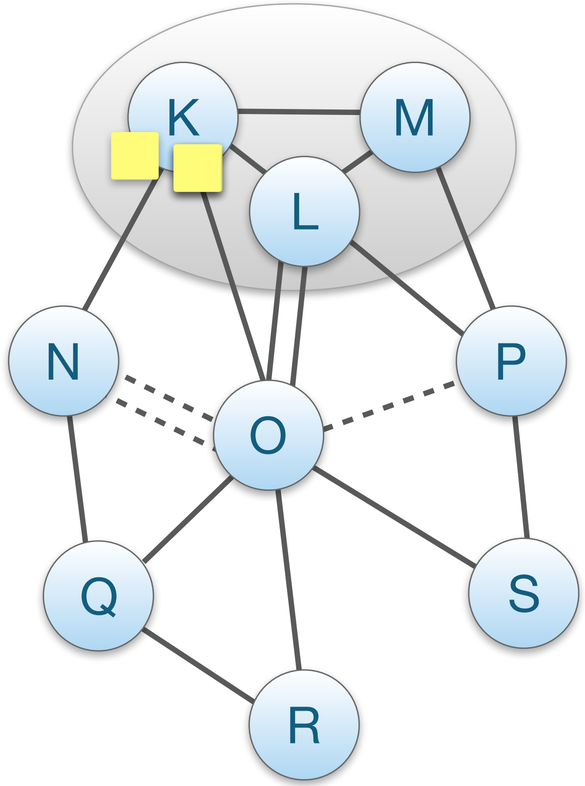

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

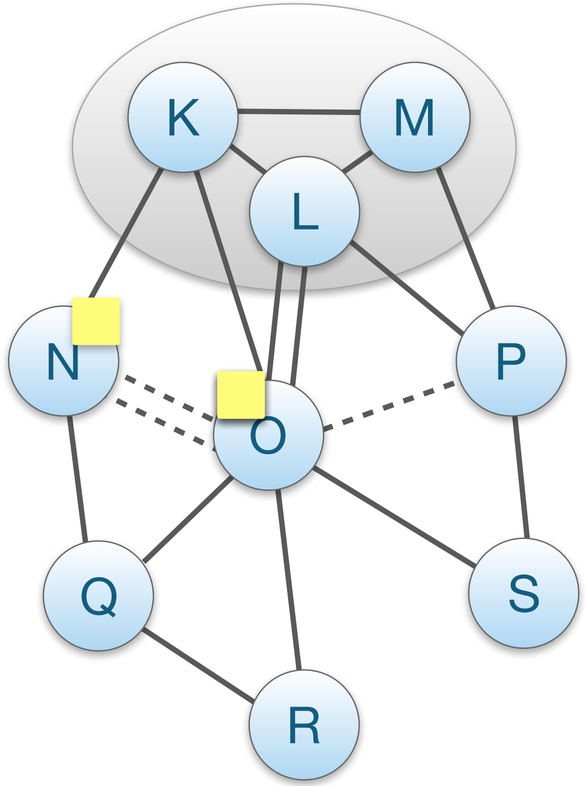

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

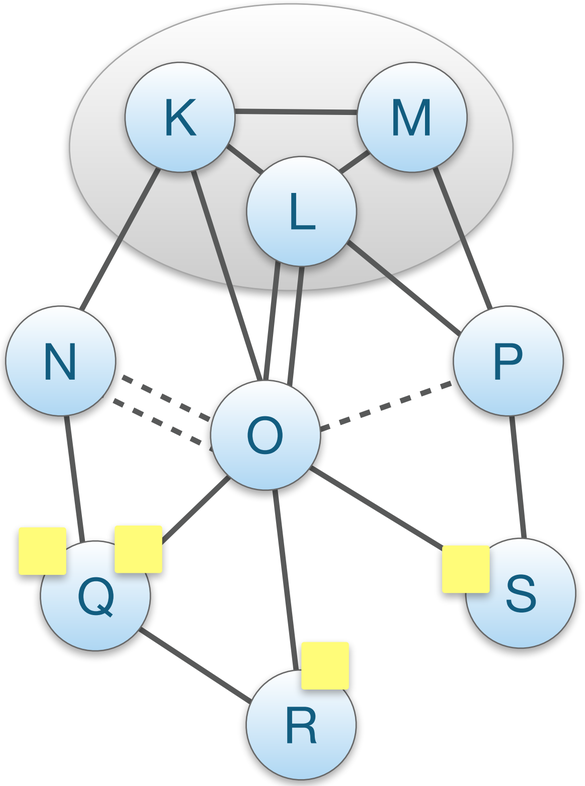

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

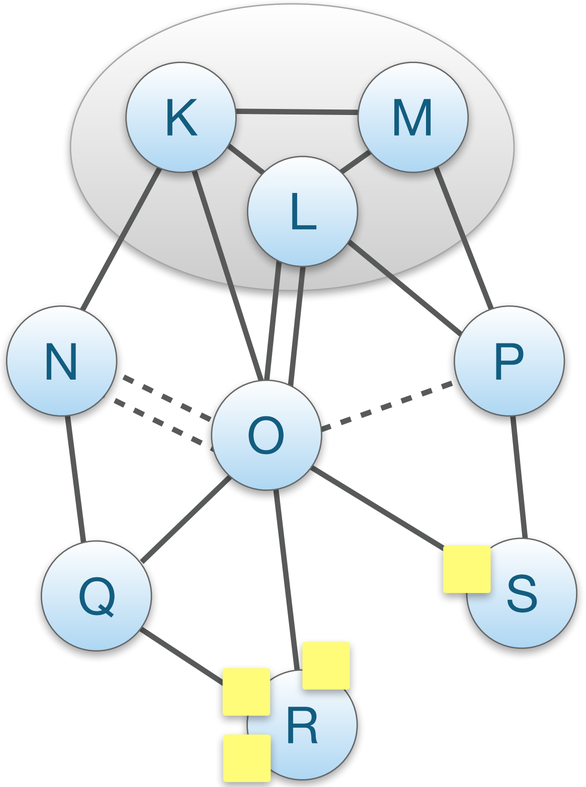

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

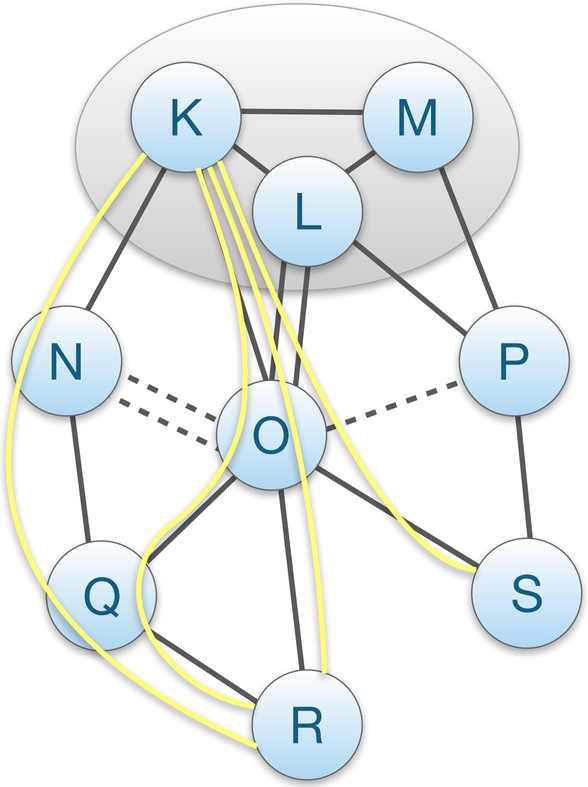

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

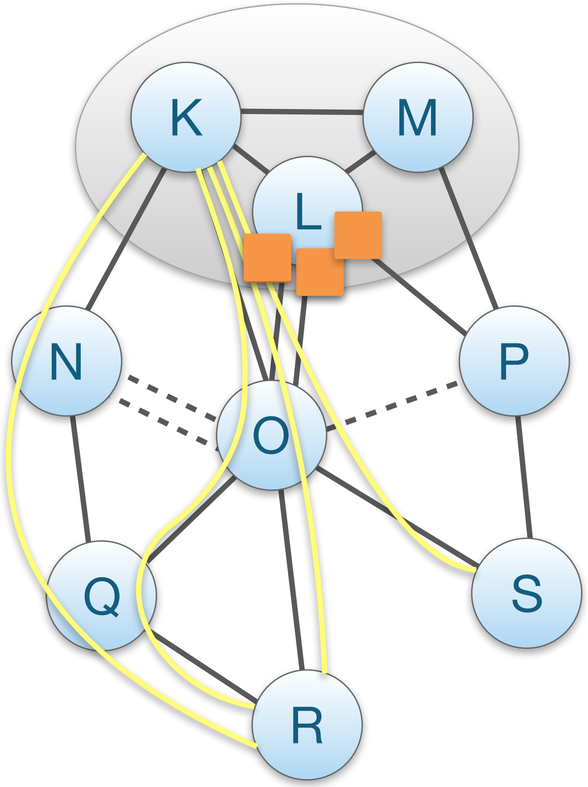

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

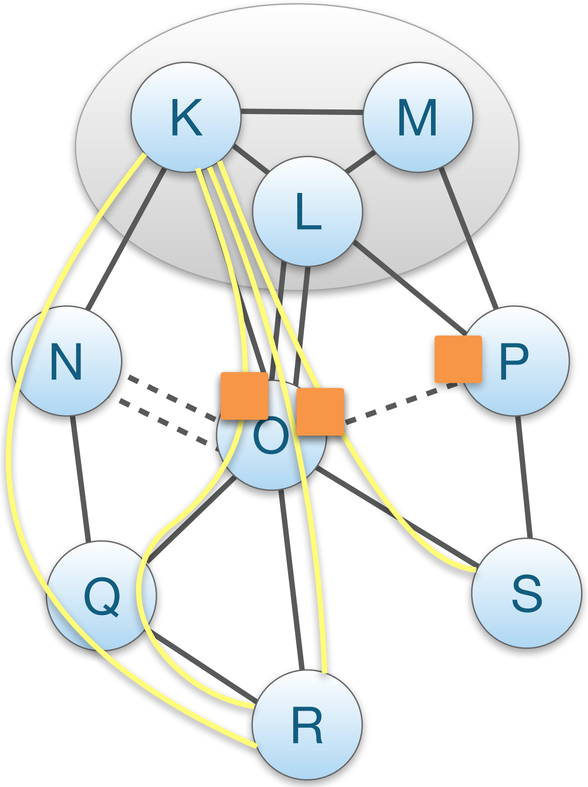

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

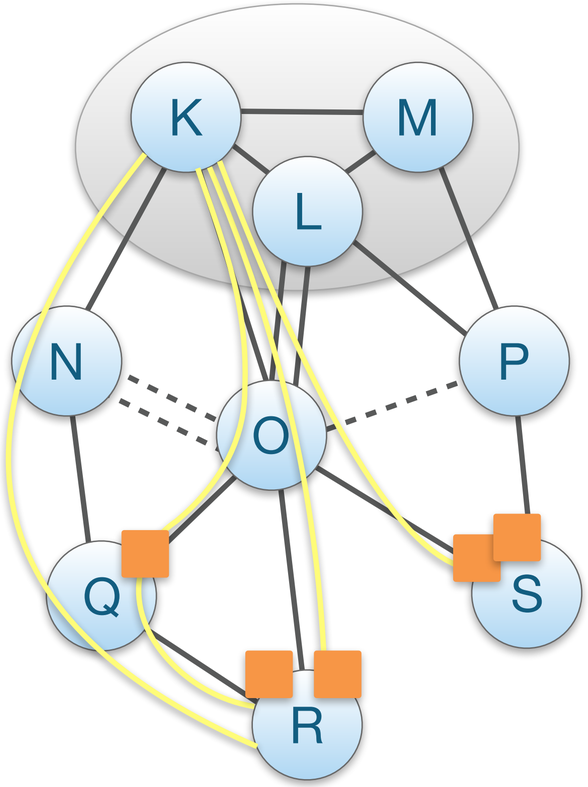

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

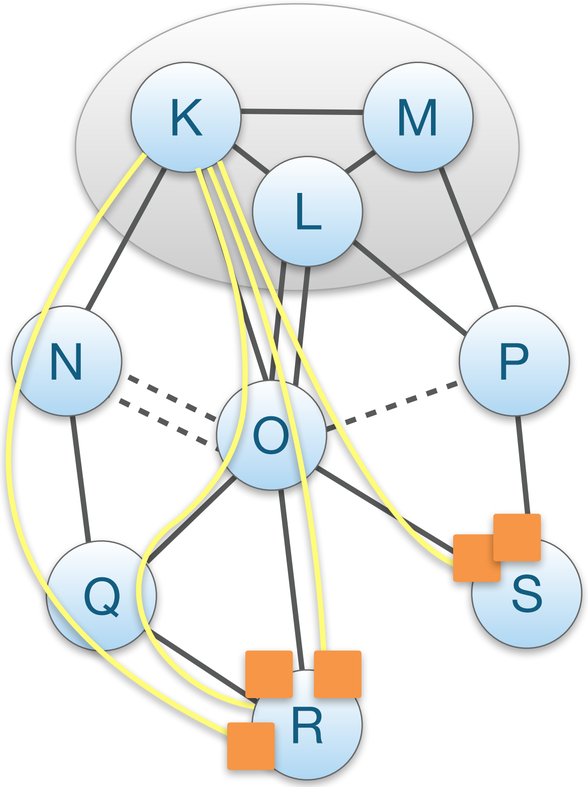

S1a: Topology Exploration through Beaconing

- “Core ASes” initiate beacons
- Beacons traverse the topology hierarchically from the core ASes and downstream
- Each AS receives multiple beacons (path diversity)

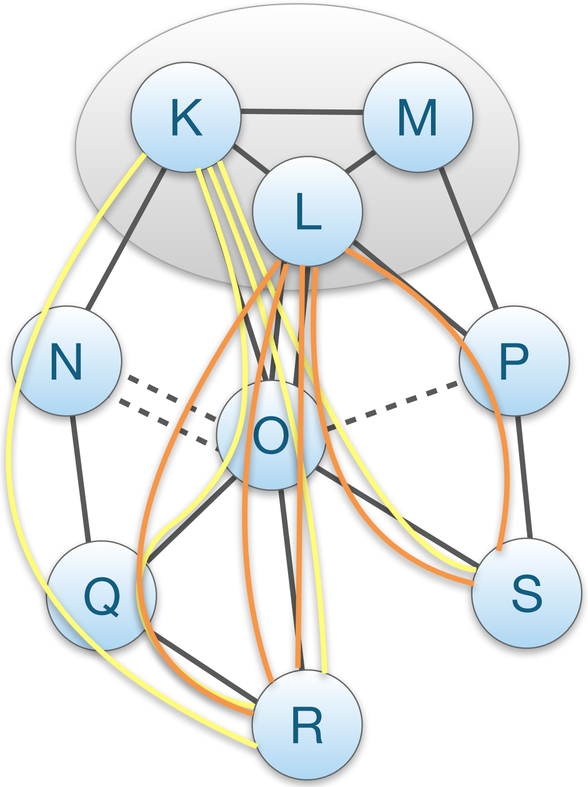

Path Property Dissemination through Beaconing

- Beacons bind properties to paths

- Per-AS (static) information included
  - MTU
  - Link capacities
  - Path policy
  - Supported services
  - Supported crypto protocols

S1b: Path-Segment Construction

- Beacons contain path segments to communicate with the core ASes that initiated it

- Up-segments: path segment from non-core AS to core AS
    - Example: R → K

- Down-segments: path segment from core AS to to non-core AS
    - Example: L → S

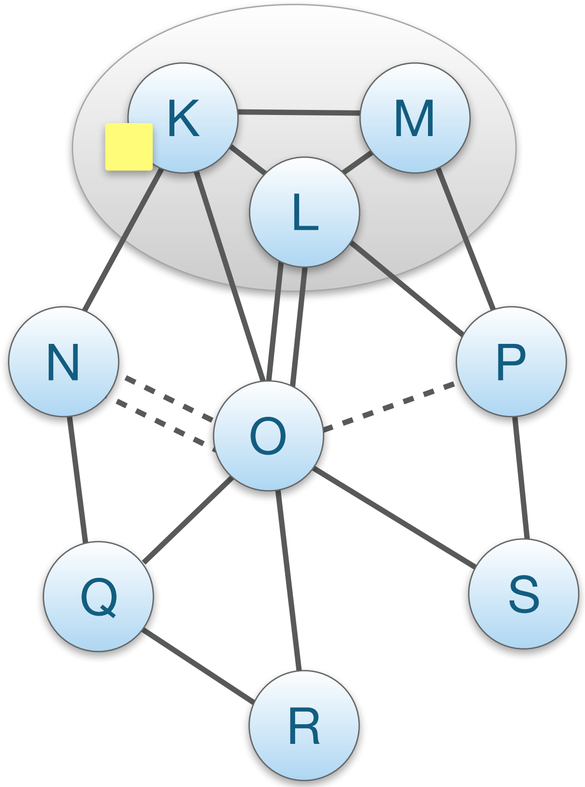

S1b: Path-Segment Construction

- Beacons contain path segments to communicate with the core ASes that initiated it

- Up-segments: path segment from non-core AS to core AS
    - Example: R → K

- Down-segments: path segment from core AS to to non-core AS
    - Example: L → S

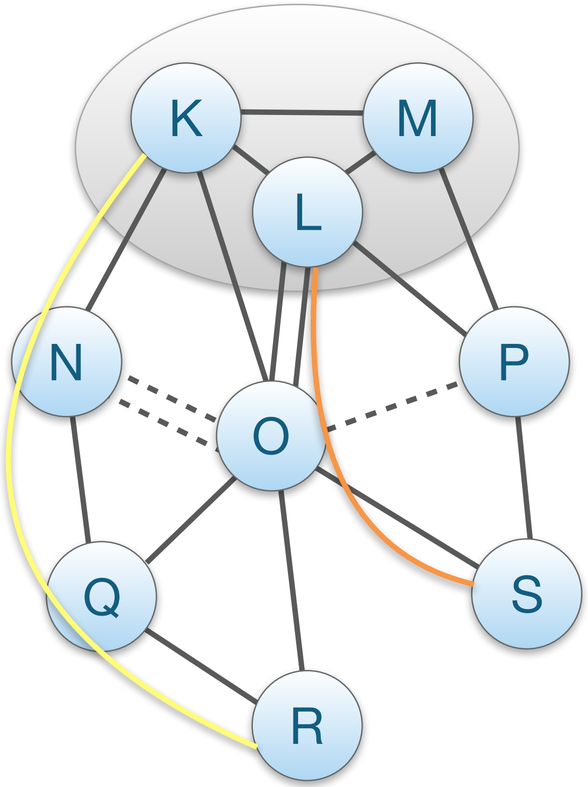

Registration of Path Segments

💡 End-to-end paths = up-segment + core-segment + down-segment

- Up/down-segments are registered and looked up

- Up-segments: AS selects subset of path segments to be used for local hosts to reach the core

- Down-segments: AS selects subset of path segments to be used by others to reach the AS

Registration of Path Segments

💡 End-to-end paths = up-segment + core-segment + down-segment

Path Servers

- Offer lookup service
- Two “types” of offered service
  - Core path servers: Core ASes serve down-segments to remote endpoints
  - Local path servers: Non-core ASes 1) serve up-segments to local endpoints, 2) resolve and cache responses for remote endpoints

S2: Path Combination and Selection

💡 Creating the final end-to-end path

1. Path-segment Lookup (2 translations needed)

   a. Endpoint contacts name server with a name

   b. Endpoint contacts local path server with a destination AS

2. Path-segment Combination

3. Path Selection

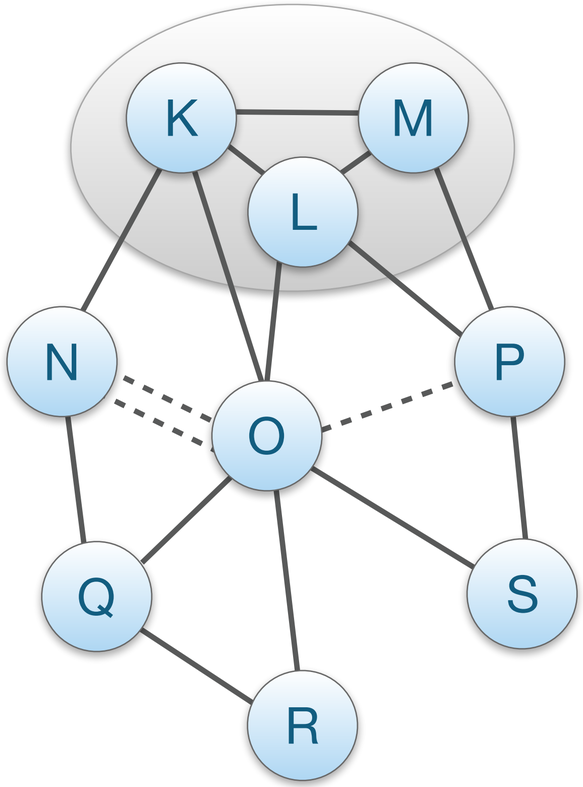

S2: Path Combination and Selection

💡 Creating the final end-to-end path

1. Path-segment Lookup (2 translations needed)

   a. Endpoint contacts name server with a name

   b. Endpoint contacts local path server with a destination AS

2. Path-segment Combination

   Up-segment + Core-segment + Down-segment

3. Path Selection

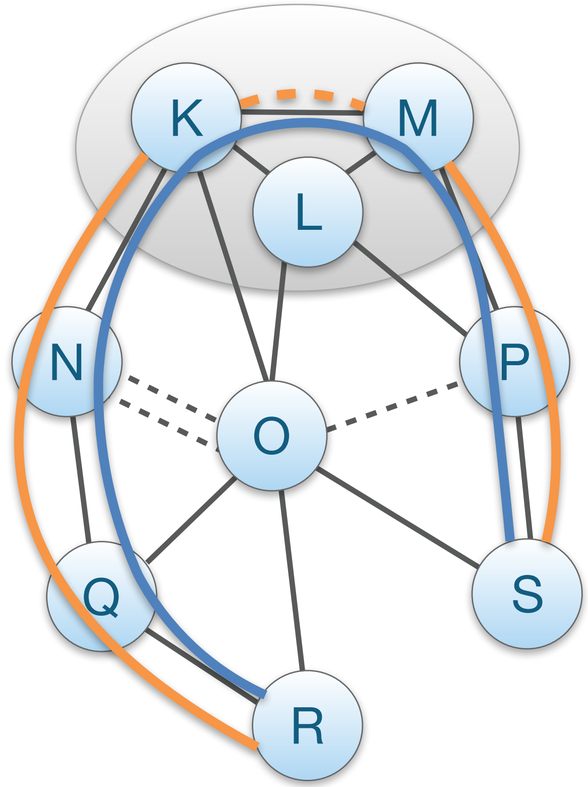

S2: Path Combination and Selection

💡 Creating the final end-to-end path

1. Path-segment Lookup (2 translations needed)

   a. Endpoint contacts name server with a name

   b. Endpoint contacts local path server with a destination AS

2. Path-segment Combination

   Up-segment + Core-segment + Down-segment

   Up-segment + Down-segment offering same peering link

3. Path Selection

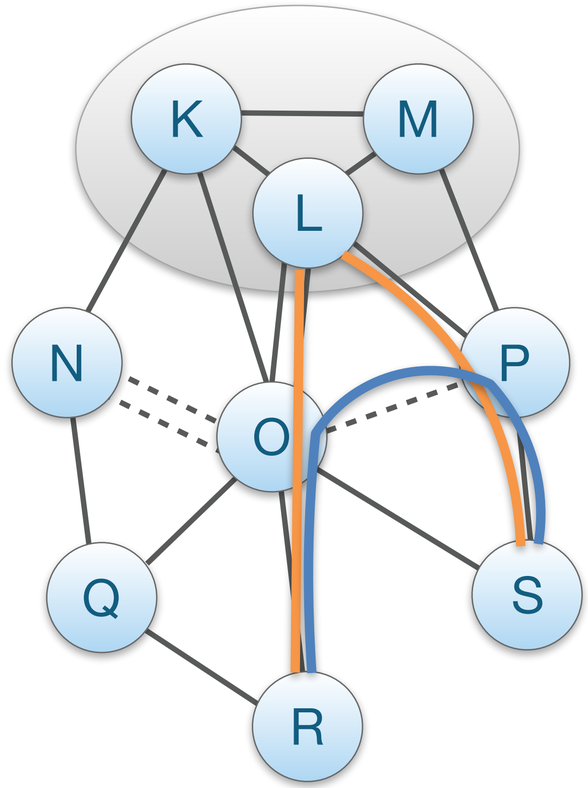

S3: Path Representation

- Packet-carried forwarding information
  - ✤ Interface-level forwarding instructions
  - ✤ No lookups at the data plane

- Beacons collect it during path exploration & then it is registered with up- and down-segments

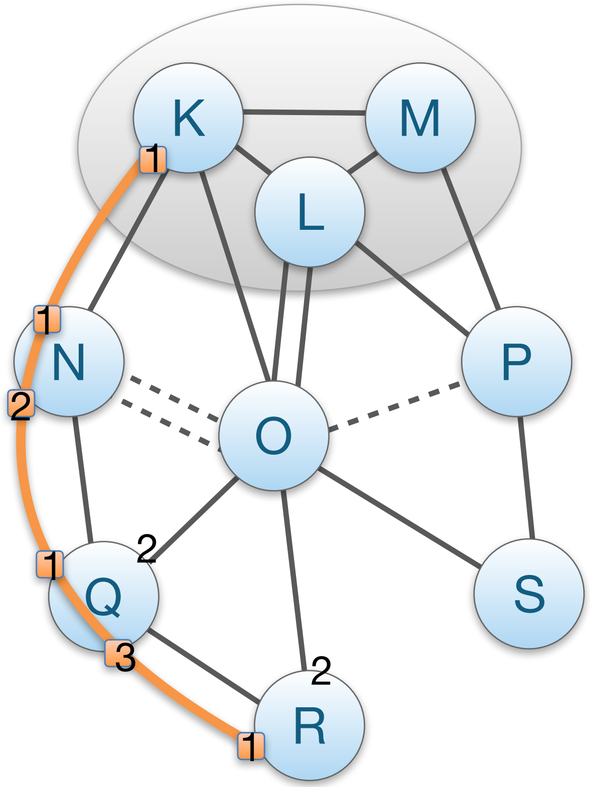

NIRA: A New Inter-domain Routing Architecture

S1: Topology Exploration & Construction

- Endpoints participate in a path-vector component and discover their up-graph (paths leading to the core)
- Selected paths are registered with a name-to-route lookup service, so that endpoints can be reached by others

S2: Path Selection

- Sender chooses a path from its up-graph
- … and combines with the looked up path for an end-to-end path

S3: Path Representation

- Through source/destination addresses (e.g., IPv6)
- Source and destination addresses together with forwarding state at routers, encode the up- and down-paths

Pathlet Routing

💡 Vnodes: virt. node created by AS, representing route structure in its network (e.g., 1 per router)
Pathlets: a sequence of vnodes along which an AS will route traffic

S1: Topology Exploration & Construction

- Routers announce their vnodes, neighbours can arbitrarily combine vnodes to create pathlets, and also combine pathlets into longer pathlets
- Path-vector to announce chosen pathlets (not to enforce routes)

S2: Path Selection - unspecified

S3: Path Representation

- Pathlets associated with flat/opaque identifiers (FIDs) that are put in packets
- Routers look up action for specific FID (reminds of MPLS)

Summary

Path Dissemination Design Choices…

	Topology Exploration & Construction	Path Selection	Path Representation
SCION	Beaconing follows business relat.	Lookup Service & then endpoint choice	Packet-carried FW Information
NIRA	Path-vector follows business relat.	Lookup Service & then endpoint choice	Address-based
Pathlet Routing	Arbitrary pathlets, path vector announced	—	MPLS like

Summary

Path Dissemination Design Choices…

	Topology Exploration & Construction	Path Selection	Path Representation
SCION	Beaconing follows business relat.	Lookup Service & then endpoint choice	Packet-carried FW Information
NIRA	Path-vector follows business relat.	Lookup Service & then endpoint choice	Address-based
Pathlet Routing	Arbitrary pathlets, path vector announced	—	MPLS like

Answered Challenges…

- multi-path support
- static path properties (e.g., capacity)

Summary

Path Dissemination Design Choices…

	Topology Exploration & Construction	Path Selection	Path Representation
SCION	Beaconing follows business relat.	Lookup Service & then endpoint choice	Packet-carried FW Information
NIRA	Path-vector follows business relat.	Lookup Service & then endpoint choice	Address-based
Pathlet Routing	Arbitrary pathlets, path vector announced	—	MPLS like

Answered Challenges…

- multi-path support
- static path properties (e.g., capacity)

Open Questions…

- how to convey dynamic path properties (e.g., congestion, latency information)
- endpoint → lookup service interface
- path information to applications?

Summary

Path Dissemination Design Choices…

	Topology Exploration & Construction	Path Selection	Path Representation
SCION	Beaconing follows business relat.	Lookup Service & then endpoint choice	Packet-carried FW Information
NIRA	Path-vector follows business relat.	Lookup Service & then endpoint choice	Address-based
Pathlet Routing	Arbitrary pathlets, path vector announced	—	MPLS like

Answered Challenges…

- multi-path support
- static path properties (e.g., capacity)

Open Questions…

- how to convey dynamic path properties (e.g., congestion, latency information)
- endpoint → lookup service interface
- path information to applications?

Thank you!