SCION: A Secure Multipath Interdomain Routing Architecture

Adrian Perrig

Network Security Group, ETH Zürich

ETH zürich

SCION

Internet Weakness: DoS and DDoS Attacks

- Expensive and difficult to protect against DoS und DDoS attacks
- Despite large investments, attacks continue to be successful
  - **November 2015: Protonmail attacked during 1 week**
  - **March 2016: CH e-commerce under attack: Digitec, Galaxus, SBB, Migros, etc.** (Hackers demanded 25 Bitcoins to stop attacks)
  - **Fall 2016: Global Mirai botnet attacks, e.g., OVH, Dyn, russian banks**
  - June 2017: Northkorea “Hidden Cobra” botnet uncovered
  - September 2017: Global airport chaos, DDoS paralyzes checkin systems
- **Can we reliably defend against DDoS attacks?**

Internet Weakness: Communication Path Hijacking

- Sender und receiver have limited control over routing paths
- Attacks can hijack and relay paths
- **How can we guarantee communication paths?**

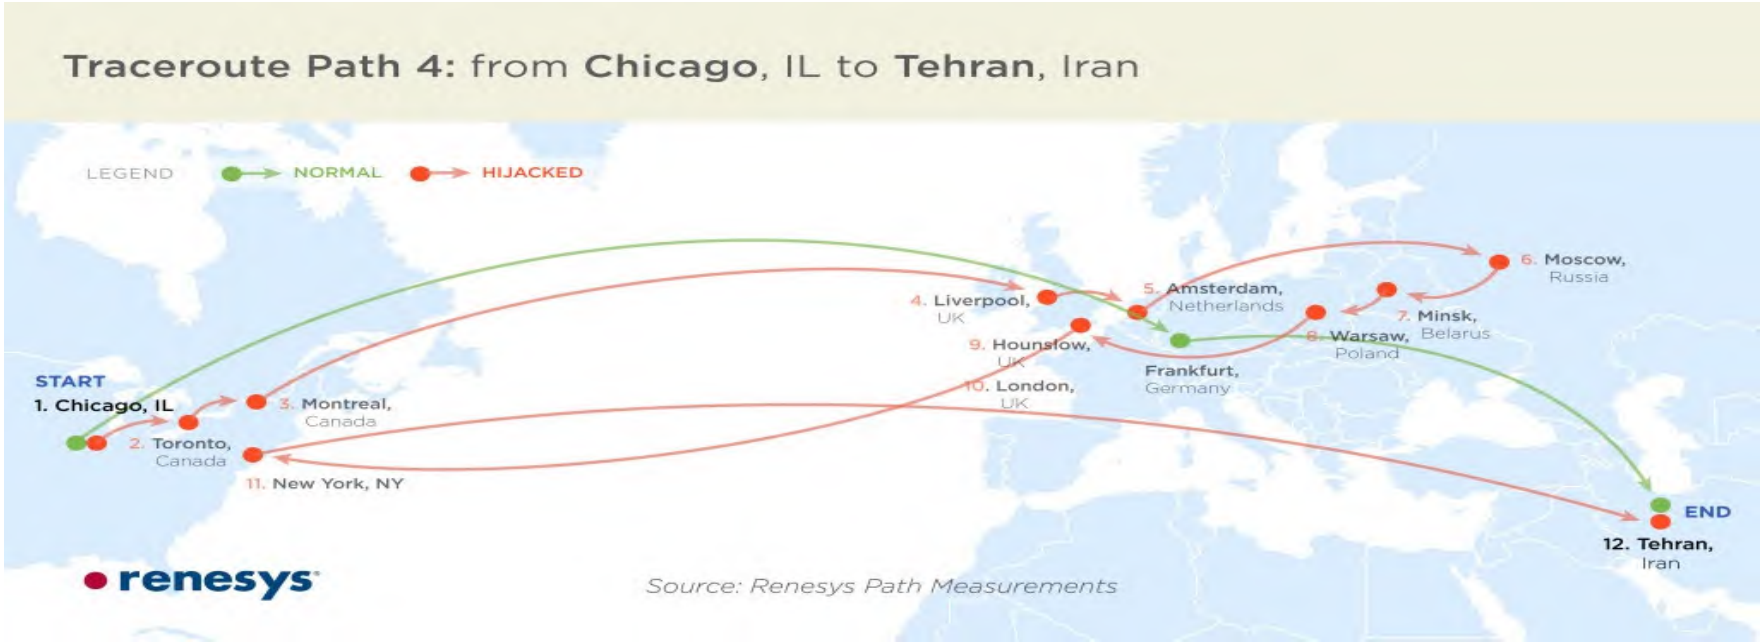

Internet Weakness: Kill Switch ruptures Sovereignty

- Current Internet suffers from several “Kill Switches”, which can halt communication within a geographical area
- Several attack avenues exist: DDoS, BGP hijacking, DNS redirection, BGPSEC / DNSSEC / TLS certificate revocation
- Example August 2017: An erroneous route injected by Google prevents communication for 50% of Internet in Japan during 40 minutes
- **Can we construct an Internet without Kill Switches?**

SCION Architecture Design Goals

- High availability, even for networks with malicious parties
  - Adversary: access to management plane of router
  - Communication should be available if adversary-free path exists
- Secure entity authentication
  that scales to global heterogeneous (dis)trusted environment
- Flexible trust: enable selection of trust roots
- Transparent operation: clear what is happening to packets and whom needs to be relied upon for operation
- Balanced control among ISPs, senders, and receivers
- Scalability, efficiency, flexibility

SCION Overview

- Control plane: How to find end-to-end paths?
  - Path exploration
  - Path registration
- Data plane: How to send packets
  - Path lookup
  - Path combination
- Deployment
- Demos

Approach for Scalability: Isolation Domain (ISD)

- Isolation Domain (ISD): grouping of ASes
- ISD core: ASes that manage the ISD
- Core AS: AS that is part of ISD core
- Control plane is organized hierarchically
  - Inter-ISD control plane
  - Intra-ISD control plane

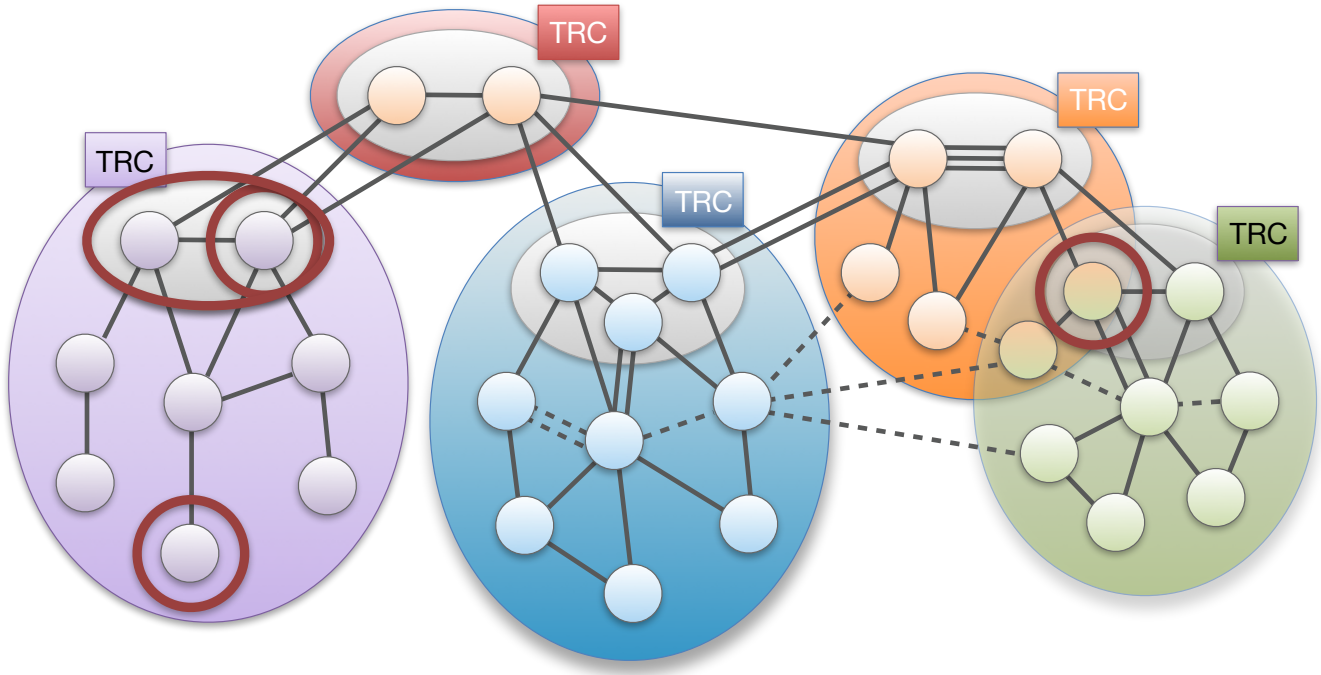

Intra-ISD Path Exploration: Beaconing

- Core ASes K, L, M initiate Path-segment Construction Beacons (PCBs), or "beacons"
- PCBs traverse ISD as a flood to reach downstream ASes
- Each AS receives multiple PCBs representing path segments to a core AS

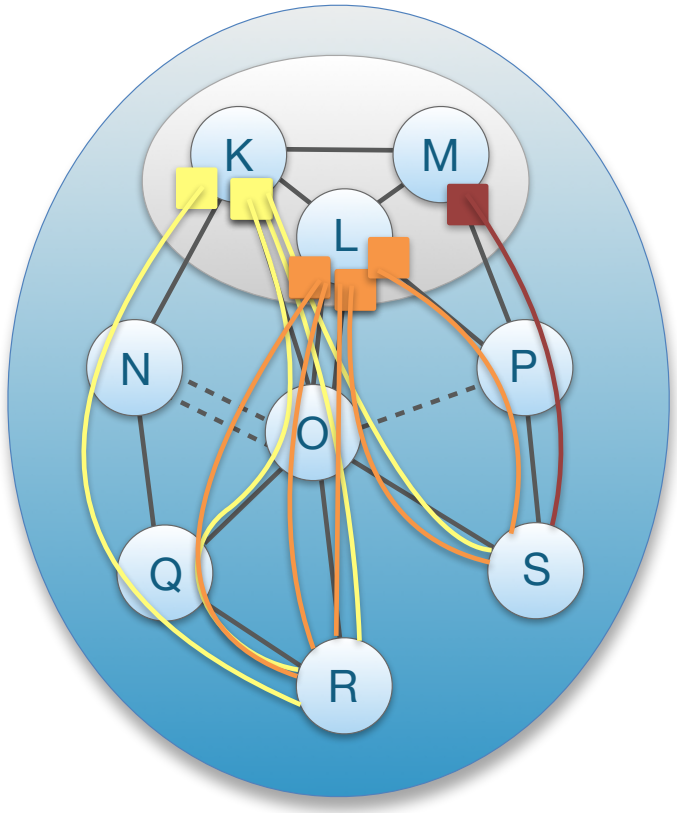

Beaconing in More Detail

- 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

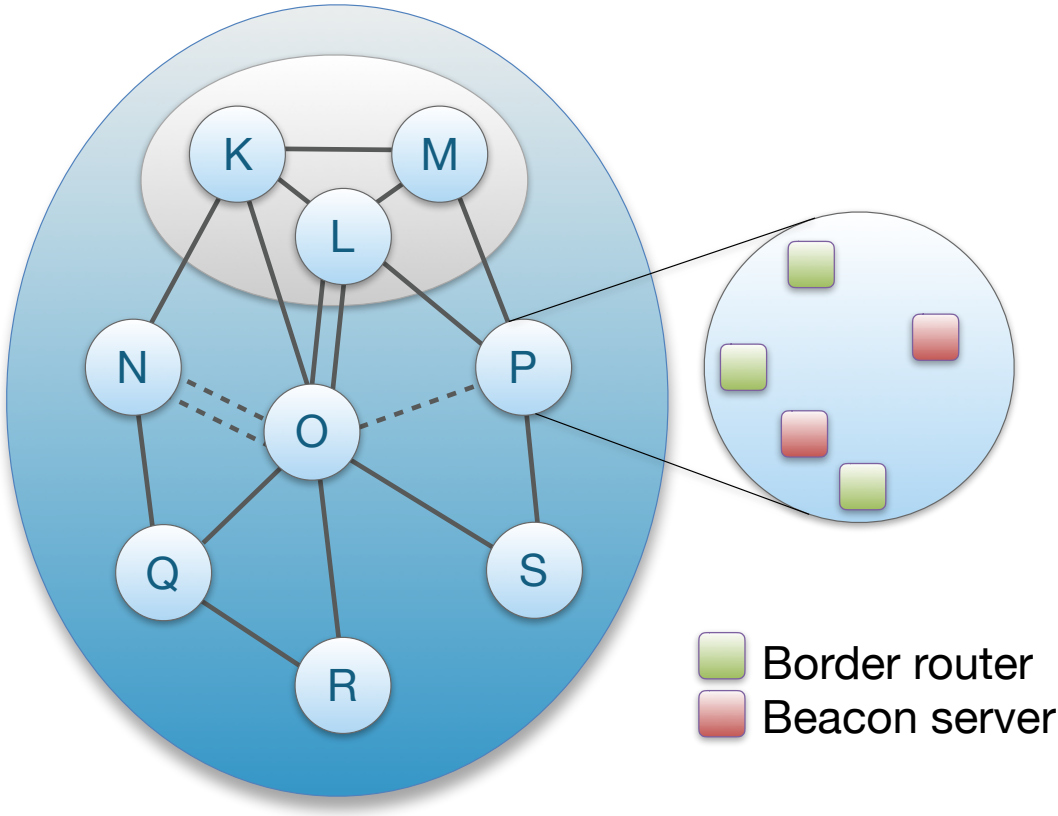

PCB Contents

- 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

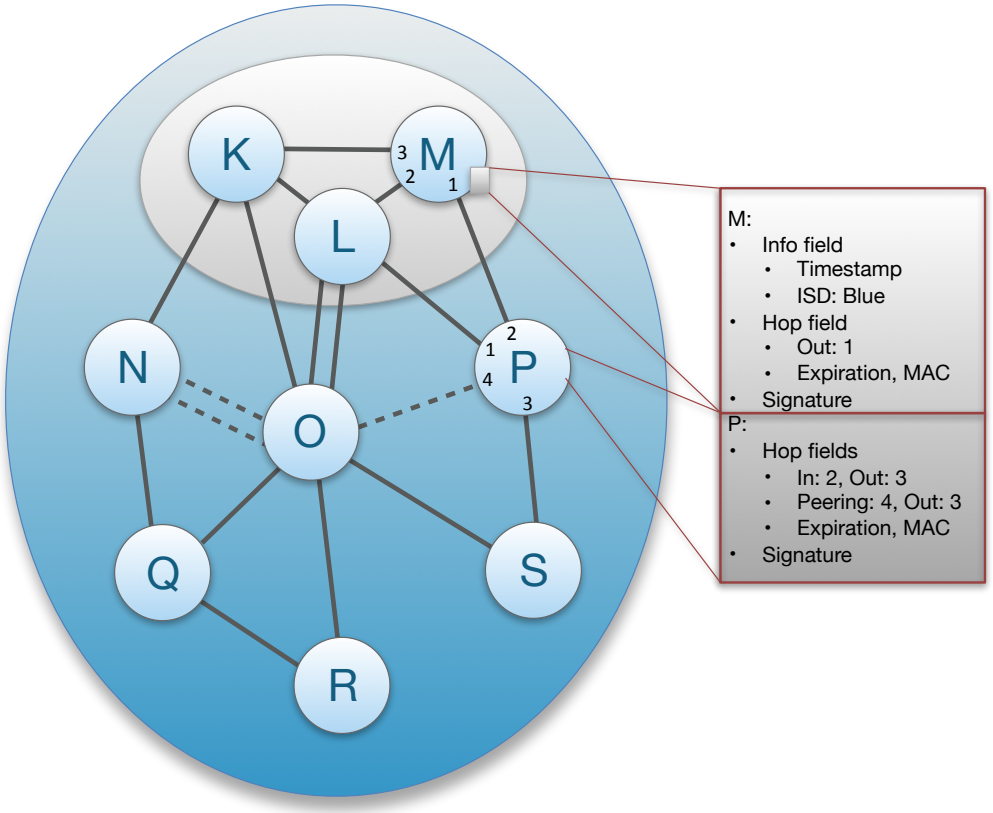

Up-Path and Down-Path Segments

- Intra-ISD beaconing process sends PCBs to ASes
- PCBs contain path segments that can be used as communication paths to communicate with the core AS that initiated it
- Up-path segment: PCB is used from AS to core AS
  - Example: R → K
- Down-path segment: PCB is used from core AS to AS
  - Example: M → S

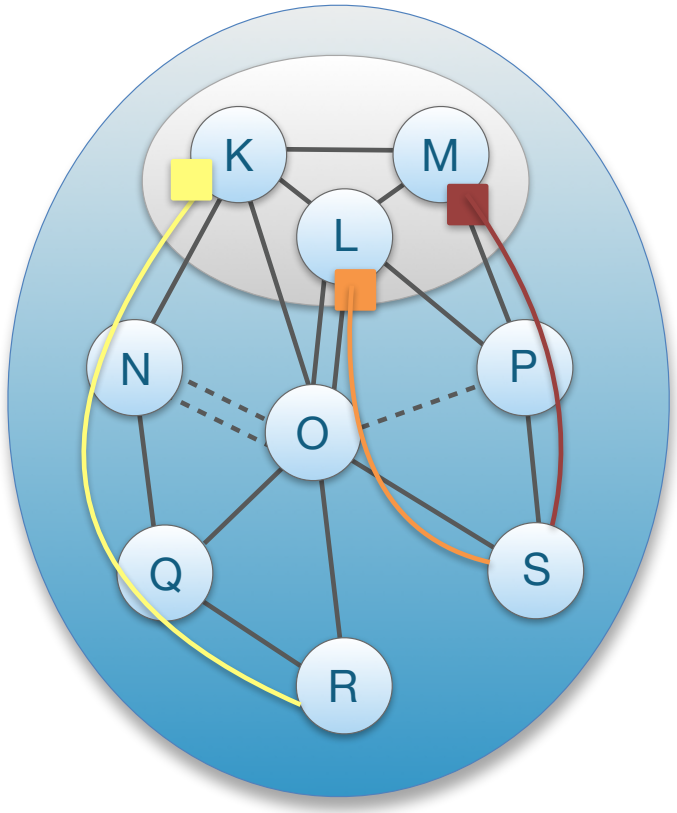

Core Beaconing for Inter-ISD Path Exploration

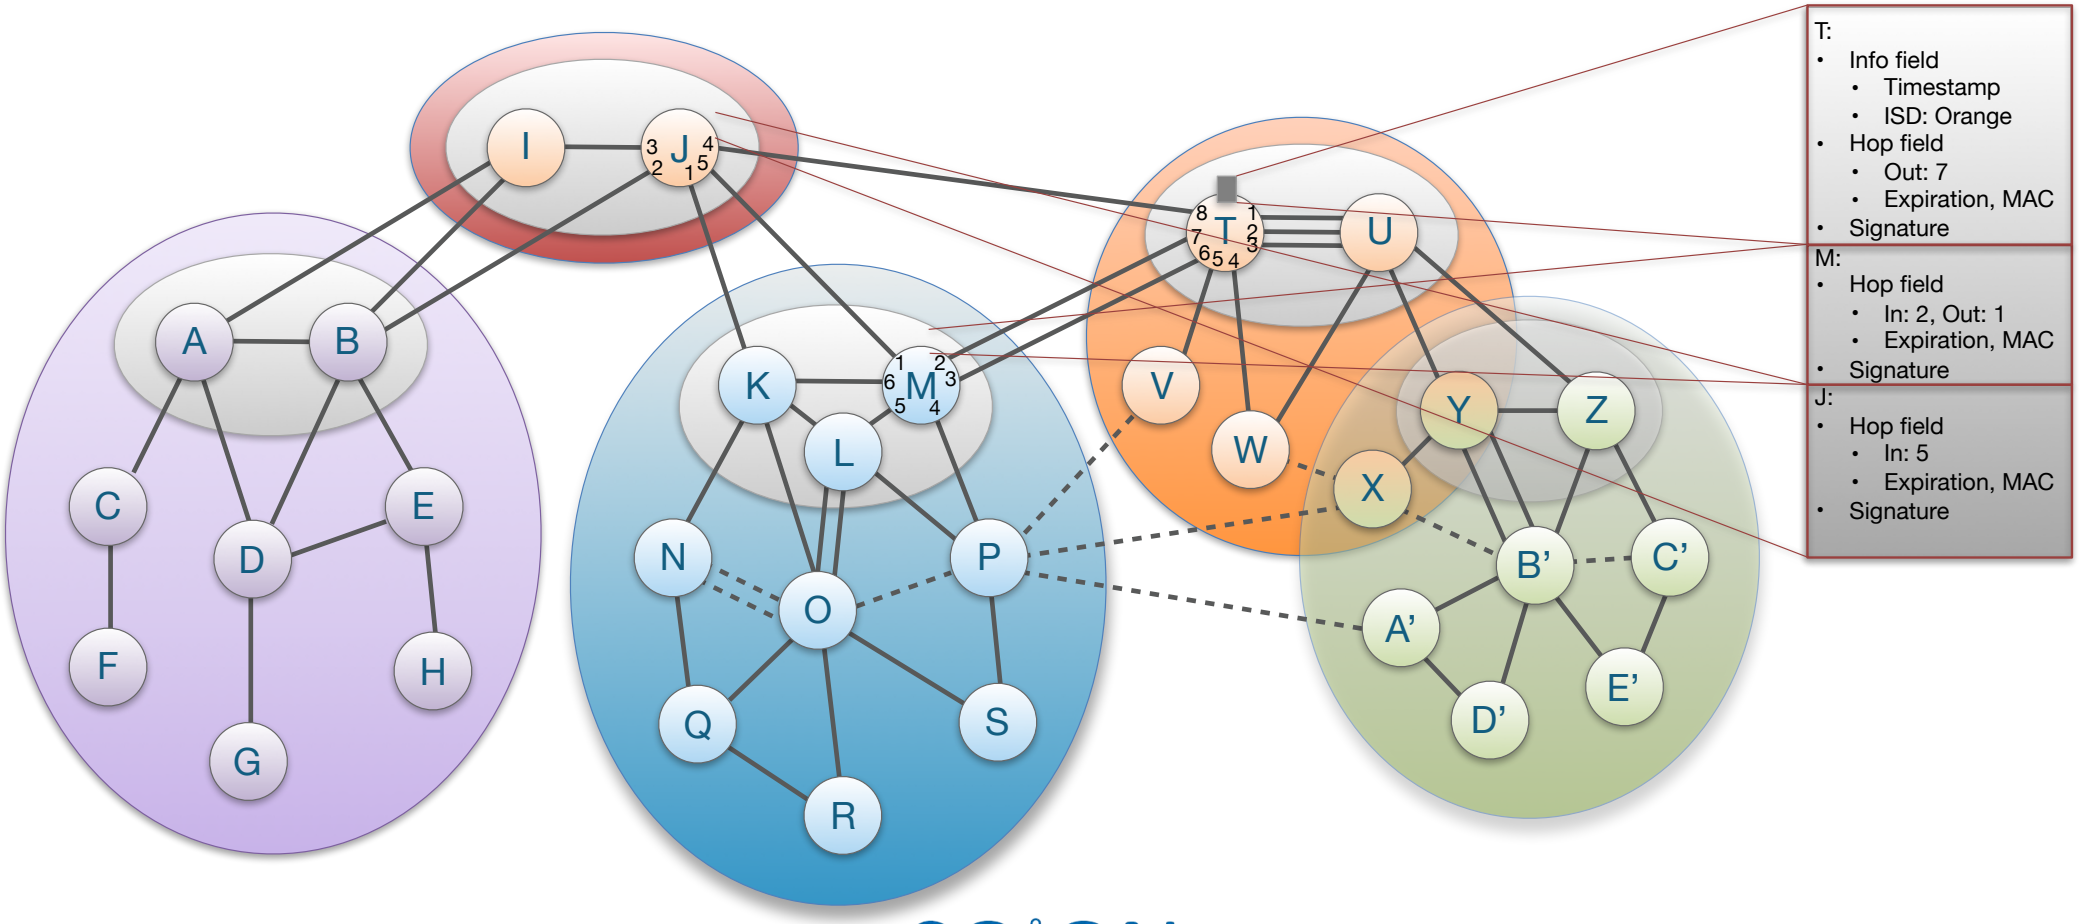

ETH zürich

SCION

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

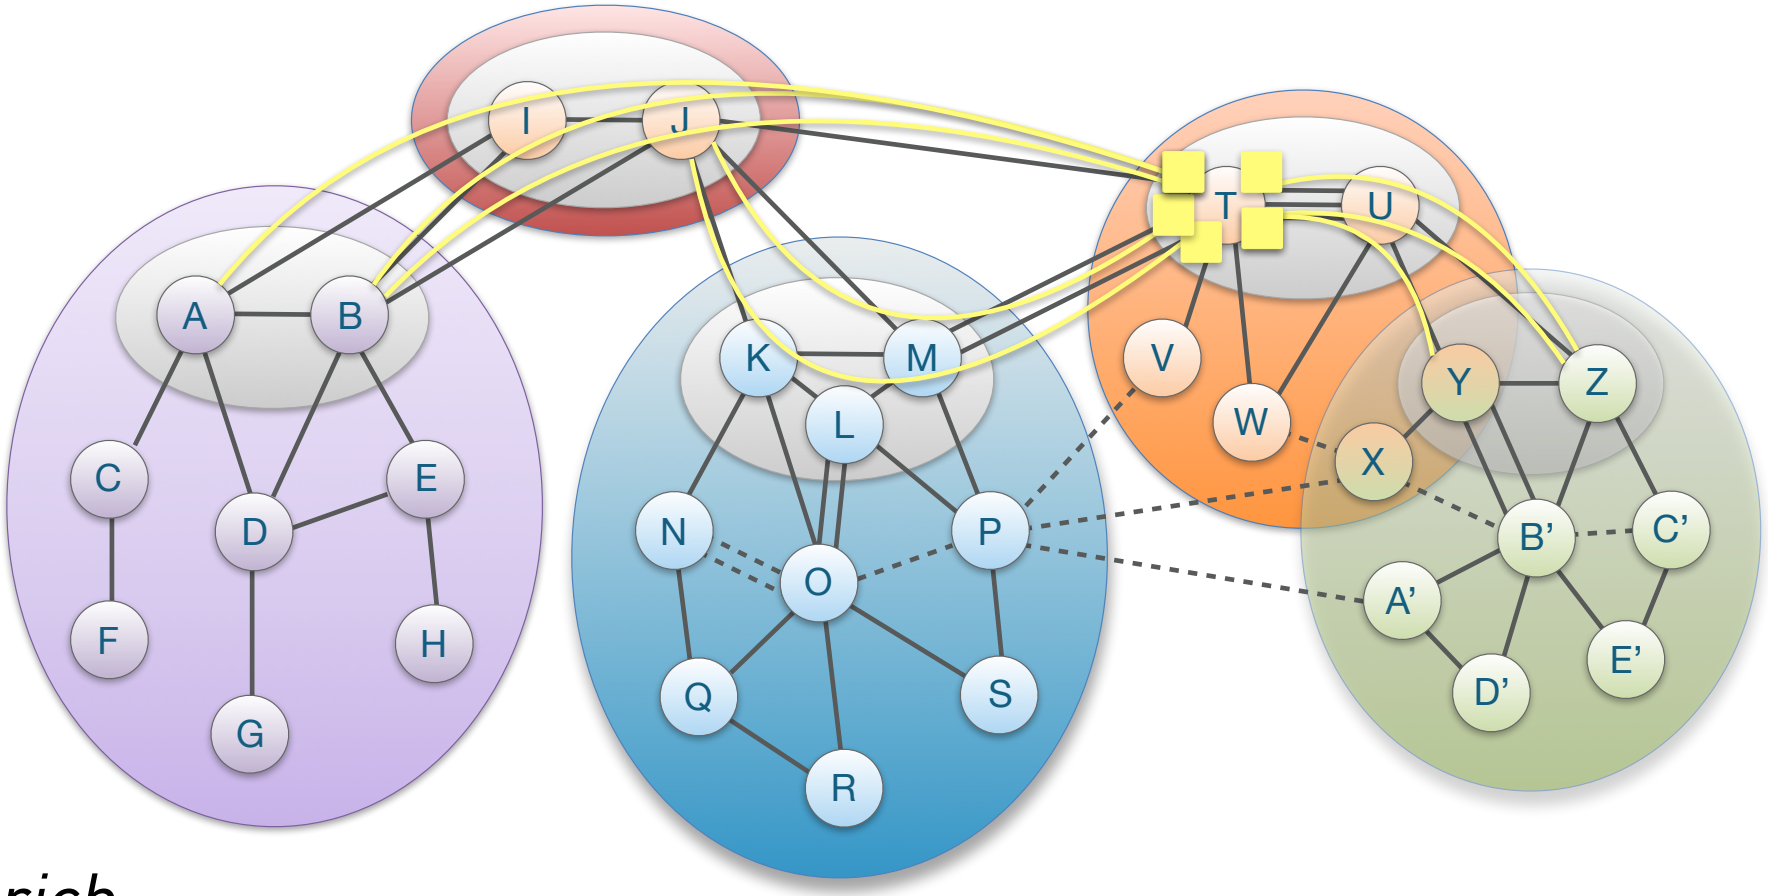

ETH zürich

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

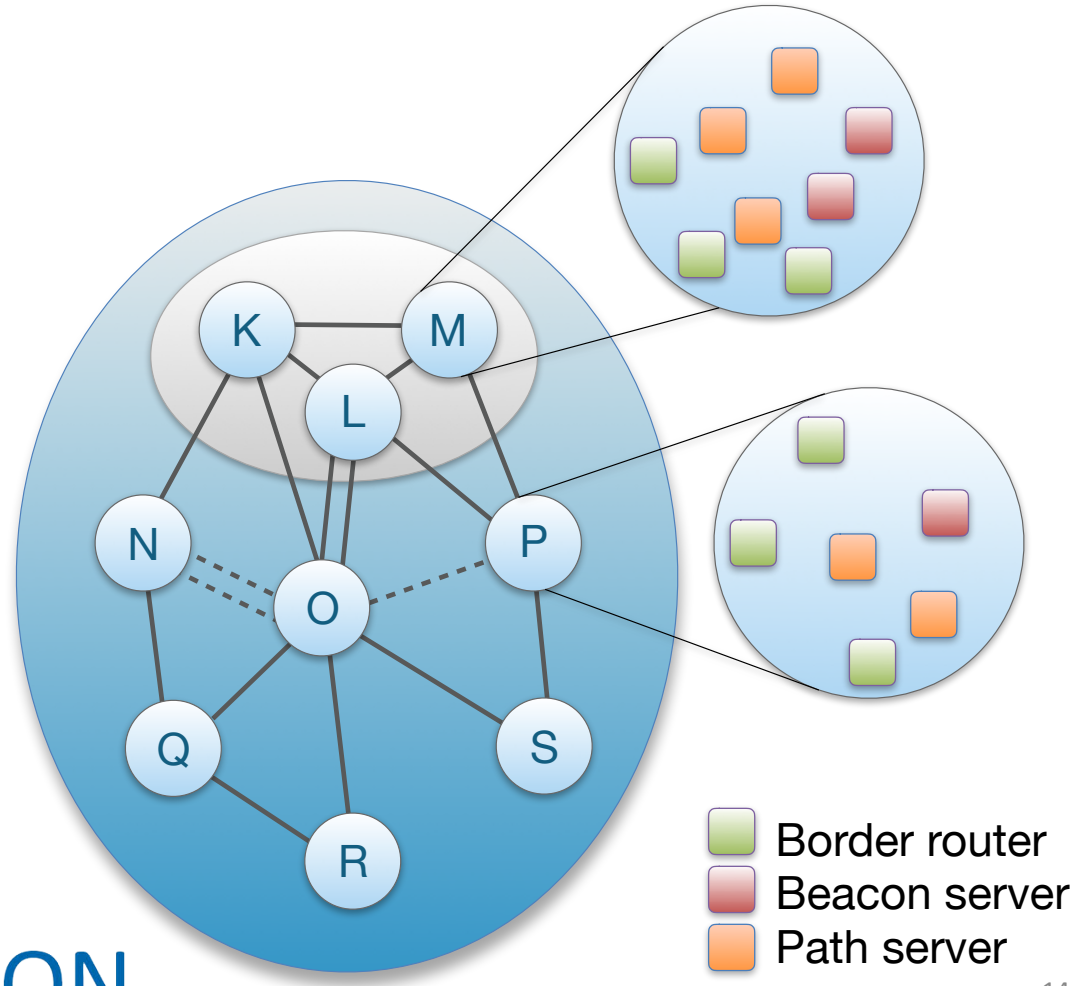

ETH zürich

SCION

Up-Path Segment Registration

- AS selects path segments to announce as up-path segments for local hosts
- Up-path segments are registered at local path servers

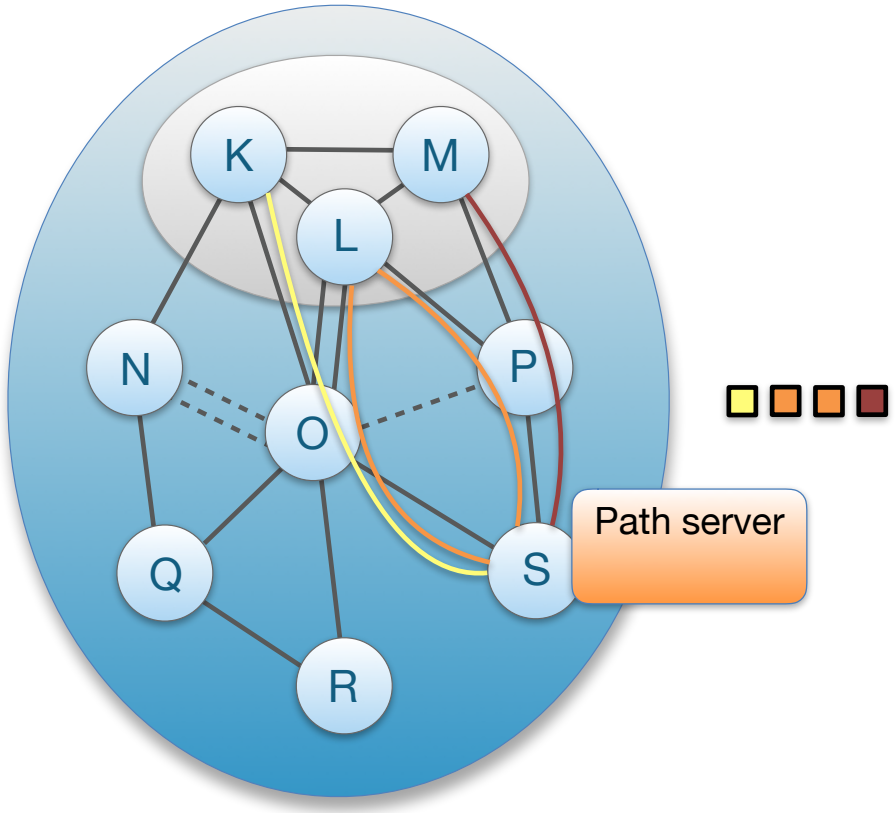

Down-Path Segment Registration

- AS selects path segments to announce as down-path segments for others to use to communicate with AS
- Down-path segments are uploaded to core path server in core AS

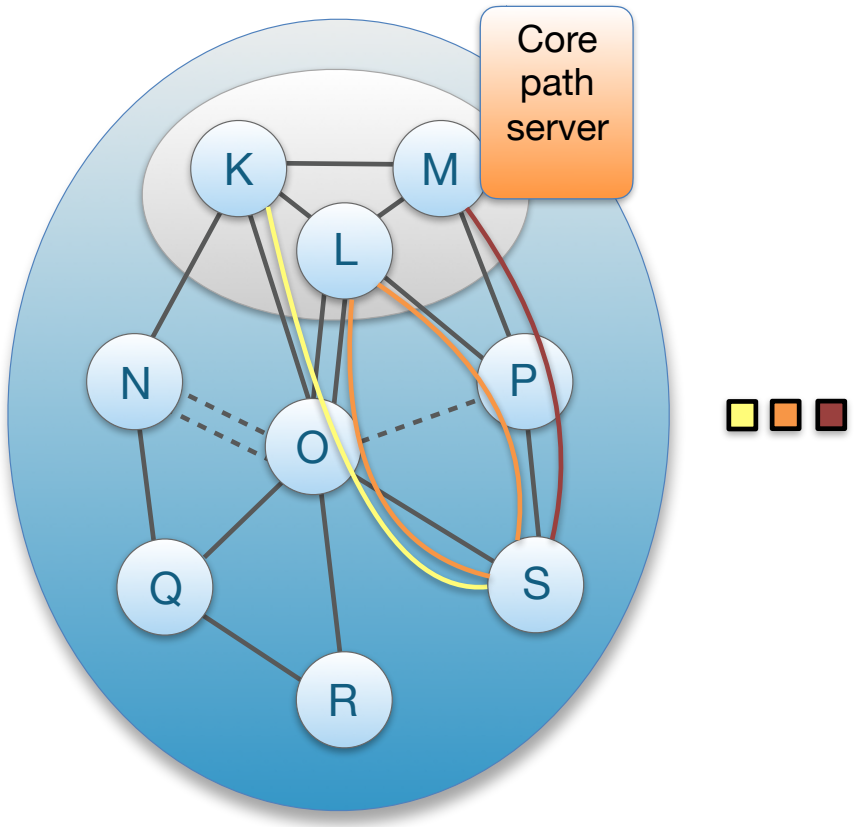

Ingress and Egress Interface Identifiers

- Each AS assigns a unique integer identifier to each interface that connects to a neighboring AS
- The interface identifiers identify ingress/egress links for traversing AS
- ASes use internal routing protocol to find route from ingress SCION border router to egress SCION border router
- Examples
  - Yellow path: L:4, O:3,6, R:1
  - Orange path: L:5, O:2,6, R:1

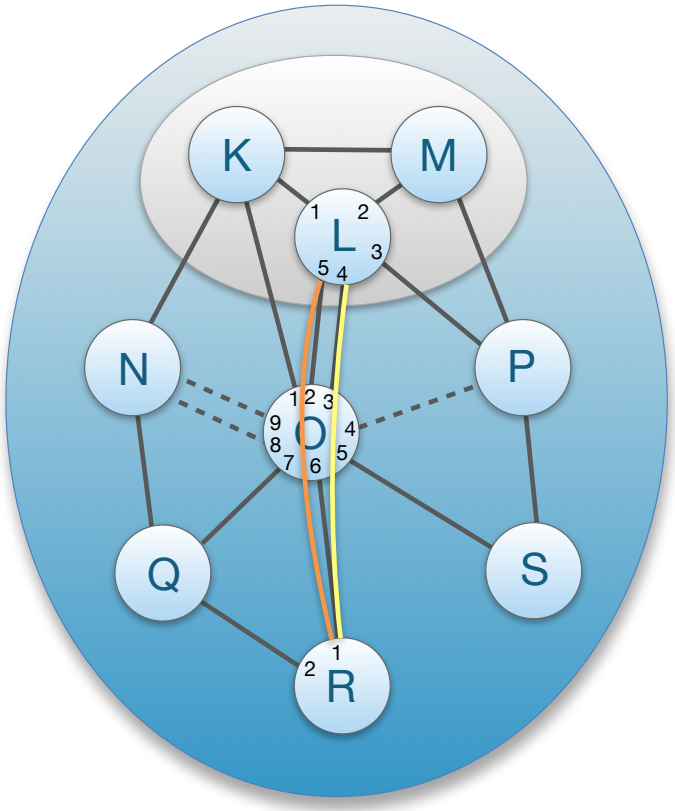

SCION Overview

- Control plane: How to find end-to-end paths?
  - Path exploration
  - Path registration
- Data plane: How to send packets
  - Path lookup
  - Path combination
- Deployment
- Demos

Path Lookup

- Steps of a host to obtain path segments
  - Host contacts RAINS server with a name
    H → RAINS: www.scion-architecture.net
    RAINS → H: ISD X, AS Y, local address Z
  - Host contacts local path server to query path segments
    H → PS: ISD X, AS Y
    PS → H: up-path, core-path, down-path segments
  - Host combines path segments to obtain end-to-end paths, which are added to packets

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

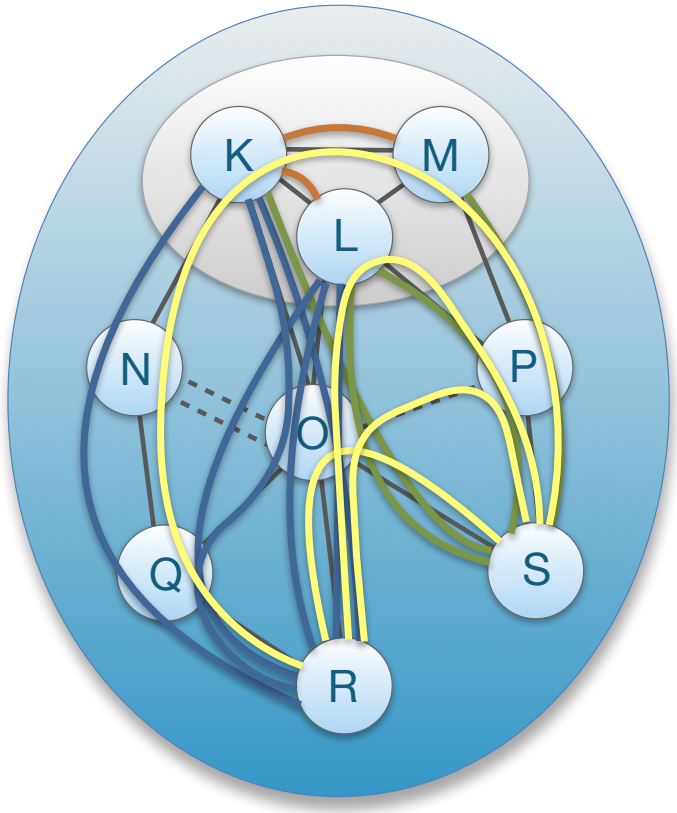

Path Lookup: Remote ISD

- 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

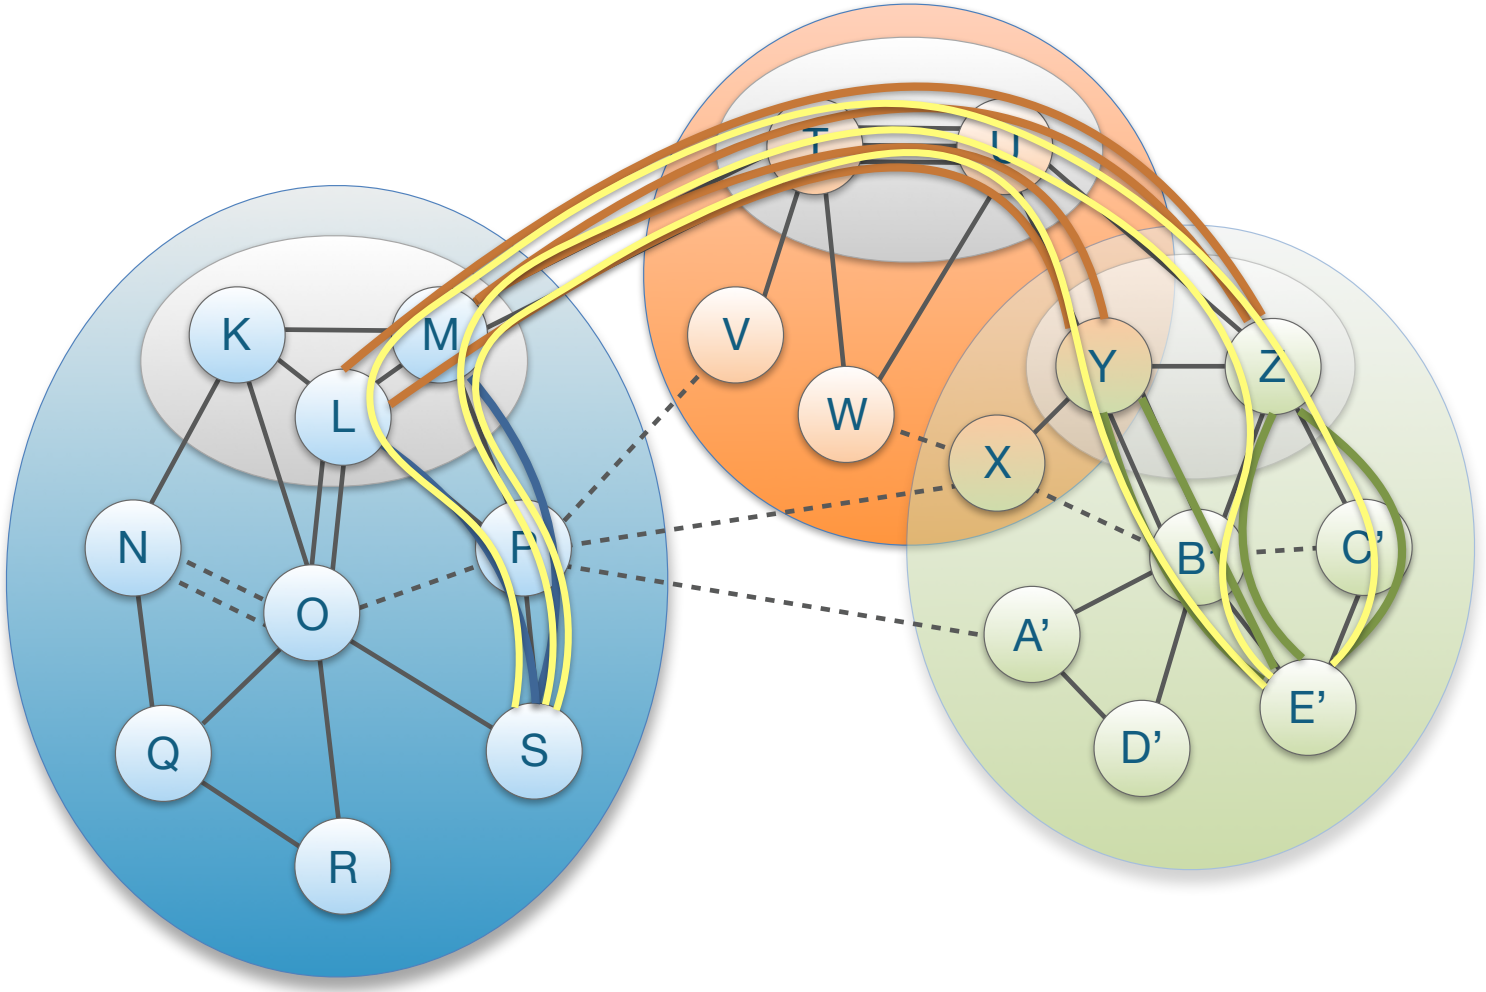

ETH zürich

Path Combination

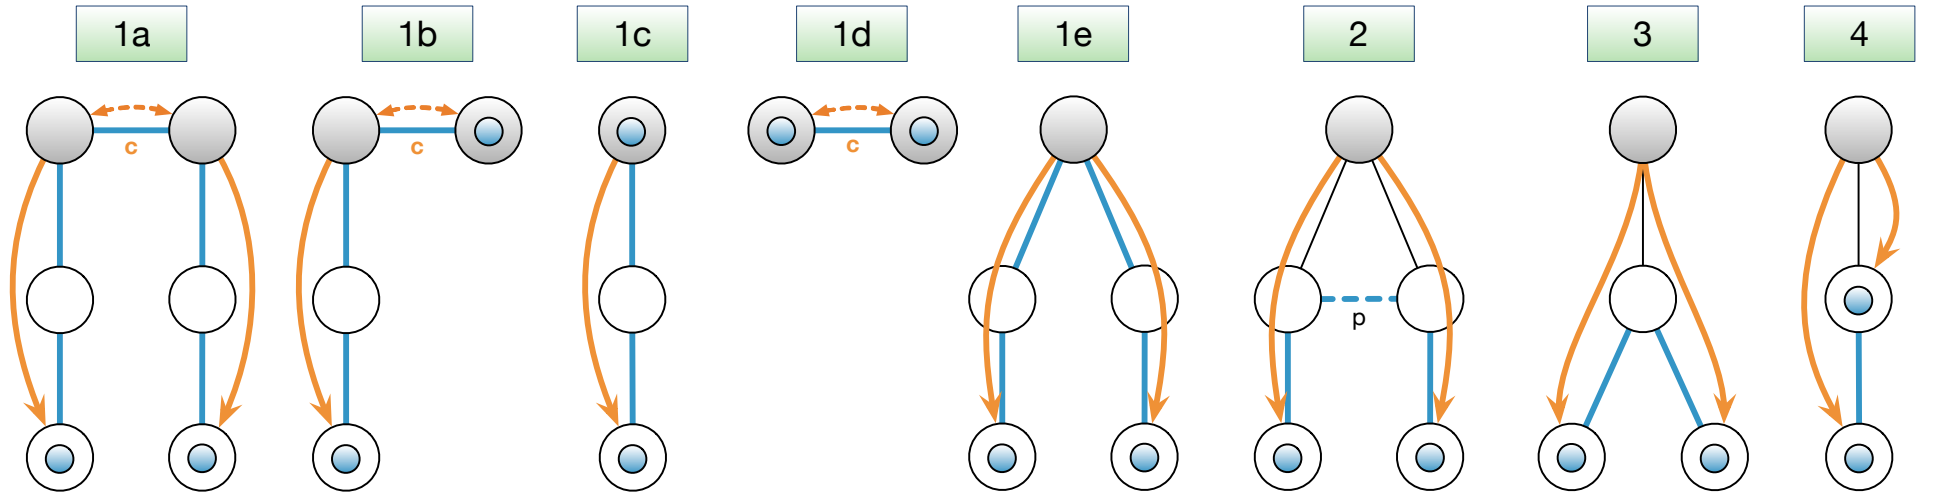

Control-plane path segments:

- Up- down-path segment
- Core-path segment

Data-plane paths:

- Regular path segment
- Peering link path segment

- Core AS
- Non-core AS
- Source/destination

ETH zürich

SCION

Path Combination Example (1)

- Core-segment combination:
  Up-path segment +
  core-path segment +
  down-path segment

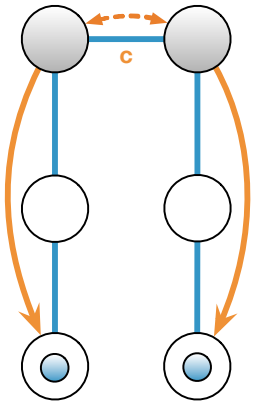

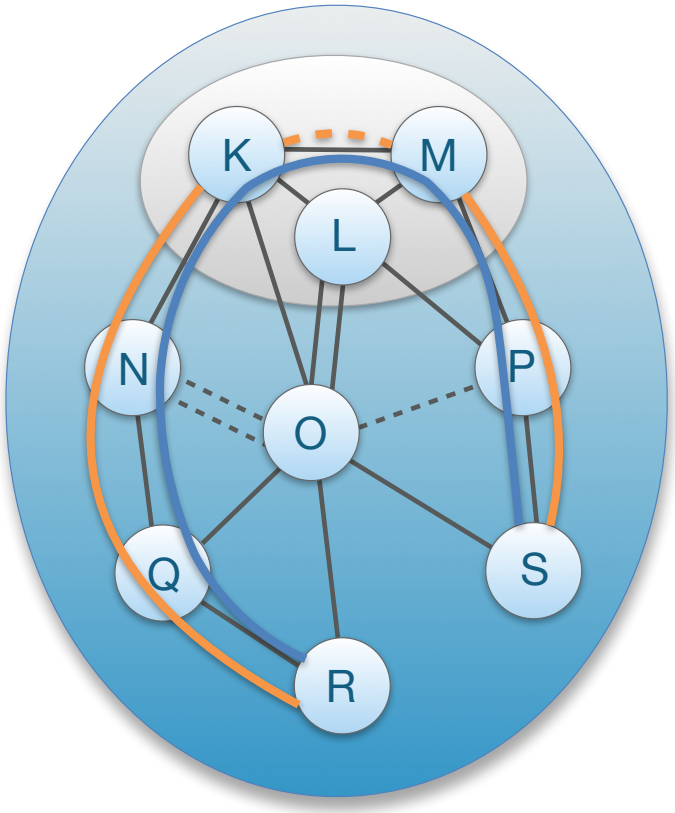

Path Combination Example (2)

- Peering shortcut: up-path segment and down-path segment offer same peering link

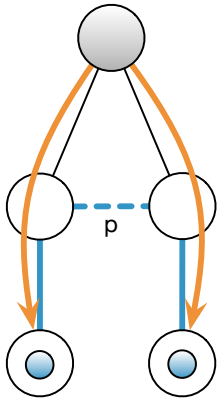

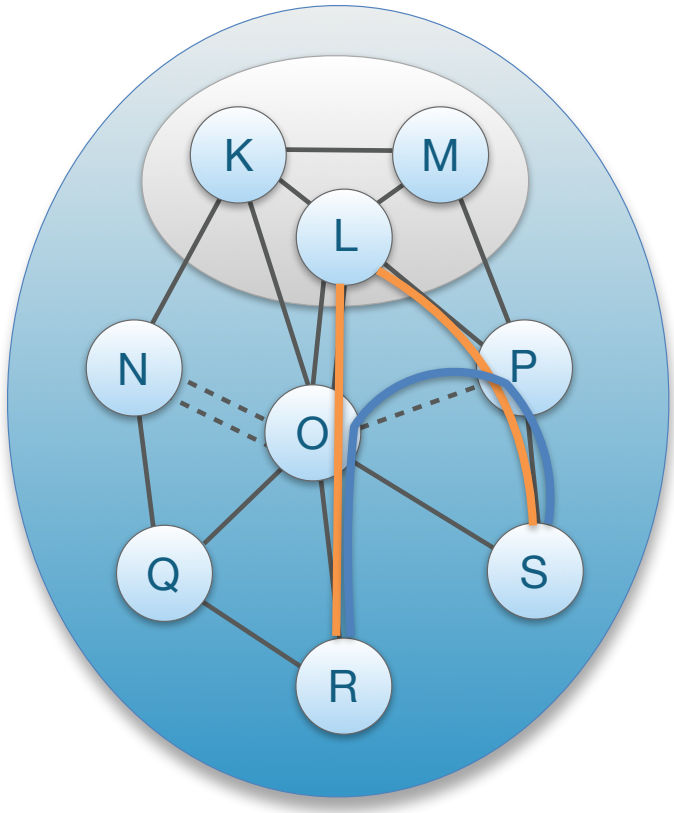

Path Combination Example (3)

- Peering shortcut: up-path segment and down-path segment offer same peering link

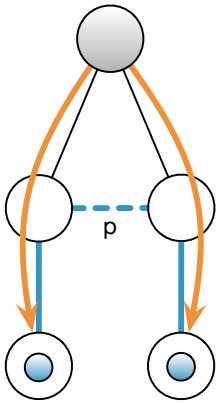

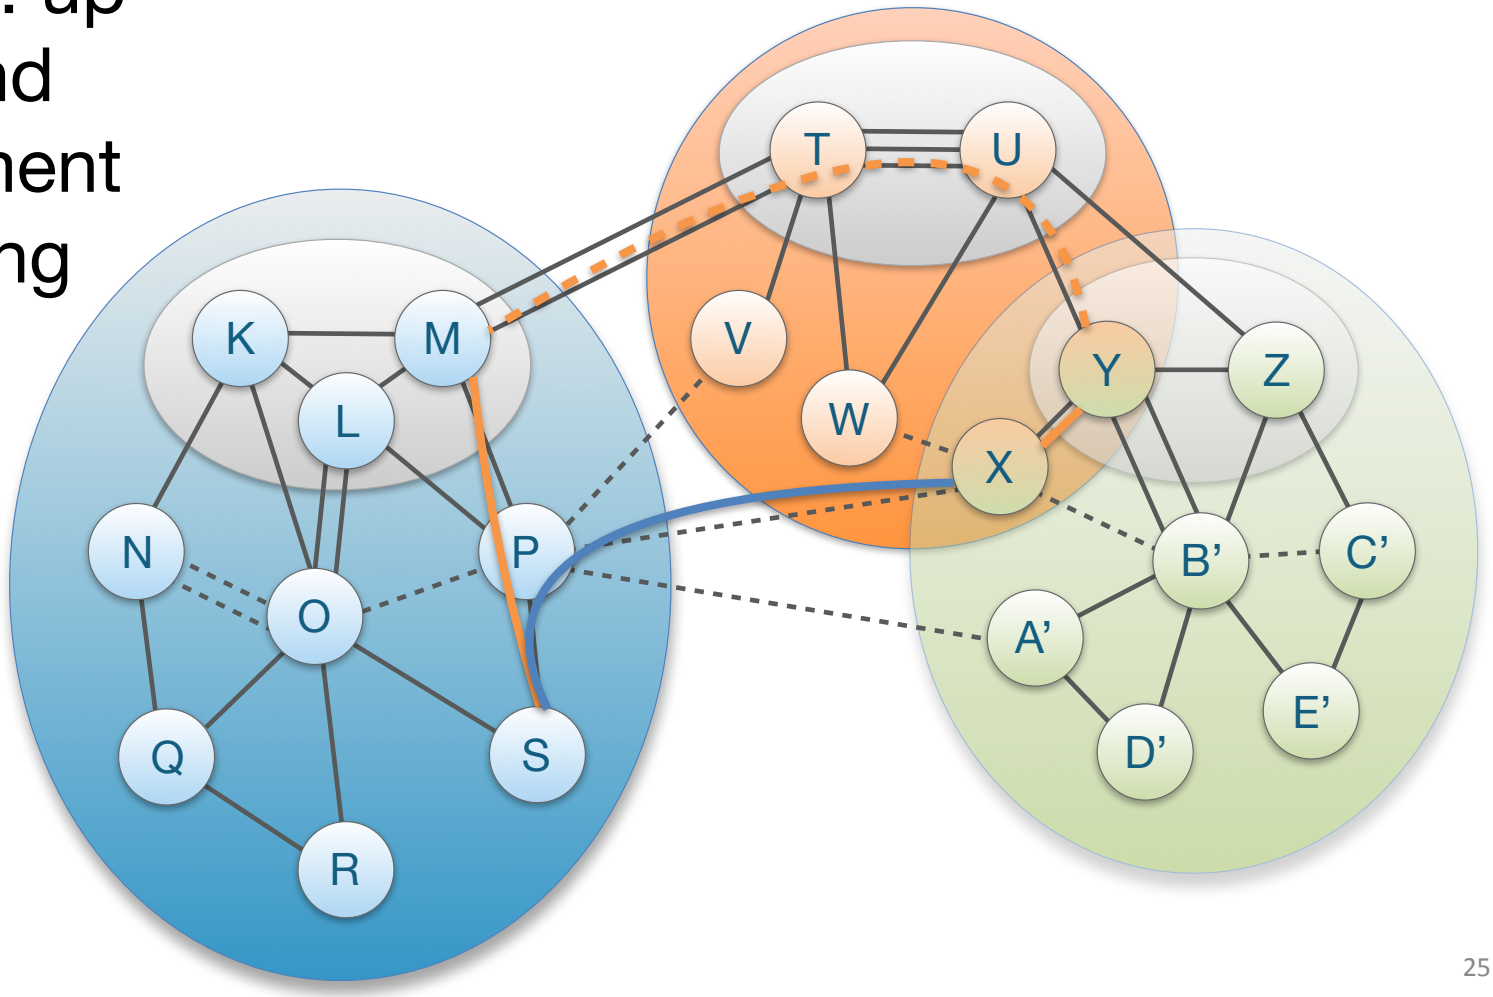

Path Combination Example (4)

- AS shortcut path through common AS on up-path and down-path segment

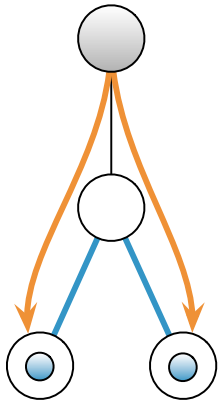

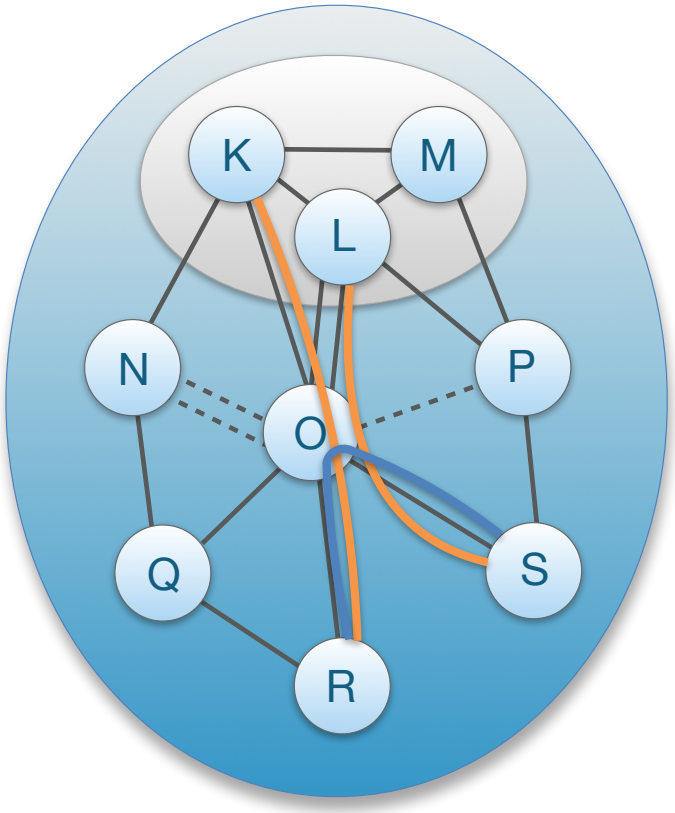

Path Construction

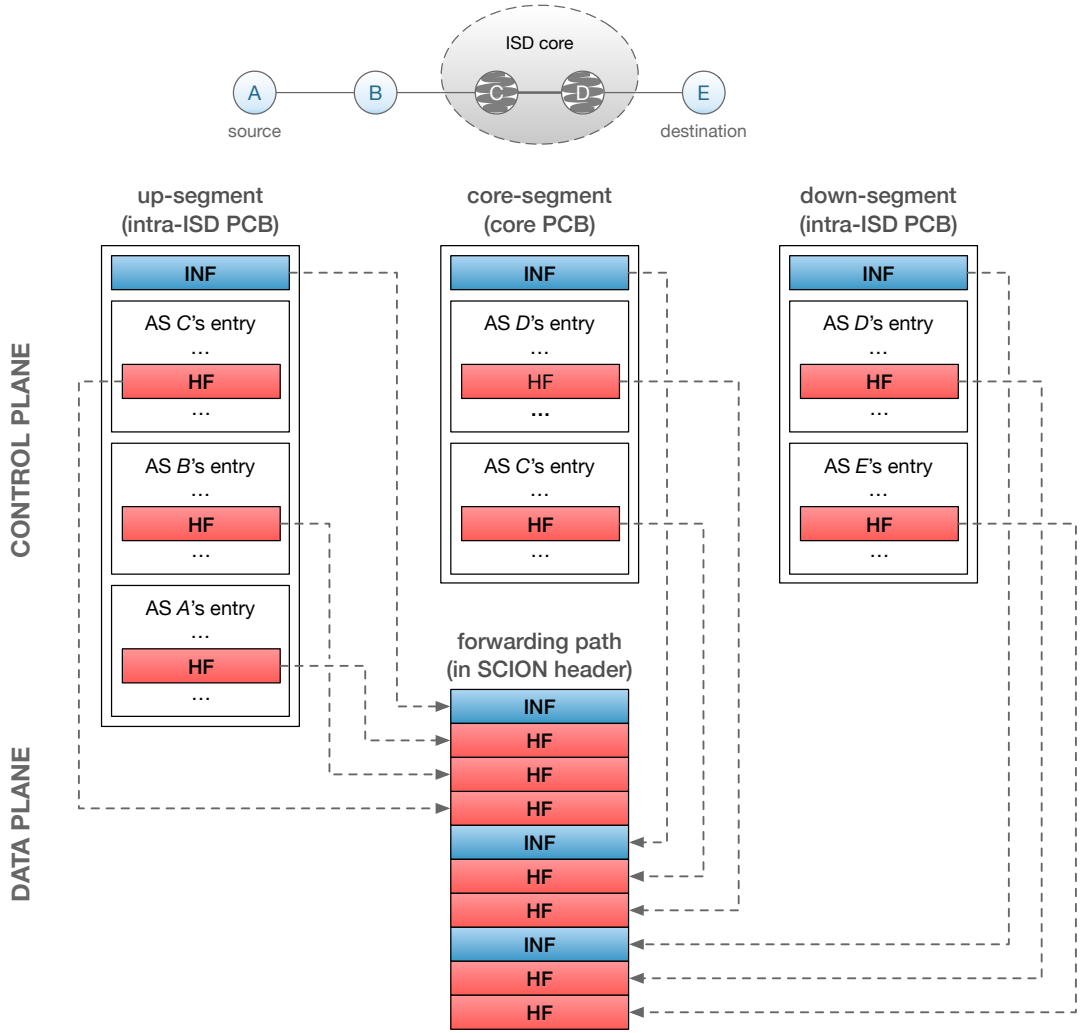

Path Encoding in Packet

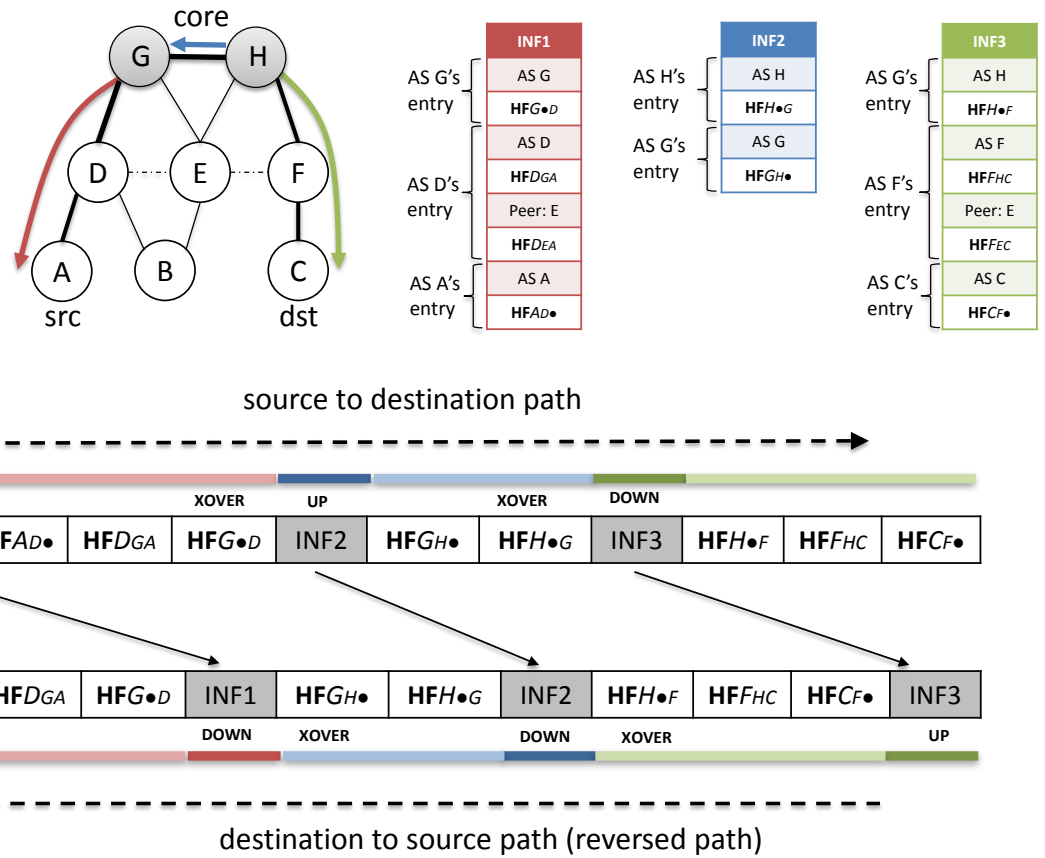

ETH zürich

SCION

Path Encoding in Packet

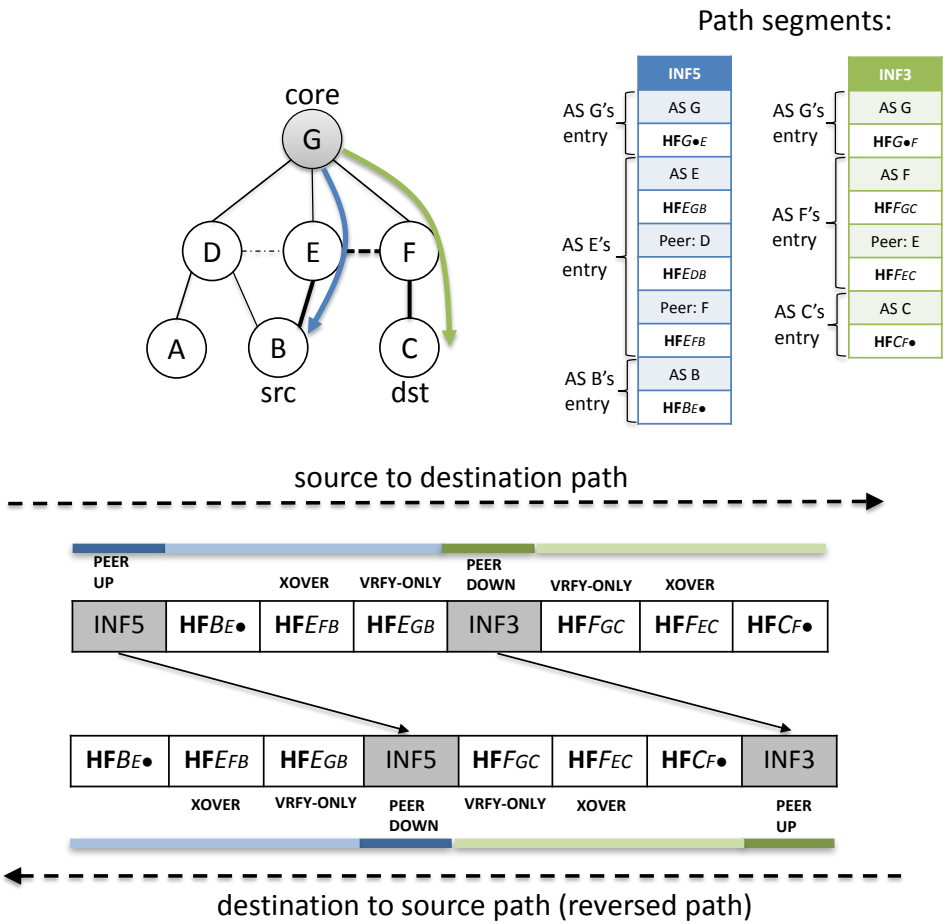

ETH zürich

SCION

Hop Field MAC Verification

- Message Authentication Code (MAC) computation and verification of Hop Field MAC value based on local AS secret key
  - Key is not shared with any external entity
- Computation: MAC_K(Timestamp, $Flags'_{HF}$, ExpTime, Ingress, Egress, HF')
  - HF' is hop field of previous AS
- In most cases, HF' size is 8 bytes, so MAC computation can be done over 128 bits: with CMAC and AES, only a single encryption operation is needed
- With AESni HW crypto, only ~50 cycles are needed to compute MAC!
  - Note that a DRAM memory lookup takes ~200 cycles
  - AES operation requires less energy than TCAM lookup
  - Thus, SCION forwarding can be faster and require less energy than IP forwarding

SCION Summary

- Complete re-design of network architecture resolves numerous fundamental problems
  - BGP protocol convergence issues
  - Separation of control and data planes
  - Isolation of mutually untrusted control planes
  - Path control by senders and receivers
  - Simpler routers (no forwarding tables)
  - Root of trust selectable by each ISD
- An isolation architecture for the control plane, but a transparency architecture for the data plane.

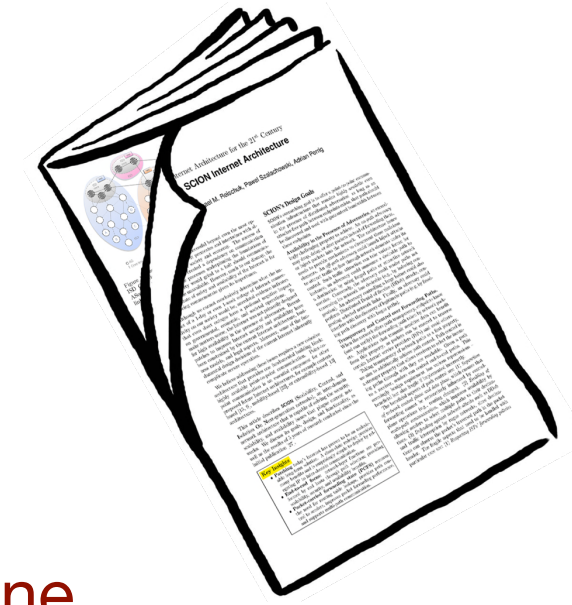

SCION Overview

- Control plane: How to find end-to-end paths?
  - Path exploration
  - Path registration
- Data plane: How to send packets
  - Path lookup
  - Path combination
- Deployment
- Demos

ISP Deployment (Core AS)

- Core AS duties
  - Manage and distribute the ISD's TRC
  - Sign TRCs of neighboring ISDs and endorse other ISDs
  - Maintain a list of all recognized ISDs
  - Issue certificates to all ASes in the ISD
  - Provide connectivity to neighboring ISDs
  - Generate and disseminate inter-ISD path-segment construction beacons (PCBs), also called core PCBs
  - Generate and disseminate intra-ISD PCBs
  - Provide highly available services: beacon, name (RAINS), path, certificate, SIBRA, and time servers

ISP Deployment (~~Core~~ AS)

- ~~Core~~ AS duties
  - ~~Manage and distribute the ISD's TRC~~
  - ~~Sign TRCs of neighboring ISDs and endorse other ISDs~~
  - ~~Maintain a list of all recognized ISDs~~
  - ~~Issue certificates to all ASes in the ISD~~
  - Provide connectivity to neighboring ~~ISDs~~ ASes
  - ~~Generate and disseminate inter-ISD path-segment construction beacons (PCBs), also called core PCBs~~
  - ~~Generate and~~ disseminate ~~intra-ISD~~ PCBs
  - Provide ~~highly~~ available services: beacon, name (RAINS), path, certificate, SIBRA, ~~and time servers~~

Leaf AS Deployment

- Leaf AS deployment tasks
  - Obtain AS certificate from core AS
  - Deploy servers: beacon, name (RAINS), path, certificate, SIBRA
- One single legacy PC suffices, e.g., attached to border router

Deployment @ ETH

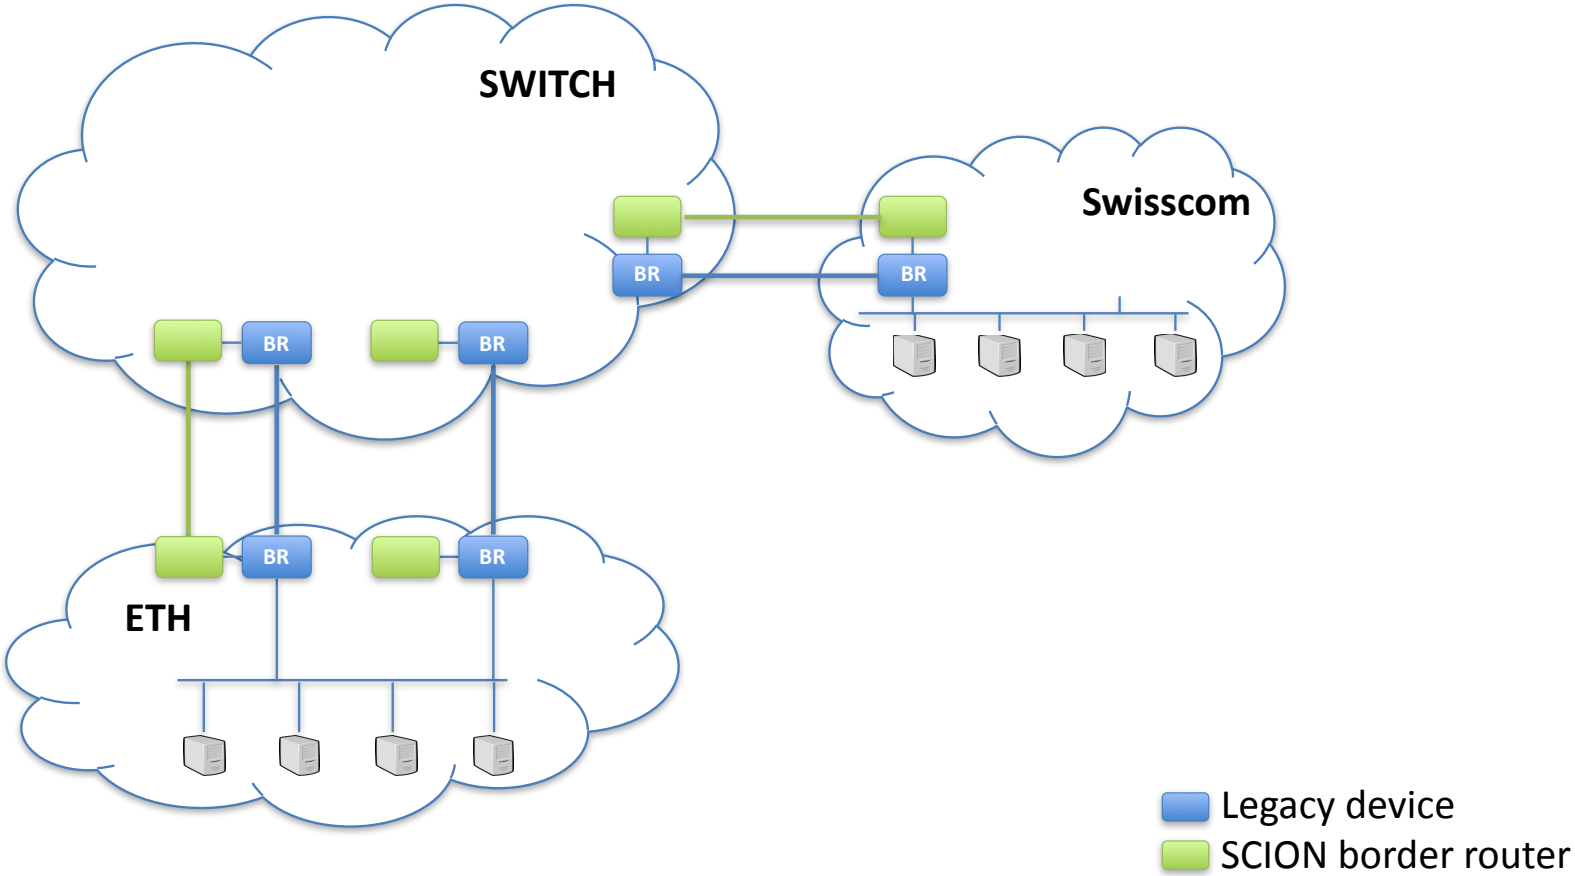

ETH zürich

SCION

Use Case: IoT Protection through Default Off

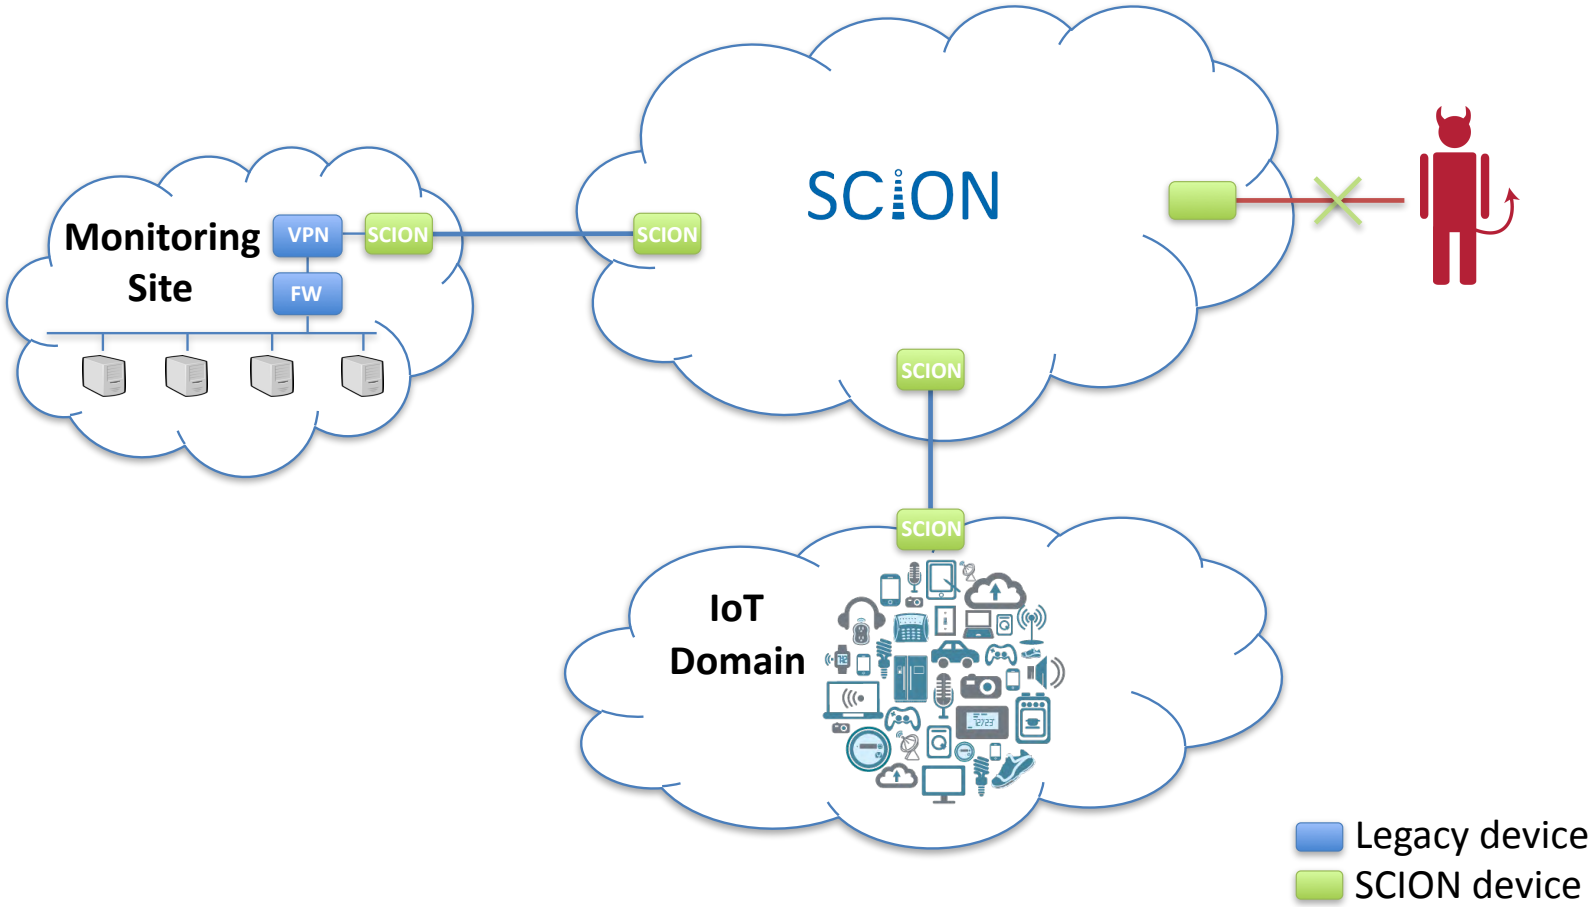

ETH zürich

SCION

Use Case: VPN-based Deployment

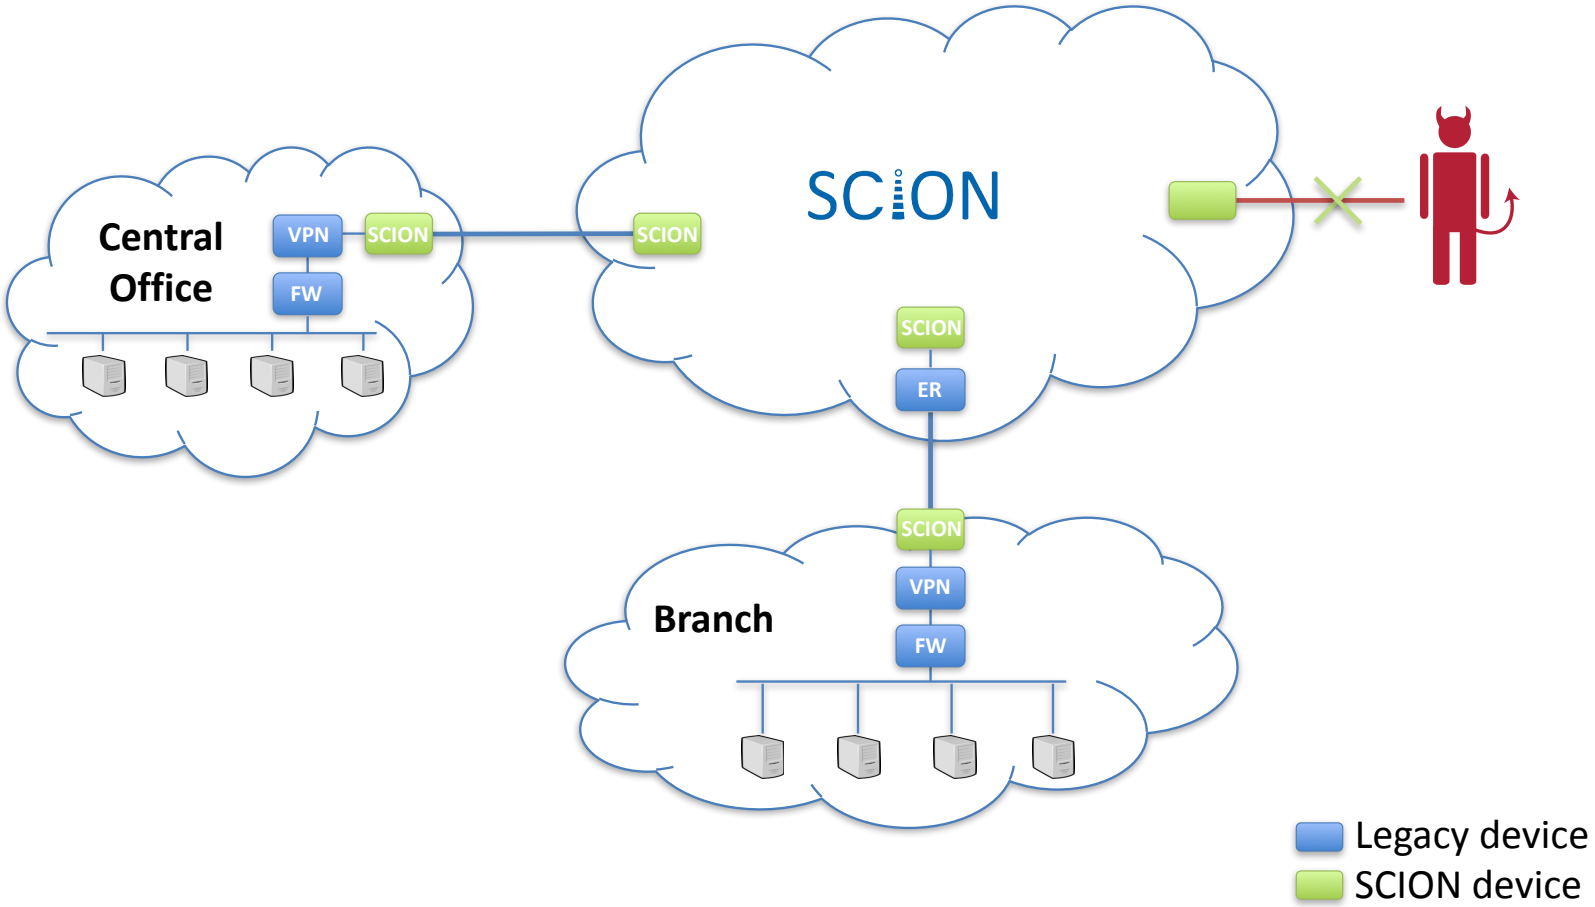

SCION-IP Gateway (SIG) Deployment

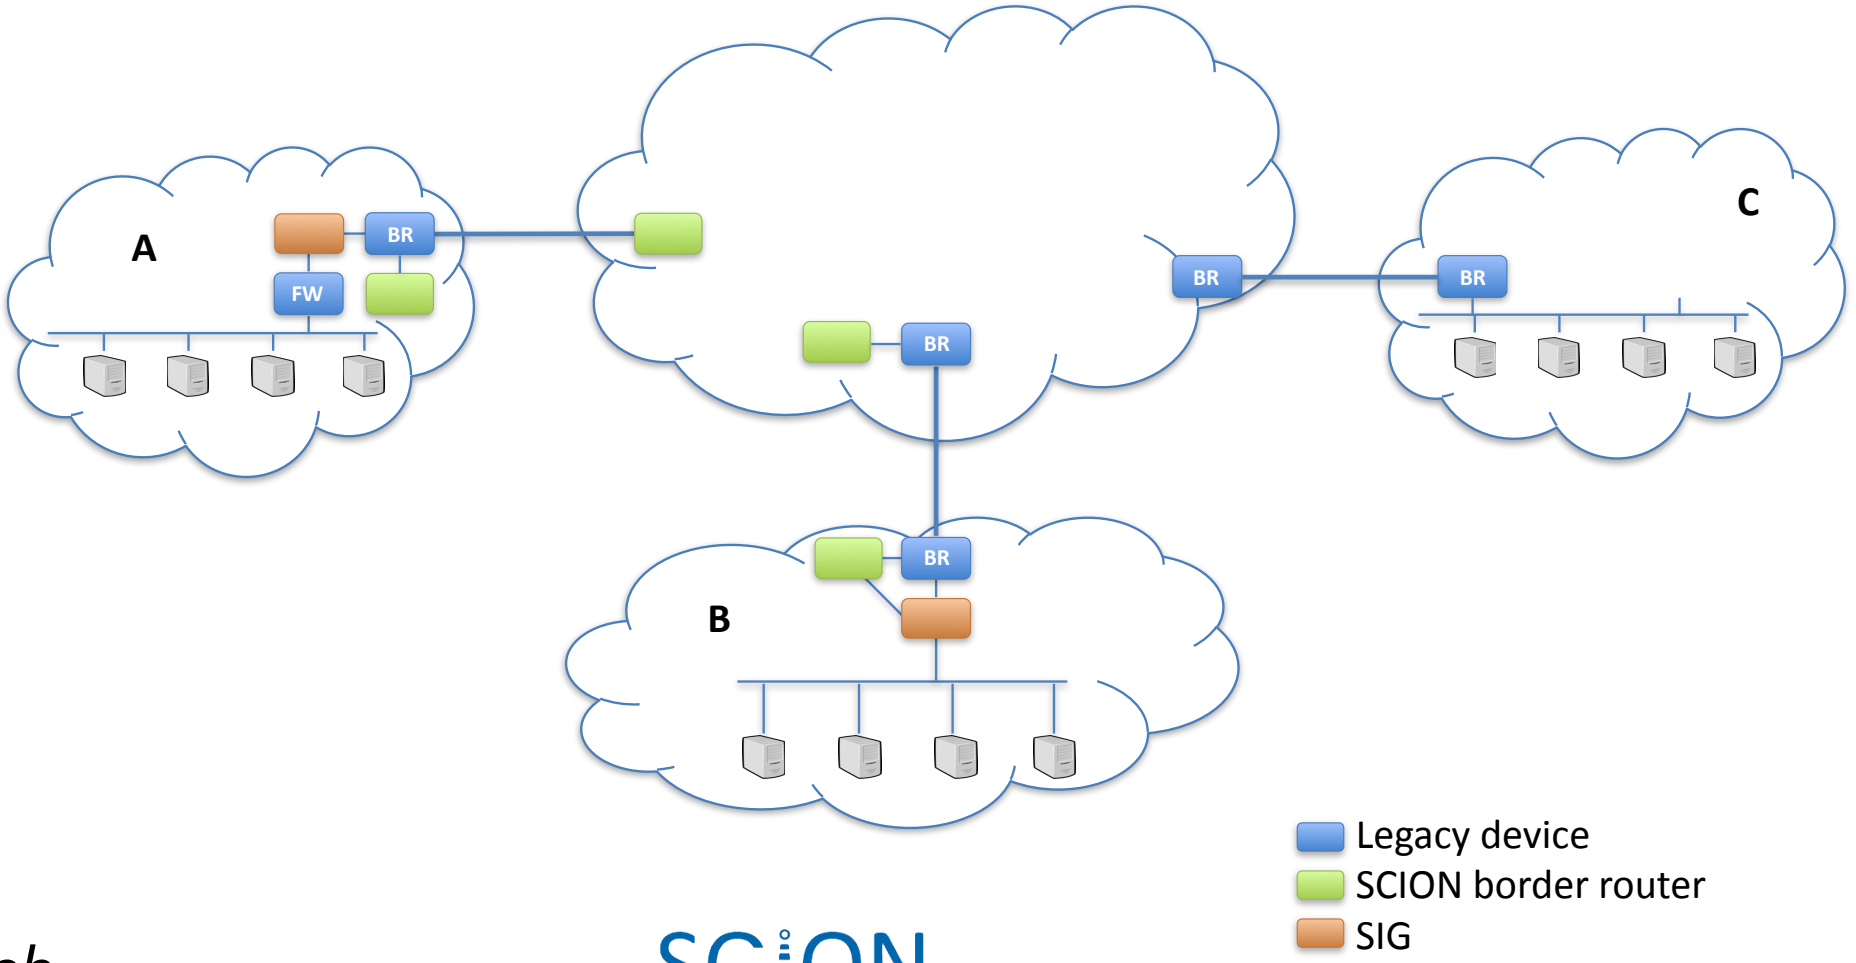

ETH zürich

SCION

SCIONLab

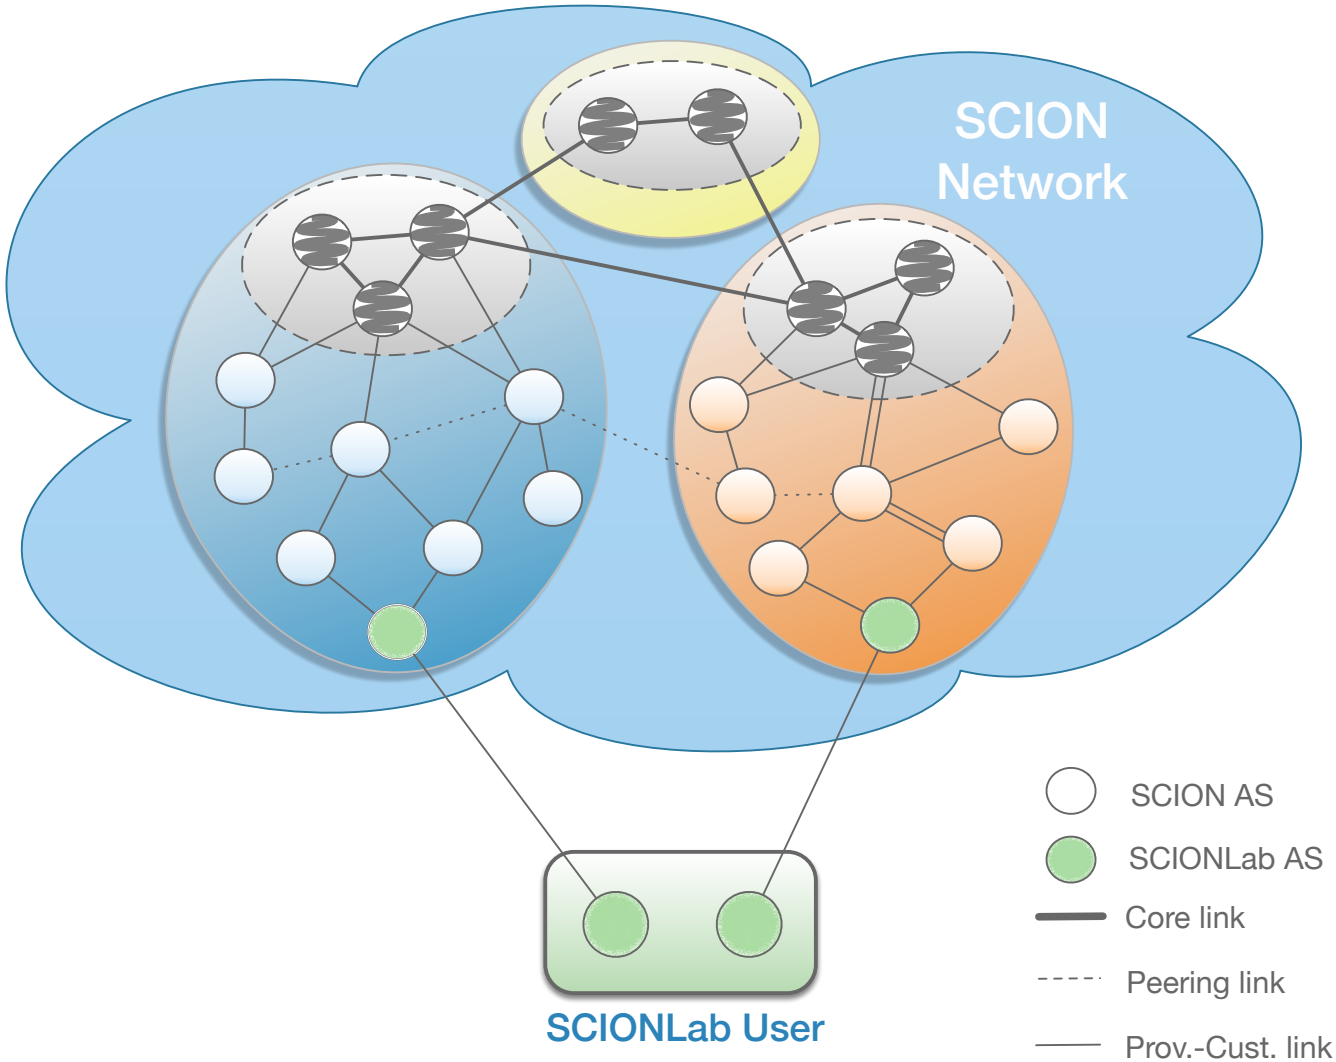
SCION
Network
SCIONLab User
SCION AS
SCIONLab AS
Core link
Peering link
Prov.-Cust. link

SCIONLab Network

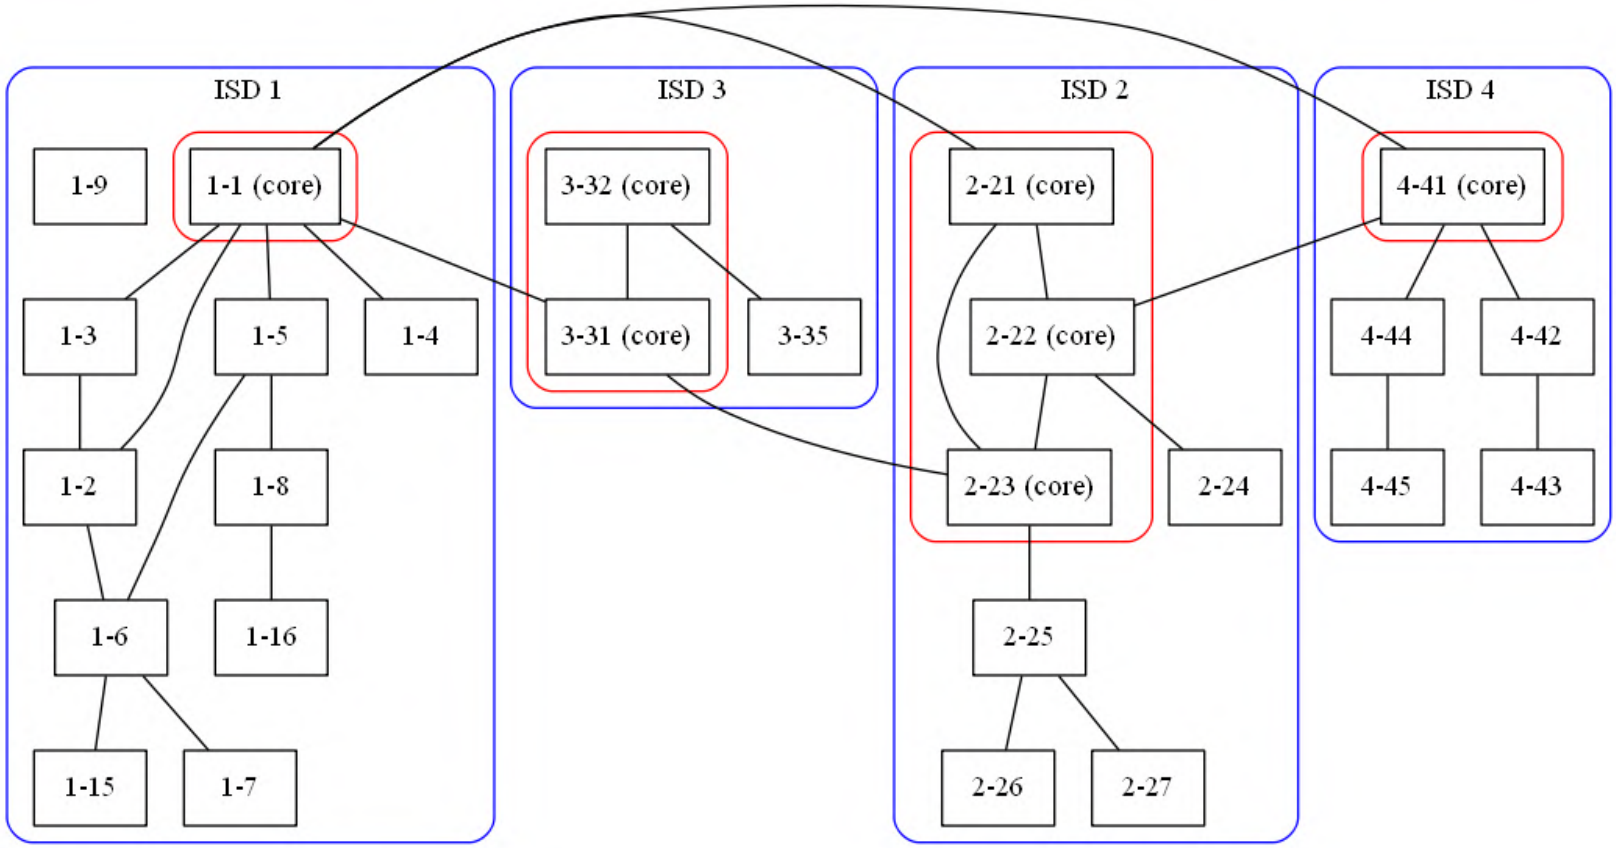

ETH zürich

SCION

SCION Visualization System

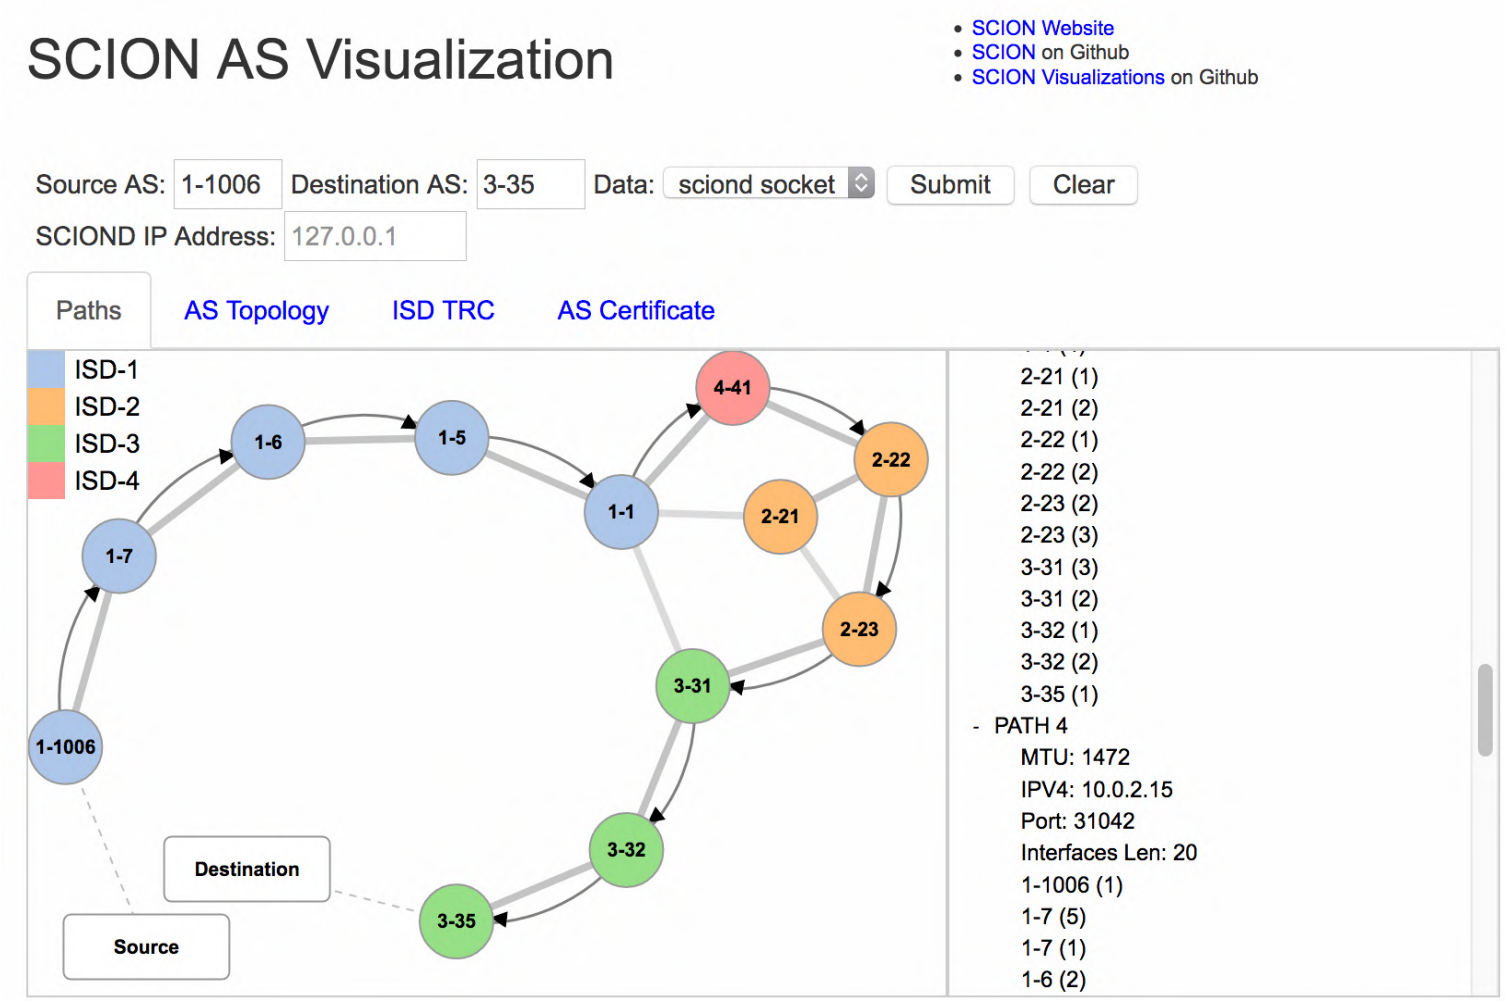

ETH zürich

SCION

Application: IoT Access

ETH zürich

SCION

AS Router Monitoring with Prometheus

Swiss SCION Network

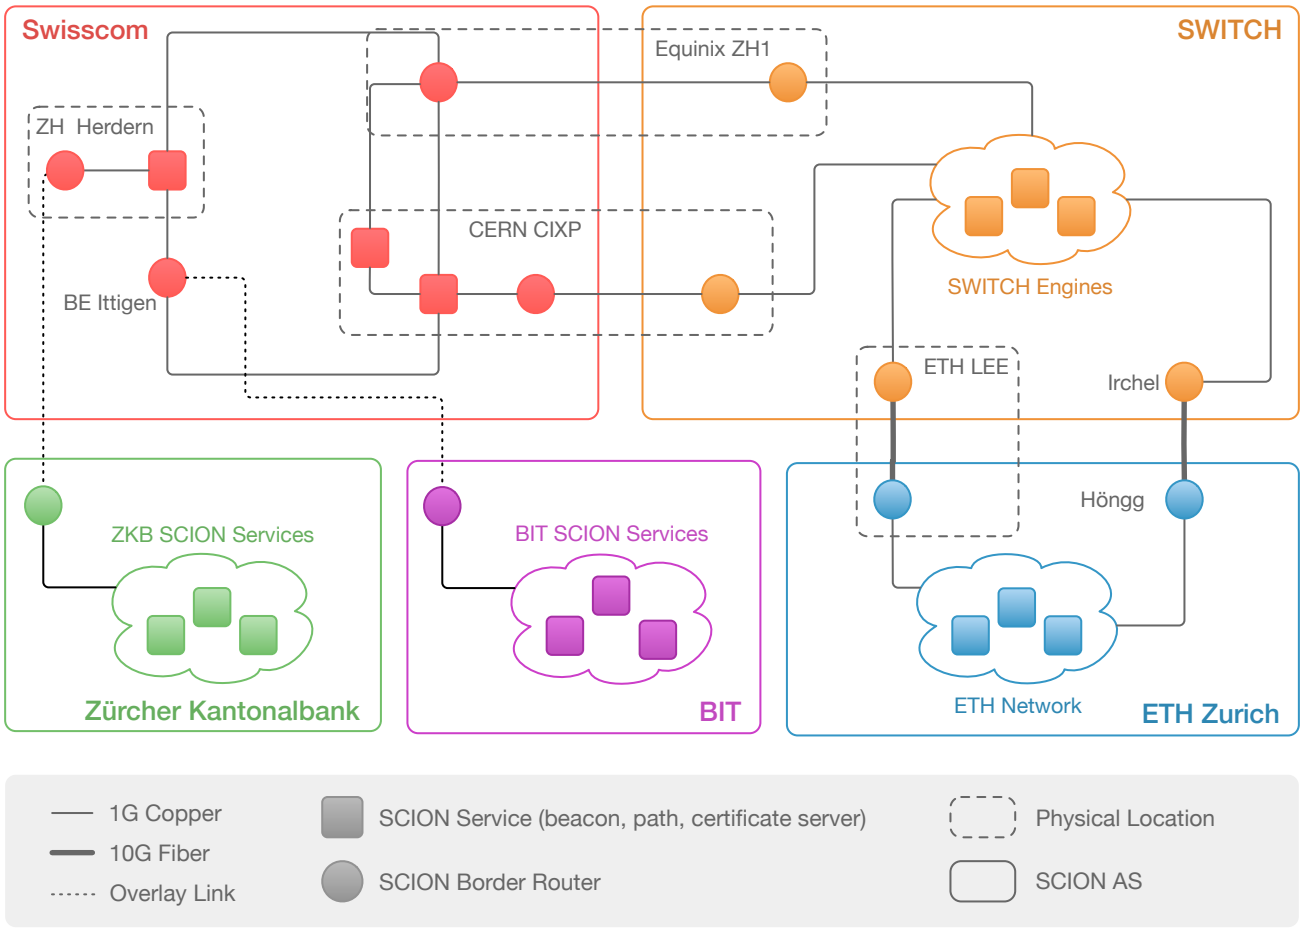

Growing Global Testbed

- Over 40 deployed SCION routers and servers

SCION AS runs on ODROID and Raspberry Pi

ODROID
Hardkernel

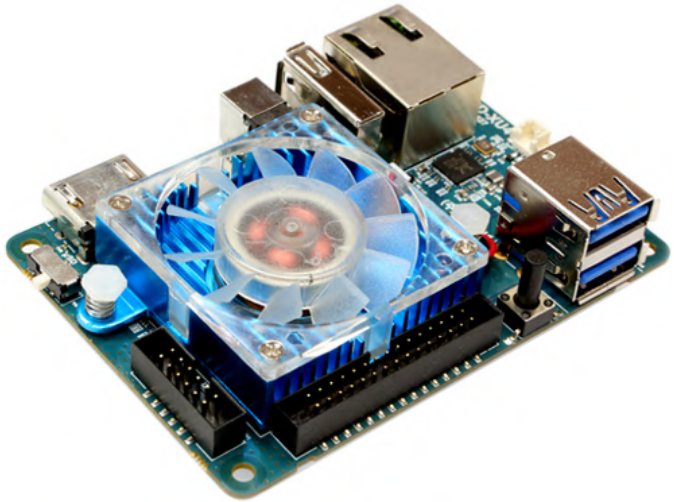

Raspberry Pi

Demos

- Fast failover and multipath demo
- Path control and geofencing demo

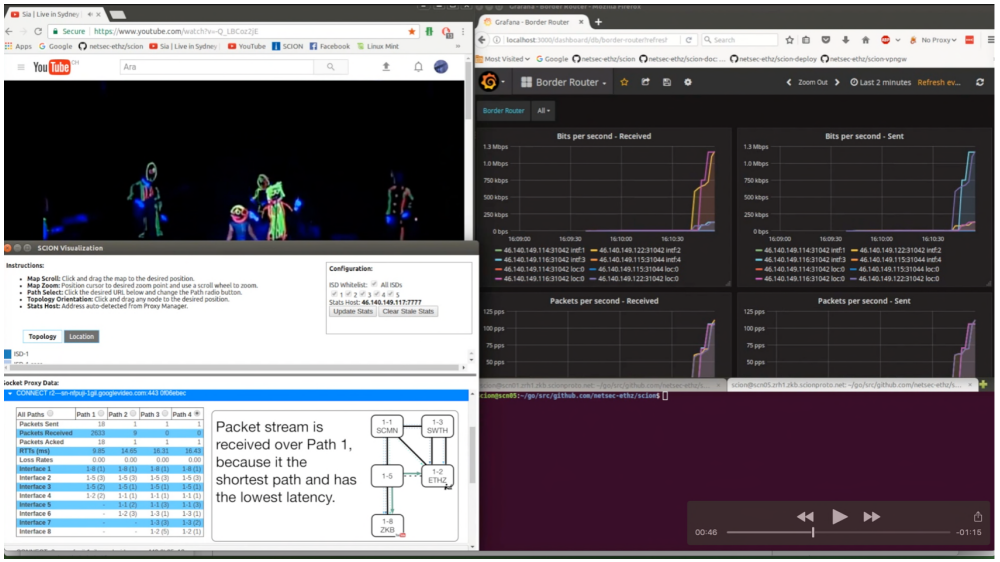

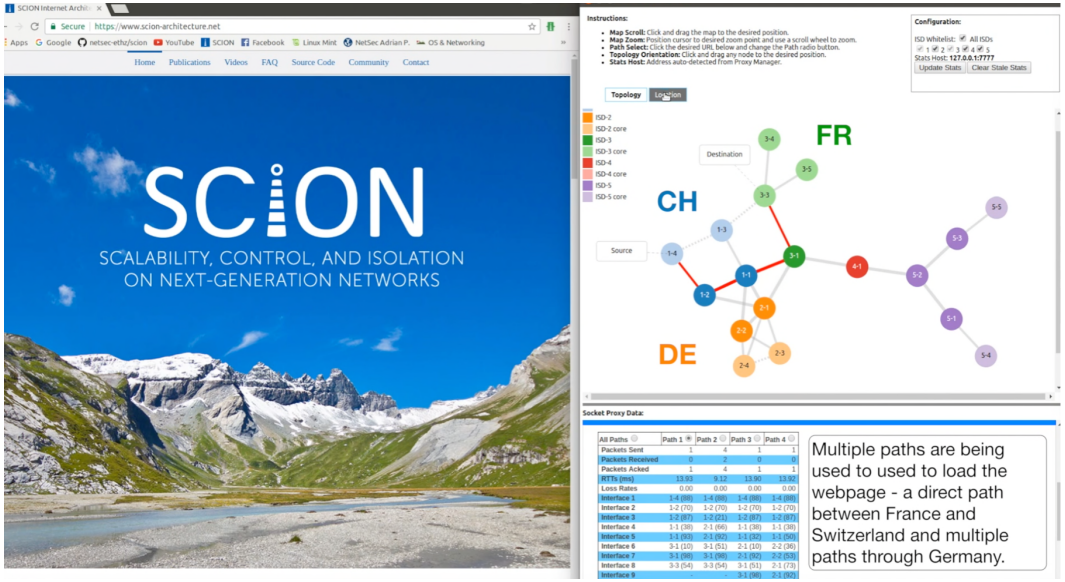

Belief that Internet is Immutable

- Evidence appears overwhelming that Internet is immutable: IPv6, BGPSEC, DNSSEC, etc.
- However, benefits are limited, esp. for early deployers
- Our goal: provide many benefits, even for early adopters, such that one cannot turn back

Conclusions

- SCION is a secure Internet architecture that we can start using today
- Open source
- Numerous opportunities for researchers
  - Multipath routing architecture offers multitude of path choices for meaningful diverse path selection
  - Security: routing, DDoS, source authentication
  - Next-generation PKI architecture
- Natural quality scalability with increasing global adoption

SCION Projekt Team

- Netsec: Daniele Asoni, Chen Chen, Laurent Chuat, Sergiu Costea, Sam Hitz, Tobias Klausmann, Tae-Ho Lee, Chris Pappas, Adrian Perrig, Benjamin Rotenberger, Stephen Shirley, Jean-Pierre Smith, Pawel Szalachowski, Brian Trammell, Ercan Ucan
- Infsec: David Basin, Tobias Klenze, Christoph Sprenger, Thilo Weghorn
- Programming Methodology: Marco Eilers, Peter Müller

ETH zürich

anapaya systems

www.anapaya.net

Additional Information

- https://www.scion-architecture.net
  - Book
  - Papers
  - Videos
  - Newsletter signup
- https://www.anapaya.net
  - Commercializing SCION equipment
- https://github.com/netsec-ethz/scion