SCION
SCALABILITY, CONTROL, AND ISOLATION
ON NEXT-GENERATION NETWORKS

A Next-generation Secure Internet for the 21st Century

Network Security Group, ETH Zürich

ETH zürich

SCION

SCION: Next-generation Internet Architecture

- Secure by design
- Path-aware networking: sender knows packet's path
  - Enables geo-fencing
- Highly available communication
- Multi-path communication
  - Caution: use is highly addictive!
- BGP-free Internet communication
- Better scalability than current Internet
- Improved network operation
  - Higher network utilization
  - Advanced traffic engineering

SCION Overview on 1 Slide

- Path-based network architecture
- Control plane: creates paths and disseminates them via path server infrastructure
- Data plane
  - Packets include path segments
  - Stateless routers verify integrity of forwarding information and send packet along to next AS hop

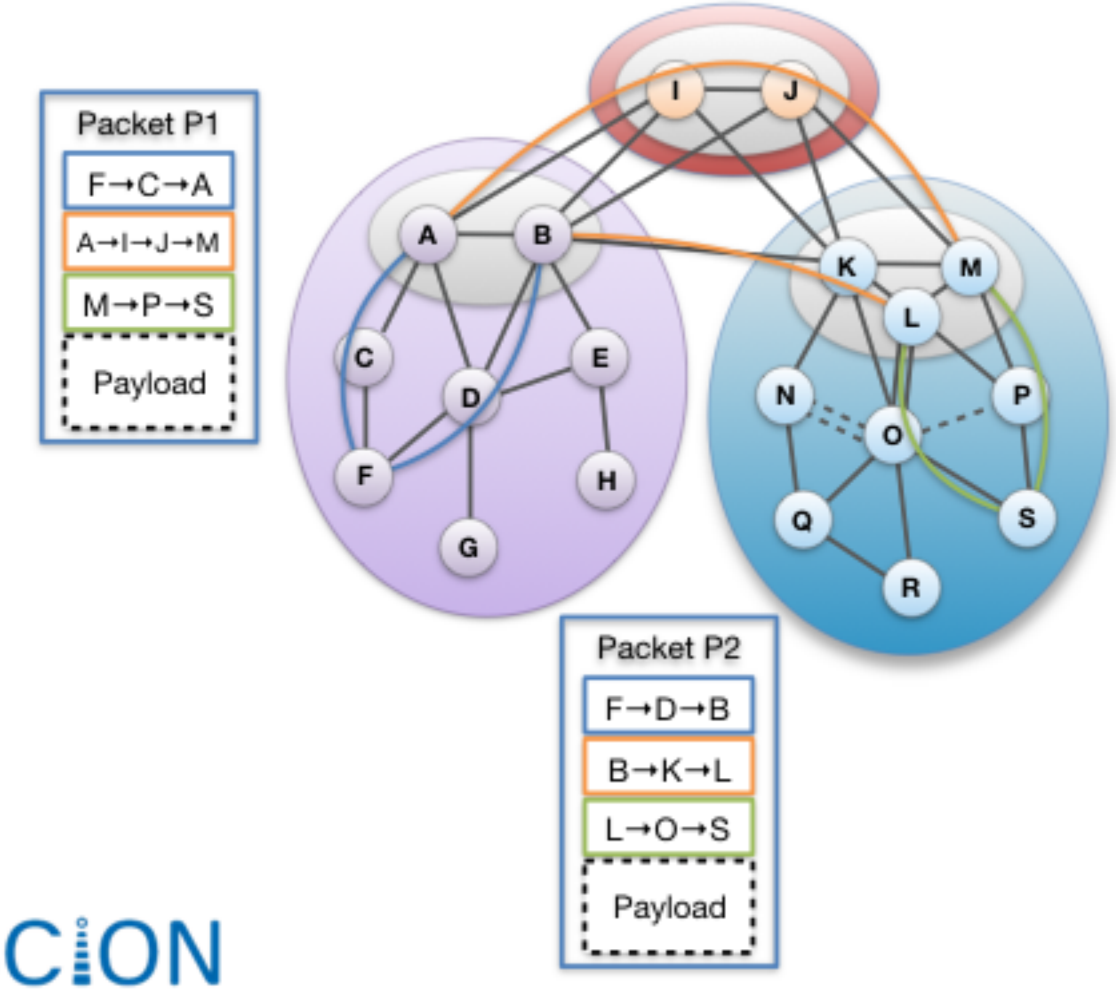

Packet P1

F→C→A

A→I→J→M

M→P→S

Payload

Packet P2

F→D→B

B→K→L

L→O→S

Payload

Deployment @ ETH, SWITCH, Swisscom

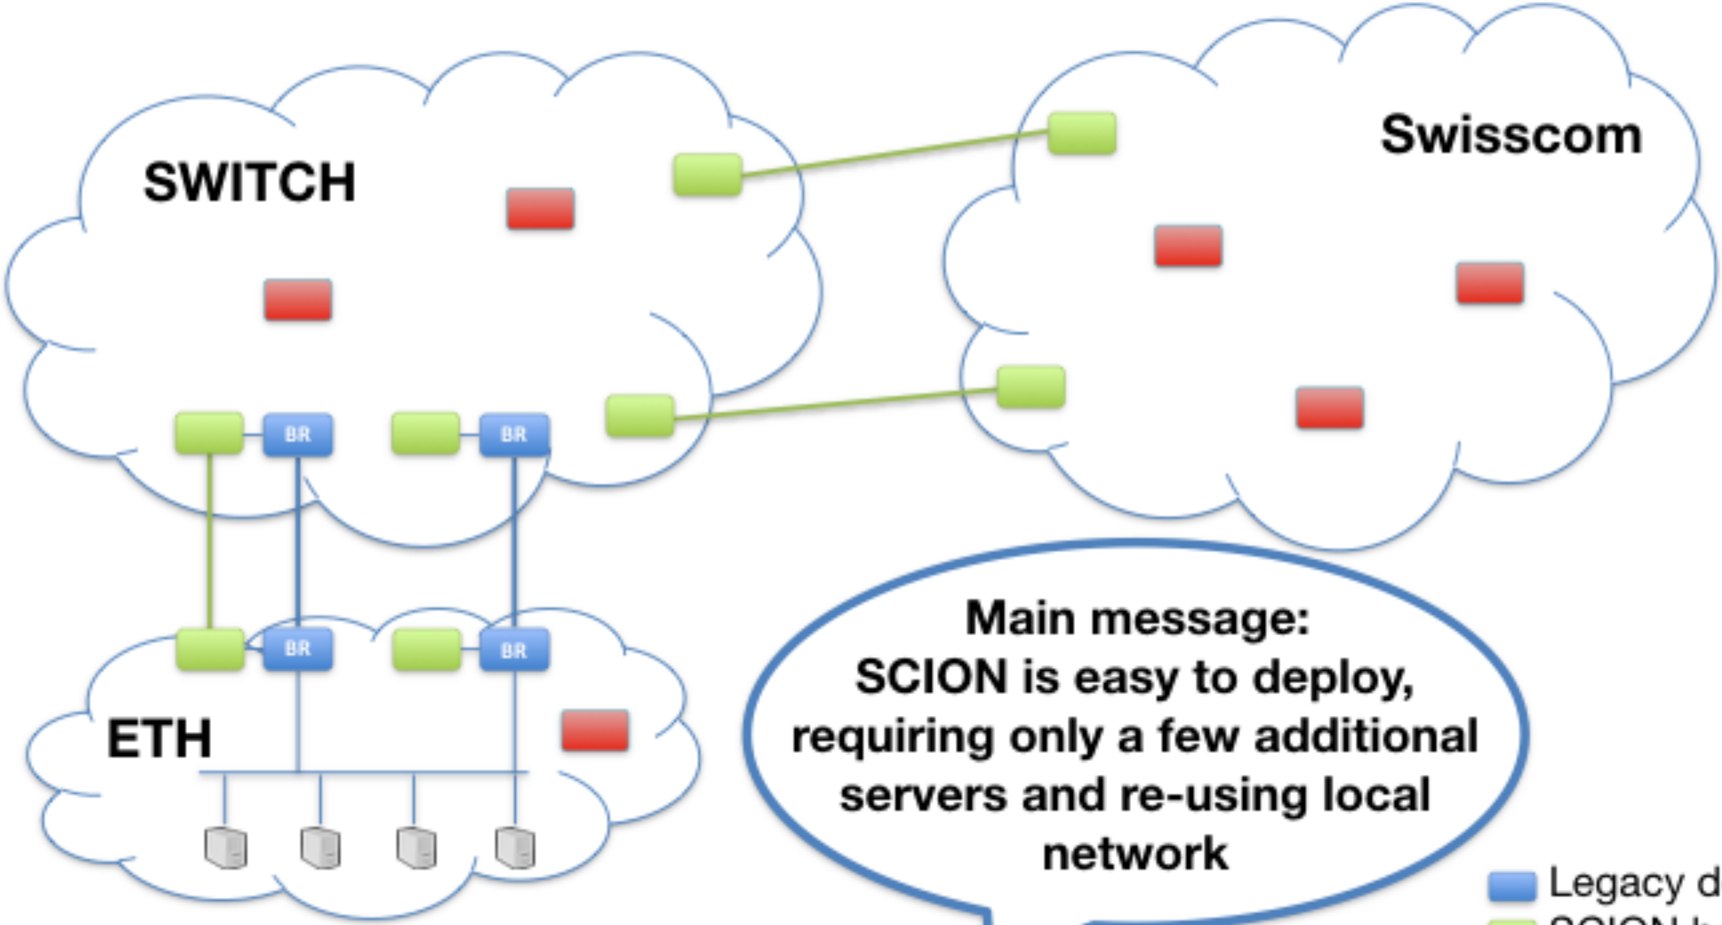

Main message:
SCION is easy to deploy, requiring only a few additional servers and re-using local network

🟦 Legacy device
🟩 SCION border router
🟥 SCION services

ETHzürich

SCION

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 aspects

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

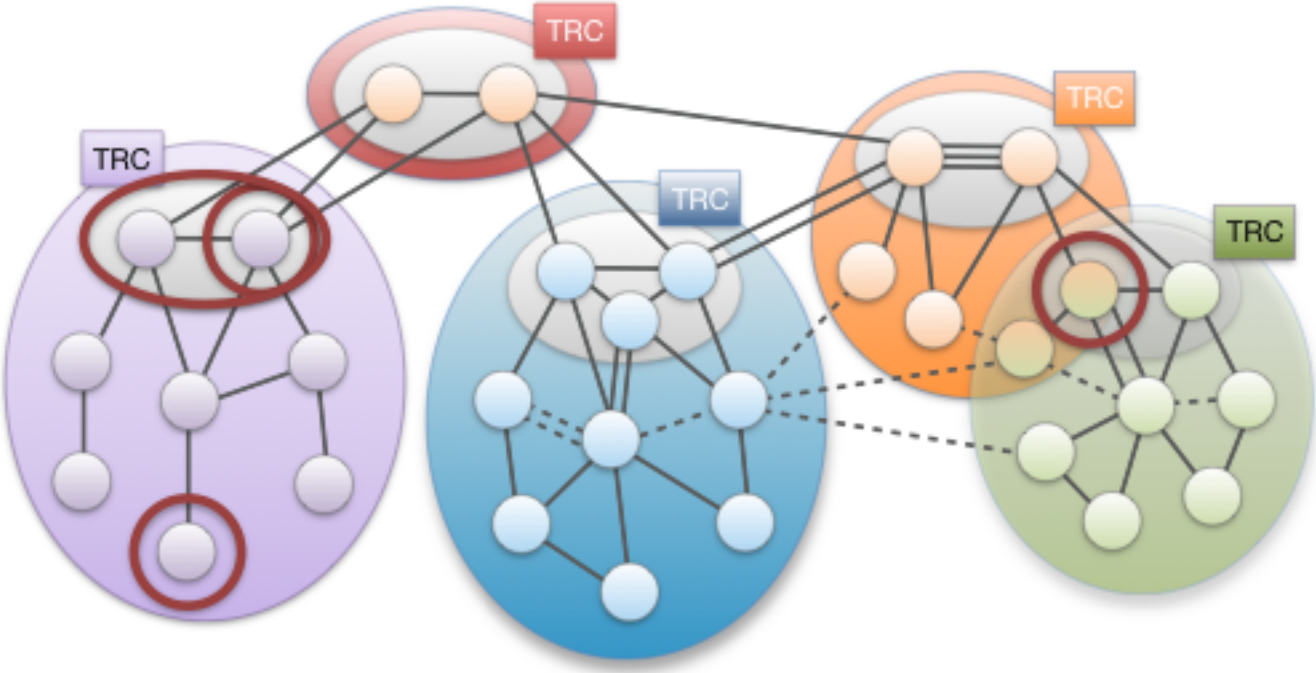

Intra-ISD Path Exploration: Beaconing

- 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

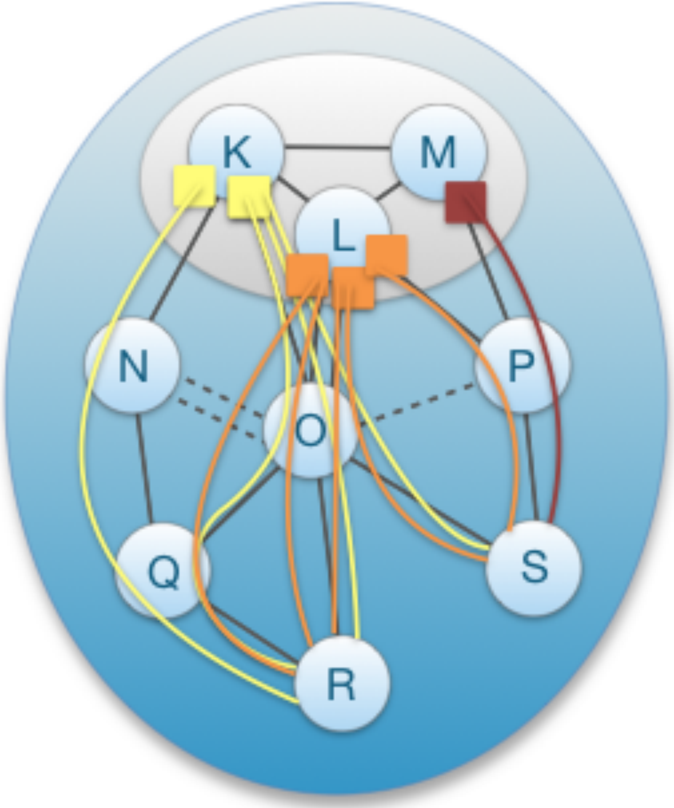

PCB Contents

- 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

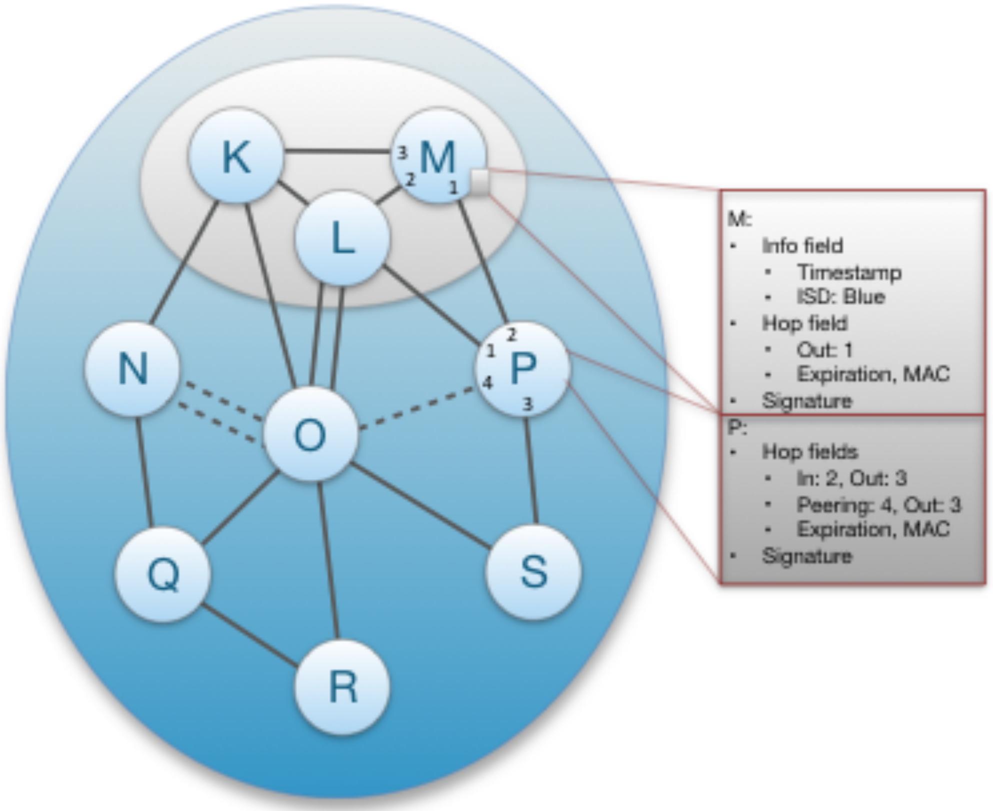

M:
- Info field
  - Timestamp
  - ISD: Blue
- Hop field
  - Out: 1
  - Expiration, MAC
- Signature

P:
- Hop fields
  - In: 2, Out: 3
  - Peering: 4, Out: 3
  - Expiration, MAC
- Signature

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

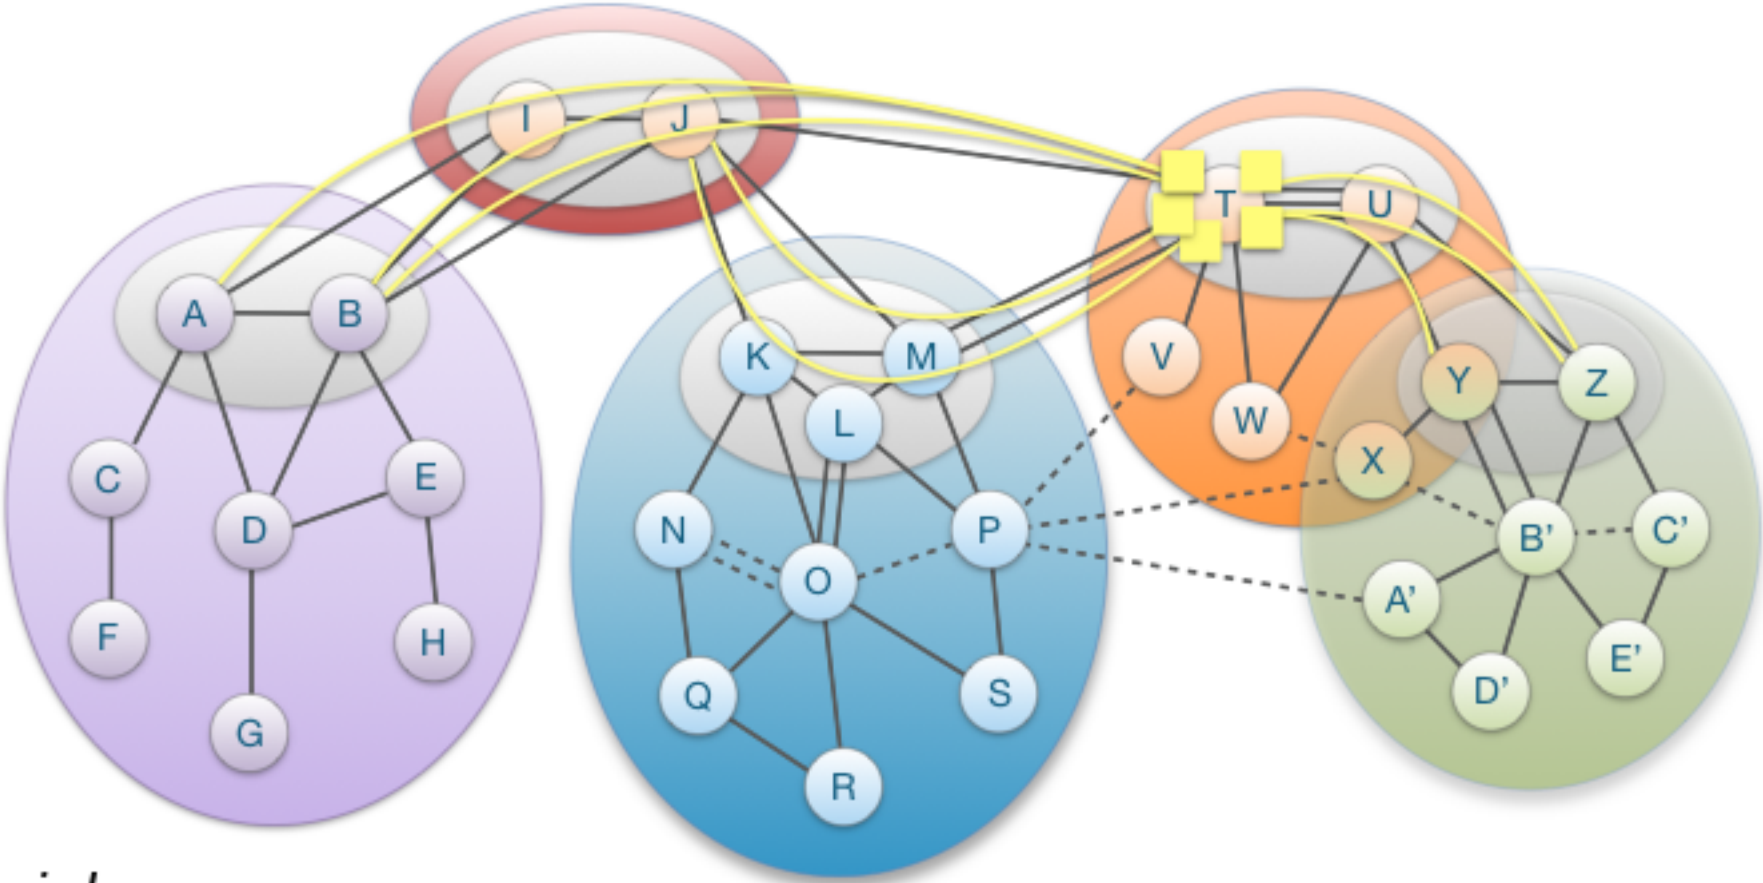

Up-Path Segment Registration

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

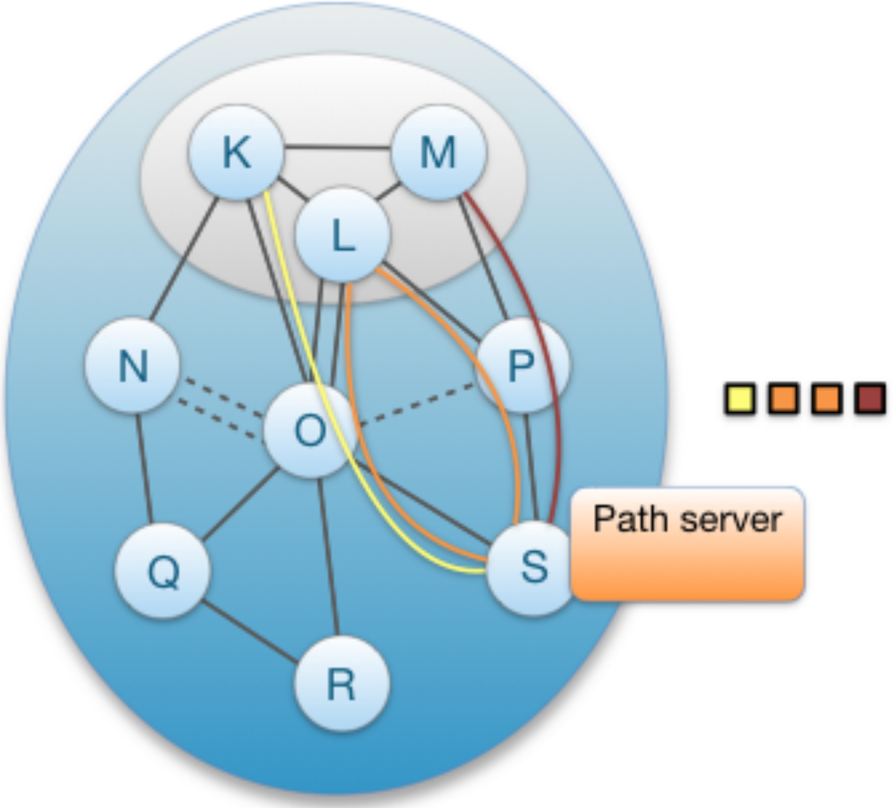

Down-Path Segment Registration

- 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

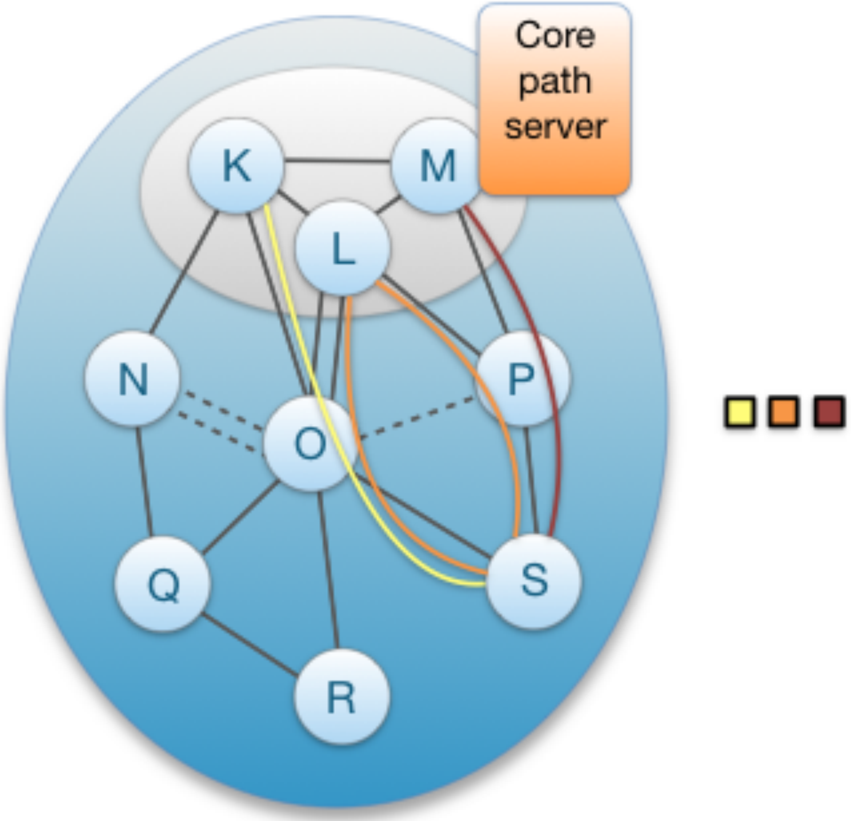

SCION Overview

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

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

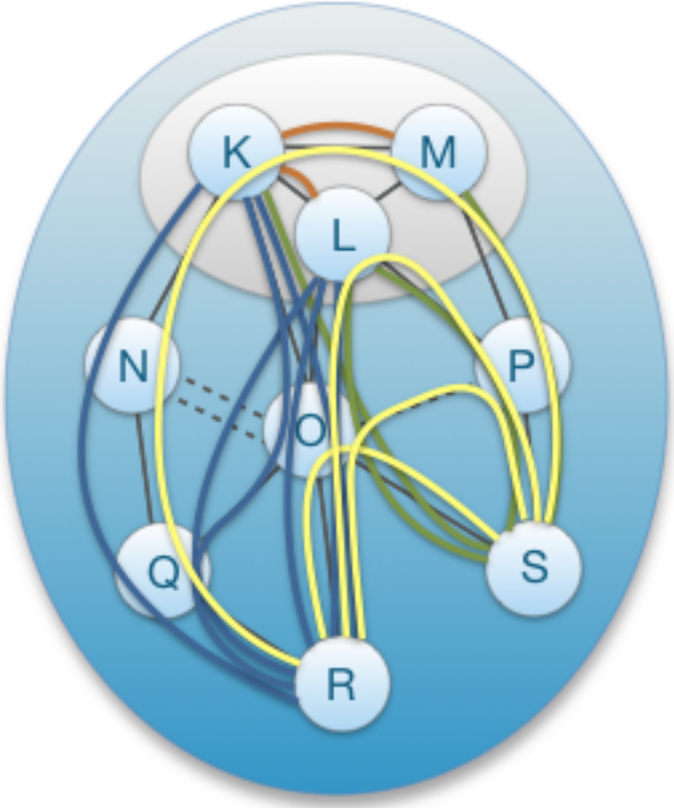

Path Lookup: Remote ISD

- 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

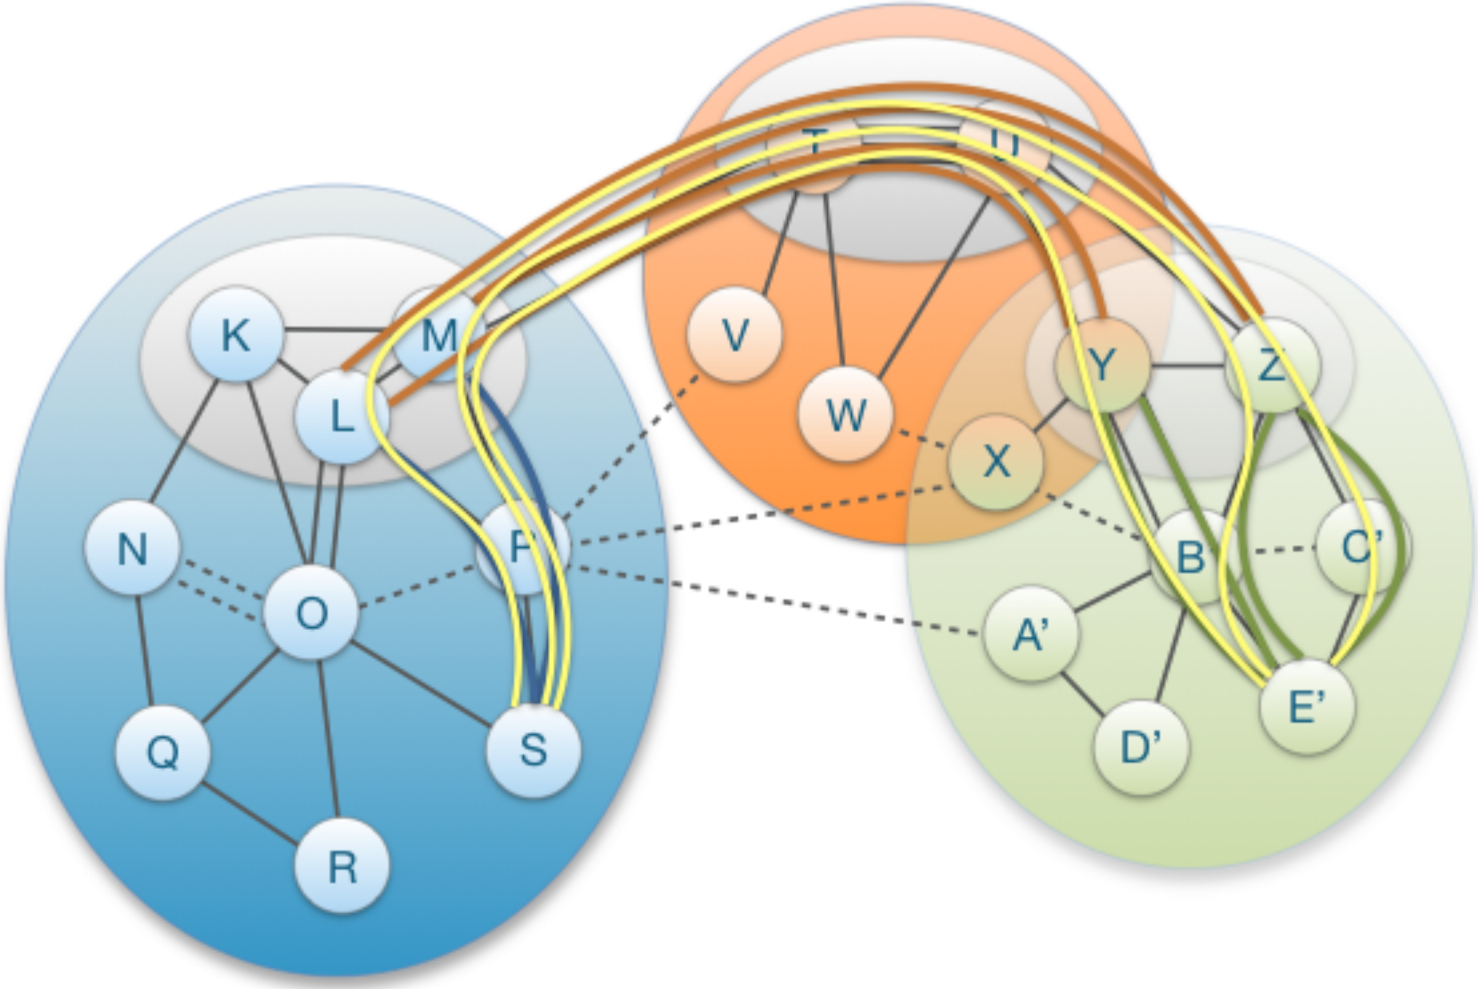

ETH zürich

Path Construction

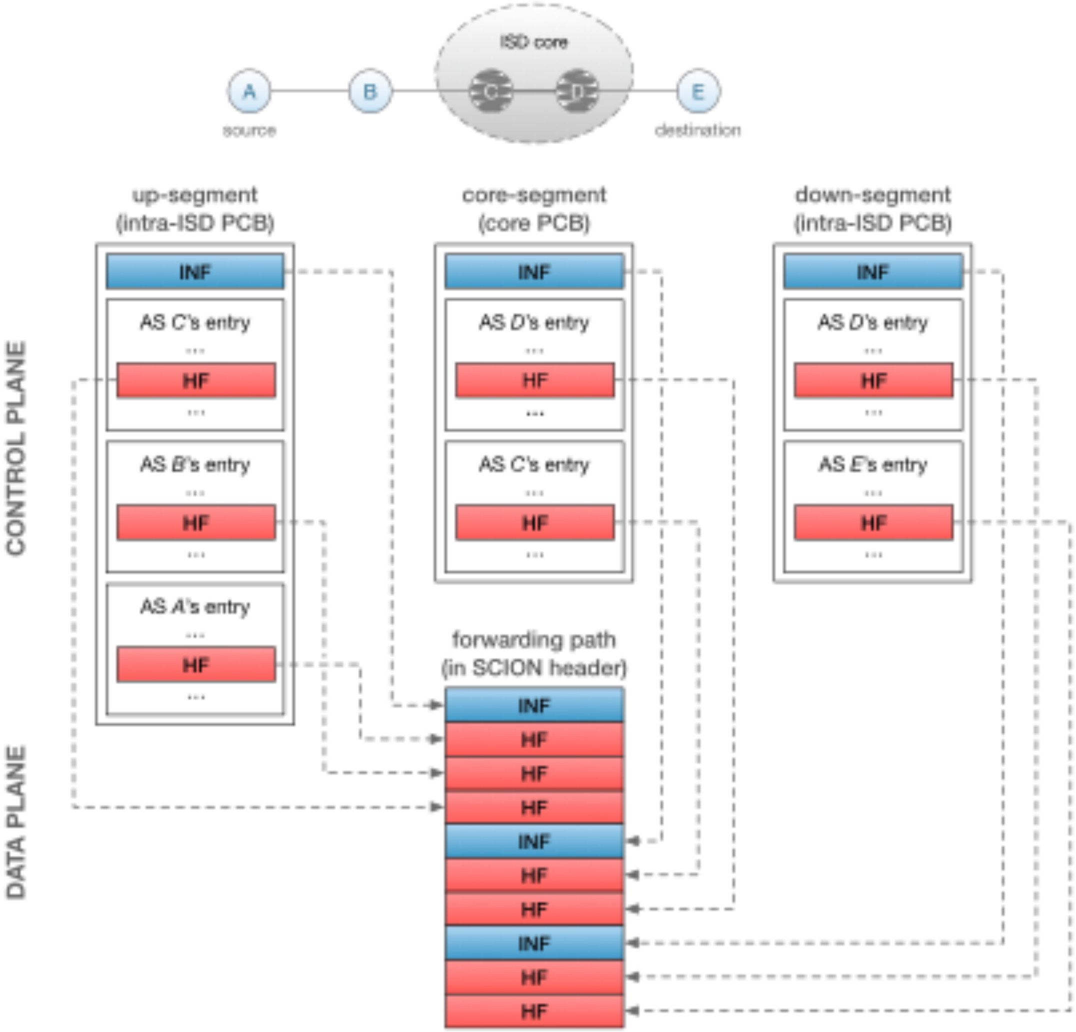

ETH zürich

SCION Drawbacks

- Longer delay until first packet can be sent, as paths need to be requested
- About 80 bytes additional overhead per packet
- Static path binding
  - No automated route failure recovery
- Additional cryptographic keys and certificates needed for communication
  - Coordination amongst core ASes required for TRC management
- New system, requires effort to learn
- Applications need to change to reap maximum benefits

SCION Overview Summary

- Complete re-design of network architecture provides numerous benefits
  - Secure operation (built-in DDoS defense, etc.)
  - Solves BGP protocol convergence issues
  - Enables dynamic path optimization
  - Multipath communication
  - 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

Information Security and Cryptography
Adrian Perrig
Pawel Szalachowski
Raphael M. Reischuk
Laurent Chuat
SCION: A Secure Internet Architecture
Springer

Outline

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

Important SCION Components and Concepts

- Hidden paths
  - Link can be converted between public and hidden
- SCION-IP Gateway (SIG)
  - Requires no upgrades to end hosts
- Carrier-grade SCION-IP Gateway (CG-SIG)
  - Enables SCION properties for customers without any extra HW or SW deployed

Exciting SCION Feature: Hidden Path

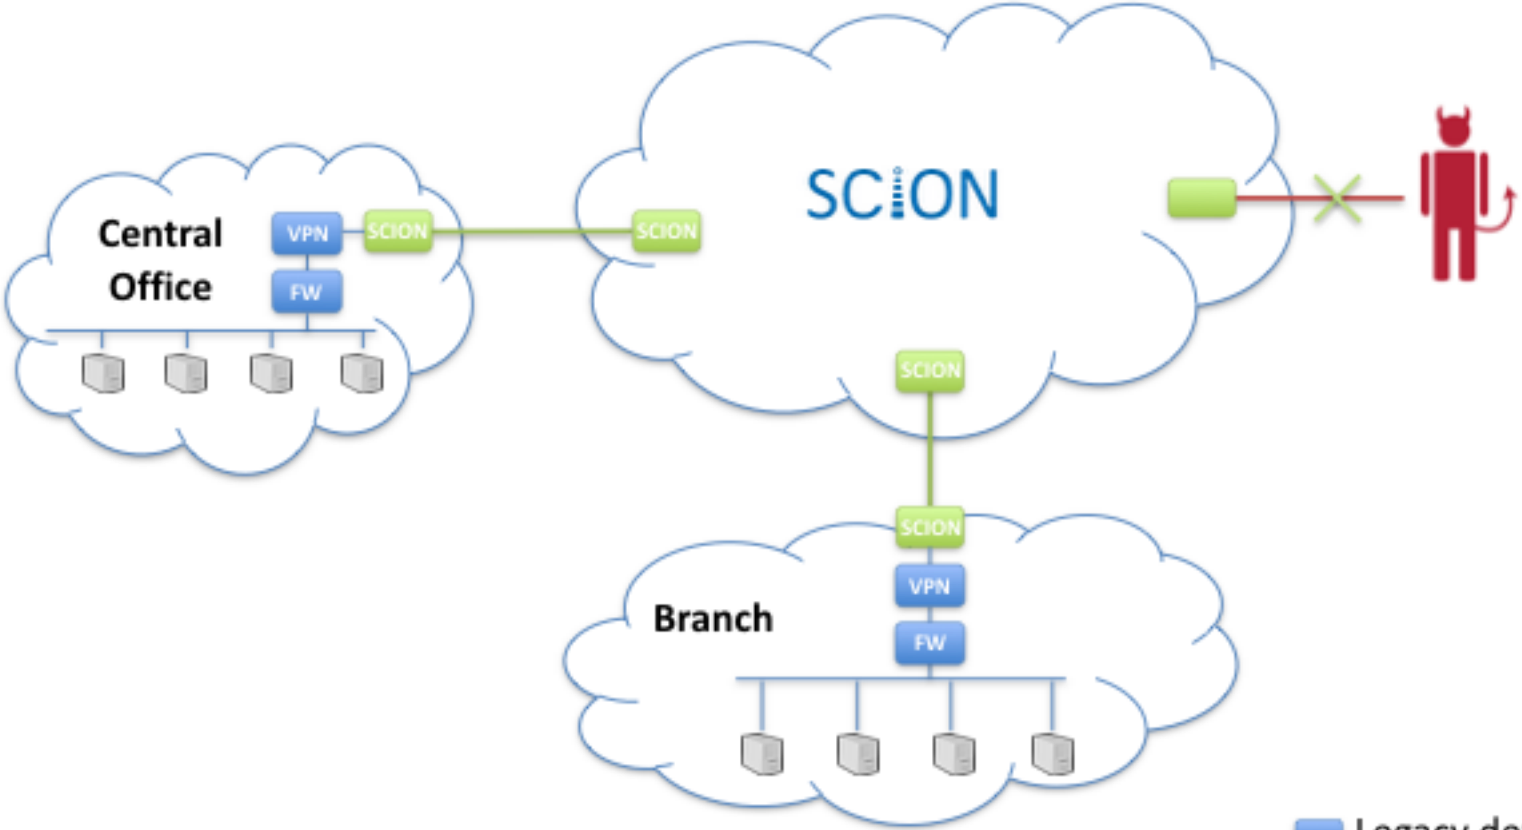

🟦 Legacy device
🟩 SCION device

SCION-IP Gateway (SIG) Deployment

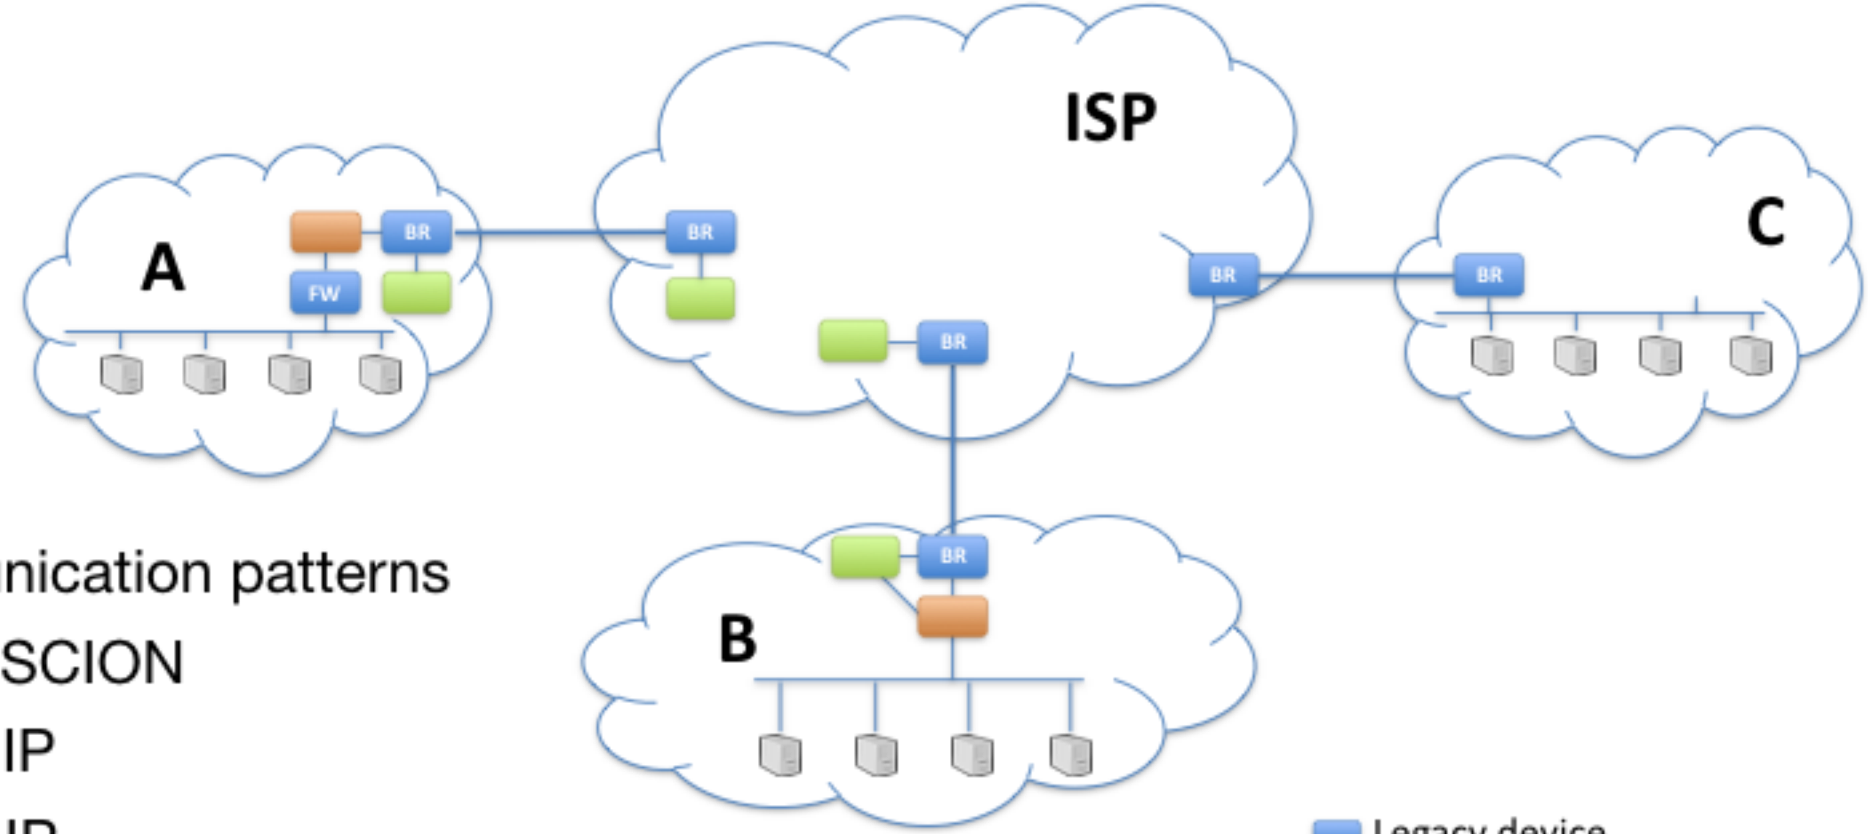

- Communication patterns
  - A - B: SCION
  - A - C: IP
  - B - C: IP

🟦 Legacy device
🟩 SCION border router
🟧 SIG

Use Cases

1. High-availability enterprise connectivity
2. High-availability small-device connectivity
3. Gaming users
4. Internet backup through SCION

Use Case 1: High-Availability Enterprise Connectivity

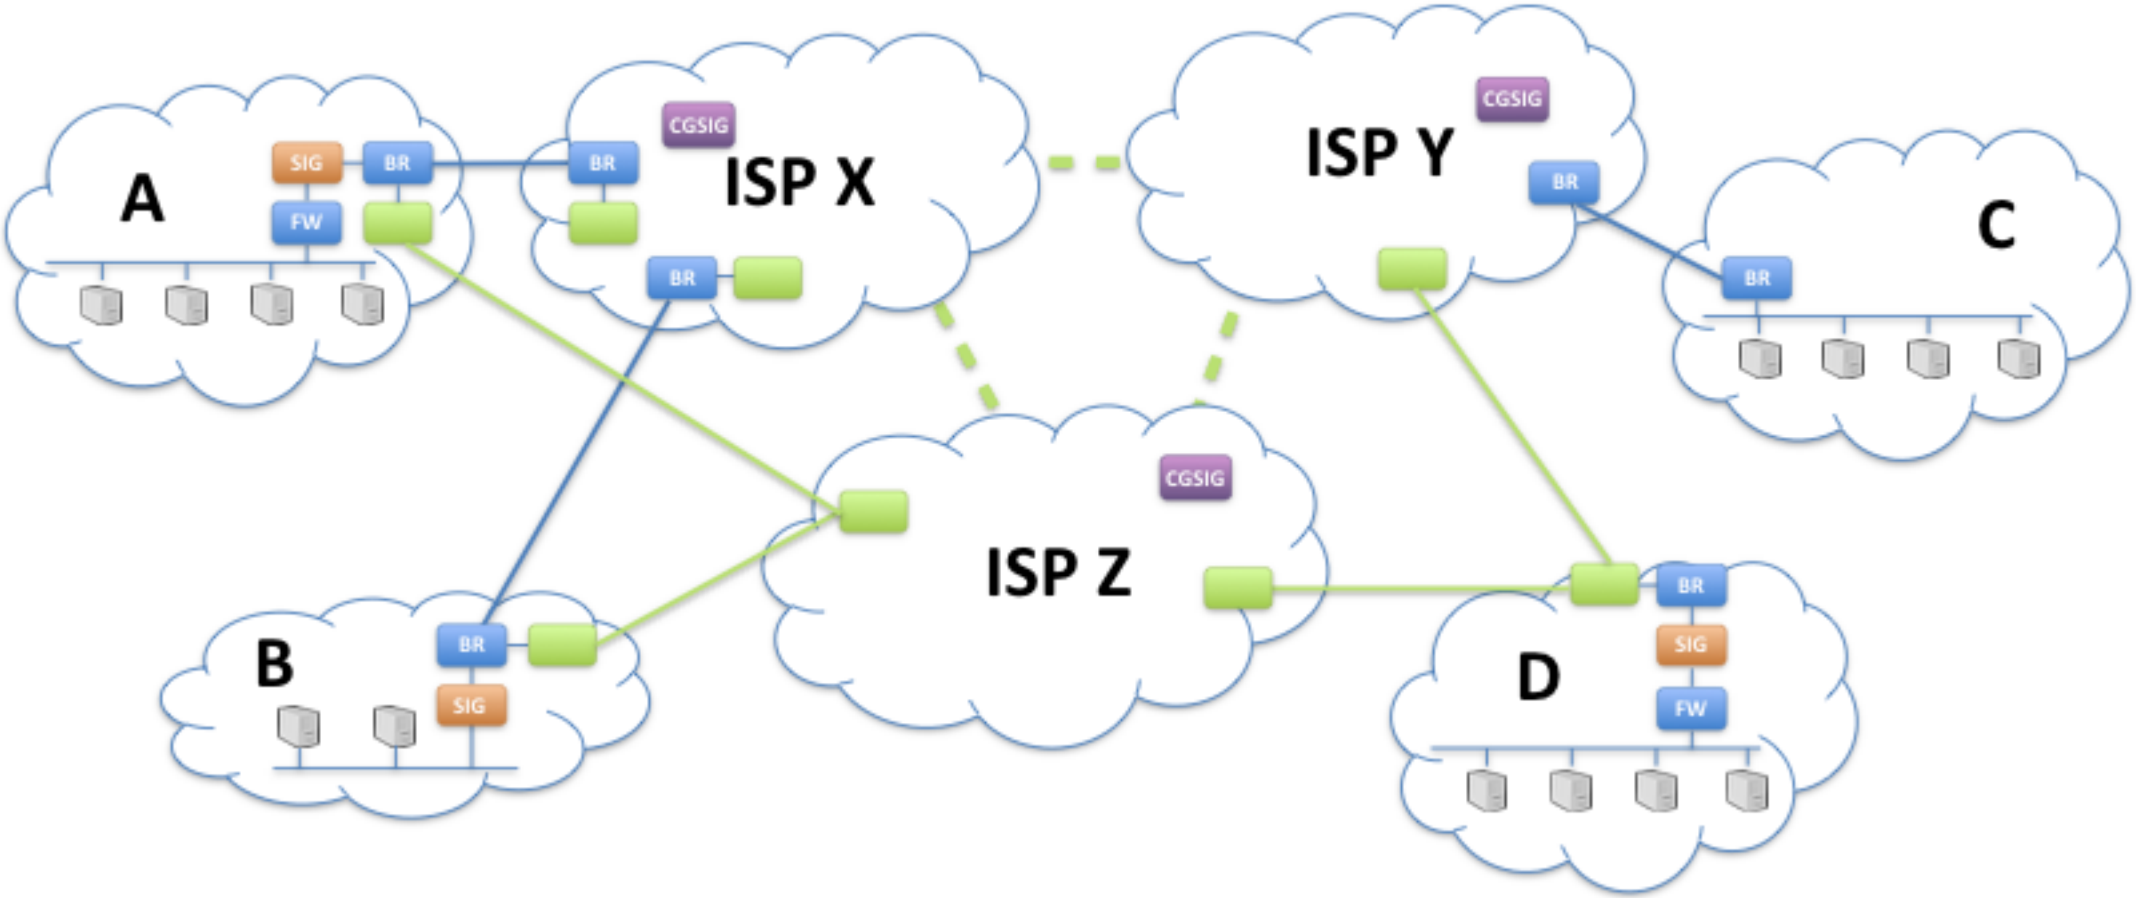

Use Case 1: Details

- ASes A, B, and D each have 2 SCION Internet connections
- ASes A and B use one regular Internet connection for both IP and SCION traffic, and a dedicated SCION connection
- AS D has two SCION-only connections, CG-SIG is used for IP traffic
- High-availability and DDoS defense is achieved through configuring SCION connections as hidden paths, enabling fine-grained traffic control
- Considering AS A, even regular IP connectivity can be retained if IP link to ISP X fails: CG-SIG in ISP X starts receiving A's traffic after link fails and sends it through AS Y to A over SCION connections
  - Low management overhead for ISPs, no need to announce A's prefix over BGP
  - Traffic management for IP traffic can be steered by A
  - Fast failover time, at the time scale of X's intra-domain protocol instead of BGP

Use Case 2: High-Availability Small-Device Connectivity

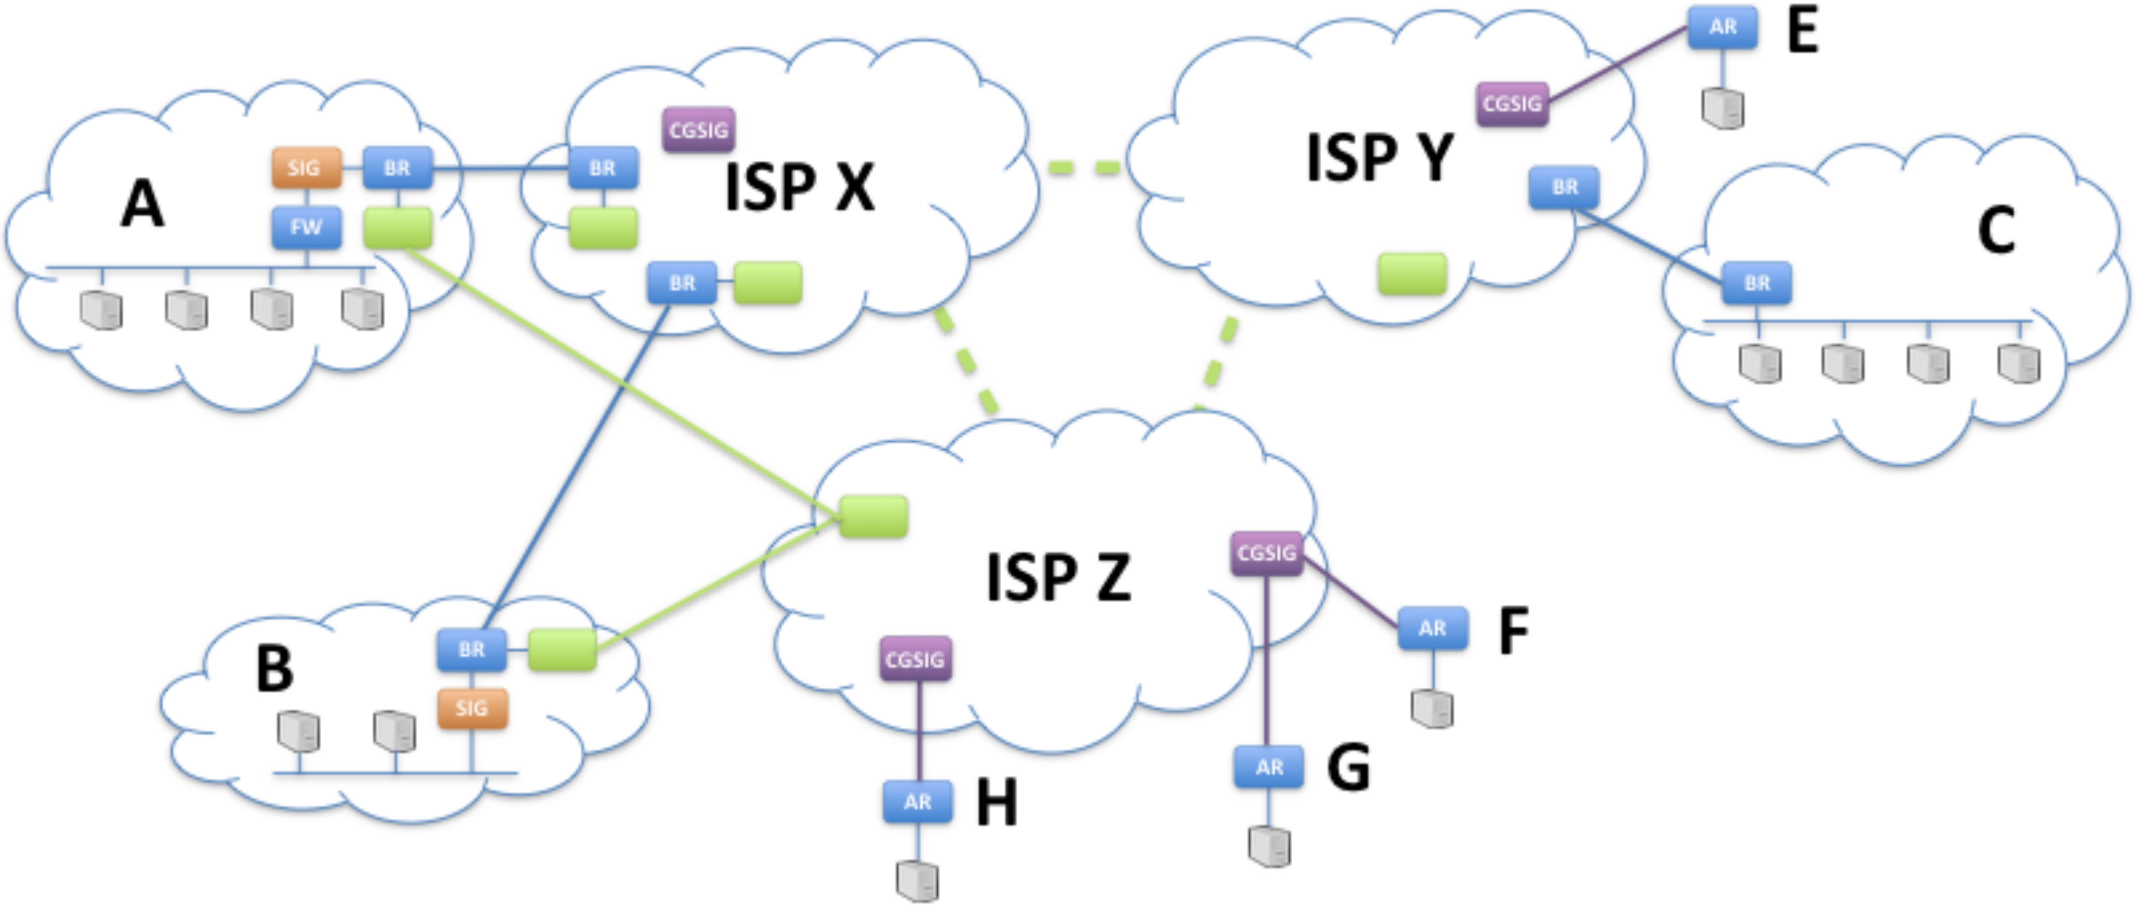

ETH zürich

SCION

Use Case 2: Details

- Since small devices cannot each be their own SCION AS, they are connected via provider's CG-SIG, with local private address space
- Desired property: devices E, F, G, and H should obtain highly available communication with AS A
- Approach: AS A registers a hidden path to ISP Z with CG-SIGs connecting its devices, for use only by its own devices
- Publicly available paths through ISP X can also be used, in case of failure of SCION link A-Z

Use Case 3: Gaming Users

- Gaming users purchase SCION Internet connection, which is realized through Carrier-Grade SIG (CG-SIG)
- Advantages
  - Latency optimization by CG-SIG
  - DoS / DDoS protection

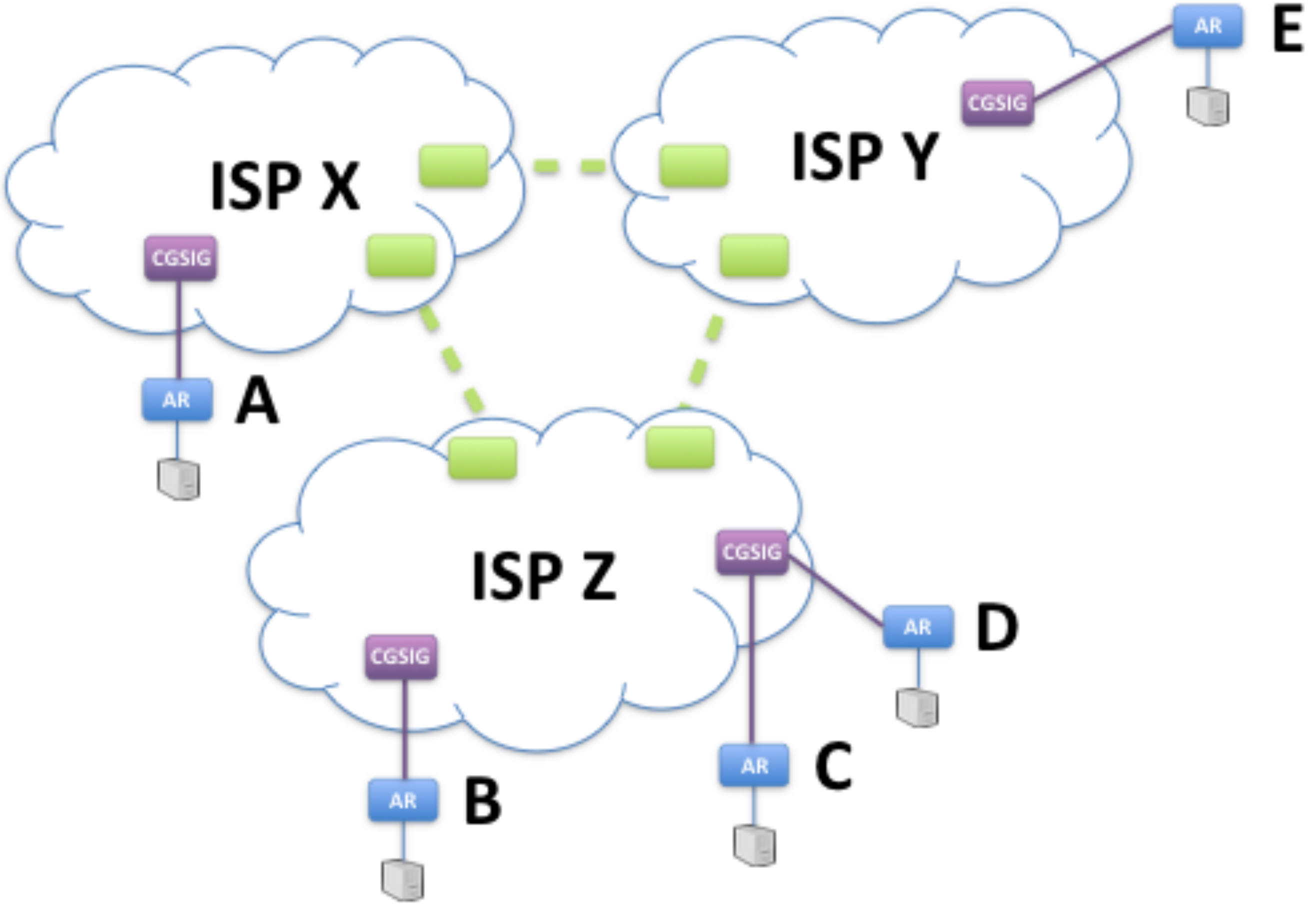

Use Case 4 Motivation:
Routing attacks can be used to partition the Bitcoin network

- See https://btc-hijack.ethz.ch

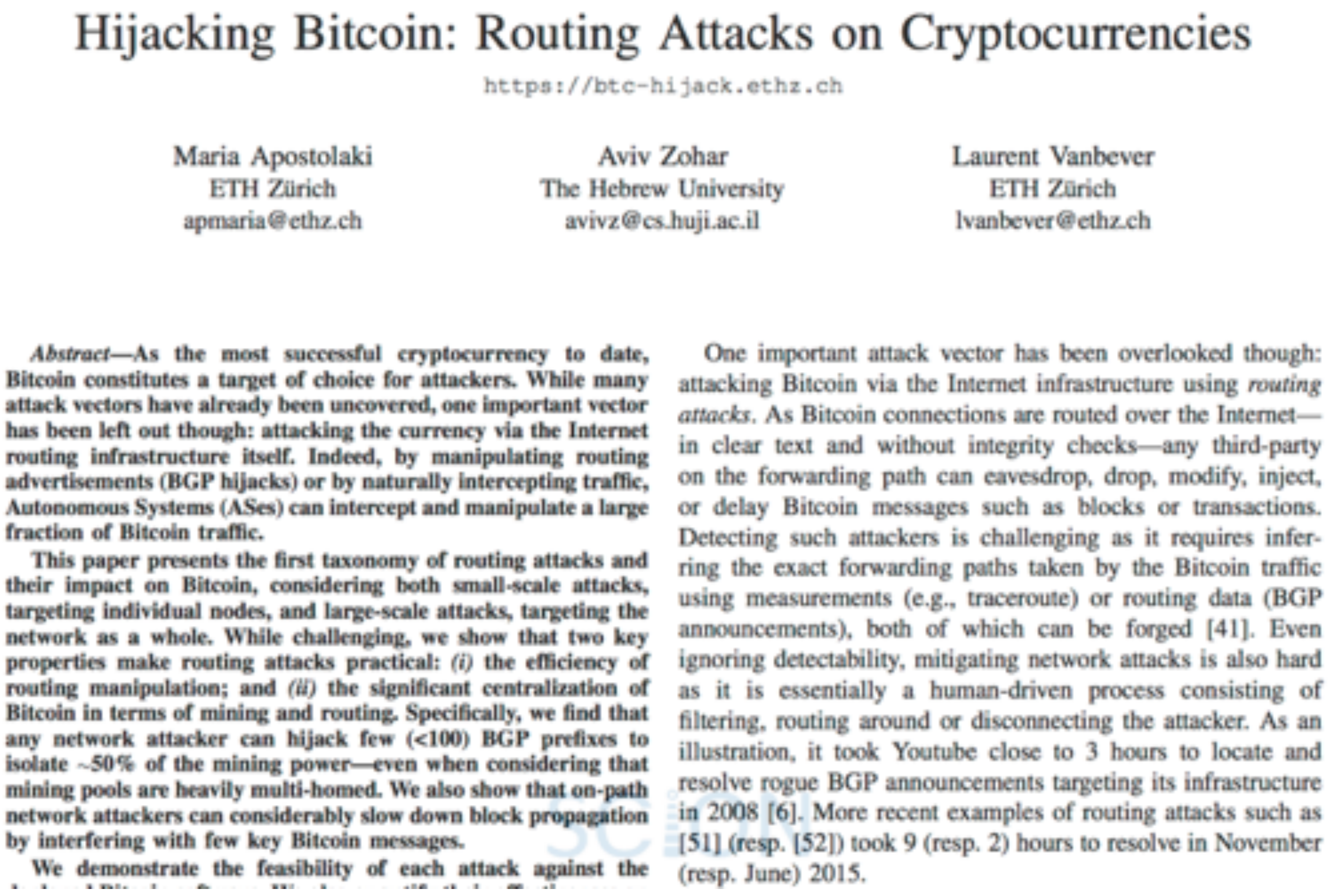

Hijacking Bitcoin: Routing Attacks on Cryptocurrencies

https://btc-hijack.ethz.ch

Maria Apostolaki
ETH Zürich
apmaria@ethz.ch

Aviv Zohar
The Hebrew University
avivz@cs.huji.ac.il

Laurent Vanbever
ETH Zürich
lvanbever@ethz.ch

***Abstract*—As the most successful cryptocurrency to date, Bitcoin constitutes a target of choice for attackers. While many attack vectors have already been uncovered, one important vector has been left out though: attacking the currency via the Internet routing infrastructure itself. Indeed, by manipulating routing advertisements (BGP hijacks) or by naturally intercepting traffic, Autonomous Systems (ASes) can intercept and manipulate a large fraction of Bitcoin traffic.**

This paper presents the first taxonomy of routing attacks and their impact on Bitcoin, considering both small-scale attacks, targeting individual nodes, and large-scale attacks, targeting the network as a whole. While challenging, we show that two key properties make routing attacks practical: *(i)* the efficiency of routing manipulation; and *(ii)* the significant centralization of Bitcoin in terms of mining and routing. Specifically, we find that any network attacker can hijack few (<100) BGP prefixes to isolate ~50% of the mining power—even when considering that mining pools are heavily multi-homed. We also show that on-path network attackers can considerably slow down block propagation by interfering with few key Bitcoin messages.

We demonstrate the feasibility of each attack against the

One important attack vector has been overlooked though: attacking Bitcoin via the Internet infrastructure using *routing attacks*. As Bitcoin connections are routed over the Internet—in clear text and without integrity checks—any third-party on the forwarding path can eavesdrop, drop, modify, inject, or delay Bitcoin messages such as blocks or transactions. Detecting such attackers is challenging as it requires inferring the exact forwarding paths taken by the Bitcoin traffic using measurements (e.g., traceroute) or routing data (BGP announcements), both of which can be forged [41]. Even ignoring detectability, mitigating network attacks is also hard as it is essentially a human-driven process consisting of filtering, routing around or disconnecting the attacker. As an illustration, it took Youtube close to 3 hours to locate and resolve rogue BGP announcements targeting its infrastructure in 2008 [6]. More recent examples of routing attacks such as [51] (resp. [52]) took 9 (resp. 2) hours to resolve in November (resp. June) 2015.

Use Case 4: Internet Backup through SCION

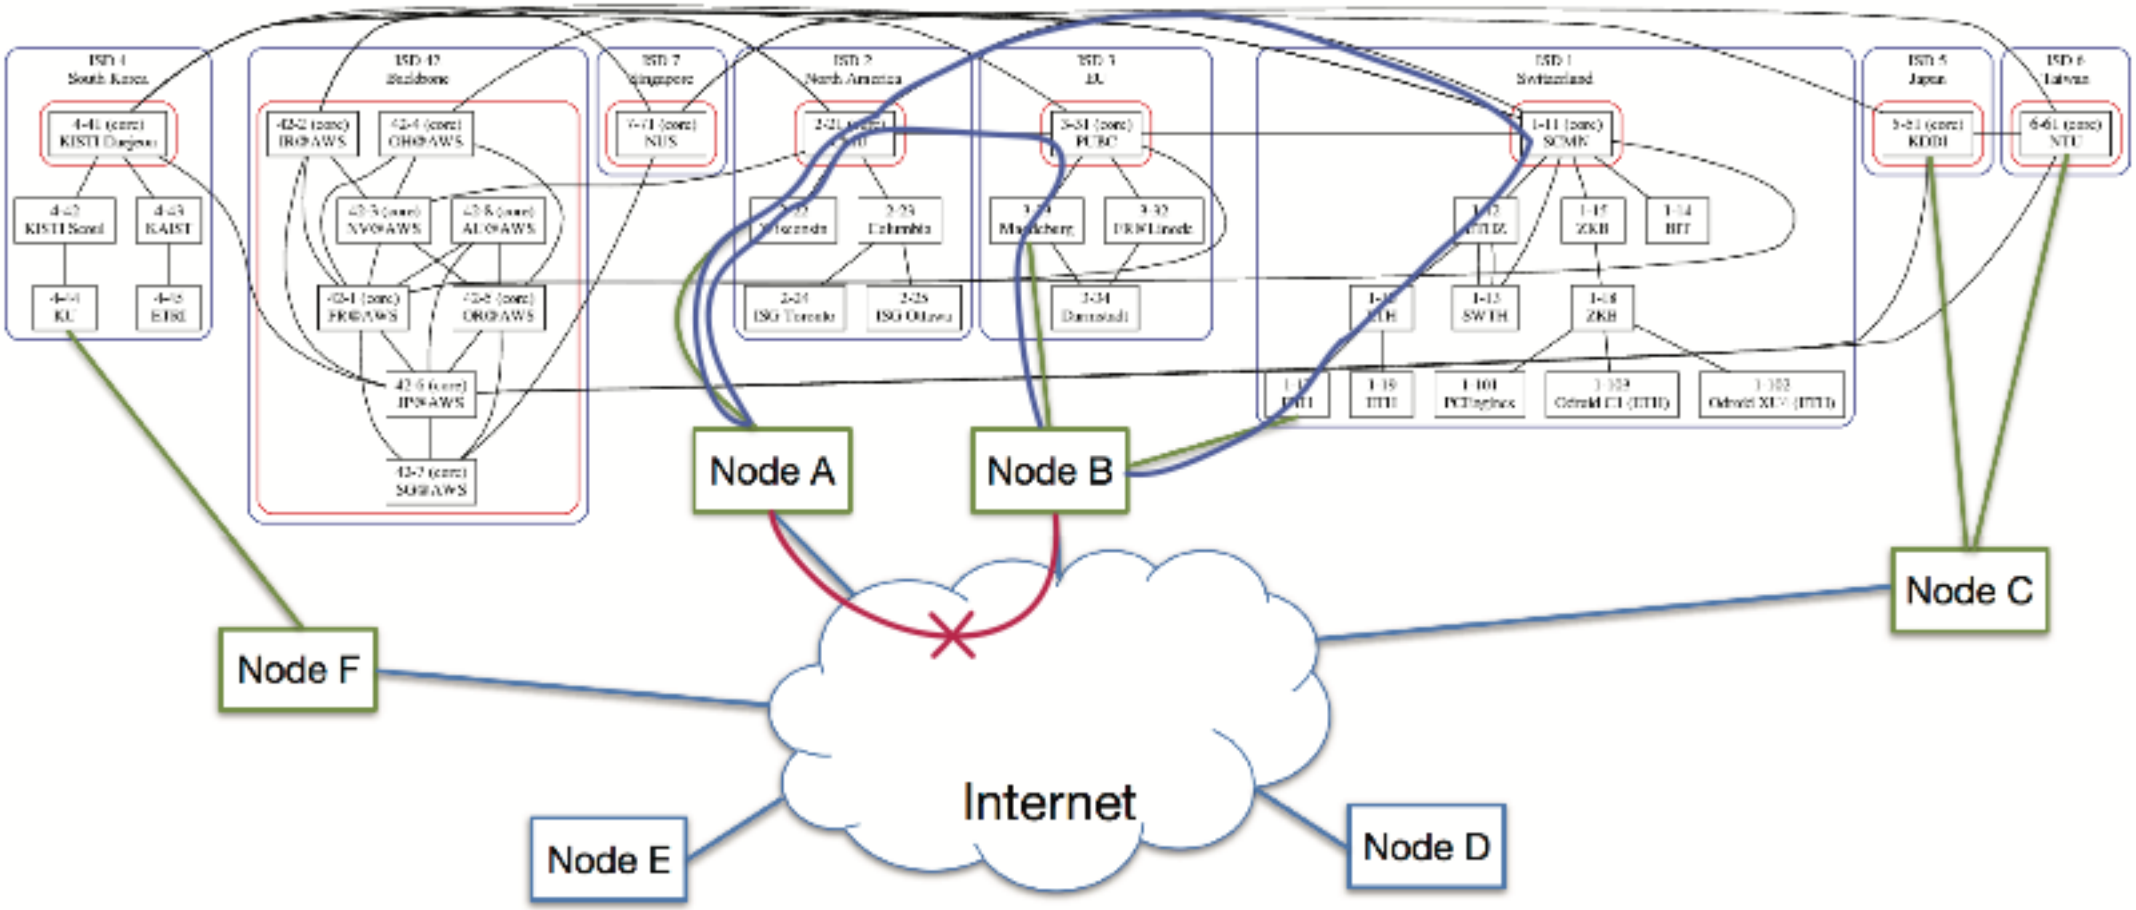

Value Proposition for Customers

- SCION offers highly secure and available Internet communication with built-in DDoS defense

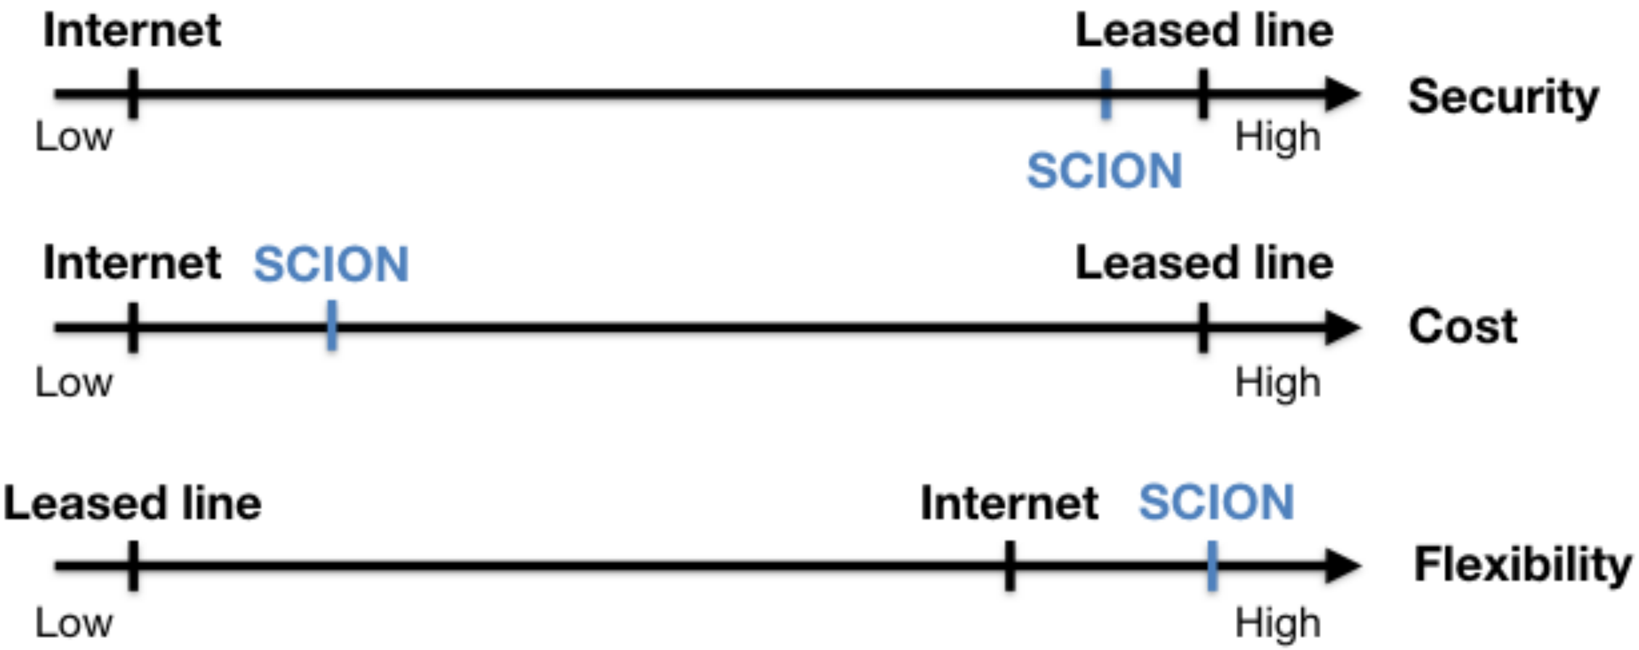

Value Proposition for ISPs

- New service offerings for customers
  - Premium link offerings
  - Geofencing, path choice
  - Business continuity (high availability / fast failover)
- Many customers want higher security / availability, but leased line cost too high for many use cases
- Customers want higher flexibility for connections
- Lower network management overhead
- Higher network utilization

SCIONLab

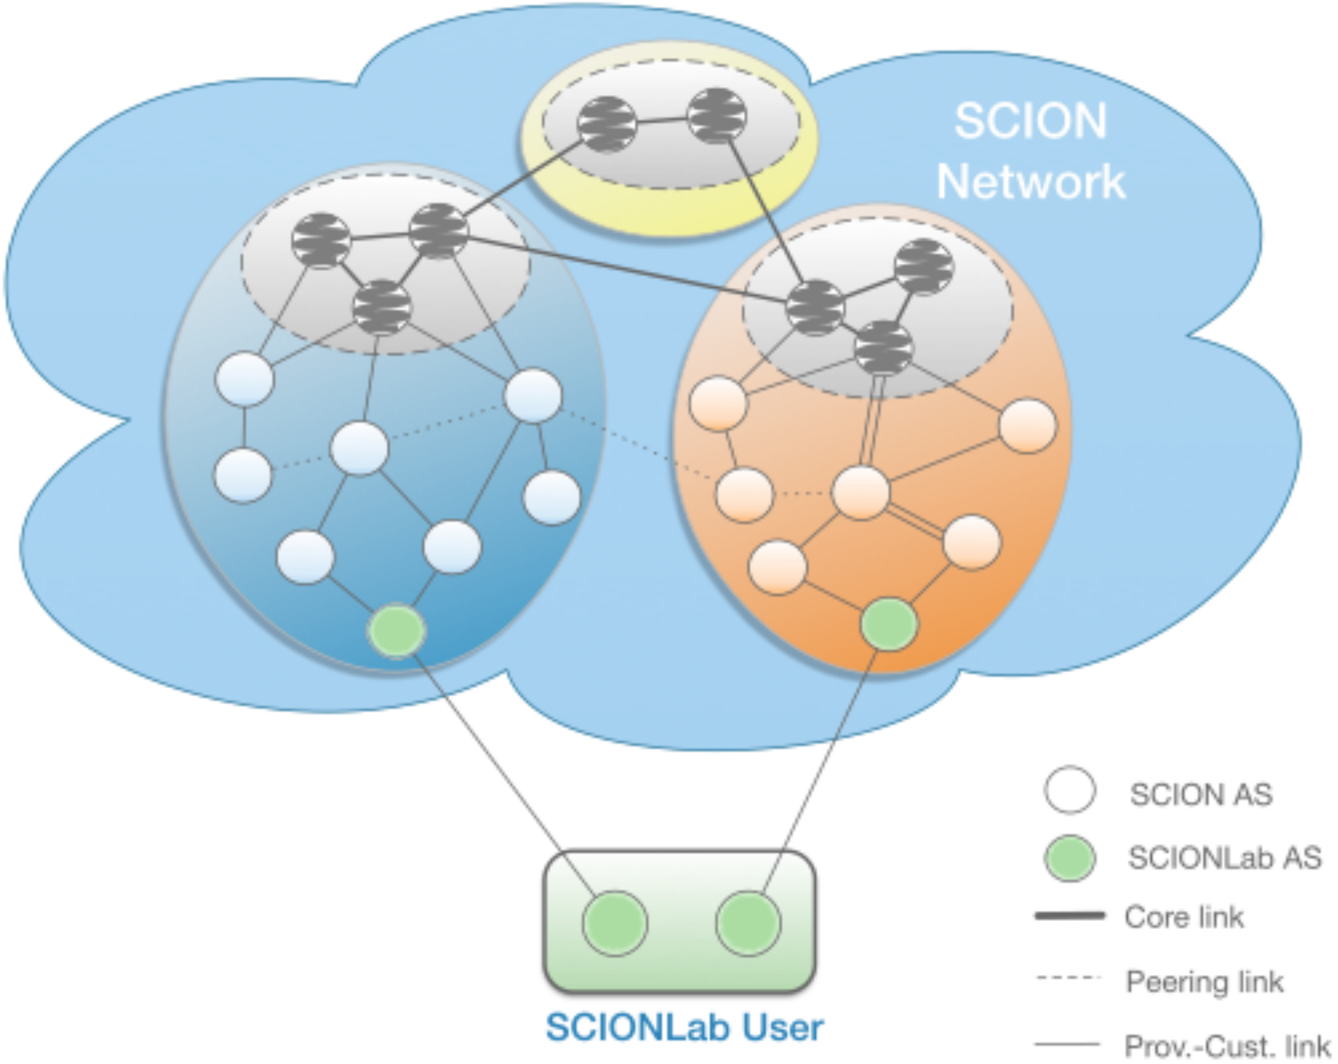
SCION Network
SCIONLab User
SCION AS
SCIONLab AS
Core link
Peering link
Prov.-Cust. link

Exciting SCIONLab Research Opportunities

- Next-generation Internet architecture research
- Users obtain real ASes with all cryptographic credentials to participate in the control plane
- ASes can use their own computing resources and attach at several points in the SCIONLab network
- Path-aware networking testbed
- Hidden paths for secure IoT operation
- Control-plane PKI in place, each AS has certificate
- Network availability and performance measurement (bandwidth and latency)
- Supported features (PKI, DDoS defense mechanisms, path selection support, end host / application support)
- (Security) Usability research
- Inter-domain routing scalability research
- Multi-path research
- Multi-path QUIC socket
- End-to-end PKI system that application developers can rely on to build highly secure TLS applications
- SIBRA inter-domain resource allocation system
- DDoS defense research using in-network defense mechanisms
- Next-generation routing architecture policy definitions

Global SCIONLab Network

- https://www.scionlab.org
- Collaboration with David Hausheer @ Uni Magdeburg

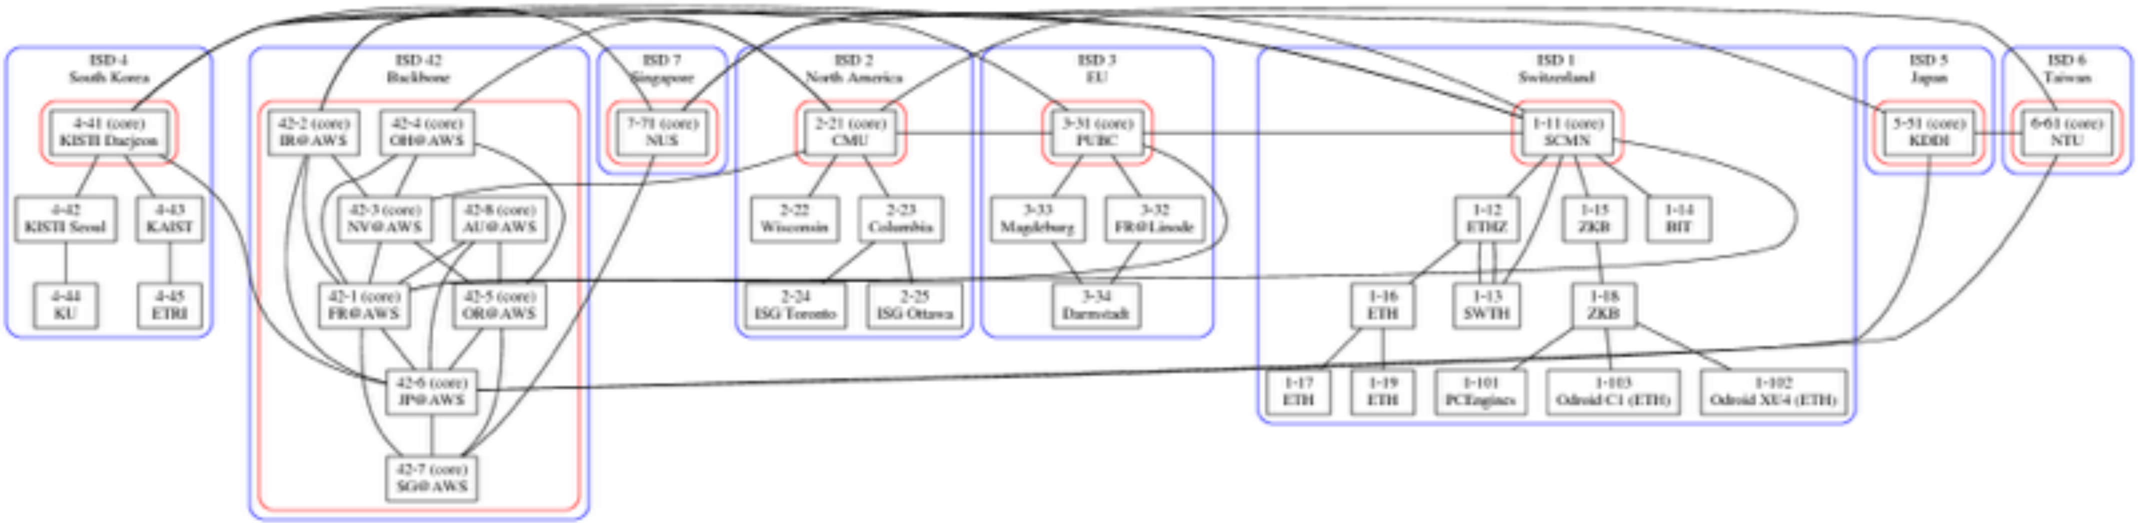

SCIONLab Path Visualization

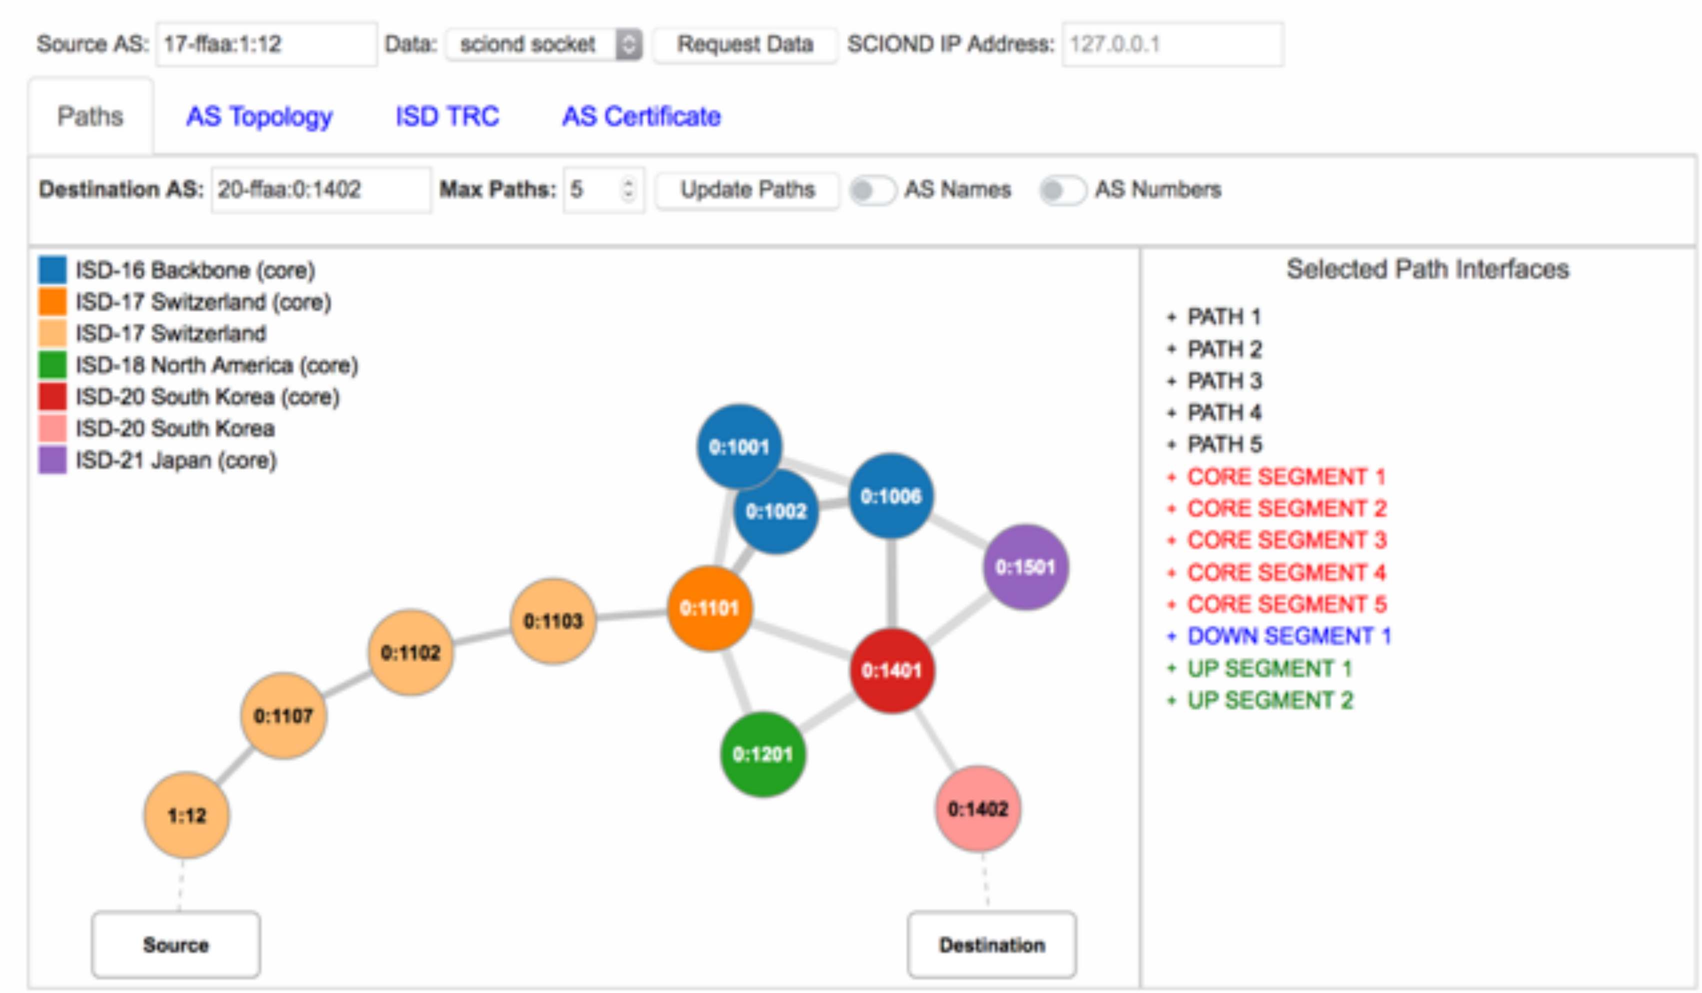

Source AS: 17-ffaa:1:12

Data: sciond socket

SCIOND IP Address: 127.0.0.1

Destination AS: 20-ffaa:0:1402

Max Paths: 5

Selected Path Interfaces

- PATH 1
- PATH 2
- PATH 3
- PATH 4
- PATH 5
- CORE SEGMENT 1
- CORE SEGMENT 2
- CORE SEGMENT 3
- CORE SEGMENT 4
- CORE SEGMENT 5
- DOWN SEGMENT 1
- UP SEGMENT 1
- UP SEGMENT 2

Commercial SCION Network

- Deutsche Telekom, Swisscom, SWITCH, Init7 offer SCION connections (as test) on a commercial SCION network
- Several banks and Swiss government are running trial deployments
  - One large bank has been running production traffic over SCION since August 2017

How to obtain a SCION Connection?

- Individual: SCIONLab https://www.scionlab.org
  - SCION AS running on VM within 10 minutes
- University, research lab
  - SWITCH, DFN can (soon) provide SCION connections
  - David Hausheer @ Uni Magdeburg has set up SCION VMs at GEANT <hausheer@ovgu.de>
  - KOREN in Korea
- Corporation, Government entity
  - Swisscom
  - Deutsche Telekom

Conclusions

- It is possible to evolve Layer 3: SCION is a secure Internet architecture that we can use today
- Strong properties for high-availability communication
  - Multipath routing architecture offers multitude of path choices for meaningful diverse path selection
  - For some cases, lower latency than in today's Internet
  - Fast failover providing business continuity
  - Prevention of routing attacks
  - Built-in DDoS defense mechanisms

SCION Commercialization

- Founded Anapaya Systems in June 2017
- 4 founders: David Basin, Sam Hitz (CEO), Peter Müller, Adrian Perrig
- Several banks and ISPs are customers
- https://www.anapaya.net

ANAPAYA SYSTEMS

Securing and Optimizing Internet Communication

anapaya systems

Online Resources

- https://www.scion-architecture.net
  - Book, papers, videos, tutorials
- https://www.scionlab.org
  - SCIONLab testbed infrastructure
- https://www.anapaya.net
  - SCION commercialization
- https://github.com/scionproto/scion
  - Source code

SCION Core Project Team

- Netsec: Daniele Asoni, Laurent Chuat, Sergiu Costea, Piet De Vaere, Sam Hitz, Mike Farb, Tobias Klausmann, Cyrill Krähenbühl, Jonghoon Kwon, Tae-Ho Lee, Sergio Monroy, Chris Pappas, Juan Pardo, Adrian Perrig, Benjamin Rothenberger, Stephen Shirley, Jean-Pierre Smith, Brian Trammell
- Infsec: David Basin, Tobias Klenze, Ralf Sasse, Christoph Sprenger, Thilo Weghorn
- Programming Methodology: Marco Eilers, Peter Müller
- Uni Magdeburg: David Hausheer

Thanks to our Collaborators!

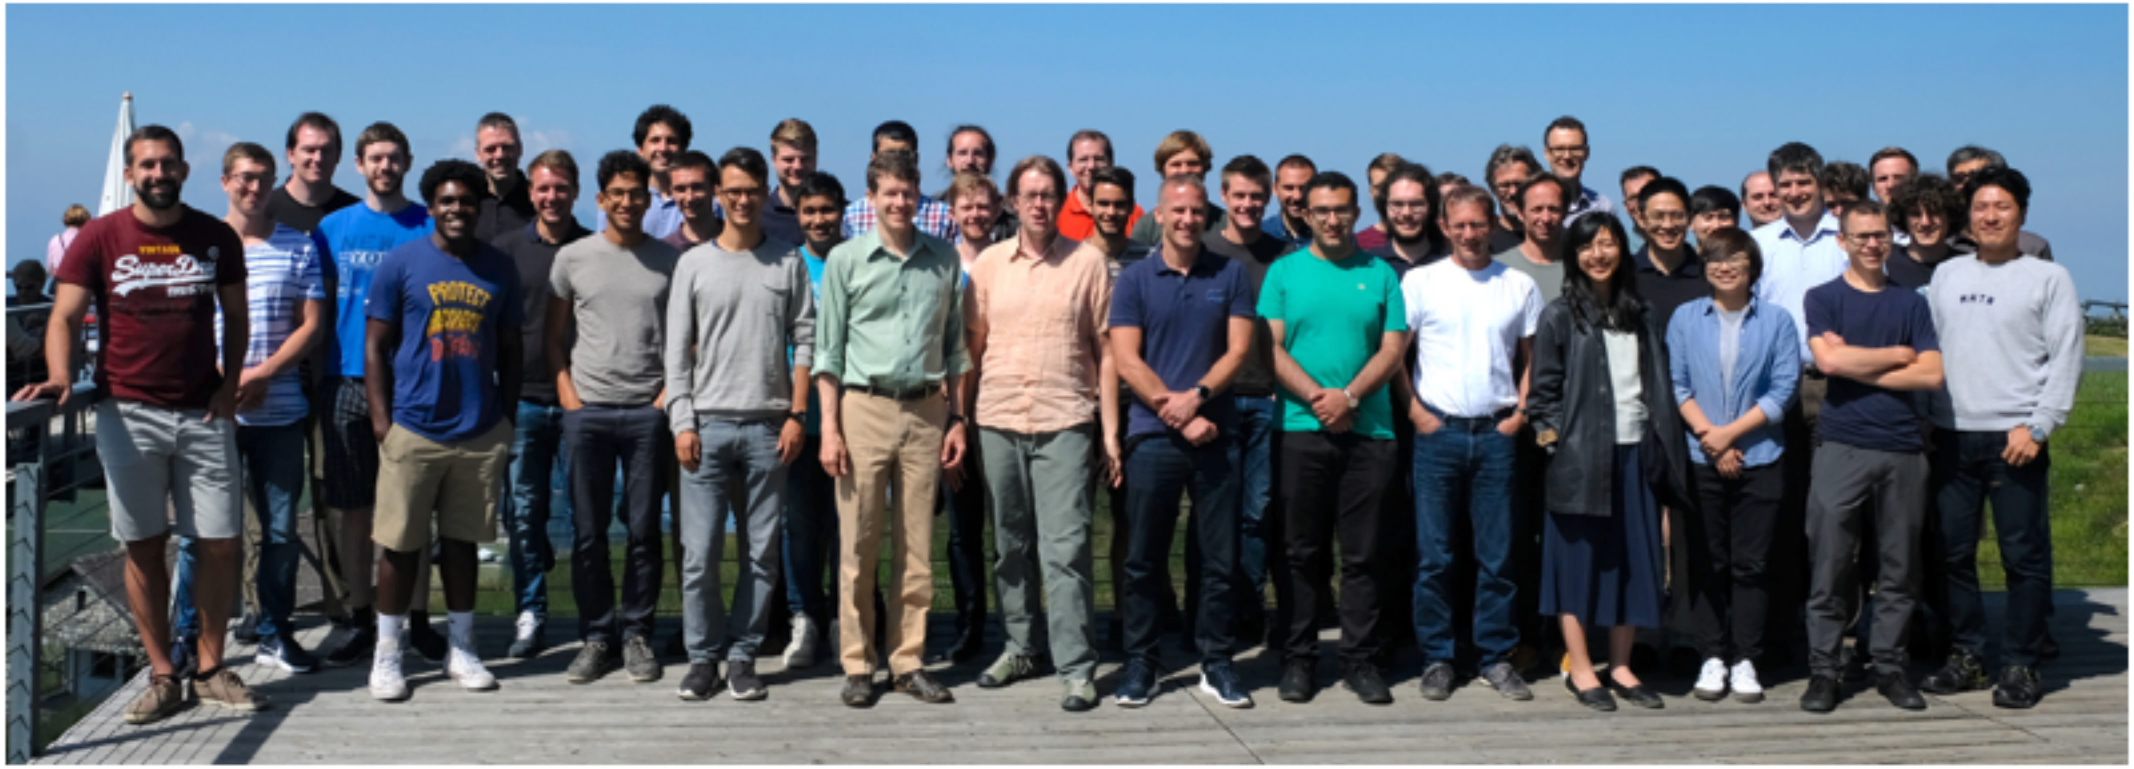

ETH zürich

SCION

Thanks to our Sponsors!

ETH zürich

ZISC Zurich Information Security & Privacy Center

erc

NSF

SIX

swisscom

Zürcher Kantonalbank

NEC

INFOSEC GLOBAL

Google

KDDI