SCION Secure Next-generation Internet Architecture

Adrian Perrig

Network Security Group, ETH Zürich

scion-architecture.net

ETH zürich

The **Internet** is perceived to be like the pyramids: **monumental structure** that has **stood the test of time** and **cannot be changed**

ETH zürich SCION

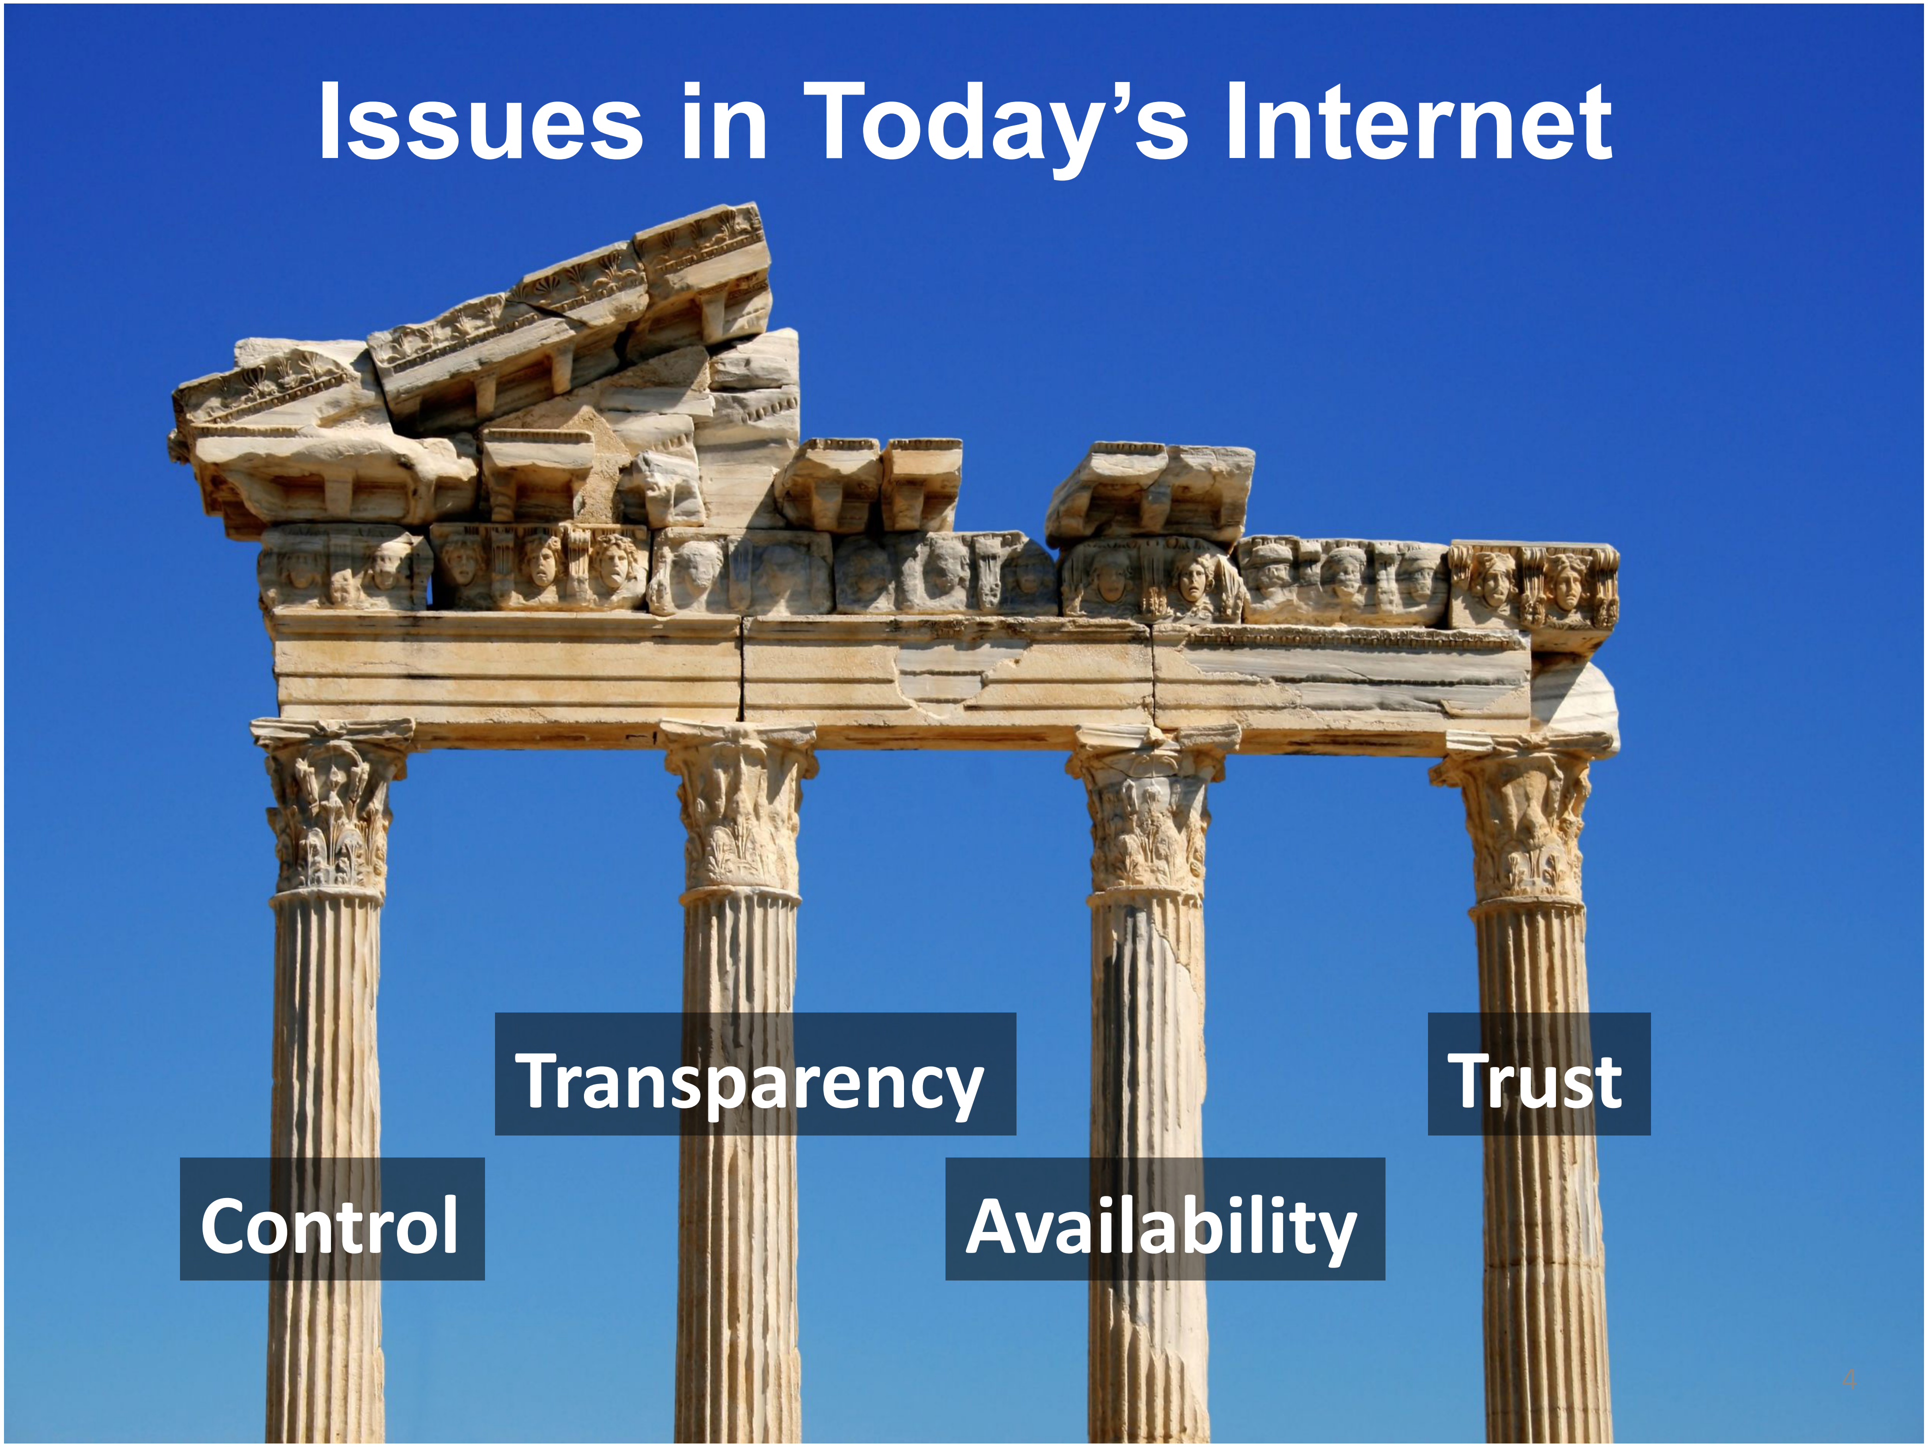
Issues in Today's Internet
Transparency
Trust
Control
Availability
4

Problem 1: Non-Scalability of Trust

Pervasive Trust in Early Internet

"There were only two other Dannys on the Internet then. I knew them both. We didn't all know each other, but we all kind of trusted each other, and **that basic feeling of trust permeated the whole network**." – Danny Hillis, about the Internet in the early 1980s, TED talk, Feb 2013.

Non-Scalability of Trust

- As the Internet has grown to encompass a large part of the global population, not everyone trusts everyone else on the Internet any more
- The heterogeneity of global environment complicates entity authentication infrastructures
  - Relevant in this context: authentication of routing updates, DNS replies, TLS certificates
- Two models for trust roots for authentication
  - Monopoly model
  - Oligarchy model

Monopoly Model for Trust Root

- Single root of trust (i.e., root public key) that is globally accepted to authenticate entities
- Examples: RPKI for BGPSEC or DNSSEC rely on a public key that forms root of trust
  - All AS certificates or DNS records are authenticated based on root of trust
- Problems
  - Entire world needs to agree on entity to hold root of trust
  - Single point of failure
  - Inefficient revocation / update mechanisms

Oligarchy Model for Trust Root

- Numerous roots of trust that are globally accepted to validate entities
- Example: TLS PKI relies on > 1000 roots of trust
  - TLS certificate accepted if signed by **any** root of trust
- Problems
  - Single point of failure: any single compromised root of trust can create any bogus TLS certificate
  - Revocation / update is handled through OS or browser software update

Proposed Approach: Isolation Domains

- Observation: subset of the Internet can agree on roots of trust → form Isolation Domain with that root of trust
- Authenticate entities within each Isolation Domain
- Users & domains can select Isolation Domain based on root of trust
- Also supports modern log-based PKI approaches: CT, AKI, ARPKI, …
- Challenge: retain global verifiability

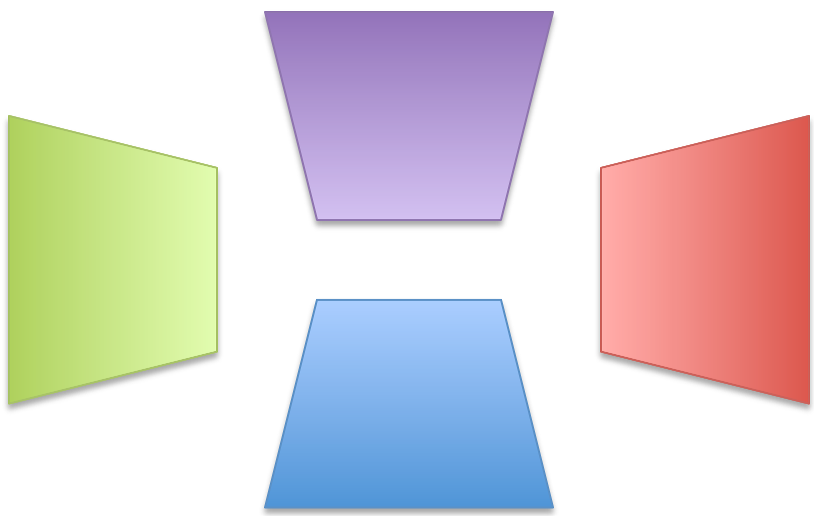

Problem 2: Control

Who controls Internet Paths?

- Current Internet offers limited control of paths
  - Border Gateway Protocol (BGP) floods announcements for destinations

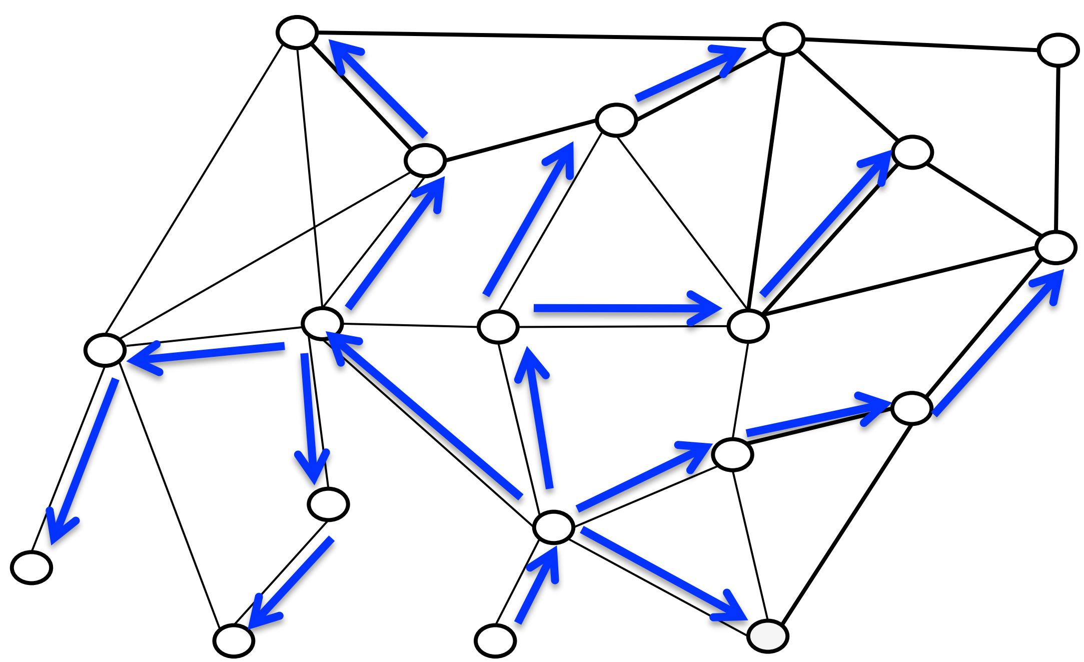

Who controls Internet Paths?

- Current Internet offers limited control of paths
  - Border Gateway Protocol (BGP) floods announcements for destinations
  - No inbound traffic control

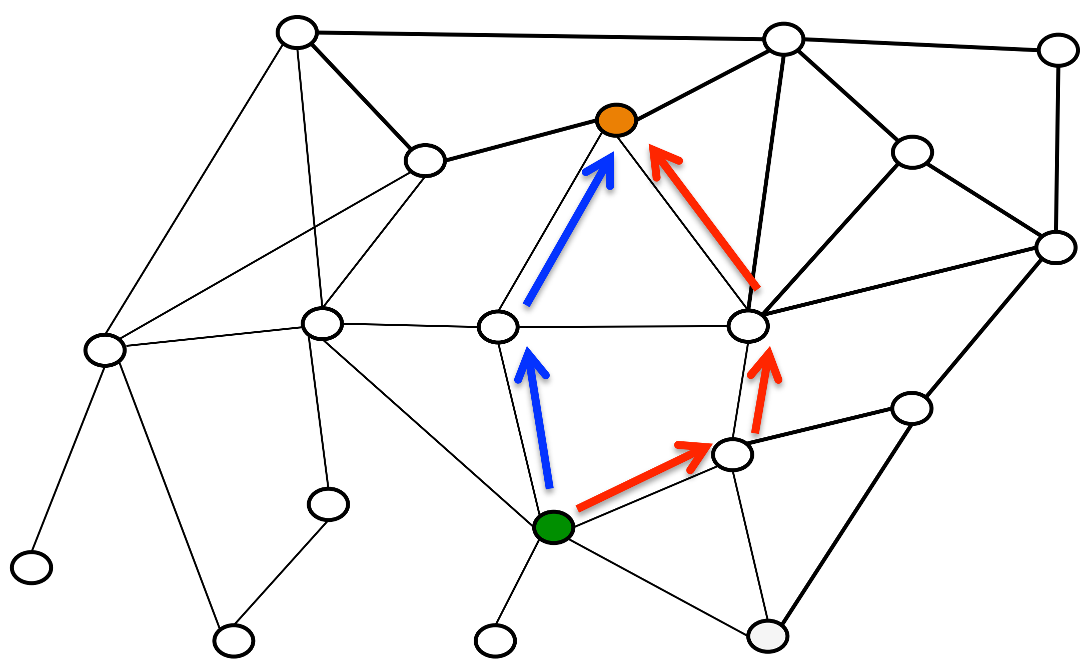

Who controls Internet Paths?

- Current Internet offers limited control of paths
- Paths can be hijacked and redirected

Who should control Paths?

- Clearly, ISPs need some amount of path control to enact their policies
- How much path control should end points (sender and receiver) have?
  - Control is a tricky issue … how to empower end points without providing too much control?

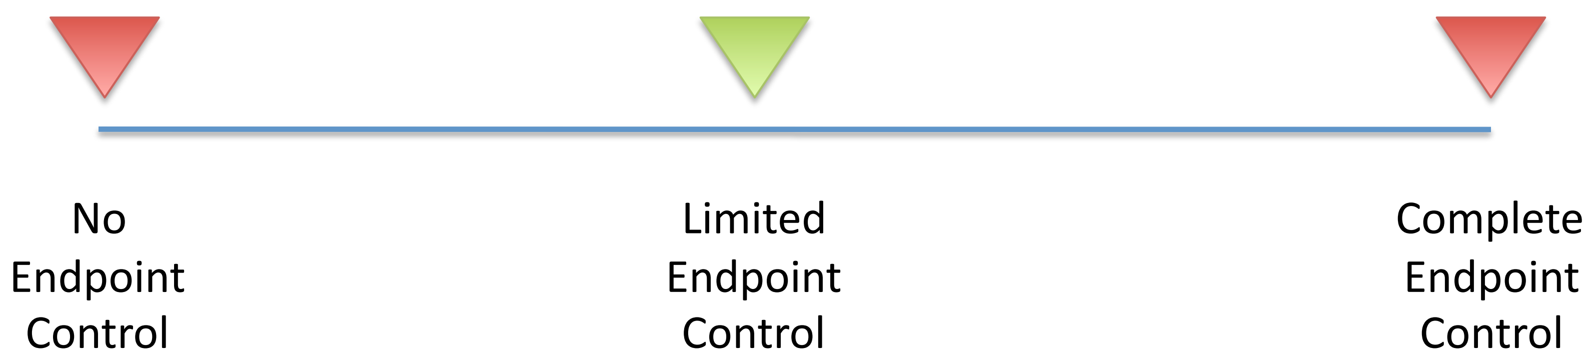

Problem 3: Transparency

Transparency: Internet Paths

- Today, sender cannot obtain guarantee that packet will travel along intended path
- Impossible to gain assurance of packet path
  - Because router forwarding state can be inconsistent with routing messages sent

Proposed Approach: Packet-Carried State

- Packets carrying forwarding information provides path transparency
  - Note: orthogonal issue to path control, as network can still define permitted paths

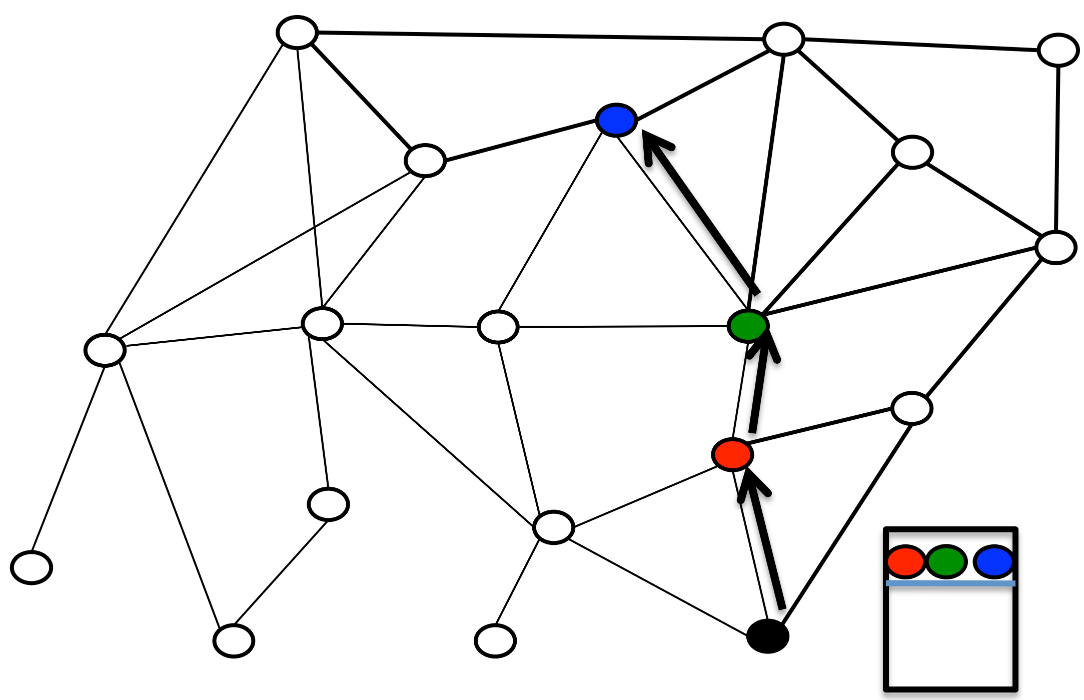

Problem 4: Availability

Poor Availability

- Well-connected entity: 99.9% availability (86 s/day unavailability) [Katz-Bassett et al., Sigcomm 2012]
- Numerous short-lived outages due to BGP route changes
  - **Route convergence** delays
- Outages due to misconfigurations
- Outages due to attacks
  - E.g., prefix hijacking, DDoS

Is a 10s Outage per Day Harmful?

- 99.99% reliability → average 8.6 s/day outage
  - Level of availability achieved by Amazon datacenter
- Insufficient for many applications
  - Critical infrastructure command and control
    - E.g., air traffic control, smart grid control
  - Internet-based business
  - Financial trading / transactions
  - Telemedicine

Proposed Approach: Replace BGP

- Border Gateway Protocol (BGP) is the inter-domain routing protocol in today's Internet
- BGP(SEC) suffers several fundamental problems
  - Trust: Uses single root of trust (RPKI / BGPSEC)
  - Control: Almost no path choice by end points
  - Transparency: Impossible to obtain path guarantee
  - Availability
    - Frequent periods of unavailability when paths change
    - Slow convergence during iterative route computation
    - Susceptible to attacks and misconfigurations

Evolutionary vs. Revolutionary Change

- Revolutionary approach is **necessary**
  - Some problems are fundamental, cannot be fixed through evolution
- Revolutionary approach is **desirable**
  - A fresh redesign can cleanly incorporate new mechanisms
- Revolutionary technology change is **easy** through evolutionary deployment
  - If IP is relegated to provide local (intra-domain) communication, only a small fraction of border routers need to change to replace BGP
  - Simultaneous operation with current Internet possible
  - Strong properties provide motivation for deployment

Proposed Future Internet Architectures

- General FIAs
  - XIA: enhance flexibility to accommodate future needs
  - MobilityFirst: empower rapid mobility
  - Nebula (ICING, SERVAL): support cloud computing
  - NIMROD: better scale and flexibility for Internet
  - NewArch (FARA, NIRA, XCP)
- Content-centric FIAs: NDN, CCNx, PSIRP, SAIL / NETINF
- Routing security: S-BGP, soBGP, psBGP, SPV, PGBGP, H-NPBR
- Path control: MIRO, Deflection, Path splicing, Pathlet, I3, Segment Routing
- Others
  - SDN: flexible intra-domain networking
  - ChoiceNet, HLP, HAIR, RBF, AIP, POMO, RINA, ANA, …

SCION Project

- SCION: Scalability, Control and Isolation On Next-Generation Networks [IEEE S&P 2011]
- Current main team: Daniele Asoni, Lorenzo Baesso, David Barrera, Cristina Basescu, Chen Chen, Laurent Chuat, Sam Hitz, Jason Lee, Tae-Ho Lee, Yue-Hsun Lin, Steve Matsumoto, Chris Pappas, Raphael Reischuk, Stephen Shirley, Pawel Szalachowski, Yao Zhang

ETH zürich

SCION Architectural 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: operate in heterogeneous trust environment
- Transparent operation: Clear *what* is happening to packets and *whom* needs to be relied upon for operation
- Balanced control among ISPs, Senders, and Receiver
- Scalability, efficiency, flexibility

SCION Isolation Domain (ISD)

- SCION Isolation Domain requirements
  - Region which can agree on a common root of trust
  - Set of ISPs to operate Isolation Domain Core to manage ISD
    - Root of trust and Autonomous Domain (AD) certificates
    - Manage core path and beacon servers
  - Other ISDs need to agree to connect as peer or as provider
- Open research issue exactly how to best structure ISDs: political and legal issues arise
  - Possible partition is along geographical regions

SCION Isolation Domain (ISD)

- SCION Isolation Domain composition
  - ISD Core with ISD Core ADs
  - Other ISP ADs or end-domain ADs

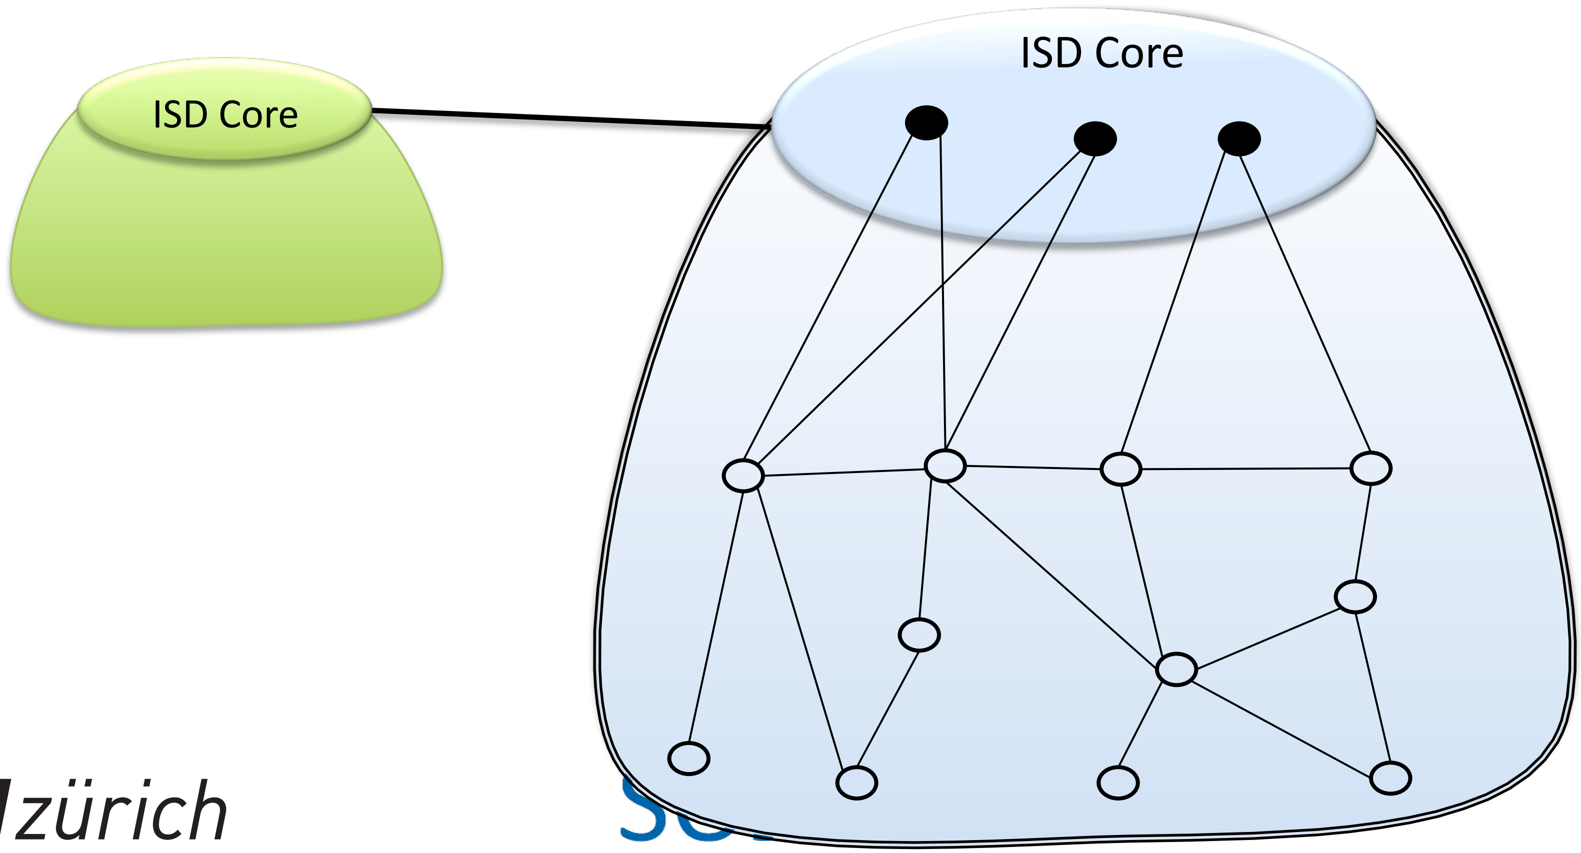

Beaconing for Route Discovery

- Periodic Path Construction Beacon (PCBs)
  - Scalable & secure dissemination of path/topological information from core to edge
  - K-wise multi-path flood to provide multiple paths

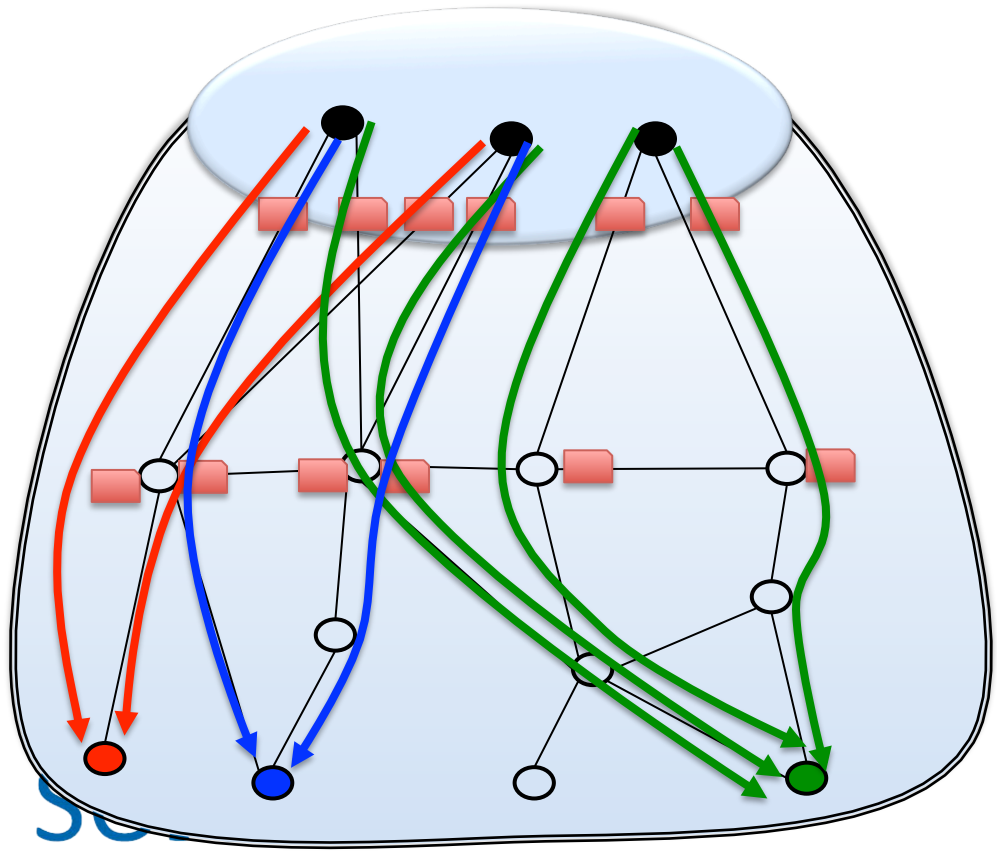

ETH zürich

SCION Forwarding (Data Plane)

- Domains register paths at DNS-like server in ISD Core
- End-to-end communication
  - Source fetches destination paths
  - Source path + destination path → end-to-end path
  - Packet contains forwarding information
- Advantages
  - Isolates forwarding from routing
  - No forwarding table at routers
  - Transparent forwarding
  - Balanced route control

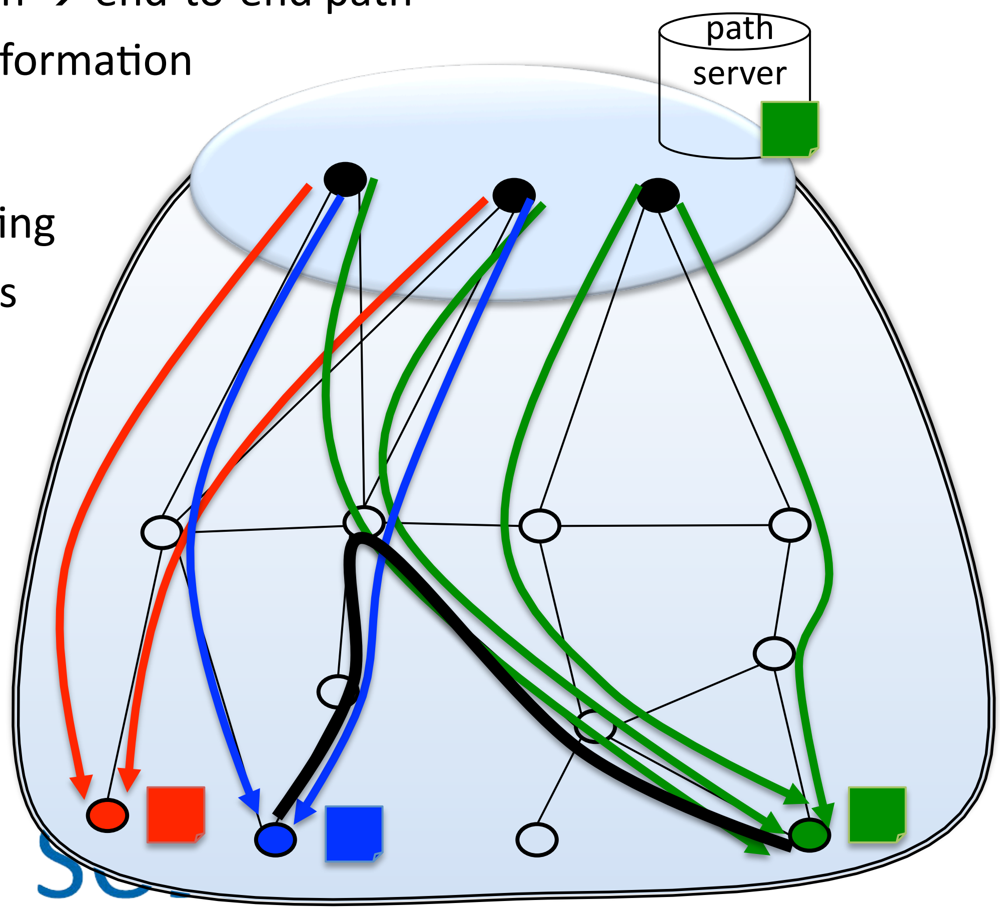

ETH zürich

Path Construction and Usage

- Path Construction Beacon (PCB) construction:
  PCB_1 = < T_{exp} Int_1 O_1 S_1 >
  Opaque field O_1 = Int_1 MAC_K(T_{exp} Int_1)
  Signature S_1 = { PCB_1 }$_{K'}$
- PCB_2 = < T_{exp} Int_1 O_1 S_1 Int_2 Int_3 O_2 S_2 >
  Opaque field O_2 = Int_2 Int_3 MAC_K(O_1 T_{exp} Int_2 Int_3)
  Signature S_2 = { PCB_2 }$_{K'}$
- AD receiving PCB_2:
  - Verify signatures
  - Use opaque fields O_1 O_2 to send packet to ISD Core

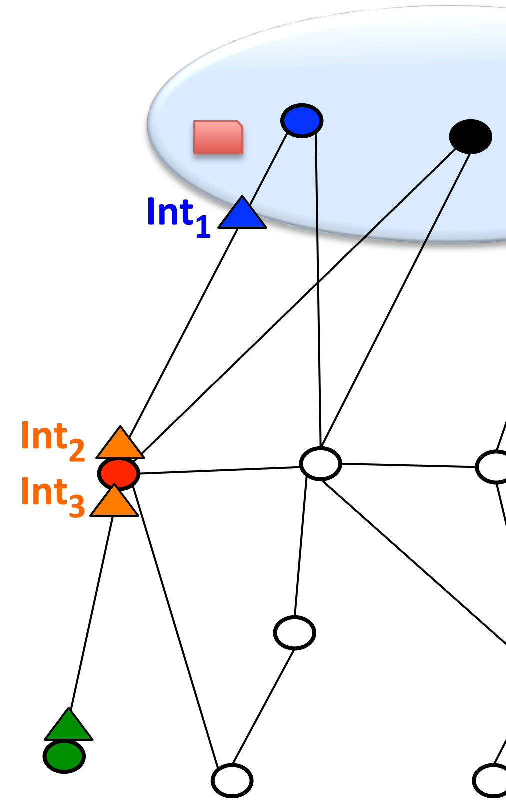

Inter-ISD Communication

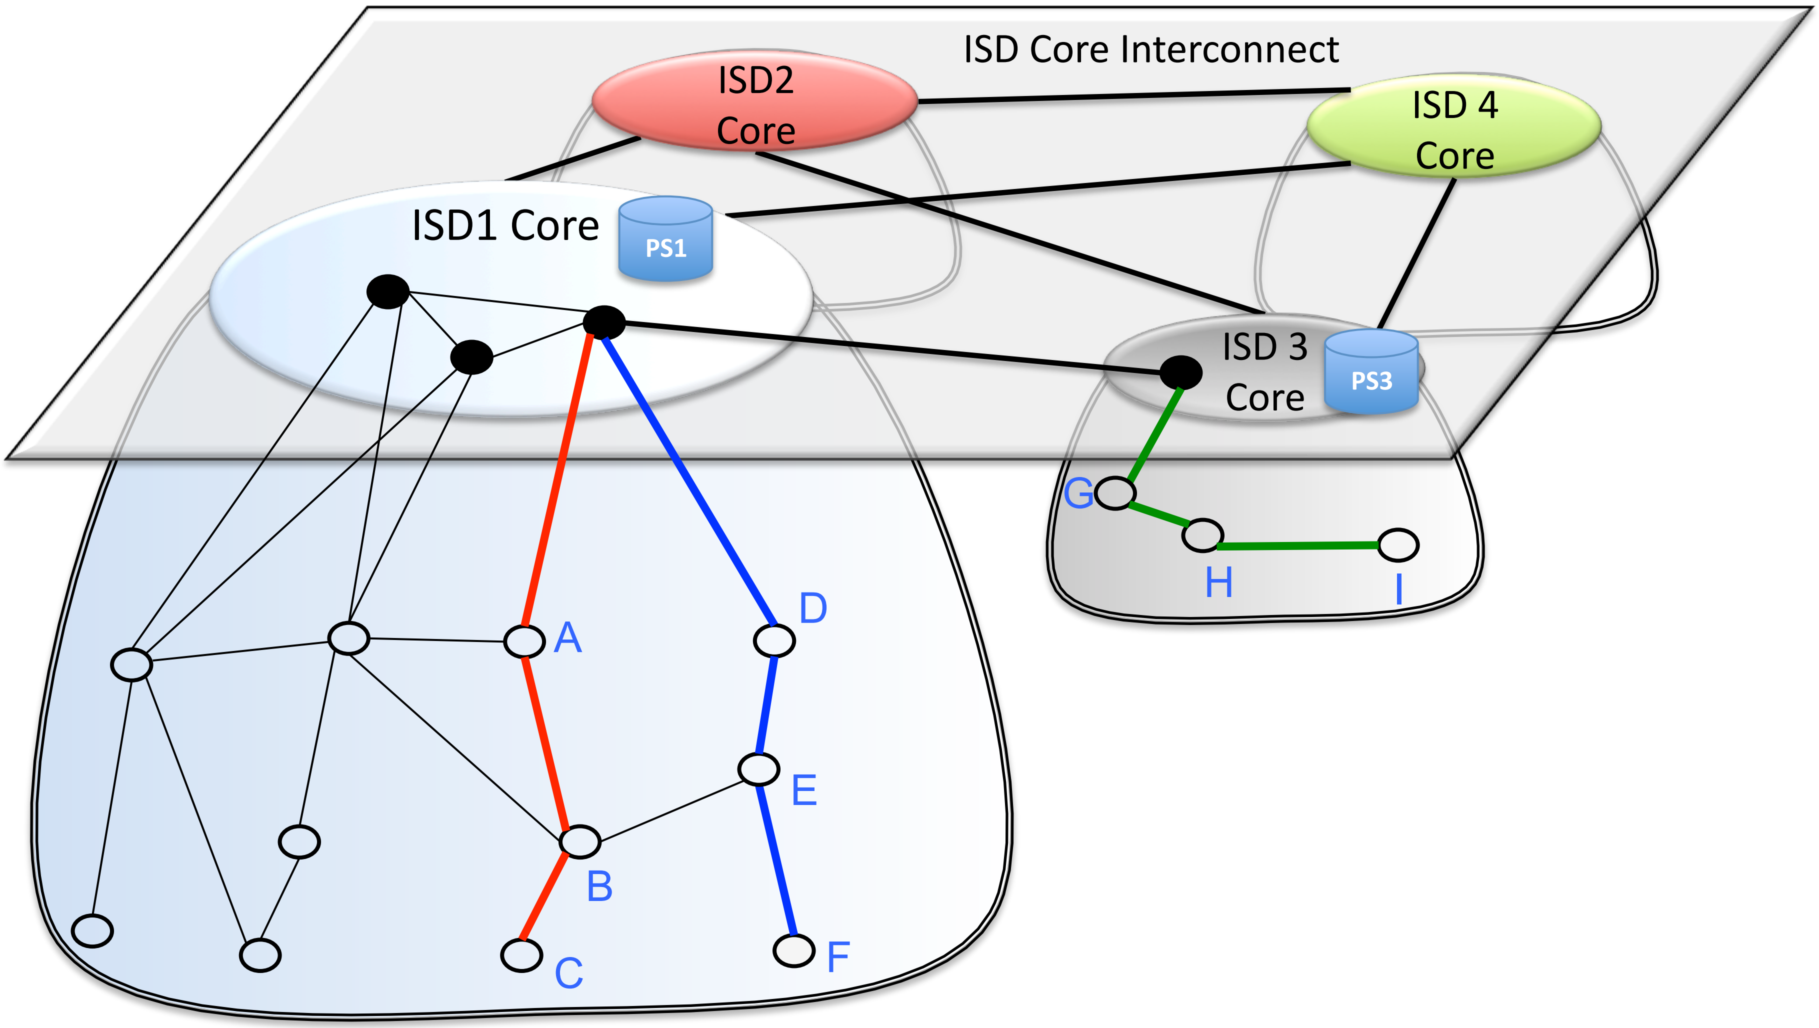

ETH zürich

SCION

Inter-ISD Communication

- ISD Cores recursively execute SCION beaconing to create paths amongst each other
  - Each ISD core initiates PCB to neighboring ISD cores
  - Propagates ISD Core PCBs to create inter-ISD-core path
- Endhosts can request path to reach any other ISD
- Endhosts combine up path + inter-ISD-core path + down path
  - Provides transparent operation, as path is known

Shortcuts through Peering Links

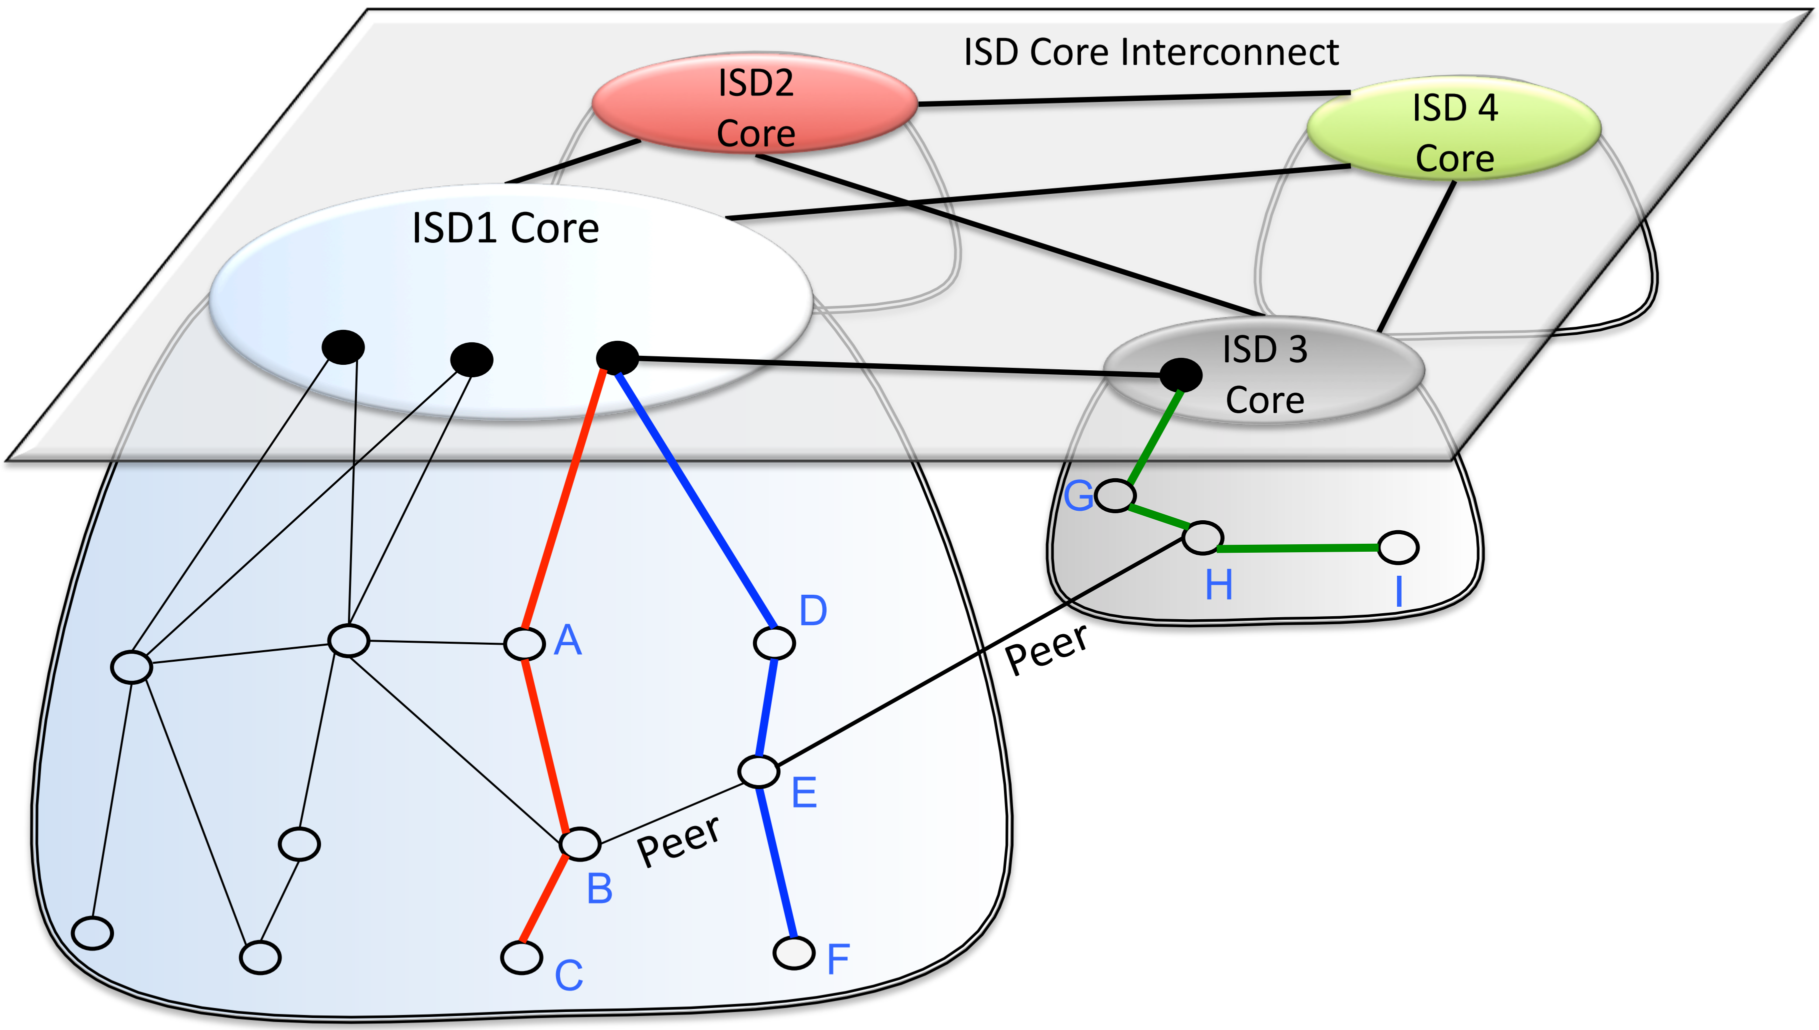

ETH zürich SCION

Handling Link Failures

- SCION clients use multi-path communication by default, other paths are likely to still function
- Path construction beacons are constantly sent, disseminating new functioning paths
- Link withdrawal message sent …
  - … upstream to cause path servers to remove paths with broken link
  - … downstream to cause beacon servers to remove paths with broken link

SCION Implementation Status

- Full V1.0 specification almost completed
- 3rd generation C/C++ implementation
- 4th generation: Python implementation
- High-speed router implementation switching 120Gbps on off-the-shelf PC
- So far ~50 person-years of effort invested
- Growing testbed

ETH zürich

SCION Packet Header

0-7	8-15	16-23	24-31	32-39	40-47	48-55	56-63
Type Vers. / Src Type	Src Type / Dst Type	Total Len		TS*	Curr OF*	Next Hdr.	HDR Len
Source Address (variable size)							
Destination Address (variable size)							
Info	EXP Time		ISD ID		hops	reserved	
Opaque Field (0)							
Next Ext.	Ext Hdr Len	extension-related data…					
… more extension-related data …							
Next Ext.	Ext Hdr Len	extension-related data…					
L4 Proto							

SCION Trust Root Management

- Each ISD manages their own trust roots
  - Used to create per-AD certificates
  - AD certificates used to verify beacon messages
- Trust Root Configuration (TRC) file serves as root of trust for ISD
  - TRC file specifies public keys of trust root and policy for TRC file update
  - Thresholds enable revocation and re-authentication of new TRC files
  - Beacon messages quickly disseminate new TRC files
- Assumption: ISDs cross-sign TRC files

Trust Root Config (TRC): ISD Root-of-Trust

- Each ISD has a TRC file
  - Each AD is verified based on trust roots in TRC

ISD EU
TRC file version N
A cert, B cert, C cert
CA1 cert
Update: at least 2
Sigs with keys of TRC
version N-1

{ D cert }K_A-1

{ CH ISD TRC }K_A-1

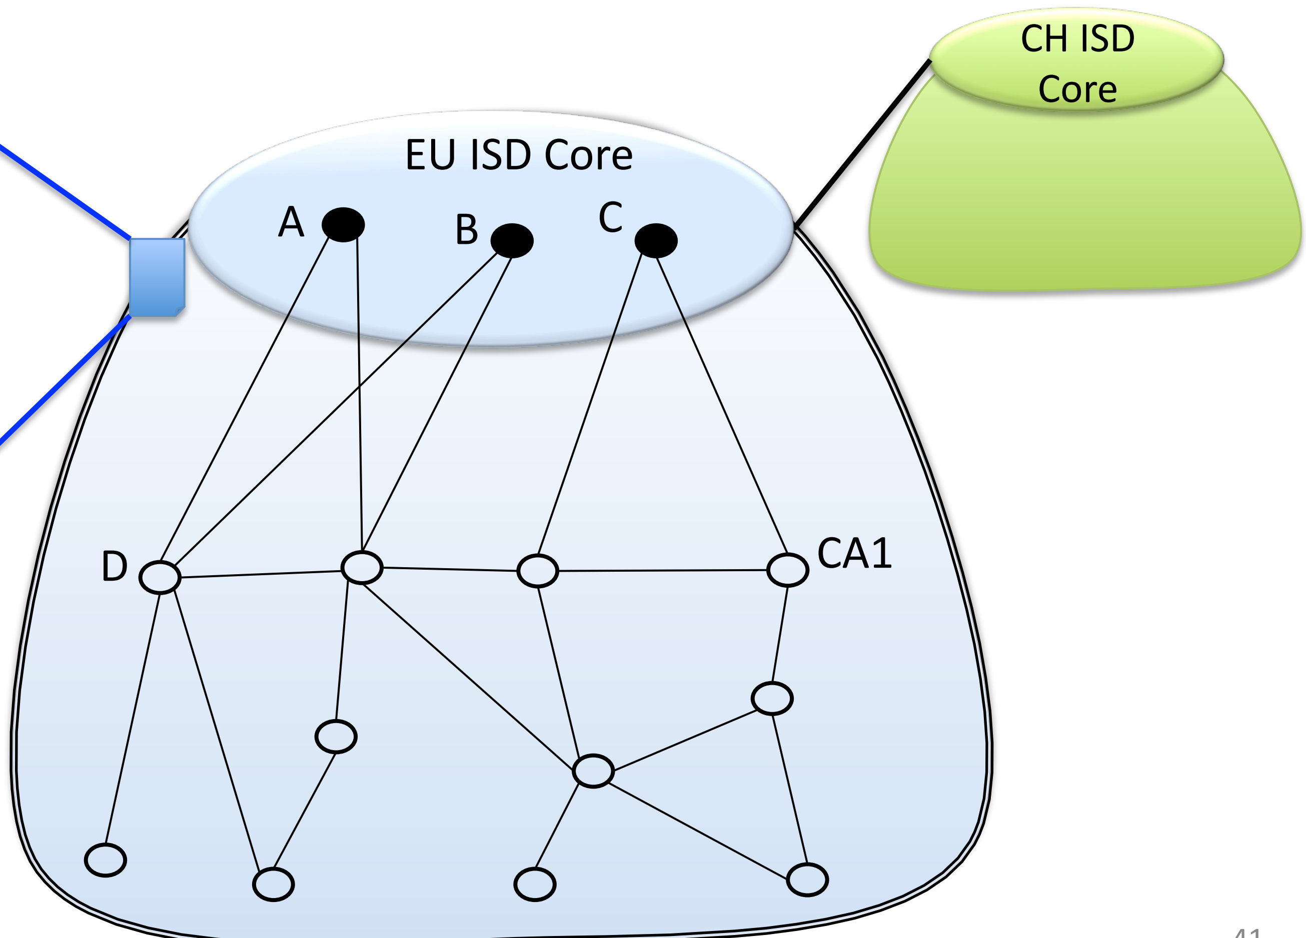

TRC File Update

- New TRC file version N+1 signed by threshold number of keys from version N
- SCION beaconing process distributes new TRC file

ISD EU
TRC file version N
A cert, B cert, C cert
CA1 cert
Update: at least 2
Sigs with keys of TRC version N-1

ISD EU
TRC file version N+1
A cert, B cert, C cert
CA1 cert
Update: at least 2
Sigs with keys of TRC version N

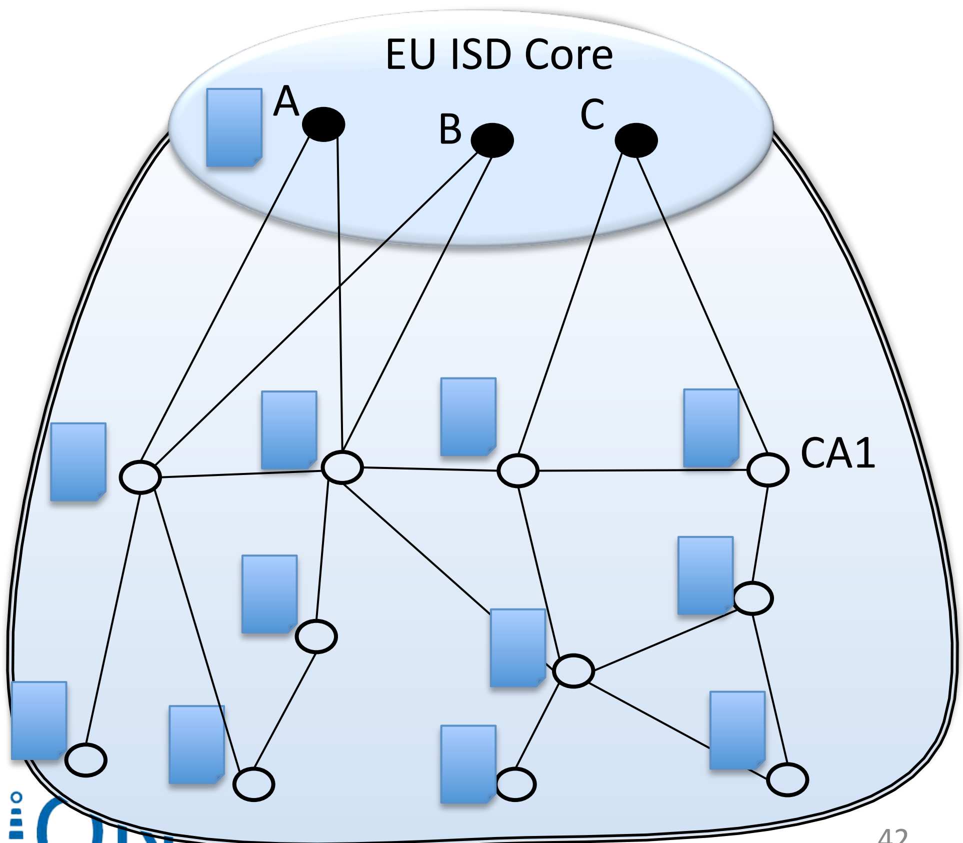

TRC File Summary

- Per-ISD TRC file enables heterogeneous trust roots
- TRC file update mechanism enables efficient update and revocation
  - Tens of seconds to update / revoke roots of trust network-wide
- Observation: network architecture should provide mechanism for updating trust roots!

Packet-Carried Forwarding State

- Observation: per-flow state on routers causes many issues
  - State exhaustion attacks [Schuchard et al., NDSS 2011]
  - State inconsistencies complicate protocol design (e.g., TTL to handle forwarding loops)
  - Complicates router design
- Mantra: **no per-flow state in the fast path**
  - Packet-carried forwarding state avoids per-flow state on routers

Uses of Packet-Carried Forwarding State

- Stable and predictable forwarding path in packet header tremendously beneficial
  - Lightweight anonymity and privacy [“LAP”, IEEE S&P 2012]
  - Stateless network capabilities for DDoS defense [“STRIDE”, AsiaCCS 2013]
  - Path validation [“OPT”, Sigcomm 2014]
  - Fault localization
  - Multipath forwarding

Incremental Deployment Aspects

- Current ISP topologies consistent with SCION ISDs
- Minor changes for ISPs
  - SCION edge router deployment
  - Beacon / certificate / path server deployment (inexpensive commodity hardware)
  - Regular MPLS/IP/SDN forwarding internally
  - IP tunnels connect SCION edge routers in different ADs
- Minor changes in end-domains
  - IP routing used for basic connectivity
  - SCION gateway enables legacy end hosts to benefit from SCION network

Incremental Deployment

- Only border routers need to adopt SCION

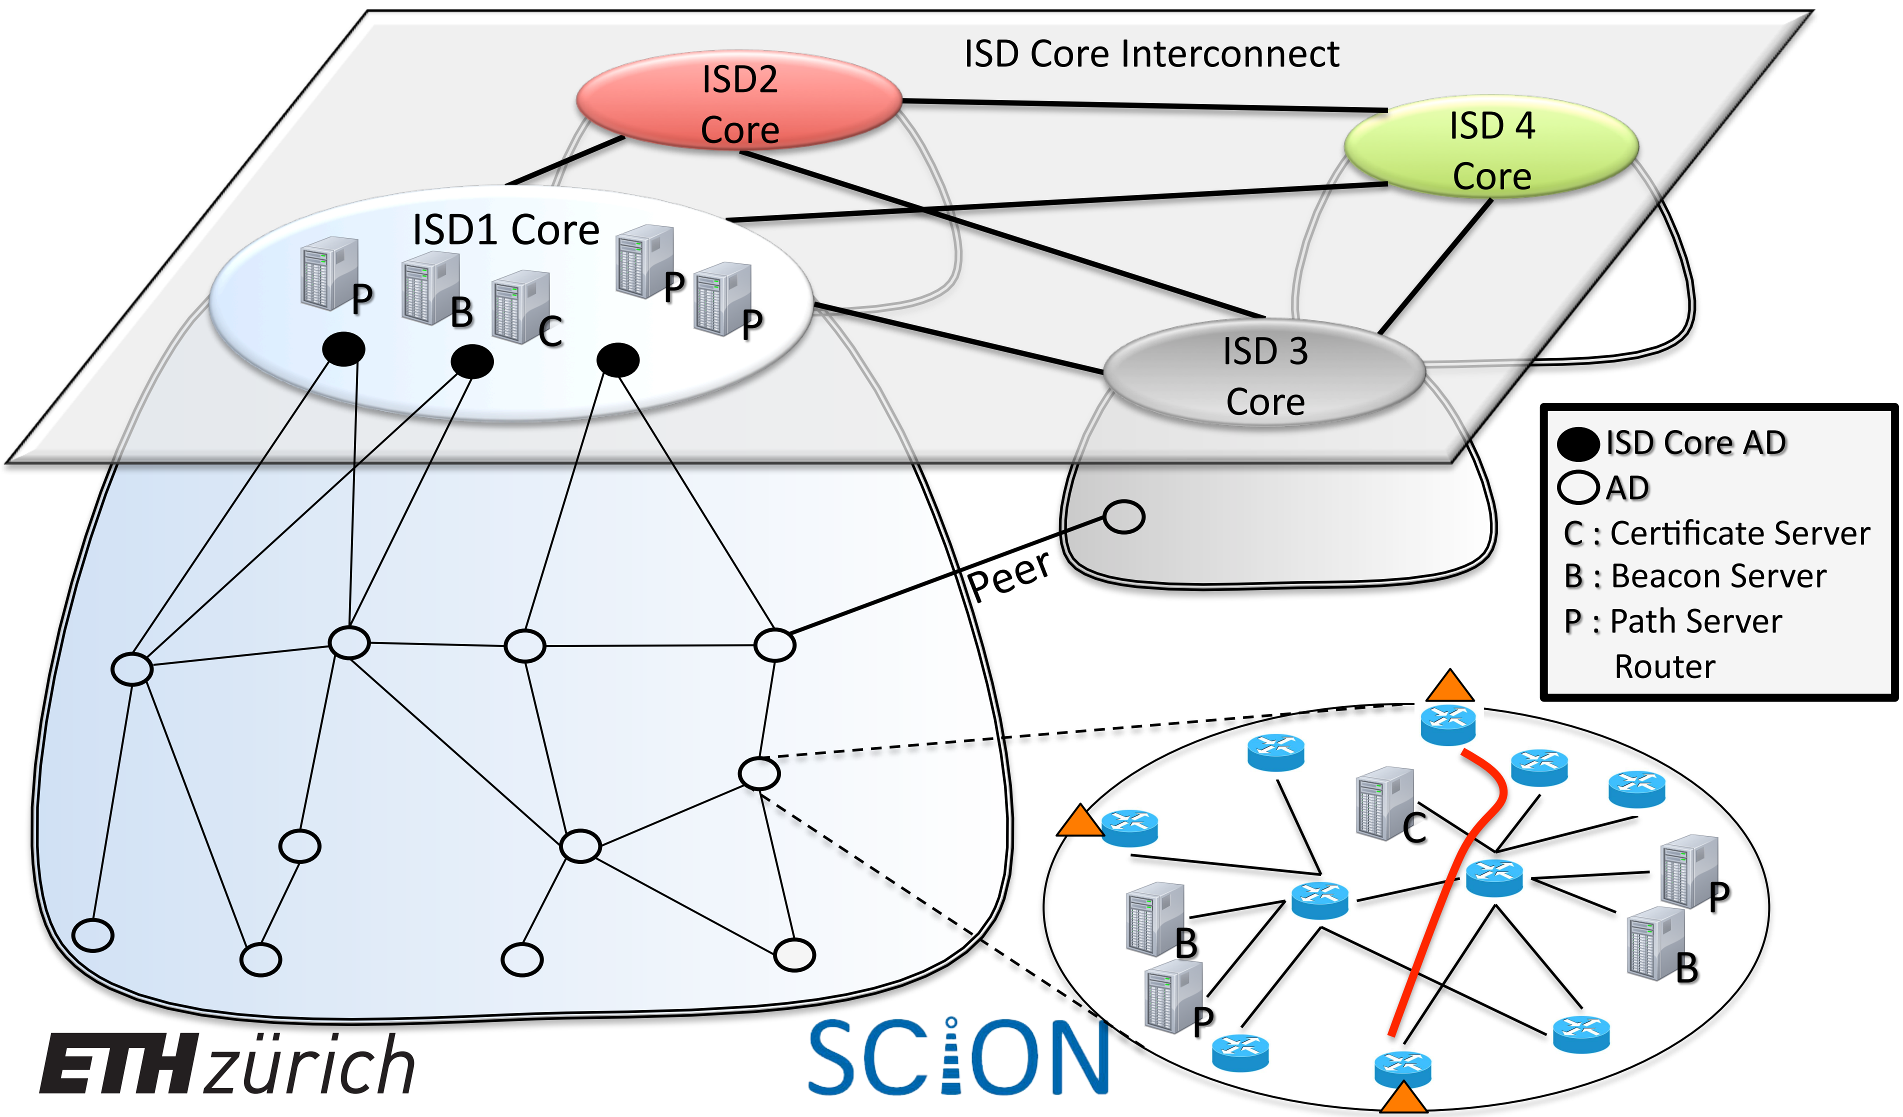

ETHzürich SCION

DENA Project

- Initial deployment without any changes to host

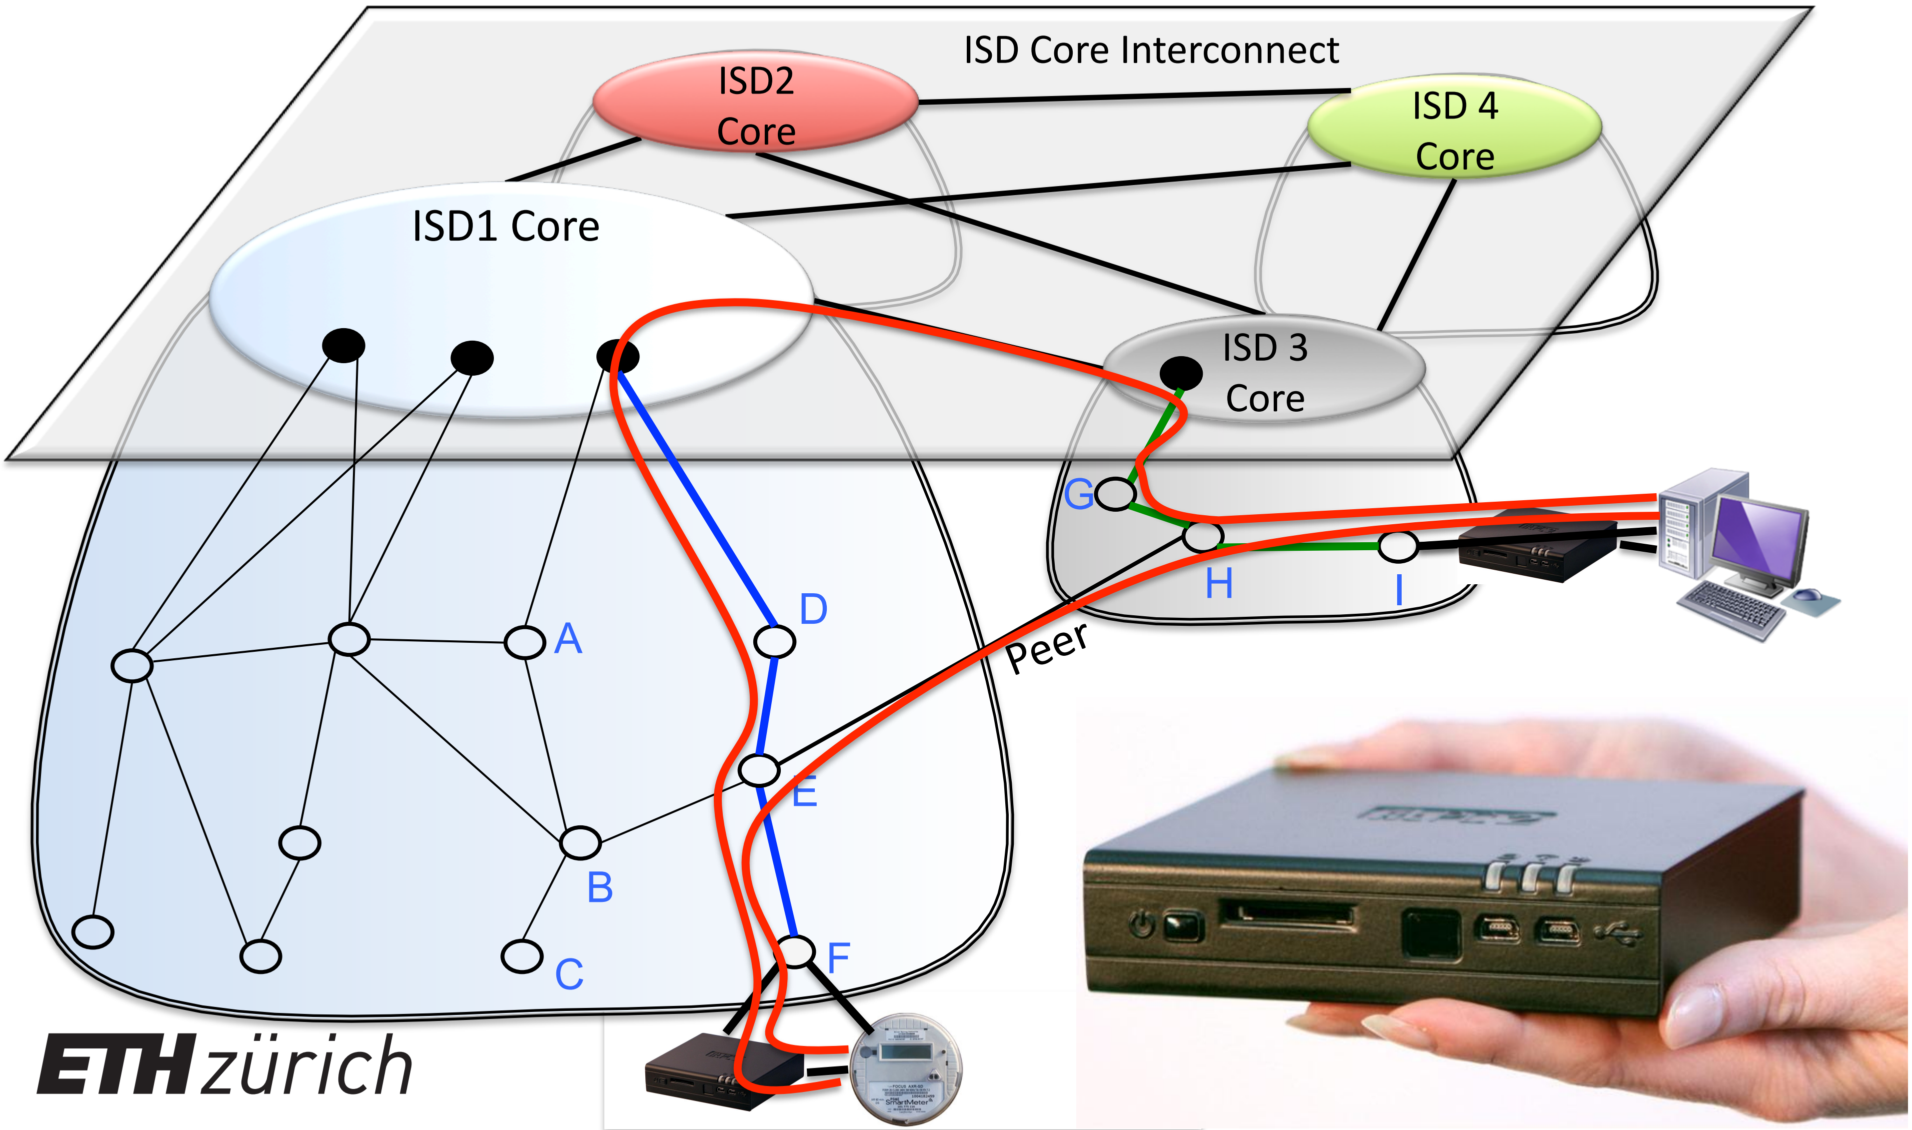

ETHzürich

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 Isolation Domain
- SCION is an isolation architecture only for the control plane, in the data plane it is a **transparency architecture**

Opportunities / Trends

- Mobility
  - SCION supports in-connection path update
  - Multipath system immediately makes use of new path
  - DNS / path server system enables dynamic updates
- SDN
  - SCION can work with SDN within domains
  - SCION has properties of an intra-domain SDN
- Content-centric communication support
- Cloud computing

SCION Dangers

- Too many top-level ISDs
  - Too many ISPs part of ISD core
- Large packet header size
  - Too many extensions used
- Higher complexity (Extensions, PKI)
- Extremely high path fluctuations, changes

SCION Stakeholder Pros and Cons

- Manufacturers
  - ✓ Sale of additional equipment
- ISPs
  - ✓ New revenue streams through service differentiation
  - ✓ High-availability service offerings, powerful DDoS defenses
  - ✓ Inter-domain Service Level Agreement (SLA)
  - ✓ Resilient to attacks and configuration errors
  - ✓ Incremental update, only new edge routers needed
  - ✓ Same business models as with BGP (peering links, customer – provider)
  - ✓ BGP routing policies can be emulated, extended
  - ⚡ Employee training: new equipment, new protocols
- Consumers
  - ✓ Faster webpage downloads
  - ✓ Efficient anonymous communication
  - ✓ Trust agility, choice of trust roots
  - ⚡ Software / HW upgrade
- Government
  - ✓ High reliability and availability for critical services
  - ✓ Selectable roots of trust, no single global root of trust
  - ✓ Verifiable router hardware

Conclusion

- Deployment of a new Internet architecture is necessary and possible
  - High-value Internet uses need strong network properties
  - New architecture can run along with current Internet
- Community effort needed to solve abundance of research challenges
  - Reliable operation with mutually untrusted operators
  - Anonymous communication
  - Network neutrality
  - DDoS attacks

ETH zürich

Thanks to SCION Team Members!