An Internet Architecture for the 21st Century

Adrian Perrig

Network Security Group, ETH Zürich

ETH zürich

SCION

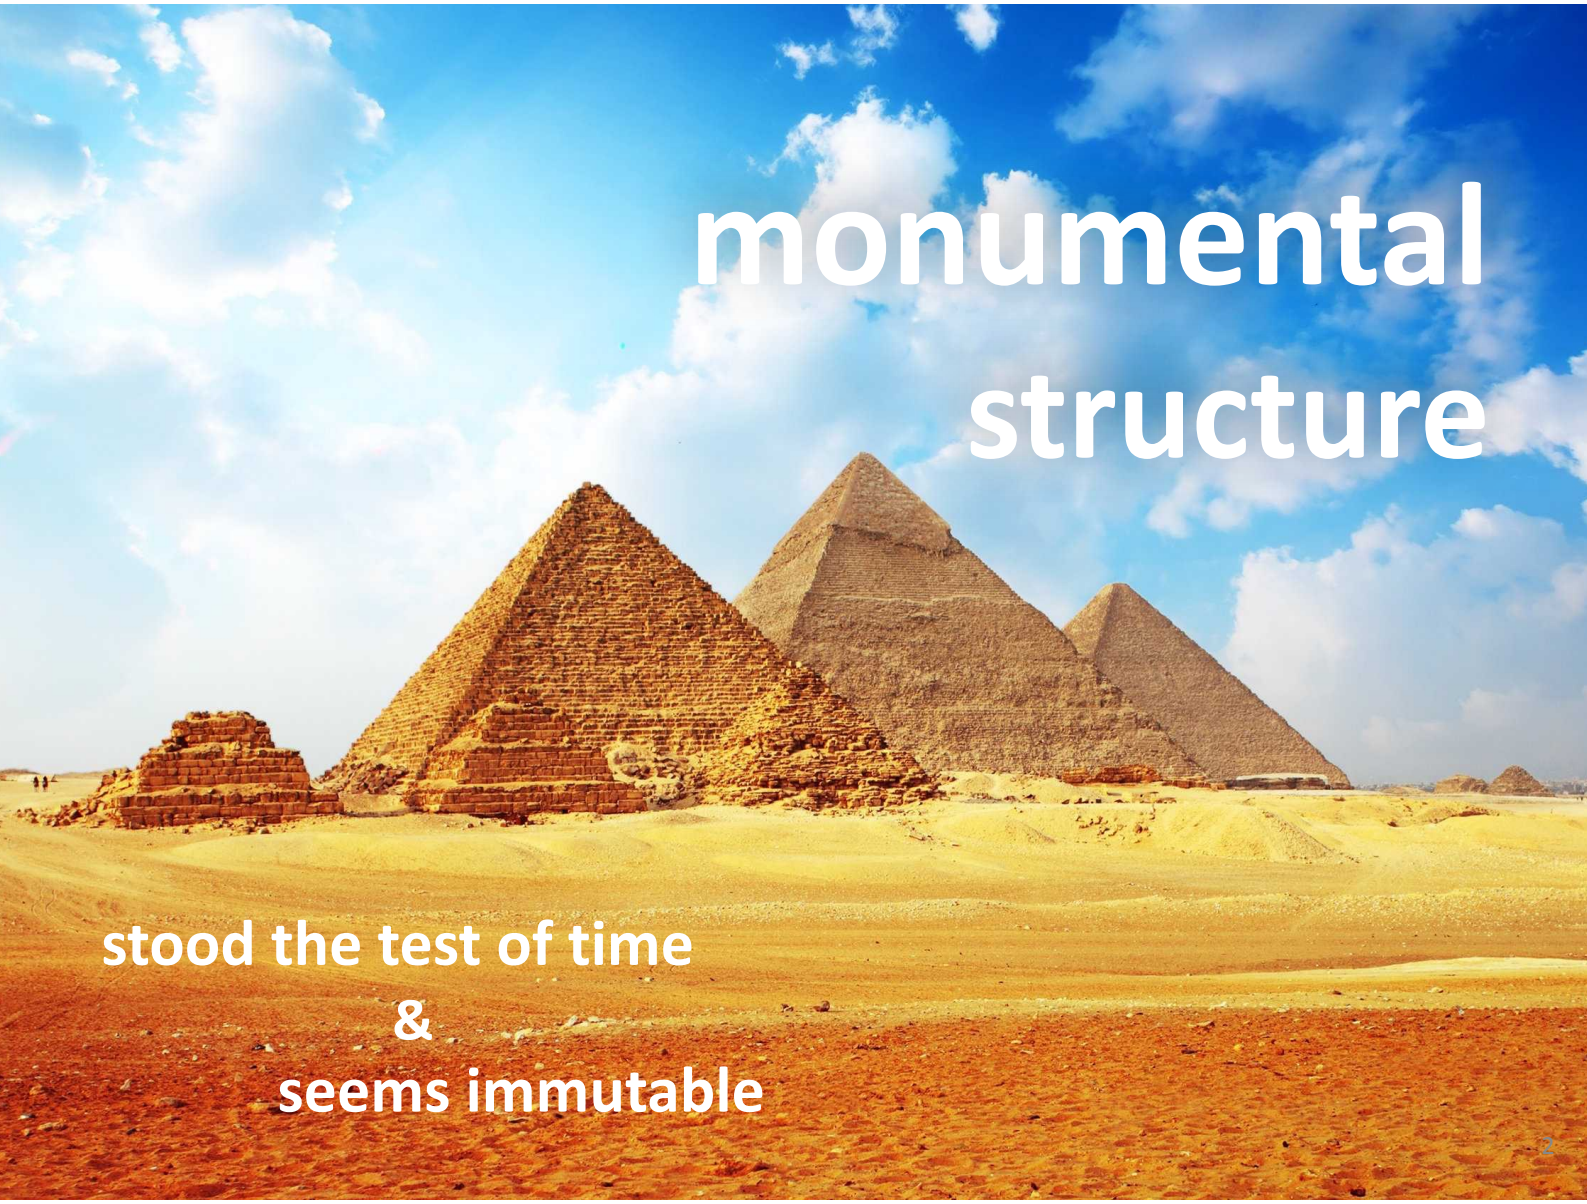
monumental
structure
stood the test of time
&
seems immutable

Just like today's Internet ?

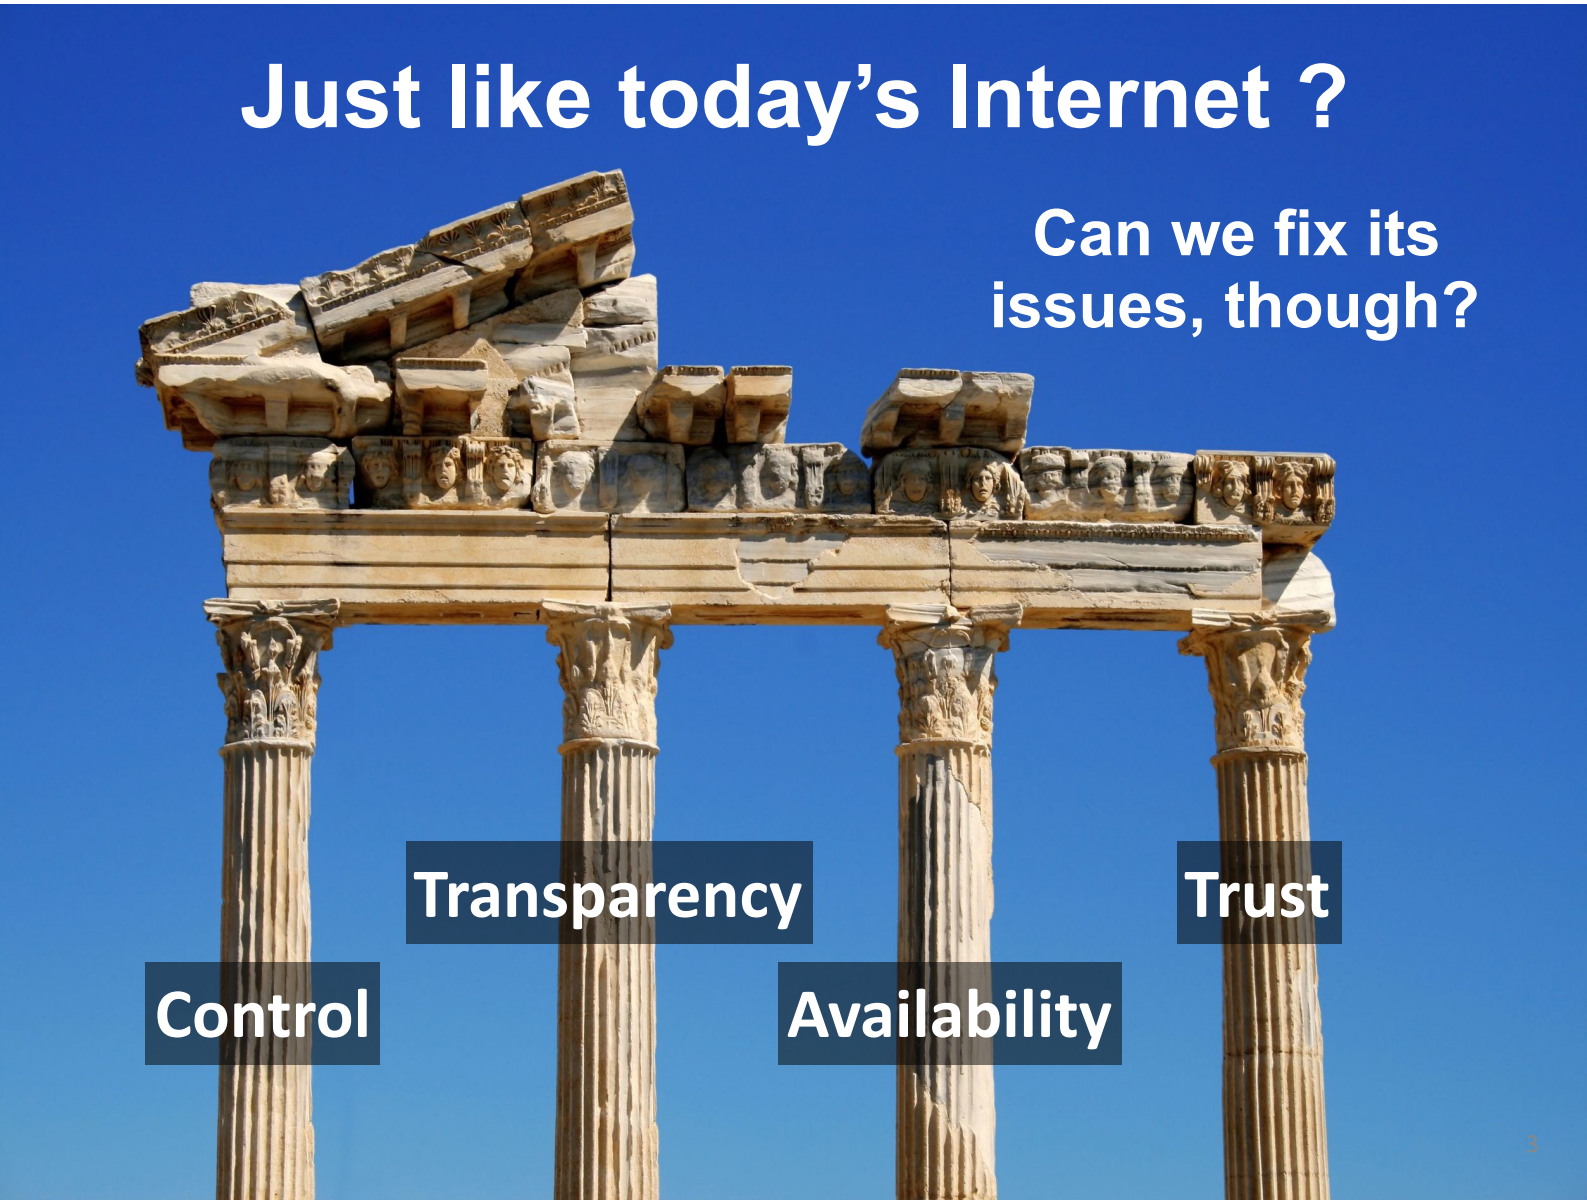

Problem 1: Non-Scalability of Trust

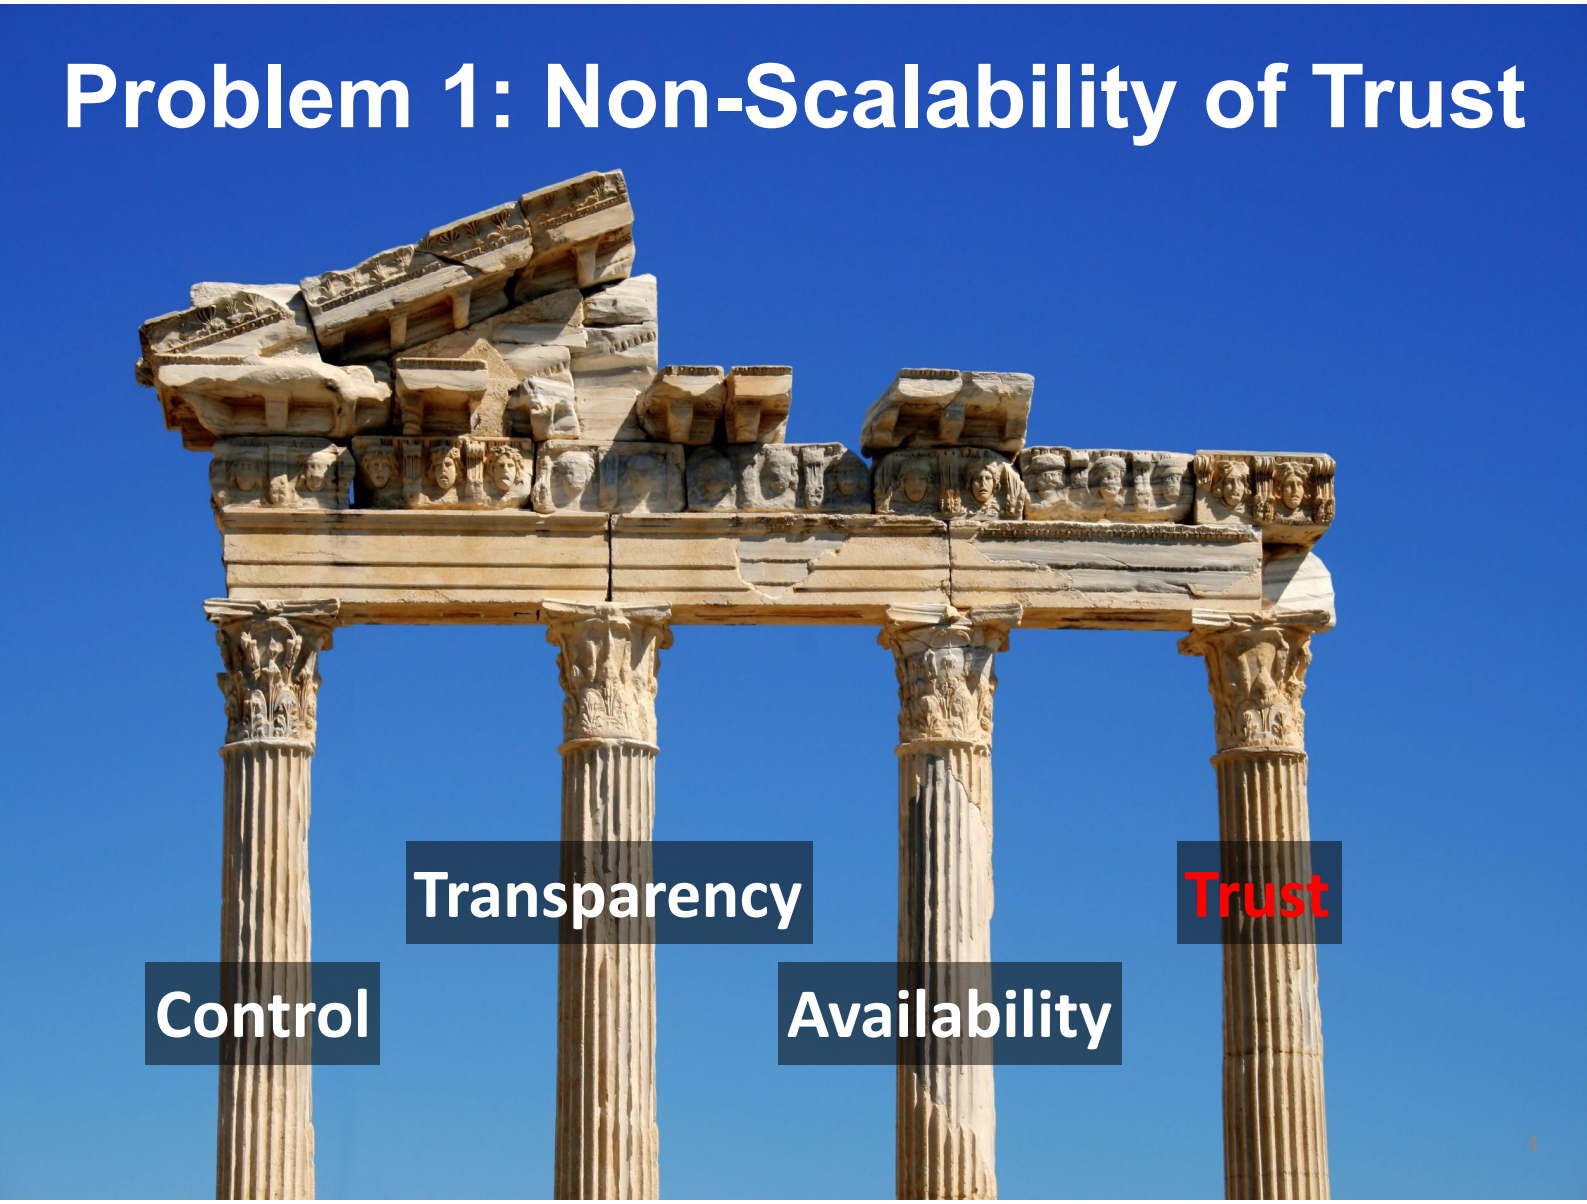

Pervasive Trust in Early Internet

" **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 anymore
- 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 that root of trust
- Problems
  - Entire world needs to agree on one entity to hold root of trust
  - Single point of failure
  - Inefficient revocation / update

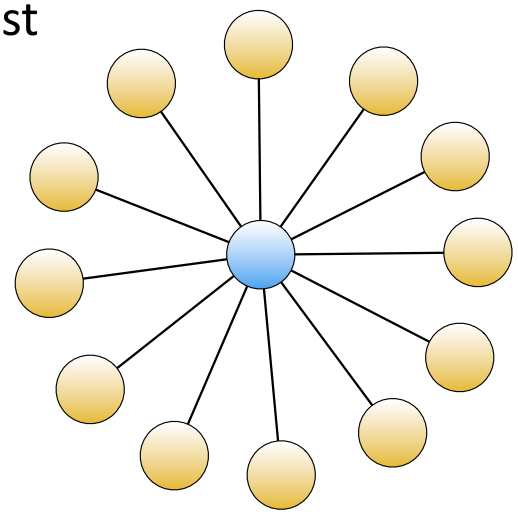

Oligarchy Model for Trust Root

- **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/updates are handled through OS or browser update

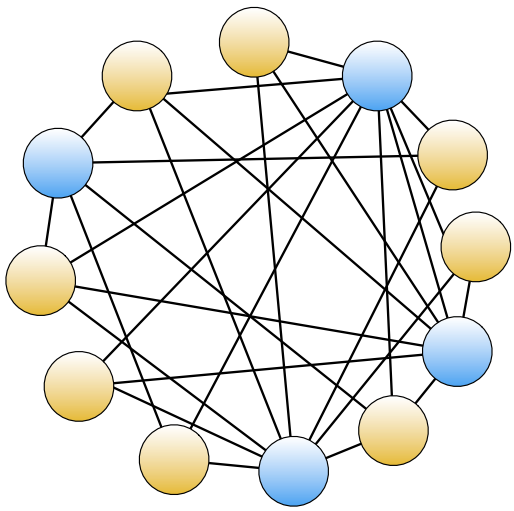

ETH zürich SCION

Global Trust leads to Kill Switch

- Current Internet has several “Kill Switches”
  - BGP: BGP hijacking
  - DNS: TLD redirection
  - BGPSEC: AS key revocation
  - DNSSEC: TLD key revocation
- Can we design networks without kill switch?

Proposed Approach: Isolation Domains

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

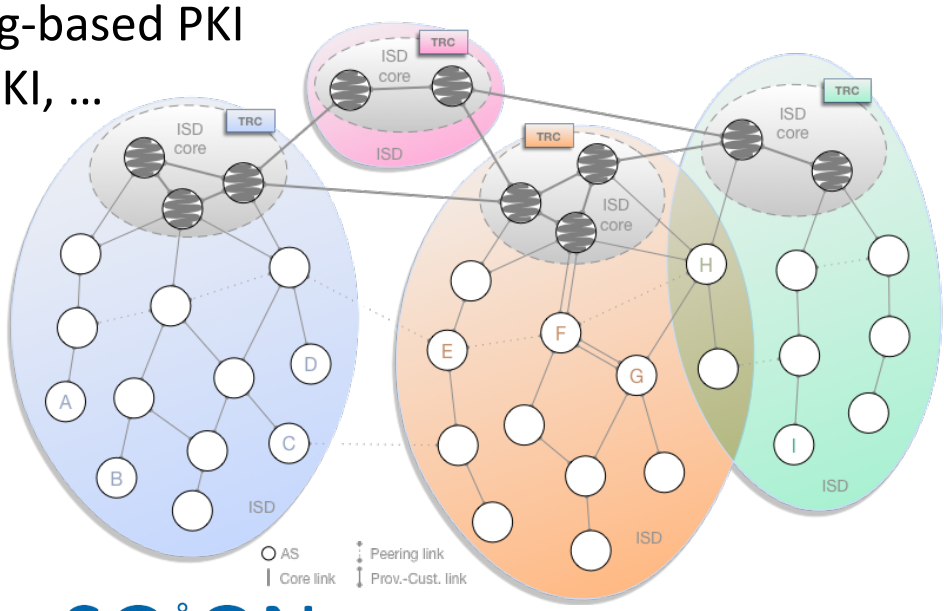

ETH zürich

SCION

Problem 2: Control

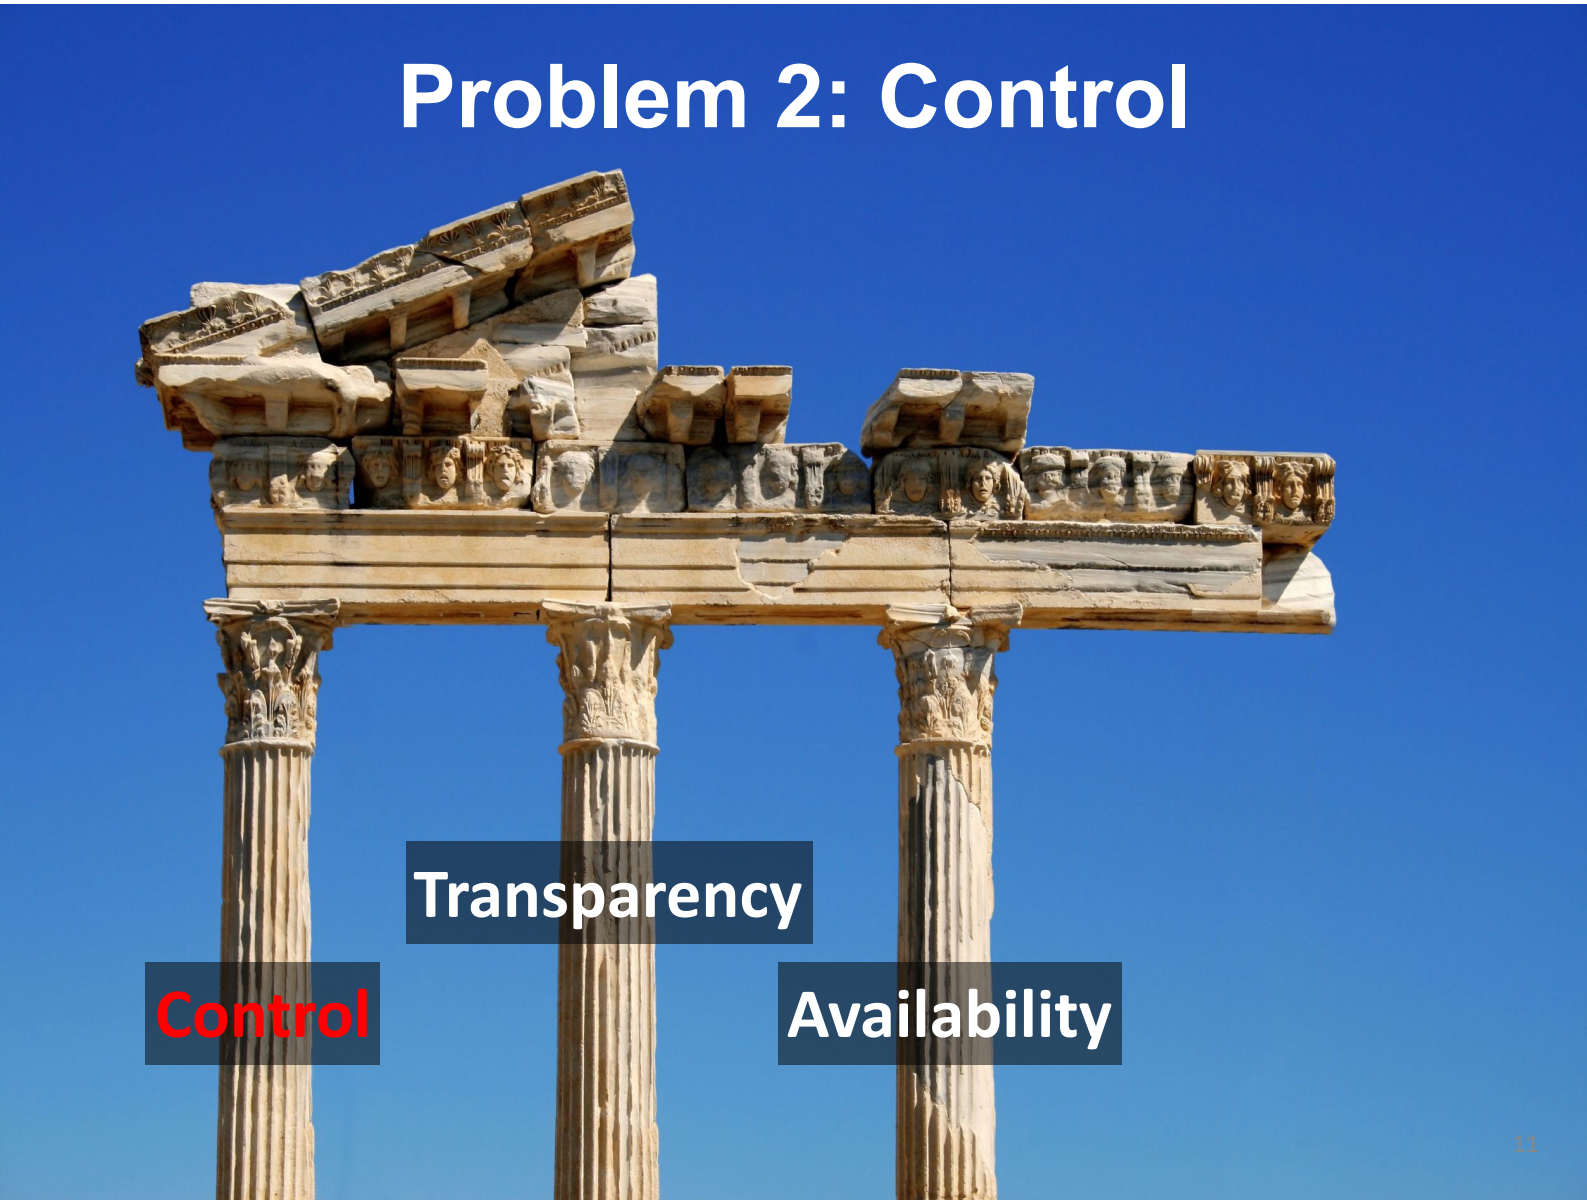

Aspects of Control

- We discuss here two aspects of control
  - Path control
  - Bandwidth control (DDoS attack defense)

Who controls Internet Paths?

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

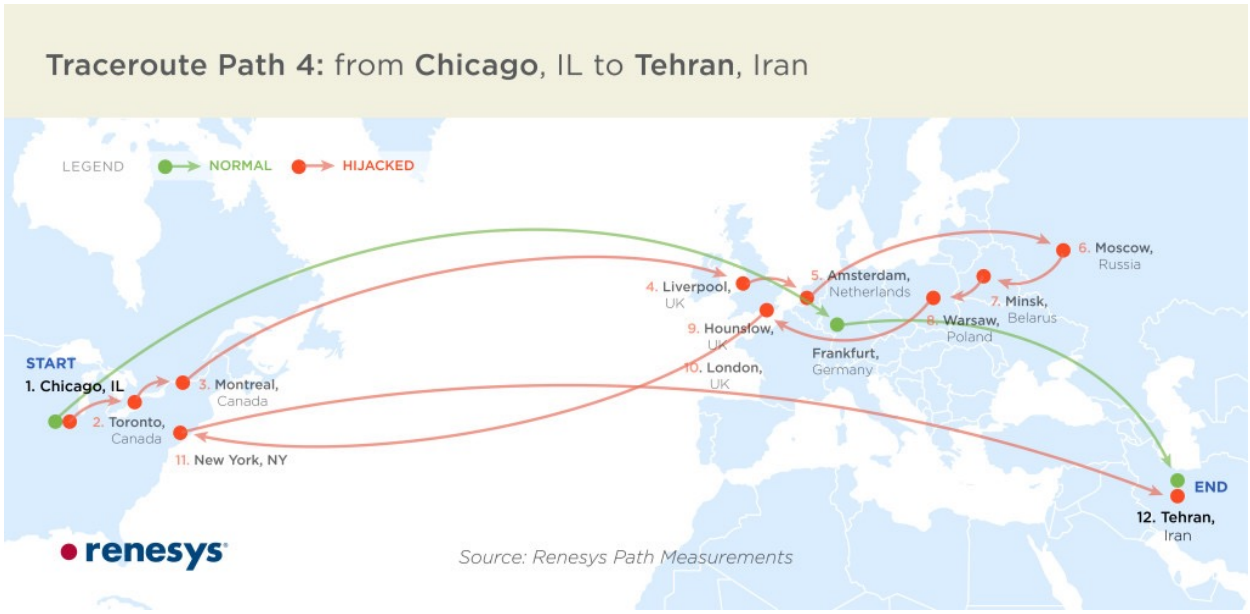

Who should control Paths?

- 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?

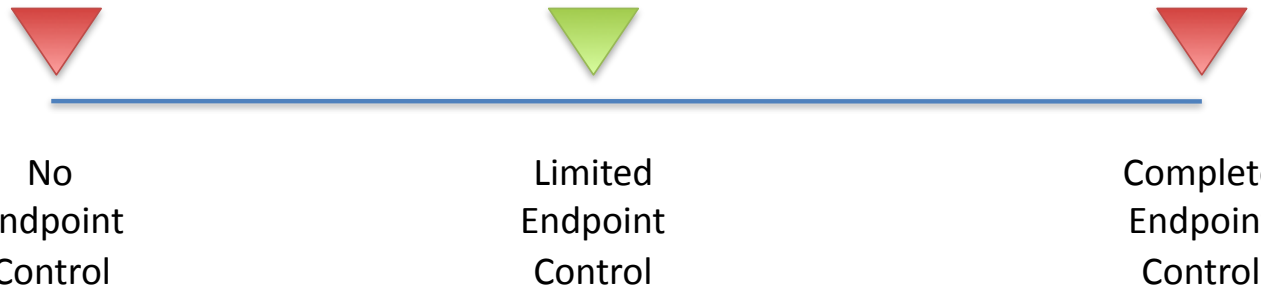

Absence of Bandwidth Control

- Today: no way to turn off malicious sender who floods victim with traffic
- Attackers use large botnets to send unwanted traffic to victims
  - Amplification attacks further increase traffic volume
  - N^2 attacks will be used in the future to evade all current DDoS defenses

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

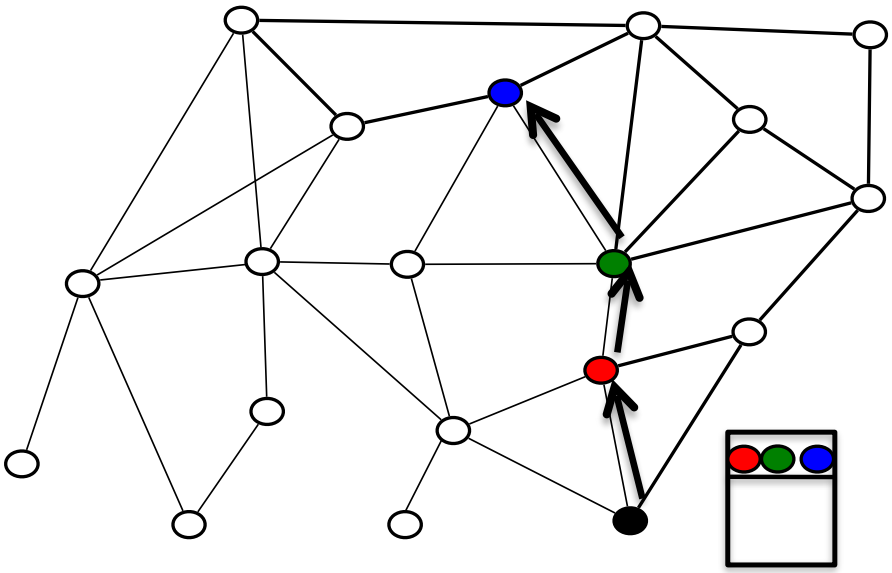

Problem 4: Availability

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

SCION Project

- **S**calability, **C**ontrol, and **I**solation **O**n **N**ext-Generation Networks [IEEE S&P 2011, CCS 2015, NDSS 2016]
- Current main team: Daniele Asoni, David Barrera, Chen Chen, Laurent Chuat, Sam Hitz, Jason Lee, Tae-Ho Lee, Steve Matsumoto, Chris Pappas, Adrian Perrig, Raphael Reischuk, Stephen Shirley, Pawel Szalachowski, Brian Trammell, Ercan Ucan

ETH zürich SCION

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**

ETH zürich

SCION

SCION Isolation Domain (ISD)

- SCION Isolation Domain **requirements**
  - **Region** that can agree on a common root of trust
    - groups a number of ASes
  - Set of ISPs to **operate** Isolation Domain Core to manage ISD
    - Certificates for roots of trust
    - Manage core path and beacon servers
  - Other ISDs need to agree to connect as peer or as provider
- Open research issue 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 ASes
  - Other ISP ASes or end-domain ASes

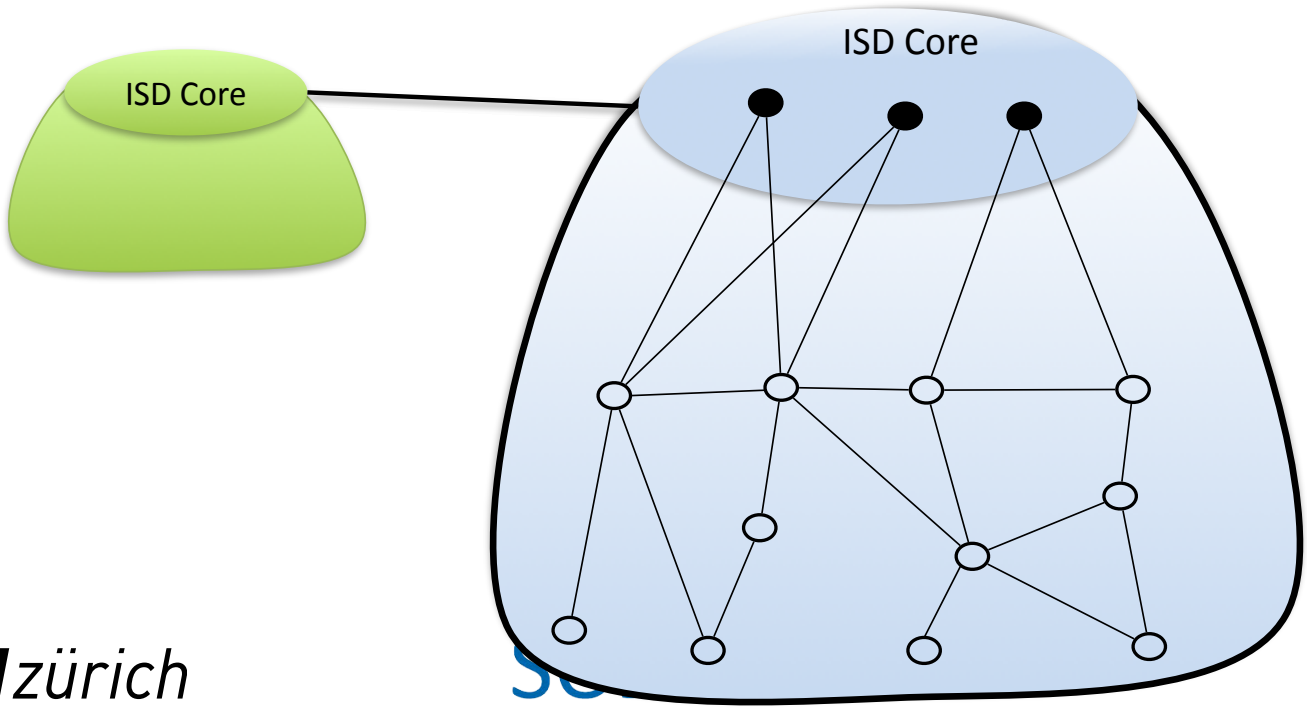

Beaconing for Route Discovery

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

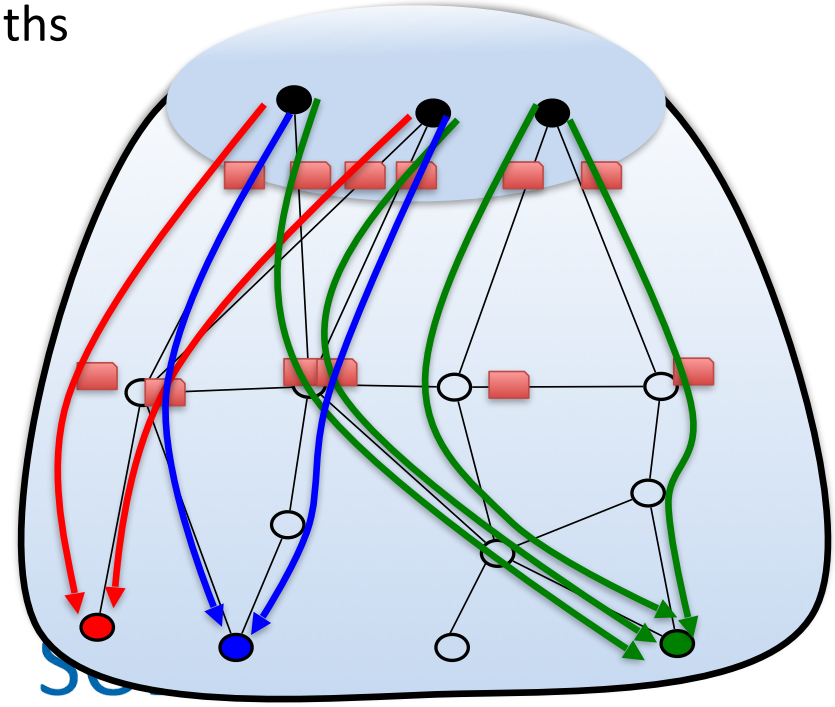

SCION Forwarding (Data Plane)

- 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**
  - **Path transparency**
  - **Balanced route control**

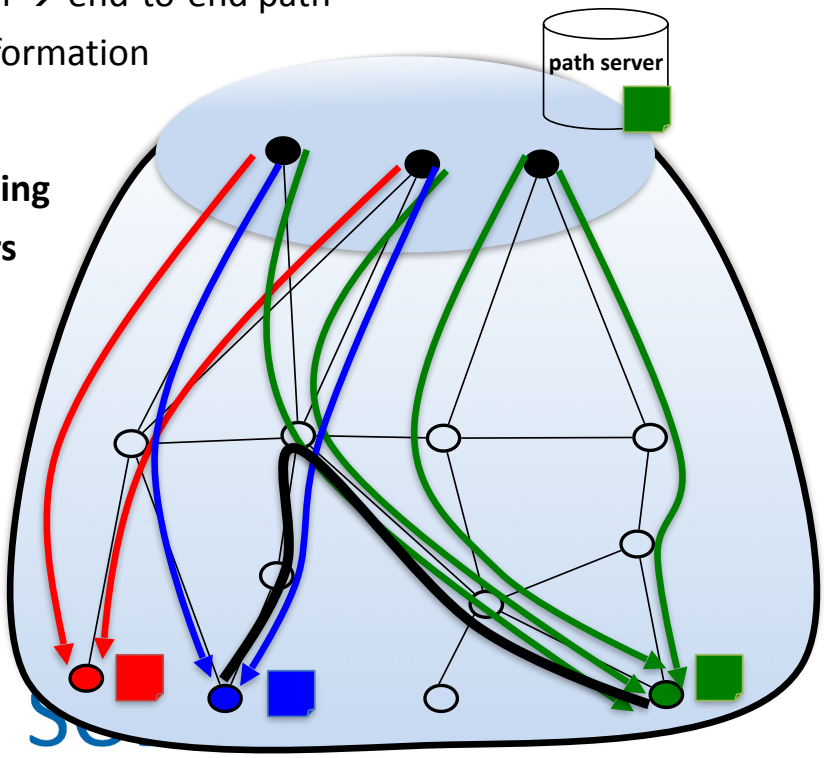

Path Construction and Usage

- Path Construction Beacon (PCB) construction:

  $\mathbf{PCB_1}$ = < T_{exp} Int_1 $\mathbf{OF_1}$ $\mathbf{S_1}$ >

  Opaque field $\mathbf{OF_1}$ = Int_1 MAC_K(T_{exp} Int_1)

  Signature $\mathbf{S_1}$ = { PCB_1 }$_{K'}$

- $\mathbf{PCB_2}$ = < T_{exp} Int_1 OF_1 S_1 Int_2 Int_3 $\mathbf{OF_2}$ $\mathbf{S_2}$ >

  Opaque field $\mathbf{OF_2}$ = Int_2 Int_3 MAC_K(T_{exp} Int_2 Int_3 OF_1)

  Signature $\mathbf{S_2}$ = { PCB_2 }$_{K'}$

- AS receiving PCB_2:
  - Verify signatures
  - Use opaque fields O_1 O_2 to send packet to ISD Core

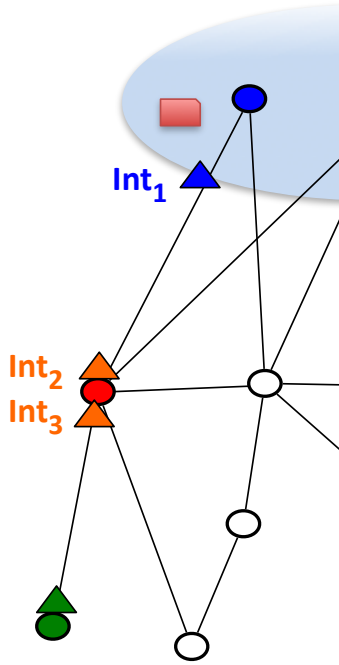

ETH zürich SCION

Inter-ISD Communication

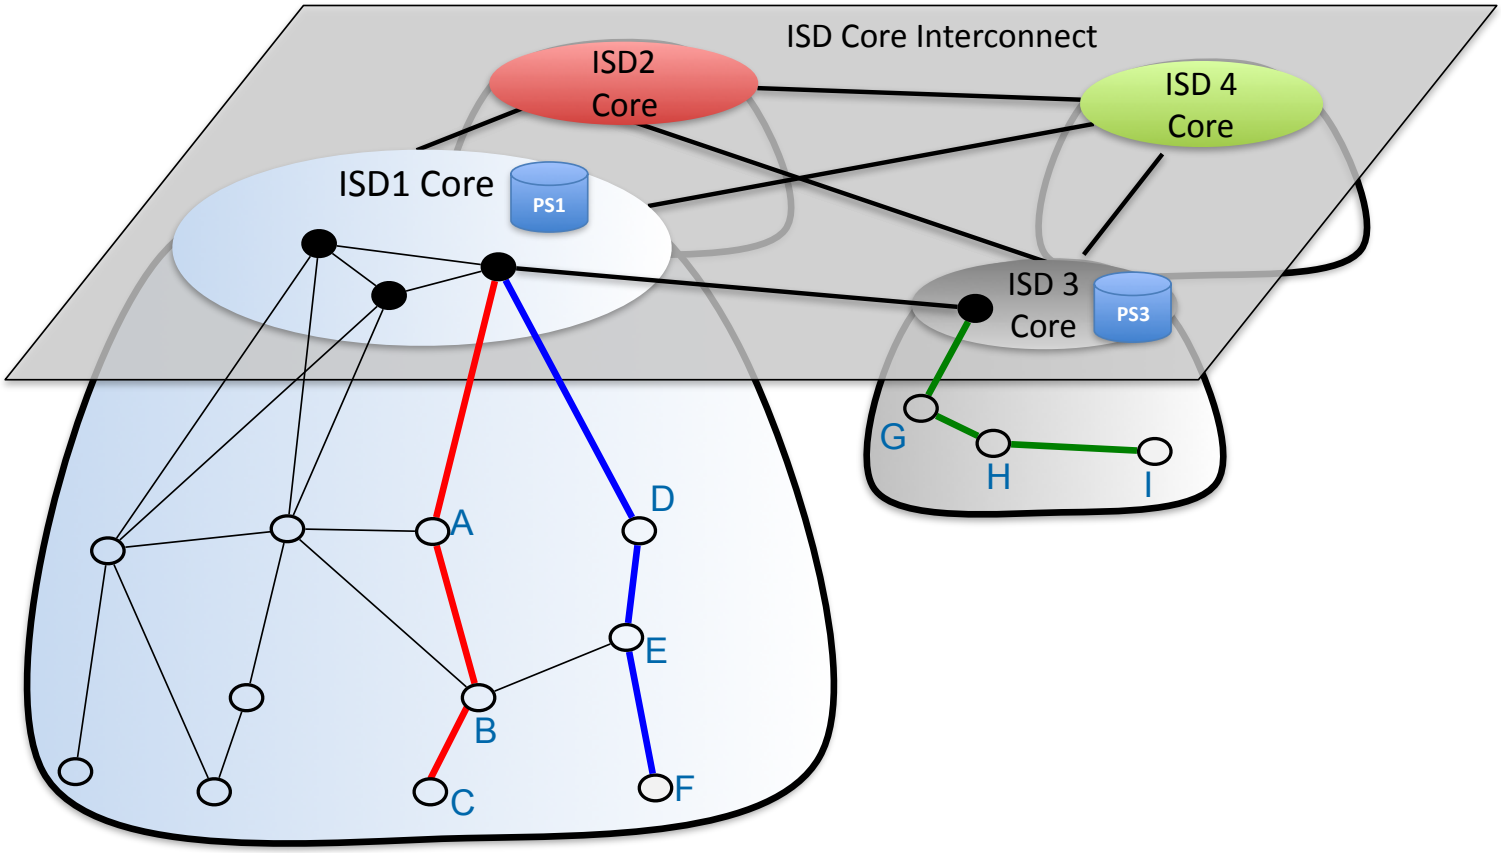

ETH zürich

SCION

Shortcuts through Peering Links

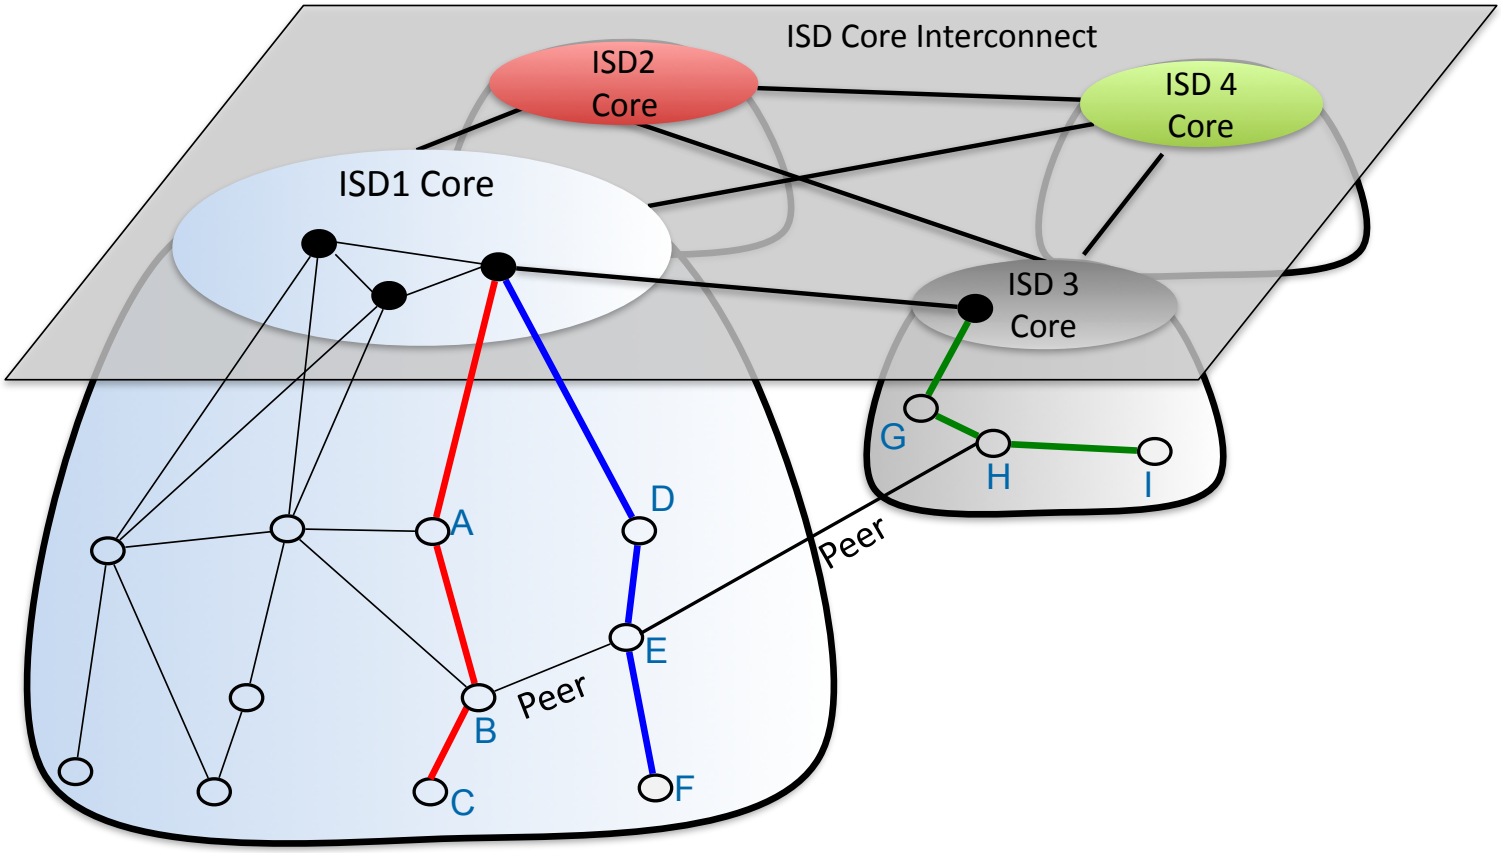

ETH zürich

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 Extensions

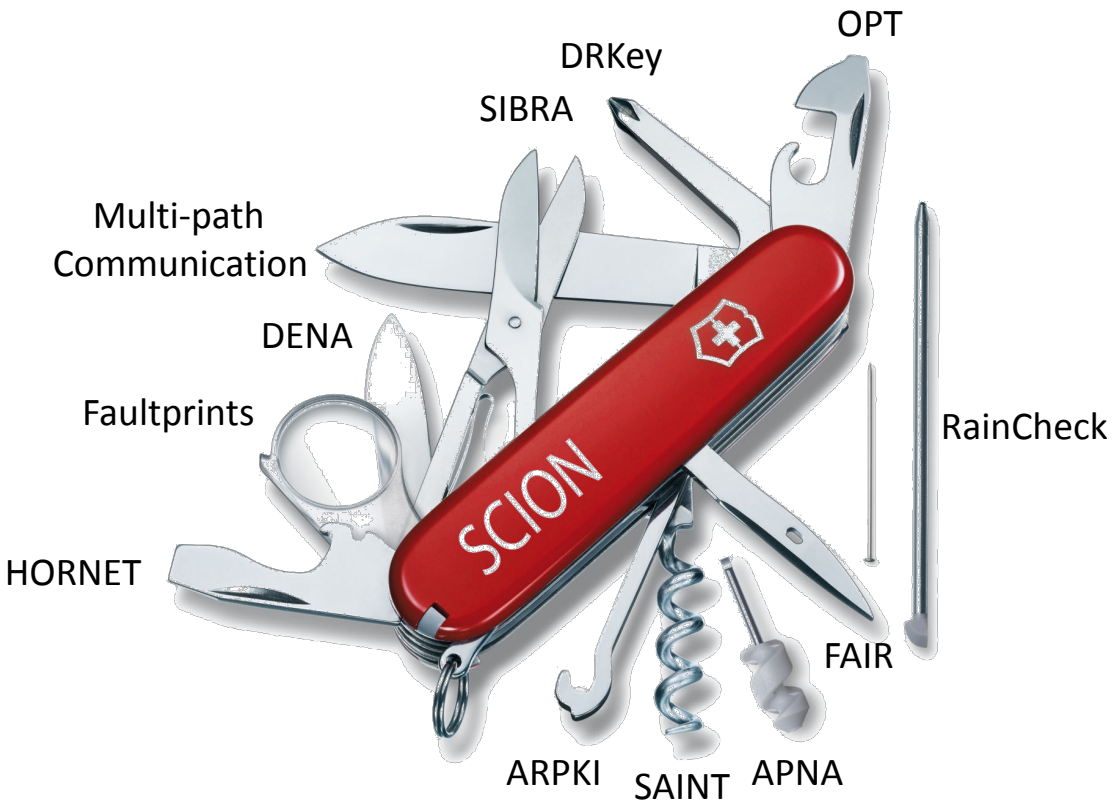

ETH zürich

SCION

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**.

Outline: Remainder of Talk

- Current implementation status
- Demos
  - Efficient forwarding
  - Multi-path communication
  - Browser-level path control
- Use cases

SCION Implementation Status

- 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 ~65 person-years of effort invested
- Growing testbed
- ISP Deployment: SWITCH, Swisscom, KDDI

ETH zürich

Demo: High-Speed Router

- Standard PC with dual Intel Xeon E5-2680 processors (~$500)
  - 8 cores per processor

- Intel 82599EB X520-DA2 NIC (2x 10Gbps) (~$600)

- Spirent SPTN4U-220 traffic generator

Multipath Communication

- SCION provides end-to-end paths to clients
- Using multiple paths can provide many benefits:
  - **Reliability** — avoid a single point of failure
  - **Bandwidth** — use more total bandwidth (subject to fairness constraints)
  - **Cost** — use cheaper links
  - **Latency** — use paths with lower propagation delays
- Different strategies are possible based on applications, path characteristics, and topology

Use Cases

- Highly efficient and available multi-path communication
- VPN link over SCION
- Inter-domain bandwidth guarantees via SIBRA and DILLs (Dynamic Inter-domain Leased Lines)
- Available and DDoS-resilient communication among banks in Switzerland
- Trustworthy network through verifiable router code

Highly Available and DDoS Resilient Communication among Banks

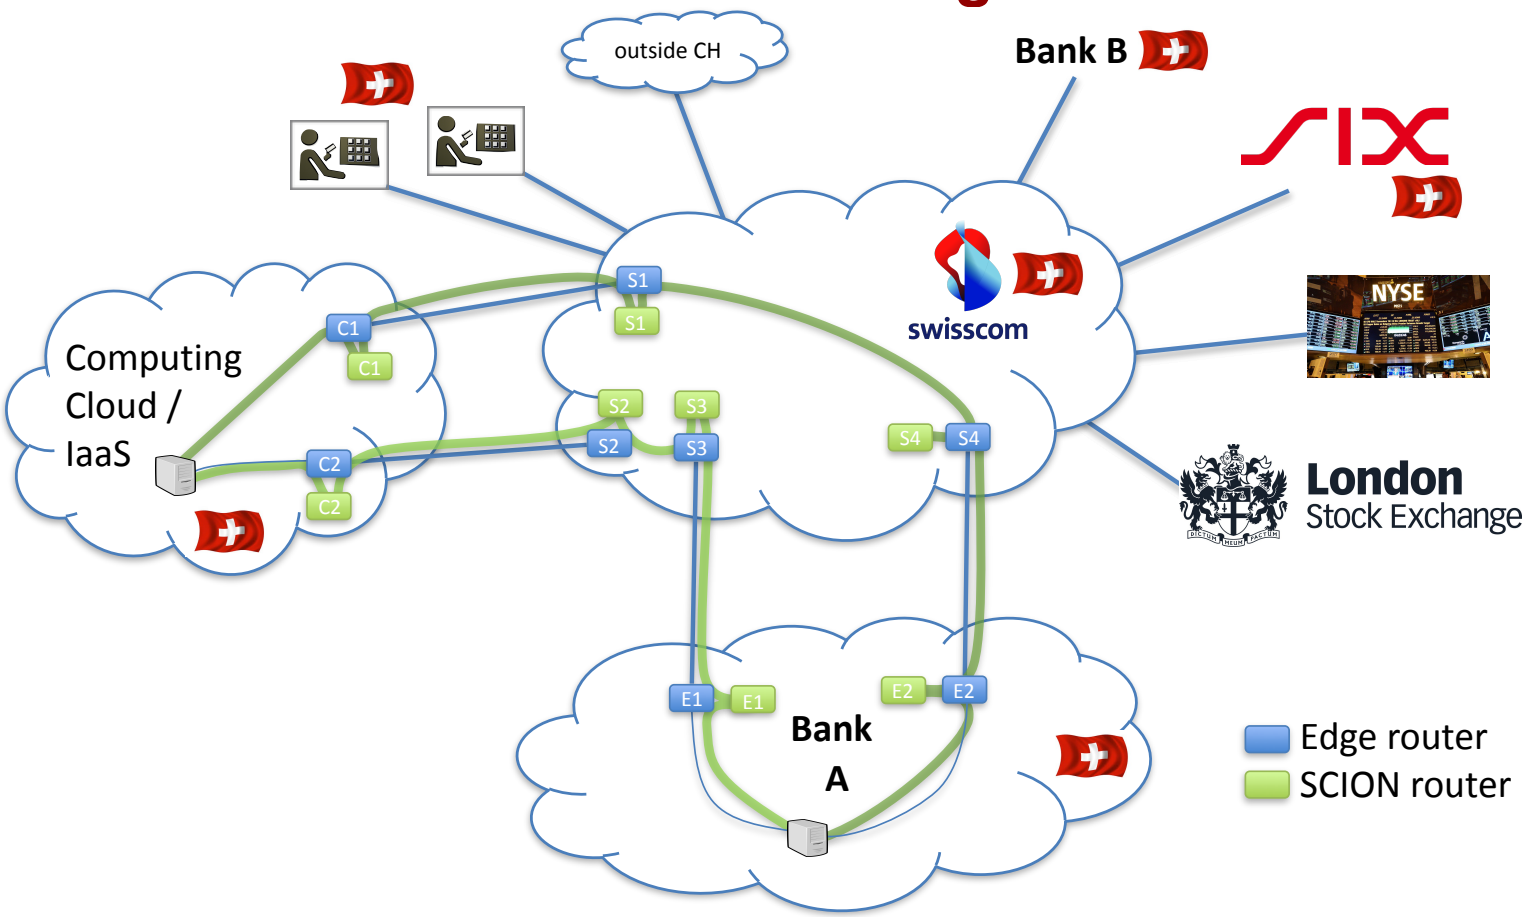

ETH zürich

SCION

Summary

- Network architecture re-design resolves 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 enables new applications and services
  - Guaranteed bandwidth, DDoS resilience
  - No Internet “Kill Switch”
  - Communication transparency
  - Verified router software
- ISPs and corporations now engaging in pilot deployment
- More information: http://www.scion-architecture.net

ETH zürich SCION