



SCiON

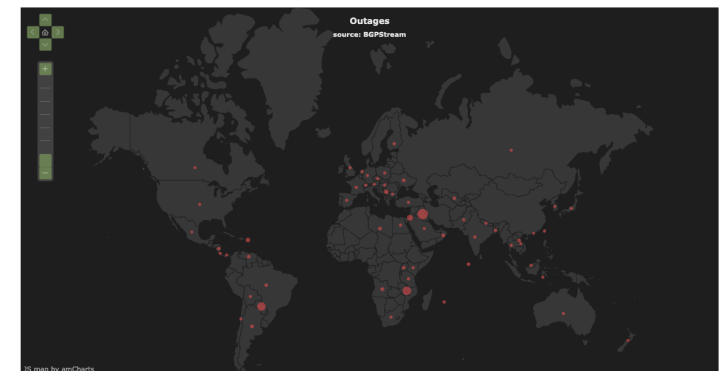
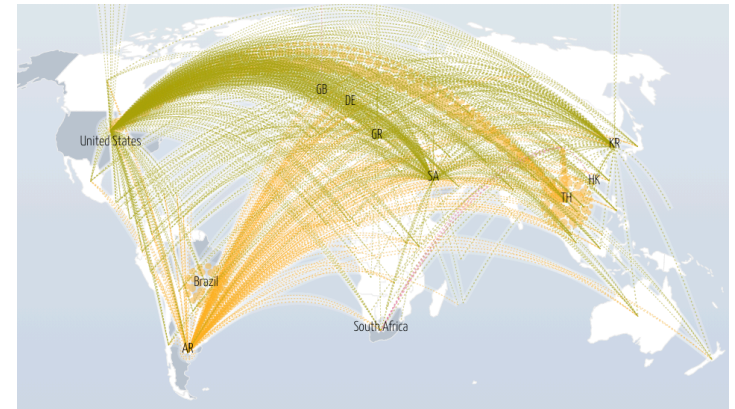
SCALABILITY, CONTROL, AND ISOLATION
ON NEXT-GENERATION NETWORKS

Experiencing a new Internet Architecture

Adrian Perrig, Network Security Group

The Internet is on Fire!

- Lack of sovereignty
- Frequent outages
 - <https://downdetector.com>
- Constant DDoS attacks
 - <https://www.digitalattackmap.com>
- Frequent routing attacks
 - <https://bgpstream.com>
- Lack of communication guarantees
- Expensive maintenance



Event type	Country	ASN	Start time (UTC)	End time (UTC)	More info
Possible Hijack		Expected Origin AS: ZOHO-EU, NL (AS 205111) Detected Origin AS: LVL-3549, US (AS 3549)	2020-10-06 01:01:28		More detail
Possible Hijack		Expected Origin AS: ZOHO-EU, NL (AS 205111) Detected Origin AS: LVL-3549, US (AS 3549)	2020-10-06 01:01:28		More detail
Outage		SWIFTNETBROADBAND-AS SWIFTNET BROADBAND PRIVATE LIMITED, IN (AS 133713)	2020-10-05 22:18:00	2020-10-05 22:22:00	More detail
Outage		U-LAN-AS, RU (AS 48128)	2020-10-05 21:24:00		More detail
Outage		TPODLASIE, PL (AS 39375)	2020-10-05 20:00:00	2020-10-05 20:52:00	More detail

Inspirations for a New Beginning

- Many exciting next-generation Internet projects over the past 25 years
- General Future Internet Architectures (FIA)
 - XIA: enhance flexibility to accommodate future needs
 - MobilityFirst: empower rapid mobility
 - Nebula (ICING, SERVAL): support cloud computing
 - NIMROD: improved scale and flexibility
 - NewArch (FARA, NIRA, XCP)
 - RINA: clean API abstractions simplify architecture
- Content-centric FIAs: NDN, CCNx, PSIRP, SAIL / NETINF
- Routing security: BGPSEC, S-BGP, soBGP, psBGP, SPV, PGBGP, H-NPBR
- Path control: MIRO, Deflection, Path splicing, Pathlet, I3
- Inter-domain routing proposals: ChoiceNet, HLP, HAIR, RBF, AIP, POMO, ANA, ...
- Intra-domain / datacenter protocols: SDN, HALO, ...



Why attempt redesigning Internet Architecture?

- We started our expedition asking the question:
How secure can a global Internet be?
 - Answer: global communication guarantees can be achieved as long as a path of benign domain exists
- During our journey we discovered that path-aware networking and multi-path communication are powerful concepts that can provide higher efficiency than a single-path Internet
 - Enables path optimization depending on application needs
 - Simultaneous use of several paths unlocks additional bandwidth
- Explore new networking concepts without the constraints imposed by current infrastructure!


Discoveries on our Journey

- During our journey, we have encountered many interesting discoveries
- Several discoveries suggest new approaches for inter-domain networking

The real voyage of discovery consists not in seeking new landscapes, but in having new eyes. Marcel Proust



SCION Ambition: A Global Next-Generation Public Internet

- 
- A globe is shown with a complex network of white lines and dots overlaid on it, representing a global communication network. The globe is set against a dark blue background with streaks of light. A yellow sticky note is pinned to the right side of the globe with two red pushpins. The sticky note contains a bulleted list of three items.
- High security and efficiency
 - Path-aware networking with multi-path communication
 - Global communication guarantees


SCION Architecture Principles

- Stateless packet forwarding (no inconsistent forwarding state)
- “Instant convergence” routing
- Path-aware networking
- Multi-path communication
- High security through design and formal verification
- Sovereignty and transparency for trust roots



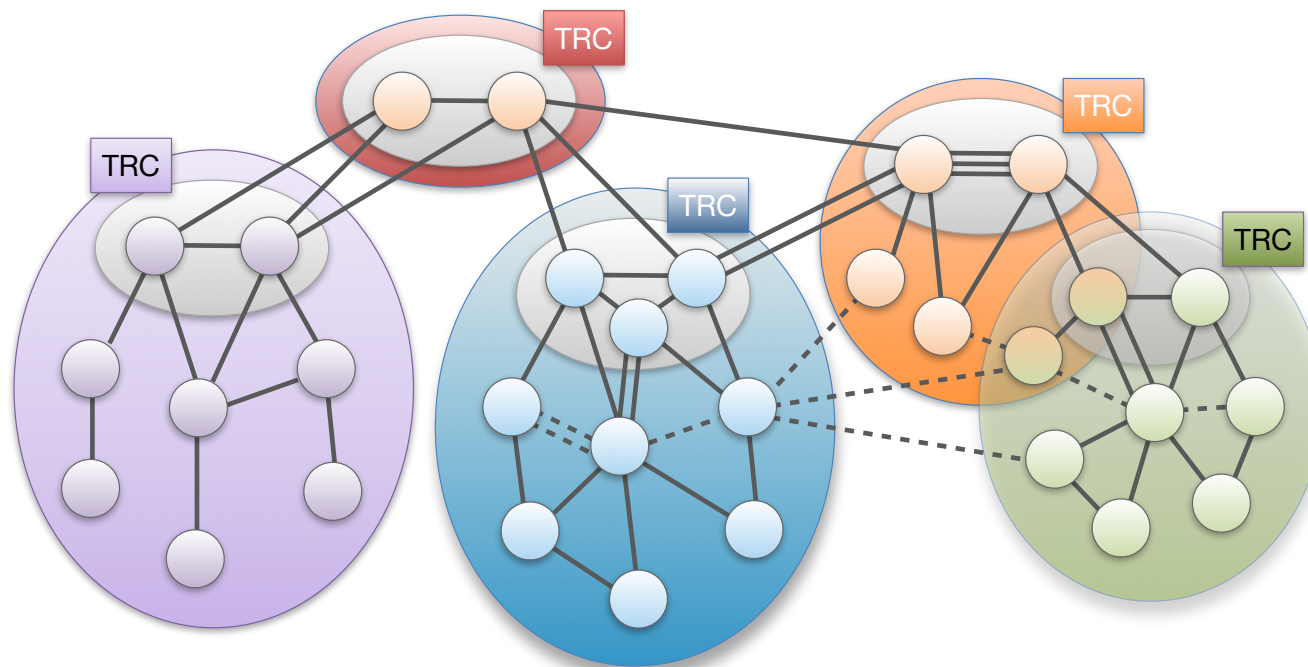
Insight: Formal Security Verification Necessary

- To achieve strong assurance for a large-scale distributed system, formal security verification is necessary
- Performing formal verification from the beginning avoids “difficult-to-verify” components
 - Many design aspects of SCION facilitate formal verification
- Collaboration with David Basin’s and Peter Müller’s teams in the VerifiedSCION project



Approach for Scalability: Isolation Domain (ISD)

- Isolation Domain (ISD): grouping of Autonomous Systems (AS)
- ISD core: ASes that manage the ISD and provide global connectivity
- Core AS: AS that is part of ISD core



SCION Overview in One Slide




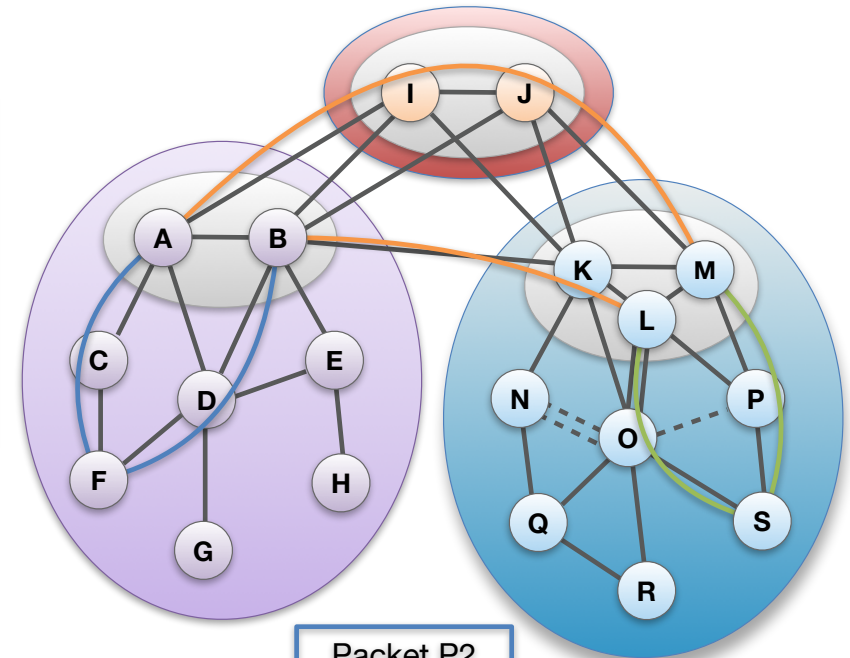

Path-based Network Architecture

Control Plane - Routing

- ❖ **Constructs** and **Disseminates** Path Segments

Data Plane - Packet forwarding

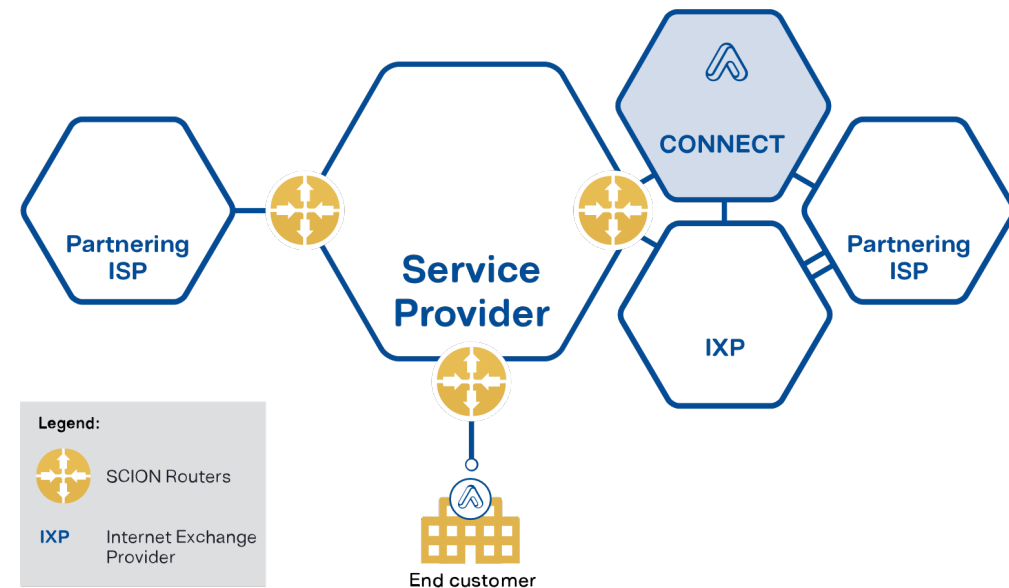
- ❖ **Combine** Path Segments to Path
- ❖ Packets contain Path
- ❖ Routers forward packets based on Path
 - ▶ Simple routers, stateless operation



SCION

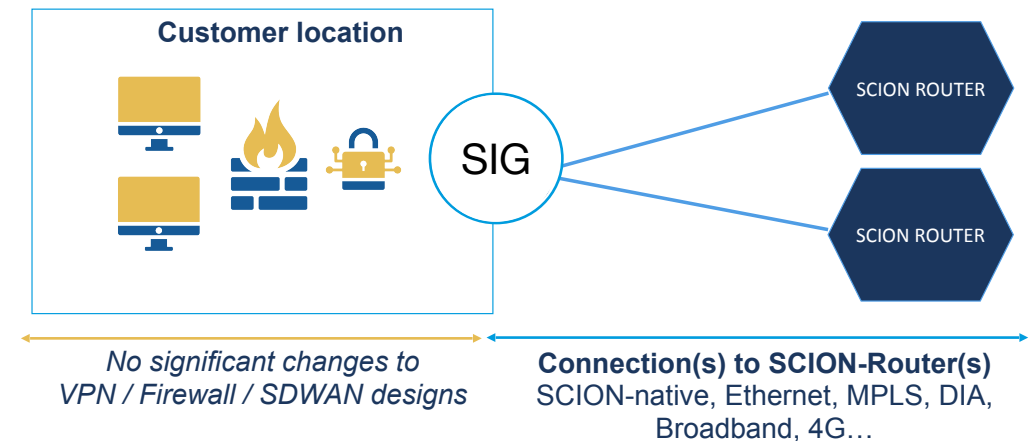
How to Deploy SCION: ISP

- CORE Routers are set up at the borders of an ISP
 - to peer with other SCION-enabled networks
 - to collect customer accesses
- No change to the internal network infrastructure of an ISP needed!




How to Deploy SCION: End Domain

- SCION IP Gateway (SIG) enables seamless integration of SCION capabilities in end-domain networks
- No upgrades of end hosts or applications needed




Insight: Incremental Deployment Possible

- Incremental deployment of a new Internet architecture is possible, operating side-by-side with BGP
- For ISPs, new architecture can be deployed with minimal effort
- For end domains, SCION-IP Gateway (SIG) offers immediate benefits without updating any end hosts
- Important: no reliance on BGP for inter-domain operation (“BGP-free”)
 - Overlay / insecure underlay should be avoided not to inherit vulnerabilities
- Re-use of intra-domain network architecture for local communication



SCIONLab


- Global SCION research testbed: <https://www.scionlab.org>
- Collaboration with David Hausheer's team at University of Magdeburg
- Open to everyone: create and connect your own AS within minutes
- ISPs: Swisscom, SWITCH, KDDI, GEANT, DFN
- Deployed 35+ permanent ASes worldwide, 600+ user ASes
 - Contact us to become an infrastructure AS, we can provide HW
- Kwon et al., “SCIONLab: A Next-Generation Internet Testbed”, ICNP 2020



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)
- 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
- Colibri inter-domain resource allocation system
- DDoS defense research using in-network defense mechanisms
- Next-generation routing architecture policy definitions


SCION Production Network

- Led by Anapaya Systems  ANAPAYA
- **BGP-free global communication**
 - Fault independent from BGP protocol
- Deployment with international ISPs
 - Goal: First **global public secure** communication network
- Construction of SCION network backbone at select locations to bootstrap adoption
- Current deployment
 - ISPs: Swisscom, Sunrise, SWITCH, + others joining soon
 - IXPs: SwissIX offers SCION peering, + others joining soon
 - Bank deployment: 4 major Swiss banks, some in production use




Global Availability of Native SCION Connectivity

- Native SCION (BGP-free) connectivity: no reliance / dependency on BGP communication
- SCION deploying ISP's networks are reaching global data centers and IXPs, offering native SCION connectivity
- Anapaya Connect: native SCION connectivity to 100+ data centers in 10+ countries
 - Further expansion next year:



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 Summary

- SCION: Next-generation Internet **you can use today!**
- **High-performance**
 - Path-aware network enables application-specific optimizations to provide **enhanced efficiency**
 - Multi-path communication enables simultaneous use of multiple paths, increasing available bandwidth
- **Secure, high assurance, high availability**
 - Per-packet authentication verification possible on routers
 - Formal verification of protocols and code
 - Immune against routing attacks, e.g., BGP prefix hijacking

Interesting Encounters on our Expedition

- Security
 - Global communication guarantees are possible
 - High-speed crypto enables line-rate processing
- Networking
 - Multi-path routing is a necessity, not a luxury
 - Improved scalability over BGP
 - Global QoS is viable



Global Communication Guarantees in the Presence of Adversaries


- Goal: If (routing policy compliant) path of benign ASes exists (with operational infrastructure), a sender can find, use, and achieve minimum bandwidth guarantees on that path
- Challenges
 - Network routing instabilities, misconfigurations, etc.
 - DoS attacks at various levels (control plane, data plane, end host)

Observation: Stable Forwarding + Multi-path Necessary

- Single-path forwarding cannot achieve strong availability guarantees
 - During routing protocol convergence, no path may be available
 - Equipment failure on path will result in unavailability until routing protocol updates and forwarding tables are adjusted
 - If forwarding path experiences high packet loss, then path is not usable for practical applications
- Approaches
 - **Stable forwarding**: packet-carried forwarding state protects forwarding from routing instabilities
 - **Multi-path** ensures presence of several paths, so as long as a single path works, end-to-end connectivity is assured

Bottleneck Routing Disrupts Availability

- Routing protocol switches route traversing a link with limited capacity (= bottleneck link)
- Bottleneck link traversal results in high packet loss
- Applications cannot operate and lose connectivity
- Since connectivity exists, often manual intervention needed to switch back to alternate path, outage typically persists for 30+ minutes
- Frequent reason for outage, caused by misconfiguration or attack



Cloudflare DNS goes down, taking a large piece of the internet with it

Devin Coldewey @techcrunch / 11:50 pm CEST • July 17, 2020

 Comment



For two hours, a large chunk of European mobile traffic was rerouted through China

It was China Telecom, again. The same ISP accused last year of "hijacking the vital internet backbone of western countries."



By Catalin Cimpanu for Zero Day | June 7, 2019 – 19:41 GMT (20:41 BST) | Topic: Security




Announcement of Failed Routes

- In some cases, networks continue to announce routes that failed
- Example: August 30 CenturyLink/Level(3) Outage
<https://blog.cloudflare.com/analysis-of-todays-centurylink-level-3-outage>
“CenturyLink/Level(3)’s network was not honoring route withdrawals and continued to advertise routes to networks like Cloudflare’s even after they’d been withdrawn”

Insight: Secure Routing Insufficient

- Secure routing protocol cannot prevent outages caused by bottleneck link or continuing announcement of failed or congested routes, as announcement is often legitimate




Global Communication Guarantees in the Presence of Adversaries

- Goal: If (routing policy compliant) path of benign ASes exists (with operational infrastructure), a sender can find, use, and achieve minimum bandwidth guarantees on that path
- Challenges
 - Network routing instabilities, misconfigurations, etc.
 - DoS attacks at various levels (control plane, data plane, end host)

Availability in a public Internet is threatened by different types of DoS attacks

Link-flooding attacks	Attacker floods network links with excessive amount of traffic
	Can target access links (last mile) or core links in the network
	Often executed using botnets and/or amplification techniques
End-system attacks	Attacker exhausts computational or memory resources of victim
	Often possible due to other defense mechanisms such as firewalls
	Examples: state exhaustion , signature flooding
Control-plane attacks	Attacker disrupts important control-plane mechanisms or access to services
	Services are essential for a functioning network
	Examples in SCION: beacon server, path server, certificate server

SCION is an Internet architecture with both *strong security properties* and *high availability*



High-Speed Packet Processing

- Current high-speed Internet links: 400Gbit/s (Gbps)
- Arrival rate for 64-byte packets: one packet every 1.3 ns
- High-speed asymmetric signature implementation: Ed25519
SUPERCOP REF10: $\sim 100\mu\text{s}$ per signature
- AES-NI instruction only requires 30 cycles: $\sim 10\text{ns}$
- Memory lookup from DRAM requires ~ 200 cycles: $\sim 70\text{ns}$
- Symmetric crypto enables high-speed processing through parallel processing and pipelining

DRKey & Control-Plane PKI

- SCION offers a global framework for authentication and key establishment for secure network operations
- Control-plane PKI
 - Sovereign operation thanks to ISD concept
 - Every AS has a public-key certificate, enabling AS authentication
- DRKey
 - High-speed key establishment (within ~20 ns), enabling powerful DDoS defense mechanisms
- PISKES: Pragmatic Internet-Scale Key-Establishment System, Rothenberger et al., ACM Asia Conference on Computer and Communications Security (ASIACCS) 2020

Dynamically Recreatable Key (DRKey)


- *Idea*: use a per-AS secret value to derive keys with an efficient Pseudo-Random Function (PRF)
- Example: AS X creates a key for AS Y using secret value SV_X
 - $K_{X \rightarrow Y} = \text{PRF}_{SV_X}(\text{"Y"})$
 - Intel AES-NI instructions enable PRF computation within 30 cycles, or 70 cycles for CMAC
Key computation is ~7 times faster than DRAM key lookup!
- Any entity in AS X knowing secret value SV_X can derive $K_{X \rightarrow *}$

EPIC: Every Packet Is Checked

- Goals
 - Per-packet source authentication by every router and destination
 - Per-packet-unique hop fields
 - Path validation by destination
- Assumption: global time synchronization ($\pm 100\text{ms}$)
- Attacks prevented
 - Malicious router replays packets or increases packet size
 - Hop field MAC is brute forced and destination attacked until expiration time
- EPIC: Every Packet Is Checked in the Data Plane of a Path-Aware Internet, Legner et al., USENIX Security Symposium 2020

Insight: Cryptographic Processing at Line Rate Possible

- Symmetric-key cryptographic operations are possible within nanoseconds, thus enabling line-rate processing
- With hardware implementation, computing an AES block cipher can be accomplished within a few nanoseconds
- DRKey + EPIC systems enable per-packet source authentication in software ~ 100 ns
- This enables new approaches for network security




Importance of Path Awareness & Multi-path

- Generally, two paths exist between Europe and Southeast Asia
 - **High latency, high bandwidth:** Western route through US, ~450ms RTT
 - **Low latency, low bandwidth:** Eastern route through Suez canal, ~250ms RTT
- BGP is a “money routing protocol”, traffic follows cheapest path, typically highest bandwidth path
- Depending on application, either path is preferred
- With SCION, both paths can be offered!



Insight: Multi-Path is a Necessity for High Availability and Performance

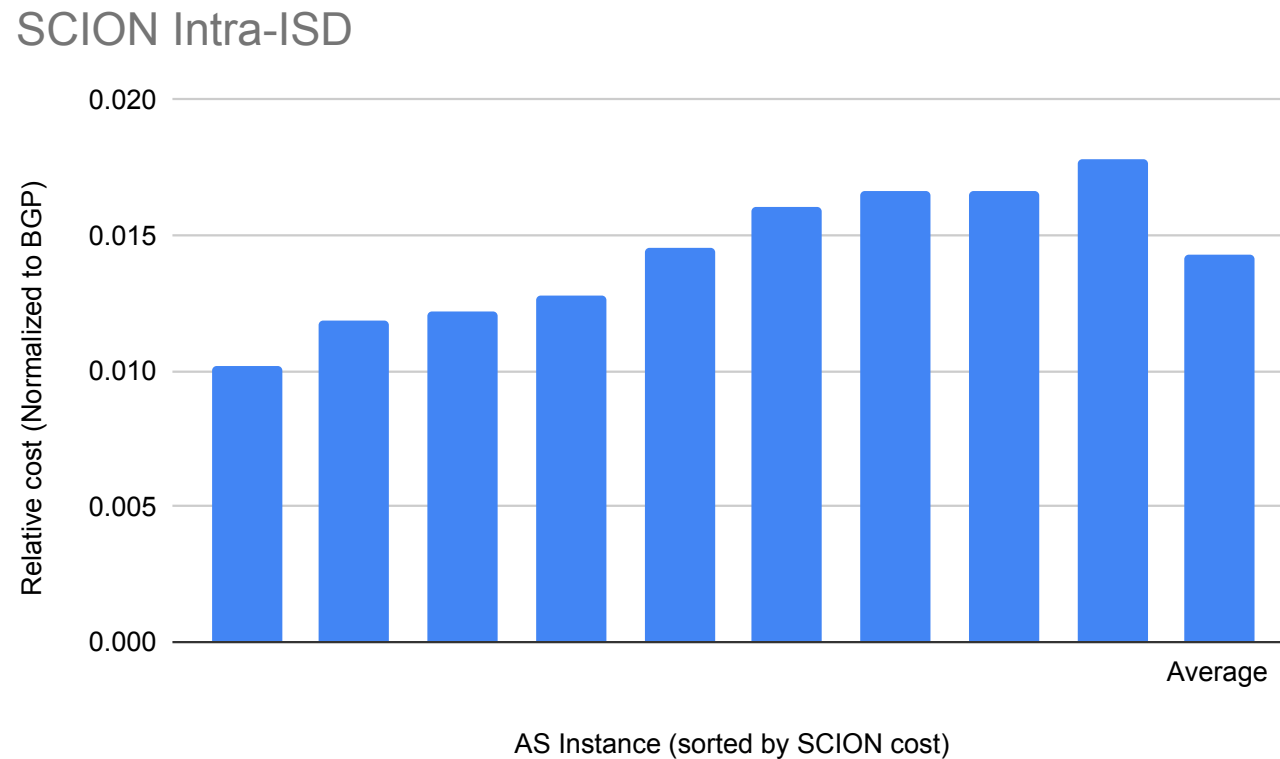
- Inter-domain multi-path is not a luxury, but a necessity to achieve high availability
- Rapid failover without routing system convergence
 - Routing bottlenecks can be avoided
- Enable higher network capacity
 - No more passive links for redundancy, all links can be active
 - Simultaneous use of several links
- Enables higher communication efficiency
 - Latency- vs. bandwidth optimal paths can be chosen
- Helps defend against DoS attacks, as adversary needs to congest all links
- QoS needs multi-path, as several alternatives need to be available to attempt resource reservations



Multi-Path Routing Approaches


- For a powerful multi-path system, we need a rich set of path choices: ideally dozens of paths if possible
- Problem: most prior multi-path routing algorithms are based on BGP, offering only 2-3 different path choices
 - Overhead increases linearly in the # of paths: hampering scalability
 - Notable exceptions: Pathlets, NIRA, HAIR
- The path segment combination of SCION provide a rich set of path choices
 - Extensible architecture: additional path segment generation algorithms can be added, and path server infrastructure can be used for dissemination

Scalability of SCION Intra-ISD Beaconsing




Scalability of SCION Core Beaconsing

SCION Core

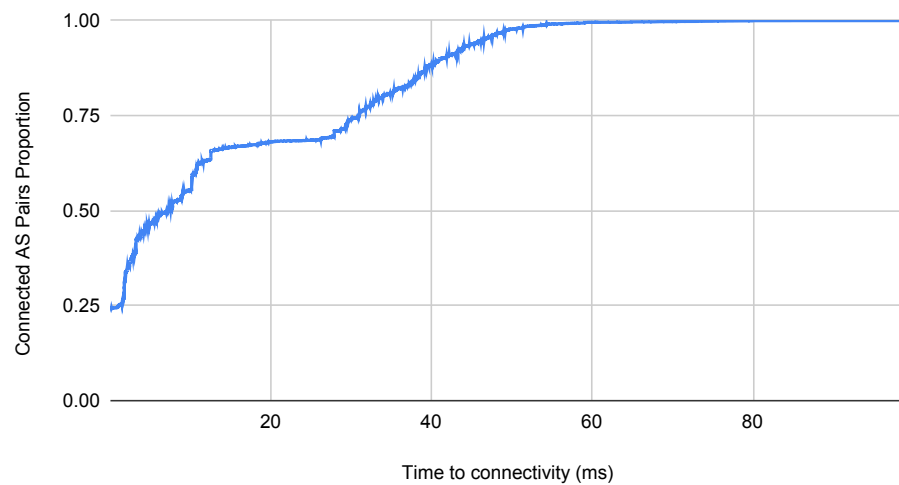


SCION Core, BGPsec and BGP

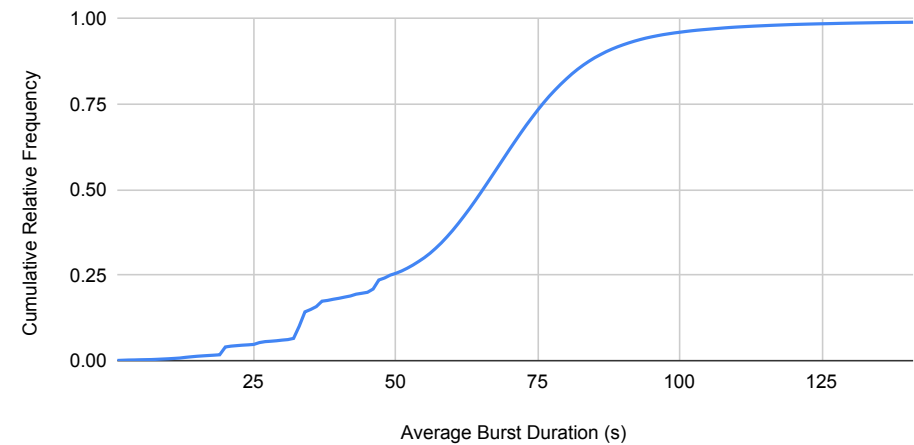


Time to Connectivity

Time to connectivity in SCION (Min Latency)




Distribution (99 %) of "average length of update bursts per prefix" during May 2020



Insight: SCION Provides Scalable Multi-Path Routing

- ISD decomposition offers scalability for segment exploration; and segment combination offers large number of path choices
- Overhead of beaconing is 10-100x lower than BGP, even though dozens of disjoint paths are being created
- Time-to-connectivity is approximately two orders of magnitude faster than BGP




Volumetric DDoS Attacks

- Attacker overloads network link to induce congestion
- Defense requires sophisticated approaches
 - EPIC dynamic hop field computation
 - COLIBRI global resource allocation and reservation


Coremelt Attack [Studer, Perrig, Esorics 2009]

- Adversary controls many bots distributed across the Internet
- Bots send traffic between each other, thus all traffic is desired by destination
 - Traffic is not sent to victim as in regular DDoS attacks
- Adversary can exhaust bandwidth on victim link
- Result: attack traffic exhausts bandwidth in per-flow fair sharing systems



Crossfire Attack [Kang, Lee, Gligor, IEEE S&P 2013]

- Adversary controls distributed bot army
- Observation: due to route optimization, few links are actually used to connect a target region to rest of Internet
- Adversary can contact selected servers to overload target links
- Result: disconnect target region from remainder of Internet

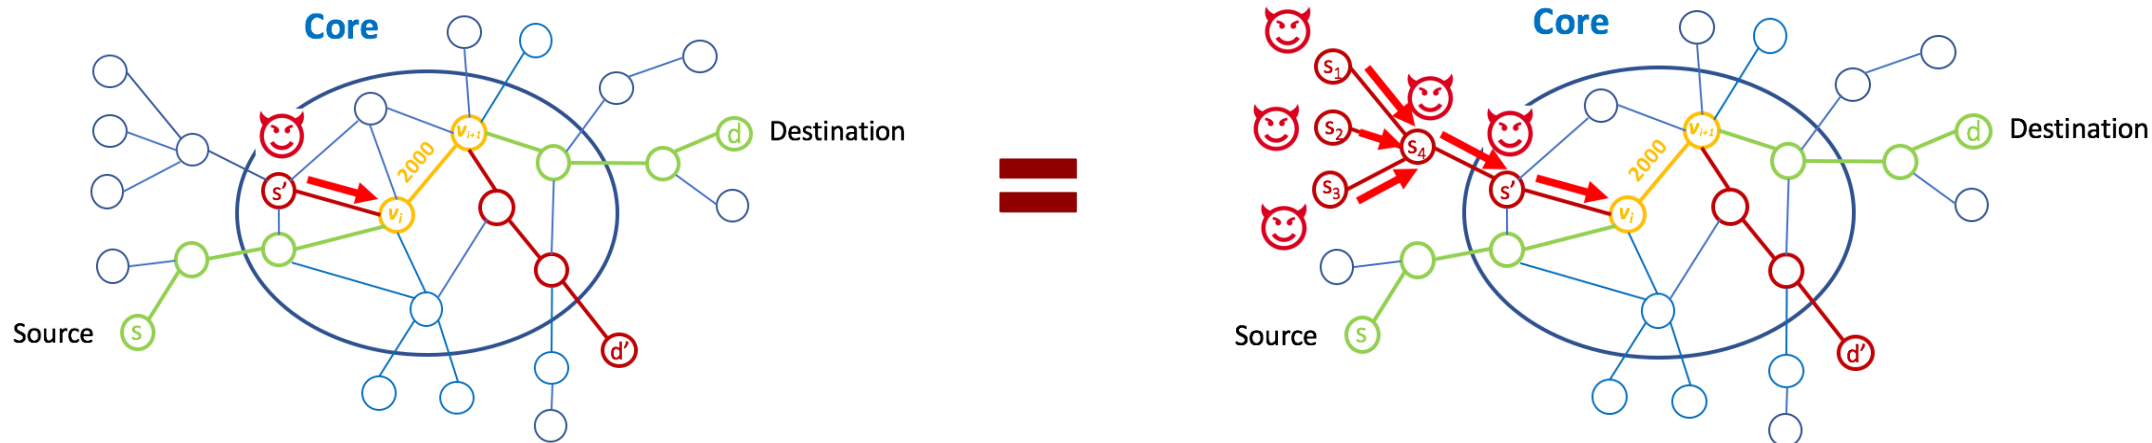


COLIBRI: Scalable Global QoS

- Thanks to several innovations, global QoS is now scalable and practical
- Stable paths ensure reservations are not affected by routing changes
- Multi-path enables searching for paths with sufficient bandwidth
- No per-flow state on routers, enabling scalability
 - DRKey enables high-speed per-packet source authentication
 - Efficient probabilistic large flow detection enable overuse detection
 - Per-flow stateful control-plane implemented on server infrastructure
- Per-neighbor fairness enables simple admission decision and configurations for ISPs


Admission Algorithm with Per-Neighbor Fairness

- Each AS defines neighbor-to-neighbor minimum bandwidth guarantees
- For any path, AS-to-AS minimum bandwidth guarantee can be computed, regardless of other demands
- Algorithm guarantees that no set of ASes can reserve a disproportionate amount of bandwidth through any link




Insight: Bandwidth Reservation Offers Many Advantages

- Explicit bandwidth admission simplifies transport layer
 - No need for sophisticated congestion control: simply use constant bitrate (CBR)
 - Reduce amount of acknowledgments due to very low loss rate
 - Fairness can be enforced at level of admissions
 - Possible reduction in energy utilization at end points
- Reserved but unused bandwidth can be used for best-effort traffic: no wasted bandwidth
- Fine-grained traffic engineering possible for ISPs
- Majority of traffic today is video: well suited for CBR traffic
 - Could simplify buffering and adaptive-bitrate algorithms



Expeditions Enable New Insights & Discoveries

- What started with the question “How secure can a global Internet be” has rewarded us with an exciting journey of insights and discoveries
- We hope to question engrained assumptions to counter Internet ossification
- Join the journey
 - <https://www.scionlab.org>
 - <https://www.scion-architecture.net>



SCION Team (2019)

