2STIC

Experimenting with the SCION Internet architecture

Caspar Schutijser, Ralph Koning (SIDN Labs)

Advanced Networking Guest Lecture, UvA/OS3

Feb 26, 2025

SIDN is the operator of the .nl TLD

- Objective: increase society's confidence in the Internet

- Provide secure and fault-tolerant registry services for .nl
  - Anycasted DNS services with DNSSEC support
  - Registration and domain protection services

- Increase the value of the Internet in the Netherlands and elsewhere
  - Enable safe and novel uses (SIDN Fonds, Yivi/IRMA)
  - Increase infrastructure security and trustworthiness (SIDN Labs)

- Not-for-profit private organization with a public role based in Arnhem

.nl = the Netherlands
17M inhabitants
6.2M domain names
3.8M DNSSEC-signed
4.0B DNS queries/day
8.6B NTP queries/day

SIDN Labs: team

- **Goal**: increase security of the Internet infrastructure, special focus on .nl and the Netherlands
- **Themes**: domain name security, infrastructure security, *emerging Internet technologies (long term)*
- Applied technical research: large-scale measurements, system design, prototyping and evaluation, standards
- Results publicly available to advance the Internet
- Bridge academia and operations/industry

2STiC program

Goal: put Dutch and European internet communities in a leading position in the field of secure, stable and transparent inter-network communication

The Internet

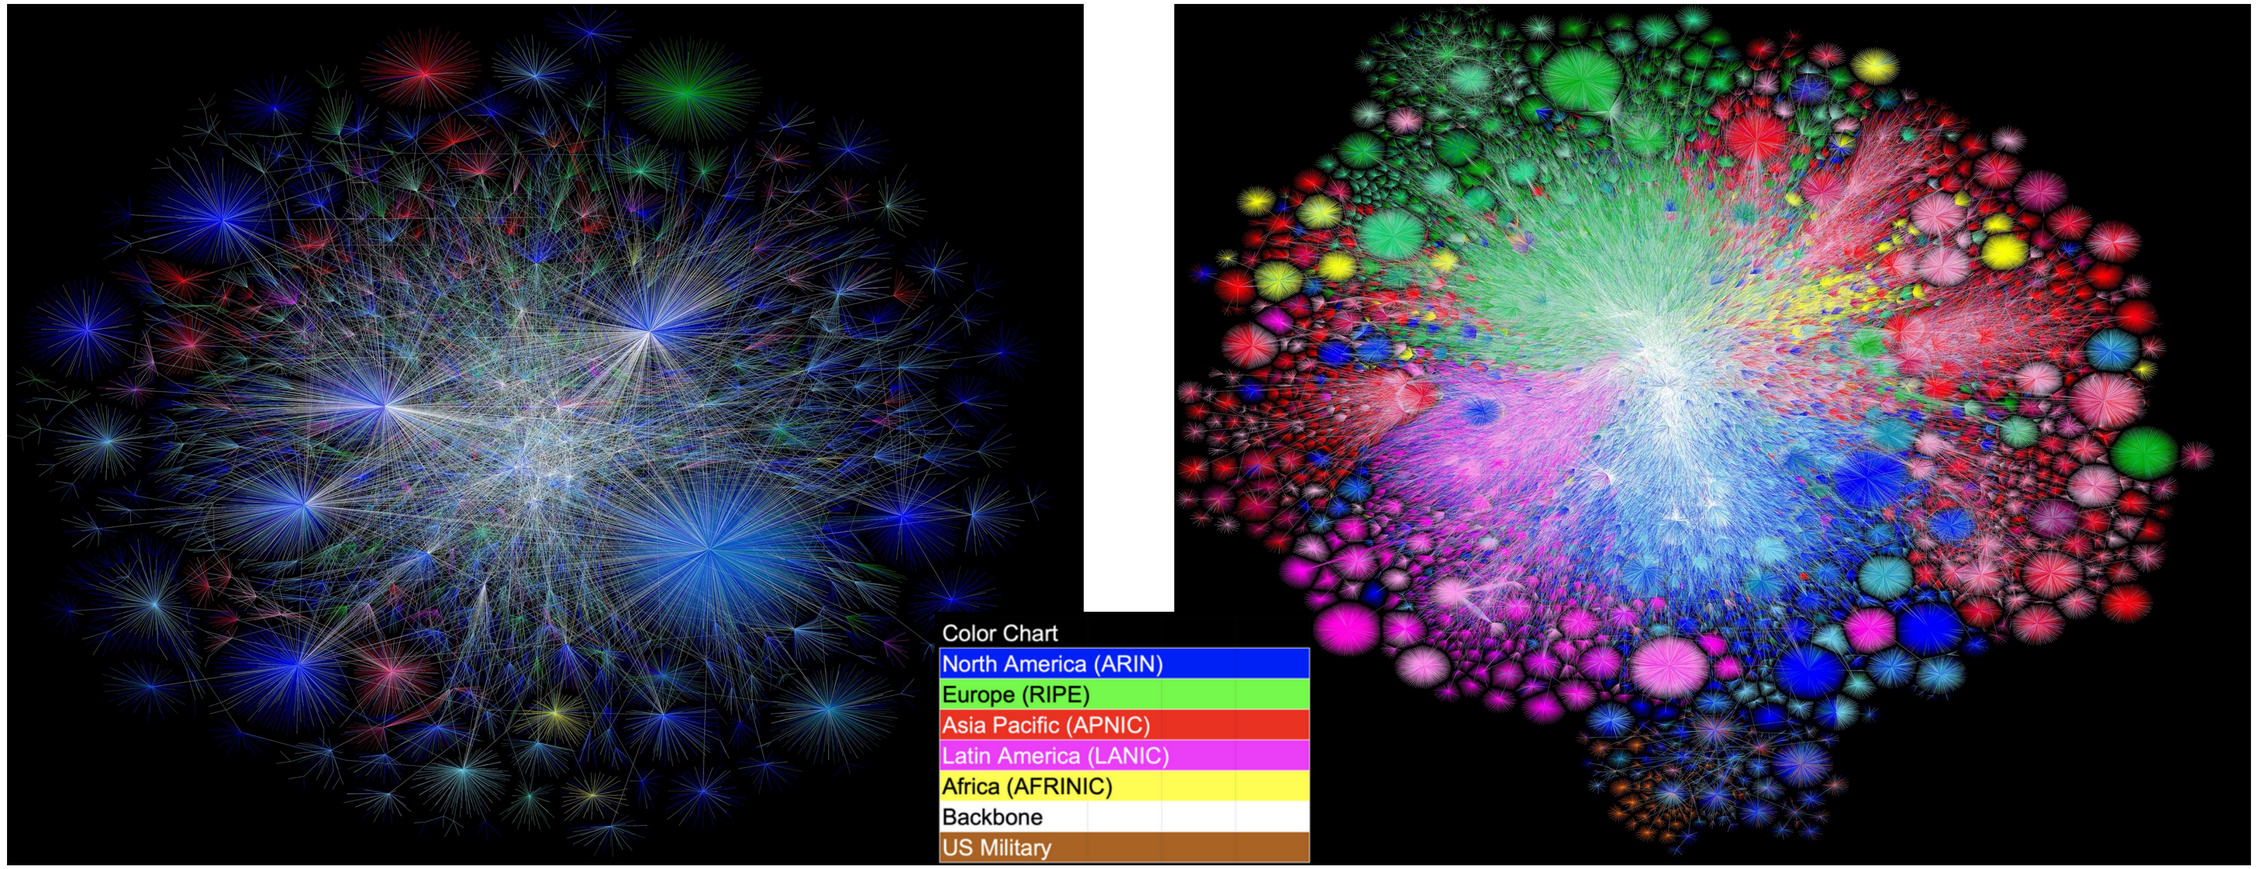

1997

2021

source: https://www.opte.org

Rate of change

Fast

Slow!

Fast

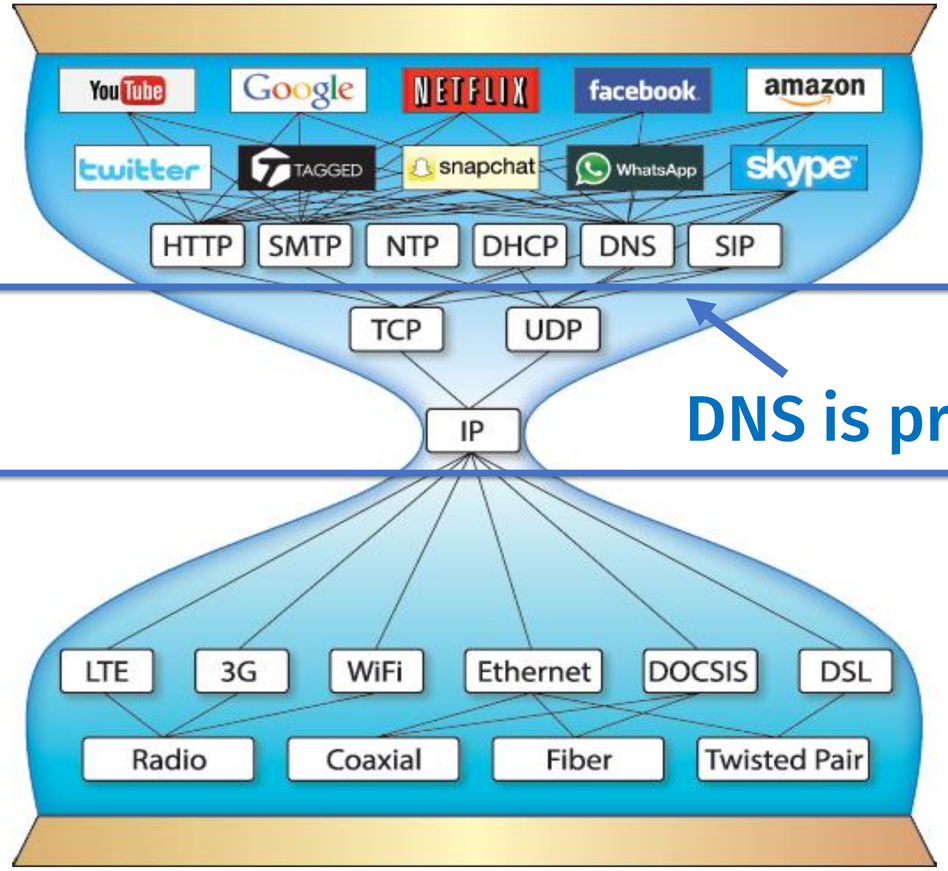

DNS is pretty slow too

New Requirements

- New applications have new **security**, **stability** and **transparency** requirements
  - More interaction with physical space (e.g., transport, smart grids, drones, remote surgery)

- To provide **trust** and **confidence** in communication we need a **responsible** internet
  - Control over routing and verification of operational behavior

SCION
NDN
RINA
ManyNets
XIA
MobilityFirst
Nebula
Service-centric networking
FII
B4
…

Some new inter-domain networked architectures

Opening up

- Adoption of new protocols in technologies was slow, but network devices are opening up.
- (Onie) Open Network Install Environment offers OS choice on network equipment.
- OpenFlow/SDN offer control plane programmability.
- P4 provides dataplane programmability.

Potentially promising clean slate architectures

- RINA
  - Everything is IPC (Inter-Process Communication)
  - WIP implementations: ProtoRINA, OpenIRATI
- NDN
  - Data centric
  - Stateful, lots of caching in the network
  - Implementation: named-data.net
- Extensible Internet
  - Introduces layer 3.5
  - No implementation yet
- **SCION**
  - **Path selection**
  - **Active community**
  - **Implementation: github.com/scionproto**

SCION

SCION

- Scalability, Control, and Isolation On Next-generation Networks
- New internet architecture
- Network Security Group, ETH Zurich
- Goal: improve security of inter-domain routing and isolation of compromise
- Scalability and security through Isolation Domains (ISDs)
  - Group of autonomous systems
  - E.g., per country or jurisdiction

SCION

- Security by design
  - Routes authenticated both in control and data plane
- Path-aware networking
  - Sender selects path
  - Enables, for example, geofencing
- Multi-path communication
  - Can be used, for example, for redundancy
- Existing application can still be used

Isolation domains

- Group of autonomous systems
  - E.g., per country or jurisdiction
- ISD core: ASes managing the ISD
- Core AS: AS part of the ISD core
- PKI organised per ISD
- Hierarchical control plane
  - Inter-ISD control plane
  - Intra-ISD control plane

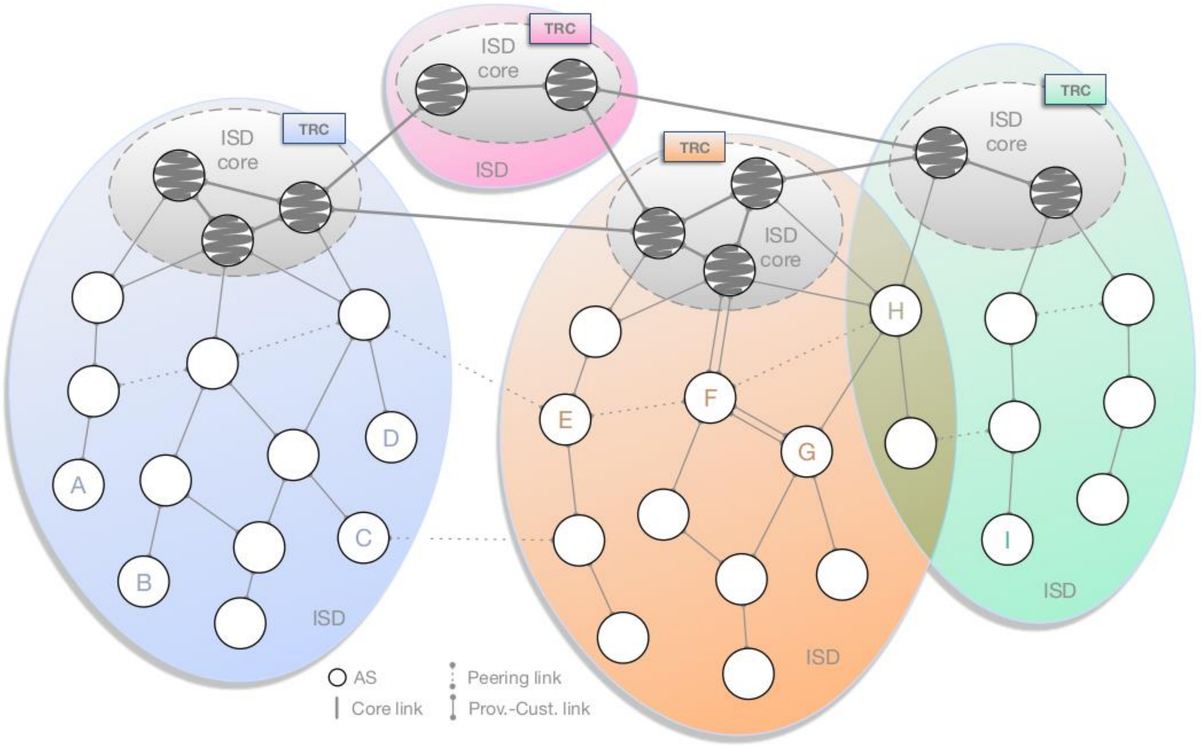

Source: The SCION Internet Architecture: An Internet Architecture for the 21st Century, Barrera et al., 2017

Routing: overview

- Control plane
  - Construct and disseminate path segments
- Data plane
  - Combine path segments to path
  - Packets contain path
  - Routers forward packets based on path (stateless)

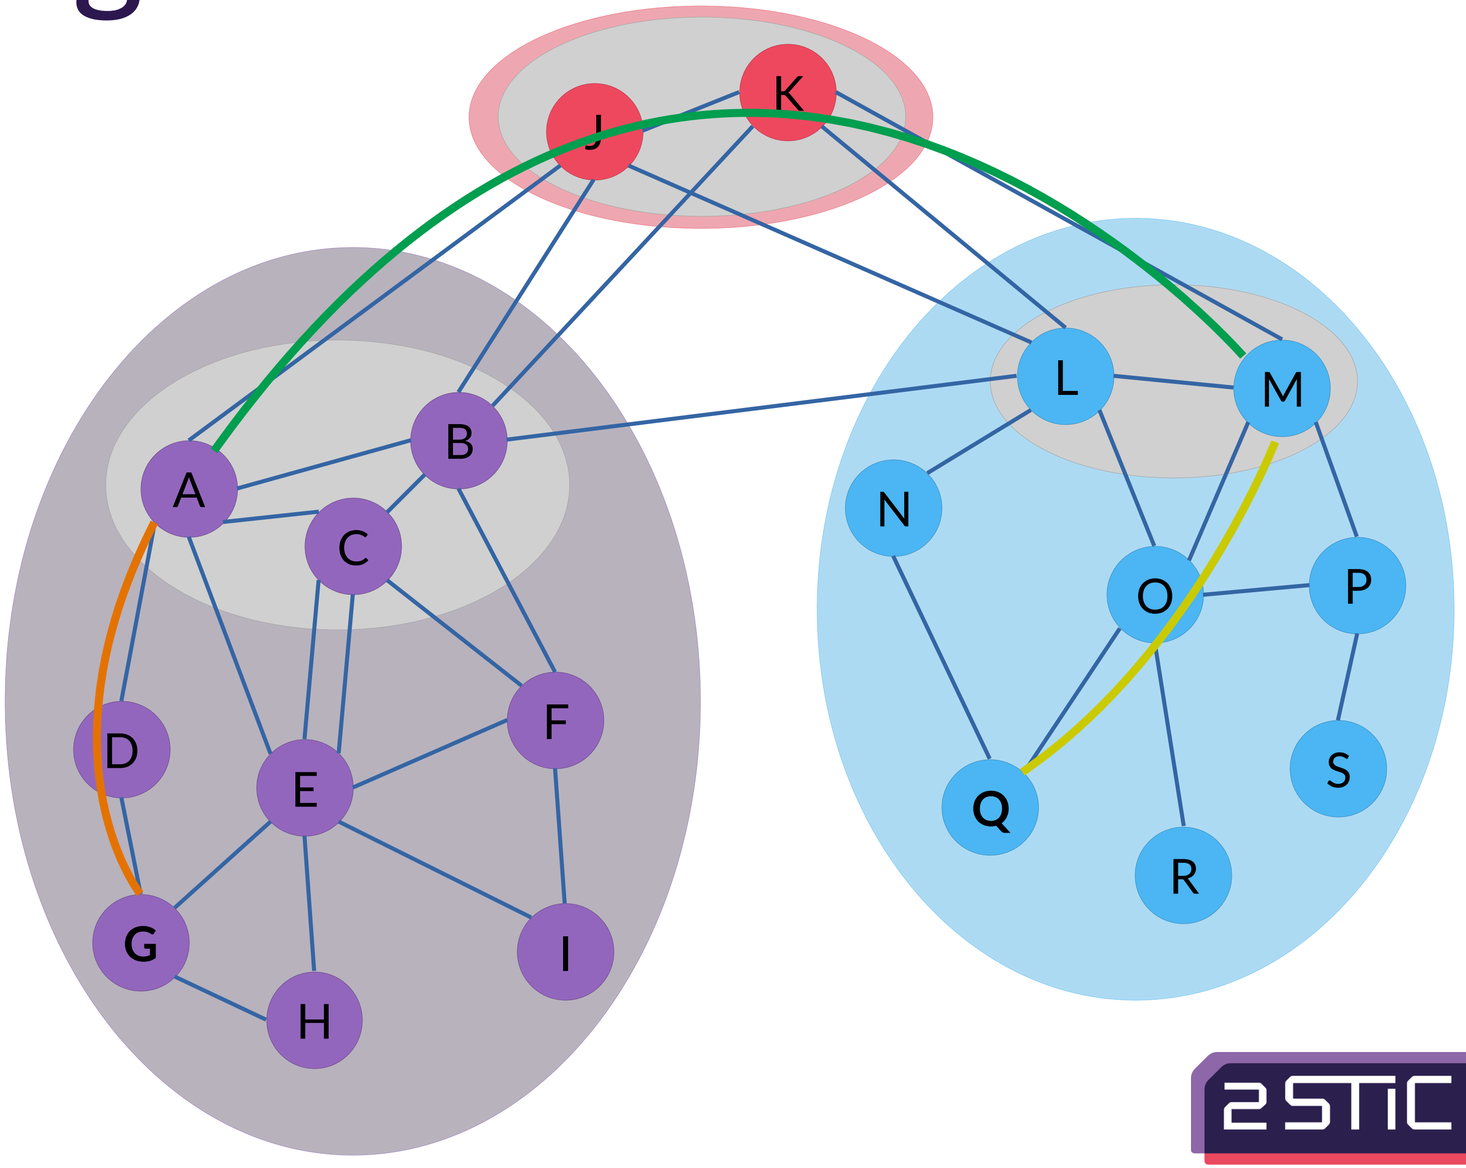

Data Plane : Segments

Routing

– Packets a contain path (PCFS)

– Routers forward packets based on path (stateless)

Path Construction

– A path consists of path segments.

– Valid segments are requested at a path server.

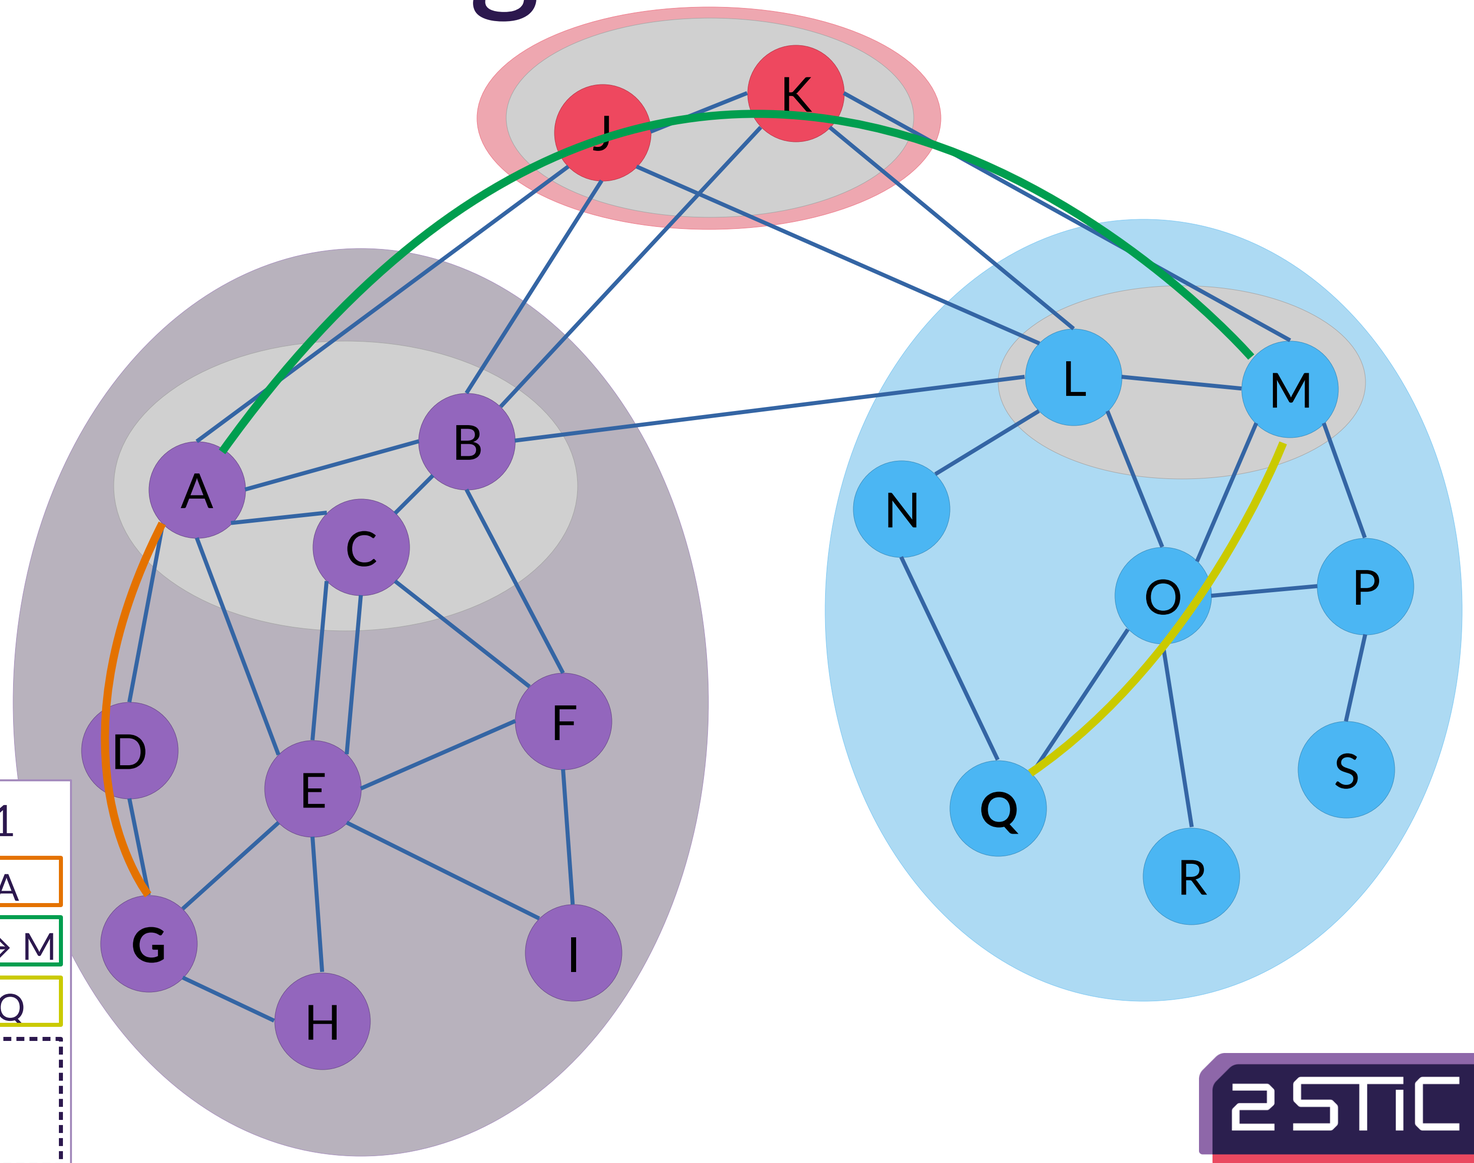

Data plane: path lookup

- Path construction performed by end hosts
- Request route to (ISD, AS) from local path server
- Local path server replies with
  - Up-path segments to local ISD core
  - Down-path segments in remote ISD from core to destination AS
  - Core-path segments needed to connect up-path and down-path segments
- End hosts pick and combine segments to determine path

Data plane: path combination

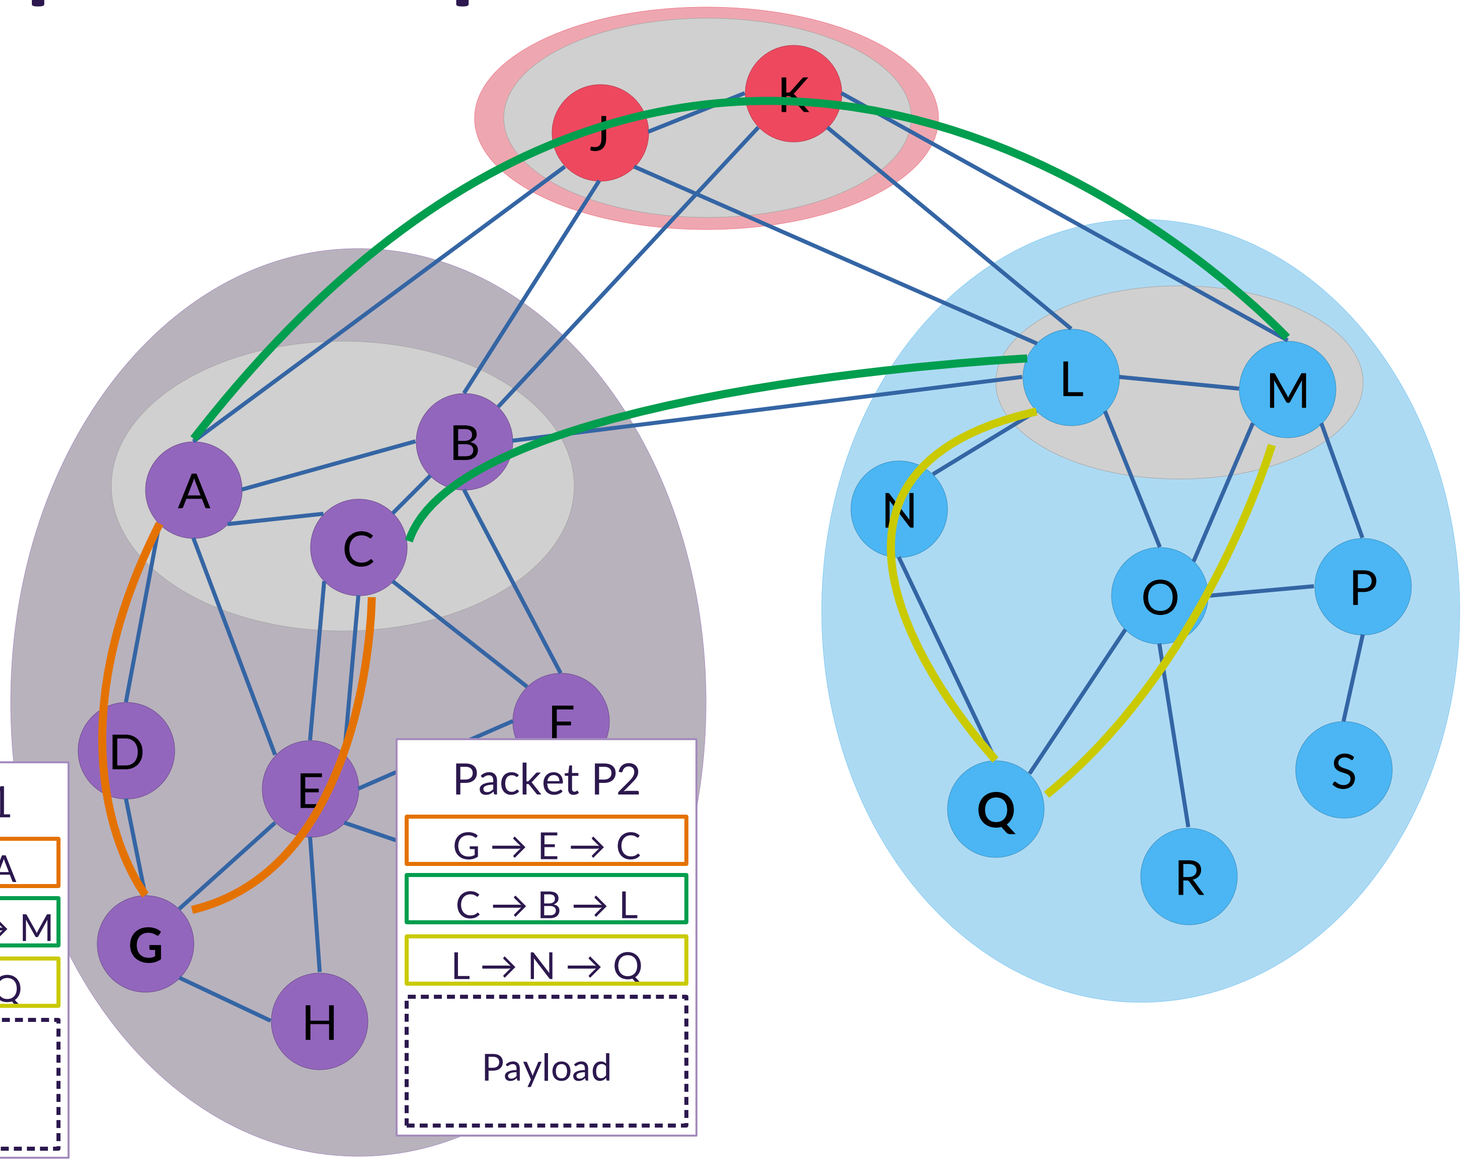

Control plane: path exploration

- Inter-ISD
  - Performed by core ASes
  - Flooding similar as with BGP
  - Less ASes involved (only core)
- Intra-ISD
  - Downstream multi-path flooding

Control plane: Intra-ISD path exploration

- Path Construction Beacons (PCBs) sent downstream using multi-path flooding
  - Initialised by core ASes
  - Extended and forwarded by receiving ASes
  - Add incoming and outgoing interface and optional peerings
- Eventually all nodes know how ISD core can be reached
- Path registration
  - Preferred down-segments (path from core to AS) with path server in the core
  - Preferred up-segments registered with local path server in AS

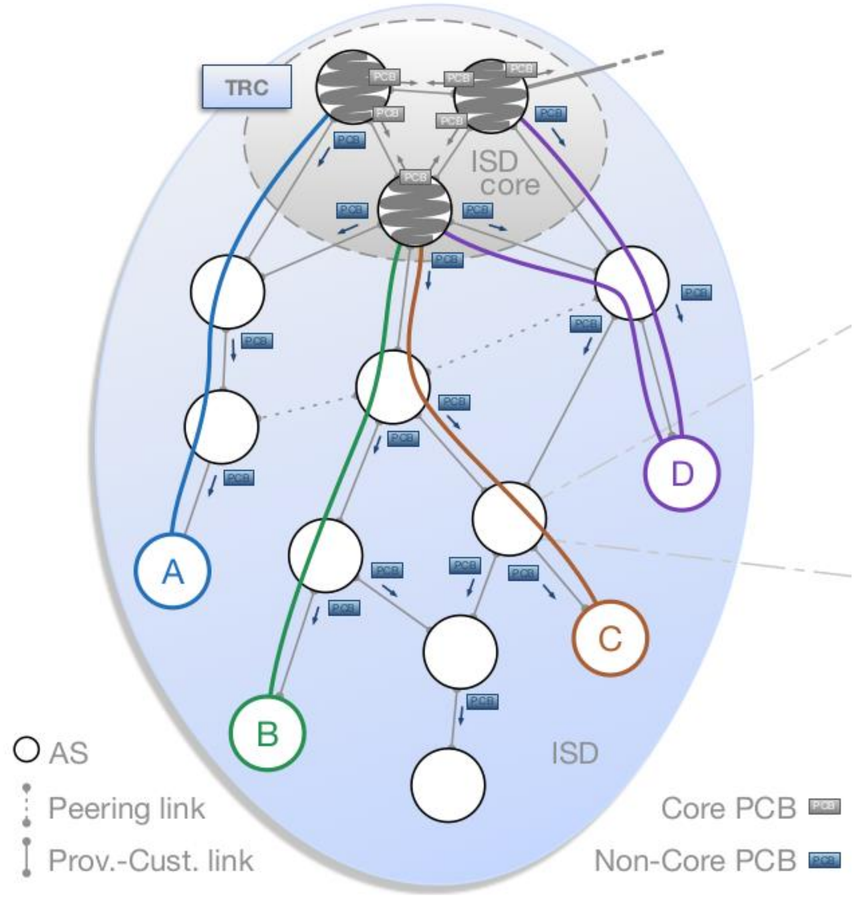

Source: The SCION Internet Architecture: An Internet Architecture for the 21st Century, Barrera et al., 2017

Control plane: Path Construction Beacons

- Path Construction Beacons are signed by every AS along the path
  - Authenticated path

- Hop fields included that can be used to later select paths
  - Contain forwarding information
  - Contain cryptographic MAC computed using hop field key
  - Only processed locally

Routing: policies

- The possible paths that a client can is determined by:
  - **Up-stream AS**, by deciding which PCBs to forward to where
  - **Core AS**, by offering path segments to path server in local AS
  - **Local AS**, by registering down-path segments with ISD core
  - **Local AS**, by offering path segments to clients
  - **Clients**, by combining path segments offered by local path server

- RP Report on: SCION Routing Policies

  A Comparative Analysis of Routing Policies in BGP and SCION by *Kaj Koole*, *Martyna Pawlus*, (2023).
  https://2stic.nl/downloads/routing_policies_koole_pawlus.pdf

Routing summary

- Path information included in packet headers
  - Corresponding hop fields included
  - No forwarding information necessary at routers
  - Packet-carried forwarding state (PCFS)
- Sender selects the path
  - Possible to use multiple paths
  - Fast failover
- Recipient address no longer used to route between autonomous systems
  - Only used by the destination AS
  - Local delivery is responsibility of destination AS

Security

- Path information authenticated in control plane and data plane
- Control plane
  - Beacons authenticated using digital signatures
  - No route hijacks
- Data plane
  - User selects path
  - Hop fields ensure only authorised paths possible

Security

- Address spoofing no longer possible on AS-level
  - Protects against reflection attacks
  - Reduces impact of DDoS attacks
- Hidden paths
  - Path information not published
  - Can only be used by parties that know the relevant hop fields

- EPIC (Every Packet Is Checked) is a dataplane extension that offers source authentication and path validation

Reliability and QoS

- Redundancy through use of multi-path communication
- Fast failover in case of link failure
  - No waiting for convergence
- Possible to add latency information to beacons
  - Path selection based on latency
- COLIBRI/Hummingbird extension
  - Minimum bandwidth reservation

Deployment

- Open source implementation available
  - https://github.com/scionproto/scion
- International testbed SCIONLab
  - https://www.scionlab.org/
- Anapaya builds SCION capable network hardware.
- In use at banks, government and hospitals (in .ch)

Transitioning to SCION

- Can be combined with existing applications using SCION-IP Gateway

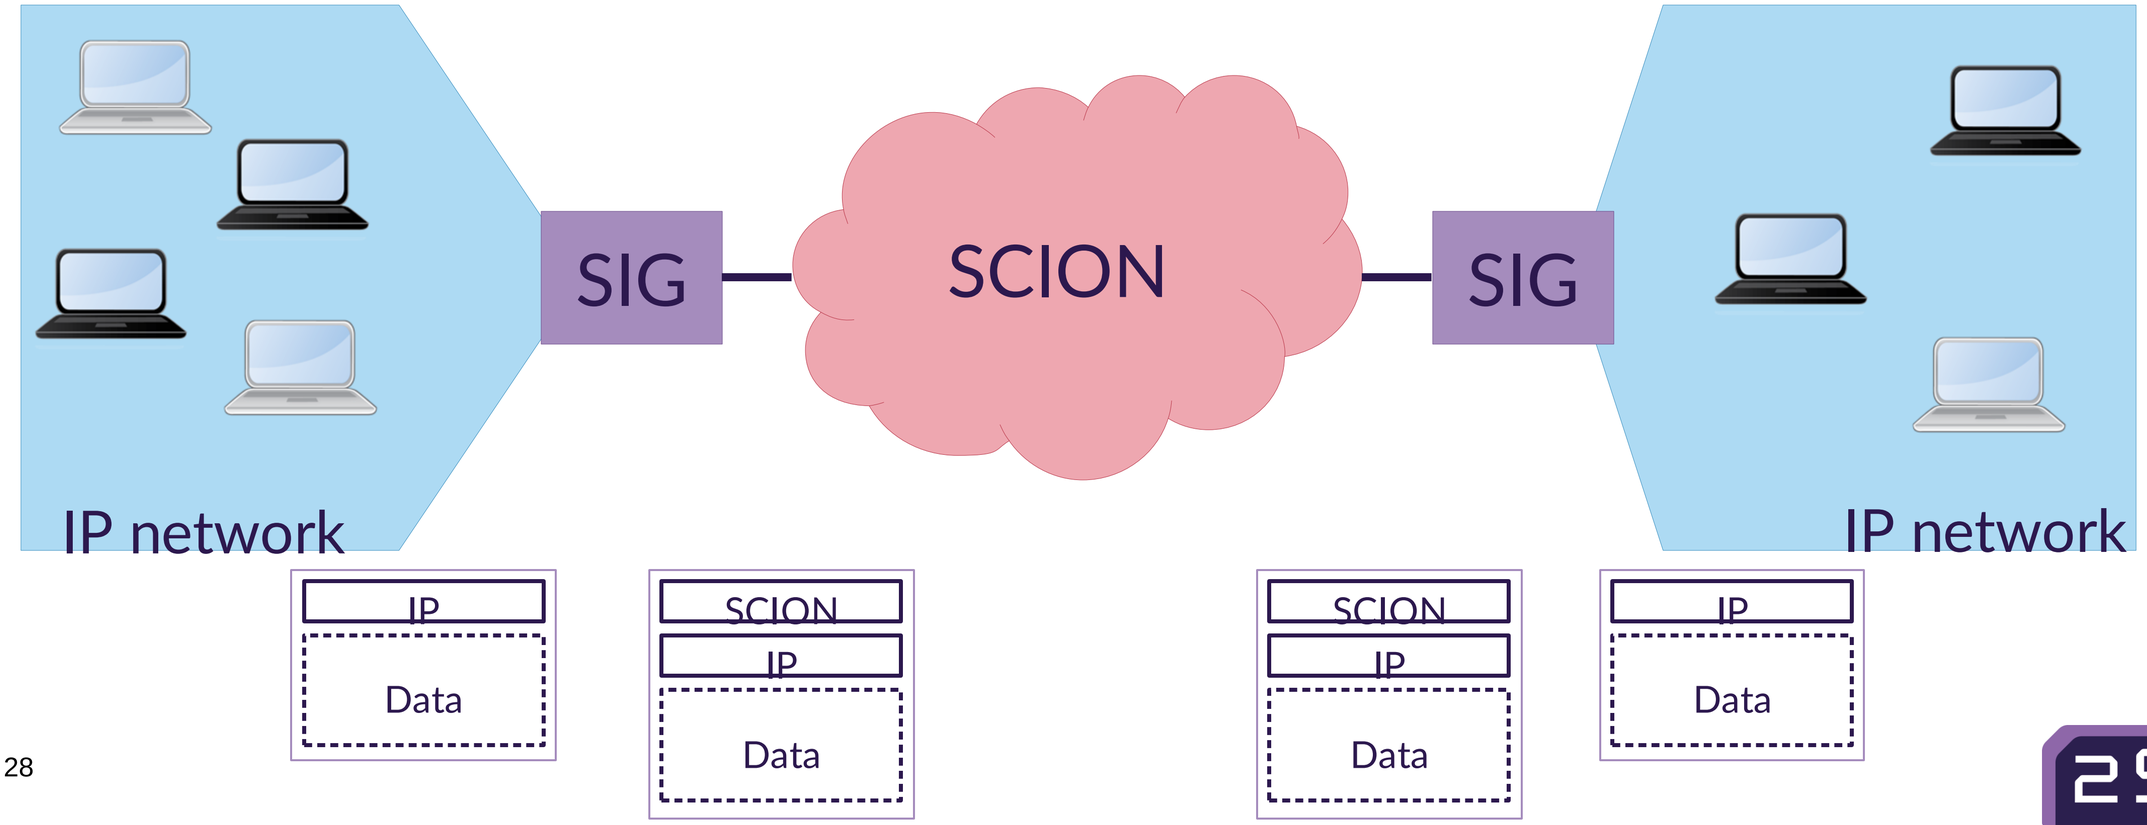

SCION in P4

A new internet architecture in P4

- We implemented the SCION internet architecture in P4 for the Intel Tofino
- Determine feasibility of running a new architecture on switch hardware and evaluate performance

P4

“Programming Protocol-independent Packet Processors (P4) is a domain-specific language for network devices, specifying how data plane devices (switches, NICs, routers, filters, etc.) process packets.”

Source: www.p4.org

Intel Tofino

- P4-programmable Ethernet switch ASIC
- Switches available with e.g., 32 or 64 100 Gbit/sec ports

Some challenges

- No support for cryptographic operations in Intel Tofino
- Protocol not designed for hardware
  - Complex headers

No cryptographic operations

- MACs verified using table containing all currently valid values
- Populated from control plane when MACs are generated
  - In the SCION control plane
  - At the switch
- Invalid entries removed

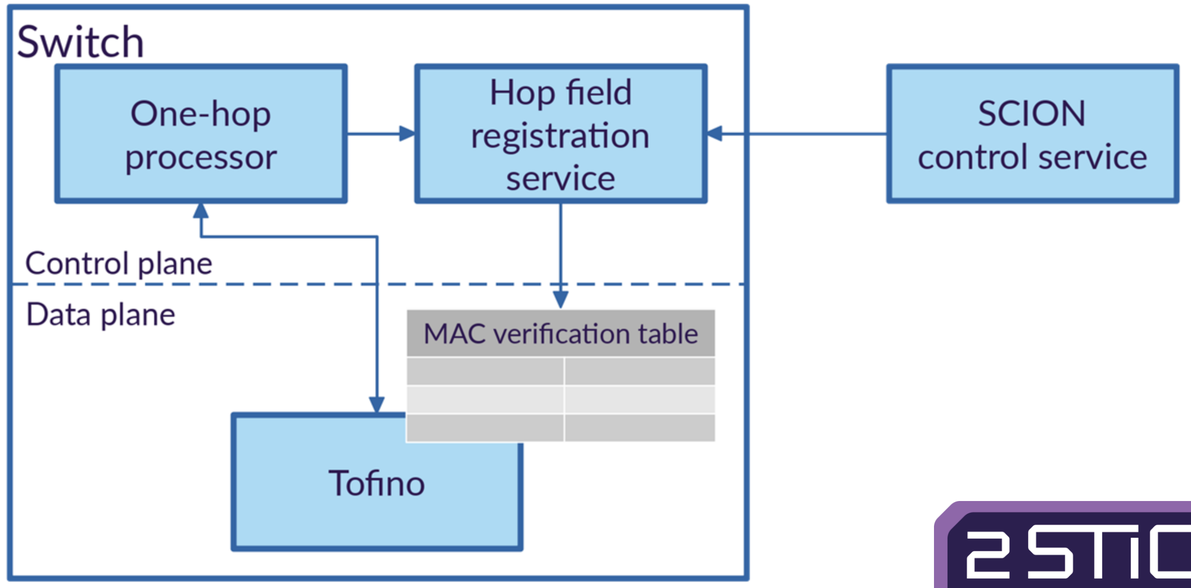

Complex header fields

- For example: forwarding path consisted of nested lists
- Flattening the structure provides for more efficient parsing

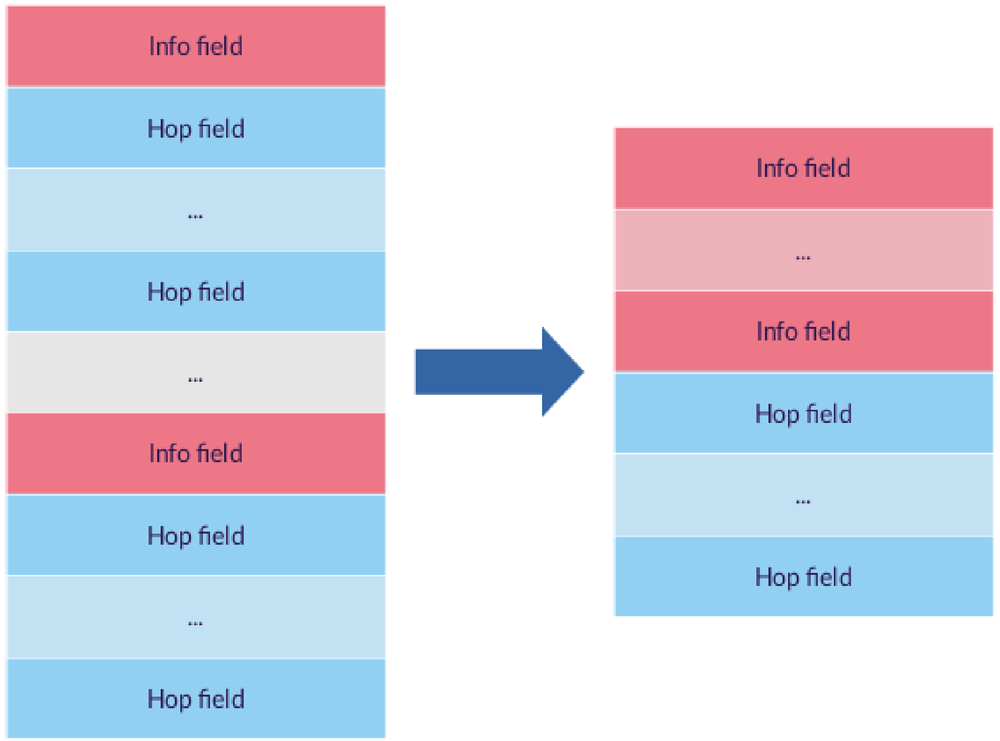

Lessons learned

- When designing a protocol with hardware in mind
  - use explicit lengths
  - do not use absolute offsets
  - limit the usage of variable length fields
  - do not use complex data structures such as nested lists

Evaluation

- Edgecore switches with 32 100 Gbps ports
- Tested functionality with topology where all border routers ran on switches
- Tested performance using packet generator for different path lengths
  - Achieved near line-rate for almost all tested path lengths

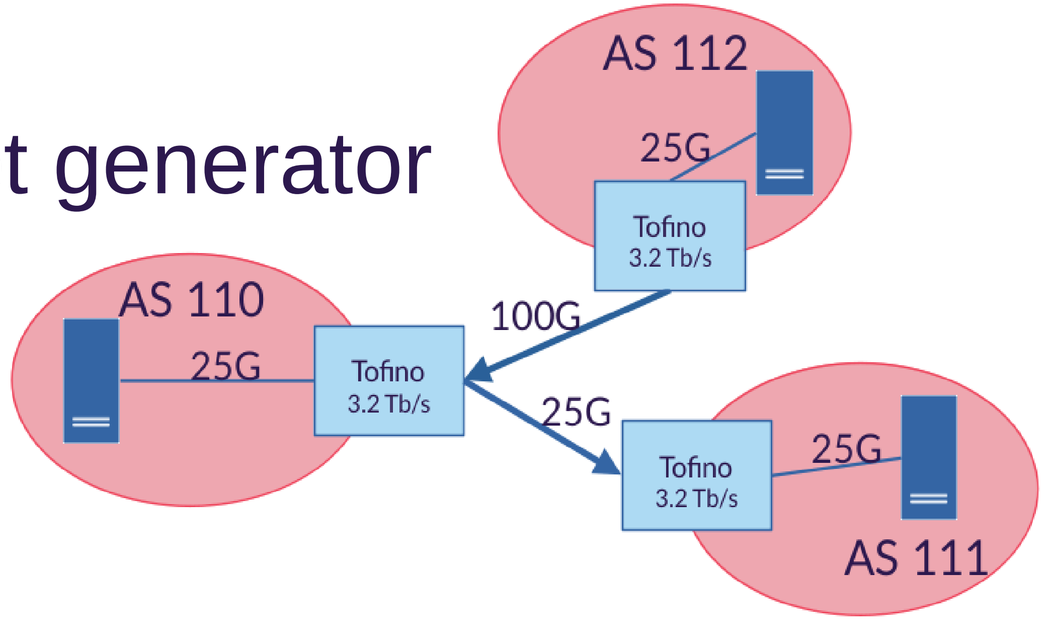

Conclusion

- SCION can be implemented for switch hardware and run on high speeds
- Several lessons learned regarding protocol design
- Future work
  - Support for protocol error handling and additional SCION-related protocols
  - More extensive performance analysis
    - Code is open source and available at github.com/SIDN/p4-scion

Deployments and Apps

ISD Assignments

ISD	Description
0	The wildcard ISD.
1 - 15	Reserved for documentation and sample code [0].
16 - 63	Private Use [1]. They can be used for testing and private deployments.
64 - 4094	Public ISDs. They should be allocated in ascending order, without gaps and "vanity" numbers.
4095 (2^{12}-1)	Reserved.
4096 - 65535	Reserved.

Well known ISDs

ISD	
64	Switserland
65	Europe
67	North Ameica
71	SCIERA
16-22 (private)	SCION Lab
59 (private)	SCION NL

https://github.com/scionproto/scion/wiki/ISD-and-AS-numbering
https://docs.anapaya.net/en/latest/resources/isd-as-assignments/

AS Assignments

AS	Size	Description
0	1	The wildcard AS.
1 - 4294967295 (~ `0:0:0/16`)	~4.3 bil	32-bit BGP AS numbers [2], formatted as decimal. If a BGP AS deploys SCION, it has the same AS number for both BGP and SCION.
`1:0:0`	1	Reserved.
`2:0:0/16`	~4.3 bil	Public SCION-only ASes (i.e. ASes that are created for SCION, and aren't existing BGP ASes). They should be allocated in ascending order, without gaps and "vanity" numbers.
`ff00:0:0/32`	65535	Reserved for documentation and test/sample code [0].
`ffaa:0:0/24`	~16.8 mil	Reserved for Private Use [1]. They can be used for testing/private deployments.
`ffff:ffff:ffff`	1	Reserved.

https://github.com/scionproto/scion/wiki/ISD-and-AS-numbering

SCION address structure

- An AS: **`ISD-AS`**
- A host inside an AS: **`ISD-AS`**`,[`**`address`**`]`

- Examples:
  - `19-ffaa:0:1305` (ISD-AS no endpoint addr)
  - `19-ffaa:0:1305,[::1]` (IPv6 endpoint)
  - `19-ffaa:0:1305,[127.0.0.1]` (IPv4 endpoint)
  - `19-ffaa:0:1305,127.0.0.1` ([] can sometimes be omitted)
  - `59-1140,[::1]` (Shorthand ASN if <=32 bits)
  - `59-1124,145.100.135.68` (Other addr than border-router)

2STiC: SCION-nl testbed

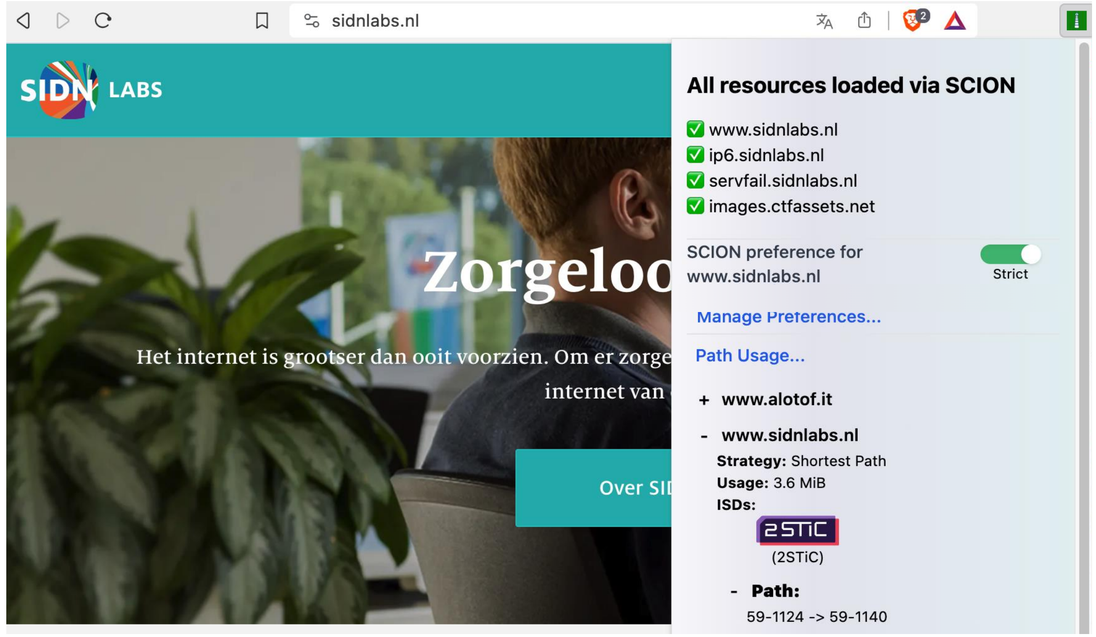


```
59-1124$ python3 path_query.py 59-1140 asis,gdpr
PATH id: 0 size: 1 path: ['59-1124 1>1']
  59-1124 ['gdpr: True']
  59-1140 ['asis: https://asis.sidnlabs.nl:18080', 'gdpr: True']

PATH id: 1 size: 1 path: ['59-1124 3>4']
  59-1124 ['gdpr: True']
  59-1140 ['asis: https://asis.sidnlabs.nl:18080', 'gdpr: True']

PATH id: 2 size: 2 path: ['59-1103 1>2', '59-1140 2>2']
  59-1124 ['gdpr: True']
  59-1103 ['gdpr: True']
  59-1140 ['asis: https://asis.sidnlabs.nl:18080', 'gdpr: True']
```

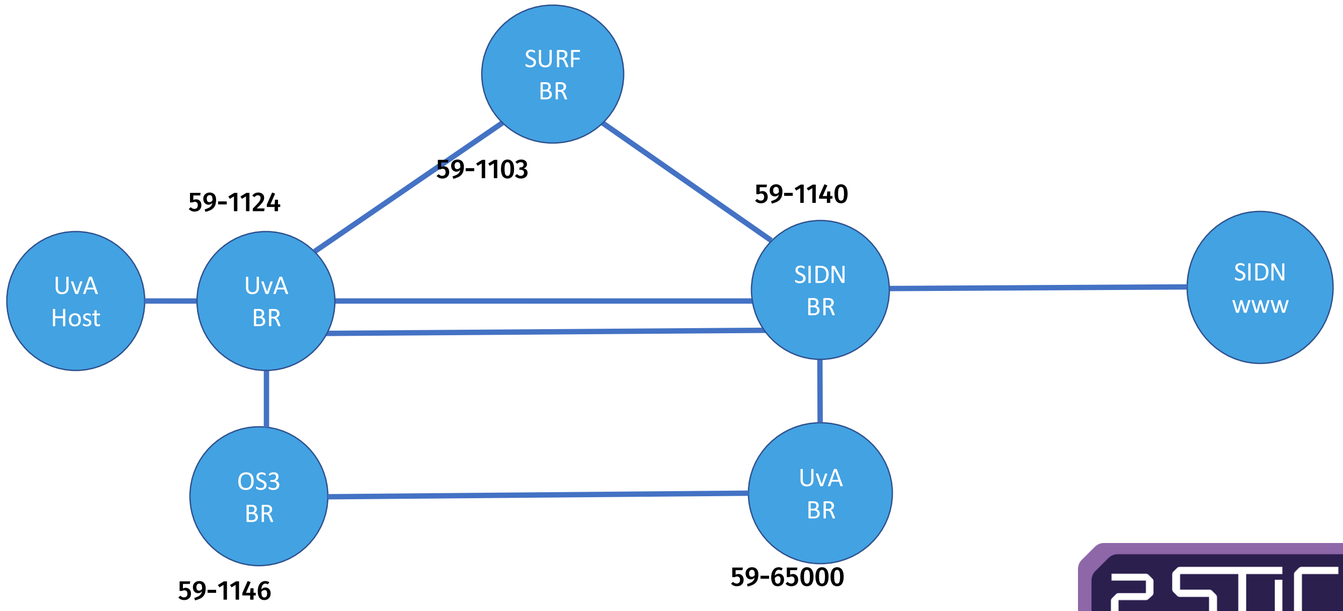

With SCION-nl we can experiment with user controllable routing:

A user can choose to *only use routes* that have the GDPR label.

https://www.2stic.nl/scion-testbed.html

SCIERA: SCION Education Research and Academic

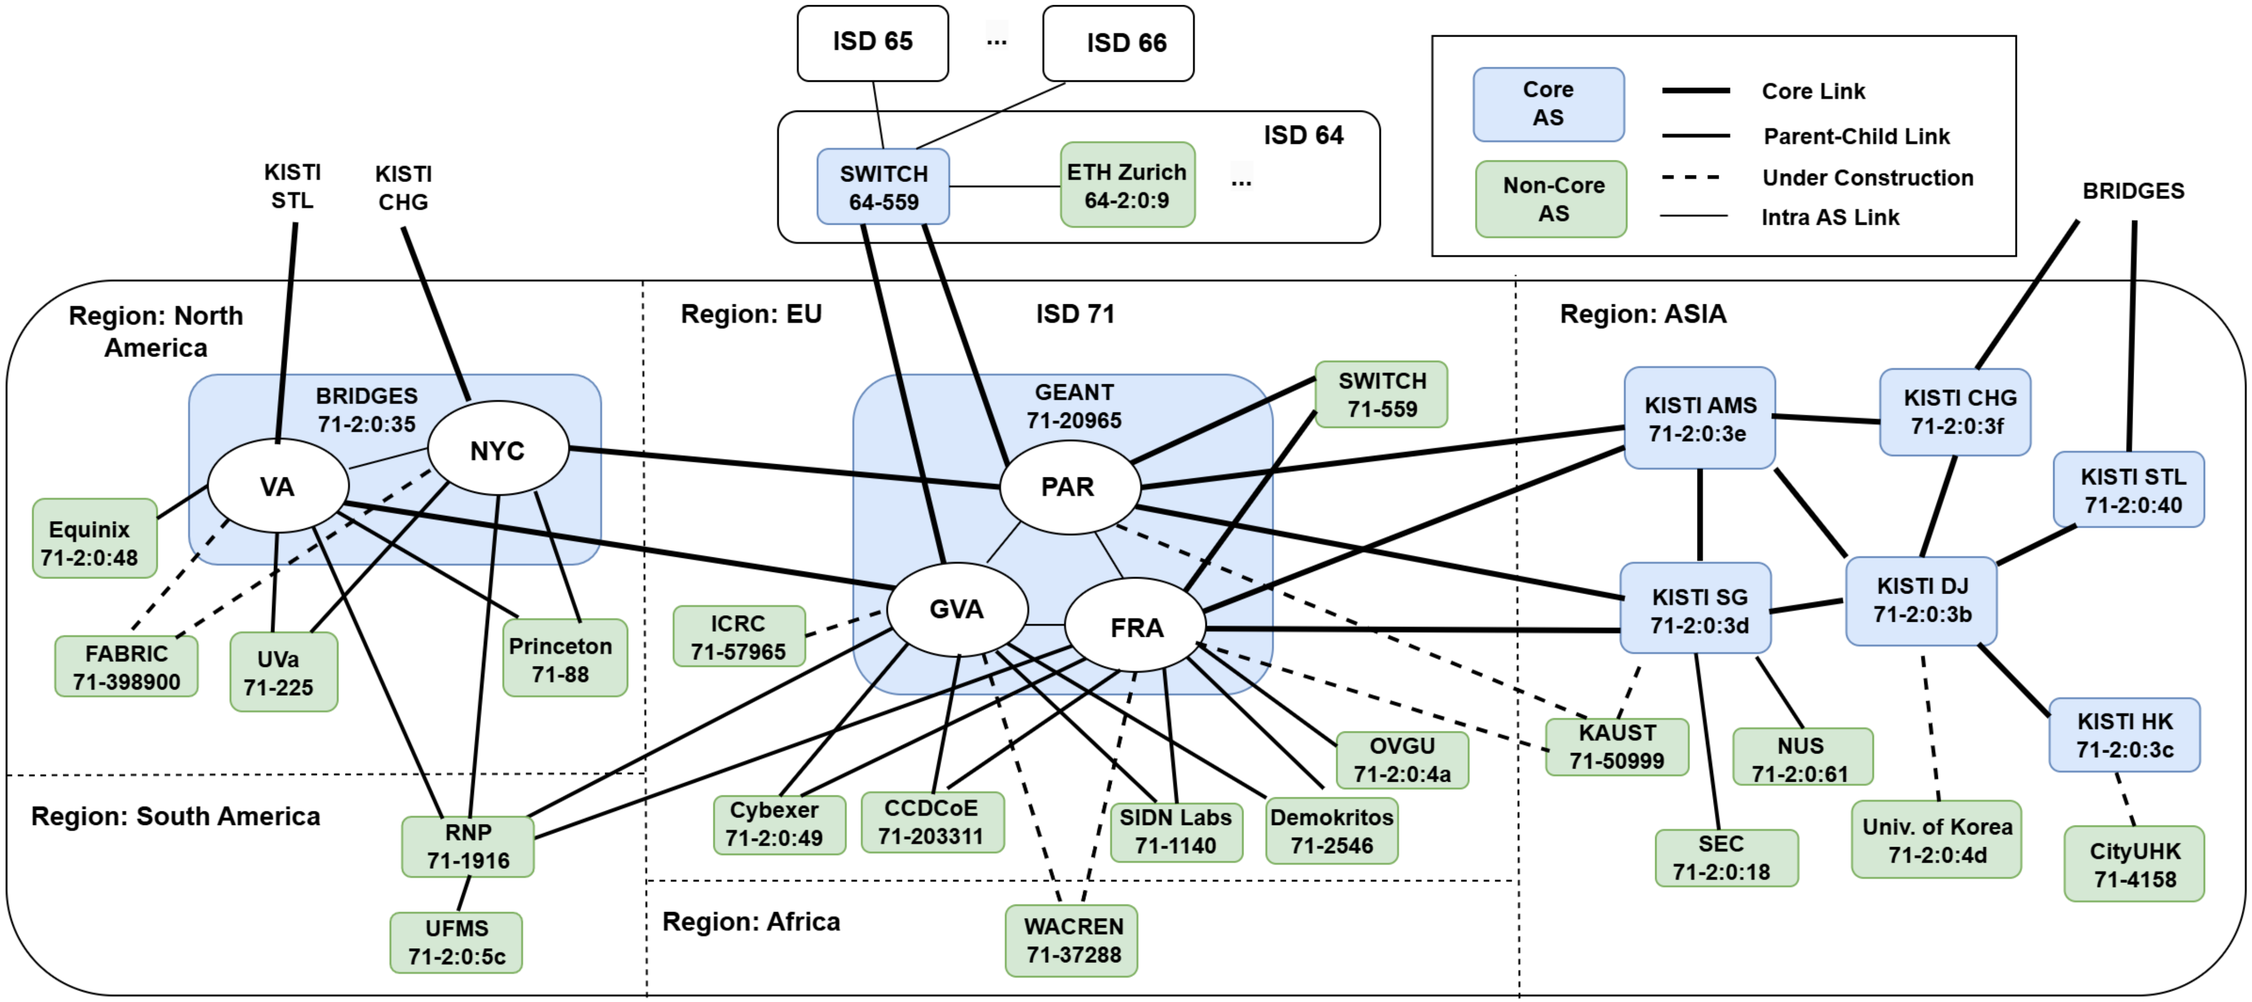

Den Haag, 30 januari 2025

Odido Business en Anapaya introduceren nieuwe security-technologie SCION in Nederland

SCION is de nieuwe generatie internetinfrastructuur, ontworpen voor bedrijven en overheden met vitale infrastructuren.

Als eerste provider in Nederland introduceert Odido de nieuwe beveiligingsoplossing SCION van Zwitserse softwarebedrijf Anapaya voor de zakelijke markt in Nederland. Dit SCION-beveiligingsprotocol, dat is ontwikkeld door het Federal Institute of Technology Zurich, is de nieuwe standaard op het gebied van veiligheid en specifiek ontworpen voor bedrijven en overheden met vitale internetinfrastructuren. Met de samenwerking met Anapaya is Odido de eerste provider in Nederland die deze nieuwe technologie aanbiedt aan klanten met kritieke infrastructuren.

Nieuwe en extra bescherming

← Voorpagina newsroom

Deel deze release

Using SCION

- SCION
  - https://github.com/scionproto/awesome-scion
  - https://github.com/scionproto/scion
- SCION Lab
  - https://www.scionlab.org
  - https://docs.scionlab.org
- Last years' SCION AN lab
  - The lab should work except for exercise 3 (just ignore the topology).
  - https://check.sidnlabs.nl/ralph/anet-lab/

DISCUSSION

Discussion

To recap, today we've seen:

- Programmable network hardware;
- A new internet architecture;
- An active community.

With these developments in mind, we wonder: how do you see the future of the Internet?

Discussion

Routing transparency

- Do users need this much transparency and control?
- 3rd parties can tell a lot about how you are connected to the internet by looking at the headers in a single packet, is this desirable?

Isolation Domains

- What is a sustainable governance structure for an ISD? Should this be government controlled or not?
- Will ISDs (and extensions) violate net neutrality?

2STIC

Thanks for your attention!

Caspar Schutijser, Ralph Koning
sidnlabs.nl
2stic.nl