



FAKULTÄT FÜR
INFORMATIK

Enabling Path Awareness for Legacy Applications through SCION-IP Translation utilizing eBPF and Tofino

M.Sc. Lars-Christian Schulz
M.Sc. Robin Wehner

M.Sc. Florian Gallrein
Prof. Dr. David Hausheer

Structure

1. Background – SCION
2. Motivation for a SCION IP Translator
3. Concepts – Translating IP addresses to SCION
4. eBPF & P4 Implementation on Tofino
5. Evaluation of Performance and Resource Usage
6. Conclusion & Future Work

1. Background: SCION Architecture




💡 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



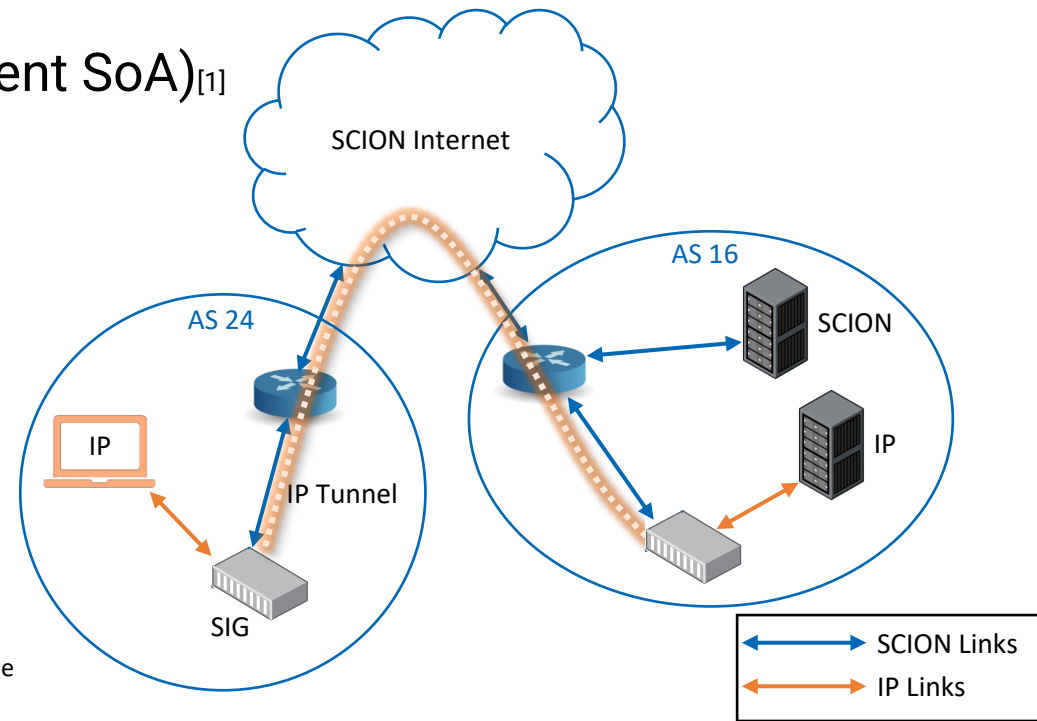
2. Motivation

- ❖ SCION solves many issues of today's Internet
- ❖ How can legacy IP applications use the SCION network?



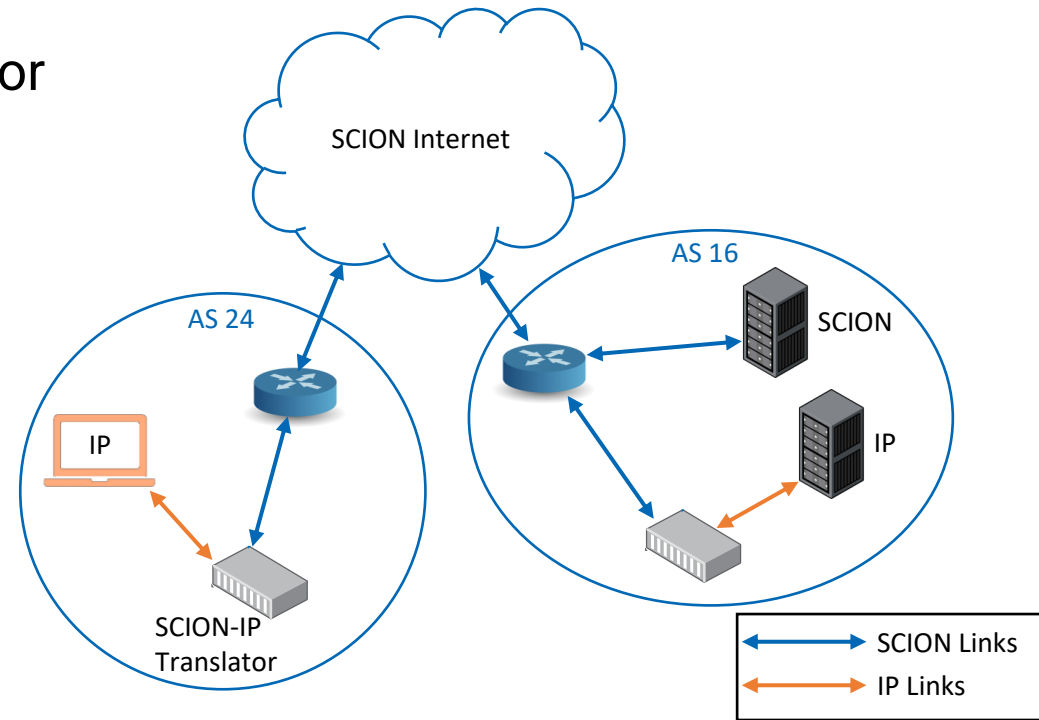
2. Motivation

- ❖ SCION solves many issues of today's Internet
- ❖ How can legacy IP applications use the SCION network?
 - Use SCION-IP-Gateway (SIG) (Current SoA)^[1]
 - Tunnels IP over SCION
 - **Complex encapsulation**
 - **Have to exchange IP routes with each other**
 - **No communication between SCION and IP apps**



2. Motivation

- ❖ SCION solves many issues of today's Internet
- ❖ How can legacy IP applications use the SCION network?
 - We propose the SCION IP Translator
 - Allows IP applications to use the SCION network
 - Enables interoperability between SCION and IP implementations of an app
 - No encapsulation overhead



3. Concepts - Translating IP addresses to SCION

- ❖ SCION addressing: **ISD** – **AS** – Host Address

Example server: **1** **18** **fd00::0d93**

- ❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
global routing prefix				subnet ID	interface ID
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix				interface ID	
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

- ❖ SCION addressing: **ISD** – **AS** – Host Address

Example server: **1** **18** **fd00::0d93**

- ❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

fc

- 8 Bit global SCION prefix

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
	global routing prefix			subnet ID	interface ID
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix				interface ID	
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

❖ SCION addressing: **ISD** – **AS** – Host Address

Example server: **1** **18** **fd00::0d93**

❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

fc00:1

➤ 12 Bit ISD number because currently assigned from range 64-4094

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
	global routing prefix			subnet ID	interface ID
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix				interface ID	
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

❖ SCION addressing: **ISD** – **AS** – Host Address

Example server: **1** **18** **fd00::0d93**

❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

fc00:1000:18

➤ 20 Bit ASN because current BGP allocation from block below 2^{19}

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
global routing prefix			subnet ID	interface ID	
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix			interface ID		
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

❖ SCION addressing: **ISD** – **AS** – Host Address

Example server: **1** **18** **fd00::0d93**

❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

fc00:1000:18xx:xxxx::0d93

➤ 64 Bit Interface ID for simple host address assignment

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
global routing prefix				subnet ID	interface ID
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

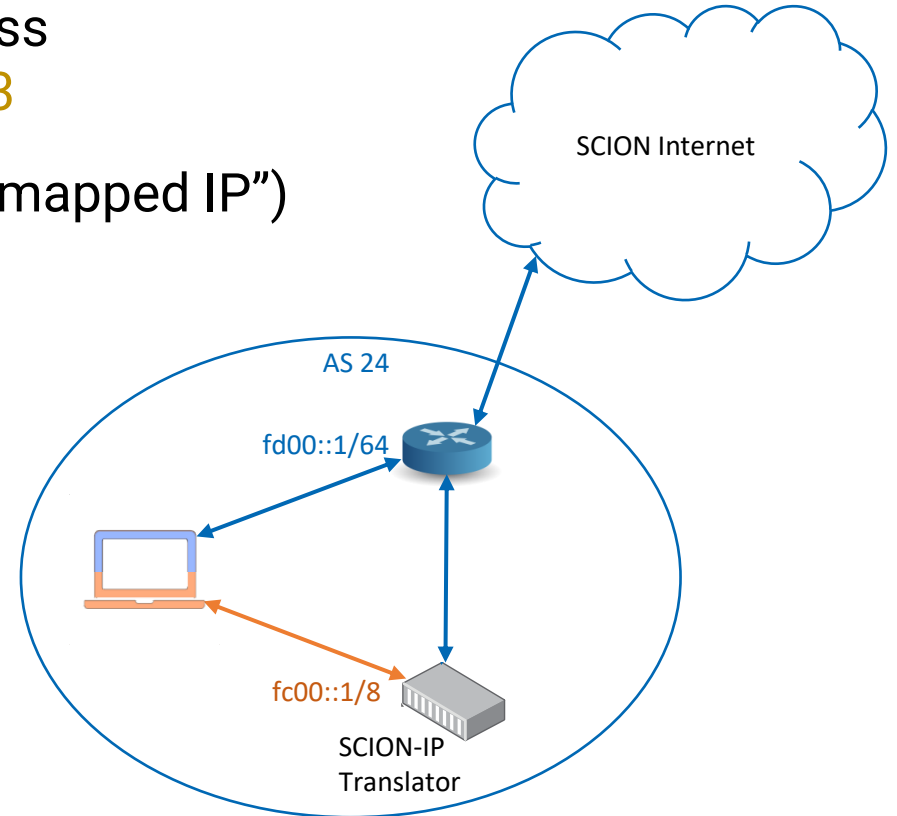
128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix				interface ID	
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

- ❖ SCION addressing: **ISD** – **AS** – Host Address
Example server: **1** **18** **fd00::0d93**

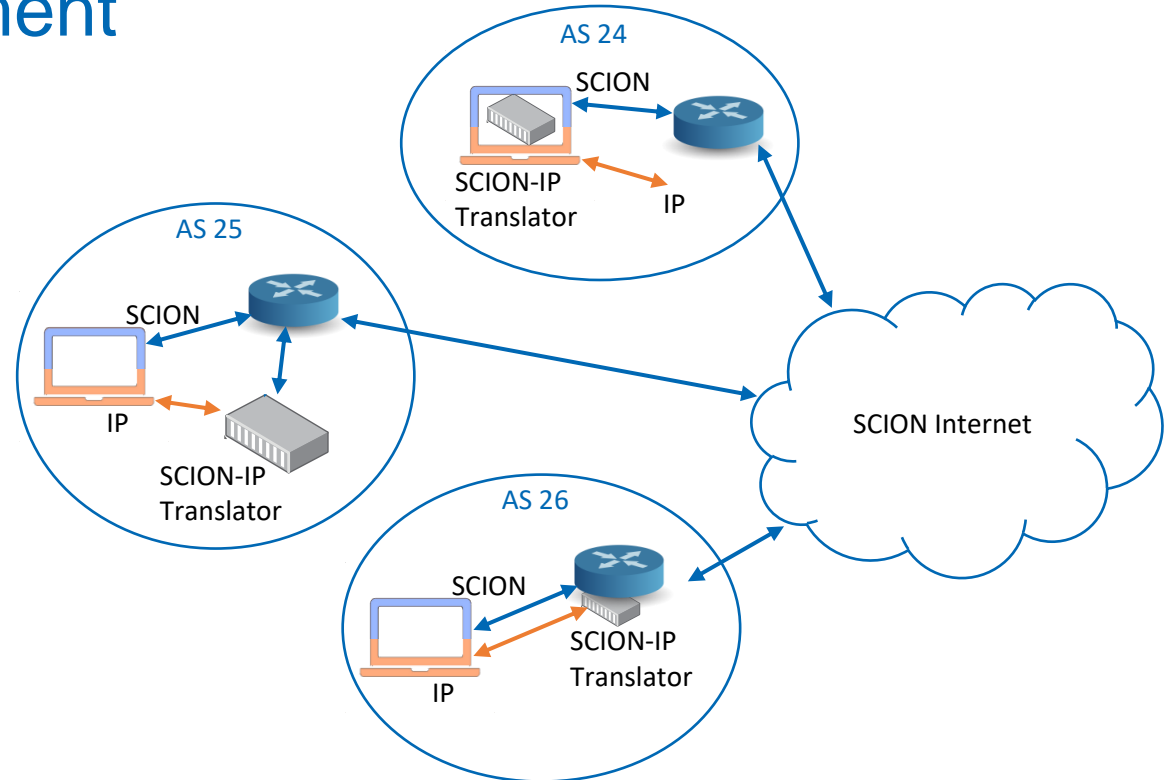
- ❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)
fc00:1000:18xx:xxxx::0d93

- ❖ All hosts need a SCION-mapped IP additionally to its existing IPv6
 - Legacy SCION traffic to `fd00::1`
 - IP traffic to SCION mapped IPv6
fc00:1000:18xx:xxxx::0d93



3. Concepts - Deployment

- ❖ At the host
 - Kernel (eBPF)
 - User Space
 - NIC
- ❖ At the Local Gateway Router
- ❖ At the Border Router/As Network Service (P4 switch)



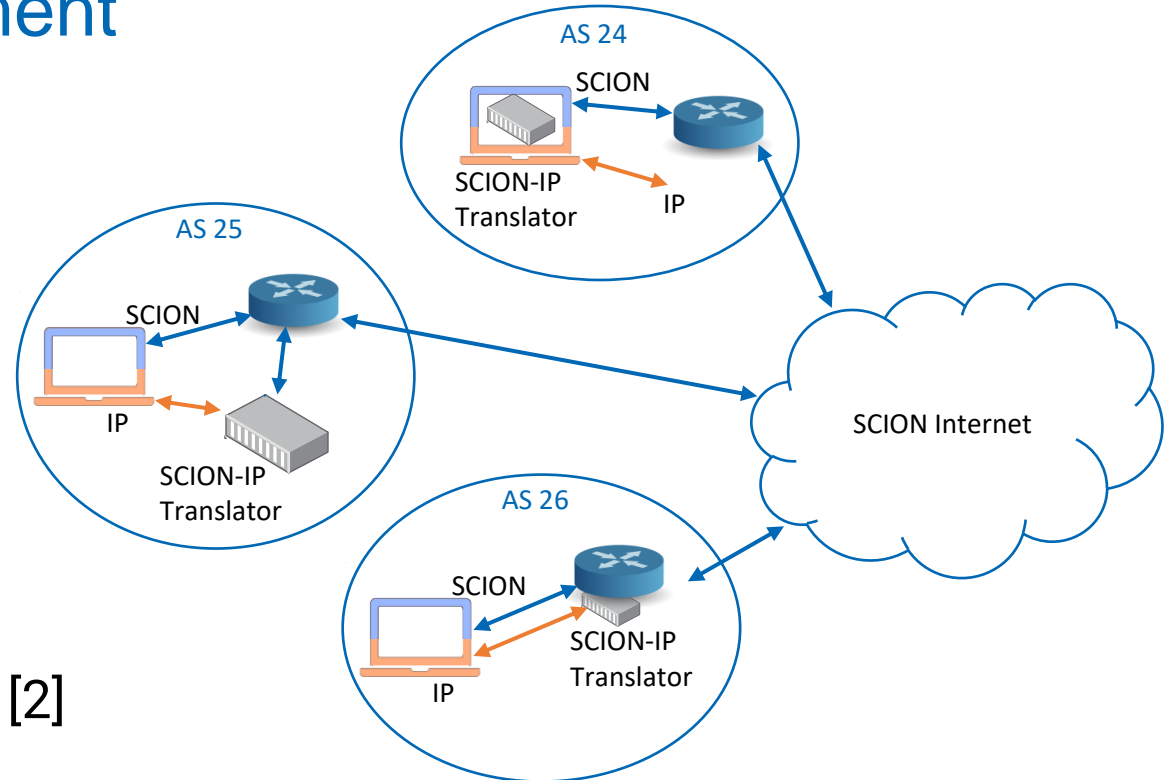
3. Concepts - Deployment

❖ At the host

- Kernel (eBPF) [1]
- User Space
- NIC

❖ At the Local Gateway Router

❖ At the Border Router/As Network Service (P4 switch) [2]

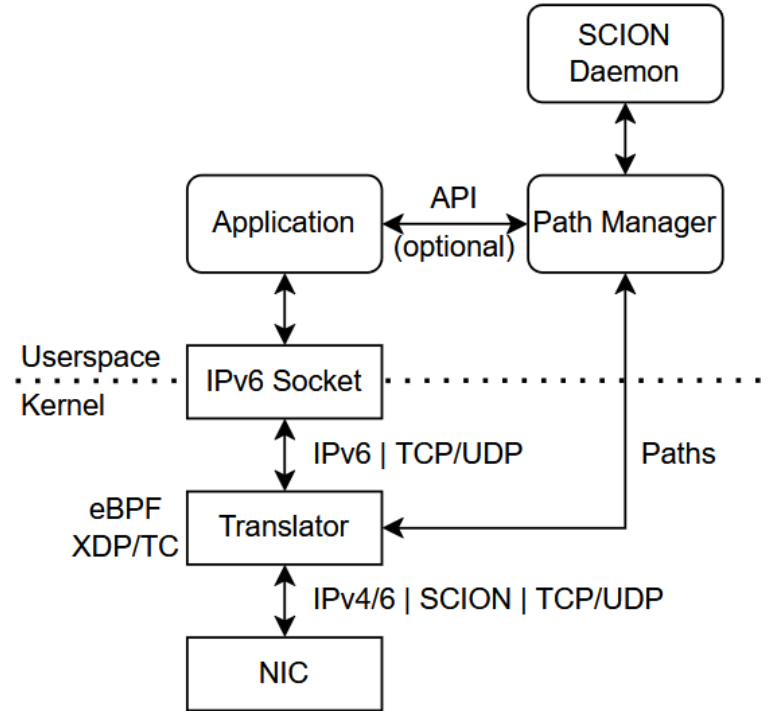


[1] Schulz, et al.: **Unlocking Path Awareness for Legacy Applications through SCION-IP Translation in eBPF.** *SIGCOMM eBPF Workshop 2024.*

[2] Schulz et al.: **Scion Edge Router for Legacy IP Applications based on Intel Tofino.** *ICNP EuroP4 Workshop 2024.*

4. Implementation – eBPF for XDP/TC

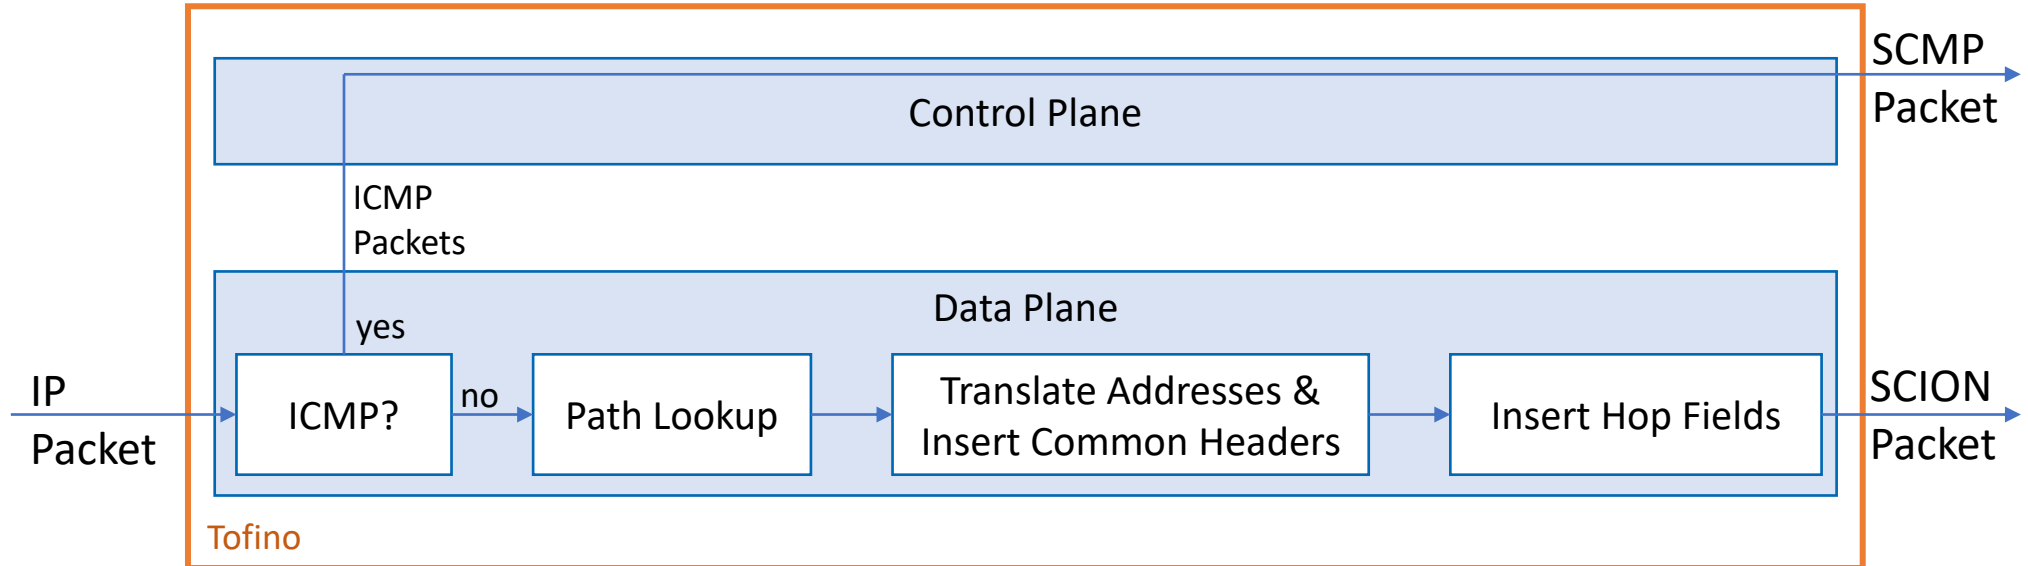
- ❖ Linux: Kernel module
- ❖ Windows: Packet interception driver (Not implemented yet)
- ❖ Control Plane in user space (path manager daemon/service)



4. Implementation – Tofino

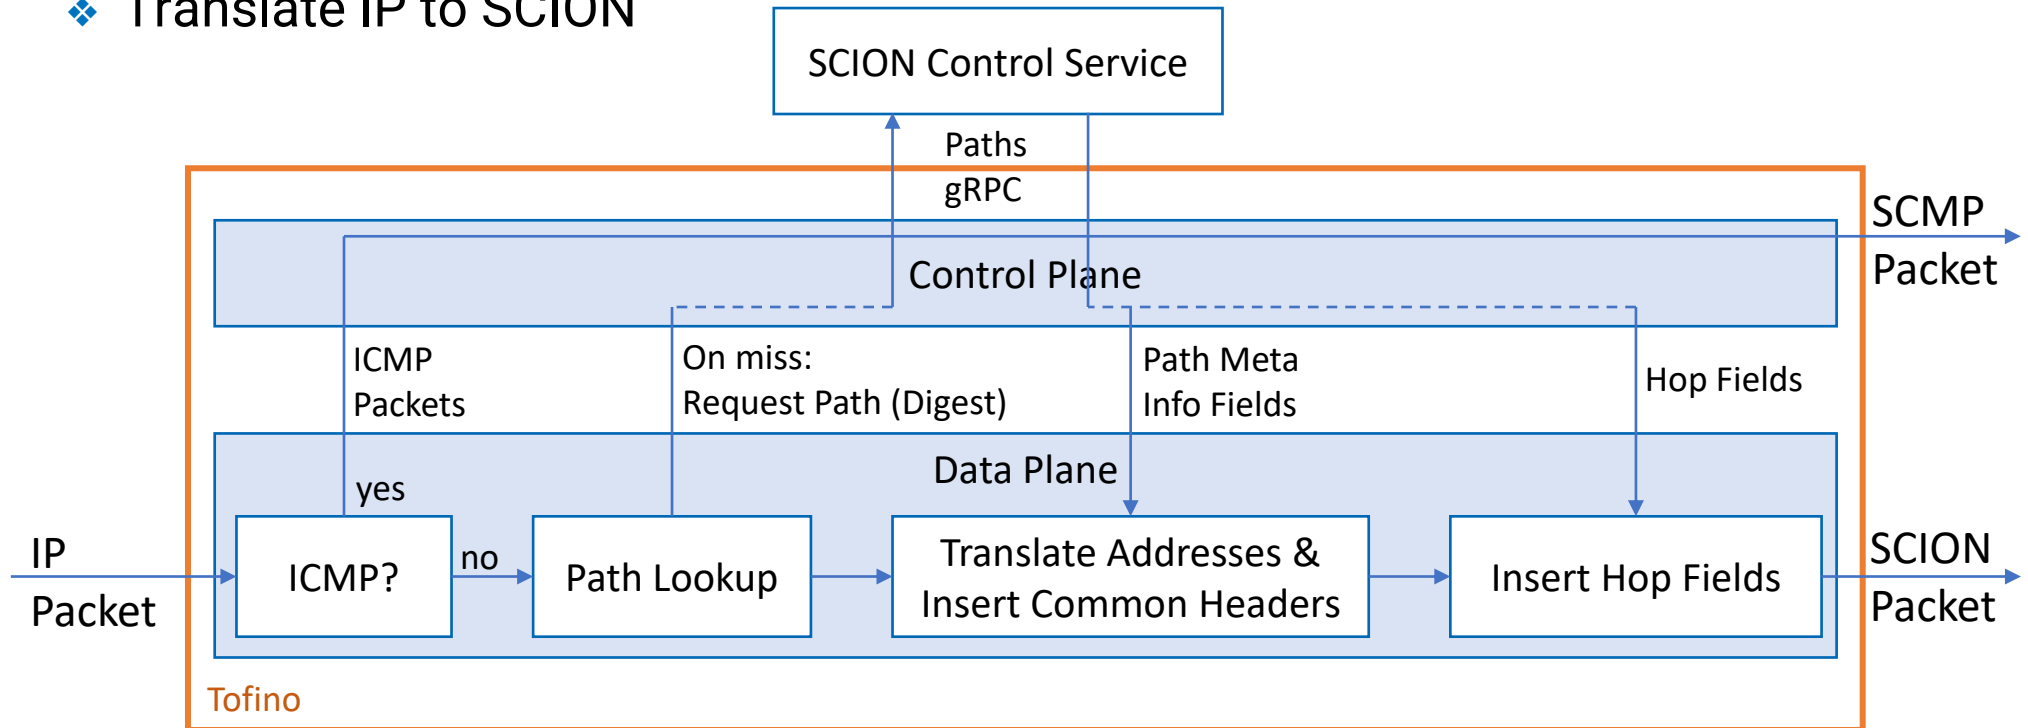
- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate IP to SCION

SCION Control Service



4. Implementation – Tofino

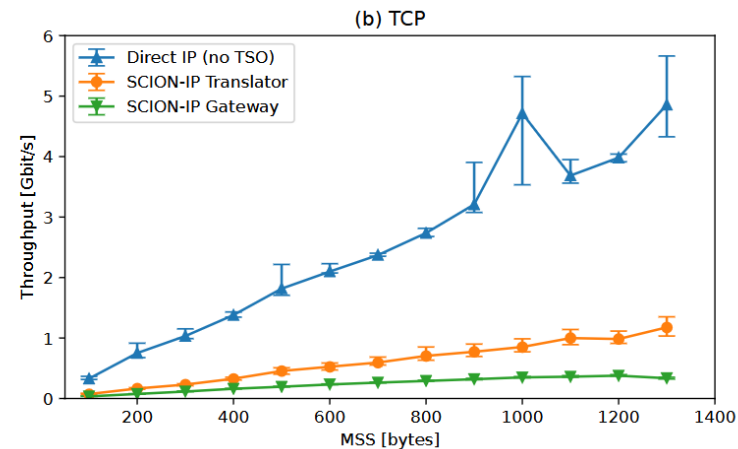
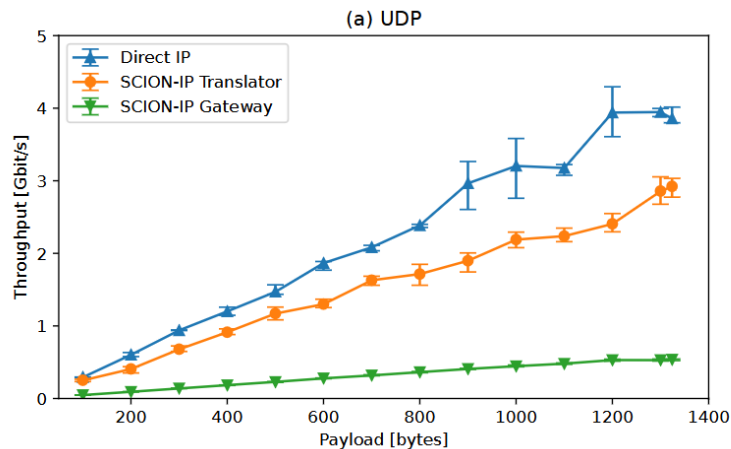
- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate IP to SCION



5. Evaluation

Throughput

- ❖ For eBPF (Server with AMD EPYC CPU (32 core) and Mellanox ConnectX-5 100GbE, traffic generated with iperf3 on 100 Gbit/s links):



- ❖ For Tofino: 400 Gbit/s (i.e. regular line rate performance)

5. Evaluation

Resource Consumption in Tofino

Translation Direction	Performance Criteria	Switch Generation	
		Tofino 1	Tofino 2
IP → SCION	Max. insertable hop fields per pass ¹	16	24
SCION → IP	Required recirculations for MAC validation	5	2

¹Only restricted by PHV registers, as every inserted hop field requires 12 byte

6. Conclusion and Future Work

The SCION-IP translator...

- ❖ Enables legacy IP applications over SCION
- ❖ Allows interconnection between IP and SCION implementations of an application
- ❖ Is easier to implement than the SIG (or rewriting applications to SCION)
- ❖ Is implemented in eBPF (user path choice) & Tofino (higher performance)

We plan to...

- ❖ Improve path selection for Tofino to meet specific applications requirements
- ❖ Implement the other translator concepts as well

Thanks for your attention!

- Papers:
- [1] Schulz, L.C., Gallrein, F., Hausheer, D.: **Unlocking Path Awareness for Legacy Applications through SCION-IP Translation in eBPF**. In: *Proceedings of the ACM SIGCOMM 2024 Workshop on EBPF and Kernel Extensions*. p. 68–70. eBPF '24 (2024) <https://doi.org/10.1145/3672197.3673437>
 - [2] Schulz, L.C., Wehner, R., Hausheer, D.: **Scion Edge Router for Legacy IP Applications based on Intel Tofino**. In: *7th European P4 Workshop (EuroP4'24)* (2024) <https://ieeexplore.ieee.org/document/10858525>

E-Mails: lschulz@ovgu.de
robin.wehner@ovgu.de