SIBRA: Scalable Internet Bandwidth Reservation Architecture

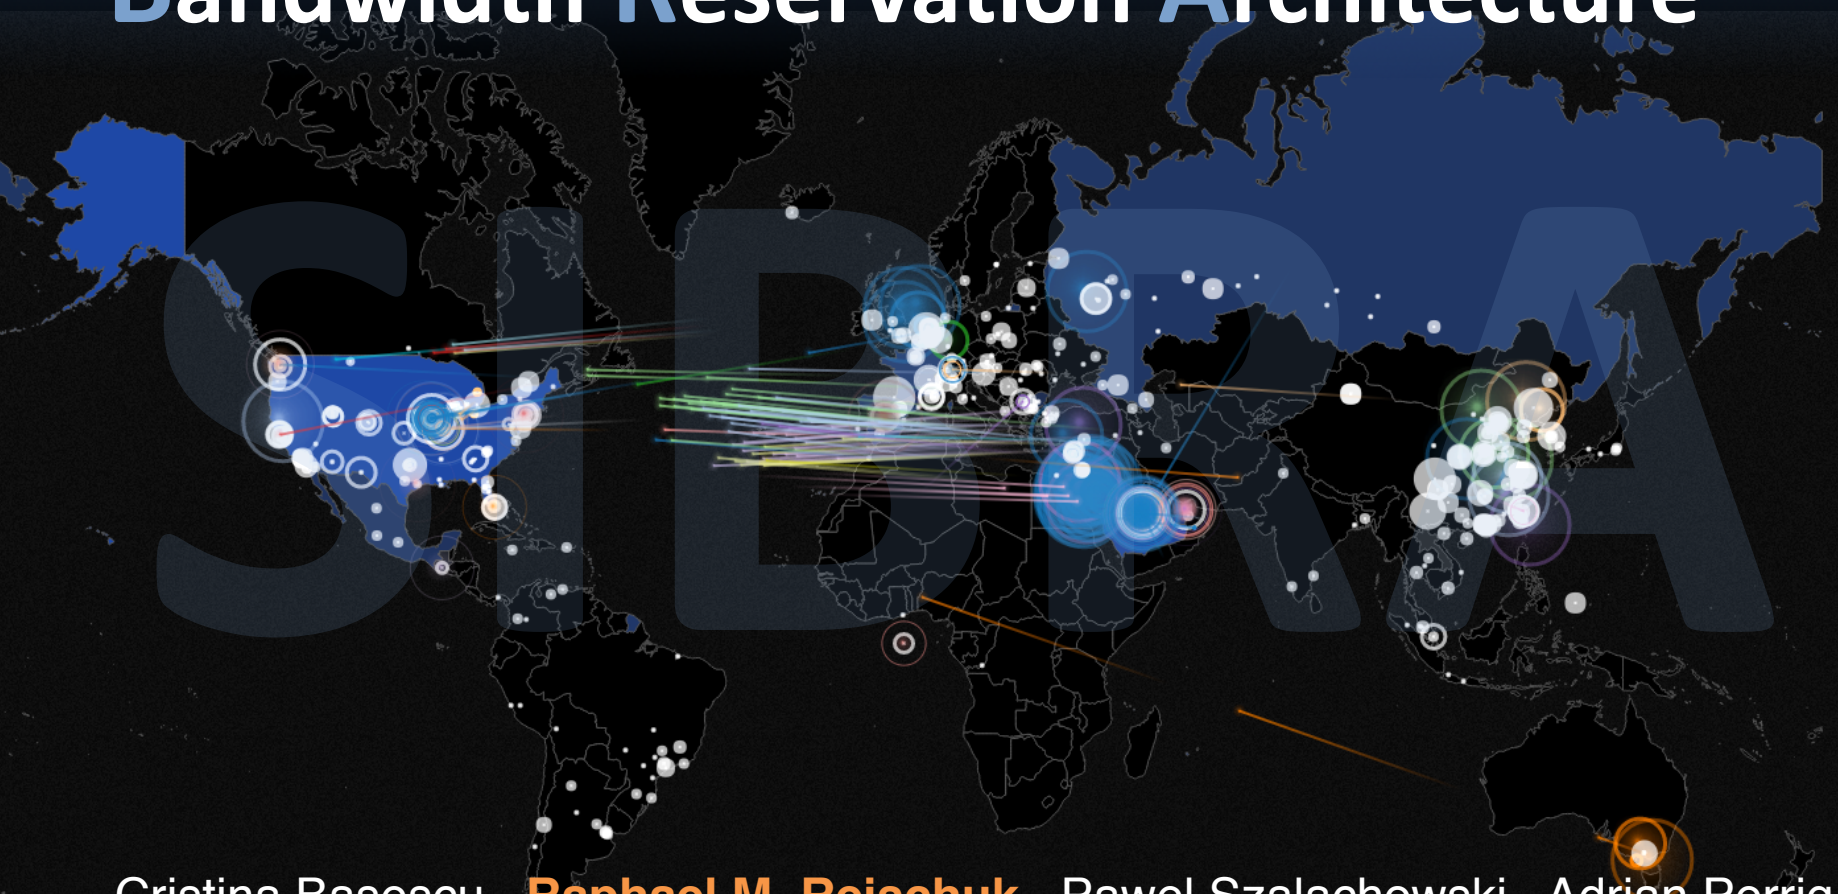

Cristina Basescu, **Raphael M. Reischuk**, Pawel Szalachowski, Adrian Perrig,
Yao Zhang, Hsu-Chun Hsiao, Ayumu Kubota, Jumpei Urakawa

ETH zürich

NDSS 2016, San Diego, CA

picture: http://map.norsecorp.com/

150 hours

180,000 IP addresses
+690,000,000 hits per day
861 user agents

2013

source: http://www.securityweek.com/ddos-attacks-cost-40000-hour-incapsula
picture: https://www.incapsula.com/blog/headless-browser-ddos.html

29 JAN 2015 | NEWS

DDoS Attacks Spike 80% in Q4 2014

CNET › Security › Record-breaking DDoS attack in Europe hits 400Gbps

Record-breaking DDoS attack in Europe hits 400Gbps

A distributed-denial-of-service attack peaked some 33 percent higher than last year's Spamhaus attack, the previous DDoS record-holder.

28 JAN 2016 NEWS

DDoS Attacks Hit Record 500 Gbps in 2015

An Internet of Treacherous Things

A zombie network of home routers highlights the importance of prioritizing smart appliance security.

By Glenn Fleishman on January 13, 2015

11 FEB 2015 | NEWS

IoT Security Systems in Alarming Security Fail

Why are current DDoS defenses inadequate?

Defense Strategies

· **Traffic Scrubbing:** clean incoming traffic from malicious flows

Useless if a link upstream is flooded

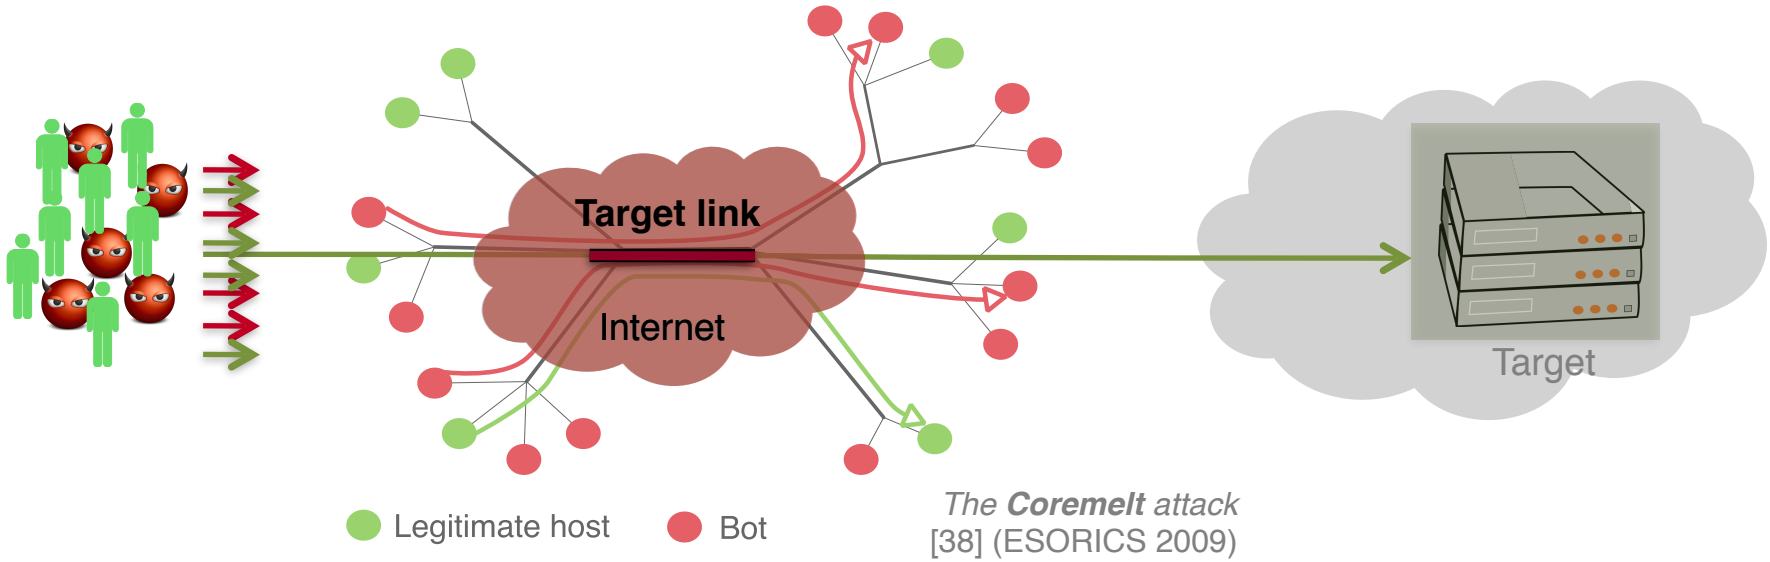

*The **Coremelt** attack*
[38] (ESORICS 2009)

Exploits a characteristic of today's Internet:
(legitimate) end hosts cannot control the path to bypass congested links

· **Network Capabilities:** isolate attack traffic from benign traffic

Useless if links are congested (DoC attacks [32])

Defense Strategies

· **Fair Resource Reservation:** guarantee exclusive usage

Useless in today's Internet since actual allocations would be too small

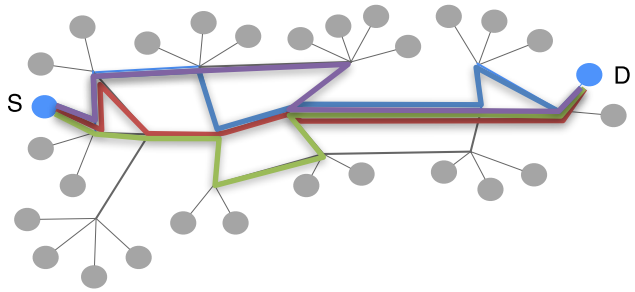

Fair share on every link too small to be useful.

Per flow fair sharing, and similar notions

Everyone has the incentive to increase their "fair share".

Tragedy of the commons, Garrett Hardin (1968)

Current defenses lack a crucial property:

Availability does not diminish — regardless of the botnet size

"Botnet-size independence"

What ingredients do we need for DDoS defense?

SIBRA: Key Ingredients

Group ASes into
Isolation Domains (ISDs)

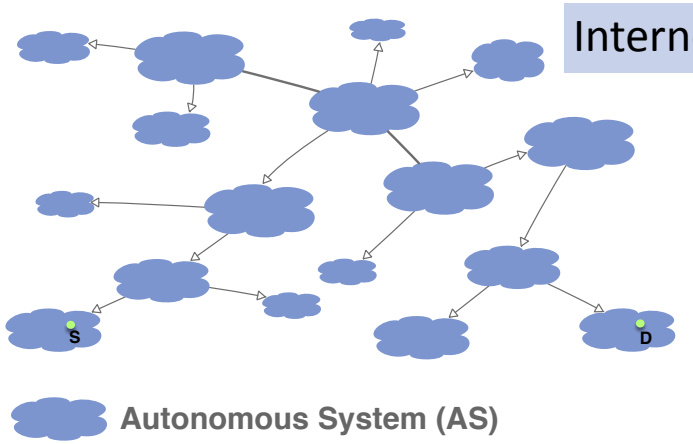

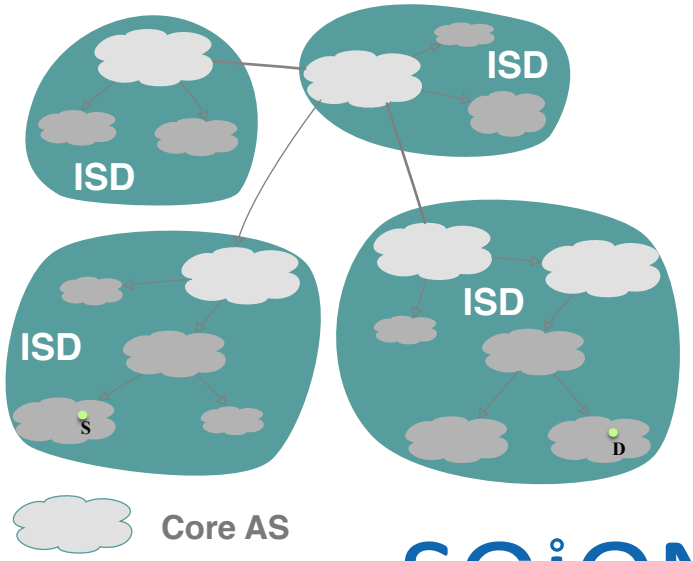

SCION
Internet Architecture

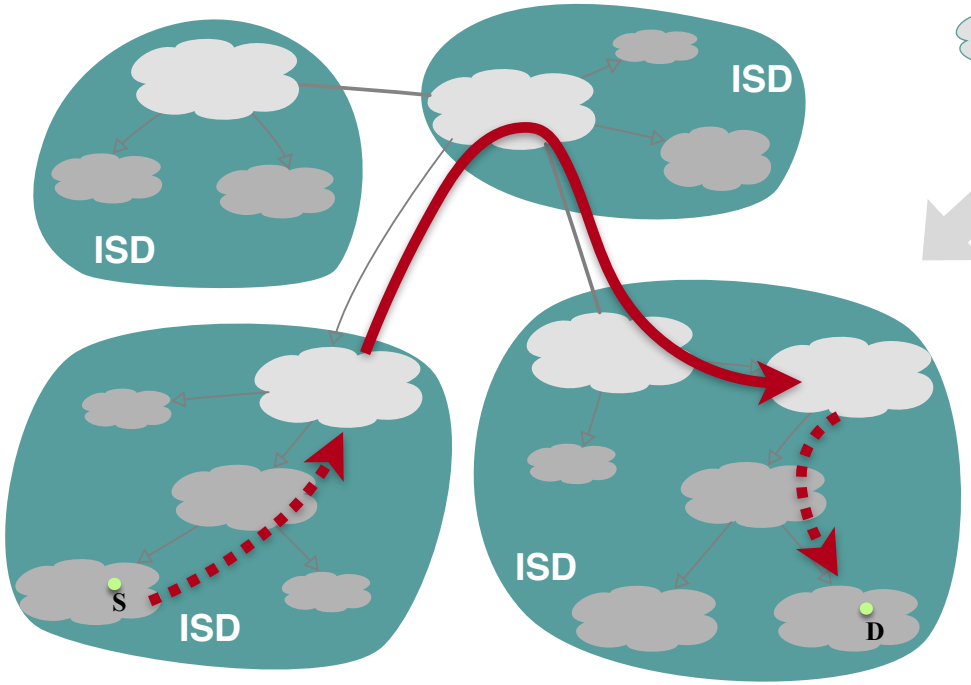

Distribute control
for path construction
& resource allocation

between
- source AS,
- destination AS,
- core ASes

Which notion of fairness is required for
botnet-size independence?

CORE SIBRA Paths

Fairness *between* ISDs: **core paths**

- between ISD Core ASes
- negotiated between direct neighbors
- initiated from destination
- according to previous traffic volumes
- long-term (months)
- optional guarantees
  e.g., 99.99% availability

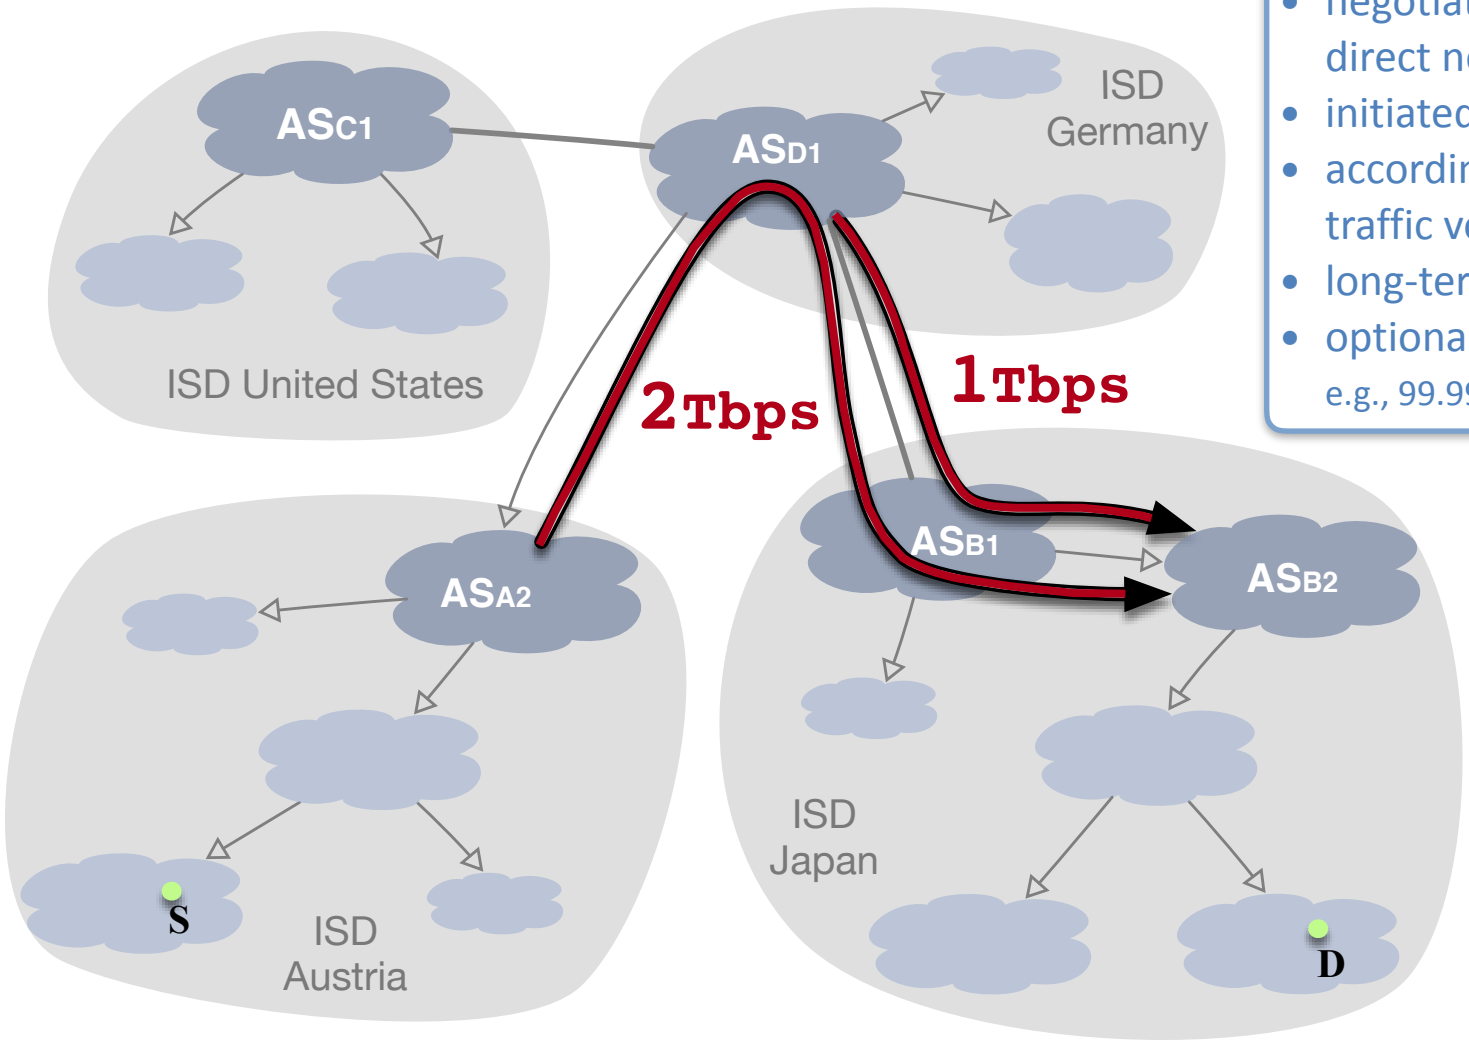

SIBRA Paths

STEADY

Fairness *between* ISDs: **core paths**

Fairness *inside* ISDs: **steady paths**

- requested by inner ASes
- low-bandwidth traffic (control traffic, DNS, ICMP)
- intermediate-term (order of minutes)
- periodically extendable
- basis for launching high-bandwidth reservations
- cryptograph. protected (using local keys)

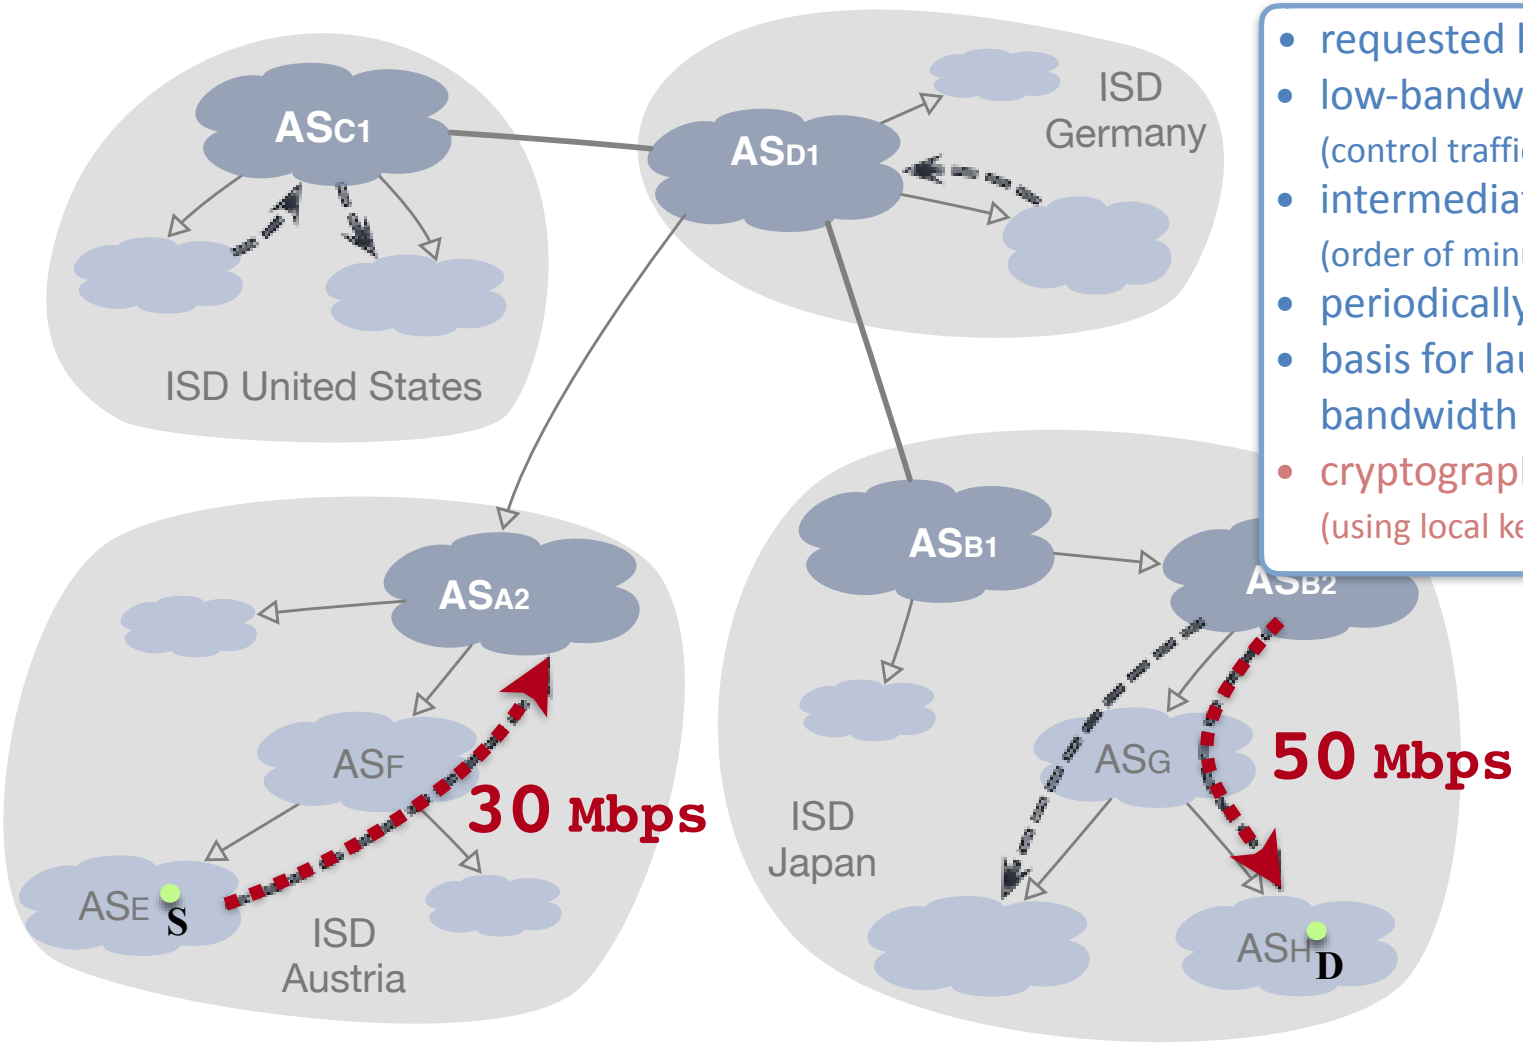

EPHEMERAL

SIBRA Paths

Fairness *between* ISDs: **core paths**

Fairness *inside* ISDs: **steady paths**

E2E reservations: **ephemeral paths**

fairness: per-source and dest. AS

bandwidth proportional to steady paths and core paths

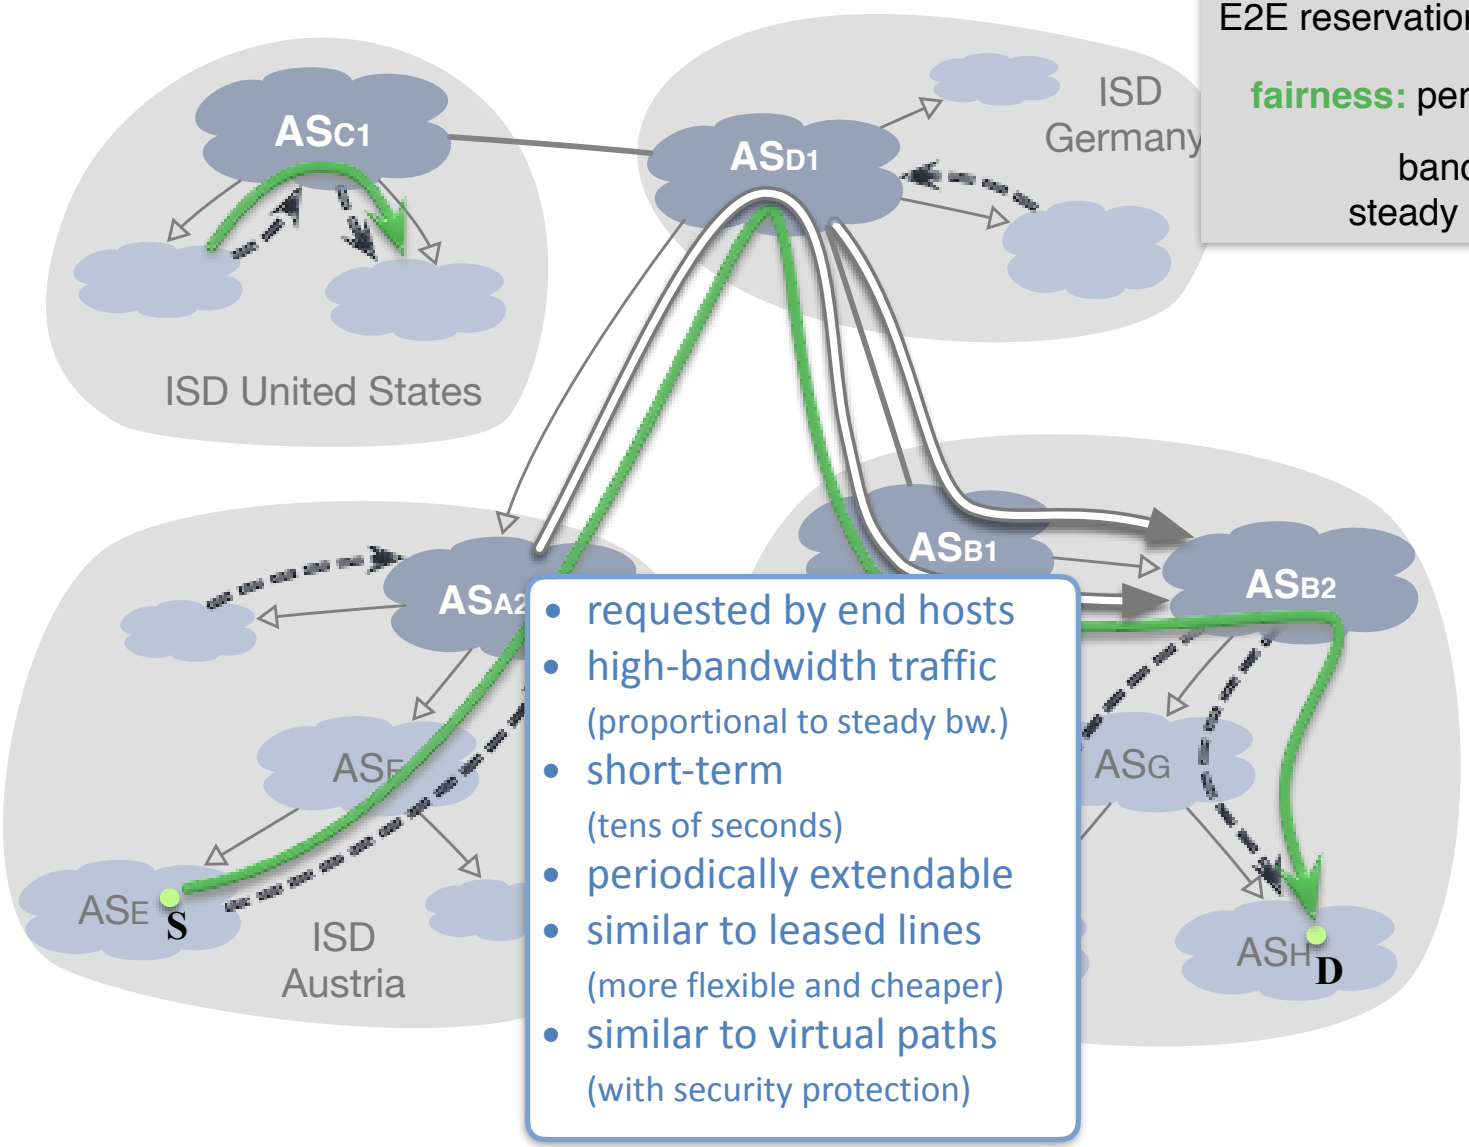

- requested by end hosts
- high-bandwidth traffic (proportional to steady bw.)
- short-term (tens of seconds)
- periodically extendable
- similar to leased lines (more flexible and cheaper)
- similar to virtual paths (with security protection)

How much bandwidth do ephemeral paths obtain?

2-Dimensional Bandwidth Decomposition

1. **vertical**
(hierarchical, per-location)

2. **horizontal**
(per-link)

2-Dimensional Bandwidth Decomposition

1. vertical
(hierarchical, per-location)

2. horizontal
(per-link)

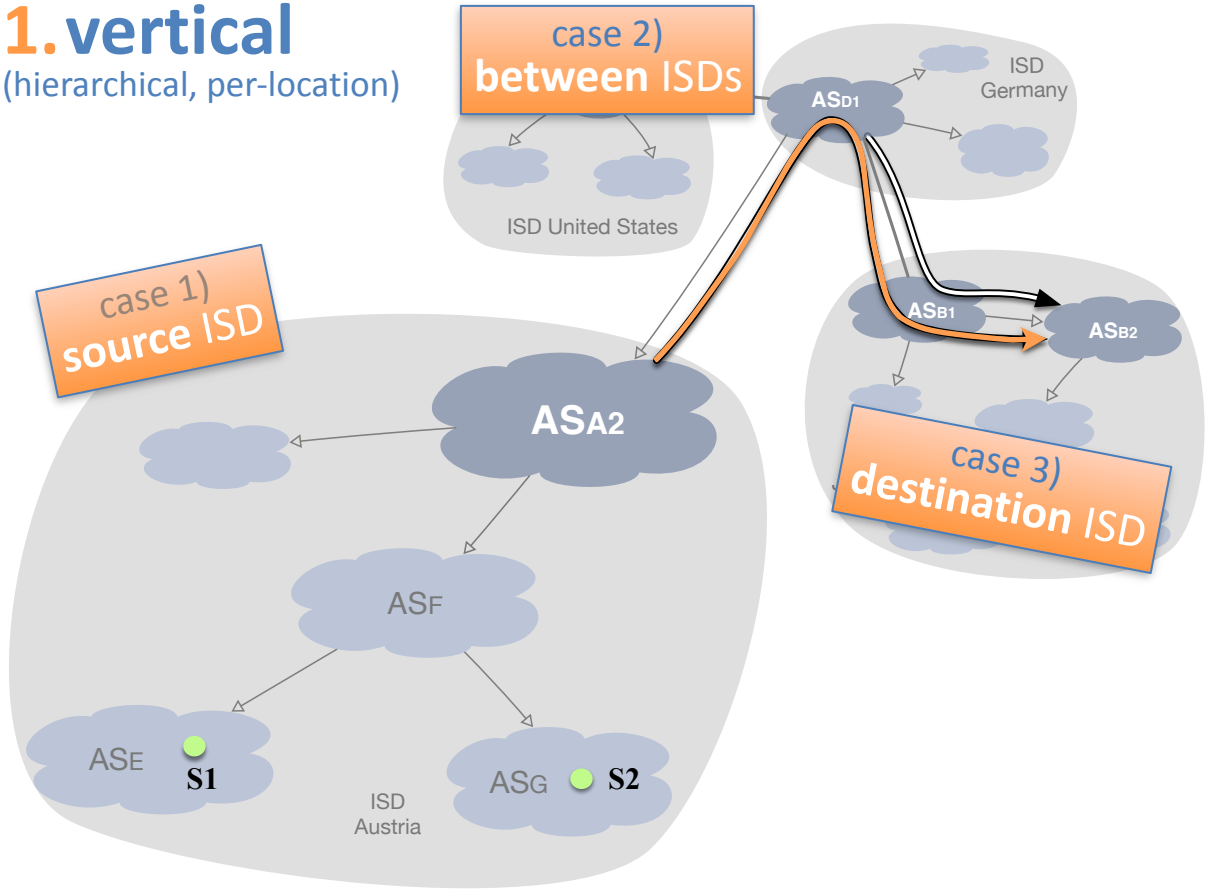

2-Dimensional Bandwidth Decomposition

1. vertical
(hierarchical, per-location)

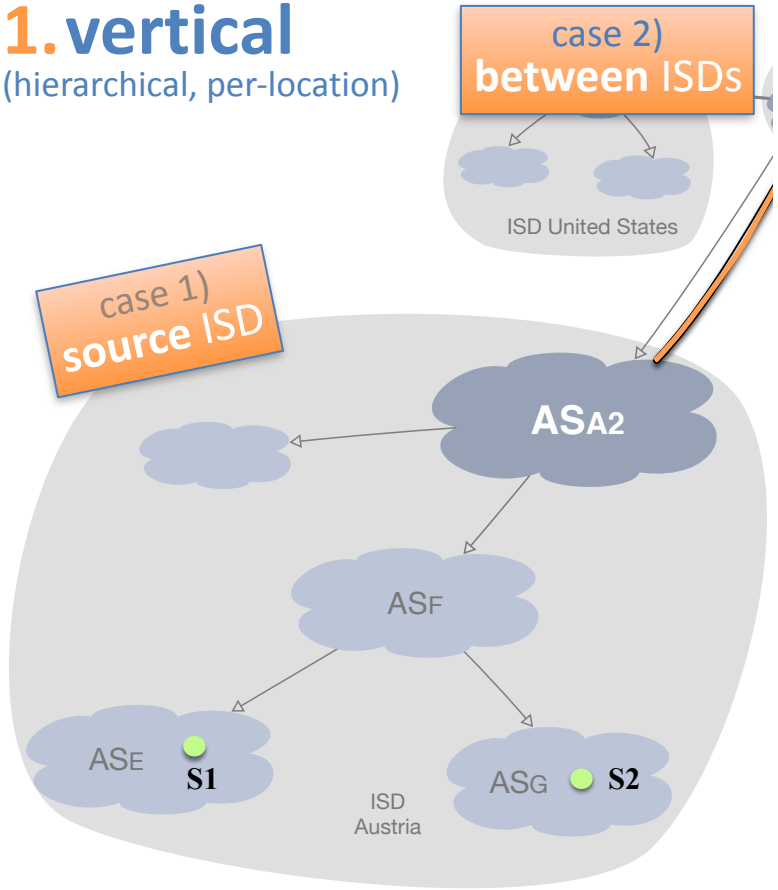

2. horizontal
(per-link)

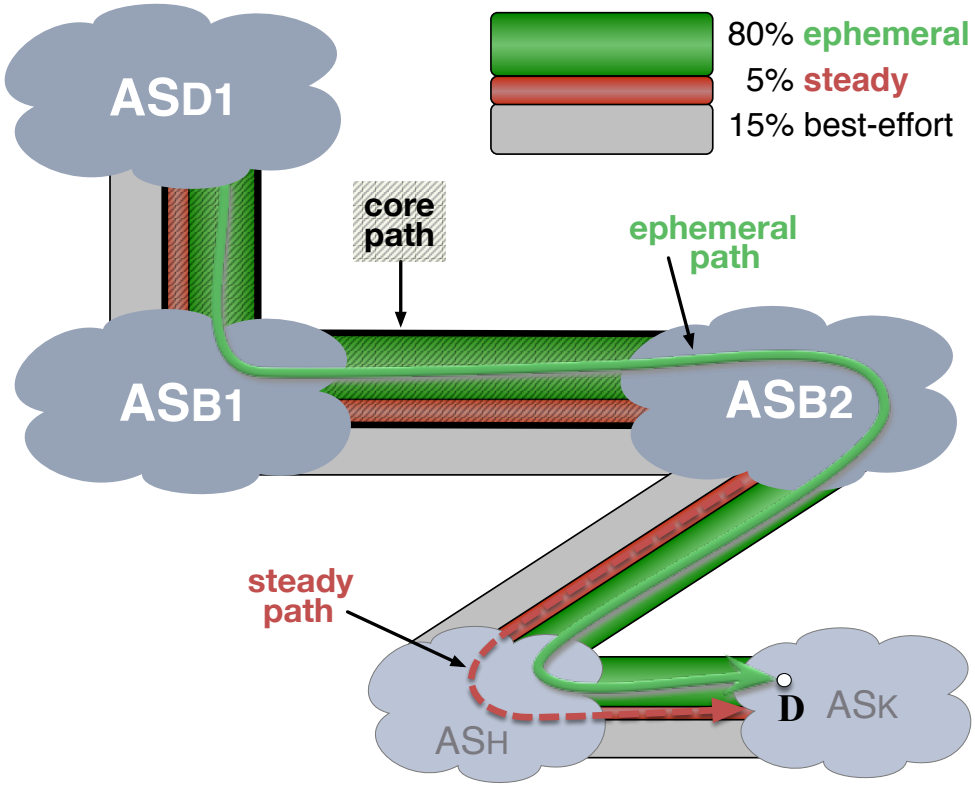

2-Dimensional Bandwidth Decomposition

1. vertical
(hierarchical, per-location)

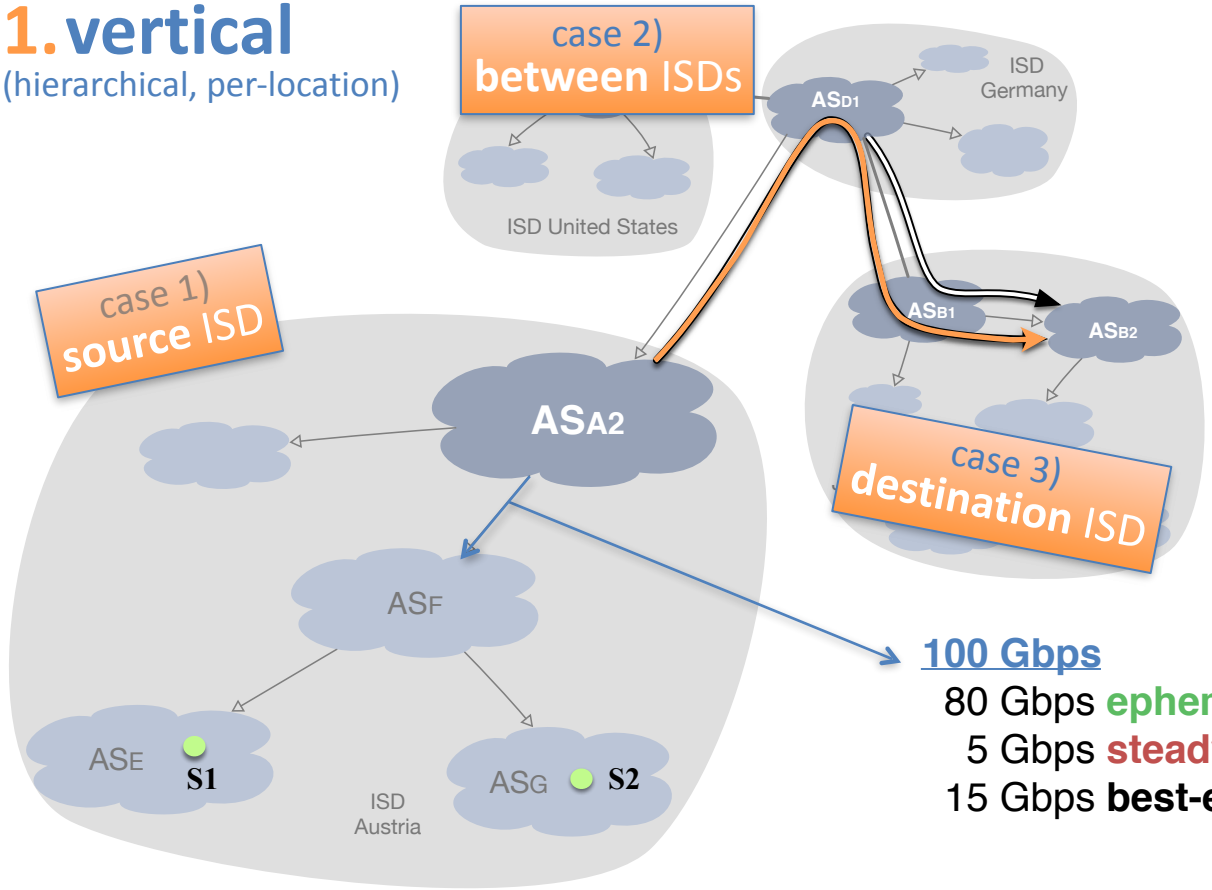

100 Gbps
80 Gbps ephemeral
5 Gbps steady
15 Gbps **best-effort**

2. horizontal
(per-link)

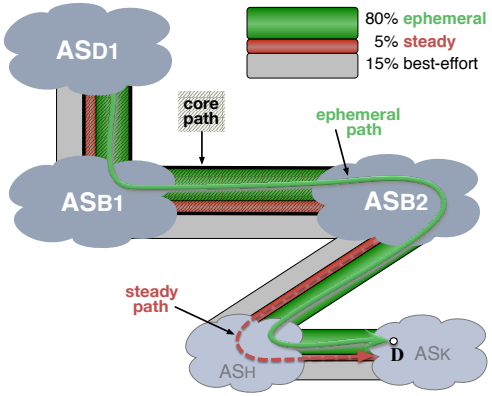

2-Dimensional Bandwidth Decomposition

1. vertical
(hierarchical, per-location)

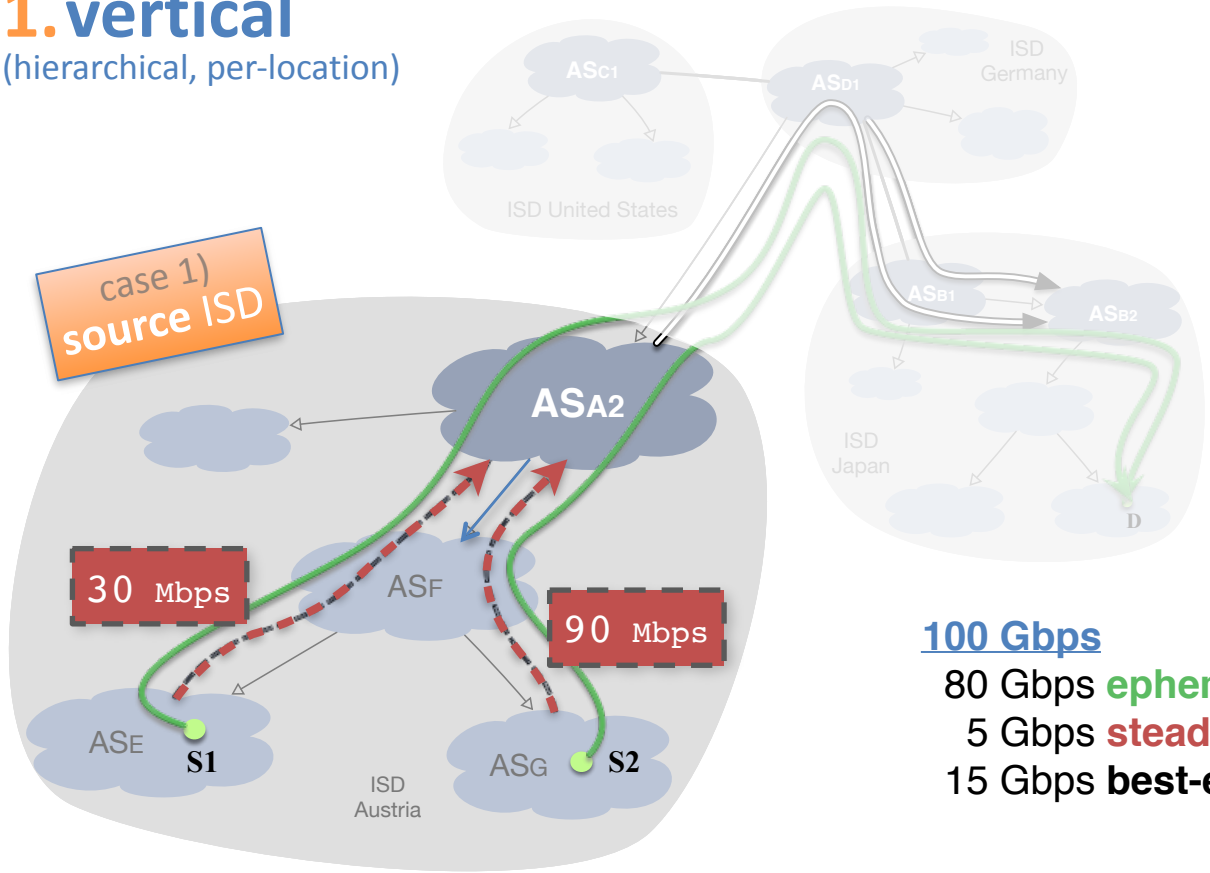

2. horizontal
(per-link)

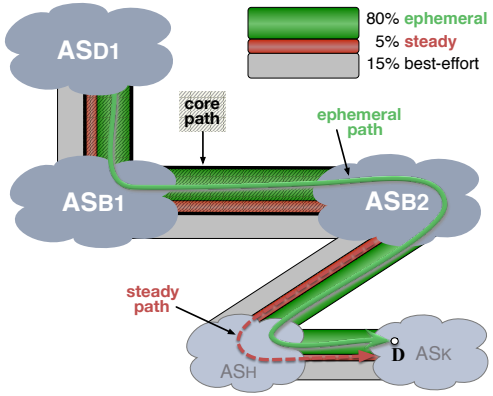

100 Gbps
80 Gbps **ephemeral**
5 Gbps **steady**
15 Gbps **best-effort**

2-Dimensional Bandwidth Decomposition

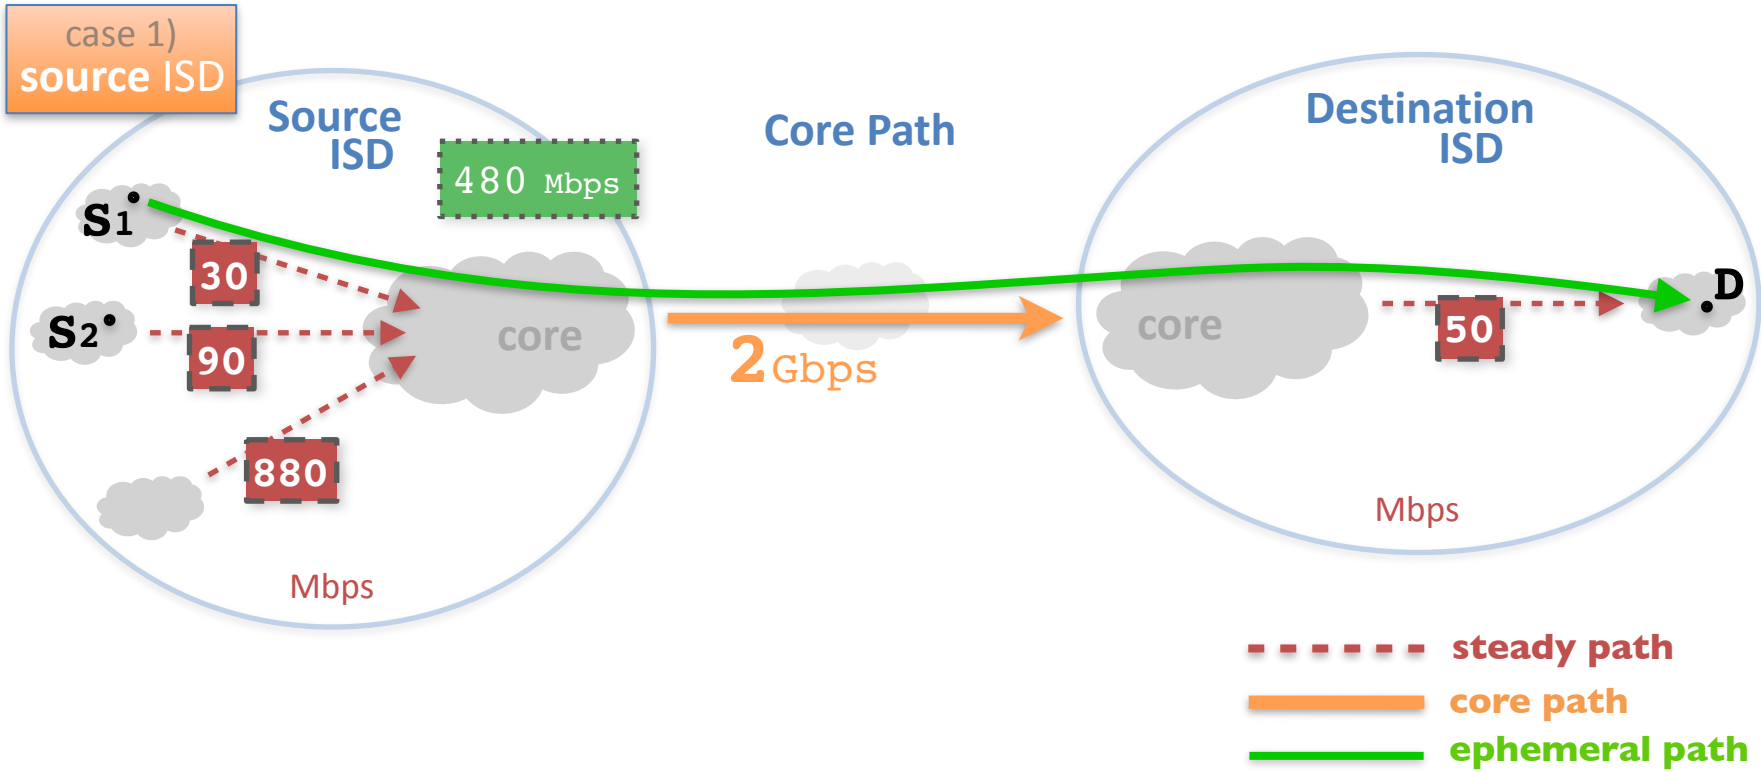

30 Mbps * 80 / 5 = 480 Mbps

2-Dimensional Bandwidth Decomposition

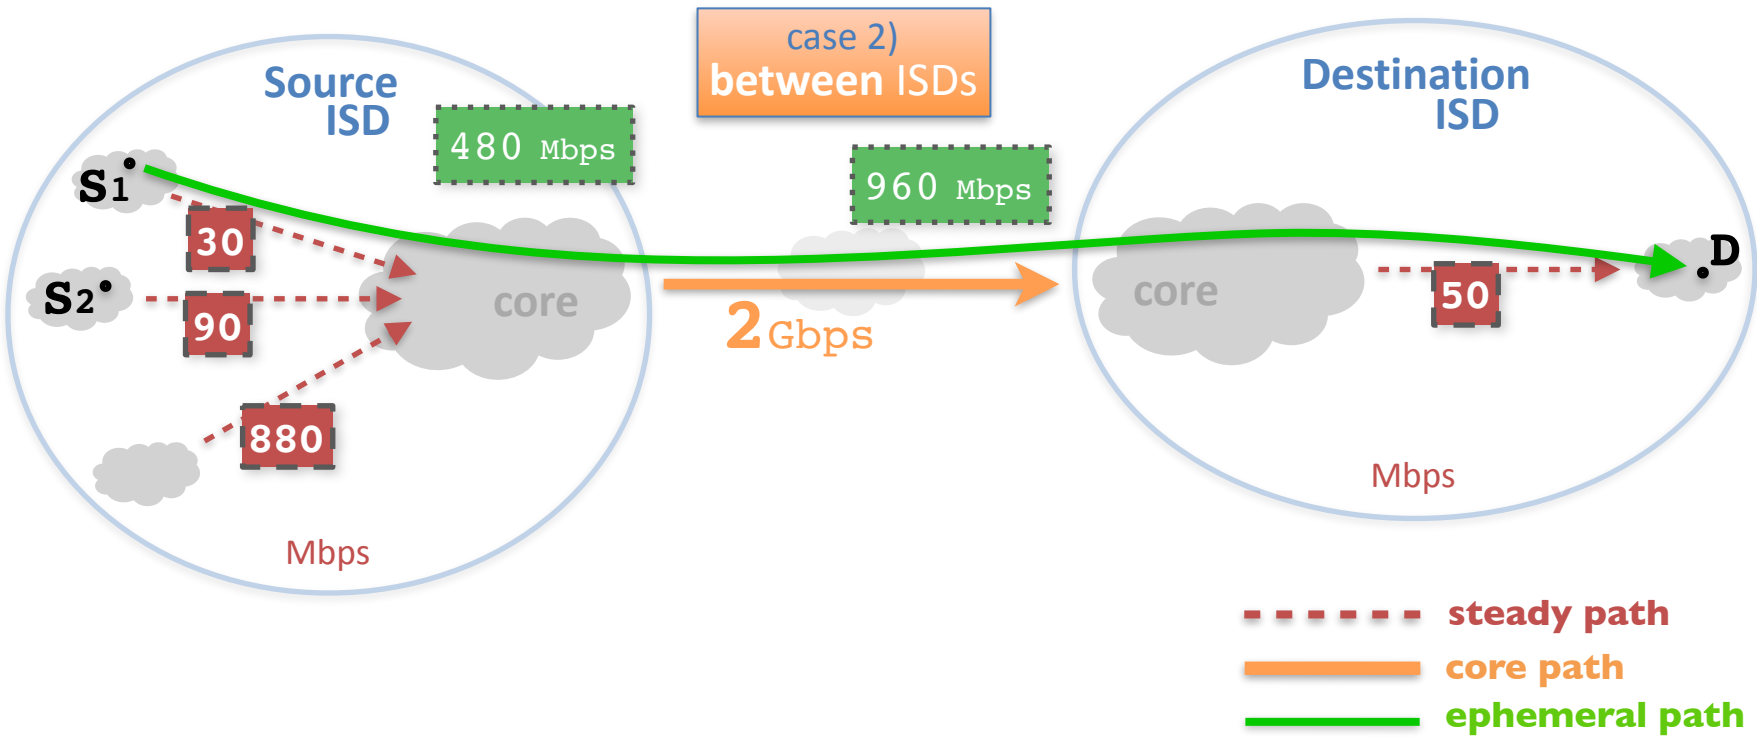

30 Mbps * 80 / 5 = 480 Mbps

30 / (30+90+880) * 2 Gbps * 80 / 5 = 960 Mbps

2-Dimensional Bandwidth Decomposition

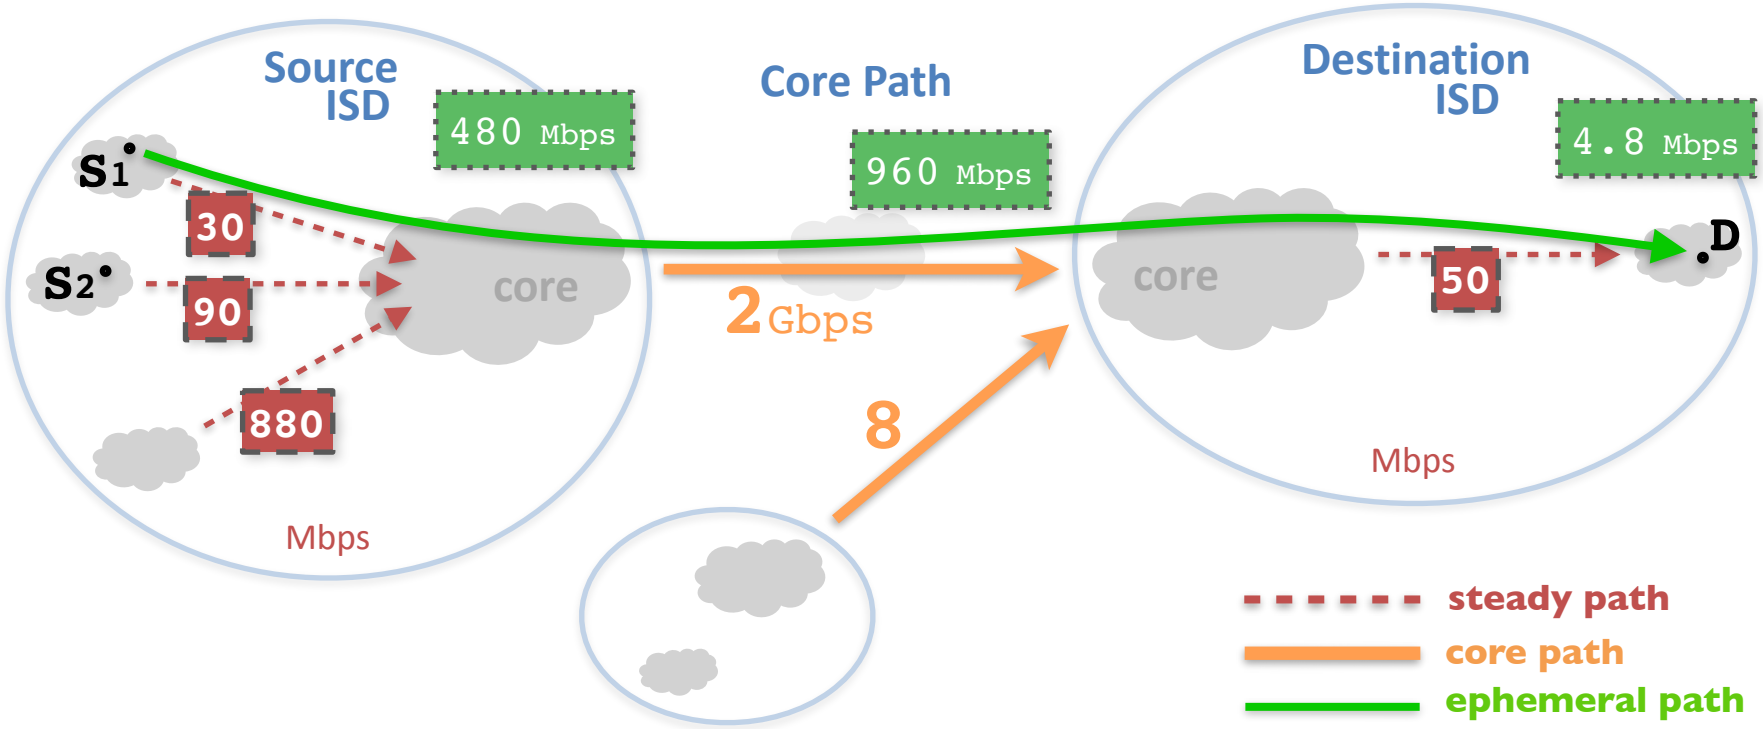


```
30 Mbps                                                  *   80 / 5   =   480 Mbps
30 / (30+90+880)     *    2 Gbps                         *   80 / 5   =   960 Mbps
30 / (30+90+880)     *    2 / (2+8)     *    50 Mbps     *   80 / 5   =   4.8 Mbps
```

2-Dimensional Bandwidth Decomposition

1. vertical
(hierarchical, per-location)

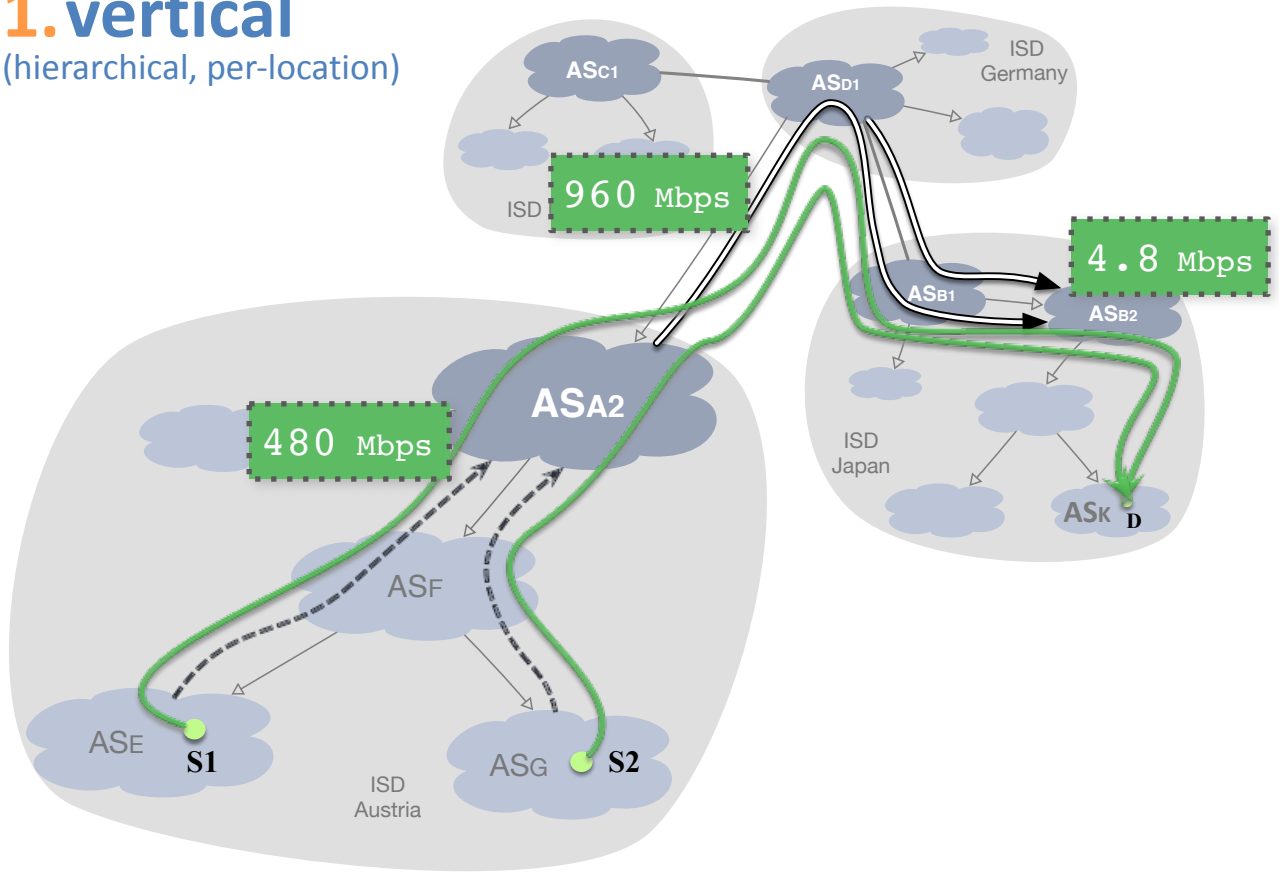

2. horizontal
(per-link)

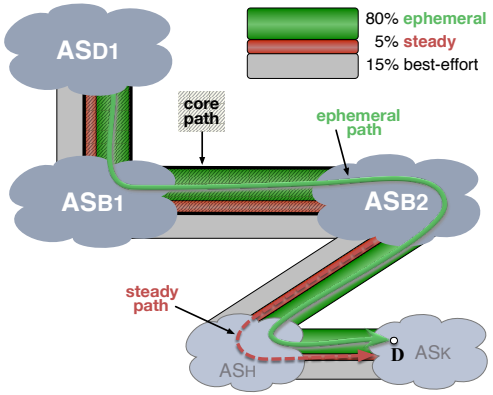

2-Dimensional Bandwidth Decomposition

1. vertical
(hierarchical, per-location)

2. horizontal
(per-link)

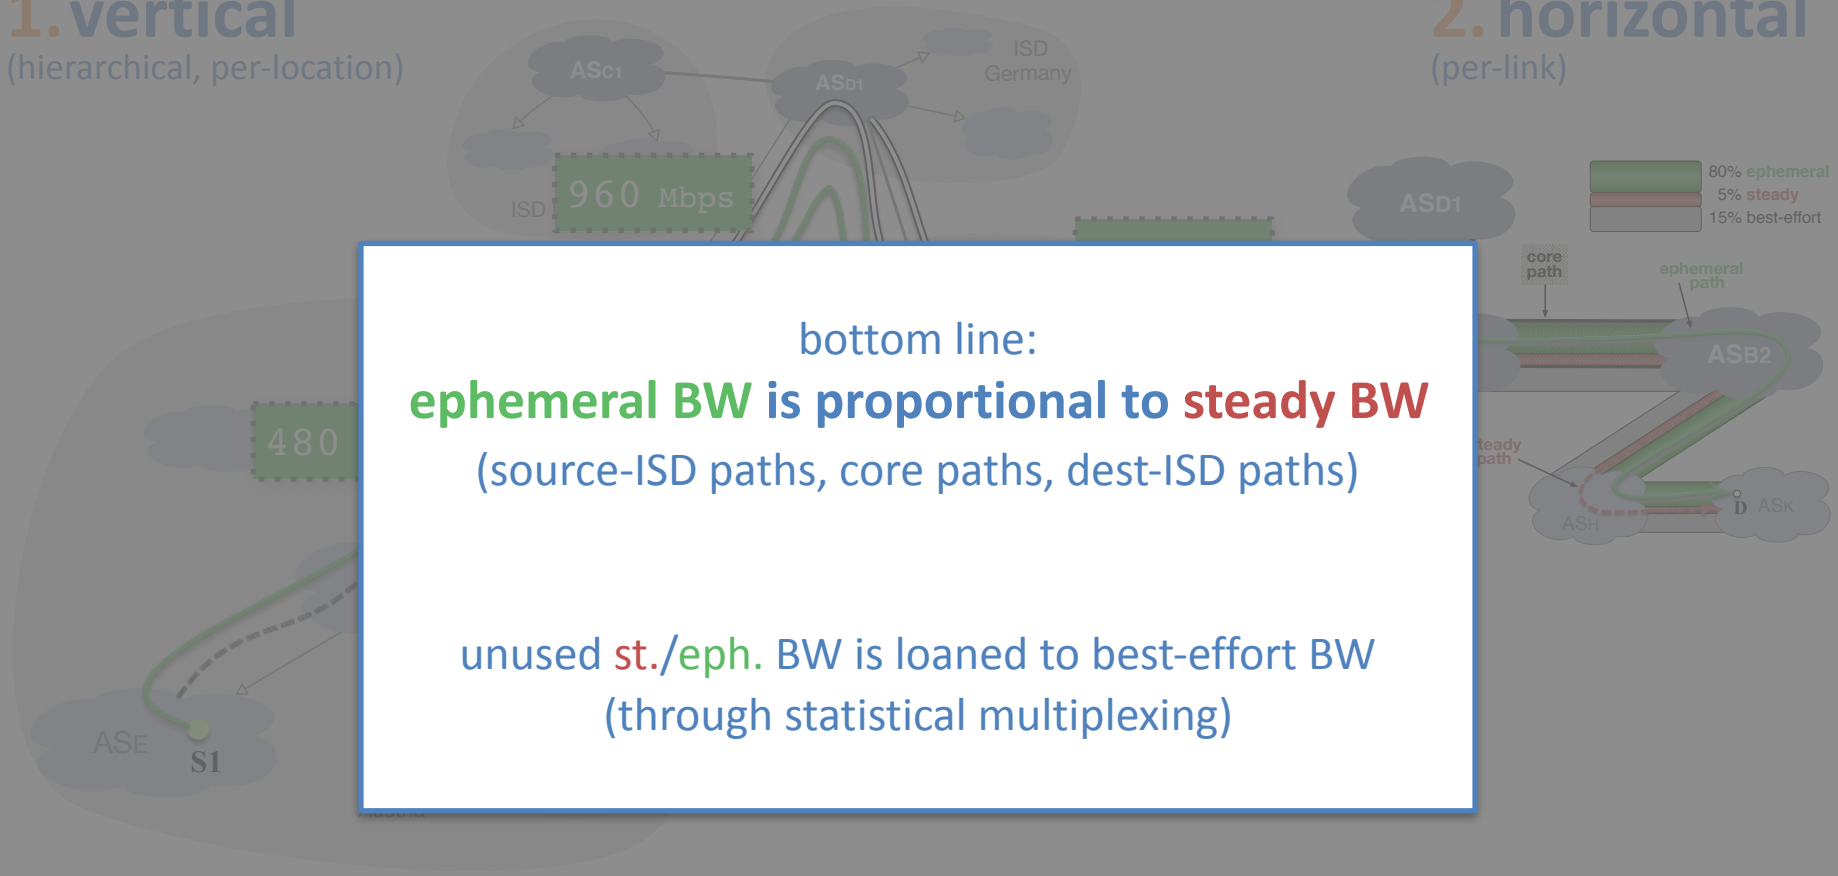

bottom line:

ephemeral BW is proportional to steady BW

(source-ISD paths, core paths, dest-ISD paths)

unused st./eph. BW is loaned to best-effort BW
(through statistical multiplexing)

SIBRA Guarantees

CBC-MAC (AES)
Intel's **AESni** [16]
4.15 cycles/byte

- Source AS S initiates a reservation. Each AS on path accepts or declines and provides **a cryptographic token:**

$$RT_{AS_i} = ingress_{AS_i} \parallel egress_{AS_i} \parallel MAC_{K_i}\left(ingress_{AS_i} \parallel egress_{AS_i} \parallel Request \parallel RT_{AS_{i-1}}\right)$$

- **Efficiency & Scalability**:
  ASes verify these **tokens**, embedded in the forwarded packets, i.e., no per-flow state.

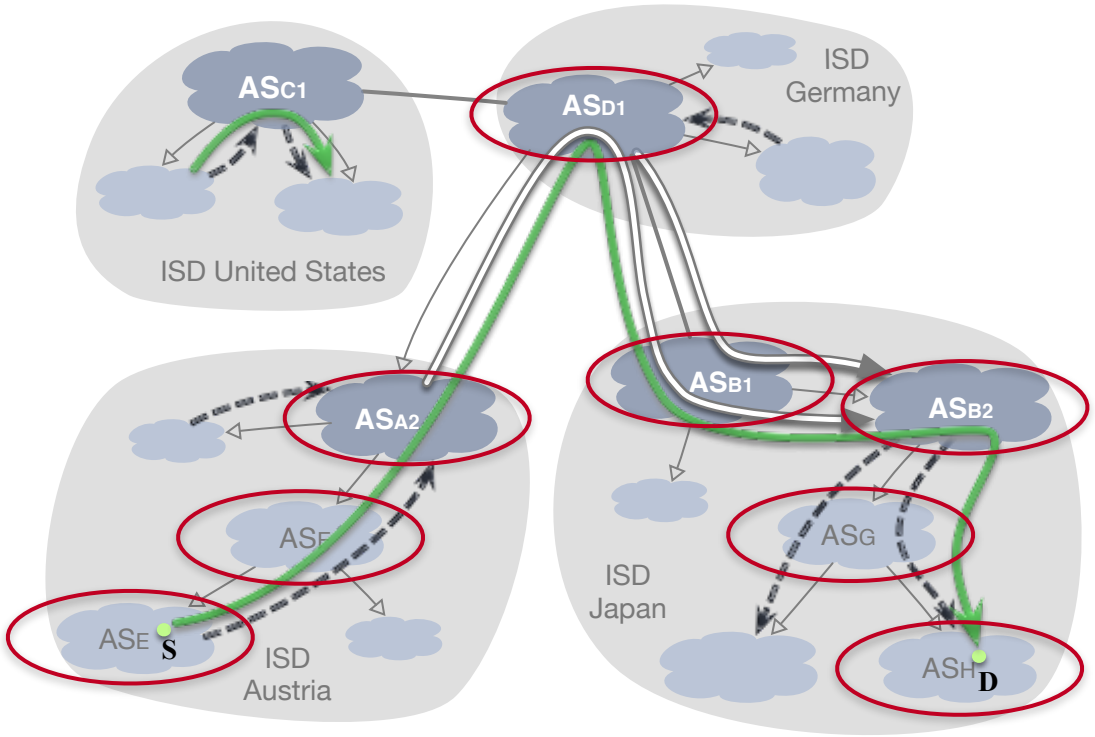

SIBRA under Attack

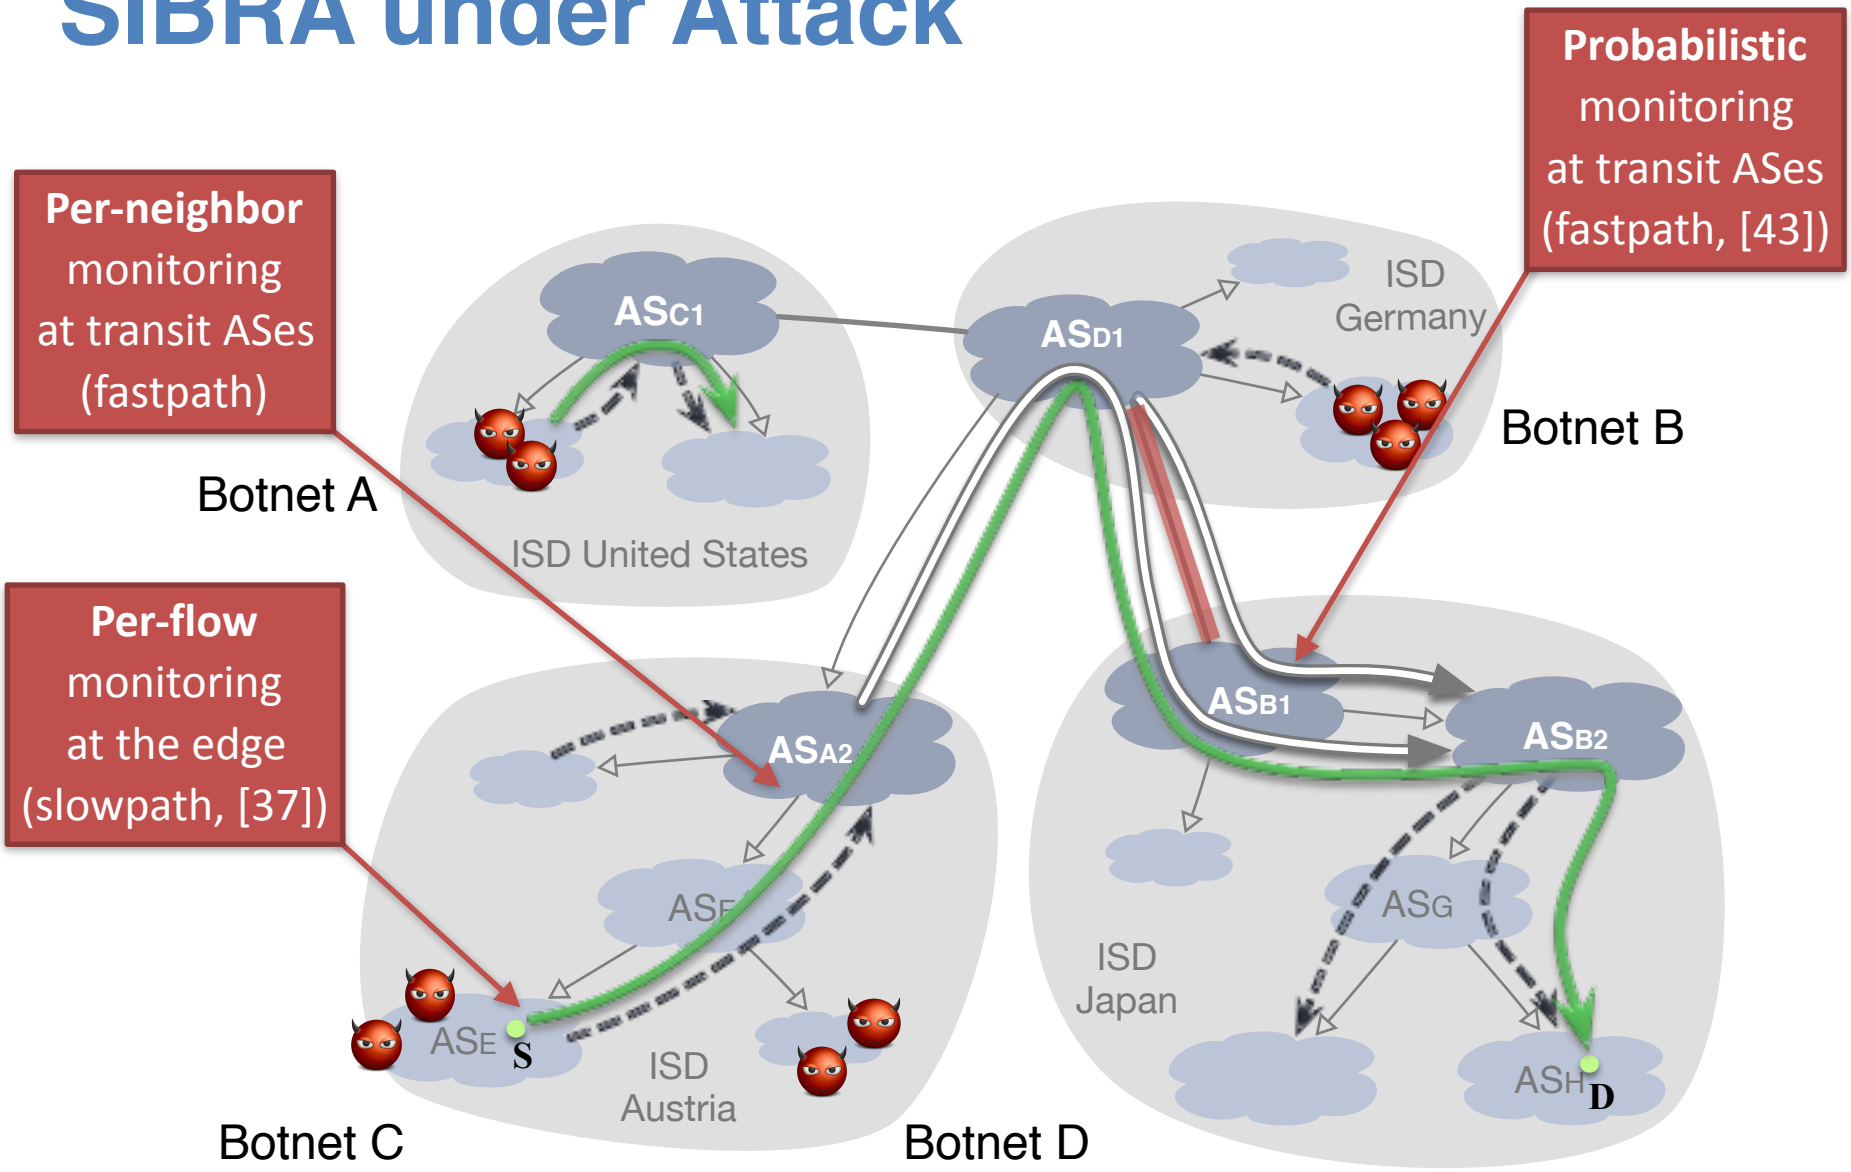

Is there enough bandwidth in today’s Internet?

Case study: core links to Australia

- The entire world connects to Australia (32 428 leaf ASes)

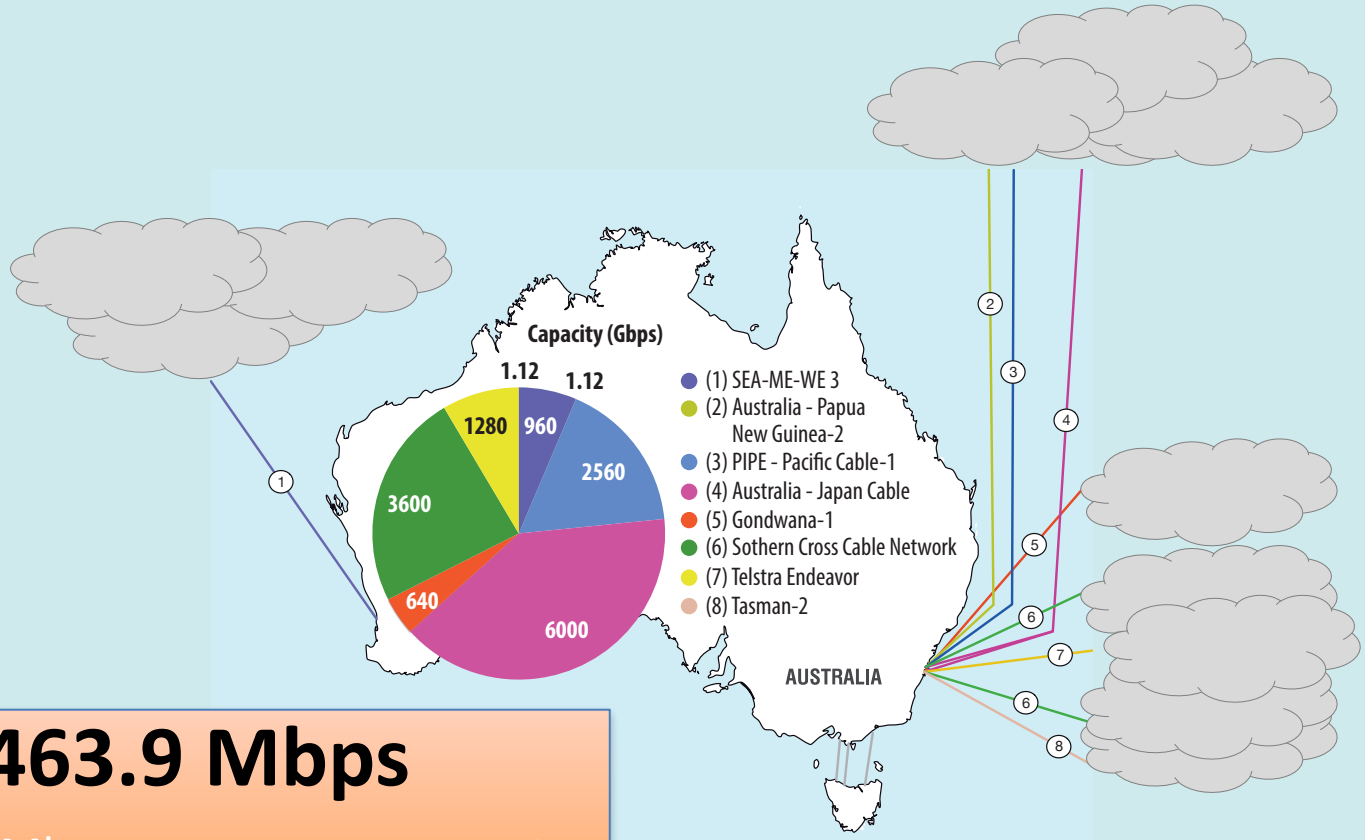

463.9 Mbps

(371.1 Mbps ephemeral bandwidth)

for each AS

5.64 Gbps

in 2018

How effective is SIBRA?

Evaluation: Defense against Coremelt

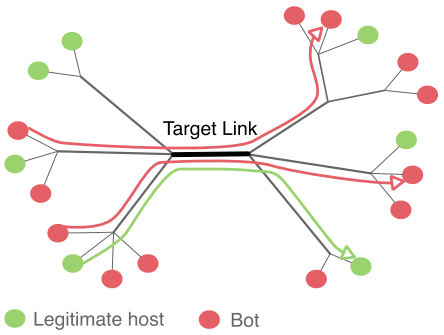

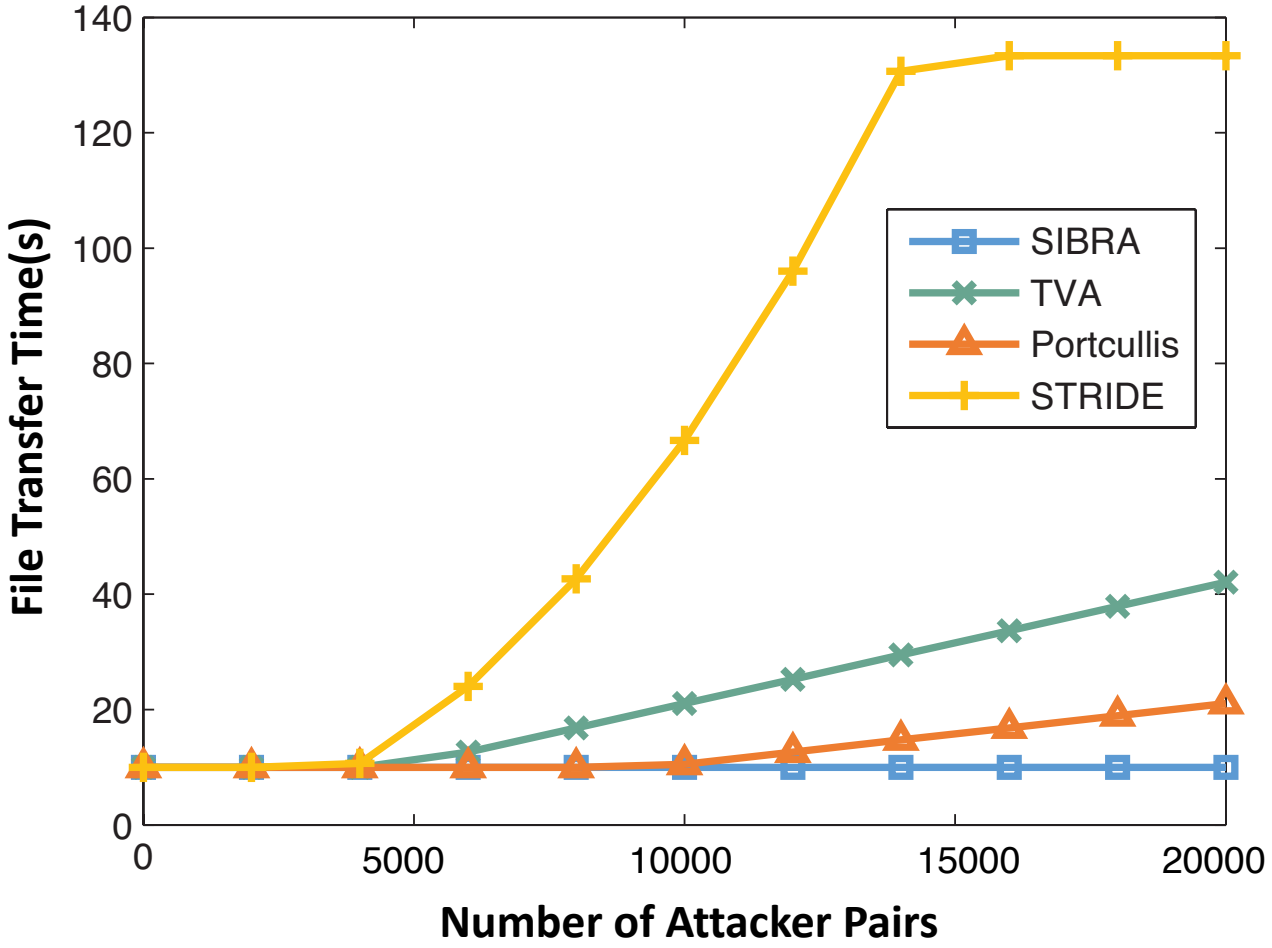

How efficient is SIBRA?

Per-flow Stateless Operations

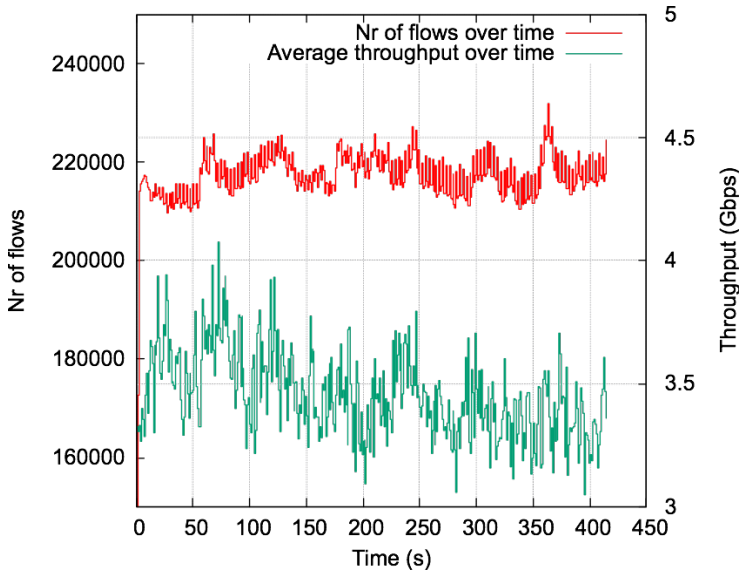

10 Gbps core link (load ~40%): 2.2×10^5 flows per second
1 Tbps core link (load ~40%): 2.2×10^7 flows per second

Storing **per-flow state** is **prohibitively expensive** — especially under attack

Router Action	Time (avg)	Per second
Processing 1 reservation request	9.10 μs	110 K
Processing 1 packet (1 500 bytes) using Intel's DPDK and AESni	0.04 μs	25 Mio

280 Gbps

Conclusions

- **Botnet-size independence** is the key property against DDoS attacks
- SIBRA is the **first bandwidth reservation architecture** to achieve botnet-size independence at Internet scale
- Two-dimensional **bandwidth decomposition**
- Very **fast operations**, per-flow stateless forwarding

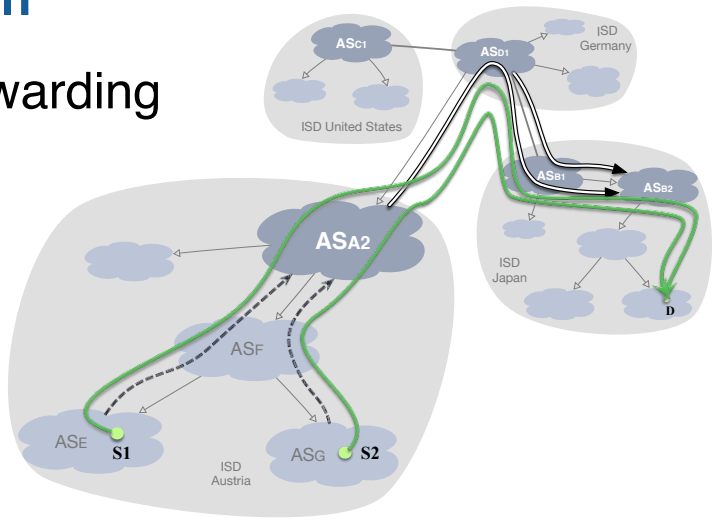

Related Work

[9] D. Barrera, R. M. Reischuk, P. Szalachowski, and A. Perrig, "*SCION five years later: Revisiting scalability, control, and isolation on next-generation networks*," arXiv, 2015.

[16] S. Gueron, "*Intel Advanced Encryption Standard (AES) New Instructions Set*," Intel, 2010, white paper 323641-001, Revision 3.

[32] B. Parno, D. Wendlandt, E. Shi, A. Perrig, B. Maggs, and Y.-C. Hu, "*Portcullis: Protecting Connection Setup from Denial-of-Capability Attacks*," in ACM SIGCOMM, 2007.

[37] I. Stoica, S. Shenker, and H. Zhang. *Core-Stateless Fair Queueing: A Scalable Architecture to Approximate Fair Bandwidth Allocations in High-Speed Networks*. IEEE/ ACM Transactions on Networking, 2003.

[38] A. Studer and A. Perrig, "*The Coremelt attack*," in ESORICS, 2009.

[43] H. Wu, H.-C. Hsiao, and Y.-C. Hu. *Efficient large flow detection over arbitrary windows: An algorithm exact outside an ambiguity region*. In ACM IMC, 2014.

ETH zürich

SCION
Internet Architecture

Backup

Parameter Choice: Traffic Types

- **ephemeral (80%)**
  - Netflix's video constitutes >50% of the entire Internet traffic
  - together with YT and FB, 70-90% are realistic for ephemeral traffic
- **steady (5%)**
  - based on a 10-day measurement of a tier-1 ISP: connection establishment (TCP-SYN) uses 0.5% of the bandwidth
  - SIBRA allocates 10x that amount
- **best-effort (15%)**
  - email, news, SSH, DNS (3.9%)
  - very short-lived flows, less than 256ms (5.6%)