Cooperative Network Systems for Protecting Internet Users

Sophia Yoo

Dissertation Defense

April 24, 2026

PRINCETON UNIVERSITY

We rely daily on the *Internet* for essential services

Communication

- Email messaging
- Voice/video conferencing
- Chat platforms
- Online forums
- File sharing

Consumer & Lifestyle

- Live streaming
- Real-time gaming
- Social networks
- News platforms
- Smart home ecosystems

Critical Infrastructure

- E-commerce
- Healthcare
- Government
- Education
- Transportation

We rely daily on the *Internet* for essential services

…but today's Internet is not up for the task!

Communication

- Email messaging
- Voice/video conferencing
- Chat platforms
- Online forums
- File sharing

Consumer & Lifestyle

- Live streaming
- Real-time gaming
- Social networks
- News platforms
- Smart home ecosystems

Critical Infrastructure

- E-commerce
- Healthcare
- Government
- Education
- Transportation

Security Issues Abound

Massive DDoS attacks disrupt the Internet

cybernews®

Home » Security

Steam, Riot Games hit by disruptions: massive DDoS attack suspected

Published: 7 October 2025 · Last updated: 7 October 2025

Ernestas Naprys, Senior Journalist

Multiplayer gamers on different platforms have experienced service outages and disruptions simultaneously. The cybersecurity community suspects a major distributed denial of service attack (DDoS) from Aisuru, a massive botnet pushing out record-breaking traffic.

Downdetector users are reporting widespread problems affecting Steam and Riot Games, two of the world's largest gaming platforms. Gamers complain about being unable to play major titles, such as Counter-Strike, Dota 2, Valorant, League of Legends, and many others.

CHRISPI
@cripto_lion · Follow

The internet is melting: Steam down. Epic Games down. AWS down. Cloudflare wobbling.

f you're refreshing, you're not alone.

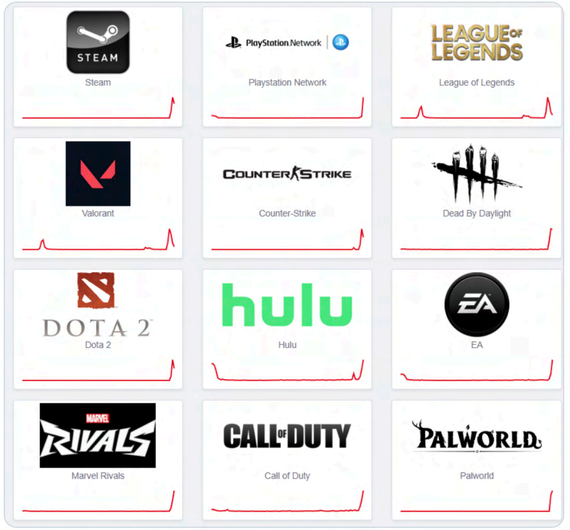

9:48 PM · Oct 6, 2025

298 Reply Copy link

Read 14 replies

*DDoS: Distributed Denial-of-Service

Security is not guaranteed

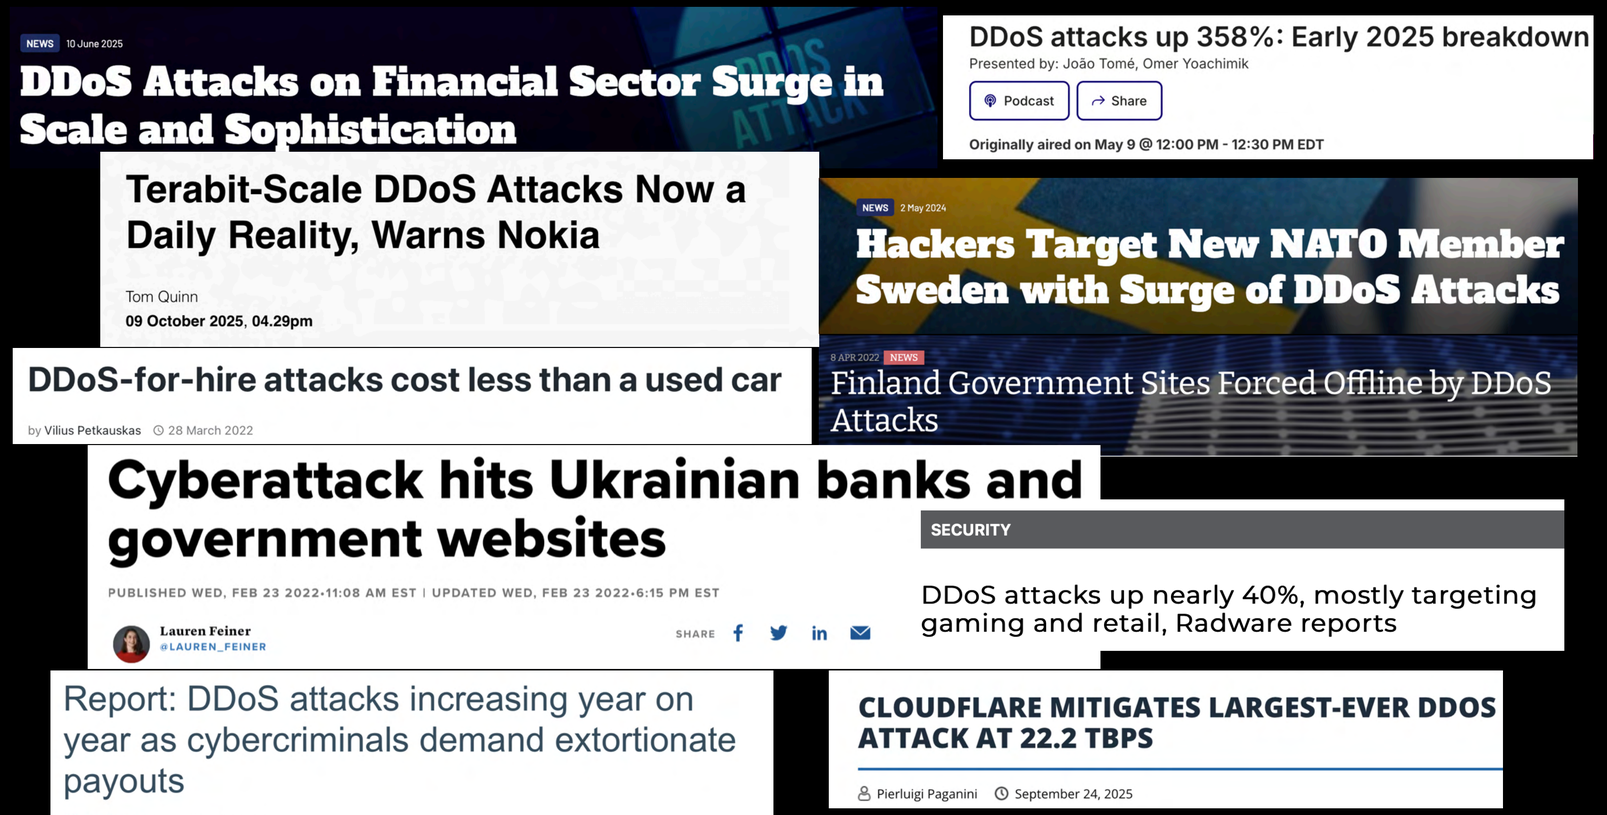

Users:
Critical service disruption in e-commerce, online banking, streaming/media, gaming, communication networks…

Providers:
Huge financial losses, data breach risk, brand/reputation damage, regulatory fines, cyber warfare escalation…

The Internet has a “plug-and-play” model

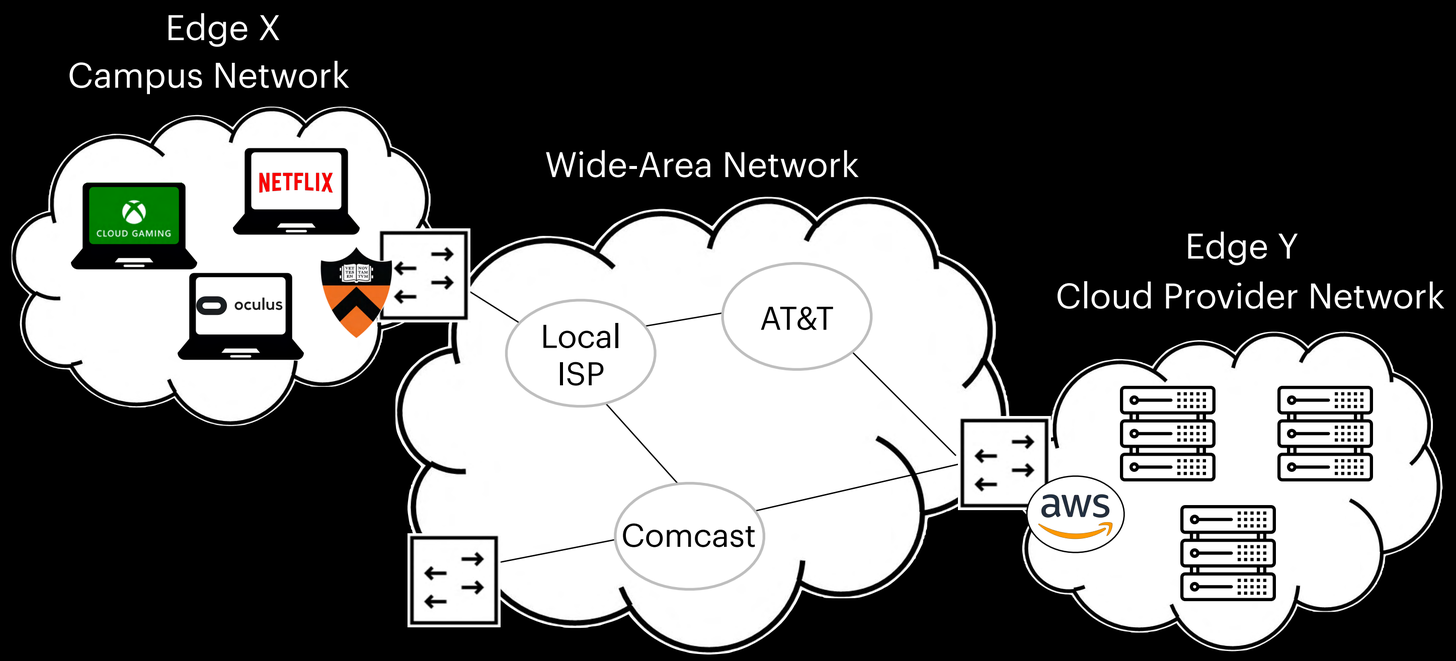

The Internet has a “plug-and-play” model

connect to one network and send traffic to anyone globally

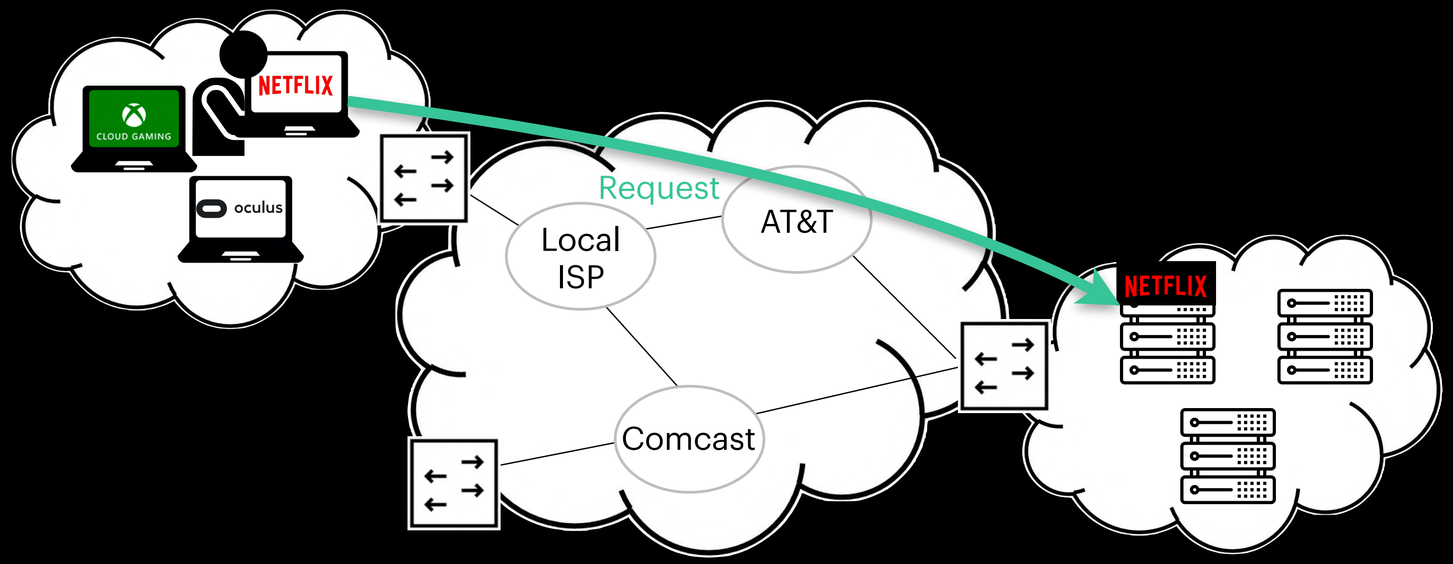

The Internet has a “plug-and-play” model

connect to one network and send traffic to anyone globally | *without dealing with underlying mechanisms (routing, etc.)*

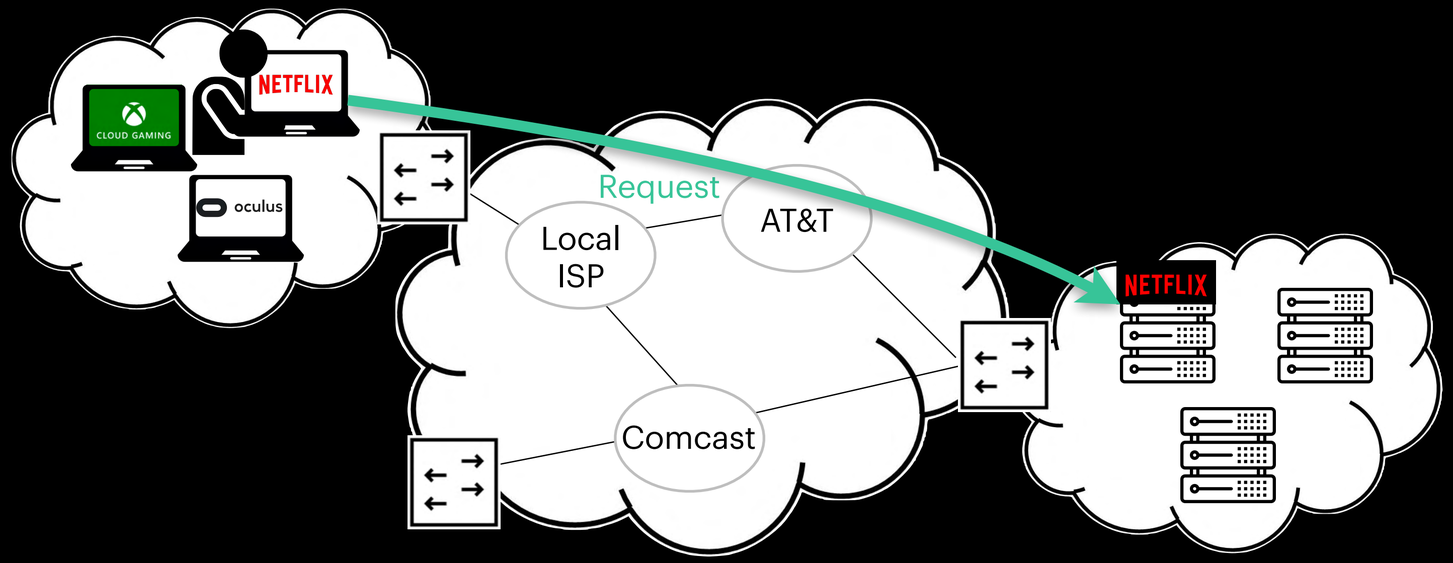

The Internet does not natively provide

The Internet does not natively provide

Ingress Control

you cannot prevent others from sending you traffic…enabling availability attacks

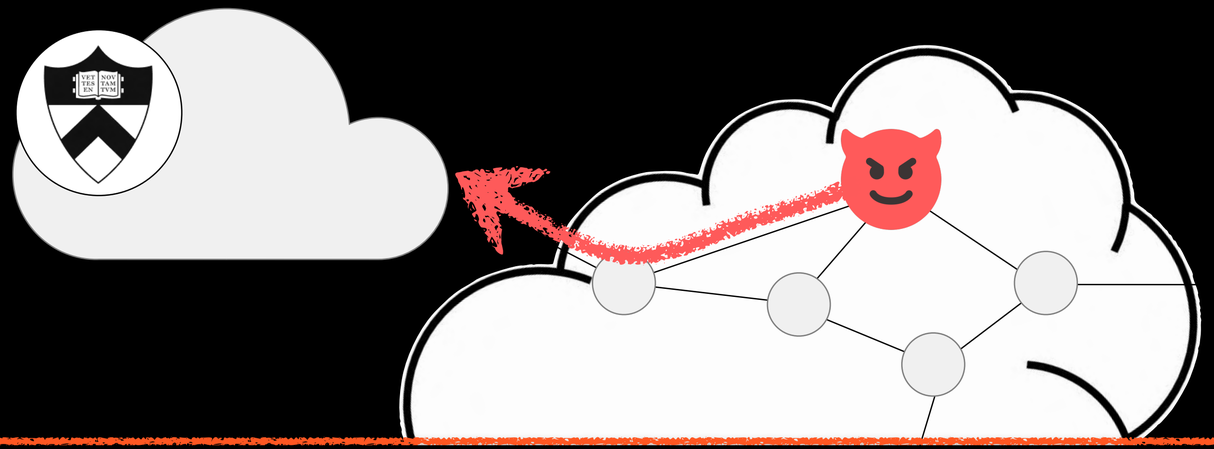

The Internet does not natively provide

Ingress Control

you cannot prevent others from sending you traffic…enabling availability attacks

Route Control

you cannot choose what path your own traffic takes…enabling privacy + performance attacks

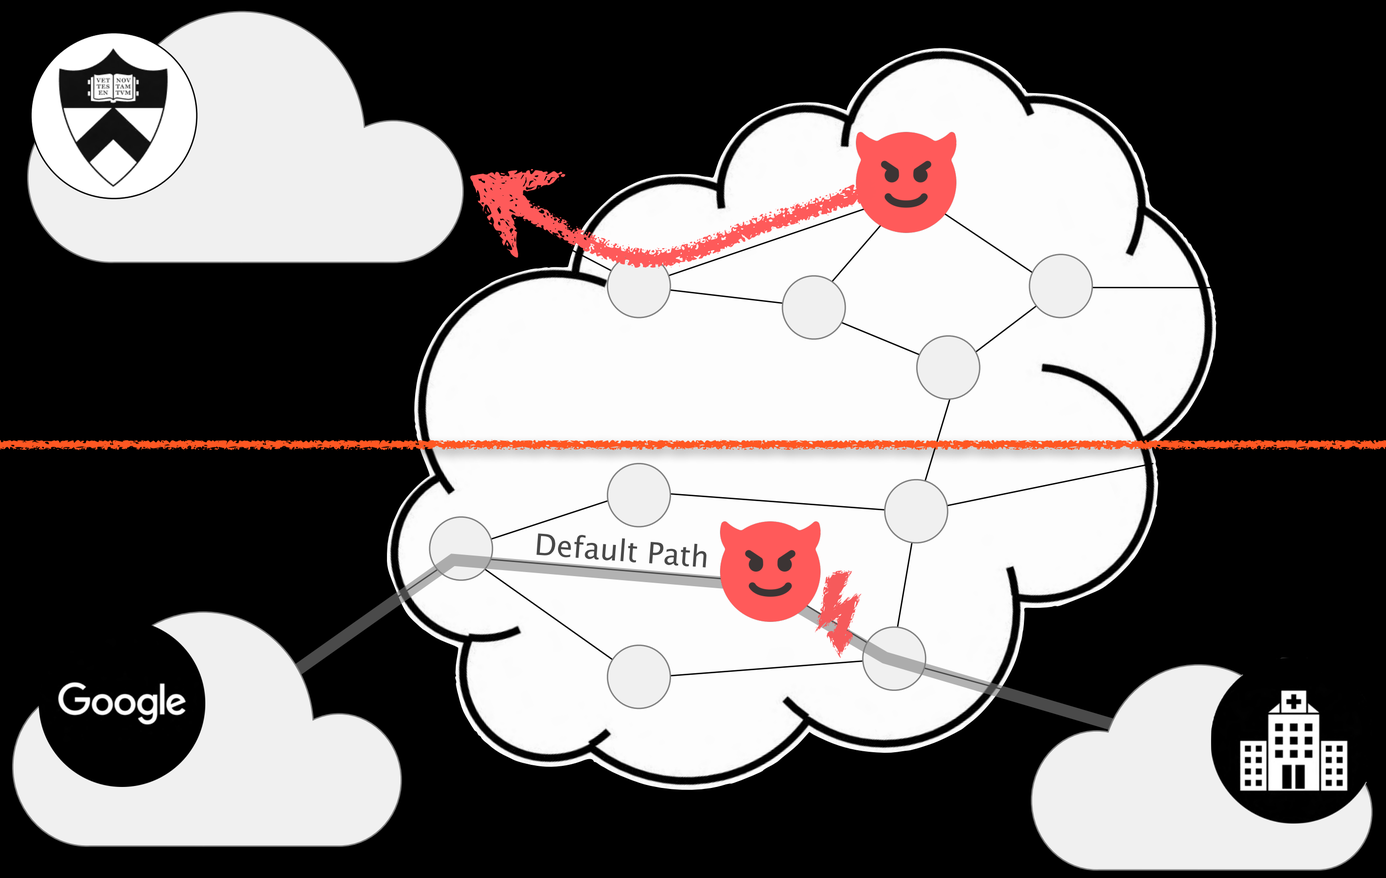

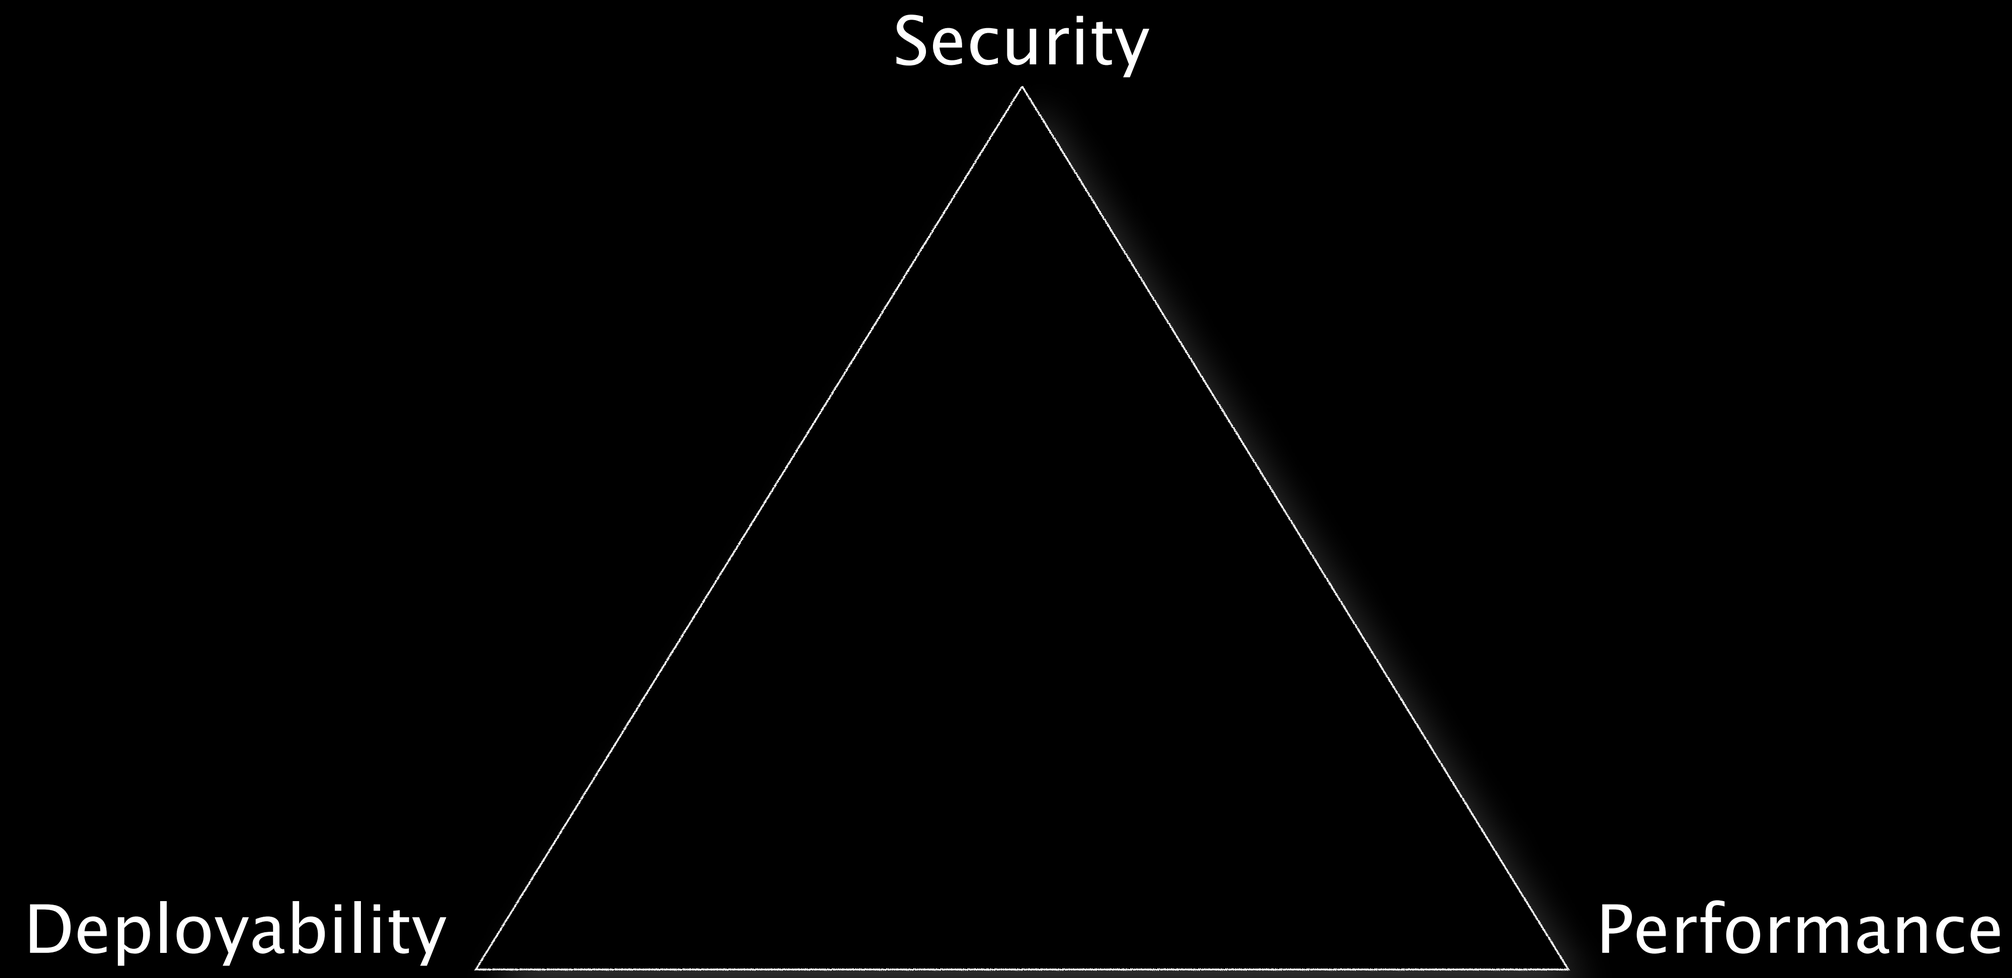
Security
Deployability
Performance

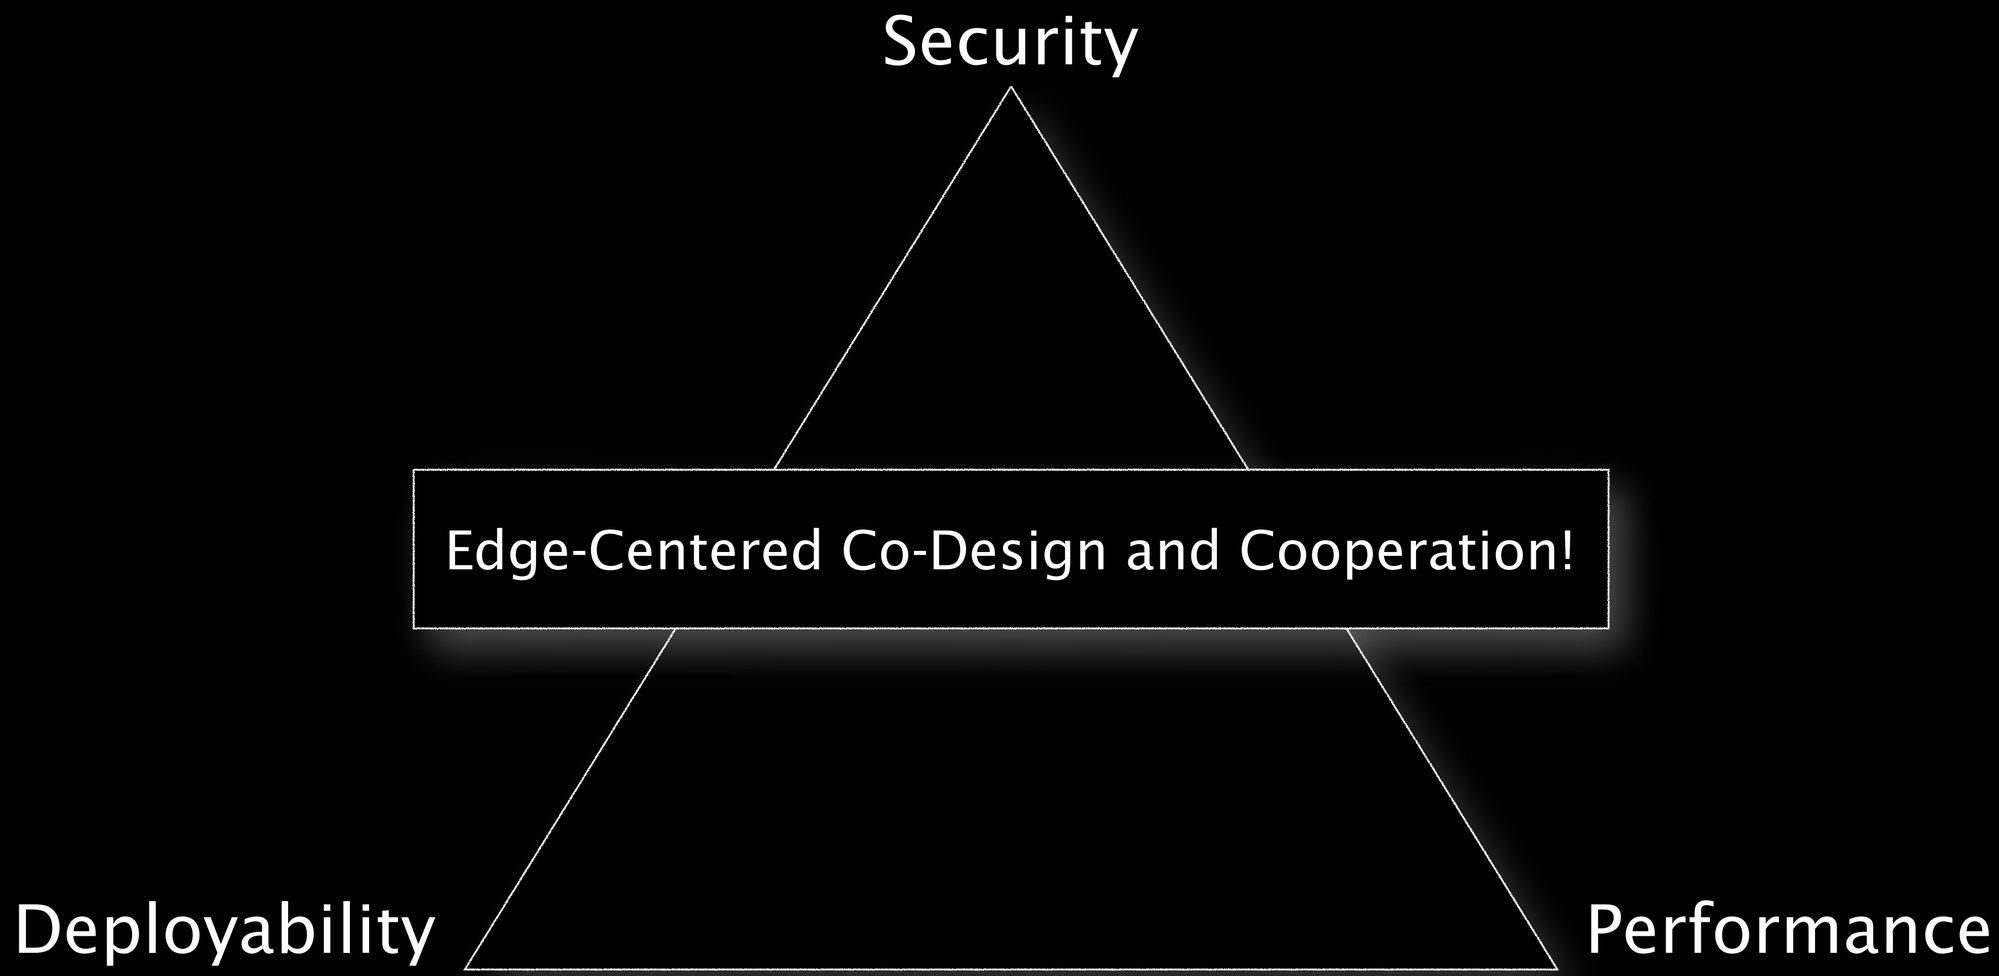
Security
Edge-Centered Co-Design and Cooperation!
Deployability
Performance

Dissertation Outline

Chapter 1: Introduction

Part I: Ingress Control via Intra-Edge Co-Design

- Chapter 2: SmartCookie (Tiered Traffic Filtering)

Part II: Route Control via Inter-Edge Cooperation

- Chapter 3: Tango (Path Exposure)
- Chapter 4: PraxiGuard (Path Exploitation)

Chapter 5: Conclusion

SmartCookie

[USENIX Sec ‘24]

Part I: Ingress Control via Intra-Edge Co-Design

Tango + PraxiGuard

[NSDI ‘24] [in preparation]

Part II: Route Control via Inter-Edge Cooperation

SmartCookie

[USENIX Sec '24]

Part I: Ingress Control via Intra-Edge Co-Design

Tango + PraxiGuard

[NSDI '24] [in preparation]

Part II: Route Control via Inter-Edge Cooperation

In order for clients to start downloading content

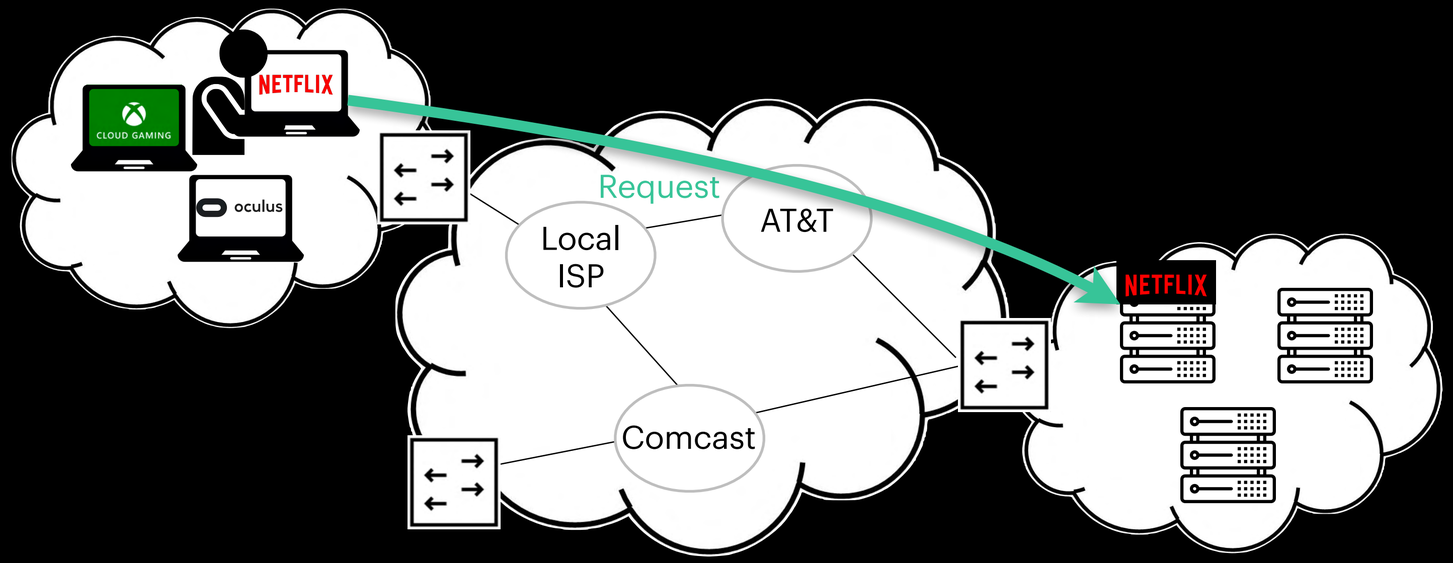

In order for clients to start downloading content

connections must first be established with a “3-way-handshake”

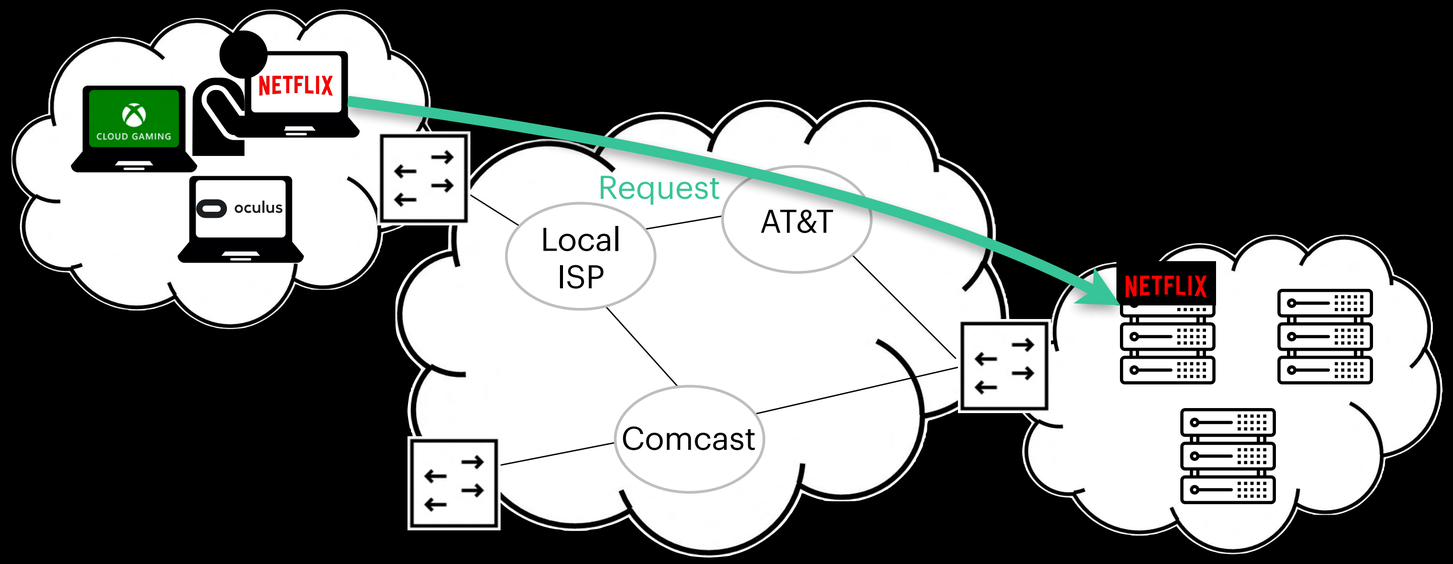

3-Way Handshake (Protocol)

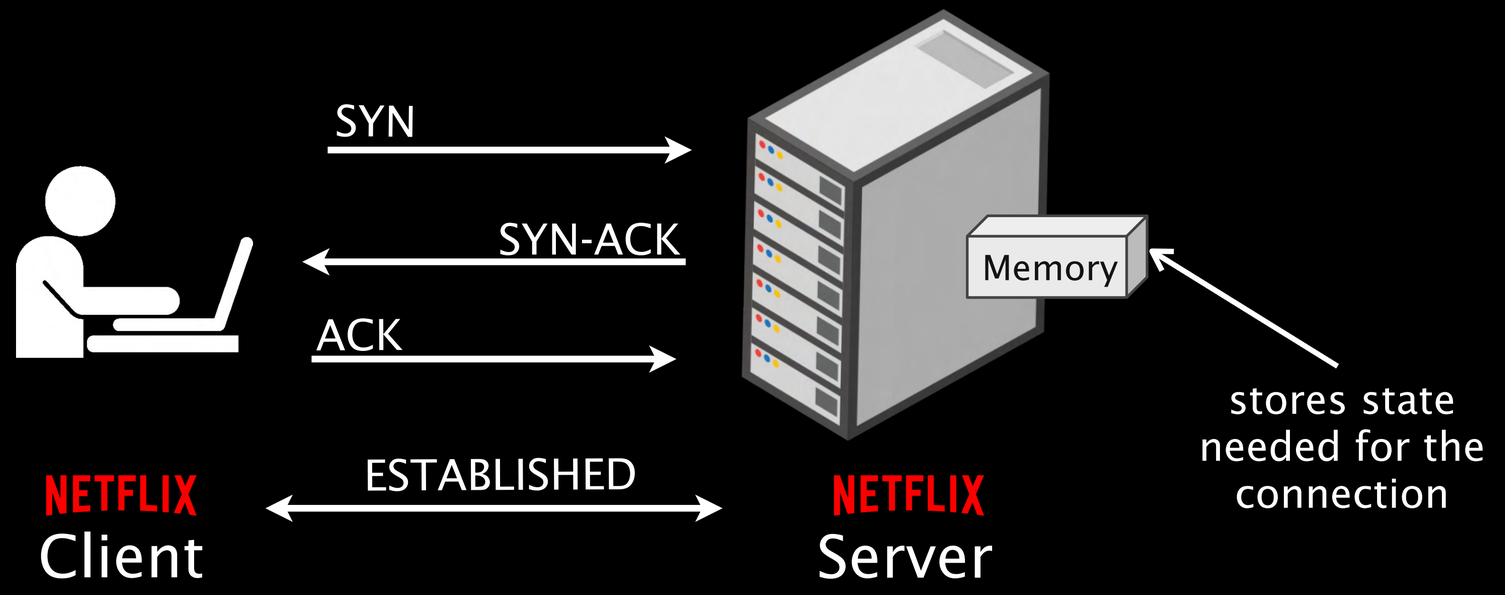

SYN: “Synchronize”
ACK: “Acknowledge”

SYN Flooding: an asymmetric attack on resources

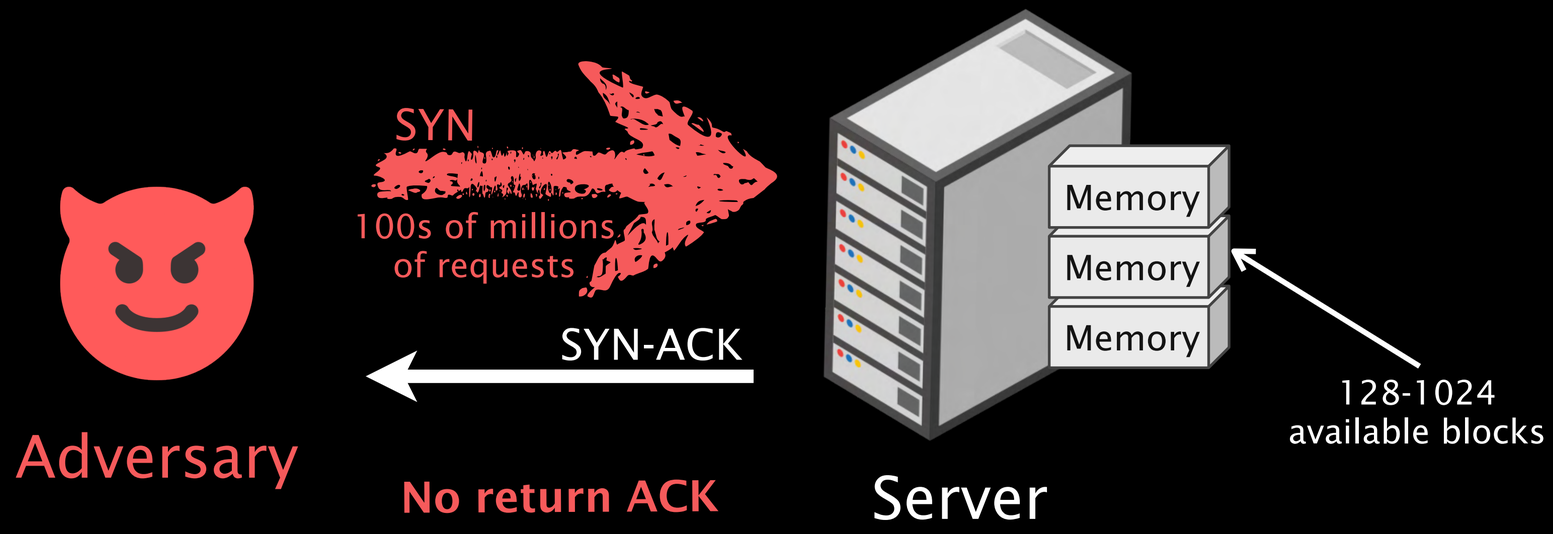

SYN Flooding: an asymmetric attack on resources

Server memory is depleted, leading to DoS for new requests

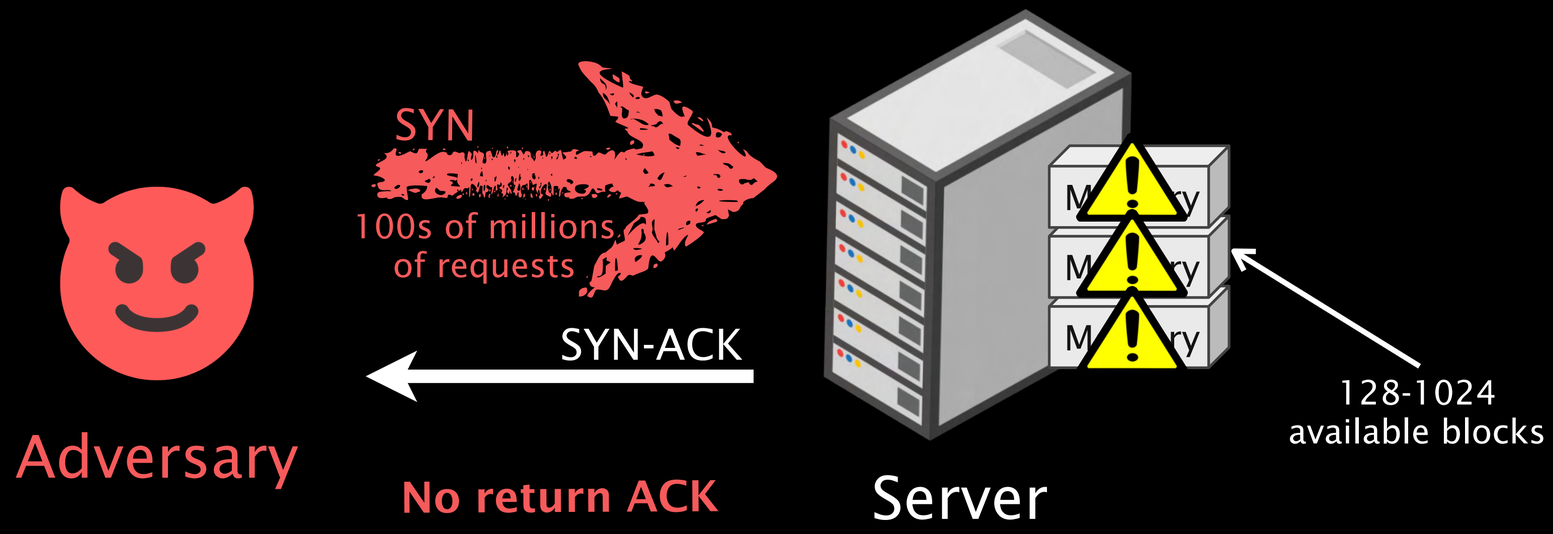

SYN Cookies: a stateless solution

Trading memory…

SYN Cookies: a stateless solution

…for compute

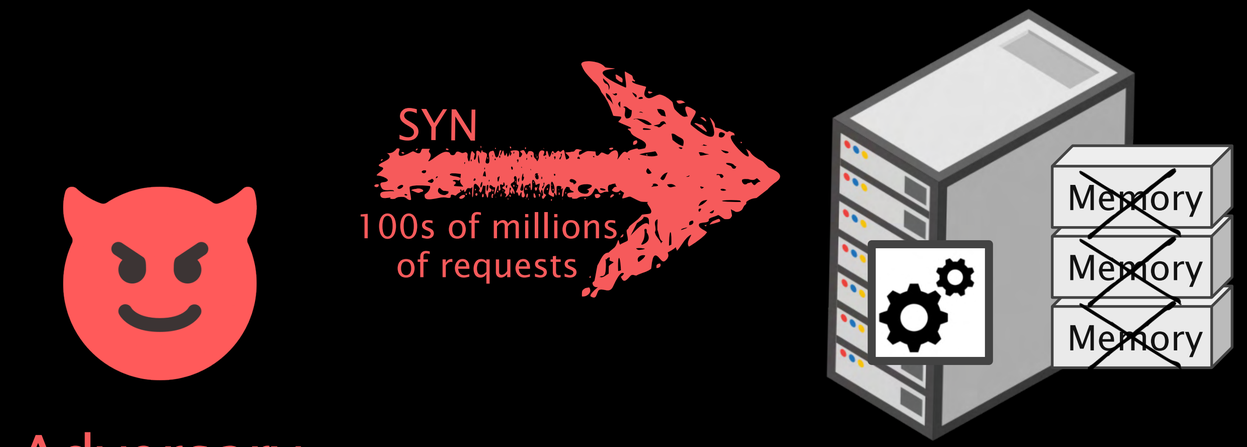

SYN Cookies: a stateless solution

Cryptographically secure "cookie" computed to store the relevant state

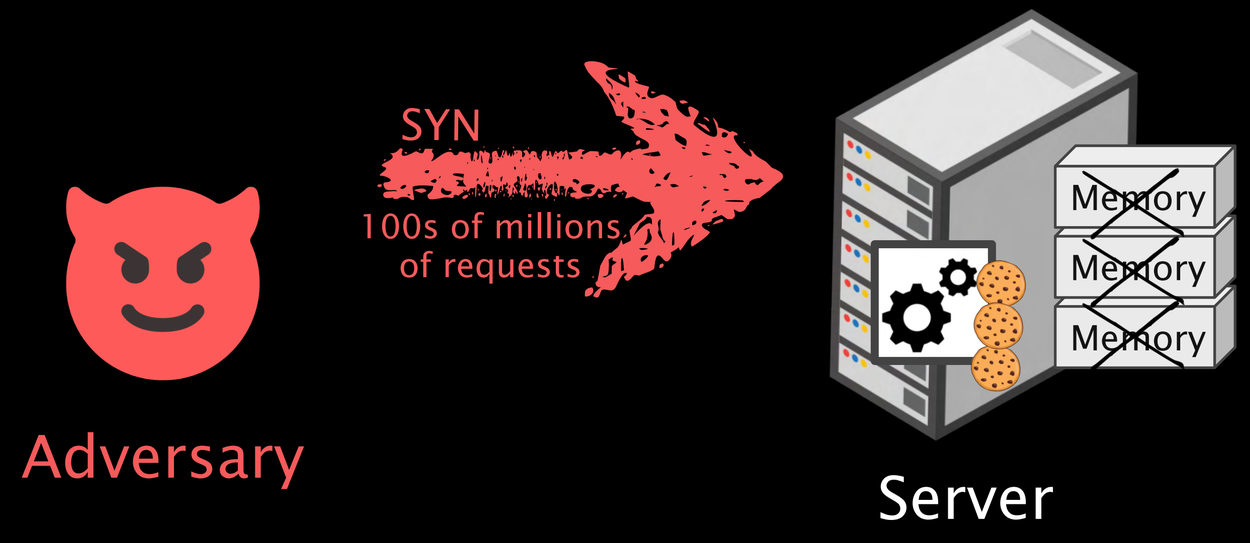

SYN Cookies: a stateless solution

Cryptographically secure "cookie" computed to store the relevant state

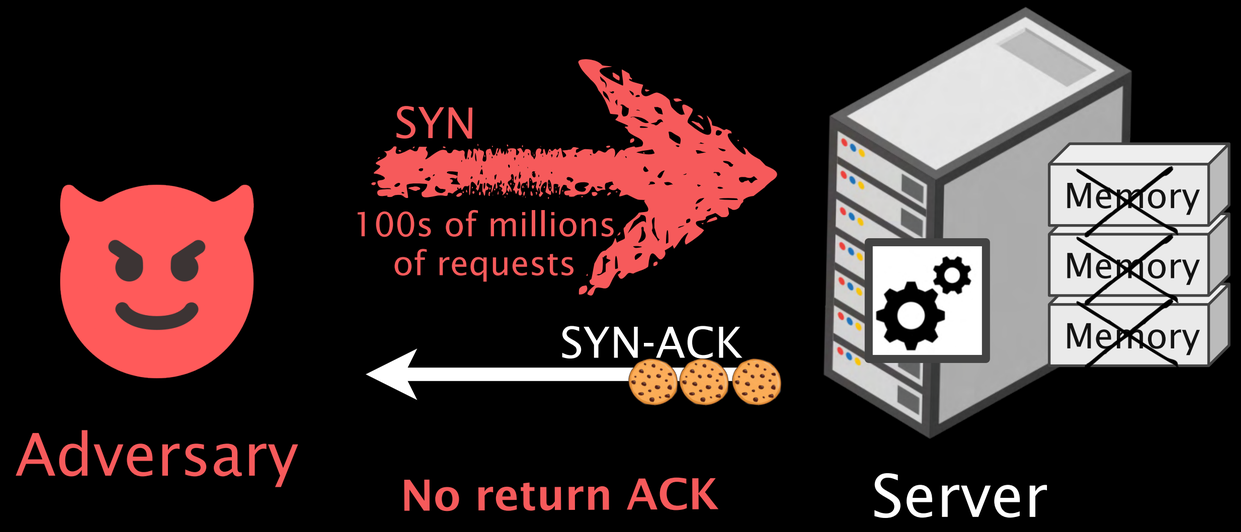

SYN Cookies: a stateless solution

Memory protected for benign connection requests

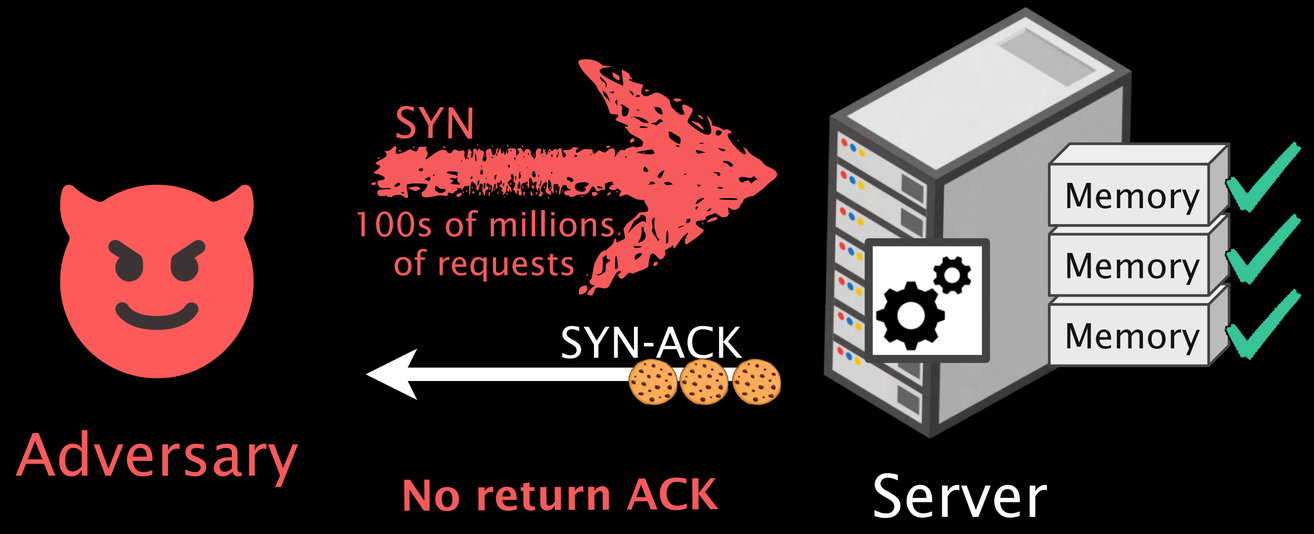

A brief timeline of SYN flooding attacks

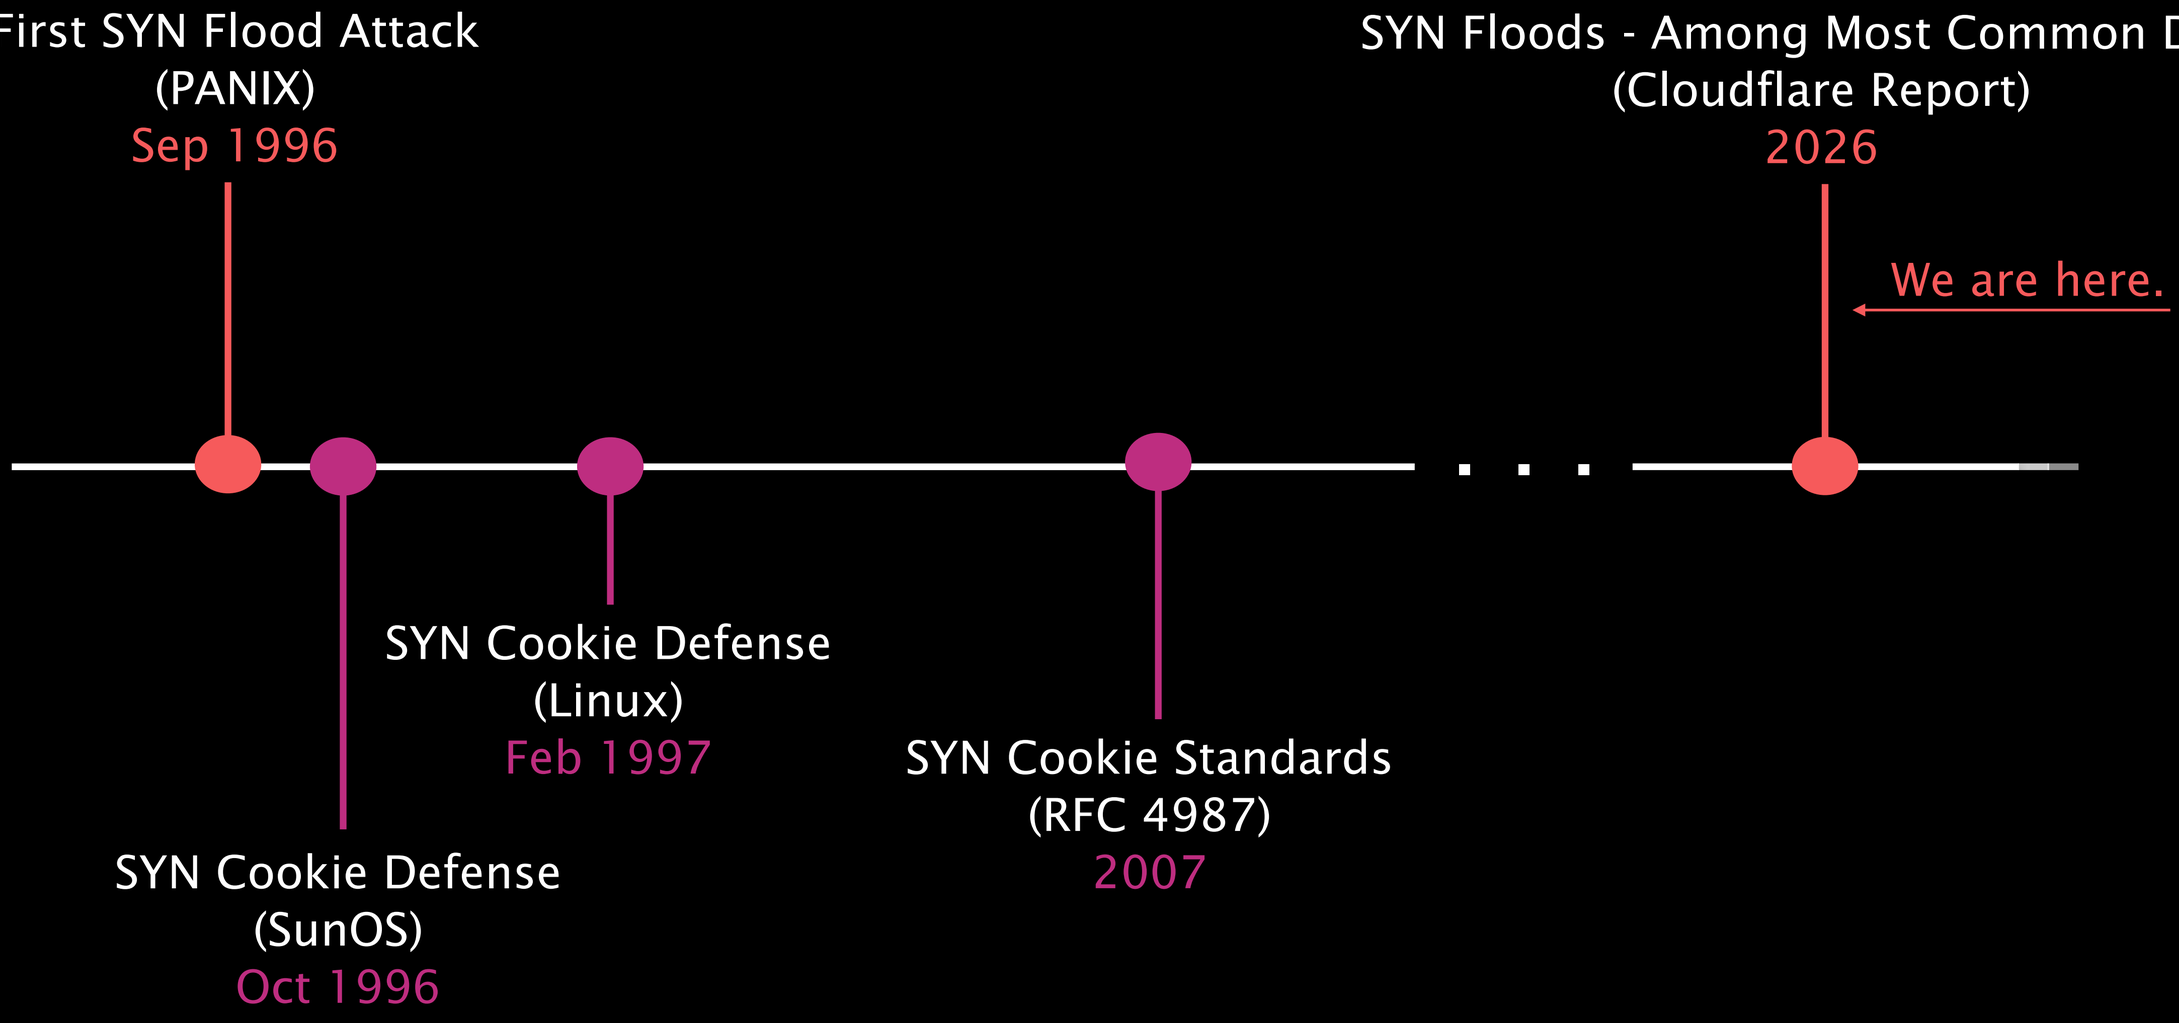

What modern SYN flood defenses really need…

!? Security

Blocking attacks from adaptive adversaries

!? Scalability

Handling large volumes of benign and attack traffic

!? Performance

Maintaining low latency for benign clients

So why is a *secure*, *scalable*, and *performant* defense so hard?

Software Solutions

Network Software

flexible packet processing in general-purpose CPU programs (servers)

Hardware Solutions

Network Hardware

high-speed packet processing in dedicated network chips (switches)

Software Solutions

Hardware Solutions

Software-Only Solutions Can’t Scale Compute

Computational cost to identify attacks leads to early CPU exhaustion.

Cryptographically computing cookies in software is costly

Cryptographic cookie computation and packet processing for every potential connection

Computational strain becomes new attack vector

Server CPU capacity easily overwhelmed under heavy loads

Application performance degraded, clients experience DoS (again)

Why not move the cookie defense to a hardware proxy?

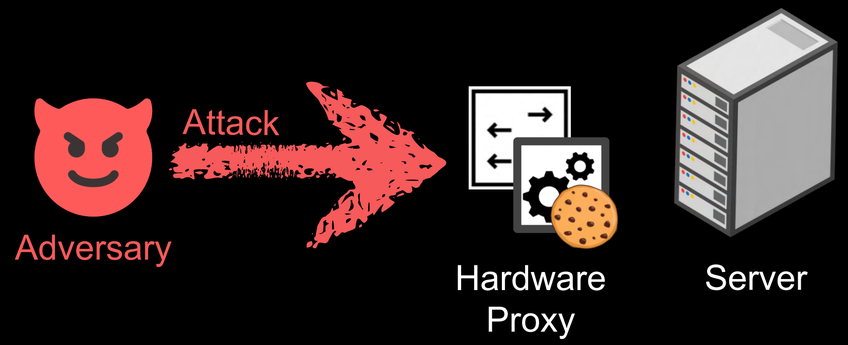

Why not move the cookie defense to a hardware proxy?

…like a high-speed programmable switch!

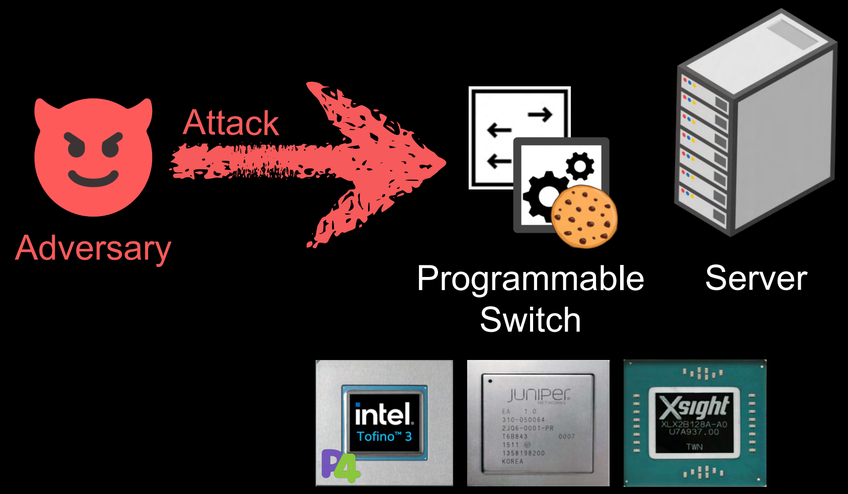

Software Solutions

Software-Only Solutions Can’t Scale Compute

Computational cost to identify attacks leads to early CPU exhaustion.

Hardware Solutions

Hardware-Only Solutions Can’t Scale Memory

Hardware-Only Solutions Are Insecure

Software Solutions

Software-Only Solutions Can’t Scale Compute

Computational cost to identify attacks leads to early CPU exhaustion.

Hardware Solutions

Hardware-Only Solutions Can’t Scale Memory

Usage of limited memory compromises performance.

Hardware-Only Solutions Are Insecure

Tracking verified connections in hardware is costly

Header translation required between client-proxy and proxy-server

Switch proxy must keep per-flow state for ongoing connections

Exhausts limited memory of high-speed switch hardware

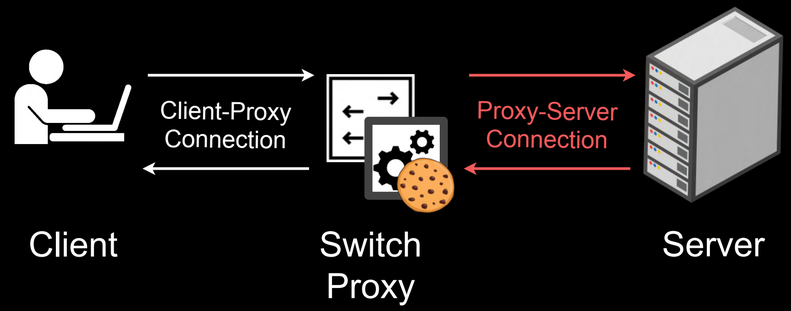

Jaqen[1] avoids memory usage, at a performance cost (extra RTT and added latency for all benign flows)

[1] Liu, et al. *Jaqen: A High-Performance Switch-Native Approach for Detecting and Mitigating Volumetric DDoS Attacks with Programmable Switches.* USENIX Security Symposium, 2021.

Software Solutions

Software-Only Solutions Can’t Scale Compute

Computational cost to identify attacks leads to early CPU exhaustion.

Hardware Solutions

Hardware-Only Solutions Can’t Scale Memory

Usage of limited memory compromises performance.

Hardware-Only Solutions Are Insecure

Weak hashing for cookie generation breaks security.

Insecure hashing breaks cookie security

Cookie = hash(4-tuple, secret)

```
4-tuple = [src_ip, dst_ip, src_port, dst_port]
```

Hash must be strong enough to withstand manipulated hash collisions

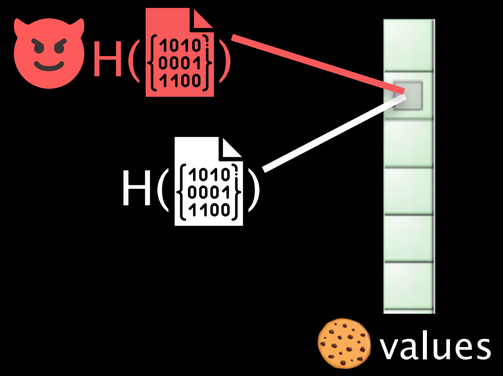

Hardware solutions (Jaqen[1], Poseidon[2]) rely on CRC Checksum for hashing - insecure!

Security abandoned, compute AND memory consumed

[1] Liu, et al. *Jaqen: A High-Performance Switch-Native Approach for Detecting and Mitigating Volumetric DDoS Attacks with Programmable Switches*. USENIX Security Symposium, 2021.
[2] Zhang, et al. *Poseidon: Mitigating volumetric DDoS attacks with programmable switches*. Network and Distributed System Security Symposium, 2020.

Software Solutions

Software-Only Solutions Can’t Scale Compute

Computational cost to identify attacks leads to early CPU exhaustion.

Hardware Solutions

Hardware-Only Solutions Can’t Scale Memory

Usage of limited memory compromises performance.

Hardware-Only Solutions Are Insecure

Weak hashing for cookie generation breaks security.

SmartCookie solves these challenges

SmartCookie solves these challenges

with *hardware + software codesign!*

so what makes SmartCookie so smart?

so what makes SmartCookie so smart?

Intelligent division of labor

what functionality should be partitioned?

how should it be partitioned?

What functionality should be partitioned?

We observe three key elements of SYN cookie proxy defenses

Cookie checks

Header translations

Keeping state for verified connections

How should this functionality be partitioned?

Cookie checks

Header translations

Keeping state for verified connections

Resources must scale for different purposes

Resources must scale for different purposes

Compute must scale to handle both benign+attack traffic
Memory must scale to handle just benign traffic

Resources must scale for different purposes

Compute must scale to handle both benign+attack traffic
Memory must scale to handle just benign traffic

	Compute	Memory
Network Hardware	✓	
Network Software		✓

Resources must scale for different purposes

Compute must scale to handle both benign+attack traffic
Memory must scale to handle just benign traffic

	Compute	Memory	
Network Hardware	✓		should **verify** all traffic (benign+attack)
Network Software		✓	should **track** verified traffic (only benign)

SmartCookie's Hardware-Software Co-Design

- Cookie checks
- Header translations
- Keeping records

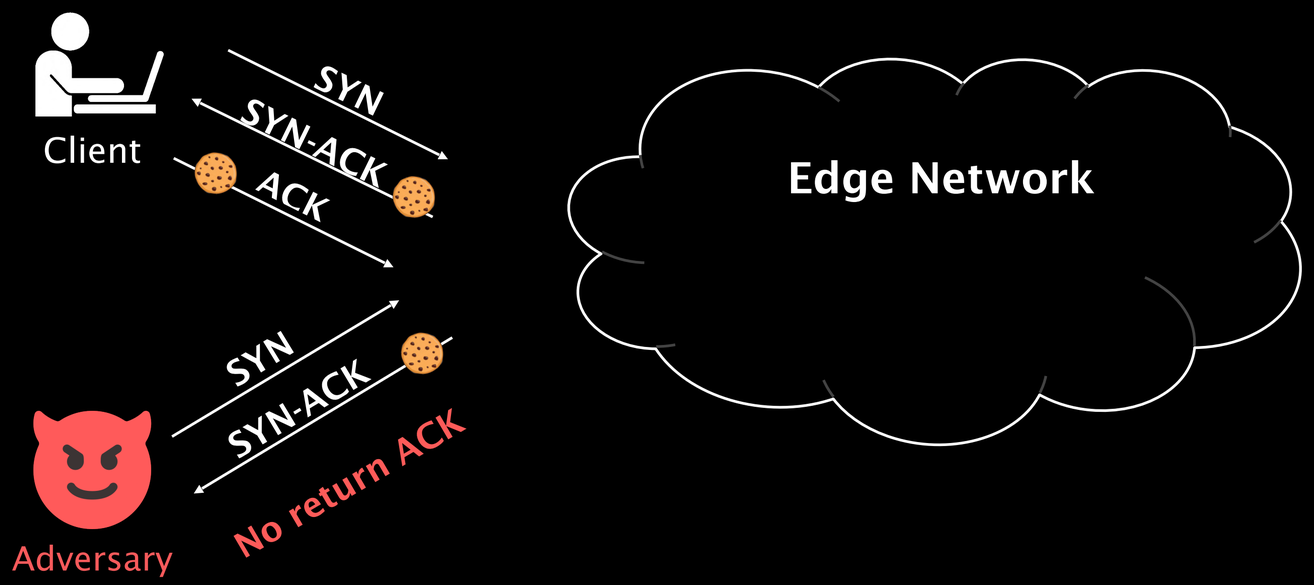

SmartCookie's Hardware-Software Co-Design

Switch agent performs cookie checks with a robust hash, not CRC*

Cookie checks

Header translations

Keeping records

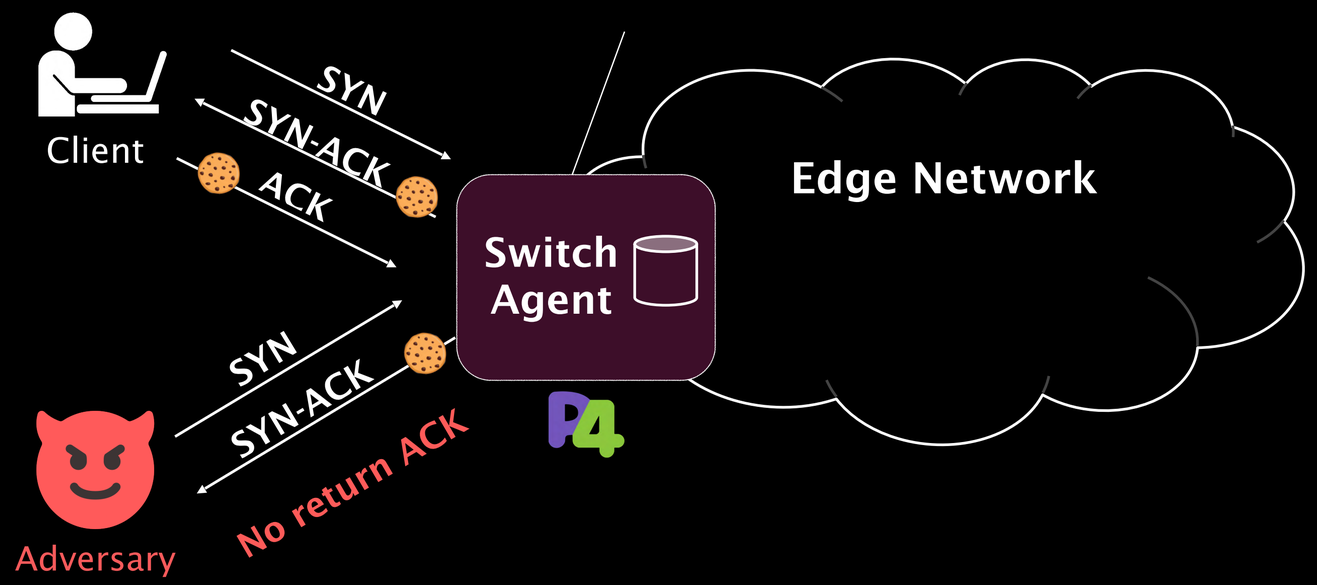

*Yoo, et al. *Secure Keyed Hashing on Programmable Switches.* ACM SIGCOMM Workshop on Secure Programmable Network Infrastructure (SPIN), 2021.

SmartCookie's Hardware-Software Co-Design

Switch agent approximately tracks verified connections

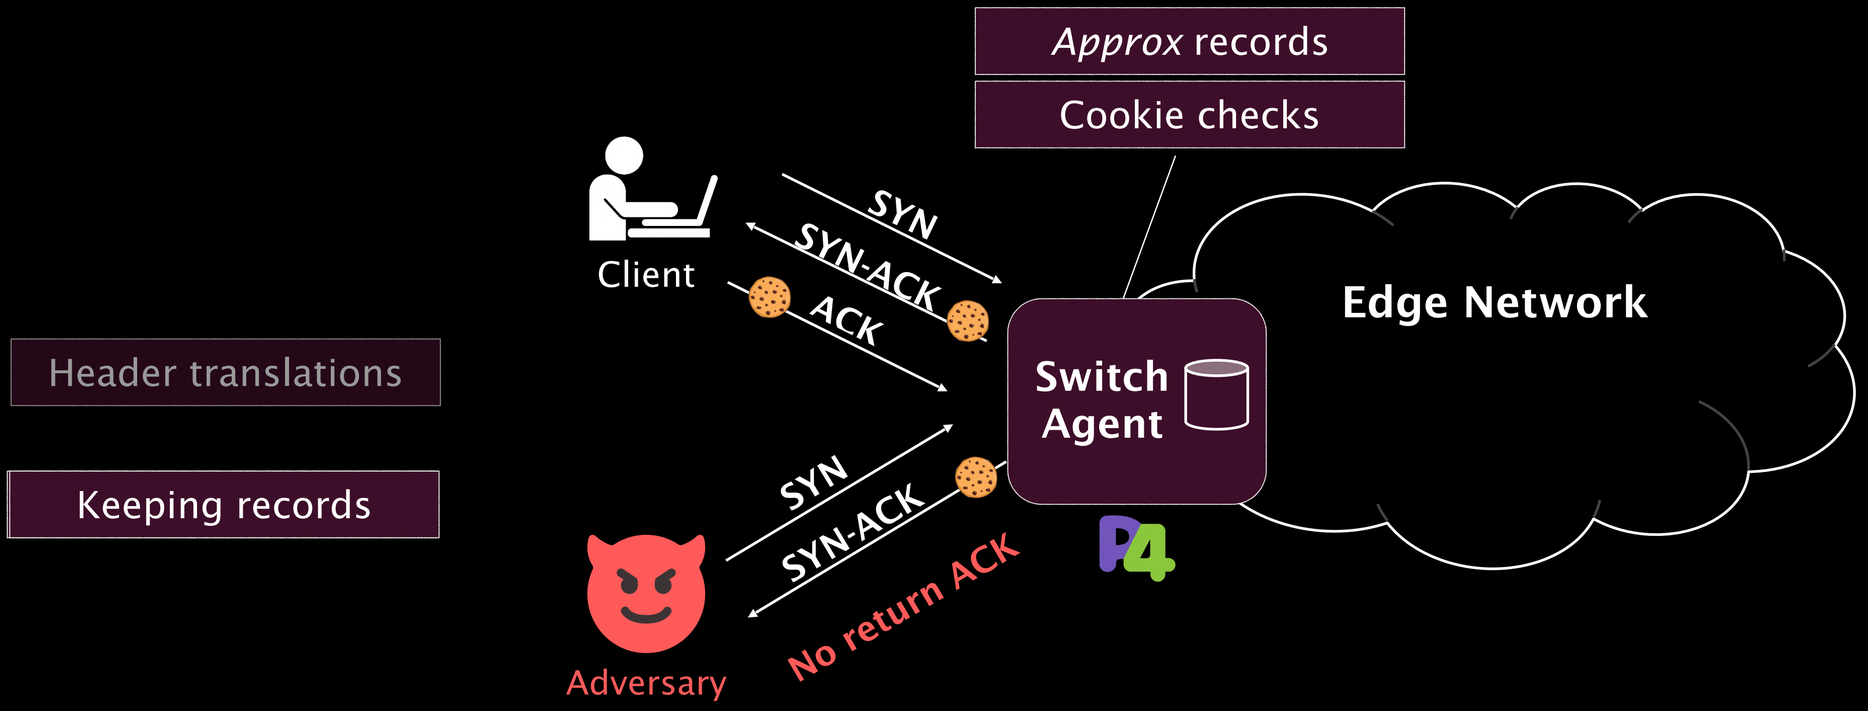

SmartCookie's Hardware-Software Co-Design

Server agent handles header translations on behalf of switch agent

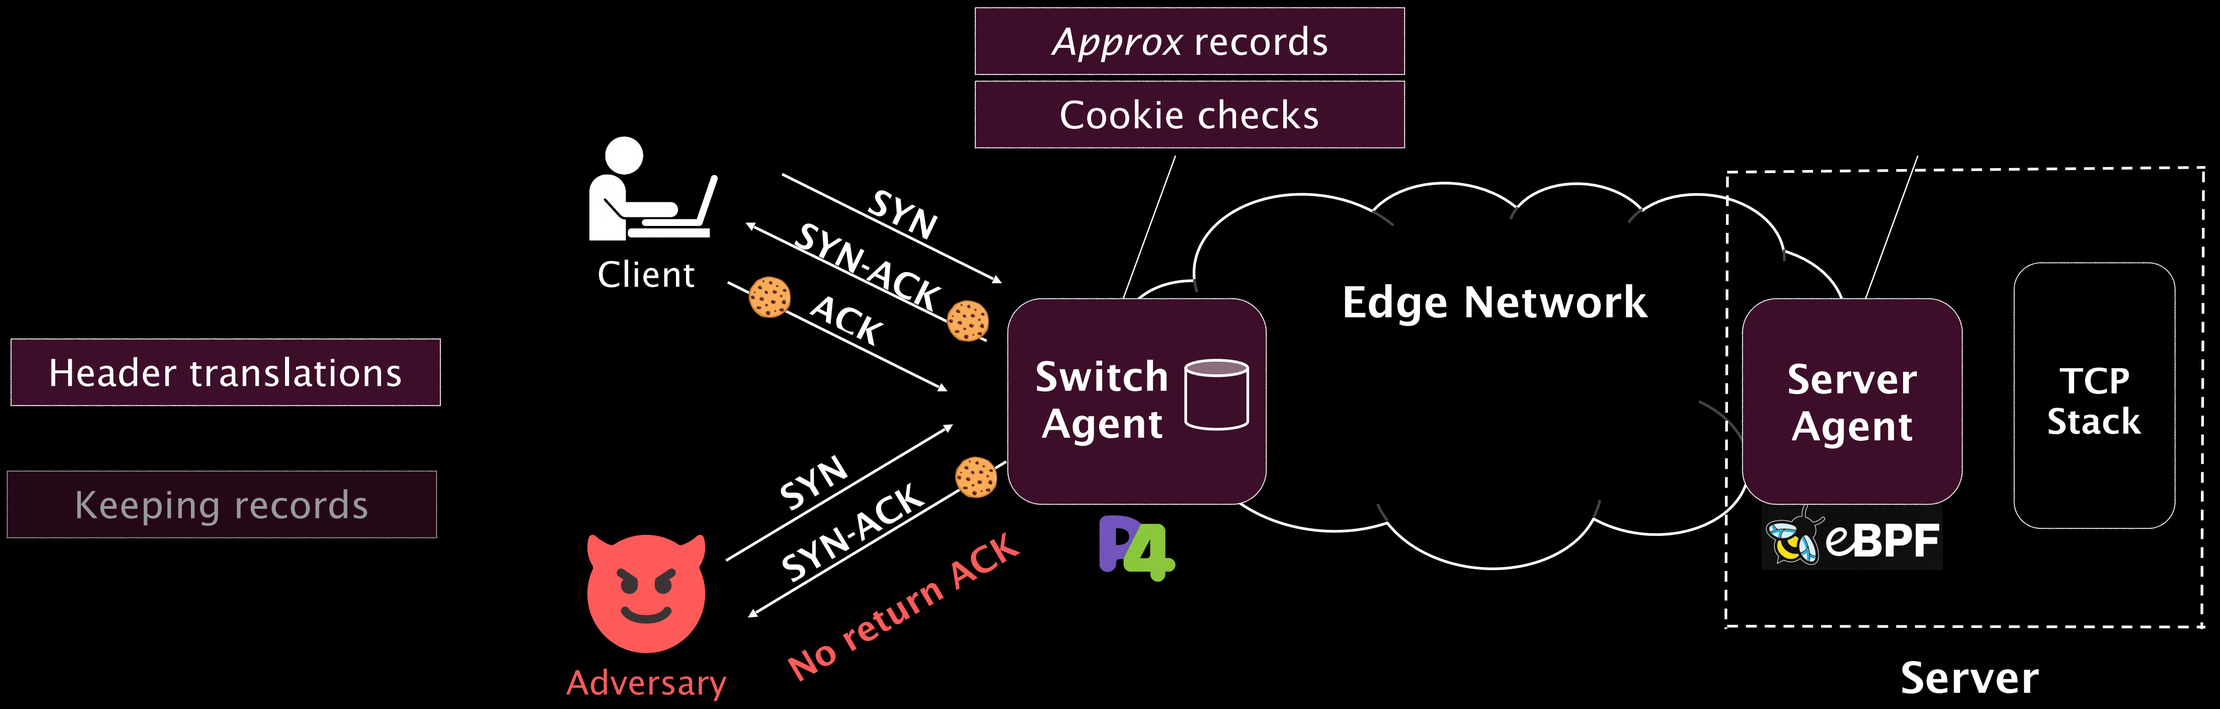

SmartCookie's Hardware-Software Co-Design

Server agent exactly tracks verified connections

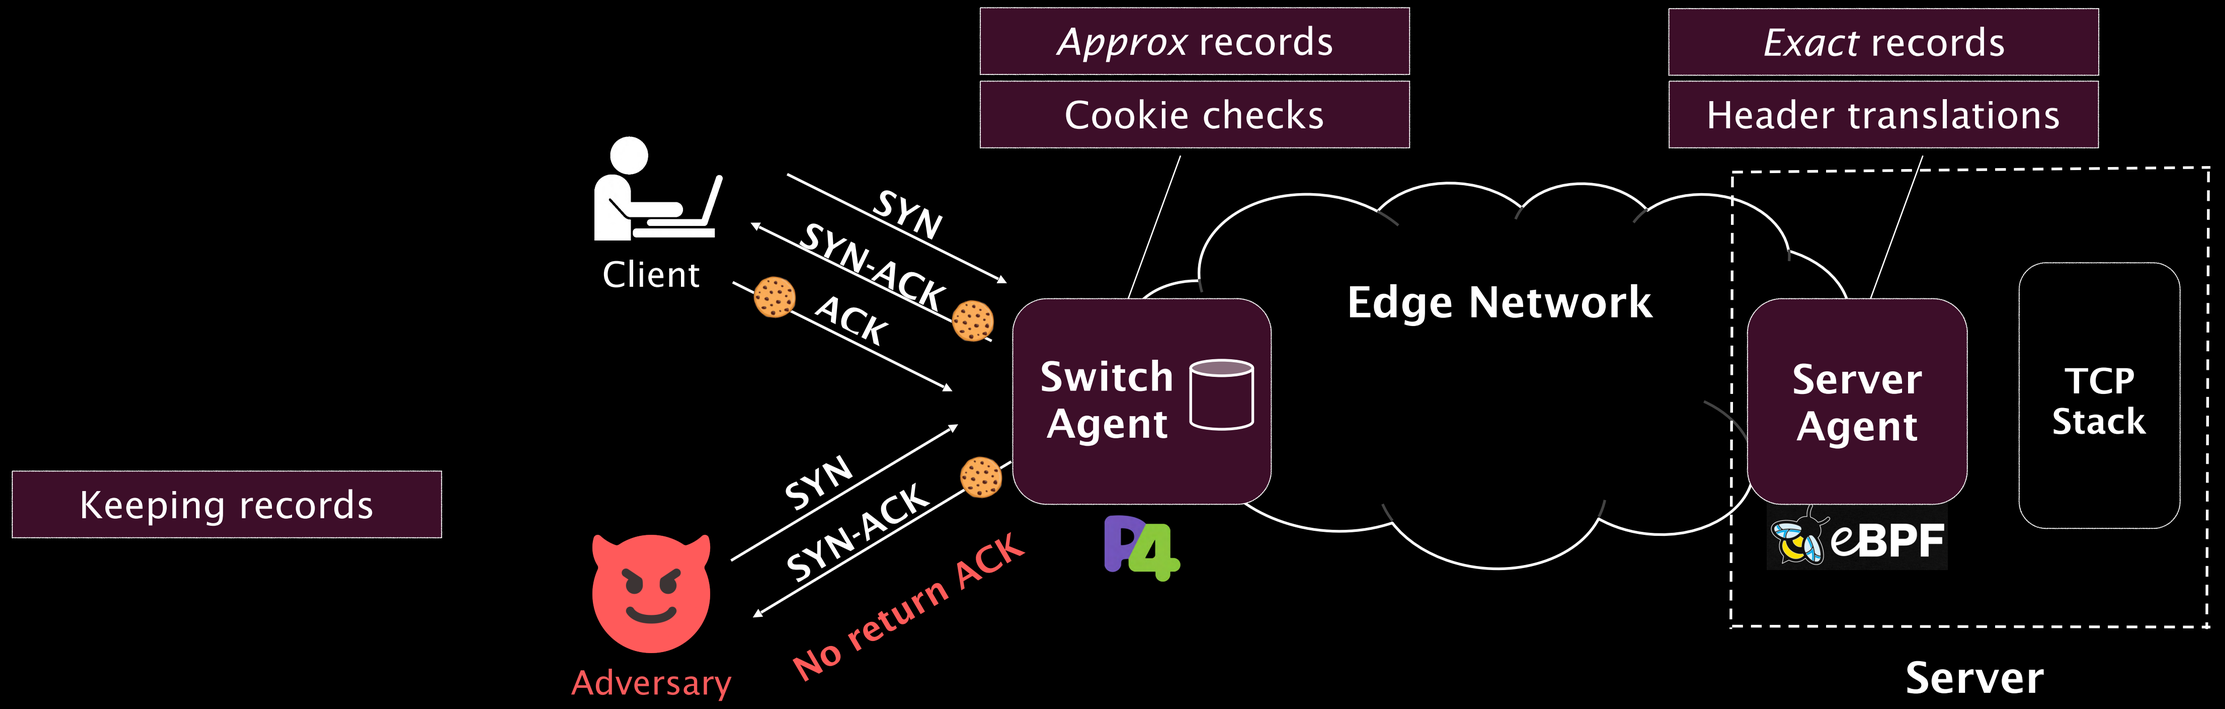

SmartCookie's Hardware-Software Co-Design

Custom collaborative protocol between SmartCookie components

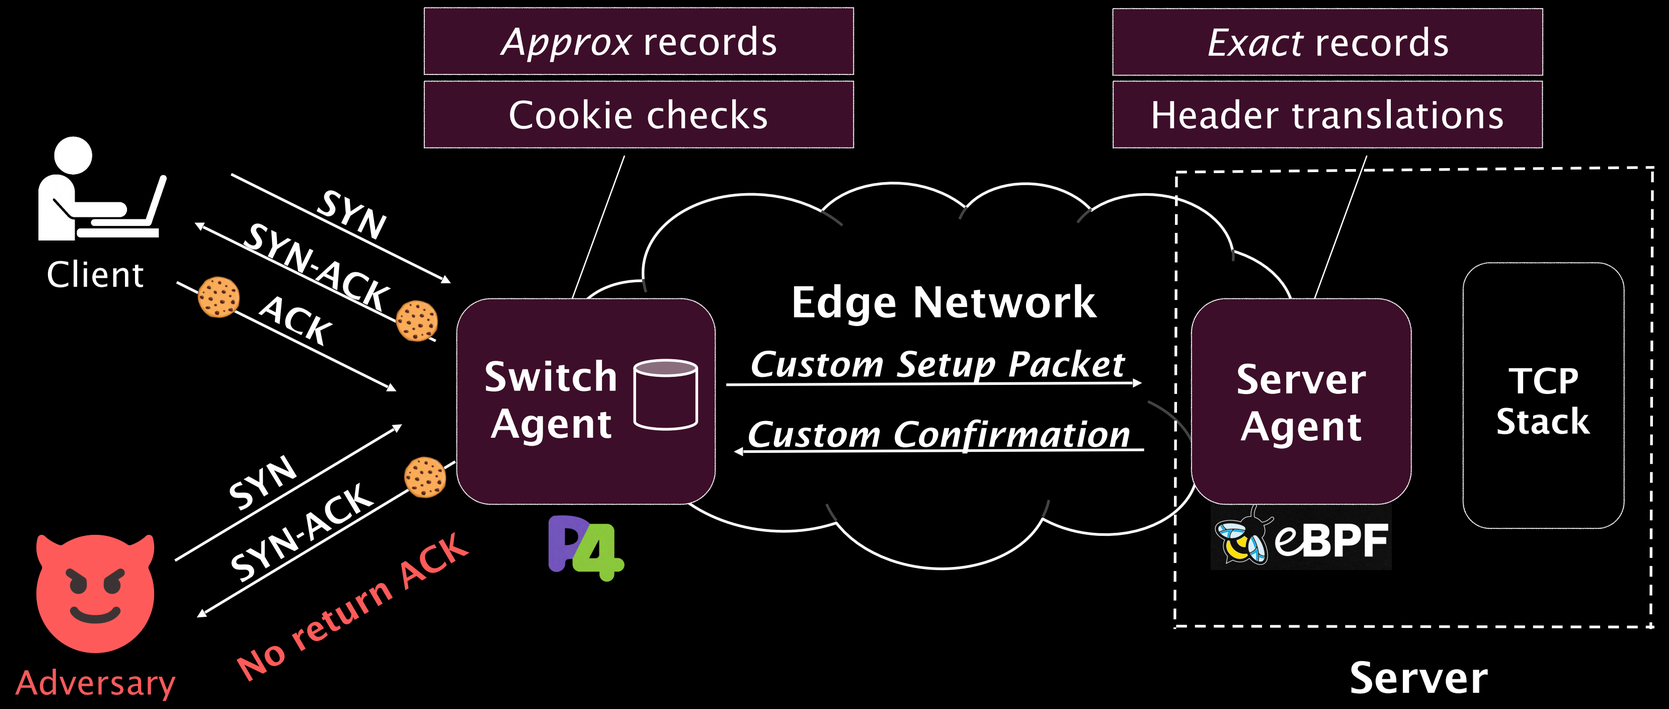

SmartCookie's Hardware-Software Co-Design

…does not require any modifications to the client or server's network stack

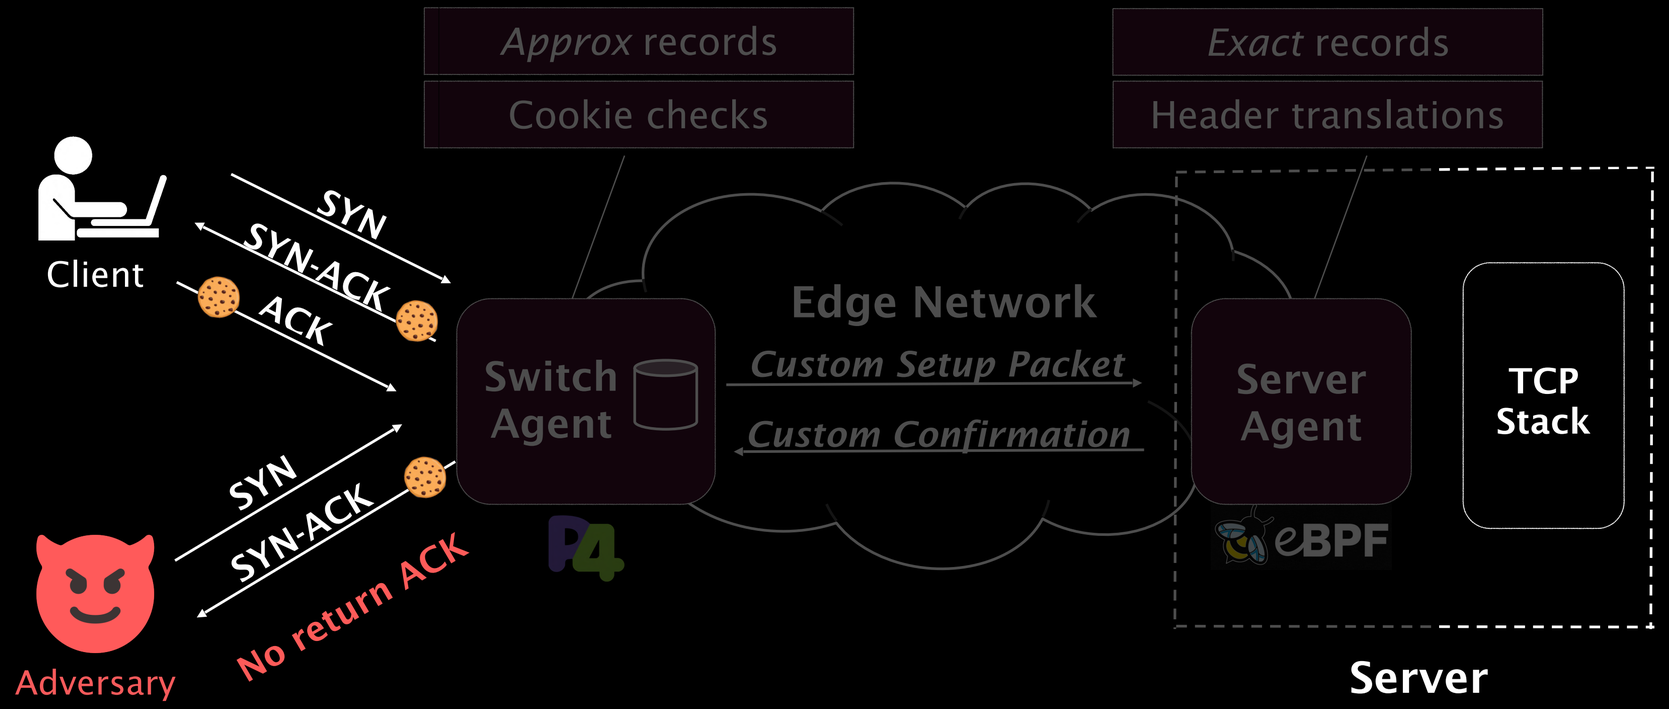

So does it work?

Evaluation

Security

Can we deliver security at high attack rates?

Scalability

Can we protect CPU capacity for scalability?

Performance

Can we maintain client performance under attack?

Evaluation

Security

Can we deliver security at high attack rates?

Scalability

Can we protect CPU capacity for scalability?

Performance

Can we maintain client performance under attack?

SmartCookie

outperforms software solutions
by two orders of magnitude

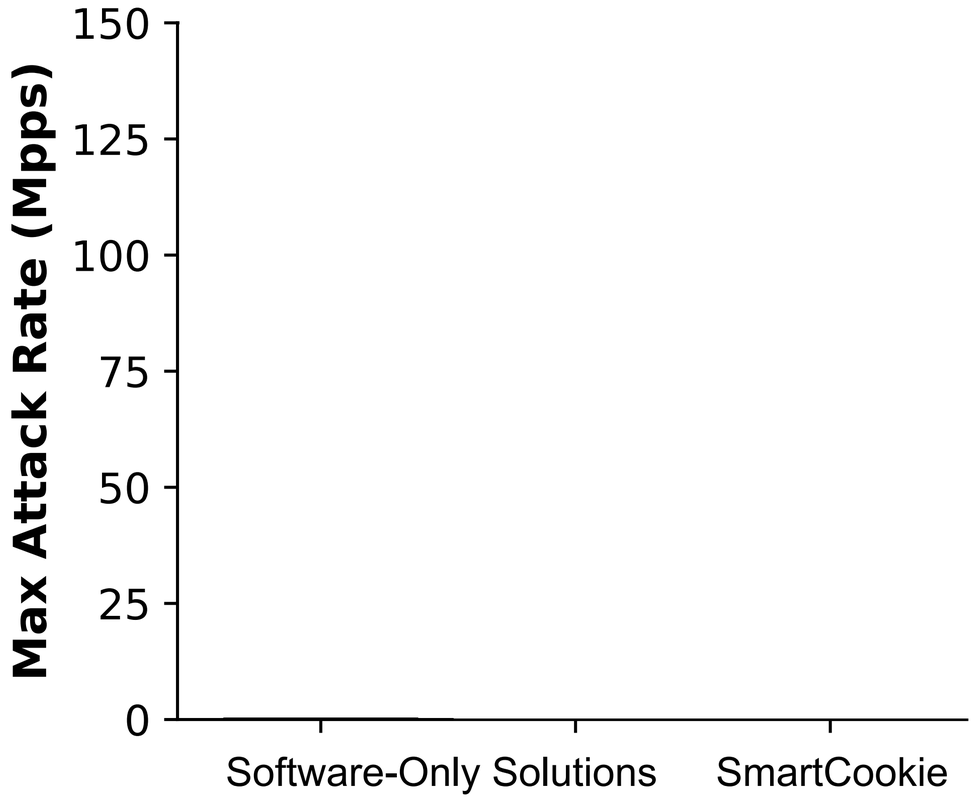

SmartCookie

outperforms software solutions
by two orders of magnitude

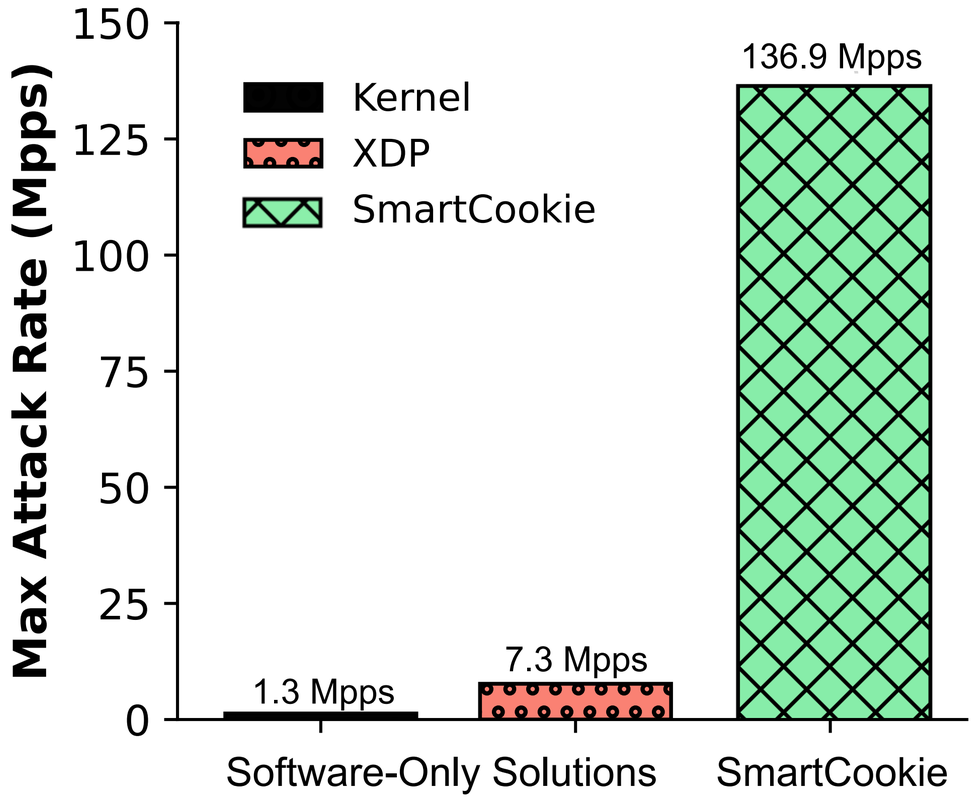

Evaluation

Security

Can we deliver security at high attack rates?

Scalability

Can we protect CPU capacity for scalability?

Performance

Can we maintain client performance under attack?

SmartCookie

protects CPU capacity
with zero overhead

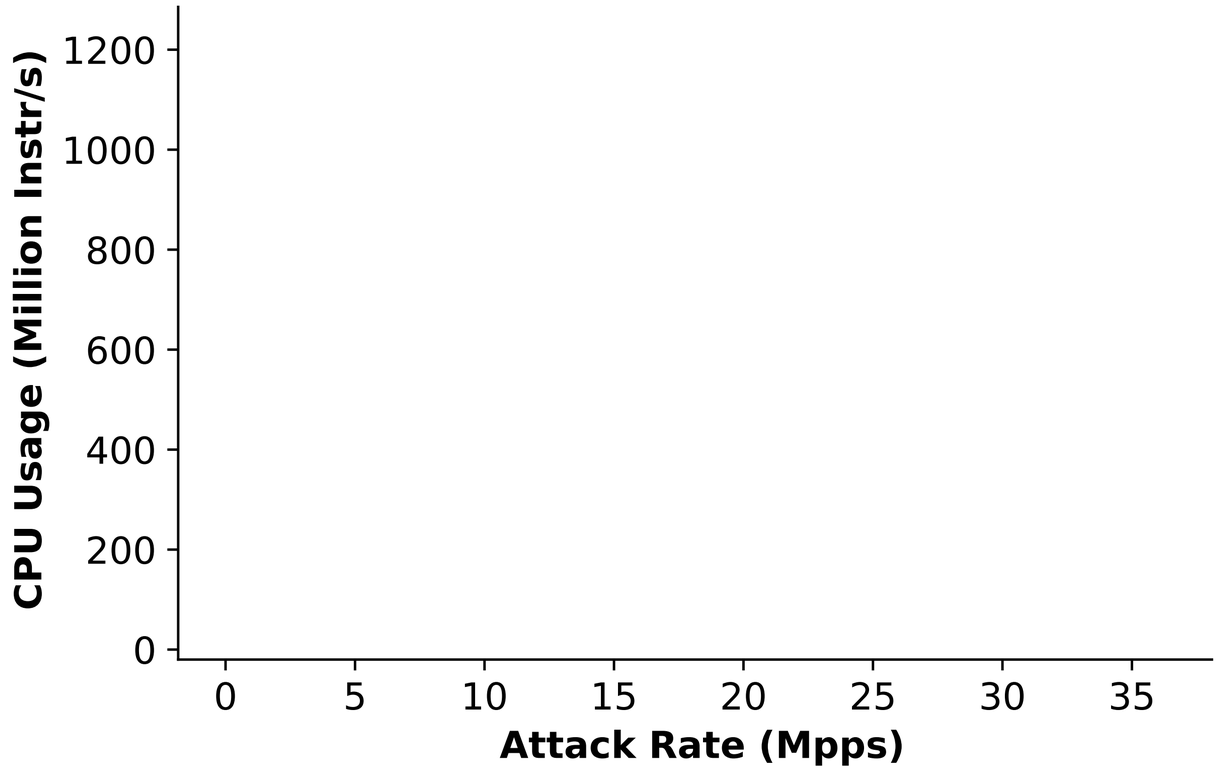

SmartCookie

protects CPU capacity
with zero overhead

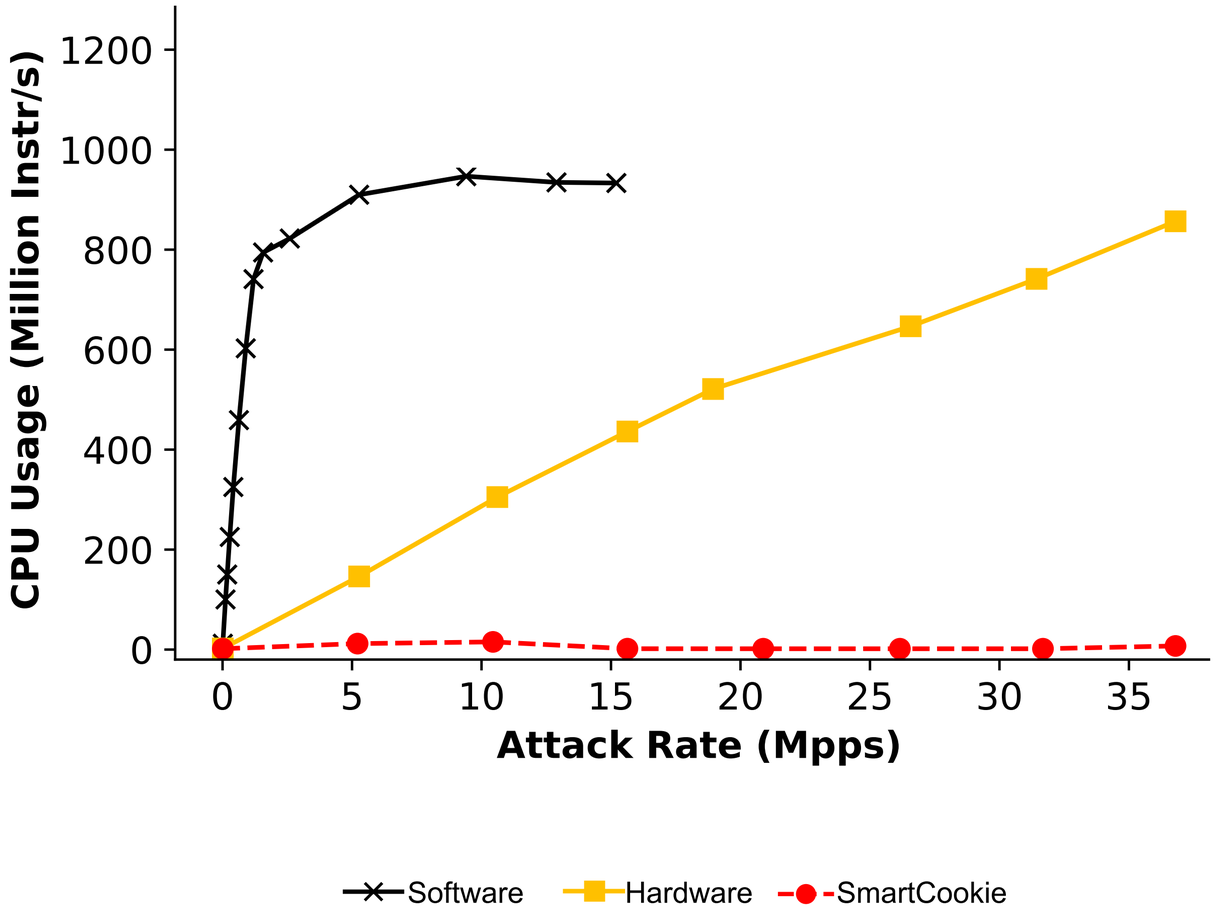

Evaluation

Security

Can we deliver security at high attack rates?

Scalability

Can we protect CPU capacity for scalability?

Performance

Can we maintain client performance under attack?

SmartCookie

reduces latency by 48-84%
vs. hardware solutions

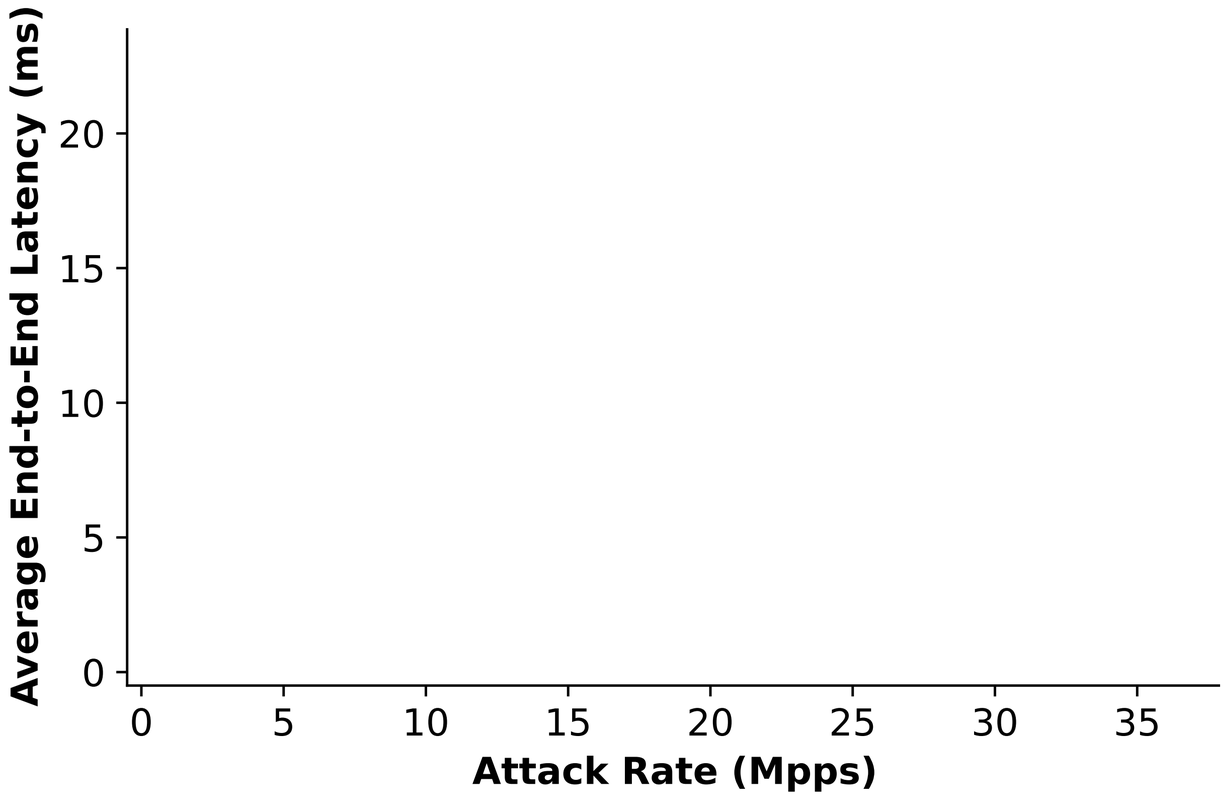

SmartCookie

reduces latency by 48-84%
vs. hardware solutions

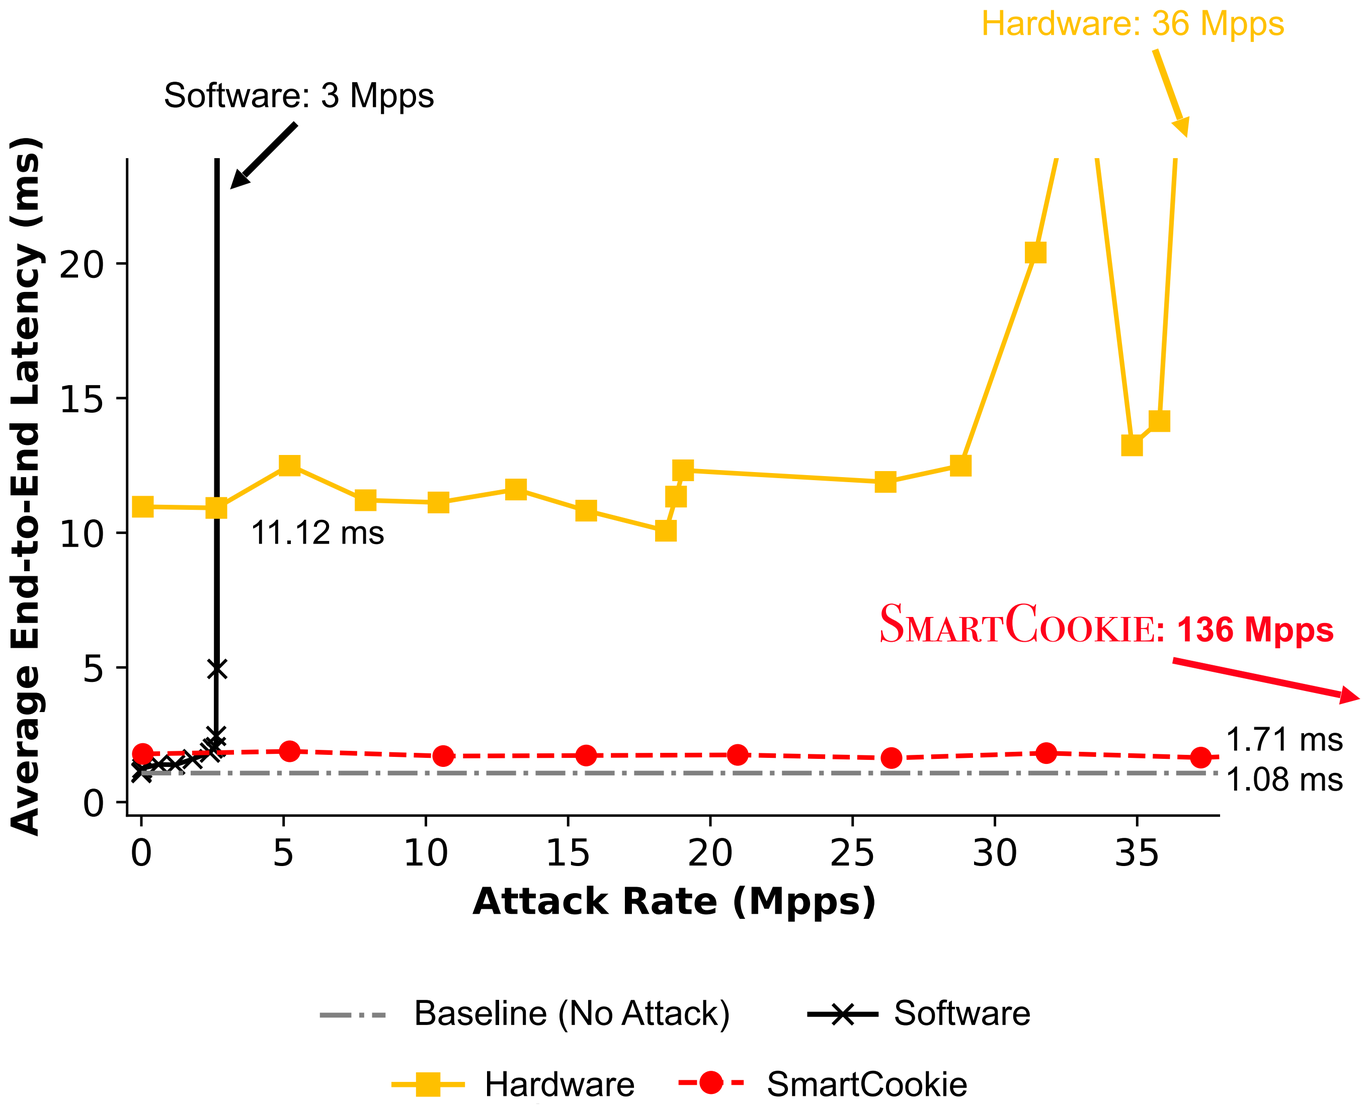

Part I: Ingress Control

Co-design within
a single edge network

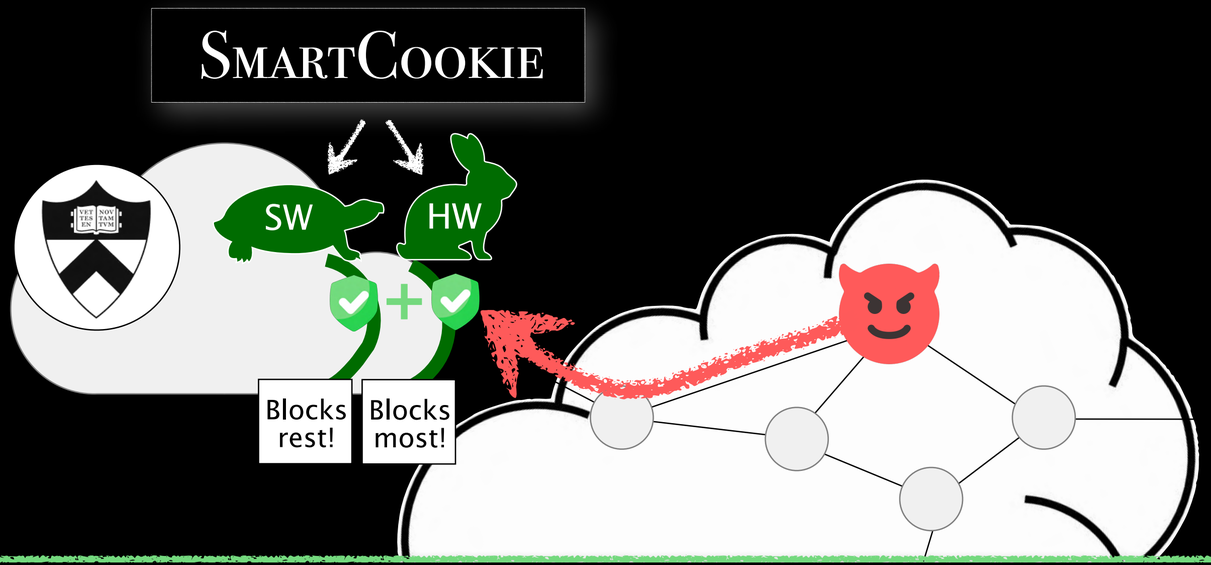

Part I: Ingress Control

Co-design within a single edge network

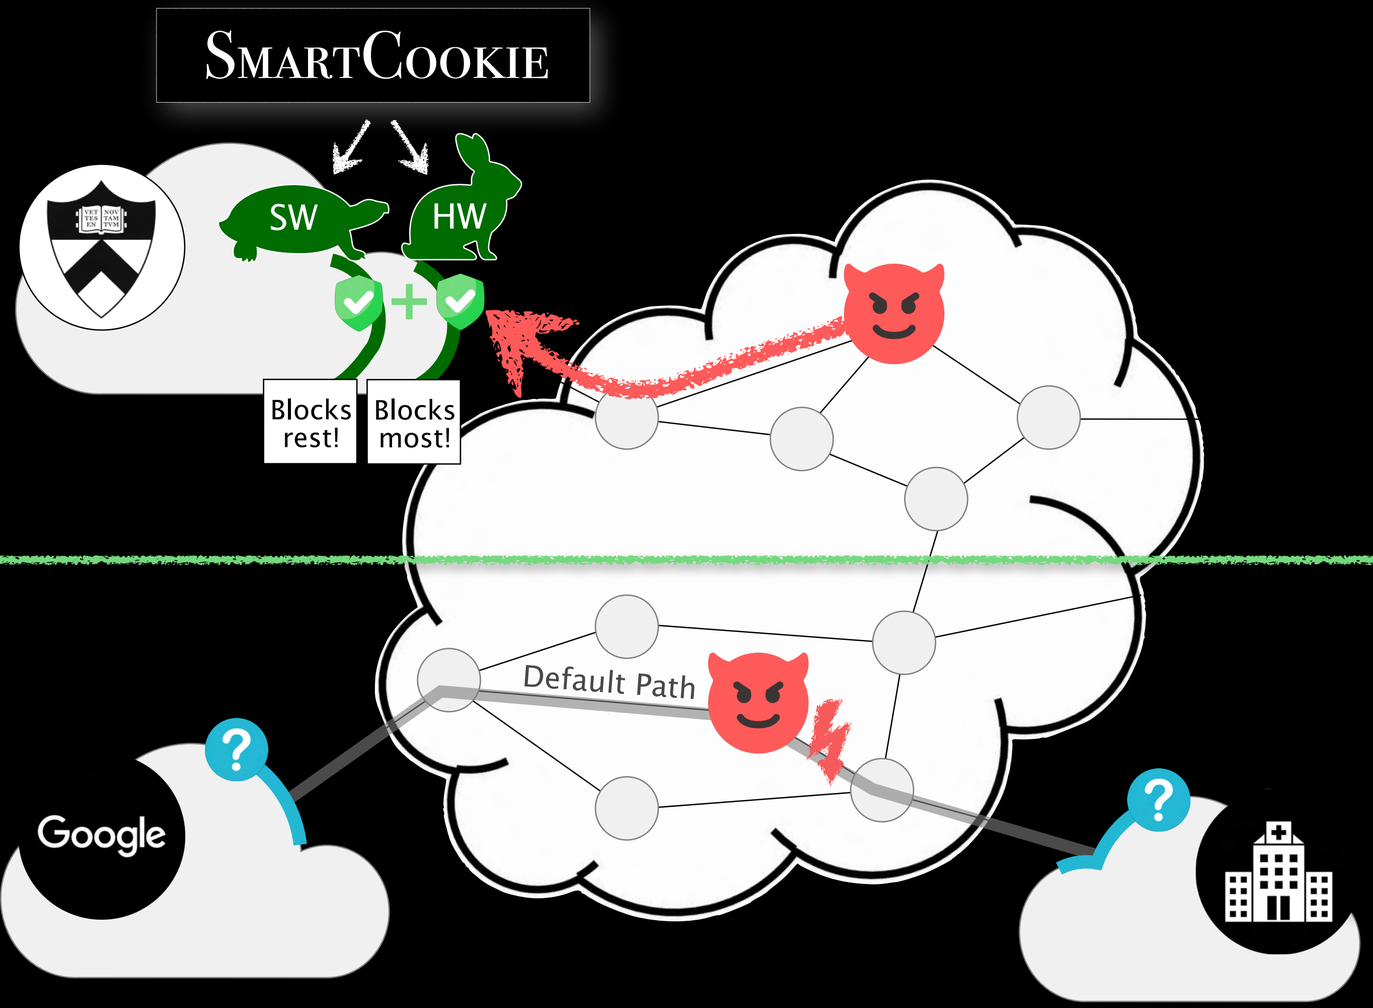

Part II: Route Control

Cooperation between multiple edge networks?

SmartCookie

[USENIX Sec '24]

Part I: Ingress Control via Intra-Edge Co-Design

Tango + PraxiGuard

[NSDI '24] [in preparation]

Part II: Route Control via Inter-Edge Cooperation

Edge networks communicate across the WAN

Edge networks communicate across the WAN

made up of ASes with individual administrative control

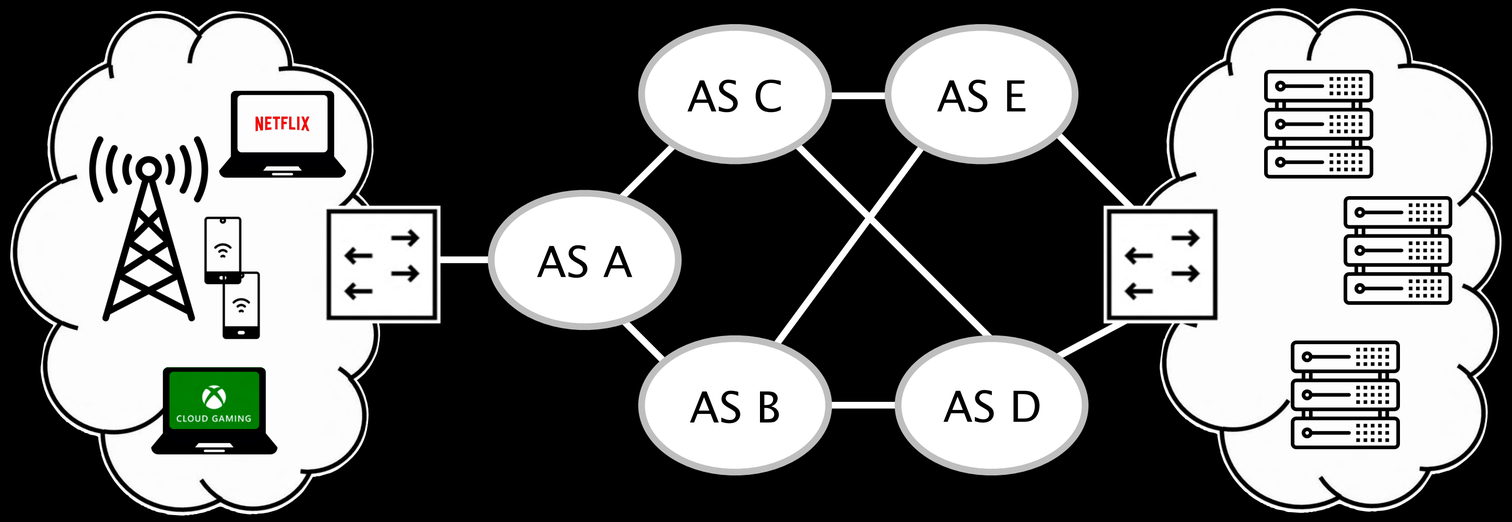

*AS: Autonomous System

Default routes are assigned by BGP

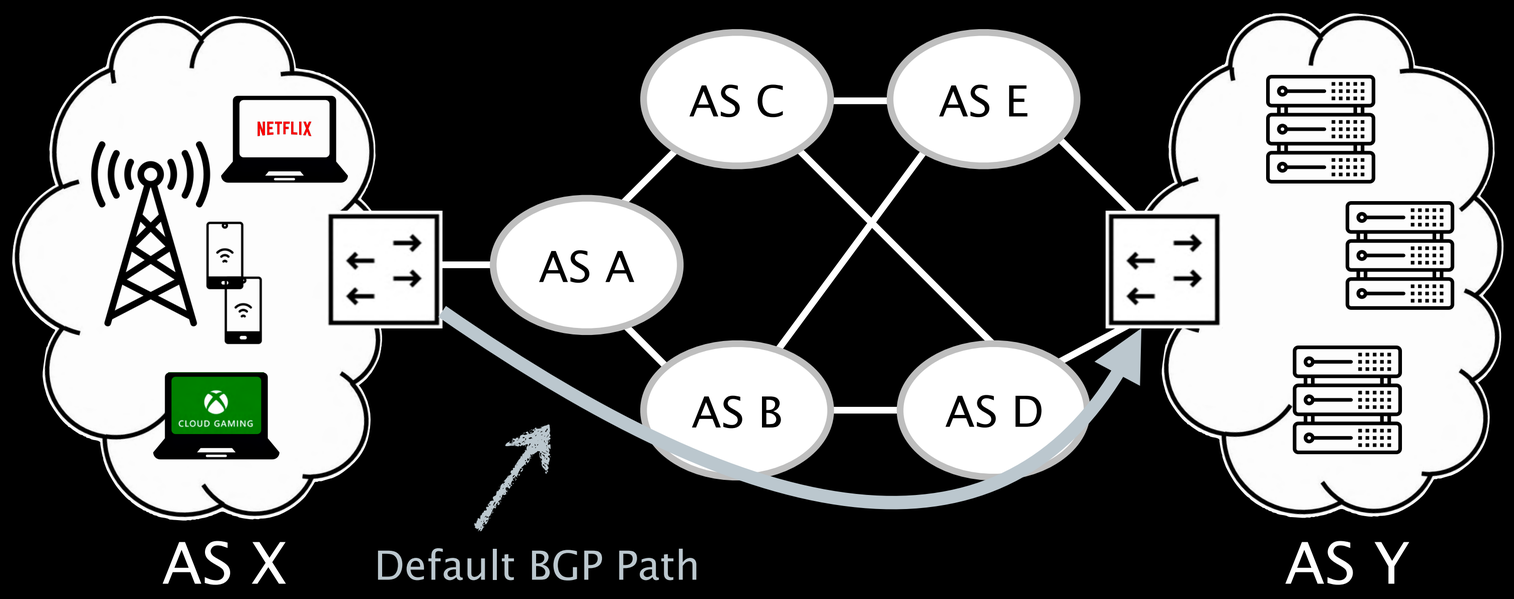

*BGP: Border Gateway Protocol

Default routes are assigned by BGP

based on economic incentives

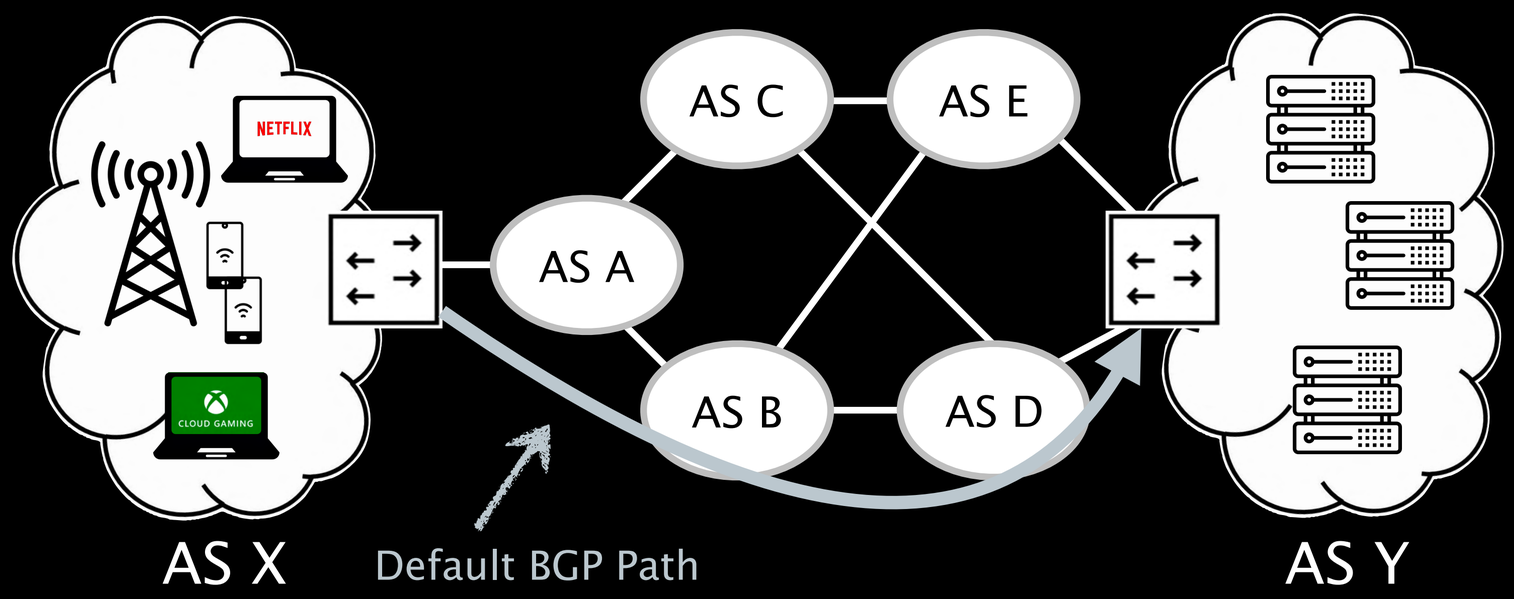

Default routes are assigned by BGP

based on economic incentives, not privacy/performance considerations

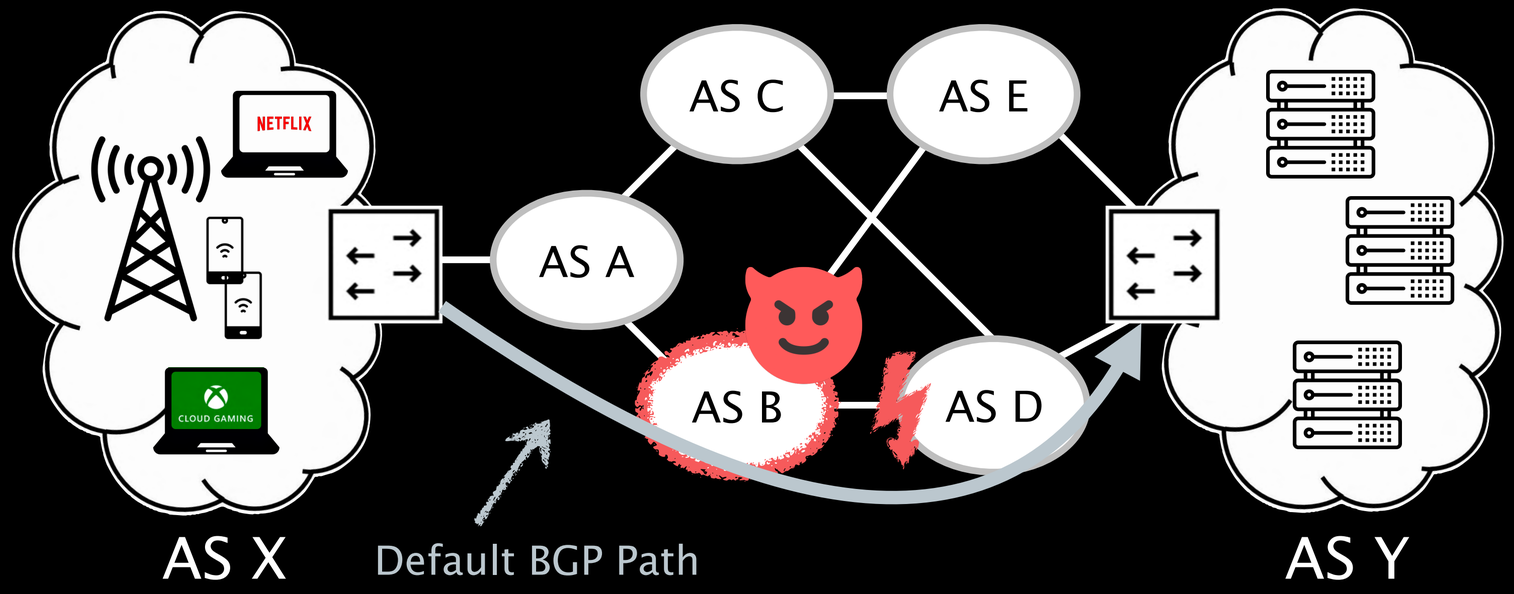

Even if *better* paths exist, edge ASes cannot use them

more performant? *more trustworthy?*

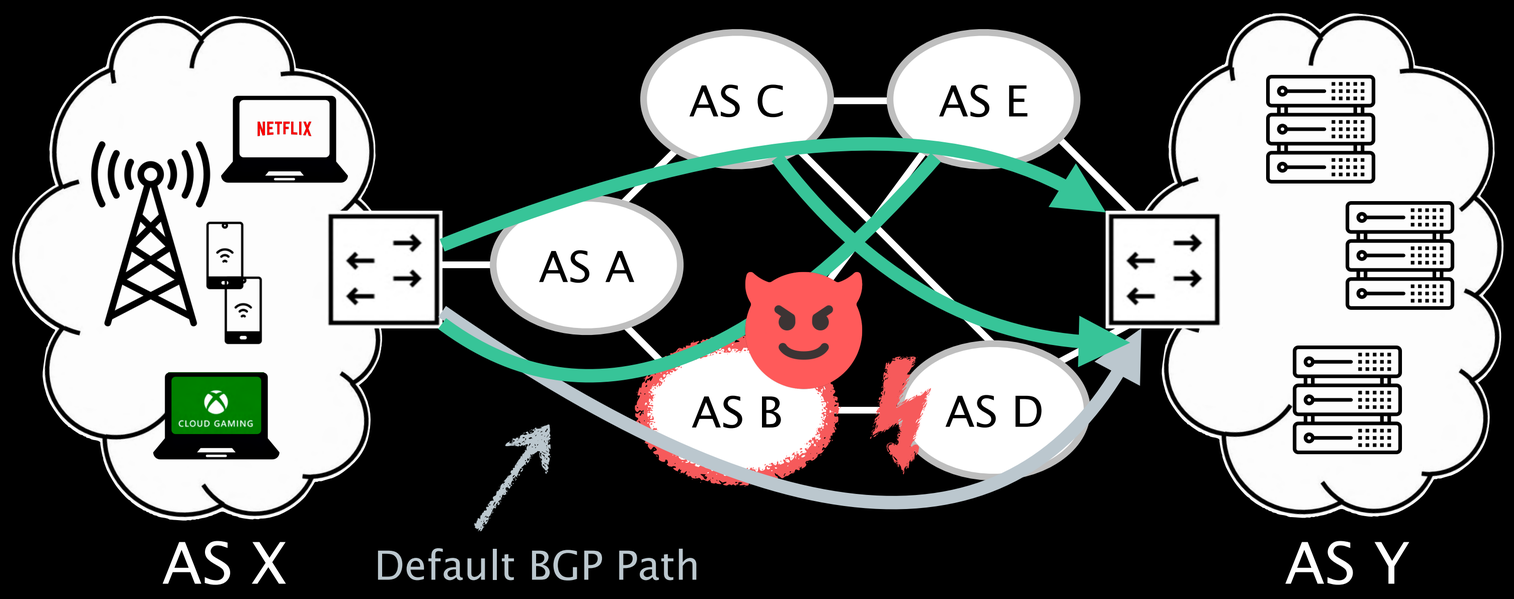

What would it take to give edges route control

over real-world Internet paths?

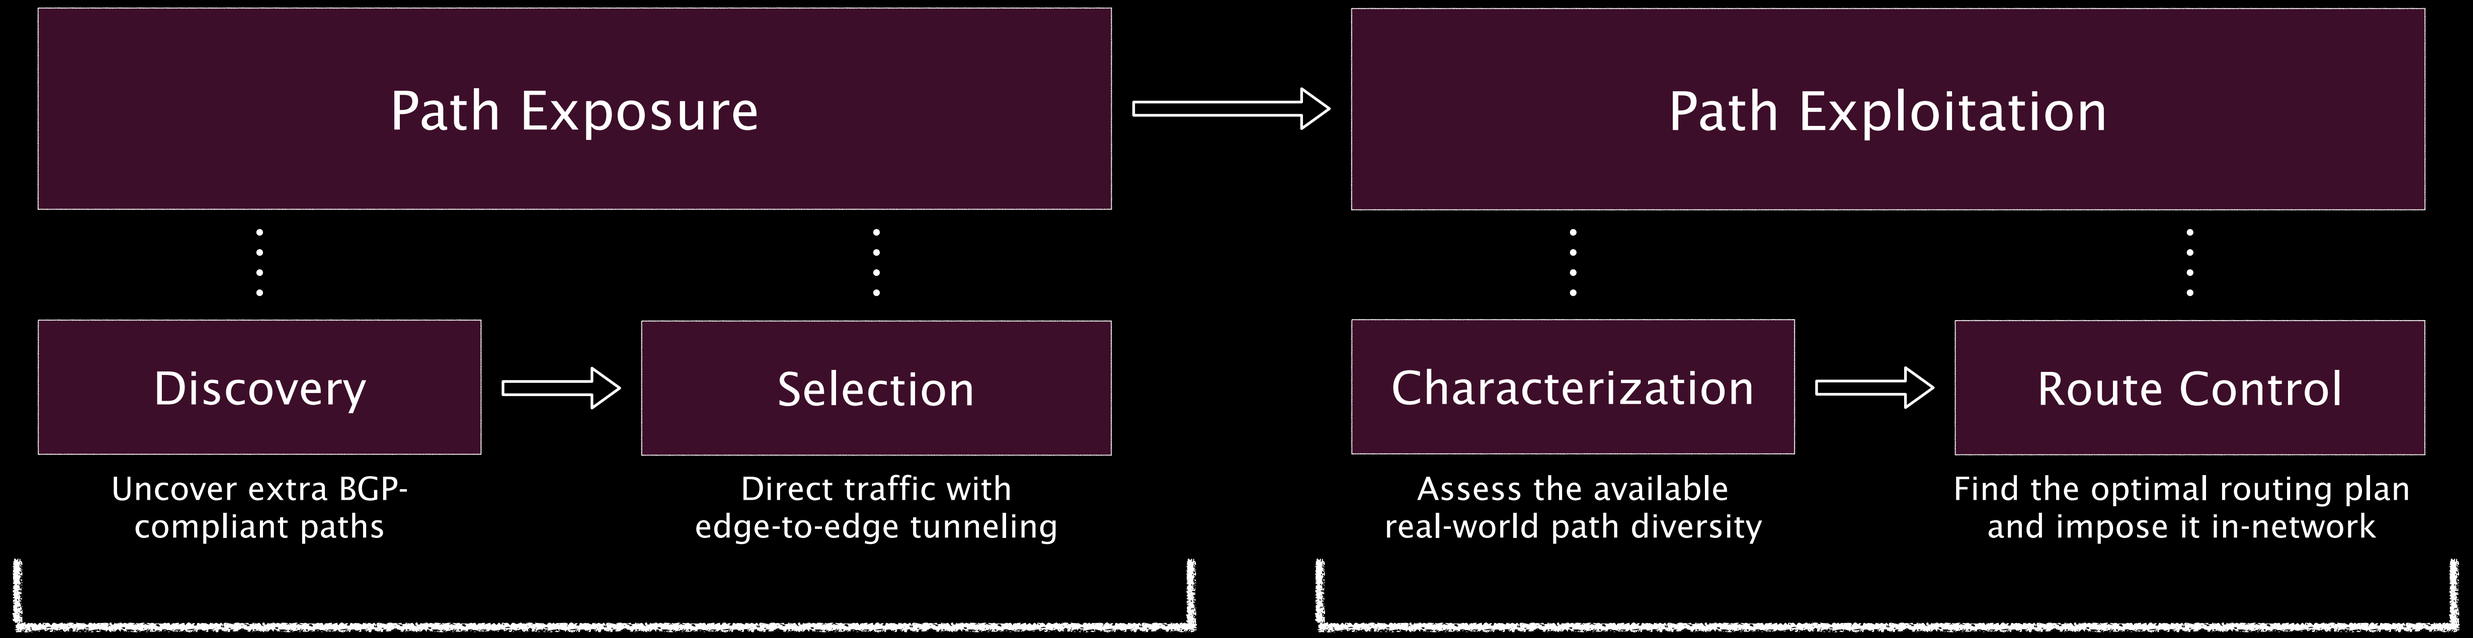

Birge-Lee, ***Yoo***, et al. *Tango: Secure Collaborative Route Control Across the Public Internet*. NSDI, 2024.

Yoo, et al. *PraxiGuard: Network-Aware Traffic Splitting for Website Fingerprinting Defenses*. In Preparation.

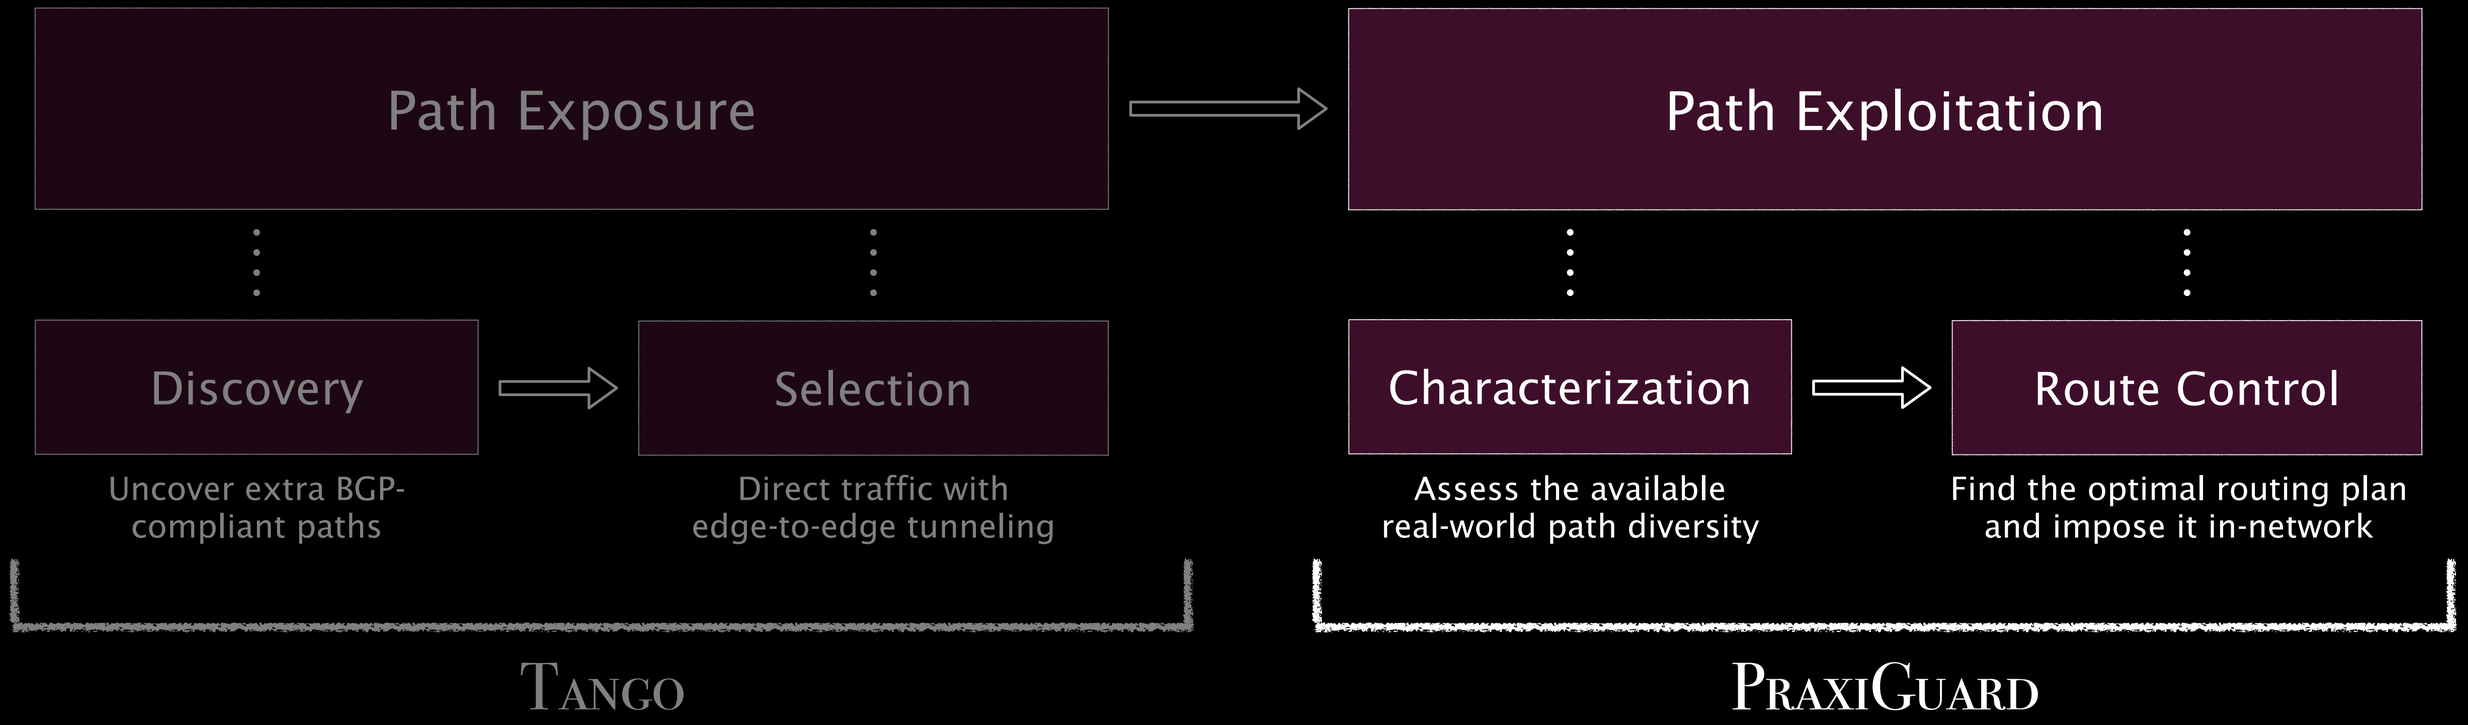
Path Exposure
Path Exploitation
Discovery
Selection
Characterization
Route Control
Uncover extra BGP-compliant paths
Direct traffic with edge-to-edge tunneling
Assess the available real-world path diversity
Find the optimal routing plan and impose it in-network
TANGO
PRAXIGUARD

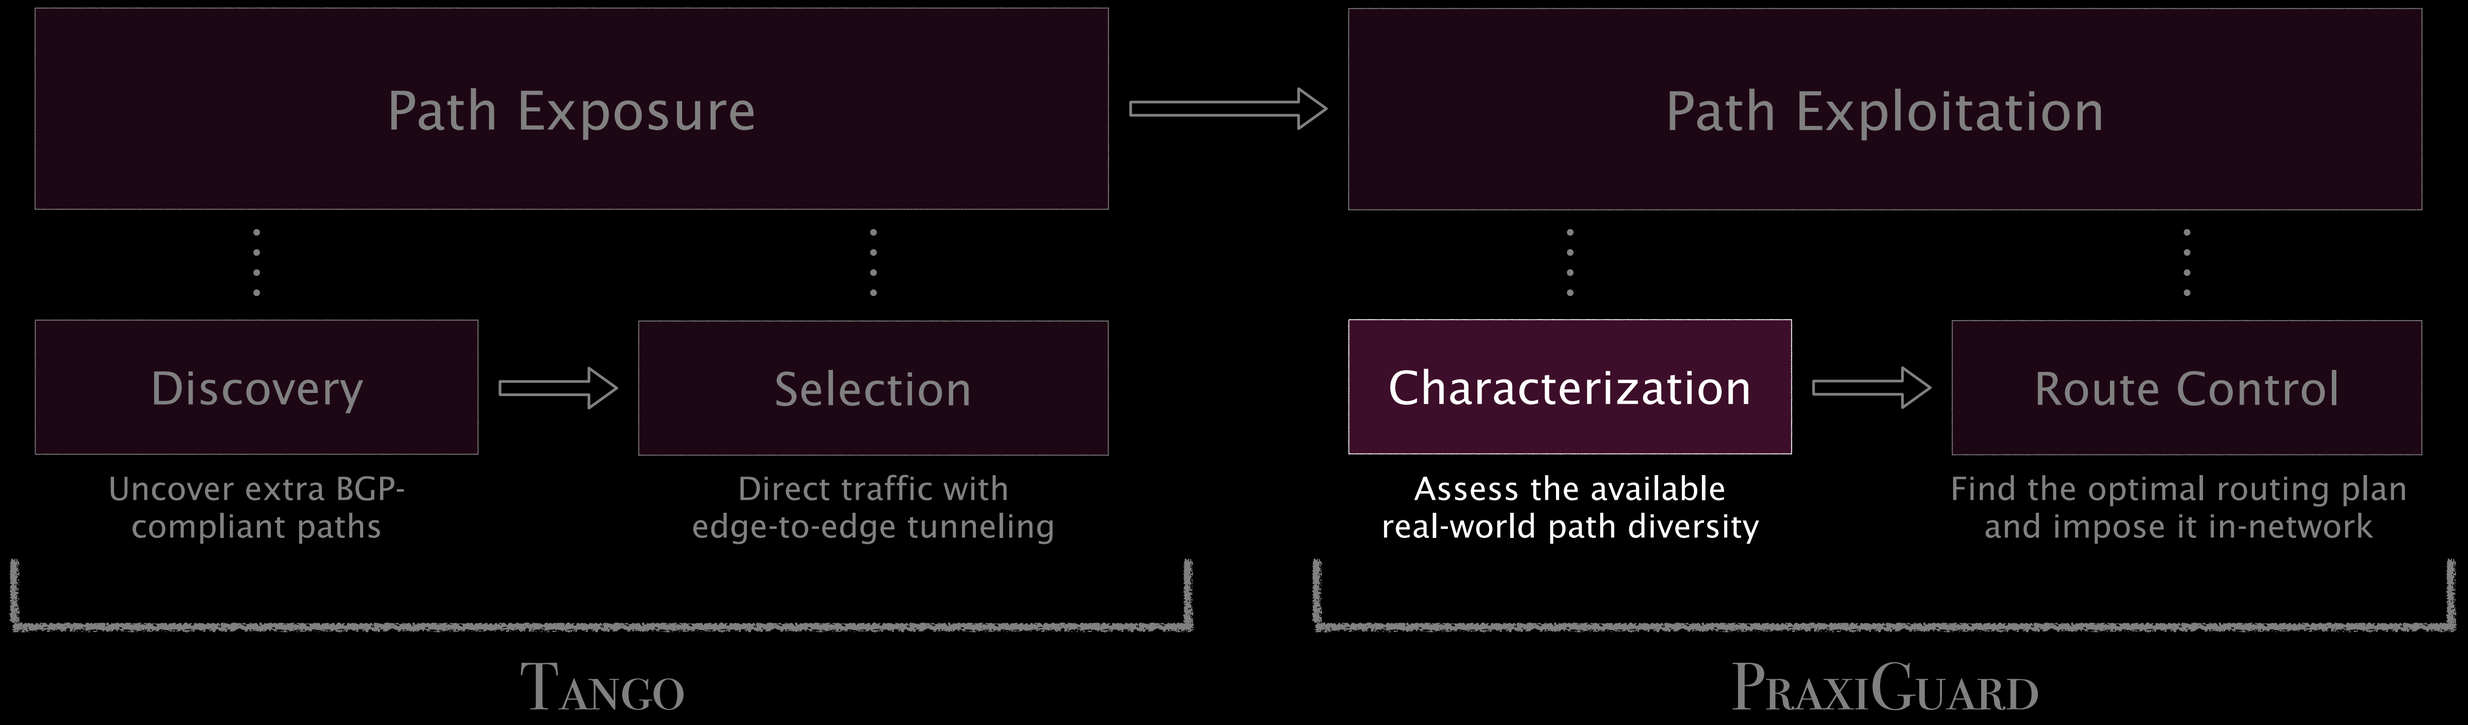
Path Exposure
Path Exploitation
Discovery
Selection
Characterization
Route Control
Uncover extra BGP-compliant paths
Direct traffic with edge-to-edge tunneling
Assess the available real-world path diversity
Find the optimal routing plan and impose it in-network
TANGO
PRAXIGUARD

For a given topology, what is the opportunity?

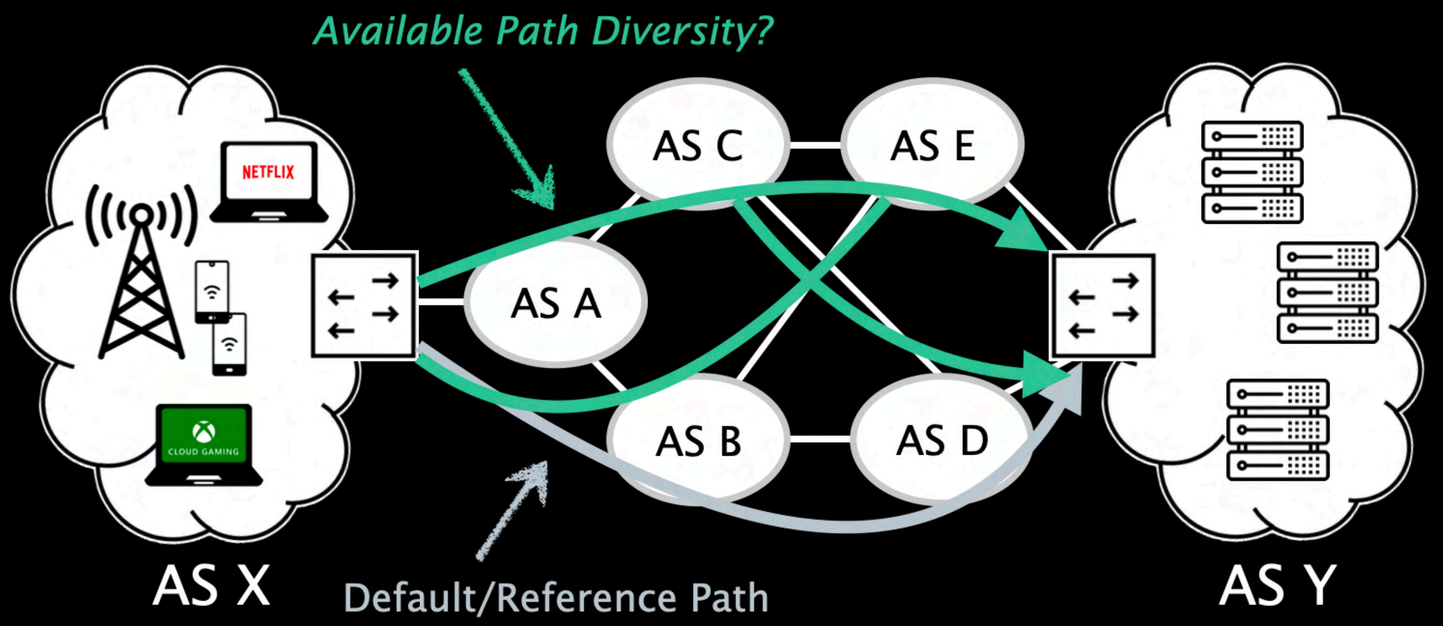

For a given topology, what is the opportunity?

Interdomain Path Options

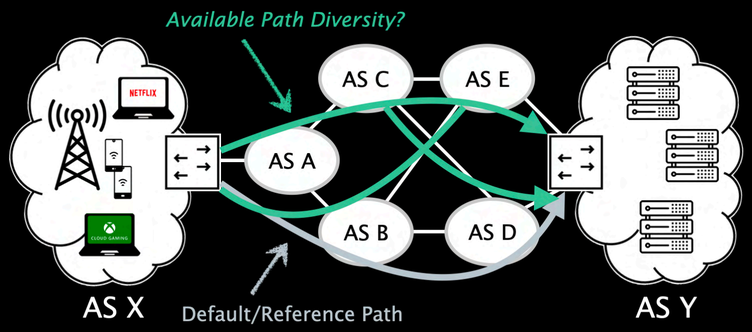

TANGO [1] - exposes paths from today's Internet

Dataset I: CAIDA, upper-bound Internet paths (simulated), 49 AS pairs

Dataset II: Vultr, real-world Internet paths (measured), 506 AS pairs

SCION [2] - exposes paths from SCION deployments

Dataset III: SCION, publicly reported paths, 111 AS pairs

All datasets are reduced to comparable internal AS-level path sets before analysis.

[1] Birge-Lee, ***Yoo***, et al. *Tango: Secure Collaborative Route Control Across the Public Internet*. NSDI, 2024.

[2] Zhang, et al. *SCION: Scalability, Control, and Isolation on Next-Generation Networks*. IEEE S&P, 2011.

Meaningful AS-level path diversity is available

1. Do we have options?

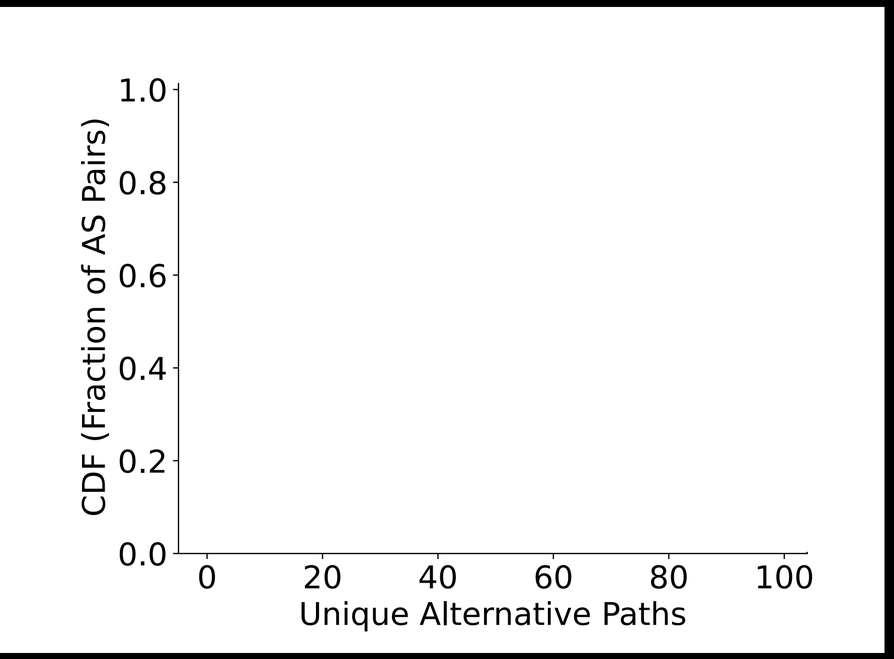

Meaningful AS-level path diversity is available

1. Do we have options?

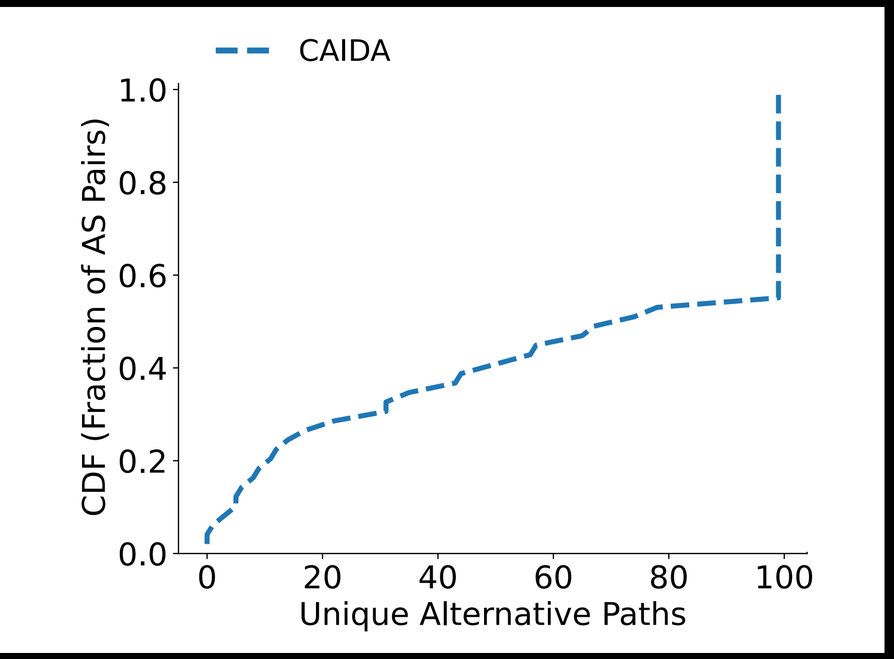

Meaningful AS-level path diversity is available

1. Do we have options?

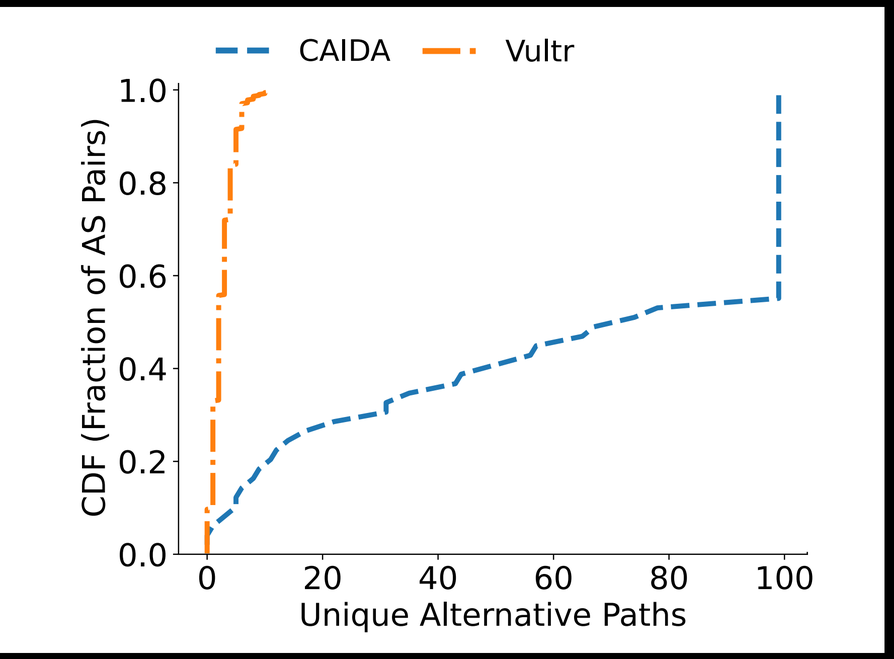

Meaningful AS-level path diversity is available

1. Do we have options?

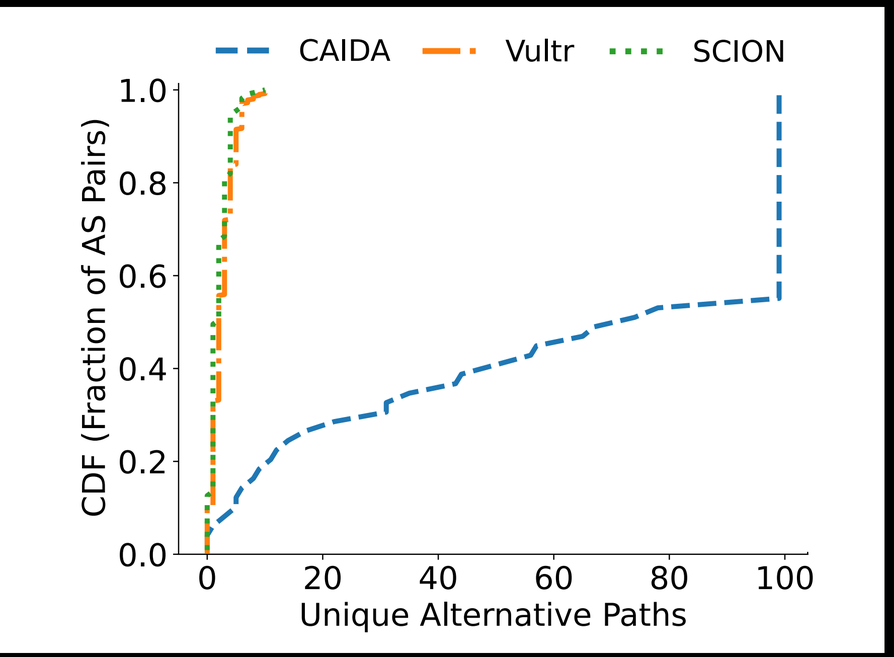

Meaningful AS-level path diversity is available

1. Do we have options?

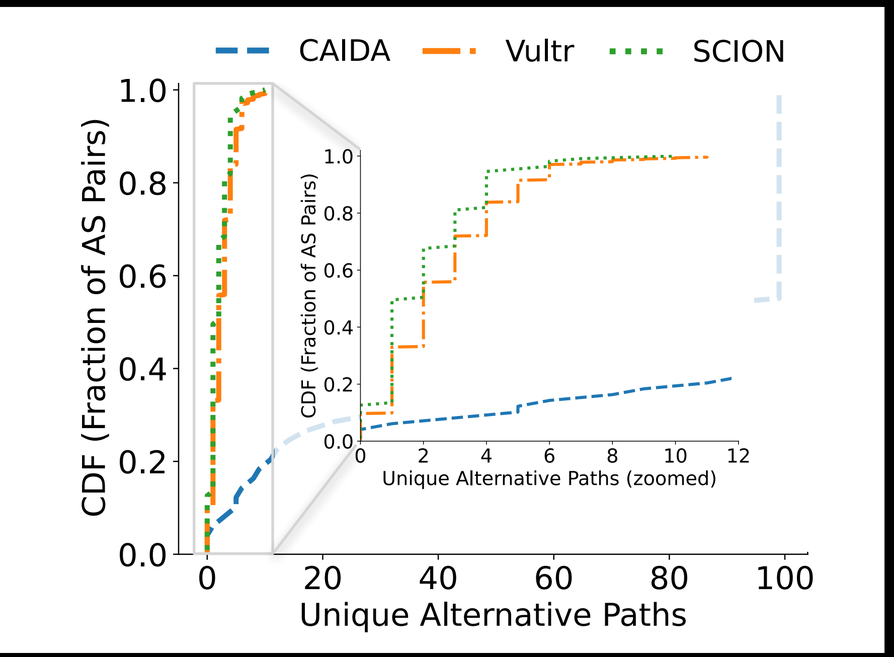

Meaningful AS-level path diversity is available

1. Do we have options?

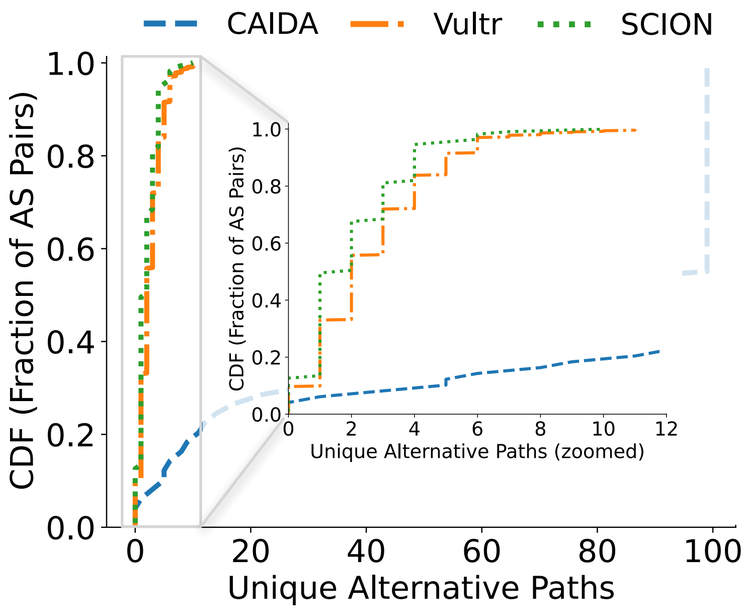

2. Are they meaningfully different?

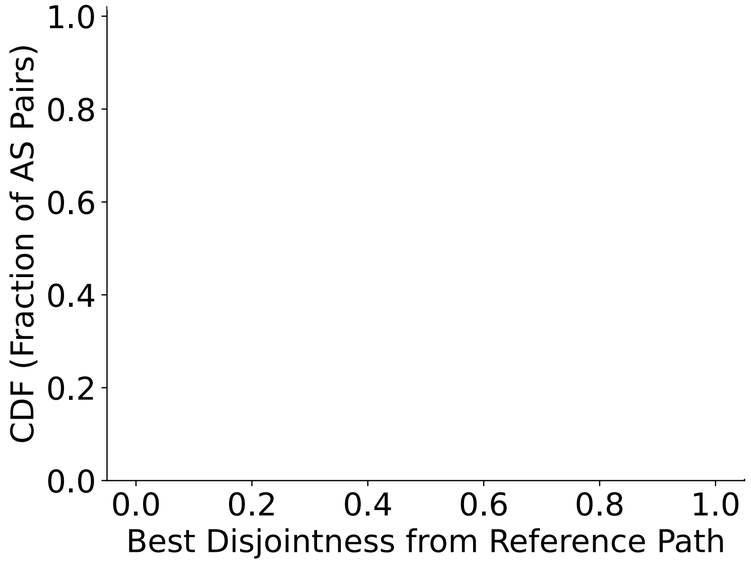

Meaningful AS-level path diversity is available

1. Do we have options?

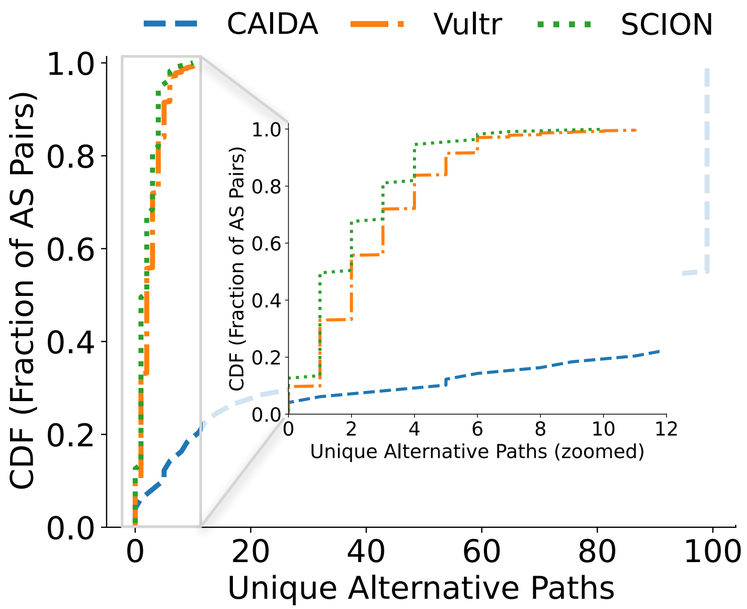

2. Are they meaningfully different?

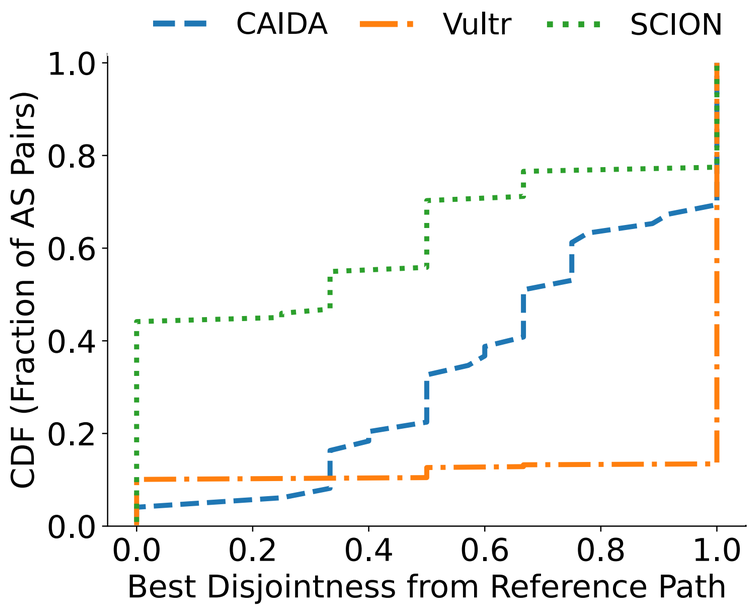

*For SCION, paths are collapsed to AS-level paths before comparison; collapsed paths with no intermediate ASes (“empty” AS-level paths) are counted as fully disjoint.

Meaningful AS-level path diversity is available

1. Do we have options?

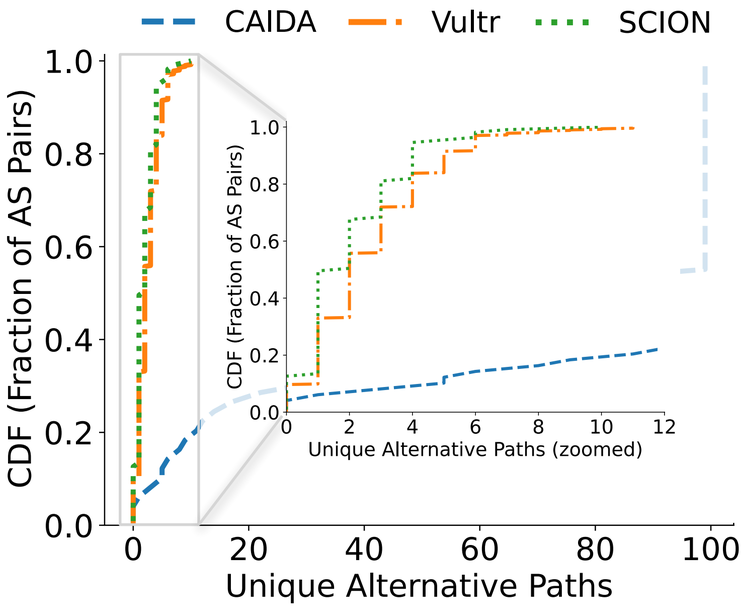

2. Are they meaningfully different?

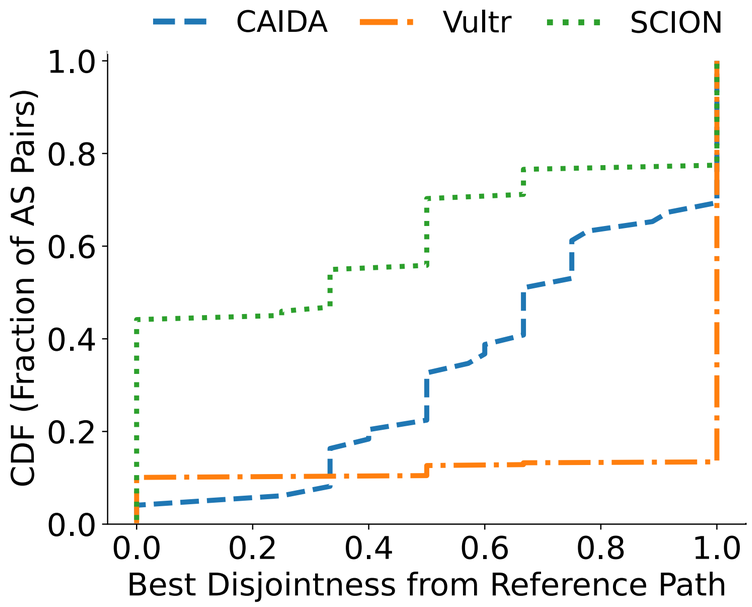

3. What do they cost in hop overhead?

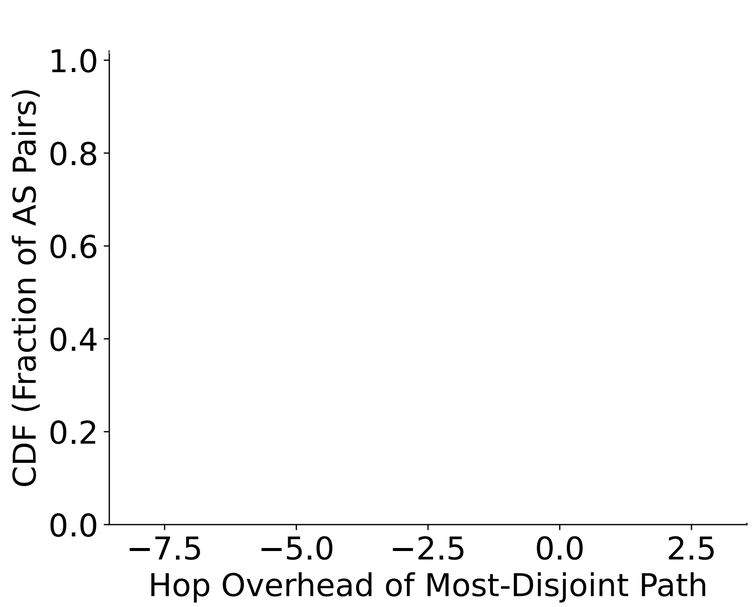

Meaningful AS-level path diversity is available

1. Do we have options?

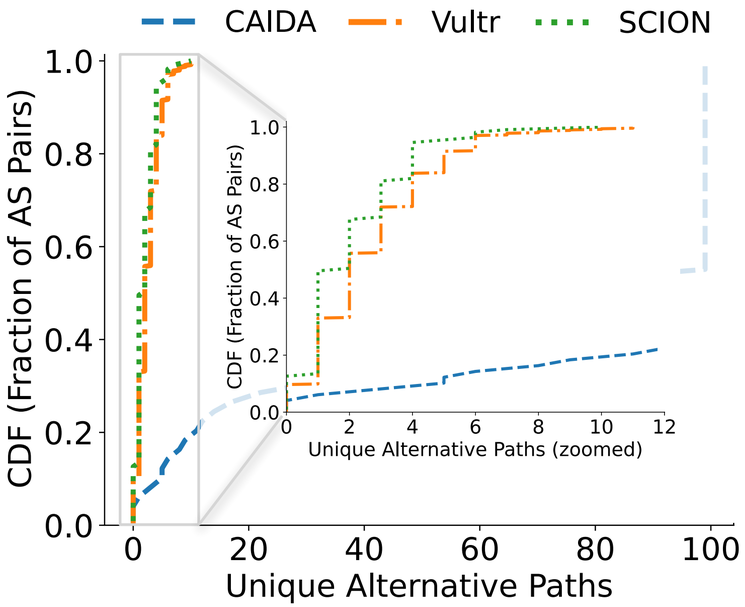

2. Are they meaningfully different?

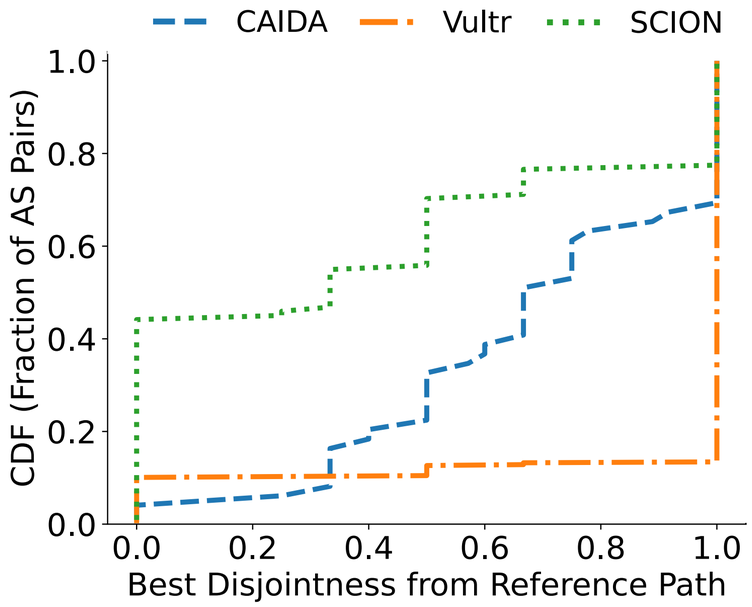

3. What do they cost in hop overhead?

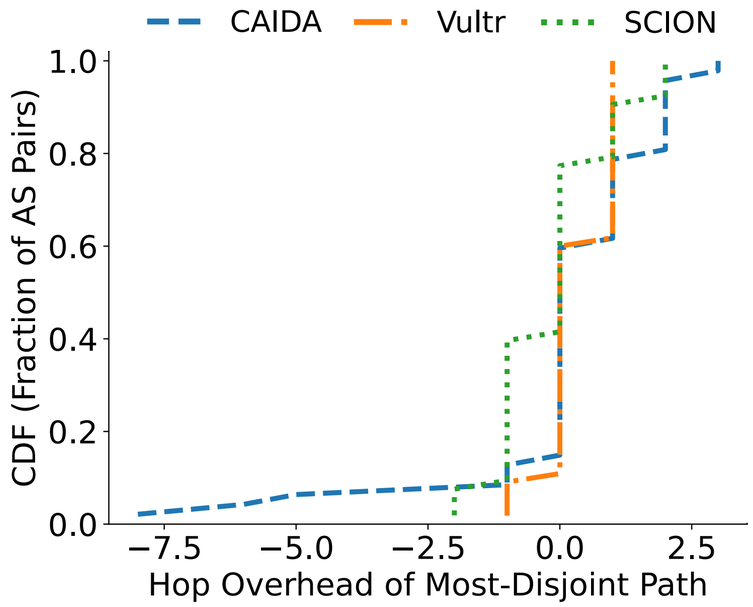

Meaningful AS-level path diversity is available

1. Do we have options?

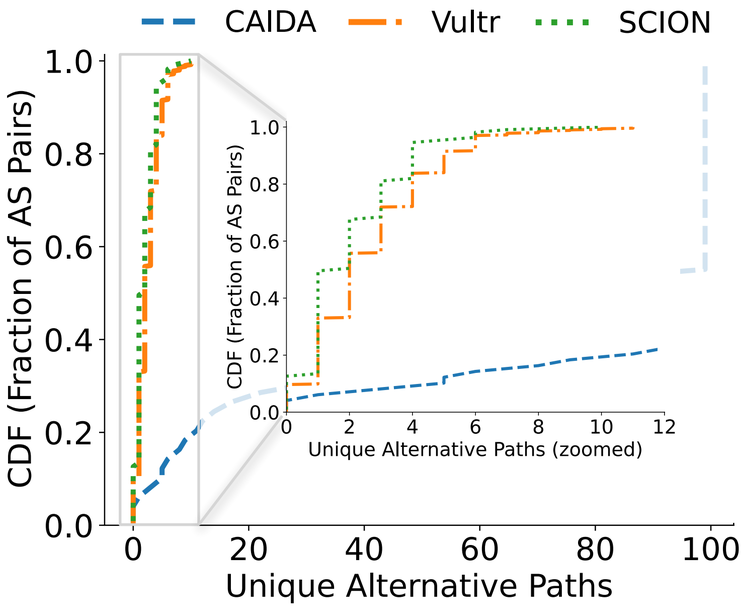

1. Many AS pairs have multiple path options

2. Are they meaningfully different?

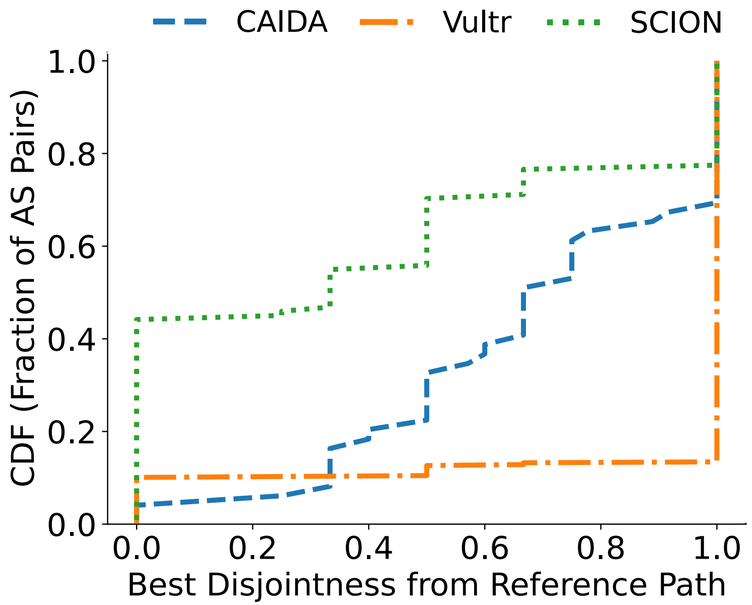

2. Alternatives are often substantially disjoint

3. What do they cost in hop overhead?

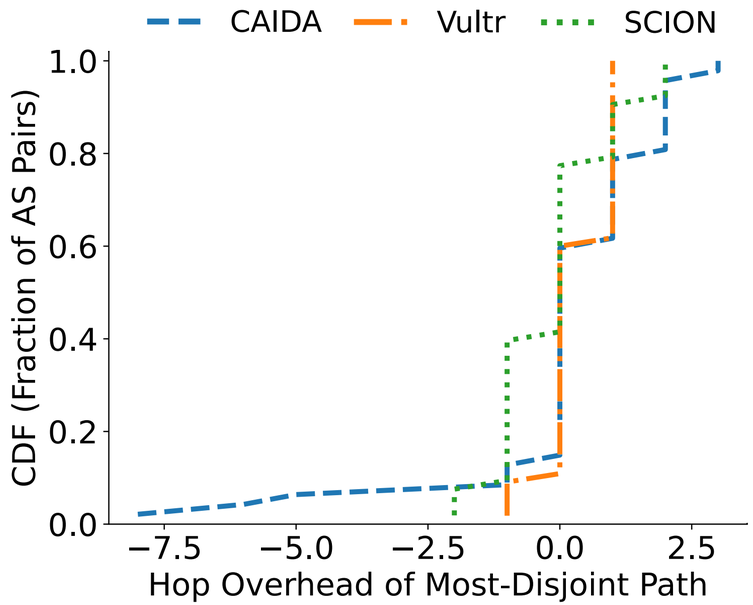

3. Most-disjoint paths add little hop overhead

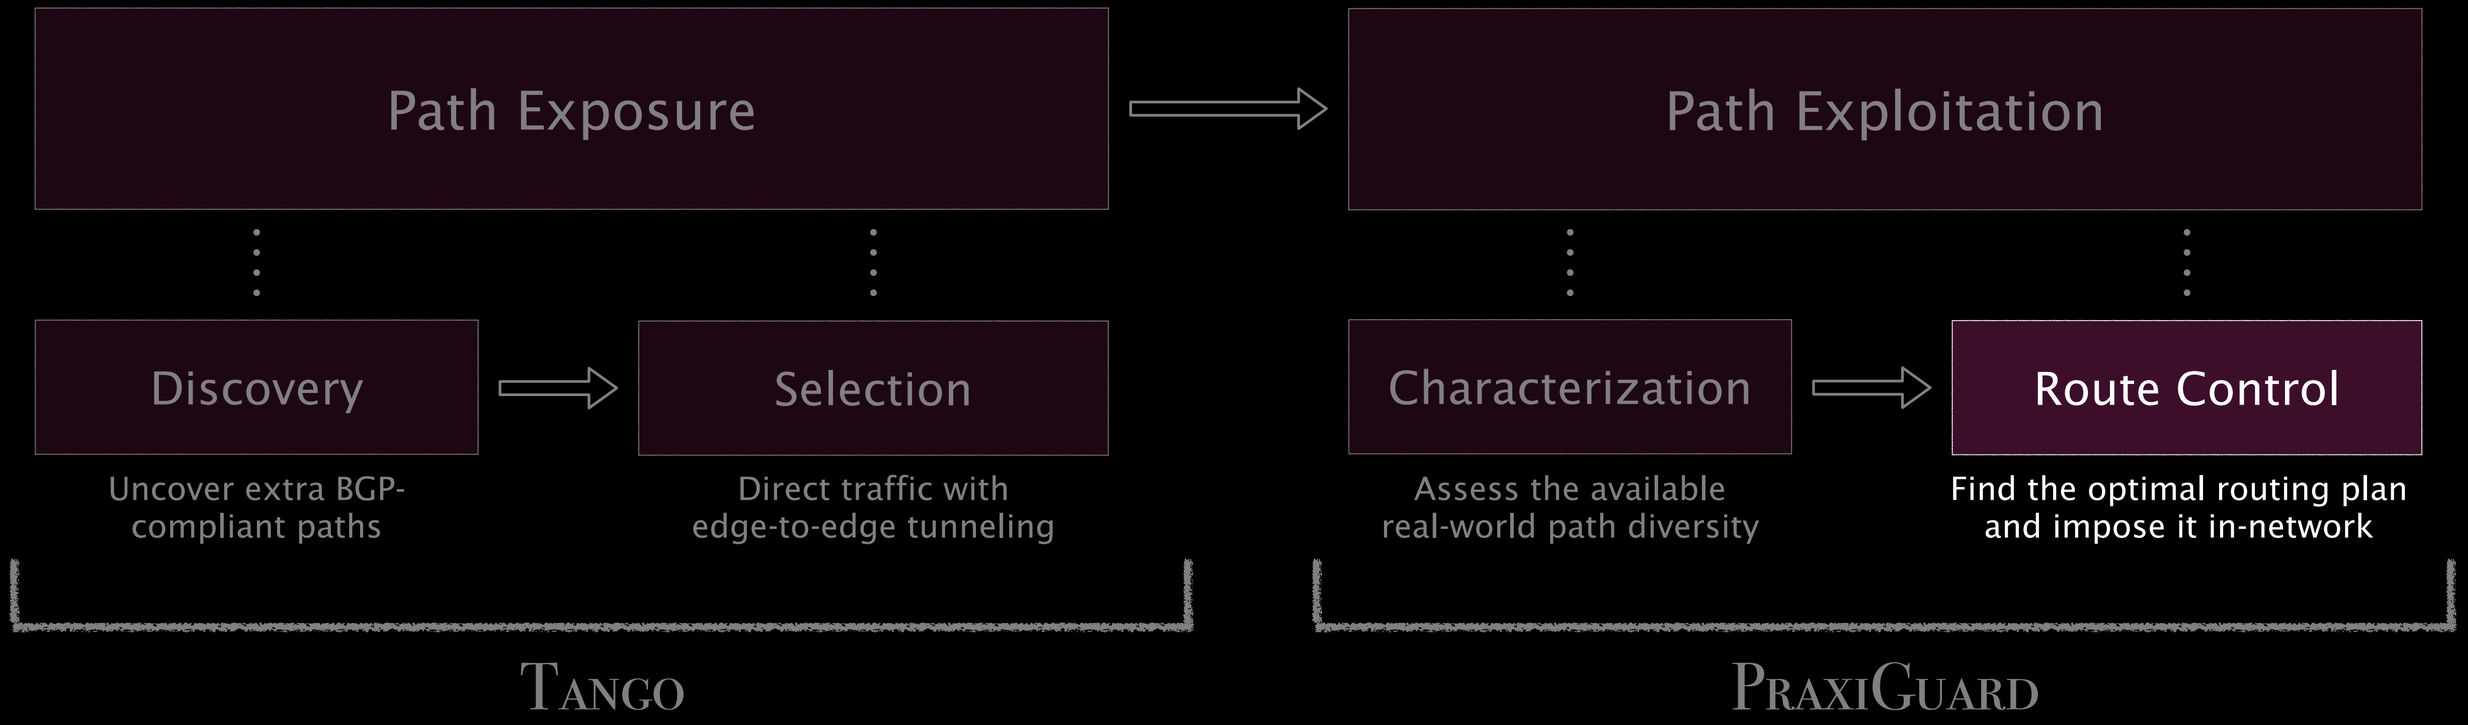
Path Exposure
Path Exploitation
Discovery
Selection
Characterization
Route Control
Uncover extra BGP-compliant paths
Direct traffic with edge-to-edge tunneling
Assess the available real-world path diversity
Find the optimal routing plan and impose it in-network
TANGO
PRAXIGUARD

How to assess the privacy offered by a path?

Our threat model assumes AS-level adversaries

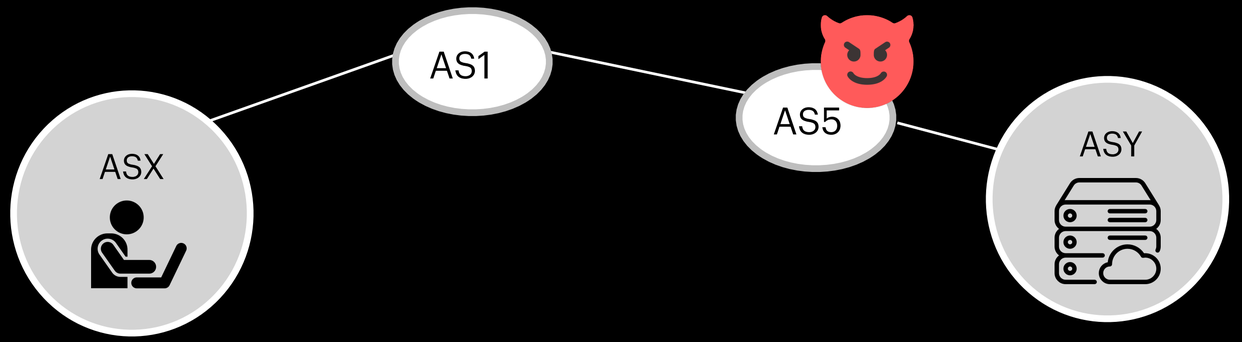

Our threat model assumes AS-level adversaries

a single-path (default) setting allows AS5 to observe all traffic

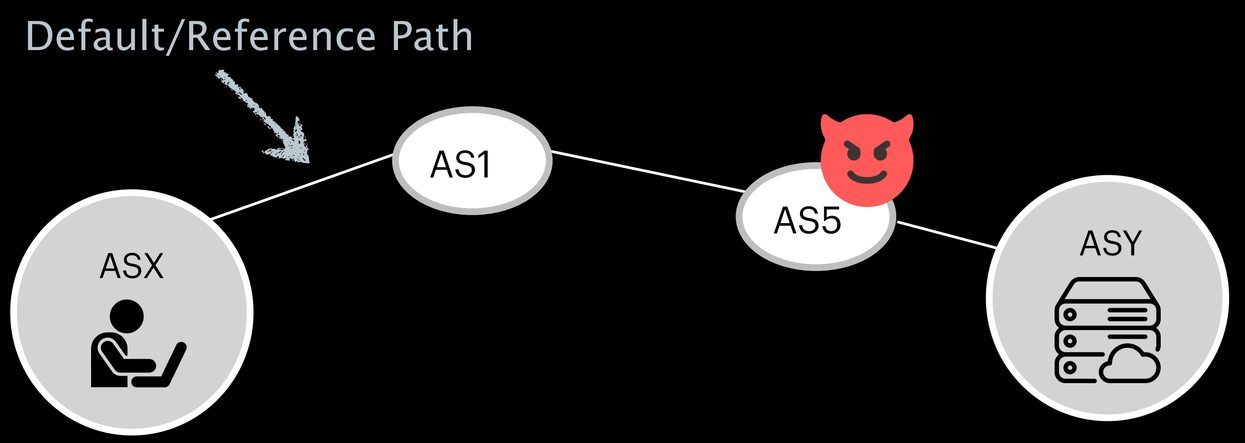

Example privacy application: website fingerprinting

classifier predicts which webpages a user visits, given network traces

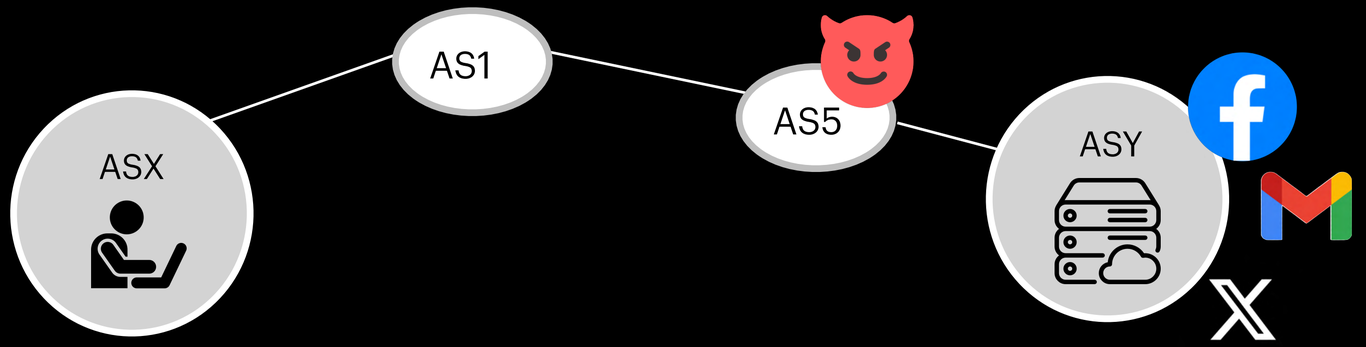

Traffic splitting at the network level?

to reduce unwanted exposure using additional available paths

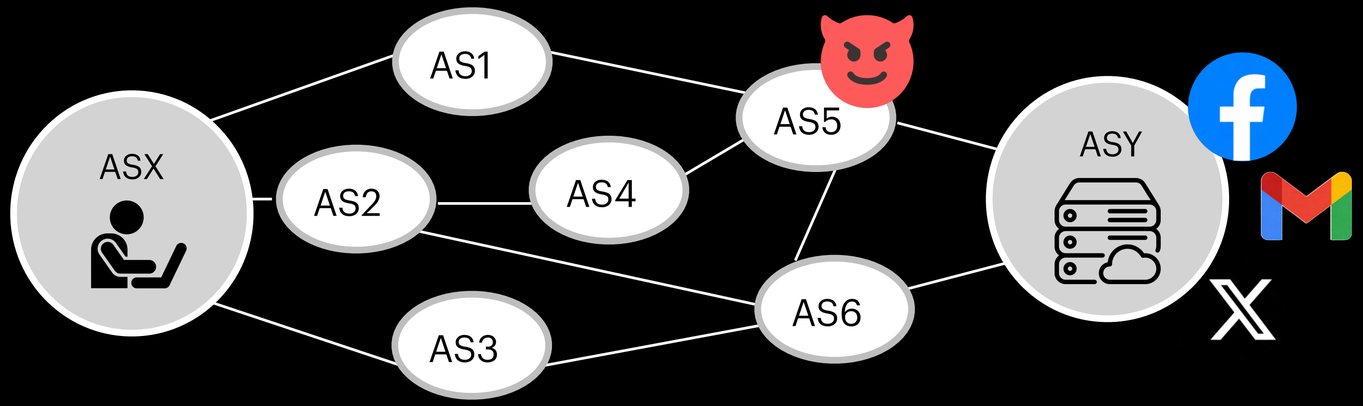

But uniform traffic splitting is not enough

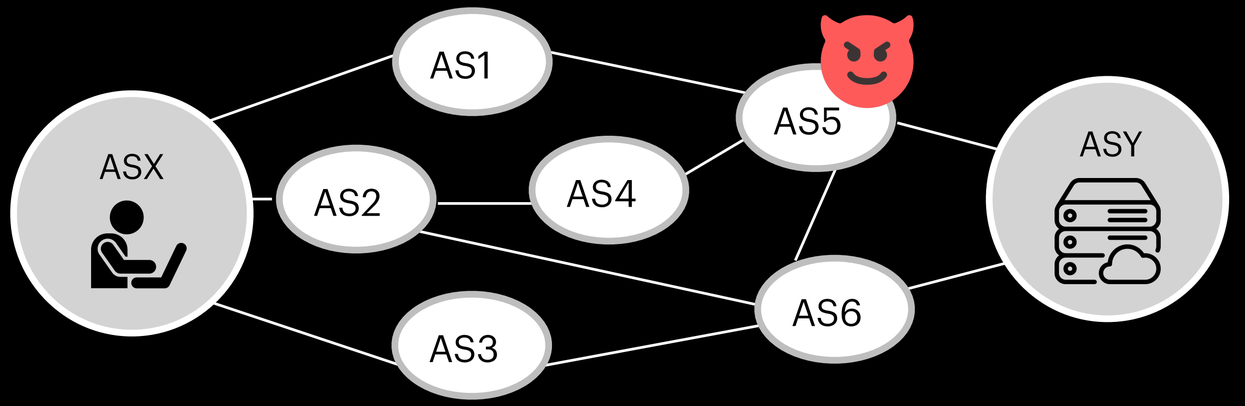

But uniform traffic splitting is not enough

since multiple paths can still traverse a single AS

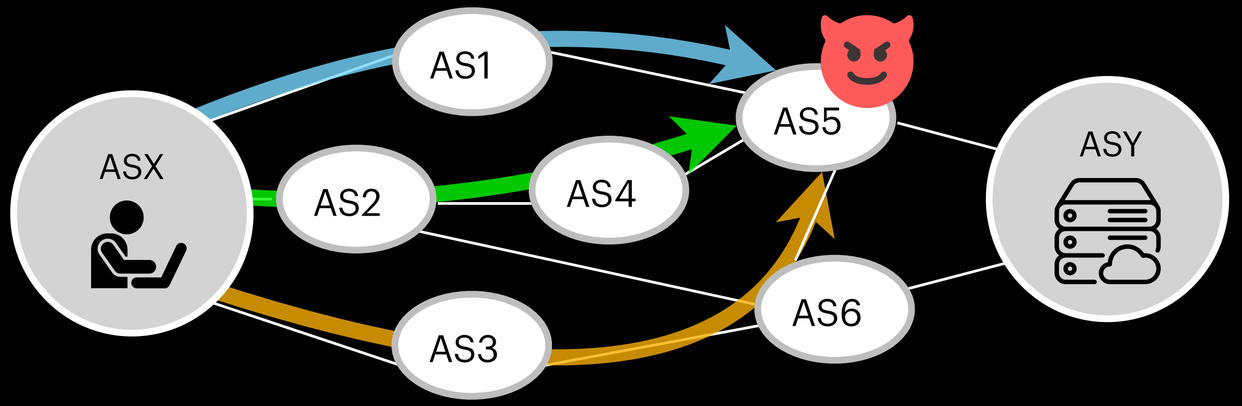

How should traffic be split over available paths?

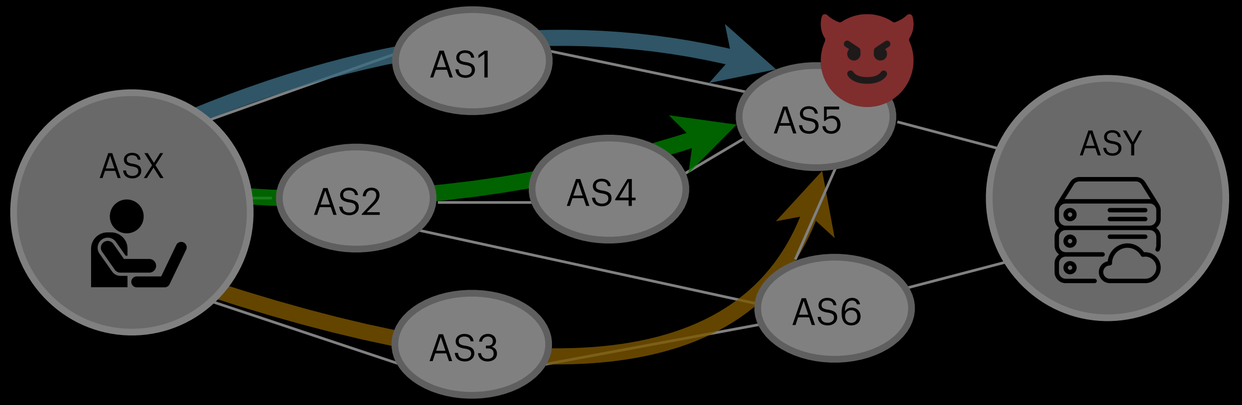

Optimization to determine the traffic split

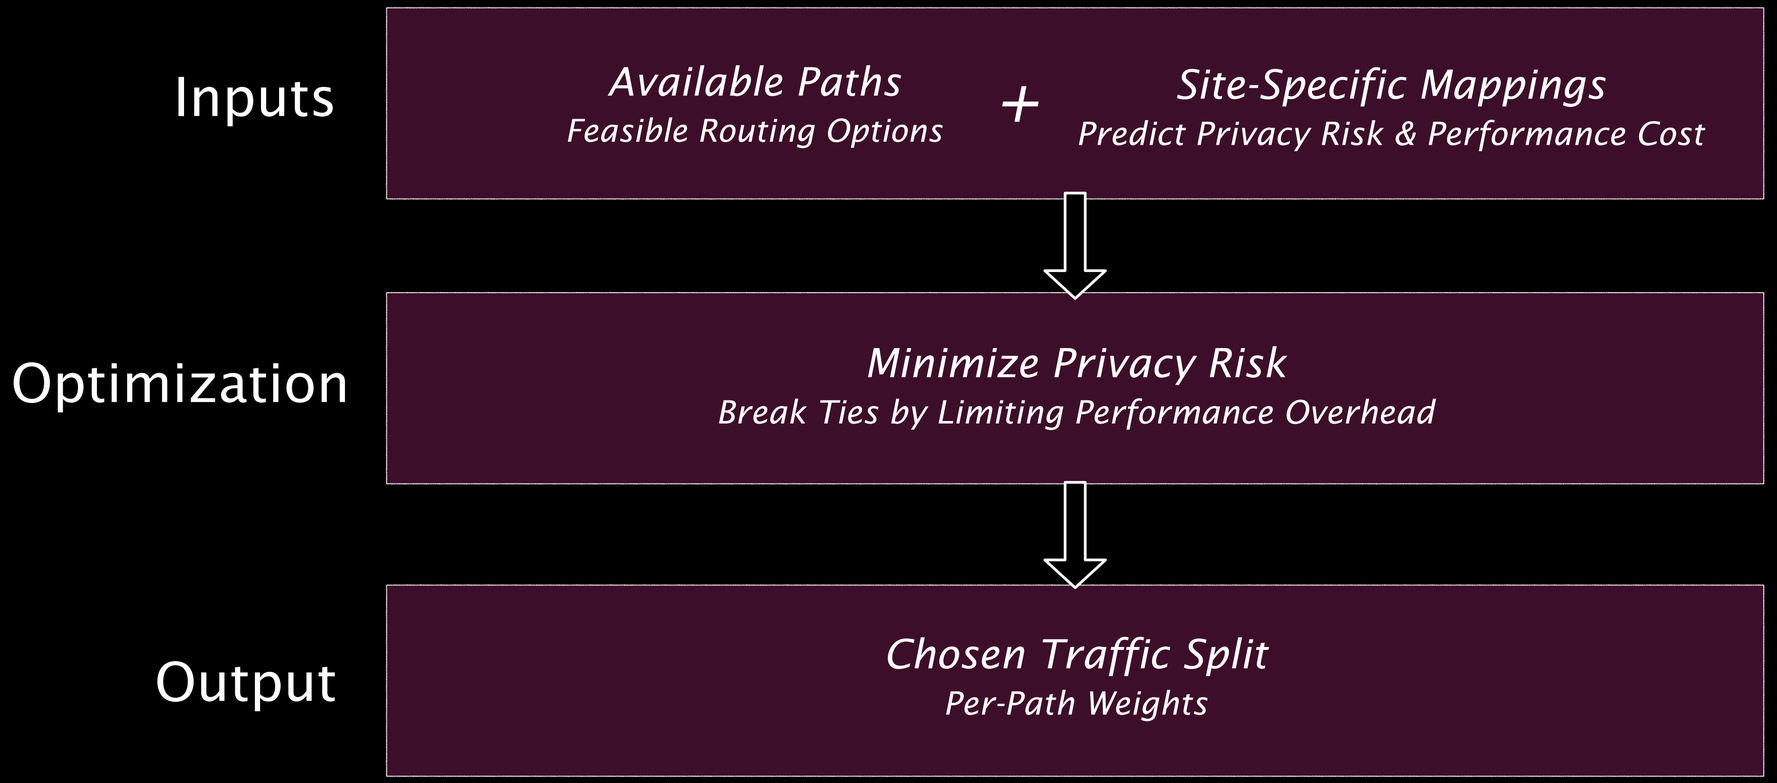

Optimization to determine the traffic split

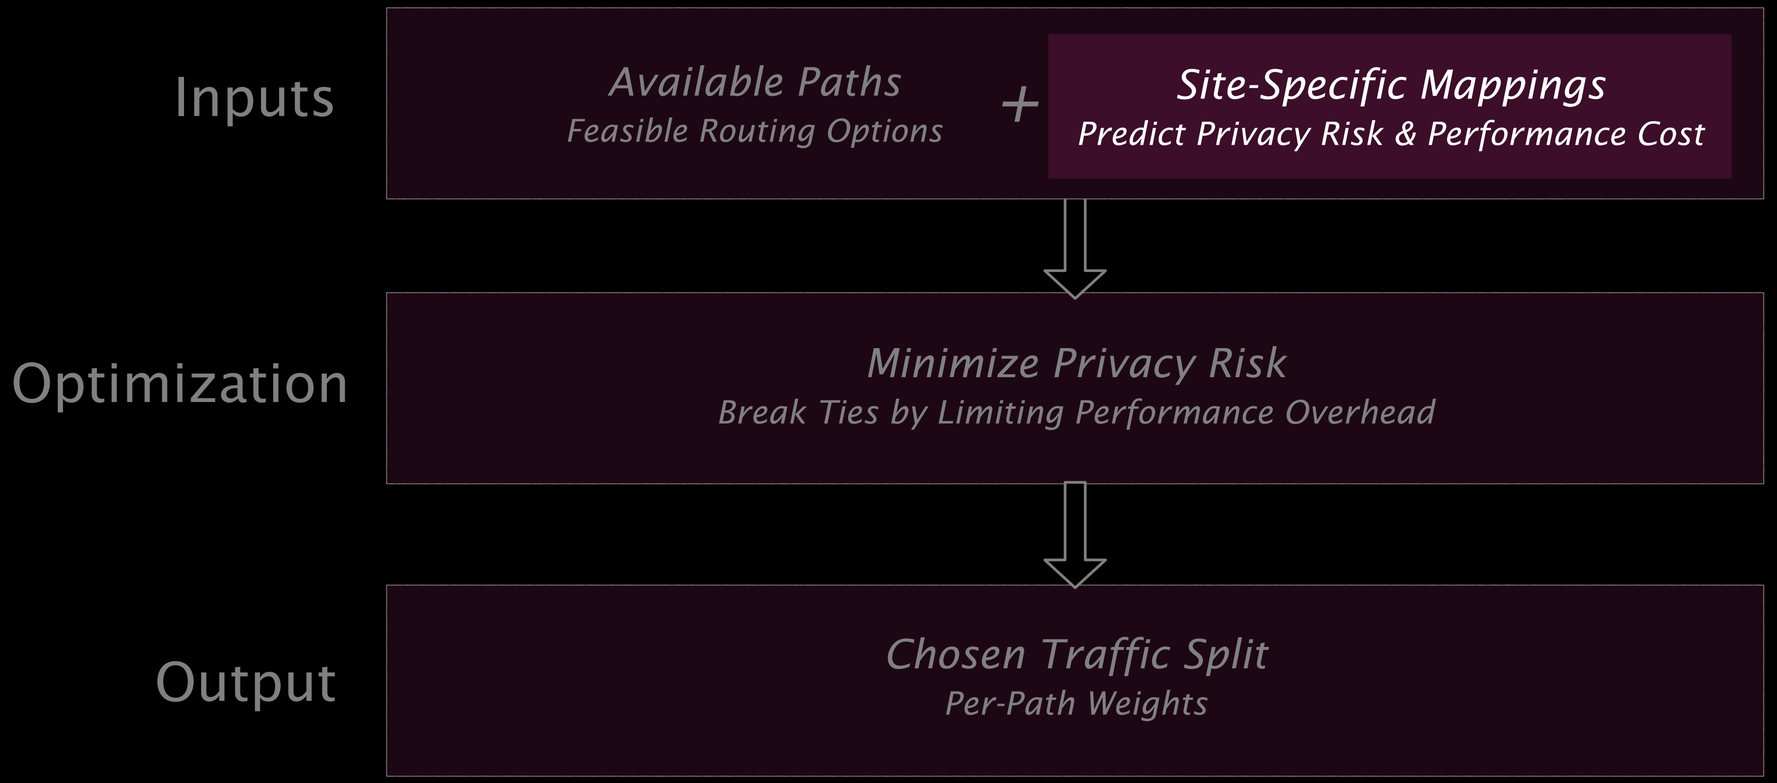

Empirically measuring site-specific mappings

from traffic splits to privacy risk and performance cost

Sites (10 of top 100)

Domain Name	Number of Flows	Baseline RTT (ms)
amazon.com	114	23.24
discord.com	11	6.31
instagram.com	14	7.31
linkedin.com	38	8.36
netflix.com	68	8.48
pinterest.com	36	7.38
spotify.com	110	8.15
twitch.tv	56	7.38
wordpress.com	37	8.05
youtube.com	22	9.02

Privacy-Risk Proxy

Traffic exposure => Classifier accuracy

Performance-Cost Proxy

Path delay => Page Load Time (PLT) overhead

Mapping traffic exposure to privacy risk

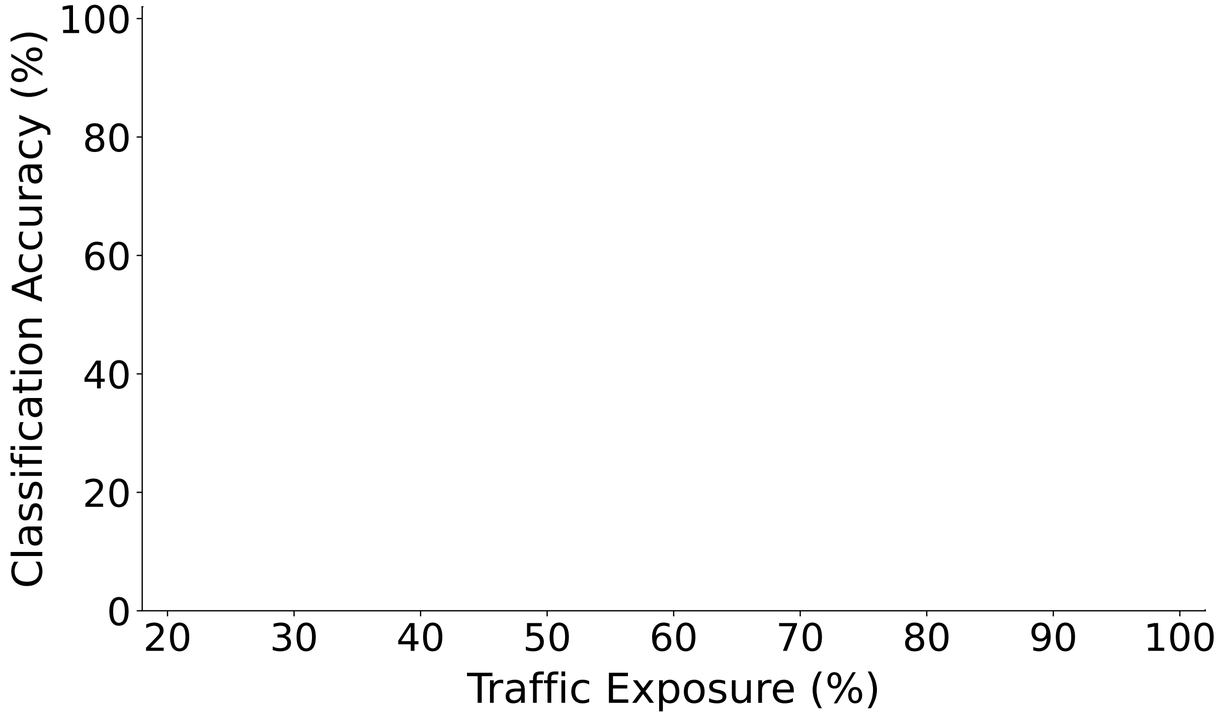

Mapping traffic exposure to privacy risk

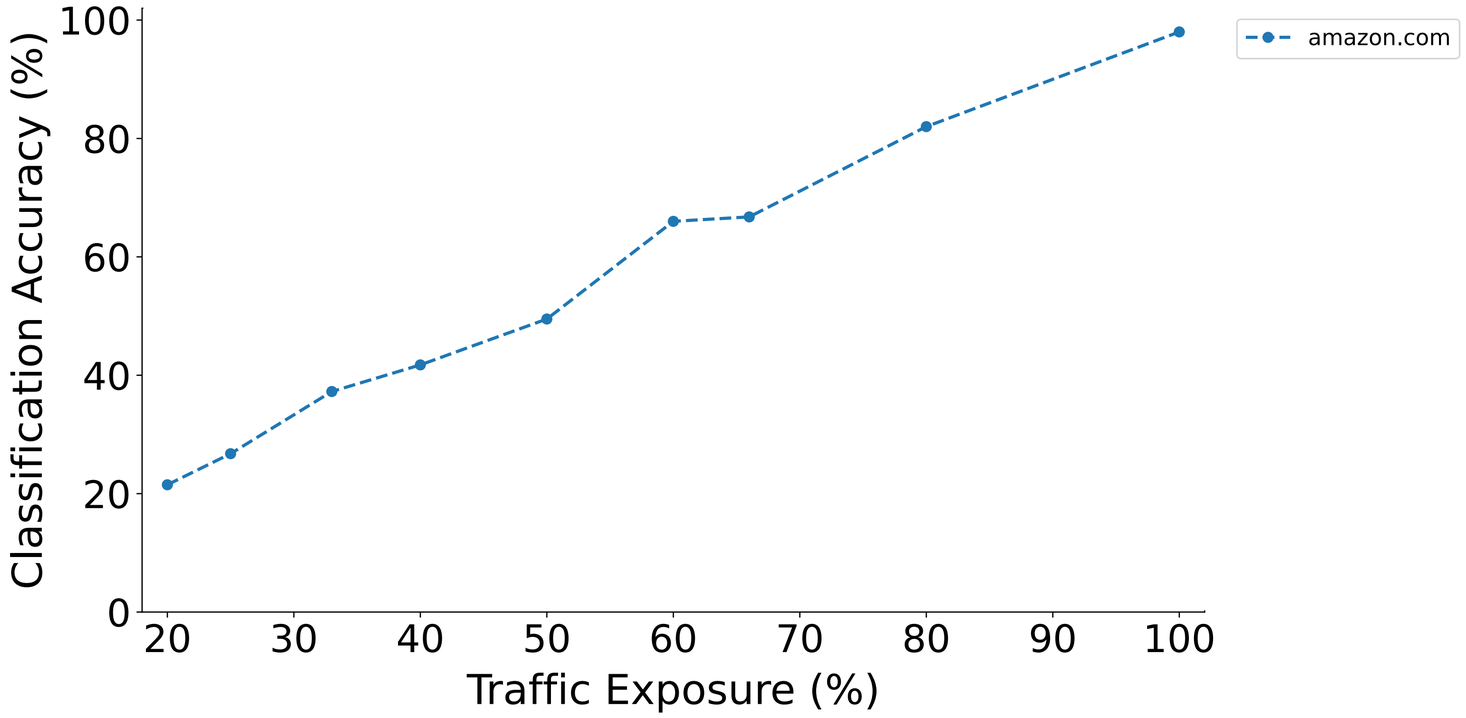

Across configurations, the *LaserBeak* [*] attack classifier shows consistent degradation in accuracy as traffic exposure decreases

[*] Matthews, et al. *Laserbeak: Evolving Website Fingerprinting Attacks With Attention and Multi-Channel Feature Representation.* IEEE TIFS, 2024.

Mapping traffic exposure to privacy risk

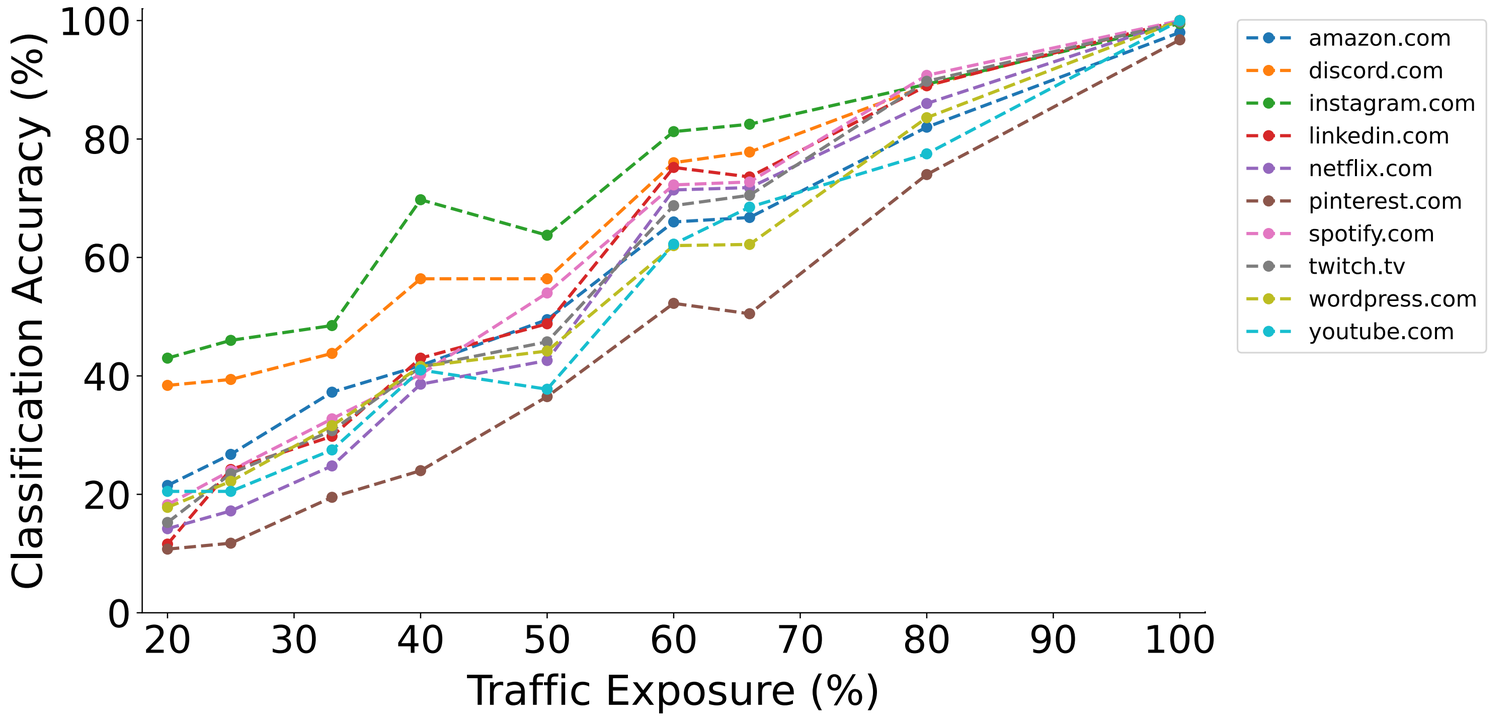

Across configurations, the *LaserBeak* [*] attack classifier shows consistent degradation in accuracy as traffic exposure decreases: slopes vary per site.

[*] Matthews, et al. *Laserbeak: Evolving Website Fingerprinting Attacks With Attention and Multi-Channel Feature Representation.* IEEE TIFS, 2024.

Mapping path delay to performance overhead

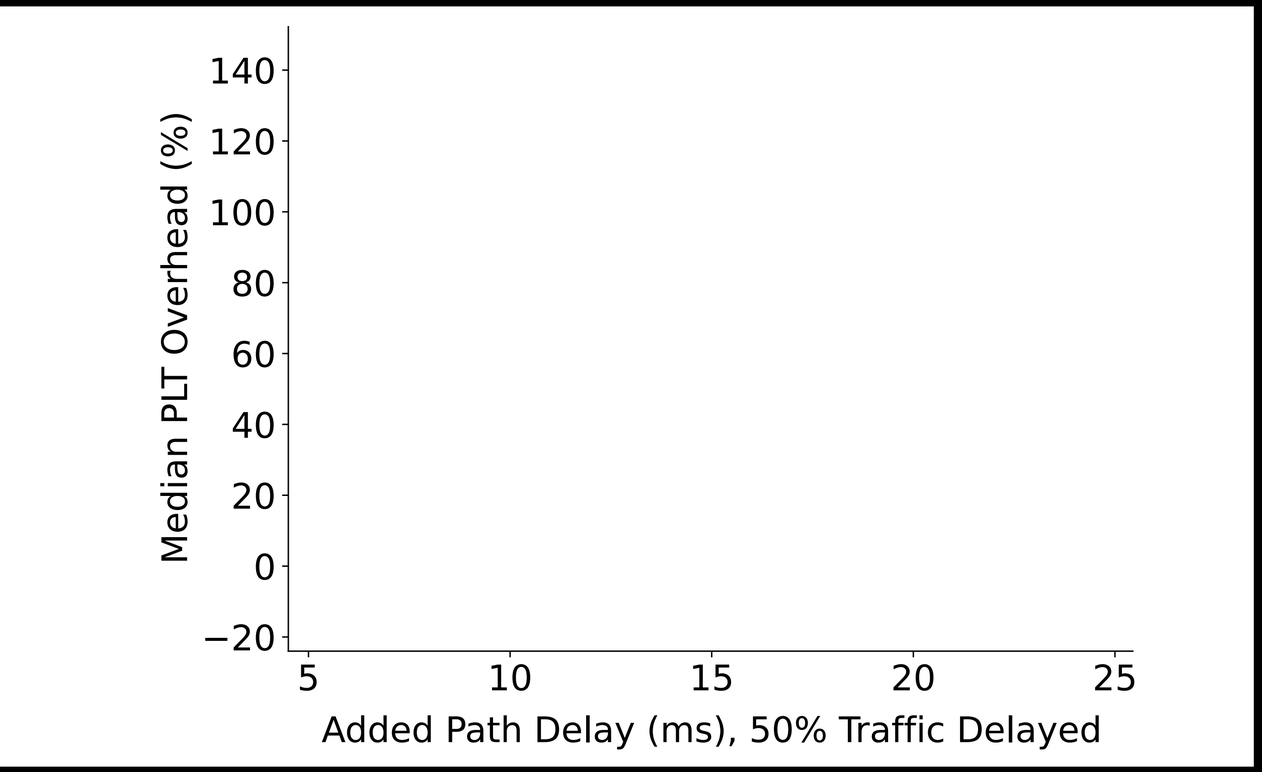

Mapping path delay to performance overhead

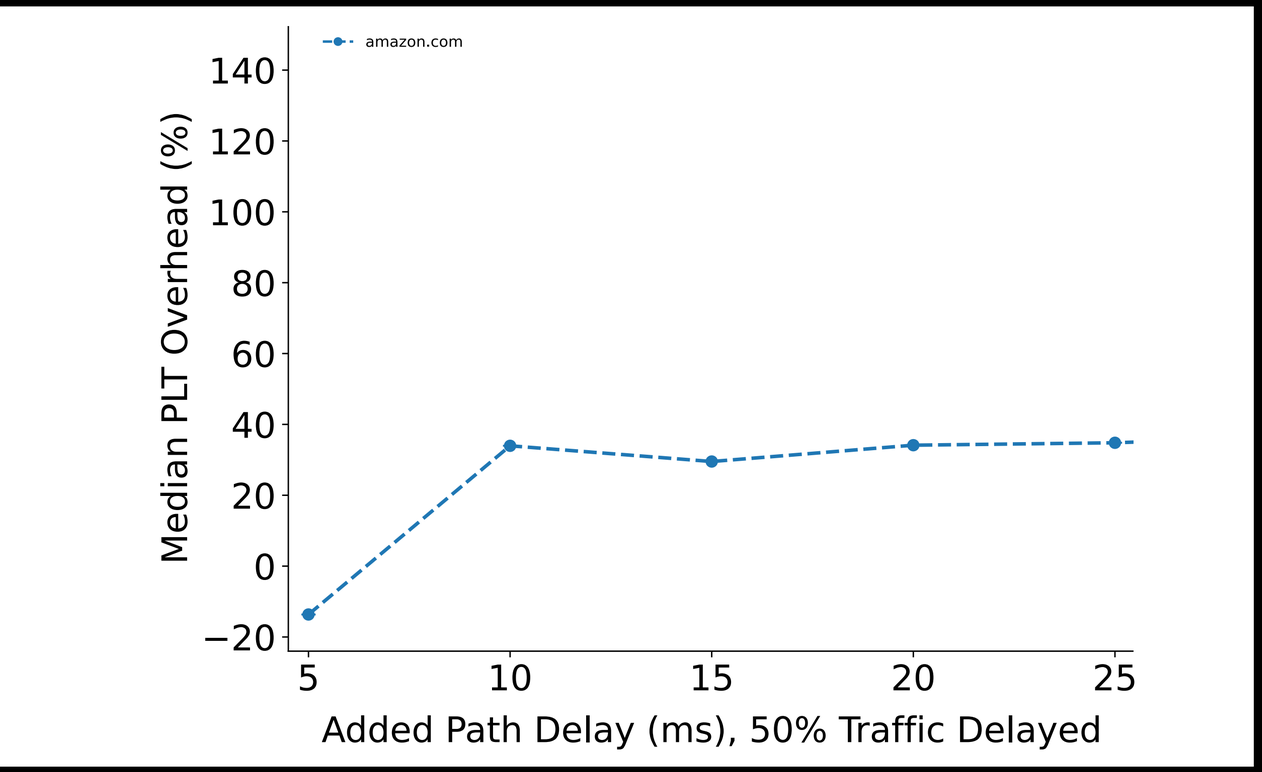

Mapping path delay to performance overhead

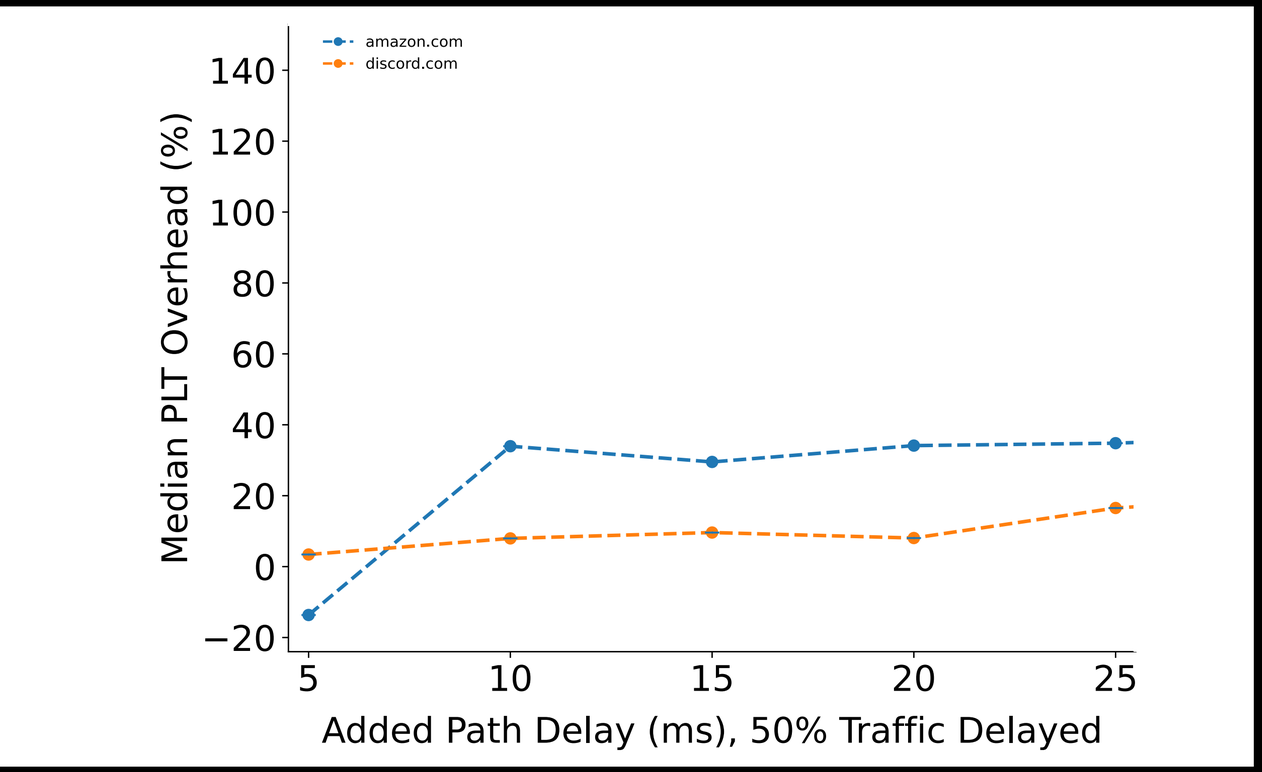

Mapping path delay to performance overhead

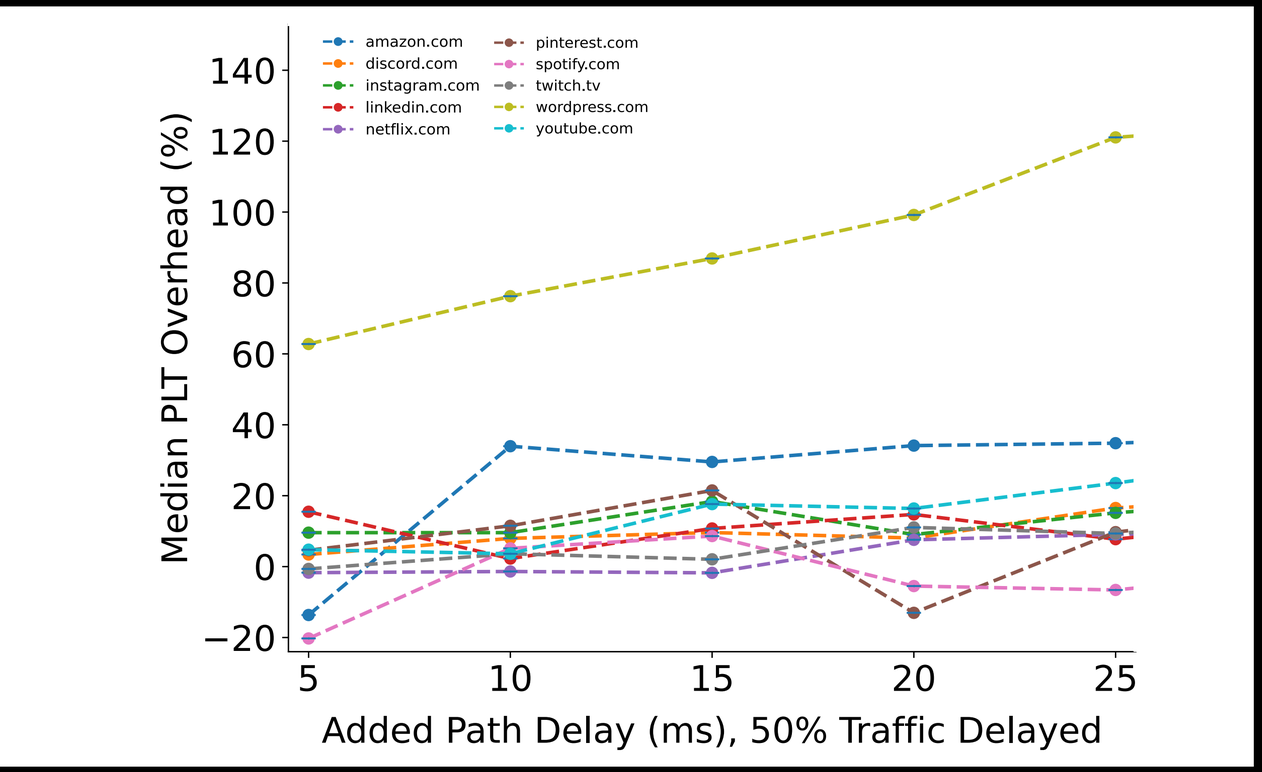

Mapping path delay to performance overhead

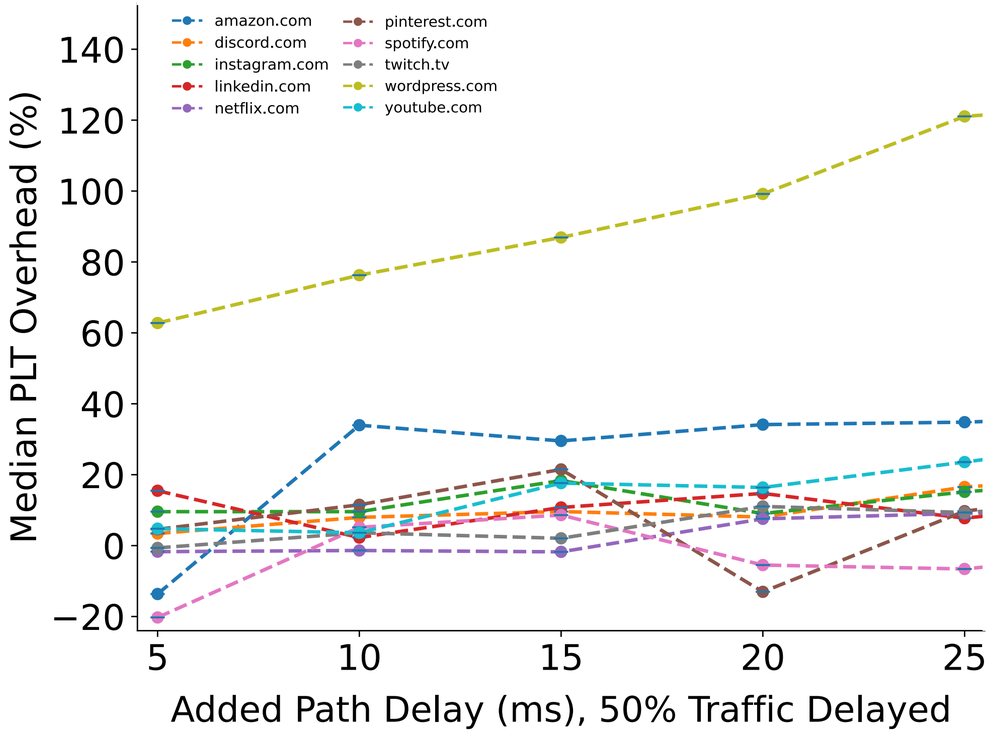

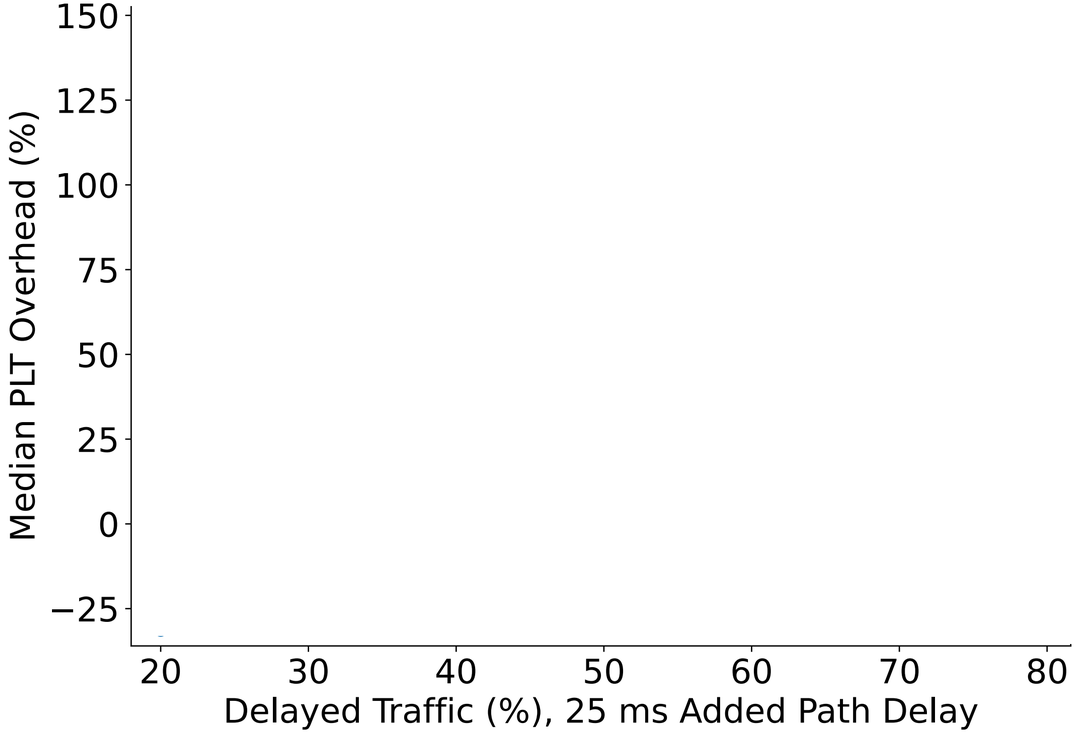

Mapping path delay to performance overhead

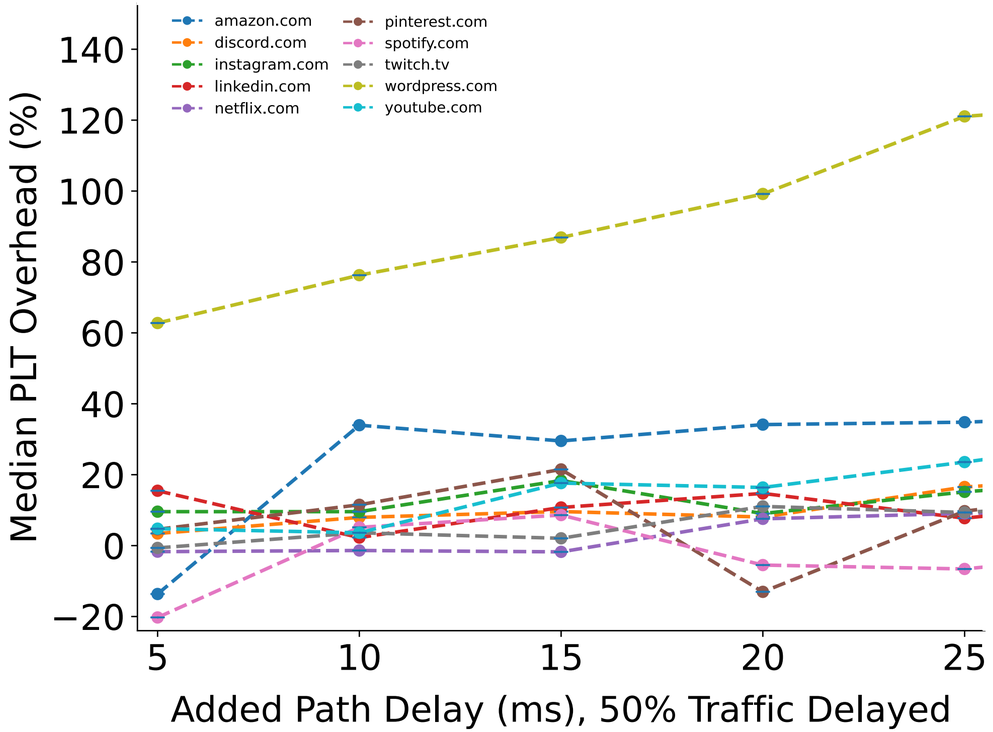

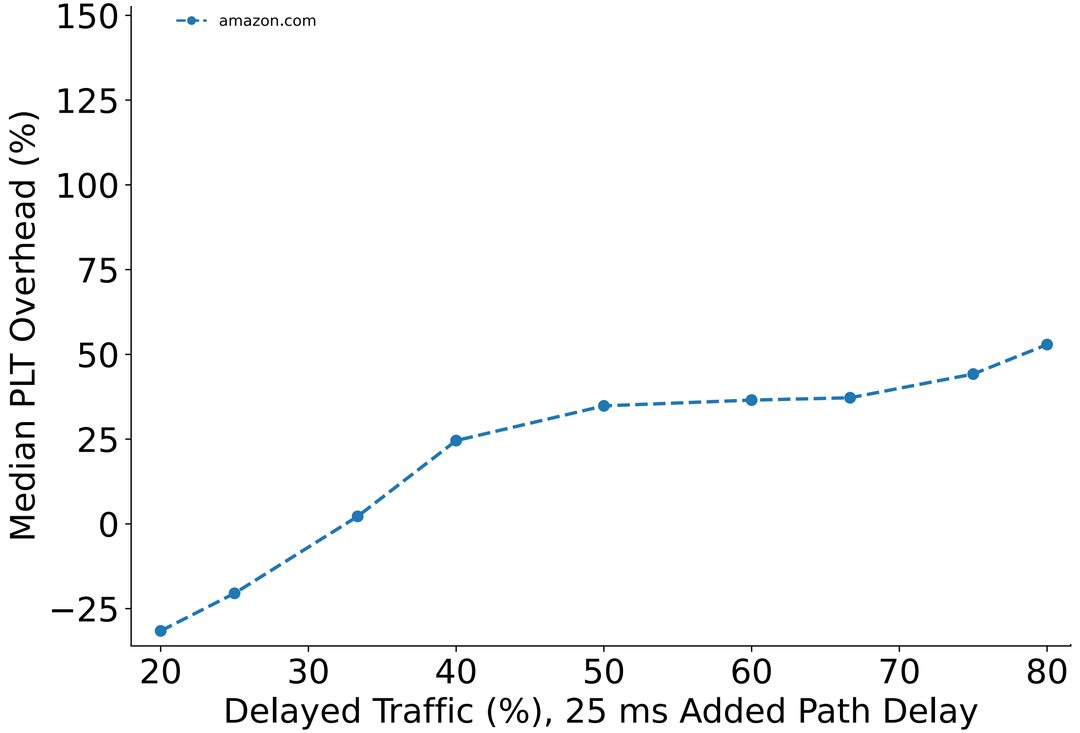

Mapping path delay to performance overhead

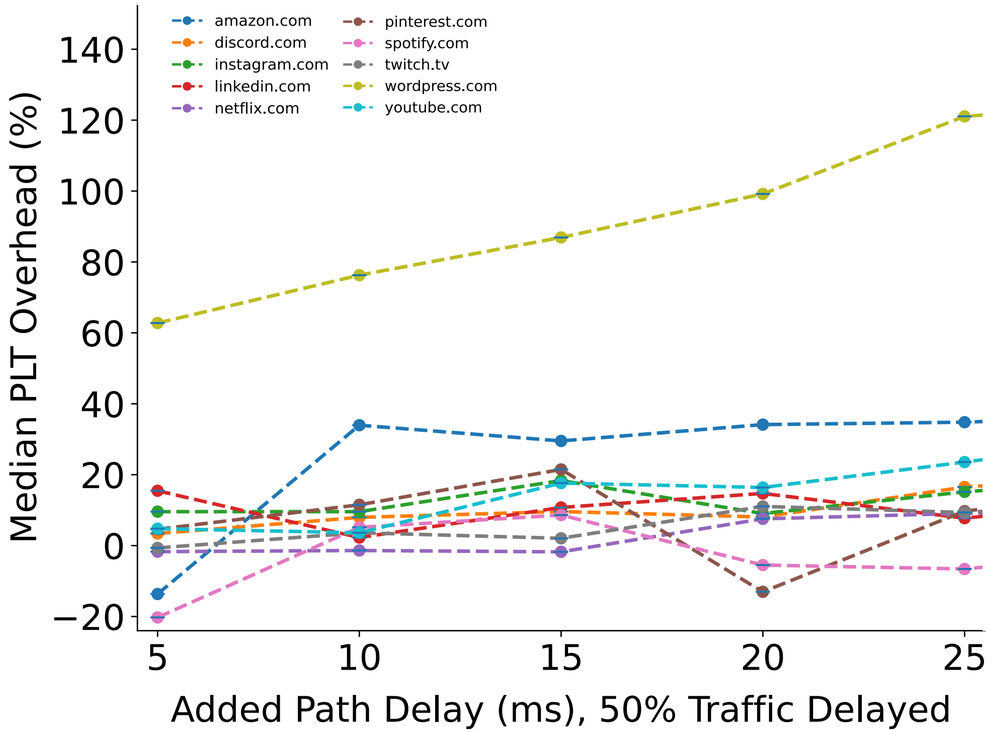

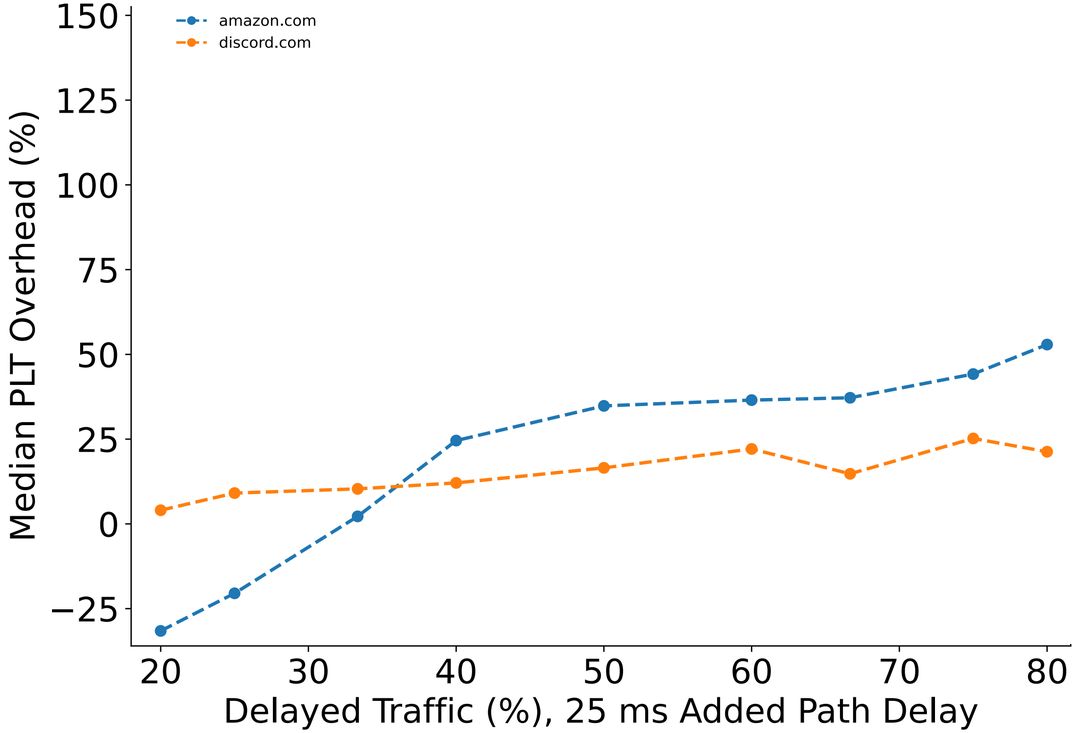

Mapping path delay to performance overhead

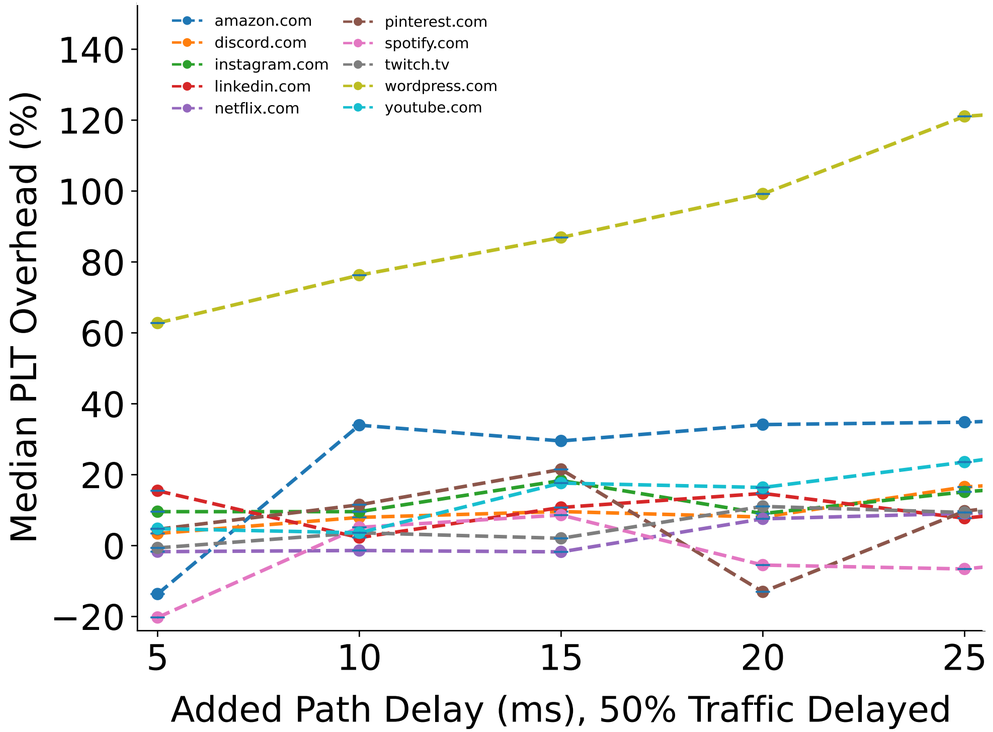

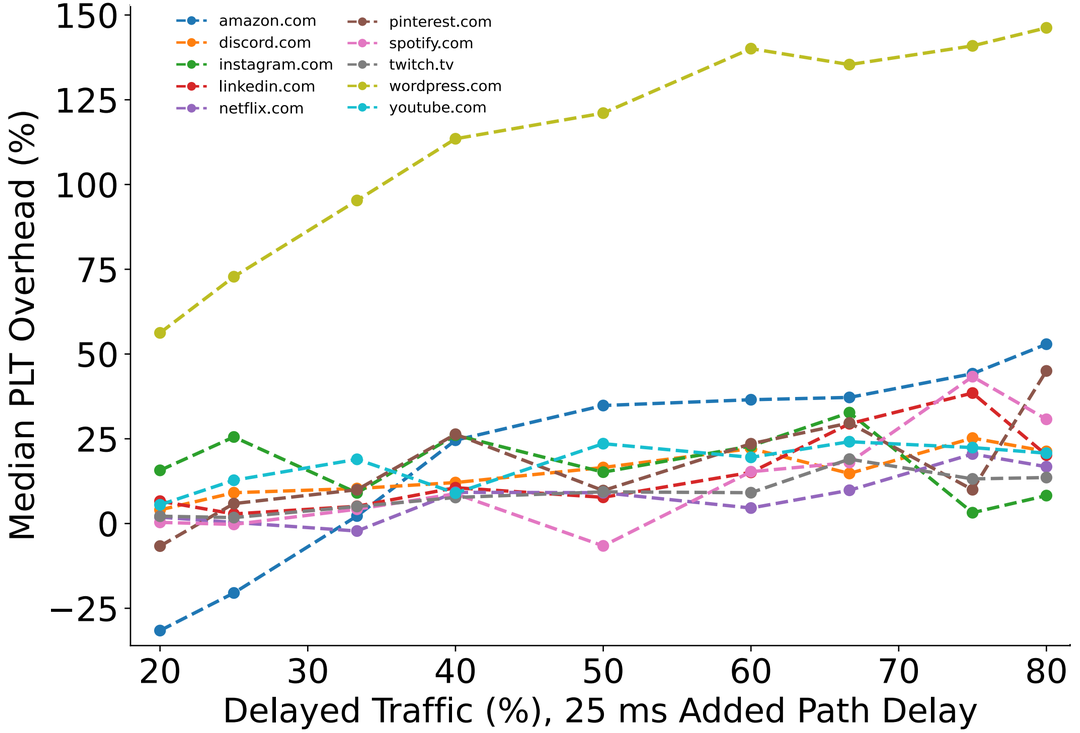

Performance overhead is driven by both delay amount and how much traffic experiences it.

So what privacy can we gain in practice,
while limiting performance overhead?

PraxiGuard

consistently delivers strong privacy gains
within a tight performance budget

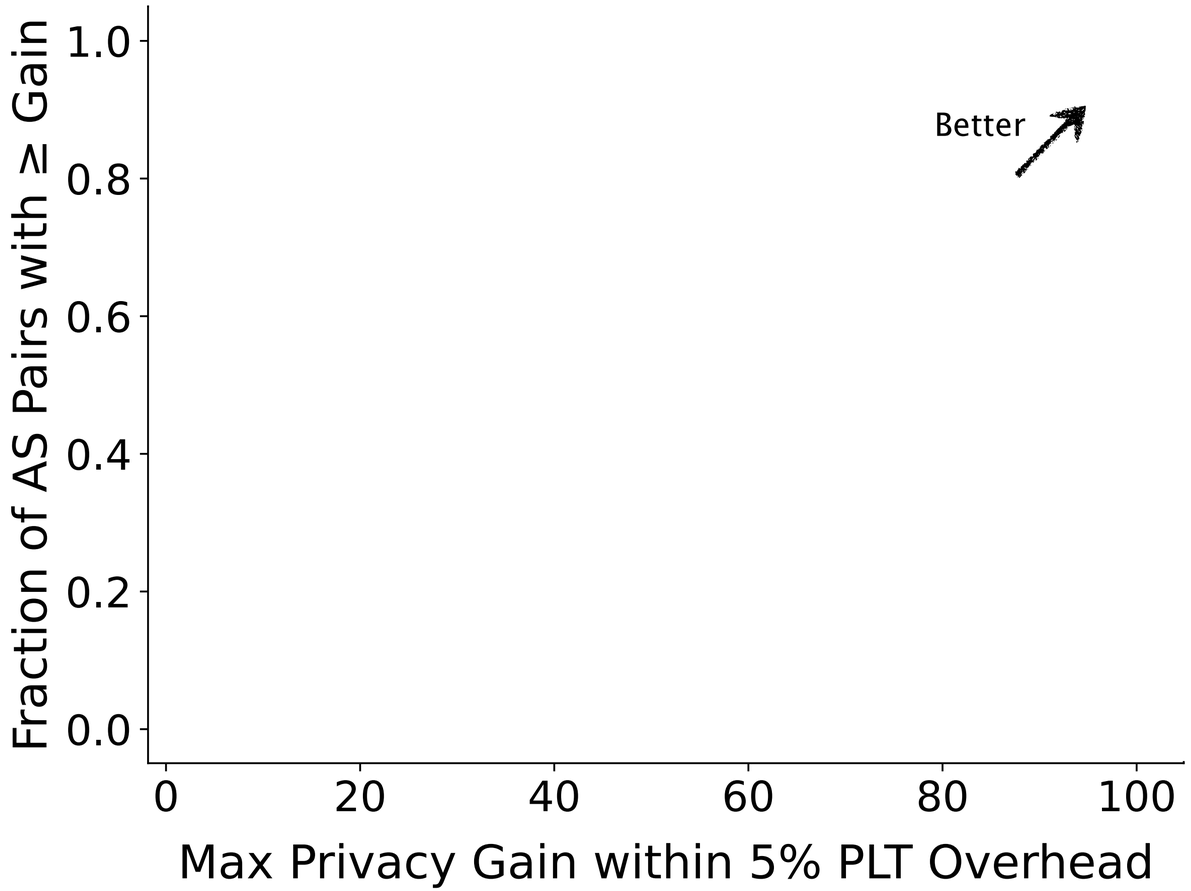

PraxiGuard

consistently delivers strong privacy gains
within a tight performance budget

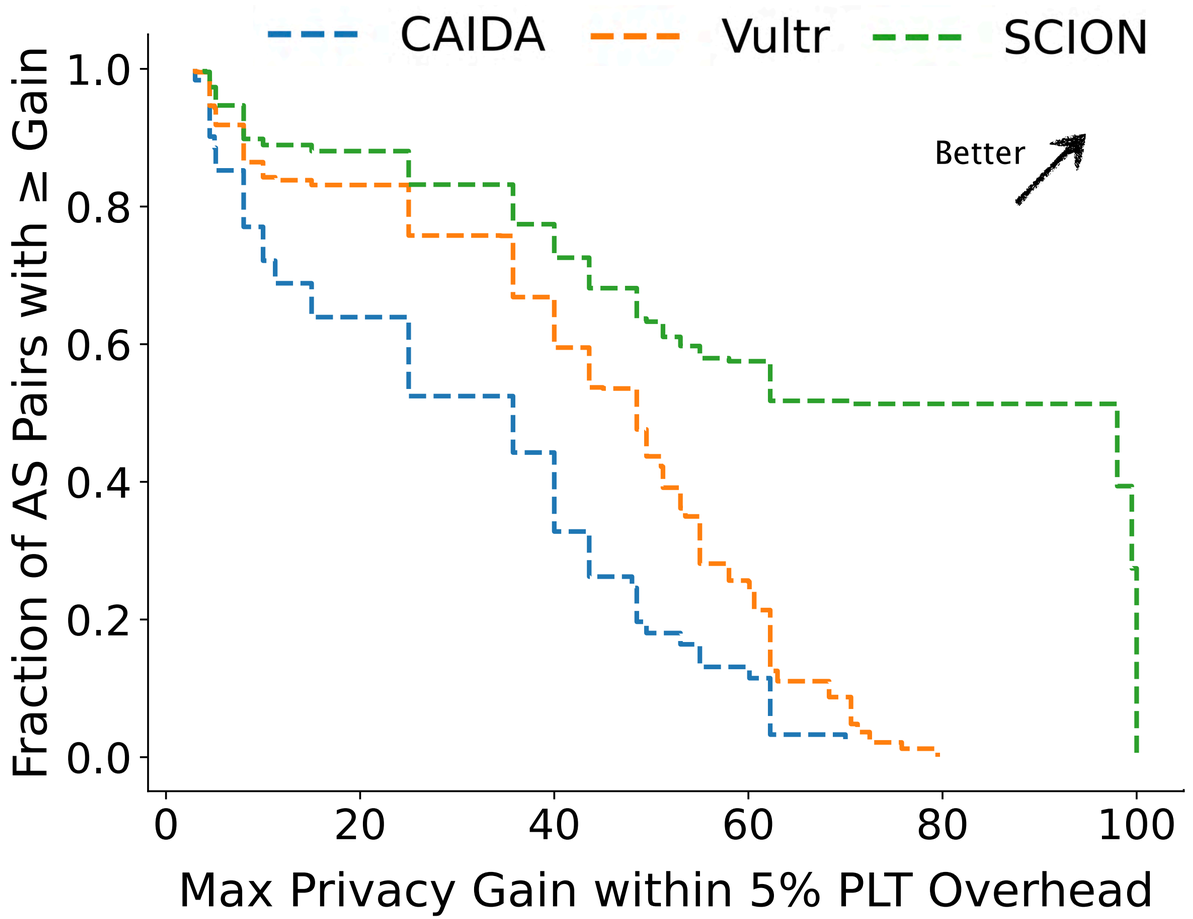

*Includes only site/AS-pair instances with at least one feasible optimized allocation under the 5% predicted PLT-overhead budget.

Part I: Ingress Control

Co-design within a single edge network

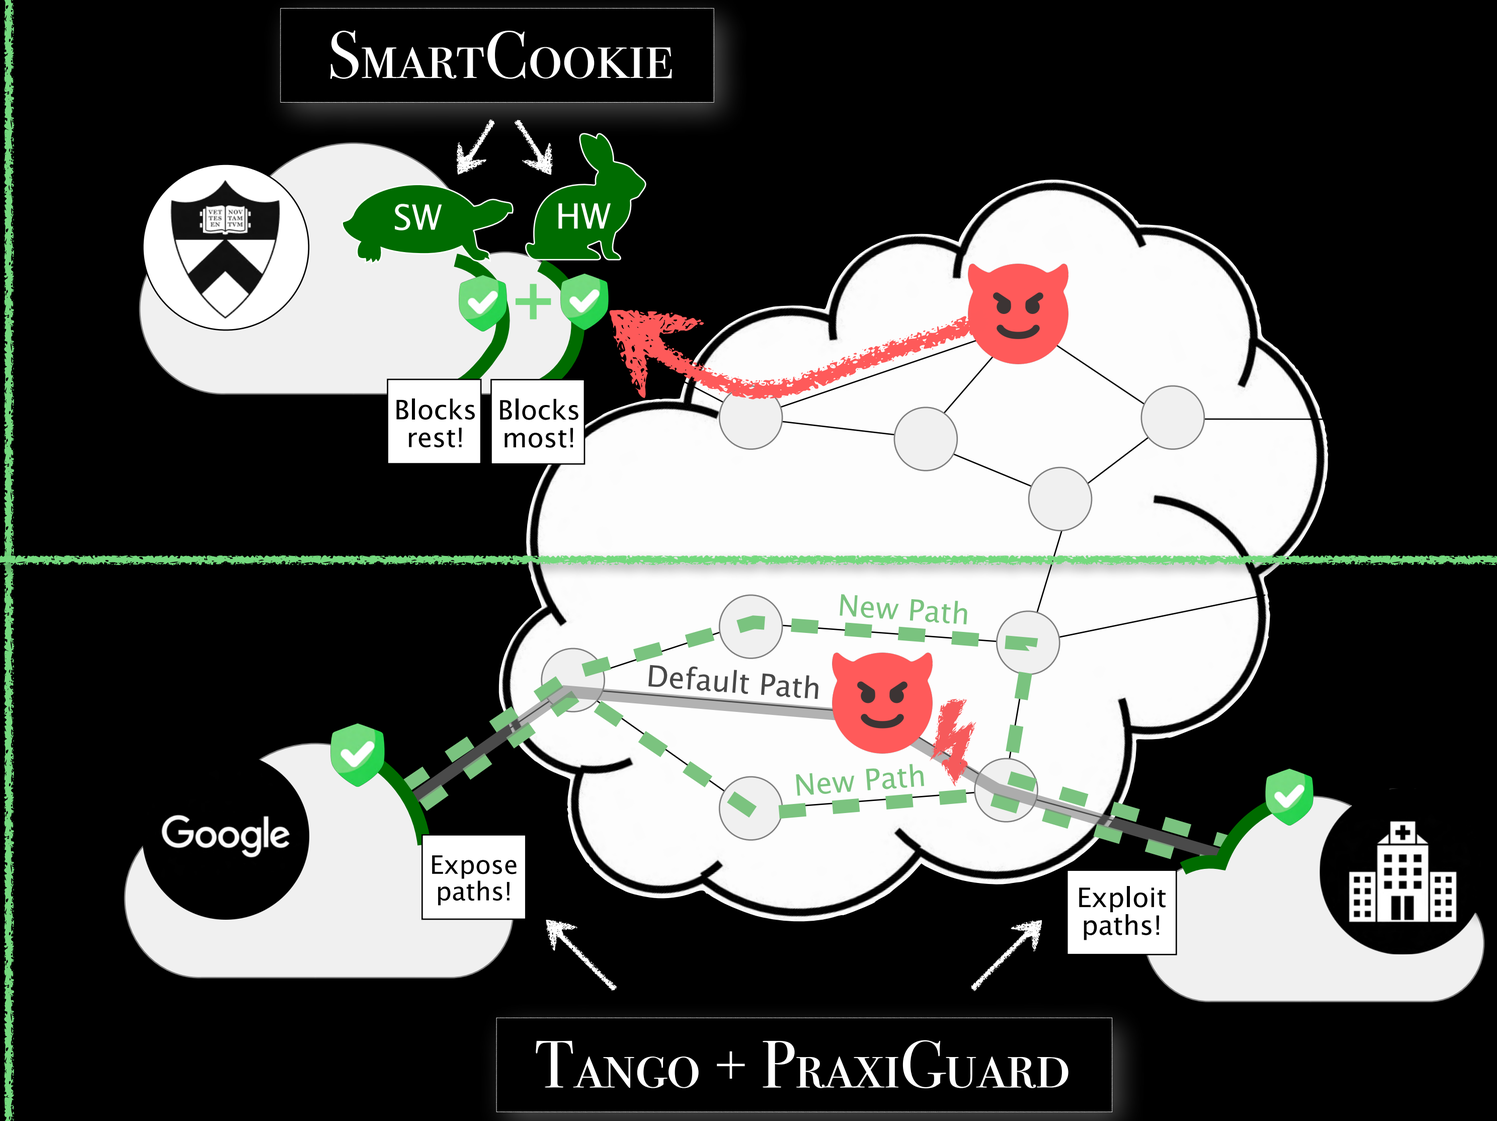

Part II: Route Control

Cooperation between multiple edge networks

Summary

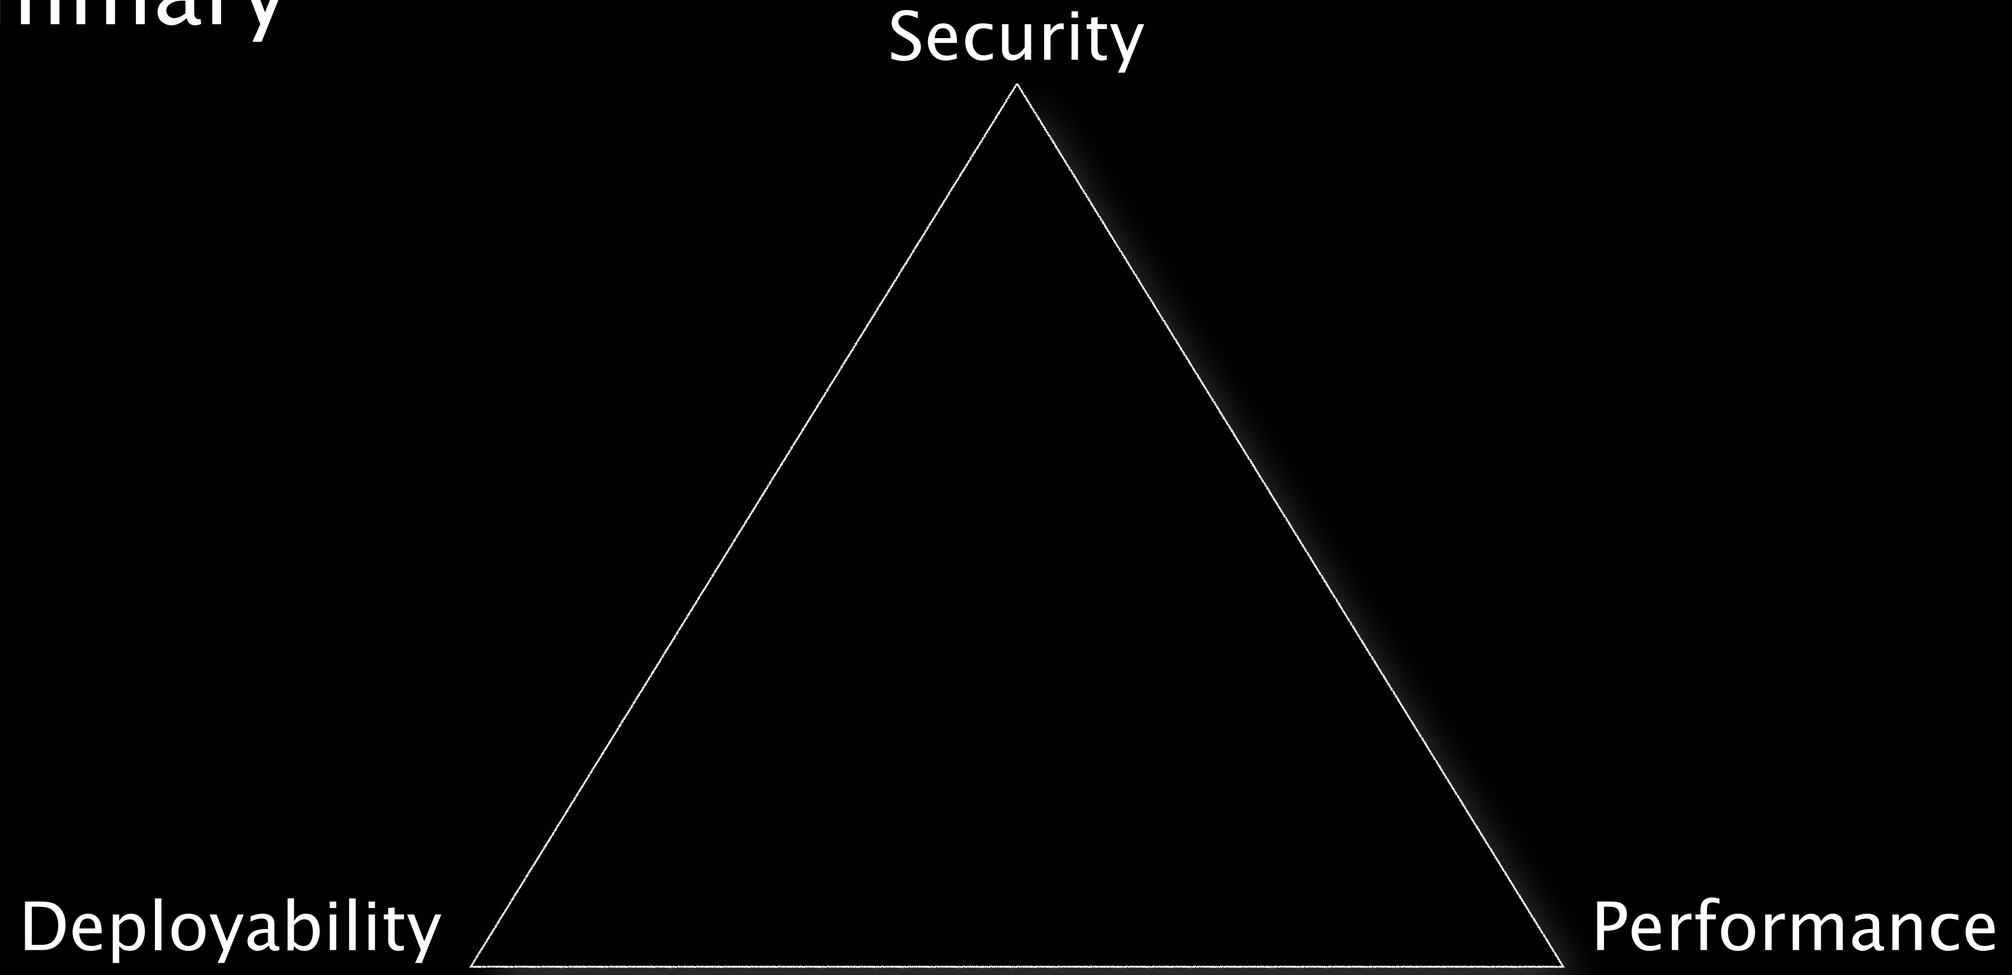

Summary

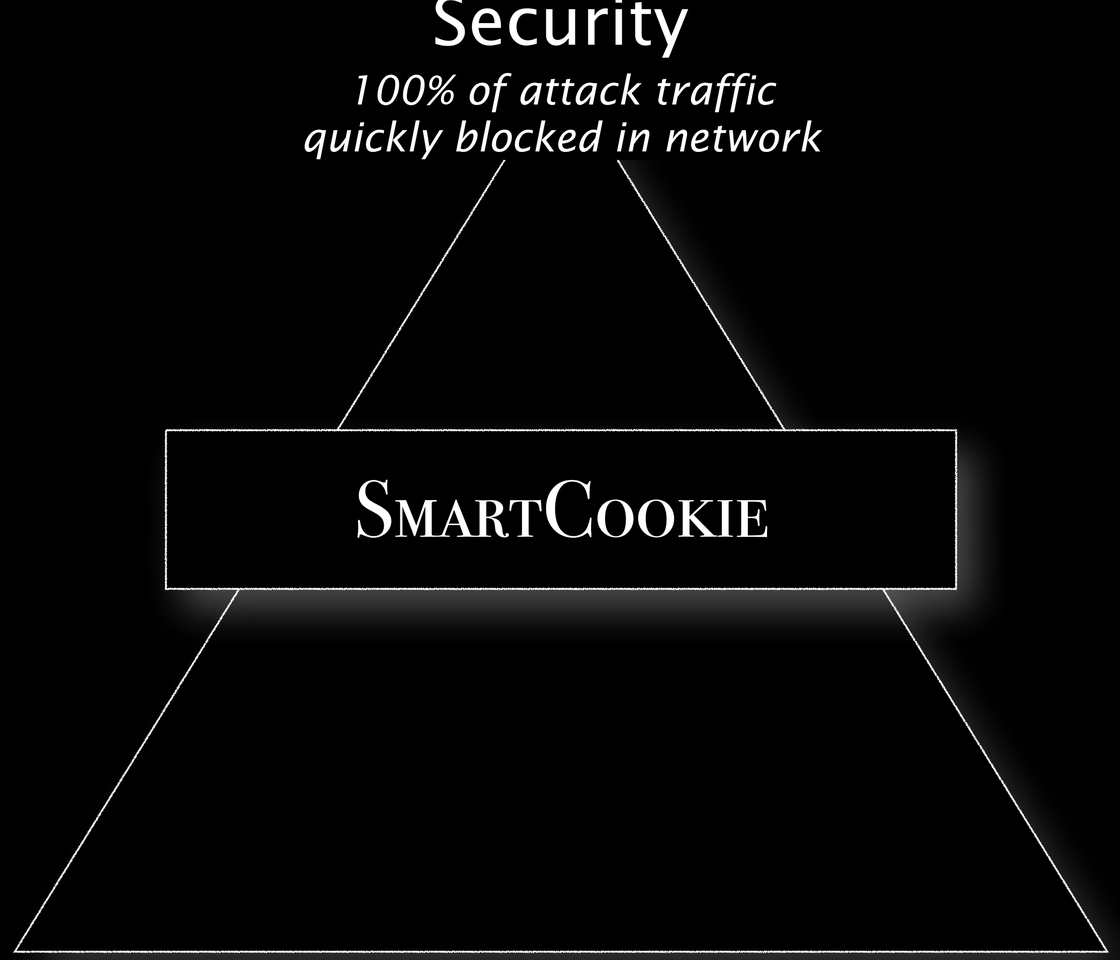
Security
100% of attack traffic
quickly blocked in network
SmartCookie
Deployability
Client behavior and end host
network stacks are unmodified
Performance
Benign client performance is
protected even under attack

Summary

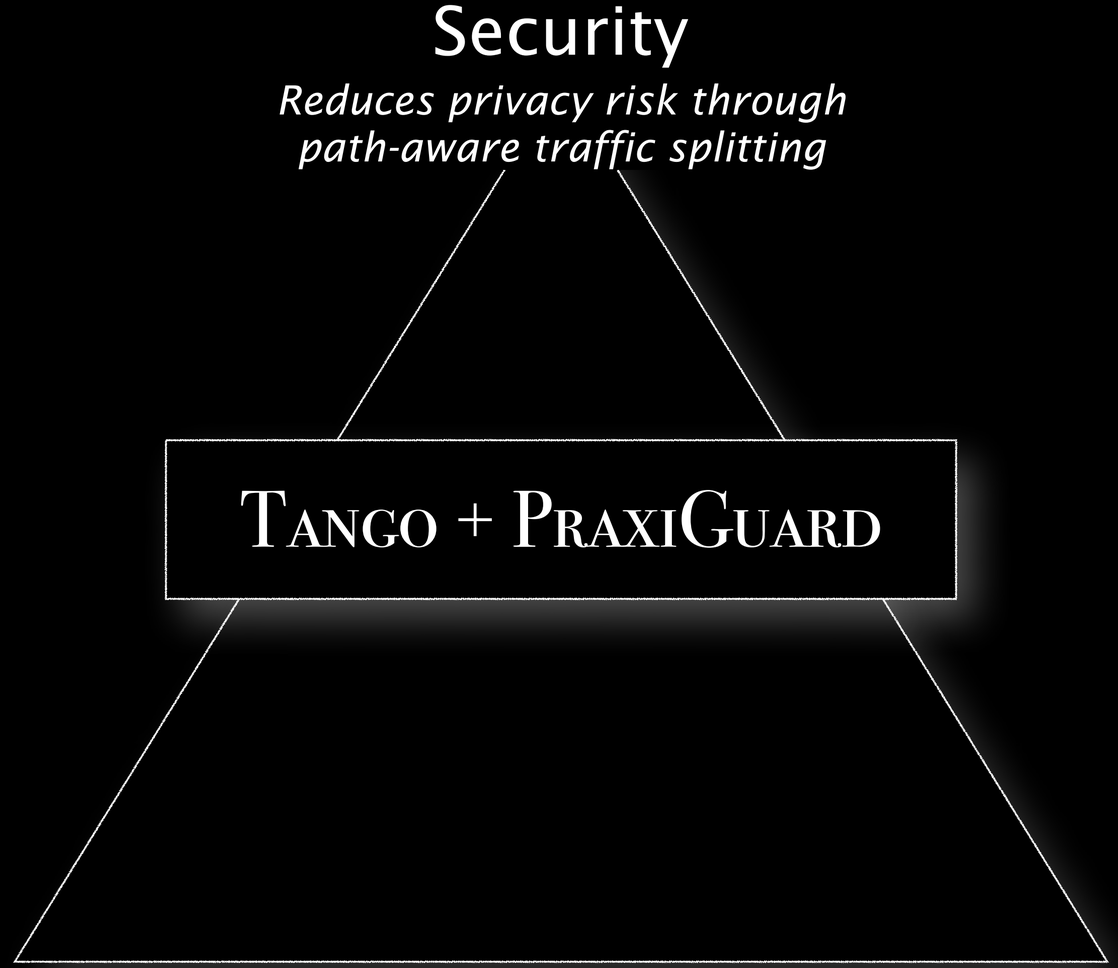
Security
Reduces privacy risk through path-aware traffic splitting
Tango + PraxiGuard
Deployability
Uses paths that are available in practice and only benefiting parties deploy solution
Performance
Limits performance overhead at the application level

Security

Identify threats and opportunities in emerging applications with Internet measurements+ML

Coming soon: joining as an Assistant Professor of Computer Science

Deployability

Build deployable systems that are compatible with today's network

Performance

Design hardware-software defenses that preserve user performance

...across a range of programmable targets

Speed *Flexibility*

Programmable switches

SmartNICs

Linux kernel

CPUs

...across a range of end point devices

Resource-Constrained *Resource-Rich*

Drones Mobile Devices Laptops Servers

...across a range of applications

Ultra Latency-Sensitive — *Latency-Sensitive*

AR/VR

Real-Time Media
(Gaming, Telemedicine)

AI/ML Inference

Streaming Media

Part I: Ingress Control

Co-design within a single edge network

Part II: Route Control

Cooperation between multiple edge networks

Thank you!
Questions?

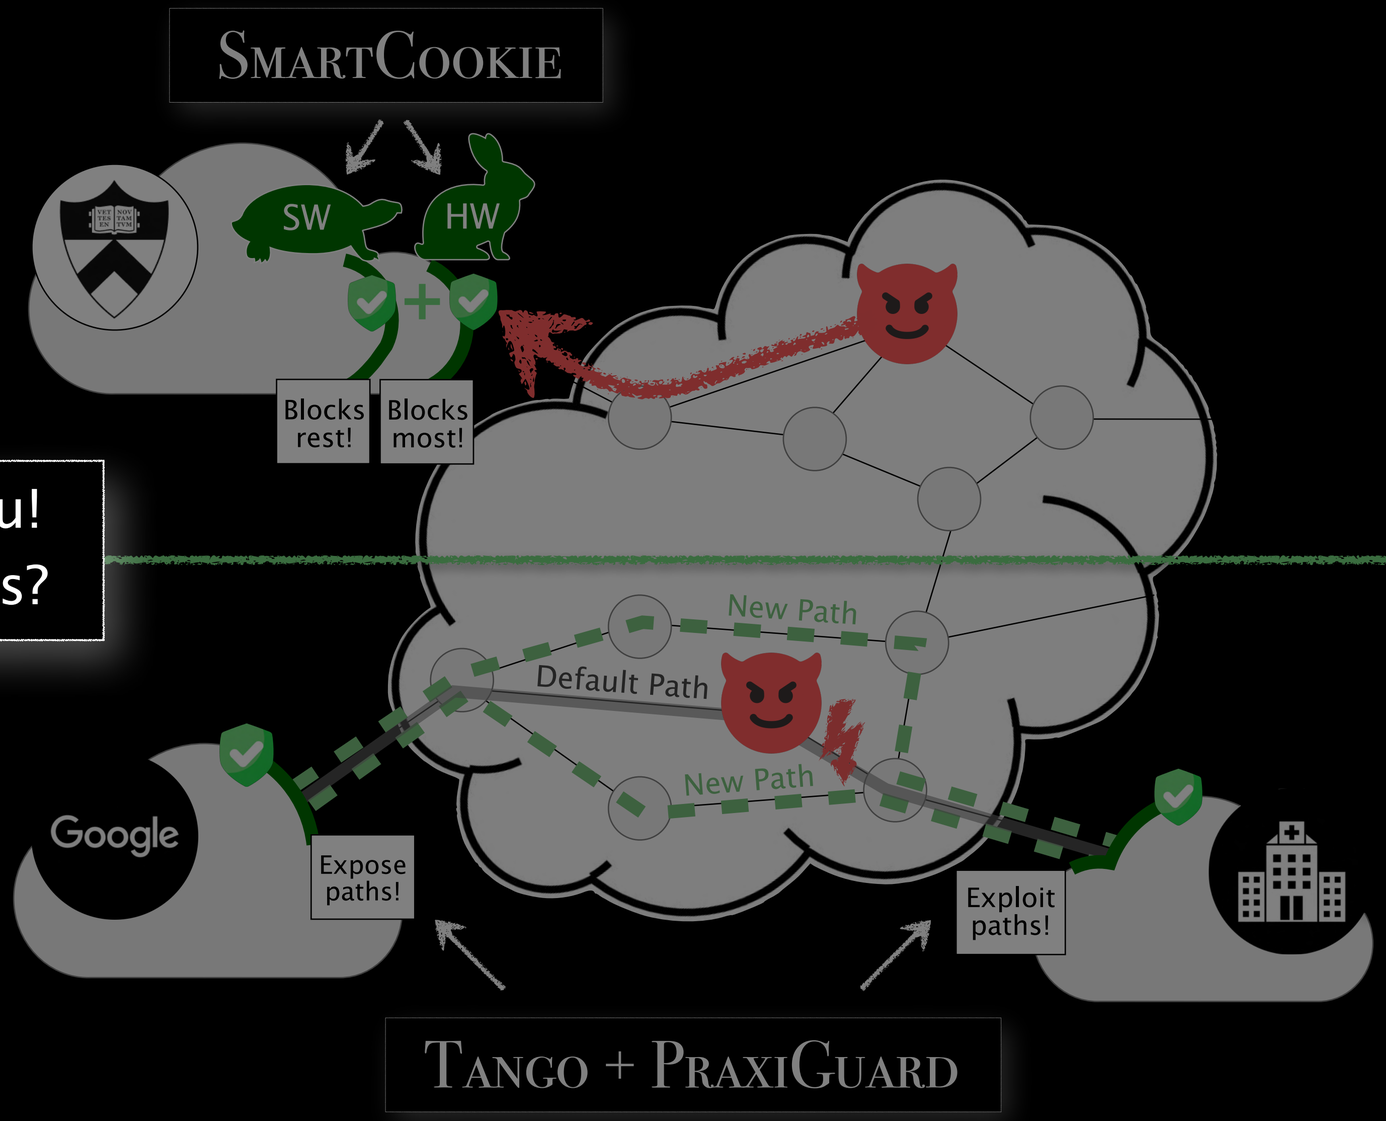

Acknowledgements

Advisors


Jennifer Rexford


Maria Apostolaki

Committee


David Hay

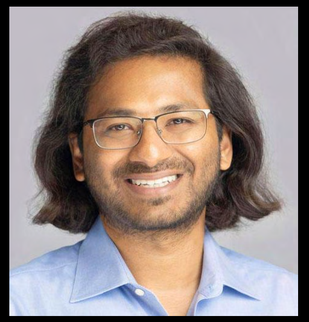
Prateek Mittal


Adrian Perrig

Collaborators

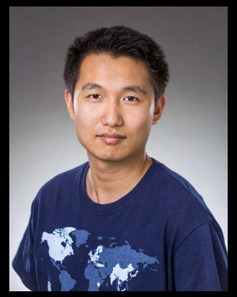
Danny Chen


Henry Birge-Lee


Sata Sengupta


Benji Herber

Sofia Marina

Academic Family (Past and Present)


Mary Hogan

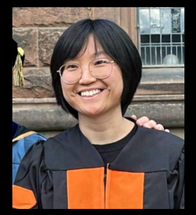
Yufei Zheng

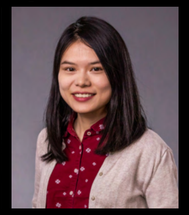
Molly Pan


Robert MacDavid


John Sonchack


Oliver Michel


Joon Kim


Fengchen Gong


Ann Zhou


Kostas Doumanidis


Minhao Jin


Hongyu Hè


Cleef Apostolaki

Cooperative Network Systems for Protecting Internet Users

Sophia Yoo

Dissertation Defense

April 24, 2026

PRINCETON UNIVERSITY