On the ineffectiveness of QUIC PADDING against website fingerprinting

Ludovic Barman, Sandra Siby, Christopher Wood, Marwan Fayed, Nick Sullivan, Carmela Troncoso

IETF PEARG, 21 March 2022

EPFL

CLOUDFLARE®

What is website fingerprinting?

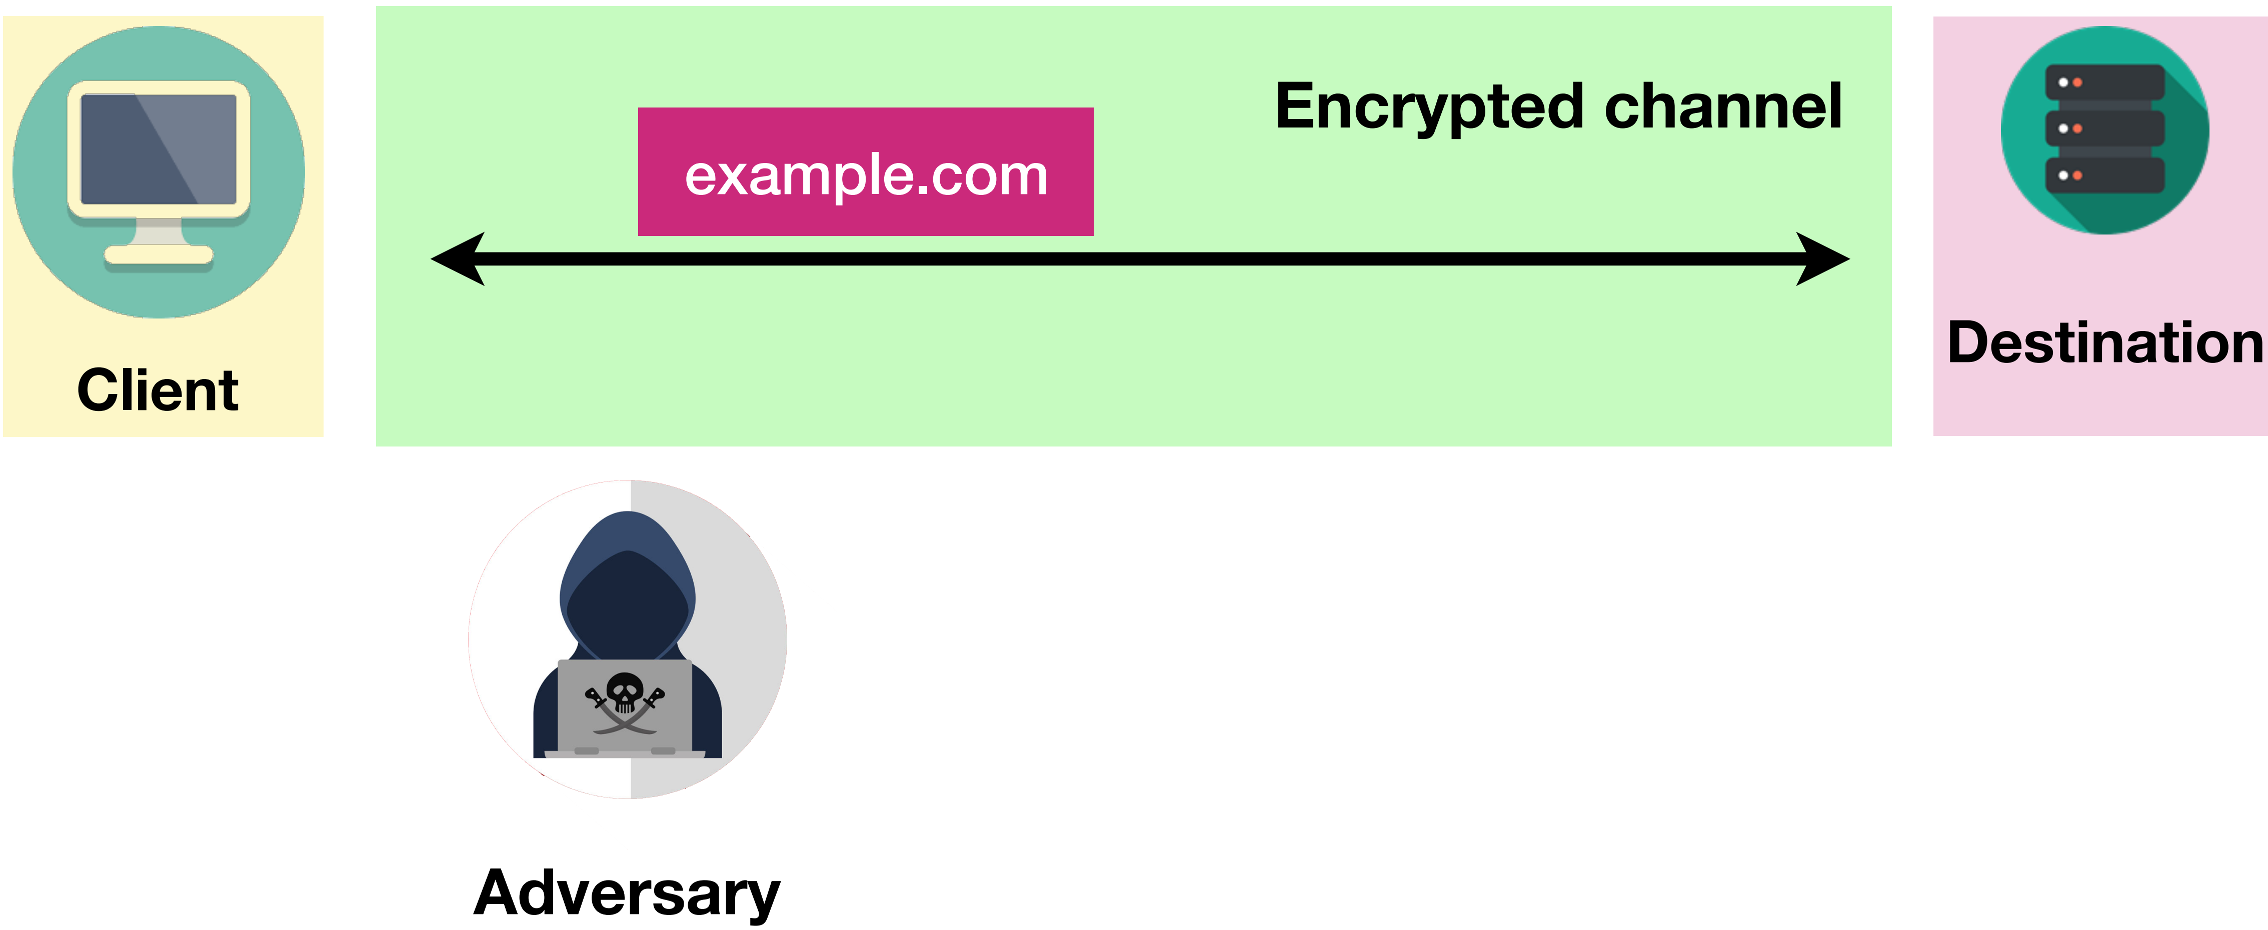

What is website fingerprinting?

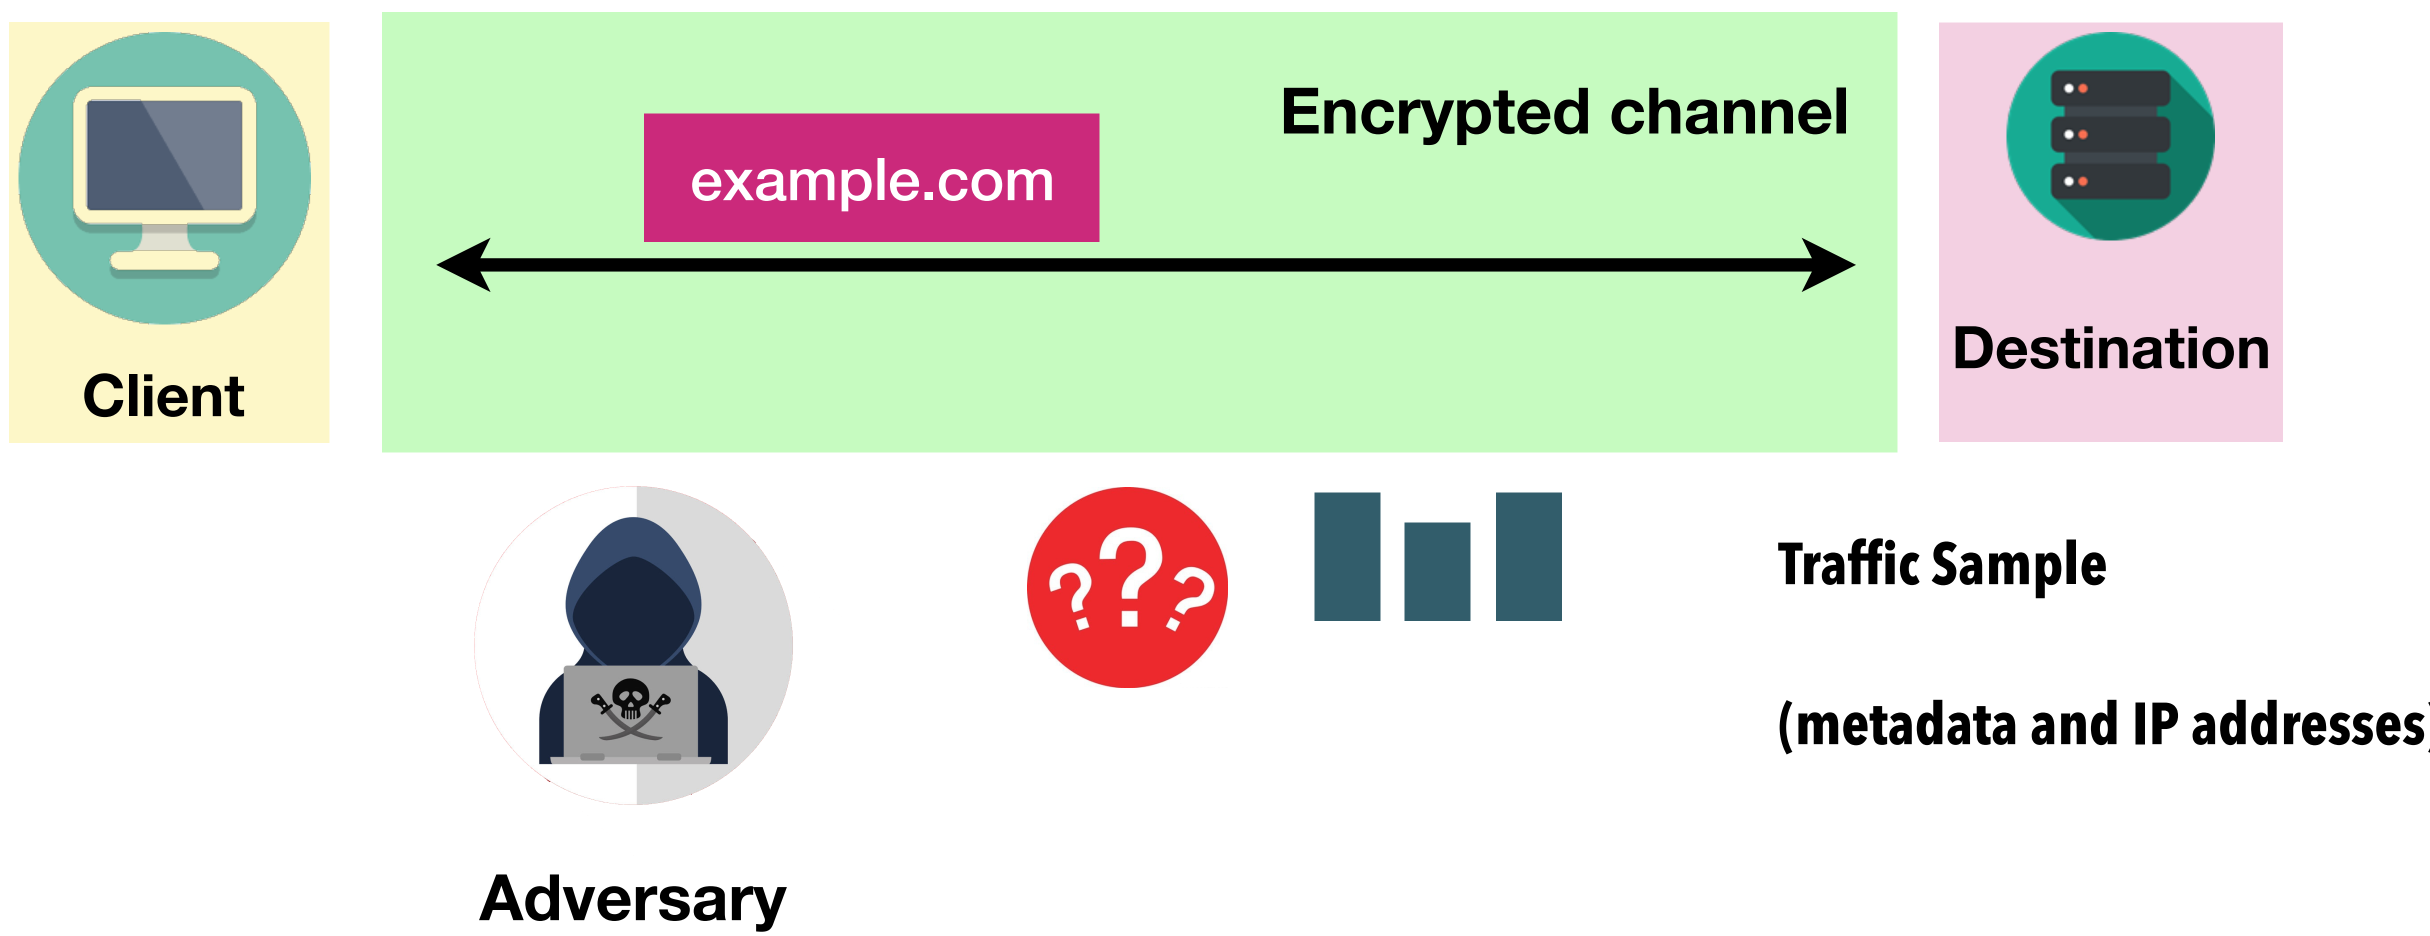

What is website fingerprinting?

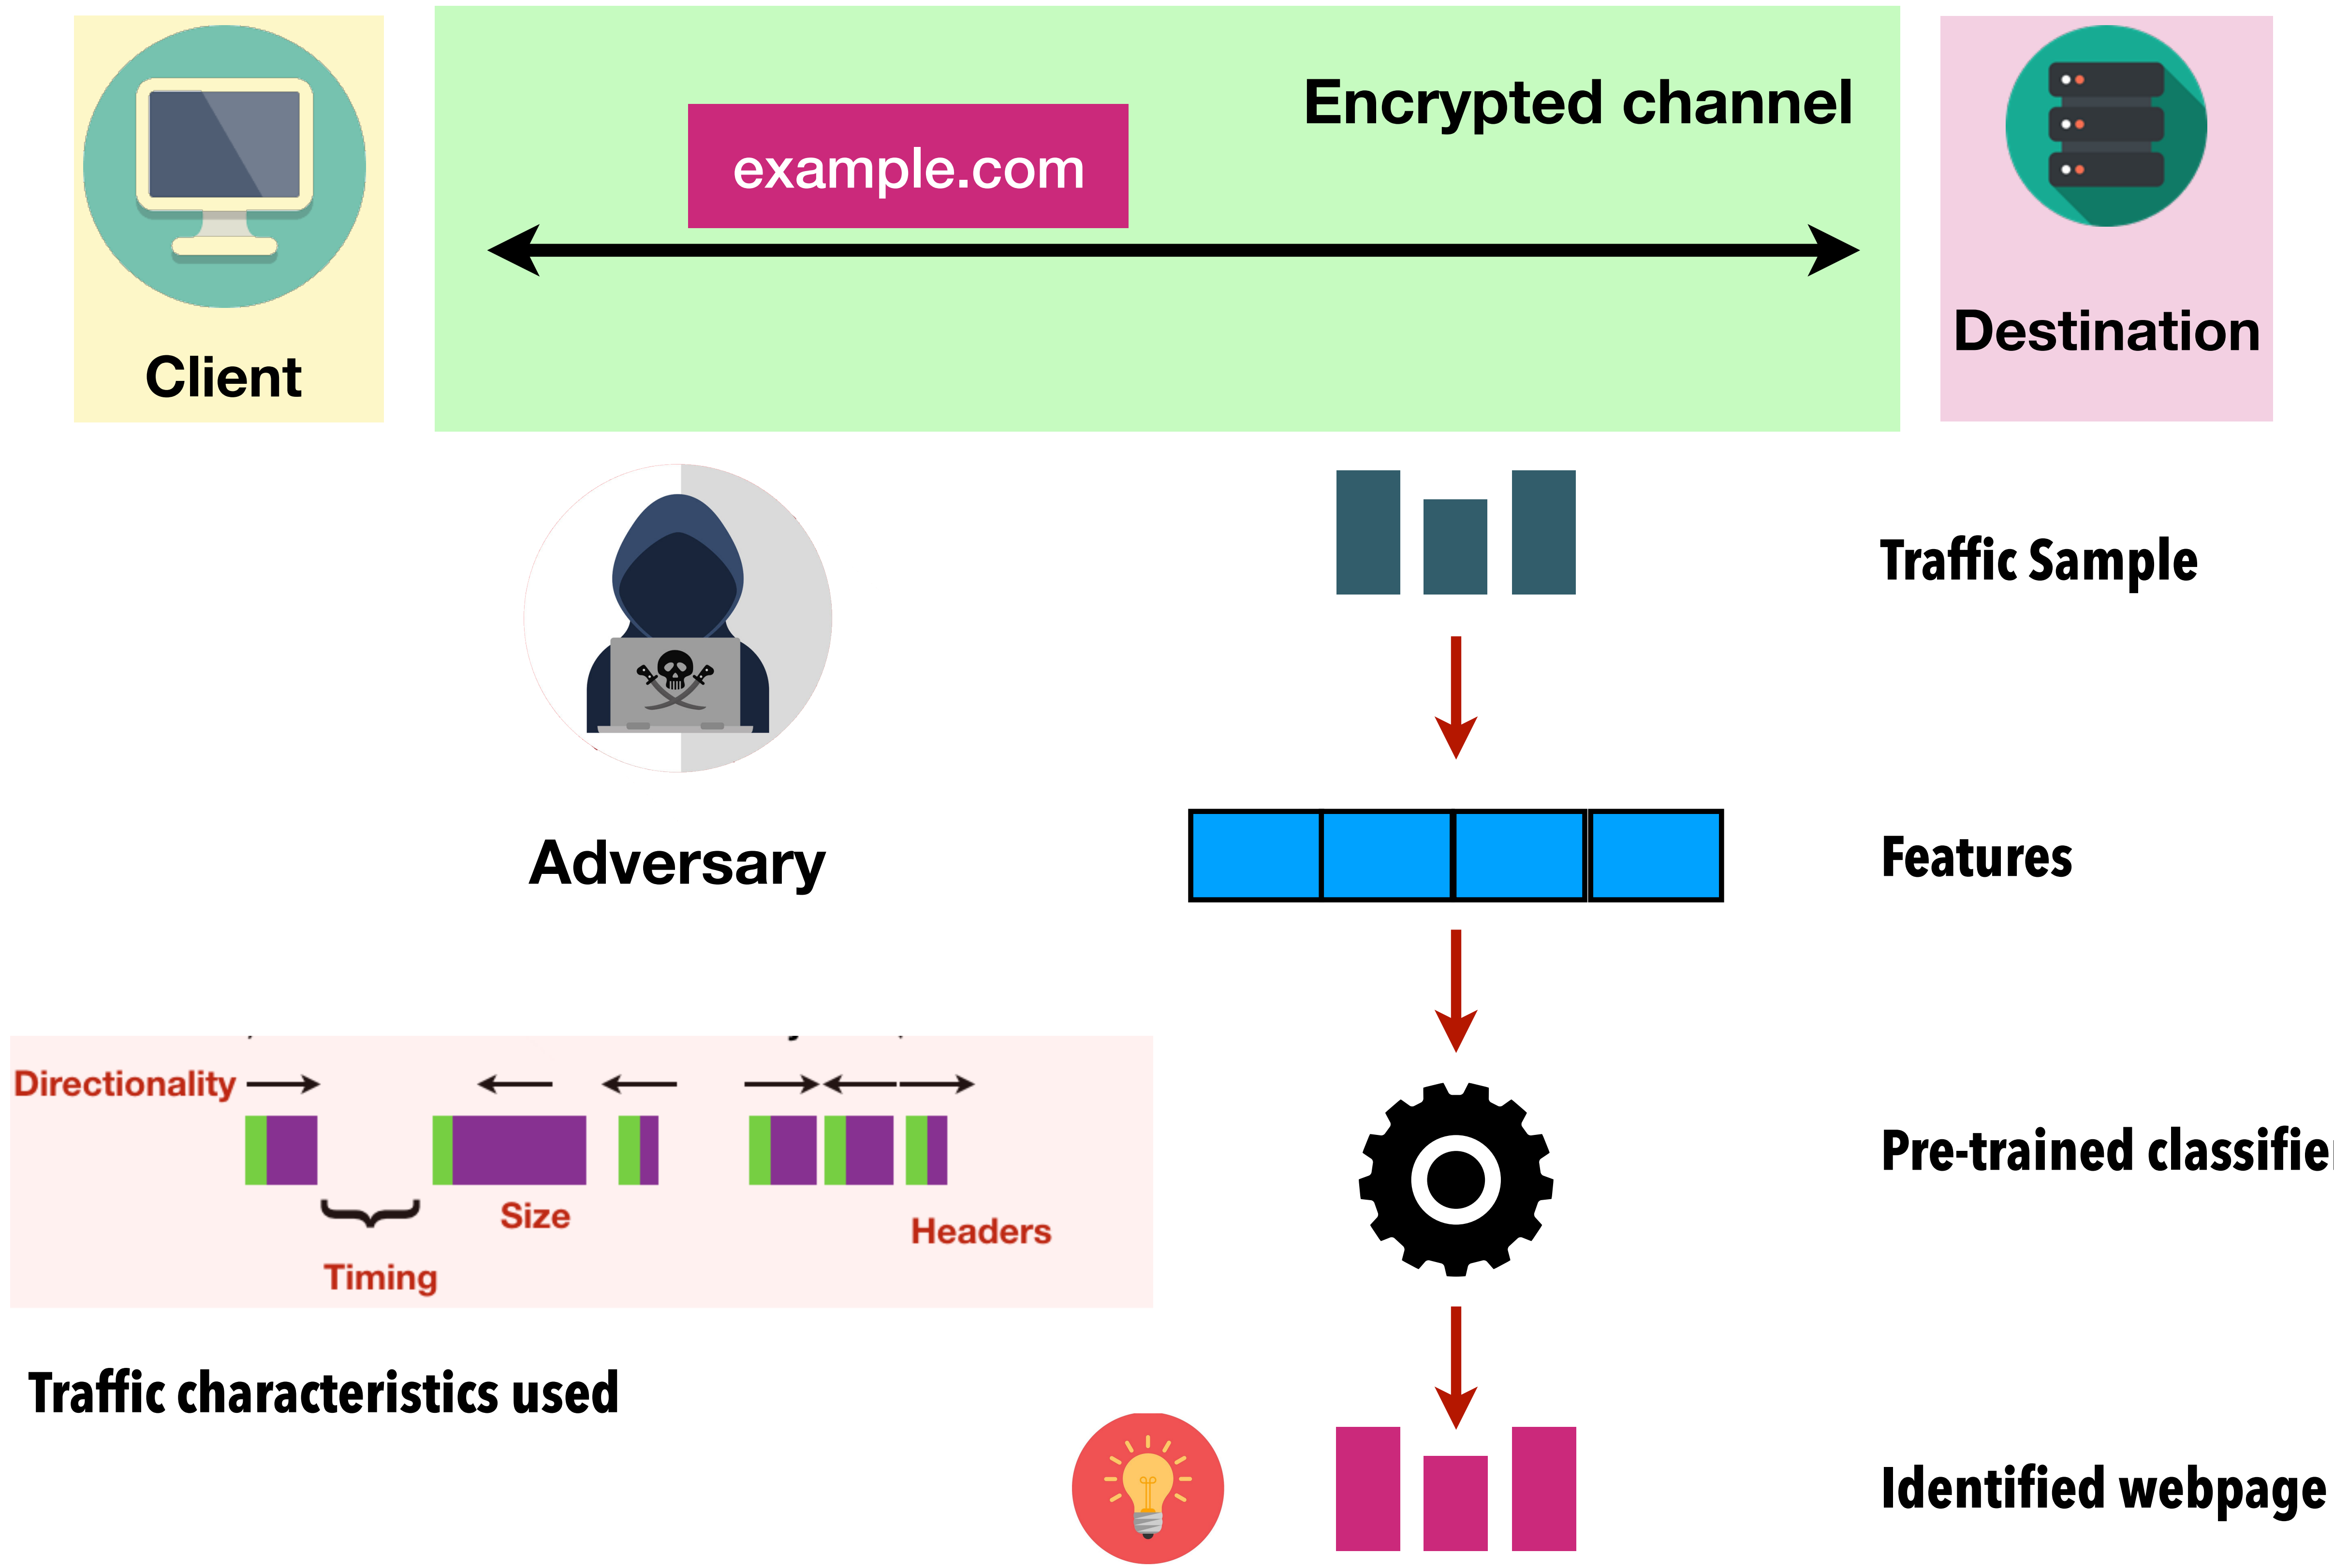

Website Fingerprinting on QUIC

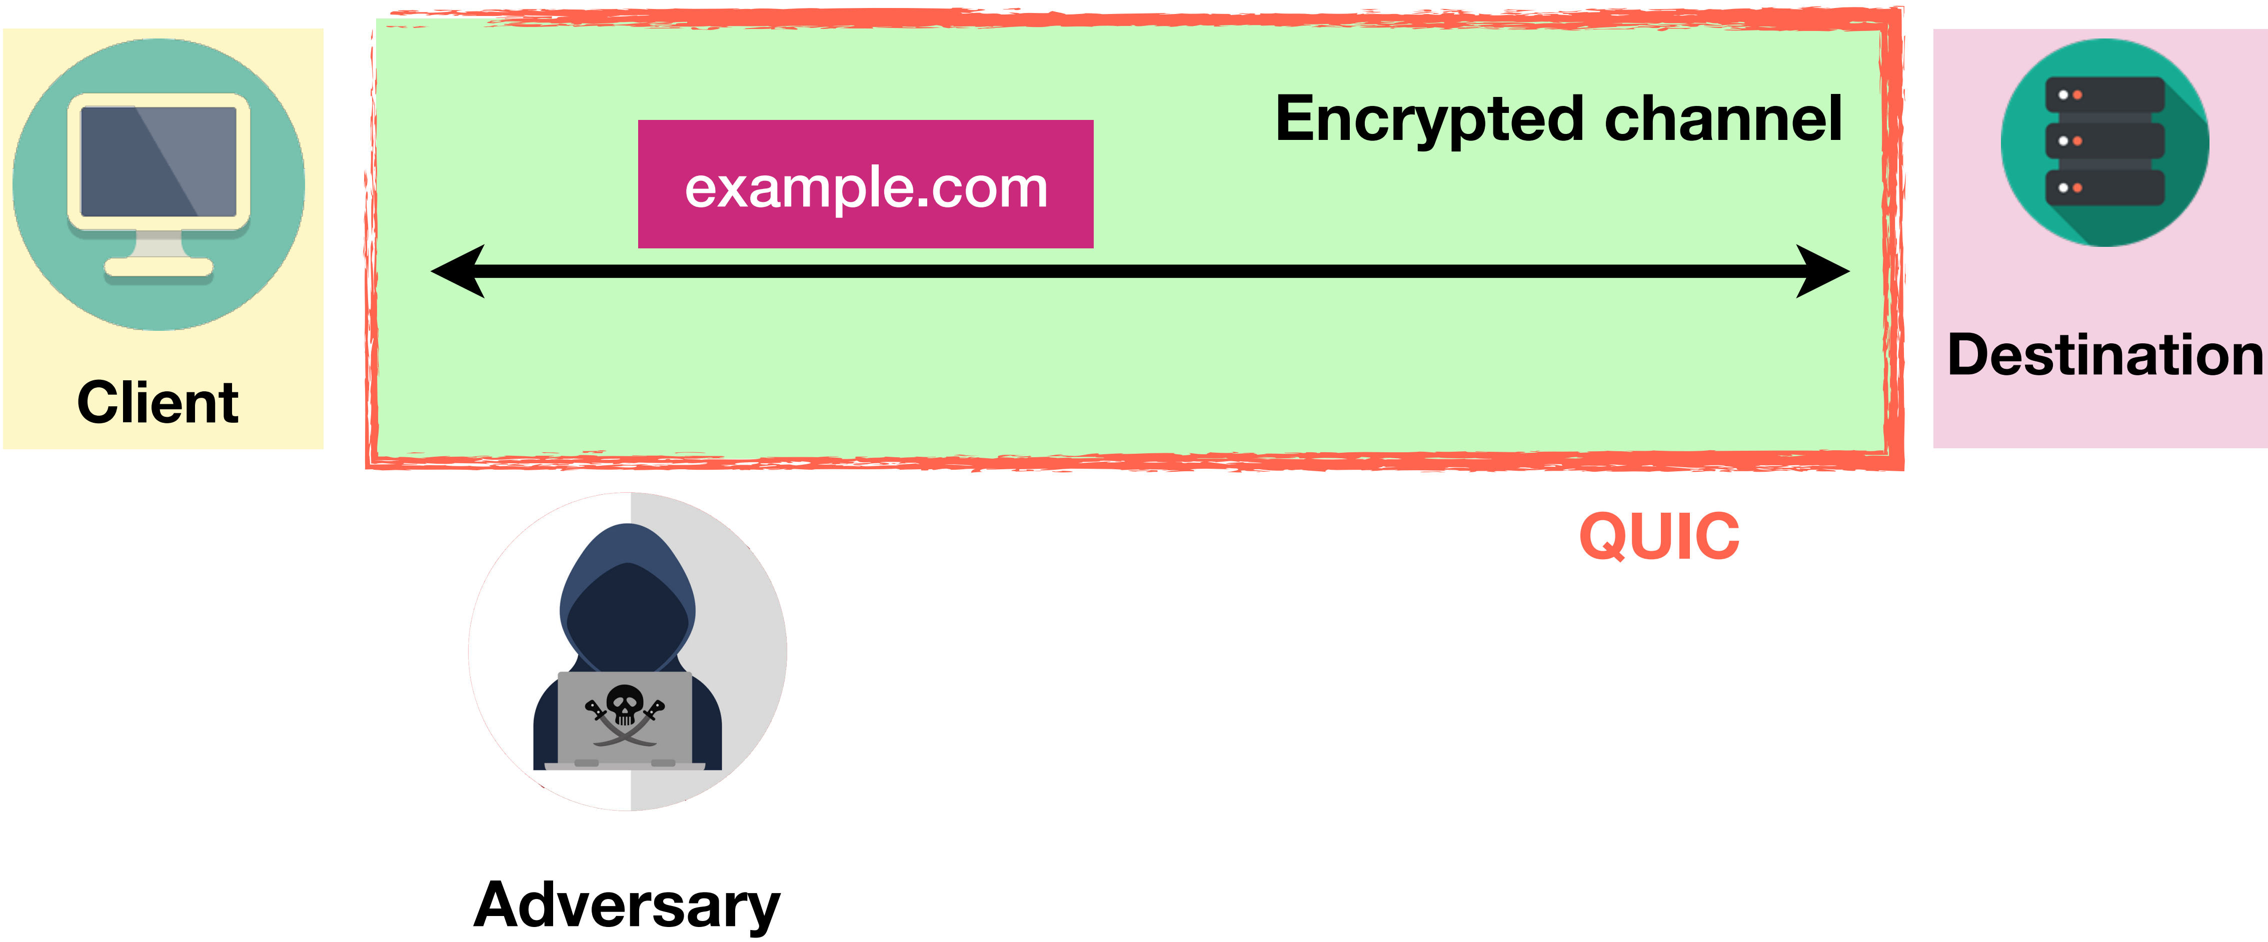

Website Fingerprinting on QUIC

Website fingerprinting on QUIC has already been studied before [1]

Conclusion: It is not harder to fingerprint QUIC as compared to TCP

[1] https://datatracker.ietf.org/meeting/111/materials/slides-111-pearg-website-fingerprinting-in-the-age-of-quic-00

Website Fingerprinting on QUIC

Website fingerprinting on QUIC has already been studied before [1]

Conclusion: It is not harder to fingerprint QUIC as compared to TCP

QUIC RFC specifies PADDING frame [2]:

"Padding can be used … to provide protection against traffic analysis …"

[1] https://datatracker.ietf.org/meeting/111/materials/slides-111-pearg-website-fingerprinting-in-the-age-of-quic-00

[2] https://www.rfc-editor.org/rfc/rfc9000.html

Adversarial Model

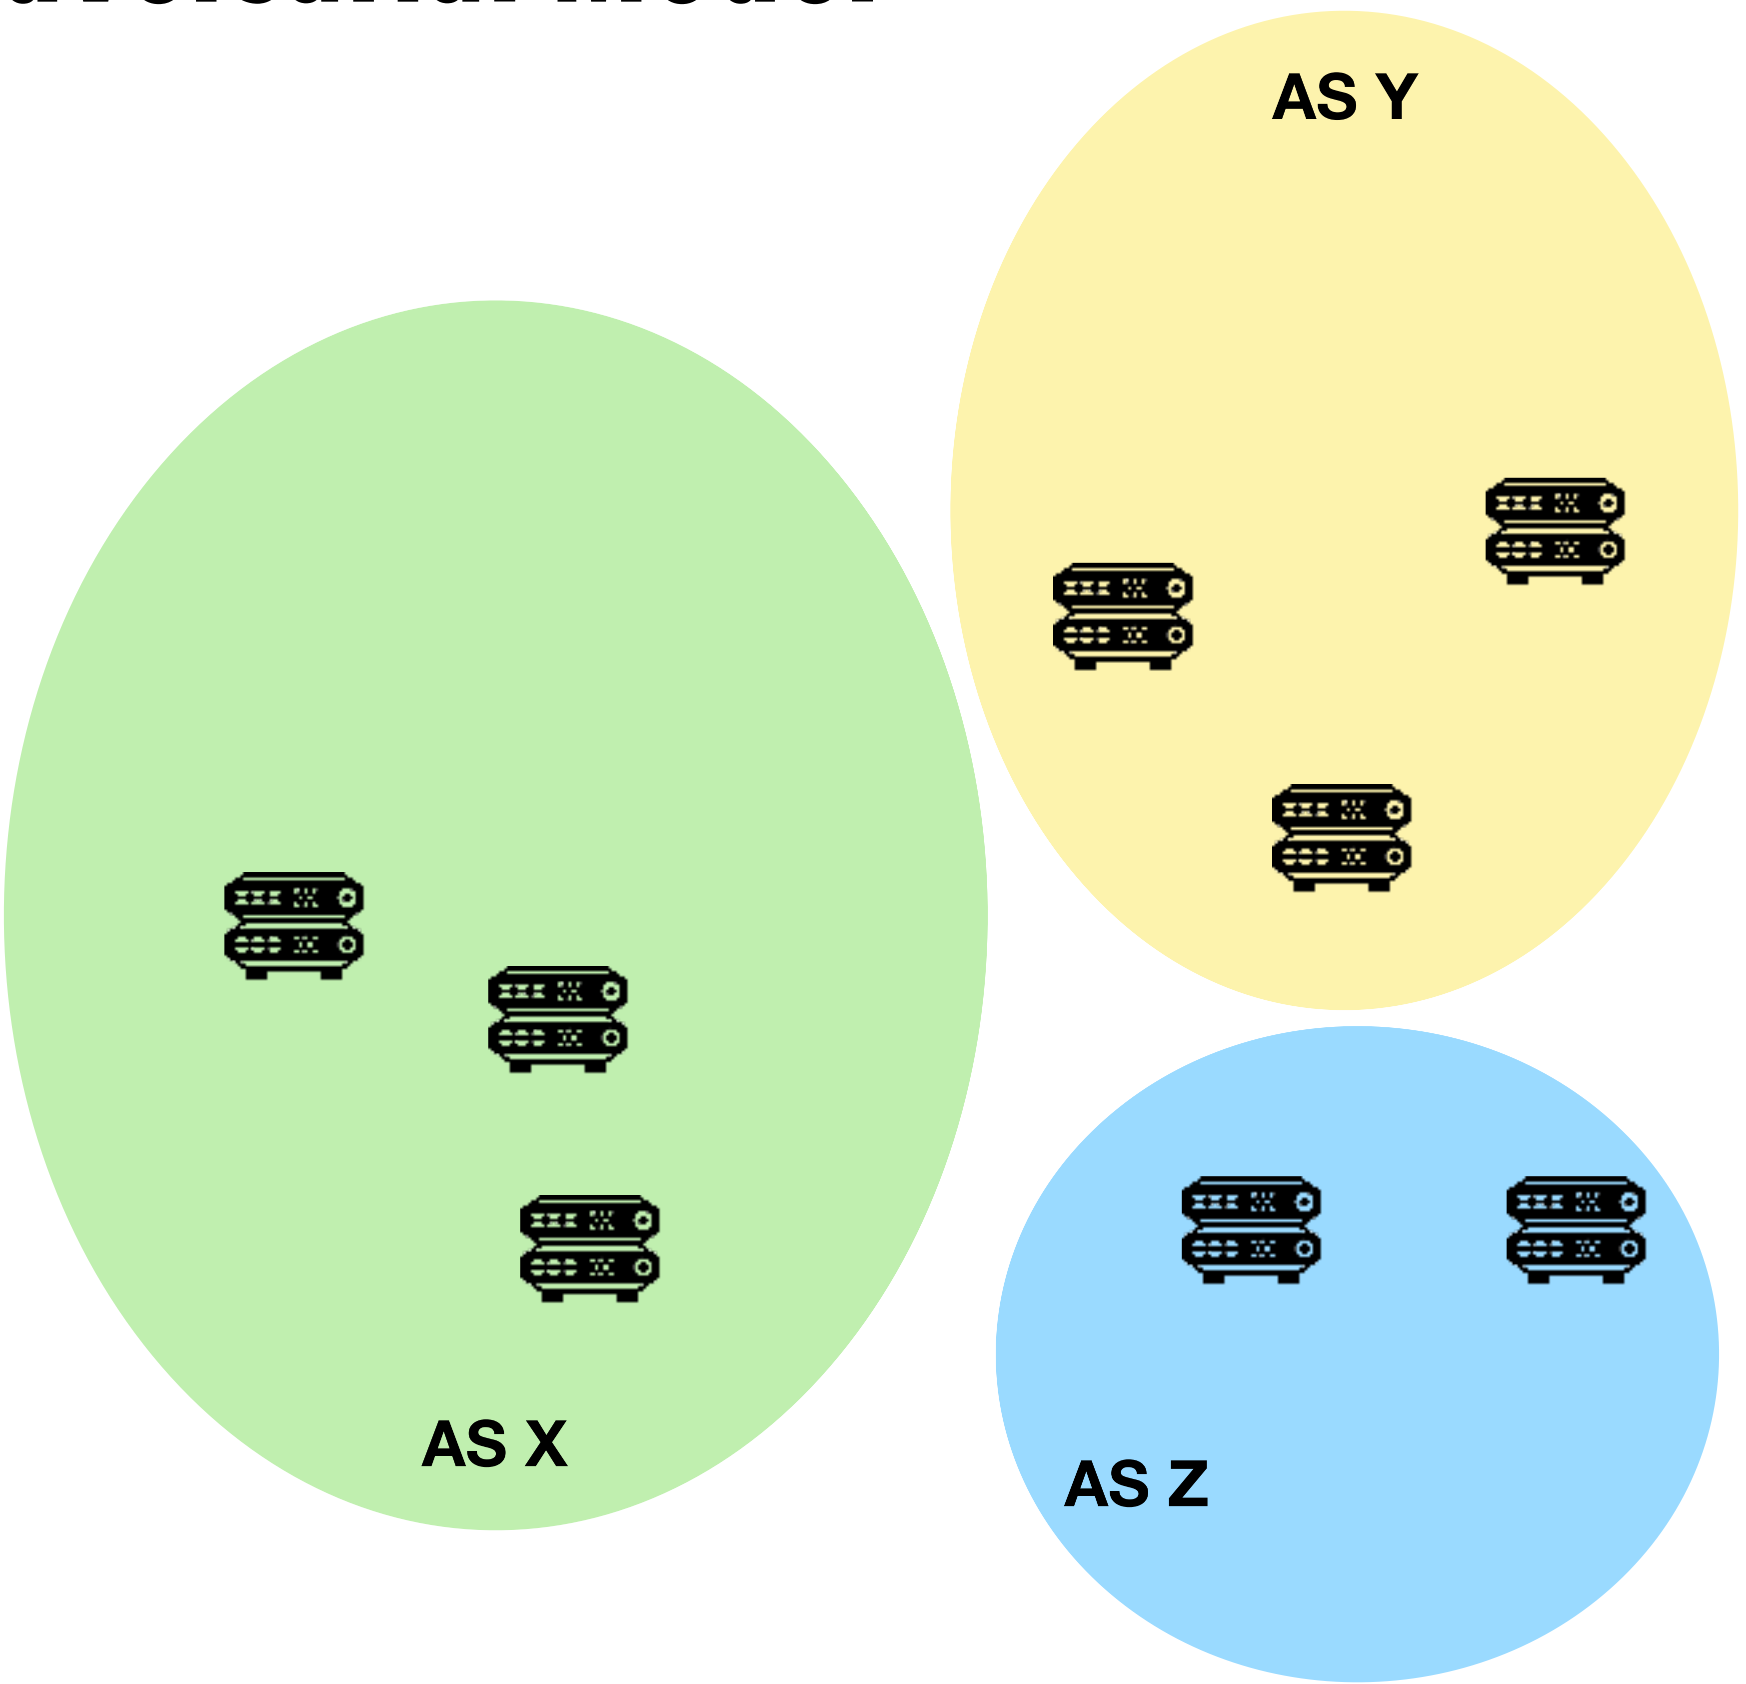

Adversarial Model

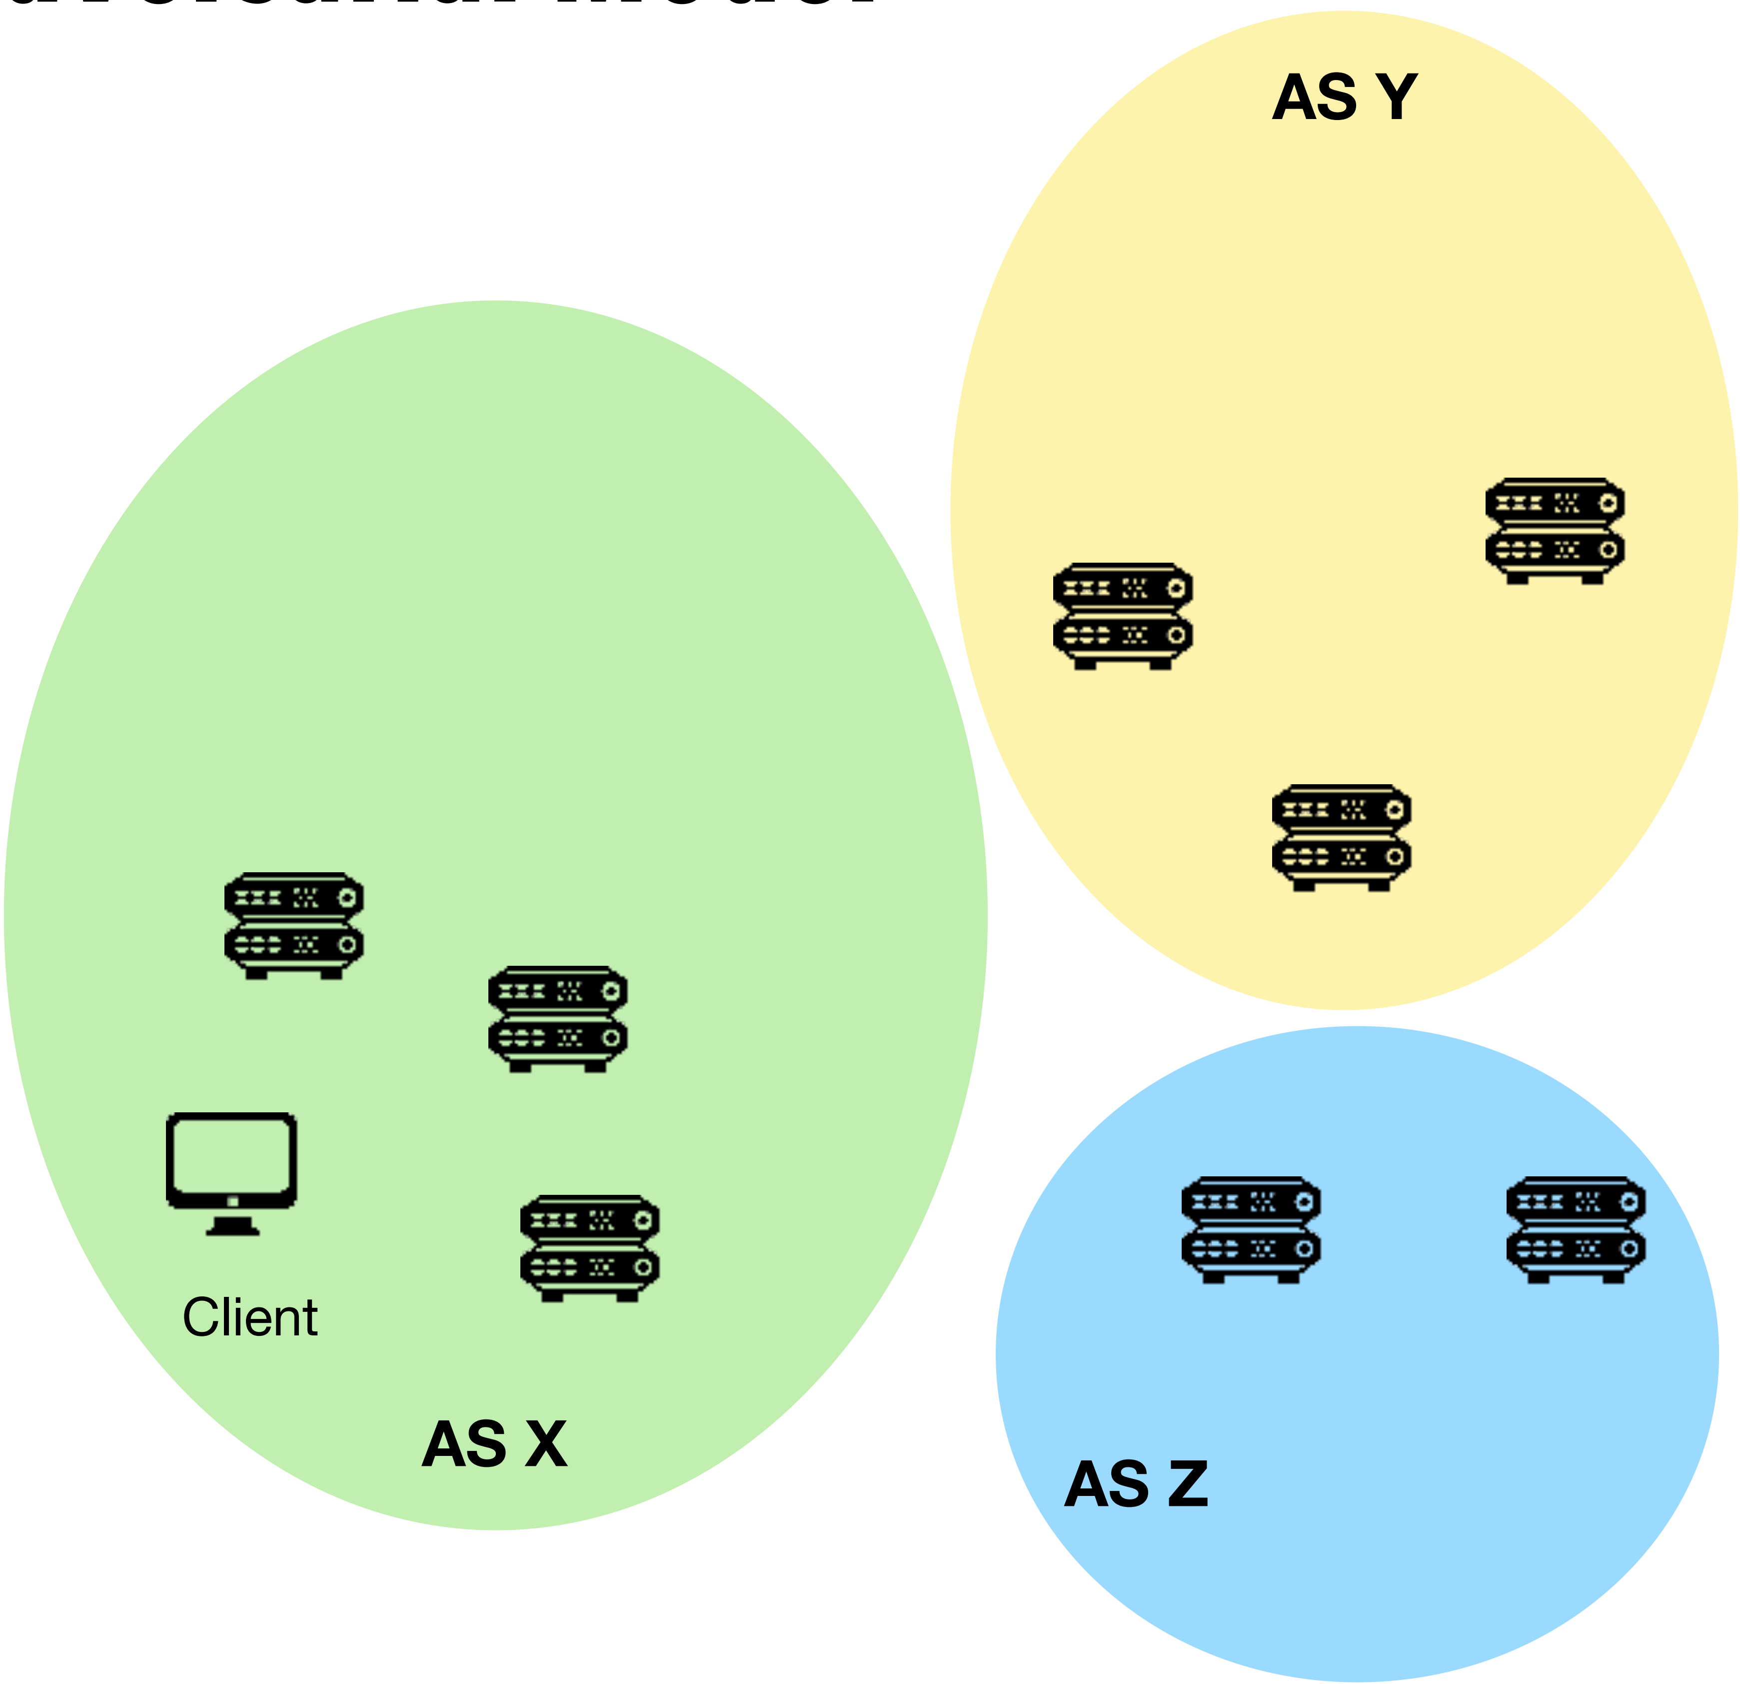

Adversarial Model

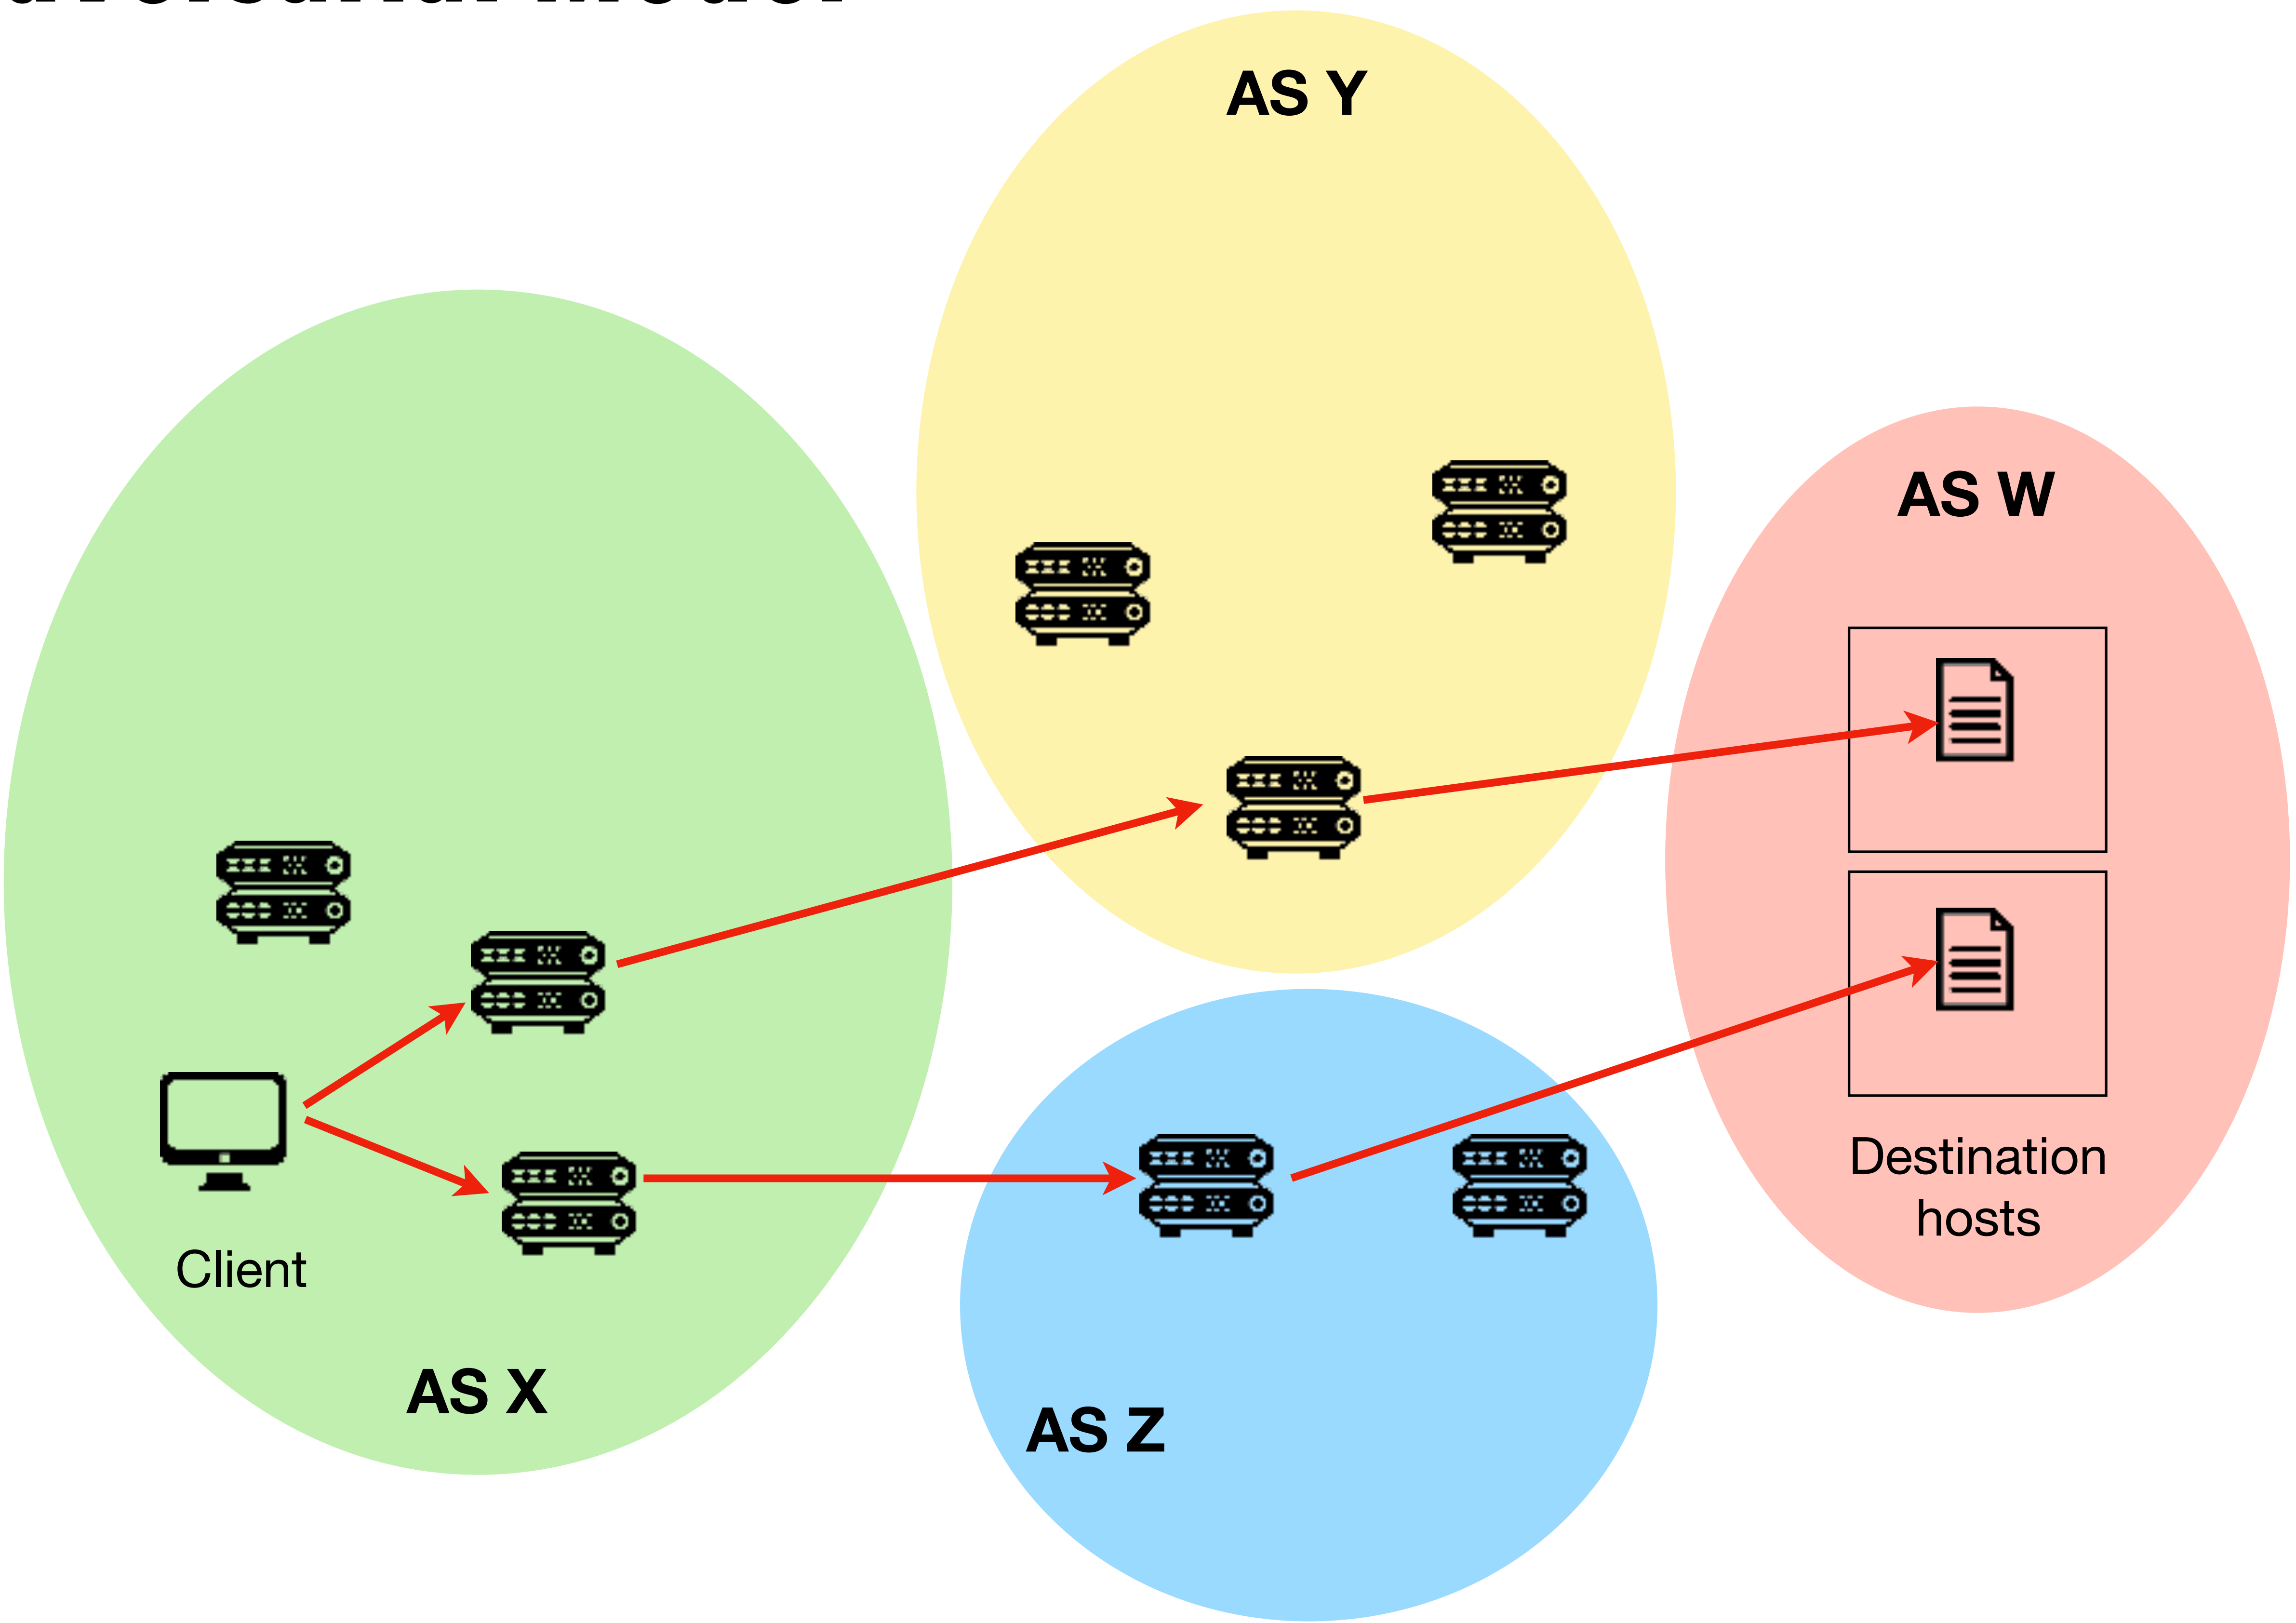

Adversarial Model

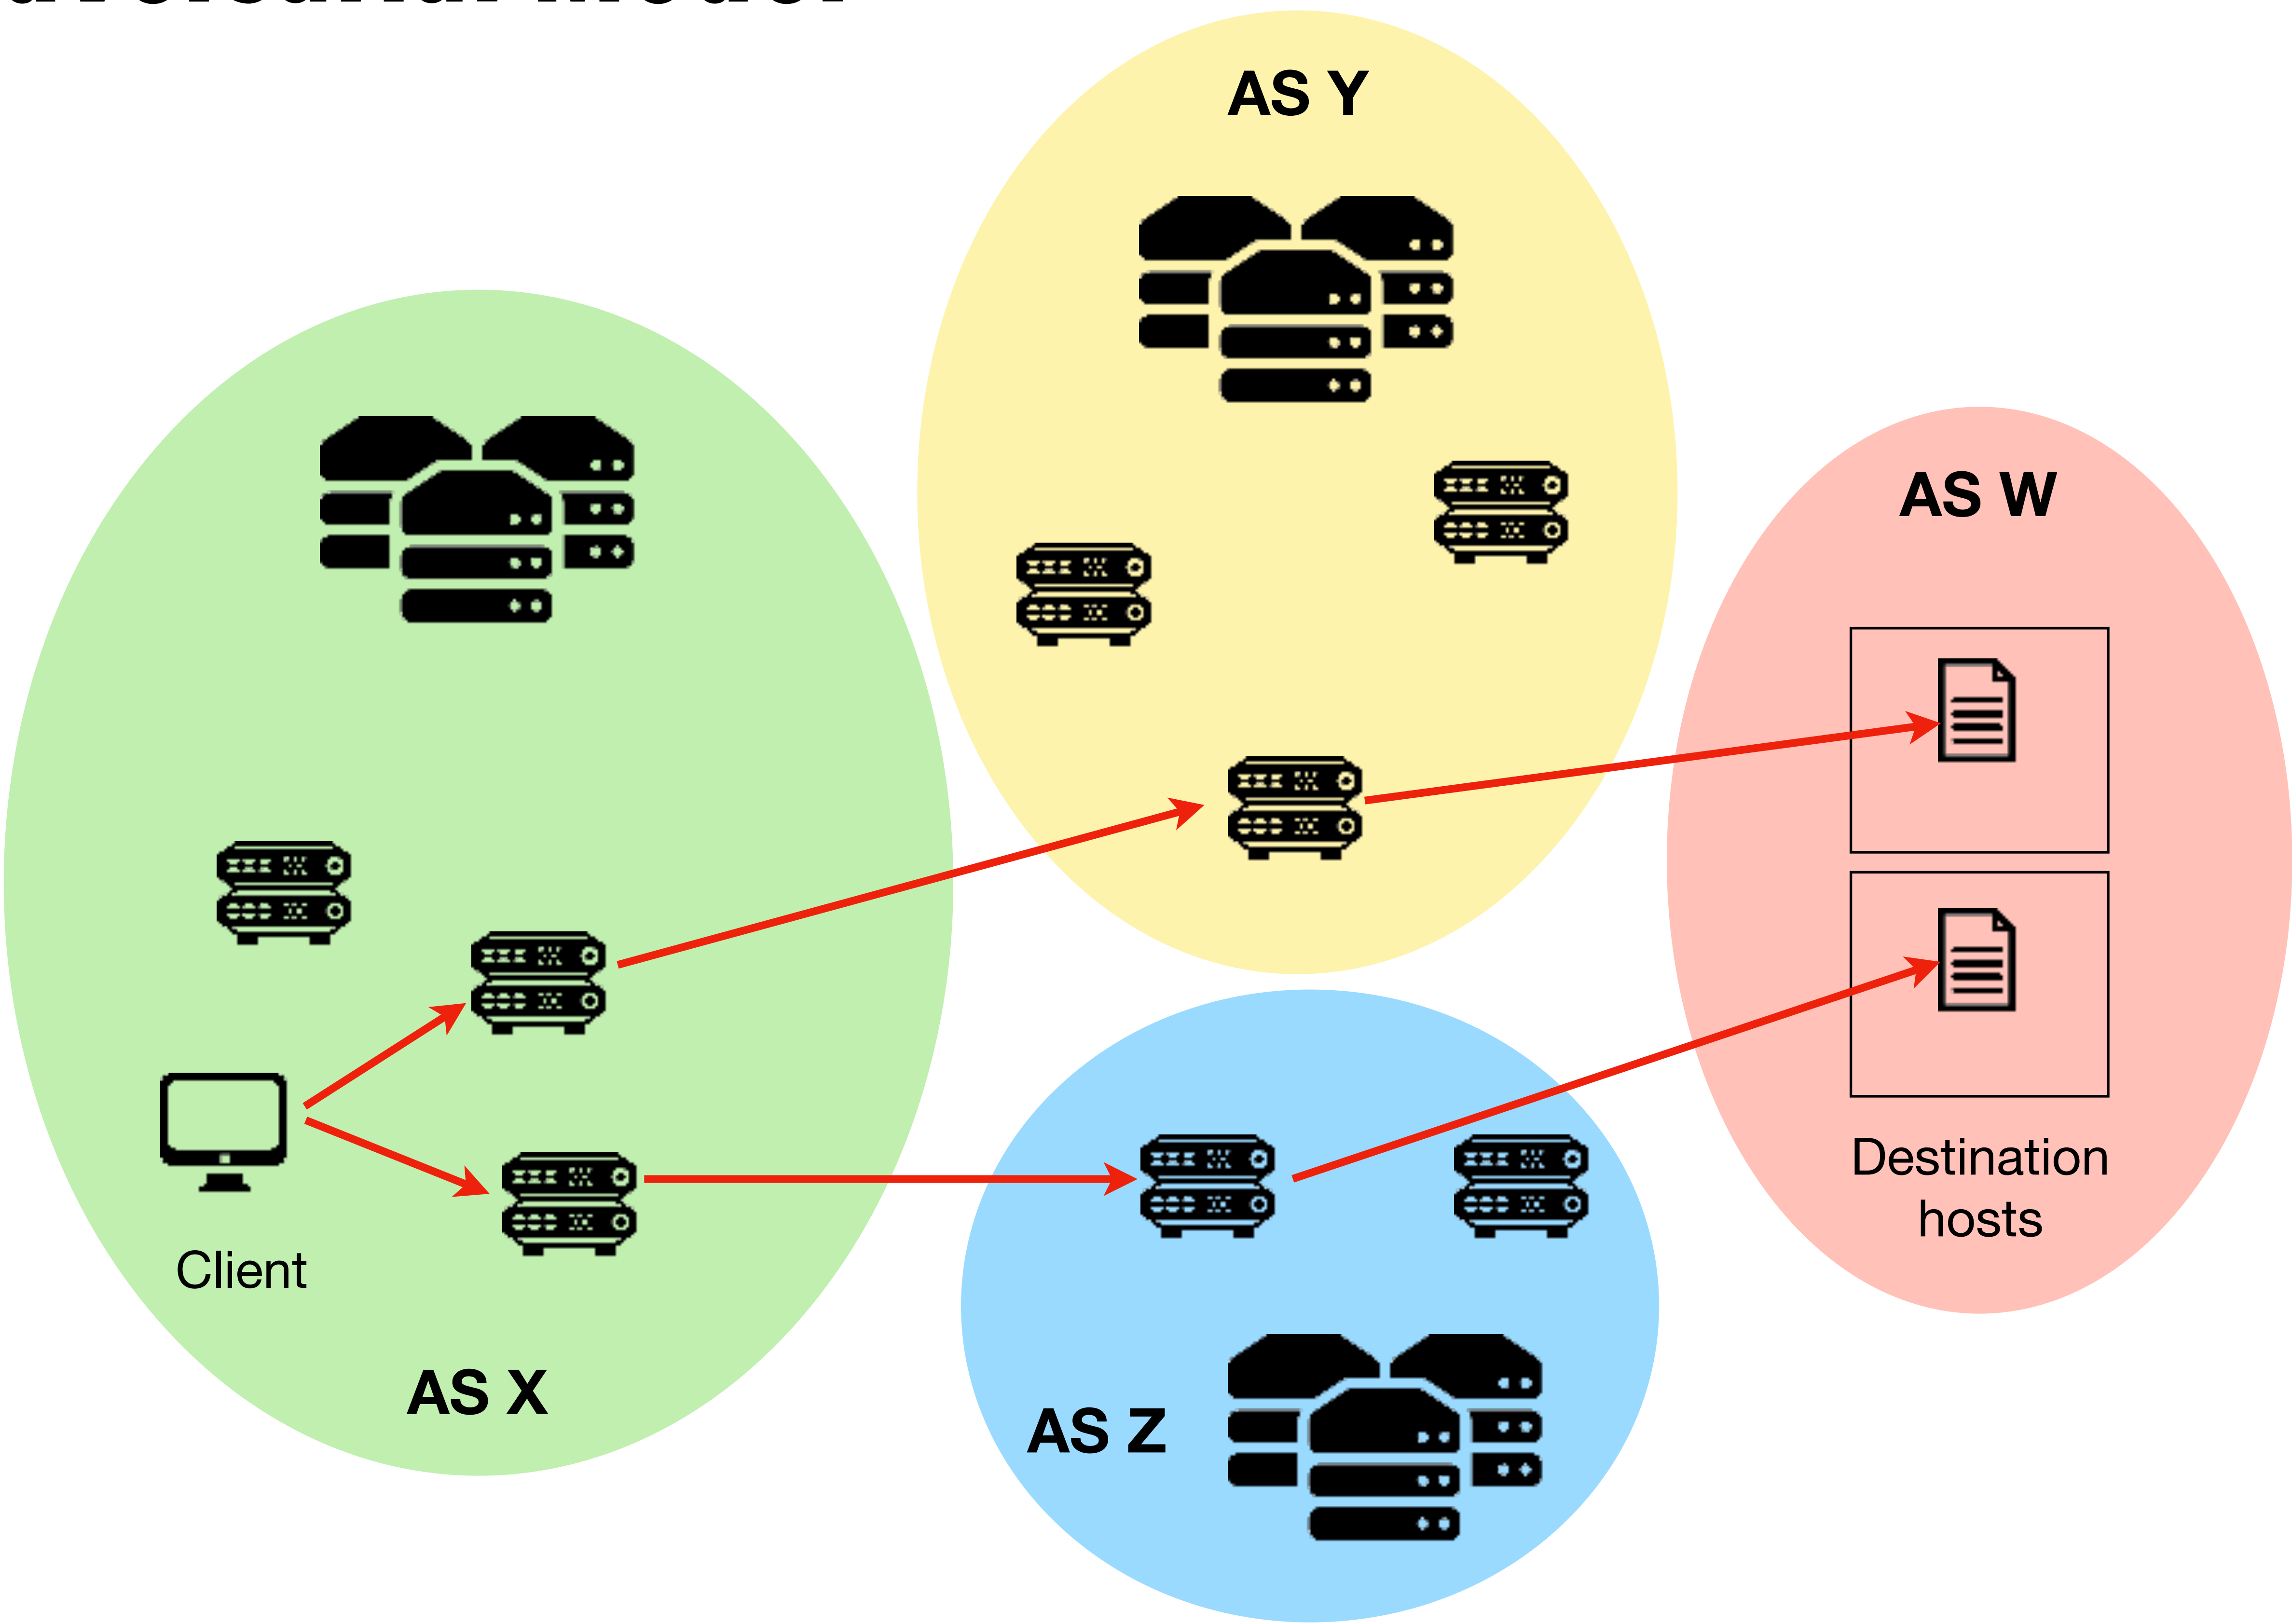

Adversarial Model

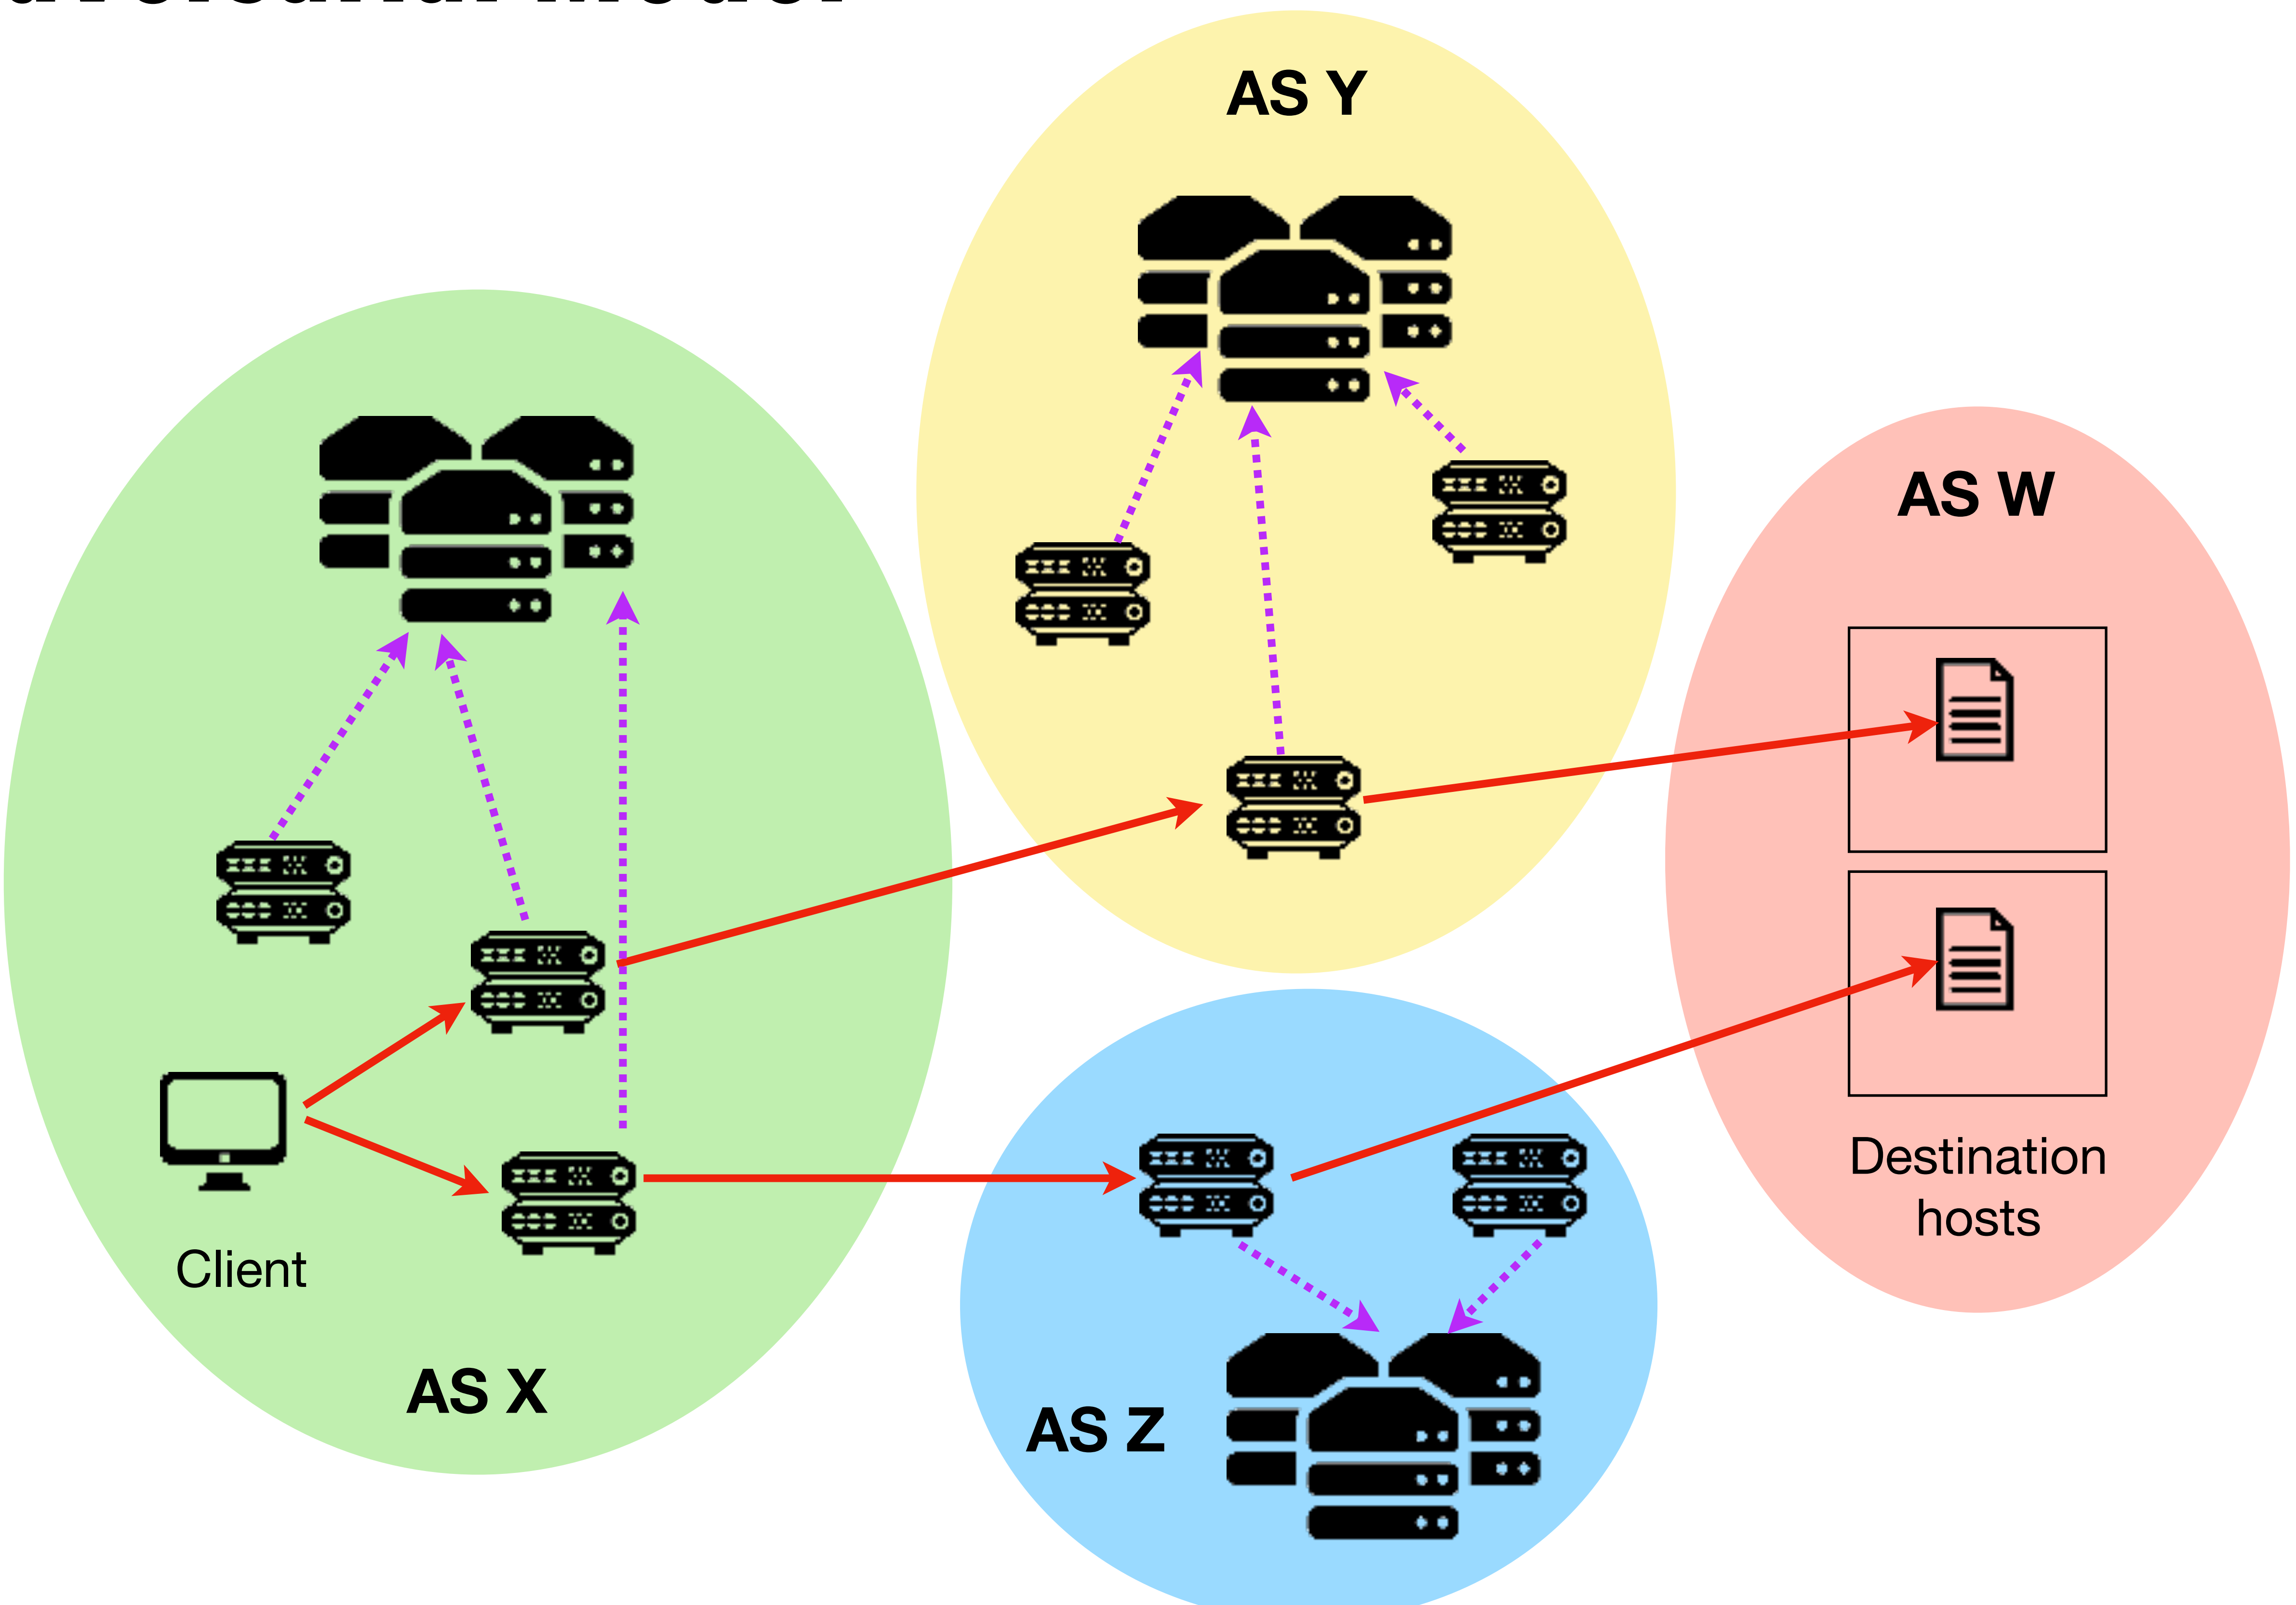

Adversarial Model

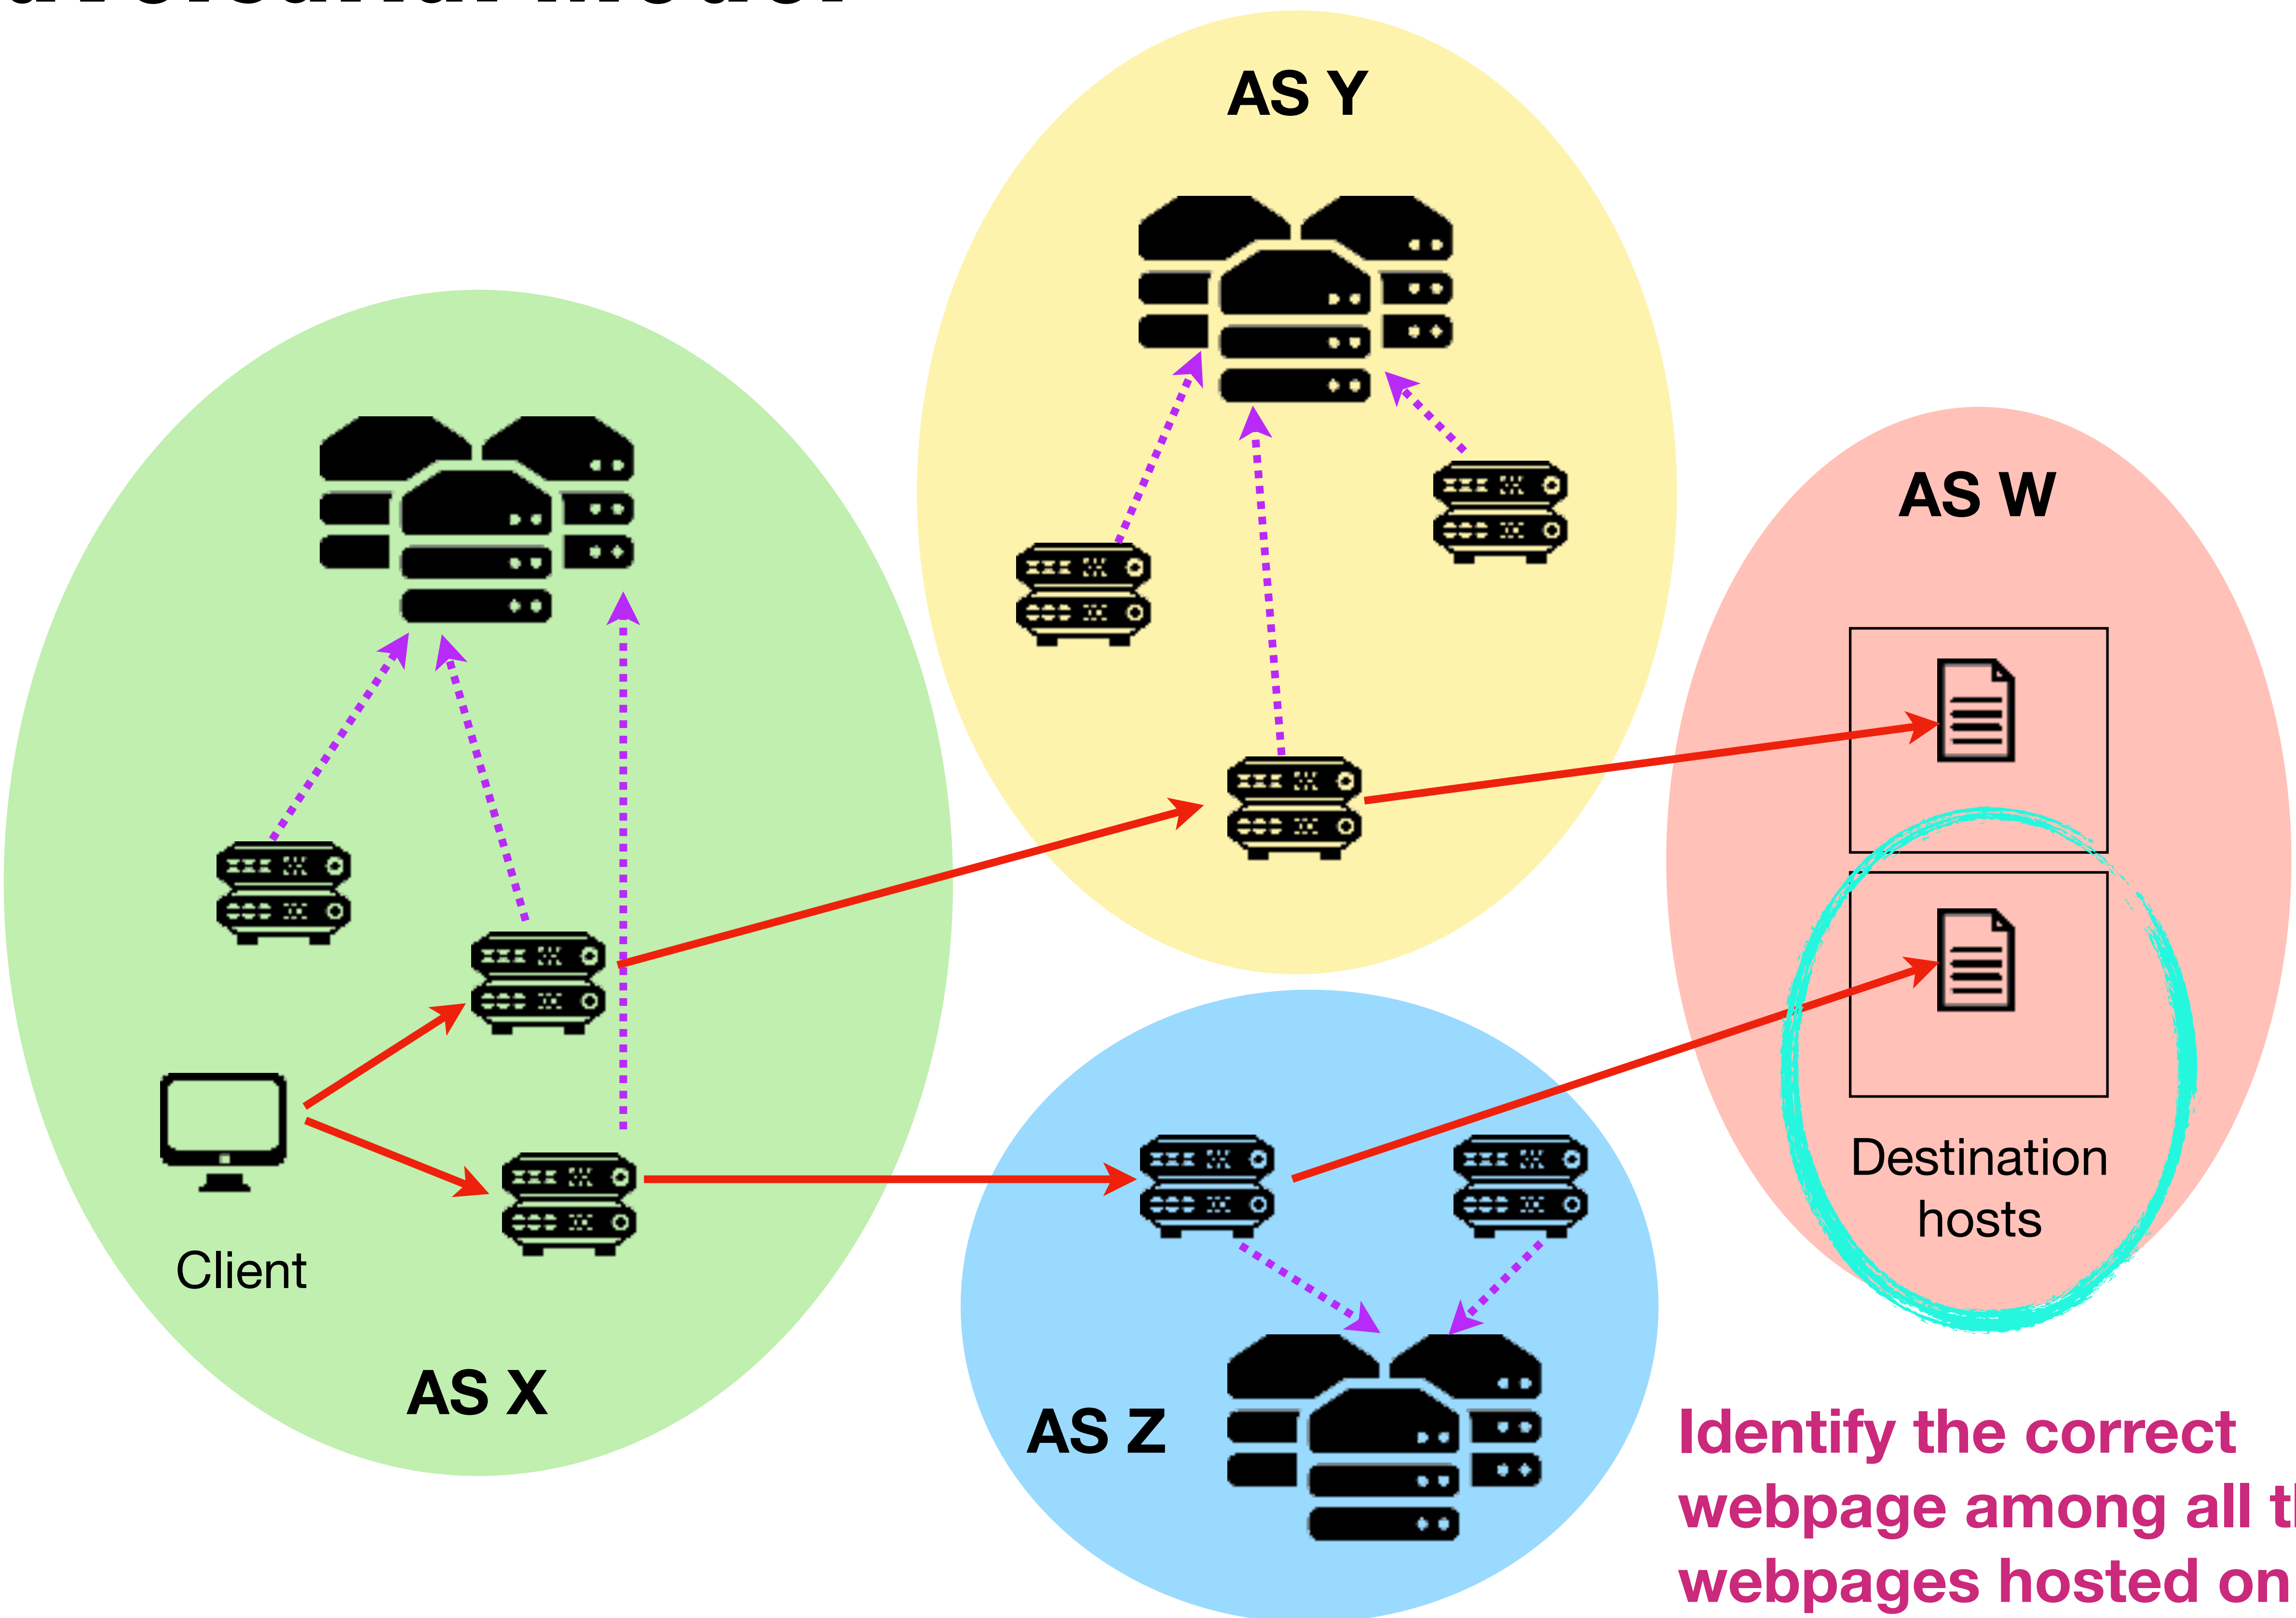

Identify the correct webpage among all the webpages hosted on an IP

Dataset

QUIC-dominant dataset (150 pages)

~70% QUIC on average

Built from lists of popular domains

Process

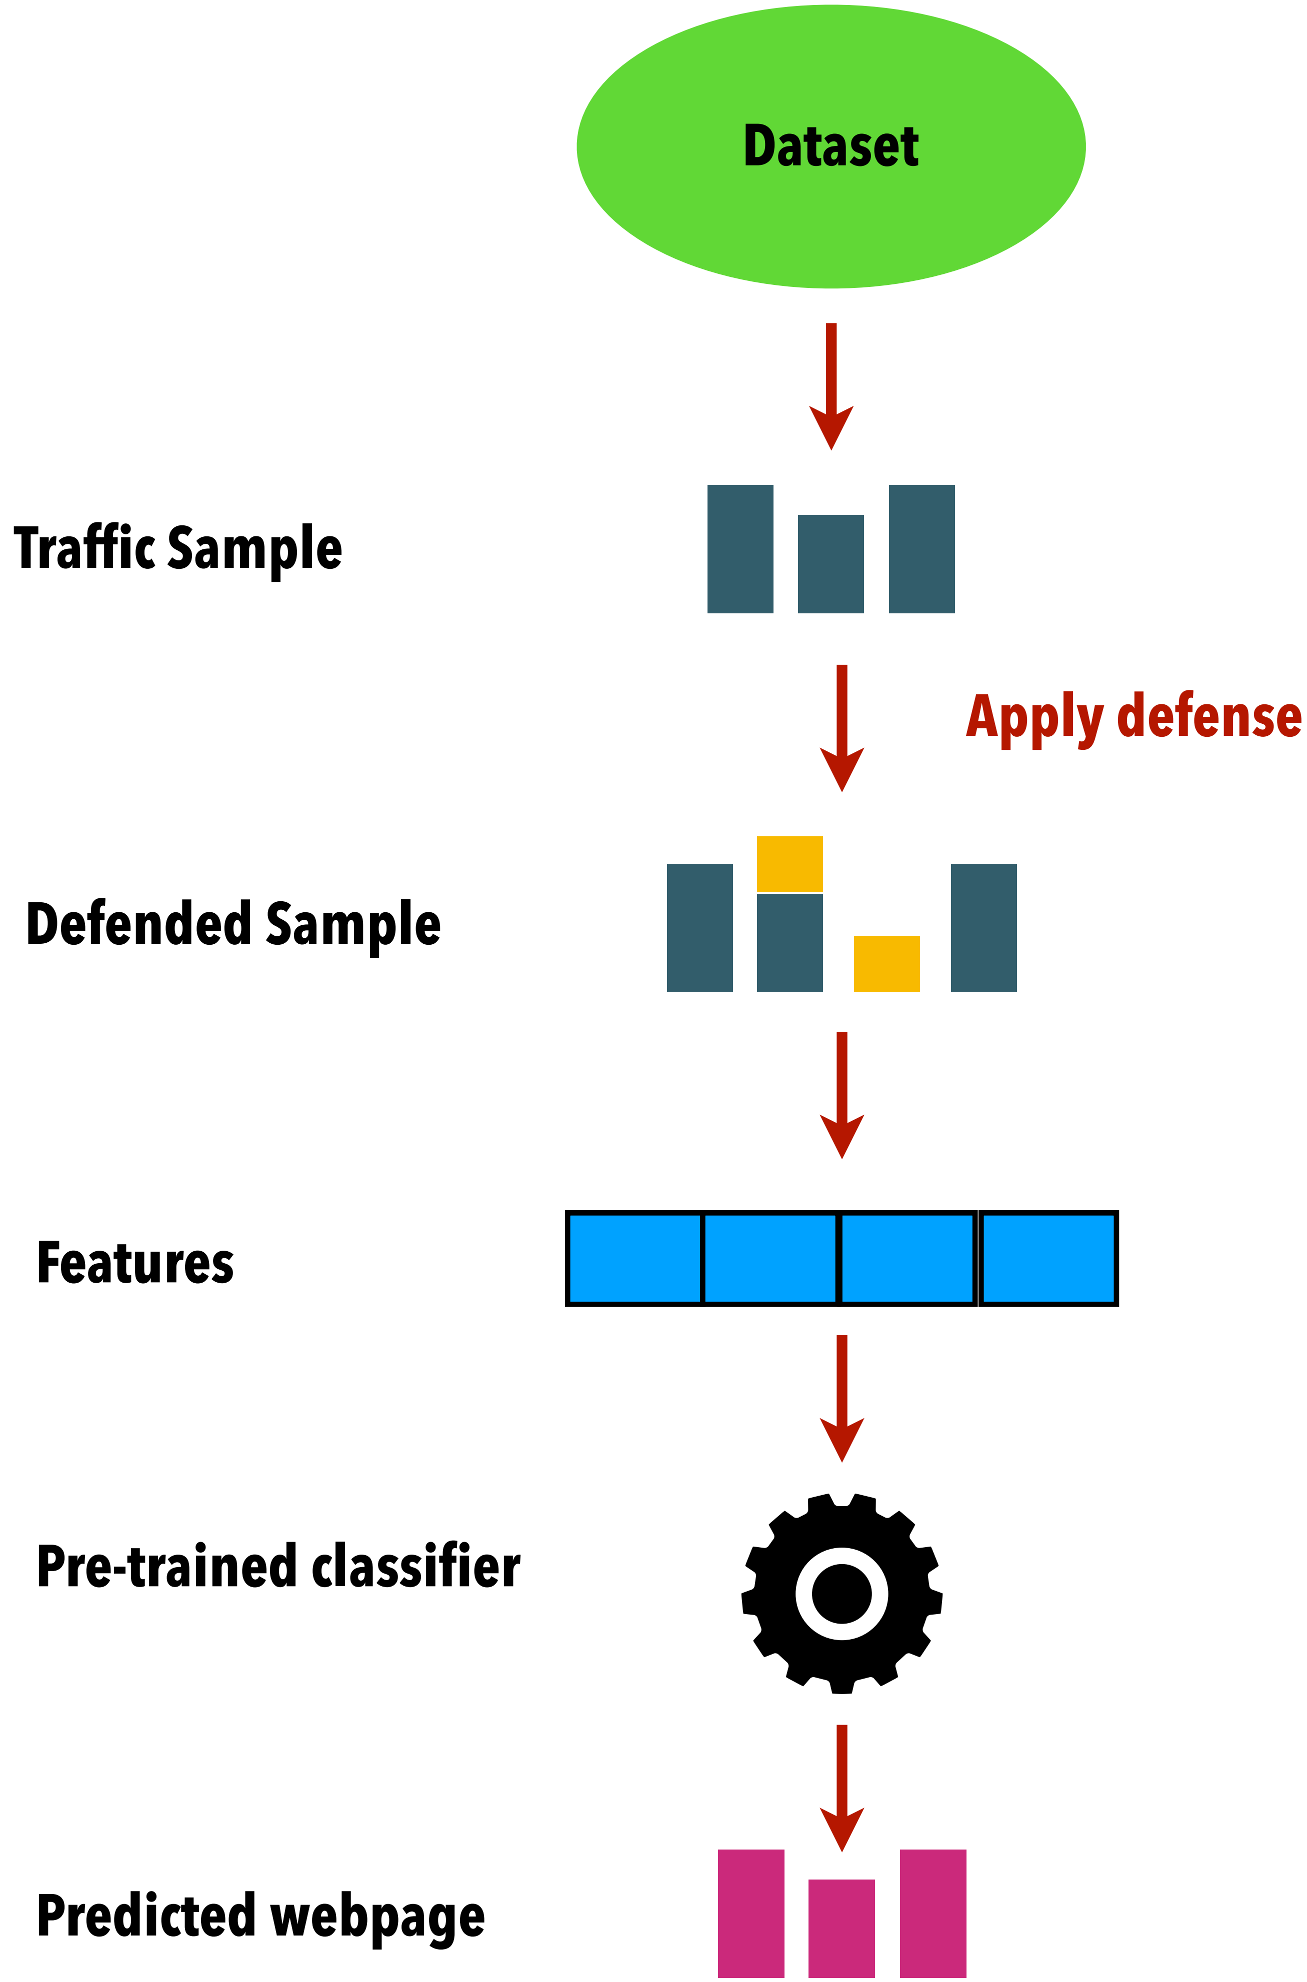

k-FP features + Random Forest / VarCNN

Unconstrained Adversary

Undefended Traffic: 96% F-Score

Unconstrained Adversary

Undefended Traffic: 96% F-Score

Random baseline: 0.67%

Unconstrained Adversary

Undefended Traffic: 96% F-Score

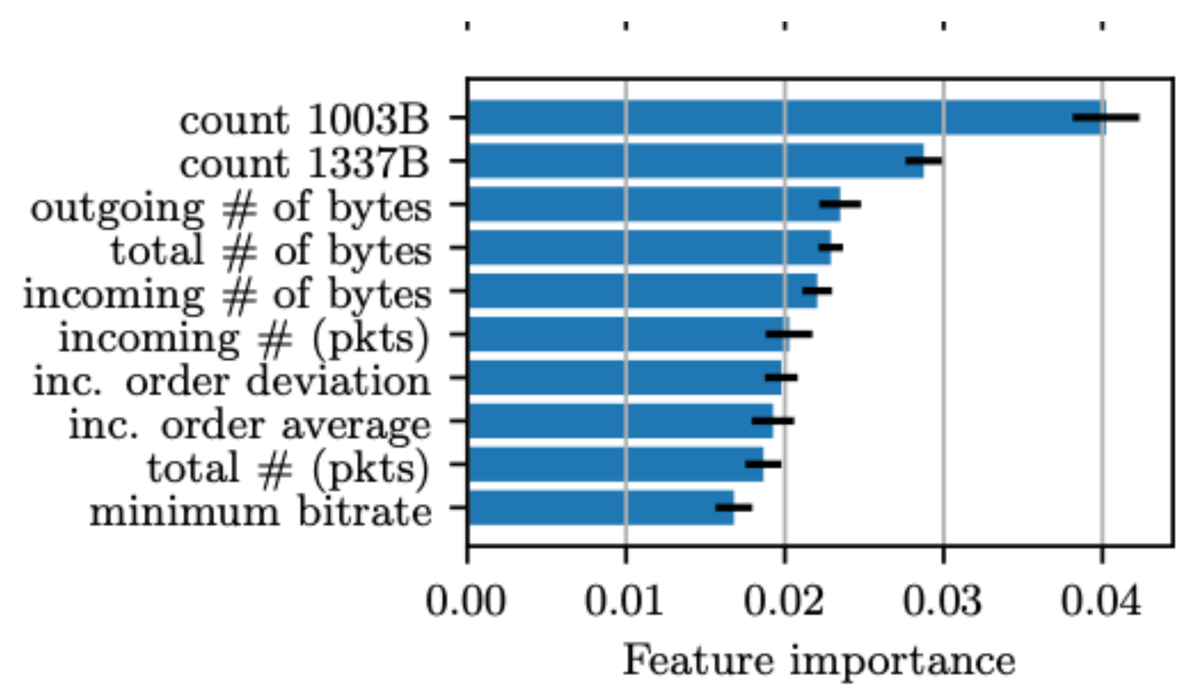

Size-based features are important

Unconstrained Adversary

Undefended Traffic: 96% F-Score

Hide packet-based features

(pad individual packets)

~94%

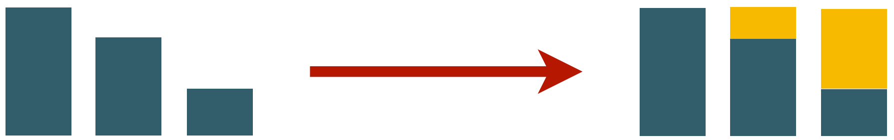

Unconstrained Adversary

Undefended Traffic: 96% F-Score

Hide packet-based features
(pad individual packets)

~94%

Hide trace-based features
(pad total size)

~92%

Unconstrained Adversary

Undefended Traffic: 96% F-Score

Hide packet-based features

(pad individual packets)

~94%

Hide trace-based features

(pad total size)

~92%

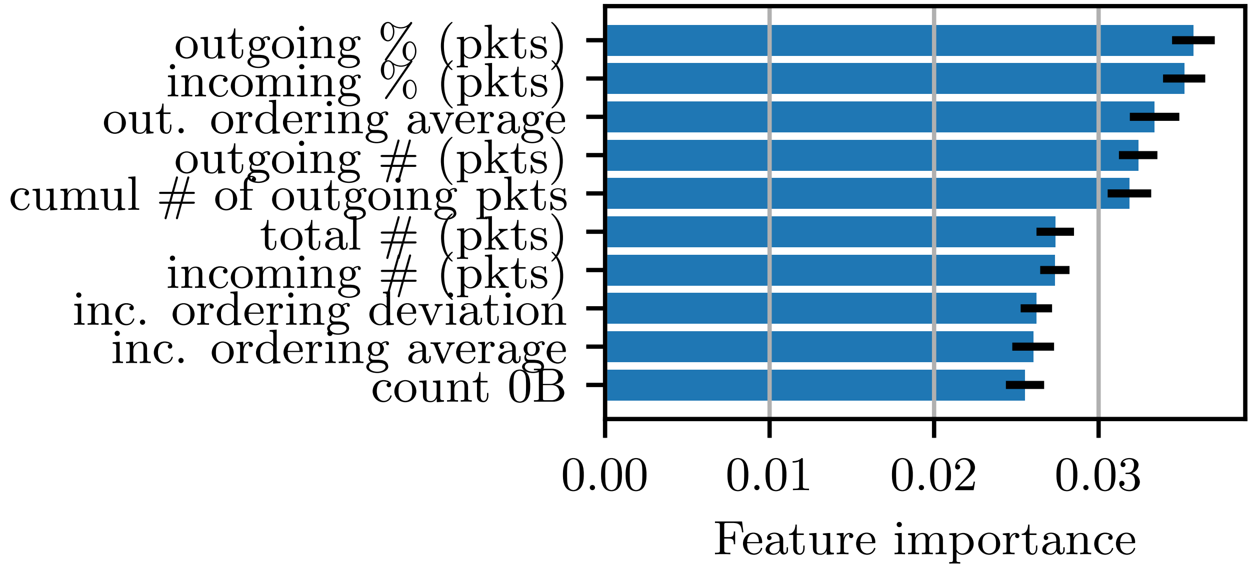

Directionality-based features become important

Unconstrained Adversary

Undefended Traffic: 96% F-Score

Dummy injection

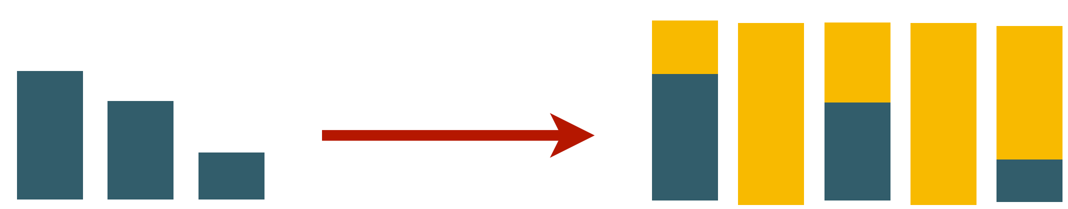

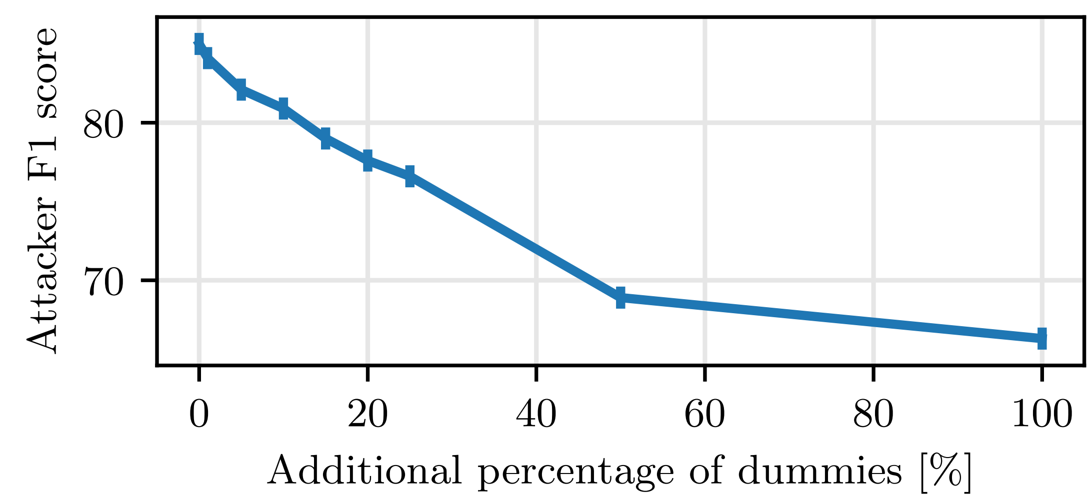

Hide packet-based features

(pad individual packets)

~94%

Hide trace-based features

(pad total size)

~92%

Unconstrained Adversary

Undefended Traffic: 96% F-Score

Network defenses offer low protection with high costs

Ex: For 10% reduction in F-Score, we need >50% overhead

~94%

Hide global features
(pad total size)

~92%

Constrained Adversary: Limited view

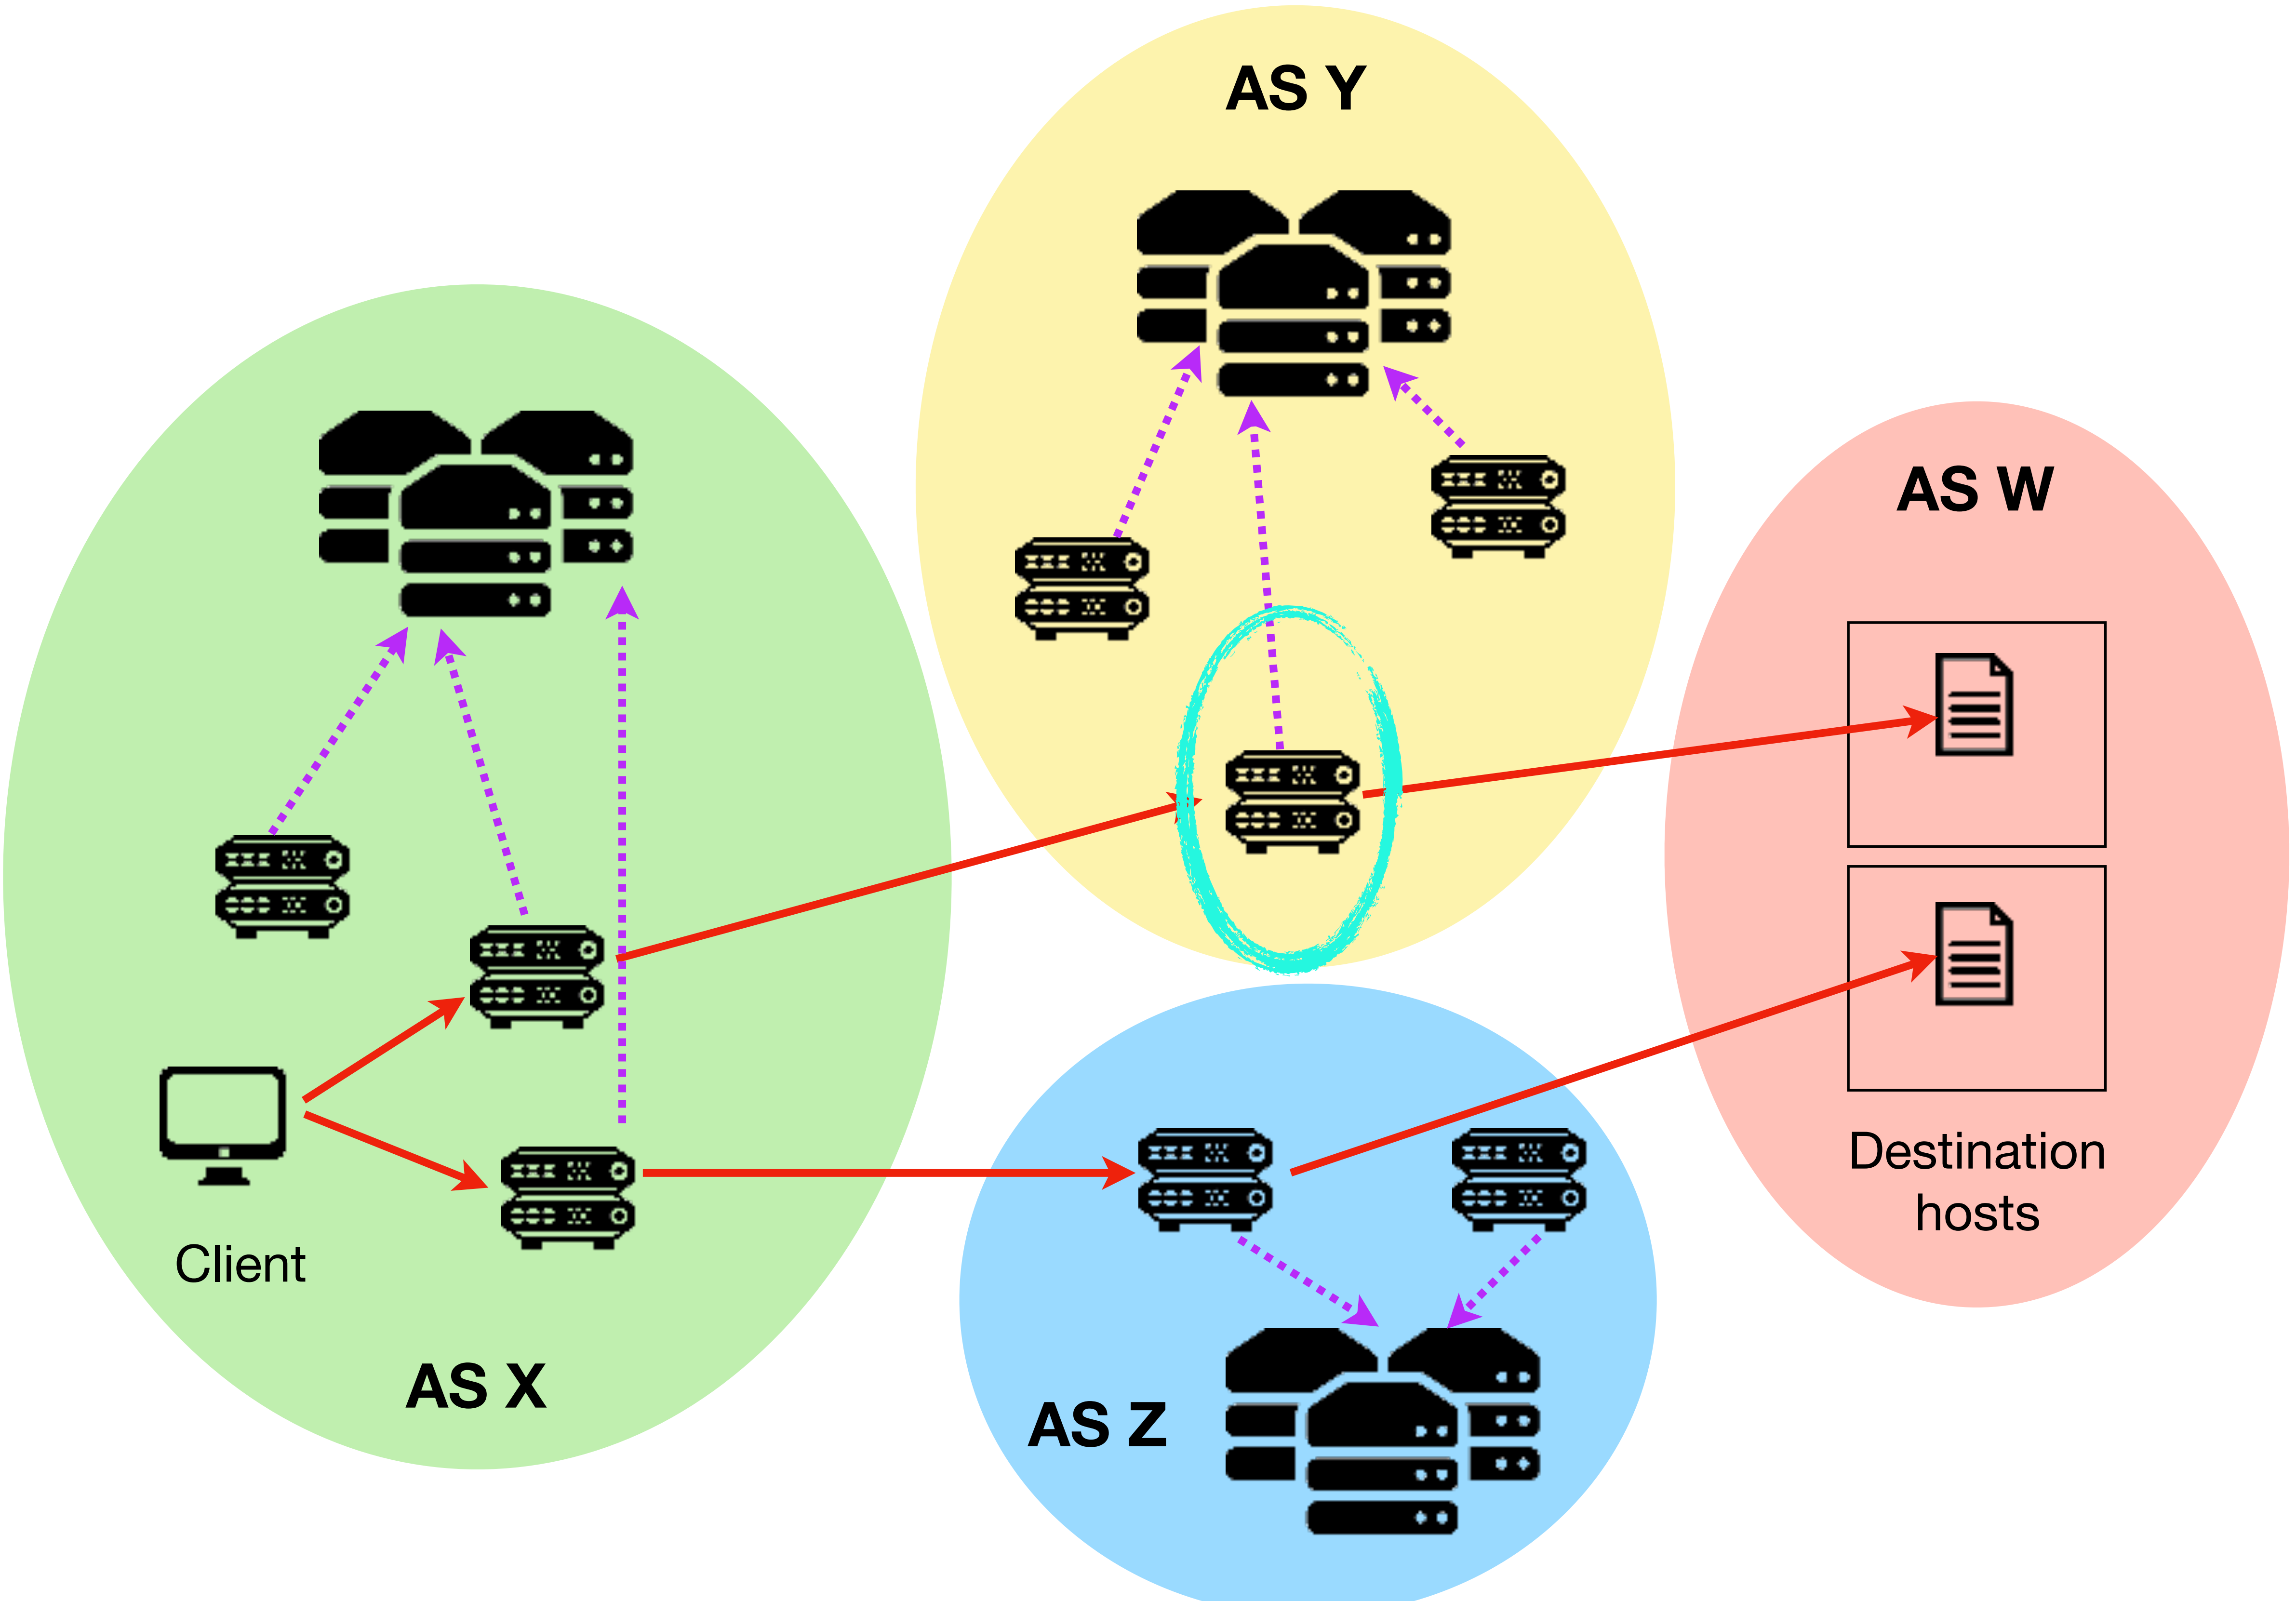

Constrained Adversary: Limited view

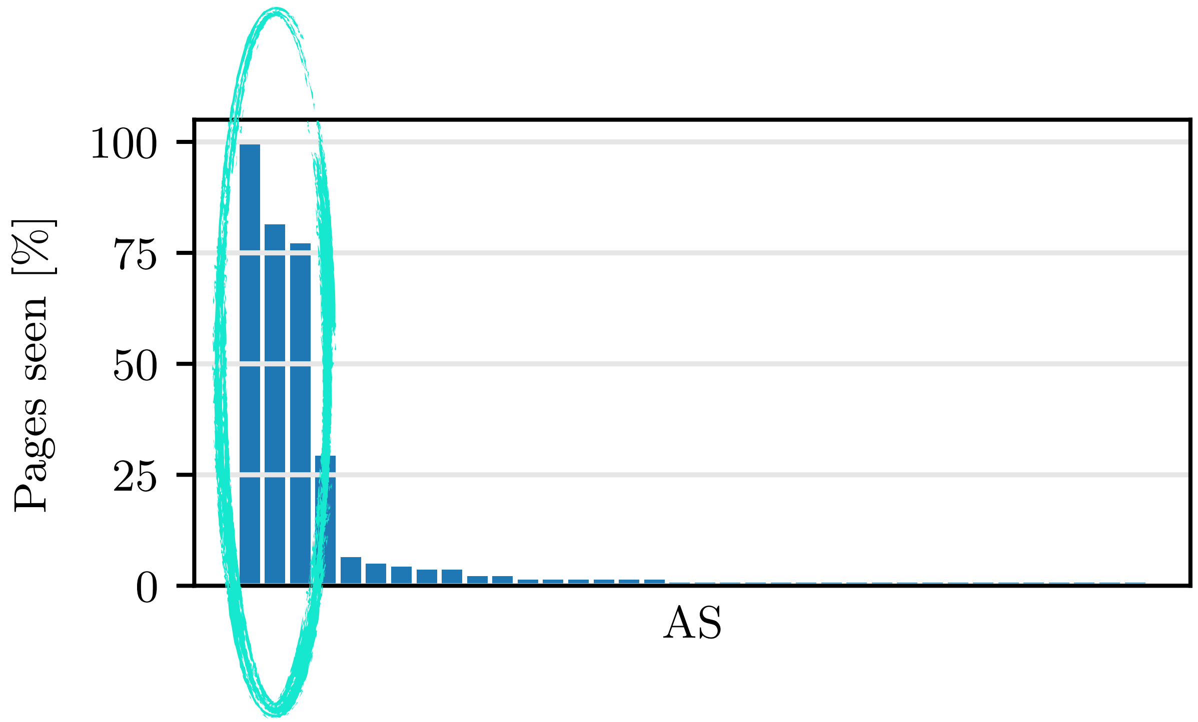

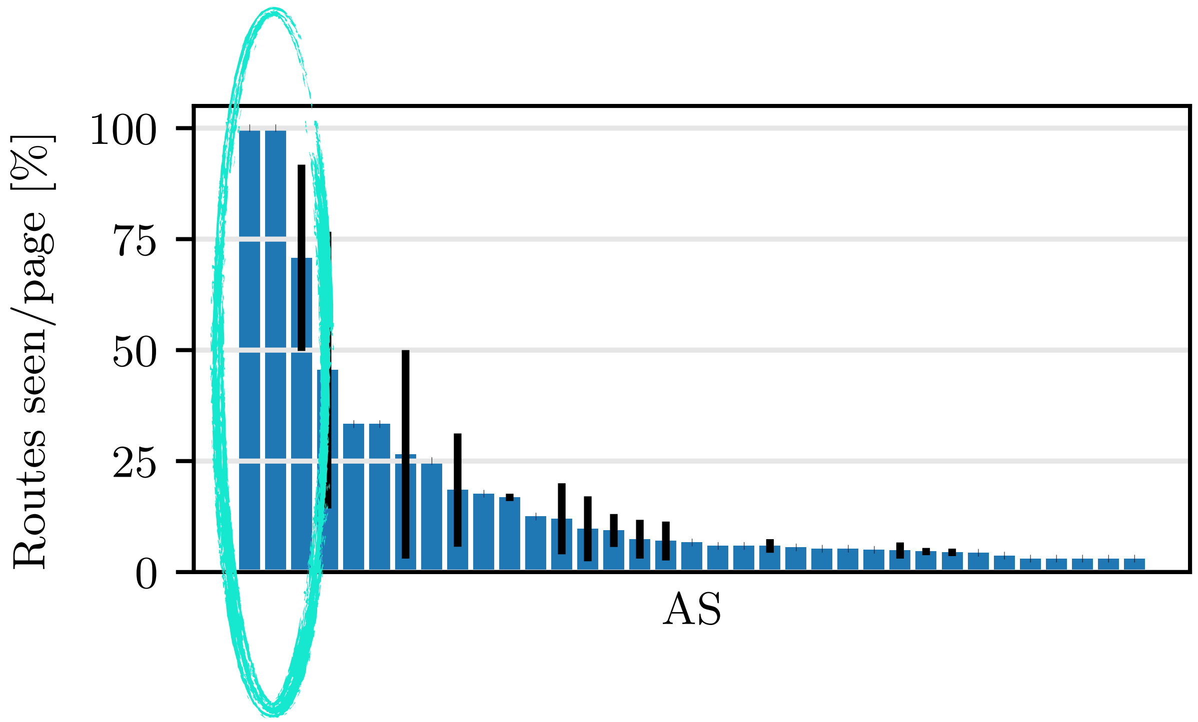

A few large ASes can successfully run attacks.

Constrained Adversary: Limited view

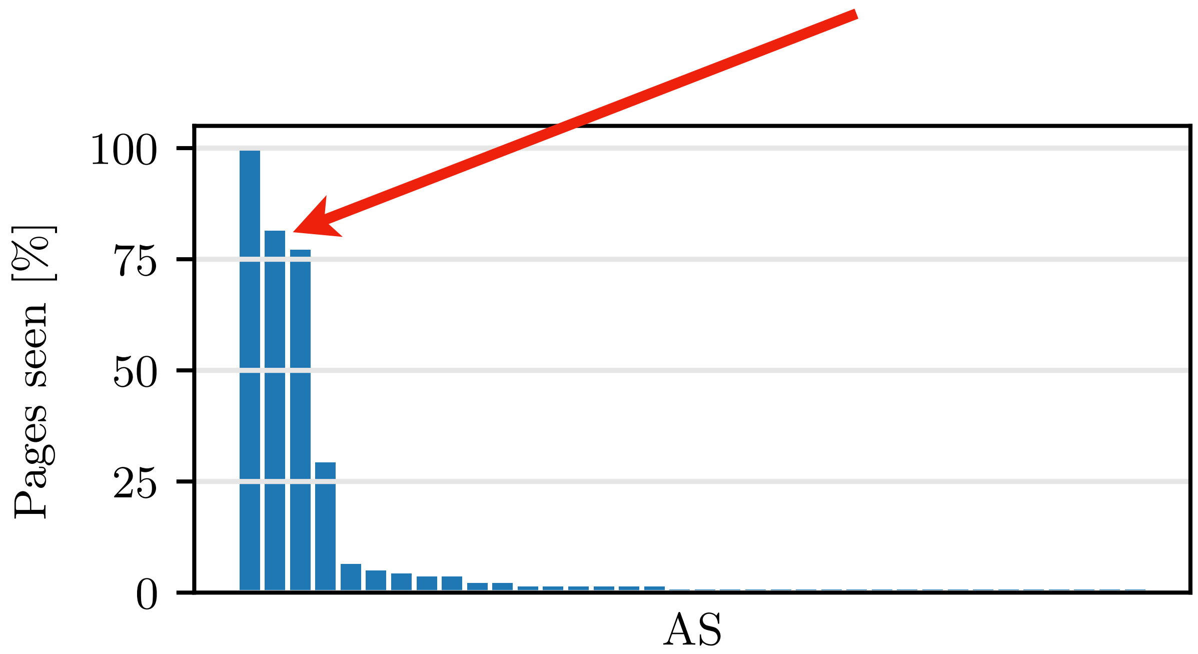

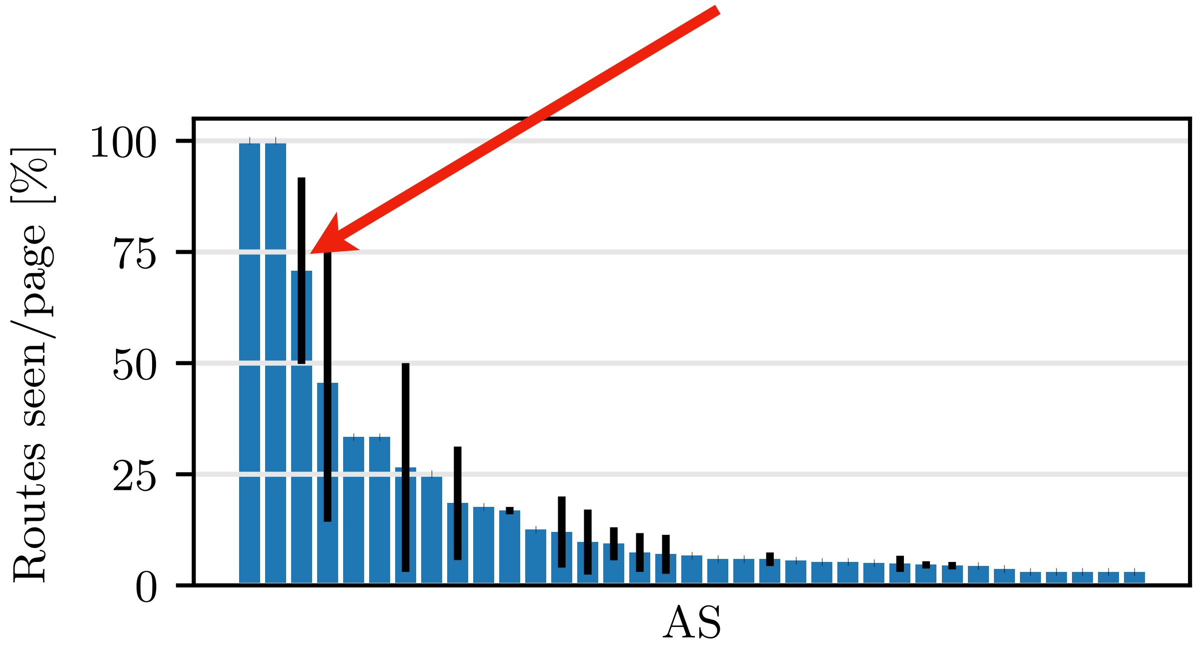

A few large ASes can successfully run attacks.

Timings to Google resources is a low-cost fingerprint: 77.9% F-score

Constrained Adversary: Limited processing

Adversary resources

- Unconstrained
- Constrained

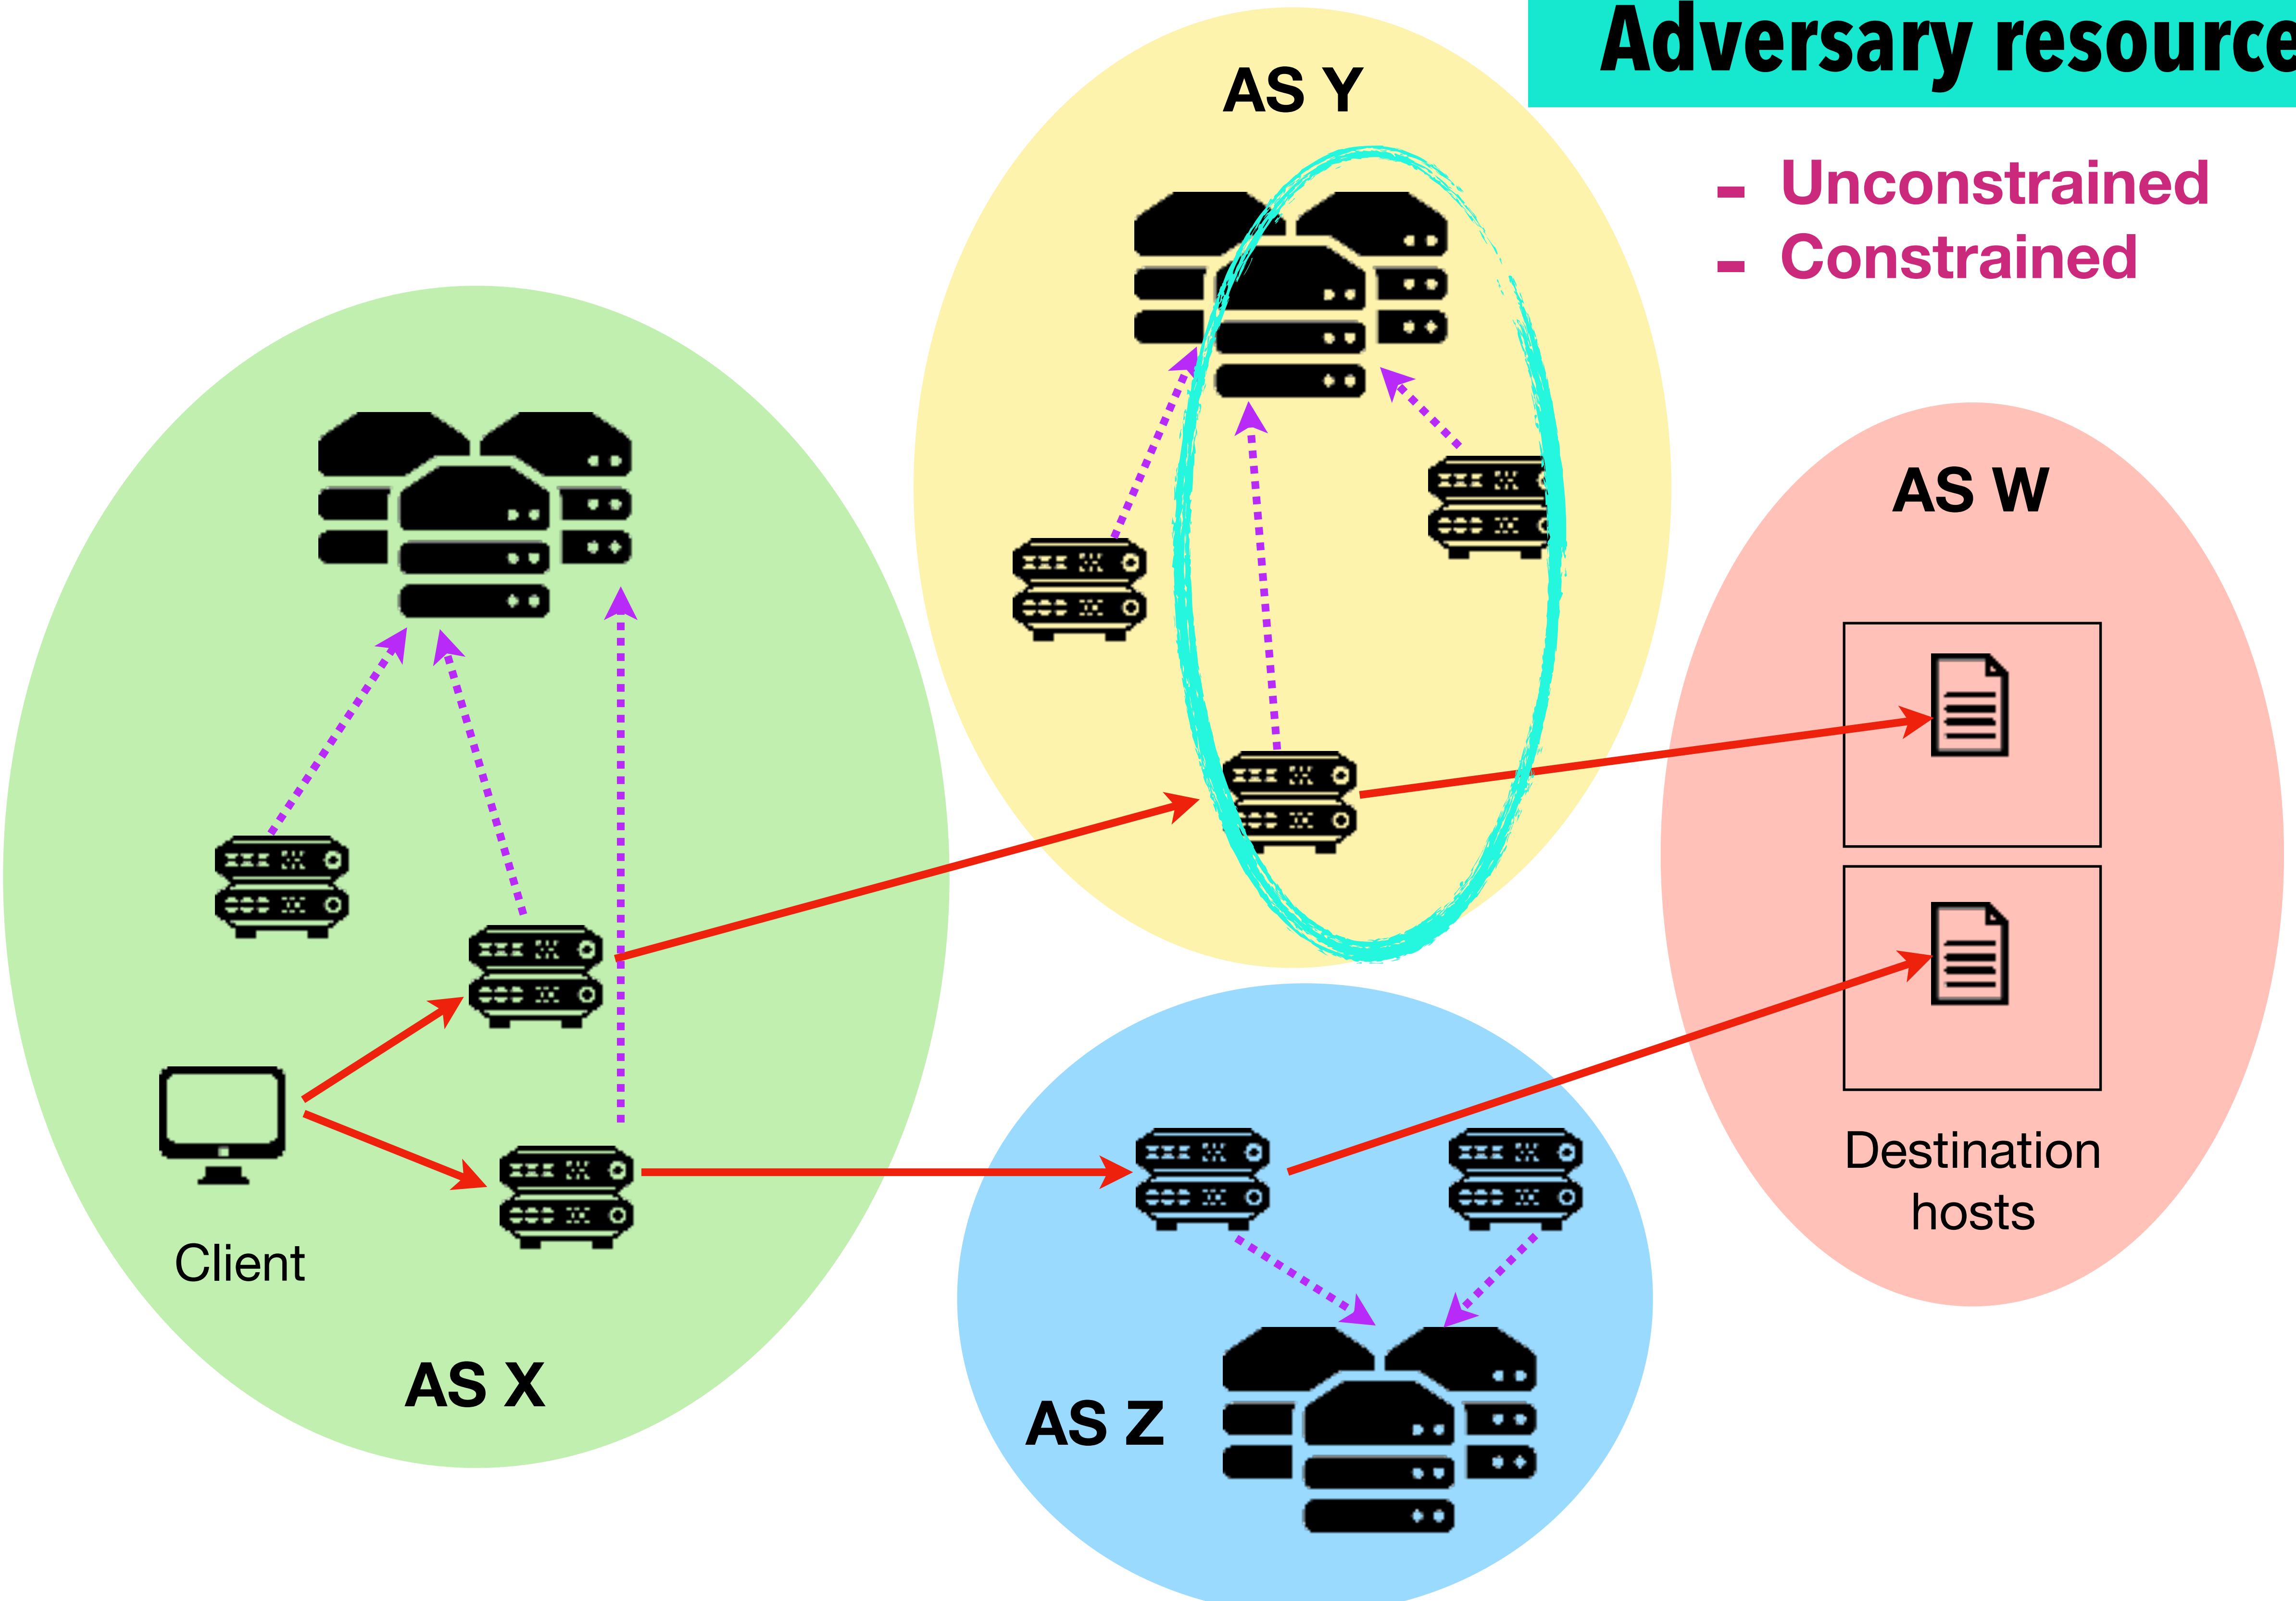

Constrained Adversary: Limited processing

Sampled NetFlow	Undefended (F-score)	Defended (F-Score)
NetFlow 100%	90.5	53.1
NetFlow 10%	66.4	33.1
NetFlow 1%	41.7	21.6
NetFlow 0.1%	16.8	8.6

Constrained Adversary: Limited processing

Sampled NetFlow	Undefended (F-score)	Defended (F-Score)
NetFlow 100%	90.5	53.1
NetFlow 10%	66.4	33.1
NetFlow 1%	41.7	21.6
NetFlow 0.1%	16.8	8.6

Random baseline: 0.67%

Constrained Adversary: Limited processing

Sampled NetFlow	Undefended (F-score)	Defended (F-Score)
NetFlow 100%	90.5	53.1
NetFlow 10%	66.4	33.1
NetFlow 1%	41.7	21.6
NetFlow 0.1%	16.8	8.6

Most of the privacy gain comes from the sampling process than the defense.

Network layer

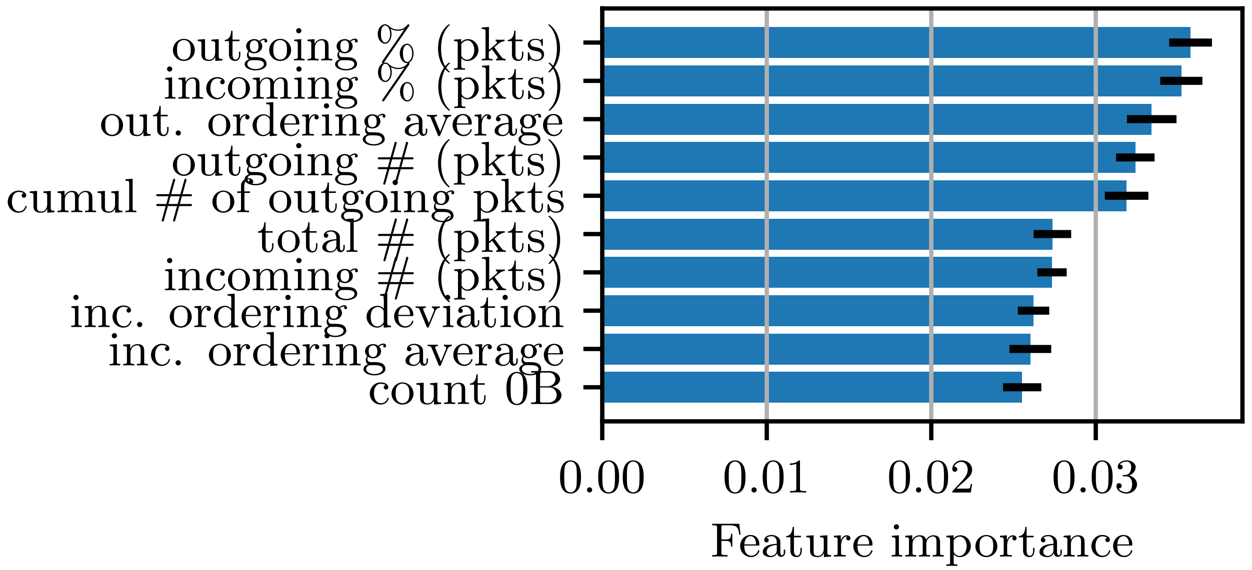

Network layer defenses cannot efficiently hide global features without application layer information.

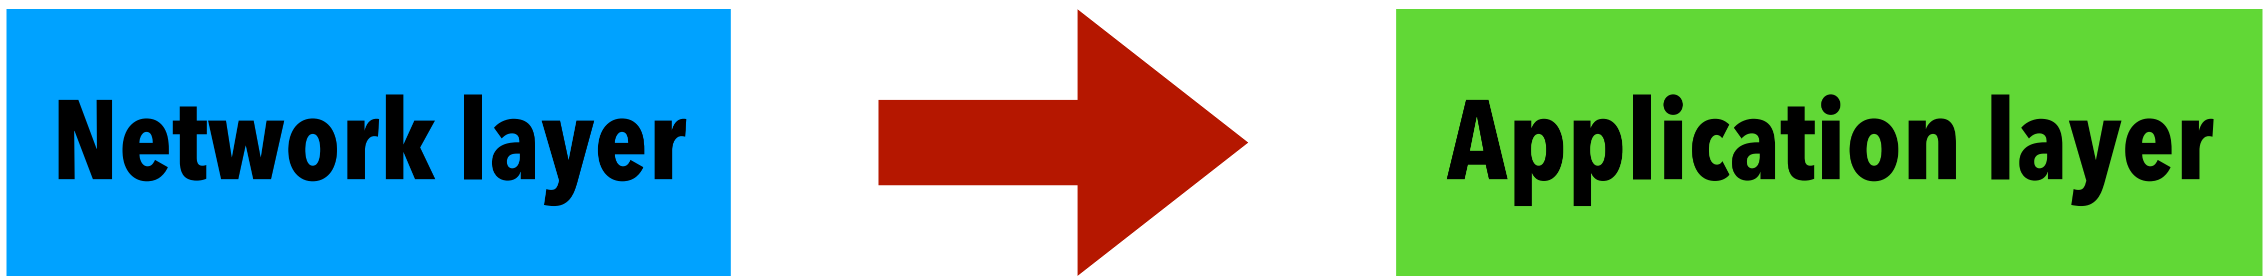
Network layer
Application layer

Analysing web pages

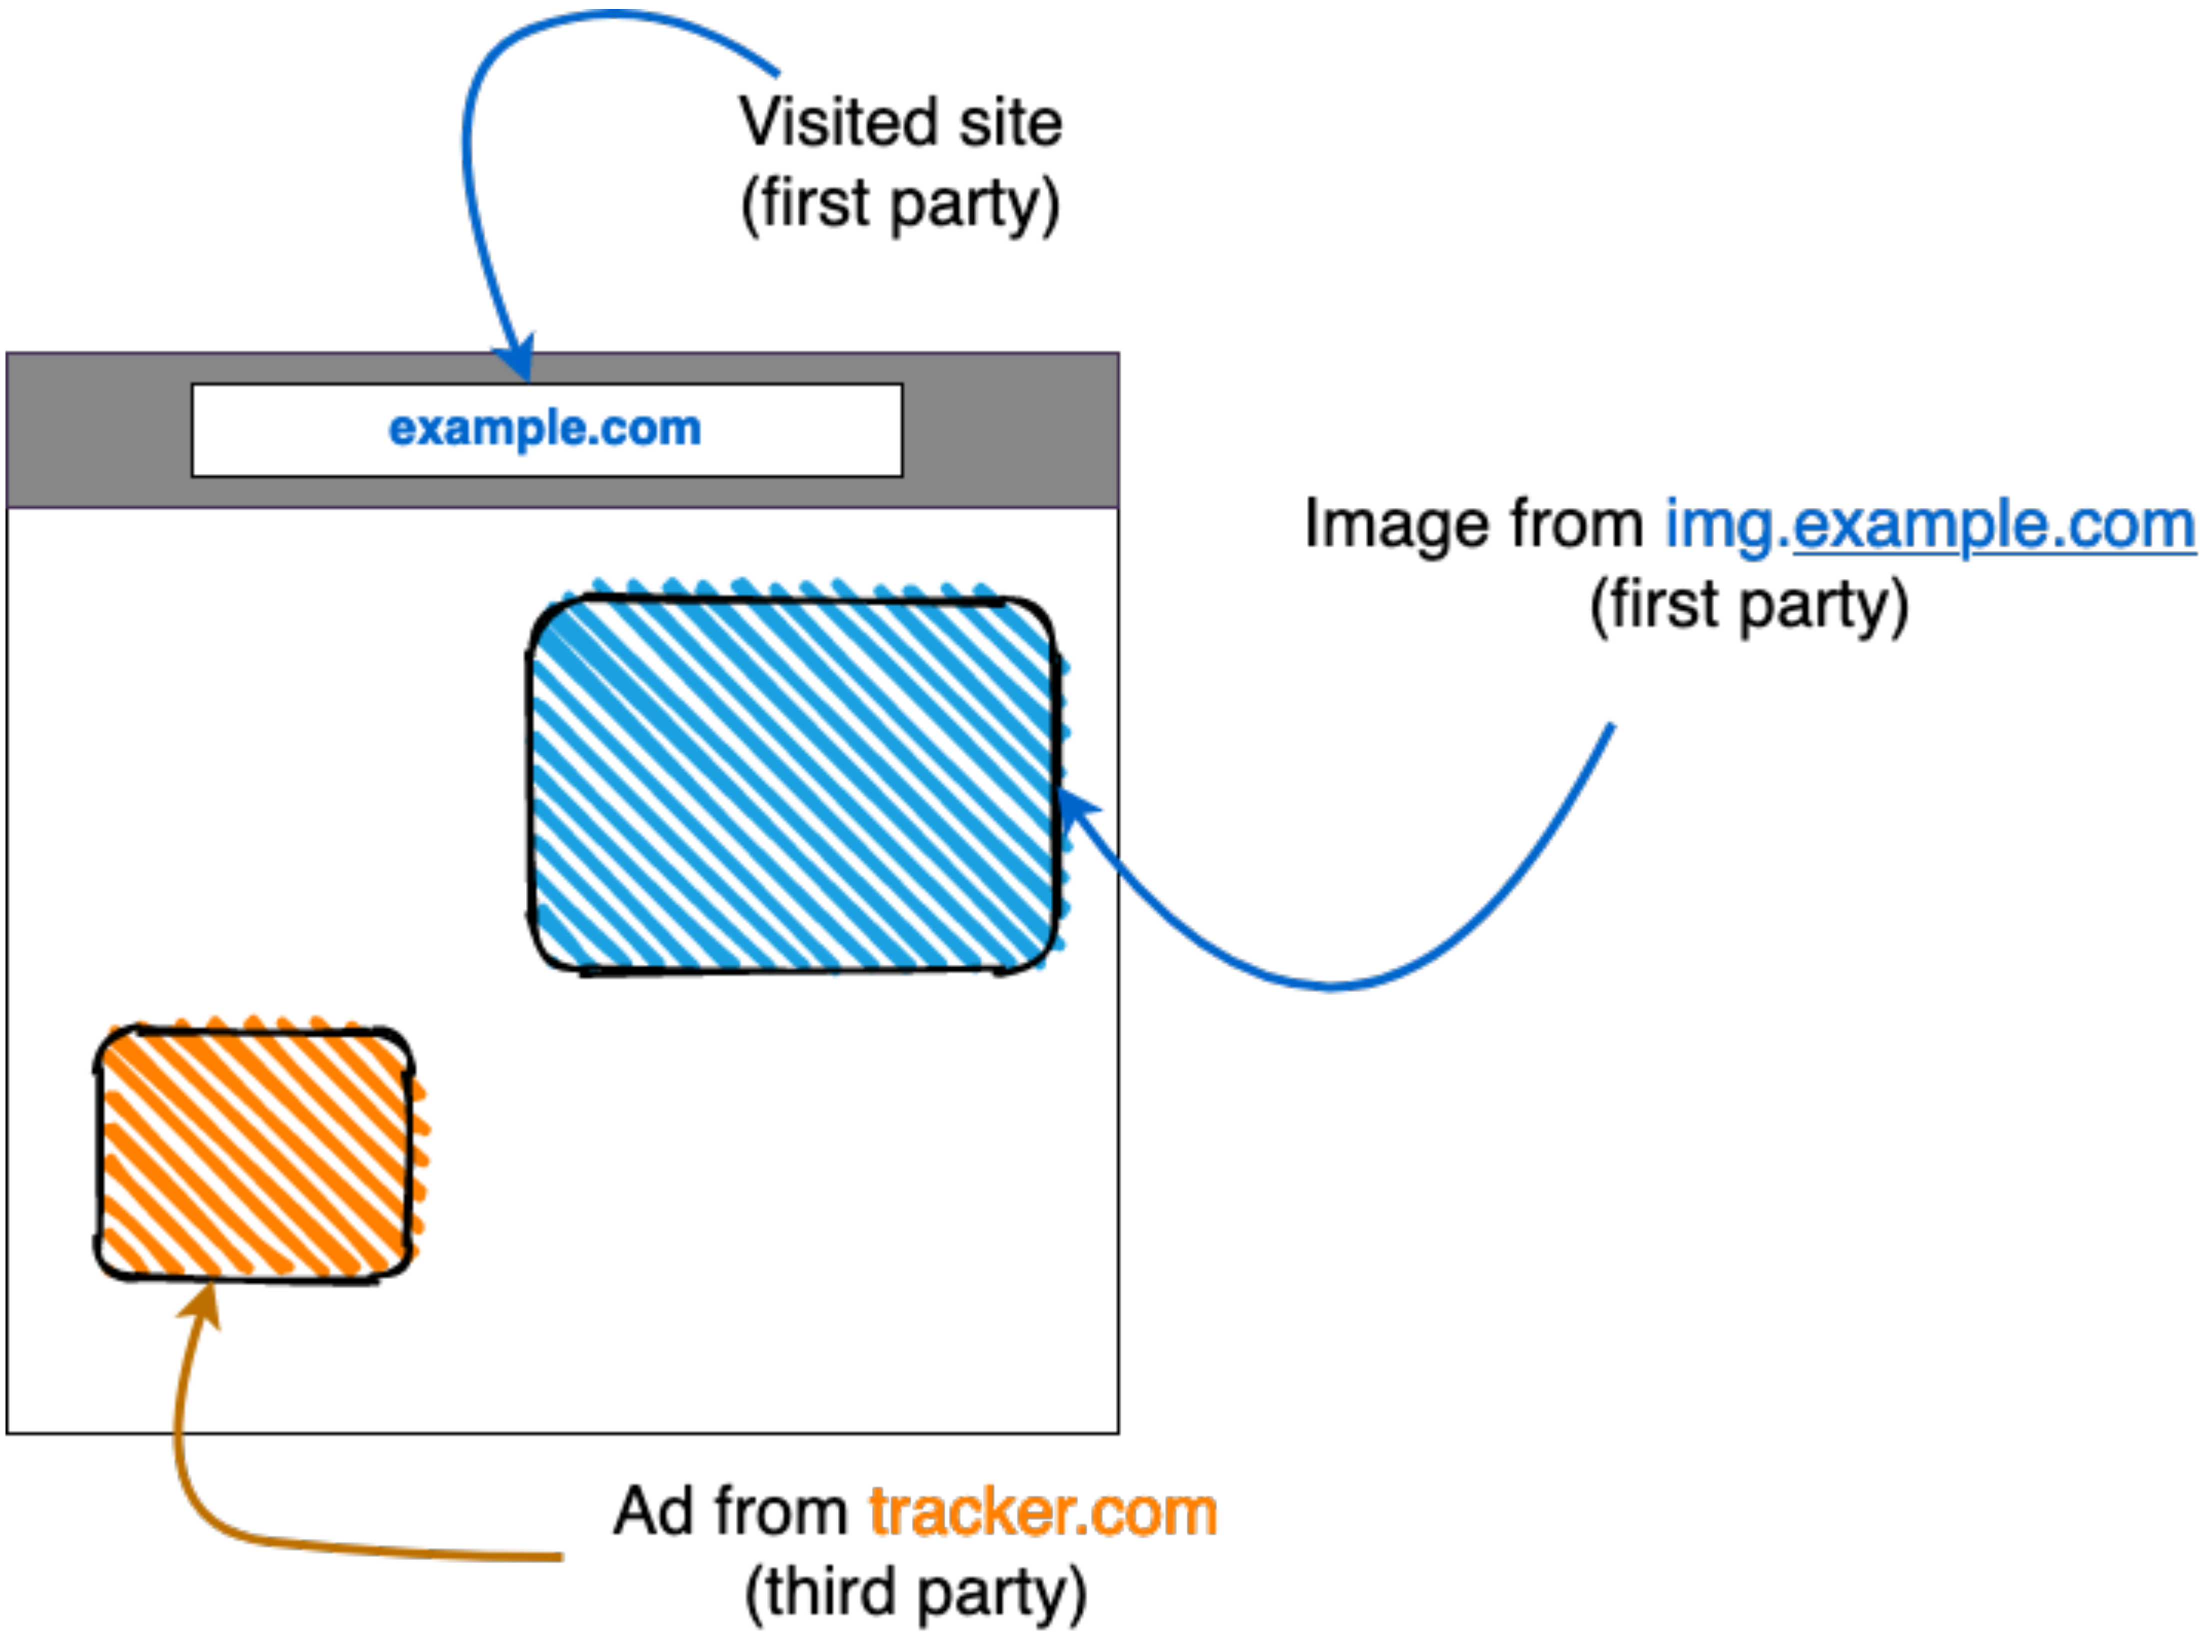

Analysing web pages

18% of pages have < 20% first party resources

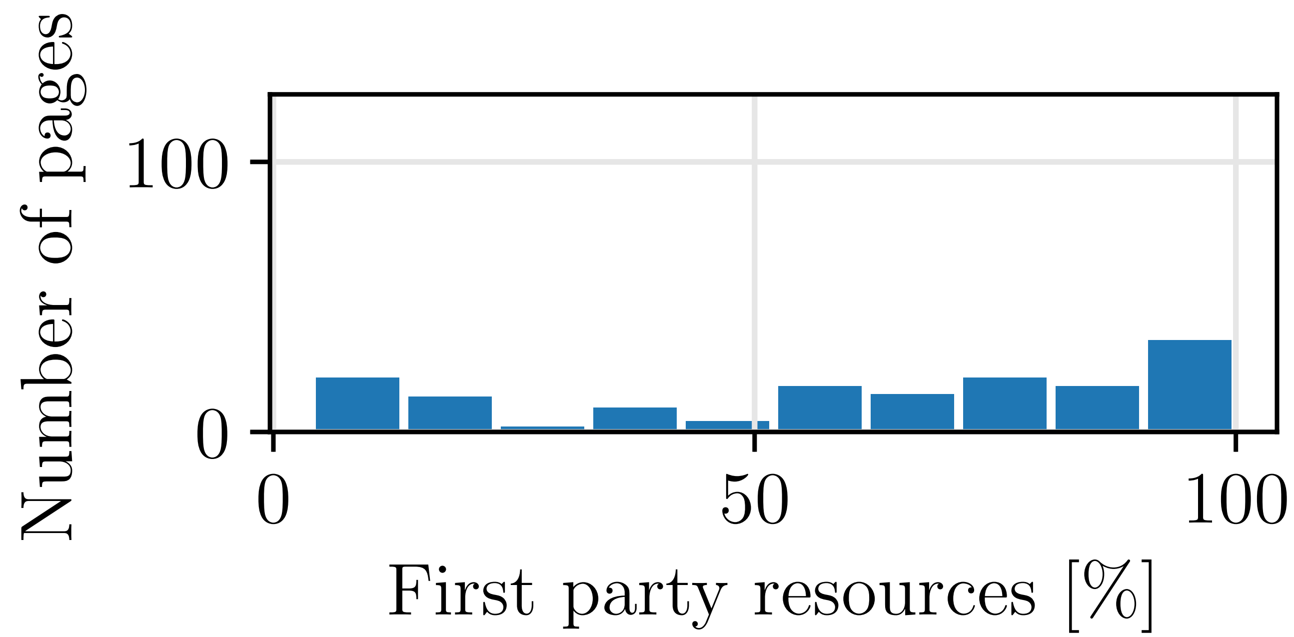

24% of pages have > 50% Google resources

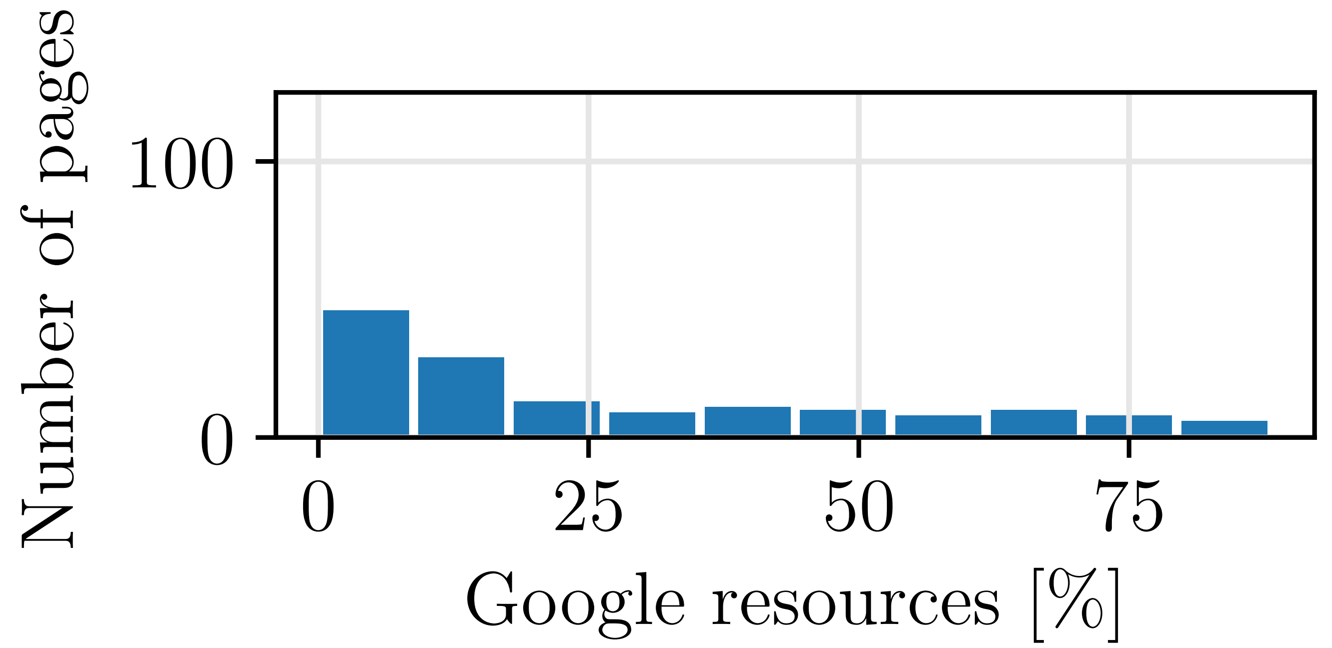

Analysing web pages

18% of pages have < 20% first party resources

24% of pages have > 50% Google resources

Third parties contribute a large proportion of resources

All parties must participate in the protection of resources

Application layer defenses

- **Packet-based and Trace-based padding**

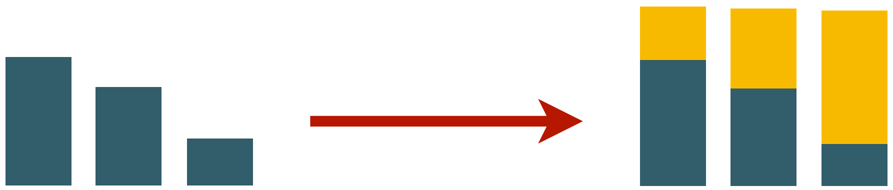

Padding is, once again, ineffective

Best case: reduces F-Score by 16% with total cost of ~8MB per resource

Application layer defenses

- **Packet-based and Trace-based padding**

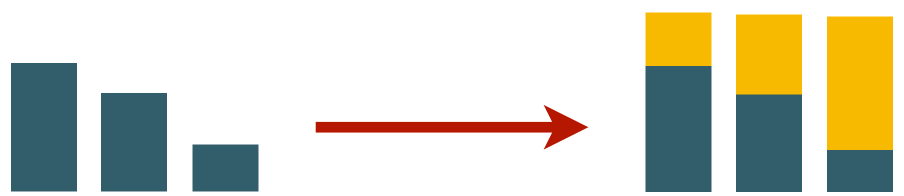

Padding is, once again, ineffective

Best case: reduces F-Score by 16% with total cost of ~8MB per resource

- **Dummy Injection**

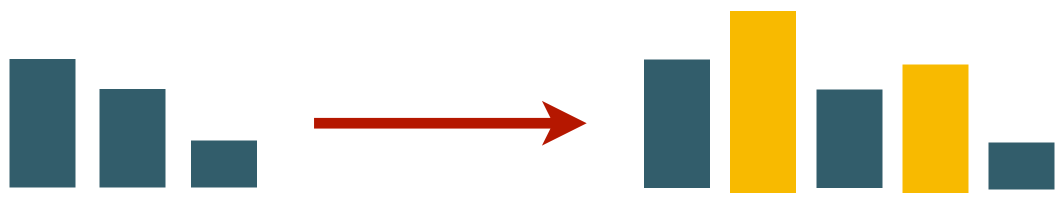

Injecting dummies is more effective, but comes with deployment complexity.

Example: Injecting 5 dummies on average reduces F-Score by 39% with total cost of ~137kB per resource

Summary

- Application-agnostic network layer defenses based on PADDING are inadequate because they fail to hide global features.
- Application-layer defenses are more effective but suffer from deployment challenges:
  - Coordination between parties
  - Developer practices
  - Client experience

Paper: https://arxiv.org/abs/2203.07806

Get in touch: sandra.siby@epfl.ch @sansib

Backup

Application layer defenses: Padding total sizes

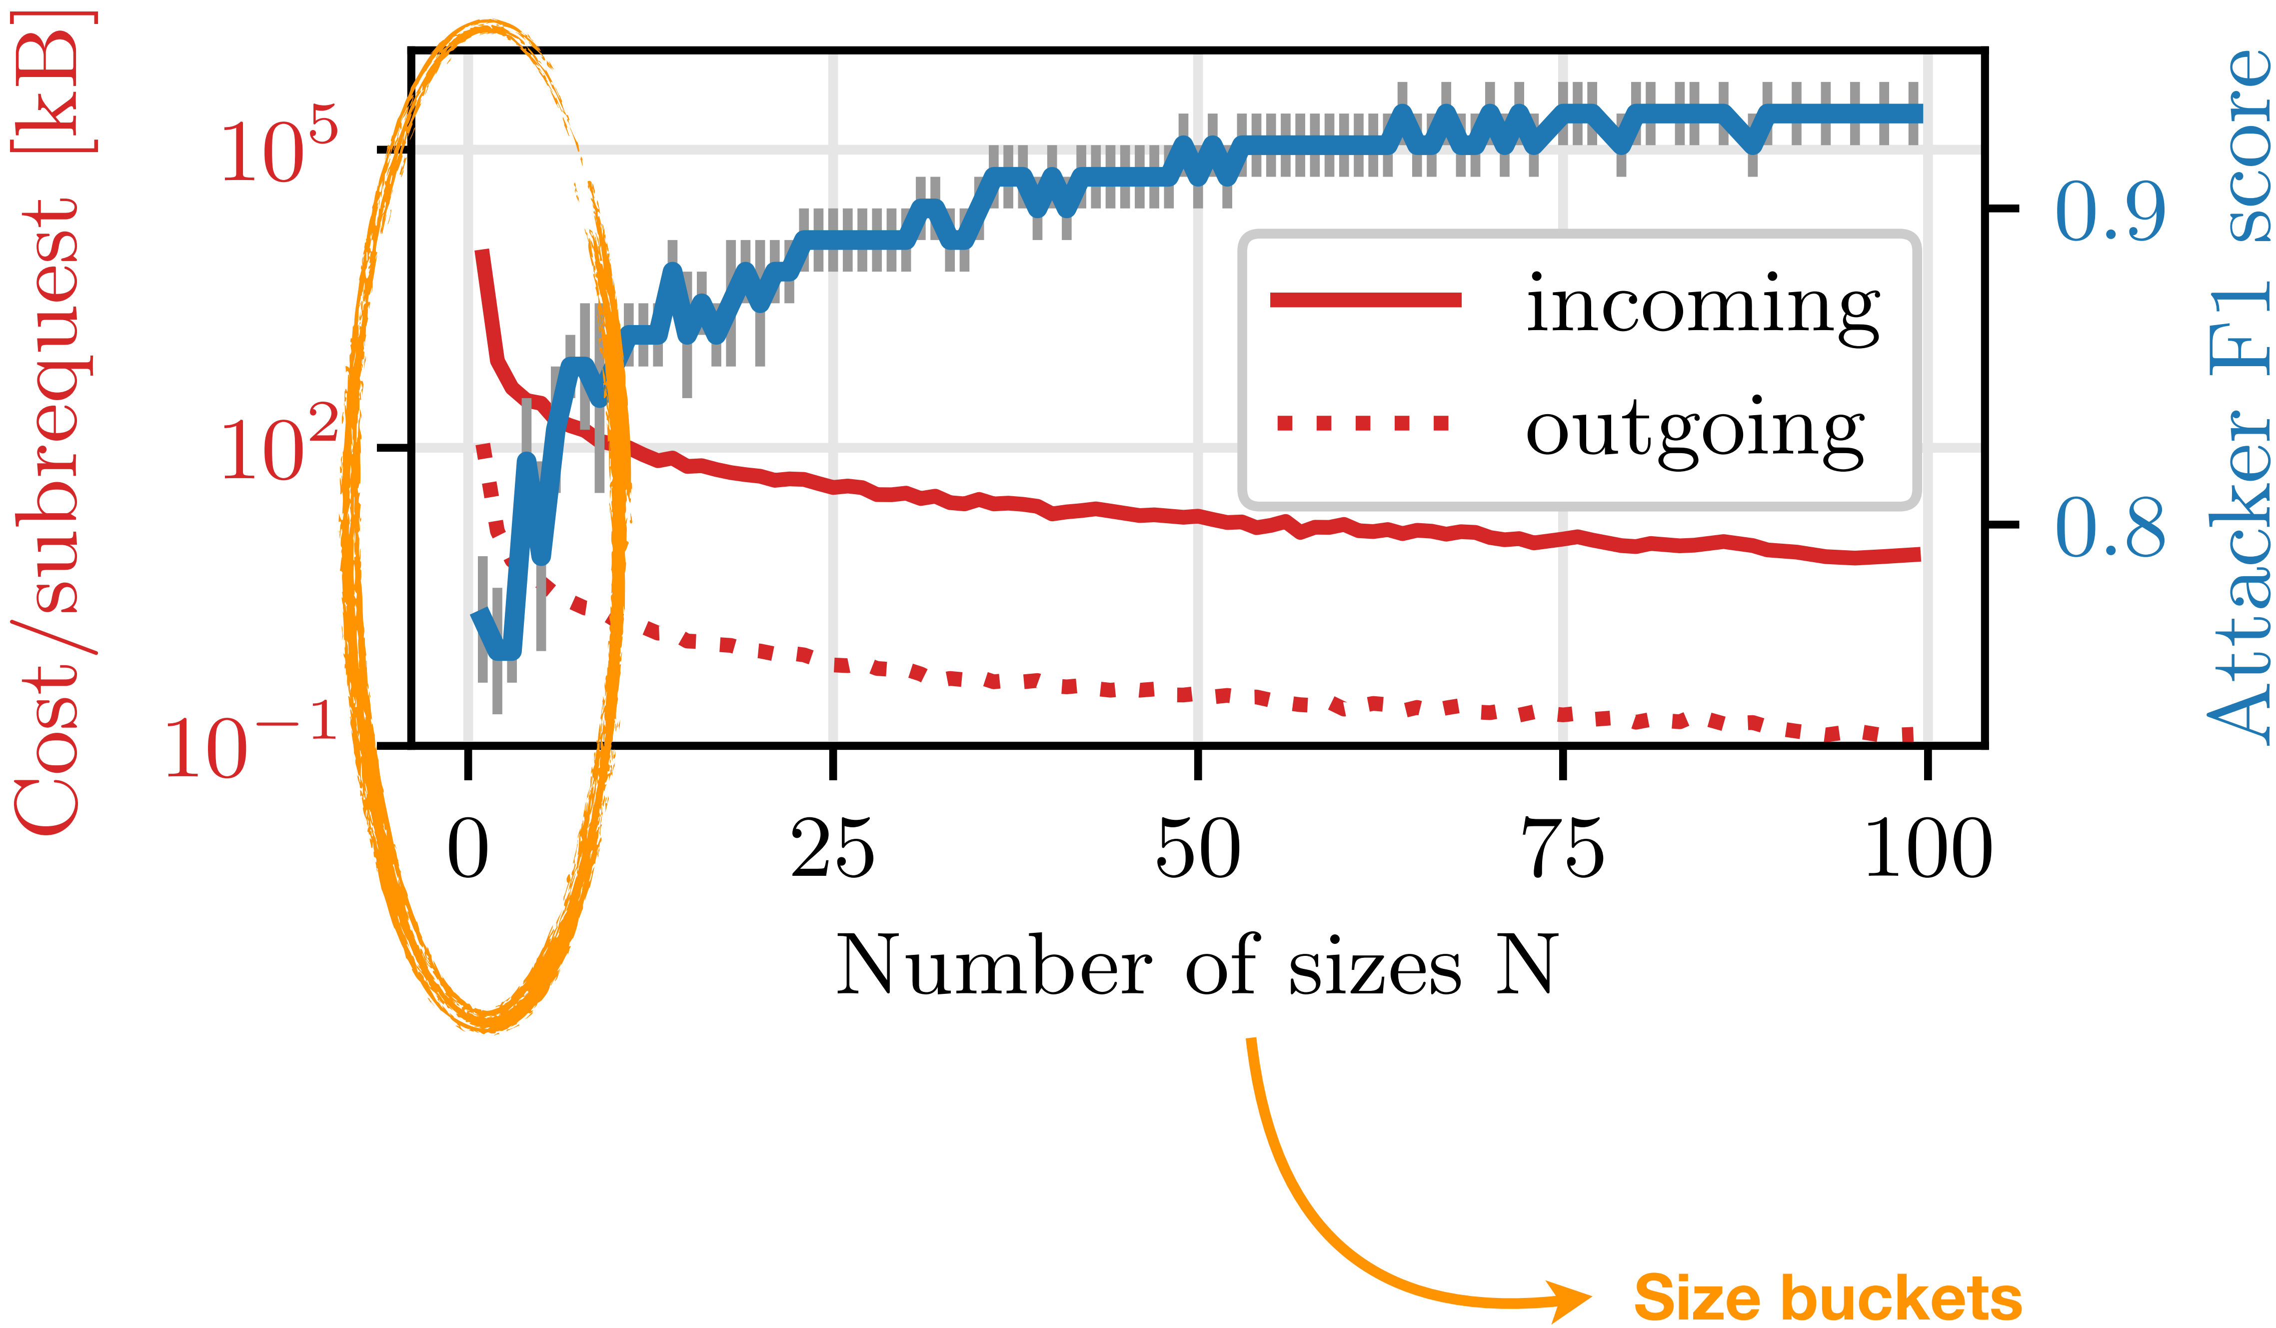

Application layer defenses: Injecting dummies

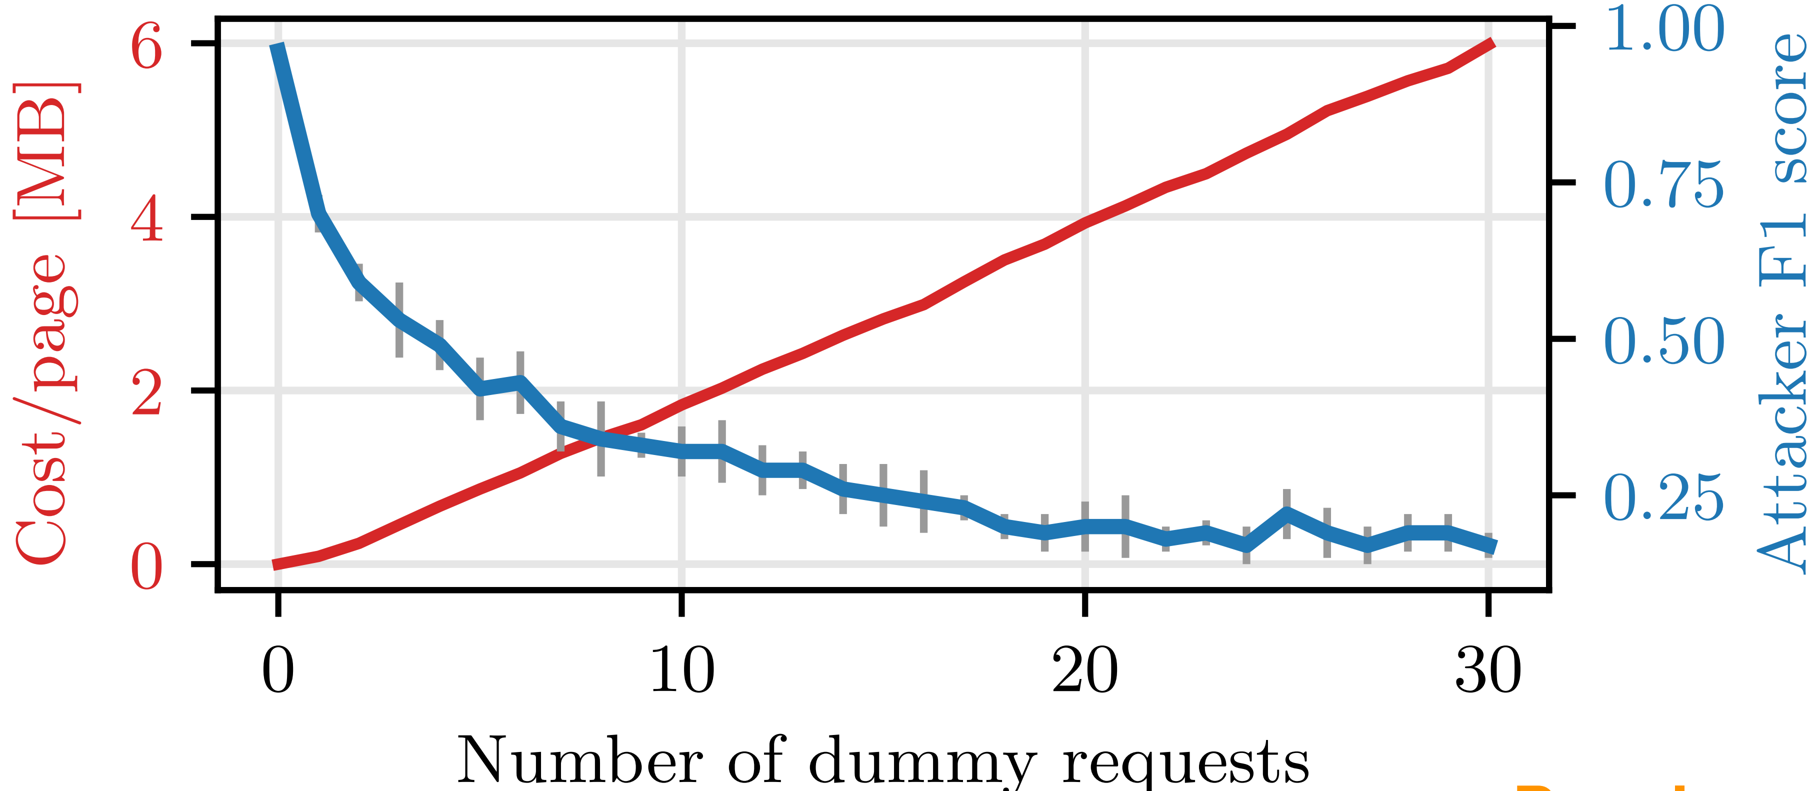

Popular requests