On the Feasibility of Rerouting-based DDoS Defenses

Muoi Tran, Min Suk Kang, Hsu-Chun Hsiao,
Wei-Hsuan Chiang, Shu-Po Tung, Yu-Su Wang

May 2019 | San Francisco, CA

National
Taiwan
University

Transit-link DDoS attack: a powerful type of volumetric DDoS attack

(distributed denial of service)

Traditional: volumetric attack traffic targeting ***end servers***

Non-traditional: volumetric attack traffic targeting ***transit links***

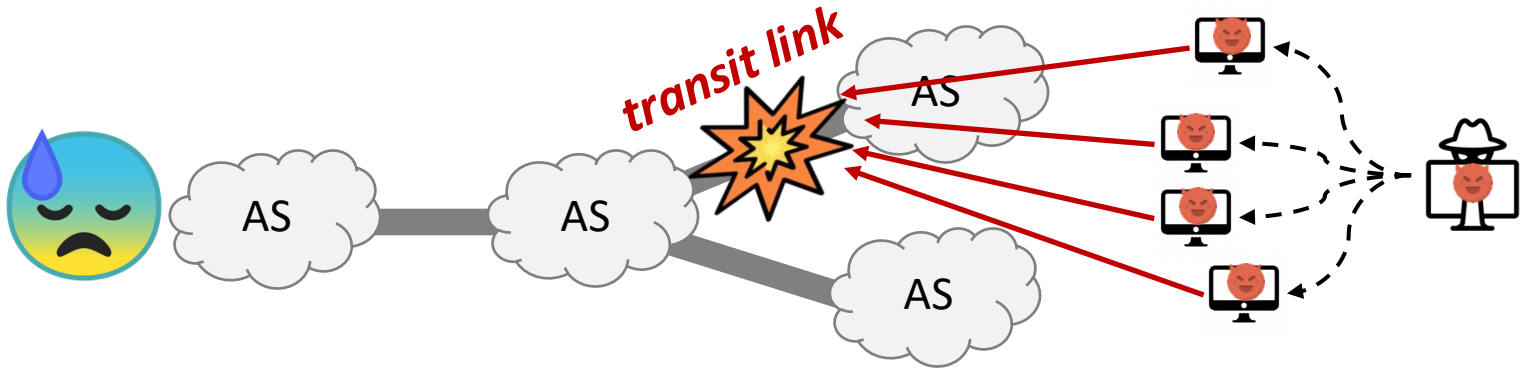

Academic studies:

Coremelt attack (ESORICS '09)

Crossfire attack (S&P '13)

Real incidents:

2013

ars TECHNICA BIZ & IT TECH SCIENCE POLICY CARS GAMING

BIZ & IT —

Spamhaus DDoS grows to Internet-threatening size

More than 300 Gb/s of traffic aimed at the anti-spam site's hosting.

PETER BRIGHT - 3/28/2013, 3:30 AM

2015

Security

ProtonMail DDoS wipeout: Day 6. Yes, we're still under attack

Maybe if you hadn't paid the ransom to the wrong attackers it would be over

By Alexander J Martin 9 Nov 2015 at 16:10 30 SHARE

Handling transit-link DDoS attack is ***challenging***

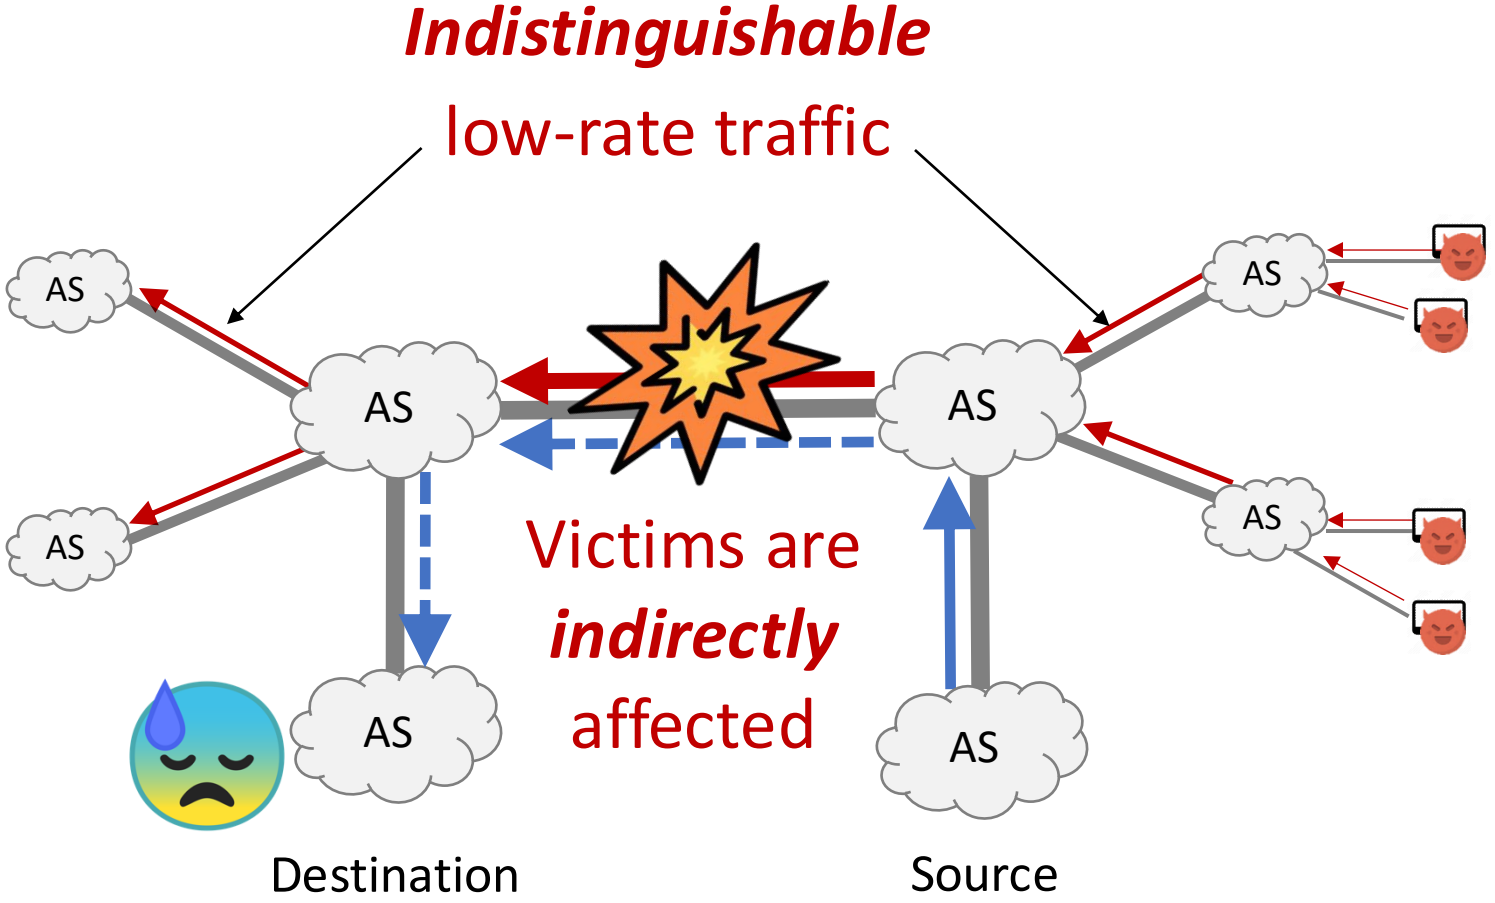

Transit-link DDoS attacks still remain an ***open problem***

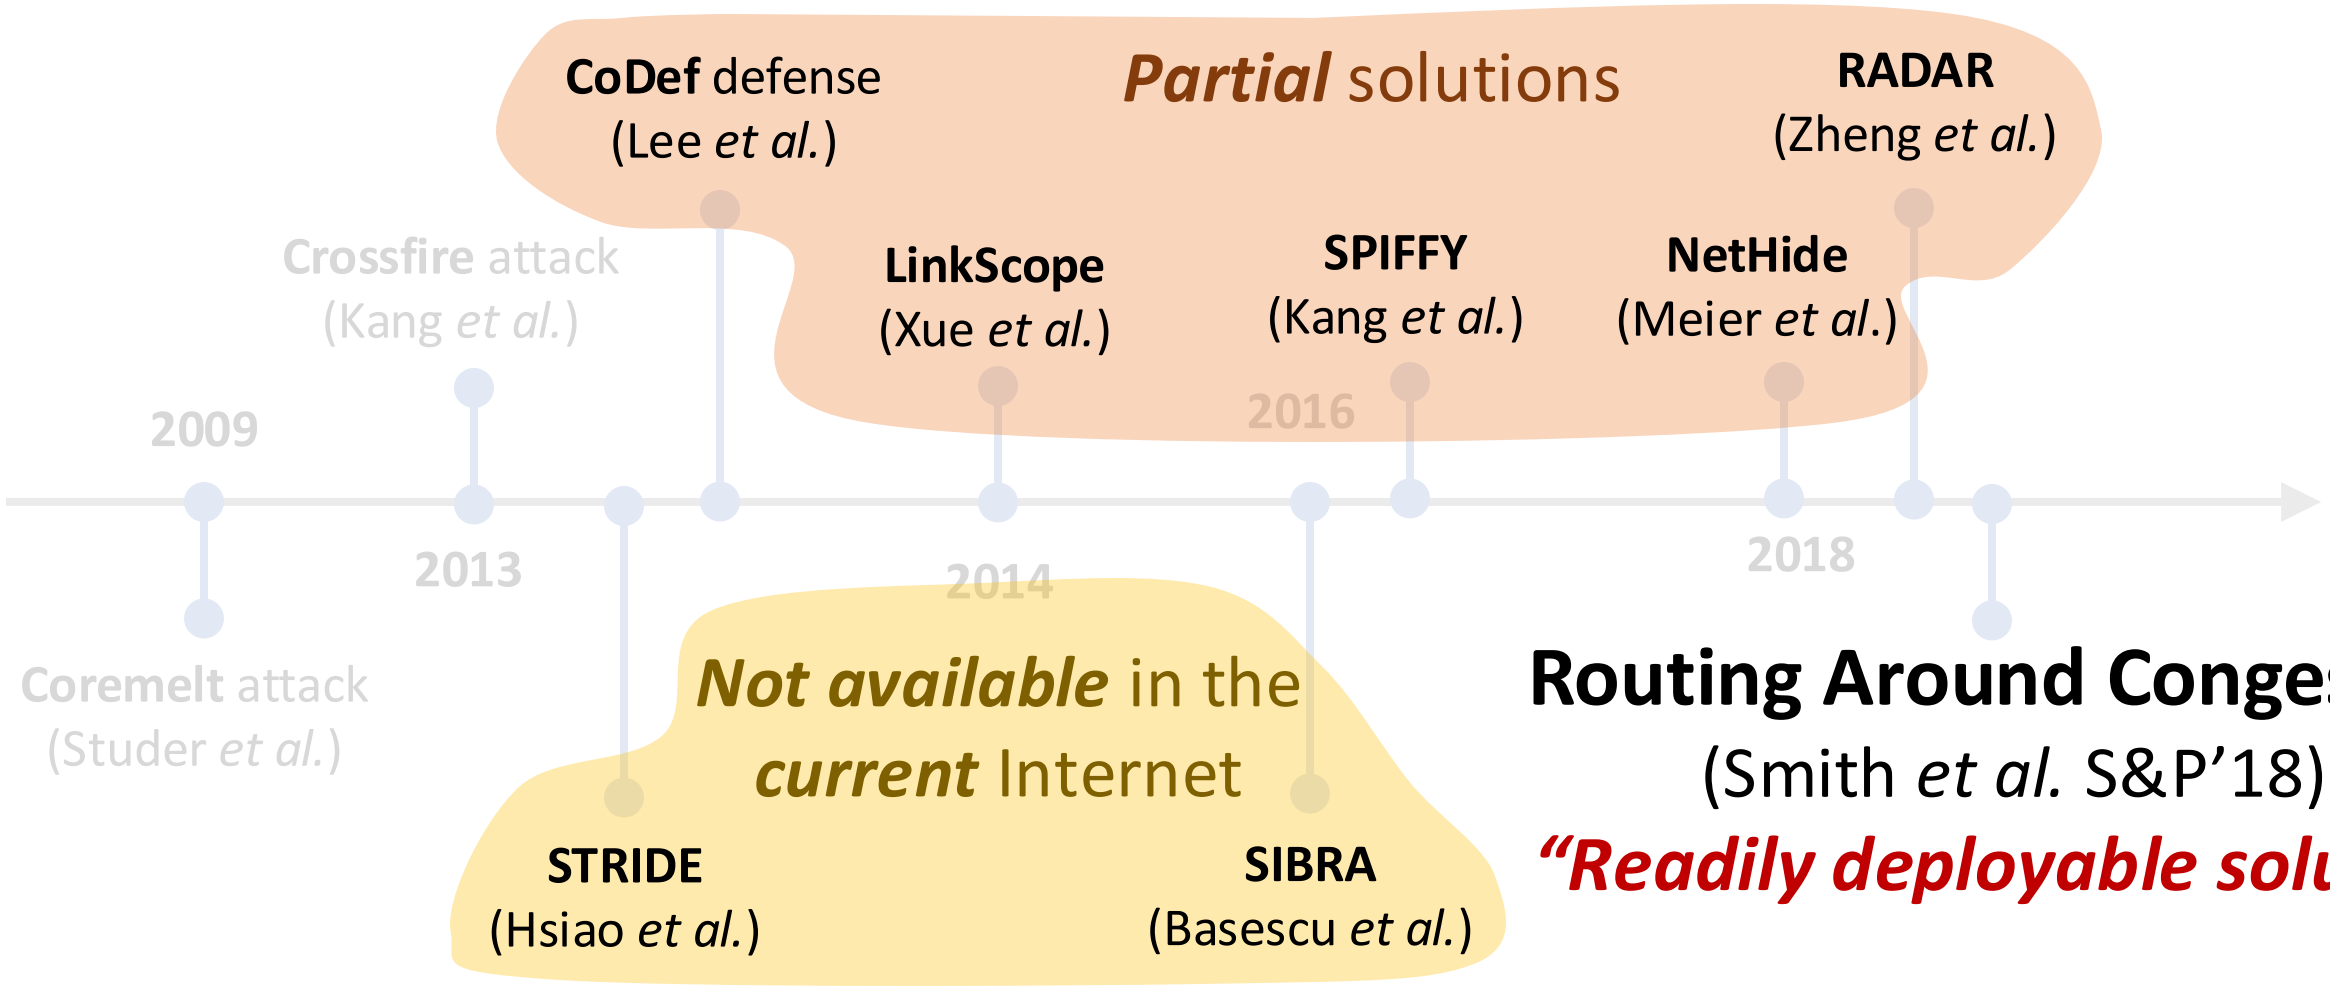

Routing Around Congestion
(Smith *et al.* S&P'18)
"Readily deployable solution"

Background: How ***BGP routing*** works?

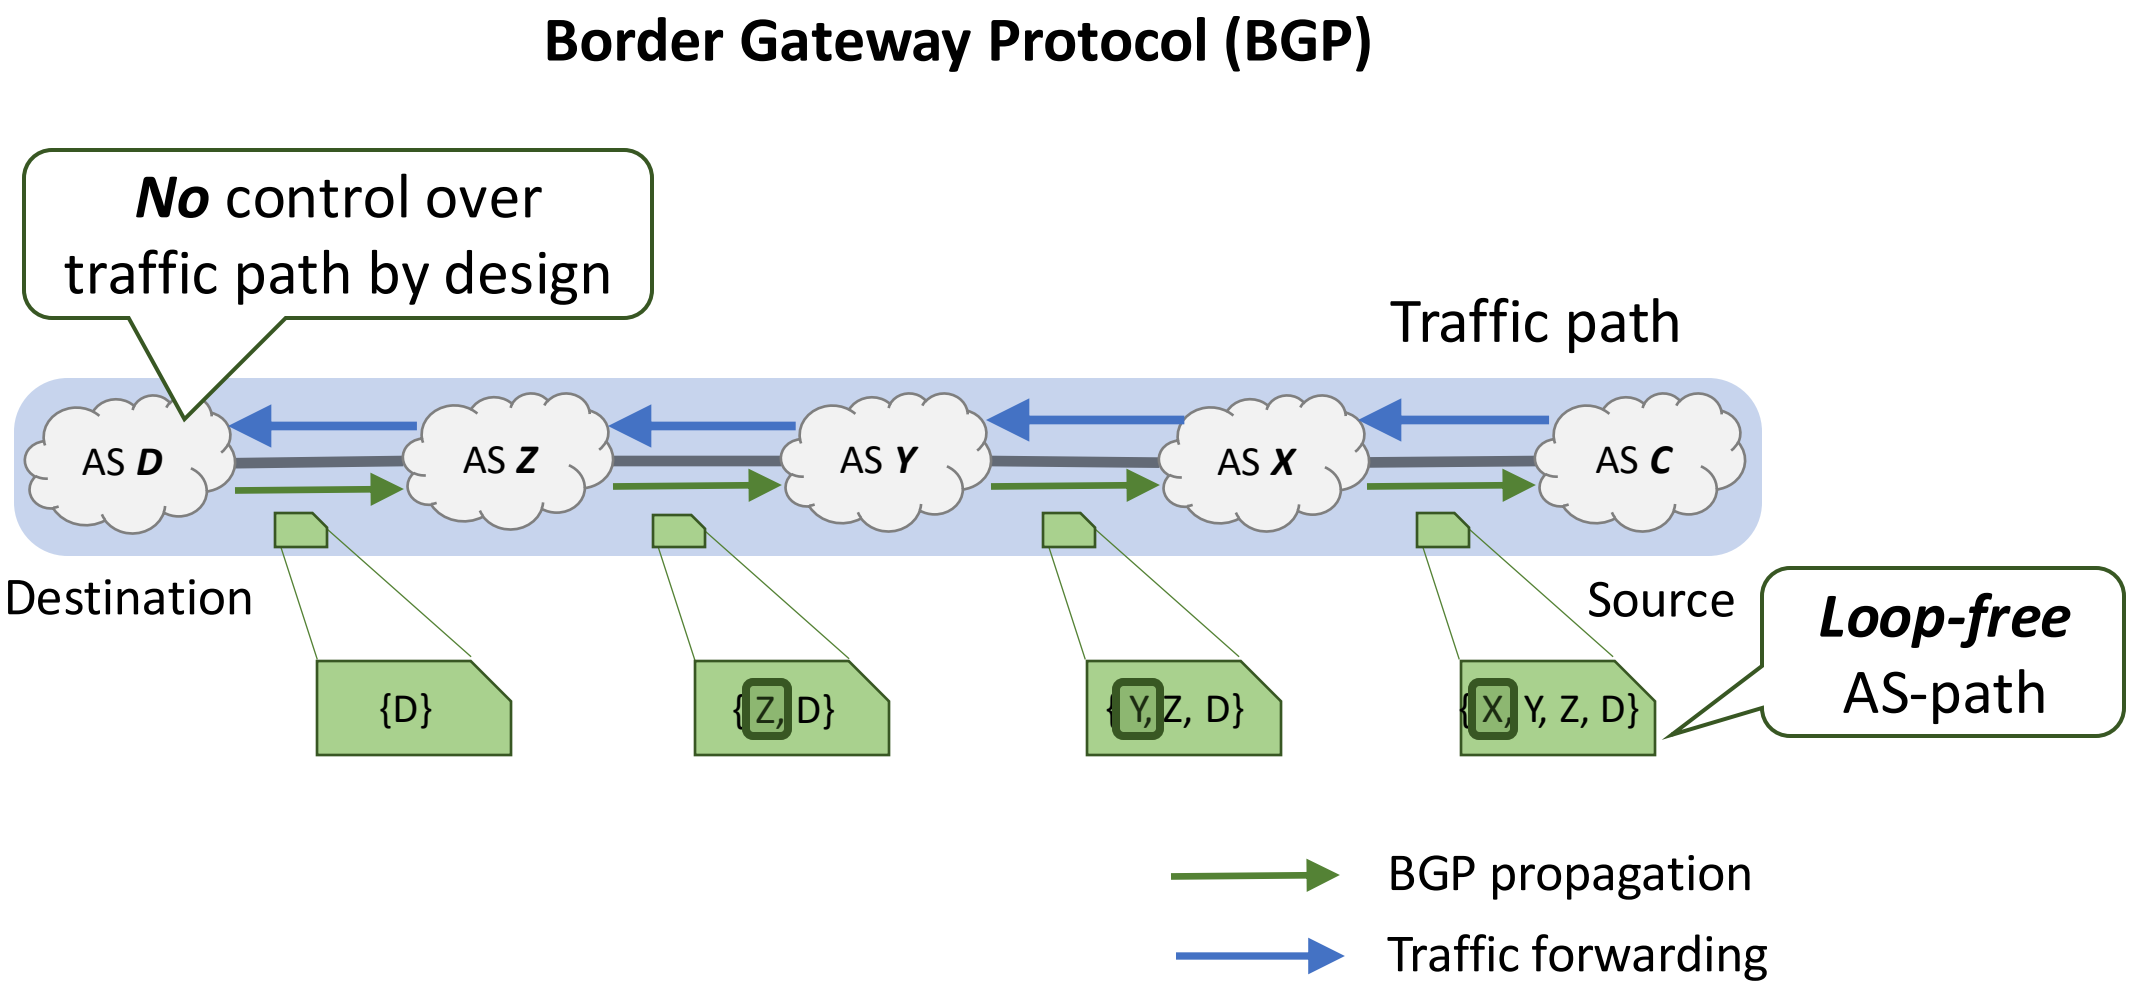

Routing Around Congestion (RAC):
Rerouting using BGP poisoning [Smith *et al.*, S&P '18]

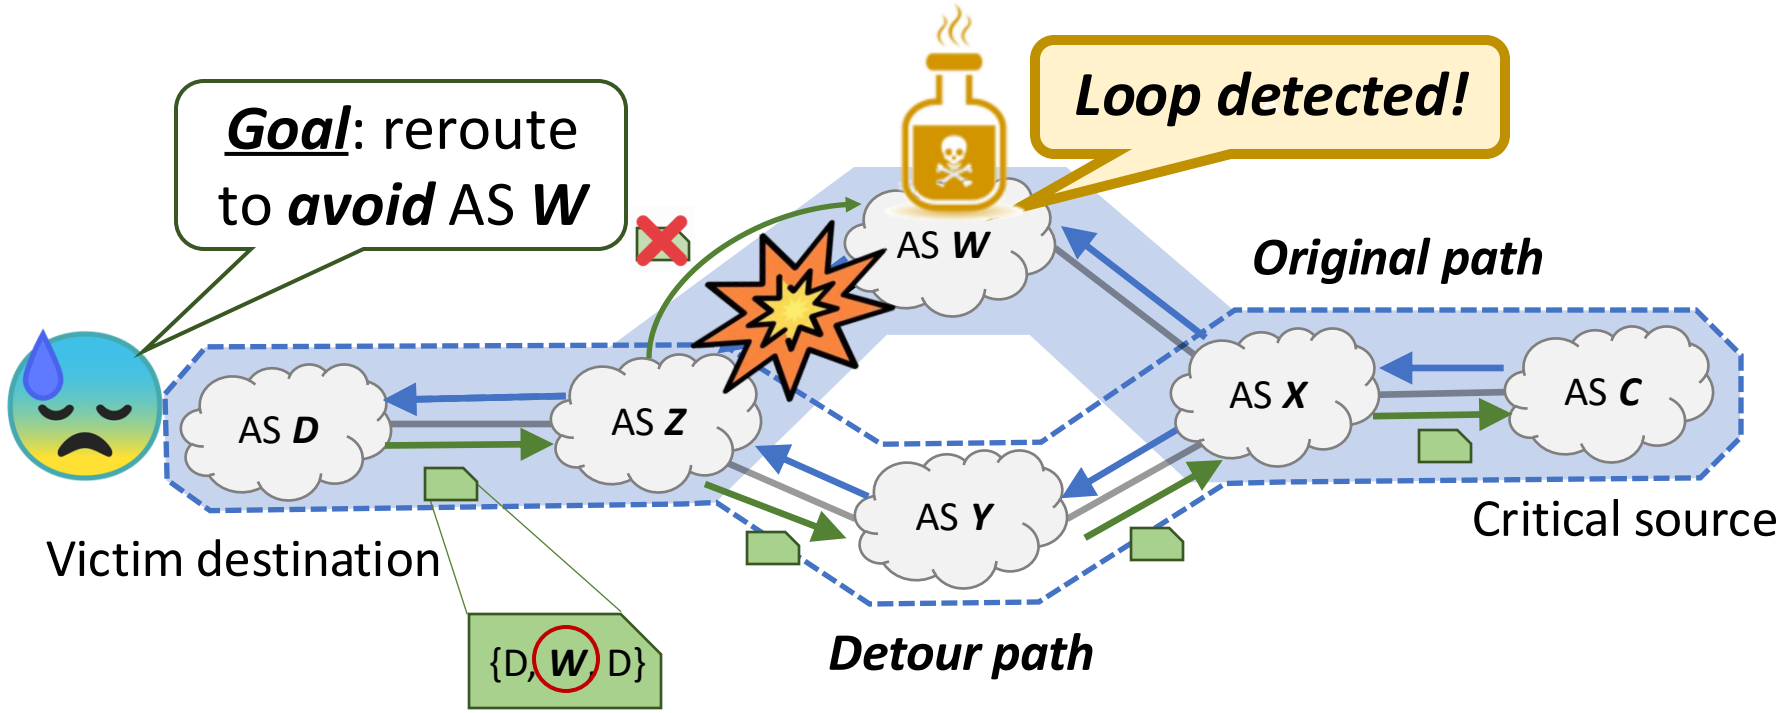

BGP ***poisoning*** message

Routing Around Congestion (RAC):
Rerouting using BGP poisoning [Smith *et al.*, S&P '18]

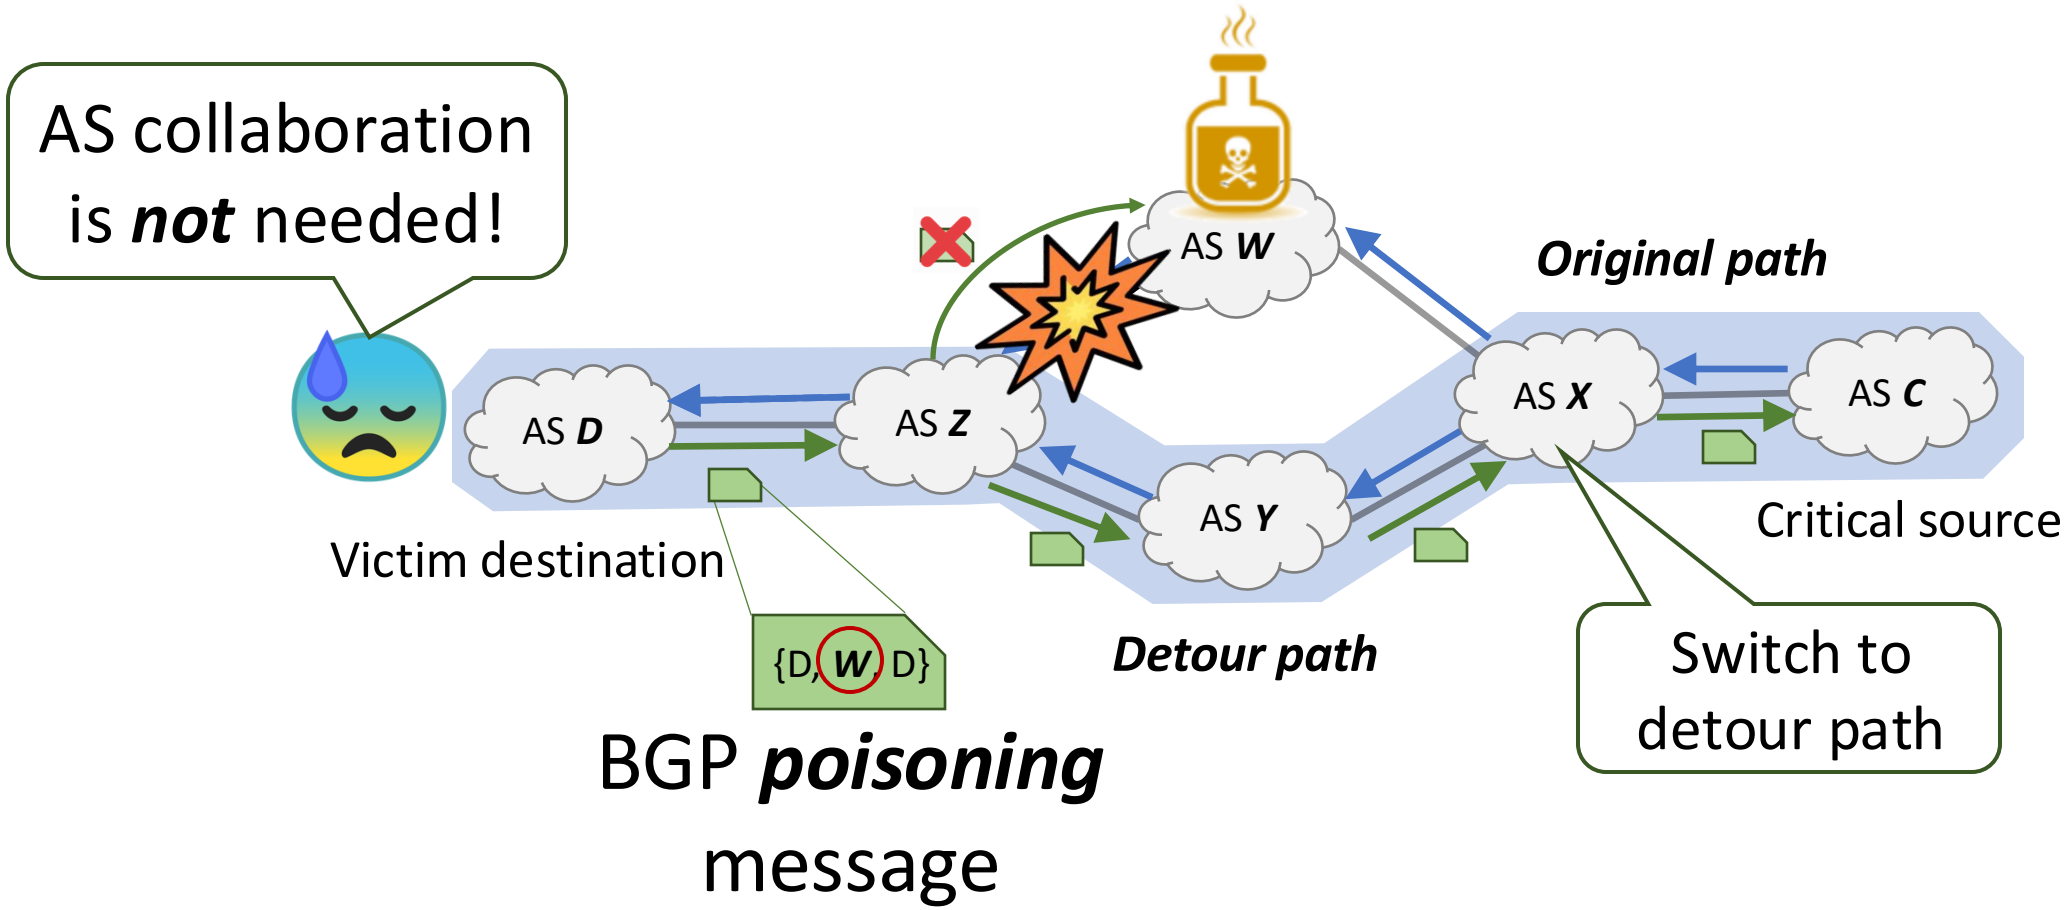

Will ***RAC*** defense still work against ***adaptive attackers***?

Our contributions

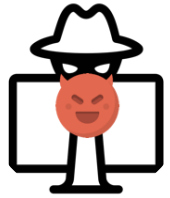

Adaptive detour-learning attack against rerouting solutions

Practical challenge of mitigating adaptive detour-learning attack

Future directions for transit-link DDoS defenses

Adaptive detour-learning attack: Threat model

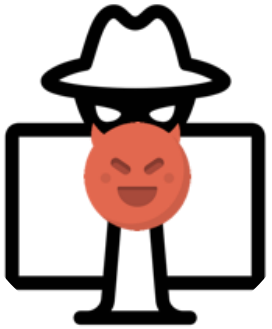

Goals:

(1) To detect rerouting in real-time

(2) To learn new detour path accurately

(3) To congest new detour path (see the paper)

Capabilities:

- Same botnets used in transit-link DDoS attack

Adaptive detour-learning attack:
(1) how to ***detect*** rerouting in ***real-time***

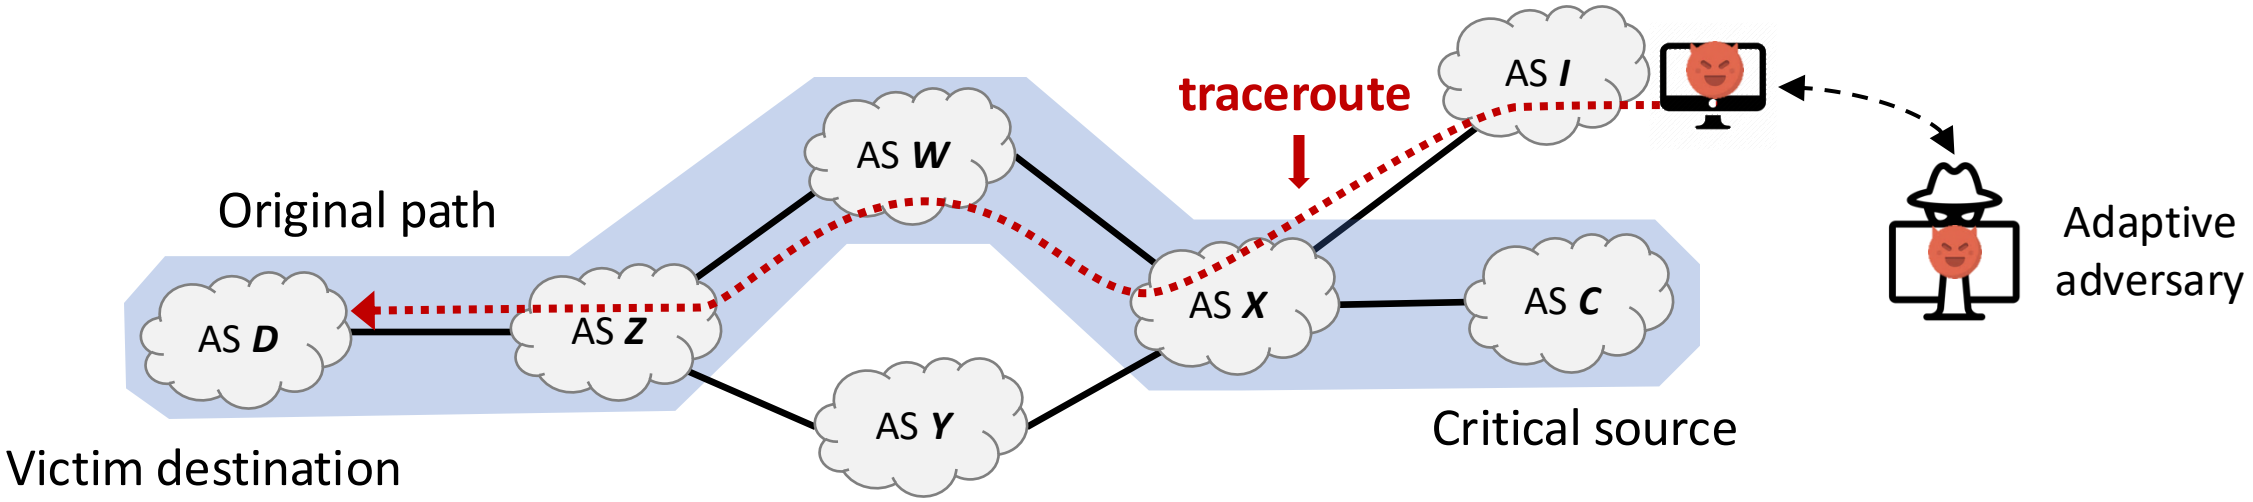

Adaptive detour-learning attack:
(1) how to ***detect*** rerouting in ***real-time***

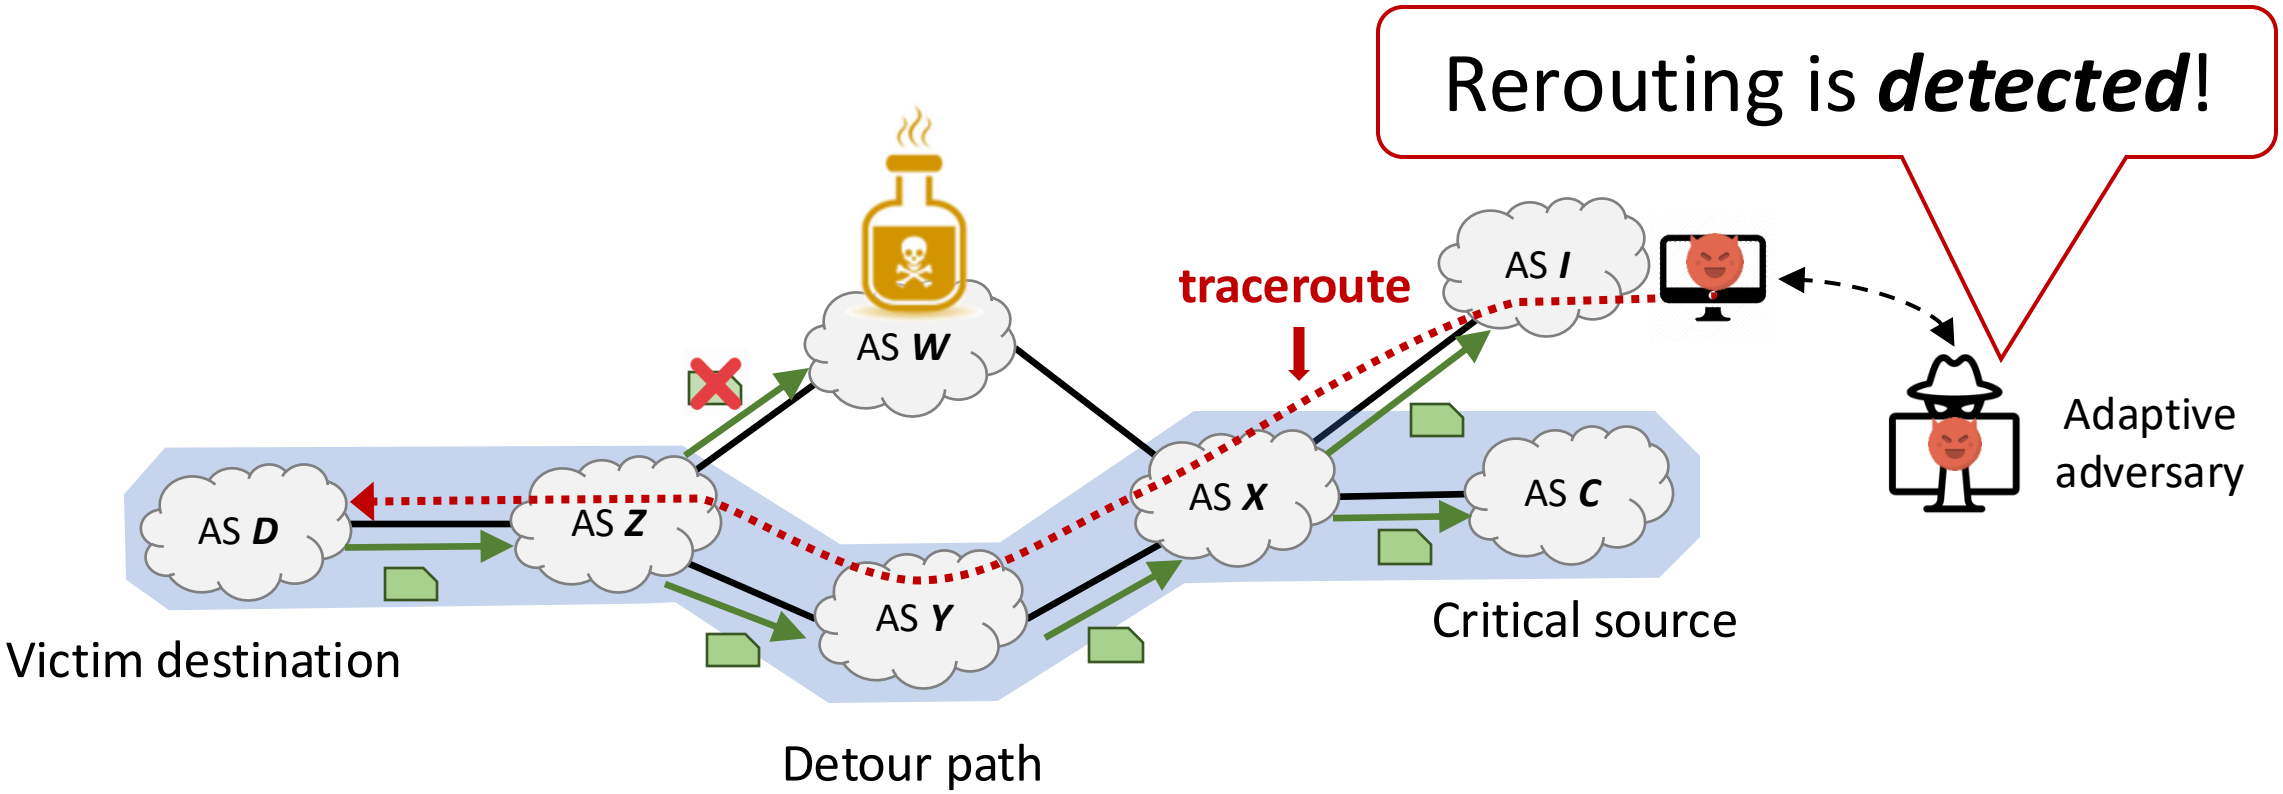

Adaptive detour-learning attack: (2) how to ***learn*** detour path ***accurately***

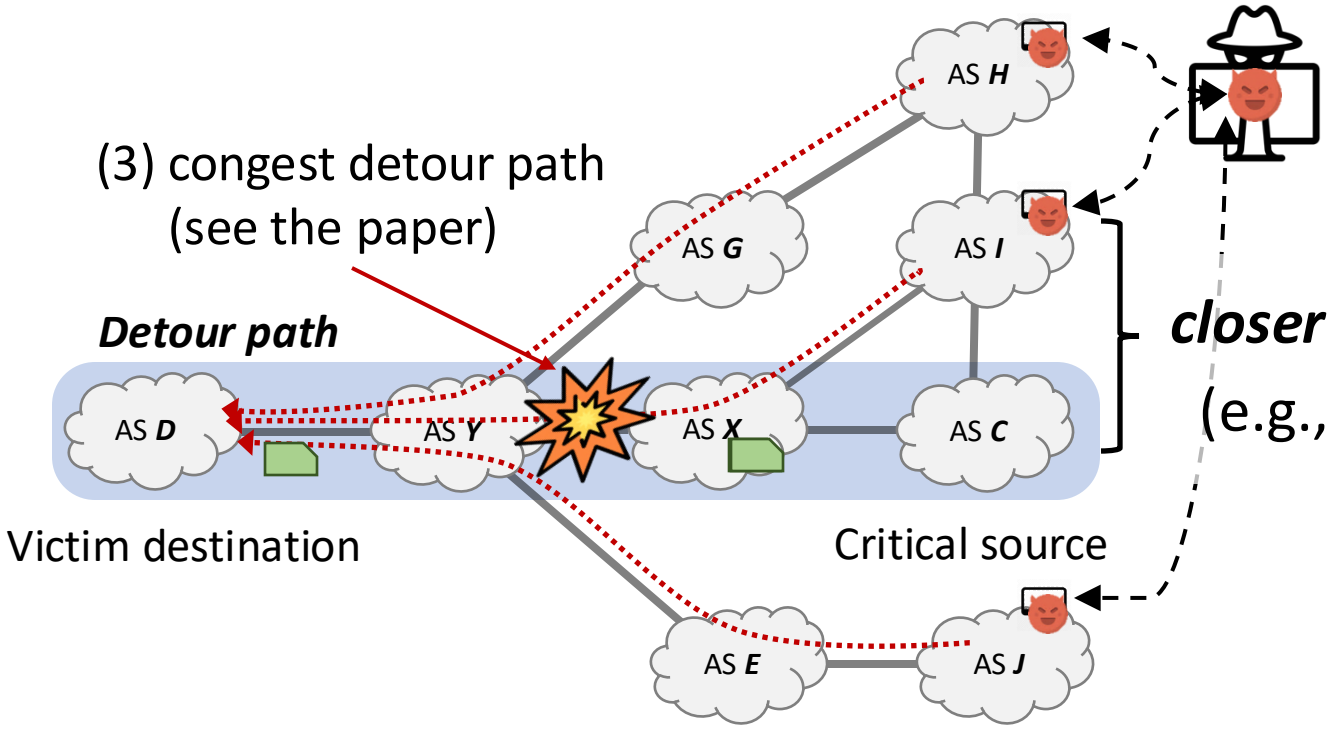

Challenge: Which is more accurate route measurement of ***actual*** detour path?

closer AS (e.g., shorter AS-path)

Solution: Prioritize measurement from bot ***closer*** to traffic source

Adaptive detour-learning attack:
(2) how to ***learn*** detour path ***accurately***

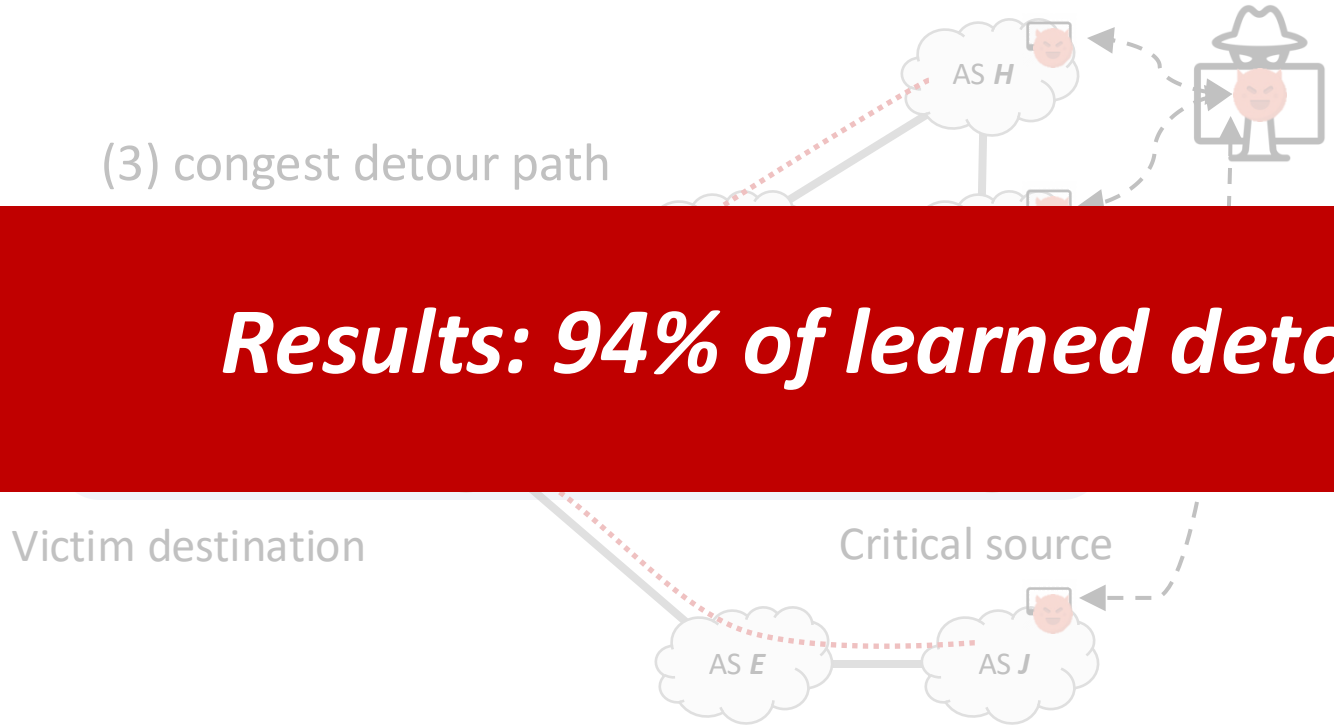

Challenge: Which is more accurate route measurement of *actual* detour path?

Results: 94% of learned detour paths are correct

Solution: Prioritize measurement from bot ***closer*** to traffic source

Our contributions

Adaptive detour-learning attack against rerouting solutions

Practical challenge of mitigating adaptive detour-learning attack

Future directions for transit-link DDoS defenses

How to defend against detour-learning attack?

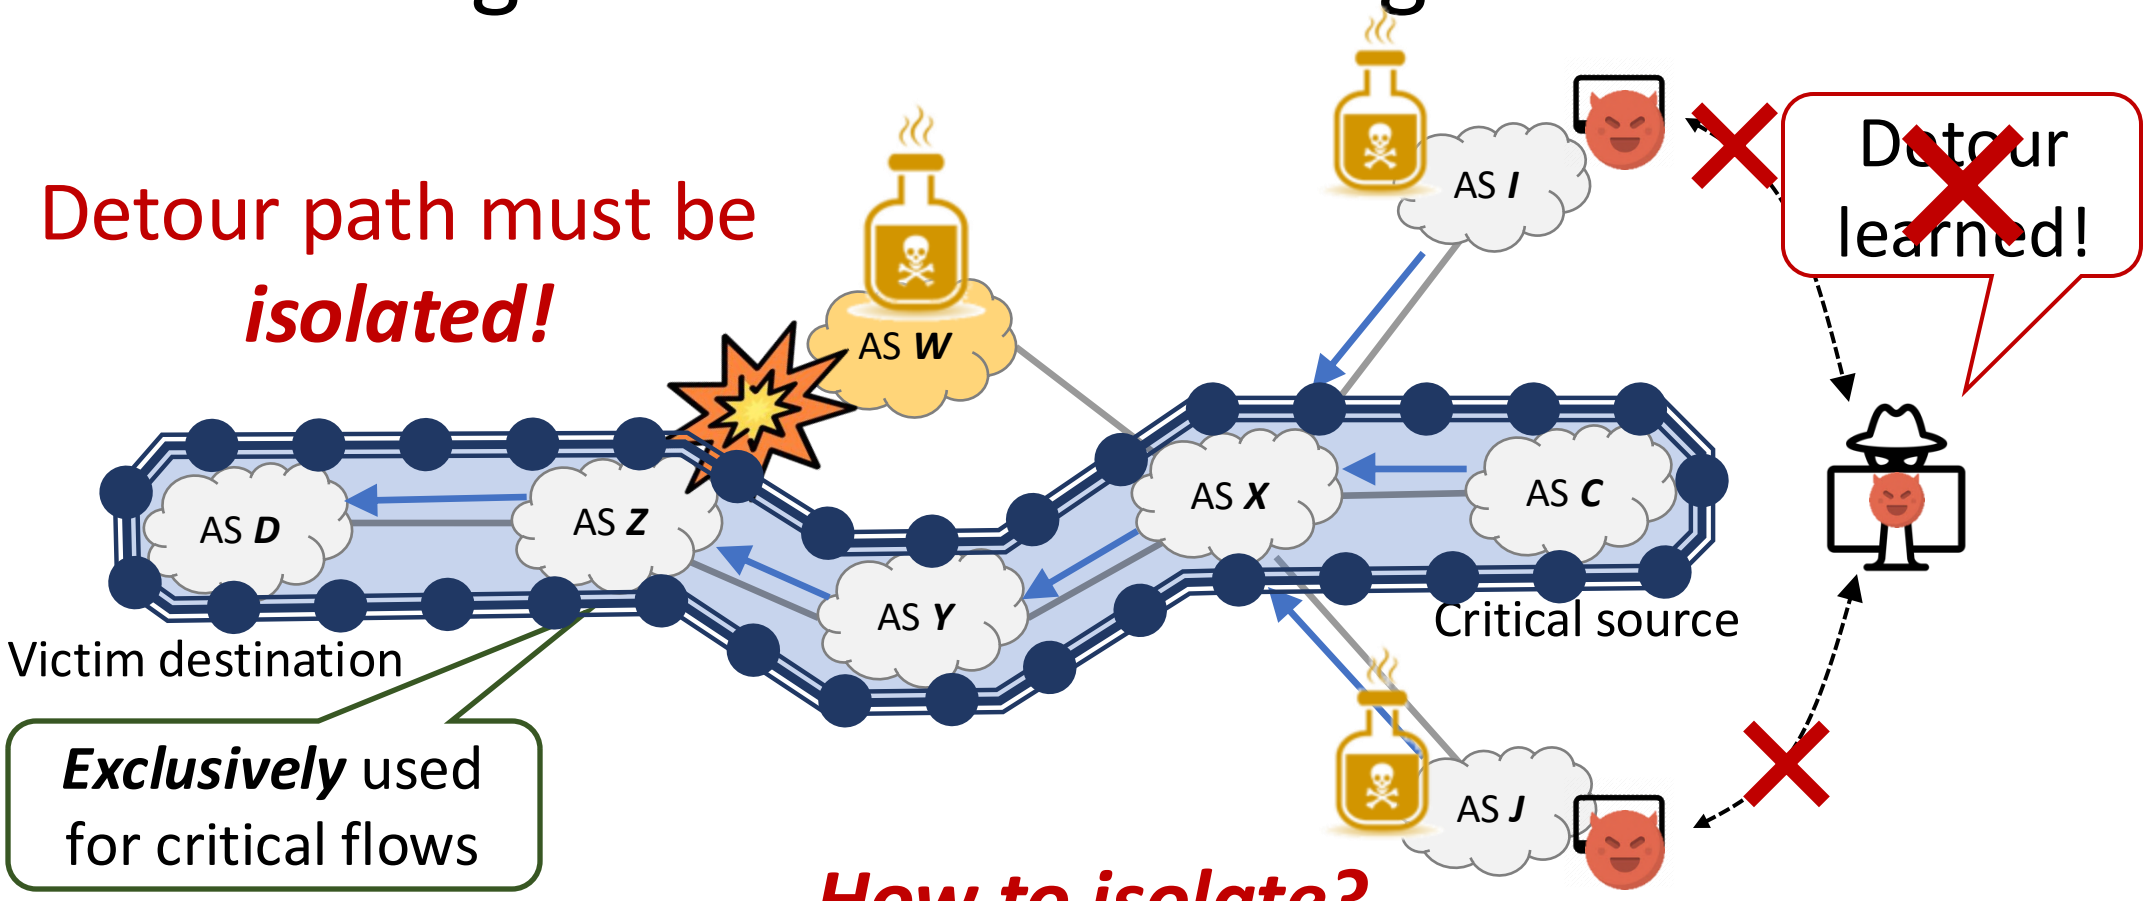

How to isolate?

Poison all peers of ASes on detour path!

Detour path isolation => poisoning ***too many*** ASes

Thousands ASes should be poisoned

But ***why***?

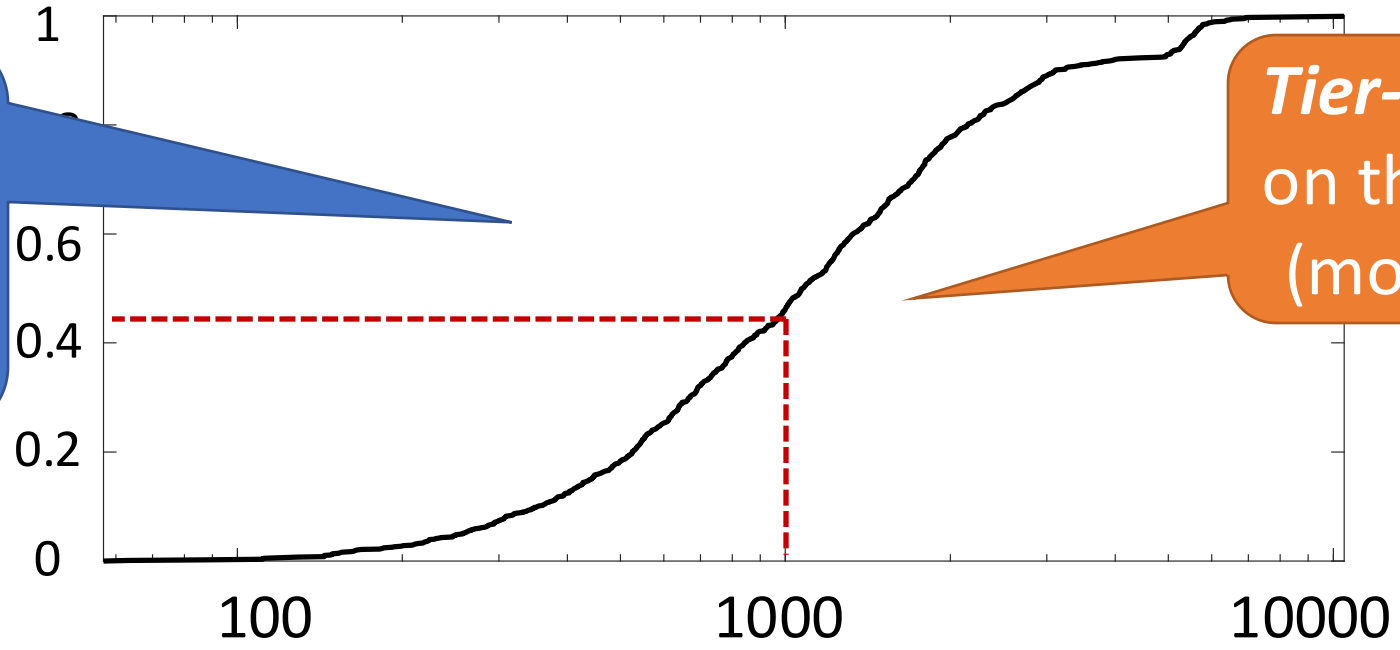

Tier-1 or large ***Tier-2*** on the detour paths (more in the paper)

Can we poison that many ASes?

Specification

Implementation

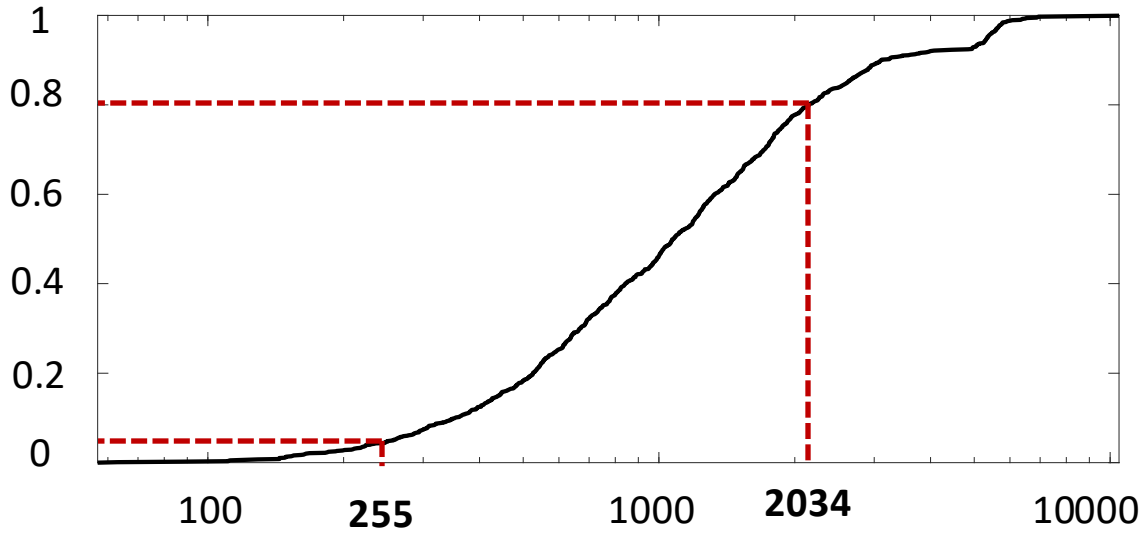

Specification

up to 2034

Implementation

up to 255

Configuration

up to 30-50

Confirmed: ISPs do not support poisoning > 255 ASes

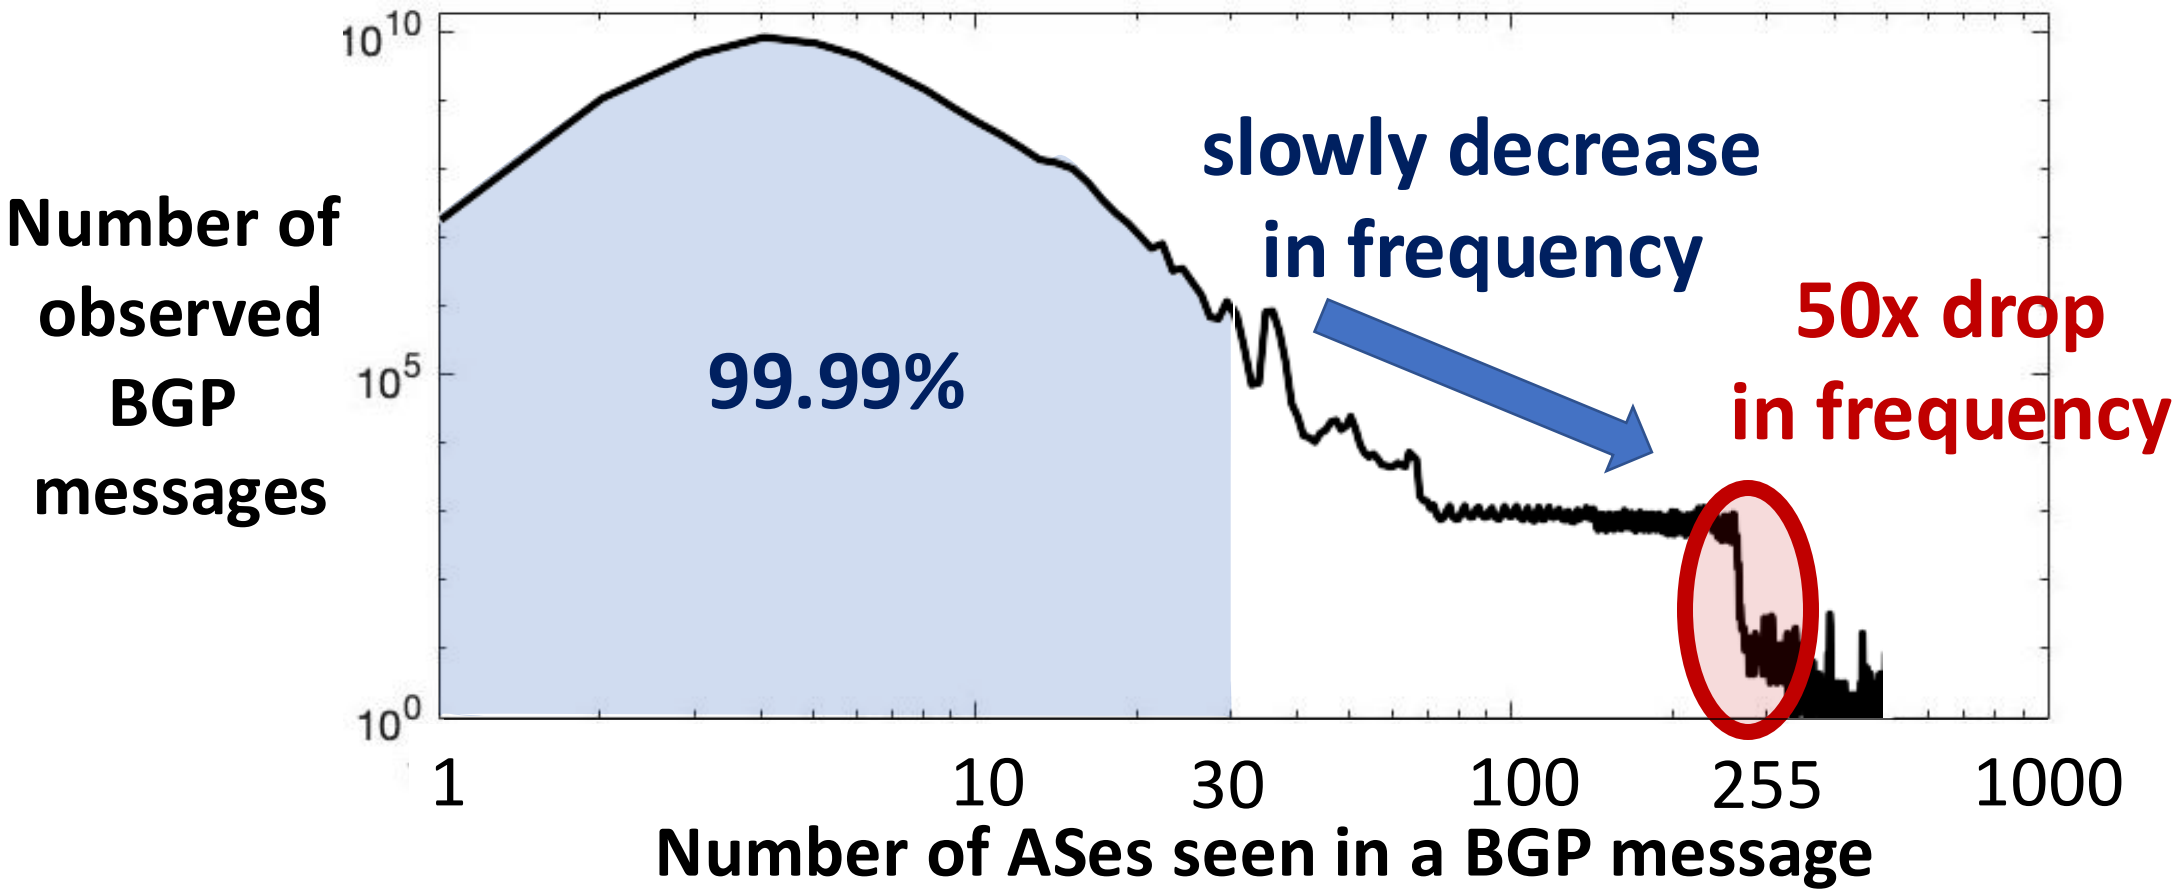

Confirmed: ISPs do not support poisoning > 255 ASes

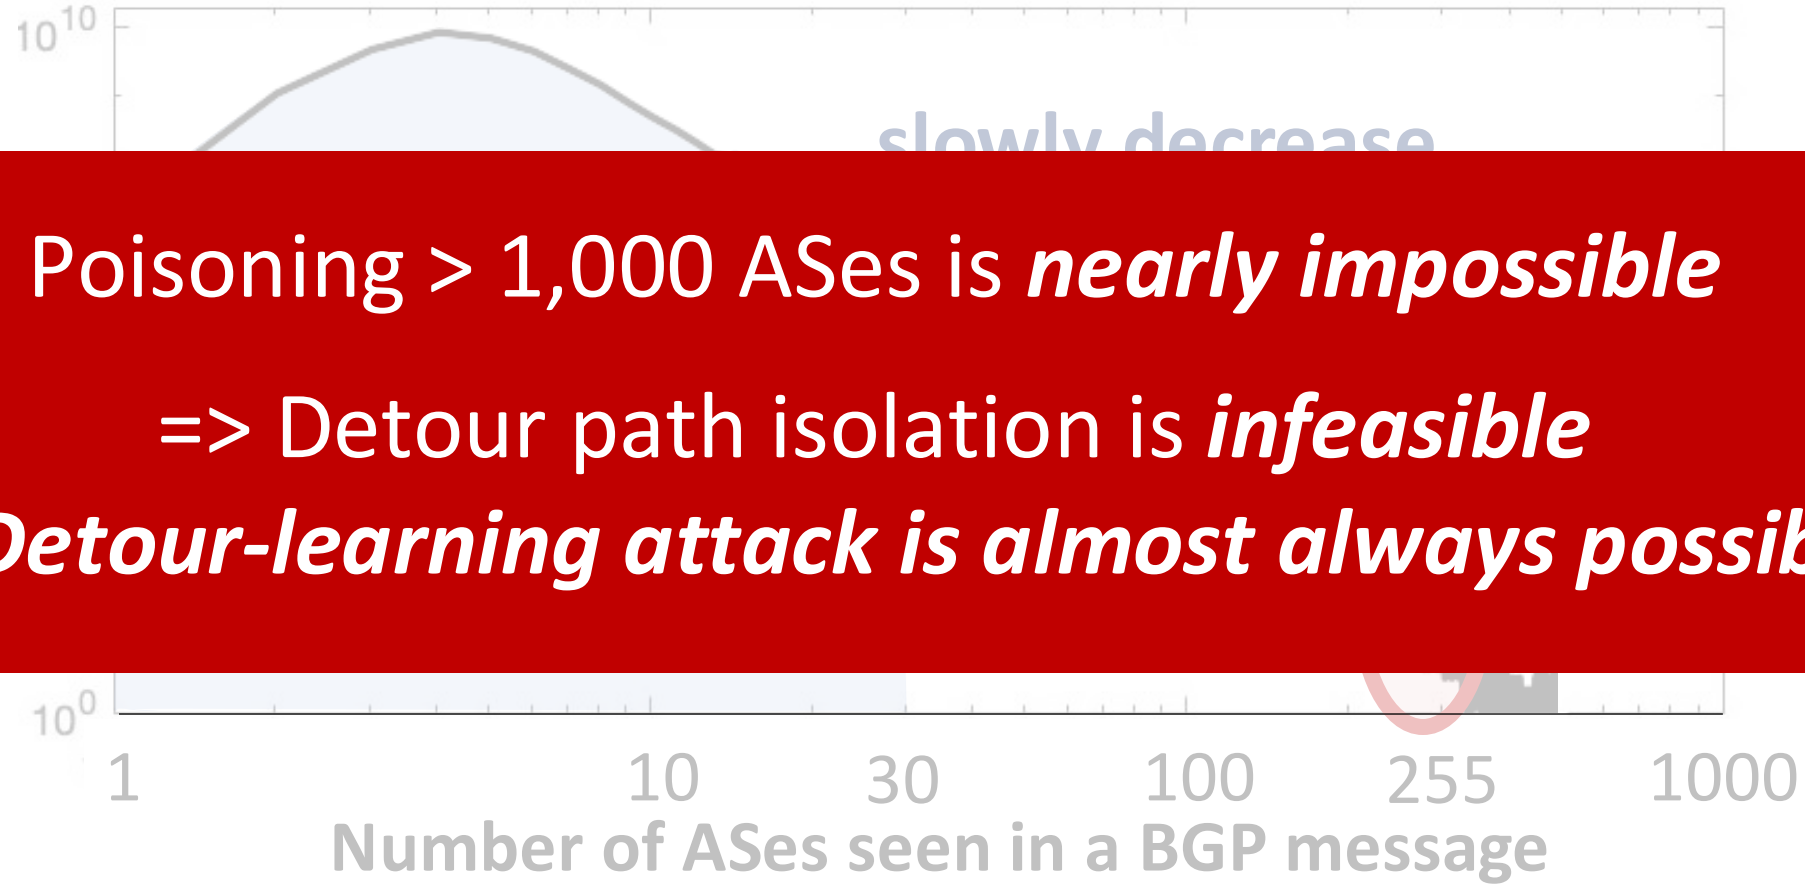

Poisoning > 1,000 ASes is ***nearly impossible***

=> Detour path isolation is ***infeasible***

=> Detour-learning attack is almost always possible

Our contributions

Adaptive detour-learning attack against rerouting solutions

Practical challenge of mitigating adaptive detour-learning attack

Future directions for transit-link DDoS defenses

Desired defense property: destination-controlled routing

Hacking BGP

e.g., Routing Around Congestion

✕ Does not work

?

e.g., ***explicit*** BGP ***rerouting*** for ***critical*** flows under emergency

Clean-slate Internet architecture

e.g., STRIDE, SIBRA

✕ Too costly to deploy

Two Lessons Learned

Lesson 1

Hacking the current Internet routing is a ***flawed*** idea!

✓Adaptive attacks are possible

✓Mitigation is hard

✓ Adaptive defense is slower than adaptive attacker (more in the paper)

Lesson 2

Analysis of protocol ***specifications alone*** is ***insufficient***!

Specification ≠ **Implementation** ≠ **Configuration**

Conclusion

- ***Detour-learning attacks*** are ***effective*** and ***hard to mitigate***
  - ✓Transit-link DDoS attacks still remain an open problem

- Suggestion on research direction
  - ✓Balance ***destination-controlled routing*** and ***deployability***

- 2 lessons learned:
  - ✓Hacking BGP for rerouting is a flawed idea
  - ✓Analysis with specification only can be dangerous

Question?

Muoi Tran

muoitran@comp.nus.edu.sg