Internet control Plane Security

Yongdae Kim

KAIST

Two Planes

- Data Plane: Actual data delivery
- control Plane
  - To support data delivery (efficiently, reliably, and etc.)
  - Routing information exchange
  - In some sense, every protocol except data delivery is considered to be control plane protocols
- Example network
  - Peer-to-peer network, cellular network, Internet, ...

Historical List of Botnet

creation	Name	# of Bots	Spam	control
2004	Bagle	230K	5.7 B/day	centralized
2007	Storm	> 1,000K	3 B/day	P2P
2008	Mariposa	12,000K	?	centralized
2008	Waledac	80K	?	centralized
2008	conficker	>10,000K	10 B/day	ctrlzd/P2P
2009?	Mega-D	4,500K	10 B/day	centralized
2009?	Zeus	>3,600K	?	
2009	BredoLab	30,000K	3.6 B/day	centralized
2010	TDL4	4,500K	?	P2P

Misconfigurations and Redirection

- 1997: AS7007
  - claimed shortest path to the whole Internet
  - causing Internet Black hole
- 2004: TTNet (AS9121)
  - claimed shortest path to the whole Internet
  - Lasted for several hours
- 2006: AS27056
  - "stole" several important prefixes on the Internet
  - From Martha Stewart Living to The New York Daily News
- 2008: Pakistan Youtube
  - decided to block Youtube
  - one ISP advertised a small part of YouTube's (AS 36561) network
- 2010: China
  - 15% of whole Internet traffic was routed through China for 18 minutes
  - including .mil and .gov domain
- 2011: China
  - All traffic from US iPhone to Facebook
  - routed through China and Korea

300Gbps DDoS

- 300 Gbps DDoS against Spamhous from Stophous
- Mitigation by cloudFlare using anycast
- Stophous turn targets to IX (Internet Exchange)
- Korea - world IX Bandwidth
  - KT: 560 Gbps, SKB: 235 Gbps, LGU+: 145 Gbps, SKT: 100 Gbps
  - Total: 1 Tbps

How to crash (or Save) the Internet?

Max Schuchard, Eugene Vasserman, Abedelaziz Mohaisen, Denis Foo Kune, Nicholas Hopper, Yongdae Kim

Losing control of the Internet

- Using the Data Plane to Attack the control Plane -

Network and Distributed System Security (NDSS) 2011

Shutting Down the Internet

- Fast propagating worm
  - codeRed, Slammer worm
- Router misconfiguration
  - AS7007
- 2011
  - Egypt, Libya: Internet Kill Switch
  - US government discussing Internet Kill Switch Bill in emergency situation

Other Internet control Plane News

- April 2008: Whole youtube traffic directed to Pakistan
- April 2010: 15% of whole Internet traffic was routed through China for 18 minutes (including .mil and .gov domain)
- March 2011: All traffic from US iPhone to Facebook was routed through China and Korea

Losing control

- Attack on the Internet's control plane
- Overwhelm routers with BGP updates
- Launched using only a botnet
- Defenses are non trivial
- Different from DDoS on web servers

Attack Model

- No router compromise or misconfiguration
  - BGPSEc or similar technologies
- Our attack model: unprivileged adversary
  - can generate only data plane events
  - does not control any BGP speakers
  - botnet of a reasonable size
    - »50, 100, 250, 500k nodes

can we shut down

the Internet only using

data plane events?

How much control plane events

can be generated by data plane events caused by

coordinated set of compromised computers?

AS, BGP and the Internet

- AS (Autonomous System)
  - core AS: High degree of connectivity
  - Fringe AS: very low degrees of connectivity, sitting at the outskirts of the Internet
  - Transit AS: core ASes, which agree to forward traffic to and from other Ases
- BGP (Border Gateway Protocol)
  - the de facto standard routing protocol spoken by routers connecting different ASes.
  - BGP is a path vector routing algorithm, allowing routers to maintain a table of AS paths to every destination.
  - uses policies to preferentially use certain AS paths in favor.

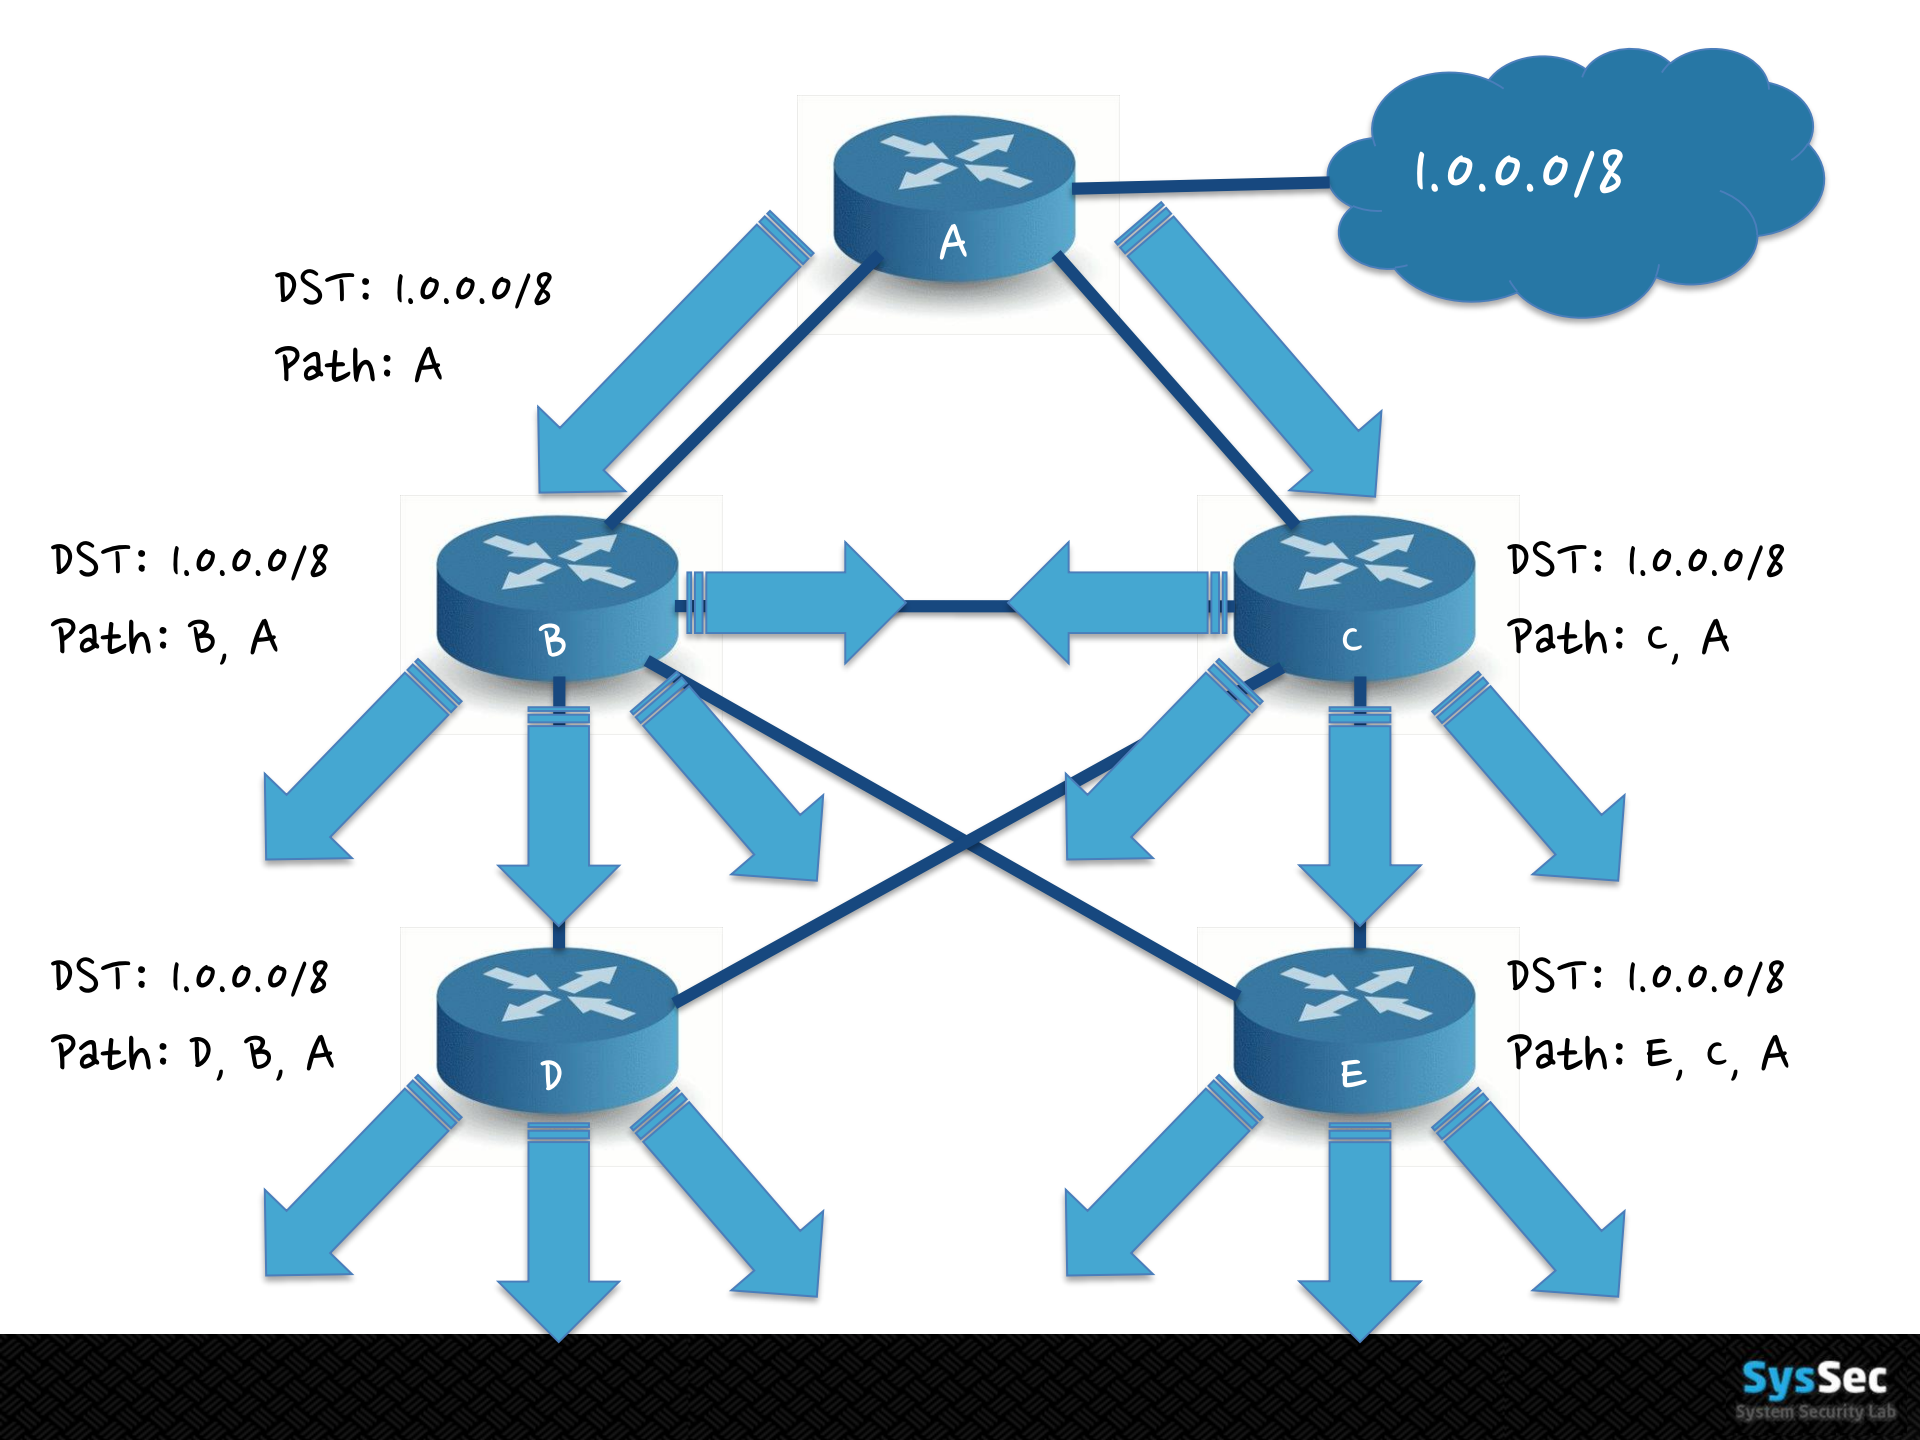
1.0.0.0/8
A
DST: 1.0.0.0/8
Path: A
B
DST: 1.0.0.0/8
Path: B, A
C
DST: 1.0.0.0/8
Path: C, A
D
DST: 1.0.0.0/8
Path: D, B, A
E
DST: 1.0.0.0/8
Path: E, C, A

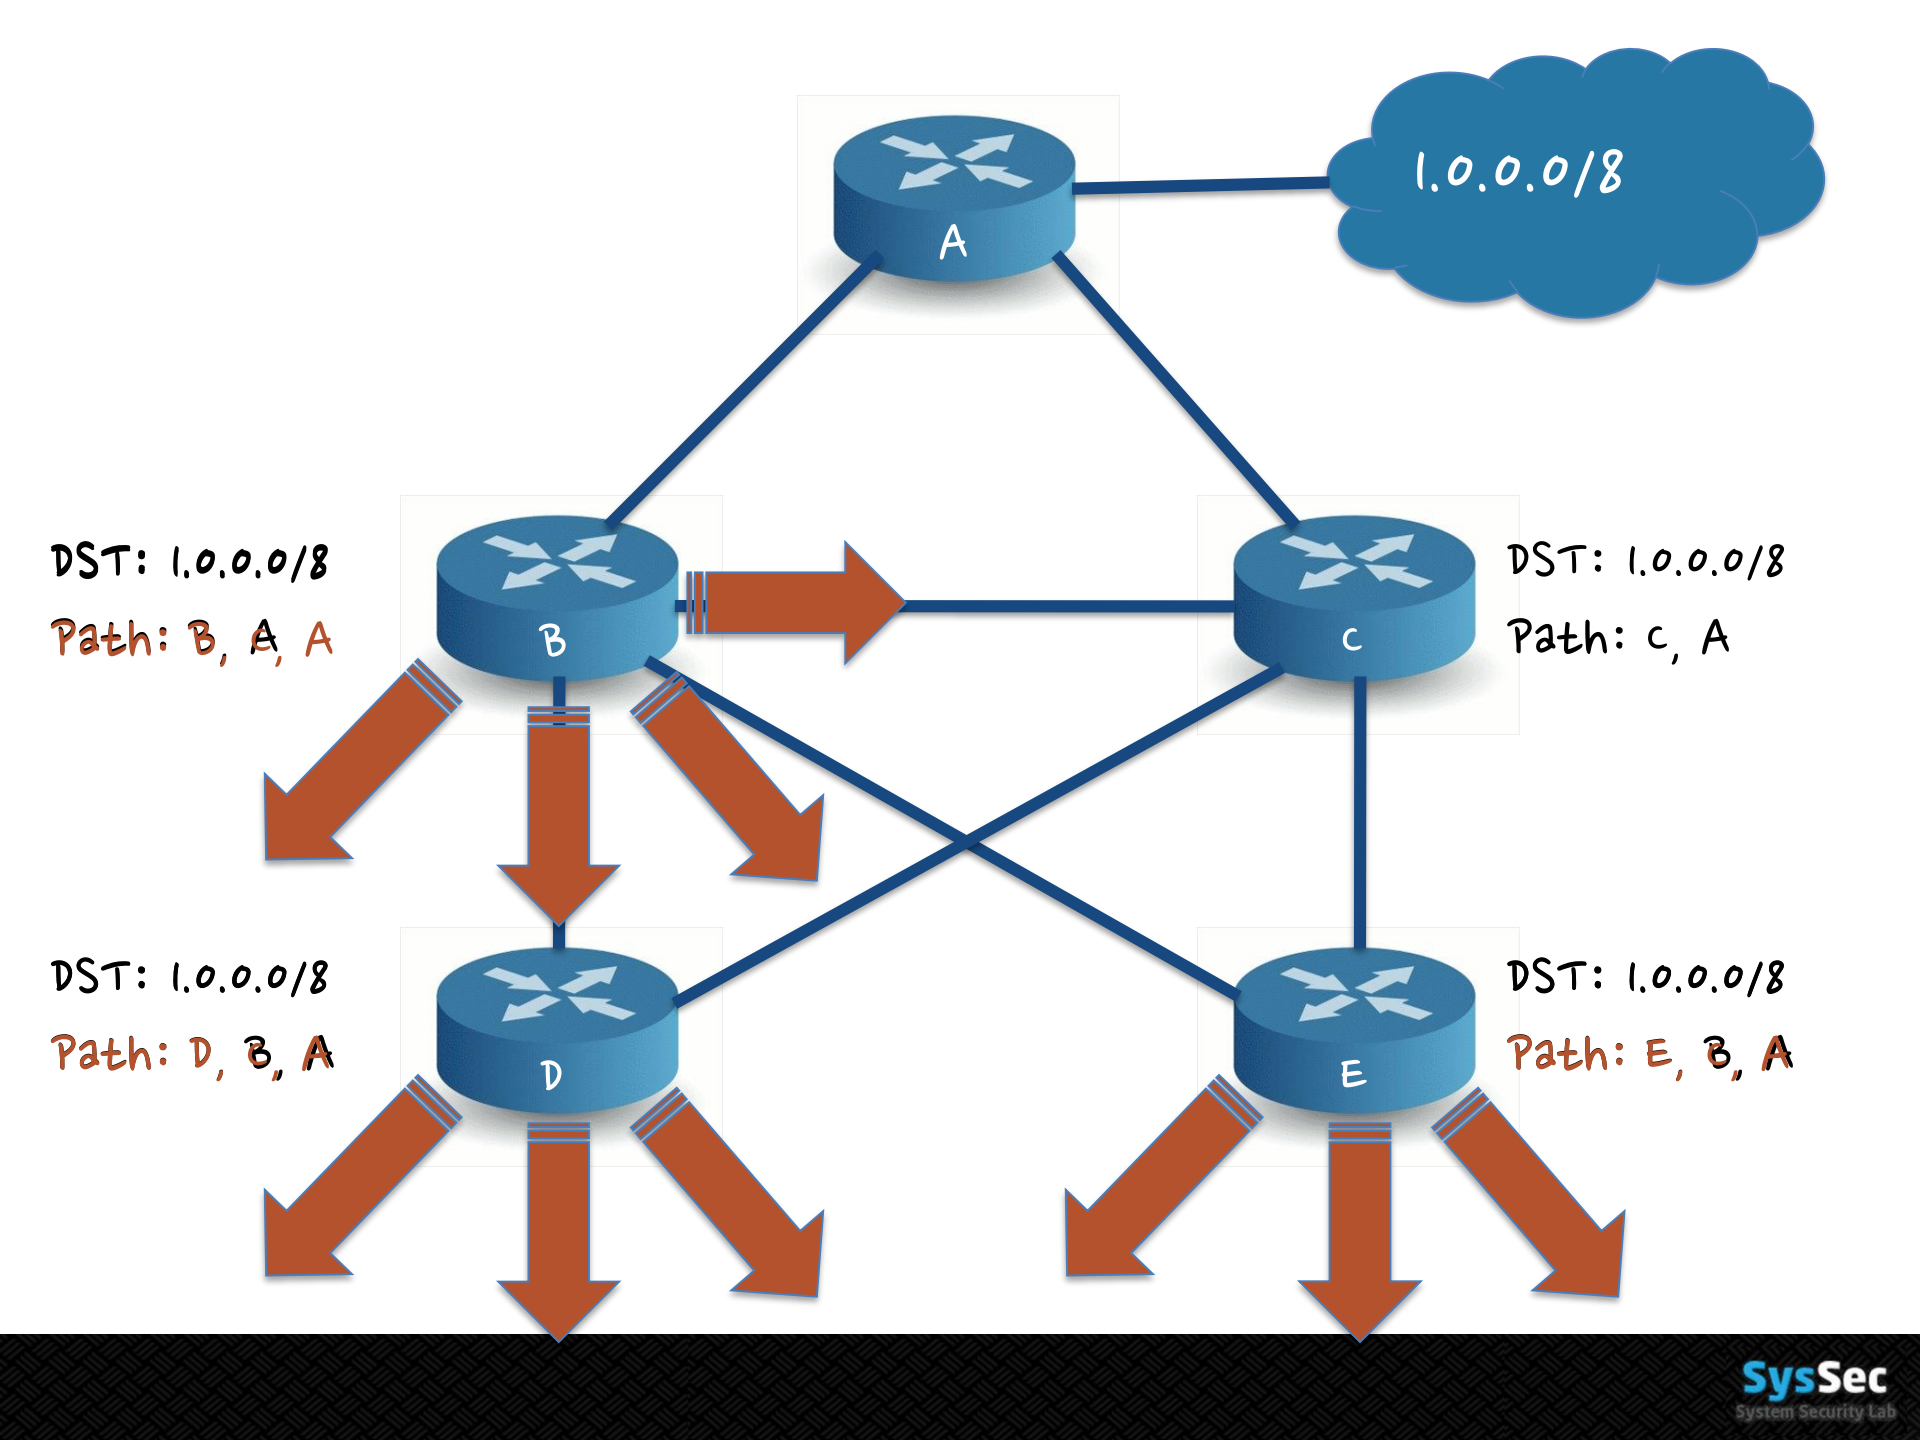
1.0.0.0/8
A
B
C
D
E
DST: 1.0.0.0/8
Path: B, A, A
DST: 1.0.0.0/8
Path: C, A
DST: 1.0.0.0/8
Path: D, B, A
DST: 1.0.0.0/8
Path: E, B, A

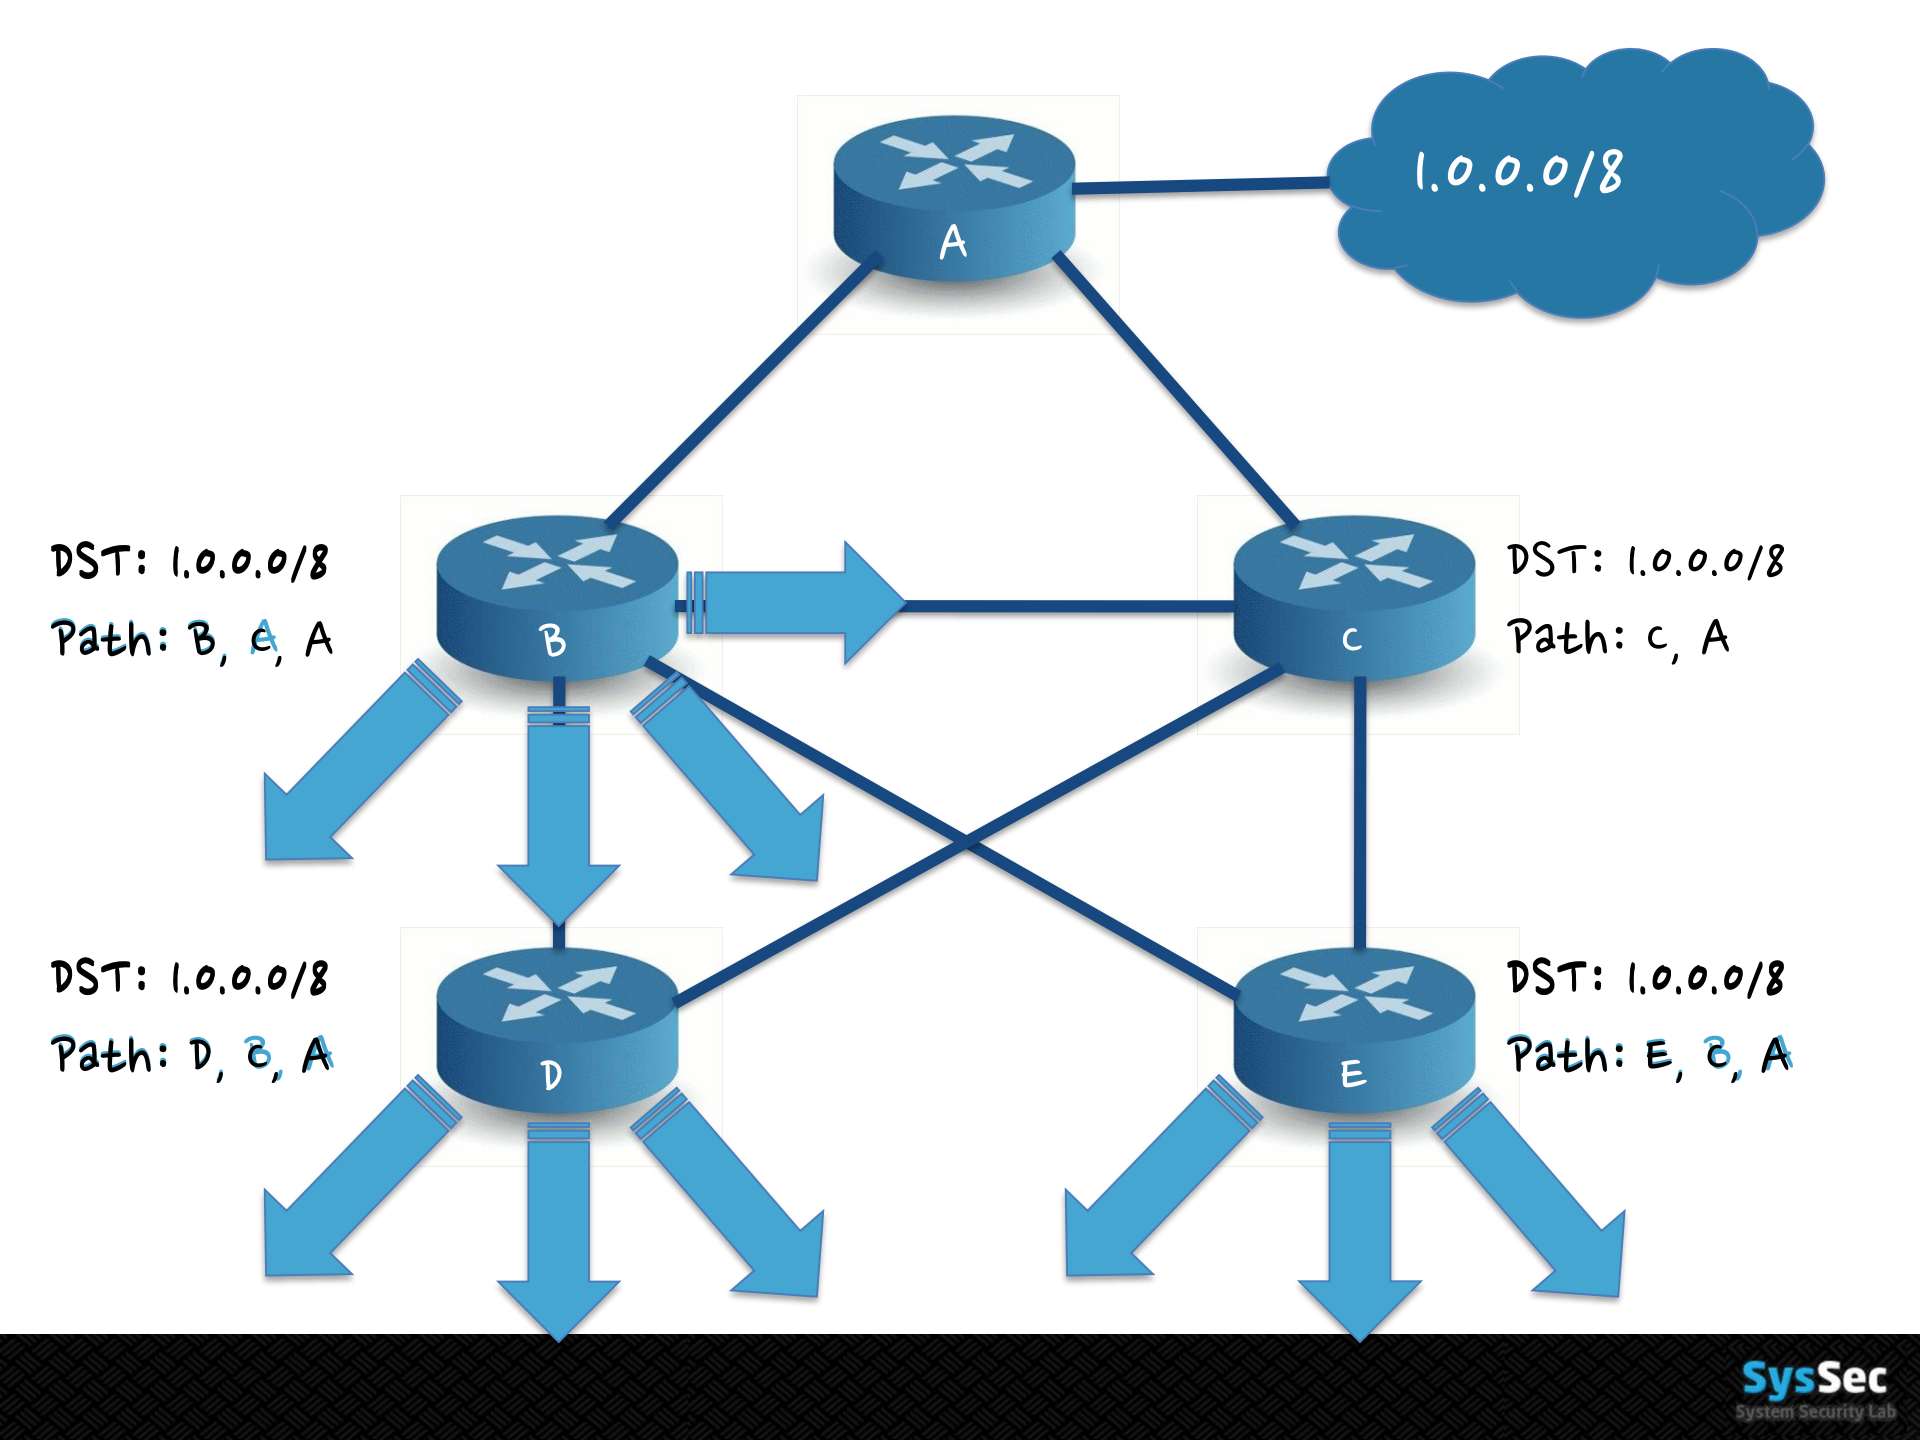
1.0.0.0/8
A
B
C
D
E
DST: 1.0.0.0/8
Path: B, A
DST: 1.0.0.0/8
Path: C, A
DST: 1.0.0.0/8
Path: D, B, A
DST: 1.0.0.0/8
Path: E, B, A

How does the attacker pick links?

How does the attacker direct traffic?

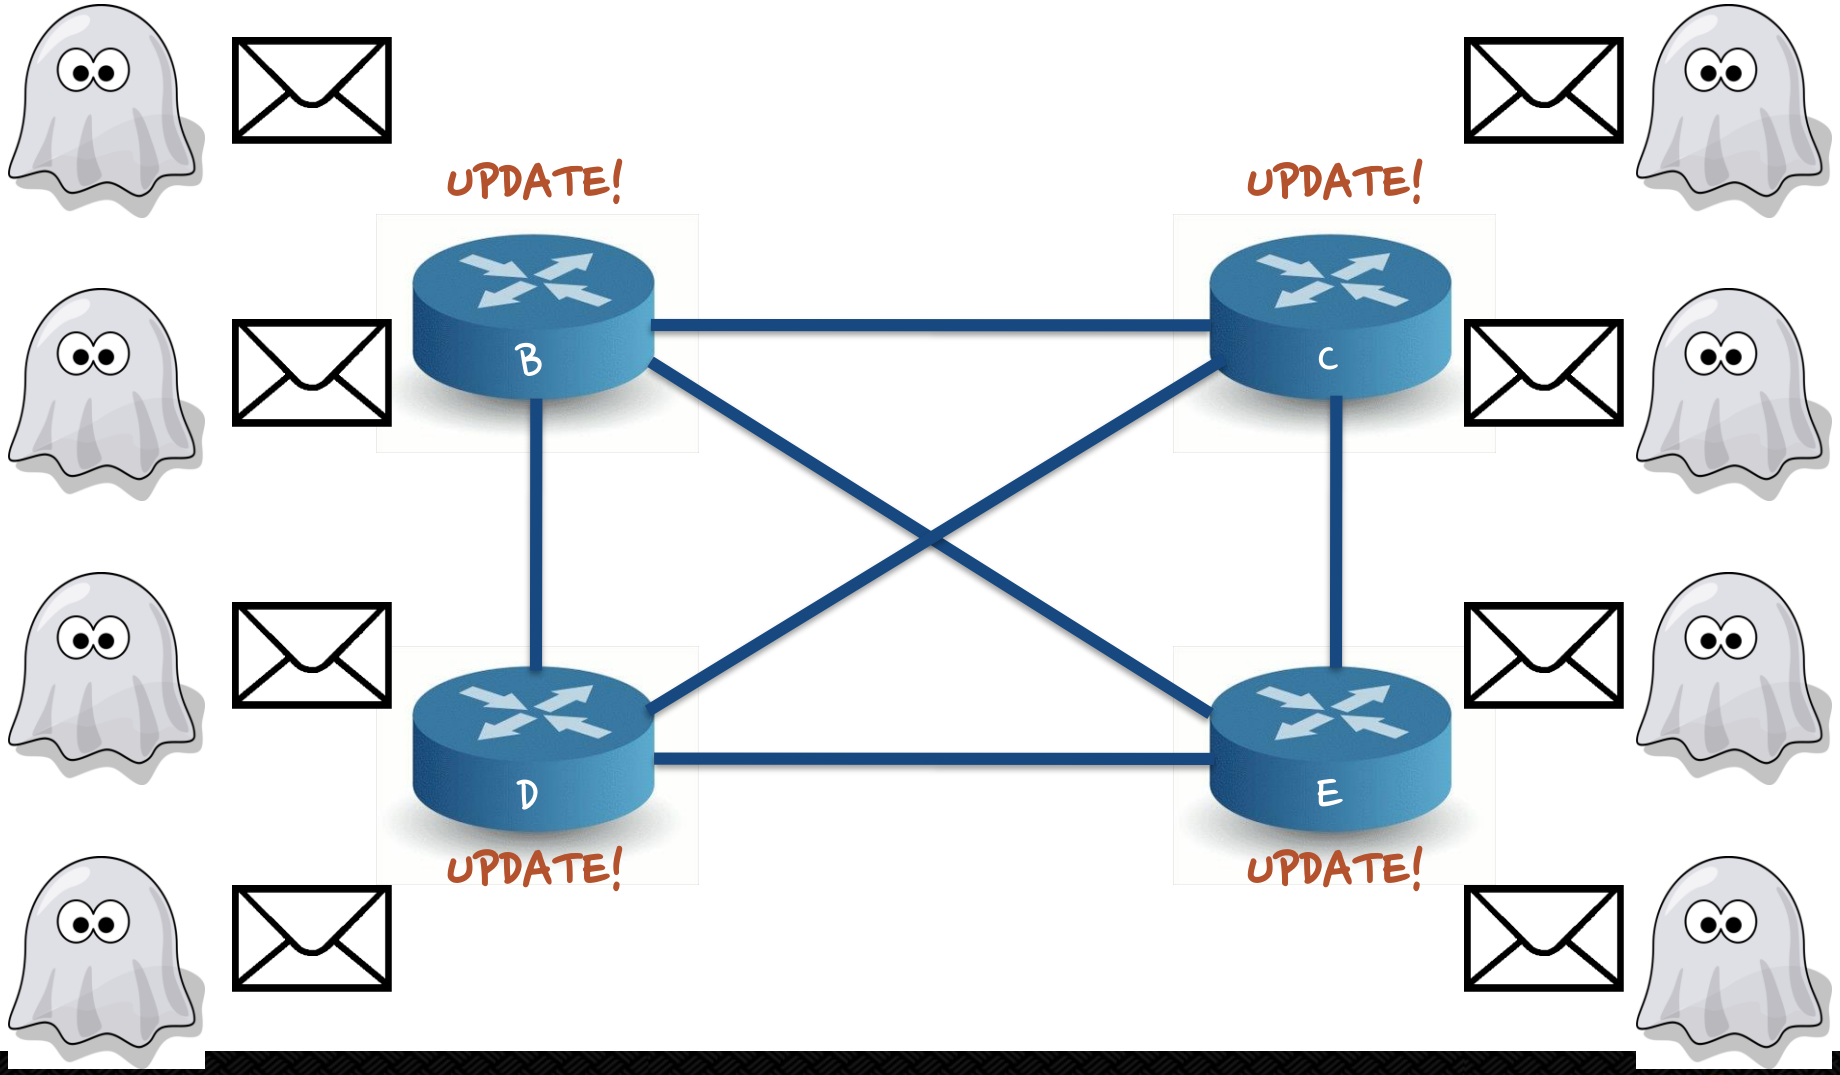

$$C_B(e) = \sum_{s \neq t \in V} \frac{s_{st}(e)}{s_{st}}$$

$$C_B(e) = \sum_{s \neq t \in V} path_{st}(e)$$

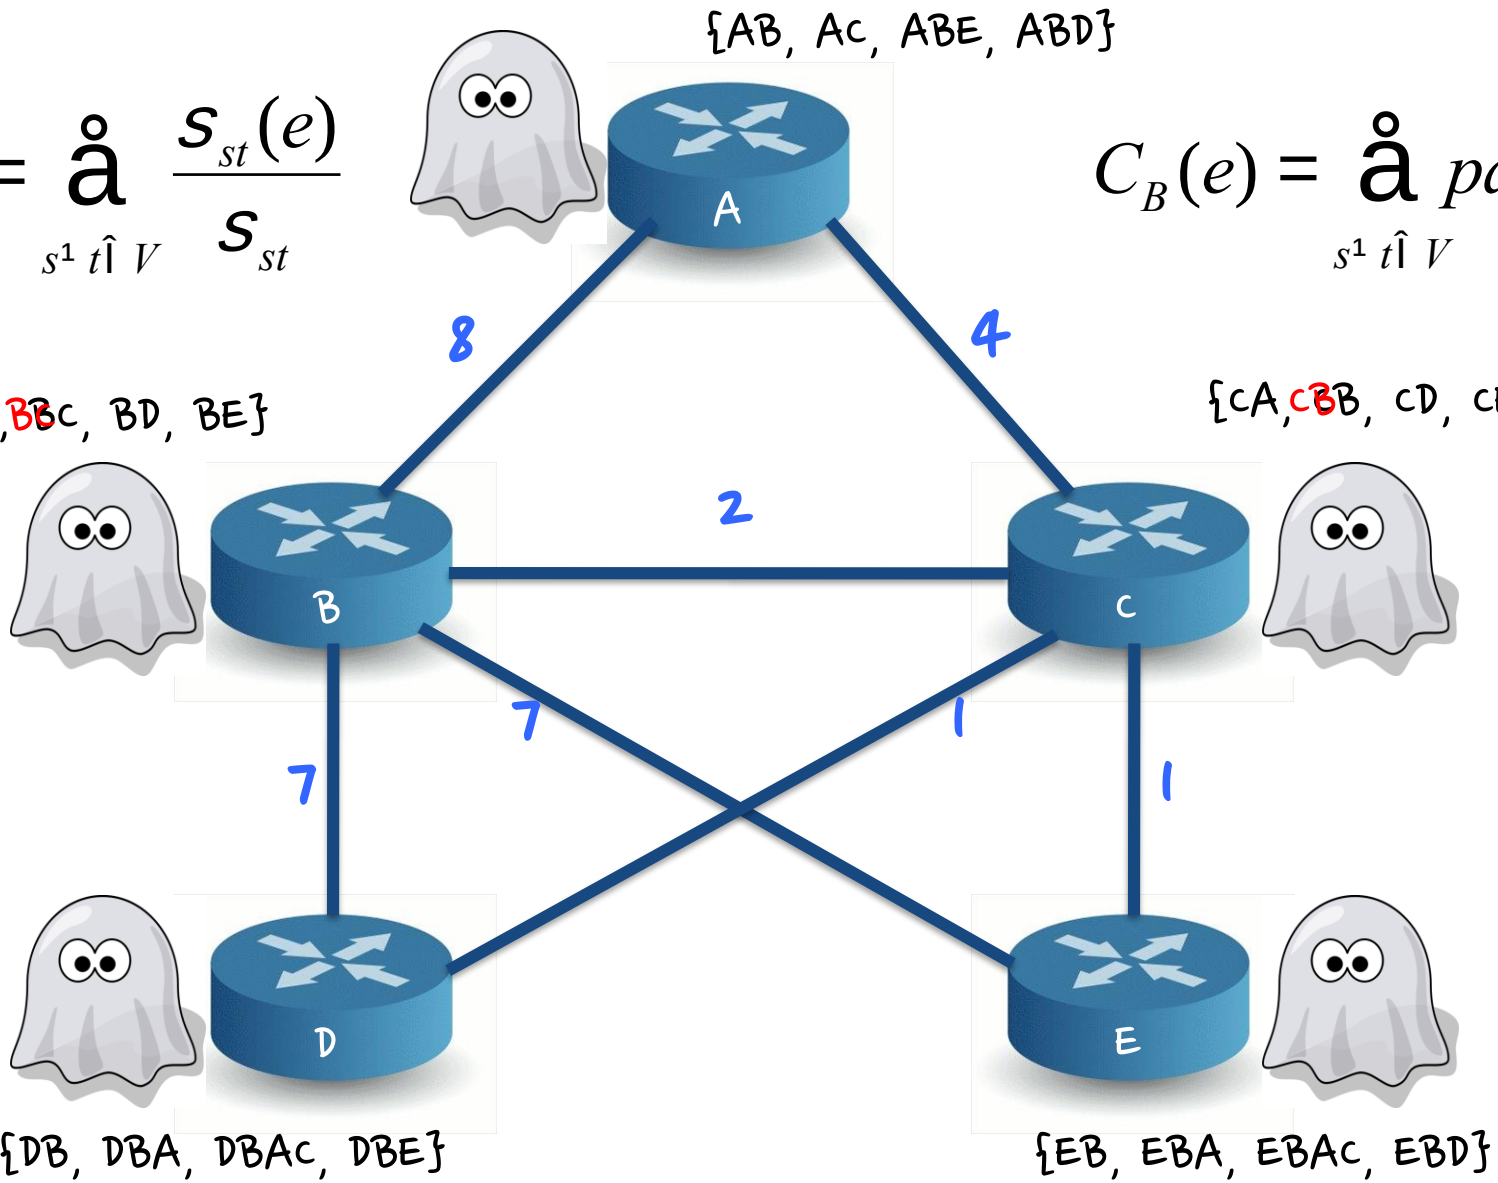

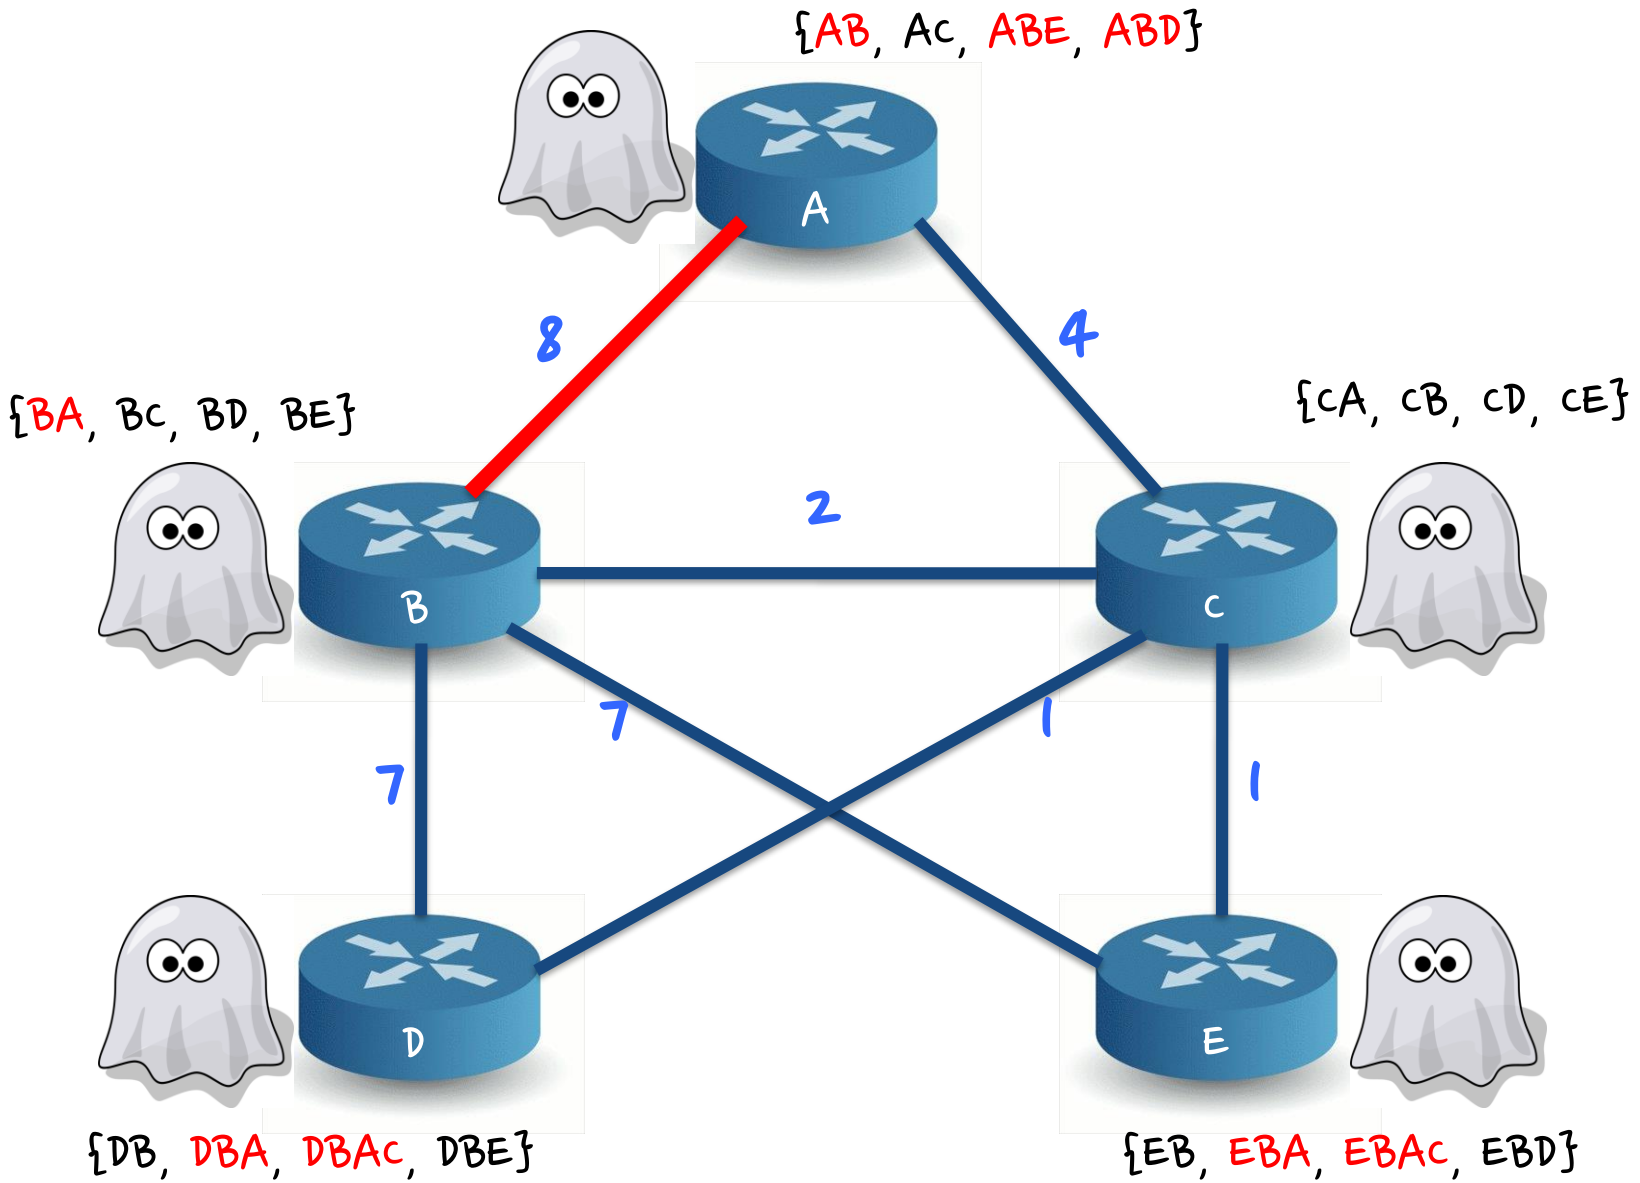
{AB, Ac, ABE, ABD}
A
8
4
{BA, Bc, BD, BE}
{cA, cB, cD, cE}
2
B
c
7
7
1
1
D
E
{DB, DBA, DBAc, DBE}
{EB, EBA, EBAc, EBD}

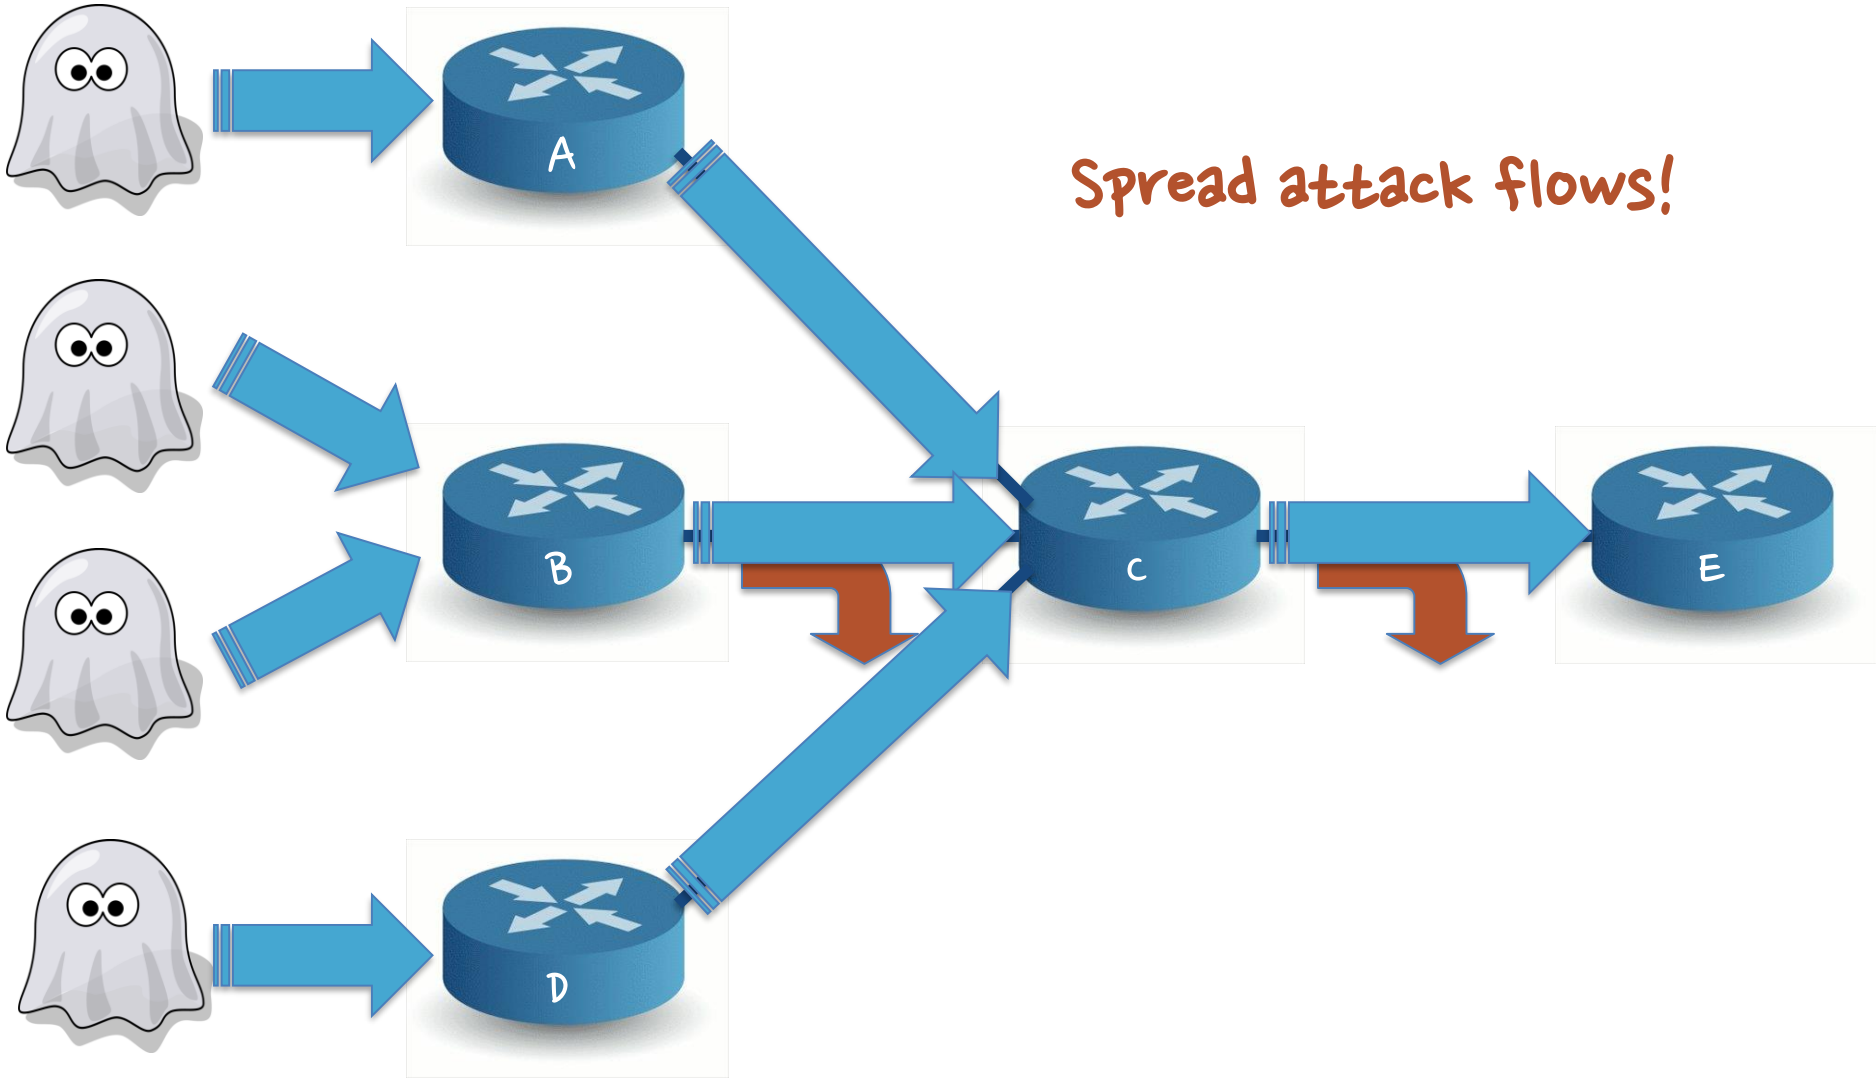
A
B
C
D
E
Spread attack flows!

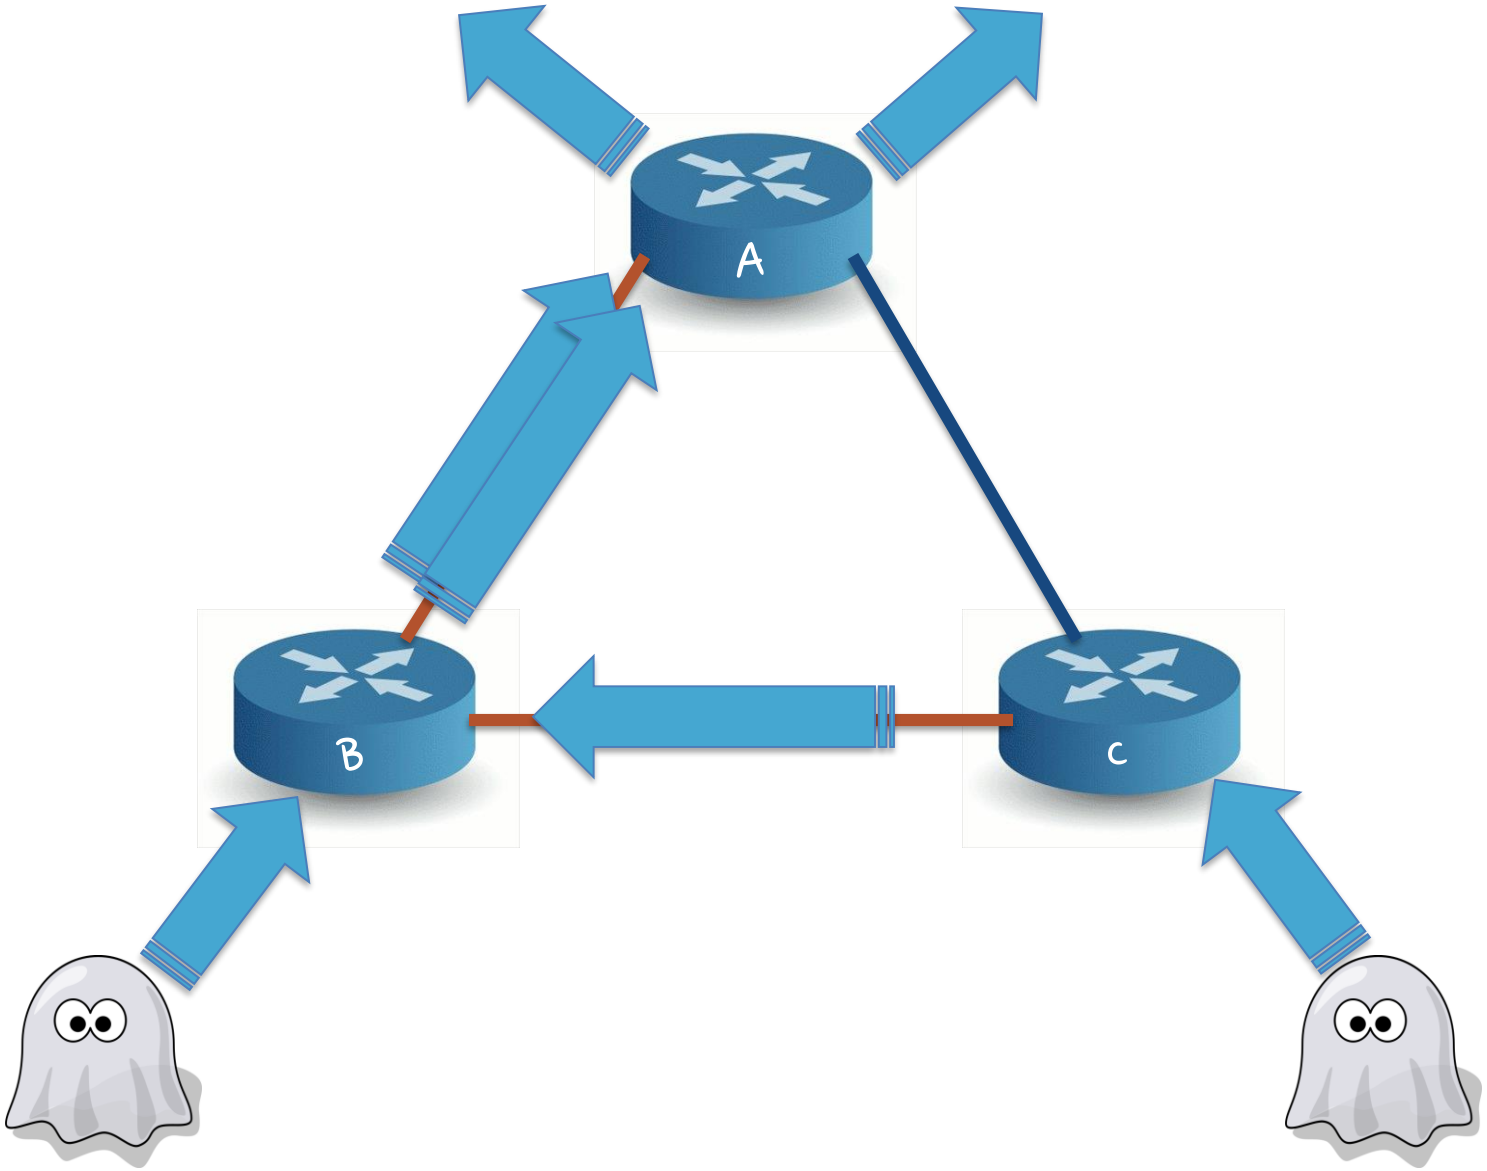
A
B
C

One Target per Attack Flow!

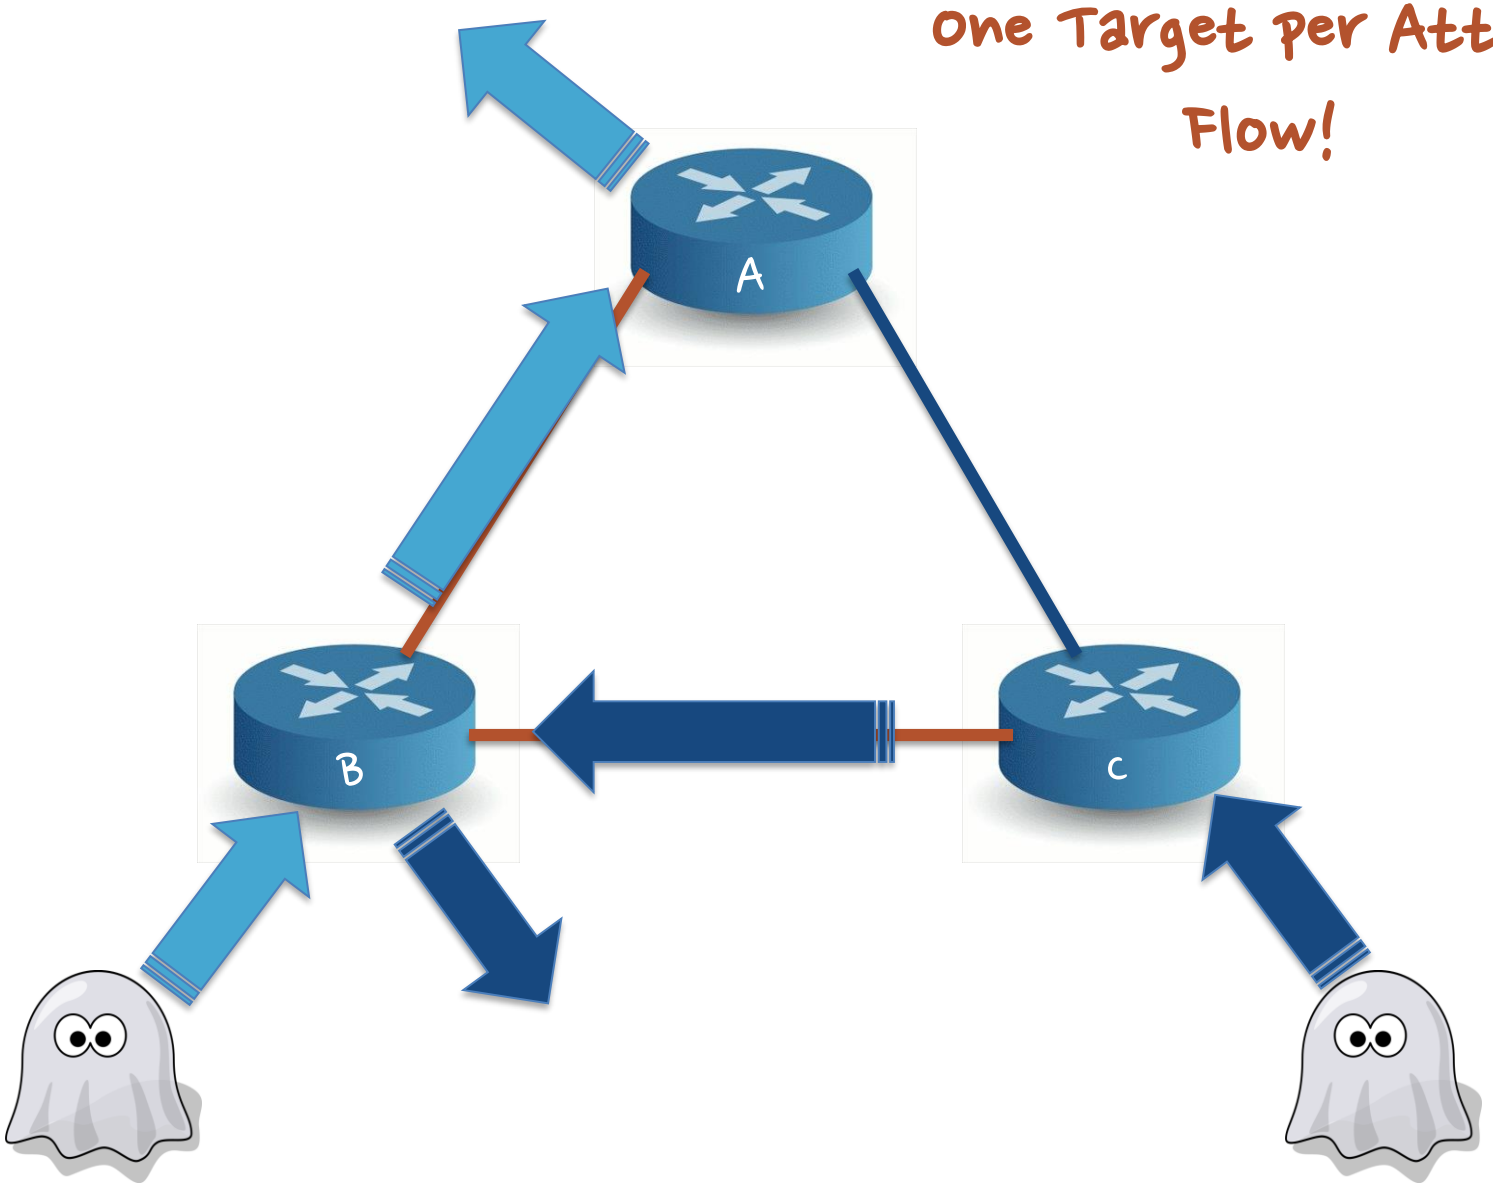

Simulation overview

- Simulator to model network dynamics
  - Topology generated from the Internet
- Routers fully functional BGP speakers
- Bot distribution from waledac
- Bandwidth model worst case for attacker

Targeted link: Any link selected for disruption

Last mile links: un-targeted links that connect fringe ASes to the rest of the network

Transit link: Any link that does not fit the other two

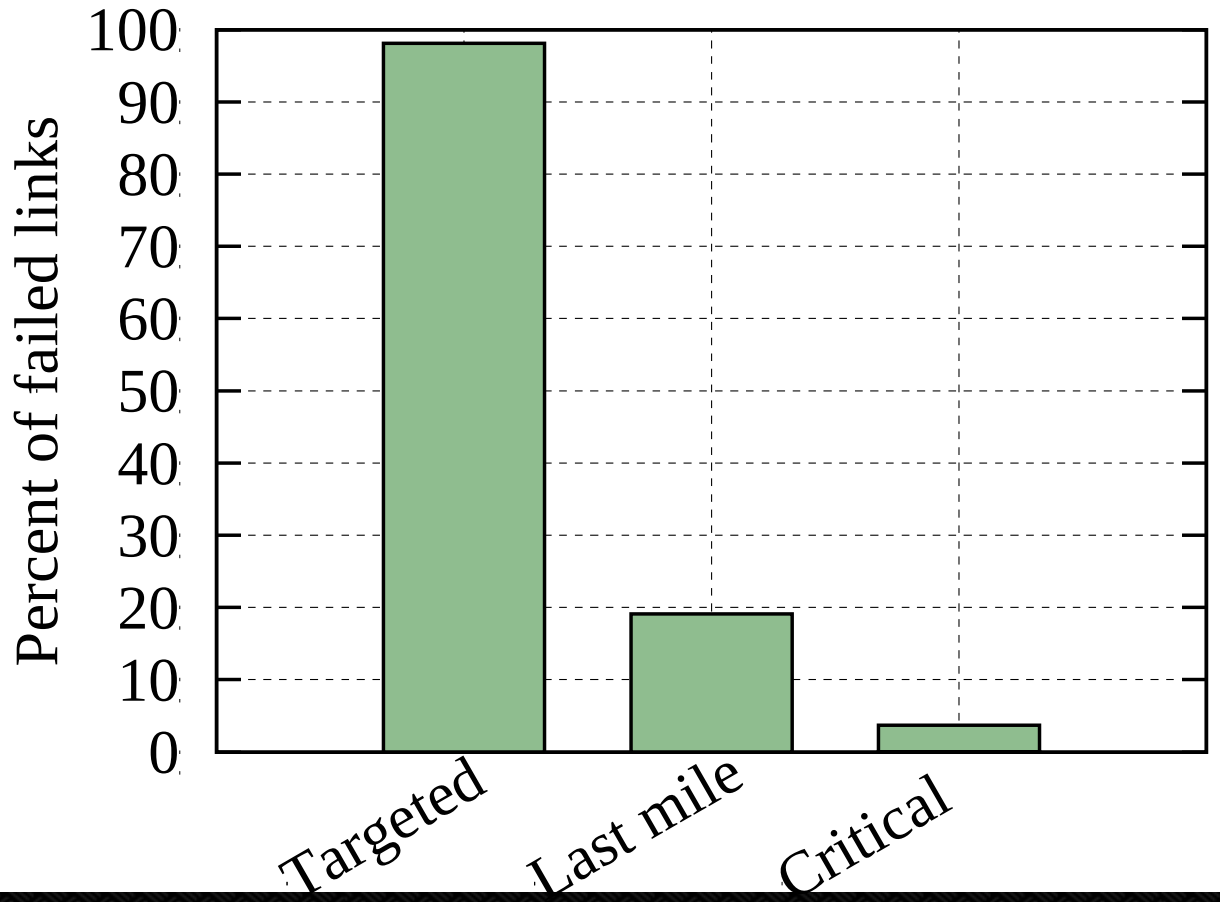

Factors of Normal Load

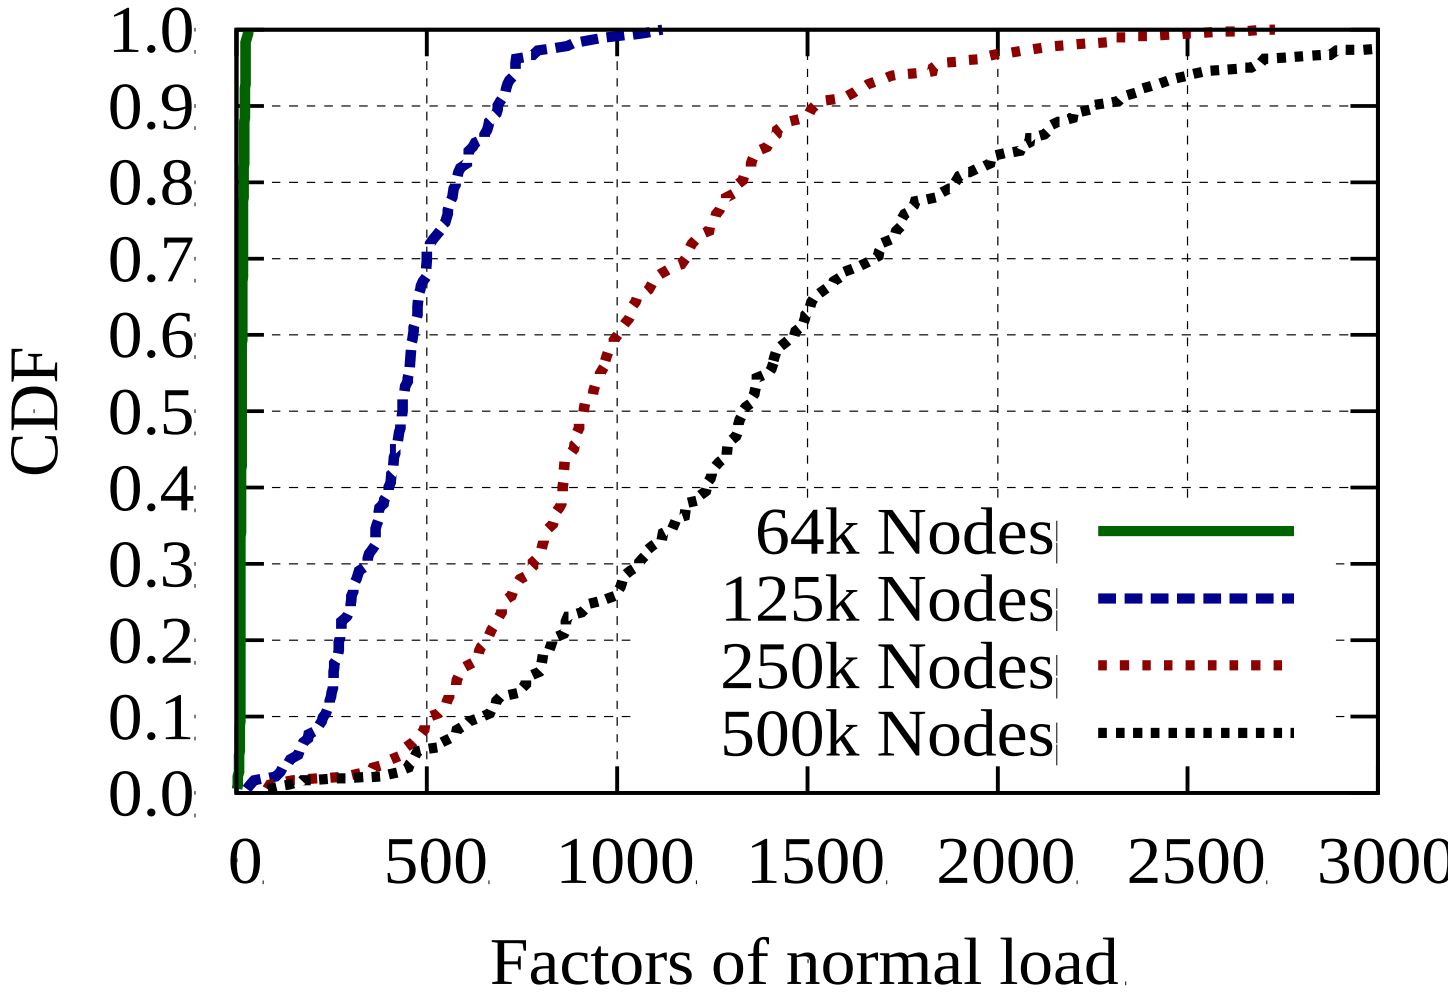

90th percentile of of message loads experienced by routers under attack

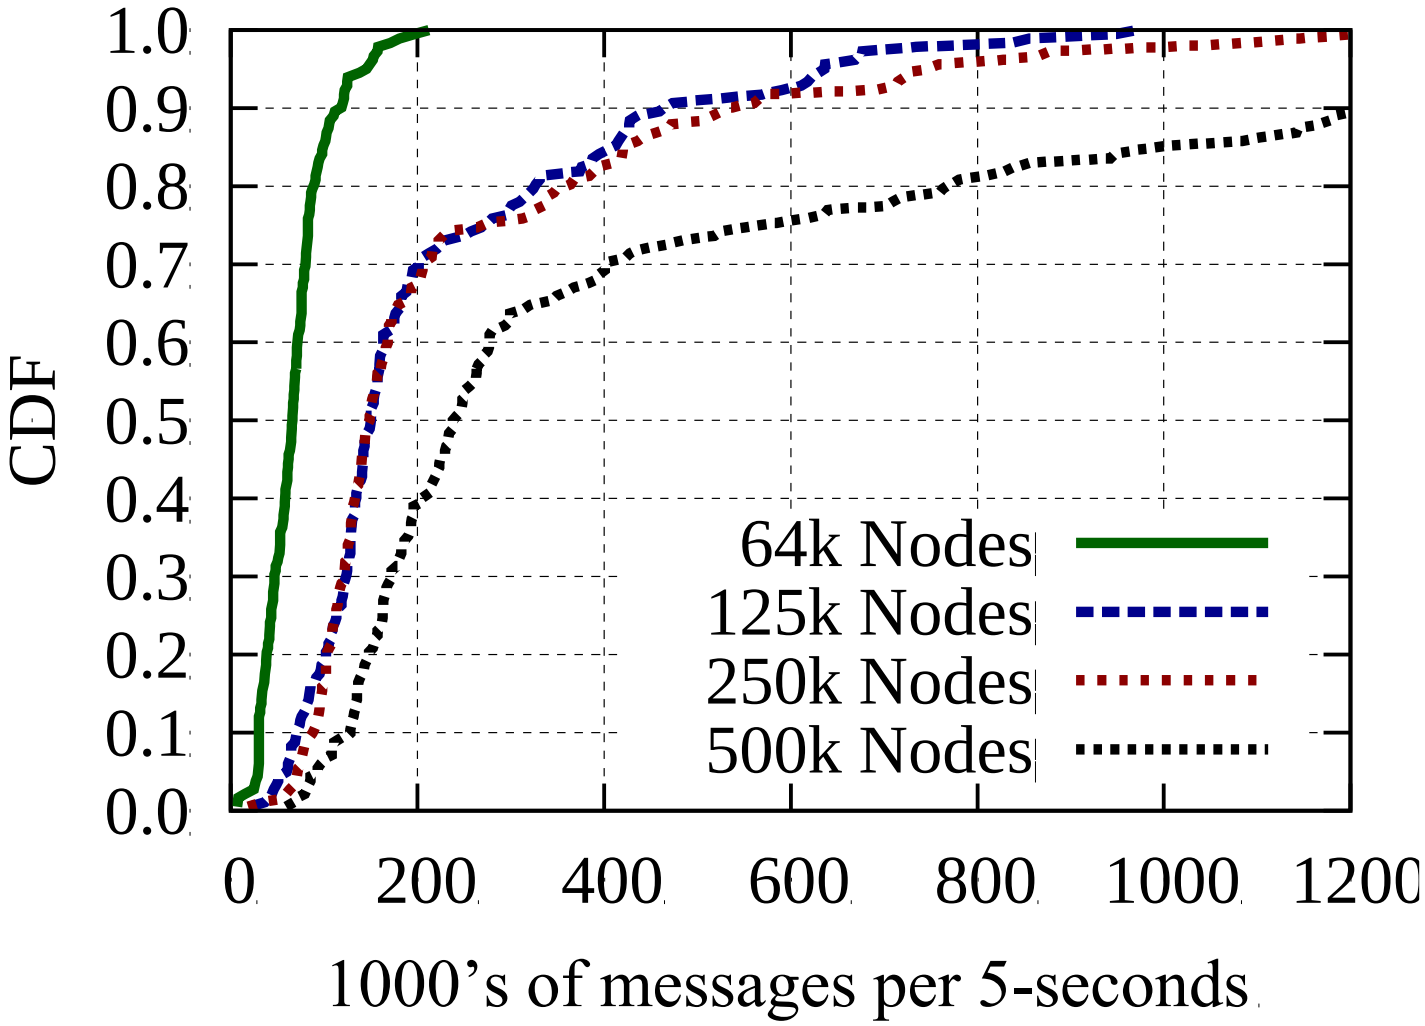

core Routers update Time

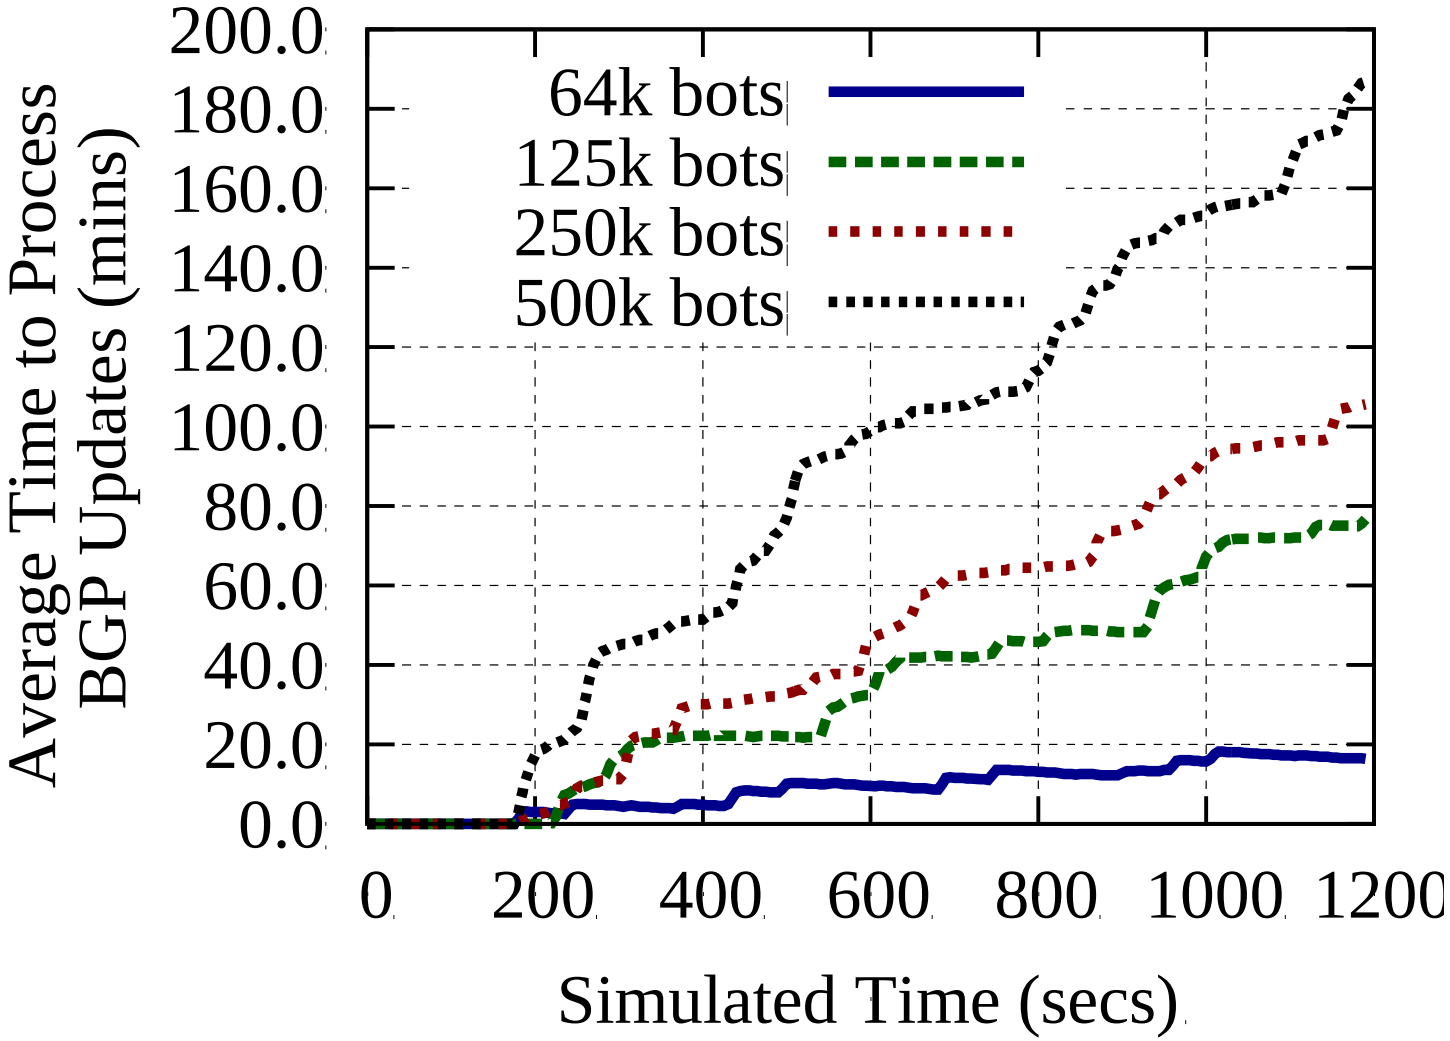

Possible Defenses

- Short Term

  Hold Time = MaxInt

- Long Term

  Perfect QOS

HoldTime = MaxInt

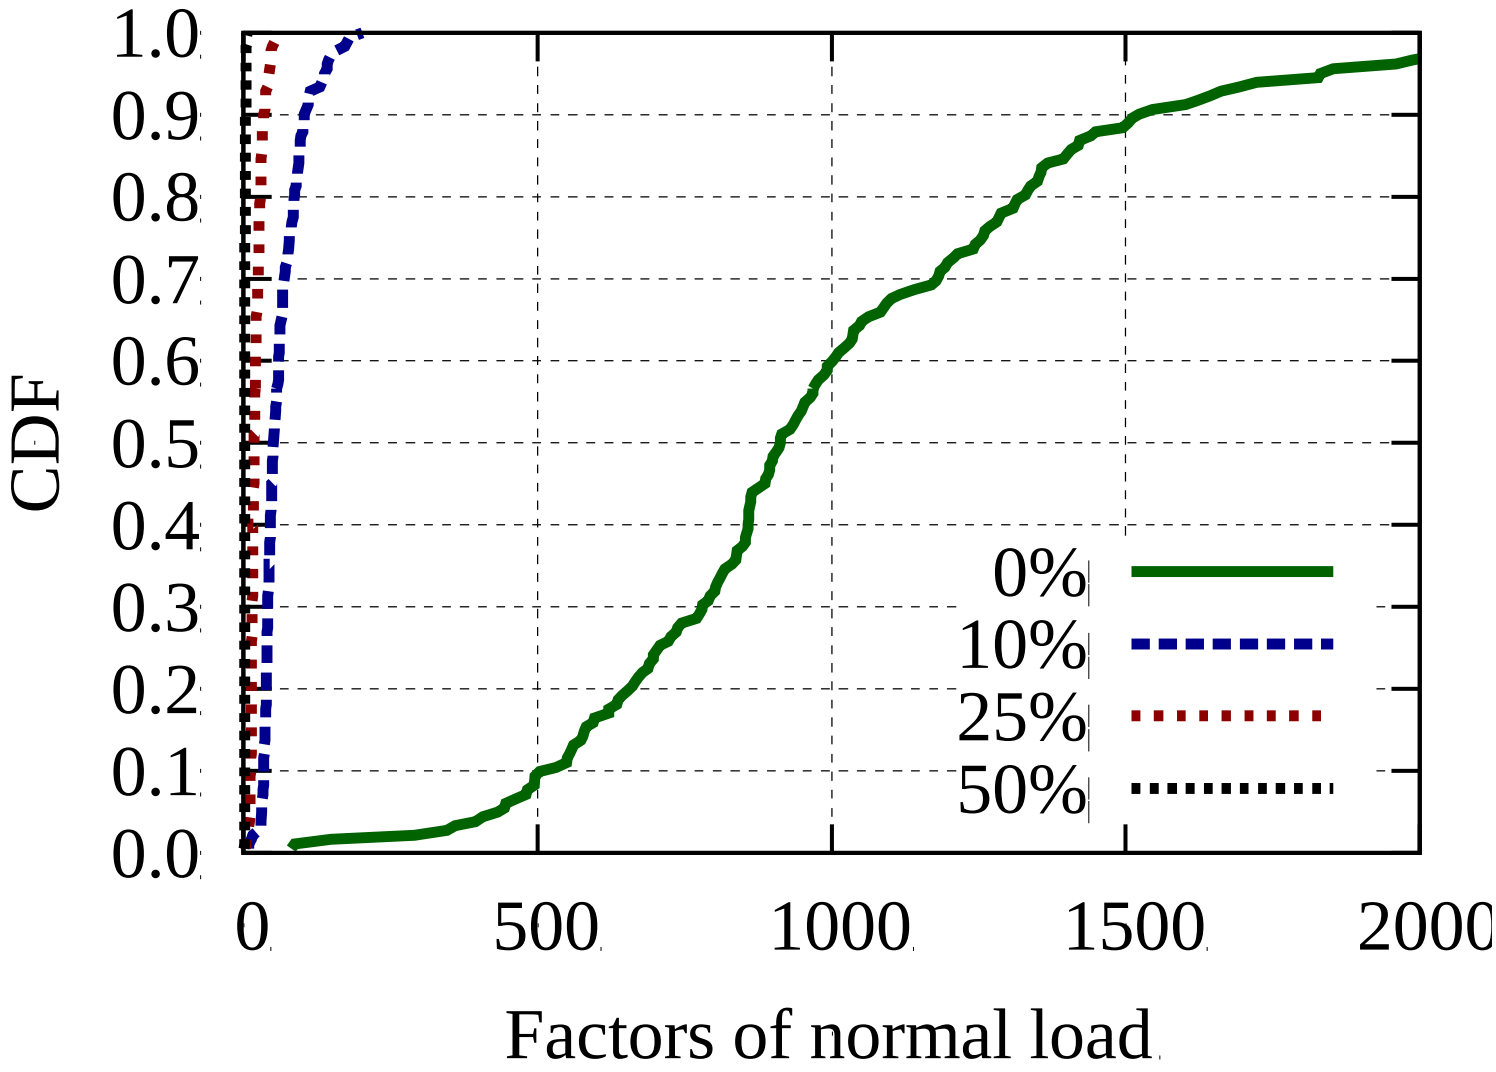

HoldTime = MaxInt

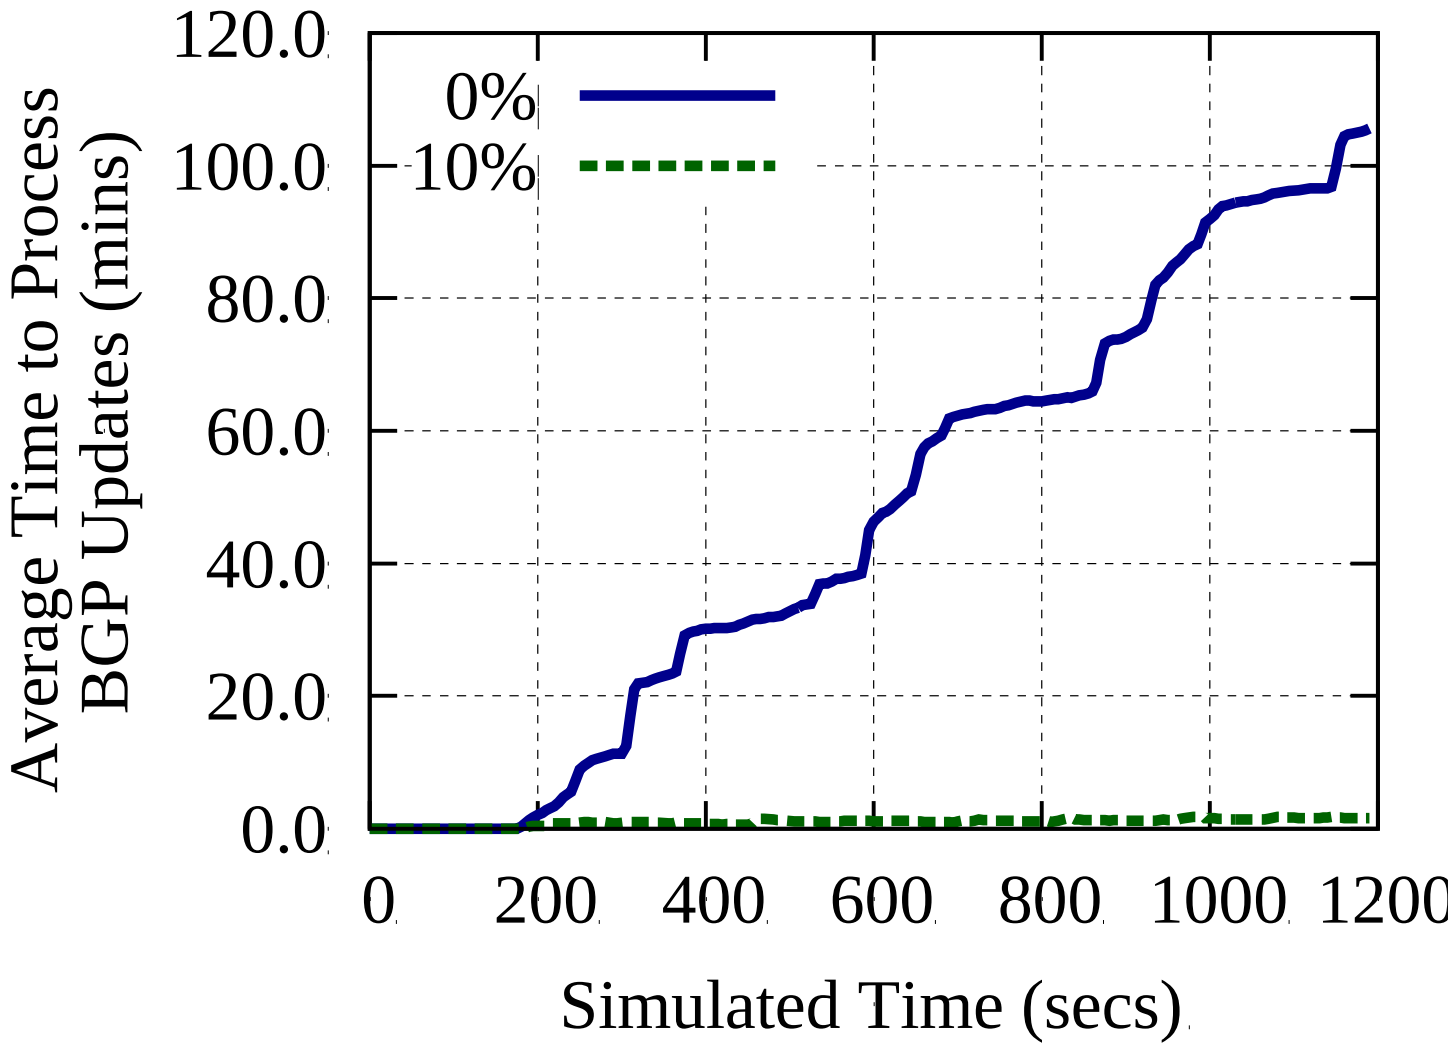

Perfect QoS

- Needs to guarantee control packets must be sent
  - Does not guarantee they will be processed due to oversubscription
- Recommendation
  - (Virtually) Separating control and data plane
  - Sender sides QoS
  - Receiving nodes must process packets in line speed

conclusion

- Adversarial route flapping on an Internet scale
- Implemented using only a modest botnet
- Defenses are non-trivial, but incrementally deployable

Future work (in progress)

- cascaded failure
  - Router failure modeling
- Attacks using remote compromised routers
  - Targeted Attack: Internet Kill Switch
- Router Design for the Future Internet
  - Software router?

BGP Stress Test

- Routers placed in certain states fail to provide the functionality they should.
- Unexpected but perfectly legal BGP messages can place routers into those states
- Any assumptions about the likelyhood of encountering these messages do not apply under adversarial conditions.

Peer Pressure: Exerting Malicious Influence on Routers at a Distance, Max Schuchard, Christopher Thompson, Nicholas Hopper and Yongdae Kim, ICDCS 2013

Attacking Neighborhood (Memory)

- How many BGP updates needed to consume 1GB memory?

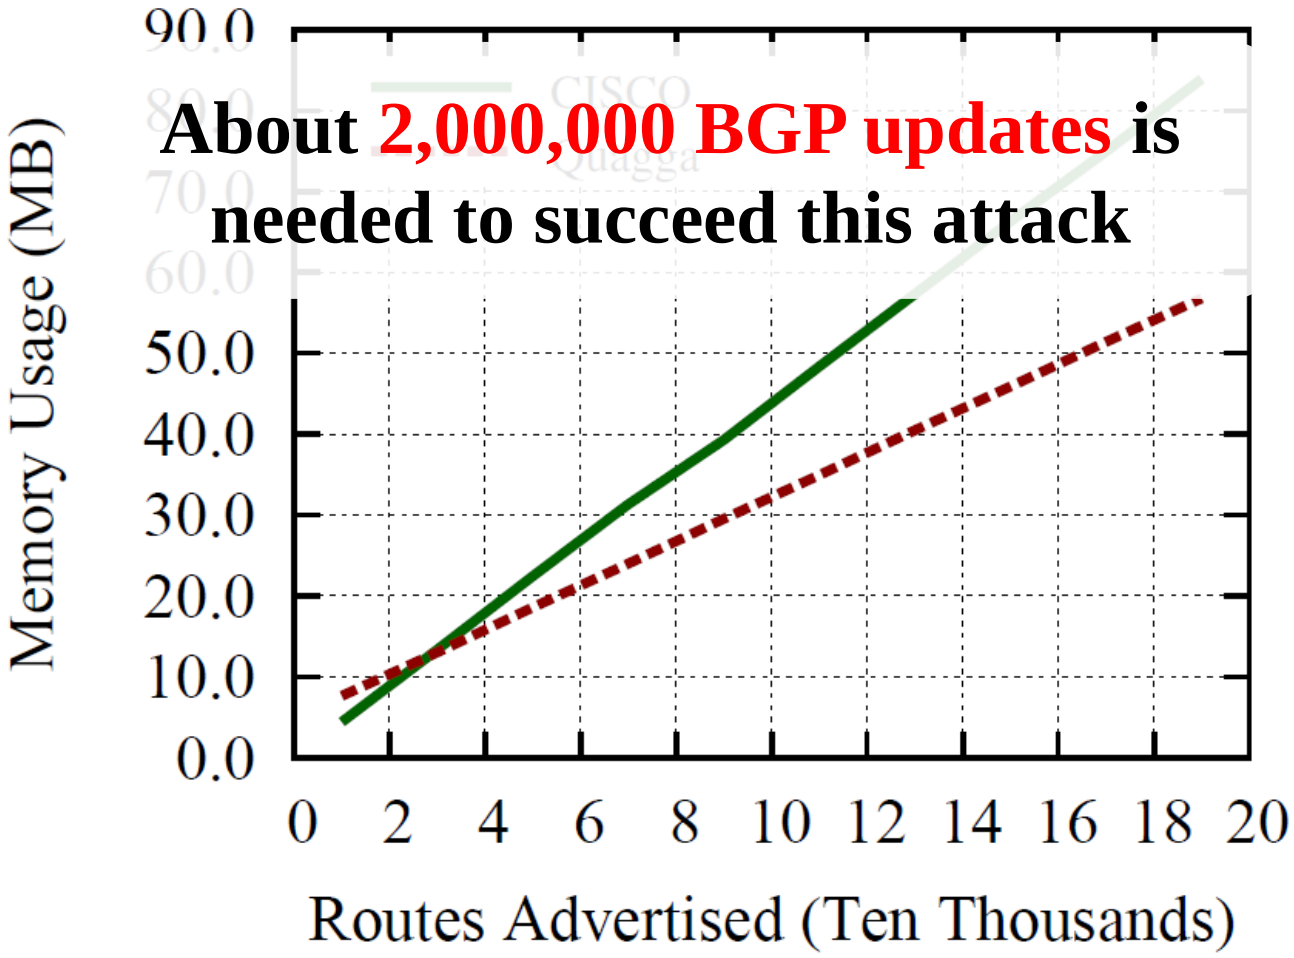

Attacking Neighborhood (Memory)

- Distinct/long length AS paths and community attribute

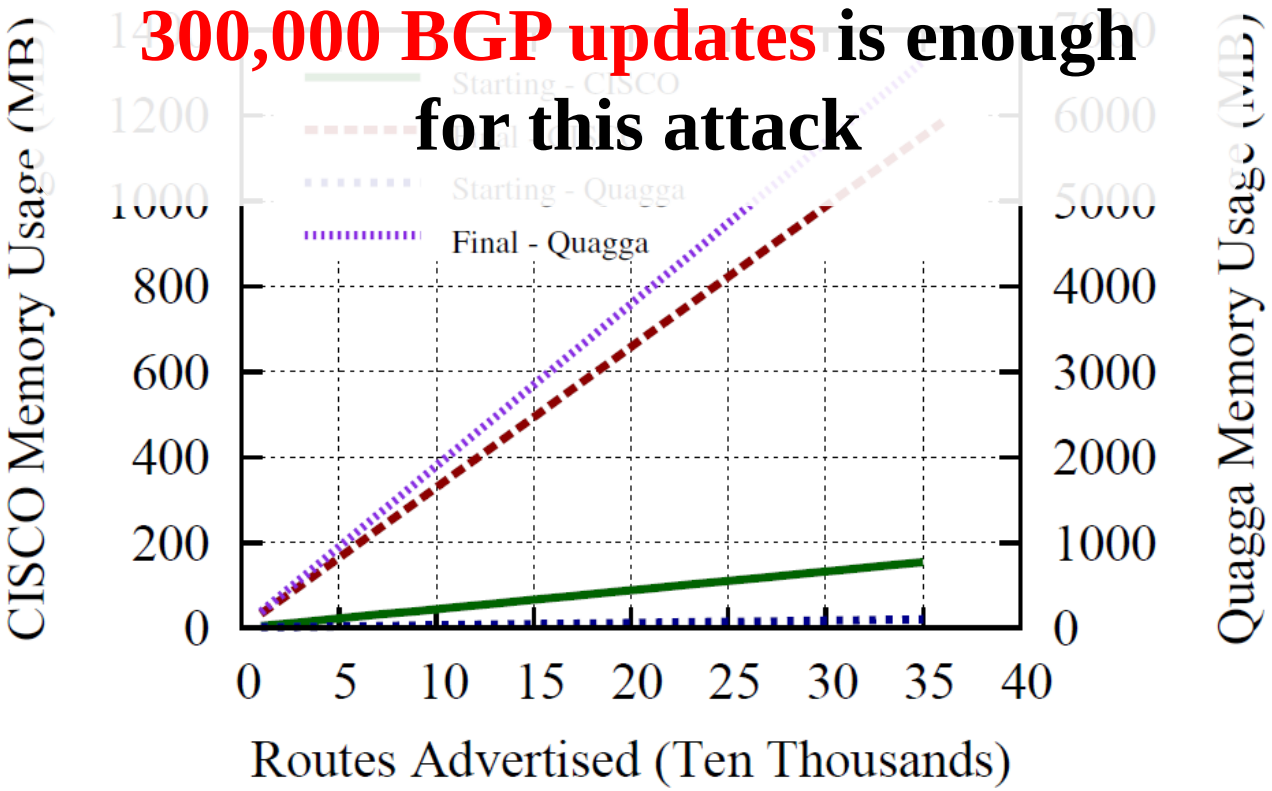

Attacking Neighborhood (CPU)

- Hash collision makes router spend more processing time

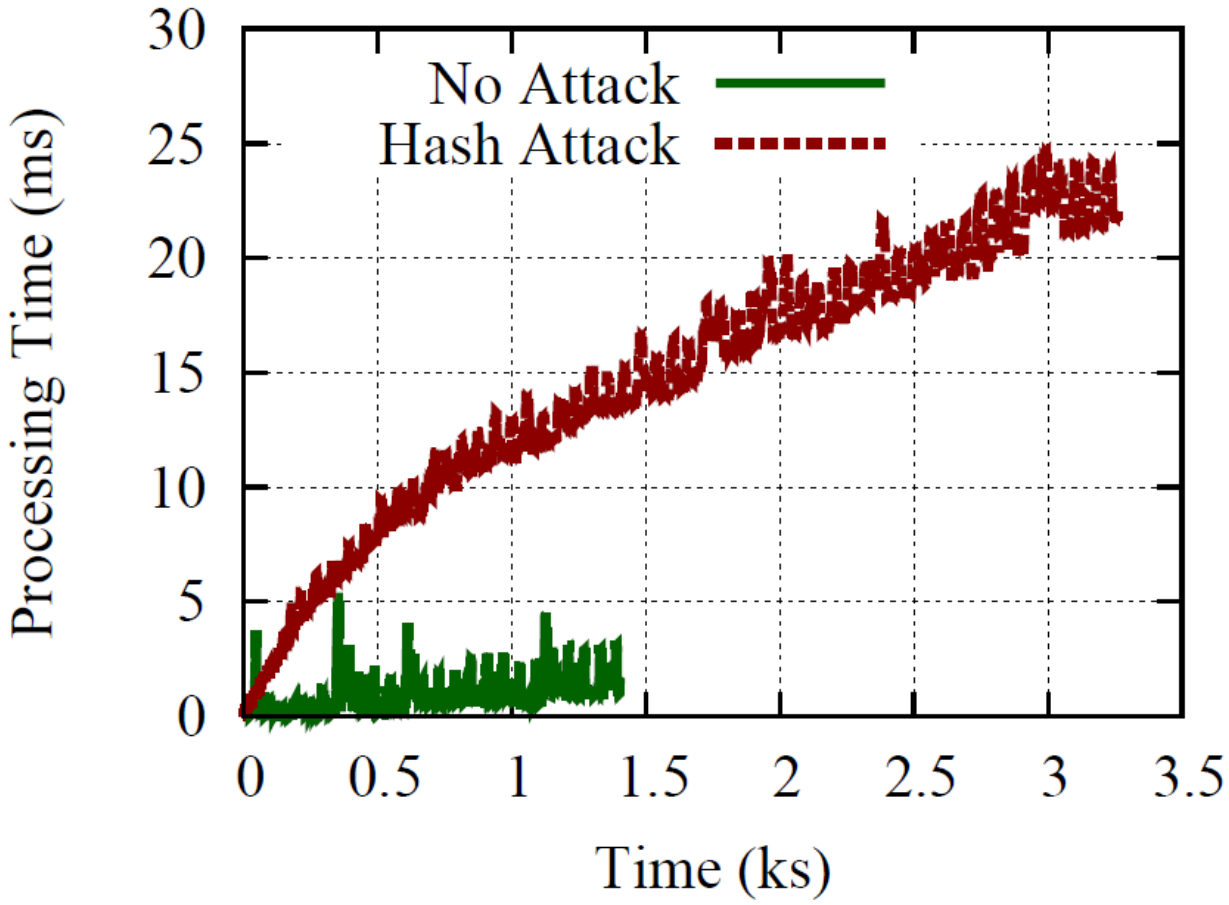

Back Pressure

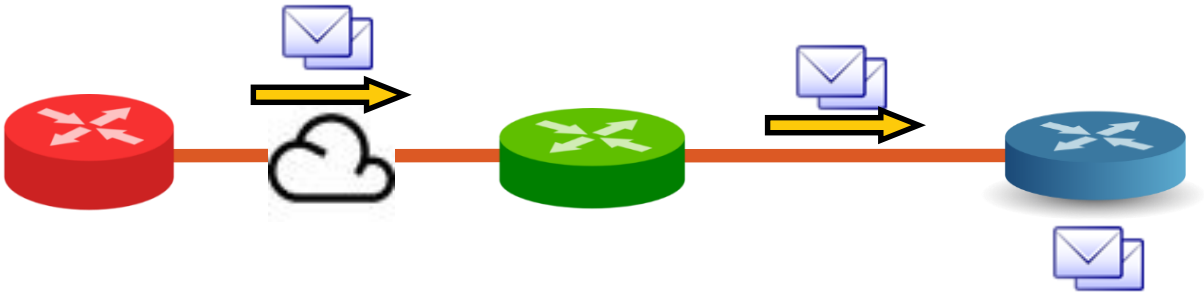

Questions?

- Yongdae Kim
  - email: yongdaek@kaist.ac.kr
  - Home: http://syssec.kaist.ac.kr/~yongdaek
  - Facebook: https://www.facebook.com/y0ngdaek
  - Twitter: https://twitter.com/yongdaek
  - Google “Yongdae Kim”