An Internet-Scale Feasibility Study of BGP Poisoning

Jared M. Smith, *Kyle Birkeland, Tyler McDaniel,* ***Max Schuchard***

AIMS 2019, 4/16/18

jms@vols.utk.edu

University of Tennessee
VOLSEC
— COMPUTER SECURITY LAB —

Full Paper: https://tiny.utk.edu/bgp

BGP Poisoning

- **Conflicting** research, **not actively** measured:
  - Smith et al. Nyx (S&P ‘18) vs. *Feasible* Nyx Tran et al. (S&P ‘19)
  - Schuchard et al. RAD (CCS ‘12) vs. Nasr et al. Waterfall of Decoys (CCS’ 17)
- **Existing** research, **limited** measurements:
  - Anwar et al. Interdomain Policies (IMC ‘15)
  - Katz-Basset et al. LIFEGUARD (SIGCOMM ’12)
- **Existing** research, **dated** measurements:
  - Bush et al. Internet Optometry (IMC ‘09)
- **Specifications** versus **reality**
  - BGP RFC best practices doc recommends filtering over 50 AS-path length
  - Community forums and BGP observations show paths over 50

Full Paper: https://tiny.utk.edu/bgp

We aim to **resolve** these issues, **highlight** discrepancies, **evaluate** accuracy of BGP simulation/emulation, and **inspire** future BGP poisoning work, with **active** measurements and analysis.

Our Approach

Detour Path Discovery System

- Executes BGP Poisoning for arbitrary steered AS
- Can be executed from any BGP router for specified prefix
- Entirely done with software
- Coordinated through globally distributed infrastructure

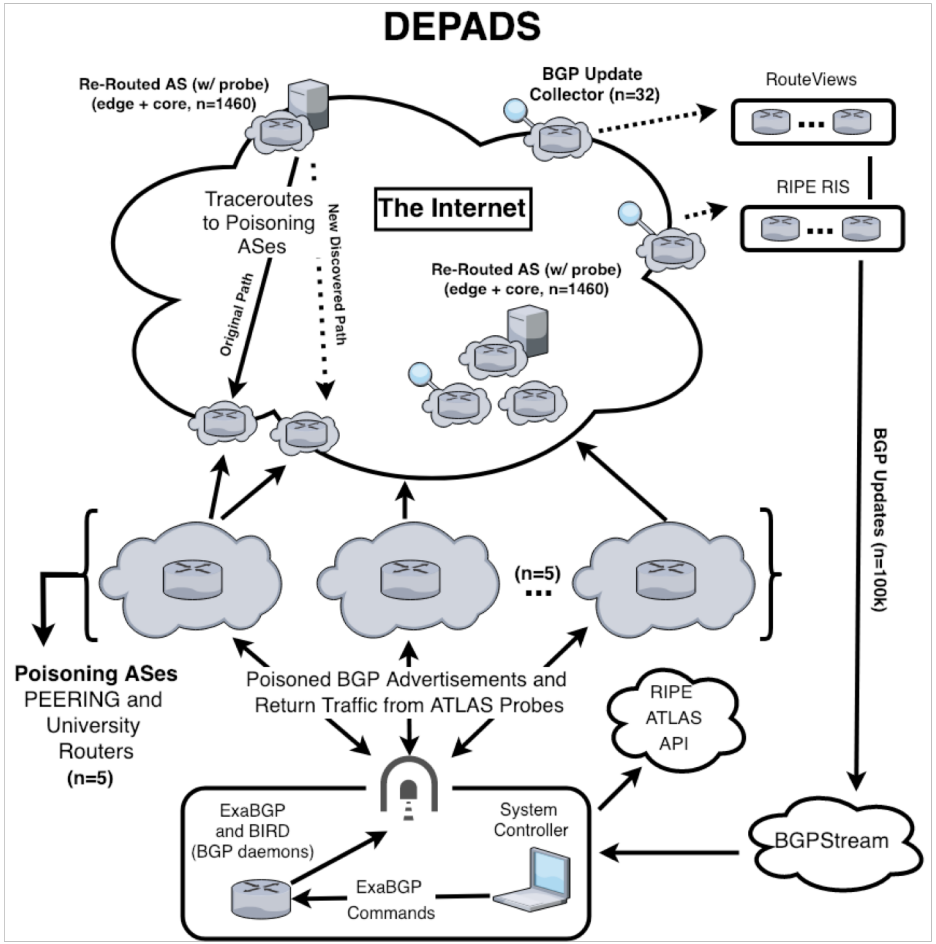

Full Paper: https://tiny.utk.edu/bgp

Infrastructure

Infrastructure	Source
5 BGP routers	PEERING and UT
8 IP prefixes	PEERING and UT
5,000+ distinct vantage points	RIPE ATLAS
3 countries	US, Amsterdam, Brazil
32 BGP collectors	CAIDA BGPStream*

*Collects BGP Updates from RouteViews and RIPE RIS

In total, we measure **1,460** instances of BGP poisoning across **3% of ASes on the Internet**.

(Largest BGP Poisoning sample size in any existing literature)

Active Measurements

- Ability to re-route across entire original AS-path
- Real-world comparison with prior simulations
- Predicting who can re-route w/ BGP poisoning

- Filtering of poisoned routes
- Routing Working Groups behavior

- Default route prevalence
- Reachability of /25's

BACKGROUND

BGP Poisoning

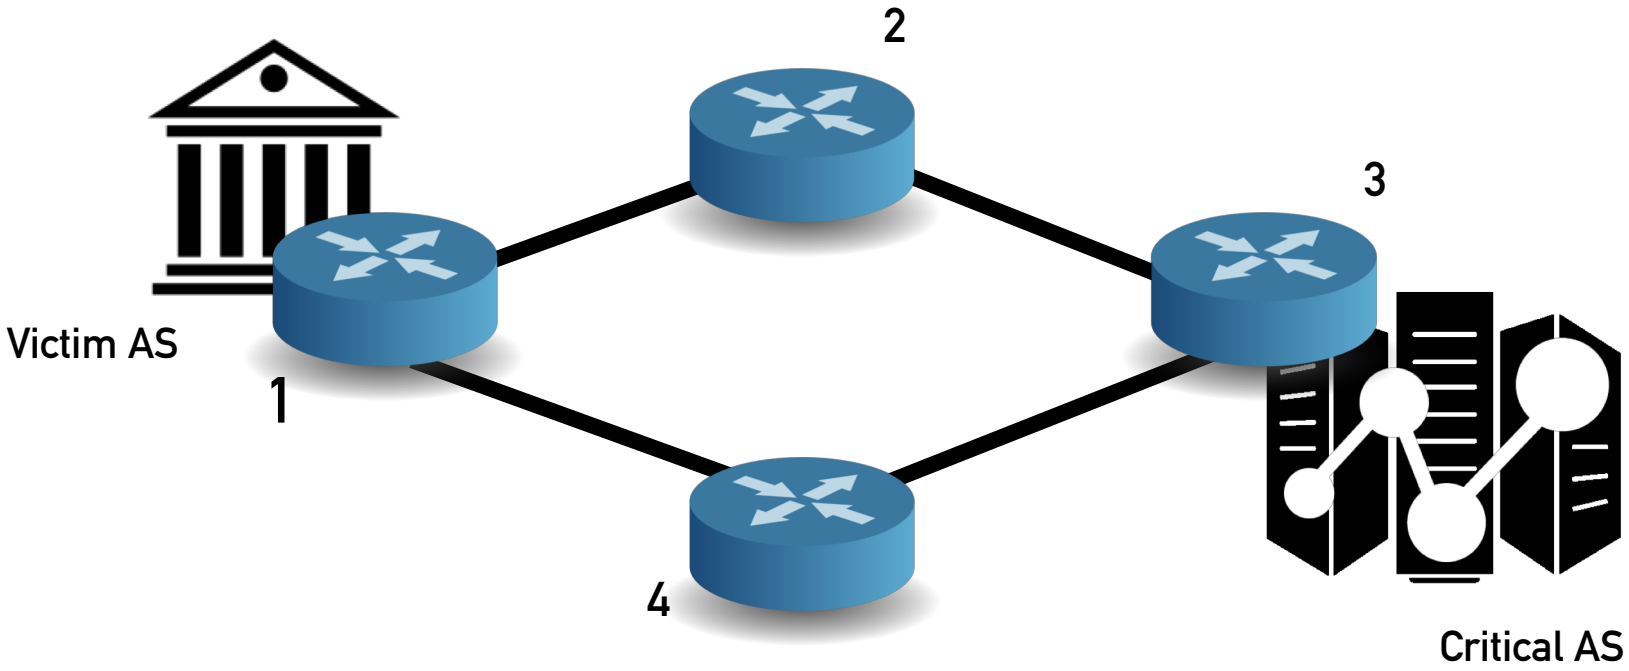

BGP Poisoning

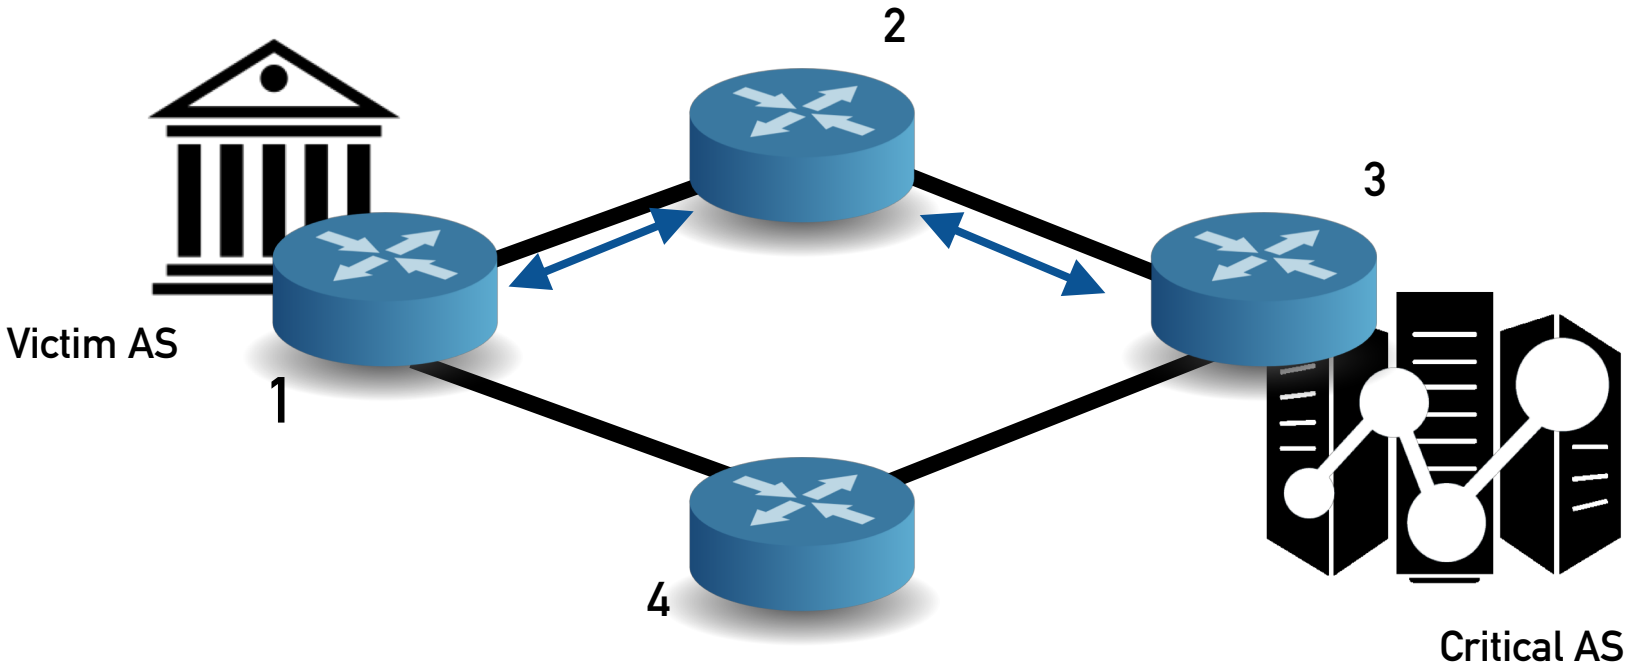

BGP Poisoning

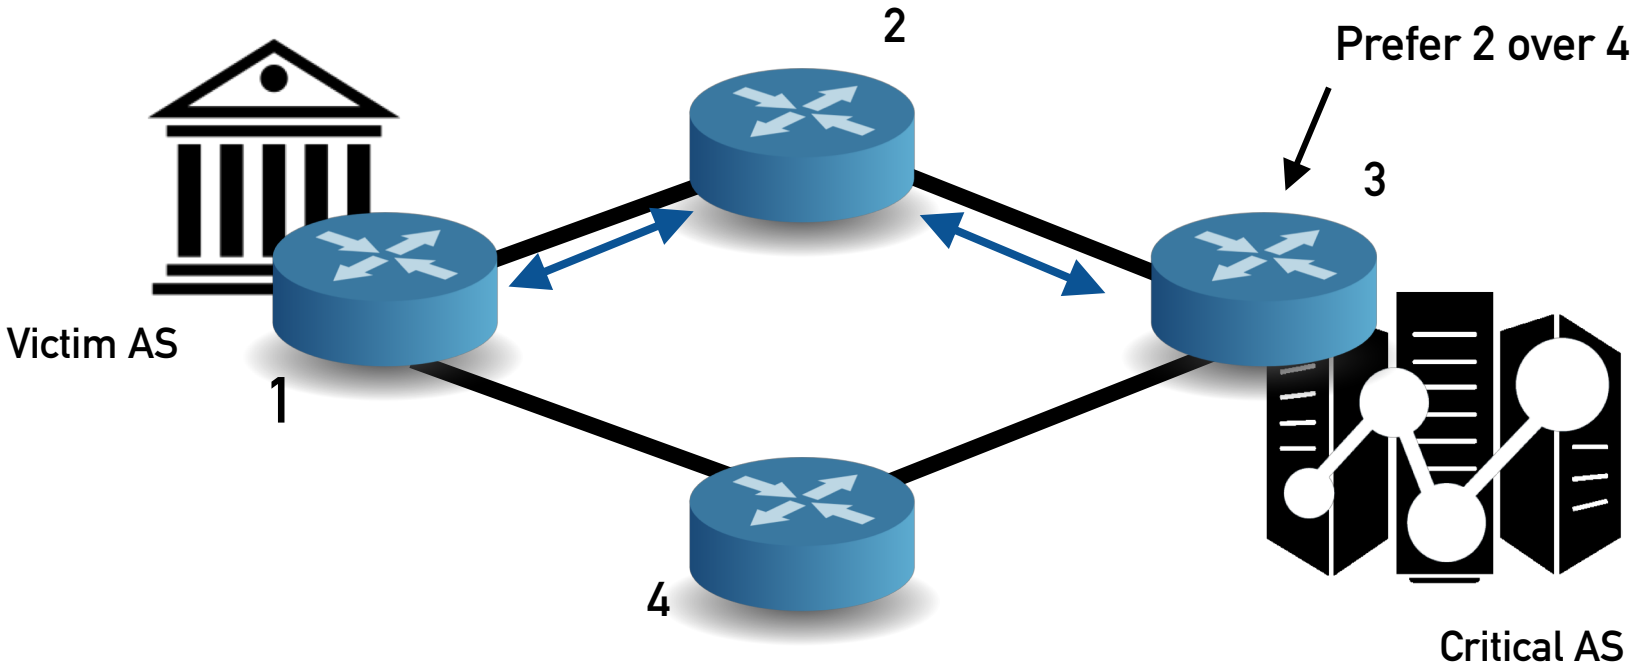

BGP Poisoning

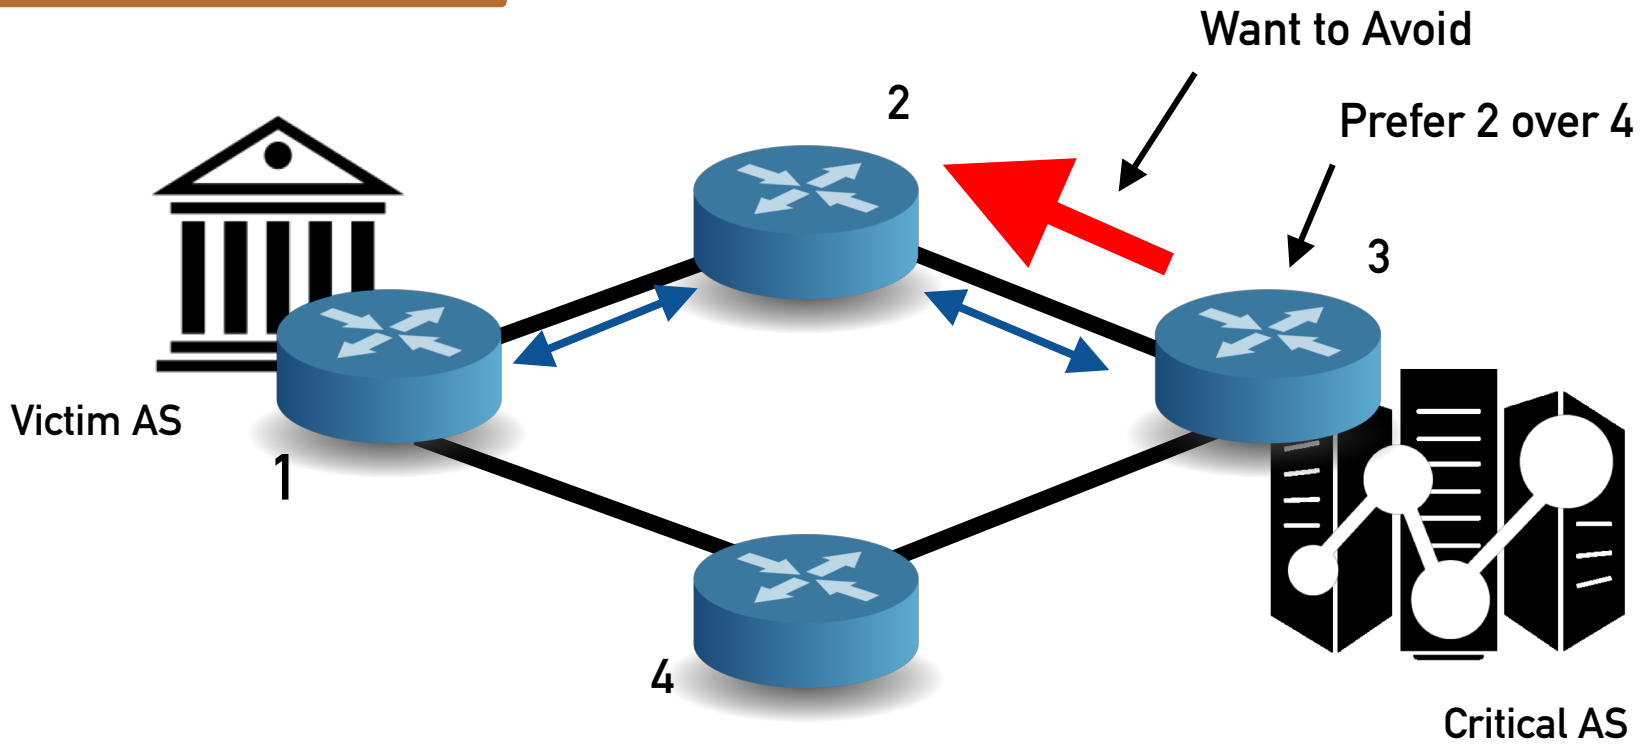

BGP Poisoning

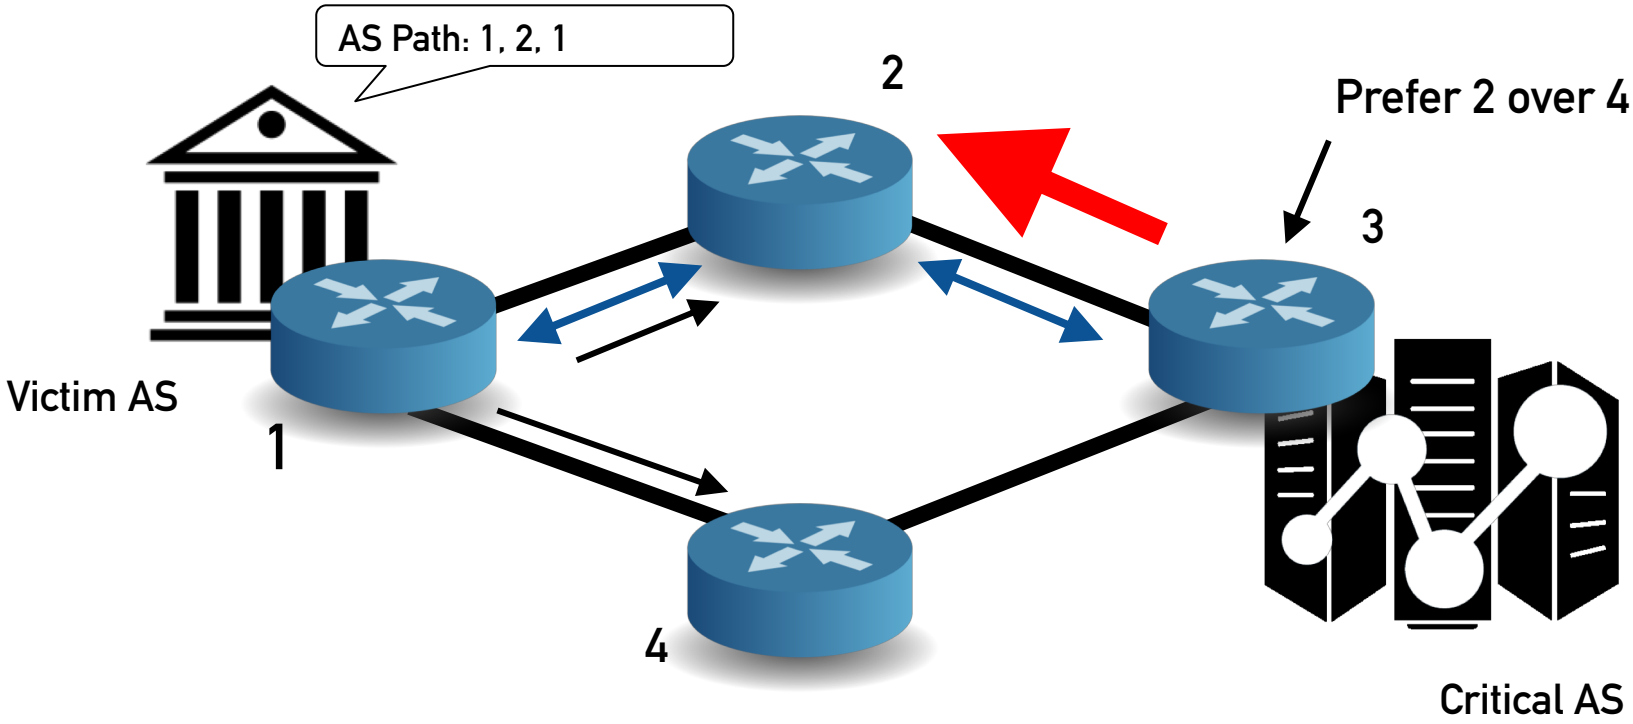

BGP Poisoning

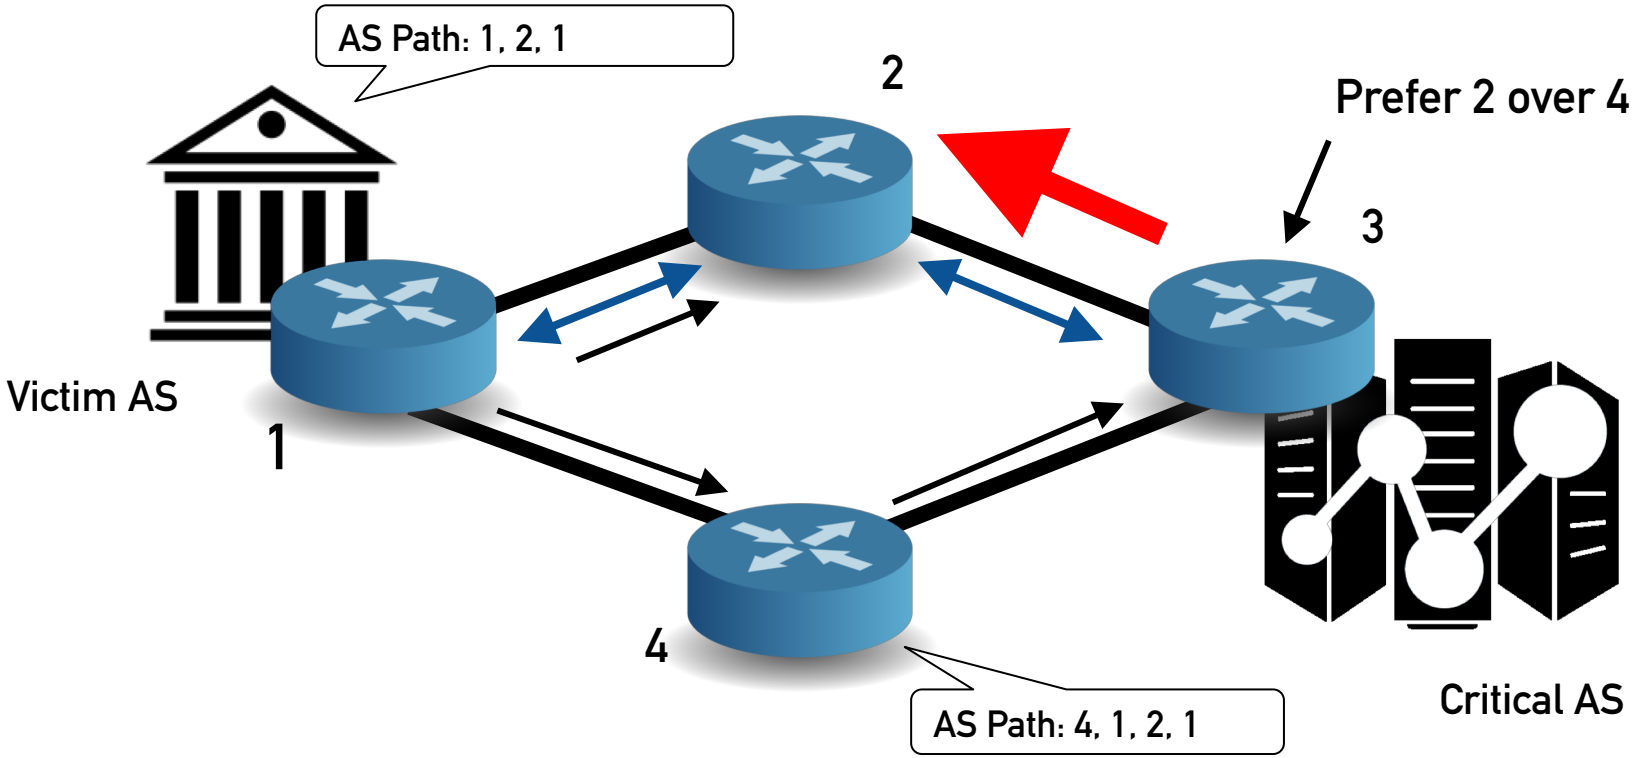

BGP Poisoning

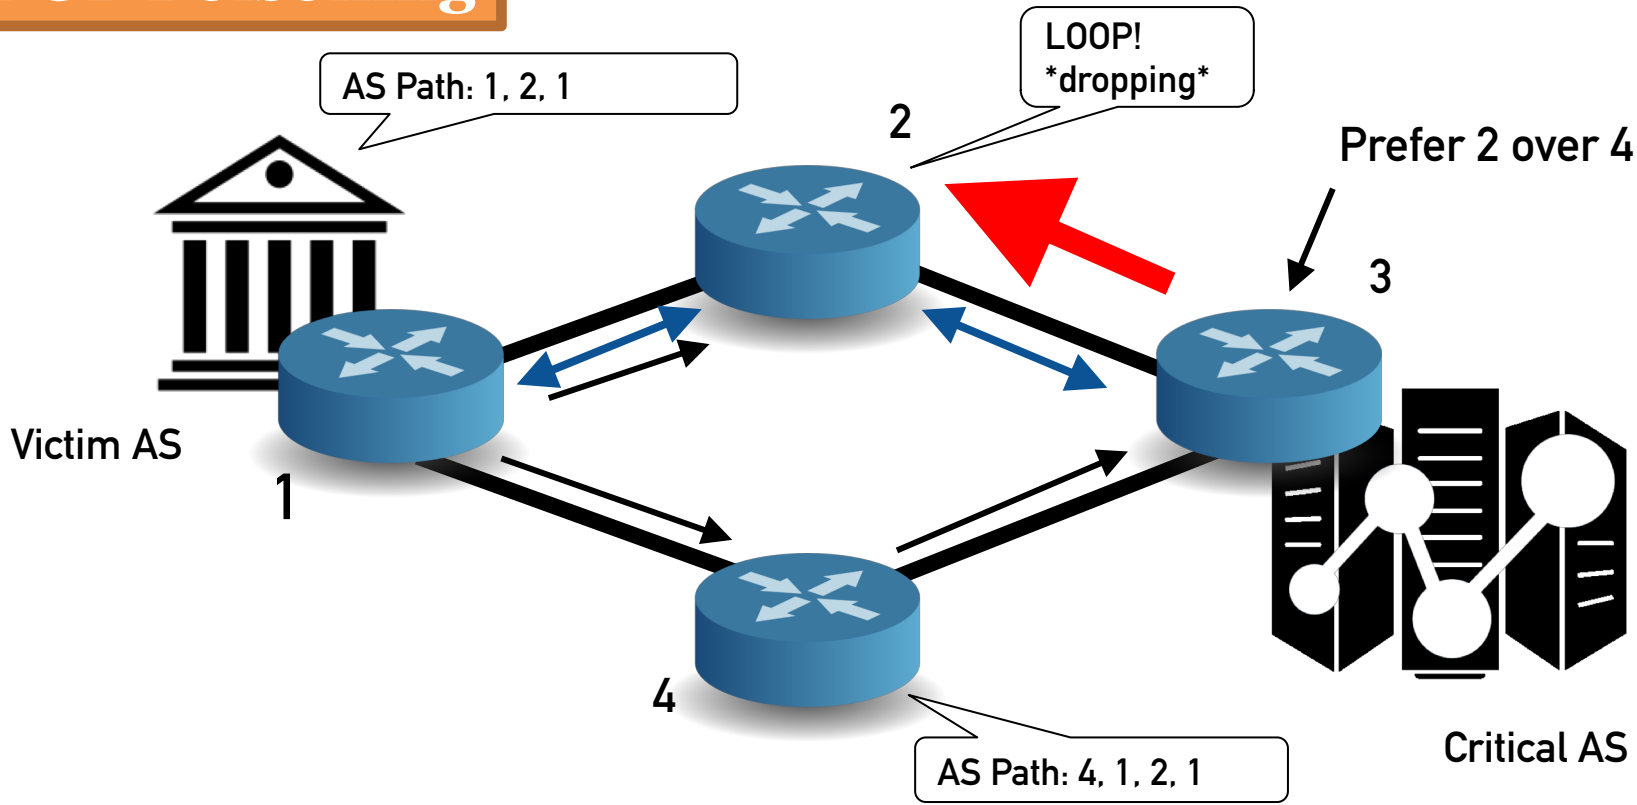

BGP Poisoning

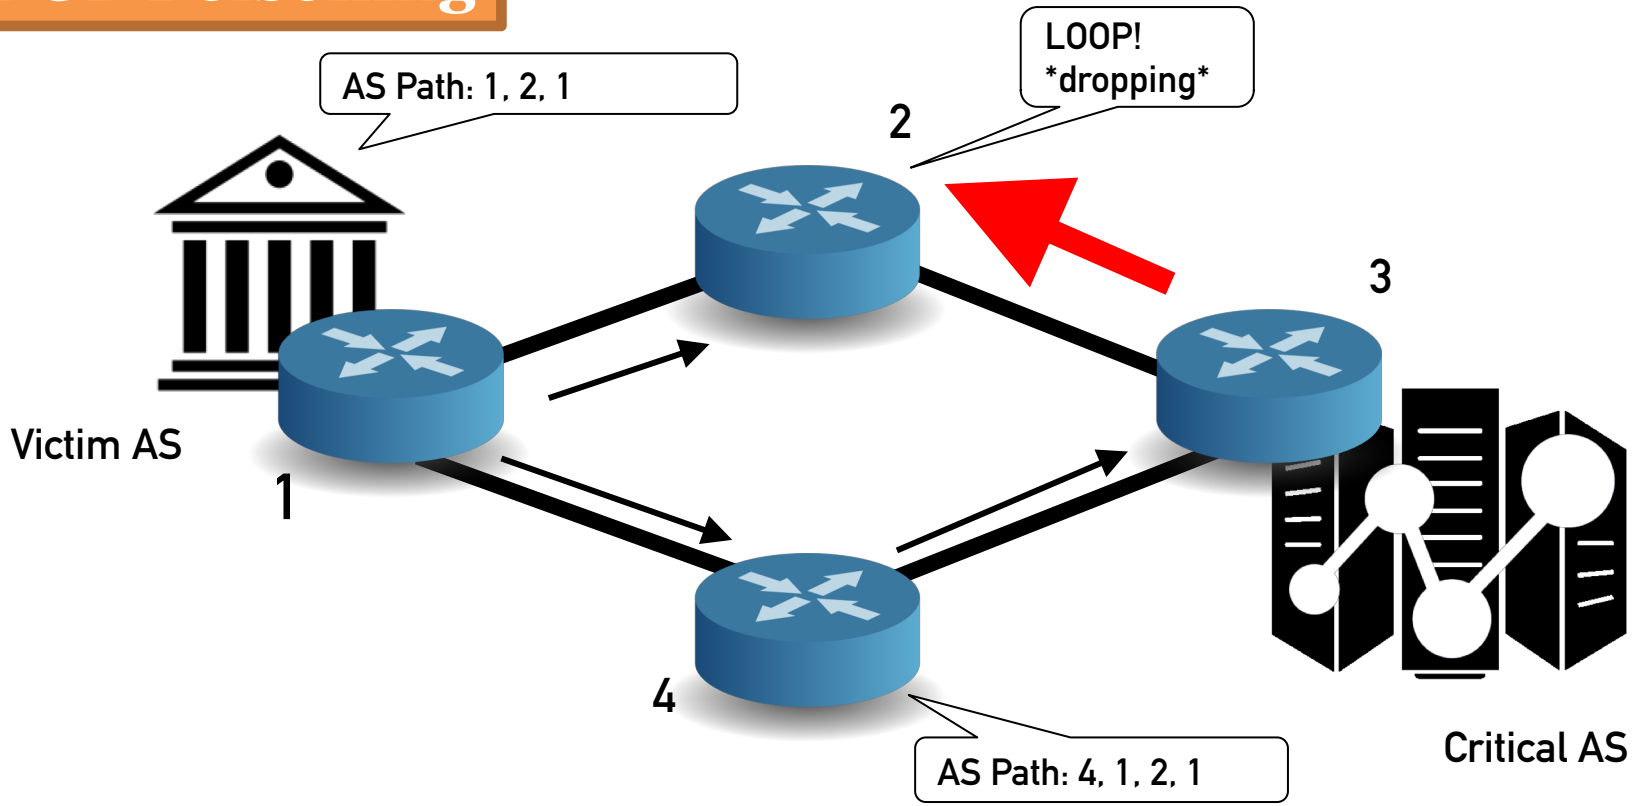

BGP Poisoning

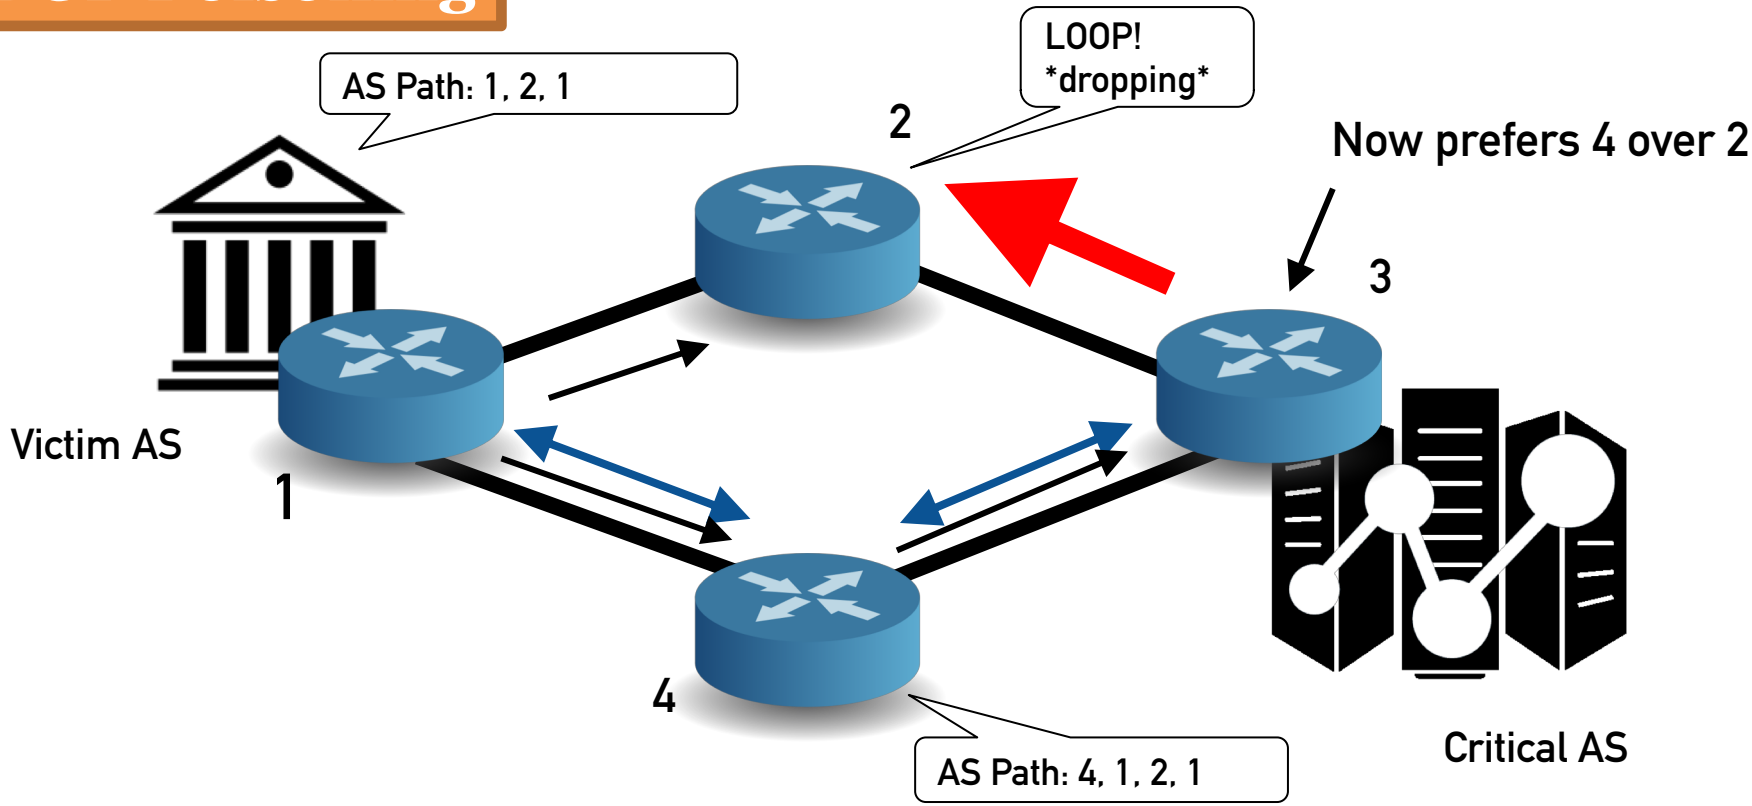

BGP Poisoning

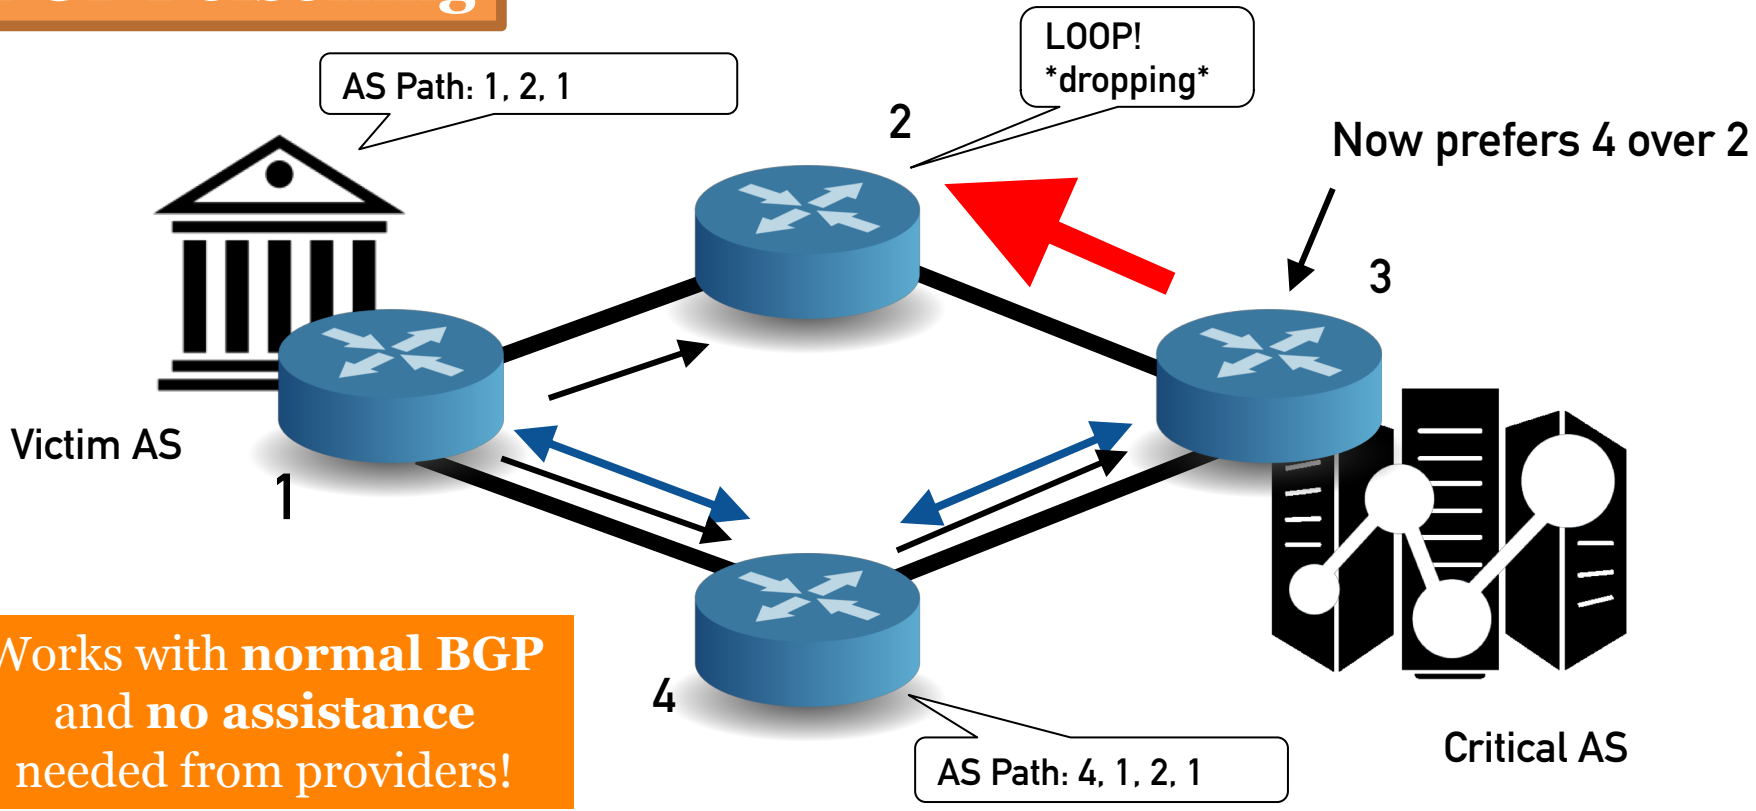

Works with **normal BGP** and **no assistance** needed from providers!

IS IT FEASIBLE?

How well can we re-route?

Metric	Result
Cases of Unsuccessful Return Path Steering	428
Cases of Successful Return Path Steering	1,460
Overall Unique Detour ASes	1369
Average Unique Detour Paths Per ATLAS AS	2.25
Average Unique Detour ASes Per ATLAS AS	6.45
Max Unique Detour Paths	19
Max Unique Detour ASes	26
Avg. Poisons Needed vs. Avg. Detour ASes	2.03/6.45
Unique Detour ASes vs. Unique Poisons Needed	1369/468

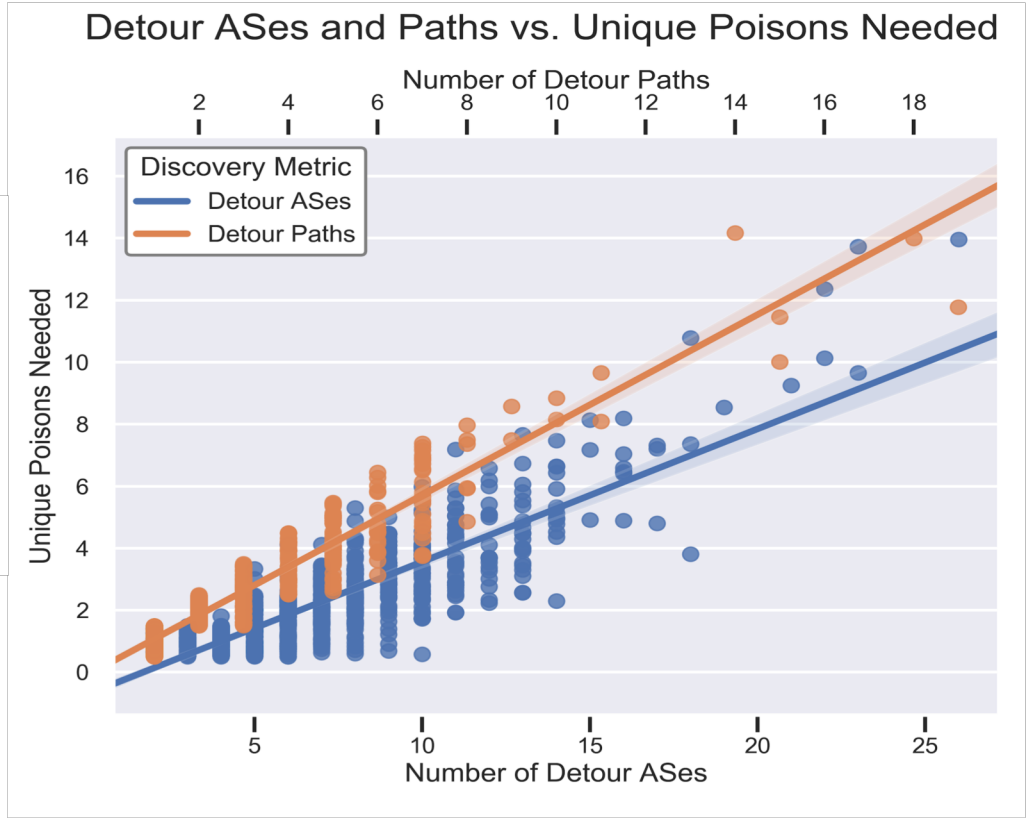

How performant are FRRP paths?

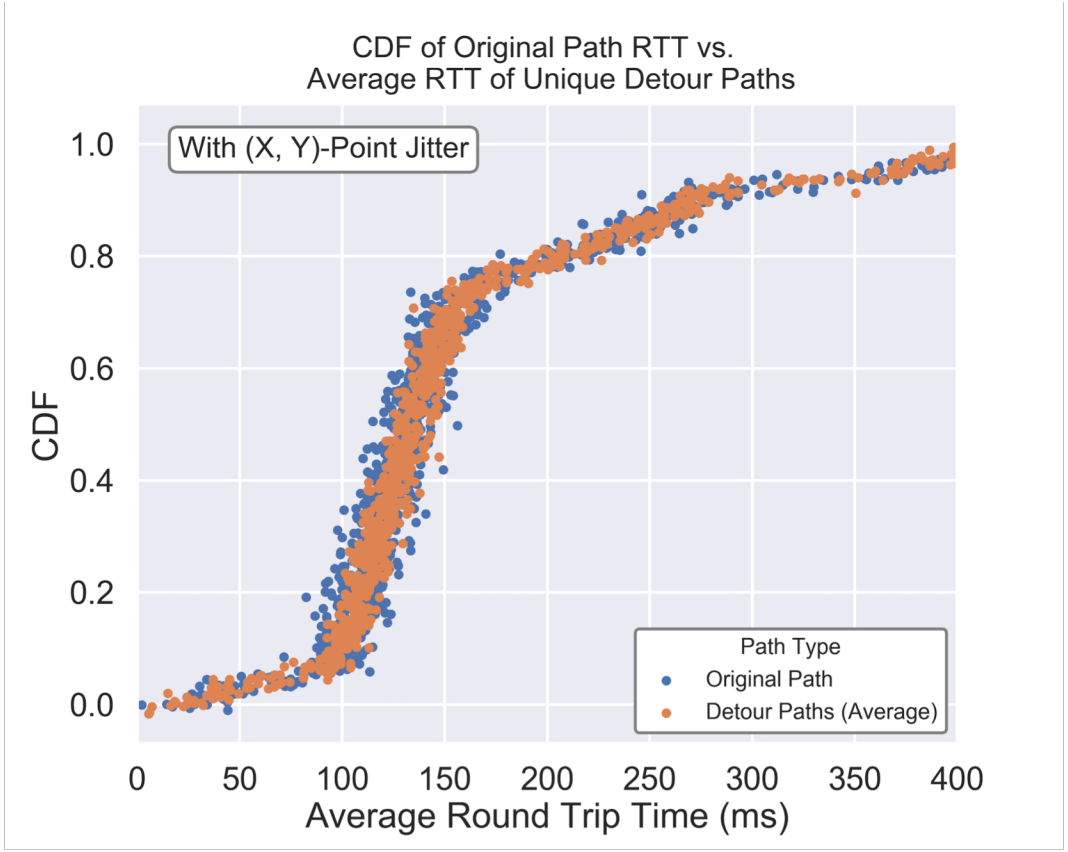

Emulation of BGP Poisoning vs. Practice

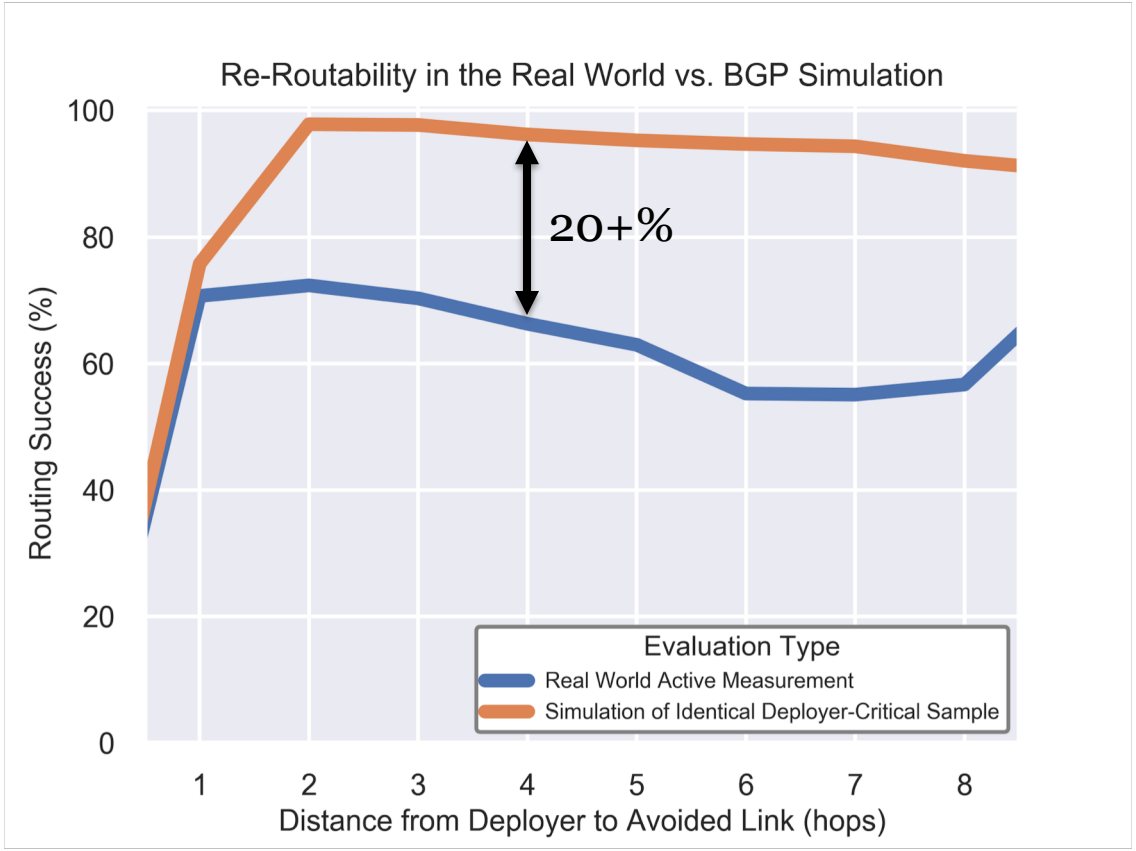

Graph-Theoretic Analysis of Return Paths

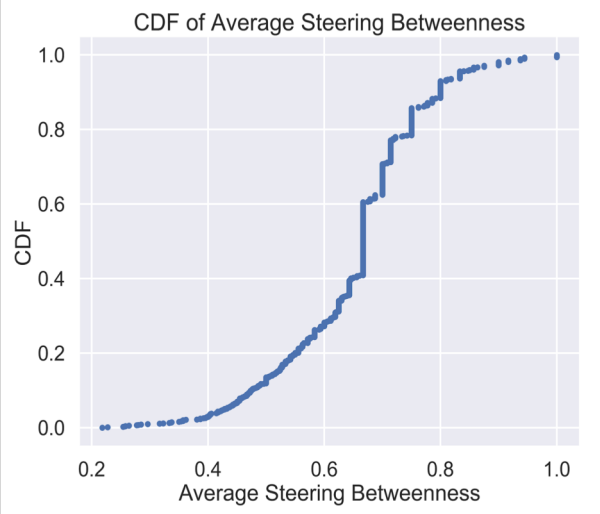

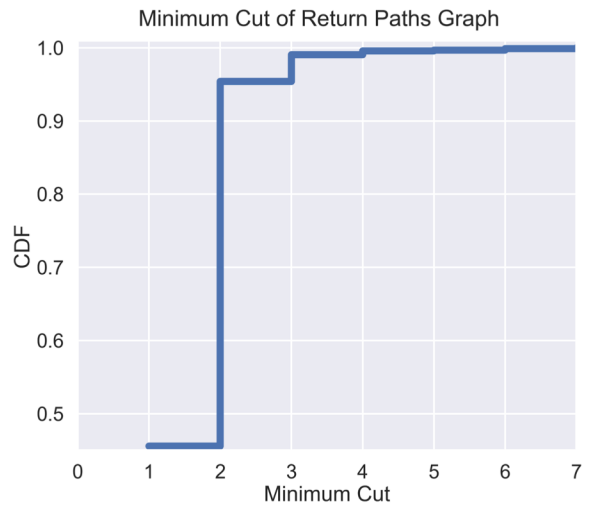

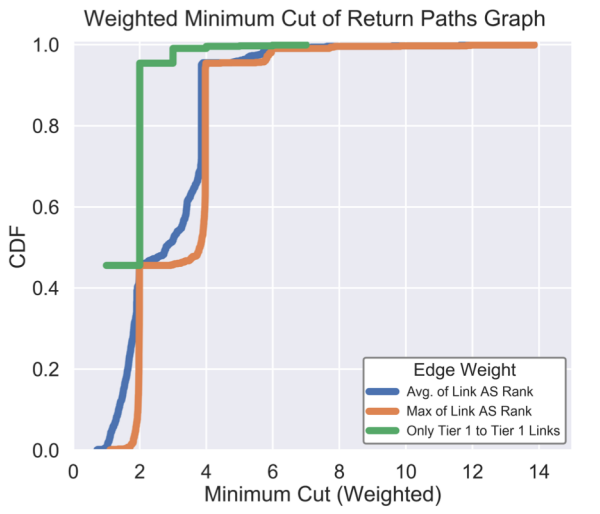

- Avg. Betweenness of 0.667
- Paths are not completely identical
- There is *some* diversity, but bottlenecks exist

- Low min. cut means bottlenecks that Nyx/RAD cannot avoid
- For 90% of links, a bottleneck of at most 2 links occurs

- Tier 1 ASes with inf. weight → bottlenecks **not** result of single unavoidable provider
- Within unweighted min cut → widely differing barriers to cut based on bandwidth

WHO CAN RE-ROUTE?

How well can we predict success with FRRP?

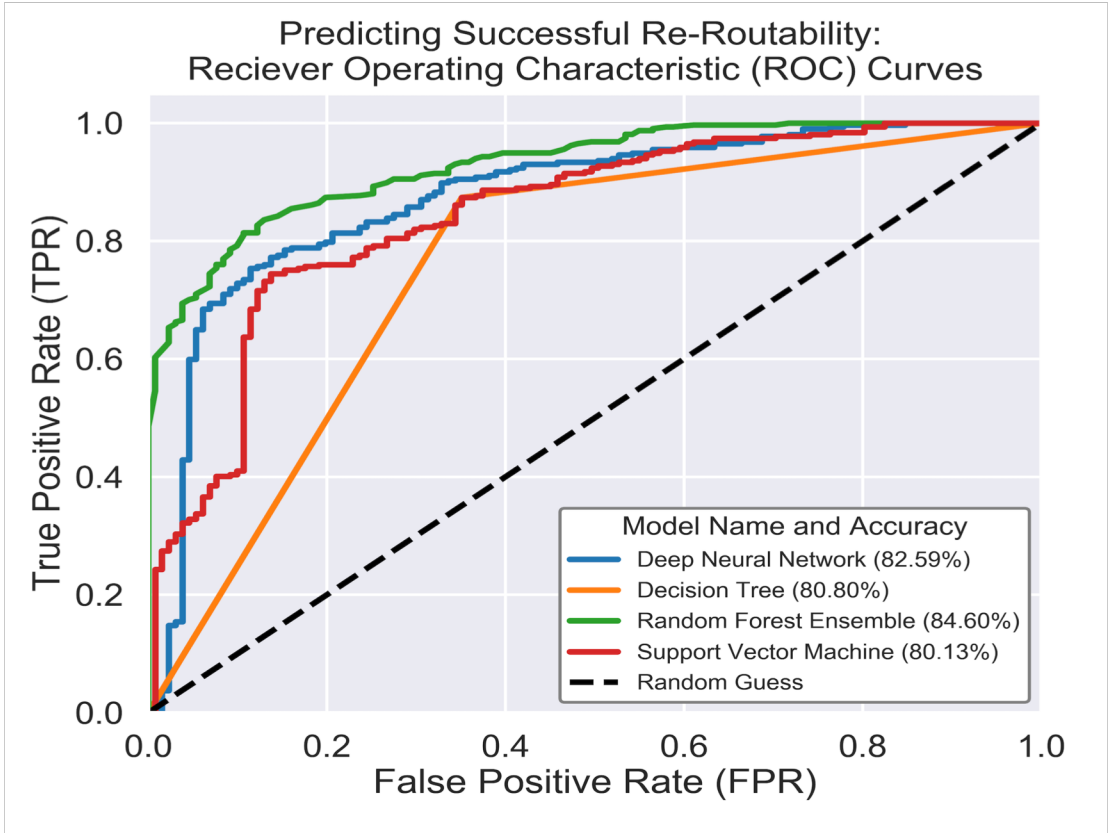

What link and AS properties are important for FRRP?

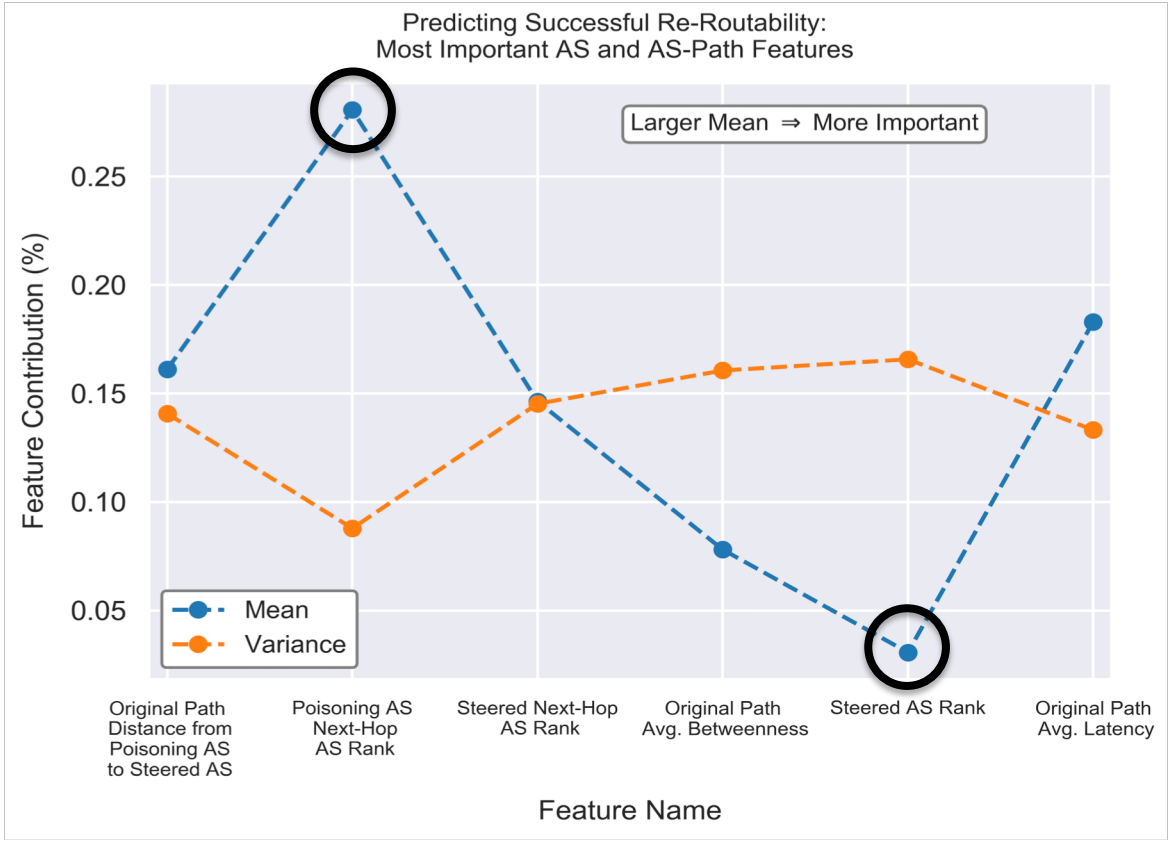

A Deeper Look at the Most Important Feature

Poisoning AS Next-Hop AS Rank

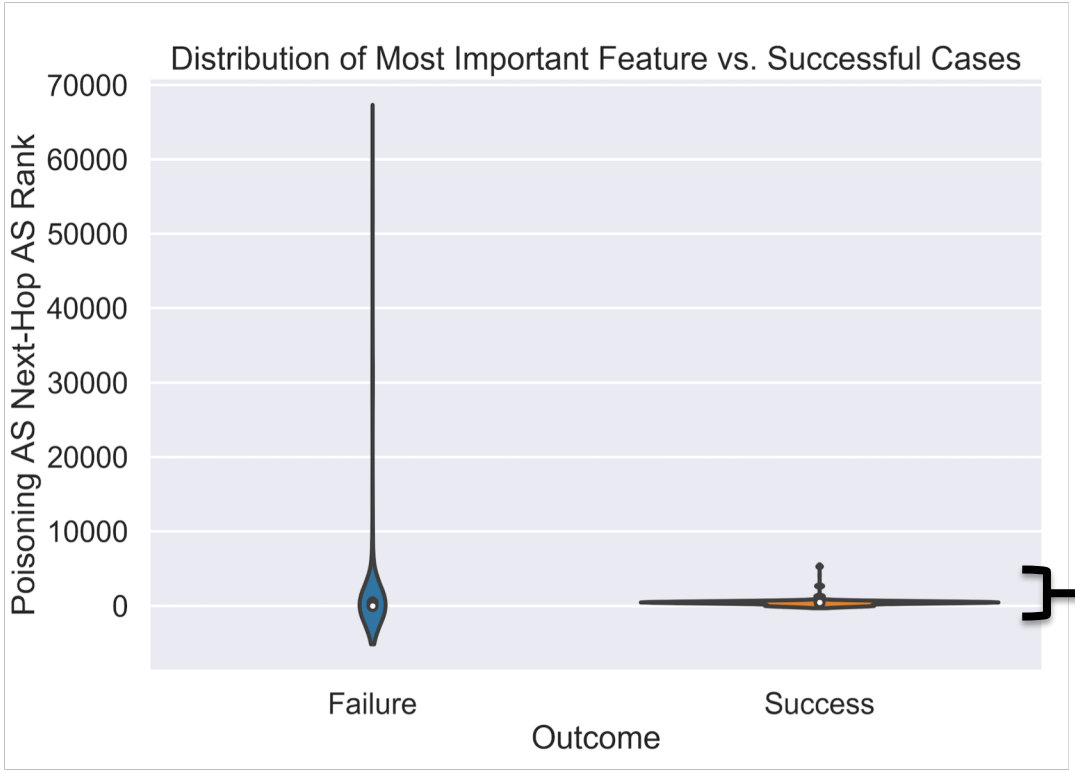

High Rank Matters

HOW MUCH CAN WE POISON?

How long can poisoned paths be?

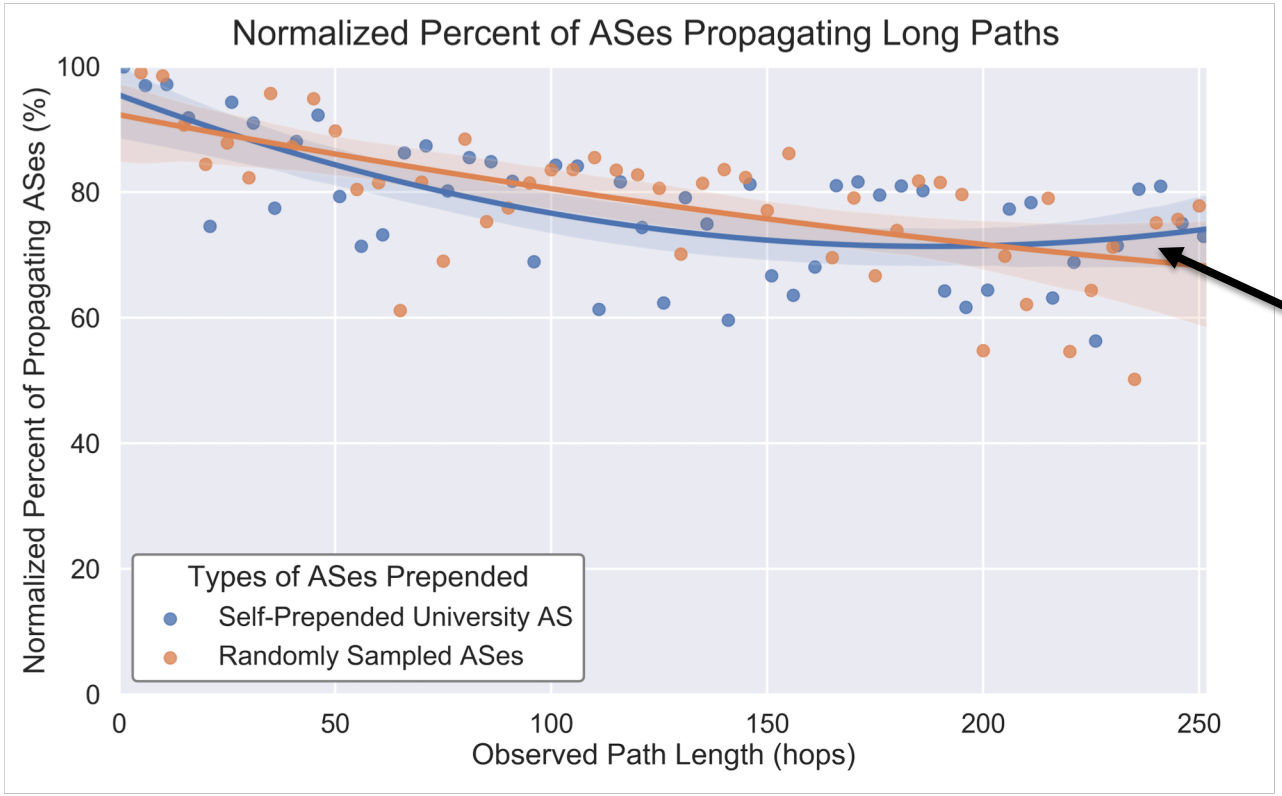

~75%

WHO FILTERS POISONS?

Filtering by Large ISPs

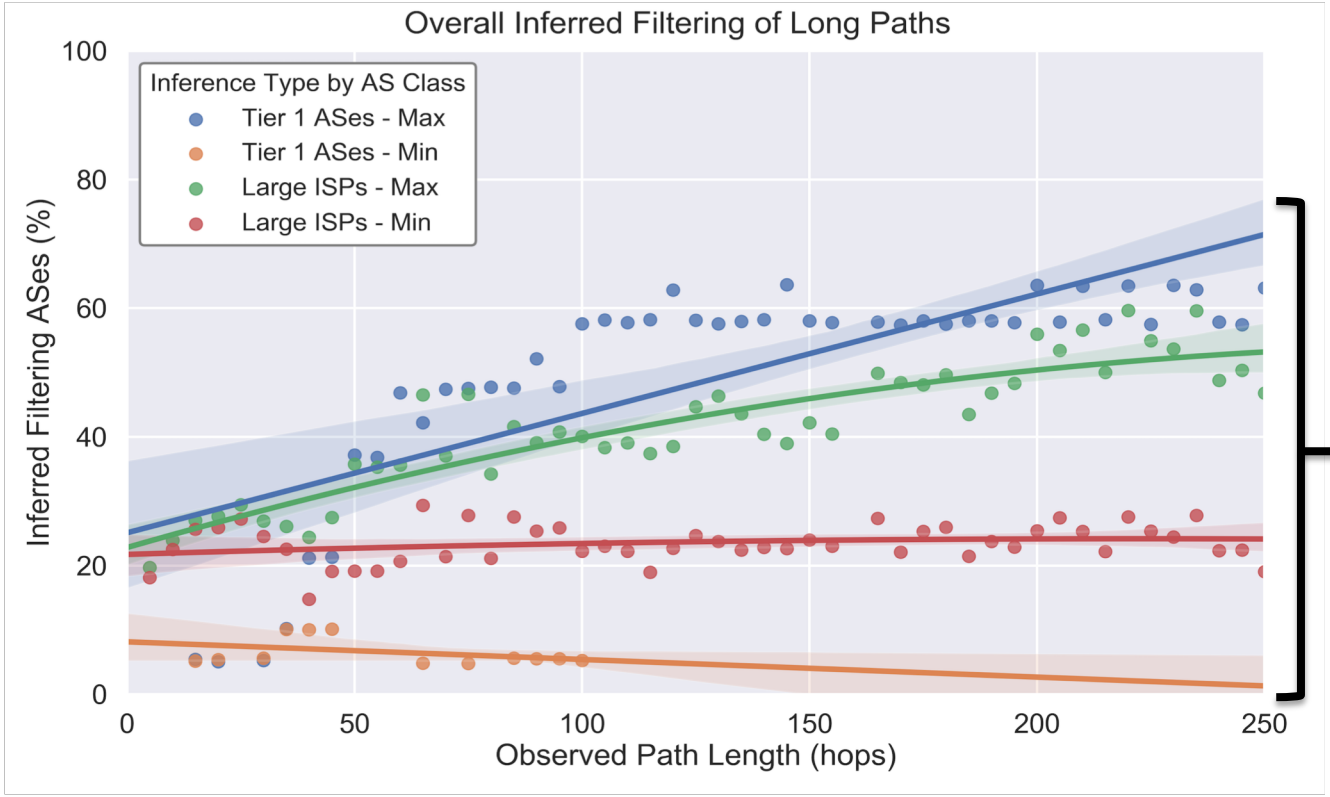

Filtering by Small ISPs + Stubs

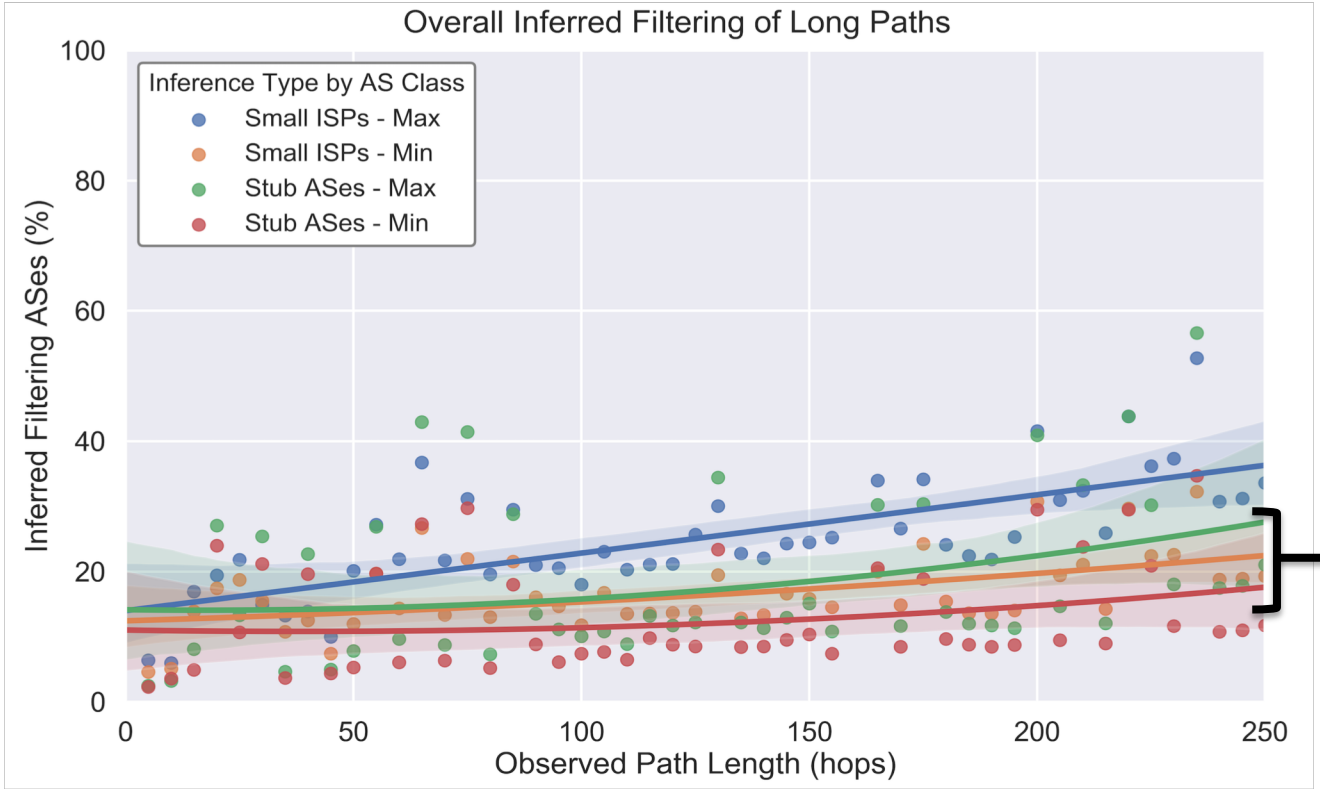

Small window

Do the Policy Leaders "*Walk the Walk*"?

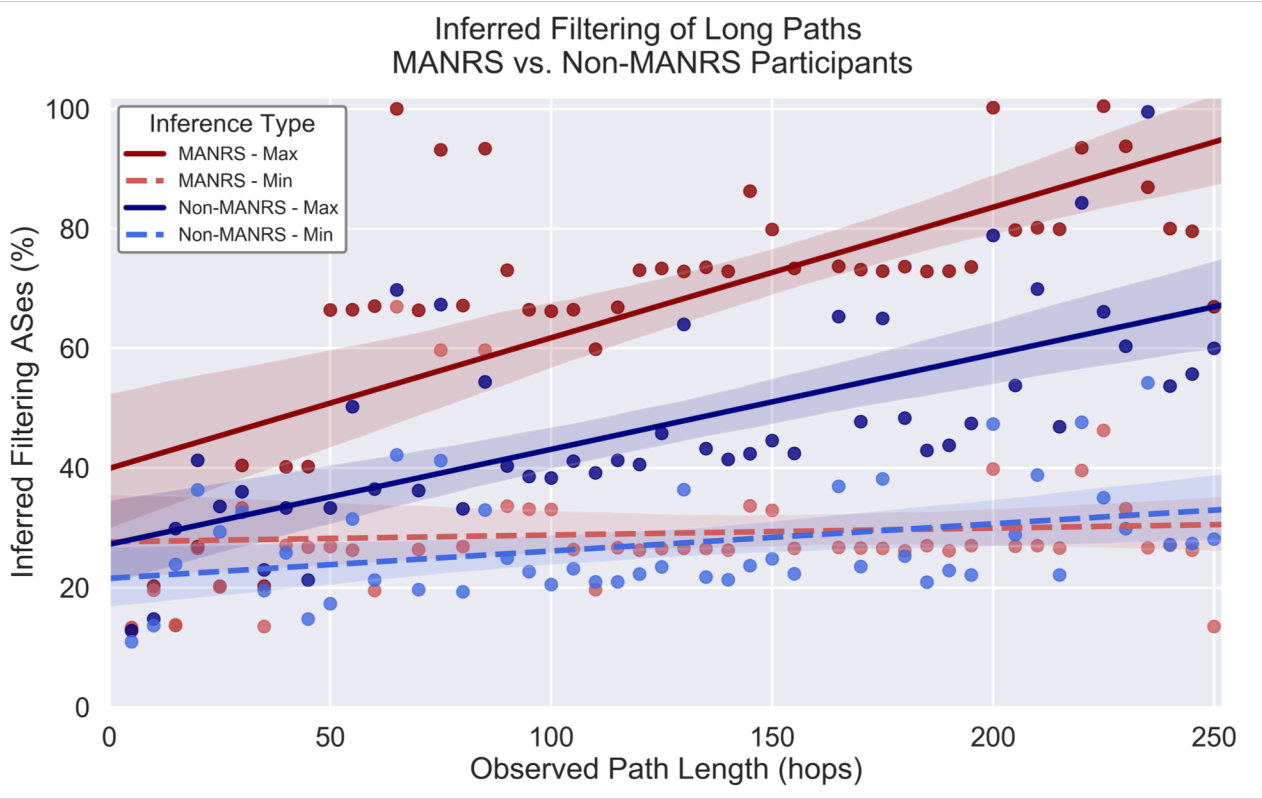

MANRS

"Mutually Agreed Norms for Routing Security"

Selected Participants (total=146):

- CenturyLink
- Charter
- Cogent
- Google
- Indiana U.
- …

Does AS-Degree of the Poisoned AS affect Filtering?

$Origin_{AS}$ $HighDegree_{AS}$ $Origin_{AS}$

...(in increments of 5)...

$Origin_{AS}$ $SmallDegree_{AS}$ $Origin_{AS}$

Rank by Degree	ASN and Name	Degree	Number of Customers	Registered Country by ASN	Normalized Propagation Percentage
1	6939 - Hurricane Electric	7064	1202	United States	11.9%
2	174 - Cogent	5352	5272	United States	11.6%
3	3356 - Level 3	4980	4898	United States	11.6%
4	24482 - SG.GS	3382	24	Singapore	96.1%
5	3549 - Level 3 GBLX	2538	2446	Unites States	11.6%
6	7018 - AT&T	2373	2330	United States	0.05%
7	58511 - Anycast	2351	13	Australia	60.1%
8	49605 - IVO	2193	11	Italy	66.7%
9	8492 - OBIT Ltd.	2153	46	Russia	71.4%
10	8220 - COLT Tech. Grp.	2143	716	United Kingdom	78.2%

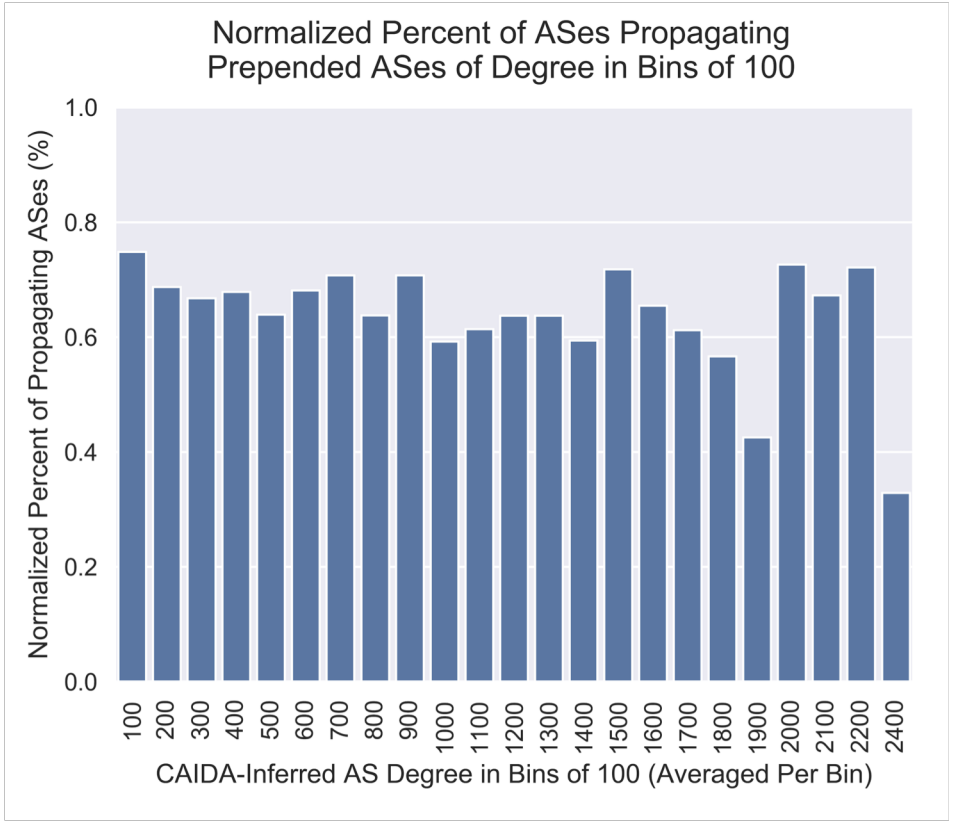

DEFAULT ROUTES AND REACHIBILITY

(NOW VS. 2009)

Default Route Metrics

Measurement	Number of Instances
Fraction of Total Samples with Only 1 Provider (not multi-homed)	28.7% (419 / 1,460 total samples)
Fraction of Total Multi-Homed Samples with Default Routes	48.6% (506 / 1,041 multi-homed samples)
Fraction of Transit ASes with Default Routes	26.8% (196 / 731 total Transit ASes)
Fraction of Stub/Edge/Fringe ASes with Default Routes	36.7% (310 / 845 total Fringe ASes)

Comparison

2009*: 77% of Stubs had default routes (out of 24,224 **with ping**)
2018: 36.7% of Stubs had default routes (out of 845 **with traceroute**)

*Bush et al. Internet Optometry, IMC 2009

Reachibility of /25 vs. /24

Prefix Length	Measurement	Findings	Timespan of Measurement
/25	BGP Observability	Seen at 21/37 (56.7%) collectors	96 hours of collection
/25	Traceroute Reachability	31% reached /25 prefix on average	7 hours; 5,000 distinct traceroutes every 1 hour
/24	BGP Observability	Seen at 34/37 (91.8%) collectors	96 hours of collection

Comparison

2009*: 1% of BGP Monitors Saw (11/615), 5% Data-Plane Reachability
2018: 50% of BGP Monitors Saw (21/37), 31% Data-Plane Reachability

*Bush et al. Internet Optometry, IMC 2009

Where do we go from here?

- **BGP poisoning** can provide helpful functionality
- Allows exertion of *unconventional behavior* with a *conventional protocol*
- **Open Questions for AIMS:**
  - *Deployment/Usage:* Where? For what?
  - *Integration*: CAIDA systems? NANOG/RIPE/etc.? MANRS?
  - *Collaboration*: Always interested in extending to new use cases/measurements.

University of Tennessee
VOLSEC
— COMPUTER SECURITY LAB —

Jared M. Smith

Twitter	jaredthecoder
Email	**jms@vols.utk.edu**
Web	volsec.org