Expanding BGP data horizons

Cristel Pelsser

UCLouvain

www.enmieux.be

Cofinancé par l'Union européenne

Wallonie

Cyber Excellence by Cyberwal

Focus on the inter-domain routing protocol BGP

The Internet is composed of Autonomous Systems (AS): one or more networks under the control of a single entity.

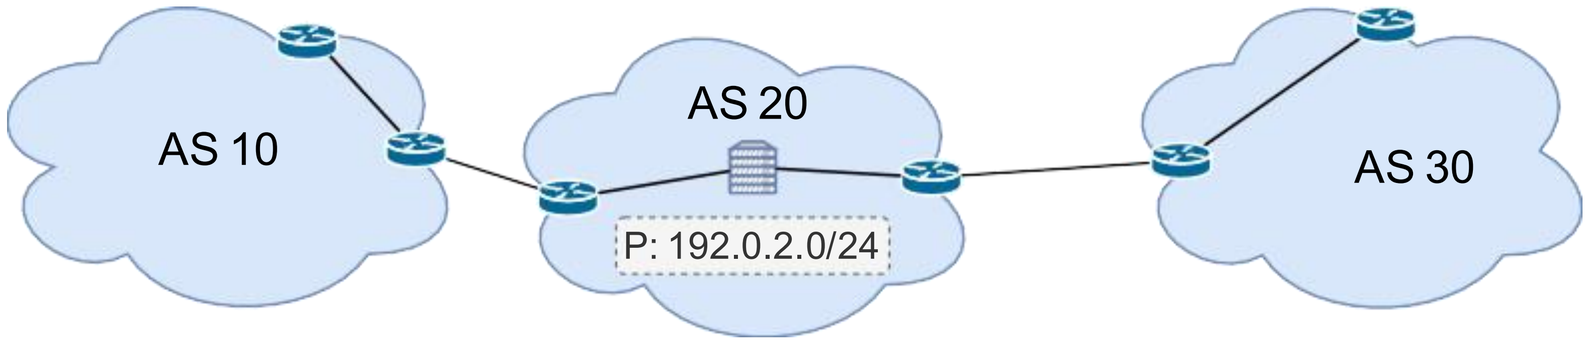

The Internet is composed of Autonomous Systems (AS): one or more networks under the control of a single entity.

Prefixes of the AS are advertised to the outside using BGP.

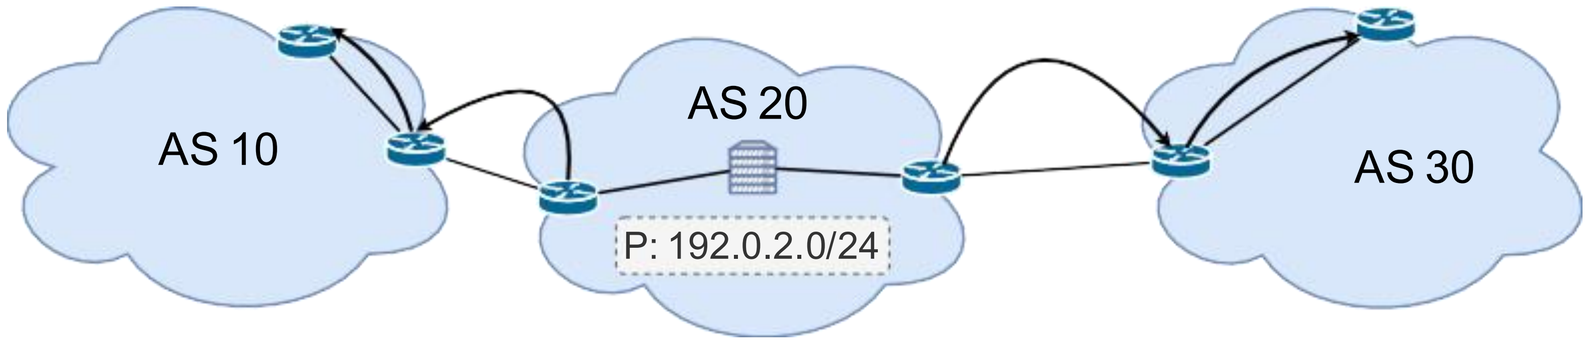

The Internet is composed of **Autonomous Systems** (AS): one or more networks under the control of a single entity.

Prefixes of the AS are advertised to the outside using BGP.

Traffic flows in the reverse direction.

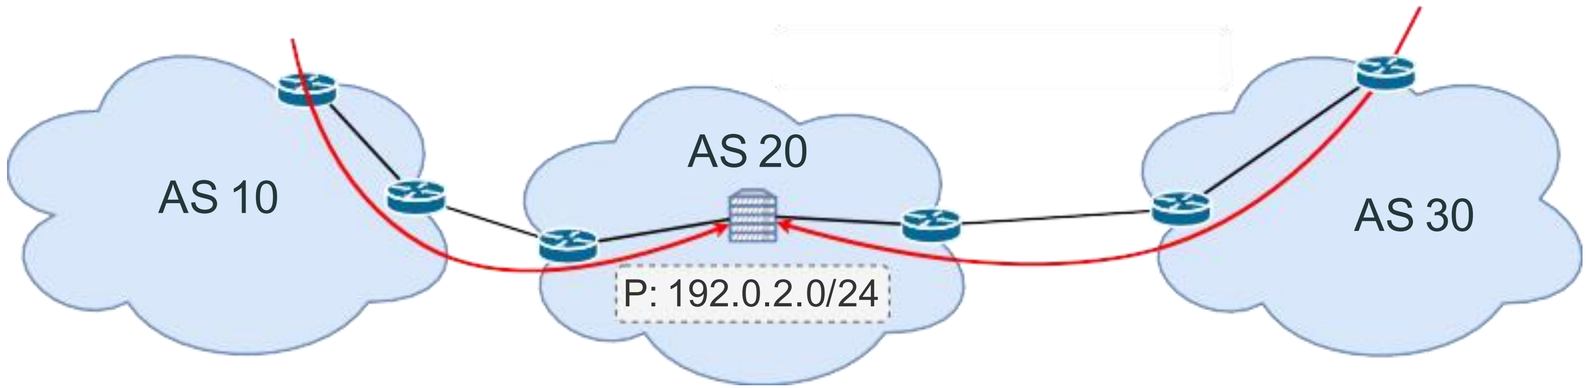

The Internet is complex and dynamic

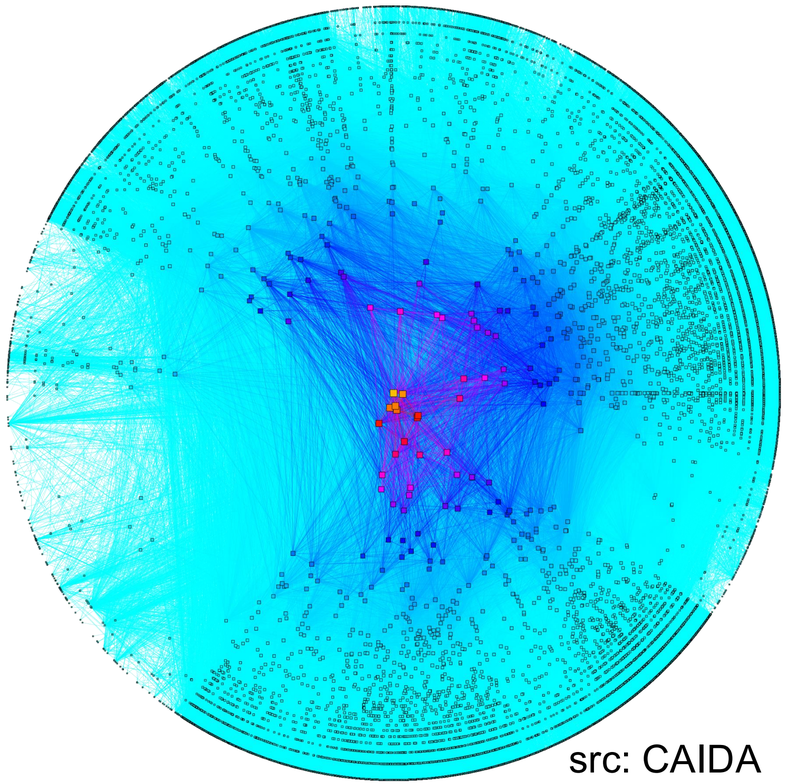

src: CAIDA

The internet is composed
of more than 75 000 ASes

The routes in the Internet are computed automatically thanks to the **BGP routing protocol**

Some vulnerabilities of BGP

Prefix hijacks

Blackjack attacks

BGP lies

BGP session injection

There is little to no security in the routing protocol used in the Internet

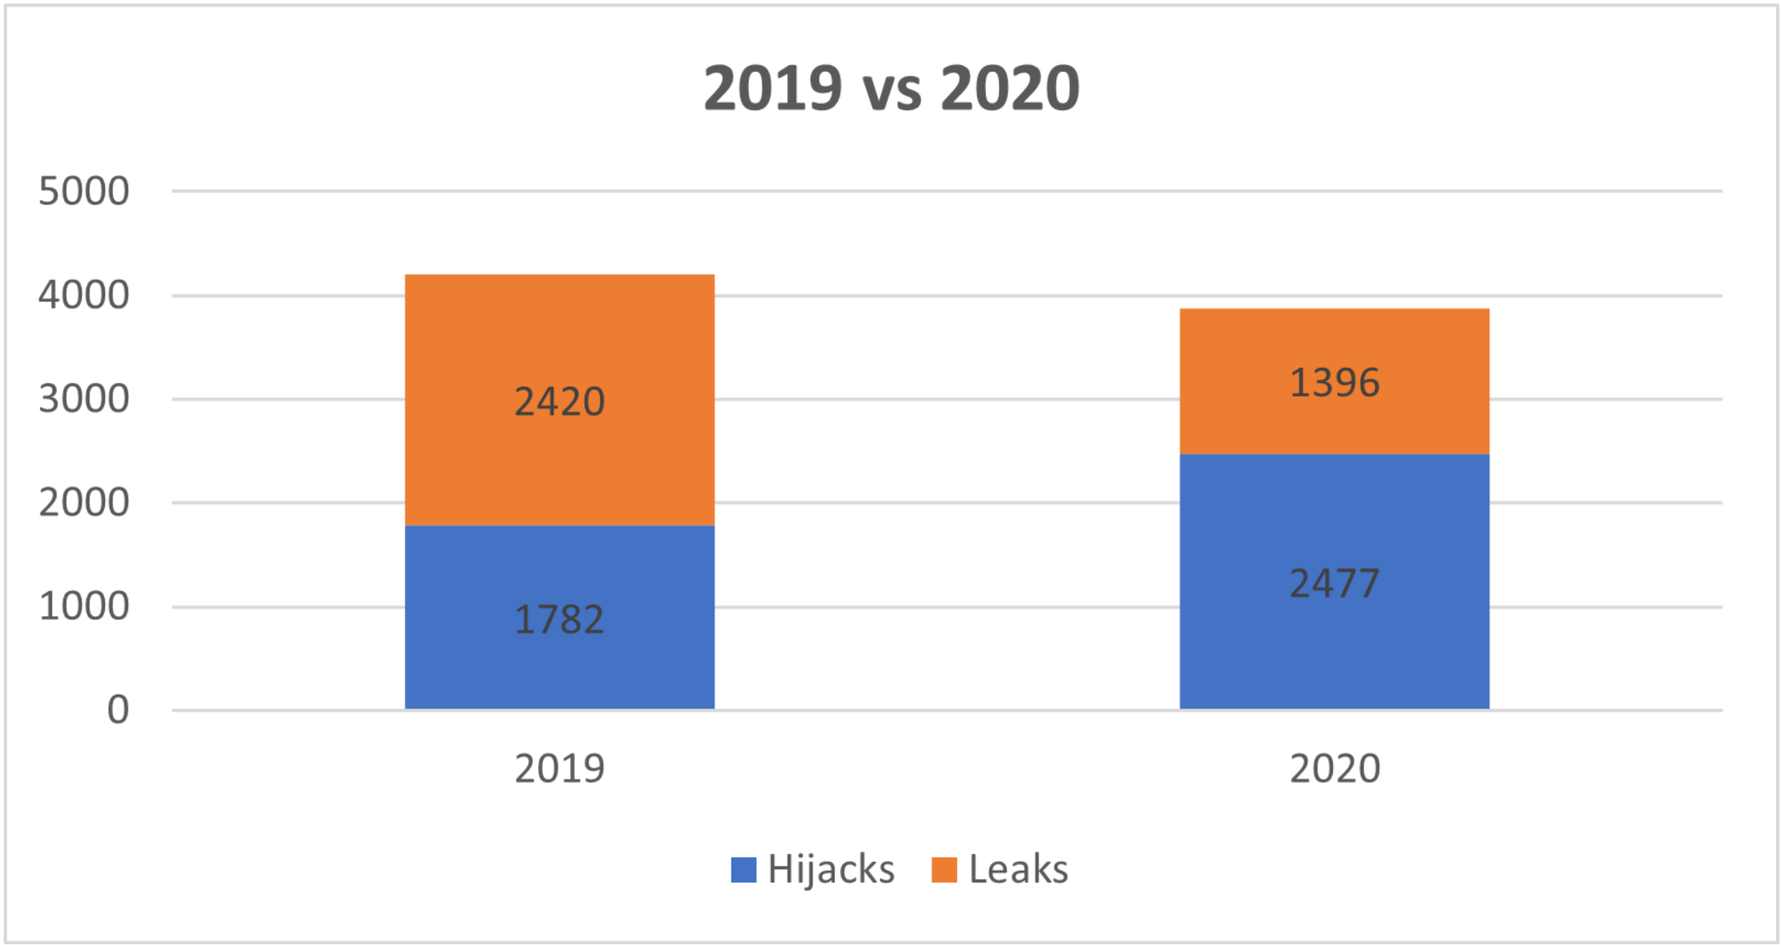

Source: https://www.manrs.org/2021/02/bgp-rpki-and-manrs-2020-in-review/

Route propagation on an example topology

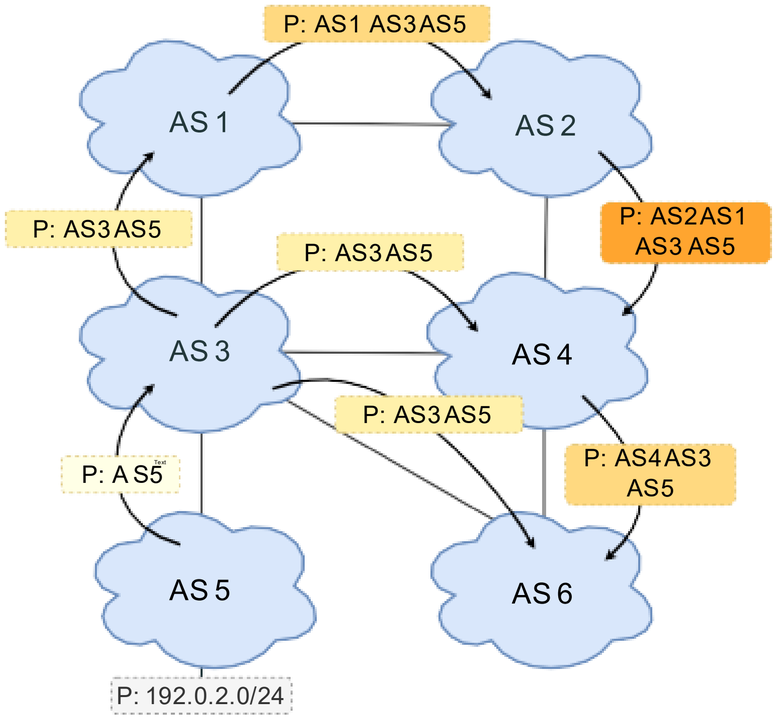

BGP update propagation

Route propagation on an example topology

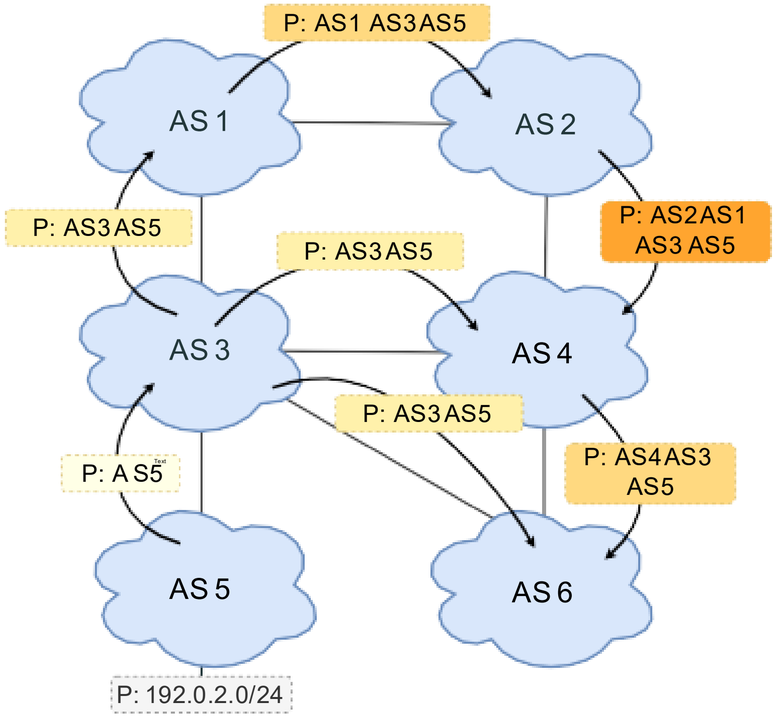

BGP update propagation

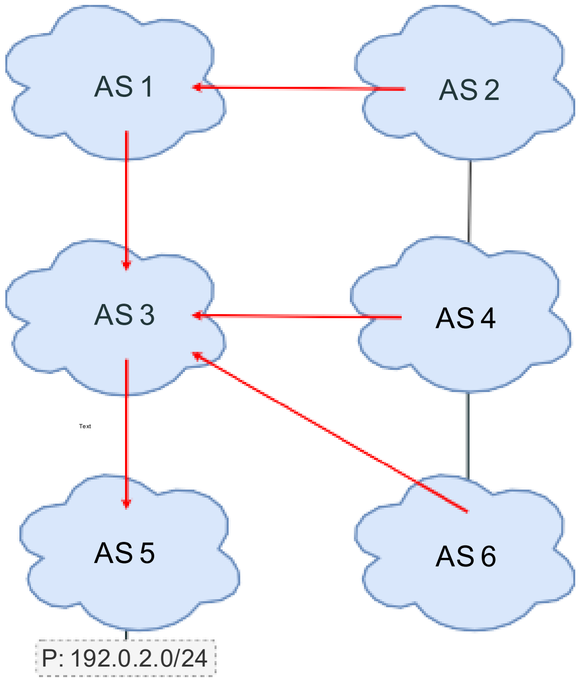

Traffic flow

Route propagation on an example topology

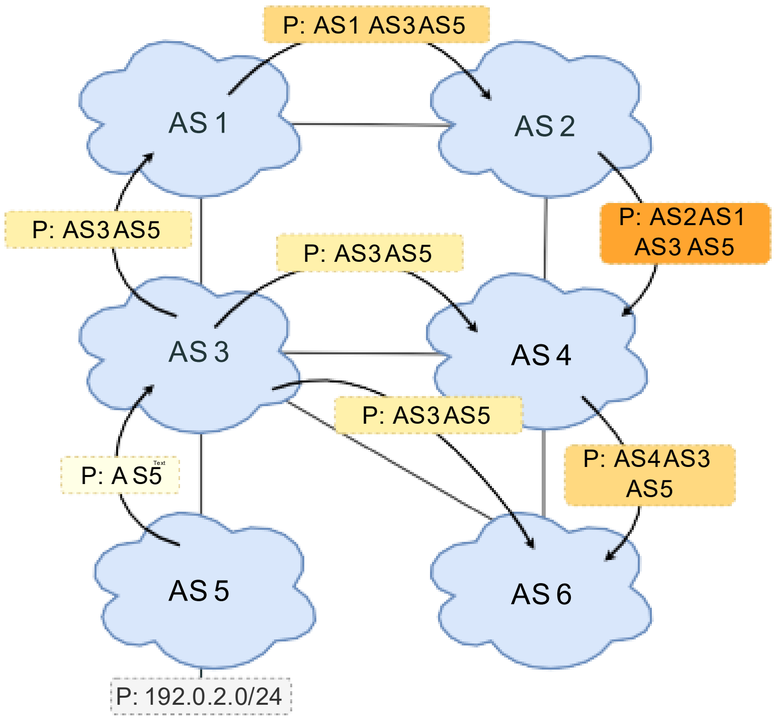

BGP update propagation

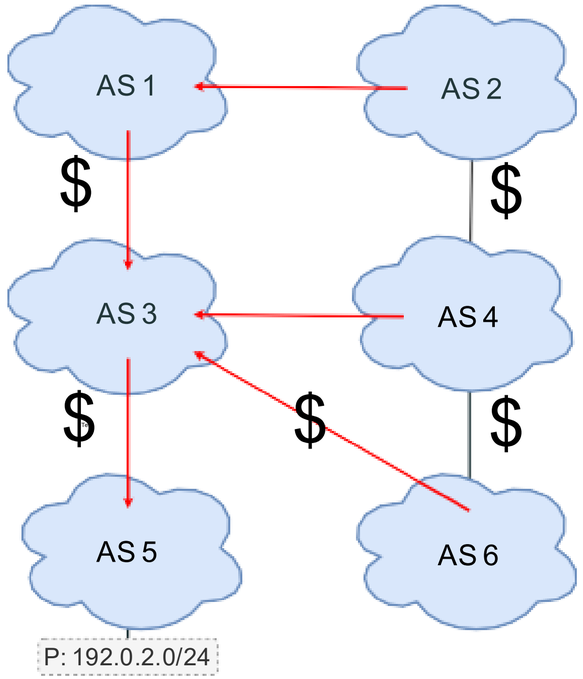

Traffic flow

BGP policies make AS2 not learn the path via AS4

Route propagation on an example topology

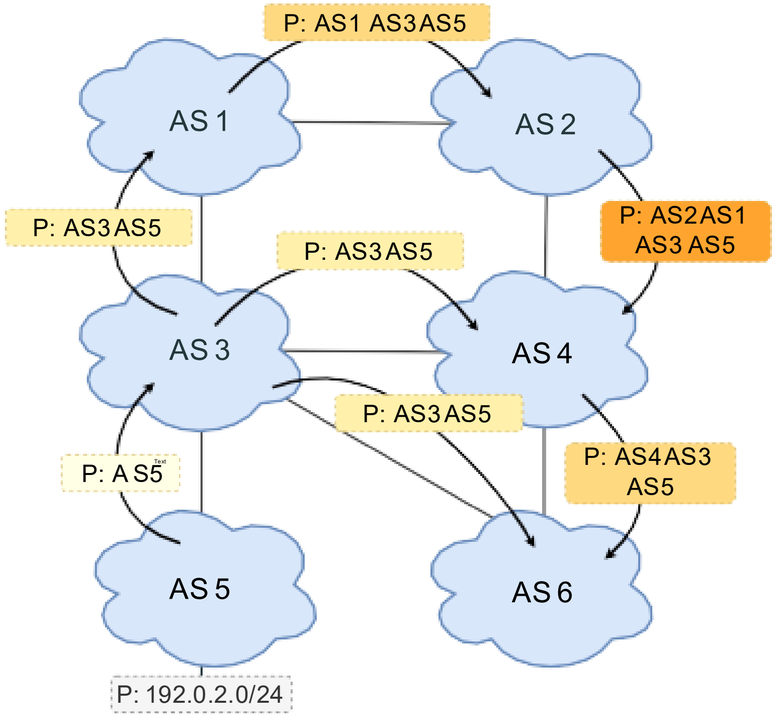

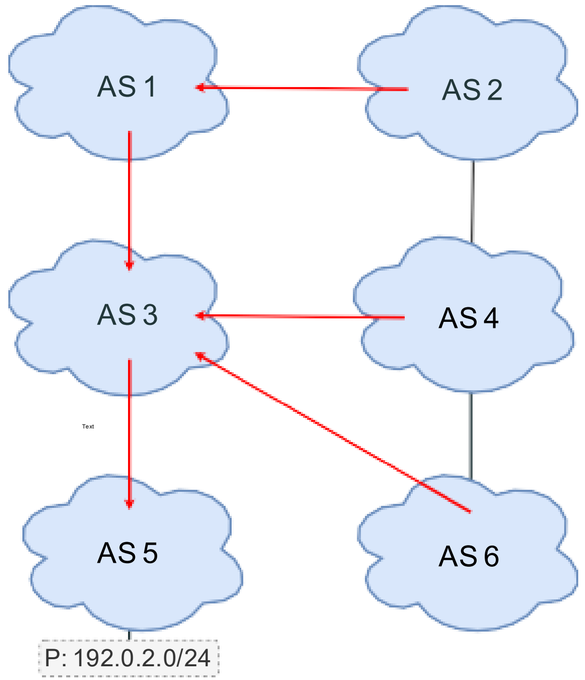

BGP update propagation

Traffic flow

BGP policies make AS2 not learn the path via AS4

In the next slides AS6 is the attacker

Hijack-0

Sermpezis 2018 (Artemis)

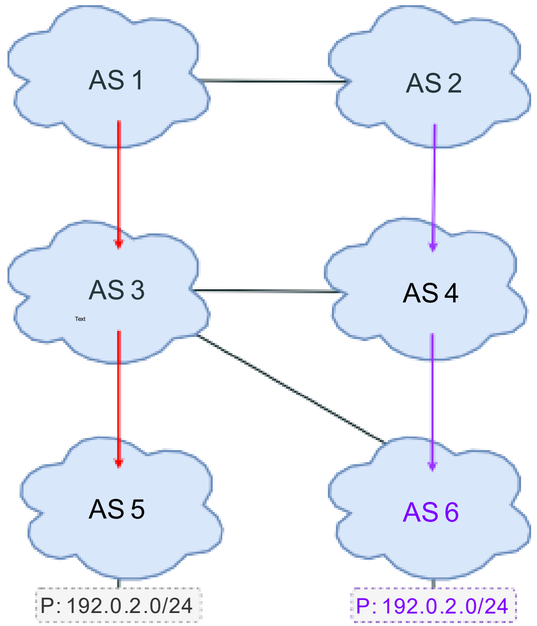

Hijack type-0

AS2 and AS4 traffic is de-routed to AS6 because the advertised path is shorter.

Hijacks can be used to divert traffic and gain inside knowledge

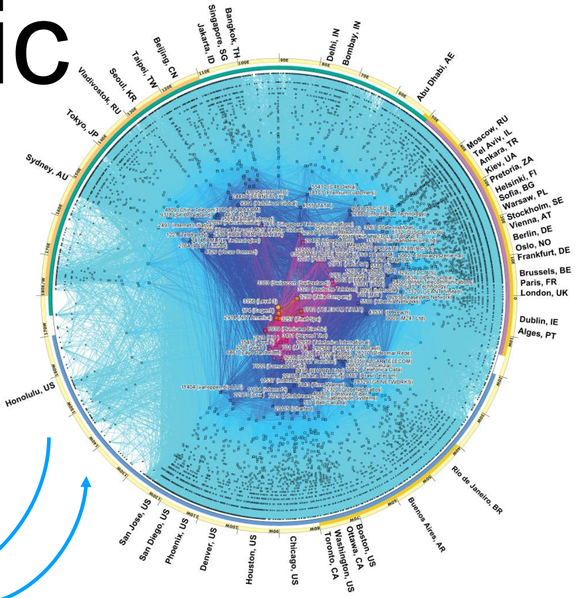

Hijacks can be used to divert traffic and gain inside knowledge

BGP hijacking

Fournisseur d'accès

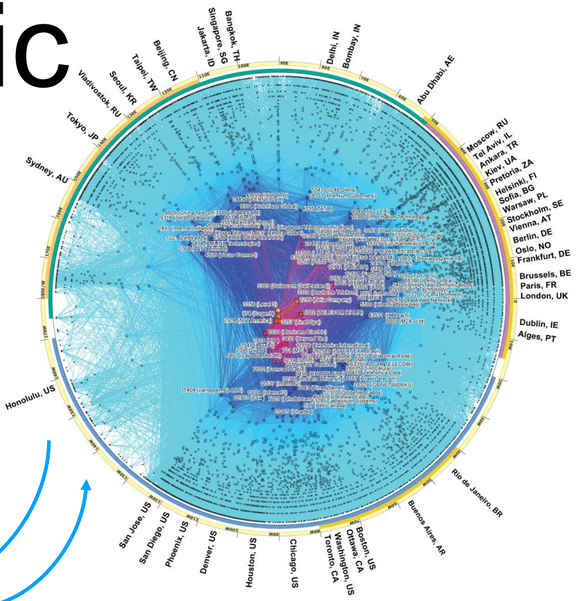

Hijacks can be used to divert traffic and gain inside knowledge

BGP hijacking

Informations Personnelles (Mail)

Fournisseur d'accès

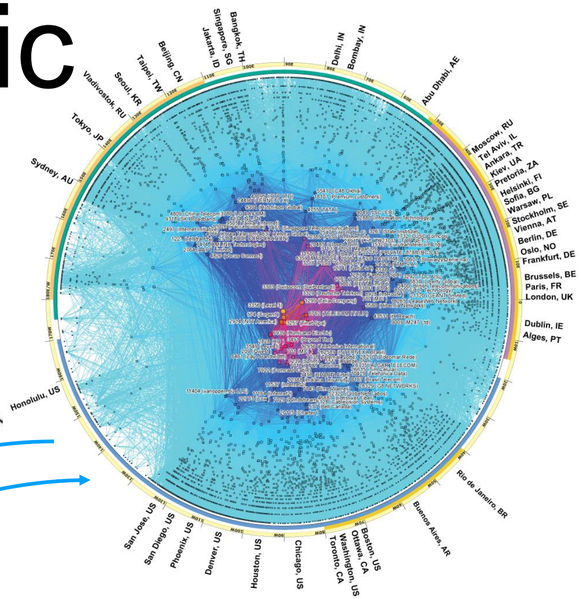

Hijacks can be used to divert traffic and gain inside knowledge

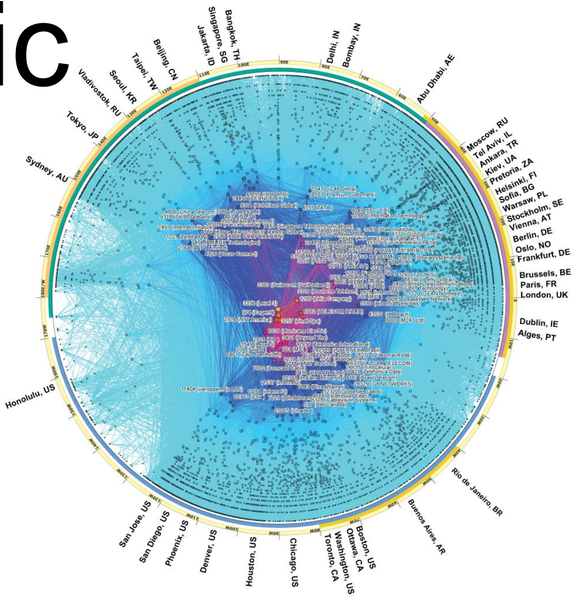

BGP hijacking

Informations Personnelles (Mail)

Phishing + Ransomware

Fournisseur d'accès

Hijacks can be used to divert traffic and gain inside knowledge

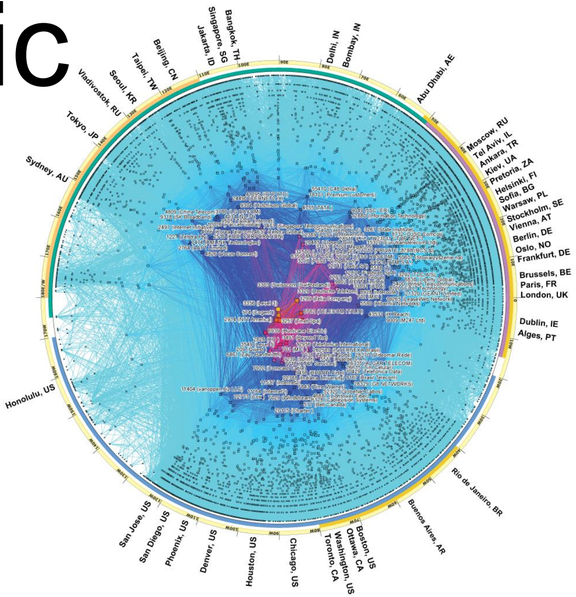

BGP hijacking

Informations Personnelles (Mail)

Fournisseur d'accès

Phishing + Ransomware

Multiple causes for hijacks

Hijacks are not always malicious

They can be the result of misconfigurations

https://www.manrs.org/2022/06/configuration-issue-penalizing-single-digit-asns/?utm_source=rss&utm_medium=rss&utm_campaign=configuration-issue-penalizing-single-digit-asns

Extract from the blog post:

“In recent years, we’ve noticed that single-digit ASNs (ASN1 through ASN9) often appear to be route hijackers. Is this true? We dug into the data and ultimately realized **no, single-digit ASNs are not hijacking address space at an alarming rate**. What’s happening is the result of a misconfiguration issue because of the “AS path prepend” command on Mikrotik routers.”

https://www.manrs.org/2022/06/configuration-issue-penalizing-single-digit-asns/?utm_source=rss&utm_medium=rss&utm_campaign=configuration-issue-penalizing-single-digit-asns

Some vulnerabilities of BGP

Prefix hijacks

Blackjack attacks

BGP lies

BGP session injection

A Backjack attack surfs on the blackholing mechanism provided to protect against DDoS

DDoS are frequent

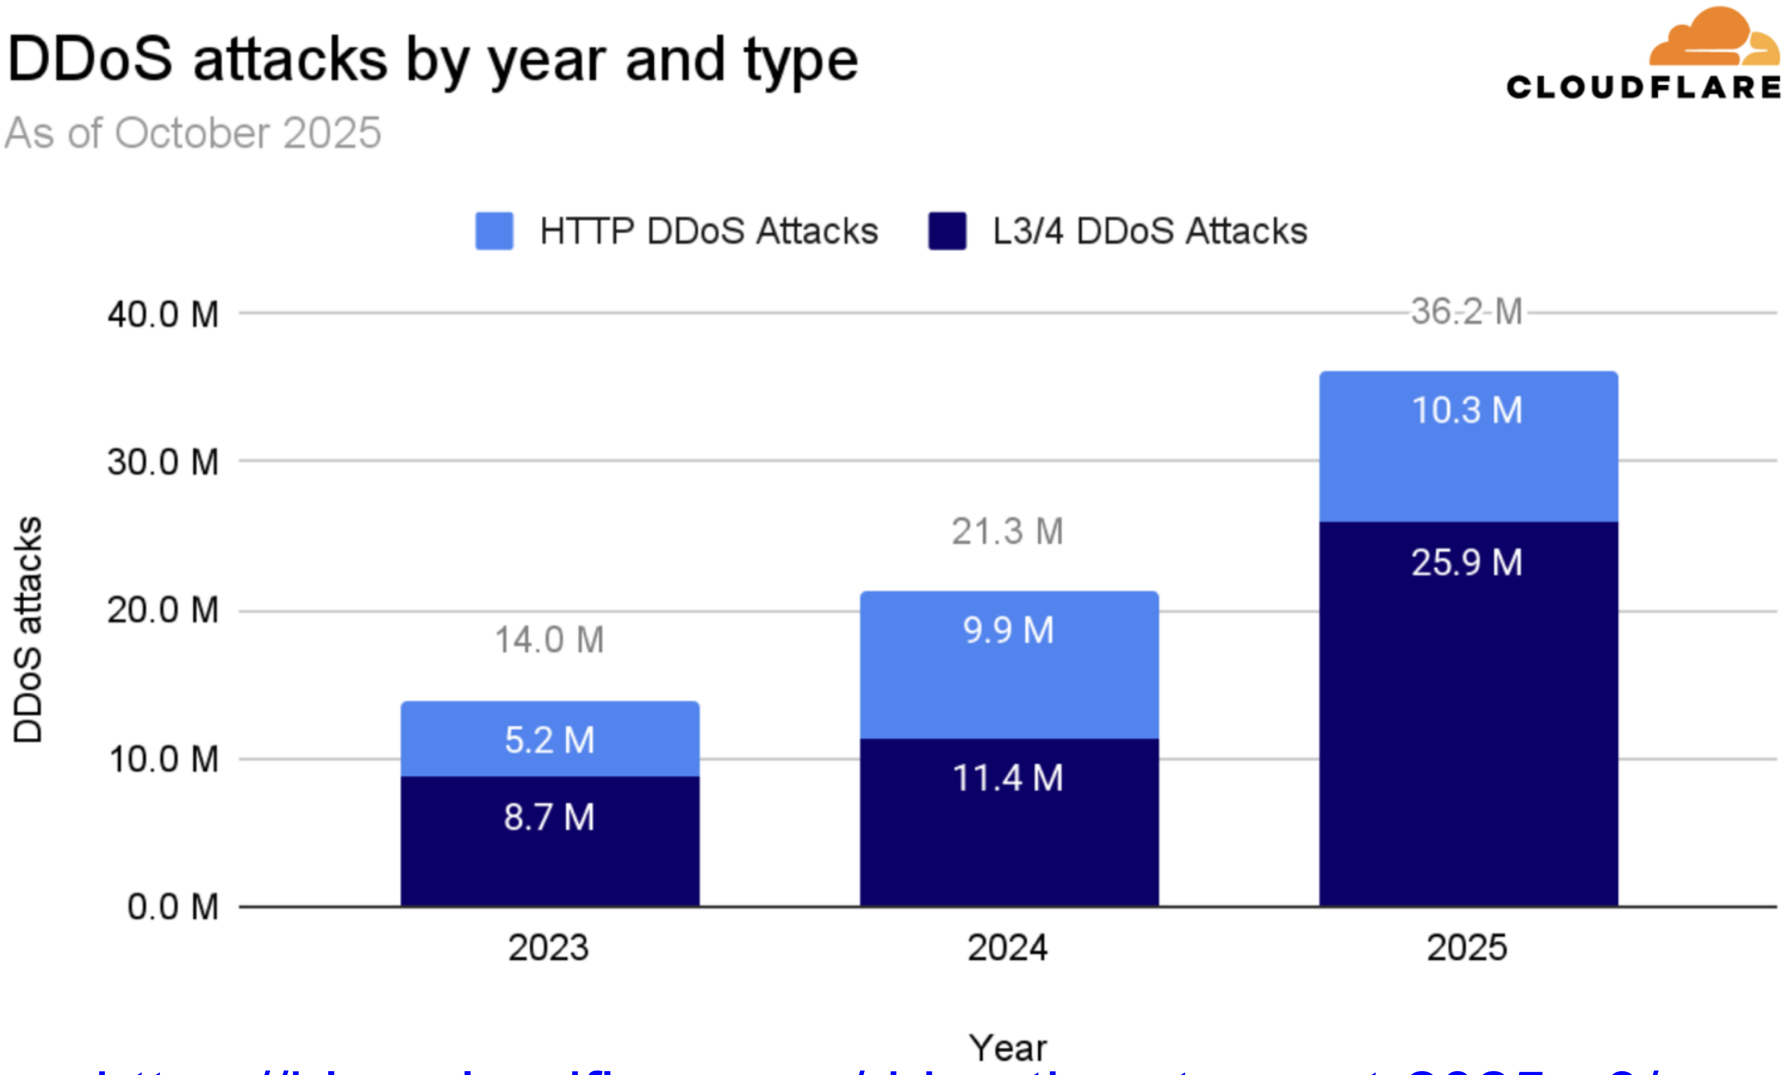

https://blog.cloudflare.com/ddos-threat-report-2025-q3/

DDoS

In a denial of service attack, the infractucture may be congested.

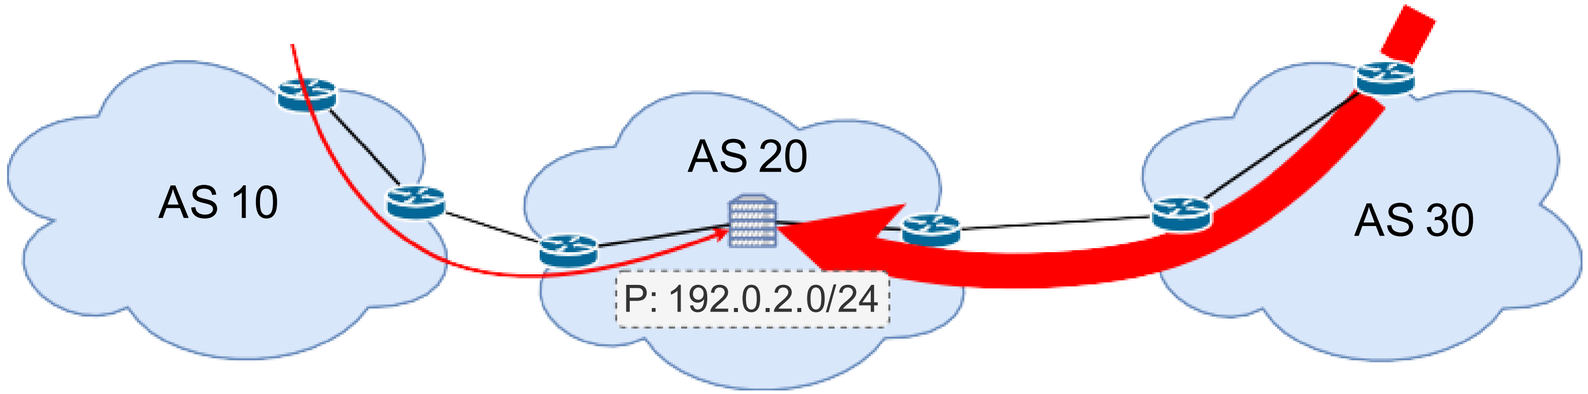

BGP blackholing

Blackholing is a **DDoS mitigation** technique signaled via **BGP** using a community.

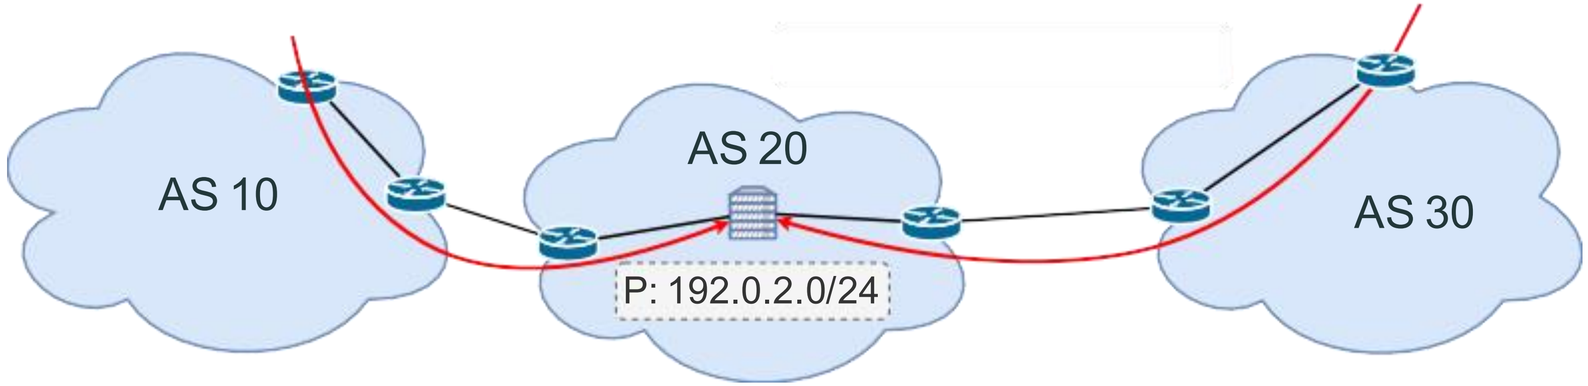

BGP blackholing

Blackholing is a **DDoS mitigation** technique signaled via **BGP** using a community.

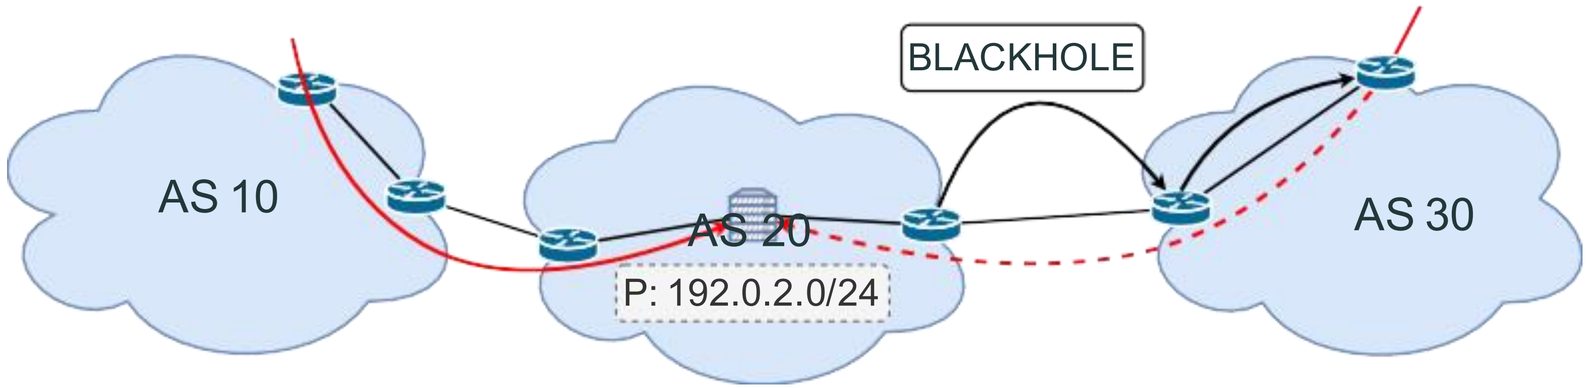

Blackholing has a double-edged sword effect: **all** traffic is dropped.

BGP community usage is increasing

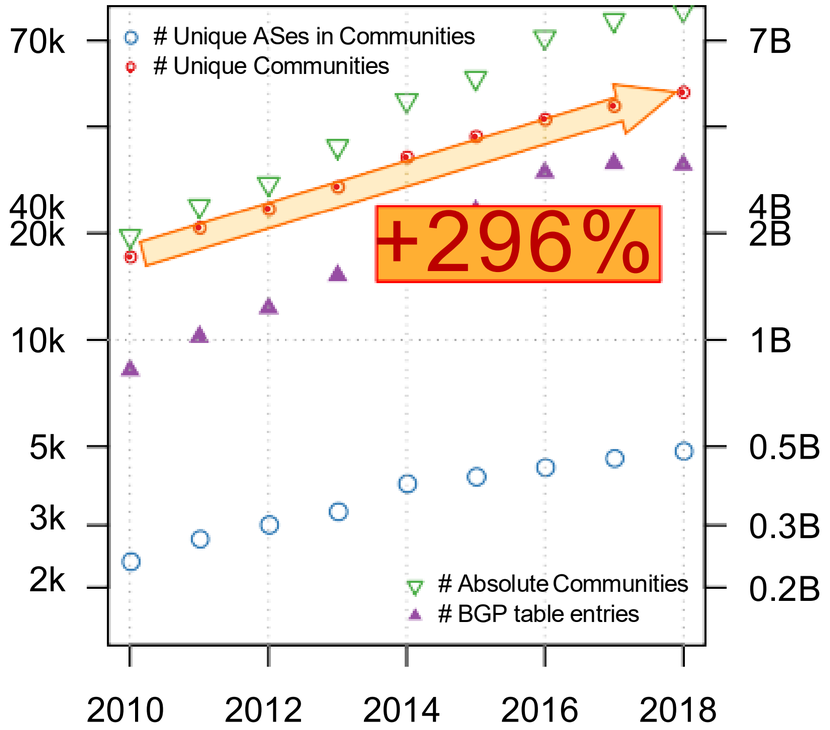

Increasing usage warrants a closer look.

Given the **increasing popularity** of BGP communities

and the ability to **trigger actions**, the first question that comes to the mind of an Internet measurement researcher is. . .

What could possibly go wrong?

Hijack-0 and Blackjack-0

Sermpezis 2018 (Artemis)

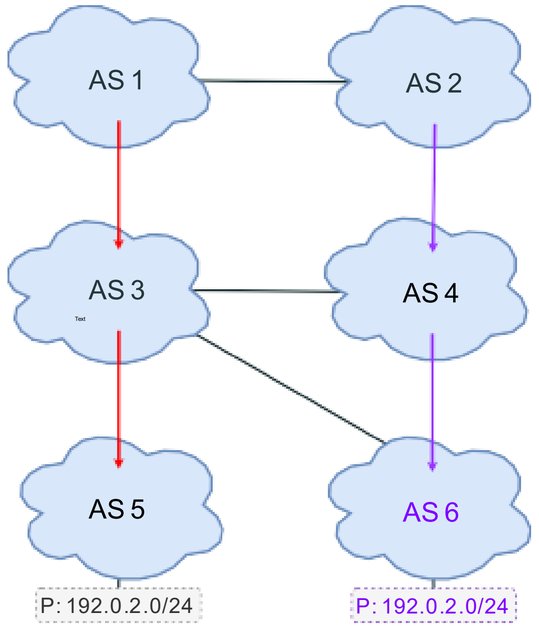

Hijack type-0

AS2 and AS4 traffic is de-routed to AS6 because the advertised path is shorter.

Hijack-0 and Blackjack-0

Sermpezis 2018 (Artemis)

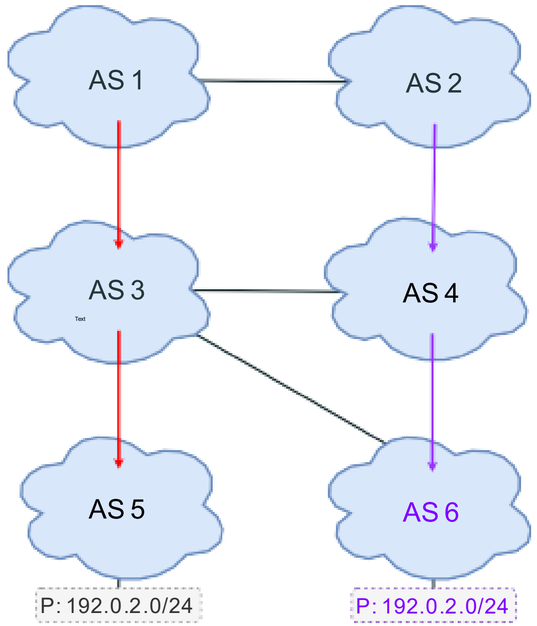

Hijack type-0

AS2 and AS4 traffic is

de-routed to AS6 because the advertised path is shorter.

Miller et Pelsser 2019

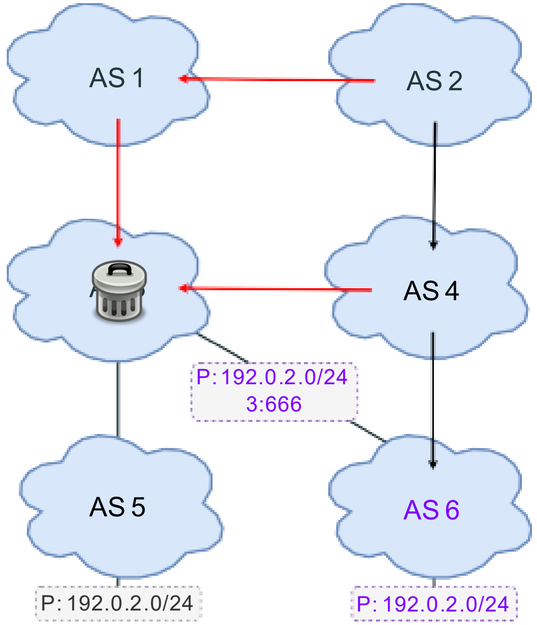

Blackjack type-0

All traffic to *P* is blackholed at AS3.

Hijacking + blackholing

Best practices for legitimate blackholing empower blackjacks

Best Practices for blackholing[4]

Give a **higher priority** to blackholing.

Do **not propagate** the advertisement across AS borders.

[4]Cisco, **Remotely Triggered Black Hole Filtering - Destination Based and Source Based**.

Best practices for legitimate blackholing empower blackjacks

Best Practices for blackholing[4]

Give a **higher priority** to blackholing.

Do **not propagate** the advertisement across AS borders.

Consequences

Reach: Precedence over AS path length. Even ASes far away are vulnerable.

Stealth: The attacker is not dropping traffic himself.

[4]Cisco, **Remotely Triggered Black Hole Filtering - Destination Based and Source Based**.

How to prevent mis-origination?

ROA Route Origin Authorizations are digitally signed objects attesting that a given AS is **authorized to originate** routes for a set of prefixes.

ROV With Route Origin validation, an AS **validates the origin** of the BGP updates with regard to the content of the RPKI Objects.

But other attacks are possible.

BGP Blackjacks - Type-N

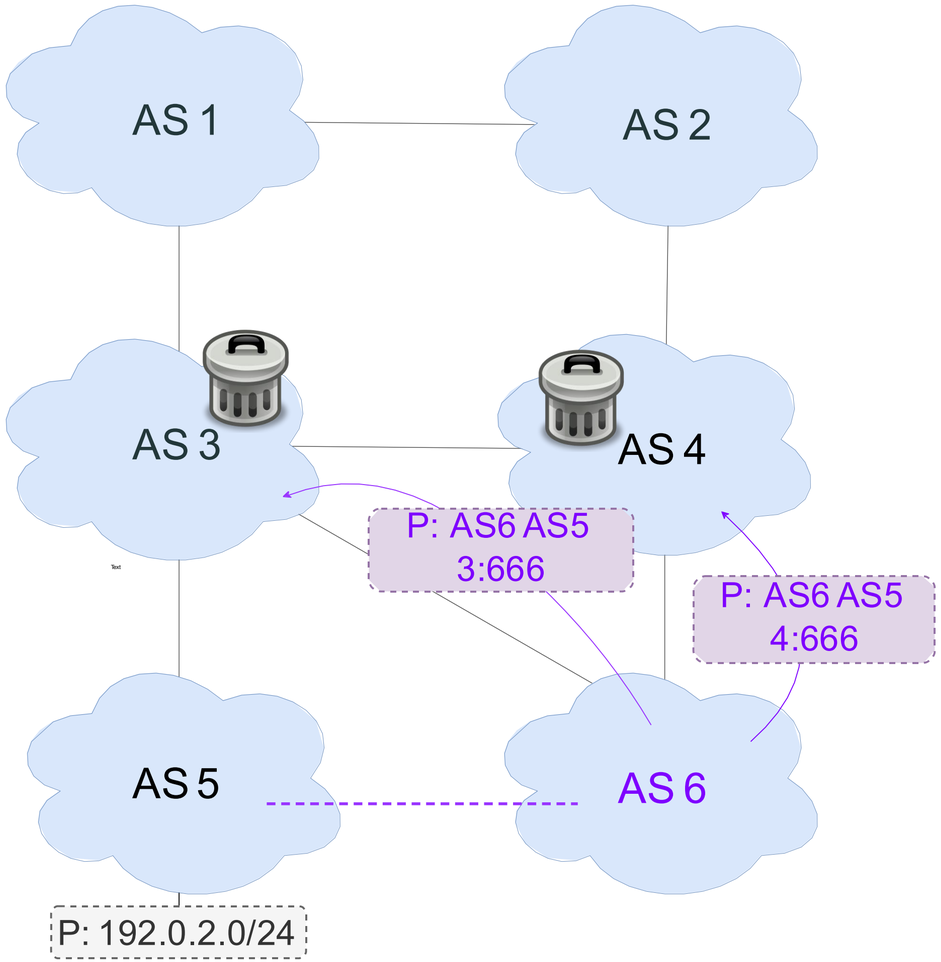

The origin AS is legit. The AS-path is not.

BGPsec[5]

BGPsec allows ASes to **sign** advertisements.

This guarantees the AS path reflects the **actual path** the advertisement went through.

But on-paths attacks are still possible.

[5]Lepinski and Sriram, **BGPsec Protocol Specification**.

Related publications

Taxonomy of Attacks using BGP Blackholing.

Loic Miller (U. Strasbourg), Cristel Pelsser (U. Strasbourg). ESORICS 2019.

BGP Communities: Even more Worms in the Routing Can.

Florian Streibelt (MPI[1]), Franziska Lichtblau (MPI), Robert Beverly (NPS[2]), Anja Feldmann (MPI), Cristel Pelsser (U. Strasbourg), Georgios Smaragdakis (TU Berlin), Randy Bush (IIJ[3]). ACM IMC 2018.

[1]Max Planck Institute for Informatics
[2]Naval Postgraduate School
[3]Internet Initiative Japan

Some vulnerabilities of BGP

Prefix hijacks

Blackjack attacks

BGP lies

BGP session injection

BGP runs on top of TCP

TCP is vulnerable to injection attacks

The attacker

- guesses the next sequence number
- sends a packet with the sequence number and forged content

The client accepts the content if it arrives before the legit packet

The recommendation is to use MD5 for session authentication.

- But there are tools able to provide payload for a given MD5 digest https://github.com/DavidBuchanan314/monomorph
- The adoption status of TCP Authentication Option (TCP-AO) for BGP is not known

Related publication

Routing over QUIC: Bringing transport innovations to routing protocols. Thomas Wirtgen, Nicolas Rybowski, Cristel Pelsser, Olivier Bonaventure (2023). Poster at NSDI 2023.

Some vulnerabilities of BGP

Prefix hijacks

Blackholing

BGP lies

BGP session injection

$\Rightarrow$BGP designed with no security in mind

Weak authentication

No integrity protection

How we may hack to live with these vulnerabilities

Prevention: some fixes

- RPKI ROA and ROV
  - State of deployment
- BGP filters
  - MANRS
- BGPsec

RPKI ROA

MANRS ROA Stats Tool

HOME | BY COUNTRY | BY ASN | Problem with the data? | MANRS

MANRS ROA Stats Tool

Search for ROA stats by country or ASN using the links above

Data last retrieved 1 day(s) ago

IPv4		
Total	922799	
Valid ROAs	357837	38.78%
Unknown ROAs	556469	60.3%
Invalid ROAs	8493	0.92%

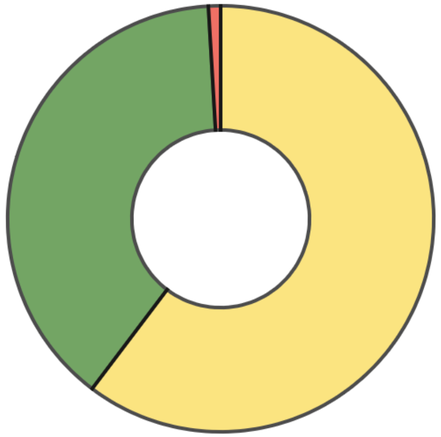

IPv6		
Total	139695	
Valid ROAs	61793	44.23%
Unknown ROAs	76338	54.65%
Invalid ROAs	1564	1.12%

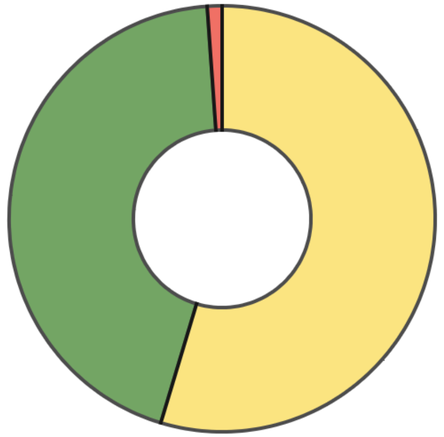

https://roa-stats.manrs.org (October 6, 2022)

RPKI ROA today ~ 60% valid prefixes

V4 Valid	Pc	V4 Invalid	Pc	V4 Unknwn	Pc
713,487	57.6%	40,701	3.3%	483,681	39.1%

V6 Valid	Pc	V6 Invalid	Pc	V6 Unknwn	Pc	V6 Total Routes
177,035	60.7%	14,498	5.0%	100,284	34.4%	291,817

https://stats.labs.apnic.net/roas 26/01/2026

RPKI ROV

75 ASs deploy ROV (certainty above 0.7) according to **rov.rpki.net** (out of > 73.5k) → Last measurement was on 2020-08-31

Only 5.9 % of user prefixes are protected according to the MANRS Observatory (May 2023)

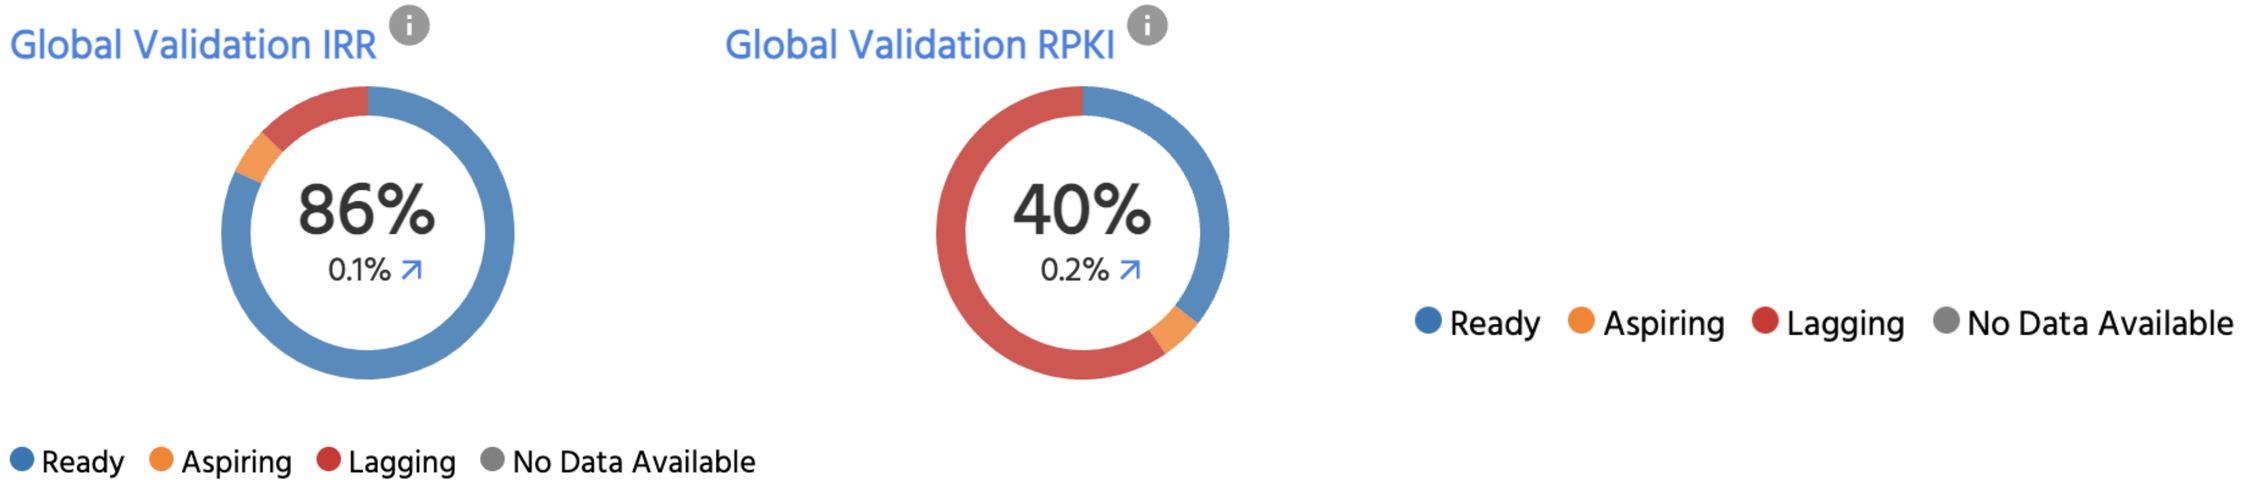

From https://observatory.manrs.org/#/overview (May, 2023)

Detecting and localizing who deploys ROV

- **Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering,** Andreas Reuter et al (2017). Computer Communications Review (CCR).

- BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping. C. Gray, M. Clemens, R. Bush, C. Pelsser, R. Matthew, T. Schmidt, M. Wählisch (2020). *Internet Measurement Conference (IMC)*.

- RoVista: Measuring and Analyzing the Route Origin Validation (ROV) in RPKI. Weitong Li, Zhexiao Lin, Md. Ishtiaq Ashiq, Emile Aben, Romain Fontugne, Amreesh Phokeer, and Taejoong Chung (2023). Internet Measurement Conference (IMC).

BGP filters and MANRS

Mutually Agreed Norms for Routing Security (MANRS)

Action	Metric	Description	Data source(s)
Filtering	M1	Route leak by the AS Calculates incidents where the AS was the culprit of BGP leakage events. In the example on Fig 1. if all pink events are route leaks by the AS, M1=3.5	bgpstream
	M2	Route misorigination by the AS Calculates incidents where the AS was the culprit of BGP misorigination (hijacking) events.	bgpstream or GRIP (*)
	M1C	Route leak by a direct customer Calculates incidents where the AS was an accomplice (the misoriginating AS was present in the AS-PATH) to BGP leakage events. Currently only incidents related to adjacent networks are taken into account.	bgpstream
	M2C	Route misorigination by a direct customer Calculates incidents where the AS was an accomplice (the leaking AS was present in the AS-PATH) to BGP hijack events. Currently only incidents related to adjacent networks are considered.	bgpstream or GRIP (*)
	M3	Bogon prefixes by the AS. Calculates incidents where the AS originated bogon address space. Note that the duration of each incident is counted per day as the data in the CIDR report is available only on a daily basis. Like with leaks and hijacks all prefixes originated by the AS on a day counted as 1 incident.	CIDR report
	M3C	Bogon prefixes propagated by the AS. Calculates incidents where the AS propagated bogon address space announcements received from its peers.	CIDR report
	M4	Bogon ASNs by the AS Calculates incidents where the AS announced bogon ASNs as adjacency. Note that the duration of each incident is counted per day as the data in the CIDR report is available only on a daily basis.	CIDR report
	M4C	Bogon ASNs propagated by the AS Calculates incidents where the AS propagated bogon ASNs announcements it received from its peers. Note that the duration of each incident is counted per day as the data in the CIDR report is available only on a daily basis.	CIDR report

BGP filters and MANRS

Mutually Agreed Norms for Routing Security (MANRS) rules for filter setting to prevent

- Leaks
- Misorigination
- Bogon prefixes
- Bogon ASs

From the AS itself and from direct customers

Deployment of protection increases but events still occur (FR)

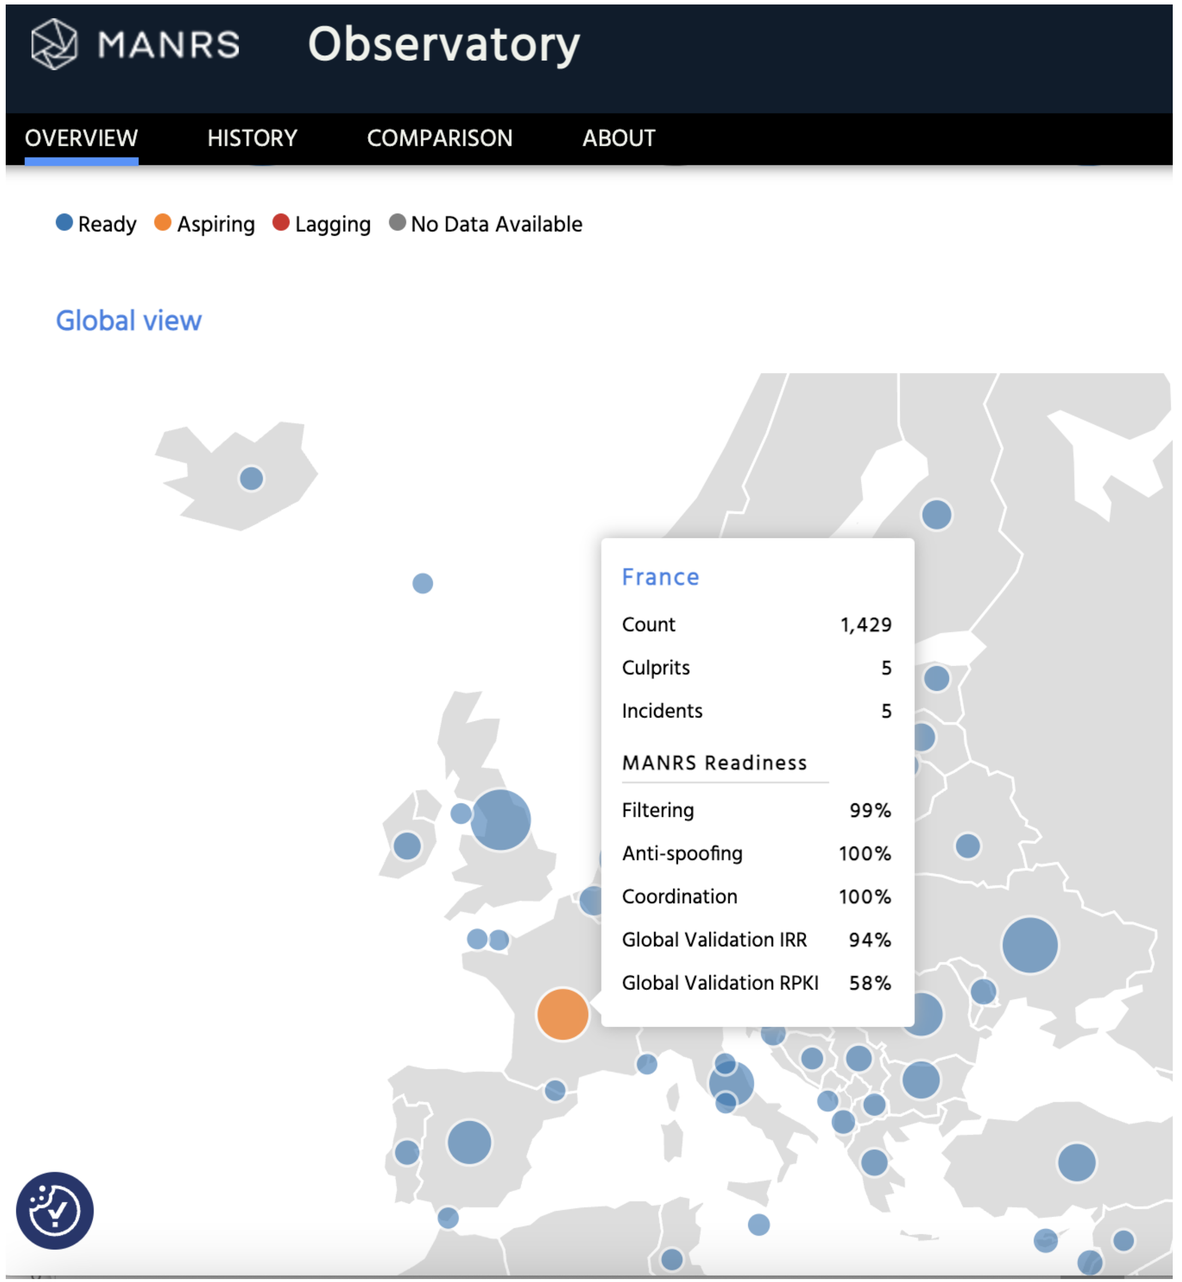

Deployement of protection increases but events still occur (US)

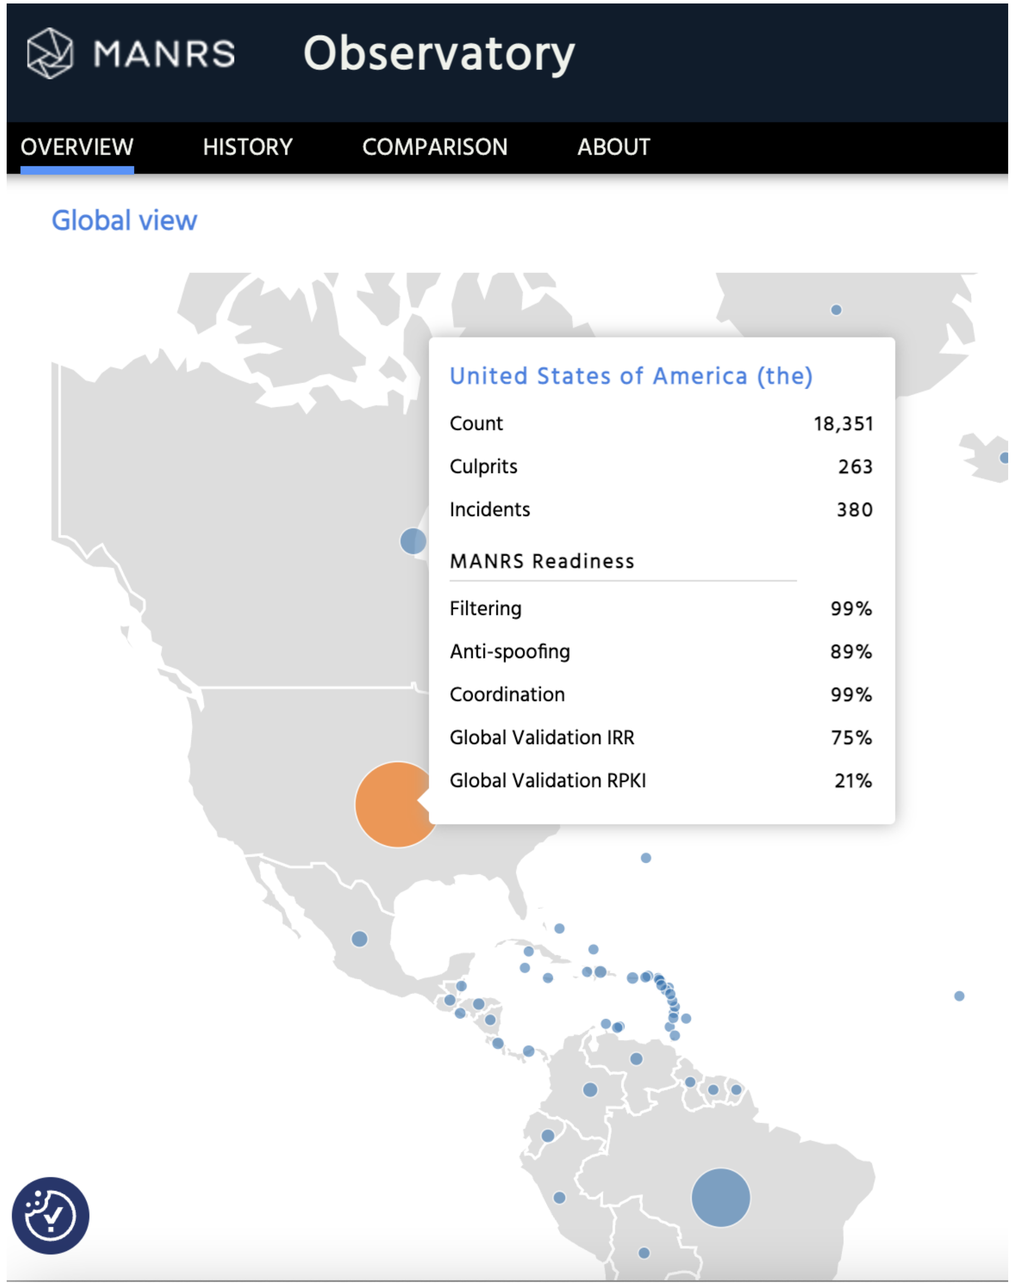

A System to Detect Forged-Origin BGP Hijacks

USENIX NSDI’24

Thomas Holterbach
Thomas Alfroy
Amreesh D. Phokeer
Alberto Dainotti
Cristel Pelsser

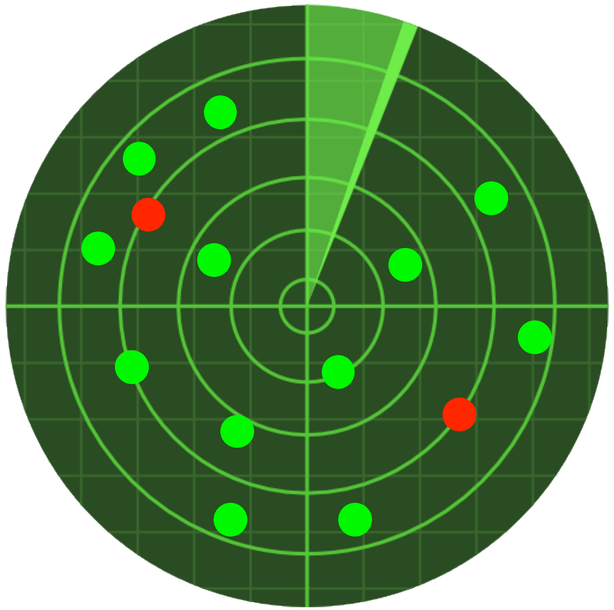

BGP is vulnerable to forged-origin hijacks

The attacker prepends the legitimate AS number to the AS path

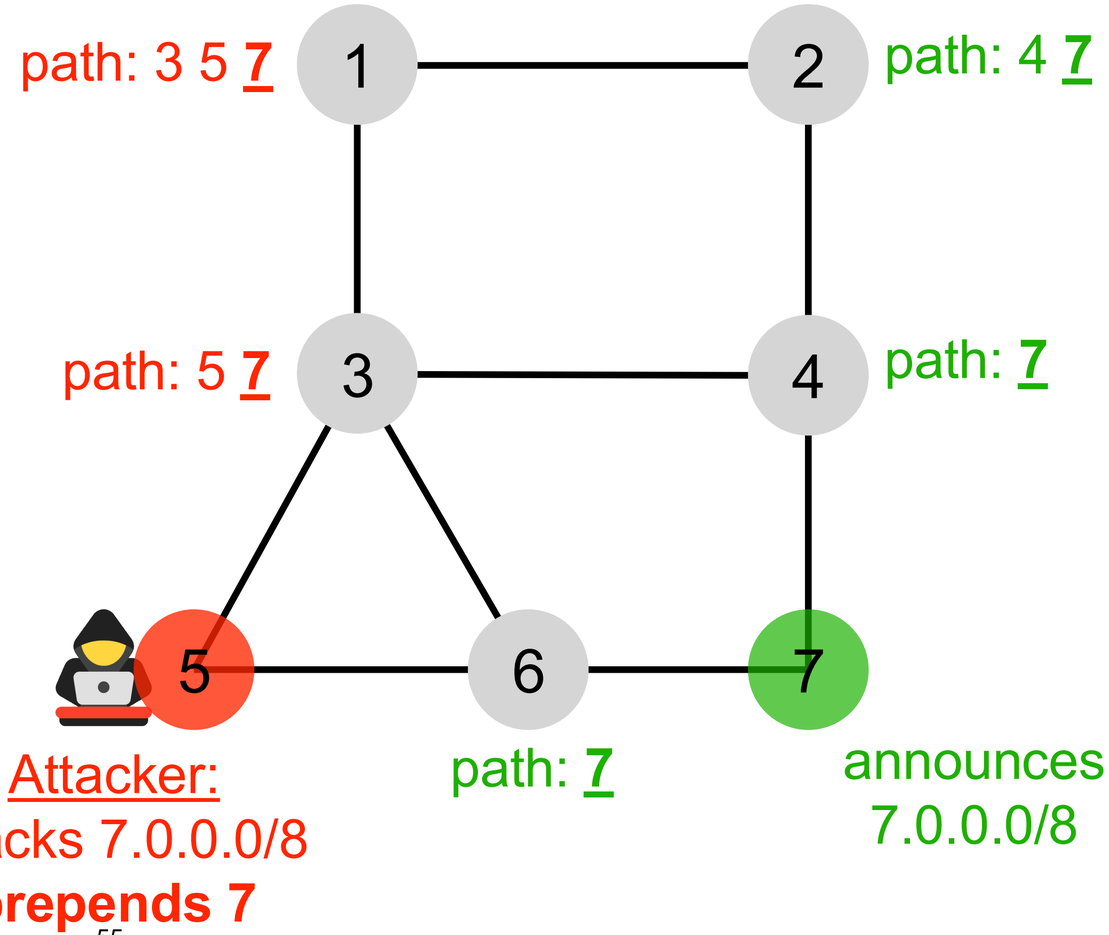

BGP is vulnerable to forged-origin hijacks

A significant fraction of the traffic is diverted to the attacker

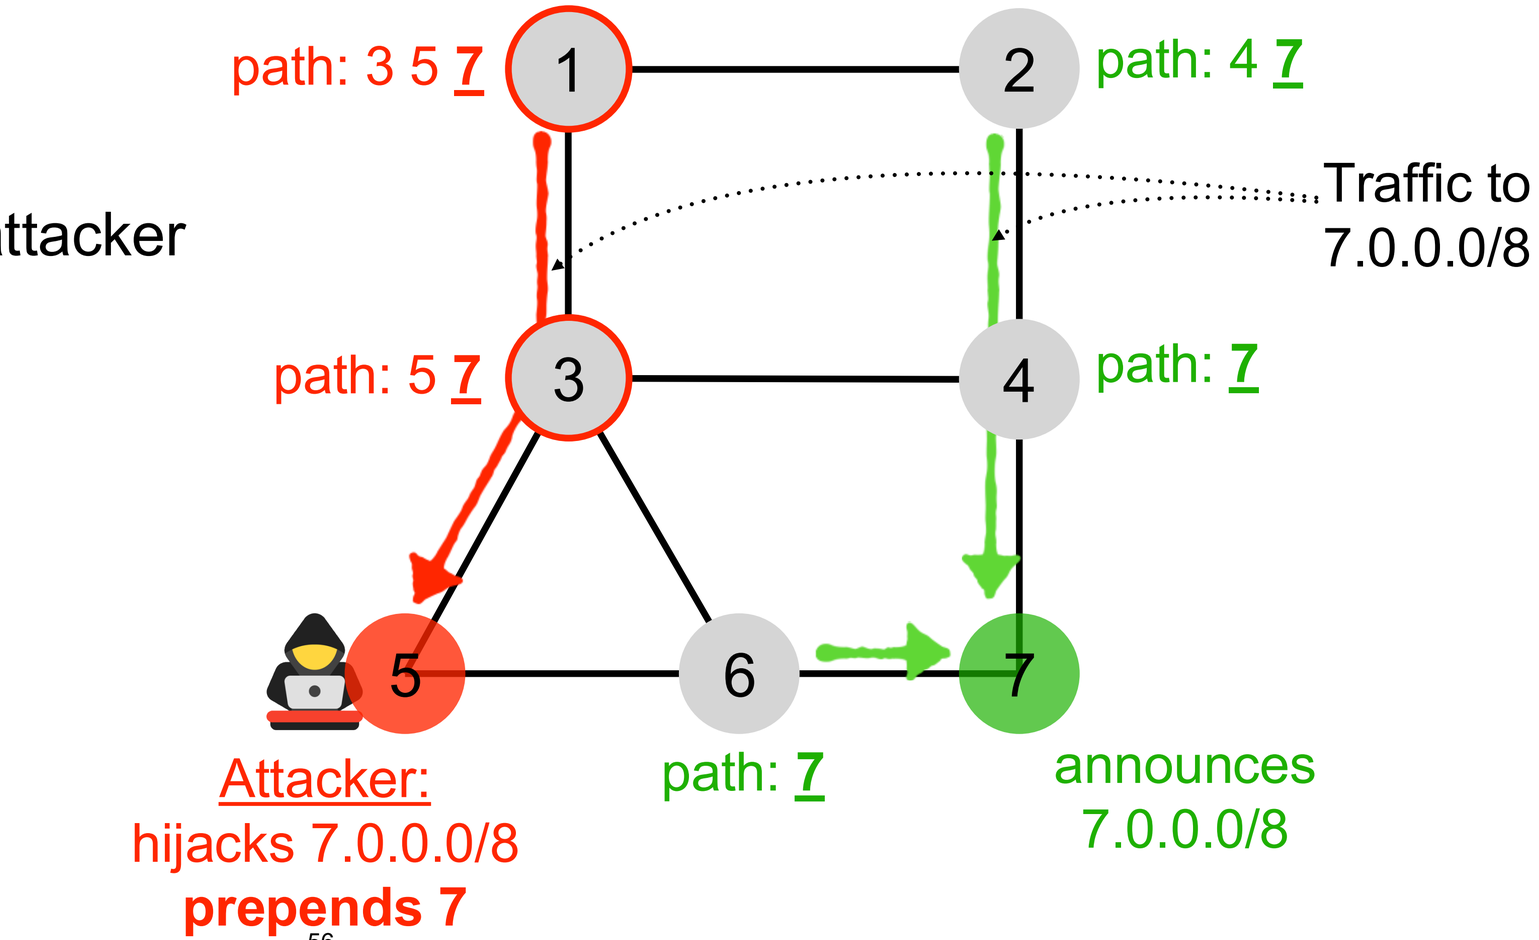

Forged-origin hijacks are actively used by attackers

August 17, 2022

The Record.
Recorded Future® News

February 3, 2022

KlaySwap crypto users lose funds after BGP hijack

Hackers have stolen roughly $1.9 million from South Korean cryptocurrency platform KLAYswap after they pulled off a rare and clever BGP hijack against the server infrastructure of one of the platform's providers.

The BGP hijack—which is the equivalent of hackers hijacking internet routes to bring users on malicious sites instead of legitimate ones—hit KakaoTalk, an instant messaging platform popular in South Korea.

The attack took place earlier this month, on February 3, lasted only for two hours, and KLAYswap has confirmed the incident last week and is currently issuing compensation for affected users.

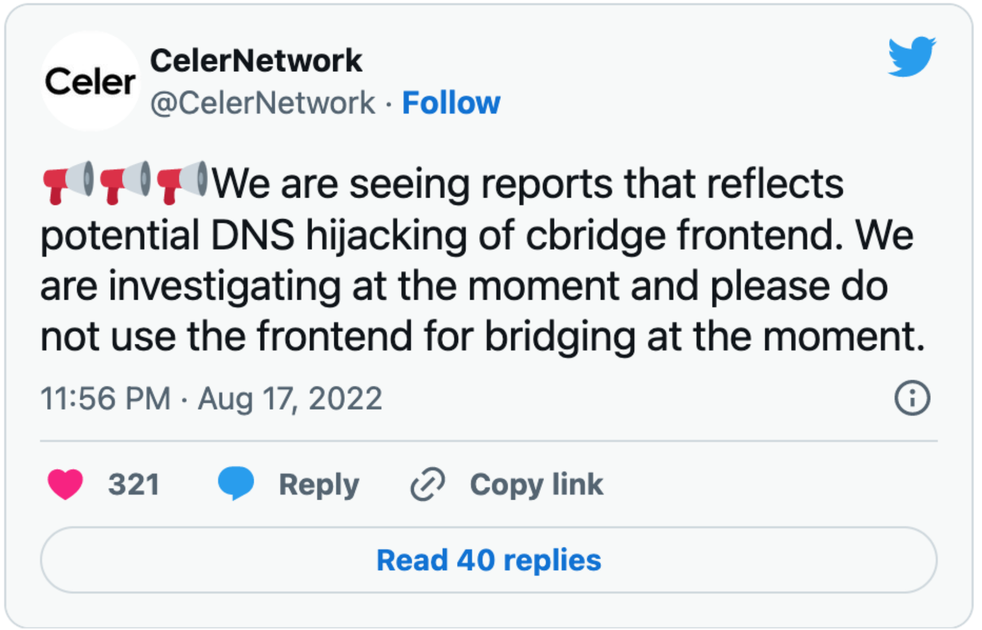

Both attacks are the result of a forged-origin hijack

DFOH: A System to **D**etect **F**orged-**O**rigin BGP **H**ijacks on the Whole Internet

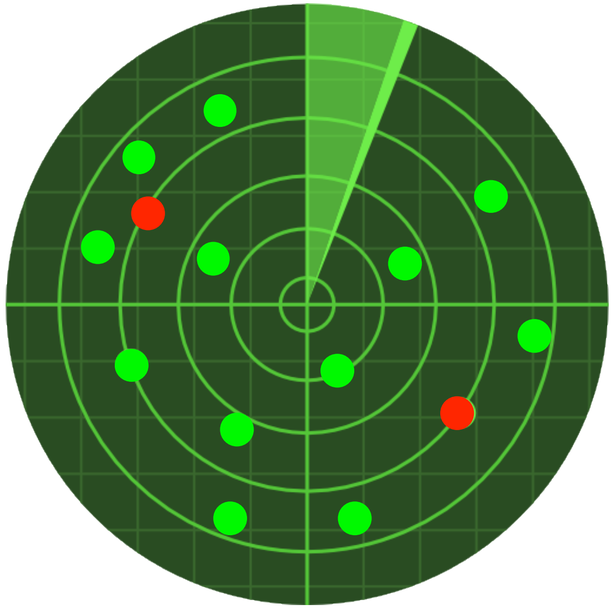

Outline

DFOH's main challenge

DFOH's inference pipeline

DFOH's inferences are accurate

DFOH *is* up and running

Outline

DFOH's main challenge is to detect fake AS links

DFOH's inference pipeline

DFOH's inferences are accurate

DFOH *is* up and running

DFOH aims to detect the fake AS links induced by forged-origin hijacks

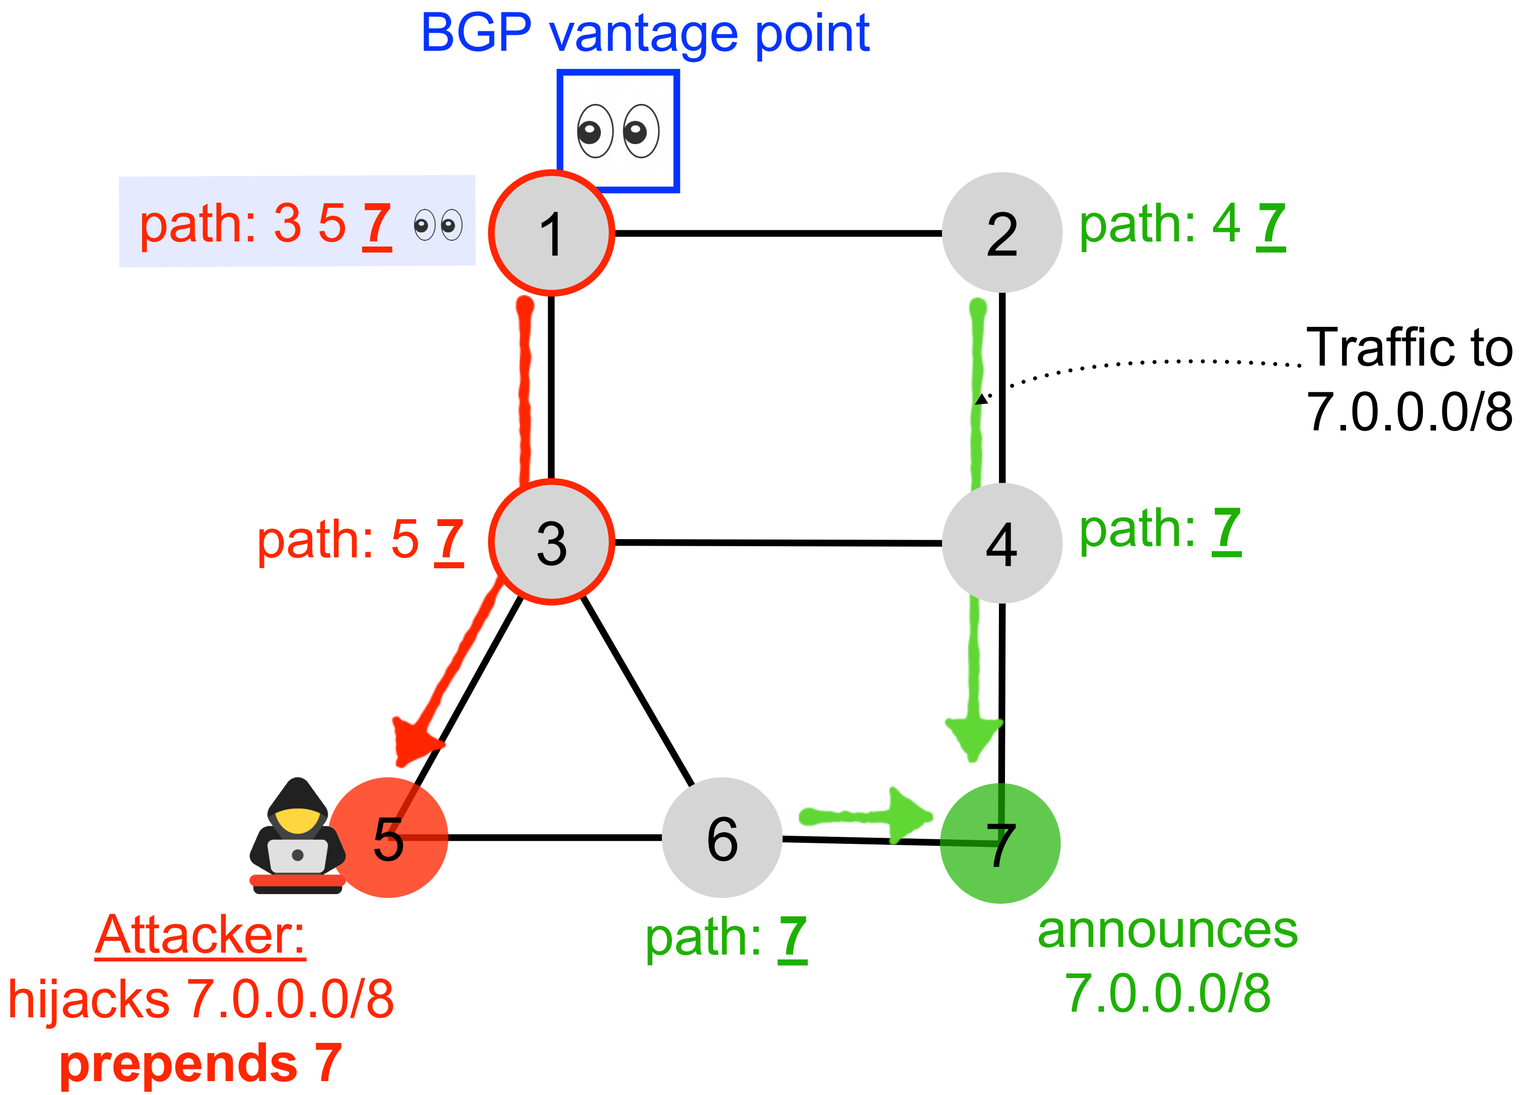

DFOH aims to detect the fake AS links induced by forged-origin hijacks

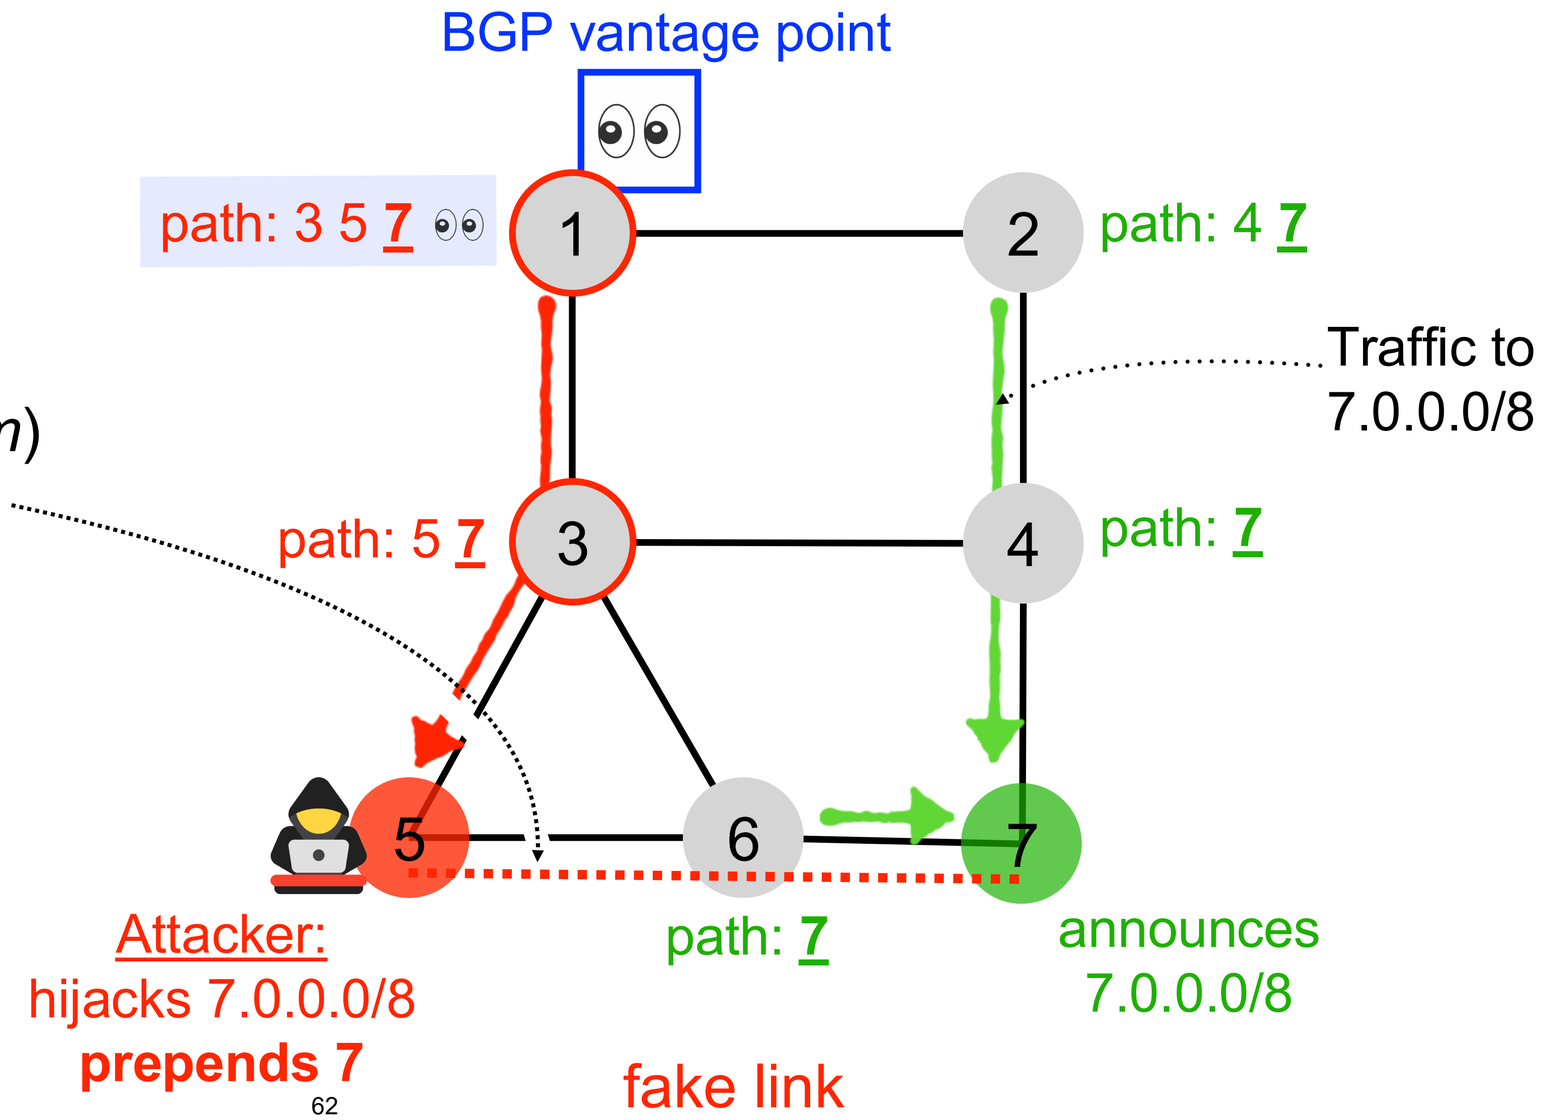

Upon the attack:
AS5 (*attacker*) and AS7 (*victim*) appear directly connected

An attacker cannot escape from creating a fake AS link without hampering the effectiveness of its attack

There is no new AS link if the attacker prepends **6 7**

But none of the ASes divert traffic to the attacker as the AS path is longer

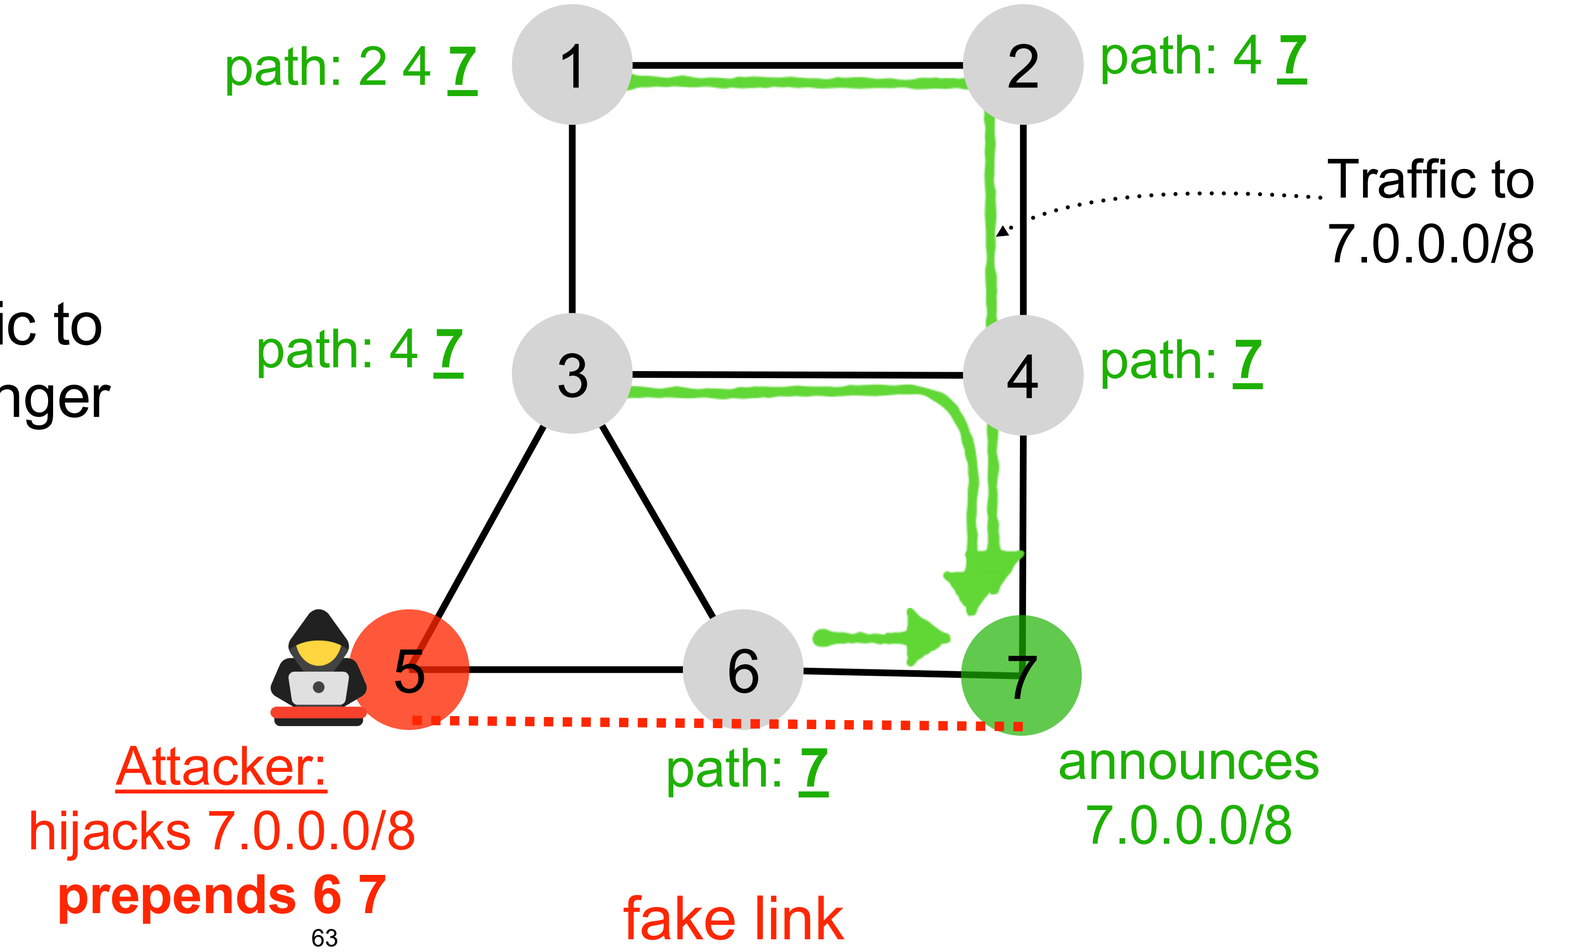

Problem: There are many new AS links every day
but no simple property that tells whether they are real or fake

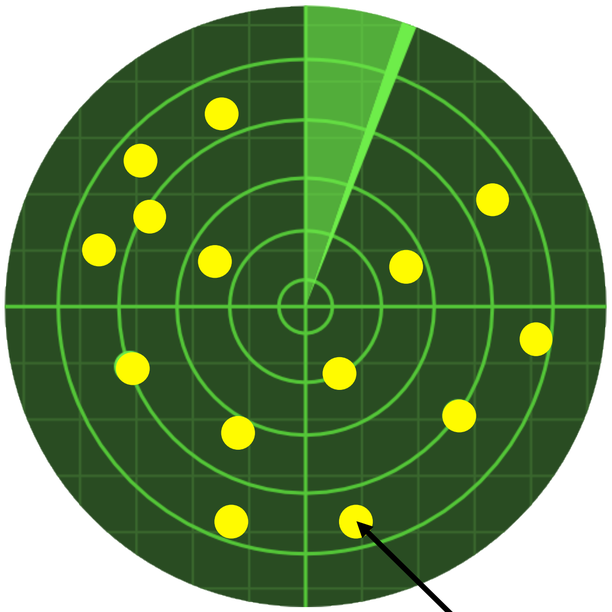

We find 166 new AS links every day (median)
and the vast majority are likely legitimate

Using the BGP data from 200 RIS and RouteViews peers and collected during ten months in 2022

Problem: There are many new AS links every day
but no simple property that tells whether they are real or fake

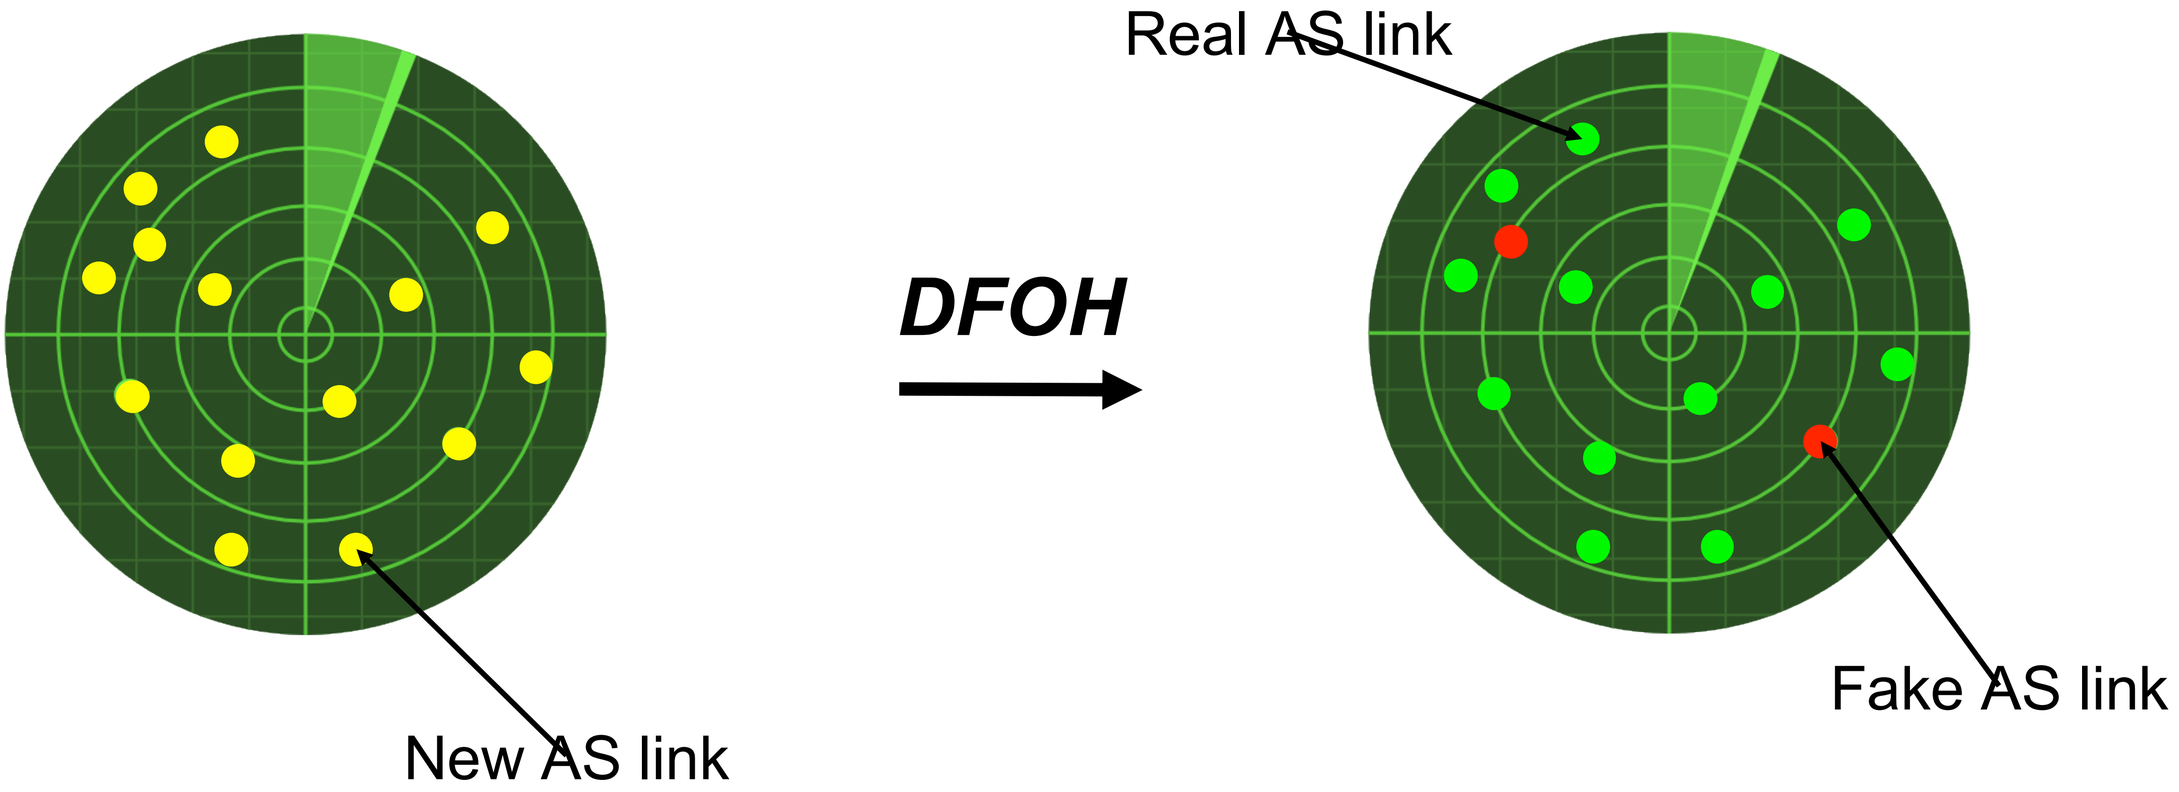

Outline

DFOH's main challenge — is to detect fake AS links

DFOH's inference pipeline — relies on domain-specific knowledge and a tailored link prediction framework

DFOH's inferences are accurate

DFOH *is* up and running

DFOH's fake AS links inference algorithm comprises three steps

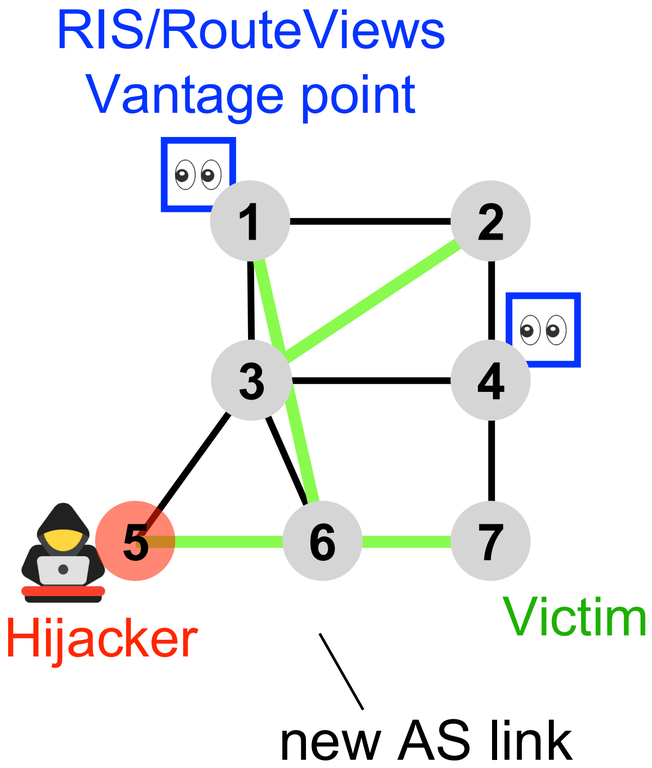

DFOH's fake AS links inference algorithm comprises three steps

RIS/RouteViews
Vantage point

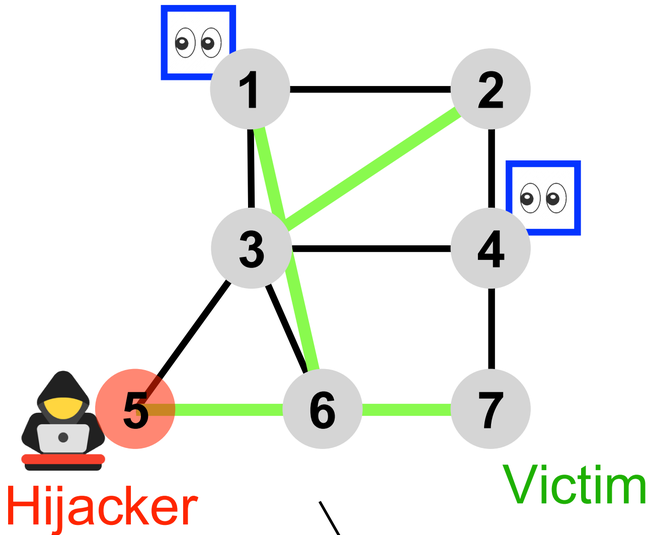

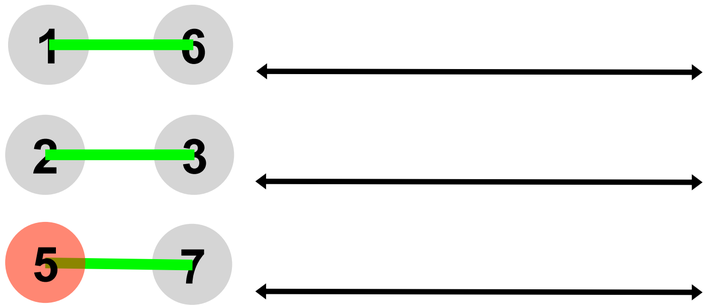

DFOH's fake AS links inference algorithm comprises three steps

RIS/RouteViews
Vantage point

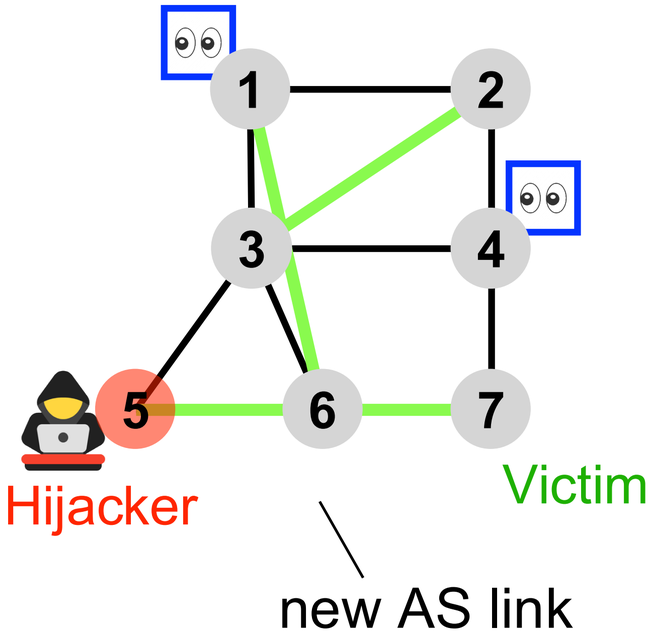

Feature categories:

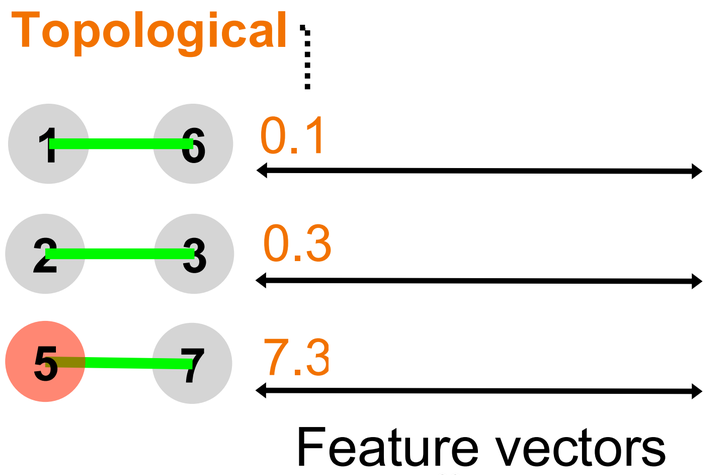

DFOH uses a total of 11 topological features that can be divided into four categories

Node centrality

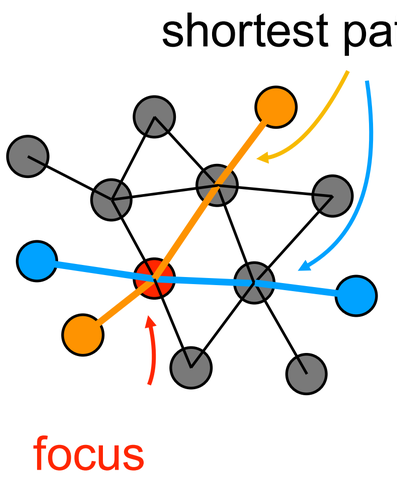

Neighborhood richness

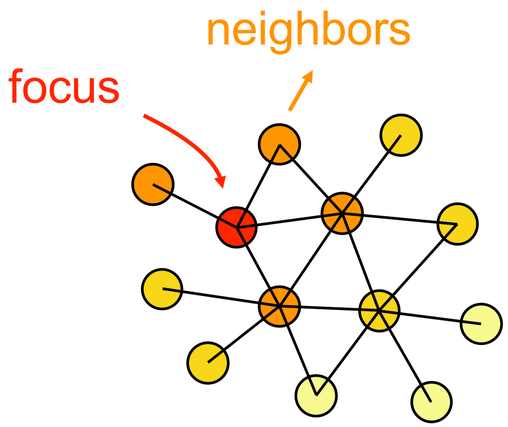

Topological patterns

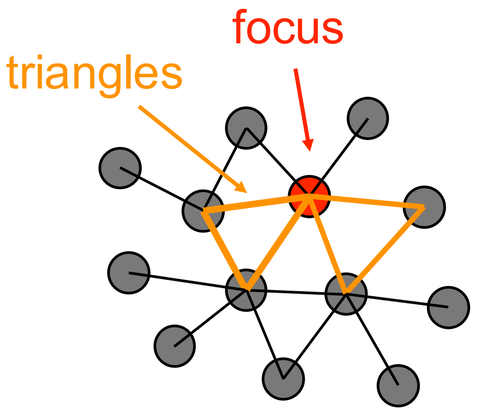

Closeness

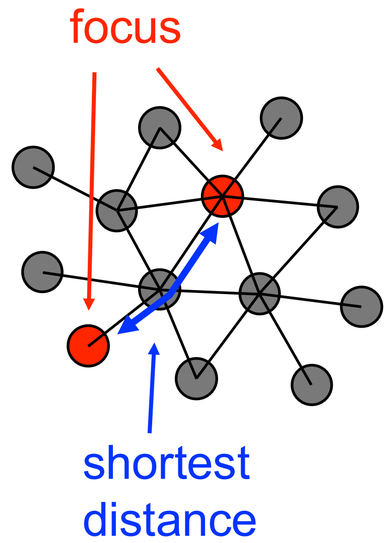

DFOH's fake AS links inference algorithm comprises three steps

RIS/RouteViews
Vantage point

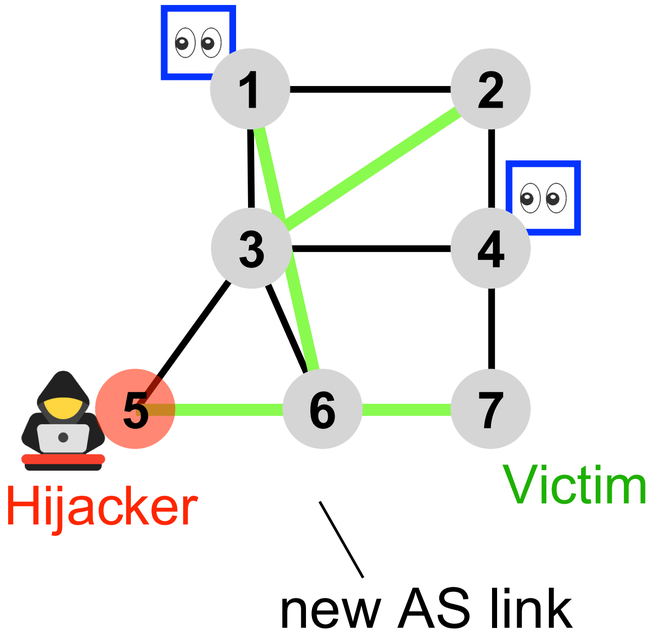

Feature categories:

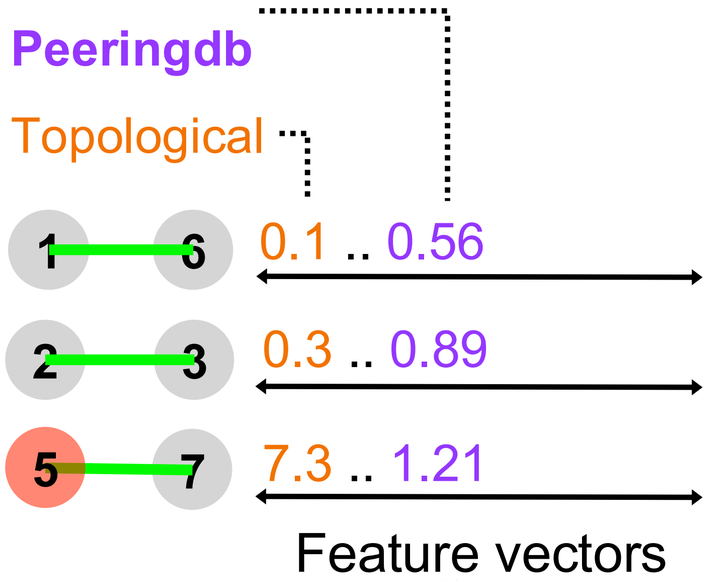

DFOH looks at public peering information
and identifies when two ASes are unlikely to peer

DFOH looks for three types
of information in PeeringDB:

1. Country
2. Public peering exchange points
3. Private peering facilities

DFOH compares the peering information of the neighbors of the hypothetical victim and attacker

Reason #1:
Protect against adversarial inputs

Reason #2:
Mitigate missing peering information

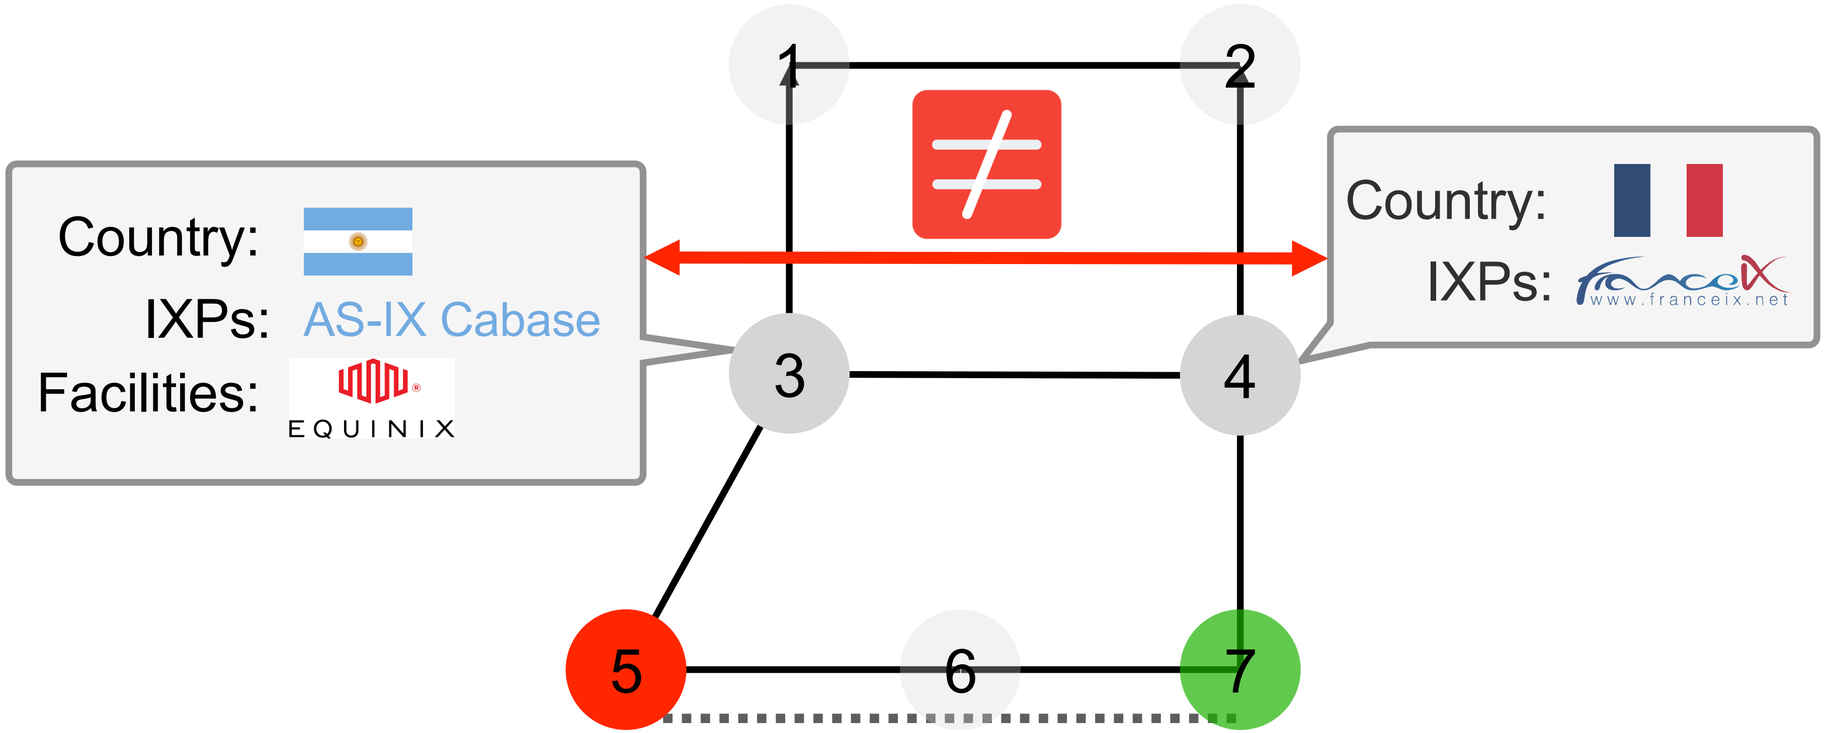

New AS link is fake

DFOH's fake AS links inference algorithm comprises three steps

RIS/RouteViews
Vantage point

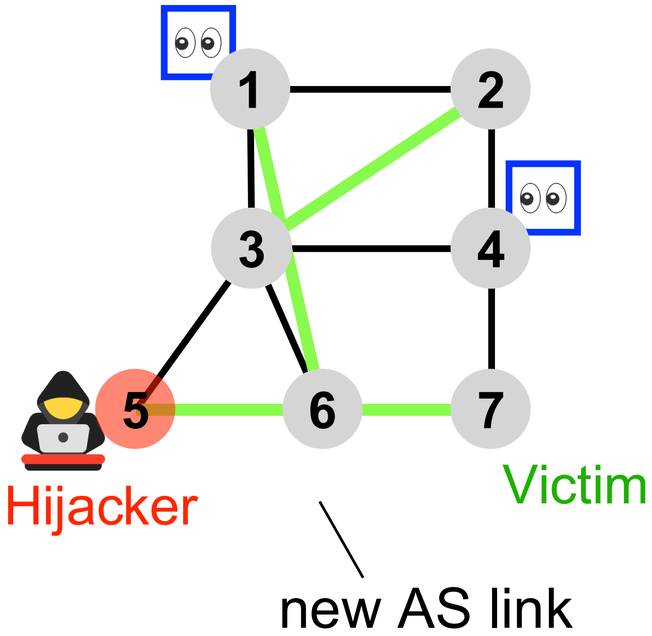

Feature categories:

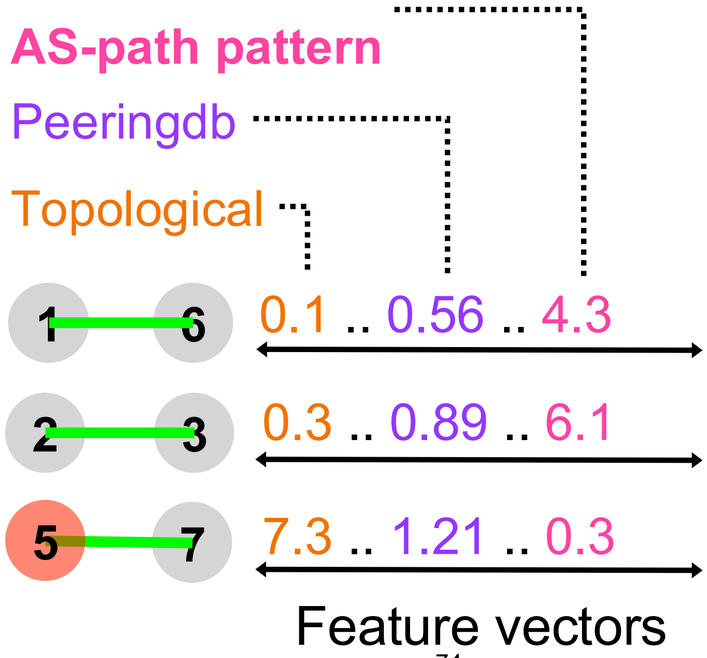

DFOH looks at the AS paths that include the new link and identifies suspicious sequence of ASes

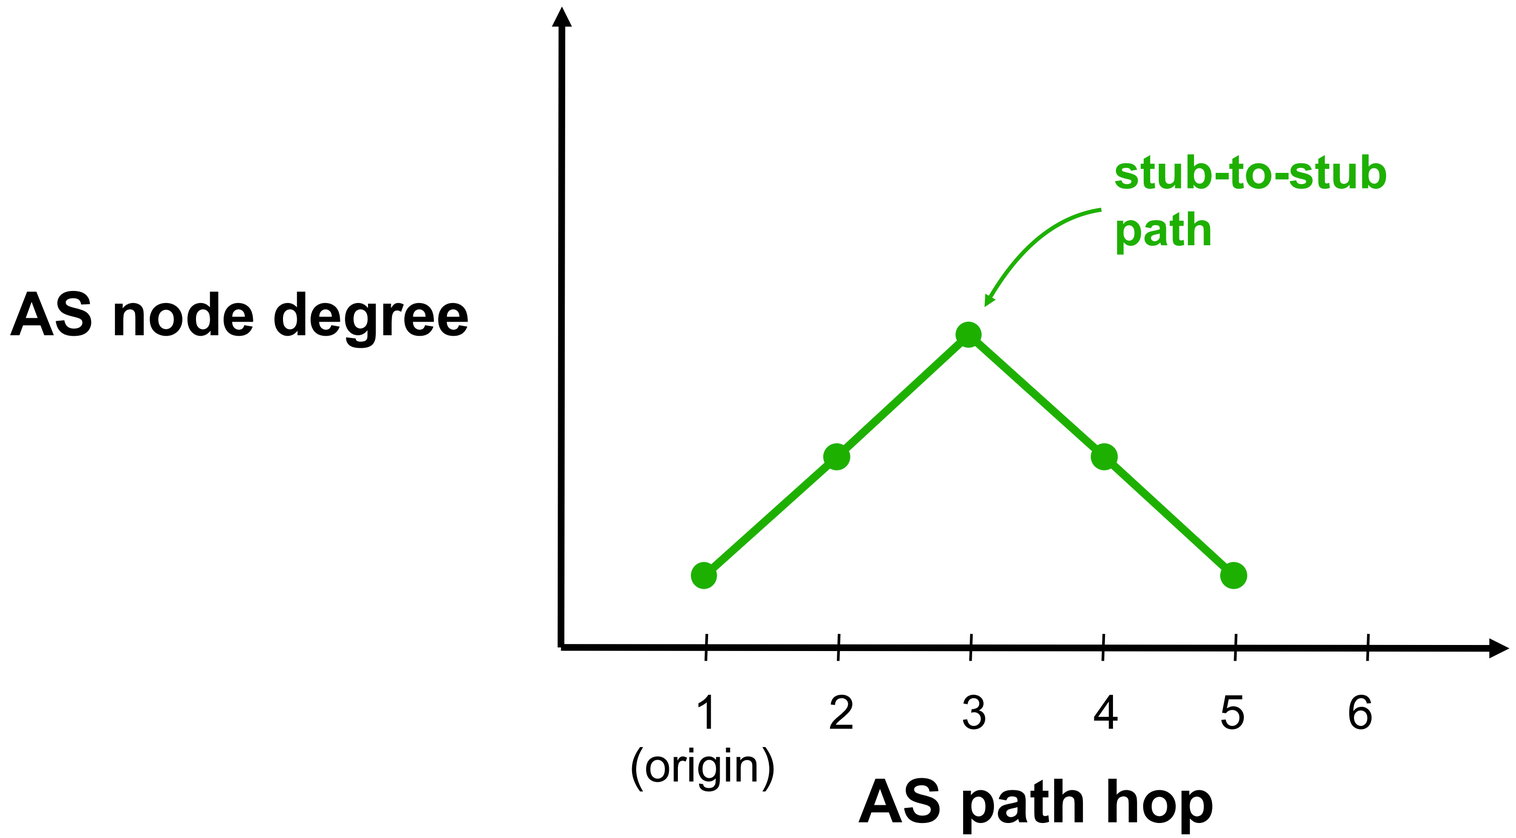

DFOH looks at the AS paths that include the new link and identifies suspicious sequence of ASes

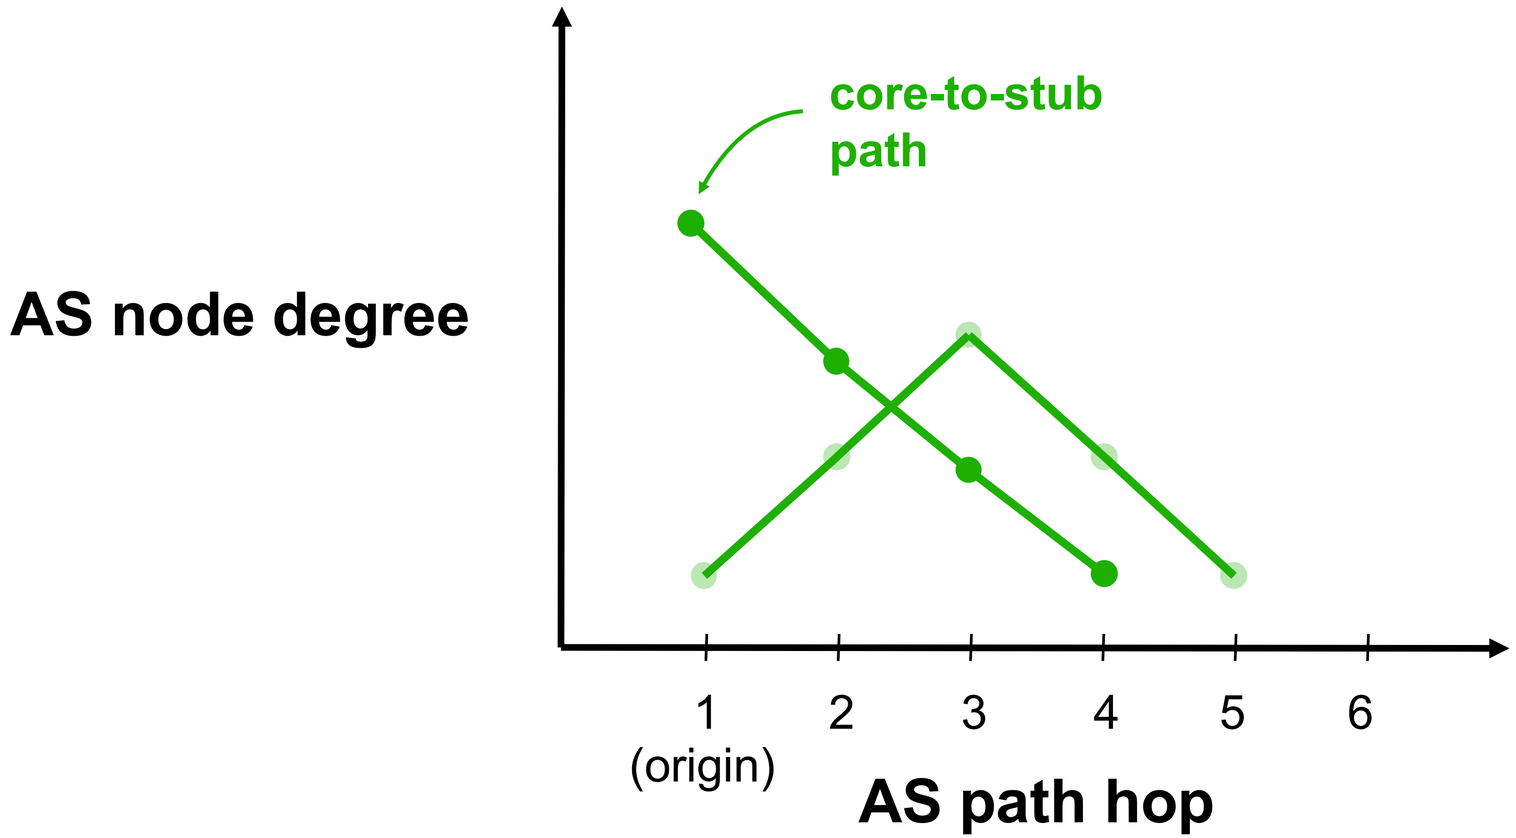

DFOH looks at the AS paths that include the new link and identifies suspicious sequence of ASes

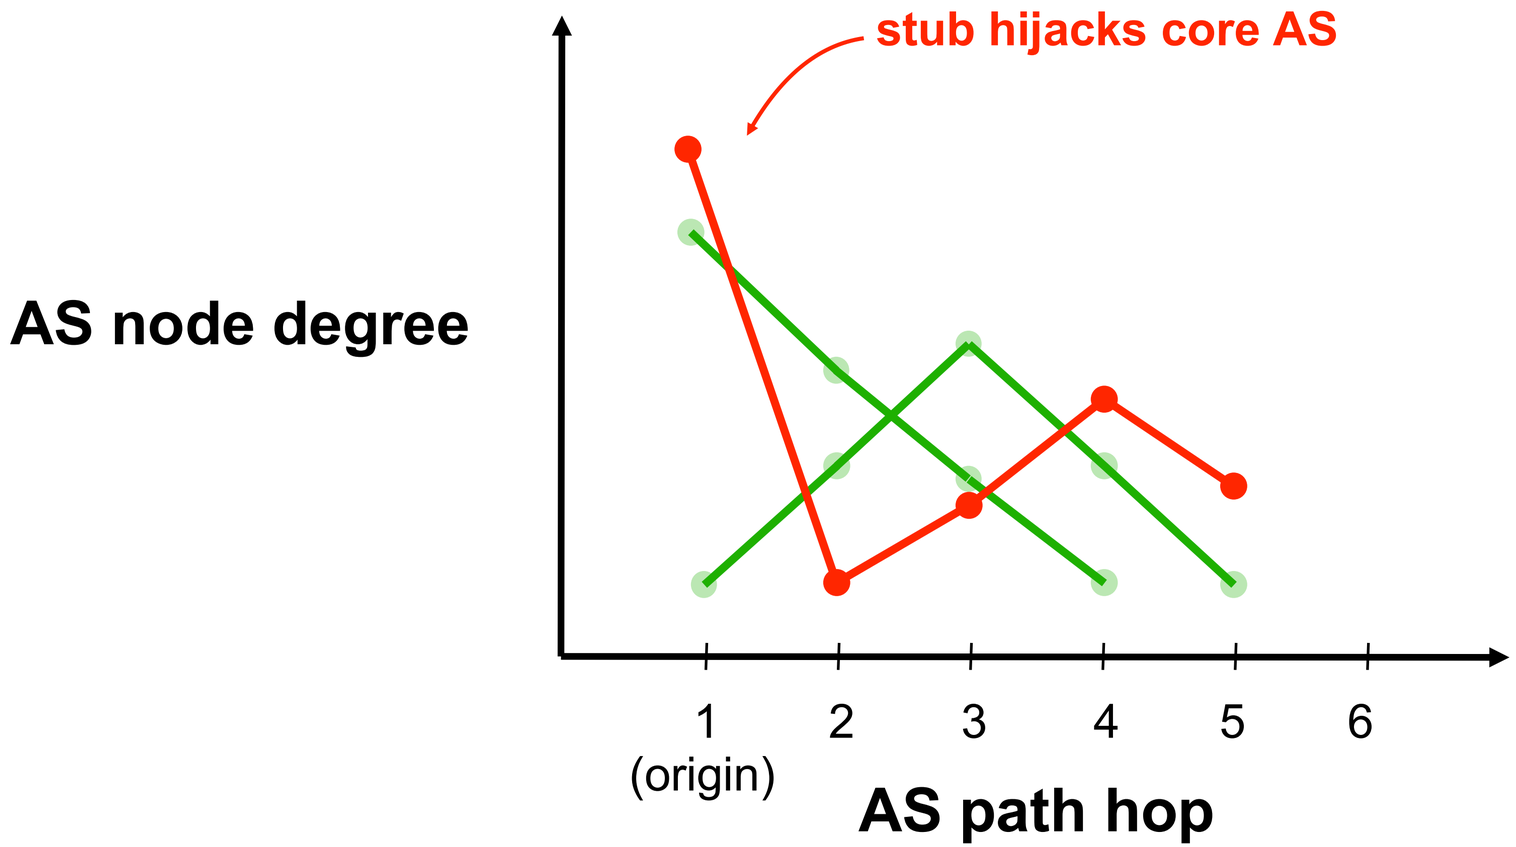

DFOH's fake AS links inference algorithm comprises three steps

RIS/RouteViews
Vantage point

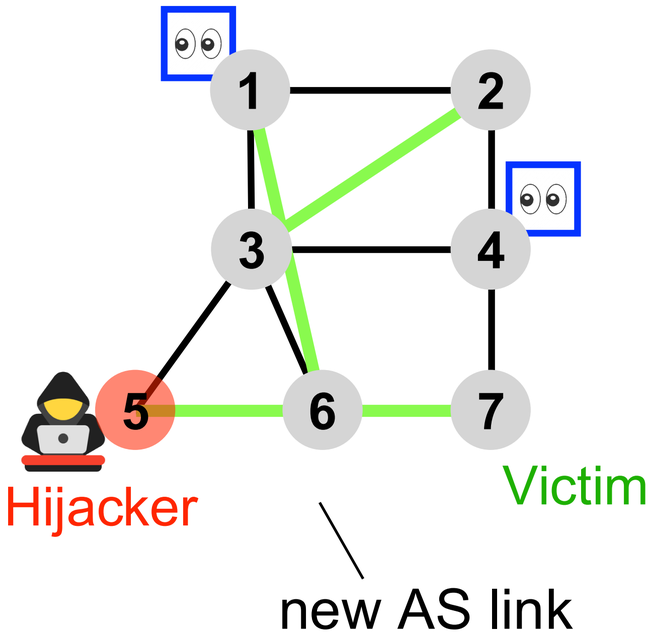

Feature categories:

Bidirectionality

AS-path pattern

Peeringdb

Topological

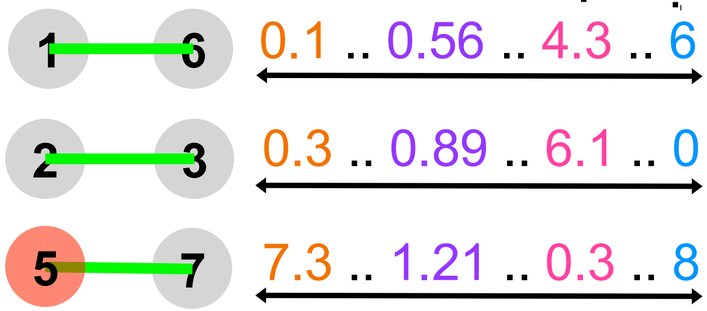

Feature vectors

DFOH's fake AS links inference algorithm comprises three steps

RIS/RouteViews
Vantage point

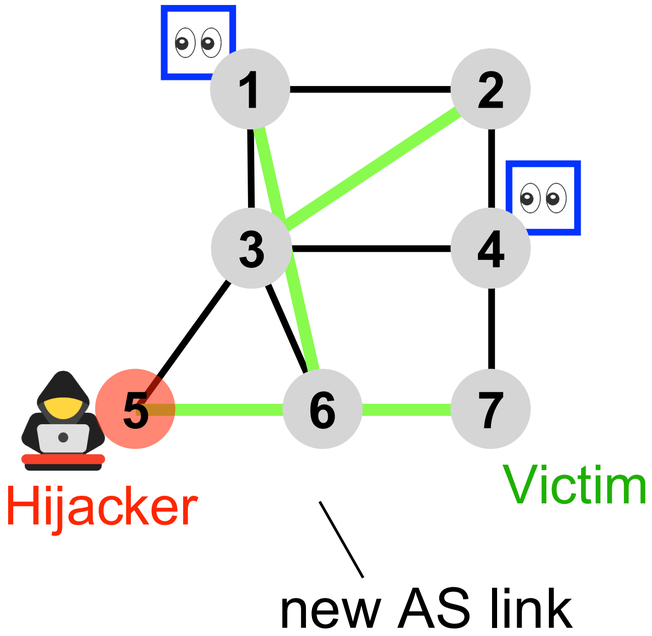

Domain-specific features that compensate each other

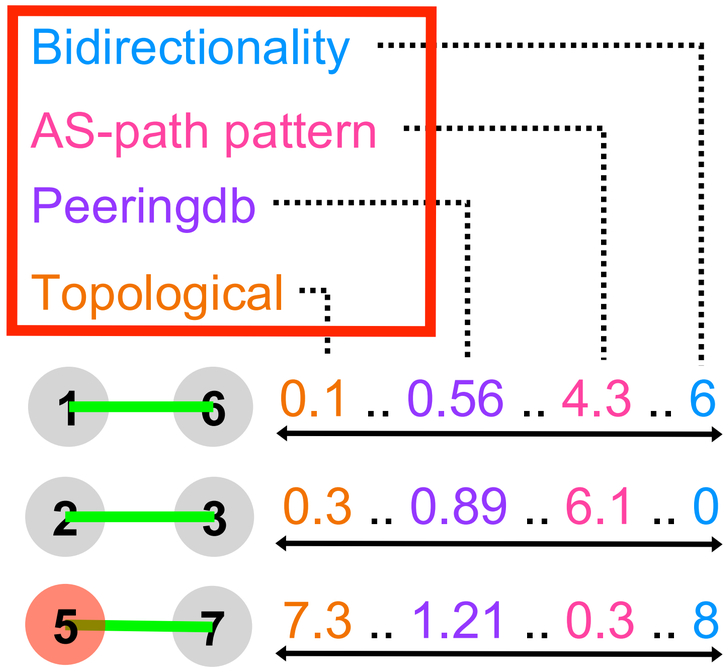

DFOH's fake AS links inference algorithm comprises three steps

RIS/RouteViews Vantage point

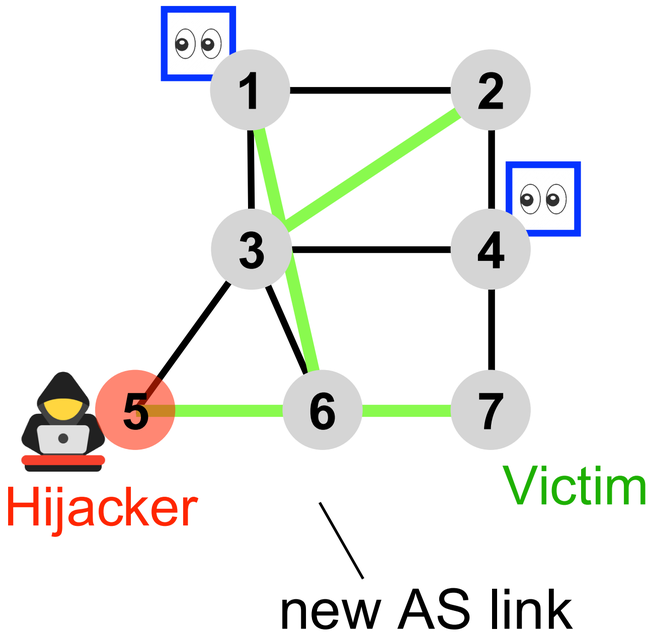

Feature categories:

- Bidirectionality
- AS-path pattern
- Peeringdb
- Topological

Feature vectors

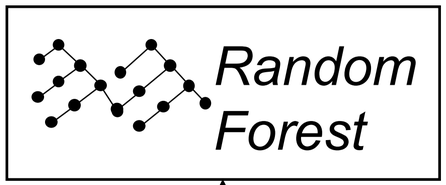

Inference

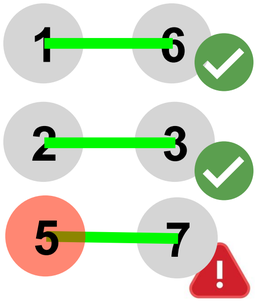

Training

Samples

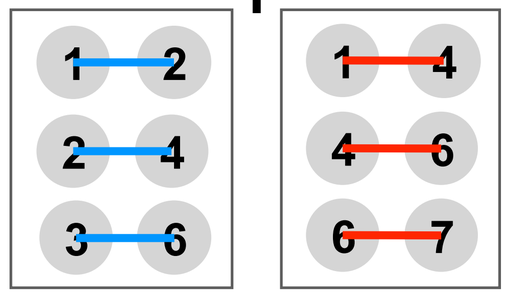

Existing links

Nonexistent links

There are several link prediction frameworks

SEAL (NIPS'18) is one example

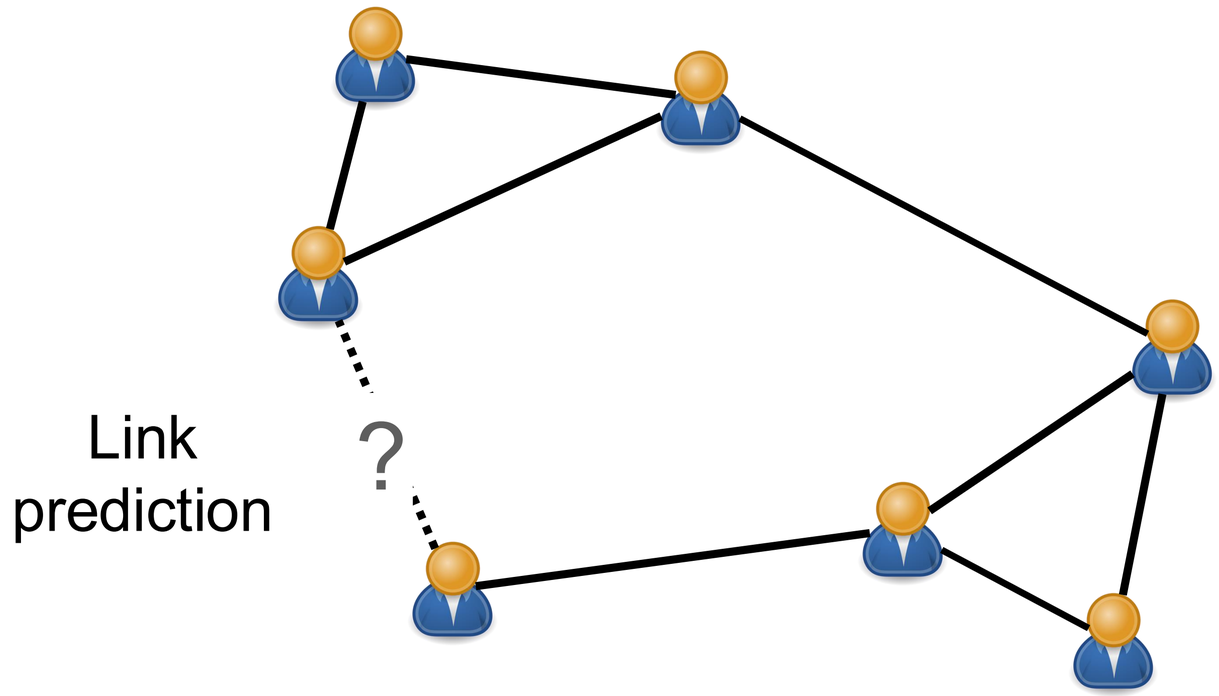

There are several link prediction frameworks
but they do not translate well for detecting fake AS links

Typical hierarchical structure
of the AS-level topology

Few tier1 ASes

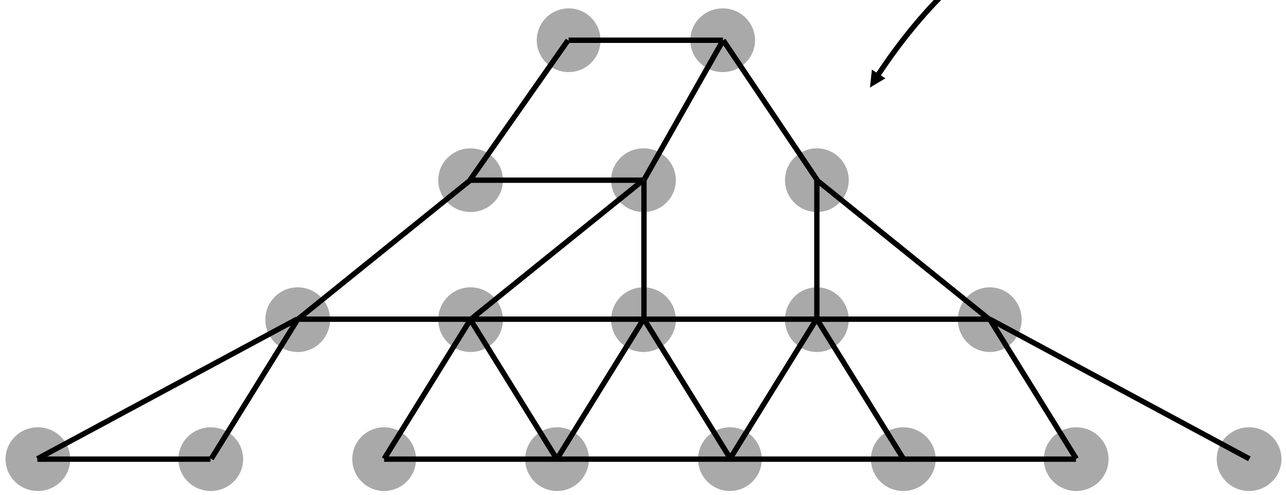

Many stub ASes

Problem: randomly sampling **nonexistent** links makes DFOH skewed towards stub-to-stub links as they are overrepresented

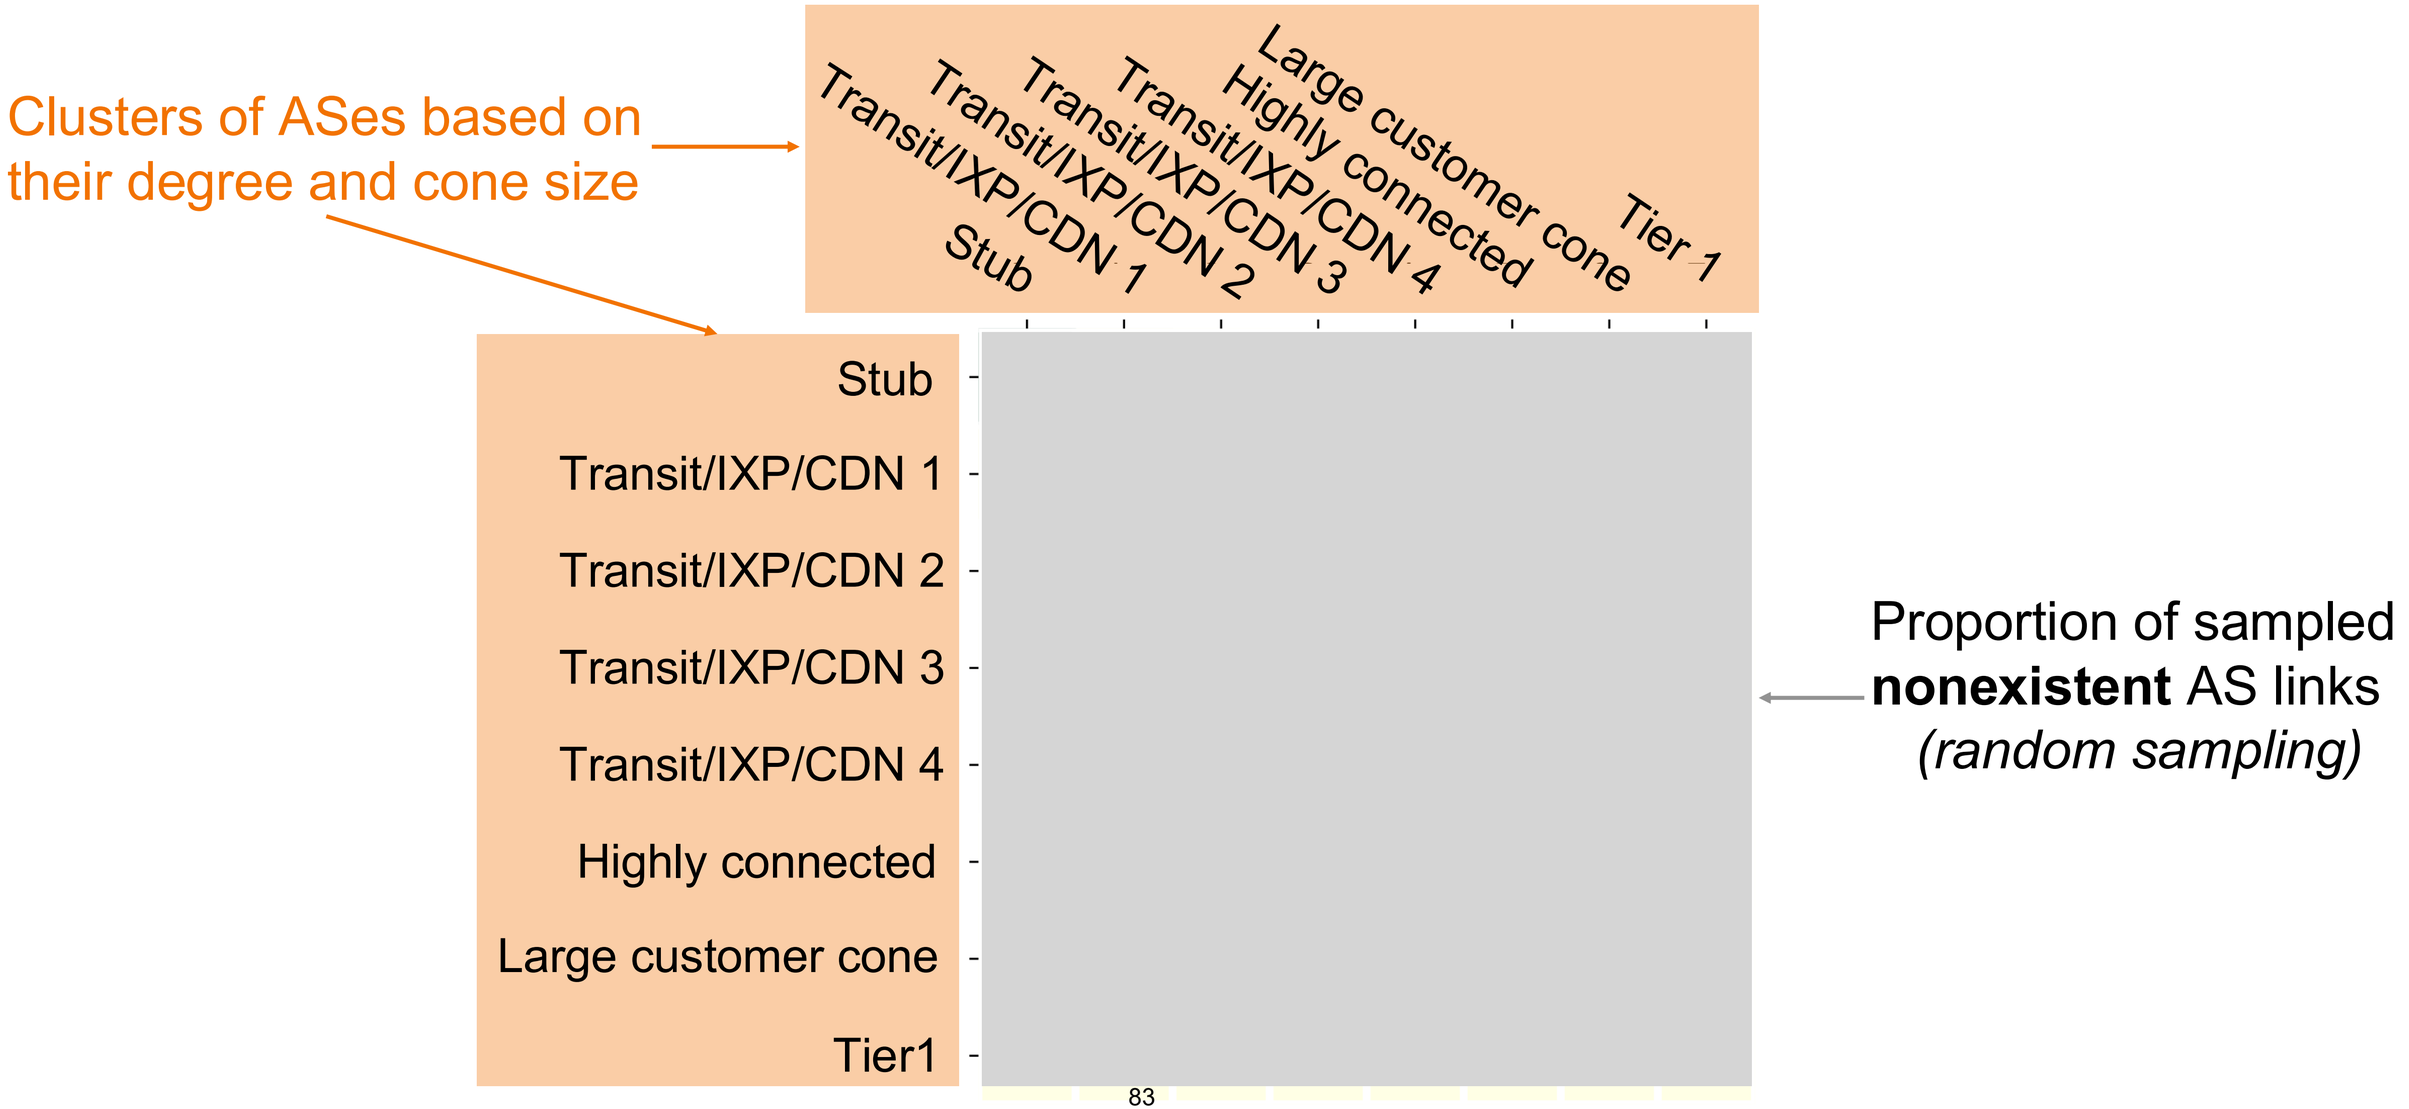

Problem: randomly sampling **nonexistent** links makes DFOH skewed towards stub-to-stub links as they are overrepresented

Clusters of ASes based on their degree and cone size

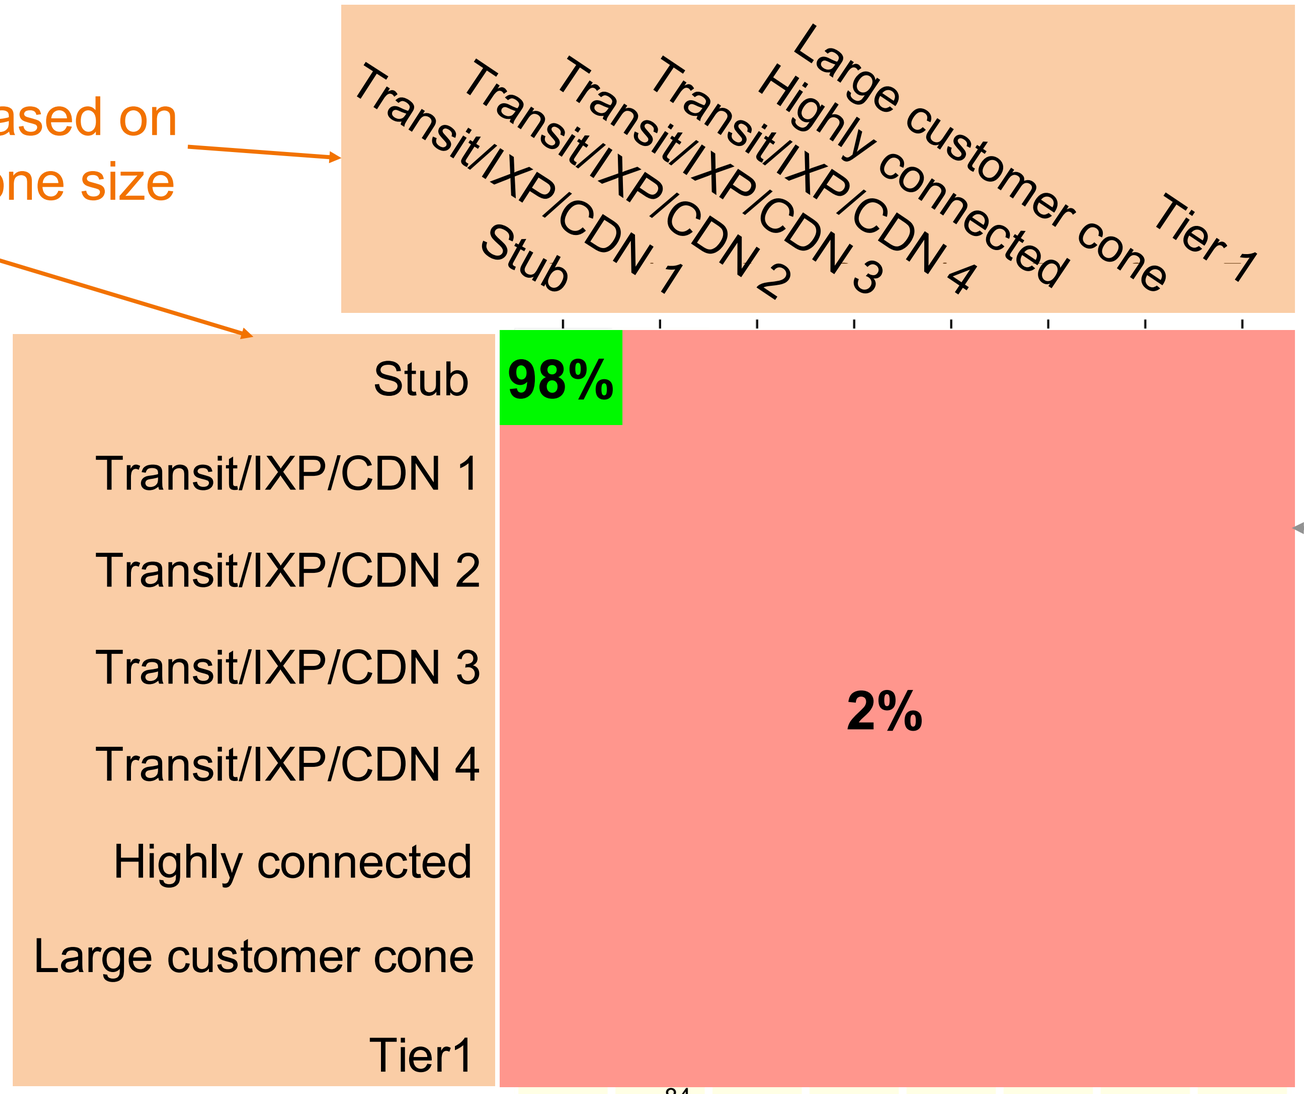

Proportion of sampled **nonexistent** AS links *(random sampling)*

Problem: randomly sampling **nonexistent** links makes DFOH skewed towards stub-to-stub links as they are overrepresented

Clusters of ASes based on their degree and cone size

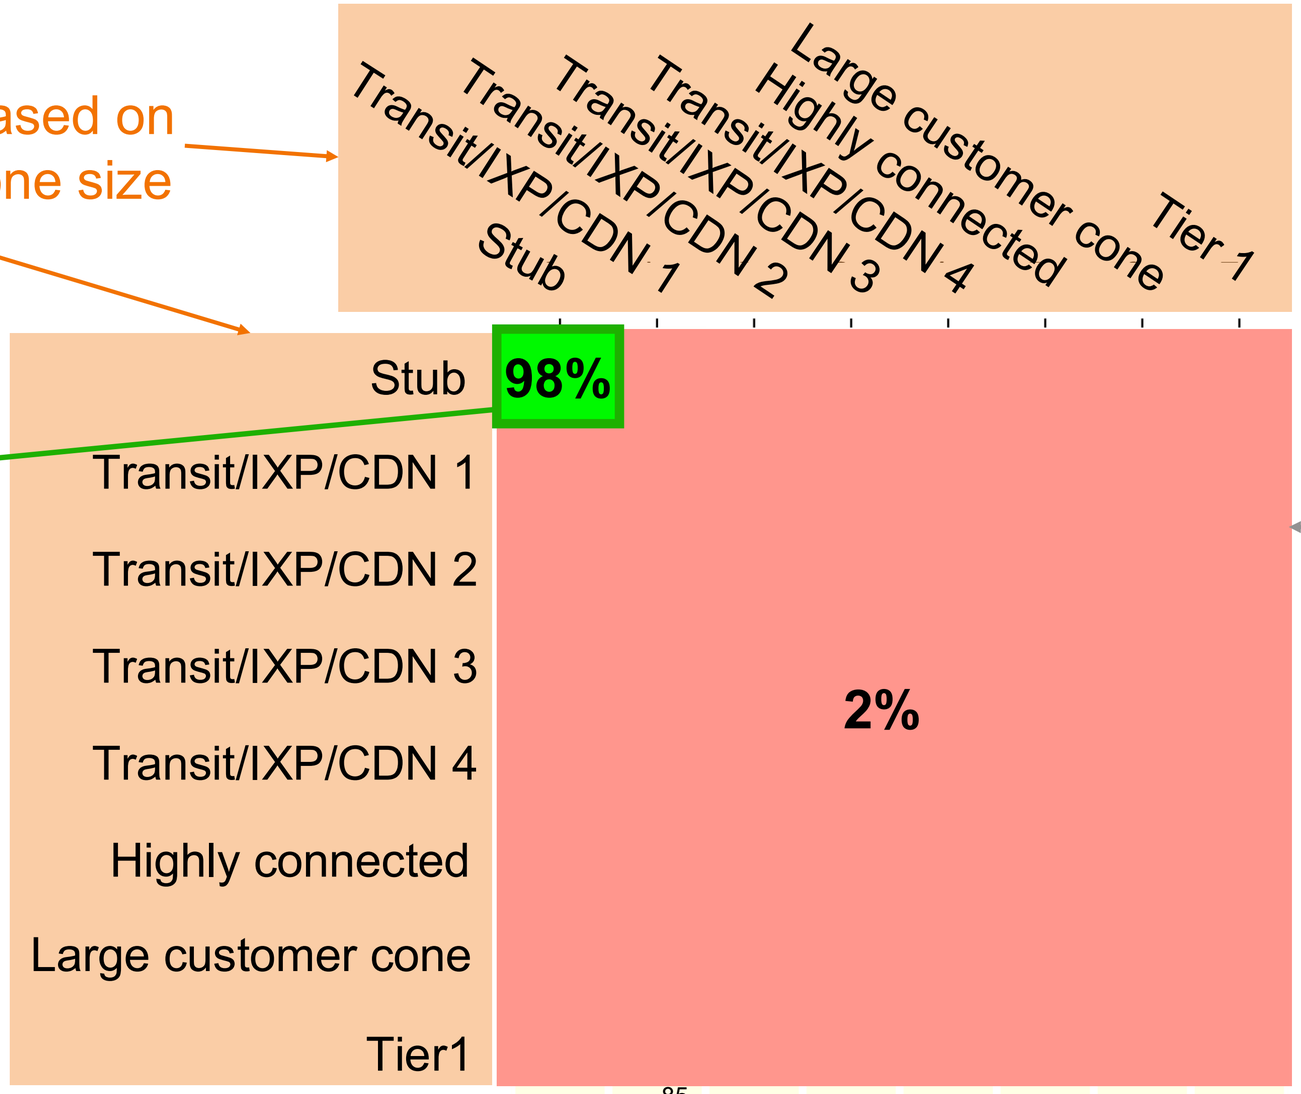

DFOH would perform well on scenarios involving two stubs

Proportion of sampled **nonexistent** AS links *(random sampling)*

Problem: randomly sampling **nonexistent** links makes DFOH skewed towards stub-to-stub links as they are overrepresented

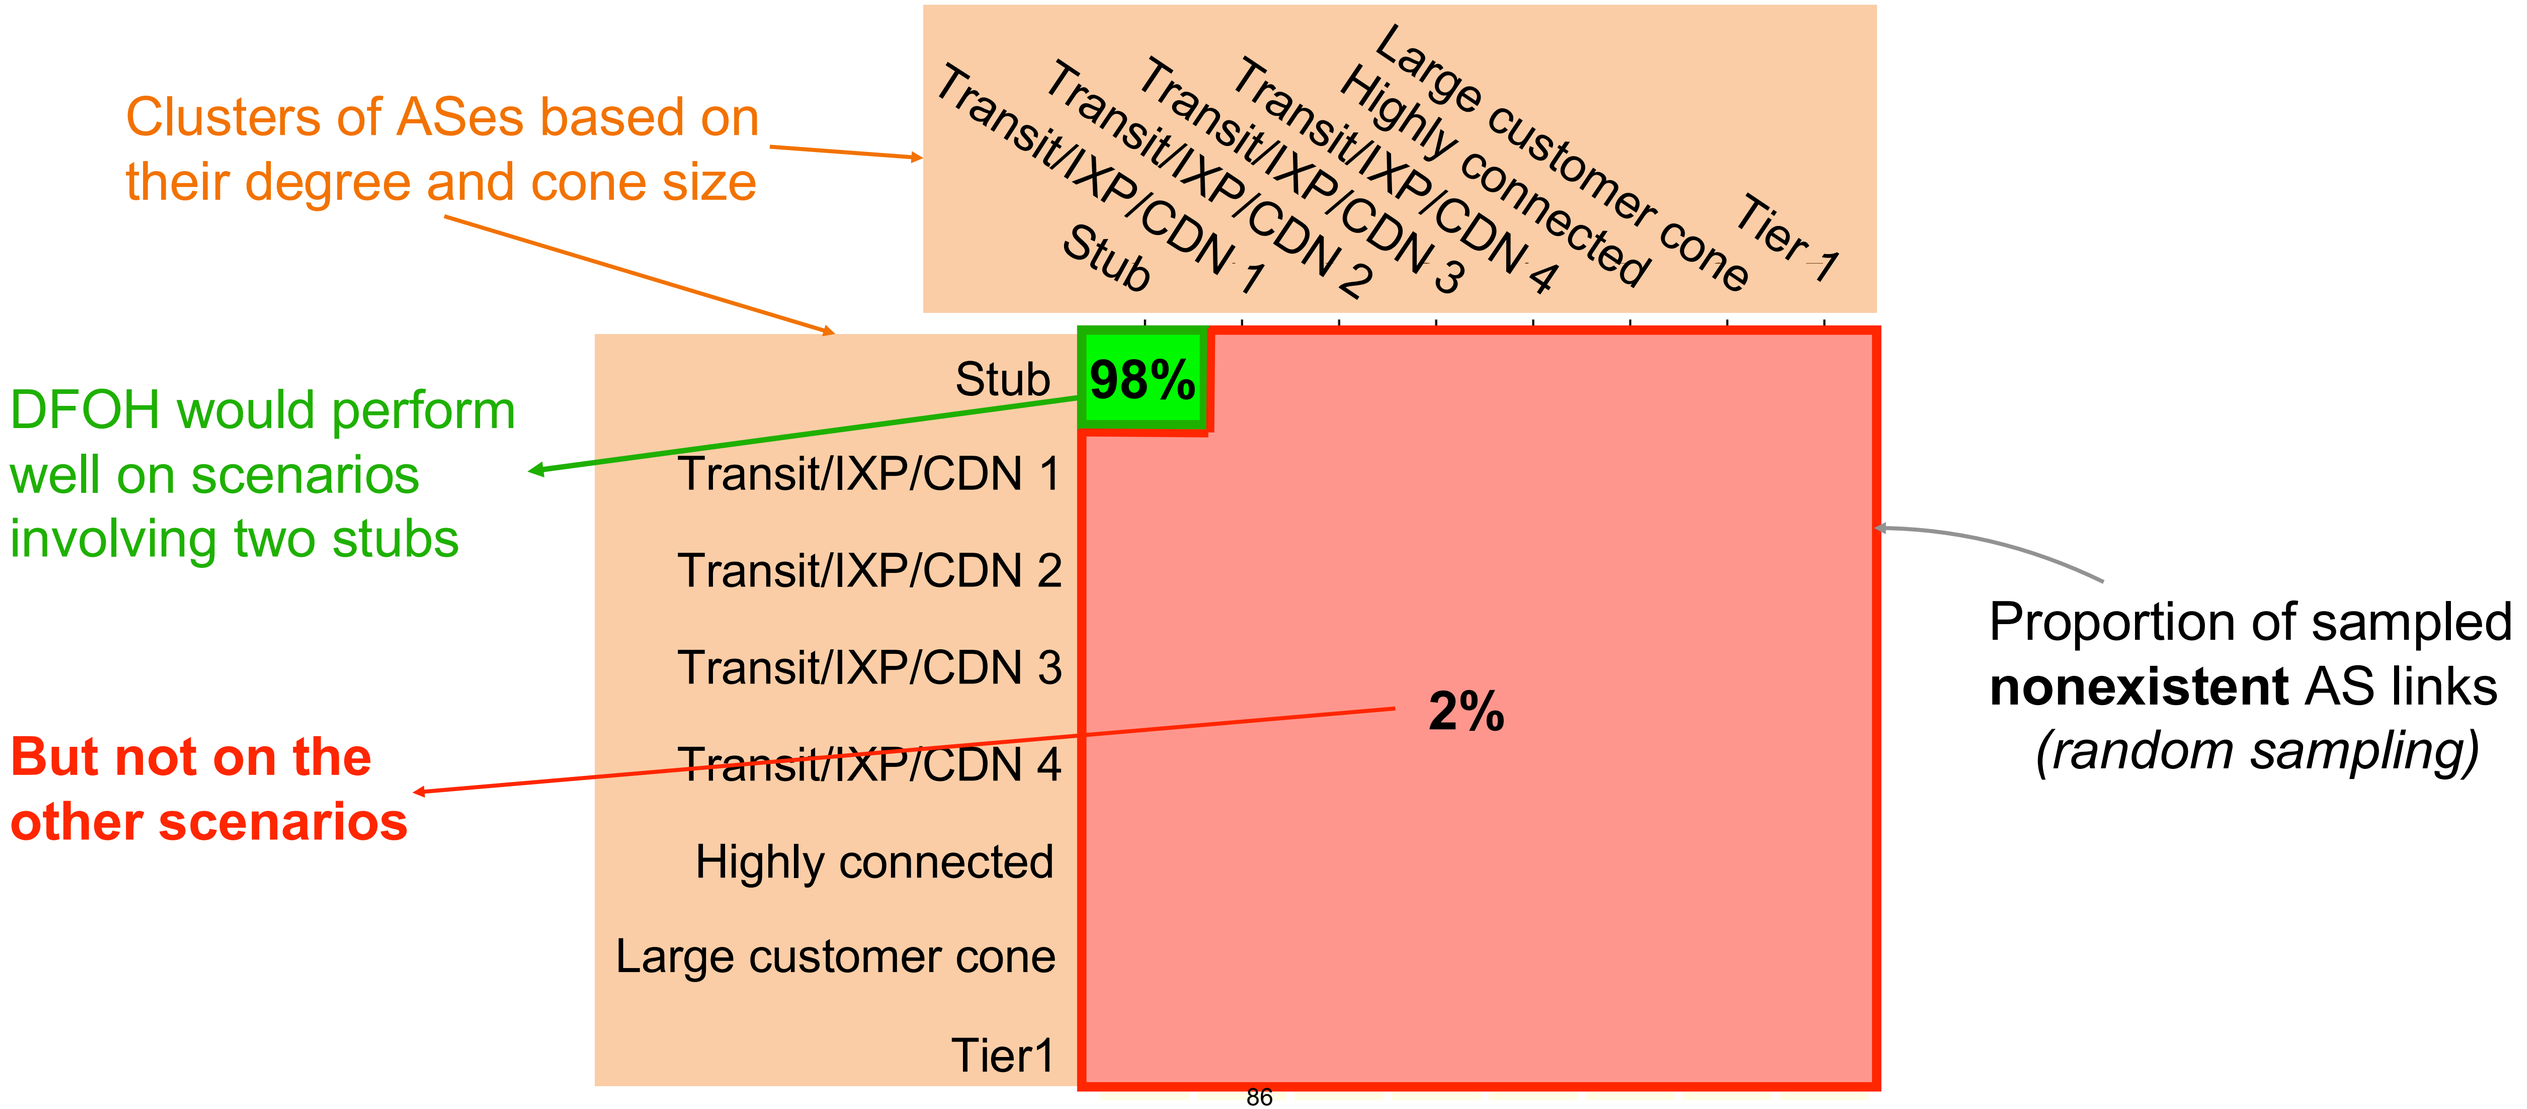

DFOH's fake AS links inference algorithm comprises three steps

Finding New Links → Computing Features → Inferring Hijacks

RIS/RouteViews Vantage point

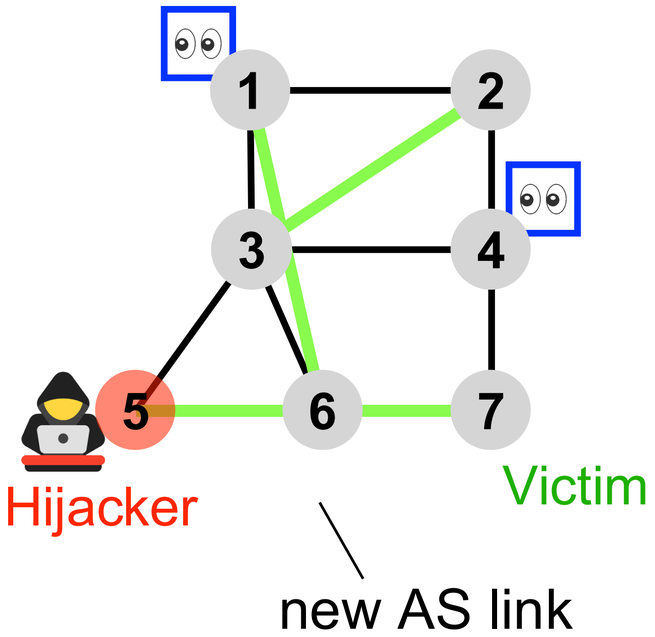

Feature categories:

- Bidirectionality
- AS-path pattern
- Peeringdb
- Topological

Feature vectors

Random Forest

Inference

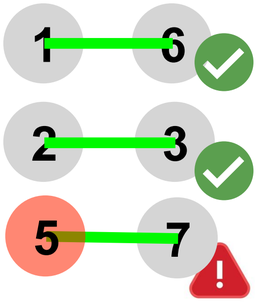

Training

Balanced sampling

- *Stub-Stub*
- *Tier2-Stub*
- *Tier1-Tier2*
- ⋮

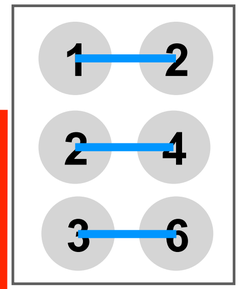

Existing links

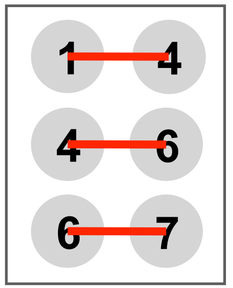

Nonexistent links

Outline

DFOH's main challenge — is to detect fake AS links

DFOH's inference pipeline — relies on domain-specific knowledge and a tailored link prediction framework

DFOH's inferences are accurate — in every attack scenario

DFOH *is* up and running

We evaluate ***DFOH*** on artificially created forged-origin hijacks
as there is no ground truth at scale

Methodology:

We take existing AS paths
and prepend a new origin to create a new link

We take 9k cases where the new link exists (*legitimate* or "*negative" cases*)
and 9k cases where the new link does not exist (*suspicious or "positive" cases*)

We evaluate ***DFOH*** on artificially created forged-origin hijacks as there is no ground truth at scale

Methodology:

We take existing AS paths
and prepend a new origin to create a new link

We take 9k cases where the new link exists (*legitimate* or "*negative" cases*)
and 9k cases where the new link does not exist (*suspicious or "positive" cases*)

We focus on the **True Positive Rate** (TPR)
and the **False Positive Rate** (FPR)

DFOH is accurate upon every attack scenario

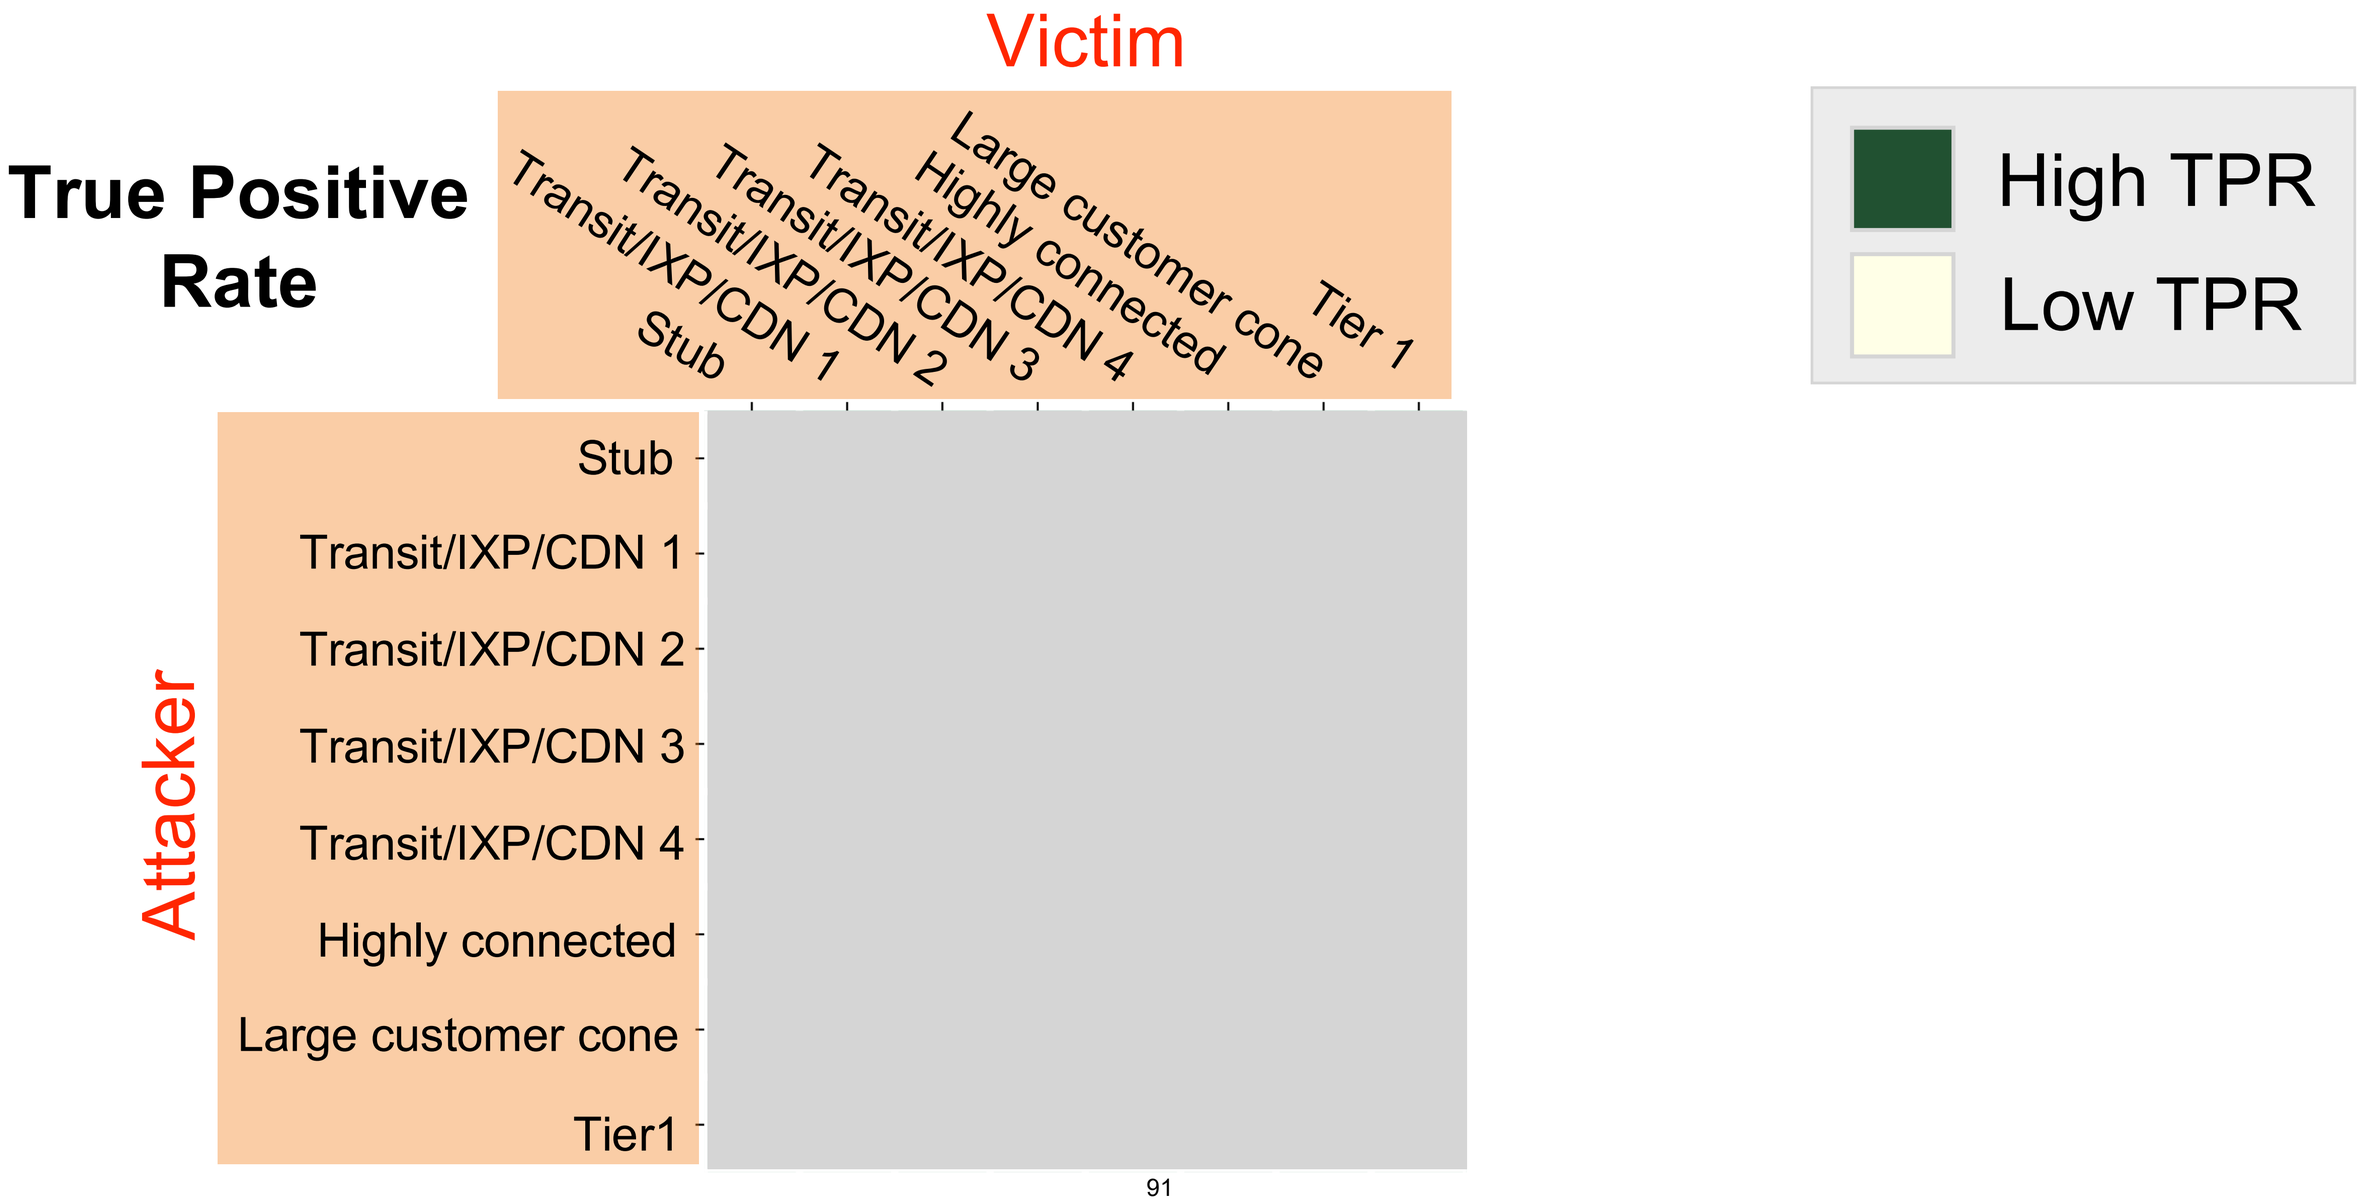

DFOH is accurate upon every attack scenario

DFOH is accurate upon every attack scenario

True Positive Rate

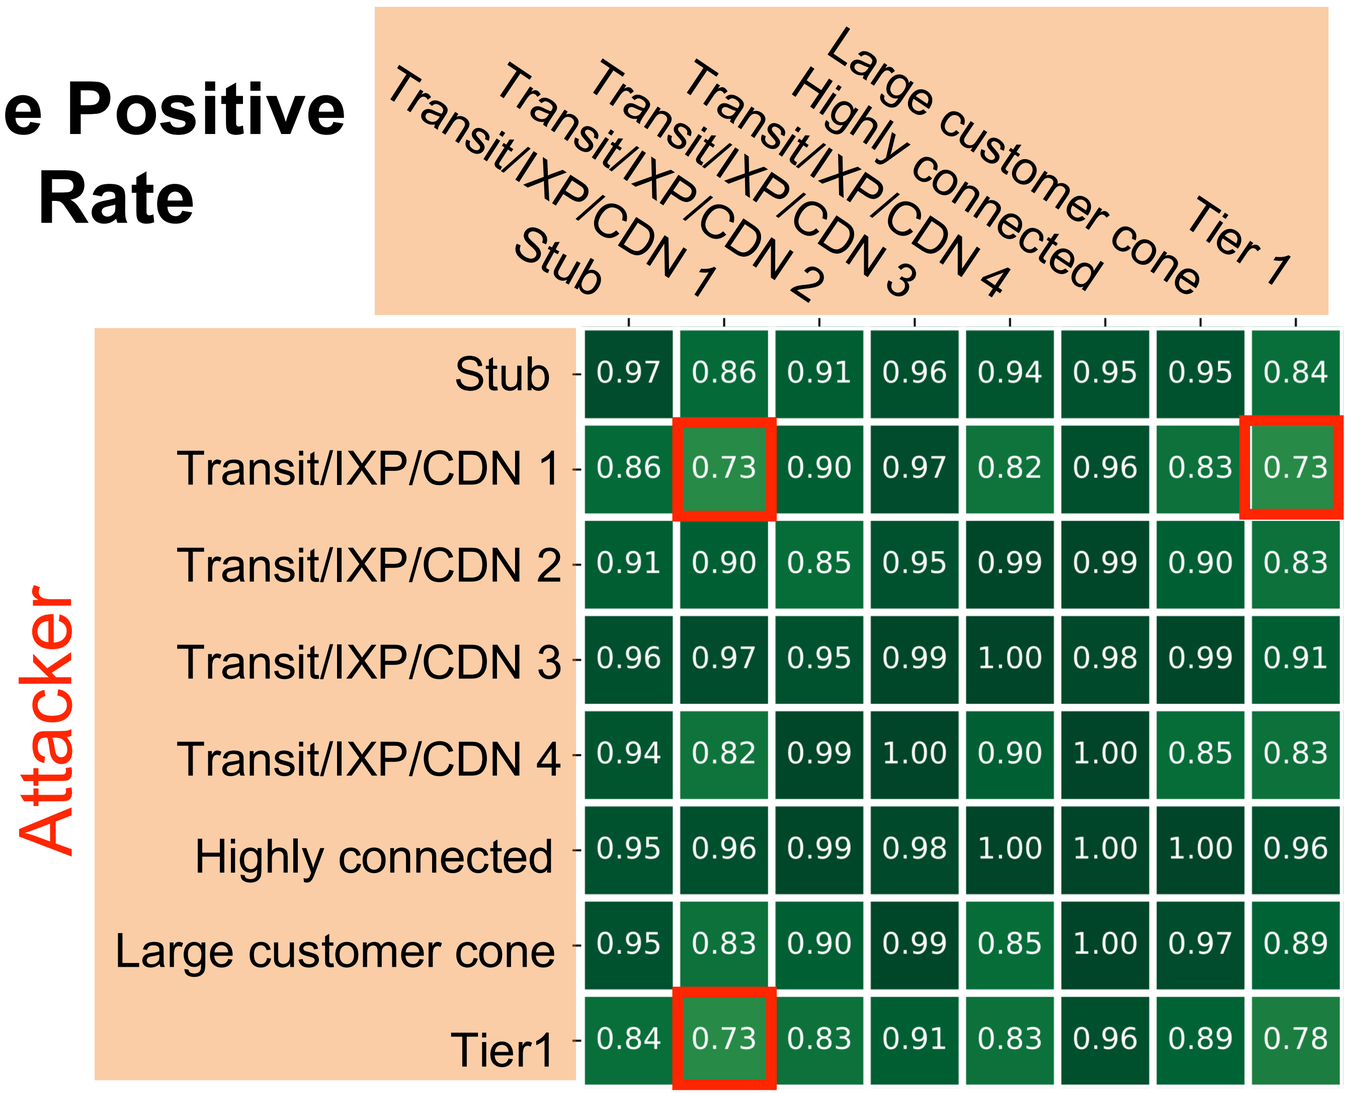

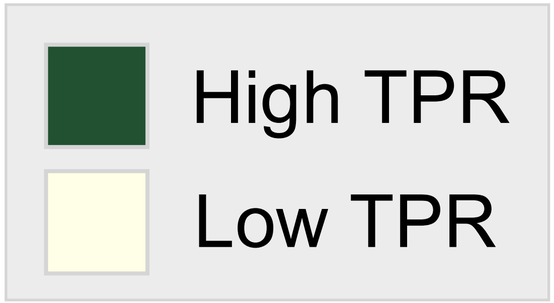

The minimum TPR is 0.73

DFOH is accurate upon every attack scenario

False Positive Rate

DFOH is accurate upon every attack scenario

False Positive Rate

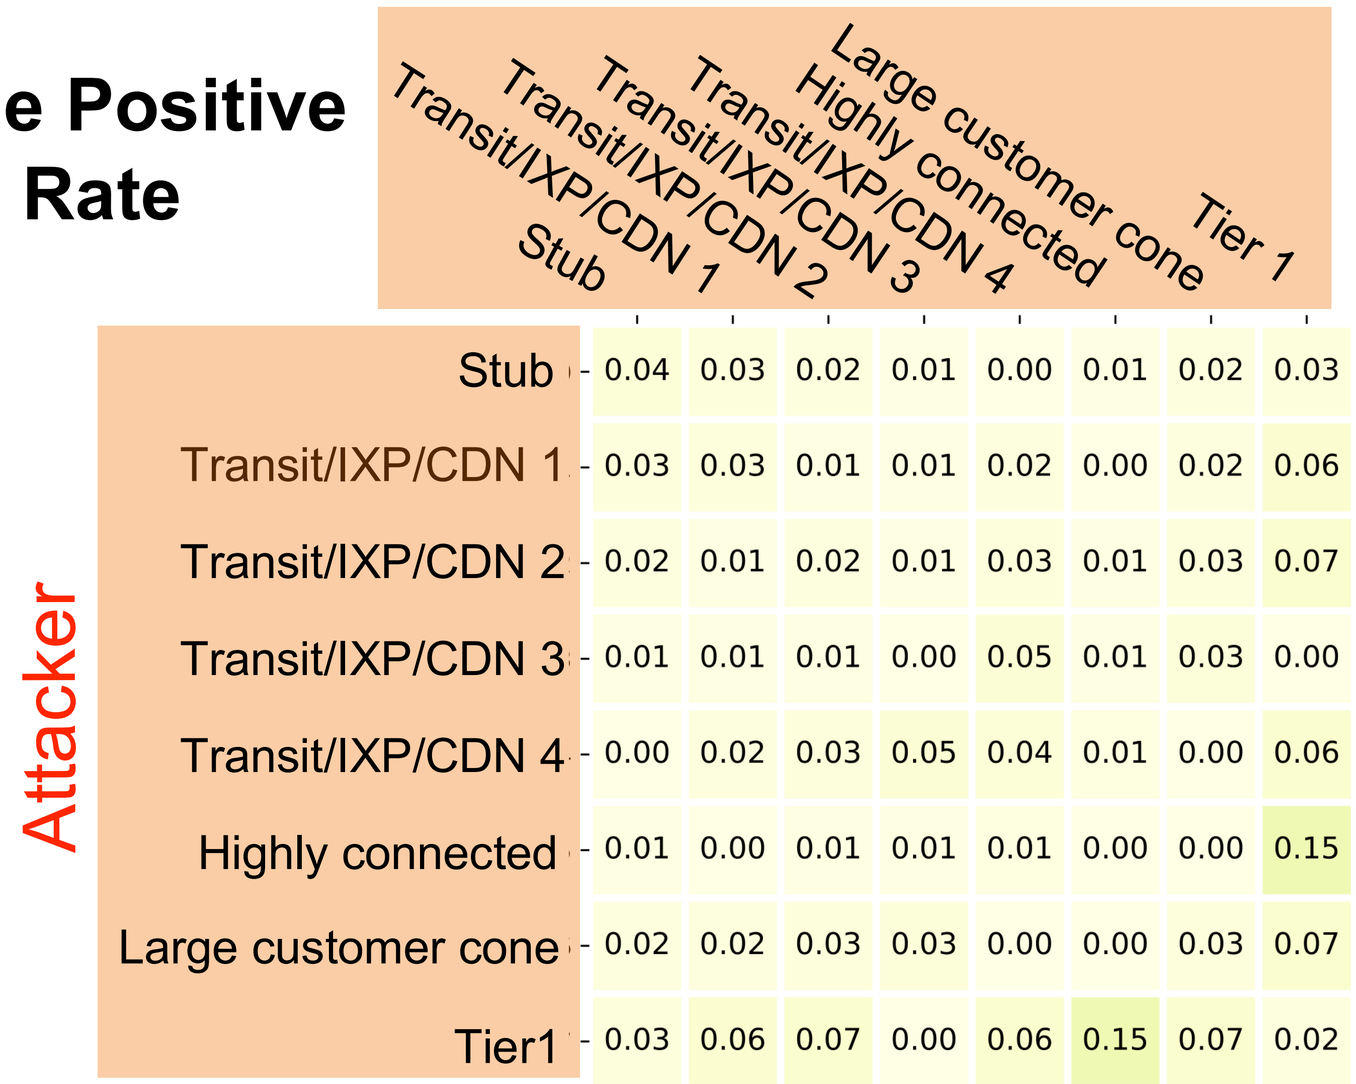

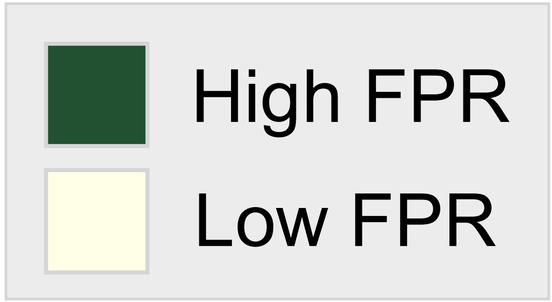

DFOH is accurate upon every attack scenario

False Positive Rate

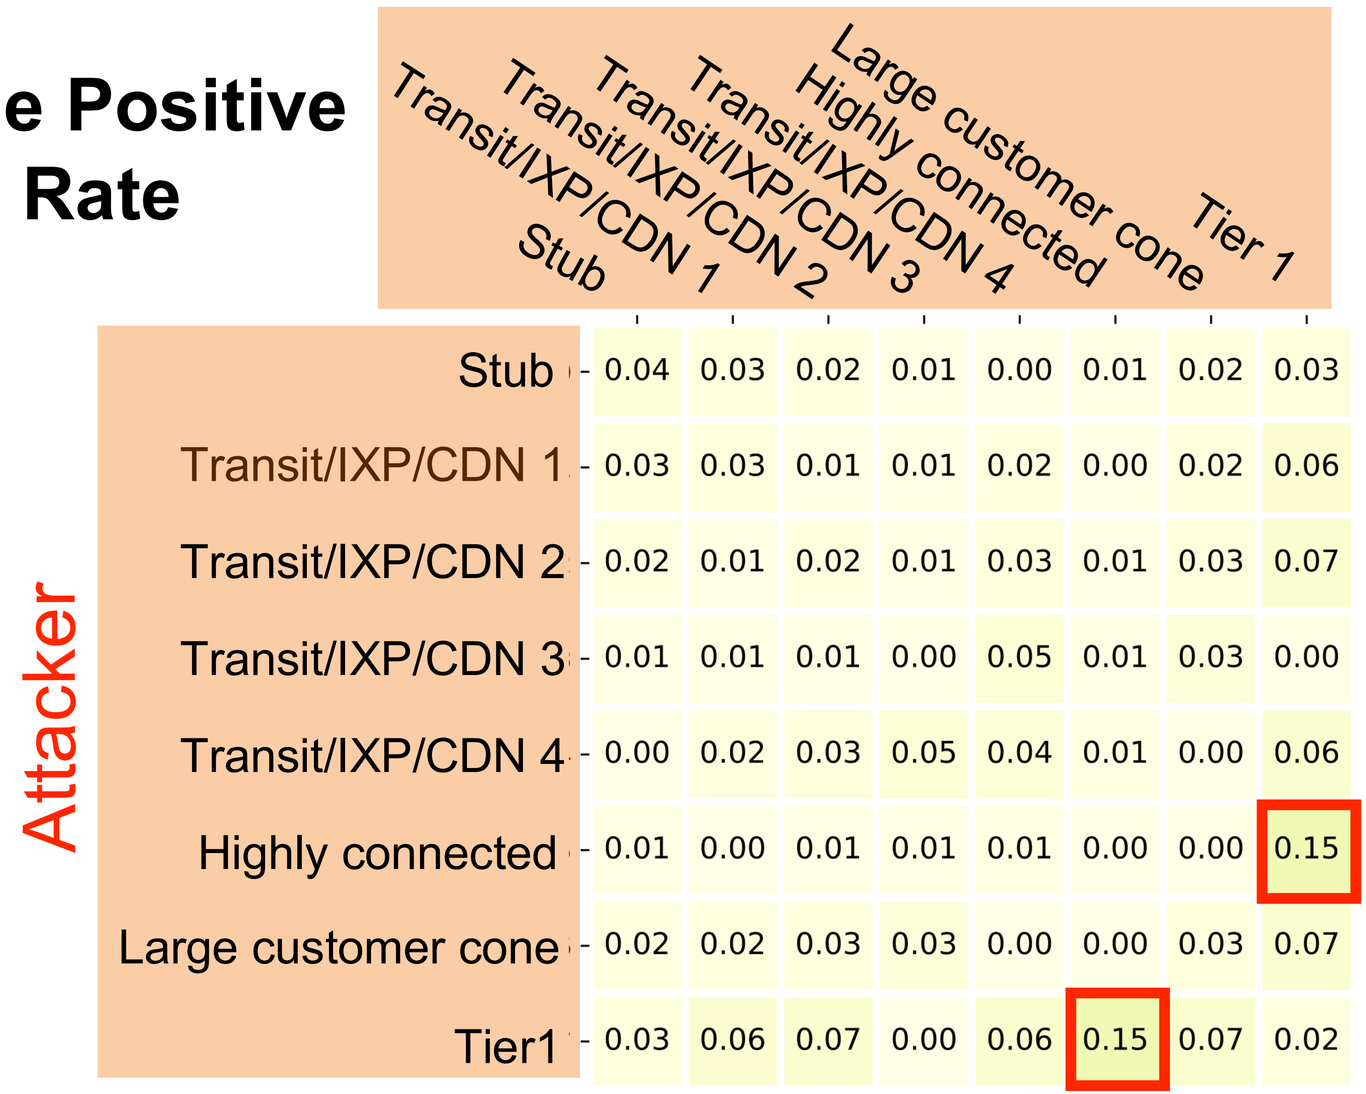

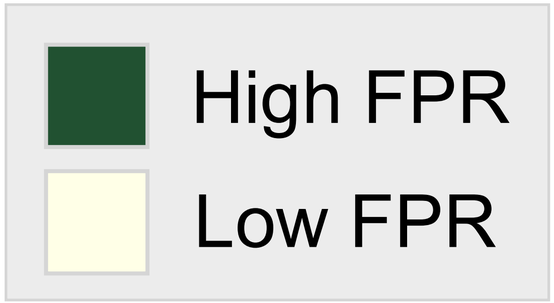

The maximum FPR is 0.15

Outline

DFOH's main challenge is to detect fake AS links

DFOH's inference pipeline discriminates fake AS links from the real ones

DFOH's inferences are accurate in every attack scenario

DFOH *is* up and running and useful for operators

DFOH is useful and practical for network operators

Useful: DFOH detects the two known forged-origin BGP hijacks
(the klayswap and cbridge attacks)

Practical: DFOH only reports zero or one case every month for 99.8% of the ASes
(worse case is 15 cases)

DFOH runs at https://dfoh.uclouvain.be

The next generation of BGP data collection platform

Best paper SIGCOMM 2024

Thomas Alfroy
Thomas Holterbach
Thomas Krenc
K. C. Claffy
Cristel Pelsser

The Internet topology observes a hierarchical structure where some ASes pay a provider to forward their traffic

Customer-to-provider

Peer-to-peer

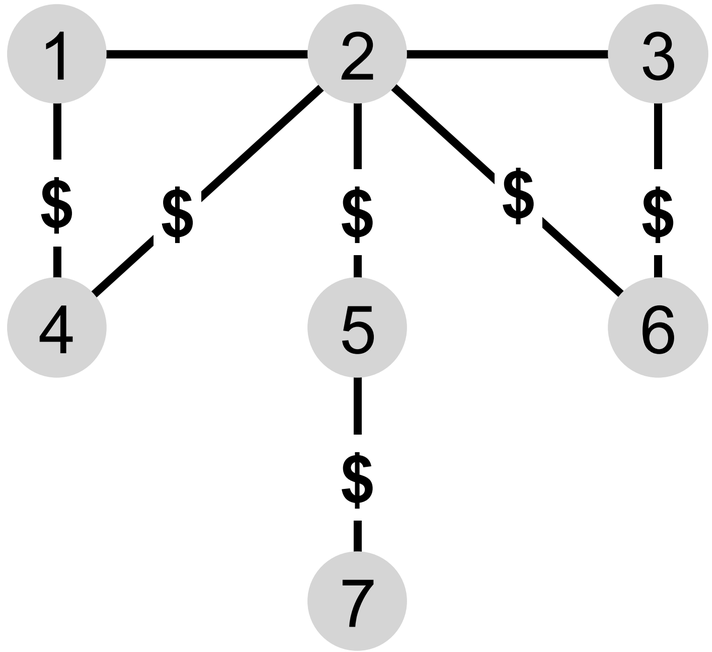

Routing is done according to these economical relationships

A route to prefix is propagated hop-by-hop
until every AS in the topology has a route toward this prefix

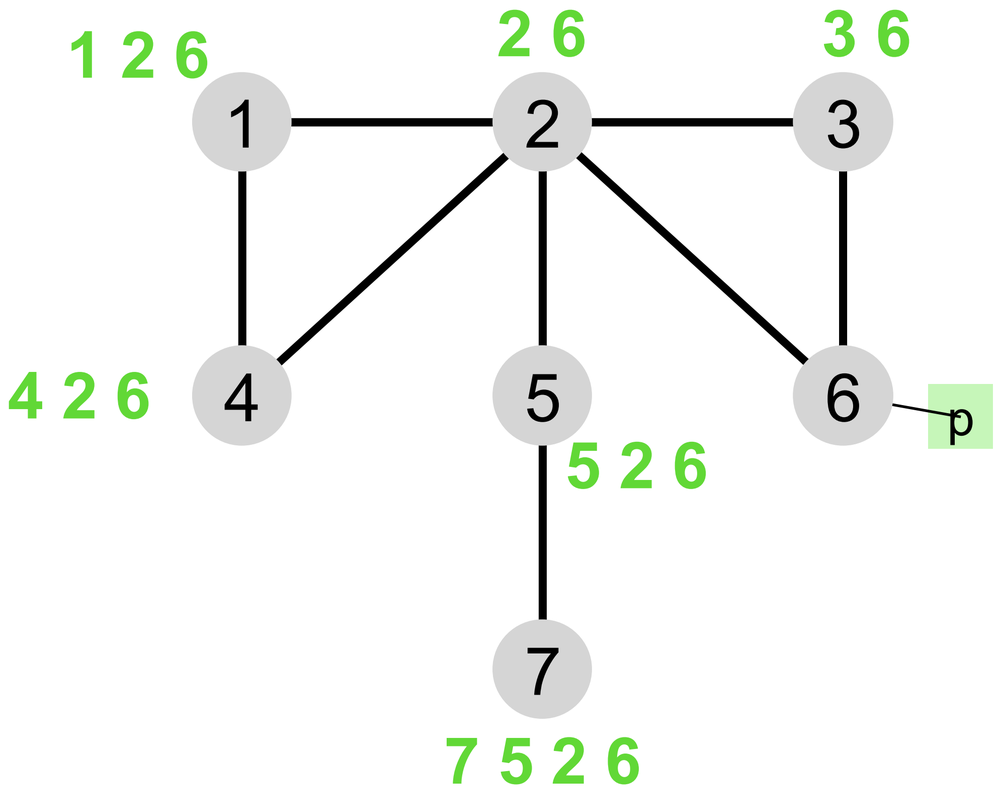

Unfortunately, the Internet is a wild place where many events can disrupt its smooth operation

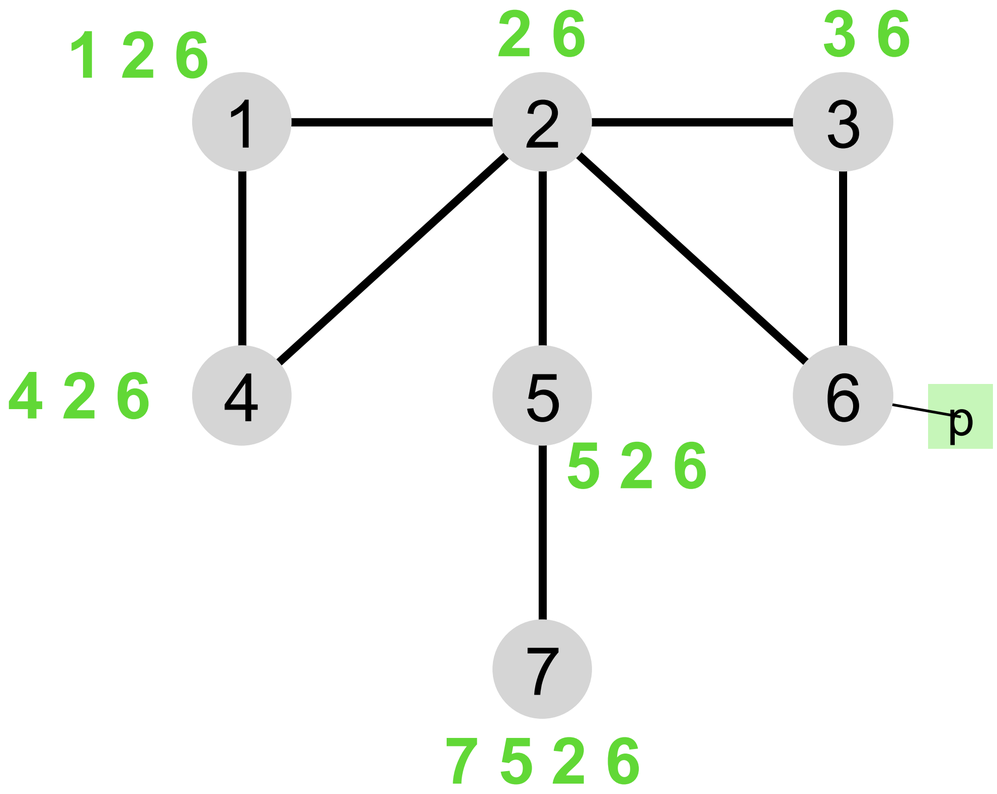

Unfortunately, the Internet is a wild place where many event can disrupt its smooth operation

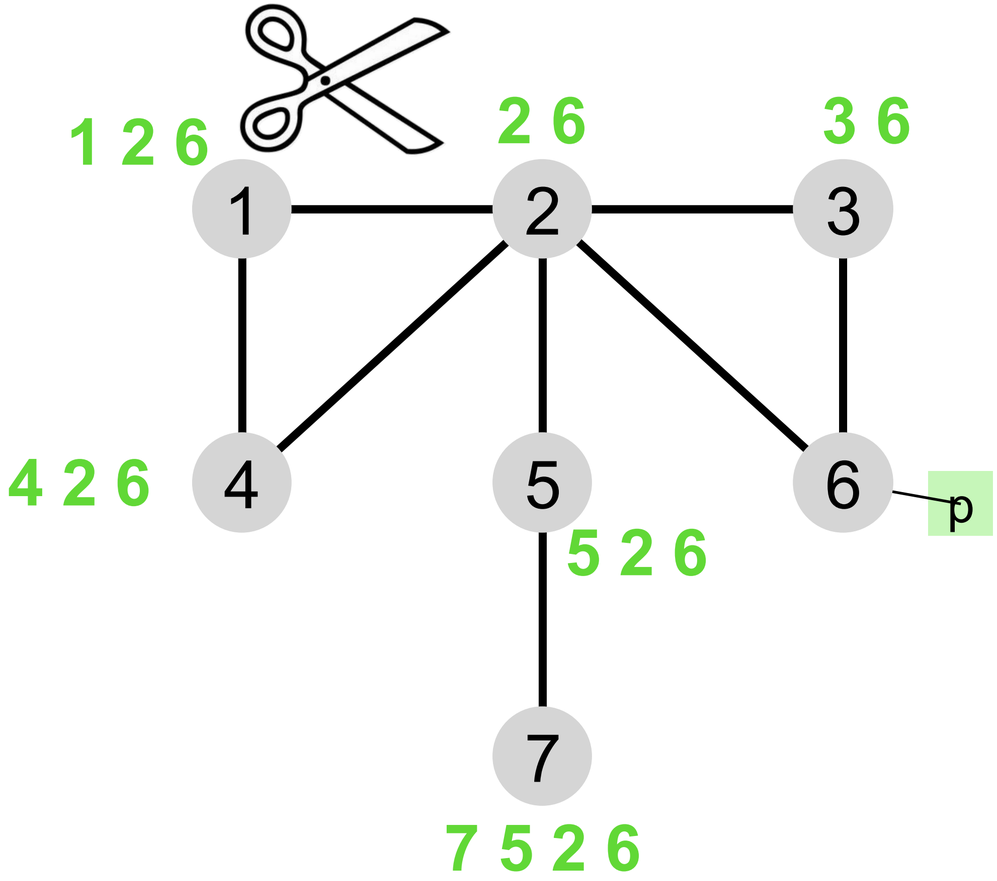

Malicious actors can disrupt the smooth operation of BGP

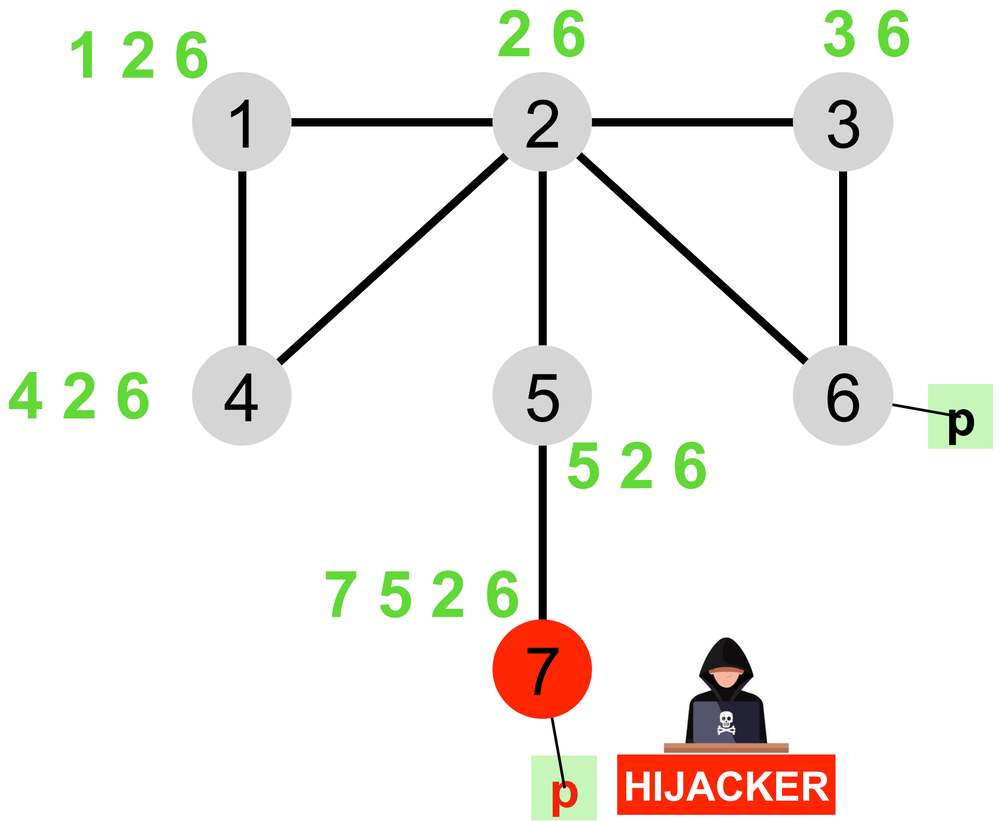

Malicious actors can disrupt the smooth operation of BGP

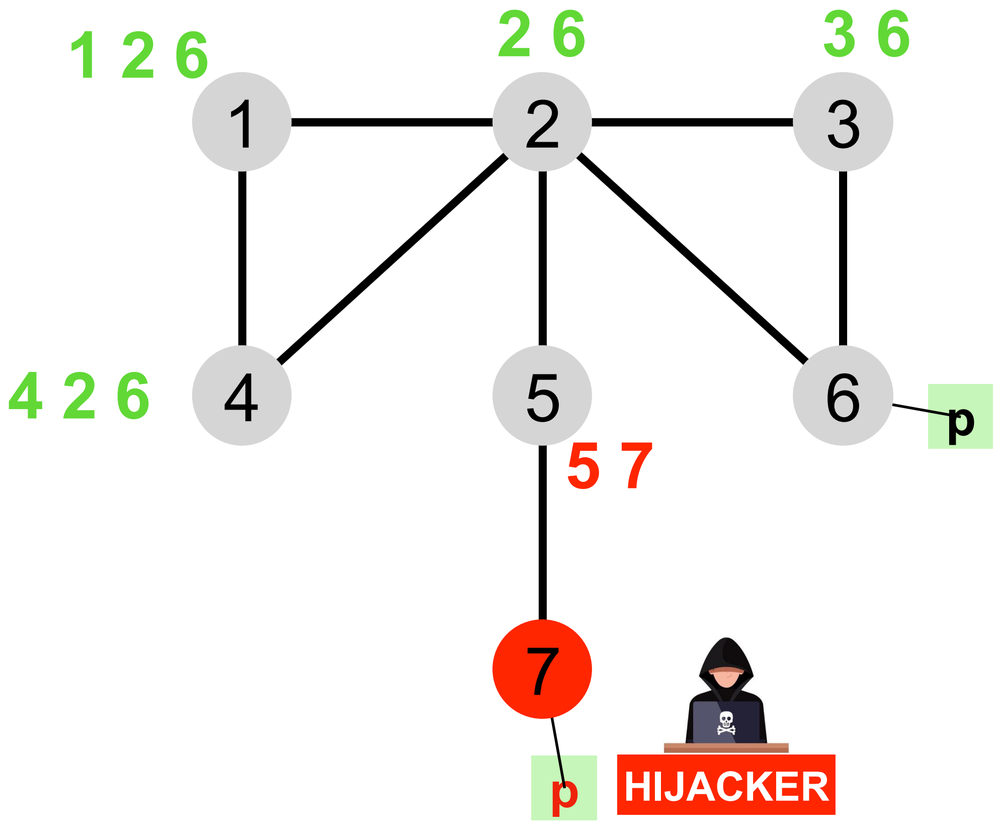

Malicious actors can disrupt the smooth operation of BGP

The victim of the attack will never notice it

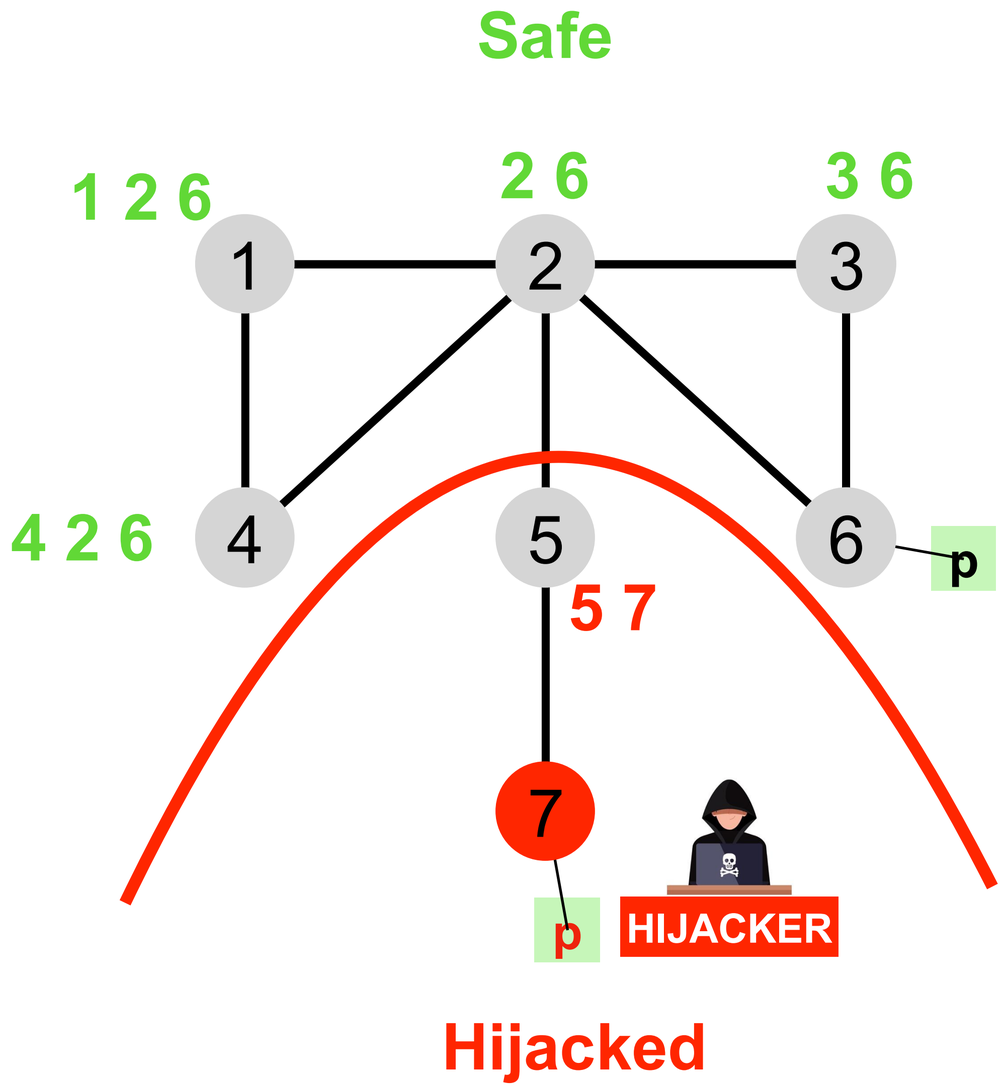

Fortunately, there exist some BGP monitoring tools

Research projects

Commercial tools

These tools rely on the collected BGP data

BGP routes are collected by public BGP data collection platforms, such as RIPE RIS or RouteViews

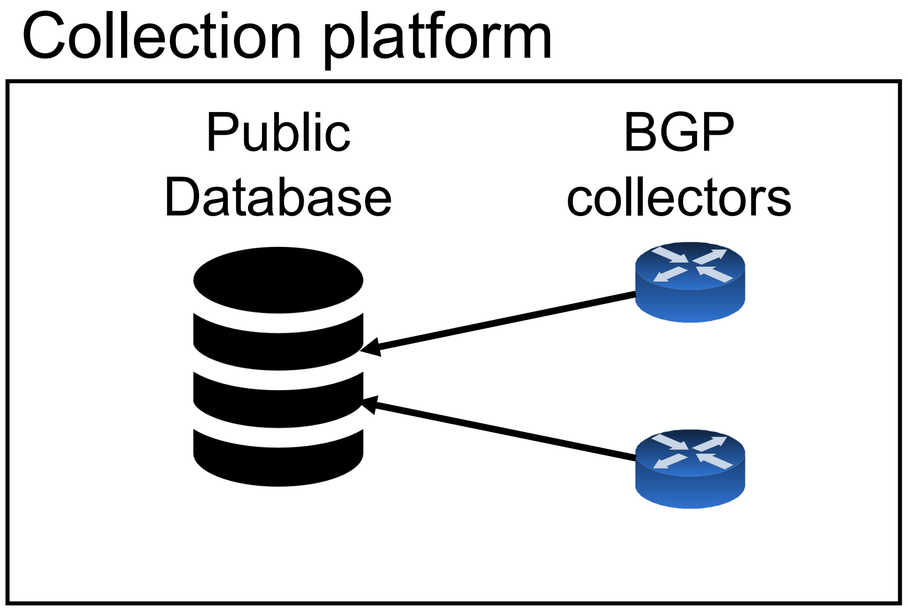

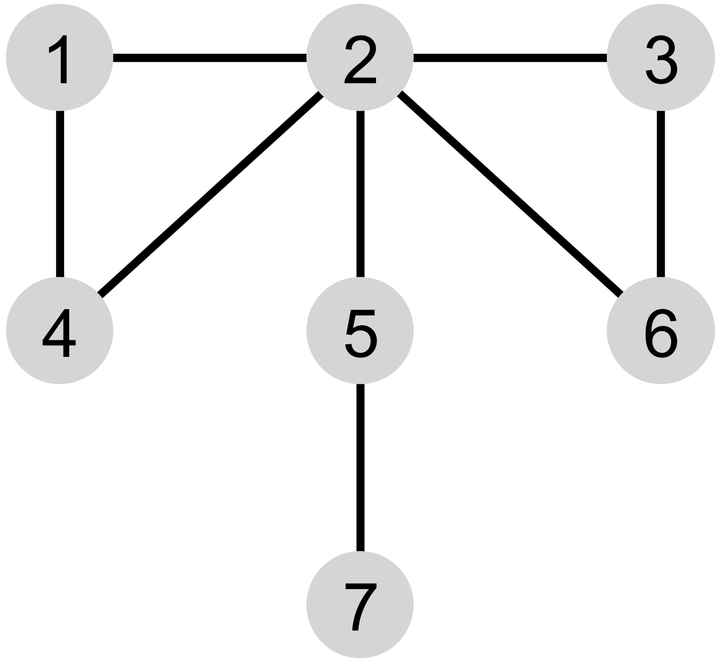

BGP routes are collected by public BGP data collection platforms, such as RIPE RIS or RouteViews

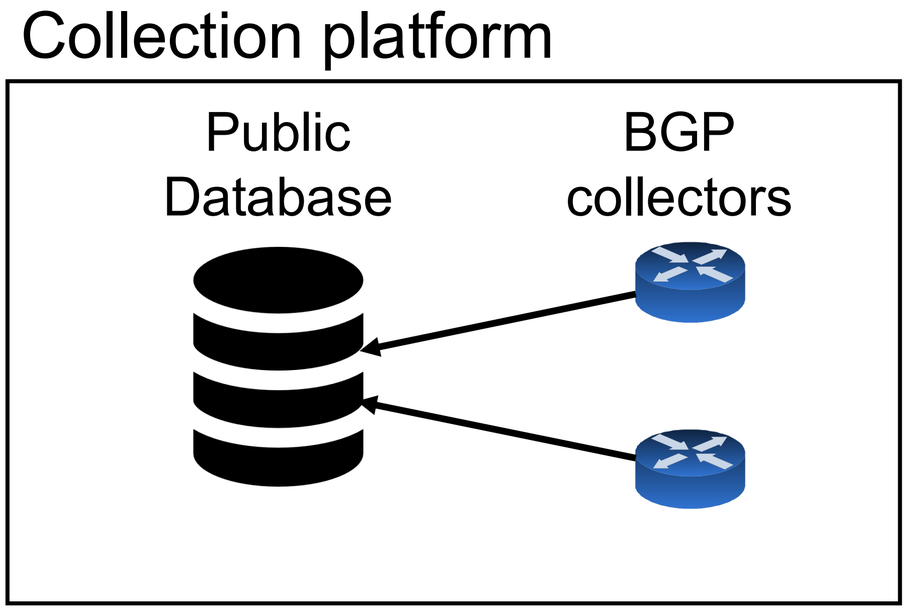

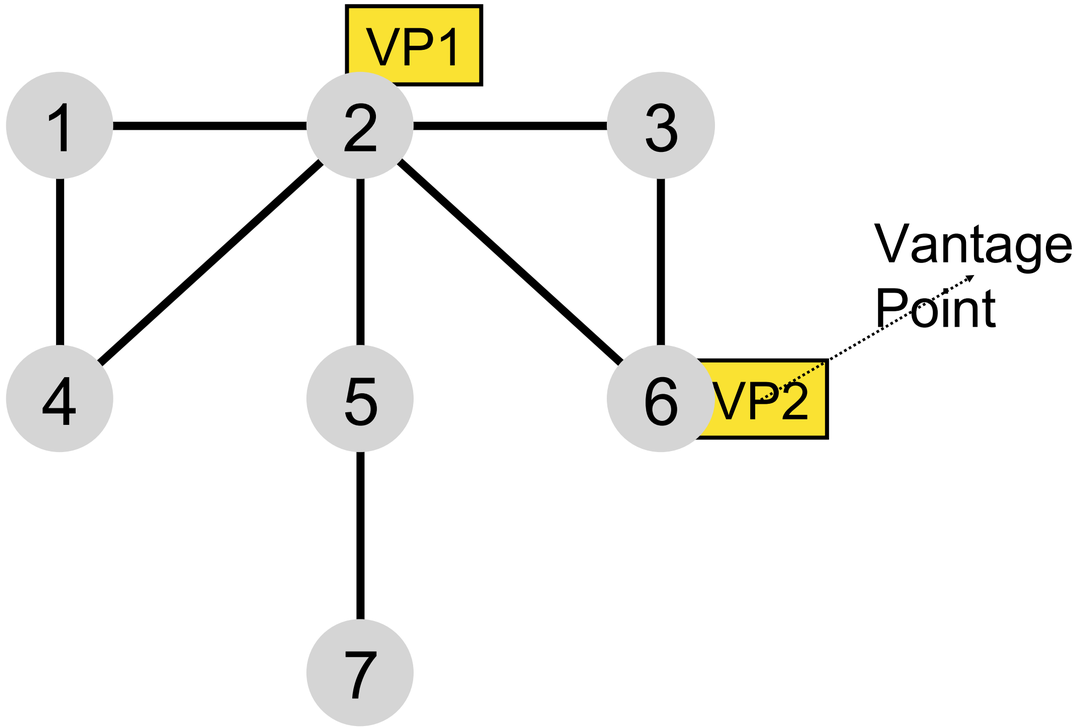

BGP routes are collected by public BGP data collection platforms, such as RIPE RIS or RouteViews

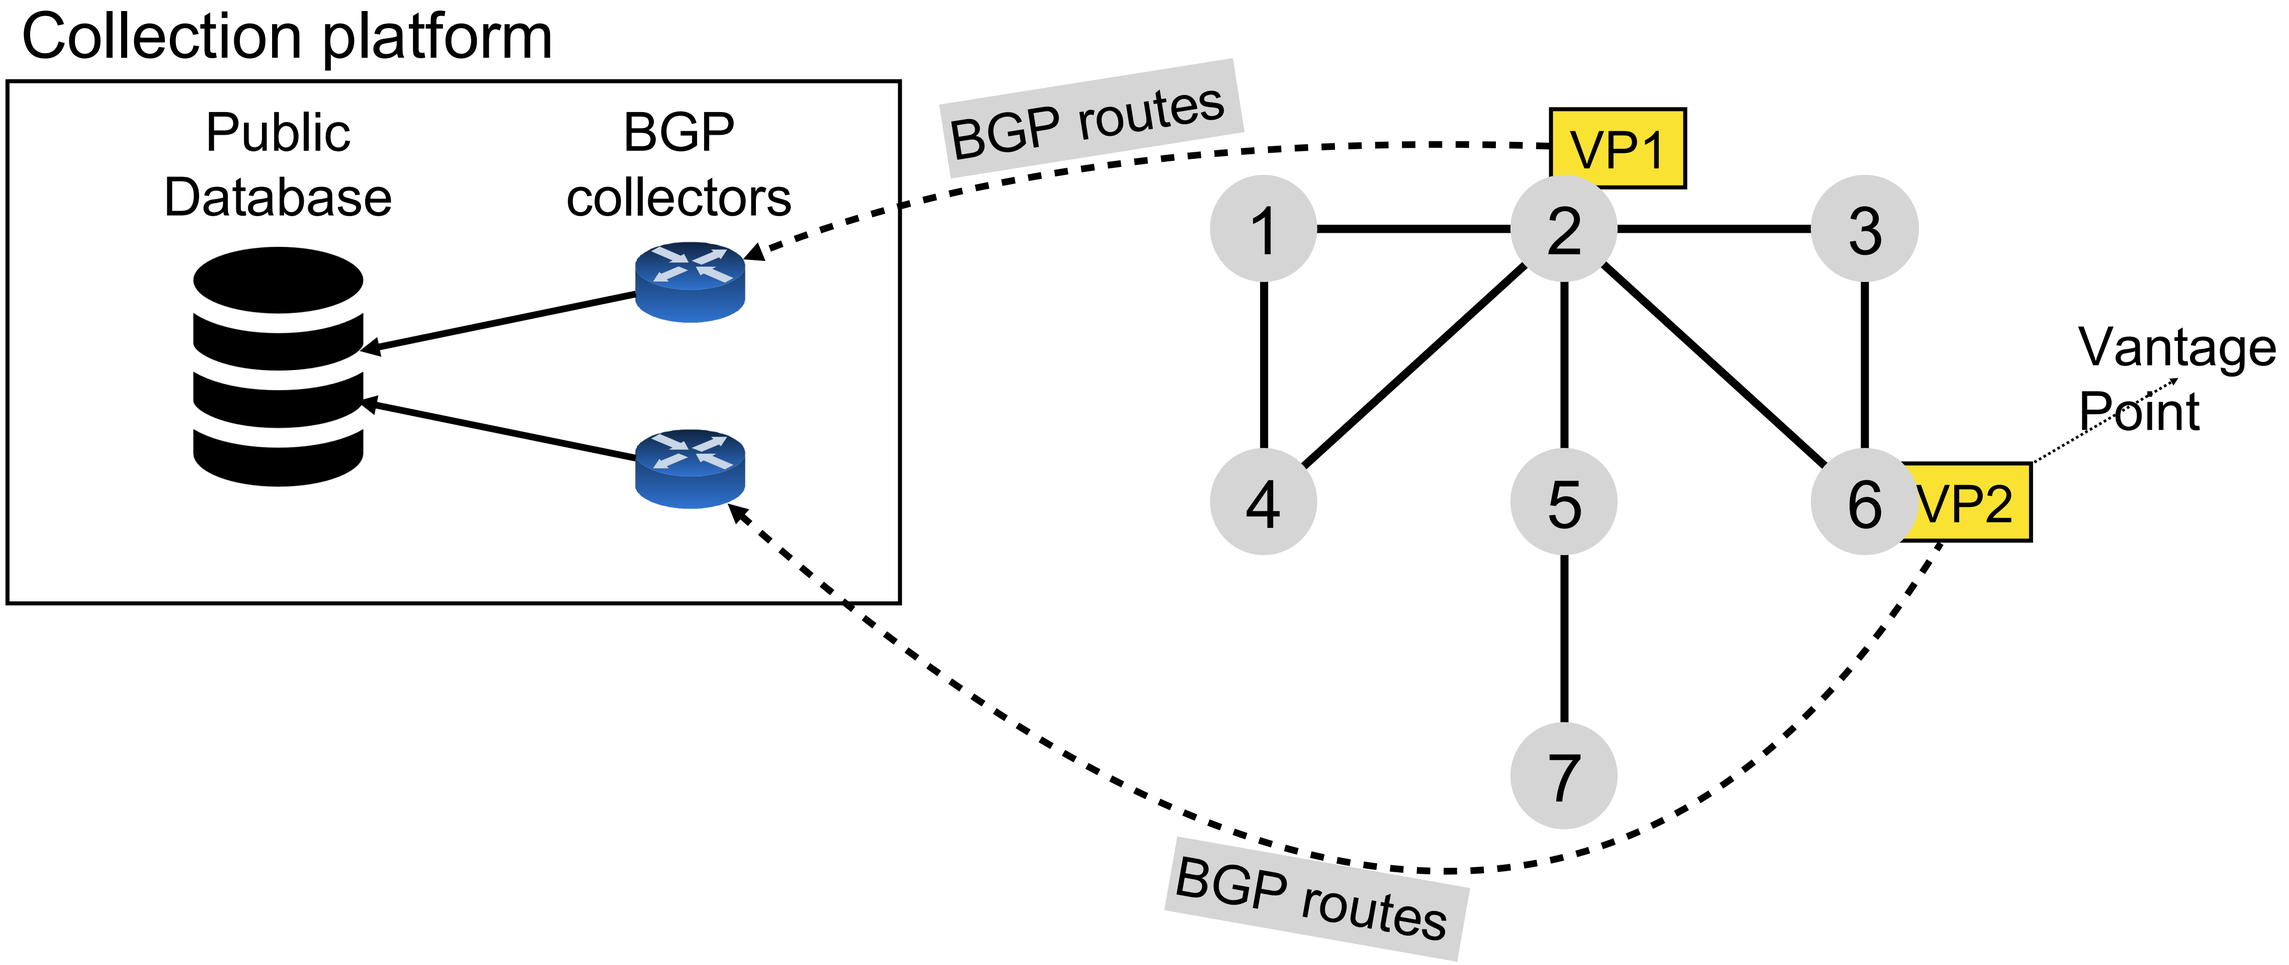

Current collection platforms are facing operational limitations

Median # of collected public BGP routes

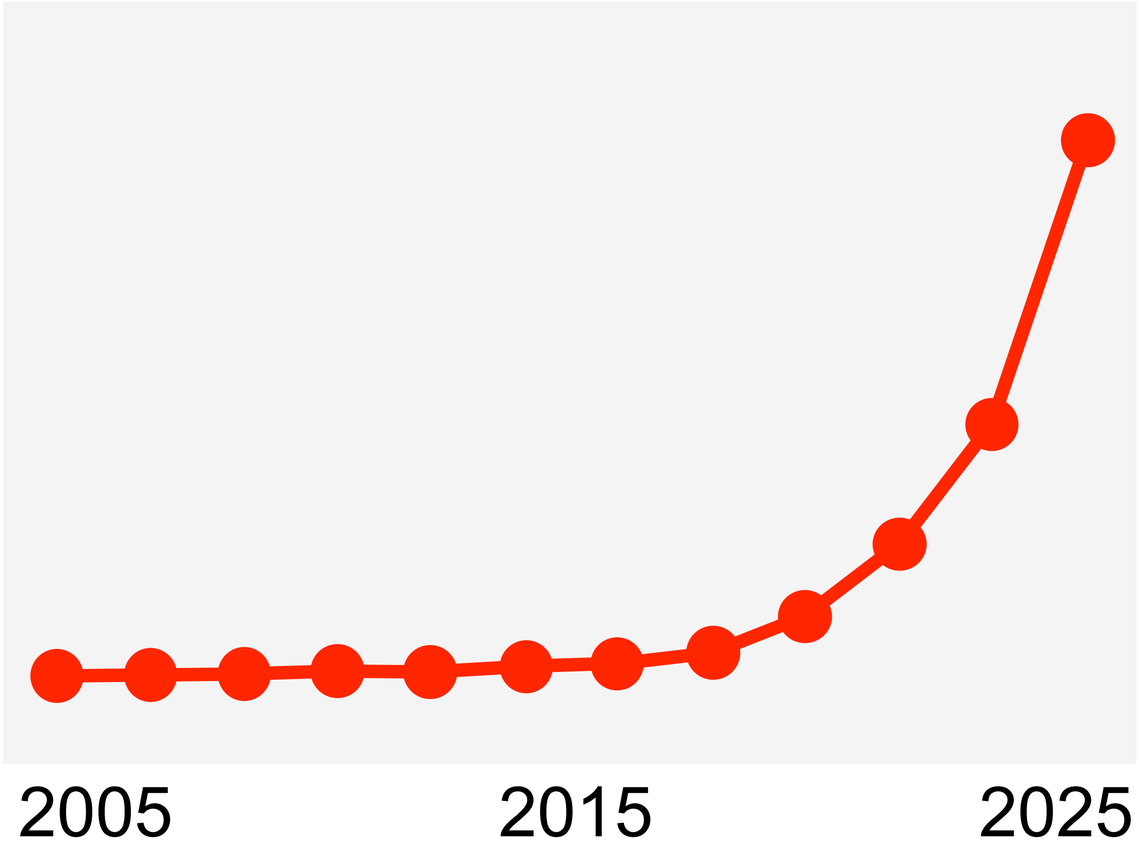

Current collection platforms are facing operational limitations

Median # of collected public BGP routes

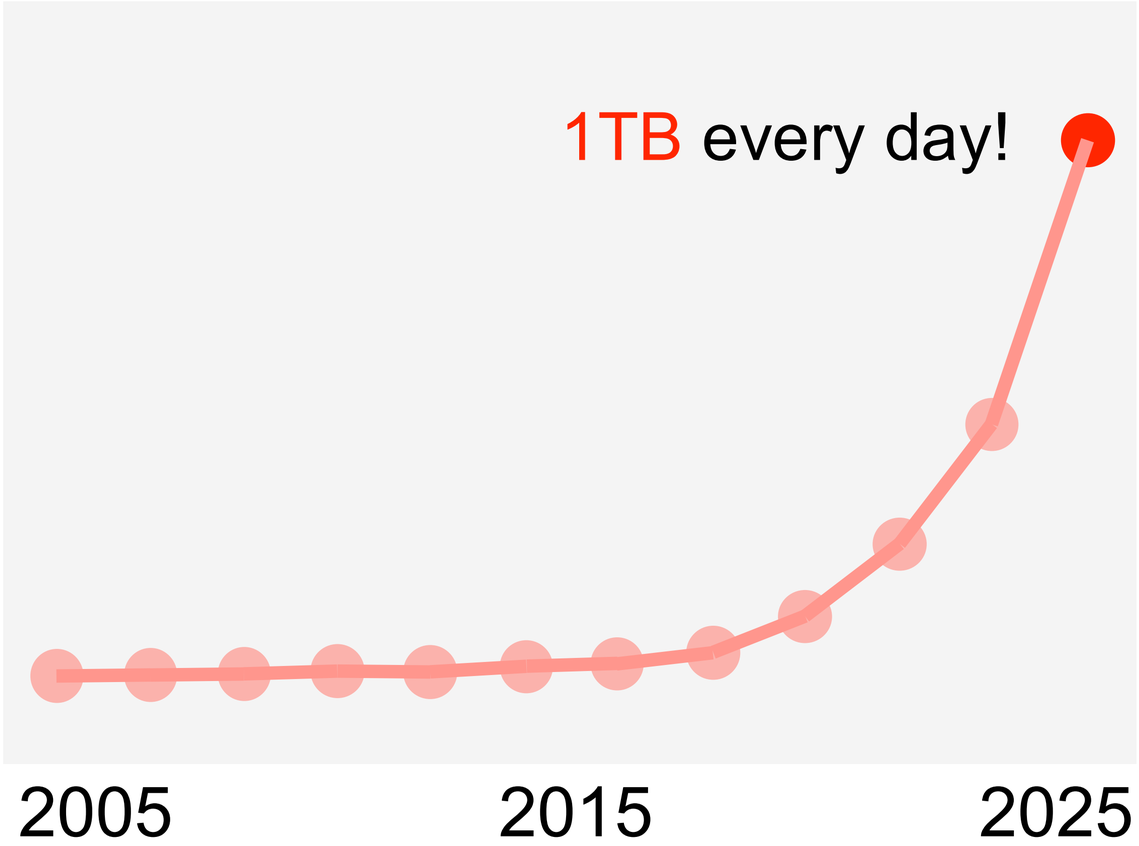

Current collection platforms are facing operational limitations

Median # of collected public BGP routes

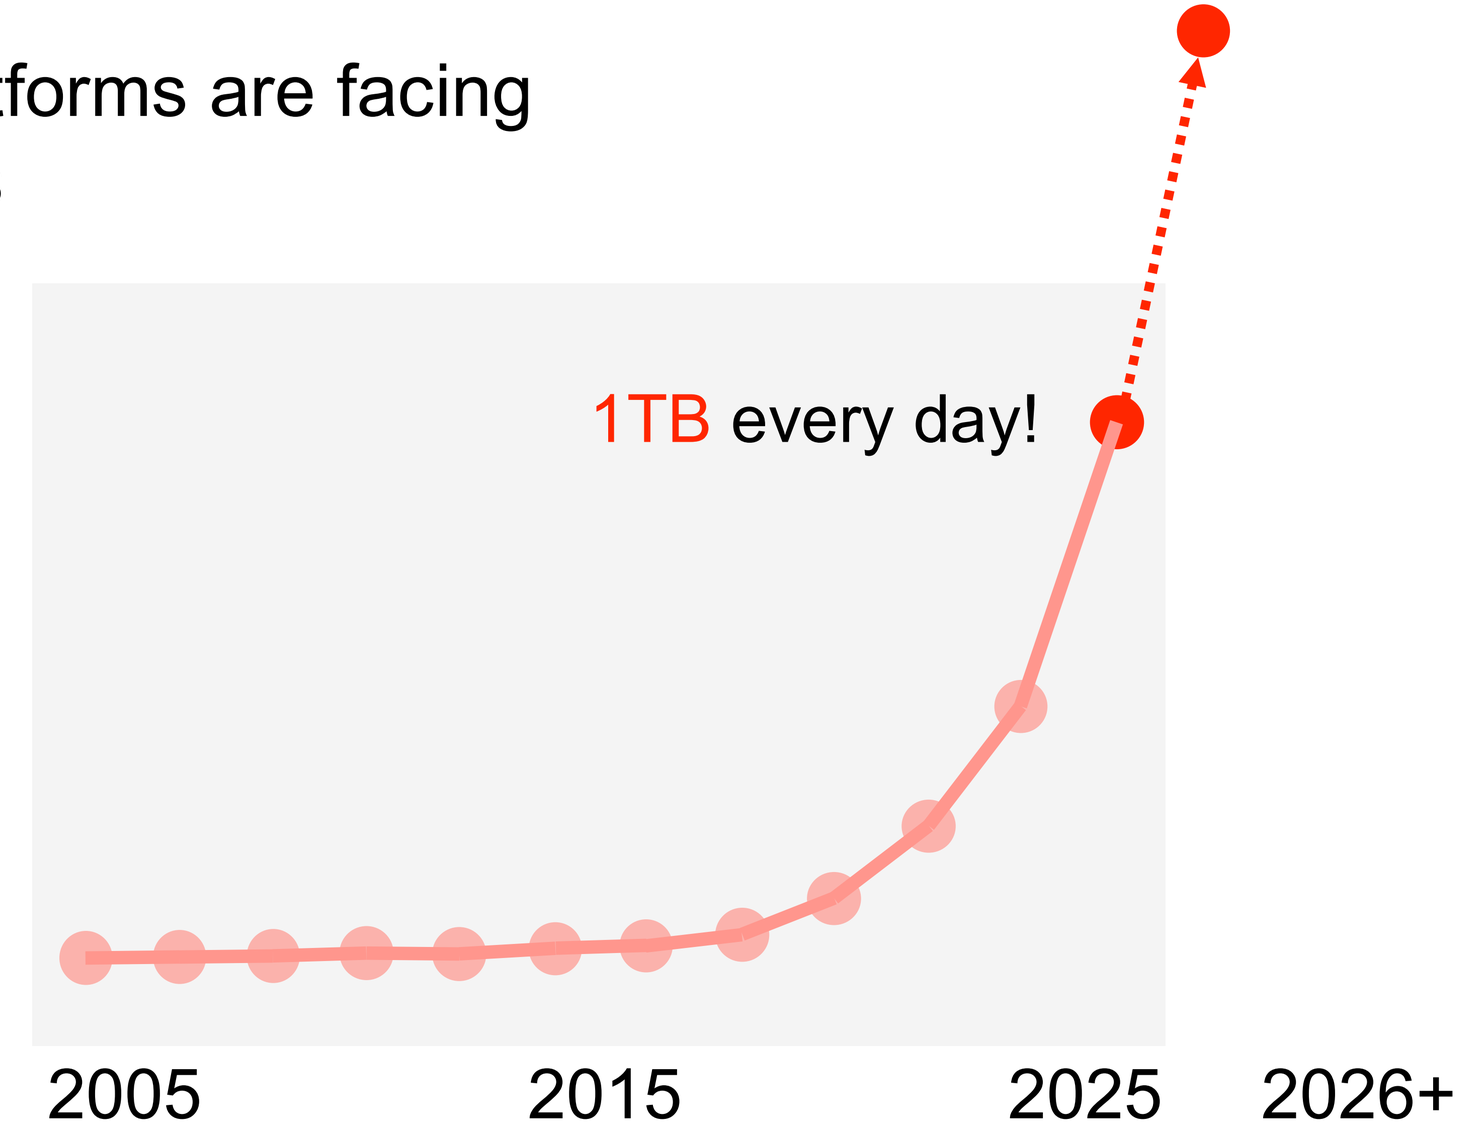

Vantage Point (VP) coverage (RIS+RouteViews) — 1.2%

% of ASes sharing their BGP data

Vantage Point (VP) coverage (RIS+RouteViews) — 1.2%

% of ASes sharing their BGP data

Vantage Point (VP) coverage (RIS+RouteViews)

1.2%

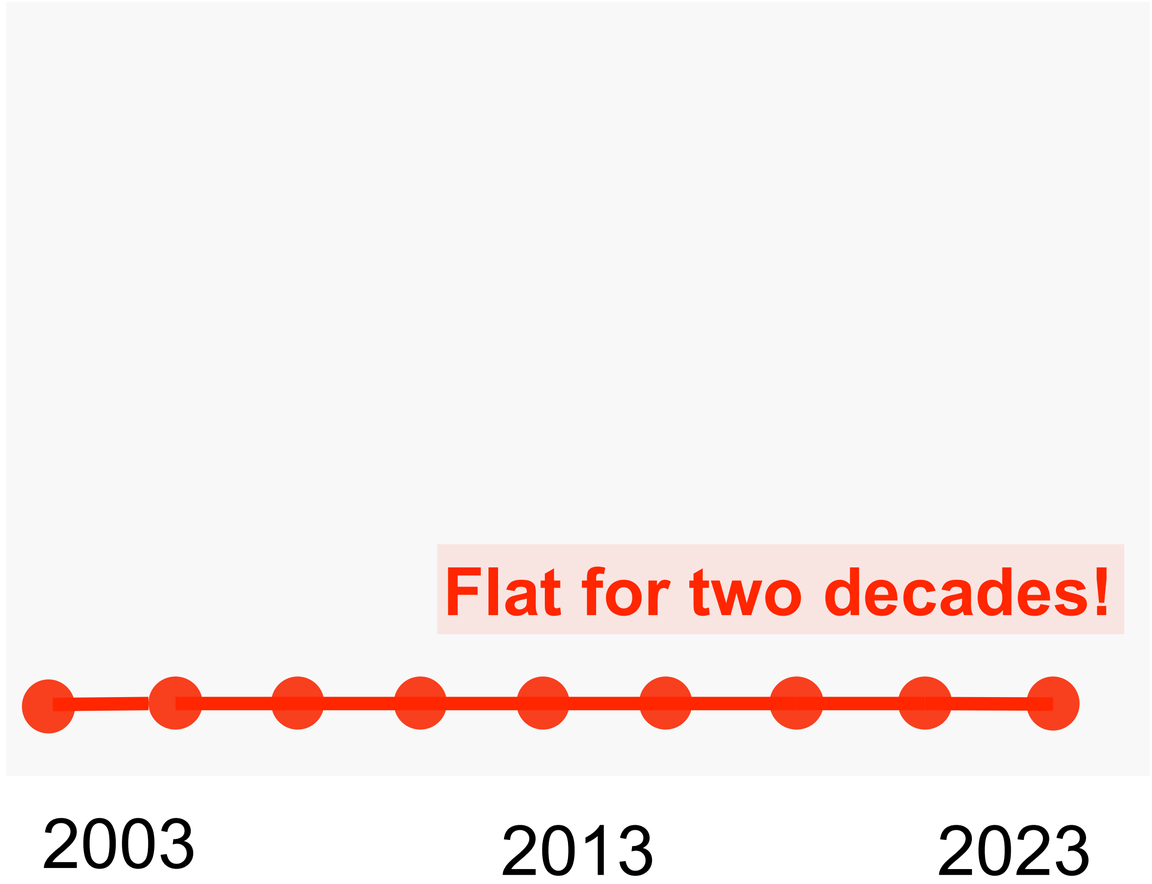

Because of the inherent path hiding nature of BGP, each VP only has a partial view over Internet routing

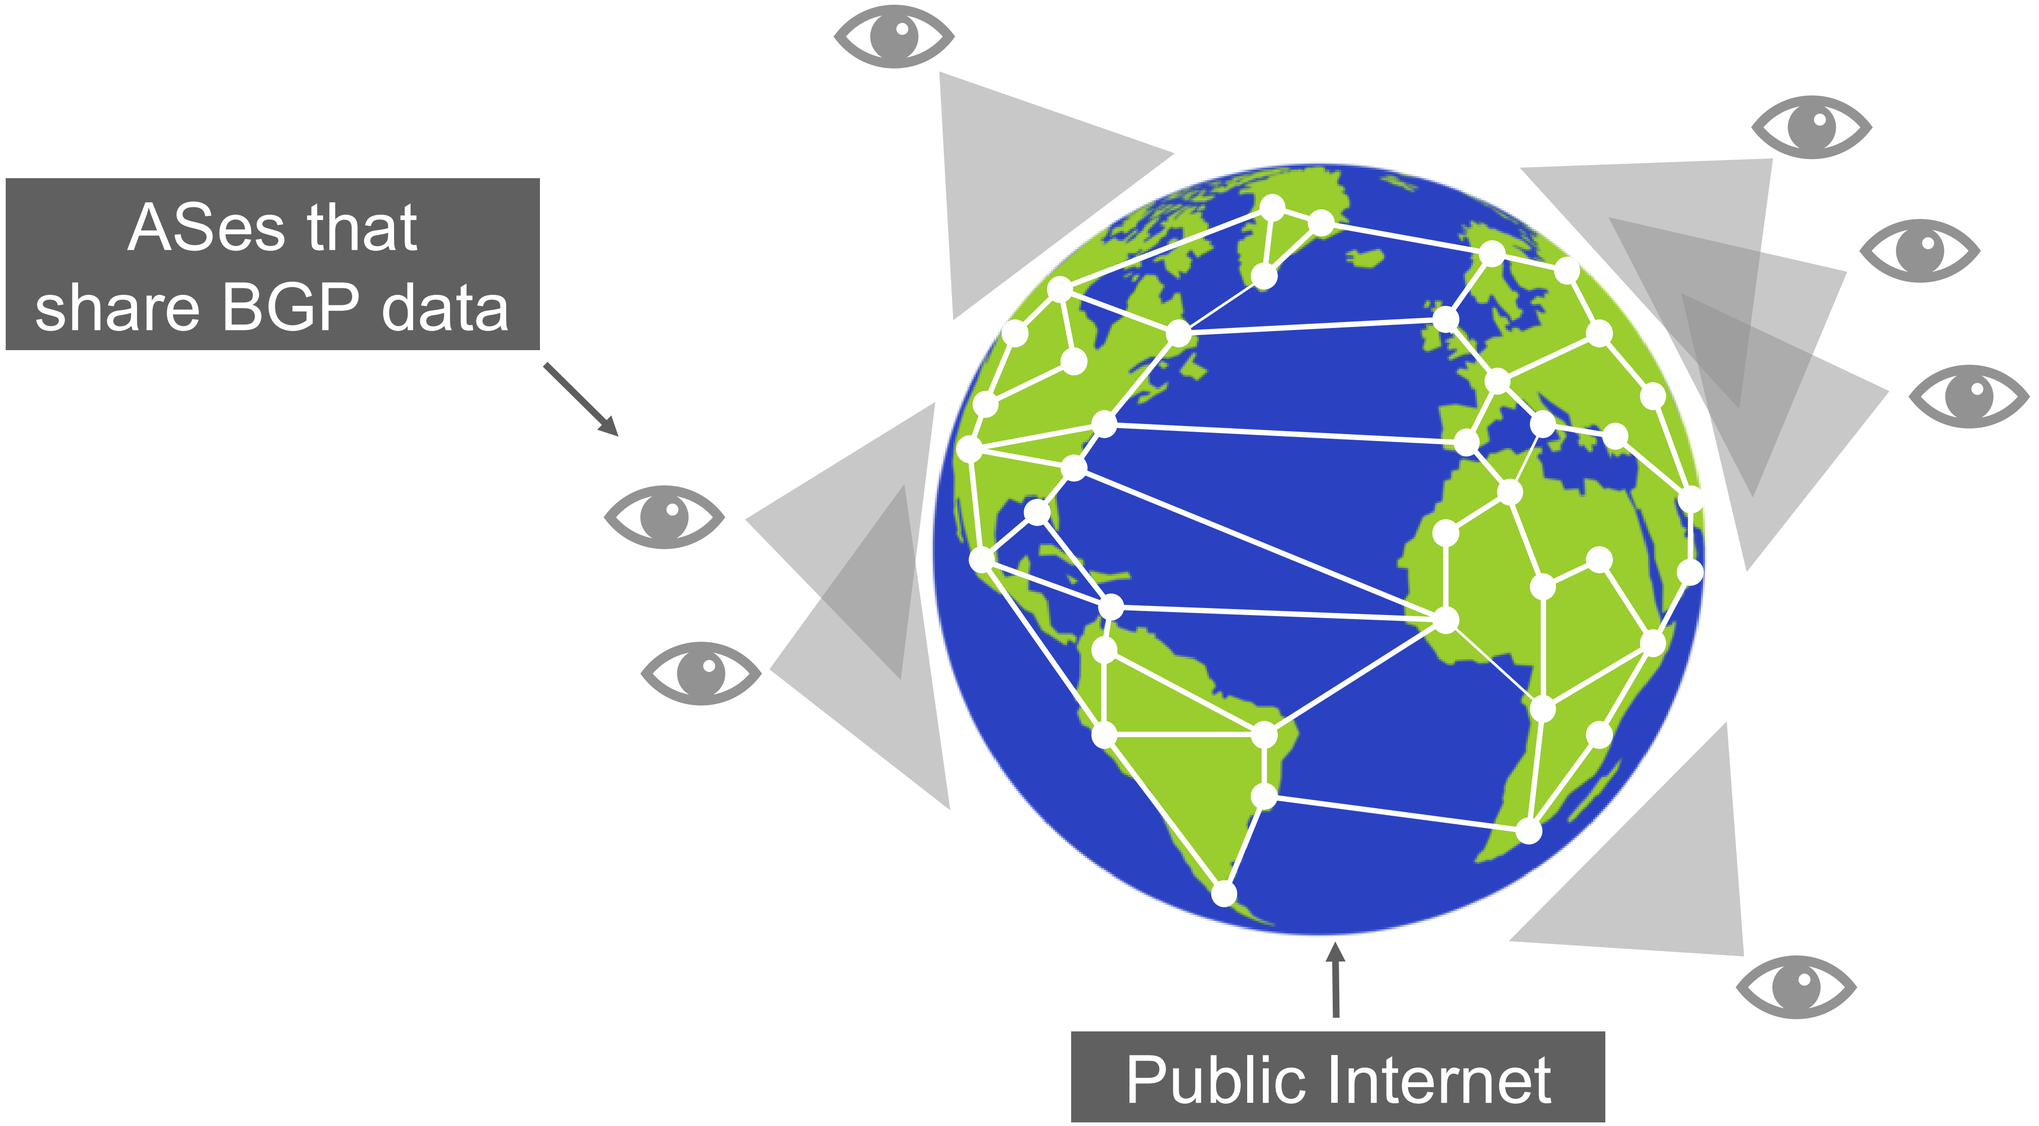

Because of the inherent path hiding nature of BGP, each VP only has a partial view over Internet routing

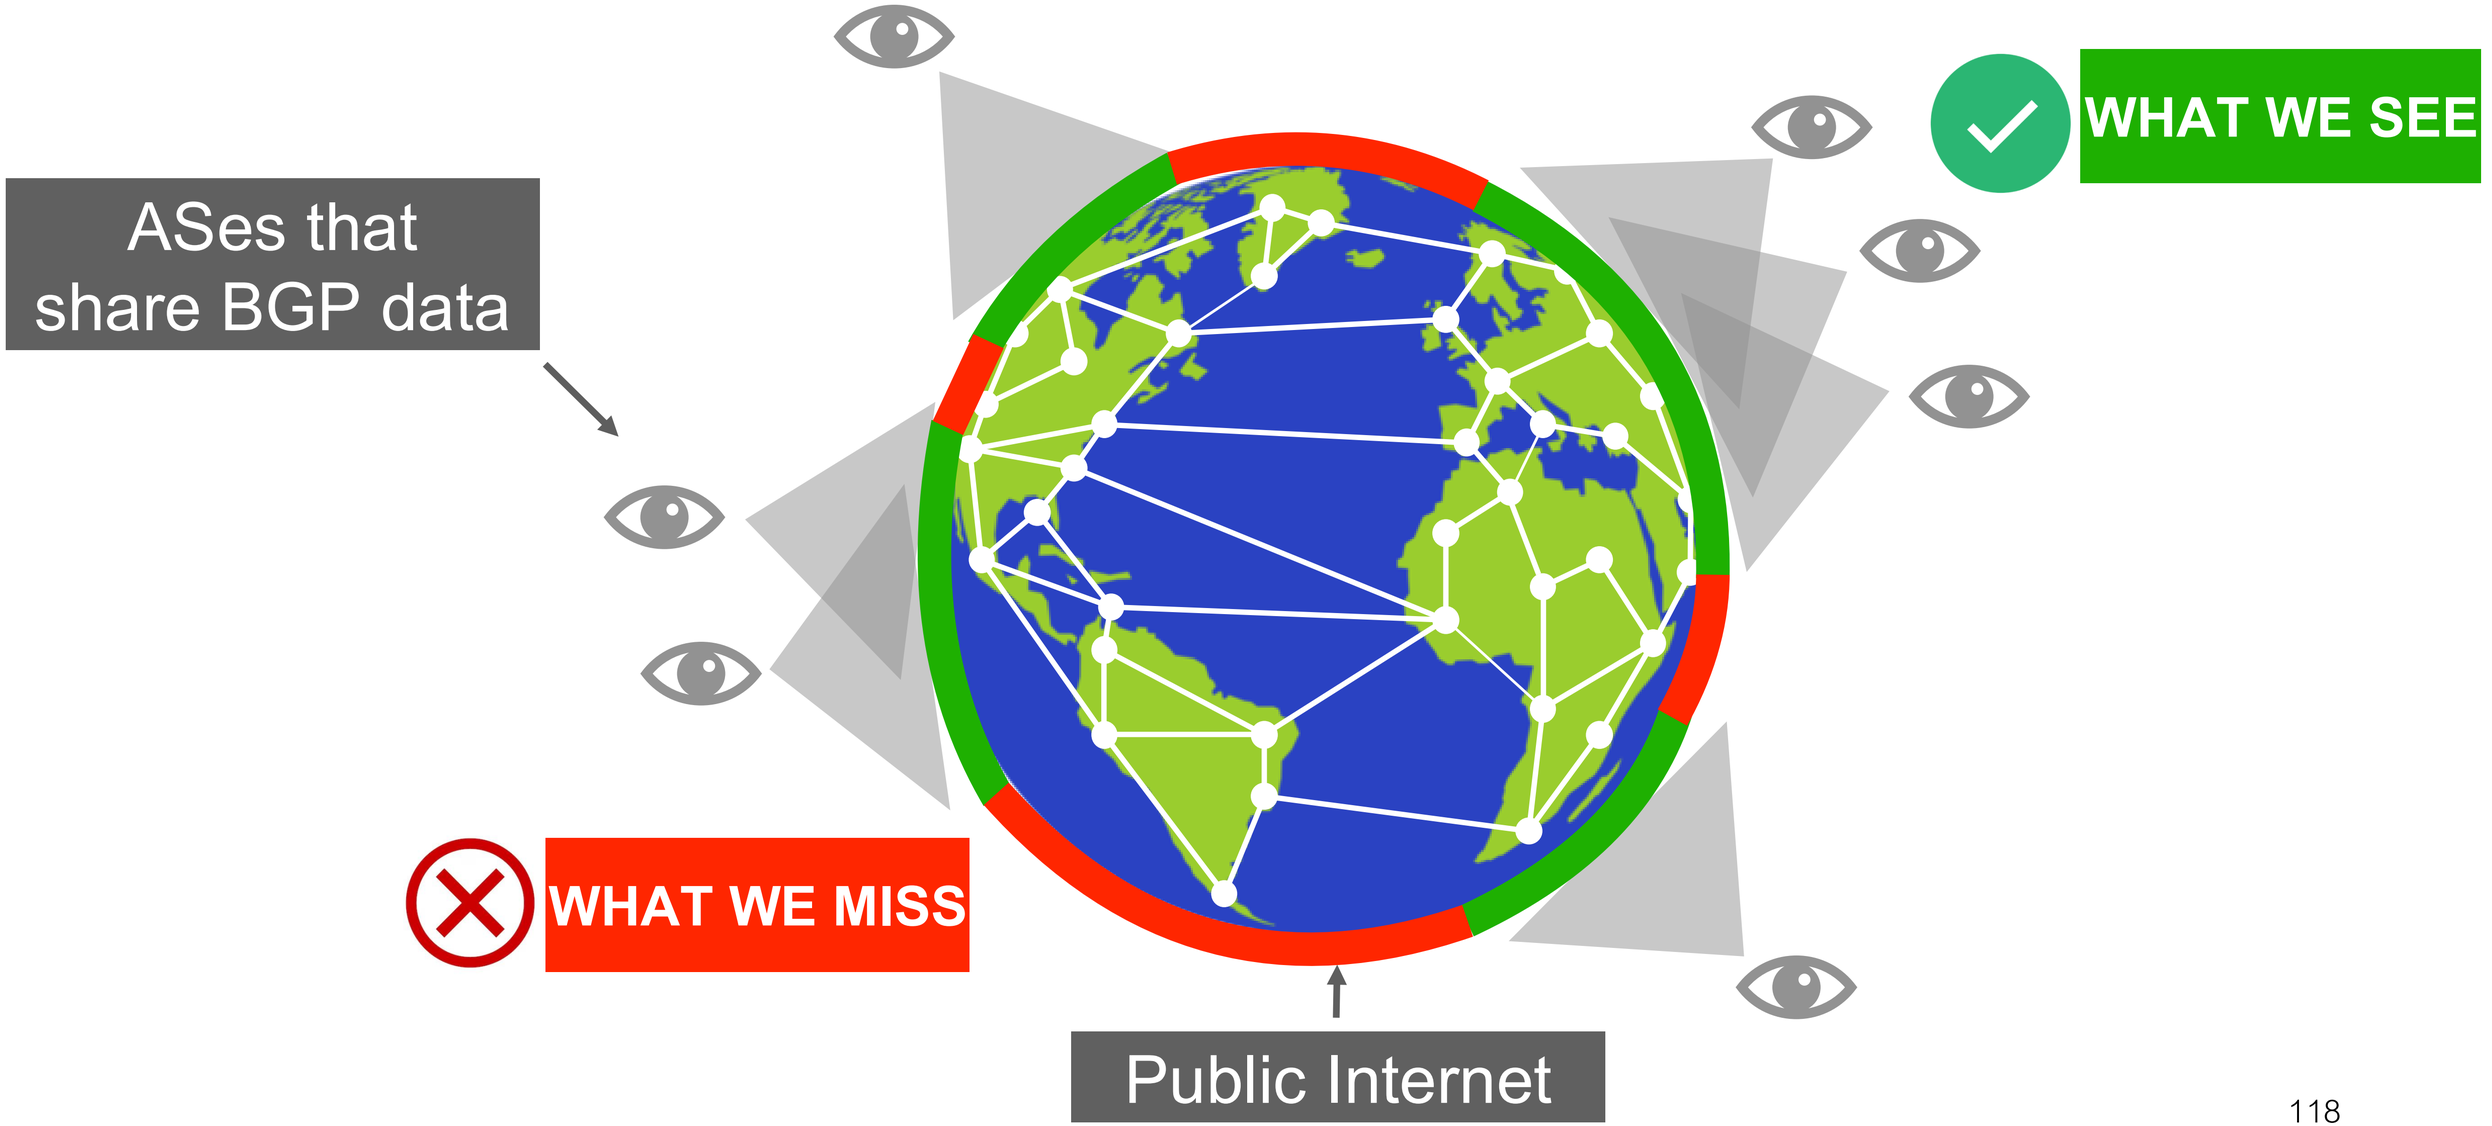

In the meantime, users often regretfully resort to sampling

Would you have used
more BGP data if you could?*

*Survey conducted among authors of eight
top research papers that sampled BGP data

In the meantime, users often regretfully resort to sampling

Would you have used
more BGP data if you could?*

Seven researchers

— Yes!

One researcher

— No

*Survey conducted among authors of eight top research papers that sampled BGP data

However: research and commercial BGP monitoring tools cannot afford visibility gaps in BGP data

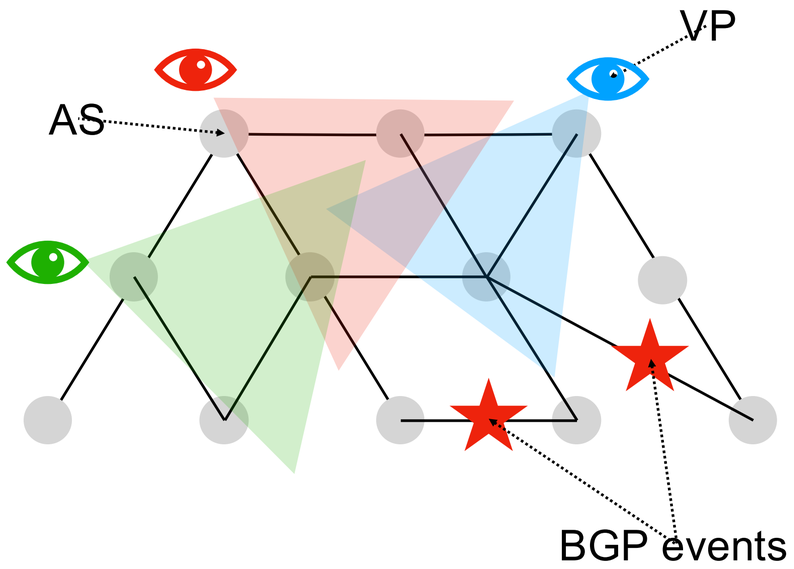

On the Effectiveness of BGP Hijackers That Evade Public Route Collectors

ALEXANDROS MILOLIDAKIS[1], TOBIAS BÜHLER[2], KUNYU WANG[1], MARCO CHIESA[1], LAURENT VANBEVER[2], AND STEFANO VISSICCHIO[3]
[1]KTH Royal Institute of Technology, 114 28 Stockholm, Sweden
[2]ETH Zürich, 8092 Zürich, Switzerland
[3]Department of Computer Science, University College London (UCL), WC1E 6BT London, U.K.
Corresponding author: Alexandros Milolidakis (miloli@kth.se)

This work was supported in part by the Swedish Foundation for Strategic Research under Grant 64455, and in part by the KTH Digital Futures.

ABSTRACT Routing hijack attacks have plagued the Internet for decades. After many failed mitigation attempts, recent Internet-wide BGP monitoring infrastructures relying on distributed route collection systems, called route collectors, give us hope that future monitor systems can quickly detect and ultimately mitigate hijacks. In this paper, we investigate the effectiveness of public route collectors with respect to future attackers deliberately engineering longer hijacks to avoid being recorded by route collectors. Our extensive simulations (and attacks we device) show that monitor-based systems may be unable to observe many carefully crafted hijacks diverting traffic from thousands of ASes. Hijackers could predict whether their attacks would propagate to some BGP feeders (i.e., monitors) of public route collectors. Then, manipulate BGP route propagation so that the attack never reaches those monitors. This observation remains true when considering plausible future Internet topologies, with more IXP links and up to 4 times more monitors peering with route collectors. We then evaluate the feasibility of performing hijacks not observed by route collectors in the real-world. We experiment with two classifiers to predict the monitors that are dangerous to report the attack to route collectors, one based on monitor proximities (i.e., shortest path lengths) and another based on Gao-Rexford routing policies. We show that a proximity-based classifier could be sufficient for the hijacker to identify all dangerous monitors for hijacks announced to peer-to-peer neighbors. For hijacks announced to transit networks, a Gao-Rexford classifier reduces wrong inferences by ≥91% without introducing new misclassifications for existing dangerous monitors.

INDEX TERMS BGP, BGP hijacking, stealthy IP prefix hijacking, inter-domain routing, routing policies, route collectors, forged AS path, BGP monitoring, BGPStream.

I. INTRODUCTION

Routing hijacks keep affecting industry (including Google, Amazon, Apple, Microsoft) [1], [2], financial platforms (e.g., the Bitcoin network) [3], security authorities [4], Internet services [5], governments [6], and citizens [7]. 775 (830) suspicious BGP hijack (route leak) incidents have been documented in 2021 [8], and more than 2200 (1200) suspicious BGP hijack (route leak) incidents have been documented in 2020, a 30% hijack increase from 2019 [9].[1] This is mainly because BGP allows attackers to inject arbitrary information in the Internet routing system. By falsely claiming ownership of IP prefixes, malicious networks can divert, eavesdrop, store, and possibly modify traffic in so-called *interception attacks*.[2]

The associate editor coordinating the review of this manuscript and approving it for publication was Salekul Islam.

[1]While the references distinguish between route leaks and BGP hijacks, we consider them analogous terms, both classified under the category of Type-N attacks, which we will introduce formally in Section III-A.
[2]BGP hijacks can be used to drop traffic as well. We focus on interception attacks as they are harder to detect for the victim.

Because of the inherent path hiding nature of BGP, each VP only has a partial view over Internet routing

We measure the impact of the low RIS and RouteViews VP coverage on three use cases using simulations*

Use case #1: Peer-to-peer link observation

Use case #2: Inter-domain failure localization

Use case #3: Forged-origin hijack detection

*Our simulations use c-bgp on topologies with 6k ASes

We measure the impact of the low RIS and RouteViews VP coverage on three use cases using simulations*

Use case #1: Peer-to-peer link observation

Use case #2: Inter-domain failure localization

Use case #3: Forged-origin hijack detection

*Our simulations use c-bgp on topologies with 6k ASes

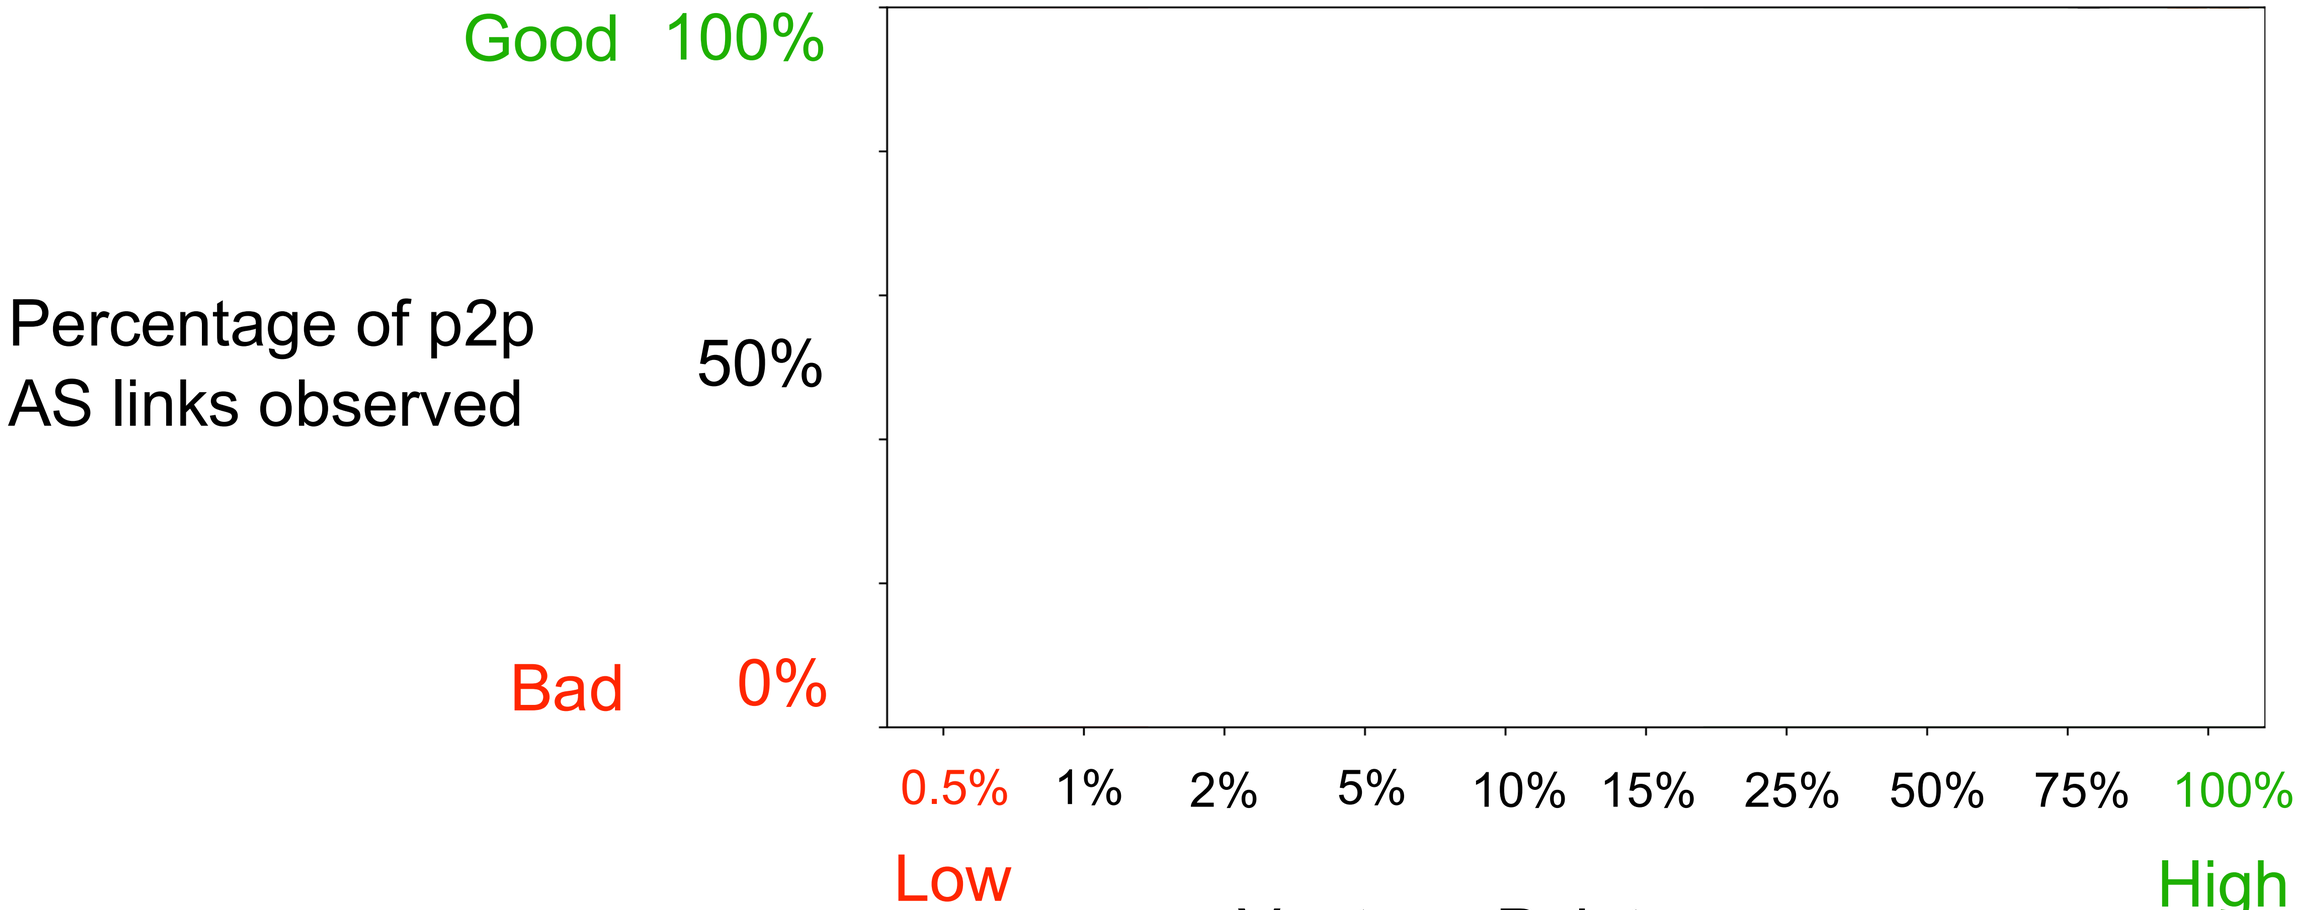
Good
100%
Percentage of p2p
AS links observed
50%
Bad
0%
0.5%
1%
2%
5%
10%
15%
25%
50%
75%
100%
Low
High
Vantage Point coverage
(custom scale)

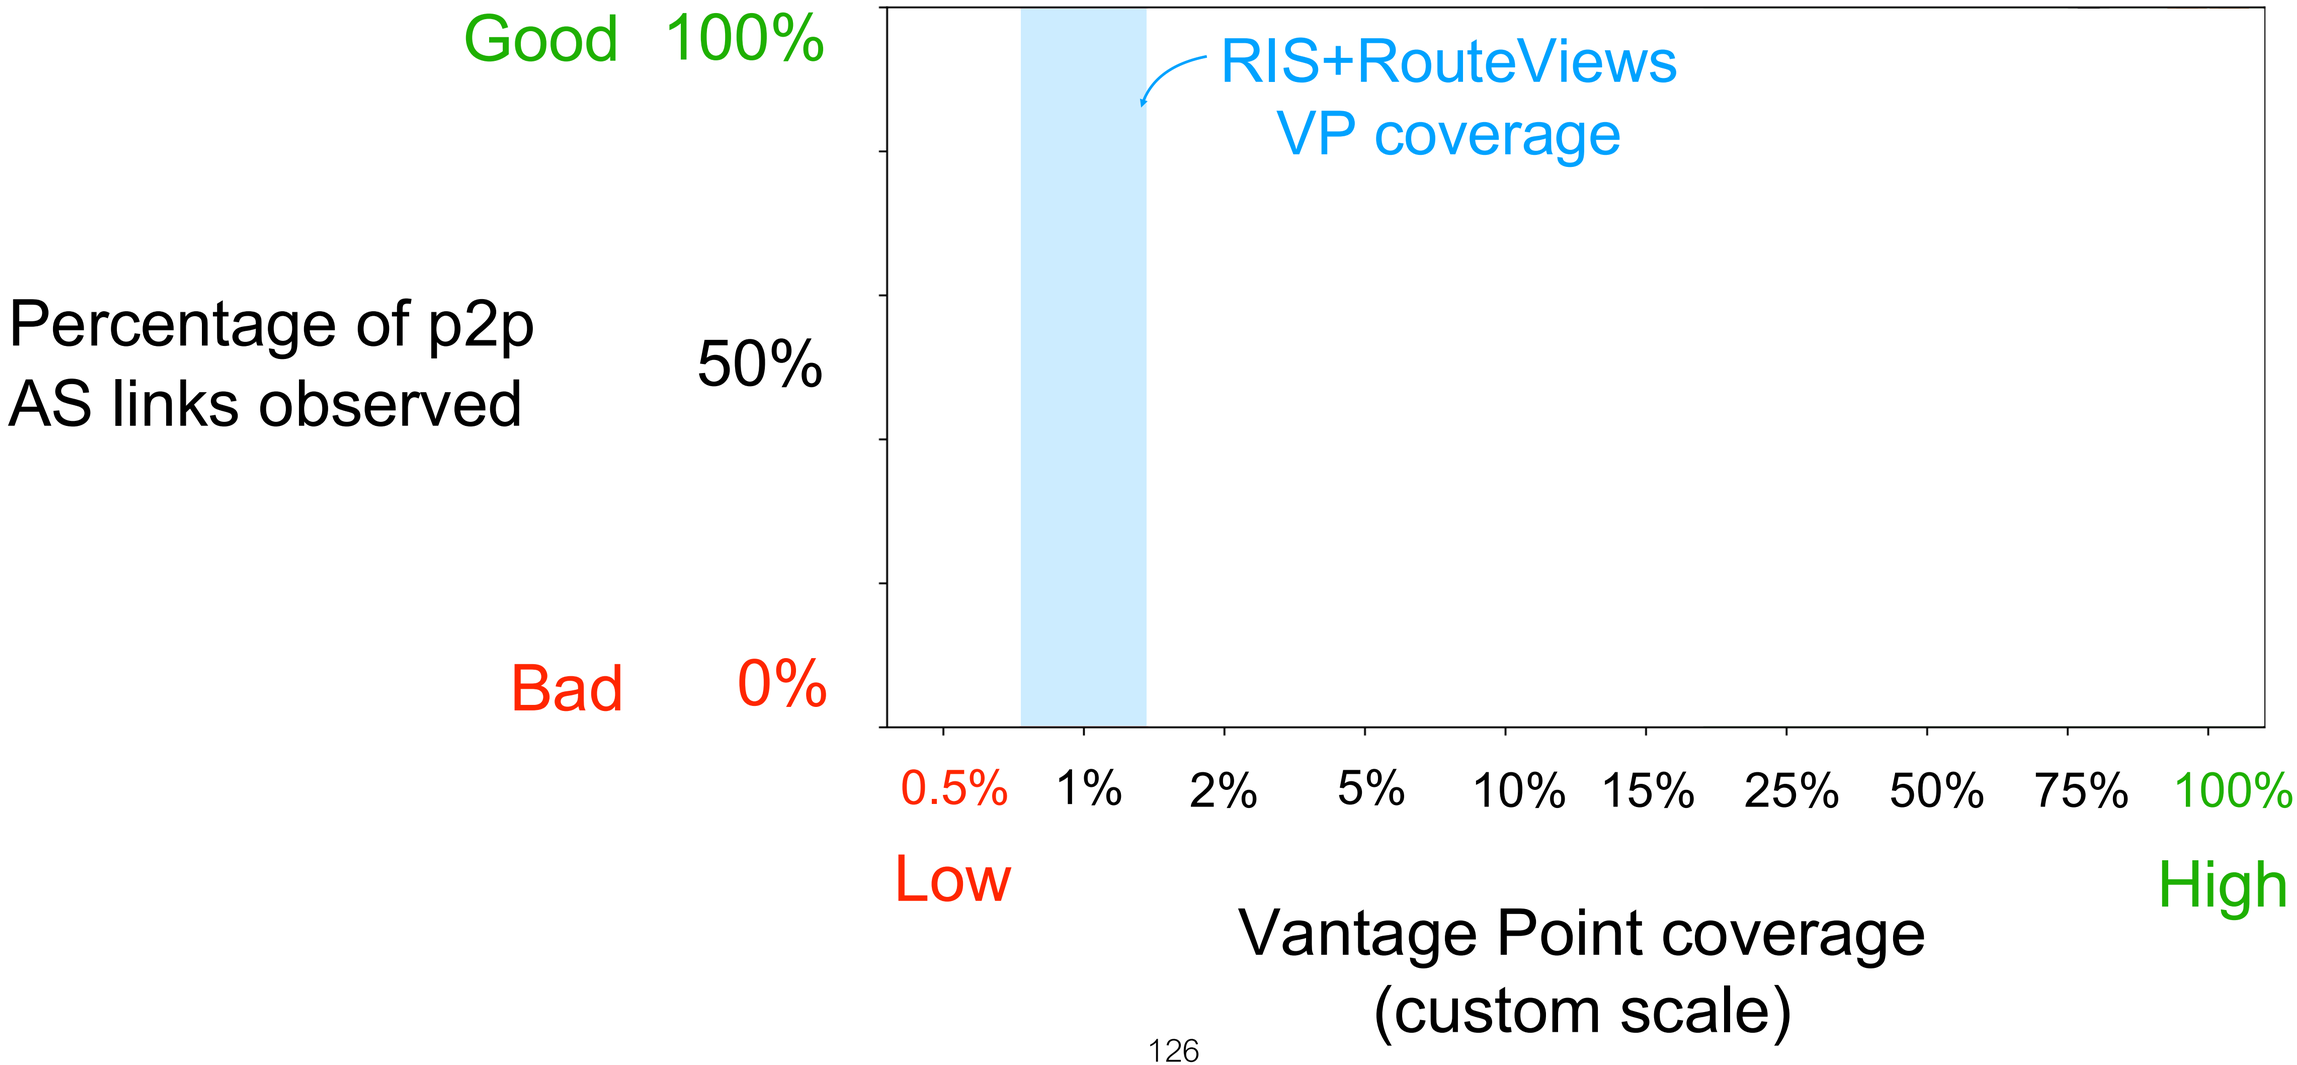
Good
100%
RIS+RouteViews
VP coverage
Percentage of p2p
AS links observed
50%
Bad
0%
0.5%
1%
2%
5%
10%
15%
25%
50%
75%
100%
Low
High
Vantage Point coverage
(custom scale)

Our simulations show that we are missing 84% of the p2p links

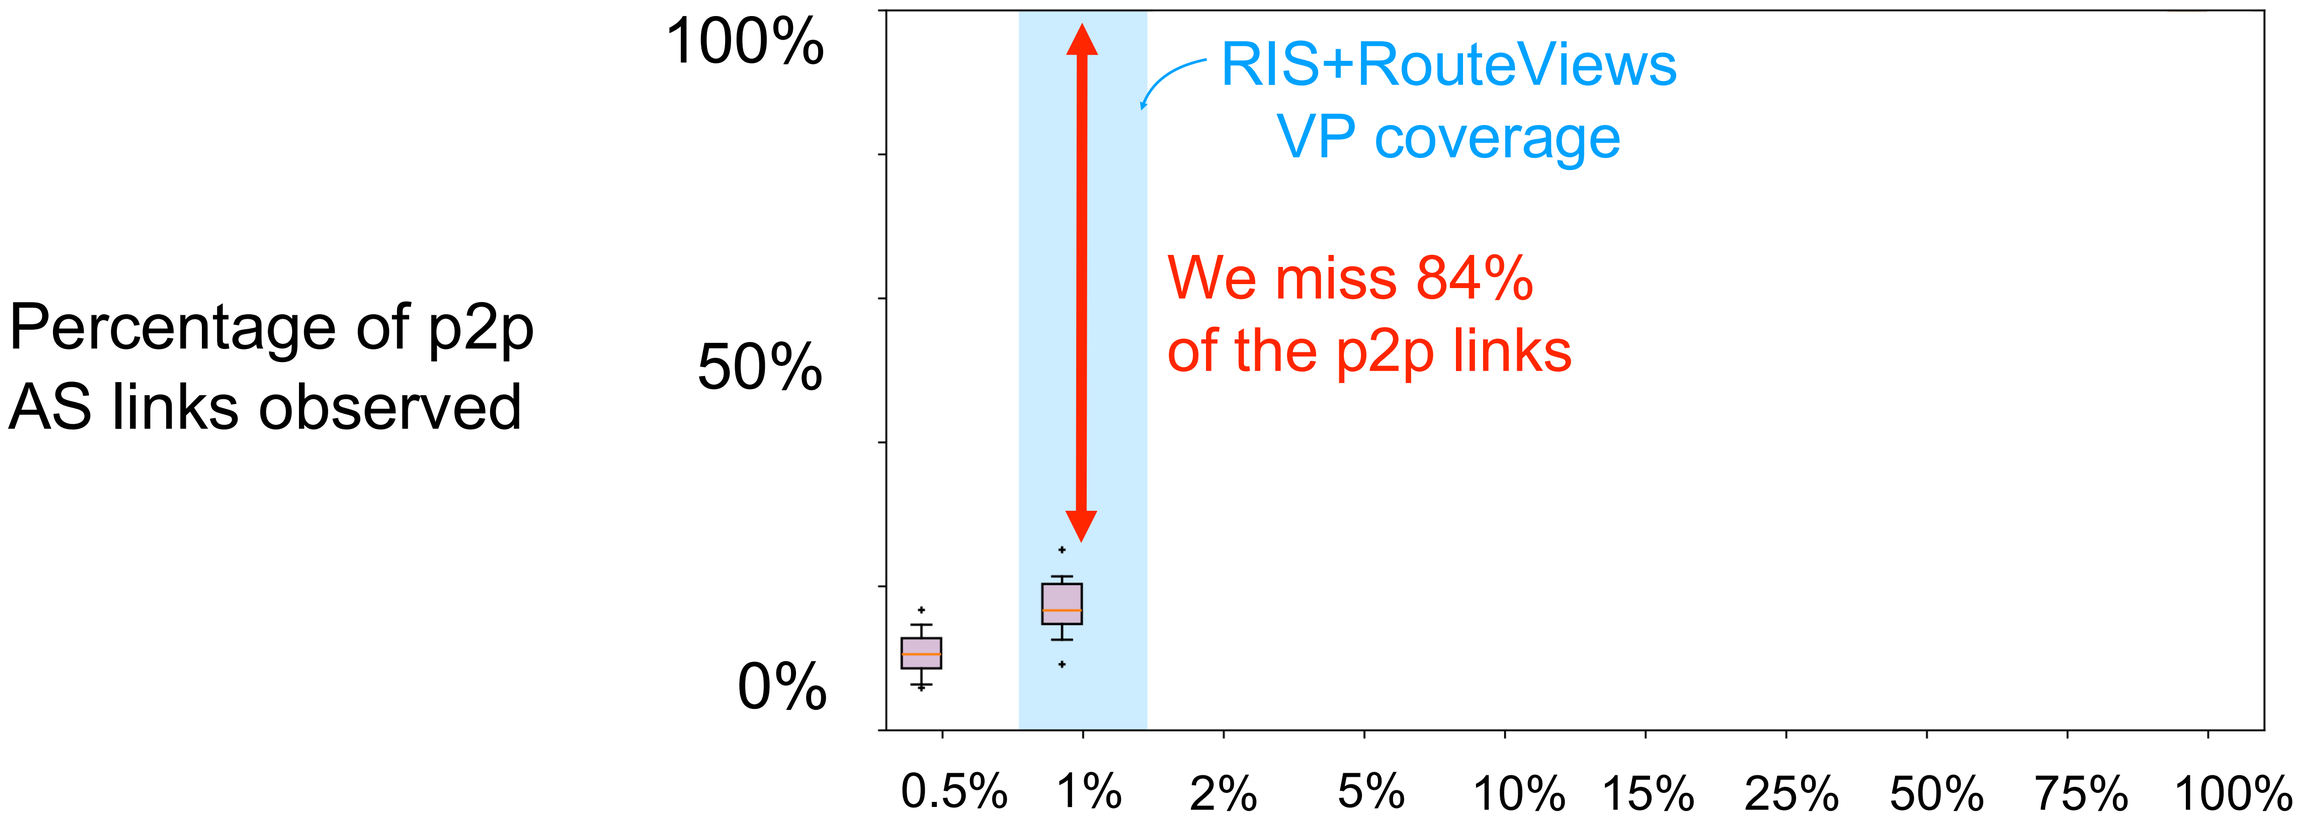

Results from 10 simulations
on “mini” Internets with 6k ASes

Our simulations show that we are missing 84% of the p2p links and suggest an order of magnitude more VPs

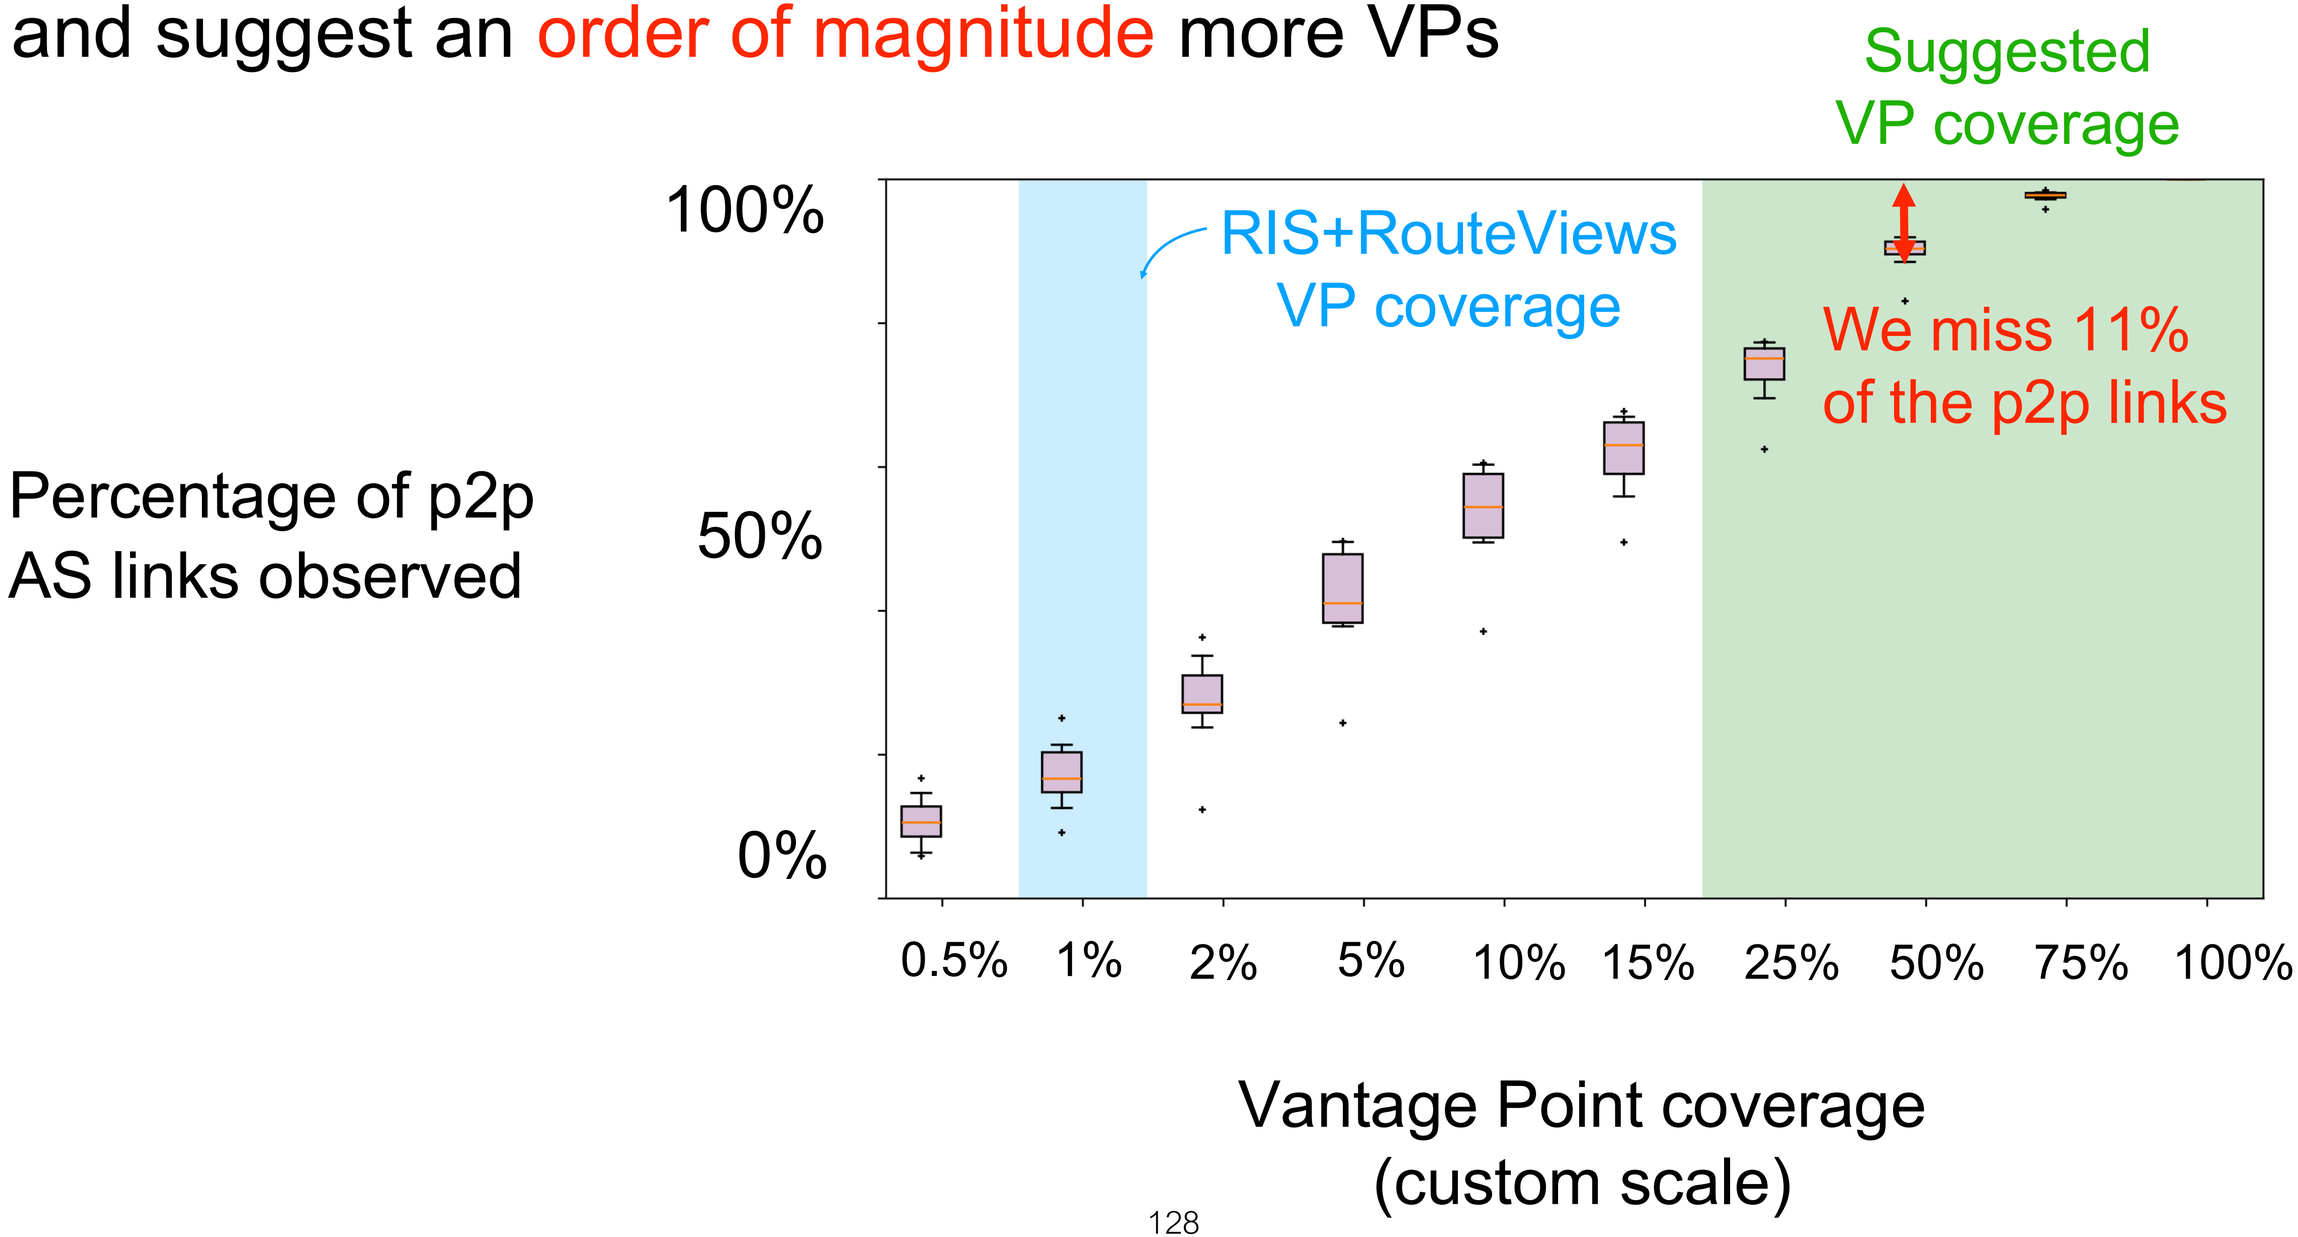

We measure the impact of the low RIS and RouteViews VP coverage on three use cases using simulations*

Use case #1: Peer-to-peer link observation
We miss 84% of the links

Use case #2: Inter-domain failure localization
We cannot localize 92% of the failures

Use case #3: Forged-origin hijack detection
We miss 24% of the Type-1 hijacks

*Our simulations use c-bgp on topologies with 6k ASes

A naive solution would be to significantly increase the VP coverage

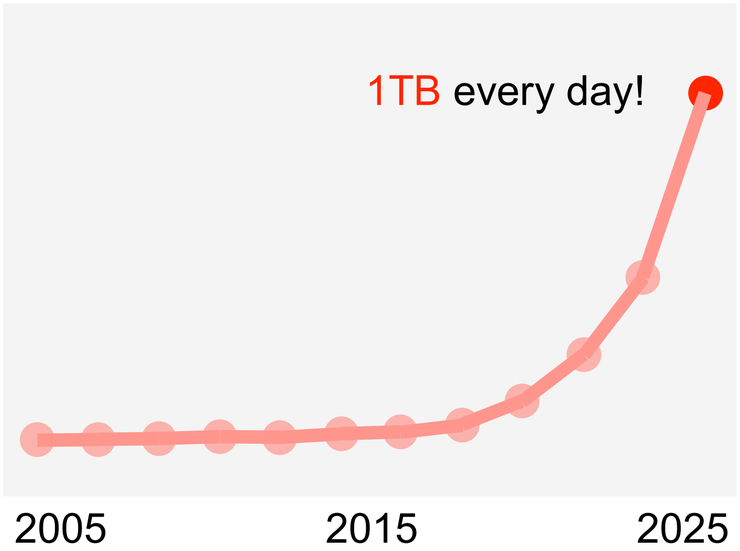

Not feasible because of the
data management constraints

Our *contribution:* A system that reduces the visibility gaps
of the public Internet routing ecosystem

GILL: The Next Generation of BGP Data Collection Platforms

Today: **all** routes
from a **few** ASes

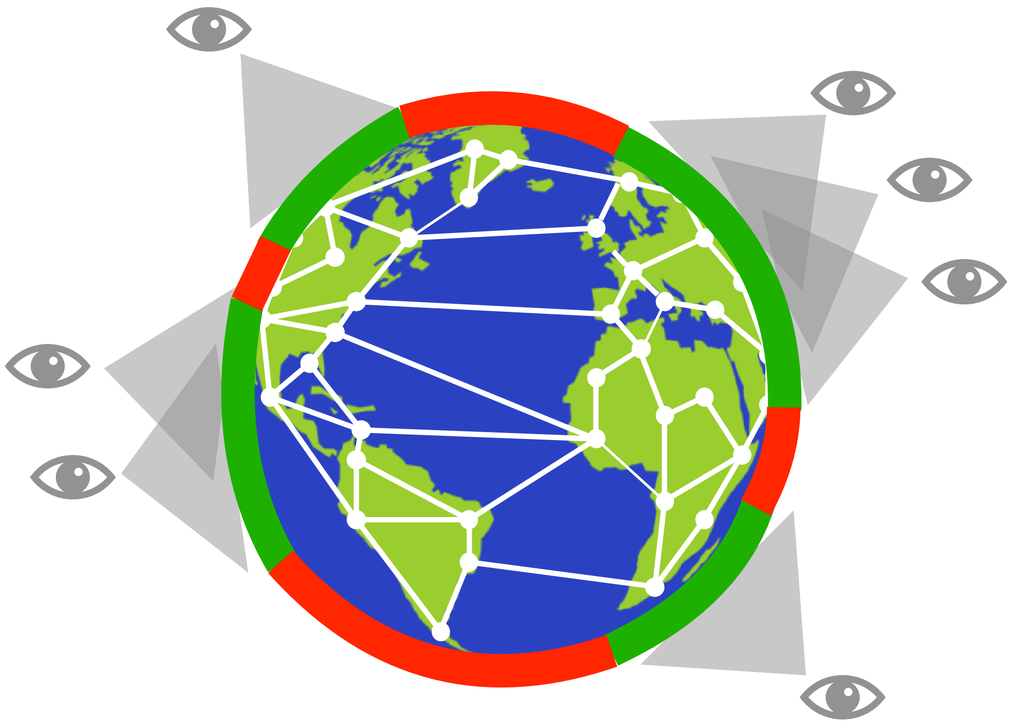

GILL: The Next Generation of BGP Data Collection Platforms

Today: **all** routes
from a **few** ASes

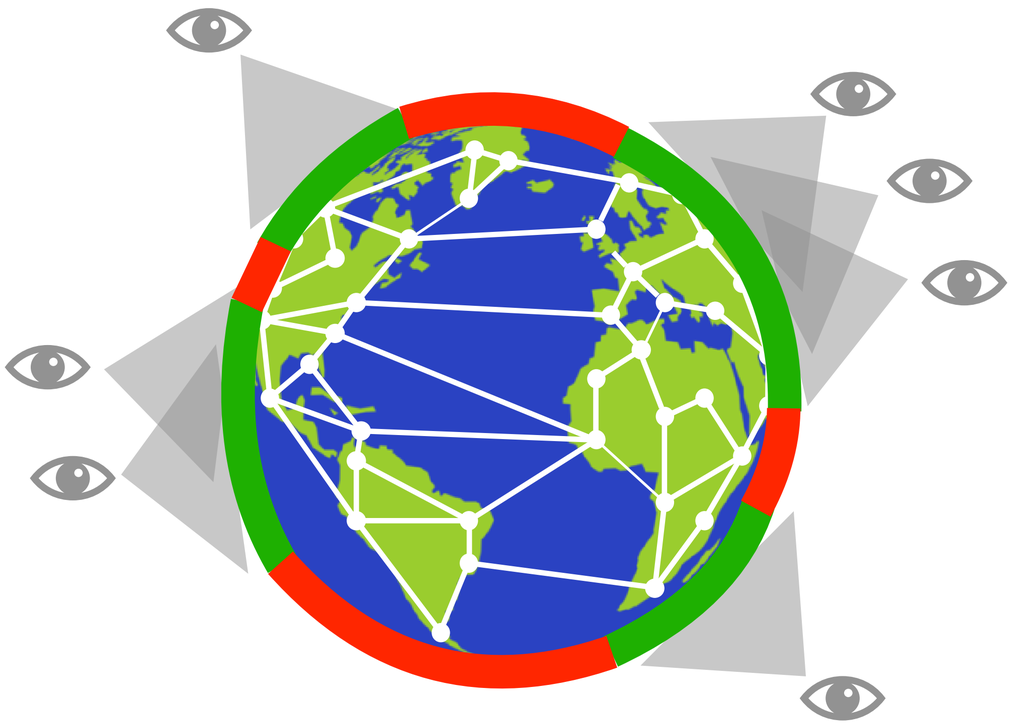

GILL: **few** routes
from **many** ASes

Outline

1. We observe that BGP routes are often redundant

BGP routes can be redundant

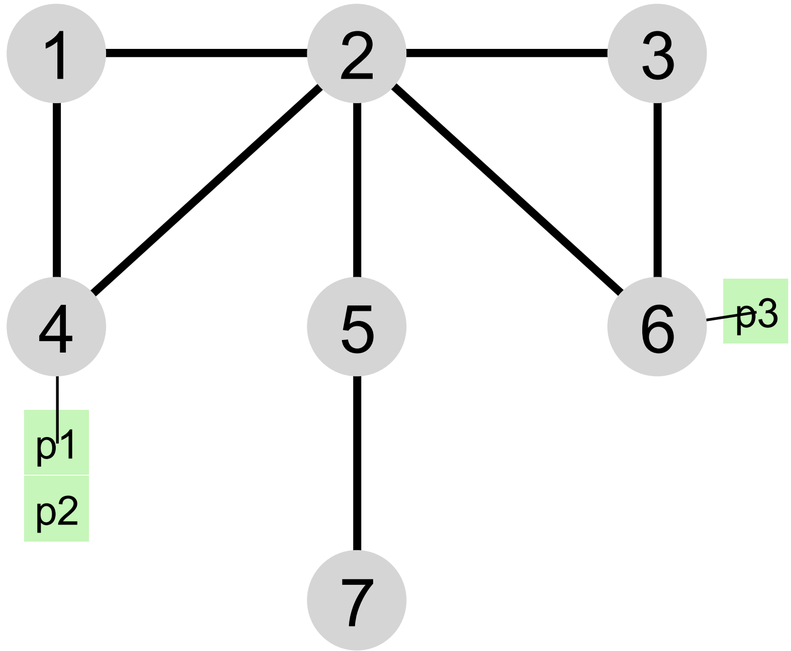

BGP routes can be redundant

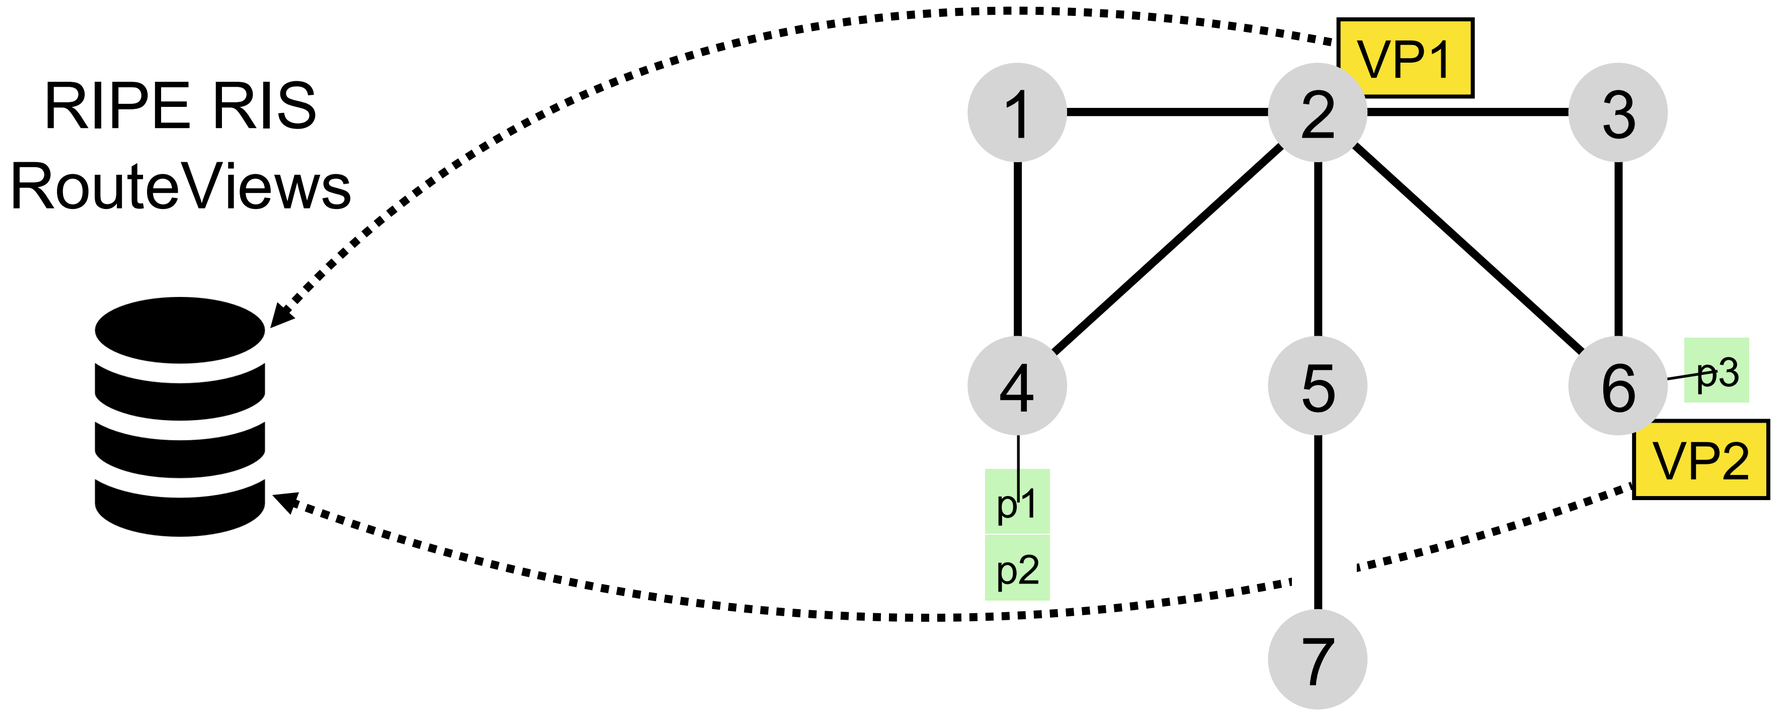

BGP routes can be redundant

Collected routes

VP	prefix	AS path

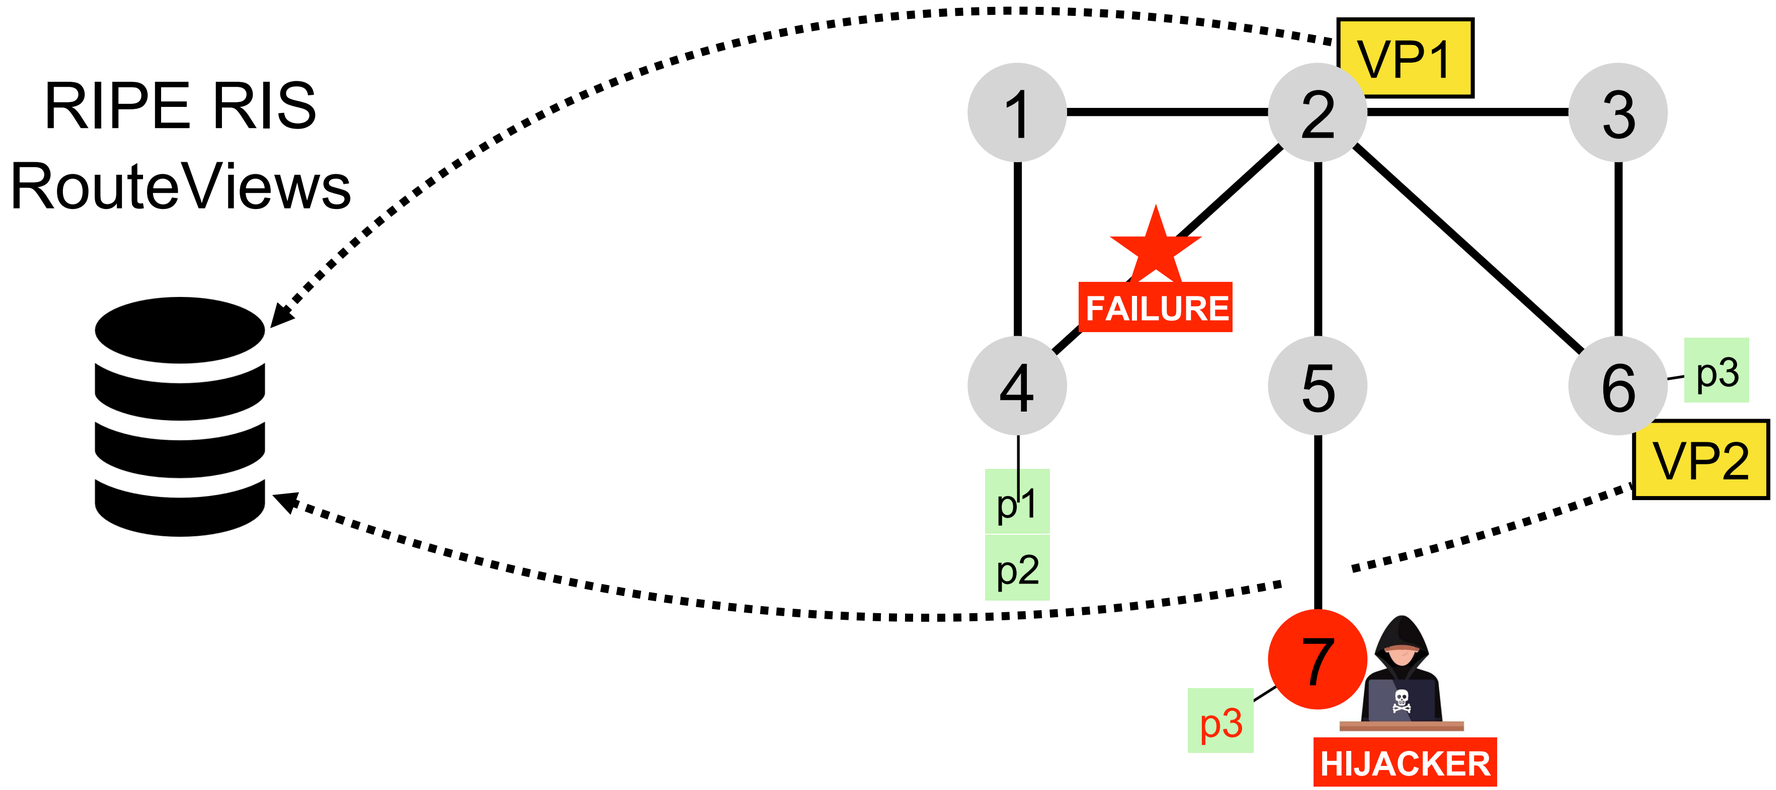

BGP routes can be redundant

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4

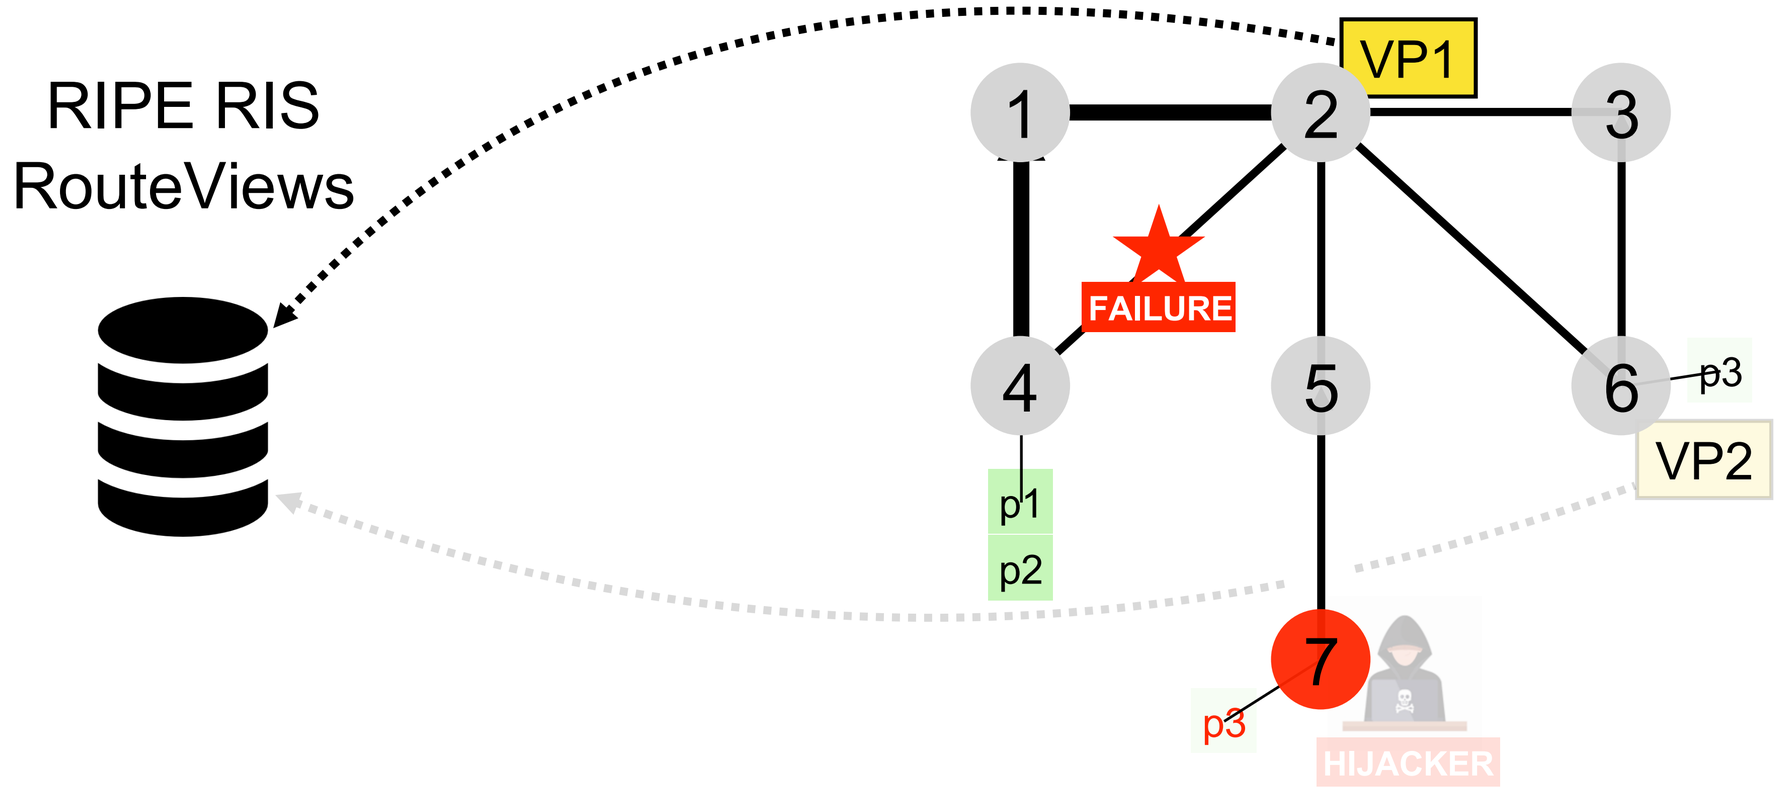

BGP routes can be redundant

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4

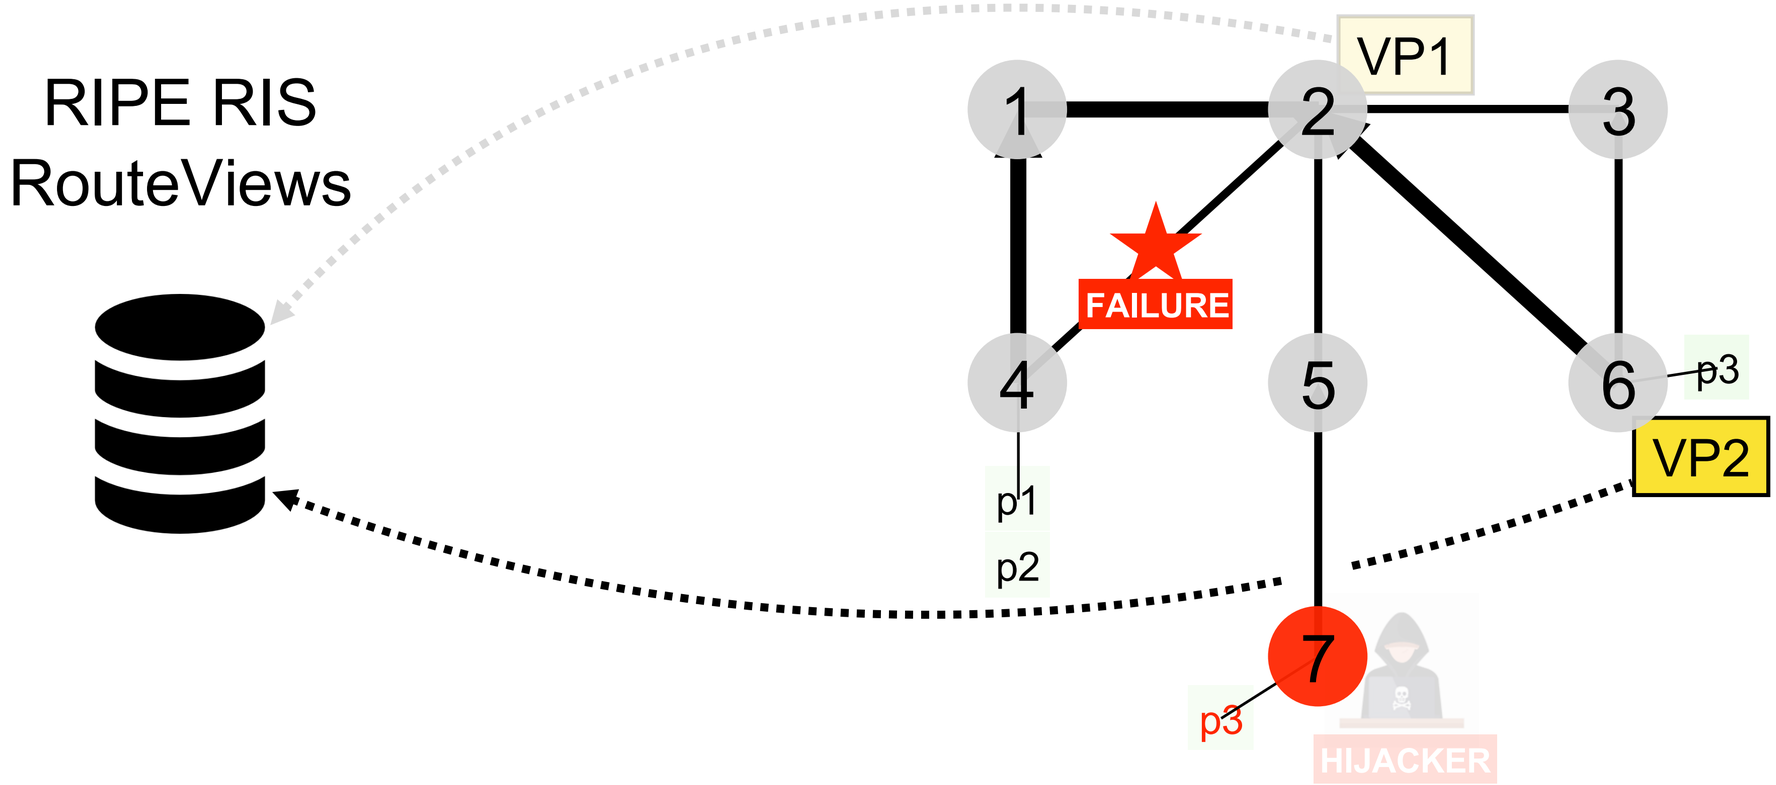

BGP routes can be redundant

Redundant routes

VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4

Redundant routes

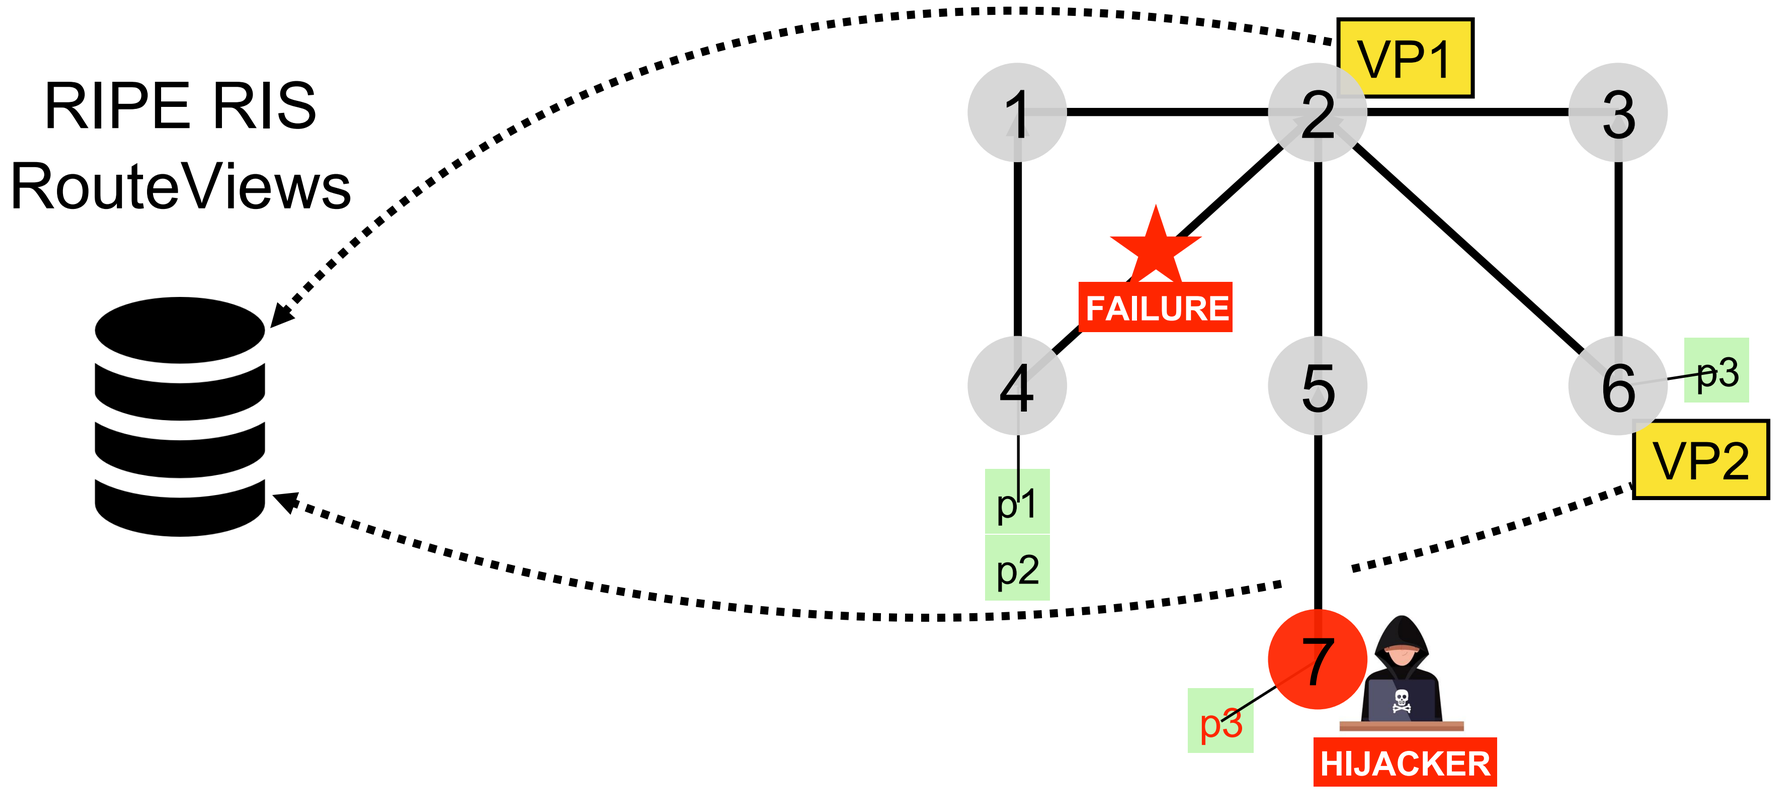

Redundant BGP routes are not so useful

Redundant routes

VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4

Redundant routes

RIPE RIS
RouteViews

The failure is visible in one direction only

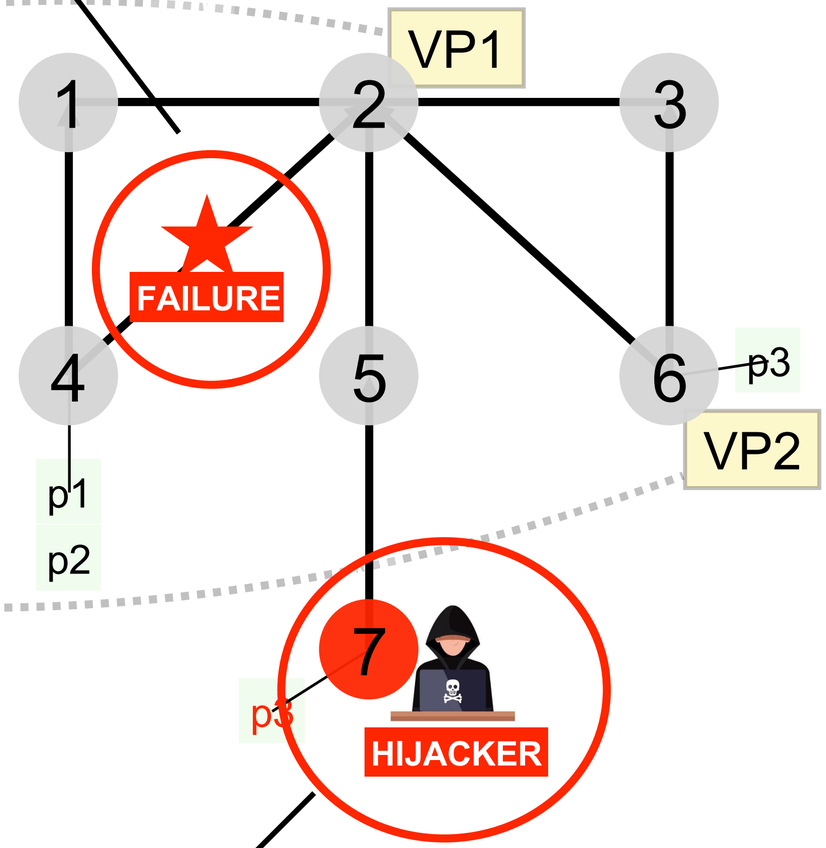

The hijack is not detected

Outline

1. We observe that BGP routes
are often redundant

2. Redundant BGP routes enable
an overshoot-and-discard collection scheme

The “overshoot-and-discard” data collection paradigm

The “overshoot-and-discard” data collection paradigm

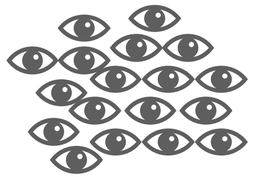

Overshoot: We collect data from as many VPs as possible

To prevent missing important information

Overshoot

Overshoot: collect data from as many VPs as possible

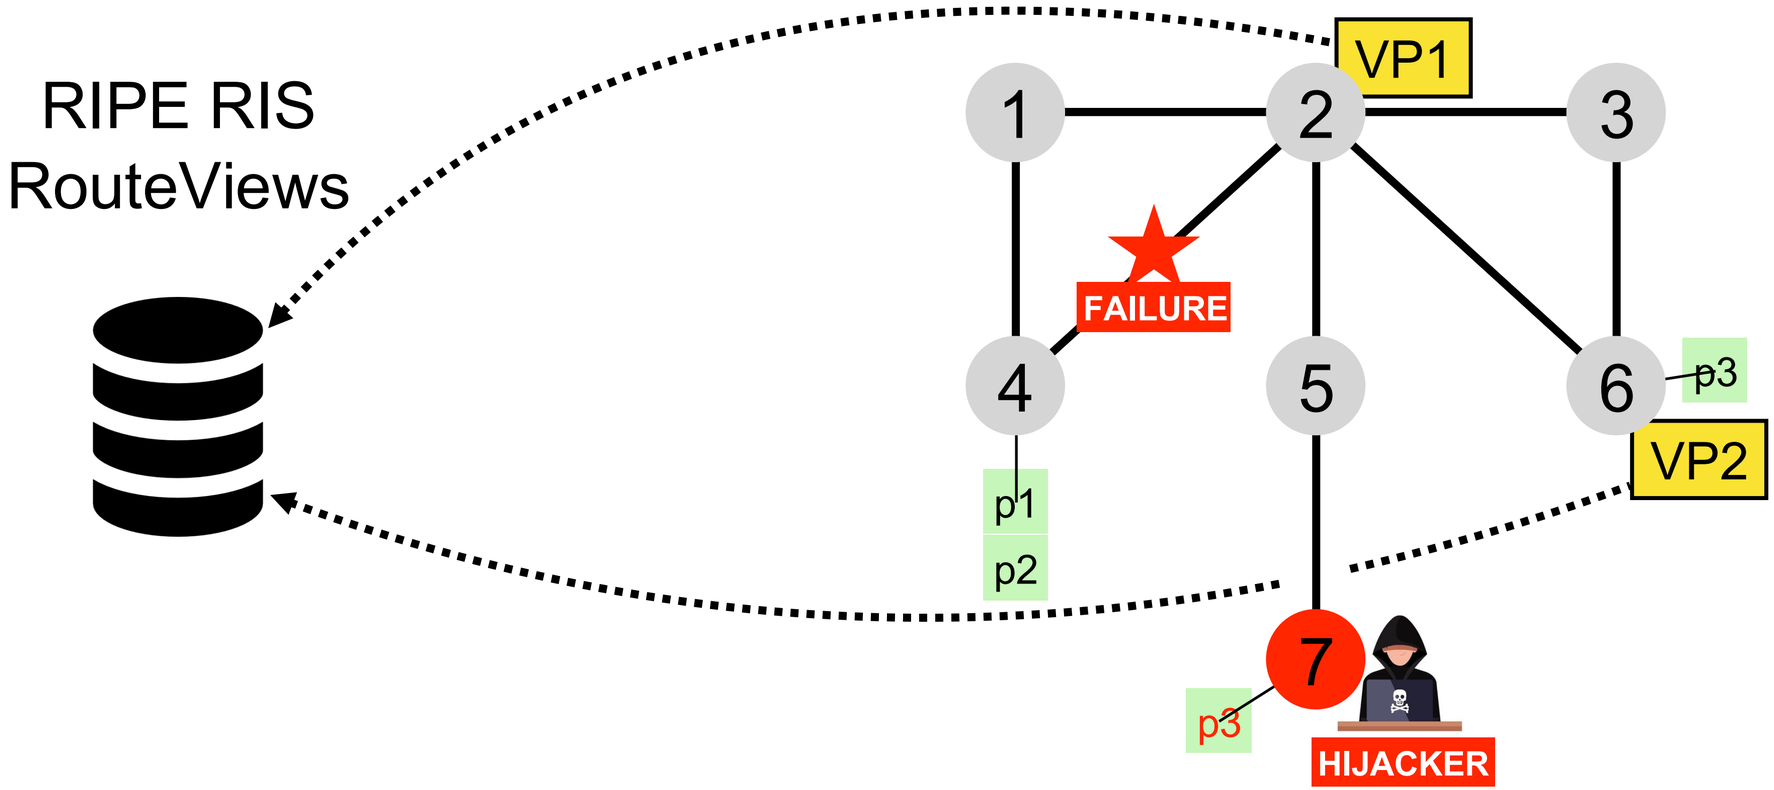

Overshoot: collect data from as many VPs as possible

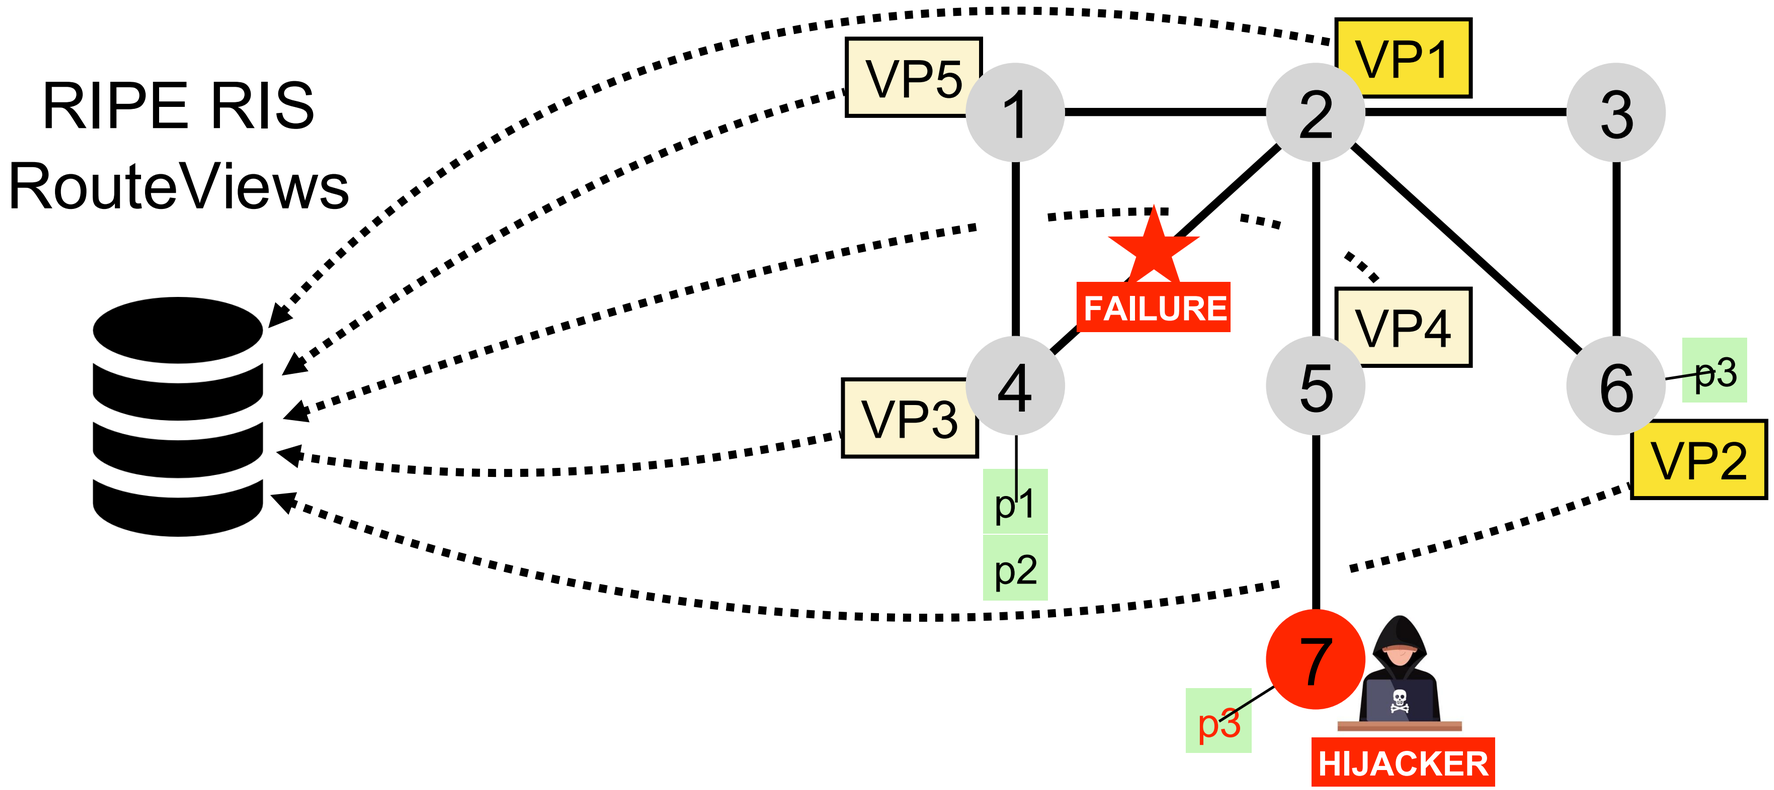

Overshoot: collect data from as many VPs as possible to prevent missing important information

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 7

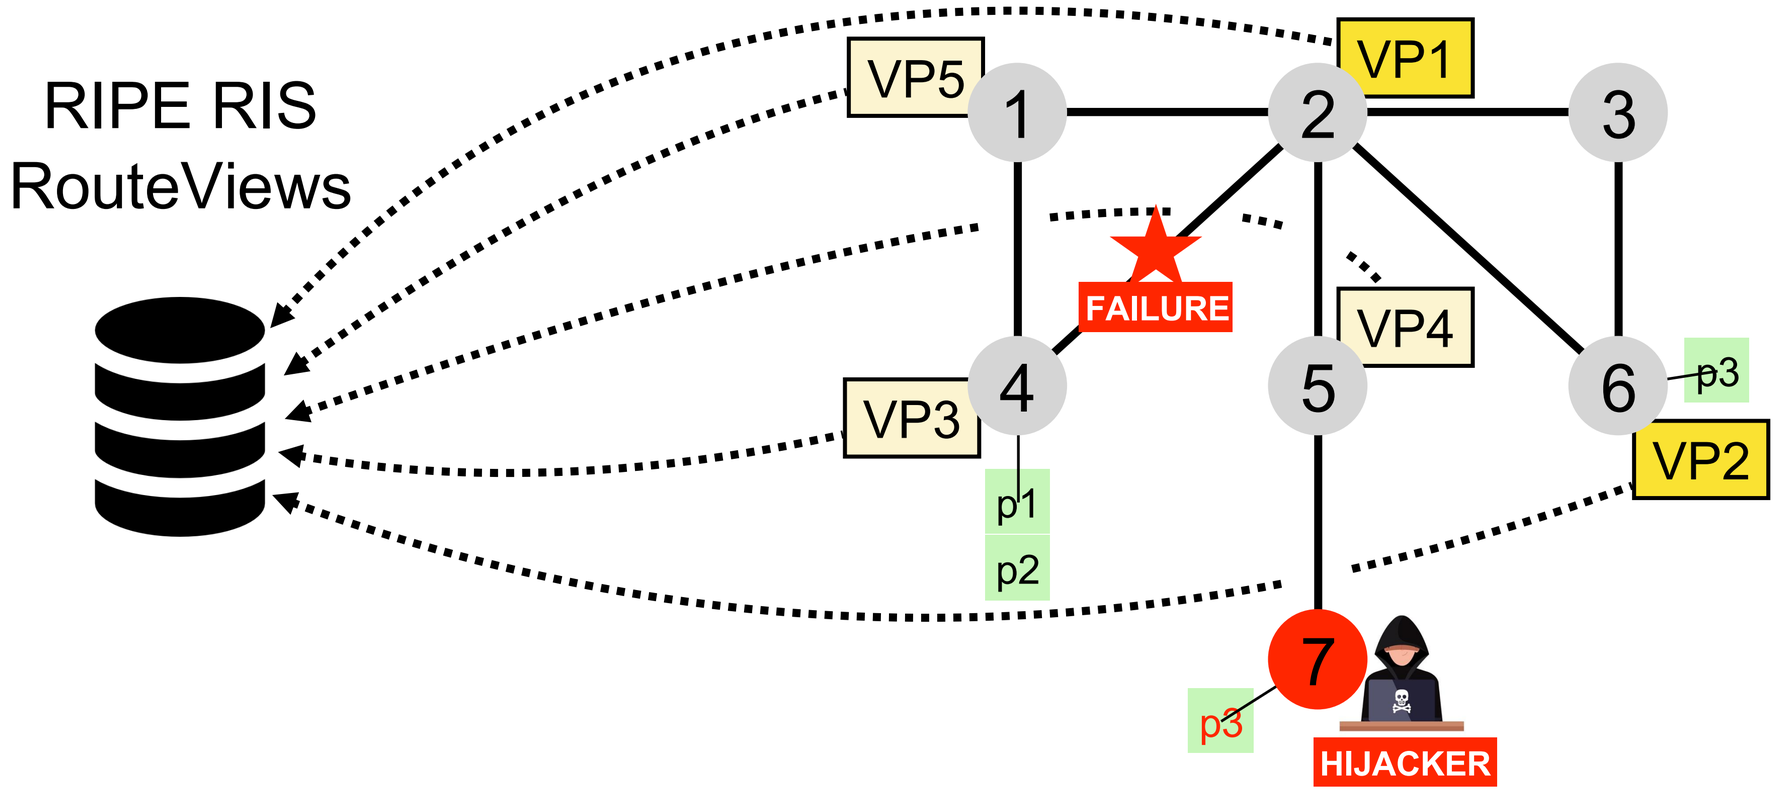

Overshoot: collect data from as many VPs as possible to prevent missing important information

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 7

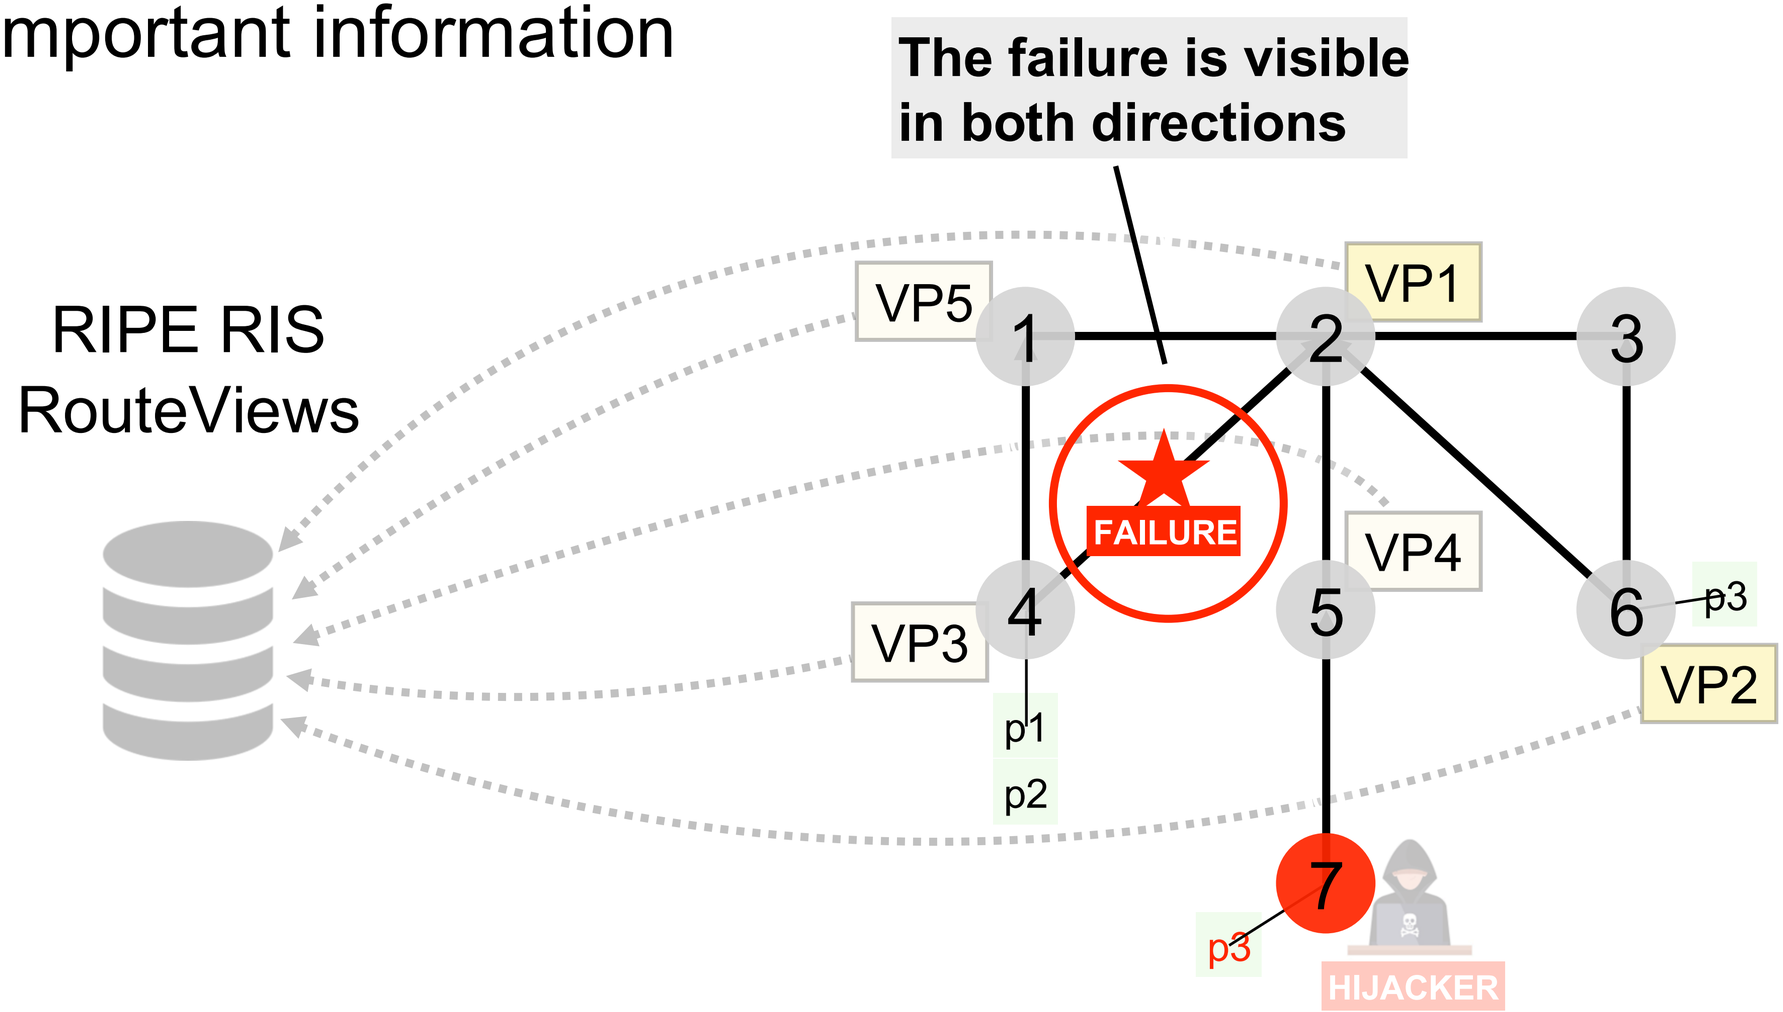

Overshoot: collect data from as many VPs as possible to prevent missing important information

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 7

RIPE RIS
RouteViews

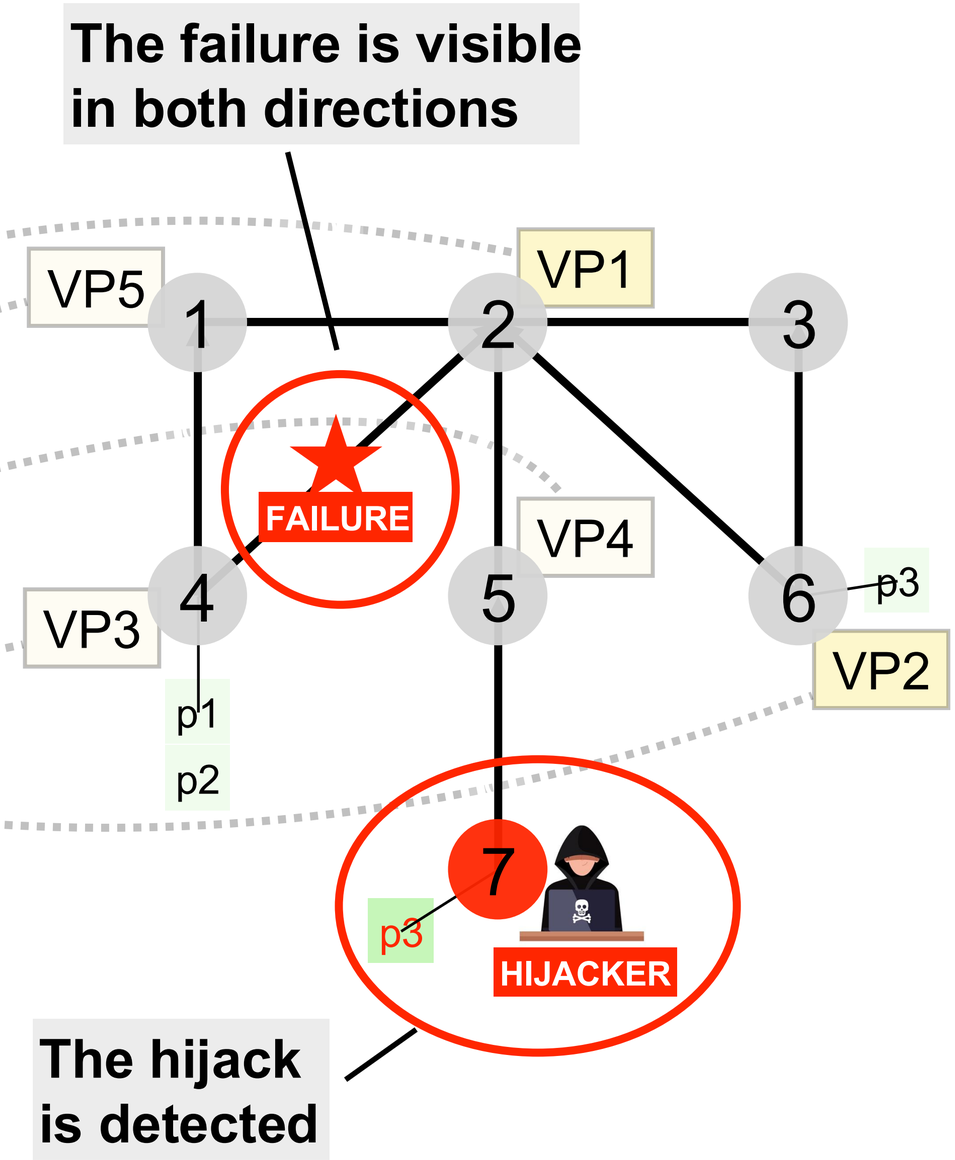

The “overshoot-and-discard” data collection paradigm

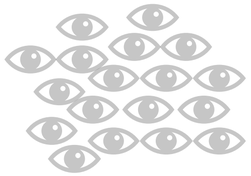

Overshoot: We collect data from as many VPs as possible

To prevent missing important information

Discard: We filter out the redundant BGP routes

To reduce the volume of data collected

Discard: redundant BGP routes are discarded using filters

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 7

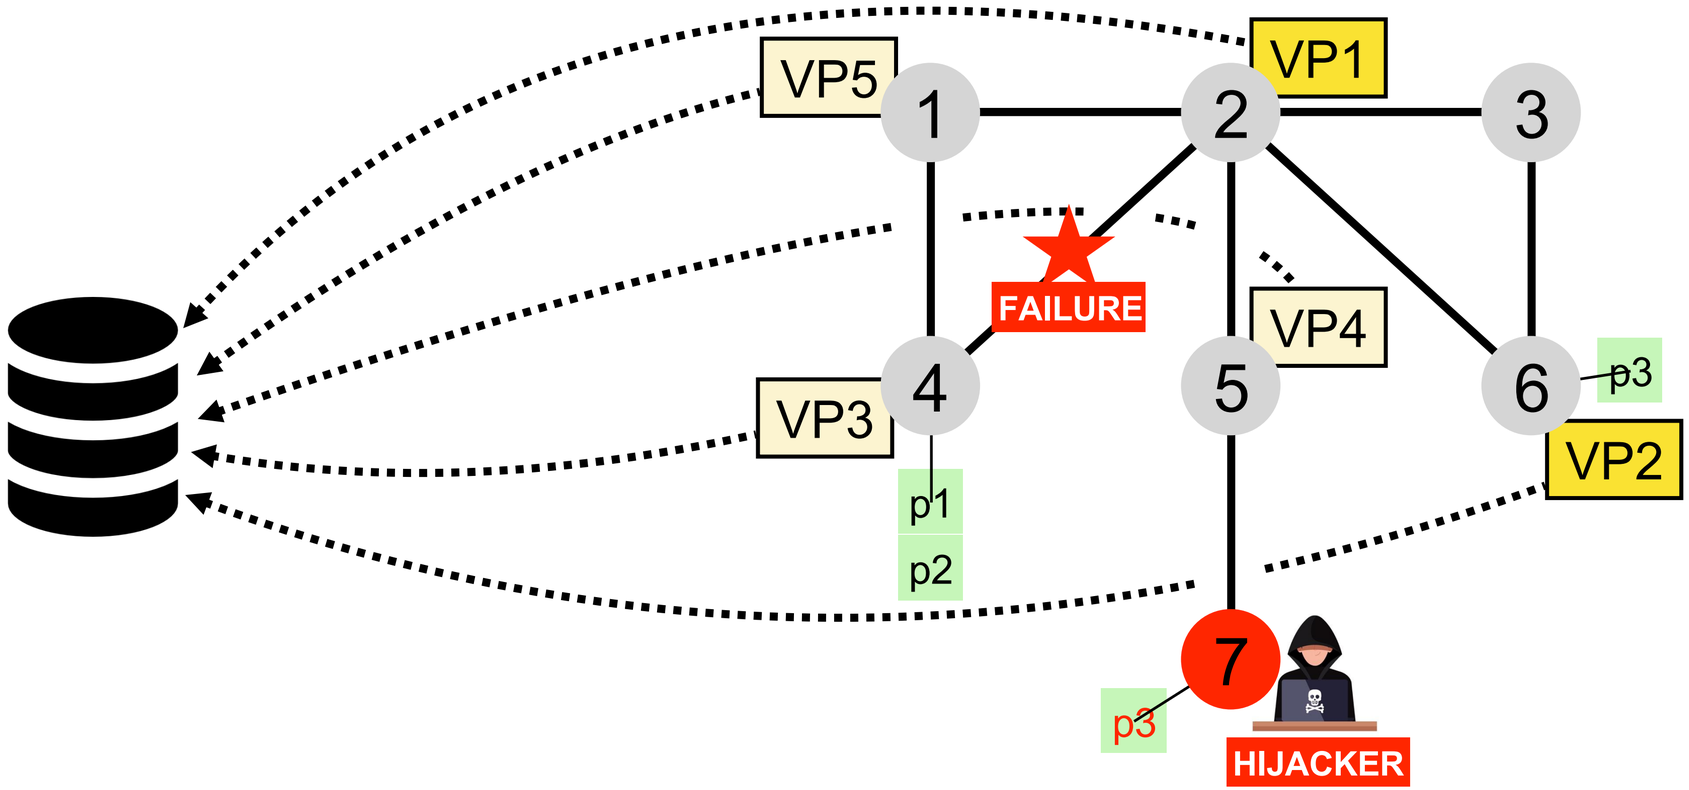

Discard: redundant BGP routes are discarded using filters

Collected routes

VP	prefix	AS path
~~VP1~~	~~p1~~	~~2 1 4~~
~~VP1~~	~~p2~~	~~2 1 4~~
~~VP2~~	~~p1~~	~~6 2 1 4~~
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 7

RIPE RIS
RouteViews

From VP1 drop p1 p2
From VP2 drop p1

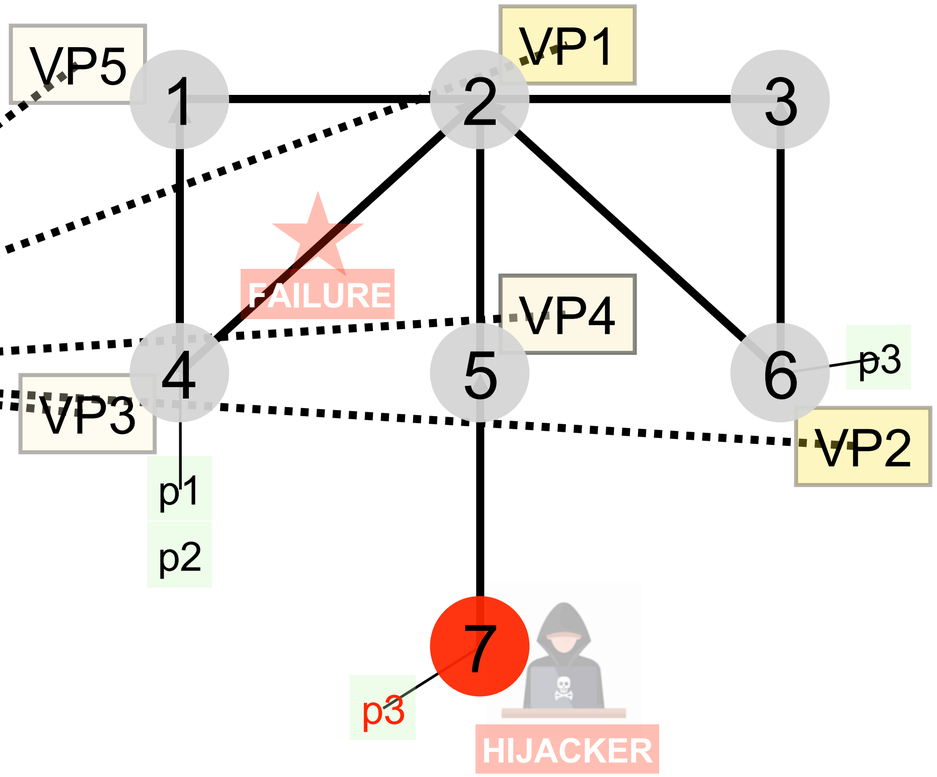

Discard: redundant BGP routes are discarded using filters

Collected routes

VP	prefix	AS path
~~VP1~~	~~p1~~	~~2 1 4~~
~~VP1~~	~~p2~~	~~2 1 4~~
~~VP2~~	~~p1~~	~~6 2 1 4~~
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 7

RIPE RIS
RouteViews

From VP1 drop p1 p2
From VP2 drop p1

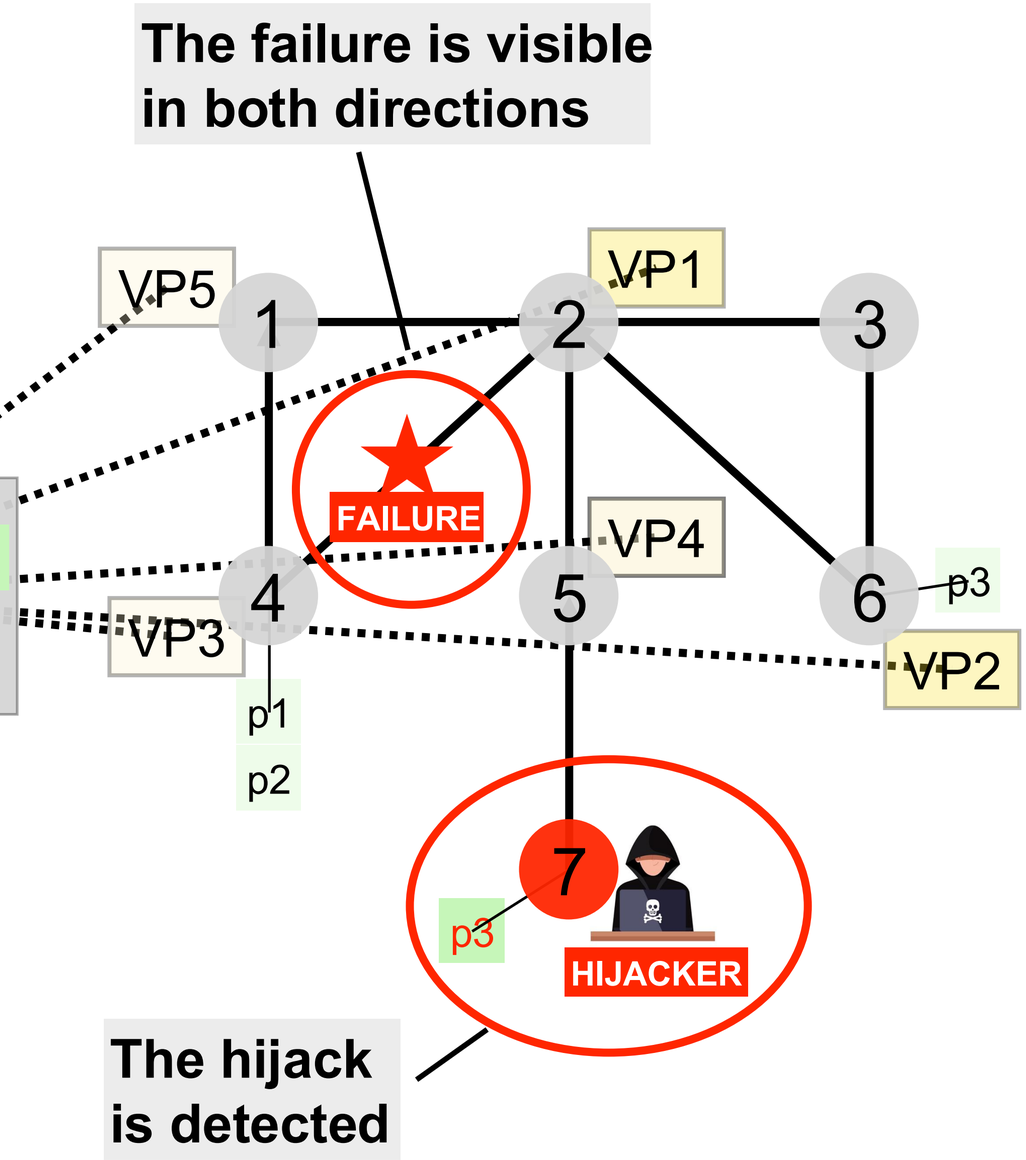

Unfortunately, some other objectives become unachievable with the deployed filters

Collected routes

VP	prefix	AS path
~~VP1~~	~~p1~~	~~2 1 4~~
~~VP1~~	~~p2~~	~~2 1 4~~
~~VP2~~	~~p1~~	~~6 2 1 4~~
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 7

RIPE RIS
RouteViews

From VP1 drop p1 p2
From VP2 drop p1

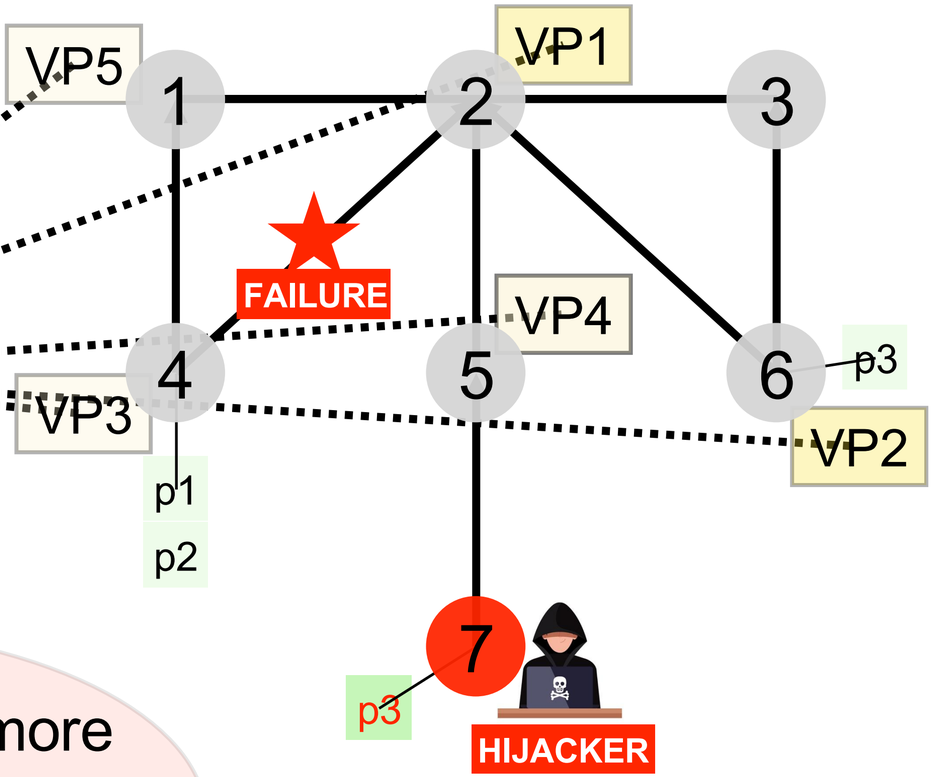

I do not know anymore if p1 p2 share the same BGP attributes

GILL identifies non redundants VPs (anchors), and keeps all data from these VPs

Collected routes

VP	prefix	AS path
~~VP1~~	~~p1~~	~~2 1 4~~
~~VP1~~	~~p2~~	~~2 1 4~~
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 7

RIPE RIS
RouteViews

From VP1 drop p1 p2

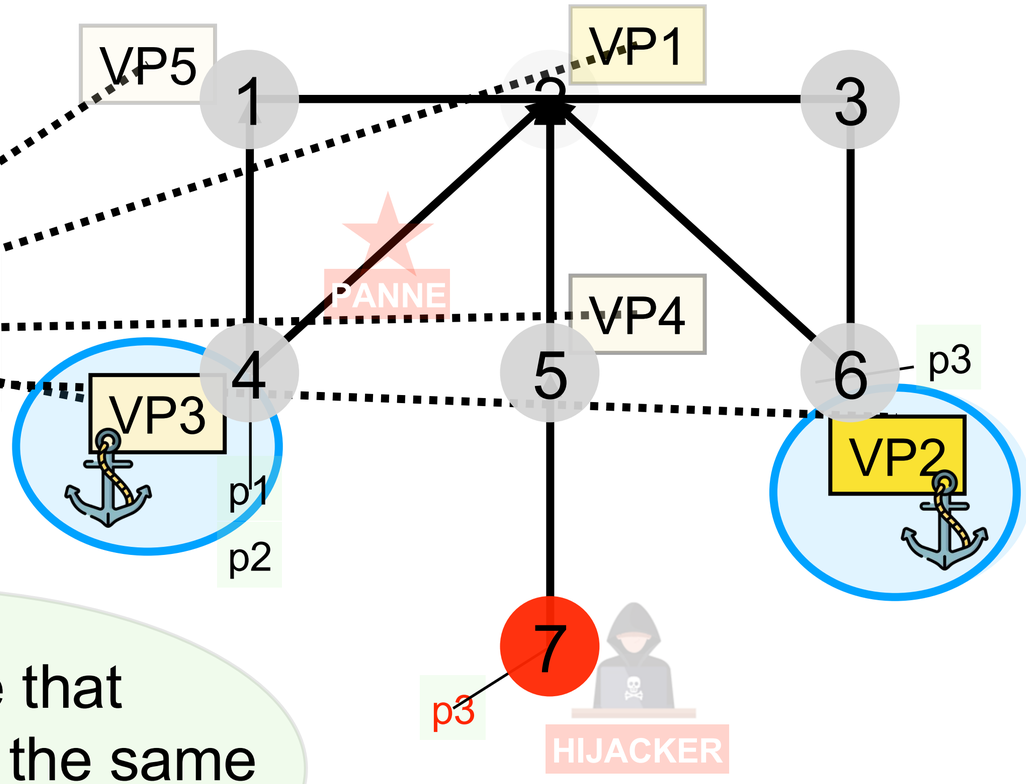

Now I can see that p1 p2 share the same BGP attributes

Outline

1. We observe that BGP routes
are often redundant

2. Redundant BGP routes enable
an overshoot-and-discard collection scheme

3. ***GILL***: a system that measures redundancy in BGP data
and uses it to generate filters that retain only the useful routes

GILL uses custom BGP collectors (one per peer) optimized to collect BGP route updates

GILL supports
10k+ BGP sessions

(see paper)

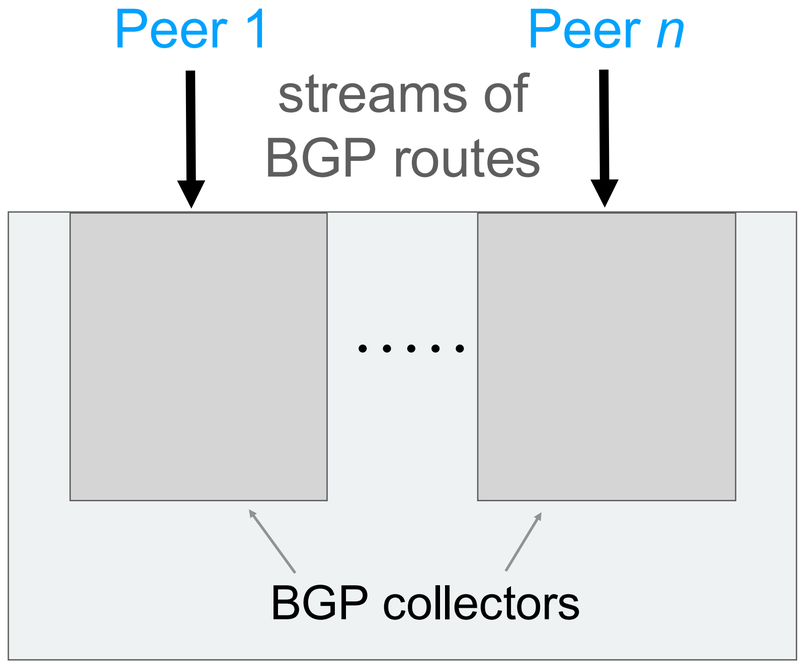

GILL discards the BGP routes that match filters and retains all the others

GILL installs the filters in the BGP collectors

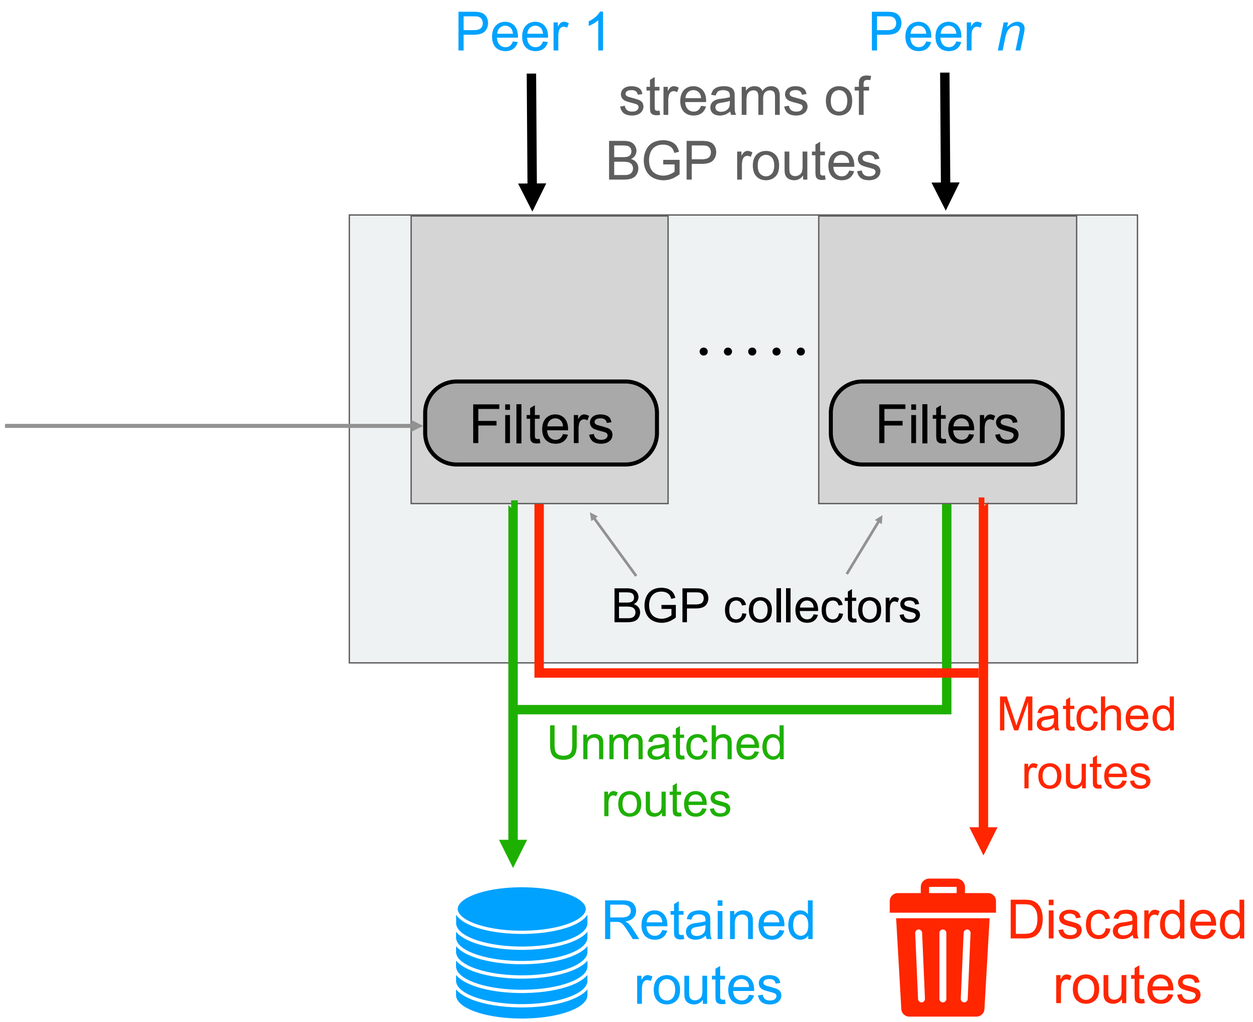

GILL's filters generation is data driven

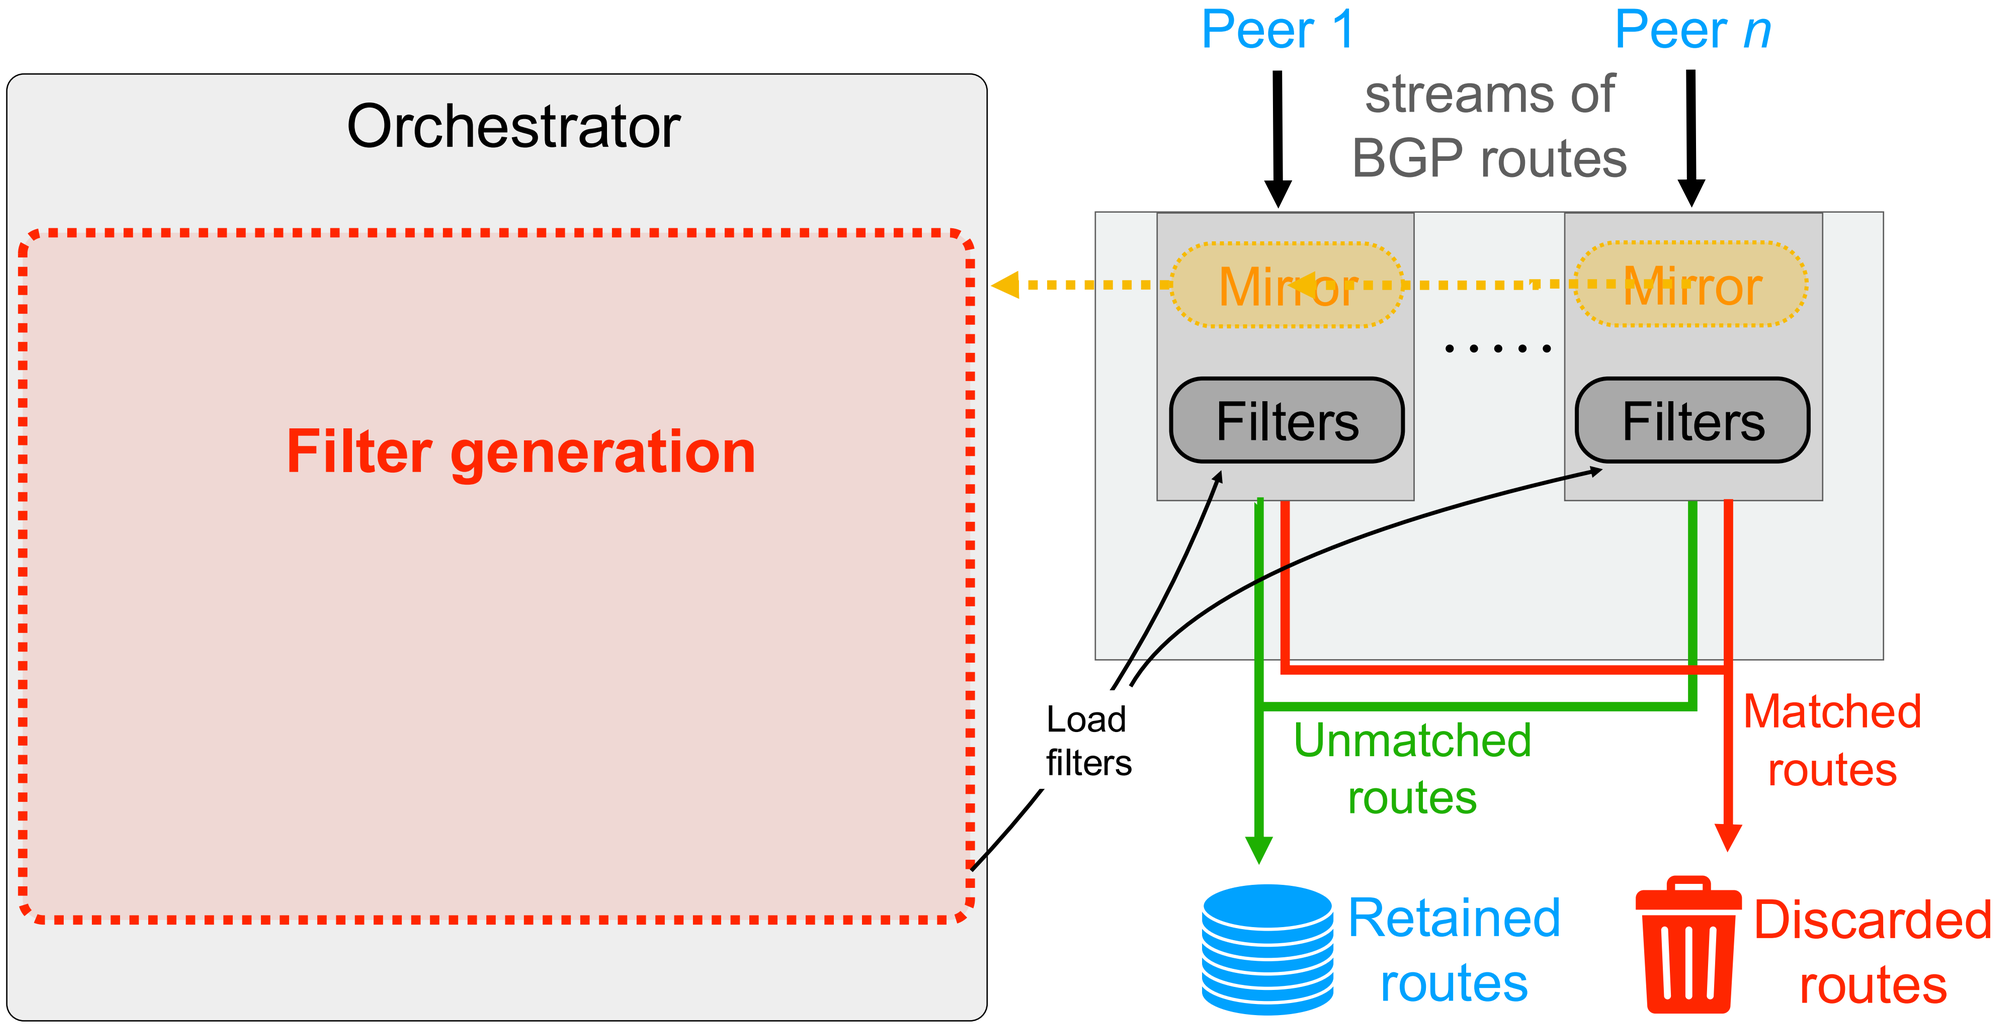

GILL's filters generation is data driven

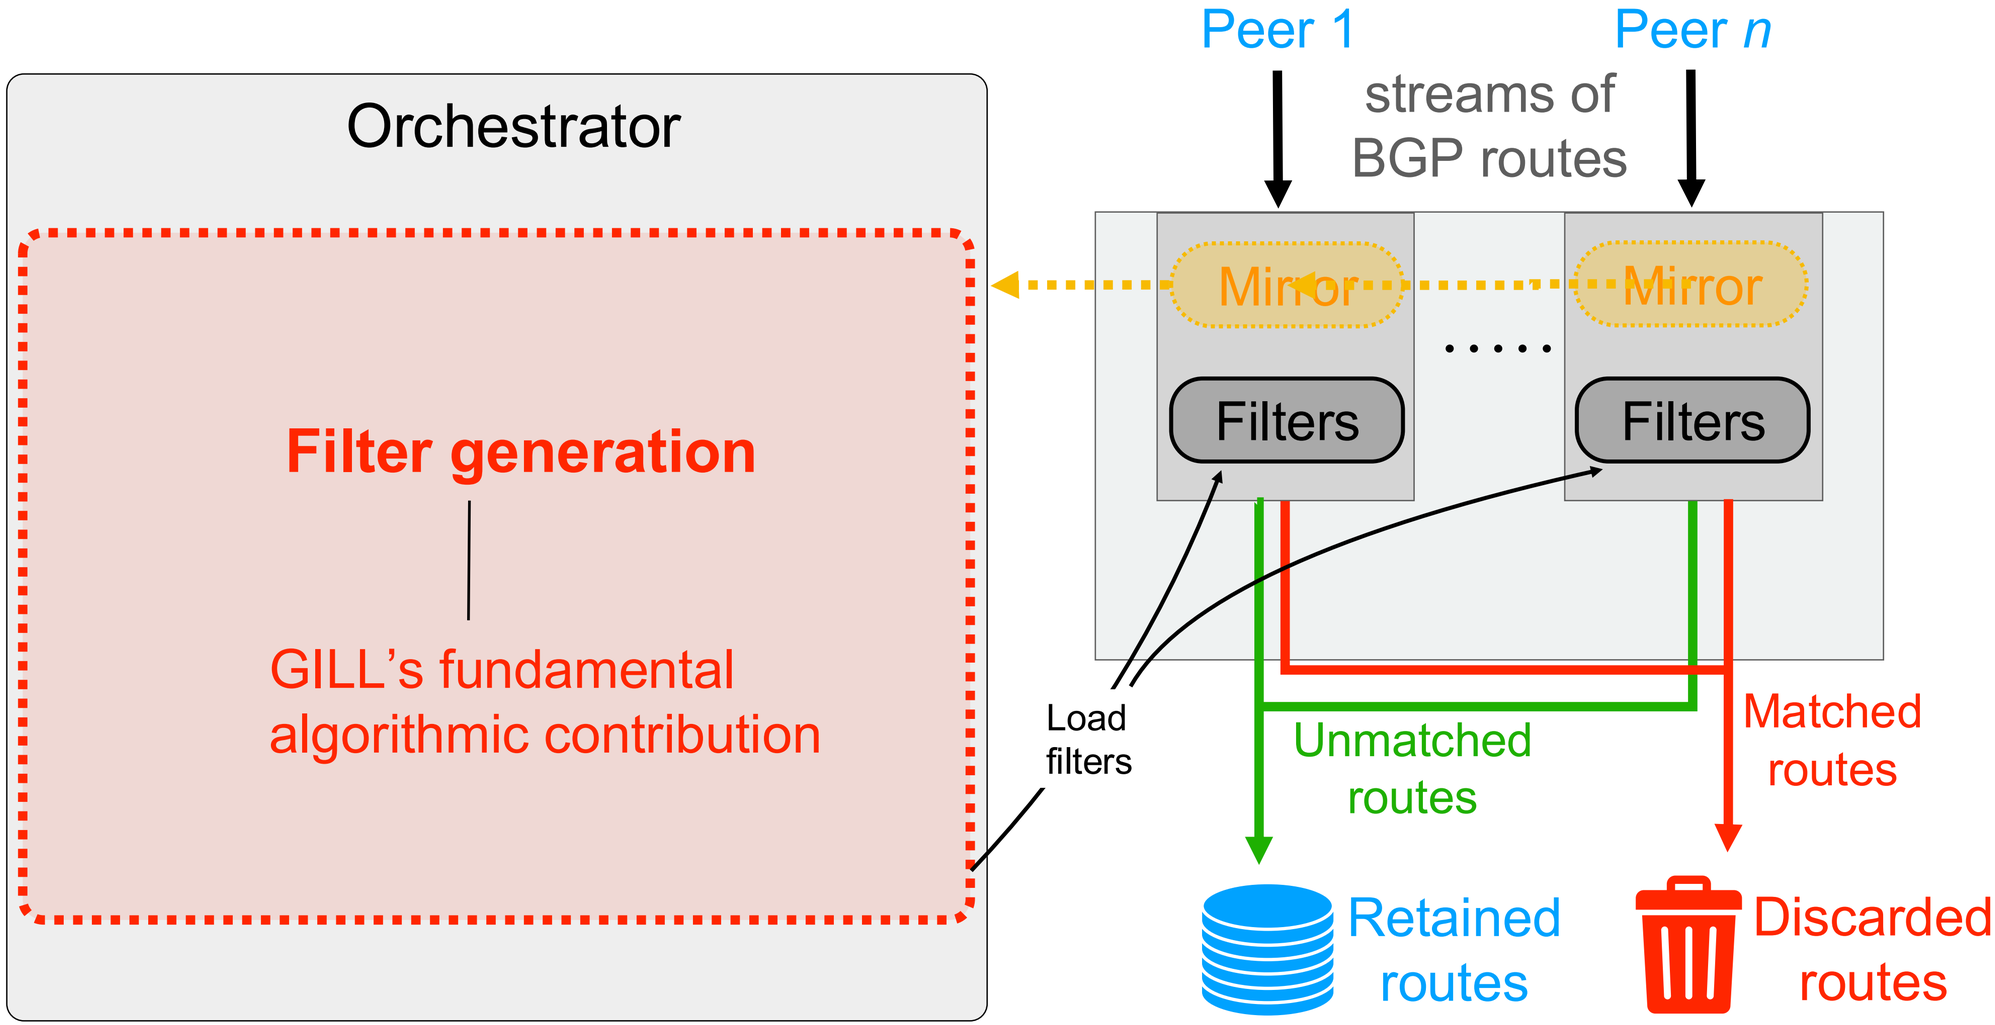

GILL's filters generation has three steps

Step#1: Identifying redundancy in past BGP data

Challenging as there is no consensus on how to define redundancy in BGP data

GILL's filters generation has three steps

Step#1: Identifying redundancy in past BGP data
Challenging as there is no consensus on how to define redundancy in BGP data

Step#2: Predict future redundancies in BGP data
Fortunately, redundancy in BGP data is stable over time (see paper)

GILL's filters generation has three steps

Step#1: Identifying redundancy in past BGP data
Challenging as there is no consensus on how to define redundancy in BGP data

Step#2: Predict future redundancies in BGP data
Fortunately, redundancy in BGP data is stable over time (see paper)

Step#3: Generate filters that retain useful bits of the BGP data
Challenging because we cannot predict what users will do with the data

GILL identifies past redundant routes, predicts the future ones, and builds filters that retains only the most useful future routes

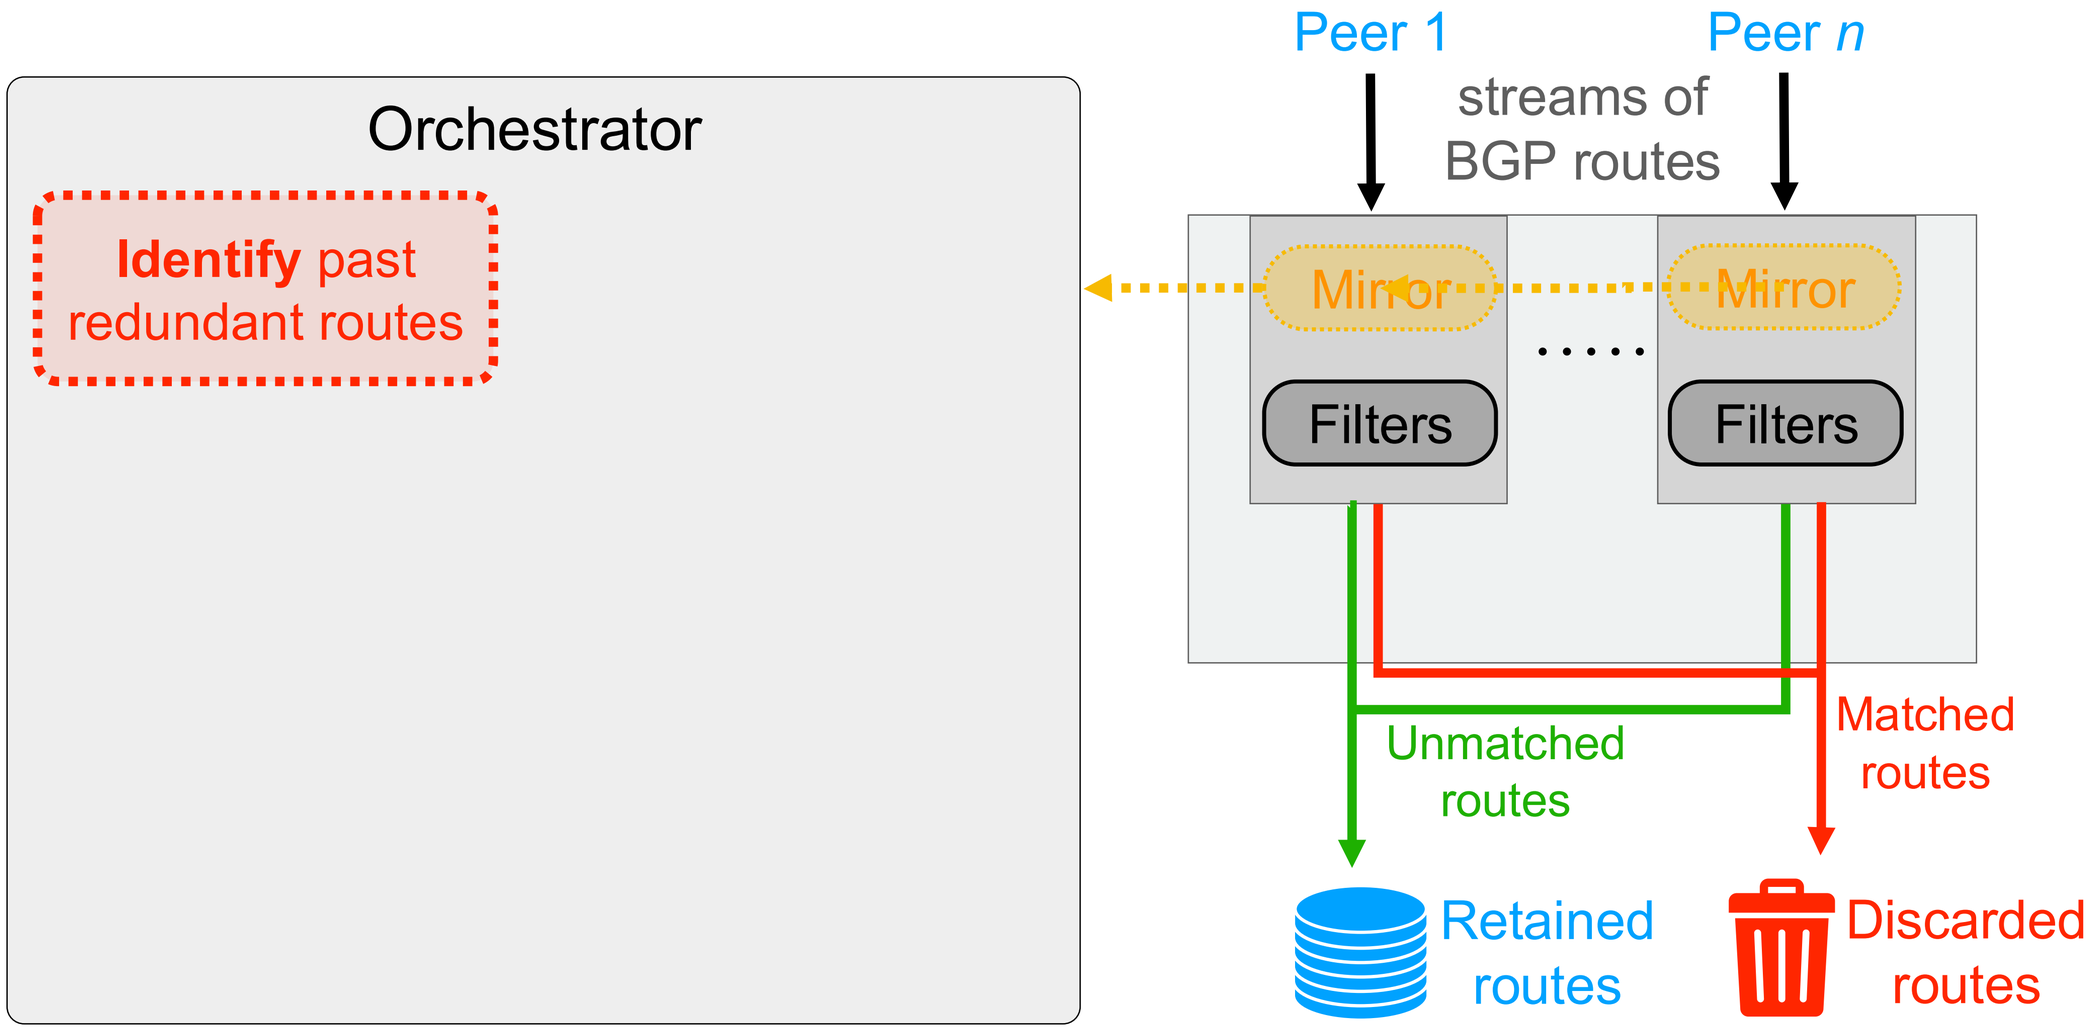

GILL identifies past redundant routes, predicts the future ones, and builds filters that retains only the most useful future routes

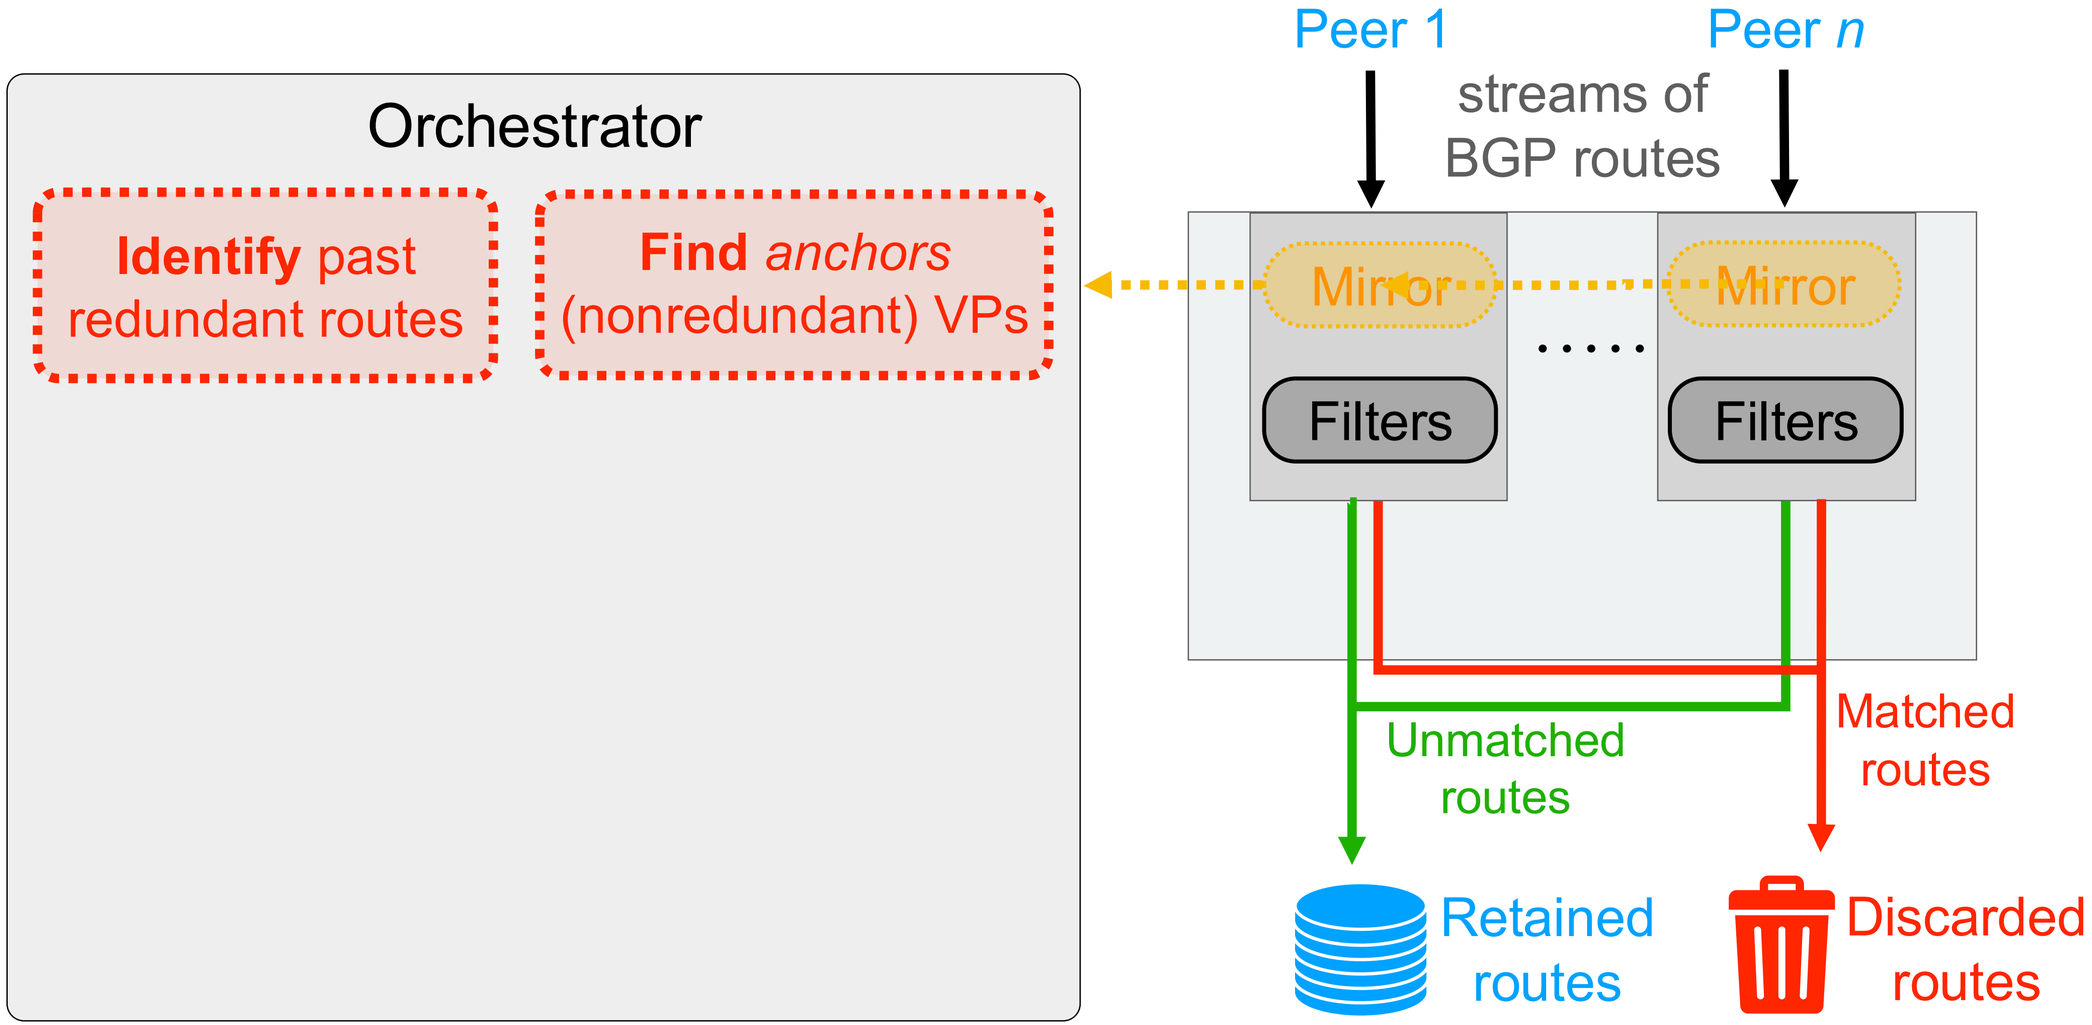

GILL identifies past redundant routes, **predicts the future ones**,
and builds filters that retains only the most useful future routes

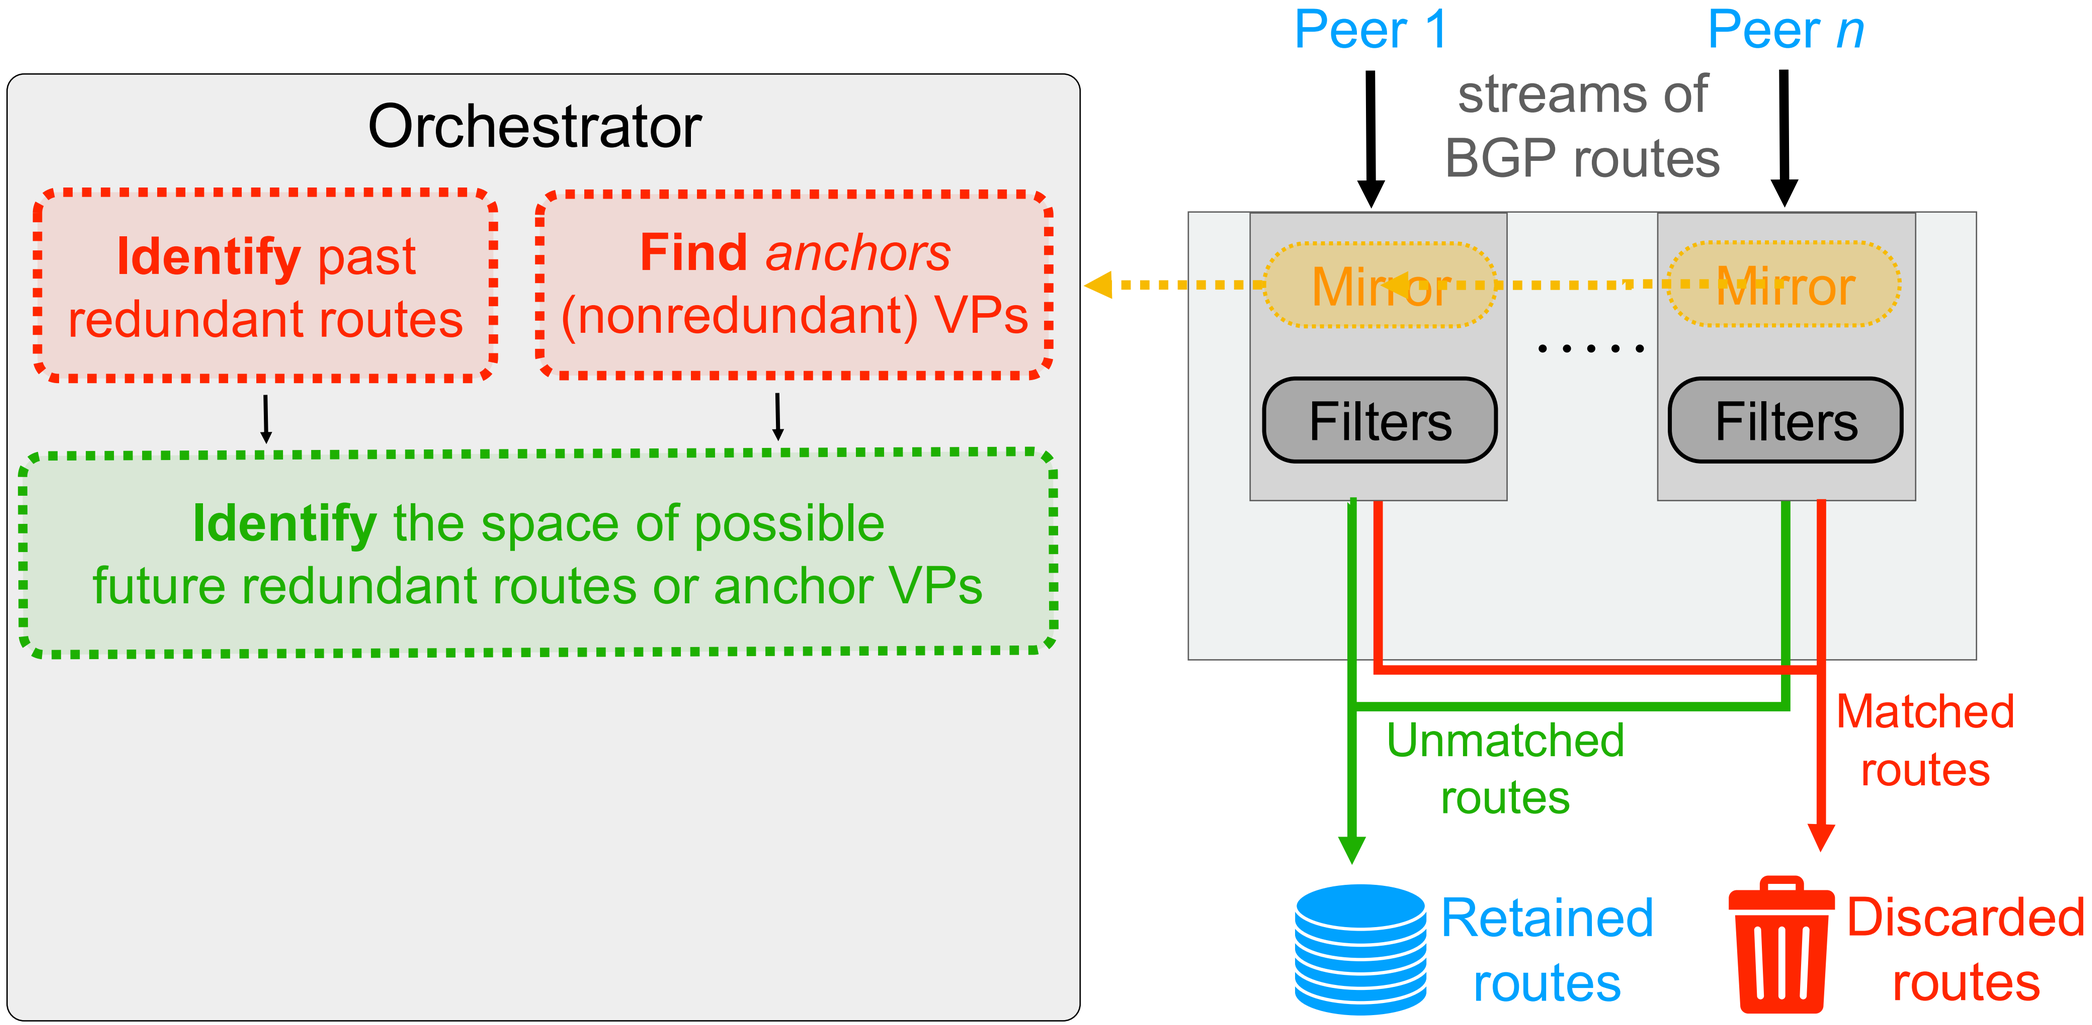

GILL identifies past redundant routes, predicts the future ones, and builds filters that retains only the most useful future routes

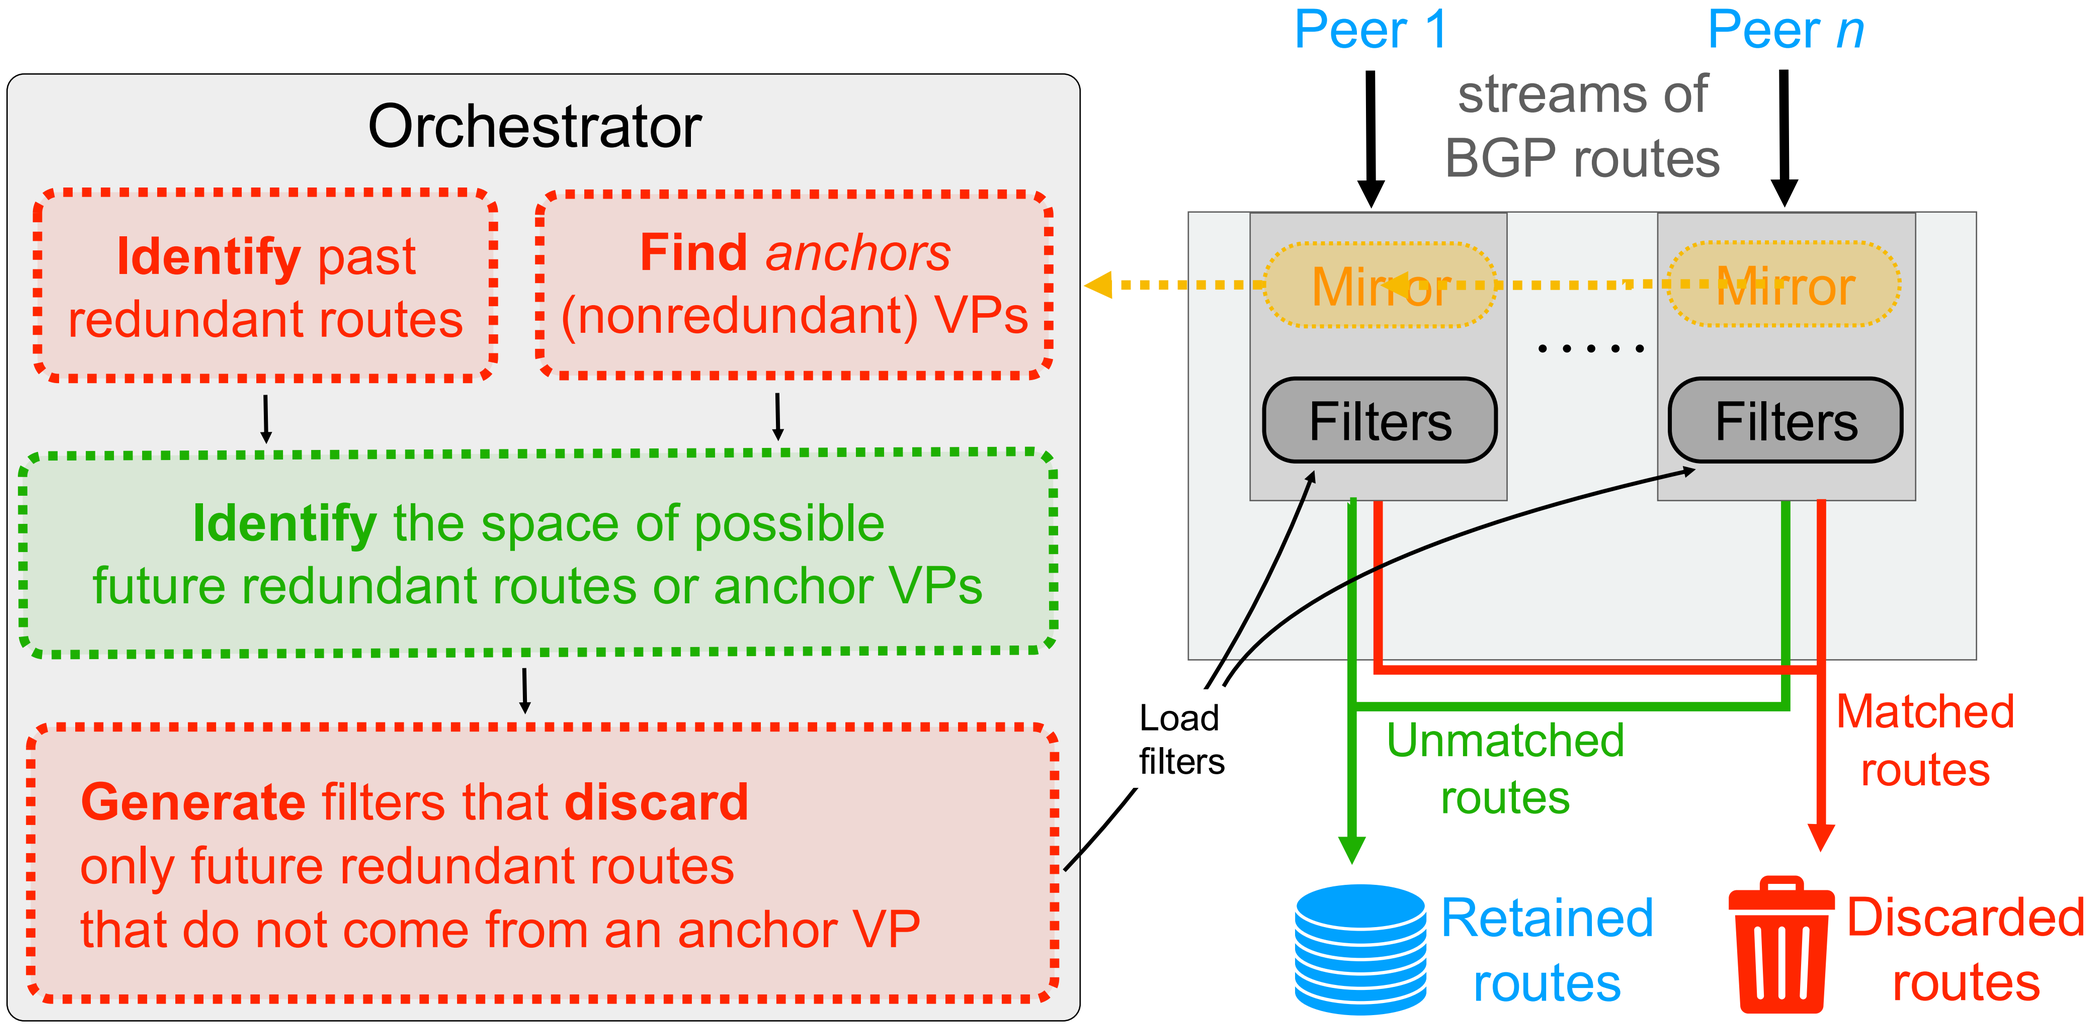

GILL regularly refresh the filters to maintain them accurate

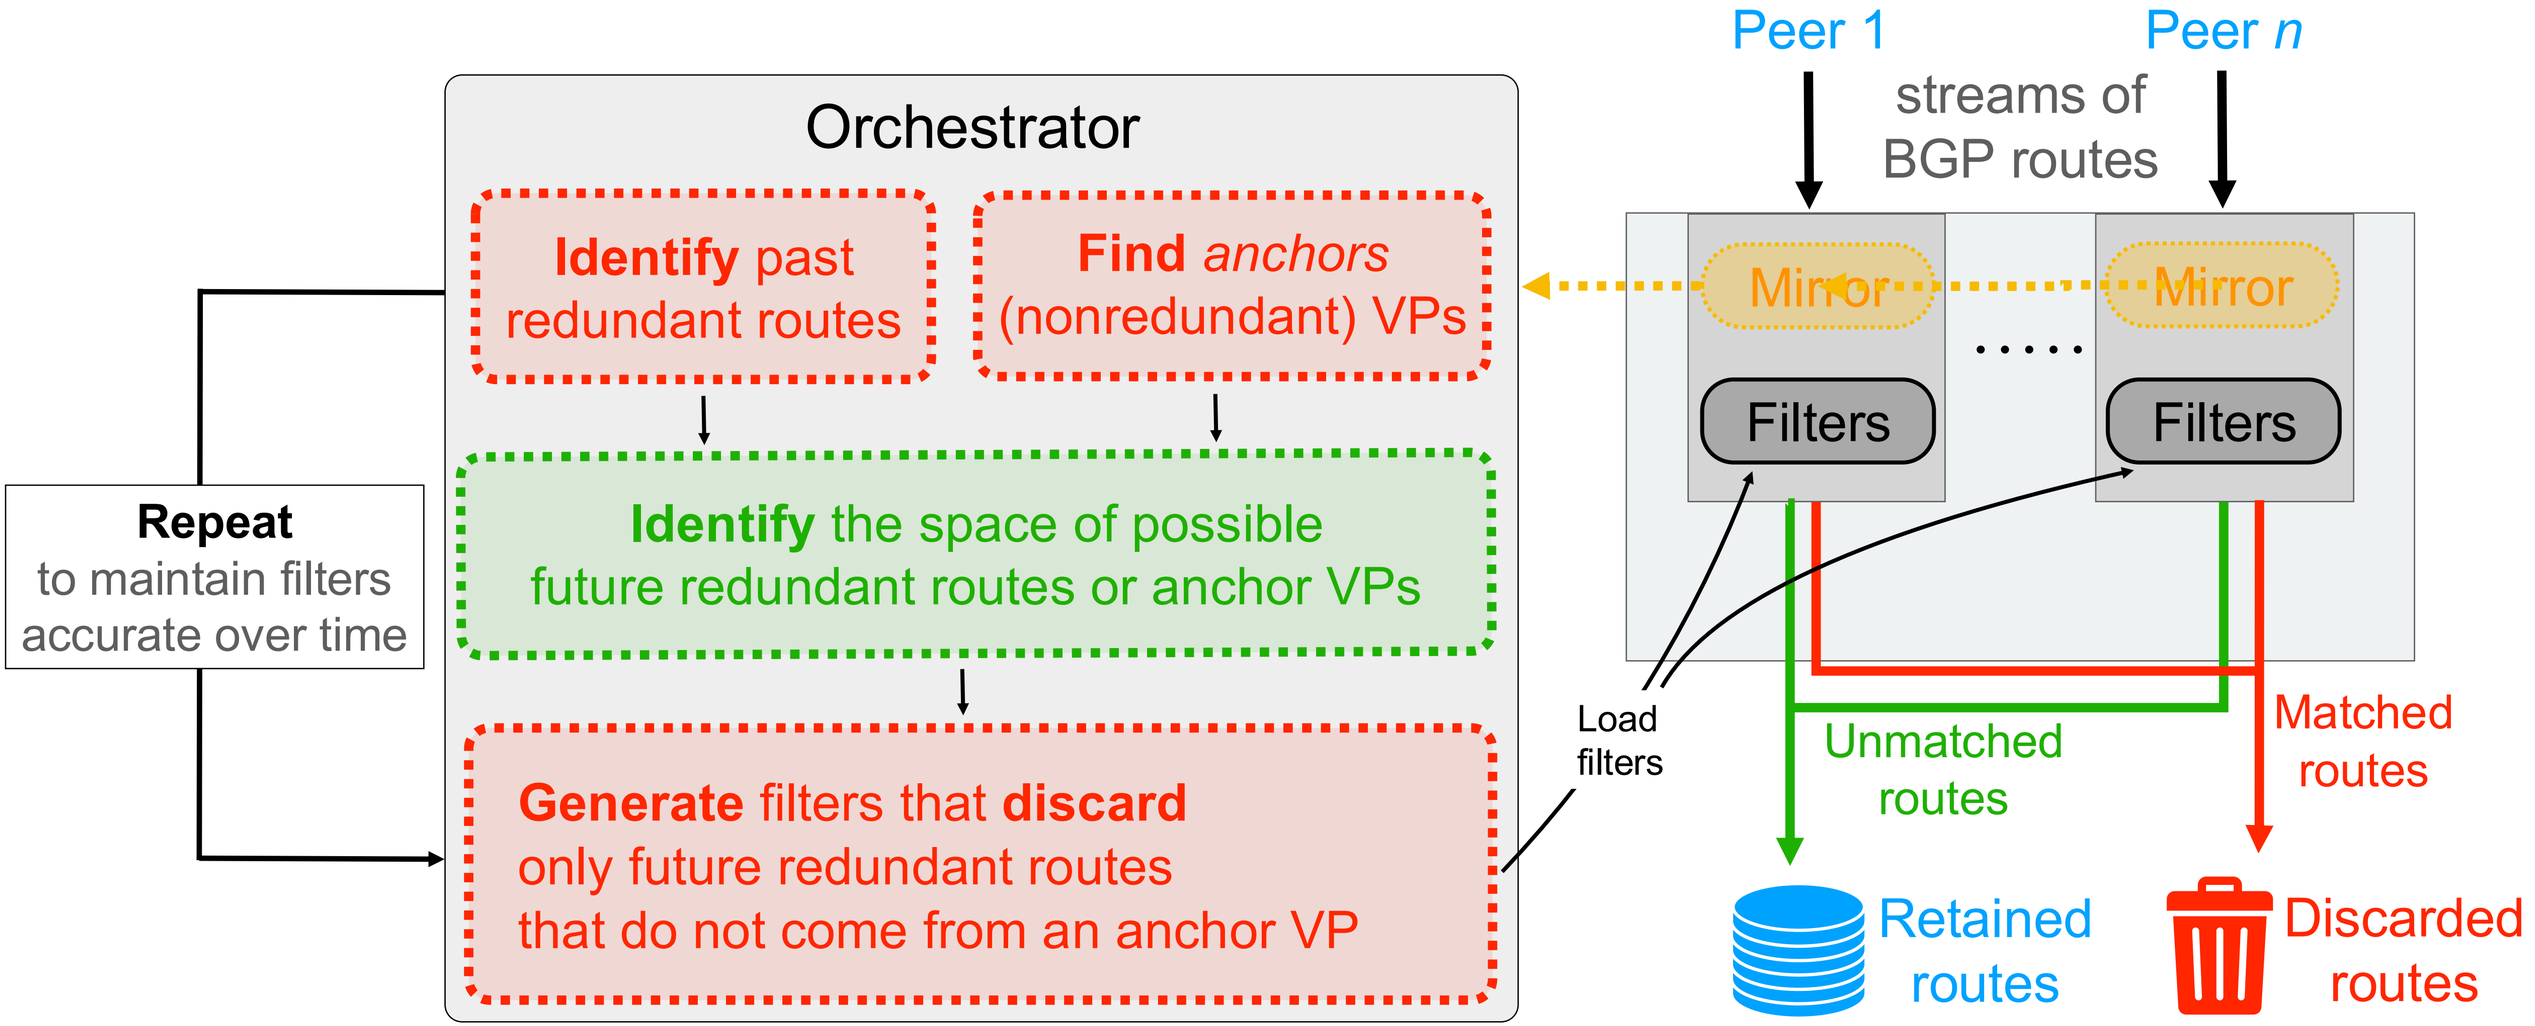

GILL regularly refresh the filters to maintain them accurate

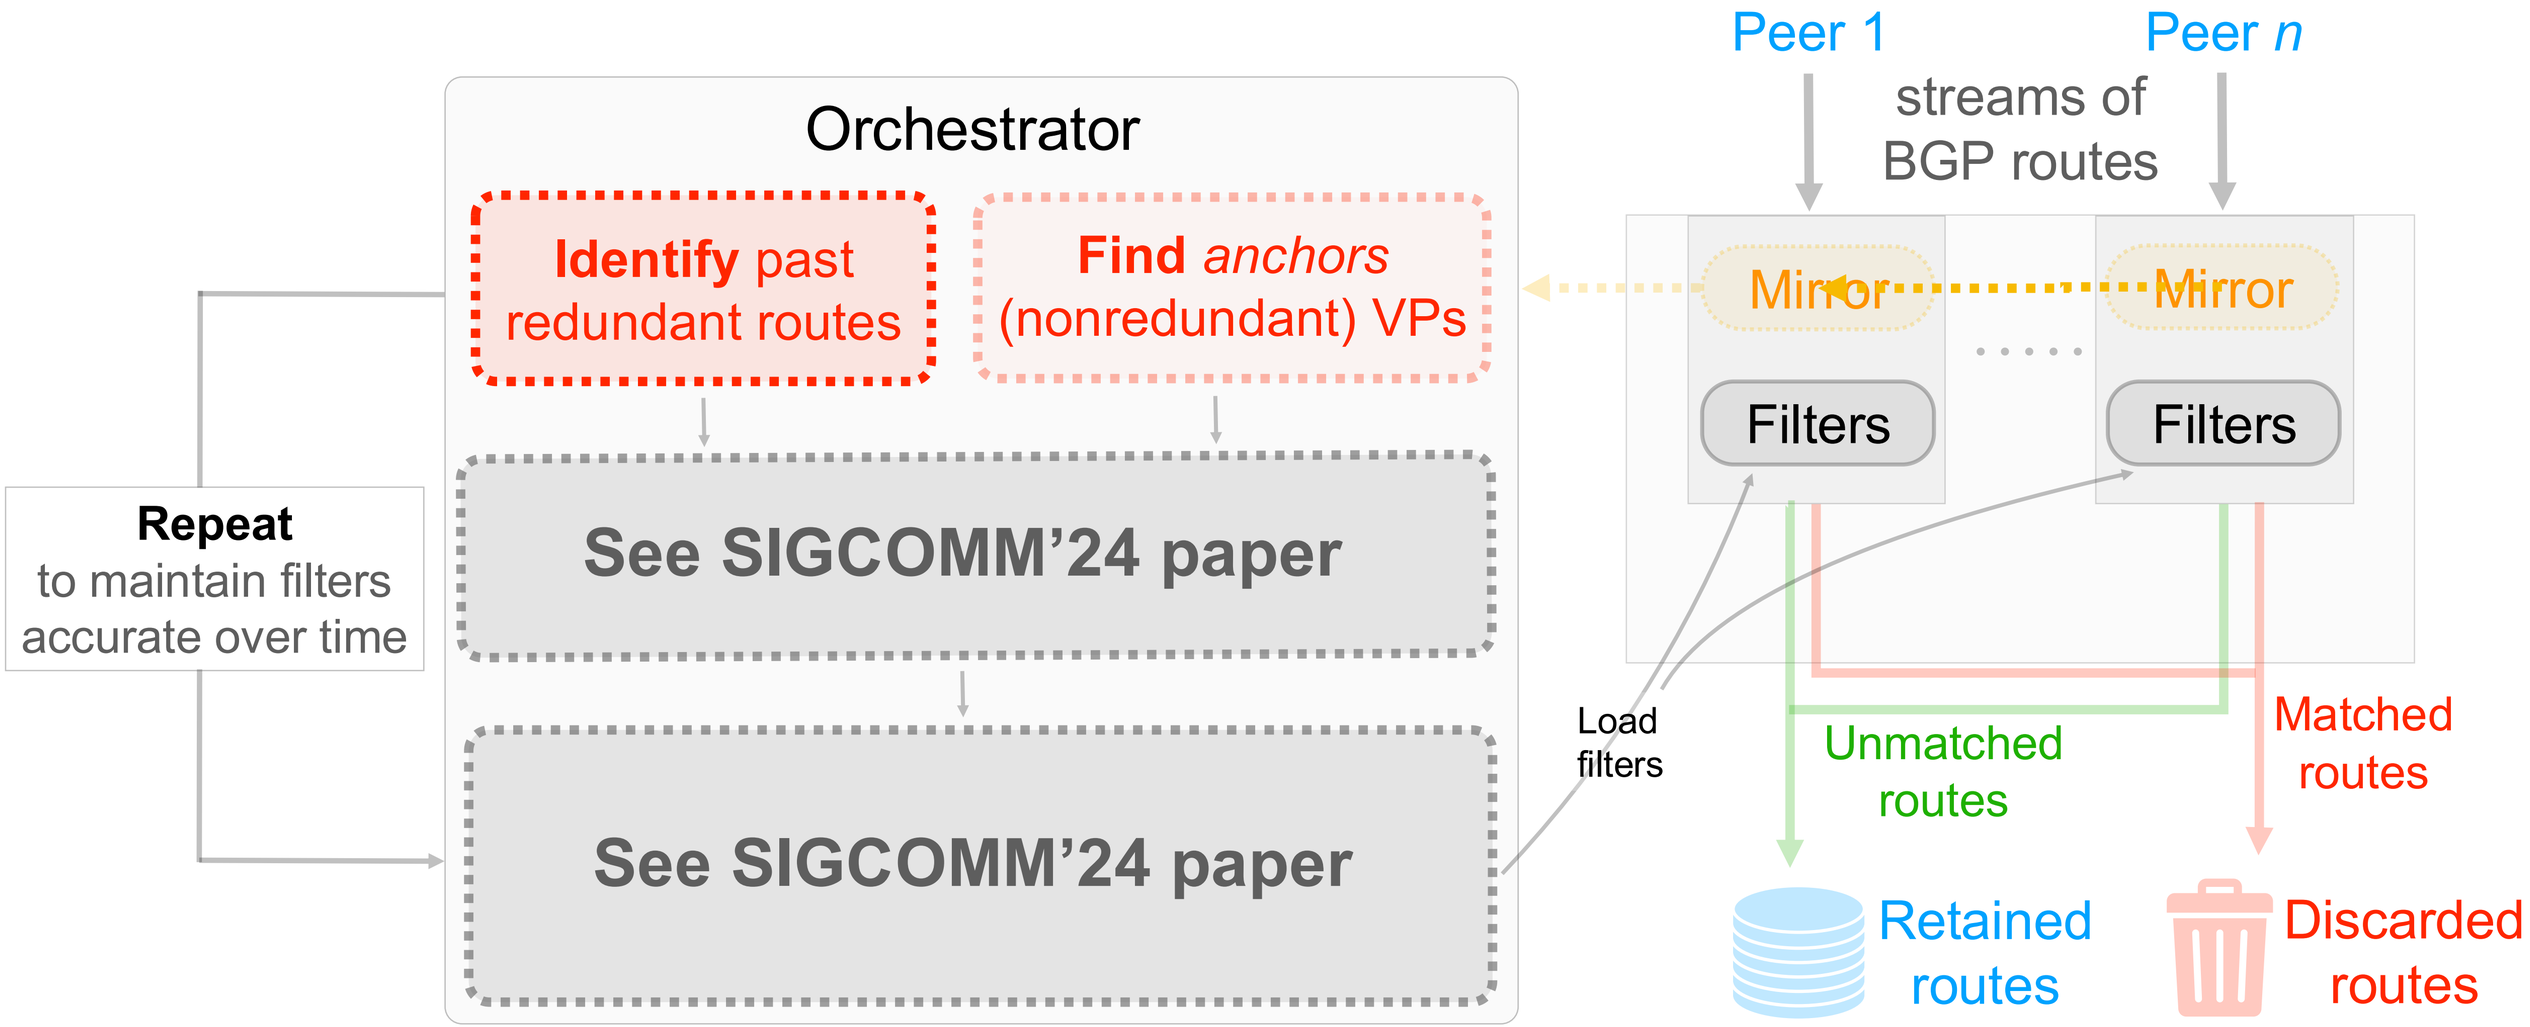

GILL *positively impacts many studies*

Long term impact

S*imulations with a 50% coverage*

x3 p2p links identified
x2 localised failures
+9% detected BGP hijacks

Short term impact

replication of past studies with our sampling algorithms

+15% AS business relationships inferred
4x precision of inferred hijacks

***GILL** positively impacts many studies*

Long term impact

***S**imulations with a 50% coverage*

x3 p2p links identifies
x2 localised failures
+9% detected BGP hijacks

Without using more data!

Short term impact

replication of past studies with our sampling algorithms

+15% AS business relationships inferred
4x precision of inferred hijacks

We developed ***GILL***, a system that collects only the most useful BGP updates

We highlighted the importance of having a higher VP coverage

Many BGP events are currently out of our radars

We designed two definitions of redundancy in BGP data

That work regardless of what the user wants to do with the data

We implemented ***GILL*** in C and python

A prototype of ***GILL*** is running at https://bgproutes.io

bgproutes.io

Also a lossless compression technique

To further reduce storage

And a fine grained API

Download only the data that you need

Aggregation of BGP data

Download the data from a single location using a unified API

You can peer with us!

Data providers

Name	Number of vantage points
bgroutes.io	**v4:** 33 **v6:** 8
RIPE RIS	**v4:** 833 **v6:** 695
RouteViews	**v4:** 745 **v6:** 706
PCH	**v4:** 2150 **v6:** 397
CGTF RIS	**v4:** 16 **v6:** 51

Thanks

- Silicon valley foundation (CG)
- ISOC grant
- RIPE RIS community fund
- Région Grand Est, ANR THIA ArtIC
- NGI Zero grant

Publications

Thomas Alfroy, Thomas Holterbach, Thomas Krenc, K. C. Claffy, Cristel Pelsser (2024). The Next Generation of BGP Data Collection Platforms 🏆. Proceedings of the ACM SIGCOMM 2024 Conference.

Thomas Holterbach, Thomas Alfroy, Amreesh D. Phokeer, Alberto Dainotti, Cristel Pelsser (2024). A System to Detect Forged-Origin Hijacks. 21th USENIX Symposium on Networked Systems Design and Implementation (NSDI 24).

Thomas Alfroy, Thomas Holterbach, Thomas Krenc, KC Claffy, Cristel Pelsser (2023). Internet Science Moonshot: Expanding BGP Data Horizons. Proceedings of the 22nd ACM Workshop on Hot Topics in Networks.