RIPE 88
Kraków, Poland
20 - 24 May 2024

20 Years of Inferring Inter-domain Routing Policies

Savvas Kastanakis, Vasileios Giotsas, Ioana Livadariu, Neeraj Suri

Lancaster University

CLOUDFLARE

simulamet
Simula Metropolitan Center for Digital Engineering AS

Questions to be addressed in this presentation

Q1. *Why would netops want to model the interdomain routing system?*

Q2. *Does the cornerstone model capture the actual interdomain routing policies?*

Q3. *What are the Selective Announced Prefixes and why should we care?*

Q4. *What can be done to enhance our inference capabilities?*

The Internet is a network of networks…

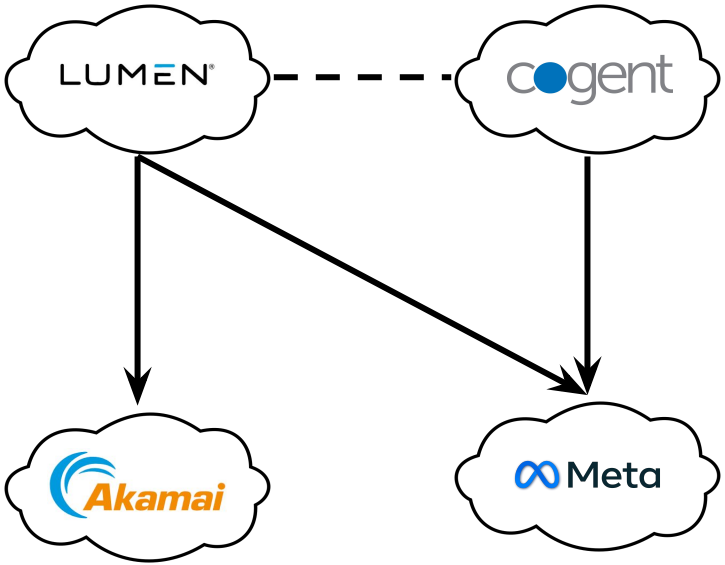

The Internet is a network of networks…

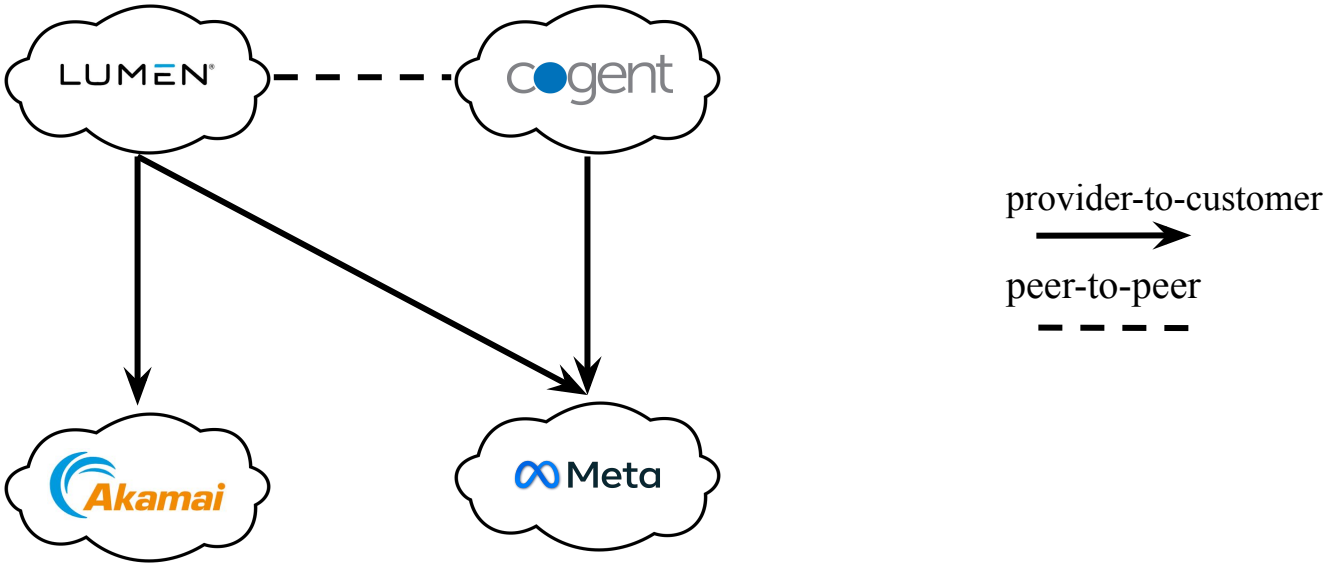

- ASes form **business relationships** to exchange traffic between their networks.

The Internet is a network of networks…

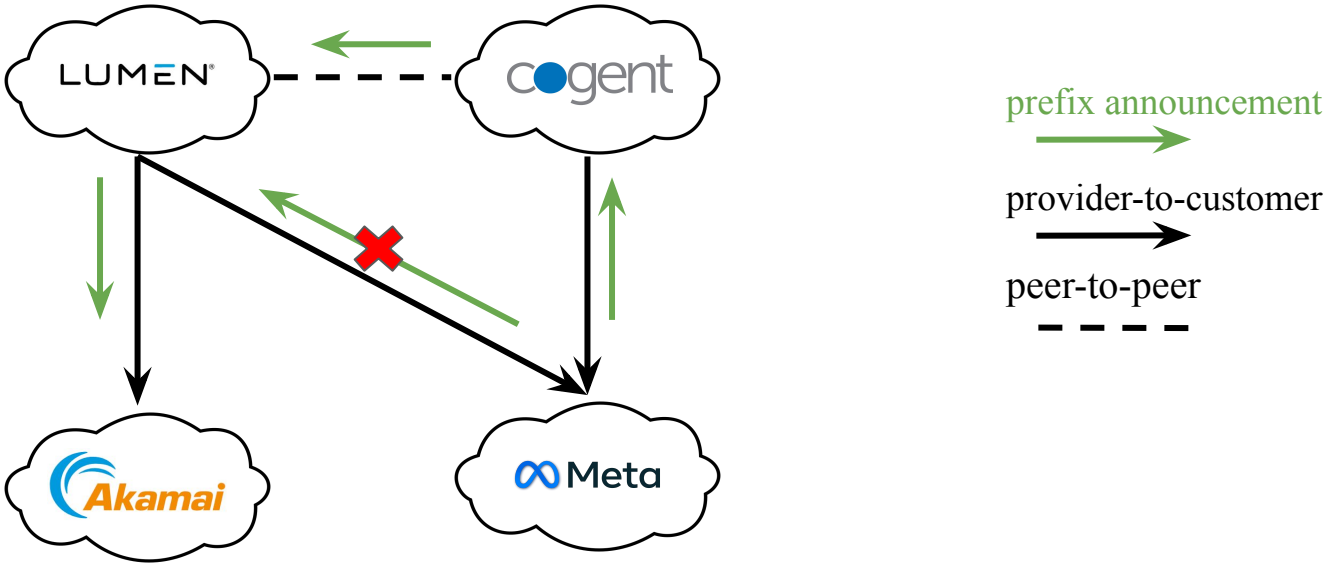

- ASes form **business relationships** to exchange traffic between their networks.
- ASes configure their **routing policies** to realize their business models.

The Internet is a network of networks…

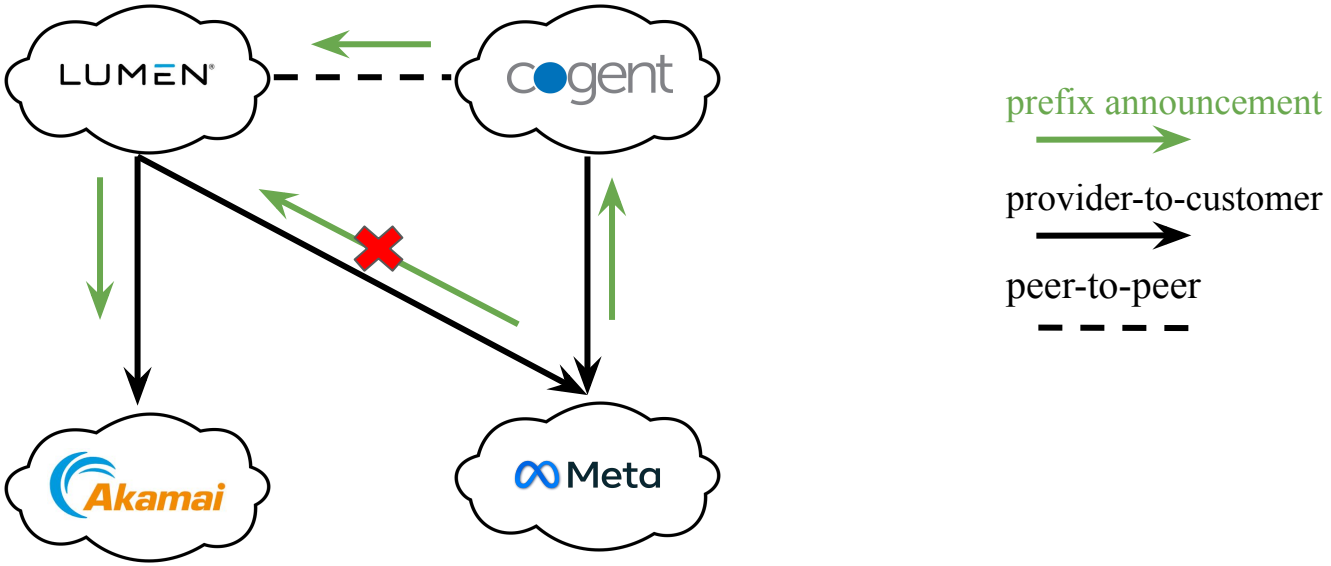

- ASes form **business relationships** to exchange traffic between their networks.
- ASes configure their **routing policies** to realize their business models.
- ASes are autonomous because they independently define their routing policies.

Why do netops need to model the interdomain routing system?

Effects of Traffic Engineering

Troubleshooting and Diagnosis

Network Planning and Design

Security Analysis

PoP Placement

Anycast Catchment

Effects of Routing Policies Changes

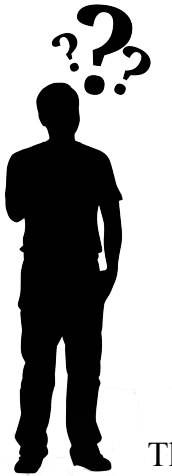

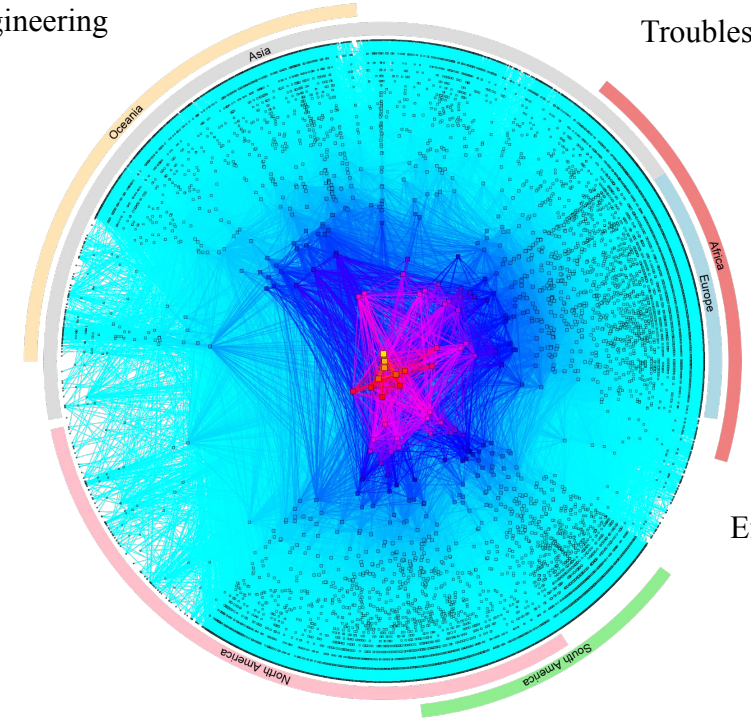

The CAIDA AS Core visualization depicts the Internet's ASes interconnections and locations
https://www.caida.org/projects/as-core/

From IMC’03 to IMC’23

From IMC'03 to IMC'23

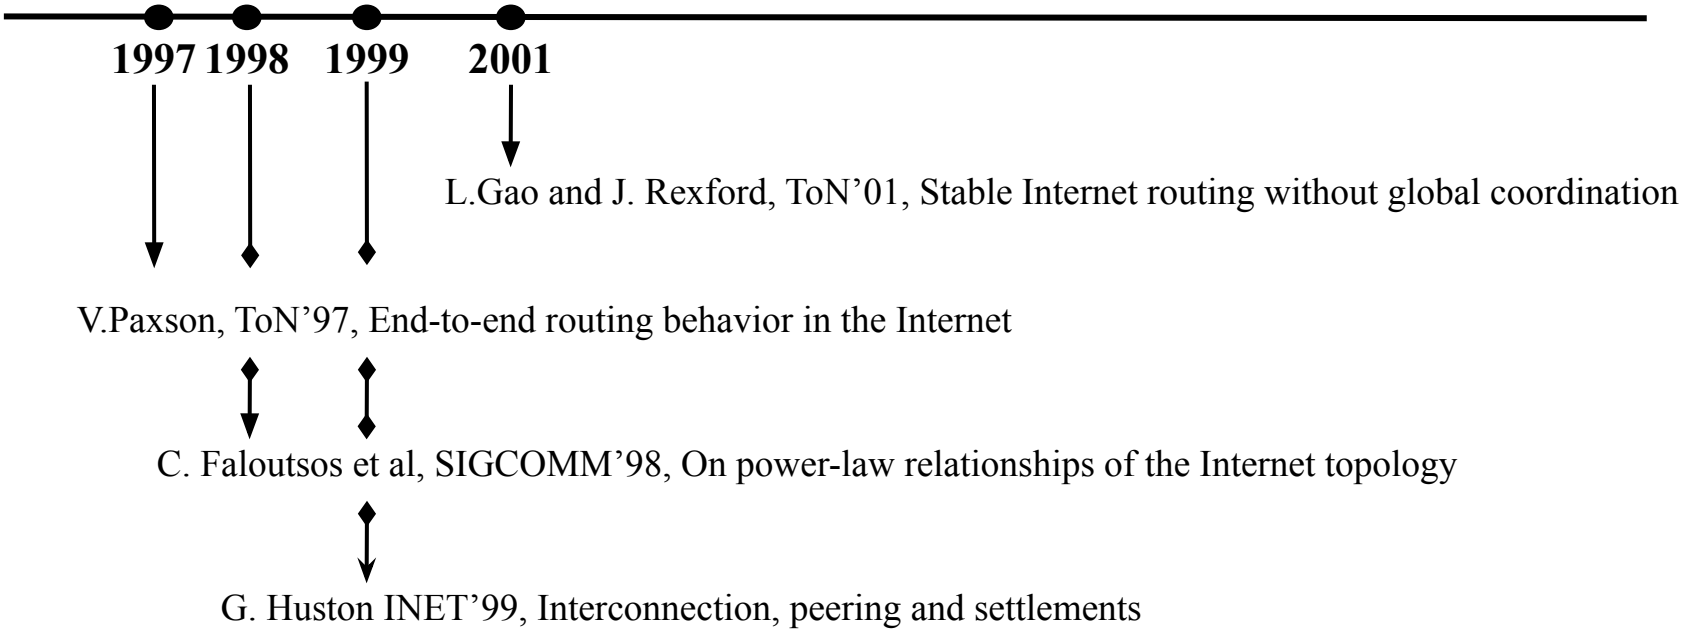

From IMC'03 to IMC'23

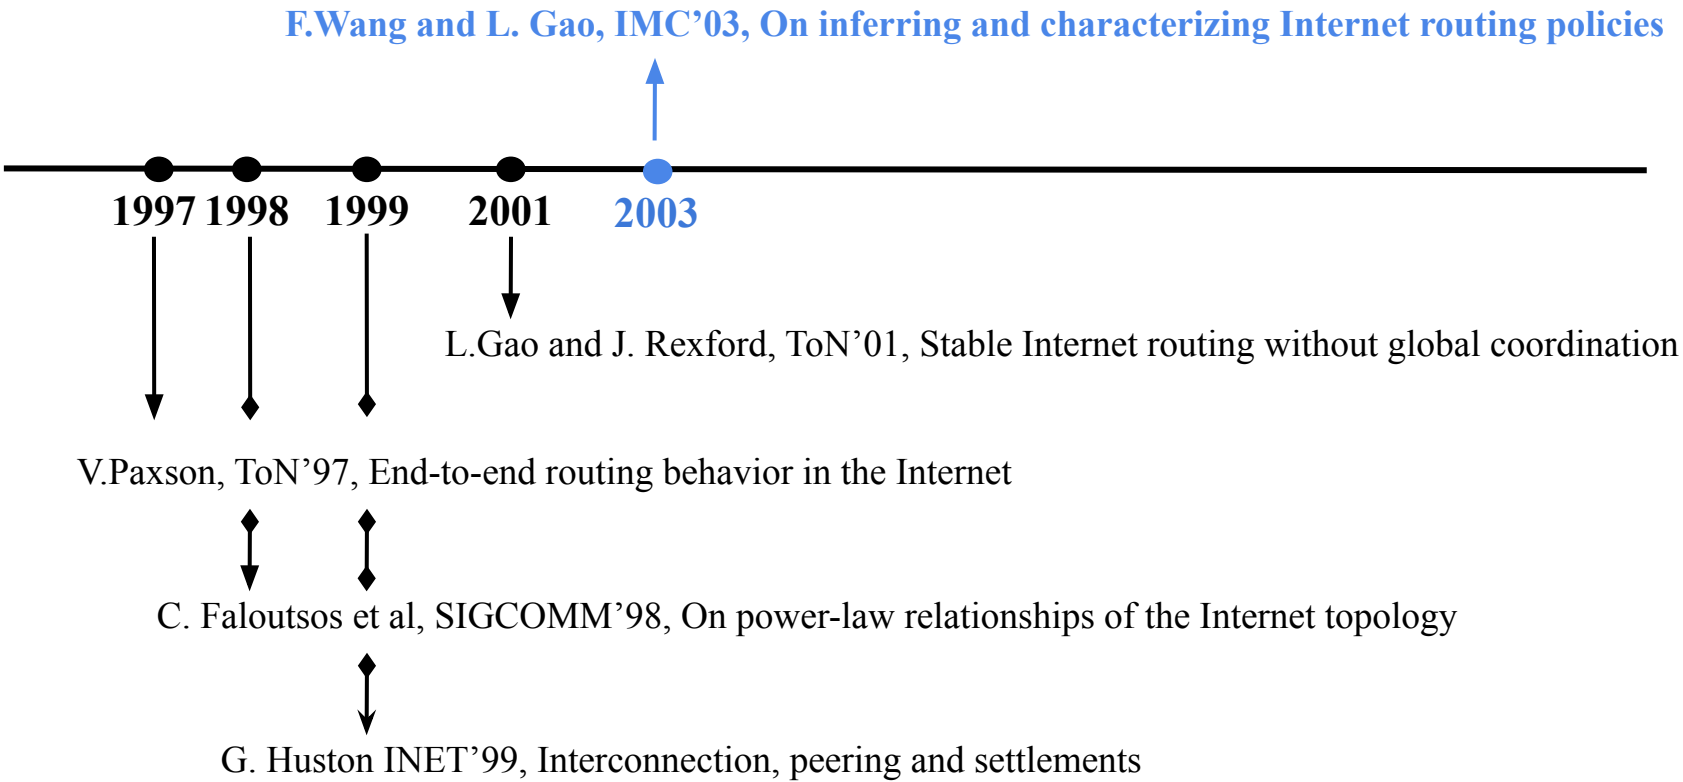

From IMC'03 to IMC'23

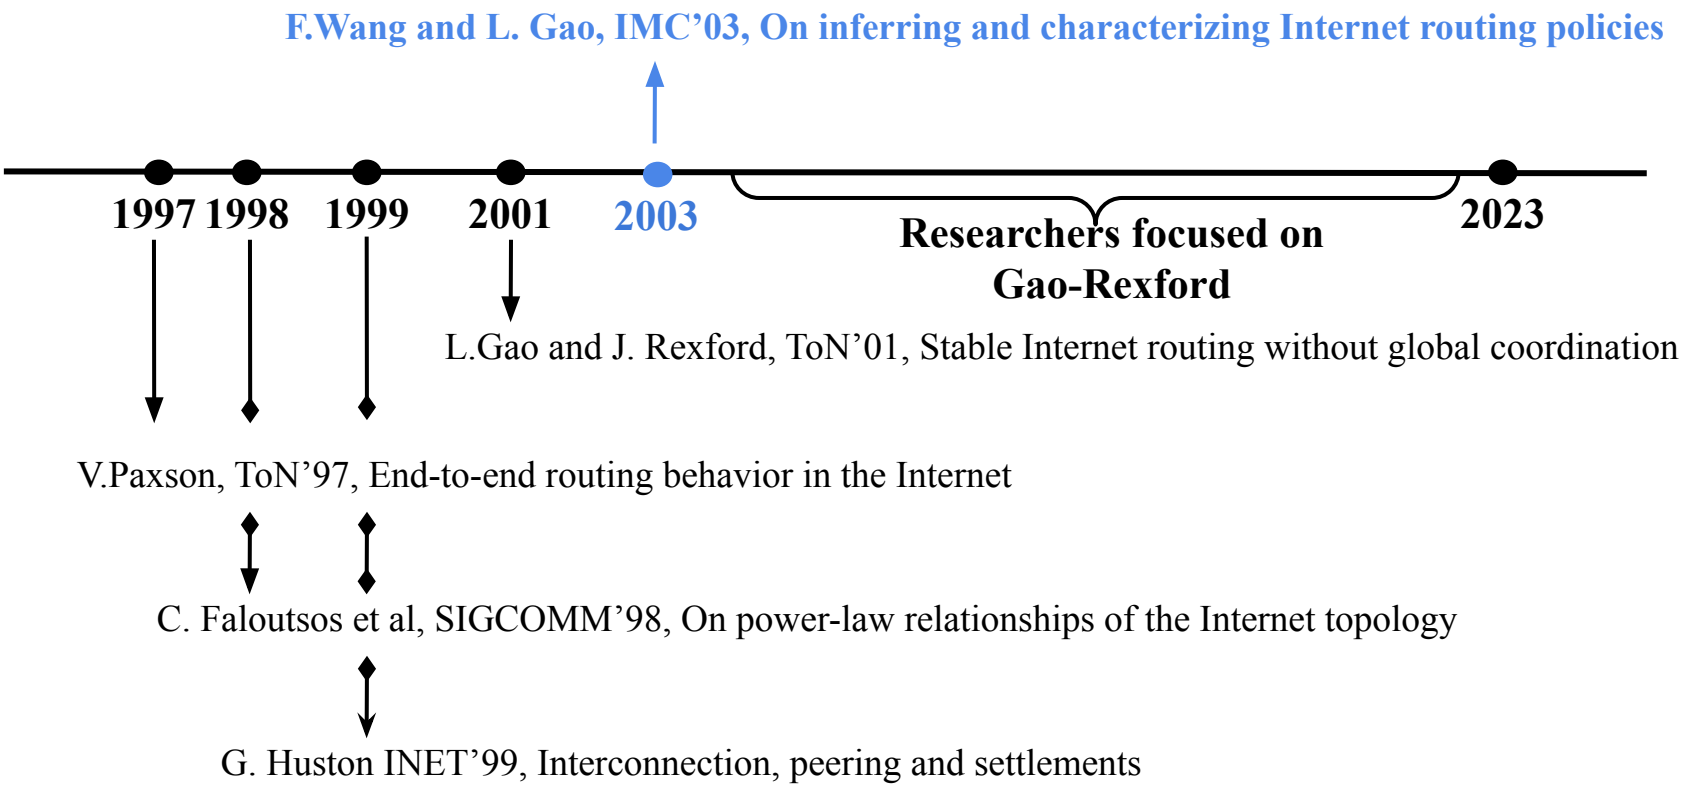

From IMC'03 to IMC'23

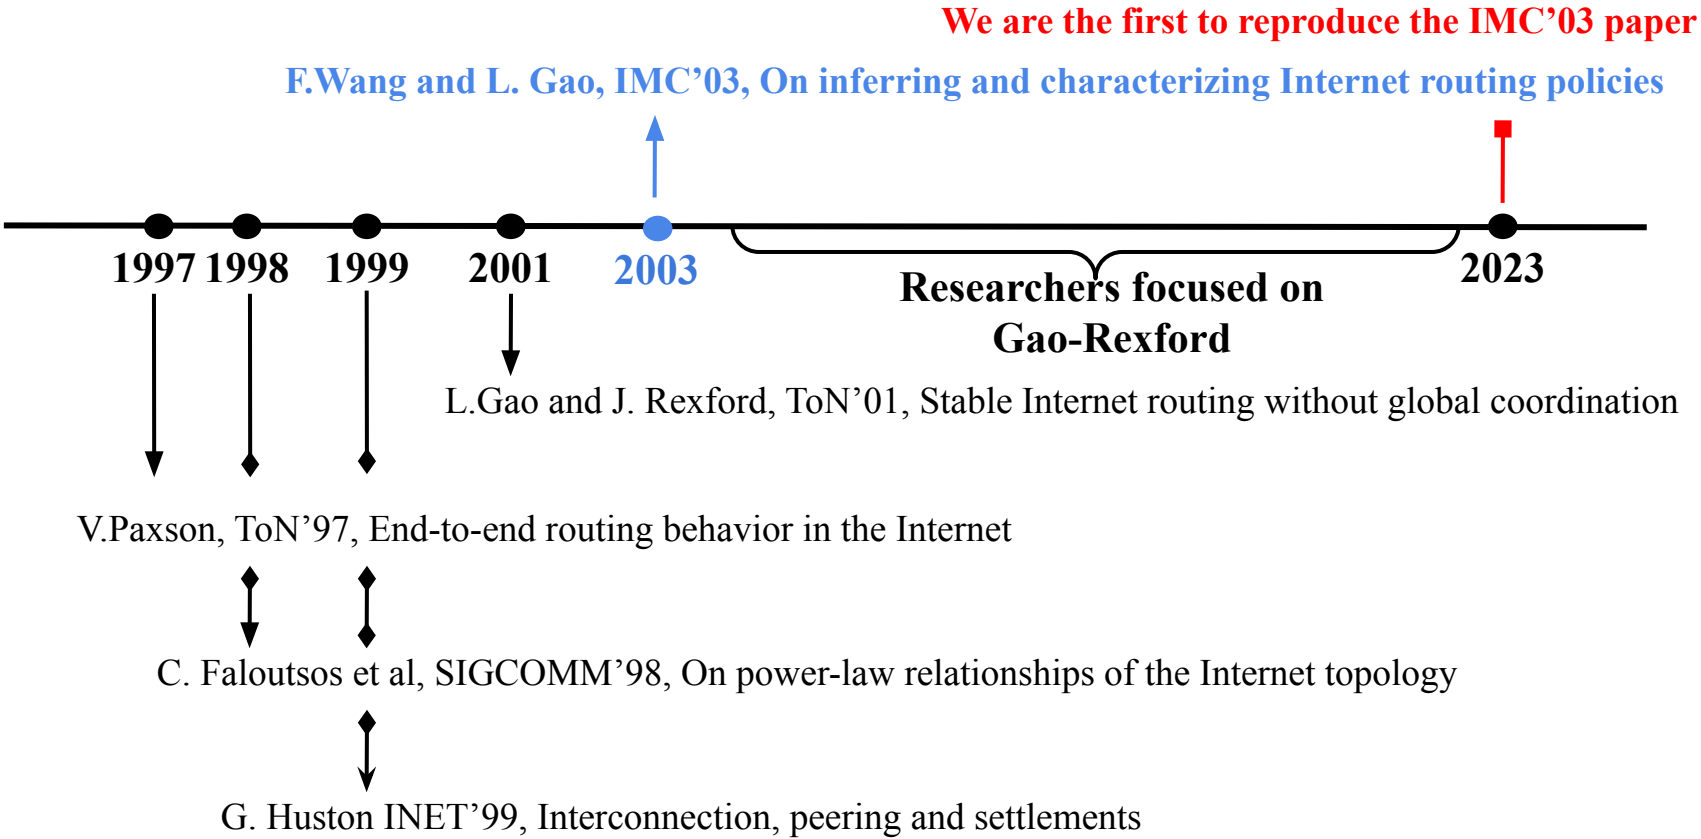

Replication overview: infer import and export routing policies

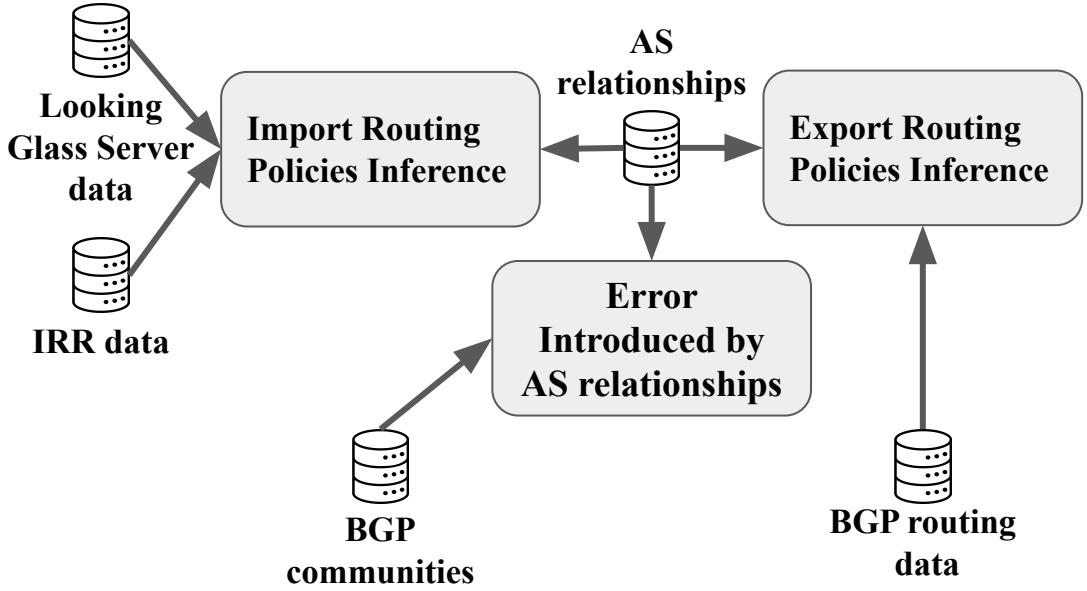

Inferring import policies

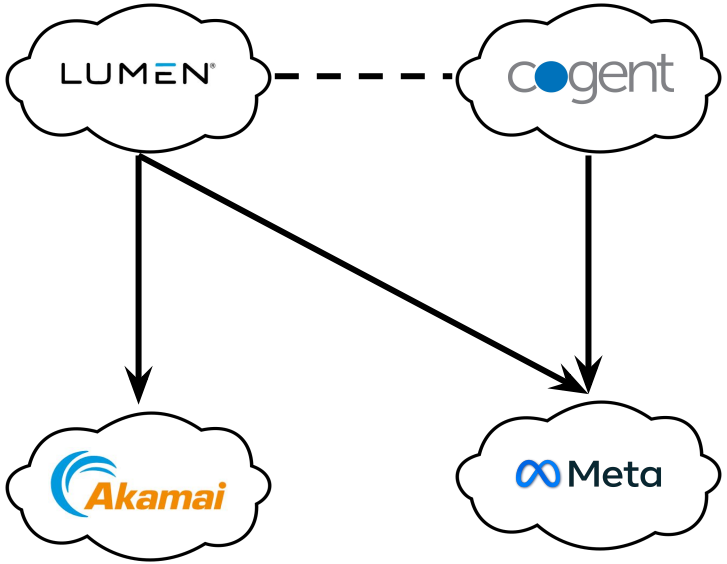

provider-to-customer

peer-to-peer

Inferring import policies

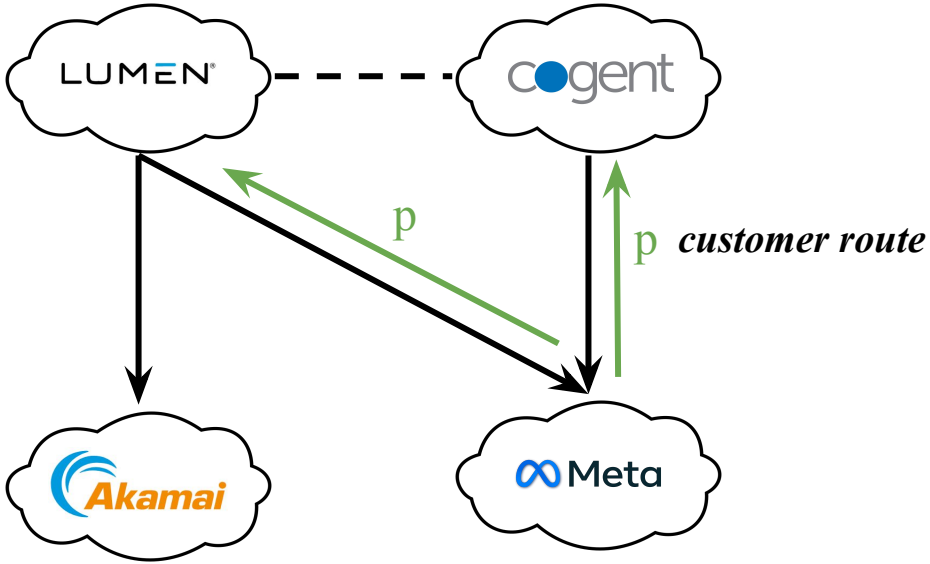

prefix announcement

provider-to-customer

peer-to-peer

Inferring import policies

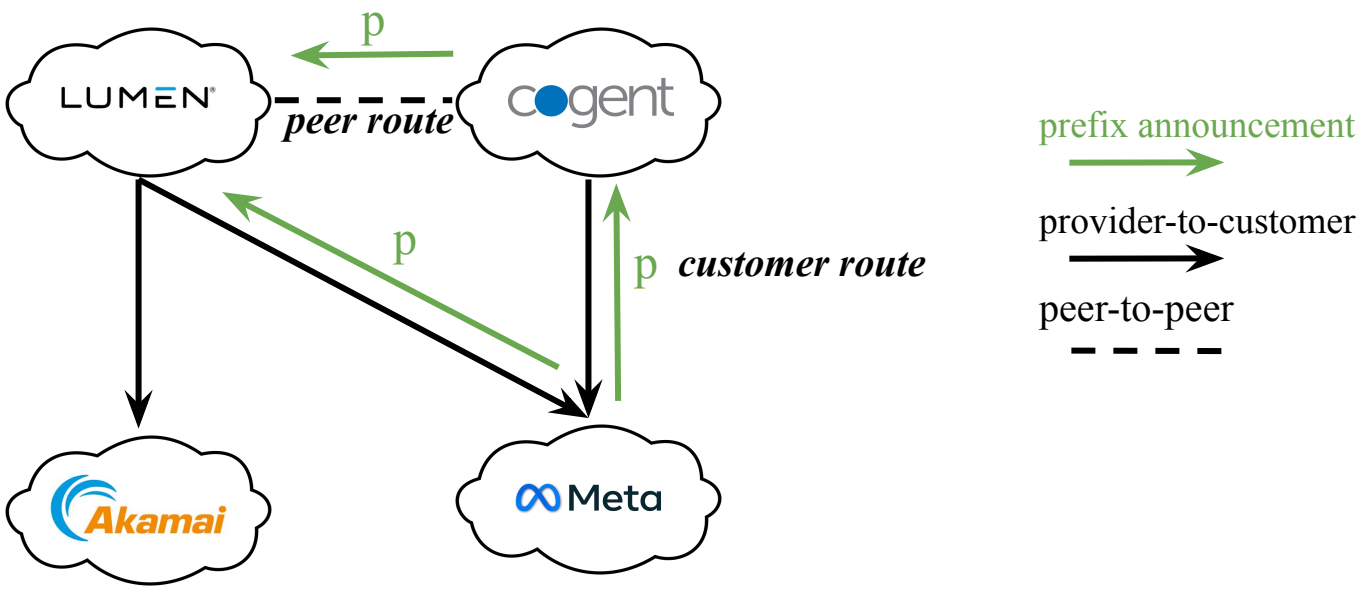

Inferring import policies

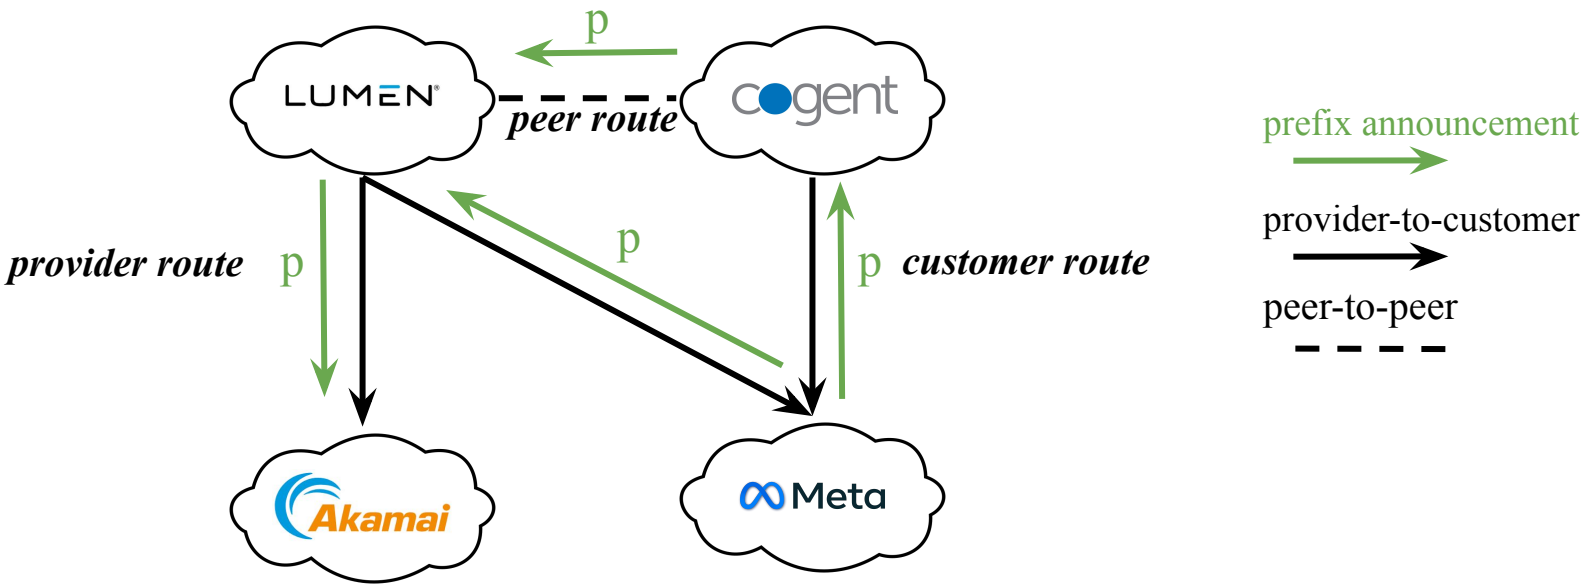

Inferring import policies

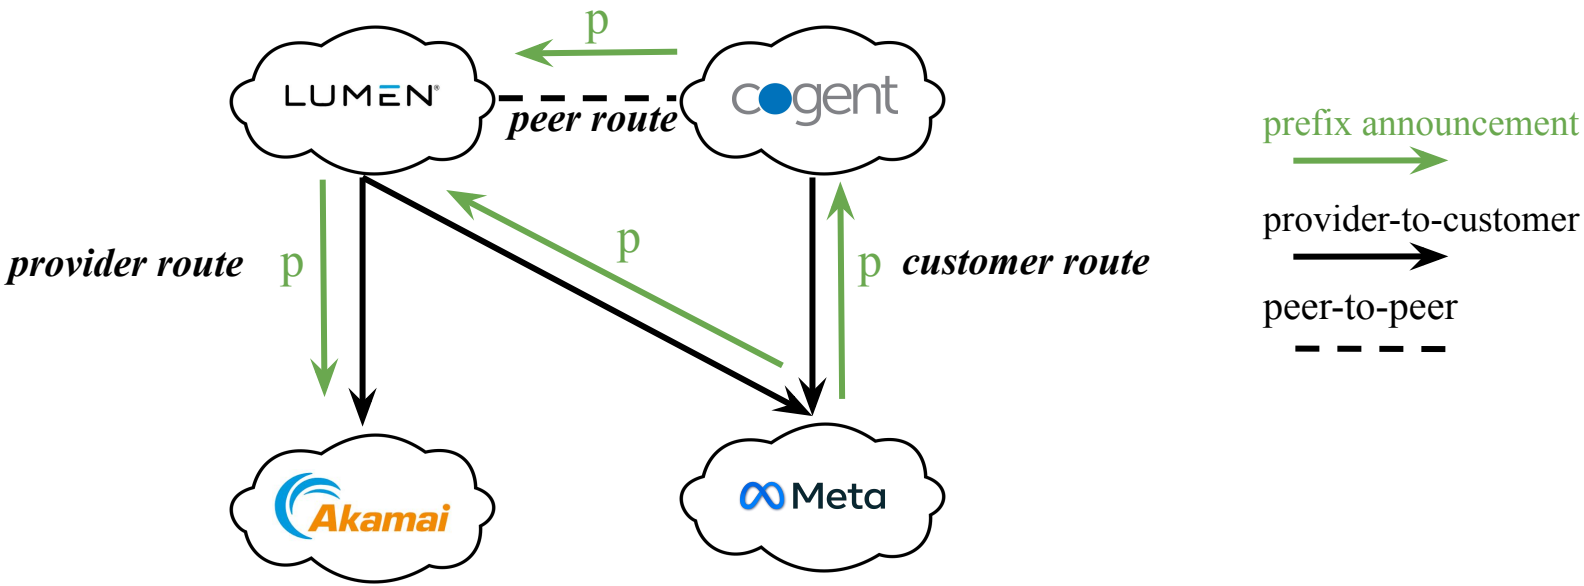

After receiving an announcement, an AS assigns a ***locpref*** value to indicate how favorable the route is:

- Customer > Peer > Provider → **Gao-Rexford model**

Inferring import policies

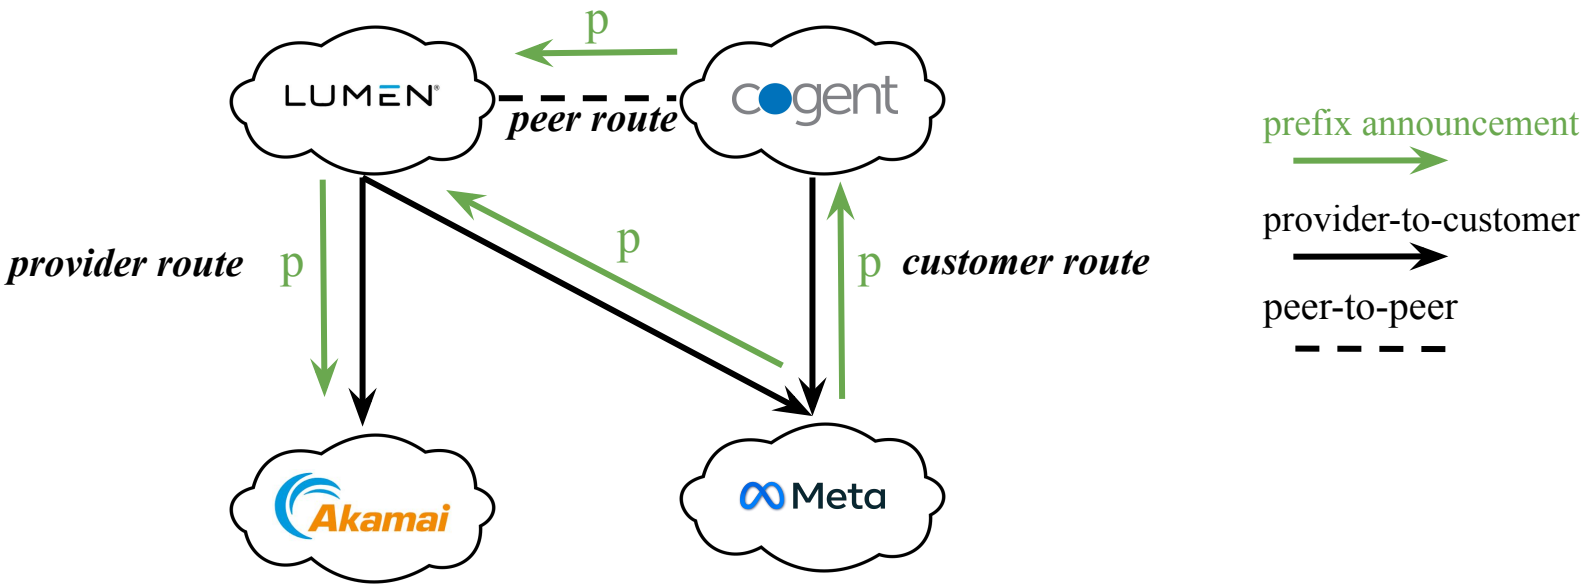

After receiving an announcement, an AS assigns a ***locpref*** value to indicate how favorable the route is:

- Customer > Peer > Provider → **Gao-Rexford model**

Does this assumption hold?

Replication overview: infer import policies

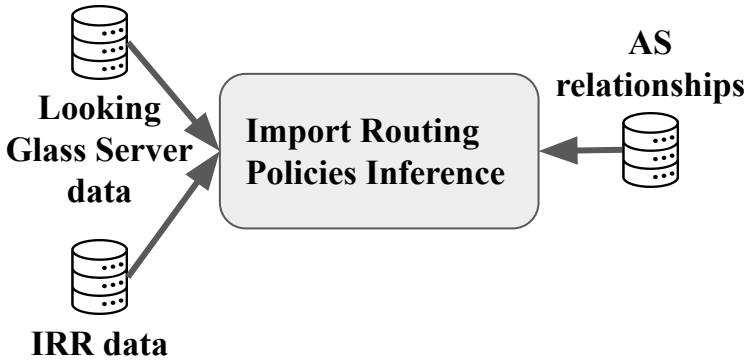

- We study public locpref settings via LG servers and IRR data.
- **We associate customer/peer/provider neighbors with their locpref values and compare them.**

Peer routes are highly inconsistent with Gao-Rexford

Peer routes are highly inconsistent with Gao-Rexford

- For all routes, ***locpref*** allocation is consistent with Gao-Rexford in 83% of the cases.

Peer routes are highly inconsistent with Gao-Rexford

- For all routes, ***locpref*** allocation is consistent with Gao-Rexford in 83% of the cases.
  - In the Wang-Gao study, the average consistency was above 99%.

Peer routes are highly inconsistent with Gao-Rexford

- For all routes, ***locpref*** allocation is consistent with Gao-Rexford in 83% of the cases.
  - In the Wang-Gao study, the average consistency was above 99%.
  - Customer/provider routes are consistent (>90%); **peer routes have 59% consistency.**

Peer routes are highly inconsistent with Gao-Rexford

- For all routes, ***locpref*** allocation is consistent with Gao-Rexford in 83% of the cases.
  - In the Wang-Gao study, the average consistency was above 99%.
  - Customer/provider routes are consistent (>90%); **peer routes have 59% consistency.**
  - IRR locpref is more consistent with AS relationships than LG locpref.

Peer routes are highly inconsistent with Gao-Rexford

- For all routes, ***locpref*** allocation is consistent with Gao-Rexford in 83% of the cases.
  - In the Wang-Gao study, the average consistency was above 99%.
  - Customer/provider routes are consistent (>90%); **peer routes have 59% consistency.**
  - IRR locpref is more consistent with AS relationships than LG locpref.

Locpref settings are not as heavily dependent to AS-relationships as they used to be, hence, we need to reconsider the Gao-Rexford model.

Error introduced by AS relationships is negligible

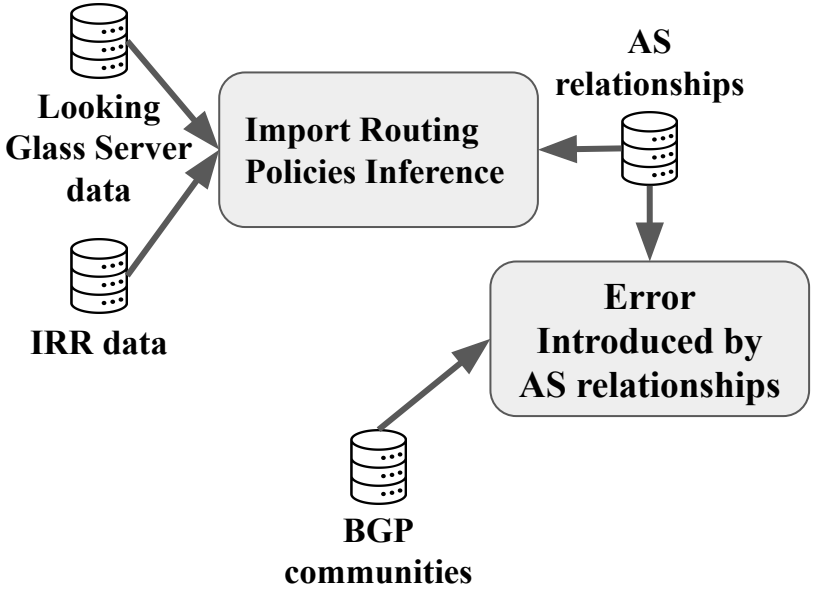

- Discrepancies in our inference can happen because:
  a. AS-relationships are inaccurate
  b. real-world local preference allocations are more complex than the Gao-Rexford model

We focus on inferring export policies to providers/peers

- As in the Wang-Gao work, we assume that a provider **has to announce** all of its prefixes to their customers based on their business agreements.

- However a customer/peer can **selectively announce** its prefixes to their provider/peer to load balance traffic/costs.

- We observe the export policies of an AS from the routing tables of its neighbors.

Inferring export policies to provider

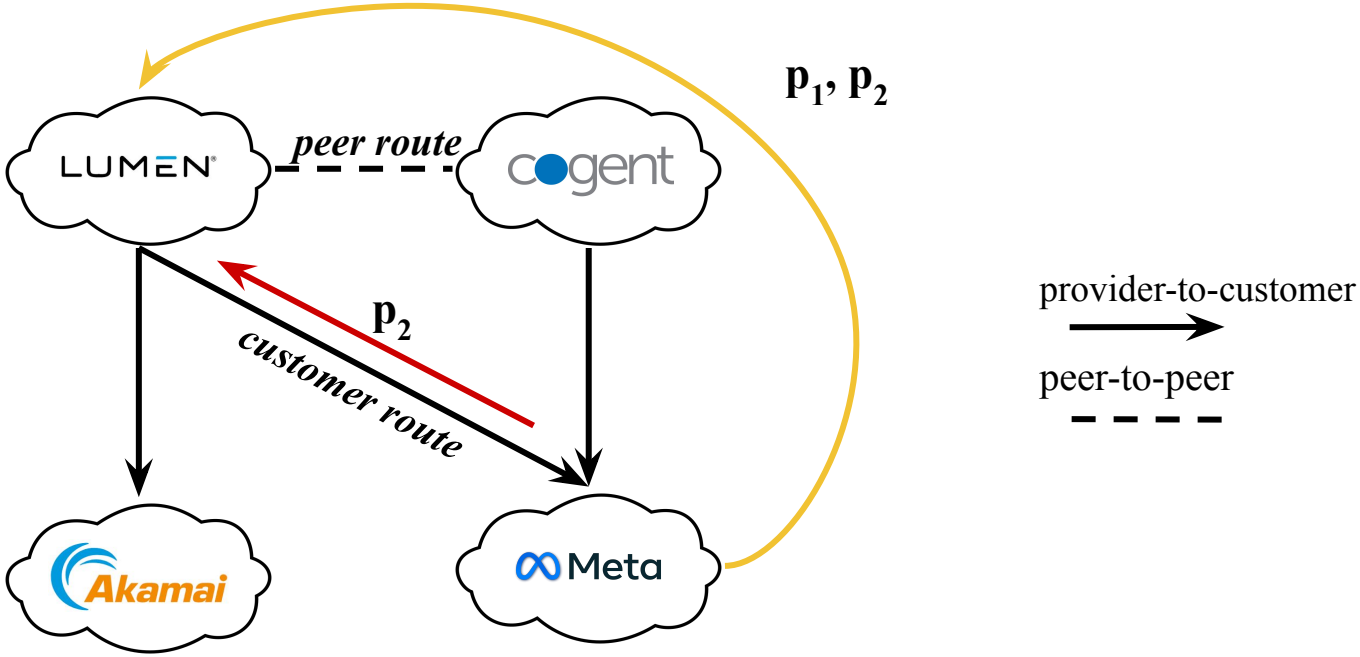

- Multihomed customers want to load balance traffic/costs, so they selectively announce prefixes.
- **Selective Announced Prefix**: a customer prefix received through a peer/provider route.

Replication overview: infer export policies

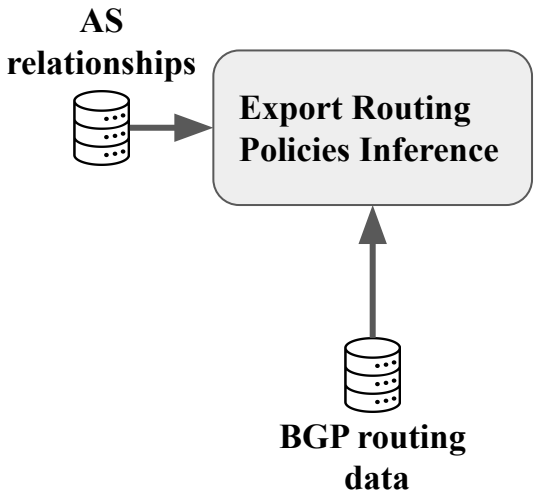

Selective announced prefixes are **still** prevalent

- We studied 21 snapshots of BGP routing tables for April 1st, 2023.
- The direct way to infer a customer's export policies is to use the BGP tables from its neighbors.
- **The top-5 ASes are high centrality networks through which we observe a large portion of SA prefixes.**
- Similar results as in the Wang-Gao study.

AS number	% of SA prefixes	% of SA origins
3303	69	45
3257	54	55
6939	44	44
3549	32	26
7018	28	17
37100	27	18
37271	21	12
3741	19	12
31027	13	09
852	13	09

Selective origins announce all their prefixes selectively

- In our study, more than 75% of SA origins advertised exactly 100% of their prefixes selectively.

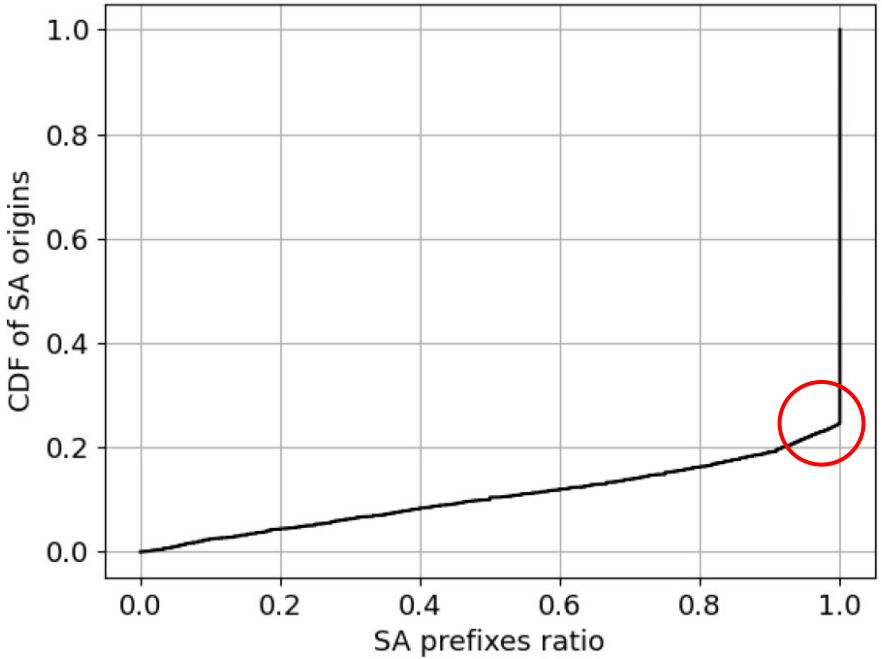

Selective announced prefixes are persistent

- We gathered snapshots of BGP routing tables from 2022-01-01 to 2023-01-31 for AS7018.
- **SA prefixes are consistent during a day/month, but become unstable during a year.**

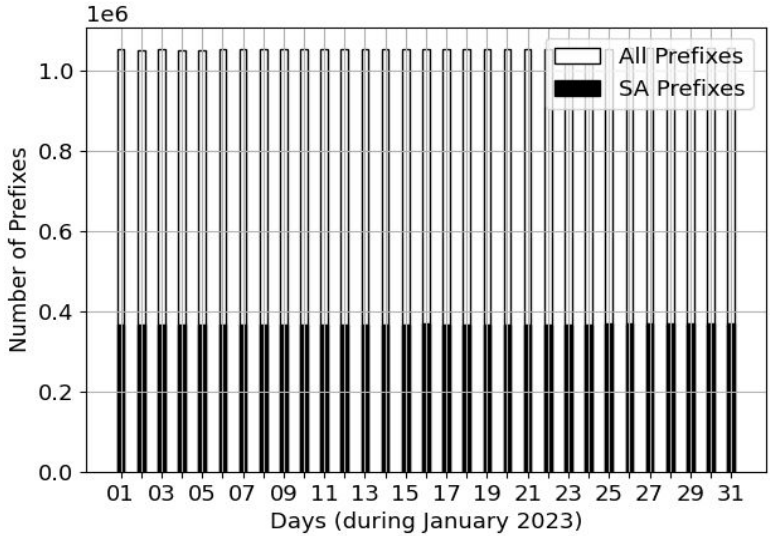

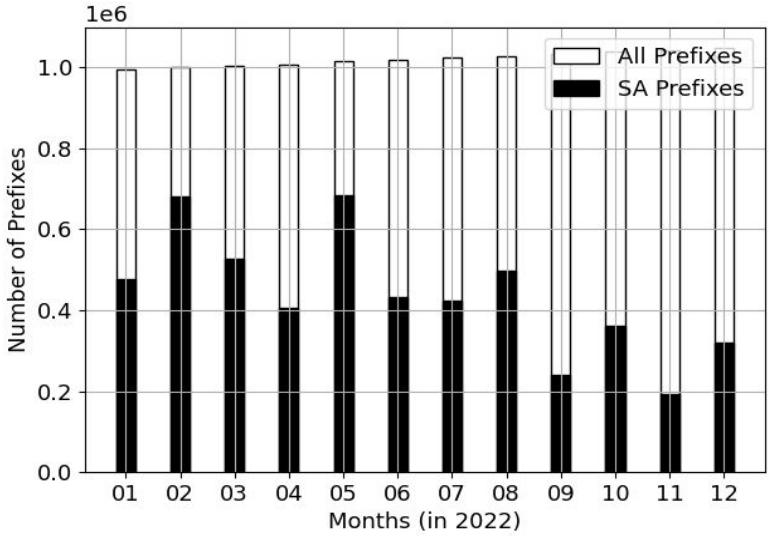

Persistence of SA prefixes over the last 20 years

- We gathered snapshots of BGP routing tables for the 1st of April of [2003, 2023].
- **SA prefixes are sensitive to topological and policy changes**.

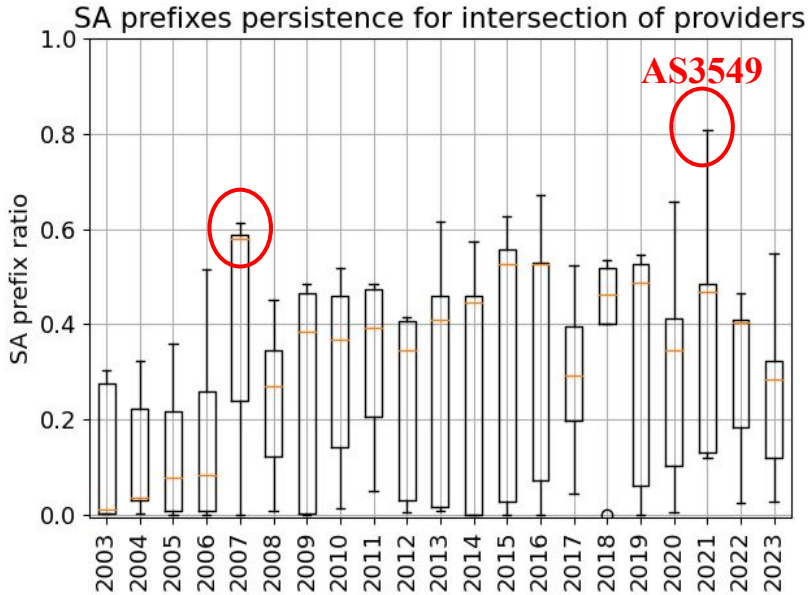

Take-away Message

Contact: s.kastanakis@lancaster.ac.uk
Github: https://github.com/kastanakis/replicating-selective-announcements-inference
Paper: https://dl.acm.org/doi/abs/10.1145/3618257.3624799

Take-away Message

A. Understanding and predicting the interdomain routing system is crucial.

Contact: s.kastanakis@lancaster.ac.uk
Github: https://github.com/kastanakis/replicating-selective-announcements-inference
Paper: https://dl.acm.org/doi/abs/10.1145/3618257.3624799

Take-away Message

A. Understanding and predicting the interdomain routing system is crucial.

B. Reachability is determined by connectivity AND routing policies; we have the tools to observe it and measure it but we don't use them!

Contact: s.kastanakis@lancaster.ac.uk
Github: https://github.com/kastanakis/replicating-selective-announcements-inference
Paper: https://dl.acm.org/doi/abs/10.1145/3618257.3624799

Take-away Message

A. Understanding and predicting the interdomain routing system is crucial.

B. Reachability is determined by connectivity AND routing policies; we have the tools to observe it and measure it but we don't use them!

C. Reconsider the Gao-Rexford model and extent it with inferential techniques on complex policies.

Contact: s.kastanakis@lancaster.ac.uk
Github: https://github.com/kastanakis/replicating-selective-announcements-inference
Paper: https://dl.acm.org/doi/abs/10.1145/3618257.3624799

Backup Slides

In simple words…

Collect BGP RIBS and extract prefixes and AS-paths

Determine if the origin AS is direct/indirect customer of the VP AS

Determine if the next-hop AS is a customer/peer/provider of the VP AS

If the first-hop relationship is *'more expensive'* than the relationship with the origin, then selective announced prefix

Replicating the IMC’03 paper was challenging…

Replicating the IMC'03 paper was challenging…

1. The ASes that provide LG servers changed; we used all available LGs.
   a. # of LG servers are too few to allow for Internet wide generalizations.
   b. LG servers allow limited querying and provide no historical data.

Replicating the IMC'03 paper was challenging…

1. The ASes that provide LG servers changed; we used all available LGs.
   a. # of LG servers are too few to allow for Internet wide generalizations.
   b. LG servers allow limited querying and provide no historical data.

2. The AS topology incompleteness problem can cause us to underestimate the number of selectively advertised prefixes (more results in the paper).

Example of BGP routing policies

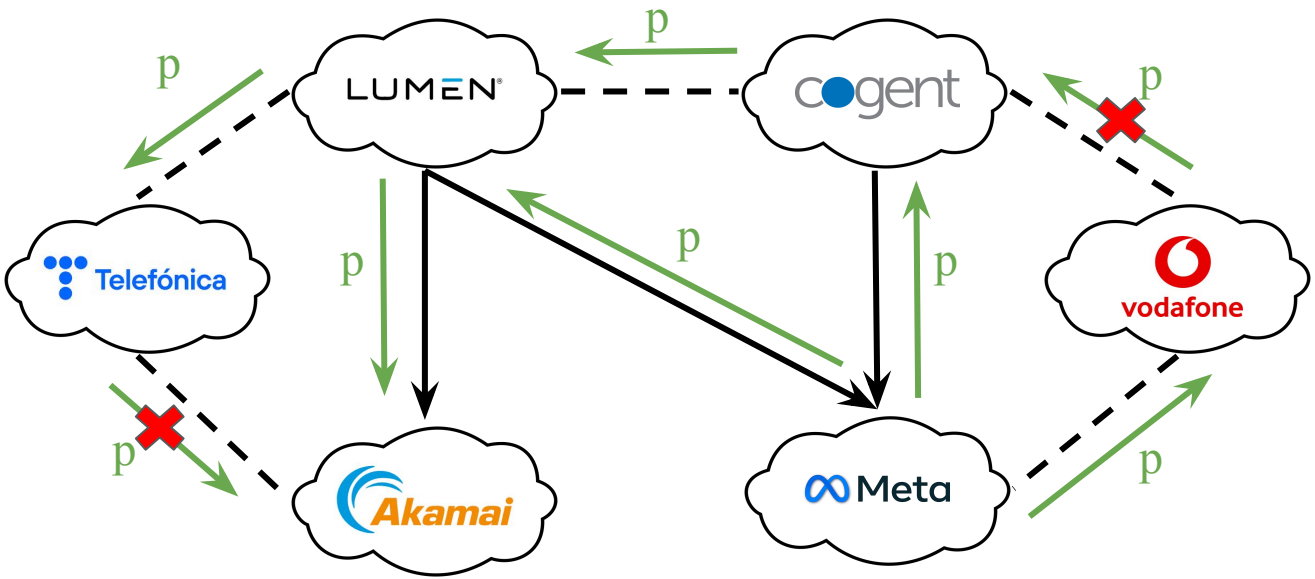

- BGP policies help an AS determine which routes are accepted from its neighbors (import policies) and which routes to be advertised to its neighbors (export policies)

IRR locpref policies are consistent with AS-relationships

- This is due to the difference between actual control-plane configurations and abstract policies described for documentation purposes.

ASN	% of typical locpref
1887	100%
2118	100%
5408	89%
6730	100%
6799	100%
8280	100%
8342	100%
8343	100%
8369	92%
8371	100%
9032	96%
12695	100%
12713	100%
15290	100%
15544	100%
16559	100%

ASN	% of typical locpref
20845	100%
20850	100%
21483	83%
24739	100%
35566	93%
39775	100%
43893	100%
44946	100%
47764	92%
49673	100%
50639	100%
52075	100%
60476	100%
199081	100%
199860	100%
396298	100%

More than 80% of SA prefixes are validated

- Find at least one AS path between AS_{vp} and AS_{origin} that traverses only p2c links for another prefix.

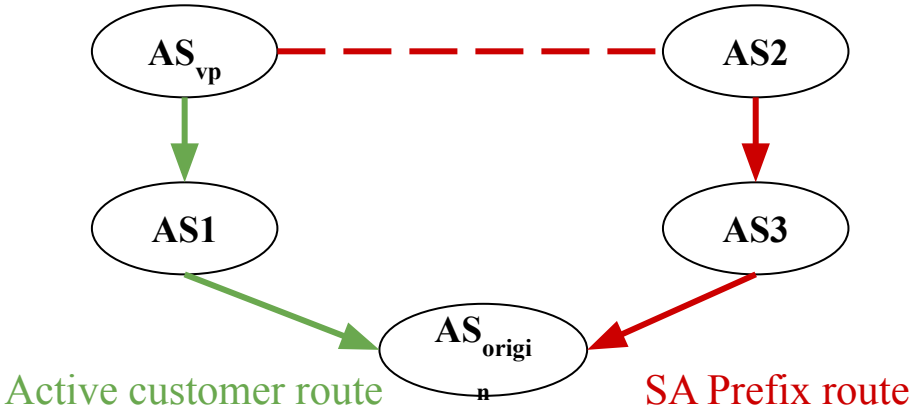

Uptime of selective announced prefixes for AS7018

- SA Prefixes Uptime: the times an SA Prefix appears.
- More than 90% of SA Prefixes are stable for the whole January, 2023, but when studying in an one year measurement window (2022), SA prefixes are active between 1 to 12 months.

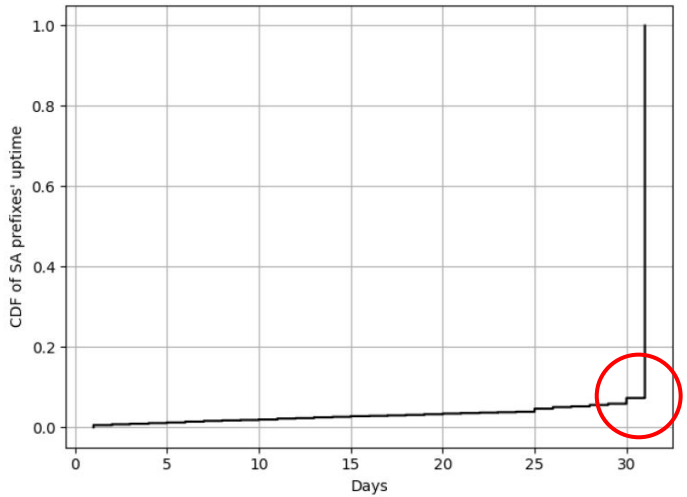

(a) CDF of SA prefixes uptime in Jan. 2023.

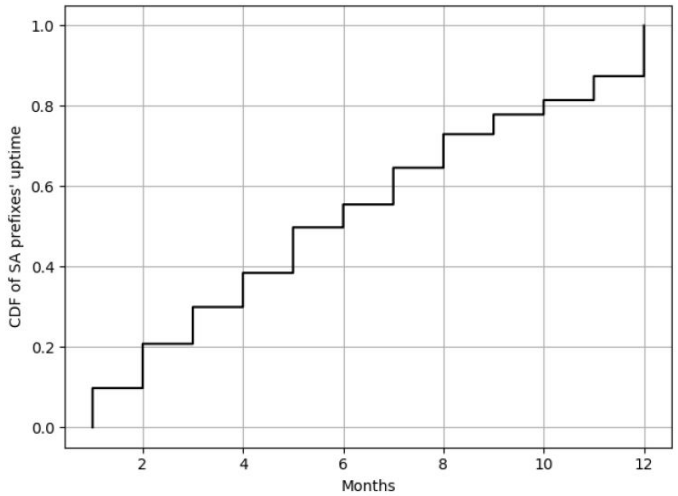

(b) CDF of SA prefixes uptime in 2022.

Possible Causes of SA Prefixes

1) Prefix Splitting

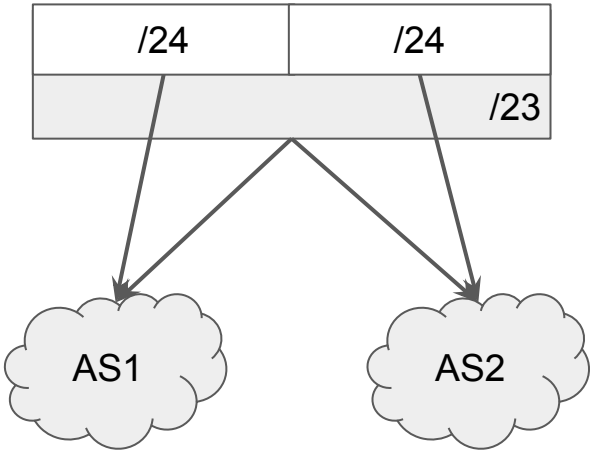

2) Prefix Aggregating

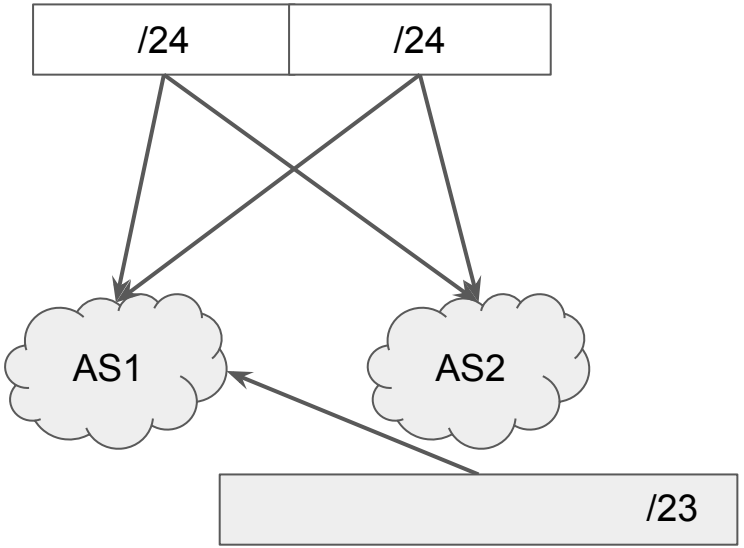

Possible Causes of SA Prefixes

- Prefix splitting and prefix aggregating are **not** the major causes of selective announcement.

AS number	% of prefix splitting	% of prefix aggregation
3257	1.3	0.03
3292	0.02	0
3549	0.4	0.1
5511	1.6	0
7018	0.3	0.3

Selective announcements is not a phenomenon prevalent among peers

- We test how many ***peer prefixes*** are announced through ***provider routes***.

AS number	% of SA prefixes	% of SA origins
AS5413	9.1	0.4
AS3741	8.2	1.1
AS3303	7.6	1
AS19653	6.6	0.8
AS37100	4.6	0.8
AS553	3.7	0.5
AS6667	2.8	0.2
AS852	2.7	0.2
AS6939	2.1	1.5
AS3292	1.3	0.1
AS37271	1.1	0.1
AS1280	1.1	0.06
AS31027	0.5	0.07