Inferring Multilateral Peering

Vasileios Giotsas[1]
Shi Zhou[1]
Matthew Luckie[2]
Kc Claffy[2]

[1]University College London
[2]UCSD/CAIDA

ACM CoNEXT 2013

Motivation and goals

- Topology sources capture only a small fraction of Autonomous System (AS) p2p links
  - **41%** p2p links missing from public BGP data (Chen 2009)
  - At least **50K** IXP links (Augustin 2009)
  - **50K** peering links in a single IXP (Ager 2012)
  - **142K** peering links (PCH survey 2011)
  - This work: **206K** peering links, **88%** missing from public BGP data

Motivation and goals

- Topology sources capture only a small fraction of Autonomous System (AS) p2p links
  - **41%** p2p links missing from public BGP data (Chen 2009)
  - At least **50K** IXP links (Augustin 2009)
  - **50K** peering links in a single IXP (Ager 2012)
  - **142K** peering links (PCH survey 2011)
  - This work: **206K** peering links, **88%** missing from public BGP data
- Goals:
  - Collect and make publicly available data
  - Low measurement cost → repeatability

IXPs facilitate peering

- 95% of missing peering links in IXPs (He 2005)
- IXP: A physical infrastructure to facilitate direct traffic exchange
- Two operational models:
  - Non-profit (European) → open data sharing
  - Commercial (N. American) → restrictive data sharing

IXPs facilitate peering

- 95% of missing peering links in IXPs (He 2005)
- IXP: A physical infrastructure to facilitate direct traffic exchange
- Two operational models:
  - Non-profit (European) → open data sharing
  - Commercial (N. American) → restrictive data sharing

Two peering paradigms

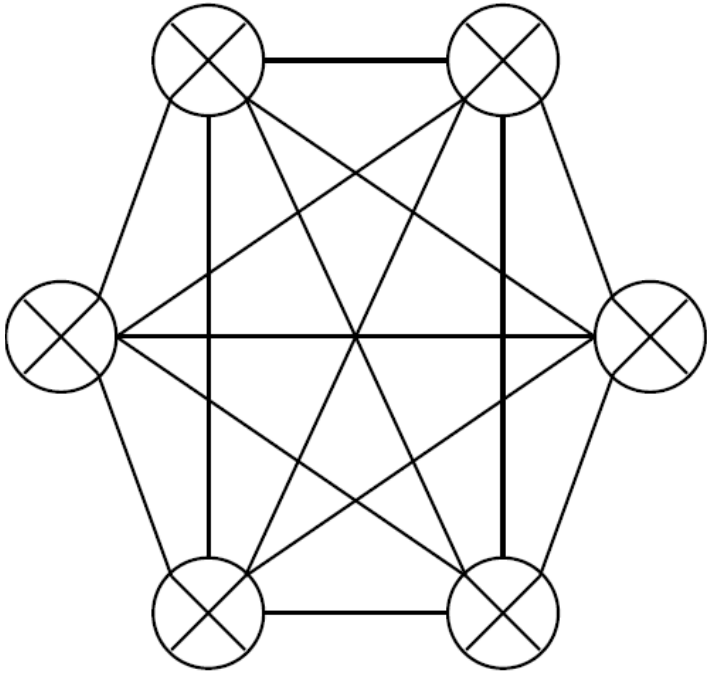

- Bilateral peering
  - Separate BGP session per peering
  - Tight control of peering
  - Poor scalability

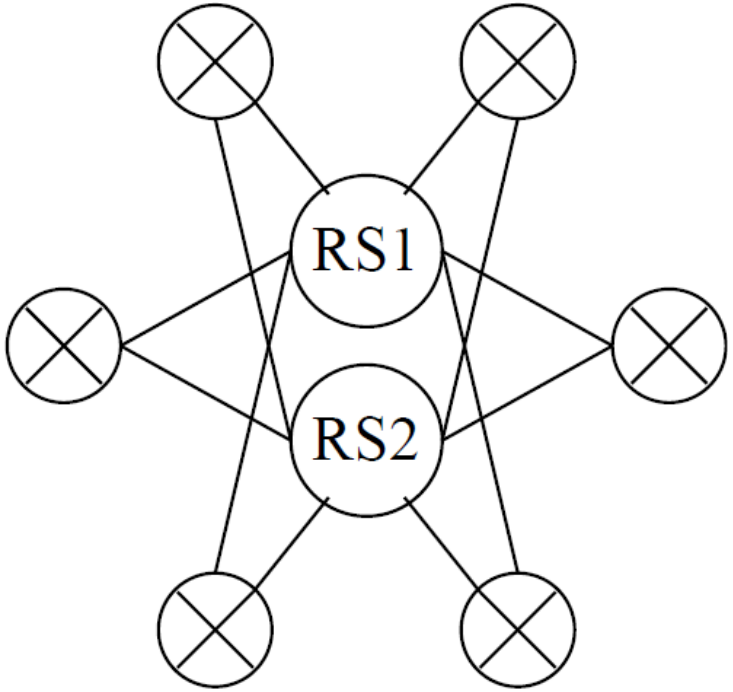

- Multilateral peering (**MLP**)
  - BGP session only with Route Servers (RS) for all links
  - Loose control of peering
  - Great scalability/flexibility

Route Servers enable dense peering

- Abundance of peering links because of multilateral agreements
  - PCH Survey (2011), Chatzis et al. (2013), Willinger et al. (2013)
- Very limited public data
  - Links
  - Policies

Route Servers enable dense peering

- Abundance of peering links because of multilateral agreements
  - PCH Survey (2011), Chatzis et al. (2013), Willinger et al. (2013)
- Very limited public data
  - Links
  - Policies

- Goal: Infer MLP links that we cannot observe

How MLP links are established?

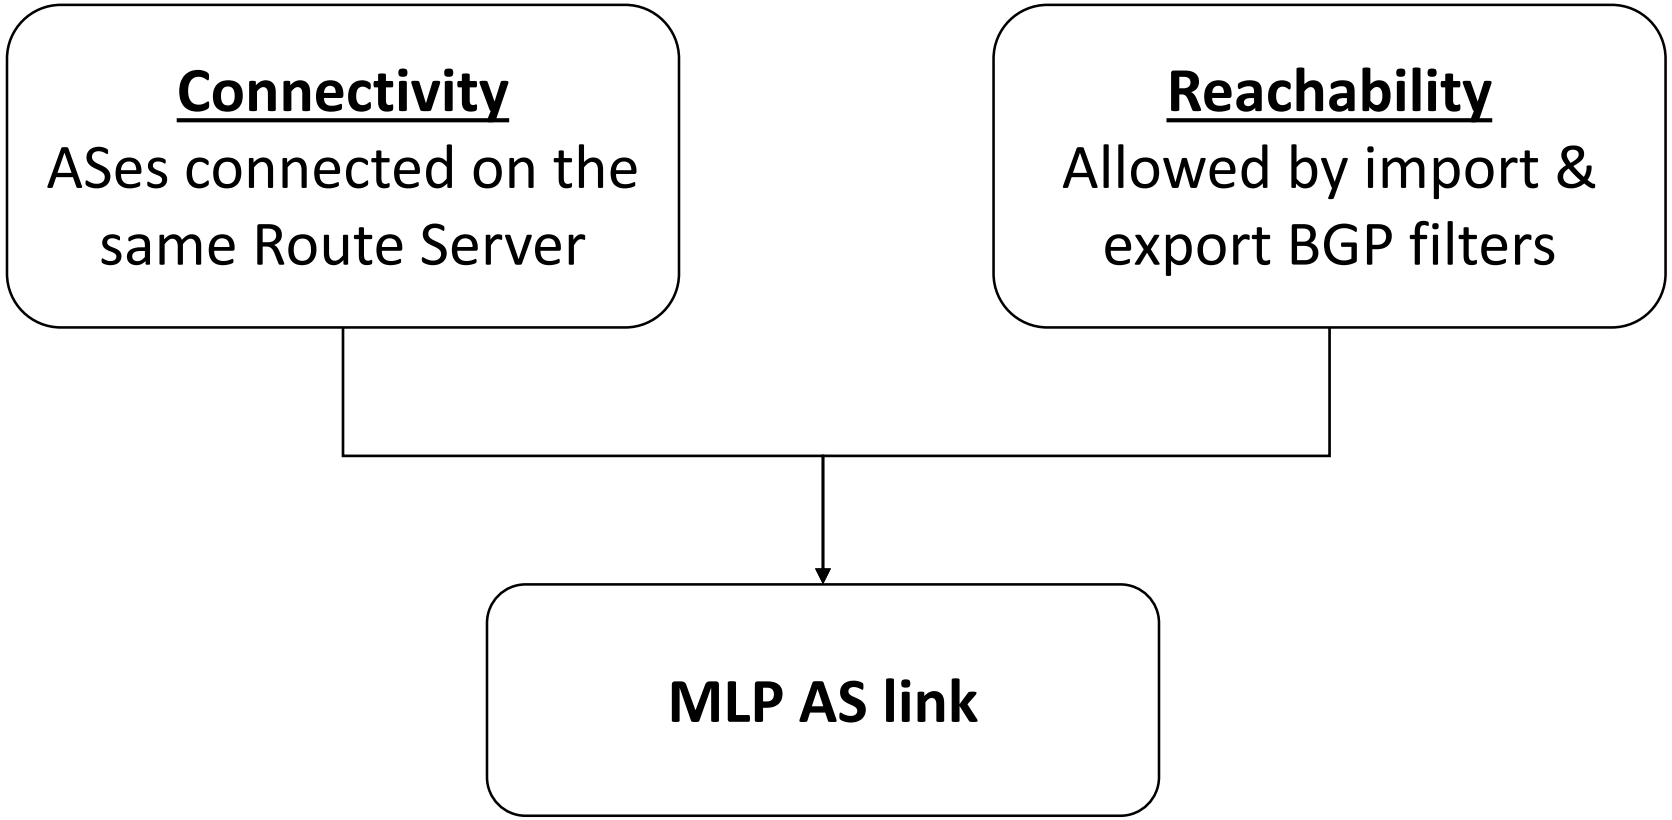

How MLP links are established?

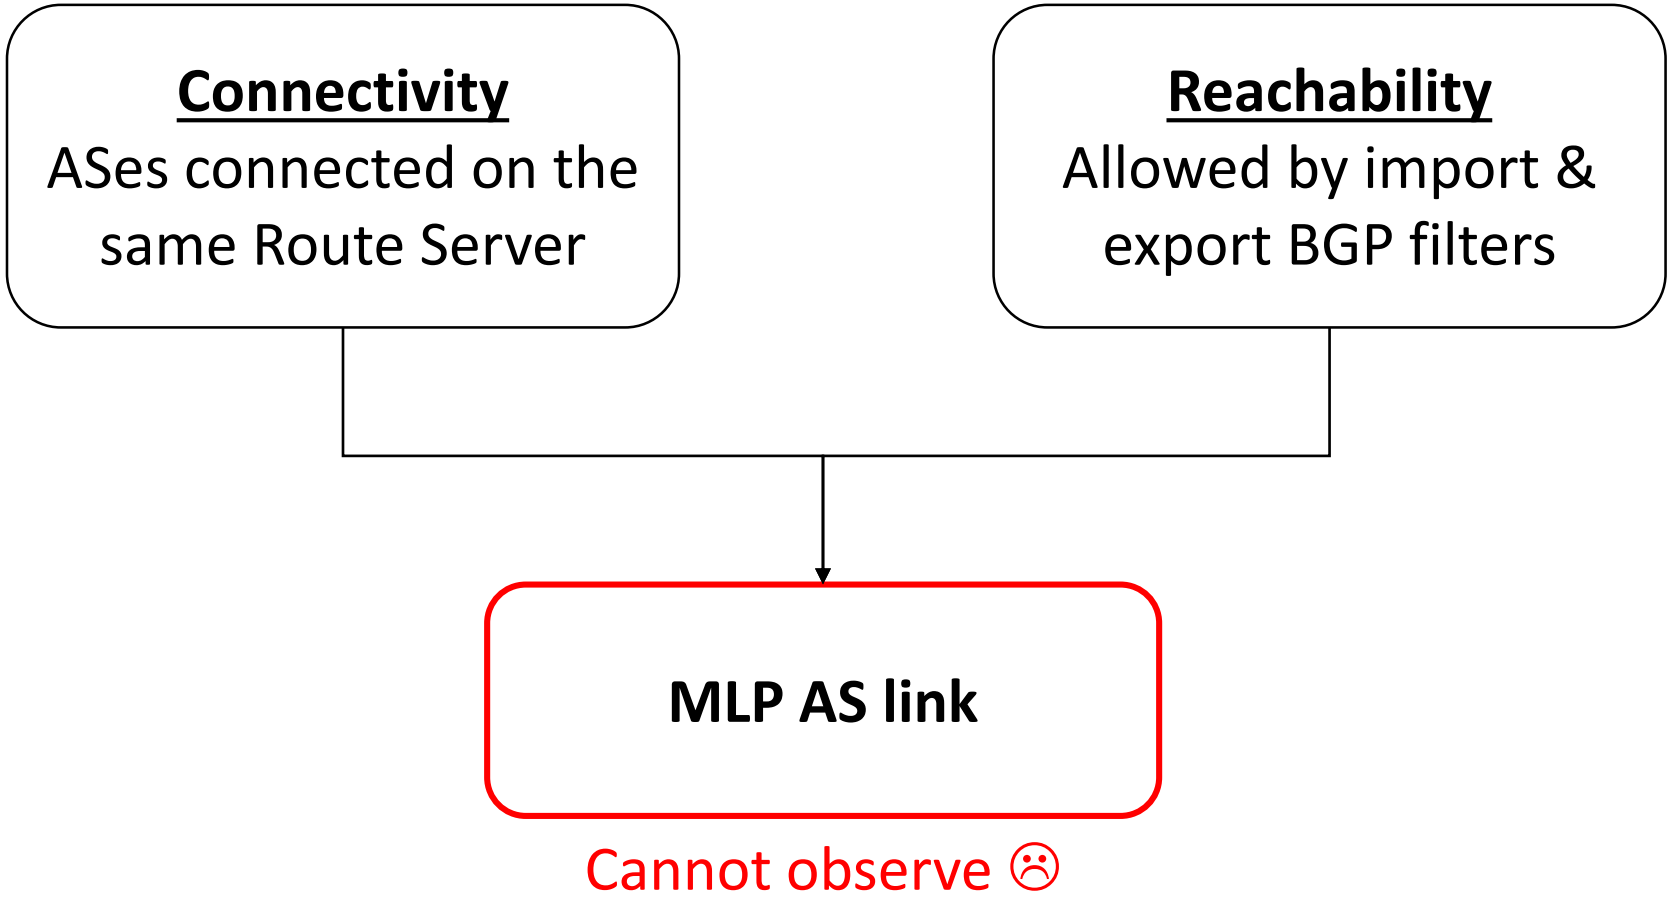

How MLP links are established?

Can observe ☺
- IXP websites
- PeeringDB
- EuroIX
- PCH
- RPSL

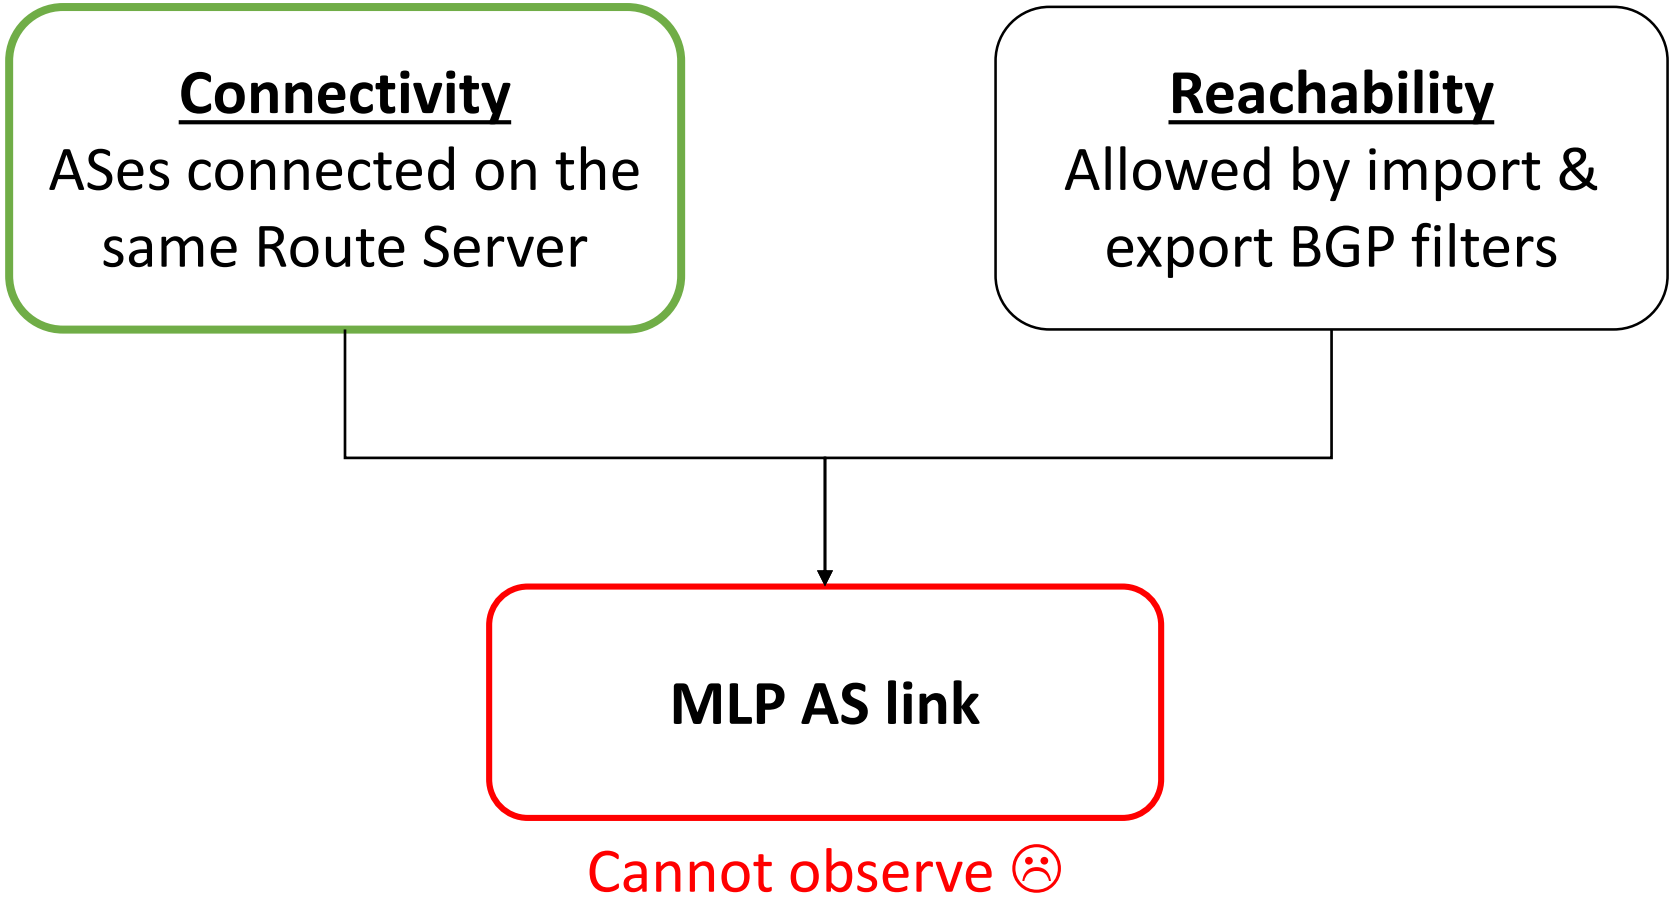

How MLP links are established?

Can observe ☺
- IXP websites
- PeeringDB
- EuroIX
- PCH
- RPSL

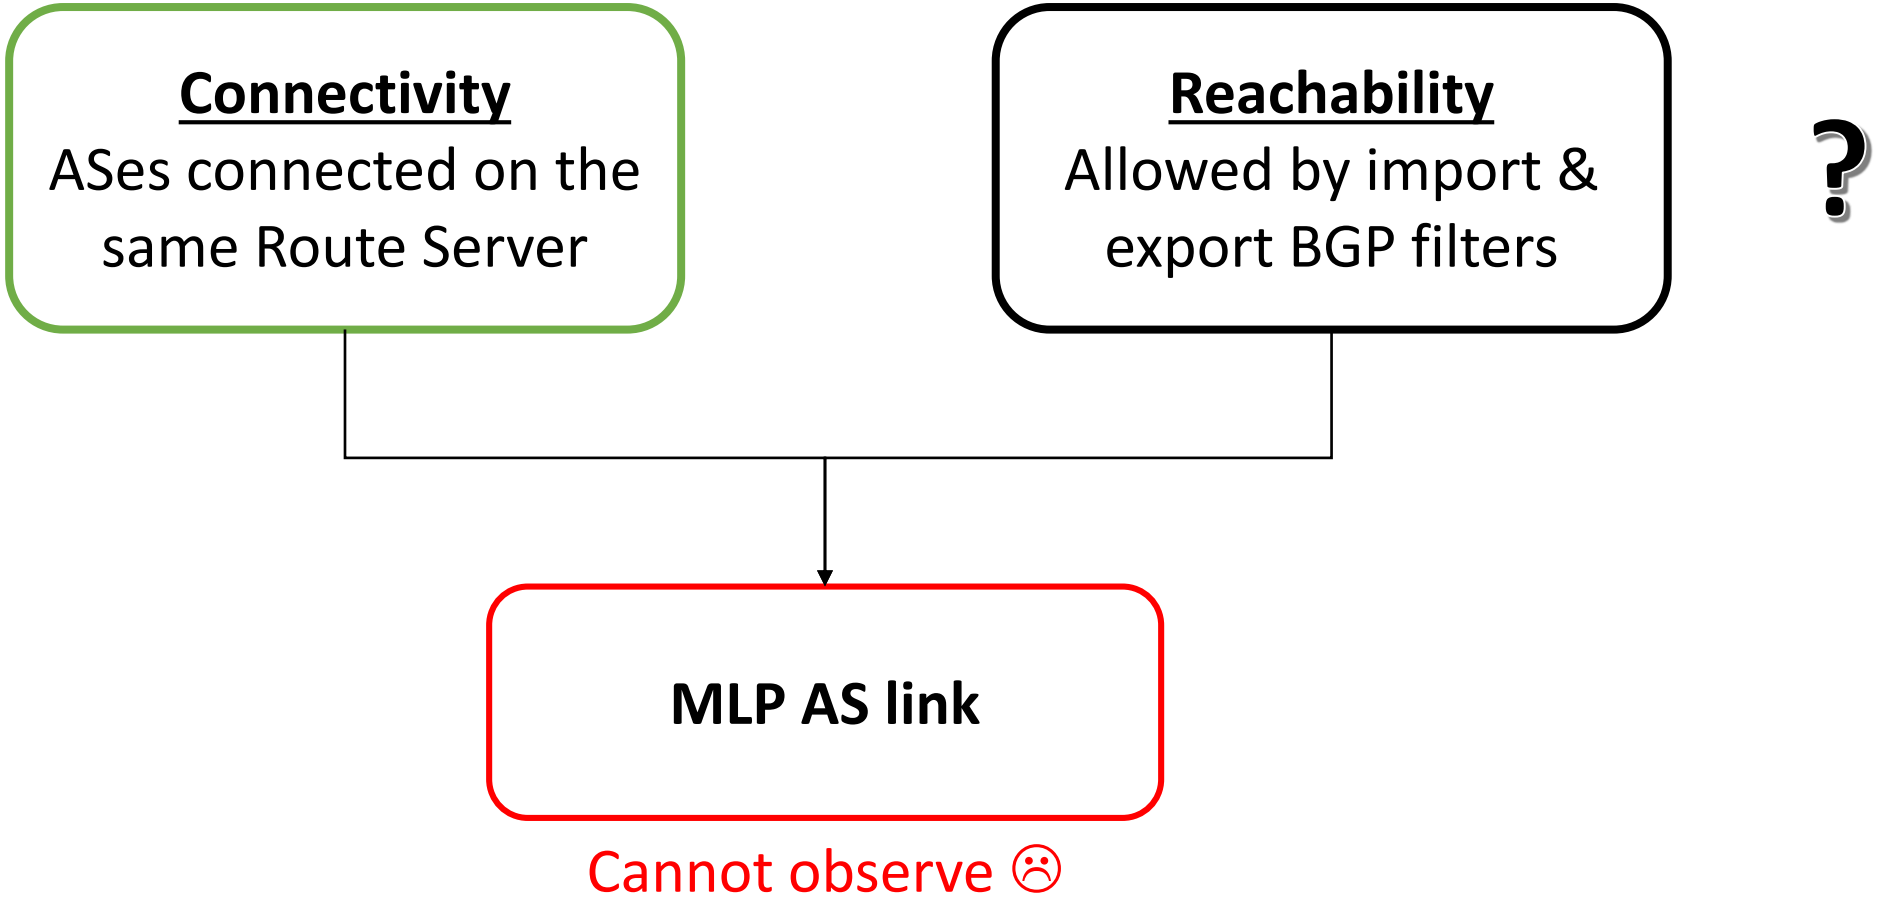

Default Route Server behaviour

Advertise to all

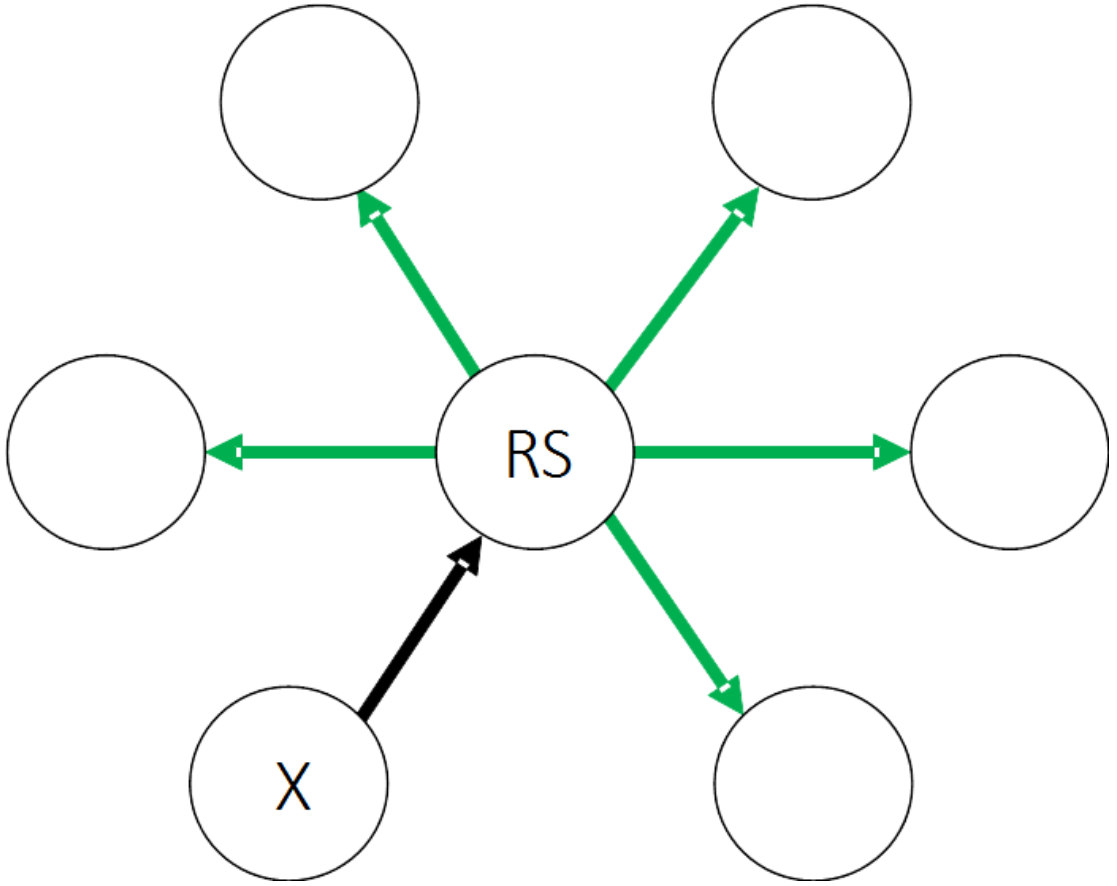

Advertisement Control BGP Communities

- **Transitive** BGP attribute
- Tags prefix advertisement with metadata
- 32-bit values divided in two parts:

ACTION:TARGET

- MSK-IX example:
  - 8635:8635 → Allow:All
  - 0:123 → Block:123
  - 0:8635 → Block:All
  - 8635:123 → Allow:123

Route Server Advertisement Control

Advertise to all except specific ASes

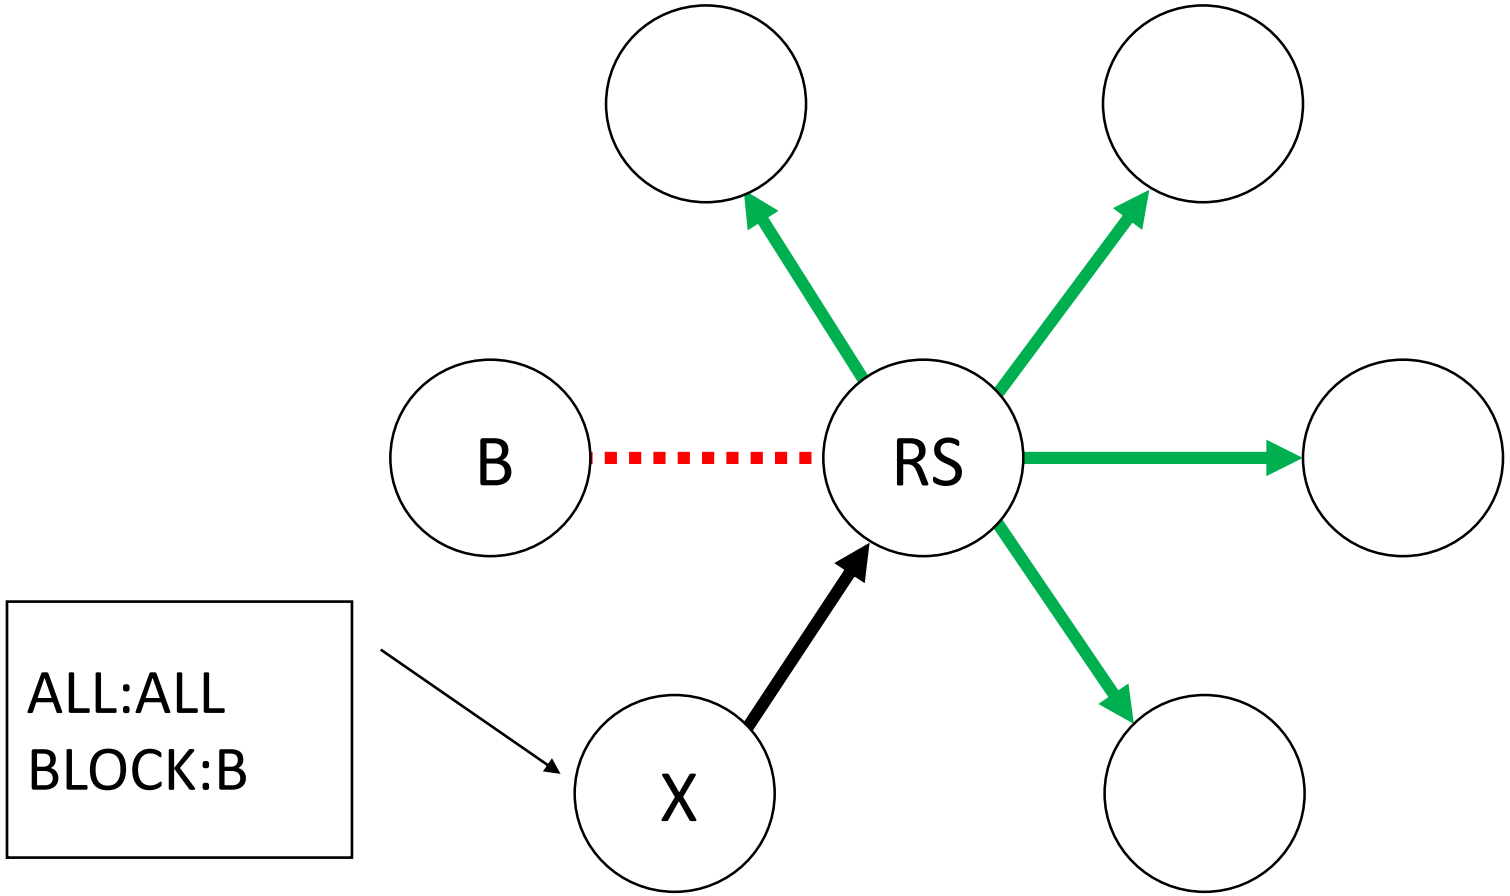

Route Server Advertisement Control

Advertise to all except specific ASes

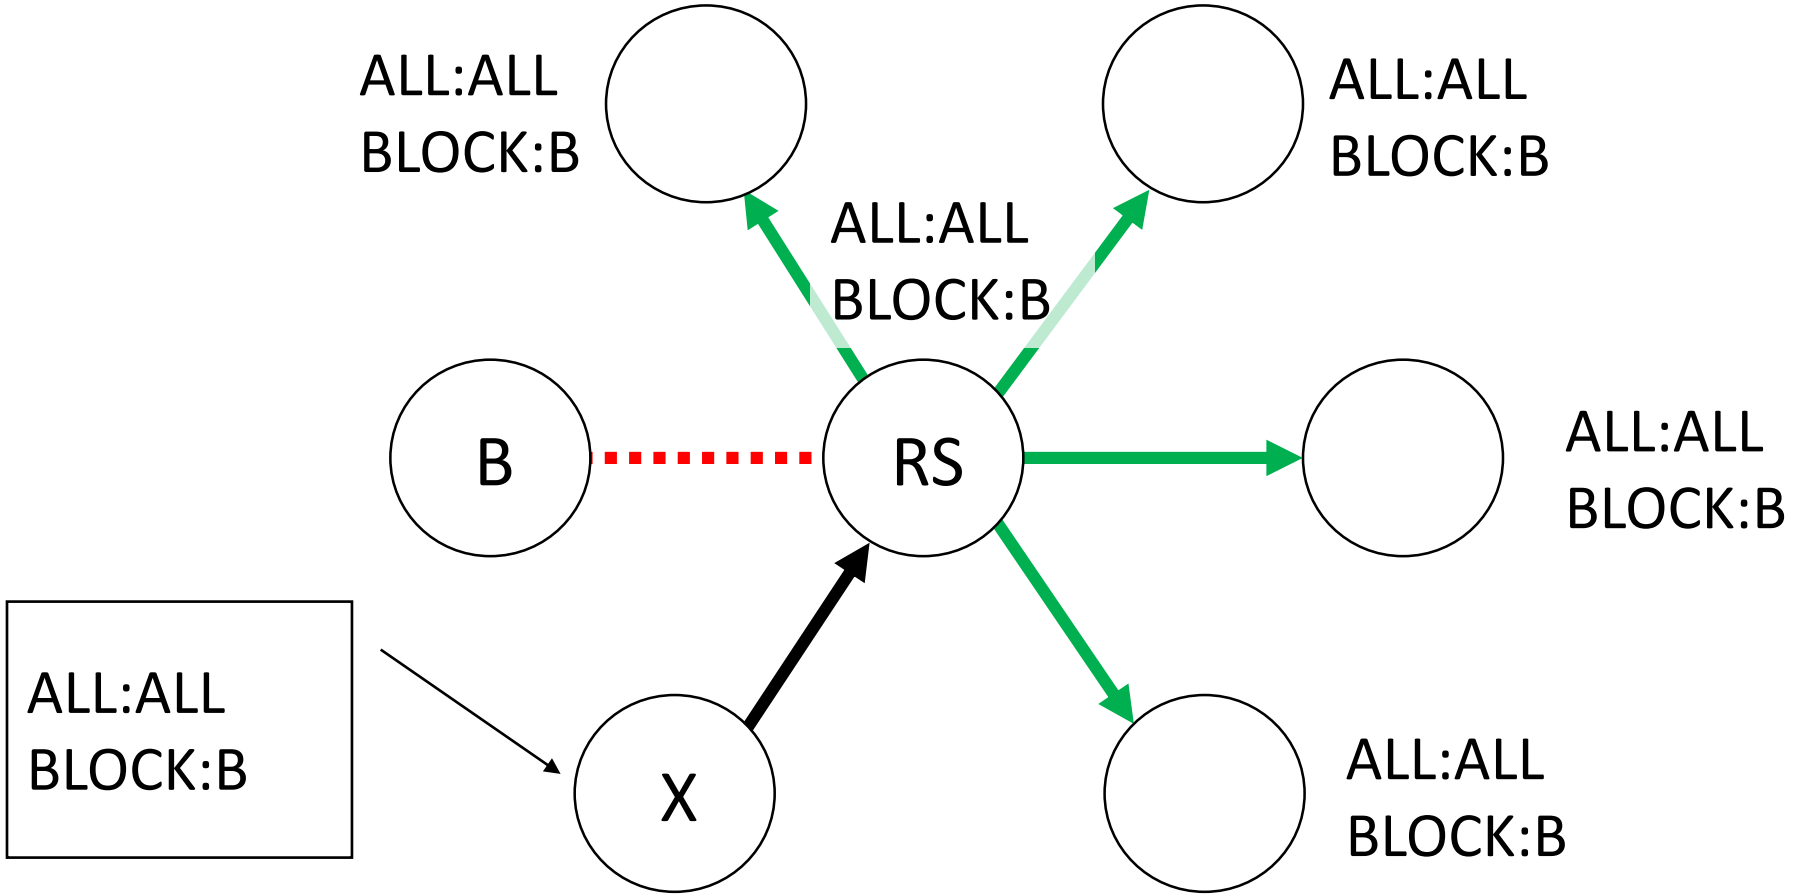

Measurement through RS Looking Glasses

- show ip bgp summary
  - Get ASes connected on Route Server
  - Query once
- show ip bgp neighbor
  - Get prefixes advertised by each Route Server member
  - Query *#members* times
- show ip bgp
  - Get BGP Communities set for a prefix
  - Query *#prefix* times

Measurement through RS Looking Glasses

- show ip bgp summary
  - Get ASes connected on Route Server
  - Query once
- show ip bgp neighbor
  - Get prefixes advertised by each Route Server member
  - Query *#members* times
- show ip bgp
  - Get BGP Communities set for a prefix
  - Query *#prefix* times

Measurement through passive BGP data

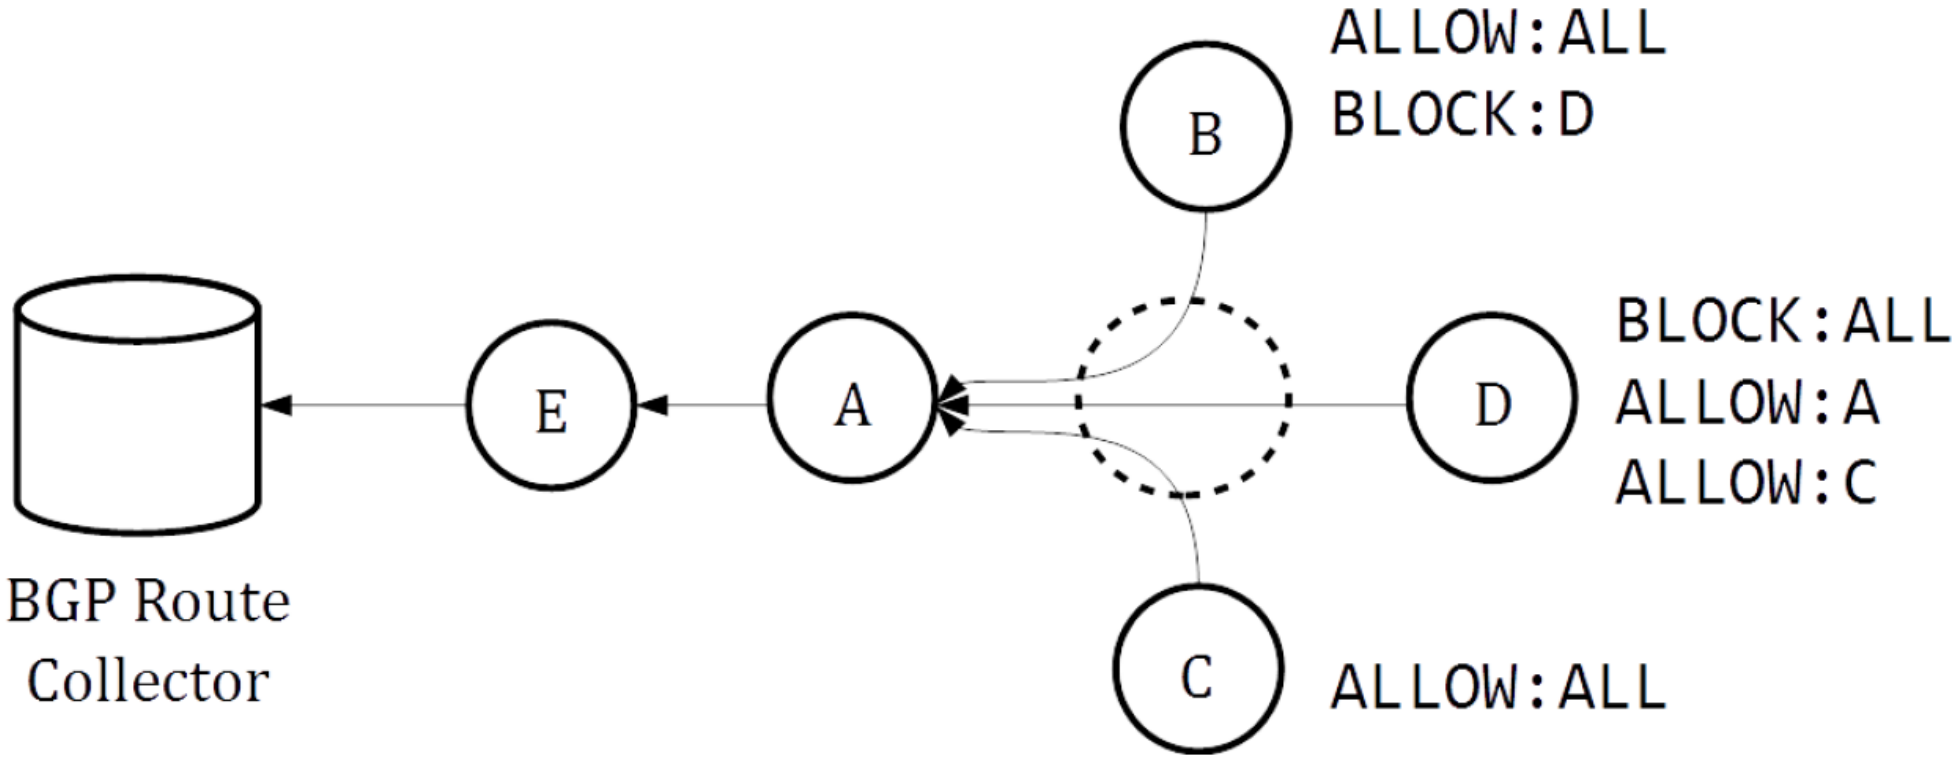

Measurement through passive BGP data

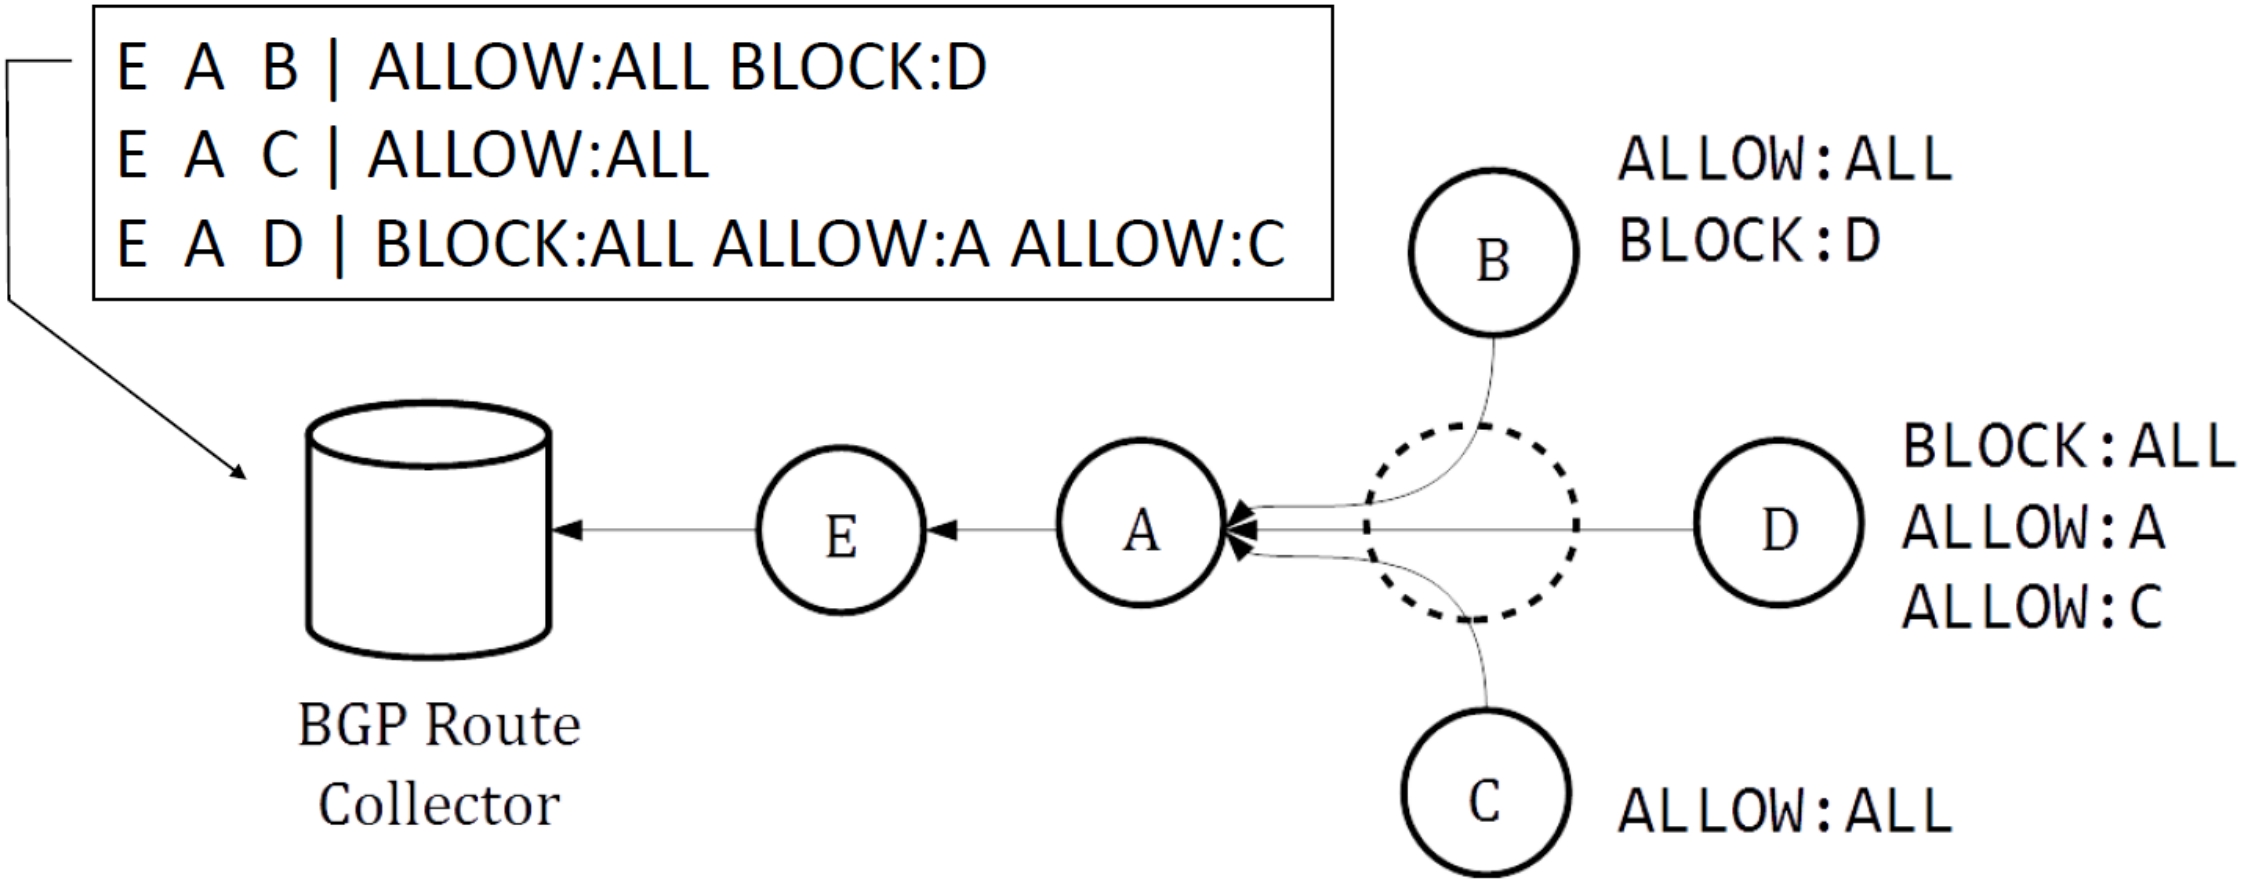

Measurement through passive BGP data

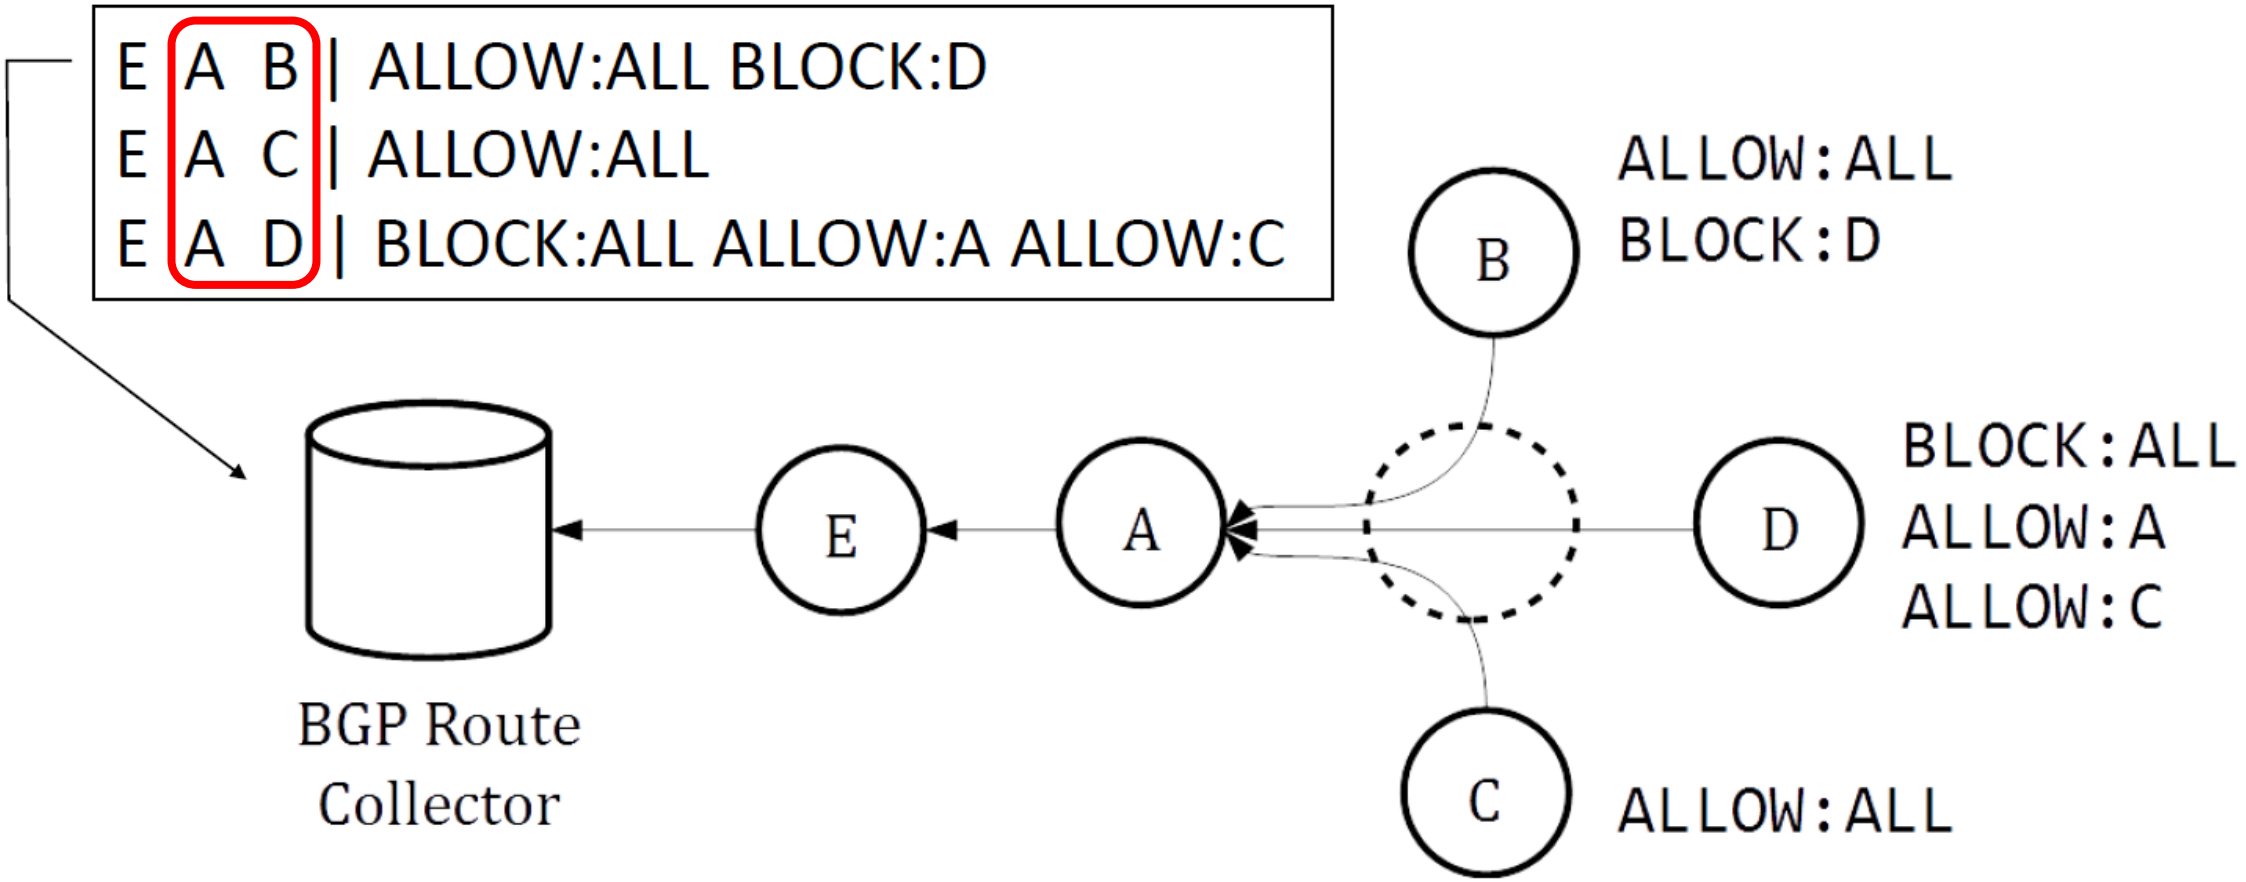

How MLP links are established?

Can observe ☺
- IXP websites
- PeeringDB
- EuroIX
- PCH
- RPSL

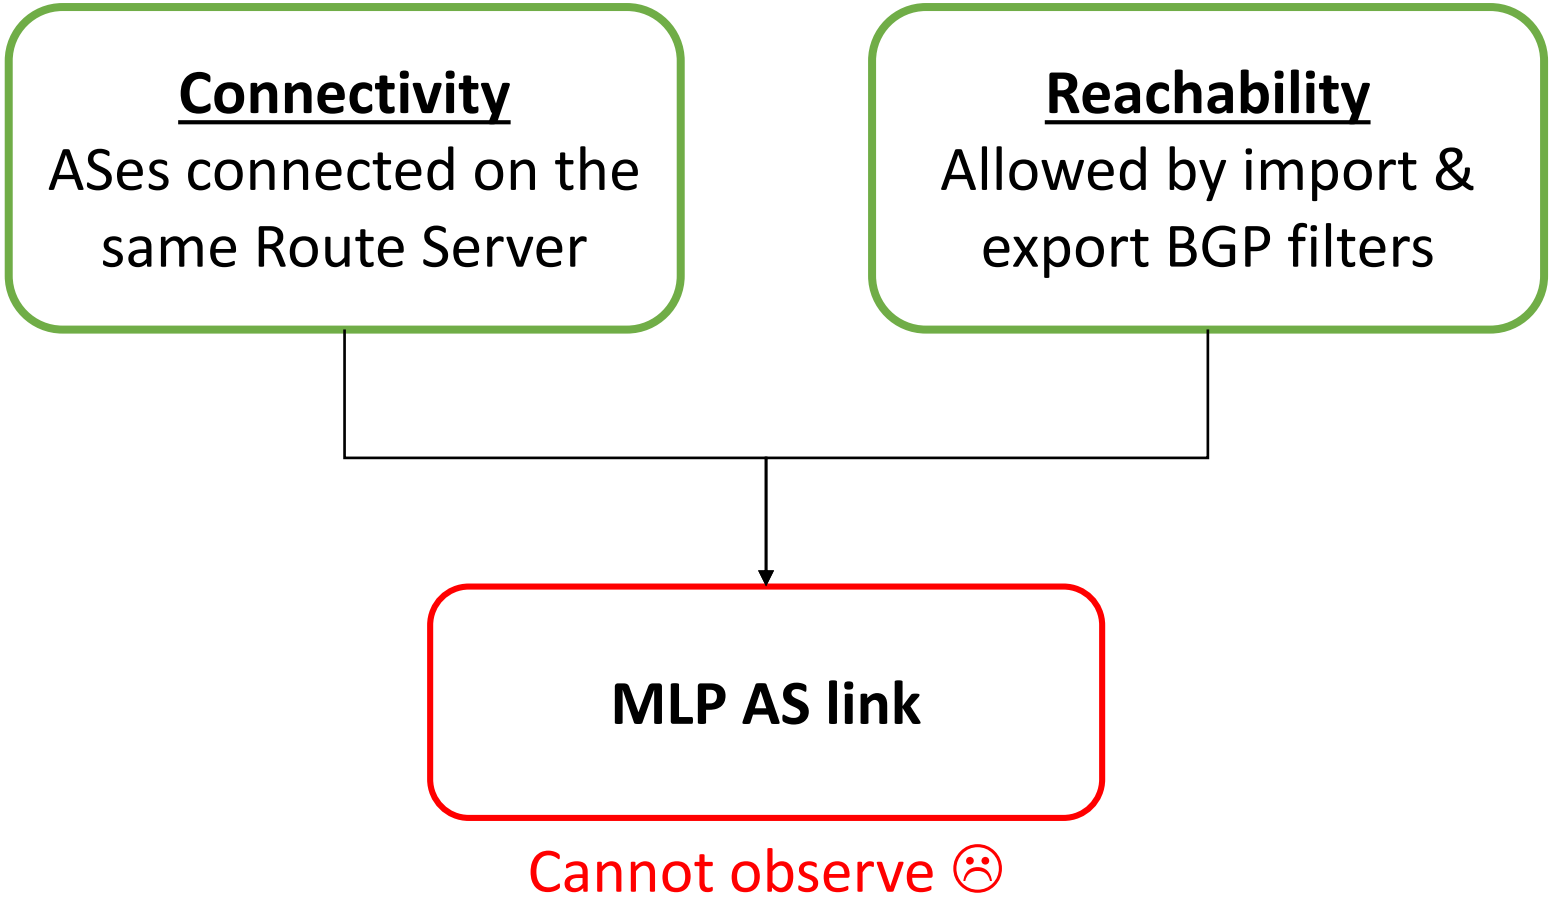

Can observe ☺
- BGP RS Redistribution Communities

How MLP links are established?

Can observe ☺
- IXP websites
- PeeringDB
- EuroIX
- PCH
- RPSL

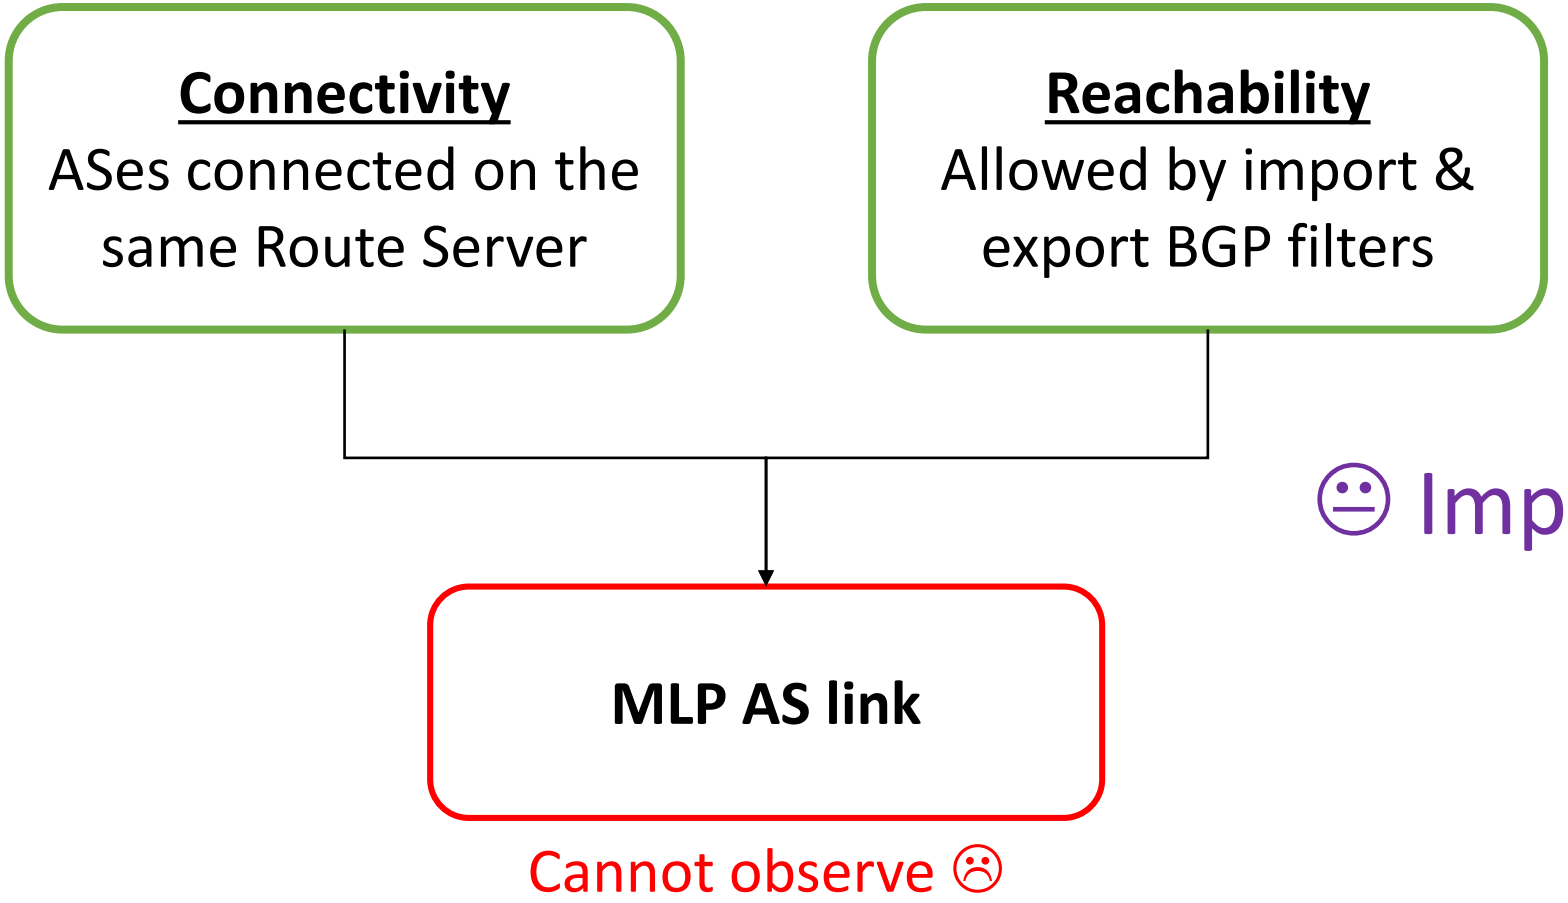

Can observe ☺
- BGP RS Redistribution Communities

😐 Import filters?

Reciprocity Assumption

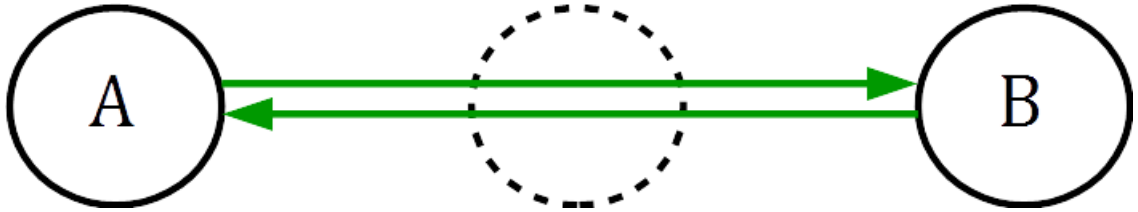

- Two ASes that allow each other to receive their traffic at the export filters will not block each other's traffic at the import filters
- Validation based on AMS-IX IRRdb filters:
  - Compared import/export filters from 230 AMS-IX Route Server members
  - Import filters never more restrictive than export filters
  - 50% of export filters more restrictive than import
    - Asymmetric (single direction) MLP peering
    - False negatives

Results

- Collected presence and permissions data for 13 large European IXPs
- 206,667 MLP links inferred between 1,362 ASes
- 14,276 links appear in more than one IXPs
  - Largest overlap between DE-CIX and AMS-IX

IXP	Members*	Links
MSK-IX	348	58,501
DE-CIX	369	54,082
AMS-IX	351	49,249
PLIX	211	21,911
LINX	176	14,759

Validation

- Tested inferences against links collected from AS paths of **70** RS member looking glasses
- Validation repeated twice, May 2013 and October 2013
- **26,392** links tested, **98.4%** successfully validated overall
- Highest validation: 100% (TOP-IX)
- Lowest validation: 96.9% (FranceIX)
- Average for all IXPs: 98.6%

Limitations of validation

- Links part of less preferred paths are hidden from looking glasses that display only the active paths
- ASes left Route Server or changed policy between inference – validation

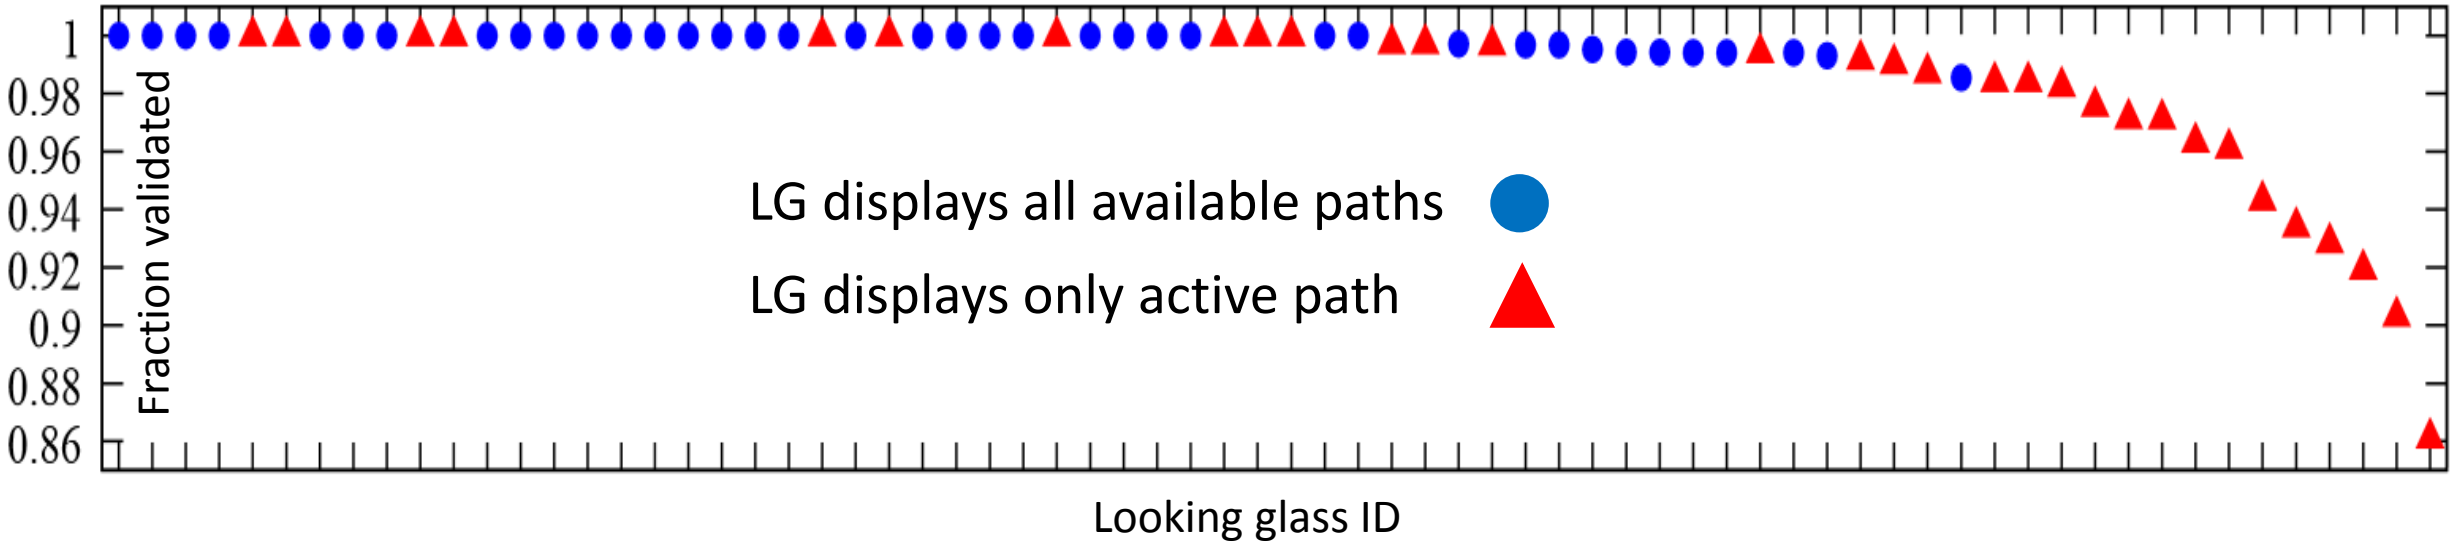

Comparison against observable p2p links

- 12% overlap with passive BGP measurements (Routeview+RIPE RIS+PCH)
- 2% overlap with active traceroute (Ark + DIMES)

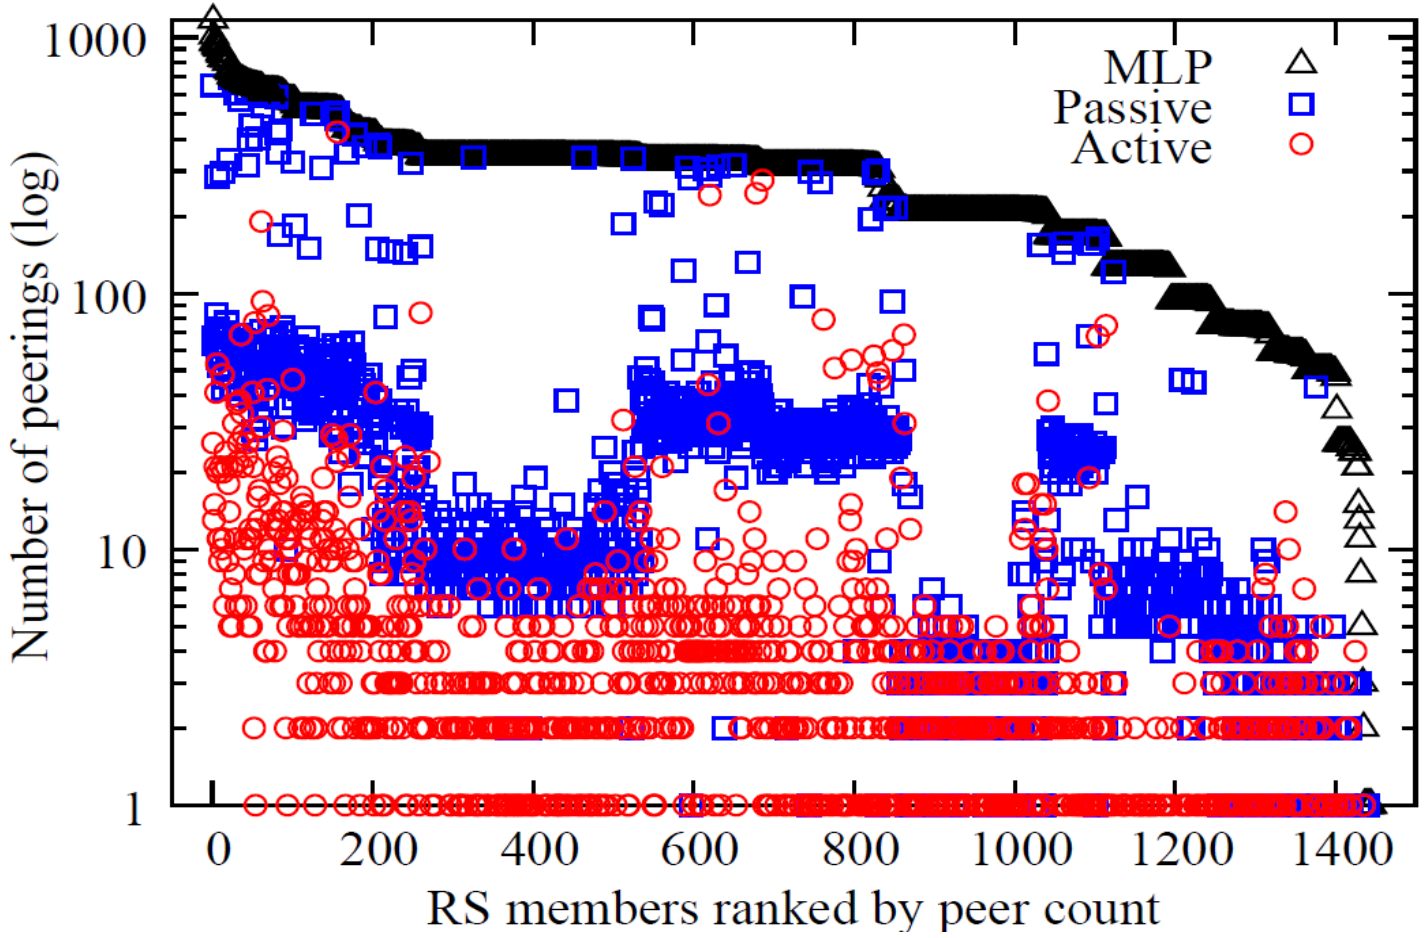

Majority of MLP links involve stub ASes

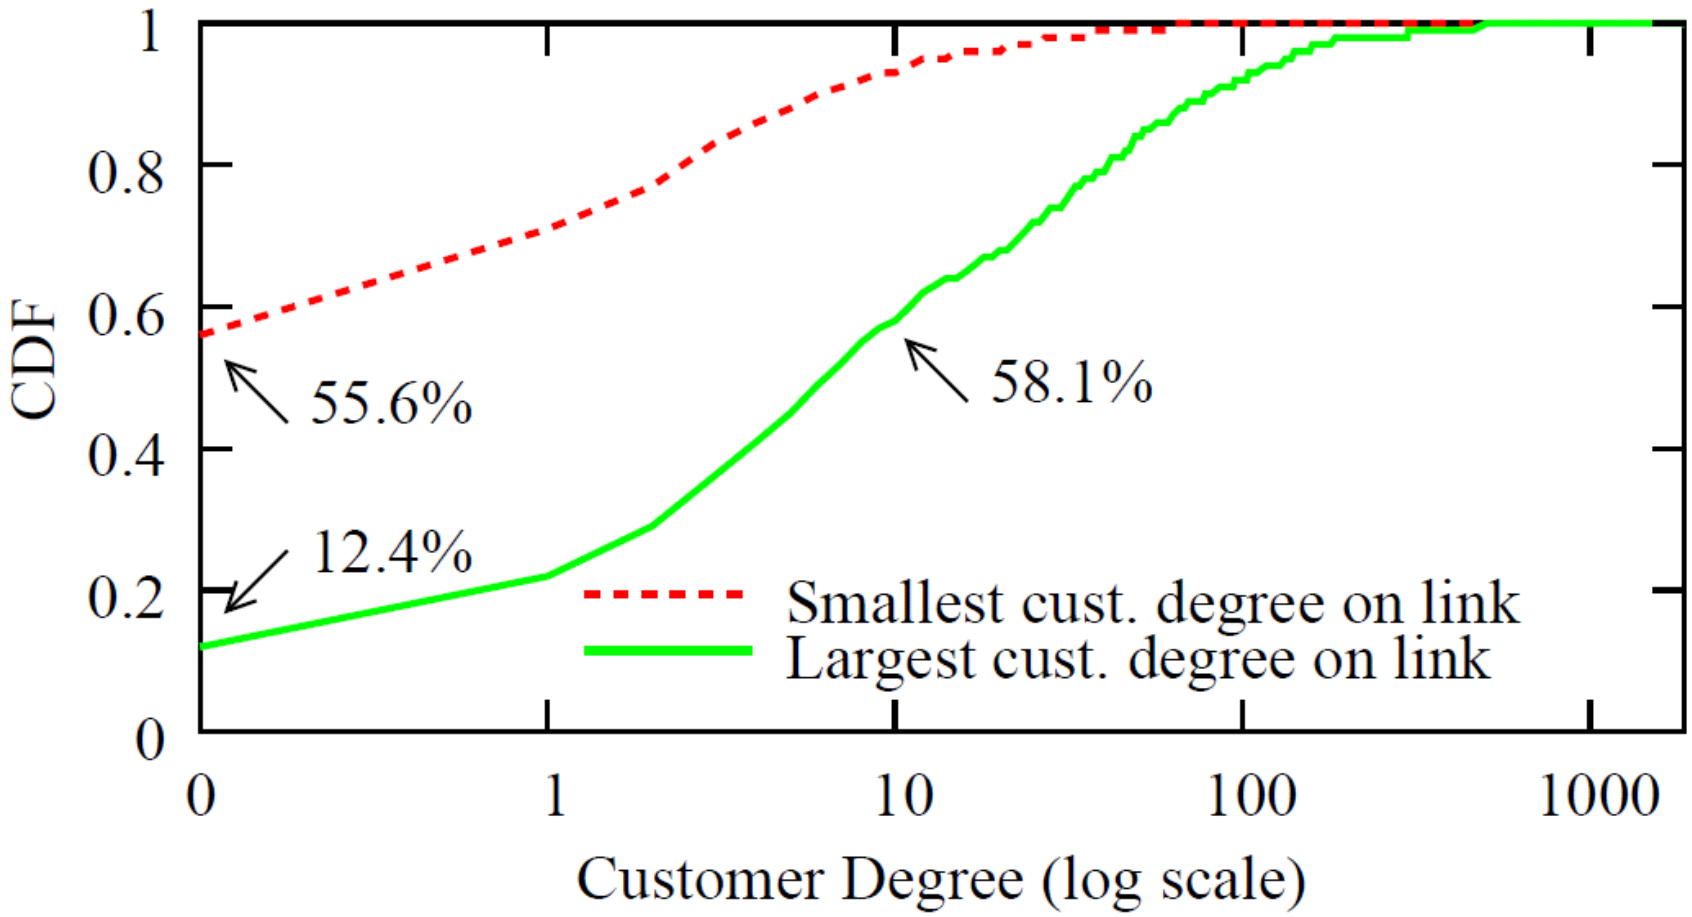

Extremely high MLP density

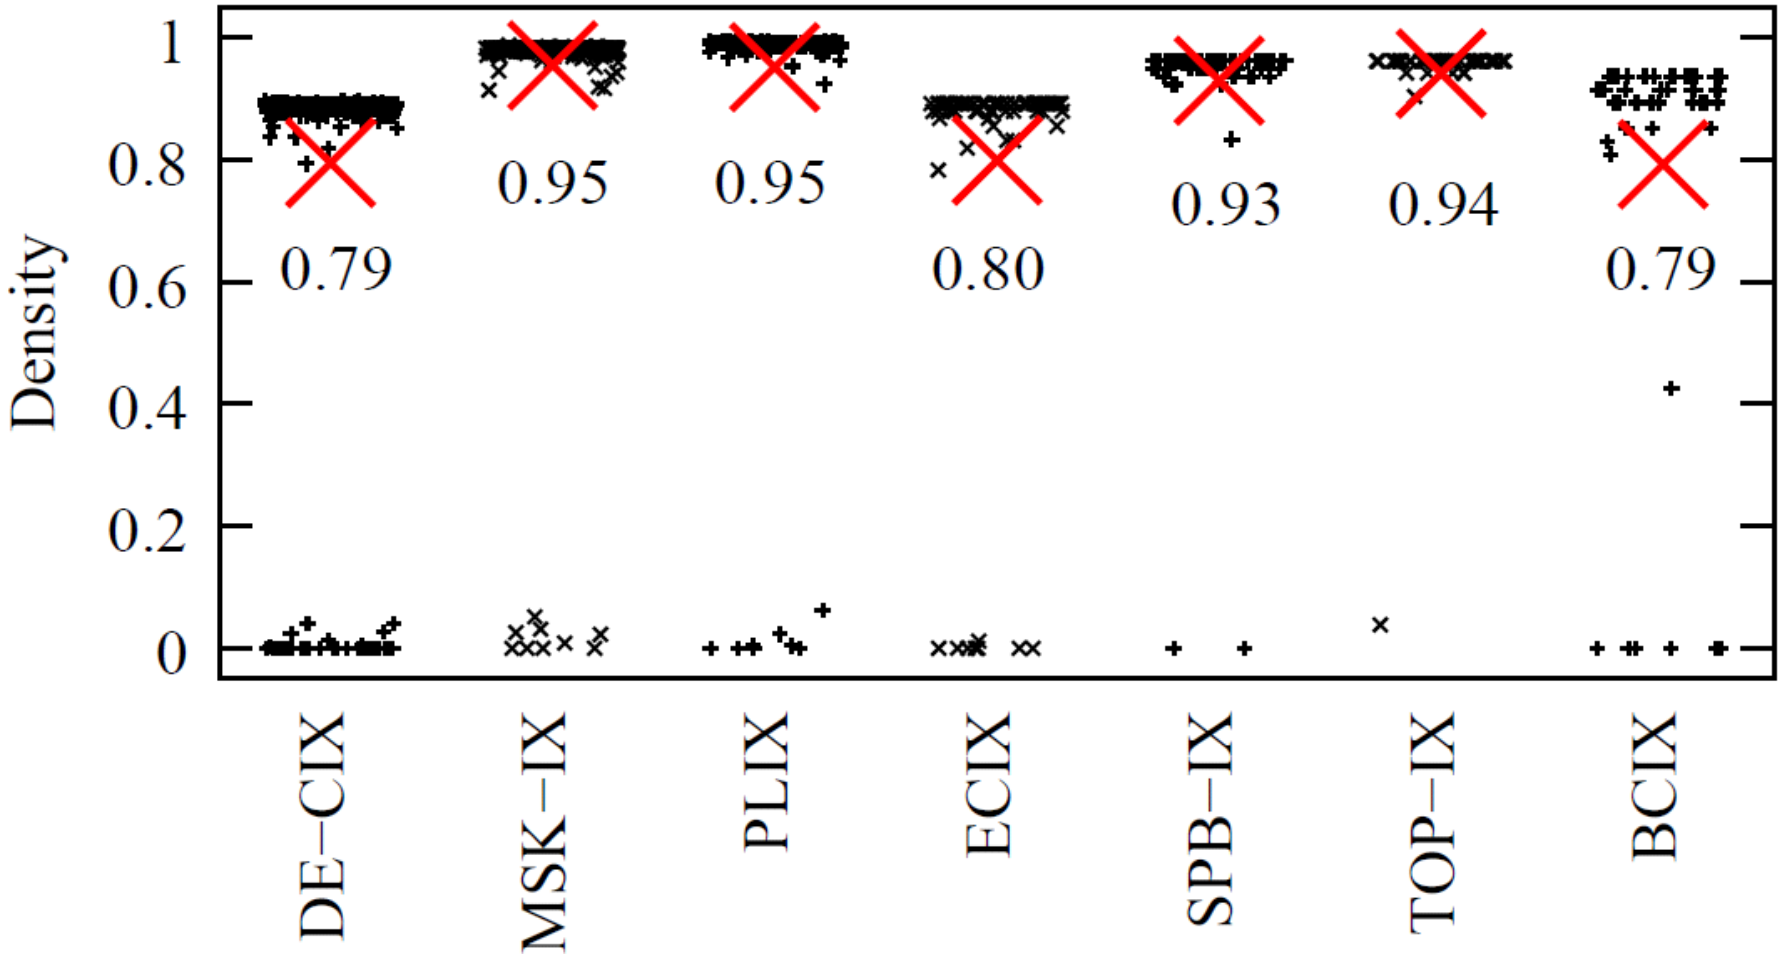

Limitations

- Cases where our methodology cannot be applied:
  - IXPs without Route Servers
  - Route Servers that do not use BGP Communities for advertisement control
  - Route Servers that strip out BGP Communities before propagating advertisements
- Coverage of European IXPs only (for the moment)

Conclusions

- Propose and implement a new approach to infer Multilateral Peering links
- Utilize only public data sources
- Apply algorithm at 13 IXPs to infer 206K links
  - 88% missing from RouteViews/RIPE RIS BGP data
- Validate 26K links with 98.4% success rate

Thank you!

Questions?

v.giotsas@cs.ucl.ac.uk

Self-reported Peering Policy can be Misleading

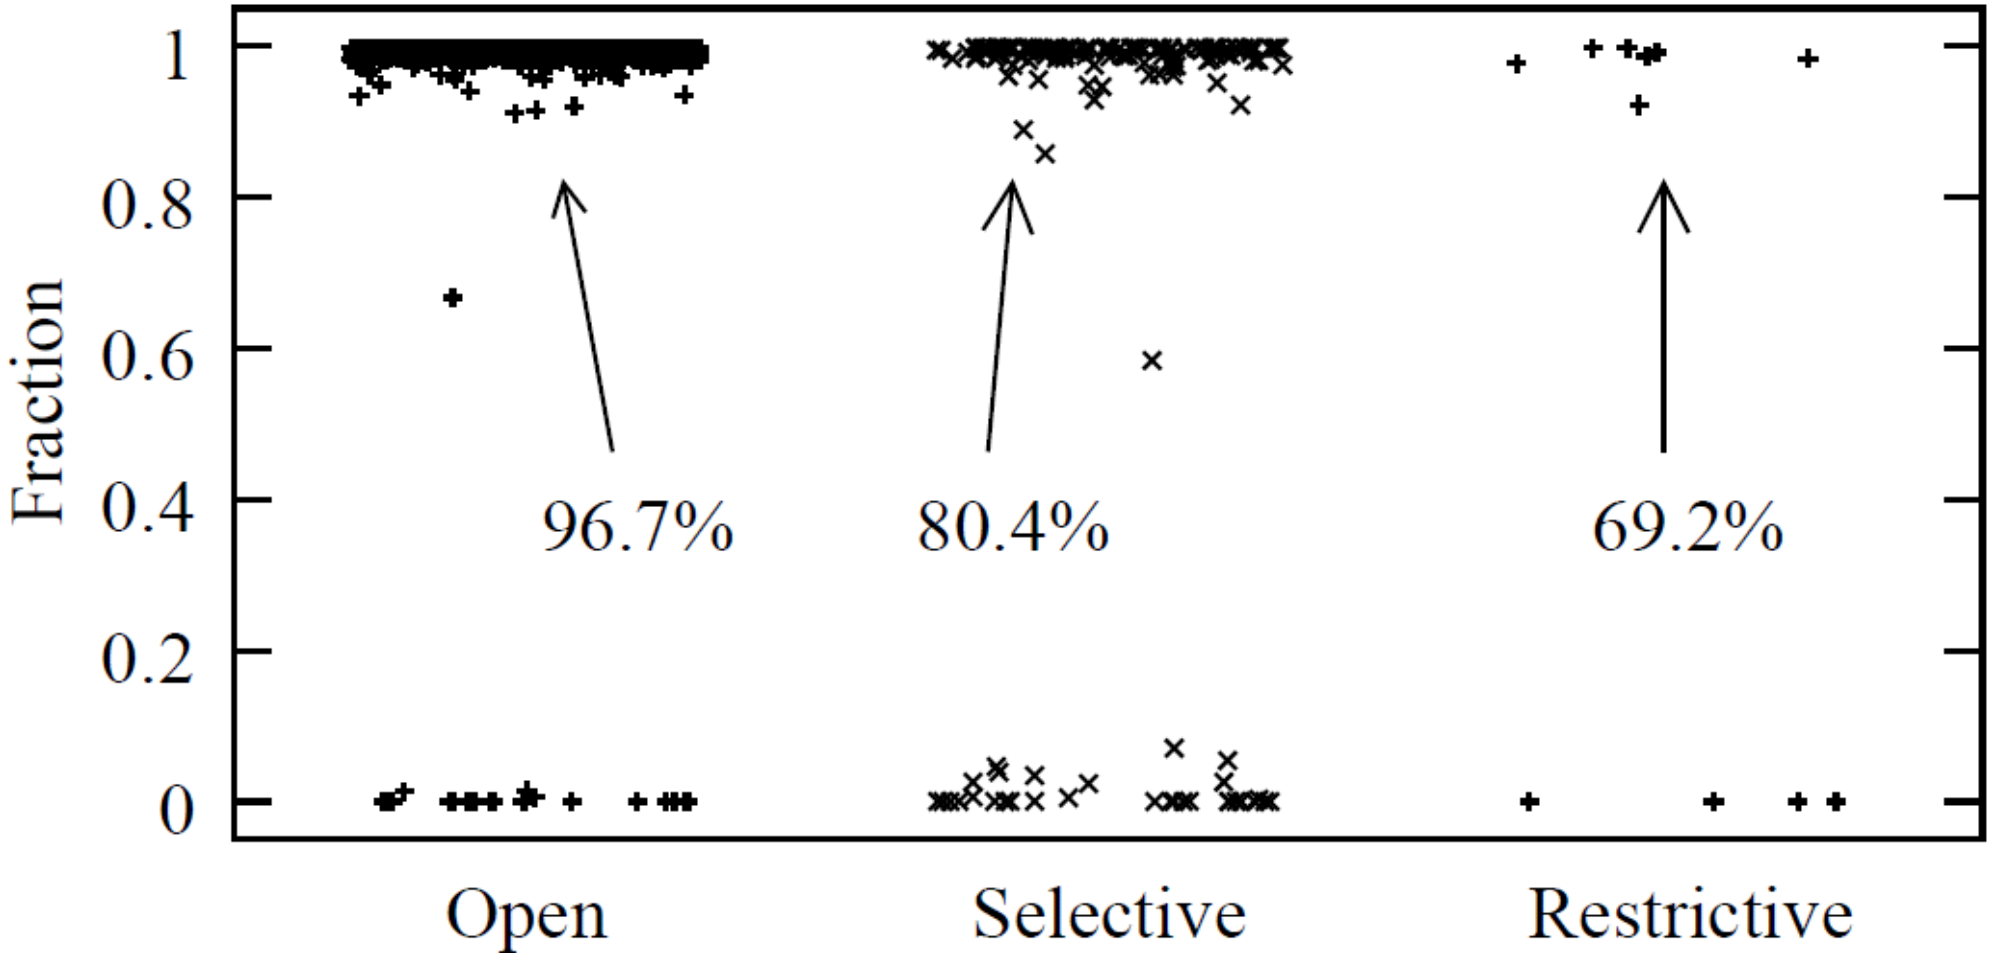

Peering policies can depend on location

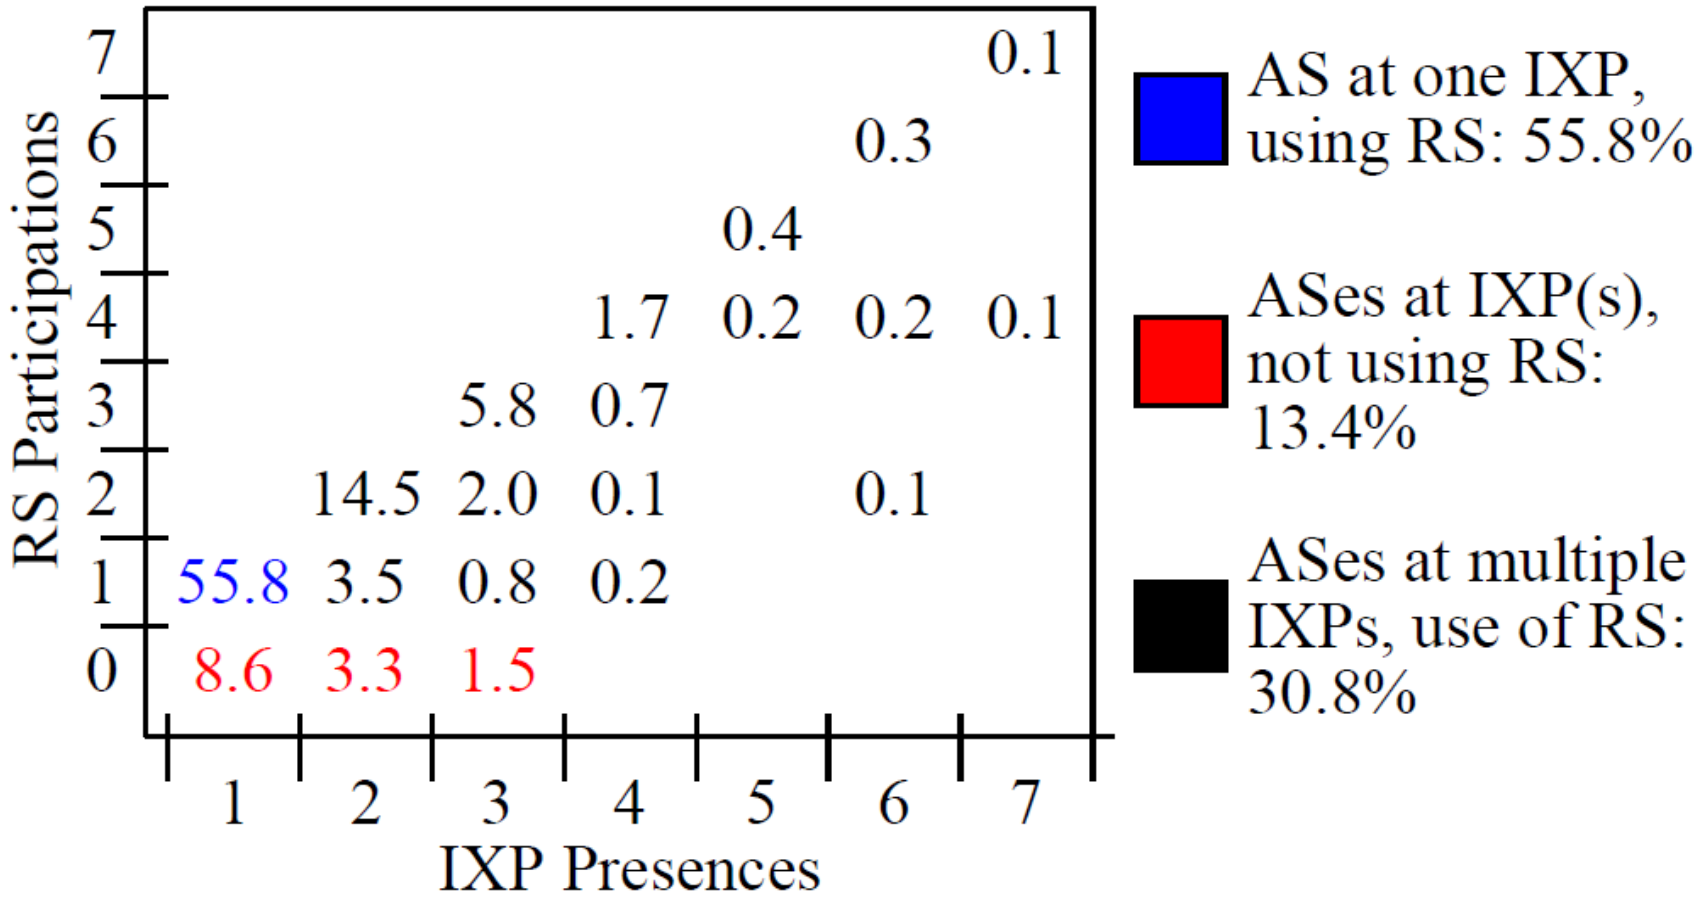

Backup slide 3: Prefixes advertised by multiple Route Server members