MAPPING PEERING INTERCONNECTIONS TO A FACILITY

Vasileios Giotsas [1]
Georgios Smaragdakis [2]
Bradley Huffaker [1]
Matthew Luckie [3]
kc claffy [1]

vgiotsas@caida.org

WIE 2015

[1] UCSD/CAIDA [2] MIT/TU Berlin [3] University of Waikato

The AS-level topology abstracts a much richer connectivity map

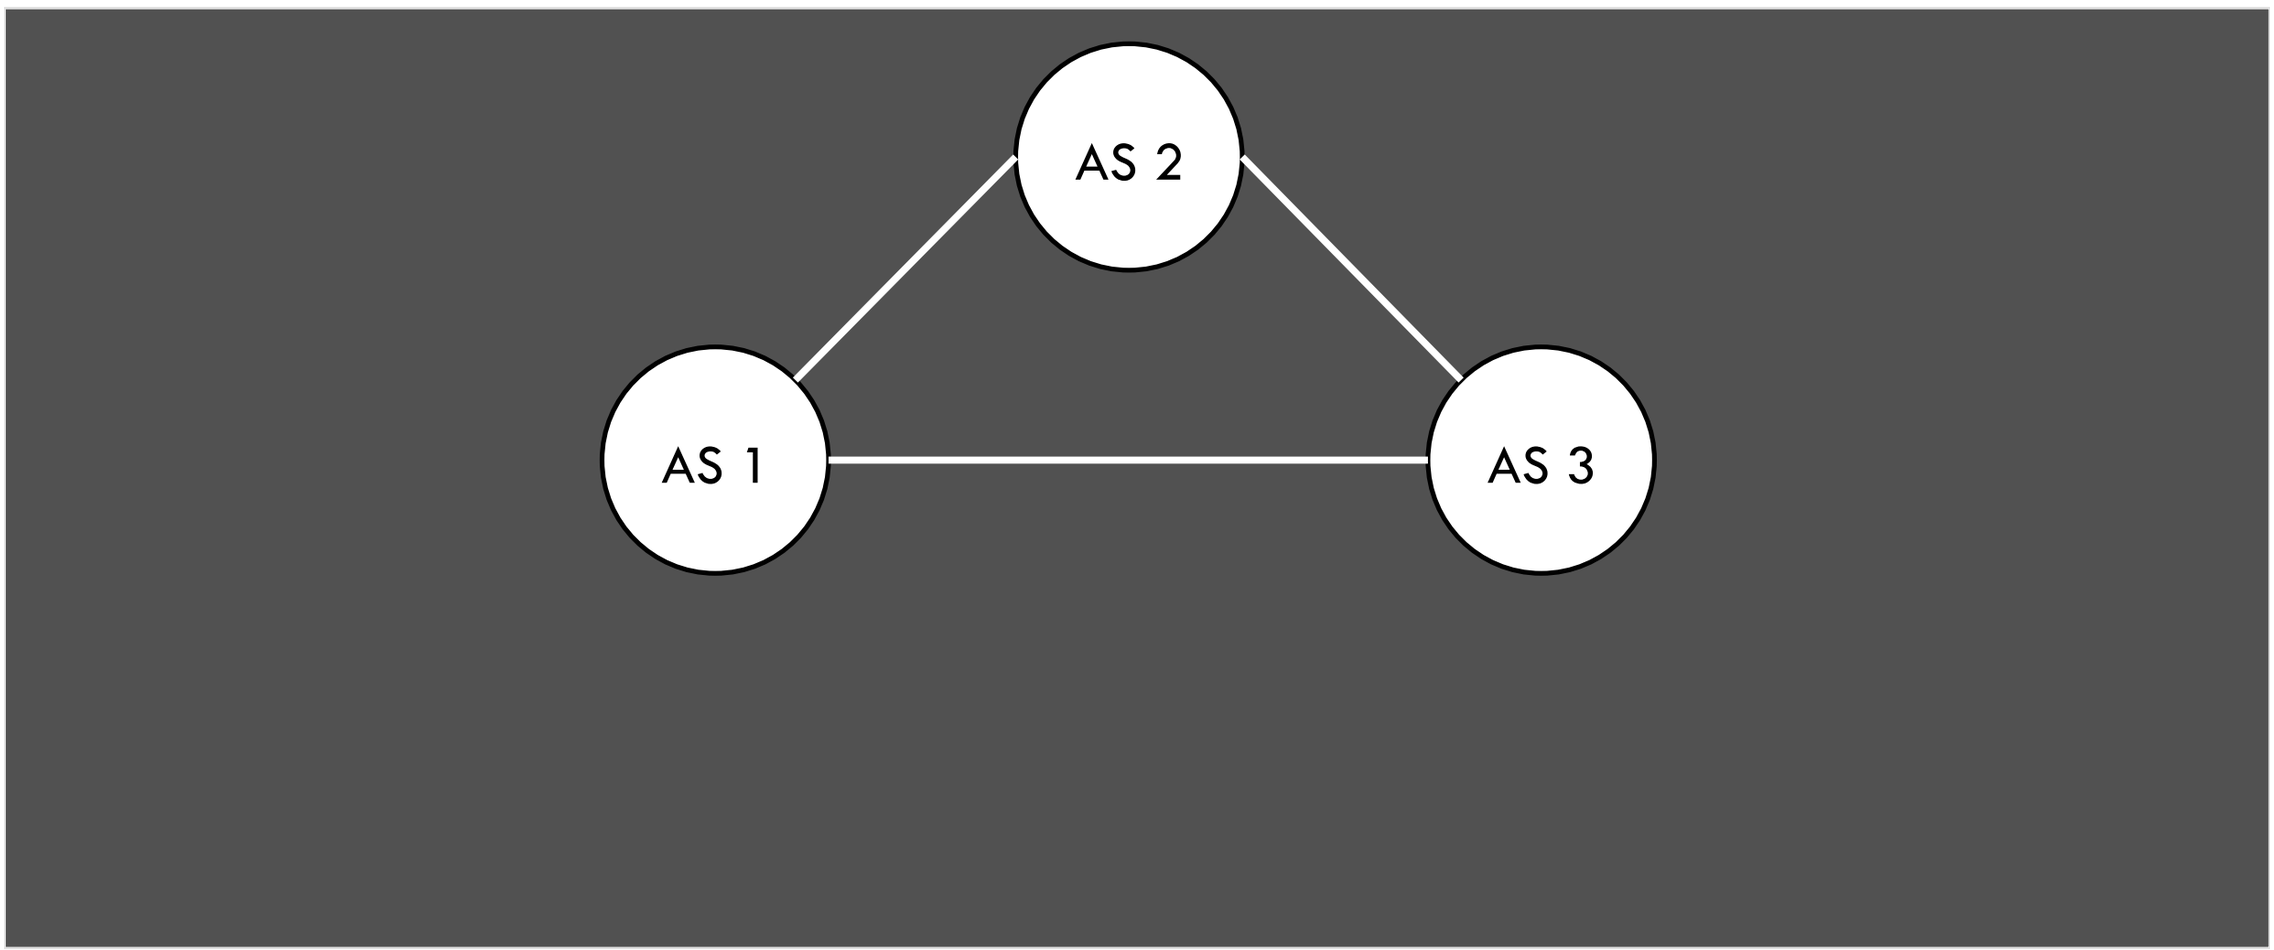

The AS-level topology abstracts a much richer connectivity map

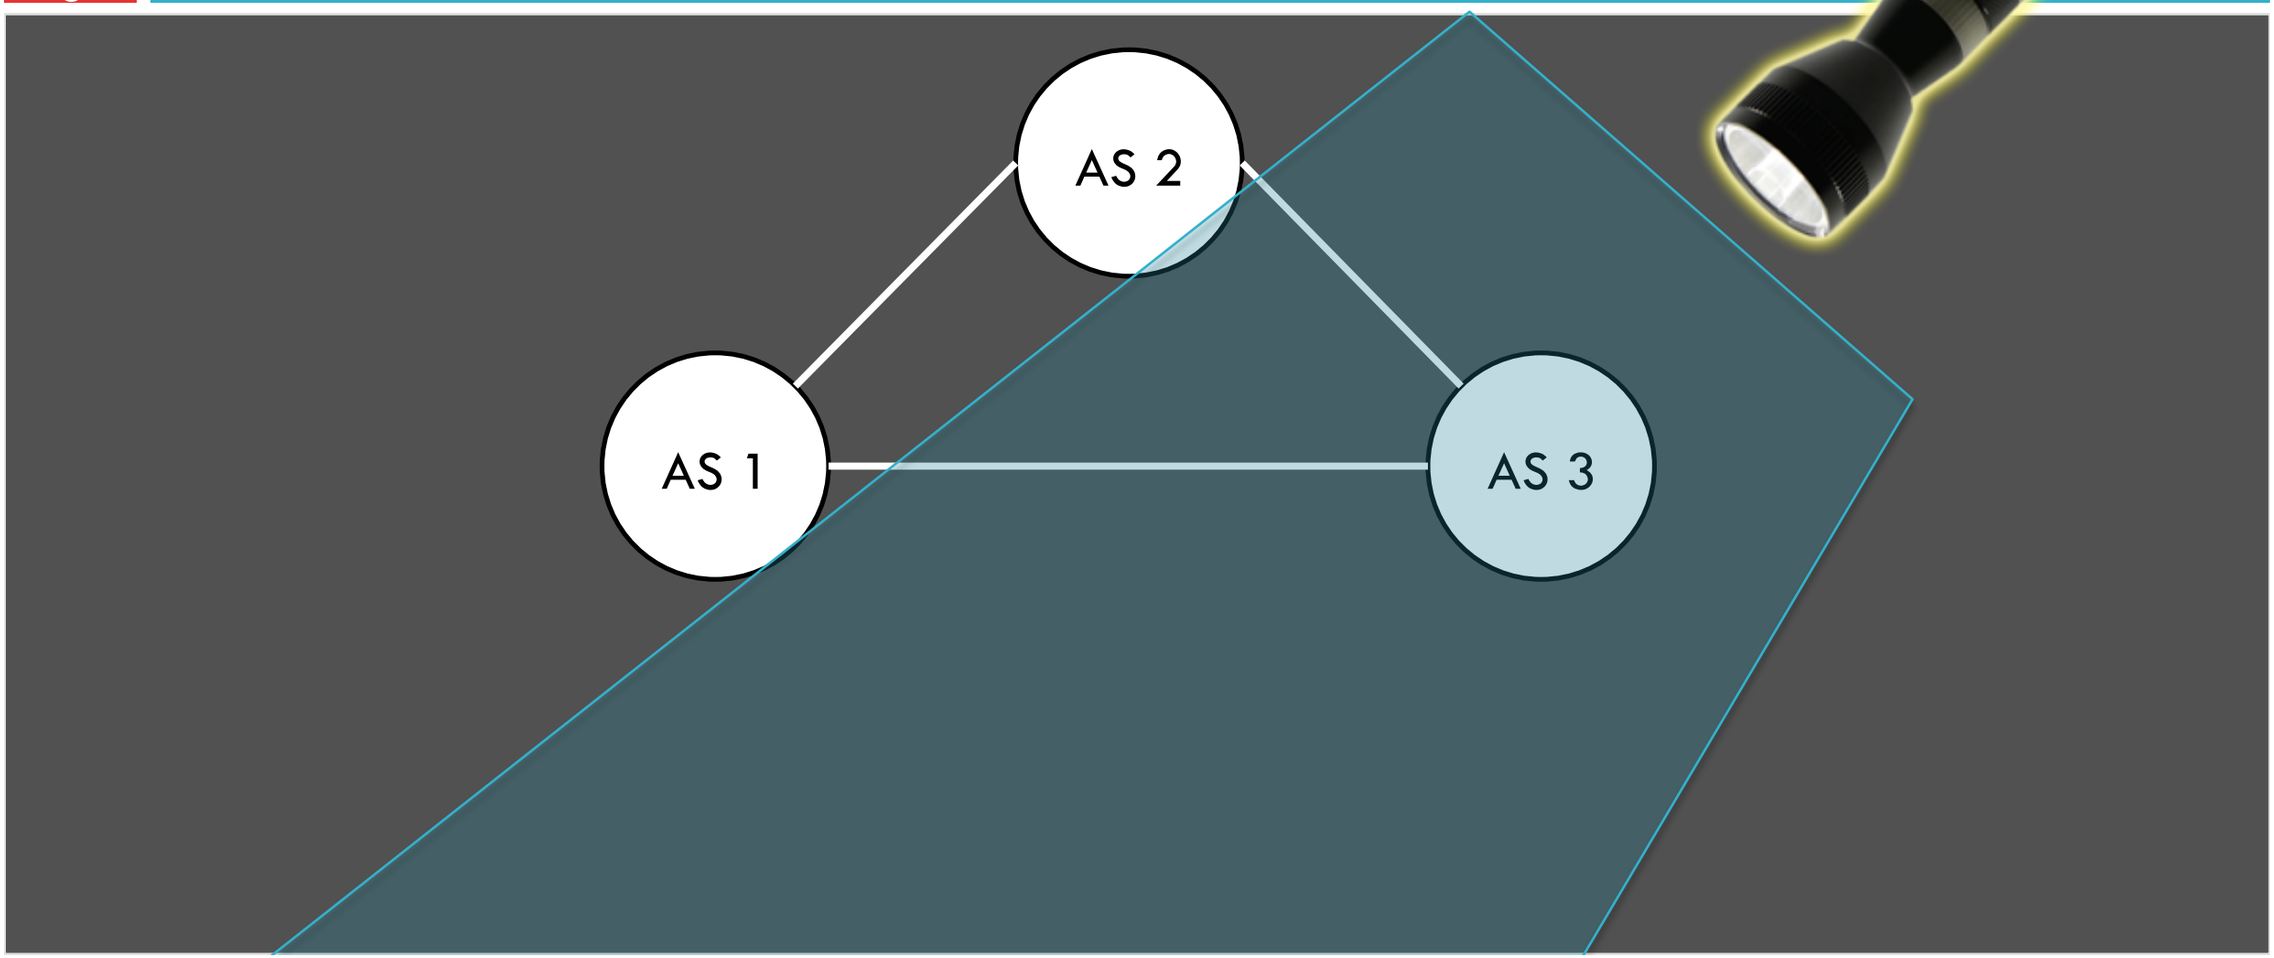

The AS-level topology abstracts a much richer connectivity map

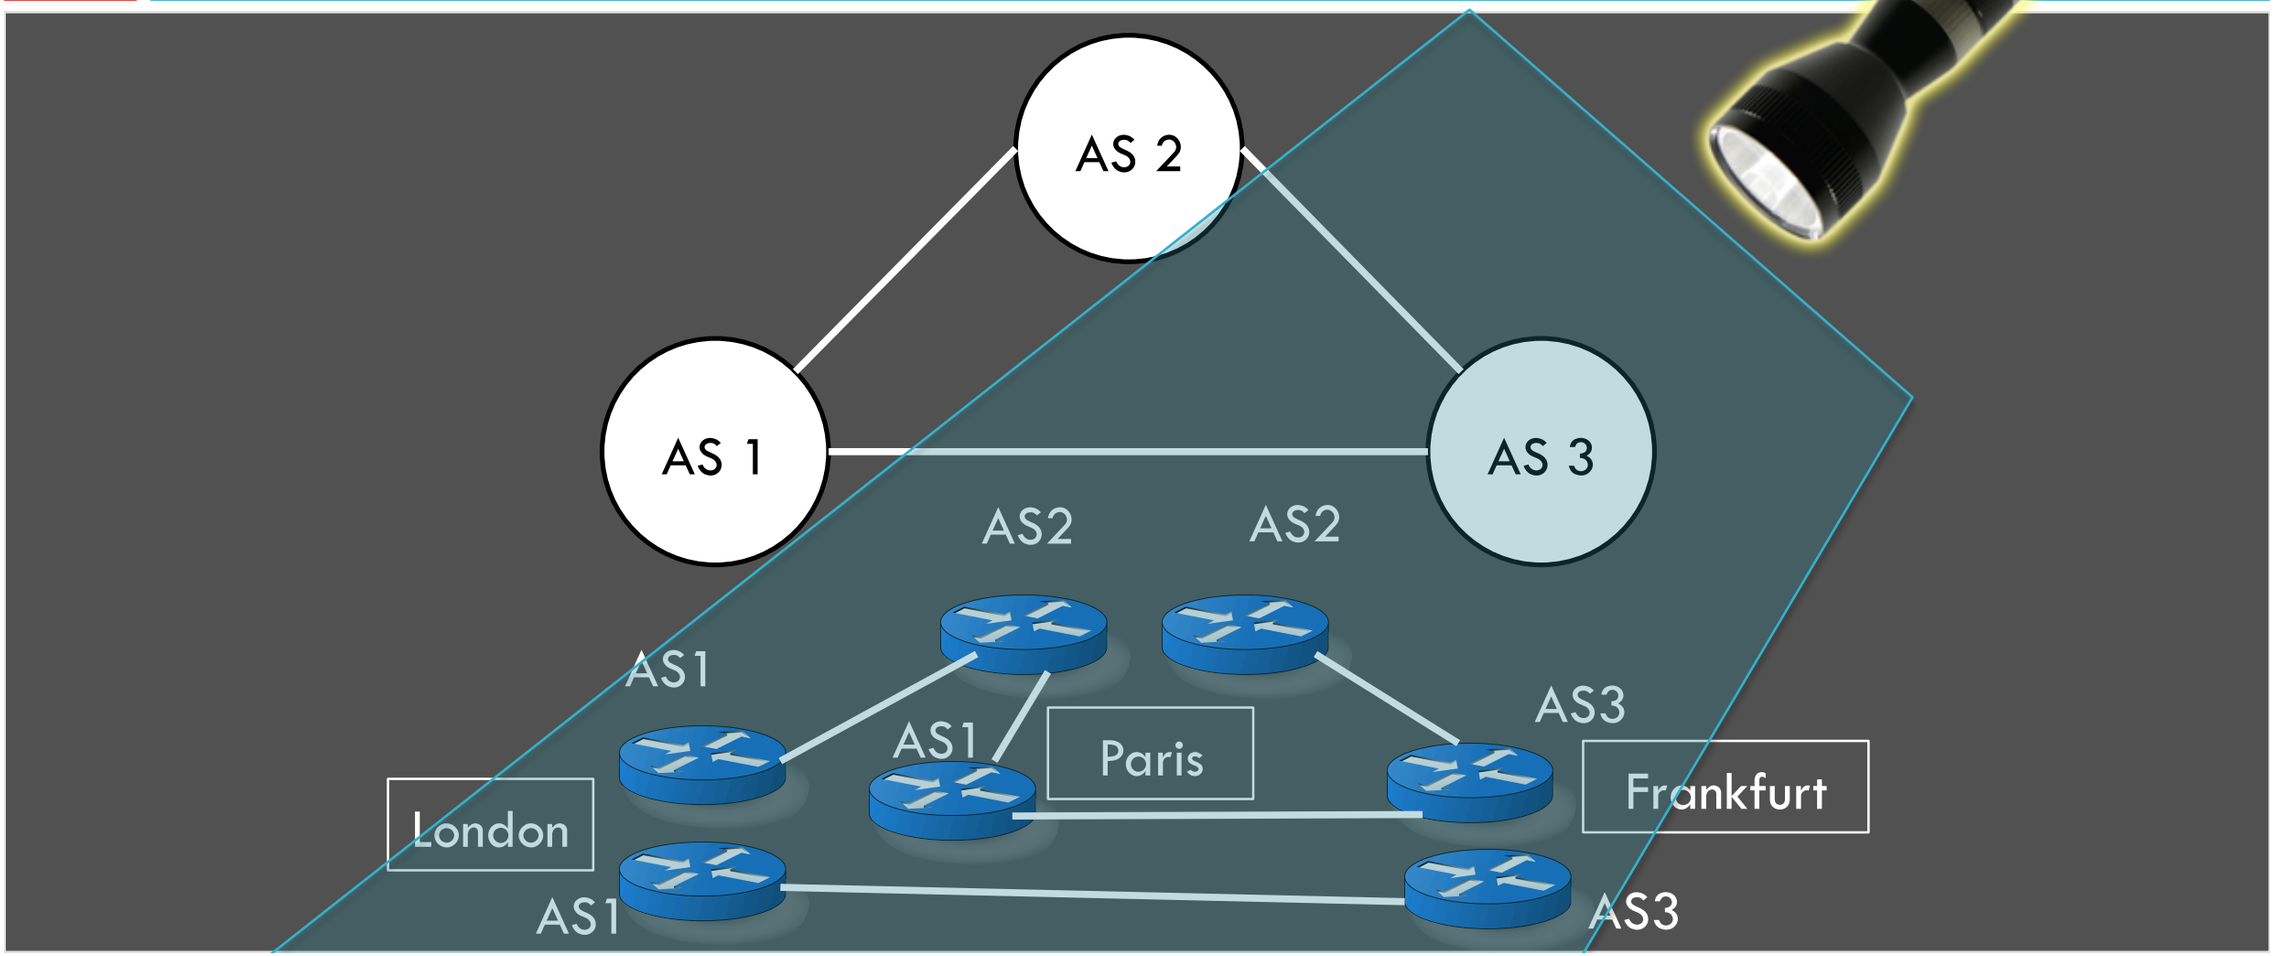

The building-level topology captures rich semantics of peering interconnections

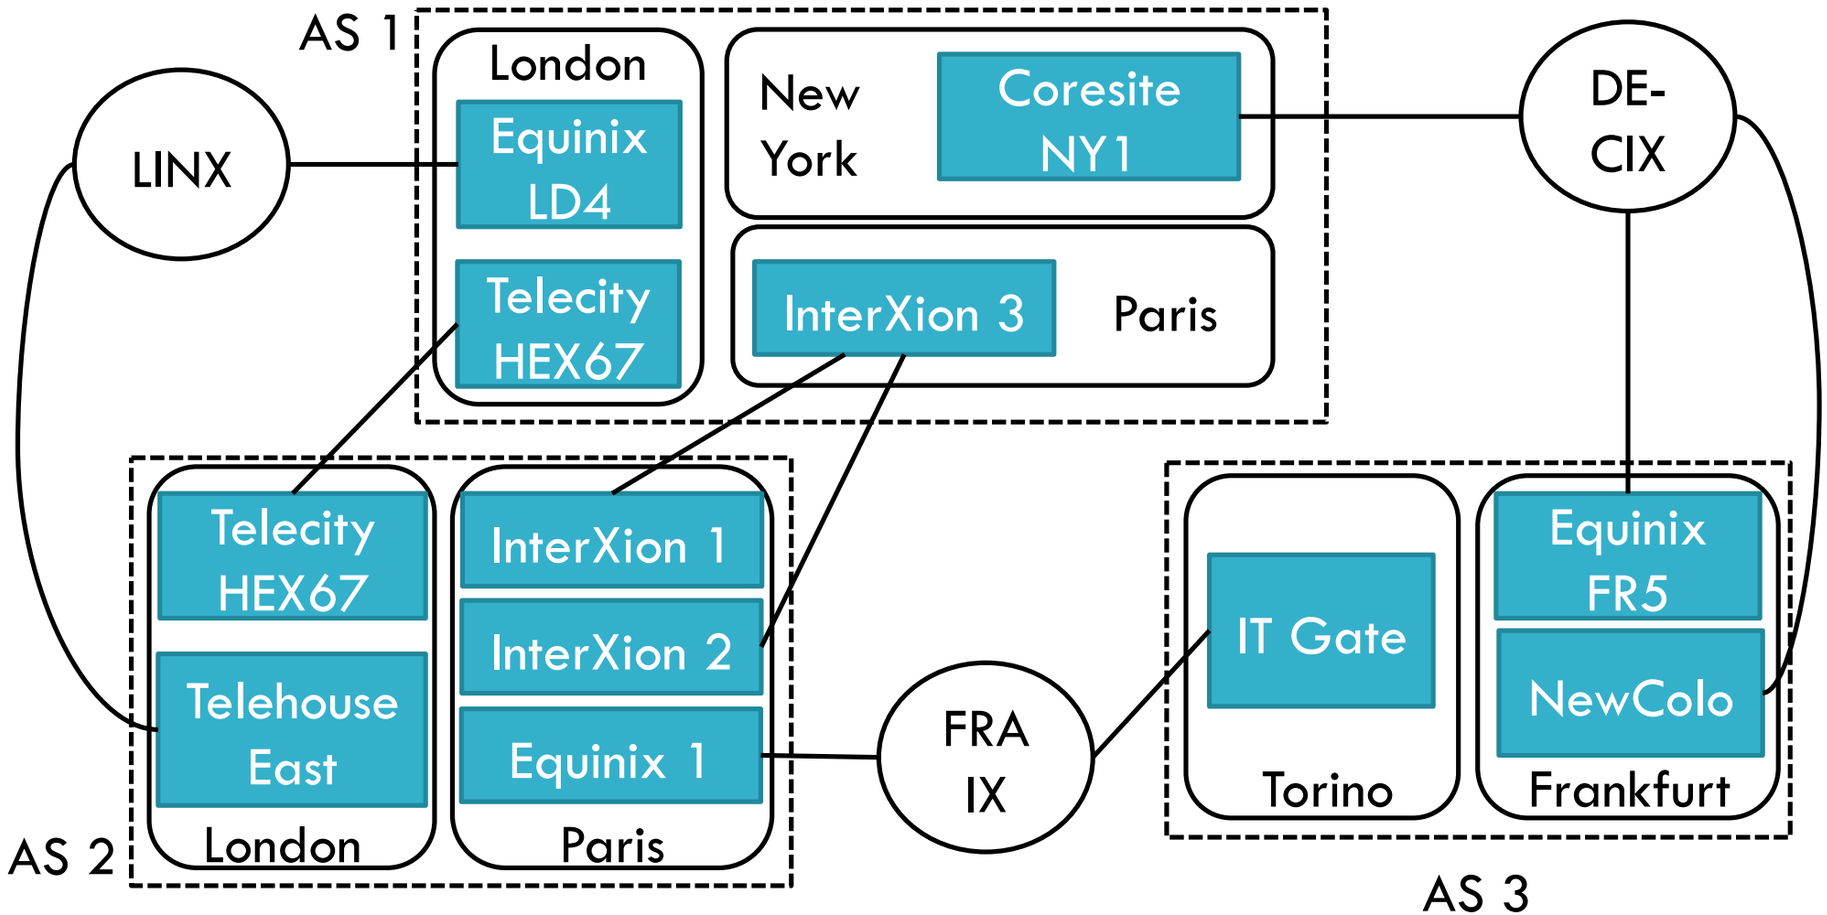

Motivation

- Increase traffic flow transparency
- Assessment of resilience of peering interconnections
- Diagnose congestion or DoS attacks
- Inform peering decisions
- Elucidate the role of colocation facilities, carrier hotels, and Internet exchange points (IXPs)

Challenges

- IP addresses are logical and region-independent
- BGP is an information hidden protocol; does not encode geographic information
- Existing methods are accurate for city-level granularity, not for finer granularities:
  - Delay-based
  - Hostname heuristics
  - Commercial IP Geolocation Databases

What buildings do we need to consider for locating peering interconnections?

- Interconnection facilities: special-purpose buildings used to co-locate routing equipment; routers have strict operational requirements

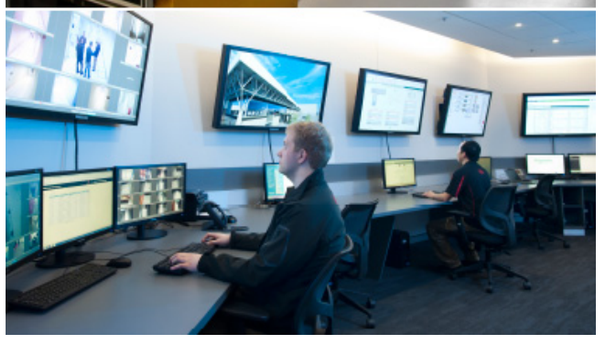

What buildings do we need to consider for locating peering interconnections?

- Interconnection facilities: special-purpose buildings used to co-locate routing equipment; routers have strict operational requirements


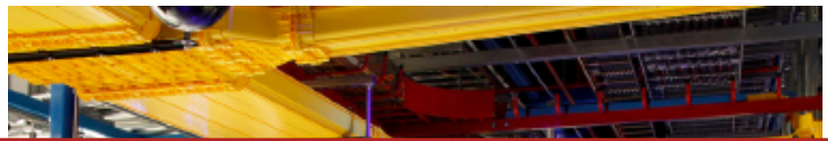

Key Intuition 1: To locate a peering interconnection, search the facilities where the peers are present

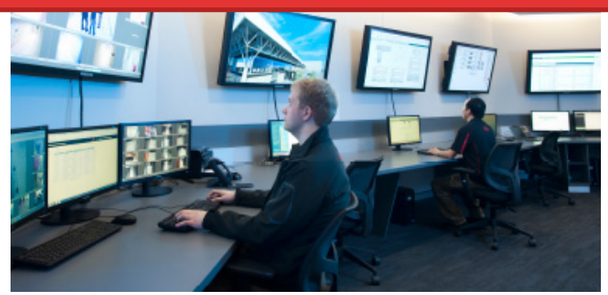


Construct a map of interconnection facilities

- Compile a list of interconnection facilities and their address
- Map ASes and IXPs to facilities
- Public data sources:
  - PeeringDB
  - AS/IXP websites

April 2015	
Facilities	1,694
ASes	3,303
AS-facility connections	13,206
IXPs	368
IXP-facility colocations	783

Facility data in PeeringDB are in many cases incomplete

- We compared the facility information between PDB and NOCs for 152 ASes:
  - 2,023 AS-to-facility connections in PDB
  - 1,424 AS-to-facility connections missing from PDB involving 61 ASes

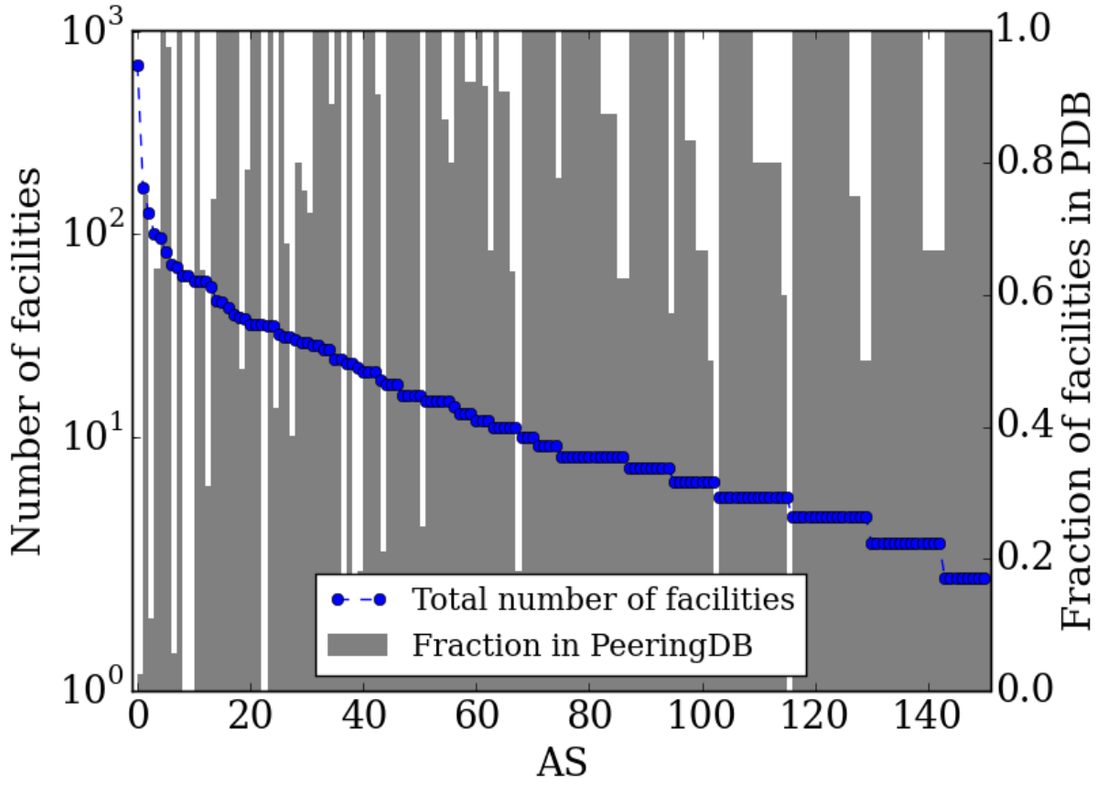

Interconnection facilities are concentrated in hub cities

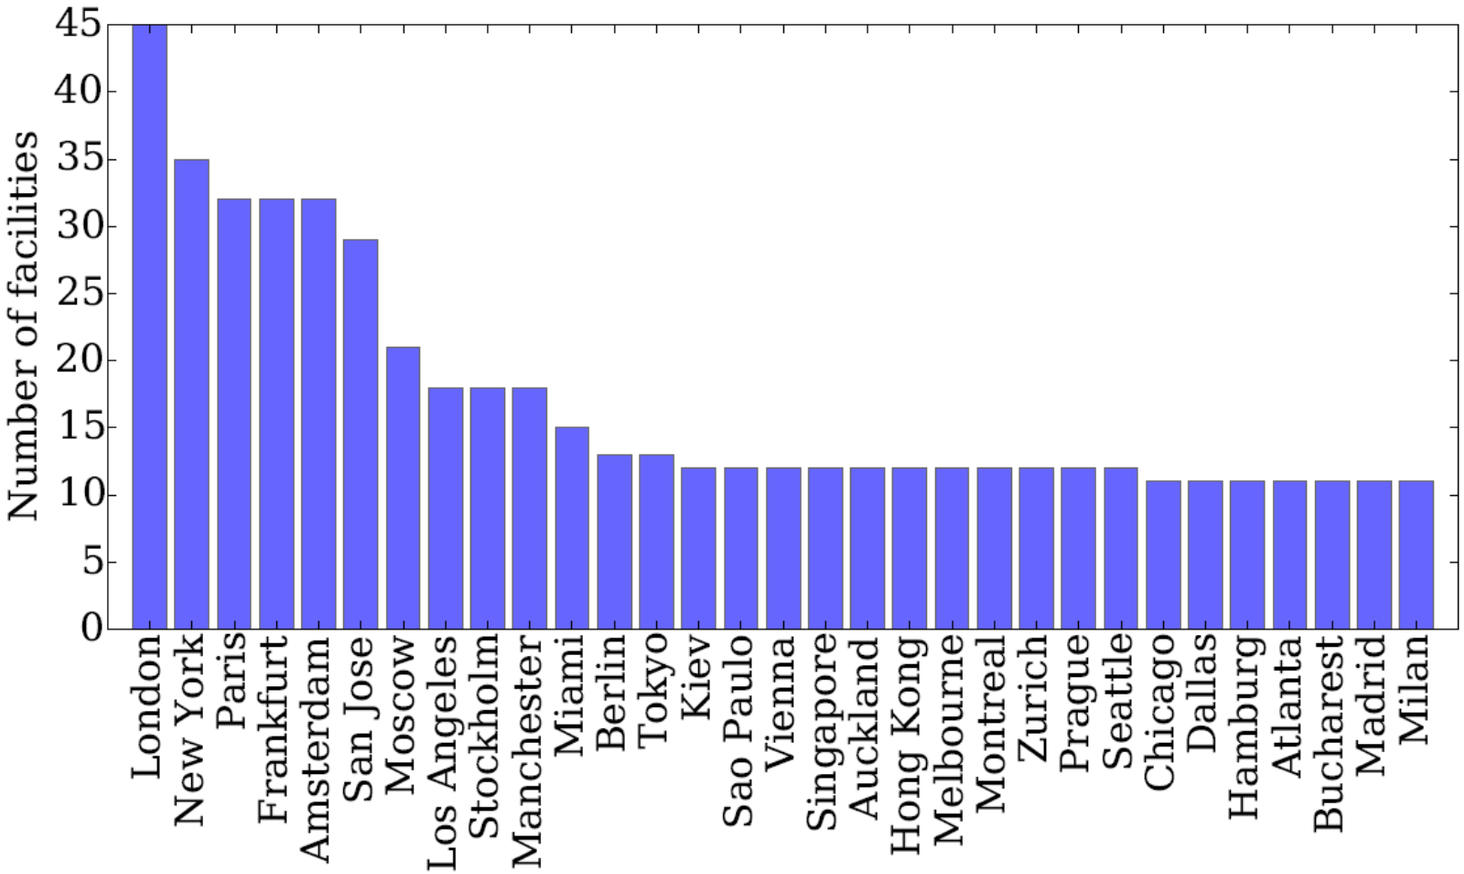

Increasing Complexity of peering interconnections

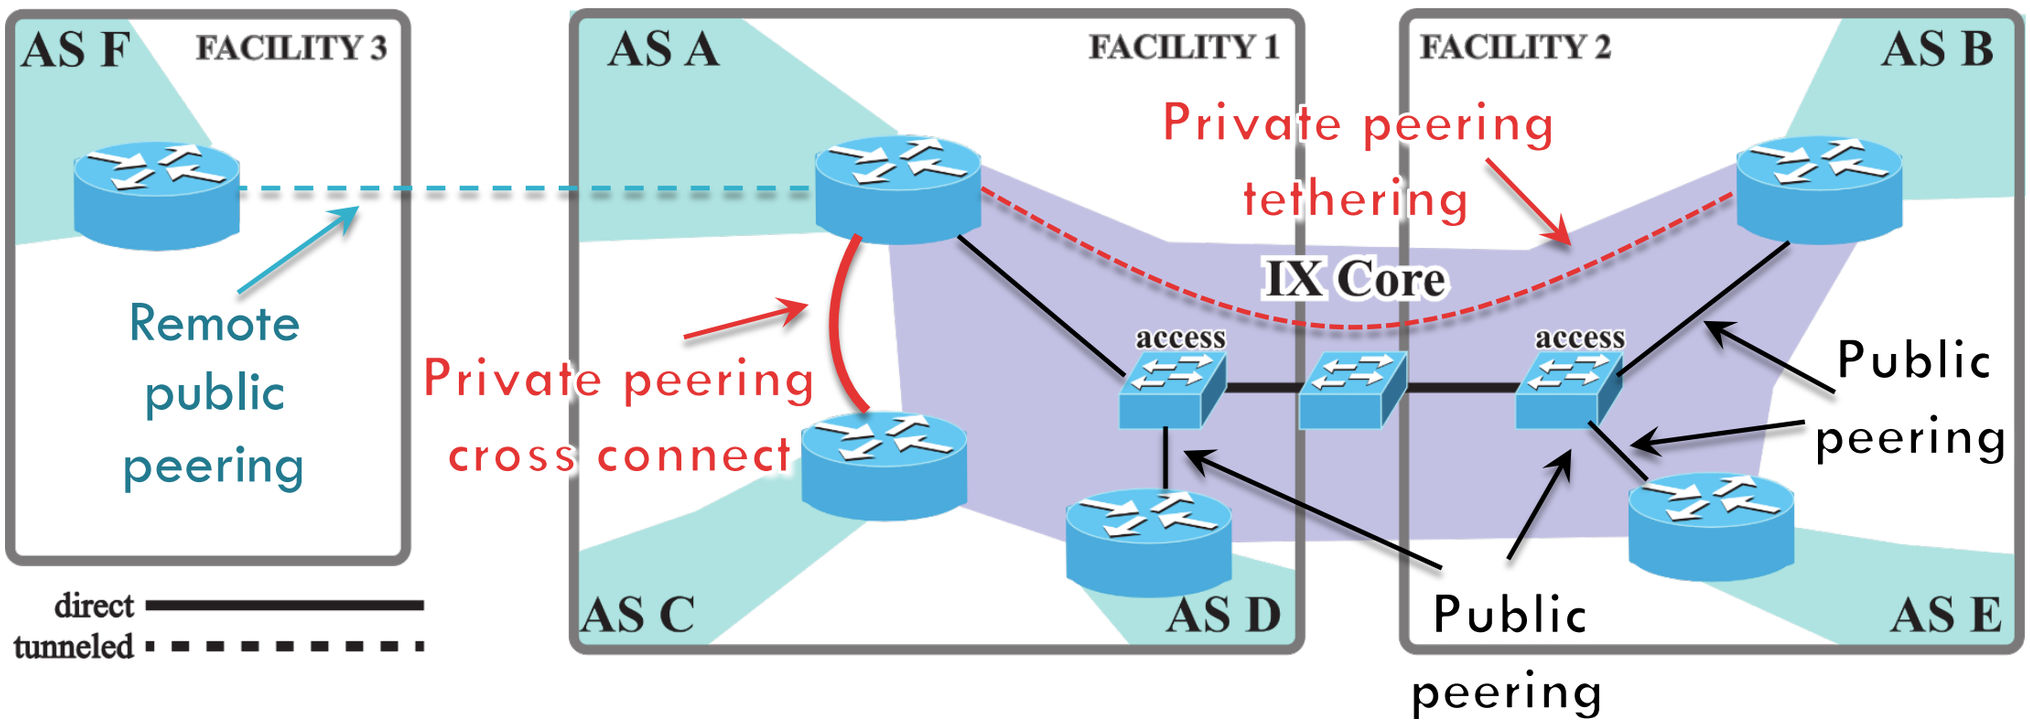

Increasing Complexity of peering interconnections

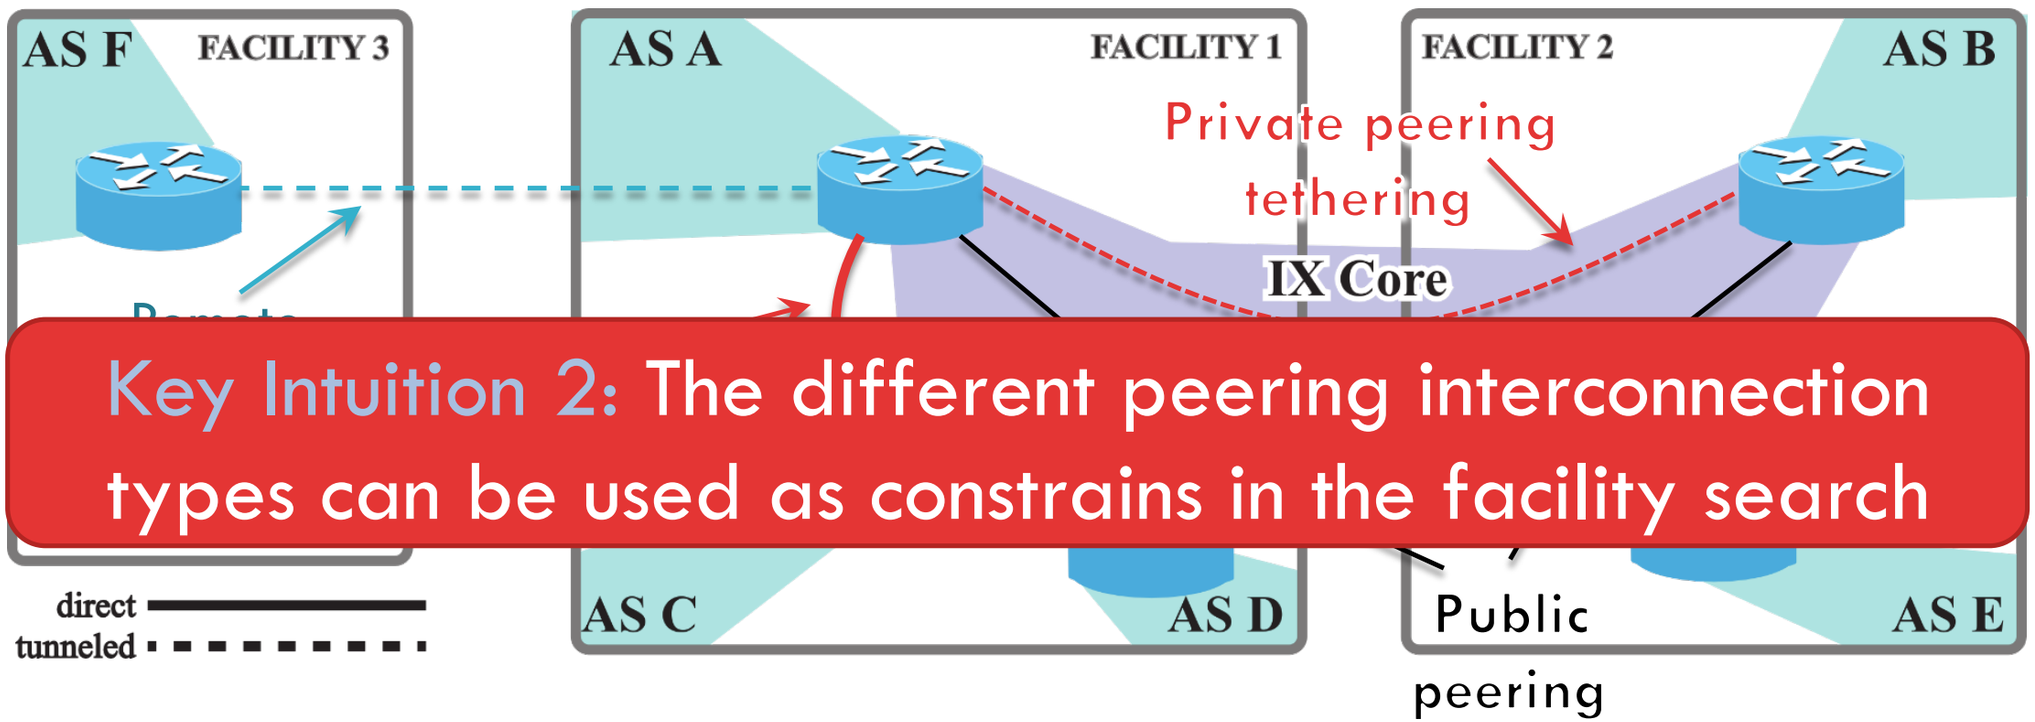

Moving Forward

Key Intuition 1: To locate a peering interconnection, search the facilities where the peers are present

Key Intuition 2: The different peering interconnection types can be used as constrains in the facility search

➔ Challenging Problem BUT Doable!
An algorithm is needed!

Algorithm: Constrained Facility Search (CFS)

For a target peering interconnection ASA - ASB:

- Step 1: Identify the type of peering interconnection
- Step 2: Initial facility search
- Step 3: Constrain facilities through alias resolution
- Step 4: Constrain facilities by repeating steps 1-3 with follow-up targeted traceroutes
- Step 5: Facility search in the reverse direction

Algorithm: Constrained Facility Search (CFS)

For a target peering interconnection ASA - ASB:

- Step 1: Identify the type of peering interconnection
- Step 2: Initial facility search
- Step 3: Constrain facilities through alias resolution
- Step 4: Constrain facilities by repeating steps 1-3 with follow-up targeted traceroutes
- Step 5: Facility search in the reverse direction

Identifying the peering type

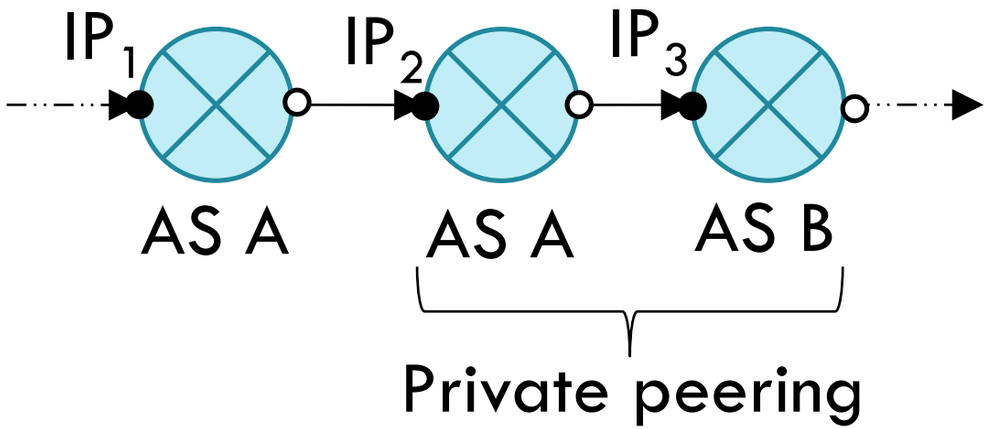

Facility search between the facilities of the peering Ases

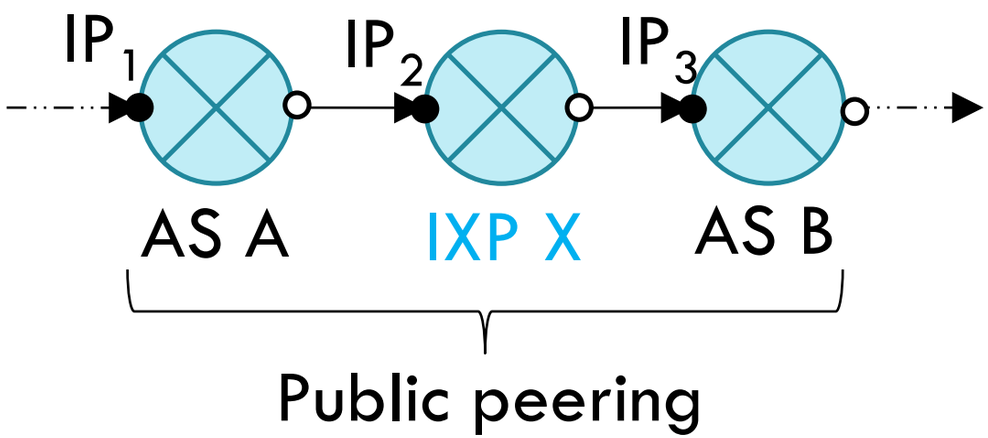

Facility search between the IXP and the peering ASes

Algorithm: Constrained Facility Search (CFS)

For a target peering interconnection ASA - ASB:

- Step 1: Identify the type of peering interconnection
- Step 2: Facility search
- Step 3: Constrain facilities through alias resolution
- Step 4: Constrain facilities by repeating steps 1-3 with follow-up targeted traceroute
- Step 5: Facility search in the reverse direction

Facility search: single common facility

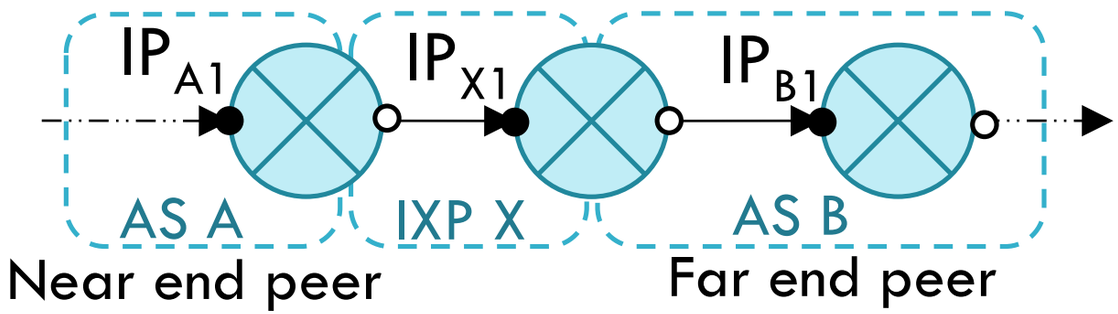

	Facilities	
AS A	F1	F2
IXP X	F4	F2

- The common facility is inferred as the location of the interface of the peer at the near end

Facility search: single common facility

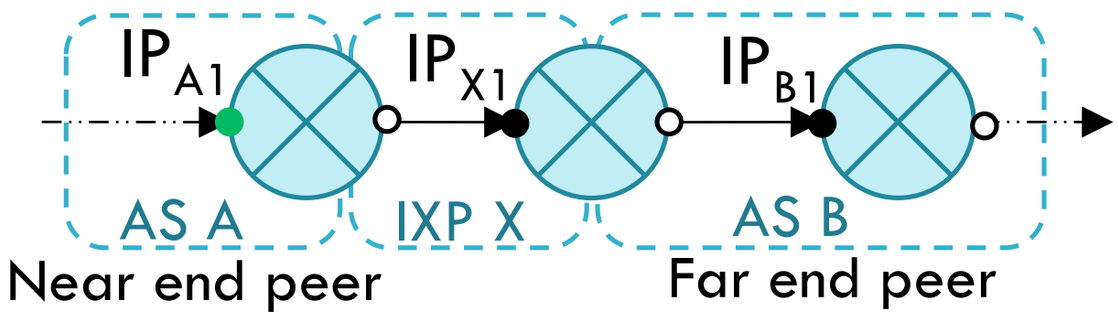

	Facilities	
AS A	F1	F2
IXP X	F4	F2

IP_{A1} facility

- The common facility is inferred as the location of the interface of the peer at the near end

Facility search: no common facility

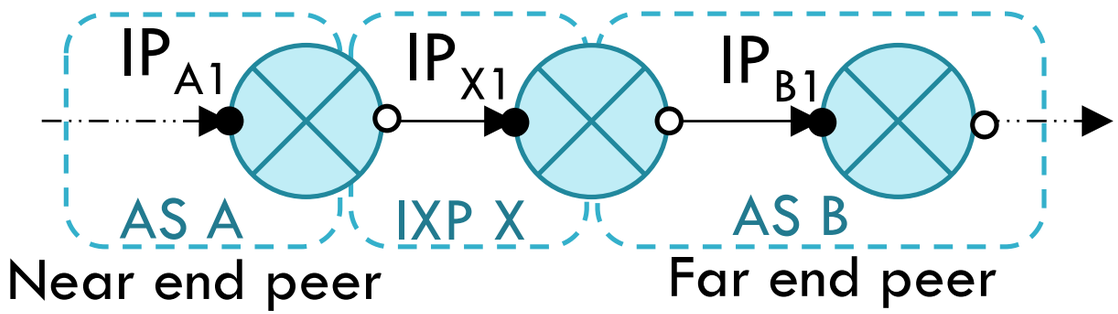

	Facilities	
AS A	F1	F2
IXP X	F4	F3

- No inference possible
  - Incomplete facility dataset or remote peering
  - Run algorithm in [Castro 2014] to detect remote peering
  - Run traceroutes changing the target peering links

Castro et al. "Remote Peering: More Peering without Internet Flattening." CoNEXT 2014

Facility search: multiple common facilities

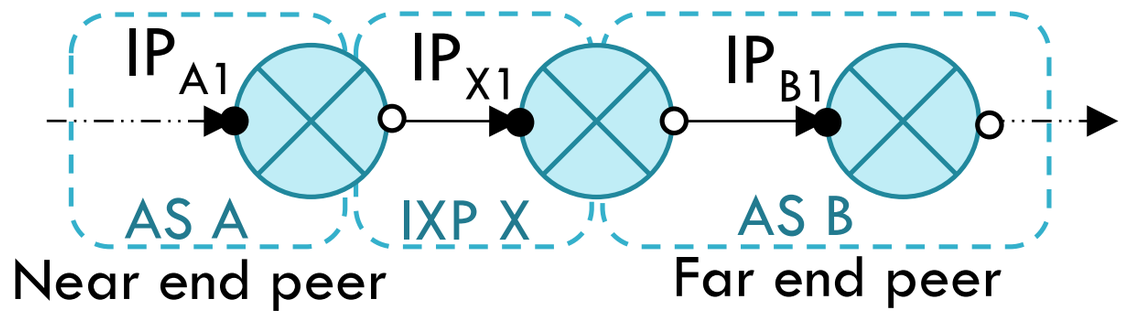

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

- Possible facilities are constrained but no inference yet

Facility search: multiple common facilities

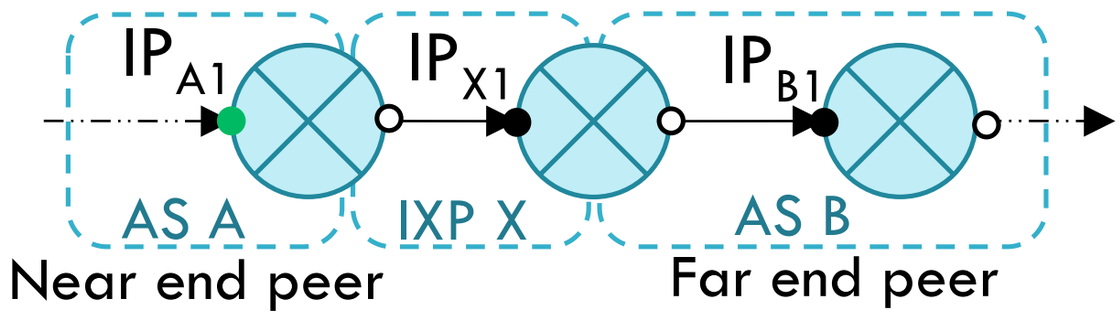

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

Possible IP_{A1} facilities

- Possible facilities are constrained but no inference yet

Algorithm: Constrained Facility Search (CFS)

For a target peering interconnection ASA - ASB:

- Step 1: Identify the type of peering interconnection
- Step 2: Initial facility search
- Step 3: Derive constrains through alias resolution
- Step 4: Constrain facilities by repeating steps 1-3 with follow-up targeted traceroutes
- Step 5: Facility search in the reverse direction

Derive constrains through alias resolution

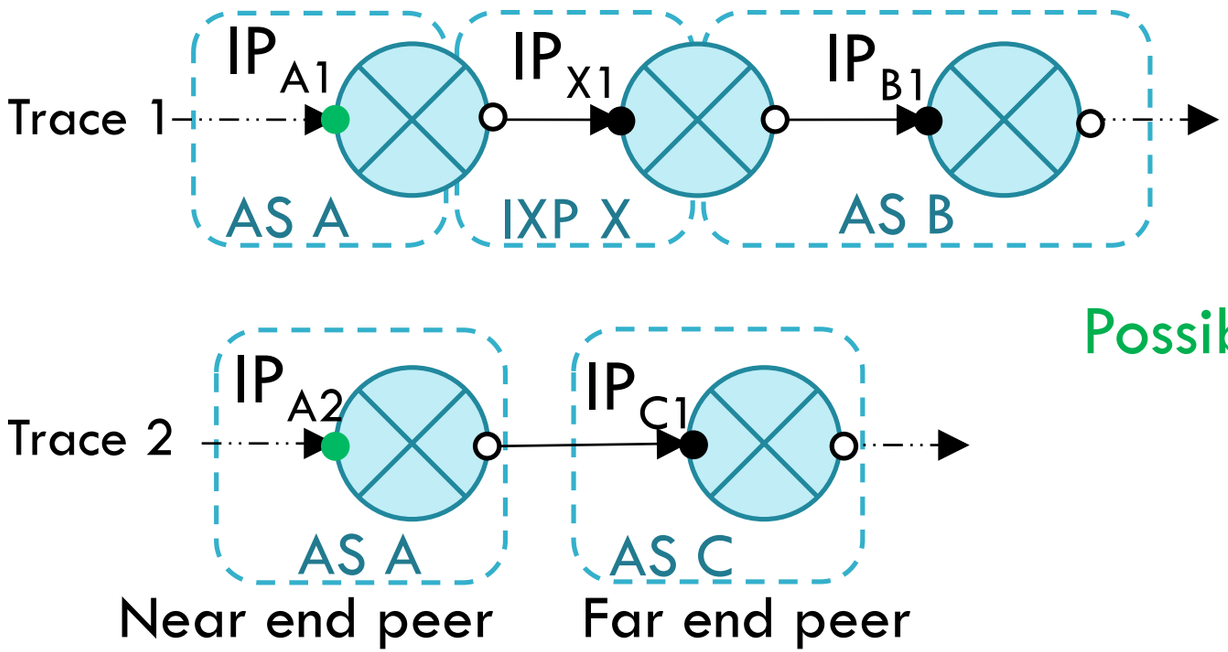

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

Possible IP_{A1} facilities

	Facilities		
AS A	F1	F2	F5
AS C	F1	F2	F6

Possible IP_{A2} facilities

- Parse additional traceroutes containing peering interconnections of the peer at the near end

Derive constrains through alias resolution

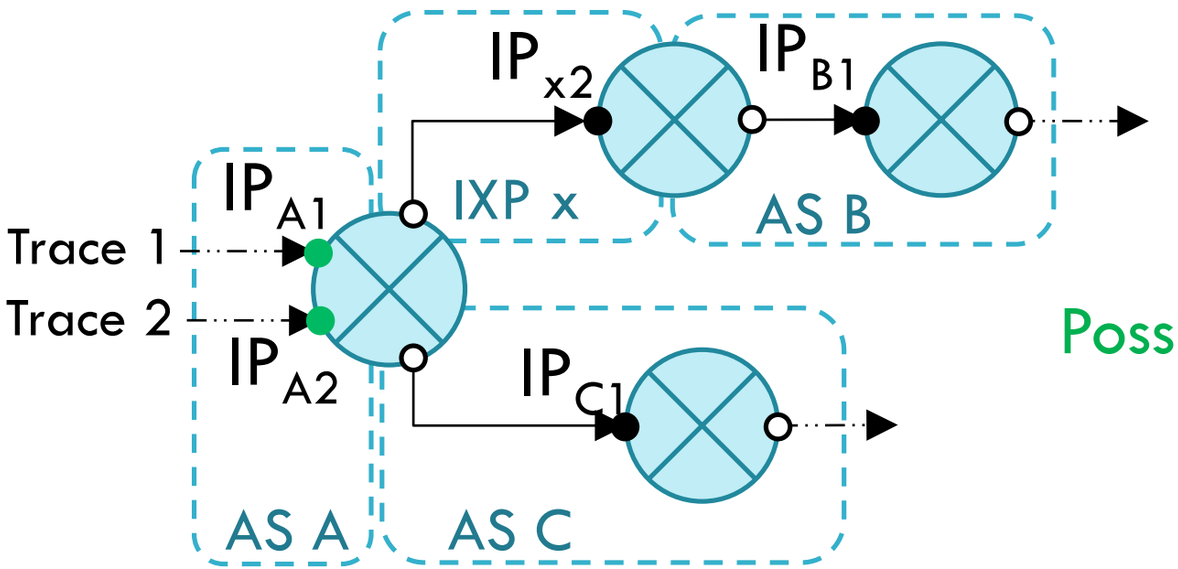

	Facilities		
AS A	F1	F2	F5
IXP x	F4	F2	F5

Possible IP_{A1} facilities

	Facilities		
AS A	F1	F2	F5
AS C	F1	F2	F6

Possible IP_{A2} facilities

- De-alias interfaces of AS A (IP_{A1}, IP_{A2})

Derive constrains through alias resolution

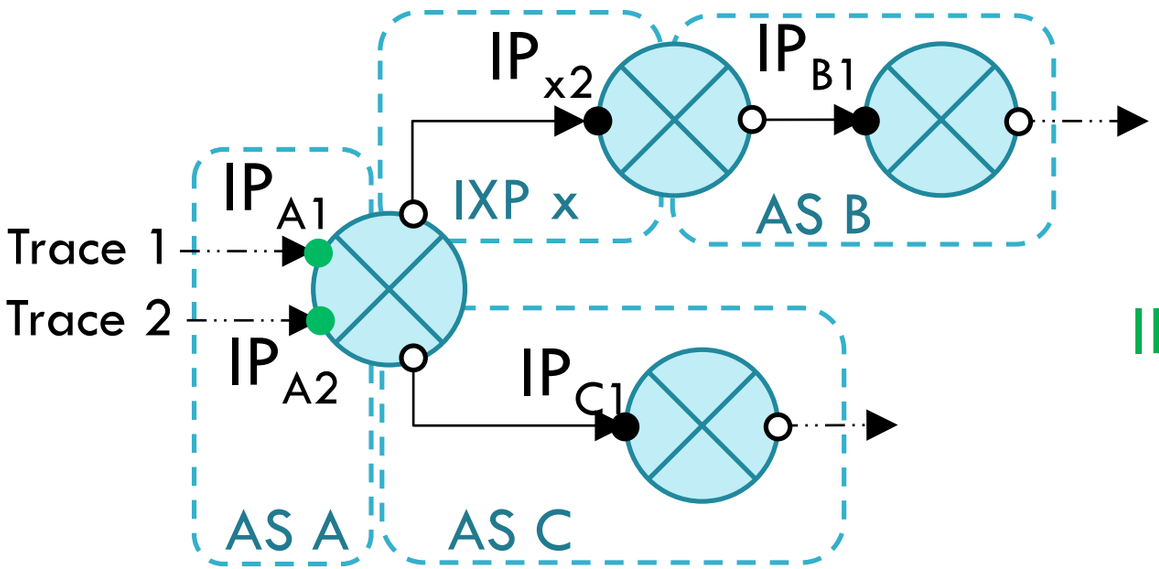

	Facilities		
AS A	F1	F2	F5
IXP x	F4	F2	F5

IP_{A1} & IP_{A2} facility

	Facilities		
AS A	F1	F2	F5
AS C	F1	F2	F3

- If two interfaces belong to the same router, find the intersection of their possible facilities

Derive constrains through alias resolution

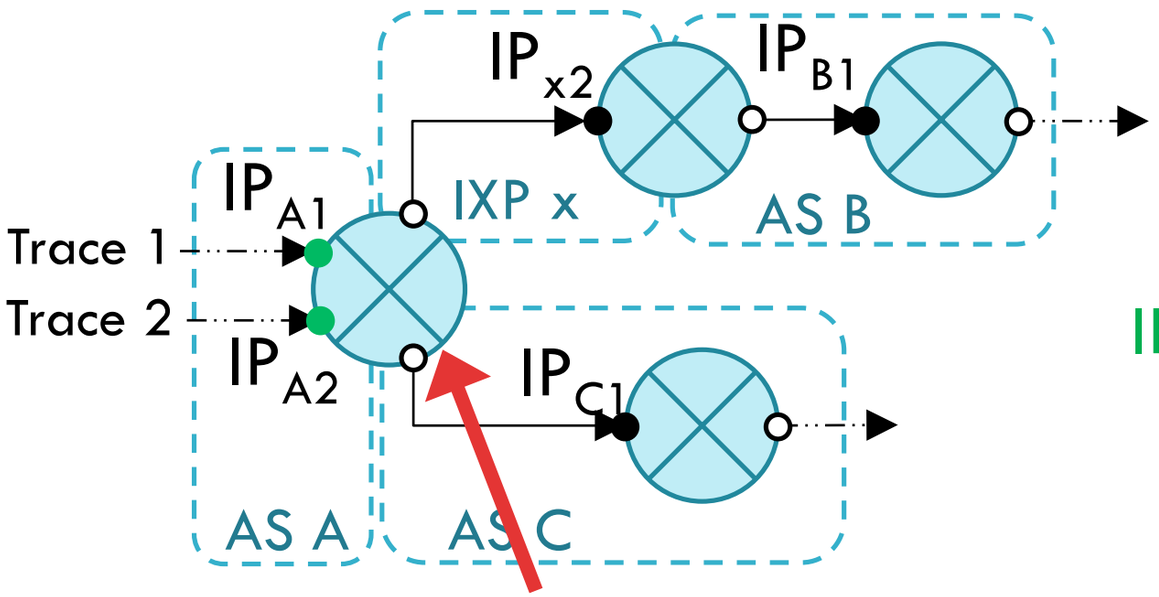

Multi-purpose router

- Used to establish both private and public peering: 40% of the routers have multi role in our study
- 12% of routers used for public peering with >1 IXP

	Facilities		
AS A	F1	F2	F5
IXP x	F4	F2	F5

IP_{A1} & IP_{A2} facility

	Facilities		
AS A	F1	F2	F5
AS C	F1	F2	F3

Algorithm: Constrained Facility Search (CFS)

For a target peering interconnection ASA - ASB:

- Step 1: Identify the type of peering interconnection
- Step 2: Initial facility search
- Step 3: Constrain facilities through alias resolution
- Step 4: Constrain facilities by repeating steps 1-3 with follow-up targeted traceroutes
- Step 5: Facility search in the reverse direction

Algorithm: Constrained Facility Search (CFS)

For a target peering interconnection ASA - ASB:

- Step 1: Identify the type of peering interconnection
- Step 2: Initial facility search
- Step 3: Constrain facilities through alias resolution
- Step 4: Constrain facilities by repeating steps 1-3 with follow-up targeted traceroutes
- Step 5: Facility search in the reverse direction

Evaluation

- Targeted the peerings of 5 CDNs and 5 Tier-1 ASes:
  - Google (AS15169), Yahoo (AS10310), Akamai (AS20940), Limelight (AS22822), Cloudflare (AS13335)
  - NTT (AS2914), Cogent (AS174), Deutsche Telekom (AS3320), Level 3 (AS3356), Telia (AS1299)
  - Queried one active IP per prefix for each of their peers

Collecting traceroute paths

- Combine various traceroute platforms to maximize coverage:
  - Active: RIPE Atlas, Looking Glasses (LGs)
  - Archived: CAIDA Ark, iPlane

	RIPE Atlas	LGs	iPlane	Ark	Total Unique
VPs	6,385	1,877	147	107	8,517
ASNs	2,410	438	117	71	2,638
Countries	160	79	35	41	170

CFS inferred the facility for 70% of collected peering interfaces

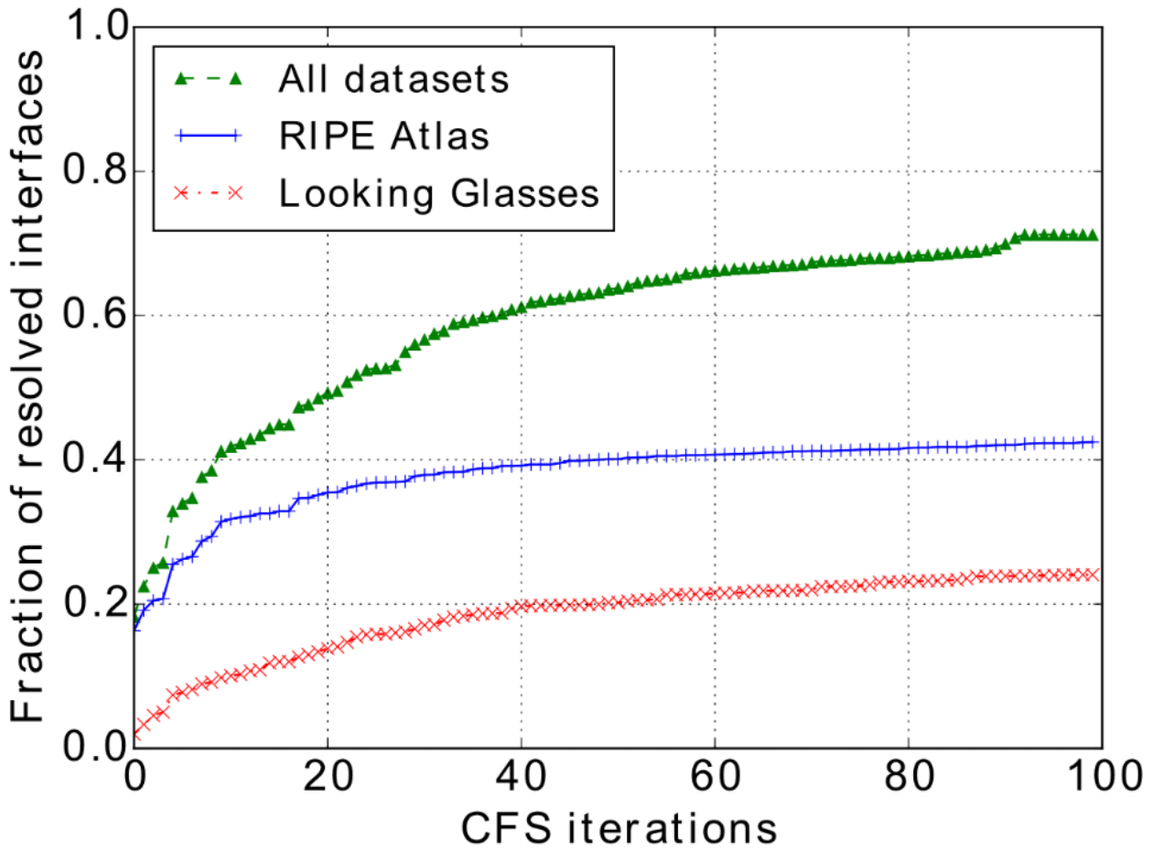

Diverse peering strategies between CDNs and Tier-1 ASes

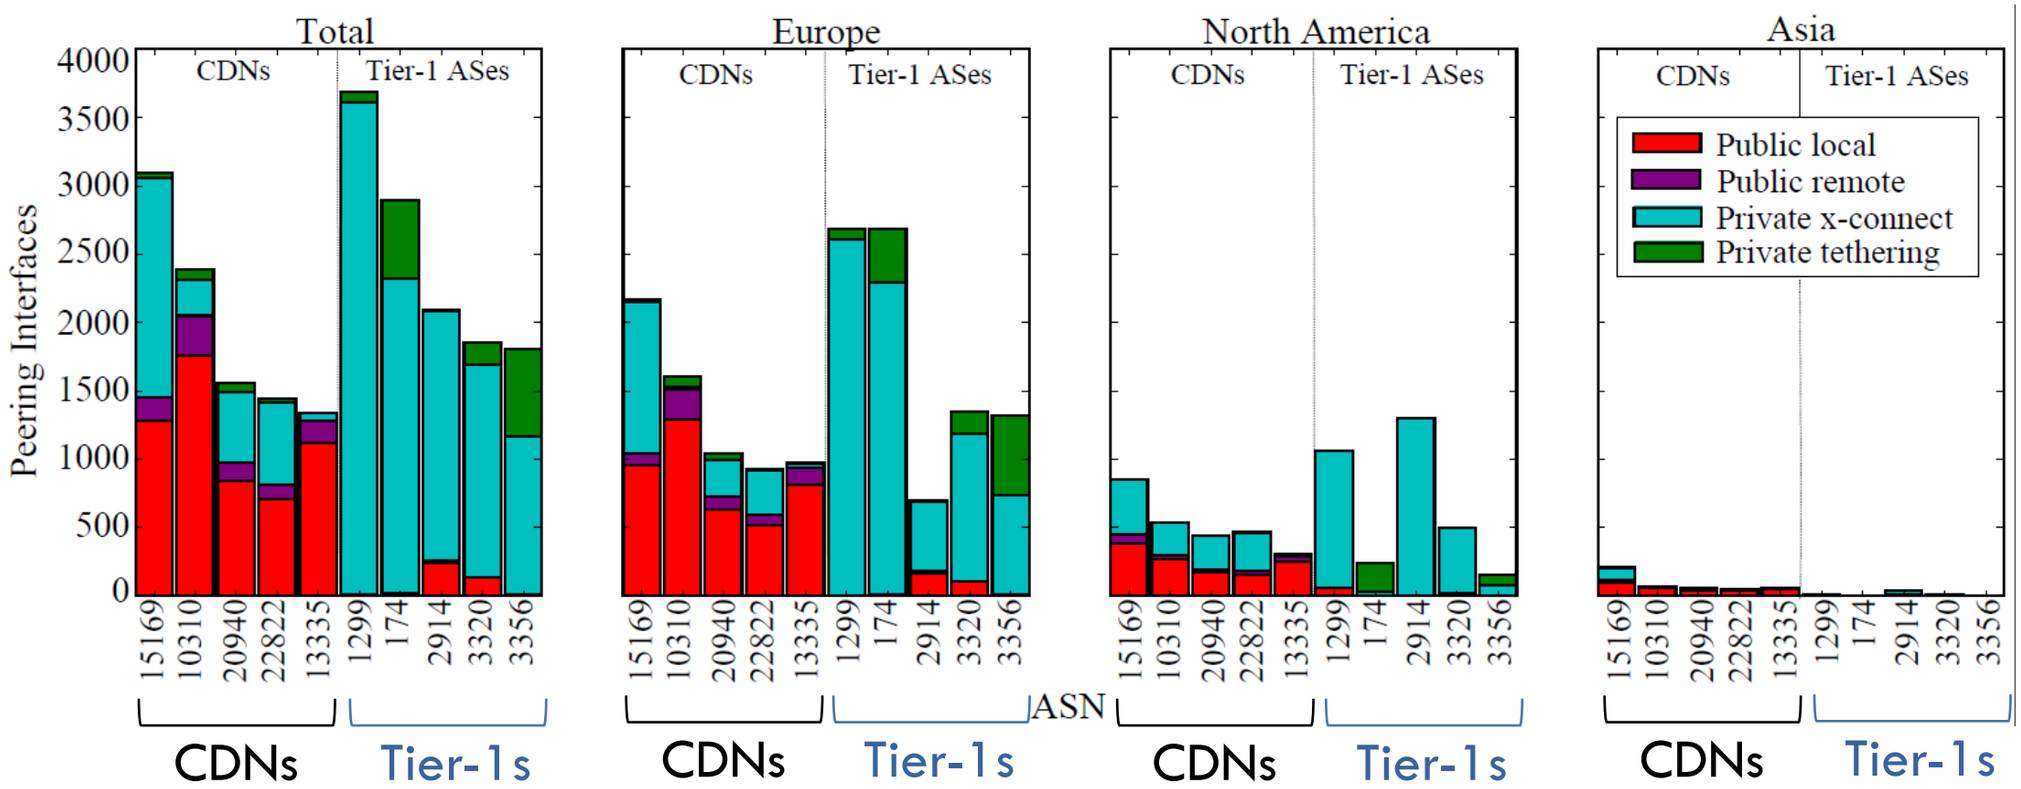

10% of the inferences validated to 90% correctness

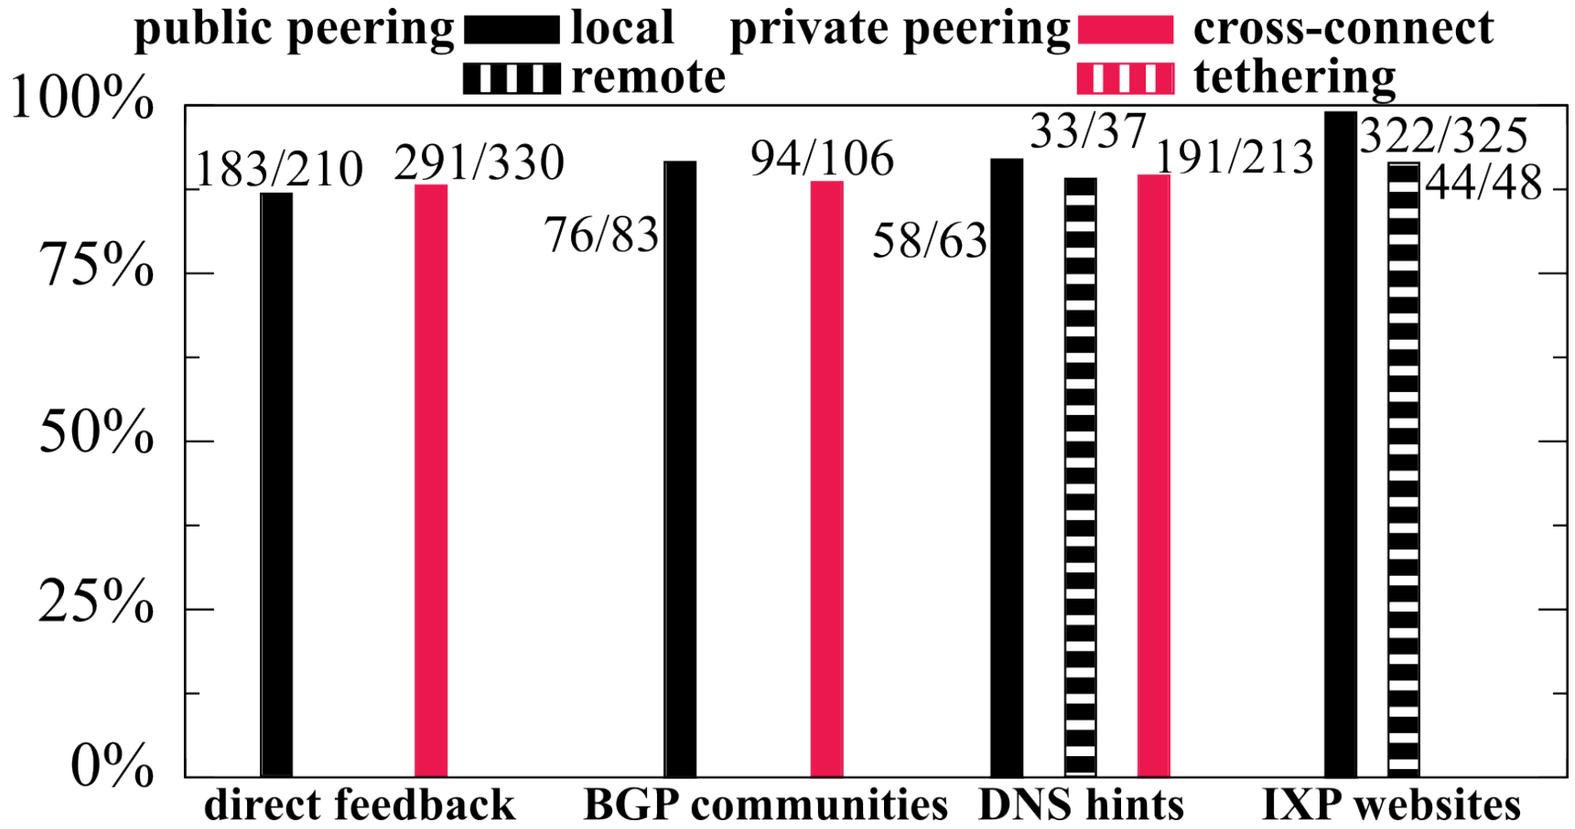

Conclusions

- Constrained Facility Search (CFS) maps peering interconnections to facilities based on public data:
  - Interconnection facility maps
  - Traceroute paths
- Evaluated CFS for 5 large CDNs and Tier-1 Ases
  - Pinpoint 70% of collected IP interfaces
  - Validated 10% of inferences to ~90% correctness

Ongoing and future work

- Extend the facility dataset
  - Collaborate with the operational community
  - Utilize third-party datasets e.g. UW Internet Atlas[1]
- Combine geolocation methods to further constrain facilities in unresolved cases
- Integrate CFS with CAIDA's Ark and Sibyl[2]

[1] SIGCOMM'15 also at http://internetatlas.org/

[2] NSDI'16 [to appear]

Thank you!

Back-up Slides

Additional results

ASes and IXPs are present at multiple facilities

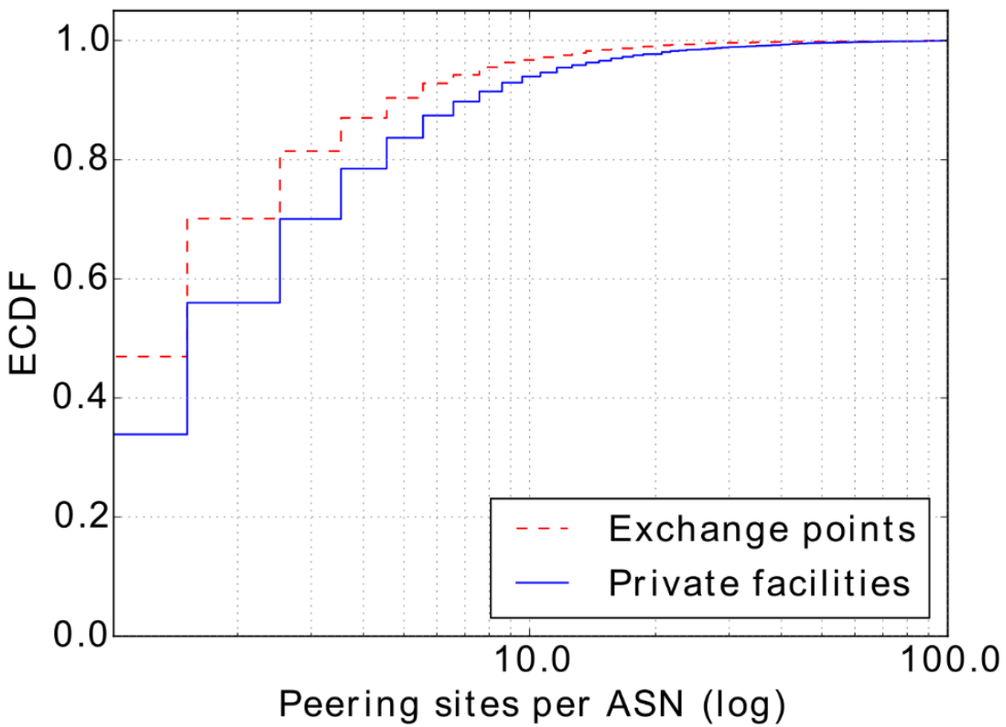

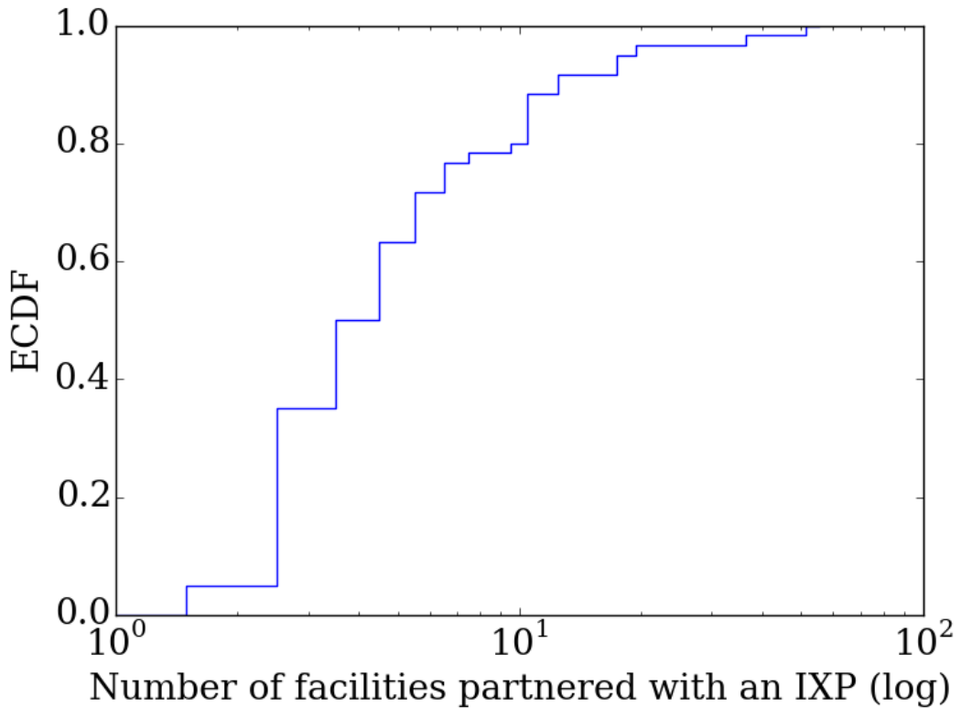

Majority of interconnection facilities are located in Europe and North America

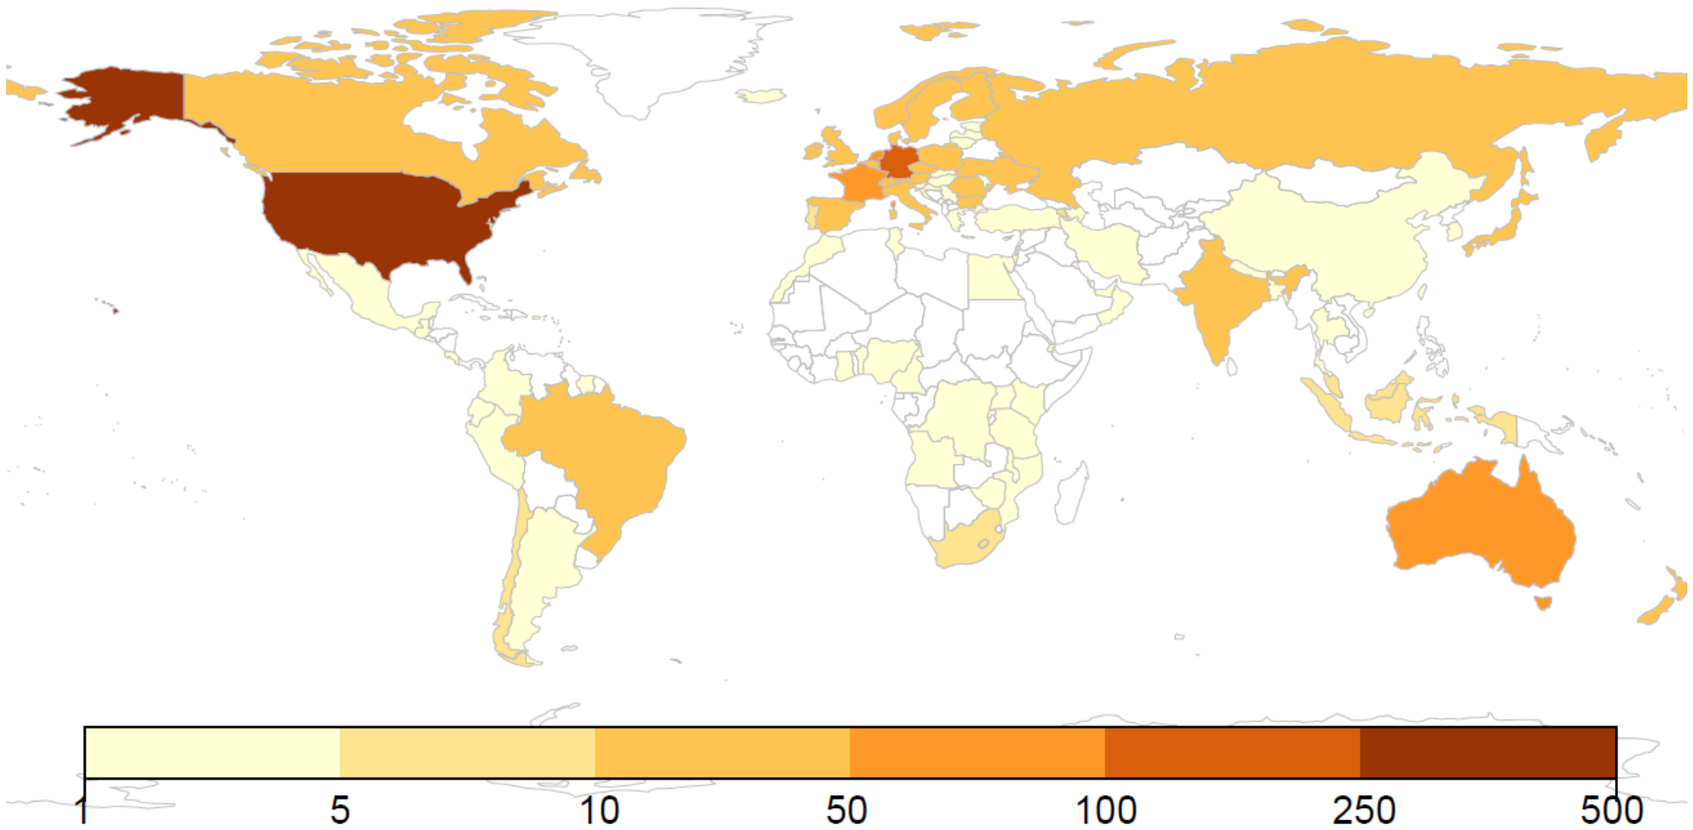

April 2015	
Europe	860
North America	503
Asia	143
Oceania	84
South America	73
Africa	31

Missing facility data affect the completeness of CFS inferences

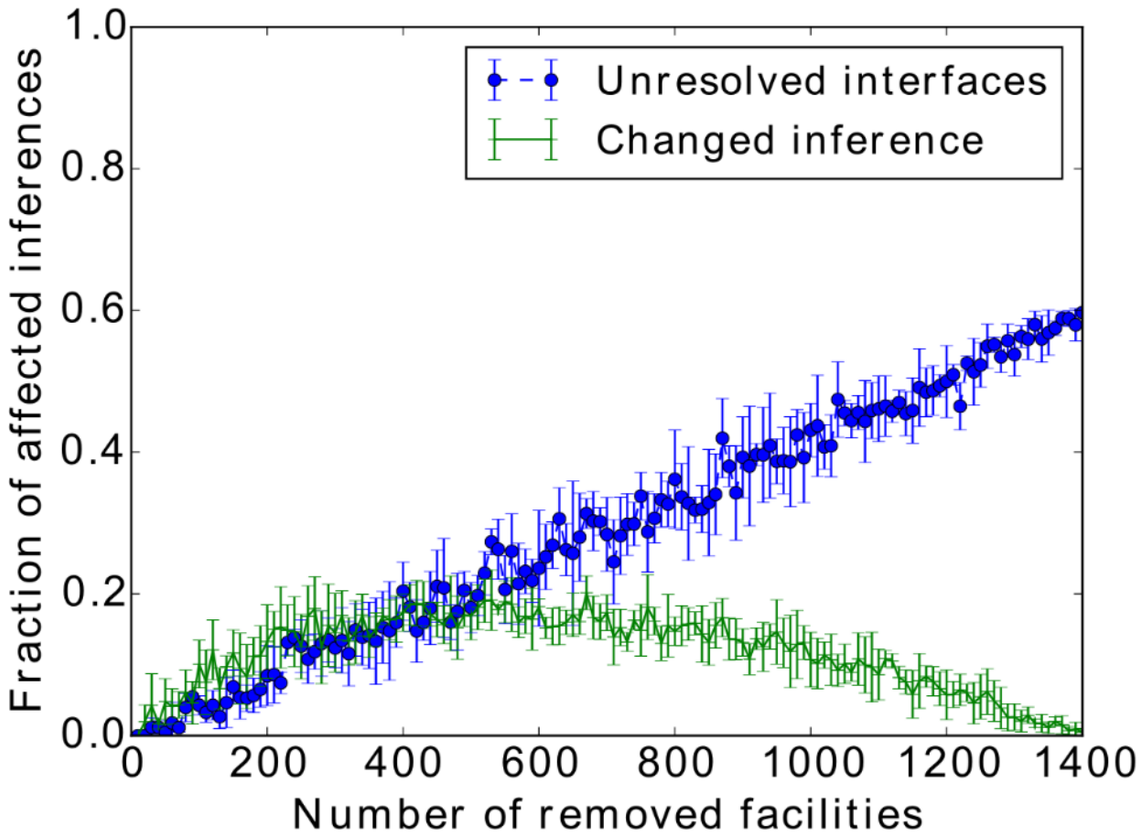

Details on Methodology

Facility inference for the far-end peer

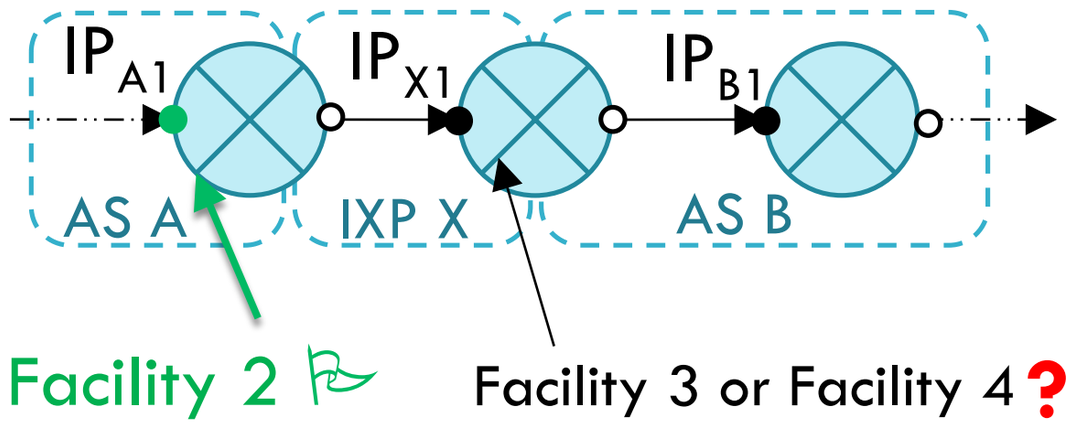

- Facility search for the peer at the far-end may not converge to a single facility
- Last resort: switch proximity heuristic

Follow-up CFS iterations

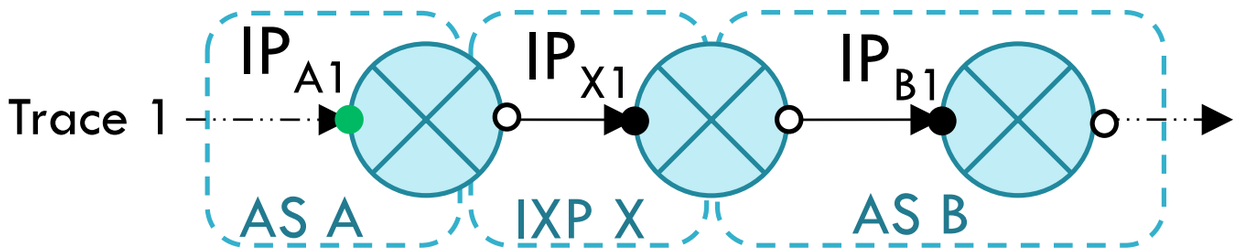

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

- If CFS has not converged to a single facility:
  - Execute a new round of traceroutes with different set of targets
  - Repeat steps 1-3 (a CFS iteration)
- 'Clever' selection of the new traceroute targets can help CFS to narrow down the facility search

Traceroute target selection

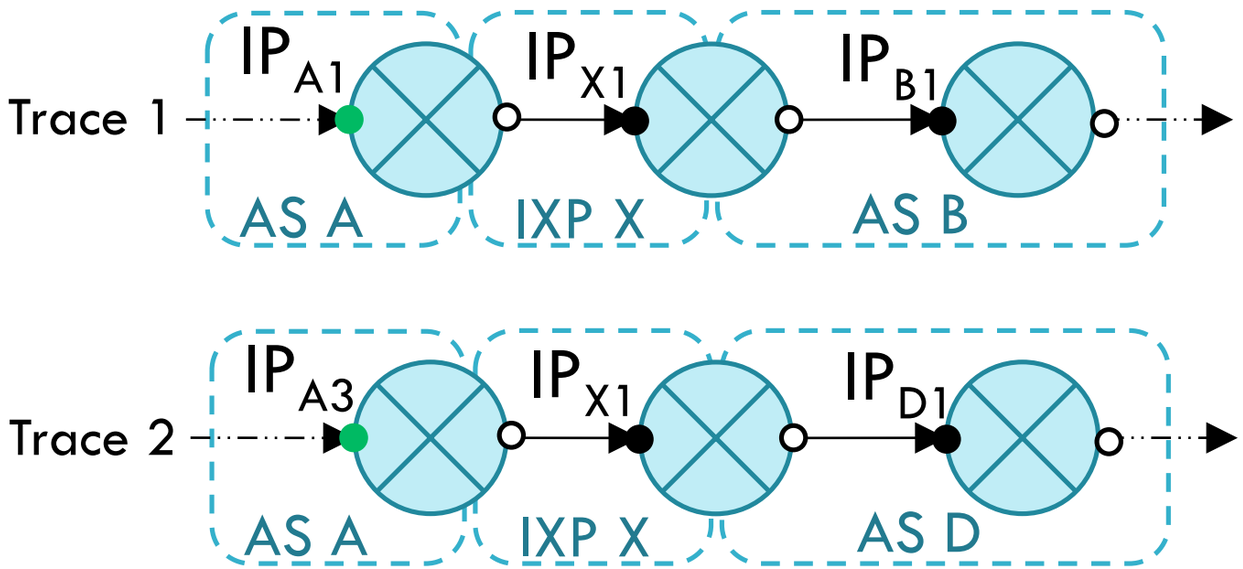

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

Traceroute target selection

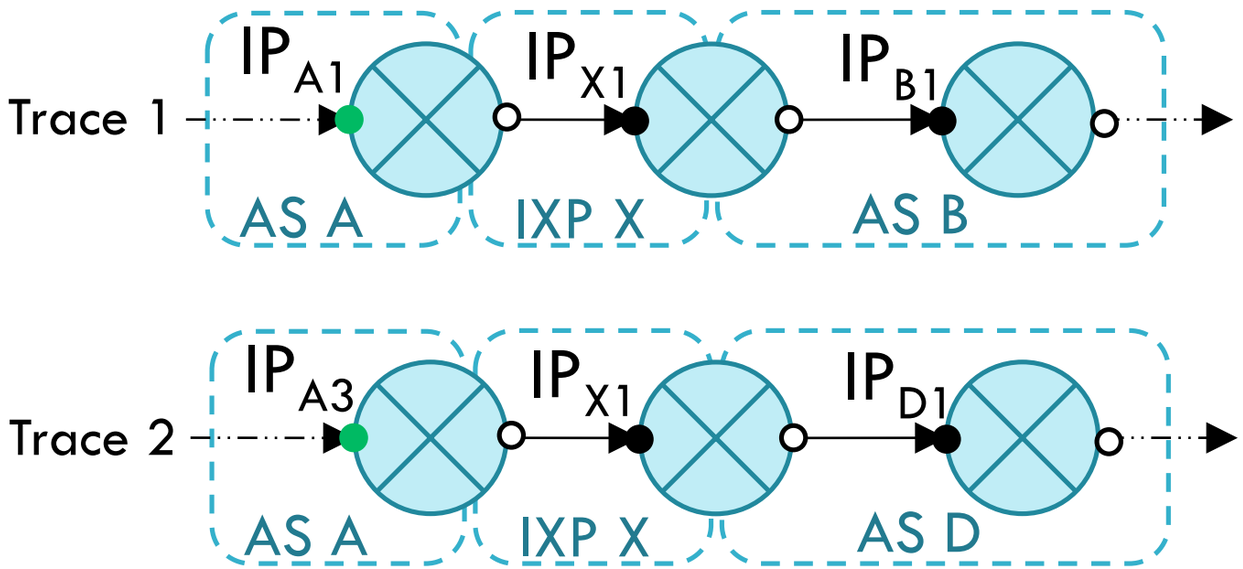

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

Targeting public peerings over the same IXP offers no additional constrains because CFS still compares the same sets of facilities

Traceroute target selection

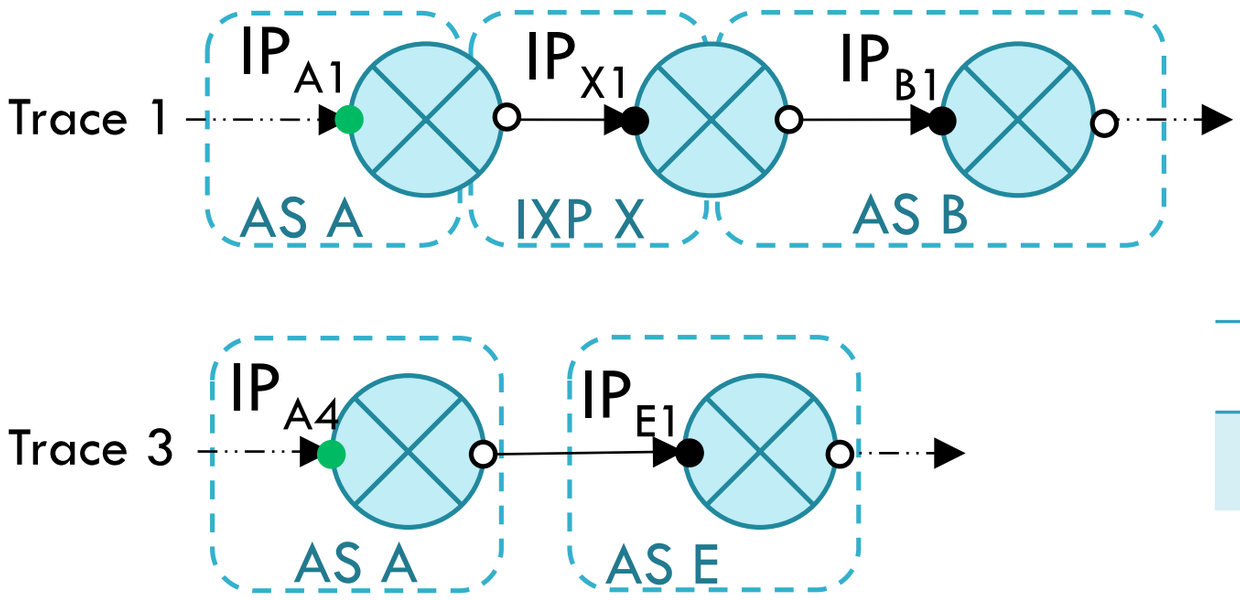

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

	Facilities		
AS A		F1	F2	F5
AS E	F9	F1	F2	F5

Traceroute target selection

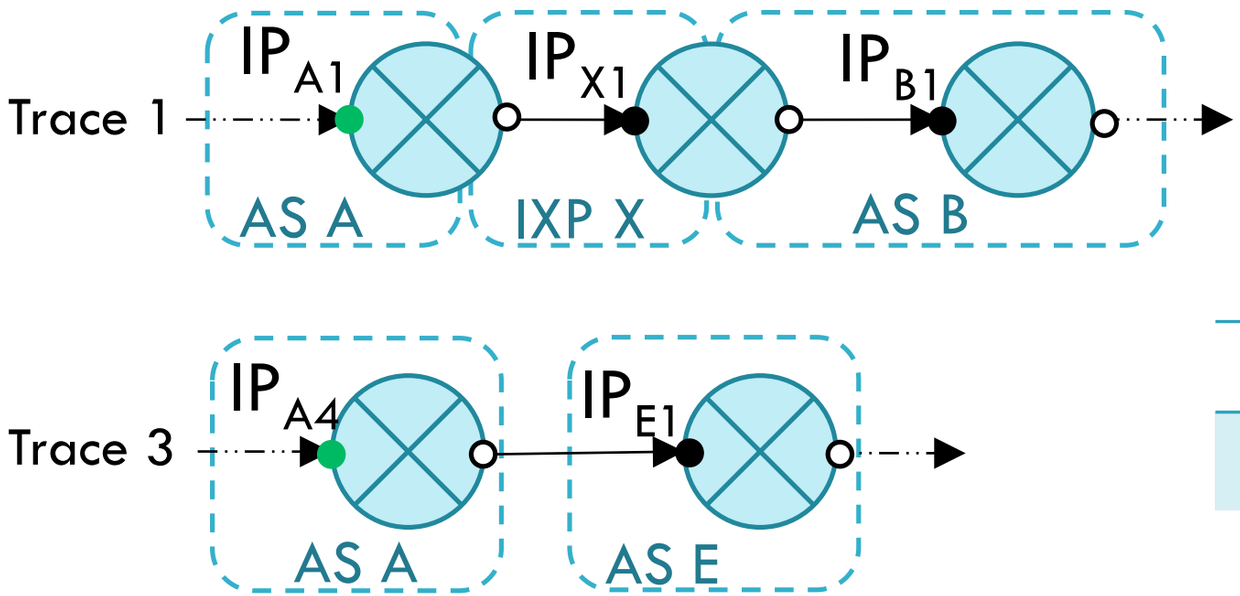

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

	Facilities		
AS A	F1	F2	F5
AS E F9	F1	F2	F5

Targeting private peers or IXPs with presence in all the possible facilities for IP_{A1} does not offer additional constrains

Traceroute target selection

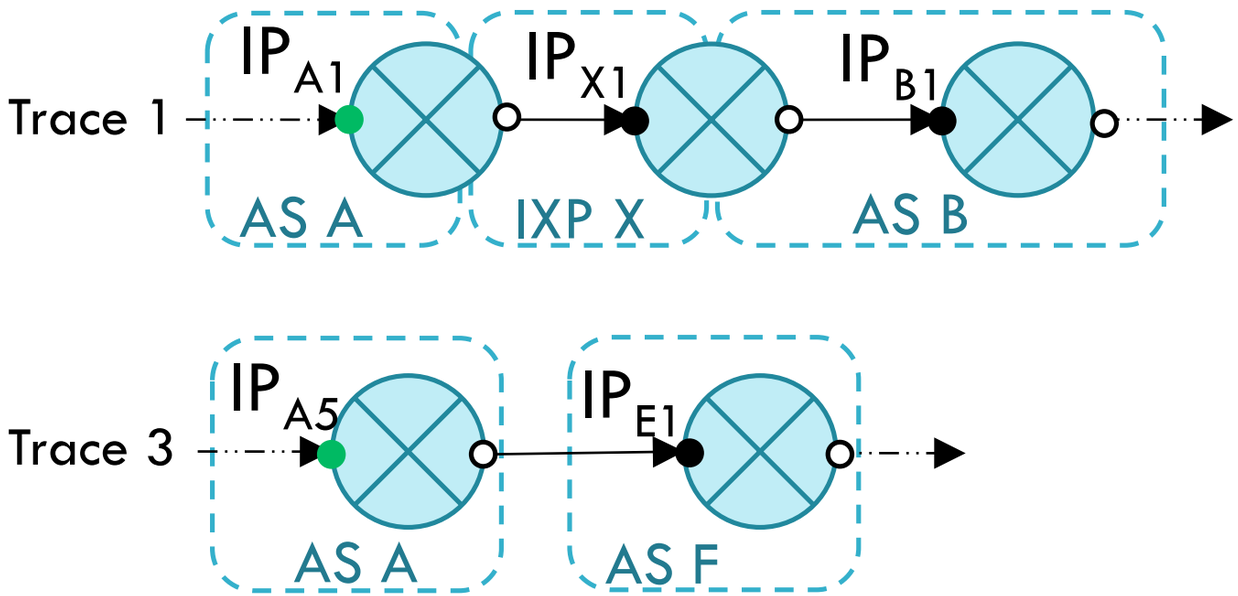

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

	Facilities		
AS A	F1	F2	F5
AS E		F2	F6

Traceroute target selection

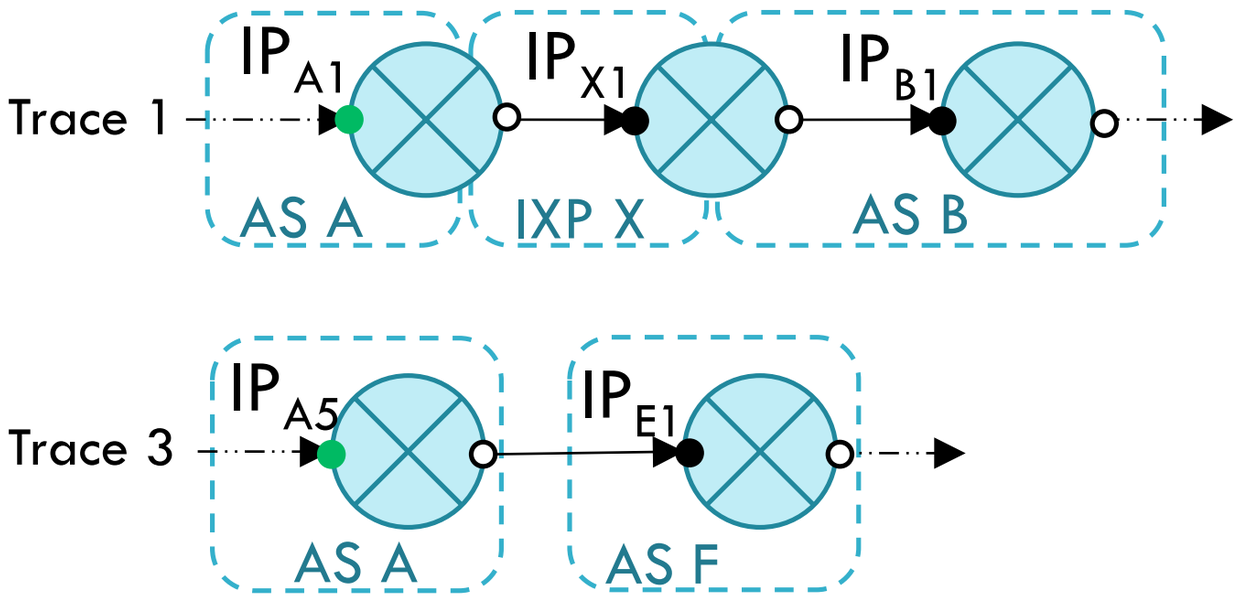

	Facilities		
AS A	F1	F2	F5
IXP X	F4	F2	F5

	Facilities		
AS A	F1	F2	F5
AS E		F2	F6

Targeting peers or IXPs with presence in **at least one but not in all** the possible facilities for IP_{A1} can offer additional constrains (depending on alias resolution)

Last Resort: Switch proximity heuristic

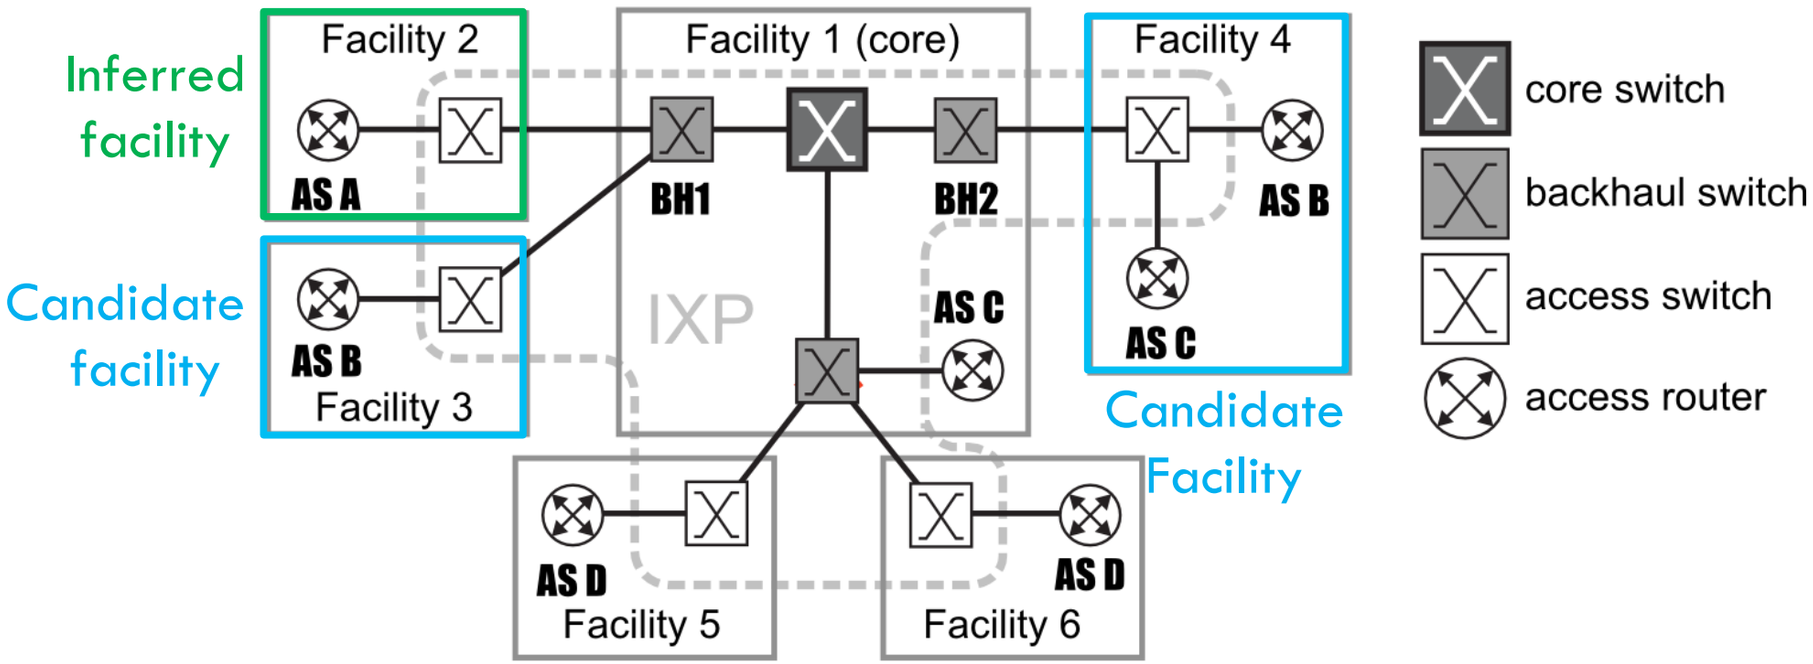

- Projecting the facilities on the IXP topology can help us reason about the actual facility of the peer at the far end

Switch proximity heuristic

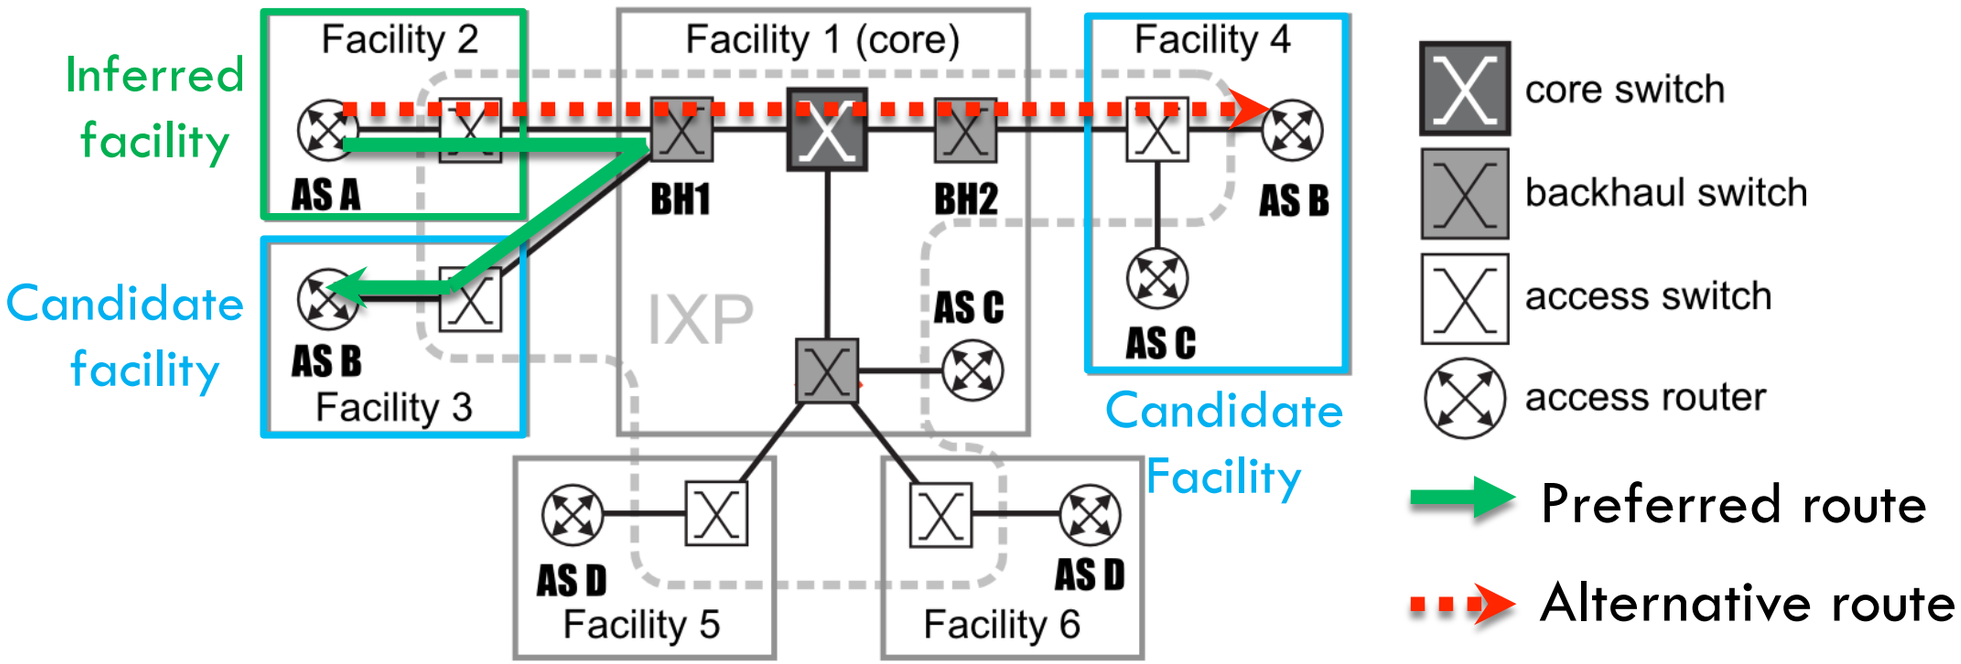

- IXPs prefer to exchange traffic over the backhaul switches instead of the core if possible

Switch proximity heuristic

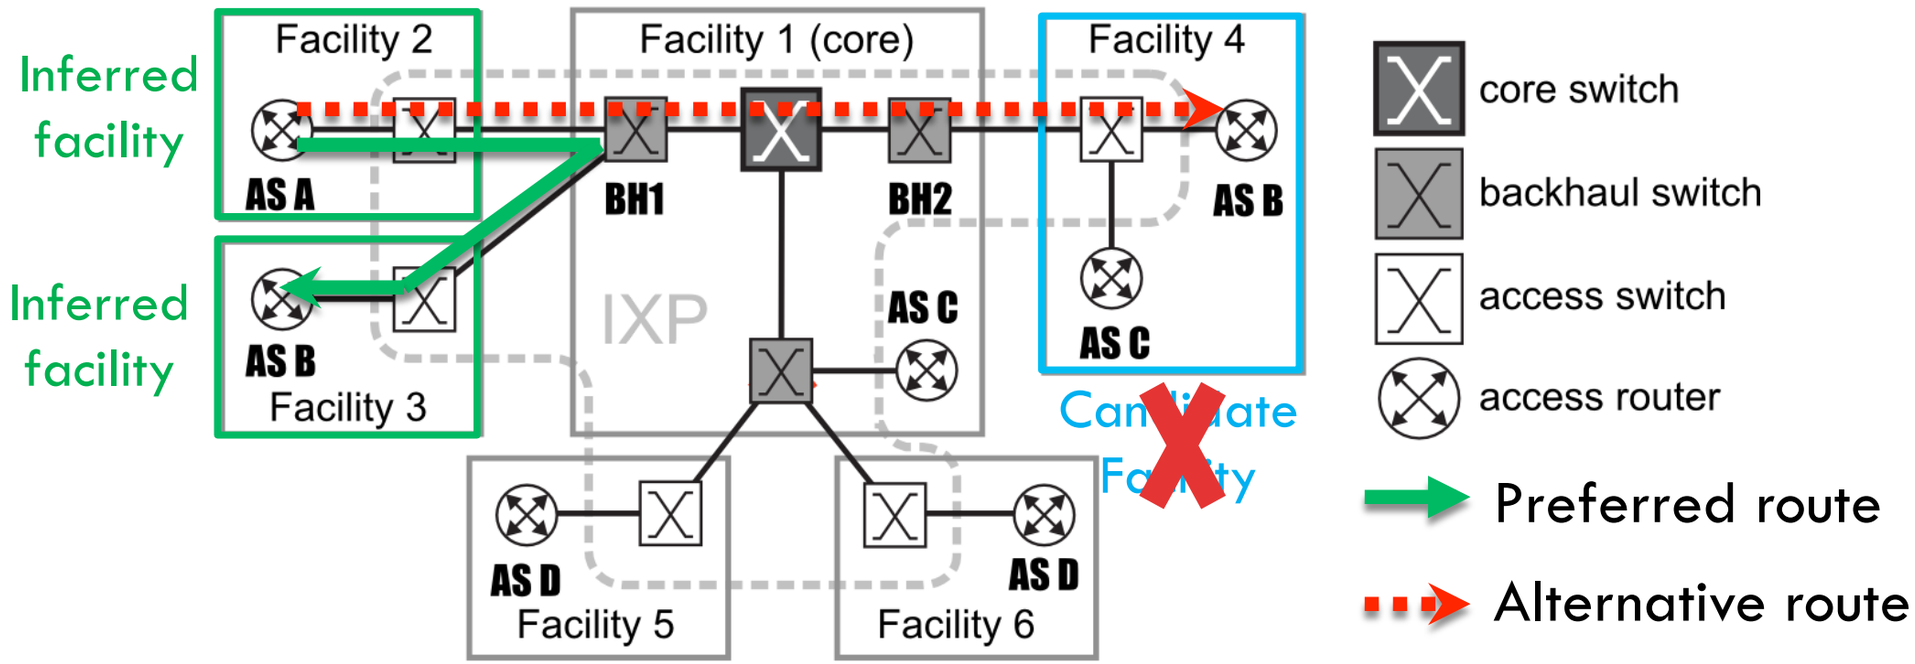

- We infer the facility of the far-end peer to be the one most proximate to the facility of the near-end peer