

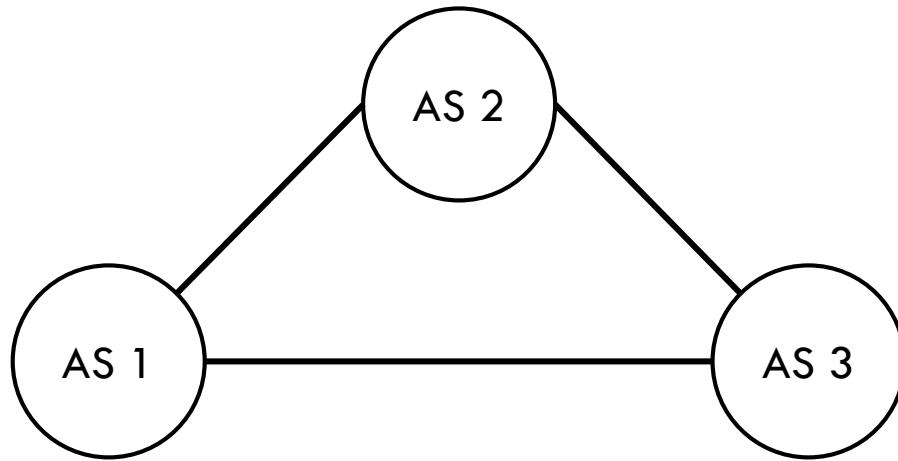
# MAPPING PEERING INTERCONNECTIONS TO A FACILITY

Vasileios Giotsas <sup>1</sup>  
Georgios Smaragdakis <sup>2</sup>  
Bradley Huffaker <sup>1</sup>  
Matthew Luckie <sup>3</sup>  
kc claffy <sup>1</sup>

[vgiotsas@caida.org](mailto:vgiotsas@caida.org)

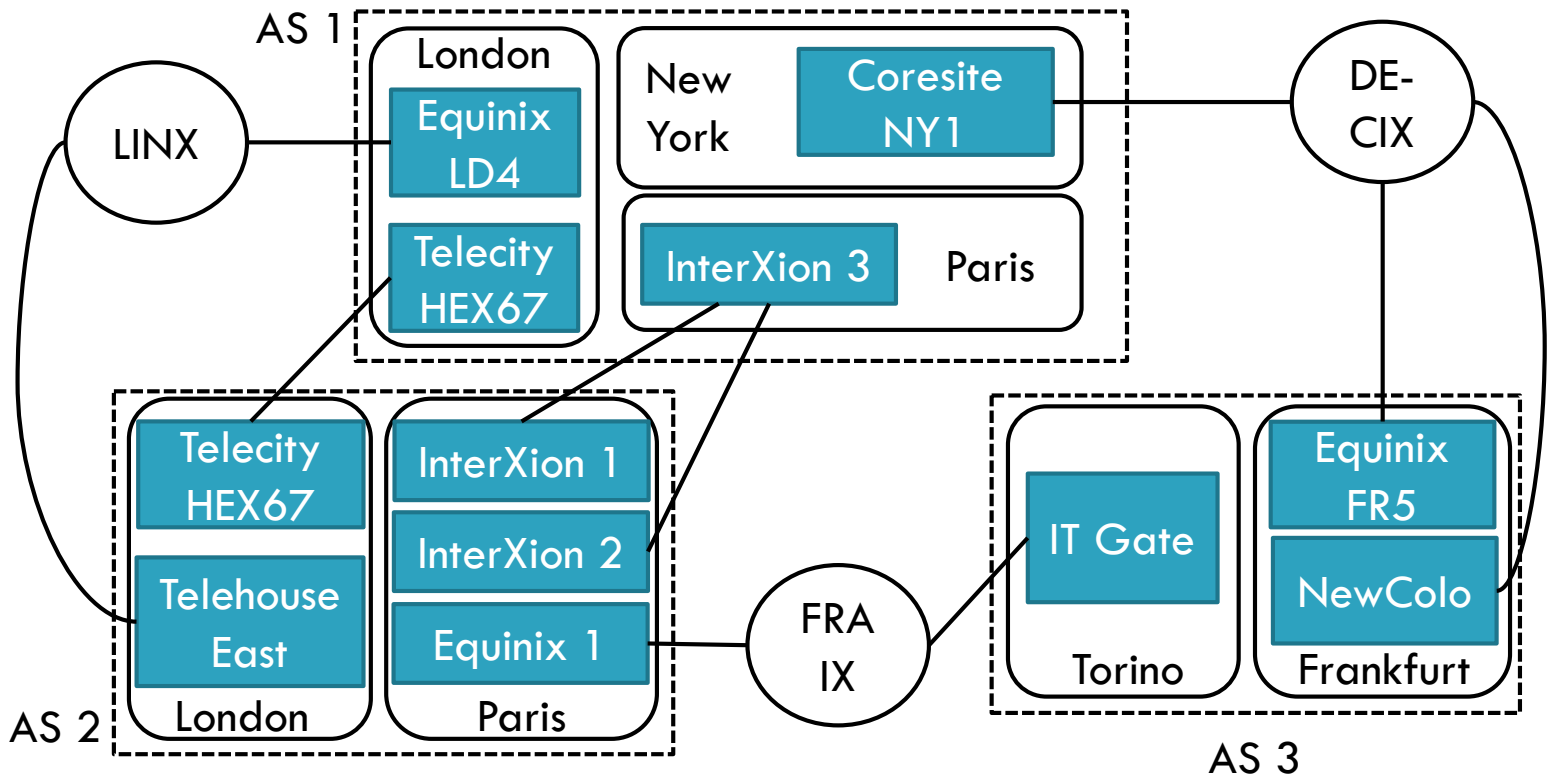
# The AS-level topology is too coarse for complex networking problems

2



# The building-level topology captures rich semantics of peering interconnections

3



# Motivation

4

- 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

5

- IP addresses are logical and region-independent
- BGP does not encode geographic information
- Existing methods are accurate for city-level granularity, not for finer granularities:
  - ▣ Delay-based
  - ▣ Hostname heuristics
  - ▣ Database-driven

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

6

- **Interconnection facilities:** special-purpose buildings used to co-locate routing equipment




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

7

- **Interconnection facilities:** special-purpose buildings used to co-locate routing equipment



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



# Develop a map of interconnection facilities


8

- ❑ 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

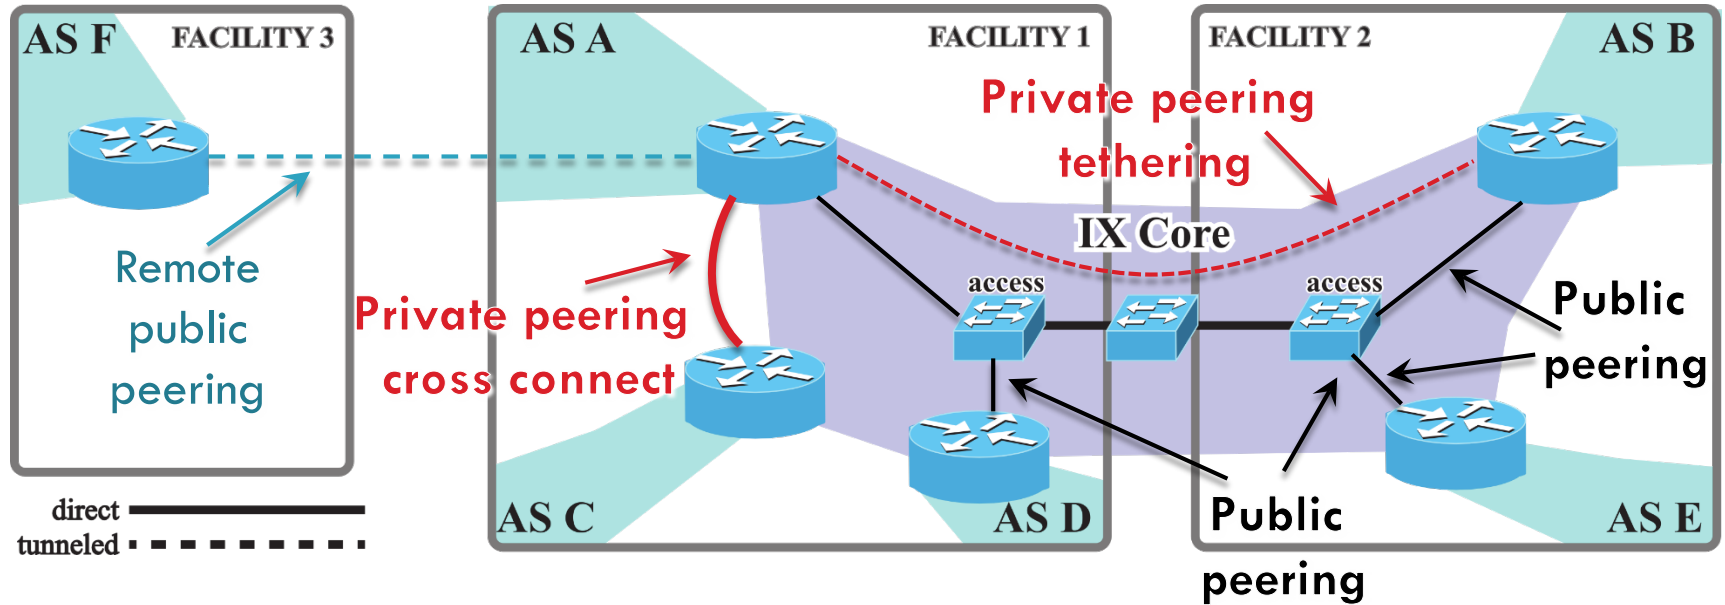


# Interconnection facilities are concentrated in hub cities




# Complexity of peering interconnections

10



# Complexity of peering interconnections

11



# Constrained Facility Search (CFS)

12

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

# Constrained Facility Search (CFS)

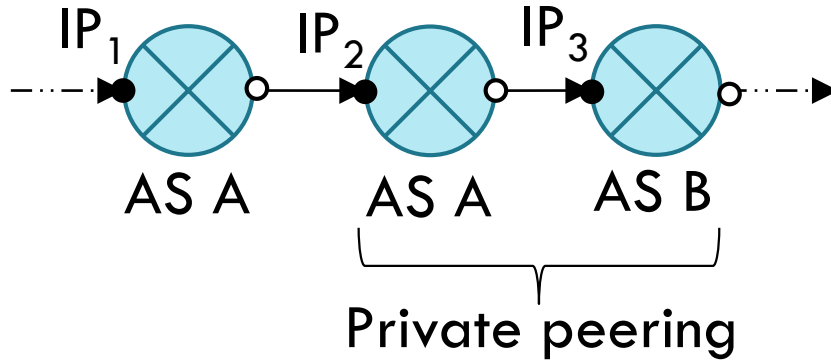
13

For a target peering interconnection ASA- ASB:

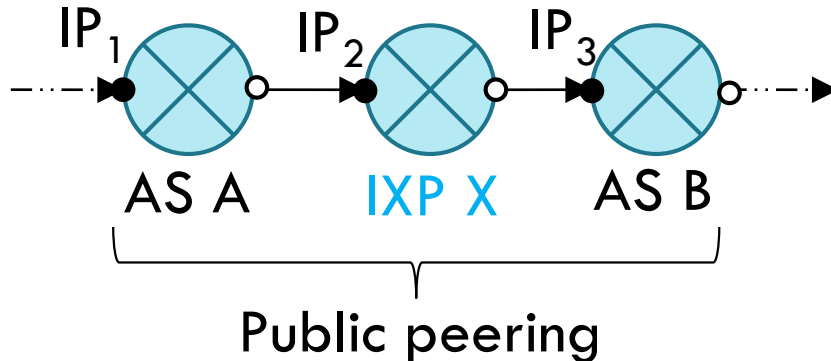
- **Step 1:** Identify the type of peering interconnection
- Step 2: Initial facility search
- Step 3: Constrain facilities through alias resolution
- Step 4: Constrain facilities by repeating steps 1-3 with follow-up targeted traceroutes
- Step 5: Facility search in the reverse direction

# Identifying the peering type

14



Facility search  
between the facilities  
of the peering Ases



Facility search  
between the IXP and  
the peering ASes

# Constrained Facility Search (CFS)


15

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

16




	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

17




Facilities		
AS A	F1	F2
IXP X	F4	F2

IP<sub>A1</sub> facility

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

# Facility search: no common facility

18




	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

# Facility search: multiple common facilities

19




Facilities			
AS A	F1	F2	F5
IXP X	F4	F2	F5

- Possible facilities are constrained but no inference yet

# Facility search: multiple common facilities

20



Facilities			
AS A	F1	F2	F5
IXP X	F4	F2	F5

Possible IP<sub>A1</sub> facilities

- Possible facilities are constrained but no inference yet

# Constrained Facility Search (CFS)



21

For a target peering interconnection ASA- ASB:

- Step 1: Identify the type of peering interconnection
- Step 2: Initial facility search
- **Step 3: Derive constraints 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 constraints through alias resolution

22



## Facilities

AS A	F1	F2	F5
IXP X	F4	F2	F5

Possible  $IP_{A1}$  facilities


## Facilities

AS A	F1	F2	Possible $IP_{A2}$ facilities
AS C	F1	F2	facilities

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

# Derive constraints through alias resolution

23



Possible IP<sub>A1</sub> facilities

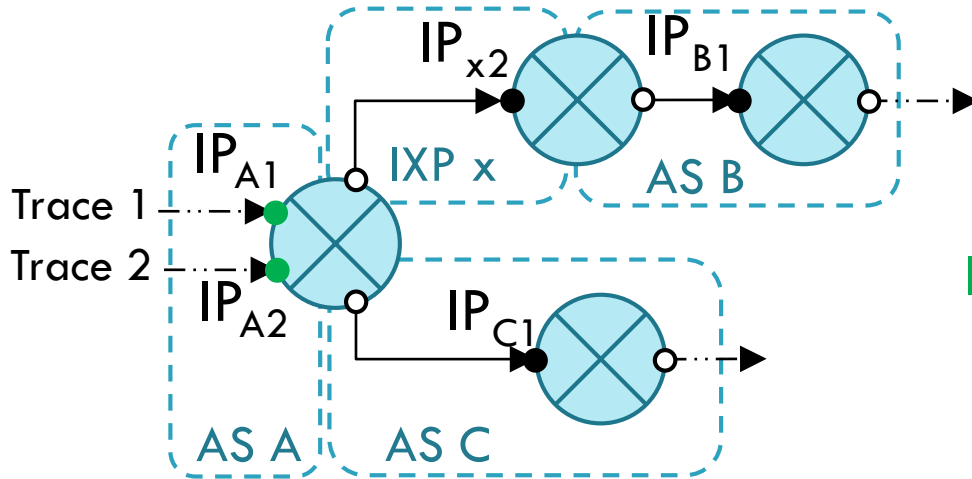
Facilities			
AS A	F1	F2	F5
IXP x	F4	F2	F5

Facilities			
AS A	F1	F2	Possible IP <sub>A2</sub> facilities
AS C	F1	F2	facilities

- De-alias interfaces of AS A (IP<sub>A1</sub>, IP<sub>A2</sub>)

# Derive constraints through alias resolution

24



IP<sub>A1</sub> & IP<sub>A2</sub> facility

		Facilities		
AS A	F1	F2	F5	
IXP x	F4	F2	F5	

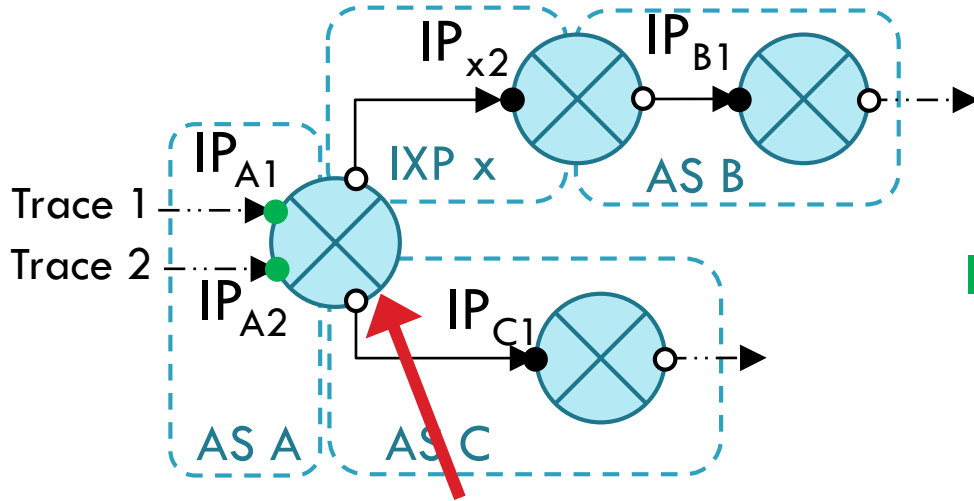
		Facilities		
AS A	F1	F2	F5	
AS C	F1	F2	F3	

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



# Derive constraints through alias resolution

25



**Multi-purpose router**

Used to establish both  
private and public peering

IP<sub>A1</sub> & IP<sub>A2</sub> facility

Facilities			
AS A	F1	F2	F5
IXP x	F4	F2	F5

Facilities			
AS A	F1	F2	F5
AS C	F1	F2	F3

# Constrained Facility Search (CFS)


26

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

# Follow-up CFS iterations

27





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

28

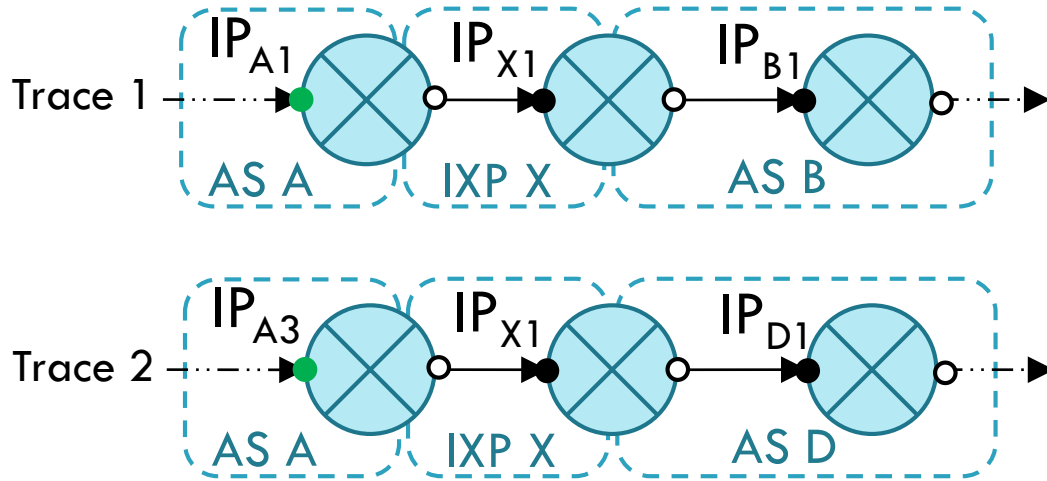


Facilities			
AS A	F1	F2	F5
IXP X	F4	F2	F5

Facilities			
AS A	F1	F2	F5
IXP X	F4	F2	F5

# Traceroute target selection

29



		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 constraints because CFS still compares the same sets of facilities

# Traceroute target selection

30





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

31





		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

32



## Facilities

AS A	F1	F2	F5
IXP X	F4	F2	F5



## Facilities

AS A	F1	F2	F5
AS E		F2	F6



# Traceroute target selection

33



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)

# Constrained Facility Search (CFS)


34

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**

# Facility inference for the far-end peer

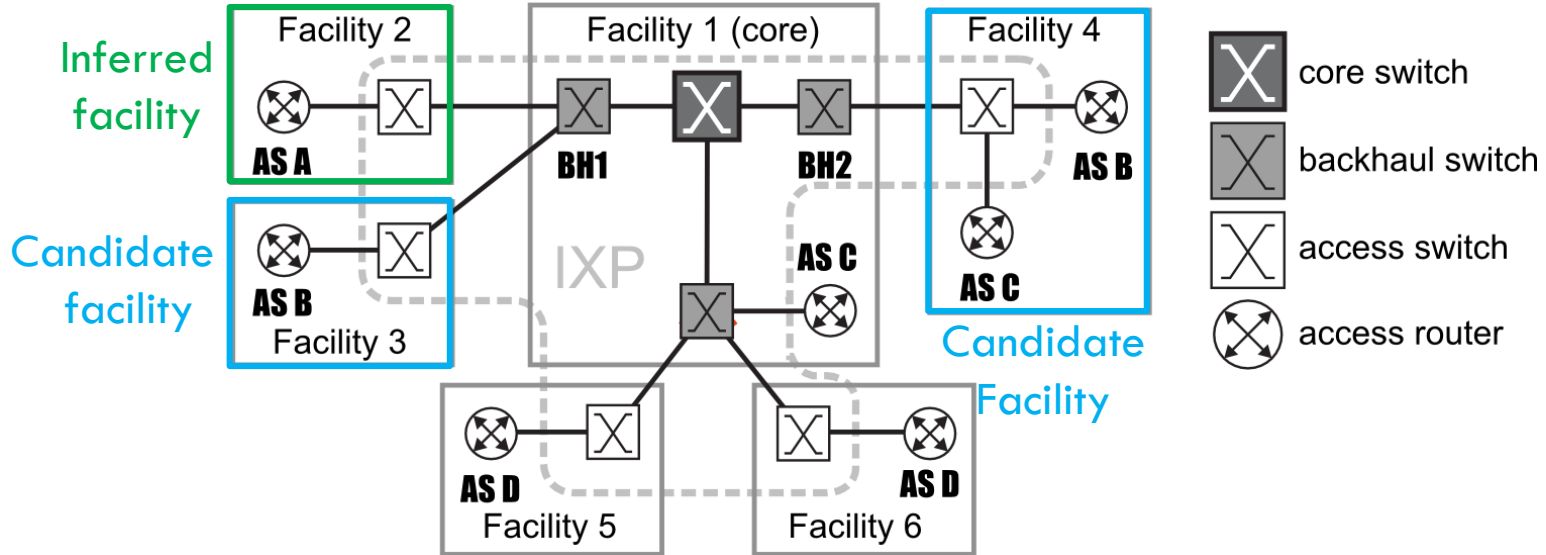
35



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

# Switch proximity heuristic

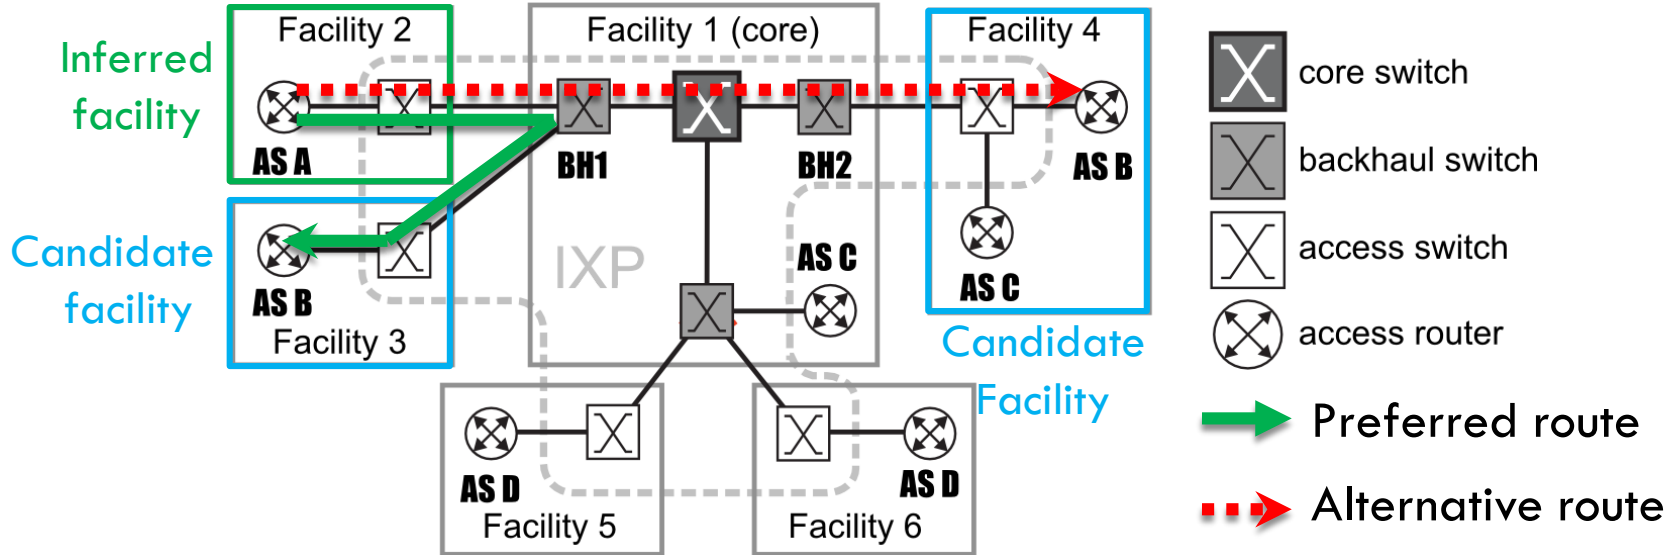
36



- 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

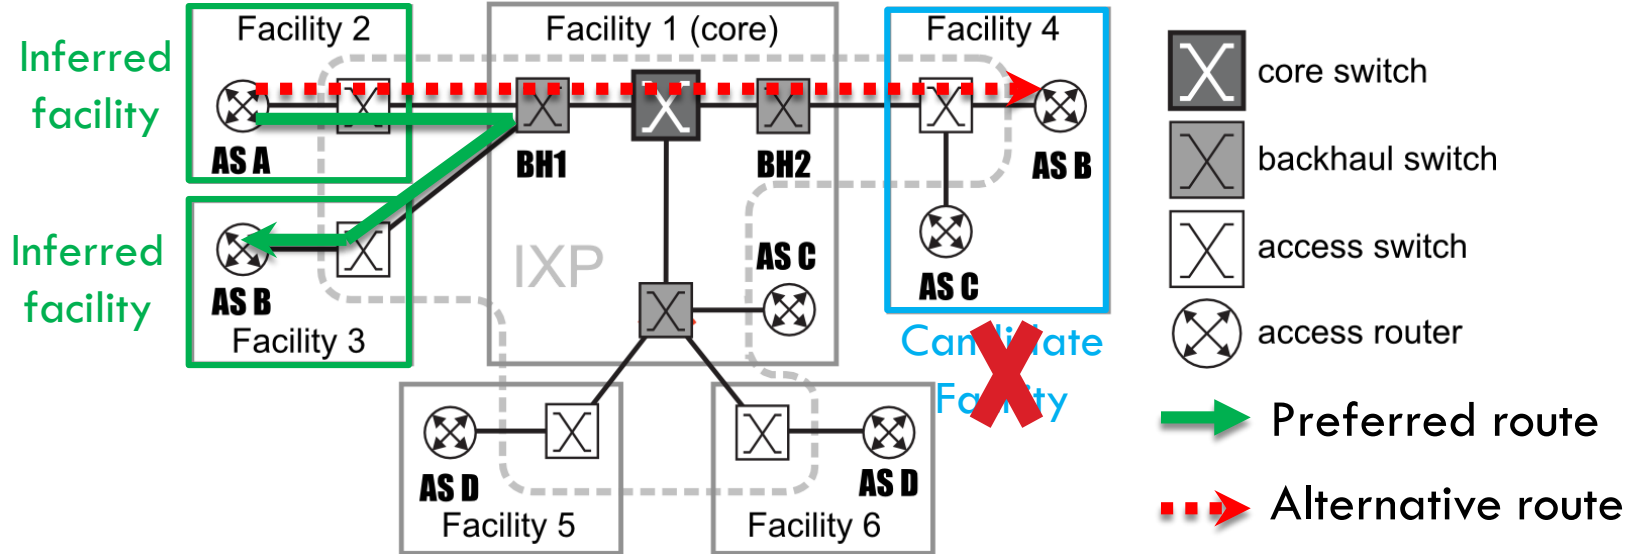
37



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

# Switch proximity heuristic

38



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

# Evaluation

39

- 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
- Executed 100 iterations of the CFS algorithm

# Collecting traceroute paths

40


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

	<b>RIPE Atlas</b>	<b>LGs</b>	<b>iPlane</b>	<b>Ark</b>	<b>Total Unique</b>
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

41



# 10% of the inferences validated to 90% correctness

42



# Ongoing and future work

43

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

<sup>1</sup> <http://internetatlas.org/>

<sup>2</sup> [https://www.caida.org/workshops/aims/1503/slides/aims1503\\_katzbasset1.pdf](https://www.caida.org/workshops/aims/1503/slides/aims1503_katzbasset1.pdf)

# Conclusions

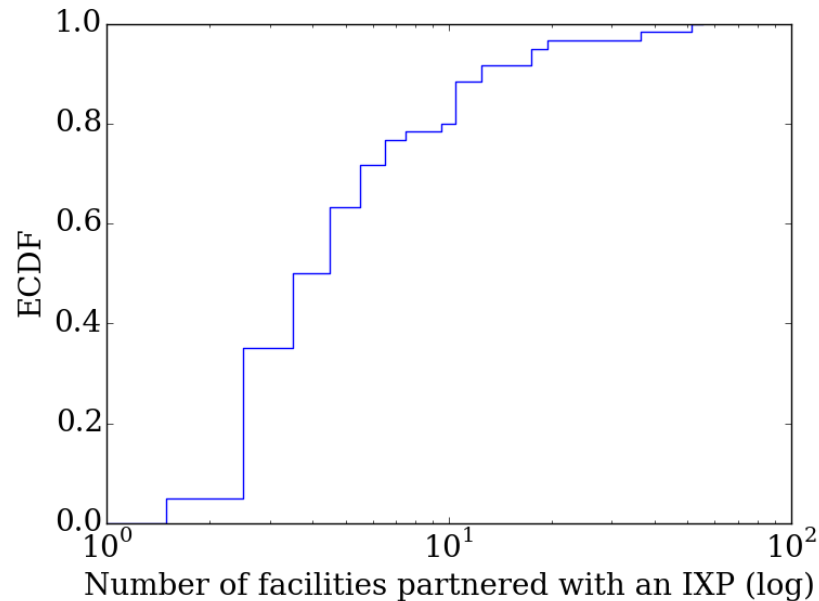
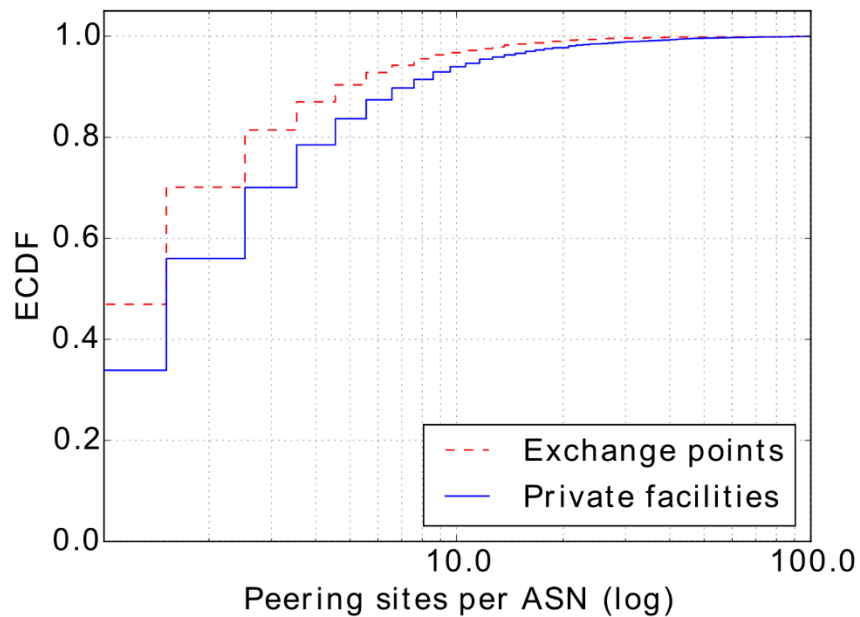
44

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

# Additional results

# ASes and IXPs are present at multiple facilities


46



# Facility data in PeeringDB are incomplete

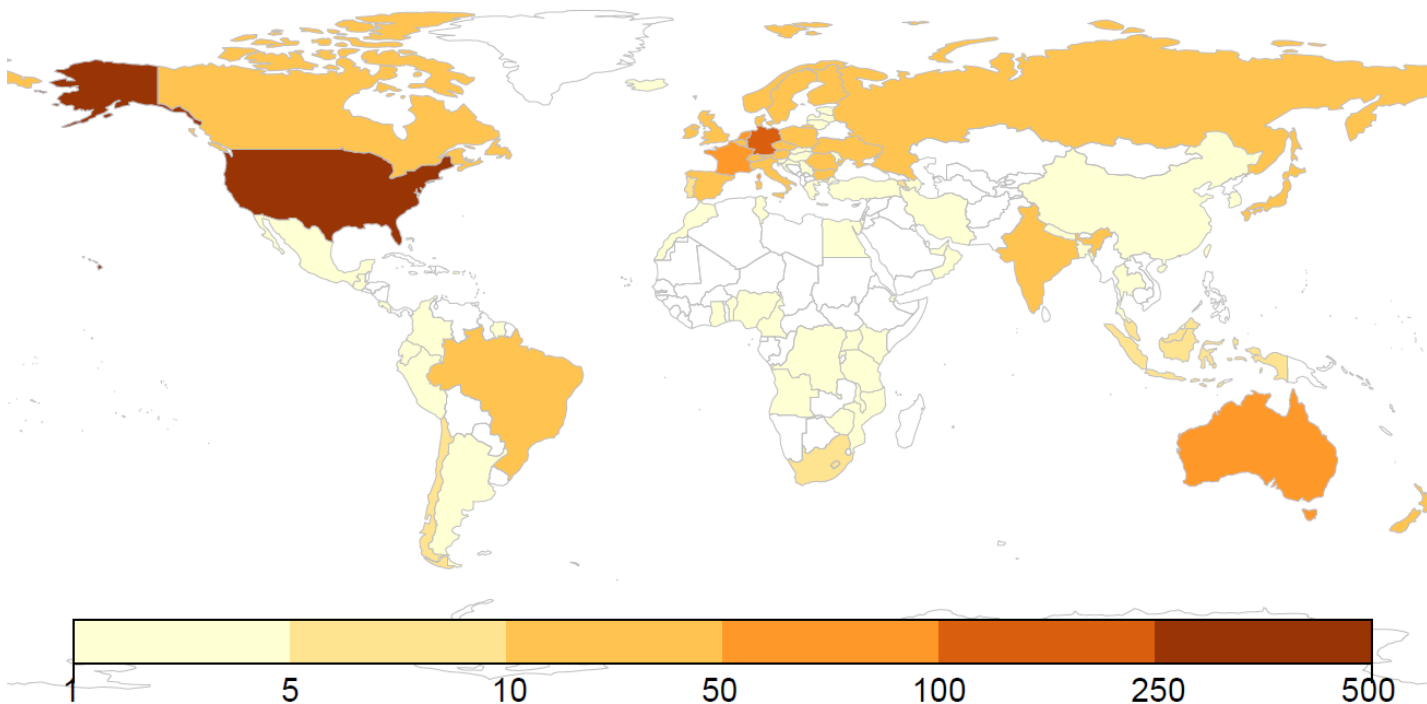
47

- 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



# Majority of interconnection facilities are located in Europe and North America

48

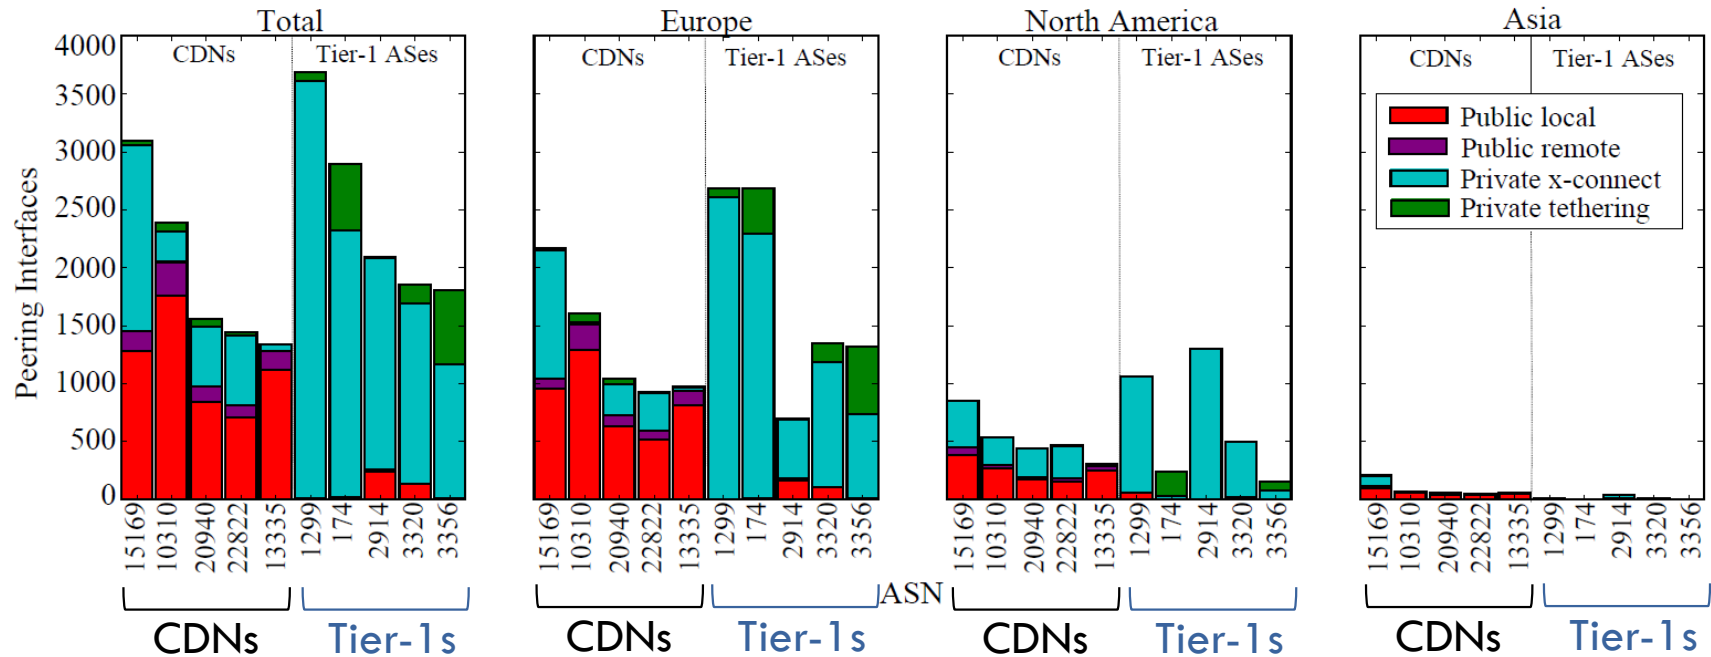


**April 2015**

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



# Diverse peering strategies between CDNs and Tier-1 ASes



# Missing facility data affect the completeness of CFS inferences

50

