Mapping Peering Interconnections to a Facility

Vasileios Giotsas[1], Georgios Smaragdakis[2], Bradley Huffaker[2], Matthew Luckie[3], Kimberly Claffy [3]

[1] CAIDA/UC San Diego, [2] MIT/TU Berlin, [3] University of Waikato

Problem definition

- While AS-level mapping has been an important step to understanding the formation and resulting structure of the Internet, it abstracts a much richer Internet connectivity map.
  - Networks may interconnect at multiple locations
- There is no mapping of AS interconnection to the location they occur
- Challenges:
  - Evolving complexity and scale of networking infrastructure
  - Information hiding properties of the routing system (BGP)
  - Security and commercial sensitivities of stackholders
  - Lack of incentives to gather or share data
- Goal:
  - A measurement and inference methodology to map a given interconnection to a physical facility.

Motivation

- Annotating peering interconnections at the level of a building facilitates:
  - Network troubleshooting and diagnosing attacks and congestion.
  - Assessment of the resilience of interconnections in the event of natural disasters, facility or router outages.
  - Peering disputes resolution.
  - Mitigating denial of service attacks.
  - Dlarify the role of emerging entities, e.g., colocation facilities, carrier hotels, and Internet exchange points (IXP)
  - Increases traffic flow transparency, e.g., to identify unwanted transit paths through specific countries.
  - Peering decisions in a competitive interconnection market.

Terminology

- Interconnection Facility
  - A physical location (a building or part of one) that supports interconnection of networks.
  - Facilities lease customers secure space to locate and operate network equipment.
  - Facilities provide power, cooling, fire protection, dedicated cabling to support different types of network connection, and in many cases administrative support.
- Internet Exchange Point
  - A physical infrastructure composed of layer-2 Ethernet switches where participating networks can interconnect their routers using the switch fabric
- Popular peering engineering options:
  - Private Peering with Cross-connect
  - Public Peering over IXP
    - Bilateral BGP connection
    - Multilateral through route server
  - Private Interconnects over IXP
  - Remote Peering / Tethering

Physical AS interconnections

- Interconnection facilities host routers of many different networks and partner with IXPs to support different types of interconnection, including:
  1. Cross-connects (private peering with dedicated medium)
  2. Public peering (peering established over shared switching fabric)
  3. Tethering (private peering using VLAN on shared switching fabric)
  4. Remote peering (transport to IXP provided by reseller).

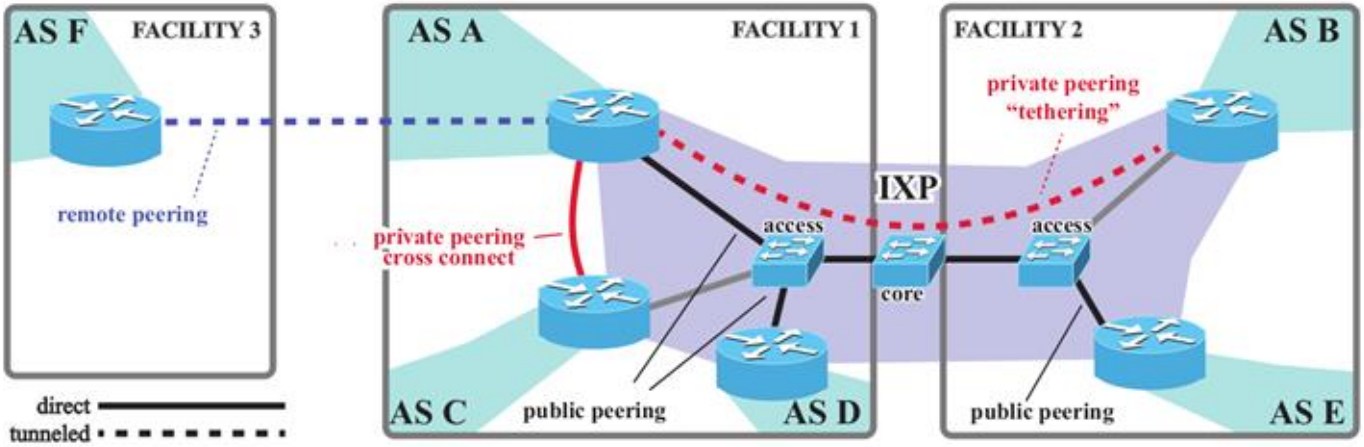

Datasets

- Facility information
  - A list of the interconnection facilities where a network is present.
  - Data source
    - PeeringDB
    - Web pages of Network Operating Centers (NOCs)
      - AS operators often document their peering interconnection facilities in this pages.
      - Only for the networks encountered in traceroutes for which PeeringDB data did not seem to reflect the geographic scope.
  - 152 ASes with PeeringDB records
    - PeeringDB misses 1,424 AS-to-facility for 61 ASes
- IXP information
  - List of IXPs
  - Their prefixes
  - Associated interconnection facilities (partner IXP and facilities)
  - Data source
    - PeeringDB
    - IXP websites
    - Regional consortia of IXPs
  - 368 IXPs in 263 cities in 87 countries

Incompleteness of PeeringDB

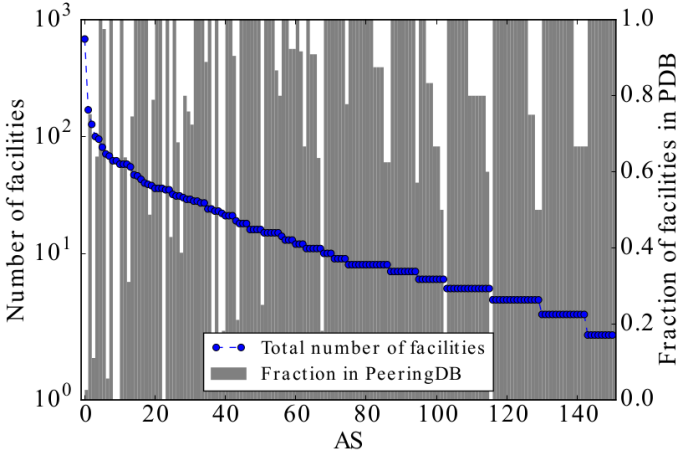

Figure 2: Number of interconnection facilities for 152 ASes extracted from their official website, and the associated fraction of facilities that appear in PeeringDB.

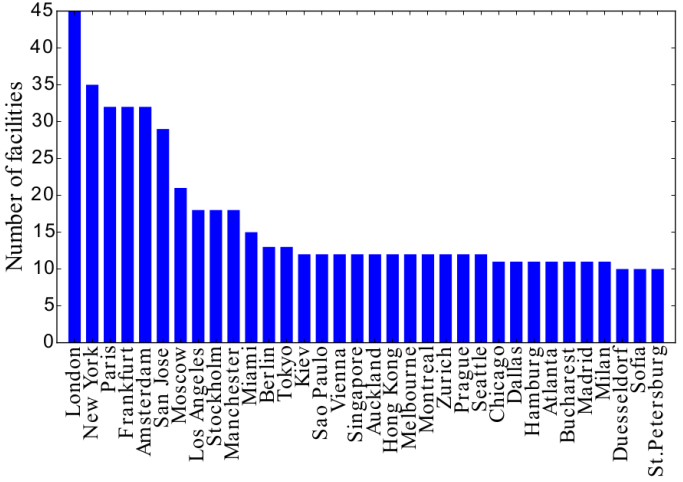

Figure 3: Metropolitan areas with at least 10 interconnection facilities.

Measurements & Vantage Points

Archived Traceroute

- iPlane
  - 300 PlanetLab nodes
  - Daily IPv4 traceroute campaigns
- CAIDA's Ark
  - 107 nodes deployed in 92 cities
  - Paris traceroutes to a randomly selected IP address in all /24 network in the advertised address space.

	RIPE Atlas	LGs	iPlane	Ark	Total unique
Vantage Pts.	6385	1877	147	107	8517
ASNs	2410	438	117	71	2638
Countries	160	79	35	41	170

Table 1: Characteristics of the four traceroute measurement platforms we utilized.

Targeted Traceroute

- RIPE Atlas
  - An open distributed Internet measurement platform.
  - Allow researchers to do ping, traceroute, and DNS lookups.
- Looking Glasses
  - Web or telnet interface to a router.
  - Allows the execution of non-privileged commands.
  - 1877 looking glasses in 438 ASes and 472 cities including many in members of IXPs and 21 offered by IXPs.

Preparation of traceroute data

- Map each IP to its AS
  - Errors may occur due to IP address sharing between siblings or neighboring ASes
- Alias resolution by MIDAR
- Map alias sets with conflicting IP interfaces to the ASN to which the majority of interfaces are mapped.
- Results in a list of
  - Public peering: (IP_A, IP_x, IP_B)
  - Private peering: (IP_A, IP_B)

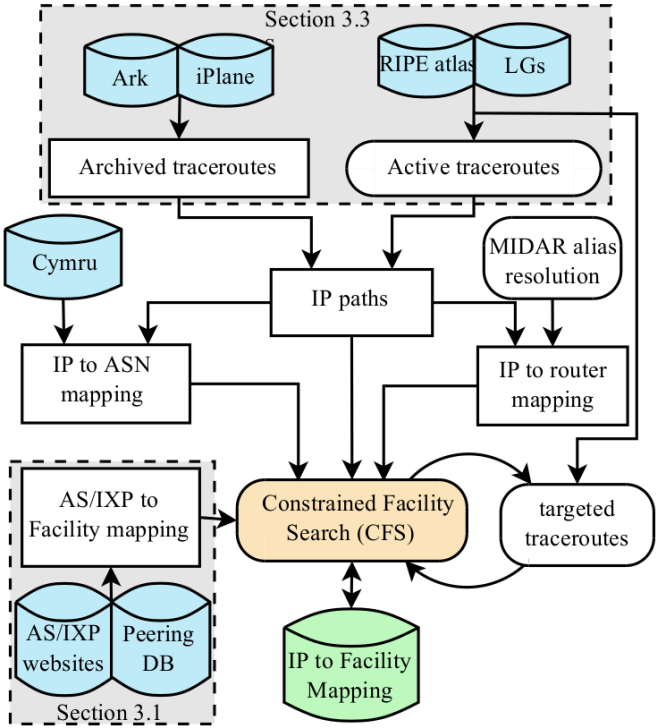

Methodology

Constrained Facility Search

- Constraint facility search: an AS link between AS A and AS B occurs in the intersection of facilities where A is present and facilities where B is present.
- Identifying public and private peering interconnections
  - IP_A, IP_x, IP_B: IXP public peering
  - IP_A, IP_B: Xconnect, tethering, remote
- Initial facility search
  - If AS A has only one common facility with the IXP
    - Resolved interface
  - If AS A has multiple common facilities with the IXP
    - Unresolved local interface
  - If AS A has no common facility with the IXP
    - Unresolved remote interface
    - Missing data
  - Same for private links
- Constraining facilities through alias resolution
  - All Aliases are located in the same facility
- Narrowing the set of facilities through follow-up targeted traceroutes

Notation	Meaning
IP_x^i	The ith IP interface that is mapped to ASN x.
(IP_x, IP_y, IP_z)	Sequence of IP hops in a traceroute path.
$\{F_A\}$	The set of interconnection facilities where ASN A is present.
$IP_x \rightarrow \{f_1, f_2\}$	The IP interface IP_x is mapped to either facility f_1 or facility f_2

Methodology

Targeted Traceroutes

- For an unresolved local peering interface IP_A, target other ASes whose facilities overlap with at least one candidate facility of A.
- The resulting traceroute will contribute constraints only if it does not cross the same IXP.
- After we launch the additional targeted traceroute
- After we launch the additional targeted traceroute
- After we launch the additional targeted traceroute Steps 2 to 4 are repeated until each interface converges to a single facility, or until a timeout set for searching expires.

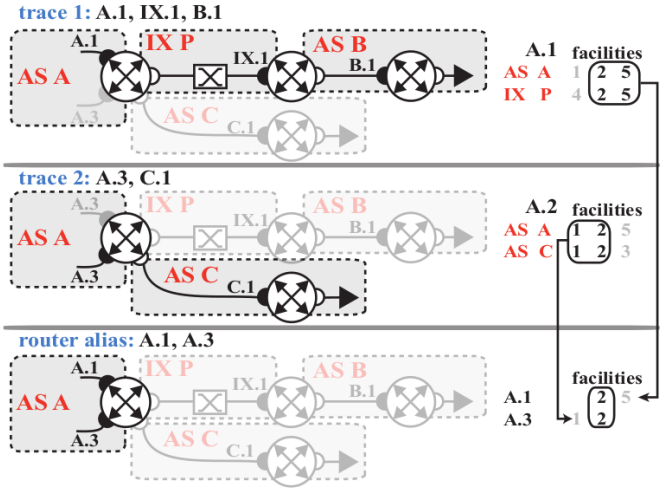

Figure 5: Toy example of how we use Constrained Facility Search (CFS) method to infer the facility of a router by probing the interconnection between peers with known lists of colocation facilities (described in detail at end of Section 4.2).

Reverse direction; Proximity Heuristic

- Networks connected to the same switch, or connected to switches attached with the same back-haul switch, exchange traffic locally and not via the core switch.
- For a public peering link $(IP_A;\ IP_{IXP;B};\ IP_B)$ for which we have already inferred the facility of IP_A, and for which IP_B has more than one candidate IXP facility, we require that IP_B is located in the facility proximate to IP_A.

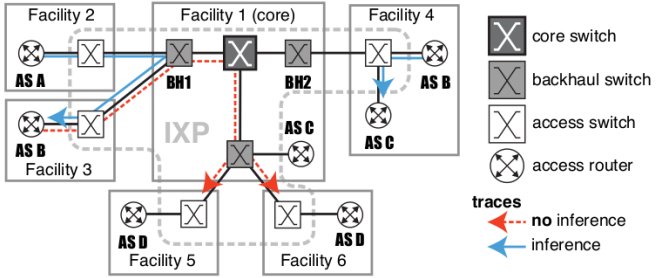

Figure 6: Toy example to illustrate the execution of the Switch Proximity Heuristic (Section 4.4) to infer the interconnection facility of the peer at the far end of an IXP peering link when the peer is connected in more than one IXP facility.

Results

- Target ASes
  - 4 large CSN: Google (AS15169), Yahoo! (AS10310), Akamai (AS20940), Limelight (AS22822) and Cloudflare (AS13335).
  - 4 large transit ASes with global footprint: NTT (AS2914), Cogent (AS174), and Deutsche Telekom (AS3320), Level3 (AS3356) and Telia (AS1299).
  - First augment the archived traceroute with active traceroute to target Ases
    - /24 prefixes of large CDNs and URLs served by these CDNs
- 9,812 router interfaces to a single interconnection facility

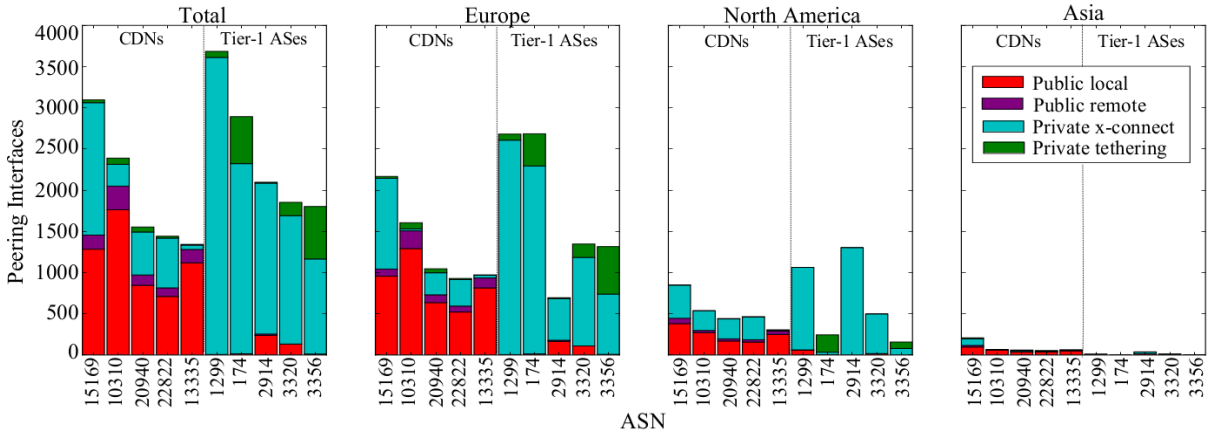

Results; Cont.

- 9,812 router interfaces to a single interconnection facility
  - 70% of all identified AS connections
- Affect of missing data
  - Iteratively executing CFS while removing 10 facilities from our dataset in random order [Figure 8]
  - Removing 850 (50%) facilities causes 30% of the previously resolved interfaces to become unresolveded interfaces become unresolved

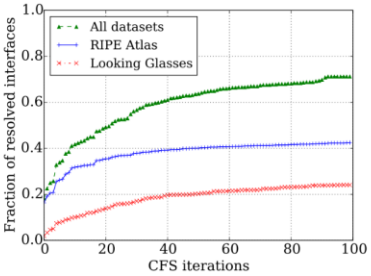

Figure 7: Fraction of resolved interfaces versus number of CFS iterations when we use all, RIPE Atlas, or LG traceroute platforms.

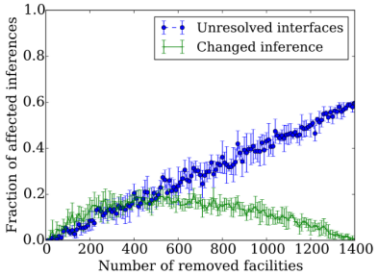

Figure 8: Average fraction of unresolved interfaces, and interfaces with changed facility inference by iteratively removing 1400 facilities.

Validation

- **Direct feedback**
  - Two CDNs. 88% (474/540) accuracy at the facility level and 95% at the city level
- **BGP communities**
  - Using 109 community values which tags the entry point of a route the network used by the large transit ASes.
  - Correctly pinpointed 76/83 (92%) of public peering interfaces and 94/106 (89%) of cross-connect interfaces.
- **DNS records**
  - Some operators encode the facility of their routers in the hostnames of the router interfaces.
  - list of naming conventions that denote interconnection facilities from 7 operators in the UK and Germany.
  - Of the interfaces validated, correctly pinpointed 91/100 (91%) of public peering interfaces and 191/213 (89%) of cross-connect interfaces.
- **IXP websites**
  - A few IXPs list on their websites the exact facilities where and the IP interfaces with which their members are connected.
  - Correctly pinpointed 322/325 (99.1%) of public peering interfaces correctly inferred 44/48 (91.7%) of remote peers.

Validation

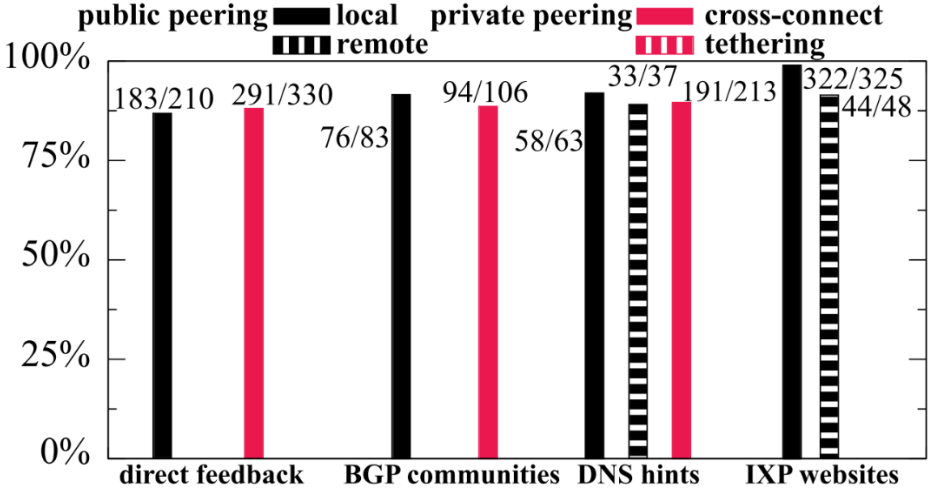

Figure 9: Fraction of ground truth locations that match inferred locations, classified by source of ground truth and type of link inferred. CFS achieves 90% accuracy overall.