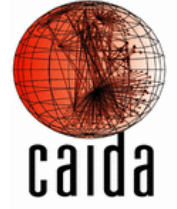

Inferring Complex AS Relationships

Vasileios Giotsas, Matthew Luckie, Bradley Huffaker, kc claffy

vgiotsas@caida.org

CAIDA – University of California San Diego

ACM SIGCOMM Internet Measurement Conference 2014, Vancouver, Canada, November 2014

Criticism on relationships inference

- Traditional relationships abstraction is *simplistic*:
  - Provider-to-customer (**p2c**)
  - Peer-to-peer (**p2p**)
  - Sibling-to-sibling (**s2s**)

- More complex relationships cannot be represented
  - Simplistic abstraction leads to artifacts / misleading results
  - *Is it possible to infer more complex relationships?*

Contributions

- Develop a new inference algorithm to infer two types of complex relationships:
  - Partial transit
  - Hybrid (dual transit/peering)
- Validate inferences against three sources of data:
  - Partial transit – 97% PPV
  - Hybrid – 93% PPV
- Infer relatively large number of complex relationships
  - ~5% of p2c

Partial transit relationships

- Provider-to-customer relationship with restricted scope
- Partial provider offers discounted transit to its customers and peers, but not its providers[1]

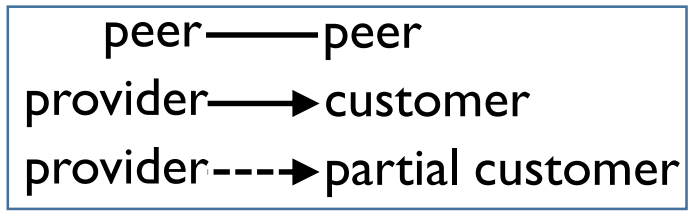

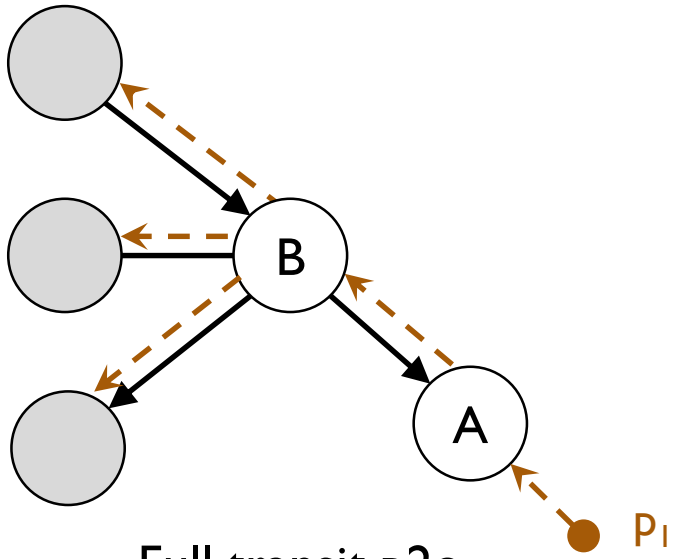

Full transit p2c
(conventional)

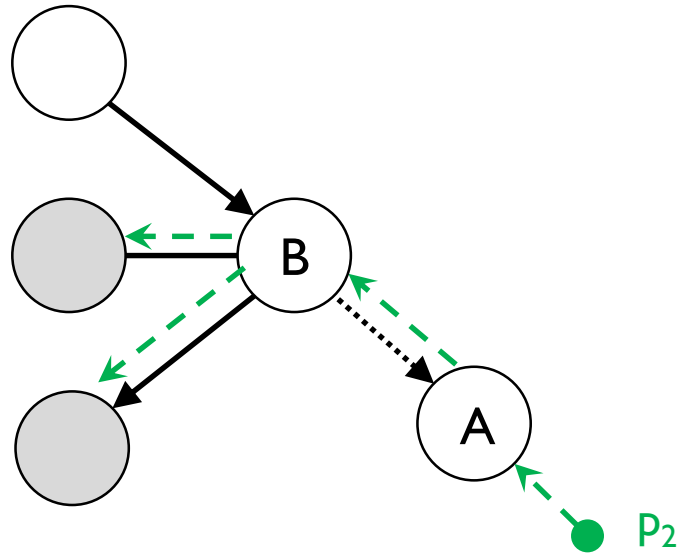

Partial transit p2c

[1] W. Norton. Partial Transit (Regional). http://drpeering.net/white-papers/Art-Of-Peering-The-Peering-Playbook.html#7

Hybrid relationships

- Two ASes establish different relationship type at different Points of Presence (PoPs)[2]

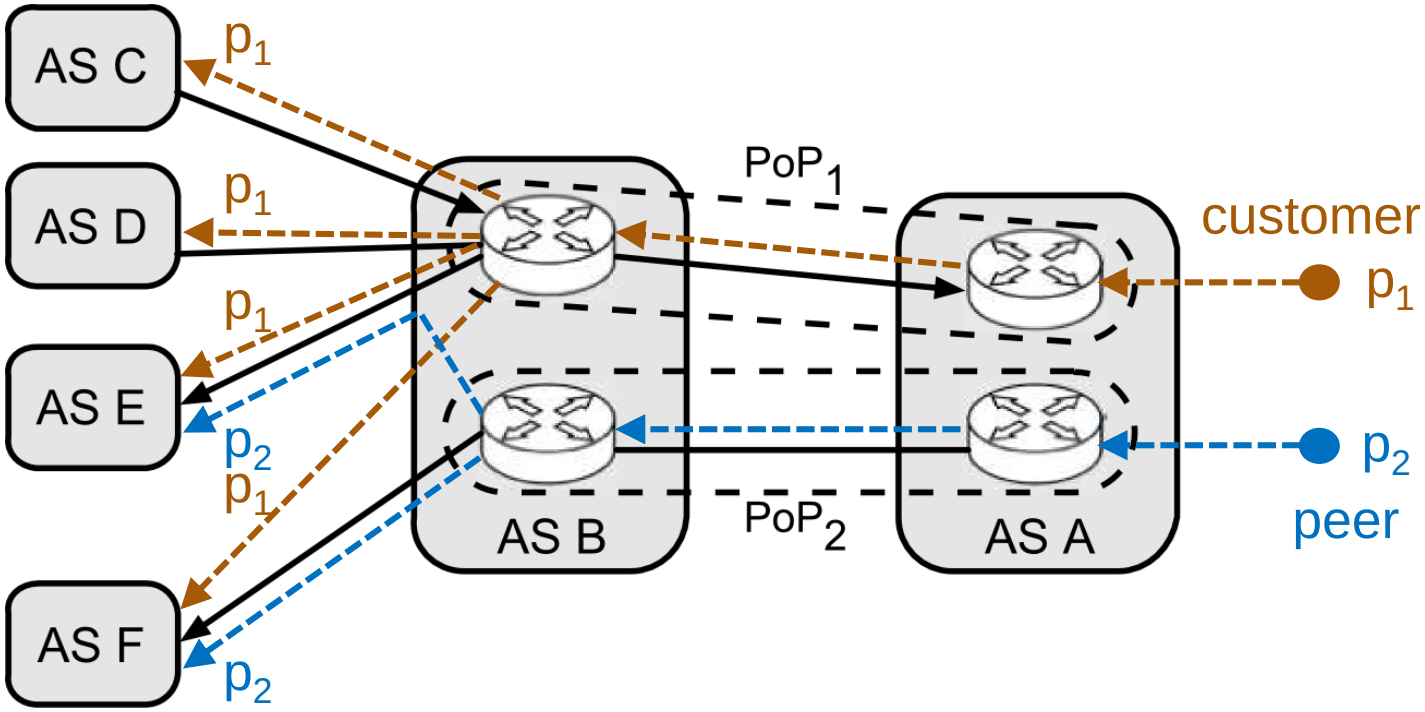

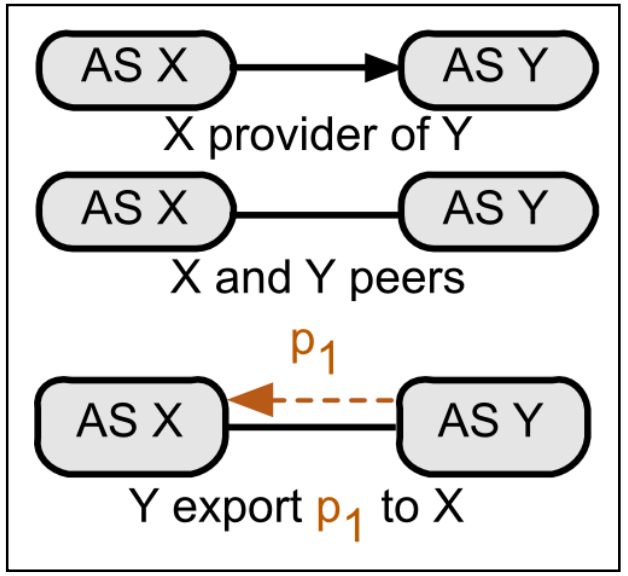

[2] W. Norton. Dual Transit/Peering. http://drpeering.net/white-papers/Art-Of-Peering-The-Peering-Playbook.html#4

Data needed to describe complex relationships

- AS Relationships
- Prefix export policies
- Locations of interconnection points

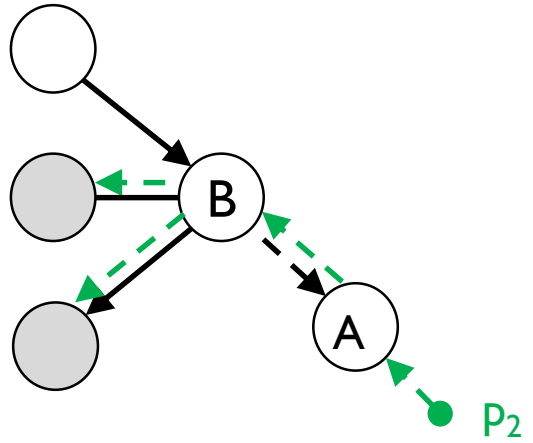

Partial transit p2c

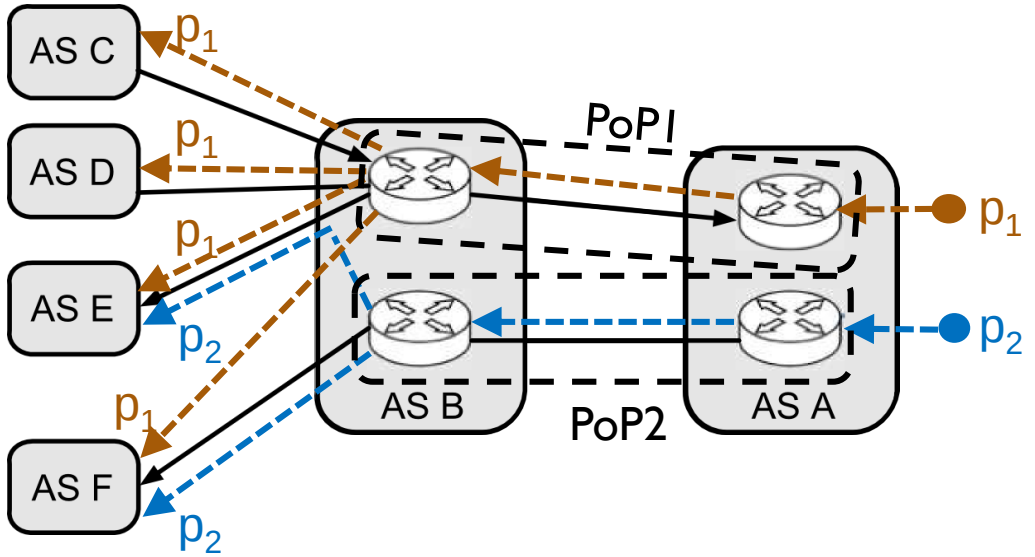

Hybrid full-transit/peering

Inference methodology

1. Use IMC 2013 algorithm[3] to infer conventional relationships
   - Both *partial transit* and *hybrid* relationships are inferred as ***p2c***

2. For each ***p2c*** link infer provider's *per-prefix* export policy
   - Full Transit
   - Partial Transit
   - Candidate Hybrid
3. For each ***candidate hybrid*** link geo-locate the ingress PoPs
4. Correlate export policies with PoPs

➔If distinct PoPs exhibit distinct export policies infer ***hybrid link***

[3] Luckie, Matthew, et al. "AS relationships, customer cones, and validation ". ACM IMC 2013.

Inference methodology

1. Use IMC 2013 algorithm[3] to infer conventional relationships
   - Both *partial transit* and *hybrid* relationships are inferred as ***p2c***

---

2. For each ***p2c*** link infer provider's *per-prefix* export policy
   - Full Transit
   - Partial Transit
   - Candidate Hybrid
3. For each ***candidate hybrid*** link geo-locate the ingress PoPs
4. Correlate export policies with PoPs

➔If distinct PoPs exhibit distinct export policies infer ***hybrid link***

[3] Luckie, Matthew, et al. "AS relationships, customer cones, and validation ". ACM IMC 2013.

Step 2: Per-prefix inference of export policy

- Find how the provider exports each prefix it receives from the customer
- Select the least restrictive policy that explains the observed behaviour

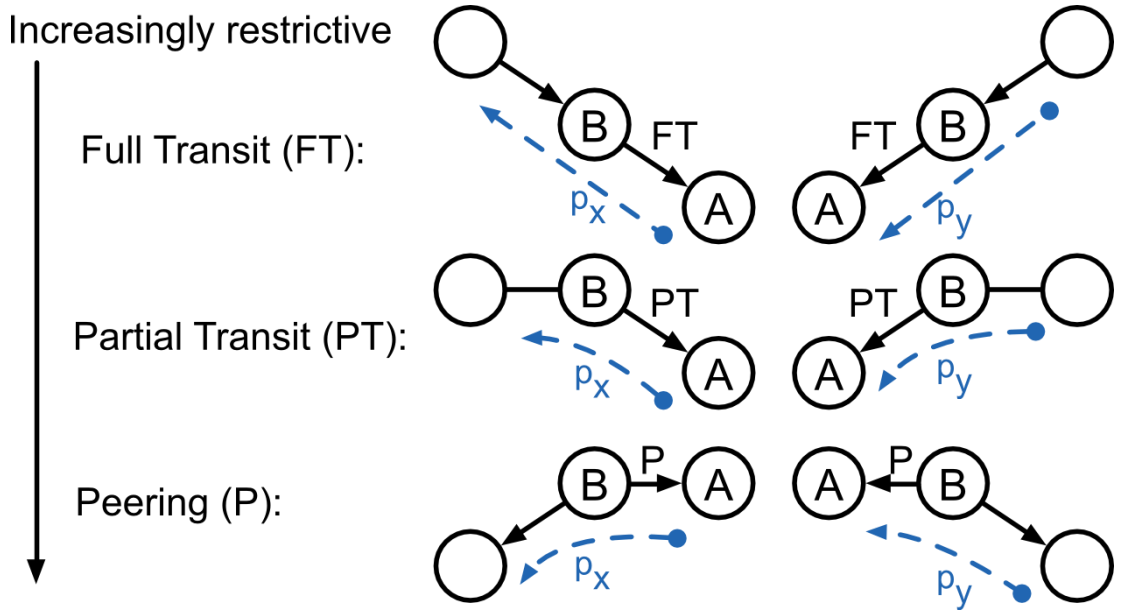

Inference methodology

1. Use IMC 2013 algorithm[3] to infer conventional relationships
   - Both *partial transit* and *hybrid* relationships are inferred as ***p2c***

2. For each ***p2c*** link infer provider's *per-prefix* export policy
   - Full Transit
   - Partial Transit

   } Same export policy for all the prefixes

   - **Candidate Hybrid**

   } Different export policy for different prefixes

3. For each ***candidate hybrid*** link geo-locate the ingress PoPs
4. Correlate export policies with PoPs

➔If distinct PoPs exhibit distinct export policies infer ***hybrid link***

[3] Luckie, Matthew, et al. "AS relationships, customer cones, and validation ". ACM IMC 2013.

Inference methodology

1. Use IMC 2013 algorithm[3] to infer conventional relationships
   - Both *partial transit* and *hybrid* relationships are inferred as ***p2c***

2. For each ***p2c*** link infer provider's *per-prefix* export policy
   - Full Transit
   - Partial Transit
   - Candidate Hybrid

3. For each ***candidate hybrid*** link geo-locate the ingress PoPs

4. Correlate export policies with PoPs

➔If distinct PoPs exhibit distinct export policies infer ***hybrid link***

[3] Luckie, Matthew, et al. "AS relationships, customer cones, and validation ". ACM IMC 2013.

Step 3: PoP Geo-location

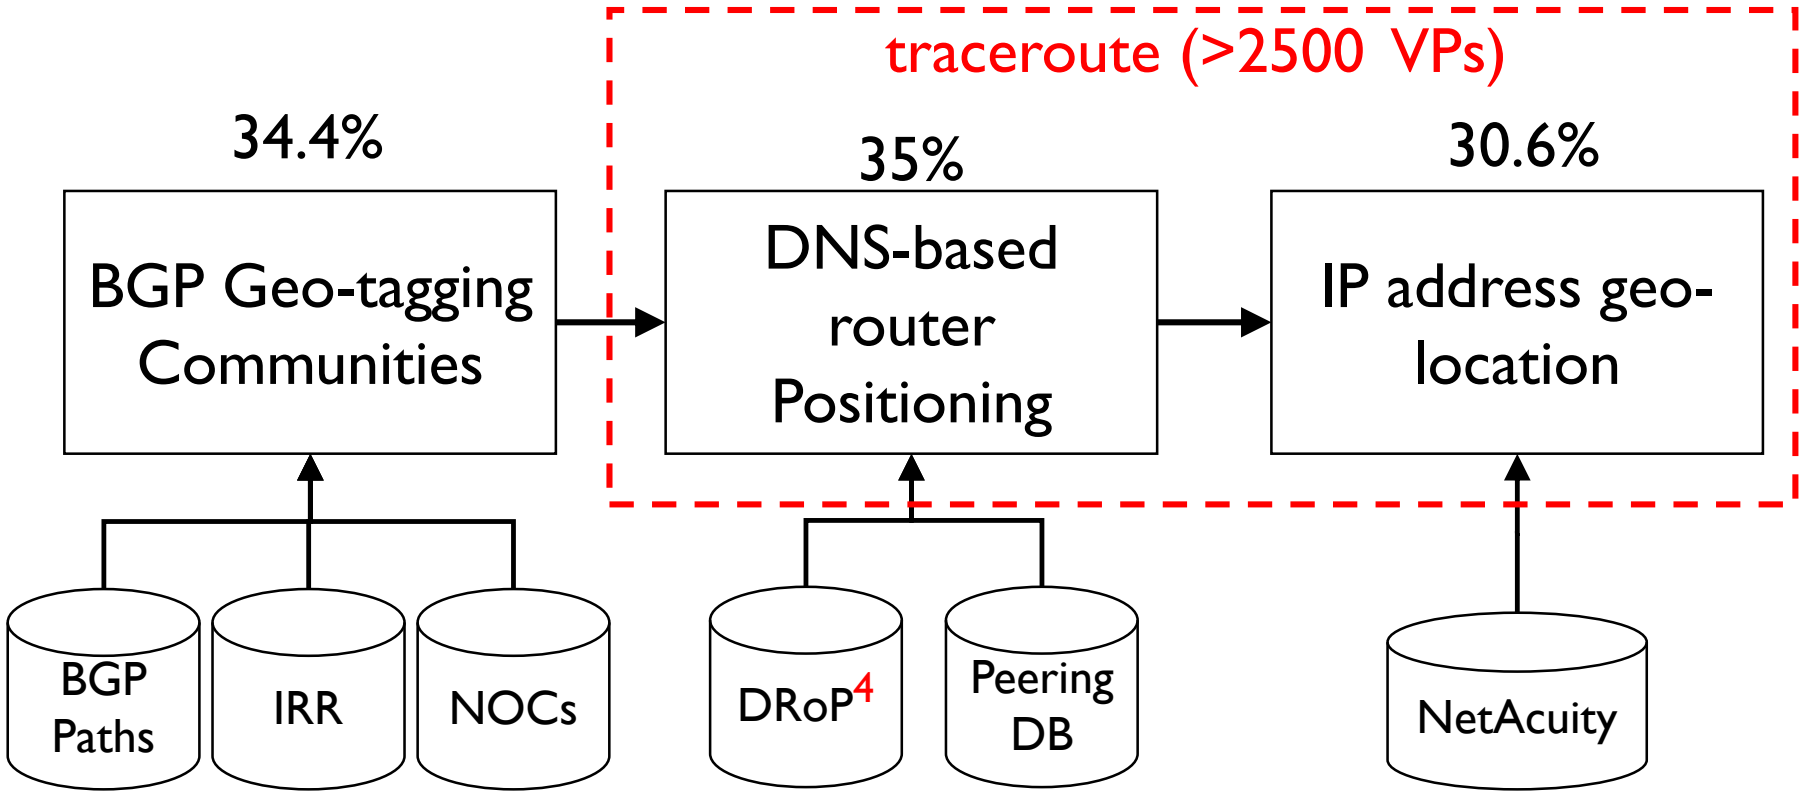

[4] B. Huffaker et al. DRoP:DNS-based router positioning. ACM SIGCOMM CCR, 44(3):6–13, July 2014

Inference methodology

1. Use IMC 2013 algorithm[3] to infer conventional relationships
   - Both *partial transit* and *hybrid* relationships are inferred as ***p2c***

2. For each ***p2c*** link infer provider's *per-prefix* export policy
   - Full Transit
   - Partial Transit
   - Candidate Hybrid

3. For each ***candidate hybrid*** link geo-locate the ingress PoPs

4. Correlate export policies with PoPs

➔If distinct PoPs exhibit distinct export policies infer ***hybrid link***

[3] Luckie, Matthew, et al. "AS relationships, customer cones, and validation ". ACM IMC 2013.

Validation

	Hybrid			Partial-Transit		
	True-Pos	False-Pos	False-Neg	True-Pos	False-Pos	False-Neg
Direct Report	33	2	1	2	0	0
Communities	124	10	4	158	5	0
RPSL	45	-	-	38	-	-
Validated	214 / 1,071: 20.0%			203 / 2,955: 6.9%		
Confirmed	202 / 1,071: 18.9%			198 / 2,955: 6.7%		
PPV	157 / 169: **92.9%**			160 / 165: **97.0%**		

Limitations

- Topology incompleteness problem
  - We can only model what we can see
- City-level geolocation granularity
  - Hybrid links within the same city can be hidden
- Difficult to neatly categorize complex relationships

Results

- 90,272 *p2c* relationships inferred for March 2014
  - 2,955 (3.3%) *partial transit* relationships
  - 1,071 (1.2%) *hybrid* relationships
- Hybrid relationships not only between large ASes
  - >50% of hybrid relationships involve AS with customer cone size < 5 ASes
  - >65% of hybrid relationships involve AS with traffic levels < 100 Gbps
- Hybrid relationships can be unintentional
  - Configuration errors
  - Open peering policies at route servers[5]

[5] Giotsas, Vasileios, "Inferring multilateral peering." *ACM CoNEXT* 2013.

Conclusion

- AS relationship inference algorithms limited by their simplistic relationship abstraction
- Implement and validate a new inference algorithm to capture partial transit and hybrid relationships with high accuracy
- Complex relationships not only among top ASes

Thank you for your attention!

Questions?

PoP Geo-location

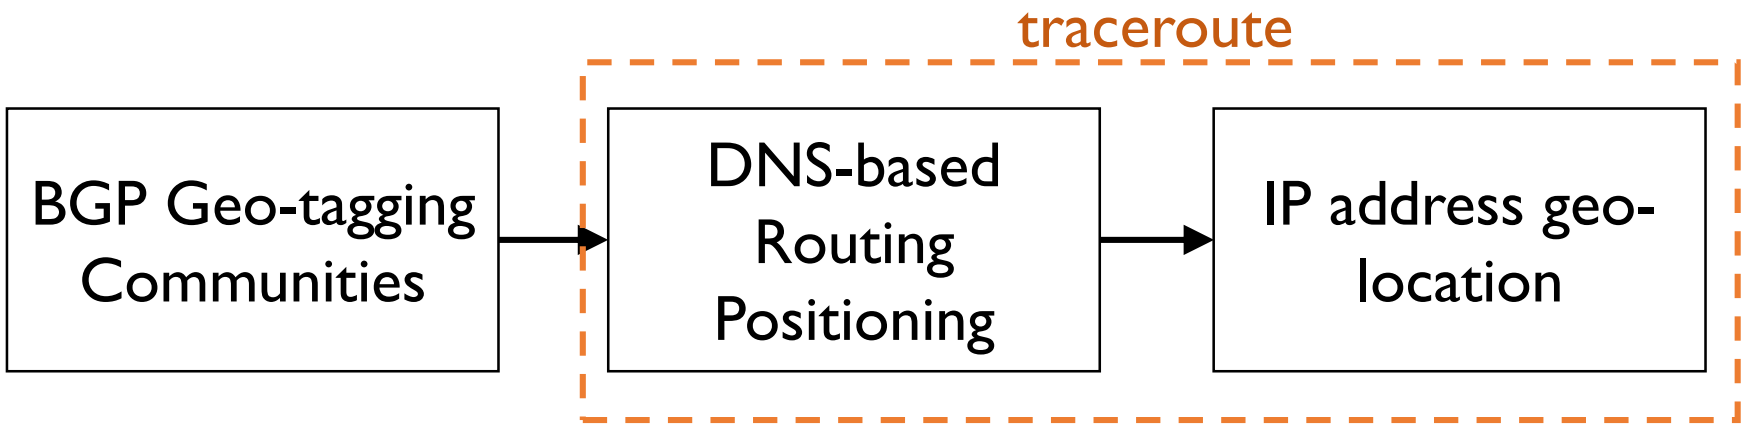

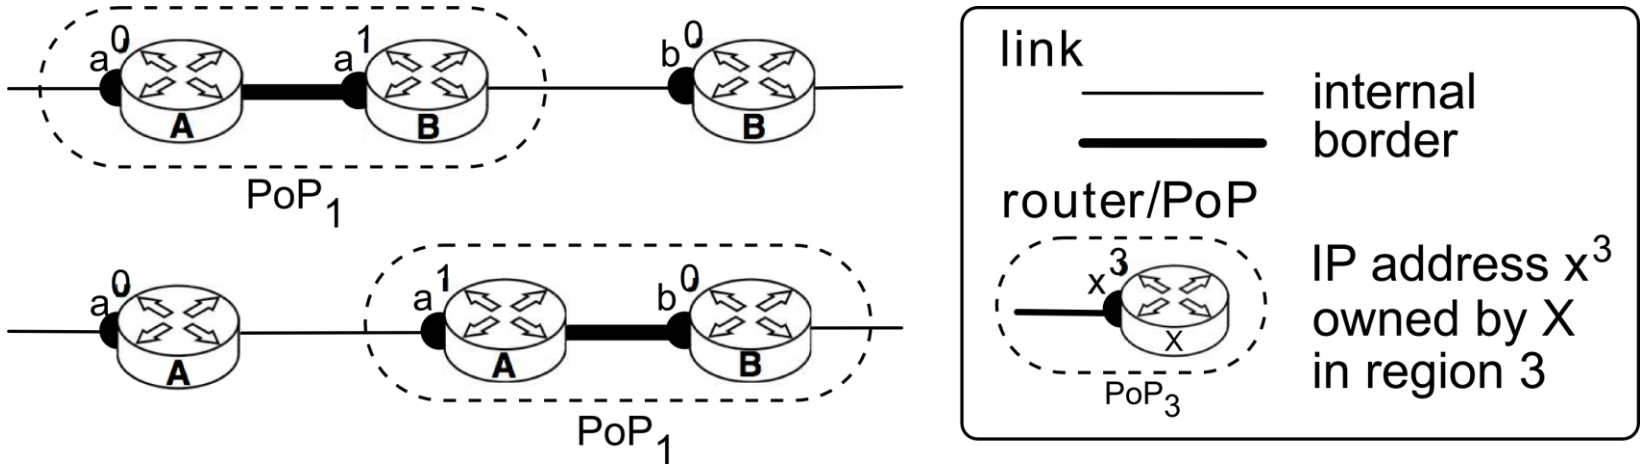

PoP Geo-location

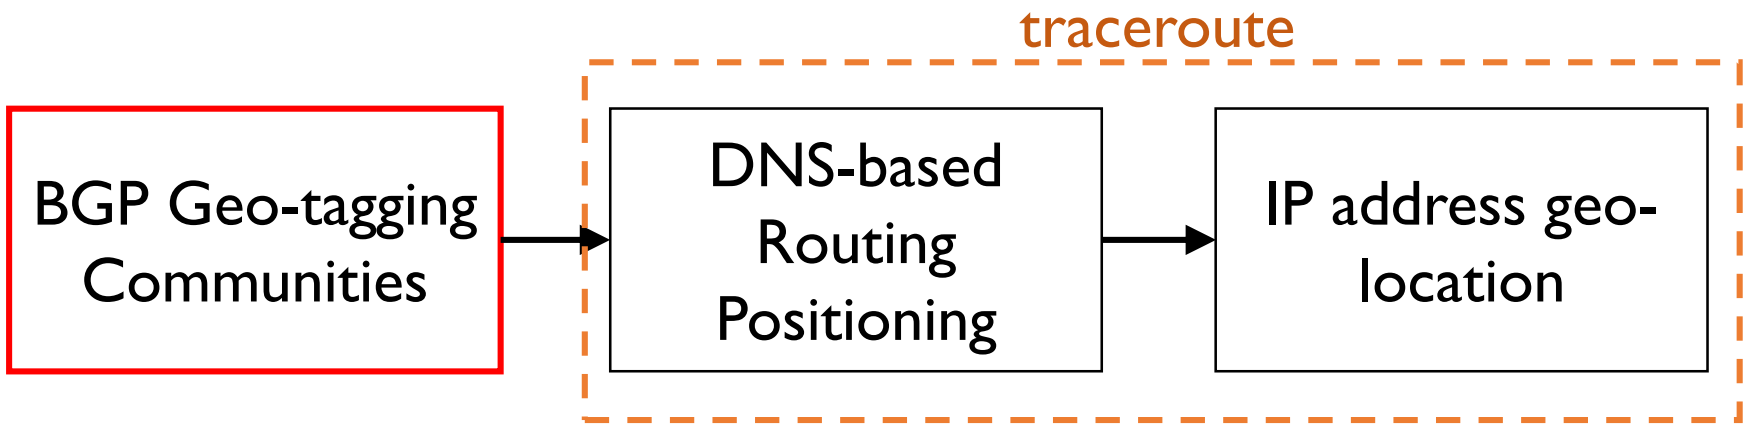


```
2914 3491 133741
    129.250.0.11 from 129.250.0.11 (129.250.0.12)
      Origin IGP, metric 6,
      localpref 100, valid, external
      Community: 2914:420 2914:1008
                 2914:2000 2914:3000
      rx pathid: 0, tx pathid: 0
```

RouteViews BGP Data

US MSA origins (2914:10--)

2914:1001	Ashburn, VA	2914:1007	Seattle, WA
2914:1001	Sterling, VA	2914:1008	Milpitas, CA
2914:1002	Atlanta, GA	2914:1008	Mountain View, CA
2914:1003	Chicago, IL	2914:1008	Palo Alto, CA
2914:1004	Dallas, TX	2914:1008	San Jose, CA
2914:1004	Houston, TX	2914:1008	Santa Clara, CA
2914:1005	Los Angeles, CA	2914:1009	New York, NY
2914:1006	Miami, FL		

NTT support center

PoP Geo-location

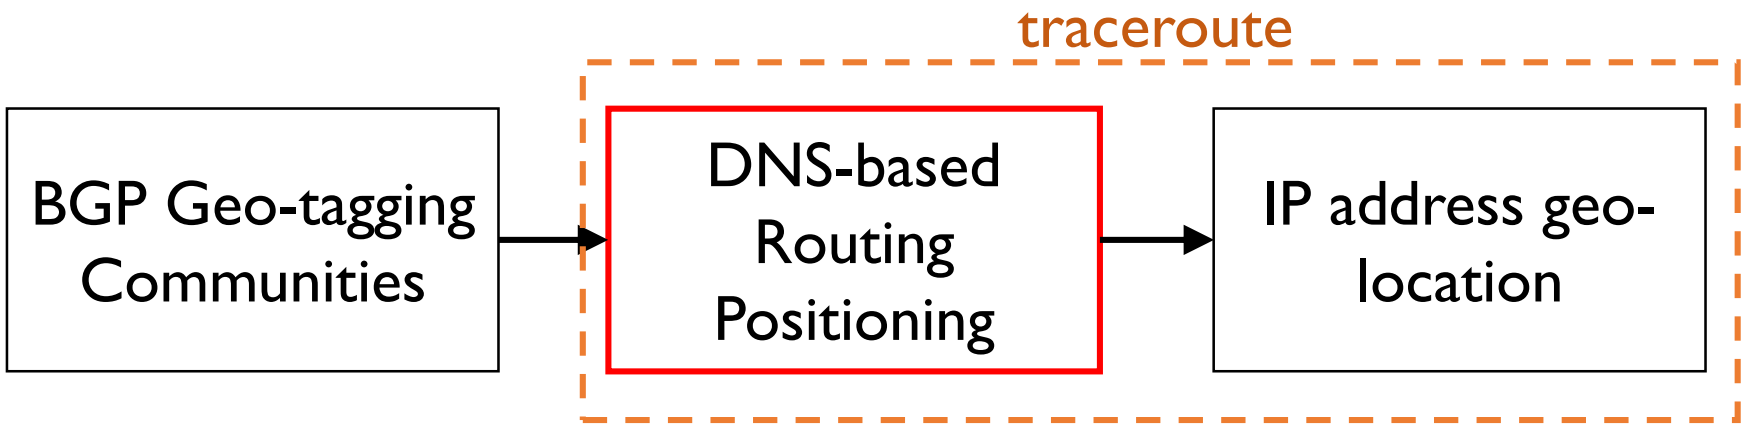


```
1   ge5-1.core1.fmt1.he.net (64.62.134.129)
2   10ge1-1.core1.pao1.he.net (184.105.213.66)
3   10ge11-6.core1.lax1.he.net (72.52.92.22)
4   lap.ln.net (198.32.146.10)
5   130.152.181.189
6   130.152.183.4
```

→ Los Angeles, CA, US

CAIDA's DNS Decoding Database: http://ddec.caida.org/

PoP Geo-location

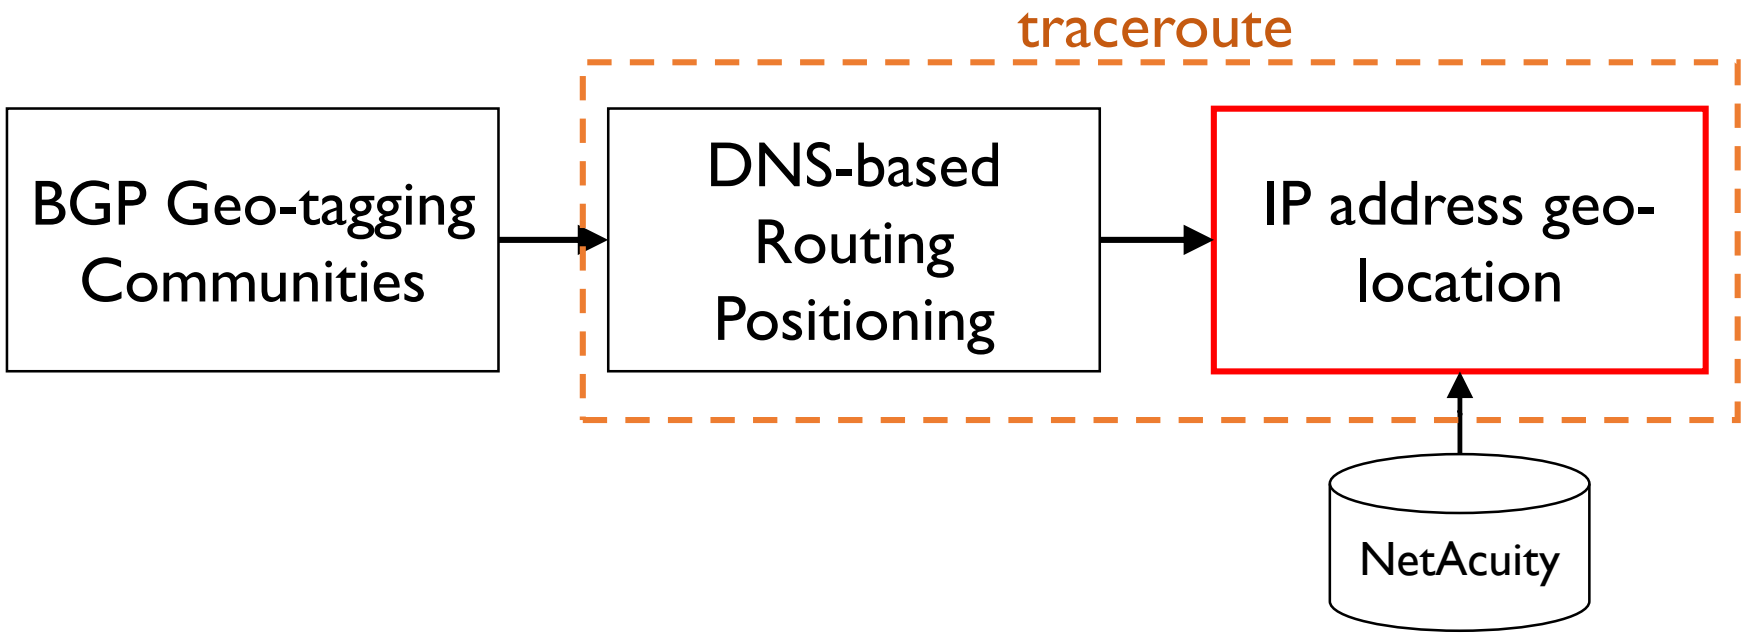

- Last-resort geo-location (optimised for end hosts)
- Neacuity found to be more accurate than other similar databases[5]

[5] B Huffaker et al. Geocompare: A comparison of public and commercial geolocation databases. Technical report, CAIDA, May 2011

Obtaining traceroute paths

- Ark:
  - 94 monitors
  - 84 Ases
  - 39 countries
- Traceroute servers
  - 2,509 public traceroute servers
  - 507 ASes
  - 77 countries

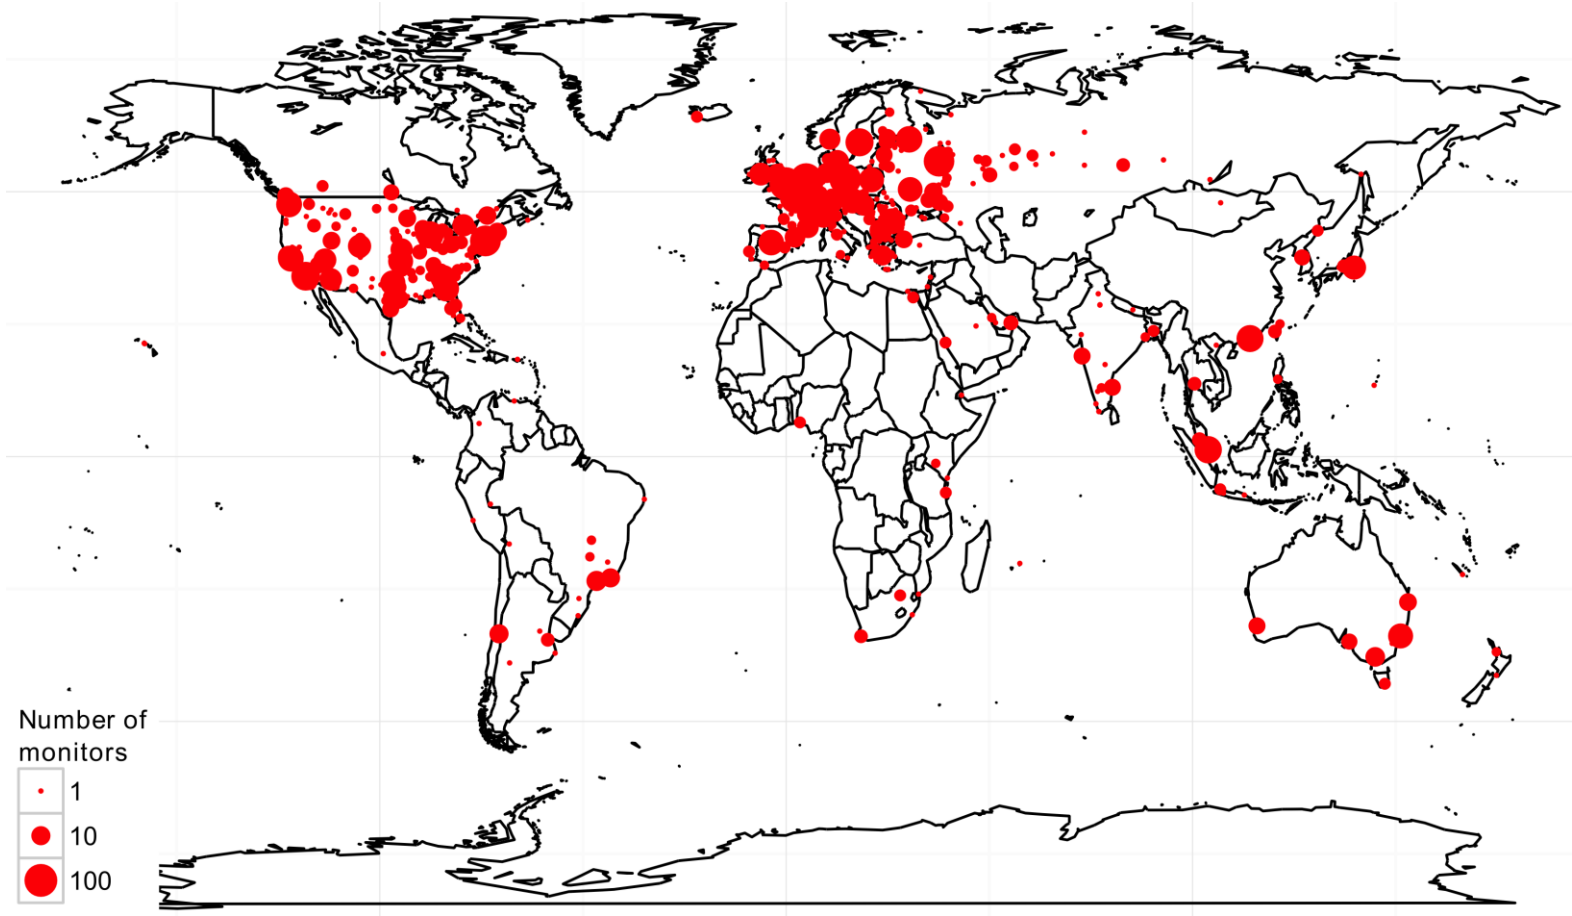

Selecting traceroute Vantage Points (VPs)

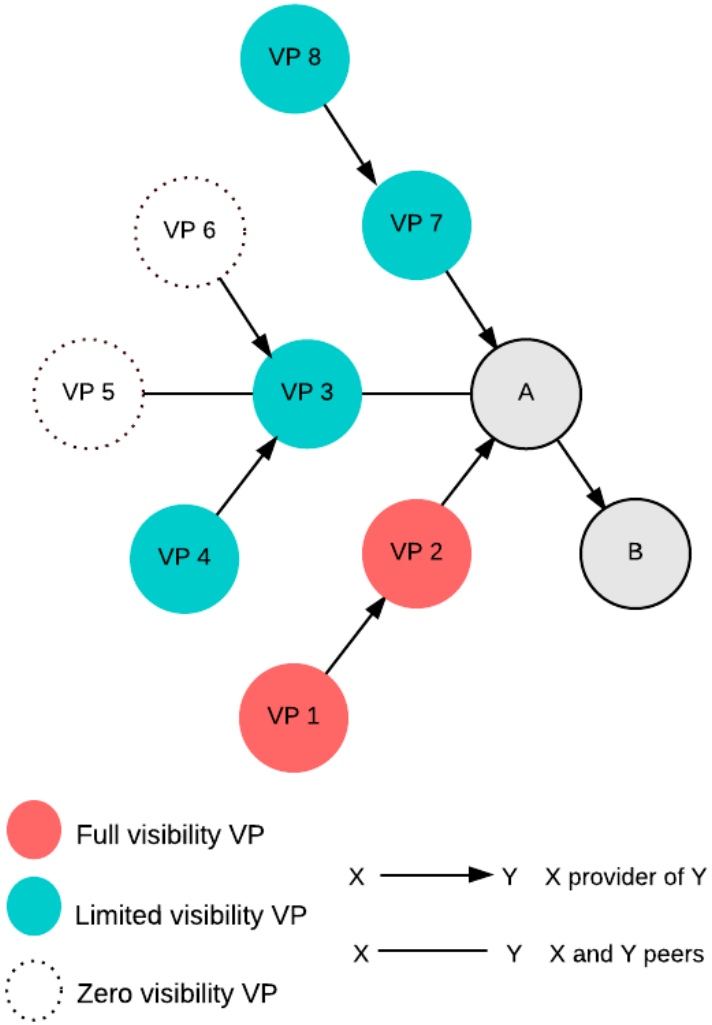

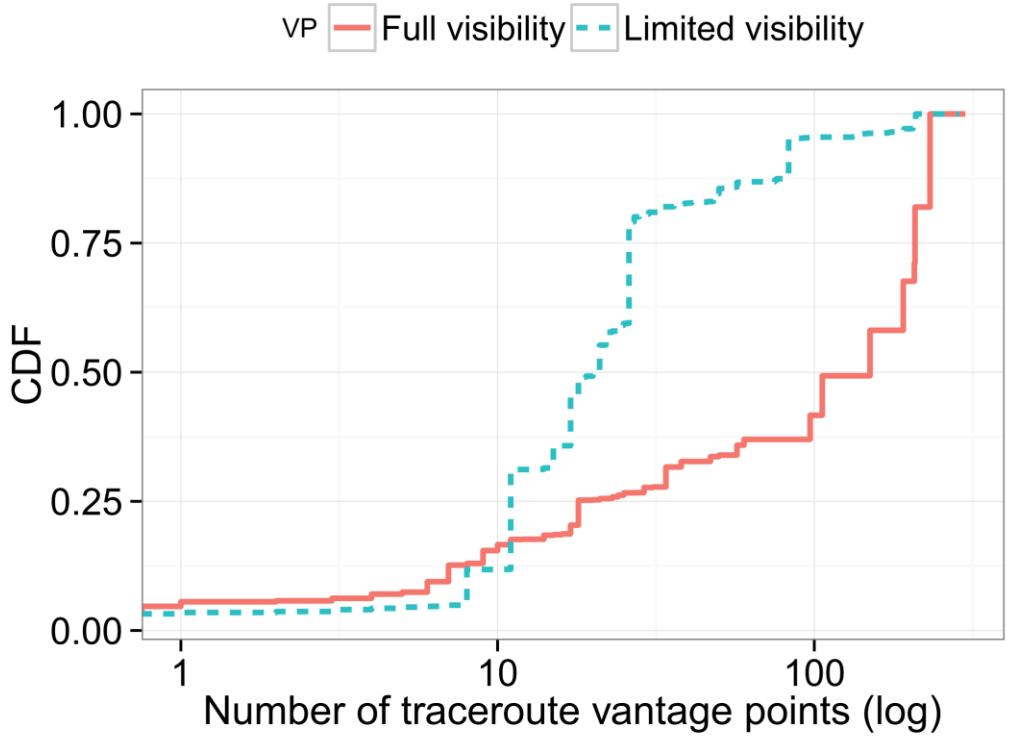

Hybrid relationships

- Two ASes establish different relationship type at different Points of Presence (PoPs)[2]

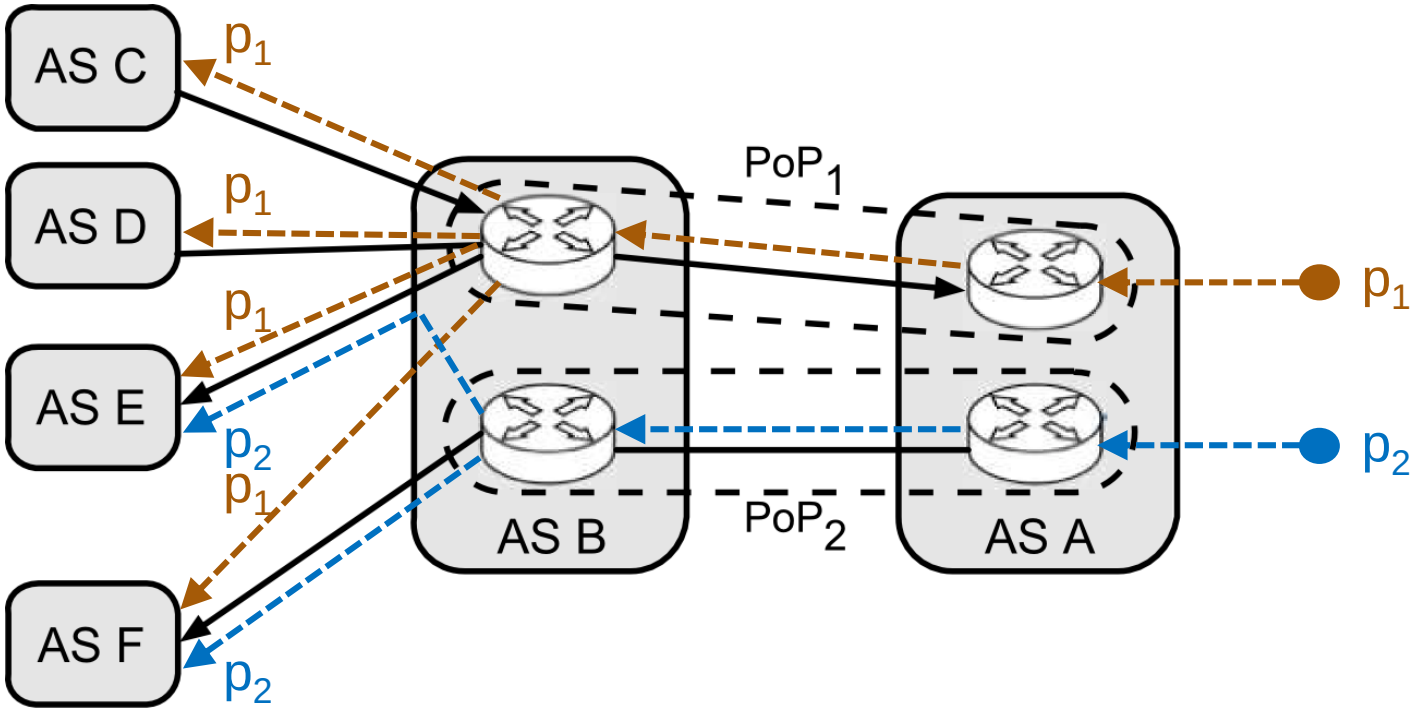

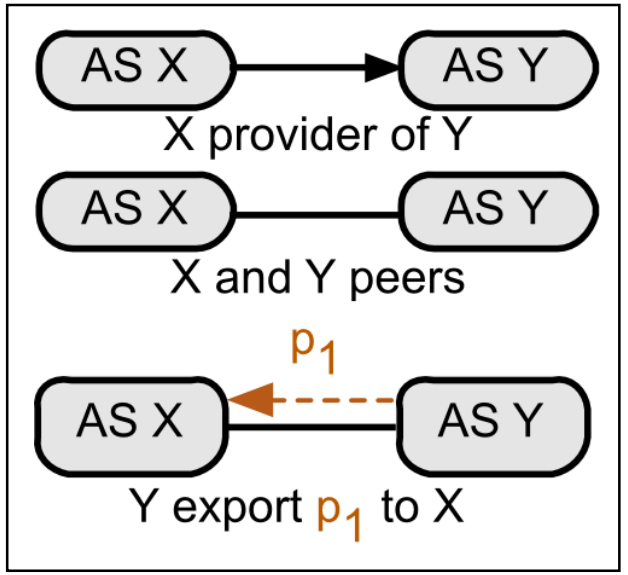

[2] W. Norton. Dual Transit/Peering. http://drpeering.net/white-papers/Art-Of-Peering-The-Peering-Playbook.html#4

Step 4: Label PoPs according to export policies

Full Transit POP

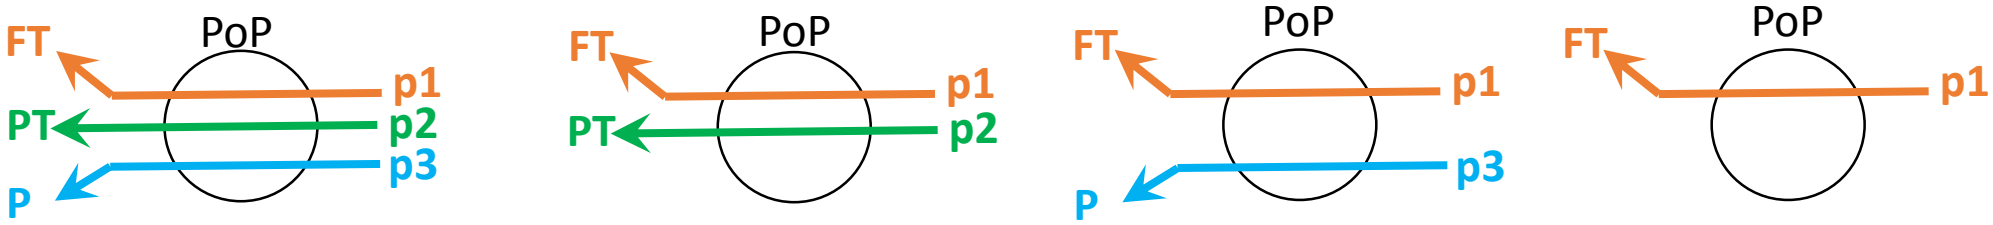

Partial Transit POP

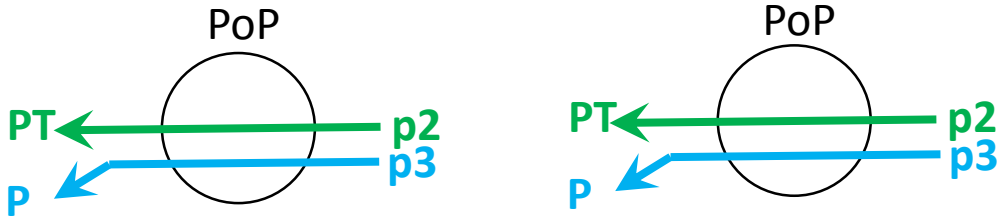

Peering POP

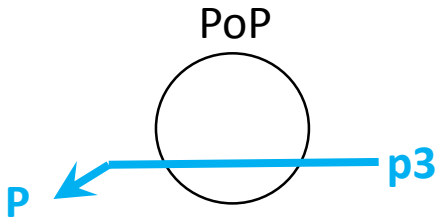

Overall inference methodology

1. Use IMC 2013 algorithm[3] to infer conventional relationships
   - Both *partial transit* and *hybrid* relationships are inferred as ***p2c***

2. For each ***p2c*** link infer provider's *per-prefix* export policy
   - Full Transit
   - Partial Transit
   - Candidate Hybrid
3. For each ***candidate hybrid*** link geo-locate the ingress PoPs
4. Correlate export policies with PoPs

➔If distinct PoPs exhibit distinct export policies infer ***hybrid link***

[3] Luckie, Matthew, et al. "AS relationships, customer cones, and validation ". ACM IMC 2013.