

Group of Border Links (GBL) Used in Internet Multipath Routing

Jie Li¹, Shi Zhou¹, Vasileios Giotsas², Yangyang Wang³

jie.li@cs.ucl.ac.uk

¹ University College London, ² Lancaster University, ³ Tsinghua University

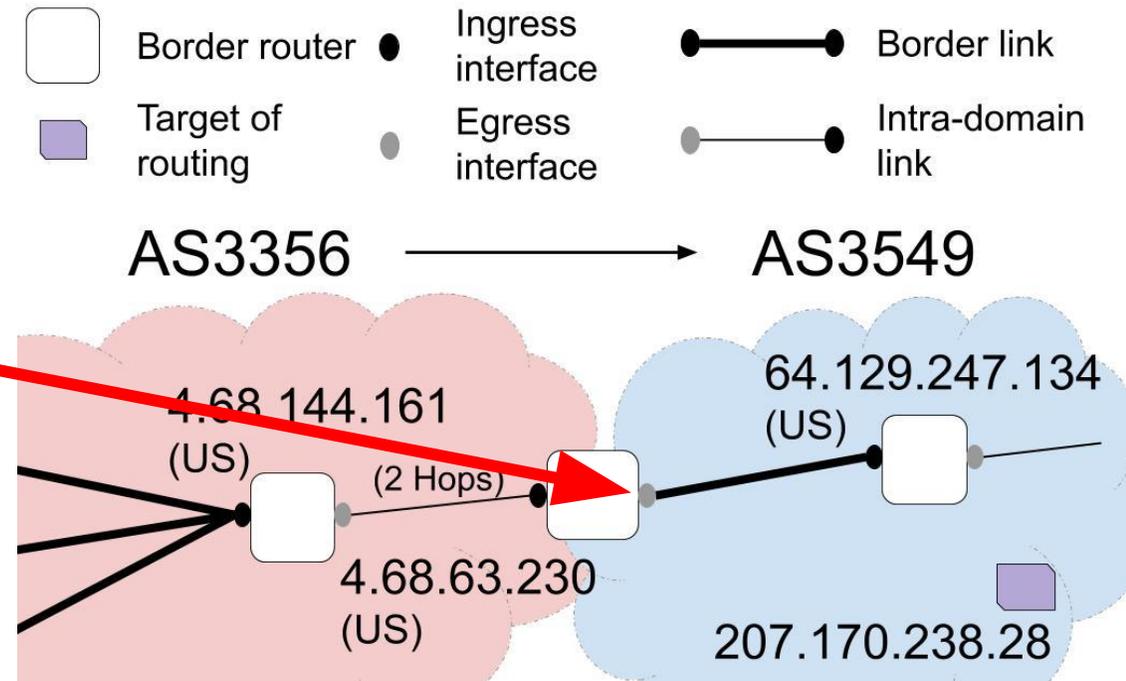


Multipath Routing

- Old assumptions
 - A single inter-domain path between a pair of hosts [1].
 - Different paths due to measurement error, misconfiguration or routing dynamics [2].
- Recent works on multipath routes
 - Multipath BGP [3]
 - Load balancing routing paths [4]
 - Periodic path changes [5]

What are Border Links?

- Physical IP-level link between two routers in different ASes
 - Problem: **Egress interfaces** are invisible in traceroute and difficult to infer.
- A border link is logically represented by two consecutive ingress IP interfaces of border routers



Terms for Group of Border Links (GBL)

- **Host Pair**

- IP addresses of source and target hosts

- **AS Pair**

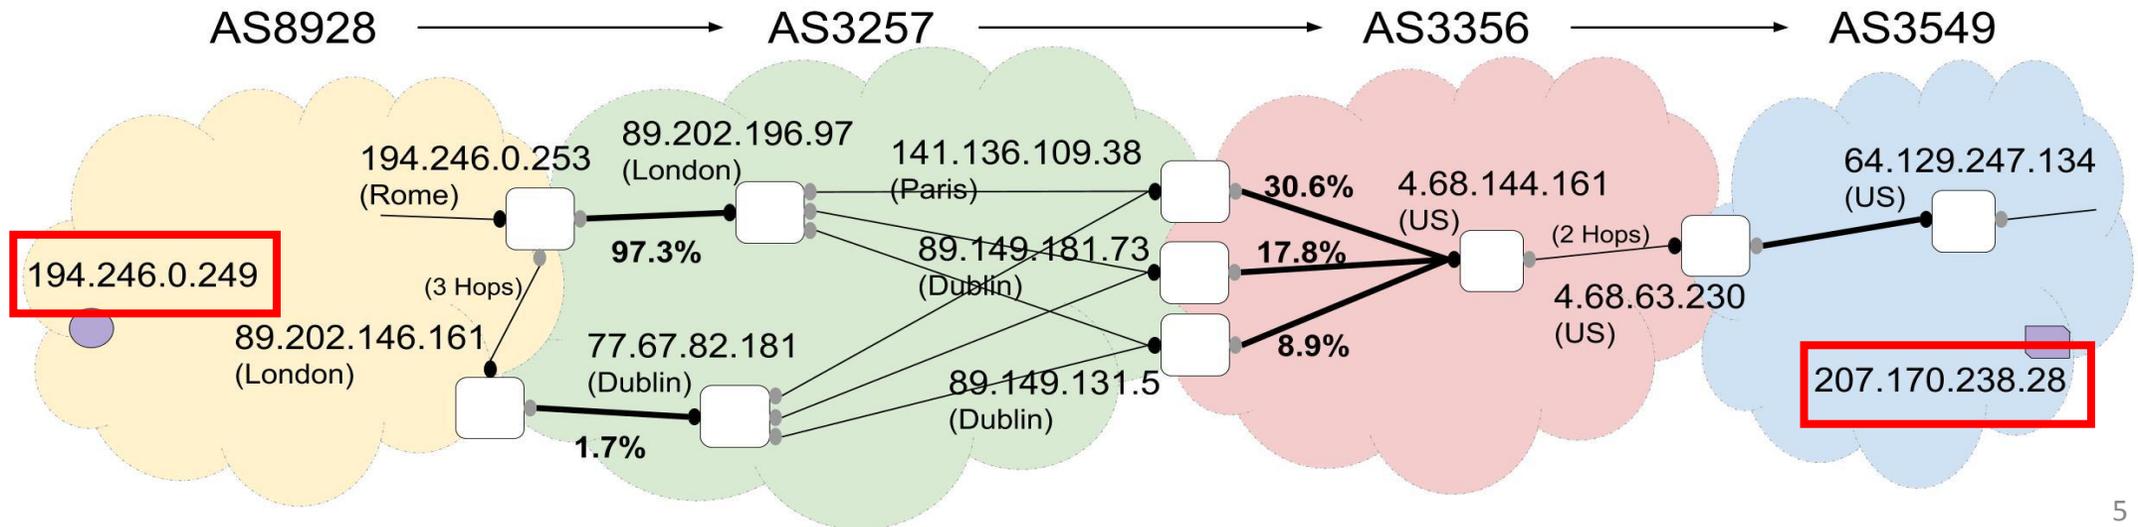
- Two adjacent ASes along an AS-level routing path between a Host Pair
- From near-side AS to far-side AS

- **GBL**

- A **group of border links** are between the **same AS Pair**
- And they are used for traceroute **routing between the same Host Pair**

Examples of GBL

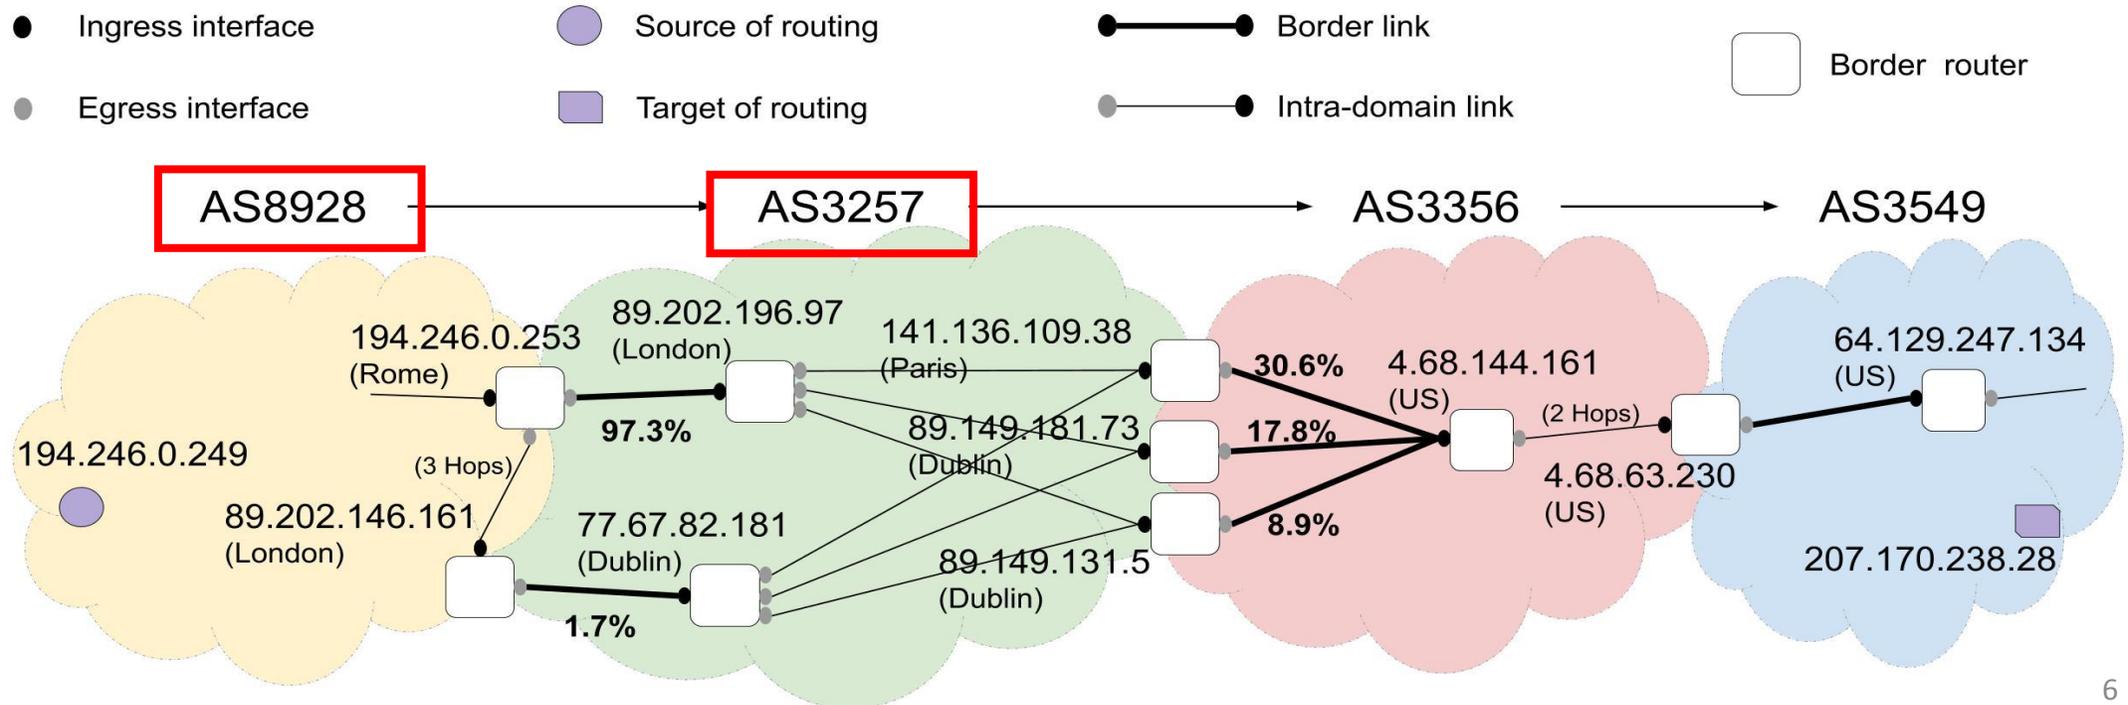
Host Pair: 194.246.0.249 - 207.170.238.28



Examples of GBL

Host Pair: 194.246.0.249 - 207.170.238.28

AS Pair: AS8928-AS3257

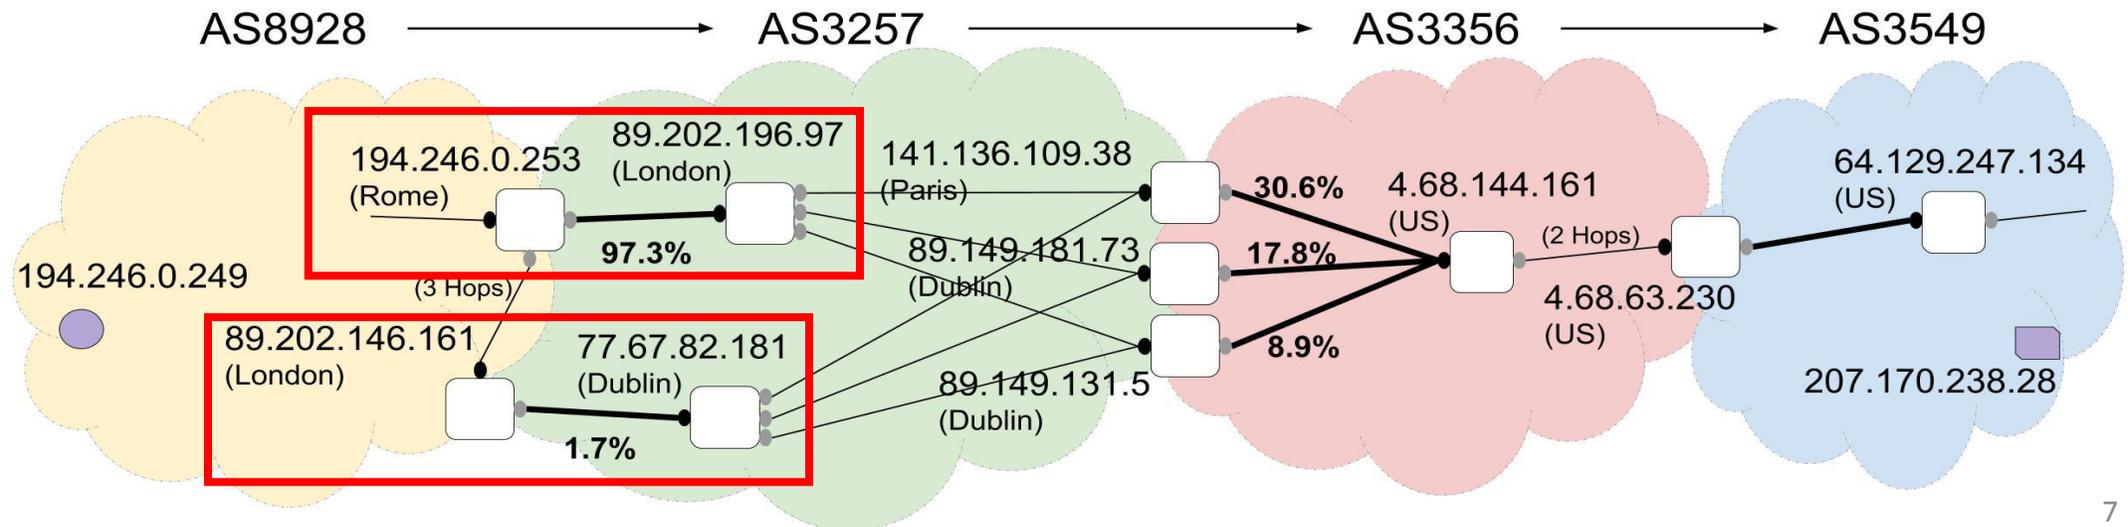


Examples of GBL

Host Pair: 194.246.0.249 - 207.170.238.28

AS Pair: AS8928-AS3257

Border links: 194.246.0.253-89.202.196.97
89.202.146.161-77.67.82.181



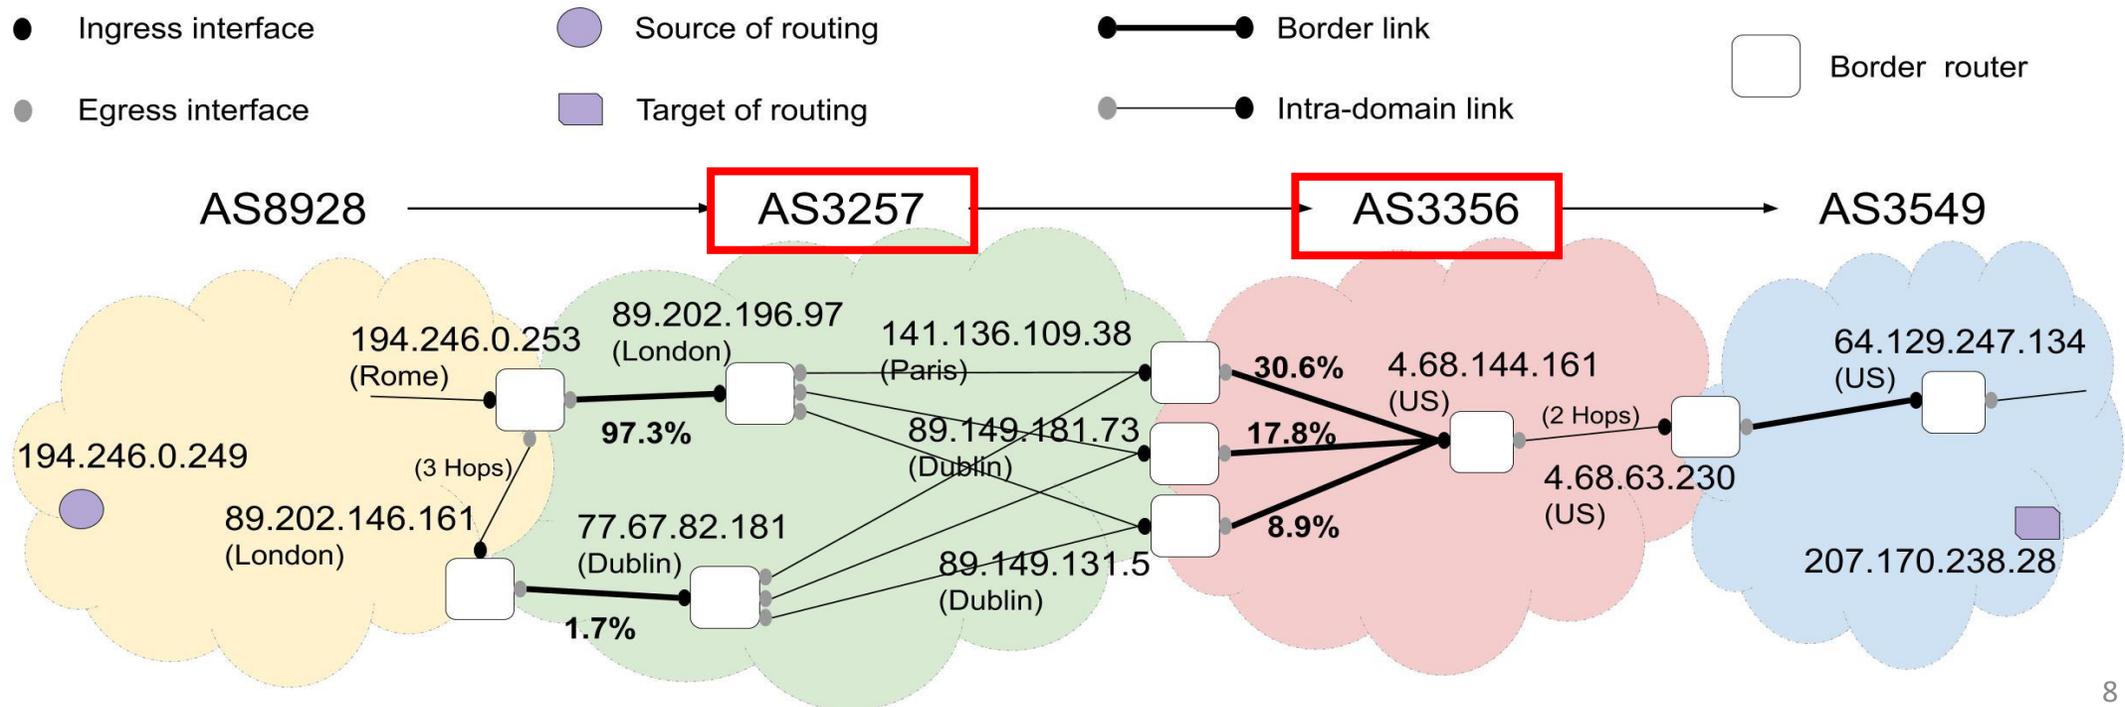
Examples of GBL

Host Pair: 194.246.0.249 - 207.170.238.28

AS Pair: AS8928-AS3257

AS Pair: AS3257-AS3356

Border links: 194.246.0.253-89.202.196.97
89.202.146.161-77.67.82.181



Examples of GBL

Host Pair: 194.246.0.249 - 207.170.238.28

AS Pair: AS8928-AS3257

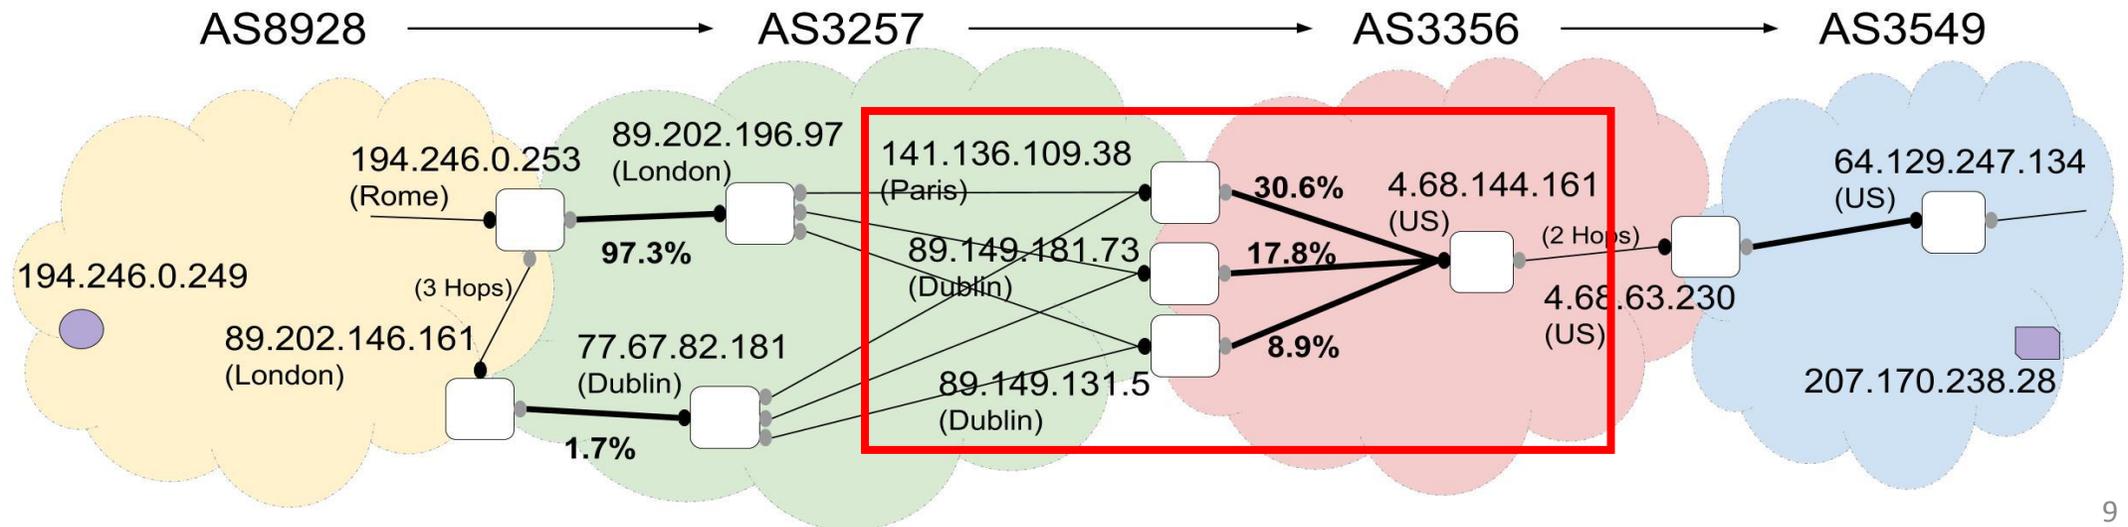
Border links: 194.246.0.253-89.202.196.97
89.202.146.161-77.67.82.181

AS Pair: AS3257-AS3356

Border links: 149.136.109.229-4.68.144.161
89.149.181.73-4.68.144.161
89.149.131.5-4.68.144.161

...

- Ingress interface
- Egress interface
- Source of routing
- Target of routing
- Border link
- Intra-domain link
- Border router



Research Questions

- How to **identify** different cases of GBL in traceroute paths?
- How **frequently** are border links in a GBL used?
- Do border links in a GBL follow certain **connection patterns**?
- Are they used **periodically**?

Our Traceroute Measurement based on RIPE Atlas

- We studied the **top-50 ASes** according to CAIDA's AS-Rank
 - 30 ASes hosted RIPE Atlas probes
 - One probe was chosen per AS
- Issue a traceroute query between each pair of the 30 chosen hosts every 5 minutes for 56 days (>16,000 measurements)



DESTINATION OPTION SIZE 0 (default) Size of an IPv6 destination option header filled with NOPs	DONT FRAGMENT false (default) Do not fragment outgoing packets	DUPLICATE TIMEOUT 10 (default) Time to wait (in milliseconds) for a duplicate response after receiving the first response
FIRST HOP 1 (default) TTL (time to live) of the first hop	HOP BY HOP OPTION SIZE 0 (default) Size of an IPv6 hop-by-hop option header filled with NOPs	MAX HOPS 32 (default) Traceroute measurement stops after the hop at which the TTL reaches this value
PACKETS 3 (default) The number of packets send in a measurement execution. Value must be between 1 and 16. Default is 3	PARIS 16 (default) The number of paris traceroute variations to try. Zero disables paris traceroute. Value must be between 0 and 64	PORT 80 (default) The target port number (TCP only). Defaults to 80
PROTOCOL ICMP (default) Protocol used in measurement	RESPONSE TIMEOUT 4000 (default) Response timeout for one packet	SIZE 48 (default) size of the data part of the packet, i.e. excluding any IP, ICMP, UDP or TCP headers. Value must be between 0 and 2048
TRAFFIC CLASS (No value set and no default defined) The traffic class (IPv6) or type of service and precedence (IPv4) value	RIPE Atlas Default settings	

Identification of Border Links: Methodology

- Step 1: IP-to-AS mapping using **bdrmapIT** [6]
 - Infer which IP hops are in the border of Ases
- Step 2: IP alias resolution of border IPs using **MIDAR** [7]
 - Group border IPs to border

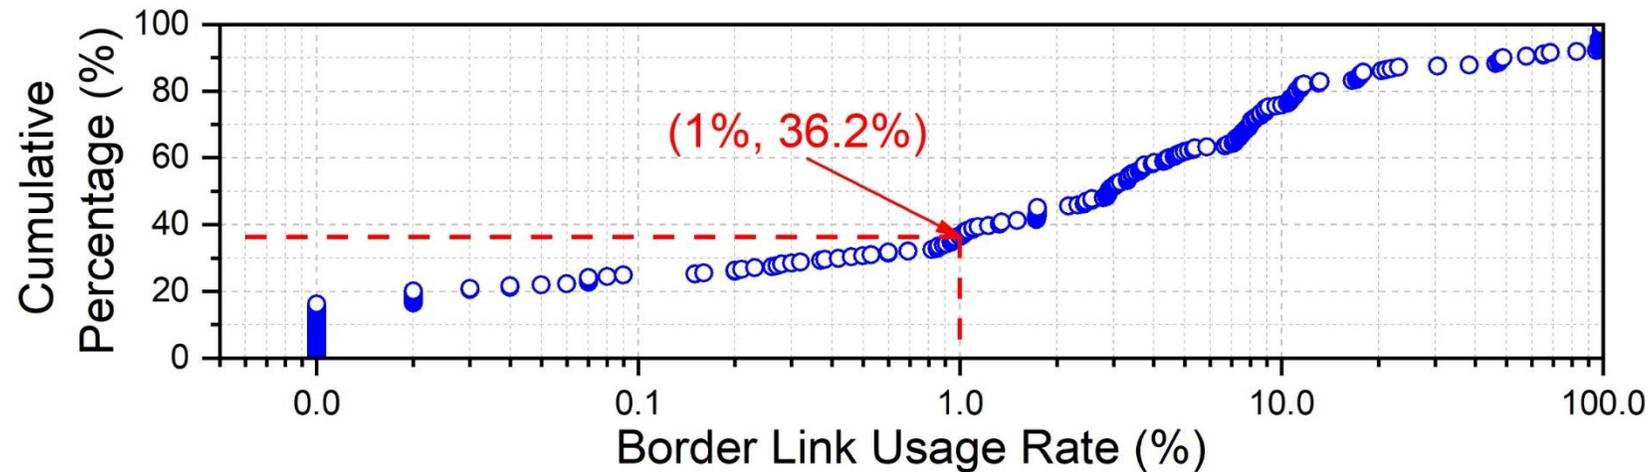
Identification of Border Links: Result

- Applied methodology on traceroute paths
 - 1,203 unique IP addresses
- Border mapping output
 - 249 border links with 267 border IP addresses
- Alias resolution result
 - 227 border links between 242 border routers

Identification of Group of Border Links (GBL)

- Usage rates of border links

- We only consider border links with usage rate $>1\%$
- 121 border links between 13 AS Pairs for routing between 13 different Host Pairs



22 cases of GBL in our data

- Each case contains 2 to 32 border links
 - Each link is with >1% usage rate
- AS3257 (GTT) is relevant to 19 cases

Number of border links	Number of cases
2	12
3	1
5	1
8	1
9	3
14	2
24	1
32	1
Total	22

22 cases of GBL in our data

- Each case contains 2 to 32 border links
 - Each link is with >1% usage rate
- AS3257 (GTT) is relevant to 19 cases

Top50 Rank	AS Number	AS Name	# of Src.	# of Tgt.	# of Near-side	# of Far-side
1	3356	LEVEL3				1
2	1299	TELIANET			1	2
5	3257	GTT			9	10
8	6939	HURRICANE		3		1
9	3491	BTN		1		1
10	3549	LVLT		2		
11	1273	CW	3		3	
12	6461	ZAYO		2		1
16	3320	DTAG		3		2
18	12389	ROSTELECOM				1
29	286	KPN		2		1
32	6830	LGI-UPC		2		1
39	8928	INTERROUTE	18	2	9	1
46	4134	CHINANET		1		
48	8220	COLT	1	2		
50	29076	CITYTELECOM		2		
		Total	22	22	22	22

Validation of GBL

- Tool: BGP Looking Glasses
 - Command: `show ip bgp <destination>`
 - Check output for multipath eBGP routes for each traceroute destination
- Validation of True Positives
 - Query the publicly accessible LGs in 3 of the ASes in our GBL cases
 - All of them confirmed the usage of multipath eBGP, which covered 13 out of the 22 cases of identified GBL.
- Validation of True Negatives
 - When querying destinations where no GBL were observed, none of the routes were denoted as multiple external.

Analysis of GBL usage and connectivity

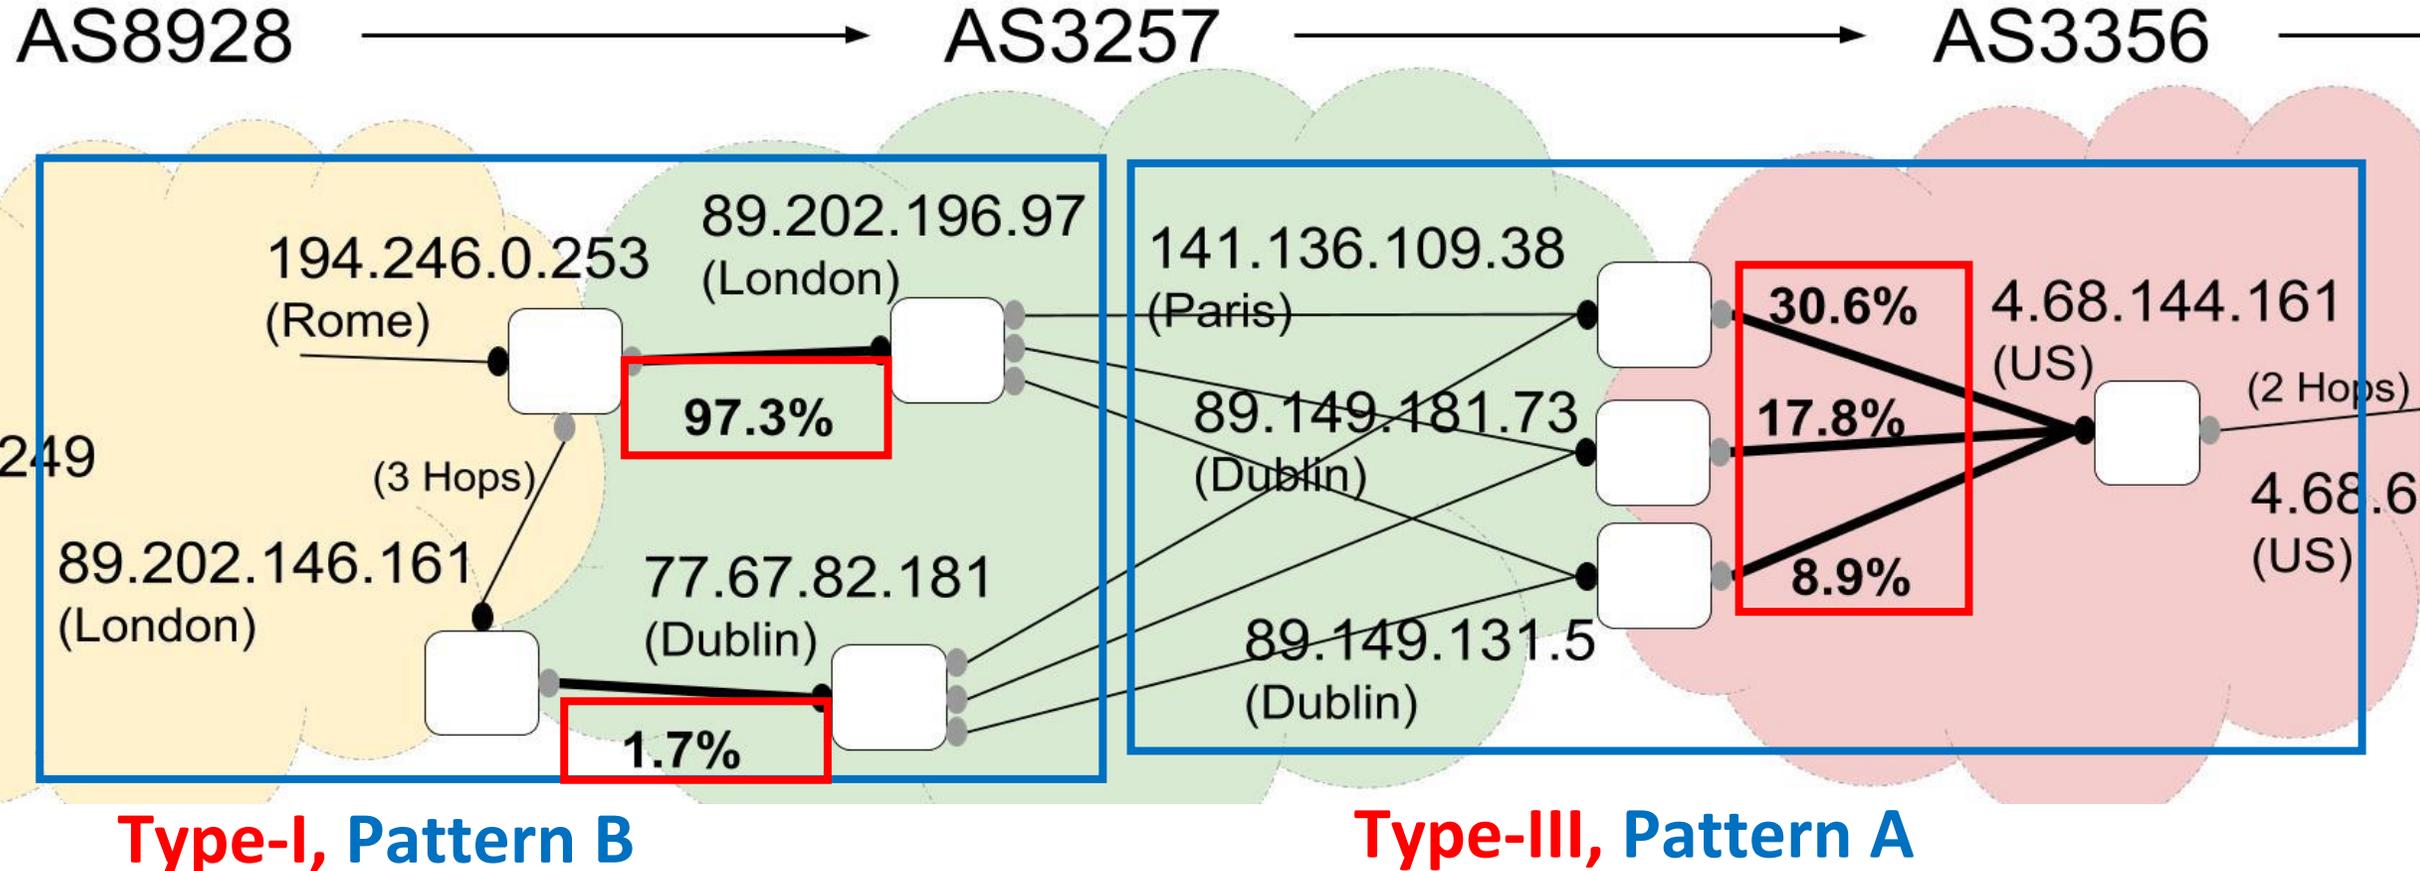
•Types of GBL by usage rate

- Type-I: 9 cases
 - Difference between the highest 2 usage rates is >20%.
- Type-II: 6 cases
 - Difference between the highest usage rate and the lowest usage rate is < 5%.
- Type-III: 7 cases
 - Other cases.

•Connection **Patterns** of GBL

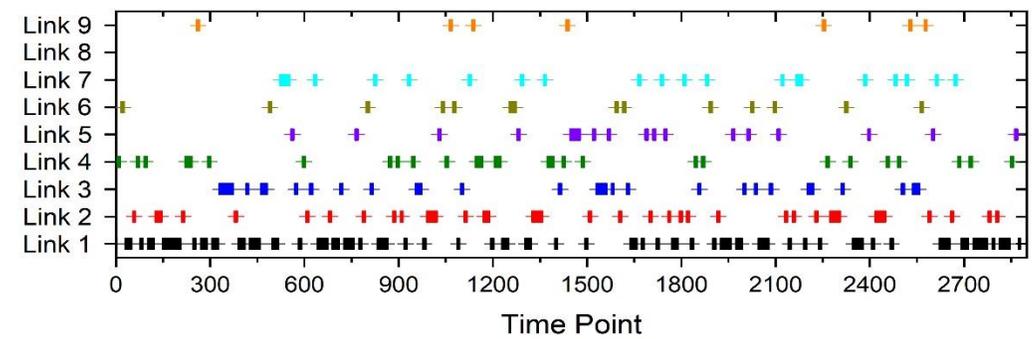
- We observed that in each case of GBL, all border links **started from different** border routers.
- Pattern A: 10 cases
 - **Ended** at the **same** border router.
- Pattern B: 12 cases
 - **Ended** at **different** border routers.

Analysis of GBL usage and connectivity (Examples)

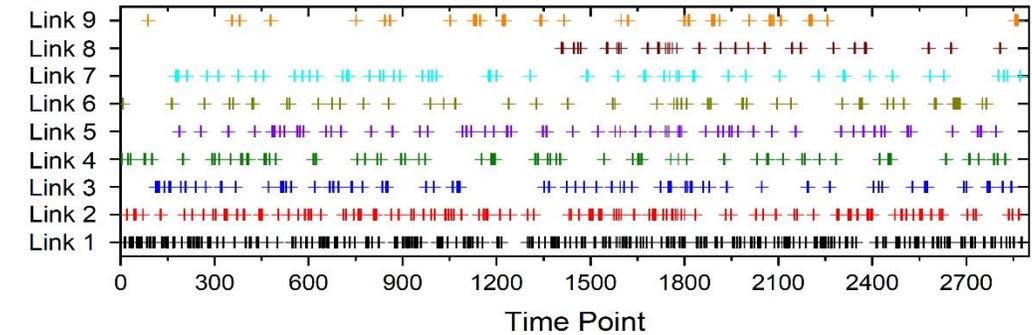


A Case Study

- Type-III, Pattern A, nine border links
- AS Pair
 - AS3257(GTT)-AS3356(Level3)
- Host Pair
 - Source: 194.246.0.249 in AS8928
 - Target: 207.170.238.28 in AS3549
- Border links were persistently and stochastically used.

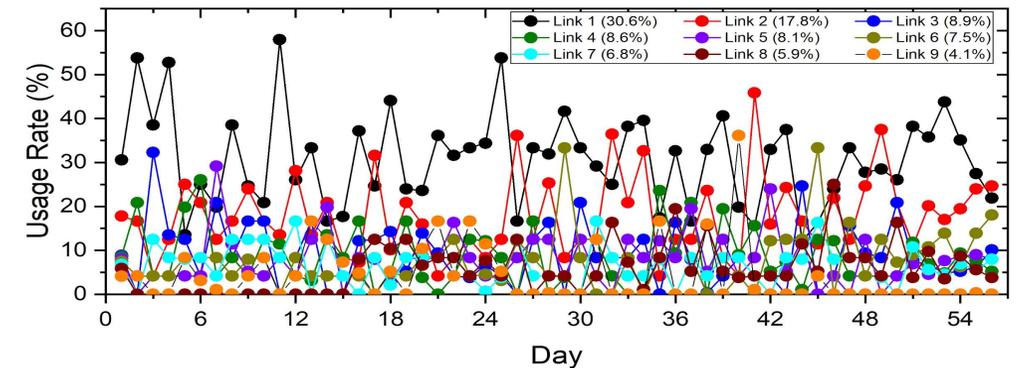


(a) 5-minute interval over 10 days



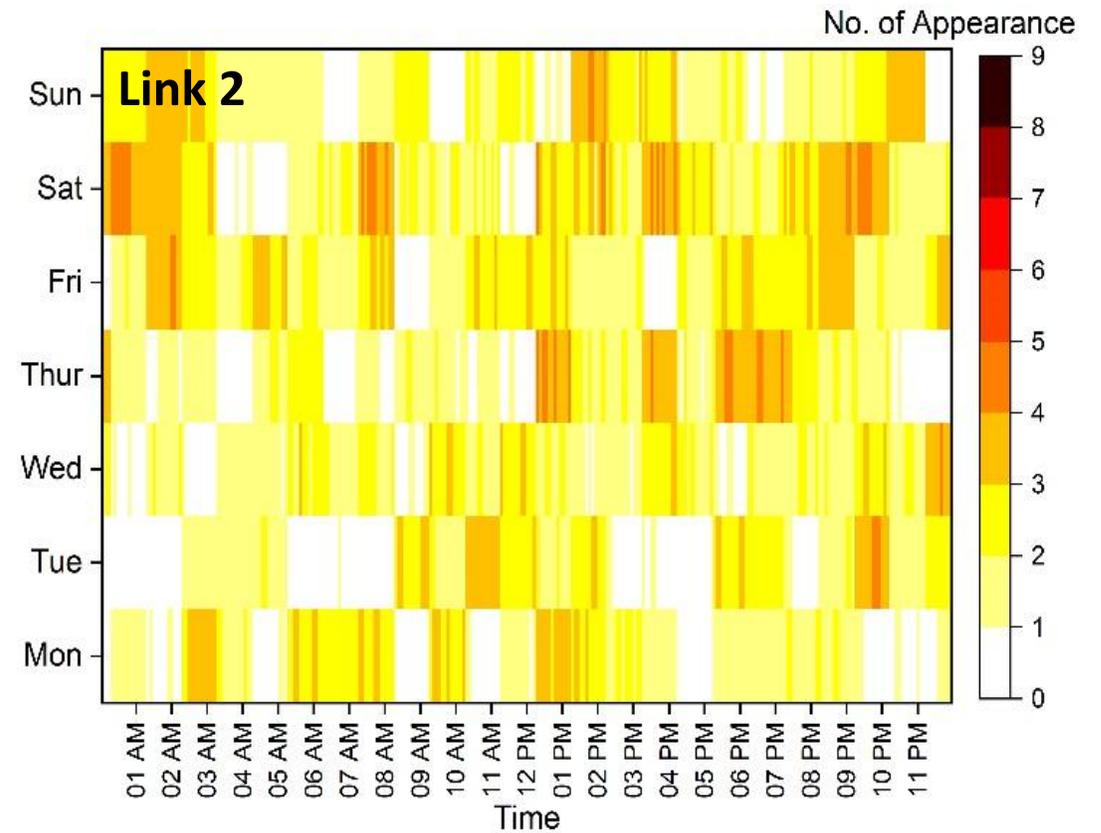
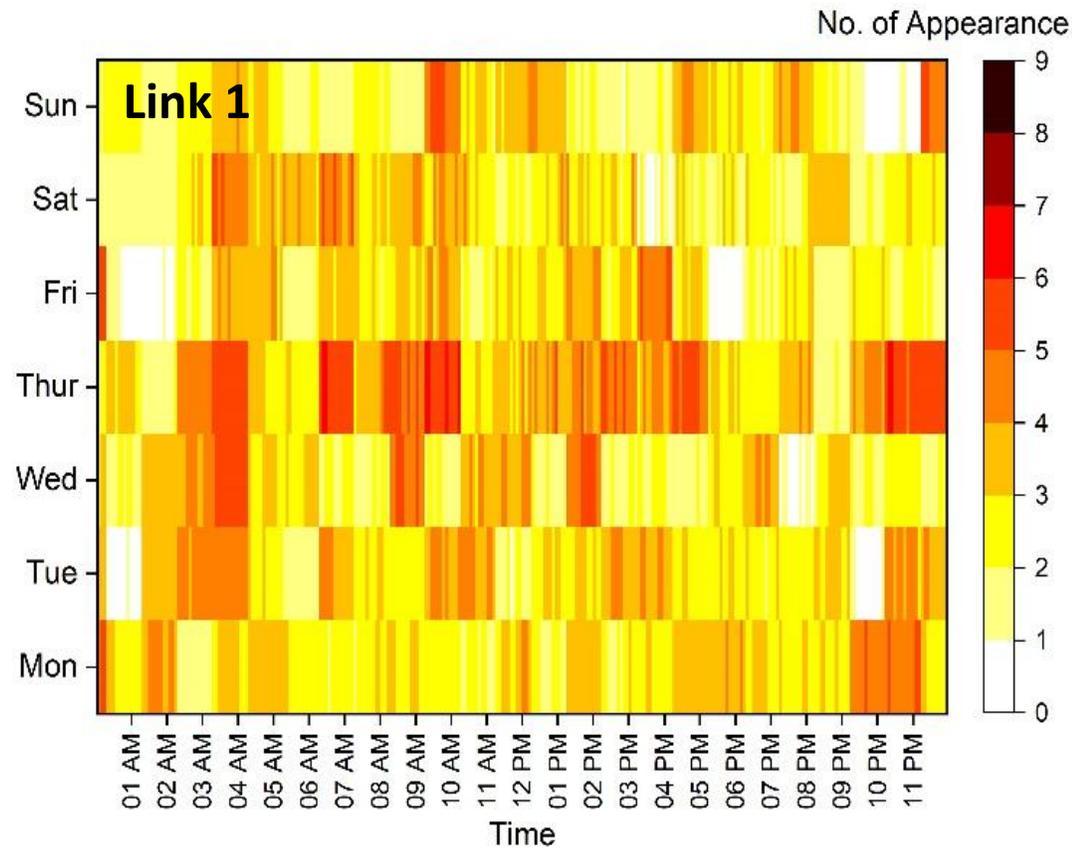
(b) 15-minute interval over 30 days

Observation of Link 1 and Link 2 in our measurements



Daily usage rates of the nine links over 56 days

A Case Study (continued)



- Appearance frequency distribution of a link as a function of the hour and the week day as measured over 56 days.
- Border links with top 2 usage rates are shown.
- There was no periodical pattern.

Discussion

- Our observation
 - Group of border links used for routing between the same Host Pair.
 - Many cases were observed in a relatively small-scale measurement.
 - 5-minute interval traceroute revealed more details.
 - Potential impact on network resilience, performance, economics, etc.
- Future works
 - More measurements for top 200 ASes
 - 15-min data from RIPE
 - Hourly data from CAIDA
 - Analysis on long time periods, e.g years
 - Per-destination, per-flow analysis

Thank you!

Jie Li
jie.li@cs.ucl.ac.uk

- [1] N. Ahmed and K. Sarac, “An experimental study on inter-domain routing dynamics using IP-level path traces,” in Proc. IEEE ICN’15, pp. 510-517.
- [2] G. Comarela, G. Gursun, and M. Crovella, “Studying interdomain routing over long timescales,” in Proc. ACM IMC’13, pp. 227–234.
- [3] Petr Lapukhov. “Equal-cost multipath considerations for BGP,” Internet Engineering Task Force. Network Working Group Internet Draft.
<https://tools.ietf.org/id/draft-lapukhov-bgp-ecmp-considerations-02.html>. July 2019.
- [4] B. Augustin, T. Friedman, and R. Teixeira, “Measuring multipath routing in the Internet,” IEEE/ACM Trans. Netw. vol. 19, no. 3, pp. 830-840, June 2011.
- [5] M. Iodice, M. Candela, and G. Di Battista, “Periodic path changes in RIPE Atlas,” IEEE Access. vol. 7, pp. 65518-65526, 2019.
- [6] A. Marder, M. Luckie, A. Dhamdhere, B. Huffaker, kc claffy, and J. M. Smith, “Pushing the boundaries with bdrmapIT: Mapping router ownership at Internet scale,” in Proc. ACM IMC’18, pp. 56-69.
- [7] K. Keys, Y. Hyun, M. Luckie, and K. Claffy, “Internet-scale IPv4 alias resolution With MIDAR,” IEEE/ACM Trans. Netw., vol. 21, no. 2, pp. 383–399, April 2013.