

# Data and analysis on CAIDA's 100G packet traces

Ricky Mok, Brendon Jones, Daniel Andersen, and kc claffy  
(CAIDA)



This work is partially supported by NSF CNS-2120399


# 100G Traffic traces

Sample 1-hour traffic (headers-only) once per month from a link between two cities in a Tier 1 ISP

IPs are anonymized using CryptoPAn

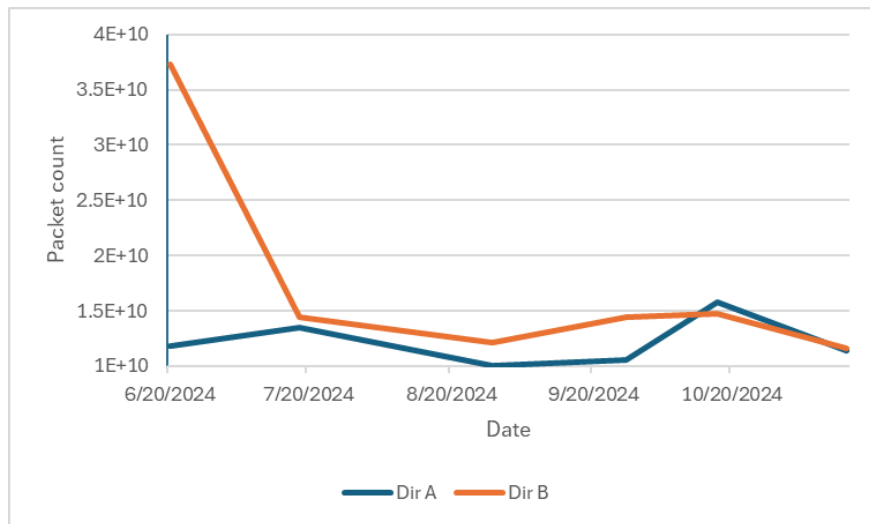
[https://catalog.caida.org/dataset/passive\\_2024\\_pcap\\_100g](https://catalog.caida.org/dataset/passive_2024_pcap_100g)

[https://catalog.caida.org/dataset/passive\\_2025\\_pcap\\_100g](https://catalog.caida.org/dataset/passive_2025_pcap_100g)



# Traffic statistics

## Los Angeles <> San Jose



## Los Angeles <> Dallas



[https://www.caida.org/catalog/datasets/100g\\_trace\\_stats/](https://www.caida.org/catalog/datasets/100g_trace_stats/)

# Combined view with other data

However, IP anonymization prohibited us from joining other datasets.

- prefix to AS
- IP geolocation
- AS-org

We are planning to provide additional metadata for the trace

- Anonymized prefix-to-AS mapping table
- IP geolocation to the anonymized addresses

# Roll back in time

CAIDA collected similar traces in the same ISP in the past

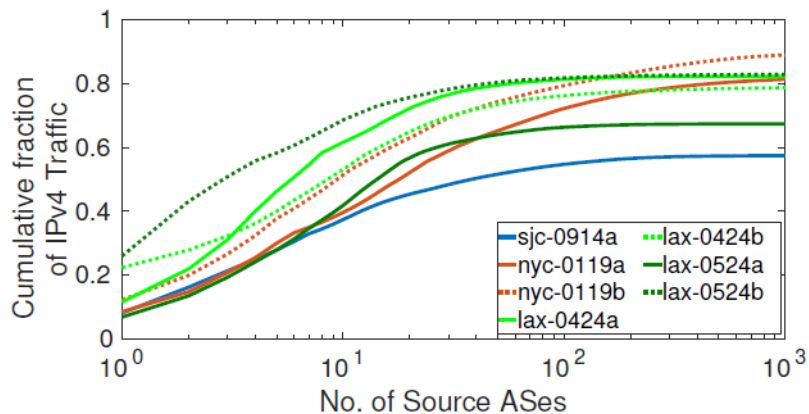
Traces	# of identified ASes (% of total IPv4/IPv6 traffic)					
	Source		Destination		Next hop	
	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6
sjc-0914a	16,924 (57.5%)	124 (24.8%)	4,120 (57.5%)	148 (24.8%)	229	75
nyc-0119a	27,481 (81.9%)	1,158 (65.0%)	671 (82.0%)	73 (65%)	65	33
nyc-0119b	4,038 (89.1%)	222 (6.63%)	41,801 (89.1%)	2,164 (6.73%)	1,401	564
lax-0424a	20,002 (82.2%)	434 (99.4%)	2668 (82.3%)	1019 (99.4%)	284	136
lax-0424b	39,822 (79.0%)	1,852 (98.4%)	5,160 (79.0%)	624 (98.4%)	405	169
lax-0524a	3,607 (67.4%)	445 (98.6%)	2,617 (67.4%)	897 (98.7%)	286	128
lax-0524b	3,361 (82.8%)	712 (99.1%)	2,031 (82.8%)	275 (99.1%)	264	94

NYC<>BR

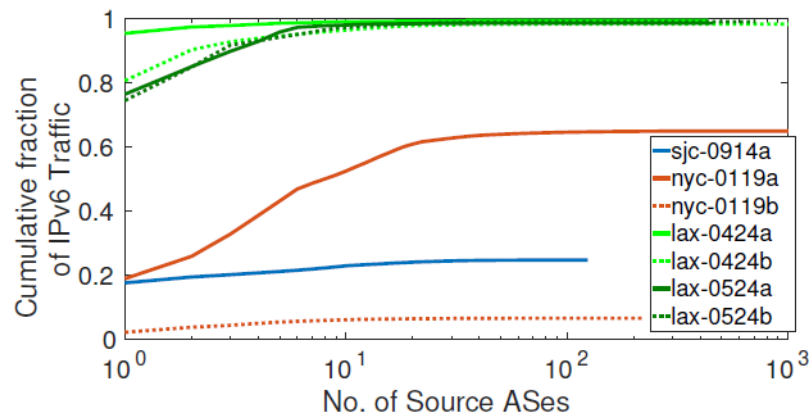
# Changes in traffic sources

The top ASes dominated in traffic volume

More prominent in IPv6



(a) Top 10 IPv4 source ASes was responsible for 41.8% to 68.6% of total IPv4 traffic of the lax traces.



(b) Top 10 IPv6 ASes was accounted for over 96.5% of the total IPv6 traffic.

# Traffic country origin

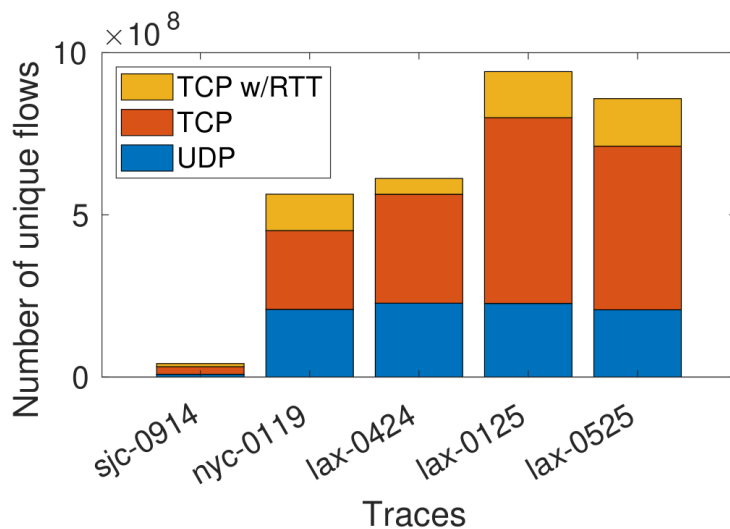
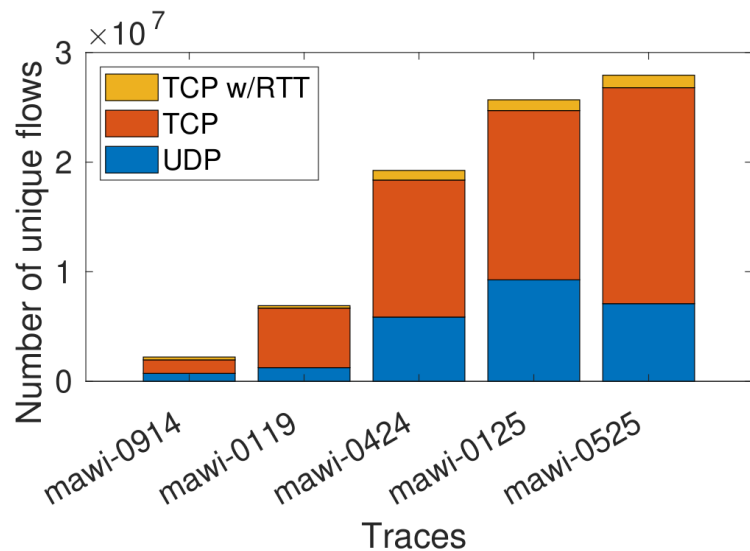
We studied the IP geolocation of endpoints between MAWI trace and CAIDA's trace (Kudo: Kenjiro Cho)

Traces	Top 3	
	Src-Dst pairs (% of connections)	Src-Dst pairs (% of traffic)
mawi-0914	JP→US (21.2%), JP→CN (13.4%), JP→DE (5.96%)	KR→JP (33.6%), JP→US (22.1%), JP→CN (20.9%)
sjc-0914a	US→US (21.4%), CN→US (16.5%), MX→US (15.0%)	CN→US (42.9%), US→US (26.8%), KR→US (8.18%)
mawi-0119	JP→US (16.5%), JP→CN (10.7%), JP→BR (6.28%)	JP→US (54.0%), JP→JP (15.5%), SG→JP (7.56%)
nyc-0119x	BR→US (31.9%), AR→US (11.5%), BR→CN (2.76%)	PY→US (47.8%), BR→BR (20.1%), US→AR (18.6%)
mawi-0424	JP→US (14.7%), US→JP (14.3%), JP→CN (8.00%)	JP→JP (19.2%), US→JP (13.8%), SG→JP (12.4%)
lax-0424	US→CN (16.0%), US→KR (14.9%), US→US (9.09%)	US→US (33.8%), US→CN (17.5%), US→KR (5.51%)
mawi-0125	US→JP (18.0%), JP→US (14.9%), JP→CN (7.64%)	JP→JP (34.8%), US→JP (16.3%), JP→US (13.6%)
lax-0125	JP→US (10.6%), US→JP (6.54%), BR→US (6.15%)	JP→US (13.4%), US→US (12.4%), US→SG (11.8%)
mawi-0525	JP→US (15.5%), US→JP (14.7%), JP→CN (6.95%)	JP→JP (45.8%), US→JP (17.9%), JP→US (13.4%)
lax-0525	JP→US (9.50%), US→JP (9.4%), US→US (4.14%)	US→MX (12.0%), US→US (9.25%), JP→US (8.69%)

# Passive RTT analysis

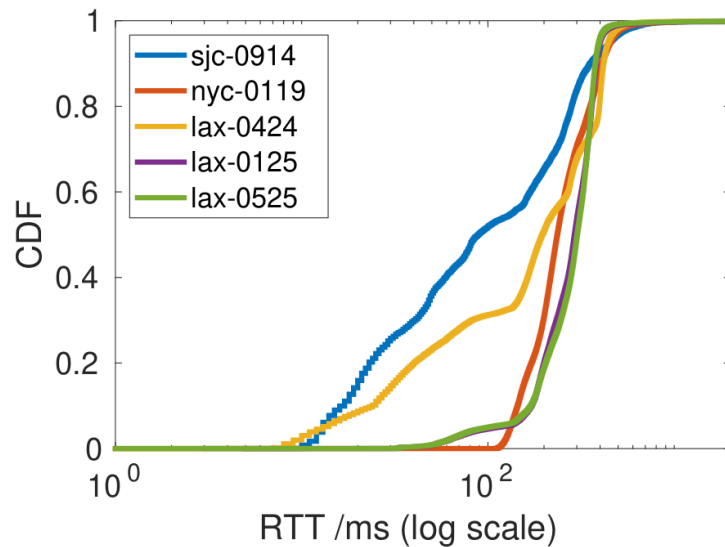
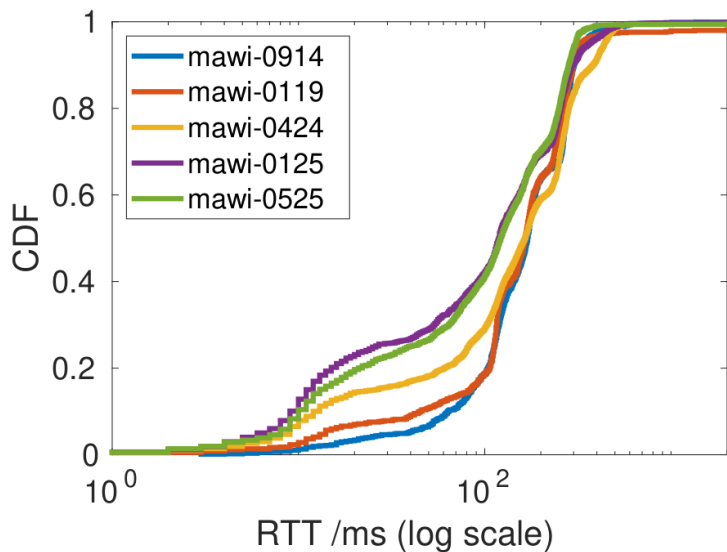
SYN-and-ACK time

CAIDA's traces are more likely to be able to capture RTTs



# Passive RTTs

RTTs observed in MAWI traces decreased over time. Recent CAIDA traces showed a long RTT, as it serves high percentage of international traffic.



# Conclusion

We are straightening out the methods and data to facilitate research with the traces (without sacrificing too much privacy)

