CAIDA passive measurement infrastructure

Emile Aben <emile@caida.org>

10th CAIDA-WIDE workshop

Aug 15-16 2008, Marina Del Rey, CA, US

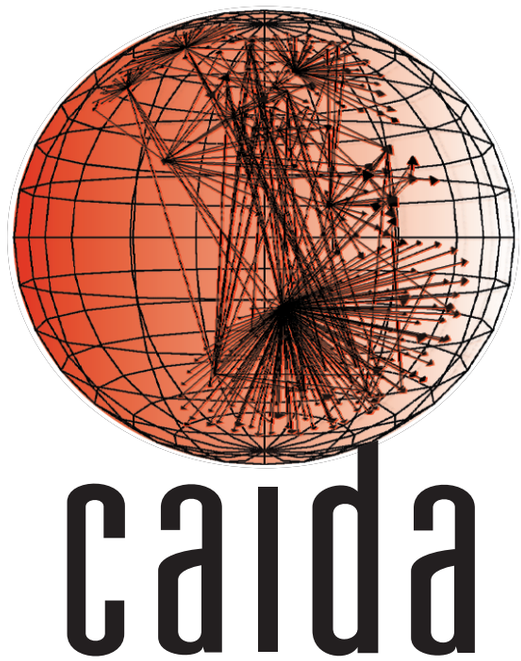

Outline

- Goal
- Problems (EOT)
- Deployments
- Measurements
- Link/Traffic characterization

Goal of passive measurement infrastructure

- Deliver needed data sets to the scientific community studying the Internet, while facing the tremendous operational, economic, and policy barriers (from CAIDA annual report 2007)

Goal of passive measurement infrastructure

- Deliver needed data sets to the scientific community studying the Internet, while facing the tremendous operational, economic, and policy barriers (from CAIDA annual report 2007)
- Data sets = traffic traces from:
  - Critical infrastructure like DNS (OARC)
  - Empty IP space: Network Telescope
    - http://www.caida.org/data/passive/network_telescope.xml
  - Commercial Internet backbone
  - R&D Internet backbone (different?)

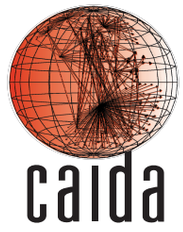

Hurdles: Economics, Ownership and Trust

Hurdles: Economics, Ownership and Trust

- Economics:
  - US Internet Backbone mostly OC192/10GE, and some OC768 already deployed
  - Expensive hardware for capturing packets
  - Lots of data to manage

Hurdles: Economics, Ownership and Trust

- Economics:
  - US Internet Backbone mostly OC192/10GE, and some OC768 already deployed
  - Expensive hardware for capturing packets
  - Lots of data to manage
- Ownership:
  - Network owners have little or no incentive to provide real data, but have legal/privacy concerns
  - Delicate: Can't name our commercial partners

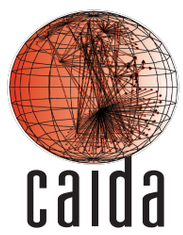

Hurdles: Economics, Ownership and Trust

- Economics:
  - US Internet Backbone mostly OC192/10GE, and some OC768 already deployed
  - Expensive hardware for capturing packets
  - Lots of data to manage
- Ownership:
  - Network owners have little or no incentive to provide real data, but have legal/privacy concerns
  - Delicate: Can't name our commercial partners
- Trust:
  - Give researchers access to data in a way that protects privacy

caida

PREDICT legal framework

- PREDICT is DHS experiment in solving EOT problems
- http://www.predict.org/
- CAIDA participates as data host and data provider
- Operational but not used much yet

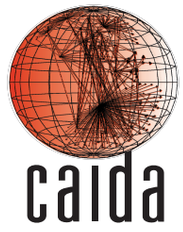

Outline

- Goal
- Problems (EOT)
- Deployments
- Measurements
- Link/Traffic characterization

Hardware testing

- Endace DAG6 cards ($$), in high-end server hardware
- Goal: less then 1% loss on a fully loaded OC192 link
- 2 DAG6 cards in single machine: heat dissipation becomes a big issue
- 2 separate machines with one DAG6 card each

Current deployments on Tier1 ISP backbone links

Current deployments on Tier1 ISP backbone links

- *equinix-chicago*:
  - March 2008
  - Seattle, WA <=> Chicago, IL
  - OC192

Current deployments on Tier1 ISP backbone links

- *equinix-chicago*:
  - March 2008
  - Seattle, WA <=> Chicago, IL
  - OC192
- *equinix-sanjose*:
  - Currently configuring/having hardware problems
  - Los Angeles, CA <=> San Jose, CA
  - 1 of 2 OC192s (flow load-balanced)

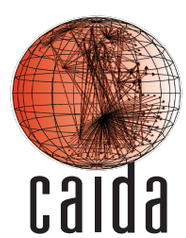

Current deployments on Tier1 ISP backbone links

- *equinix-chicago*:
  - March 2008
  - Seattle, WA <=> Chicago, IL
  - OC192

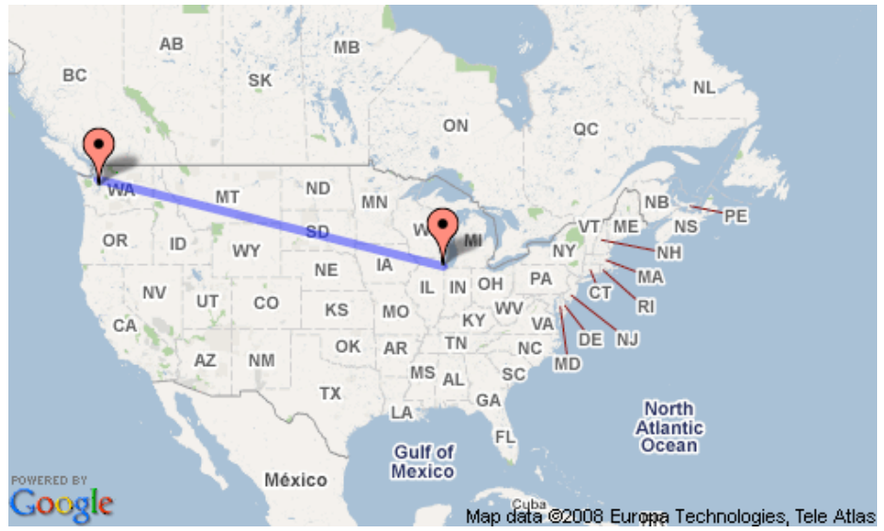

- *equinix-sanjose*:
  - Currently configuring/having hardware problems
  - Los Angeles, CA <=> San Jose, CA
  - 1 of 2 OC192s (flow load-balanced)

caida

Future deployment: internet2-chicago

- ... in the works
- Infinera (layer 1 magic) might allow us to switch between links without touching fiber
- Internet2 backbone link and/or Internet2 peering link
- Only aggregated reporting for now (Report generator)

Measurements / Data

- Report generator
  - Provides insight into current status / trends
- Monthly traces
  - Allows for more detailed analysis

Report Generator (1)

August 06 2008 - August 13 2008 UTC

Application	Min	Avg	Max
HTTP	302.13M	1.68G	3.57G
UNKNOWN_TCP	111.29M	498.49M	942.76M
UNKNOWN_UDP	35.74M	149.87M	340.80M
NNTP	8.44M	54.93M	139.63M
HTTPS	5.45M	22.24M	51.99M
ABACAST	3.01M	20.53M	44.05M
RTSP	1.90M	14.18M	39.17M
GNUTELLA	2.96M	13.20M	29.90M
SMTP	2.87M	13.16M	31.40M
BITTORRENT	2.72M	11.55M	22.03M
IPSEC	1.67M	11.01M	22.50M
other	6.69	110.97M	225.53M

generated 2008-08-14 19:16 UTC

Report generator (2)

- Example: http://www.caida.org/data/realtime/passive/?monitor=equinix-chicago-isp1-B
- Part of CoralReef
- List of installations at: http://www.caida.org/data/realtime
- On OC192 uses adaptive netflow (paper:Building a Better Netflow)

Monthly passive traffic traces

- Capture 64 bytes per packet => 1 hr = 30-120 GB of data per direction, in DAG format
- Strip payload from trace
  - This also removes packets with unknown encapsulation (0.001 % of pkts)
  - Non-payload trace transferred and kept at CAIDA
- Anonymize trace (Crypto-PAn prefix preserving anonymization), strip layer2 and convert to PCAP
  - Anonymized traces available to external researchers under conditions

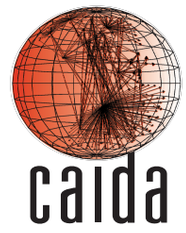

Distribution to external researchers

- Academics and CAIDA members can get access to anonymized passive traffic traces
- Have to fill out data request form
  - Approve of AUP
    - Do not reverse engineer anonymization
- Data requests are vetted
  - US export restrictions

Current strategy for monthly traces

- 1 hour, simultaneously at all locations
- Anonymize all with same key (per year)
- Try to keep at same day/same time-of-day each month
  - + compare month-to-month
  - - can't compare different time-of-day
  - - hardware doesn't always comply ...

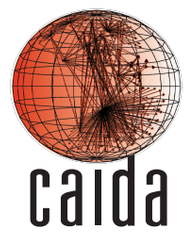

What is a good measurement strategy?

- What does heavy variation in traffic volume imply about appropriate measurement strategies?
  - appropriate length (hour)
  - frequency (weekly, monthly?)
- What meta-data is needed?
  - packet loss at measurement
  - high precision timestamps
- thoughts ... ?

Outline

- Goal
- Problems (EOT)
- Deployments
- Measurements
- Link/Traffic characterization

equinix-chicago asymmetry in flows and routing

Seattle

↓

Chicago

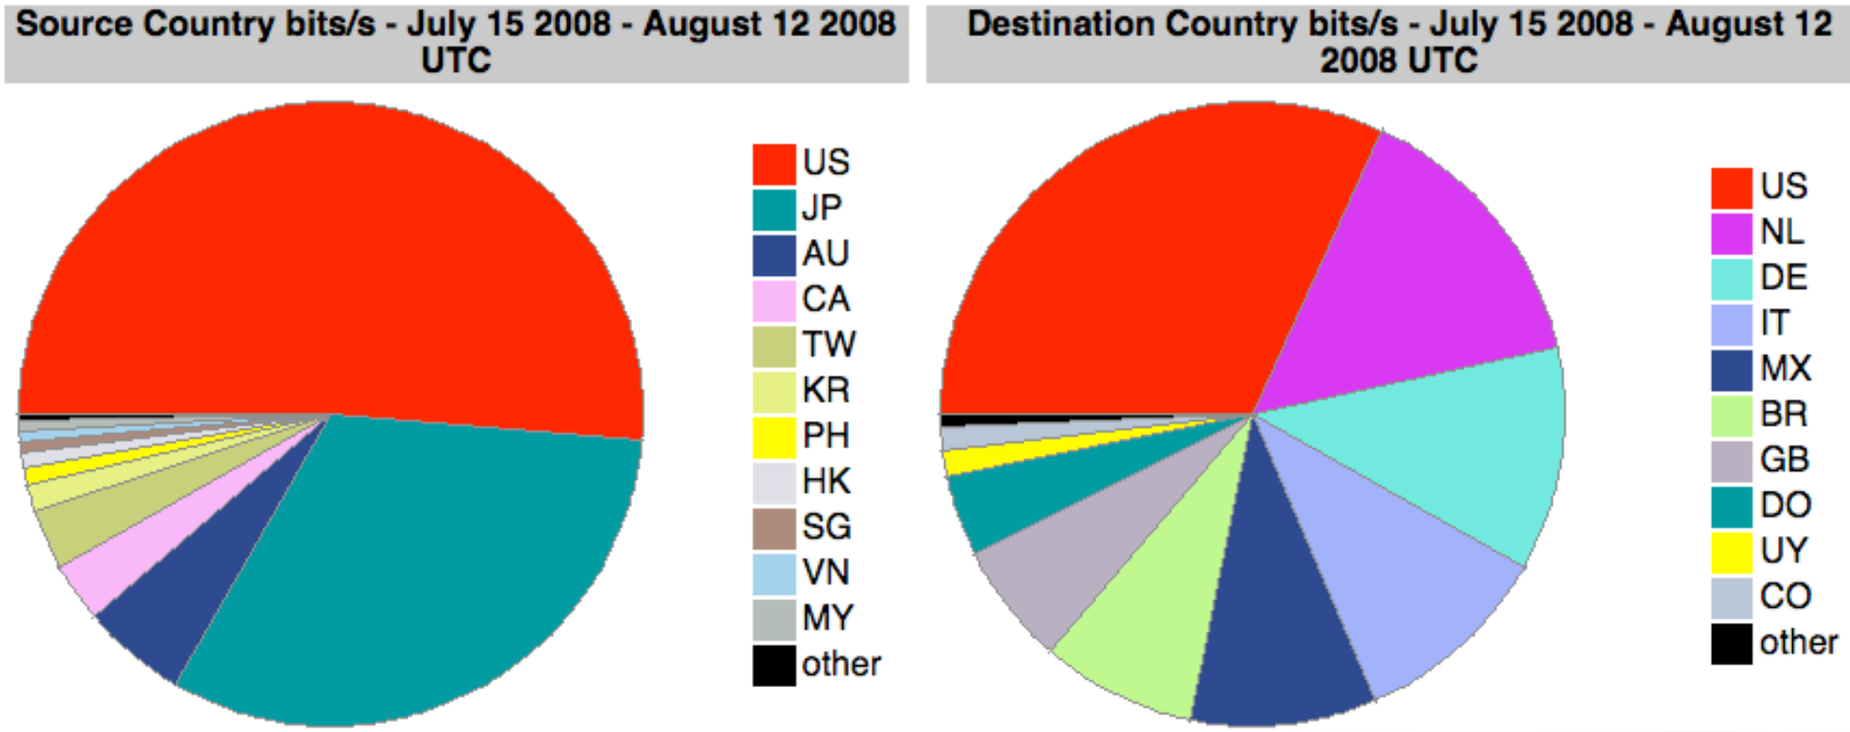

equinix-chicago asymmetry in flows and routing

Seattle

↓

Chicago

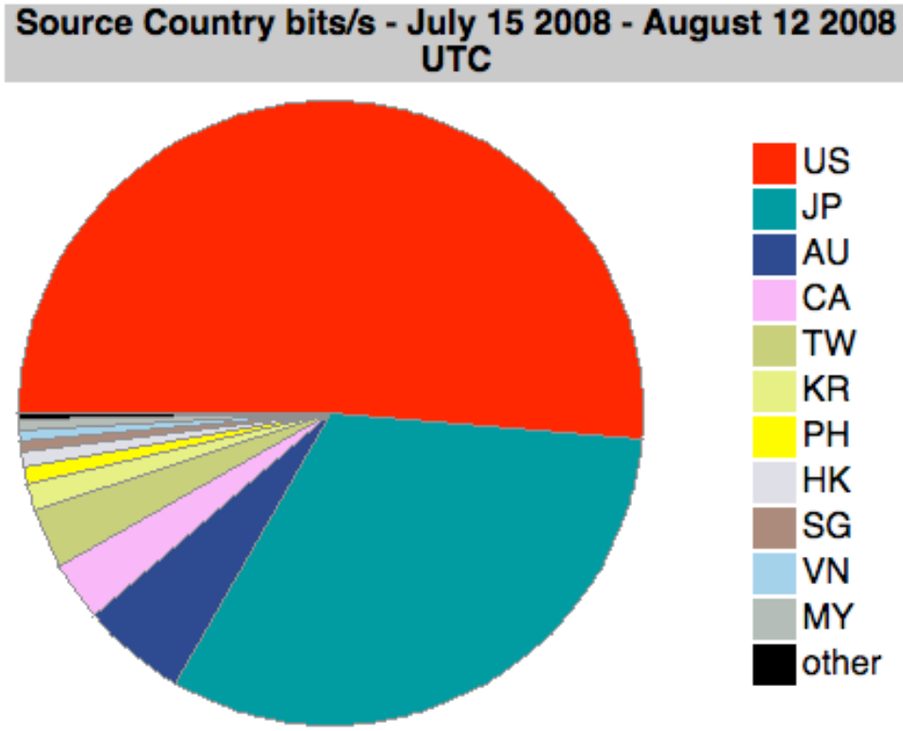

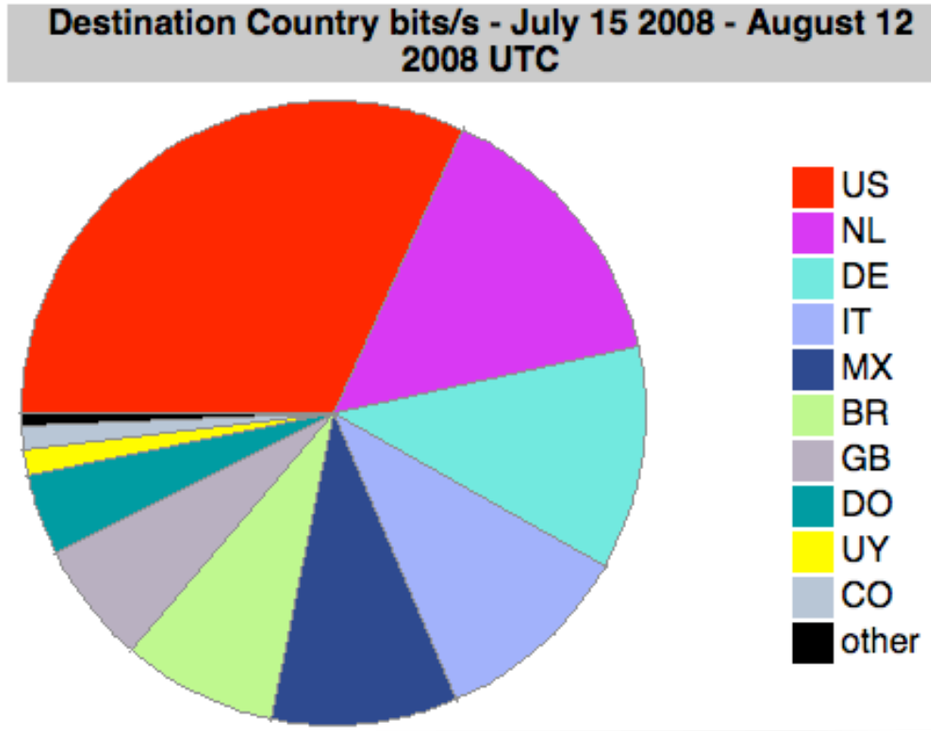

Seattle

↑

Chicago

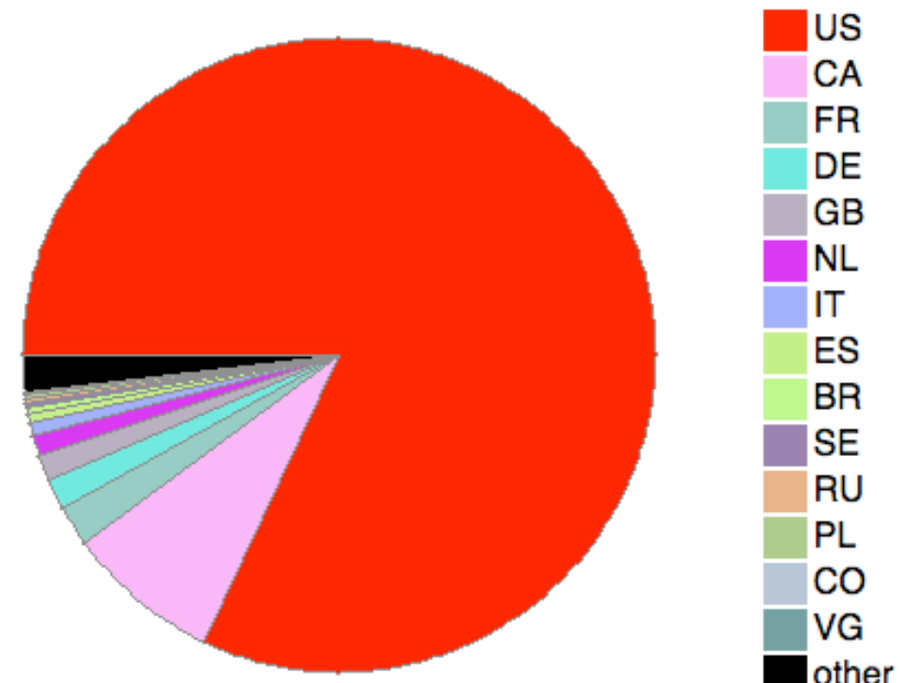

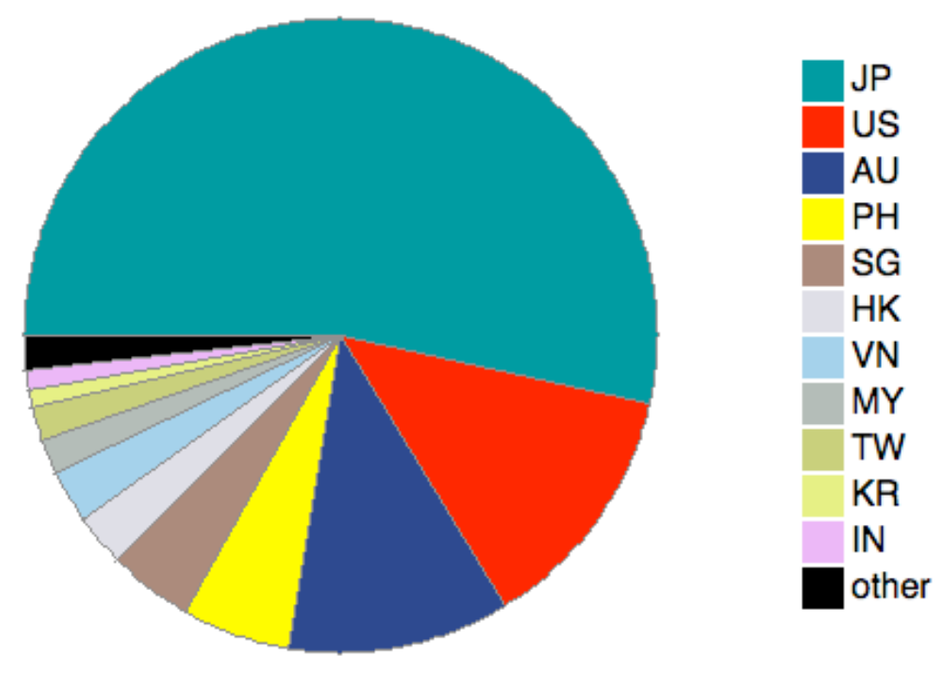

18

Comparison against other links

dataset	link	date	duration	IPv4 pkts	IPv6 pkts	IPv6 fraction
FIX-west	OC3	1998-03-12	7m 15s	5.7 M	-	-
equinix-chicago (westbound only)	OC192	2008-03-19	1h 2m	1.75 G	76 k	0.004%
WIDE MAWI samplepoint F	150 Mbps	2008-03-17 - 2008-03-21	3d 15m	3.82 G	11 M	0.3%

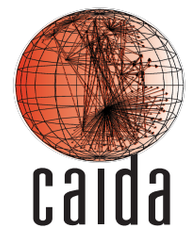

IPv4 packet size distribution

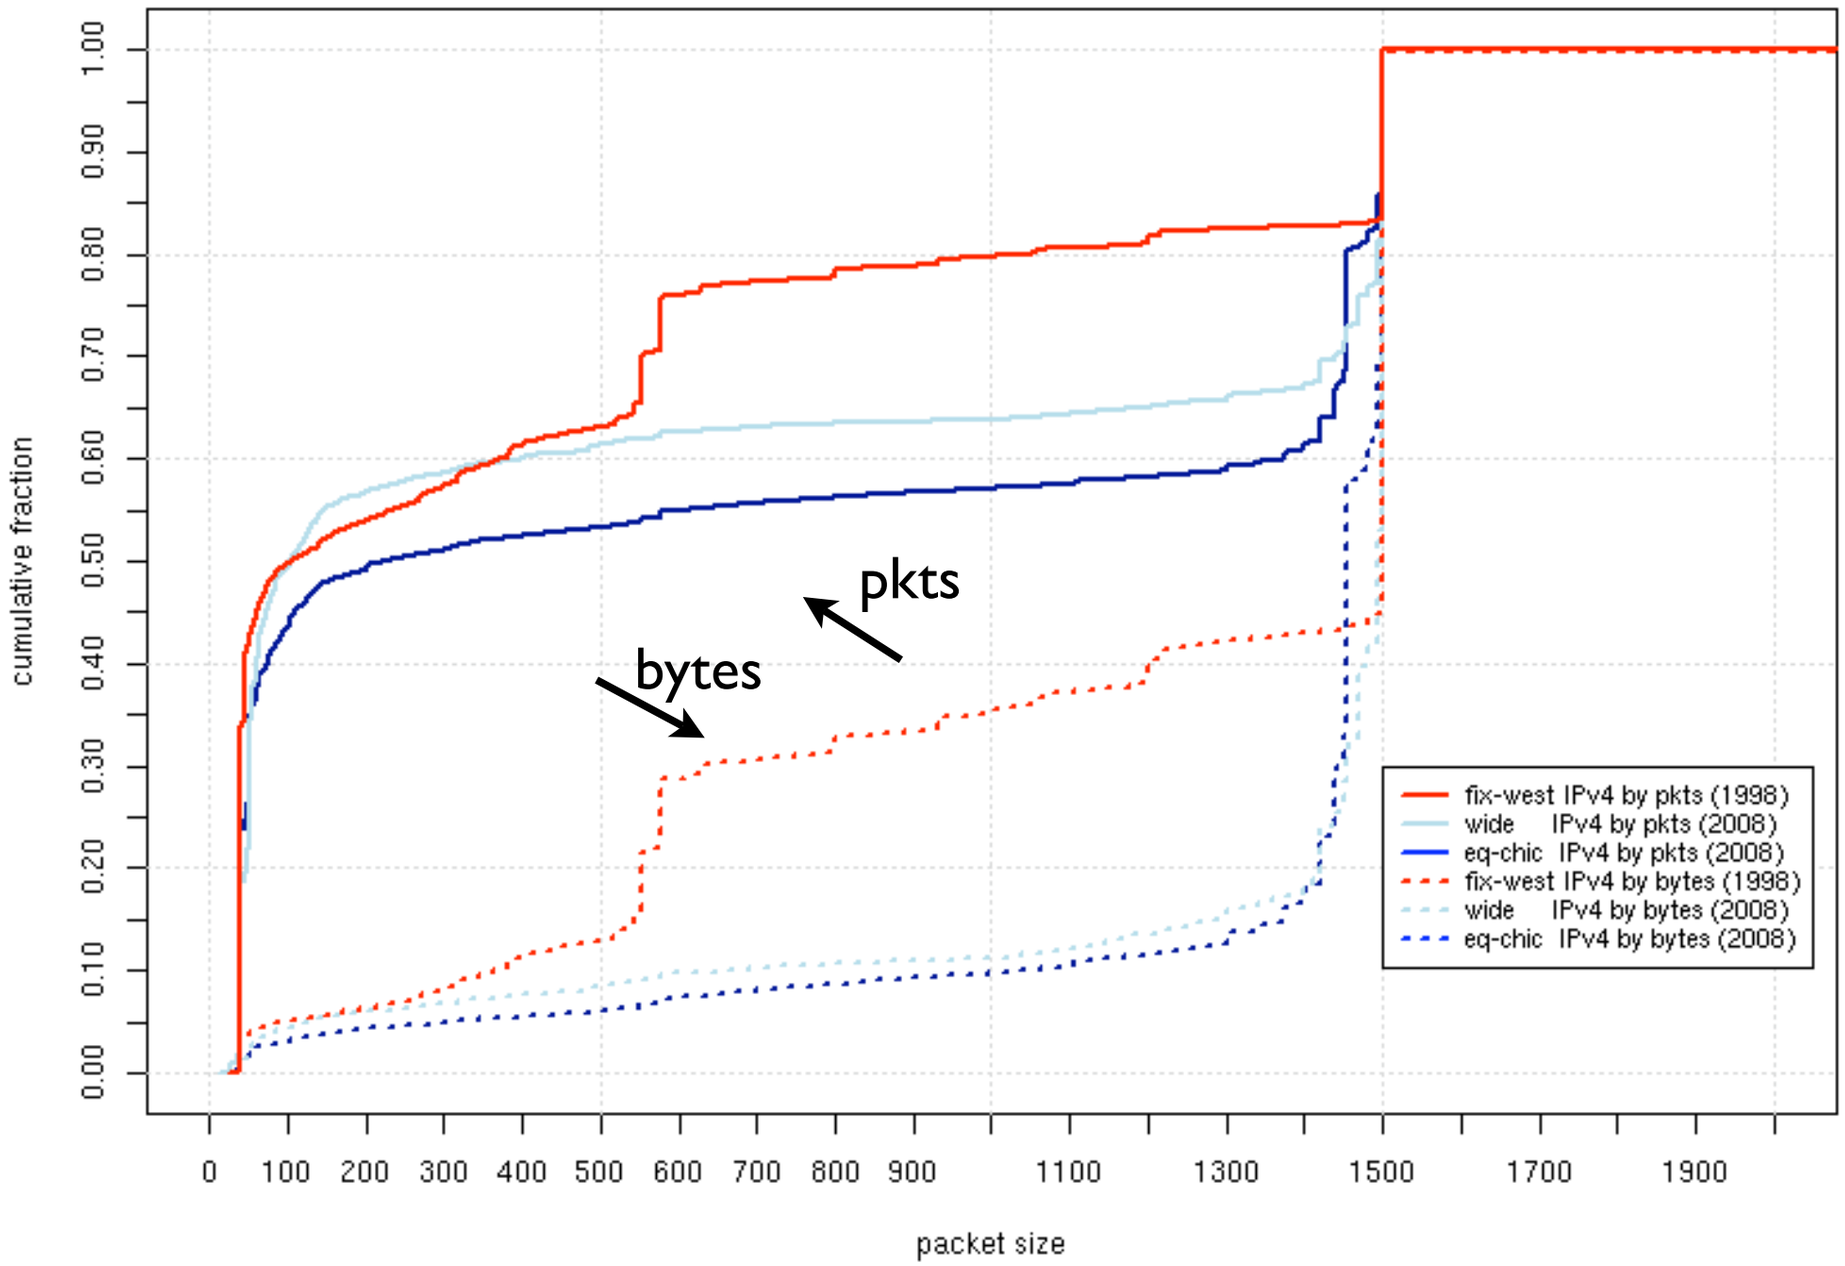

caida

IPv4 packet size distribution

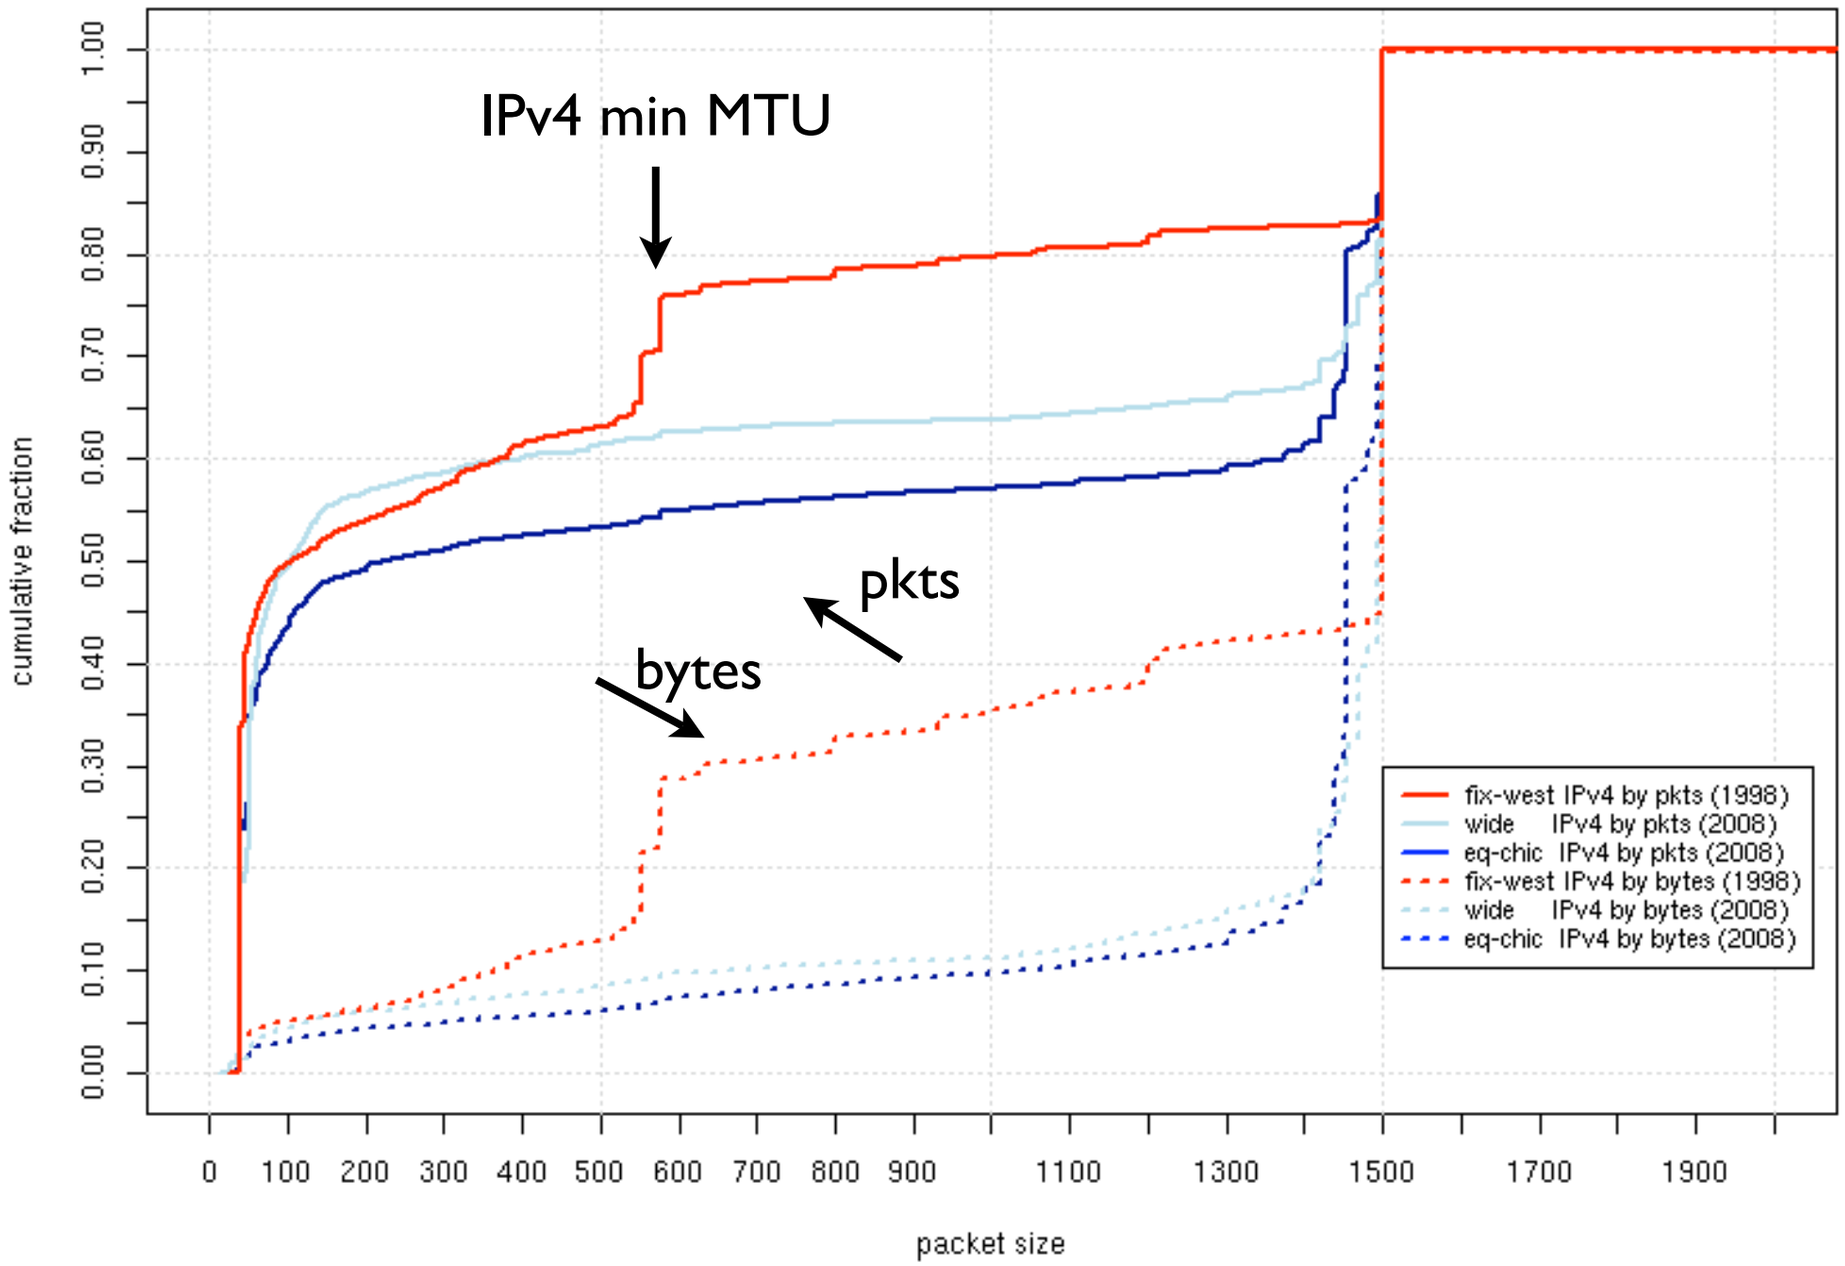

caida

IPv4 vs IPv6 packet size distribution

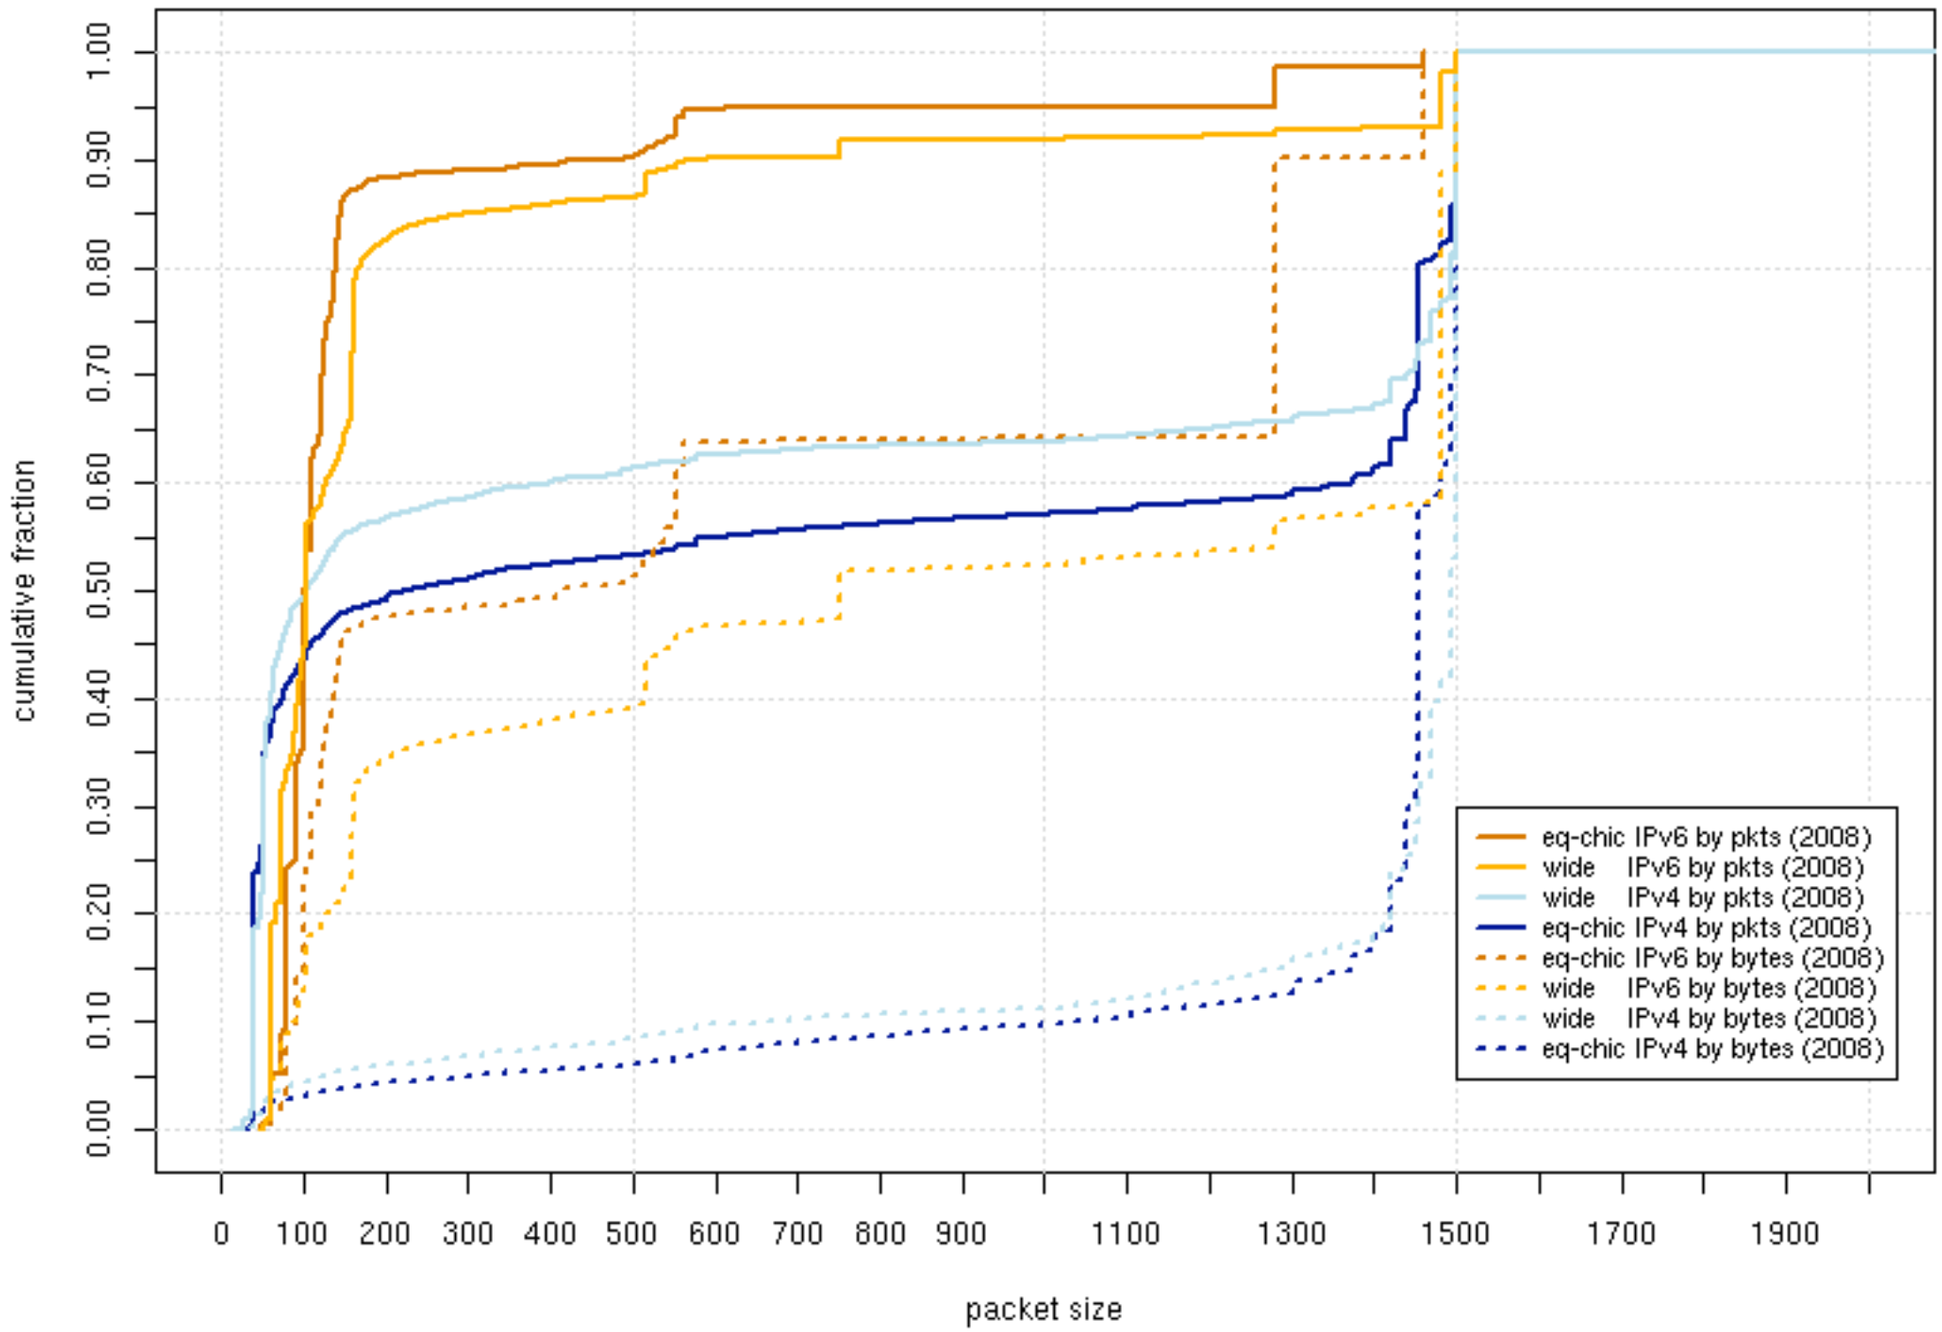

caida

IPv4 vs IPv6 packet size distribution

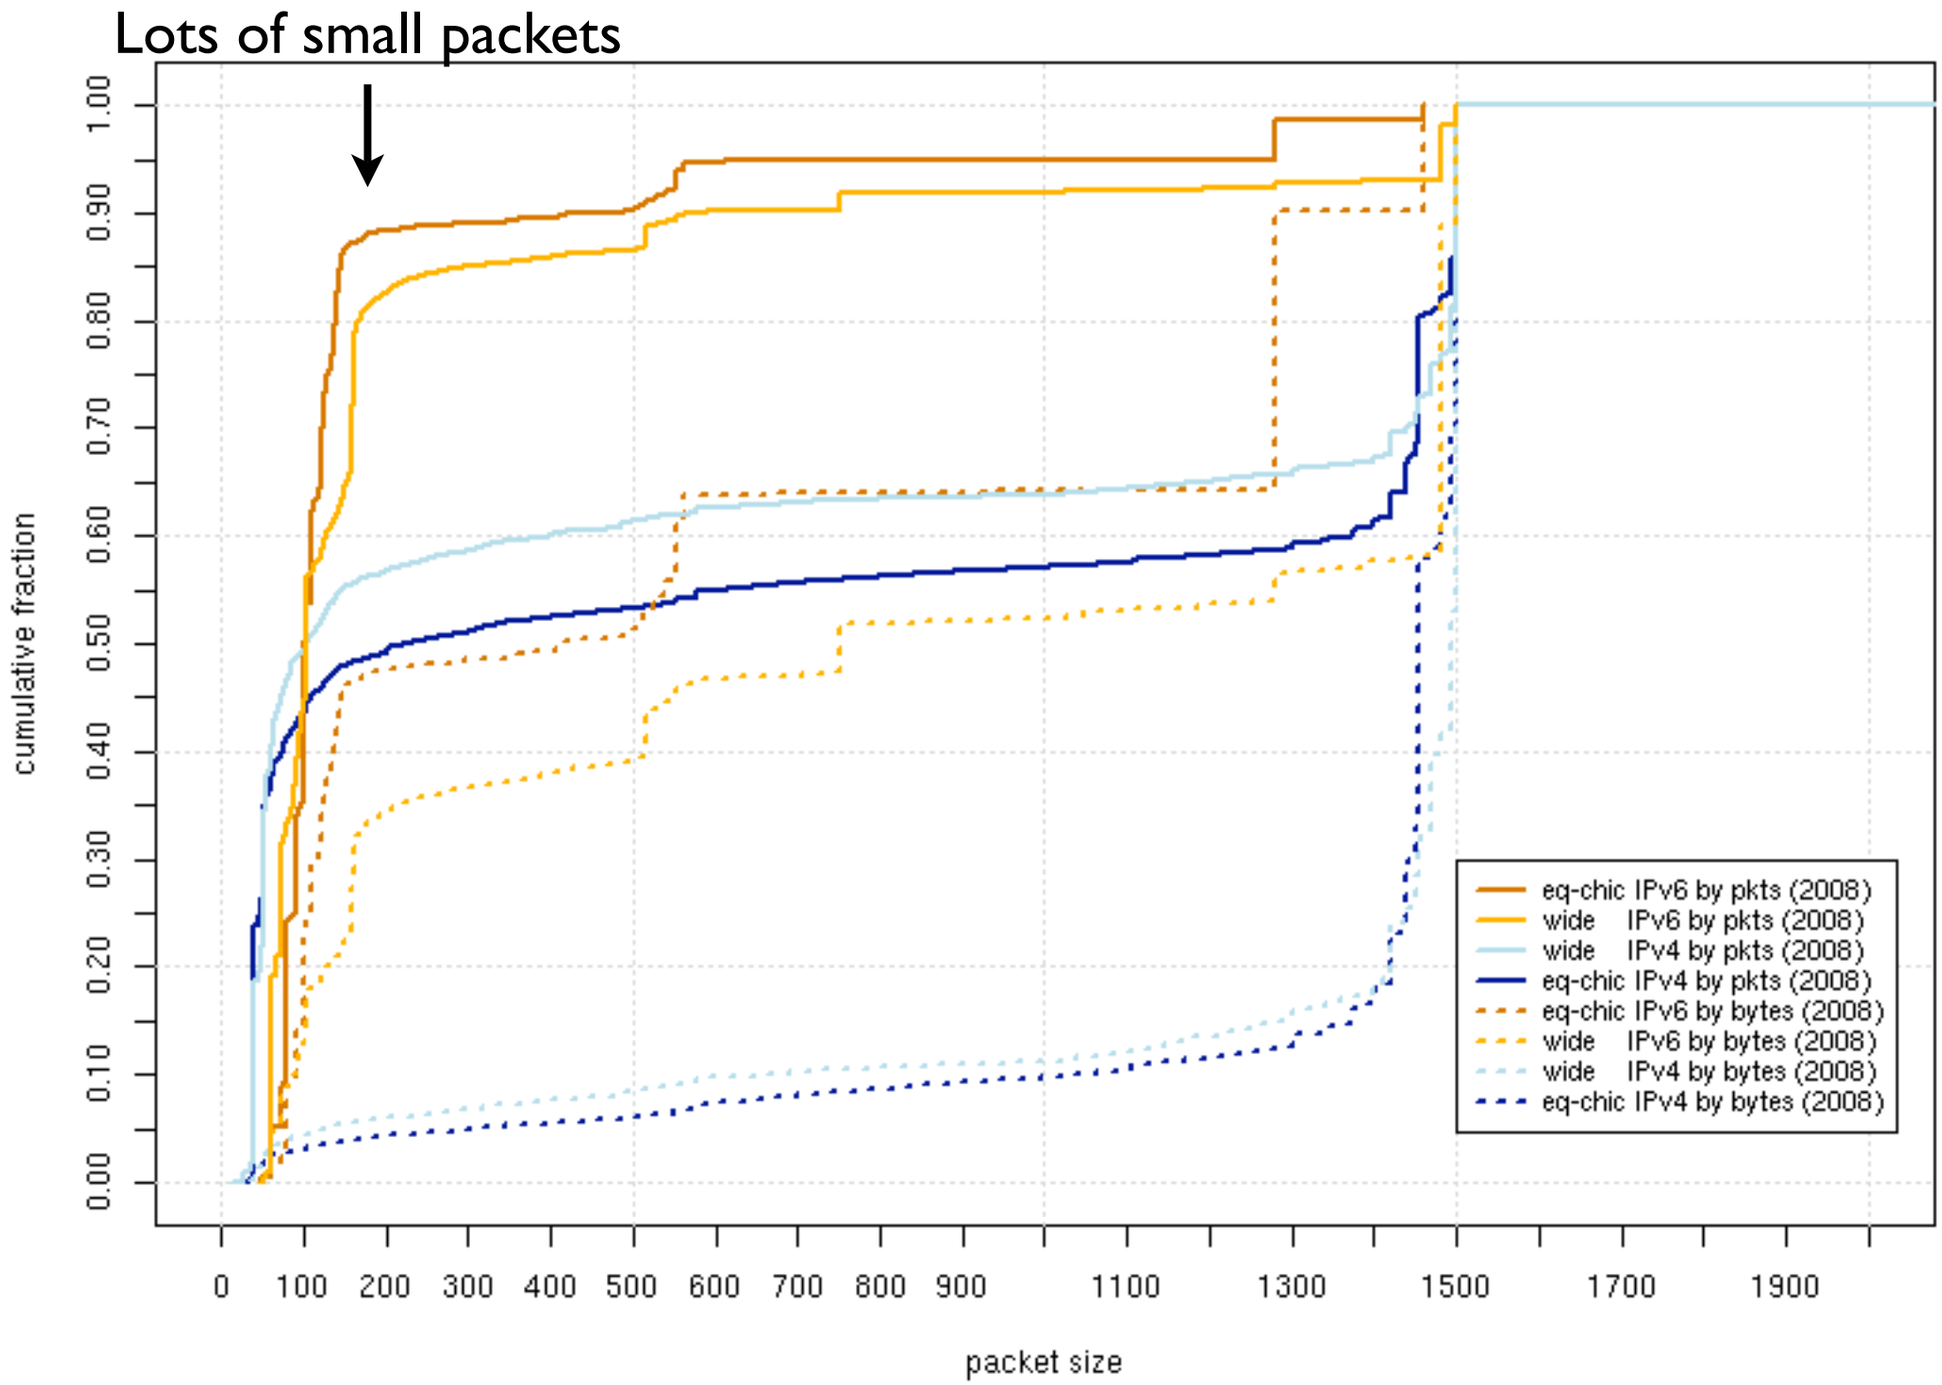

caida

Links

CAIDA passive data overview	http://www.caida.org/data/passive
passive data access	http://www.caida.org/data/passive/anon_internet_traces_request.xml
equinix-chicago monitor	http://www.caida.org/data/passive/monitors/equinix-chicago.xml
packet size distribution	http://www.caida.org/research/traffic-analysis/pkt_size_distribution/graphs.xml
PREDICT	http://www.predict.org