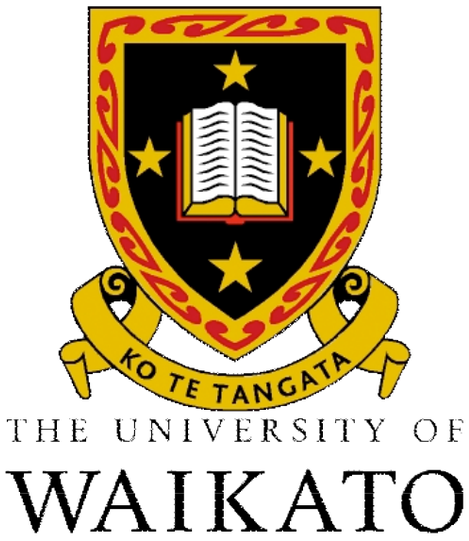



# Traceroute Probe Method and Forward IP Path Inference

Matthew Luckie
mjl@wand.net.nz

Department of
Computer Science
University of Waikato
Hamilton, New Zealand

Young Hyun
youngh@caida.org

Bradley Huffaker
bradley@caida.org

CAIDA
University of California
at San Diego, La Jolla, CA

IMC Greece -- October 2008

# Goals

- Determine which traceroute technique is the most effective
  - most reachable destinations
  - most complete paths
  - most IP links discovered
  - most AS links discovered
  - fewest gap limits (5 consecutive unresponsive hops)
  - fewest loops
  - fewest obviously spoofed responses
- … depending on the destination type
  - 261,530 routable IP addresses selected at random
  - top 500 webservers as ranked by alexa (422 IPs)
  - 2000 routers selected at random
- will focus mostly on random routable IP addresses



2

# Methodology

- conduct six traceroutes for each destination in random order
- 5 second cool-down between methods
- conduct traceroutes at 100pps from *.ark.caida.org
  - 8 vantage points
  - 2 attempts per hop
  - 5 hop gaplimit
  - halt on first loop
  - prove past time exceed message from dst
  - prove past zero-TTL forwarding



3

# Traceroute methods surveyed

- UDP
  - probe id: dport (unused); ephemeral sport;
- **UDP-Paris**
  - probe id: UDP checksum field; ephemeral sport; unused dport;
- ICMP
  - probe id: ICMP sequence field;
- **ICMP-Paris**
  - probe id: ICMP sequence field;
- **TCP (port 80)**
  - probe id: IP ID; dport 80, ephemeral sport
- **UDP-Paris DNS**
  - probe id: UDP checksum field; 5-tuple constant; sport 53; unused dport; valid DNS payload



4

# vantage points

| Host | Location |
|---|---|
| cbg-uk | University of Cambridge Cambridge, England |
| nrt-jp | Asia-Pacific Advanced Network (APAN) Tokyo, Japan |
| syd-au | AARNet Sydney, Australia |
| bcn-es | Universitat Polit`ecnica de Catalunya Barcelona, Spain |
| hel-fi | Helsinki University of Technology (TKK) Espoo, Finland |
| cjj-kr | KREONet2 Daejeon, Korea |
| iad-us | ARIN Bethesda, Maryland |
| san-us | CAIDA San Diego, California |



5

# destination lists

| list | size |
|---|---|
| Random IP address selected from Routeviews' prefixes<br>http://www.routeviews.org | 261,530 |
| Alexa top 500 websites<br>http://www.alexa.com/site/ds/top_sites?ts_mode=global&lang=none | 500 |
| Routers selected at random from with in previous traces | 2,000 |



6

# halt reasons*

| | |
|---|---|
| **reached** | the destination was successfully reached |
| **ICMP-Unreach** | an ICMP unreachable packet was received |
| **Loop** | an IP address was repeated in the collected path<br>this does not include zero-TTL forwarding |
| **gap limit** | the maximum number of non-responsive hops (5) |

*Why did scamper stop probing.



7

# Random routable IP addresses

- 257,504 prefixes observed at routeviews for week of 19-25 March 2008 (median snapshot per day)

- 255,981 prefixes observed in at least 3 snapshots
  - one random address per prefix if prefix is more specific than /16
  - one per /16 otherwise
  - never select more than 1 address per /24, addresses in team cymru bogon list, do-not-probe (1.14 /8s)

- 261,530 addresses selected

- use unique list per vantage point



8

# 261,530 routable IP addresses: cbg-uk

Gap Limit was reached roughly 75% of traces.

All monitors saw roughly the same ranges for halt reasons, these graphs are just for cbg-uk.

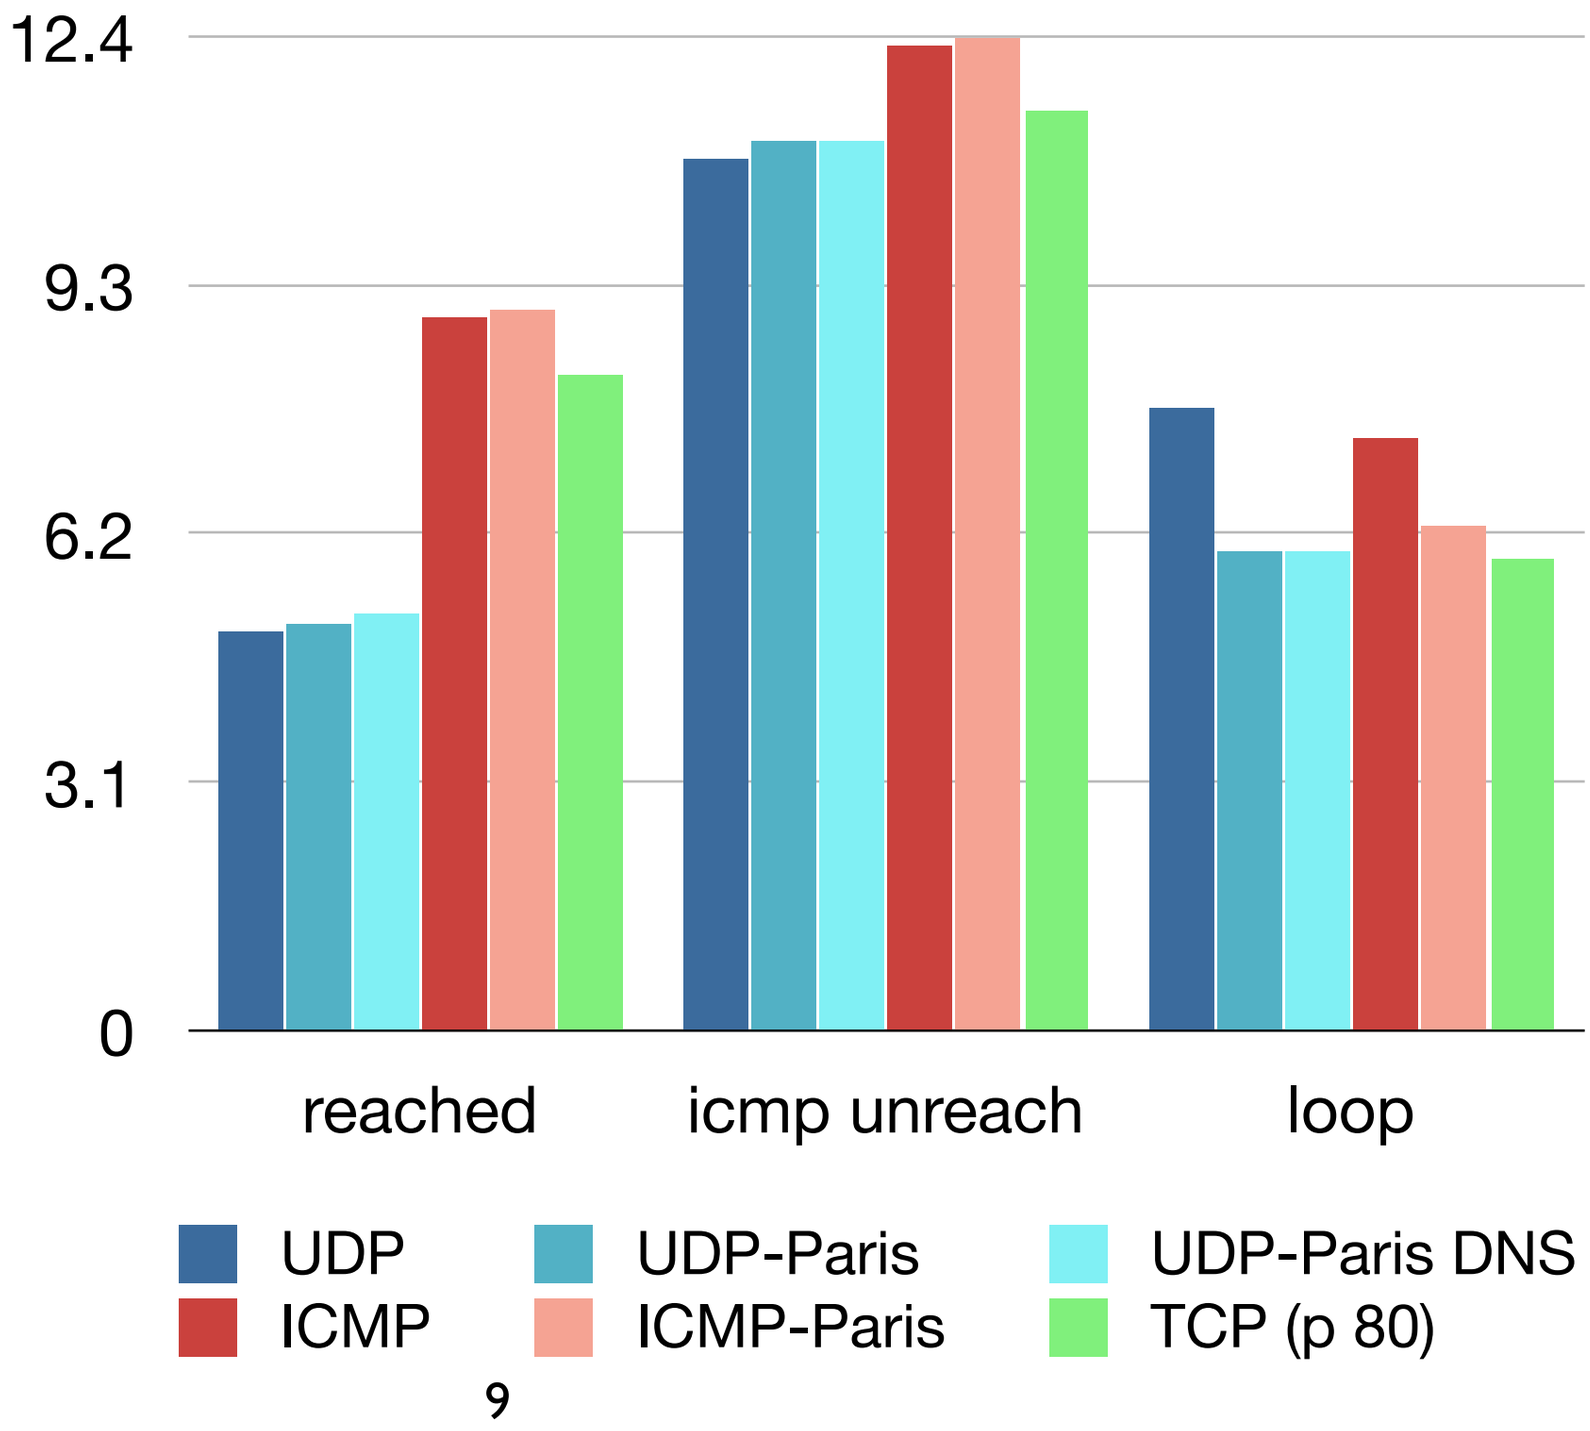

# Comments

- ICMP-Paris reaches most destinations
  - –also obtains most ICMP unreachables, which is better than having your probe silently discarded

- UDP reaches the least
  - –But it and the ICMP technique are known to produce invalid IP paths more frequently than their Paris counterpart

- UDP-Paris DNS performs about the same as vanilla UDP-Paris

# Comments

- Reachability results very similar across other ten vantage points
  - –despite different IP lists

- Some variation in ICMP-Unreach, Loops, Gaplimit
  - –vantage point a factor

11

# reachable dest.: cbg-uk

- Total reachable: 31,439 (12.0%)

- ICMP-paris by itself yields the most:
  - 2,3638 (9.0%)

- ICMP-paris and TCP together get:
  - 30,726 (11.7%)

- Not using UDP misses 2.3% of destinations reachable with the three methods

# through the firewall

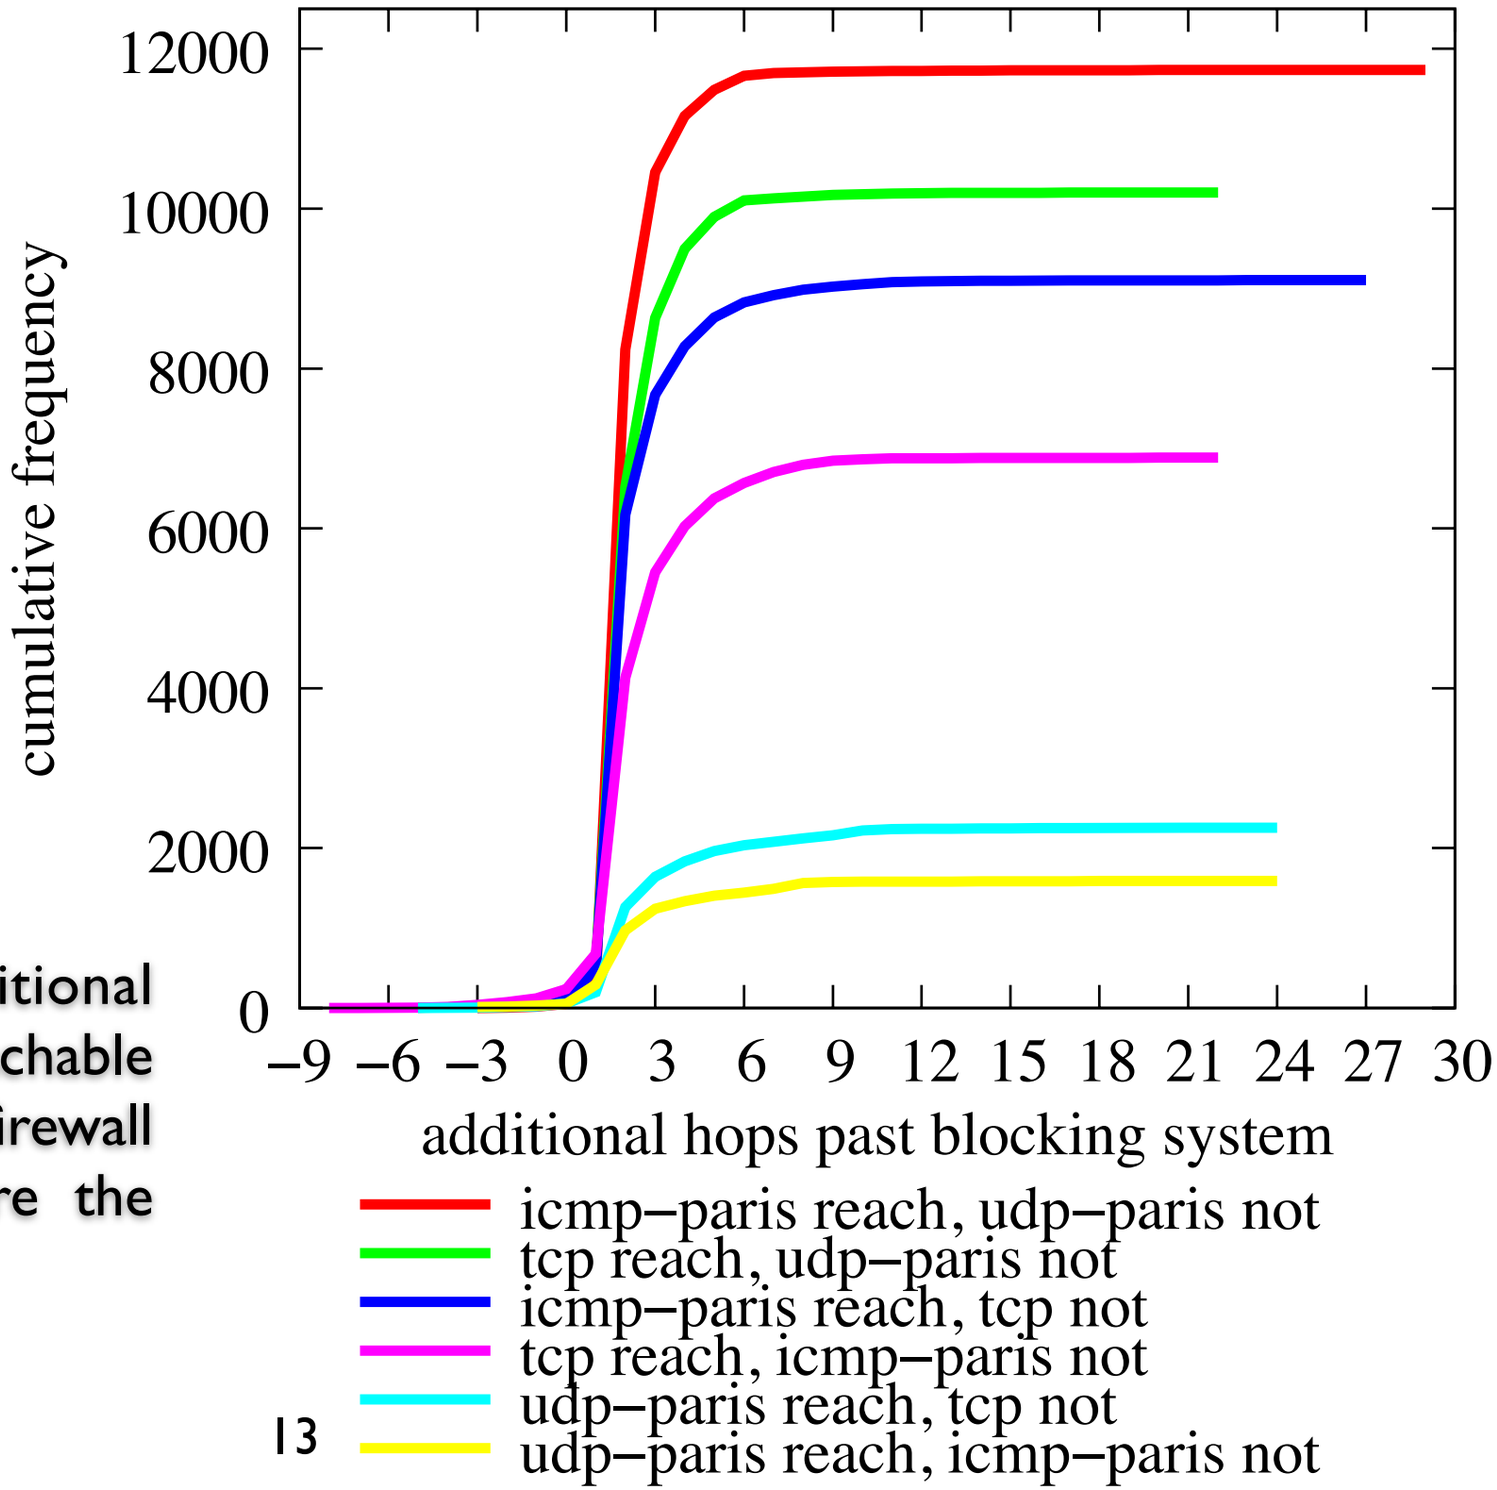


Cumulative frequency of additional hops past blocking system to reachable destination from cbg-uk. Most firewall were observed two hops before the destination.

# Complete Unique Paths

Set of destination where all three had responses from every TTL including the destination, complete, and counted the number of unique paths to those destinations.

UDP-Paris saw the most different set of paths to the same set of destinations.

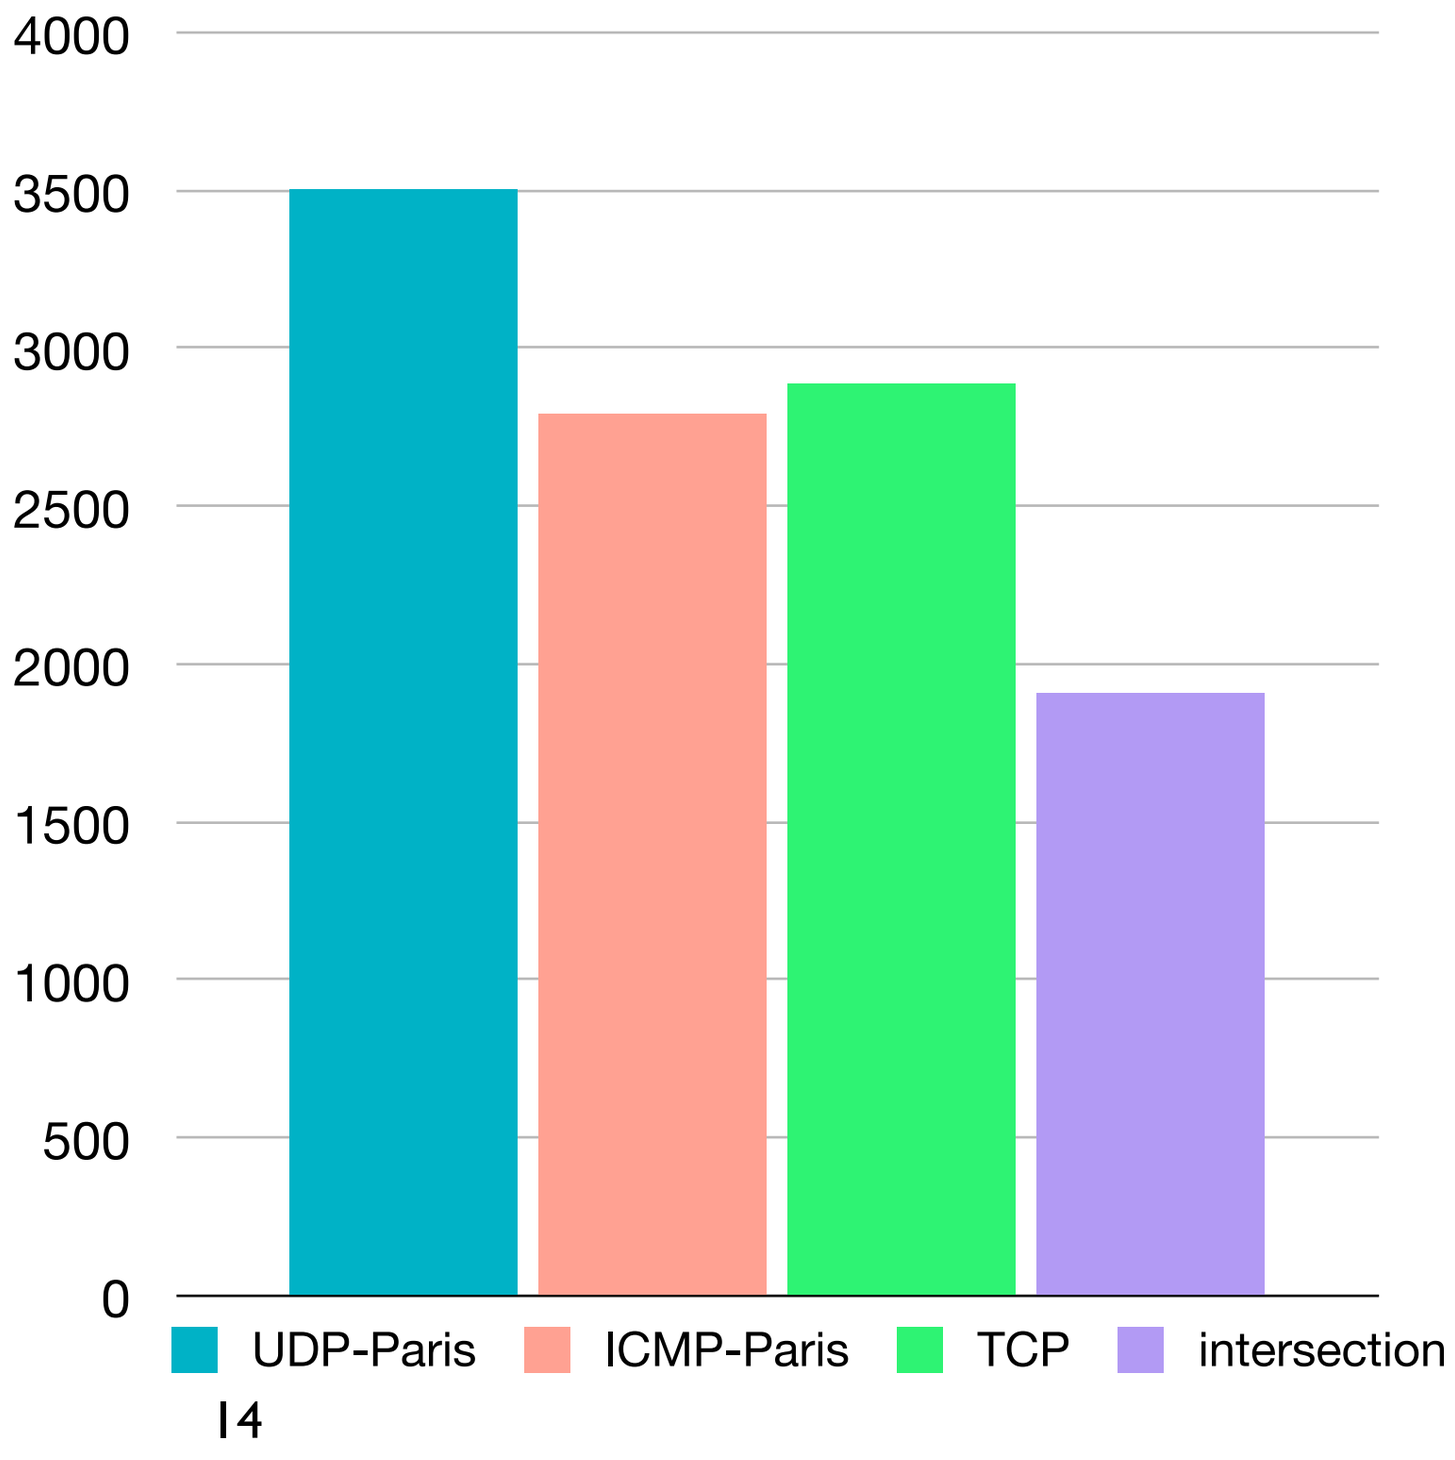

# Complete Unique Paths: cbg-uk

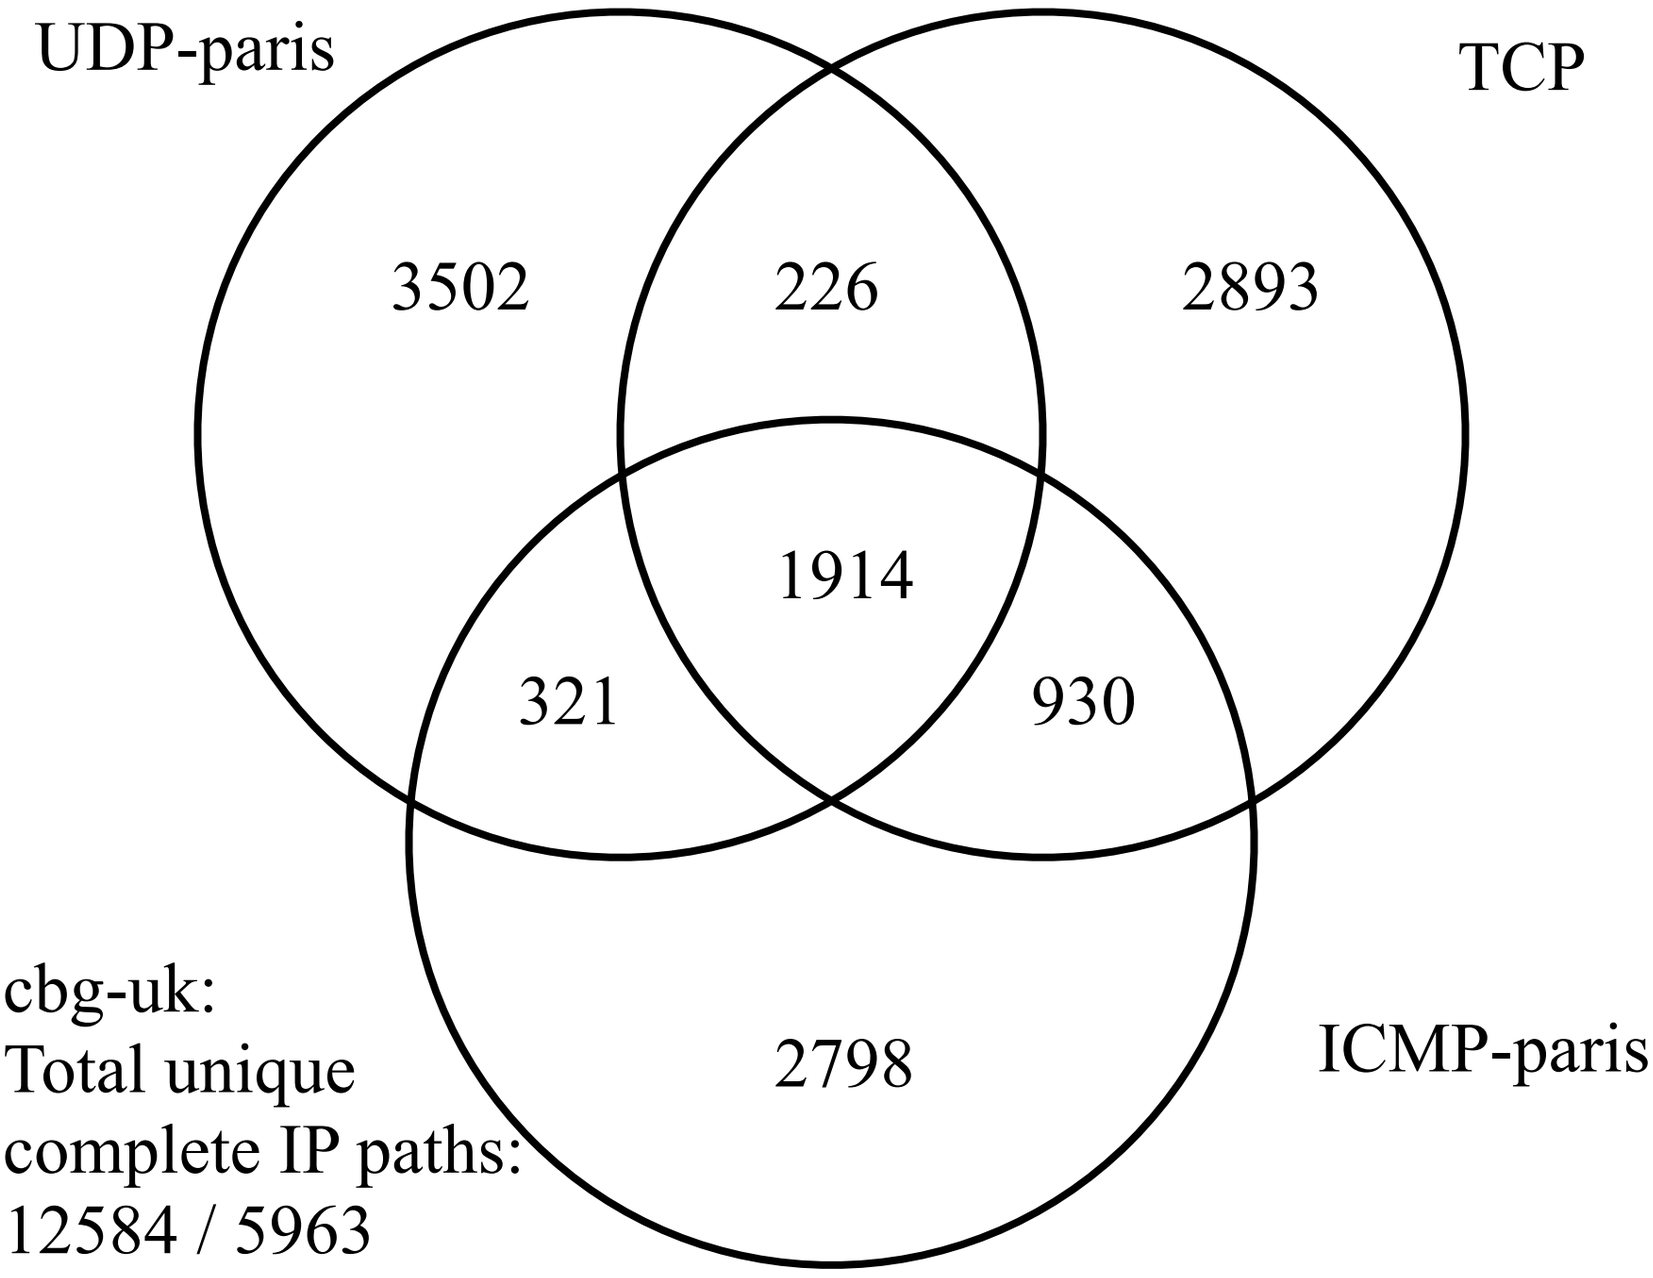




15

# IP links: cbg-uk

UDP-Paris infers the most IP links despite reaching the fewest destinations.

96.6% of links are seen between UPD-Paris and ICMP-Paris.

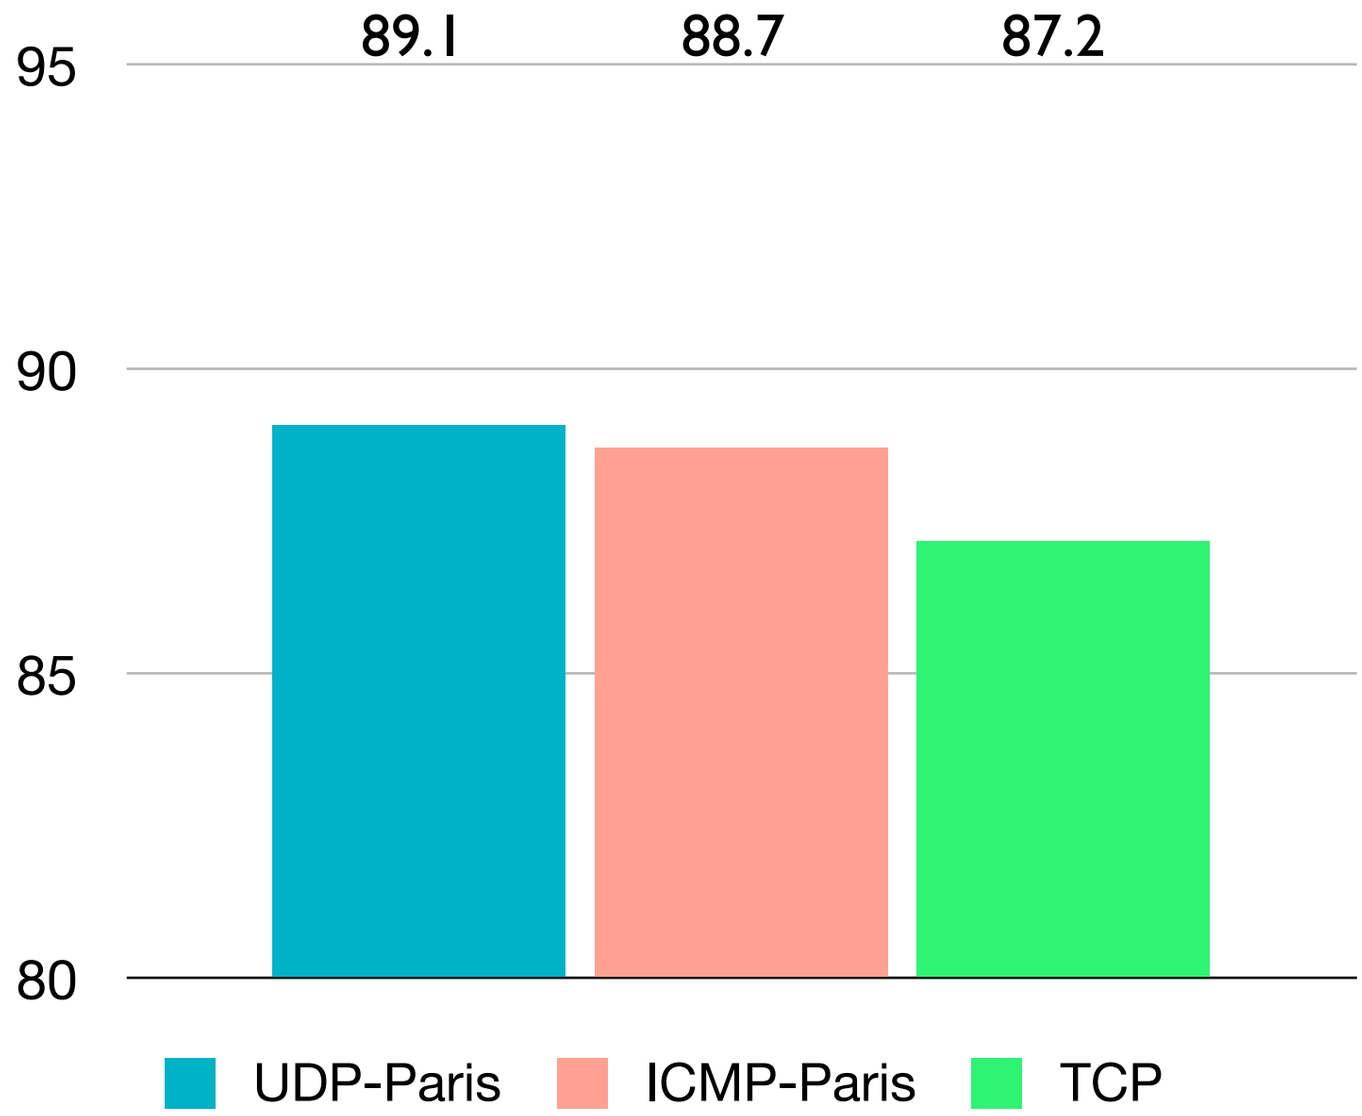


16

# AS links: cbg-uk

UDP-Paris inferred the fewest AS links despite inferring the most IP links, suggesting it found more IP links inside an AS then between them.

99.4% of links are seen between ICMP-Paris and TCP.

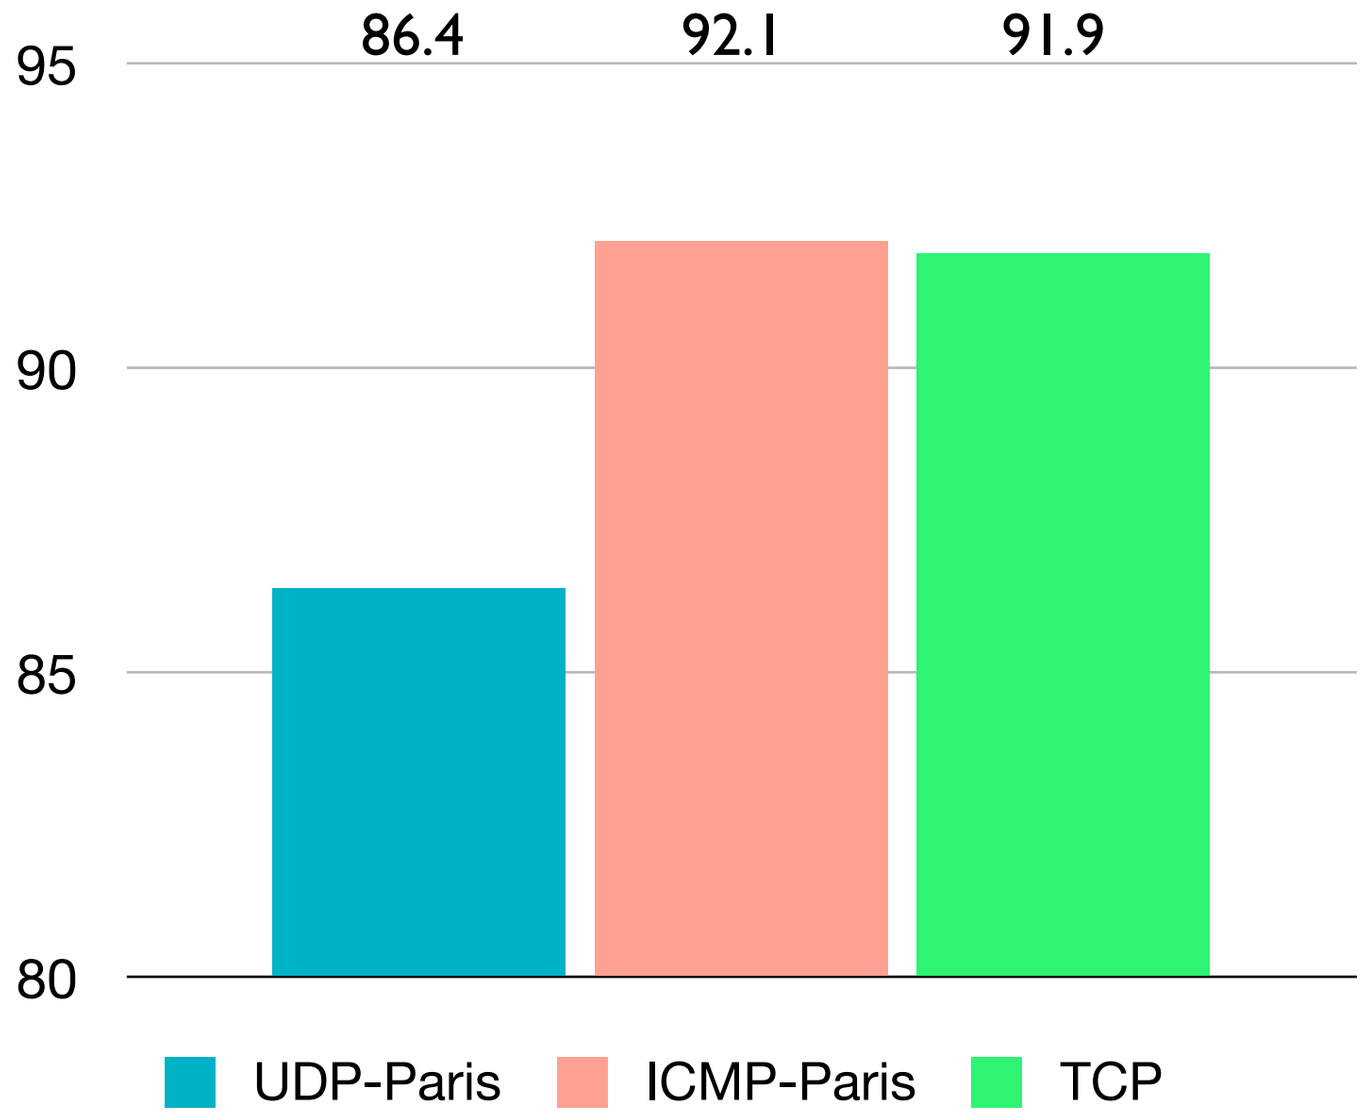


17

# Summary so far

- ICMP-paris reaches most destinations, infers most AS links
  - TCP not far behind
- UDP-paris infers most IP links
  - TCP least
- TCP and ICMP IP paths appear to be the most similar
  - vantage point has an effect, but trend is there
- Firewalls are most commonly two TTLs from the target.



18

# Inferring Spoofed Destinations #1

- ICMP destination unreachable: port unreachable

  – RFC 792: Indicated port is not running an active process

  – Source address may vary, but supposed to be from destination

  – Used in alias resolution



19

# Inferring Spoofed Destinations #1

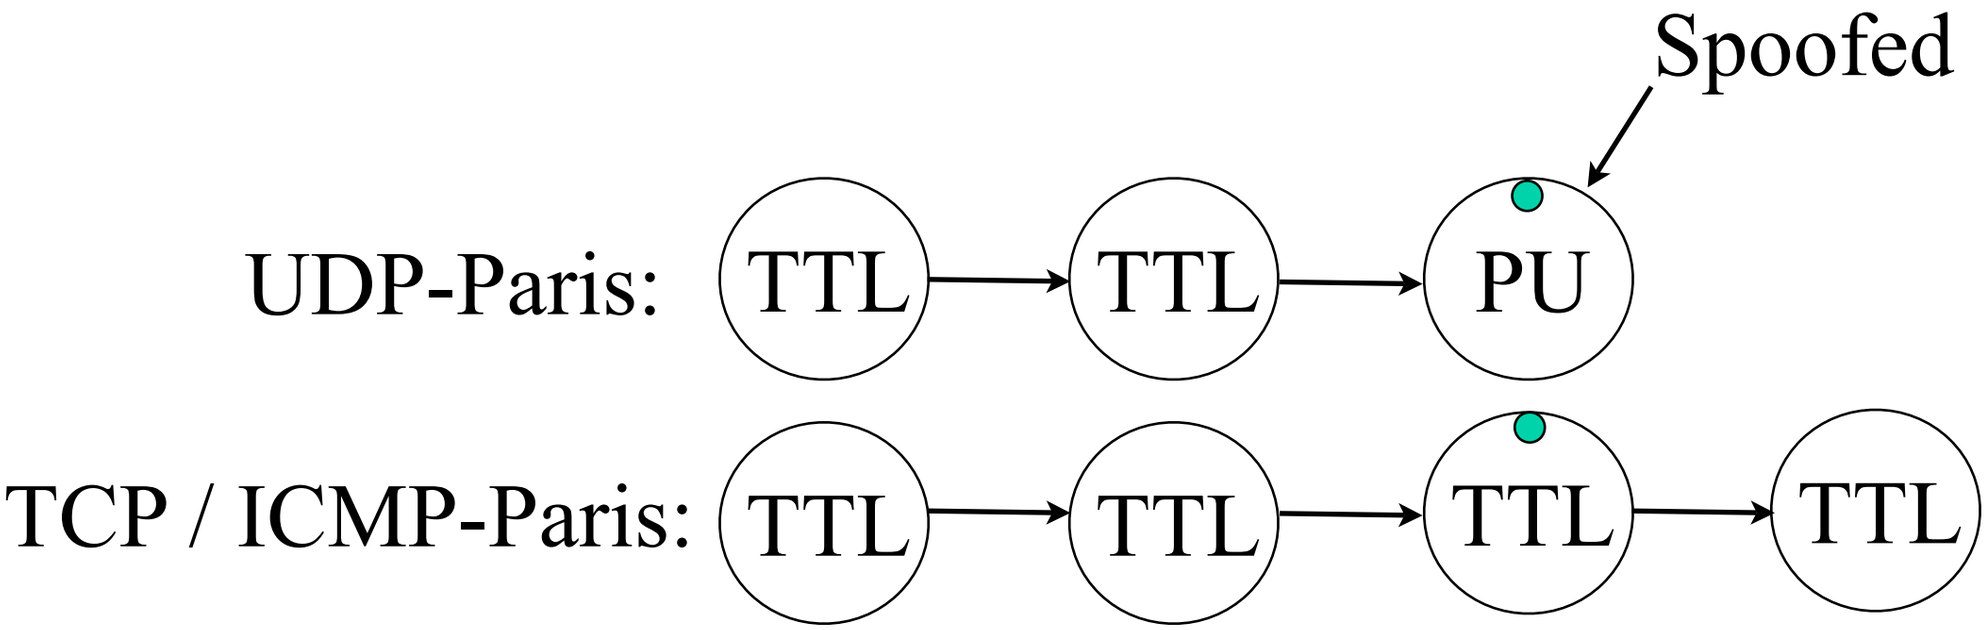


Of 12875 port unreachable for UDP-Paris, 27 were spoofed



20

# Inferring Spoofed Destinations #2

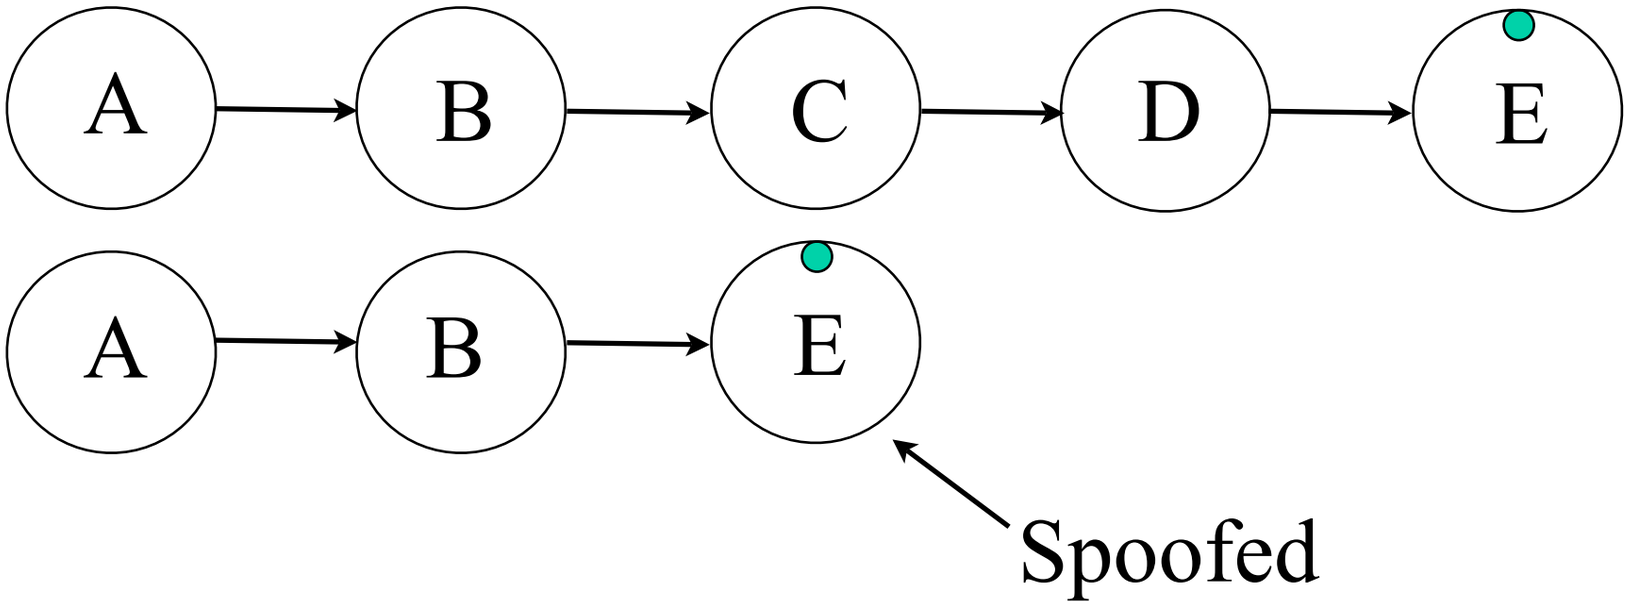


Of 21576 destinations reached with TCP, 221 were spoofed.
UDP-Paris: 21 destinations spoofed, 14 for ICMP-Paris

221 SYN/ACK
61 RST/ACK

# Packet counts

- ICMP-Paris: 6,943,071
- TCP: 7,033,384
- UDP-Paris: 7,122,459

ICMP-Paris sends 2.5% fewer packets than UDP-Paris



22

# Router List

- 2000 IP addresses selected at random
- Previously observed in traceroute path:
  - –to send time exceeded message
  - –at least one additional ICMP time exceeded past the address, from a different IP



23

# 2000 random routers

caida

THE UNIVERSITY OF WAIKATO

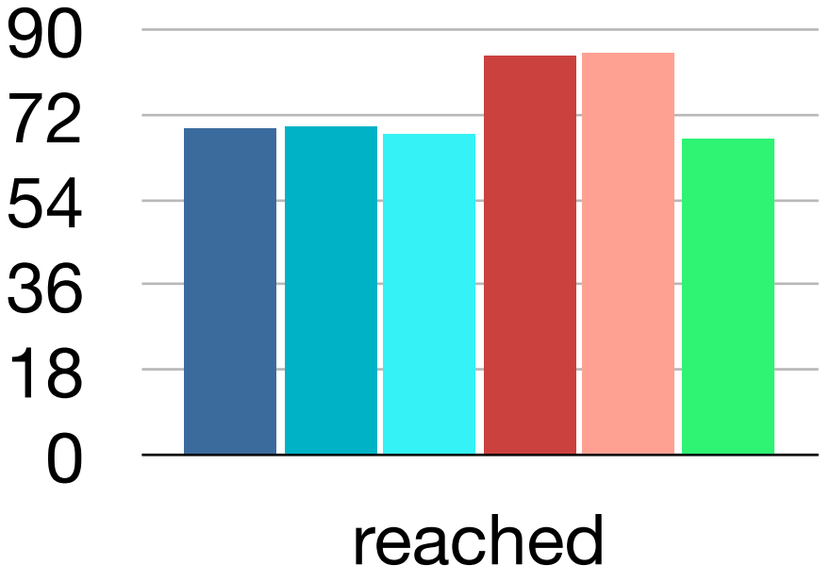


reached

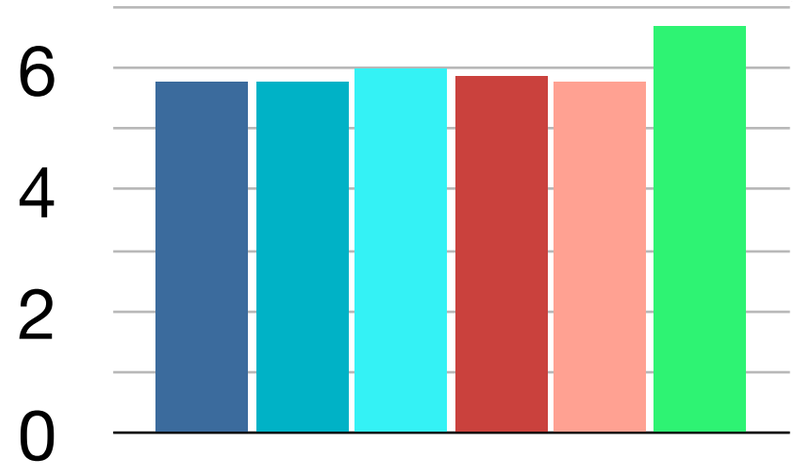


icmp unreach

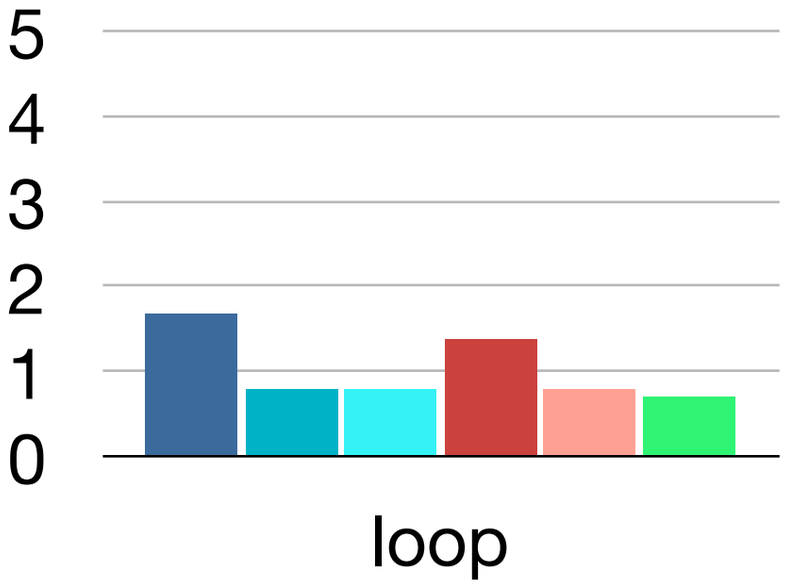


loop

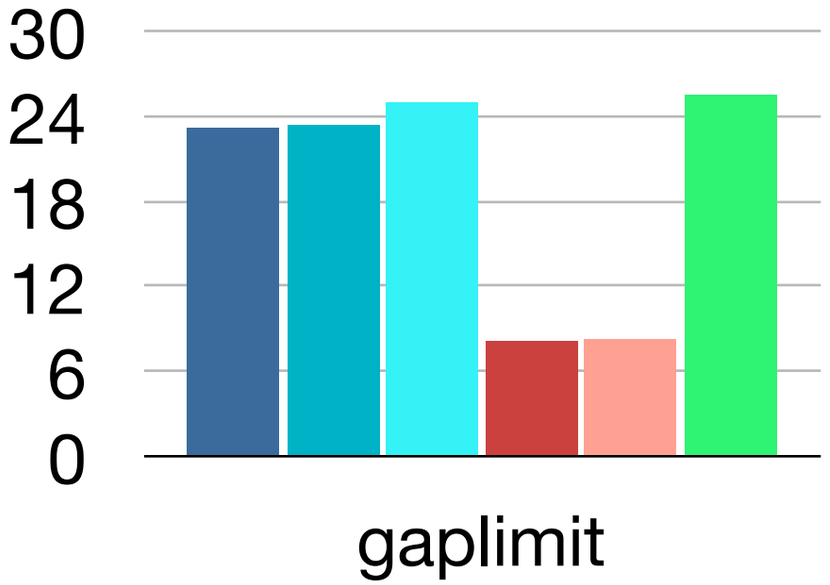


gaplimit

UDP
UDP-Paris
UDP-Paris DNS
ICMP
ICMP-Paris
TCP (p 80)

24

# Webserver list

- Screen scrape of alexa.com top 500
- Resolved from san-us.ark.caida.org
- 422 IP addresses selected
  - –58 Google ccTLD instances => 4
  - –Ebay ccTLD instances
  - –Akamai



25

# 422 webservers

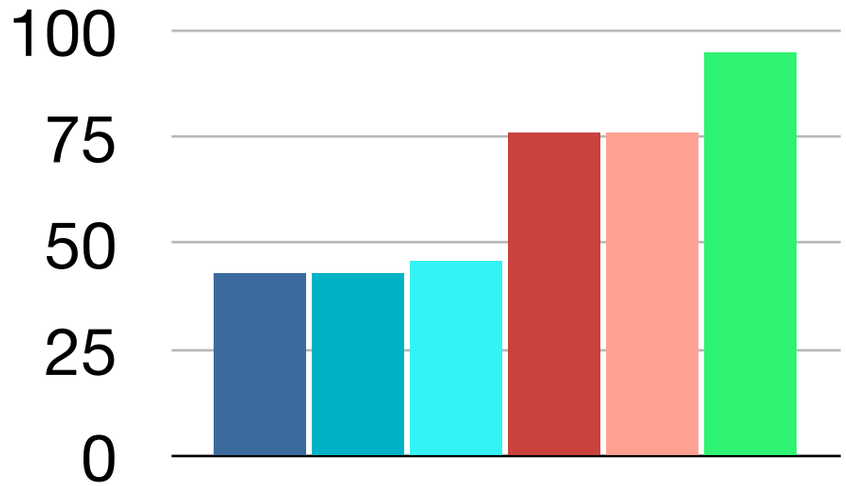


reached

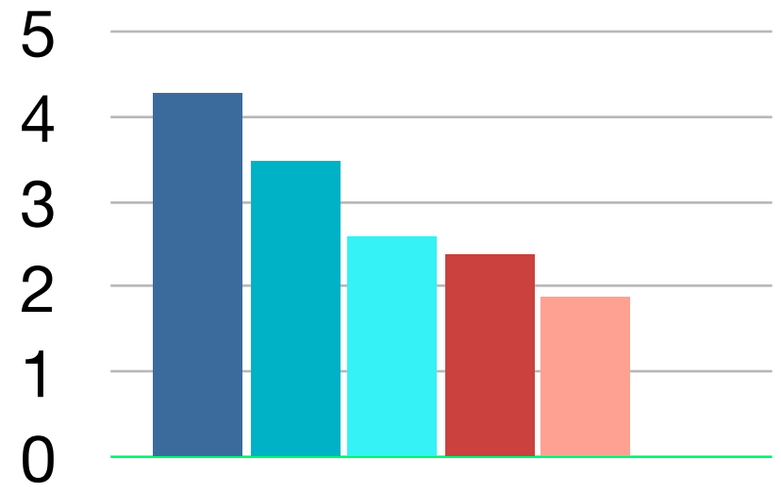


icmp unreach

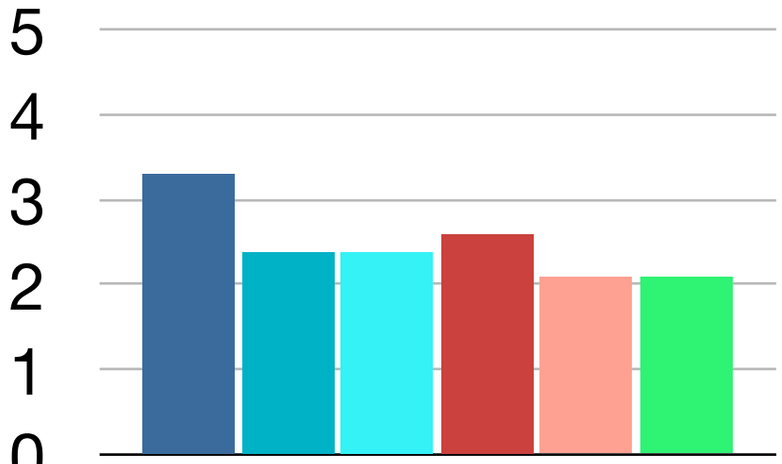


loop

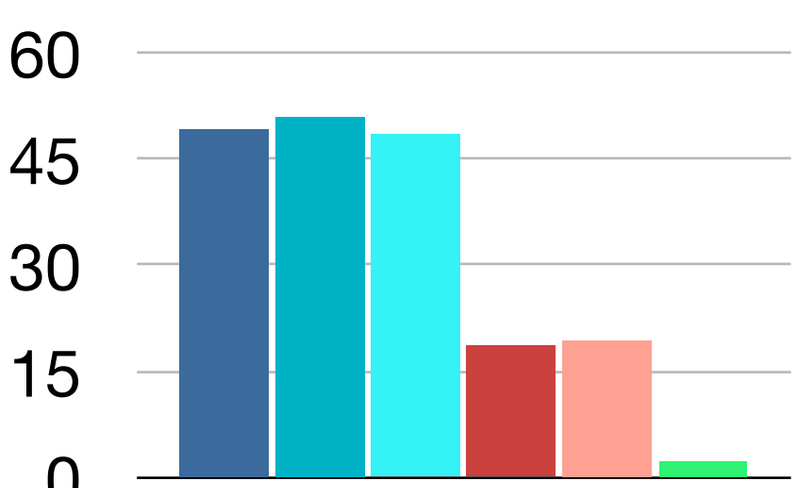


gaplimit

UDP
UDP-Paris
UDP-Paris DNS
ICMP
ICMP-Paris
TCP (p 80)



26

# Conclusion

- ICMP-Paris is superior in destinations reached and AS links found

- UDP-Paris finds more IP links inside an AS then between them.

- Using multiple probe methods improves coverage
  - Also allows integrity of IP paths to be tested

- UDP-Paris DNS bit of a flop



27