Internet Path Transparency Measurements using RIPE Atlas

Brian Trammell and Mirja Kühlewind, ETH Zürich

(with thanks to Emile Aben, RIPE NCC)

Measurements and Tools WG, RIPE 72 Copenhagen, 25 May 2016

mami

measurement and architecture for a middleboxed internet

measurement | architecture | experimentation

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 688421.The opinions expressed and arguments employed reflect only the authors' view. The European Commission is not responsible for any use that may be made of that information.

Supported by the Swiss State Secretariat for Education, Research and Innovation under contract number 15.0268. The opinions expressed and arguments employed herein do not necessarily reflect the official views of the Swiss Government.

path transparency (in one slide) (From MAT at RIPE 71)

- The Internet is not end-to-end...
  - some of this is policy, but a lot of it is accident
  - deployment of new protocols over IP, transport extensions difficult or impossible
- ...but some paths are worse than others.
  - Goal: data on "how bad" and "where" to guide future protocol design
  - In operations: another tool for troubleshooting connectivity dependency for unusual traffic

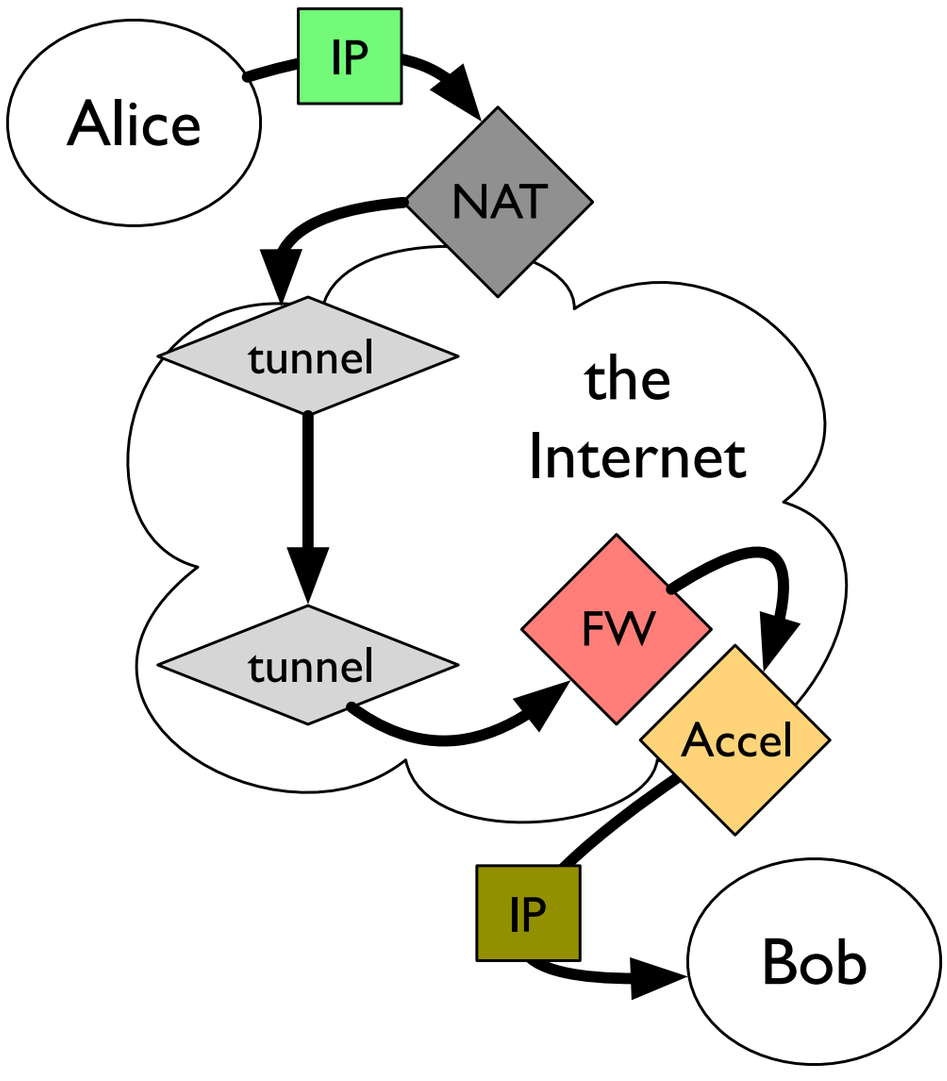

Background: Active Measurement of Path Transparency

- Basic methodology:
  1. throw a bunch of packets at the Internet
  2. see what happens.
- Ideal: two-ended A/B testing
- Scalable: one-ended A/B testing
- Multiple sources: isolate on-path from near-target impairment

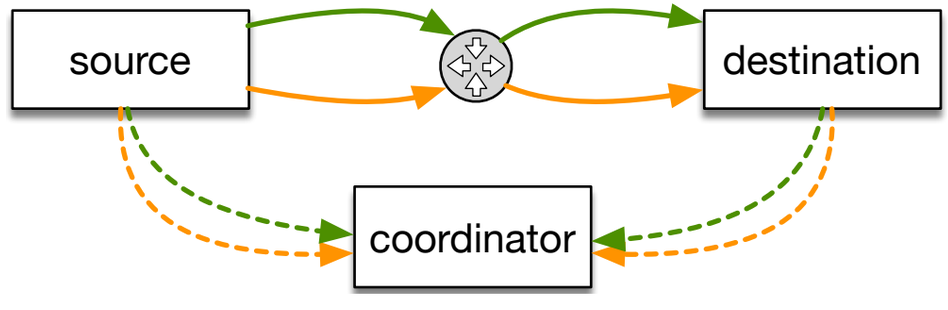

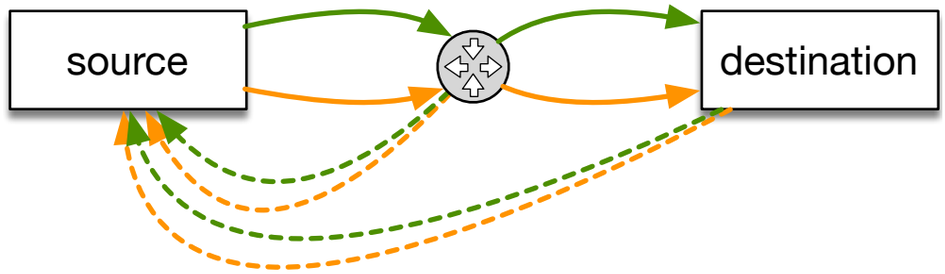

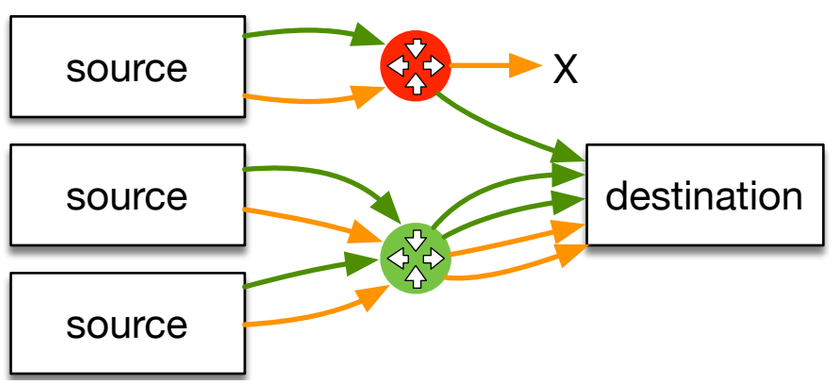

“Can we run the Internet over UDP?”

- UDP encapsulation attractive for new transport protocols
  - (mostly) NAT- and middlebox-compatible header
  - wide availability of APIs in userland
- Lots of current work:
  - WebRTC data channel: SCTP/DTLS/UDP
  - QUIC: new HTTP/2 new transport over UDP
  - ~~SPUD~~ PLUS: universal shim for explicit cooperation
- ***Is this safe?***
  - Widespread operational practice may hinder UDP

RIPE Atlas to the rescue

- No arbitrary TCP/UDP on Atlas…
- …but: traceroute!
  - basic connectivity and first-packet latency with high TTL
- Many probes to many anchors
  - How many probes on UDP blocked networks?
  - Is blocking path- or access-network dependent?

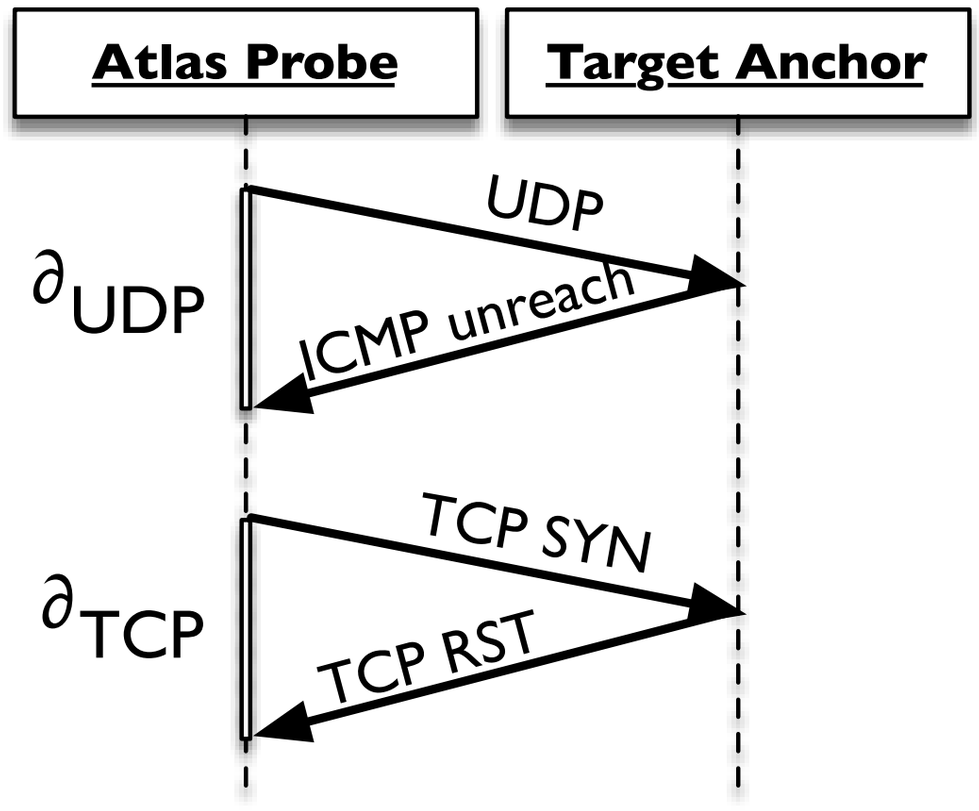

TCP appears more impaired than UDP

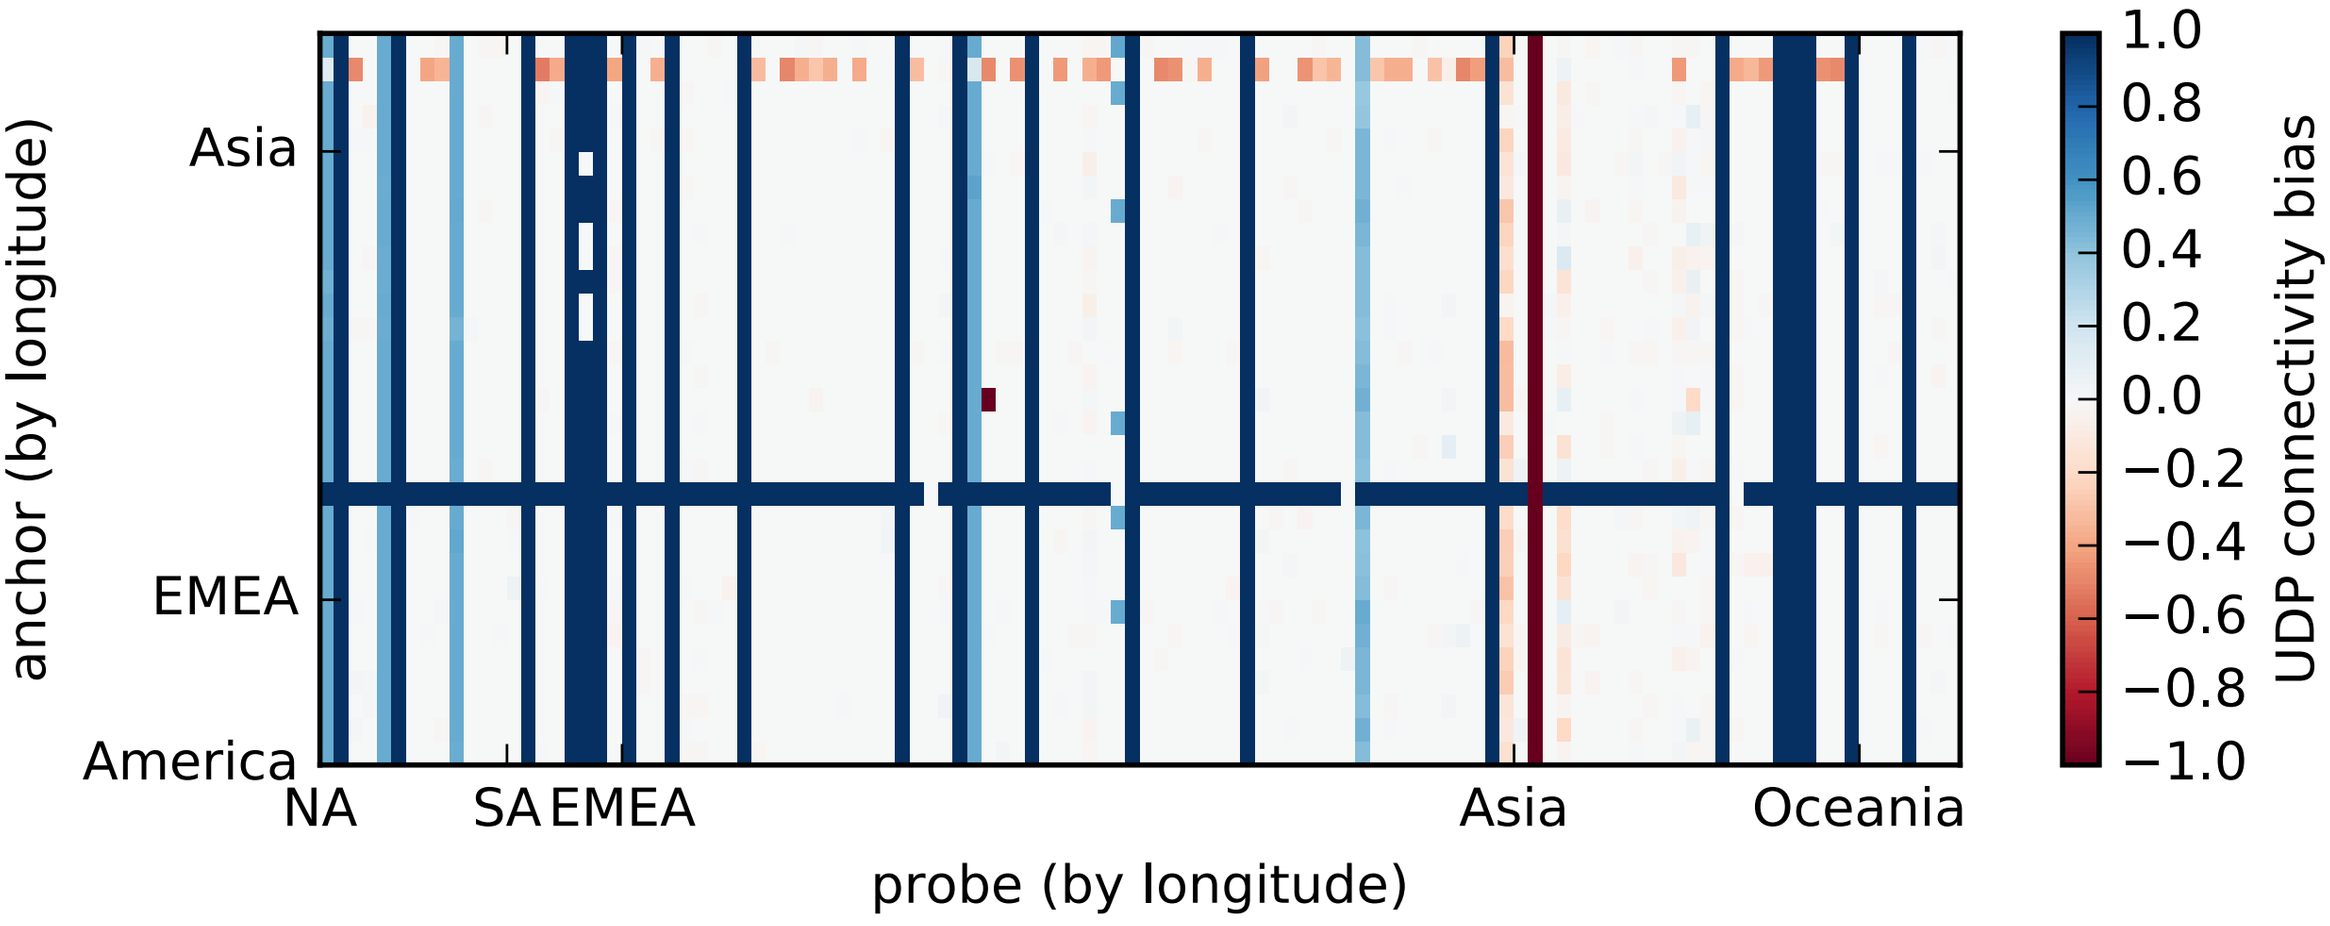

Connectivity, UDP/33435 vs TCP/33435, <= 19 trials, 128 probes to 32 anchors
September 2015

RTT bias mostly probe-dependent

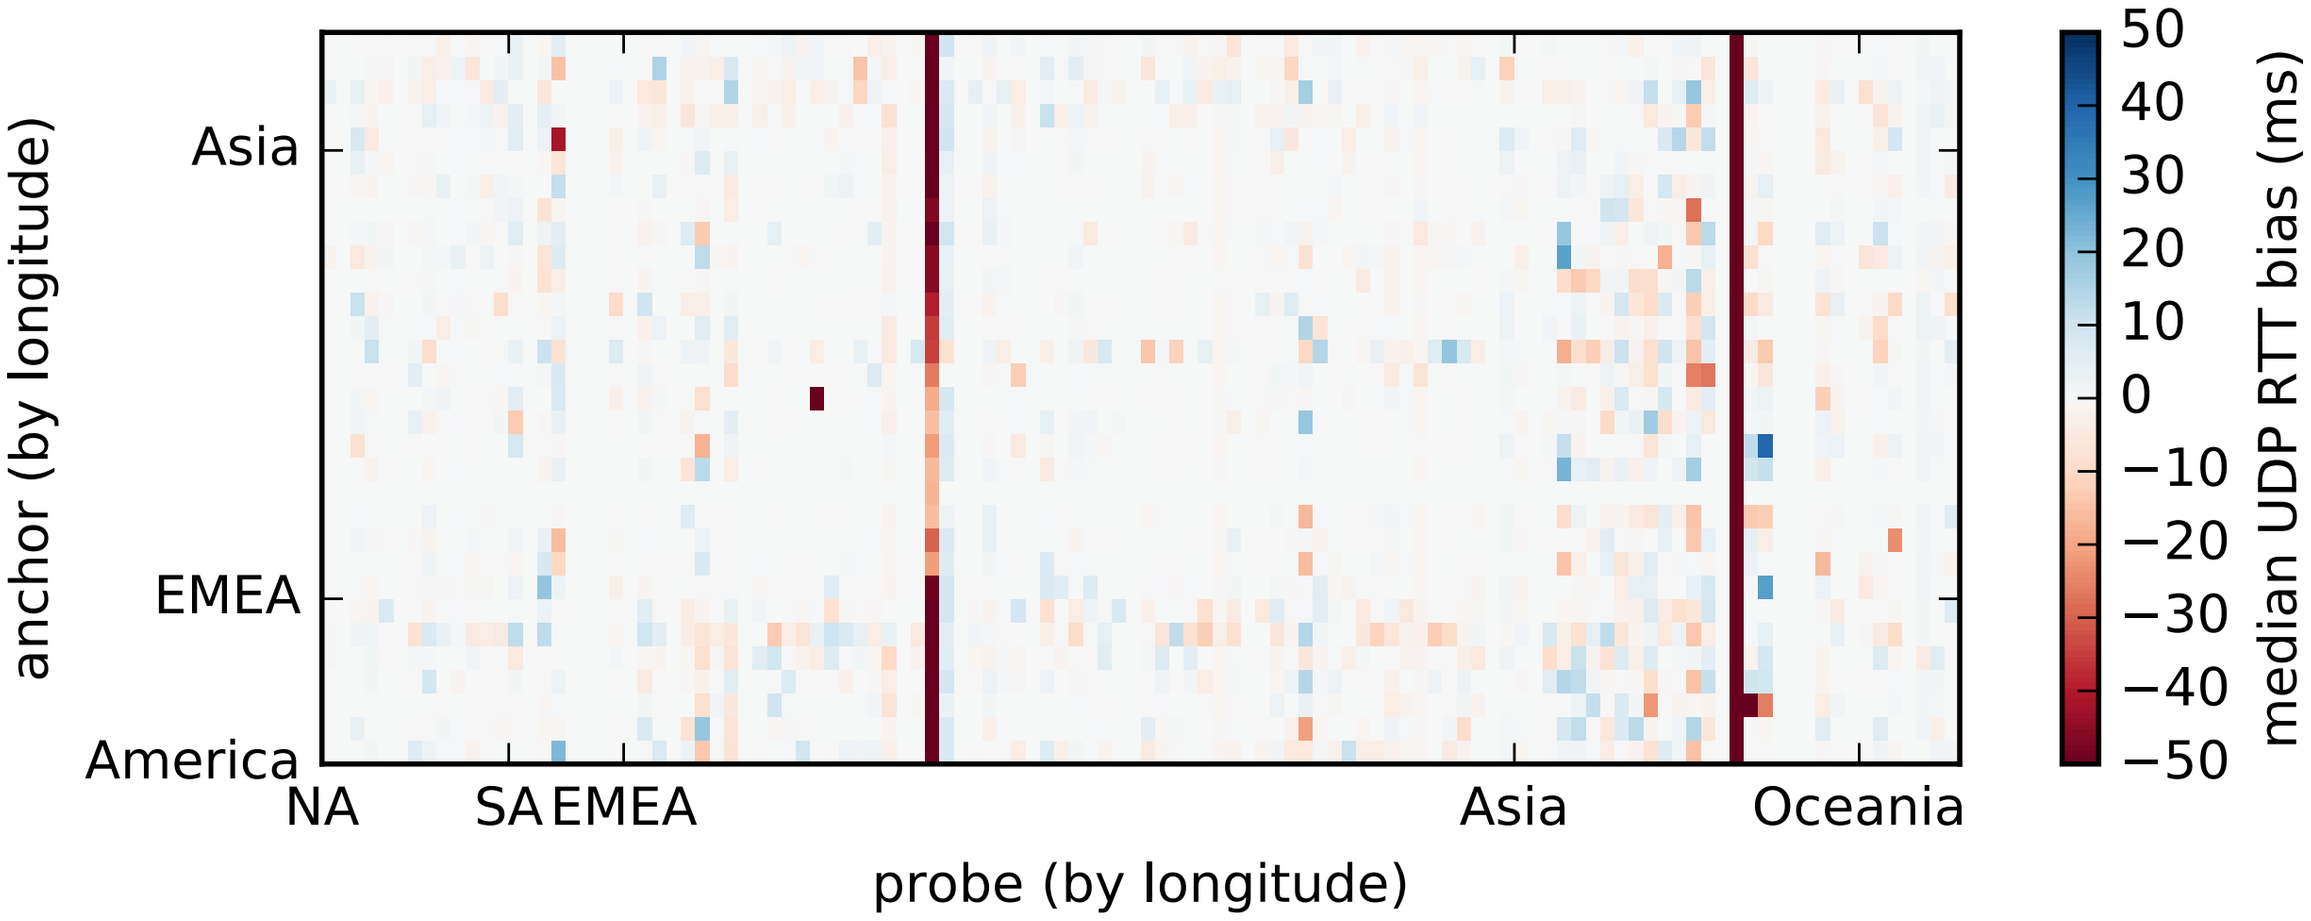

Median RTT bias, UDP/33435 vs TCP/33435, <= 19 trials, 128 probes to 32 anchors
September 2015

More interference with TCP/80

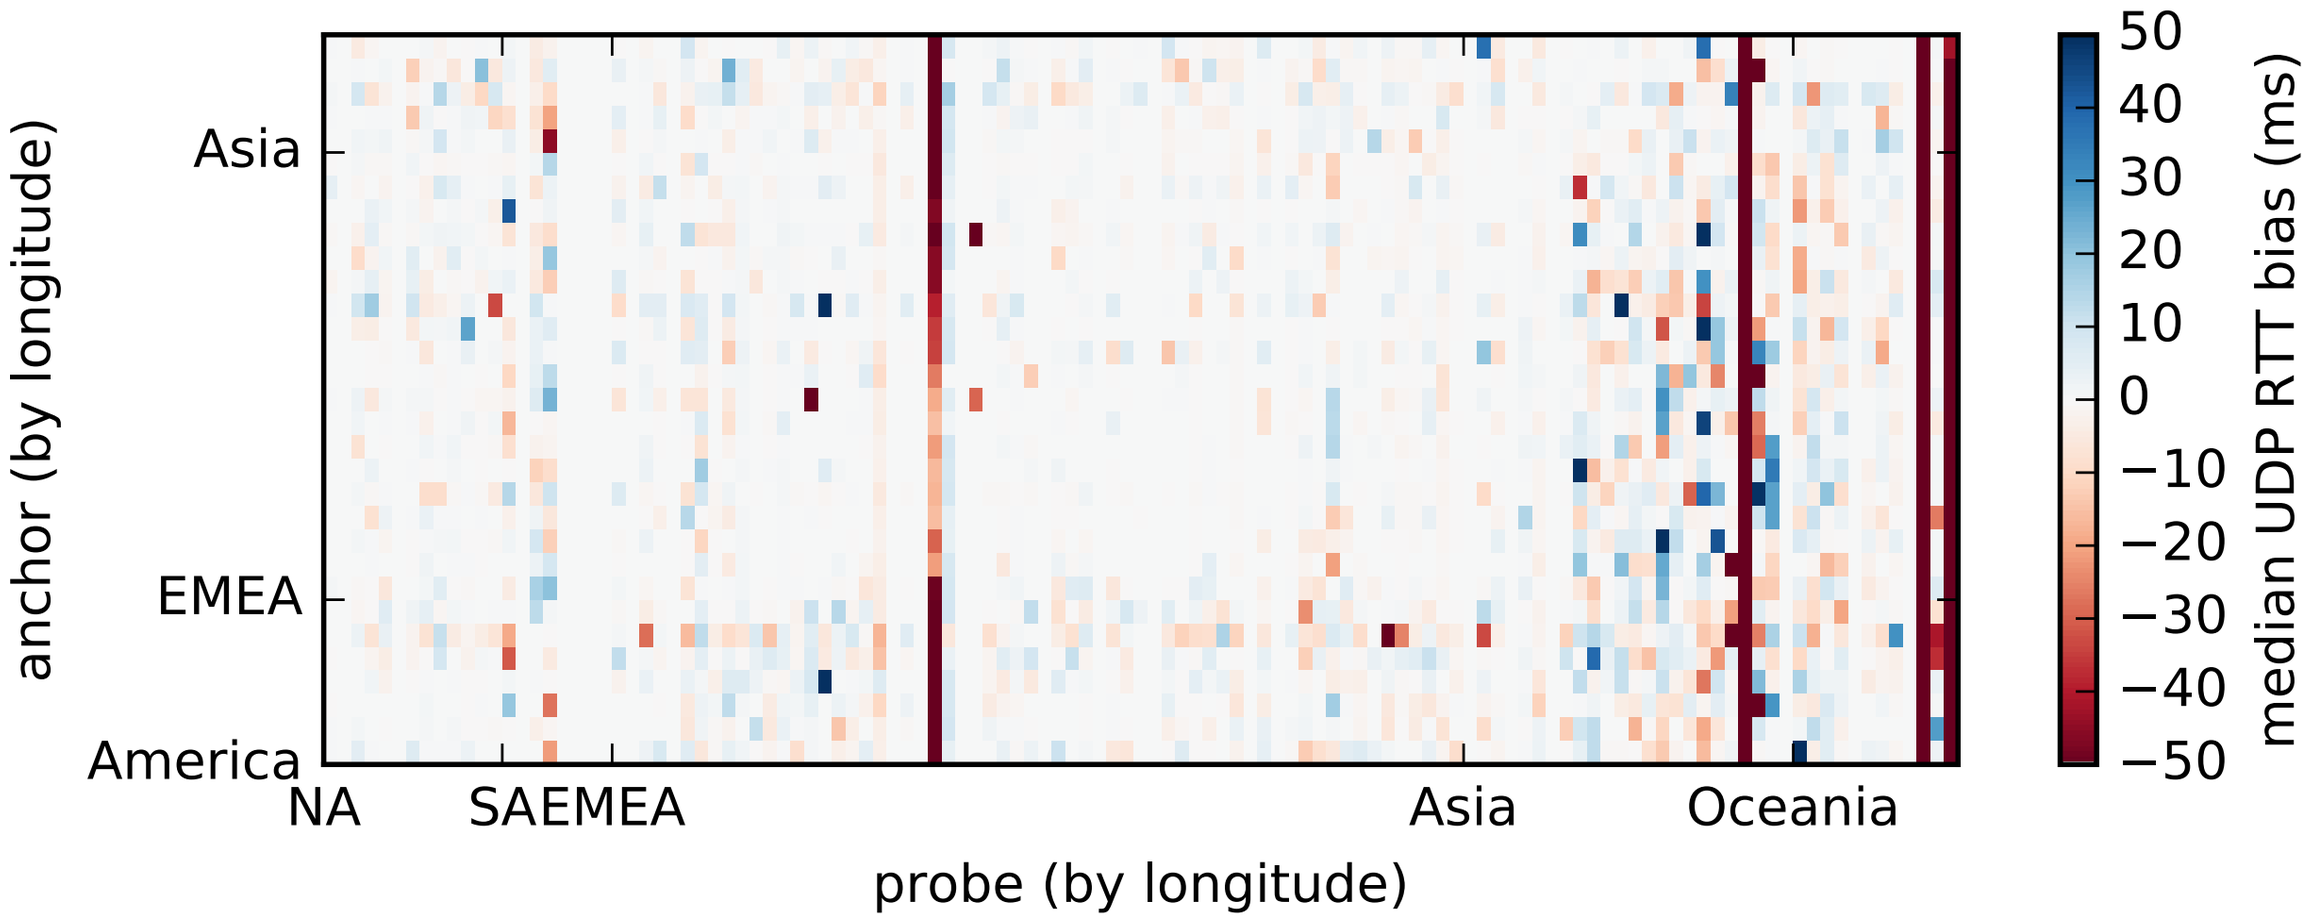

Median RTT bias, UDP/33435 vs TCP/80, <= 19 trials, 128 probes to 32 anchors
September 2015

RTT bias spread tighter on IPv6 than IPv4

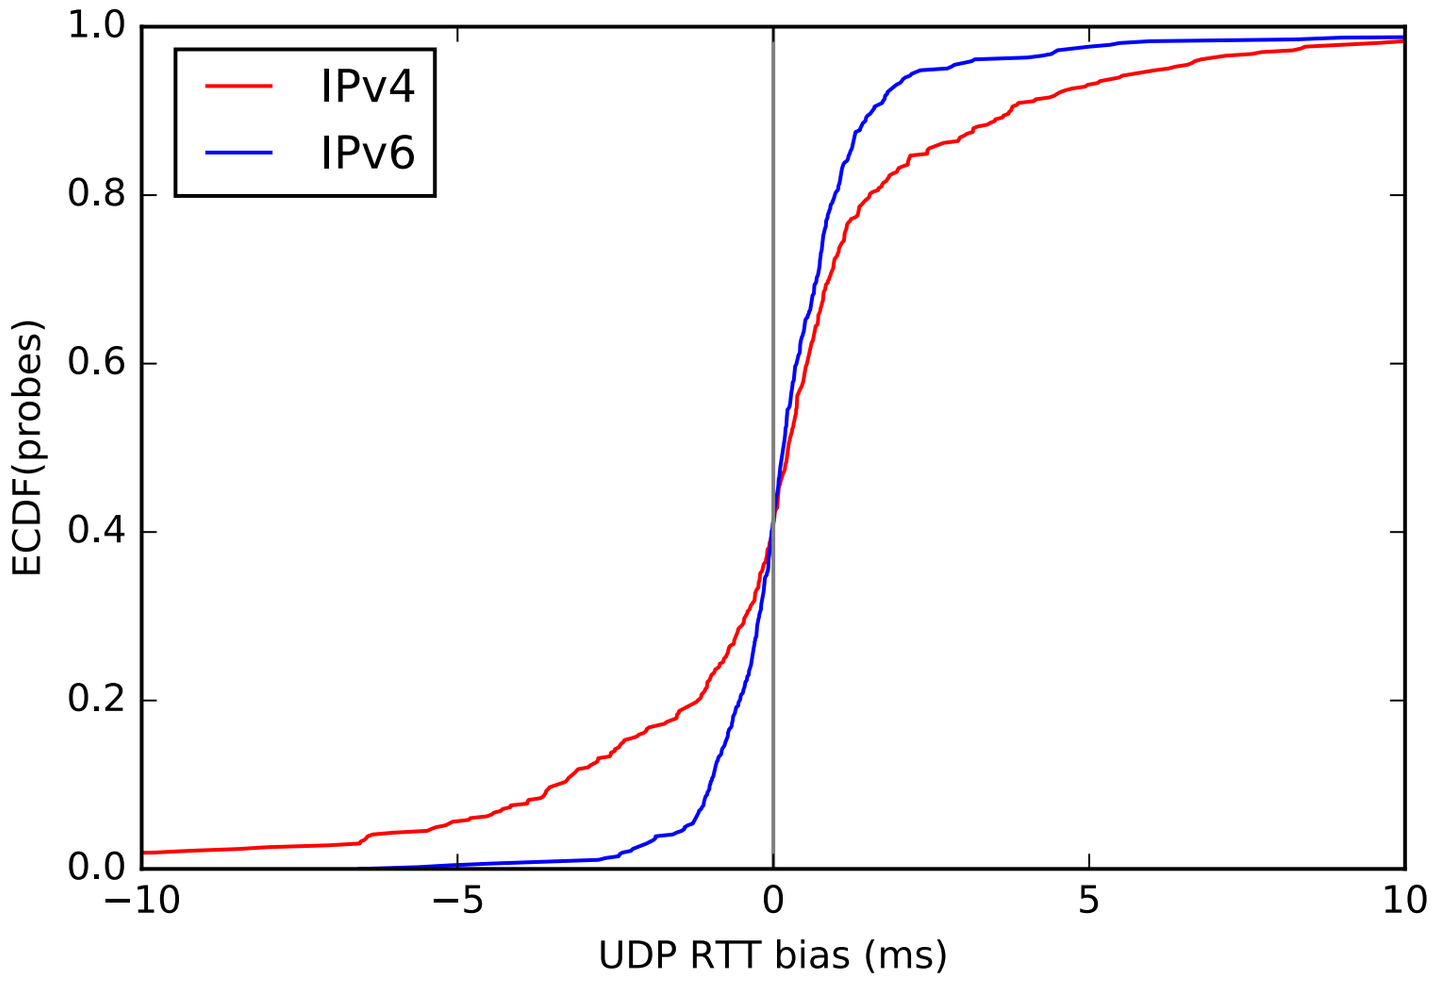

Median RTT bias, UDP/33435 vs TCP/33435, 464 probes to APNIC anchor
Feburary 2016

…not so fast: UDP blocked on one in thirty Atlas probe networks

- Methodology: find all probes
  - that tried to do at least 9 UDP traceroutes in 2015
  - to targets that were up at the time
  - and that showed connectivity via TCP or ICMP
- **2240** probes meet this criterion
  - How many of these never succeeded via UDP?
- **82** probes, largely on networks with marginal connectivity

- Running the Internet over UDP needs a backup for this 3.6%
  - (In line with a 6-7% “QUIC doesn’t work” reported in HOPSRG)

Are larger UDP packets blocked?

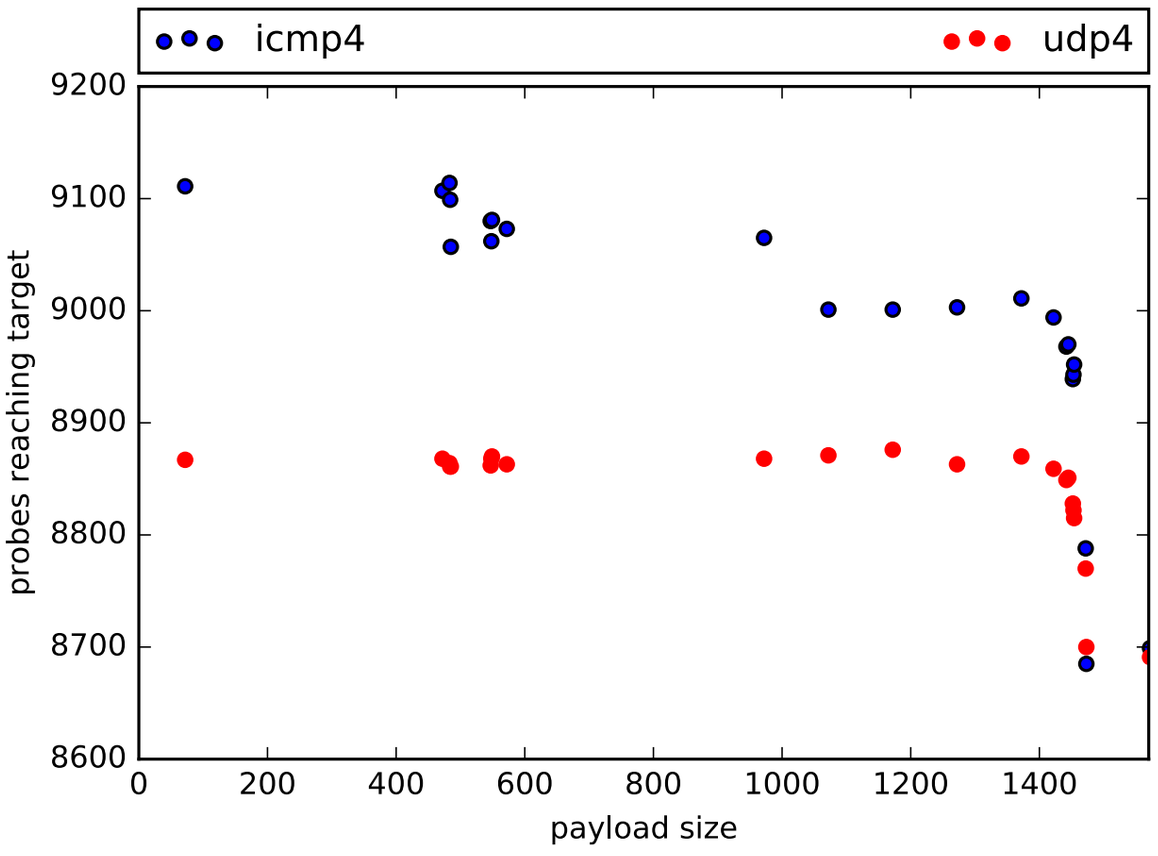

- Apparently not
  - one-off measurement, Mar '16, 9396 probes to one anchor
- No additional blocking after 512, 1024 for IPv4
- (In this short campaign, **296** of **9262** probes (3.2%) may block UDP)

Conclusions

- Atlas useful for estimating UDP connectivity
  - it’s a hack, but it’s a nice one
- Basic UDP connectivity not very broken
  - Works on 29 in 30 (RIPE Atlas) access networks
  - Easy to find out when you’re on the other one
- Running the internet over UDP not prevented by blocking
  - 3% failure is a lot, but fallback helps.

Bonus slide:
Adding new layers to the stack for fun and profit

Why care so much about UDP connectivity?

Path Layer UDP Substrate (PLUS):

- BoF at IETF 96, Berlin, 17-22 July

Enables in-protocol performance measurement headers

- See Mirja Kühlewind's RACI talk (y'day)
- (maybe coming soon to a RIPE BoF near you?)