Can we run the Internet over UDP?

Brian Trammell, ETH Zürich (with Emile Aben, RIPE Labs)

Measurement and Analysis for Protocols (MAP) proposed RG
Monday 4 April 2016 — IETF 95 — Buenos Aires, Argentina

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.

Yes.

Yes.

(Mostly. Probably. Carefully.)

Why would we want to?

- UDP encapsulation attractive for new transport protocols
  - (mostly) NAT- and middlebox-compatible header
  - wide availability of APIs in userland
  - few other real options for evolving the stack
- Lots of current work in the IETF
  - WebRTC (draft-ietf-rtcweb-data-channel)
  - QUIC (draft-tsvwg-quic-protocol)
  - SPUD (draft-trammell-spud-req)
- Many other examples
  - RTP, uTP, basically every gaming application-layer protocol

Reframing the question

- What kinds of different treatment can we expect based *just* on the presence of a UDP header?
  - Connectivity risk (firewall rules, etc.)
  - Differential latency and traffic shaping
  - Reordering, MTU, etc.

- Today's answers: RIPE Atlas measurements
  - Snapshot of work in progress
  - A few more data points to add to others:
    "works fine in (much) more than 90% of cases"

Using Atlas for TCP/UDP connectivity testing

- No arbitrary TCP/UDP packets...
- …but traceroute uses ICMP, TCP, or UDP on the forward path
  - can measure basic connectivity and first-packet latency
- Many-to-many measurements: isolate path- from access-impairment
- Many-to-one measurements: find probes on UDP-blocking networks
- Not perfect, but better than nothing

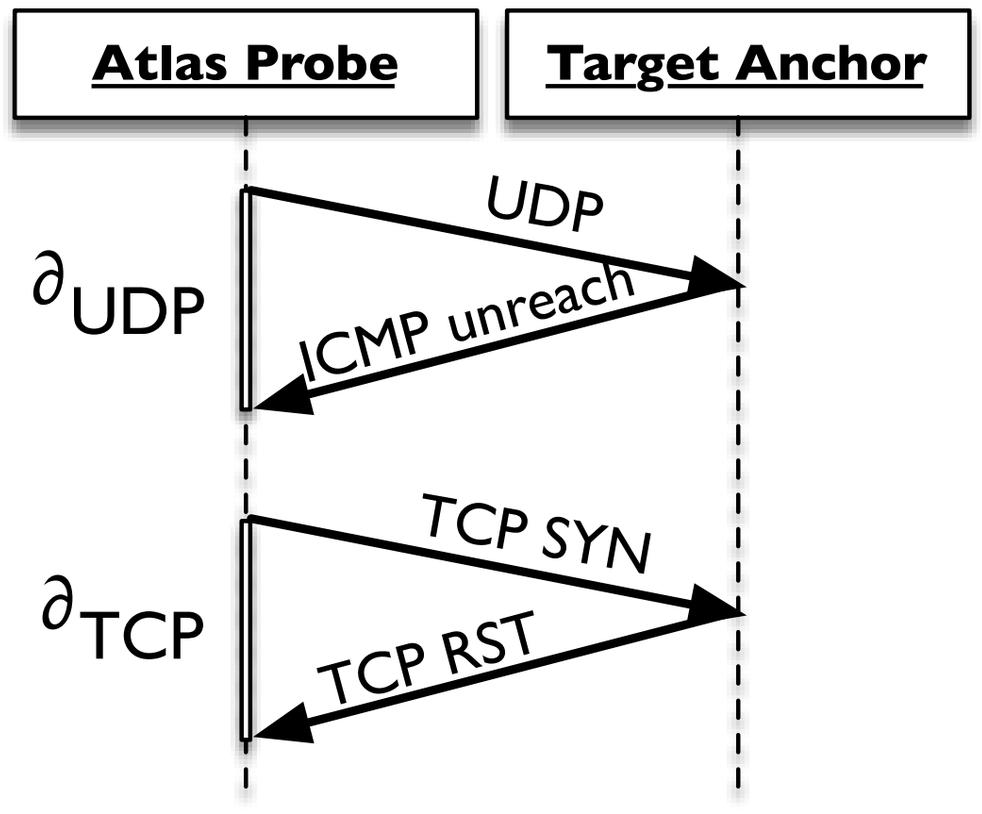

How many networks block UDP?

- **2240** probes did UDP traceroutes* in 2015
- **82** (3.6%) never succeeded using UDP: probable blocking
  - Many of these on networks with marginal connectivity
- Selection bias: Atlas probe hosts tend to be network geeks or network geek adjacent.
  - Enterprise networks under-represented
  - Few mobile/wireless-connected probes

* more than 9 samples to targets that were up, excluding non-ICMP/non-TCP connected probes.

How much slower is UDP?

- No systematic first-packet penalty
- Some probes slower
  - 128 probes to 32 anchors, Sep '15
- IPv6 less variable
  - 461 probes to APNIC, Feb '16
- Widespread meddling with TCP (esp. 80)

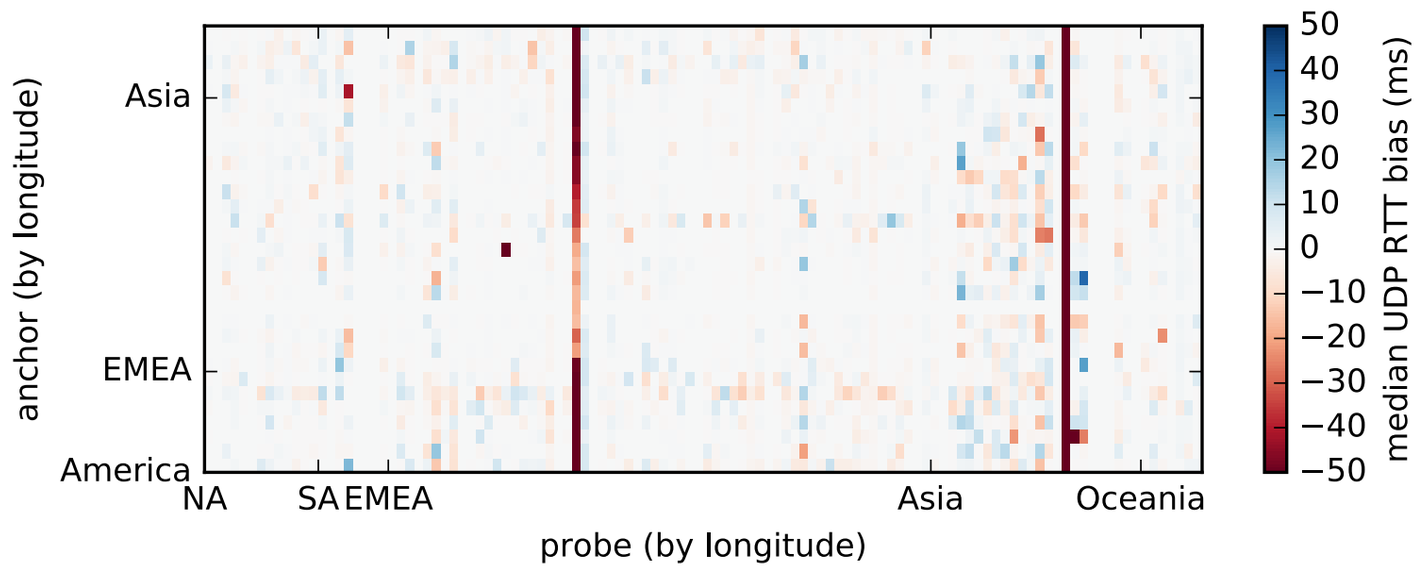

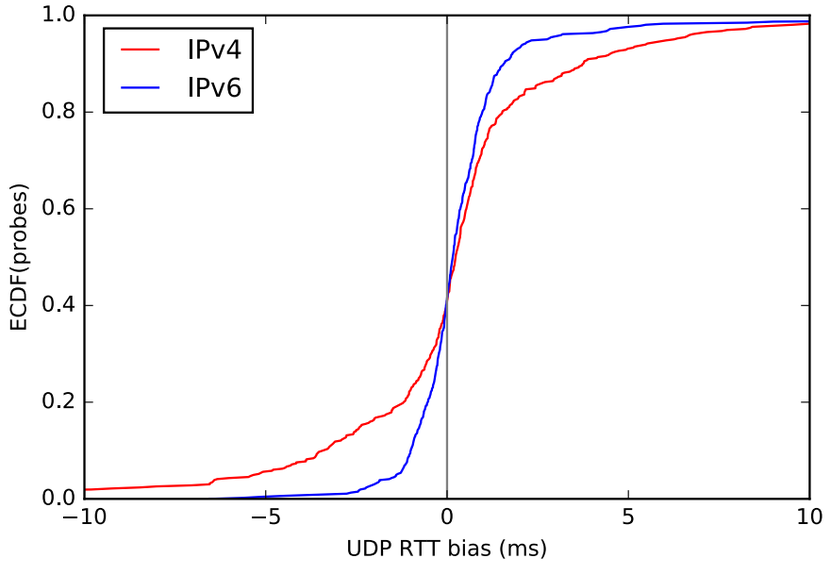

Are larger UDP packets blocked?

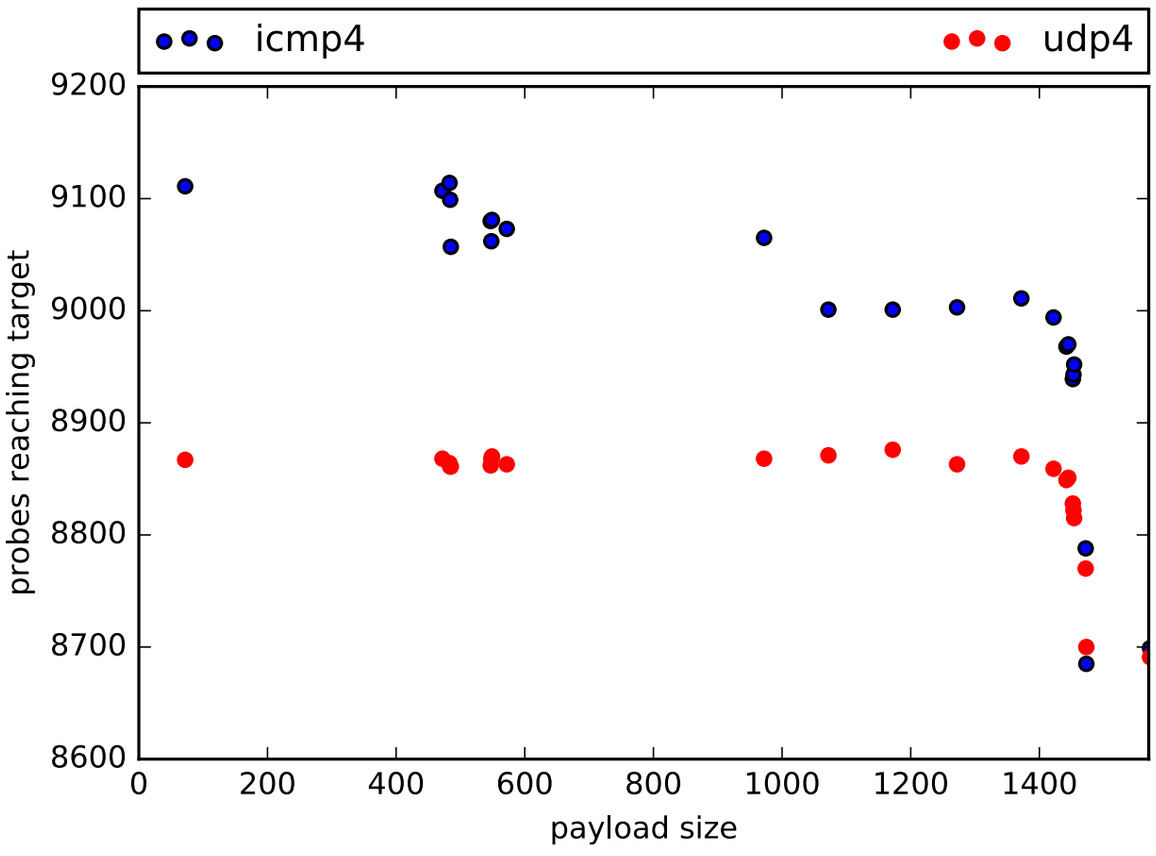

- Apparently not
  - one-off measurement, Mar '16, 9396 probes to one anchor
- No additional blocking after 512, 1024 for IPv4
- (In this short campaign, **296** of **9262** probes (3.2%) may block UDP)

Can we run the internet over UDP?

- Two more datapoints: 82/2240 (3.6%) and 296/9262 (3.2%) of Atlas probes may to be on UDP-blocked networks.
- No latency difference or protocol-dependent MTU
- Works on twenty-nine access networks in thirty
  - It's easy to tell when you're on the other one:
    trivial fallback mechanisms are useful for UDP encapsulations

- Work in progress: full-mesh measurements for loss rates and achievable bandwidth comparison.
- Watch https://mami-project.eu for more