Tracing the Path to YouTube -
A Quantification of Path Lengths and Latencies towards Content Caches

Accepted for publication in IEEE Communications Magazine
(Pre-print: `http://in.tum.de/~doan/2018-yt-traces.pdf`)

Trinh Viet Doan, Ljubica Pajević,
Vaibhav Bajpai, Jörg Ott

Chair of Connected Mobility
Technical University of Munich

RIPE 77, Amsterdam
October 17, 2018

Introduction

Introduction

Motivation

Previous work [2]:

- Measuring YouTube performance for popular videos
- Performance over IPv6 is worse than over IPv4
- **Speculation: Content caches not dual-stacked?**

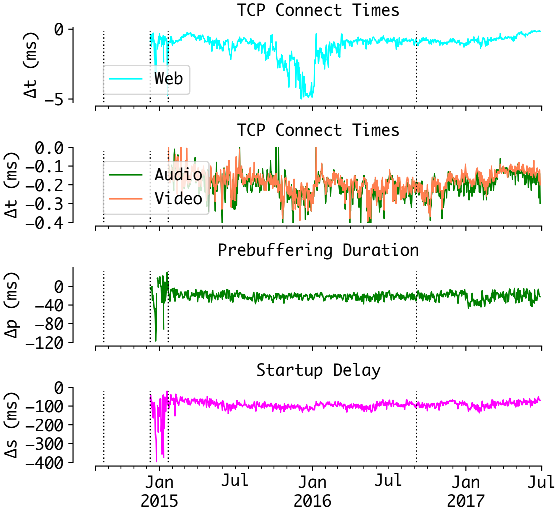

Figure 1: Difference of YouTube performance metrics over IPv4 and IPv6

Introduction

Research Questions

1. **How far are content caches from users?**
2. **How much benefit do these caches provide?**
3. **How do these metrics compare quantitatively over IPv4 and IPv6?**

Methodology

Methodology

Measurement Setup

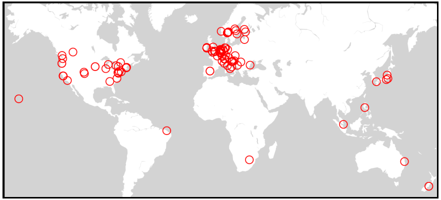

Figure 2: Map of SamKnows probes

Figure 3: Example of measurement probe: SamKnows Whitebox 8.0[1]

- $\approx$ 100 probes deployed around the world since 2014
- Deployed in **dual-stacked** residential networks, NRENs, business networks, research labs, data centers, IXPs, ...
- Active measurement studies from **fixed-line networks**

1 https://blog.samknows.com/new-testing-superfast-broadband-27a7abcf1303 [accessed 2018-08-07]

Methodology

Targets and Metrics

- Hourly `traceroute` measurements over IPv4 & IPv6
  - Using `scamper` [3] for `paris traceroute` over ICMP
- Targets: YouTube media servers
  - Media servers identified by `youtube` test [1] that mimics video streaming from YouTube
  - DNS resolution for this streaming directly on the probe
    ⇒ Redirected to best/closest cache, determined by YouTube
  - Identified IP addresses of media servers to `scamper` for measurements
  - Time period: since May 2016

Analysis

Analysis

Paths

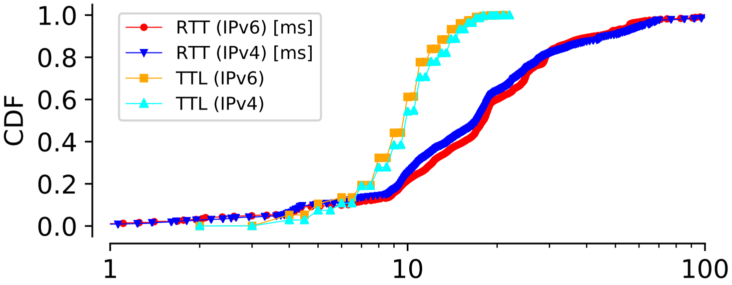

Figure 4: CDF of median IP path TTL and RTT

- Comparable number of paths observed
  - 78% with TTL $\leq$ 12 (IPv4), $\leq$ 11 (IPv6)
    $\rightarrow$ IPv6 paths more often shorter
  - 74% with RTT $\leq$ 25 ms (IPv4), 72% over (IPv6)
    $\rightarrow$ IPv6 more often slower

However, no direct comparison possible ⇒ look at *destination pairs*

unit_id	dtime	source	destination	status	ttl	endpoint	rtt
239416	2016-06-07 16:45:35	2001:67c:_:_:_:_:fef0:d612	2a00:1450:400f:f::a	COMPLETED	9	2a00:1450:400f:f::a	10.522
239416	2016-06-07 16:45:36	10.0.1.3	83.255.235.81	COMPLETED	7	83.255.235.81	13.178

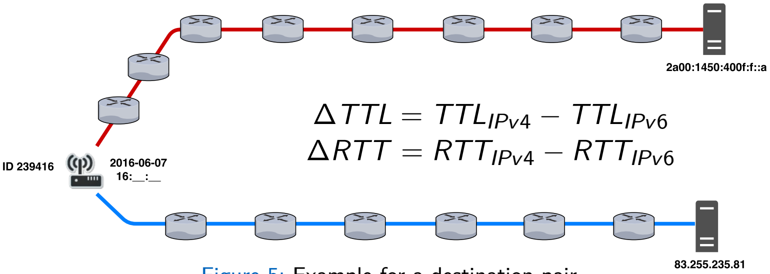

$$\Delta TTL = TTL_{IPv4} - TTL_{IPv6}$$
$$\Delta RTT = RTT_{IPv4} - RTT_{IPv6}$$

Figure 5: Example for a destination pair

Analysis

Destination Pairs: General

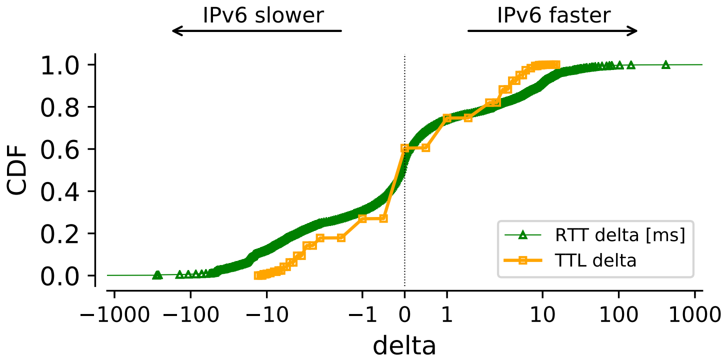

Figure 6: CDF of median destination pair deltas

- TTL:
  - 27% with $\Delta TTL < 0$
  - 33% with $\Delta TTL = 0$
  - 40% with $\Delta TTL > 0$
- RTT:
  - $\approx$ 50% with $\Delta RTT < 0$
  - $\approx$ 50% with $\Delta RTT > 0$

Analysis

Destination Pairs: General

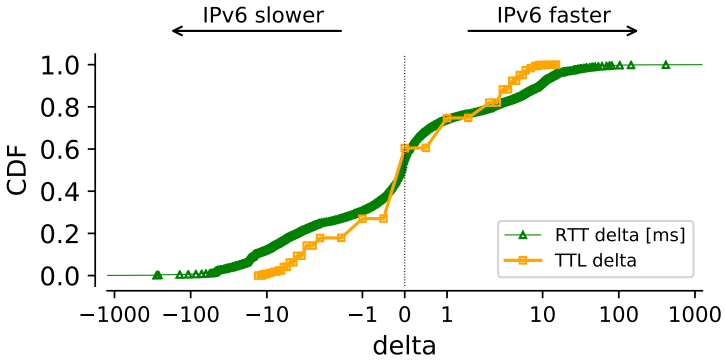

Figure 6: CDF of median destination pair deltas

- Overall:
  - TTL: 91% within [-5; +5]
  - RTT: 91% within [-20; +20] ms

Analysis
Content Caches

- Content caches usually deployed within ISP networks
- In close proximity to users to reduce latency

Analysis

Content Caches

- Content caches usually deployed within ISP networks
- In close proximity to users to reduce latency
- How to identify caches?
  - Matching AS numbers for source and destination
    $\rightarrow$ src ASN == dst ASN
  - Reverse DNS lookups of destination IP addresses to retrieve human-readable hostnames
    $\rightarrow$ keywords: `cache` or `ggc`
  - Lookups using RIPEstat[2]

[2] https://stat.ripe.net/

Analysis

Destination Pairs: Caches

Possible scenarios for identification of caches when comparing between different address families.

IPv4 \ IPv6	Cache	No Cache
Cache	both O	IPv4 only △
No Cache	IPv6 only □	neither ◇

Analysis

Destination Pairs: Caches

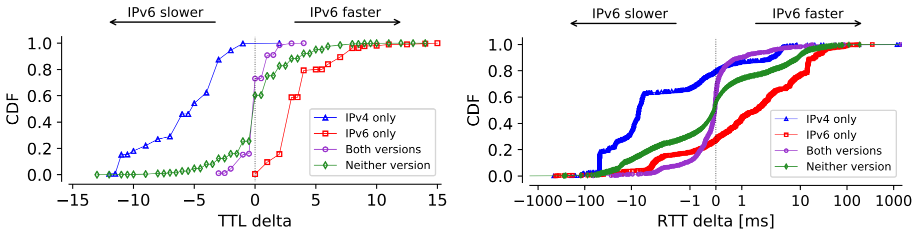

Figure 7: CDF of median destination pair deltas (split)

Analysis

Destination Pairs: Caches

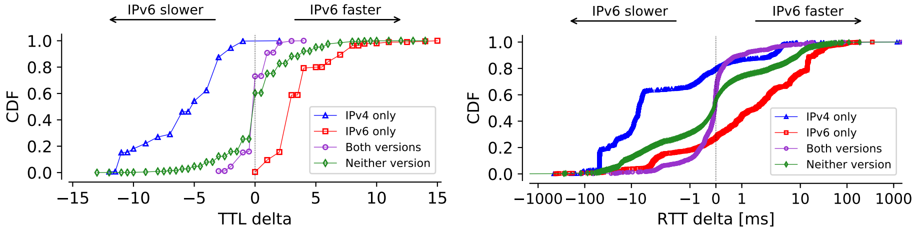

Figure 7: CDF of median destination pair deltas (split)

- IPv4 cache only (△): shifted to left side; RTT lower over IPv4 for ≈ 80%

Analysis

Destination Pairs: Caches

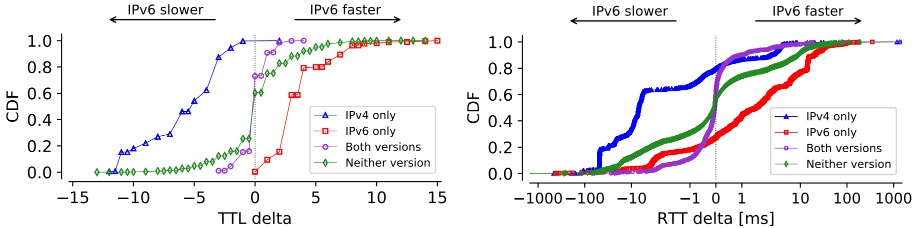

Figure 7: CDF of median destination pair deltas (split)

- IPv4 cache only (△): shifted to left side; RTT lower over IPv4 for ≈ 80%
- IPv6 cache only (□): paths shorter to IPv6 caches compared to IPv4 no-cache destinations; yet still higher RTTs in most cases

Analysis

Destination Pairs: Caches

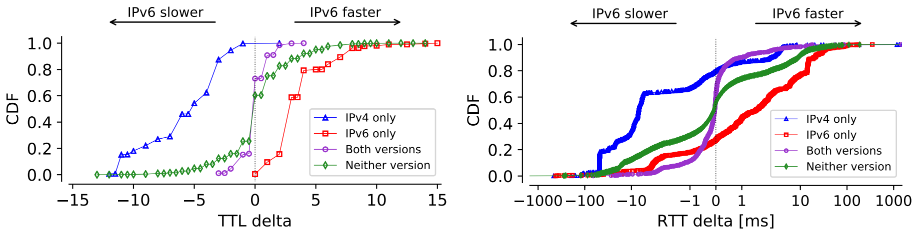

Figure 7: CDF of median destination pair deltas (split)

- IPv4 cache only (△): shifted to left side; RTT lower over IPv4 for ≈ 80%
- IPv6 cache only (□): paths shorter to IPv6 caches compared to IPv4 no-cache destinations; yet still higher RTTs in most cases
- Both (**O**): deltas converging towards zero; 60% of the time faster over IPv4, 40% of the time faster over IPv6, however ≈ 80% within [-1,+1] ms

Analysis

Content Caches: Distributions

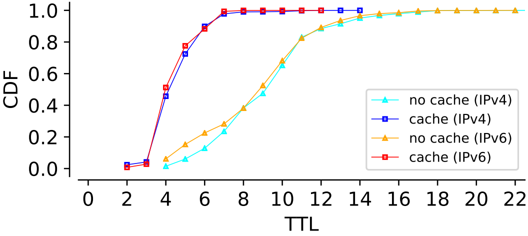

Figure 8: CDF of cache vs no cache path values for all traces (TTL)

- $\approx$ 100% of ISP caches reachable within 7 IP hops
- Cache vs no cache
  - $\leq$ 6 IP hops for $\approx$ 90% of the cache measurements
  - $\leq$ 12 IP hops for $\approx$ 89% of the no cache measurements

Analysis

Content Caches: Distributions

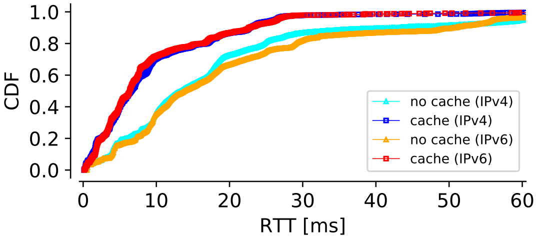

Figure 9: CDF of cache vs no cache path values for all traces (RTT)

- Majority of caches reachable within 20 ms (87%)
- For 80% of the measurements (no cache → cache)
  - IPv4: 25 ms → 17 ms; ≈ 1/3 improvement
  - IPv6: 29 ms → 16 ms; ≈ 1/2 improvement

Conclusion

Conclusion

1. **Distance of caches?**
   - Caches within 6 IP hops and 20 ms over both IPv4 and IPv6

Conclusion

1. **Distance of caches?**
   - Caches within 6 IP hops and 20 ms over both IPv4 and IPv6

2.&3. **Benefits of caches? Performance over IPv4 vs IPv6?**
   - *IP path length:* up to 6 hops lower (i.e. ½) for both IPv4 and IPv6
   - *Latency:* up to $\approx$ 10 ms lower; relative improvement of IPv6 caches higher
     - IPv4: up to 8 ms (⅓); IPv6: up to 13 ms (½)
   - *Surprise:* IPv6 caches higher RTT than IPv4 non-caches despite lower TTL

Conclusion

1. **Distance of caches?**
   - Caches within 6 IP hops and 20 ms over both IPv4 and IPv6

2.&3. **Benefits of caches? Performance over IPv4 vs IPv6?**
   - *IP path length:* up to 6 hops lower (i.e. ½) for both IPv4 and IPv6
   - *Latency:* up to ≈ 10 ms lower; relative improvement of IPv6 caches higher
     - IPv4: up to 8 ms (⅓); IPv6: up to 13 ms (½)
   - *Surprise:* IPv6 caches higher RTT than IPv4 non-caches despite lower TTL

Takeaways: Room for improvement regarding IPv6 content delivery:

- Ensure caches are dual-stacked within ISP networks (see △ and □ cases),
- Optimize delivery regarding performance, routing, forwarding, ...
- Caches are not the end of the story regarding IPv4 and IPv6 discrepancy

Conclusion

1. **Distance of caches?**
   - Caches within 6 IP hops and 20 ms over both IPv4 and IPv6

2.&3. **Benefits of caches? Performance over IPv4 vs IPv6?**
   - *IP path length:* up to 6 hops lower (i.e. ½) for both IPv4 and IPv6
   - *Latency:* up to ≈ 10 ms lower; relative improvement of IPv6 caches higher
     - IPv4: up to 8 ms (⅓); IPv6: up to 13 ms (½)
   - *Surprise:* IPv6 caches higher RTT than IPv4 non-caches despite lower TTL

Takeaways: Room for improvement regarding IPv6 content delivery:

- Ensure caches are dual-stacked within ISP networks (see △ and □ cases),
- Optimize delivery regarding performance, routing, forwarding, ...
- Caches are not the end of the story regarding IPv4 and IPv6 discrepancy

Dataset and code publicly available at:
https://github.com/tv-doan/youtube-traceroutes

✉ doan@in.tum.de

References

[1] Ahsan, S., Bajpai, V., Ott, J., and Schönwälder, J.
Measuring YouTube from Dual-Stacked Hosts.
In *PAM* (2015), vol. 8995 of *Lecture Notes in Computer Science*, Springer, pp. 249–261.
https://doi.org/10.1007/978-3-319-15509-8_19.

[2] Bajpai, V., Ahsan, S., Schönwälder, J., and Ott, J.
Measuring YouTube Content Delivery over IPv6.
Computer Communication Review 47, 5 (2017), 2–11.
http://doi.acm.org/10.1145/3155055.3155057.

[3] Luckie, M. J.
Scamper: a Scalable and Extensible Packet Prober for Active Measurement of the Internet.
In *Internet Measurement Conference* (2010), ACM, pp. 239–245.
http://doi.acm.org/10.1145/1879141.1879171.

Backup Slides

Analysis

Temporal View

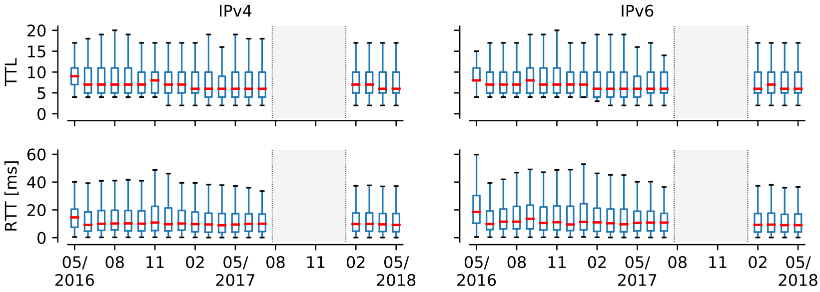

Figure 10: Boxplots of path TTL and RTT values, aggregated by month

- Median TTL across all months: 7 IP hops (both IPv4 and IPv6)
- Median RTT across all months: 9.9 ms (IPv4), 10.7 ms (IPv6)

Analysis

Intermediate IP Hops

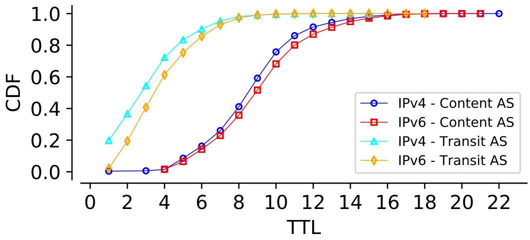

Figure 11: CDF of all TTL values by version and AS type[3]

$TTL \approx 7$ as a separator for both IPv4 and IPv6:

- Transit/Access ASes: $TTL \leq 7$ for 93%
- Content ASes: $TTL \geq 7$ for 85%

[3] CAIDA AS Classification: https://www.caida.org/data/as-classification/

Analysis

Intermediate IP Hops

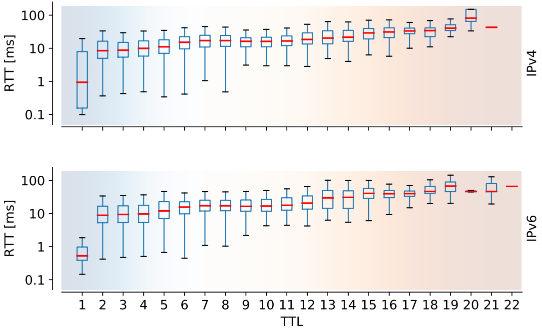

Figure 12: Boxplots of RTT by TTL

- Destination reached in $TTL < 7$ (blue gradient):
  ISP cache in Transit/Access AS
- Destination reached in $TTL > 7$ (orange gradient):
  origin content server in Content AS