Hochschule für Angewandte Wissenschaften Hamburg
Hamburg University of Applied Sciences

Freie Universität Berlin

On the Interplay between TLS Certificates and QUIC Performance

Marcin Nawrocki, Pouyan Fotouhi Tehrani, Raphael Hiesgen,
Jonas Mücke, Thomas C. Schmidt, Matthias Wählisch

{marcin.nawrocki, jonas.muecke, m.waehlisch}@fu-berlin.de
pouyan.fotouhi.tehrani@fokus.fraunhofer.de
{raphael.hiesgen, t.schmidt}@haw-hamburg.de

QUIC handshake design goal 1: Reduced round-trips.

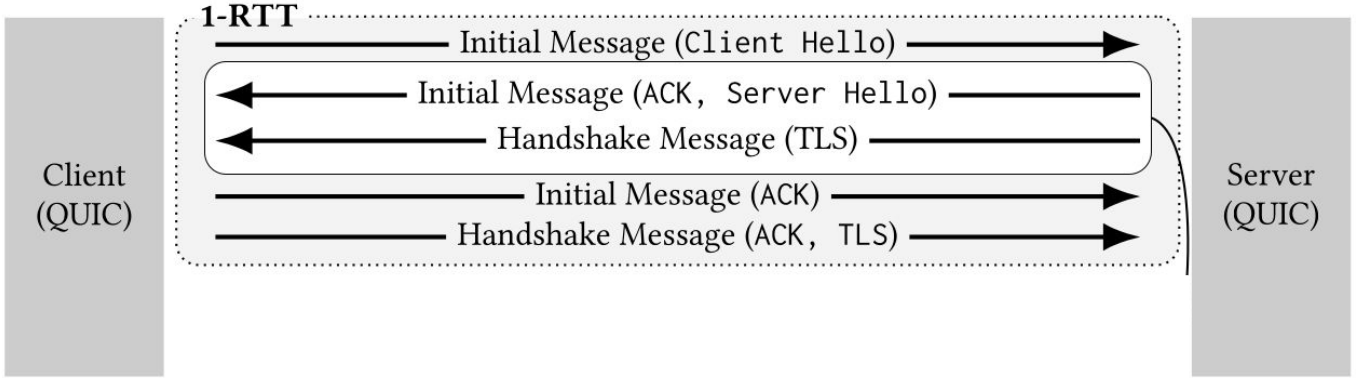

QUIC handshake design goal 2: Reduced amplification.

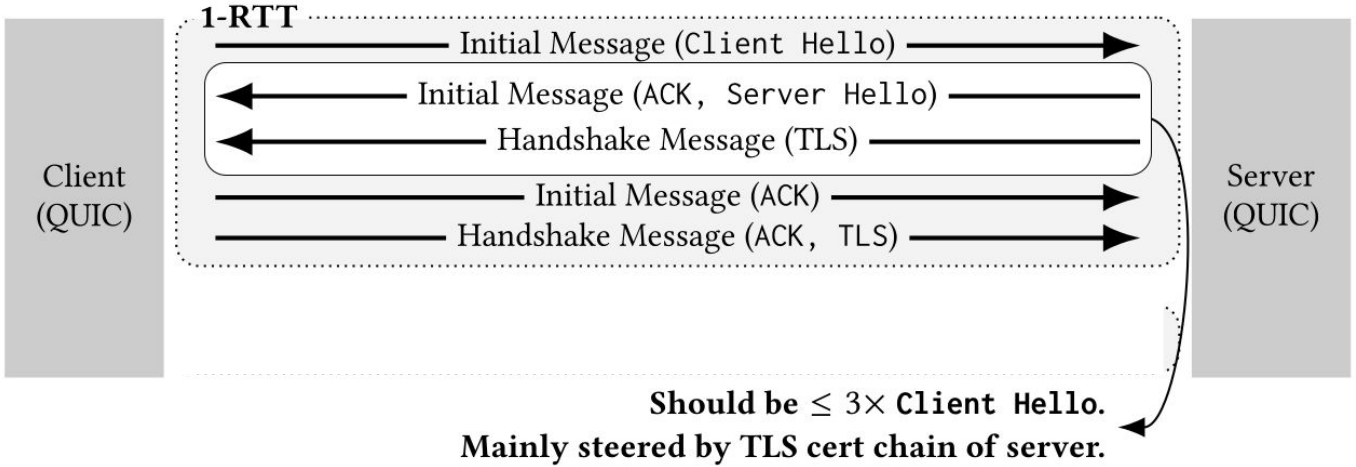

Multi-RTT handshakes validate clients but are inefficient.

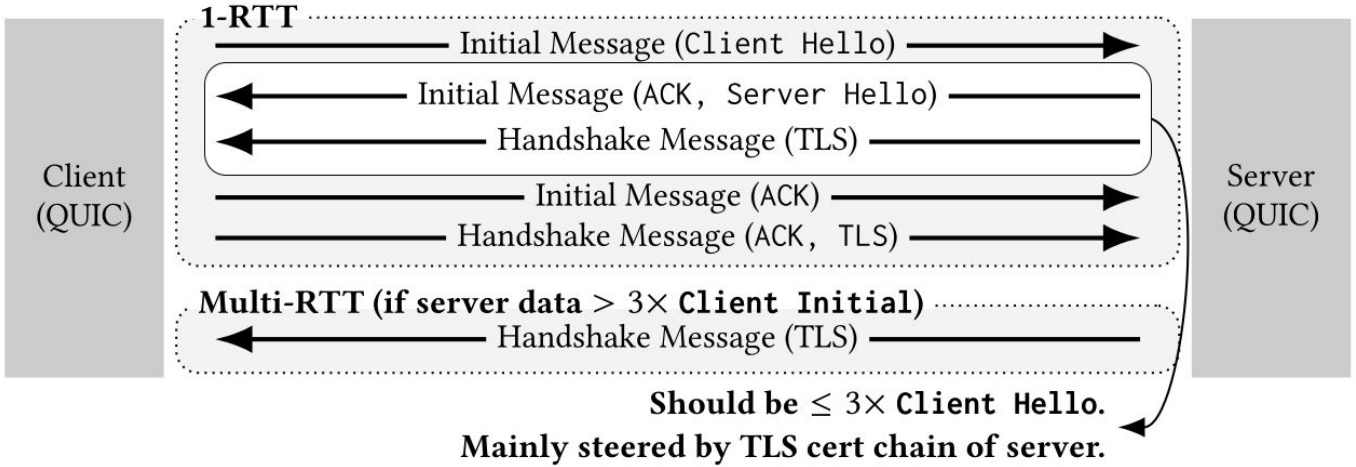

A lot of TLS data? Certificates are delivered as a chain.

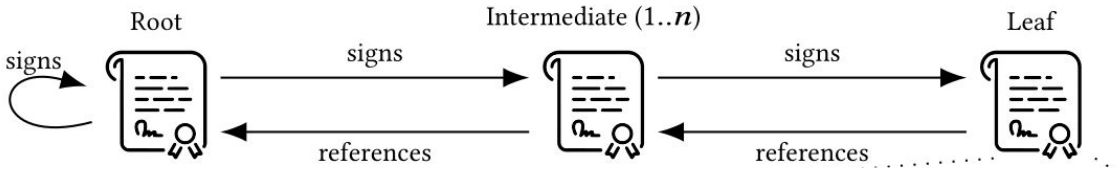

A lot of TLS data? Large keys, alternative names, *etc.*

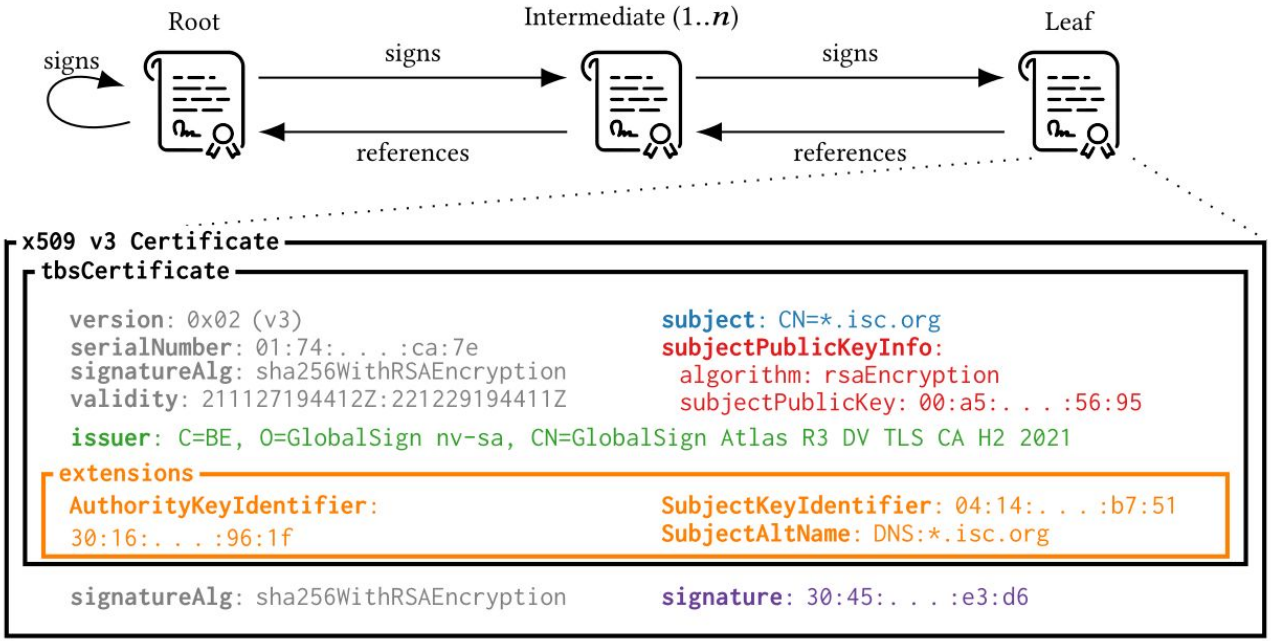

Agenda

Hypergiants purposefully ignore the anti-amplification.

This enables clients to estimate a precise RTT.

TLS data still interferes with QUIC performance.

Improvements such as compression hard to integrate.

Incomplete QUIC handshakes amplify up to 45x.

Server retransmissions can lead to adverse effects.

Methodology: Active scans with open-source tools.

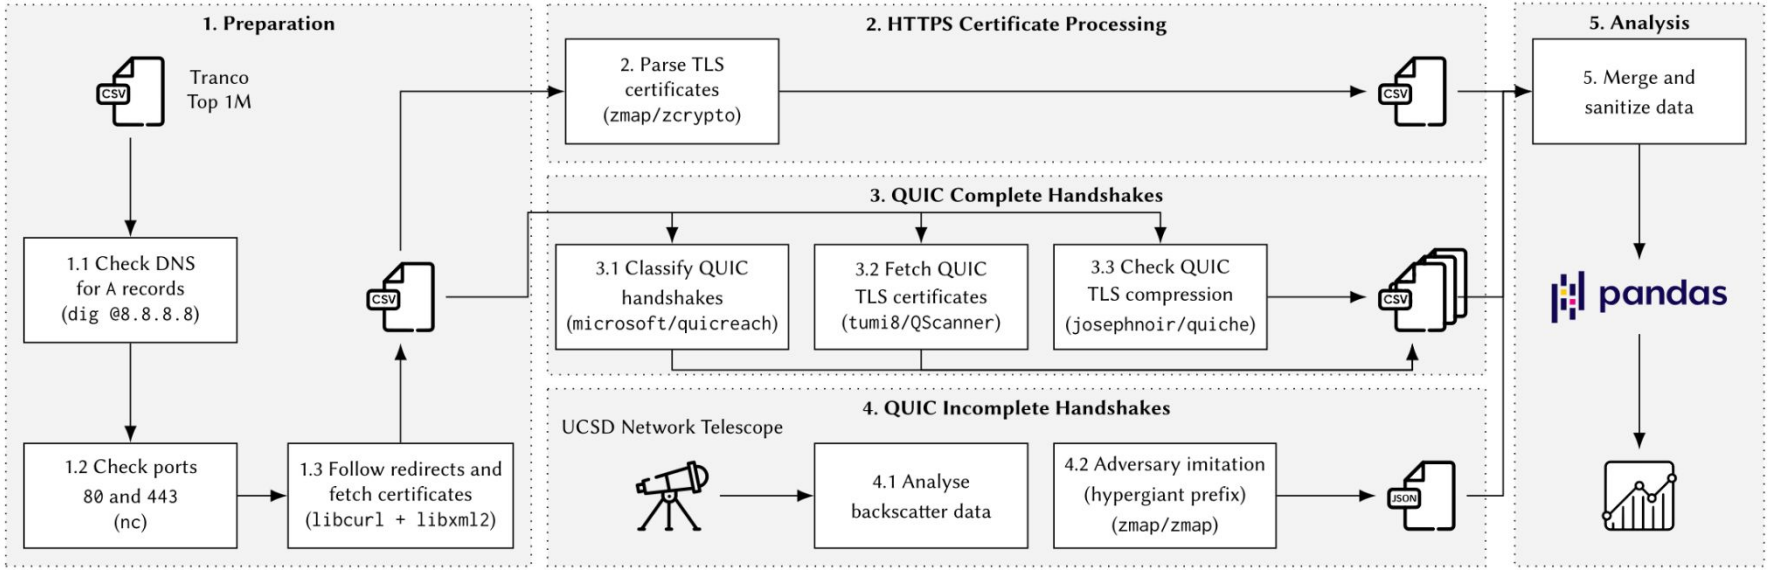

Methodology: Active scans with open-source tools.

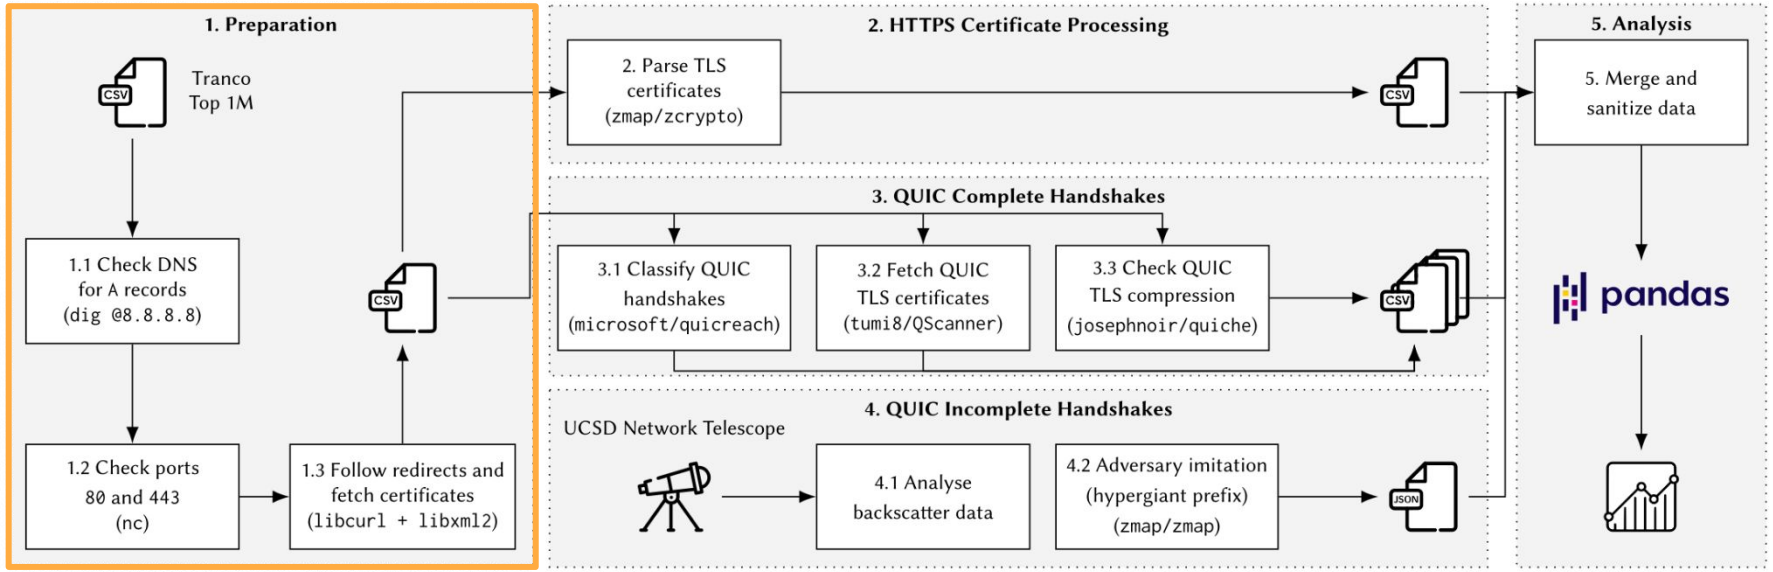

Methodology: Active scans with open-source tools.

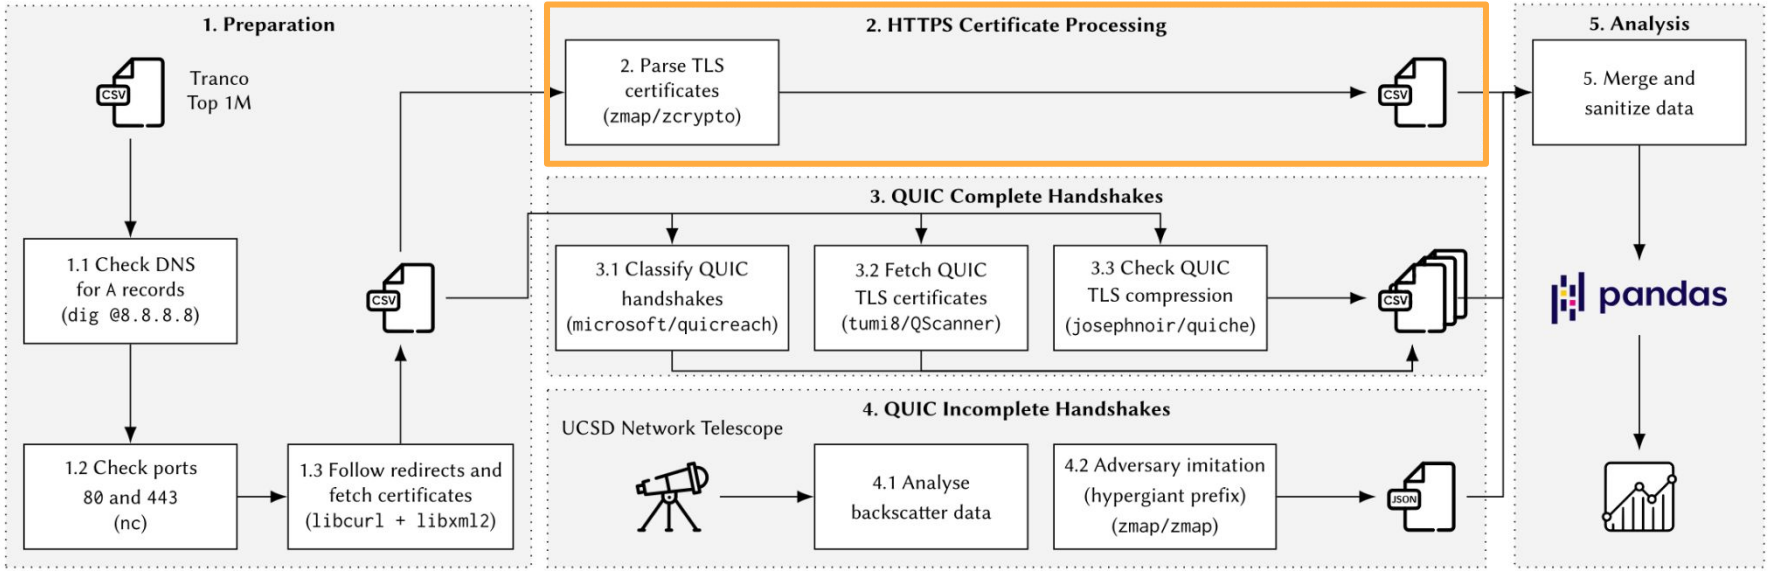

Methodology: Active scans with open-source tools.

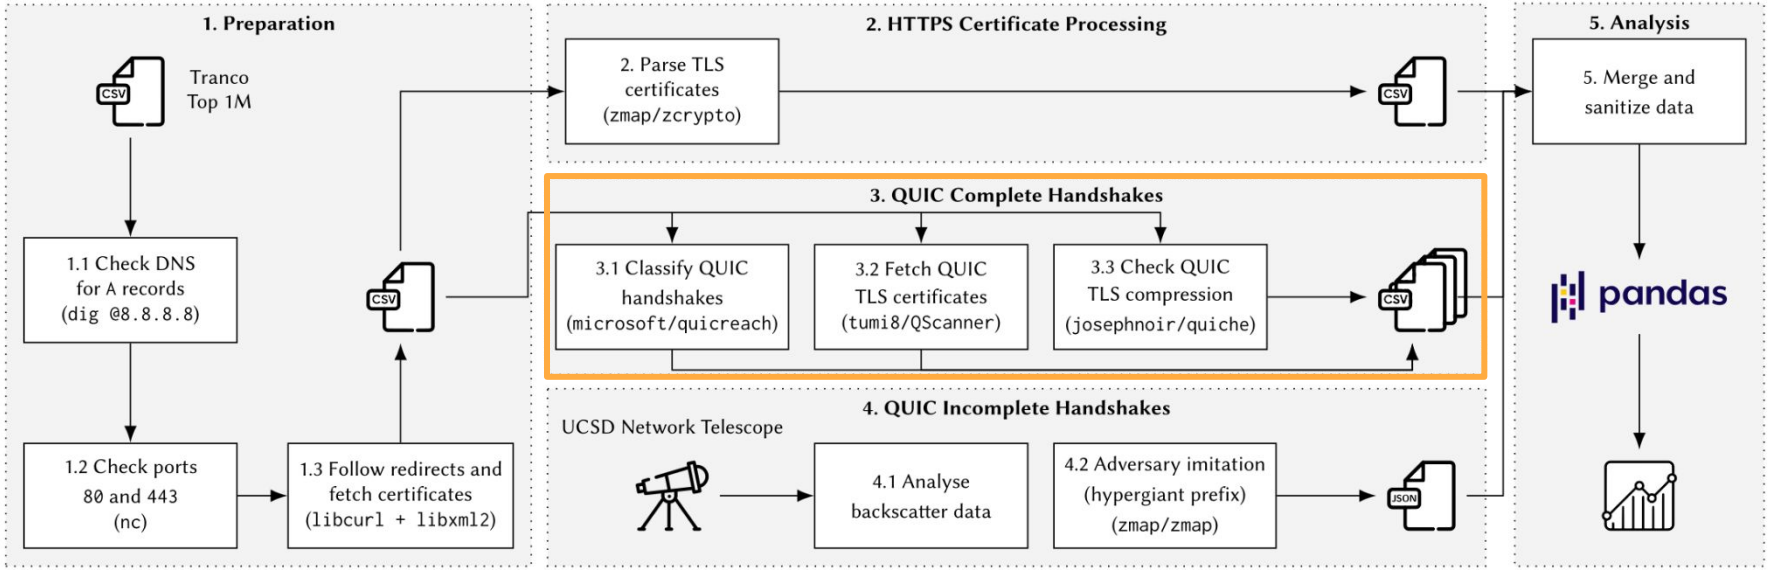

Complete handshakes enable the assessment of real-world performance.

Methodology: Active scans with open-source tools.

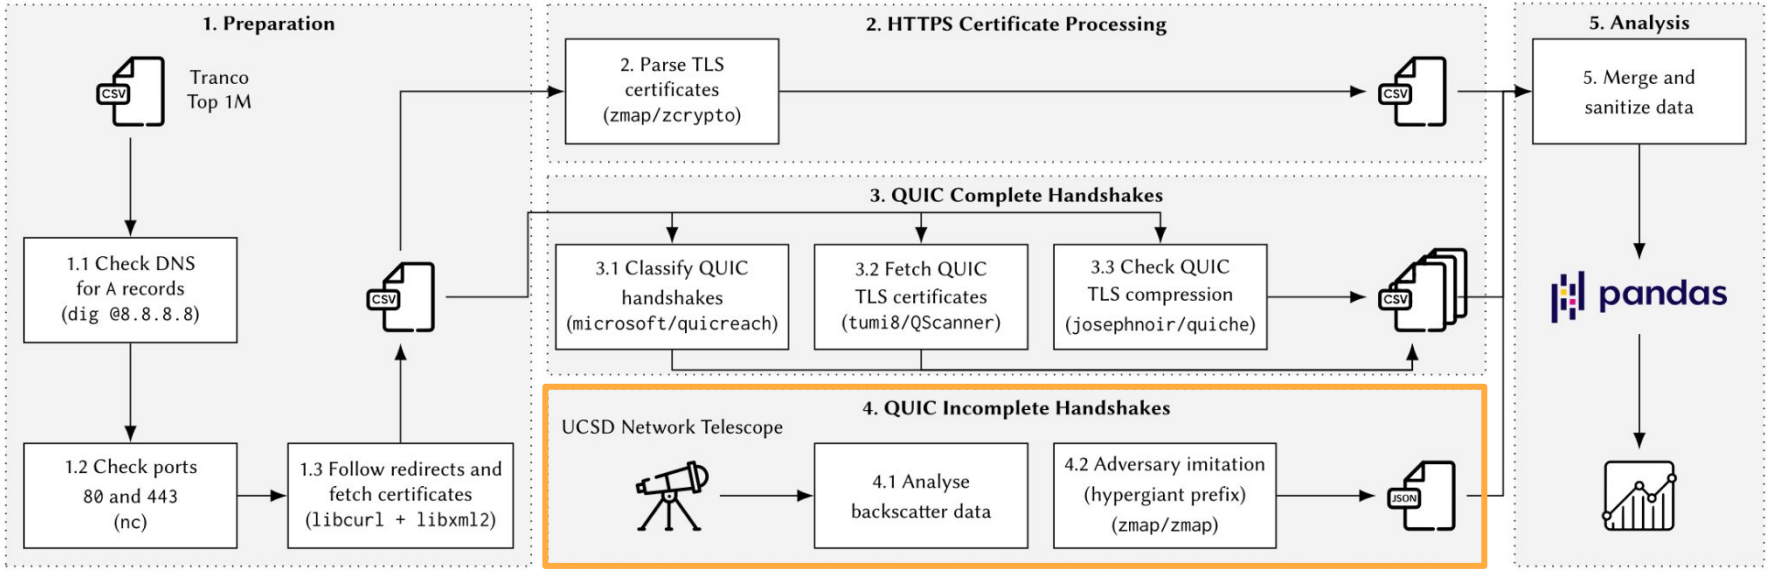

Incomplete handshakes unveil total susceptibility to reflective DDoS attacks.

Classifying QUIC complete handshakes.

(1) **1-RTT (optimal):** Handshakes that complete within 1-RTT and comply with the anti-amplification limit.
(2) **RETRY (less efficient):** Handshakes that require multiple RTTs because the `Retry` option is used [23, §8.1.].
(3) **Multi-RTT (unnecessary):** Handshakes that do not use `Retry` but require multiple RTTs because of large certificates.
(4) **Amplification (not RFC-compliant):** Handshakes that complete within 1-RTT but exceed the anti-amplification limit.

RFC-compliant 1-RTT handshakes are rare!

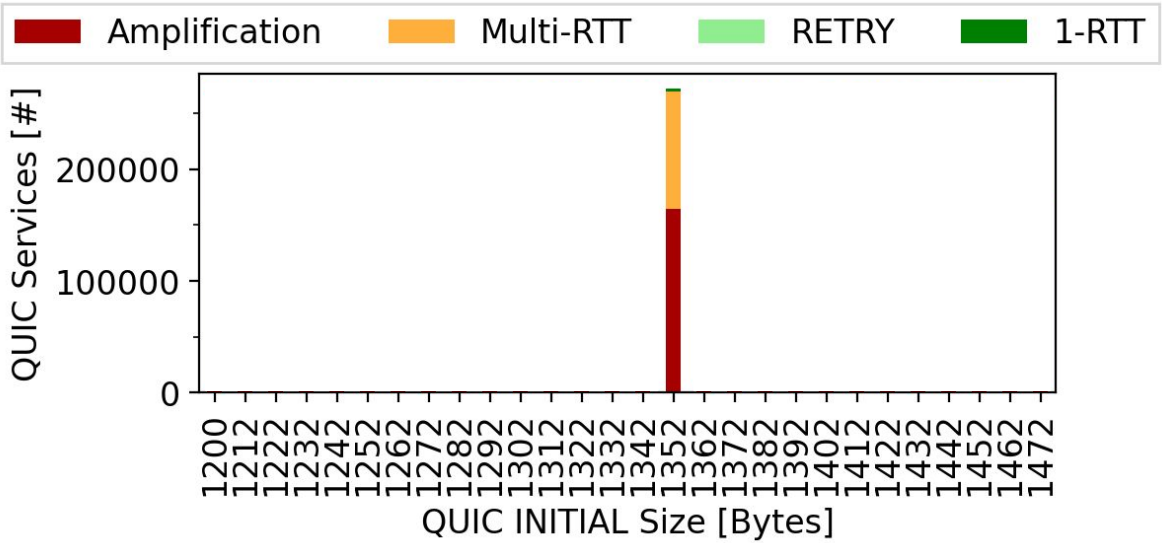

RFC-compliant 1-RTT handshakes are rare!

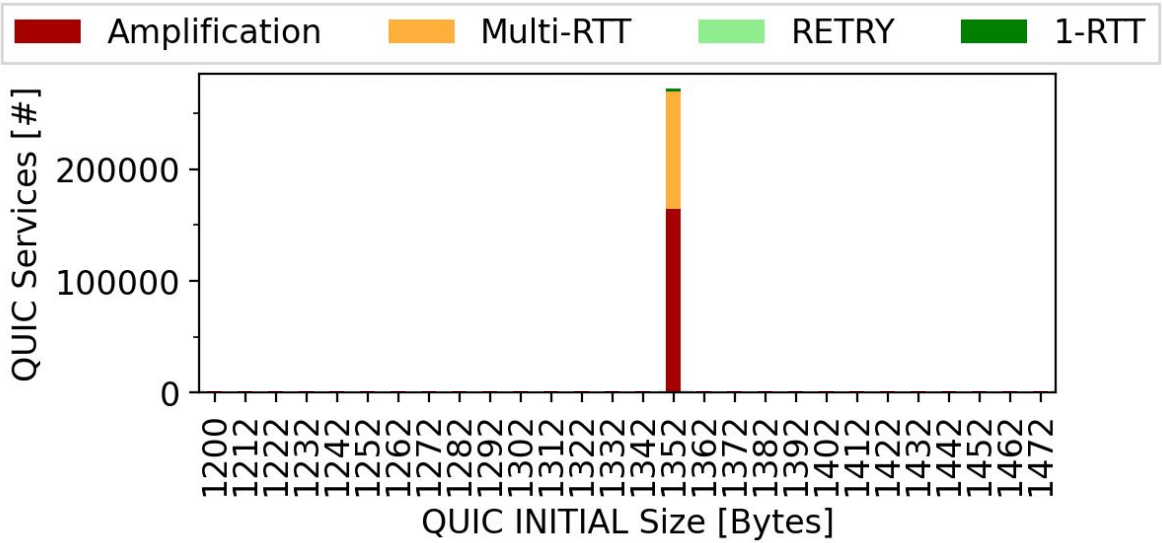

Browser Defaults

Smaller client INITIALs lead to multiple RTTs.

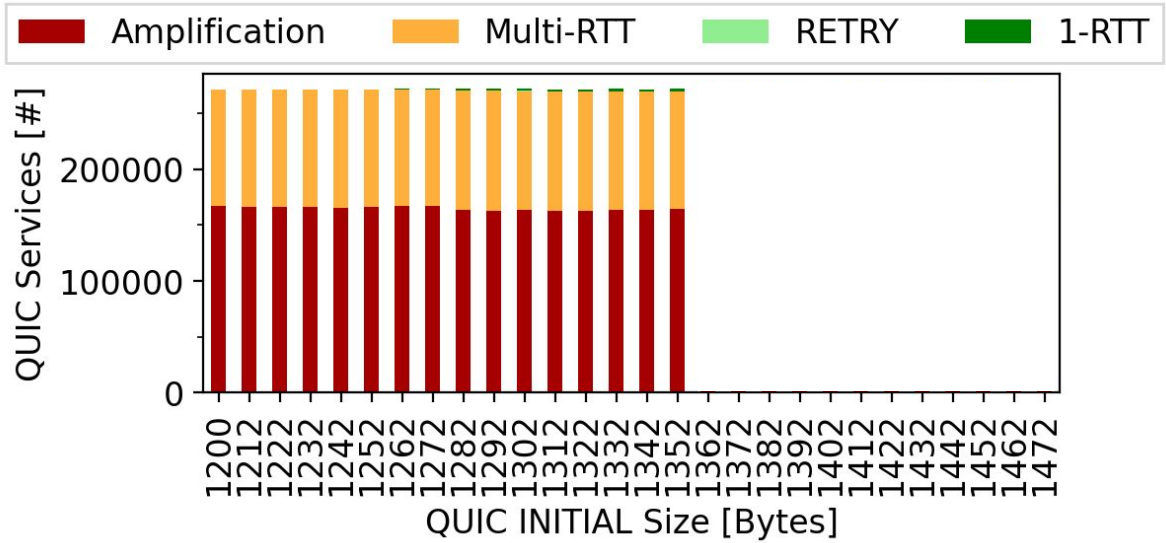

Browser Defaults

Very large client INITIALs reduce reachability.

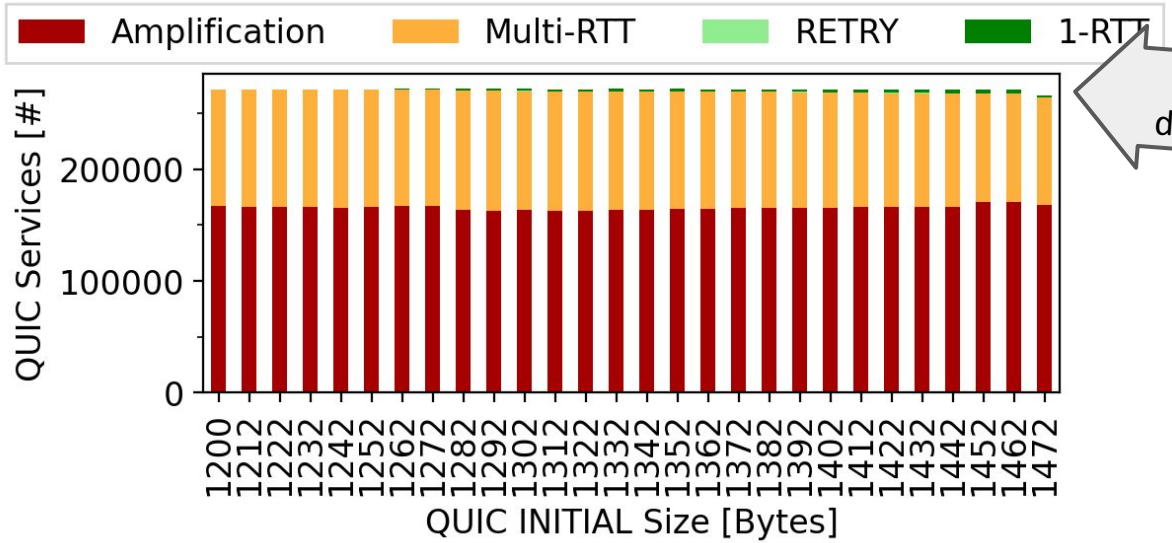

25% of the top 1k domains unreachable!

Browser Defaults

Agenda

Hypergiants willingly ignore the anti-amplification.

This enables clients to estimate a precise RTT.

TLS data still interferes with QUIC performance.

Improvements such as compression hard to integrate.

Incomplete QUIC handshakes amplify up to 45x.

Server retransmissions can lead to adverse effects.

How bad are the amplifying handshakes? Not bad.

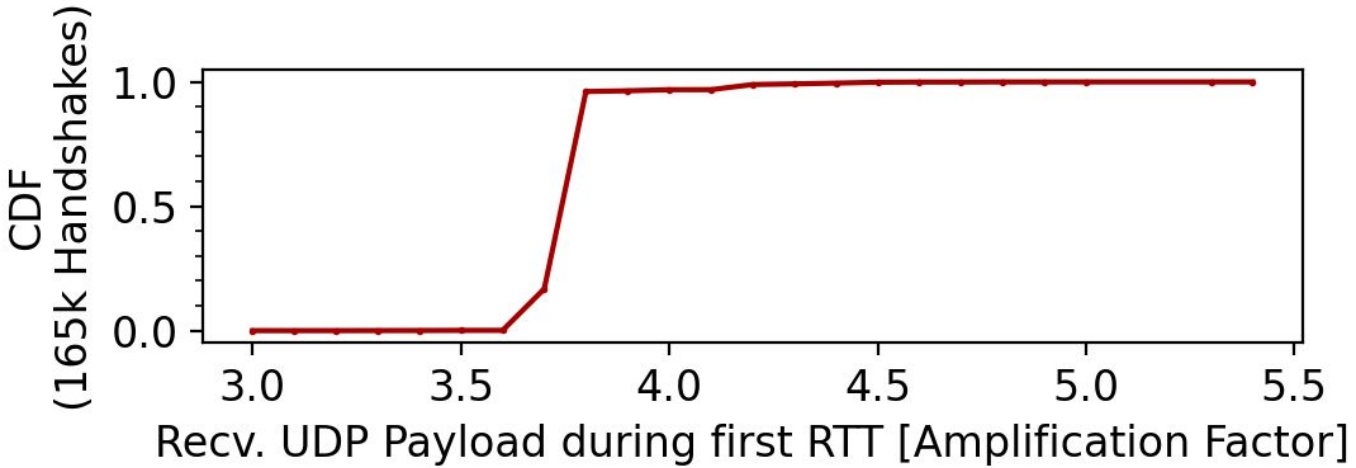

How bad are the amplifying handshakes? Not bad.

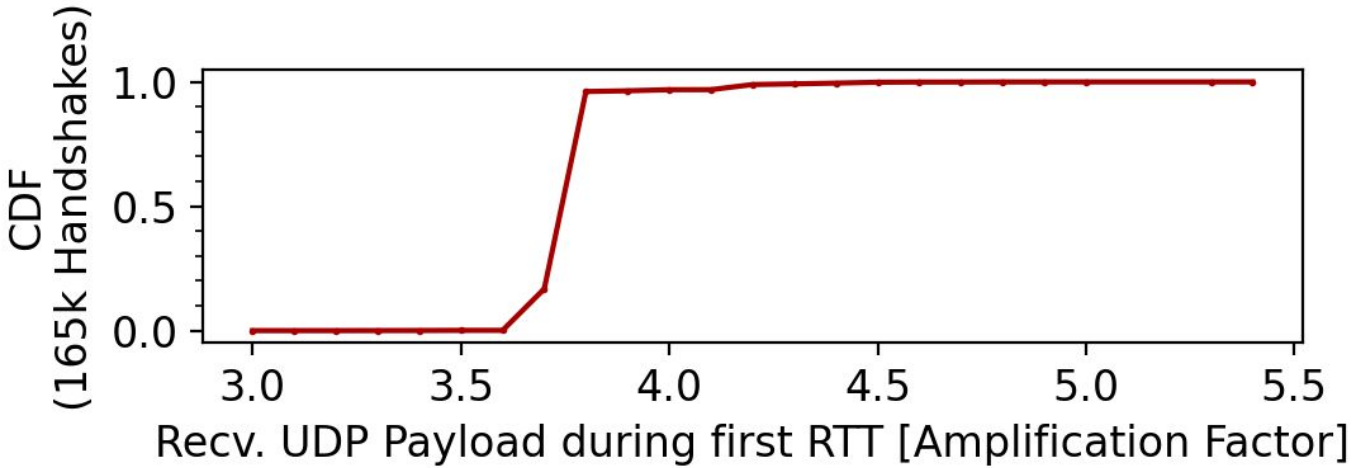

96% of the amplifying handshakes are completed with Cloudflare servers.

In several CDN deployments, the QUIC server can be separate from the process that has access to TLS material. This may add delay and disturb the client RTT estimation.

CDNs deal with this by splitting server Initials …

CDNs deal with this by splitting server Initials … and responding instantly only with the `ACK` …

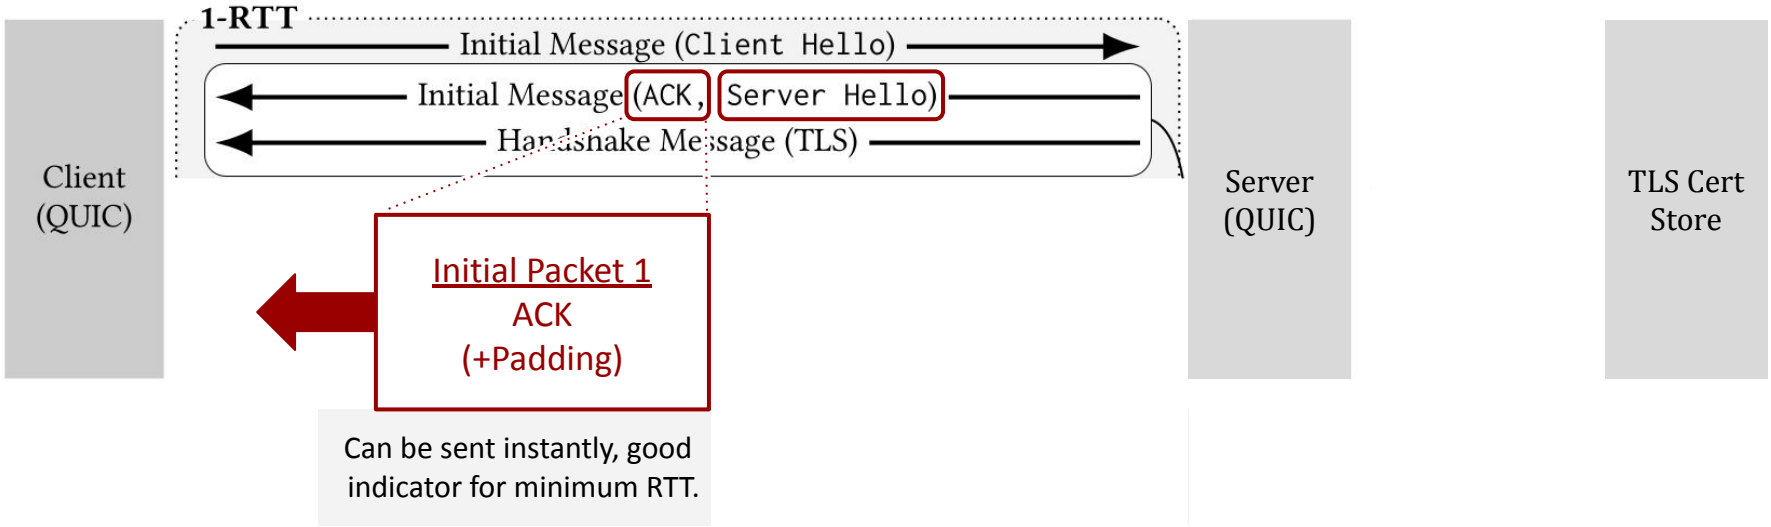

… and then retrieve and deliver the certificate.

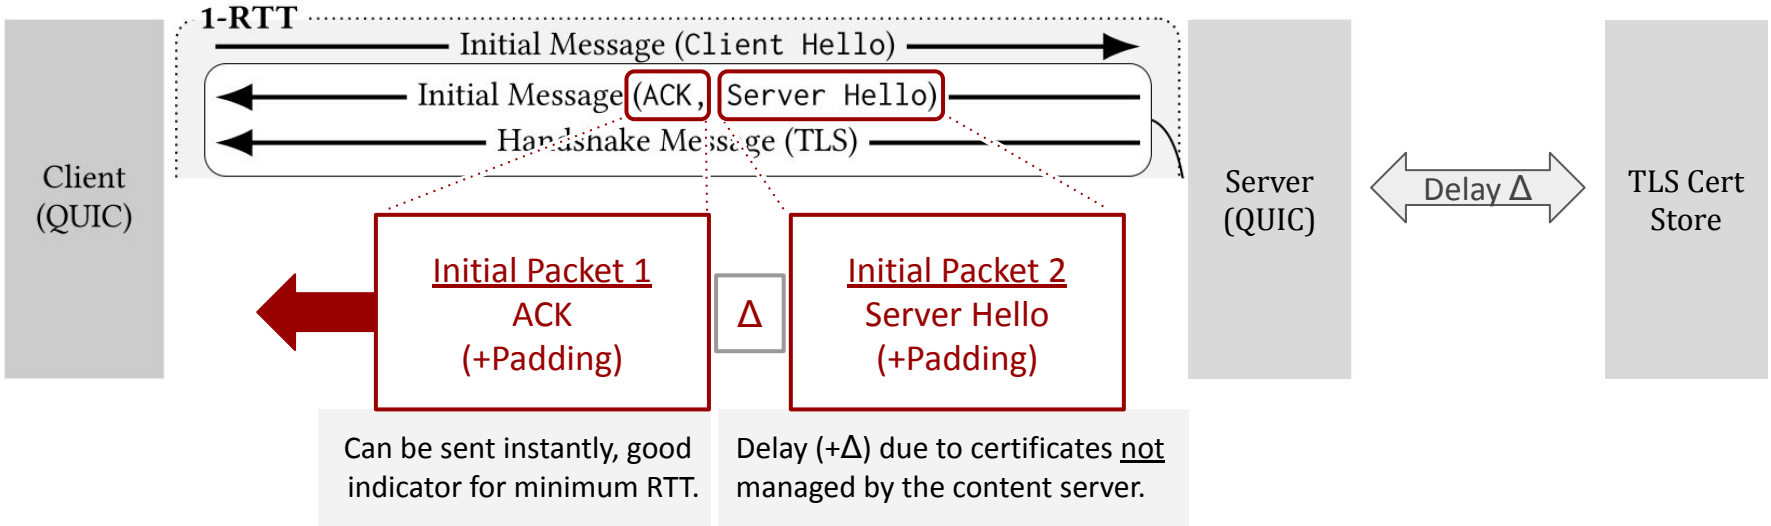

… and then retrieve and deliver the certificate.

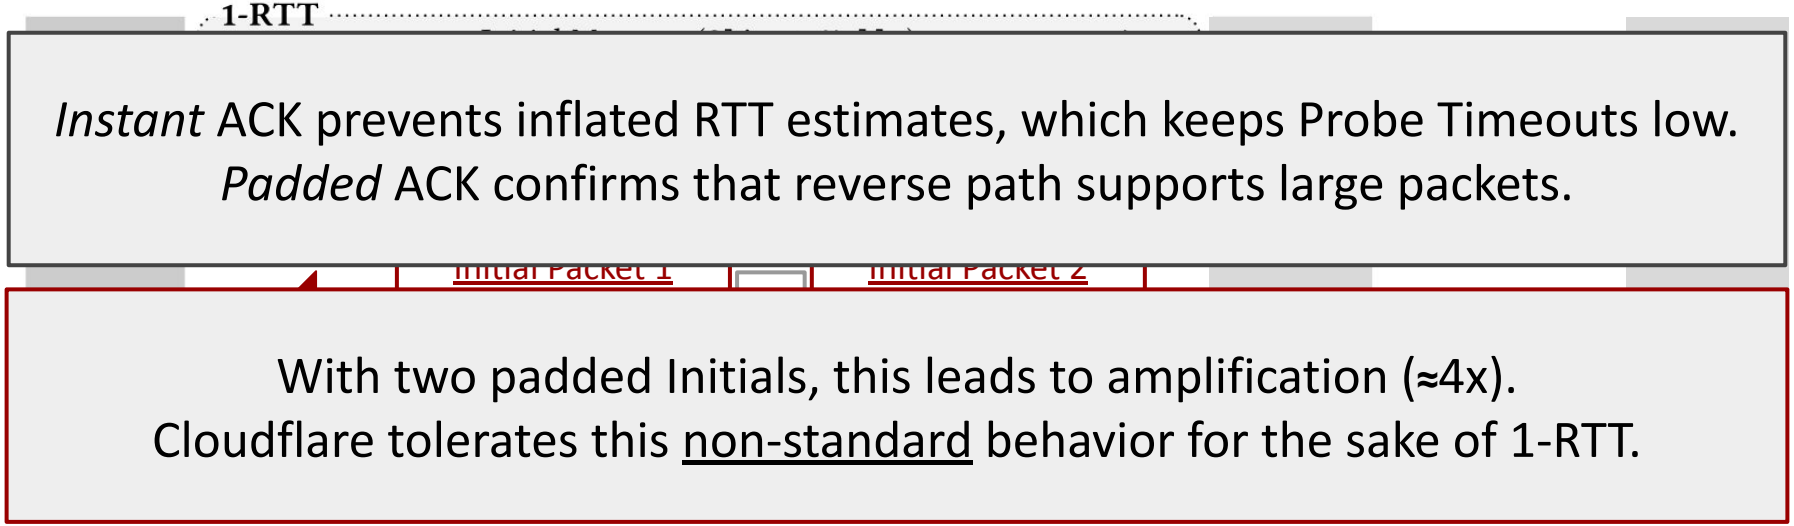

Instant ACK prevents inflated RTT estimates, which keeps Probe Timeouts low.
Padded ACK confirms that reverse path supports large packets.

With two padded Initials, this leads to amplification (≈4x).
Cloudflare tolerates this non-standard behavior for the sake of 1-RTT.

Agenda

Hypergiants willingly ignore the anti-amplification.

This enables clients to estimate a precise RTT.

TLS data still interferes with QUIC performance.

Improvements such as compression hard to integrate.

Incomplete QUIC handshakes amplify up to 45x.

Server retransmissions can lead to adverse effects.

What causes **multiple RTTs**?

DDoS prevention
(RETRY tokens)

< 200 domains.

Large TLS certificates
(that challenge the 3x limit)

The majority!

For multi-RTT handshakes, TLS bytes almost always (87%) exceed the limit but padding also has a significant impact.

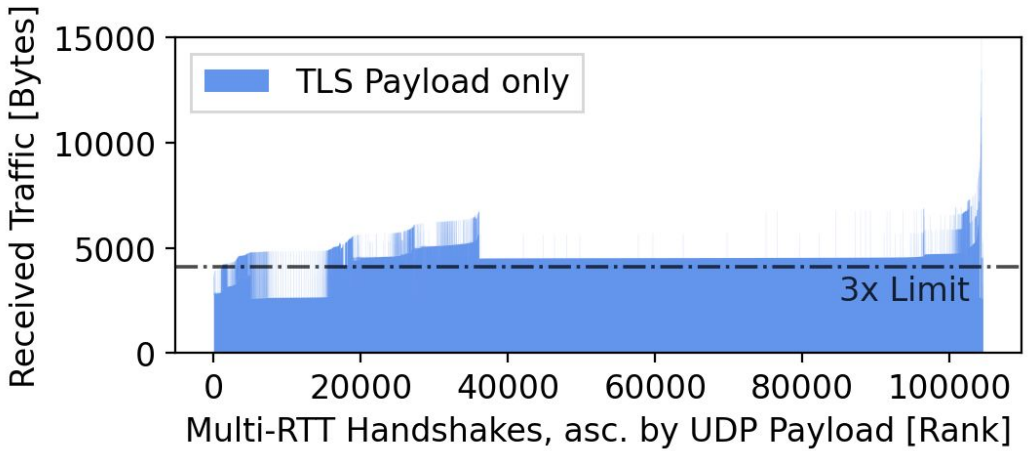

For multi-RTT handshakes, TLS bytes almost always (87%) exceed the limit but padding also has a significant impact.

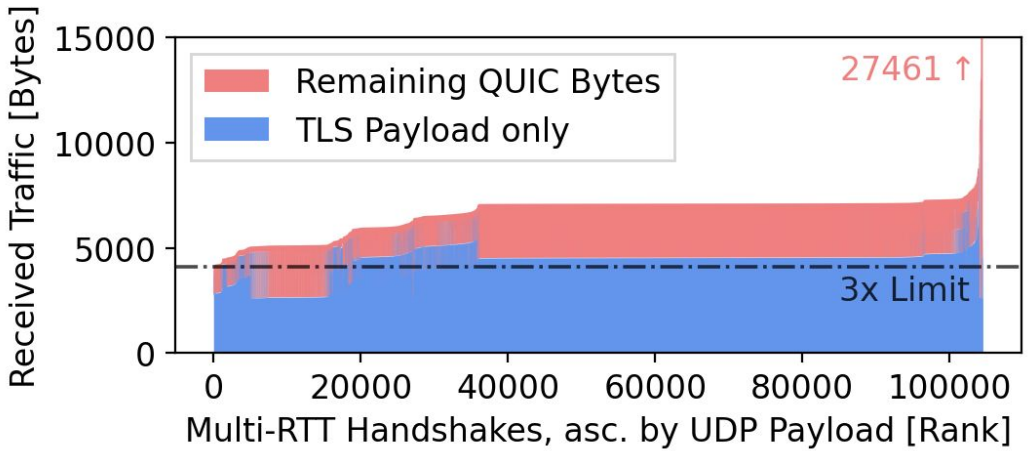

QUIC certificate chains. We look at non-leafs

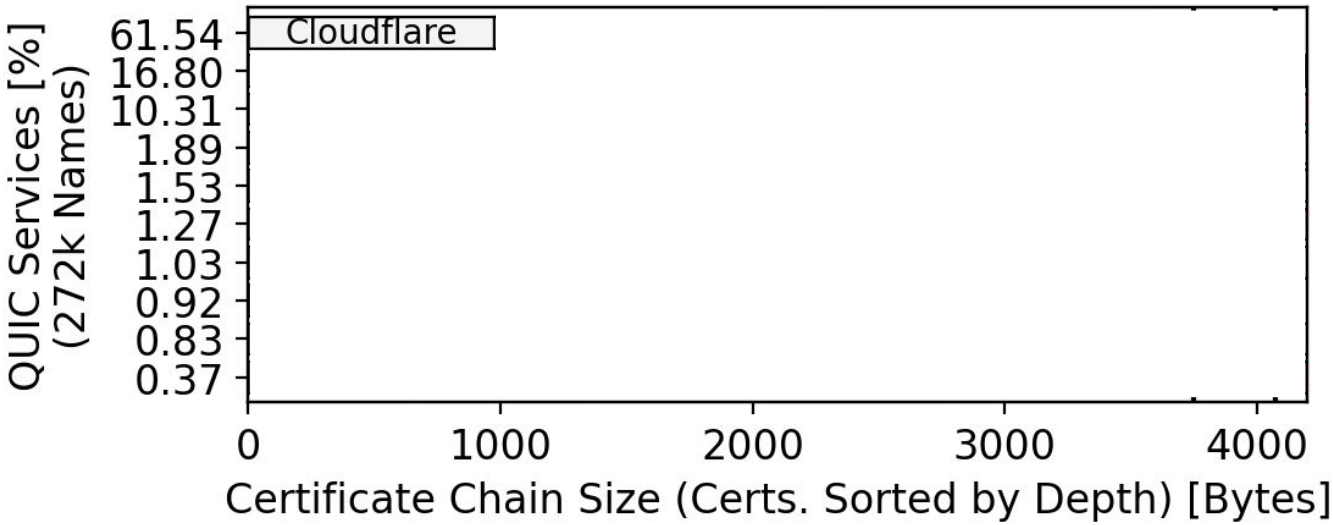

QUIC certificate chains. We look at non-leafs, median leaf sizes

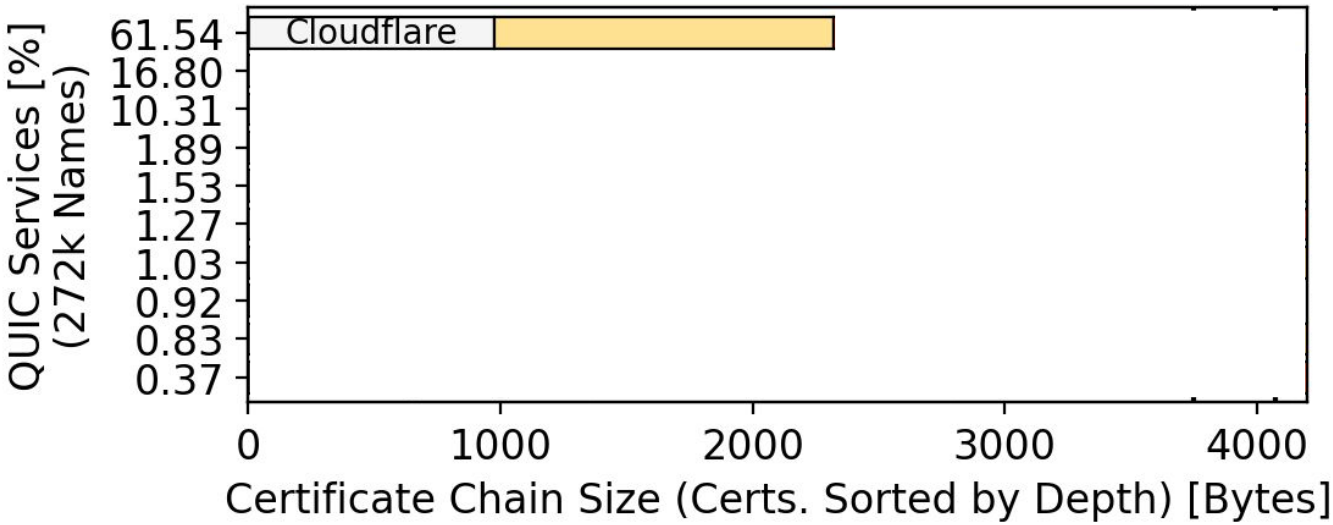

QUIC certificate chains. We look at non-leafs, median leaf sizes, extra bytes for maximum leaf

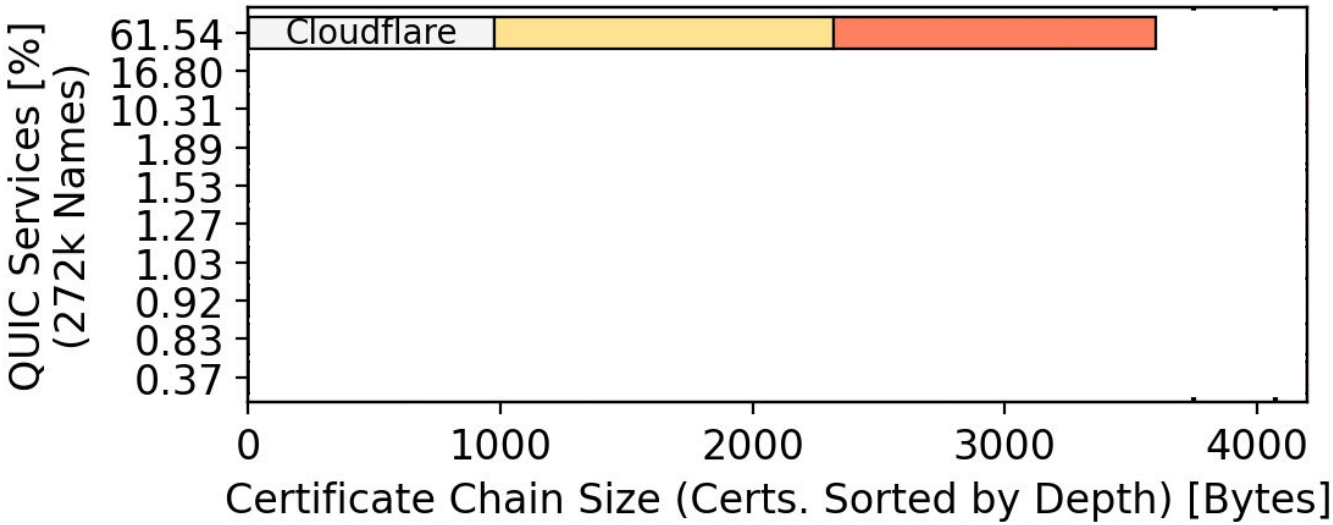

QUIC certificate chains. We look at non-leafs, median leaf sizes, extra bytes for maximum leaf, and common limits.

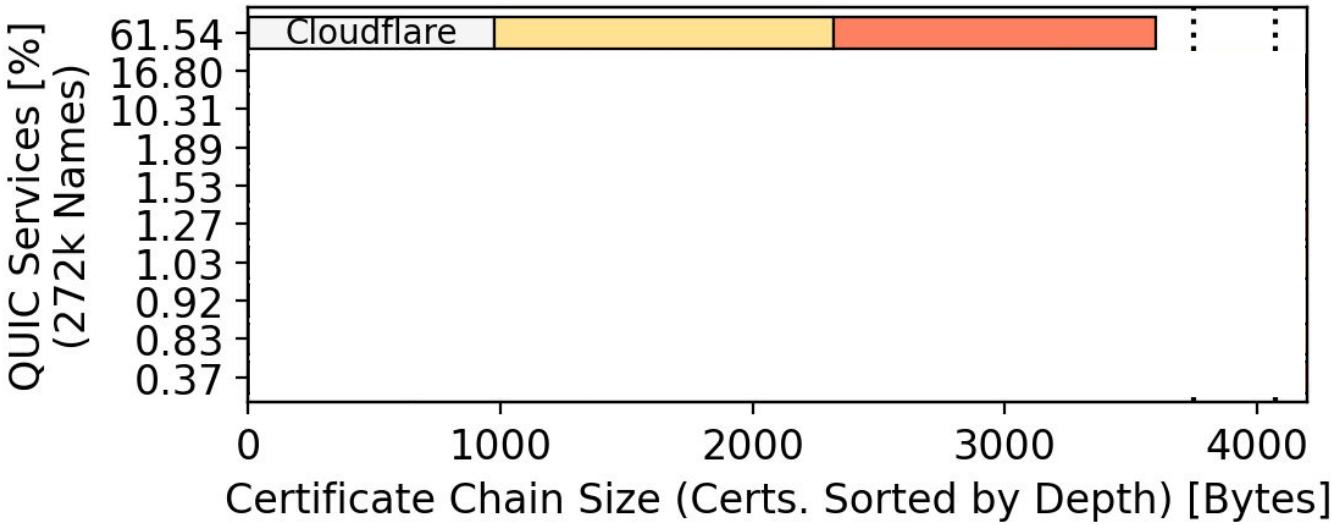

QUIC certificate chains. Median chains are likely to exceed common anti-amplification limits.

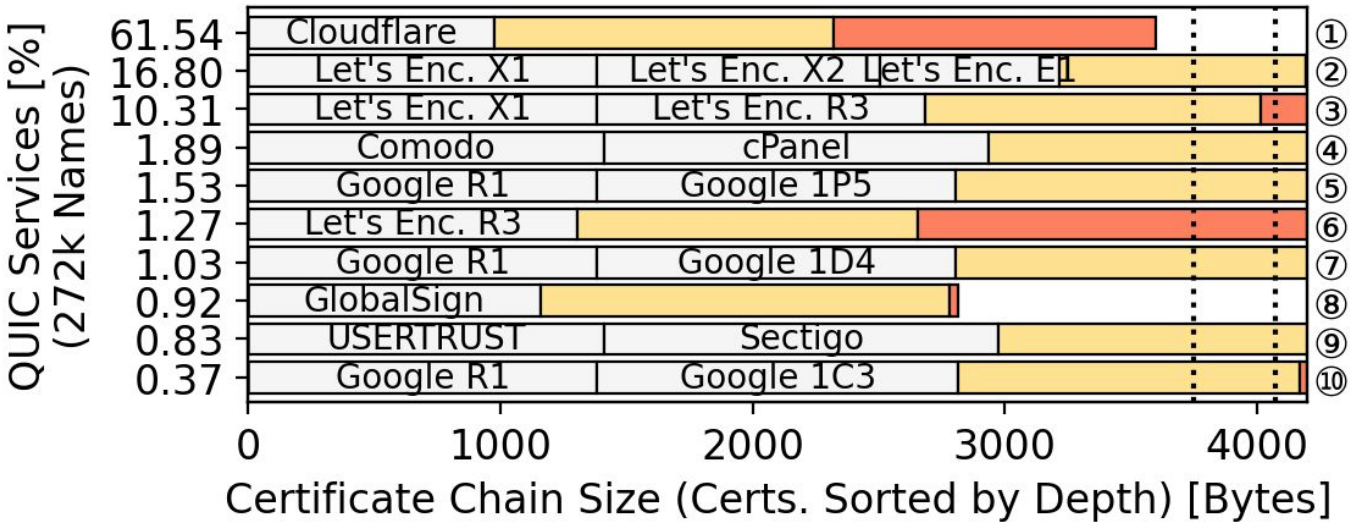

QUIC certificate chains. Median chains are likely to exceed common anti-amplification limits.

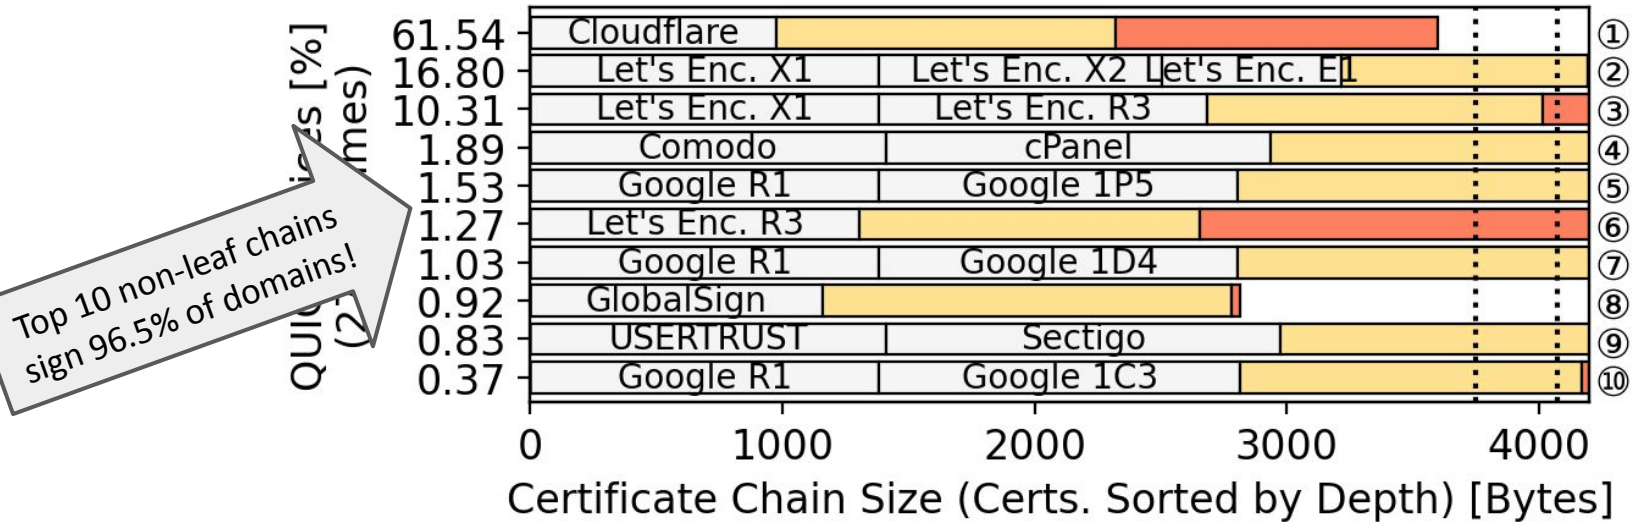

TCP/HTTPS-only services are less consolidated but still exceed the common limits.

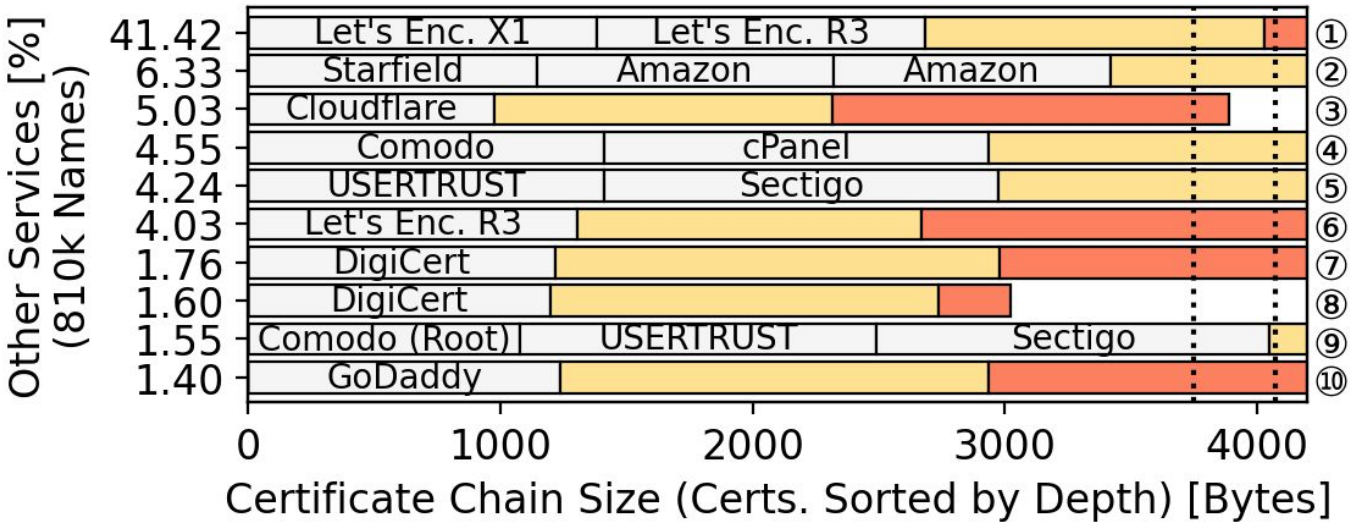

How to compensate for large certificates?

Updating non-leafs (RSA → ECDSA) would have beneficial cascading effects.

How to compensate for large certificates?

Updating non-leafs (RSA → ECDSA) would have beneficial cascading effects.

TLS certificate compression keeps 99% of data below anti-amplification limits.
Although we see high server support, clients and libraries struggle.

Agenda

Hypergiants willingly ignore the anti-amplification.

This enables clients to estimate a precise RTT.

TLS data still interferes with QUIC performance.

Improvements such as compression hard to integrate.

Incomplete QUIC handshakes amplify up to 45x.

Server retransmissions can lead to adverse effects.

Amplification factors increase drastically for incomplete handshakes because of server retransmissions.

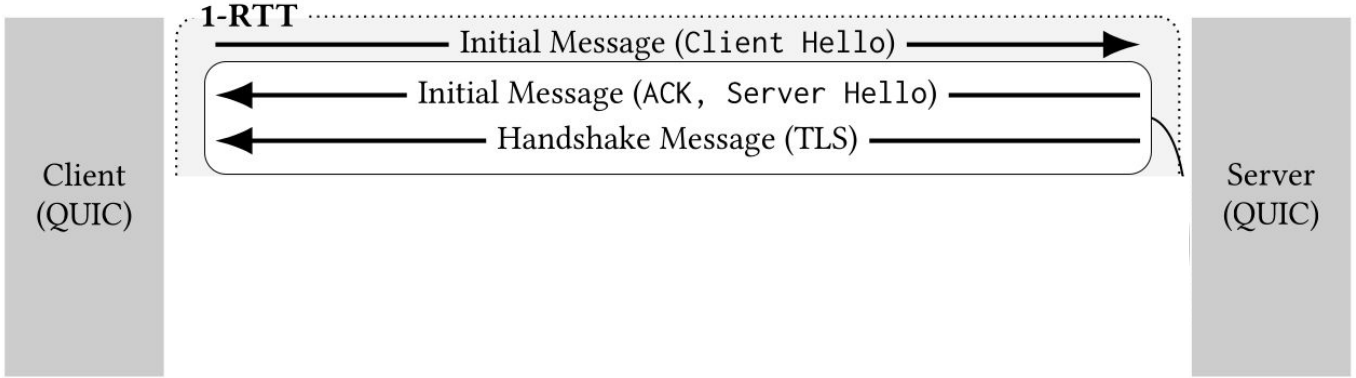

Amplification factors increase drastically for incomplete handshakes because of server retransmissions.

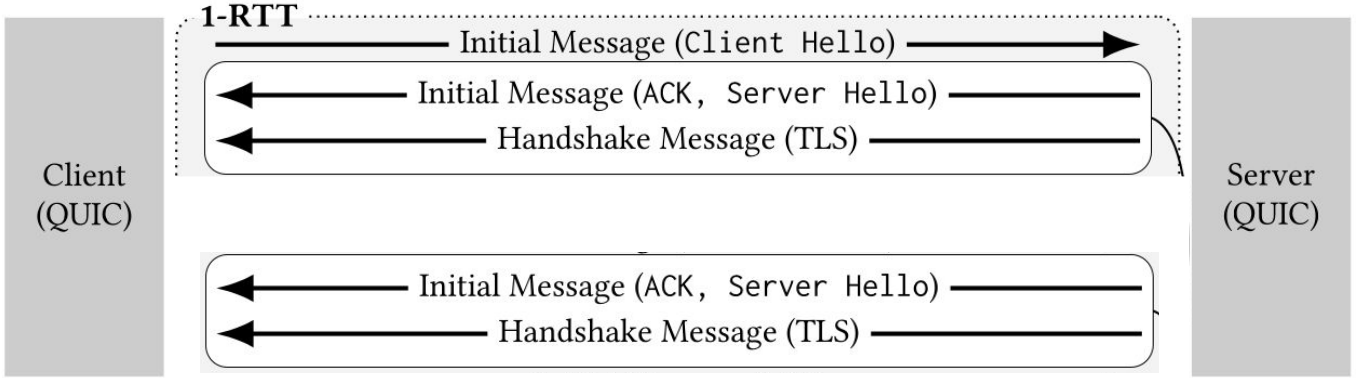

Amplification factors increase drastically for incomplete handshakes because of server retransmissions.

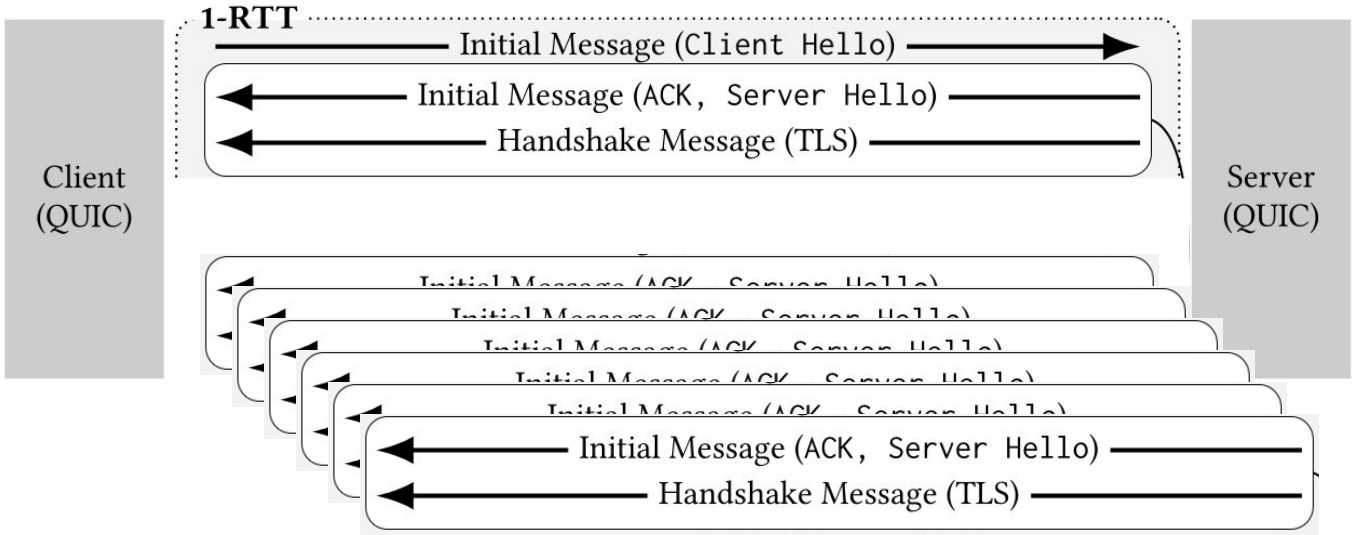

Amplification factors increase drastically for incomplete handshakes because of server retransmissions.

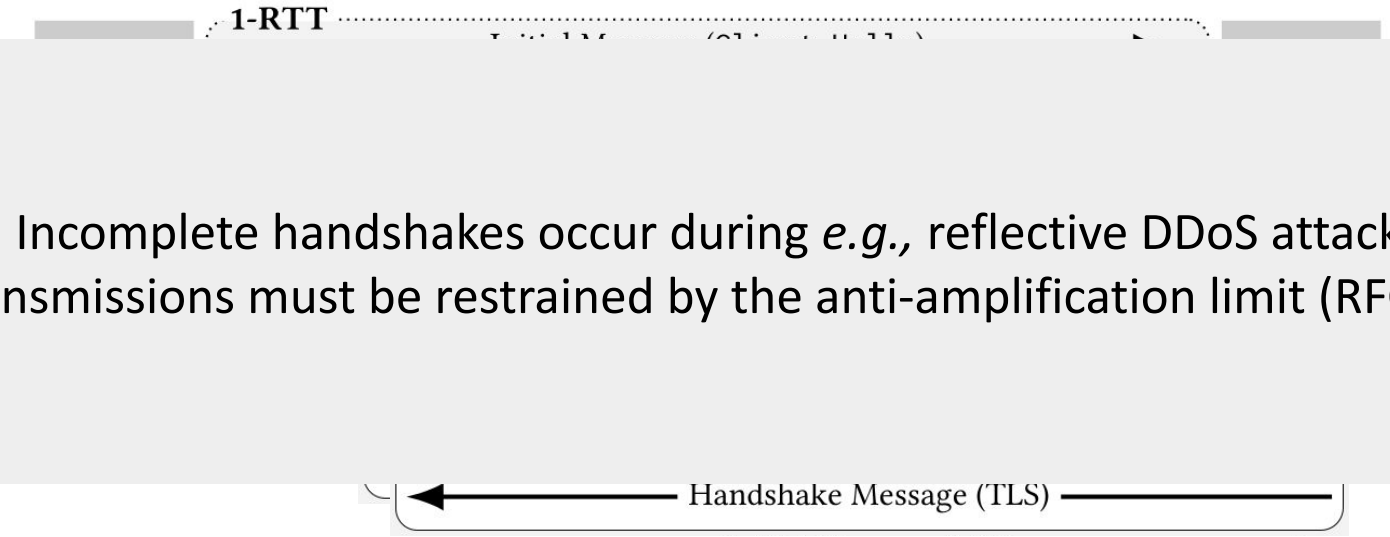

Incomplete handshakes occur during *e.g.,* reflective DDoS attacks.
Retransmissions must be restrained by the anti-amplification limit (RFC 9002).

Amplification for incomplete handshakes with Meta PoPs.

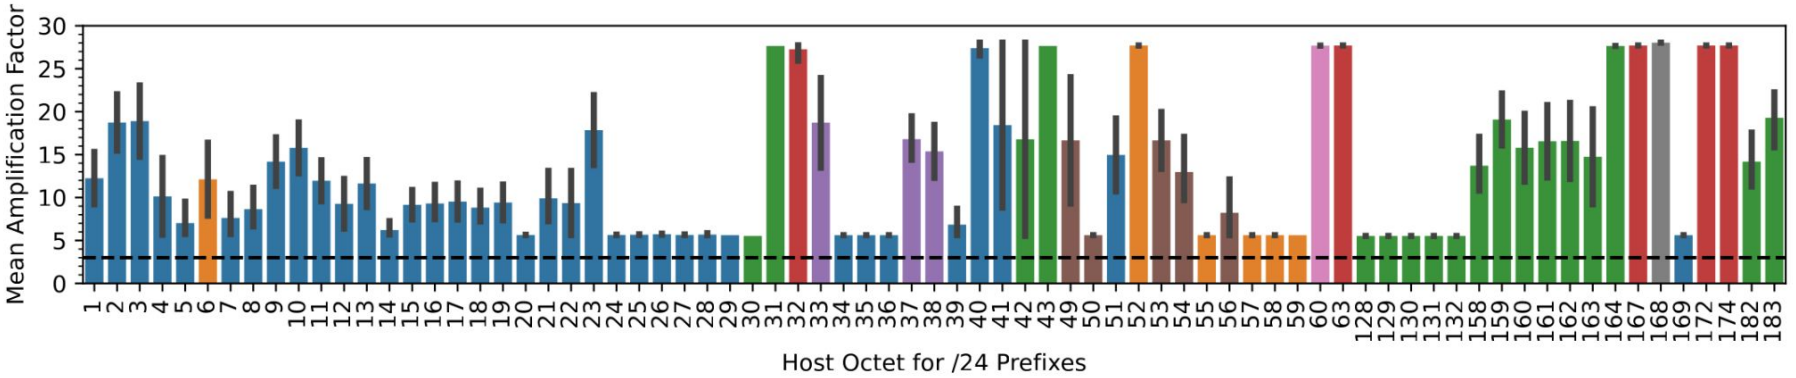

Amplification factors vary across different services.

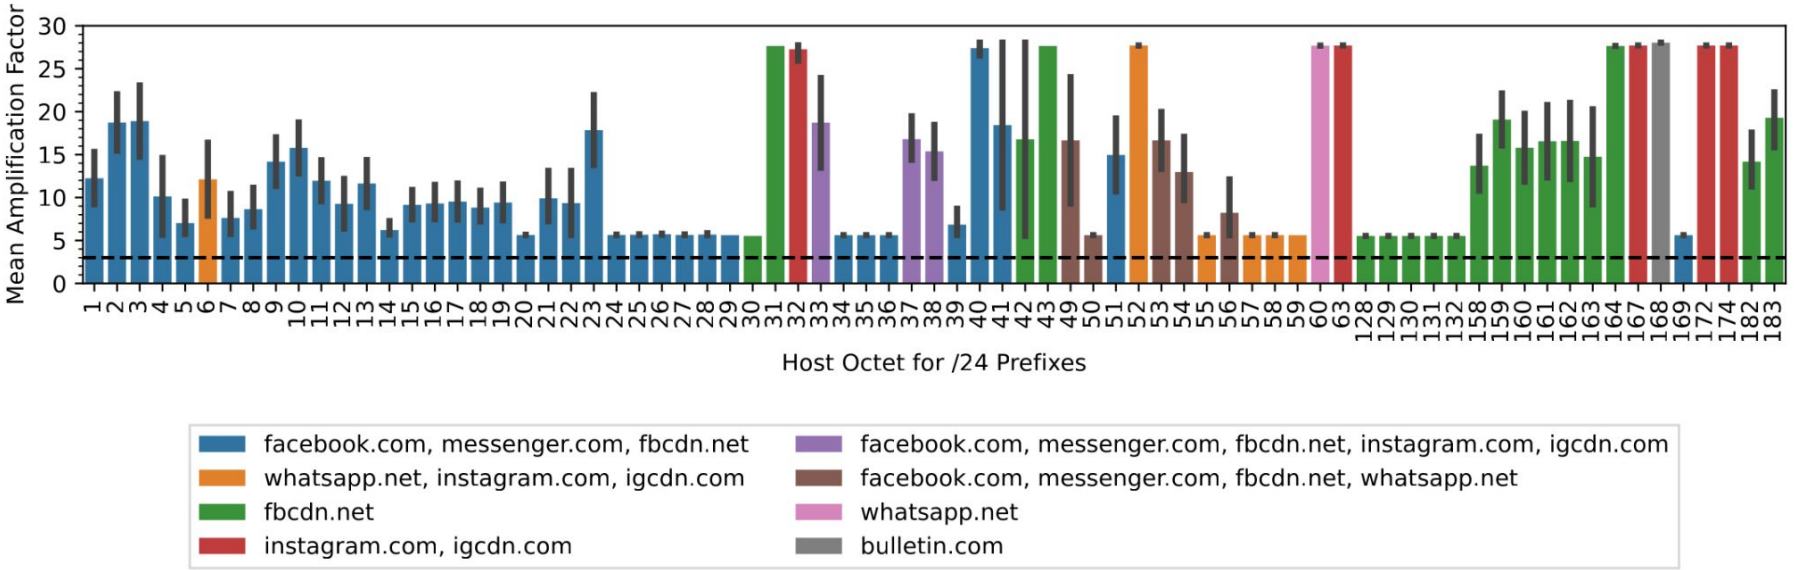

Follow-up scans show improvement, but still >3x.

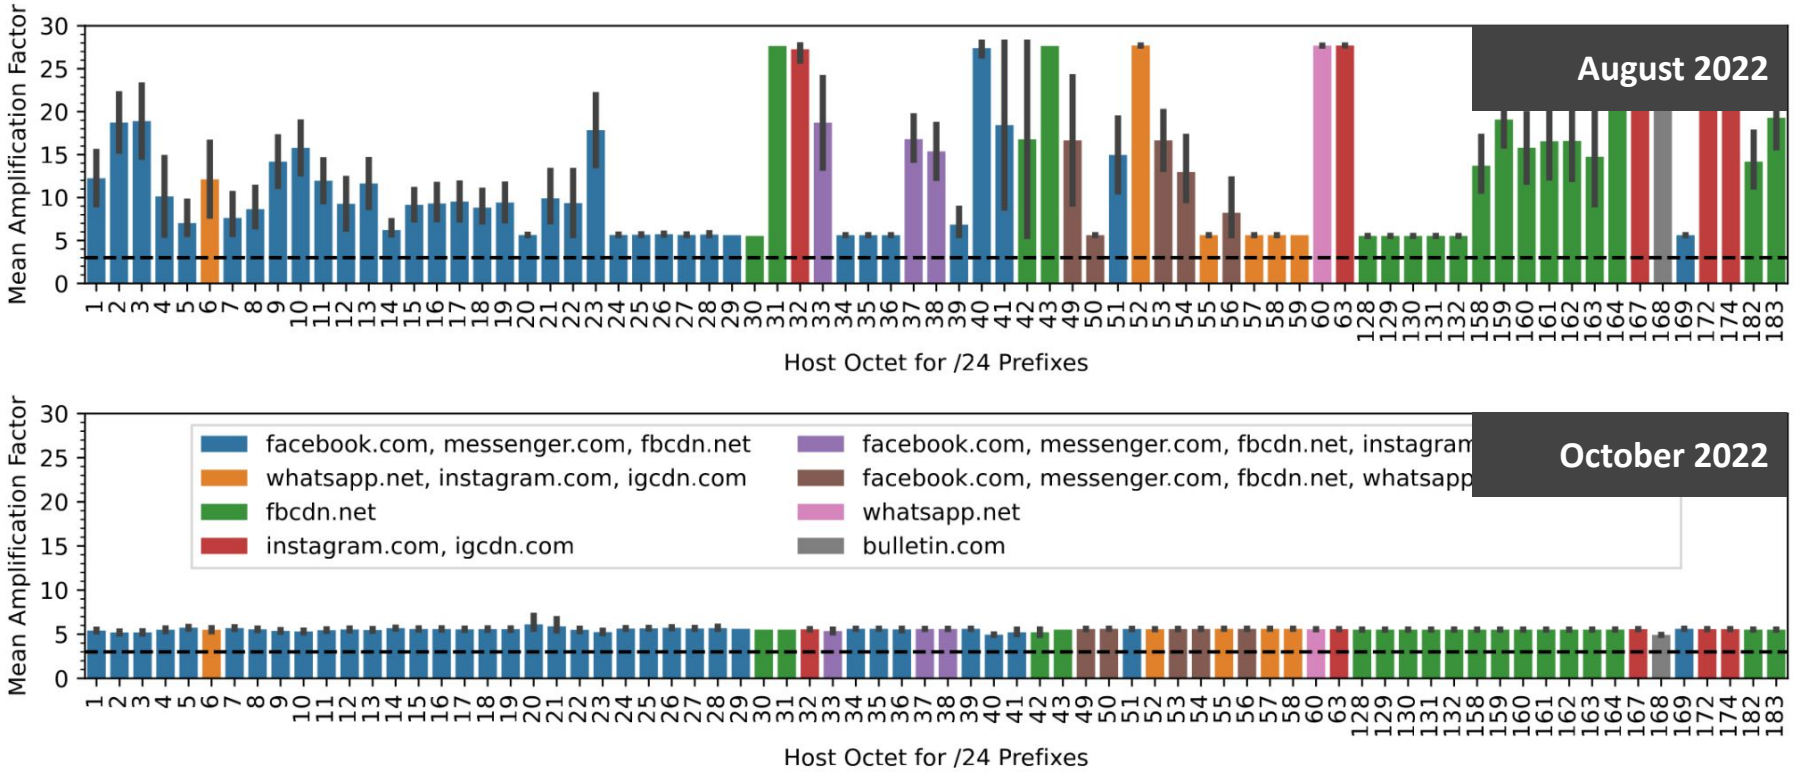

Follow-up scans show improvement, but still >3x.

Large TLS data leads to large retransmits. Respecting the anti-amplification limit decreases the chances of loss correction.

Open challenge: How to deal with packet loss during the QUIC connection setup in a secure but efficient way?

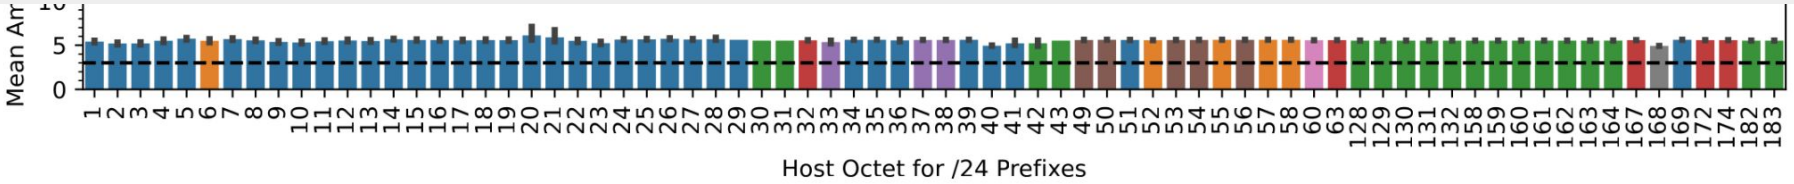

Conclusion

TLS Certificate Ecosystem

TLS configurations have now direct impact on transport layer performance.

ECDSA certificates lead to substantially smaller certificates chains.

Updates to non-leaf certificates would have beneficial cascading effects.

Conclusion

TLS Certificate Ecosystem

TLS configurations have now direct impact on transport layer performance.

ECDSA certificates lead to substantially smaller certificates chains.

Updates to non-leaf certificates would have beneficial cascading effects.

QUIC Deployments

Design goals (1-RTT, 3x anti-amplification limit) have been not met in the wild.

Trade-off during the handshake: Space efficiency (packet coalescence) vs. delay.

Padding and retransmissions significantly exacerbate the amplification factor.

QUIC Handshake Classification API (IETF 115 Hackathon)

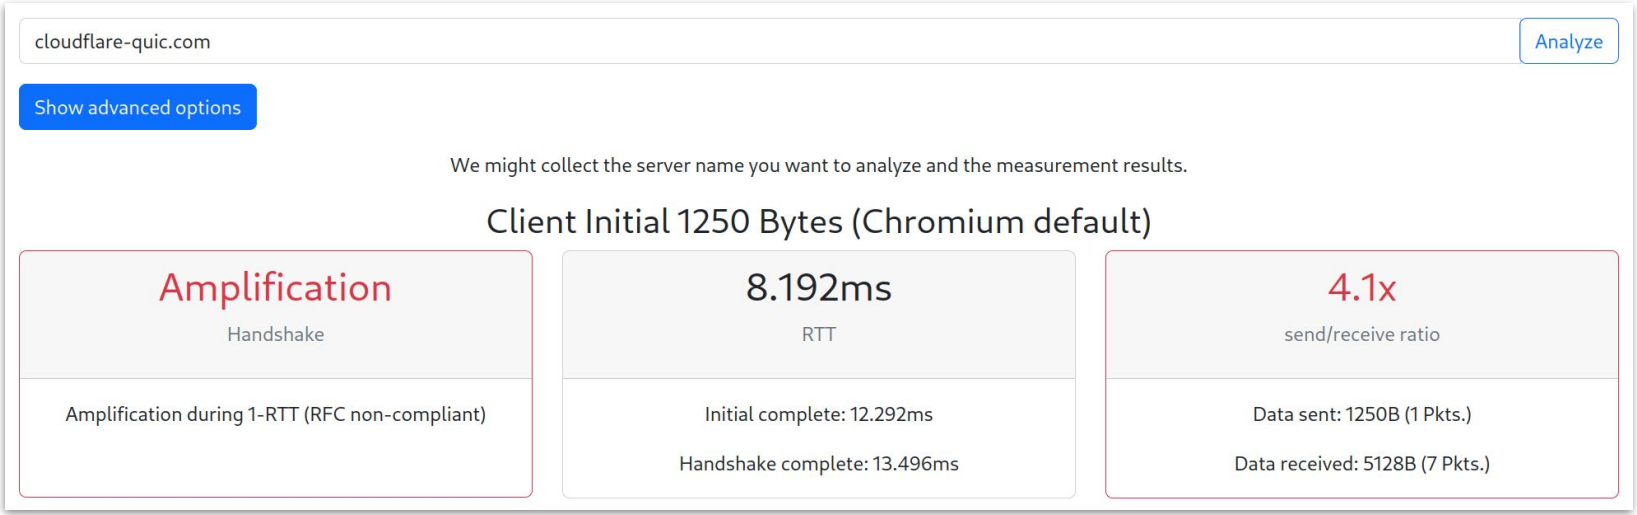

[understanding-quic.net]

Backup


Let's make
QUIC even
better!

TLS certificate fields and sizes

x509 v3 Certificate

tbsCertificate

version: 0x02 (v3)
serialNumber: 01:74:. . . :ca:7e
signatureAlg: sha256WithRSAEncryption
validity: 211127194412Z:221229194411Z

subject: CN=*.isc.org
subjectPublicKeyInfo:
algorithm: rsaEncryption
subjectPublicKey: 00:a5:. . . :56:95

issuer: C=BE, O=GlobalSign nv-sa, CN=GlobalSign Atlas R3 DV TLS CA H2 2021

extensions

AuthorityKeyIdentifier:
30:16:. . . :96:1f

SubjectKeyIdentifier: 04:14:. . . :b7:51
SubjectAltName: DNS:*.isc.org

signatureAlg: sha256WithRSAEncryption

signature: 30:45:. . . :e3:d6

(a) X.509 certificate

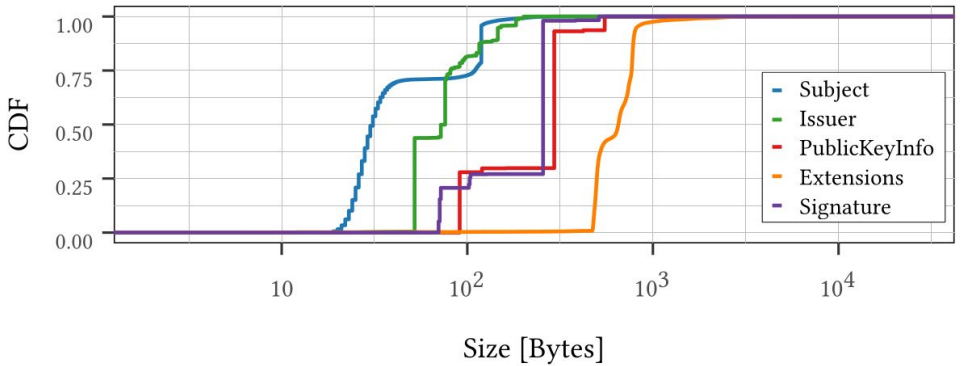

(b) Size distribution

Non-leafs contribute most bytes to large chains (QUIC).

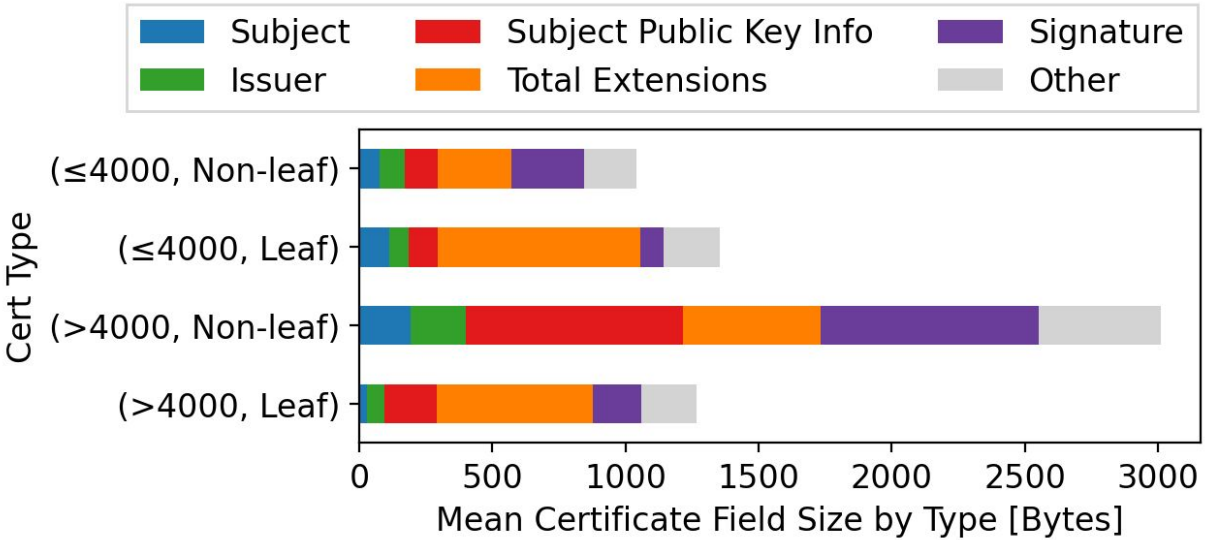

QUIC domains use smaller certificates.

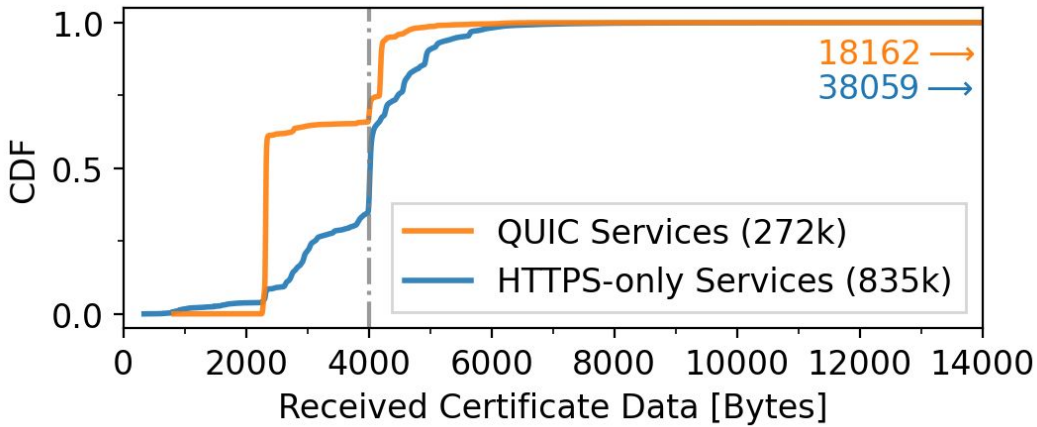

HTTPS-only domains depend heavily on RSA.

		RSA		ECDSA	
Service	Certificate	2048	4096	256	384
QUIC	Non-leaf	15.1%	22.4%	40.4%	22.1%
	Leaf	19.2%	1.4%	78.9%	0.0%
HTTPS-only	Non-leaf	63.3%	32.1%	2.7%	1.6%
	Leaf	81.4%	8.1%	7.8%	1.9%

Client Initial sizes and TLS compression of web browsers.

Browser	Version	Init. Size [Bytes]	Compression Algorithm[3]	Compression Rate[4]	Compression Services[5]
Firefox	`101.x`	1357	–	–	–
Chromium-based[1]	`105.x`	1250[2]	brotli	73%	96%
Safari (macOS)	`15.5`	no QUIC	`zlib`	74%	0.05%
			`zstd`	72%	0.05%

[1] Chrome `102.x`, Brave `V1.39`, Vivaldi `5.3.x`, Edge `102.x`, Opera `88.0.x`.
[2] Recently reduced from 1350 [13]. [3] Tested with TLS 1.3 in TCP.
[4] Mean rate observed by our Quiche client. [5] Out of 272k QUIC services.

Anti-amplification limit in the IETF QUIC Internet drafts.

IETF Spec	Date	Proposed Limit
Draft 09	01/2018	"A server MAY send a CONNECTION_CLOSE frame with error code PROTOCOL_VIOLATION in response to an Initial packet smaller than 1200 octets."
Draft 10 – 12	03/2018 – 05/2018	"Servers MUST NOT send more than three **Handshake** packets without receiving a packet from a verified source address."
Draft 13 – 14	06/2018 – 08/2018	"Servers MUST NOT send more than three **datagrams** including Initial and Handshake packets without receiving a packet from a verified source address."
Draft 15 – 32	10/2018 – 10/2020	"Servers MUST NOT send more than three times as many **bytes** as the number of bytes received prior to verifying the client's address."
Draft 33 – 34, RFC 9000	12/2020 – 01/2021, 05/2021	"[…] an endpoint MUST limit the amount of **data** it sends to the unvalidated address to three times the amount of data received from that address."

QUIC and HTTPS deployment rates are stable across rank groups.

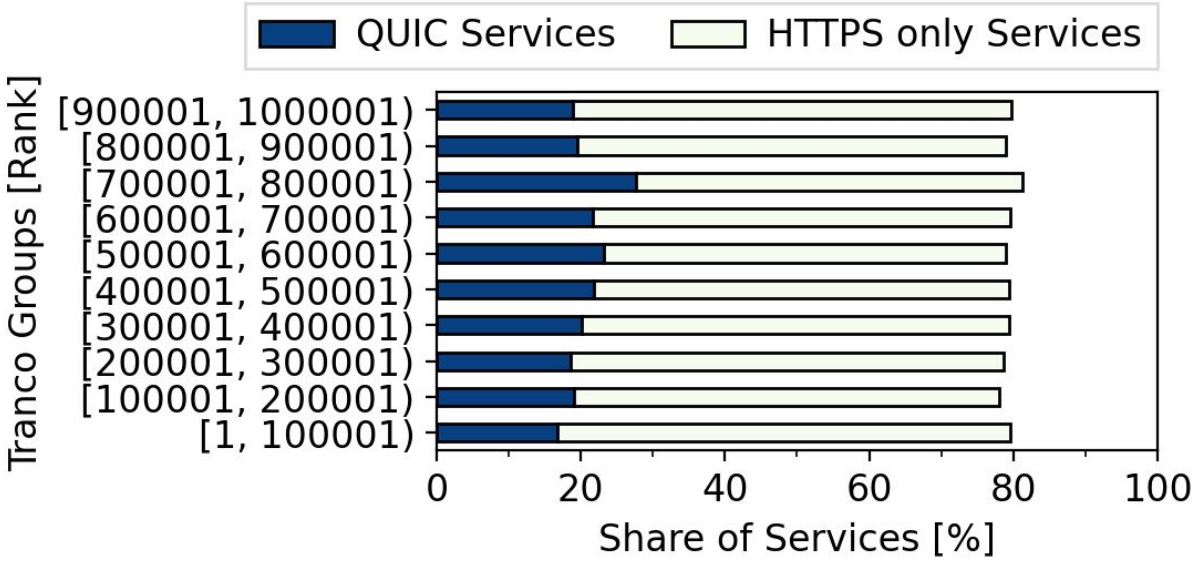

Handshake types are mostly stable across rank groups.

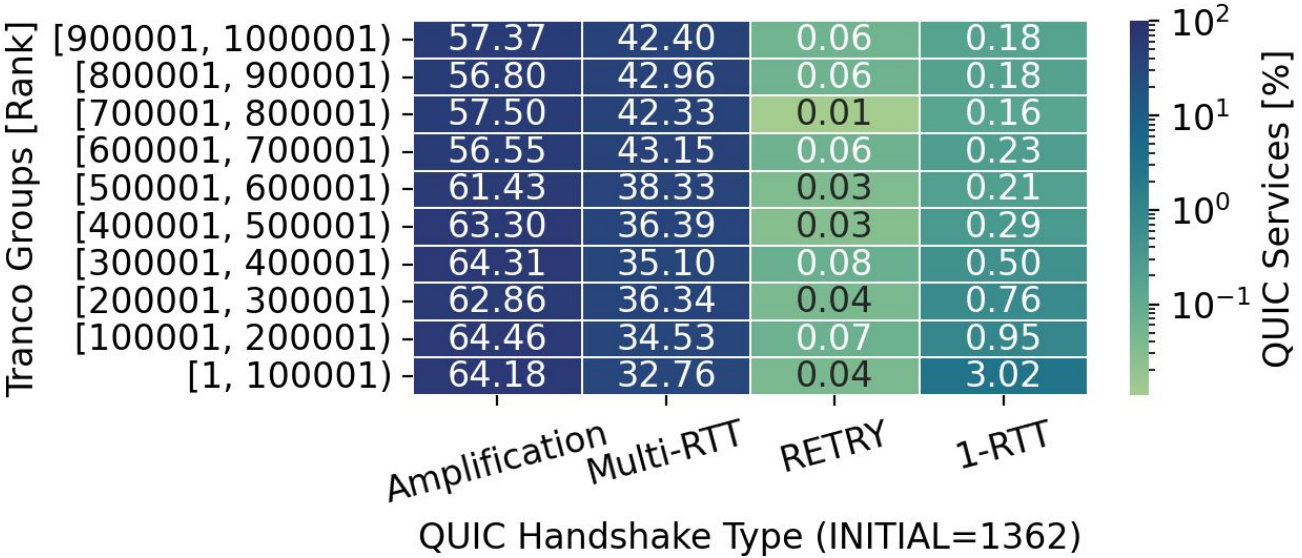

Cruise-liner certificates are rare for QUIC services.

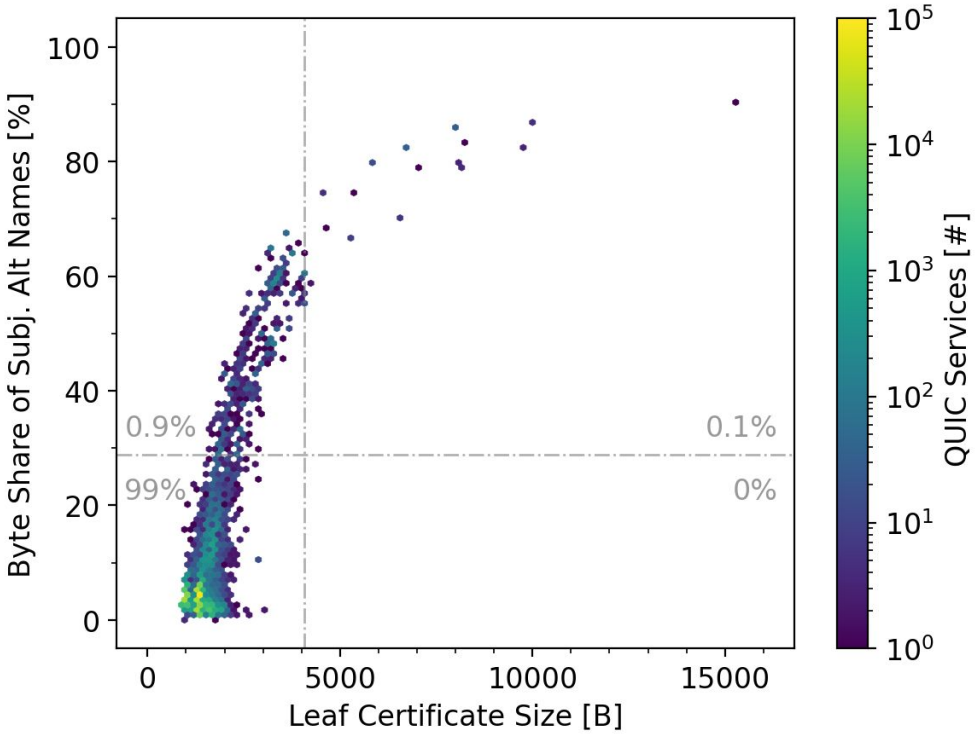

Telescopes passively observe incomplete handshakes. Especially Meta fails to comply with the limit.