CLOUDFLARE

Respect the ORIGIN! A Best-case Evaluation of Connection Coalescing

Sudheesh Singanamalla
Muhammad Talha Paracha
Suleman Ahmad
Jonathan Hoyland
Luke Valenta
Yevgen Safronov
Peter Wu
Andrew Galloni
Kurtis Heimerl
Nick Sullivan
Christopher A. Wood
Marwan Fayed

CLOUDFLARE

What is connection coalescing?

Same IP addresses but results in multiple **possibly blocking** DNS queries.

example.com

images.example.com

content.example.com

cdn.external.com

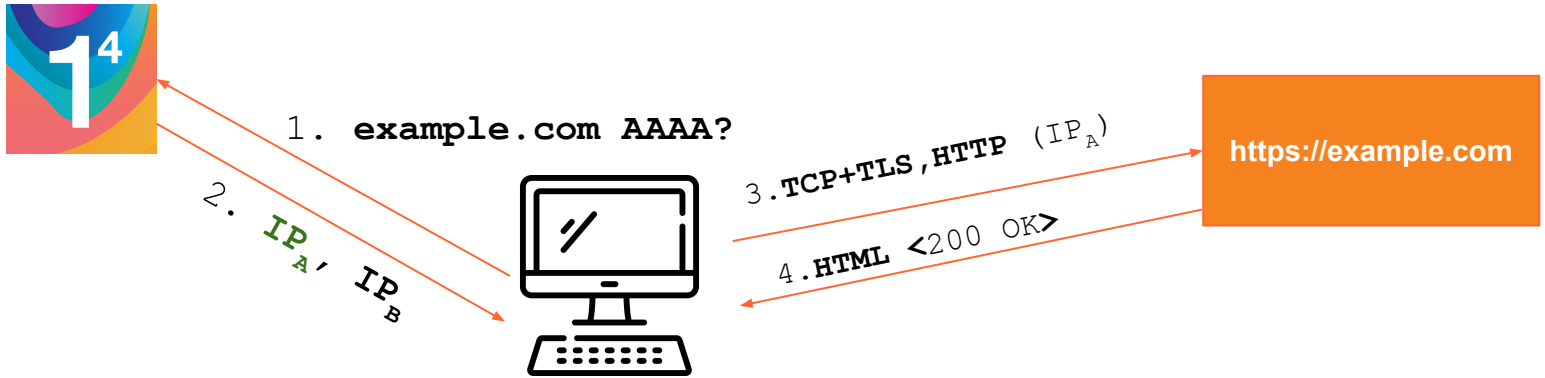

Next: What happens for subresources?

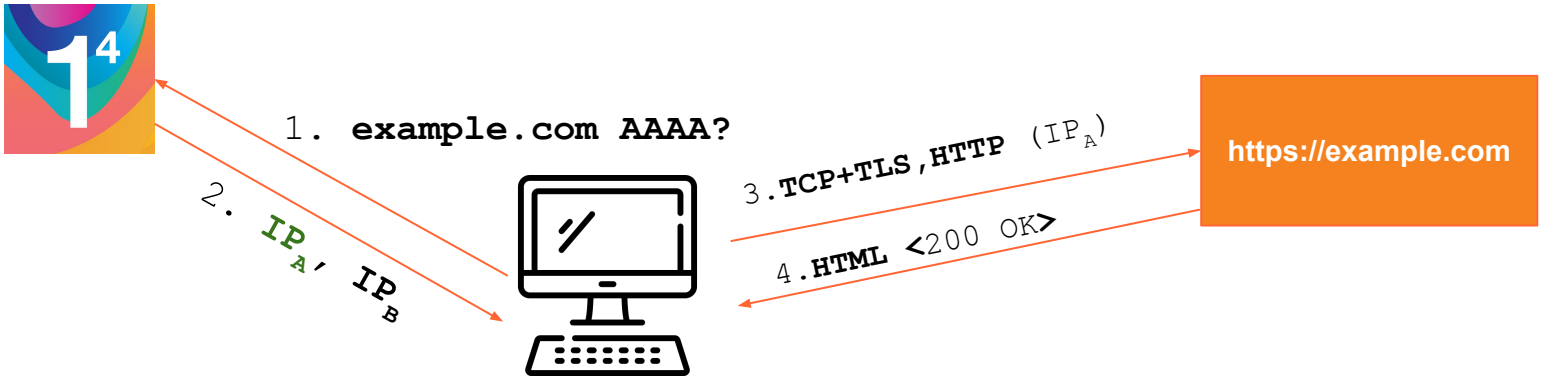

Chrome's Approach: IP addresses for different hostnames must match

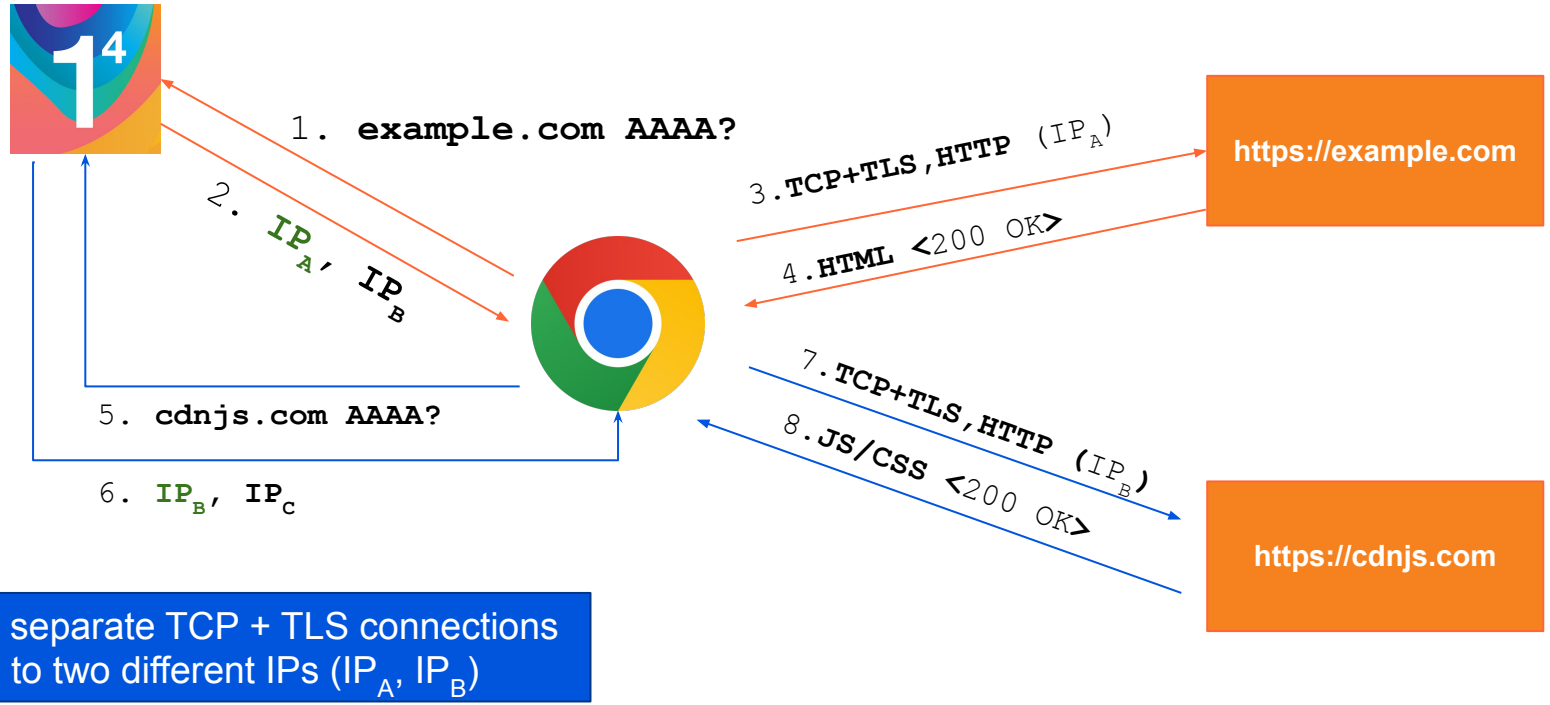

Two separate TCP + TLS connections to two different IPs (IP_A, IP_B)

CLOUDFLARE

Firefox's Approach: Transitivity between sets of IPs

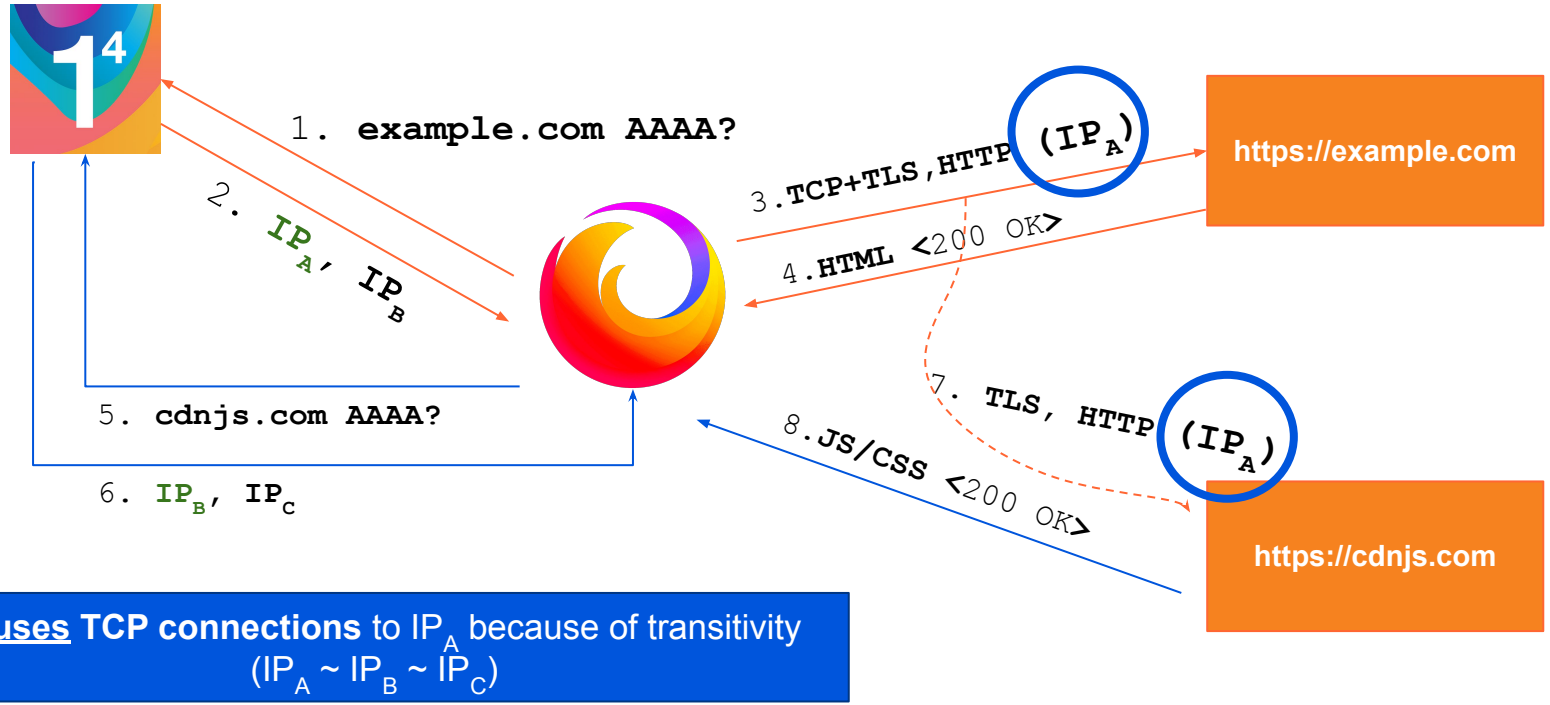

Reuses TCP connections to IP_A because of transitivity (IP_A ~ IP_B ~ IP_C)

CLOUDFLARE

Key Research Questions

A. How much of the Internet is coalescable?
   a. Where are the subresources located?
   b. How are coalescable sub-resources distributed?
B. What changes are required to enable missed opportunities?
C. Can this be done (and at scale)?

CLOUDFLARE

Where are the subresources located?

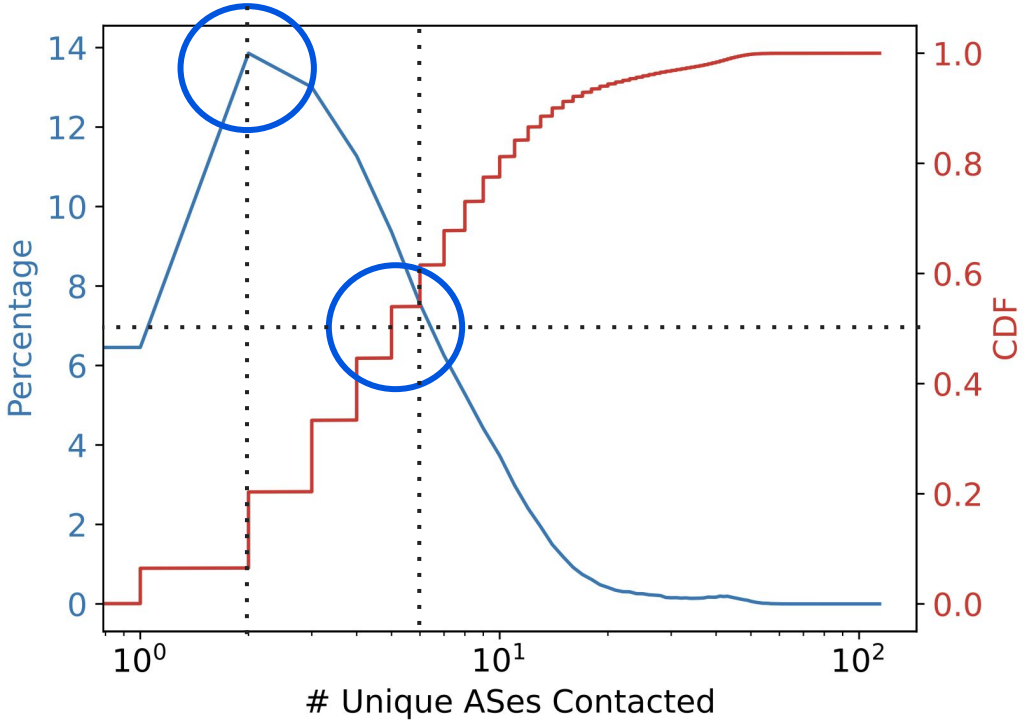

Insights:

1. 14% of web pages have a dependency on resources from one other AS.

2. More than 50% of webpages need no more than ***6 ASes*** for all subresources.

CLOUDFLARE

Where are the most coalescable sub-resources?

Rank	AS Number	Org. Name	#Req	%
1	AS 15169	Google	7932198	22.10
2	AS 13335	Cloudflare	4937395	13.75
3	AS 16509	Amazon 02	3017176	8.40
4	AS 14618	Amazon AES	2019308	5.62
5	AS 54113	Fastly	1281402	3.57
6	AS 16625	Akamai AS	1087172	3.02
7	AS 32934	Facebook	998685	2.78
8	AS 20940	Akamai Intl. B.V.	583700	1.62
9	AS 16276	OVH SAS	548107	1.52
10	AS 24940	Hetzner Online GmbH	469293	1.30
Total				**63.68**

Note: Coalescing opportunities exist because of CDNs!

Insights:

1. The top 10 ASes handle more than 60% of all web requests for subresources

2. Connection re-use potential (Min. number of connections) **could be approximated** to number of unique ASes contacted (further supported by ***Fayed et al.*** SIGCOMM '21)

CLOUDFLARE

Challenges with ORIGIN Frames (RFC 8336) (*Standardized in 2018*)

1. Default ORIGIN Frame standard allows any hostname(s) to be sent by the server (***lack of authority***).

2. Clients validate the hostnames in the ORIGIN frame for authenticity
   a. Firefox is the only client which supports ORIGIN Frame
   b. Clients resolve DNS queries and retrieve TLS Certificates
      i. If the IP addresses match IP based coalescing results.
      ii. Else, new TCP+TLS connections are made.

3. Lack of server software support for ORIGIN Frames.

4. Not widely adopted … yet, despite standardization!

CLOUDFLARE

Authoritative ORIGIN Frames (RFC 8336) could preclude DNS

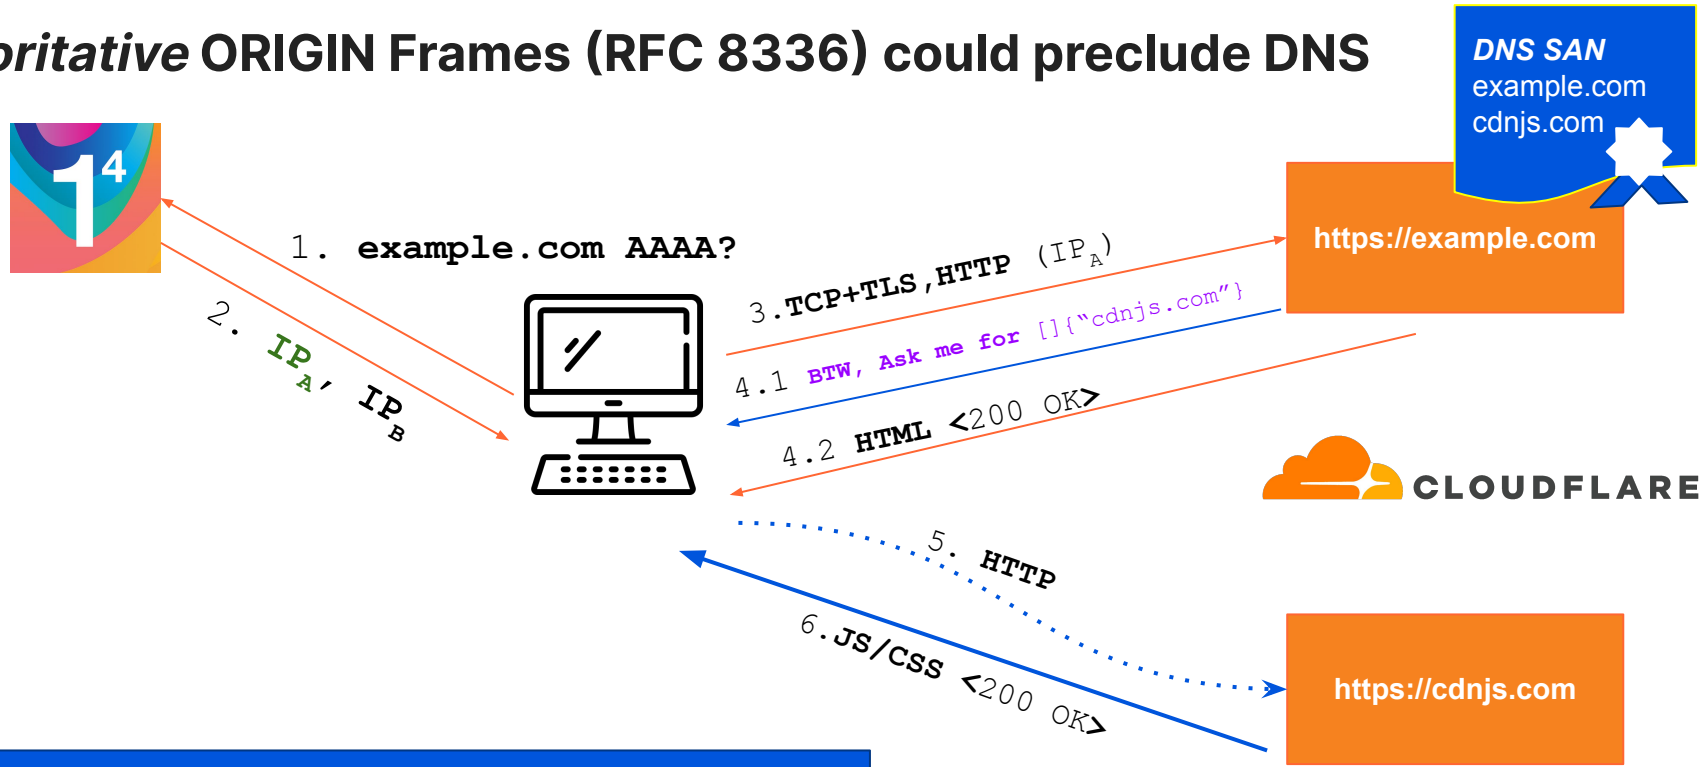

Could Prevent unwanted DNS queries if authority established

CLOUDFLARE

Authoritative ORIGIN Frames (RFC 8336) could preclude DNS

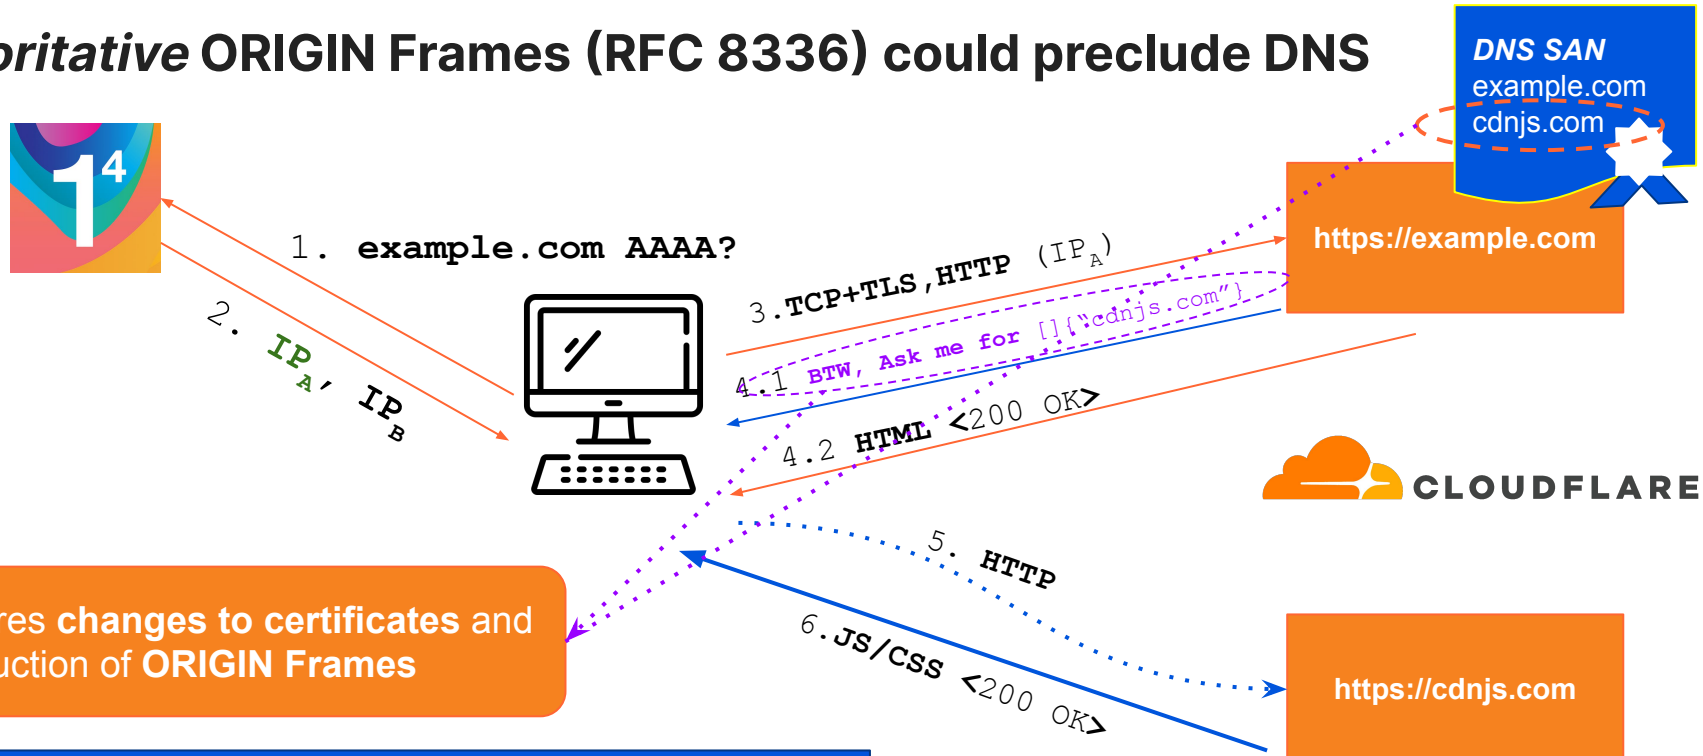

Requires **changes to certificates** and introduction of **ORIGIN Frames**

Could Prevent unwanted DNS queries if authority established

Modelling: > 60% improvement in Number of DNS and TLS connections

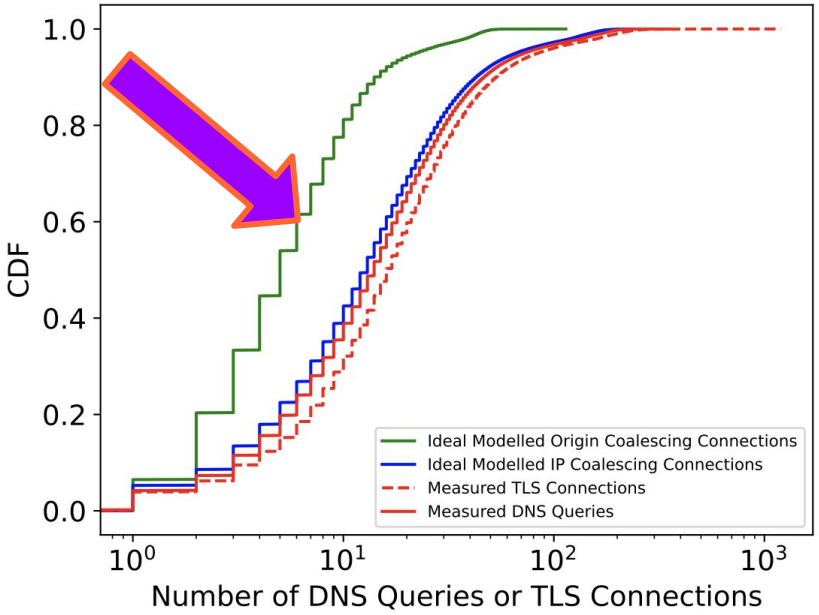

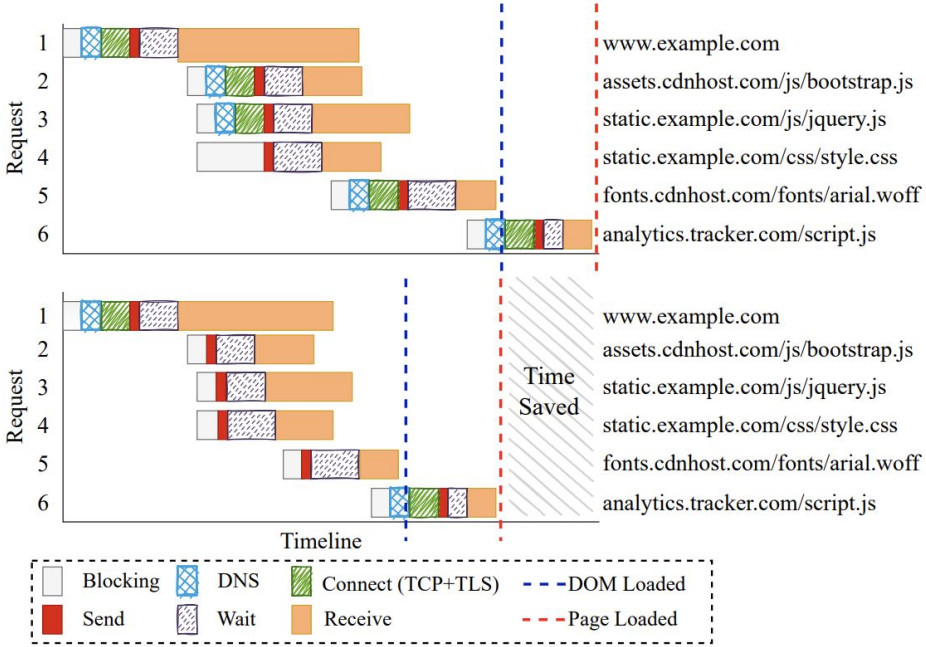

Modelled Timeline reconstruction when `*.example.com` is proxied by the CDN network also serving `cdnhost.com`

Real World – IP coalescing ties services, hard to coordinate

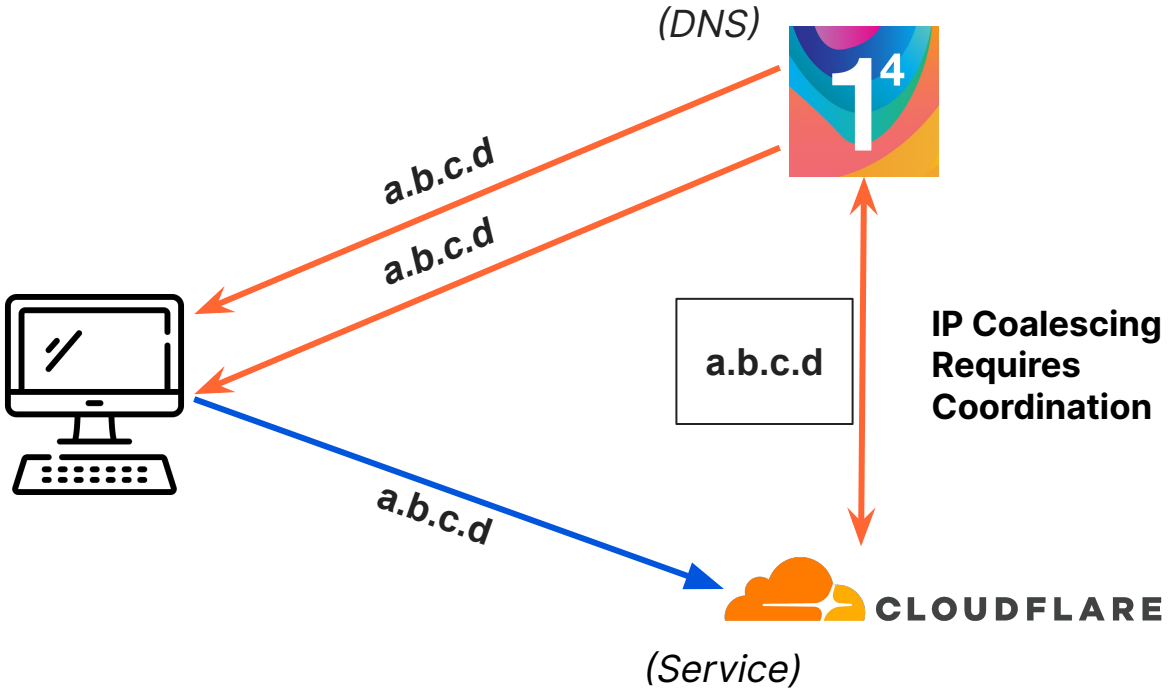

Always returns the same IP address for all services

Browser reuses connection because of same IP

Real World – ORIGIN…

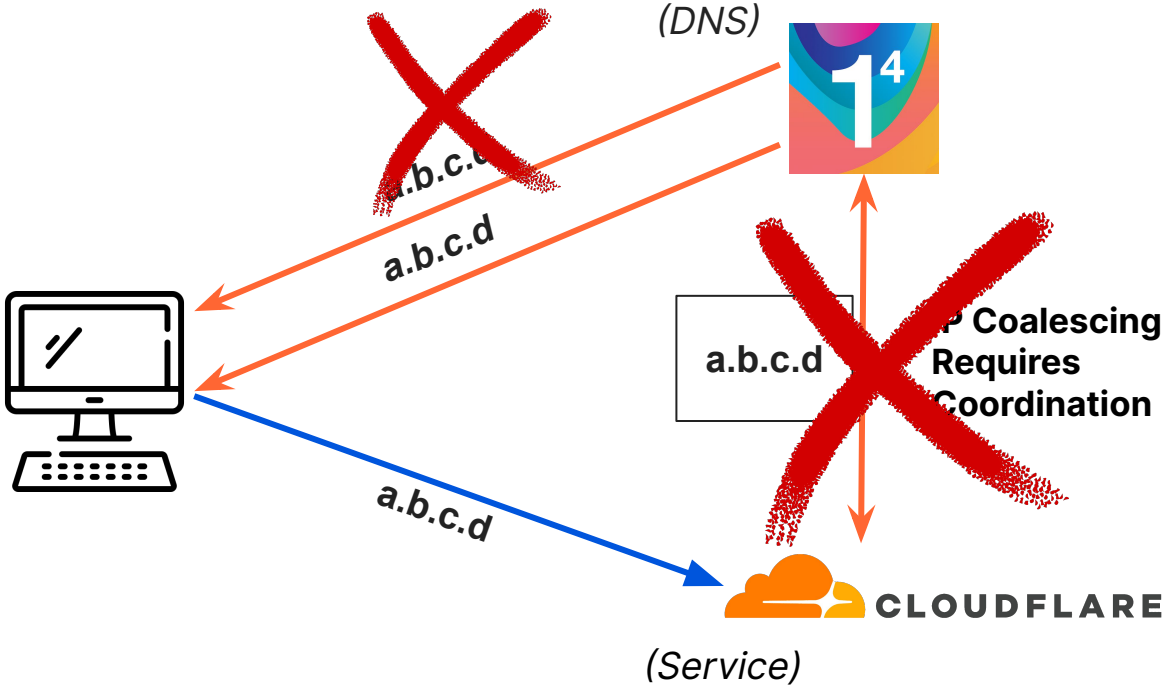

No changes to DNS or service coordination necessary!

Service sends ORIGIN Frames

Advantage: Does not disrupt existing traffic engineering practices

Real World – ORIGIN… makes coalescing practical.

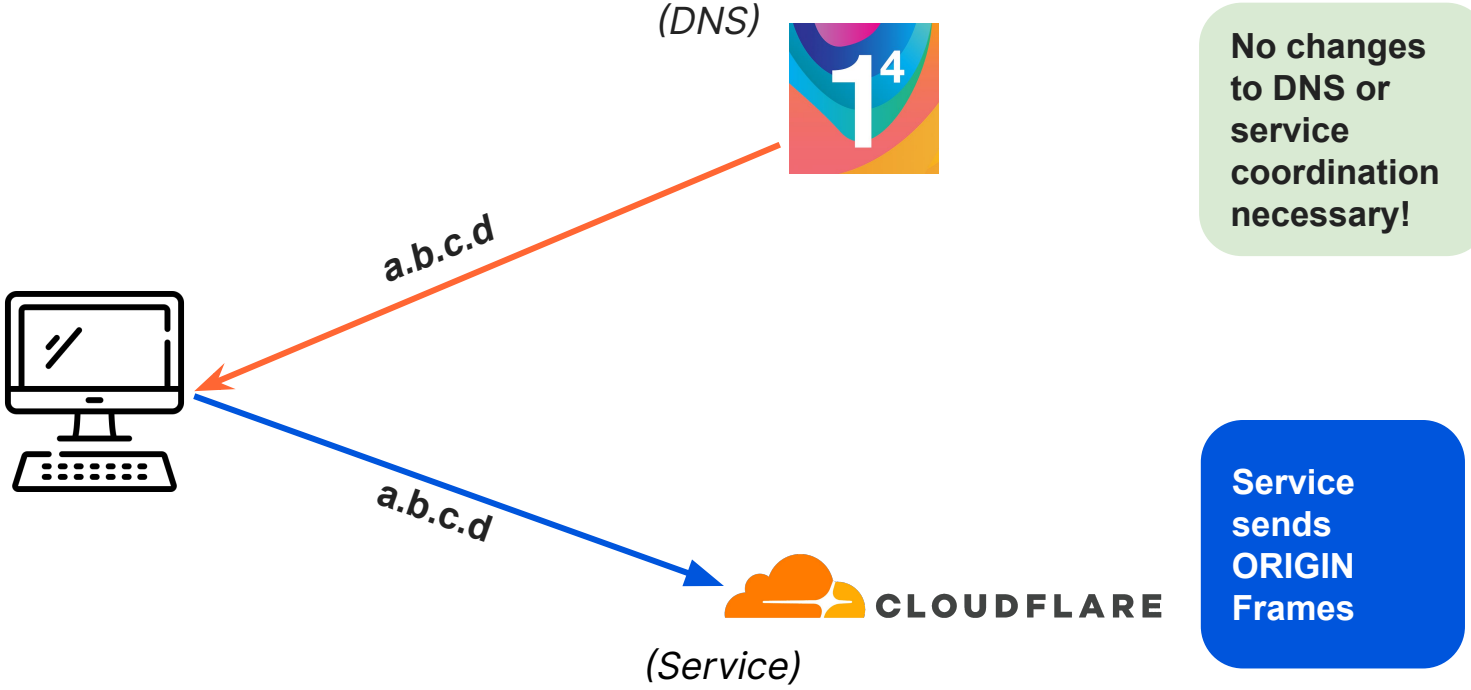

No changes to DNS or service coordination necessary!

Service sends ORIGIN Frames

Advantage: Little difference to ‘wire-line’ activities

Takeaway 1: Connection Coalescing works in practice!

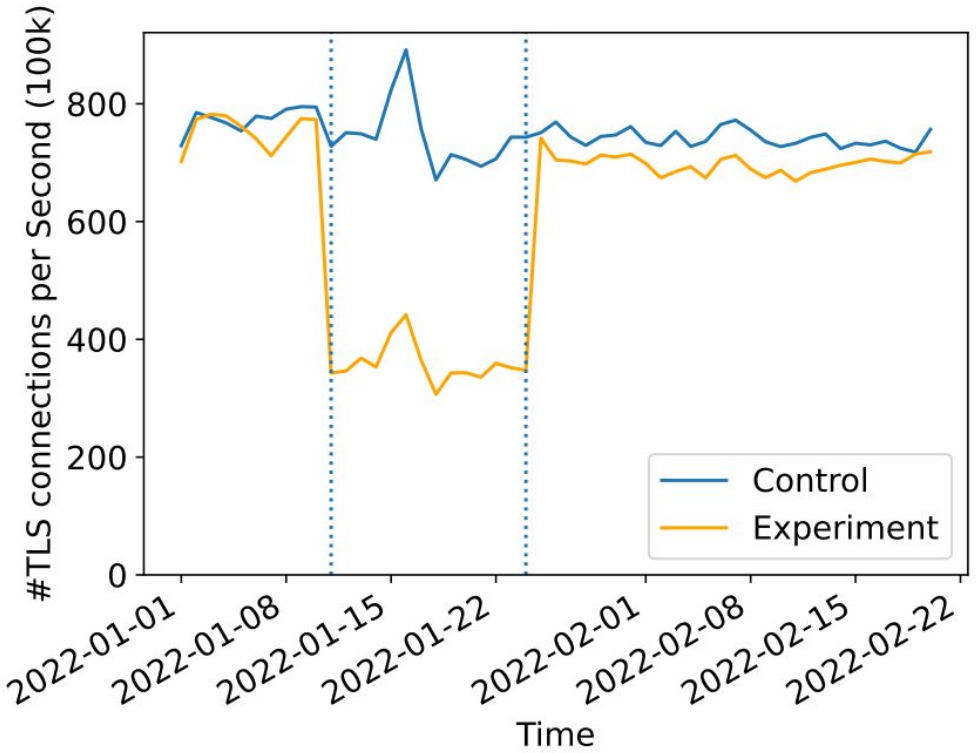

~50% reduction in number of new connections to the cdnjs hostname we attempted coalescing to.

Client: Reduced Number of Cryptographic Certificate Validations.

Client: Active measurements show ~65-70% connections coalesced.

Server: Reduced number of connections → allow more client connections

CLOUDFLARE

Takeaway 2: No-worse performance, almost immeasurable improvement

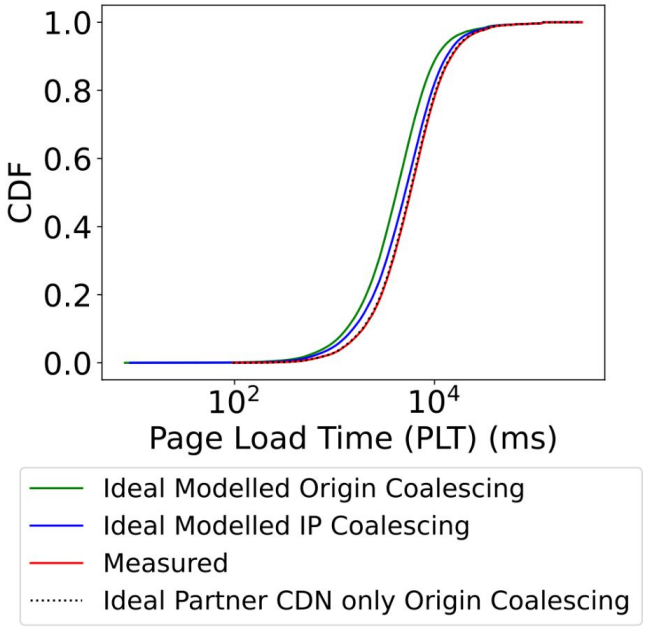

(a) Measured and modelled.

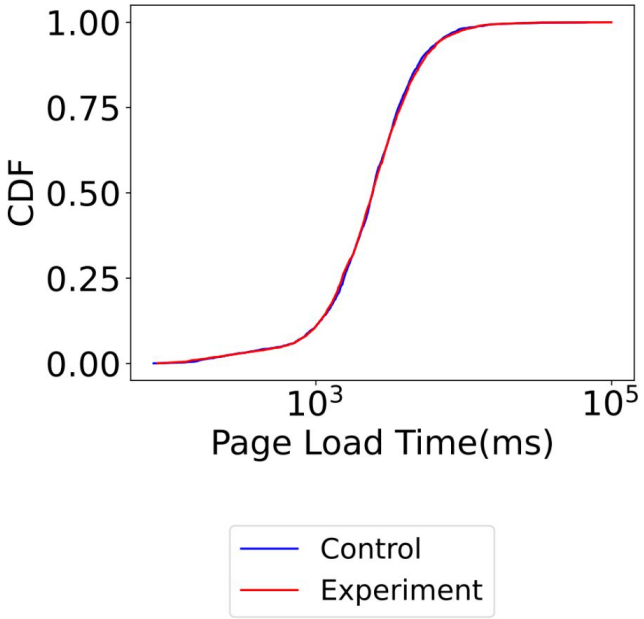

(b) IP and ORIGIN

Wide impact depends on **path characteristics,** AND **bottleneck share,** AND **numbers of operators that support ORIGIN.**

CLOUDFLARE

Takeaway 2: No-worse performance, almost immeasurable improvement

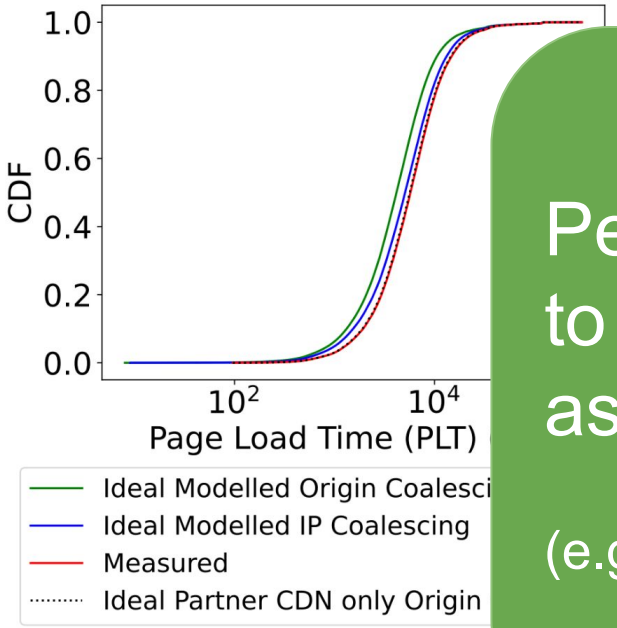

(a) Measured and mode

Performance cannot be assumed to improve, and **should** be avoided as primary motivation.

(e.g. fewer bottleneck connections, and more)

Open Source: We contribute a public large scale server implementation

ORIGIN Frames are yet to see large scale adoption, no public server implementation exists.

We contribute (to our knowledge) the first public server side implementation.

https://github.com/cloudflare/net-originframe

https://github.com/cloudflare/go-originframe

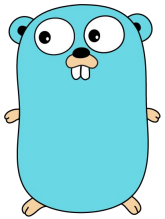

Golang

CLOUDFLARE

Needs Careful Deployment – Non RFC Compliant network stacks exist

Deployment of ORIGIN Frames resulted in uncovering **compliance issues** in popular antivirus and Internet security software which did not drop unknown HTTP/2 frames and instead resulted in connection tear-down.

CLOUDFLARE

Claim: ORIGIN Frame based Coalescing improves privacy

Whose privacy?
What does it mean?

Each coalesced connection hides an otherwise exposed plaintext SNI and prevents at-least one additional plaintext DNS query-response.

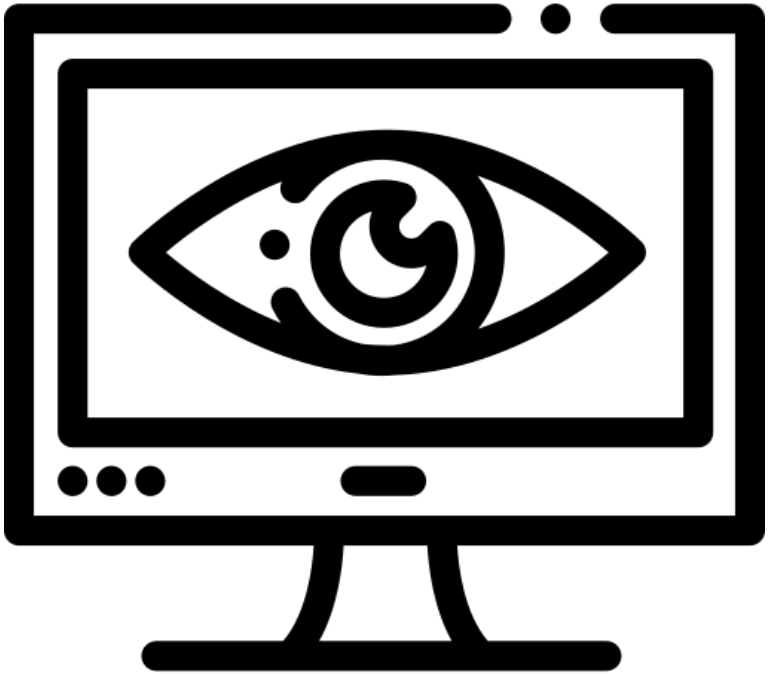

Resource Scheduling Opportunities at the Endpoints!

Extra Connections → non-deterministic path characteristics

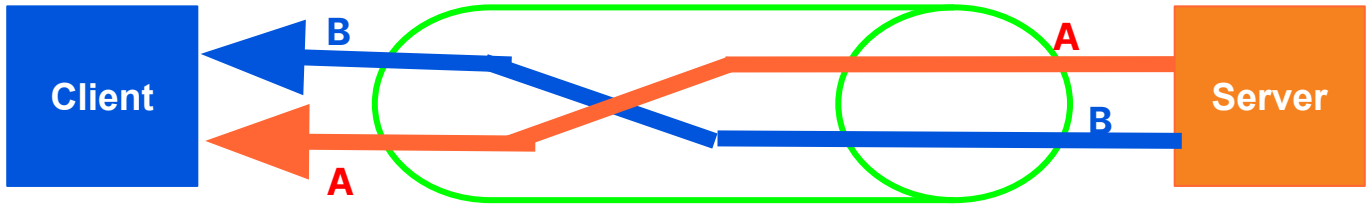

CLOUDFLARE

Resource Scheduling Opportunities at the Endpoints!

Extra Connections → non-deterministic path characteristics

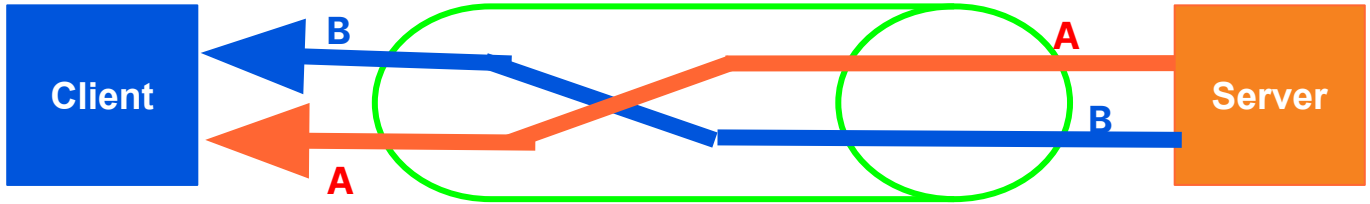

Call to action: Implement ORIGIN Frames!

CLOUDFLARE

Thank You!

sudheesh@cs.washington.edu
marwan@cloudflare.com

Link to paper: https://dl.acm.org/doi/10.1145/3517745.3561453

We'd like to thank Larry Archer, Michel Bamps, Petros Gigis, Vasileios Giotsas, Vânia Gonçalves, John Graham-Cumming, Mihir Jham, Lucas Pardue, Kyle Schomp, and Avani Wildani.