

Demystifying SPDY and HTTP/2

Jonathan Klein
@jonathanklein
April 21, 2015

Slides, Links
jkle.in/http2

H/2 == HTTP/2

*What Every Web Developer Should Know About
Networking and Browser Performance*




High Performance

Browser Networking

O'REILLY®

Ilya Grigorik

Figures from
this book



About Me

About Me

- Formerly a performance engineer at Etsy

About Me

- Formerly a performance engineer at Etsy
- Started the Boston Web Perf Meetup Group

About Me

- Formerly a performance engineer at Etsy
- Started the Boston Web Perf Meetup Group
- Former CTO at a pre-product startup in Boston

About Me

- Formerly a performance engineer at Etsy
- Started the Boston Web Perf Meetup Group
- Former CTO at a pre-product startup in Boston
- Planned on launching with HTTPS everywhere and H/2 support

SPDY

Experimental protocol
developed at Google

SPDY

Experimental protocol
developed at Google

H/2

Official second major version of
HTTP, based on SPDY

SPDY's Future, H/2 Status

SPDY's Future, H/2 Status

- H/2 approved by the IETF for publication as standards-track RFCs

SPDY's Future, H/2 Status

- H/2 approved by the IETF for publication as standards-track RFCs
- Chrome removing SPDY support in 2016

mnot's blog

"Design depends largely on constraints." — Charles Eames

Wednesday, 18 February 2015

HTTP/2 is Done


The IESG has formally approved the HTTP/2 and HPACK specifications, and they're on their way to the RFC Editor, where they'll soon be assigned RFC numbers, go through some editorial processes, and be published.

We should have an official entry about this up soon over at the IETF blog, but I want to recognise a few people specifically:

- Mike Belshe and Roberto Peon brought SPDY to the IETF to standardise it. While a few have painted Google as forcing the protocol upon us, anyone who actually interacted with Mike and Roberto in the group knows that they came with the best of intent, patiently explaining the reasoning behind their design, taking in criticism, and working with everyone to evolve the protocol. Roberto also served alongside Herve Ruellan as editor of HPACK, after they merged their competing proposals for header compression.
- Martin Thomson is a fantastic editor, and we couldn't have had such a high-quality document (read the IESG reviews) without his tireless work.
- Barry Leiba is our Area Director and supported this work within the IESG, doling out wisdom and advice to me along the way.



Hi, I'm Mark Nottingham. I currently chair the [IETF HTTP Working Group](#) and am a member of the [W3C TAG](#). I usually write here about the Web, protocol design, HTTP, and caching. Once in a while, I'll write about economics, travel or living in Australia. [Find out more.](#)

 Follow @mnot 4,590 followers

og, but I want to recognise a

dise it. While a few have
y interacted with Mike and
patiently explaining the
everyone to evolve the
IPACK, after they merged their

a high-quality document (read

IESG, doling out wisdom and



Hi, I'm Mark Nottingham. I currently chair the [IETF HTTP Working Group](#) and am a member of the [W3C TAG](#). I usually write here about the Web, protocol design, HTTP, and caching. Once in a while, I'll write about economics, travel or living in Australia.

[Find out more.](#)



Follow @mnot


4,590 followers

HTTP/2 protocol / SPDY - OTHER

Global

72.15% + 7.98% = 80.13%

Networking protocol for low-latency transport of content over the web. Originally started out from the SPDY protocol, now standardized as HTTP version 2.



Global

$$72.15\% + 7.98\% = 80.13\%$$

ari *

Opera Mini *

Android
Browser *

Chrome for
Android

		34	
		36	
		37	
		38	
8			
9	31	39	
10	35	40	
11	36	41	
TP	37	42	
	38	43	
	39	44	

TCP/IP Network Model

TCP/IP Layers

Application Layer
Transport Layer
Network Layer
Network Interface Layer

TCP/IP Protocols

HTTP	FTP	Telnet	SMTP	DNS
TCP			UDP	
IP		ARP	ICMP	IGMP
Ethernet	Token Ring		Other Link-Layer Protocols	


TCP/IP Network Model

TCP/IP Layers

Application Layer
Transport Layer
Network Layer
Network Interface Layer

TCP/IP Protocols

HTTP	FTP	Telnet	SMTP	DNS
TCP			UDP	
IP		ARP	ICMP	IGMP
Ethernet	Token Ring		Other Link-Layer Protocols	



TCP/IP Network Model

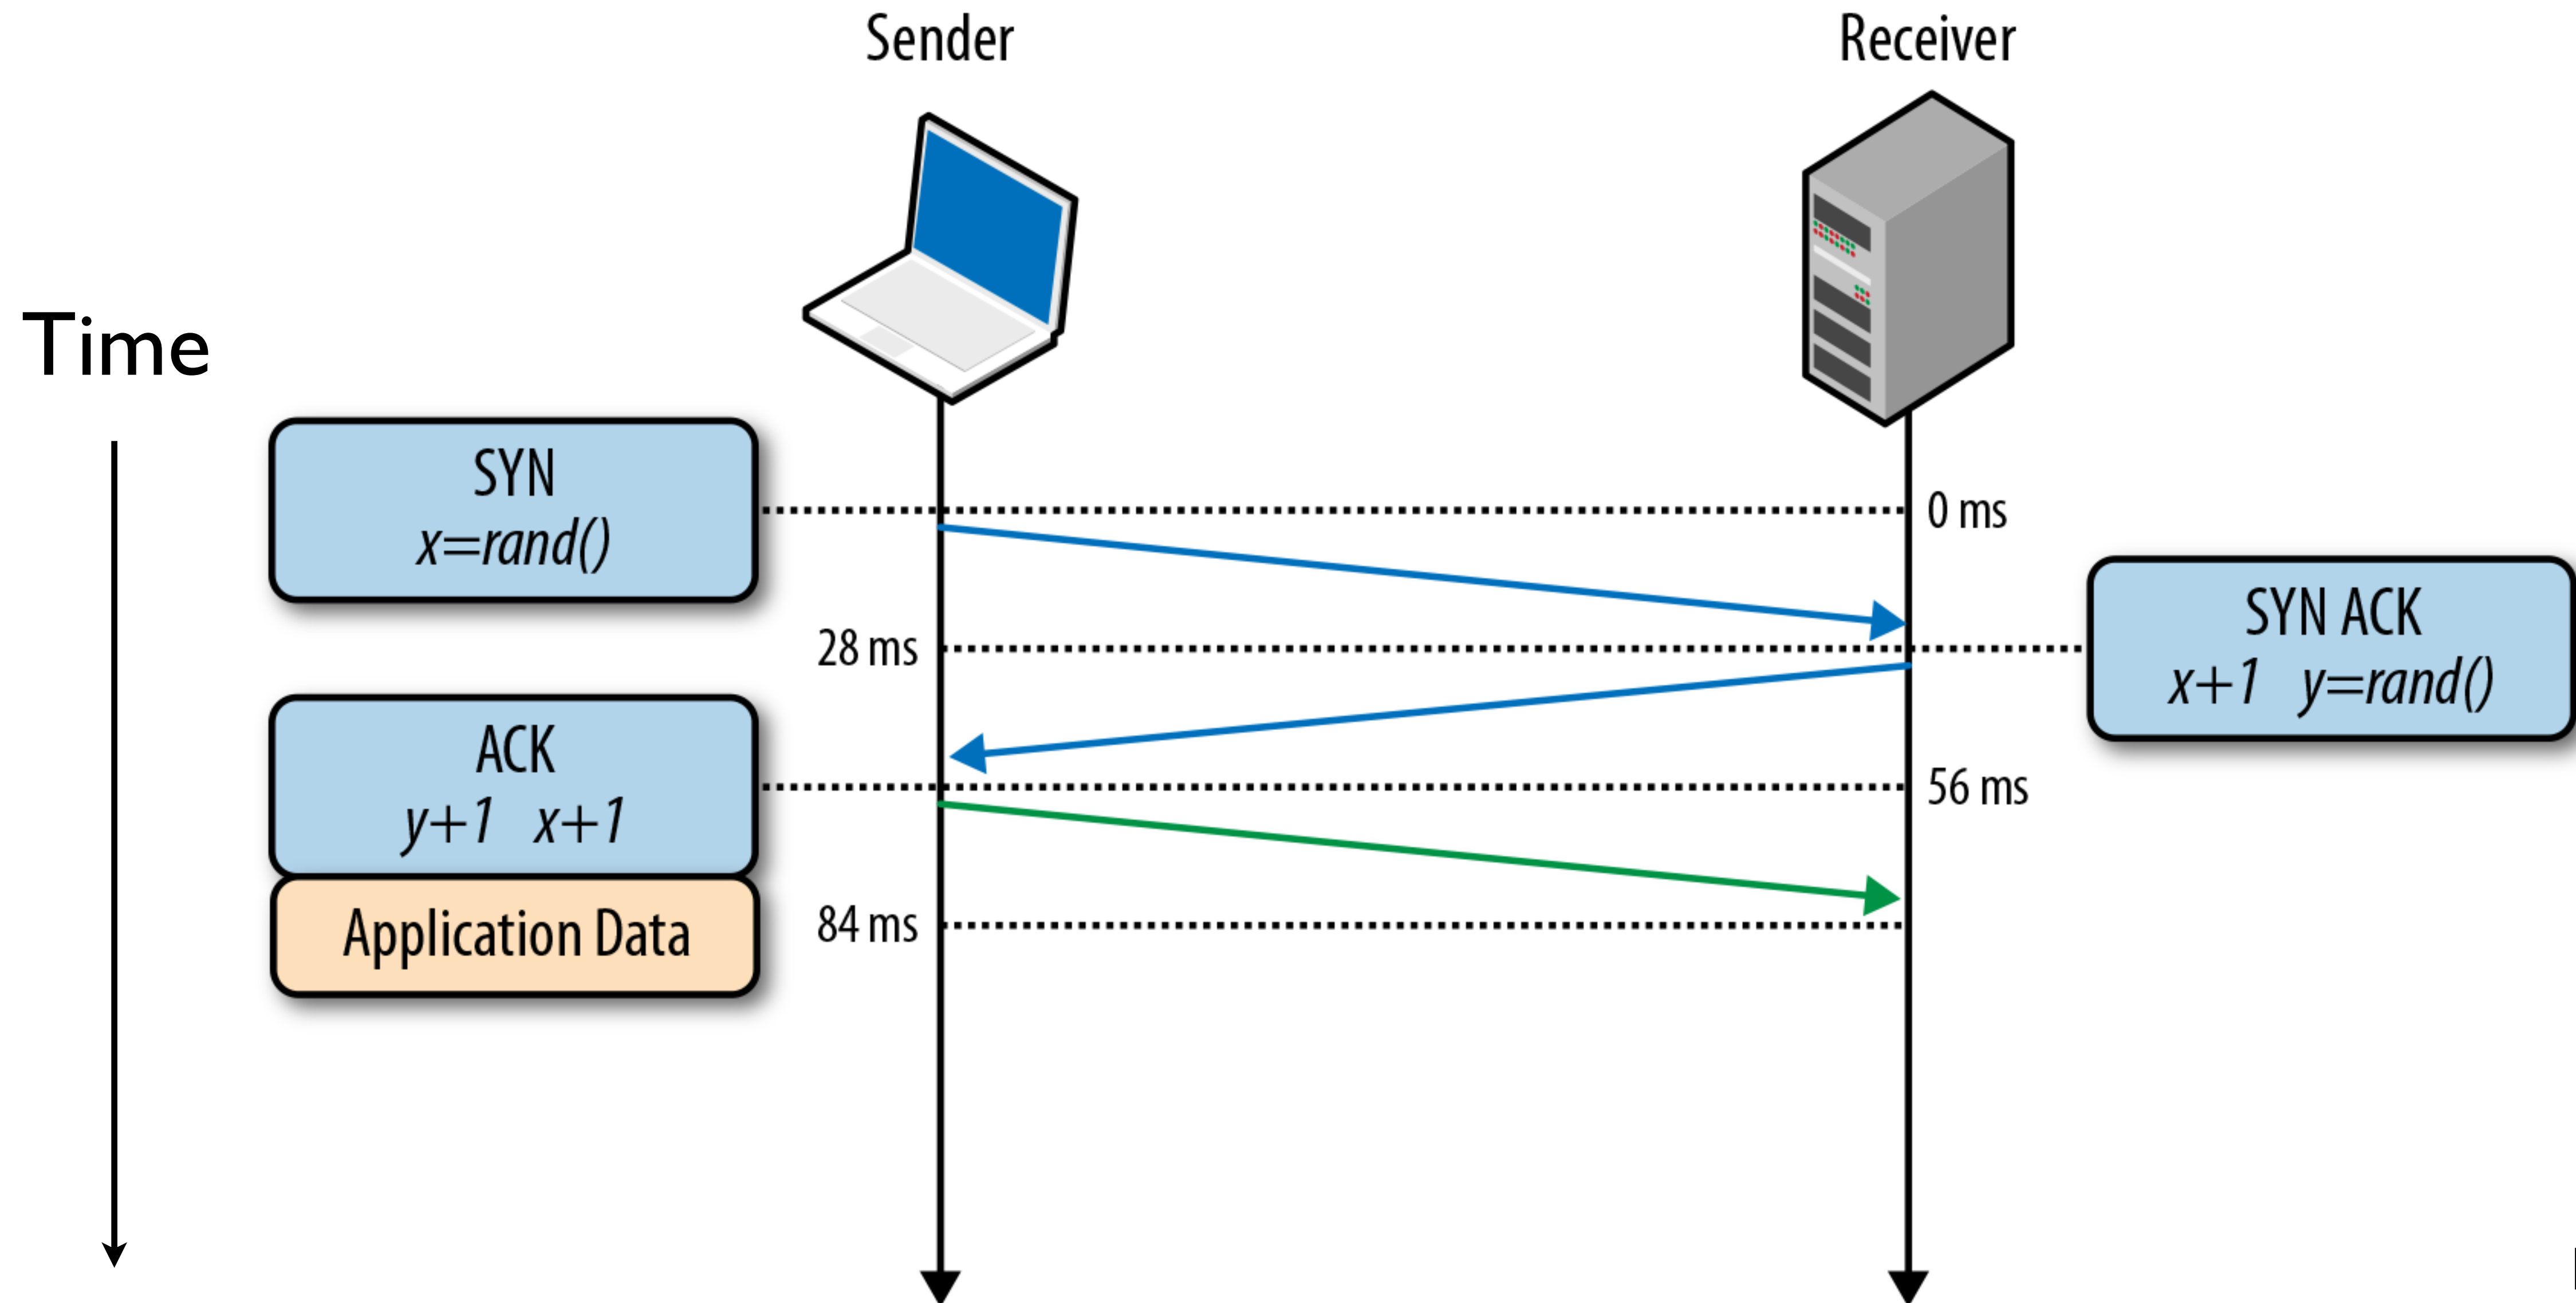
TCP/IP Layers

Application Layer
Transport Layer
Network Layer
Network Interface Layer

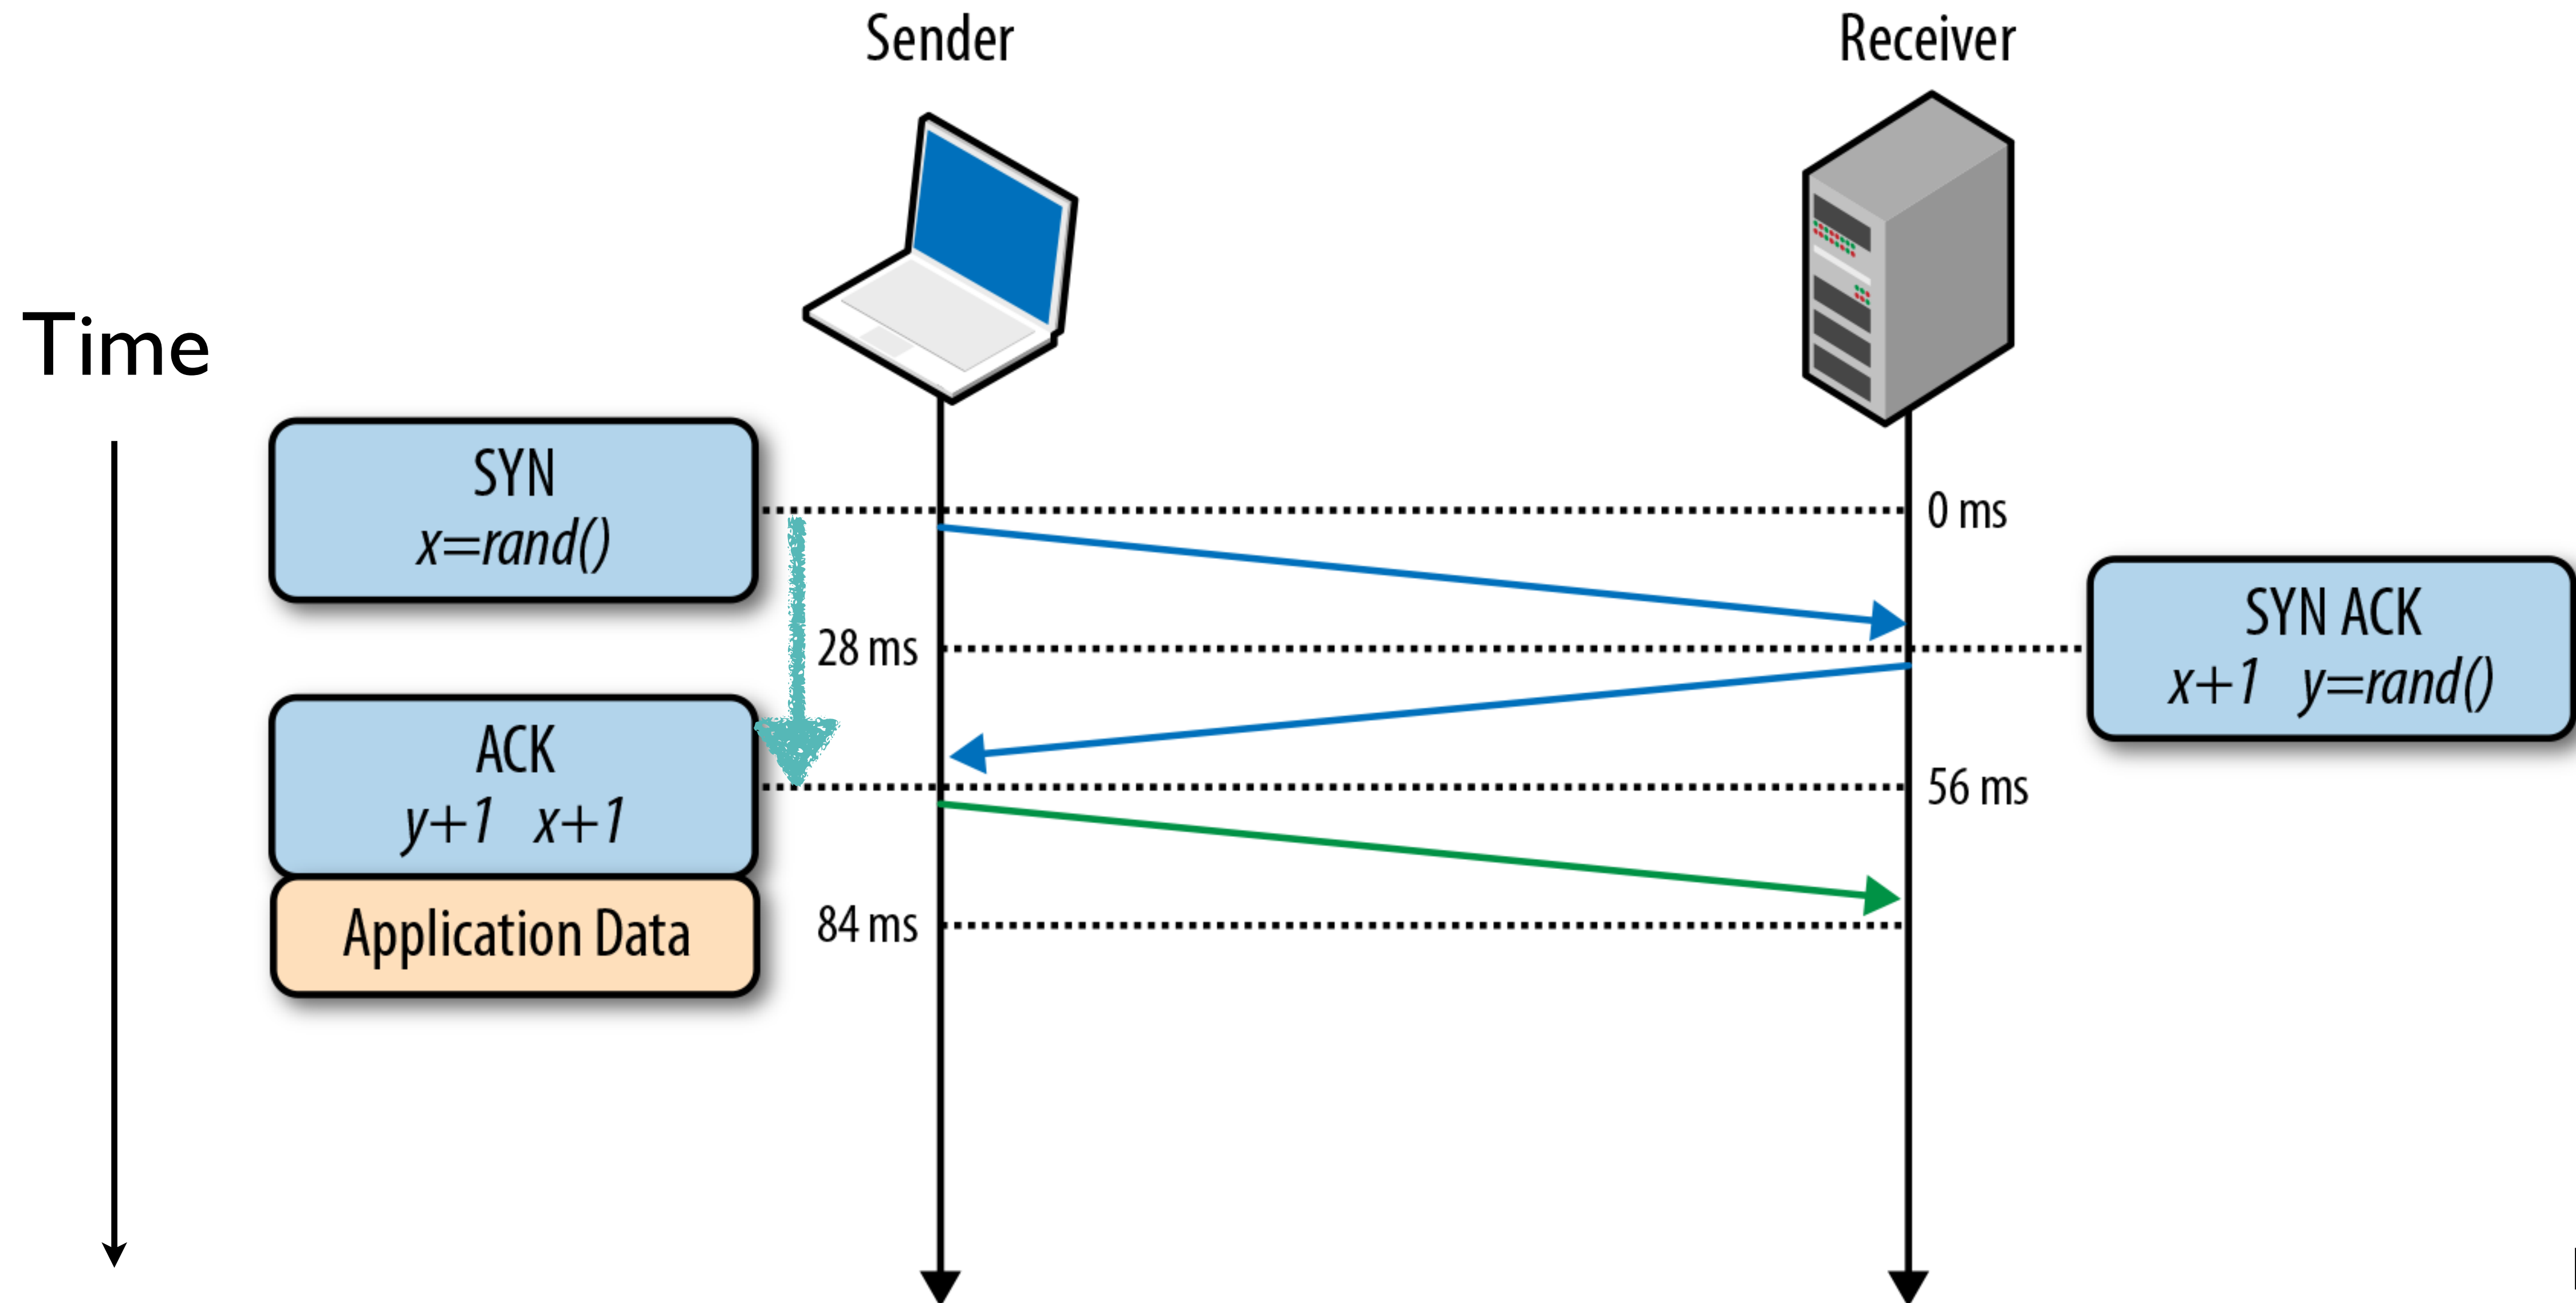
TCP/IP Protocols

HTTP	FTP	Telnet	SMTP	DNS
TCP			UDP	
IP		ARP	ICMP	IGMP
Ethernet	Token Ring		Other Link-Layer Protocols	

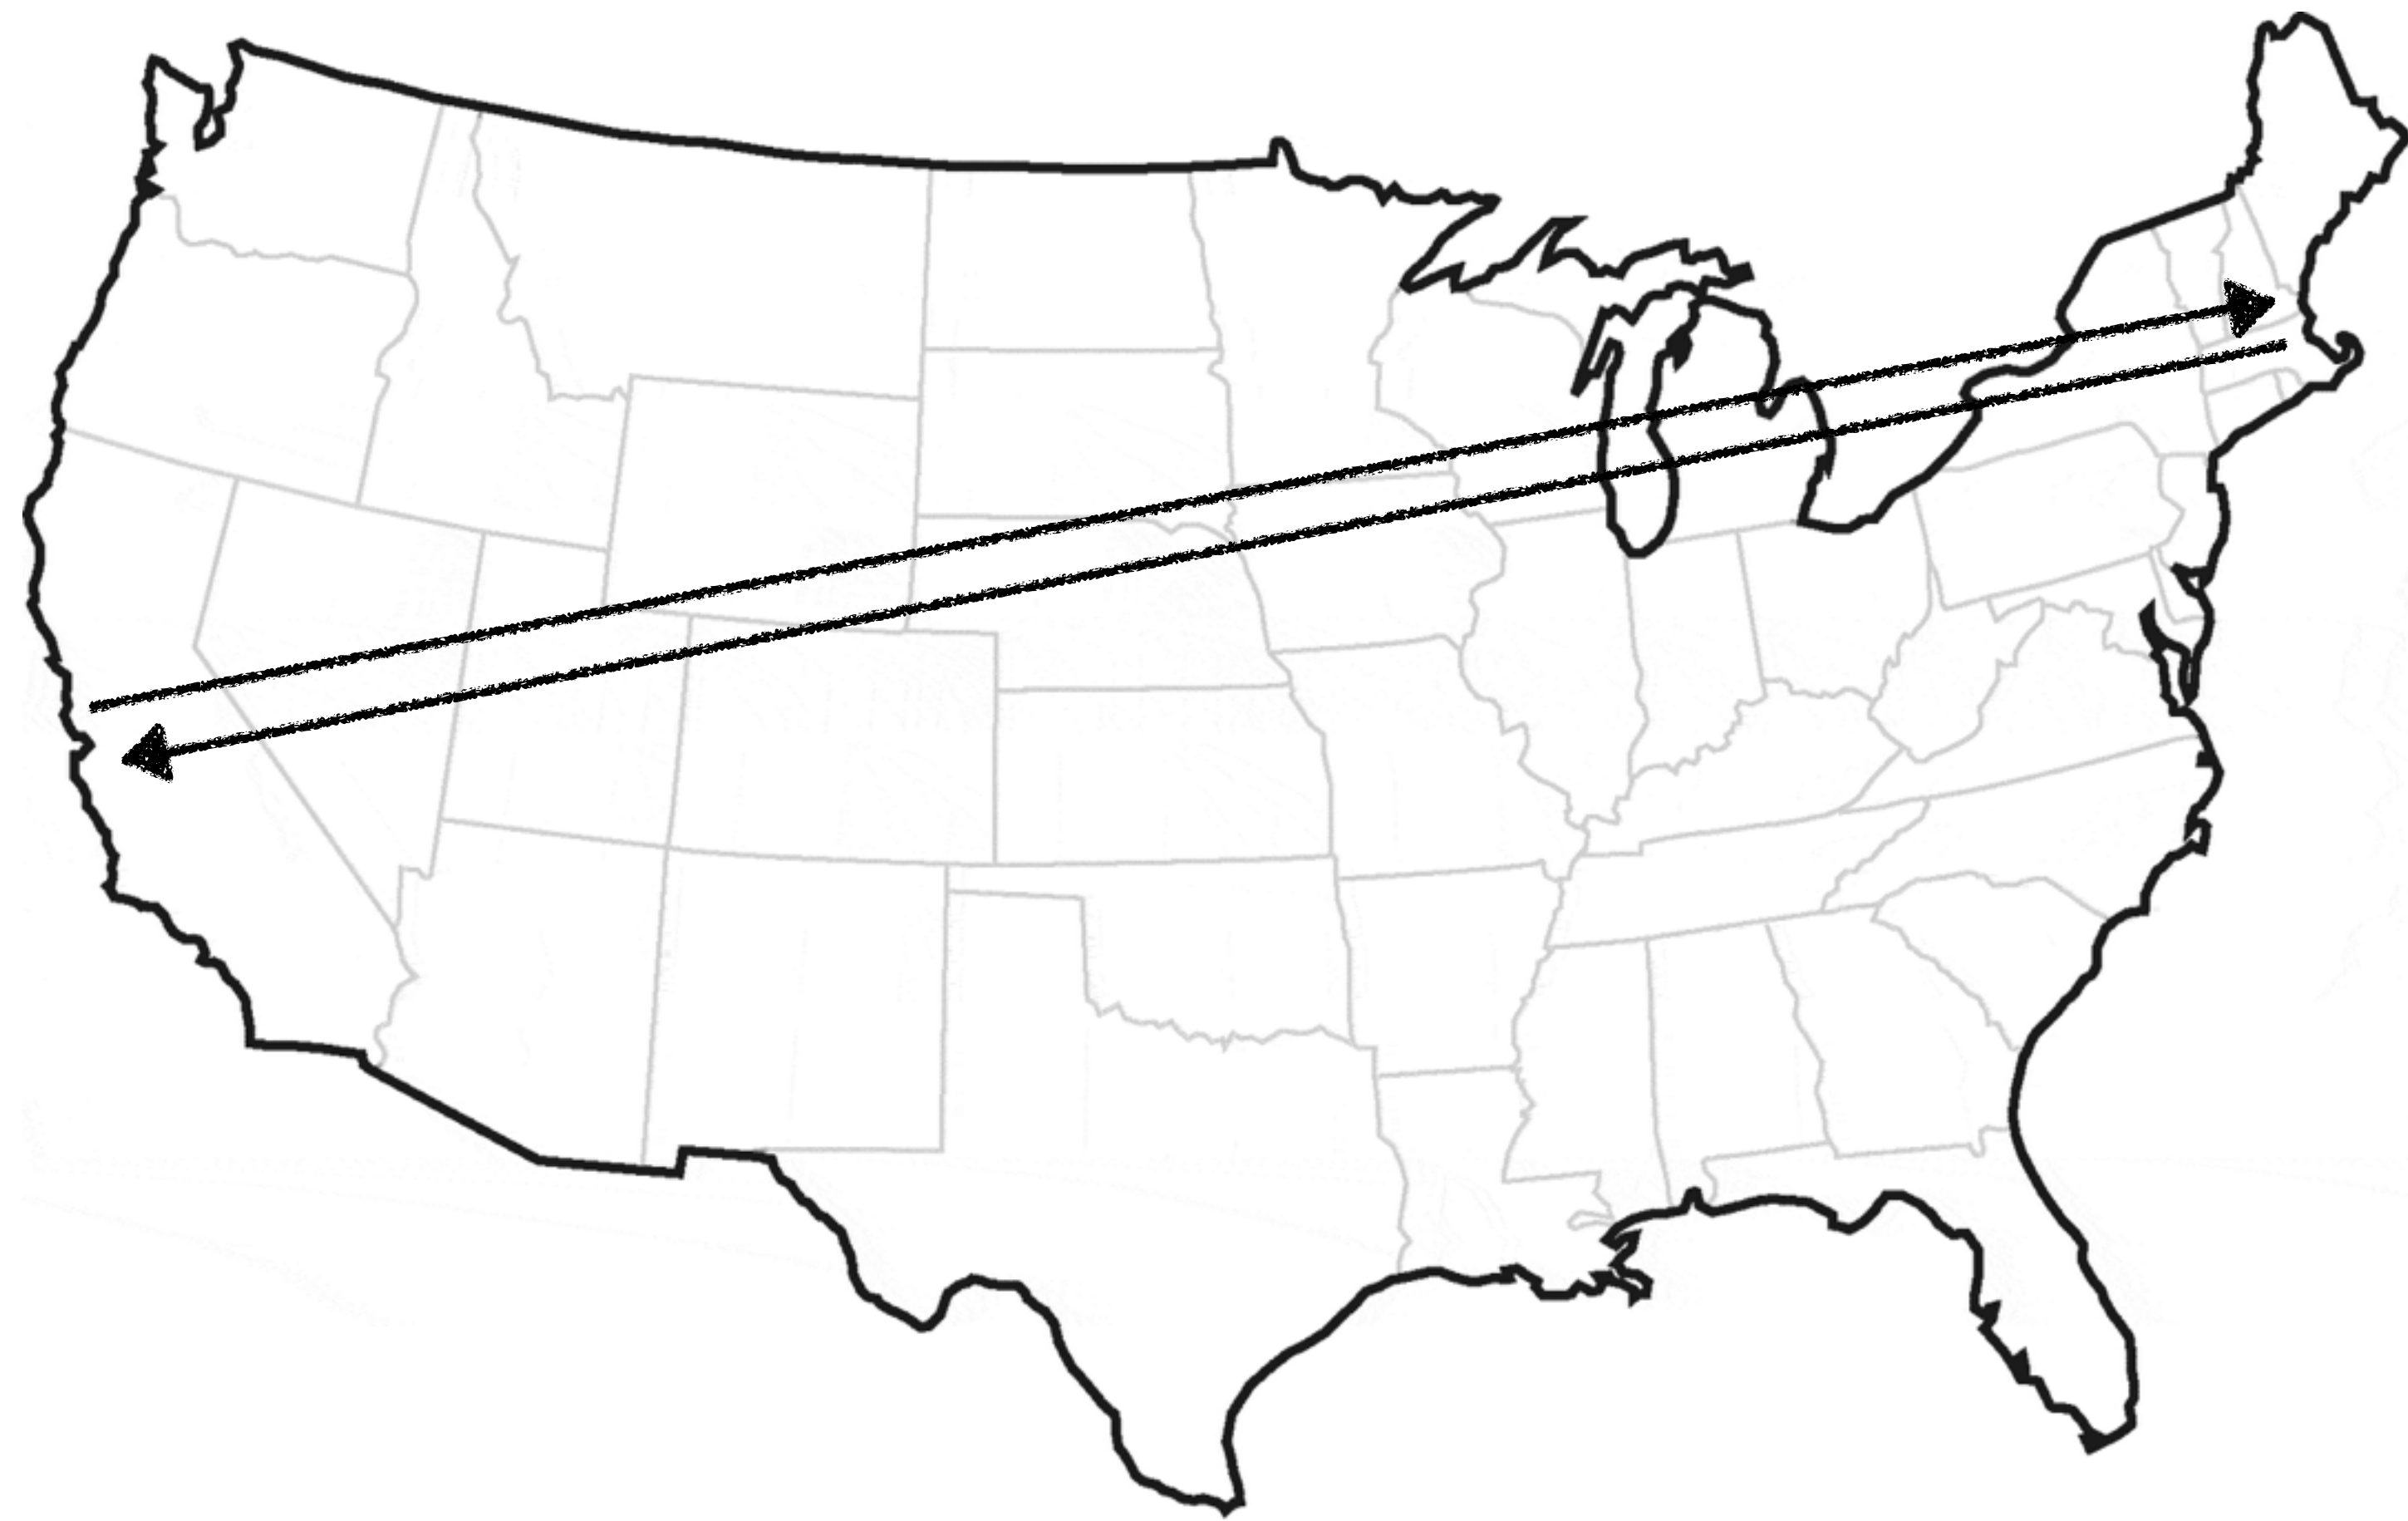
TCP Three Way Handshake



TCP Three Way Handshake





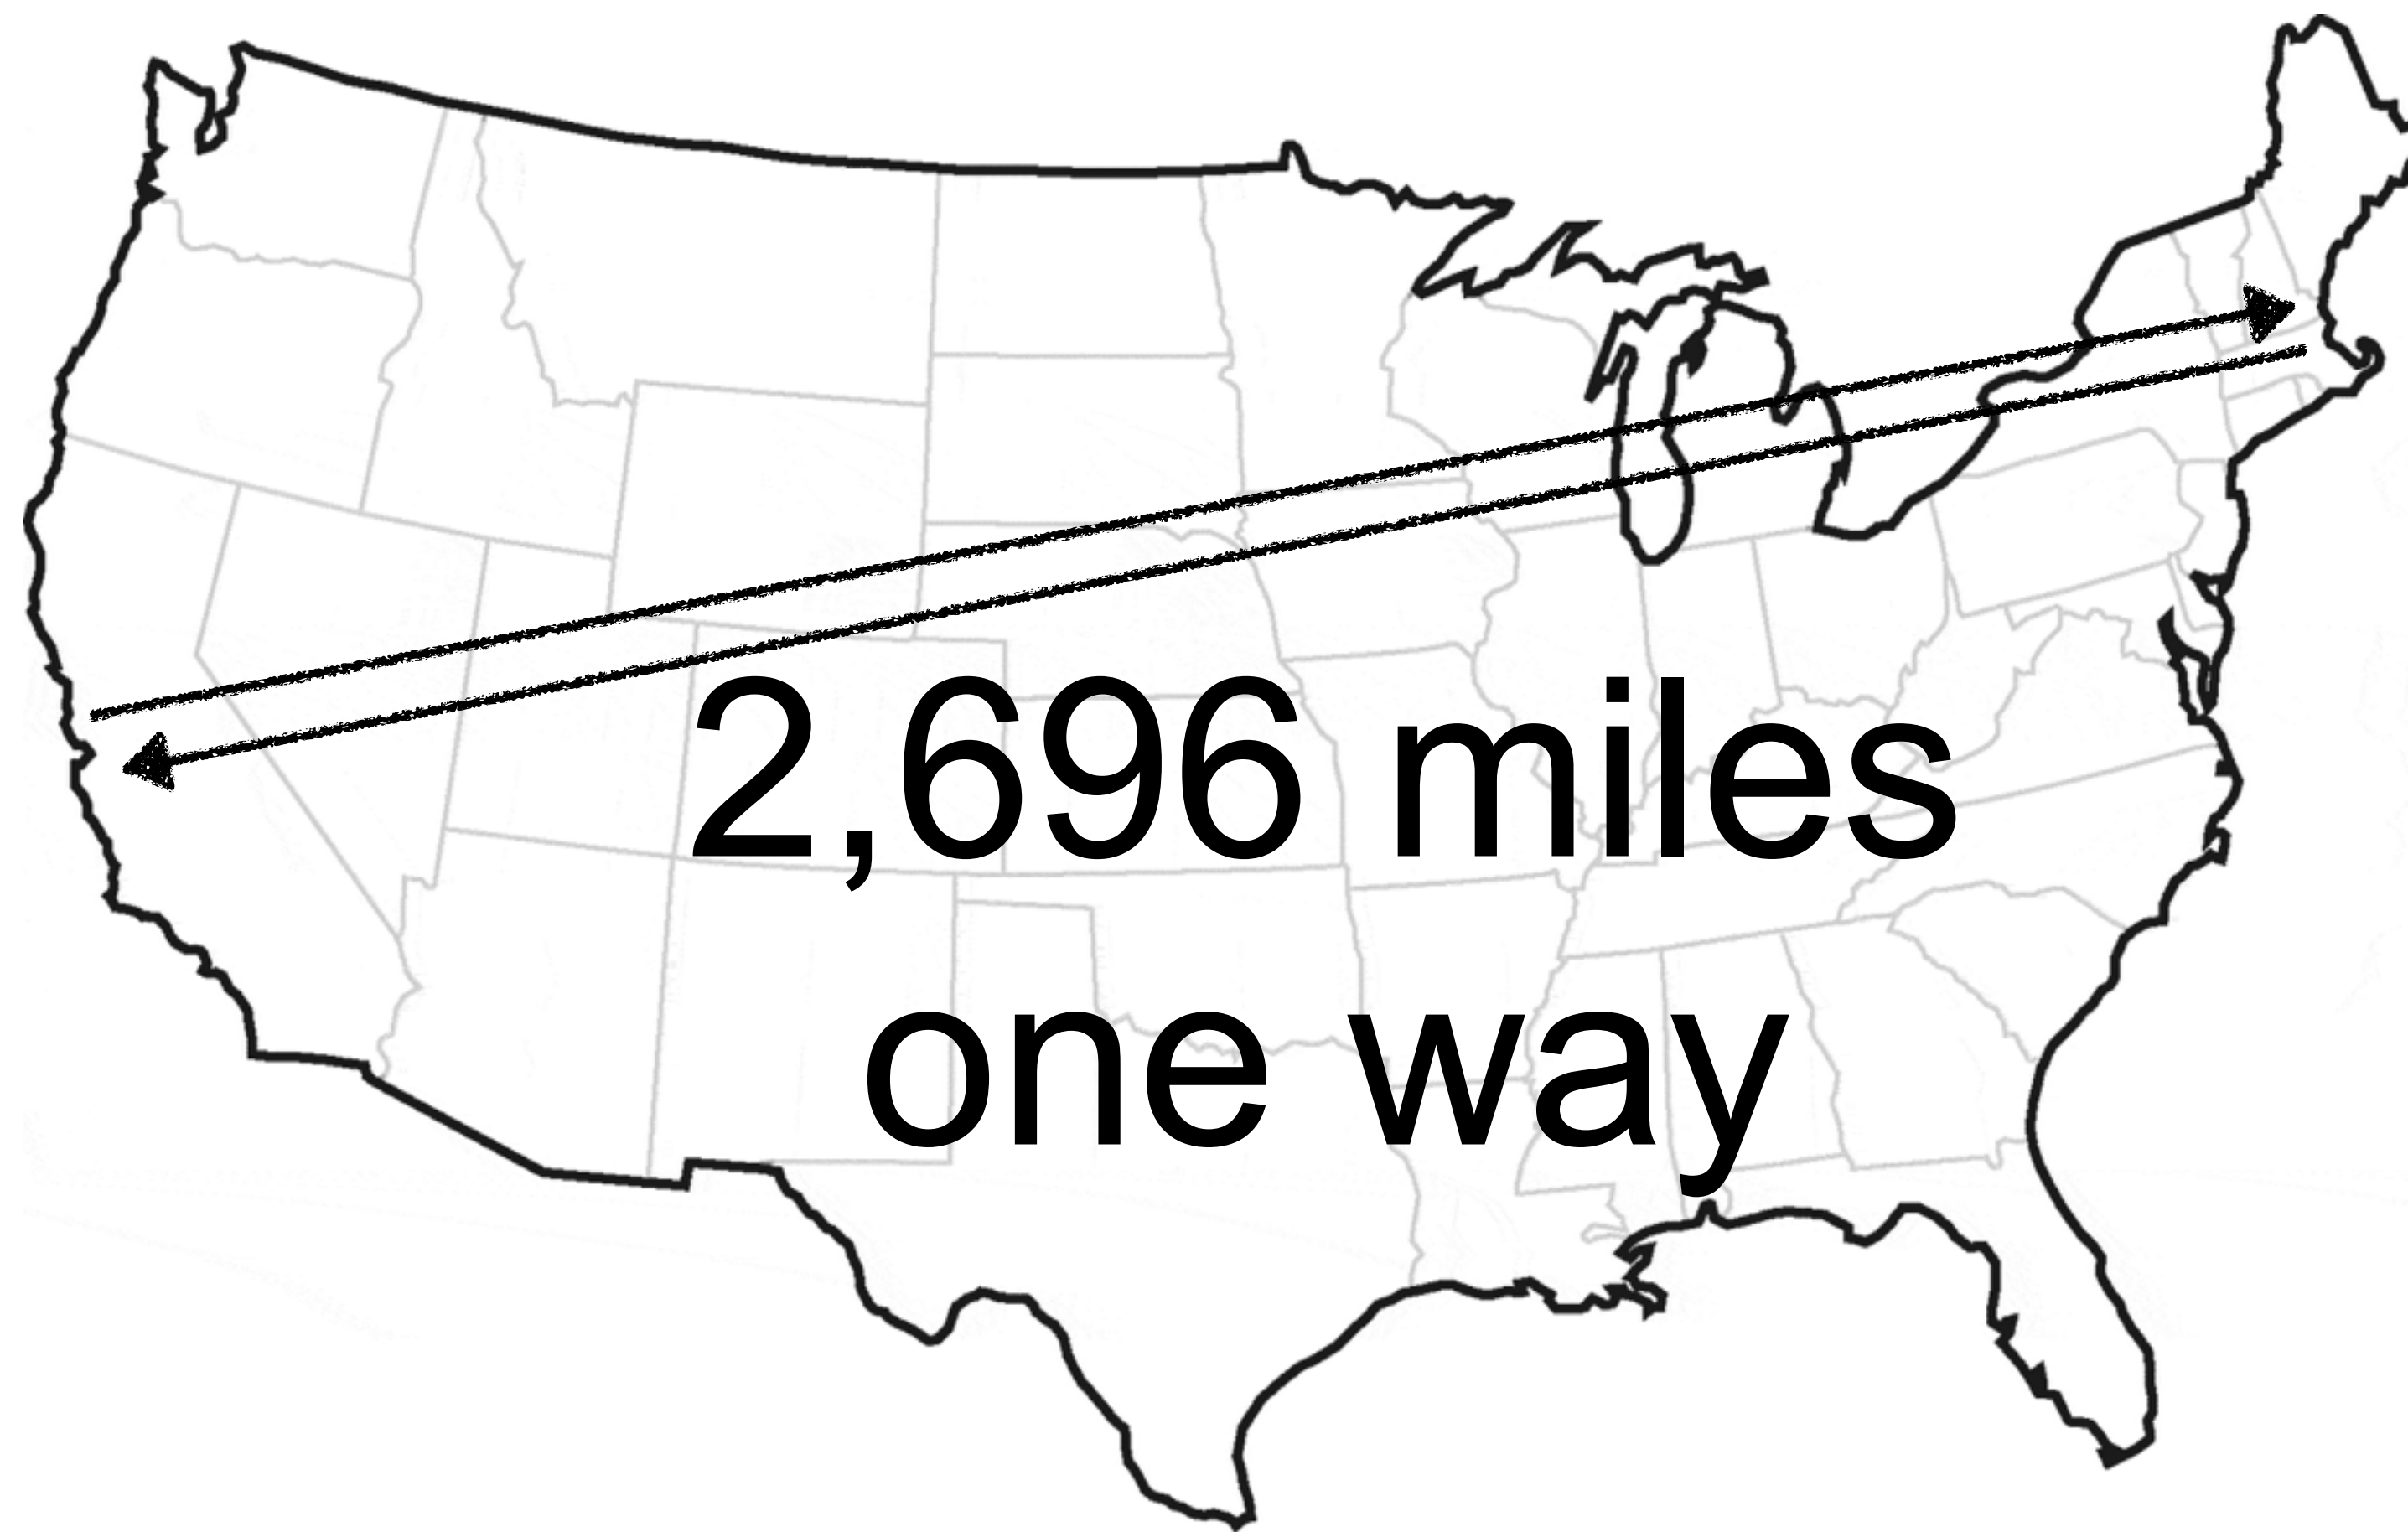




2,696 miles
one way



Speed of Light = 186,000 miles / sec




Speed of Light = 186,000 miles / sec

Assumed RTT = 28ms


📶 Global Ping Statistics → Boston and San Francisco

Showing historical ping data between major cities. Bar shows Average \pm Median Deviation.

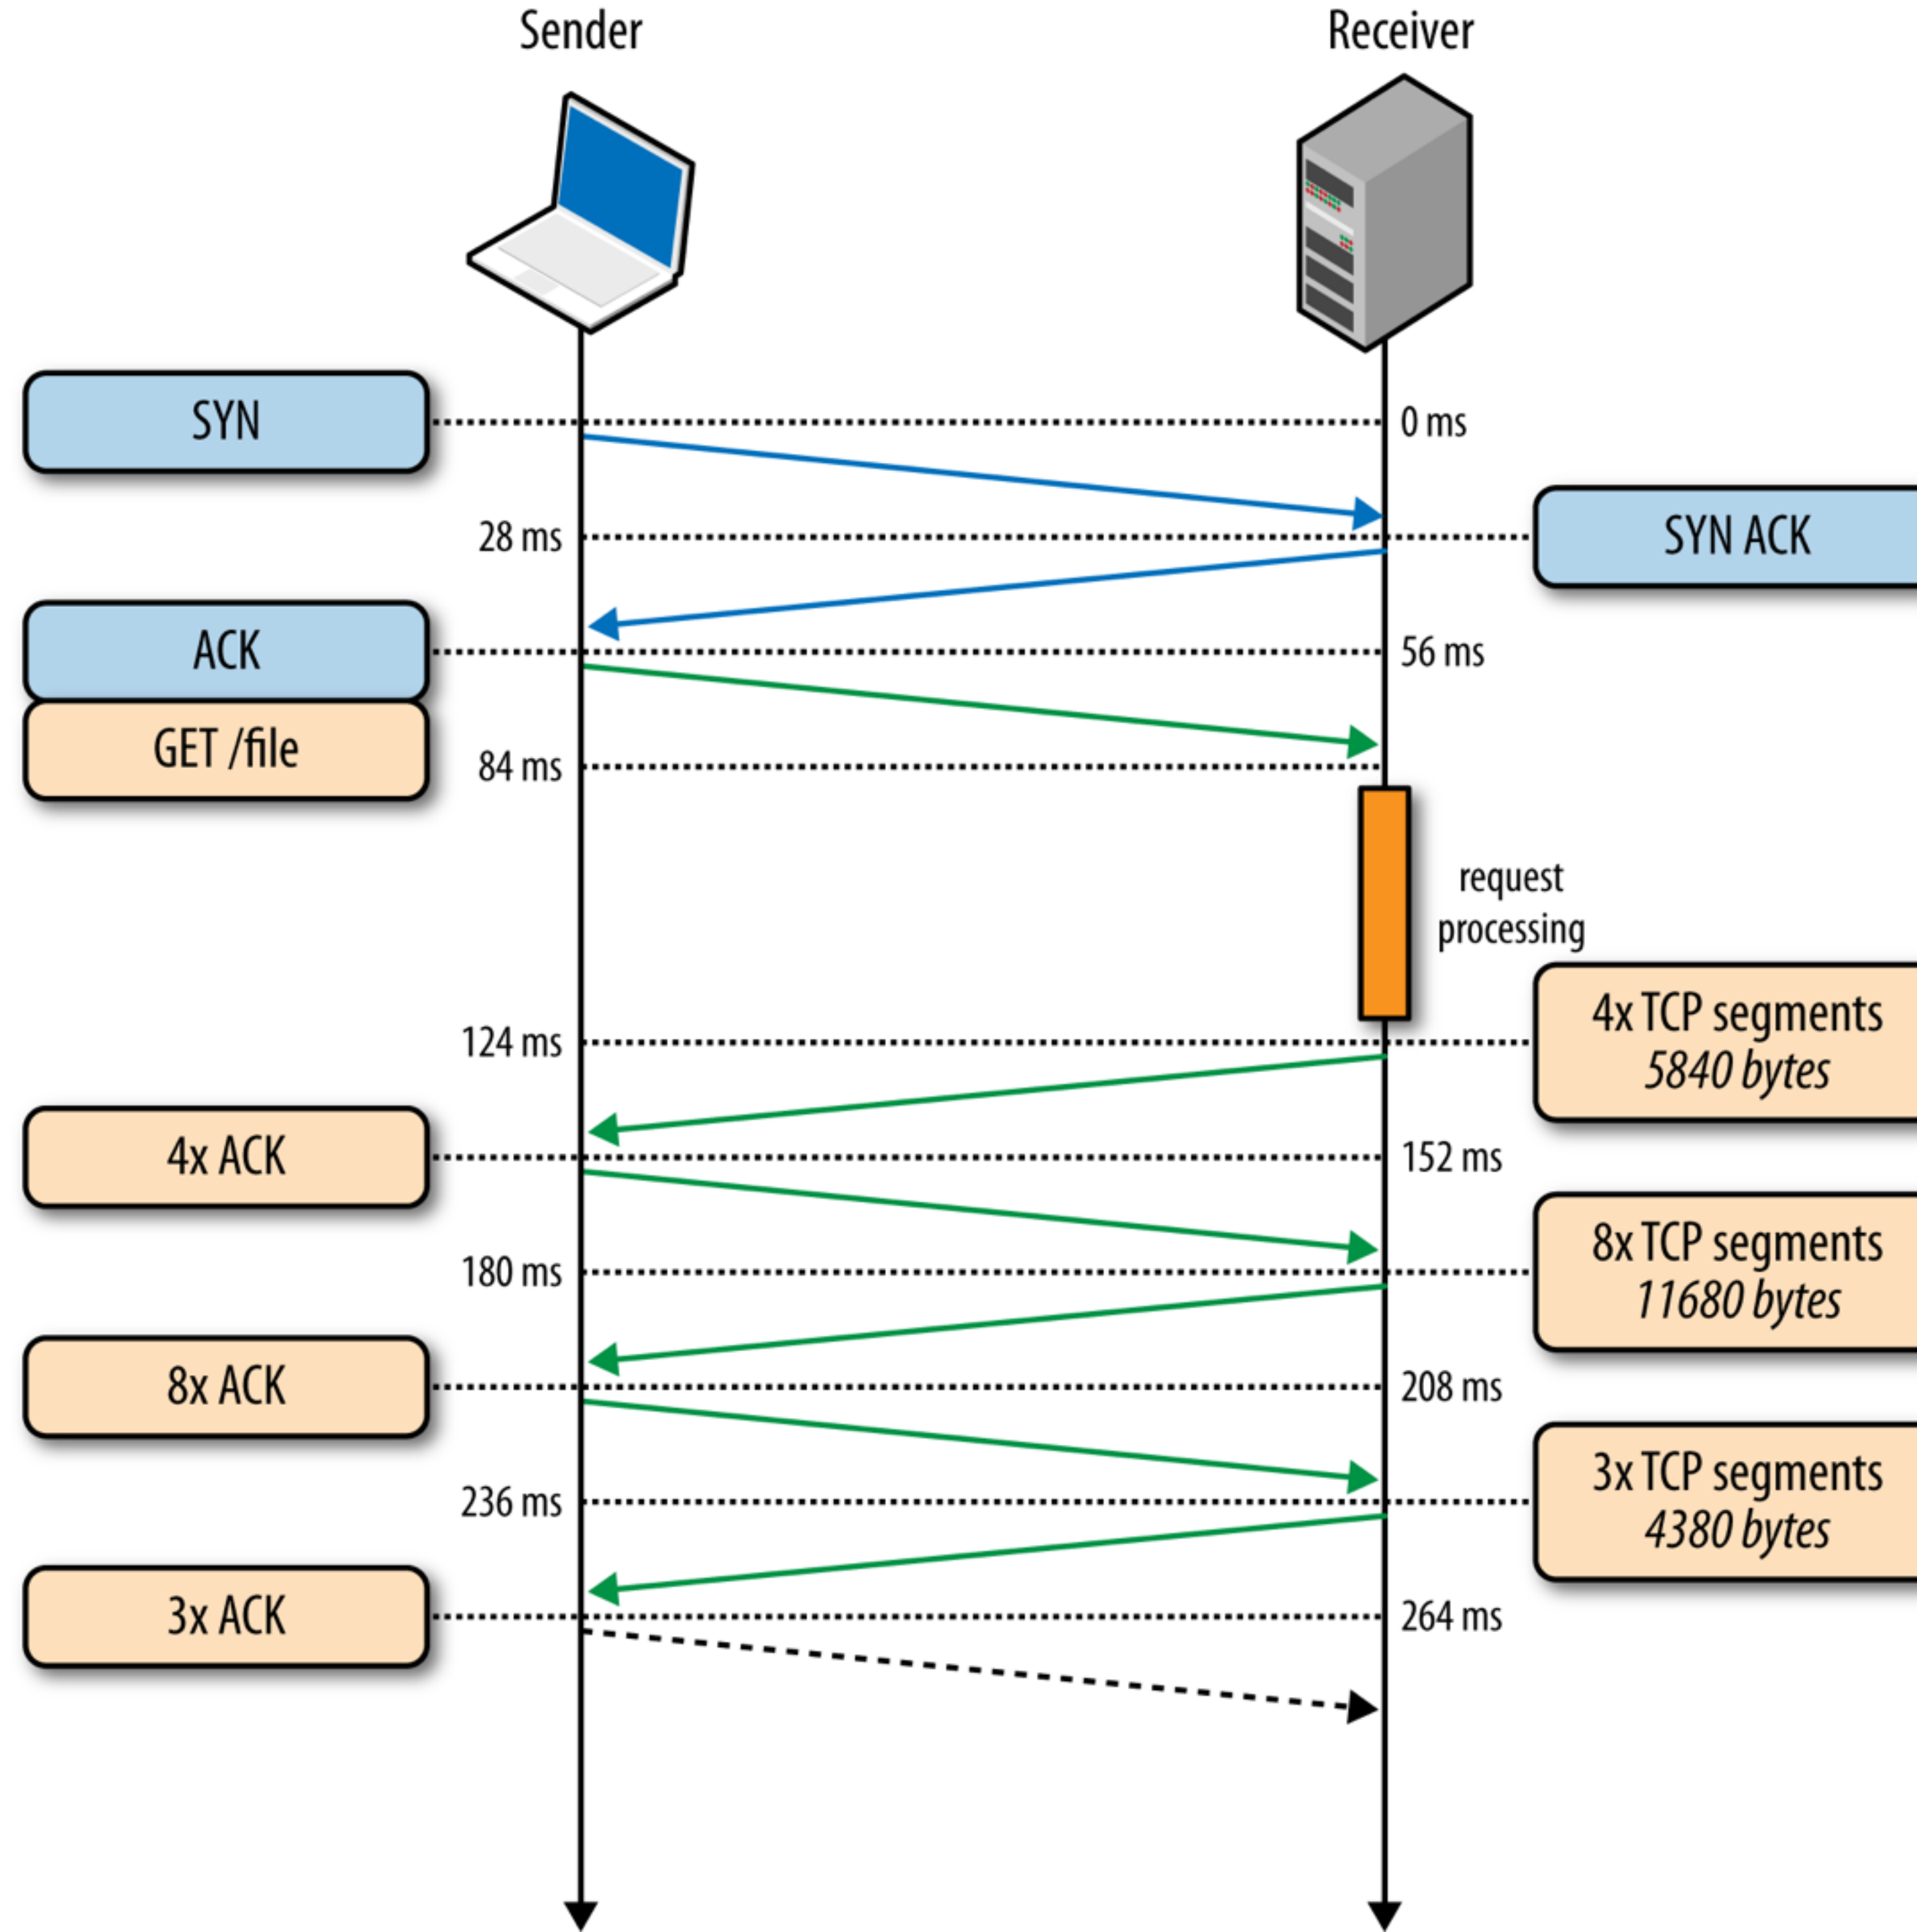


Actual RTT ~ 75ms

TCP SLOW START



STAY




Congestion Collapse


“This condition is stable. Once the saturation point has been reached, if the algorithm for selecting packets to be dropped is fair, the network will continue to operate in a degraded condition.”

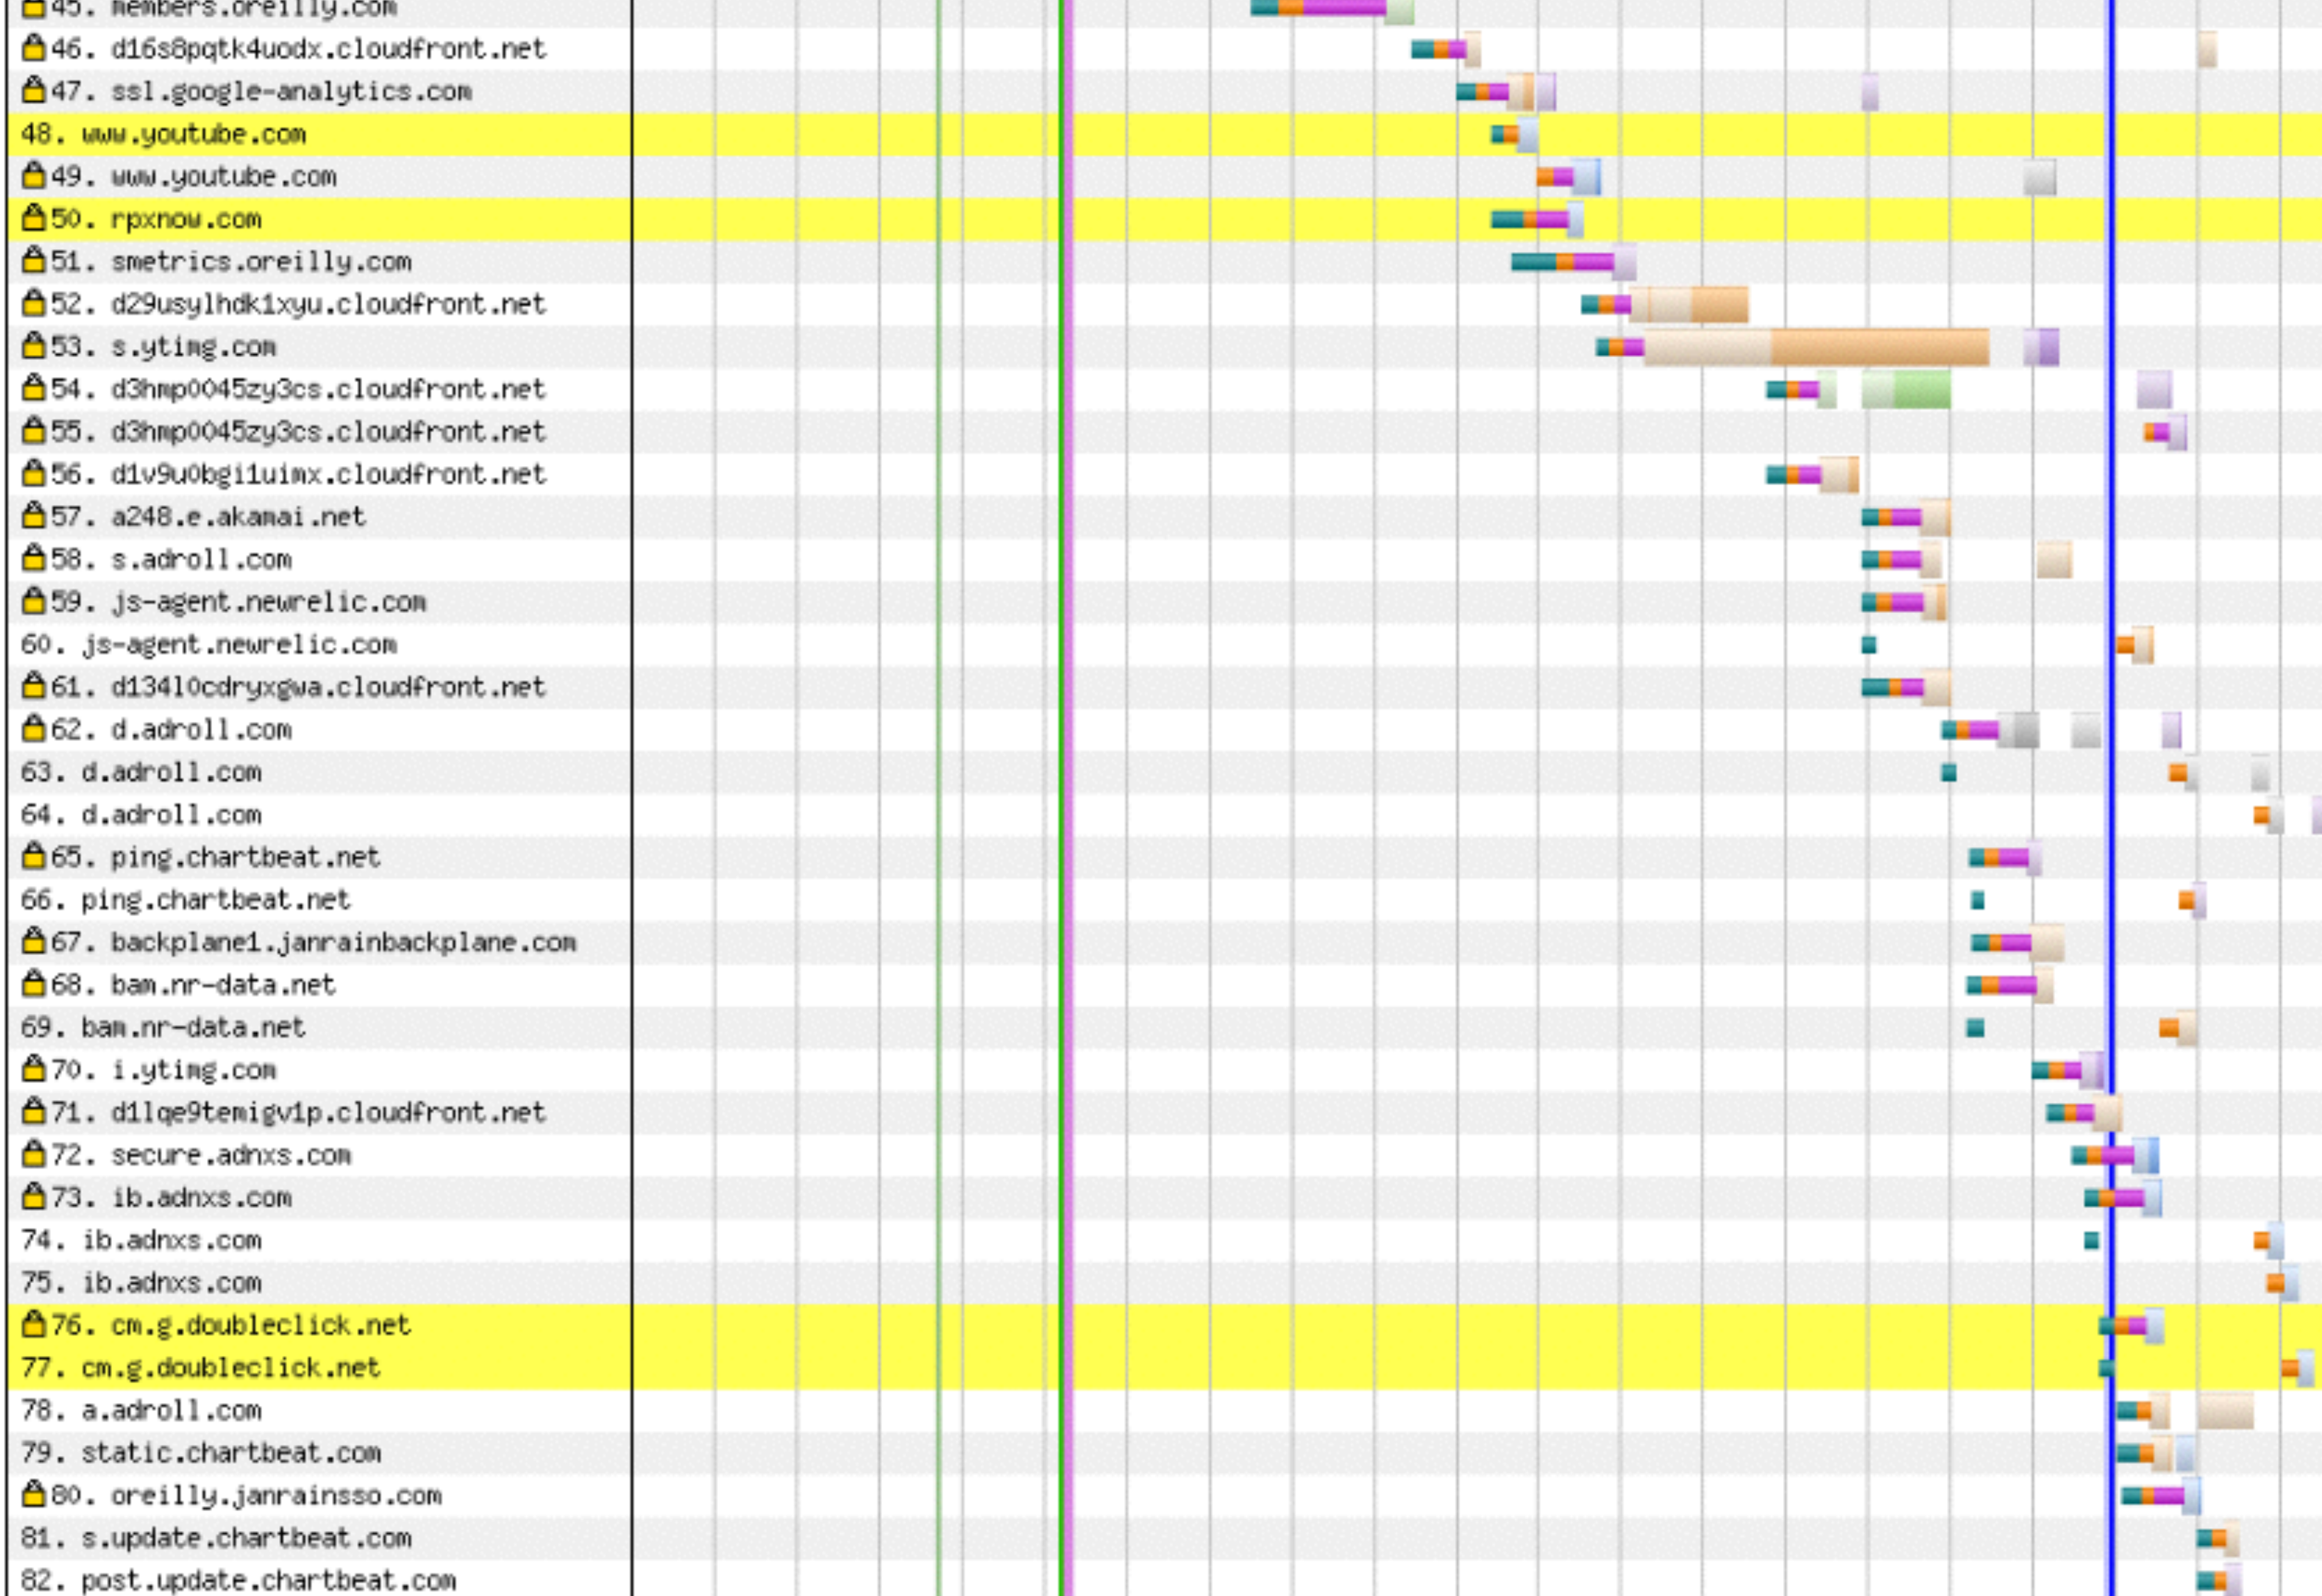
- RFC 896



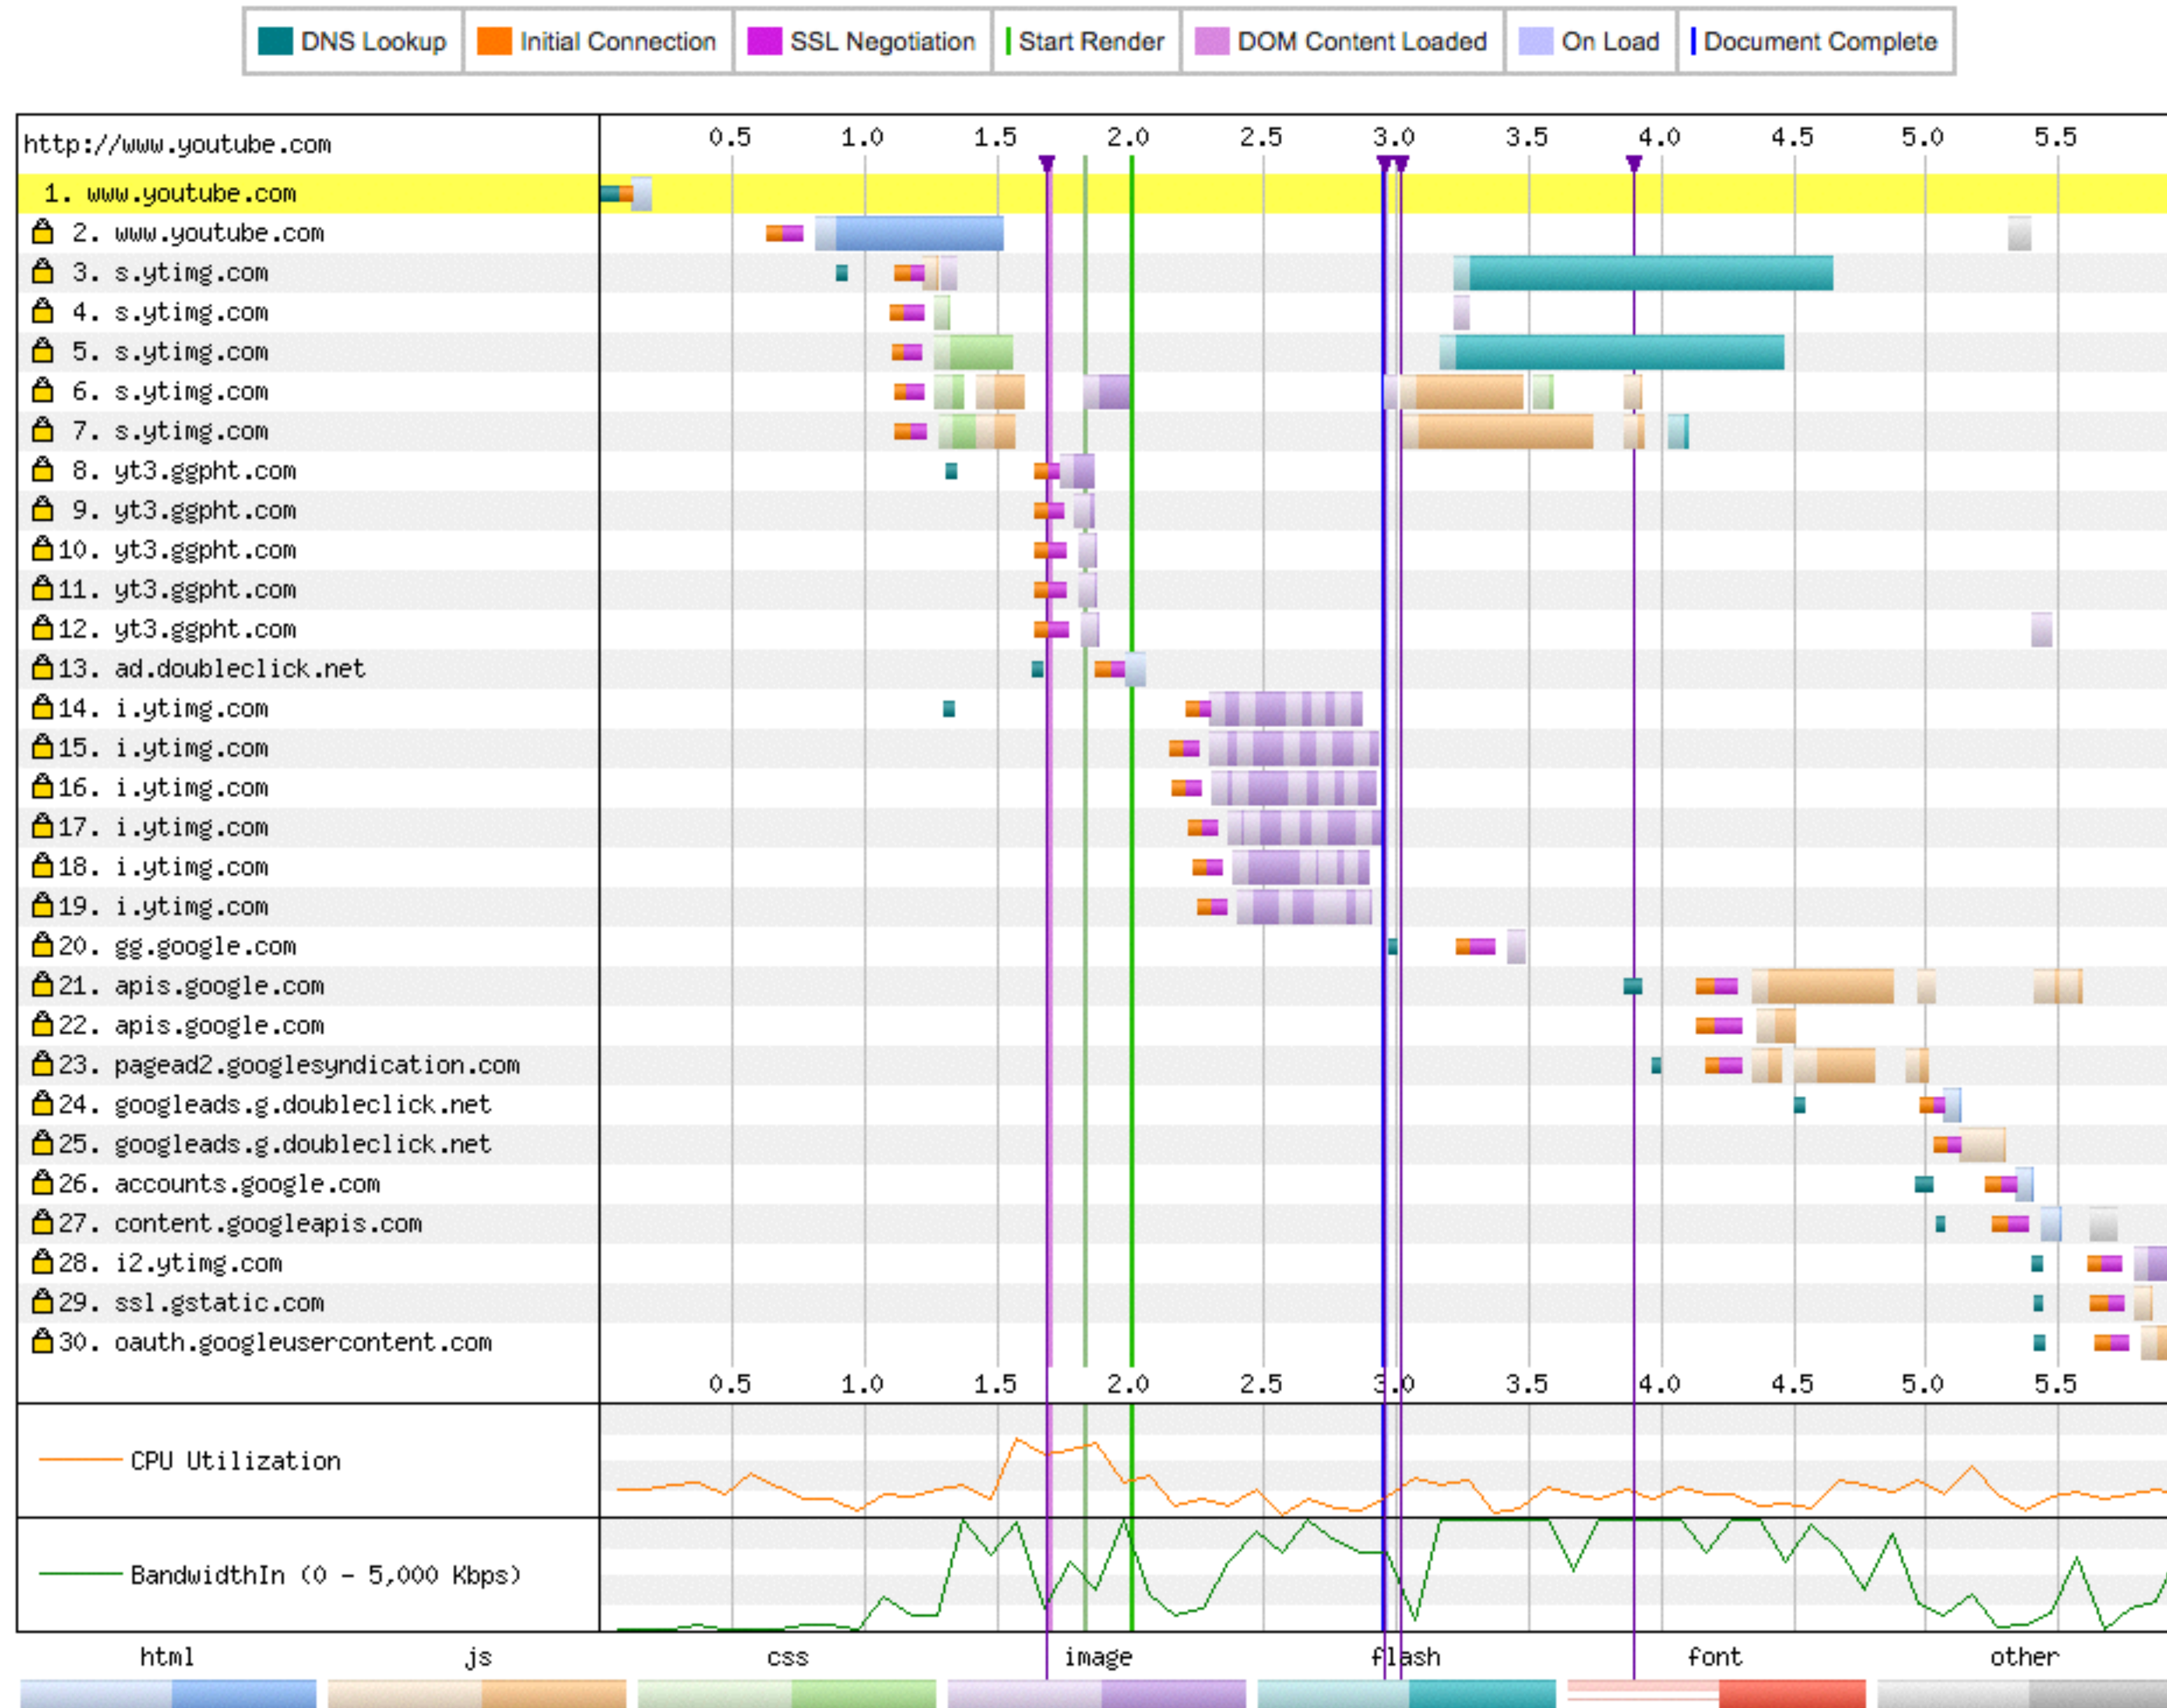


Connection View





Connection View




6. s.ytimg.com
7. s.ytimg.com
8. yt3.ggpht.com
9. yt3.ggpht.com
10. yt3.ggpht.com
11. yt3.ggpht.com
12. yt3.ggpht.com
13. ad.doubleclick.net

14. i.ytimg.com
15. i.ytimg.com
16. i.ytimg.com
17. i.ytimg.com
18. i.ytimg.com
19. i.ytimg.com

20. gg.google.com
21. apis.google.com
22. apis.google.com
23. pagead2.googleadsyndication.com
24. googleads.g.doubleclick.net
25. googleads.g.doubleclick.net
26. accounts.google.com
27. content.googleapis.com

6 connections



HTTP Head-of-line Blocking

A single slow response blocks
all requests behind it.

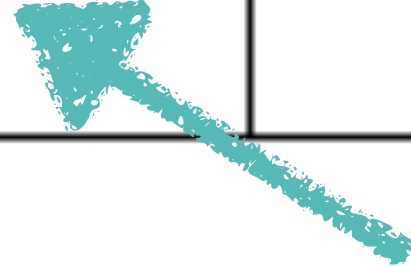


TCP/IP Network Model

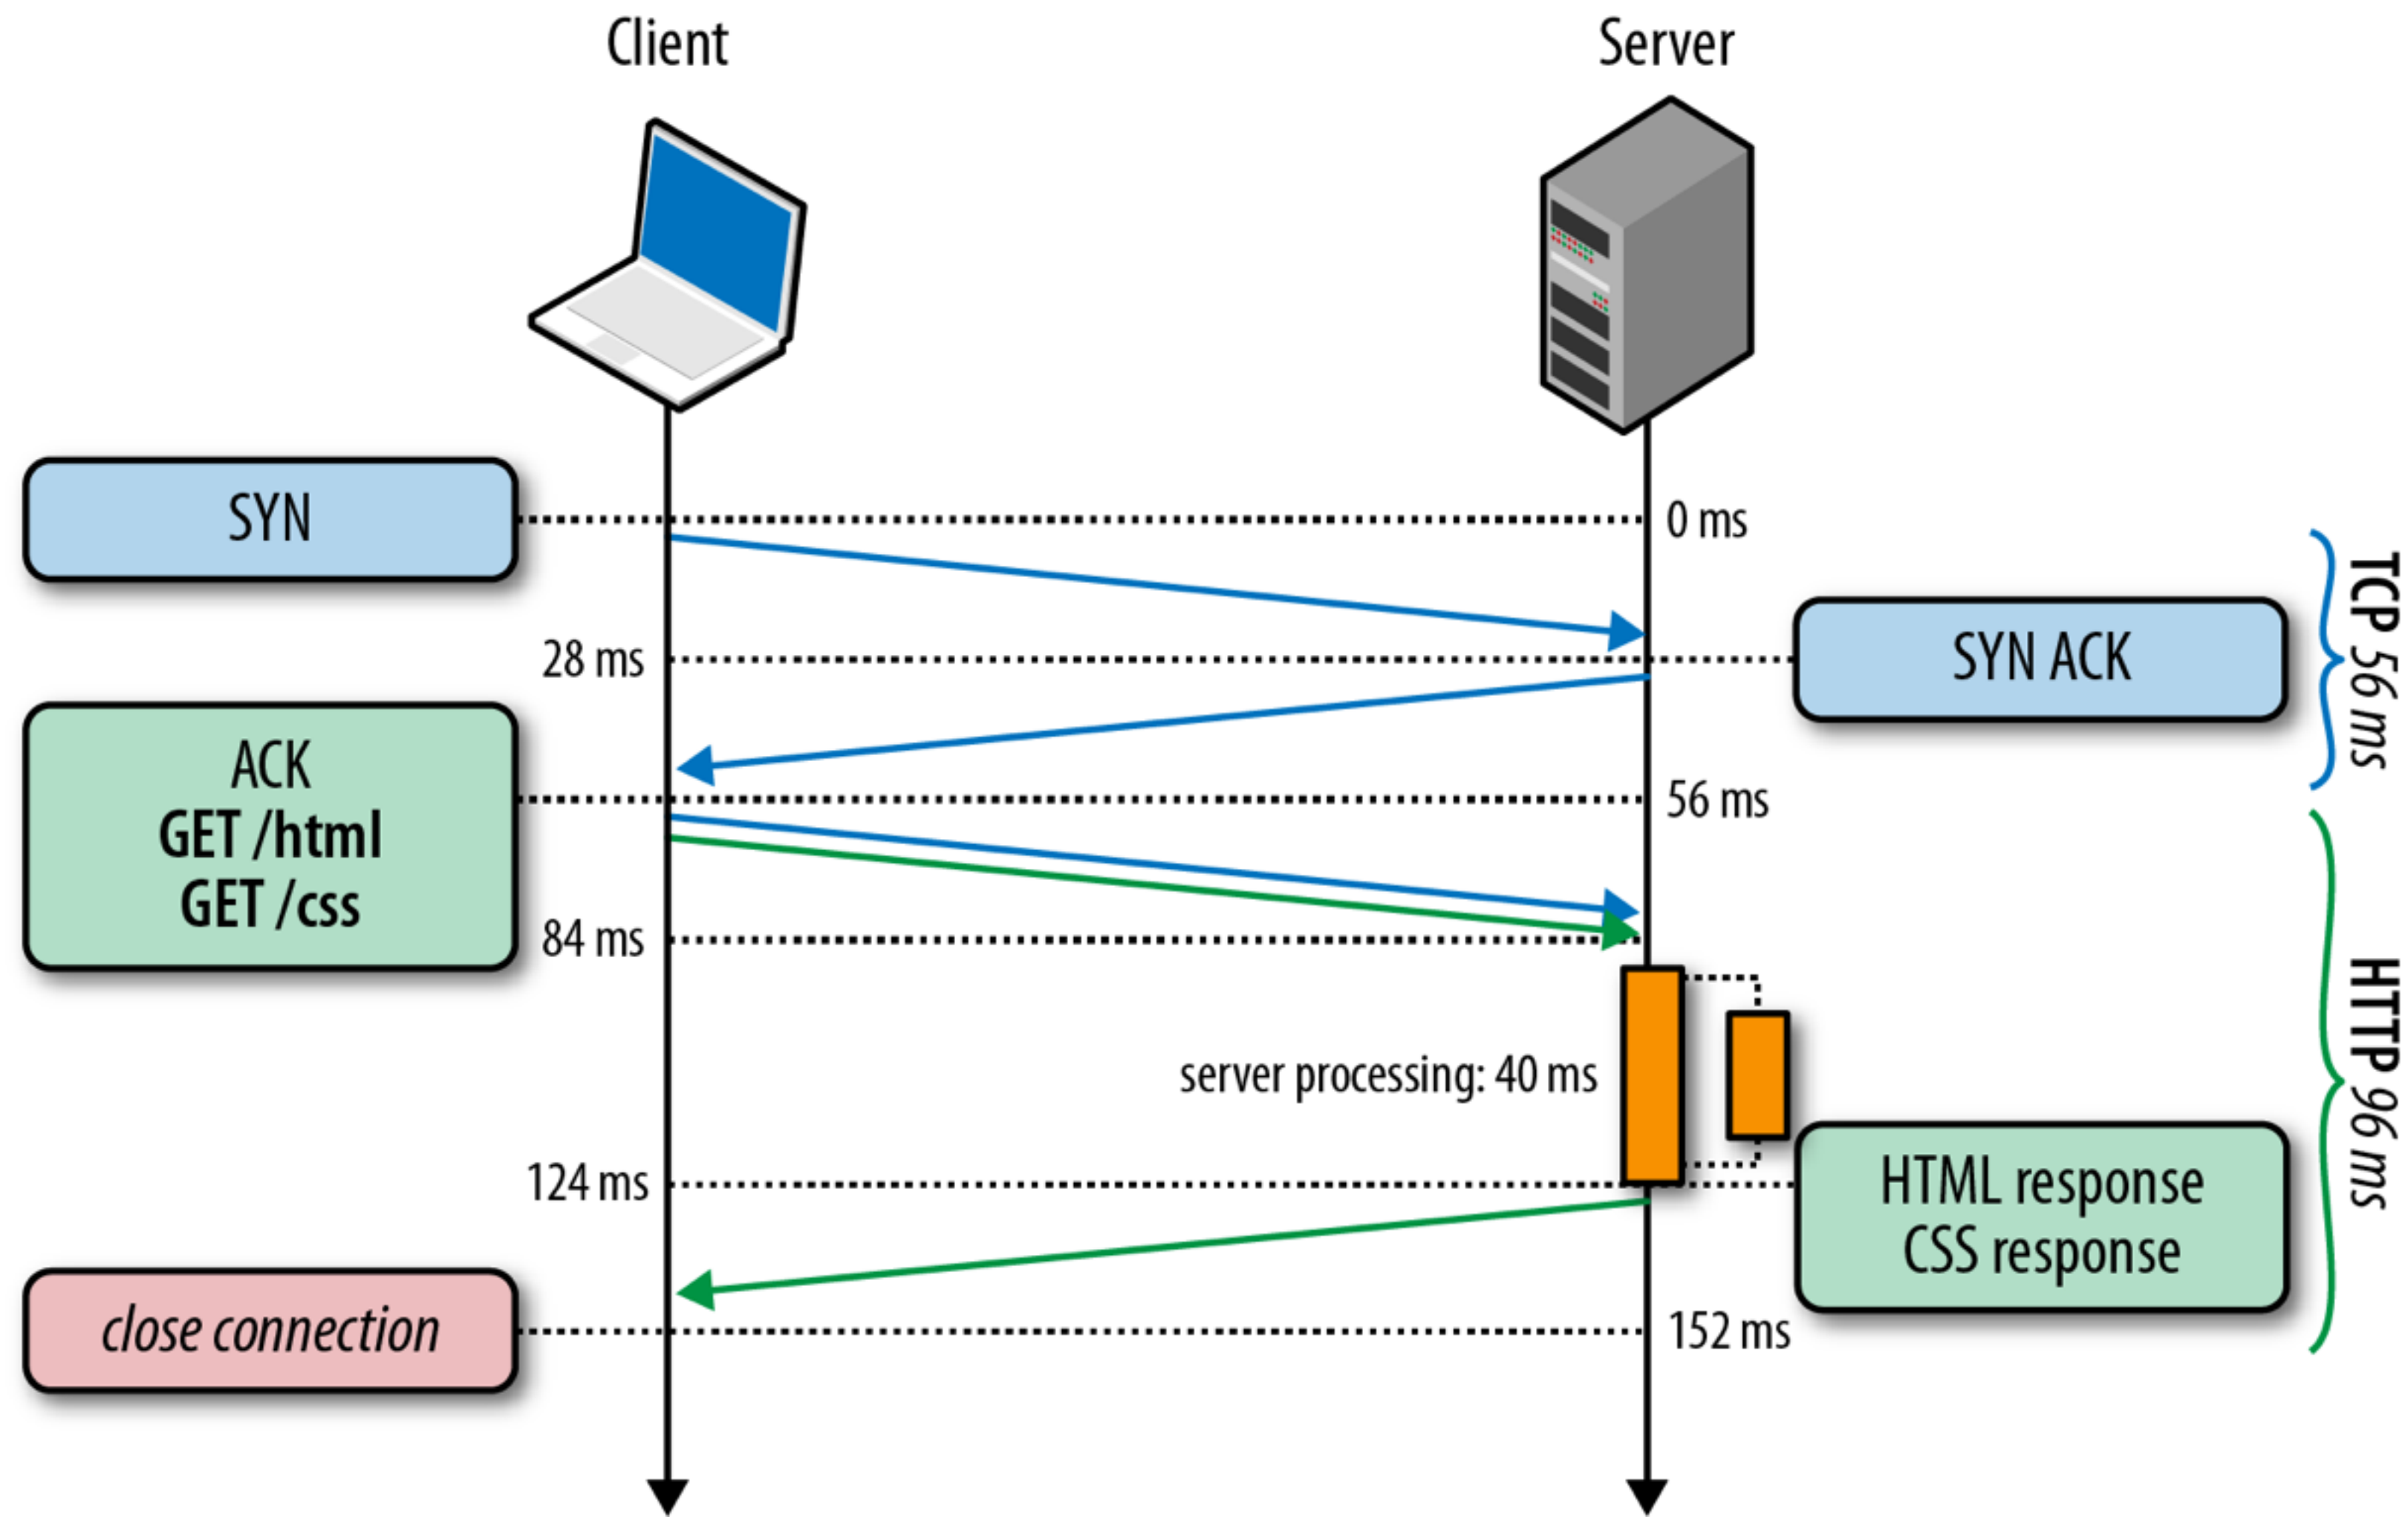
TCP/IP Layers

TCP/IP Protocols

Application Layer	HTTP	FTP	Telnet	SMTP	DNS
Transport Layer	TCP			UDP	
Network Layer	IP		ARP	ICMP	IGMP
Network Interface Layer	Ethernet		Token Ring		Other Link-Layer Protocols



HTTP Head-of-line Blocking



What About HTTP Pipelining?

“...if the web browser is the primary delivery vehicle for your web application, then we can't count on HTTP pipelining to help with performance.”


- High Performance Browser Networking

6. s.ytimg.com
7. s.ytimg.com
8. yt3.ggpht.com
9. yt3.ggpht.com
10. yt3.ggpht.com
11. yt3.ggpht.com
12. yt3.ggpht.com
13. ad.doubleclick.net

14. i.ytimg.com
15. i.ytimg.com
16. i.ytimg.com
17. i.ytimg.com
18. i.ytimg.com
19. i.ytimg.com

20. gg.google.com
21. apis.google.com
22. apis.google.com
23. pagead2.googleadsyndication.com
24. googleads.g.doubleclick.net
25. googleads.g.doubleclick.net
26. accounts.google.com
27. content.googleapis.com

6 connections



Index	File Name	Size (KB)	Transfer Time (ms)
18.	i.ytimg.com - mqdefault.jpg	237	237 ms
19.	i.ytimg.com - mqdefault.jpg	205	205 ms
20.	i.ytimg.com - mqdefault.jpg	189	189 ms
21.	i.ytimg.com - mqdefault.jpg	161	161 ms
22.	i.ytimg.com - hqdefault.jpg	351	351 ms
23.	i.ytimg.com - mqdefault.jpg	211	211 ms
24.	i.ytimg.com - mqdefault.jpg	181	181 ms
25.	i.ytimg.com - mqdefault.jpg	261	261 ms
26.	i.ytimg.com - mqdefault.jpg	177	177 ms
27.	i.ytimg.com - mqdefault.jpg	146	146 ms
28.	i.ytimg.com - mqdefault.jpg	136	136 ms
29.	i.ytimg.com - mqdefault.jpg	114	114 ms
30.	i.ytimg.com - mqdefault.jpg	123	123 ms
31.	i.ytimg.com - mqdefault.jpg	95	95 ms
32.	i.ytimg.com - mqdefault.jpg	108	108 ms
33.	i.ytimg.com - mqdefault.jpg	78	78 ms
34.	i.ytimg.com - mqdefault.jpg	87	87 ms
35.	i.ytimg.com - mqdefault.jpg	166	166 ms
36.	i.ytimg.com - 1.jpg	62	62 ms
37.	i.ytimg.com - mqdefault.jpg	139	139 ms
38.	i.ytimg.com - mqdefault.jpg	100	100 ms
39.	i.ytimg.com - mqdefault.jpg	90	90 ms
40.	i.ytimg.com - mqdefault.jpg	92	92 ms
41.	i.ytimg.com - mqdefault.jpg	100	100 ms
42.	i.ytimg.com - mqdefault.jpg	91	91 ms
43.	i.ytimg.com - mqdefault.jpg	115	115 ms
44.	i.ytimg.com - mqdefault.jpg	88	88 ms
45.	i.ytimg.com - 1.jpg	58	58 ms
46.	i.ytimg.com - mqdefault.jpg	87	87 ms

Workarounds and Downsides

Workarounds and Downsides

- Multiple TCP Connections

Workarounds and Downsides

- Multiple TCP Connections
 - Extra resource utilization

Workarounds and Downsides

- Multiple TCP Connections
 - Extra resource utilization
 - Bandwidth competition

Workarounds and Downsides

- Multiple TCP Connections
 - Extra resource utilization
 - Bandwidth competition
- Domain Sharding

Workarounds and Downsides

- Multiple TCP Connections
 - Extra resource utilization
 - Bandwidth competition
- Domain Sharding
 - Extra DNS Lookup(s)

Workarounds and Downsides

- Multiple TCP Connections
 - Extra resource utilization
 - Bandwidth competition
- Domain Sharding
 - Extra DNS Lookup(s)
 - Implementation complexity

H/2 Multiplexing

H/2 Multiplexing

From HPBN

#fluentconf

@jonathanklein

jkle.in/http2

O'REILLY®
Fluent

H/2 Multiplexing

- Interleave multiple requests/responses on the same connection

From HPBN

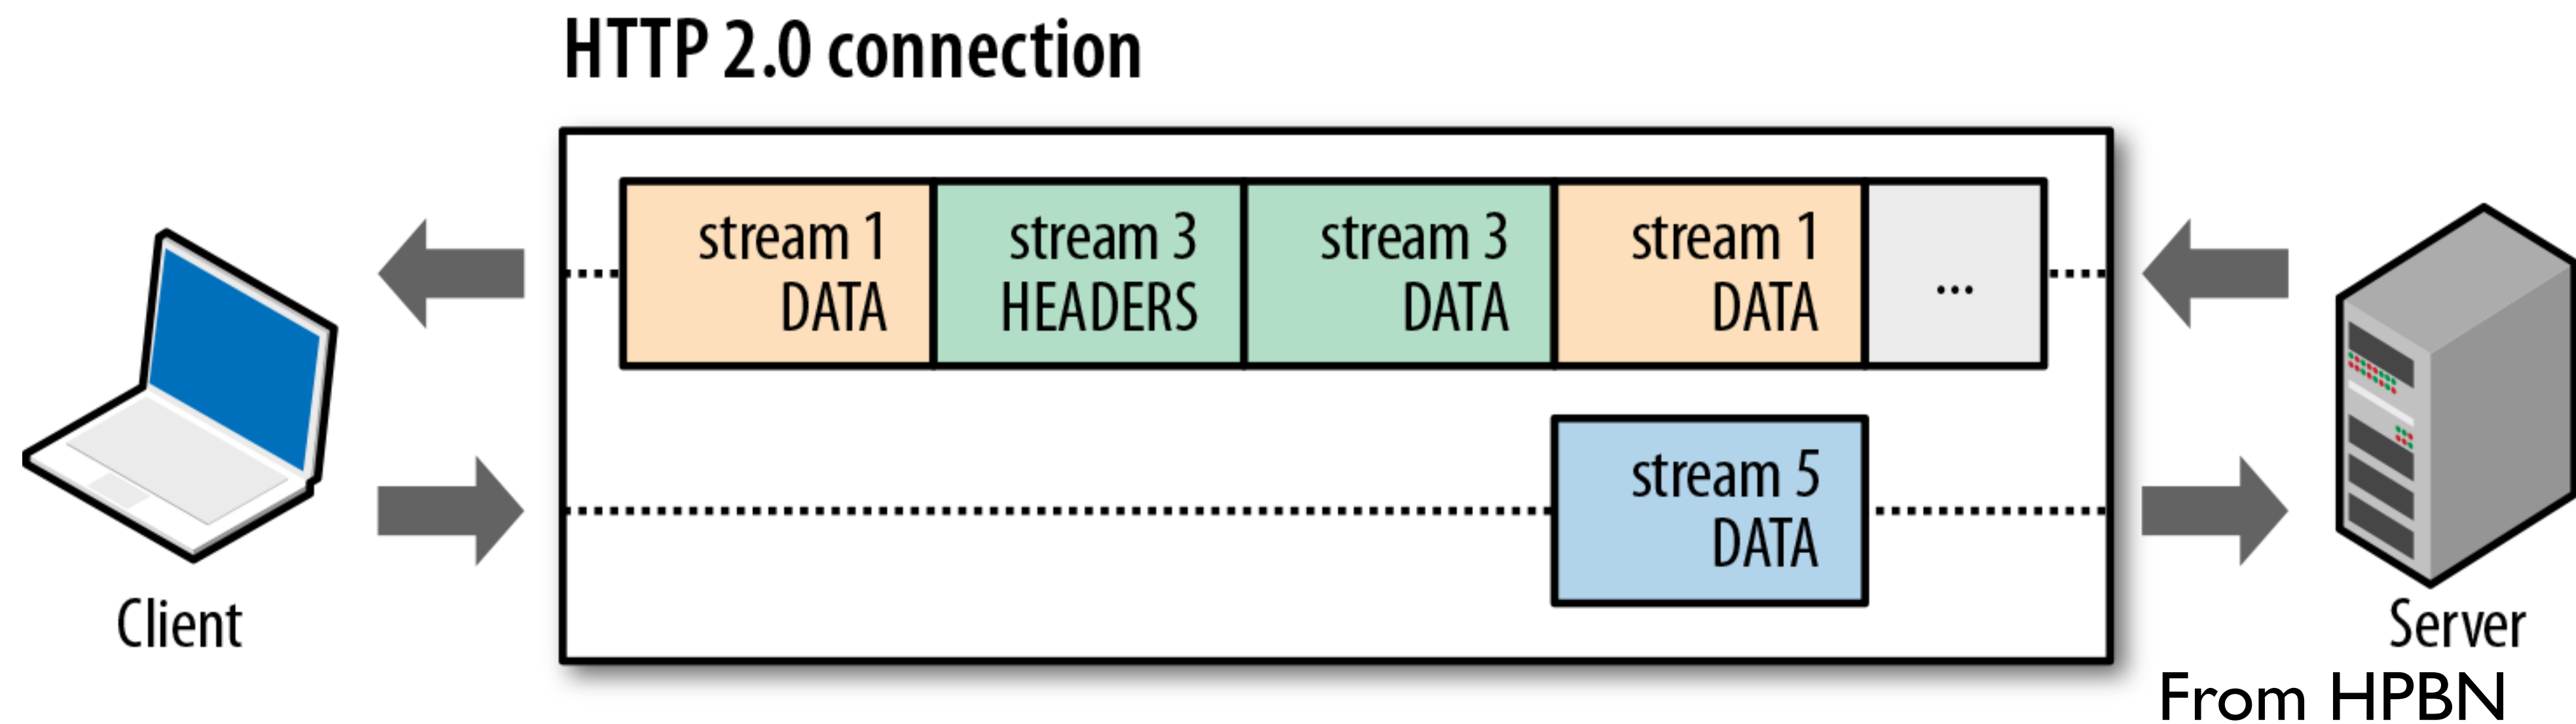
H/2 Multiplexing

- Interleave multiple requests/responses on the same connection
- No more head-of-line blocking!

From HPBN

H/2 Multiplexing

- Interleave multiple requests/responses on the same connection
- No more head-of-line blocking!



Forget about...

Forget about...

- Combining CSS and JS files

Forget about...

- Combining CSS and JS files
- Inlining small CSS/JS

Forget about...

- Combining CSS and JS files
- Inlining small CSS/JS
- Using image sprites

Forget about...

- Combining CSS and JS files
- Inlining small CSS/JS
- Using image sprites
- Using DataURIs

Forget about...

- Combining CSS and JS files
- Inlining small CSS/JS
- Using image sprites
- Using DataURIs
- Domain sharding

What About
HTTP Headers?

Request Headers

```
GET /ajax/libs/jquery/1.7.2/jquery.min.js HTTP/1.1
Host: cdn.oreillystatic.com
Connection: keep-alive
Pragma: no-cache
Cache-Control: no-cache
Accept: */*
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_2)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/41.0.2272.118 Safari/
537.36
Referer: http://fluentconf.com/javascript-html-2015
Accept-Encoding: gzip, deflate, sdch
Accept-Language: en-US,en;q=0.8
```

Response Headers

HTTP/1.1 200 OK

Server: Apache

Vary: Accept-Encoding

Last-Modified: Thu, 02 Aug 2012 22:41:31 GMT

Accept-Ranges: bytes

Content-Encoding: gzip

Content-Length: 33622

Content-Type: application/javascript

Cache-Control: max-age=2236992

Expires: Thu, 30 Apr 2015 17:21:11 GMT

Date: Sat, 04 Apr 2015 19:57:59 GMT

Connection: keep-alive

Request Headers

```
GET /en/assets/1/eventprovider/1/lazyload_min_v3.js HTTP/1.1
Host: cdn.oreillystatic.com
Connection: keep-alive
Pragma: no-cache
Cache-Control: no-cache
Accept: */*
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_2)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/41.0.2272.118 Safari/
537.36
Referer: http://fluentconf.com/javascript-html-2015
Accept-Encoding: gzip, deflate, sdch
Accept-Language: en-US,en;q=0.8
```


Request Headers

```
GET /en/assets/1/eventprovider/1/lazyload_min_v3.js HTTP/1.1
Host: cdn.oreillystatic.com
Connection: keep-alive
Pragma: no-cache
Cache-Control: no-cache
Accept: */*
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_2)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/41.0.2272.118 Safari/
537.36
Referer: http://fluentconf.com/javascript-html-2015
Accept-Encoding: gzip, deflate, sdch
Accept-Language: en-US,en;q=0.8
```

HPACK

Request headers

:method	GET
:scheme	https
:host	example.com
:path	/resource
user-agent	Mozilla/5.0 ...
custom-hdr	some-value



Static table

1	:authority	
2	:method	GET
...
51	referer	
...
62	user-agent	Mozilla/5.0 ...
63	:host	example.com
...

Dynamic table




Encoded headers

2	
7	
63	
19	Huffman("/resource")
62	
	Huffman("custom-hdr")
	Huffman("some-value")

Why HPACK and not zlib?

- CRIME attack
 - zlib + encryption = session hijacking
- HPACK: New, header specific compression algorithm



<https://www.mnot.net/talks/http2-expectations/#/tsv>

Request Prioritization and Server Push

Server Push

Server Push

- Resources can be...

Server Push

- Resources can be...
 - Cached by client

Server Push

- Resources can be...
 - Cached by client
 - Declined by client

Server Push

- Resources can be...
 - Cached by client
 - Declined by client
 - Multiplexed

Server Push

- Resources can be...
 - Cached by client
 - Declined by client
 - Multiplexed
 - Prioritized by server

Server Push

- Resources can be...
 - Cached by client
 - Declined by client
 - Multiplexed
 - Prioritized by server
- Must obey same-origin policy

**IS IT REALLY
FASTER?**



Data From Google

	Google News	Google Sites	Google Drive	Google Maps
Median	-43%	-27%	-23%	-24%
5th percentile (fast connections)	-32%	-30%	-15%	-20%
95th percentile (slow connections)	-44%	-33%	-36%	-28%

time from first request byte to the onload event



Cart (empty)

Welcome | Login

[Home](#) ▾[iPods](#)[Accessories](#)

TAGS

ipod apple [superdrive](#) [ipod](#)
[touch](#) [nano](#) [shuffle](#)

CATEGORIES

- ▶ [iPods](#)
- ▶ [Accessories](#)
- ▶ [Laptops](#)

MANUFACTURERS

- ▶ [Apple Computer, Inc](#)
- ▶ [Shure Incorporated](#)

[All manufacturers](#) ▾

SUPPLIERS

- ▶ [AppleStore](#)
- ▶ [Shure Online Store](#)

[All suppliers](#) ▾

INFORMATION

- ▶ [Delivery](#)
- ▶ [Legal Notice](#)



FEATURED PRODUCTS

**iPod Nano**

New design. New features. Now in 8GB and 16GB. iPod nano...

**iPod shuffle**

iPod shuffle, the world's most wearable music player, now...

**MacBook**

MacBook makes it easy to hit the road thanks to its tough...

**iPod touch**

Revolutionary Multi-Touch interface 3.5-inch widescreen...


SPECIALS



iPod Nano
~~\$166.39~~
\$158.07

-5%[» All specials](#)

OUR STORE(S)!

[» Discover our store\(s\)!](#)


CONTACT US

Our support hotline is available 24/7.

 **0123-456-789****CONTACT OUR EXPERT
SUPPORT TEAM!**


H/1

Connection View



H/2

Connection View



Key Stats

	H/1	H/2	Difference
Bytes	387KB	387KB	-
Start Render	1019ms	972ms	-5%
Speed Index	1278ms	1108ms	-14%

What About HTTPS?

Not required by spec, but
Firefox and Chrome will require it

istlsfastyet.com

Criticism

Criticism

Criticism

“It was rushed, we wanted more”

Criticism

“It was rushed, we wanted more”

- Something good now > something perfect in 5 years

Criticism

“It was rushed, we wanted more”

- Something good now > something perfect in 5 years

“Why does it require TLS?”

Criticism

“It was rushed, we wanted more”

- Something good now > something perfect in 5 years

“Why does it require TLS?”

- The spec doesn't, some browsers do

HATERS GONNA HATE



Implementation

Webservers

	Session identifiers	Session tickets	OCSP stapling	Dynamic record sizing	ALPN / NPN	Forward secrecy	SPDY & HTTP/2
<u>Apache</u>	<u>yes</u>	<u>yes</u>	<u>yes</u>	no	<u>no*</u>	<u>yes</u>	<u>mod_spdy</u>
<u>ATS</u>	<u>yes</u>	<u>yes</u>	<u>yes</u>	<u>dynamic</u>	<u>yes</u>	<u>yes</u>	<u>spdy/3.1</u>
<u>bud</u>	no	<u>yes</u>	<u>yes</u>	static	<u>yes</u>	<u>yes</u>	no
<u>F5 BIG-IP</u>	<u>yes</u>	<u>yes</u>	<u>yes</u>	no	<u>yes</u>	<u>yes</u>	<u>spdy/3.1, http/2.0</u>
<u>HAProxy</u>	<u>yes</u>	<u>yes</u>	<u>yes</u>	<u>dynamic</u>	<u>yes</u>	<u>yes</u>	no
<u>IIS</u>	<u>yes</u>	<u>yes</u>	<u>yes</u>	no	<u>yes</u>	<u>yes</u>	no
<u>NetScaler</u>	<u>yes</u>	no	no	no	<u>yes</u>	<u>yes</u>	<u>spdy/3.0</u>
<u>NGINX</u>	<u>yes</u>	<u>yes</u>	<u>yes</u>	static	<u>yes</u>	<u>yes</u>	<u>spdy/3.1</u>
<u>node.js</u>	<u>yes</u>	no	no	no	<u>yes</u>	no	<u>spdy/3.1</u>

<https://istlsfastyet.com/>

CDN Support

	Session identifiers	Session tickets	OCSF stapling	Dynamic record sizing	ALPN / NPN	Forward secrecy	SPDY & HTTP/2
Akamai	yes	yes	yes	configurable (static)	yes	yes	spdy/3 spdy3.1 (opt-in)
CloudFlare	yes	yes	yes	4KB (static)	yes	yes	spdy/3.1
AWS ELB	yes	yes	no	no	no	yes	no
AWS CloudFront	no	yes	yes	no	yes	yes	no
EdgeCast	no	yes	yes	no	no	yes	no
Fastly	yes	yes	yes	no	no	yes	no
Google App Engine	yes	yes	no	dynamic	yes	yes	spdy/3.1
Heroku	yes	yes	no	no	no	yes	no
Instart Logic	yes	yes	no	configurable (static)	yes	yes	spdy/3.1
Limelight	yes	yes	no	no	no	yes	no
MaxCDN	yes	yes	no	no	yes	no	spdy/3.1

<https://istlsfastyet.com/>

Handling Both H/1 and H/2 Clients

Okay to Keep	Stop
Combining CSS/JS	Domain Sharding
Spriting	Inlining Content

Global

$$72.15\% + 7.98\% = 80.13\%$$

Wrapping Up

Wrapping Up

- H/2 is here

Wrapping Up

- H/2 is here
- H/2 is fast

Wrapping Up

- H/2 is here
- H/2 is fast
- H/2 is supported

Wrapping Up

- H/2 is here
- H/2 is fast
- H/2 is supported

... go implement it!

Thanks!

<http://jkle.in/http2>

@jonathanklein