

# Mind the Delay: The Adverse Effects of Delay-Based TCP on HTTP

Neil Agarwal\*, Matteo Varvello<sup>^</sup>, Andrius Aucinas<sup>¶</sup>, Fabián Bustamante<sup>†</sup>, Ravi Netravali\*

\*UCLA    <sup>^</sup>Nokia, Bell Labs    <sup>¶</sup>Brave Software    <sup>†</sup>Northwestern University

Application



Transport



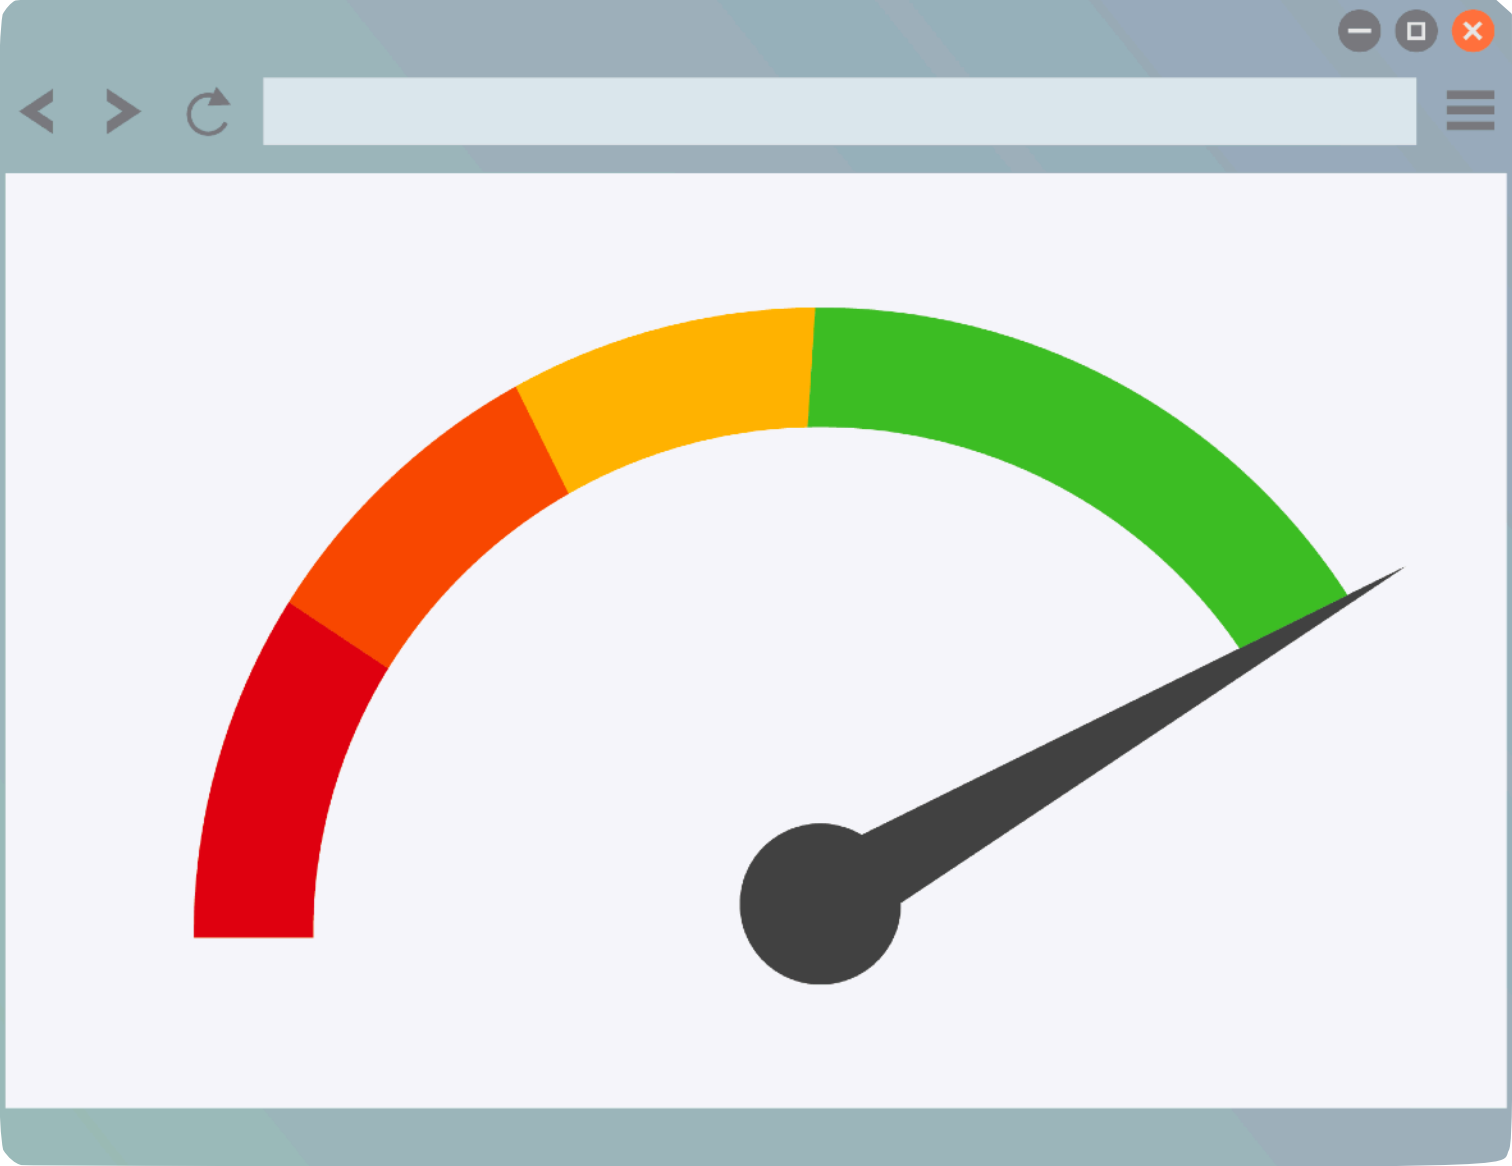
Network



Data Link



Physical



Web Performance


# Application

HTTP

HTTP/1.0, HTTP/1.1

HTTP/2, HTTP/3

Multiple connections



Single multiplexed connection



# Transport

TCP Congestion Control

Tahoe

Vegas 🕒

Illinois

Cubic w/ HyStart 🕒

BBR 🕒

Reno

Cubic

YeAH 🕒

New Reno

Time  
(Not to scale)


*A shift from loss-based -> delay-based* 🕒

# Application

# Transport

**Multiple**  
connections

**Loss-based**  
Congestion  
Control



**Single**  
multiplexed  
connection

**Delay-based**  
Congestion  
Control



Effect on Web Performance???

de facto approach:  
**use HTTP/2** 

# Synthetic Experiments

artificially **trigger delay-based behavior** by gradually increasing network latency



for pages with objects of at least 500 KB, **H1** outperformed **H2** on the order of **seconds**

BBR

Cubic HyStart


YeAH

# What's Happening?

**Packet Delay**  
**NOT** Caused by  
Network Congestion



**Underestimated**  
Network Capacity



HTTP/1.1



HTTP/2



**H1** can use its **multiple connections** to **better utilize** the **network** than **H2**.

# Delay-based CC Underestimate *cwnd*?

**Studies** have shown **delay-based** algorithms often **underestimate** the optimal congestion window in **mobile, jittery** networks.



@ **12** locations.

Loaded synthetic web page on **real**  network



congestion control

**Result:**

**H1** is **faster** than **H2** for **75%** of **runs**

Potential for **underestimating** the optimal **congestion window** in mobile networks



**H1** outperforming **H2**

**Next step?** Run experiments on **real web pages**.




# Evolving the Web Performance Toolbox



1. congestion control algorithm fingerprinting

2. fine-grained control over HTTP selection



# CC Fingerprinting

existing tool: **Gordon**

**POST**-loss  
behavior



small objects



cubic hystart



our tool: **Igor**

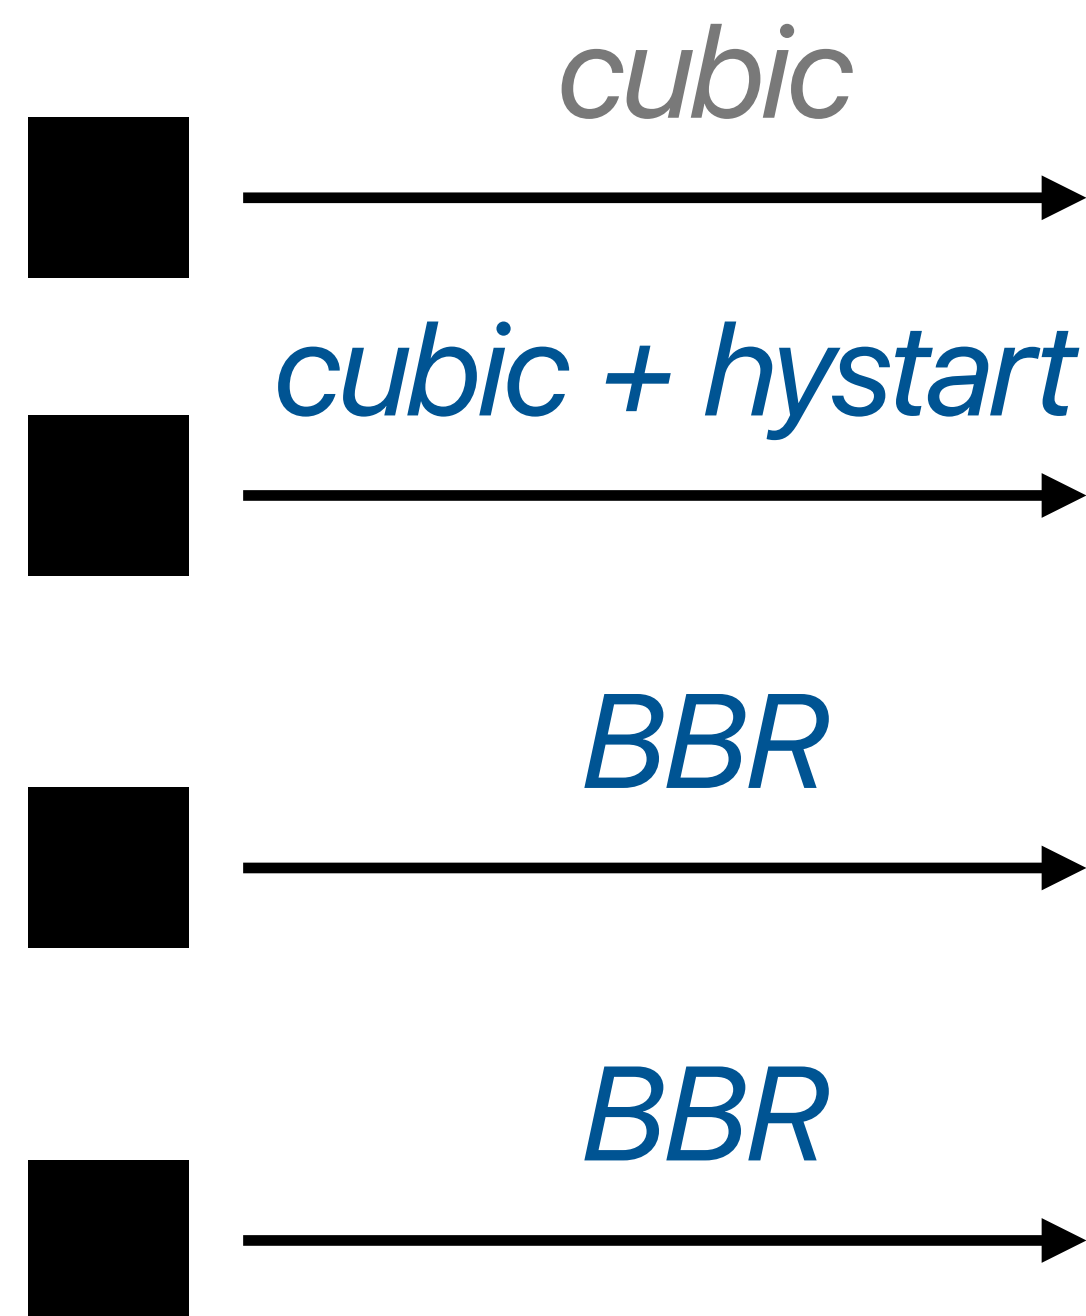
**PRE**-loss  
behavior




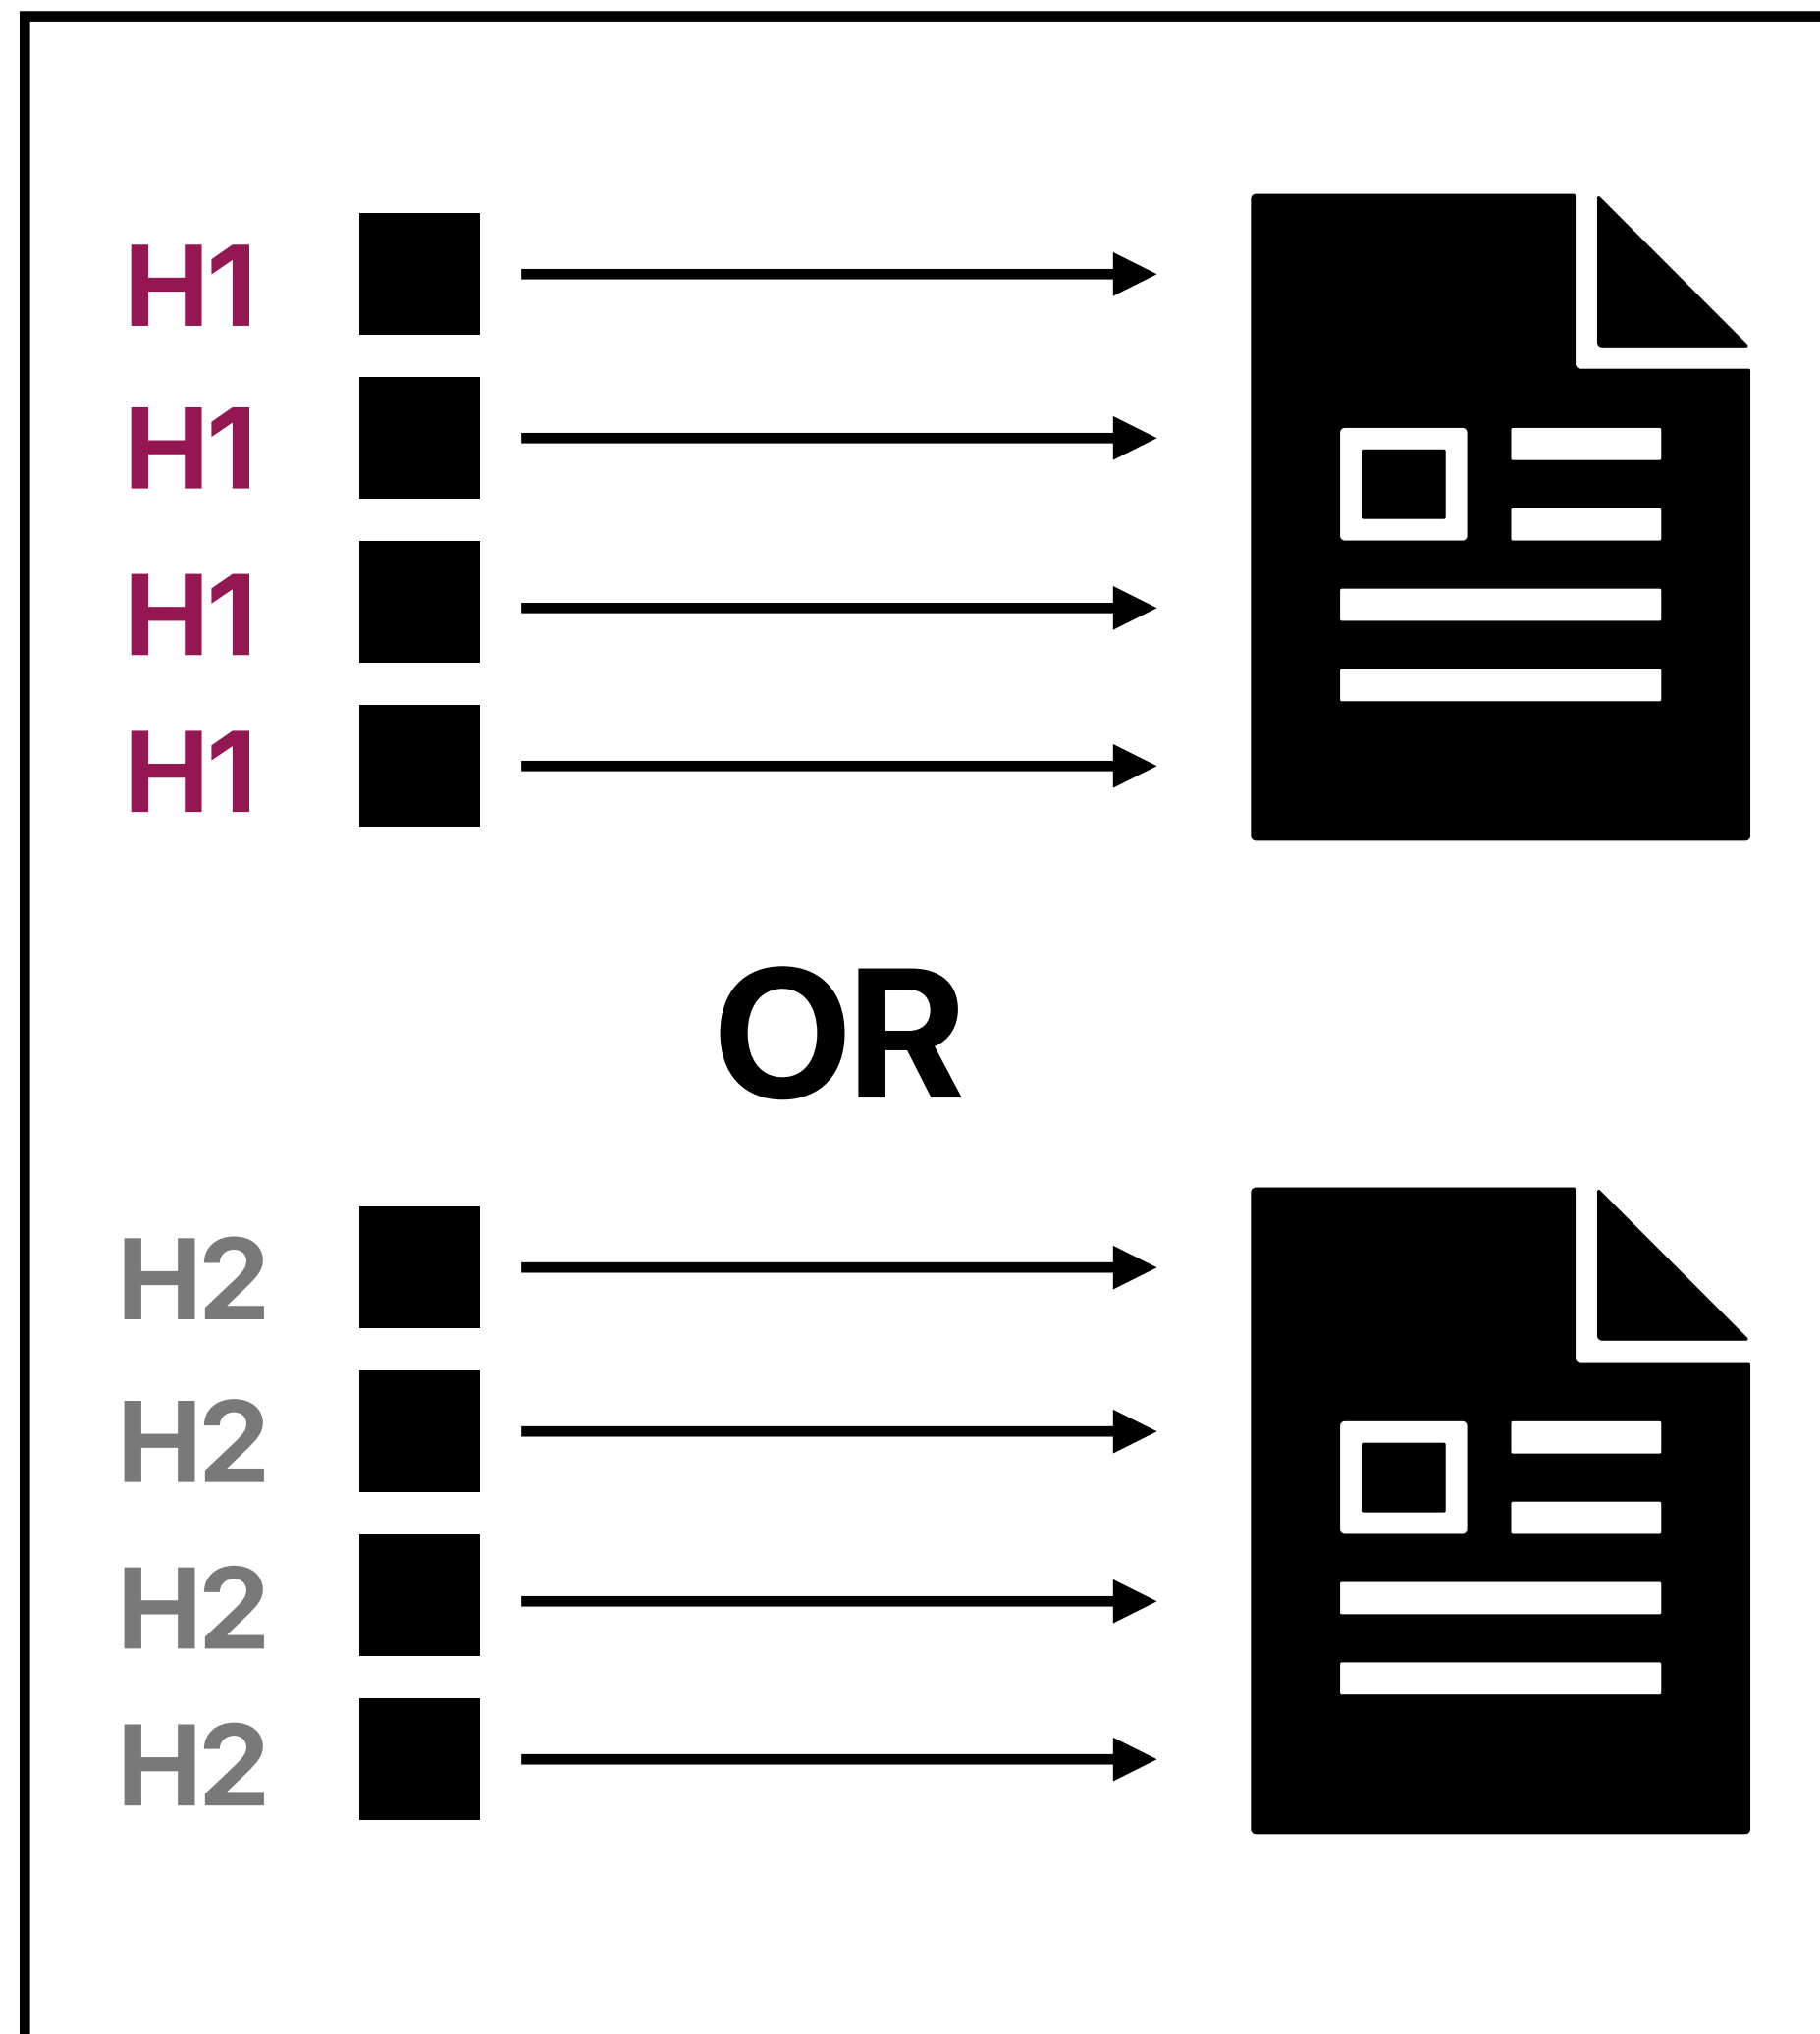

small objects



delay or loss?

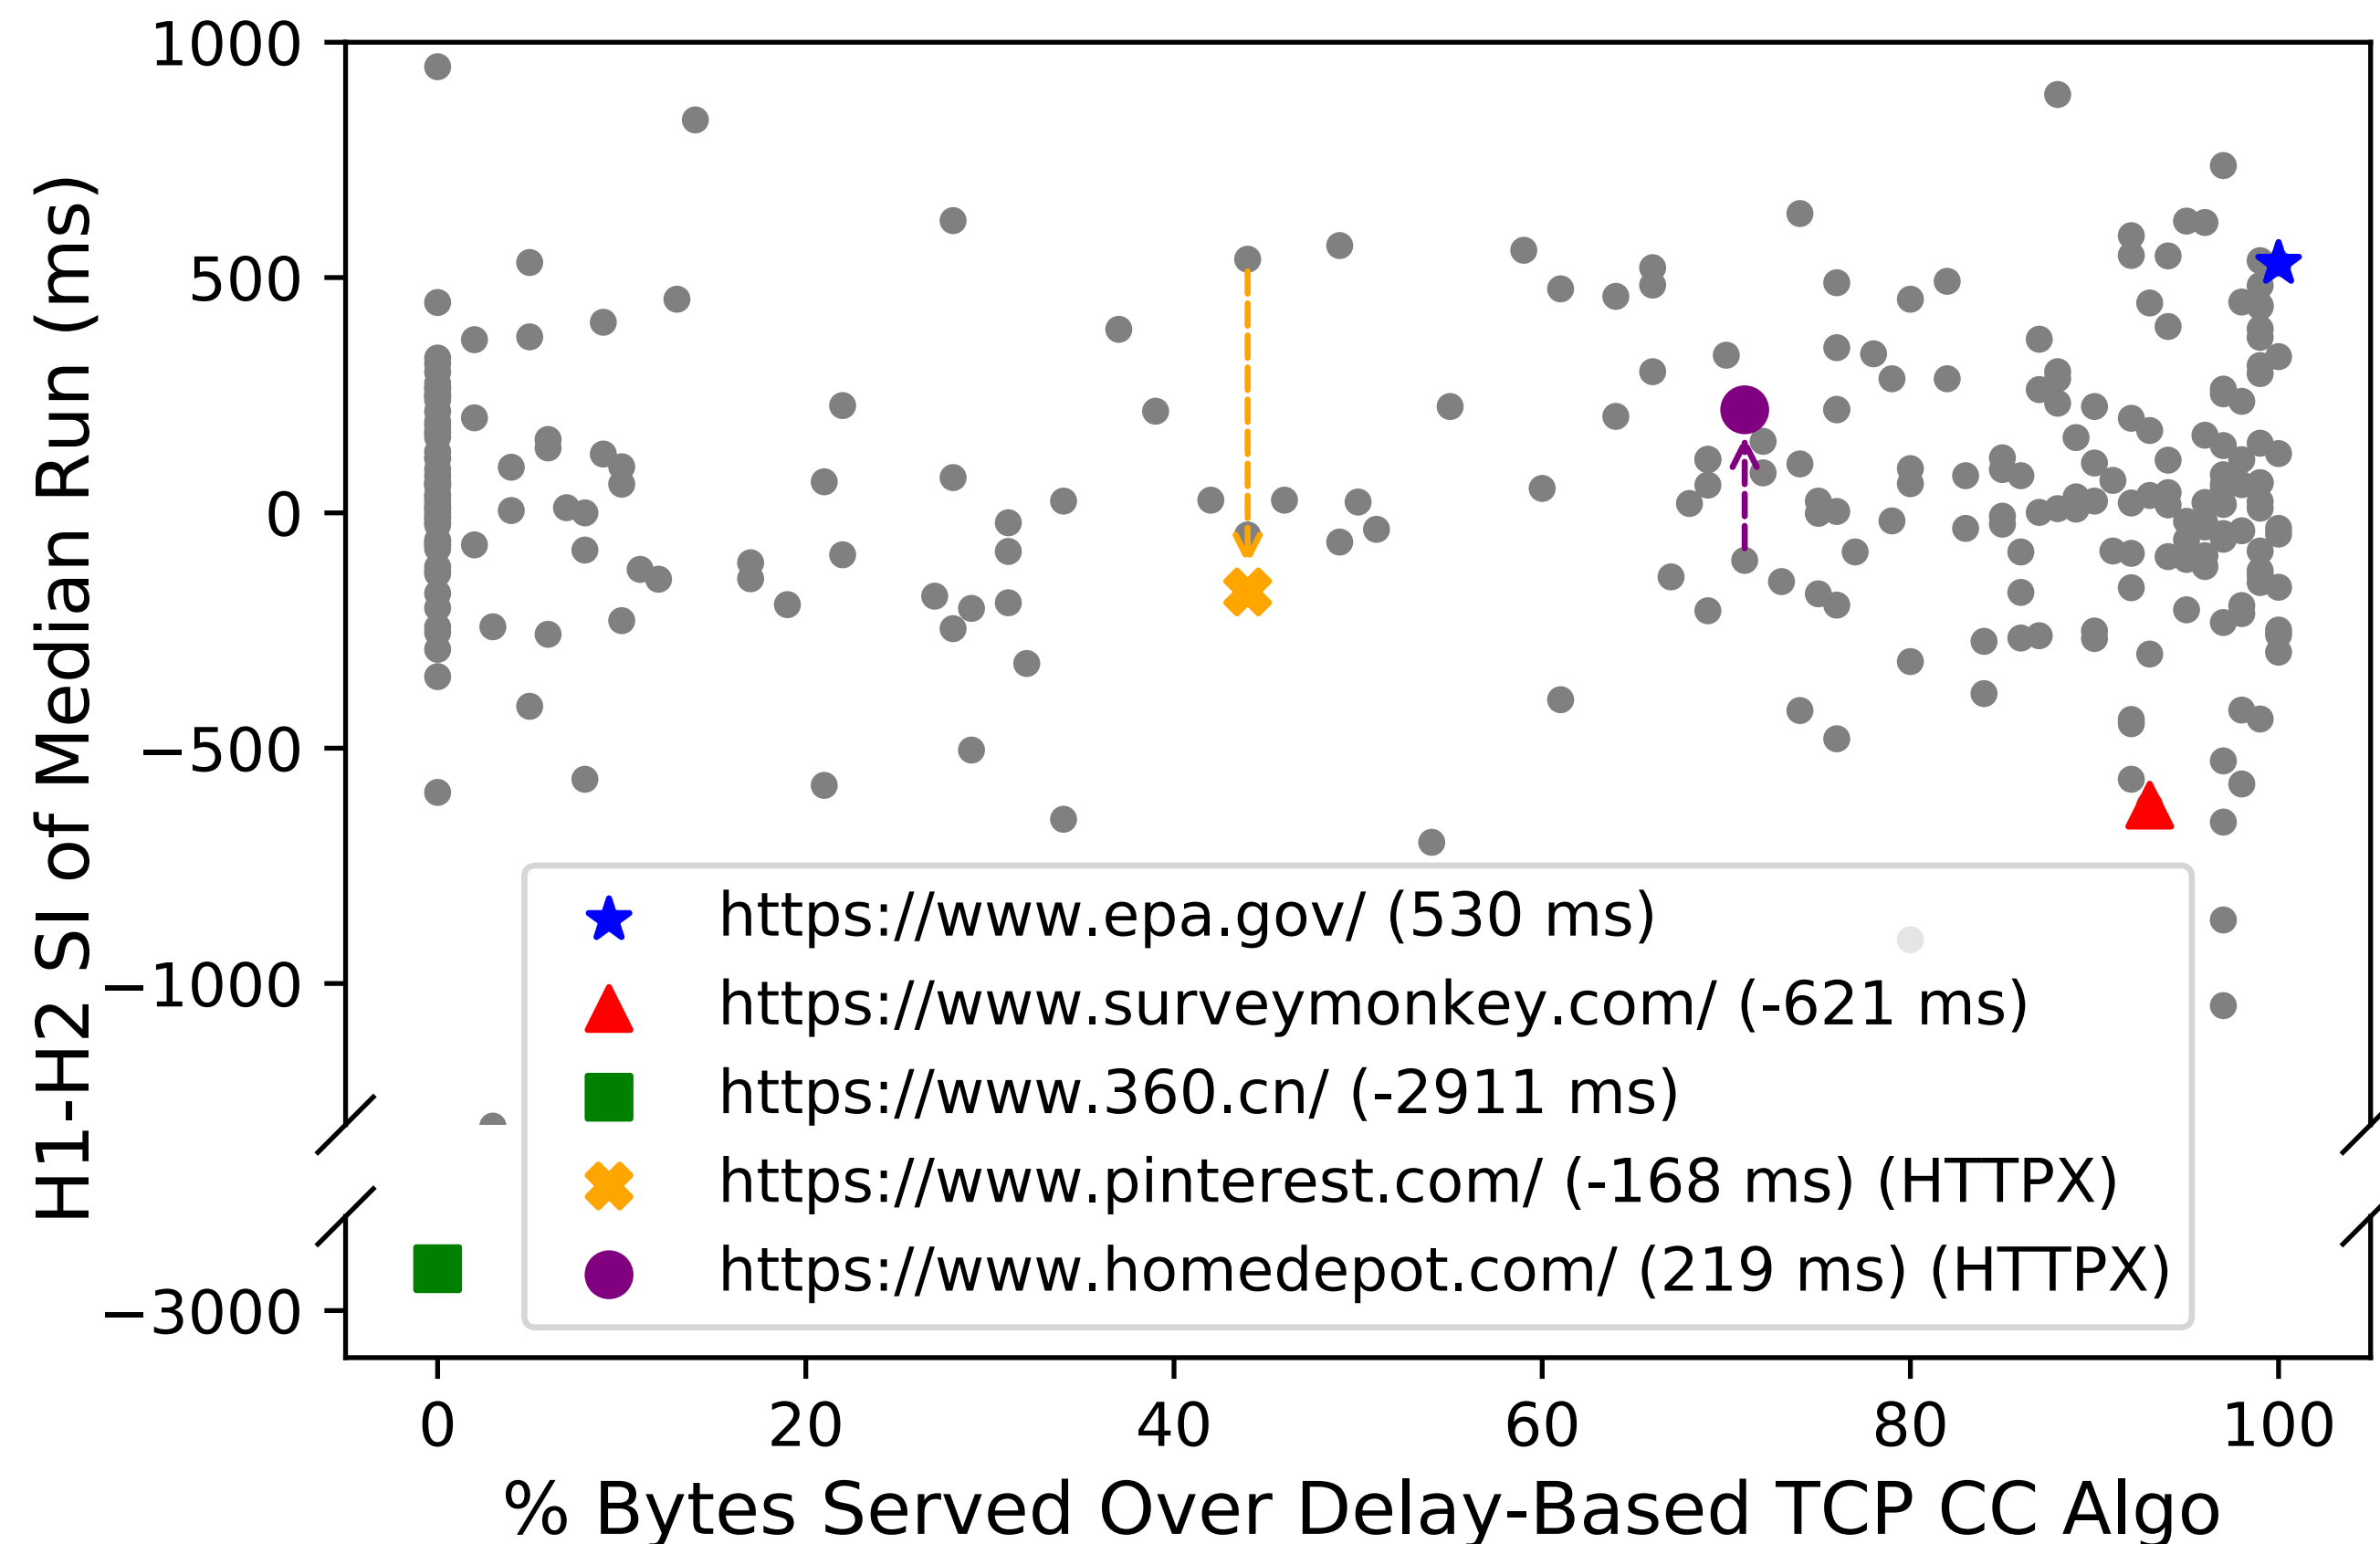


# Fine-Grained HTTP Control



Developed a **Chromium Patch**


# Toolbox + 300 Web Pages



- HTTP/2 not always the best → goes against de facto approach of always using HTTP/2
- Paper presents a few case studies that dig into details of page load and explain effect on web performance
- Motivates a deeper and more sophisticated study


# Conclusion

HTTP variants



delay-based CC

evolved the web  
performance  
toolbox



Thank you. Available for questions at [neilagarwal@cs.ucla.edu](mailto:neilagarwal@cs.ucla.edu).