

Measuring and Mitigating Web Performance Bottlenecks in Broadband Access Networks

Presented by Angela Jiang

Reducing Web page load time

1.2% decrease in Bing revenue with 500ms delay

1% decrease in Amazon revenue with 100ms delay

Every millisecond counts!

Performance optimizations

Client-side

- SPDY, QUIC

- Browser caching

ISP edge

- Content and DNS caching

Server side

- Persistent connections

- TCP ICW

In modern broadband access networks

What are the bottlenecks of Web response time?

How can we mitigate them?

Page load time

Time from initial request to all objects retrieved

DNS lookup time

Time to first byte

- TCP connection initiation

- Server response

Object download time

Measuring page load time

Measurements from 5000+ homes

SamKnows and Bismark

Deployed on routers across US

Uses Mirage

Mirage

Used to estimate page load time

Shows breakdown of page load time

Only processes statics objects

Waits to download home page before processing it

After parsing home page, performs all DNS queries before downloading any

Mirage vs Phantomjs



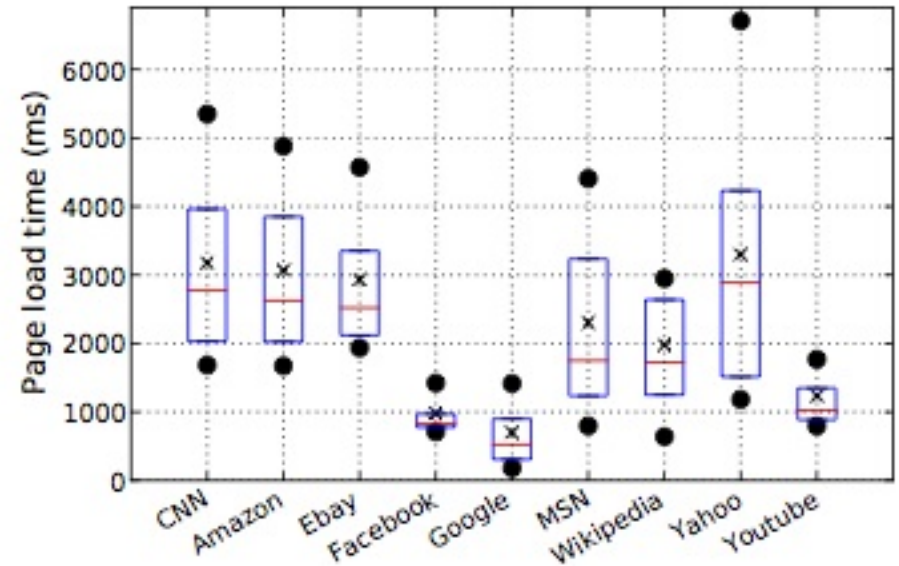


Figure 2: Comparison of Mirage to Phantomjs. We see that Mirage underestimates the page load times: real load times are higher than what we measure.

Higher page load times outside US

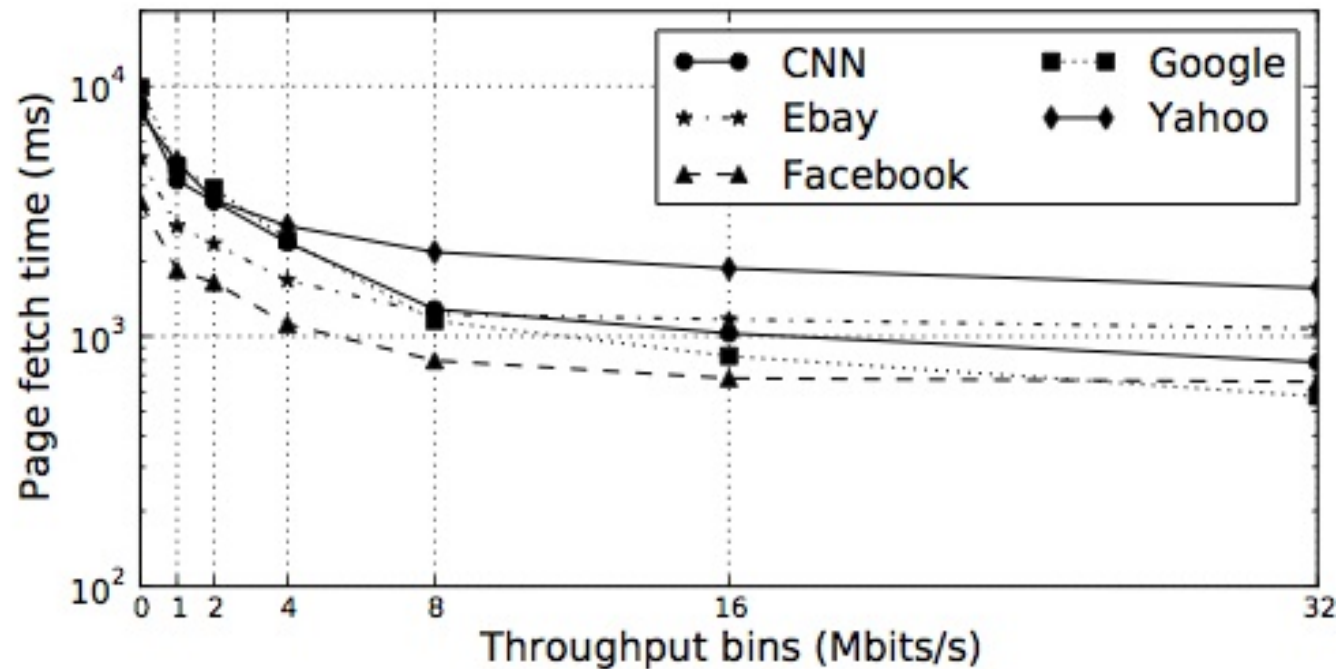


(b) *BISmark-US*



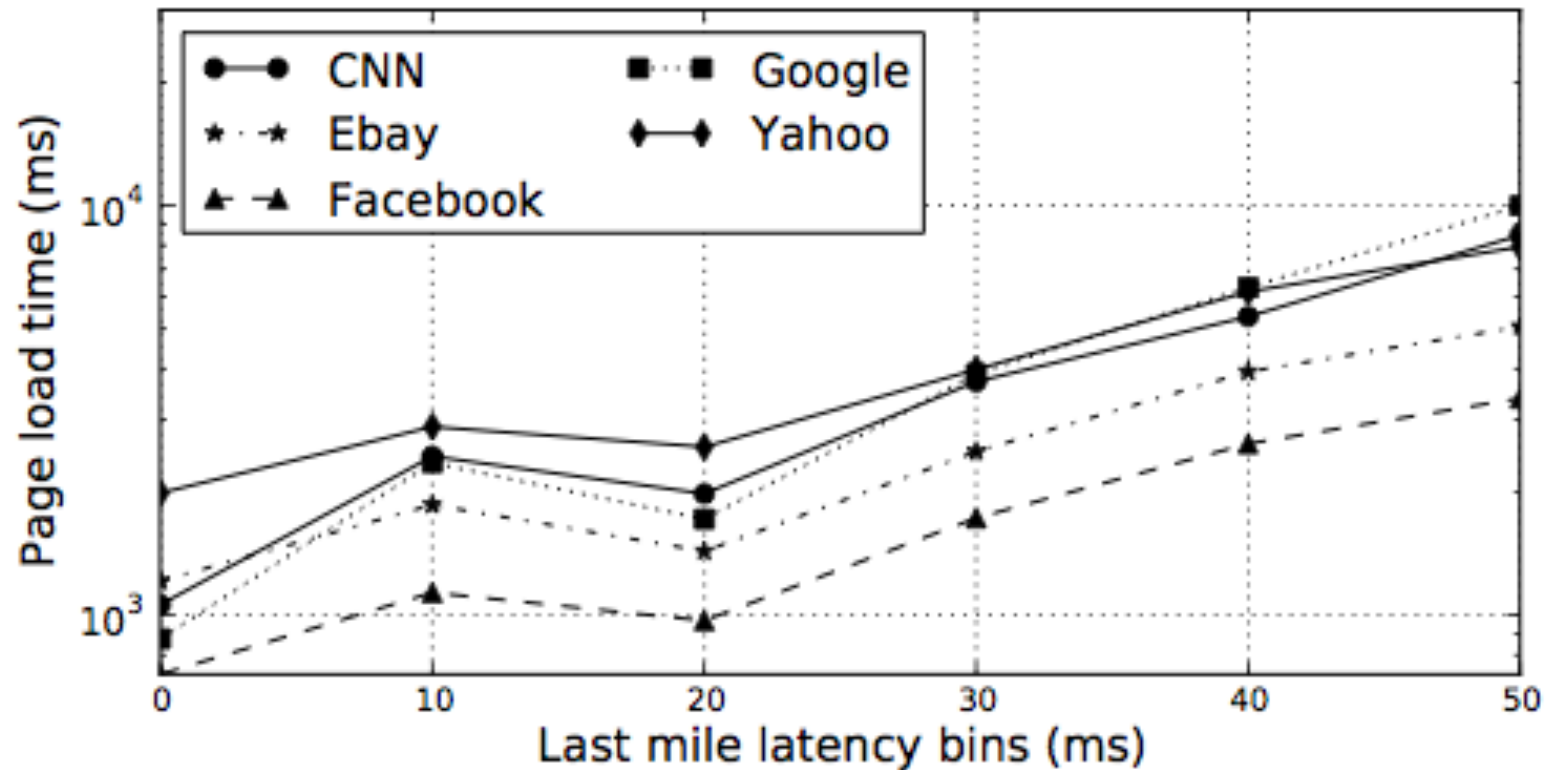
(c) *BISmark-nonUS*

Throughput not a bottleneck after 16Mbps



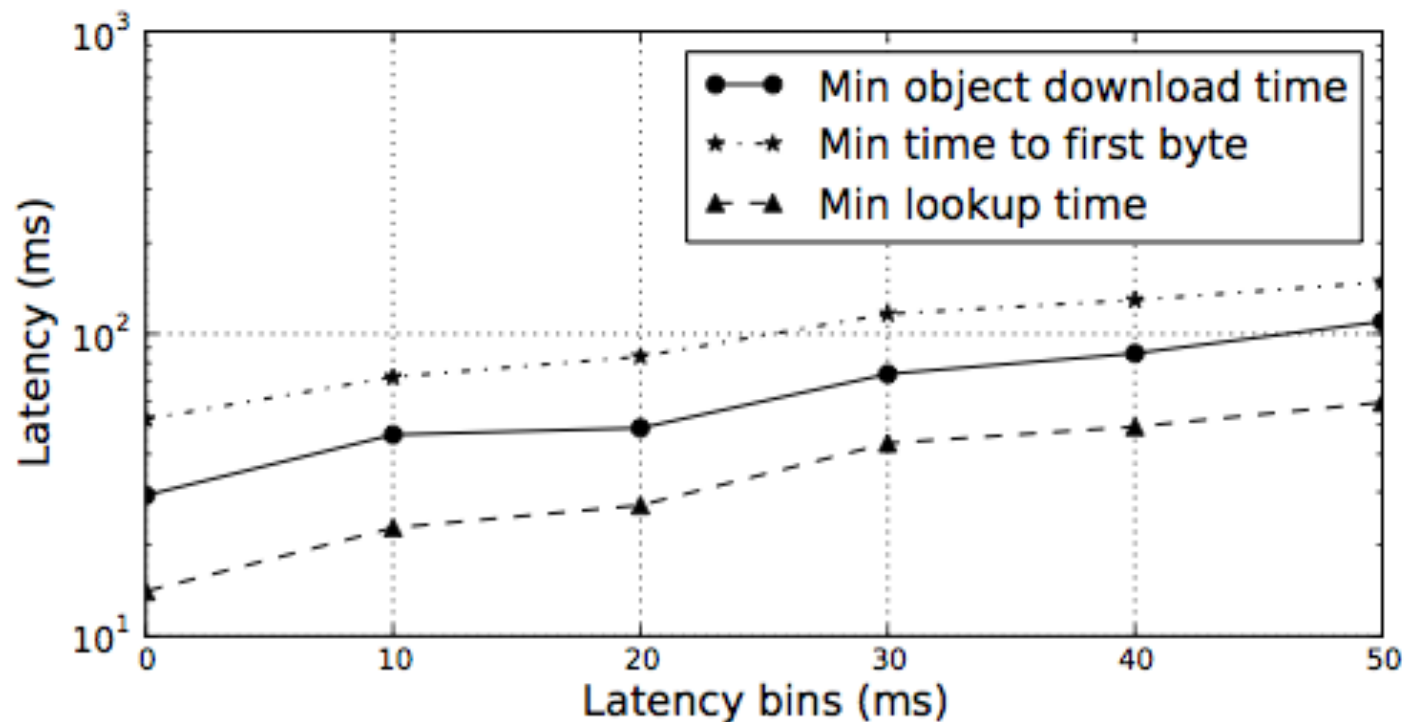
(a) Page load times

Last mile latency has multiplicative effect



(a) Page load times

TCP overhead is large on small objects



(b) *Components of page load time*

Where are the bottlenecks?

Low bandwidth (under ~16Mbps)

Throughput -> cache content

High bandwidth (over ~16Mbps)

Latency -> cache DNS responses, TCP connections

How to optimize for latency

Preemptive, persistent TCP connections

DNS caching

Content caching

DNS resolvers and CDNs stop at ISP edge

Doesn't solve last-mile latency problem!

Caching must be done inside the home

TCP connections already setup




Figure 8: *Connection caching in the home can reduce median page load times by 100–750 ms. (Home Connection Proxy vs. Home DNS Measurements)*

DNS Caching in home vs ISP

Maximum times

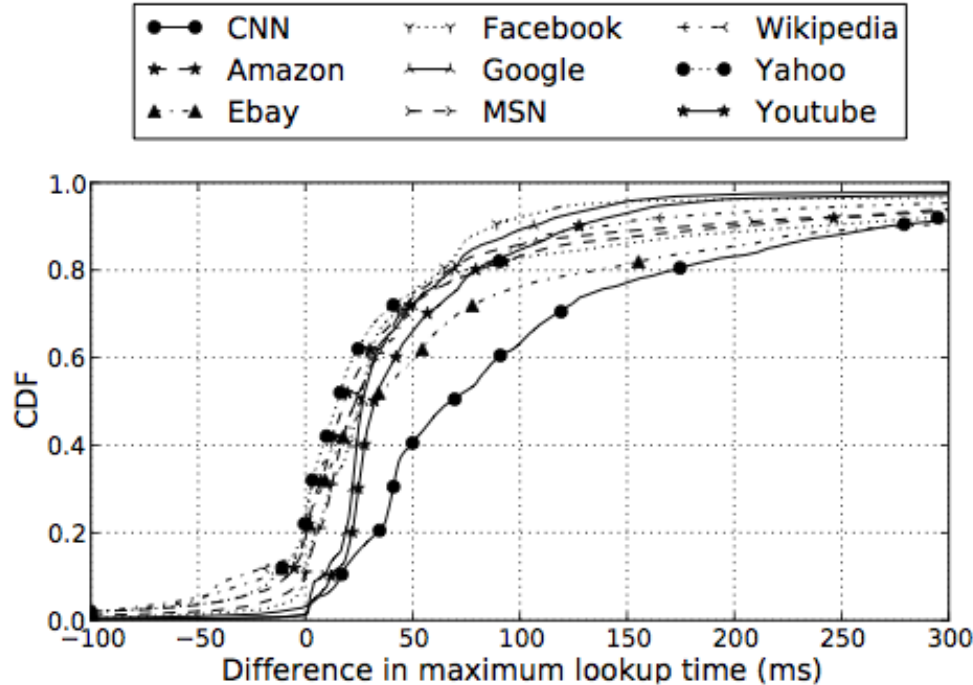


Figure 7: *Caching DNS in the home can reduce the maximum DNS lookup time by 15–50 ms. (Home DNS Measurement vs. No Proxy, ISP DNS Measurement)*

Content caching vs. DNS+Connection




Figure 9: *Content caching reduces the median page load time by 75–400 ms over connection caching alone. For sites with more cacheable content, the benefit is greater (Home Proxy vs. Home Connection Caching Measurements)*

Content caching in home vs ISP


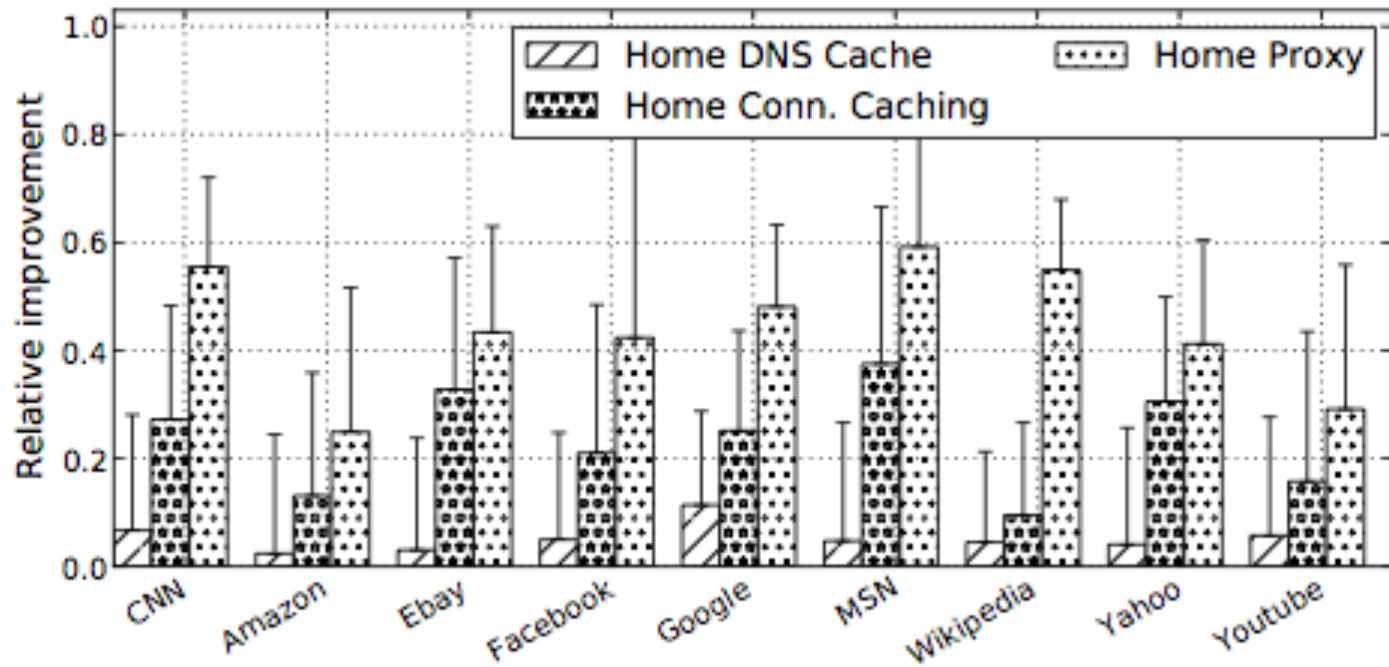


Figure 10: *Running a proxy in the home improves median page load times by 150–600 ms versus running a proxy in the ISP. (Home Proxy vs. ISP Proxy Measurement)*

Relative improvement of page load time



(a) *BISmark-US*

Popularity based prefetching

Cache size: 20 domains

Timeout threshold: 2 minutes

Improved DNS hit ratios: 19-93%

Improved connection hit ratios: 6-21%

Browser using popularity based prefetching

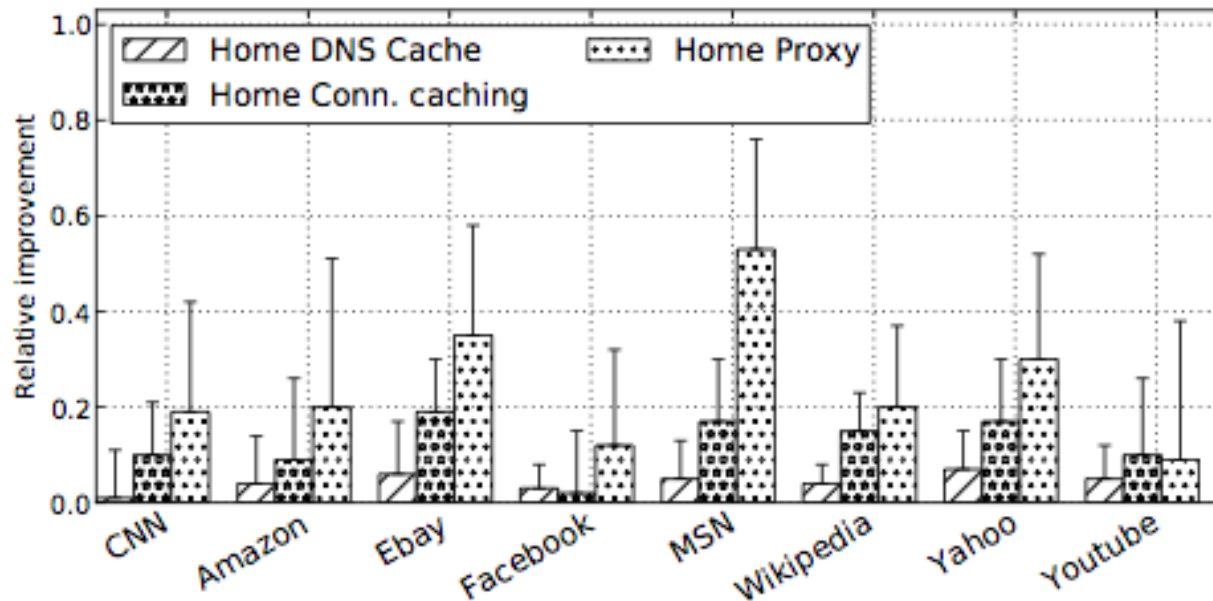


Figure 12: Average relative improvement in page load times for various optimizations, as observed from the browser. Error bars denote standard deviation.

Conclusion and discussion

Confirms and quantifies

- Last-mile latency problem

- Benefits of caching at the home

Anything novel or surprising?

- Persistent connections and DNS caching at the router happens

- Reason to cache at the router instead of browser?

- Server load from maintaining longer persistent connections?