Characterizing and Mitigating Web Performance Bottlenecks in Home Networks

Srikanth Sundaresan, Nazanin Magharei, Nick Feamster
Georgia Tech
Renata Teixeira
CNRS and UPMC Sorbonne Universités
Sam Crawford
SamKnows

Broadband and the Web

- Broadband speeds are increasing
  - Do applications get the benefit of higher speeds?

spdy://

Which factors affect Web page load time?

Reducing page load time is important

- Server side
  - Make TCP work better: higher ICW
  - Newer protocols: SPDY
- Network
  - Move content closer: CDNs
  - Increase throughput
  - Reduce lookup time

Page load times are still high

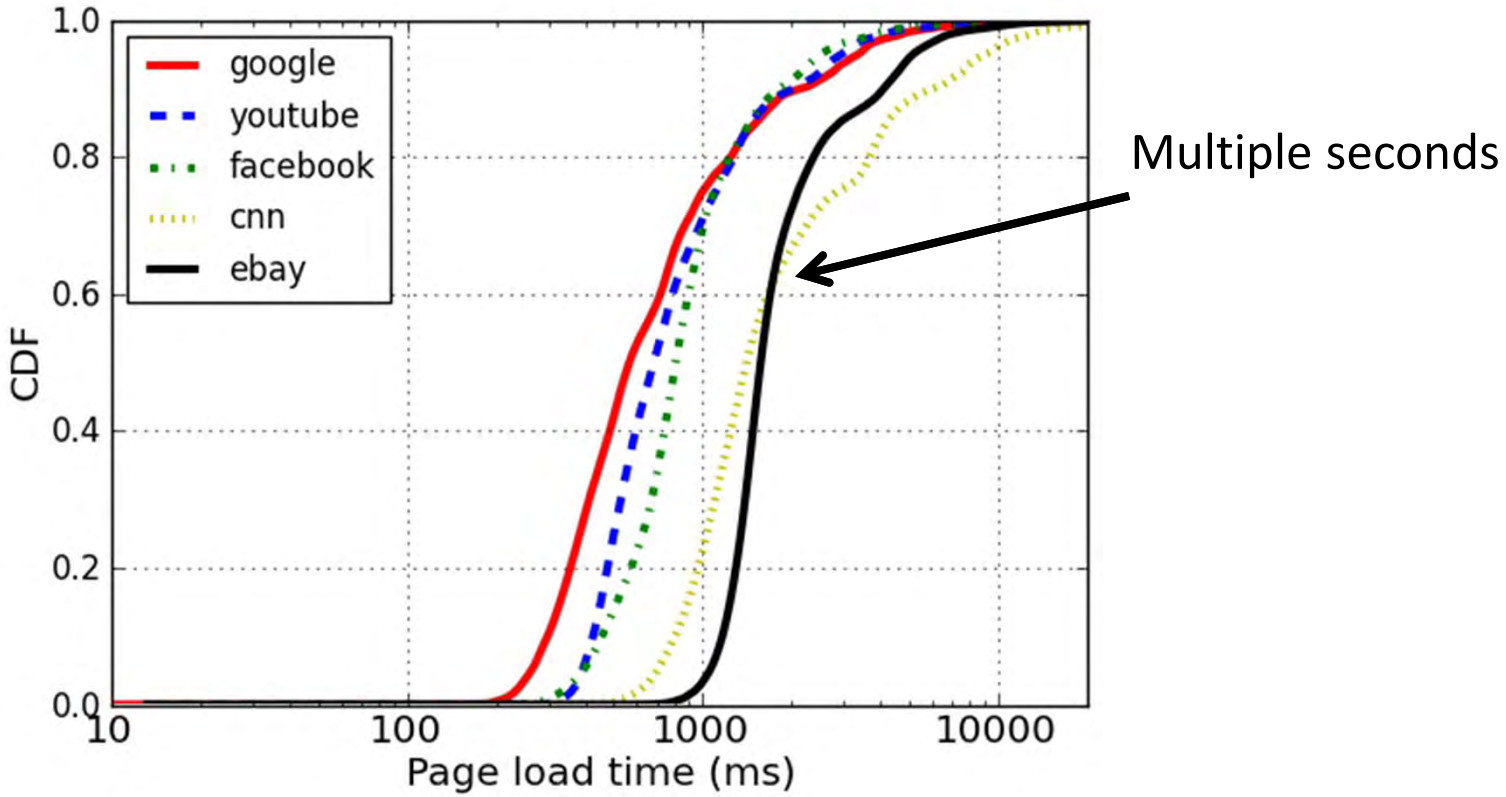

Web Performance Bottlenecks:

Throughput, latency, DNS, TCP overhead

Outline

- Characterizing Web Performance Bottlenecks
  - Throughput bottlenecks
  - Latency bottlenecks
- Analyzing benefit of acceleration techniques from the home
  - DNS caching
  - TCP connection caching
  - Content caching

The Data

- Collected from the FCC/SamKnows deployment
  - 11 ISPs, 5500 homes, 1 month
- Web test client
  - 9 popular sites
  - Up to 8 concurrent threads
  - Breakdown of lookup time, time to first byte, time to fetch objects

Higher throughput does not always improve load time

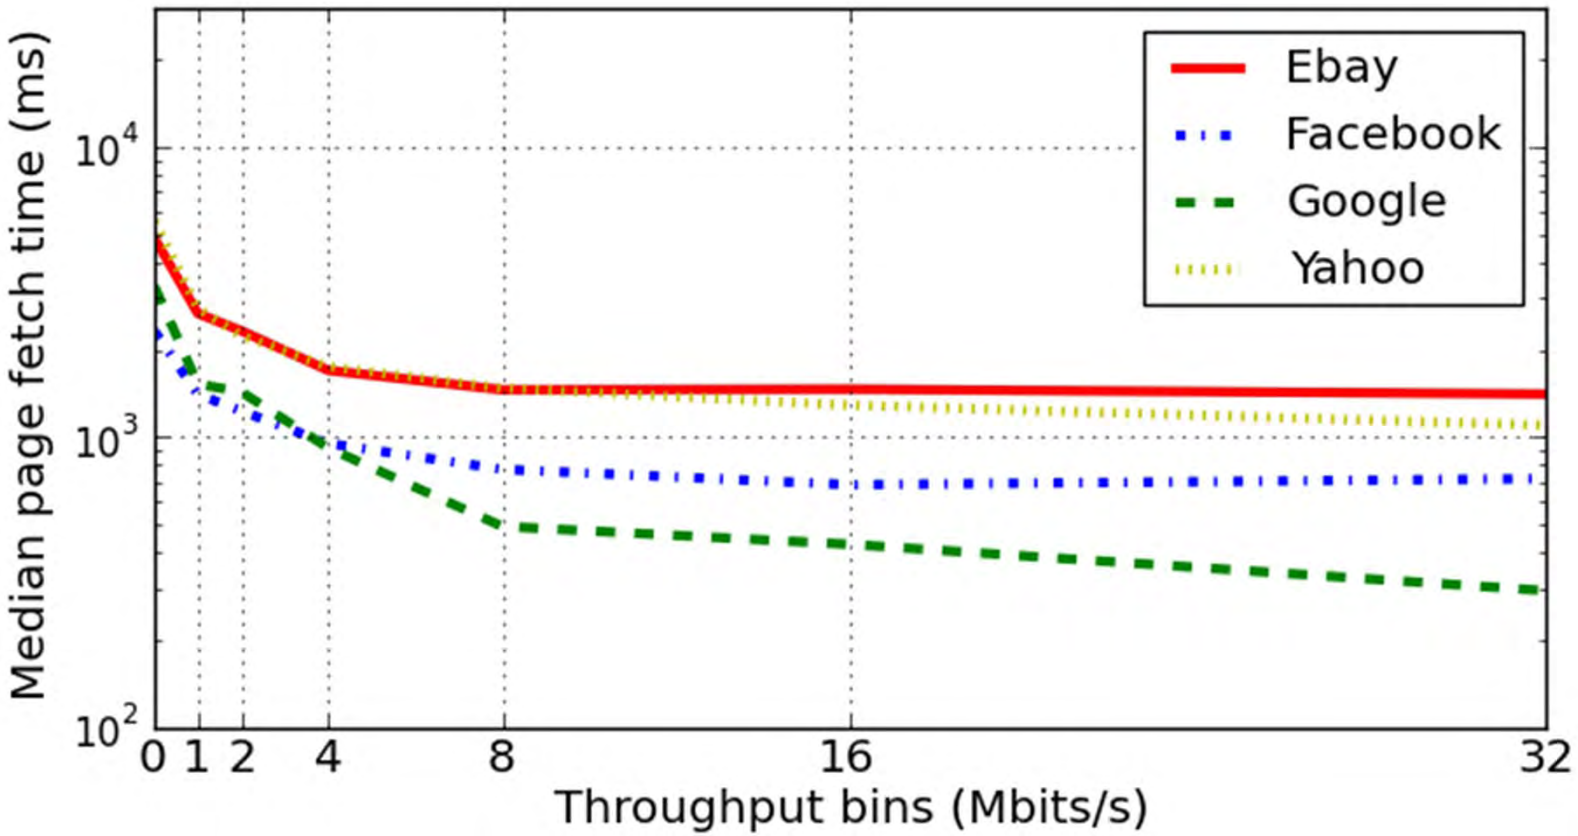

Load time flat-lines after 8 Mbits/s

The access link is a bottleneck

- Both throughput *and* latency affect load time
  - Last mile latencies can exceed 40 ms
  - Increases load times by many times that number
- Reducing impact of the last mile
  - Apply acceleration techniques in the home!

Outline

- Characterizing Web Performance Bottlenecks
  - Throughput bottlenecks
  - Latency bottlenecks
- Analyzing benefit of acceleration techniques from the home
  - DNS caching
  - TCP connection caching
  - Content caching

Experiment set up

- Study impact of:
  - DNS caching
  - TCP connection caching
  - Content caching
- Web client fetches objects serially
  - *curl*, *dnsmasq, polipo*
- BISmark test bed
  - 27 homes, 13 days

Summary of Results

- Any sort of caching helps!
  - DNS caching saves 10-200ms
  - Connection caching saves 100ms-1s+
  - Content caching saves 100ms-1s+ *over connection caching*
- Improvements depend on site characteristics
  - Server location, number of domains
  - Number of objects
  - Cacheability of DNS records, content

Relative impact of different schemes

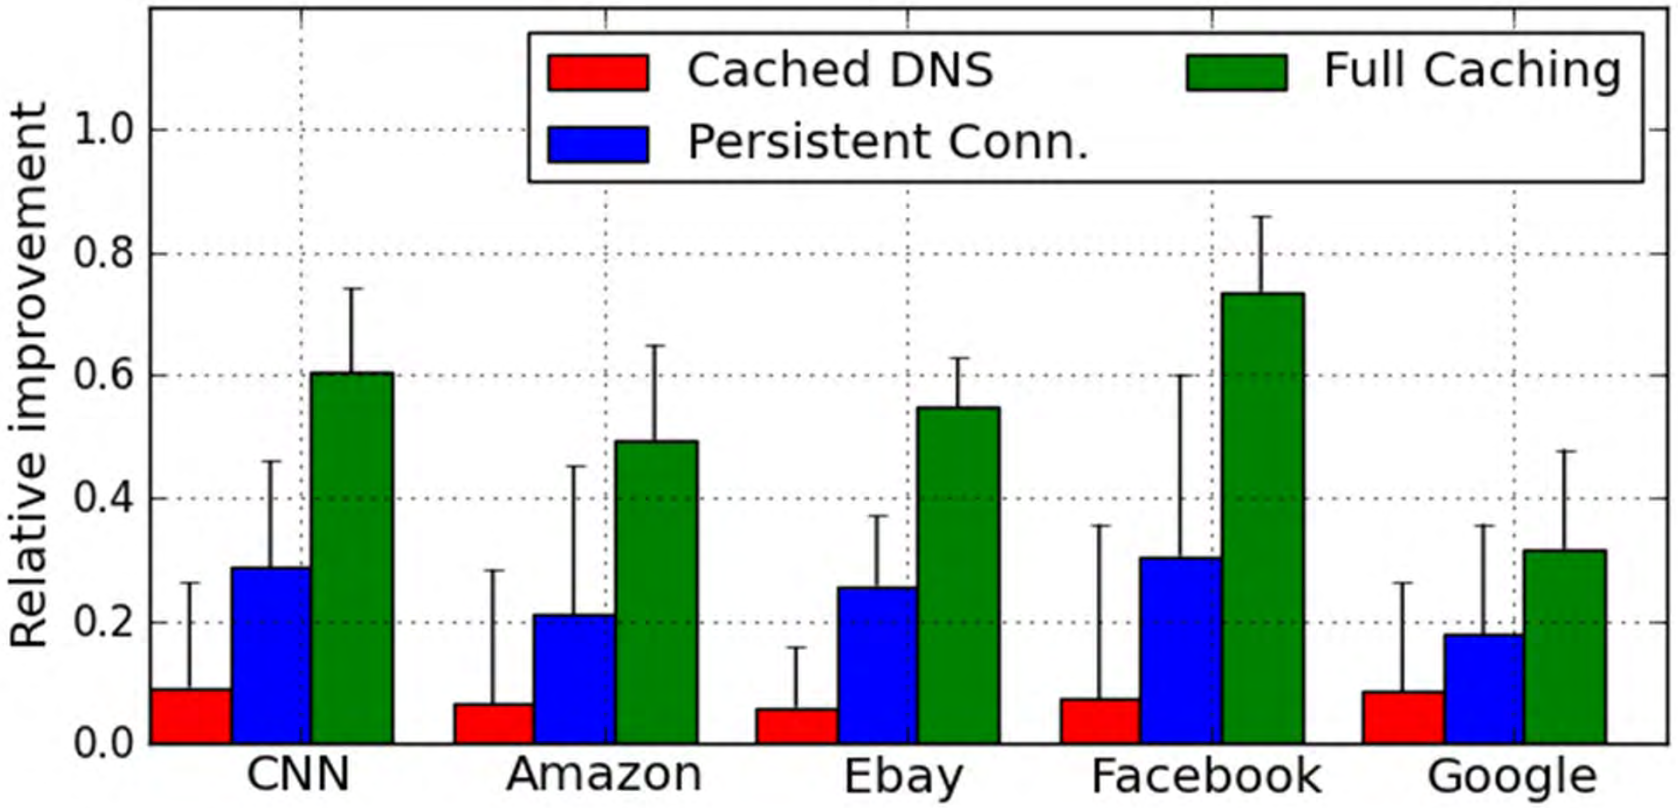

Impact varies across sites, but connection and content caching offer significant benefits

Is it better than ISP-edge caching?

Improvement in load time by caching in home vs nearby server

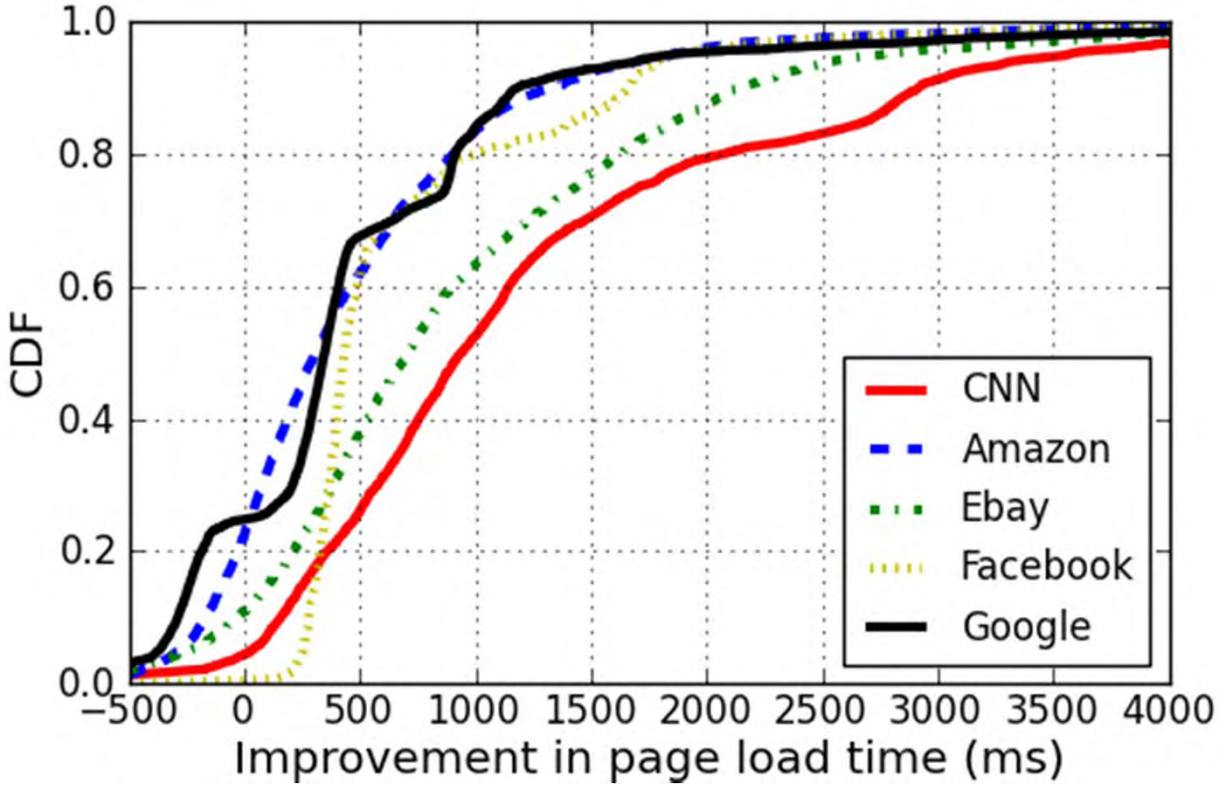

Optimizations in the home have significant impact

Future Work

- Implement system to do “popularity based” caching
  - Pre-fetch popular DNS domains
  - Keep connections open to popular Web sites
- Initial implementation on BISmark router
  - Understand potential improvement over end-host implementation

Conclusion

- Latency becomes major bottleneck at higher throughput
- Network-side optimizations help, but last-mile presents fundamental problem
- Moving standard optimizations inside the home results in major improvements

srikanth@gatech.edu