

Does More Bandwidth Really Not Matter (Much)?

Seraj Al Mahmud Mostafa¹, Mike P. Wittie¹, and Utkarsh Goel²

¹ Montana State University, Bozeman, MT, USA

² Akamai

Abstract. The prevailing wisdom is that more network bandwidth does not matter much and that website performance is primarily limited by network latency. However, as mobile websites become more complex and mobile network performance improves, does this adage continue to hold? To understand the effects of small changes in network bandwidth and latency on website performance, we propose a novel webpage characterization metrics – Critical Path of Improvement (CPI). We compute CPI for 45 websites and analyze it against the network performance of four mobile ISPs in 57 US cities. Our results show that 18% of websites are primarily limited by bandwidth with others limited by bandwidth to some extent. These results show that contrary to accepted wisdom, insufficient bandwidth is a limiting factor in some website/network combinations. We also offer a discussion of approaches website developers and mobile network administrators can follow to understand and mitigate bandwidth limitations to website performance.

Keywords: Mobile web, measurement, web performance metrics

1 Introduction

The evolution of mobile web services goes hand-in-hand with the improvements in mobile network performance. As mobile networks offer lower latency and higher bandwidth, mobile websites may transfer more content to offer a richer experience, while maintaining good response times. But what are there instances where website design and network performance do not match up well?

The prevailing wisdom is that network latency is the limiting factor to website performance. Two key articles driving this idea came from Mike Belshe [1] and Ilya Grigorik [2]. They observe that beyond some point, adding more bandwidth results in diminishing marginal returns of website load time. Incremental reduction of latency, on the other hand, results in corresponding improvement of website performance. It would seem then that beyond a certain minimal bandwidth website performance is limited by network latency.

While that may have been true a decade ago for landline internet and desktop websites, much of today's web traffic involves mobile networks and smartphones. The mobile networks themselves have undergone significant improvement in the last several years. OpenSignal reports that from February, 2016 to January, 2020

4G latency (averaged across providers) improved from 75.5 ms to 52.9 ms (by 29.9%) and bandwidth from 9.7Mbps to 25.6Mbps (by 163%) [3, 4]. At the same time, there remain pronounced regional differences in mobile network performance. Specifically, in January, 2020 the variation of mean latency between US cities was as large as 37.8ms and the variation of mean bandwidth more than 28.9Mbps [3]. Further, as websites become more complex they may act differently in networks of varying performance. Considering these differences, does the adage of network latency being the limiting factor website performance continue to hold?

To answer this question, we take a fresh look at characterizing mobile website network performance requirements. We propose a new website characterization method we dub the Critical Path of Improvement (CPI), illustrated in Figure 1. CPI allows us to understand how small changes to network bandwidth and latency affect website performance as measured by the Perceptual Speed Index (PSI) [5], or any other web performance metric [6]. We base CPI on the observation that at any given point of network performance on the bandwidth/latency plane, PSI will improve on average with small increases to bandwidth, or small reductions to latency. Figure 1 illustrates PSI improvement by lighter points along the CPI. We evaluate the impact on PSI of these small changes in network performance and record CPI as the path of the fastest PSI improvement (details in Section 3). Effectively, CPI allows us to understand when a website performance benefits more from additional bandwidth or from lower latency.

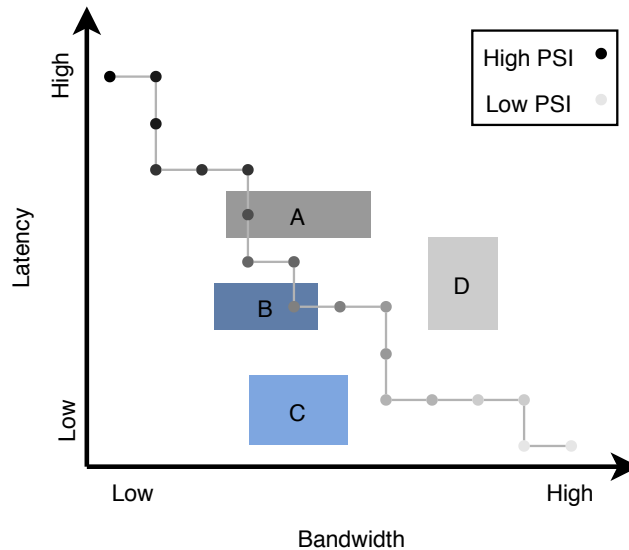


Fig. 1. Illustration of CPI and network conditions.

To understand whether latency or bandwidth of mobile network is the limiting factor for a given website, we analyze the relationship between a network’s performance and the CPI. In Figure 1, the rectangular regions represent network performance envelopes in terms of minimum and maximum latency and bandwidth measurements. For a network whose performance envelope lies in region D, the path to better PSI points downwards to the CPI and so via lower latency. On the other hand, if the network performance envelope corresponds to region C, the path to the CPI and better PSI lies towards higher bandwidth. We want to understand how common are scenarios where a website could benefit from higher bandwidth, or lower latency, than what is available in a given network deployment. In other words, how often is website performance limited by insufficient bandwidth versus by high latency.

This paper offers the following contributions:

- We propose CPI – a new website performance characterization method. We show how developers may use CPI to understand the impact of small changes in network performance on the performance of their websites.
- We analyse the CPI of 45 websites from the Alexa-100 list in four mobile networks across 57 US cities to identify cases where websites are bandwidth- and latency-limited.
- We analyse regional variation in network performance to show areas of the US where network bandwidth tends to be the limiting factor to website performance.
- We point to factors of webpage design that contribute to websites being limited by bandwidth.
- We offer a discussion of how developers and network operators may use these results to improve the match between website design and network performance.

The rest of this paper is organized as follows. In Section 2 we discuss related work on website performance bottlenecks. Section 3 details CPI measurements and comparisons with network performance. In Section 4 we describe our measurement methodology and discuss measurement results. Finally, we offer a discussion of the applicability of our findings to website developers and network administrators in Section 5 and conclude in Section 6.

2 Related Work

A number of papers quantified performance bottlenecks in mobile websites. Ne-jati and Balasubramanian found that device processing speed as well as delays in transferring objects on the critical path of rendering are the main causes of lower website performance on mobile devices [7]. While some researchers have focused on reducing the computational load of rendering [8, 9], others have focused on eliminating loading bottlenecks.

One approach to eliminating loading bottlenecks is to reorder objects on the critical path [6, 10]. Even with object reordering, Furtak et al. have shown that

increased bandwidth improves page Speed Index by permitting large images to load more quickly [6]. However, others have observed that the impact of extra bandwidth is less pronounced on Page Load Time (PLT) [11]. Han et al. have shown that extra bandwidth improves the performance of large pages [12]. They loaded pages using MultiPath TCP (MPTCP) over cellular and WiFi connections in parallel and showed that the loading process can use the bandwidth of the additional link effectively, when the performance of the links is comparable in terms of latency.

These results, however, focus on the impact of network performance in general, but do not consider the correspondence of specific web pages loaded in specific mobile networks. While researchers agree that additional network bandwidth may be helpful in some cases, its not clear to what degree such cases are prevalent. In this paper we aim to answer that question.

3 Critical Path of Improvement

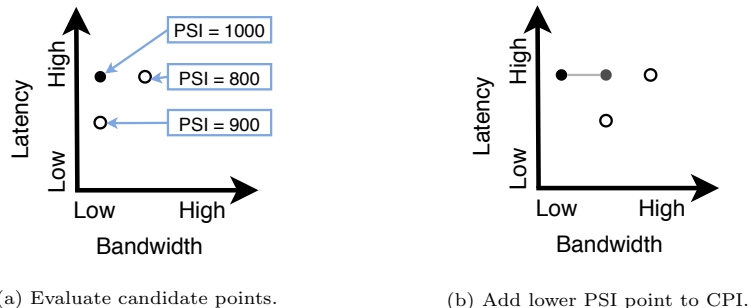


Fig. 2. CPI creation process.

CPI measures how small changes to network performance affect website performance. To obtain the CPI, we load a given website under different combinations of latency and download bandwidth controlled using NetEm [13]. For each set of network conditions we measure website PSI using PWMetrics [14].

Using these tools, we measure CPI as follows:

- As a starting point we chose a set of atrocious network conditions, specifically bandwidth of 256 Kbps and latency of 180 ms. We load a website over thus constrained link and measure its PSI. For example, in Figure 2a these low bandwidth, high latency conditions might result in the PSI of 1000.
- We then proceed to find the next CPI point. We measure the PSI of two candidate points. For the first candidate point we improve network latency by 10 ms, but keep the bandwidth of the starting point. For the second candidate point, we improve network bandwidth, but keep the latency of

the starting point. Since improvements in network bandwidth have a greater impact on website performance in the lower ranges [1, 2], we double the network bandwidth up to 8 Mbps and then increase it by a constant interval of 8 Mbps. We then measure the PSI of the candidate points. In Figure 2a the candidate points have the PSI of 900 and 800 respectively.

- Finally, we select the next point on the CPI as the candidate point with the lower PSI of the two, as shown in Figure 2b. Its important to note that website PSI varies somewhat across loads. When computing the CPI of the starting and candidate points we average the PSI over seven trials.

Figure 1 shows an example of a complete CPI measurement. As we move along the CPI from the starting point, we can see improvements to PSI represented as lighter points. The CPI shows how small changes to network performance improve website performance. The CPI also shows a website’s *preference* for higher bandwidth, or lower latency at any point. A CPI with more points further away from the x-axis show a preference for more bandwidth, while those with more points close to the x-axis show preference for lower latency.

We also want to consider the relationship between the CPI and the performance of a network to understand whether a website is bandwidth- or latency-limited *in a given network*. In Figure 1 we plot four areas: A, B, C, and D. These areas represent some possible envelopes of a networks performance represented as minimum and maximum latency and bandwidth. While a network’s performance varies during a given time period, it will in general be bounded some maximum and minimum values, as reported by OpenSignal measurements [3, 15, 16]. As such, a network’s performance constrains page PSI to a particular region of the bandwidth/latency region.

Referring to the regions in Figure 1 identify four cases of these constraints.

- Case A: the CPI exits network performance bounds through the minimum latency boundary of area A. This means that page PSI would improve more if A extended to better (lower) latency, rather than towards higher bandwidth.
- Case B: the CPI exits network performance bounds through the maximum bandwidth boundary of area B. This means that page PSI would improve more if B extended to better (higher) bandwidth, rather than towards lower latency.
- Case C: the CPI lies to the right of the network performance bounds of area C. This means that page PSI would improve more if C extended to better (higher) bandwidth, rather than towards lower latency. The rationale for this claim is that a CPI measurement started in area C tends towards the shown CPI and so by improving network bandwidth to the right.
- Case D: the CPI lies below the network performance bounds of area D. This means that page PSI would improve more if D extended to better (lower) latency, rather than towards higher bandwidth. The rationale for this claim is that a CPI measurement started in area D tends towards the shown CPI and so by improving network latency downward.

Our goal is to understand how common are the B and C cases. Their presence indicates that web pages could benefit from additional network bandwidth, which

stands in opposition to the accepted wisdom that web performance is constrained by insufficient latency.

4 Evaluation

Our goal is to understand the relationship between page CPI and network performance envelopes. To frame our results, we first describe our measurement methodology.

4.1 Measurement Methodology

We chose 43 pages from the list of Alexa top 100 pages, excluding adult content [17]. To satisfy our curiosity, we also added major news sites from the authors home countries: `prothom-alo.com` and `wyborcza.pl`. For each site, we measure the CPI using the process described in the previous section. We load each page from within the Montana State University’s wired network, while constraining the performance of the internet link with NetEm [13]. Using `speedtest.net` we measured the performance of our campus network to the regional ISP hosting the Content Distribution Network (CDN) servers that serve web content. Our network is limited to 18.14 ms of latency and 307.64 Mbps of bandwidth, and so NetEm constrains our network in real terms up to these numbers.

We then find the network performance regions of network providers using data from OpenSignal. Specifically we extract 4G mobile network performance for four major network providers (AT&T, Sprint, T-Mobile, and Verizon) across the US for January 2020, 2019, and July 2018 [3, 15, 16] using WebPlotDigitizer [18]. These data contain the mean latency and bandwidth measurements in a 57 cities along with measurement confidence intervals [19]. These confidence intervals describe the likely upper and lower bounds on network performance and we treat them as the upper and lower bounds of the network performance envelopes.

4.2 Results

Figure 3 shows a summary of our dataset. The x and y-axis respectively show the link bandwidth and latency constraints of our CPI measurements. The figure includes CPIs for all the websites in our dataset. Each CPI is represented by a semi-transparent gray line. Segments where these lines are darker represent overlaps between CPIs of different websites. The points along CPI lines represent normalized PSI values – we illustrate improvements in PSI more clearly further down in this section. Finally, the figure represents network performance envelopes for the four provider networks.

Overall we observe that as the network performance improves, so does page PSI, as CPI points become lighter in the low right quadrant of Figure 3. We also see that CPIs tend to concentrate in the area 80 ms - 180 ms and 256 Kbps - 4 Mbps, as overlapping CPI lines become darker, but spread out at latencies

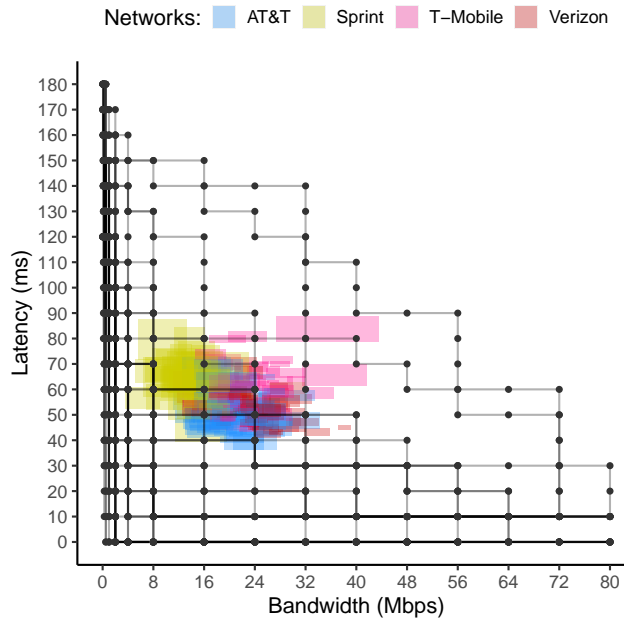


Fig. 3. CPI and network conditions.

lower than 60 ms. This tells us that while the performance of many websites initially benefits primarily from lower latency, once latency becomes sufficiently low, the benefits from improving bandwidth and latency become comparable.

We also observe that the performance of network providers tends to concentrate in a region of 40 ms-80 ms and 8 Mbps-32 Mbps. The CPIs of different websites cross this region, and so we expect to see that these websites are limited by latency and by bandwidth with respect to the different networks. In the following subsections we delve deeper into this phenomenon by characterizing the relationships between individual CPIs and individual network performance regions into A, B, C, and D cases.

How often are individual websites limited by bandwidth? Figure 4 shows the frequency of the A, B, C, and D cases on the y-axis against website URL on the x-axis. The figure shows cases based on the 2019 network performance data across all network providers and regions. We observe that most of the cases are in the A and D categories, which implies that website performance would improve with lower network latency. However, there is also a significant percentage of B and C cases. Specifically, B and C cases represent 5.69% and 9.02% of cases respectively. For 8 out of 45 (18%) of the websites B and C cases represent the majority of cases, which means that these websites are limited by insufficient bandwidth in most network/region combinations. This result shows




Fig. 4. A, B, C, D cases across websites.

a contradiction to the accepted wisdom that website performance is limited by insufficient latency in general. While that is true for most of the measured websites, it is not true for all. In fact 15 out of 45 (33%) of these websites are limited by network bandwidth in at least some network/region performance envelopes.

Our recommendation is that developers evaluate the CPI for their websites and compare it against network performance to determine whether the websites are bottlenecked on latency or bandwidth in a given network. In Section 5 we discuss techniques they may apply, depending on its limiting factor in a given website, to take advantage of available network performance.


Are certain providers more bandwidth limited? We also wanted to understand whether the frequency of B and C cases varies by network provider, or over time. Figure 5 shows the number of A, B, C, and D cases on the y-axis and network providers on the x-axis. The three panels show the number of occurrences for the different years of network performance data.

In general, we observe that for all providers most of the cases are in the D and A categories. At the same time all providers show a number of B and C cases. The frequency of B and C cases does not vary substantially across the years and in fact has decreased over the years as network providers upgrade their infrastructure. Still in the 2020 network data, B and C cases represent 12.12% of cases. The number of B and C cases does not differ substantially across providers, with the difference between 371 on T-Mobile with the most B and C cases in 2020 and 314 on Verizon with the fewest being less than 16%.


Do bandwidth limitations vary geographically? We also wanted to understand whether website performance is limited by bandwidth more often in certain areas of the country. In Figure 6 each point shows the ratio of the occurrence of B and C cases to A and D cases for different US cities. Each point shows the fraction of $\frac{B+C}{A+D}$ cases with darker/larger points showing higher values, in other words the prevalence of bandwidth-limited websites. The rows of maps shows data for different years of network performance, while the columns show data for different network providers.

In general we observe that instances of websites being bandwidth limited (darker points showing higher $\frac{B+C}{A+D}$ ratios) as local phenomena that vary by year and provider. In 2018, websites were bandwidth limited primarily in Oklahoma in the Sprint network and in different regions in the Verizon network. In 2019, websites were bandwidth limited to a larger degree across all providers and regions. In 2020, as network performance improved, bandwidth bottlenecks were more prevalent in North East and the South West regions of the country for all providers.


While providers in strive to improve network performance both in terms of latency and bandwidth, they may not be aware how different websites perform in their networks regionally. By understanding the prevalence of the different



(a) Network data from 2018 [15].



(b) Network data from 2019 [16].



(c) Network data from 2020 [3].

Fig. 5. A, B, C, and D cases across network providers.

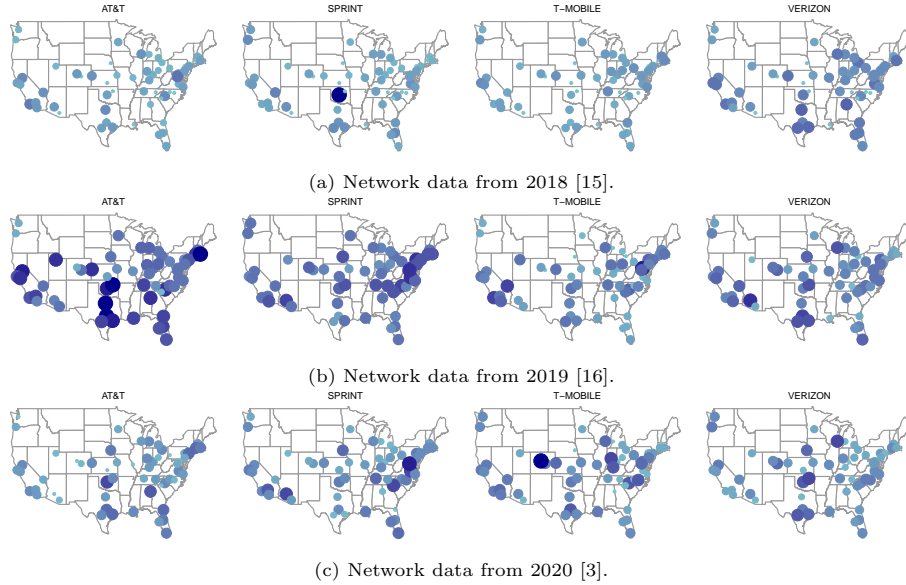


Fig. 6. Ratio of $\frac{B+C}{A+D}$ cases in US cities.

types of cases in different regional network deployments, providers may adjust network resource allocation to better meet website needs.

Do small changes to bandwidth and latency improve PSI? We also wanted to understand how small changes to network latency and bandwidth alone improve website performance. In Figure 7 we show the effect of network latency on the x-axis on PSI on the y-axis. In Figure 8 we show the effect of network bandwidth on the x-axis on PSI on the y-axis. The lines in the graphs show page CPIs.

In general we observe that as network performance improves so does the PSI. Somewhat contrary to results observed by Belshe and Grigorik [1, 2] we observe that both improvements in latency and improvements in bandwidth have a steady effect on page PSI. In other words we do not observe the diminishing improvements to page performance of additional bandwidth for all pages. This result is interesting in that as modern pages become more complex, there is no clear path to improving page performance by simply decreasing network latency. Instead a careful analysis of page CPI and network performance gives more insight.

What makes websites limited by bandwidth? We also wanted to understand what aspects of a page make it more likely to fall into the B and C cases. To do so, we use our measurement data to predict the ratio of B and C cases to A and D cases ($\frac{B+C}{A+D}$) using page characteristics. To construct

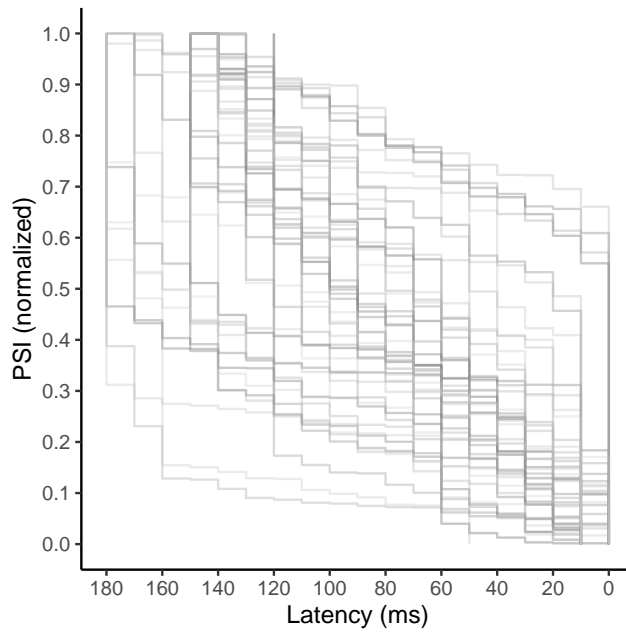


Fig. 7. PSI versus latency.




Fig. 8. PSI versus bandwidth.

the data set for our model we compute the $\frac{B+C}{A+D}$ ratio for each page, provider network, and year. Next, we compute page characteristics as the number of script (.js), Cascading Style Sheets (CSS) (.css), image (.png, .jpg, .gif, .svg), and text (.html, .json) files as well as their total size in KB, from HTTP Archive (HAR) files for each page. Based on these data we train a linear regression model (`sklearn.linear_model.LinearRegression`) using an 80/20 training/testing split in a 100 trials.

Factor	Coefficient Value
CSS count	0.261606
image count	0.041035
script count	0.037109
CSS KB	0.025595
text count	0.019713
text KB	0.002346
script KB	0.000432
image KB	0.000123

Table 1. Page factors in linear regression of the $\frac{B+C}{A+D}$ ratio.

On average we are able to predict the $\frac{B+C}{A+D}$ ratio withing about a factor of 2.4 on average. While the behavior of pages is complex and depends on their structure, this prediction nevertheless gives us some insight into the influence of the page characteristics that results in particular $\frac{B+C}{A+D}$ ratios. Table 1 shows the model factors (page characteristics) and their sorted coefficient values. We observe that the number of CSS, image, and script objects have the greatest effect on the $\frac{B+C}{A+D}$ ratio. When comparing page $\frac{B+C}{A+D}$ ratio to the number of CSS, image, and script objects on a page we observe that page/network combinations with high $\frac{B+C}{A+D}$ ratios tend to have more of these objects. However, there are also pages which fall into D cases in all networks with similar number of objects. Thus, we conclude that a more in-depth analysis of page structure, which results in higher $\frac{B+C}{A+D}$ ratios, should be a part of future work. In the meantime, we reiterate, that developers should measure the performance of their pages in different network conditions directly to understand if there are network scenarios where their pages are limited by bandwidth.

5 Discussion

Our results show a mismatch between website needs in terms of network performance, expressed through the CPI, and the network performance offered by network operators. Within each network performance envelope a website is limited by either insufficient bandwidth or latency. If a website is bandwidth-limited within a network performance envelope, it could improve user experience by re-

shaping its CPI through page optimization. Alternatively, if a website is latency-limited it could be redesigned to change its CPI and take advantage of the extra bandwidth available in a given network. Below we offer a discussion of techniques website developers could apply to their sites to change their CPI to better fit available network performance. While the optimization approaches discussed below are already used by website developers and CDNs, we think that grouping them by impact is of value to design strategies for reshaping page CPIs.

If a website is limited by high latency, website developers could leverage different CDNs to serve content from replica servers closer to users in a given network [20]. If possible, they could move origin logic onto the CDN for even faster delivery. Developers could also use Transport Layer Security 1.3 for faster connection setup times [21]. Finally, they could deliver their content using HTTP/3 over the QUIC protocol to tackle lossy networks and retransmission [22].

If a website is limited by low bandwidth, developers could use HTTP link headers to preload CSS and script assets critical for rendering embedded objects without having to wait for the browser to identify these resources and prioritize their HTTP requests [23]. This approach helps to reduce network idle time as the browser parses page HTML. Developers could also defer loading of non-critical assets after the onload event, which will improve page interactivity, or defer below-the-fold assets, since they do not contribute to lower the PSI. Developers may also adapt image sizes to the quality of the network. For example, Akamai Image Manage and Adaptive Image Compression dynamically adjusts image sizes in response to network quality [24, 25].

6 Conclusions and Future Work

We investigated the degree to which latency continues to be the limiting factor to mobile website performance. Through a novel webpage performance characterization method, Critical Path of Improvement (CPI), we show the sensitivity of websites to small changes in network bandwidth and latency performance. Comparing against measurements of mobile network performance we show instances where websites do suffer from a lack of bandwidth in mobile networks. Finally, we suggest techniques that web developers and network administrators can use to create better matches between website requirements and network performance by reshaping page CPI.

To be sure, CPI should be used to investigate factors other than bandwidth/latency on website performance. Two interesting directions are to investigate the impact of website design features on CPI and of web server configurations, including different network protocols used to serve web content. While we consider such investigations a part of future work and hope that CPI will be a useful tool there, we believe they are out of scope for this paper, which focuses on taking a fresh look at the predominance of latency as the limiting factor in website performance.

Another interesting direction for future investigations is to incorporate the effects of jitter and packet loss in CPI characterization. Jitter and loss could arise

due to packet queuing, mobile network scheduler decisions, and route switching. To keep our discussion focused, we do not explicitly control for these factors, choosing instead to observe them in wild though measurements of page PSI in live networks. However, CPI could be measured on axes other than latency/bandwidth and opens a path to including jitter and loss in future work.

Disclosure

The positions, strategies, or opinions reflected in this article are those of the authors and do not necessarily represent the positions, strategies, or opinions of Akamai.

References

1. Belshe, M.: More Bandwidth Doesn't Matter (Much). <http://bit.ly/2TxHv84> (Apr 2010)
2. Ilya Grigorik: Latency: The New Web Performance Bottleneck. <http://bit.ly/2PEfeeP> (Jul 2012)
3. USA Mobile Network Experience Report January 2020. <https://www.opensignal.com/reports/2020/01/usa/mobile-network-experience> (Jan 2020)
4. State of Mobile Networks: USA (February 2016). <https://www.opensignal.com/reports/2016/02/usa/state-of-the-mobile-network> (Feb 2016)
5. Gao, Q., Dey, P., Ahammad, P.: Perceived Performance of Top Retail Webpages In the Wild: Insights from Large-Scale Crowdsourcing of Above-the-Fold QoE. In: Workshop on QoE-Based Analysis and Management of Data Communication Networks (Aug 2017), <https://doi.org/10.1145/3098603.3098606>
6. Furtak, M., Wittie, M.P.: The ODDness of Webpages. In: Network Traffic Measurement and Analysis Conference (TMA) (Jun 2019)
7. Nejati, J., Balasubramanian, A.: An in-depth study of mobile browser performance. In: World Wide Web (Apr 2016), <https://doi.org/10.1145/2872427.2883014>
8. Pourghassemi, B., Amiri Sani, A., Chandramowliswaran, A.: What-if analysis of page load time in web browsers using causal profiling. *ACM Measurement and Analysis of Computing Systems* 3(2) (Jun 2019), <https://doi.org/10.1145/3341617.3326142>
9. Ren, J., Wang, X., Fang, J., Feng, Y., Zhu, D., Luo, Z., Zheng, J., Wang, Z.: Proteus: Network-aware web browsing on heterogeneous mobile systems. In: CoNEXT (Dec 2018), <https://doi.org/10.1145/3281411.3281422>
10. Liu, X., Ma, Y., Wang, X., Liu, Y., Xie, T., Huang, G.: SWAROVsky: Optimizing Resource Loading for Mobile Web Browsing. *IEEE Transactions on Mobile Computing* 16(10), 2941–2954 (Oct 2017)
11. Rajiullah, M., Lutu, A., Khatouni, A.S., Fida, M.R., Mellia, M., Brunstrom, A., Alay, O., Alfredsson, S., Mancuso, V.: Web Experience in Mobile Networks: Lessons from Two Million Page Visits. In: The World Wide Web Conference (2019), <https://doi.org/10.1145/3308558.3313606>
12. Han, B., Qian, F., Ji, L.: When Should We Surf the Mobile Web Using Both WiFi and Cellular? In: All Things Cellular (Oct 2016), <https://doi.org/10.1145/2980055.2980060>

13. NetEm. <https://www.linux.org/docs/man8/tc-netem.html> (Nov 2011)
14. Irish, P.: PWMetrics. <https://github.com/paulirish/pwmetrics> (Accessed Jan 2020)
15. State of Mobile Networks: USA (January 2018). <https://www.opensignal.com/reports/2018/01/usa/state-of-the-mobile-network> (Jan 2018)
16. USA Mobile Network Experience Report January 2019. <https://www.opensignal.com/reports/2019/01/usa/mobile-network-experience> (Jan 2019)
17. The top 500 sites on the web. <https://www.alexa.com/topsites> (Mar 2020)
18. Rohatgi, A.: WebPlotDigitizer. <https://automeris.io/WebPlotDigitizer> (Apr 2019)
19. Boyland, P.: Why confidence intervals are vital in analyzing mobile network experience. <https://www.opensignal.com/2019/02/13/why-confidence-intervals-are-vital-in-analyzing-mobile-network-experience> (Feb 2019)
20. Nygren, E., Sitaraman, R.K., Sun, J.: The akamai network: A platform for high-performance internet applications. SIGOPS Oper. Syst. Rev. 44(3), 2–19 (Aug 2010), <https://doi.org/10.1145/1842733.1842736>
21. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3. <https://tools.ietf.org/html/rfc8446> (Aug 2018)
22. Bishop, M.: Hypertext Transfer Protocol Version 3 (HTTP/3). <https://tools.ietf.org/html/draft-ietf-quic-http-27> (Feb 2020)
23. Osmani, A.: Preload, Prefetch And Priorities in Chrome. <https://medium.com/reloading/preload-prefetch-and-priorities-in-chrome-776165961bbf> (Mar 2017)
24. Akamai: Image Manager. <https://learn.akamai.com/en-us/webhelp/ion/oca/GUID-449AD93C-4E69-4C07-BA8B-FE93BFD4C39D.html> (Accessed Mar 2020)
25. Akamai: Adaptive Image Compression (AIC). <https://learn.akamai.com/en-us/webhelp/ion/oca/GUID-03E4FAD1-10EC-4496-AD44-812D9BD94DBF.html> (Accessed Mar 2020)