Effect of Web Caching on Load Time

http://dx.doi.org/10.3991/ijet.v8i6.2938

Emma Tang[1], Tae (Tom) Oh[1], Young B Choi[2] and Swapna Kalpagam Subramaniam[1]
[1]Rochester Institute of Technology, Rochester, USA
[2]Regent University, Virginia Beach, VA, USA

Abstract—**This paper explores the effect of web caching on load times in terms of browser choice and website type. By comparing the load time of web pages using different browsers and website types it was possible to determine the relationship between the variables. Load time differs between browsers based on how the browsers parse different kinds of data. It terms of website type, the traffic volume as well as frequency of updates by administrators and website users is responsible for the varying load times.**

Index Terms—**Browser, load, load-time, cache, non-cache, stale, speed, website, web, type, platform, measurement, documentation, performance, theory, reliability, experimentation, verification, human factors, standardization.**

I. INTRODUCTION

Web caching is the process of storing frequently visited website pages in memory to reduce bandwidth and load on a server. By caching website pages, the load time of these websites is significantly decreased. It is no longer necessary to retrieve frequently viewed files from the website server because they can be loaded directly from the cache [1]. Web caching is an important tool that allows high traffic websites or dynamically generated websites to run as though they were static. The required bandwidth for these websites can be reduced as well. As a result, the load times for each page are reduced significantly.

There are two main types of caches: proxy caches and browser caches [2]. Proxy caches retrieve content that is saved by the server and can be accessed by users accessing the server. This reduces the bandwidth of organizations working within the same server, accessing frequently requested resources. This paper focuses on browser caches as well as the Time to Live (TTL) of a website. A browser cache is lowest level cache that can be manipulated directly by the user. All content that is requested from a website either comes from the server itself or from an associated cache that contains copies of the original data. The risk that is taken when information is retrieved from a cache rather than from the original source is that the data will be 'stale' [3]. Stale data is data that does not reflect the most current state of the webpage. If the data on the server has been recently updated, but the cache still stores previous data, the users will end up loading a page with outdated data. The amount of time that has elapsed since the information was obtained from the server is considered the age of the webpage or information [2]. The greater the age, the more likely it is that the information is stale. In order to prevent stale web pages, all data has a distinct TTL [4]. Once the TTL has expired, the information will be cleared from the cache and will be reloaded from the server. The TTL of information differs from website to website based on the cache settings determined by the website developer.

The cache settings are set appropriately based on the type of website. For example, a website based on RSS Feeds that constantly combs the web for news content needs an extremely short TTL so that stale data does not stay in the cache but instead is replaced with more recent data. Websites such as this, have TTL values of a little as several seconds [5]. A personal website, which is updated rarely, probably has a longer TTL because regardless of the amount of time that the data has been in the cache, it will not be stale as long as the web developer has not updated the website.

This paper explores the load time for cached and non cached websites with a series of variables including browser and website type. In order to stay within a reasonable scope, this paper targets browsers that receive at least 3% of internet traffic based on the July, 2010 Browser Statistics collected by W3Schools and from Alexa's top 500 sites. Because website type plays a major role in the load time of websites, this paper looks at the load time for different types of websites such as a personal website; RSS Feed based website, and social networking website. By focusing on this range of websites, it is possible to determine the effect of website type on load time.

II. EXPERIMENTATION

A. Equipment

The equipment used in this experiment includes a stopwatch, a Windows XP platform desktop computer, and a Macintosh OSX laptop computer with add-ons installed to support the measurement of the caching time.

This experiment tests browser speeds on both a Windows and Macintosh platform; hence we need two different computers that make available the most popular browsers. As a result it is necessary to understand not only the differences in browsers but also in computer hardware. The speed of the computer is associated with several components of the computer including the CPU (Central Processing Unit), RAM, and the state of the hard drive. The CPU has four main tasks: fetch, decode, execute and write back; essentially retrieving the instructions, executing the instructions, and relaying feedback to make it clear that the instructions have been completed. The clock speeds as well as the l1 and l2 caches are responsible for the effect of the CPU on the speed of the machine. Because the computers that were used for this experiment varied in all of the components that reflect the speed of the computer, it is necessary to compare results within the category of platform rather than comparing all of the results at once. Also it is to be considered that each of the different browsers cache and store information in different ways. Each type of information such as photos, text, videos and others are handled independently.

B. Procedure

The experimentation presented in this paper utilizes the ability of the user to manually clear cache contents. For each variable tested, including browser and website type, there were fifty trials for both the cached and non-cached state. To test the cached state, a stopwatch was used to manually time the load time from the moment that the refresh button was pushed until the page was done loading. To reduce error associated with human reaction time, the average and standard deviations were calculated based on the fifty trials conducted. To test the non-cached state, the procedure was repeated, but the cache was manually cleared in between each trial. The data collected is displayed in the results section below.

C. Variables

In this experiment two different variables were chosen that would affect the load time. The first variable is the browser on which the user views the webpage. To narrow down the scope of this paper, web pages were chosen that are used by at least 3% of the population based on w3schools and Alexa's top 500 global sites. The browser statistics showed that the most popular browsers in 2012 with at least 3% of users included: Browser1, Browser2, Browser3, Browser4, and Browser5 [7]. Browser names are not revealed to protect the company's reputation. Because these browsers are used on different platforms, Browser1, Browser2 and Browser3 were tested on a PC platform [Windows Vista] and Browser4 and Browser5 were tested on a Macintosh platform.

Few major website types were chosen including a personal website, an RSS Feed based website, and a social networking website and video hosting site. The exact websites that were chosen were http://www.emmatang.com, http://www.nytimes.com, and http://www.facebook.com respectively.

Various pages of each of the sites were chosen to study and analyzing the caching and its performance effects. The cached and non-cached states of these websites were tested on Browser2 on a windows platform to maintain consistency.

These websites cover a broad spectrum of website types. http://www.emmatang.com is a personal website with basic HTML programming that does not receive much traffic and has a long TTL. http://www.nytimes.com is based on an RSS Feed and is constantly automatically being updated with breaking news and therefore has a relatively short TTL. http://www.facebook.com is a social networking site that is updated constantly by the users. Users are able to update profile information, statuses, and photo albums among other areas of updates at any time. Owing to this frequent and unmonitored system of updating by the user, the TTL is almost nonexistent because the amount of time before information is stale is random, based on the millions of facebook users. Also www.youtube.com is a site that hosts information in the form of videos and is accessed regularly by millions of users. Other sites that were considered include www.google.com and www.espn.com. The list of websites that were considered was sorted into certain categories based on their content and usage.

III. RESULTS

Tables 1 and 2 show the results based on the two variables: browser and website type. Within the browser variable, two different computers were used in order to test the browsers that are only available on one type of platform including Browser1 (Windows) and Browser3 (Windows). In order to evaluate them separately, each was tested on individual systems. The three tables discussed in the paper together have classified the data by the following categories: **Browser for Macintosh platform, Browser for Windows platform, Website Type**.

A. Browser Results

To keep a minimum number of variables, the load time was recorded using the same website page. The page that that was used for this experiment was the NY Times home page, www.nytimes.com. This page has dynamically generated content from an RSS feed as well as blogs and advertisements. From Figure 1 and Figure 2, it is clear that the cached pages loaded faster than the non-cached pages regardless of browser type or platform. The cached page stores the data, making it possible for the page to load without returning to the server to gather data. Therefore, these results are consistent with the original expectations. When comparing the Macintosh browsers to each other, the average load time for Browser4 was the fastest, followed by Browser5. For the Windows browsers, Browser3 was the fastest, followed by Browser1 and then Browser2. Table 1 and Table 2 show the average of the trials, the standard deviation of the trials, and the ratio of cached time to non cached time for Macintosh and Windows platforms respectively. By calculating the average, it gives a clearer picture of the load times for the cached versus non cached status of the web page. The standard deviation shows the range of data and how much the data has strayed from the average. The ratio of cached to non cached load times makes it possible to compare the differences in load time between browsers.

1) Browser Results: Average

Looking at the results of the average load times for both Macintosh and Windows platforms, it is clear to see that the cached load times were measurably faster than the non-cached load times. Overall, on the Macintosh platform, Browser4 showed the fastest load times while on the Windows platform, Browser3 showed the fastest load times. These results apply for both cached and non-cached load times. Browser4 loads the document along with the required scripts and style information ahead of time, so that they are available faster, with significantly less waiting time. With regards to the Windows platform, one of the explanations for the slower load time for Browser2 and Browser1 is that both browsers have limits regarding the number of connections per hostname when talking to HTTP/1.1 servers. Browser2 in general is laden with additional add-ons and developer tools which significantly increase the load time for the browser on either platform.

2) Browser Results: Standard Deviation

The standard deviation of the load times shows the consistency of the data. For the Macintosh platform, as the load times increased, so did the standard deviation. For the Windows platform results, the category with the highest

TABLE I.
CALCULATIONS FOR VARIABLE: BROWSER, PLATFORM MACINTOSH

	Browser4(mac) cached(sec)	Browser4(mac) non-cached(sec)	Browser2(mac) cached(sec)	Browser2(mac) non-cached(sec)	Browser5(mac) cached(sec)	Browser5(mac) non-cached(sec)
average	3.69	4.06	4.65	4.96	5.13	6.7
standard dev.	0.11	0.18	0.35	0.33	0.49	0.45
ratio	0.91	...	0.94	...	0.77	...

TABLE II. CALCULATIONS FOR VARIABLE: BROWSER, PLATFORM WINDOWS

	Browser1(win) cached(sec)	Browser1(win) non-cached(sec)	Browser2(win) cached(sec)	Browser2(win) non-cached(sec)	Browser3(win) cached(sec)	Browser3(win) non-cached(sec)
average	3.71	4.71	4.86	5.84	3.11	3.26
standard dev.	0.49	0.28	0.21	0.22	0.12	0.12
ratio	0.79	...	0.83	...	0.95	...

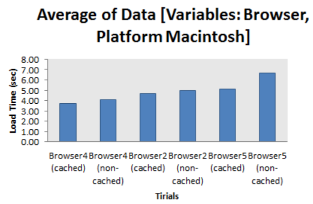

Figure 3: Average Macintosh

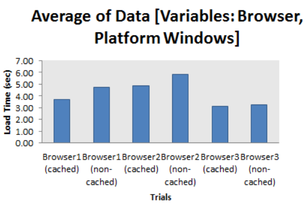

Figure 4: Average Windows

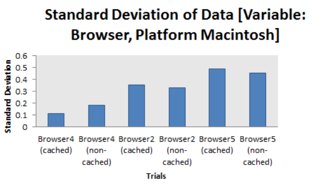

Figure5. Standard Deviation: Macintosh

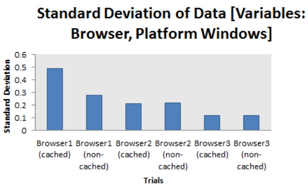

Figure 6. Standard Deviation: Windows

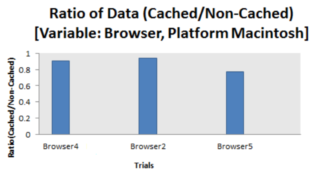

Figure 7. Ratio: Macintosh

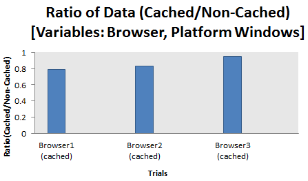

Figure 8. Ratio: Windows

standard deviation was the cached Browser1 data which nearly doubled the standard deviation of the non-cached Browser1 data. There is no consistent conclusion regarding the standard deviation being higher or lower based on cached or non-cached data. This data shows that in some categories, the cached load times had a higher standard deviation than the non-cached load times while in other categories the cached load times had a lower standard deviation than the non-cached load times.

3) *Browser Results: Ratio*

The ratio table shows the relationship between the cached to non-cached data with the following relationship:

ratio = cached/non-cached

The ratio data ranges from 0.77 to 0.95. This is based on the sample testing performed across Browser1 and Browser5, the ratio of cached to non-cached data correlates that the performance variation is justified given the latent differences.

B. *Website Type Results*

Load time is significantly affected by the type of website. If a website contains many photos with large file sizes or receive a high amount of traffic, the load time will be increased. In order to test the effect of website type on load time, three different types of websites were chosen and tested for the load times of those websites on one browser. Only one browser was tested in order to minimize the number of variables and to isolate website type as the sole cause of load time differences.

TABLE III.
CALCULATIONS FOR VARIABLE: WEBSITE TYPE

Website type	Website	Empty cache	Average	Primed cache	Average
Video-intensive	http://www.youtube.com	123	123	115	115
	http://www.ted.com	1652	1652	991	991
	http://www.dailymotion.com	4277	4277	3230	3230
Photo-intensive	http://www.flickr.com	480	480	320	320
	http://www.photo.net	5964	5964	2392	2392
Blog and basic sites	http://www.emmatang.com	1200	1200	197	197
	http://www.wordpress.com	3004	3004	591	591
Mail sites	http://www.gmail.com	503	503	174	174
	http://www.yahoomail.com	780	780	321	321
File sharing sites	http://www.mediashare.com	3086	3086	447	447
	http://www.fileshare.com	1371	1371	555	555
Search engine and web hosting sites	http://www.wix.com/lifeatrit/swapnaks	762	762	461	461
	http://www.google.com	142	142	135	135
E-commerce sites	http://www.bankofamerica.com	1355	1355	447	447
	http://www.amazon.com	1581	1581	1010	1010
RSS feeds sites	http://www.espn.com	2643	2643	2479	2479
	http://www.washingtonpost.com	3978	3978	1343	1343
	http://www.scribd.com	2177	2177	2068	2068
Social networking sites	http://www.facebook.com/userprofile	7536	7536	3935	3935
	http://www.linkedin.com	3071	3071	2776	2776

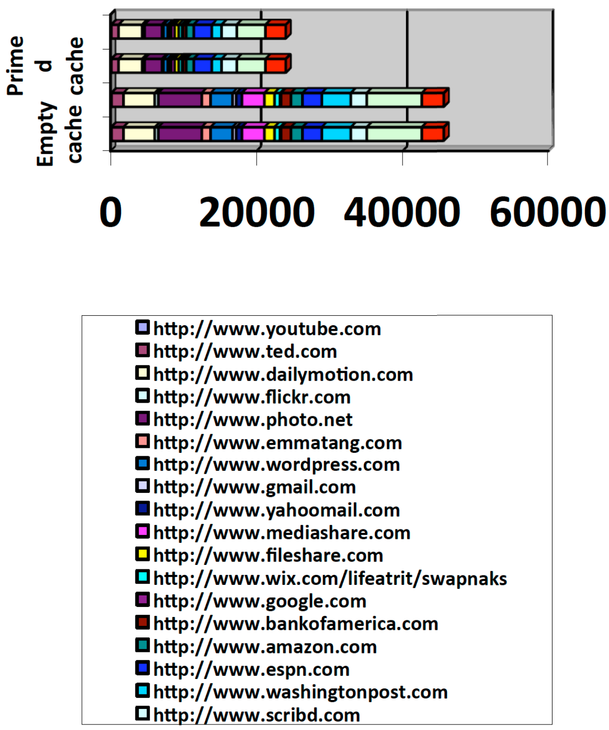

Figure 10: Variable website type

Various websites were chosen based on specific categories and they were tested in order to ensure a clear observation of the performance. The websites that were considered included personal website; http://www.emmatang.com, a website based on an RSS feed; http://www.nytimes.com, and a social networking website; http://www.facebook.com, a website dedicated to images http://www.photo.net, e-commerce website such as http://www.amazon.com among others. Table 3 represents the data from these three website that were tested on the Browser2 browser on a Windows platform. Table 4 represents the average, standard deviation, and ratio of the data.

There is a clear observation of the sites which are the fastest to load in cached and non-cached scenarios owing to the components they support during load and the contents that are cached. These results can be seen graphically in Figure 9.

The websites have been classified into the following areas in order to have a collective observation of the various frames, contents and application that websites support during load and the kind of information that gets cached deciding the load time of the site. They are Basic sites: sites that contain blog information, or personal sites, website hosting sites, File hosting sites: sites like www.filestube.com and others that support content sharing, Social Networking sites: Sites like www.facebook.com, www.linkedin.com that comprise of various type of user information and is updated frequently,

Image-intensive sites: Sites like www.photo.net and www.flickr.com that host images and majorly dedicated to them, E-commerce sites:www.amazon.com and banking sites that generate lot of traffic, RSS feeds: sites with regular content update through RSS like www.nytimes.com, www.espn.com that keep updating, Video-intensive sites: sites like www.youtube.com, www.dailymotion.com that have huge user base and load video content in cache, e-mail sites: google and yahoo mails and other random sites.

The TTL settings are set individually by the website developer in each of these sites. The settings determine how much time passes before the data within the cache is updated. The parameters for website caching are heavily time dependent. The TTL specifications can be fixed and independent - for example, every hour. With this capacity, it does not matter how many times the user visits the site or opens up a new browser. If the user wants to open an updated version of the site every five minutes, he would need to manually clear the cache. The developer can set the TTL specifications to be dependent on the user - for example, every time the user opens a new browser.

This TTL setting varies depending on user tendencies. Setting TTL specifications is not trivial, and determines careful consideration by the developer depending on the use of the website. For instance, because the NY Times website is updated frequently, the TTL is shorter than http://www.emmatang.com, which is rarely updated. http://www.facebook.com is updated by millions of users every second. Every time the facebook page is refreshed, there is different content on the page. As a result, the TTL is extremely minimal. There is content that is cached in the system, including user profile information, friends, and applications, but updates including the new feed and status updates are not cached, to allow for the continual flow of new information. If that content were only refreshed every day, there is no doubt that the content would not be stale.

Based on the expected traffic, the allotted bandwidth for websites differs. According to Nielson's study, http://www.facebook.com had the most traffic in a certain month, with over sixty million visitors [8]. Nielson's study also showed that http://www.nytimes.com received over eighteen million visitors [9]. Google Analytics showed that http://www.emmatang.com received over fifty visitors in same month [10]. This significant traffic difference plays a major role in the load time of websites, with higher traffic responsible for slower load times. That said, traffic is not the only factor for load times. The data in Table 3 shows that despite the higher traffic of http://www.facebook.com, it had a higher average load time than http://www.nytimes.com. The content being loaded also play an important role in determining the load time, the sites which supported all types of contents relatively take more time to load which in an inherent behavior.

1) Website Type Results: Average

The average load times for the various website types show that the website type with the longest load time is the RSS Feed based website follows by the social networking website and lastly the personal website. Because the traffic of the website is not the only cause for different load times, as seen previously with the traffic versus load time discrepancy, it must be attributed to other factors. These factors could include the information on the page including the number and size of images, the way that the browser parses the JavaScript code or PHP code on the page or the dynamic quality of the data.

2) Website Type Results: Ratio

The ratio table shows the relationship between the cached to non-cached data with the following relationship:

$$\text{ratio} = \text{cached}/\text{non-cached}$$

The ratio data ranges from 0.78 to 0.98.

The ratio 0.98 for http://www.facebook.com shows that the cached versus non-cached load times were extremely similar. Because there are hundreds of millions of users who update their profile information every day, the TTL for data is close to zero. Within seconds of loading a page on the social networking website, the information is stale and needs to be reloaded from the server. It is probable that the cache settings for http://www.facebook.com are set to reload the page from the server rather than from information saved in a cache. As a result, the load time for the cached versus non-cached data was virtually the same. Although the cache was not cleared manually in between trials for the cached category, the cache settings for the website cleared it automatically.

In contrast, the ratio results for http://www.emmatang.com was 0.78, indicating that there was a large difference in load time between cached versus non-cached settings. Because the website is so rarely updated, the cache settings are not set to automatically clear the cache frequently. A large amount of data is stored in the cache and remains there for a much longer period of time than either http://www.facebook.com or http://www.nytimes.com. Therefore, when the cache is cleared and the data must be retrieved from the original server, it takes much longer for the process to be completed.

The ratio for http://www.nytimes.com was in the middle of the other two website types at 0.94. This makes sense because while the website does need to be updated daily to keep current, there are a significantly smaller number of users with administrative access as well as a smaller number of daily return users.

C. Error Analysis

In order to precisely analyze the data from this experiment, it is necessary to take into account the human error associated with the procedure. This experiment was based on recording the load time of websites from different browsers and websites of different types. The experiment relied on human reaction time and used a stop watch to record the load times. Based on an online reaction time test, the average reaction time calculated for determining error is .26s [11].

In order to account for the error associated with the initial reaction time as well as the end reaction time, the following formula was used [12]:

$$\partial q = \sqrt{[(\partial x)^2+\ldots+(\partial z)^2+\ldots+(\partial w)^2]}$$

Because there were two instances of human error - at the beginning and the end - there are two independent variables. Taking into account the following function:

$$F(x,y) = x - y$$

$$\Delta F = \sqrt{[(\partial f/\partial x)^2 \Delta x^2 + (\partial f/\partial y)^2 \Delta y^2]}$$

(ΔF is the error associated with each individual trial)

$$\Delta F = \sqrt{[(1)^2(.26)^2 + (-1)^2(.26)^2]}$$

$$\Delta F = \sqrt{[.1352]}$$

$$\Delta F = 0.37s$$

Based on this error analysis, there is a window of error of 0.37seconds for the individual data taken.

To determine the precision of the averages that were calculated from the individual data, another formula must be followed [10]:

$$\sigma_{avg(x)} = \sigma_x / \sqrt{N}$$

(N is the number of data points)

This formula calculates the standard deviation of the average value of the data taken. It is safe to assume that the average of the data is more reliable than any one individual piece of data. Following that logic, the standard deviation of the average value should be more precise than the standard deviation of the individual trials. Table 4 shows the average standard deviation for each data category.

TABLE IV.
STANDARD DEVIATION OF THE AVERAGE

	σ_x	$\sigma_{avg(x)}$
Browser4 - mac (cached)	0.11	**0.02**
Browser4 - mac (non-cached)	0.18	**0.03**
Browser5 - mac (cached)	0.49	**0.07**
Browser5 - mac (non-cached)	0.45	**0.06**
Browser1 - win (cached)	0.49	**0.07**
Browser1 - win (non-cached)	0.28	**0.04**
Browser2 - win (cached)	0.21	**0.03**
Browser2 - win (non-cached)	0.22	**0.03**
Browser3 - win (cached)	0.12	**0.02**
Browser3 - win (non-cached)	0.12	**0.02**

Another source of error associated with the experiment was, regarding the browser as the variable. Each browser has its own graphical user interface and its own indication of when the browser is done loading. In some cases, the browser waits until the page is completely done loading before indicating that it is done. In other cases, the browser indicates that it is done loading at different stages of loading. Browser1 on Windows platform for example, clearly states the word "Done" at the bottom of the browser when one phase of the page is done loading, and then continues to load, removing the word "Done" from the bottom of the browser, and then shows it again when the page is completely finished. There is no way to differentiate between the completion of a phase and the completion of the entire phase. As a result, the reaction time associated with the load time of the Browser1 browser may be higher. The same is true with Browser5 on a Macintosh platform.

Another source of possible error can be attributed to the tools/add-ons that were used for measurement of the time. The same tool/add-on was not used across the different browser platforms, hence there is scope for the measurement being slight off by 10-15 milliseconds for a same site. It was however verified by manually collecting the information again through stopwatch. The Browser2 browser supported firebug which comprised of Hammerhead and Yslow add-on tools that helped collect empty cache and primed cache data for a website along with statistics of cache content through YSlow. This helped attain a clear information about the caching being done and how it affects the load time. It showed that load time considerably reduced for primed cache entry and how the statistics affirmed the same based on which content was refreshed from the server and not from local cached entry. Each browser had its own way of caching data of sites and this has an impact on the load time as well. Also automated stopwatch macro was used for measurement of load time on mac platform offered by www.numion.com website.

IV. DISCUSSION

All browsers are in the race for the title of 'Fastest Browser.' The data to back up this claim, however, has not definitively been released by any of the browsers. Browser5 claims to be the "Faster and Safer Internet" while Browser4 has self-titled itself "The World's Fastest Browser" [14]. But what does "fastest" really mean? In this paper, the word "fastest" was examined in the context of cached and non-cached web pages from a variety of browsers and website types. Based on the results, it is clear that browsers are not equal in terms of load time. When looking at a Macintosh platform, Browser4 was the fastest, followed by Browser5. This held true for both the cached and non-cached web pages. For a windows platform, the fastest browser was Browser3, followed by Browser1 and then Browser2.

When looking at website types. the results showed that the pages with minimal dynamic content loaded the fastest. Also, because the personal website was rarely updated and the cache settings were set to have a relatively long TTL, the cached load times were significantly shorter than the non-cached load times.

While the personal website loaded faster, the probability that the information displayed on the webpage would be stale was higher. This introduces the question of speed versus reliability. Is it worth it to have a faster load time if the information on the page might not be the most recent? This question is best answered by looking at the data recorded in this paper. While it is not possible to generalize the worth of current data, it can be assumed that for websites such as social networking sites and RSS Feed based sites, current information trumps speed.

A. Mobile browsers:

This experiment does not take into account the browsers running on mobiles, although they are same, they perform differently owing to the processor speed, operating system, display restrictions and cache memory available based on the device. A casual verification of this experiment on mobiles also showed similar results with safari showing considerable speed on iphones, while android based phones support Fennec(aka Browser2) and Browser5 among others, symbian or windows based phone support Browser1 and other browsers. There is a lot of scope for improvement of web browsing speed specific to load time concerning the caching efficiency. There has been numerous research works that focus on such specifics.

Thus it can be inferred that the web caching techniques currently deployed need more organization and the current browsers keep working on it to improve the overall performance efficiency. However, the load time is majorly dependent on the type of the site, content shared, device of use and its technical stipulations not to mention the traffic on the system.

REFERENCES

[1] Wang, Jia. “A Survey of Web Caching Schemes for The Internet.” ACM SIGCOMM Computer Communication Review 29.5 (1999): 36-46. ACM Digital Library. Web. 1 Dec. 2009. <http://portal.acm.org.ezproxy.rit.edu/citation.cfm?id=505696.505701&coll=portal&dl=ACM&CFID=66193779&CFTOKEN=27115967>.

[2] Yu, Haobo, Lee Breslau, and Scott Shenker. “A Scalable Web Cache Consistency Architecture.” Applications, Technologies, Architectures, and Protocols for Computer Communication (1999): 163-74. ACM Digital Library. Web. 22 Nov. 2009. http://portal.acm.org.ezproxy.rit.edu/citation.cfm?id=316188.316219&coll=portal&dl=ACM&CFID=66193779&CFTOKEN=27115967&

[3] Cohen, Edith, and Haim Kaplan. “Aging Through Cascaded Caches: Performance Issues in the Distribution of Web Content.” Applications, Technologies, Architectures, and Protocols for Computer Communication (2001): 41-53. ACM Digital Library. Web. 27 Nov. 2009. http://portal.acm.org.ezproxy.rit.edu/citation.cfm?id=383059.383063&coll=portal&dl=ACM&CFID=66193779&CFTOKEN=27115967

[4] Yin, Jian, Lorenzo Alvisi, Mike Dahlin, and Arun Lyengar. “Engineering Web Cache Consistency.” ACM Transactions on Internet Technology (TOIT) 2.3 (2002): 224-59. ACM Digital Library. Web. 23 Nov. 2009. http://portal.acm.org.ezproxy.rit.edu/citation.cfm?id=572326.572329&coll=portal&dl=ACM&CFID=66193779&CFTOKEN=27115967

[5] Jung, Jaeyeon, Emil Sit, Hari Balakrishnan, and Robert Morris. “DNS Performance and the Effectiveness of Caching.” IEEE/ACM Transactions on Networking (TON) 10.5 (2002): 589-603.\ ACM Digital Library. Web. 3 Dec. 2009. <http://portal.acm.org.ezproxy.rit.edu/citation.cfm?id=581876.581877&coll=portal&dl=ACM&CFID=66193779&CFTOKEN=27115967>.

[6] “Browser Statistics: Web Statistics and Trends” W3Schools.com. Web. July 2010. <http://www.w3schools.com/browsers/browsers_stats.asp>

[7] “Browser Statistics.” W3Schools Online Web Tutorials. 2009. Web. 02 Dec. 2009. <http://www.w3schools.com/browsers/browsers_stats.asp>

[8] “Web Traffic to Career Development Sites Increases 20 Percent Year-Overyear in January, According to Nielson Online.” (2009). Nielson Online. The Nielson Company. Web. 1 Dec. 2009.

[9] “Web Traffic To Top 10 Online Newspapers Grows 16 Percent In December.” NYTimes.com | Nielsen Wire. 27 Jan. 2009. Web. 02 Dec. 2009. <http://blog.nielsen.com/nielsenwire/tag/nytimescom/>.

[10] “Overview: www.emmatang.com.” Google Analytics. Web. 02 Dec. 2009. <https://www.google.com/analytics/settings/home?scid=11943266>.

[11] “Human Benchmark - Reaction Time Test.” Human Benchmark. 2009. Web. 05 Dec. 2009. <http://www.humanbenchmark.com/tests/reactiontime/index.php>.

[12] Taylor, John R. An Introduction to Error Analysis: The Study of Uncertainties in Physical Measurements. 2nd ed. Sausalito: University Science Books, 1997. Print.

[13] Opera browser | Faster & safer internet | Free download. Web. 3 Dec. 2009. <http://www.opera.com>.

[14] “Safari - Download Safari - Download worlds fastest and most innovative browser for Mac and PC.” Apple. Web. 07 Dec. 2009. <http://www.apple.com/safari>.

[15] Numion stopwatch for measuring website load time. <http://numion.com/Stopwatch/index.html>

[16] Add-on for web browsers, Firebug and Hammer head < http://stevesouders.com/hammerhead/> <http://getfirebug.com/ >

[17] Ramya Karri, "Client side page element web-caching", May 7 2009 <http://scholarworks.sjsu.edu/cgi/viewcontent.cgi?article=1136&context=etd_projects >

AUTHORS

Emma Tang is with the Department of Information Technology, B. Thomas Golisano College of Computing and Information Technology, Rochester Institute of Technology, One Lomb Memorial Drive, Rochester, NY 14623-5603, ext8244@rit.edu

Tae (Tom) Oh is with the Dept. of Information Sciences and Technologies, B. Thomas Golisano College of Computing and Information Services, Rochester Institute of Technology, Rochester, NY 14623-5603, USA tom.oh@rit.edu

Young B Choi is Associate Professor of Information Systems Technology, Department of Natural Science, Mathematics & Technology, School of Undergraduate Studies, RH 464, Regent University, Virginia Beach, Virginia 23464-9800 , USA.

Swapna Kalpagam Subramaniam is with the Dept. of Networking, Security and Systems Administration, B. Thomas Golisano College of Computing and Information Services, Rochester Institute of Technology, Rochester, NY 14623-5603, USA, E-mail: sxk8430@rit.edu

This article is an extended and modified version of a paper presented at the 2011 Interdisciplinary Conference of AHLiST, held May 20-22, 2011, at College of Technology, University of Houston, TX, USA. Manuscript received 09 October 2005. Published as re-submitted by the authors 08 November 2013.