



- (51) **International Patent Classification:**
G06Q 50/10 (2012.01)
- (21) **International Application Number:**
PCT/US2012/042186
- (22) **International Filing Date:**
13 June 2012 (13.06.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/497,056 14 June 2011 (14.06.2011) US
13/175,115 1 July 2011 (01.07.2011) US
- (71) **Applicant (for all designated States except US):**
GOOGLE INC. [US/US]; 1600 Amphitheatre Parkway,
Mountain View, CA 94043 (US).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **BURKARD, Timo**
[DE/US]; 55 Page Street, Apt. 726, San Francisco, CA
94102 (US). **JAIN, Arvind** [US/US]; 996 Miramonte Av-
enue, Los Altos, CA 94024 (US). **BENTZEL, Christoph-**
er [US/US]; 27 Wesson Avenue, Quincy, MA 02169 (US).
- (74) **Agents:** **BRABHAM, Cicero, H.** et al.; Lerner, David
Littberg, Krumholz & Mentlik, LLP, 600 South Avenue
West, Westfield, NJ 07090 (US).
- (81) **Designated States (unless otherwise indicated, for every
kind of national protection available):** AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD,
SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR,
TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) **Designated States (unless otherwise indicated, for every
kind of regional protection available):** ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) **Title:** IMPROVING ACCESS TO NETWORK CONTENT




FIGURE 2

(57) **Abstract:** A method 300 and system 200 for improving access to network content are described. Aspects of the disclosure minimize the delay between a navigation event and a network response by prerendering the next navigation event. For example, the method and system may predict 302 a likely next uniform resource locator during web browsing to preemptively request 304 content from the network 306 before the user selects 308 the corresponding link on a web page 312. The methods and systems describe a variety of manners for prerendering content 400 and managing 500 and configuring prerendering operations 600, 700.

IMPROVING ACCESS TO NETWORK CONTENT

CROSS-REFERENCE TO RELATED APPLICATION

[0001] The present application is a continuation of U.S. Patent Application No. 13/175,115, filed on July 1, 2011, which claims the benefit of the filing date of U.S. Provisional Patent Application No. 61/497,056 filed June 14, 2011, the disclosure of which is hereby incorporated herein by reference.

BACKGROUND

[0002] The advent of the World Wide Web has placed more information at the fingertips of today's users than ever before. Various websites cater to nearly every need and interest, providing access to reference information, business and financial documents, social networking, and more. Widespread broadband Internet access provides faster access to these sites than ever before.

[0003] However, as fast as current high-speed Internet services are, the act of browsing the web is not instantaneous. When a user selects a link on a page or enters a uniform resource locator (URL) in a text field, there is a delay while data is requested from the host, sent to the client, and rendered in the browser. The user is typically idle while waiting for their requested site to load. While high-speed Internet access may limit this delay to a few seconds, even this short delay can add up to thousands of man-hours of lost productivity each year.

BRIEF SUMMARY

[0004] A method and system for improving access to network content are described. Aspects of this disclosure minimize the delay in accessing web content by identifying (e.g. predicting) a user navigation event on a web page. Aspects of the disclosure then request content associated with the next navigation event and prerender the content in a hidden application instance. When the user actually selects the navigation event, the hidden application instance is displayed to the user, minimizing the delay the user experiences.

[0005] An aspect of the disclosure describes a computer-implemented method for managing website analytics for prerendered content. The method includes receiving notification that a network content request is a prerender request for a set of network content, redacting, using a processor, a subset of content from the network content to create a set of partial content such that the set of partial content comprises the network content with the redacted content removed, providing the set of partial content in response to the network content request, receiving notification that the redacted content was accessed by a user device, and sending the redacted content in response to the notification that the redacted content was accessed. In some aspects, the method further includes updating an impression statistic in response to the notification that the redacted content was accessed. The redacted content may include one or more advertisements. The method may further include inserting one or more hypertext transfer protocol (HTTP) error codes within the set of partial content in place of the redacted content. The set of redacted content may include at least one image, script, cookie, or executable instruction. The redacting stage may be performed by not providing the redacted content in response to a request from a client device for the redacted content.

[0006] Another aspect of the disclosure provides a computer-implemented method for managing website analytics for prerendered content. The method includes receiving notification that a network content request is a prerender request for a set of network content, providing the network content in response to the network content request, receiving notification that the network content was viewed after prerendering, and updating an impression statistic in response to the notification that the network content was viewed such that the impression statistic reflects a count of users who have actually viewed the content. The network content may include at least one advertisement, and the impression statistic may be associated with the at least one

advertisement. The impression statistic may be associated with at least one image, script, cookie, or executable instruction.

[0007] Another aspect of the disclosure provides a processing system for managing website analytics for prerendered content. The processing system includes at least one processor, a prerender module associated with the at least one processor, and a memory coupled to the at least one processor. The prerender module is configured to receive notification that a network content request is a prerender request for a set of network content, redact a subset of content from the network content to create a set of partial content, such that the set of partial content comprises the network content with the redacted content removed, provide the set of partial content in response to the network content request, receive a notification that the redacted content was accessed, and send the redacted content in response to the notification that the redacted content was accessed. The redacted content may include at least one advertisement. The prerender module may be further configured to update an impression statistic in response to the notification. The impression statistic may be associated with the redacted content. The prerender module may be further configured to provide a hypertext transfer protocol error code in place of the redacted content.

[0008] Another aspect of the invention provides a processing system for serving network content and managing prerender requests. The processing system includes at least one processor, a prerender module associated with the at least one processor, and a memory coupled to the at least one processor. The prerender module is configured to determine whether an incoming network request is a prerender request. In some aspects, the prerender module is further configured to, in response to the incoming network request being identified as a prerender request, not provide access to the network content. The prerender module may be further configured to provide a hypertext markup language metatag to indicate that the

processing system does not provide prerender content. The prerender module may be further configured to measure a system utilization, and, in response to the system utilization being below a threshold value, provide access to the network content requested by the incoming network request. The system utilization may include at least one of a processor utilization, a network bandwidth, or an available system memory. The prerender module may be further configured to receive a notification that a user wishes to access content previously requested via a prerender request, and to provide web content in response to the notification.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] Figure 1 is a system diagram depicting an example of a server in communication with example client devices in accordance with aspects of the disclosure.

[0010] Figure 2 is a block diagram depicting an example of a computing device in accordance with aspects of the disclosure.

[0011] Figure 3 is a flow diagram depicting an example method for prerendering a web page in accordance with aspects of the disclosure.

[0012] Figure 4 is a flow diagram depicting an example method for indicating a navigation event to prerender in accordance with aspects of the disclosure.

[0013] Figure 5 is a flow diagram depicting an example method for managing web impression statistics during a prerender operation in accordance with aspects of the disclosure.

[0014] Figure 6 is a flow diagram depicting an example method for managing a prerender operation in accordance with aspects of the disclosure.

[0015] Figure 7 is a flow diagram depicting an example method for configuring a prerender operation in accordance with aspects of the disclosure.

DETAILED DESCRIPTION

[0016] Embodiments of a system and method for improving access to network content are described herein. Aspects of the disclosure minimize the delay between a user's selection of a navigation event and display of the content associated with the navigation event by prerendering content associated with the navigation event prior to the user's selection of the event. For example, the method and system may predict a likely next uniform resource locator during web browsing to preemptively request content from the network before the user selects the corresponding link, thus minimizing the wait time when a user selects a hyperlink on a web page. Various methods for performing the prerendering operation, configuring the prerendering operation, and managing the prerendering operation are described. For the purposes of this application, the term "prerendering" generally refers to the act of requesting all resources (e.g., any executable code, scripting language, files, and/or interactive objects) necessary to load the content of a web address, and loading the content in a web browser instance.

[0017] As shown in Figure 1, an example system 102 in accordance with one aspect includes a server 104 in communication with one or more client devices 106, 108, 110 displaying web browser interfaces 114, 116, 118. The user of the client device 106 may browse to a web page located at "www.a.com" as displayed on the web browser interface 114. That page includes content selectable by the user.

[0018] The client devices may be any device capable of managing data requests via a network 112. Examples of such client devices include a personal computer (PC) 108, a mobile device 110, or a server 106. The client devices 106, 108, 110 are operable to perform prerendering operations during the execution of a web browser application. The server 104 may transmit navigation history data to the client devices 106, 108, 110, to enable prediction of a next navigation event. In some aspects, the client devices 106, 108, 110 determine a next

navigation event using a local navigation history and generate a web request to the server 104 to prerender the content associated with the next navigation event. The prerender operation may also be determined dynamically by, for example, a JAVASCRIPT code. For example, a prerendering operation may occur when the user places a cursor over a link. For example, the user of the client device 106 may browse to a web page located at "www.a.com" as displayed on the web browser interface 114. That page includes content selectable by the user. Based on the user's navigation history, the client device 106 may determine which of the selectable content the user is likely to select, and then prerender the content associated with the selectable content by requesting the content from the server 104.

[0019] As another example, the client device 108 may display www.a.com within a browser interface 116. The client device 108 may receive an aggregate set of navigation statistics from the server 104, and then determine which selectable content the user is likely to select based upon the aggregate set of navigation statistics. As yet another example, the client device 110 may display www.a.com within a browser interface 118. The client device 108 may determine which selectable content the user is likely to select based upon a cursor position within the browser interface 116.

[0020] While the concepts described herein are generally described with respect to a web browser interface 114, 116, 118 executing on a client device 106, 108, 110, aspects of the disclosure can also be applied to any computing node capable of managing navigation events over a network, including a server 104.

[0021] The client devices 106, 108, 110 may also comprise personal computers, personal digital assistants ("PDA"): tablet PCs, netbooks, laptops, mobile phones, etc. Indeed, client devices in accordance with the systems and methods described herein may comprise any device operative to process instructions

and transmit data to and from humans and other computers including general purpose computers, network computers lacking local storage capability, etc.

[0022] The client devices 106, 108, 110 may be operable to predict navigation events to assist in data access on the network 112. For example, the client devices may predict a likely navigation event to facilitate prerendering of a web page in order to improve the user's browsing experience. In some aspects, the server 104 provides navigation data which may be used by the client devices 106, 108, 110 to predict a likely navigation event. In some aspects, the client devices 106, 108, 110 predict a likely navigation event using local data. In some aspects, a likely navigation event is indicated as a hypertext markup language (HTML) tag embedded within a web page.

[0023] The network 112, and the intervening nodes between the server 104 and the client devices 106, 108, 110, may comprise various configurations and use various protocols including the Internet, World Wide Web, intranets, virtual private networks, local Ethernet networks, private networks using communication protocols proprietary to one or more companies, cellular and wireless networks (e.g., Wi-Fi), instant messaging, hypertext transfer protocol ("HTTP") and simple mail transfer protocol ("SMTP"), and various combinations of the foregoing. It should be appreciated that a typical system may include a large number of connected computers.

[0024] Although certain advantages are obtained when information is transmitted or received as noted above, other aspects of the system and method are not limited to any particular manner of transmission of information. For example, in some aspects, information may be sent via a medium such as an optical disk or portable drive. In other aspects, the information may be transmitted in a non-electronic format and manually entered into the system.

[0025] Although some functions are indicated as taking place on the server 104 and other functions are indicated as taking

place on the client devices 106, 108, 110, various aspects of the system and method may be implemented by a single computer having a single processor. It should be appreciated that aspects of the system and method described with respect to the client may be implemented on the server, and vice-versa.

[0026] Figure 2 is a block diagram depicting an example of a computing device 200, such as one of the client devices 106, 108, 110 described with respect to Figure 1. The computing device 200 may include a processor 204, a memory 202 and other components typically present in general purpose computers. The memory 202 may store instructions and data that are accessible by the processor 204. The processor 204 may execute the instructions and access the data to control the operations of the computing device 200.

[0027] The memory 202 may be any type of tangible memory operative to store information accessible by the processor 204, including a computer-readable medium, or other medium that stores data that may be read with the aid of an electronic device, such as a hard-drive, memory card, read-only memory ("ROM"), random access memory ("RAM"), digital versatile disc ("DVD") or other optical disks, as well as other write-capable and read-only memories. The system and method may include different combinations of the foregoing, whereby different portions of the instructions and data are stored on different types of media.

[0028] The instructions may be any set of instructions to be executed directly (such as machine code) or indirectly (such as scripts) by the processor 204. For example, the instructions may be stored as computer code on a computer-readable medium. In that regard, the terms "instructions" and "programs" may be used interchangeably herein. The instructions may be stored in object code format for direct processing by the processor 204, or in any other computer language including scripts or collections of independent source code modules that are interpreted on demand or compiled in advance. Functions, methods and routines of the

instructions are explained in more detail below (see figures 3-7).

[0029] Data may be retrieved, stored or modified by processor in accordance with the instructions. For instance, although the architecture is not limited by any particular data structure, the data may be stored in computer registers, in a relational database as a table having a plurality of different fields and records, Extensible Markup Language ("XML") documents or flat files. The data may also be formatted in any computer readable format such as, but not limited to, binary values or Unicode. By further way of example only, image data may be stored as bitmaps comprised of grids of pixels that are stored in accordance with formats that are compressed or uncompressed, lossless (e.g., BMP) or lossy (e.g., JPEG), and bitmap or vector-based (e.g., SVG), as well as computer instructions for drawing graphics. The data may comprise any information sufficient to identify the relevant information, such as numbers, descriptive text, proprietary codes, references to data stored in other areas of the same memory or different memories (including other network locations) or information that is used by a function to calculate the relevant data.

[0030] The processor 204 may be any suitable processor, such as various commercially available general purpose processors. Alternatively, the processor 204 may be a dedicated controller such as an application-specific integrated circuit ("ASIC").

[0031] Although Figure 2 functionally illustrates the processor and memory as being within a single block, it should be understood that the processor 204 and memory 202 may actually comprise multiple processors and memories that may or may not be stored within the same physical housing. Accordingly, references to a processor, computer or memory will be understood to include references to a collection of processors, computers or memories that may or may not operate in parallel.

[0032] The computing device 200 may be at one node of a network and be operative to directly and indirectly communicates

with other nodes of the network. For example, the computing device 200 may comprise a web server that is operative to communicate with client devices via the network such that the computing device 200 uses the network to transmit and display information to a user on a display of the client device.

[0033] The system can provide privacy protections for the client data including, for example, anonymization of personally identifiable information, aggregation of data, filtering of sensitive information, encryption, hashing or filtering of sensitive information to remove personal attributes, time limitations on storage of information, and/or limitations on data use or sharing. Data can be anonymized and aggregated such that individual client data is not revealed. A user may also be provided with an opportunity to opt in/out to enable the collection of sharing of data.

[0034] In order to facilitate the prerendering operations of the computing device 200, the memory 202 may further comprise a browser 206, a navigation prediction module 208, a prerender module 210, a client navigation history 212, and an aggregate navigation history 214. Although a number of discrete modules (e.g., 206, 208, 210, 212 and 214) are identified in connection with Fig. 2, the functionality of these modules may overlap and/or exist in a fewer or greater number of modules than what is shown, with such modules residing at one or more processing devices, which may be geographically dispersed.

[0035] The browser 206 provides for the display of a web page 216 to a user of the client device by sending and receiving data across a computer network. The web page 216 may be received in response to a network request, such as a Hypertext Transfer Protocol (HTTP) GET request. The web page 216 may be provided in a markup language, such as Hypertext Markup Language (HTML). The web page 216 may also include various scripts, data, forms, and the like, including interactive and executable content such as ADOBE FLASH content, JAVASCRIPT content, and the like.

[0036] The browser 206 may further comprise a prerendered web page 218. The prerendered web page 218 represents a web page that was requested and accessed by the prerender module 210 in response to a predicted navigation event provided by the navigation prediction module 208. The prerendered web page 218 may be stored in a hidden rendering environment such that the prerendered web page 218 is not visible to a user until the user selects a navigation link with which the prerendered web page 218 is associated. The hidden rendering environment refers to any representation that does not have an immediately apparent interface for the user. For example, the prerendered web page 218 may be stored within a browser "tab" that does not have an associated selection control in a web browser tab interface. The hidden rendering environment may be accessible via a menu or interface command, and in response to the selection of this interface command an interface for the prerendered content is generated, such as by adding a selection control for the prerendered content to the browser tab interface. The hidden rendering environment may also be provided by a variety of other methods suitable to receive and render the page environment while avoiding distraction of the user with potentially irrelevant content. The prerendered web page 208 may be stored in a completely executed and rendered format. For example, any executable code, scripting language, multimedia files, and/or interactive objects may be fully requested and loaded in the prerendering environment. In the event the user selects the navigation event associated with the prerendered web page 218, the browser 206 may swap the prerendered web page 218 with the web page 216, thus providing the content associated with the navigation event without the need to send another network request.

[0037] The memory 202 may further comprise a prerender module 210 to perform fetching of a next web page as identified by the navigation prediction module 208. The prerender module 210 sends a network request for the web page identified to be the

likely next navigation destination that the user will select. The web page received in response to this request is then stored in the browser 206 as the prerendered web page 218. In some aspects, the web page request generated by the prerender module 210 is identical to a standard web page request. In some aspects, the web page request generated by the prerender module 210 comprises certain features to facilitate the prerender process as described further below (see Figures 3-7).

[0038] In some aspects, the computing device 200 may determine a likely navigation event using a client navigation history, such as the client navigation history 212. The client navigation history 212 comprises a set of navigation events associated with past activity of the browser 206. The client navigation history 212 may track a set of visited URLs, also known as a "clickstream," which tracks an order in which the user typically visits URLs (e.g., when the user visits a news website, they tend to next select a URL corresponding to the top story of the day), a set of access times associated with the URLs, and the like. The aggregate navigation history 214 may comprise similar data, but keyed to multiple users rather than a single user. The client navigation history 212 and aggregate navigation history 214 may represent data collected using one or more browser add-ons, scripts, or toolbars. In some aspects, the client navigation history 212 and/or aggregate navigation history 214 are maintained on a remote server, such as the server 104, and provided to the computing device 200. The computing device 200 may maintain separate records to facilitate the predicting of a next likely navigation event, or it may act in concert with remotely stored data. In some aspects, only aggregate navigation history 214 pertaining to the particular web page the user is currently viewing is provided to the computing device 200.

[0039] As described above, the aggregate navigation history data 214 can be maintained in an anonymous fashion, with privacy protections for the individual client data that comprises the

aggregate navigation history, including, for example, anonymization of personally identifiable information, aggregation of data, filtering of sensitive information, encryption, hashing or filtering of sensitive information to remove personal attributes, time limitations on storage of information, and/or limitations on data use or sharing. The aggregate navigation history 214 data can be anonymized and aggregated such that individual client data is not revealed.

[0040] Figure 3 is a flow diagram depicting an example method 300 for prerendering a web page in accordance with aspects of the disclosure. Aspects of the method 300 operate to minimize load times for network content by requesting the network content prior to the user navigating to said network content. The method 300 may be performed by a computing device, such as the computing device 200, to eliminate delays in the user web browsing experience by prerendering web pages that are identified as likely navigation targets by the user. For example, the method 300 may be performed by elements of the browser 206, the navigation prediction module 208, and the prerender module 210 acting together. While aspects of the method 300 are described with respect to the computing device 200, the method 300 may also be performed by the server 104, or any device with hardware and/or software designed to accept instructions. Example aspects of the method 300 as described relate to prerendering of a single page associated with a single navigation event, but the method 300 is also broadly applicable to prerendering multiple pages associated with one or more navigation events, such as prerendering multiple pages in sequence or in a plurality of parallel prerendering instances.

[0041] At stage 302, the computing device 200 determines a next navigation event within a web browser. The next navigation event may be received from a navigation prediction element, such as the navigation prediction module 208, or the next navigation event may be embedded within the web page displaying within the web browser. For example, a search engine may embed a tag

within a set of search results to indicate that a particular link associated with the most relevant result should be prerendered by the web browser. A method for embedding prerendering tags within a web page is described below (see Figure 4).

[0042] At stage 304, the computing device 200 requests the content associated with the navigation event identified at stage 302. The request may comprise a hypertext transfer protocol (HTTP) GET request, a file transfer protocol (FTP) request, or any other network data access request as known in the art. A first set of web content may be provided in response to a HTTP GET request, which may then lead to additional HTTP GET requests in accordance with the first set of web content. For example, a HTTP GET request may result in a base web page being sent to the web browser, which may trigger other page loading operations, such as additional GET requests. The base web page may have multiple embedded images, which are then fetched and rendered using HTTP GET requests for each image. The prerendering process as described may request the full web page, including associated scripts, and then execute those scripts. In response to the request, the web content associated with the navigation event is sent to the web browser. For example, the remote server 104 may send the content to the web browser. To return to the search engine example, upon receiving a web page with an embedded prerender tag, the computing device 200 requests and renders the web page associated with the URL associated with the top search result as identified by the search engine.

[0043] At stage 306, the computing device 200 renders content associated with the navigation event determined at stage 302. The content may be rendered in an alternate instance of the web browser that is hidden from the user. For example, the content may be rendered by an alternate process or in a hidden web browser tab. In some aspects, the computing device 200 may render content associated with multiple navigation events. For example, the navigation events may be associated with a priority

value, with higher or lower values associated with more likely navigation events. As such, the computing device 200 may render the top "n" most likely events, depending upon the capabilities of the computing device 200 (see Figure 7).

[0044] In some aspects, the rendered content may include a redirection operation. An example of a redirection operation is when the act of loading a first page, for instance, www.a.com, causes the browser to load a second page, such as www.b.com. Redirection may occur in response to a HTML tag, a JAVASCRIPT navigation command, or the like. In the case a prerendered page requests a redirection operation, said redirection operation would also occur in the alternate instance.

[0045] In some aspects, a first redirection operation may lead to subsequent redirections. For example, www.a.com may lead to www.b.com, which leads to www.c.com. Aspects of the system and method may account for this fact, such as by storing a history of redirection operations and prerendering the last page in a series in response to a first page being identified to prerender. In some aspects, the method 300 may short-circuit redirection operations by not prerendering intermediate pages when the destination result page (i.e. the page at the end of the redirection operations) is known, and only prerendering the final page in the series.

[0046] At stage 308, the computing device 200 may determine whether the user has selected the link identified at stage 302 within a particular time period. For example, the method 300 may wait for 10 seconds, 20 seconds, 30 seconds, one minute, 5 minutes, or any other time frame. In some aspects, the time frame may be configurable within an options menu associated with the web browser. The time frame may also be specified by the server 104. For example, the server 104 may embed a timeout value within a prerender tag, indicating the length of time the content associated with the tag should be held before being discarded. If the user has selected the link within the time

period, the method 300 proceeds to stage 312. Otherwise, the method 300 proceeds to stage 310.

[0047] At stage 310, the computing device 200 discards the prerendered content. The prerendered content is allowed to expire in this manner in order to free up system resources associated with the prerender operation and prevent the user from being presented with stale data. In some aspects, a different expiration period may be specified for different types of data. For example, data that is likely to rapidly change, such as breaking news articles, stock quotes, syndication feeds, and the like, is more likely to grow stale quickly, and might be associated with a shorter expiration timer, such as 10 seconds, 20 seconds, or one minute. Data that is less likely to change, such as reference articles, archives, box scores, and the like, might be associated with a longer expiration timer, such as an hour, 6 hours, 24 hours, a week, or the like. In some aspects, the wait time is based on a page type. For example, if the user is viewing a lengthy news article, a longer wait time may be specified when prerendering a next page of the article to provide the user time to read the current page of the article before selecting the next page. In some aspects, the wait time is determined by a time-to-live (TTL) parameter embedded within the page. In some aspects, the wait time is hard coded within the browser. The method 300 then ends after the prerendered content is discarded, though in some aspects the method 300 may repeat indefinitely as the user browses web content.

[0048] At stage 312, the computing device 200 swaps the alternate instance containing the prerendered content into the current user instance, in response to the user selecting the navigation event determined at stage 302. In this manner, the computing device 200 displays the content associated with the navigation event instantly, without a delay while the data is requested from and received from the host server. To return to the search results example, if the user selects the most relevant result, the associated page is already loaded in the

alternate instance and thus is available instantly as soon as the user clicks the link. In some aspects, the user may select the navigation event prior to the prerendering operation completing. In such cases, the partially loaded page would be swapped into the current user instance. The partially loaded page would then continue to load in the current user instance as normal. The method 300 then ends after making the prerendered content available to the user.

[0049] Figure 4 is a flow diagram depicting an example method 400 for indicating a navigation event to prerender. In some aspects, the server hosting the content, such as the server 104, can specify the navigation event for the web browser to prerender. Allowing the server to specify the prerendered link or links may be optimal for a variety of reasons, such as because the server has the best information on which navigation event is most likely, or because the server wishes to manage bandwidth associated with particular hosted links, such as for load balancing operations. As such, the server 104 may embed a prerender tag or tags within a hosted page, instructing the prerender module 210 of the computing device 200 in which content to render.

[0050] At stage 402, the server 104 identifies one or more links within a hosted web page. For example, the server 104 may generate a list of web sites in response to a query from the computing device 200. Each search result may be associated with a particular link on the page that is to be provided in response to the query.

[0051] At stage 404, the server 104 selects one or more links to prerender, chosen from the links identified at stage 402. The links to prerender may be chosen by a variety of methods for a variety of purposes, such as by examining a user navigation history, an aggregate navigation history, a most relevant link, and the like. In some aspects, the links are chosen by a search engine analyzing the links most frequently clicked in response to a particular query. For example, the search engine may

collect "click-through" data indicating which links are clicked on in response to a particular query, and select one or more links to prerender based upon the frequency with which the particular links are selected. To continue the search engine example, the server 104 might identify the top search result, or the top "n" search results as optimal links to prerender. In some aspects, the server 104 may further identify a ranking of links to prerender, with the most optimal link to prerender being ranked number 1, the next most optimal link being ranked number 2, and so on. The server 104 may also identify a likelihood that each link will be selected, rather than a rank. For example, the server 104 may indicate that a particular link has a 50% likelihood, a 25% likelihood, or a 80% likelihood of being selected by a user. Prerender operations performed by the client may elect to prerender links associated with a threshold level of likelihood, for example, such as at least 50% likelihood, at least 80% likelihood, or any other threshold value.

[0052] At stage 406, the server 104 embeds one or more prerender tags with the one or more links identified at stage 404. For example, the server 104 may include a hypertext markup language (HTML) tag for identifying the optimal link or links to prerender. As described above, the prerender tag may also include a ranking value and a TTL value specifying a time to keep the prerender. To continue the search engine example, the server 104 may identify each search result with a prerender tag, along with ranking the prerender order of the results by the relevance of the search result.

[0053] At stage 408, the page with the embedded prerender tags is sent to the client that requested the page. The method 400 then ends after sending the page.

[0054] Figure 5 is a flow diagram depicting an example method 500 for managing web impression statistics during a prerender operation in accordance with aspects of the disclosure. Web sites often track what content is viewed, which pages are

accessed, in what order those pages are accessed, and how much time is spent on each page. These capture and analysis of such data is broadly referred to as "website analytics". These statistics provide data that may be used by site operators for a variety of purposes. For example, one such use of this data is to calculate rates for advertisements displayed along with the content. However, in some instances, the user may not select the navigation event associated with prerendered content, thus possible interfering with collection of this data. For example, the traditional model of tracking user impressions via the number of times a given page is sent to the user may no longer provide accurate results, because the user may not view the content every time it is sent. As such, the server, such as the server 104, may provide a method for accurate tracking of impressions that takes into account prerendered content. One such method is described below. Although specific examples of the method may be provided with respect to managing impression statistics of advertisement content, it should be appreciated that aspects of the method are broadly applicable to ensuring accurate website analytics for pages served in response to prerender requests.

[0055] At stage 502, the server 104 receives a prerender request from a computing device, such as the computing device 200. In some aspects, the prerender request may be identified using a scripting application programming interface (API), such as by the computing device 200 indicating the content request is for a prerender operation via a JAVASCRIPT API transmitted with the content request. The prerender request may also be accompanied by a modified HTTP header, indicating that the request is a prerender request. In some aspects, the server 104 is configured to receive an event notification from the computing device 200 as provided by the API, such as by registering for the event notification using an application executing on the server 104. In some aspects, the server 104 may elect to not send content in response to a prerender

request, such as in instances where the server is bandwidth constrained, or otherwise does not wish to allow prerendering. Providing notification of the prerender request in this manner also allows for the server 104 to "opt-out" of processing the prerender request. For example, the server 104 may be bandwidth constrained and thus need to focus on hosting non-prerendered content. In some aspects, the server 104 may also opt-out by including a HTML metatag indicating as such, or an entry in a "robots.txt" file hosted on the server to prevent prerendering of content by a search engine, such as indicated in a list of search results.

[0056] At stage 504, the server 104 may send a partial web page, with the certain content redacted. For example, the server 104 may not provide images associated with advertisement content in response to a prerender request, as in many instances advertisement impressions are determined based upon the number of times the advertisement image is sent. A similar method may apply to advertisements provided in other manners, such as via scripting or executable content. As an example, the server 104 might respond to a prerender request by providing a base web page. The base web page may result in the computing device 200 initiating one or more additional HTTP GET requests for content identified by the base web page. The server 104 may opt not to send content in response to HTTP GET requests for ad images identified within the base web page.

[0057] In some aspects, the redacted version of the content may include a special HTTP error codes associated with redacted images, scripts, and the like, indicating to the receiving computing device that certain content has been redacted. In this manner, the computing device 200 may be informed that certain content has been redacted and will be supplied later, rather than that the content is unavailable. As such, the computing device 200 may opt to indicate redacted content with a particular graphic or temporary display while the content loads upon access, or the redacted content may appear blank.

[0058] In some cases, the computing device 200 may not be aware of how to properly display redacted content, in which case the computing device 200 may abort the prerender operation. Furthermore, when the user accesses the content, the computing device 200 may notify the server 104 that the prerendered page has been viewed. In response, the server 104 may send the redacted content to the computing device 200.

[0059] At stage 506, the server 104 receives confirmation that the prerender of the content was viewed by the user. As above, the server 104 may be notified of the viewing of the content via an API provided through the web browser. For example, a script may communicate an event notification to the server 104. If the server 104 is registered for such an event notification, it may take appropriate action, such as sending the redacted content to the computing device 200. By informing the server 104 that the prerender content was viewed, the server 104 knows to send any redacted content and update the impression statistics associated with any advertisements.

[0060] At stage 508, the server 104 sends the content that was redacted at stage 504. The content can now be sent because the server 104 has ensured that the user has accessed the content. For example, the server 104 may send images associated with advertisements, scripting code, or the like.

[0061] At stage 510, the server 104 updates the impression statistics associated with the content as sent at stage 508. In some aspects, stages 508 and 510 occur simultaneously as the act of hosting the image/sending an advertisement automatically updates the page impression statistics. In this manner the server 104 maintains accurate reporting of website analytics data while also allowing the computing device 200 to perform prerender operations.

[0062] Figure 6 is a flow diagram depicting an example method 600 for managing a prerender operation in accordance with aspects of the disclosure. In order to provide an optimal prerendering experience, the computing device 200 can account

for different factors that may affect the browsing experience. For example, web pages commonly use small files known as cookies to provide various page services via the web browser. For example, a news site may employ a tracking cookie to allow a user to access one news article per day without registering on the page. A prerender operation associated with that news site could then interfere with the user's browsing experience, as the user may prerender an article he does not intend to read, thus using up the one free article per day on a page he has not viewed.

[0063] Some web pages have embedded audio or video content that begins to play as soon as the page loads. In some implementations, the prerender method and system as described herein can allow such content to begin playing before the user navigates to the content, thus when the user begins viewing the page, the content is already playing. Some web pages may include additional content that might not be suitable for prerendering, such as pop-up pages, scripted programs, downloads, authentication prompts, translation bars, and the like. As such, the prerender module 210 may employ a method, such as the method 600, for managing these special cases.

[0064] At stage 602, the computing device 200 requests and receives prerender content from a server, such as the server 104. As described above, the request may be presented in or accompanied by an API call, indicating to the server that the request is a prerender request. The page is prerendered based on a response received from the server. During the prerendering process, the computing device 200 monitors for several particular cases as described below. Although presented in a particular order, these cases could be encountered in any order, such as by identifying an embedded video prior to experiencing a cookie state change, for example.

[0065] At stage 604, the computing device 200 determines whether the requested prerender content includes a cookie state change or another special case. For example, if a user is

accessing a banking web page, if an account page is prerendered and then the user logs out of their account, selecting the prerender link would incorrectly report data as if the user was still logged in. As such, the prerender would be invalid as soon as the user has logged out. The computing device 200 therefore monitors the state of any cookies associated with prerender content and may take appropriate action. Appropriate action may include attempting to properly propagate cookie state changes, ignoring cookie state changes and continuing with the prerender operation, or giving up and aborting the prerender operation. The computing device 200 further determines if the prerender content includes a special case, such as a pop-up page, a scripted application, a download, an authentication menu, or a translation bar. If a cookie state change has occurred or the prerender page includes a special case, the method proceeds to stage 608. Otherwise the method 600 continues to stage 606.

[0066] At stage 608, the computing device 200 may abort the prerender operation. The computing device 200 may abort the prerender operation because the cookie state change or special case identified at stage 604 makes the data associated with the prerendered page stale or otherwise unviewable to the user. In some aspects, the computing device 200 may ignore the cookie state change. The method 600 then ends.

[0067] As an alternate method of managing cookie data, the computing device 200 may track any cookie state change associated with the prerendered content, and apply the state change when the prerendered content is accessed by the user. In some aspects, the computing device 200 may also monitor for conflicts in cookie state change, and abort the prerendering operation in the event of a conflict. For example, a cookie state change may occur after the prerender operation is complete, thus causing a conflict. In such a case, the prerendered version of the page could be discarded to avoid presenting incorrect page data to the user.

[0068] At stage 606, if the prerender content does not contain a special case and has not had a cookie state change, the computing device 200 determines whether there is multimedia content on the page, such as audio or video. For example, the page may include an embedded video programmed with ADOBE FLASH, HTML5 media, or another type of embedded media. In such a situation, it may be desirable to pause the multimedia until the user actually navigates to the page, thus saving bandwidth and ensuring the user may view the entire multimedia file. If the page contains embedded multimedia, the method 600 proceeds to stage 612. Otherwise the method 600 proceeds to stage 610.

[0069] At stage 610, the computing device 200 has determined that the content contains no multimedia content, cookie state changes, or special cases, and thus the computing device 200 prerenders the page. The prerendering process may be performed by receiving the page from the server 104, and rendering the page in an application instance that is hidden from the user. Rendering the page generally relates to loading all of the data received from the server 104, and executing any instructions contained therein, as they would otherwise be executed when viewing the page in a non-hidden application instance.

[0070] At stage 612, the computing device 200 has determined that the content contains multimedia content. Thus the computing device 200 prerenders the content of the page other than the multimedia content, and may pause playback of the multimedia content, or defer the instantiation of a plug-in associated with the content. For example, the computing device 200 might pause an embedded audio or video file after prerendering the rest of the page, or a FLASH video might not execute until the prerender content is shown to the user.

[0071] At stage 614, the computing device 200 determines whether the user has accessed the prerendered content. If the user accesses the prerendered content, the content is displayed as described above, and the method 600 proceeds to stage 616. If the user does not access the prerendered content, for

example, if a prerender timer expires as described above with respect to Figure 3, then the method 600 ends.

[0072] At stage 616, the computing device 200 sends notification that the user accessed the prerender to the server 104. As described above, the notification may be sent using a scripted API from which the server is configured to receive event notifications. Providing notification in this manner allows for the management of website impressions by the server (see Figure 5).

[0073] Figure 7 is a flow diagram depicting an example method 700 for configuring a prerender operation in accordance with aspects of the disclosure. Different computing devices may possess different capabilities, such as network bandwidth, system memory, and processor speed. As such, various devices may be able to support different prerender operations. A device with more capability may be able to support multiple prerender operations, or the prerendering of more complicated pages. As such, the method 700 provides for configuration of these prerender operations based upon the capabilities of the system.

[0074] At stage 702, the computing device 200 conducts a measurement of system capabilities. For example, the computing device 200 may query the operating system for the amount of free system memory, the computing device 200 may request the processor speed from the basic input-output system (BIOS), or the computing device 200 may conduct a network speed test to determine free bandwidth.

[0075] At stage 704, the computing device 200 determines a maximum number "n" prerenders based upon the capabilities as determined at stage 702. In some aspects, the computing device 200 may determine a set of system requirements to support a single prerender instance. For example, a single prerender instance may require 10 percent processor utilization, 16 megabytes of system memory, at least 10 kilobytes/second of network bandwidth, or the like. The system requirements may also include multiple capabilities associated with different

system components, such as each of a minimum memory requirement, a processor utilization, and a network bandwidth. These system requirements may be determined by measuring the system utilization of a single prerender instance executing on the client device 200.

[0076] For example, the computing device 200 may determine that each prerender operation is likely to require the system processor while prerendering. Such processor use may be "bursty" in that the entire processor is needed when the page is first loading, with reduced processor need thereafter. In such cases, the computing device 200 may assign a lower processor priority to the prerender instance to ensure that prerendering of a page does not impact the user's browsing experience. In some aspects, the computing device 200 may monitor a system memory to determine a number of prerender instances. The computing device 200 may determine an amount of memory used by a prerender instance and then compare this amount to the total available memory capacity. In some aspects, the computing device 200 may reserve a particular amount of system memory for prerendering, and generate prerender instances until the reserved memory is filled. In some aspects, a maximum network bandwidth of the system is determined, and prerender instances are generated based on the percentage of available bandwidth. In some aspects, the computing device 200 may request a bandwidth measurement from a remote server to identify a local maximum bandwidth, and use the measurement to determine a prerender setting.

[0077] At stage 706, the computing device 200 is configured to perform as many simultaneous prerenders as the system capabilities can support, as determined at stage 704. In some instances, the number of prerenders may be configurable by the user, such as in an options menu.

[0078] The stages of the illustrated methods are not intended to be limiting. The functionality of the methods can exist in a fewer or greater number of stages than what is shown and, even

with the depicted methods, the particular order of events may be different from what is shown in the figures.

[0079] The systems and methods described herein advantageously provide for an improved browsing experience. By prerendering the next navigation event, the system and methods minimize the amount of time users wait for network content to load. The various management and configuration methods provide for a robust experience that ensures accurate impression counts for advertisers hosting content.

[0080] As these and other variations and combinations of the features discussed above can be utilized without departing from the disclosure as defined by the claims, the foregoing description of the embodiments should be taken by way of illustration rather than by way of limitation of the disclosure as defined by the claims. It will also be understood that the provision of examples of the disclosure (as well as clauses phrased as "such as," "e.g.," "including" and the like) should not be interpreted as limiting the disclosure to the specific examples; rather, the examples are intended to illustrate only some of many possible embodiments.

CLAIMS

1. A computer-implemented method for managing website analytics for prerendered content, the method comprising:

receiving notification that a network content request is a prerender request for a set of network content;

redacting, using a processor, a subset of content from the network content to create a set of partial content such that the set of partial content comprises the network content with the redacted content removed;

providing the set of partial content in response to the network content request;

receiving notification that the redacted content was accessed by a user device; and

sending the redacted content in response to the notification that the redacted content was accessed.

2. The method of claim 1, further comprising updating an impression statistic in response to the notification that the redacted content was accessed.

3. The method of claim 1, wherein the redacted content comprises one or more advertisements.

4. The method of claim 1, further comprising inserting one or more hypertext transfer protocol (HTTP) error codes within the set of partial content in place of the redacted content.

5. The method of claim 1, wherein the set of redacted content comprises at least one image, script, cookie, or executable instruction.

6. The method of claim 1, wherein the redacting stage is performed by not providing the redacted content in response to a request from a client device for the redacted content.

7. A computer-implemented method for managing website analytics for prerendered content, the method comprising:

receiving notification that a network content request is a prerender request for a set of network content;

providing the network content in response to the network content request;

receiving notification that the network content was viewed after prerendering; and

updating an impression statistic in response to the notification that the network content was viewed such that the impression statistic reflects a count of users who have actually viewed the content.

8. The method of claim 7, wherein the network content comprises at least one advertisement, and the impression statistic is associated with the at least one advertisement.

9. The method of claim 7, wherein the impression statistic is associated with at least one image, script, cookie, or executable instruction.

10. A processing system for managing website analytics for prerendered content comprising:

at least one processor;

a prerender module associated with the at least one processor; and

a memory coupled to the at least one processor;

wherein the prerender module is configured to receive notification that a network content request is a prerender request for a set of network content, redact a subset of content from the network content to create a set of partial content, such that the set of partial content comprises the network content with the redacted content removed, provide the set of partial content in response to the network content request, receive a notification that the redacted content was accessed, and send the redacted content in response to the notification that the redacted content was accessed.

11. The processing system of claim 10, wherein the redacted content comprises at least one advertisement.

12. The processing system of claim 10, wherein the prerender module is further configured to update an impression statistic in response to the notification.

13. The processing system of claim 10, wherein the impression statistic is associated with the redacted content.

14. The processing system of claim 10, wherein the prerender module is further configured to provide a hypertext transfer protocol error code in place of the redacted content.

15. A processing system for serving network content and managing prerender requests comprising:

- at least one processor;
- a prerender module associated with the at least one processor; and
- a memory coupled to the at least one processor;

wherein the prerender module is configured to determine whether an incoming network request is a prerender request.

16. The processing system of claim 15, wherein the prerender module is further configured to, in response to the incoming network request being identified as a prerender request, not provide access to the network content.

17. The processing system of claim 16, wherein the prerender module is further configured to provide a hypertext markup language metatag to indicate that the processing system does not provide prerender content.

18. The processing system of claim 15, wherein the prerender module is further configured to measure a system utilization, and, in response to the system utilization being below a threshold value, provide access to the network content requested by the incoming network request.

19. The processing system of claim 15, wherein the system utilization comprises at least one of a processor utilization, a network bandwidth, or an available system memory.

20. The processing system of claim 15, wherein the prerender module is further configured to receive a notification that a user wishes to access content previously requested via a prerender request, and to provide web content in response to the notification.




FIGURE 1




FIGURE 2


300

FIGURE 3


400

FIGURE 4


500

FIGURE 5


600

FIGURE 6


700

FIGURE 7