

(51) International Patent Classification:
G06Q 50/10 (2012.01)(21) International Application Number:
PCT/US2011/066002(22) International Filing Date:
20 December 2011 (20.12.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/424,940 20 December 2010 (20.12.2010) US
13/330,935 20 December 2011 (20.12.2011) US(71) Applicant (for all designated States except US):
AKAMAI TECHNOLOGIES, INC. [US/US]; 8 Cambridge Center, Cambridge, MA 02142 (US).(72) Inventors: **FLACK, Martin, T.**; c/o Akamai Technologies, Inc., 8 Cambridge Center, Cambridge, MA 02142 (US). **KOBRIN, Eric, L.**; c/o Akamai Technologies, Inc., 8 Cambridge Center, Cambridge, MA 02142 (US). **LUDIN, Stephen, L.**; c/o Akamai Technologies, Inc., 8 Cambridge Center, Cambridge, MA 02142 (US). **KAUFMAN, David**; c/o Akamai Technologies, Inc., 8 Cambridge Center, Cambridge, MA 02142 (US).(74) Agents: **MATT, Joshua, T.** et al.; Akamai Technologies, Inc., 8 Cambridge Center, Cambridge, MA 02142 (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,

[Continued on next page]

(54) Title: METHODS AND SYSTEMS FOR DELIVERING CONTENT TO DIFFERENTIATED CLIENT DEVICES




FIG. 4 – DEC Caching Process

(57) Abstract: Methods and systems are disclosed for delivery of tailored content to differentiated devices, such as desktop, mobile, and tablet devices, over a computer network. In one embodiment, a proxy cache server has a content cache for storing previously retrieved objects like web pages or multimedia files. For at least some objects, several versions are stored, each version representing an object suited for a given set of client device characteristics. A device-equivalency data structure maintained at the proxy facilitates a determination of whether such cached versions can be used to service a current request. The versions might represent, for example, modified versions created using, e.g., mobile device transcoding techniques, in response to prior requests. They may also represent a set of alternate content created by a content provider and available from an origin server. Such methods and systems may be implemented in a distributed computing networks, e.g., a content delivery network.



UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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 the identity of the inventor (Rule 4.17(i))

— 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))

METHODS AND SYSTEMS FOR DELIVERING CONTENT TO DIFFERENTIATED CLIENT DEVICES

REFERENCE TO RELATED APPLICATIONS

- 5 This application is a continuation of and claims the benefit of priority of U.S. Application No. 13/330,935, filed December 20, 2011, and also claims the benefit of priority of U.S. Provisional Application No. 61/424,940, filed December 20, 2010, the teachings of which are hereby incorporated by reference in their entirety.

BACKGROUND

10 Technical Field

The present application generally relates to the delivery of content to client devices, including in particular the delivery of content to a variety of devices with different characteristics.

Brief Description of the Related Art

- As is known in the art, a distributed computer system may be used to deliver content to client
15 devices in both wireless and other fixed-line environments. One such distributed computer system is a “content delivery network” or “CDN” that is operated and managed by a service provider. The service provider typically provides the service on behalf of third parties. A “distributed system” of this type typically refers to a collection of autonomous computers linked by a network or networks, together with the software, systems, protocols and
20 techniques designed to facilitate various services, such as content delivery or the support of outsourced site infrastructure. Typically, “content delivery” refers to the storage, caching, or transmission of content, or streaming media and applications on behalf of content providers, and ancillary technologies used therewith including, without limitation, DNS request handling, provisioning, data monitoring and reporting, content targeting, personalization, and
25 business intelligence.

Content delivery to mobile devices such as a smartphone or tablet, presents unique challenges. Often such devices have limited capabilities in terms of their display or processing power, and they typically receive content via wireless connections that are of limited bandwidth in comparison to fixed-line environments.

One approach for delivering content to such devices so is through the use of transcoders, which take in web content (e.g., a web page) and transform or re-author that content into a format suited to the display and data transmission capabilities of the device. A variety of transcoding algorithms are known in the art. Exemplary approaches for transcoding content are described in U.S. Patent No. 7,047,033 and US Patent Publication 2003/0115365, the teachings of both of which are hereby incorporated by reference. Transcoding of image, video, and other content is described in Bharadvaj, Joshi et al., "An Active Transcoding Proxy To Support Mobile Web Access," Reliable Distributed Systems, Proceedings of 17th IEEE Symposium on Reliable Distributed Systems, pp. 118-123 (1998), the teachings of which are hereby incorporated by reference.

Alternatively, a content provider might author several versions of content in the first instance, each version being adapted for a particular device. The number such versions can rapidly escalate, however, and the availability of multiple versions does not address how to efficiently deliver an appropriate version in response to a particular request.

While the foregoing approaches relate to adaptation techniques, they do not fully address the problems of delivering tailored content at scale. Hence, there is a need to provide methods and systems for providing device-adapted content in distributed computing systems and for advantageously managing adapted content in such systems. The techniques disclosed herein address these and other needs that will become apparent in view of the following disclosure.

SUMMARY

Described herein are methods and systems for improving the delivery of content that is suited for different client devices, such as content adapted for particular mobile devices. For example, described herein is a proxy cache server that has a content cache for storing previously retrieved objects, such as web pages or multimedia files. Several versions of at least some objects (including a default version, potentially) can be stored in the content cache, each of the versions representing an object suited for a particular set of one or more client devices. These versions may have been adapted on-the-fly by a content adaptation engine (e.g., in response to previous client-device requests) or may represent several alternate versions of an object authored and published by a content provider. A device-equivalency-class data structure (referred to as DEC cache in some examples) is maintained at the proxy.

This data structure facilitates a determination of which, if any, of the cached versions can be used to service a given request from a particular client device. Typically, the determination turns on whether the device characteristics that were used to adapt the cached content match the characteristics of the particular client device that is subsequently requesting the content.

- 5 In many cases, the DEC cache may reflect only some (not necessarily all) of the characteristics that were used for adaptation, improving the chances of a cache hit. In effect this creates device equivalency classes, whose breadth is configurable. Furthermore, the characteristics may be expressed in ranges, such a match occurs when a given client device has characteristics that fall within the specified ranges for a DEC.
- 10 If a particular version of content can be used, then it can be delivered to the client directly from the content cache. In many instances, a default version will be available (e.g., from origin or from the cache) if all other versions fail to match. Alternatively, if there is no matching version, the content can be adapted at request time by a content adaptation engine. The result of this real-time adaptation can be stored in the DEC cache and associated with a
- 15 newly defined class definition, enabling it to be re-used for future client requests that fall within the class.

- By way of further illustration, in one non-limiting aspect of the invention, a content server (or more than one content server, acting cooperatively) operates as follows to facilitate the delivery of content to a client. The content server receives a request for web content from a
- 20 client device, e.g., a mobile device such as a smartphone, or a tablet, or another connected device. Typically, the content server is acting as a proxy cache server. The content server determines that the client device has certain characteristics from information in the request, such as the user_agent or other header information. It determines that a set of one or more of the client device's characteristics match one of a plurality of sets of stored characteristics for
- 25 client devices. This set of characteristics may be all of the device's characteristics that were determined from the request, but it is typically less than that.

- The stored sets of device characteristics are preferably specified in a control file, in a header received from an origin server, in information received from a content adaption engine, or from another source. Each of the sets of stored characteristics is associated with a version of
- 30 the web content.

The content server obtains the version of the web content associated with the matching set of

stored characteristics, e.g., retrieving that version from cache. If not cached, or depending on the implementation, the version might be obtained from an origin server, or from a content adaptation engine. Whatever the source, the content server serves that version to the client device in response to the request for web content.

- 5 If it was not cached, then the content server caches the matching version of the web content in the local cache. Preferably, it uses a cache key associated with the matching set of stored characteristics. To manage access to the local cache, the content server preferably maintains a data structure, sometimes referred to as a DEC cache herein, that maps a universal resource locator (URL) or other web content identifier to the plurality of versions of the web content.
- 10 The characteristics referred to above can take many forms, but preferably are represented as key/value pairs, where the key represents a property of the device, e.g., screen_size=320px. The value may be a range, e.g., 100-300px. The values and ranges are not limited to numerical values, but may also include Booleans, strings, etc., such as for an operating system: OS={OS_name1, OS_name2}. Thus, the content server determines whether a given
- 15 characteristic of the first client device matches a given stored characteristic by determining whether the value of the given property of the first client device matches the value, or falls within the range of values, of the given stored property.

In another non-limiting aspect of the invention, a content server delivers content to a client over a computer network using the following approach. The content server receives a request

20 for web content from a first client device, typically a mobile device such as a smartphone or tablet, and so on. The content server obtains the web content from local cache (if available) or from a remote server. The remote server is often an origin server associated with a content provider customer who has offloaded delivery of content to the content server, which is operated by a content delivery network service provider.

- 25 The content server modifies the web content to create web content adapted for the first client device. This modification may be performed in a variety ways known in the art, but it takes into account one or more characteristics of the first client device. The characteristics of the first device can be determined by mapping a user_agent string, or other header information in the request from the first client device, to set of characteristics. The content server caches the
- 30 modified web content.

Next, the content server receives a request for the web content from a second client device. The content server determines whether the modified web content is valid to be served to the second client device, that determination including a determination of whether the second client device has at least one of the one or more characteristics. The content server sends the second client device the modified web content from the cache, if it is valid to be served to the second client device. If not, the content server can obtain a version modified for the second client device by adapting the content, as it did for the first client device.

As noted above, the determination of whether the modified web content is valid to be served to the second client device typically includes a determination of whether the second client device has some or all of the one or more characteristics that were used to modify the content for the first client device. This may be achieved by specifying a configurable range of values, as was described earlier. Note that another approach involves defining a tolerance range for the values, e.g., match if the second device's screen width is within +/-10% of the size of the first device's screen width.

The content server typically consults a data structure that maps web content identifiers to a plurality of versions of web content. Each version is identified by a cache key denoting a device equivalency class (DEC). The DEC identifies one or more characteristics of the client device that were taken into account in adapting the web content to create the given version.

Another, non-limiting aspect of the invention relates to a content server that receives a request for web content from a first client device and obtains the requested web content from a remote server or from a local cache. The content server obtains a version of the web content modified for the first client device. The content server assigns a class identifier (e.g., device equivalency class, or DEC) to that version. The class identifier can be thought of as an equivalency bucket: other devices that qualify for the class can be served with the same version of the content. Typically, the content server caches the subject version using the class identifier as part or all of a key.

Assume that the content server then receives a request for the web content from a second client device. The content server determines whether the second client device falls within the class identified by the class identifier. If so, the content server sends the second client device that version of the web content from the cache. If not, then the content server obtains another version of the content by having the content modified for the second client device.

The number of characteristics constituting a particular class is typically configurable. It may represent fewer than the number of characteristics actually taken into account to create a given version of the web content.

As those skilled in the art will recognize, the foregoing description merely refers to examples of the invention. It is not limiting and the teachings hereof may be realized in a variety of computer systems, methods, apparatus, and non-transitory computer-readable media.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic diagram illustrating one embodiment of a known distributed computer system configured as a content delivery network;

FIG. 2 is a schematic diagram illustrating one embodiment of a machine on which a content delivery server in the system of FIG. 1 may be implemented;

FIG. 3 is a schematic diagram illustrating one embodiment of a communication flow amongst a client, a content server, and an origin server;

FIG. 4 is a flow diagram illustrating one embodiment of a process for caching according to the teachings hereof;

FIG. 5 is a schematic diagram illustrating one embodiment of a device equivalency class cache and a proxy's content cache;

FIG. 6 is a flow diagram illustrating one embodiment of a process for determining whether a given requesting client device is a member of a device equivalency class; and

FIG. 7 is a block diagram that illustrates hardware in a computer system which may be used to implement the teachings hereof.

DETAILED DESCRIPTION

The following description sets forth embodiments to provide an overall understanding of the principles of the structure, function, manufacture, and use of the methods and apparatus disclosed herein. The methods and apparatus described herein and illustrated in the accompanying drawings are non-limiting examples; the scope of the present invention is

defined solely by the claims. The features described or illustrated in connection with one exemplary embodiment may be combined with the features of other embodiments. Such modifications and variations are intended to be included within the scope of the present invention. All publications and references cited herein are expressly incorporated herein by
5 reference in their entirety.

1.0 Content Delivery Networks

The methods and systems disclosed herein may be implemented in the context of a distributed computer system, e.g., a content delivery network ("CDN") as illustrated in Figures 1-2; although they are not limited to such implementations.

- 10 With respect to FIG. 1, a distributed computer system 100 is configured as a CDN and is assumed to have a set of machines 102a-n distributed around the Internet. Typically, most of the machines are servers located near the edge of the Internet, i.e., at or adjacent end user access networks. A network operations command center (NOCC) 104 manages operations of the various machines in the system. Third party sites affiliated with content providers, such
15 as web site 106, offload delivery of content (e.g., HTML, embedded page objects, streaming media, software downloads, and the like) to the distributed computer system 100 and, in particular, to the CDN's content servers (also called "edge servers" in light of the possibility that some may be located near an Internet "edge"). Such content servers may be grouped together into a point of presence (POP) 107.
- 20 Typically, content provider customers offload their content delivery by aliasing (e.g., by a DNS CNAME) given content provider domains or sub-domains to domains that are managed by the service provider's authoritative domain name service, more details of which are set forth in U.S. Patent No. 7,293,093 and 7,693,959, the disclosures of which are incorporated by reference. End users operating client machines 122 that desire the content are directed to
25 the distributed computer system, and more particularly to one of the content servers on the machines 102, to obtain that content more reliably and efficiently. The content servers service the client requests, e.g., by retrieving requested content from a content server cache, retrieving the requested content from another content server, or by retrieving the requested content from the origin server 106 (and, potentially, storing that content in the cache to
30 service future requests).

Although not shown in detail in FIG. 1, the distributed computer system may also include other infrastructure, such as a distributed data collection system 108 that collects usage and other data from the edge servers, aggregates that data across a region or set of regions, and passes that data to other back-end systems 110, 112, 114 and 116 to facilitate monitoring, logging, alerts, billing, management and other operational and administrative functions. Distributed network agents 118 monitor the network as well as the server loads and provide network, traffic and load data to a DNS query handling mechanism 115, which is authoritative for content domains being managed by the CDN. A distributed data transport mechanism 120 may be used to distribute control information (e.g., metadata to manage content, to facilitate load balancing, and the like) to the CDN servers.

As illustrated in FIG. 2, a given machine 200 comprises commodity hardware (e.g., an Intel Pentium or other processor) 202 running an operating system kernel (such as Linux or variant) 204 that supports one or more applications 206a-n. To facilitate content delivery services, for example, given machines typically run a set of applications, such as an HTTP proxy 207 (sometimes referred to as a “global host” or “ghost” process), a name server 208, a local monitoring process 210, a distributed data collection process 212, and the like. For streaming media, the machine typically includes one or more media servers, such as a Windows Media Server (WMS) or Flash server, as required by the supported media formats.

Client machines 122 include conventional personal computers, laptops, other digital data processing devices. Client machines 122 may also include mobile devices, which may include any of a variety of smartphones or personal digital assistants (PDAs), running a client application (e.g., a web browser) that issues requests for content from the servers, receives responses from the servers, and processes and displays the received content for a user.

A given content server in the CDN is typically configured to provide one or more extended content delivery features, preferably on a domain-specific, customer-specific basis, preferably using configuration files that are distributed to the content servers using a configuration system. A given configuration file preferably is XML-based and includes a set of content handling rules and directives that facilitate one or more advanced content handling features. The configuration file may be delivered to a content server via the data transport mechanism. U.S. Patent No. 7,111,057, the disclosure of which is incorporated herein by reference, illustrates a useful infrastructure for delivering and managing content control information, and this and other server control information can be provisioned by the CDN service provider

itself, or (via an extranet or the like) the content provider customer who operates the origin server.

The CDN may include a network storage subsystem, such as described in U.S. Patent No. 7,472,178, the disclosure of which is incorporated herein by reference.

- 5 The CDN may operate a server cache hierarchy (Cache-H) to provide intermediate caching of customer content; one such cache hierarchy subsystem is described in U.S. Patent No. 7,376,716, the disclosure of which is incorporated herein by reference.

- 10 The CDN may provide secure content delivery among a client browser, content server and customer origin server in the manner described in U.S. Publication No. 2004/0093419, the disclosure of which is incorporated herein by reference. Secure content delivery as described therein enforces SSL-based links between the client and the content server process, on the one hand, and between the content server process and an origin server process, on the other hand. This enables an SSL-protected web page and/or components thereof to be delivered via the content server.

15 2.0 Delivery Of Device-Adapted Content

- At the outset, it should be understood that while for purposes of illustration the following embodiments sometimes refer to mobile devices, the methods and systems described herein are not limited to the adaptation/caching of content for mobile devices. The methods and apparatus disclosed herein are broadly applicable to the delivery of content to client devices that are differentiated in terms of their capabilities, limitations, or other hardware and/or software characteristics. Examples of differentiating characteristics include display, processing, installed software (operating system, browser), support for certain technologies (e.g., Flash), form factor, intended deployment, and so on.

- 25 A mobile device such as a smartphone or tablet represent particular examples of client devices, and can be contrasted with the traditional desktop PC. They are therefore convenient to use for illustrative purposes. In accordance with the teachings hereof, content can also be adapted and delivered to other client devices, e.g., such as gaming systems, televisions, e-readers, connected devices, and so on, in accordance with the teachings hereof. Furthermore, content may be adapted for different browsers running on the same device.

FIG. 3 is a block diagram illustrating a flow of messages between a mobile client device 300, a content server 301, and an origin server 304. For ease of illustration, in FIG. 3, these components are shown separate from the larger context of the content delivery network 100, but they correspond in functionality to the apparatus described above in connection with

FIGS. 1-2.

In this embodiment, content server 301 contains a content adaptation engine (CAE) 302a component. This component may be a routine or set of routines incorporated into an HTTP proxy process running on the content server 301, or a separate application executing in the content server machine. Alternatively, the content adaptation engine can be implemented in the origin server, as shown by CAE 302b. Generally, the content adaptation engine 302a-b adapts content for delivery to and display by the particular mobile device 300 that requested it, using for example conventional transcoding and adaptation techniques.

As shown in FIG. 3, the mobile device 300 makes requests for content to the content server 301. That content may include, for example, an HTML document representing a web page on the origin site. In many cases, the mobile device 300 will have been directed to the content server 301 by the DNS system of a CDN after the mobile device makes a request for content associated with a domain name associated with an origin server 304, as described above in connection with FIGS. 1-2.

Turning to the information flow illustrated in Fig. 3, in step 1, the content server 301 receives a request for content from the mobile device 300, typically an HTTP 'get' request for content at a particular URL. The content server 301 may check its cache to determine if it already has a copy of the requested content. Assume for purposes of illustration that the requested content has not yet been cached, or was previously cached but has expired due to the lapsing of a time-to-live value for example, then the content server 301 experiences a cache miss. In step 2, the content server 301 issues a request to the origin server 304 for the content, and the origin server 304 responds with the requested content (unadapted) in step 3. To be able to respond to the mobile device 300 with adapted content, the content server 301 makes a request to the content adaptation engine 302a to adapt the content for the particular mobile device 300, based on the characteristics of that mobile device 300. Alternatively, the origin server 304 may use content adaptation engine 302b to return adapted content in step 3, obviating the need for the content server 301 itself to perform the adaptation.

In performing such adaption, which is also referred to as transcoding or "mobilizing" the content, the content adaption engine (whether 302a or b) considers the characteristics of the client device 300.

The client device characteristics can be determined in many different ways, but in one embodiment can be determinate at least in part from the 'user_agent' string supplied by the mobile device 300 in the HTTP/HTTPS header of its request in step 1. In some cases, other request headers or other information from the request may be used instead of or in addition to the user_agent string. In one example, the device characteristics are obtained from a device description resource that maps a user_agent or other information to a set of characteristics.

Known such resources include WURFL (wireless universal resource file) and Device Atlas. It is noted that if the origin server 304 is performing the adaptation with content adaptation engine 302b, then in step 2 (FIG. 3) the content server 301 is preferably configured to send the client device characteristics to the origin server 304 with the request for content.

The process for adapting (sometimes referred to as transcoding or mobilizing) the content may take a variety of forms. It is not limited to a particular technique; it may be performed in any conventional manner, as modified by the teachings hereof. For example, the layout of a requested web page may be adjusted and inline images resized/resampled/recolored in accordance with the display size and capabilities of the target device/browser. Video content may be replaced by representative frame captures. Unsupported content or tags may be omitted, or modified to a supported type (as certain devices/browsers do not support certain types of content, e.g., certain video formats). Content exceeding a threshold size may be blocked. In some cases, a single web page may be broken into multiple pages. Additionally, compression or other data reduction techniques may be employed in order to compensate for a slower network connection to the mobile device.

Following adaptation, the content adaption engine 302a-b returns the adapted content and in step 4 the content server 301 serves it to the mobile device 300. That delivery may trigger the mobile device to make (e.g., in the case of an adapted HTML page) requests for additional content referenced therein, causing steps 1-4 to repeat.

2.1 Caching -- Overview

In one embodiment, to improve the responsiveness of the content server 301, the content

server's HTTP proxy process caches the adapted content returned from content adaptation engine 302a-b. By doing so, the content server 301 need not retrieve the content from the origin server 304, nor call the content adaptation engine 302a to adapt that content, in order to service the second request for the same content. Rather, as previously mentioned, at step 1 of FIG. 3 the content server may check its cache to determine if it already has the adapted content (and, e.g., that it has not expired or is otherwise valid for delivery). If so, it returns the cached adapted content to the mobile device 300.

As a result of this caching of adapted content, the proxy's content cache may hold multiple entries for a given piece of content. These entries correspond to different versions of the content that were adapted for different requesting devices. Thus, for example, multiple versions of a web page may be present in the cache. Preferably, each version is valid to serve not just for an individual device, but a given set or class of requesting devices. Additionally, in this embodiment, the content server 301 maintains a data structure to track content that has been adapted for particular requesting devices. This data structure is used to manage access to the content cache so as to determine, among other things, whether to use cached content to respond to subsequent requests from different requesting devices, based on whether they have characteristics similar to those client devices that were previously serviced. For convenience of description, the data structure is referred to herein as a device equivalency class (DEC) cache, to distinguish it from the proxy's main content cache. With the foregoing by way of introduction, the approach will be described in more detail below.

2.1.1 Caching Device-Adapted Content with DEC Cache

A content server 301 running an HTTP proxy contains consolidated rule-sets for fast matching of user-agent strings against known client devices, upon receipt of a client request. The output of this process may be a simple true/false to determine if the responsive content should be forwarded to a content adaptation engine for creation of adapted content for the given device, such as a mobilized page for a mobile device, or whether the content can be served "as-is" from the content cache. If the latter, then an identification of which cached version of the content to serve can be provided.

Many implementations are possible. In one embodiment, when the content adaptation engine 302a-b returns adapted content to the content server 301, the content server can cache the output in its content cache using the user-agent string in the cache key. This ensures correct

operation but reduces the positive cache hits, since any difference in a user_agent string in a subsequent request results in cache miss. As an improvement, in an alternate embodiment, the content server may use a canonical serialization of the device characteristics for the mobile device, as this would allow identical devices to make cache hits, but given the hundreds of characteristics per mobile device, and the fact that only one minor value difference is necessary to be different with this approach, it is likely that not many would be identical.

Another approach is to enable the content server to know which device characteristics and values were taken into consideration by the content adaptation engine 302a-b in creating the adapted content, so that only such characteristics relevant to the adaptation are considered in device equivalency for caching purposes. For mobile adaptation, this means that characteristics relevant to the mobilization process are used caching the content – typically this is fewer than all of the known characteristics for the device. For example, for a given adaptation, if the only considerations were for example that a device screen was 320 pixels wide and the device supports Javascript, then the content could be cached in the content server in such a way as to allow any device that was also 320px wide and supported Javascript to be immediately returned the same adapted content from the cache (subject to other normal time to live(TTL)/no-store metadata rules), even if it had been originally adapted for another client device with other characteristics not relevant to the adaptation process.

In the context of FIG. 3, above, when the content adaption engine 302a returns the adapted content, it also includes information indicating what characteristics were taken into account. If the content adaptation is being done by the origin server with content adaptation engine 302b, then the origin server returns this information in step 4 along with the adapted content. Note that the origin server might specify a range of characteristics, to indicate that not only that certain characteristics were taken into account, but that the adapted content can be used to serve other devices in the future that fall within the range. In one implementation, this information can be returned as part of a Vary header in HTTP. Hence, the approach may be further generalized to allow content provider customers (via origin servers) to emit headers to the content server to express equivalencies.

This approach takes advantage of the fact that for many groups of mobile and other devices, such as related smartphone versions/models, the output of the adaptation may be identical. This technique is not restricted to naturally related product groups/families and can in fact

match equivalencies across brands and product lines – it comes down to the technical mobilization-relevant characteristics alone.

It should be understood that all adaptation-relevant characteristics need not be considered for adaptation-equivalency amongst devices. For example, in the incremental case, the

5 performance improvement resulting from defining broader equivalency classes (and hence enabling more positive cache hits) may be more desirable than the display improvement that results from using narrower classes defined by a wider array of adaptation-relevant characteristics. Weighting such factors provides a way to tradeoff response speed (from cache hits) versus optimization granularity (from consideration of wider/narrower numbers of
10 characteristics). In some cases, a content provider customer of a CDN may be able to provide weights to such factors and thereby control the foregoing response characteristics of the content servers. In short, number of characteristics to be used can be controlled by configuration parameters.

Hence, according to one embodiment, there is provided a modification to a content server that
15 permits it to take advantage of equivalencies calculated for the first client request above to the content adaptation engine in serving other requests from the content cache. Such equivalencies may be used for cache management as device equivalency classes (DECs). An example of use of a DEC, an example of which is provided below.

2.1.2 Caching Logic Example

20 Abstracted Terminology & Concepts For Purposes of Exemplary Logic

Proxy = Content Server Running Web Proxy Process

CAE = Content Adaptation Engine

Request Attributes = user-agent string

25 Device Characteristics = key/value pairs from a device description resource, such as WURFL or Device Atlas

DEC = Device Equivalency Class

DEC Cache = a data structure that exists on the Proxy which maps each URL (key) to a (value) list of sets of key/value pairs.

30 Example DEC Cache with one URL and two DEC's for it:

/url1 => [[key1=value1,key2=value2], [keyX=valueX,keyY=valueY]]

Examples of key/value pairs are device characteristics relevant to adaptation, e.g., screen width=320px; javascript_support=false. The key typically represents a property of a device, and the value specifies the value or a range of values for that property. Thus, screen_width could be equated to a range such as 100-300px, or 200px[+/-100px], etc. The range may be continuous or discrete. For example, a property be specified with reference to a range of discrete values forming a set, such as OS=iOS, Android, etc, or browser=IE, Firefox,Opera, etc.

FIG. 4 is a flow chart illustrating one embodiment of a process for caching with device equivalency classes. The corresponding steps are described below.

In Step 400, a request is received at the Proxy. Typically the request is an HTTP request and includes user_agent or other client identifying information in a header. In Step 401, the Proxy decides if the URL is in the DEC Cache as a key. FIG. 5 illustrates the DEC cache in a table that contains, for each URL for which there is cached content, a list of DEC's associated therewith that represent the various versions of the content that have been stored in the Proxy's content cache. Note that keying the cache using URLs is typical, but a web content identifier other than a URL could be used.

If the requested URL is found in the DEC Cache, then the collection of DEC's are retrieved in Step 408 and passed to a function that determines whether the given client device (as identified by user_agent string or otherwise) is a member of any DEC. That is, Step 410 represents a membership function that takes as input the DEC Cache value (a list of DEC's) and request attributes (user_agent or otherwise) and gives as output a DEC match or 'null' for no match. The membership function is illustrated in FIG. 6.

Turning to FIG. 6, in Step 600, the Device Characteristics are determined, typically using a service that is more local than the CAE but may be inside or outside the Proxy. This may involve consulting a device description database. In step 601, iterate over the list of DEC's. For each DEC constituent, check the value of the same key from the Device Characteristics with the DEC value and confirm that the values are identical. (Steps 602, 604). Note that a configuration may specify that only a certain level of correspondence is necessary for a

"match." (E.g., device characteristics may not be exact matches but allowed to fall within predetermined tolerances of the DEC, such as a screen width of 330 px being "matched" to a screen width of 320 px in the DEC.) In Steps 606, 608, return either a matching DEC or the absence of a matching DEC.

5 Returning to FIG. 4., in Step 412, determine if the requested device matched to a DEC cached for the URL. If yes, then serve content from main content cache under key URL+DEC. (Step 414 and shown with the dotted lines in FIG. 5.) If not, then move to Step 402, which is the same as if no URL has been found in the DEC cache in the first place at Step 401. In Step 402 and 404, the Proxy goes forward to the CAE for the content. When CAE replies,
10 expect a header in the response such as X-Equivalency-Class: key1=value1,key2=value2. This header has been computed by the CAE as any characteristic that was taken into consideration during the construction of the adapted content. Conversely any characteristic not taken into consideration should be omitted. Canonical order for keys is unnecessary.

In Step 406, store the DEC in the DEC Cache under the key URL. The DEC represents a
15 new DEC to take into account for future requests. Store the content in the main content cache under key URL+DEC. The two stored cached entries in the two caches expire together (the same TTLs). In step 414, serve content from main content cache under key URL+DEC.

It is noted that Steps 402-406 of the above flow illustrate the process of building a table in the 'DEC Cache' at the content server. The DEC Cache is constructed in this case by the CAE
20 reporting back in an HTTP header, although in other embodiments equivalency classes can be emitted by the origin server 304. One potential implementation is to use the HTTP 'Vary' header for this purpose.

Note that in the foregoing implementation, the content cache key of URL + DEC can be used to purge the content cache by URL and/or by device equivalency class. Incorporating the
25 URL+DEC approach into the DEC Cache allows the DEC Cache to be purged by, for example URL (master base purge for all URL) or by specific device equivalency class (using the DEC). The CAE may initiate such a purge, for example, upon determining that it will begin to take into account a different set of adaptation characteristics for given devices previously within a single DEC (in other words, splitting one DEC into multiple DEC's), or
30 vice versa.

2.2 Ambiguity in Matching DEC's

In some cases, a given requesting device may match to more than one DEC in the DEC cache. A variety of algorithms can be used to resolve such an ambiguity. For example, of the matching DEC's, the DEC to use could be (i) the most specific DEC or (ii) the most recent
5 DEC. Alternatively, the upon hitting an ambiguity, the match process could be aborted and the content server could go back to the content adaptation engine to obtain a newly adapted piece of content. Preferably, the approach taken is configurable so that an administrator and/or content provider can select the desired approach for these cases.

3.0 Specifying DEC's In A Control File

- 10 In the above embodiments, the characteristics defining a DEC (e.g., screen_size=320px and so on) were specified by the content adaptation engine, be it a CAE in the content server or the origin server. However, this is not a limitation. The characteristics can be specified by a content provider customer (or a content server operator) separately from the content adaptation process.
- 15 In one embodiment, such characteristics are specified in a control file that is sent to the content server as part of its configuration for a given content provider customer. The control file may be XML-based metadata, such as is described in U.S. Patent No. 7,111,057, the contents of which are hereby incorporated by reference. The control file specifies the DEC's that apply for a given site, a given URL or range of URLs. Upon receiving a request from a
20 given client device, the content server matches the requesting device to one of the DEC's. It then uses the URL+DEC to locate the appropriate cached version of the content. If no such version is cached, then the content server obtains the appropriate version from a content adaptation engine, e.g., CAE 302a-b, as described above. The content server may send to the CAE the device characteristics from the DEC that are to be taken into account when adapting
25 the content.

If none of the DEC's match (or only a catch-all "default" DEC is matching), then the content server can use a default version of the web content from cache or from origin. Preferably, such options are configurable.

- 30 In other embodiments, DEC functionality is provided by a process external to the content server. For example, a DEC service can be provided by another CDN server, or a server

associated with a content provider customer or third-party provider. The content server makes requests to the service to obtain DEC definitions, or to invoke logic that performs the DEC matching function and returns the appropriate DEC.

3.1 Content Provider Configuration

- 5 As noted above, the granularity of the DEC cache function can be configurable on a content provider by content provider basis. For example, a given content provider may desire that device equivalency classes be narrowly defined (defined by more device characteristics), resulting in more finely adapted content for client devices, but potentially reducing responsiveness because of the reduced chance of a cache hit. Another content provider may
10 desire that that device equivalency classes be broader (defined by fewer device characteristics), decreasing the cache footprint and increasing responsiveness because of the increased likelihood of a cache hit, but reducing the precision of the content adaptation.

- Furthermore, in one embodiment, a given content provider can configure each of its sites, or even its pages, with different DEC granularity settings. Hence, for example, a content
15 provider's home page may be associated with broadly defined equivalency classes, perhaps only based on one or two important characteristics. The home page will have a relatively high chance of a cache hit, resulting in higher performance. Pages deeper down in the site may be configured with more narrowly defined equivalency classes. These are merely examples.

20 4.0 Authoring Several Versions of Adapted Content

- Note that the use of a request-time content adaption engine to create adapted versions of content is not limiting. In some cases, the content provider authors several versions of certain URLs. For example, for a given page, the origin server may publish foo_v1.html, foo_v2.html, and so on. One may represent a 'default' page useable when all other DEC's do
25 not apply, but the other versions correspond to a DEC and are adapted for devices defined by those DEC's as part of the authoring process. Upon receiving a request, the content server matches the requesting device to a DEC and then uses that DEC to either retrieve from cache (as previously described, using URL+DEC) or to request that version of the page from the origin. In this case, typically each DEC points to a unique URL on the origin server that
30 holds the appropriate version of the content. Since an origin server typically will not

understand requests for URL+DEC, logic can be employed to convert such a request into an standard URL according to a suitable naming convention. For example, assuming a URL for a given object is www.customer.com/foo.html, URL+DEC1 is mapped to a URL specifying www.customer.com/foo_v1.html, while URL+DEC2 maps to

- 5 www.customer.com/foo_v2.html, and the content server goes forward to request these objects from origin. Alternatively, the metadata control file can specify the forward URL to use for a given DEC.

- In these implementations, preferably the DEC's are specified in the control file (metadata), as previously described. However, it should be understood that the origin server could also be
 10 modified to send the DEC information during operation of the system, thereby building the DEC cache as described earlier in connection with FIG. 4 (e.g., steps 402-406). In effect, instead of adapting content to create the appropriate version at request time with CAE 302b, the origin server can simply select one of the pre-authored, adapted content versions and return that version, along with a list of device characteristics, indicating the range of devices
 15 for which the returned version is valid. The device characteristics returned by the origin may thus be broader than the specific values sent to the origin. For example:

Content server sends characteristics: Device: OS=Android, OSver=3.1.16

Origin server returns with vary header: OS=Android, OSver=3.*.*-4.*.*

Content server sends characteristics : Device_Model=BlackBerry Curve 3200

- 20 Origin server returns with vary header: Model=BlackBerry Curve*

In light of the foregoing, it should be understood from that the teachings hereof are not limited to caching applications but also extend to non-cacheable content, the DEC cache facilitating the forward requests from the proxy for adapted, non-cacheable content at origin.

25 5.0 Further Extensions

The mechanism for determining the list of device characteristics from the user_agent or other information in an HTTP request may be implemented as a lookup in one or more device databases such as WURFL, Device Atlas, or otherwise; it may also be a different mechanism provided by the content delivery provider, the origin, or a third party; the mechanism may or

may not include device database lookup in its operation. A content server may run customer-expressed logic (e.g., via metadata) to facilitate the determination.

Moreover, the teachings hereof may be applied to non-device characteristics gleaned from the client request as well. Example non-device functions include fraud rank mapping,

5 geographic request tagging, or authentication detection.

Hence, a content provider may desire that one version of content be served and cached for requests that exhibit characteristics leading to a low likelihood of fraud, while another is served for request that exhibit characteristics leading to a high likelihood of fraud, as determined by known fraud-detection techniques. Thus a "Fraud-Equivalency Class" or FEC
10 may be employed.

Similarly, a content provider may desire one version of content to be served and cached for requests from authenticated users (e.g., "logged-in" requests) and another version for non-authenticated requests.

For the geographic applications, conventional geo-location functionality can be used to map
15 information in the request, e.g., the client IP address, to an array of geographic characteristics, such as country, state, city, zip, and so on. These geographic characteristics can be organized into 'Geographic Equivalency Classes' and treated as above.

Example :

Content server sends characteristics gleaned from request: Country=US, State=FL

20 Origin server returns with vary header: Country=US, State=FL,GA,SC

Thus, this response indicates that the response is valid for future requests from several different states in the southeast US.

6.0 Computer-Based Implementation

The clients, servers, and other devices described herein may be implemented with
25 conventional computer systems, as modified by the teachings hereof, with the functional characteristics described above realized in software, special-purpose hardware, general-purpose hardware modified or reconfigured by software stored therein for special purposes,

or a combination thereof.

Software may include one or several discrete programs. Any given function may comprise part of any given module, process, execution thread, or other such programming construct. Generalizing, each function described above may be implemented as computer code, namely, as a set of computer instructions, executable in one or more processors to provide a special purpose machine. The code may be executed using conventional apparatus – such as a processor in a computer, digital data processing device, or other computing apparatus – as modified by the teachings hereof. In one embodiment, such software may be implemented in a programming language that runs in conjunction with a DNS-compliant name server (e.g., BIND).

While in some cases above a particular order of operations performed by certain embodiments is set forth, it should be understood that such order is exemplary, as alternative embodiments may perform the operations in a different order, combine certain operations, overlap certain operations, or the like. References in the specification to a given embodiment indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic.

FIG. 7 is a block diagram that illustrates hardware in a computer system 700 upon which such software may run in order to implement embodiments of the invention. The computer system 700 may be embodied in a client device, server, personal computer, workstation, tablet computer, wireless device, mobile device, network device, router, hub, gateway, or other device. Representative machines on which the subject matter herein is provided may be Intel Pentium-based computers running a Linux or Linux-variant operating system and one or more applications to carry out the described functionality.

Computer system 700 includes a processor 704 coupled to bus 701. In some systems, multiple processor and/or processor cores may be employed. Computer system 700 further includes a main memory 710, such as a random access memory (RAM) or other storage device, coupled to the bus 701 for storing information and instructions to be executed by processor 704. A read only memory (ROM) 708 is coupled to the bus 701 for storing information and instructions for processor 704. A non-volatile storage device 706, such as a magnetic disk, solid state memory (e.g., flash memory), or optical disk, is provided and

coupled to bus 701 for storing information and instructions. Other application-specific integrated circuits (ASICs), field programmable gate arrays (FPGAs) or circuitry may be included in the computer system 700 to perform functions described herein.

A peripheral interface 712 communicatively couples computer system 700 to a user display 714 that displays the output of software executing on the computer system, and an input device 715 (e.g., a keyboard, mouse, trackpad, touchscreen) that communicates user input and instructions to the computer system 700. The peripheral interface 712 may include interface circuitry, control and/or level-shifting logic for local buses such as RS-485, Universal Serial Bus (USB), IEEE 1394, or other communication links.

Computer system 700 is coupled to a communication interface 716 that provides a link (e.g., at a physical layer, data link layer, or otherwise) between the system bus 701 and an external communication link. The communication interface 716 provides a network link 718. The communication interface 716 may represent a Ethernet or other network interface card (NIC), a wireless interface, modem, an optical interface, or other kind of input/output interface.

Network link 718 provides data communication through one or more networks to other devices. Such devices include other computer systems that are part of a local area network (LAN) 726. Furthermore, the network link 718 provides a link, via an internet service provider (ISP) 720, to the Internet 722. In turn, the Internet 722 may provide a link to other computing systems such as a remote server 730 and/or a remote client 731. Network link 718 and such networks may transmit data using packet-switched, circuit-switched, or other data-transmission approaches.

In operation, the computer system 700 may implement the functionality described herein as a result of the processor executing code. Such code is typically read from or provided by a non-transitory computer-readable medium, such as memory 710, ROM 708, or storage device 706. Other forms of non-transitory computer-readable media include disks, tapes, magnetic media, CD-ROMs, optical media, RAM, PROM, EPROM, and EEPROM. Any other non-transitory computer-readable medium may be employed. Executing code may also be read from network link 718 (e.g., following storage in an interface buffer, local memory, or other circuitry).

It is noted that any trademarks appearing herein are the property of their respective owners

and used for identification and descriptive purposes only, given the nature of the subject matter at issue, and not to imply endorsement or affiliation in any way.

Having described our invention, we claim:

CLAIMS

1. A method operable in at least one content server for delivering content to a client over a computer network, the at least one content server having at least one processor and memory
5 holding instructions for execution by the at least one processor, the method comprising:
- receiving a request for web content from a first client device;
- determining that the first client device has a set of characteristics that match one of a
10 plurality of sets of stored characteristics for client devices, each of said plurality of sets of stored characteristics being associated with a version of said web content;
- obtaining the version of the web content associated with the matching set of stored characteristics;
- 15 serving the obtained version of the web content to the first client device, in response to the request for web content.
- 20 2. The method of claim 1, further comprising: determining the set of characteristics of the first client device from information in the request.
3. The method of claim 1, further comprising caching, in a cache local to the at least one content server, the version of the web content associated with the matching set of stored
25 characteristics.
4. The method of claim 3, wherein said caching uses a cache key associated with the matching one of the plurality of sets of stored characteristics.
- 30 5. The method of claim 1, wherein
- (i) a given characteristic of the first client device has an property and a value;
- (ii) a given stored characteristic in the matching set of stored characteristics has a property associated with any of a value and a range of values,

and the method further comprises determining whether the given characteristic of the first client device matches the given stored characteristic by: determining whether the value of the given property of the first client device matches the value, or falls within the range of values, of the given stored property.

5

6. The method of claim 5, wherein a given value is any of number, a Boolean, a string.

7. The method of claim 1, wherein obtaining the version of web content associated with the matching set of stored characteristics comprises any of: retrieving that version from
10 cache, retrieving that version from an origin server, retrieving that version from a content adaptation engine.

8. The method of claim 1, wherein at least one of the plurality of sets of stored characteristics are specified in metadata executable at the at least one content server.

15

9. The method of claim 1, wherein at least one of the plurality of sets of stored characteristics is specified in a header received from an origin server.

10. The method of claim 1, wherein at least one of the plurality of sets of stored
20 characteristics is specified in information sent from a content adaptation engine in response to sending the content adaptation engine a request to modify web content to create a version of the web content.

11. The method of claim 10, wherein the stored characteristics represent those
25 characteristics taken into account to modify the web content.

12. The method of claim 1, wherein the first client device comprises a mobile device.

13. The method of claim 1, wherein at least one of the plurality of sets of stored
30 characteristics includes a characteristic specifying a value or range of values for any of the following properties: display size, browser, javascript support, operating system, user_agent.

14. The method of claim 1, further comprising providing a data structure that maps a universal resource locator (URL) to a plurality of versions of that web content, each version being associated with one of the plurality of sets of stored characteristics.

5 15. The method of claim 1, wherein the set of characteristics of the first client device that matches the one of the plurality of sets of stored characteristics represents fewer than all of the known characteristics of the first client device.

10 16. The method of claim 1, wherein the set of characteristics of the first client device that matches the one of the plurality of sets of stored characteristics includes only one characteristic.

17. Computer apparatus for delivering content to a client device over a computer network, the apparatus comprising:

15 circuitry forming at least one processor and memory holding instructions that, upon execution by the at least one processor, provide a special purpose machine that:

(i) receives a request for web content from a first client device;

20 (ii) determines that the first client device has a set of characteristics that match one of a plurality of sets of stored characteristics for client devices, each of said plurality of sets of stored characteristics being associated with a version of said web content;

25 (iii) obtains the version of the web content associated with the matching set of stored characteristics;

(iv) serves the obtained version of the web content to the first client device, in response to the request for web content.

30 18. The apparatus of claim 17, further including instructions that, when executed by the at least one processor, cause the special purposes machine to: determine the set of characteristics of the first client device from information in the request.

19. The apparatus of claim 17, further including instructions that, when executed by the at least one processor, cause the special purposes machine to: cache, in a local cache, the version of the web content associated with the matching set of stored characteristics.

20. The apparatus of claim 19, wherein said caching uses a cache key associated with the matching one of the plurality of sets of stored characteristics.

21. The apparatus of claim 17, wherein

(i) a given characteristic of the first client device has an property and a value;

(ii) a given stored characteristic in the matching set of stored characteristics has a property associated with any of a value and a range of values,

and the apparatus further includes instructions that, when executed by the at least one processor, cause the special purposes machine to: determine whether the given characteristic of the first client device matches the given stored characteristic by: determining whether the value of the given property of the first client device matches the value, or falls within the range of values, of the given stored property.

22. The apparatus of claim 21, wherein a given value is any of number, a Boolean, a string.

23. The apparatus of claim 17, wherein to obtain the version of web content associated with the matching set of stored characteristics, the apparatus retrieves that version from any of: cache, an origin server, and a content adaptation engine.

24. The apparatus of claim 17, wherein at least one of the plurality of sets of stored characteristics are specified in metadata executable by the apparatus.

25. The apparatus of claim 17, wherein at least one of the plurality of sets of stored characteristics is specified in a header received from an origin server.

26. The apparatus of claim 17, wherein at least one of the plurality of sets of stored characteristics is specified in information sent from a content adaptation engine in response to

sending the content adaptation engine a request to modify web content to create a version of the web content.

27. The apparatus of claim 26, wherein the stored characteristics represent those characteristics taken into account to modify the web content.

28. The apparatus of claim 17, wherein the first client device comprises a mobile device.

29. The apparatus of claim 17, wherein at least one of the plurality of sets of stored characteristics includes a characteristic specifying a value or range of values for any of the following properties: display size, browser, javascript support, operating system, user_agent.

30. The apparatus of claim 17, further including instructions that, when executed by the at least one processor, cause the special purposes machine to: provide a data structure that maps a universal resource locator (URL) to a plurality of versions of that web content, each version being associated with one of the plurality of sets of stored characteristics.

31. The apparatus of claim 17, wherein the set of characteristics of the first client device that matches the one of the plurality of sets of stored characteristics represents fewer than all of the known characteristics of the first client device.

32. The apparatus of claim 17, wherein the set of characteristics of the first client device that matches the one of the plurality of sets of stored characteristics includes only one characteristic.

33. The apparatus of claim 17, wherein the apparatus is a proxy server in a content delivery network (CDN) operated by a CDN service provider.

34. A method operable in a content server for delivering content to a client device over a computer network, the content server having a processor and memory holding instructions for execution by the processor, the method comprising:

receiving a request for web content from a first client device;

obtaining a version of the web content that is modified from its original form for the first client device, the modification taking into account one or more characteristics of the first client device;

5

caching the modified web content;

receiving a request for the web content from a second client device;

10 determining whether the modified web content is valid to be served to the second client device, the determination including a determination of whether the second client device has at least one of the one or more characteristics;

if the modified web content is valid to be served to the second client device, sending the
15 second client device the modified web content from the cache.

35. The method of claim 34, wherein the first client device and the second client device each comprise a mobile device.

20 36. The method of claim 34, wherein the determination of whether the modified web content is valid to be served to the second client device includes at least a determination of whether the second client device has all of the one or more characteristics.

25 37. The method of claim 34, wherein each of one or more characteristics comprises a property associated with any of a value and a range of values.

38. The method of claim 34, wherein the one or more characteristics comprise at least one of: of display size, installed browser, javascript support, operating system, user_agent string.

30 39. The method of claim 34, wherein obtaining the modified web content comprises modifying the web content with the content server.

40. The method of claim 34, comprising obtaining the modified web content from a remote server, the remote server being an origin server associated with a content provider customer of a content delivery network service provider who operates the content server.

5 41. The method of claim 34, further comprising mapping information in a header of the request to set of characteristics of the first client device that include the one or more characteristics.

10 42. The method of claim 34, consulting a data structure that maps web content identifiers to a plurality of versions of modified web content, each version being identified by one or more characteristics of a client device that were taken into account in modifying the web content to create the given version.

15 43. The method of claim 34, wherein caching the modified web content comprises caching the modified web content using a cache key based on the one or more characteristics of the first client device taken into account to create the modified web content.

20 44. Computer apparatus for delivering content to a client over a computer network, the apparatus comprising:

at least one processor and memory holding instructions that, upon execution by the at least one processor, provide a special purpose machine that:

25 (i) receives a request for web content from a first client device;

(ii) obtains a version of the web content that is modified from its original form for the first client device, the modification taking into account one or more characteristics of the first client device;

30 (iv) caches the modified web content in a local cache of the apparatus;

(v) receives a request for the web content from a second client device;

(vi) determines whether the modified web content is valid to be served to the second client device, the determination including a determination by the apparatus of whether the second client device has at least one of the one or more characteristics;

(vii) if the modified web content is valid to be served to the second client device sends the second client device the modified web content from the cache.

45. The apparatus of claim 44, wherein the first client device and the second client device each comprise a mobile device.

46. The apparatus of claim 44, wherein the determination of whether the modified web content is valid to be served to the second client device includes at least a determination of whether the second client device has all of the one or more characteristics.

47. The apparatus of claim 44, wherein each of one or more characteristics comprises a property associated with any of a value and a range of values.

48. The apparatus of claim 44, wherein the one or more characteristics comprise at least one of: display size, installed browser, javascript support, operating system, user_agent string.

49. The apparatus of claim 44, wherein upon execution the instructions provide a special purposes machine that obtains the modified web content by modifying the web content with the special machine.

50. The apparatus of claim 44, wherein upon execution the instructions provide a special purposes machine that obtains the modified web content from a remote server, the remote server being an origin server associated with a content provider customer of a content delivery network service provider who operates the content server.

51. The apparatus of claim 44, wherein upon execution the instructions provide a special purposes machine that maps information in a header of the request to set of characteristics of the first client device that include the one or more characteristics.

5 52. The apparatus of claim 44, wherein the apparatus further comprises a data structure that maps web content identifiers to a plurality of versions of modified web content, each version being identified by one or more characteristics of the client device taken into account in modifying the web content to create the given version.

10 53. The apparatus of claim 44, wherein the apparatus comprises a cache for the modified web content, the cache having a cache key based on the one or more characteristics of the first client device taken into account to create the modified web content.

54. A method operable in a content server for delivering content to a client over a
15 computer network, the content server having a processor and memory holding instructions for execution by the processor, the method comprising:

receiving a request for web content from a first client device;

20 obtaining a version of the web content that is modified from its original form for the first client device, the modification taking into account one or more characteristics of the first client device, to create modified web content;

25 assigning a class identifier to the modified web content, the class identifier representing one or more characteristics of a client device taken into account to create the modified web content;

caching the modified web content;

30 receiving a request for the web content from a second client device;

determining whether the second client device falls within the class identified by the class identifier;

sending the second client device the modified web content from the cache, if the second client device is within the class.

5 55. The method of claim 54, wherein the first client device and the second client device each comprise a mobile device.

56. The method of claim 54, comprising caching the modified web content using a cache key reflecting the class identifier.

10

57. The method of claim 54, wherein the number of characteristics included in assigning the class is configurable.

58. The method of claim 54, wherein the number of characteristics included in assigning
15 the class is configurable and represents fewer than the number of characteristics actually taken into account to create the modified web content.

59. The method of claim 54, wherein the one or more characteristics of the first client device comprise at least one of: of display size, browser feature, javascript support, operating
20 system, browser characteristic, user_agent string.

60. The method of claim 54, comprising obtaining the modified web content from a remote server, the remote server being an origin server associated with a content provider customer of a content delivery network service provider who operates the content server.

25

61. The method of claim 54, further comprising modifying the web content with the content server.

62. Computer apparatus for delivering content to a client over a computer network, the
30 apparatus comprising:

at least one processor and memory holding instructions that, upon execution by the at least one processor, provide a special purpose machine that:

(i) receives a request for web content from a first client device;

5 (ii) obtaining a version of the web content that is modified from its original form for the first client device, the modification taking into account one or more characteristics of the first client device, to create modified web content;

10 (iii) assigns a class identifier to the modified web content, the class identifier representing one or more characteristics of a client device taken into account to create the modified web content;

(iv) caches the modified web content in a local cache of the apparatus;

15 (v) receives a request for the web content from a second client device;

(vi) determines whether the second client device falls within the class identified by the class identifier;

20 (vii) sends the second client device the modified web content from the cache, if the second client device is within the class.

63. The apparatus of claim 62, wherein the first client device and the second client device each comprise a mobile device.

25 64. The apparatus of claim 62, wherein the local cache of the comprises a cache key reflecting the class identifier.

65. The apparatus of claim 62, wherein the number of characteristics included in assigning the class is configurable.

30 66. The apparatus of claim 1, wherein the number of characteristics included in assigning the class is configurable and represents fewer than the number of characteristics actually taken into account to create the modified web content.

67. The apparatus of claim 62, wherein the one or more characteristics of the first client device comprise at least one of: of display size, browser feature, javascript support, operating system, browser characteristic, user_agent string.


5


68. The apparatus of claim 62, wherein upon execution the instructions provide a special purpose machine that obtains the web content from a remote server, the remote server being an origin server associated with a content provider customer of a content delivery network service provider who operates the content server.

10

69. The apparatus of claim 62, wherein upon execution the instructions provide a special purposes machine that obtains the modified web content by modifying the web content with the special machine.

15



*FIG. 2*




FIG. 3

4 / 7




FIG. 4 – DEC Caching Process




FIG. 5




FIG. 6 – DEC Matching Service

7 / 7




FIG. 7