



- (51) International Patent Classification: *G06F 15/16* (2006.01)
- (21) International Application Number: PCT/US2012/052162
- (22) International Filing Date: 23 August 2012 (23.08.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 13/237,722 20 September 2011 (20.09.2011) US
- (71) Applicant (for all designated States except US): **IN-START LOGIC, INC.** [US/US]; 2307 Leghorn Street, Mountain View, CA 94040 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **KOLAM, Hariharan** [IN/US]; 480 San Antonio Road, Suite 210, Mountain View, CA 94040 (US). **MITAL, Manav Ratan** [IN/US]; 480 San Antonio Road, Suite 210, Mountain View, CA 94040 (US). **VENKAT, Raghu Batta** [IN/US]; 480 San

Antonio Road, Suite 210, Mountain View, CA 94040 (US).

(74) Agent: **TONG, Grace**; Van Pelt, Yi & James LLP, 10050 N. Foothill Blvd., Suite 200, Cupertino, CA 95014 (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, BN, 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,

[Continued on next page]

(54) Title: APPLICATION ACCELERATION WITH PARTIAL FILE CACHING

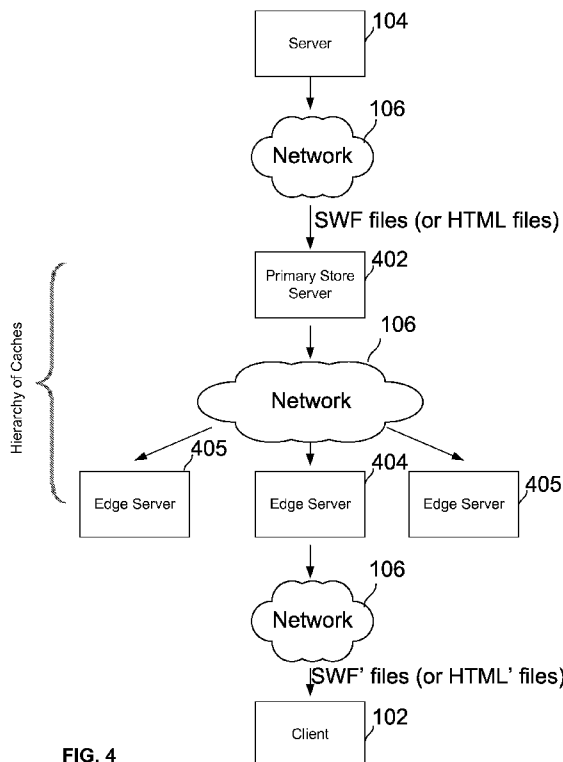


FIG. 4

(57) Abstract: A method of distributing content is disclosed. The method includes receiving by an edge server an indication indicating that a client has sent a request for a file that includes content related to executing an application. The method includes analyzing interaction with the client by the edge server to determine a first portion of the file that is likely to be required by the client for executing the application sooner than a second portion of the file. The method further includes generating by the edge server a modified file that includes the first portion of the file. The method further includes transmitting the modified file to the client.

EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, **Published:**
LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, — *with international search report (Art. 21(3))*
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

APPLICATION ACCELERATION WITH PARTIAL FILE CACHING

BACKGROUND OF THE INVENTION

[0001] Some computer applications require the downloading of resources through a client-server connection over a network before the applications can start. The long startup time may deter some end users from using the applications. Therefore, improved techniques for downloading the resources for these applications would be desirable.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Various embodiments of the invention are disclosed in the following detailed description and the accompanying drawings.

[0003] Figure 1 is a block diagram illustrating an embodiment of a client-server connection through a network for downloading files or content for starting and loading an application on the client side.

[0004] Figure 2 is a diagram illustrating an embodiment of a Small Web Format (SWF) file.

[0005] Figure 3 is a diagram illustrating an embodiment of a webpage described by an HTML file.

[0006] Figure 4 is a diagram illustrating an embodiment of a client-server connection through a network for downloading files or content for starting or loading an application on the client side.

[0007] Figure 5 is a diagram illustrating an embodiment of a process 500 for downloading a file from a server to a client, wherein the file is for starting and loading an application on the client side.

[0008] Figure 6 is a diagram illustrating an embodiment of a profiling process 600 for determining the initial portions of original startup SWF files that are sufficient to permit initial functioning of the application in process 500.

DETAILED DESCRIPTION

[0009] The invention can be implemented in numerous ways, including as a process; an apparatus; a system; a composition of matter; a computer program product embodied on a computer readable storage medium; and/or a processor, such as a processor configured to execute instructions stored on and/or provided by a memory coupled to the processor. In this specification, these implementations, or any other form that the invention may take, may be referred to as techniques. In general, the order of the steps of disclosed processes may be altered within the scope of the invention. Unless stated otherwise, a component such as a processor or a memory described as being configured to perform a task may be implemented as a general component that is temporarily configured to perform the task at a given time or a specific component that is manufactured to perform the task. As used herein, the term 'processor' refers to one or more devices, circuits, and/or processing cores configured to process data, such as computer program instructions.

[0010] A detailed description of one or more embodiments of the invention is provided below along with accompanying figures that illustrate the principles of the invention. The invention is described in connection with such embodiments, but the invention is not limited to any embodiment. The scope of the invention is limited only by the claims and the invention encompasses numerous alternatives, modifications and equivalents. Numerous specific details are set forth in the following description in order to provide a thorough understanding of the invention. These details are provided for the purpose of example and the invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the invention is not unnecessarily obscured.

[0011] Figure 1 is a block diagram illustrating an embodiment of a client-server connection through a network for downloading files or content for starting and loading an application on the client side. As shown in Figure 1, a client 102 is connected to a server 104 through a network 106. However, any number of clients and servers may be connected through network 106. Network 106 may be any combination of public or private networks, including intranets, LANs, WANs, and the Internet. Server 104 sends one or more files or other content through network 106 to client 102. In some embodiments, the one or more files

or content sent by server 104 to client 102 are files or content required by client 102 before a process can start to run.

[0012] In some embodiments, the process that requires the one or more files from server 104 is a game application. For example, the game application may be a browser game, i.e., a computer game played over the Internet using a web browser. Browser games include social network games, such as FarmVille, Mafia Wars, FrontierVille, and the like. These social network games are online games that are distributed through social networks, e.g., Facebook, and may feature multiplayer gameplay mechanics. The popular browser-based social network game, FarmVille, is used hereinafter as an example for the process or application described above. Farmville is selected for illustration purposes only; accordingly, the present application is not limited to this specific example only.

[0013] FarmVille is a farm simulation game available as an application on the social networking website Facebook. The game allows members of Facebook to manage a virtual farm by planting, growing, and harvesting virtual crops and trees or raising livestock. Before the application can start to run on the client side, one or more SWF (Small Web Format) files are sent by server 104 to client 102 through network 106. These SWF files provide the application on the client side with resources it needs for starting and running the game. SWF can be used to present vector-based objects and images, audio, video, text, and other forms of interaction with the end user. After the one or more SWF files are received on the client side, the SWF files can be played by an Adobe Flash Player as a browser plugin. For example, text can be displayed, audio can be played, and movie clips can be played as the game begins.

[0014] Figure 2 is a diagram illustrating an embodiment of a SWF file. A SWF file has different sections, including the character, symbol table, and timeline sections. The character section is used to store contents, such as images, movie clips, audio, and the like. For example, the character section of a SWF file sent by a FarmVille server may include image files (e.g., bitmap or jpeg files) and movie clips files (e.g., mpeg files) of various crops that may be planted on the farm, farm animals that can be raised on the farm, and the like. The timeline section has code to define how the animation of the game should progress over time. For example, the timeline may specify that a particular farm animal, e.g., a sheep, should be displayed at time t_0 at a particular position on the screen. The symbol section includes code, e.g., actionscripts, for loading the image file (e.g., sheep.jpg or sheep.bmp) of the farm animal from the character section. The loaded image can then be displayed on the

screen at the specified time and location. The symbol table section also includes character ID and class name information.

[0015] The number of SWF files for starting up FarmVille on the client side and the content of those SWF files are different for different player levels. As the player level increases, the number of SWF files for starting up FarmVille on the client side and the size of those SWF files may increase. This is because some characters are only made available to players at or above a certain player level.

[0016] Although each player level has a different set of startup SWF files, players at the same player level receive an identical set of SWF files. The set of startup SWF files for a particular player level is updated periodically to incorporate new features or new characters for the game. Since the set of startup SWF files for a particular player level is not dynamically generated for a particular game player by the game publisher (i.e., Zynga), some game players may receive characters in the set of SWF files that are never used during their respective games. As the size (in bytes) of each image or movie clip file is typically large, receiving unnecessary characters embedded in the set of startup SWF files on the client side over a network translates to a long startup or loading time for the game, especially for higher player levels. In some cases, the long startup time can even deter some game players from playing the game. Therefore, a technique to speed up the startup time for a process or application on the client side, without compromising the game player's experience, would be desirable.

[0017] In some embodiments, the process that requires the one or more files or content from server 104 (as shown in Figure 1) is a web browser accessing a webpage. The webpage may be described by different markup languages, including Hypertext Markup Language (HTML), Extensible Markup Language (XML), and the like. The webpage may also be described by different scripting languages, including JavaScript Object Notation (JSON), and the like. HTML is used hereinafter as an example of the various languages for describing webpages. Note that the examples of HTML are selected for illustration purposes only; accordingly, the present application is not limited to these specific examples only.

[0018] Figure 3 is a diagram illustrating an embodiment of a webpage described by an HTML file. To display the webpage, a web browser sends a Hypertext Transfer Protocol (HTTP) request message to a content server requesting the HTML webpage file. After the

content server locates the requested HTML webpage file, the content server returns the requested HTML webpage file in an HTTP response message to the web browser. The web browser parses the received webpage file and begins to render a portion of the webpage, e.g., a text portion, on a screen.

[0019] As shown in Figure 3, the webpage file may include one or more scripts. For example, the webpage may include a number of icons or buttons for a webpage viewer to click on. A script associated with a specific icon or button is executed on the client side only if the webpage viewer clicks on the corresponding icon or button. Since some of the icons or buttons on the webpage are never clicked by some webpage viewers when they are accessing the webpage, the webpage file received by these webpage viewers include unnecessary information and content that is sent through network 106. Receiving unnecessary information and content embedded in the HTTP response message on the client side over a network translates to a longer loading time for the webpage. Therefore, a technique to speed up the loading time for a process or application on the client side, without compromising the webpage viewer's experience, would be desirable.

[0020] Figure 4 is a diagram illustrating an embodiment of a client-server connection through a network for downloading files or content for starting or loading an application on the client side. Continuing with the FarmVille illustrative example above, client 102 no longer receives the startup files for running the FarmVille browser game directly from server 104 (as shown in Figure 1) but from an edge server 404 other than server 104. In some embodiments, the edge server 404 that the client is receiving the startup files from is one of a plurality of edge servers (404 and 405) in a content delivery network or content distribution network (CDN). For example, edge server 404 that is connected to client 102 is the edge server that is located geographically closest to client 102.

[0021] With continued reference to Figure 4, server 104 has access to the original startup SWF files provided by the game publisher of FarmVille (i.e., Zynga). These original startup SWF files are the SWF files that are normally downloaded by the user's browser when the game is started. As shown in Figure 4, a primary store server 402 periodically obtains the most updated version of these original startup SWF files from server 104. In some embodiments, primary store server 402 stores a copy of the original startup SWF files in a local cache/memory of primary store server 402. In some embodiments, primary store server 402 breaks down the SWF files into portions, e.g., into one or more partial SWF files.

The one or more partial SWF files can be non-overlapping or overlapping. The degree of granularity of the partial SWF files can be different. For example, a partial SWF file may be an image file included in the character section of a SWF file. In another example, a partial SWF file is the entire SWF file itself. In yet another example, a partial SWF file is a skeleton SWF file (referred to hereinafter as a skeleton) with most or all of the characters removed. These one or more partial SWF files can be distributed among the local cache/memory of primary store server 402 and the local caches of one or more of the plurality of edge servers (404 and 405), as will be described in greater detail below.

[0022] With continued reference to Figure 4, edge server 404 connected to client 102 includes a transformation module for generating modified startup SWF files (denoted SWF') for a particular player level and transmitting the SWF' files to client 102 upon the client's request. For example, if server 104 provides three startup SWF files (SWF1, SWF2, and SWF3) for starting the game at a particular player level, then three modified startup SWF' files (SWF1', SWF2', and SWF3') are dynamically generated by edge server 404 and sent to client 102. In some embodiments, the modified startup SWF' files are generated based on one or more partial SWF files. For example, a modified startup SWF' file is generated by combining multiple partial SWF files together, e.g., by combining a skeleton with a plurality of image files together.

[0023] The modified startup SWF' files are smaller in size (measured in bytes) than the original startup SWF files, thereby reducing the downloading time of the files and thus the startup time of the application. The modified startup SWF' files contain portions of the original startup SWF files. These initial portions provide resources for the initial functioning of the application. For example, only a portion of the characters embedded in the original startup SWF files are required for the initial functioning of the application. The initial portions for each player level and each player (or a group of players) may be determined empirically. For example, the initial portions may be determined based on statistics collected by the edge servers (404 and 405) in a profiling process, which will be described in greater detail below.

[0024] As described above, partial SWF files may be distributed among the local cache/memory of primary store server 402 and the local caches of one or more of the plurality of edge servers (404 and 405). The local cache of primary store server 402 and the local caches of the plurality of edge servers (404 and 405) form a hierarchy of caches for

storing the partial SWF files. In some embodiments, a particular partial SWF file is determined to be stored in a particular local cache at least in part based on the profiling process. For example, if a partial SWF file is determined by the profiling process that the partial SWF file is frequently used for generating modified SWF' files for different clients, then a copy of the partial SWF file may be stored lower in the hierarchy of caches with relatively shorter latency paths to the clients, e.g., at one or more of the plurality of edge servers (404 and 405). In this case, since multiple copies are stored, more cache/memory space is required. However, the time needed for generating modified SWF' files by an edge server will be reduced because the partial SWF file does not have to be fetched from a remote location. Conversely, if a partial SWF file is determined by the profiling process that the partial SWF file is infrequently used for generating modified SWF' files for different clients, then a copy of the partial SWF file may be stored higher in the hierarchy of caches with relatively higher latency paths to the clients, e.g., at primary store server 402 only. In this case, since only a single copy is stored, less cache space is required. However, the time for generating modified SWF' files by an edge server will be increased because edge server 404 needs to fetch the partial SWF file prior to the generation of the modified SWF' files.

[0025] In some embodiments, the game publisher may include code in the game application to be executed on the game player's machine (e.g., a laptop or smartphone) such that when the game is started, the browser will request (e.g., using HTTP requests) for startup files from primary store server 402 or from one of the plurality of edge servers (404 and 406), instead of from server 104. However, other ways to redirect the file requests may be used as well.

[0026] Figure 5 is a diagram illustrating an embodiment of a process 500 for downloading a file from a server to a client, wherein the file is for starting and loading an application on the client side. At 502, an indication indicating that a client has sent a request for a file that includes content related to executing an application is received by an edge server. For example, the request may be an HTTP request for a startup SWF file containing content that is related to executing an application, such as a game application. In another example, the request is an HTTP request for an HTML webpage file. In some embodiments, the request is received by the edge server, and the indication is the request itself. In some embodiments, the request is received by a primary store server (e.g., primary store server 402

in Figure 4) and an indication is sent from the primary store server to the edge server indicating that the client has sent the request for the file.

[0027] At 504, the edge server analyzes its interaction with the client to determine a first portion of the file that is likely to be required by the client for executing the application sooner than a second portion of the file. For example, the edge server analyzes its interaction with the client as part of a profiling process, as will be described in greater detail below. By analyzing its interaction with the client, the edge server may determine that the first portion of the file includes a subset of the characters embedded in an original startup SWF file provided by the game publisher. Continuing with the FarmVille illustrative example above, the first portion of the file may include the image or movie clip files of the most commonly planted crops or most commonly raised farm animals, background image files, the FarmVille logo, and the like. In other embodiments, the first portion of the file may include the image or movie clip files of the crops or animals that were previously planted or raised by a game player associated with the client.

[0028] At 506, the edge server generates a modified file that includes the first portion of the file, but not the second portion of the file. In some embodiments, the modified file is generated dynamically for a particular client. In some embodiments, the modified file is generated for a group of clients sharing similar characteristics. In the FarmVille example, the image or movie clip files of some of the less commonly planted crops or less commonly raised farm animals are not embedded in the character section of the modified startup SWF' file. In some embodiments, the first portion of the file is sufficient to permit initial functioning of the application. At 508, the generated modified file is transmitted to the client by the edge server.

[0029] With continued reference to Figure 5, after the modified file is received, the application can be executed at 510. Continuing with the FarmVille illustrative example above, a game player associated with the client may begin to plant crops on the virtual farm, and so forth. For example, if the game player chooses to plant a particular crop and the image file for the crop is embedded in the character section of the modified startup SWF' file, then the application can begin to access the image file locally by accessing the modified startup SWF' file. If, however, the game player chooses to plant a particular crop and the image file for the crop is not embedded in the character section of the modified startup SWF' file, then the image file is accessed using a placeholder in the modified startup SWF' file.

[0030] In particular, the modified file is formatted with at least one placeholder for the second portion of the file. The at least one placeholder allows the application to access the second portion of the file. In some embodiments, the placeholder may be program code, data structures, interfaces, or a combination thereof, that allow the application to access the second portion of the file. As will be described in greater detail below, the second portion of the file may be fetched from one of the following: 1) one of the servers (the primary store server or the edge servers) through the network, or 2) a local cache, file system, or memory on the client device.

[0031] In some embodiments, the modified startup SWF' file includes program code, data structures, or interfaces from the original SWF file and also new code, data structures, or interfaces that are added by the edge server. In a SWF or SWF' file, a character is represented by a class. The class of a character is instantiated by calling the function "new()" before the character is used in the application. Because a particular character may not be actually embedded in the character section of the modified SWF' file, a new class for the character is created. The new class may import the same interfaces of the original class. Program code for loading the image includes program code from the original SWF file and also new code added by the edge server. For example, the original program code that uses the character will call into the new code such that before a character is loaded, the character is looked up in one of the following three locations: 1) the character section of the modified SWF' files, 2) one of the servers (the primary store server or the edge server) through the network, or 3) a local cache, file system, or memory associated with the client.

[0032] The first location is the character section of the modified SWF' files. In the FarmVille example, image or movie clip files of some of the most commonly planted crops or most commonly raised farm animals, background image files, the FarmVille logo, and the like, can be embedded in the character section of the modified startup SWF' file.

[0033] The second location is one of the servers (the primary server or the edge servers). In the FarmVille example, image or movie clip files of some of the less commonly planted crops or less commonly raised farm animals are not embedded in the character section of the modified startup SWF' file. Therefore, if a game player chooses a character that falls into this category, the client will not be able to access its content from the modified startup SWF' file. Instead, the client needs to explicitly request (e.g., using HTTP requests) one of the servers to send the content for the character over the network. In this case, the

game player enjoys a shorter startup time at the beginning of the game, but may experience some network latency at a later stage. After the content of the character is received, it is loaded by the application for immediate use. The content is also stored locally in the cache, in a file, or memory associated with the client for later retrieval.

[0034] The third location is the local cache, file system, or memory on the client device. In the FarmVille example, sometimes the contents for some of the characters are already stored in the local cache, file system, or memory of the game player's device (e.g., a laptop or smartphone) before a game is started. For example, if the game has been played on the device before, contents for some of the characters that have been previously fetched from one of the servers can be found in the local cache, file system, or memory on the device.

[0035] Figure 6 is a diagram illustrating an embodiment of a profiling process 600 for determining the initial portions of original startup SWF files that are sufficient to permit initial functioning of the application in process 500. At 602, a subset of client-server connections are selected for determining the initial portions of the original startup SWF files that are sufficient to permit initial functioning of the application. In some embodiments, only client-server connections that correspond to a specific client are selected, and the profiling process is for profiling a single client/game player. The results of the profiling process may be stored in a user metadata file or log. In some embodiments, only client-server connections that correspond to a specific group of clients are selected, and the profiling process is for profiling a single group of clients/game players. The results of the profiling process may be stored in a group-user metadata file or log. In some embodiments, the subset of client-server connections is in part selected based on the connection speed of those connections. For example, high-speed connections may be preferred in some embodiments. In some embodiments, the profiling process 600 is performed during a trial period, and the subset of client-server connections corresponds to the connections made during the trial period.

[0036] At 604, for each connection, the player level is determined in part based on the number of startup SWF files requested by the client. Using the FarmVille example, the number of SWF files for starting up the game on the client side and the contents of those SWF files are different for different player levels. Therefore, based on the number of startup SWF files requested, the server may determine the player level.

[0037] At 606, for each connection, modified SWF' files are sent with no characters embedded in the character sections of the modified SWF' files. Since the modified SWF' files do not have any characters embedded in the character sections of the files, the client needs to request for the characters that the application needs explicitly.

[0038] At 608, for each connection, requests for characters from the client are received. For example, if the game player chooses to plant a particular crop and the image file for the crop is not embedded in the character sections of the modified startup SWF' files, then the image file is accessed using a placeholder in one of the modified startup SWF' files. Assuming that the image file is not found locally, a request for the image file is sent through the network to the server.

[0039] At 610, modified startup SWF' files are generated based on the received requests and the determined player level. For example, if a particular crop is requested by 70% or more of the connections at player level #2, then the content corresponding to the crop is included in the character section of one of the modified startup SWF' files for player level #2; otherwise, the content corresponding to that crop is not included. In this way, modified SWF' files for different player levels can be generated based on the collected statistics. In some embodiments, the statistics collected by process 600 above can additionally be used for determining the distribution of partial SWF files in the hierarchy of caches as shown in Figure 4. For example, if a partial SWF file is determined by process 600 that the partial SWF file is frequently used for generating modified SWF' files for different clients, then a copy of the partial SWF file may be stored lower in the hierarchy of caches, e.g., at one or more of the plurality of edge servers (404 and 405).

[0040] In some embodiments, after the modified startup SWF' files are downloaded and the application is up and running, additional resources can be pushed from edge server 404 to client 102. These additional resources are pushed to the client even if the client has not specifically requested for them. In some embodiments, the profiling process 600 may be used to determine the resources that are pushed and the order in which the resources are pushed. For example, characters that are frequently requested by game players at a particular player level but are not included in the modified start SWF' files can be pushed to the client using this method. For example, characters that are frequently requested at an earlier stage of the player level are pushed earlier. Pushing resources to the client trades increased bandwidth requirements for reduced network latency. Therefore, the determination of

whether to push resources to the client this way or what kind of resources to push can be based on different factors. For example, the determination can be based on the preferences of the game developer and/or the game player, the speed of the connection, and the device (e.g., a smartphone or laptop) on which the application is run.

[0041] As described above, the local caches of the various servers can be used to store partial files in different content granularity. For example, compared to other techniques in which the entire SWF file is cached, the servers may cache at the level of characters. The content granularity of the CDN is a character. If a character changes, only the character needs to be distributed, as opposed to the entire SWF file. In some embodiments, the quality of the characters can be adapted based on the network bandwidth. For example, when the available bandwidth is limited, the image file of a character sent in a modified startup SWF file can be a compressed image file.

[0042] Distributing content dynamically as described above reduces the quantity and the volume of the content that is distributed to the end-users. As a result, the downloading time and startup time are reduced. Because of the reduced network activities, battery consumption by the user devices (e.g., phones and tablets) is also reduced. In some embodiments, the client is installed in the firmware of the user devices, thereby providing faster and improved responses.

[0043] Although the foregoing embodiments have been described in some detail for purposes of clarity of understanding, the invention is not limited to the details provided. There are many alternative ways of implementing the invention. The disclosed embodiments are illustrative and not restrictive.

[0044] WHAT IS CLAIMED IS:

CLAIMS

1. A method of distributing content, comprising:
 - receiving by an edge server an indication indicating that a client has sent a request for a file that includes content related to executing an application;
 - 5 analyzing interaction with the client by the edge server to determine a first portion of the file that is likely to be required by the client for executing the application sooner than a second portion of the file;
 - generating by the edge server a modified file that includes the first portion of the file;
 - and
 - 10 transmitting the modified file to the client.
2. The method of claim 1, wherein the first portion of the file includes a partial file corresponding to the file, and wherein the partial file is stored in a cache within a hierarchy of caches, and wherein a level of the cache within the hierarchy of caches is determined at least in part based on the analyzed interaction with the client.
- 15 3. The method of claim 2, wherein the determining of the level of the cache within the hierarchy of caches further comprises determining the level of the cache at least in part based on how likely the partial file is required by the client and other clients for executing the application.
4. The method of claim 2, wherein a higher level within the hierarchy of caches
- 20 corresponds to a higher latency path to the client.
5. The method of claim 2, wherein the cache is located at one of the following: a primary store server, the edge server, or other edge servers.
6. The method of claim 5, wherein the primary store server has access to the file, and wherein the partial file is distributed to the cache from the primary store server.
- 25 7. The method of claim 1, wherein the first portion of the file is sufficient to permit initial functioning of the application, and wherein the modified file is formatted with at least one placeholder for the second portion of the file, and wherein the application can begin to access the first portion of the file by accessing the modified file, and wherein the application can access the second portion of the file using the at least one placeholder.
- 30 8. The method of claim 7, wherein accessing the second portion of the file using the at least one placeholder comprises accessing the portion from one of the following: a cache

located on a device on which the client is run, a file system located on the device on which the client is run, or a memory located on the device on which the client is run.

9. The method of claim 7, wherein accessing the second portion of the file using the at least one placeholder comprises sending a request for the portion by the client.

5 10. The method of claim 1, wherein the first portion of the file is based on a state of the application.

11. The method of claim 10, wherein the application is a game application, and wherein the state of the application is a player level.

10 12. The method of claim 1, wherein the first portion of the file comprises likely to be used portions of the file.

13. The method of claim 1, wherein the file is a Small Web Format (SWF) file.

14. The method of claim 1, wherein the file includes a script, and wherein the modified file includes a modified script.

15 15. The method of claim 1, wherein the file includes a first class, and wherein the modified file includes a second class that imports the first class.

16. The method of claim 1, further comprising transmitting a subset of the second portion of the file to the client prior to the client requesting the subset.

17. The method of claim 16, wherein the subset is in part based on how likely the subset is used by the application.

20 18. An edge server for distributing content, comprising:
a processor configured to:

receive an indication indicating that a client has sent a request for a file that includes content related to executing an application;

25 analyze interaction with the client to determine a first portion of the file that is likely to be required by the client for executing the application sooner than a second portion of the file;

generate a modified file that includes the first portion of the file; and

transmit the modified file to the client; and

30 a memory coupled to the processor and configured to provide the processor with instructions.

19. The edge server of claim 18, wherein the first portion of the file includes a partial file corresponding to the file, and wherein the partial file is stored in a cache within a hierarchy of caches, and wherein a level of the cache within the hierarchy of caches is determined at least in part based on the analyzed interaction with the client.
- 5 20. The edge server of claim 19, wherein the determining of the level of the cache within the hierarchy of caches further comprises determining the level of the cache at least in part based on how likely the partial file is required by the client and other clients for executing the application.
21. The edge server of claim 19, wherein a higher level within the hierarchy of caches
10 corresponds to a higher latency path to the client.
22. The edge server of claim 19, wherein the cache is located at one of the following: a primary store server, the edge server, or other edge servers.
23. The edge server of claim 22, wherein the primary store server has access to the file, and wherein the partial file is distributed to the cache from the primary store server.
- 15 24. The edge server of claim 18, wherein the first portion of the file is sufficient to permit initial functioning of the application, and wherein the modified file is formatted with at least one placeholder for the second portion of the file, and wherein the application can begin to access the first portion of the file by accessing the modified file, and wherein the application can access the second portion of the file using the at least one placeholder.
- 20 25. The edge server of claim 24, wherein accessing the second portion of the file using the at least one placeholder comprises accessing the portion from one of the following: a cache located on a device on which the client is run, a file system located on the device on which the client is run, or a memory located on the device on which the client is run.
26. The edge server of claim 24, wherein accessing the second portion of the file using
25 the at least one placeholder comprises sending a request for the portion by the client.
27. The edge server of claim 18, wherein the first portion of the file is based on a state of the application.
28. The edge server of claim 27, wherein the application is a game application, and wherein the state of the application is a player level.
- 30 29. The edge server of claim 18, wherein the first portion of the file comprises likely to be used portions of the file.

30. The edge server of claim 18, wherein the file is a Small Web Format (SWF) file.
31. The edge server of claim 18, wherein the file includes a script, and wherein the modified file includes a modified script.
32. The edge server of claim 18, wherein the file includes a first class, and wherein the
5 modified file includes a second class that imports the first class.
33. The edge server of claim 18, further comprising transmitting a subset of the second portion of the file to the client prior to the client requesting the subset.
34. The edge server of claim 33, wherein the subset is in part based on how likely the subset is used by the application.
- 10 35. A computer program product for distributing content, the computer program product being embodied in a computer readable storage medium and comprising computer instructions for:
- receiving by an edge server an indication indicating that a client has sent a request for a file that includes content related to executing an application;
 - 15 analyzing interaction with the client by the edge server to determine a first portion of the file that is likely to be required by the client for executing the application sooner than a second portion of the file;
 - generating by the edge server a modified file that includes the first portion of the file;
 - and
 - 20 transmitting the modified file to the client.




FIG. 1

<u>Characters</u>
Sheep.jpg
Strawberry.jpg
Cat.mpeg
· · ·
<u>Symbol Table</u> actionscripts
<u>Timeline</u> new sheep() new cat() · · · screen.display(sheep, x, y)

FIG. 2

300

```
<html>
  <head>
    <title>Welcome</title>
  </head>
  <body>
    <h1>HelloWorld </h1>
    .
    .
    .
    <img src = "url for image"/>
    <video>
      <source src = "url for video" type = "video/ogg" />
    </video>
    <script type = "text/javascript">
      <!--script
      *** Some javascript code is placed here ***
      -->
    </script>
  </body>
</html>
```

Script

FIG. 3

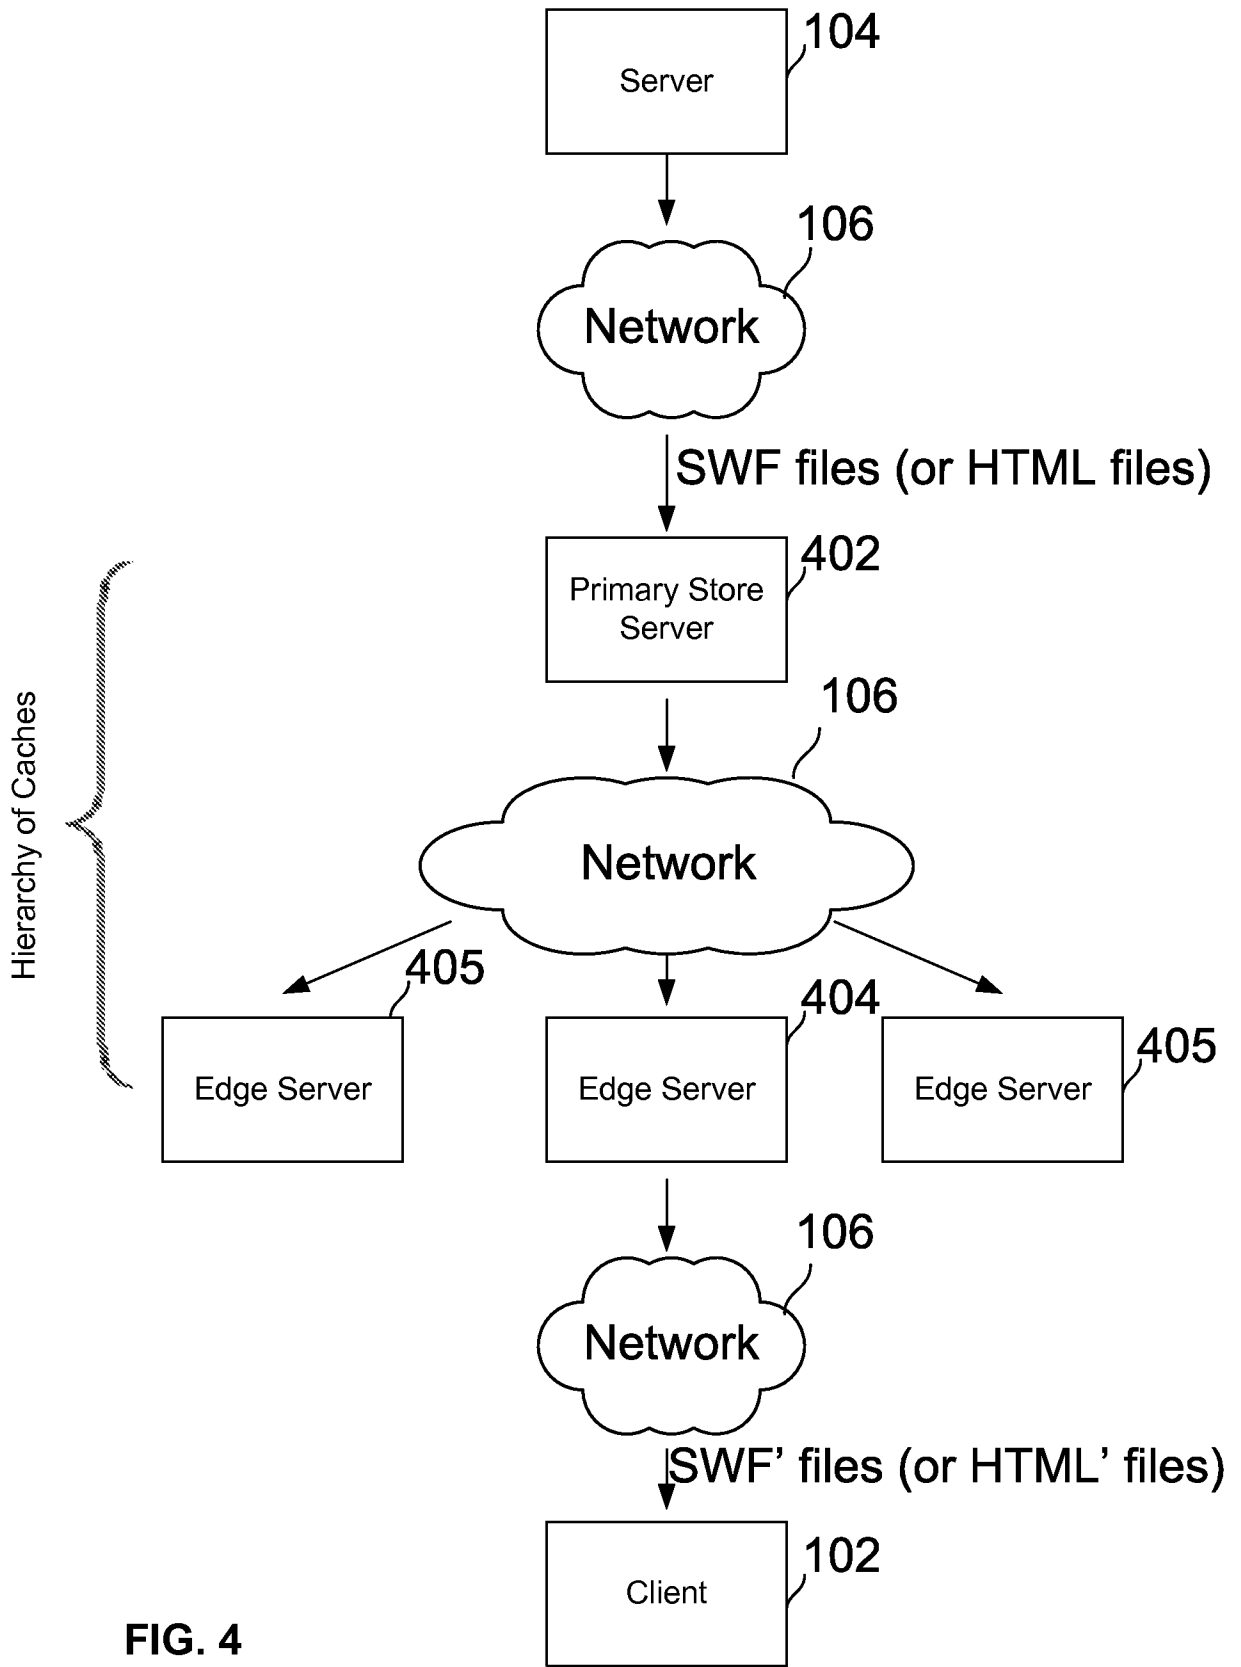
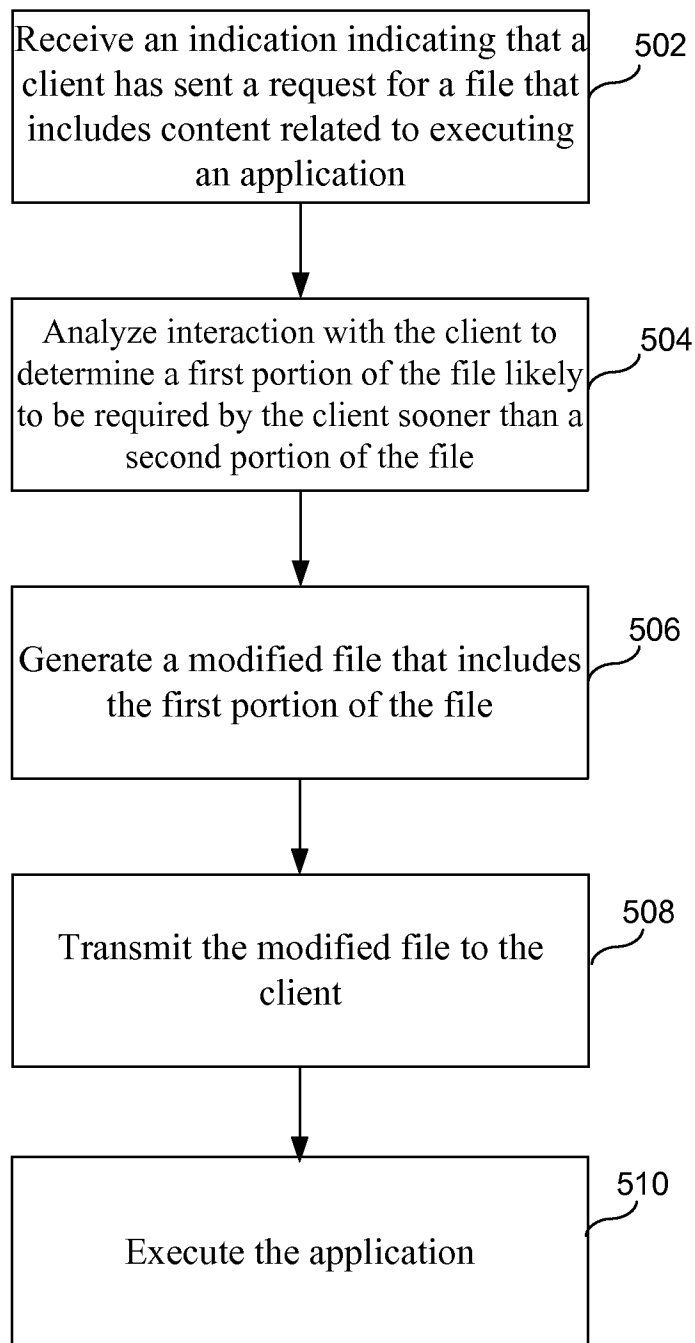
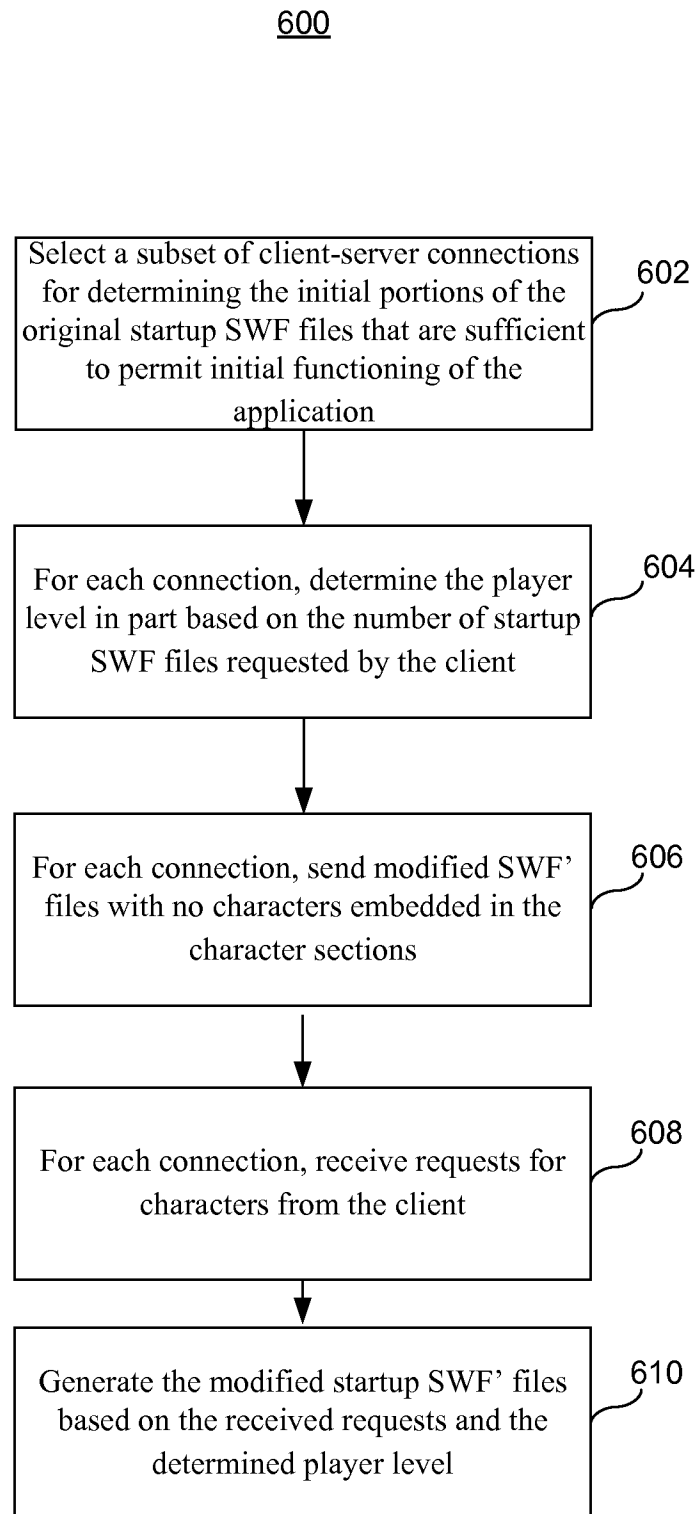


FIG. 4

500**FIG. 5**

**FIG. 6**

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 12/52162

A. CLASSIFICATION OF SUBJECT MATTER
 IPC(8) - G06F 15/16 (2012.01)
 USPC - 709/203
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 IPC(8): G06F 15/16 (2012.01)
 USPC: 709/203

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
 IPC(8): G06F 15/16 (2012.01)
 USPC: 709/201, 202, 203, 217 (keyword limited; terms below)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 PatBase(All); PubWEST(PGPB,USPT,EPAB,JPAB); Google (Scholar, Patents); FreePatentsOnline; Espacenet. Search terms used:
 edge-server distributed-content cache-content, probable likely anticipate most-likely anticipate-next-access anticipate-next-content, client
 -access client-file-access, file file-request database portion-file etc.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2002/0010798 A1 (BEN-SHAUL et al.) 24 January 2002 (24.01.2002) entire document, especially Abstract; para [0029]-[0031], [0037]-[0040], [0051], [0170], [0426], [0427], [0478]	1 - 35
Y	US 2004/0044731 A1 (CHEN et al.) 04 March 2004 (04.03.2004) entire document, especially Abstract; para [0008], [0056], [0061], [0064], [0069], [0090]-[0092], [0099]	1 - 35
Y	US 7,695,370 B2 (LIU et al.) 13 April 2010 (13.04.2010) entire document, especially Abstract; Fig. 10; col 2, ln 33-44; col 5, ln 25-41; col 15, ln 65 to col 16, ln 22; col 17, ln 4-12	10, 11, 13, 27, 28, 30
Y	US 2008/0301135 A1 (ALVES et al.) 04 December 2008 (04.12.2008) entire document, especially Abstract; para [0018], [0339], [0347]	15, 32
A	US 2005/0010653 A1 (MCCANNE) 13 January 2005 (13.01.2005) entire document	1 - 35
A	US 2002/0129123 A1 (JOHNSON et al.) 12 September 2002 (12.09.2002) entire document	1 - 35

Further documents are listed in the continuation of Box C.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier application or patent but published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 10 October 2012 (10.10.2012)	Date of mailing of the international search report 26 OCT 2012
---	--

Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774
---	--