

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
1 November 2007 (01.11.2007)

PCT

(10) International Publication Number  
**WO 2007/124144 A2**

(51) International Patent Classification:  
**G06F 7/00** (2006.01)

(74) Agent: **KOKKA, Scott, S.**; Kokka & Backus, PC, 200  
Page Mill Road, Suite 103, Palo Alto, CA 94306 (US).

(21) International Application Number:  
PCT/US2007/009816

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, 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, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(22) International Filing Date: 19 April 2007 (19.04.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
11/408,199 20 April 2006 (20.04.2006) US

(71) Applicant (*for all designated States except US*): **DATAS-COUT, INC.** [US/US]; 1125 South 103rd Street, Suite 550, Omaha, Nebraska 68124 (US).


(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(72) Inventor; and

(75) Inventor/Applicant (*for US only*): **KAMINSKI, Jr., Charles, F.** [US/US]; 15021 Hawthorne Circle, Omaha, Nebraska 68154-1822 (US).

[Continued on next page]

(54) Title: SURROGATE HASHING



(57) Abstract: Surrogate hashing is described, including running a hashing algorithm against a portion of a file to generate a hash value, determining whether the hash value is substantially similar to a stored hash value associated with another portion of another file, the portion and the another portion being standardized, and identifying a location of the another file if the hash value is substantially similar to the stored hash value associated with the another portion of the another file.



**Published:**

— without international search report and to be republished upon receipt of that report

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

## **SURROGATE HASHING**

### **FIELD OF THE INVENTION**

[0001] The present invention relates generally to software architecture. More specifically, surrogate hashing is described.

### **BACKGROUND OF THE INVENTION**

[0002] The Internet, World Wide Web, and other types of data networks may be used to find information. Specific information is typically sought using these sources by conducting a search. Searches are conducted for various reasons such as research, education, personal interest, rights management, and others. However, while a large amount of information is available from various sources and services on these networks, the approach used by search service providers and the amount of data (either raw or returned in searches) renders conventional search techniques problematic with regard to accuracy, efficiency, and latency.

[0003] Conventional search techniques are problematic because information is identified and found by analyzing text associated with a file. "File" may refer to a physical or logical grouping of data and as such, the file may or may not exist physically. Files may also refer to directory structures or data. A file can have text associated with it such as a reference on a web page (e.g., link, in-line image, and the like), metadata attached to the file, or another resource with text in proximity to or associated with the file reference. If a search is performed using keywords that correspond to the associated text of the file, then the file or file location is delivered as a search result.

[0004] This conventional approach is used when searching for files (such as an image file) on the Internet. The service provider's search engine has no knowledge of the contents of the file searched for. Instead, numerous results are returned based on text associated with the file intending to return files that accurately match a search request. However, the file is neither analyzed nor checked to ensure that it matches a user's desired search.

[0005] For example, if an intellectual property rights management organization (e.g., law firm, agency) is determining whether a particular image of a popular singer such as Madonna has been copied illegally, the organization may use a conventional search engine to search a network such as the Internet for the image in question. Conventional techniques typically associate the word "Madonna" with an image file. If text is found, automatic search solutions then attempt to analyze the text to determine whether the text indicates the image is similar to the image being sought. The analysis of text associated with a file (image or otherwise) is

neither accurate nor efficient. With each search result returned, a user must download the file in its entirety and manually evaluate the file. In the example cited, this approach forces the user to wade through thousands of pictures of other Madonnas such as the biblical Mary. When images of the pop singer Madonna are found, the image files often require additional manual review to determine which image files match a protected image of the popular singer. If a match is determined, then the image is identified as a copy and rights may be enforced. However, there may be additional copies of the protected image online, but if the indicated text is not found associated with the file, then a match can not be determined and rights may not be enforced.

[0006] In yet another example, a company may be trying to determine if its computer program is being distributed illegally on a network. Leveraging conventional solutions, the company would search based on text possibly associated with the computer program (e.g., "Get ABC's computer program here for free"). Once again, the files returned in the search are neither analyzed nor checked by the search engine to ensure that they match a user's desired search. There may be copies of the computer program that are never returned in the search results because the copies are not associated with text or because the associated text does not match the search request. For returned search results, manual review of a large amount of data is again required to determine if the files found in a search match those of the proprietary computer application.

[0007] Further, conventional solutions that identify files based on content are inefficient for all but comparatively small file sizes (e.g., HTML text, extremely small programs, pictures, or data files) because downloading larger files (e.g., picture files, music files, movie files, executables, and others) requires prohibitive amounts of bandwidth, data storage space, and processing power, which can be expensive and difficult to scale for implementation. Even if the required resources were obtained, the systems on the other side of the network providing the data would quickly become overloaded and may also exceed their allotted data transfer limits. Conventional solutions are also inefficient because analysis of the complete file is required, thus requiring large data storage facilities (e.g., data warehouses, arrays, and the like) and prohibitive amounts of processing power.

[0008] Conventional hashing algorithms or "hashing" techniques use an algorithm to generate a unique hash value for a file. However, this technique is problematic, as discussed above and because conventional solutions must first process an entire file to assign a hash value for the file. Subsequently, each file in the search results must have also been processed completely in order to generate a comparable hash value. If the hash value is the same, the files

are determined to match. However, using conventional techniques, the same hash value could be calculated for two different files (i.e., collisions may occur), leading to error-prone results. Other conventional hashing solutions require pre-processing of the entire data file, which requires large amounts of storage, processor capability, and bandwidth availability to perform the pre-processing, which is unduly burdensome, slow, and expensive. Conventional solutions are inefficient, inaccurate, labor and time-intensive, and expensive.

[0009] Thus, what is needed is for searching for data without the limitations of conventional techniques.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0010] Various embodiments of the invention are disclosed in the following detailed description and the accompanying drawings:

[0011] FIG. 1 illustrates an exemplary system for surrogate hashing, in accordance with an embodiment;

[0012] FIG. 2 illustrates an exemplary application architecture for surrogate hashing, in accordance with an embodiment;

[0013] FIG. 3 illustrates an exemplary overall process for surrogate hashing, in accordance with an embodiment;

[0014] FIG. 4A illustrates an exemplary overall process for surrogate hashing, in accordance with an embodiment;

[0015] FIG. 4B illustrates exemplary processing of a URL from a Local URL collection, in accordance with an embodiment;

[0016] FIG. 4C illustrates an exemplary process for parsing a URL, in accordance with an embodiment;

[0017] FIG. 4D illustrates an alternative exemplary overall process for surrogate hashing, in accordance with an embodiment; and

[0018] FIG. 5 illustrates an exemplary computer system suitable for surrogate hashing, in accordance with an embodiment.

### **DETAILED DESCRIPTION**

[0019] Various embodiments or examples may be implemented in numerous ways, including as a system, a process, an apparatus, or a series of program instructions on a computer readable medium such as a computer readable storage medium or a computer network where the program instructions are sent over optical, electronic, or wireless communication links. In

general, operations of disclosed processes may be performed in an arbitrary order, unless otherwise provided in the claims.

[0020] A detailed description of one or more examples is provided below along with accompanying figures. The detailed description is provided in connection with such examples, but is not limited to any particular example. The scope is limited only by the claims and numerous alternatives, modifications, and equivalents that are encompassed. Numerous specific details are set forth in the following description in order to provide a thorough understanding. These details are provided as examples and the described techniques may be practiced according to the claims without some or all of the accompanying details. For clarity, technical material that is known in the technical fields related to the embodiments has not been described in detail to avoid unnecessarily obscuring the description.

[0021] Surrogate hashing may be performed by evaluating a sampling or portion (“portion”) of a file’s data contents. In some embodiments, surrogate hashing may refer to the selection of a standardized portion of a file to determine whether, based on hash values, a selected file is similar to another file. Standardization may be performed systematically and repeatedly to ensure the same portion is taken the next time an identical file is encountered so that hashes are comparable. A portion may be selected from one or multiple parts of a file, including the beginning, middle, or end of a file, or a combination thereof. The data chosen to comprise a portion may be sequential or non-sequential. In some examples, other data outside of the file (e.g., application date, file metadata, and others) may be included in the portion. The data comprising the portion may also be modified before it is hashed. If a file is small (e.g., approximately 5 kilobytes or a comparably-sized file that has a substantially insignificant impact on supporting computing systems), a portion may also include the whole file. In some examples, surrogate hashing may refer to hashing a portion of a file to determine if another file has the same hash value or set of values. One or more hash values may be generated from a portion to determine whether a given file matches another file. A file may be a group of data for various types of computing systems, including binary, tertiary, quantum, textual, hexadecimal, octal, and others. The group of data may represent an image, photo, graphic, video, audio, computer program or application (“application”), text, or some other data structure. A file may refer to a physical or logical grouping of data and as such, the file may or may not exist physically. In some examples, a portion of a file may be analyzed to generate multiple (e.g., two (2) or more) hash values to identify a given file without the risk of collision. And in still other examples, multiple hash values may be concatenated together. More than one hash may be used

to minimize the risk of collisions (i.e., a different file having the same hash value) and to avoid mistakenly identifying a file. By analyzing a portion of a file instead of text or other information associated with a file, file identification may be performed quickly and accurately. Functions such as image searching, rights management, and others, may be performed without delay or omission errors (i.e., failing to return a match when a match should be indicated), and with few or no matching errors (i.e., mistakenly matching two different images). Surrogate hashing may be performed in various environments and is not limited to the use of Hosts, Uniform Resource Locators ("URLs"), crawlers, or the other exemplary environments described herein.

[0022] FIG. 1 illustrates an exemplary system for surrogate hashing, in accordance with an embodiment. Here, system 100 includes crawlers 102-106, network 108, content servers 110-118, and storage system 120. The number, type, configuration, and implementation of system 100 and the elements shown may be varied and are not limited to the examples given. In some examples, system 100 may be used to implement the described file identification techniques but may be varied in design, implementation, configuration, and other aspects and features.

Crawlers 102-106 may be implemented on computers and processors, including networked computing devices, notebook computers (i.e., laptops), mobile computing devices such as personal digital assistants, smart phones, or other wired or wireless computing devices. Content servers 110-118 may be implemented as application, web, or other types of servers that, when connected to a network, provide information at various locations and addresses (e.g., uniform resource locators (URLs)) accessible from network 108. Crawlers 102-106 may be configured to process domains or hosts ("hosts"), web pages, or other data files (collectively referred to as "files") located on content servers 110-118, which is described in greater detail below in connection with FIGs. 4A-4D. URLs may be addresses or indicators of a file location regardless of system, network, or application protocol. Links may be references to URLs and are not limited to the example used.

[0023] In some examples, crawlers 102-106 may be computer programs or applications ("applications") that are designed to search for content by processing files located at a given address and, in some examples, traversing links to other files at the given address according to various types of data processing techniques and structures (e.g., processing pages and links using a tree-structure, and others). Network 108 may be implemented as the Internet, a LAN, WAN, MAN, WLAN, or other type of data network over which data may be exchanged, transferred, downloaded, sent, received, and the like. The techniques described herein are not

limited to the type of data network from which files are retrieved or the protocols used to support those networks and may be varied without limitation to the example shown. Storage 120 may be implemented using one or more physical or logical data stores, databases, storage arrays (e.g., SAN), redundant arrays of independent disks (e.g., RAID), data warehouses, clustered storage systems, storage systems using volatile and/or non-volatile storage, storage networks, or other type of data storage formats or facilities and may be varied without limitation to the example shown. In some examples, a database management system may be used. In still other examples, relational database structures and languages may be implemented to enable files, portions of files, hashes, hash values, and other data relating to file searching, indexing, and management to be stored on storage 120. Further, techniques described herein may be implemented as software, hardware, circuitry, or a combination thereof. In some examples, software may be implemented using various programming, scripting, formatting, or other computer programming languages, including C, C++, Java, machine code, assembly, Fortran, XML, HTML, and others. The techniques described herein are not limited to any particular language or format and may be varied accordingly.

[0024] FIG. 2 illustrates an exemplary application architecture for surrogate hashing, in accordance with an embodiment. Here, application 200 may include logic module 202, input module 204, crawler interface (I/F) 206, hash module 208, and database system I/F 210. In some examples, application 200 may be implemented as software, hardware, circuitry, or a combination thereof. In some examples, software may be implemented using various programming, scripting, formatting, or other computer programming languages, including C, C++, Java, machine code, assembly, Fortran, XML, HTML, and others. Application 200 is not limited to any particular language or format and its design, architecture, implementation, and operation may be varied apart from the given description.

[0025] Here, logic module 202 may guide the operation of application 200, receiving user input via input module 204, sending/receiving data over crawler I/F 206 from crawlers 102-106 processing files found on content servers 110-118 (FIG. 1), running hashing algorithms to generate hash values for files identified, and storing/retrieving data from storage 120 (FIG. 1) using database system (DBS) I/F 210. Logic module 202 may also provide some, all or none of the applications, structure, or functionality of crawlers 102-106. As an example, a search may be initiated by providing a copy of the file desired to be found via input module 204. Once received, a portion of the file is hashed (i.e., hash algorithms are run against the data in the portion of the file) to generate one or more hash values. In some examples, more than one



hashing algorithm may be run in order to reduce collisions (i.e., different values having the same hash value or set of values). In other examples, multiple hash values are concatenated together to produce a stronger hash value. Once generated, the hash values are compared to those stored in storage 120. If the hash values generated for the file being sought match hash values of a file stored in storage 120, a location for the file associated with the hash values stored in memory is provided. Thus, other copies of a file (i.e., authorized, unauthorized, copyrighted, or otherwise protected or unprotected) may be found.

[0026] In some examples, hash values stored in storage 120 are generated from portions of files found by crawlers 102-106. Here, crawlers 102-106 are directed to a location (e.g., website, URL, or other type of file address) and begin processing and traversing directories, links, URLs, and files associated with the given location. In some examples, crawlers 102-106 (via crawler I/F 206) may continuously or non-continuously process and traverse directories, links, URLs, and files at various locations to continue to store hash values associated with files and locations (e.g., addresses, URLs, and the like) on storage 120. Files may be manually or automatically provided using various types of interfaces (e.g., graphical user interface (GUI), a system administration interface, command line interface (CLI), and others).

[0027] Here, a copy of the file to be sought is provided to logic module 202 using input module 204. Logic module 202 may be configured to run one or more hashes (i.e., hashing algorithms) to generate one or more hash values associated with the file. In some examples, two, three, or more hashes may be run instead of a single hash in order to minimize collisions (i.e., to avoid generating the same hash value for different files). In other words, to reduce the risk that files with different binary data found at different locations (i.e., on the Internet or another data networks) may have the same hash value, multiple hashing algorithms (i.e., hashes) may be run to generate a hash value that is individually assigned to a given file.

[0028] In some examples, if different files on different hosts have the same hash value, a new hash value may be generated using one or more hashing algorithms that individually identify the different files without conflict. Further, by generating individualized hash values associated with a given value, a file may be accurately matched to a copy of the file. For example, storage 120 may have 80 billion hashes and locations (e.g., URLs). If a file is sought, a hash value is generated for the file, which is then used for a search of storage 120 to determine whether the same hash is found. If a match of the hash value or set of values for the file is found, the location is returned, which identifies the location of the file associated with the hash values stored in storage 120.

[0029] FIG. 3 illustrates an exemplary process for surrogate hashing, in accordance with an embodiment. File identification may be performed using the below-described process, which may also be varied and is not limited to the description provided. Here, a file is received for a search (302). In some examples, a file may be submitted using a user interface (UI), command line interface, or other application for providing the file to application 200 (FIG. 2). Once a file is provided, a portion of the file is selected for analysis (304). In some examples, portions are “standardized,” which refers to identifying a consistent set, part, or sub-set of data that is selected from a file. Standardized portions may be identical in size and location (e.g., 128 bits of data selected from the first (i.e., “front end”) 128 bits of a file) or may be identical to other files. The use of standardized portions ensures that substantially similar portions or segments of data are selected for evaluation to help enhance finding a match. In other examples, “standardized” may be different and is not limited to the example given above.

[0030] Here, after a standardized portion of data has been selected, one or more hashing algorithms are run against the standardized portion to generate one or more hash values (306). If one hashing algorithm is run, a single hash value may be produced. However, if multiple hashing algorithms are run, then multiple hash values are produced, which may be used individually or in combination to identify a given file. In some examples, multiple hashing algorithms are run to minimize collisions. Here, minimizing collisions refers to the process of generating one or more hash values to individually identify a file without the risk of another, different file having the same set of hash values. After generating the one or more hash values, stored hash values are searched to determine whether a match exists (308). An example of developing hash values for storage and use in searches is described below in connection with FIGs. 4A-4F. In other examples, different techniques for finding, generating, and storing hash values may be implemented apart from those described in connection with FIGs. 4A-4F.

[0031] Referring back to FIG. 3, a search is performed to determine if the same hash value or set of hash values exist (310). If the same hash value or set of hash values are not found in storage 120, then the process ends. If the same hash value or set of hash values are found in storage 120, then the location for the file associated with the hash value or set of hash values is returned (312). In other examples, the above-described process may be varied and is not limited to the description given.

[0032] FIG. 4A illustrates an exemplary overall process for surrogate hashing, in accordance with an embodiment. Here, a crawler instance (i.e., an instantiation of a web crawler, bot, or substantially similar application) is registered with a storage facility, database, data warehouse,

or the like (402). Local variables are initialized, including hosts, Local URLs (i.e., URLs that link to other internal files of a host), and Foreign URLs (i.e., URLs that link to files on other hosts) collections (404). In some embodiments, initialization of local variables may include other variables and collections used to decide if a URL should be processed currently or stored (i.e., in storage 120) for later processing instead of processing Local URLs or Foreign URLs. In still other embodiments, initialization of local variables may include variables and collections which support URLs being processed currently or URLs being stored for later processing. Initialization may be performed to make collections of local variables (e.g., Local URLs, Foreign URLs, hosts) available to determine whether a URL is included in a collection. In other embodiments, initialization of local variables may be performed differently. After local variables are initialized, a host is retrieved, including associated local URLs (e.g., links that lead to other pages associated with the location, URL, or website), for processing (406). The retrieved URL is then processed (408). Processing a URL against a Local URLs collection is described in greater detail below in connection with FIG. 4B.

[0033] Referring back to FIG. 4A, once a URL has been processed, a determination is made as to whether another URL exists to be processed (410). If another URL is available for processing, then it is processed from the Local URLs collection (408). However, if no further URLs are detected for processing, then the local URLs are stored (in storage 120 (FIG. 1)) along with the hashed values associated with each local URL (412). Foreign URLs are also stored for future processing in storage 120 (414). The process then repeats with initializing local variables prior to retrieving another Host to process (404). In some embodiments, the above-described process may be performed repeatedly on some, none, or all URLs found by registered crawlers as directed. In other embodiments, the above-described process may be varied in design, implementation, execution, and is not limited to the example provided.

[0034] FIG. 4B illustrates exemplary processing of a URL from a Local URL collection, in accordance with an embodiment. Here, a file found at a given URL may be retrieved and hashed. In some examples, a determination is made as to whether a file indicates there are additional files that need to be downloaded (420). If no further files are available for download, then a determination is made to download a standardized (i.e., as described above) portion of a file to be hashed (422). However, if a file contains data indicating other additional files for download (i.e., html, directory listing, or other), then the remainder of the file is downloaded (424). URLs are parsed to capture additional file location data indicated in 420, as described in greater detail below in connection with FIG. 4C (FIG. 426).

[0035] Referring back to FIG. 4B, after parsing URLs from a file to identify additional file locations (426) or after downloading a standardized portion of a file (422), the file is hashed to calculate hash values (428). The calculated hash values are then stored locally with the given URL for later storage in storage 120 (430). In other examples, the above-described process may be varied and is not limited to the description provided above.

[0036] FIG. 4C illustrates an exemplary process for parsing a URL, in accordance with an embodiment. A more detailed process is provided for describing parsing URLs as mentioned above in connection with FIG. 4B. Here, a URL is parsed out to break up an address into constituent parts in order to standardize the URL into a standard address form that can be checked against a collection (440). Once parsed, the URL is standardized into a given format for an address that can be checked against a collection (442). A determination is made as to whether the URL is in an existing collection (444). If the URL is not found in an existing collection (e.g., Local URLs, Foreign URLs, and others), then a determination is made as to whether the URL is Local or Foreign (446). If the URL is a local URL (444), then it is added to a Local URLs collection (448). If the URL is a foreign URL, then it is added to a Foreign URLs collection (450). After adding the URL to either a Local or a Foreign URLs collection or if the URL is found in an existing collection (444), then a further determination is made as to whether there is another URL in the file (452). If another URL is found, then the process is repeated. If another URL is not found, then the process ends. In some embodiments, the decision to process a URL currently or at a later time may be based on information other than if the URL is Local or Foreign. In yet other embodiments, URLs may be processed currently or stored for later processing. Other data or collections may be used to support this decision. In other embodiments, the above-described process may be varied and is not limited to the example shown and described.

[0037] FIG. 4D illustrates an alternative exemplary overall process for surrogate hashing, in accordance with an embodiment. Here, a first portion of a first file is hashed to generate (i.e., calculate) a first hash value (460). The hash value is stored (e.g., in storage 120 (FIG. 1)) (462). A URL is received and processed (464), from which a second file is retrieved (466). A second portion of the second file is hashed to generate (i.e., calculate) a second hash value (468). The first hash value and the second hash value are compared to determine whether they are substantially similar (470). In some embodiments, determining whether the first hash value and the second hash value are substantially similar may include determining whether the first hash value and the second hash value are the exact same value. In other embodiments, determining

whether the first and the second hash value are substantially similar may include the first and second hash values being different, albeit slightly.

[0038] FIG. 5 illustrates an exemplary computer system suitable for surrogate hashing, in accordance with an embodiment. In some examples, computer system 500 may be used to implement computer programs, applications, methods, processes, or other software to perform the above-described techniques. Computer system 500 includes a bus 502 or other communication mechanism for communicating information, which interconnects subsystems and devices, such as processor 504, system memory 506 (e.g., RAM), storage device 508 (e.g., ROM), disk drive 510 (e.g., magnetic or optical), communication interface 512 (e.g., modem or Ethernet card), display 514 (e.g., CRT or LCD), input device 516 (e.g., keyboard), and cursor control 518 (e.g., mouse or trackball).

[0039] According to some examples, computer system 500 performs specific operations by processor 504 executing one or more sequences of one or more instructions stored in system memory 506. Such instructions may be read into system memory 506 from another computer readable medium, such as static storage device 508 or disk drive 510. In some examples, hard-wired circuitry may be used in place of or in combination with software instructions for implementation.

[0040] The term “computer readable medium” refers to any medium that participates in providing instructions to processor 504 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as disk drive 510. Volatile media includes dynamic memory, such as system memory 506. Transmission media includes coaxial cables, copper wire, and fiber optics, including wires that comprise bus 502. Transmission media can also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communications.

[0041] Common forms of computer readable media includes, for example, floppy disk, flexible disk, hard disk, magnetic tape, any other magnetic medium, CD-ROM, any other optical medium, punch cards, paper tape, any other physical medium with patterns of holes, RAM, PROM, EPROM, FLASH-EPROM, any other memory chip or cartridge, carrier wave, or any other medium from which a computer can read.

[0042] In some examples, execution of the sequences of instructions may be performed by a single computer system 500. According to some examples, two or more computer systems 500 coupled by communication link 520 (e.g., LAN, PSTN, or wireless network) may perform the

sequence of instructions in coordination with one another. Computer system 500 may transmit and receive messages, data, and instructions, including program (i.e., application code) through communication link 520 and communication interface 512. Received program code may be executed by processor 504 as it is received, and/or stored in disk drive 510, or other non-volatile storage for later execution.

[0043] The foregoing examples have been described in some detail for purposes of clarity of understanding, but are not limited to the details provided. There are many alternative ways and techniques for implementation. The disclosed examples are illustrative and not restrictive.

What is claimed:

1. A method for file identification, comprising:

running a first hashing algorithm against a first portion of a first file to generate a first hash value, and running a second hashing algorithm against the first portion of the first file to generate a second hash value;

determining whether the first hash value and the second hash value are substantially similar to one or more stored hash values associated with a second portion of a second file, the second portion of the second file being substantially similar to the first portion of the first file; and

identifying a location of the second file if the first hash value and the second hash value are substantially similar to the one or more stored hash values associated with the second portion of the second file.

2. The method of claim 1, wherein the location further comprises an address associated with the second file.

3. The method of claim 1, wherein the first file is an image

4. The method of claim 1, wherein the second file is an image.

5. The method of claim 1, wherein the first file comprises video data.

6. The method of claim 1, wherein the second file comprises video data.

7. The method of claim 1, wherein the first file is an application.

8. The method of claim 1, wherein the second file is an application.

9. The method of claim 1, wherein the first portion is the first file.

10. The method of claim 1, wherein the second portion is the second file.

11. The method of claim 1, wherein the first portion further comprises binary data.

12. The method of claim 1, wherein the second portion further comprises binary data.

13. The method of claim 1, wherein the one or more stored hash values are developed using a crawler.

14. The method of claim 1, wherein the one or more stored hash values are developed using crawlers.

15. The method of claim 1, wherein the first portion and the second portion are standardized.

16. A system for file identification, comprising:
- a database configured to store data associated with a first file and a second file; and
  - a processor configured to run a first hashing algorithm against a first portion of a first file to generate a first hash value, and running a second hashing algorithm against the first portion of the first file to generate a second hash value, to determine whether the first hash value and the second hash value are substantially similar to one or more stored hash values associated with a second portion of a second file, the second portion of the second file being substantially similar to the first portion of the first file, and to identify a location of the second file if the first hash value and the second hash value are substantially similar to the one or more stored hash values associated with the second portion of the second file.
17. A method for file identification, comprising:
- running a hashing algorithm against a portion of a file to generate a hash value;
  - determining whether the hash value is substantially similar to a stored hash value associated with another portion of another file, the portion and the another portion being standardized; and
  - identifying a location of the another file if the hash value is substantially similar to the stored hash value associated with the another portion of the another file.
18. The method recited in claim 17, further comprising minimizing collisions by running another hashing algorithm against the file.
19. The method recited in claim 17, wherein the another hashing algorithm is used to modify the hash value.
20. A computer program product embodied in a computer readable medium and comprising computer instructions for:
- running a first hashing algorithm against a first portion of a first file to generate a first hash value, and running a second hashing algorithm against the first portion of the first file to generate a second hash value;
  - determining whether the first hash value and the second hash value are substantially similar to one or more stored hash values associated with a second portion of a second file, the second portion of the second file being substantially similar to the first portion of the first file; and
  - identifying a location of the second file if the first hash value and the second hash value are substantially similar to the one or more stored hash values associated with the second portion of the second file.



21. A computer program product embodied in a computer readable medium and comprising computer instructions for:

running a hashing algorithm against a portion of a file to generate a hash value;

determining whether the hash value is substantially similar to a stored hash value associated with another portion of another file, the portion and the another portion being standardized; and

identifying a location of the another file if the hash value is substantially similar to the stored hash value associated with the another portion of the another file.

22. A method for file identification, comprising:

running a first hashing algorithm against a first portion of a first file to generate a first hash value, and running a second hashing algorithm against the first portion of the first file to generate a second hash value;

determining whether the first hash value and the second hash value are substantially similar to one or more stored hash values associated with a second portion of a second file, wherein the second portion is identified by one or more attributes that are substantially similar to one or more corresponding attributes associated with the first portion; and

identifying a location of the second file if the first hash value and the second hash value are substantially similar to the one or more stored hash values associated with the second portion of the second file.

23. The method of claim 22, wherein at least one of the one or more attributes is used to standardize the first portion and the second portion.

24. The method of claim 22, wherein at least one of the one or more attributes is used to identify a standardized region of the first file, the standardized region comprising the first portion.

25. The method of claim 22, wherein at least one of the one or more corresponding attributes is used to identify a standardized region of the second file, the standardized region comprising the second portion.

26. The method of claim 22, wherein at least one of the one or more corresponding attributes is used to standardize the first portion and the second portion.

27. The method of claim 22, wherein the first portion comprises the first file.

28. The method of claim 22, wherein the second portion comprises the second file.

29. A method for file identification, comprising:

selecting a standardized first portion of a first file, and a standardized second portion of a second file;

running a first hashing algorithm against the standardized first portion of the first file to generate a first hash value, and running a second hashing algorithm against the standardized first portion of the first file to generate a second hash value;

determining whether the first hash value and the second hash value are substantially similar to one or more stored hash values associated with the standardized second portion of the second file; and

identifying a location of the second file if the first hash value and the second hash value are substantially similar to the one or more stored hash values associated with the standardized second portion of the second file.

30. The method of claim 29, wherein selecting the standardized first portion of the first file, and the standardized second portion of the second file further comprises identifying a set of data that is substantially similar in both the standardized first portion and the standardized second portion.

31. The method of claim 29, wherein selecting the standardized first portion of the first file, and the standardized second portion of the second file further comprises identifying a location of the standardized first portion that is substantially similar to another location of the standardized second portion.

32. The method of claim 29, wherein the standardized first portion comprises the first file.

33. The method of claim 29, wherein the standardized second portion comprises the second file.

34. A computer program product embodied in a computer readable medium and comprising computer instructions for:

running a first hashing algorithm against a first portion of a first file to generate a first hash value, and running a second hashing algorithm against the first portion of the first file to generate a second hash value;

determining whether the first hash value and the second hash value are substantially similar to one or more stored hash values associated with a second portion of a second file, wherein the second portion is identified by one or more attributes that are substantially similar to one or more corresponding attributes associated with the first portion; and

identifying a location of the second file if the first hash value and the second hash value are substantially similar to the one or more stored hash values associated with the second portion of the second file.

35. A computer program product embodied in a computer readable medium and comprising computer instructions for:

selecting a standardized first portion of a first file, and a standardized second portion of a second file;

running a first hashing algorithm against the standardized first portion of the first file to generate a first hash value, and running a second hashing algorithm against the standardized first portion of the first file to generate a second hash value;

determining whether the first hash value and the second hash value are substantially similar to one or more stored hash values associated with the standardized second portion of the second file; and

identifying a location of the second file if the first hash value and the second hash value are substantially similar to the one or more stored hash values associated with the standardized second portion of the second file.

36. A system, comprising:

a database configured to store data associated with a first file and a second file; and

a processor configured to run a first hashing algorithm against a first portion of a first file to generate a first hash value, and running a second hashing algorithm against the first portion of the first file to generate a second hash value, to determine whether the first hash value and the second hash value are substantially similar to one or more stored hash values associated with a second portion of a second file, wherein the second portion is identified by one or more attributes that are substantially similar to one or more corresponding attributes associated with the first portion, and to identify a location of the second file if the first hash value and the second hash value are substantially similar to the one or more stored hash values associated with the second portion of the second file.

37. The system of claim 36, further comprising a hash module.

38. The system of claim 36, wherein the hash module is configured to generate the first hashing algorithm.

39. The system of claim 36, wherein the hash module is configured to generate the second hashing algorithm.

40. The system of claim 36, wherein the hash module is configured to generate the first hashing algorithm and the second hashing algorithm.
41. The system of claim 36, wherein the processor is further configured to select a standardized portion.
42. The system of claim 36, wherein the standardized portion is the first portion.
43. The system of claim 36, wherein the standardized portion is the second portion.
44. A system, comprising:  
a memory configured to store data associated with a standardized first portion of a first file and a standardized second portion of a second file; and  
logic configured to select a standardized first portion of a first file, and a standardized second portion of a second file, to run a first hashing algorithm against the standardized first portion of the first file to generate a first hash value, and running a second hashing algorithm against the standardized first portion of the first file to generate a second hash value, to determine whether the first hash value and the second hash value are substantially similar to one or more stored hash values associated with the standardized second portion of the second file, and to identify a location of the second file if the first hash value and the second hash value are substantially similar to the one or more stored hash values associated with the standardized second portion of the second file.
45. The system of claim 44, wherein the standardized first portion and the standardized second portion are selected from substantially similar sets of data associated with the first file and the second file.
46. The system of claim 44, wherein the standardized first portion and the standardized second portion are selected from substantially similar locations within the first file and the second file.
47. The system of claim 44, wherein the standardized first portion and the standardized second portion are substantially similar in size.
48. The system of claim 47, wherein the size is an amount of data.
49. The system of claim 44, wherein the standardized first portion comprises the first file.
50. The system of claim 44, wherein the standardized second portion comprises the second file.
51. A method, comprising:  
initializing one or more local variables;  
retrieving a host, the host being evaluated to determine if an address is associated with the host;

detecting an address, the address being further processed to download a file associated with the address if the file is associated with the host or storing the address if the address points to another host;

running a first hashing algorithm against a portion of the file to generate a first hash value, and running a second hashing algorithm against the portion of the file to generate a second hash value; and

storing the address associated with the file, the first hash value, and the second hash value.

52. The method of claim 51, further comprising retrieving the address if the address points to another host.

53. The method of claim 51, further comprising:

retrieving the address if the address points to another host; and  
processing the address to download a file from the another host.

54. The method of claim 51, further comprising:

retrieving the address if the address points to another host;  
processing the address to download another file from the another host;  
running a first hashing algorithm against another portion of the another file to generate a first hash value, and running a second hashing algorithm against the another portion of the another file to generate a second hash value; and  
storing the address associated with the another file, the first hash value, and the second hash value.

55. The method of claim 51, further comprising registering a crawler with a storage facility.

56. The method of claim 51, further comprising evaluating the host to determine if another file should be downloaded.

57. The method of claim 51, further comprising evaluating the host to identify one or more other files, the one or more other files being downloaded.

58. The method of claim 51, wherein the one or more other files being downloaded are used to generate a plurality of hash values by running a first hashing algorithm against a first portion of each of the one or more other files to generate a first hash value for each of the one or more other files, and running a second hashing algorithm against a second portion of each of the one or more other files to generate a second hash value for each of the one or more other files.

59. The method of claim 58, wherein first hash value and the second hash value for each of the one or more other files are stored in a storage facility.

60. The method of claim 51, wherein the portion is standardized.
61. The method of claim 51, wherein storing the address further comprises storing the address with the first hash value and the second hash value.
62. The method of claim 51, further comprising comparing the address to a collection of one or more other addresses.
63. A method for file identification, comprising:  
    initializing one or more variables in a collection;  
    evaluating an address associated with a host;  
    comparing the address to the collection to determine if the address is stored in the collection; and  
    processing the address to hash a file identified by the address if the address is not stored in the collection or determining if another address is indicated by the address if the address is stored in the collection.
64. The method of claim 63, wherein processing the address further comprises downloading a file associated with the address.
65. The method of claim 63, wherein processing the address further comprises storing the address if the address is associated with another host.
66. The method of claim 63, wherein processing the address further comprises running a first hashing algorithm against a portion of the file to generate a first hash value, and running a second hashing algorithm against the portion of the file to generate a second hash value.
67. The method of claim 66, further comprising storing the address, the first hash value, and the second hash value.

68. A computer program product embodied in a computer readable medium and comprising computer instructions for:

initializing one or more local variables;

retrieving a host, the host being evaluated to determine if an address is associated with the host;

detecting an address, the address being further processed to download a file associated with the address if the file is associated with the host or storing the address if the address points to another host;

running a first hashing algorithm against a portion of the file to generate a first hash value, and running a second hashing algorithm against the portion of the file to generate a second hash value; and

storing the address associated with the file, the first hash value, and the second hash value.

69. A computer program product embodied in a computer readable medium and comprising computer instructions for:

initializing one or more variables in a collection;

evaluating an address associated with a host;

comparing the address to the collection to determine if the address is stored in the collection; and

processing the address to hash a file identified by the address if the address is not stored in the collection or determining if another address is indicated by the address if the address is stored in the collection.

70. A system, comprising:

a memory configured to store data associated with one or more variables; and

a processor configured to retrieve a host, the host being evaluated to determine if an address is associated with the host, to detect an address, the address being further processed to download a file associated with the address if the file is associated with the host or storing the address if the address points to another host, to run a first hashing algorithm against a portion of the file to generate a first hash value, and running a second hashing algorithm against the portion of the file to generate a second hash value, and to store the address associated with the file, the first hash value, and the second hash value.

71. The system of claim 70, further comprising a crawler instance registered with the memory.

72. The system of claim 70, wherein the memory is a database.

73. The system of claim 70, wherein the processor is further configured to process the address if the address points to another host and to download a file from the another host.
74. The system of claim 70, wherein the processor is further configured to process the address if the address points to another host, to download another file from the another host, to run a first hashing algorithm against another portion of the another file to generate a first hash value, and running a second hashing algorithm against the another portion of the another file to generate a second hash value, and to store the address associated with the another file, the first hash value, and the second hash value.
75. The system of claim 70, wherein the processor is further configured to evaluate the host to determine if another file should be downloaded.
76. The system of claim 70, wherein the portion is standardized.
77. The system of claim 70, wherein the memory is configured to store the address further comprises storing the address with the first hash value and the second hash value.
78. The system of claim 70, wherein the processor is further configured to compare the address to a collection of one or more other addresses.
79. A system, comprising:  
a repository configured to store data associated with a host, an address, and one or more variables in a collection;  
a logic module configured to initialize one or more variables in a collection, to evaluate an address associated with a host, to compare the address to the collection to determine if the address is stored in the collection, and to process the address to hash a file identified by the address if the address is not stored in the collection or determining if another address is indicated by the address if the address is stored in the collection.
80. The system of claim 79, wherein the logic module is configured to download a file associated with the address.
81. The system of claim 79, wherein the logic module is configured to store the address if the address is not associated with the host.
82. The system of claim 79, wherein the logic module is configured to run a first hashing algorithm against a portion of the file to generate a first hash value, and to run a second hashing algorithm against the portion of the file to generate a second hash value.
83. The system of claim 79, wherein the memory is configured to store the address, the first hash value, and the second hash value.



84. A system, comprising:

one or more databases configured to store data associated with one or more addresses;

a crawler instance configured to crawl the one or more addresses, wherein the crawler is registered with the one or more databases; and

a distributed processor network configured to initialize one or more variables, to evaluate an address associated with a host, to compare the address to another address stored in a collection, and to process the address to hash a file identified by the address if the address is not stored in the collection.

85. A method, comprising:

parsing an address associated with a host;

generating a standardized address from the address;

determining if the standardized address is identified in a collection; and

processing the standardized address to retrieve a file, wherein a first hashing algorithm is run against a portion of the file to generate a first hash value, and a second hashing algorithm is run against the portion to generate a second hash value.

86. The method of claim 85, further comprising adding the standardized address to the collection if the standardized address is not identified in the collection.

87. The method of claim 85, wherein the collection comprises one or more local URLs.

88. The method of claim 85, wherein the collection comprises one or more foreign URLs.

89. The method of claim 85, wherein determining if the standardized address is identified in the collection further comprises determining if the standardized address is a local URL.

90. The method of claim 85, wherein determining if the standardized address is identified in the collection further comprises determining if the standardized address is a foreign URL.

91. The method of claim 85, wherein generating the standardized address further comprises identifying one or more parts of the address.

92. The method of claim 91, wherein the one or more parts of the address is a header.

93. The method of claim 91, wherein the one or more parts of the address is data associated with a host and a file located at the address.

94. The method of claim 85, wherein the portion is standardized.

95. A method, comprising:

parsing an address associated with a host to generate a standardized address in a format used to compare the standardized address to one or more addresses stored in a collection;

determining if the standardized address is listed in the collection; and

processing the standardized address.

96. The method of claim 95, wherein processing the standardized address further comprises adding the standardized address to the collection if the standardized address is not listed in the collection.

97. The method of claim 95, wherein processing the standardized address further comprises determining if another address is found at a location indicated by the standardized address.

98. The method of claim 95, wherein processing the standardized address further comprises downloading a file associated with the address.

99. The method of claim 95, wherein processing the address further comprises retrieving a file using the standardized address.

100. The method of claim 99, further comprising running a first hashing algorithm against a portion of the file to generate a first hash value, and running a second hashing algorithm against the portion to generate a second hash value.

101. A computer program product embodied in a computer readable medium and comprising computer instructions for:

- parsing an address associated with a host;

- generating a standardized address from the address;

- determining if the standardized address is identified in a collection; and

- processing the standardized address to retrieve a file, wherein a first hashing algorithm is run against a portion of the file to generate a first hash value, and a second hashing algorithm is run against the portion to generate a second hash value.

102. A computer program product embodied in a computer readable medium and comprising computer instructions for:

- parsing an address associated with a host to generate a standardized address in a format used to compare the standardized address to one or more addresses stored in a collection;

- determining if the standardized address is listed in the collection; and

- processing the standardized address.

103. A system, comprising:

- a memory configured to store data associated with an address; and

- a logic module configured to parse an address associated with a host, to generate a standardized address from the address, to determine if the standardized address is identified in a collection, and to process the standardized address to retrieve a file, wherein a first hashing

algorithm is run against a portion of the file to generate a first hash value, and a second hashing algorithm is run against the portion to generate a second hash value.

104. The system of claim 103, wherein the logic module is configured to add the standardized address to the collection if the standardized address is not identified in the collection.

105. The system of claim 103, wherein the collection comprises one or more local URLs.

106. The system of claim 103, wherein the collection comprises one or more foreign URLs.

107. The system of claim 103, wherein the logic module is configured to determine if the standardized address is a local URL.

108. The system of claim 103, wherein the logic module is configured to determine if the standardized address is a foreign URL.

109. The system of claim 103, wherein the logic module is configured to identify one or more parts of the address.

110. The system of claim 109, wherein the one or more parts of the address is a header.

111. The system of claim 109, wherein the one or more parts of the address is data associated with a host and a file located at the address.

112. The system of claim 103, wherein the portion is standardized.

113. A system, comprising:

a repository configured to store a standardized address and data associated with the standardized address; and

a processor configured to parse an address associated with a host to generate a standardized address in a format used to compare the standardized address to one or more addresses stored in a collection, to determine if the standardized address is listed in the collection, and to process the standardized address.

114. The system of claim 113, wherein the processor is configured to add the standardized address to the collection if the standardized address is not listed in the collection.

115. The system of claim 113, wherein the processor is configured to determine if another address is found at a location indicated by the standardized address.

116. The system of claim 113, wherein the processor is configured to download a file associated with the address.

117. The system of claim 113, wherein the processor is configured to retrieve a file using the standardized address.

118. The system of claim 117, wherein the processor is configured to run a first hashing algorithm against a portion of the file to generate a first hash value, and to run a second hashing algorithm against the portion to generate a second hash value.

1/8




FIG. 1

2/8

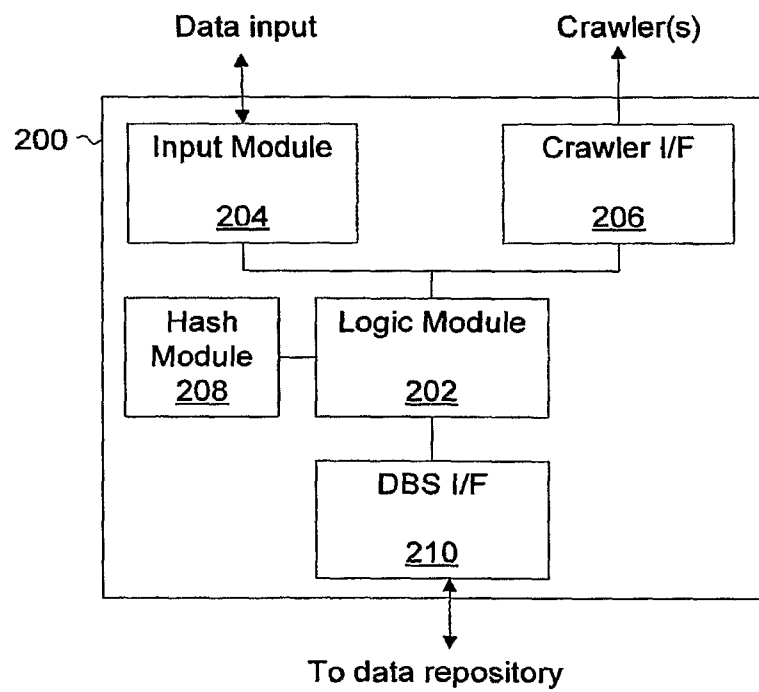


FIG. 2

3/8

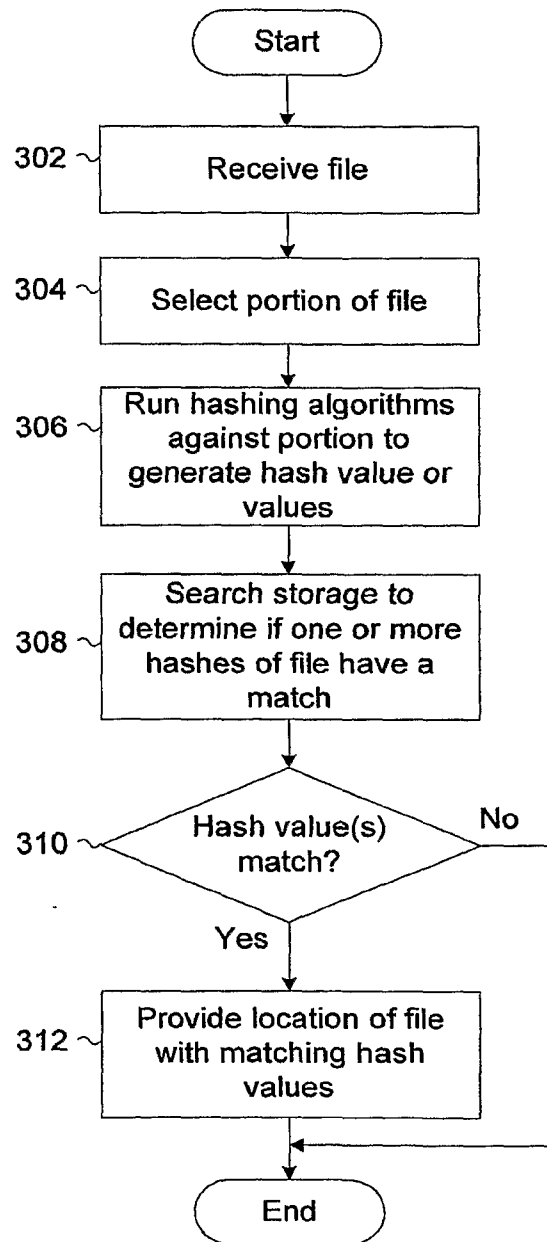


FIG. 3

4/8

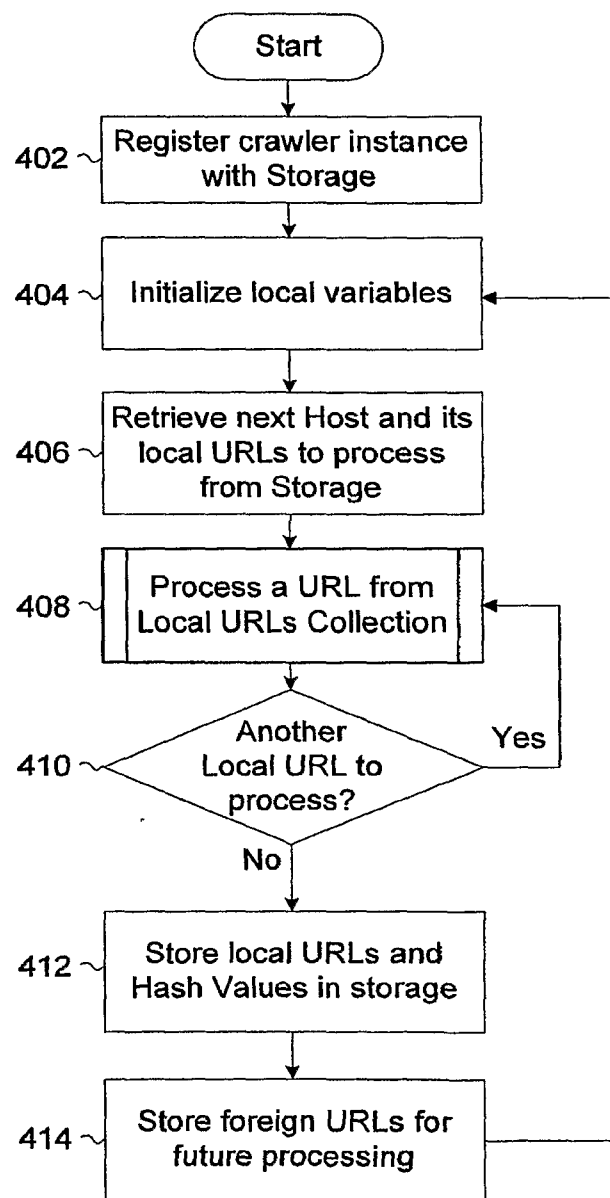


FIG. 4A



5/8

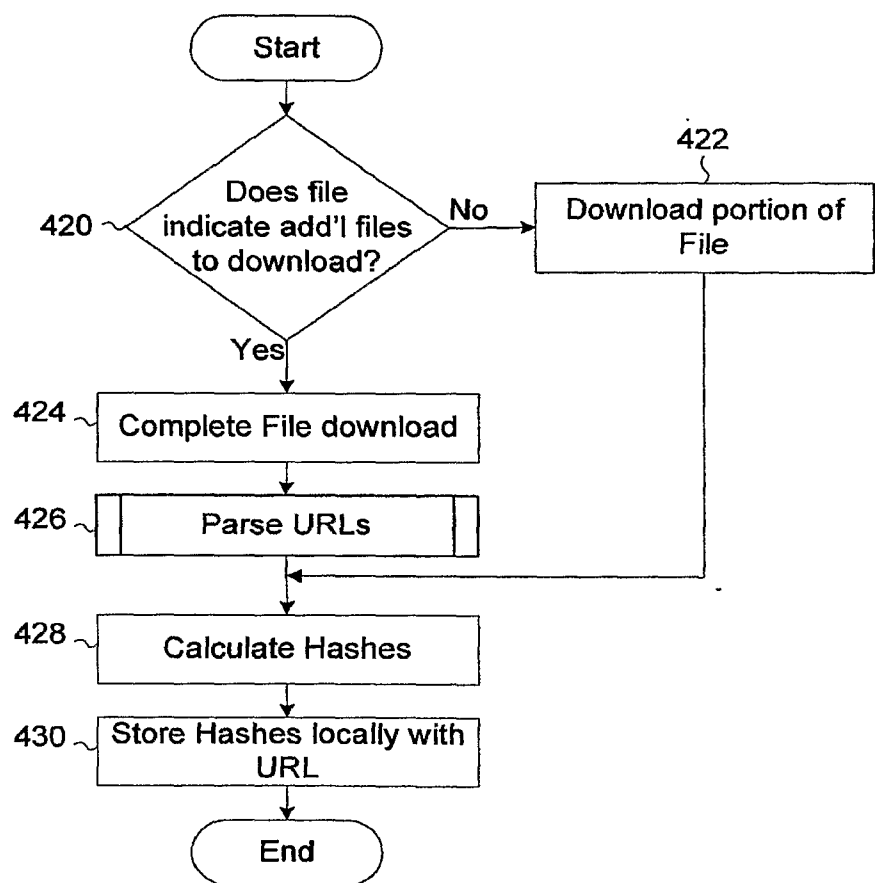


FIG. 4B

6/8

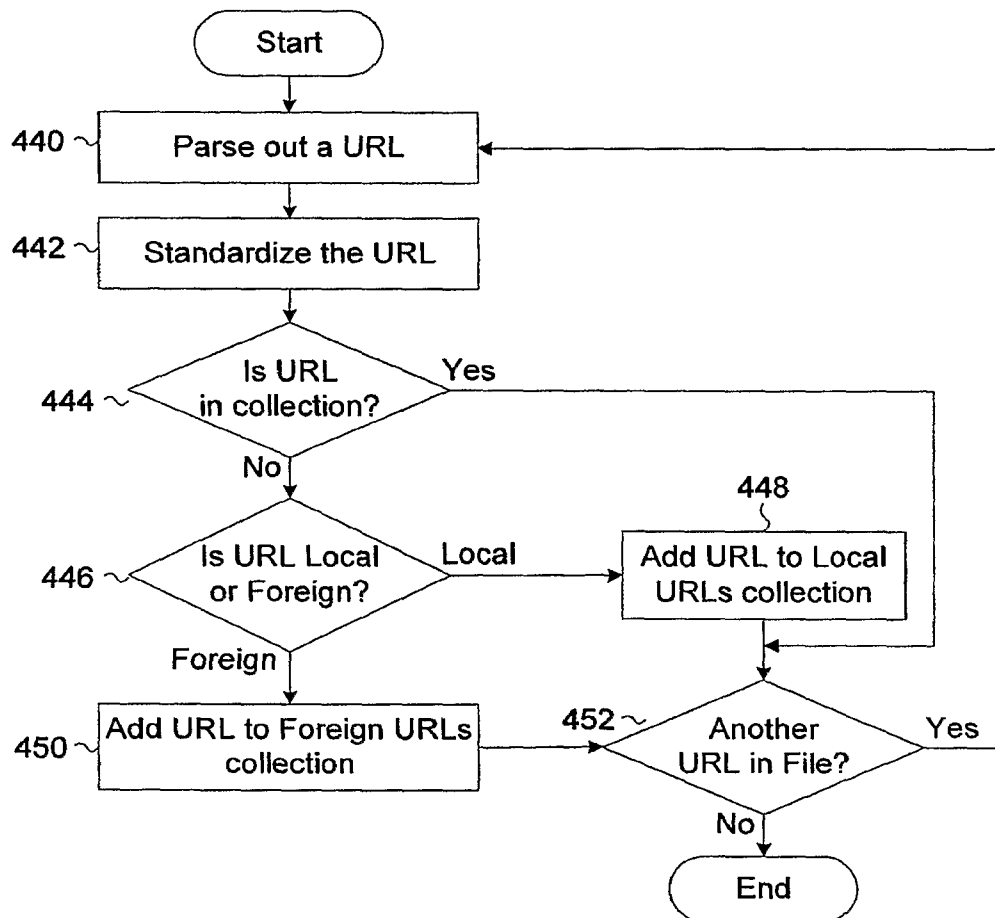


FIG. 4C

7/8




FIG. 4D

8/8

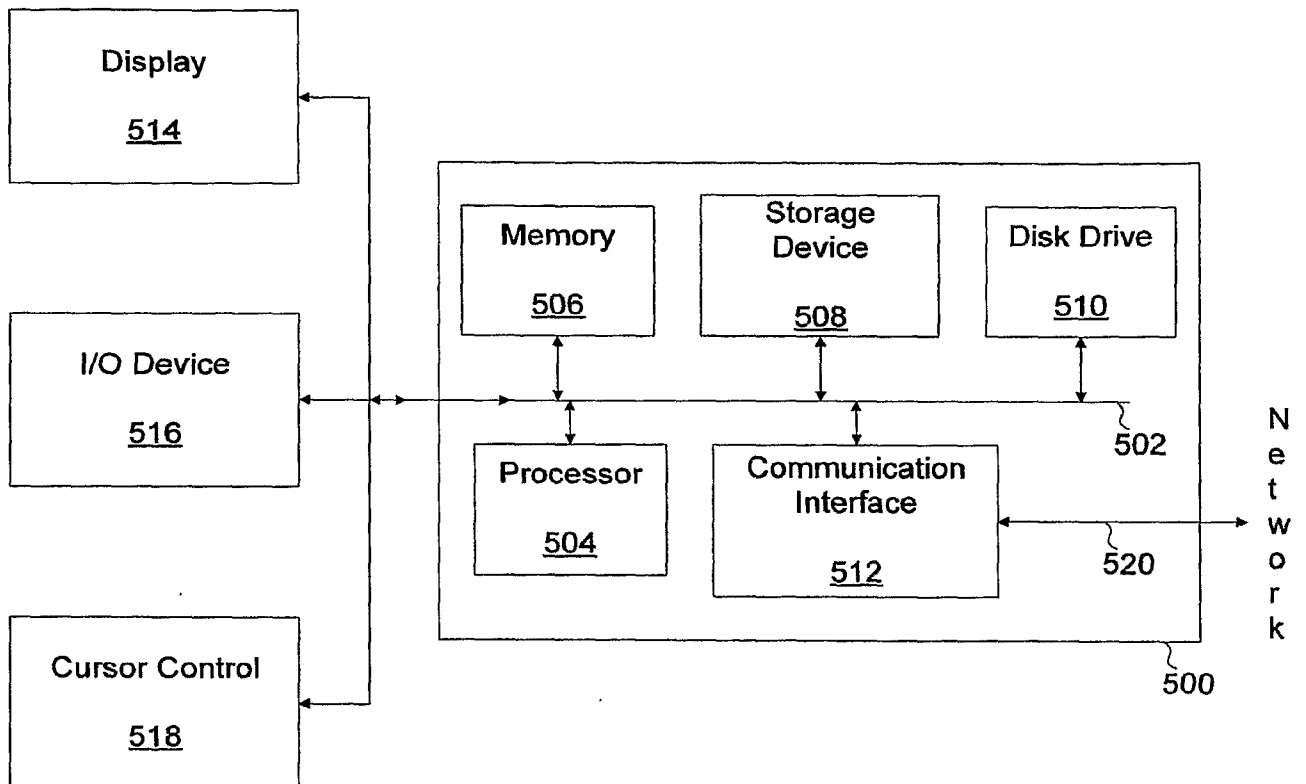


FIG. 5