

(12) **UK Patent Application** (19) **GB** (11) **2 366 965** (13) **A**

(43) Date of A Publication 20.03.2002

(21) Application No **0021371.0**

(22) Date of Filing **01.09.2000**

(71) Applicant(s)
NCR International Inc.
 (Incorporated in USA - Delaware)
1700 South Patterson Boulevard, Dayton,
Ohio 45479, United States of America

(72) Inventor(s)
Richard Cudd
Matthew Ashley Ward
Roderick James McMillan Macleod
Michael Waller

(74) Agent and/or Address for Service
NCR Limited
International IP Department, 206 Marylebone Road,
LONDON, NW1 6LY, United Kingdom

(51) INT CL⁷
G06F 17/30

(52) UK CL (Edition T)
H4P PPG

(56) Documents Cited
EP 0993163 A1 **EP 0877326 A2**
WO 98/22891 A1 **US 5864854 A**

(58) Field of Search
 UK CL (Edition S) **H4P PPG**
 INT CL⁷ **G06F 17/30 , H04L 29/06**
 Online: **WPI, EPODOC, JAPIO**

(54) Abstract Title
Downloading data to a requesting client from the local cache of another client

(57) A method of downloading or uploading data via a client-server communications network, which network comprises a plurality of clients 8, 9, 10 each having a local cache storing data downloaded via the network. The method comprises responding to a data request made to the network by a first client 8 by uploading data from the cache of a second client 9, 10 and transmitting that data across the network to the first client 8. The second clients 9, 10 act as proxy server clients. Preferably a proxy list is provided, listing which clients store particular web pages or other data. This proxy list may be held as a look up table in the requesting client 8, a server 3 or one or more of the proxy server clients 9, 10. The proxy list can be used to select from which client the requested data should be downloaded. The selection may take into account the speed of connection to an addresses in the proxy list, which may be assessed using ping signals.




FIG. 4

FIG. 1a
PRIOR ART

A CLIENT REQUEST TO A SERVER





FIG. 1b
PRIOR ART

CLIENT REQUEST VIA A PROXY SERVER



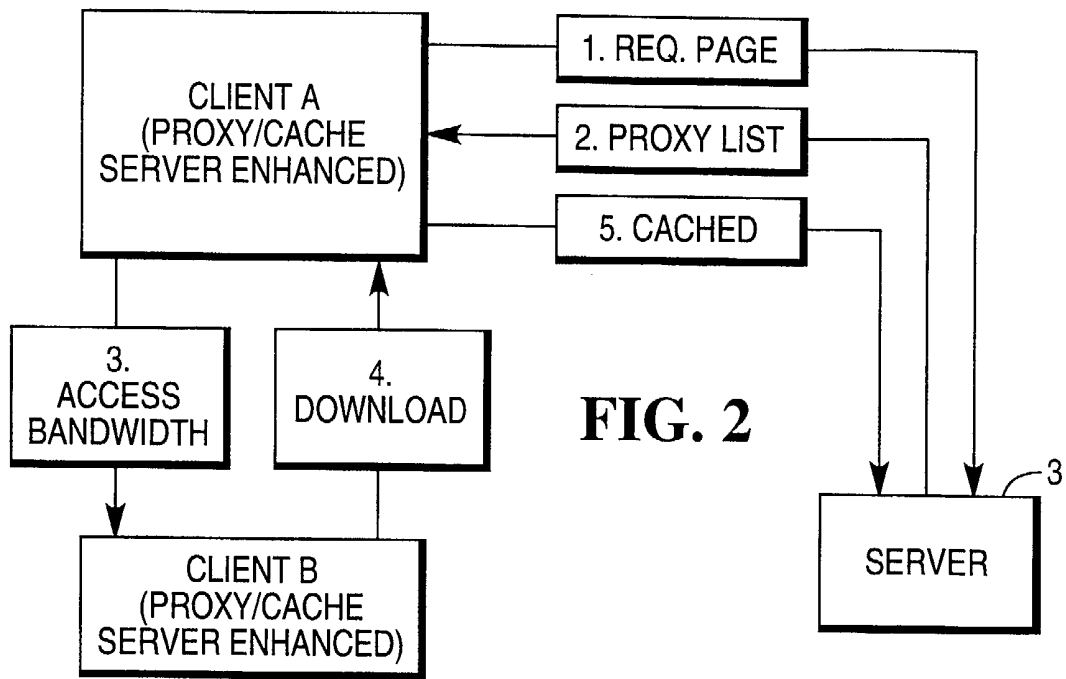


FIG. 2




FIG. 3




FIG. 4

DOWNLOADING AND UPLOADING DATA
IN INFORMATION NETWORKS

This invention relates to downloading data from or uploading data to information
5 sources via information networks and, in preferred embodiments, relates to techniques
for retrieving files such as web pages and other web content in an Internet
environment.

10 Currently, the Internet operates under the hypertext transfer protocol (HTTP) and
embodies a client-server architecture. The vast majority of Internet access - about
99% - is achieved via web browser programs, predominantly Netscape or Microsoft
Internet Explorer, whose trade marks are acknowledged.

15 Existing download techniques will be discussed later with reference to Figures 1(a)
and 1(b) but, typically, the client is a user's terminal such as a PC, a suitably-adapted
(e.g. Wireless Access Protocol or WAP) mobile telephone or other communications
device running a browser program. This terminal downloads and displays a desired
HTML web page held on a web server by using a communications network to send a
20 request for that web page across the Internet to the appropriate server. The server
responds by sending the requested web page back across the Internet and from there
to the client via the communications network to which the user's terminal is
connected.

25 Whilst a web page is mentioned by way of example, other web content files such as
.gif, .jpg or .mpg files can be downloaded in the same way.

The client and server can be in direct contact across the Internet via the
communications network or can be connected via a proxy server acting between the
client and the server. The purpose of the proxy server is to cache some web pages,
30 usually as a result of previous user requests, so that future user requests for the cached
web pages can be satisfied without connecting to the server. If the user requests a web
page that is not cached on the proxy server, the proxy server forwards the request to
the server and receives and forwards the requested page from the server to the client.

However, in general, less traffic needs to connect to the server and so the average download time is decreased.

Cache techniques are, of course, commonplace in the Internet art. Most commonly,
5 when a server or a proxy server responds to a user's request by sending a web page
back to the client, that page may be cached on the user's terminal so that future user
requests for the same web page can be satisfied immediately without having to
connect to the server or the proxy server at all. Nevertheless, the user's terminal
cannot cache every page that the user ever downloads, and the user will naturally wish
10 to update cached web pages and to download new web pages from time to time. This
means that efficient downloading remains paramount.

Despite ongoing efforts to speed Internet usage with faster modems and high-speed
network technologies such as ADSL and optical cable, the majority of Internet users
15 are burdened with slow download times. Even if an Internet user invests heavily in a
fast modem and in subscribing to a high-speed communications network, the user may
still suffer delays due to the architecture of the Internet itself and the nature of its
components. Particular problems arise due to the limited speed with which servers can
operate and the restricted bandwidth of the numerous communications channels that
20 lie between the server and the client. There is also the problem of unreliability,
meaning that if a server is down and no cached copy of the desired web page is
accessible elsewhere, the user may have to wait until the server is operational again.

The slowness and unreliability of downloads makes the Internet less useful and
25 appealing than it could and should be, to the detriment of users and also those who
seek to provide information to users. Recent research suggests that, on average, a user
will wait just eight seconds for a web page to download before moving on elsewhere.
If that happens, the user misses information that could have been beneficial and the
provider of the web page misses an opportunity to convey that information, possibly
30 resulting in lost business and decreased advertising revenues. The problem is likely to
get worse until efforts to upgrade the Internet and its associated communications
technologies begin to outweigh the explosion of new Internet users and the move
towards 'always-on' Internet access.

Broadly, this invention contemplates a method of downloading data via a client-server communications network, which network comprises a plurality of clients each having a local cache storing data downloaded via the network. The method comprises responding to a data request made to the network by a first client by uploading data
5 from the cache of a second client and transmitting that data across the network to the first client.

In use of a first architecture, the invention may be defined as a method of downloading or uploading data via a client-server communications network that
10 includes a server and a plurality of clients, each client having a local cache storing data downloaded via the network, the method comprising a requesting client (8) sending a request for data to the server, and the server responding by sending the requested data to the requesting client (8) or referring the requesting client (8) to a proxy server client that holds the requested data in its local cache, the requesting
15 client (8) then downloading the requested data from the cache of the proxy server client across the network.

This first architecture is embodied in a client-server communications network including a server and a plurality of clients, each client having a local cache storing
20 data downloaded via the network and the server having means for responding to a client that sends a data request to the server, wherein the server is adapted to send the requested data to the requesting client (8) or to refer the requesting client (8) to a proxy server client that holds the requested data in its local cache.

25 The server of the first architecture includes means for responding to a client that sends a data request to the server and is adapted to send the requested data to the requesting client or to refer the requesting client to a proxy server client that holds the requested data in a local cache of data downloaded via the network. The first architecture also involves a client terminal for connection to a server, or a browser for such a client
30 terminal, including selection means for choosing between a plurality of proxy server clients if the server is unable to respond to a data request from the client within a target period or at all, and means for downloading the requested data from a chosen proxy server client.

In use of a second architecture, the invention may be defined as a method of downloading or uploading data via a client-server communications network that includes a server and a plurality of clients, each client having a local cache storing data downloaded via the network, the method comprising a requesting client
5 broadcasting a data request over the network to the server and/or one or more other clients or connecting to at least one client whose address is on a proxy list held by the requesting client, the requesting client then downloading the requested data across the network from the cache of a proxy server client that is caching the requested data.

10 This second architecture is embodied in a client-server communications network including a server and a plurality of clients, each client having a local cache storing data downloaded via the network, wherein a requesting client is adapted to broadcast a data request over the network to the server and/or one or more other clients or to
15 connect to at least one client whose address is on a proxy list held by the requesting client, and includes means for downloading the requested data across the network from the cache of a proxy server client that is caching the requested data.

The second architecture involves a client terminal for connection to a server, or a browser for such a client terminal, including selection means for choosing among a
20 plurality of proxy server clients the proxy server client from which it will download the requested data, and means for downloading the requested data from a chosen proxy server client.

In an Internet context, the network is the Internet, the clients are user terminals
25 running web browsers and the respective local caches are associated with the browsers on the user terminals that act as proxy server clients. The invention capitalises upon the facts that (i) the vast majority of Internet access is done through web browsers that have become a *de facto* standard, such as Netscape and Microsoft Internet Explorer and (ii) those millions of web browsers cache a great deal of
30 information about the web sites that people most often look at.

A key advantage of the invention, in preferred embodiments, is that it can be deployed and propagated among a large number of users as a plug-in for their existing browsers. So, whilst a simple module would have to be added to the server, users

would not have to install completely new browsers but instead could upgrade their existing systems using simple download techniques with which most users are familiar. The invention therefore extends to software plug-ins for a client terminal or for a browser loaded on that client terminal and being programmed to adapt a terminal or a browser in accordance with the invention.

In order that this invention can be more readily understood, reference will now be made, by way of example, to the accompanying drawings in which:

10 Figure 1(a) and Figure 1(b) are block diagrams showing existing ways in which a web page may be downloaded via the Internet, Figure 1(a) showing a simple client-server architecture and Figure 1(b) showing a variant in which a proxy server acts between the client and the server;

15 Figure 2 is a block diagram showing how, in a preferred embodiment of the invention, a web page can be downloaded by one client from another client;

 Figure 3 is a block diagram showing how the bandwidths of connections between various clients can be compared;

20

 Figure 4 is a block diagram of a second embodiment of the invention that can be used in isolation or, preferably, as a fall-back for the preferred embodiment.

 In the simplest existing client-server architecture of Figure 1(a), a client 1 such as a PC running a browser makes an HTTP request 2 for a particular downloadable file to the server 3 and the server 3 responds by sending the requested file to the client 1. The file could, for example, be an HTML web page 4, a .gif, .jpg or .mpg file, or other web content. Web pages will be used as an example of such files throughout the description that follows.

30

 Figure 1(b) shows a variant in which a proxy server 5 acts between the client 1 and the server 3 to speed the average download time by reducing traffic connecting to the server 3. The client 1 makes an HTTP request for a particular web page via the proxy server 5. The proxy server 5 may have a cached copy of the requested web page as a

result of a previous request for that web page. If it does, then it returns that cached web page to the client 1, without connecting to the server 3. If not, the proxy server 5 requests and downloads the requested web page from the server 3, forwards it to the client 1 and advantageously keeps a copy in its cache in case a client 1 requests that page in the future.

In both of the above variants, the client 1 may cache a downloaded web page so that if a user requests the same page in future, it is not necessary to download that web page again from either a server 3 or a proxy server 5. However, the user may set the browser so that when a web page in cache is requested, the cached web page is compared to the corresponding web page then available from the server 3 or the proxy server 5. If there is a difference between the ostensibly corresponding web pages, the latest version of the web page can be downloaded, displayed and cached in place of the previous version.

Referring now to Figure 2 which illustrates a preferred embodiment of the invention, two clients, Client A and Client B, can access a server 3 via the Internet. There would of course be many more clients in practice, but just two clients are necessary to illustrate the broad inventive concept. Both Client A and Client B run browsers that have been enhanced in accordance with the invention, preferably by downloading and executing a suitable plug-in on the respective client terminals.

In the manner of the prior art, each respective terminal of Client A and Client B holds web pages in cache as a result of previous downloads. However, by virtue of the invention, the enhanced browsers open the caches of their respective client terminals for access by other network users. So, as the client terminals can act like proxy servers, the entire network can offer faster and more reliable downloads. In effect, the invention creates a network in which there are almost as many proxy servers as there are clients. These clients that emulate proxy servers will be referred to herein as proxy server clients.

In the embodiment of Figure 2, each proxy server client reports its cached web pages to the server 3 so that the server 3 can store in a look-up table the client location(s) at which a particular web page is cached. In use, a user at Client A requests a web page

from the server 3 with an HTTP request in the usual way. The server 3 either fulfils that request or, if it cannot fulfil the request quickly enough, looks up where else that web page is cached and responds to Client A with a short acknowledgement that points the browser of Client A to the appropriate proxy server client location. As will
5 be explained, the acknowledgement to the requesting client can be little more than a list of IP addresses constituting a proxy list of clients to identify the proxy server client terminal(s) at which the requested web page is cached.

If the server 3 tells Client A that the web page requested by Client A is cached at
10 Client B, then Client A contacts Client B and downloads that web page from the cache in the client terminal of Client B. Preferably, however, Client A firstly assesses and compares the bandwidths available at that time in the connections between itself and the server 3 on the one hand and between itself and Client B on the other hand. The aim is to determine which of the available connections would be the faster to use,
15 and then to select that connection so as to minimise download times and maximise the efficient use of network resources.

In Figure 2, the bandwidth available between Client A and Client B has been assessed as being greater than the bandwidth available between Client A and the server 3, so
20 Client A requests the desired web page from the cache of Client B. However, if the bandwidth comparison was instead in favour of the connection between Client A and the server 3, that connection would be used to download the web page from the server 3 instead of from Client B.

25 By extension, the technique of comparing bandwidths can be applied to the connections between Client A and proxy server clients other than Client B. So, if the desired web page is cached at other network resources such as other clients, the web page can be downloaded from one of those other resources if it would be more efficient, bandwidth-wise, to do so than to download the same web page from either
30 Client B or the server 3.

The principle of resource selection is shown in Figure 3, in which Client A has received a proxy list 6 of client IP addresses 7 from the server and then assesses the speed of the connection to each of the proxy server clients identified by the proxy list

6. This is done by a simple PING (Packet INternet Groper) operation that attempts to contact each specified IP address 7 and returns the times taken to connect to the terminals at those addresses 7. Specifically, a PING utility sends a packet to each IP address 7 and waits for, and times, the reply from each address 7. It will be noted that
5 Client A performs the PING operation rather than the server 3 because it is the connections between respective clients that matter in this context, not the connections between the server 3 and its clients.

The response times from each pinged proxy server client are recorded by Client A so
10 that, once an appropriate number of proxy server clients have been pinged, Client A can compare the recorded response times and select the proxy server client with the fastest response for the purpose of downloading the desired web page. It is also possible for a target response time to be set and for Client A to select from the first
15 pinged proxy server client to meet that target. This saves Client A continuing the process of pinging all of the proxy server clients on the proxy list when it has already found a proxy server client whose response time meets the target and so is deemed to be adequate.

It will be apparent that the pinging process set out in Figure 3 is preceded by an
20 assessment of whether it is faster to download the desired web page directly from the server 3 or to download from one of the proxy server clients that are caching the desired web page. This assessment can be done in various ways. In a first technique, for example, the server 3 may initially attempt to respond with the desired web page but if it cannot respond within a predetermined period deemed to be acceptably quick,
25 it instead responds with the aforementioned proxy list and leaves the requesting client to find and download the desired web page from another client. This has the advantage that if the server 3 can respond and upload the desired web page to the requesting client quickly enough, there is no need to go through the delay of pinging other clients at which the desired web page is cached.

30 The predetermined period in which the server 3 is challenged to respond need not be a fixed period of time: that period could change dynamically in accordance with the average download time for the requesting client terminal concerned. Clearly, all else being equal, a client terminal connected to the Internet via an ordinary 56k modem

will expect slower downloads than a terminal using a modem that can exploit a 128k ISDN or ADSL connection. In those circumstances, it is appropriate that the server 3 senses the speed of the connection from the requesting client and responds to the sensed speed by tailoring the predetermined period accordingly. Specifically, the server 3 should shorten the predetermined period when the connection is relatively fast, and should lengthen that period when the connection is relatively slow.

A second technique for assessing the speed of server response involves the server 3 invariably and immediately responding to web page requests with a proxy list 6 but including in that proxy list 6 the IP address of the server 3 itself. As a result, the server 3 is treated like the other resources (i.e. proxy server clients) identified by the proxy list 6 and so will be pinged along with the proxy server clients identified by that list. If the server 3 happens to be the fastest resource to respond, or if it is the first resource to respond within a predetermined target time, then the desired web page is downloaded from the server 3. Otherwise, the desired web page is downloaded from a proxy server client whose response is found to be fastest among the various resources identified by the proxy list 6 or whose response time is the first to meet the target.

This second assessment technique is currently less preferred than the first assessment technique because although it is more elegant in terms of architecture, it incurs the overhead of downloading the proxy list 6 from the server 3 and then pinging the IP addresses 7 on that list.

The various proxy server client locations, expressed as their respective IP addresses, can be stored on the server 3 as simple text files and a proxy list file is associated with every web page in the look-up table held on the server. The size of the proxy list file should obviously be kept within reasonable bounds, for example limited to a maximum of twenty IP addresses for each web page. This is due to considerations of memory capacity and download time but also has implications for efficient bandwidth assessment, in which it is desirable not to ping too many client terminals.

Another approach of the invention is shown in Figure 4, in which a client 8 making a web page request holds a proxy list of IP addresses defining a server 3 and also a group of other clients 9. The request is made to the IP addresses in the proxy list, so

being made to the server 3 in the normal way but also being broadcast to the group of other proxy server clients 9, to enquire as to whether they hold any of the requested information in cache. If they do, they can report back to the requesting client 8 and the web page can then be downloaded by that client 8 from any of the proxy server clients 5 9 that are caching the requested web page. Each of the client terminals on the proxy list may in turn be connected to other proxy server client terminals 10 to which they can forward the request, thus forming the chain or tree structure shown in simplified form in Figure 4.

10 In the Figure 4 approach, the server 3 is not relied upon to return a list of proxy terminals. The requesting client terminal merely needs to have the IP address of one, or the IP addresses of a few, of the proxy server client terminals 9 for the chain or tree to begin. The necessary IP address(es) could be downloaded from a web site or distributed with software.

15 It would also be possible to broadcast a short request over the network to make an initial connection with a proxy server client terminal 9 that has the requested web page in cache and responds to the broadcast. This operation would only need to be performed once since the IP address of the responding proxy server client 20 terminal could then be stored by the requesting client terminal for later use.

An advantage of the approach of Figure 4 is that it is possible to find web pages even when the server 3 is down. It could thus be a fall-back to the architecture of Figure 2, to be used only when there is no response from the original server in 25 operation of the Figure 2 embodiment.

If appropriate, the choice between potential proxy server client sources 9 of the cached web page can be made after the abovementioned bandwidth comparison between the various connections to those proxy server clients 9. As before, this 30 involves pinging all the proxy server client terminals 9 on the proxy list, including the server 3, and downloading via the fastest connection.

Again, a decision is required about whether to ping each IP address on the proxy list held by the requesting client 8, or simply to download the web page from the

server 3. If the Figure 4 architecture is used as a fall back, then a timeout can be set so that if no response has been received from the server 3 after a predetermined (but not necessarily fixed) period of time, then the chain or tree process is followed.

5

Another issue with pinging through a chain or tree structure is that each proxy server client terminal 9 at each level of the structure will ping to proxy server client terminals 10 in the next level of the structure. So, the originally-requesting client terminal 8 will not necessarily ping directly to the potential source of the requested web page if that source is more than one level down the structure. In those circumstances, it is necessary to add another step into the process to check the speed between the originally-requesting client terminal 8 and the potential source 10A, as shown in Figure 4, to be sure that there is a fast connection for downloading from the potential source 10A.

15

In all cases, a client advantageously reports to the server 3 upon downloading and caching a web page retrieved from cache among proxy server clients. In this way, the server can add that proxy server client location to its look-up table as a further potential source of that web page. It is similarly advantageous that a proxy server client reports to the server 3 removal of a web page from its cache, for example during a routine automated cache clean-up or in response to a user command. The server 3 can therefore delete that web page location from its look-up table and so knows to point requesting clients elsewhere if they request the deleted web page.

25 The invention requires extra messaging but it is expected that, in most practical cases, the overhead of that extra messaging in terms of download time would still allow shorter aggregate download times than can be achieved directly from the server.

The inventors recognise the need to ensure that no proxy server clients are overloaded with incoming requests, broadcasts and the resulting uploads, as this would unacceptably reduce the bandwidth available to those clients for other, unrelated communications. Accordingly, the invention contemplates means for monitoring the proxy workload of proxy server clients and preventing overload. This can be achieved at the client end by refusing to serve certain requests over a specified workload limit,

and/or at the server end by omitting an overloaded proxy server client from the proxy lists sent in response to requests from other clients. Also, in the embodiment of Figure 4, the invention contemplates limiting the number of proxy server clients in the group of clients that are polled and, possibly, changing the members of that group from request to request.

Many other variations are possible within the inventive concept. For example, the server in the Figure 2 embodiment can periodically update the look-up table from which the proxy lists are derived so to ensure that the lists are optimal. This can be achieved by pinging the IP addresses in the table from time to time, comparing their response times and discarding the slowest for a given item of data or those that fail to meet a target threshold. A possible problem with this approach is that the connection between a server and a proxy server client does not equate to the connection between one client and another client, so the server is not always best placed to assess client-client bandwidth. Another and possibly better approach is that when a client reports to the server upon downloading and caching a web page received from a resource such as a proxy server client, that client tells the server the IP address of the proxy server client that provided the web page. The server can therefore assemble a list of the most commonly accessed and hence fastest resources and can discard the less commonly accessed and hence slowest resources, like a voting scheme in which only those resources proven by various requesting clients to be fastest will continue to survive in the proxy lists held by the server.

The present invention may be embodied in other specific forms without departing from its essential attributes. Accordingly, reference should be made to the appended claims rather than to the foregoing specific description as indicating the scope of the invention.

CLAIMS

1. A method of downloading or uploading data via a client-server communications network that includes a server (3) and a plurality of clients (8, 9, 10), each client (8, 9, 10) having a local cache storing data downloaded via the network, the method comprising a requesting client (8) broadcasting a data request over the network to the server (3) and/or one or more other clients (9, 10) or connecting to at least one client whose address is on a proxy list held by the requesting client (8), the requesting client (8) then downloading the requested data across the network from the cache of a proxy server client (9, 10) that is caching the requested data.
2. The method of Claim 1, wherein the network is the Internet, the clients (8, 9, 10) are user terminals running web browsers and the respective local caches are associated with the browsers on the user terminals.
3. The method of Claim 1 or Claim 2, wherein the requesting client (8) chooses from a plurality of proxy server clients (9, 10) the proxy server client (9, 10) from which it will download the requested data.
4. The method of any preceding Claim, wherein the proxy list contains the respective addresses of the plurality of proxy server clients (9, 10).
5. The method of any preceding Claim, wherein the proxy list also contains the address of the server (3).
6. The method of any preceding Claim, wherein the requesting client (8) assesses the speed of at least one connection to an address in the proxy list.
7. The method of Claim 6, wherein the requesting client (8) assesses the speed of the connection by pinging the associated address taken from the proxy list.
8. The method of Claim 6 or Claim 7, wherein the requesting client (8) assesses the speeds of connections to more than one address in the proxy list, records and

compares the measured speeds, and downloads the requested data from the address with the fastest connection.

5 9. The method of any of Claims 6 to 8, wherein the requesting client (8) sets a target connection speed and downloads the requested data from the first address in the proxy list to meet that target.

10 10. The method of any preceding Claim, wherein the requesting client (8) maintains a look-up table correlating items of data with the addresses of proxy server clients (9, 10) that are caching that data.

11. The method of Claim 10, wherein proxy server clients (9, 10) report changes in their cache status to the requesting client (8).

15 12. The method of Claim 10 or Claim 11, wherein the requesting client (8) updates the look-up table by assessing connection speeds to the proxy server client addresses in the table and discarding or demoting the addresses of proxy server clients (9, 10) to which connections are relatively slow.

20 13. The method of Claim 12, wherein the addresses in the table are pinged and their response times are compared with each other or with a target threshold.

25 14. The method of Claim 16, wherein the requesting client (8) stores the address of the proxy server client (9, 10) that provided the requested data, and assembles an address list of the proxy server clients (9, 10) most commonly accessed to obtain the requested data.

30 15. The method of any preceding Claim, wherein the data request cascades through levels of a tree or chain client structure to a proxy server client (10A) whose address is not on the proxy list held by the requesting client (8).

16. The method of Claim 15, wherein the requesting client (8) checks that the speed of the connection to the proxy server client (10A) meets a target connection speed before downloading the requested data from that proxy server client (10A).

17. The method of Claim 16, wherein the requesting client (8) assesses the speed of the connection by pinging the proxy server client (10A) and comparing the measured connection speed with the target connection speed.

5

18. The method of any preceding Claim, wherein proxy server clients (9, 10) maintain a workload limit and do not serve requests that would exceed that workload limit.

19. The method of any preceding Claim, wherein the requesting client (8) monitors
10 the workload of proxy server clients (9, 10) and does not contact proxy server clients (9, 10) whose workload is above a workload limit.

20. A client-server communications network including a server (3) and a plurality of
15 clients (8, 9, 10), each client (8, 9, 10) having a local cache storing data downloaded via the network, wherein a requesting client (8) is adapted to broadcast a data request over the network to the server (3) and/or one or more other clients (9, 10) or to connect to at least one client (9) whose address is on a proxy list held by the requesting client (8), and includes means for downloading the requested data across the network from the cache of a proxy server client (9, 10) that is caching the
20 requested data.

21. The network of Claim 20, wherein the network is the Internet, the clients (8, 9, 10) are user terminals running web browsers and the respective local caches are associated with the browsers on the user terminals.

25

22. The network of Claim 20 or Claim 21, wherein the requesting client (8) includes selection means for choosing from a plurality of proxy server clients (9, 10) the proxy server client (9, 10) from which it will download the requested data.

30 23. The network of Claim 22, wherein the selection means selects from a proxy list containing the respective addresses of the plurality of proxy server clients (9) and optionally also containing the address of the server (3).

24. The network of Claim 23, wherein the selection means comprises bandwidth assessment means for measuring the speed of connection to an address in the proxy list.

5 25. The network of Claim 24, wherein the bandwidth assessment means includes means for pinging an address taken from the proxy list.

26. The network of Claim 24 or Claim 25, wherein the bandwidth assessment means includes means for recording speeds of connections to addresses on the proxy list,
10 means for comparing the recorded speeds, and means for downloading the requested data from the address with the fastest recorded speed.

27. The network of any of Claims 23 to 26, wherein the selection means includes means for comparing a measured connection speed to an address with a target
15 connection speed and downloading the requested data from that address if the measured connection speed meets the target connection speed.

28. The network of any of Claims 20 to 27, wherein the requesting client (8) stores a look-up table correlating items of data with the addresses of proxy server clients (9,
20 10) that are caching that data.

29. The network of Claim 28, wherein proxy server clients (9, 10) include means for reporting changes in their cache status to the requesting client (8).

25 30. The network of Claim 28 or Claim 29, wherein the requesting client (8) includes means for updating the look-up table, the requesting client (8) including means for assessing connection speeds to the proxy server client addresses in the table and means for discarding or demoting the addresses of proxy server clients (9, 10) to which connections are relatively slow.

30 31. The network of Claim 30, wherein the requesting client (8) includes means for pinging addresses in the table, means for recording their response times and means for comparing the recorded response times with each other or with a target threshold.

32. The network of any of Claims 20 to 31, wherein the requesting client (8) includes timeout means responsive to failure by the server (3) to respond within a predetermined target period, and means responsive to the timeout means to broadcast a data request over the network to at least one other client (9, 10) or to connect to at least one proxy server client (9) whose address is on an auxiliary proxy list held by the requesting client (8).

33. The network of any of Claims 20 to 32, wherein the proxy server clients (9, 10) are in a tree or chain structure comprising a plurality of levels.

10

34. The network of any of Claims 20 to 33, wherein a proxy server client (9, 10) includes workload limit means that prevents the proxy server client (9, 10) serving requests that would exceed a workload limit.

35. The network of any of Claims 20 to 34, wherein the server (3) includes workload limit means that prevents referral of the requesting client (8) to proxy server clients (9, 10) whose workload is above a workload limit.

36. A server (3) adapted for use in the client-server communications network of any of Claims 20 to 35.

37. The server (3) of Claim 36 and including workload limit means that prevents referral of the requesting client (8) to proxy server clients (9, 10) whose workload is above a workload limit.

25

38. A client terminal (8, 9, 10) for connection to a server (3), or a browser for such a client terminal, including selection means for choosing between a plurality of proxy server clients (9, 10) the proxy server client (9, 10) from which it will download the requested data, and means for downloading the requested data from a chosen proxy server client (9, 10).

30

39. The terminal (8, 9, 10) or browser of Claim 38, wherein the selection means selects from a proxy list containing the respective addresses of the plurality of proxy server clients (9, 10) and optionally also containing the address of the server (3).

40. The terminal (8, 9, 10) or browser of Claim 38 or Claim 39, wherein the selection means comprises bandwidth assessment means for measuring the speed of connection to an address in the proxy list.

5

41. The terminal (8, 9, 10) or browser of Claim 40, wherein the bandwidth assessment means includes means for pinging an address taken from the proxy list.

42. The terminal (8, 9, 10) or browser of Claim 40 or Claim 41, wherein the
10 bandwidth assessment means includes means for recording speeds of connections to addresses on the proxy list, means for comparing the recorded speeds, and means for downloading the requested data from the address with the fastest recorded speed.

43. The terminal (8, 9, 10) or browser of any of Claims 39 to 42, wherein the selection
15 means includes means for comparing a measured connection speed to an address with a target connection speed and downloading the requested data from that address if the measured connection speed meets the target connection speed.

44. The terminal (8, 9, 10) or browser of any of Claims 39 to 43, including means for
20 storing a look-up table correlating items of data with the addresses of proxy server clients (9, 10) that are caching that data.

45. The terminal (8, 9, 10) or browser of Claim 44, including means for reporting
25 changes in their cache status to a requesting client (8).

25

46. The terminal (8, 9, 10) or browser of Claim 44 or Claim 45, including means for updating the look-up table, means for assessing connection speeds to the proxy server client addresses in the table and means for discarding or demoting the addresses of proxy server clients (9, 10) to which connections are relatively slow.

30

47. The terminal (8, 9, 10) or browser of Claim 46, including means for pinging addresses in the table, means for recording their response times and means for comparing the recorded response times with each other or with a target threshold.

48. The terminal (8, 9, 10) or browser of any of Claims 39 to 47, including timeout means responsive to failure by the server (3) to respond within a predetermined target period, and means responsive to the timeout means to broadcast a data request over the network to at least one other client (9, 10) or to connect to at least one proxy server client (9) whose address is on an auxiliary proxy list held by the requesting client (8).

49. The terminal (8, 9, 10) or browser of any of Claims 39 to 48, including workload limit means that prevents the terminal serving requests that would exceed a workload limit.

50. A software plug-in for a client terminal (8, 9, 10) or for a browser loaded on that client terminal and being programmed to adapt the terminal (8, 9, 10) or the browser in accordance with any of Claims 39 to 49.



INVESTOR IN PEOPLE

Application No: GB 0021371.0
Claims searched: 1-37

20

Examiner: John Cullen
Date of search: 26 April 2001

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
UK Cl (Ed.S): H4P (PPG)
Int Cl (Ed.7): G06F 17/30; H04L 29/06
Other: ONLINE:WPI, EPODOC, JAPIO - via EPOQUE

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	EP 0993163 A1 (BACKWEB) Note in particular paras. 7, 8, 37 and 73 but also see Abstract, Figs 1A, 1B and 4 and paras. 68-77.	1, 20 at least
X	EP 0877326 A2 (AT & T) See whole document, but particularly line 54 of col. 6 to line 16 of col. 7, lines 34-37 of col. 7, lines 21-26 of col. 8 and Fig. 1. Note directory 160 acts as proxy list held in a server.	1, 2, 4, 21, 36, 20, 21, 23
A	WO 98/22891 A1 (MANGOSOFT) See Figs. 11 and 12 and pages 38-42.	---
X	US 5864854 (LSI LOGIC) See whole document, but particularly lines 32-40 of col. 3, lines 34-57 of col. 4, lines 16-18 of col. 5, line 55-57 of col. 5, lines 33-35 of col. 7 and all Figs.	1, 2-4, 6, 8, 10, 11, 15, 20-24, 26, 28, 29, 32, 33, 36

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.