(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
25 March 2010 (25.03.2010)

PCT

(10) International Publication Number
WO 2010/033620 A1

(51) International Patent Classification:
G06F 15/16 (2006.01)

(21) International Application Number:
PCT/US2009/057192

(22) International Filing Date:
16 September 2009 (16.09.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/235,409 22 September 2008 (22.09.2008) US

(71) Applicant *(for all designated States except US)*: **SONY COMPUTER ENTERTAINMENT AMERICA INC.** [US/US]; 919 East Hillsdale Boulevard, 2nd Floor, Foster City, California 94404 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **ROY, Ronald, J.** [US/US]; 919 East Hillsdale Boulevard, 2nd Floor, Foster City, California 94404 (US). **JACOB, Mark, L.** [US/US]; 9781 Paseo Montril, San Diego, California 92129 (US).

(74) Agent: **ISENBERG, Joshua, D.**; JDI Patent, 809 Corporate Way, Fremont, California 94539 (US).

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

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, 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, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

(54) **Title**: METHOD FOR HOST SELECTION BASED ON DISCOVERED NAT TYPE

FIG. 1

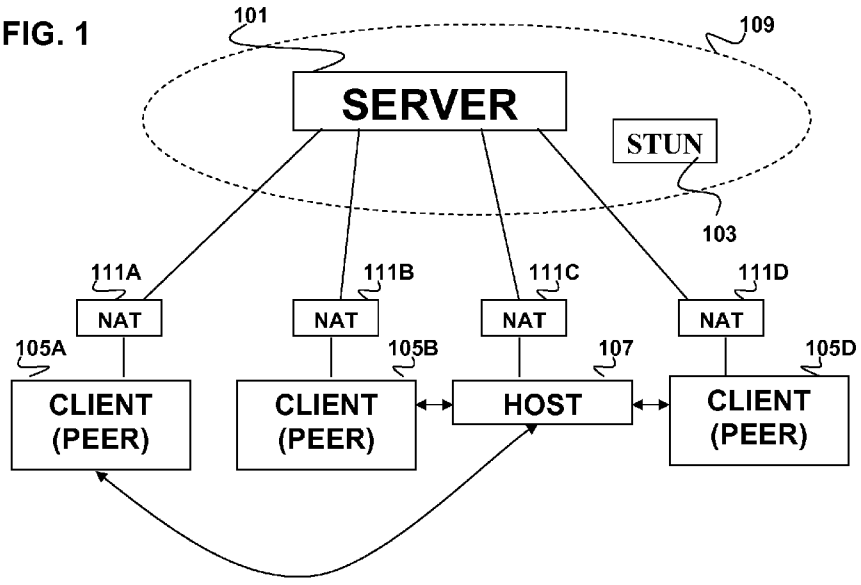

(57) **Abstract**: A host may be selected for a peer-to-peer grid based on discovered NAT type. NAT profile information may be collected by each of the peers and shared with the other peers. Each peer receives NAT profile information for the other peers. Each peer may determine from the NAT profile information for the peer and the NAT profile information for the other peers which of the two or more peers to designate as the host.

WO 2010/033620 A1

METHOD FOR HOST SELECTION BASED ON DISCOVERED NAT TYPE

PRIORITY CLAIM

This application claims the benefit of priority of commonly assigned U.S. Patent Application 12/235,409, filed September 22, 2008, the entire disclosures of which are incorporated herein by reference.

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is related to commonly-assigned US Patent Application Number 12/235,438 to Mark Lester Jacob et al. entitled "SEAMLESS HOST MIGRATION BASED ON NAT TYPE" (Attorney Docket Number PA4712US) , filed September 22, 2008, the entire contents of which are incorporated herein by reference.

This application is related to U.S. Patent Application Number 12/049,954 to Mark Lester Jacob et al. entitled "SYSTEM AND METHODS FOR SEAMLESS HOST MIGRATION", filed March 17, 2008, the entire contents of which are incorporated herein by reference. This application is also related to U.S. Provisional Patent Application Number 60/997,918 filed October 5, 2007 and entitled "Systems and Methods for Seamless Host Migration", the disclosure of which is incorporated herein by reference.

This application is related to commonly-assigned US Patent Application Publication Number 20070076729 to Yutaka Takeda et al. entitled "PEER-TO-PEER COMMUNICATION TRAVERSING SYMMETRIC NETWORK ADDRESS TRANSLATORS", filed October 4, 2005, the entire contents of which are incorporated herein by reference.

FIELD OF THE INVENTION

This invention is related to computer networks and more specifically to determination of a host for peer-to-peer communication among several clients of a computer network.

BACKGROUND OF THE INVENTION

Presently, during game play amongst several clients of a computer network, clients communicate directly with the server. The central server processes data from every client relaying this data to all other clients so that several clients in the network can engage in game play with each other through a central server. The central server's ability to communicate data amongst several clients is limited by bandwidth, and so communication outside of game play data is limited.

Aside from game play data; clients may want to communicate other information to each other without having to go through a bandwidth-limited central server. Examples of such communication may include voice over Internet protocol (VoIP), bit torrents, video data, file sharing, and data streaming. By allowing clients to communicate amongst each other without the need of a central server, some of the central server's responsibilities may be offloaded to the clients.

A peer-to-peer network is established when a host is determined amongst the numerous clients participating in the P2P network. The host takes the duty of a central server and directs communication amongst clients (hereinafter referred to as peers when referencing non-hosts in a P2P network).

It is within this context the embodiments of the present invention arise.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic diagram illustrating a P2P network according to an embodiment of the present invention.

FIG. 2 is a flow diagram illustrating a method for host selection based on discovered NAT type according to a specific embodiment of the present invention from the client's perspective.

FIG.3 is a flow diagram illustrating the basic method for host selection based on discovered NAT type from the client's perspective.

FIG. 4 is a block diagram illustrating an example of a client interface according to an embodiment of the current invention.

DESCRIPTION OF THE SPECIFIC EMBODIMENTS

INTRODUCTION

One problem that arises when trying to establish communication amongst clients through peer-to-peer communication (P2P) is the issue of network address translation (NAT). Most clients connected to the central server are situated behind a NAT. NAT is an Internet standard that enables a local area network (LAN) to use of one set of private IP addresses for internal traffic and a second set of global IP addresses for external traffic. It is therefore desirable for a host in a P2P network to have a favorable NAT profile in order to create an optimal P2P network.

Thus, there is need in the art, for a method of determining a host with a favorable NAT profile amongst several clients connected to a central server.

There are 4 types of NATs available: Full Cone NAT, Restricted Cone NAT, Port Restricted Cone NAT, and Symmetric NAT. A full cone NAT is one where all requests from the same internal IP address and port are mapped to the same external IP address and port. Furthermore, any external host can send a packet to the internal host by sending a packet to the mapped external address.

In a restricted cone NAT, all requests from the same internal IP address and port are mapped to the same external IP address and port. Unlike a full cone NAT, an external host (with IP address X) can send a packet to the internal host only if the internal host had previously sent a packet to IP address X.

A port restricted NAT is like a restricted cone NAT, but the restriction includes port numbers. Specifically, an external host can send a packet with source IP address X and source port P to the internal host only if the internal host had previously sent a packet from IP address X and port P.

In a symmetric NAT, all requests from the same internal IP address and port, to a specific destination IP address and port, are mapped to the same external IP address and port. If the same host sends a packet with the same source address and port, but to a different destination, a different mapping is used. Furthermore, only the external host that receives a packet can send a user data protocol (UDP) packet back to the internal host.

HOST SELECTION BASED ON NAT TYPE

FIG. 1 is a schematic diagram illustrating a P2P network according to an embodiment of the present invention. Client devices **105A, 105B, 107**, and **105D** on the network are initially connected to a server **101** in an external network **109**. By way of example, the server **101** may monitor data for game play between clients **105A, 105B, 107**, and **105D** connected to the external network **109**. The clients **105A, 105B, 107**, and **105D** may be situated behind corresponding network address translators (NAT) **111A**, **111B**, **111C** and **111D**. Each NAT may be configured according to an Internet standard that enables a local area network (LAN) to use one set of private IP addresses for internal traffic and a second set of global IP addresses for external traffic. Most NATs perform dynamic translation on IP addresses such that external networks have no way of reaching the internal network prior to the internal network initiating communication. However, because the clients **105A, 105B, 107**, and **105D** establish server-client relationships, communication between the server **101** and clients **105A, 105B, 107**, and **105D** is not limited by the NATs **111A**, **111B**, **111C** and **111D**.

In establishing a P2P network in which clients can communicate directly with each other, one client may be established as a host **107** by which other peers **105** (clients connected to the P2P network who are not the host) may establish direct connections with each other. By way of example and not by way of limitation, the peers **105** may be connected in a configuration known as a fully connected grid (FCG). Such a configuration prevents any one peer from being a bottle neck. In embodiments of the present invention, the host **107** may be determined based on each client's NAT profile. The NATs **111A**, **111B**, **111C** and **111D** may be of one of four different configurations generally referred to as Full Cone NAT, Restricted Cone NAT, Port Restricted Cone NAT, and Symmetric NAT.

A full cone NAT is one where all requests from the same internal IP address and port are mapped to the same external IP address and port. Any client can send a packet to the client behind a full cone NAT by sending a packet to the mapped external address for the client.

In a restricted cone NAT, all requests from the same internal IP address and port are mapped to the same external IP address and port. However, unlike a full cone NAT, an external client (with IP address X) can send a packet to the client behind the full cone NAT only if the internal host had previously sent a packet to IP address X.

A port restricted NAT is similar to a restricted cone NAT, but the restriction includes port numbers. Specifically, an external client can send a packet with a source IP address X and source port P to the client behind the port restricted NAT only if the client behind the port restricted NAT had previously sent a packet from IP address X and port P.

In a symmetric NAT, all requests from the same internal IP address and port, to a specific destination IP address and port, are mapped to the same external IP address and

port. If the same host sends a packet with the same source address and port, but to a different destination, a different mapping is used. Furthermore, only the external host that receives a packet can send a UDP packet back to the internal host.

Traversal of full cone, restricted cone and port restricted cone NATs is relatively straightforward and somewhat more complicated for symmetric NATs. NAT traversal may be implemented if a client is behind a symmetric NAT, e.g., as described in commonly-assigned US Patent Application Publication Number 20070076729. In particular, the client may perform a port prediction involving construction of a list of predicted transport addresses on the NAT that the client is behind. The client may then send an INVITE message containing the list of predicted transport addresses from the first node to a second client. The client behind the symmetric NAT may then perform a connectivity check with the second node using the predicted transport addresses. Connectivity checks may be performed by e.g., by sending simple traversal of user data protocol (UDP) through NAT (STUN) requests to each predicted transport address, in parallel. When the client behind the symmetric NAT receives these, it sends a STUN response to the second client. If the second client receives the STUN response, it can then begin to send media to that address.

There may be NAT types other than the four described above. In some cases it may be possible to traverse such NATs using standard techniques and in other cases the NAT behavior may be so unpredictable or unstable that communication with a client behind such a NAT is unreliable.

Because the duty of the host **107** is to communicate information between other peers **105A**, **105B**, and **105D** it is essential that the host **107** be behind a NAT of a type that does not interfere with its ability to communicate. In embodiments where the peers **105A**, **105B**, **105D** and the

host **107** are connected in a fully connected grid (FCG) it is particularly desirable to use a host behind a NAT that does not interfere with its ability to communicate in order to provide the highest level of service to the greatest number of peers. There have been cases where a NAT identified as having questionable (unknown) support for P2P has actually performed quite well against the other NAT types in use by the existing peers in the P2P network. By selecting a host **107** with the most favorable NAT profile, more reliable peer-to-peer communication may be obtained. As used herein, the term P2P communication generally refers to direct communication between client devices connected to a network. Examples of P2P applications include, but are not limited to, the voice over internet protocol (VoIP), bit torrent transmission, video transmission, file sharing, data sharing, and other types of direct data transfer between clients that does not exceed the bandwidth capabilities of an individual client. Once the host **107** has been established, peers **105A**, **105B**, **105D** can communicate with each other by initially transmitting information to the host **107** which then relays that information to the respective recipient peers **105A**, **105B**, **105D**, or they may transmit information directly after using host **107** to establish this direct communication path.

In certain embodiments, clients **105A**, **105B**, **105D**, and **107** may obtain their NAT profile information through a simple traversal of user data protocol (UDP) through NAT (STUN) server **103** that is associated with the external network **109**. STUN server **103** is a lightweight protocol proposed by Internet Engineering Task Force (IETF) that allows an IP enabled client to discover the presence and types of NAT that the client is behind. STUN **103** works with most NAT types and does not depend on any special behavior of the NAT. The STUN server **103** acts like a mirror held up to a client **105A** so that the client **105A** can see how its local transport address gets mapped to a public transport address. The client **105A** can also determine the type of

the NAT **111A** that the client **105A** is behind through communication with the STUN server **103**.

By way of example and not by way of limitation, each client **105A**, **105B**, **107**, **105D** may use a STUN server **103** to obtain NAT profile information, that it may then relay that information to the central server **101** in order for the central server **101** to determine which client would make the most favorable host **107**. Likewise, a client **105A**, **105B**, **107**, **105D** may use a STUN server **103** to obtain NAT profile information that may then be relayed to all other clients **105** attempting to communicate through the P2P network in order for the clients **105A**, **105B**, **107**, **105D** to best determine the host **107**.

FIG. 2 is a flow chart illustrating the basic method for host selection based on discovered NAT type from a client's perspective. Each client connected to the server **101** may collect its own NAT profile information for use with the P2P network, as indicated at **201**. This profile information may include information regarding the type that the client is behind, information regarding whether the NAT supports universal plug and play (UPnP), information regarding whether the NAT exhibits port preservation, and information regarding whether the NAT supports port predictability. As used herein, port preservation means that once an internal IP address mapped to a particular external port it will be consistently mapped to that particular port. As used herein, port predictability means that it is possible to predict the external port to which an internal IP address will be mapped, even if it is not always the same port. For example, the external port number may be consistently incremented with each attempted mapping of the internal IP address.

In addition to NAT behavior, each client's profile information may include quality of service (QoS) information. As used herein, the term QoS information, includes information regarding a client device that is

related to how well the client device can communicate with other client devices. By way of example, but not by way of limitation, such information may relate to how fast the client can communicate, how reliably the client can communicate, or some combination of both. Particular examples of QoS information include, but are not limited to ping time, bandwidth behavior, latency, geography, IP provider.

The profile information for each client may be used to create a priority list amongst all the clients connected to the server in order to best select the host for P2P communication. If there is a tie in priority between potential hosts, ordinal numbers may be assigned to determine which potential host is selected as the actual host. In some embodiments, such ordinal numbers may be assigned by a server as indicated at **203**, e.g., in the order in which the clients connected to the server. Alternatively, a distributed arbitration algorithm may be used to select a host from amongst two or more equally suitable potential hosts. If the initial host decides to leave the P2P network or is somehow disconnected from the network, this information may be used to select the next host for the P2P network.

Once the profile information has been collected by a given client, the NAT profile for that client is shared with other clients that are connected to the server. Profile information includes an expected count of the number of clients connected to the server, in order to determine whether an optimized P2P network has been reached. At this point, each client waits until every other client has submitted their profile information regarding the NAT they are behind as indicated at **205**. Once all NAT profiles have been submitted by the clients, a determination is made as to whether the client should be a host, peer or if they fail to meet the requirements for either as described at **207**. This determination may be made based on the profile information obtained for

each client described above. For example, by assigning a priority based on a number of factors, a host can be selected from amongst the available clients based on the client having the most favorable profile. The rest of the clients may be assigned as peers or may fail to be recognized as a peer or host based on their profile information. For example, a client behind a non-traversable NAT may not meet the requirements necessary to connect to the network as a peer or host. By way of example and not by way of limitation, the client that is unable to be recognized as a peer or host may update his profile to notify the rest of the P2P network of his status, and the expected count may be decremented to exclude this client from the P2P network.

Table 1 illustrates an example of a prioritization scheme that may be used to decide host assignments in a P2P network. By way of example, but not by way of limitation, the prioritization scheme may be broken down into 5 separate distinctions: Active, Likely, Unknown, In Progress, and Inactive. These are listed in host priority order in Table 1. An active tag indicates that a client is a very good candidate for host. A likely tag indicates that a client is a good candidate for host, but priority is still given to a client with an active tag. An unknown tag indicates that the network is unable to decide whether that particular client is a good candidate for host. An in progress tag indicates that the network is still deciding whether the client is a good candidate for host. Lastly, an inactive tag indicates that a client is unable to assume the duties of a host. In the example illustrated in Table 1, the priority tags may be based on 4 separate criteria: NAT type, universal plug and play (UPnP) capability, port preservation, and port predictability. However, several other factors may also be used in determining priority. These factors fall under a client's quality of service profile and may include QoS information including, but not limited to

ping time, bandwidth behavior, geography, latency, and IP provider.

TABLE 1

Prioritization Tag	Active	Likely (Scenario 1)	Likely (Scenario 2)	Unknown	Inactive
NAT Type	Cone (Full Cone, Restricted Cone, Port Restricted Cone)	Not Cone (Symmetric or anything else)	Symmetric	Symmetric	Not Cone (Symmetric or anything else)
UPnP	Doesn't matter	Yes	No	No	No
Port Preservation	Doesn't matter	Doesn't matter	Yes	Cannot determine	No
Port Predictability	Doesn't matter	Doesn't matter	Yes	Cannot determine	No

FIG. 3 is a flow chart illustrating a method for host selection based on discovered NAT type according to a specific embodiment of the present invention from the client's perspective. A client initially obtains profile information about the network address translator (NAT) that it is behind as described at **301**. This profile information includes information regarding the NAT type, whether the NAT supports universal plug and play (UPnP), whether the NAT supports port preservation, and whether the NAT supports port predictability. In addition to this information, profile information may also include quality of service information including, but not limited to ping time and bandwidth behavior, geography, latency, IP provider, and the like. This information may be used to create a priority amongst several clients in order to determine a host for supporting peer-to-peer (P2P) communication amongst these several different clients. If the initial host decides to leave the P2P network or is somehow disconnected from the network, this information may be used to select the next host for the P2P network.

Once the client has obtained the profile information regarding the NAT it is behind, the client then shares this NAT profile with other clients that are connected to the same server as indicated at **303**. Not only are these clients now aware of the profile information described above for all other clients, but they are also notified of the expected count of clients currently connected to the server. At this point, each client may wait until every other client has submitted their profile information regarding the NAT they are behind as indicated at **305**. Once all NAT profiles have been submitted by the clients, a determination is made as to whether the client should be a host, peer or if they fail to meet the requirements for either as indicated at **307**. This determination may be made based on the profile information obtained for each client described above. By assigning a priority based on a number of factors, a host can be selected amongst the clients based on having the most favorable profile. The rest of the clients may be assigned as peers or may fail to be recognized as a peer or host based on their profile information. For example, a client behind a non-traversable NAT may not meet the requirements necessary to connect to the network as a peer or host. By way of example and not by way of limitation, the client who is unable to be recognized as a peer or host may update his profile to notify the rest of the P2P network of his status, and the expected count may be decremented to exclude this client from the P2P network.

The client that is assigned the duty of being a host may then wait for the rest of the clients on the network to connect and also updates their profiles as indicated at **309**. Once all the clients have connected to the host, the expected count may be checked to determine whether all clients that are connected to the server have connected to the host as indicated at **311**. If the count is as expected, P2P communication is enabled as indicated at **317**. If however, the count is less than expected due

to one or more clients failing to meet the requirements of being a host or peer, then the count is reduced and the process may begin again at **303**.

If the client is determined to be a peer, it may attempt to connect to the host once the host has been determined as indicated at **313**. A determination is made as to whether the peer is able to connect to the host or fails to connect to the host as indicated at **315**. If the client is able to connect to the host, then it waits for all other peers to connect to the host before P2P communication is enabled as described at **317**. If the client is unable to connect to the host, then the client's profile information is updated and the process begins again at **305**. Ultimately, if the client is unable to connect to the host on multiple occasions, the count may be decremented leaving that individual client out of the P2P network.

If a peer failure occurs while collecting the NAT information, then this may be recorded as a failure and shared with the other peers. It may be assumed that sharing the data or waiting for the data will not fail, e.g., if the communication of this data is done through a reliable communication channel, such as the server **101**. If the communication to the server **101** fails at any point during this process then the entire process may be aborted and the remaining clients are notified of this disconnect as specified by the protocol used to communicate with the server **101**. The remaining clients will no longer be waiting for a response from the disconnected client.

Once a host has been identified from among the client devices, the server **101** may offload one or more functions to the P2P grid managed by the peer designated as the host. Such functions may be assumed by the host and/or any of the peers within the P2P grid. By way of example, but not by way of limitation, the P2P grid maybe a fully connected grid or a star topology where by the all

communication goes through the host. The topology of the grid is not strictly relevant to embodiments of this invention. What is relevant is that the primary point of connection (the host) to a P2P grid of any topology is established based on the behavior of a NAT associated with that primary point of connection.

A client device **400** may be configured to implement a method for host determination according to an embodiment of the present invention as shown in FIG. 4. By way of example, and without loss of generality, the client device **400** may be implemented as part of a digital television set, personal computer, video game console, personal digital assistant, mobile or handheld device such as a cellular phone or personal digital assistant, handheld video game device, portable email device and the like, or other digital device.

The device **400** may include a central processing unit (CPU) **402** and a memory **404** coupled to the CPU **402**. The CPU **402** may be configured to run software applications and, optionally, an operating system. Some embodiments of the present invention may take advantage of certain types of processor architecture in which the CPU **402** includes a main processor **402A** and one or more auxiliary processors **402B**. Each auxiliary processor **402B** may have its own associated local data storage. One example, among others of such processor architecture is a Cell Processor. An example of Cell Processor architecture is described in detail, e.g., in Cell Broadband Engine Architecture, copyright International Business Machines Corporation, Sony Computer Entertainment Incorporated, Toshiba Corporation August 8, 2005 a copy of which may be downloaded at http://cell.scei.co.jp/, the entire contents of which are incorporated herein by reference.

Referring again to FIG. 4, the receiving device **400** may also include well-known support functions **410**, such as input/output (I/O) elements **411**, power supplies (P/S) **412**, a clock (CLK) **413** and cache **414**. The device **400** may

further include a fast data storage device **415** such as a hard disk drive that provides non-volatile storage for applications and data. The fast storage device **415** may be used for temporary or long-term storage of files **416** retrieved from a slower data storage device. By way of example, the storage device **415** may be a fixed disk drive, removable disk drive, flash memory device, tape drive. Alternatively, the storage device **415** may be, e.g., a CD-ROM, DVD-ROM, Blu-ray, HD-DVD, UMD, or other optical storage devices. Files **416** from a slower storage device may be temporarily stored in a faster storage device in a hardware cache for quick loading into the memory **404**.

One or more user input devices **420** may be used to communicate user inputs from one or more users to the system **400**. By way of example, one or more of the user input devices **420** may be coupled to the client device **400** via the I/O elements **411**. Examples of suitable input devices **420** include keyboards, mice, joysticks, touch pads, touch screens, remote control units, light pens, still or video cameras, and/or microphones.

The client device **400** may communicate with one or more other client devices in a peer-to-peer network through a network interface **425** that facilitates communication via the electronic communications network **427**. The network interface **425** may be configured to implement wired or wireless communication over local area networks and wide area networks such as the Internet. The system **400** may send and receive data and/or requests for files via one or more message packets **426** over the network **427**. By way of example and not by way of limitation, the electronic communication network **427** may be a local area network, wide area network, or personal area network (e.g., blue-tooth) that can communicate between the device **400** and one or more other client devices.

The memory **404** may store applications and data for use by the CPU **402**. The memory **404** may be in the form of an

integrated circuit (e.g., RAM, DRAM, ROM, and the like). A computer program **401** may be stored in the memory **404** in the form of instructions that can be executed on the processor **402**.

The program **401** may include instructions that when executed by the processor cause it to implement a method for determining a host from among a group of two or more peer client devices in a peer-to-peer network. By way of example, and without loss of generality, the program **401** may, upon execution, cause the device **400** to implement the method **200** illustrated in FIG. 2 or the method **300** illustrated in FIG. 3. In particular the program **401** may cause the device **400** to a) collect NAT profile information for the device **400**; b) share the NAT profile information for the device with one or more other client devices **400'**, **400"** connected to the network **427**; c) receive NAT profile information for the one or more other client devices; and d) determine from the NAT profile information for the client device **400** and the NAT profile information for the one or more other client devices which of the client devices **400**, **400'**, **400"** to designate as the host.

NAT profile information **406** for the client device **400** and the other client devices **400'**, **400"** may be stored in the memory **404** for use in determining the host. Each of the client device **400**, **400'**and **400"** may be behind a corresponding NAT **403**, **403'** and **403"**. The client device **400** may reside behind a network address translator (NAT) **403**, which translates an internal IP address for the client device **400** to a public IP address that is seen by other devices. It is noted that the NAT **403** is usually not a part of the client device **400** although some types of client device may include a NAT. Furthermore, embodiments of the present invention may be implemented if any or all of the client devices **400**, **400'**, **400"** is not behind a NAT at all.

The NAT profile information **406** may include, but is not limited to, information about the type of NAT (if any) that the client device **400** is behind, the ability of the NAT to engage in universal plug and play (UPnP), the NAT's ability to maintain port preservation, and the NAT's port predictability. The program **401** may implement selection may be based on the criteria described above, with respect to Table 1. Each of the other client devices **400', 400"** may be similarly configured and may implement the same host selection process. If all client devices **400**, **400'**, and **400"** use the same program **401** and the profile information **406** it is reasonable to expect that they should all make the same host determination.

In some situations, it may be desirable for the program **401** to apply a supplemental arbitration filter based on additional information **408**, such as quality of service information, to arbitrate a determination the host device from among two or more equally likely candidate devices. Quality of service information may include factors such as the client's ping time, bandwidth behavior, geography, latency, IP provider, etc, and other factors as described above. Such additional information **408** may also be stored in the memory **404**.

The client device **400** may further comprise a graphics subsystem **430**, which may include a graphics processing unit (GPU) **435** and graphics memory **437**. The graphics memory **437** may include a display memory (e.g., a frame buffer) used for storing pixel data for each pixel of an output image. The graphics memory **437** may be integrated in the same device as the GPU **435**, connected as a separate device with GPU **435**, and/or implemented within the memory **404**. Pixel data may be provided to the graphics memory **437** directly from the CPU **402**. Alternatively, the CPU **402** may provide the GPU **435** with data and/or instructions defining the desired output images, from which the GPU **435** may generate the pixel data of one or more output images. The data and/or

instructions defining the desired output images may be stored in memory **406** and/or graphics memory **437**. In an embodiment, the GPU **435** may be configured (e.g., by suitable programming or hardware configuration) with 3D rendering capabilities for generating pixel data for output images from instructions and data defining the geometry, lighting, shading, texturing, motion, and/or camera parameters for a scene. The GPU **435** may further include one or more programmable execution units capable of executing shader programs.

The graphics subsystem **430** may periodically output pixel data for an image from the graphics memory **437** to be displayed on a video display device **450**. The video display device **450** may be any device capable of displaying visual information in response to a signal from the device **400**, including CRT, LCD, plasma, and OLED displays that can display text, numerals, graphical symbols or images. The digital broadcast receiving device **400** may provide the display device **450** with a display driving signal in analog or digital form, depending on the type of display device. In addition, the display **450** may be complemented by one or more audio speakers that produce audible or otherwise detectable sounds. To facilitate generation of such sounds, the client device **400** may further include an audio processor **455** adapted to generate analog or digital audio output from instructions and/or data provided by the CPU **402**, memory **404**, and/or storage **415**.

The receiving device **400** may optionally include a position location device **470**. Such a device may be based on any suitable technology capable of providing information on the geographic location of a device. Examples of existing technology include global positioning satellite (GPS) technology, inertial guidance technology, and the like. Information from such devices may be used in digital broadcast data applications such as navigation for mobile or hand-held devices.

According to certain embodiments, it may be useful to determine a geographic location of the client device **400**. Certain QoS considerations, such as bandwidth, ping time and latency may be affected by the location of the device. By way of example and not by way of limitation, to facilitate host determination, the position-locating device **470** may provide geographic location information that the program **401** may use in determining whether the client device **400** is a good candidate to be the host for a peer-to-peer service.

The components of the device **400**, including the CPU **402**, memory **404**, support functions **410**, data storage devices **415** user input devices **420**, network interface **425**, graphics unit **430**, audio processor **455** and position location device **470** may be operably connected to each other via one or more data buses **460**. These components may be implemented in hardware, software or firmware or some combination of two or more of these.

Embodiments of the present invention allow for cooperative host determination and host migration in peer-to-peer network situations.

While the above is a complete description of the preferred embodiment of the present invention, it is possible to use various alternatives, modifications and equivalents. Therefore, the scope of the present invention should be determined not with reference to the above description but should, instead, be determined with reference to the appended claims, along with their full scope of equivalents. Any feature described herein, whether preferred or not, may be combined with any other feature described herein, whether preferred or not. In the claims that follow, the indefinite article "A", or "An" refers to a quantity of one or more of the item following the article, except where expressly stated otherwise. The appended claims are not to be interpreted as including means-plus-function limitations, unless such

a limitation is explicitly recited in a given claim using the phrase "means for".

WHAT IS CLAIMED IS:

1. In a peer, a method for determining which of two or more peers to designate as a host for a service within a peer to peer grid, the method comprising:
a) collecting NAT profile information for the peer;
b) sharing the NAT profile information for the peer with one or more other peers of the two or more peers;
c) receiving NAT profile information for the one or more other peers; and
d) determining from the NAT profile information for the peer and the NAT profile information for the one or more other peers which of the two or more peers to designate as the host.

2. The method of claim 1, wherein a) includes collecting NAT profile information through use of a STUN server.

3. The method of claim 1, wherein b) includes relaying the NAT profile information for the peer to the one or more other peers via a pre-established connection between the peer and the one or more other peers.

4. The method of claim 3, wherein the pre-established connection includes use of a server for an application common to the two or more peers.

5. The method of claim 4, wherein the server offloads one or more functions to the peer to peer grid, which is managed by the peer designated as the host.

6. The method of claim 3, wherein the pre-established connection includes one of the one or more other peers acting as a host.

7. The method of claim 3, wherein c) includes receiving the NAT profile information for the one or more other peers via the pre-established connection.

8. The method of claim 1, wherein d) includes assigning priority values to each of the two or more peers based

on the NAT profile information for the peer and the NAT profile information for the on one or more other peers and selecting the host based on the priority values.

9. The method of claim 8 wherein the NAT profile information includes a NAT type for a NAT that the peer is behind, information regarding whether the NAT supports universal plug and play (UPnP), port predictability information for the NAT, and port preservation information for the NAT.

10. The method of claim 8 wherein the NAT profile information for the one or more other peers includes a NAT type for a NAT that another peer is behind, information regarding whether the NAT supports universal plug and play (UPnP), port predictability information for the NAT, and port preservation information for the NAT.

11. The method of claim 8, wherein, if two or more of the two or more peers are of equal priority values, d) includes resolving which of the two or more potential hosts of equal priority values to designate as the host.

12. The method of claim 11, wherein resolving which of the two or more potential hosts of equal priority values to designate as the host includes designating as the host one of the two or more potential hosts of equal priority value having a higher ordinal value, wherein a different ordinal value is assigned to each of the two or more potential hosts.

13. The method of claim 11, wherein resolving which of the two or more potential hosts of equal priority values to designate as the host includes a distributed arbitration algorithm.

14. The method of claim 1, further comprising, if the peer is designated as the host, notifying the one or

more other peers of the peer's designation as host and waiting for the one or more other peers to connect.

15. The method of claim 14, further comprising comparing a number of peers connected to the host to an expected count and sharing the number of peers connected to the host with the one or more other peers if the number of peers connected to the host does not equal the expected count.

16. The method of claim 1, further comprising attempting to connect the peer to the host if the peer is not designated as the host.

17. The method of claim 14, further comprising, if the peer fails to connect to the host, notifying the two or more other peers that the peer has failed to connect to the host.

18. A peer device, comprising:
a computer processor;
a computer memory coupled to the processor;
a network interface coupled to the processor, wherein the network interface is configured to facilitate communication between the processor and one or more other peer devices connected in a peer-to-peer grid; and
a set of computer readable instructions embodied in the memory, wherein the computer-readable instructions are configured to implement a method for determining which of peer device in the peer-to-peer grid to designate as a host for a service within the peer to peer grid, the method comprising:
a) collecting NAT profile information for the peer;
b) sharing the NAT profile information for the peer with one or more other peers of the two or more peers;
c) receiving NAT profile information for the one or more other peers; and
d) determining from the NAT profile information for the peer and the NAT profile information for the one or

more other peers which of the two or more peers to designate as the host.

19. A computer readable medium having a set of computer readable instructions embodied therein, wherein the computer-readable instructions are configured to implement a method for determining which of two or more peer devices in a peer-to-peer grid to designate as a host for a service within the peer to peer grid, the method comprising:
a) collecting NAT profile information for the peer;
b) sharing the NAT profile information for the peer with one or more other peers of the two or more peers;
c) receiving NAT profile information for the one or more other peers; and
d) determining from the NAT profile information for the peer and the NAT profile information for the one or more other peers which of the two or more peers to designate as the host.

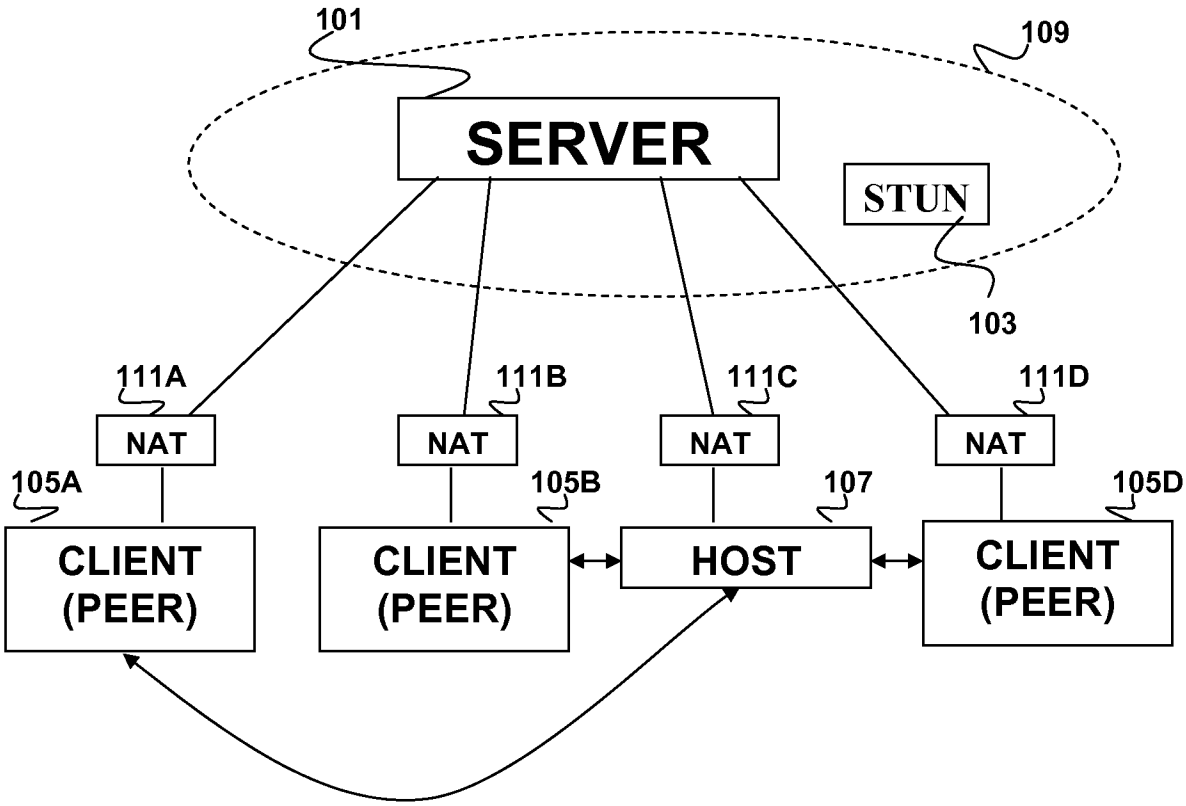
101
109
SERVER
STUN
103
111A
111B
111C
111D
NAT
NAT
NAT
NAT
105A
105B
107
105D
CLIENT (PEER)
CLIENT (PEER)
HOST
CLIENT (PEER)

FIG. 1

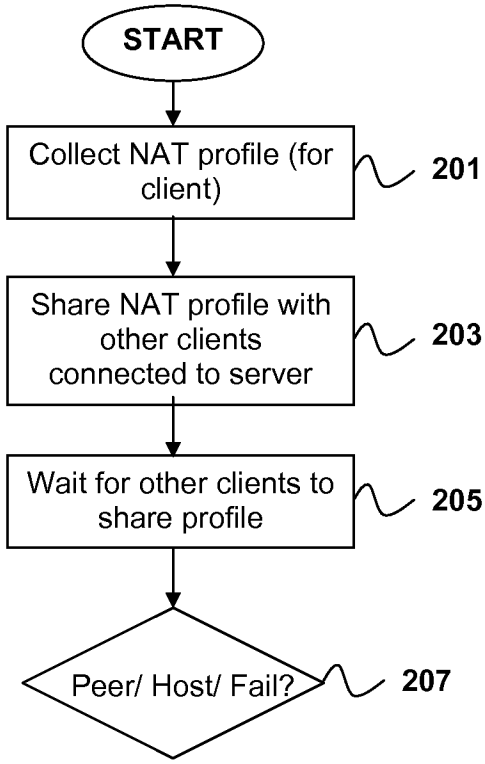
START
Collect NAT profile (for client)
201
Share NAT profile with other clients connected to server
203
Wait for other clients to share profile
205
Peer/ Host/ Fail?
207

FIG. 2

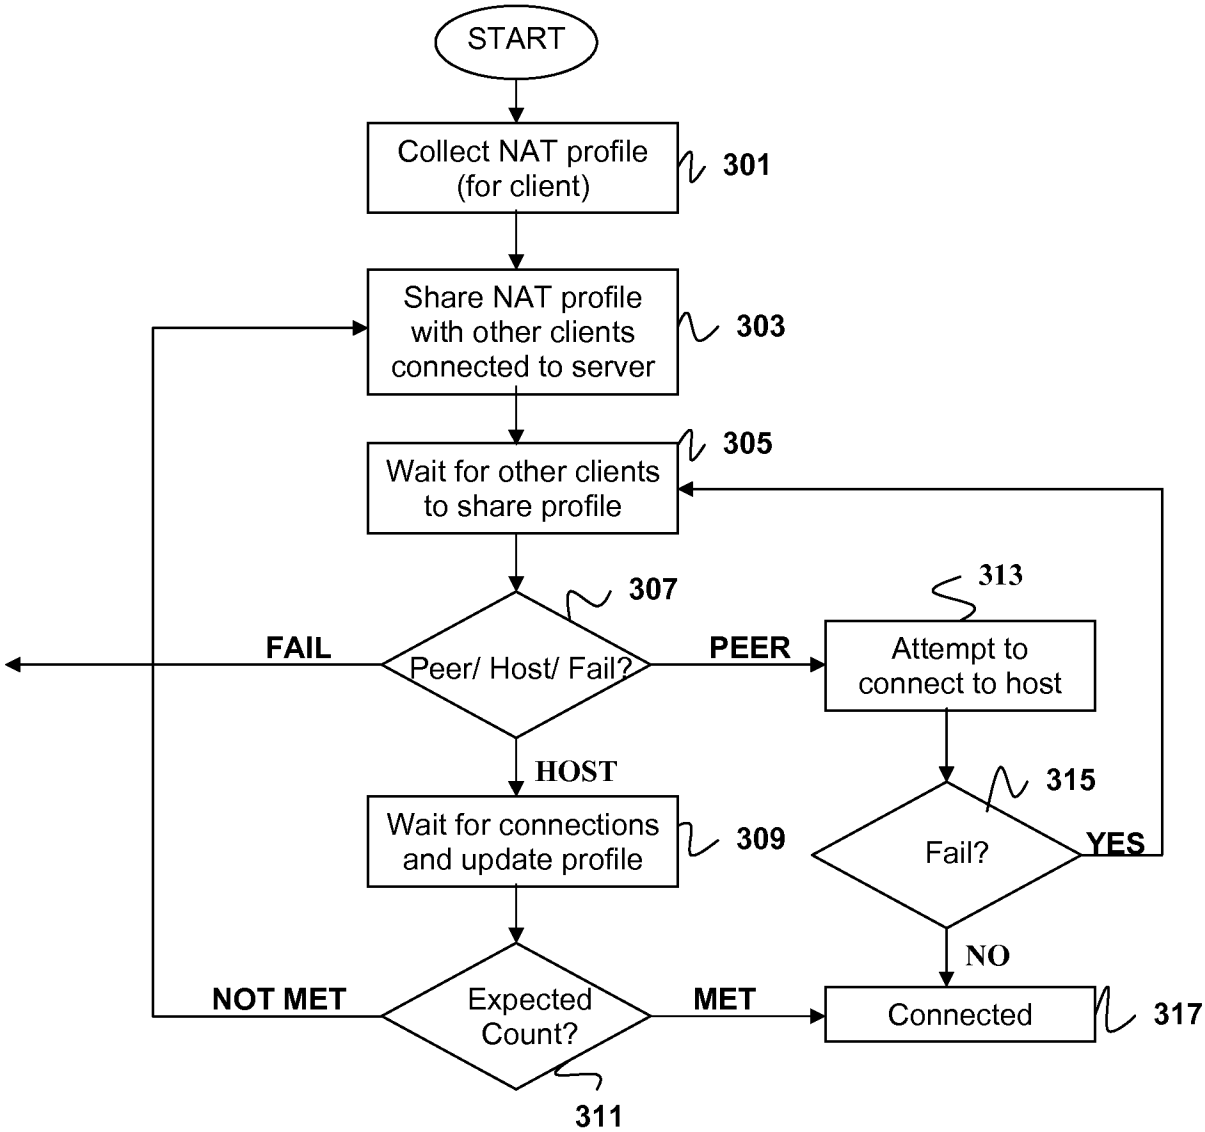
START
Collect NAT profile (for client)
301
Share NAT profile with other clients connected to server
303
Wait for other clients to share profile
305
307
Peer/ Host/ Fail?
FAIL
PEER
HOST
313
Attempt to connect to host
315
Fail?
YES
NO
Wait for connections and update profile
309
Expected Count?
NOT MET
MET
Connected
317
311

FIG. 3

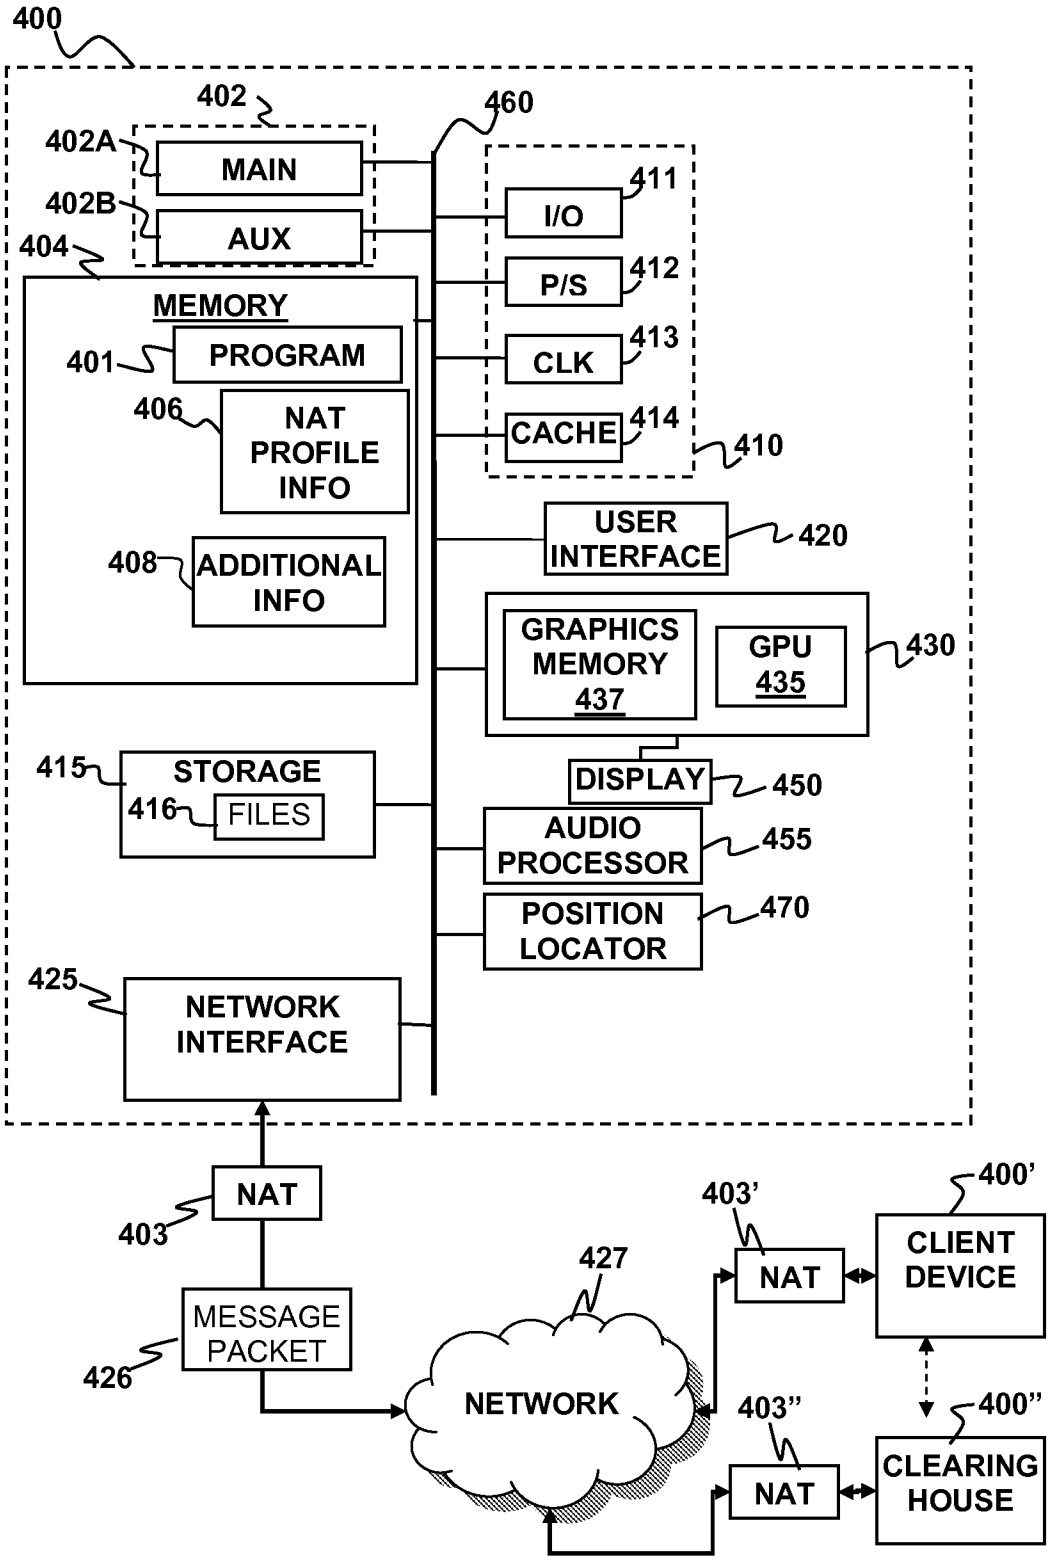
400
402
460
402A
MAIN
402B
AUX
404
MEMORY
401
PROGRAM
406
NAT PROFILE INFO
408
ADDITIONAL INFO
411
I/O
412
P/S
413
CLK
414
CACHE
410
USER INTERFACE
420
GRAPHICS MEMORY 437
GPU 435
430
DISPLAY
450
AUDIO PROCESSOR
455
POSITION LOCATOR
470
415
STORAGE
416
FILES
425
NETWORK INTERFACE
NAT
403
MESSAGE PACKET
426
427
NETWORK
403'
NAT
400'
CLIENT DEVICE
403"
NAT
400"
CLEARING HOUSE

FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2009/057192

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8) - G06F 15/16 (2009.01)
USPC - 709/245
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC(8) - G06F 15/16 (2009.01)
USPC - 709/245

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
USPTO WEST System (US, USPG-PUB, EPO), Google Patent, PatBase

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2007/0076729 A1 (TAKEDA) 05 April 2007 (05.04.2007) entire document	1-8, 11-19
---		---
Y		9, 10
Y	US 7,243,141 B2 (HARRIS) 10 July 2007 (10.07.2007) entire document	9, 10
A	US 2006/0218624 A1 (RAVIKUMAR et al) 28 September 2006 (28.09.2006) entire document	1-19

☐ Further documents are listed in the continuation of Box C. ☐

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

Date of the actual completion of the international search	Date of mailing of the international search report
27 October 2009	06 NOV 2009

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

Form PCT/ISA/210 (second sheet) (July 2009)