

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
7 September 2001 (07.09.2001)

PCT

(10) International Publication Number
WO 01/65357 A1

(51) International Patent Classification⁷: **G06F 9/00**

(21) International Application Number: PCT/US01/04734

(22) International Filing Date: 14 February 2001 (14.02.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/516,084 1 March 2000 (01.03.2000) US

(71) Applicant (for all designated States except US): **INTER-
ACTIVE INTELLIGENCE, INC.** [US/US]; Suite 300,
8909 Purdue Road, Indianapolis, IN 46268 (US).

(74) Agents: **PAYNTER, L., Scott** et al.; Woodward, Emhardt,
Naughton, Moriarty & McNett, Suite 3700, Bank One Cen-
ter/Tower, 111 Monument Circle, Indianapolis, IN 46204
(US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ,
DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR,
HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR,
LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ,
NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM,
TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian
patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European
patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,
IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF,
CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

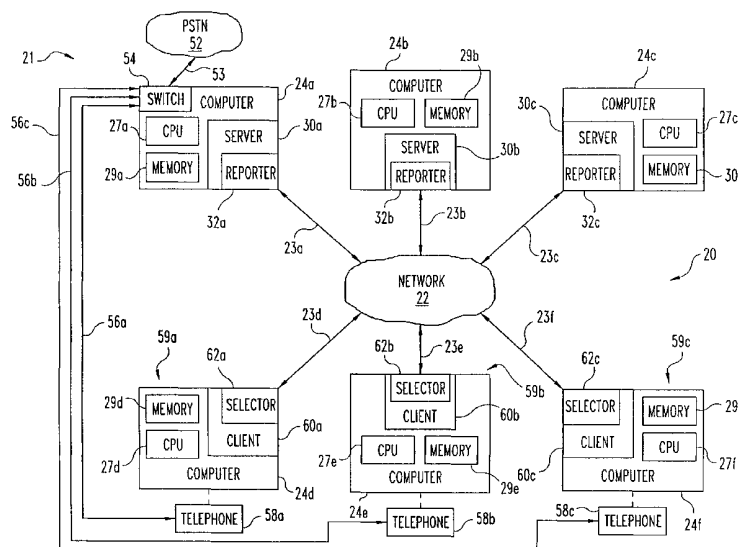
(72) Inventor; and

(75) Inventor/Applicant (for US only): **FLORMAN, Bruce**
[US/US]; 214 Bays Drive, Noblesville, IN 46060 (US).

Published:
— with international search report

[Continued on next page]

(54) Title: PROCESSING LOAD DISTRIBUTION TECHNIQUES FOR CALL CENTERS AND OTHER PROCESSING SYSTEMS



(57) Abstract: A system of network computers is disclosed having a number of servers (30a-30c) and a number of clients (60a-60c). The clients are each provided with a server selector routine (62a-62c) and the servers are each provided with a server reporter routine (32a-32c). The server selector routine of any of the clients is operable to obtain server processing load information from the server reporter routine of any of the servers. The server selector routine is responsive to a client activation request for a software task designated for server execution to select two or more servers to interrogate for server processing load information and designate one of the servers to execute the software task based in the server processing load information obtained by this interrogation.



WO 01/65357 A1



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

PROCESSING LOAD DISTRIBUTION TECHNIQUES FOR CALL CENTERS AND OTHER PROCESSING SYSTEMS

5

BACKGROUND

The present invention relates to computerized processing, and more particularly, but not exclusively, relates to distributing processing tasks among a number of networked computers.

10 With the advent of computer networks, interest has grown in the efficient distribution of processing tasks among the resources coupled to the network. The ability to balance processing load among these resources has become one focus of distributed processing efforts. This interest is especially acute in the area of computer networks utilized to manage various communications, such as telephone calls, voice mail, e-mail, telefaxes, electronic chats, and the like.

15 Many load balancing approaches introduce a centralized component between resources requesting services and those resources capable of providing such services. Unfortunately, central components can prove to be a processing bottleneck, impeding the speed with which processes are distributed and executed. Also, centralized components can adversely impact overall reliability of the
20 network system.

Furthermore, it is occasionally desired to apply processing load balancing to legacy applications and components for which load balancing was not provided. Unfortunately, the addition of current distributed processing and corresponding load balancing mechanisms to legacy applications often requires undesirable
25 modifications. Thus, there is a demand for further contributions in this area of technology.

SUMMARY

One form of the present invention is a unique processing load distribution technique. Other forms include unique systems and methods to balance processing load among a number of servers of a computer network.

5 A further form includes operating a computer system that has several clients and servers, and generating an activation request with one of the clients for a software task that has been designated for server execution. One or more of the servers are interrogated by this client and the client is provided information corresponding to processing load for each of the servers interrogated. One of the
10 servers is selected to execute the task as a function of this information.

Still further forms, objects, features, aspects, benefits, advantages, and embodiments of the present invention shall become apparent from the detailed description and drawings provided herewith.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic view of a computer system according to one embodiment of the present invention.

5 FIG. 2 illustrates a flow chart of a distribution procedure performed by the system of FIG. 1.

FIGs. 3A and 3B illustrate a flow chart for a client-side distribution/load balancing routine executed as part of the server selection operation of FIG. 2.

10 FIG. 4 illustrates a flow chart of a server-side distribution/load balancing routine executed as part of the selection operation of FIG. 2.

FIG. 5 illustrates a flow chart of a start-up and maintenance routine to support the procedure of FIG. 2.

DETAILED DESCRIPTION OF SELECTED EMBODIMENTS

For the purpose of promoting an understanding of the principles of the invention, reference will now be made to the embodiments illustrated in the drawings and specific language will be used to describe the same. It will nevertheless be understood that no limitation of the scope of the invention is thereby intended. Any alterations and further modifications in the described embodiments, and any further applications of the principles of the invention as described herein are contemplated as would normally occur to one skilled in the art to which the invention relates.

FIG. 1 schematically illustrates computer system 20 of one embodiment of the present invention. System 20 includes computer network 22 coupling together a number of computers 24a, 24b, 24c, 24d, 24e, 24f (collectively designated computers 24) by corresponding network pathways 26a, 26b, 26c, 26d, 26e, 26f (collectively designated pathways 26). Network 22 may be in the form of a Local Area Network (LAN), Wide Area Network (WAN), or other network type as would occur to those skilled in the art. Computers 24 may all be of the same type, or a heterogeneous combination of different computing devices coupled to network 22.

Each of computers 24 is further designated as including a client or server. For computers 24a, 24b, 24c; servers 30a, 30b, 30c are respectively designated. For computers 24d, 24e, 24f; clients 60a, 60b, 60c are respectively designated. While computers 24 are each illustrated as being a server or client, it should be understood that any of computers 24 may be arranged to include both a client and server.

In the case of servers 30a, 30b, 30c (collectively designated servers 30), each has a broker/reporter entity 32a, 32b, 32c (collectively designated brokers 32), respectively. In the case of clients 60a, 60b, 60c (collectively designated clients 60), each has a corresponding selector agent 62a, 62b, 62c (collectively designated agents 62), respectively.

Each computer 24 includes a corresponding programmable processor or Central Processing Unit (CPU) 27a, 27b, 27c, 27d, 27e, 27f (collectively designated processors 27) and computer-readable memory 29a, 29b, 29c, 29d, 29e,

29f (collectively designated memories 29), respectively. It should be understood that each computer 24 may include more than one processor or CPU and more than one type of memory, each of memories 29 being representative of one or more types associated with the corresponding computer 24. Furthermore, it should be
5 understood that while six computers 24 are illustrated, more or fewer may be utilized in alternative embodiments.

Processors 27 may each be comprised of one or more components configured as a single unit. Alternatively, when of a multi-component form, processors 27 may each have one or more components located remotely relative to
10 the others. One or more components of each processor 27 may be of the electronic variety defining digital circuitry, analog circuitry, or both. In one embodiment, processors 27 each have a conventional microprocessor architecture provided in the form of one or more integrated circuit chips. Each processor 27 may include any oscillators, control clocks, interfaces, signal conditioners, filters, limiters,
15 converters, communication ports, and other types of operators as would occur to those skilled in the art to implement the present invention.

Memories 29 may each include one or more types of solid-state electronic memory, magnetic memory, or optical memory, just to name a few. By way of nonlimiting example, each memory 29 may include solid-state electronic Random
20 Access Memory (RAM), Sequentially Accessible Memory (SAM) (such as the First-In, First-Out (FIFO) variety or the Last-In First-Out (LIFO) variety), Programmable Read Only Memory (PROM), Electrically Programmable Read Only Memory (EPROM), or Electrically Erasable Programmable Read Only Memory (EEPROM); an optical disc memory (such as a DVD or CD ROM); a
25 magnetically encoded hard disc, floppy disc, tape, or cartridge media; or a combination of any of these memory types. Also, memory 29 may be volatile, nonvolatile or a hybrid combination of volatile and nonvolatile varieties.

System 20 further illustrates Public Switched Telephone Network (PSTN) 52 coupled to computer-controlled telephone switch 54 of computer 24a by
30 pathway 53. Switch 54 is also coupled to telephones 58a, 58b, 58c (collectively designated telephones 58) by pathways 56a, 56b, 56c (collectively designated pathways 56), respectively. Telephones 58 may be in the form of a handset,

headset, or other arrangement as would occur to those skilled in the art.

Telephones 58a, 58b, 58c are each associated with one of computers 24d, 24e, 24f, respectively, to provide stations 59a, 59b, 59c (collectively designated stations 59).

Switch 54 may be arranged in the form of a Private Branch Exchange (PBX),
5 predictive dialer, Automatic Call Distributor (ACD), a combination of these, or another switching configuration as would occur to those skilled in the art.

In one embodiment, system 20 operates as a call center with computer 24a being configured as a call center server host and computers 24d, 24e, 24f each arranged as a call center client host. Accordingly, system 20 is designated as
10 including call center 21. For such an arrangement, additional telephones 58 may be connected to switch 54 that each correspond to an additional client host to provide more stations 59 (not shown). Typically call center applications of system 20 would include many more stations of this type, but only a few have been illustrated in FIG. 1 to preserve clarity. Also, more than one computer 24 may be
15 configured as a call center server host.

Alternatively or additionally, system 20 may be arranged to provide for the coordination, conversion, and distribution of a number of different forms of communication, such as telephone calls, voice mail, faxes, e-mail, web chats, web call backs, and the like. Furthermore, business/customer data associated with
20 various communications may be selectively accessed with system 20. This data may be presented to an operator (call center agent) at each station 59 by way of a station display monitor operably coupled to the corresponding computer 24. In addition to a display monitor, each computer 24 of station 59 may be operably coupled to one or more operator input devices such as a keyboard or mouse, just to
25 name a few.

It should be recognized that network 22 may include one or more elements of PSTN 52. Indeed, in an alternative embodiment, network 22 and PSTN 52 are provided as a common network. Still other alternative embodiments of the present invention are not directed to telephone or communication operations.

30 Correspondingly, such alternatives need not include PSTN 52, switch 54, or telephones 58.

System 20 is arranged to distribute software tasks from clients 60 among servers 30 via network 22. Such tasks include, but are not limited to instances of an object class, dedicated code segments, device interface operations, or any other forms of software algorithms, routines, or modules as would occur to those skilled in the art. Processing load balancing is performed through communications
5 between agents 62 and brokers 32. Generally, when one of clients 60 has a task that can be distributed among servers 30, its corresponding agent 62 communicates over network 22 with the brokers 32 to select a particular server 30 to perform the task. During this communication, agent 62 for the requesting client 60 interrogates
10 each broker 32 to receive a bid corresponding to its server's ability to promptly perform the task to be distributed. The interrogating agent 62 accumulates the bids from the interrogated brokers 32 and compares the bids to select a server 30. Connection/distribution information is provided by the broker 32 of the selected server 30 to establish remote performance of the task.

Referring additionally to FIG. 2, one embodiment for implementation with system 20 is illustrated in flow chart form as software task distribution procedure 70. As shown, procedure 70 illustrates management of only one software task requested by one of clients 60 for clarity; however, it should be understood that procedure 70 may be executed multiple times by one or more of clients 60.
15 Procedure 70 begins with an activation request for a software task by one of clients 60 in stage 72. Procedure 70 then proceeds to conditional 74 to determine whether the activation request is for a software task of the type to be distributed to one of servers 30. If the software task is not to be distributed (corresponding to a negative outcome of the test of conditional 74), then procedure 70 continues with stage 76
20 to execute the task locally or with a designated remote server. Procedure 70 halts after execution of stage 76. If the test of conditional 74 is positive, then distribution branch 80 of procedure 70 is engaged, as indicated by the downwardly pointing arrow in FIG. 2. Branch 80 begins with stage 82. In stage 82, the activation request is directed to the agent 62 corresponding to the requesting client
25 30 60.

In one nonlimiting example, conditional 74 is not implemented as a dedicated software or hardware test, but rather is executed for a given task based

on the nature of a corresponding entry in a system registry. For this example, the registry entry for a task that is to be distributed references an intermediate agent, such as agent 62, which performs distribution and/or load balancing as more fully described hereinafter. In contrast, tasks of a conventional type may be registered
5 in the usual manner, without an intervening agent.

In a further nonlimiting example, another more specific implementation of stages 72, 76, 82 and conditional 74 of procedure 70 is described with respect to an arrangement in which computers 24 each include the MICROSOFT WINDOWS NT 4.0 operating system to perform client and server operations. This
10 arrangement is particularly suited for operation of a call center with system 20 as previously described. The Distributed Component Object Model (DCOM) feature of the NT 4.0 operating system provides a mechanism to distribute class object activation requests from one of clients 60 to one of servers 30. Common functions used for distributed requests are the functions CoCreateInstance and
15 CoCreateInstanceEx which are wrappers for calling the CoGetClassObject and IClassFactory::CreateInstance functions. The CoGetClassObject function serves as a starting point for activation of a software task, which for this embodiment may be in the form of an instance of a desired class object. The CoGetClassObject function takes a set of class context flags (CLSCTX) and server information
20 structure (COSERVERINFO) as parameters. Unfortunately, the NT 4.0 operating system does not include a mechanism to regulate processing load between multiple servers. As a result, these functions are typically set up to operate according to a default behavior without load balancing among available servers.

Generally, when the DCOM mechanism on the client side (hereinafter
25 designated a DCOM client) wishes to initiate an instance of an object, it first checks an internal table to see if there is a running instance of the class object already available and returns a pointer to (a proxy for) that class object if one is found. If there is not a running instance of the class object, the DCOM client looks in the NT 4.0 operating system registry to determine how to locate or start an
30 appropriate server. If the class is under the registry index or "key" InprocServer 32 of NT 4.0, the Dynamic Link Library (DLL) named by this key's value is loaded with the LoadLibrary function, the address of the _DllGetClassObject entry point

is queried with the GetProcAddress function , and that entry point is invoked to get the necessary pointer. In contrast, if the class is registered with the LocalServer32 key, the executable program (EXE) named by that key's value is started with the CreateProcess function, and the DCOM mechanism waits up to 30 seconds for the new process to register its class objects. Once the requested class object is registered, the DCOM mechanism returns the pointer to that object. Additional registry settings can be used to cause the server to be activated on a different host, in which case the DCOM client contacts the DCOM mechanism operating on the remote host which takes over the responsibility of finding or starting an appropriate DCOM server.

It has been found that this existing DCOM mechanism may be utilized to provide for load balancing of existing applications and components running under the NT 4.0 operating system. For this procedure, when a client activation request for an object instance corresponding to stage 72 of procedure 70 occurs, a DLL is inserted under the InprocServer32 key for a corresponding DCOM class to be distributed across multiple servers 30. For this embodiment, the inserted DLL intercepts and re-directs activation requests from any client 60 making a CoGetClassObject call that includes CLSCTX_INPROC_SERVER. Moreover, for a DLL that can run in a surrogate process under the NT 4.0 operating system, it can also be arranged to intercept and re-direct activation requests from any client specifying an object to be executed by an out-of-process server.

Accordingly, for this embodiment, the client activation request stage 72 and conditional 74 test may be accomplished through this interception and re-direction procedure using the DCOM tools of the NT 4.0 operating system. Further, the inserted DLL may be used to direct the activation request to the agent 62 of the requesting client 60 as illustrated in stage 82; where agent 62 may be in the form of an NT compatible software object class. Naturally, in another embodiment of the present invention in which computers 24 do not include the NT 4.0 operating system, stages 72, 76, 86 and/or conditional 74 are implemented differently. Indeed, in other embodiments, different techniques may additionally or alternatively be utilized to implement stage 72, conditional 74, stage 76, and/or stage 82, regardless of the type of operating system or systems used.

Once stage 82 is performed, stage 84 is encountered to designate one of servers 30 to perform the task. This designation takes place by performing a server selection operation with the agent 62 for the requesting client 60 and brokers 32. FIGs. 3A and 3B illustrate a client-side distribution/load balancing routine 160
5 executed by the agent 62 as part of this operation. Routine 160 starts with stage 162. In stage 162, the agent 62 for the requesting client 60 excludes any of servers 30 that are known to lack the particular software task desired to be distributed. This capability and other data about servers 30 is maintained by each agent 62 and is updated at various stages as further described hereinafter. Typically, each agent
10 62 maintains server information in memory 29 of the computer 24 on which is resides.

Conditional 164 is next encountered which tests whether the number of servers known to have the capability to perform the task exceed a maximum limit designated by "LIMIT". If this limit is exceeded, then stage 166 is encountered to
15 perform a reduction of the number of server candidates. This reduction is performed by accessing established server data maintained by the agent 62 of the requesting client 60. This data includes estimated bids expected from the various servers 30 that were not excluded as lacking the capability to perform the requested task in stage 162. In stage 166, the estimates of the eligible servers 30 are
20 compared and sorted to identify a subset of servers 30 having the lowest estimated bids and a total membership of LIMIT. This subset is provided as the candidate servers in subsequent stages of routine 160. The bid estimates are determined as a function of bid information from prior executions of routine 160, timing, and other factors as explained in more detail in connection with stages 194 and 196 of
25 routine 160 (see FIG. 3B).

The test of conditional 164 and reduction of stage 166 is directed to maintaining the number of servers 30 involved in routine 160 below a manageable level. Typically, the maximum number of candidate servers 30, corresponding to the value of LIMIT, is selected so that the burden presented to system 20 by these
30 servers does not overtake efficiency gains realized by performing task distribution. The value selected for LIMIT may be based on one or more factors, including but not limited to: the total number of servers, any network performance differences

of the servers, the type and complexity of the tasks to be distributed, and the degree of loading of the servers 30 and network 22. The value of LIMIT may be static or dynamically adjusted in accordance with one or more parameters, may be uniformly or individually assigned with respect to each client 60 and/or task, or
5 otherwise assigned as would occur to those skilled in the art. It should be appreciated that for a first time distribution of a given software task or a first time activation of a given agent 62, stages 162 and 166 may not result in any exclusion or reduction, respectively, of candidate servers. Alternatively, a randomized reduction may be performed in stage 166 when anticipated or estimated bids have
10 not yet been determined. After performing stage 166, routine 160 proceeds to stage 168. Likewise, if the test of conditional 164 is negative, routine 160 bypasses stage 166 to proceed directly to stage 168. Stage 168 sets index counter N to 1 (N=1).

Next, server interrogation loop 170 is entered starting with stage 172. In
15 stage 172, the server 30 corresponding to index N (alternatively designated "server N") is interrogated by the agent 62 of the requesting client 60 to solicit a bid corresponding to server N's ability and/or availability to perform the task. The response from server N is communicated by its corresponding broker 32 to the interrogating agent 62 via network 22. For the NT 4.0 operating system
20 embodiment previously described, broker 36 may be an instance of an object class established through the DCOM protocol for distributing tasks to remote server hosts. As reflected in conditional 174, if the response of the broker 32 indicated that the task is not available for the particular server N, then the data maintained by the interrogating agent 62 is updated in stage 175 to indicate server N lacks the
25 requisite task capability, and routine 160 proceeds to stage 178 to exclude server N as a candidate. Accordingly, future interrogations by the agent 62 of the requesting client 60 to distribute the pending type of task will exclude the current server N during performance of stage 162.

If the test of conditional 174 is affirmative, routine 160 continues with
30 conditional 176. Conditional 176 times how long it takes the interrogated broker 36 to respond. If a response is not received within a predefined time interval, then server N is excluded as a candidate in stage 178. This exclusion corresponds to a

server and/or broker 32 that is too busy to respond to an interrogating agent 62 within a reasonable amount of time. This time limit may be selected based on one or more of the factors described in connection with the selection of LIMIT.

Likewise, the time limit may be statically or dynamically assigned; may be specific
5 to the client, server, or task type; or may be the same for all clients 60, servers 30, and/or task types. Typically, the failure of a given broker 32 to respond within the time limit does not exclude that server from future agent interrogations unlike the case when it is determined a task is unavailable via conditional 174.

After performance of stage 178, loop 170 continues with stage 180
10 illustrated in FIG. 3B. Likewise, stage 180 is reached if the test of conditional 176 is negative. In stage 180, data maintained by the interrogating agent 62 is updated for the particular server N undergoing interrogation in loop 170. Next, conditional 182 is encountered to determine if the index counter N equals the value previously described in connection with conditional 164. If N does not equal LIMIT
15 ($N \neq \text{LIMIT}$), routine 160 continues at stage 184 to increment index counter N ($N=N+1$). From stage 184 routine 160 returns to stage 172 of loop 170 as illustrated in FIG. 3A. Correspondingly, loop 170 is again performed with a different one of the candidate servers 30 designated as server N. In this manner, loop 172 is repeated, interrogating a different broker 32 with each repetition, until
20 a total number of servers 30 interrogated equals LIMIT.

Once N reaches LIMIT ($N=\text{LIMIT}$), conditional 182 is satisfied and loop 170 is exited, as shown in FIG. 3B. Routine 160 then proceeds to termination branch 190. In stage 192 of branch 190, agent 62 of the requesting client 60 selects the candidate server 30 with the lowest bid, as provided by that server's
25 broker 32 in stage 172. The broker 32 of the selected server 30 sends distribution/connection information to the requesting client 60 to facilitate distributed execution of the subject task. In the case of the previously described NT 4.0 embodiment, distribution may be established by returning an interface (proxy) pointer back to the interrogating agent 62 of the requesting client 60.
30 Branch 190 proceeds from stage 192 to stage 194 to adjust the bid estimate for the selected server 30 upward in correspondence with its increased processing load as a result of being designated for task execution. In stage 196, estimates for other

servers 30 may be adjusted downward if a predefined time interval has passed since the last interrogation for bids from such servers 30. Other factors may alternatively or additionally be utilized to adjust bid estimates as would occur to those skilled in the art. Also, it should be appreciated that adjustments of bid
5 estimates can be performed more than once during routine 160 or in a different sequence as would occur to those skilled in the art. In one alternative embodiment, bid estimate adjustments of one or more servers 30 are made as part of stage 166 in addition or as an alternative to stage 196. After stage 196, routine 160 halts.

Referring additionally to FIG. 4, server-side distribution/load balancing
10 routine 230 is illustrated. Routine 230 is performed as part of the server selection operation of stage 84 shown in FIG. 2. It should be appreciated that routine 230 operates as part of each broker 32 and selectively communicates with interrogating agents 62 executing routine 160. Routine 230 begins with conditional 232. Conditional 232 tests whether the desired task is available in response to an
15 interrogation by an agent 62 in stage 172 of routine 160. If the task is not available, the broker 32 reports to the interrogating agent 62 in stage 234 of routine 230. In response, conditional 174 of routine 160 is negative, leading to stage 175. If the task is available as tested by conditional 232, routine 230 proceeds to stage 236 to calculate processing load for the corresponding server 30. The processing
20 load calculation may include a number of factors such as the available percentage of idle time for the corresponding processor 29, an average load level indication for the processor 29, and/or a relative processing index or capability of the given computer 24. Also, specifics relative to the performance of the desired task may be utilized in this determination. For example, some tasks may require the use of
25 certain unique hardware resources that will result in a disproportionate degree of loading compared to other computers and/or server arrangements. Indeed, there are numerous factors that may be considered in calculating processing load.

Once processing load is determined in stage 236, control flows to conditional 238 to compare the processing load information to a threshold value
30 illustrated as variable "THRESHOLD". The value of THRESHOLD corresponds to a lightly loaded or minimally loaded processor 29 for which there is a high probability that the bid of the responding broker 32 will be the lowest.

Accordingly, when conditional 238 is satisfied, routine 230 proceeds to stage 240 to initiate the task and report corresponding distribution/connection information necessary to utilize the task at the same time that the bid corresponding to processing load information is provided in subsequent stage 242. For the NT 4.0
5 operating system embodiment, the distribution connection information may be an interface pointer (proxy) for an instance of the requested class object running on the responding server 30. As a result, if the given server is selected, then fewer communication exchanges will be necessary between the interrogating agent 62 and broker 32 to establish task distribution; thereby reducing the overhead imposed
10 by the distribution procedure 70. If the server 30 for the broker 32 responding according to stage 240 of routine 230 is not selected in stage 192 of routine 160, the distribution/connection information is disregarded or ignored by the interrogating agent 62.

If conditional 238 is not satisfied, then routine 230 proceeds directly to
15 stage 242 to report processing load information in the form of a bid by the interrogated broker 32. It should be appreciated that when a task is unavailable it may be reported in the form of an infinite or excessively high bid recognized by the interrogating agent 62 as being an unavailable task. Once stage 242 is performed, process 230 halts.

20 While procedure 70 and routines 160, 230 utilize a "lowest bid" protocol to select a server 30 to perform a given task, in other embodiments other protocols may alternatively or additionally be utilized. For example, the processing load information may be directly reported by brokers 32 and/or one or more parameters corresponding to processing load may be reported to interrogating agents 62 by
25 brokers 32. In one alternative, agents 62 include the capability to vary the type of processing load information and/or processing load parameters requested. In another embodiment, the type of processing load information reported may be server, client, and/or task type specific. In still other embodiments, the interrogating agent 62 collects raw data about the performance of each interrogated
30 server 30 through a corresponding broker 32 and partially or completely calculates processing load. Indeed, many techniques to exchange and determine server

processing load information as would occur to those skilled in the art are contemplated.

Returning to FIG. 2, once the server selection operation of stage 84 has been performed by executing routine 160 with the agent 62 of the requesting client 60, and routine 230 with each broker 32 for the candidate servers 30; procedure 70 continues with stage 86. In stage 86, execution of the distributed task with the designated server 30 begins. Stage 86 may include establishing and exchanging task connection information between the requesting client 60 and selected server 30 via network 22. Alternatively, it should be appreciated that some or all of this information may have already been provided to the requesting client 60. Once execution is started, procedure 70 then halts.

Referring next to FIG. 5, start-up and maintenance routine 320 for system 20 is illustrated. Routine 320 is directed to the establishment and maintenance of information necessary for brokers 32 to communicate with each other and for active agents 62 to identify active brokers 32. Stage 322 corresponds to an initial start-up or activation of one of brokers 32, one of agents 62, or both. Each broker 32 and agent 62 maintains a prioritized broker list in the memory 29 of one of computers 24 corresponding to the broker 32 or agent 62 being activated. In stage 322, the starting broker 32 or agent 62 interrogates the first entry on the list across network 22. If contact cannot be established with the first entry, subsequent entries are contacted until one is found. Once communication with a listed entry is established, the broker list is updated with list entries of the contacted broker 32 that are indicated to be active. This active broker list is then saved by the contacting broker 32 or agent 62. Routine 320 then proceeds to conditional 324.

In conditional 324, if the contacting entity is an agent 62, routine 320 proceeds to branch 360. On the other hand, if the contacting entity is not an agent 62, it defaults to a broker 32 designation according to conditional 324 and routine 320 proceeds to branch 330. Branch 330 starts with stage 332. In stage 332, starting broker 32 continues to contact each new active broker 32 entered in its updated list, and the active entries of each contacted broker 32 are added to the active broker list of the contacting broker 32, until all its entries have been contacted. If an entry fails to respond within a specified time-out period during

stage 332, such nonresponding brokers 32 are regarded as inactive. With the conclusion of stage 332, all active brokers 32 are listed with the starting broker 32, and the start-up phase is complete. After the start-up phase, each broker 32 enters a maintenance loop starting with stage 334. In stage 334, the active broker 32 re-
5 establishes contact with each member of its active broker list from time to time, and modifies the list as needed. After stage 334, conditional 338 is encountered to determine whether the active broker 32 is to be terminated. If not, control returns to stage 334, closing the maintenance loop to once again establish contact with each member of the active broker list. The time interval between contacts is
10 selected such that its impact is negligible on the performance of system 20, and may be uniform or variable. If the subject broker 32 has been terminated as tested in conditional 336, then routine 320 halts.

Returning to branch 360, in the case of an active agent 62, its broker list is periodically updated as it executes routine 160 in stage 362. This update may not
15 occur with every execution, but may be timed in a manner to provide negligible impact to overall performance of system 20. In stage 364 of segment 360, there is an interrogation from time to time of servers 30 that were previously indicated to lack a given task. This interrogation is performed to determine if that task has since been loaded on the given server. The time interval between actions in stages
20 362 and 364 should be relatively infrequent to provide negligible impact on the performance of system 20, and may be uniform or variable. After stage 364, conditional 366 tests whether the active agent 62 is to terminate. If not, routine 320 loops back to stage 362 and subsequently stage 364 until the given agent 62 is terminated. The return loops provided from conditionals 336, 366 each function as
25 background routines for active brokers 32 and agents 62, respectively. Although only described in terms of a particular broker 32 and agent 62, it should be understood that routine 320 is performed for each active broker 32 and agent 62. Routine 320 is invoked upon start-up or activation of a corresponding broker 32 or agent 62 and continues to maintain the broker list of each entity as well as other
30 data in the case of agents 62.

For the NT 4.0 operating system embodiments earlier described, the agents 62 and brokers 32 can be provided on each participating computer 24 in the form

of a distributor module with two logical classes, a first class corresponding to agents 62 and a second class corresponding to brokers 32. Indeed, either of these classes may involve one or more different objects. Furthermore, the distributor module may be provided as an NT service to facilitate the start-up aspects of routine 320. Naturally, in other embodiments, a different implementation may be utilized regardless of operating system type. Indeed, in one alternative embodiment, at least a portion of the agent 62 and/or the broker 32 is provided in the form of hardware.

It should be recognized that the allocation of the process distribution/load balancing function among all servers 30 and clients 60, provides a robust mechanism to continue distribution, even if several servers 30 and corresponding brokers 32 unexpectedly go off-line. Such off-line brokers 32 can be detected by agents 62 or other entities based on the failure of the off-line broker 32 to send a response when expected. Off-line brokers 32 may be reported to other agents 62 and brokers 32 once detected, or each agent 62 may be left to make the determination independently. When one or more brokers 32 go off-line, the remaining brokers 32 can continue to participate in routines 160, 230 even if they are left as two or more disjoint groups by the disorderly loss of the off-line brokers 32. Indeed, the re-establishment of contact with other brokers in stage 334 of routine 320 from time to time eventually rejoins active broker groups that have become separated. As a result, recovery from a catastrophic loss of several brokers 32 all at once is possible.

Many alternative embodiments are envisioned for the present invention. In one example, a technique to migrate remote server execution of a distributed task is included. In other alternatives, the general availability of all tasks with all servers excludes the need to test whether a given task is available and routines 160, 230, and 320 may be modified accordingly. In still other alternatives, the particular sequence and timing of various background loops and checks may be adjusted depending on processing load or through other dynamic means as would occur to those skilled in the art. Indeed, the various operations, stages, conditionals, procedures, and routines of the present invention may be rearranged,

substituted, combined, deleted, reordered, or otherwise modified as would occur to those skilled in the art.

In one embodiment, a method according to the present invention includes operating a computer system including several clients and several servers;
5 generating an activation request with one of the clients for a software task that is being designated for distributed server execution; selecting two or more servers in response to the request based on data maintained about the servers by the one of the clients; interrogating the two or more servers with the one of the clients; providing the one of the clients server information corresponding to processing
10 load for each of the two or more servers in response to this interrogation; and selecting one of the servers to execute the task as a function of the server information.

In another embodiment, a system of network computers includes a number of servers and a number of clients. The clients are each provided with a server
15 selector and the servers are each provided with a server reporter. The server selector of any of the clients is operable to obtain server processing information from the server reporter of any of the servers. The server selector is responsive to an activation request for a software object designated for server execution by a corresponding one of the clients to select two or more servers to interrogate for the
20 server processing load information based on data maintained by the server selector about the servers and designate one of the servers to execute the software object.

In a further embodiment, a computer system is operated including a number of clients and a number of servers. An activation request is generated with one of the clients for a software object to be executed by one of the servers and the
25 servers are interrogated with the one of the clients in response to the request. An object pointer and server processing load information is provided to the one of the clients from one of the servers in response to this interrogation.

In yet a further embodiment, a computer system including a number of clients and a number of servers is operated and an activation request is generated
30 by one of the clients for a routine to be executed by one of the servers. The servers are interrogated by this client in response to the request, and provided information from each of the servers in response. This information corresponds to availability

of the routine. One of the servers is selected to execute the routine in accordance with the information.

All publications, patents, and patent applications cited in this specification are herein incorporated by reference as if each individual publication, patent, or
5 patent application were specifically and individually indicated to be incorporated by reference and set forth in its entirety herein. While the invention has been illustrated and described in detail in the drawings and foregoing description, the same is to be considered as illustrative and not restrictive in character, it being understood that only the preferred embodiment has been shown and described and
10 that all changes, equivalents, and modifications that come within the spirit of the inventions defined by following claims are desired to be protected.

What is claimed is:

1. A method, comprising:
 - operating a computer system including several clients and several servers;
 - 5 generating an activation request with a first one of the clients for a software task, the software task being designated for distributed server execution;
 - selecting two or more servers in response to the request based on data maintained about the servers by the first one of the clients;
 - interrogating the two or more servers with the first one of the clients;
 - 10 providing the first one of the clients information corresponding to processing load for each of the two or more servers in response to said interrogating; and
 - selecting one of the servers to execute the task as a function of the information.
- 15 2. The method of claim 1, wherein the data corresponds to one or more prior server interrogations.
3. The method of claim 1, further comprising:
 - 20 requesting server execution of a software object with a second one of the clients;
 - selecting a subset of the servers with the second one of the clients based on server data maintained by the second one of the clients, the server data being determined from one or more prior server interrogation responses;
 - 25 providing the second one of the clients a number of server characteristics each corresponding to server processing load for a different member of the subset in response to said requesting; and
 - designating one member of the subset to execute the software object in accordance with the server characteristics.

4. The method of claim 1, wherein the computer system includes a host computer with a telephone switch coupled to a telephone network, and further comprising:
- controlling the telephone switch with the host computer; and
- 5 directing a telephone call to a telephone associated with one of the clients during said controlling.
5. The method of claim 1, further comprising measuring response time of each of the number of servers to said interrogating to at least partially determine
- 10 processing load.
6. The method of claim 1, wherein the data maintained by the first one of the clients includes a plurality of estimates each corresponding to processing load for a different server.
- 15
7. The method of claim 6, further comprising adjusting one of the estimates corresponding to the one of the servers selected to execute the task to reflect a processing load increase.
- 20
8. The method of claim 6, further comprising adjusting one of the estimates to reflect a decrease in processing load in response to passage of a selected time interval determined relative to a most recent interrogation of a corresponding one of the servers.
- 25
9. The method of claim 1, further comprising:
- executing an application program with the first one of the clients, the software task being an object belonging to an object class selected for distributed processing among the servers;
- intercepting the activation request from a nominal processing flow;
- 30 directing the activation request to a distribution routine of the first one of the clients; and

wherein the information corresponds to a number of current bids to execute the object, the current bids each being from a different one of the two or more servers, the data corresponds to bid estimates for the servers, and the two or more servers selected for said interrogating correspond to the lowest bid estimates.

5

10. A system of networked computers, comprising:

a number of servers; and

a number of clients, the clients each being provided with a server selector and the servers each being provided with a server reporter, the server selector of

10 any of the clients being operable to obtain server processing load information from the server reporter of any of the servers; and

wherein the server selector is responsive to a client activation request for a software object designated for server execution by a corresponding one of the clients to select two or more servers to interrogate for the server processing load
15 information based on data maintained by the server selector about the servers and designate one of the servers to execute the software object.

11. The system of claim 10, wherein the server reporter of each of the servers includes means for bidding on the execution of the software object.

20

12. The system of claim 10, wherein the server selector of each of the clients includes means for generating a number of estimates each representative of server processing load and each corresponding to a different one of the servers.

25 13. The system of claim 10, wherein the server selector of each of the clients is operable to maintain the data to indicate which of the servers has the capability to execute the software object.

14. The system of claim 10, wherein the server selector of each of the clients is
30 operable to base the data on one or more prior interrogations of one or more of the servers.

15. The system of claim 10, wherein the clients and the servers each reside on a different one of the networked computers, and at least one of the networked computers includes a telephone switch coupled to a telephone network and a telephone associated with one of the clients.

5

16. The system of claim 10, wherein the server reporter is operable to return a pointer to an instance of the software object if a corresponding one of the servers has a processing load below a predefined threshold.

10

17. A method, comprising:

operating a computer system including a number of clients and a number of servers;

generating an activation request with a first one of the clients for a software object to be executed by one of the servers;

15

interrogating the servers with the first one of the clients in response to the request;

providing an object pointer and server processing load information to the first one of the clients from a first one of the servers in response to said interrogating.

20

18. The method of claim 17, further comprising:

providing another object pointer and further server processing load information to the first one of the clients from a second one of the servers in response to said interrogating; and

25

selecting one of the first and second one of the servers to execute an instance of said software object based on the server processing load information from the first one of the servers and the further server processing load information from the second one of the servers.

30

19. The method of claim 17, further comprising maintaining data at the first one of the clients corresponding to one or more prior interrogations of the servers.

20. The method of claim 17, wherein the object pointer is provided to the first one of the clients by the first one of the servers only if the server processing load information indicates a processing load level for the first one of the servers that is below a predefined threshold.

5

21. The method of claim 17, further comprising:
generating a number of estimates each relating to a processing load of a different one of the servers;

upwardly adjusting one of the estimates in response to selection of a
10 corresponding one of the servers to execute an instance of the software object; and
downwardly adjusting another of the estimates after a predetermined time interval passes without interrogating another of the servers.

22. The method of claim 17, further comprising maintaining data at the first
15 one of the clients relating to capability of the servers to generate an instance of the software object.

23. The method of claim 17, further comprising:
executing an application program with a second one of the clients, said
20 executing including generating a different activation request for a different software object for server execution;
intercepting the different activation request from a nominal processing flow;
directing the different activation request to a distribution routine of the
25 second one of the clients;
selecting a subset of the servers with the second one of the clients based on server data maintained by the second one of the clients, the server data being determined from one or more prior server interrogation responses;
providing the second one of the clients a number of server characteristics
30 each corresponding to server processing load for a different member of the subset;
and

designating one member of the subset to execute the different software object in accordance with the server characteristics.

24. The method of claim 17, further comprising:

- 5 determining the server processing load information corresponds to a processing load level below a predefined threshold with the first one of the servers; and
returning the object pointer in response to said determining.

10 25. The method of claim 17, further comprising operating a call center with the computer system.

26. A method, comprising:

- 15 operating a computer system including a number of clients and a number of servers;
generating an activation request with a first one of the clients for a routine to be executed by one of the servers;
interrogating each of the servers with the first one of the clients in response to the request;
20 providing the first one of the clients information from each of the servers in response to said interrogating, the information corresponding to availability of the routine; and
selecting a first one of the servers to execute the routine with the first one of the clients in accordance with the information from each of the servers.

25

27. The method of claim 26, wherein the information further corresponds to server processing load for each of the servers.

28. The method of claim 26, wherein the information corresponds to a bid by
30 each of the servers to execute the routine, the bid being proportional to processing load for each of the servers.

29. The method of claim 26, further comprising maintaining data at the first one of the clients corresponding to one or more prior interrogations of the servers.

30. The method of claim 26, further comprising generating a number of
5 estimates each relating to a processing load of a different one of the servers.

31. The method of claim 26, wherein the computer system includes a telephone switch coupled to a telephone network and further comprising operating a call center with the computer system.

10

32. The method of claim 31, further comprising:
executing an application program with a second one of the clients, said
executing including generating a different request for a different routine for server
execution;

15

intercepting the different request from a nominal processing flow;

directing the different request to a distribution routine of the second one of
the clients;


20

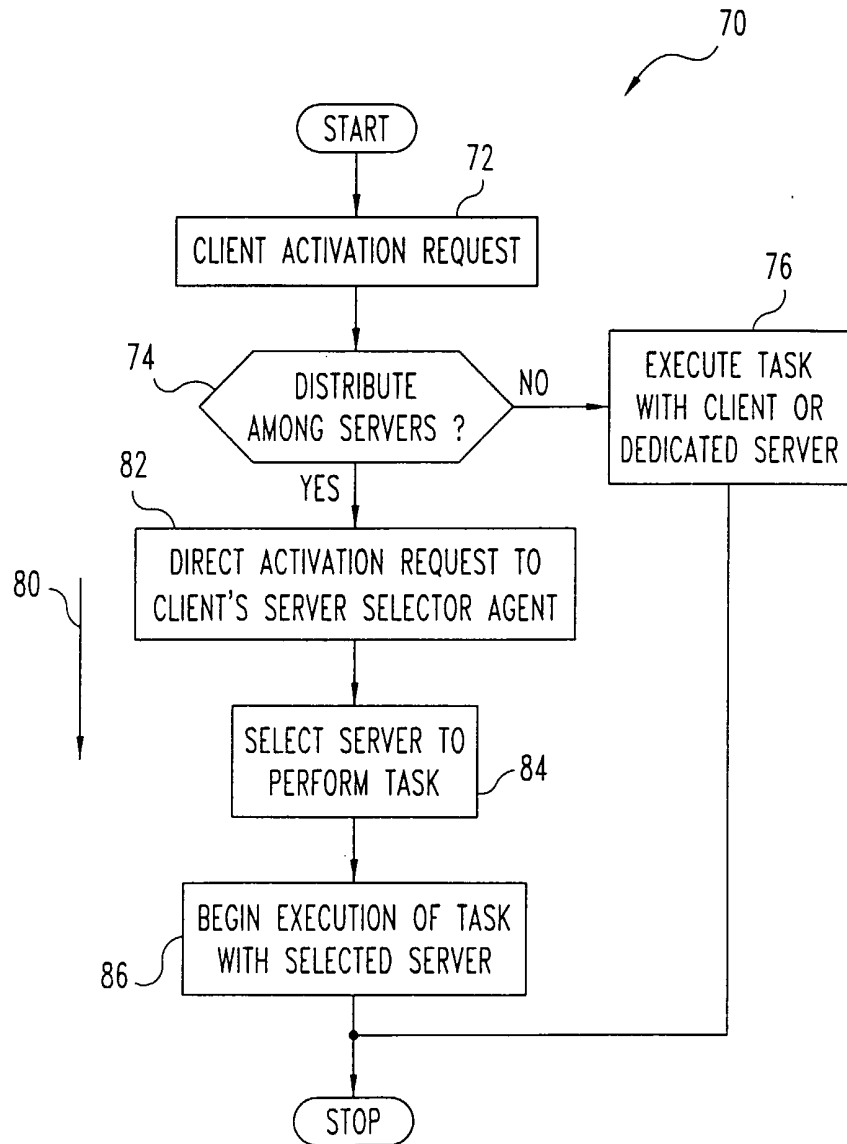
selecting a subset of the servers with the second one of the clients based on
server data maintained by the second one of the clients, the server data being
determined from one or more prior server interrogation responses;

providing the second one of the clients a number of server characteristics
each corresponding to server processing load for a different member of the subset;
and

25

designating one member of the subset to execute the different routine in
accordance with the server characteristics.

**Fig. 1**

**Fig. 2**

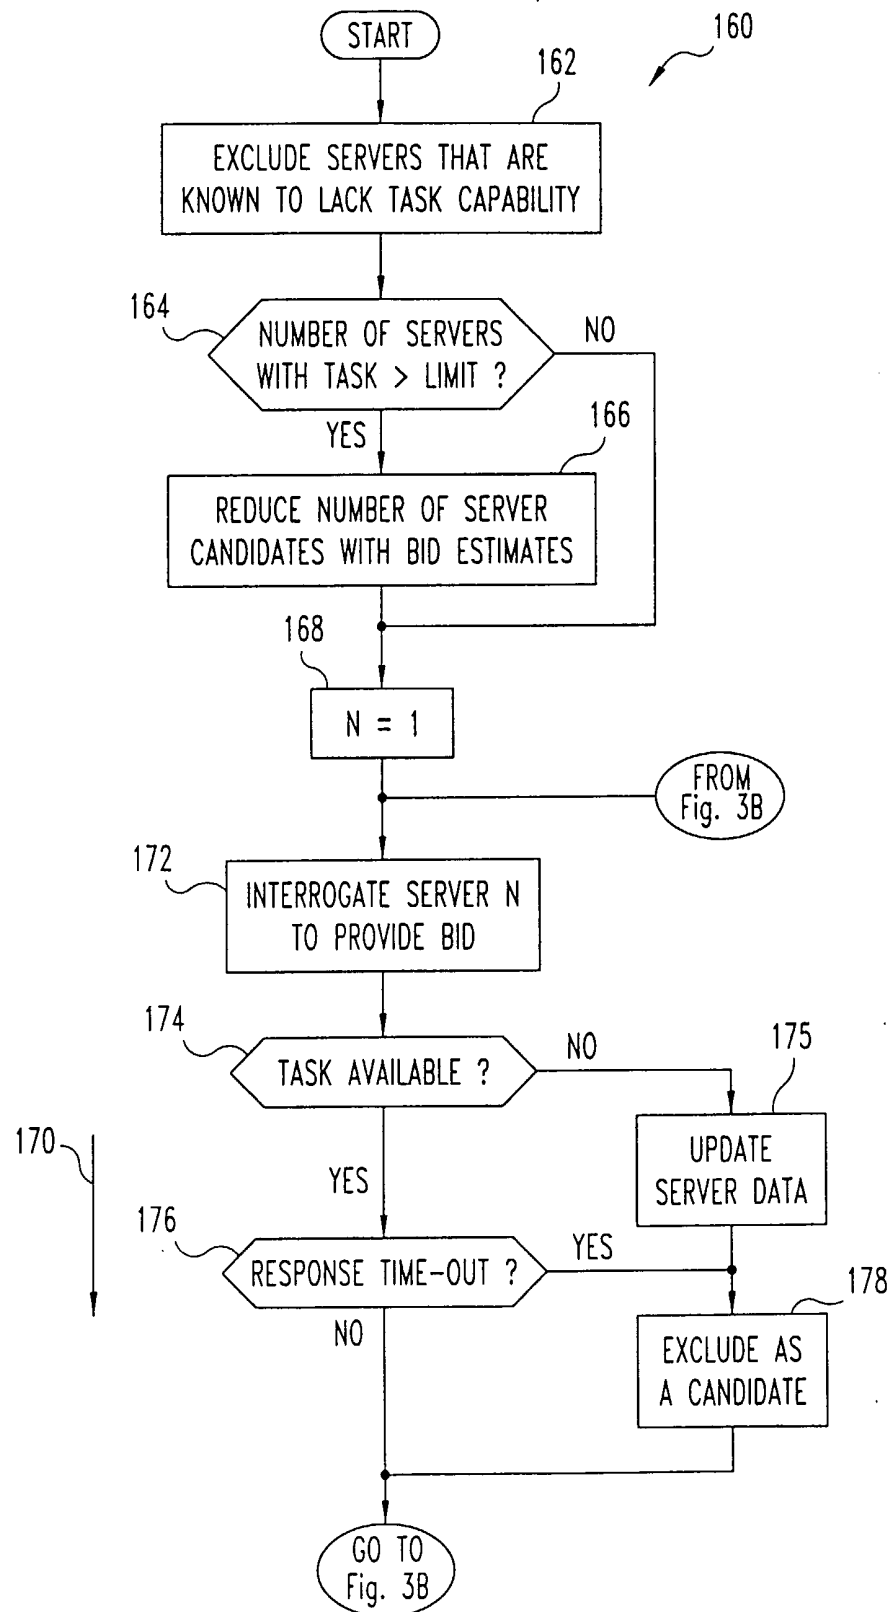
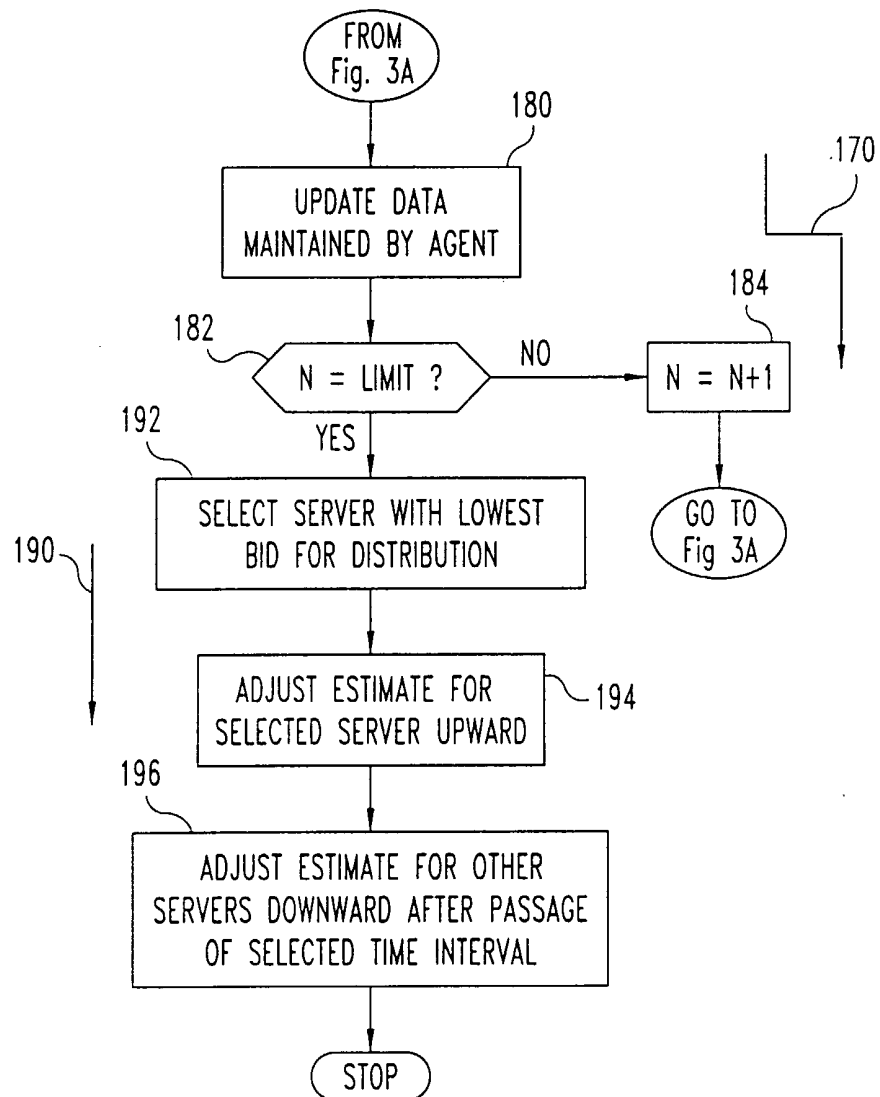
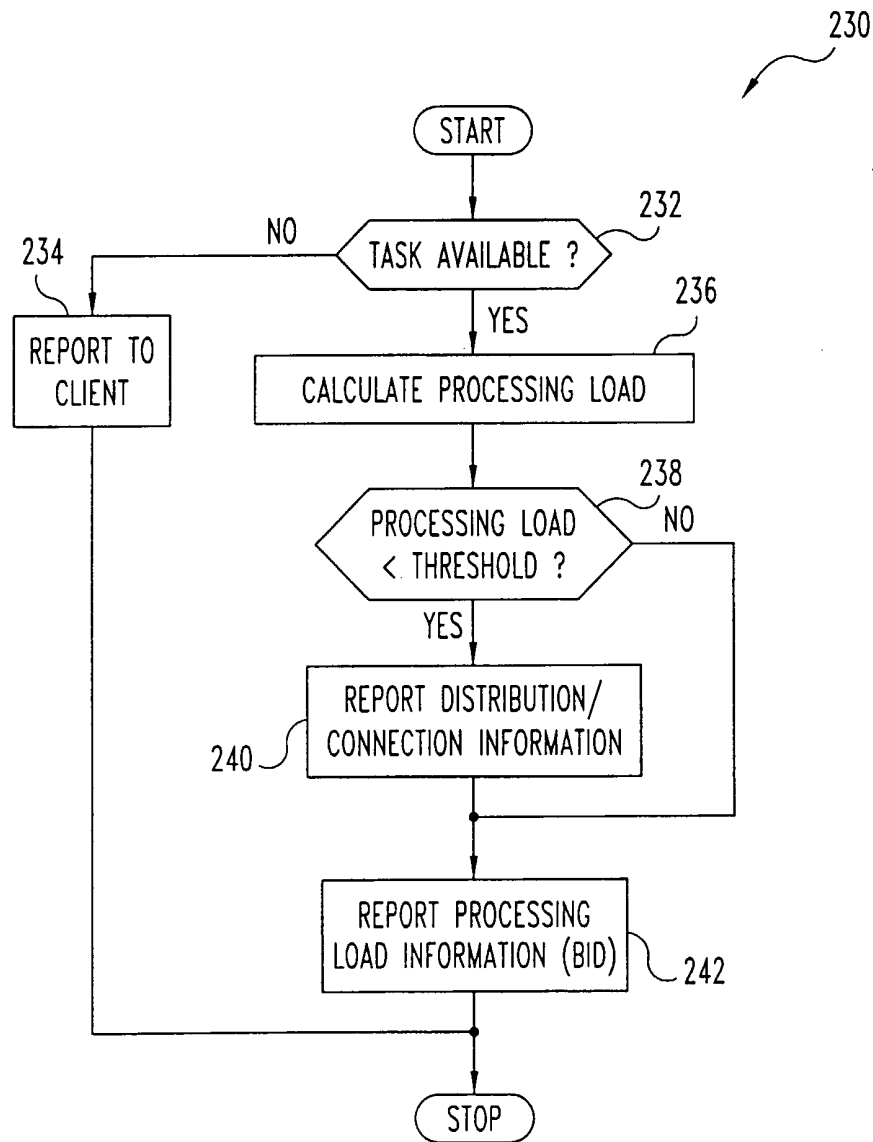
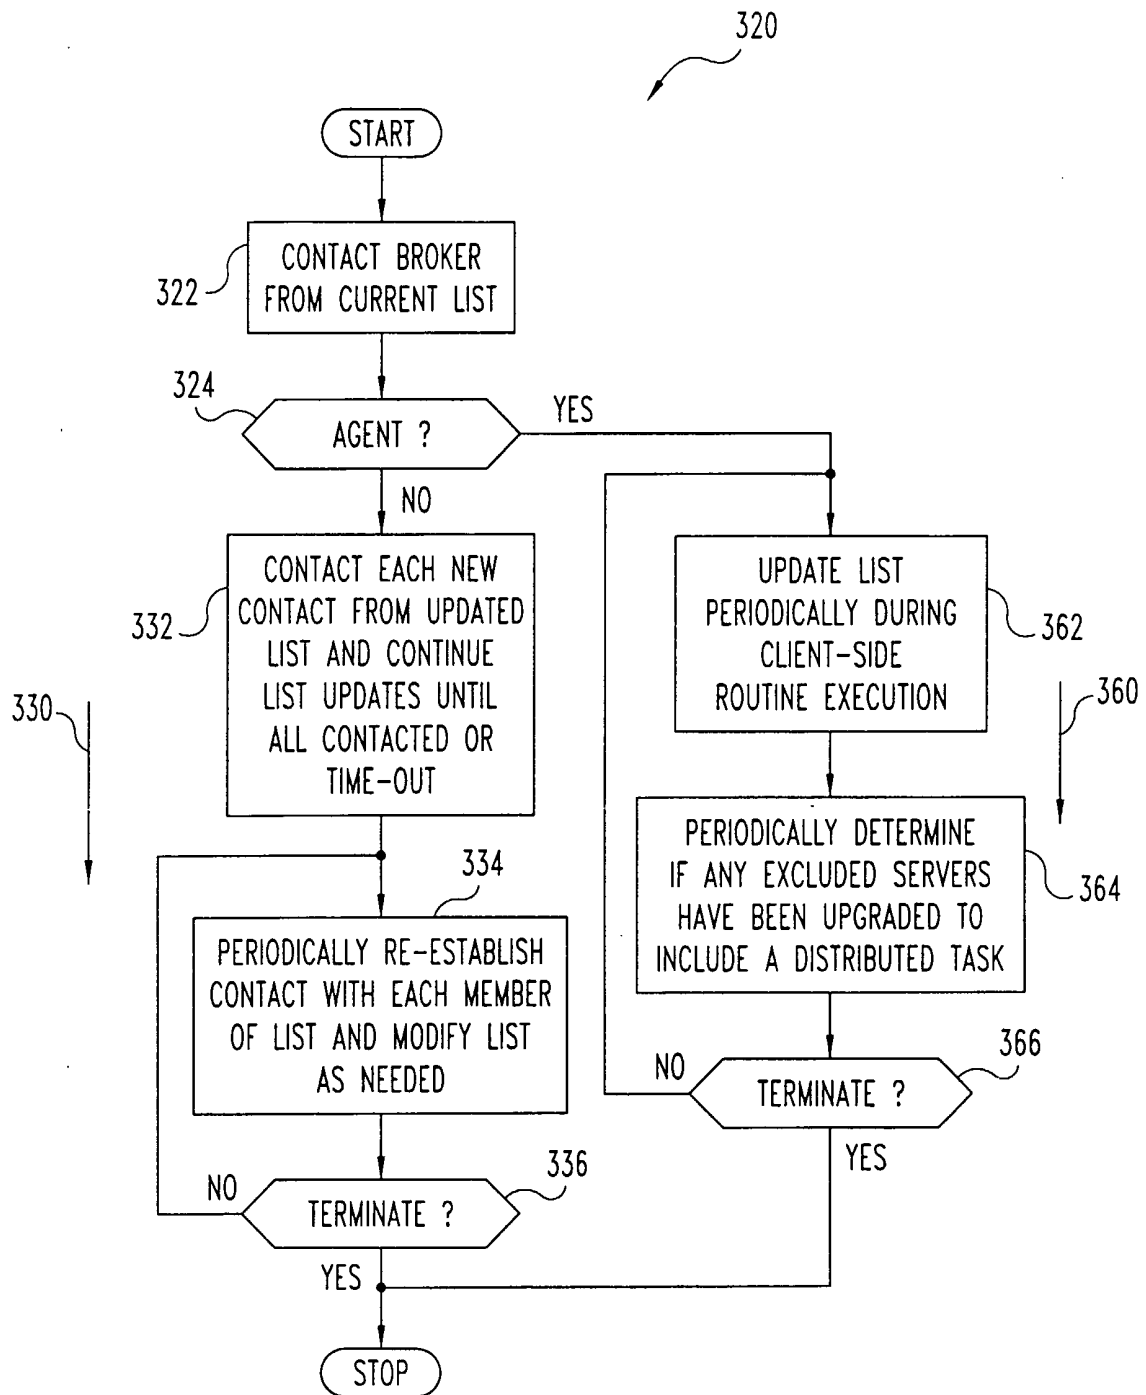


Fig. 3A

**Fig. 3B**

**Fig. 4**

**Fig. 5**

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US01/04734

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F/9/00

US CL : 709/201, 202, 203, 223, 226, 229; 714/15; 395/610

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)

U.S. : 709/201, 202, 203, 223, 226, 229; 714/15; 395/610

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)

East search terms: processing load, load balancing, multiple servers and clients and response time.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|---|-----------------------|
| Y | US 5,774,668 A (CHOQUIER et al.) 30 June 1998, abstract, figs. 1, 5, 7, col. 2 lines 14-65, col. 8 lines 14-35, col. 12 lines 55-col. 13 lines 3. | 1-32 |
| Y | US 5,951,694 A (CHOQUIER et al.) 14 September 1999, abstract, figs 1,5, col. 1 lines 45-col. 2 lines 57, col. 25 lines 49-col. 26 lines 23. | 1-32 |
| Y,P | US 6,167,427 A (RABINOVICH et al.) 26 December 2000, abstract, col. 3 lines 9-45 and col. 4 lines 7-53. | 1-32 |

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

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

Date of the actual completion of the international search

20 MARCH 2001

Date of mailing of the international search report

26 APR 2001

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

AYAZ R. SHEIKH

Telephone No. (703) 305-9648