

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 September 2003 (12.09.2003)

PCT

(10) International Publication Number
WO 03/075152 A1

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

(21) International Application Number: PCT/US03/06177

(22) International Filing Date: 27 February 2003 (27.02.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/087,055 1 March 2002 (01.03.2002) US

(71) Applicant: **VERITY, INC.** [US/US]; 894 Ross Drive,
Sunnyvale, CA 94089 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, 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, OM, PH, PL, PT, RO, RU, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.

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

Published:

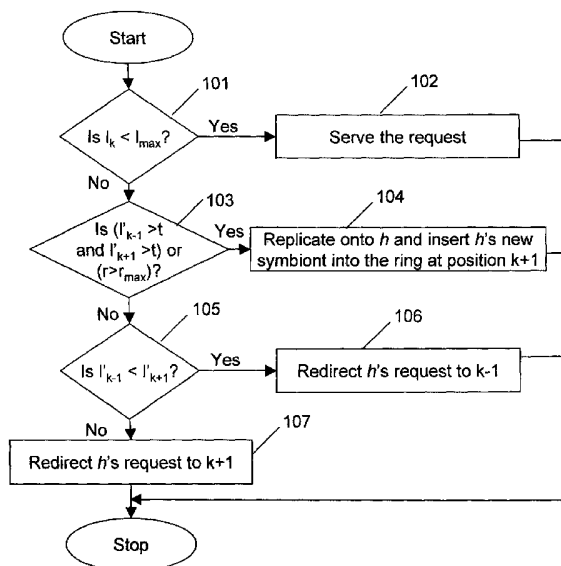
— with international search report

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

(72) Inventor: **CHOO, Kiam**; 500 W. Middlefield Rd. #134,
Mountain View, CA 94043 (US).

(74) Agent: **BOTJER, William, L.**; P.O. Box 478, Center
Moriches, NY 11934 (US).

(54) Title: AUTOMATIC NETWORK LOAD BALANCING USING



(57) **Abstract:** The present invention provides a method, system and computer program to balance the computational and network load in networked computers using self-replicating programs, referred to as symbionts. The method presented here reduces hostspots by encapsulating a resource in a symbiont, and having a user access that symbiont through programs that host symbionts, referred to as hosts. When a host accesses a symbiont, it may replicate a copy of that symbiont resource on itself (104) or may be redirected to some other replicate of the same symbiont (106, 107). The host then offers the replicated resource on the network to alleviate the load experienced by the original symbiont's computer. If the load on a symbiont falls below a threshold, it is removed from the host on which it was hosted (102).



WO 03/075152 A1

AUTOMATIC NETWORK LOAD BALANCING USING SELF-REPLICATING RESOURCES

BACKGROUND

Field of the Invention

5 The present invention relates to load balancing in a computer network, and deals more particularly with a method, system and computer program for load balancing of network traffic, computation and data resources through the use of replicating programs.

Description of the Related Art

10 Networked computer systems are rapidly growing as the means for storage and exchange of information. These days, a large number of resources are available on computer networks; these resources exist at the hardware, software and at networking levels. For example, at the hardware level, these resources usually include disk space, Random Access Memory, and computational power, whereas at
15 the software level, these may include compilers and/or databases. One fundamental advantage of networking computers together is that one computer (or a user) can often access and use the resources of another. However, if a large number of users access any one of these resources simultaneously, there would be a sharp increase in network traffic, which in turn would result in the slowing down the entire network.
20 Furthermore, if a resource is accessed from many computers at the same time, then such an overload may slow down the computer encapsulating that resource, to the extent of even essentially shutting it down. This would especially be true when the computer contains a very popular resource and when other computers access this resource frequently. When a computer or a node on a network thus becomes
25 overloaded, it is commonly referred to as a "hot-spot". Thus, a need arises to balance various loads on the network so that overloading of computers is avoided and the number of hot spots is reduced.

 One method that addresses this problem deploys a powerful central server, that is, a server that has a powerful Central Processing Unit (CPU) and large
30 memory space. However, confining distributed information to servers ignores the

fact that substantial processing and storage power may be available on many smaller computers and these computers may constitute a majority of nodes in the network. Further, this method has a drawback in that the central server is incapable of meeting sudden upsurges in demands and the entire system is not easily
5 scaleable. Finally, this has a major disadvantage in that this server may act as a single point of failure, and failure of this server may render the entire network essentially incapable of accessing all resources that reside on this computer. Therefore, there is a need to effectively balance the resources within a computer network in such a manner that the resources are easily and effectively accessed by
10 all authorized users on the network, while at the same time ensuring that there is no hindrance to the performance and functioning of any computer on the network.

There exist various methods for reduction of hotspots on a network. In one such method, multiple servers may offer identical resources and the client may be connected to any of the multiple servers in order to satisfy the client's request. This
15 method involves replication of popular resources (including data or computational services) on several other nodes of the network. However, this would typically involve an increase in hardware requirements and may even require additional servers. Further, the replication of resources from one server to another usually requires manual supervision. Moreover, this method is not dynamic in nature;
20 indeed, if there is a sudden upsurge in demand, this method will not be able to replicate such resources automatically. Finally, in this method, even if a given resource is not accessed for a long time, it may still continue to consume precious storage space on the server or use its computational power.

Another widely used method for reducing hotspots is replication of data using
25 "ftp mirrors". The "File-Transfer-Protocol (ftp) mirroring" is generally used where the traffic is typically very high and the number of resources is very large. Examples of such networks include large Local Area Networks (LANs), Wide Area Networks (WANs), and of course, the Internet. For example, suppose there is a single server that is located in California, USA and it hosts MP3 files on the Internet. Clearly, such
30 a server would be overloaded by requests from different locations of the world. Moreover, it would be more time consuming to access these resources from a distant location such as Singapore than a nearby location in the USA. Hence, in the

ftp mirroring method, another server – that contains a “replicated image” of the first server -- is deployed to minimize the traffic and reduce access time. However, even in this method, if there is a sudden increase in traffic, then one server does not have the capability to automatically replicate the resources, data and program of the other server (in order to reduce the load of the overloaded server). This is because ftp mirroring requires manual intervention to select “mirror servers” that are appropriate for replication and to select servers that are best suited to download data (by taking into account the incoming traffic and the proximity of server). In addition, it is worth noting that manual supervision is also required to setup these servers; this comprises installation of a server and uploading of resources. Hence, the installation and maintenance of a server proves to be a cumbersome exercise. Moreover, when a resource is not in use for a long time, there is no provision to automatically erase it from server.

Various other methods exist in the literature that are related to load balancing. “Artificial Life Applied to Adaptive Information Agents” Spring Symposium on Information Gathering from Distributed, Heterogeneous Databases, AAAI Press, 1995 by Filippo Menczer, Richard K. Belew and Wolfram Willuhn describes a method that uses agents to retrieve information from a large, distributed collection of documents. When the agents obtain high quality results to their search queries, they replicate. This does not address the issue of load balancing in a network, and is primarily focused on retrieval of relevant documents. “Building Peer-to-Peer Systems With Chord, a Distributed Lookup Service” by Frank Dabek, Emma Brunskill, M. Frans Kaashoek, David Karger, Robert Morris, Ion Stoica and Hari Balakrishnan discloses a method of locating documents while placing few constraints on the applications that use it (<http://pdos.lcs.mit.edu/chord>). This addresses the issue of locating documents in a decentralized network that can be used as a basis for general-purpose peer-to-peer systems. Agoric systems use an economic paradigm to allocate distributed resources according to free market principles. The programs and computers in these systems become buyers and sellers of resources, much like a real-life marketplace and do not explicitly include replication. In consistent hashing, Freenet and other distributed hash systems, users of these systems have little or no control over the kind of data that may come and reside on their computers. These are more like file replication systems rather than systems meant for load balancing.

All of these talk about either load balancing or file replication, but they do not discuss using file or program replication for load balancing.

In addition to the aforementioned means for load balancing, various patents have been granted during the last few years. These are discussed below.

5 US Patent No. 6,279,001 titled "Web Service", and International Patent Application Nos. WO 98/57275 titled "Arrangement for Load Sharing in Computer Networks", WO 00/28713 titled "Internet System and Method for Selecting a Closest Server from a Plurality of Alternative Servers" and WO 01/31445 titled "System and Method for Web Mirroring" disclose and describe selection of a mirror server based
10 on certain heuristics such as availability of a resource, load on the server and geographical proximity of the server to the client. The mirror server has a replicated resource that may be a web page, one or more software programs, media files, or other such items. In all these inventions, the resource replication is performed manually. Replication in these inventions is not dynamic, i.e., even if a mirror server
15 is not accessed very frequently, the replicated resource continues to reside on the same. Conversely, where the resource requirement witnesses an increase, the current methods do not have a provision for automatically replicating the resource onto an appropriate server since the replication is predetermined and it requires manual supervision. In other words, various heuristics given in these inventions do
20 not contain any 'birth' and 'death' rules. Further, all these deal with replication of data only and not computational services.

International Patent Application No. WO 00/14634, titled "Load Balancing for Replicated Services", deals with providing load balancing for replicated services or applications among a plurality of servers. A central server receives request for a
25 service from a client and then directs it to the appropriate server, based upon its operational characteristics (such as its load and its proximity to the client). However, this invention does not undertake the actual replication of services; rather it deals with choosing the most appropriate server for a particular service request.

US Patent No. 5,963,944, titled "System and Method for Distributing and
30 Indexing Computerized Documents Using Independent Agents", uses autonomous agents to manage the distribution of data and to index information among the nodes

of a computer network. Each network node includes a data storage device and an agent interface for execution of autonomous agents. The autonomous agents move independently among different network nodes and for each node they visit, they use the agent interface to execute their functions. However, this invention does not explicitly deal with replication of resources. It uses Balance Agents to break large files into smaller sub-files, and tries to alleviate any overload on any node in the network. Further, the load balancing mechanism is external to the resources i.e. the agents that manage replication are not embedded in the resources that are to be replicated.

Therefore, what is needed is a method and system for effectively balancing the load in a computer network by means of replication of resources, without the need for additional dedicated hardware. Indeed, such a replication should be dynamic, i.e. the resources should be automatically replicated depending upon its current demand; if the demand falls below a predetermined threshold, then such a replicated resource should be removed from the node onto which it had been originally copied. In other words, the replicated resource should 'die' so that various resources (such as computational power, storage space, networking ports and software) of a computer are not unnecessarily used up. Additionally, in order to avoid "single point failures," it is desired that the replication of resources be decentralized.

SUMMARY OF THE INVENTION

An object of the present invention is to provide a method, system and computer program for balancing computational and network loads in a network of computers using self-replicating programs.

Another object of the present invention is to provide a method, system and computer program to reduce the number of computers in a networked environment, that are under heavy usage.

Another object of the present invention is to provide a method, system and computer program that provides a solution for balancing either data or computational services resources, in a network of computers.

Another object of the present invention is to provide a method, system and computer program that manages the replication of a resource, when the need for that resource arises, in a fully automatic and dynamic way, in a network of computers.

5 Another object of the present invention is to provide a method, system and computer program that manages the deletion of a resource from a computer, when its need expires, in a fully automatic and dynamic way, in a network of computers.

10 Another object of the present invention is to provide a method, system and computer program that connects replicates of resources in a manner so as to minimize their frequent replication and deletion from the network of computers.

Still another object of the present invention is to provide a method, system and computer program that provides a self-replicating program (symbiont) that encapsulates the resource.

15 A further object of the present invention is to provide a method, system and computer program that provides a program (host) that provides a suitable living environment for symbionts to function and exposes the network's symbionts to applications on its computer.

20 Yet another object of the present invention is to provide a method, system and computer program that provides for genetic evolution of symbionts wherein each symbiont has a chromosome embedded in it.

25 To achieve the foregoing objects, and in accordance with the purpose of the present invention as broadly described herein, the present invention provides for a method, system and computer program to balance the computational and network load on networked computers using replicating programs. The invention reduces the hotspots by encapsulating a resource in a replicating program called a symbiont. When a host contacts a symbiont on behalf of an application, it may acquire and host a replicate of the resource. Further, when a host contacts a symbiont resource it may be redirected to another copy of the same resource. This redirection and replication, is done by the symbiont using the following algorithm: A host h contacts a symbiont s for a resource. If the symbiont encapsulating the resource is not "too busy", it serves

30

the request. If not, *s* checks out the load on its neighbors and if they are also “too busy”, *s* replicates the resources on to *h*. It also replicates the resource onto *h* if it has been redirected more than a predetermined number of times. In case, any of *s*’s neighbors is not “too busy”, the one with less load serves the request. If *h* acquires a new symbiont, it joins the pool of available copies of the resource by letting some number of symbionts know about its existence. This is done so as to make sure that future requests to *s* are redirected to the symbiont on *h*. Finally, all the symbionts keep checking their own loads at regular “sufficiently large” time intervals. If they find that their load is below a threshold, they “die”.

The present invention will now be described with reference to the following drawings, in which like reference numbers denote the same elements throughout.

BRIEF DESCRIPTION OF THE DRAWINGS

The preferred embodiments of the invention will hereinafter be described in conjunction with the appended drawings provided to illustrate and not to limit the invention, where like designations denote like elements, and in which:

FIG. 1 is a block diagram of a computer workstation environment in which the present invention may be practiced;

FIG. 2 is a diagram of a networked computing environment in which the present invention may be practiced;

FIG. 3 is a diagram showing replicates of a symbiont in a multiply-connected ring; and

FIG. 4 is a flowchart that illustrates the algorithm used by a symbiont when a host contacts it.

DESCRIPTION OF PREFERRED EMBODIMENTS

FIG. 1 illustrates a representative workstation hardware environment in which the present invention may be practiced. The environment of FIG. 1 comprises a representative single user computer workstation 10, such as a personal computer, including related peripheral devices. Workstation 10 includes a microprocessor 12 and a bus 14 employed to connect and enable communication between

microprocessor 12 and the components of workstation 10 in accordance with known techniques. Workstation 10 typically includes a user interface adapter 16, which connects microprocessor 12 via bus 14 to one or more interface devices, such as a keyboard 18, mouse 20, and/or other interface devices 22, which can be any user interface device, such as a touch sensitive screen, digitized entry pad, etc. Bus 14 also connects a display device 24, such as an LCD screen or monitor, to microprocessor 12 via a display adapter 26. Bus 14 also connects microprocessor 12 to memory 28 and long-term storage 30 which can include a hard drive, diskette drive, tape drive, etc.

Workstation 10 communicates via a communications channel 32 with other computers or networks of computers. Workstation 10 may be associated with such other computers in a local area network (LAN) or a wide area network, or workstation 10 can be a client in a client/server arrangement with another computer, etc. All of these configurations, as well as the appropriate communications hardware and software, are known in the art.

FIG. 2 illustrates a data processing network 40 in which the present invention may be practiced. Data processing network 40 includes a plurality of individual networks, including LANs 42 and 44, each of which includes a plurality of individual workstations 10. Alternatively, as those skilled in the art will appreciate, a LAN may comprise a plurality of intelligent workstations coupled to a host processor.

Still referring to FIG. 2, data processing network 40 may also include multiple mainframe computers, such as a mainframe computer 46, which may be preferably coupled to LAN 44 by means of a communications link 48.

Mainframe computer 46 may also be coupled to a storage device 50, which may serve as remote storage for LAN 44. Similarly, LAN 44 may be coupled to a communications link 52 through a subsystem control unit/communication controller 54 and a communications link 56 to a gateway server 58. Gateway server 58 is preferably an individual computer or intelligent workstation that serves to link LAN 42 to LAN 44.

Those skilled in the art will appreciate that mainframe computer 46 may be located a great geographic distance from LAN 44, and similarly, LAN 44 may be located a substantial distance from LAN 42.

Software programming code, which embodies the present invention, is typically accessed by microprocessor 12 of workstation 10 from long-term storage media 30 of some type, such as a CD-ROM drive or hard drive. In a client-server environment, such software programming code may be stored with storage associated with a server. The software programming code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users from the memory or storage of one computer system over a network of some type to other computer systems for use by users of such other systems. Alternatively, the programming code may be embodied in memory 28, and accessed by microprocessor 12 using bus 14. The techniques and methods for embodying software programming code in memory, on physical media, and/or distributing software code via networks are well known and will not be further discussed herein.

The preferred embodiments of the present invention will now be discussed with reference to FIGS. 3-5. In the preferred embodiments, the present invention is implemented as a computer software program. The software may execute on the user's computer or on a remote computer that may be connected to the user's computer through a LAN or a WAN that is part of a network owned or managed internally to the user's company, or the connection may be made through the Internet using an ISP. What is common to all applicable environments is that the user accesses a public network, such as the Internet, through his computer, thereby accessing the computer software that embodies the invention.

An embodiment of the present invention is hereinafter described in detail. The invention provides a method for balancing the computational and network load on networked computers using replicating programs. The invention balances resources that could be either data or computational services. A resource is something that on receiving a request from a host sends back a reply based on its current state. A data

resource could be a database or a document or an article. A computational service could be a software program that runs on a computer.

The two essential software components in this invention are symbiont and host. A symbiont is a software program that replicates and dies based on certain birthing and death rules. These rules could either be hardwired into the system or could be specified when the system is being installed/ used. These rules are formulated so that as soon as a computer on the network is overloaded (according to some threshold), the symbiont takes "birth" on another computer, to share this computer's load. Further, all symbionts keep checking loads on themselves at regular "long enough" time intervals, and if the symbiont experiences load less than some predetermined threshold, it "dies". Moreover, the rules are such that there is not too much "churning" i.e. symbionts do not keep dying and taking birth at a high frequency. These rules could vary depending upon the embodiment of the present invention i.e. there could be several different rule based systems depending upon the embodiment used.

A host is a program that provides a suitable living environment for the symbiont to run i.e. it provides memory, storage, script interpretation, and other services necessary for the symbiont to function i.e. the symbiont runs within the host. In contrast, a symbiont is a self-replicating program that encapsulates a given resource and it does it in a manner that minimizes its frequent replication and deletion (from various computers in a given network).

The host may contain more than one symbiont. The host exposes its symbionts on the computer network as resources that others can use. It also exposes the network's symbiont resources to applications on its computer. It is through this host layer that applications connect to and send messages to symbiont resources on the network.

When the host contacts the symbiont on behalf of an application, it may acquire and host a copy (a replicate) of the resource. Further, when another host contacts the same symbiont resource, it may be redirected to this replicated copy of the same resource.

In the preferred embodiment of the present invention, all the replicates of a particular resource are arranged in the form of a multiply connected ring, by which we mean a graph whose vertices (labeled 0 through $n-1$) are arranged in a circle, with each vertex connected to m neighbors on either side, so that the replicates can communicate with each other. Let us assume that the ring has n replicates of a particular resource. Also, let us assume that each replicate is 'connected' to m other replicates on both sides (i.e. each node is connected to $2m$ other nodes) to make the entire design scalable. If one says that two replicates are 'connected', it means that they can know each other's loads and other characteristics of the nodes. Consider the example network in FIG. 3. The figure illustrates a network with 8 nodes numbered 1 through 8. Each node, in turn is connected to two other nodes on each side. For example, node 6 is connected to nodes 8 and 7 on its left and nodes 4 and 5 on its right. Similarly, node 3 is connected to nodes 4 and 5 on its left and nodes 1 and 2 on its right. This way, each node can keep track of four other nodes. This information will be useful in case any of the nodes wants to redirect a request to any of its neighbors. Also, knowing the loads of only a certain number of neighbors makes the entire design scalable.

Now let us consider a hypothetical situation wherein there is a ring of n nodes with each node connected to just one neighbor on each side. Let the load on the k 'th replicate be l_k . Load here refers to the computational load on the node: the exact way in which it is to be represented depends on the implementation of the system. In first embodiment, the computational load is defined as the number of instructions per second that is executed by a given processor. In an alternate embodiment, the computational load may be defined as the number of requests that are handled by the processor; often, since the processor may take different amount of time to handle to different requests, yet another alternate embodiment may be used where each request has a weight associated with it (which corresponds to the time that will be taken by the processor to service it) and the computational load can be defined as the cumulative sum of the weighted requests that are handled by the processor in one second.

Connecting the replicates in the ring allows replicate k to acquire l_{k-1} and l_{k+1} at regular time intervals, which it stores as the last known loads l'_{k-1} and l'_{k+1} at $k-1$ and

$k+1$. When a host h accesses replicate k , it specifies how many times r it has been redirected. k then runs the algorithm as illustrated as a flow chart in FIG. 4, as follows:

```

    if  $l_k < l_{max}$  at 101 then
5         serve the request 102
    else
        if  $(l'_{k-1} > t \text{ and } l'_{k+1} > t) \text{ or } r > r_{max}$  at 103 then
            replicate on to  $h$  and
            insert  $h$ 's new symbiont into the ring at position  $k+1$  at 104.
10        else
            if  $l'_{k-1} < l'_{k+1}$  at 105 then
                redirect  $h$ 's request to  $k-1$  at 106
            else
                redirect  $h$ 's request to  $k+1$  at 107
15        end if
    end if
end if
end if

```

In the above, $t \leq l_{max}$. In words, the algorithm does the following: when k 's load exceeds the threshold l_{max} , it chooses between replicating the symbiont on h and
20 redirecting h to one of its neighbors. If the last known loads on both k 's left and right neighbors exceed the threshold t , it chooses to replicate symbiont onto h rather than burden its neighbors with an additional request. It also replicates symbiont onto h if h

has already suffered from more than r_{max} redirections. The new symbiont on h then joins the ring as k 's left neighbor, i.e., at position $k+1$.

For services that are deemed essential, the reproductive threshold, l_{max} can be lowered so that its replicates become more abundant.

5 In addition to the above decision made by the symbiont, h ensures that, if it is redirected, subsequent requests are directed at the new target. Once a symbiont has been replicated onto h , it directs future requests at itself. Thus, the load on k is eased. Furthermore, workload has the tendency to diffuse out from busy areas of the ring.

10 Finally, at regular time intervals not triggered by requests, each symbiont checks its own loads. If it is below the threshold l_{min} , it dies, i.e., it makes itself inoperable and ceases to exist, thereafter. This time interval must not be too short or it may lead to churning. Specifically, it should not be comparable to the time scale of the natural fluctuations in load seen by a symbiont. Moreover, one of the replicates
15 of the resource can be encapsulated in a symbiont that is immortal i.e. it never dies. This is important so that even when all the replicates of a particular resource have died, at least one original copy remains. Further, the communication between the replicates can be improved by having some non-local connections between the replicates in the ring.

20 It is worth pointing out that in the preferred embodiment of the present invention, all replicates of a particular resource are arranged in the form of a multiply connected ring, i.e., a graph whose vertices (labeled 0 through $n-1$) are arranged in a circle, with each vertex connected to m neighbors on either side. In an alternative embodiment of the present invention, the replicates can be arranged in the form of a
25 "tree." In a tree, a one vertex (or a replicate) forms the root of the tree, this root is connected to several other vertices (called its children), and each of its children are, in turn, connected to several of their own children, and so on, until the "end children vertices" form the "leaves" of this tree. In yet another alternate embodiment, the vertices (or the replicates) may be all connected to each other, thereby, forming a
30 "complete graph." Indeed, it is easy to create other embodiments wherein the vertices (or the replicates) are connected to each other in any given, specified

manner; such a specification is referred to as a simple graph (in Computer Science and the Mathematics' literature).

5 In another alternative embodiment of the present invention, the host in which the symbiont is residing can also perform some of the functions performed by the symbionts. For example, the host can perform the function of redirection that is presently encapsulated in the symbiont. As another example, the hosts may, for security reasons, have control over what is done by a symbiont that encapsulates a program. In that case, the "program" carried by the symbiont could be relegated to an integer that chooses between a few possible actions, each of which is actually
10 implemented in the host although they might be thought of as computations that have been performed by the symbiont.

In another alternative embodiment of the present invention, genetic evolution of symbionts is possible wherein each symbiont has a "chromosome" embedded in it that is simply a piece of software code which is "distinctive" of that symbiont (just like
15 a chromosome is distinctive of a living thing). The "chromosome" contains certain features of the symbiont that distinguish it from others. Moreover, these "chromosomes" decide the "superiority" of symbionts, i.e. symbionts with "better chromosomes" are considered better. This can be deduced from the access preference of hosts, as well as from the symbiont's performance. Using these
20 "chromosomes", and genetic operations like mutations and crossover, the system can come up with better quality symbionts. Further, if one needs to upgrade a symbiont, one just needs to introduce a higher version symbiont in the symbiont pool (with the heuristics that a higher version symbiont is a better one), so that it can be used from there on (only if it performs better than the previous versions!).

25 In another alternative embodiment of the present invention, the redirection is done on to the replicate that is "closest" (geographically or on the basis of some other user preferences) to the host that has requested for the resource.

In another alternative embodiment of the present invention, heavyweight resources may be broken up so that the smaller units can be run on different
30 computers. A heavyweight resource is a file that is too large or a computation that is too intensive. In these cases, special non-birthing symbionts may be used, as hosts

may refuse to host heavyweight symbionts. Also, the replication of data that is of a proprietary or sensitive nature needs to be carefully controlled.

While the preferred embodiment of the present has been described, additional variations and modifications in that embodiment may occur to those skilled in the art
5 once they learn of the basic inventive concepts. Therefore, it is intended that the appended claims shall be construed to include both the preferred embodiment and all such variations and modifications as fall within the spirit and scope the invention.

1 What is claimed is:

- 2 1. A method for serving requests for resources by applications running on a computer,
3 the computer being part of a network of computers, each computer on said network
4 comprising a host program, each said host program comprising a symbiont, each
5 said symbiont encapsulating one data processing resource, said method comprising
6 the steps of:
 - 7 a. said host receiving a request for said resource from an application running on said
8 host's computer;
 - 9 b. said host contacting said symbiont that encapsulates said resource; and
 - 10 c. said symbiont either serving said request, or redirecting it to another replicate of
11 itself, or replicating itself onto said host.
- 1 2. The method according to claim 1, wherein said host provides information relating to
2 said symbionts available on said network to applications running on said host's
3 computer.
- 1 3. The method according to claim 1, wherein said host provides information relating to
2 said symbionts available on said host's computer to said network.
- 1 4. The method according to claim 1, wherein various replicates of said symbiont is
2 connected together, to support a measure of communication among said replicates.
- 1 5. The method according to claim 4, wherein said various replicates of said symbiont
2 are connected together in a multiply connected ring.
- 1 6. The method according to claim 1, wherein said step of said symbiont either serving
2 said request, or redirecting it to another replicate of itself, or replicating itself onto
3 said host, said step further comprising the steps of:
 - 4 a. determining load on said symbiont, if load on said symbiont is less than its
5 threshold, I_{\max} , said symbiont serving said request;

- 6 b. determining load on said symbiont, if load on said symbiont is more than its
7 threshold, I_{\max} , and if load on all said connected replicates of said symbiont, is
8 also more than their threshold, t , said symbiont replicating itself on said host;
- 9 c. determining load on said symbiont, if load on said symbiont is more than its
10 threshold, I_{\max} , and if said host has been redirected more than a predetermined
11 number of times, said symbiont replicating itself on said host; and
- 12 d. determining load on said symbiont, if load on said symbiont is more than its
13 threshold, I_{\max} , and if at least one of said connected replicates of said symbiont,
14 has a load less than their threshold, t , one of said connected replicates with load
15 less than its threshold serving said request.
- 1 7. The method according to claim 6, wherein said threshold, I_{\max} , of said symbiont,
2 evolves with time according to some probabilistic measure.
- 1 8. The method according to claim 6, wherein said threshold, t , of said replicate of said
2 symbiont is less than said threshold, I_{\max} of said symbiont.
- 1 9. The method according to claim 6, wherein said threshold, t , of said replicate of said
2 symbiont, evolves with time according to some probabilistic measure.
- 1 10. The method according to claim 6, wherein said step of one of said connected
2 replicates with load less than its threshold serving said request, further comprises
3 said replicate with least load serving said request.
- 1 11. The method according to claim 6, wherein said step of one of said connected
2 replicates with load less than its threshold serving said request, further comprises
3 said replicate closest to said host serving said request.
- 1 12. A system for serving requests for resources by applications running on a computer,
2 the computer being part of a network of computers, each computer on said network
3 comprising a host program, each said host comprising a symbiont, each said
4 symbiont encapsulating one data processing resource, said system comprising:

- 5 a. means for said host receiving a request for said resource from an application
6 running on said host's computer;
- 7 b. means for said host contacting said symbiont that encapsulates said resource;
8 and
- 9 c. means for said symbiont handling said request.

1 13.The system according to claim 12, wherein said host provides information relating to
2 said symbionts available on said network to applications running on said host's
3 computer.

1 14.The system according to claim 12, wherein said host provides information relating to
2 said symbionts available on said host's computer to said network.

1 15.The system according to claim 12, wherein said various replicates of said symbiont
2 are connected together, to support some measure of communication among said
3 replicates.

1 16.The system according to claim 15, wherein said various replicates of said symbiont
2 are connected together in a multiply connected ring.

1 17.The system according to claim 12, wherein said means for said symbiont handling
2 said request, further comprises:

- 3 a. means for said symbiont serving said request,
- 4 b. means for said symbiont replicating itself on said host,
- 5 c. means for one of said connected replicates with load less than its threshold
6 serving said request.

1 18.The system according to claim 17, wherein said means for one of said connected
2 replicates with load less than its threshold serving said request, further comprises
3 means for said replicate with least load serving said request.

1 19. The system according to claim 17, wherein said means for one of said connected
2 replicates with load less than its threshold serving said request, further comprises
3 means for said replicate closest to said host serving said request.

1 20. A method for managing hosts and symbionts in a network of computers, each
2 computer on said network comprising a host program, each said host program
3 comprising a symbiont, each said symbiont encapsulating one data processing
4 resource, said method comprising the steps of:

- 5 a. initializing a set of hosts and symbionts on said network;
- 6 b. adding a new symbiont for an existing resource to said network, whenever there is
7 a need for one;
- 8 c. adding a new symbiont for a new resource to said network whenever said new
9 resource is to be added; and
- 10 d. deleting said symbiont from said network of computers whenever certain
11 conditions are met.

1 21. The method according to claim 20, wherein said host provides information relating to
2 said symbionts available on said network to applications running on said host's
3 computer.

1 22. The method according to claim 20, wherein said host provides information relating to
2 said symbionts available on said host's computer to said network.

1 23. The method according to claim 20, wherein various replicates of said symbiont are
2 connected together, to support some measure of communication among said
3 replicates.

1 24. The method according to claim 23, wherein said various replicates of said symbiont
2 are connected together in a multiply connected ring.

1 25. The method according to claim 20, wherein said initializing step further comprises
2 the steps of:

- 3 a. initializing a host on each computer of said network;
- 4 b. encapsulating said resources that are to be initialized in one said symbiont each;
- 5 c. marking original copy of each of said symbiont encapsulating said resource, as
6 immortal so that they are always present in said network; and
- 7 d. initializing said symbionts on computers in said network, wherein said symbiont
8 runs in said host.

1 26. The method according to claim 25, wherein a symbiont run in said host.

1 27. The method according to claim 20, wherein said step of adding a new symbiont for
2 an existing resource to said network, whenever there is a need for one, further
3 comprises the steps of:

- 4 a. determining load on said symbiont, if load on said symbiont is more than its
5 threshold, I_{\max} , and if load on all said connected replicates of said symbiont, is
6 also more than their threshold, t , said symbiont replicating itself on said host;
- 7 b. determining load on said symbiont, if load on said symbiont is more than its
8 threshold, I_{\max} , and if said host has been redirected more than a predetermined
9 number of times, said symbiont replicating itself on said host; and
- 10 c. determining load on said symbiont, in either case, connecting said new symbiont
11 to other said symbionts of said existing resource.

1 28. The method according to claim 27, wherein said threshold, I_{\max} , of said symbiont,
2 evolves with time according to some probabilistic measure.

1 29. The method according to claim 27, wherein said threshold, t , of said replicate of said
2 symbiont is less than said threshold, I_{\max} of said symbiont.

1 30. The method according to claim 27, wherein said threshold, t , of said replicate of said
2 symbiont, evolves with time according to some probabilistic measure.

1 31. The method according to claim 20, wherein said step of adding a new symbiont for a
2 new resource to said network whenever a new resource is to be added, further
3 comprises the steps of:

4 a. encapsulating said new resource to be initialized in a new symbiont;

5 b. marking original copy of said new symbiont encapsulating said new resource, as
6 immortal so that it is always present in said network; and

7 c. initializing said new symbiont on a computer in said network, wherein said new
8 symbiont runs in said host.

1 32. The method according to claim 20, wherein said step of deleting said symbiont from
2 said network of computers whenever certain conditions are met, further comprises
3 the steps of:

4 a. said symbionts checking their loads at regular time intervals; and

5 b. said symbionts dying if their load is less than a threshold, I_{min} .

1 33. The method according to claim 32, wherein said time intervals evolve with time.

1 34. The method according to claim 32, wherein said threshold, I_{min} , evolves with time.

1 35. The method according to claim 32, wherein said symbionts marked immortal are
2 never deleted from said network.

1 36. A system for managing hosts and symbionts in a network of computers, each
2 computer on said network comprising a host, each said host comprising a symbiont,
3 each said symbiont encapsulating one data processing resource, said system
4 comprising:

5 a. means for initializing a set of hosts and symbionts on said network;

- b. means for adding a new symbiont for an existing resource to said network;
- c. means for adding a new symbiont for a new resource to said network; and
- d. means for deleting said symbiont from said network of computers.

37. The system according to claim 36, wherein said host provides information relating to said symbionts available on said network to applications running on said host's computer.

38. The system according to claim 36, wherein said host provides information relating to said symbionts available on said host's computer to said network.

39. The system according to claim 36, wherein various replicates of said symbiont are connected together, to support some measure of communication among said replicates.

40. The system according to claim 39, wherein said various replicates of said symbiont are connected together in a multiply connected ring.

41. The system according to claim 36, wherein said initializing means further comprises:

- a. means for initializing a host on each computer of said network;
- b. means for encapsulating said resources that are to be initialized in one said symbiont each;
- c. means for marking original copy of each of said symbiont encapsulating said resource, as immortal so that they are always present in said network; and
- d. means for initializing said symbionts on computers in said network, wherein said symbiont runs in said host.

42. The system according to claim 41, wherein zero or more symbionts run in said host.

1 43. The system according to claim 36, wherein said means for adding a new symbiont for
2 an existing resource to said network, whenever there is a need for one, further
3 comprises:

4 a. means for said symbiont replicating itself on said host as a new symbiont; and

5 b. means for connecting said new symbiont to other said symbionts of said existing
6 resource.

1 44. The system according to claim 36, wherein said means for adding a new symbiont for
2 a new resource to said network whenever a new resource is to be added, further
3 comprises:

4 a. means for encapsulating said new resource to be initialized in a new symbiont;

5 b. means for marking original copy of said new symbiont encapsulating said new
6 resource, as immortal so that it is always present in said network; and

7 c. means for initializing said new symbiont on a computer in said network, wherein
8 said new symbiont runs in said host.

1 45. The system according to claim 36, wherein said means for deleting said symbiont
2 from said network of computers whenever certain conditions are met, further
3 comprises:

4 a. means for said symbionts checking their loads at regular time intervals; and

5 b. means for said symbionts dying if their load is less than a threshold, I_{\min} .

1 46. The system according to claim 45, wherein said time intervals evolve with time.

1 47. The system according to claim 45, wherein said threshold, I_{\min} , evolves with time.

1 48. The system according to claim 45, wherein said symbionts marked immortal are
2 never deleted from said network.





FIG. 1






FIG. 3

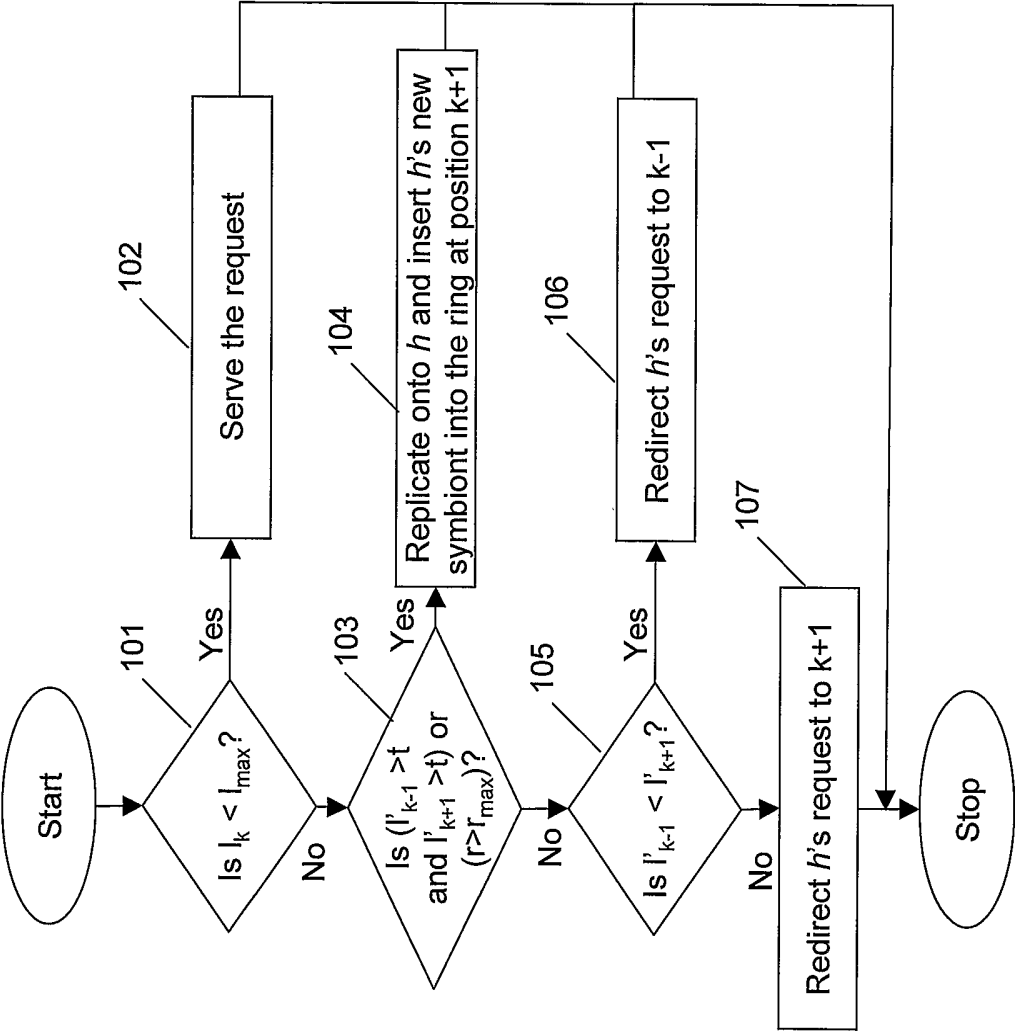


FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US03/06177

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 9/00

US CL : 709/104

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. : 709/103, 105; 711/170; 712/214; 345/804, 805

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 BRS: EAST:

search terms: network load, application resources, encapsulating data, threshold, computer program, agents, routines, symbiont

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y, P	US 6,473,791 B1 (AL-GHOSEIN et al.) 29 October 2002, col. 4, line 30 - col. 6, line 62	1-48
Y	US 2001/0034752 A1 (KREMIEN) 25 October 2001, page 4, par. [0024] - page 5, par. [0040]	1-48
A	US 4,928,252 A (GABBE et al.), 22 May 1990, col. 3, line 5 - col. 4, line 10.	1-48

☐ 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

16 MAY 2003

Date of mailing of the international search report

04 JUN 2003

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

MARK POWELL

Telephone No. (703) 305-9703