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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
13 July 2006 (13.07.2006)

PCT

(10) International Publication Number
WO 2006/072547 A1

(51) International Patent Classification:
G06F 11/34 (2006.01)

(21) International Application Number:
PCT/EP2005/056931

(22) International Filing Date:
19 December 2005 (19.12.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/031,490 6 January 2005 (06.01.2005) US

(71) Applicant *(for all designated States except US)*: **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, NY 10504 (US).

(71) Applicant *(for MG only)*: **IBM UNITED KINGDOM LIMITED** [GB/GB]; PO Box 41 North Harbour, Portsmouth Hampshire PO6 3AU (GB).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **FELLENSTEIN, Craig, William** [US/US]; Six Hunting Ridge Road, Brookfield, Connecticut 06804 (US). **HAMILTON, Rick, Allen, II** [US/US]; 1532 Dairy Road, Charlottesville, Virginia 22903 (US). **JOSEPH, Joshy** [IN/US]; 5812 Cherry Hill Drive, Poughkeepsie, New York 12603 (US). **SEAMAN, James** [US/US]; 2759 Marshall Street, Falls Church, Virginia 22042 (US).

(74) Agent: **WILLIAMS, Julian, David**; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester Hampshire SO21 2JN (GB).

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

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

[Continued on next page]

(54) **Title:** FACILITATING OVERALL GRID ENVIRONMENT MANAGEMENT BY MONITORING AND DISTRIBUTING GRID ACTIVITY

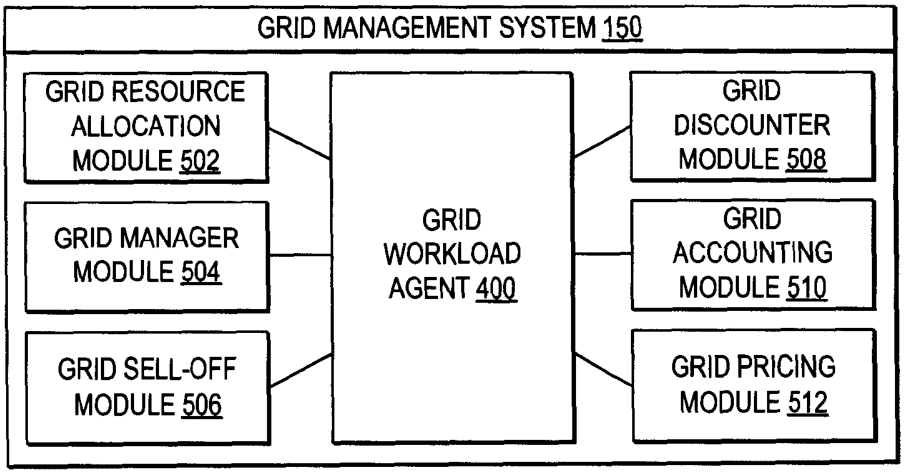

(57) **Abstract:** A method, system, and program for facilitating overall grid environment management by monitoring grid activity across disparate grid resources and distributing grid activity to decisional grid modules (502,504,506,508,510,512) are provided. A grid workload controller within a computational grid environment monitors real-time grid activity at an application level from multiple disparate grid application environments . The grid workload controller (402,408) then determines a selection of grid modules within the computational grid environment that require the real-time grid activity to make decisions about the management of the computational grid environment. The grid workload controller distributes the real-time grid activity to the selection of grid modules, wherein the selection of grid modules then make automated decisions within the grid environment to maintain performance requirements .

WO 2006/072547 A1

Published:
— *with international search report*
— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

FACILITATING OVERALL GRID ENVIRONMENT MANAGEMENT BY MONITORING AND DISTRIBUTING GRID ACTIVITY

BACKGROUND OF THE INVENTION

Technical Field

The present invention relates in general to improved grid computing and in particular to coordinating automated workload performance controllers within a grid computing environment. Still more particularly, the present invention relates to facilitating automated grid workload performance maintenance by multiple decisional grid modules that make decisions based on grid activity gathered from disparate types of grid resource groups.

Description of the Related Art

Ever since the first connection was made between two computer systems, new ways of transferring data, resources, and other information between two computer systems via a connection continue to develop. In typical network architectures, when two computer systems are exchanging data via a connection, one of the computer systems is considered a client sending requests and the other is considered a server processing the requests and returning results. In an effort to increase the speed at which requests are handled, server systems continue to expand in size and speed. Further, in an effort to handle peak periods when multiple requests are arriving every second, server systems are often joined together as a group and requests are distributed among the grouped servers. Multiple methods of grouping servers have developed such as clustering, multi-system shared data (sysplex) environments, and enterprise systems. With a cluster of servers, one server is typically designated to manage distribution of incoming requests and outgoing responses. The other servers typically operate in parallel to handle the distributed requests from clients. Thus, one of multiple servers in a cluster may service a client request without the client detecting that a cluster of servers is processing the request.

Typically, servers or groups of servers operate on a particular network platform, such as Unix or some variation of Unix, and provide a hosting environment for running applications. Each network platform may provide functions ranging from database integration, clustering services, and security to workload management and problem determination. Each network platform typically offers different implementations, semantic behaviors, and application programming interfaces (APIs).

Merely grouping servers together to expand processing power, however, is a limited method of improving efficiency of response times in a network. Thus, increasingly, within a company network, rather than just grouping servers, servers and groups of server systems are organized as distributed resources. There is an increased effort to collaborate, share data, share cycles, and improve other modes of interaction among servers within a company network and outside the company network. Further, there is an increased effort to outsource nonessential elements from one company network to that of a service provider network. Moreover, there is a movement to coordinate resource sharing between resources that are not subject to the same management system, but still address issues of security, policy, payment, and membership. For example, resources on an individual's desktop are not typically subject to the same management system as resources of a company server cluster. Even different administrative groups within a company network may implement distinct management systems.

The problems with decentralizing the resources available from servers and other computing systems operating on different network platforms, located in different regions, with different security protocols and each controlled by a different management system, has led to the development of Grid technologies using open standards for operating a grid environment. Grid environments support the sharing and coordinated use of diverse resources in dynamic, distributed, virtual organizations. A virtual organization is created within a grid environment when a selection of resources, from geographically distributed systems operated by different organizations with differing policies and management systems, is organized to handle a job request.

While clusters or other groups of servers can be grouped within a grid environment, Grid technologies do not solve all the problems to provide communication between groups of resources managed by different management systems with different standards. In particular, a current problem with Grid technology is the limitations of tools and systems that already monitor each group of systems. In particular, a limitation of standard performance monitors is that these monitors group resources according to the type of hardware resource. For example, a first monitoring tool may monitor pSeries machines while a second monitoring tool monitors systems grouped as zSeries machines (pSeries and zSeries are registered trademarks of International Business Machines Corporation). As a result of grouping resources according to hardware resource, these monitoring tools are limited to monitoring results of performance at a hardware levels. In addition, as a result of grouping resources according to hardware resource, these monitoring tools are limited to using the protocols implemented on the hardware resource and therefore typically not supporting communication directly between monitoring tools.

As clusters and other groups of systems are decentralized into grid environments, it would be advantageous to provide for tracking the grid based activity across disparate hardware and software platforms at an application level, rather than just a hardware level, and balancing grid jobs and grid workload across an entire grid environment including hardware, software, and network resources, rather than just a particular hardware environment. Thus, within a grid environment, it would be advantageous to monitor performance and other activity across the entire grid environment and coordinate distribution of that grid activity to modules affected by current grid activity.

SUMMARY OF THE INVENTION

In view of the foregoing, the present invention in general provides for improved grid computing and in particular to coordinating automated workload performance controllers within a grid computing environment. Still more particularly, the present invention relates to coordinating automated grid workload performance by multiple decisional grid modules according to grid activity gathered from disparate grid resource groups.

According to one embodiment, a grid workload controller within a computational grid environment monitors real-time grid activity at an application level from multiple disparate grid application environments. The grid workload controller then determines a selection of grid modules within the computational grid environment that require the real-time grid activity to make decisions about the management of the computational grid environment. The grid workload controller distributes the real-time grid activity to the selection of grid modules, wherein the selection of grid modules then make automated decisions within the grid environment to maintain performance requirements.

The real-time grid activity may include data such as a newly scheduled grid job, a currently executing grid job, a future job schedule, a current pricing schedule, a future pricing schedule, a current resource availability, and a predicted resource availability. In addition, the real-time grid activity may include commands or instructions decided by a decisional grid module. Grid modules may include, for example, a grid resource allocation module, a grid resource manager module, a grid sell-off module, a grid discounter module, a grid accounting module, and a grid pricing module.

In monitoring the real-time grid activity from multiple disparate grid application environments, the grid workload controller may receive grid activity in multiple protocols requiring adaptation to a particular protocol used by the grid workload controller. In addition to adapting the grid activity, the grid workload controller may filter out grid activity based on the type of grid characteristics within the grid activity and based on which grid application environment or module sends the grid activity. In addition, the grid workload controller may receive the grid activity automatically or the grid workload controller may interrogate resource manager modules for the real-time grid activity, where each of the resource manager modules tracks at least one real-time performance characteristic of a grouping of resource nodes allocated for one from among the multiple disparate grid application environments.

In determining the selection of grid modules that require the real-time grid activity, the grid workload controller may filter the grid activity according to distribution rules that designate which types of grid activity may affect which types of grid modules. The distribution rules may be set by a system administrator or an automated decision controller, for example.

Viewed form a first aspect the present invention provides a computer-implemented method for managing workload within a grid environment, comprising: monitoring, at a grid workload controller, real-time grid activity at an application level from a plurality of disparate grid application environments within a computational grid environment; determining a selection of grid modules that require said real-time grid activity to make decisions about the management of said computational grid environment; and distributing, from said grid workload controller, said real-time grid activity to said selection of grid modules, wherein said selection of grid modules make automated decisions within said grid environment to maintain workload performance requirements within said computational grid environment.

Preferably, the present invention provides a method for managing workload, wherein monitoring, at a grid workload controller, real-time grid activity at an application level from a plurality of disparate grid application environments within a computational grid environment, further comprises: monitoring, at said grid workload controller, said real-time grid activity received in a plurality of protocols requiring adaptation to a particular protocol used by said grid workload controller.

Preferably, the present invention provides a computer-implemented method for managing workload, wherein monitoring, at a grid workload controller, real-time grid activity at an application level from a plurality of disparate grid application environments within a computational grid environment, further comprises: receiving, at said grid workload controller, a plurality of types of grid characteristics within said real-time grid activity; and filtering said plurality of types of grid characteristics within said real-time grid activity according to a plurality of monitoring rules, wherein said plurality of monitoring rules specify a selection of said plurality of types of grid characteristics and a selection of said plurality of disparate grid application environments from which said real-time grid activity is stored.

Preferably, the present invention provides a computer-implemented method for managing workload, wherein monitoring, at a grid workload controller, real-time grid activity at an application level from a plurality of disparate grid application environments within a computational grid environment, further comprises: interrogating a plurality of manager modules for said real-time grid activity, wherein each of said plurality of manager modules tracks at least one real-time performance characteristic of a grouping of resource nodes allocated to one of said plurality of disparate grid application environments within said computational grid environment.

Preferably, the present invention provides a computer-implemented method for managing workload, wherein determining a selection of grid modules that require said real-time grid activity to make decisions about the management of said computational grid, further comprises: filtering said real-time grid activity through a plurality of distribution rules set by at least one from among a system administrator and an automated decision controller.

Preferably, the present invention provides a computer-implemented method for managing workload , wherein said real-time grid activity comprise data about at least one from among a receipt of a job request, a pending job offer, a newly scheduled grid job, a currently executing grid job performance, a future job schedule, a current pricing schedule, a future pricing schedule, a current resource availability, and a predicted resource availability.

Preferably, the present invention provides a computer-implemented method for managing workload, wherein said selection of grid modules comprises at least one from among a grid resource allocation module, a grid sell-off module, a grid discounter module, a grid accounting module, and a grid pricing module.

Preferably, the present invention provides a method for managing workload, wherein said grid workload controller is a grid service.

Viewed from a second aspect, the present invention provides a system for managing workload within a grid environment, comprising: a grid workload controller implemented within a computational grid environment; said grid workload controller comprising: means for monitoring real-time grid activity at an application level from a plurality of disparate grid application environments within said computational grid environment; means for determining a selection of grid modules that require said real-time grid activity to make decisions about the management of said computational grid environment; and means for distributing said real-time grid activity to said selection of grid modules, wherein said selection of grid modules make automated decisions within said grid environment to maintain workload performance requirements within said computational grid environment.

Preferably, the present invention provides a system for managing workload wherein said means for monitoring real-time grid activity at an application level from a plurality of disparate grid application environments within said computational grid environment, further comprises: means for monitoring said real-time grid activity received in a plurality of protocols requiring adaptation to a particular protocol used by said grid workload controller.

Preferably, the present invention provides a system for managing workload, wherein said means for monitoring real-time grid activity at an application level from a plurality of disparate grid application environments within said computational grid environment, further comprises: means for receiving, at said grid workload controller, a plurality of types of grid characteristics within said real-time grid activity; and means for filtering said plurality of types of grid characteristics within said real-time grid activity according to a plurality of monitoring rules, wherein said plurality of monitoring rules specify a selection of said plurality of types of grid characteristics and a selection of said plurality of disparate grid application environments from which said real-time grid activity is stored.

Preferably, the present invention provides a system for managing workload, wherein said means for monitoring real-time grid activity at an application level from a plurality of disparate grid application environments within said computational grid environment, further comprises: means for interrogating a plurality of manager modules for said real-time grid activity, wherein each of said plurality of manager modules tracks at least one real-time performance characteristic of a grouping of resource nodes allocated to one of said plurality of disparate grid application environments within said computational grid environment.

Preferably, the present invention provides a system for managing workload, wherein said means for determining a selection of grid modules that require said real-time grid activity to make decisions about the management of said computational grid, further comprises: means for filtering said real-time grid activity through a plurality of distribution rules set by at least one from among a system administrator and an automated decision controller.

Preferably, the present invention provides a system for managing workload, wherein said real-time grid activity comprise data about at least one from among a receipt of a job request, a pending job offer, a newly scheduled grid job, a currently executing grid job performance, a future job schedule, a current pricing schedule, a future pricing schedule, a current resource availability, and a predicted resource availability.

Preferably, the present invention provides a system for managing workload, wherein said selection of grid modules comprises at least one from among a grid resource allocation module, a grid sell-off module, a grid discounter module, a grid accounting module, and a grid pricing module.

Preferably, the present invention provides a system for managing workload, wherein said grid workload controller is a grid service.

Viewed from a third aspect the present invention provides a computer program product loadable into the internal memory of a digital computer, comprising software code portions for performing, when said product is run on a computer, to carry out the invention as described above.

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed aspect of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use, further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figure 1 depicts one embodiment of a computer system which may be implemented in a grid environment and in which the present invention may be implemented;

Figure 2 is block diagram illustrating one embodiment of the general types of components within a grid environment;

Figure 3 is a block diagram depicting one example of an architecture that may be implemented in a grid environment;

Figure 4 is a block diagram depicting the components of a grid workload agent for coordinating the monitoring of grid activity from multiple disparate grid modules and distributing the grid activity to those decisional grid modules within a computational grid environment potentially affected by the grid activity;

Figure 5 is a block diagram depicting an example of a grid workload agent interacting with other grid modules within a grid environment in accordance with the method, system, and program of the present invention;

Figure 6 is a block diagram depicting multiple grid manager modules managing disparate groups of grid resources allocated as nodes in accordance with the method, system, and program of the present invention; and

Figure 7 is a high level logic flowchart of a process and program for coordinating the monitoring and distribution of grid activity across disparate monitoring and decisional grid modules within a grid environment in accordance with the method, system, and program of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring now to the drawings and in particular to **Figure 1**, there is depicted one embodiment of a computer system which may be implemented in a grid environment and in which the present invention may be implemented. As will be further described, the grid environment includes multiple computer systems managed to provide resources. Additionally, as will be further described, the present invention may be executed in a variety of computer systems, including a variety of computing systems, mobile systems, and electronic devices operating under a number of different operating systems managed within a grid environment.

In one embodiment, computer system **100** includes a bus **122** or other device for communicating information within computer system **100**, and at least one processing device such as processor **112**, coupled to bus **122** for processing information. Bus **122** may include low-latency and higher latency paths connected by bridges and adapters and controlled within computer system **100** by multiple bus controllers. When implemented as a server system, computer system **100** typically includes multiple processors designed to improve network servicing power.

Processor **112** may be a general-purpose processor such as IBM's PowerPC™ processor that, during normal operation, processes data under the control of operating system and application software accessible from a dynamic storage device such as random access memory (RAM) **114** and a static storage device such as Read Only Memory (ROM) **116**. The operating system may provide a graphical user interface (GUI) to the user. In one embodiment, application software contains machine executable instructions that when executed on processor **112** carry out the operations depicted in the flowcharts of **Figure 7** and others operations described herein. Alternatively, the steps of the present invention might be performed by specific hardware components that contain hardwired logic for performing the steps, or by any combination of programmed computer components and custom hardware components.

The present invention may be provided as a computer program product, included on a machine-readable medium having stored thereon the machine executable instructions used to program computer system **100** to perform a process according to the present invention. The term "machine-readable medium" as used herein includes any medium that participates in providing instructions to processor **112** or other components of computer system **100** for execution. Such a medium may take many forms including, but not limited to, non-volatile media, volatile media, and transmission media. Common forms of non-volatile media include, for example, a floppy disk, a flexible disk, a hard disk, magnetic tape or any other magnetic medium, a compact disc ROM (CD-ROM) or any other optical medium, punch cards or any other physical medium with patterns of holes, a programmable ROM (PROM), an erasable PROM (EPROM), electrically EPROM (EEPROM), a flash memory, any other memory chip or cartridge, or any other medium from which computer system **100** can read and which is suitable for storing instructions. In the present embodiment, an example of a non-volatile medium is mass storage device **118** which as depicted is an internal component of computer system **100**, but will be understood to also be provided by an external device. Volatile media include dynamic memory such as RAM **114**. Transmission media include coaxial cables, copper wire or fiber optics, including the wires that comprise bus **122**. Transmission media can also take the form of acoustic or light waves, such as those generated during radio frequency or infrared data communications.

Moreover, the present invention may be downloaded as a computer program product, wherein the program instructions may be transferred from a remote virtual resource, such as a virtual resource **160**, to requesting computer system **100** by way of data signals embodied in a carrier wave or other propagation medium via a network link **134** (e.g. a modem or network connection) to a communications interface **132** coupled to bus **122**. Virtual resource **160** may include a virtual representation of the resources accessible from a single system or systems, wherein multiple systems may each be considered discrete sets of resources operating on independent platforms, but coordinated as a virtual resource by a grid manager. Communications interface **132** provides a two-way data communications coupling to network link **134** that may be connected, for example, to a local area network (LAN), wide area network (WAN), or an Internet Service Provider (ISP) that provide access to network **102**. In particular, network link **134** may provide wired and/or wireless network communications to one or more networks, such as network **102**, through which use of virtual resources, such as virtual resource **160**, is accessible as provided within a grid environment **150**. Grid environment **150** may be part of multiple types of networks, including a peer-to-peer network, or may be part of a single computer system, such as computer system **100**.

As one example, network **102** may refer to the worldwide collection of networks and gateways that use a particular protocol, such as Transmission Control Protocol (TCP) and Internet Protocol (IP), to communicate with one another. Network **102** uses electrical, electromagnetic, or optical signals that carry digital data streams. The signals through the various networks and the signals on network link **134** and through communication interface **132**, which carry the digital data to and from computer system **100**, are exemplary forms of carrier waves transporting the information. It will be understood that alternate types of networks, combinations of networks, and infrastructures of networks may be implemented.

When implemented as a server system, computer system **100** typically includes multiple communication interfaces accessible via multiple peripheral component interconnect (PCI) bus bridges connected to an input/output controller. In this manner, computer system **100** allows connections to multiple network computers.

Additionally, although not depicted, multiple peripheral components and internal/external devices may be added to computer system **100**, connected to multiple controllers, adapters, and expansion slots coupled to one of the multiple levels of bus **122**. For example, a display device, audio device, keyboard, or cursor control device may be added as a peripheral component.

Those of ordinary skill in the art will appreciate that the hardware depicted in **Figure 1** may vary. Furthermore, those of ordinary skill in the art will appreciate that the depicted example is not meant to imply architectural limitations with respect to the present invention.

With reference now to **Figure 2**, a block diagram illustrates one embodiment of the general types of components within a grid environment. In the present example, the components of a grid environment **150** include a client system **200** interfacing with a grid management system **240** which interfaces with server clusters **222**, servers **224**, workstations and desktops **226**, data storage systems **228**, and networks **230**. For purposes of illustration, the network locations and types of networks connecting the components within grid environment **150** are not depicted. It will be understood, however, that the components within grid environment **150** may reside atop a network infrastructure architecture that may be implemented with multiple types of networks overlapping one another. Network infrastructure may range from multiple large enterprise systems to a peer-to-peer system to a single computer system. Further, it will be understood that the components within grid environment **150** are merely representations of the types of components within a grid environment. A grid environment may simply be encompassed in a single computer system or may encompass multiple enterprises of systems. In addition, it will be understood that a grid vendor may provide grid environment **150**, where the grid vendor may calculate a cost for use of resources within grid environment **150** based on the amount of time required for a grid job to execute or the actual amount of resources used, for example.

The central goal of a grid environment, such as grid environment **150** is organization and delivery of resources from multiple discrete systems viewed as virtual resource **160**. Client system **200**, server clusters **222**, servers **224**, workstations and desktops **226**, data storage systems **228**, networks **230** and the systems creating grid management system 240 may be heterogeneous and regionally distributed with independent management systems, but enabled to exchange information, resources, and services through a grid infrastructure enabled by grid management system **240**. Further, server clusters **222**, servers **224**, workstations and desktops **226**, data storage systems **228**, and networks **230** may be geographically distributed across countries and continents or locally accessible to one another. It will be understood that mechanisms for discovery of grid resources within virtual resource **160** are not depicted herein, however, client system **200** may discover the resources within virtual resource **160** as advertised from local and global directories available within grid environment **150**.

In the example, client system **200** interfaces with grid management system **240**. Client system **200** may represent any computing system sending requests to grid management system **240**. In particular, client system **200** may send virtual job requests and jobs to grid management system **240** and grid management system **240** may respond with a grid offer and controls processing of grid jobs. Further, while in the present embodiment client system **200** is depicted as accessing grid environment **150** with a request, in alternate embodiments client system **200** may also operate within grid environment **150.**

While the systems within virtual resource **160** are depicted in parallel, in reality, the systems may be part of a hierarchy of systems where some systems within virtual resource **160** may be local to client system **200**, while other systems require access to external networks. Additionally, it is important to note, that client system **200** may physically encompass the systems depicted within virtual resources **160.**

To implement grid environment **150**, grid management system **240** facilitates grid services. Grid services may be designed according to multiple architectures, including, but not limited to, the Open Grid Services Architecture (OGSA). In particular, grid management system **240** refers to the management environment which creates a grid by linking computing systems into a heterogeneous network environment characterized by sharing of resources through grid services.

In one example, grid management system **240** may include multiple grid modules that perform grid services for monitoring the grid activity within grid environment **150.** Grid activity may include, but is not limited to, receipt of grid job requests, grid job offers sent out to clients, a newly scheduled grid job, a currently executing grid job, current and future grid pricing, adjustments to grid pricing, current and future predicted workload within virtual resource **160**, and current and future allocation and de-allocation of the resources within virtual resource **160.** In one embodiment, grid activity indicates the real-time status of grid environment **150**, however, in another embodiment, grid modules monitoring and reporting grid activity may delay reporting. Grid modules may include, but are not limited to, monitoring services that monitor current and future performance, workload, and pricing and decisional services that automate decision making within grid environment **150** to maintain performance requirements.

According to an advantage of the invention, grid management system **240** may include a centralized grid workload agent that collects grid activity from across the multiple grid modules and then distributes the grid activity to those grid modules affected by the grid activity. For example, the grid workload agent may collect the data indicating the current workload across a selection of hardware resources from one monitoring module and distribute the current workload to grid modules that determine pricing and availability of resources based on the current workload.

Referring now to **Figure 3**, a block diagram illustrates one example of an architecture that may be implemented in a grid environment. As depicted, an architecture **300** includes multiple layers of functionality. As will be further described, the present invention is a process which may be implemented in one or more layers of an architecture, such as architecture **300**, which is implemented in a grid environment, such as the grid environment described in **Figure 2**. It is important to note that architecture **300** is just one example of an architecture that may be implemented in a grid environment and in which the present invention may be implemented. Further, it is important to note that multiple architectures may be implemented within a grid environment.

Within the layers of architecture **300**, first, a physical and logical resources layer **330** organizes the resources of the systems in the grid. Physical resources include, but are not limited to, servers, storage media, and networks. The logical resources virtualize and aggregate the physical layer into usable resources such as operating systems, processing power, memory, I/O processing, file systems, database managers, directories, memory managers, and other resources.

Next, a web services layer **320** provides an interface between grid services **310** and physical and logical resources **330**. Web services layer **320** implements service interfaces including, but not limited to, Web Services Description Language (WSDL), Simple Object Access Protocol (SOAP), and eXtensible mark-up language (XML) executing atop an Internet Protocol (IP) or other network transport layer. Further, the Open Grid Services Infrastructure (OSGI) standard **322** builds on top of current web services **320** by extending web services **320** to provide capabilities for dynamic and manageable Web services required to model the resources of the grid. In particular, by implementing OGSI standard **322** with web services **320**, grid services **310** designed using OGSA are interoperable. In alternate embodiments, other infrastructures or additional infrastructures may be implemented a top web services layer **320**.

Grid services layer **310** includes multiple services which together provide at least one management function of grid management system **240**. For example, grid services layer **310** may include grid services designed using OGSA, such that a uniform standard is implemented in creating grid services. Alternatively, grid services may be designed under multiple architectures. Grid services can be grouped into four main functions. It will be understood, however, that other functions may be performed by grid services.

First, a resource management service **302** manages the use of the physical and logical resources. Resources may include, but are not limited to, processing resources, memory resources, and storage resources. Management of these resources includes scheduling jobs, distributing jobs, and managing the retrieval of the results for jobs. Resource management service **302** monitors resource loads and distributes jobs to less busy parts of the grid to balance resource loads and absorb unexpected peaks of activity. In particular, a user may specify preferred performance levels so that resource management service **302** distributes jobs to maintain the preferred performance levels within the grid.

Second, information services **304** manages the information transfer and communication between computing systems within the grid. Since multiple communication protocols may be implemented, information services **304** manages communications across multiple networks utilizing multiple types of communication protocols.

Third, a data management service **306** manages data transfer and storage within the grid. In particular, data management service **306** may move data to nodes within the grid where a job requiring the data will execute. A particular type of transfer protocol, such as Grid File Transfer Protocol (GridFTP), may be implemented.

Finally, a security service **308** applies a security protocol for security at the connection layers of each of the systems operating within the grid. Security service **308** may implement security protocols, such as Open Secure Socket Layers (SSL), to provide secure transmissions. Further, security service **308** may provide a single sign-on mechanism, so that once a user is authenticated, a proxy certificate is created and used when performing actions within the grid for the user.

Multiple services may work together to provide several key functions of a grid computing system. In a first example, computational tasks are distributed within a grid. Data management service **306** may divide up a computation task into separate grid services requests of packets of data that are then distributed by and managed by resource management service

302. The results are collected and consolidated by data management system **306.** In a second example, the storage resources across multiple computing systems in the grid are viewed as a single virtual data storage system managed by data management service **306** and monitored by resource management service **302.**

An applications layer **340** includes applications that use one or more of the grid services available in grid services layer **310.** Advantageously, applications interface with the physical and logical resources **330** via grid services layer **310** and web services **320,** such that multiple heterogeneous systems can interact and interoperate.

With reference now to **Figure 4,** there is depicted a block diagram of the components of a grid workload agent for coordinating the monitoring of grid activity from multiple disparate grid modules and distributing the grid activity to those decisional grid modules within a computational grid environment potentially affected by the grid activity. As depicted, a grid workload agent **400** includes multiple components, which may be implemented within hardware and software across multiple disparate server systems, or within a single server system, within grid environment **150.**

In the example, grid workload agent **400** includes a module database **404.** Module database **404** stores information about each grid module that grid workload agent **400** may monitor for grid activity or to which grid workload agent **400** may distribute grid activity. In particular, module database **404** may include monitoring rules **405** that specify a selection of grid modules and types of grid activity to monitor. In addition, module database **404** may include distribution rules **406** that specify which grid modules should receive grid activity and the types of grid activity to distribute to each grid module. It will be understood that monitoring rules **405** and distribution rules **406** may be provided by a system administrator or may be automatically determined based on the type of grid activity information used by grid modules for making decisions within grid environment **150.**

In addition, in the example, grid workload agent **400** includes a grid activity database **410** for storing real-time grid activity collected by grid workload agent **400** and for storing records of the distributions of the grid activity to other grid modules. It will be understood that grid activity database **410** may also be referenced by a grid accounting module that calculates job cost.

In the example, grid workload agent **400** includes a module monitoring controller **402** for managing the inquiry to grid modules for grid activity and the receipt of grid activity from grid modules. In particular, module monitoring controller **402** may query those grid modules specified in monitoring rules **405** and may filter received grid activity according to monitoring rules **405.** In one example, module monitoring controller **402** may include a queuing system to manage the receipt and distribution of grid activity. For example, a queuing system, such as the MQ Series (MQ Series is a registered trademark of International Business Machines Corporation), may manage the receipt of grid activity packaged using multiple disparate communication protocols and control the distribution of each packaged message to an adapter that is enabled to map the message into a generic protocol used by grid workload agent **400.**

In one example, module monitoring controller **402** queries a job offer controller for grid activity indicating grid offers made to client systems requesting resource usage for a particular job. In another example, module monitoring controller **402** queries a resource manager for grid activity indicating the current workload and performance characteristics for a group of resources managed by the resource manager. In yet another example, module monitoring controller **402** queries a pricing controller for grid activity, where the grid activity indicates current pricing for a particular selection of grid resources.

Further, in the example, grid workload agent **400** includes a distribution controller **408** for managing the distribution of grid activity to grid modules. In particular, distribution controller **408** may distribute grid activity according to distribution rules **406.** Further, distribution controller **408** may include a mapping controller to map grid activity into a protocol understood by the grid module receiving the grid activity.

Referring now to **Figure 5**, there is depicted a block diagram of an example of a grid workload agent interacting with other grid modules within a grid environment in accordance with the method, system, and program of the present invention. As illustrated, grid management system **240** for an entire grid environment **150** may include multiple grid modules, including a grid workload agent **400.** Grid workload agent **400** tracks real-time grid activity received from monitoring and decisional grid modules within grid environment **150** and distributes the grid activity to selections of decisional modules within grid management system **240** that are likely to be affected by the grid activity or require the grid activity to make decisions.

In one example, grid workload agent **400** receives real-time grid activity indicating resource performance and workload from a grid manager module **504**. Grid manager module **504** monitors the grid activities for a grid application environment including resources of a same or similar type and manages the nodes within the grid application environment. In particular, grid manager module **504** may submit grid activities that indicate resource usage of a selection of grid resources and the status of grid jobs and workload for the selection of grid resources. Grid workload agent **400** receives the grid activity from grid manager module **504** and may distribute this performance and workload data to other decisional grid modules, such as a grid accounting module **510** that calculates the current cost of a job, a grid resource allocation module **502** that determines whether additional resource nodes should be allocated or de-allocated for particular purposes based on current workload and performance, or a grid sell-off module **506** that attempts to sell-off grid jobs if the current or future predicted workload is too much for the current grid resources.

In another example, grid resource allocation module **502** monitors resource node allocation and de-allocation, including allocation and de-allocation of on-demand grid resources accessed from grid environments outside of grid environment **150**. Grid workload agent **400** receives this grid activity from grid resource allocation module **502** and may distribute the allocation activity information to a grid sell-off module **506** that determines whether grid load is too heavy and grid jobs need to be transferred to grid farms or vendors outside grid environment **150**.

In yet another example, grid workload agent **400** distributes any grid activity indicating low resource usage levels, currently or predicted for the future, to a grid discounter module **508**. Grid discounter module **508** determines whether to offer grid resources as discounted pricing. Decisions by grid discounter module **508** to discount pricing are received by grid workload agent **400** and distributed, for example, to a grid pricing module **512** that determines pricing offers for grid job requests received by grid management system **240**.

It will be understood that the grid modules illustrated within grid management system **240** are merely examples of the types of grid modules that may monitor grid activity at an application, hardware, and network level, across grid resources and are merely examples of the types of grid modules that provide automated decisional services within grid management system **240**. In particular, additional decisional modules that predict current workload requirement, further workload requirements, and potential changes within the grid environment may receive grid activity.

With reference now to **Figure 6**, there is depicted a block diagram of multiple grid manager modules managing disparate groups of grid resources allocated as nodes in accordance with the method, system, and program of the present invention. As depicted, multiple grid manager modules **602**, **604**, **606**, and **608** manage the usage of groups of grid resources **610**, **612**, **614**, and **616**, respectively.

Each group of grid resources may include multiple nodes, where a node is logical representation of a hardware, software, or network resource assigned to a particular purpose. For example, group **610** includes DB2EEE nodes **620**, **622**, **624**, and **626**, monitored and managed within grid environment **150** by grid manager module **402** (DB2 is a registered trademark of International Business Machines Corporation). It will be understood that a group of grid resources may include nodes of the same type of resources or of disparate resources.

In the example, each of groups of grid resources **610**, **612**, **614**, and **616** includes nodes of resources operating on disparate platforms or for different purposes. For purposes of illustration, each group of grid resources represents a different grid application environment. For example, group **610** includes DB2EEE grid nodes **620**, **622**, **624**, and **626** and other associated resources which are dedicated to executing jobs required a parallel DB2 database environment. In another example, group **612** includes DB2EE grid nodes **630**, **632**, **634**, and **636** and other associated resources currently dedicated to executing jobs requiring a standard DB2 database environment. In addition, in another example, group **614** includes six compute nodes, which are dedicated to executing logic which does not require specific applications or data, but rather simply requires powerful CPU's and a high compiled C code mode. Further, in another example, group **616** includes idle nodes **640** and **642**, which are included in a pool of resources not currently assigned to execution of any computational tasks and available for incorporation into another grid application environment, such as group **610**.

In the example, grid workload agent **400** can track application level grid activity for four separate grid application system environments by communicating with grid manager modules **602, 604, 606,** and **608.** The grid activity reported by each of grid manager modules **602, 604, 606,** and **608** may effect the decisions made by the other grid manager modules and by other grid modules to which grid workload agent **400** distributes grid activity.

In one example, grid manager module **602** may report grid activity to grid workload agent **400** that indicates the DB2EEE nodes within group **610** are only using 50% of the available CPU resources within group **610** and that there no future database jobs requiring a parallel DB2 database environment are scheduled. Grid manager module **602** may distribute this grid activity to grid discounting module **510** that then determines whether to offer discounted pricing for database jobs. In addition, grid manager module **602** may distribute this grid activity to grid resource allocation module **502** that then determines whether to de-allocate a node within group **610** and reallocate the node within another group with a heavier job load. In particular, grid workload agent **400** would manage the communication of grid activity between grid manager module **602** and grid resource allocation module **502** to inform grid manager module **602** that a node will be removed from use for database jobs, work scheduled for the node would be rerouted to other nodes within group **610,** and all current transactions would be completed. Once node **622,** for example, is idle from DB2EEE transactions, grid workload agent **400** would update the grid modules with grid activity indicating that node **622** is idle and grid manager module **608** would begin monitoring and management of node **622.**

In another example, grid manager module **604** detects that the nodes in group **612** are running at 100% capacity and updates grid workload agent **400** with the current capacity. Grid workload agent **400** determines that the capacity percentage should be distributed to grid resource allocation module **502** and to grid pricing module **512.** Grid resource allocation module **502** returns grid activity information to grid workload agent **400** indicating whether a new node is allocated to grid group **612.** Grid workload agent **400** distributes the new database node grid activity to grid manager module **602.** In addition, grid pricing module **512** may determine that the rate for database jobs needs increase for the next ten hours and returns the price increase to grid workload agent **400.** Grid workload agent **400** may determine to distribute the price increase to grid accounting module **510** so that grid accounting module **510** can properly calculate a cost of database jobs using grid group **612** for the next ten hours.

It is important to note that while the present example is described with reference to nodes grouped according to application environment, it alternate examples, nodes may be grouped according to other functional criteria, such as hardware environment, network environment, and other groupings. In addition, it will be understood that resource nodes within a grouping may be physically grouped together or may be located in disparate physical locations.

Referring now to **Figure 7**, there is depicted a high level logic flowchart of a process and program for coordinating the monitoring and distribution of grid activity across disparate monitoring and decisional grid modules within a grid environment in accordance with the method, system, and program of the present invention. It will be understood that the process depicted may be executed within grid workload agent 400, controlled by applicable policies and rules, but that additional manual or automated administrative decisions may be required for coordinating the monitoring and distribution of grid activity.

As depicted, the process starts at block **700** and thereafter proceeds to block **702.** Block **702** depicts a determination whether grid activity indicators are detected. If grid activity indicates are not detected, then the process passes to block **704.** Block **704** depicts periodically archiving grid activity to the grid account module to maintain historical data for statistical analysis and accounting activities. Next, block **706** depicts interrogating the grid manager modules for currently executing and scheduled jobs and other grid modules monitoring grid performance, and the process returns to block **702.**

Otherwise, at block **702**, if grid activity is detected, then the process passes to block **708.** Block **708** depicts filtering the grid activity according to monitoring rules. Next, block **710** depicts a determination whether the activity indicates that a new application specific job is scheduled. If the activity does not indicate that a new application specific job is scheduled, then the process passes to block **714.** Block **714** depicts interrogating the application environment for current workload statistics or reading the current archived snapshot of grid activity, and the process passes to block **716.** Otherwise, at block **710,** if the grid activity indicates a new application specific job is scheduled for performance within the grid environment, then the process passes to block **712.** Block **712** depicts interrogating the general grid platform for current workload statistics or the current archived snapshot of grid activity is read, and the process passes to block **716.**

Block **716** depicts performing predictive analysis to estimate the impact of the job submission within the grid environment. Next, block **718** depicts storing the predictive data for comparison with actual job impact for more accurate future prediction. Thereafter, block **720** depicts distributing grid activity data to potentially effected decisional grid modules, and the process ends.

CLAIMS

1. A method for managing workload within a grid environment, comprising:

monitoring, at a grid workload controller, real-time grid activity at an application level from a plurality of disparate grid application environments within a computational grid environment;

determining a selection of grid modules that require said real-time grid activity to make decisions about the management of said computational grid environment; and

distributing, from said grid workload controller, said real-time grid activity to said selection of grid modules, wherein said selection of grid modules make automated decisions within said grid environment to maintain workload performance requirements within said computational grid environment.

2. A method for managing workload according to claim 1, wherein monitoring, at a grid workload controller, real-time grid activity at an application level from a plurality of disparate grid application environments within a computational grid environment, further comprises:

monitoring, at said grid workload controller, said real-time grid activity received in a plurality of protocols requiring adaptation to a particular protocol used by said grid workload controller.

3. A method for managing workload according to claim 1, wherein monitoring, at a grid workload controller, real-time grid activity at an application level from a plurality of disparate grid application environments within a computational grid environment, further comprises:

receiving, at said grid workload controller, a plurality of types of grid characteristics within said real-time grid activity; and

filtering said plurality of types of grid characteristics within said real-time grid activity according to a plurality of monitoring rules, wherein said plurality of monitoring rules specify a selection of said plurality of types of grid characteristics and a selection of said plurality of disparate grid application environments from which said real-time grid activity is stored.

4. A method for managing workload according to claim 1, wherein monitoring, at a grid workload controller, real-time grid activity at an application level from a plurality of disparate grid application environments within a computational grid environment, further comprises:

interrogating a plurality of manager modules for said real-time grid activity, wherein each of said plurality of manager modules tracks at least one real-time performance characteristic of a grouping of resource nodes allocated to one of said plurality of disparate grid application environments within said computational grid environment.

5. A method for managing workload according to claim 1, wherein determining a selection of grid modules that require said real-time grid activity to make decisions about the management of said computational grid, further comprises:

filtering said real-time grid activity through a plurality of distribution rules set by at least one from among a system administrator and an automated decision controller.

6. A method for managing workload according to claim 1, wherein said real-time grid activity comprise data about at least one from among a receipt of a job request, a pending job offer, a newly scheduled grid job, a currently executing grid job performance, a future job schedule, a current pricing schedule, a future pricing schedule, a current resource availability, and a predicted resource availability.

7. A method for managing workload according to claim 1, wherein said selection of grid modules comprises at least one from among a grid resource allocation module, a grid sell-off module, a grid discounter module, a grid accounting module, and a grid pricing module.

8. A method for managing workload according to claim 1, wherein said grid workload controller is a grid service.

9. A system for managing workload within a grid environment, comprising:

a grid workload controller implemented within a computational grid environment;

said grid workload controller comprising:

means for monitoring real-time grid activity at an application level from a plurality of disparate grid application environments within said computational grid environment;

means for determining a selection of grid modules that require said real-time grid activity to make decisions about the management of said computational grid environment; and

means for distributing said real-time grid activity to said selection of grid modules, wherein said selection of grid modules make automated decisions within said grid environment to maintain workload performance requirements within said computational grid environment.

10. A system for managing workload according to claim 9, wherein said means for monitoring real-time grid activity at an application level from a plurality of disparate grid application environments within said computational grid environment, further comprises:

means for monitoring said real-time grid activity received in a plurality of protocols requiring adaptation to a particular protocol used by said grid workload controller.

11. A system for managing workload according to claim 9, wherein said means for monitoring real-time grid activity at an application level from a plurality of disparate grid application environments within said computational grid environment, further comprises:

means for receiving, at said grid workload controller, a plurality of types of grid characteristics within said real-time grid activity; and

means for filtering said plurality of types of grid characteristics within said real-time grid activity according to a plurality of monitoring rules, wherein said plurality of monitoring rules specify a selection of said plurality of types of grid characteristics and a selection of said plurality of disparate grid application environments from which said real-time grid activity is stored.

12. A system for managing workload according to claim 9, wherein said means for monitoring real-time grid activity at an application level from a plurality of disparate grid application environments within said computational grid environment, further comprises:

means for interrogating a plurality of manager modules for said real-time grid activity, wherein each of said plurality of manager modules tracks at least one real-time performance characteristic of a grouping of resource nodes allocated to one of said plurality of disparate grid application environments within said computational grid environment.

13. A system for managing workload according to claim 9, wherein said means for determining a selection of grid modules that require said real-time grid activity to make decisions about the management of said computational grid, further comprises:

means for filtering said real-time grid activity through a plurality of distribution rules set by at least one from among a system administrator and an automated decision controller.

14. A system for managing workload according to claim 9, wherein said real-time grid activity comprise data about at least one from among a receipt of a job request, a pending job offer, a newly scheduled grid job, a currently executing grid job performance, a future job schedule, a current pricing schedule, a future pricing schedule, a current resource availability, and a predicted resource availability.

15. A system for managing workload according to claim 9, wherein said selection of grid modules comprises at least one from among a grid resource allocation module, a grid sell-off module, a grid discounter module, a grid accounting module, and a grid pricing module.

16. A system for managing workload according to claim 9, wherein said grid workload controller is a grid service.

17. A computer program product loadable into the internal memory of a digital computer, comprising software code portions for performing, when said product is run on a computer, to carry out the invention as claimed in claims 1 to 11.

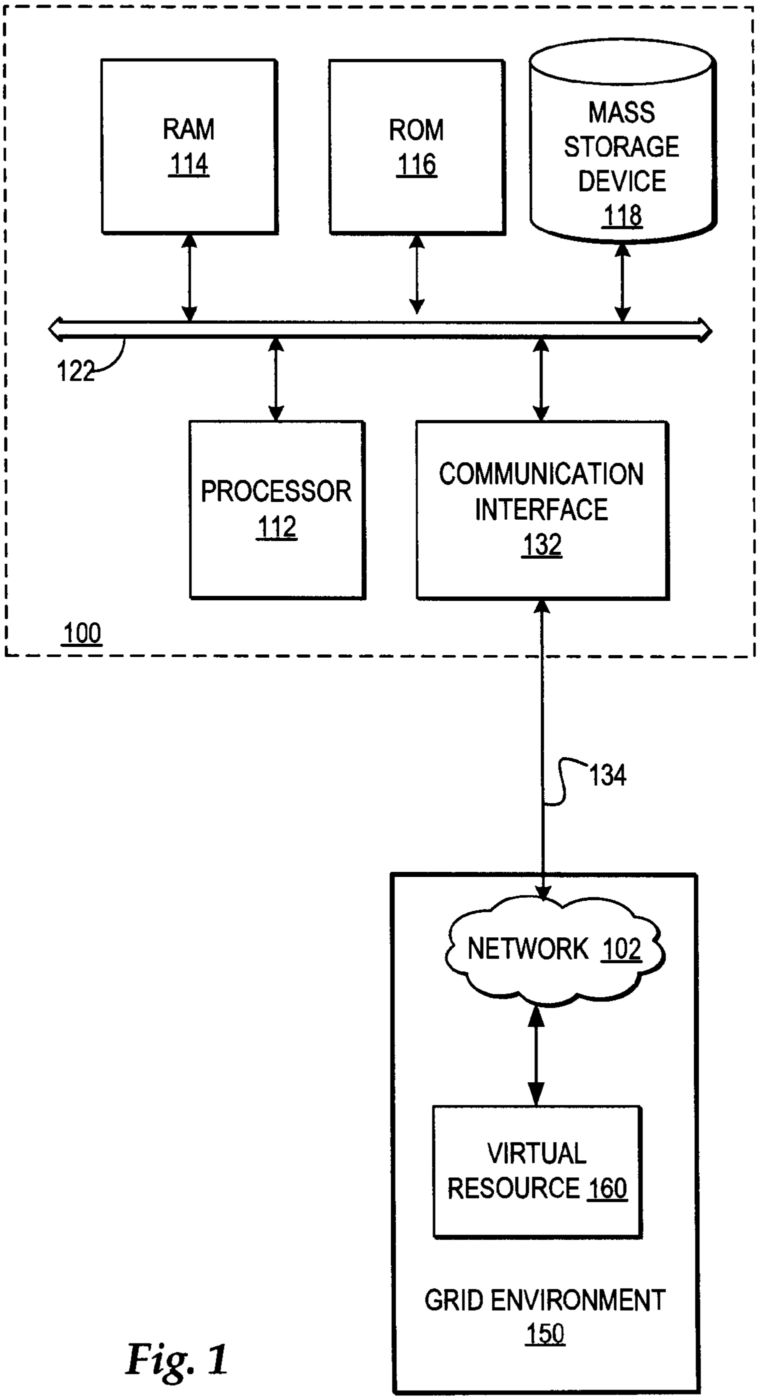
RAM
114
ROM
116
MASS
STORAGE
DEVICE
118
122
PROCESSOR
112
COMMUNICATION
INTERFACE
132
100
134
NETWORK 102
VIRTUAL
RESOURCE 160
GRID ENVIRONMENT
150

Fig. 1

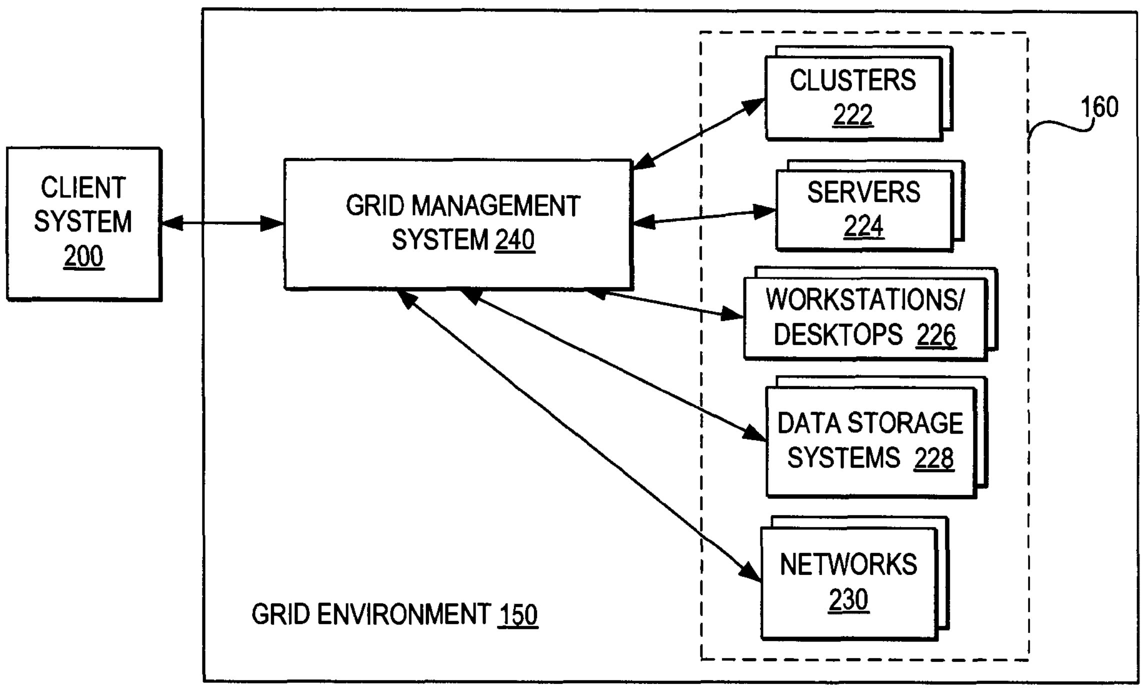
CLIENT SYSTEM 200
GRID MANAGEMENT SYSTEM 240
CLUSTERS 222
SERVERS 224
WORKSTATIONS/ DESKTOPS 226
DATA STORAGE SYSTEMS 228
NETWORKS 230
160
GRID ENVIRONMENT 150

Fig. 2

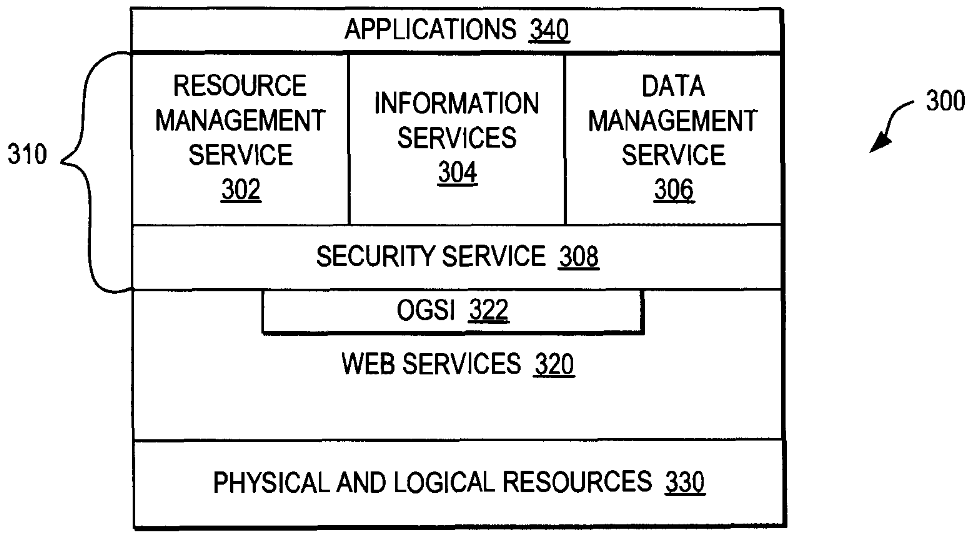
APPLICATIONS 340
RESOURCE MANAGEMENT SERVICE 302
INFORMATION SERVICES 304
DATA MANAGEMENT SERVICE 306
SECURITY SERVICE 308
310
300
OGSI 322
WEB SERVICES 320
PHYSICAL AND LOGICAL RESOURCES 330

Fig. 3

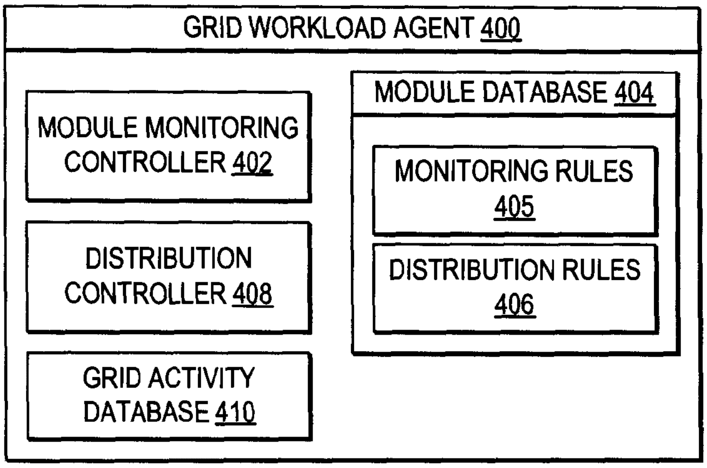
GRID WORKLOAD AGENT 400
MODULE MONITORING CONTROLLER 402
DISTRIBUTION CONTROLLER 408
GRID ACTIVITY DATABASE 410
MODULE DATABASE 404
MONITORING RULES 405
DISTRIBUTION RULES 406

Fig. 4

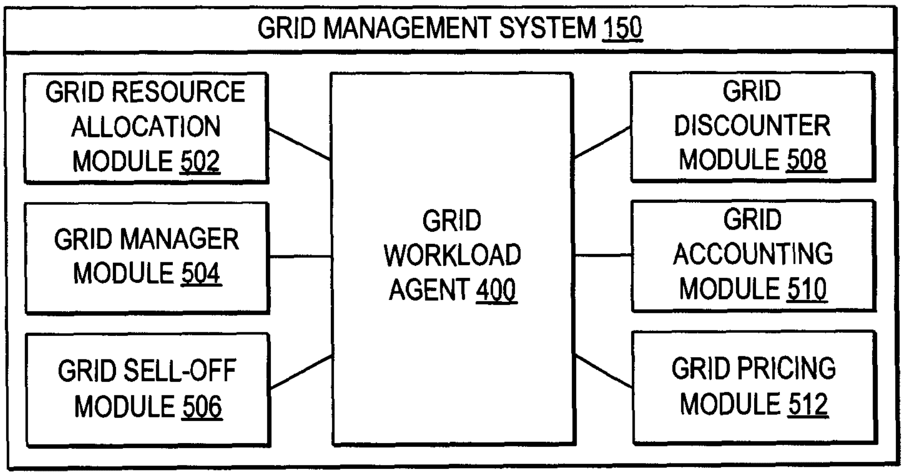
GRID MANAGEMENT SYSTEM 150
GRID RESOURCE ALLOCATION MODULE 502
GRID MANAGER MODULE 504
GRID SELL-OFF MODULE 506
GRID WORKLOAD AGENT 400
GRID DISCOUNTER MODULE 508
GRID ACCOUNTING MODULE 510
GRID PRICING MODULE 512

Fig. 5

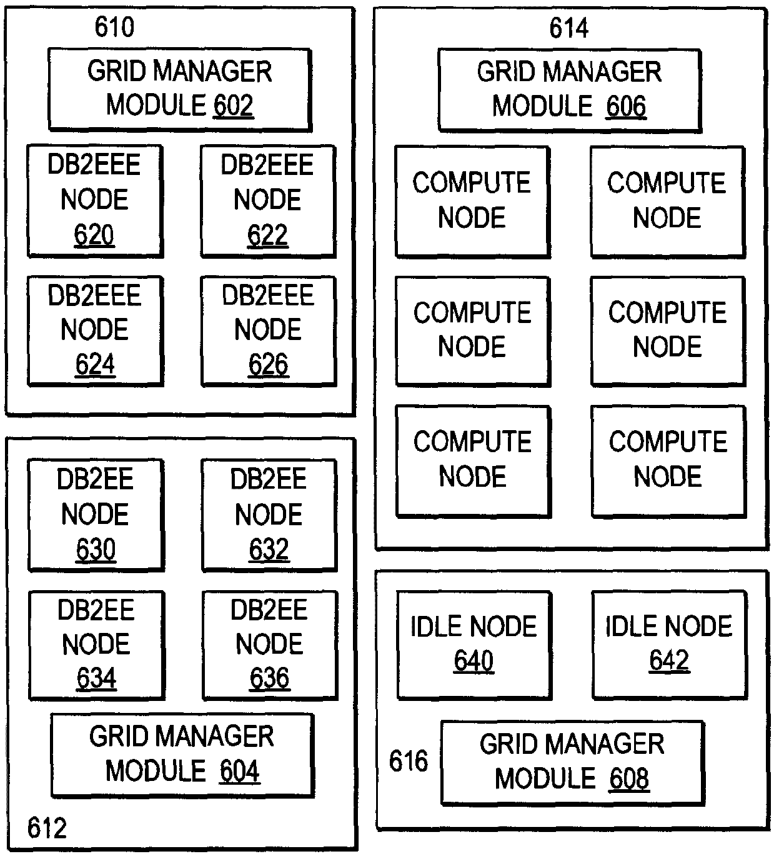
610
GRID MANAGER MODULE 602
DB2EEE NODE 620
DB2EEE NODE 622
DB2EEE NODE 624
DB2EEE NODE 626
614
GRID MANAGER MODULE 606
COMPUTE NODE
COMPUTE NODE
COMPUTE NODE
COMPUTE NODE
COMPUTE NODE
COMPUTE NODE
DB2EE NODE 630
DB2EE NODE 632
DB2EE NODE 634
DB2EE NODE 636
GRID MANAGER MODULE 604
612
IDLE NODE 640
IDLE NODE 642
616
GRID MANAGER MODULE 608

Fig. 6

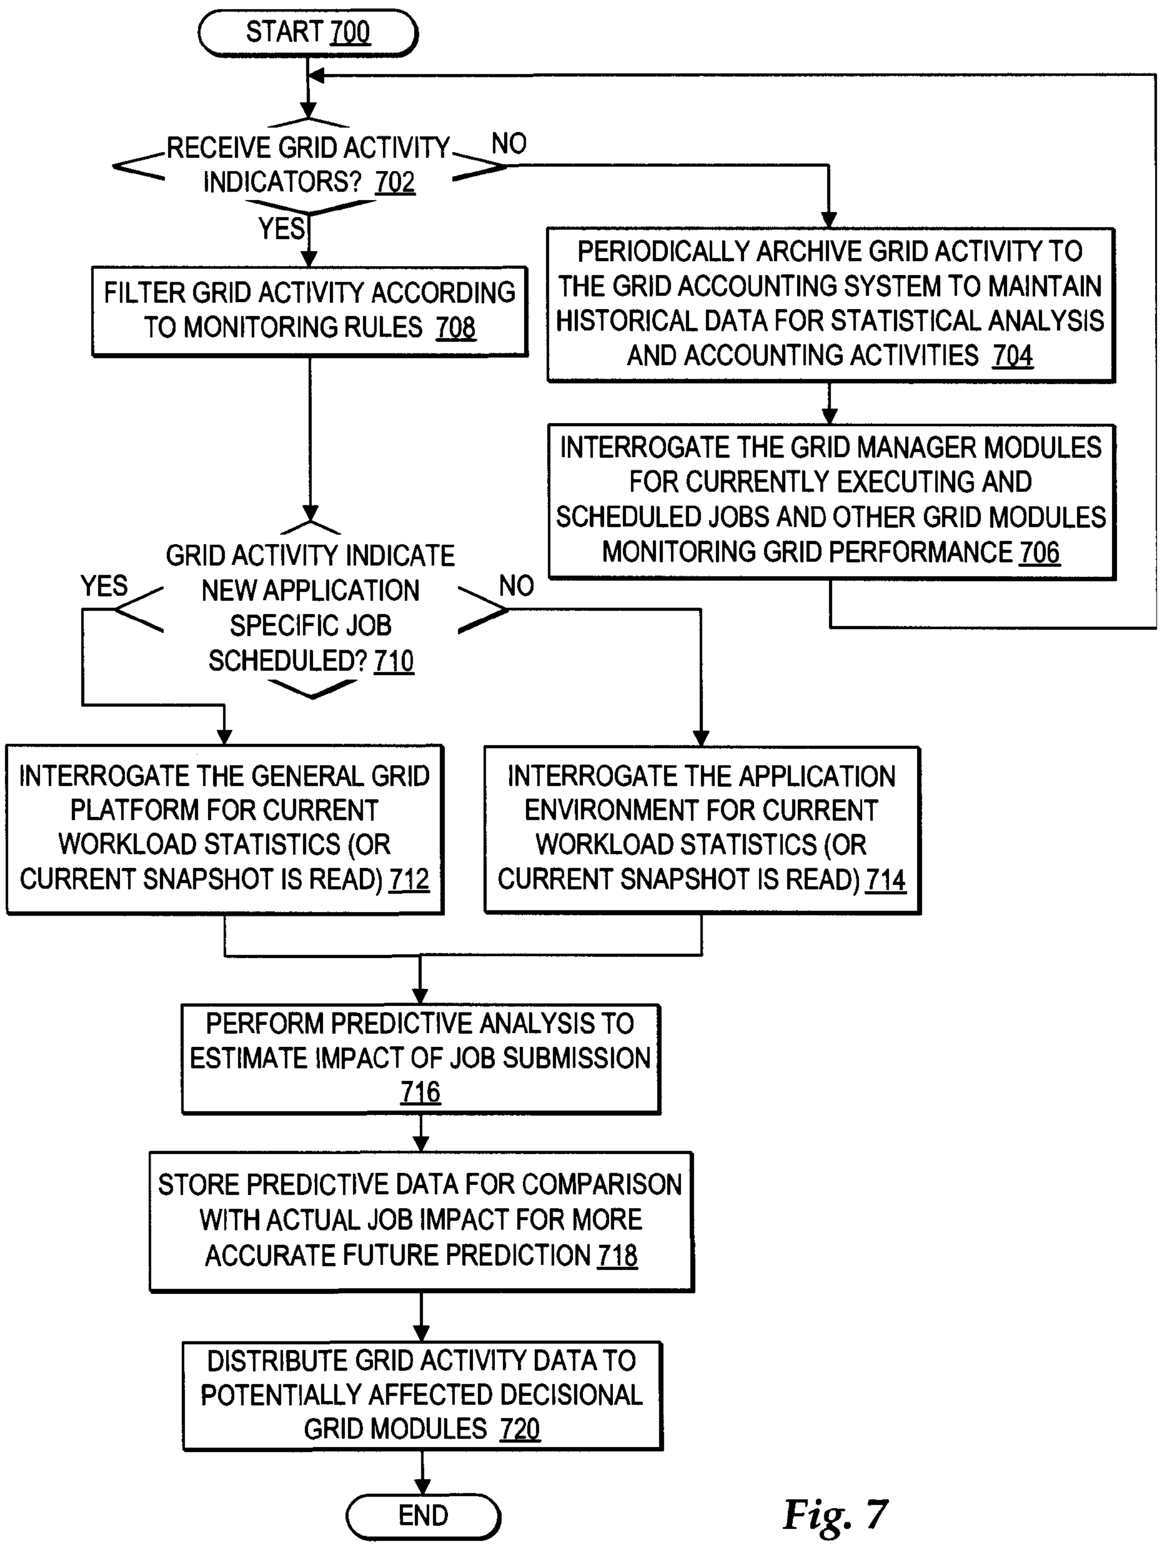
START 700
RECEIVE GRID ACTIVITY INDICATORS? 702
NO
YES
FILTER GRID ACTIVITY ACCORDING TO MONITORING RULES 708
PERIODICALLY ARCHIVE GRID ACTIVITY TO THE GRID ACCOUNTING SYSTEM TO MAINTAIN HISTORICAL DATA FOR STATISTICAL ANALYSIS AND ACCOUNTING ACTIVITIES 704
INTERROGATE THE GRID MANAGER MODULES FOR CURRENTLY EXECUTING AND SCHEDULED JOBS AND OTHER GRID MODULES MONITORING GRID PERFORMANCE 706
GRID ACTIVITY INDICATE NEW APPLICATION SPECIFIC JOB SCHEDULED? 710
YES
NO
INTERROGATE THE GENERAL GRID PLATFORM FOR CURRENT WORKLOAD STATISTICS (OR CURRENT SNAPSHOT IS READ) 712
INTERROGATE THE APPLICATION ENVIRONMENT FOR CURRENT WORKLOAD STATISTICS (OR CURRENT SNAPSHOT IS READ) 714
PERFORM PREDICTIVE ANALYSIS TO ESTIMATE IMPACT OF JOB SUBMISSION 716
STORE PREDICTIVE DATA FOR COMPARISON WITH ACTUAL JOB IMPACT FOR MORE ACCURATE FUTURE PREDICTION 718
DISTRIBUTE GRID ACTIVITY DATA TO POTENTIALLY AFFECTED DECISIONAL GRID MODULES 720
END

Fig. 7

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2005/056931

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F11/34

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)
G06F

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 practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	I FOSTER C KESSELMAN: "The Grid 2, pp 333-340, sections 20.3 - 20.4" 2004, ELSEVIER , SAN FRANCISCO , XP002379532 Section 20.3 "Sensors and Sensor Management": page 333 - page 334, line 15 Section 20.4: page 336, line 1 - line 10 Section 20.4.3: page 337 - page 338; figure 20.5 Section 20.4.5: page 339 - page 340	1,3,4, 6-9,11, 12,14-17
Y		2,5,10, 13

	-/--	

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
4 May 2006	24/05/2006

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Ross, C

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2005/056931

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	EP 1 267 552 A (WTS WIRELESS TECHNOLOGY SYSTEMS OY) 18 December 2002 (2002-12-18) abstract; figure 1 paragraph [0016] paragraph [0028]	2,5,10, 13
A	US 5 729 472 A (SEIFFERT ET AL) 17 March 1998 (1998-03-17) abstract column 10, line 37 - column 11, line 36; figure 1	3,5,11, 13
A	EP 0 790 559 A (HITACHI, LTD; HITACHI ULSI ENGINEERING CORP) 20 August 1997 (1997-08-20) abstract; figure 1	1-17

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/EP2005/056931

Patent document cited in search report		Publication date	Patent family member(s)			Publication date
EP 1267552	A	18-12-2002	FI	20011273	A	16-12-2002
US 5729472	A	17-03-1998	NONE			
EP 0790559	A	20-08-1997	DE	69712552	D1	20-06-2002
			DE	69712552	T2	09-01-2003
			US	5961596	A	05-10-1999