



- (51) International Patent Classification:  
*G06F 1/00* (2006.01)
- (21) International Application Number:  
PCT/US2012/027221
- (22) International Filing Date:  
1 March 2012 (01.03.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
13/039,214 2 March 2011 (02.03.2011) US
- (72) Inventor; and
- (71) Applicant : RALLO, Aaron J. [US/US]; 15127 NE 24th Street, Number 479, Redmond, WA 98052 (US).
- (74) Agent: TAN, Carina; Reed Smith LLP, P.O. Box 488, Pittsburgh, PA 15230 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ,

CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:  
— with international search report (Art. 21(3))

WO 2012/118934 A1

(54) Title: NON-INTRUSIVE POWER MANAGEMENT




FIG. 1

(57) Abstract: A method and system for managing power consumption of a pool of computing devices that are logically grouped to provide a common set of functionality is disclosed. One aspect of certain embodiments includes predicting resource utilization for each device without installing customized software, firmware or hardware on the device.

## Non-Intrusive Power Management

### TECHNICAL FIELD

[0001] The disclosed embodiments relate generally to power management in any pool of computing devices that are logically grouped to provide a common set of functionality. More particularly, the disclosed embodiments relate to power management in server pools, interchangeably referred to as server clusters, typically found in large computing establishments like data centers.

### BACKGROUND

[0002] The proliferation of the Internet, devices that access it, and consequently, Internet based services are driving an insatiable thirst for computational power. To meet this need, large data centers have been set up. Typical data centers house hundreds, maybe even thousands of servers, and serve as the backbone for a variety of Internet services. The services hosted by data centers typically come with the requirement of high availability, close to 99.9% up time, which is usually supported by replicating servers and maintaining spare capacity. Furthermore, data centers are designed for a peak loads which are both occasional and short lived. As a result, data centers tend to consume large amounts of power. In phases that the data center is not fully loaded, idle servers can be shutdown without substantial loss in performance. When the load increases, powered off servers can be booted-up to service the requests and maintain Quality of Service (QoS).

[0003] Reducing the power consumption of a data center contributes to reduced operational expense, and allows the data center operator to invest in newer hardware and supporting infrastructure, to save money and/or to provide improved services to customers. Prior studies have reported that servers can draw close to 60% of their peak power consumption when idle, and that the global electricity costs for data centers have been reported as

running into the billions. Therefore, substantial reduction in power consumption can be achieved by shutting down idle servers.

[0004]

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0005] For a better understanding of the aspects of the invention as well as embodiments thereof, reference should be made to the description of embodiments below, in conjunction with the following drawings in which like reference numerals refer to corresponding parts throughout the figures.

[0006] FIG. 1 is a high-level block diagram illustrating power management of a pool of computing devices that are logically grouped to provide a common set of functionality, according to certain embodiments of the invention.

[0007] FIG. 2 is a block diagram showing some of the high-level steps for obtaining correlation information associated with the servers in the server pool, according to certain embodiments of the invention.

[0008] FIG. 3 is a block diagram that illustrates a power management method, according to certain embodiments of the invention.

[0009] FIG. 4 illustrates the class diagram of the central classes used for implementing the power manager, according to certain embodiments of the invention.

[0010] FIG. 5 illustrates the class diagram for the LoadInformation class hierarchy, according to certain embodiments of the invention.

[0011] FIG. 6 illustrates the class diagram for the UtilizationPredictor class hierarchy, according to certain embodiments of the invention.

[0012] FIG. 7 illustrates the class diagram for the ResourcesMeasureMethod class hierarchy, according to certain embodiments of the invention.

[0013] FIG. 8 illustrates the class diagram for the LoadBalancer class hierarchy, according to certain embodiments of the invention.

## DESCRIPTION OF EMBODIMENTS

[0014] Methods, systems and other aspects of the invention are described. Reference will be made to certain embodiments of the invention, examples of which are illustrated in the accompanying drawings. While the invention will be described in conjunction with the embodiments, it will be understood that it is not intended to limit the invention to these particular embodiments alone. On the contrary, the invention is intended to cover alternatives, modifications and equivalents that are within the spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

[0015] Moreover, in the following description, numerous specific details are set forth to provide a thorough understanding of the present invention. However, it will be apparent to one of ordinary skill in the art that the invention may be practiced without these particular details. In other instances, methods, procedures, components, and networks that are well known to those of ordinary skill in the art are not described in detail to avoid obscuring aspects of the present invention.

[0016] The embodiments described herein are in reference to servers in data centers. However, the embodiments apply to any pool of computing devices that are logically grouped to provide a common set of functionality.

[0017] According to certain embodiments, the problem associated with power consumption in data centers can be effectively managed by turning off or turning on servers in response to the load experienced by the data center. Servers are turned on when the load increases and turned off when load decreases. Load can be defined by the number and/or size of requests that are being received by the server pool per unit time interval, for example.

[0018] According to certain embodiments, a new server or device is characterized to understand how the resource utilization changes as the number of requests being serviced changes by the server/device. The characterization, using statistical analysis techniques, can be used to predict the utilization of the server/device for a given load. The correlation function

associated with the characterization is stored in a database, for example. A power management server can retrieve the correlation function during initialization. The power management server takes decisions at regular time intervals to shutdown a server/device, power-on a server/device or maintain status quo in the pool of servers/devices based on the predicted utilization.

[0019] According to certain embodiments, a non-intrusive mechanism is used to power down servers or devices. In contrast, existing power management solutions typically require that the data center operators install software, firmware or hardware on the servers/devices for power management. Such power management decisions are taken by a centralized administrative software component which communicates with the software installed in the individual servers, which then initiate the action. The custom software typically sends information that the centralized entity can use for decision making. Such an approach is intrusive unlike the non-intrusive approach as described in the embodiments herein. The embodiments described herein do not require any such additional software, firmware or hardware installation on each server/device in the data center.

[0020] According to certain embodiments, a centralized entity takes power management decisions and initiates them on the servers/devices without the need for custom software, hardware or firmware. The centralized entity uses information exported by the OS only of the servers/devices. Such an approach requires little or no downtime for installation, does not require custom software to be installed, or require any major system reconfiguration.

[0021] Further, unlike vendor specific solutions, the embodiments are not restricted to hardware vendors (processor or OEM) or to operating systems.

[0022] FIG. 1 is a high-level block diagram illustrating power management of a pool of computing devices that are logically grouped to provide a common set of functionality, such as servers in a data center, according to certain embodiments of the invention. In FIG. 1, system 100

includes an application delivery controller 104 that receives HTTP requests 102 from client devices, and a computer 106 that executes the power manager. Application delivery controller 104 sends the HTTP requests 102 to the server pool 108 and also receives the responses to the HTTP requests from server pool 108. The power manager implemented in computer 106 receives information from application delivery controller 104 and information from the server pool 108 to make power management decisions. The power manager may be implemented on multiple computers as in a distributed computer system, according to certain embodiments. Application delivery controller 104 may be a commercial off-the-shelf load balancer, according to certain embodiments. Similarly, computer 106 can be an off-the-shelf computer on which the power management solution is installed and executes. Server pool 108 or server cluster comprises server machines or nodes that service requests from client devices via application delivery controller 104. An application delivery controller is hardware or software that manages requests received from client devices and distributes such requests to the computing devices in the server pool. A non-limiting example of an application delivery controller is a load balancer.

[0023] HTTP requests initiated by client devices reach application delivery controller 104 which redirects the requests to an appropriate server in the server pool 108. According to certain embodiments, application delivery controller 104 is configured to use a round robin policy. Consequently, server nodes in server pool 108 service a comparable number of requests. The power manager interacts with application delivery controller 104 to obtain information including but not limited to:

- Information on the number of requests being executed by each server in server pool 108,
- the average response time by each server in server pool 108, and
- information on server state.

[0024] The power manager does not service any requests from client devices. The power manager's job is to make power management decisions

and initiate such decisions, while maintaining consistency between actual server state and information at application delivery controller 104.

[0025] According to one aspect of certain embodiments, each server of at least a subset of servers in the server pool is characterized for the utilization behaviour of that particular server. Characterization involves measuring on the server to be characterized, the utilization of various resources as the number of requests being executed by the server varies. Such measurement information is utilized to draw correlations between the number of requests being serviced by the server that is being characterized and its utilization of resources, according to certain embodiments. The power manager (computer 106) can remotely query the servers in server pool 108 to obtain resource utilization information using standardized protocols like Simple Network Management Protocol (SNMP) for any OS or Windows Management Instrumentation (WMI) for MS Windows. The correlation drawn can be used to predict the utilization of a given server for any given number of HTTP requests being serviced per minute, according to certain embodiments. According to certain embodiments, the characterization is performed using the same application that the server to be characterized is expected to execute in production because a server can be expected to show differences in behaviour with different application types.

[0026] According to certain embodiments, correlation information is obtained using well established statistical analysis techniques such as linear regression. The statistical analysis can be performed using any commercially/freely available statistical analysis software such as R statistical software. According to certain embodiments, the correlation information is an expression that correlates the number of requests to the CPU utilization. According to some embodiments, this correlation information is XML serialized and inserted into a database along with other information that the power management solution requires. XML serialization is the process of converting a binary object in memory into an XML representation that can then be stored on disk (files or database). For purposes of simplicity, the statistical analysis is done in the background and the results are stored in the

database. The process of deriving correlations can be made real time, according to certain embodiments.

[0027] FIG. 2 is a block diagram showing some of the high-level steps for obtaining correlation information associated with the servers in the server pool, according to certain embodiments of the invention. At block 202, a fixed workload is executed against a given server that is to be characterized. At block 204, the information on resource utilization, workload and other related information is logged for analysis. At block 206, statistical analysis is performed on the information to obtain correlation information. At block 208, the correlation information for the given server is stored in the database. If the database already contains correlation information for the particular server, then the correlation information is updated. At block 210, the power manager retrieves correlation information for making power management decisions.

[0028] The power manager runs at regular intervals. For example, the power manager can run every 10 seconds. At each iteration of the power manager solution, a decision is taken as to whether a server needs to be powered on or powered off. The power manager also identifies which server must be powered on or off based on a server selection policy. The server selection policy is described herein.

[0029] FIG. 3 is a block diagram that illustrates the power management method, according to certain embodiments of the invention. After initialization at block 302, correlation data is retrieved from the database at block 304. At block 308, on each iteration, the power manager checks if all the servers in the server pool are above a pre-configured utilization threshold called the overload threshold, according to certain embodiments. According to certain other embodiments, the utilization threshold is determined dynamically rather than being pre-configured. If all the servers are above the utilization threshold, then at block 310, the power manager determines if all the servers in the server pool are powered on. If all the servers are powered on, then at block 306, the status quo of the server pool is maintained. If not all servers in the server pool are powered on, then at block 314, the power manager identifies which server is to be powered on, if more than one server is not powered on

in the server pool. At block 316, the power manager initiates power-on process for the selected server. At block 318, the power manager waits for the resume duration. At block 320, the power manager updates the state information for the selected server that was just powered on. At block 322, the server that was just powered on is marked on the application delivery controller as available for servicing requests.

[0030] If at block 308, it is determined that not all the servers in the server pool are below the utilization threshold, then at block 312 a check is made to identify if any server in the server pool can be powered off safely. If none of the servers in the server pool can be powered off safely, then the status quo is maintained at block 334.

[0031] If there are servers in the server pool can be powered off, then at block 324, the power manager identifies a server to be powered off. The server identified to be powered off is referred to as a candidate server. A decision to power off is taken only if the load on the candidate server can be migrated to the remaining power-on servers in the server pool without causing such remaining power-on servers to cross an overload threshold associated with a given server. At block 326, the server identified to be powered off is marked as unavailable on the application delivery controller. At block 328, the state information of the server to be powered off is updated. At block 330, the power manager waits for the number of requests sent to the server to be powered off drops to zero. At block 332, the power manager initiates the power-off process for the server to be powered off.

[0032] Powering servers on or off can be done using existing mechanisms supported by operating systems of the servers. For example, Windows Management Instrumentation (WMI) on Microsoft Windows or ssh based remote command execution on Linux/Solaris can be used for powering servers on or off.

[0033] According to certain embodiments, a staggered suspend and boot up process is used at a given point in time. In other words, exactly one server is suspending or resuming at a given time. The staggered suspend ensures that there is capacity in the server pool to handle any spikes in the

load and thus is a conservative approach. Staggered resume ensures that the load on the power supply for the server does not go high because computers typically draw higher power during the boot up phase.

[0034] According to certain embodiments, the power management method can include the following features:

- Predicting the demand: Historical data can be analysed to predict the demand that the server pool will experience in the next time interval. The prediction can augment the decisions taken by the power manager. Existing statistical methods like Auto Regressive Moving Average can be used for the time based trend analysis and prediction.
- Predict the number of servers or devices required to support a given workload.
- Chart the response time and performance of a server or a device under a given workload.
- Moving server nodes across pools: The power management method described herein can be extended to multiple pools using a global power management scheme. In such a global power management scheme, it is possible to move servers across pools to serve the needs of various pools. Depending on the demand, servers can be either moved across pools or turned on/off.

[0035] The Advanced Configuration and Power Interface (ACPI) specification defines the following server states, according to certain embodiments. Other suitable standards for defining server states can also be used. The embodiments are not restricted to the ACPI standard.

**Table 1: ACPI Server States**

<b>Server state</b>	<b>Global State</b>	<b>Description</b>
S0	G0	Server is powered on and operational.
S1 and S2	G1	Undefined and unused.

S3	G1	Suspended to RAM – Operating system context stored in RAM and most components powered down. Typically RAM and NIC are active in this state.
S4	G1	Suspend to Disk – Operating system context is written to disk and server is powered down.
S5	G2	Soft off – Server is powered down, no OS context is retained.
S5	G3	Mechanical off – Server is powered down and main power supply is cut off.

According to certain embodiments, servers are switched between S0 and S5.

[0036] If all the servers in the server pool have similar properties like operating frequency, RAM, disk space etc, the choice of server to shutdown/resume become trivial because any server can be chosen. However, if the server pools are heterogeneous pools, where servers differ in their properties, then a server selection policy is needed in order to select an appropriate server to power on or off. According to certain embodiments, policies that can be used to select servers if multiple servers are available for shutdown/resume are described below:

[0037] Policies for server power off include but are not limited to:

1. Lowest Frequency: Power off the server that operates at the lowest frequency.
2. Highest power: Power off the server that consumes the highest power.
3. Max post-utilization: Power off the server that will result in other servers having high utilization.
4. Types of applications running on the system (application capabilities).

[0038] The policies for server power on include but are not limited to:

1. Lowest power: Power on the server that consumes lowest power.

2. Highest frequency: Power on the server that runs at the highest frequency.
3. Shortest Resume Time: Power on the server that takes the shortest time to boot up.

[0039] As a non-limiting example, suspend policy 3 (max post-utilization) and resume policy 2 (highest frequency) can be used, according to certain embodiments. It is possible to support any combination of policies, but the power management mechanism must ideally be configured to use the ones that provide high power savings without significant loss in performance. Further, different combinations of suspend and resume policies will show different power/performance characteristics.

[0040] At any point in time, at least one server will be active in the pool. The reasoning behind having at least one server active is to have available computational capacity to handle requests while other servers are resuming.

[0041] As a non-limiting example, turning servers off is achieved by issuing a remote shutdown command using WMI (as our cluster is currently MS Windows based). Remote command execution requires that appropriate services are enabled on the server and appropriate ports are kept on in the firewall. Alternate techniques can be used for Linux and Solaris. Servers are turned on using Wake-On-LAN (WoL), an industry standard technique to resume computers that are currently suspended. A WoL packet is a specially crafted network packet which contains a WoL header and the MAC address of the target server repeated 16 times. WoL packet definition is standardized. WoL must be supported by the network interface card (NIC) and also enabled by the operating system driver. Modern NICs typically support WoL.

[0042] Such a non-intrusiveness approach does not require any additional software components to be installed on the individual servers in the server pool for the power manager to work. At most, it requires certain standard operating system services which might be turned off by default (like ssh, snmp) to be turned on.

[0043] FIG. 4 illustrates the class diagram of the central classes used for implementing the power manager, according to certain embodiments. FIG. 4 shows ServerMachine class 402, ServerLoadInformation class 404, Resource class 406, ResourceMeasureMethod class 408, ImmutableServerProperties class 410, UtilizationPredictor class 412, NetworkResource class 414, DiskResource class 416, and CPUResource class 418. The central data structure to the application is a ServerMachine class 402 that holds information about a server in the server cluster. The ServerMachine class contains the immutable server properties (like MAC address, maximum operating frequency, power consumption etc) and dynamically created objects for measuring resource utilization (see FIG. 7), predicting the utilization (FIG. 6), storing load information (FIG. 5) etc. A server contains resource objects – CPU, Disk, network, and memory, and is a resource in itself. The utilization predictor for each server is read from on disk storage (a database, for example) as an XML serialized stream and then de-serialized to get the object.

[0044] Some of the hierarchies of other classes used in the implementation of the power manager are described herein with reference to Figs. 5-8.

[0045] FIG. 5 illustrates the class diagram for the LoadInformation class hierarchy, according to certain embodiments. LoadInformation class defines classes that are used to store information on connections/requests queried at regular intervals from the load balancer or server. FIG. 5 shows that LoadInformation class 502 includes LocalHTTPLoadInfo class 504, PoolLoadInformation class 506, and ServerLoadInformation class 508. ServerLoadInformation class 508 includes VirtualServerLoadInformation class.

[0046] FIG. 6 illustrates the class diagram for the UtilizationPredictor class hierarchy, according to certain embodiments. UtilizationPredictor class 602 includes LinearRegressionBased class 604.

[0047] FIG. 7 illustrates the class diagram for the ResourcesMeasureMethod class hierarchy, according to certain embodiments. ResourcesMeasureMethod class 702 includes WMIAdaptor class 704 and SNMPAdaptor class 706.

[0048] FIG. 8 illustrates the class diagram for the LoadBalancer class hierarchy, according to certain embodiments. LoadBalancer class 802 includes F5Adaptor class 804. The load balancer class hierarchy is used to define classes that can be used to query and control the load balancer.

[0049] According to certain embodiments, a simple database with a single table is used to store information about individual servers.

[0050] The characterization phase requires utilization information to be gathered from servers for later analysis. According to certain embodiments, this information is stored in a database. The utilization information of each resource is stored in a separate file with the format shown in Table 2 Utilization Information, as non-limiting example.

Table 2: Utilization information

Date-Time stamp	Resource utilization (varying from 0-100%)	Weighted Moving average utilization (0-100%)
-----------------	--	--

[0051] The weighted moving average is used to help smoothing any sharp fluctuations in the measured utilization. An example for CPU utilization on a dual core machine, measured using WMI is given below.

Date-Time stamp	Core 0	Core 1	Total	Moving Avg (Core 0)	Moving Avg (Core )	Moving Avg (Total)
-----------------	--------	--------	-------	---------------------	--------------------	--------------------

[0052] The level of detail – per core utilization – is not provided by SNMP implementations. However, overall system utilization is available and the power manager implementation uses the overall utilization for analysis and decision making.

We Claim:

1. A method of power management, the method comprising:  
predicting resource utilization associated with each device from at least a subset of devices without installing customized software, firmware or hardware on the device for predicting the resource utilization;  
and  
selecting a candidate device from the subset of devices for powering on or off based on the predicted resource utilization and a set of power utilization criteria.
2. The method of Claim 1, wherein predicting resource utilization further comprises:  
measuring resource utilization on the device for a given number of requests executed by the device; and  
obtaining correlation information between the measured resource utilization and the number of requests executed by the device.
3. The method of Claim 1, further comprises:  
determining an overload threshold for the subset of devices; and  
determining if all power-on devices in the subset of devices are above the overload threshold;  
selecting a device from the subset of devices to power on if all power-on devices in the subset of are above the overload threshold;

selecting a device from the subset of devices to power off if all power-on devices in the subset of are below the overload threshold.

4. The method of Claim 3, wherein determining an overload threshold can be determined dynamically or can be pre-defined.
5. The method of Claim 3, wherein selecting a device to power off further comprises:
  - identifying a device, if any, to power off based on the set of power utilization criteria;
  - marking the identified device as unavailable;
  - updating state information for the identified device; and
  - initiating shut down for the identified device.
6. The method of Claim 3, wherein selecting a device to power on further comprises:
  - identifying a device, if any, to power on based on the set of power utilization criteria;
  - updating state information for the identified device; and
  - initiating power-on for the identified device.
7. The method of Claim 1, wherein the set of power utilization criteria comprises:
  - lowest operating frequency;

highest power consumption;  
maximum post-utilization;  
lowest power consumption;  
highest operating frequency;  
shortest resume time; and  
application capabilities.

8. The method of Claim 1, wherein the set of power utilization criteria provides high power savings without significant loss in performance.
9. The method of Claim 1, wherein the subset of devices are logically grouped to provide a common set of functionality.
10. The method of Claim 1, further comprising:  
predicting usage demand associated with the subset of devices; and  
using the predicted demand to augment selection of a candidate device  
from the subset of devices for powering on or off.
11. The method of Claim 10, further comprising analyzing historical usage data associated with the subset of devices.
12. The method of Claim 10, further comprising moving devices from the subset of devices to one or more other subsets of devices to optimize the power management across the subsets of devices.

13. The method of Claim 3, wherein determining if all power-on devices in the subset of devices are above the overload threshold is performed at periodic intervals.
14. A system for power management of a subset of devices, the system comprising:
  - at least one application delivery controller for receiving requests from client devices;
  - at least one power management machine for predicting resource utilization associated with each device from the subset of devices, wherein the subset of devices do not include customized software, firmware or hardware for predicting the resource utilization.
15. The system of Claim 14, wherein the least one power management machine interacts with the application delivery controller to obtain information on the total number of requests received by the application controller and the number of requests executed by each device in the subset of devices.
16. The system of Claim 14, wherein the least one power management machine remotely queries each device for information on the number of requests and resource utilization information in the subset of devices.

17. The system of Claim 14, further comprising at least one database server for storing correlation information that correlates the number of requests executed by each device and the resource utilization information associated with the execution of the requests for each device.
18. The system of Claim 14, wherein the at least one power management machine is used for:
  - determining an overload threshold for the subset of devices; and
  - determining if all power-on devices in the subset of devices are above the overload threshold;
  - selecting a device from the subset of devices to power on if all power-on devices in the subset of are above the overload threshold;
  - selecting a device from the subset of devices to power off if all power-on devices in the subset of are below the overload threshold.
19. The system of Claim 14, wherein the at least one power management machine uses a set of power utilization criteria for selecting a device from the subset of devices to power on or off, the power utilization criteria:
  - lowest operating frequency;
  - highest power consumption;
  - maximum post-utilization;
  - lowest power consumption;
  - highest operating frequency;

shortest resume time; and  
application capabilities.




FIG. 1

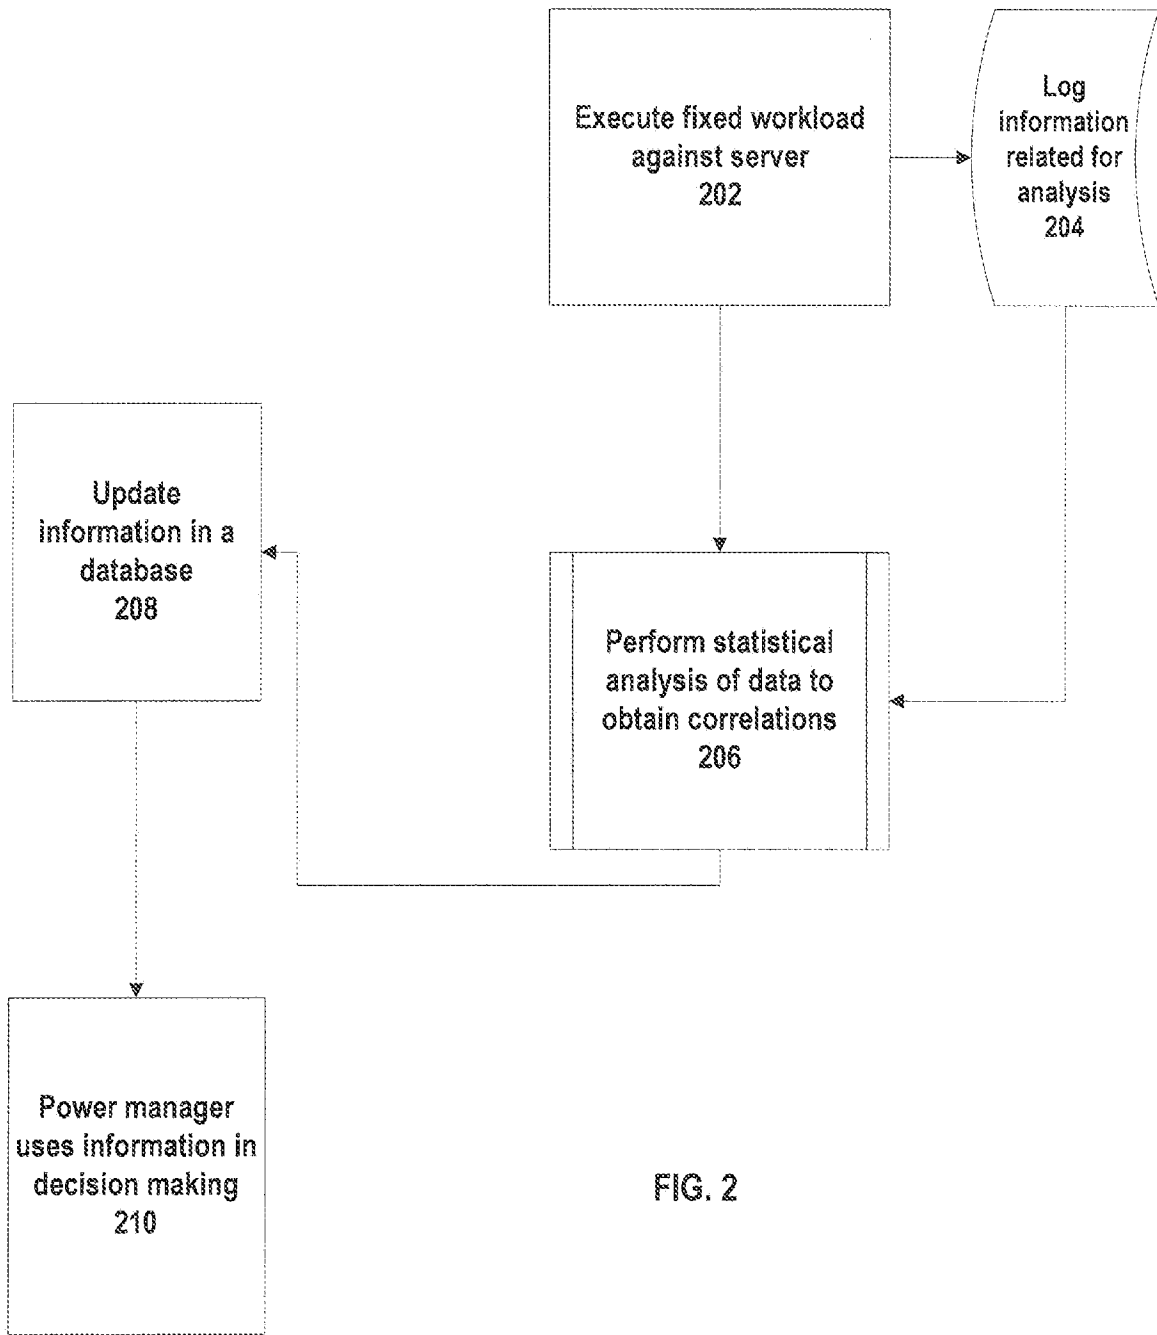


FIG. 2

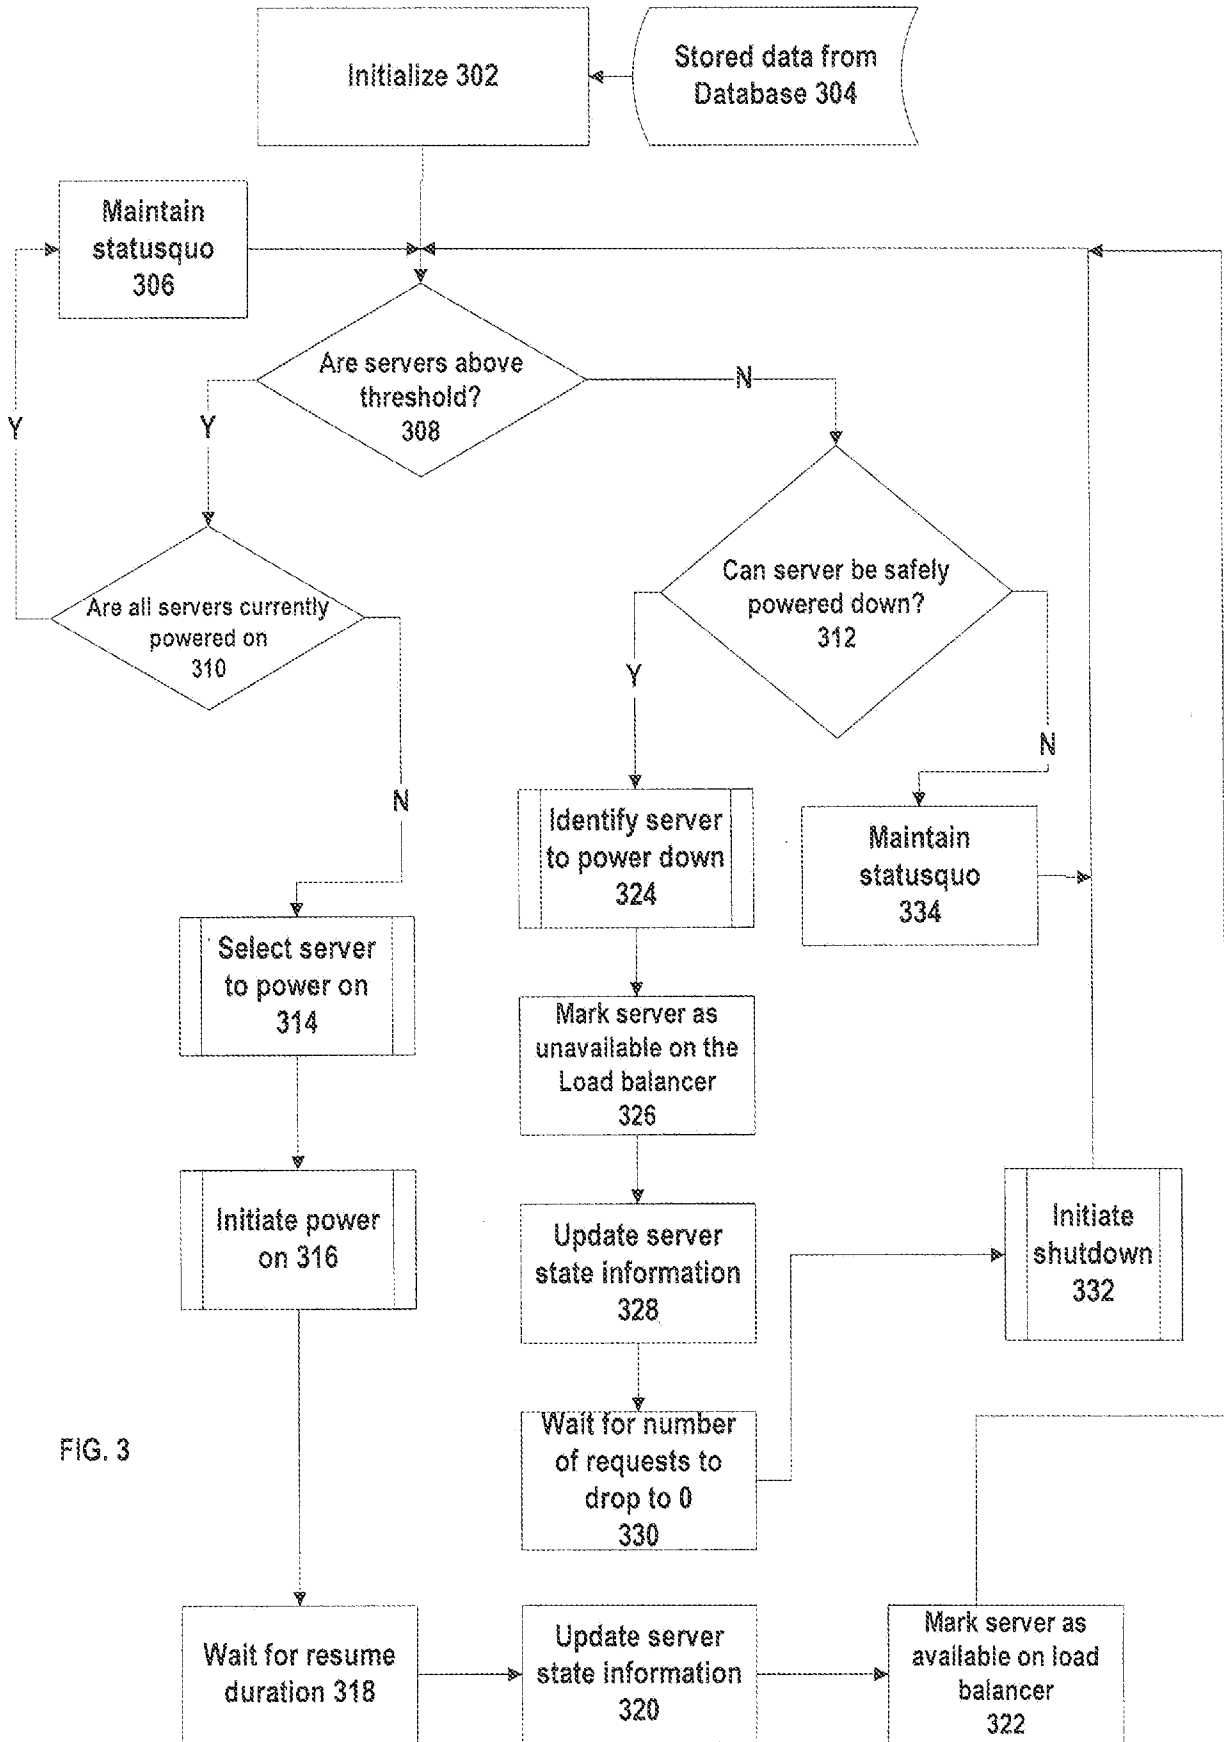
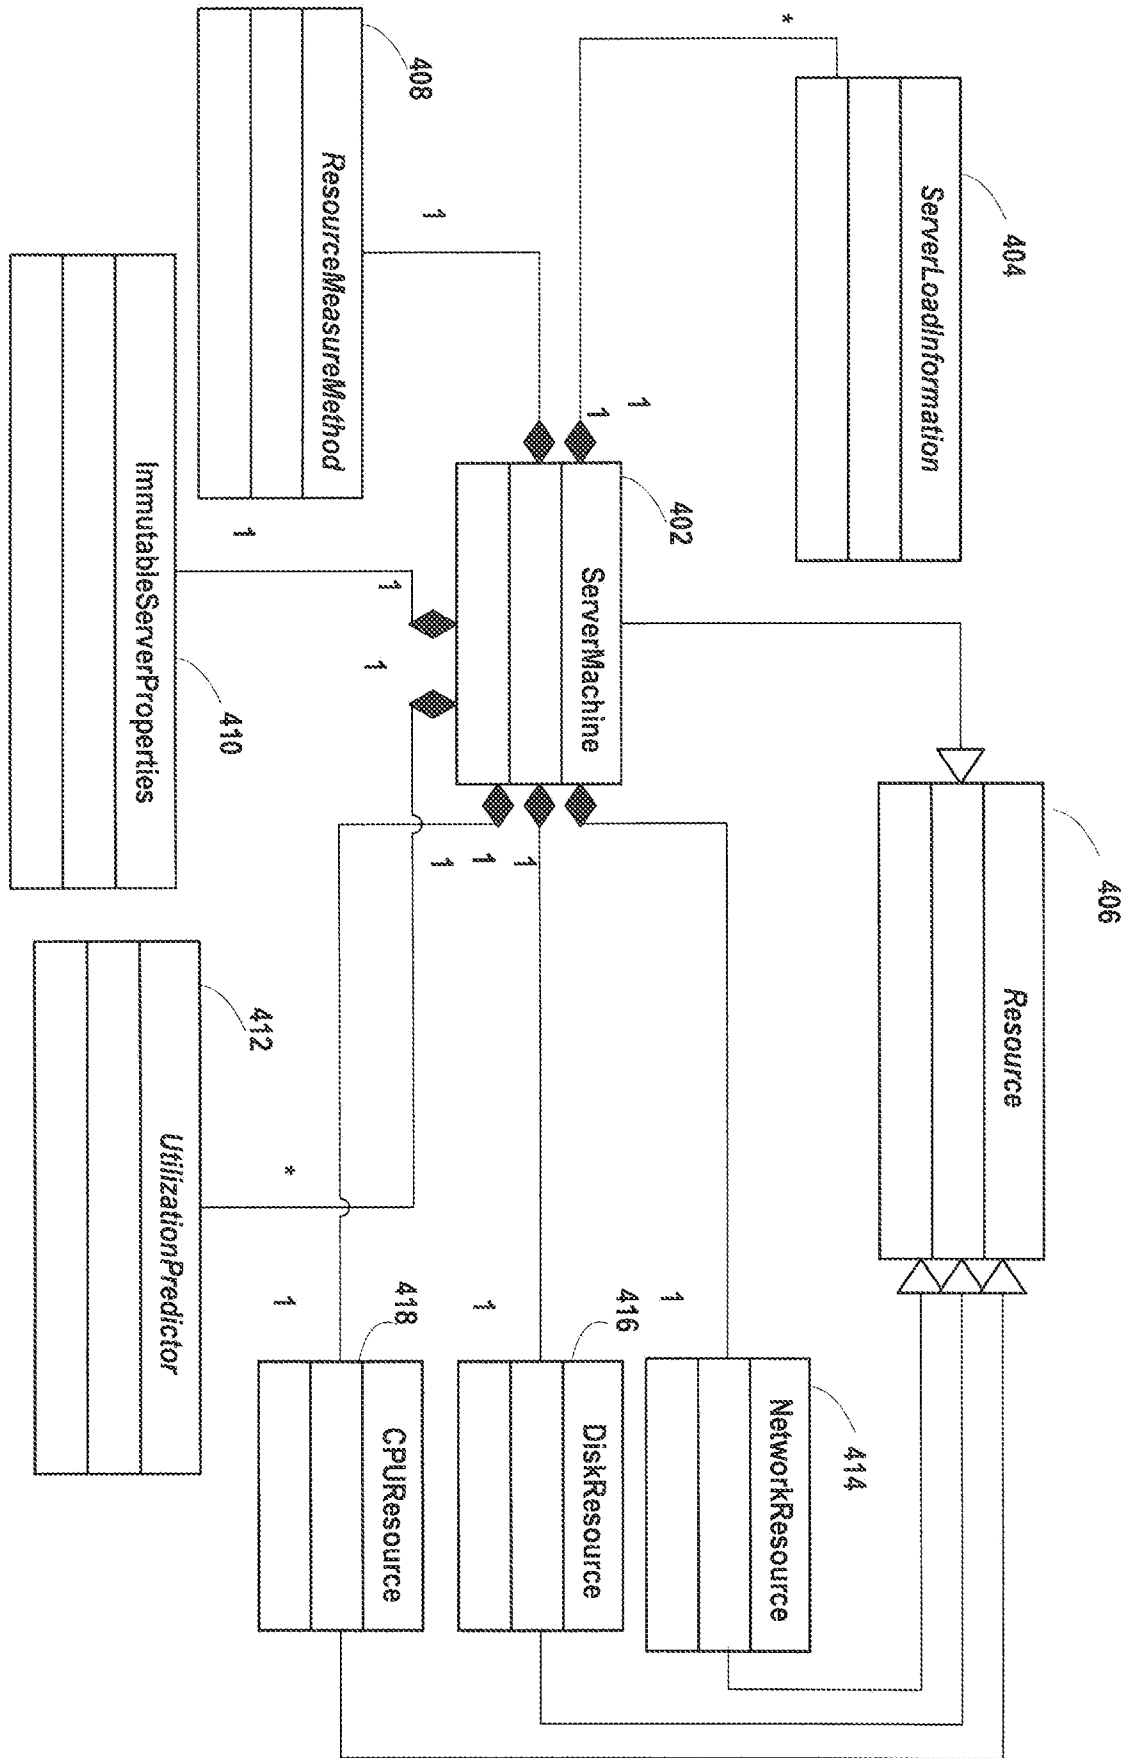
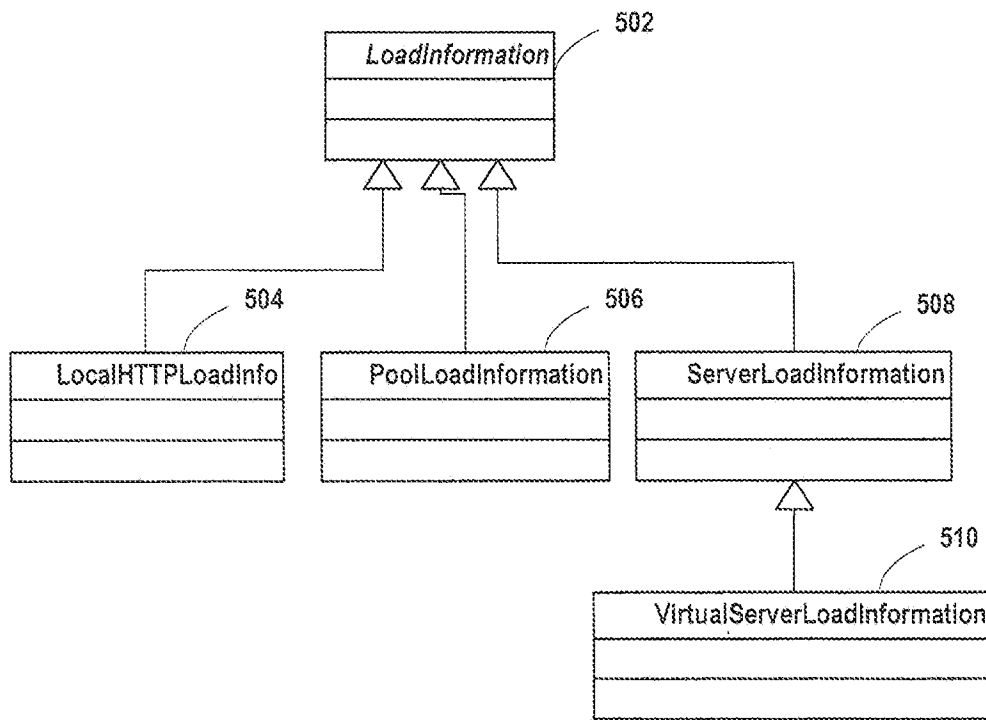


FIG. 3



Class Diagram


FIG. 4



LoadInformation Class Hierarchy


FIG. 5

FIG. 6




Utilization Predictor Class Hierarchy

FIG. 7



Resource Query Mechanism Class Hierarchy

FIG. 8



Load Balancer Class Hierarchy

**INTERNATIONAL SEARCH REPORT**

International application No.

PCT/US2012/027221

<p><b>A. CLASSIFICATION OF SUBJECT MATTER</b>                  IPC(8) - G06F 1/00 (2012.01)                  USPC - 713/300                  According to International Patent Classification (IPC) or to both national classification and IPC</p>																							
<p><b>B. FIELDS SEARCHED</b></p> <p>Minimum documentation searched (classification system followed by classification symbols)                  IPC(8) - G06F 1/00, G06F 1/26, G06F 1/32 (2012.01)                  USPC - 713/300, 713/340, 718/105</p> <p>Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched</p> <p>Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)                  MicroPatent, Google Patents, ACM</p>																							
<p><b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b></p> <table border="1" style="width:100%; border-collapse: collapse;"> <thead> <tr> <th style="width:10%;">Category*</th> <th style="width:70%;">Citation of document, with indication, where appropriate, of the relevant passages</th> <th style="width:20%;">Relevant to claim No.</th> </tr> </thead> <tbody> <tr> <td>X — Y</td> <td>US 2006/0184287 A1 (BELADY et al) 17 August 2006 (17.08.2006) entire document</td> <td>1, 8, 10-11, 14, and 17 <hr/>2-7, 9, 12-13, 15-16, and 18-19</td> </tr> <tr> <td>Y</td> <td>US 7,177,901 B1 (DUTTA) 13 February 2007 (13.02.2007) entire document</td> <td>2-6, 13, and 15</td> </tr> <tr> <td>Y</td> <td>US 2005/0273456 A1 (REVANURU et al) 08 December 2005 (08.12.2005) entire document</td> <td>3-6, 13, 16, and 18</td> </tr> <tr> <td>Y</td> <td>US 2007/0143454 A1 (MA et al) 21 June 2007 (21.06.2007) entire document</td> <td>9 and 12</td> </tr> <tr> <td>Y</td> <td>US 2011/0022868 A1 (HARCHOL-BALTER et al) 27 January 2011 (27.01.2011) entire document</td> <td>7 and 19</td> </tr> <tr> <td>Y</td> <td>US 2008/0082604 A1 (MANSOUR et al) 03 April 2008 (03.04.2008) entire document</td> <td>7 and 19</td> </tr> </tbody> </table>			Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.	X — Y	US 2006/0184287 A1 (BELADY et al) 17 August 2006 (17.08.2006) entire document	1, 8, 10-11, 14, and 17 <hr/> 2-7, 9, 12-13, 15-16, and 18-19	Y	US 7,177,901 B1 (DUTTA) 13 February 2007 (13.02.2007) entire document	2-6, 13, and 15	Y	US 2005/0273456 A1 (REVANURU et al) 08 December 2005 (08.12.2005) entire document	3-6, 13, 16, and 18	Y	US 2007/0143454 A1 (MA et al) 21 June 2007 (21.06.2007) entire document	9 and 12	Y	US 2011/0022868 A1 (HARCHOL-BALTER et al) 27 January 2011 (27.01.2011) entire document	7 and 19	Y	US 2008/0082604 A1 (MANSOUR et al) 03 April 2008 (03.04.2008) entire document	7 and 19
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.																					
X — Y	US 2006/0184287 A1 (BELADY et al) 17 August 2006 (17.08.2006) entire document	1, 8, 10-11, 14, and 17 <hr/> 2-7, 9, 12-13, 15-16, and 18-19																					
Y	US 7,177,901 B1 (DUTTA) 13 February 2007 (13.02.2007) entire document	2-6, 13, and 15																					
Y	US 2005/0273456 A1 (REVANURU et al) 08 December 2005 (08.12.2005) entire document	3-6, 13, 16, and 18																					
Y	US 2007/0143454 A1 (MA et al) 21 June 2007 (21.06.2007) entire document	9 and 12																					
Y	US 2011/0022868 A1 (HARCHOL-BALTER et al) 27 January 2011 (27.01.2011) entire document	7 and 19																					
Y	US 2008/0082604 A1 (MANSOUR et al) 03 April 2008 (03.04.2008) entire document	7 and 19																					
<p><input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/></p>																							
<p>* Special categories of cited documents:</p> <table style="width:100%;"> <tr> <td style="width:50%;"> <p>“A” document defining the general state of the art which is not considered to be of particular relevance</p> <p>“E” earlier application or patent but published on or after the international filing date</p> <p>“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)</p> <p>“O” document referring to an oral disclosure, use, exhibition or other means</p> <p>“P” document published prior to the international filing date but later than the priority date claimed</p> </td> <td style="width:50%;"> <p>“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</p> <p>“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</p> <p>“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</p> <p>“&amp;” document member of the same patent family</p> </td> </tr> </table>			<p>“A” document defining the general state of the art which is not considered to be of particular relevance</p> <p>“E” earlier application or patent but published on or after the international filing date</p> <p>“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)</p> <p>“O” document referring to an oral disclosure, use, exhibition or other means</p> <p>“P” document published prior to the international filing date but later than the priority date claimed</p>	<p>“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</p> <p>“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</p> <p>“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</p> <p>“&amp;” document member of the same patent family</p>																			
<p>“A” document defining the general state of the art which is not considered to be of particular relevance</p> <p>“E” earlier application or patent but published on or after the international filing date</p> <p>“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)</p> <p>“O” document referring to an oral disclosure, use, exhibition or other means</p> <p>“P” document published prior to the international filing date but later than the priority date claimed</p>	<p>“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</p> <p>“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</p> <p>“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</p> <p>“&amp;” document member of the same patent family</p>																						
<p>Date of the actual completion of the international search</p> <p>17 May 2012</p>		<p>Date of mailing of the international search report</p> <p align="center"><b>30 MAY 2012</b></p>																					
<p>Name and mailing address of the ISA/US</p> <p>Mail Stop PCT, Attn: ISA/US, Commissioner for Patents                  P.O. Box 1450, Alexandria, Virginia 22313-1450                  Facsimile No. 571-273-3201</p>		<p>Authorized officer:</p> <p align="center">Blaine R. Copenheaver</p> <p>PCT Helpdesk: 571-272-4300                  PCT OSP: 571-272-7774</p>																					