



(51) International Patent Classification:

G06F 9/455 (2006.01) *G06F 17/50* (2006.01)
G06F 9/50 (2006.01) *H04W 72/08* (2009.01)
G06F 15/173 (2006.01) *H04W 72/10* (2009.01)

(21) International Application Number:

PCT/US2016/041426

(22) International Filing Date:

8 July 2016 (08.07.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

201510404937.X 10 July 2015 (10.07.2015) CN
15/204,623 7 July 2016 (07.07.2016) US

(71) Applicant: ALIBABA GROUP HOLDING LIMITED

[—/US]; Fourth Floor, One Capital Place, P. O. Box 847,
Grand Cayman (KY).

(72) Inventors: ZHANG, Yunyang; c/o Alibaba Group Legal

Department, 5/f, Building 3, No. 969 West Wen Yi Road,
Yu Hang District, Hangzhou, 311121 (CN). YAN, Huin-

ing; c/o Alibaba Group Legal Department, 5/f, Building 3,
No. 969 West Wen Yi Road, Yu Hang District, Hangzhou,
311121 (CN). LI, Meng; c/o Alibaba Group Legal Depart-
ment, 5/f, Building 3, No. 969 West Wen Yi Road, Yu
Hang District, Hangzhou, 311121 (CN). XU, Hongbin; c/o
Alibaba Group Legal Department, 5/f, Building 3, No. 969
West Wen Yi Road, Yu Hang District, Hangzhou, 311121
(CN).

(74) Agent: DeCARLO, James J.; Greensberg Traurig, LLP,
200 Park Avenue, 38th fl., New York, NY 10166 (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, BN, 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, IR, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG,
MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,
PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, 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.

[Continued on next page]

(54) Title: METHOD AND DEVICE FOR COMPUTING RESOURCE SCHEDULING




FIG. 1

(57) Abstract: A method, apparatus and device for scheduling resources of a cluster comprising a plurality of hosts, each running at least one instance, including acquiring a resource parameter of the cluster; calculating the number of predicted hosts in the cluster according to the resource parameter; determining to-be-migrated hosts and target hosts from the current hosts in the cluster when the number of current hosts in the cluster is greater than the number of predicted hosts; and migrating instances running on the to-be-migrated host to the target host.



(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, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, 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, KM, ML, MR, NE, SN, TD, TG).

Published:

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

METHOD AND DEVICE FOR COMPUTING RESOURCE SCHEDULING

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of priority from Chinese Patent Application No. 201510404937 filed on July 10, 2015, entitled “Method, Apparatus and Equipment for Resource Distribution,” and U.S. Patent Application No. 15/204,623 filed July 7, 2016, entitled “Method and Device for Computing Resource Scheduling,” both of which are incorporated herein in their entirety by reference.

BACKGROUND

Field of the Disclosure

[0002] The present application relates to the field of cloud computing, and in particular, relates to a method and device for resource scheduling.

Description of the Related Art

[0003] Cloud computing is an Internet-based computing approach. Generally, a cloud computing service provider may divide the resources on a host (which may also be referred to as a cloud server) according to different demands of different users, so as to allocate different sizes of resources to different users. The resources may comprise a central processing unit (CPU), a memory, a storage space, and the like. The resources allocated to a user form an instance of the user, which may also be referred to as a virtual machine. Typically, resources on one host may be divided into a plurality of instances for a plurality of users to use.

[0004] In the prior art, the cloud computing service provider generally configures a cluster comprising a plurality of hosts. At least one instance of the user runs on each host in the cluster. However, since the resource utilization rates of different instances are different, the resources of the host are not evenly allocated. As a result, the overall utilization rate of the cluster is imbalanced.

SUMMARY

[0005] According to some embodiments, the present disclosure provides a method and device for resource scheduling. According to some embodiments, the disclosed methods and devices solve the problem that the overall resource utilization of a cluster is imbalanced in the prior art.

[0006] In a non-limiting embodiment of the disclosure, a method for resource scheduling is provided. The method is used for scheduling resources of a cluster comprising a plurality of current hosts, wherein each of the plurality of current hosts runs at least one instance, and the method includes: acquiring a resource parameter of the cluster; calculating a number of predicted hosts in the cluster according to the resource parameter; determining a to-be-migrated host and a target host from the current hosts in the cluster when the number of current hosts in the cluster is greater than the number of predicted hosts; and migrating an instance running on the to-be-migrated host to the target host.

[0007] According to another non-limiting embodiment of the disclosure, an apparatus for resource scheduling is provided. The apparatus is used for scheduling resources of a cluster including a plurality of current hosts, wherein each of the plurality of current hosts runs at least one instance, and the apparatus includes: an acquiring unit, configured to acquire a resource parameter of the cluster; a calculating unit, configured to calculate the number of predicted hosts in the cluster according to the resource parameter; a determining unit, configured to determine a to-be-migrated host and a target host from the current hosts in the cluster when the number of current hosts in the cluster is greater than the number of predicted hosts; and a scheduling unit, configured to migrate an instance running on the to-be-migrated host to the target host.

[0008] In yet another non-limiting embodiment of the disclosure, a system for resource scheduling is provided, wherein the system includes: a processor; and a memory storing instructions executable by the processor that cause the processor to: acquire a resource parameter of the cluster; calculate the number of predicted hosts in the cluster according to the resource

parameter; determine a to-be-migrated host and a target host from the current hosts in the cluster when the number of current hosts in the cluster is greater than the number of predicted hosts; and migrate an instance running on the to-be-migrated host to the target host.

[0009] According to some embodiments of the present disclosure, during scheduling of resources in a cluster, by acquiring a resource parameter of the cluster, calculating the number of predicted hosts in the cluster, determining to-be-migrated hosts and target hosts from the current hosts in the cluster when the current number of hosts in the cluster is greater than the number of predicted hosts, and migrating instances running on the to-be-migrated hosts to the target hosts, each target host achieves a maximum improvement on a resource utilization rate, such that the overall resource utilization rate of the cluster is more balanced while the resource demands of the cluster are satisfied with a minimum migration cost.

[0010] It shall be appreciated that the above general description and the detailed description hereinafter are only illustrative and interpretative, and are not intended to limit the present disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The foregoing and other objects, features, and advantages of the disclosure will be apparent from the following description of embodiments as illustrated in the accompanying drawings, in which reference characters refer to the same parts throughout the various views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating principles of the disclosure.

[0012] FIG. 1 is a schematic structural diagram illustrating architecture of a cloud system for resource scheduling according to certain embodiments of the disclosure.

[0013] FIG. 2 is a flow diagram illustrating a method for resource scheduling according to certain embodiments of the disclosure.

[0014] FIG. 3 is a flow diagram illustrating a method for resource scheduling according to certain embodiments of the disclosure.

[0015] FIG. 4 is a schematic structural diagram illustrating hardware of a device for resource scheduling according to certain embodiments of the disclosure.

[0016] FIG. 5 is a block diagram illustrating a device for resource scheduling according to certain embodiments of the disclosure.

[0017] FIG. 6 is a block diagram illustrating a device for resource scheduling according to certain embodiments of the disclosure.

[0018] FIG. 7 is a block diagram illustrating an acquiring unit according to certain embodiments of the disclosure.

[0019] FIG. 8 is a block diagram illustrating a calculating unit according to certain embodiments of the disclosure.

[0020] FIG. 9 is a block diagram illustrating a determining unit according to certain embodiments of the disclosure.

[0021] FIG. 10 is a block diagram illustrating a scheduling unit according to certain embodiments of the disclosure.

DETAILED DESCRIPTION

[0022] The terminology used in the present application is for the purpose of describing particular embodiments, and is not intended to limit the scope of the present disclosure, which is defined by the appended claims. As used in the present application and the appended claims, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It shall also be understood that the term "and/or" used herein is intended to signify and include any or all possible combinations of one or more of the associated listed items.

[0023] It shall be understood that, although the terms first, second, third, etc. may be used herein to describe various information, the information should not be limited by these terms.

These terms are only used to distinguish one category of information from another. For example, without departing from the scope of the present disclosure, first information may be termed as second information; and similarly, second information may also be termed as first information. As used herein, the term "if" may be understood to mean "when" or "upon" or "in response to determining" or "in response to detecting" depending on the context. FIG. 1 is a schematic structural diagram illustrating an architecture of a cloud system for resource scheduling according to certain embodiments of the disclosure.

[0024] The architecture of the cloud system as illustrated in FIG. 1 is described by using one cluster as an example. In practical application, a cloud computing server provider may be configured with a plurality of clusters. The cluster as illustrated in FIG. 1 includes M hosts. The resources on each host can be described by using storage space as an example. The resources on each host can be allocated to different users. Using host 1 as an example, the storage space on host 1 is divided into N parts, which are provided to N users. The resource allocated to each user from host 1 is referred to as an instance of the user. As illustrated in FIG. 1, N instances, for example, can run on host 1.

[0025] With reference to FIG. 1, according to some embodiments of the present disclosure, during scheduling of resources in a cluster, by acquiring a resource parameter of the cluster, calculating the number of predicted hosts in the cluster, determining one or more to-be-migrated hosts and one or more target hosts from the current hosts in the cluster when the current number of hosts in the cluster is greater than the number of predicted hosts, and migrating instances running on the to-be-migrated hosts to the target hosts, each target host achieves the maximum improvement on its resource utilization rate, such that the overall resource utilization rate of the cluster is more balanced while the resource demands of the cluster are satisfied with a minimum migration cost. The present disclosure is described in further detail hereinafter with reference to a number of specific, interchangeable embodiments.

[0026] FIG. 2 is a flow diagram illustrating a method for resource scheduling according some embodiments of the disclosure.

[0027] At step S201, a resource parameter of the cluster is acquired. The resource parameter of the cluster in this step may include, for example, a resource configuration quantity of the host and actual resource utilization quantities of the instances running on the host.

[0028] To ensure consistency of the hosts in the cluster, generally the hosts in the cluster have the same resource configuration quantity. The resource configuration quantity is an inherited configuration parameter of the host. According to the resource types, the resource configuration quantity may comprise a CPU parameter, a memory capacity, a storage space capacity, and the like. When the resources of the cluster need to be scheduled, the configuration parameter of the hosts in the cluster may be directly collected as the resource configuration quantity.

[0029] Corresponding to different types of resources on the hosts, the actual resource utilization quantity of each instance running on each host may be collected. For example, an actual storage space utilization quantity, an actual memory utilization quantity, and the like, of each instance may be collected.

[0030] At step S202, the number of predicted hosts in the cluster is calculated according to the resource parameter of the cluster.

[0031] With respect to the resource configuration quantity of the host and the actual resource utilization quantity of the instances running on the host that are acquired in step S201, in this step, a predicted resource utilization quantity of each instance may be acquired by performing a linear regression analysis on the actual resource utilization quantity of each instance. Linear regression analysis is a mathematical statistics approach, and is an analytical method to determine a mutually-dependent quantitative relationship between two or more variables. In these embodiments, by means of the linear regression analysis, a maximum value of the resource

utilization quantity in a prospective month, for example, may be predicted according to the resource utilization quantity of the instance on each day within a past month, for example, and used as the predicted resource utilization quantity of the instance.

[0032] After the predicted resource utilization quantity of each instance is acquired, a sum of the predicted resource utilization quantities of all the instances on each host may be used as the predicted resource utilization quantity of each host. A sum of the predicted resource utilization quantity of all the hosts in the cluster may be calculated, and then used as the predicted resource utilization quantity of the cluster. Next, the number of predicted hosts in the cluster is calculated according to the resource configuration quantities of the hosts and the predicted resource utilization quantity of the cluster. Still using storage space as an example, assume that the storage space of a host is 128 GB and a predicted storage space utilization quantity of the cluster is 300 GB, in order to satisfy the 300 GB storage space utilization quantity, at least three hosts are needed, and the number of predicted hosts in the cluster is 3.

[0033] At step S203, to-be-migrated hosts and target hosts are determined from the current hosts in the cluster when the number of current hosts in the cluster is greater than the number of predicted hosts.

[0034] The current number of hosts in the cluster is the number of hosts online in the cluster. When the current number of hosts is greater than the number of predicted hosts calculated in step S202, one or more of the hosts in the cluster are redundant, and host resources may be scheduled. In this case, a difference between the current number of hosts in the cluster and the number of predicted hosts may be calculated, and the difference may be used as a number of to-be-migrated hosts in the cluster.

[0035] Next, a migration cost value of each host in the cluster may be determined. For example, the number of instances running on each host and the total storage space of each instance may be acquired, a notification migration cost value of each host may be calculated according to the number

of instances, a data migration cost value of all the instances on each host may be calculated according to the total storage space, and a sum of the notification migration cost value and the data migration cost value is used as the migration cost value of each host. Finally, the calculated number of to-be-migrated hosts is acquired according to an order of the migration cost values, and a same number of remaining hosts as the number of predicted hosts are used as the target hosts.

[0036] In step S204, instances running on the to-be-migrated hosts are migrated to the target hosts.

[0037] After the to-be-migrated hosts and the target hosts are determined, the to-be-migrated instances running on the to-be-migrated hosts may be sequenced in a descending order of the predicted resource utilization quantities of the instances. To-be-migrated instances may be sequentially migrated to a target migration host in the target hosts according to the sequence, wherein the target migration host is a target host whose predicted resource utilization quantity achieves a maximum improvement after the to-be-migrated instance is migrated. The predicted resource utilization rate is a ratio of the predicted resource utilization quantity on the target host to the resource configuration quantity. The improvement of the predicted resource utilization rate is an improvement in the proportion of the predicted resource utilization rate of the host after a to-be-migrated instance is migrated to the target host.

[0038] As seen from the above, in these embodiments, by acquiring a resource parameter of the cluster, calculating the number of predicted hosts in the cluster, determining to-be-migrated hosts and target hosts from the current hosts in the cluster when the current number of hosts in the cluster is greater than the number of predicted hosts, and migrating instances running on the to-be-migrated hosts to the target hosts, each target host achieves a maximum improvement on resource utilization rate, such that the overall resource utilization rate of the cluster is more balanced while the resource demands of the cluster are satisfied with a minimum migration cost.

[0039] FIG 3 is a flow diagram illustrating a method for resource scheduling according to some embodiments of the disclosure. Resource scheduling is described in detail in these embodiments.

[0040] At step S301, a resource configuration quantity of the host in the cluster is collected.

[0041] To ensure consistency of the hosts in the cluster, generally the hosts in the cluster have the same resource configuration quantity. The resource configuration quantity is an inherited configuration parameter of the host, generally referring to a hardware configuration parameter. According to the resource types, the resource configuration quantity may comprise a CPU parameter, a memory capacity, a storage space capacity, and the like. When the resources of the cluster need to be scheduled, the configuration parameter of the hosts in the cluster may be directly collected as the resource configuration quantity.

[0042] At step S302, an actual resource utilization quantity of the instance running on each host in the cluster is collected.

[0043] Corresponding to different types of resources on the hosts, the actual resource utilization quantity of each instance running on each host may be collected. For example, an actual storage space utilization quantity, an actual memory utilization quantity, and the like, of each instance may be collected.

[0044] At step S303, the number of predicted hosts in the cluster is calculated according to the resource configuration quantities of the hosts and the actual resource utilization quantities of the instances.

[0045] With respect to the resource configuration quantity of the host and the actual resource utilization quantity of the instances running on the host that are acquired in step S301 and step S302, respectively, in this step, a predicted resource utilization quantity of each instance may be acquired by performing a linear regression analysis on the actual resource utilization quantity of each instance. Using storage space utilization as an example, a storage space utilization on each day within a past one month may be acquired, and a linear regression analysis may be performed

on the acquired storage space utilization to calculate a maximum value of the storage space utilization in a prospective month. The maximum value may be used as the predicted storage space utilization quantity of the instance. Using memory utilization as another example, a memory utilization on each day within a past month may be acquired, and a linear regression analysis may be performed on the acquired memory utilization to calculate a maximum value of the memory utilization in a prospective month. The maximum value may be used as the predicted memory utilization quantity of the instance.

[0046] After the predicted resource utilization quantity of each instance is acquired, a sum of the predicted resource utilization quantities of all the instances on each host may be used as the predicted resource utilization quantity of each host. A sum of the predicated resource utilization quantities of all the hosts in the cluster may be calculated and then used as the predicted resource utilization quantity of the cluster. Next, the number of predicted hosts in the cluster is calculated according to the resource configuration quantities of the hosts and the predicted resource utilization quantities of the cluster. Still using storage space as an example, assume that the storage space of a host is 128 GB and a predicted storage space utilization quantity of the cluster is 300 GB. In order to satisfy the 300 GB storage space utilization quantity, at least three hosts are needed, and the number of predicted hosts in the cluster is 3. It should be noted that when there are a plurality of resource types, the number of predicted hosts may be calculated with respect to each resource type, and a maximum number may be selected from a plurality of predicted numbers of hosts as the number of predicted hosts in the cluster. For example, if the number of predicted hosts acquired according to the storage space is 3, the number of predicted hosts acquired according to the CPU is 2, and the number of predicted hosts acquired according to the memory is 4, then it is determined that the number of predicted hosts in the cluster is 4.

[0047] At step S304, it is determined whether the current number of hosts in the cluster is greater than the number of predicted hosts. If the current number of hosts is greater than the number of predicted hosts, then step S305 is performed; otherwise, step S308 is performed.

[0048] In these embodiments, the current number of hosts in the cluster is the number of all hosts online in the cluster. When the current number of hosts is greater than the number of predicted hosts calculated in step S303, one or more hosts in the cluster are redundant, and the host resources may be scheduled. When the current number of hosts is smaller than the number of predicted hosts, the hosts in the cluster are insufficient, and new hosts need to be added.

[0049] At step S305, the to-be-migrated hosts and the target hosts are determined from the current hosts in the cluster.

[0050] When it is judged that the number of current hosts is greater than the number of predicted hosts, the difference between the number of current hosts in the cluster and the number of predicted hosts may be calculated. The difference may be used as the number of to-be-migrated hosts in the cluster. Next, a migration cost value of each host in the cluster may be determined. For example, the number of instances running on each host and the total storage space of each instance may be acquired, a notification migration cost value of each host may be calculated according to the number of instances, a data migration cost value of all the instances running on each host may be calculated according to the total storage space, and a sum of the notification migration cost value and the data migration cost value may be used as the migration cost value of each host.

[0051] In an example of calculating the migration cost value of a host, the notification cost and data migration cost may be firstly determined, wherein the notification cost is the manpower cost consumed to notify the user corresponding to each to-be-migrated instance of a migration event, and the data migration cost is a sum of the network traffic transmission cost consumed to migrate each 1 GB of data and the loss cost caused due to input/output (I/O) operations performed for the hard disk when reading and writing data in the process of migrating the data on the to-be-migrated host to the target host. Assume that the notification cost is 1 yuan/instance, the number of instances is b , the data

migration cost is c yuan/GB and the total storage space is d GB, then the migration cost of the host is $a*b+c*d$.

[0052] After the migration cost value of each host is determined, the to-be-migrated hosts may be acquired according to an order of the migration cost values. The number of acquired to-be-migrated hosts is the above calculated number of to-be-migrated hosts. Then, the remaining hosts are used as the target hosts, wherein the number of target hosts is consistent with the number of predicted hosts.

[0053] At step S306, the instances running on the to-be-migrated hosts are migrated to the target hosts, such that each target host achieves a maximum improvement on the resource utilization rate.

[0054] After the to-be-migrated hosts and the target hosts are determined, the to-be-migrated instances running on the to-be-migrated host may be sequenced in a descending order of the predicted resource utilization quantities of the instances. To-be-migrated instances may be sequentially migrated to a target migration host in the target hosts according to the sequence, wherein the target migration host is a target host whose predicted resource utilization quantity achieves maximum improvement after the to-be-migrated instance is migrated to the target host. The predicted resource utilization rate of the target host is a ratio of the predicted resource utilization quantity on the target host to the resource configuration quantity. The improvement of the predicted resource utilization rate is an improvement in the proportion of the predicted resource utilization rate of the host after a to-be-migrated instance is migrated to the target host.

[0055] In some embodiments of the disclosure, to ensure normal running of each host, a resource utilization rate threshold may be preset for the host to ensure that the actual resource utilization rate of the host is smaller than the preset utilization rate threshold. Therefore, in these embodiments, with respect to the target hosts, the actual resource utilization rate of each target host may be calculated in

advance, the target host whose actual resource utilization rate exceeds the preset utilization rate threshold is determined as a first target host, and the remaining target hosts are determined as second target hosts; a first to-be-migrated instance is acquired from the first target host, such that the actual resource utilization rate of the first target host when the first to-be-migrated instance is migrated is smaller than the preset utilization rate threshold; and the first to-be-migrated instance is migrated to the second target host, such that the second target host achieves maximum improvement rate on the resource utilization rate.

[0056] At step S307, after the instances running on the to-be-migrated host are migrated to the target host, the to-be-migrated host is controlled to go offline. Then, the method ends.

[0057] At step S308, new hosts are added to the cluster, wherein the number of new hosts is equal to the difference between the number of predicted hosts and the number of current hosts.

[0058] When it is determined that the number of current hosts is smaller than the number of predicted hosts, the difference between the number of predicted hosts in the cluster and the number of current hosts may be calculated and then used as the number of new hosts to be added to the cluster.

[0059] At step S309, a to-be-migrated host whose actual resource utilization rate exceeds a preset utilization rate threshold is acquired from the current hosts in the cluster, and at least one instance running on the to-be-migrated host is migrated to the new host such that neither of the actual resource utilization rates of the to-be-migrated host and the new host exceeds the preset utilization rate threshold. Then, the method ends.

[0060] As seen from the above, in these embodiments, by acquiring a resource parameter of the cluster, calculating the number of predicted hosts in the cluster, determining to-be-migrated hosts and target hosts from the current hosts in the cluster when the current number of hosts in the cluster is greater than the number of predicted hosts, and migrating instances running on the to-be-migrated hosts to the target hosts, each target host achieves maximum improvement on

resource utilization rate, such that the overall resource utilization rate of the cluster is more balanced while the resource demands of the cluster are satisfied with a minimum migration cost.

[0061] Corresponding to the embodiments illustrating the method for resource scheduling according to the present application, the present application further provides embodiments illustrating an apparatus and device for resource scheduling.

[0062] The embodiment illustrating the device for resource scheduling is used for describing scheduling of resources of a cluster comprising a plurality of hosts, wherein at least one instance runs on each of the plurality of hosts. The apparatus embodiment may be practiced via software, or may be practiced via a combination of hardware and software. Using practice via the software as an example, as a logical apparatus, the apparatus for resource scheduling is formed by reading the corresponding computer program instructions from a non-volatile memory to a memory for execution by using a processor of the device where the apparatus is configured. From the perspective of hardware, FIG. 4 is a schematic structural diagram illustrating hardware of a device where the apparatus for resource scheduling is configured according to the present application. In addition to the processor, memory, network interface, and the non-volatile memory illustrated in FIG. 4, the device where the apparatus according to the embodiment is configured may typically further comprise other hardware according to the actual functions of the device, which is not further listed in FIG. 4.

[0063] FIG. 5 is a block diagram illustrating an apparatus for resource scheduling according to some embodiments of the disclosure. The apparatus includes an acquiring unit 510, a calculating unit 520, a determining unit 530, and a scheduling unit 540.

[0064] The acquiring unit 510 is configured to acquire a resource parameter of the cluster.

[0065] The calculating unit 520 is configured to calculate the number of predicted hosts in the cluster according to the resource parameter.

[0066] The determining unit 530 is configured to determine a to-be-migrated host and a target host from the current hosts in the cluster when the number of current hosts in the cluster is greater than the number of predicted hosts.

[0067] The scheduling unit 540 is configured to migrate an instance running on the to-be-migrated host to the target host.

[0068] FIG. 6 is a block diagram illustrating an apparatus for resource scheduling according to some embodiments of the disclosure. The apparatus includes an acquiring unit 610, a calculating unit 620, a determining unit 630, a scheduling unit 640, a control unit 650, and an adding unit 660.

[0069] The acquiring unit 610 is configured to acquire a resource parameter of the cluster;

[0070] The calculating unit 620 is configured to calculate the number of predicted hosts in the cluster according to the resource parameter.

[0071] The determining unit 630 is configured to determine to-be-migrated hosts and target hosts from the current hosts in the cluster when the number of current hosts in the cluster is greater than the number of predicted hosts.

[0072] The scheduling unit 640 is configured to migrate an instance running on the to-be-migrated host to the target host.

[0073] The control unit 650 is configured to control the to-be-migrated host to go offline from the cluster when the instance running on the to-be-migrated host is migrated to the target host.

[0074] The adding unit 660 is configured to add a new host to the cluster when the current number of hosts in the cluster is smaller than the number of predicted hosts, the number of new hosts being a difference between the number of predicted hosts and the current number of hosts.

[0075] The scheduling unit 640 is further configured to acquire a third host whose actual resource utilization rate exceeds a preset utilization rate threshold from the current hosts in the cluster,

and migrate at least one instance running on the third host to the new host, such that neither of the actual resource utilization rates of the third host and the new host exceeds the preset utilization rate threshold.

[0076] According to some embodiments, the acquiring unit 610 may include (FIG. 7):

[0077] a first collecting unit 611, configured to collect a resource configuration quantity of the host in the cluster; and

[0078] a second collecting unit 612, configured to collect an actual resource utilization quantity of the instance running on each host in the cluster.

[0079] According to some embodiments, the calculating unit 620 may include (FIG. 8):

[0080] a prediction analyzing subunit 621, configured to perform linear regression analysis on the actual resource utilization quantity of each instance running on the host, to acquire a predicted resource utilization quantity of each instance;

[0081] a utilization quantity calculating subunit 622, configured to calculate a sum of the predicted resource utilization quantities of all the instances running on each host, and use the sum as the predicted resource utilization quantity of each host; and calculate a sum of the predicted resource utilization quantities of all the hosts in the cluster, and use the sum as the predicted resource utilization quantity of the cluster; and

[0082] a predicted host number calculating subunit 623, configured to calculate the number of predicted hosts in the cluster according to the resource configuration quantities of the hosts and the predicted resource utilization quantity of the cluster.

[0083] According to some embodiments, the determining unit 630 may include (FIG. 9):

[0084] a migration host number calculating subunit 631, configured to calculate the difference between the current number of hosts in the cluster and the number of predicted hosts, and use the difference as the number of to-be-migrated hosts in the cluster;

[0085] a migration cost determining subunit 632, configured to determine a migration cost value of each host in the cluster; and

[0086] a host determining subunit 633, configured to acquire a same number of to-be-migrated hosts as the number of the migrated hosts according to a descending order of the migration cost values, and use a same number of remaining hosts as the number of predicted hosts as the target hosts.

[0087] According to some embodiments, the migration cost determining circuitry may be specifically configured to: acquire the number of instances running on each host and a total storage space of each instance; calculate a notification migration cost value of each host according to the number of instances, and calculate a data migration cost value of all the instances on each host according to the total storage space; and calculate a sum of the notification migration cost value and the data migration cost value, and use the sum as the migration cost value of each host.

[0100] According to some embodiments, the scheduling unit 640 may include (FIG. 10):

[0101] an instance sequencing subunit 641, configured to sequence all the to-be-migrated instances running on the to-be-migrated host according to a descending order of the predicted resource utilization quantities of the instances; and

[0102] a sequentially migrating subunit 642, configured to sequentially migrate a to-be-migrated instance to a target migration host in the target hosts according to the sequence, the target migration host being a target host whose predicted resource utilization quantity achieves maximum improvement after the to-be-migrated instance is migrated to the target host.

[0103] According to some embodiments, the scheduling unit 640 may further include:

[0104] a utilization rate calculating subunit 643, configured to calculate an actual resource utilization rate of each target host;

[0105] a target host determining subunit 644, configured to determine the target host whose actual resource utilization rate exceeds a preset utilization rate threshold as a first target host, and determine the remaining target hosts as second target hosts;

[0106] a to-be-migrated instance acquiring subunit 645, configured to acquire a first to-be-migrated instance from the first target host, such that the actual resource utilization rate of the first target host is smaller than the preset utilization rate threshold when the first to-be-migrated instance is migrated; and

[0107] an instance migrating subunit 646, configured to migrate the first to-be-migrated instance to the second target host, such that the second target host achieves maximum improvement on the resource utilization rate.

[0108] Implementation of functions and functions of various units in the above apparatus can be specifically referred to the implementation of the corresponding steps in the above method in detail, which is not detailed herein.

[0109] Since the apparatus embodiments substantially correspond to the method embodiments, the apparatus embodiments are described simply, and the relevant part may be obtained with reference to the part of the description in the method embodiments. The above described apparatus embodiments are merely for illustration purposes only. The units which are described as separate components may be physically separated or may be not physically separated, and the components which are illustrated as units may be or may not be physical units, that is, the components may be located in the same position or may be distributed into a plurality of network units. A part or all the modules may be selected according to the actual needs to achieve the objectives of the technical solutions of the embodiments of the present application. Persons of ordinary skill in the art may understand and implement the present application without paying any creative effort.

[0110] As seen from the above embodiment, when scheduling the resources in the cluster in this embodiment, by acquiring a resource parameter of the cluster, calculating a predicted number of hosts in the cluster, determining a to-be-migrated host and a target host from the current hosts in the cluster when the number of current hosts in the cluster is greater than the number of predicted hosts, and migrating an instance running on the to-be-migrated host to the target host, each target host achieves maximum improvement on resource utilization rate, such that the overall resource utilization rate of the cluster is more balanced while the resource demands of the cluster are satisfied with a minimum migration cost.

[0111] For the purposes of this disclosure a module or unit is a software, hardware, or firmware (or combinations thereof) system, process or functionality, or component thereof, that performs or facilitates the processes, features, and/or functions described herein (with or without human interaction or augmentation). A module can include sub-modules. Software components of a module may be stored on a computer readable medium for execution by a processor. Modules may be integral to one or more servers, or be loaded and executed by one or more servers. One or more modules may be grouped into an engine or an application.

[0112] For the purposes of this disclosure the term “user”, “subscriber” “consumer” or “customer” should be understood to refer to a consumer of data supplied by a data provider. By way of example, and not limitation, the term “user” or “subscriber” can refer to a person who receives data provided by the data or service provider over the Internet in a browser session, or can refer to an automated software application which receives the data and stores or processes the data.

[0113] Those skilled in the art will recognize that the methods and systems of the present disclosure may be implemented in many manners and as such are not to be limited by the foregoing exemplary embodiments and examples. In other words, functional elements being performed by single or multiple components, in various combinations of hardware and software

or firmware, and individual functions, may be distributed among software applications at either the client level or server level or both. In this regard, any number of the features of the different embodiments described herein may be combined into single or multiple embodiments, and alternate embodiments having fewer than, or more than, all of the features described herein are possible.

[0114] Functionality may also be, in whole or in part, distributed among multiple components, in manners now known or to become known. Thus, myriad software/hardware/firmware combinations are possible in achieving the functions, features, interfaces and preferences described herein. Moreover, the scope of the present disclosure covers conventionally known manners for carrying out the described features and functions and interfaces, as well as those variations and modifications that may be made to the hardware or software or firmware components described herein as would be understood by those skilled in the art now and hereafter.

[0115] Furthermore, the embodiments of methods presented and described as flowcharts in this disclosure are provided by way of example in order to provide a more complete understanding of the technology. The disclosed methods are not limited to the operations and logical flow presented herein. Alternative embodiments are contemplated in which the order of the various operations is altered and in which sub-operations described as being part of a larger operation are performed independently.

[0116] While various embodiments have been described for purposes of this disclosure, such embodiments should not be deemed to limit the teaching of this disclosure to those embodiments. Various changes and modifications may be made to the elements and operations described above to obtain a result that remains within the scope of the systems and processes described in this disclosure.

[0117] Other embodiments of the present application will be apparent to those skilled in the art through consideration of the specification and practice disclosed herein. The present application is intended to cover any variations, uses, or adaptations of the present disclosure following the general principles thereof and including such departures from the present application as coming within common knowledge or customary technical means in the art. It is intended that the specification and examples shall be considered as exemplary only, with a true scope and spirit of the present application being indicated by the appended claims.

CLAIMS

What is claimed is:

1. A method for resource scheduling comprising:

acquiring, by a resource scheduling device, a resource parameter of the cluster, the cluster including a plurality of hosts, the plurality of hosts including a plurality of current hosts each running one or more instances;

calculating, by the resource scheduling device, a number of predicted hosts in the cluster according to the resource parameter;

determining, by the resource scheduling device, one or more to-be-migrated hosts and one or more target hosts from the current hosts in the cluster when a number of current hosts in the cluster is greater than the number of predicted hosts; and

migrating, by the resource scheduling device, an instance of the one or more instances running on a to-be-migrated host of the one or more to-be-migrated hosts to a target host of the one or more target hosts when the number of current hosts is greater than the number of predicted hosts.

2. The method according to Claim 1, wherein the acquiring the resource parameter of the cluster comprises:

collecting a resource configuration quantity of each of the plurality of hosts in the cluster; and

collecting an actual resource utilization quantity of each of the one or more instances running on each of the plurality of current hosts in the cluster.

3. The method according to Claim 2, wherein the calculating the number of predicted hosts in the cluster comprises:

performing a linear regression analysis on the actual resource utilization quantity of each of the one or more instances running on each of the plurality of current hosts to acquire a predicted resource utilization quantity of each of the one or more instances;

calculating a sum of the predicted resource utilization quantities of the one or more instances running on each of the plurality of current hosts as a predicted resource utilization quantity of each of the plurality of current hosts;

calculating a sum of the predicted resource utilization quantities of the plurality of current hosts as a predicted resource utilization quantity of the cluster; and

calculating the number of predicted hosts in the cluster according to the resource configuration quantities of the plurality of hosts and the predicted resource utilization quantity of the cluster.

4. The method according to Claim 3, wherein the determining the one or more to-be-migrated hosts and the one or more target hosts comprises:

calculating a difference between the number of current hosts and the number of predicted hosts as a number of to-be-migrated hosts in the cluster;

calculating a migration cost value for each of the plurality of current hosts; and

determining the to-be-migrated hosts in the cluster according to an order of the migration cost values of the plurality of current hosts and the calculated number of to-be-migrated hosts, and a number of remaining current hosts in the cluster as the one or more target hosts.

5. The method according to Claim 4, wherein the calculating the migration cost value comprises:

acquiring a number of instances and a total storage space for each of the one or more instances running on each of the plurality of current hosts;

calculating a notification migration cost value for each of the plurality of current hosts according to the number of instances;

calculating a data migration cost value for each of the plurality of current hosts according to the total storage space for each of the one or more instances running; and

calculating a sum of the notification migration cost value and the data migration cost value as the migration cost value for each of the plurality of current hosts.

6. The method according to Claim 3, further comprising:

sequencing, by the resource scheduling device, all the one or more instances running on the one or more to-be-migrated hosts according to a descending order of the predicted resource utilization quantities of the respective instances when the number of current hosts in the cluster is greater than the number of predicted hosts; and

sequentially calculating, by the resource scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, a corresponding target migration host of the one or more target hosts for each instance running on the one or more to-be-migrated hosts, the target migration host being a host of the one or more target hosts for which a predicted resource utilization quantity achieves a maximum improvement after the respective instance is migrated to the target migration host, and migrating each instance running on the one or more to-be-migrated hosts to the respective target migration host,

wherein the migrating the instance of the one or more instances running on the to-be-migrated host to the target host is migrating a first instance running on a first to-be-migrated host to the corresponding target migration host.

7. The method according to Claim 1, further comprising:

controlling, by the resource scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, the to-be-migrated host of the one or more hosts to go offline from the cluster when the instance running on the to-be-migrated host is migrated to the target host.

8. The method according to Claim 1, further comprising:

calculating, by the resource scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, an actual resource utilization rate corresponding to each of the one or more target hosts;

determining, by the resource scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, a first target host of the one or more target hosts for which an actual resource utilization rate exceeds a preset utilization rate threshold;

determining, by the resource scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, a first to-be-migrated instance from the one or more instances running on the first target host, such that the actual resource utilization rate of the first target host is determined to be smaller than the preset utilization rate threshold after the first to-be-migrated instance is migrated; and

migrating, by the resource scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, the first to-be-migrated instance to a second target host of the one or more target hosts, the second target host being determined to achieve a maximum improvement on the resource utilization rate after migrating the first to-be-migrated instance.

9. The method according to Claim 1, further comprising:

determining, by the resource scheduling device when the number of current hosts is less than the number of predicted hosts, a first current host of the plurality of current hosts for which an actual resource utilization rate exceeds a preset utilization rate threshold;

determining, by the resource scheduling device when the number of current hosts is less than the number of predicted hosts, a group of instances from the one or more instances running on the first current host, such that the actual resource utilization rate of the first current host is determined to be smaller than the preset utilization rate threshold after the group of instances is migrated; and

adding, by the resource scheduling device, one or more new hosts from the cluster, when the number of current hosts is less than the number of predicted hosts, a number of the new hosts being equal to a difference between the number of predicted hosts and the number of current hosts, and migrating the group of instances to a second host of the one or more new hosts, the second host being determined such that an actual resource utilization rate of the second host does not exceed the preset utilization rate threshold after migration.

10. A scheduling device for resource scheduling on a cluster including a plurality of hosts, the device comprising:

one or more processors;

a network interface; and

a non-transitory memory storing computer-executable instructions executable by the processor, the instructions causing the device to:

acquire a resource parameter of the cluster, the cluster including a plurality of hosts, the plurality of hosts including a plurality of current hosts each running one or more instances;

calculate a number of predicted hosts in the cluster according to the resource parameter;

determine one or more to-be-migrated hosts and one or more target hosts from the current hosts in the cluster when a number of current hosts in the cluster is greater than the number of predicted hosts; and

migrate an instance of the one or more instances running on a to-be-migrated host of the one or more to-be-migrated hosts to a target host of the one or more target hosts when the number of current hosts is greater than the number of predicted hosts.

11. The scheduling device according to Claim 10, wherein the acquiring instruction further causes the device to:

collect a resource configuration quantity of each of the plurality of hosts in the cluster;
and

collect an actual resource utilization quantity of each of the one or more instances running on each of the plurality of current hosts in the cluster.

12. The scheduling device according to Claim 11, wherein the calculating instruction further causes the device to:

perform a linear regression analysis on the actual resource utilization quantity of each of the one or more instances running on each of the plurality of current hosts to acquire a predicted resource utilization quantity of each of the one or more instances;

calculate a sum of the predicted resource utilization quantities of the one or more instances running on each of the plurality of current hosts as a predicted resource utilization quantity of each of the plurality of current hosts;

calculate a sum of the predicted resource utilization quantities of the plurality of current hosts as a predicted resource utilization quantity of the cluster; and

calculate the number of predicted hosts in the cluster according to the resource configuration quantities of the plurality of hosts and the predicted resource utilization quantity of the cluster.

13. The scheduling device according to Claim 12, wherein the determining instruction further causes the device to:

calculate a difference between the number of current hosts and the number of predicted hosts as a number of to-be-migrated hosts in the cluster;

calculate a migration cost value for each of the plurality of current hosts; and

determine the to-be-migrated hosts in the cluster according to an order of the migration cost values of the plurality of current hosts and the calculated number of to-be-migrated hosts, and a number of remaining current hosts in the cluster as the one or more target hosts.

14. The scheduling device according to Claim 13, wherein the calculating the migration cost value instruction further causes the device to:

acquire a number of instances and a total storage space for each of the one or more instances running on each of the plurality of current hosts;

calculate a notification migration cost value for each of the plurality of current hosts according to the number of instances;

calculate a data migration cost value for each of the plurality of current hosts according to the total storage space for each of the one or more instances running; and

calculate a sum of the notification migration cost value and the data migration cost value as the migration cost value for each of the plurality of current hosts.

15. The scheduling device according to Claim 12, wherein the instructions further cause the device to:

sequence all the one or more instances running on the one or more to-be-migrated hosts according to a descending order of the predicted resource utilization quantities of the respective instances when the number of current hosts in the cluster is greater than the number of predicted hosts; and

sequentially calculate, by the scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, a corresponding target migration host of the one or more target hosts for each instance running on the one or more to-be-migrated hosts, the target migration host being a host of the one or more target hosts for which a predicted resource utilization quantity achieves a maximum improvement after the respective instance is migrated to the target migration host, and migrating each instance running on the one or more to-be-migrated hosts to the respective target migration host, the migrating the instance of the one or more instances running on the to-be-migrated host to the target host comprising migrating a first instance running on a first to-be-migrated host to the corresponding target migration host.

16. The scheduling device according to Claim 10, wherein the instructions further cause the device to:

control, by the scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, the to-be-migrated host of the one or more hosts to go offline from the cluster when the instance running on the to-be-migrated host is migrated to the target host.

17. The scheduling device according to Claim 10, wherein the instructions further cause the device to:

calculate, by the scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, an actual resource utilization rate corresponding to each of the one or more target hosts;

determine, by the scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, a first target host of the one or more target hosts for which an actual resource utilization rate exceeds a preset utilization rate threshold;

determine by the scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, a first to-be-migrated instance from the one or more instances running on the first target host, such that the actual resource utilization rate of the first target host is determined to be smaller than the preset utilization rate threshold after the first to-be-migrated instance is migrated; and

migrate, by the scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, the first to-be-migrated instance to a second target host of the one or more target hosts, the second target host being determined to achieve a maximum improvement on the resource utilization rate after migrating the first to-be-migrated instance.

18. The scheduling device according to Claim 10, wherein the instructions further cause the device to:

determine, by the scheduling device when the number of current hosts is less than the number of predicted hosts, a first current host of the plurality of current hosts for which an actual resource utilization rate exceeds a preset utilization rate threshold;

determine, by the scheduling device when the number of current hosts is less than the number of predicted hosts, a group of instances from the one or more instances running on the first current host, such that the actual resource utilization rate of the first current host is determined to be smaller than the preset utilization rate threshold after the group of instances is migrated; and

add, by the scheduling device, one or more new hosts from the cluster, when the number of current hosts is less than the number of predicted hosts, a number of the new hosts being equal

to a difference between the number of predicted hosts and the number of current hosts, and migrating the group of instances to a second host of the one or more new hosts, the second host being determined such that an actual resource utilization rate of the second host does not exceed the preset utilization rate threshold after migration.

19. A non-transitory computer-readable storage medium storing computer-executable instructions thereon that, when executed by a processor of a scheduling device, cause the scheduling device to perform resource scheduling on a cluster including a plurality of hosts, the scheduling comprising:

acquiring, by the scheduling device, a resource parameter of the cluster, the cluster including a plurality of hosts, the plurality of hosts including a plurality of current hosts each running one or more instances;

calculating, by the scheduling device, a number of predicted hosts in the cluster according to the resource parameter;

determining, by the scheduling device, one or more to-be-migrated hosts and one or more target hosts from the current hosts in the cluster when a number of current hosts in the cluster is greater than the number of predicted hosts; and

migrating, by the scheduling device, an instance of the one or more instances running on a to-be-migrated host of the one or more to-be-migrated hosts to a target host of the one or more target hosts when the number of current hosts is greater than the number of predicted hosts.

20. The non-transitory computer-readable storage medium according to Claim 19, the scheduling further comprising

controlling, by the scheduling device when the number of current hosts in the cluster is greater than the number of predicted hosts, the to-be-migrated host of the one or more hosts to go

offline from the cluster when the instance running on the to-be-migrated host is migrated to the target host.




FIG. 1

2/9





FIG. 2

FIG. 3



4/9




FIG. 4

5/9




FIG. 5

6/9




FIG. 6

7/9




FIG. 7




FIG. 8

8/9




FIG. 9

9/9




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2016/041426

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 9/455, G06F 9/50, G06F 15/173, G06F 17/50, H04W 72/08, H04W 72/10 (2016.01)

CPC - G06F 9/4856, G06F 9/5016, G06F 9/5077, G06F 11/008, G06F 11/3409, G06F 11/3452 (2016.08)

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC - G06F 9/455, G06F 9/50, G06F 15/173, G06F 17/50, H04W 72/08, H04W 72/10

CPC - G06F 9/4856, G06F 9/5016, G06F 9/5077, G06F 11/008, G06F 11/3409, G06F 11/3452, G06F 2201/815, H04L 41/0896, H04L 41/147, H04L 41/22, H04L 47/14, H04L 47/623, H04L 47/626, H04L 47/6265, H04L 47/6275, H04W 72/121

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

USPC - 455/512, 455/513, 703/13, 709/223, 709/224, 709/226, 718/1 (keyword delimited)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Orbit, Google Patents, Google Scholar, Google

Search terms used: resource scheduling, acquire resource parameter of cluster, resource scheduling device

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2014/0019966 A1 (NEUSE et al) 16 January 2014 (16.01.2014), entire document	1-20
Y	US 2013/0346969 A1 (SHANMUGANATHAN et al) 26 December 2013 (26.12.2013), entire document	1-20
Y	US 2014/0337442 A1 (ORACLE INTERNATIONAL CORPORATION) 13 November 2014 (13.11.2014), entire document	3-6, 12-15
Y	US 2015/0169354 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 18 June 2015 (18.06.2015), entire document	4, 5, 13, 14
Y	US 2012/0036515 A1 (HEIM) 09 February 2012 (09.02.2012), entire document	6, 8, 15, 17
A	US 2009/0300173 A1 (BAKMAN et al) 03 December 2009 (03.12.2009), entire document	1-20
A	US 2013/0040681 A1 (CELLCO PARTNERSHIP D/B/A VERIZON WIRELESS) 14 February 2013 (14.02.2013), entire document	1-20
A	US 2012/0096468 A1 (CHAKRAVORTY et al) 19 April 2012 (19.04.2012), entire document	1-20
A	US 2014/0344810 A1 (HUAWEI TECHNOLOGIES CO., LTD.) 20 November 2014 (20.11.2014), entire document	1-20



Further documents are listed in the continuation of Box C.



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 application or patent but published on or after the international filing date

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

02 September 2016

Date of mailing of the international search report

14 SEP 2016

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774