




- (51) International Patent Classification:
G06F 17/00 (2006.01)
- (21) International Application Number:
PCT/US2016/041500
- (22) International Filing Date:
8 July 2016 (08.07.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
201510398194.X 8 July 2015 (08.07.2015) CN
- (71) Applicant: **ALIBABA GROUP HOLDING LIMITED**
[—/US]; Fourth Floor, One Capital Place, P. O. Box 847,
George Town, Grand Cayman (KY).
- (72) Inventors: **WANG, Heng**; Alibaba Group Legal Depart-
ment, 5/f, Building 3, No. 969 West Wen Yi Road, Yu
Hang District, Hangzhou, 311121 (CN). **CHEN, Xu**;
Alibaba Group Legal Department, 5/f, Building 3, No. 969
West Wen Yi Road, Yu Hang District, Hangzhou, 311121
(CN).
- (74) Agent: **CAPRON, Aaron, J.**; Finnegan, Henderson, Fara-
bow, Garrett & Dunner LLP, 901 New York Avenue, NW,
Washington, DC 20001-4413 (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.
- (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).

[Continued on next page]

(54) Title: APPARATUSES AND METHODS FOR FLEXIBLE SCHEDULING IN A DATABASE SYSTEM

600



(57) Abstract: Apparatuses and methods are disclosed for database schedul-
ing. An exemplary method may include obtaining a performance parameter
associated with a database server. The method may also include determining
whether the performance parameter matches a preset condition. After it is
determined that the performance parameter matches the preset condition, the
method may further include selecting a database instance operating on the
database server and migrating the database instance to a receiving database
server.

FIG. 6



Published:

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

APPARATUSES AND METHODS FOR FLEXIBLE SCHEDULING IN A DATABASE SYSTEM

CROSS REFERENCE TO RELATED APPLICATION

[001] The present application claims the benefits of priority to Chinese Application No. 201510398194.X, filed July 8, 2015, the entire contents of which have been incorporated herein by reference.

TECHNICAL FIELD

[002] The present application relates to database technologies, and more particularly, to flexible scheduling methods and apparatuses in a database system.

BACKGROUND

[003] The development of cloud computing brings new challenges to database technologies. Traditional database services suitable for servicing a single client or workstation have to evolve into cluster-level or even datacenter-level services. In order to manage and operate the vast amount of database instances involved in such high-level services, database flexible scheduling technology has been introduced and become an important technical solution for supporting a robust clouding computing environment.

[004] In general, database flexible scheduling involves computing a distribution of database instances among database servers according to certain triggering conditions, constrains, and a predetermined scheduling goal. Database instances are then migrated to target database servers according to the computed distribution.

[005] Existing database flexible scheduling technologies normally adopt an optimization approach, such as methods based on minimum cost. In such methods, one or more objective functions may be generated as the first step. Based on the objective

functions and various constraints, an optimized scheduling strategy may be obtained through computing and searching for an optimized solution.

[006] Existing database flexible scheduling technologies, however, have several disadvantages. First, the computational jitter is relatively large, which may cause instability of the database service. This is because the optimization-based methods may be sensitive to performance changes of database instances, causing recomputation of the objective functions and the corresponding scheduling strategy, which in turn may cause local or even global jitters, affecting the stability of the database service. Second, the cost of migration is normally quite significant. This is because the optimization-based scheduling methods aim to achieve an optimal distribution based on pure mathematical consideration, which would normally cause migration of a large number of database instances. Third, the computation is normally very complex and inefficient. As the number of the objective functions increases, the computation cost increases almost exponentially.

SUMMARY

[007] In one aspect, the present disclosure is directed to a method for performing database scheduling. The method may include obtaining a performance parameter associated with a database server. The method may also include determining whether the performance parameter matches a preset condition. After it is determined that the performance parameter matches the preset condition, the method may further include selecting a database instance operating on the database server and migrating the database instance to a receiving database server.

[008] In another aspect, the present disclosure is directed to an apparatus for performing database scheduling. The apparatus may include a memory device storing computer-readable instructions. The apparatus may also include a processor device

communicatively coupled to the memory device. The computer-readable instructions, when executed by the processor device, may cause the processor device to perform various operations. The operations may include obtaining a performance parameter associated with a database server. The operations may also include determining whether the performance parameter matches a preset condition. After it is determined that the performance parameter matches a preset condition, the operations may further include selecting a database instance operating on the database server and migrating the database instance to a receiving database server.

[009] In a further aspect, the present disclosure is directed to a non-transitory computer-readable medium that stores a set of instructions that is executable by at least one processor of a device to cause the device to perform a method for performing database scheduling. The method may include obtaining a performance parameter associated with a database server. The method may also include determining whether the performance parameter matches a preset condition. After it is determined that the performance parameter matches the preset condition, the method may further include selecting a database instance operating on the database server and migrating the database instance to a receiving database server.

[010] Additional objects and advantages of the present disclosure will be set forth in part in the following detailed description, and in part will be obvious from the description, or may be learned by practice of the present disclosure. The objects and advantages of the present disclosure will be realized and attained by means of the elements and combinations particularly pointed out in the appended claims.

[011] It is to be understood that the foregoing general description and the following detailed description are exemplary and explanatory only, and are not restrictive of the invention, as claimed.

BRIEF DESCRIPTION OF THE DRAWINGS

[012] The accompanying drawings, which constitute a part of this specification, illustrate several embodiments and, together with the description, serve to explain the disclosed principles.

[013] Fig. 1 is a flow chart of an exemplary database scheduling method, according to some embodiments of the present application.

[014] Fig. 2 is a flow chart of an exemplary database instance migration method, according to some embodiments of the present application.

[015] Fig. 3 is a flow chart of another exemplary database instance migration method, according to some embodiments of the present application.

[016] Fig. 4 is a flow chart of another exemplary database scheduling method, according to some embodiments of the present application.

[017] Fig. 5 is a flow chart of another exemplary database scheduling method, according to some embodiments of the present application.

[018] Fig. 6 is a schematic diagram of an exemplary database scheduling apparatus, according to some embodiments of the present application.

[019] Fig. 7 is a flow chart of another exemplary database scheduling method, according to some embodiments of the present application.

[020] Fig. 8 is a schematic diagram of another exemplary database scheduling apparatus, according to some embodiments of the present application.

[021] Fig. 9 is a flow chart of another exemplary database scheduling method, according to some embodiments of the present application.

[022] Fig. 10 is a schematic diagram of another exemplary database scheduling apparatus, according to some embodiments of the present application.

DETAILED DESCRIPTION

[023] Reference will now be made in detail to exemplary embodiments of the invention, examples of which are illustrated in the accompanying drawings. When appropriate, the same reference numbers are used throughout the drawings to refer to the same or like parts.

[024] The present application discloses methods and apparatuses for performing database scheduling. Database scheduling refers to a technology for determining a distribution of database instances among database servers and migrating one or more database instances into one or more target database servers according to the distribution. Advanced database scheduling methods that provide enhanced flexibility and scalability may often be referred to as flexible database scheduling methods or elastic database scheduling methods. For simplicity, the term database scheduling is used to refer to such advanced database scheduling methods, which are the subject of this application.

[025] A database instance is a set of memory structures that manage database files. A database is a set of physical files stored on one or more storage devices such as hard drive disks. A database server is a computer server that provides database services by hosting one or more database instances.

[026] Fig. 1 is a flow chart of an exemplary method 100 for performing database scheduling. Method 100 may include several steps.

[027] In step 101, method 100 may include obtaining performance parameters associated with database servers and/or database instances that are subject to scheduling.

[028] A database cluster may include a plurality of database servers. Each database server may host a plurality of database instances. Each database instance may provide database service to one or more users.

[029] In some embodiments, a range of performance parameters may be predetermined for each of the database servers and/or instances. A database server/instance may be deemed operating properly when its performance parameter falls within the predetermined range.

[030] When the performance parameter of a database server exceeds an upper threshold of the predetermined range, such as a capability threshold associated with that database server, the database server, as well as the database instances operating on the database server, is considered overloaded and may not be able to provide stable service. Under this situation, technical solutions disclosed in the present application may be used to reduce the workload of the database server by migrating one or more database instances to another database server according to a flexible scheduling strategy. The disclosed method is based on considerations of the intrinsic characteristics of the database services. After migration, the performance parameter of the database server may be reduced to a level within the predetermined range, avoiding risks associated with overloading the database server.

[031] On the other hand, when the performance parameter of a database server is below a lower threshold of the predetermined range, the database server is considered underutilized. In this case, the CPU usage rate and the memory usage rate may be relatively low, causing a waste of system resources. Under this situation, technical solutions disclosed in the present application may be used to migrate all database instances operating on the database server to one or more other database servers. After migration, the underutilized database server no longer hosts any database instance and can be taken offline (e.g., be placed in an offline mode) to save system resources.

[032] A performance parameter may be a performance indicator reflecting the operating condition of a database server or a database instance. In some embodiments, each

database server and each database instance operating on the database server may have a corresponding performance parameter. A performance parameter may include one or more performance indicators, such as at least one of: CPU usage (e.g., in percentage), memory usage (e.g., in percentage), disk I/O speed, network transmission speed, or disk usage (e.g., in percentage). Other suitable performance parameter may also be used.

[033] In some embodiments, the performance parameter may include one or more dimensions to indicate a single or multiple operation conditions of a database server/database instance. For example, the performance parameter may include two, three, four, or more dimensions, each corresponding to a performance indicator.

[034] The predetermined range refers to a predetermined performance parameter values defining a range including a lower threshold and an upper threshold. The predetermined range may be set according to the physical configurations of the database servers.

[035] The performance parameter may be obtained using tools provided by the operating system on which the database servers operate.

[036] In step 102, method 100 may include determining whether the performance parameter of a database server matches a preset condition, such as whether the performance parameter falls within or outside of a predetermined range. In some embodiments, the preset condition may include a condition that the performance parameter falls outside of a predetermined range. In other embodiments, the preset condition may include a condition that the performance parameter falls within a predetermined range.

[037] When the performance parameter falls within the predetermined range, no further scheduling operation is needed. The predetermined range can provide a certain tolerance to small-scale performance fluctuations, thereby reducing computational jittering.

[038] When the performance parameter falls outside of the predetermined range, method 100 proceeds to step 103 and/or step 104, depending on whether the performance parameter is above the upper threshold or below the lower threshold of the predetermined range.

[039] When the performance parameter is above (e.g., exceeds) the upper threshold, method 100 proceeds to step 103, in which one or more database instances operating on the database server are migrated to other one or more database servers to reduce the workload of the database server originally hosting these database instances.

[040] The one or more database servers receiving the one or more migrating database instances may be referred to as receiving database servers. In some embodiments, a receiving database server may be a database server whose performance parameter is below the lower threshold of the predetermined range (e.g., the receiving database server is underutilized and receiving database instances helps to boost the utilization rate of the receiving database server). In other embodiments, a receiving database server may be a database server whose performance parameter is below the upper threshold of the predetermined range (e.g., the receiving database server is operating properly and is neither underutilized nor overloaded).

[041] In some embodiments, the receiving database server may be selected such that after receiving the migrating database instance, the performance parameter of the receiving database server remains under the upper threshold of the predetermined range. In this way, the receiving database server will not become overloaded solely because of the acceptance of the migrating database instance.

[042] Fig. 2 shows a flow chart of an exemplary method for migrating database instances.

[043] In step 103-1, the method may include determining workloads for those database servers whose performance parameters exceed the upper threshold of the predetermine range. The workload of a database server may be determined based on the performance parameter of that database server. In some embodiments, the workload may refer to the load of the database server. In other embodiments, other values may also be used to indicate the workload of the database server.

[044] In step 103-2, the method may include sorting the database servers whose performance parameters exceed the upper threshold according to their workloads. The sorted list of database servers may be arranged in a descending order (e.g., from the heaviest workload to the lightest workload).

[045] In step 103-3, the method may include selecting a database server from the sorted list. In some embodiments, database servers having heavier workloads may be preferably selected over database servers having lighter workloads. Database servers having heavier workloads may have a higher probability to encounter problems than those having lighter workloads. Therefore, preferably selecting database servers having heavier workloads may reduce the overall risk of failure for the database server cluster. In some embodiments, the selection may be conducted sequentially down the list, from the database server having the heaviest workload to the database server having the lightest workload.

[046] In step 103-4, the method may include selecting one or more database instances from all database instances operating on the selected database server and migrating the one or more database instances to one or more receiving database servers, such that after migration, the performance parameter of the selected database server no longer exceeds the upper threshold of the predetermined range.

[047] Specifically, step 103-4 may include selecting a database instance that satisfies a preset condition. The preset condition may include ranking relatively high in a

list of database instances operating on the selected database server that are arranged according to their performance parameters in an ascending order (e.g., from the smallest parameter value to the highest parameter value). The performance parameter may include any performance indicator selected from CPU usage, memory usage, disk I/O speed, network transmission speed, and disk usage. The performance parameter may also be calculated as a combined performance indicator based on any combination of the above performance indicators.

[048] In some embodiments, a database instance having a relatively small performance parameter (e.g., the smallest performance parameter) may be selected as the migrating database instance.

[049] After being selected, the database instance may be migrated to another database server (e.g., a receiving database server) whose performance parameter, after accepting the migrating database instance, would remain under the upper threshold of the predetermined range so that the receiving database server would not become overloaded by accepting the migrating database instance.

[050] Specifically, the migration process can be conducted as follows:

[051] 1) requesting initialization of a database instance in the receiving database server;

[052] 2) loading and executing the migrating database instance on the initialized database instance; and

[053] 3) taking the migrating database instance offline from the database server originally hosting the migrating database instance.

[054] Once a database instance has been emigrated, step 103-4 may include determining whether the performance parameter of the database server (e.g., the database server originally hosting the migrating database instance, hereinafter referred to as the

hosting database server) offloading the migrating database instance still exceeds the upper threshold of the predetermined range. If yes, then the process repeats the database instance selection and migration operations until the performance parameter of the hosting database server no longer exceeds the upper threshold of the predetermined range. If no, then it indicates that the hosting database server is no longer overloaded. The process then proceeds to step 103-5.

[055] In step 103-5, the method may include determining whether there is any remaining database server whose performance parameter exceeding the upper threshold of the predetermined range. If yes, then the process loops back to step 103-3 to select the next database server to offload one or more database instances. If no, then step 103 completes, indicating that all overloaded database servers subject to scheduling have been processed.

[056] Referring back to Fig. 1, step 103 involves operations for reducing workloads of those overloaded database servers, while step 104 involves operations for taking those underutilized database servers offline to save system resources. Steps 103 and 104 may be performed in any order, or may be performed concurrently in separate threads. In step 104, the method may include selecting and migrating all database instances operating on an underutilized database server to one or more other database servers (e.g., receiving database servers), such that the underutilized database server no longer hosts any database instance and can be taken offline.

[057] As described above, a database server may be deemed underutilized when its performance parameter is below the lower threshold of the predetermined range. The one or more receiving database servers that accept the migrating database instances may be selected according to the same conditions as in step 103. For example, a receiving database server can be another underutilized database server whose performance parameter is below the lower threshold. In another example, a receiving database server can be a properly

operating database server that is neither underutilized nor overloaded (e.g., the performance parameter falls within the predetermined range).

[058] Fig. 3 shows an exemplary database instance migration method implementing step 104. The method shown in Fig. 3 is similar to the method shown in Fig. 2. Therefore, the duplicated portions will be omitted, which can be referred to the corresponding description of Fig. 2.

[059] In step 104-1, the method may include determining workloads for those database servers whose performance parameters are below the lower threshold of the predetermine range.

[060] In step 104-2, the method may include sorting the database servers whose performance parameters are below the lower threshold according to their workloads. The sorted list of database servers may be arranged in an ascending order (e.g., from the lightest workload to the heaviest workload).

[061] In step 104-3, the method may include selecting a database server from the sorted list. In some embodiments, database servers having lighter workloads may be preferably selected over database servers having heavier workloads. Database servers having lighter workloads may waste more system resources than those having heavier workloads. Therefore, preferably selecting database servers having lighter workloads may take offline these resource wasters earlier, thereby saving the overall system resources. In some embodiments, the selection may be conducted sequentially down the list, from the database server having the lightest workload to the database server having the heaviest workload.

[062] In step 104-4, the method may include selecting one or more database instances from all database instances operating on the selected database server and migrating the one or more database instances to one or more receiving database servers.

After all database instances have been migrated, the selected database server no longer hosts any database instance and can be taken offline.

[063] Specifically, step 104-4 may include selecting a database instance that satisfies a preset condition. The preset condition may include ranking relatively high in a list of database instances operating on the selected database server that are arranged according to their performance parameters in an ascending order (e.g., from the smallest parameter value to the highest parameter value).

[064] In some embodiments, a database instance having a relatively small performance parameter (e.g., the smallest performance parameter) may be selected as the migrating database instance.

[065] After being selected, the database instance may be migrated to another database server (e.g., a receiving database server) whose performance parameter, after accepting the migrating database instance, would remain under the upper threshold of the predetermined range so that the receiving database server would not become overloaded by accepting the migrating database instance.

[066] In some embodiments, the migration may also help to increase the utilization rate of another underutilized database server. For example, the receiving database server may be an underutilized database server whose performance parameter is below the lower threshold of the predetermined range.

[067] After a database instance has been emigrated, step 104-4 may include determining whether the hosting database server still hosts any database instance. If yes, then the process repeats the database instance selection and migration operations until the hosting database server no longer hosts any database instance. If no, then it indicates that the hosting database server no longer hosts any database instance and the database server can be taken offline. The process then proceeds to step 104-5.

[068] In step 104-5, the method may include determining whether there is any remaining database server whose performance parameter is below the lower threshold of the predetermined range. If yes, then the process loops back to step 104-3 to select the next database server to offload all its database instances. If no, then step 104 completes, indicating that all underutilized database servers subject to scheduling have been processed.

[069] Method 100 of Fig. 1 conducts database scheduling through setting a predetermined range. Database servers whose performance parameters exceed the upper threshold of the predetermined range are considered overloaded, and their database instances can be migrated to other database servers to reduce their workload back to a level within the predetermined range, thereby enhancing the stability of the services. On the other hand, database servers whose performance parameters are below the lower threshold of the predetermined range are considered underutilized, and their database instances can be migrated to other database servers such that they can be taken offline to save system resources. In addition, by setting a predetermined range and foregoing any migration operation when the performance parameter of a database server falls within the range, method 100 is resilient to insignificant performance fluctuations during real time operations, avoiding system jittering, and reducing scheduling cost.

[070] Fig. 4 shows a flow chart of another exemplary method 400 for performing database scheduling. Some steps shown in Fig. 4 are similar to those shown in Fig. 1. Therefore, duplicative description is omitted.

[071] In step 401, method 400 may include obtaining performance parameters associated with database servers and/or database instances that are subject to scheduling.

[072] A database cluster may include a plurality of database servers. When the database servers within a database cluster are non-uniform, such as having different hardware and/or software configurations, the scheduling strategy may be affected. For

example, the complexity of the scheduling algorithm may be increased to account for the non-uniformity. In addition, during events such as holding promotions, adding new functions, or adding new businesses, the fluctuations of database server workloads may increase, causing system jittering.

[073] To solve the above problem, method 400 introduces the concept of expectation factor. The expectation factor is a customized configuration, which can provide some customized features. The expectation factor can be implemented by setting a performance parameter expectation for a database server or a database instance subject to scheduling.

[074] In step 402, method 400 may include determining whether any database server or database instance is associated with a corresponding performance parameter expectation. If yes, then method 400 proceeds to step 403, in which the performance parameter of the database server/instance obtained in step 401 is substituted with the corresponding performance parameter expectation. If no, then the performance parameter of the database server/instance obtained in step 401 remains the same.

[075] In step 404, method 400 may include determining whether the performance parameter of a database server falls within a predetermined range. Here, the performance parameter may be the original performance parameter obtained in step 401 or the performance parameter expectation, depending on the outcome of step 402.

[076] If the outcome of step 404 is Yes, the method 400 may end because no scheduling operations are needed. For simplicity and clarity, this branch is omitted from Fig. 4.

[077] If the outcome of step 404 is No, then method 400 proceeds to step 405/406, which are similar to 103/104.

[078] In some embodiments, the performance parameter expectation may be associated with a validity period. For example, multiple time periods may be set for a database server/instance, each time period may correspond to a separate performance parameter expectation value.

[079] For each database server subject to scheduling, the validity period can be implemented as follows.

[080] To determine the performance parameter expectation of the database server, it is first determined whether the current time is within a particular validity period.

[081] If yes, then substitute the obtained performance parameter with the performance parameter expectation corresponding to the validity period.

[082] If no, then keep the obtaining performance parameter unchanged.

[083] Similarly, for each database instance subject to scheduling, the validity period can be implemented as follows.

[084] To determine the performance parameter expectation of the database instance, it is first determined whether the current time is within a particular validity period.

[085] If yes, then substitute the obtained performance parameter with the performance parameter expectation corresponding to the validity period.

[086] If no, then keep the obtaining performance parameter unchanged.

[087] Performance parameter expectations can be used to customize the expectation values for particular database servers/instances, reducing the complexity due to the non-uniformity in the database cluster.

[088] In addition, during a promotion event, the performance parameter expectations for those involved database servers/instances may be temporarily increased, avoiding overloading the servers due to data surge. After the promotion event, the

performance parameter expectations can be reduced to regular levels, avoiding resource waste.

[089] When new functions or new businesses are introduced, the performance parameter expectations for those involved database servers/instances may be adjusted, meeting the service demands and reducing the number of migrations, thereby reducing the scheduling cost.

[090] Fig. 5 shows a flow chart of another exemplary method 500 for performing database scheduling. Some steps shown in Fig. 5 are similar to those shown in Fig. 1. Therefore, duplicative description is omitted.

[091] In step 501, method 500 may include obtaining performance parameters associated with database servers and/or database instances that are subject to scheduling.

[092] In general, databases should have a certain level of disaster recovery ability, referring to the ability to handover functionalities among two or more databases in case one of the databases experiences problems. This way, the backup database(s) may take over the functionalities to ensure continued availability of the database services. For example, in MySQL, databases are often prepared in pairs, such as Master <--> Slave pairs (wherein the master database and the slave database are mutual backups of each other) and Master --> Slave pairs (where the slave database is a backup of the master database). In some cases, multiple slave databases may be cascaded, such as Slave --> Master <-- Slave <-- Slave.

[093] In some embodiments, master and slave databases provide different services. For example, master database serves read and write requests, while slave database does not provide any services or only provide a limited amount of services. Because the differences in services, the performance parameters associated with master and slave databases may also be different. In general, the performance parameter of a master database is far greater than the performance parameter of a slave database.

[094] When the master database experiences problems, the disaster recovery mechanism may be initialized to switch the slave database to the role of the master database. The master database, after repair, may be switched to the role of the slave database. In this case, because the workload of the former slave database is relatively light before the switching, during database scheduling, some database instances may be migrated to the former slave database. In addition, after switching, the former slave database may experience much heavier workload for performing the role of a master database. Therefore, it is possible or even likely that overloading conditions occur to the former slave database after switching. Under certain situations, the switching may risk the possibility of failure.

[095] To avoid the above problem, the present application introduces the concept of master-slave consistency. The consistency refers to a configuration that the performance parameters for the master and slave databases are kept the same, shielding the actual performance differences between the two, thereby avoiding the problems caused by the performance differences.

[096] In step 502, method 500 may include determining the largest performance parameter among a plurality of performance parameters respectively associated with a plurality of interrelated database instances, and substituting the plurality of performance parameters with the largest performance parameter.

[097] For example, the interrelated database instances may include a master database instance and a slave database instance, which is a backup of the master database instance and is therefore normally associated with a much smaller performance parameter than the master database instance. More interrelated database instances may form a group according to particular disaster recovery mechanisms and backup relationship. In this case, the largest performance parameter, normally associated with the master database instance, may be used as the common performance parameter for the entire group.

[098] For each group of interrelated database instances, the largest performance parameter can be determined and selected based on the performance parameters obtained in step 501. The largest performance parameter may then be used as the performance parameter for each of the database instances in the group.

[099] The relationship among the interrelated database instances in the group may include a master-slave or master-backup relationship, such as one master - one backup or one master - multiple backup. Other relationship may also be used to establish the connections among the database instances to form a group.

[0100] In step 503, method 500 may include determining whether the performance parameter of a database server falls within a predetermined range, similar to step 102.

[0101] If no, method 500 may proceed to step 504/505, similar to steps 103/104.

[0102] Using a common performance parameter for a group of interrelated database instances may reduce the risk of overloading during master-slave switching, enhancing the stability of services.

[0103] The above passages are described in connection with exemplary methods for performing database scheduling. The present application also discloses apparatuses for performing the database scheduling, corresponding to the methods. Fig. 6 shows a functional block diagram of an exemplary apparatus. Because the functions of each block shown in Fig. 6 correspond to the steps shown in Fig. 1, detail description of Fig. 6 is omitted. The following description is for illustrative purposes.

[0104] As shown in Fig. 6, apparatus 600 may include a performance parameter obtaining unit 601 configured to obtain performance parameters associated with database servers and/or database instances subject to scheduling. Apparatus 600 may also include a performance parameter determination unit 602 configured to determine whether the performance parameter of a database server falls within a predetermined range. Apparatus

600 may also include a first migration unit 603 configured to select one or more database instances operating on a database server whose performance parameter exceeding an upper threshold of the predetermined range and migrate the one or more database instances to one or more receiving database servers. Apparatus 600 may further include a second migration unit 604 configured to migrate one or more database instances operating on a database server whose performance parameter is below a lower threshold of the predetermined range. After migration, the database server no longer hosting any database instances may be placed in an offline mode to save system resources.

[0105] Fig. 7 shows a flow chart of another exemplary database scheduling method 700.

[0106] In step 701, method 700 may include obtaining a predetermined performance parameter expectation corresponding to a database server of a database instance, and use the performance parameter expectation as the performance parameter of the database server/instance.

[0107] In step 702, method 700 may include performing scheduling operations to the database server/instance based on the performance parameter expectation.

[0108] Step 701 of method 700 is similar to step 403 of method 400. In method 700, the performance parameter expectation concept is used as a standalone step, which can be coupled to any suitable database scheduling methods to improve the stability of the services.

[0109] Fig. 8 shows a functional block diagram of an exemplary apparatus 800 configured to perform method 700. Apparatus 800 may include a performance parameter obtaining unit configured to obtain a predetermined performance parameter expectation corresponding to a database server or a database instance, and use the performance parameter expectation as the performance parameter of the database server/instance.

Apparatus 800 may also include a scheduling unit 802 configured to perform scheduling operations to the database server/instance based on the performance parameter expectation.

[0110] Fig. 9 shows a flow chart of another exemplary database scheduling method 900.

[0111] In step 901, method 900 may include selecting, from performance parameters respectively associated with a group of database instances, a performance parameter satisfying a preset condition, and using the selected performance parameter as the common performance parameter for the entire group. The preset condition may include ranking relatively high in a sorted list of database instances arranged in a descending order. For example, the largest performance parameter may be set as the common performance parameter for the group.

[0112] In step 902, method 900 may include performing scheduling operations to the database instances based on the common performance parameter.

[0113] Step 901 of method 900 is similar to step 502 of method 500. In method 900, the common performance parameter concept is used as a standalone step, which can be coupled to any suitable database scheduling methods to improve the stability of the services.

[0114] Fig. 10 shows a functional block diagram of an exemplary apparatus 1000 configured to perform method 900. Apparatus 1000 may include a performance parameter determination unit 1001 configured to obtain a predetermined performance parameters of a group of database instances; select, from performance parameters respectively associated with a group of database instances, a performance parameter satisfying a preset condition; and use the selected performance parameter as the common performance parameter for the entire group. Apparatus 1000 may also include a database instance scheduling unit 1002 configured to perform scheduling operations to the database instances based on the common performance parameter.

[0115] Apparatuses 600, 800, and 1000 may be implemented by a computer device including a memory device and a processor device. The memory device may store computer-readable instructions. The instructions, when executed by the processor device, may cause the processor device to perform operations and functions of the various units shown in Figs. 6, 8, and 10, as well as the functions and operations of the method shown in Figs. 1-5, 7, and 9.

[0116] The specification has described apparatuses and methods for performing database scheduling. The illustrated steps are set out to explain the exemplary embodiments shown, and it should be anticipated that ongoing technological development will change the manner in which particular functions are performed. Thus, these examples are presented herein for purposes of illustration, and not limitation. For example, steps or processes disclosed herein are not limited to being performed in the order described, but may be performed in any order, and some steps may be omitted, consistent with disclosed embodiments. Further, the boundaries of the functional building blocks have been arbitrarily defined herein for the convenience of the description. Alternative boundaries can be defined as long as the specified functions and relationships thereof are appropriately performed. Alternatives (including equivalents, extensions, variations, deviations, etc., of those described herein) will be apparent to persons skilled in the relevant art(s) based on the teachings contained herein. Such alternatives fall within the scope and spirit of the disclosed embodiments.

[0117] While examples and features of disclosed principles are described herein, modifications, adaptations, and other implementations are possible without departing from the spirit and scope of the disclosed embodiments. Also, the words “comprising,” “having,” “containing,” and “including,” and other similar forms are intended to be equivalent in meaning and be open ended in that an item or items following any one of these words is not

meant to be an exhaustive listing of such item or items, or meant to be limited to only the listed item or items. It must also be noted that as used herein and in the appended claims, the singular forms “a,” “an,” and “the” include plural references unless the context clearly dictates otherwise.

[0118] Embodiments of the present application may be implemented using hardware, software, firmware, or any combination thereof for allowing a specialized device to perform the functions described above. One or more steps, operations, functions, and modules described herein may be implemented by software instructions or codes stored in one or more memory devices and executed by one or more hardware processor devices. Exemplary hardware processor devices include logic gate circuitry configured or made to perform data processing and/or logic operations, integrated circuits (ICs) designed and made to perform specific functions, programmable gate arrays (PGAs), field-programmable gate arrays (FPGAs), etc.

[0119] Multiple function modules may be integrated into a single physical device, or may be provided in separate physical devices.

[0120] Furthermore, one or more computer-readable storage media may be utilized in implementing embodiments consistent with the present disclosure. A computer-readable storage medium refers to any type of physical and/or non-transitory memory on which information or data readable by a processor may be stored. Thus, a computer-readable storage medium may store instructions for execution by one or more processors, including instructions for causing the processor(s) to perform steps or stages consistent with the embodiments described herein. The term “computer-readable medium” should be understood to include tangible and/or non-transitory items and exclude carrier waves and transient signals, i.e., be non-transitory. Examples of non-transitory computer-readable medium include RAM, ROM, a volatile memory, a nonvolatile memory, a hard drive, a CD

ROMs, a DVD, a flash drive, a flash memory, a cache, a register, a disk, and any other known storage media.

[0121] It is intended that the disclosure and examples be considered as exemplary only, with a true scope and spirit of disclosed embodiments being indicated by the following claims.

WHAT IS CLAIMED IS:

1. A method for performing database scheduling, the method comprising:
obtaining a performance parameter associated with a database server;
determining whether the performance parameter matches a preset condition;
after it is determined that the performance parameter matches the preset condition,
selecting a database instance operating on the database server; and
migrating the database instance to a receiving database server.
2. The method of claim 1, wherein:
the performance parameter includes at least one of: CPU usage, memory usage, disk I/O speed, network transmission speed, or disk usage.
3. The method of claim 1 or 2, wherein the preset condition includes:
the performance parameter falls within a predetermined range; or
the performance parameter does not fall within a predetermined range.
4. The method of claim 3, wherein:
the predetermined range includes an upper threshold; and
a performance parameter associated with the receiving database server is below the upper threshold.
5. The method of claim 3, wherein:
the predetermined range includes a lower threshold; and
a performance parameter associated with the receiving database server is below the

lower threshold.

6. The method of any one of claims 1 to 5, further comprising:
determining whether the database server is associated with a corresponding performance parameter expectation; and
after it is determined that the database server is associated with the corresponding performance parameter expectation, substituting the obtained performance parameter with the performance parameter expectation.

7. The method of claim 6, wherein the performance parameter expectation is associated with a validity period, and the method further comprises:
determining whether a current time is within the validity period; and
after it is determined that the current time is within the validity period, substituting the obtained performance parameter with the performance parameter expectation.

8. The method of any one of claims 1 to 7, further comprising:
obtaining a performance parameter associated with the database instance;
determining whether the database instance is associated with a corresponding performance parameter expectation; and
after it is determined that the database instance is associated with a corresponding performance parameter expectation, substituting the obtained performance parameter associated with the database instance with the performance parameter expectation.

9. The method of claim 8, wherein the performance parameter expectation is associated with a validity period, and the method further comprises:

determining whether a current time is within the validity period; and

after it is determined that the current time is within the validity period, substituting the obtained performance parameter associated with the database instance with the performance parameter expectation.

10. The method of any one of claims 1 to 9, wherein the database server includes a plurality of database instances interrelated to one another, and the method further comprises:

determining a performance parameter associated with each of the plurality of database instances;

determining a largest performance parameter from the plurality of performance parameters respectively determined for the plurality of database instances; and

substituting the performance parameter associated with each of the plurality of database instances with the largest performance parameter.

11. The method of claim 10, wherein the plurality of database instances include a master database instance and a backup of the master database instance.

12. The method of any one of claims 3 to 5, further comprising:

obtaining a performance parameter associated with each of a plurality of database servers;

determining a list of one or more database servers whose performance parameters exceed an upper threshold of the predetermined range;

selecting a first database server from the list;

selecting one or more database instances operating on the first database server for

migration, such that after emigrating the selected one or more database instances from the first database server, the performance parameter of the first database server falls within the predetermined range; and

migrating the one or more database instances to one or more receiving database servers.

13. The method of claim 12, wherein selecting the first database server includes: sorting the one or more database servers in the list according to their respective workloads indicated by the corresponding performance parameters; and selecting the first database server from the sorted one or more database servers, wherein the first database server has a heaviest workload.

14. The method of claim 12 or 13, wherein selecting the one or more database instances includes:

sorting database instances operating on the first database server according to their respective workloads;

selecting a first database instance from the sorted database instances, wherein the first database instance has a lowest workload; and

selecting additional one or more database instances from the sorted database instances for migration according to an ascending order of workload until the performance parameter of the first database server falls within the predetermined range after migration.

15. The method of any one of claims 3 to 5 and 12 to 14, further comprising: obtaining a performance parameter associated with each of a plurality of database servers;

determining a list of one or more database servers whose performance parameters are below a lower threshold of the predetermined range;

selecting a first database server from the list;

migrating all database instances operating on the first database server to one or more receiving database servers; and

placing the first database server in an offline mode after migration.

16. The method of claim 15, further comprising:

sorting the one or more database servers in the list according to their respective workloads indicated by the corresponding performance parameters; and

selecting the first database server from the sorted one or more database servers, wherein the first database server has a lightest workload.

17. An apparatus for performing database scheduling, comprising:

a memory device that stores computer-readable instructions; and

a processor device communicatively coupled to the memory device, wherein the computer-readable instructions, when executed by the processor device, cause the processor device to perform operations including:

obtaining a performance parameter associated with a database server;

determining whether the performance parameter matches a preset condition;

and

after it is determined that the performance parameter matches the preset condition:

selecting a database instance operating on the database server; and

migrating the database instance to a receiving database server.

18. The apparatus of claim 17, wherein:
the performance parameter includes at least one of: CPU usage, memory usage, disk I/O speed, network transmission speed, or disk usage.

19. The apparatus of claim 17 or 18, wherein the preset condition includes:
the performance parameter falls within a predetermined range; or
the performance parameter does not fall within a predetermined range.

20. The apparatus of claim 19, wherein:
the predetermined range includes an upper threshold; and
a performance parameter associated with the receiving database server is below the upper threshold.

21. The apparatus of claim 19, wherein:
the predetermined range includes a lower threshold; and
a performance parameter associated with the receiving database server is below the lower threshold.

22. The apparatus of any one of claims 17 to 21, wherein the operations further comprise:

determining whether the database server is associated with a corresponding performance parameter expectation; and

after it is determined that the database server is associated with the corresponding performance parameter expectation, substituting the obtained performance parameter with

the performance parameter expectation.

23. The apparatus of claim 22, wherein the performance parameter expectation is associated with a validity period, and the operations further comprise:

determining whether a current time is within the validity period; and

after it is determined that the current time is within the validity period, substituting the obtained performance parameter with the performance parameter expectation.

24. The apparatus of any one of claims 17 to 23, wherein the operations further comprise:

obtaining a performance parameter associated with the database instance;

determining whether the database instance is associated with a corresponding performance parameter expectation; and

after it is determined that the database instance is associated with a corresponding performance parameter expectation, substituting the obtained performance parameter associated with the database instance with the performance parameter expectation.

25. The apparatus of claim 24, wherein the performance parameter expectation is associated with a validity period, and the operations further comprise:

determining whether a current time is within the validity period; and

after it is determined that the current time is within the validity period, substituting the obtained performance parameter associated with the database instance with the performance parameter expectation.

26. The apparatus of any one of claims 17 to 25, wherein the database server includes a plurality of database instances interrelated to one another, and the operations further comprise:

determining a performance parameter associated with each of the plurality of database instances;

determining a largest performance parameter from the plurality of performance parameters respectively determined for the plurality of database instances; and

substituting the performance parameter associated with each of the plurality of database instances with the largest performance parameter.

27. The apparatus of claim 26, wherein the plurality of database instances include a master database instance and a backup of the master database instance.

28. The apparatus of any one of claims 19 to 21, wherein the operations further comprise:

obtaining a performance parameter associated with each of a plurality of database servers;

determining a list of one or more database servers whose performance parameters exceed an upper threshold of the predetermined range;

selecting a first database server from the list;

selecting one or more database instances operating on the first database server for migration, such that after emigrating the selected one or more database instances from the first database server, the performance parameter of the first database server falls within the predetermined range; and

migrating the one or more database instances to one or more receiving database

servers.

29. The apparatus of claim 28, wherein the operation of selecting the first database server includes:

sorting the one or more database servers in the list according to their respective workloads indicated by the corresponding performance parameters; and

selecting the first database server from the sorted one or more database servers, wherein the first database server has a heaviest workload.

30. The apparatus of claim 28 or 29, wherein the operation of selecting the one or more database instances includes:

sorting database instances operating on the first database server according to their respective workloads;

selecting a first database instance from the sorted database instances, wherein the first database instance has a lowest workload; and

selecting additional one or more database instances from the sorted database instances for migration according to an ascending order of workload until the performance parameter of the first database server falls within the predetermined range after migration.

31. The apparatus of any one of claims 19 to 21 and 28 to 30, wherein the operations further comprise:

obtaining a performance parameter associated with each of a plurality of database servers;

determining a list of one or more database servers whose performance parameters are below a lower threshold of the predetermined range;

selecting a first database server from the list;
migrating all database instances operating on the first database server to one or more receiving database servers; and
placing the first database server in an offline mode after migration.

32. The apparatus of claim 31, wherein the operations further comprise:
sorting the one or more database servers in the list according to their respective workloads indicated by the corresponding performance parameters; and
selecting the first database server from the sorted one or more database servers, wherein the first database server has a lightest workload.

33. A non-transitory computer-readable medium that stores a set of instructions that is executable by at least one processor of a device to cause the device to perform a method for performing database scheduling, the method comprising:

obtaining a performance parameter associated with a database server;
determining whether the performance parameter matches a preset condition; and
after it is determined that the performance parameter matches the preset condition:
selecting a database instance operating on the database server; and
migrating the database instance to a receiving database server.

34. The computer-readable medium of claim 33, wherein:
the performance parameter includes at least one of: CPU usage, memory usage, disk I/O speed, network transmission speed, or disk usage.

35. The computer-readable medium of claim 33 or 34, wherein the preset condition includes:

the performance parameter falls within a predetermined range; or

the performance parameter does not fall within a predetermined range.

36. The computer-readable medium of claim 35, wherein:

the predetermined range includes an upper threshold; and

a performance parameter associated with the receiving database server is below the upper threshold.

37. The computer-readable medium of claim 35, wherein:

the predetermined range includes a lower threshold; and

a performance parameter associated with the receiving database server is below the lower threshold.

38. The computer-readable medium of any one of claims 33 to 37, wherein the set of instructions that is executable by the at least one processor of the device to cause the device to further perform:

determining whether the database server is associated with a corresponding performance parameter expectation; and

after it is determined that the database server is associated with the corresponding performance parameter expectation, substituting the obtained performance parameter with the performance parameter expectation.

39. The computer-readable medium of claim 38, wherein the performance parameter expectation is associated with a validity period, and wherein the set of instructions that is executable by the at least one processor of the device to cause the device to further perform:

determining whether a current time is within the validity period; and

after it is determined that the current time is within the validity period, substituting the obtained performance parameter with the performance parameter expectation.

40. The computer-readable medium of any one of claims 33 to 39, wherein the set of instructions that is executable by the at least one processor of the device to cause the device to further perform:

obtaining a performance parameter associated with the database instance;

determining whether the database instance is associated with a corresponding performance parameter expectation; and

after it is determined that the database instance is associated with a corresponding performance parameter expectation, substituting the obtained performance parameter associated with the database instance with the performance parameter expectation.

41. The computer-readable medium of claim 40, wherein the performance parameter expectation is associated with a validity period, and wherein the set of instructions that is executable by the at least one processor of the device to cause the device to further perform:

determining whether a current time is within the validity period; and

after it is determined that the current time is within the validity period, substituting the obtained performance parameter associated with the database instance with the

performance parameter expectation.

42. The computer-readable medium of any one of claims 33 to 41, wherein the database server includes a plurality of database instances interrelated to one another, and wherein the set of instructions that is executable by the at least one processor of the device to cause the device to further perform:

determining a performance parameter associated with each of the plurality of database instances;

determining a largest performance parameter from the plurality of performance parameters respectively determined for the plurality of database instances; and

substituting the performance parameter associated with each of the plurality of database instances with the largest performance parameter.

43. The computer-readable medium of claim 42, wherein the plurality of database instances include a master database instance and a backup of the master database instance.

44. The computer-readable medium of any one of claims 35 to 37, wherein the set of instructions that is executable by the at least one processor of the device to cause the device to further perform:

obtaining a performance parameter associated with each of a plurality of database servers;

determining a list of one or more database servers whose performance parameters exceed an upper threshold of the predetermined range;

selecting a first database server from the list;

selecting one or more database instances operating on the first database server for migration, such that after emigrating the selected one or more database instances from the first database server, the performance parameter of the first database server falls within the predetermined range; and

migrating the one or more database instances to one or more receiving database servers.

45. The computer-readable medium of claim 44, wherein selecting the first database server includes:

sorting the one or more database servers in the list according to their respective workloads indicated by the corresponding performance parameters; and

selecting the first database server from the sorted one or more database servers, wherein the first database server has a heaviest workload.

46. The computer-readable medium of claim 44 or 45, wherein selecting the one or more database instances includes:

sorting database instances operating on the first database server according to their respective workloads;

selecting a first database instance from the sorted database instances, wherein the first database instance has a lowest workload; and

selecting additional one or more database instances from the sorted database instances for migration according to an ascending order of workload until the performance parameter of the first database server falls within the predetermined range after migration.

47. The computer-readable medium of any one of claims 35 to 37 and 44 to 46, wherein the set of instructions that is executable by the at least one processor of the device to cause the device to further perform:

obtaining a performance parameter associated with each of a plurality of database servers;

determining a list of one or more database servers whose performance parameters are below a lower threshold of the predetermined range;

selecting a first database server from the list;


migrating all database instances operating on the first database server to one or more receiving database servers; and


placing the first database server in an offline mode after migration.


48. The computer-readable medium of claim 47, wherein the set of instructions that is executable by the at least one processor of the device to cause the device to further perform:


sorting the one or more database servers in the list according to their respective workloads indicated by the corresponding performance parameters; and


selecting the first database server from the sorted one or more database servers, wherein the first database server has a lightest workload.


100**FIG. 1**


**FIG. 2**


**FIG. 3**


**FIG. 4**


500**FIG. 5**

600**FIG. 6**

700**FIG. 7**

800**FIG. 8**

900**FIG. 9**

1000**FIG. 10**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 16/41500

A. CLASSIFICATION OF SUBJECT MATTER IPC(8) - G06F 17/00 (2016.01) CPC - G06F 17/30578; G06F 12/0804; G06F 12/0875 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(8): G06F 17/00 (2016.01) CPC: G06F 17/30578; G06F 12/0804; G06F 12/0875 Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched IPC(8): G06F 17/00 (2016.01); CPC: G06F 17/30578; G06F 12/0804; G06F 12/0875; G06F 17/30; G06F 17/30176; USPC: 707/618; 707/E17.032; 707/966 (keyword limited; terms below) Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) PatBase; Google(Web); Search terms used: scheduling database performance parameter server match preset condition instance migrating														
C. DOCUMENTS CONSIDERED TO BE RELEVANT <table border="1"> <thead> <tr> <th>Category*</th> <th>Citation of document, with indication, where appropriate, of the relevant passages</th> <th>Relevant to claim No.</th> </tr> </thead> <tbody> <tr> <td>X</td> <td>US 2013/0268734 A1 (Hotz et al.) 10 October 2010 (10.10.2010), entire document especially abstract; Fig. 2, 3; para [0024]-[0033], [0064]-[00077], [0081]-[0094]</td> <td>1-5, 17-21, 33-37</td> </tr> <tr> <td>A</td> <td>US 2014/0180664 A1 (Kochunni et al.) 26 June 2014 (26.06.2014), entire document</td> <td>1-5, 17-21, 33-37</td> </tr> <tr> <td>A</td> <td>US 2009/0106571 A1 (Low et al.) 23 April 2009 (23.04.2009), entire document especially</td> <td>1-5, 17-21, 33-37</td> </tr> </tbody> </table>			Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.	X	US 2013/0268734 A1 (Hotz et al.) 10 October 2010 (10.10.2010), entire document especially abstract; Fig. 2, 3; para [0024]-[0033], [0064]-[00077], [0081]-[0094]	1-5, 17-21, 33-37	A	US 2014/0180664 A1 (Kochunni et al.) 26 June 2014 (26.06.2014), entire document	1-5, 17-21, 33-37	A	US 2009/0106571 A1 (Low et al.) 23 April 2009 (23.04.2009), entire document especially	1-5, 17-21, 33-37
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.												
X	US 2013/0268734 A1 (Hotz et al.) 10 October 2010 (10.10.2010), entire document especially abstract; Fig. 2, 3; para [0024]-[0033], [0064]-[00077], [0081]-[0094]	1-5, 17-21, 33-37												
A	US 2014/0180664 A1 (Kochunni et al.) 26 June 2014 (26.06.2014), entire document	1-5, 17-21, 33-37												
A	US 2009/0106571 A1 (Low et al.) 23 April 2009 (23.04.2009), entire document especially	1-5, 17-21, 33-37												
☐ Further documents are listed in the continuation of Box C. ☐														
* Special categories of cited documents: <table border="0"> <tr> <td style="vertical-align: top;"> "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 </td> <td style="vertical-align: top;"> "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 </td> </tr> </table>			"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										
"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 11 September 2016		Date of mailing of the international search report 04 OCT 2016												
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-8300		Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774												

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 16/41500

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☒ Claims Nos.: 6-16, 22-32, 38-48
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.