



(51) International Patent Classification:

G06F 9/46 (2006.01) G06F 17/30 (2006.01)
G06F 9/50 (2006.01) H04L 12/24 (2006.01)
G06F 15/173 (2006.01)

(21) International Application Number:

PCT/US2017/036894

(22) International Filing Date:

09 June 2017 (09.06.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

201610410884.7 13 June 2016 (13.06.2016) CN

(71) Applicant: ALIBABA GROUP HOLDING LIMITED

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

(72) Inventors: ZHANG, Yang; Alibaba Group Legal Department,

5/F, Building 3, No. 969 West Wen Yi Road, Yu Hang District, Hangzhou, 311121 (CN). FENG, Yihui; Alibaba Group Legal Department, 5/F, Building 3, No. 969 West Wen Yi Road, Yu Hang District, Hangzhou, 311121 (CN). OUYANG, Jin; Alibaba Group Legal Department, 5/F, Building 3, No. 969 West Wen Yi Road, Yu Hang

District, Hangzhou, 311121 (CN). HAN, Qiaohuan; Alibaba Group Legal Department, 5/F, Building 3, No. 969 West Wen Yi Road, Yu Hang District, Hangzhou, 311121 (CN). WANG, Fang; Alibaba Group Legal Department, 5/F, Building 3, No. 969 West Wen Yi Road, Yu Hang District, Hangzhou, 311121 (CN).

(74) Agent: NELSON, Brett, L. et al.; Lee & Hayes, PLLC, 601 W. Riverside Ave, Suite 1400, Spokane, WA 99201 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, 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,

(54) Title: PRIORITY-BASED RESOURCE ALLOCATION




FIG. 2

(57) Abstract: A priority-based resource allocation method, includes accepting a resource application submitted by a job, the resource application including resource demand information and job priority information; determining, according to the resource demand information of the resource application, whether remaining resources of a system meet the resource application, and traversing, in an allocated resource application queue when the remaining resources do not meet the resource application, allocated resource applications having job priorities lower than that of the resource application; using the sum of system resources occupied by all traversed resource applications plus the remaining resources as available resources; and stopping traversing when the available resources meet the resource application, and allocating the available resources to the resource application. The technical solution of the present disclosure enables a resource application having a high job priority to preempt resources of a resource application having a low job priority.

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))*

PRIORITY-BASED RESOURCE ALLOCATION

Cross Reference to Related Patent Applications

This application claims priority to Chinese Patent Application No. 201610410884.7,
5 filed on June 13, 2016, entitled “PRIORITY-BASED RESOURCE ALLOCATION
METHOD, APPARATUS AND DEVICE,” which is hereby incorporated by reference in its
entirety.

Technical Field

10 The present disclosure relates to the field of computer technologies, and, in particular,
to a priority-based resource allocation method, and corresponding to the above method, the
present disclosure further relates to a priority-based resource allocation apparatus and device.

Background

15 In a distributed system, there are generally three roles: a resource scheduler, a job
manager and subordinate job nodes, and machine nodes. When a user needs to submit a job,
the job manager may apply for a certain amount of machine resources, such as CPU and
memory, from the resource scheduler for a job node to use. After receiving a job resource
application, the resource scheduler may calculate, according to the remaining available
20 resources in the distributed system, a list of available resources that can be allocated to the
job manager, such that a 1-core CPU and 1 GB memory are available on a machine A, and a
2-core CPU and 2 GB memory are available on a machine B. Then, the resource scheduler
may send the resource list to the job manager and a corresponding machine node.

In a distributed system, there may be tens of thousands of job managers, and job
25 priorities of resource applications of these job managers are generally different. Some job
managers have resource applications with higher job priorities, and need to be allocated with
resources preferentially by the resource scheduler. Some job managers have resource
applications with lower job priorities, indicating that they can tolerate that the resource
scheduler preferentially allocates resources to other resource applications having higher job
30 priorities. Such a situation that resource applications queue up due to exhaustion of overall
cluster resources may usually occur in a distributed system. A common solution is queuing

requests according to priorities, and then periodically extracting a resource request from the queue for resubmission. As shown in FIG. 1, in a waiting queue, a resource request having a high job priority may be queued in the front, and a resource request having a low priority may be queued behind. In FIG. 1, for example, the waiting queue includes a resource application 1 with priority 3 (102), a resource application 2 with priority 2 (104), and a resource application 3 with priority 1 (106).

However, the above method cannot solve the following problems: the resource scheduler receives a resource application having an extremely high priority in the case where the overall cluster resources are exhausted, and this resource application needs to obtain resources immediately; however, because it needs to take quite a long period of time to execute all resource applications that have been met, even though the job priority of the resource application is very high, resources cannot be allocated to the resource application until other resource applications having low job priorities return resources. In this case, the resource scheduler needs to support that a resource application having a high job priority can acquire resources that have been obtained by a resource application having a low job priority.

Summary

This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify all key features or essential features of the claimed subject matter, nor is it intended to be used alone as an aid in determining the scope of the claimed subject matter. The term “technique(s) or technical solution(s)” for instance, may refer to apparatus(s), system(s), method(s) and/or computer-readable instructions as permitted by the context above and throughout the present disclosure.

The present disclosure provides a priority-based resource allocation method, enabling a resource application having a high job priority to preempt resources of a resource application having a low job priority, such that system resources are effectively utilized.

The present disclosure provides a priority-based resource allocation method, including:

accepting a resource application submitted by a job, the resource application including resource demand information and job priority information;

determining, according to the resource demand information of the resource application, whether remaining resources of a system meet the resource application, and sequentially traversing, in an allocated resource application queue when the remaining resources do not meet the resource application, allocated resource applications having job priorities lower than that of the resource application, the allocated resource applications being resource applications that have been completely or partially allocated with resources;

forming traversed resource applications by using all the traversed allocated resource applications, and using the sum of system resources occupied by all the traversed resource applications plus the remaining resources as available resources; and

determining, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application, stopping traversing when the available resources meet the resource application, and allocating the available resources to the resource application.

Optionally, in the process of determining, according to the resource demand information of the resource application, whether remaining resources of a system meet the resource application, if the remaining resources meet the resource application, corresponding resources are allocated to the job according to the resource demand information of the resource application; and the resource application is recorded into the allocated resource application queue as the allocated resource application.

Optionally, the determining, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application further includes: if the available resources cannot meet the resource application, continuing traversing until the available resources meet the resource application.

Optionally, the determining, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application further includes: allocating partial resources to the resource application if the available resources still cannot meet the resource application after all the allocated resource applications having low job priorities are traversed.

Optionally, the stopping traversing when the available resources meet the resource application, and allocating the available resources to the resource application includes:

stopping traversing when the available resources meet the resource application, and allocating the available resources to the resource application; and allocating, according to a

sequence of the job priorities from high to low, second remaining resources after the allocation to the traversed resource applications, the second remaining resources being system resources that remain after the available resources are allocated to the resource application.

5 Optionally, the stopping traversing when the available resources meet the resource application, and allocating the available resources to the resource application includes:

 stopping traversing if the available resources meet the resource application, and further determining whether the second remaining resources meet any one of the traversed resource applications, if yes, allocating the available resources to the resource application,
10 and allocating the second remaining resources to any one of the traversed resource applications.

 Optionally, the stopping traversing when the available resources meet the resource application, and allocating the available resources to the resource application includes:

 stopping traversing if the available resources meet the resource application, and
15 further determining whether the second remaining resources meet any one of the traversed resource applications, if yes, allocating remaining available resources, excluding resources occupied by any one of the traversed resource applications, to the resource application.

 Optionally, the allocating partial resources to the resource application includes:

 allocating the remaining resources to the resource application, an unmet part of the
20 resource application entering a to-be-allocated resource application queue to wait for next resource allocation; and

 allocating the remaining part of the available resources after the available resources are allocated to the resource application to the traversed resource applications according to a sequence of the job priorities from high to low.

25 Optionally, the allocating partial resources to the resource application includes:

 keeping system resources occupied by the allocated resource applications unchanged, and allocating the remaining resources to the resource application when the available resources still cannot meet the resource application after all the allocated resource applications having low job priorities are traversed, an unmet part of the resource application
30 entering a to-be-allocated resource application queue to wait for next resource allocation.

 Optionally, the resource application includes a resource application name, a resource application volume, a resource application quantity, and a job priority;

the resource application name is a label name of the resource application;

the resource application volume is the minimum resource occupation needed for running the job corresponding to the resource application;

the resource application quantity is the quantity of the resource application volumes that the resource application applies for; and

the job priority is a definite quantity of priority levels defined by a user; if the priority level is higher, it indicates that the corresponding job is more important, and the resource application should be met more preferentially; if the priority levels are the same, the resource application submitted earlier in a time sequence should be met more preferentially.

Optionally, the job priorities are clustered into different priority intervals according to a set dimension; and the sequentially traversing, in an allocated resource queue, allocated resource applications having job priorities lower than that of the job resource application includes:

traversing, in the allocated resource queue according to a sequence of the job priorities from low to high, all the allocated resource applications of which priority intervals are lower than a priority interval of the resource application.

Optionally, the sequentially traversing, in an allocated resource queue, allocated resource applications having job priorities lower than that of the resource application includes:

traversing, in the allocated resource queue according to a sequence of the job priorities from low to high, all allocated resource applications having job priorities lower than that of the resource application.

Optionally, the allocated resource queue is a set of allocated resource applications that have been completely or partially allocated with resources and are arranged according to the sequence of the job priorities.

Optionally, the allocated resource application includes the resource application name, the resource application volume, an allocated resource quantity, allocated resources, and the job priority;

the allocated resource quantity is the quantity of the resource application volumes that have been allocated to the resource application; and

the allocated resources are a product of the resource application volume and the allocated resource quantity.

Optionally, the to-be-allocated resource application queue is a set of to-be-allocated resource applications that are completely or partially not allocated with resources and arranged according to the job priorities; the application includes the resource application name, the resource application volume, the resource application quantity, a to-be-allocated resource quantity, and the job priority.

A priority-based resource allocation apparatus includes: a resource application receiving unit, and a resource preemption unit;

the resource application receiving unit is configured to accept a resource application submitted by a job, the resource application including resource demand information and job priority information; and

determine, according to the resource demand information of the resource application, whether remaining resources of a system meet the resource application, and if the remaining resources do not meet the resource application, trigger the resource preemption unit.

Optionally, the priority-based resource allocation apparatus further includes a resource allocation unit;

it is determined, according to the resource demand information of the resource application, whether the remaining resources of the system meet the resource application, and if the remaining resources meet the resource application, the resource allocation unit is entered; and

the resource allocation unit is configured to allocate, according to the resource demand information of the job, corresponding resources to the resource application; and record the resource application into an allocated resource application queue as an allocated resource application.

Optionally, the resource preemption unit includes a traversing sub-unit, and a traversing post-processing sub-unit;

the traversing sub-unit is configured to traverse, in the allocated resource queue according to a sequence of the job priorities from low to high, the allocated resource applications having job priorities lower than that of the resource application; determine, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application; stop traversing if the available resources meet the resource application, and allocate the available resources to the job; and continue traversing, if the

available resources cannot meet the resource application, until the available resources meet the resource application; and

the traversing post-processing sub-unit is configured to process a situation in which the available resources still cannot meet the resource application after the traversing.

5 Optionally, in the traversing sub-unit, the stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application includes:

 stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application; and allocating, according to a
10 sequence of the job priorities from high to low, second remaining resources after the allocation to the traversed resource applications.

 Optionally, in the traversing sub-unit, the stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application includes:

15 stopping traversing if the available resources meet the resource application, and further determining whether the second remaining resources meet any one of the traversed resource applications, if yes, allocating the available resources to the resource application, and allocating the second remaining resources to any one of the traversed resource applications.

20 Optionally, in the traversing sub-unit, the stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application includes:

 stopping traversing if the available resources meet the resource application, and further determining whether the second remaining resources meet any one of the traversed
25 resource applications, if yes, allocating remaining available resources, excluding resources occupied by any one of the traversed resource applications, to the resource application.

 Optionally, the traversing post-processing sub-unit is configured to:

 allocate the available resources to the resource application when the available resources still cannot meet the resource application after all the allocated resource
30 applications having low job priorities are traversed, an unmet part of the resource application entering a to-be-allocated resource application queue to wait for next resource allocation; and

allocate the remaining part of the available resources after the available resources are allocated to the resource application to the allocated traversed resource applications, of which resources are preempted, according to a sequence of the job priorities from high to low.

Optionally, the traversing post-processing sub-unit is configured to:

5 keep system resources occupied by the allocated resource applications unchanged, and allocate the remaining resources to the resource application when the available resources still cannot meet the resource application after all the allocated resource applications having low job priorities are traversed, an unmet part of the resource application entering the to-be-allocated resource application queue to wait for next resource allocation.

10 Optionally, the resource application submitted by the resource application receiving unit includes: a resource application name, a resource application volume, a resource application quantity, and a job priority.

An electronic device is provided, where the electronic device includes:

a display;

15 a processor; and

a memory configured to store a priority-based resource allocation program, the program, when being read and executed by the processor, performing the following operations:

20 accepting a resource application submitted by a job, the resource application including resource demand information and job priority information;

25 determining, according to the resource demand information of the resource application, whether remaining resources of a system meet the resource application, and if yes, allocating corresponding resources to the job according to the resource demand information of the resource application; and recording the resource application into an allocated resource application queue as an allocated resource application;

if no, traversing, in the allocated resource application queue according to a sequence of job priorities from low to high, the allocated resource applications having job priorities lower than that of the resource application;

30 forming traversed resource applications by using all the traversed allocated resource applications, and using the sum of system resources occupied by all the traversed resource applications plus the remaining resources as available resources; determining, every time one of the allocated resource applications is traversed, whether the available resources meet the

resource application; stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application; continuing traversing, if the available resources cannot meet the resource application, until the available resources meet the resource application; or continuing traversing till all the allocated resource applications having low job priorities are traversed, and allocating resources to the resource application.

Optionally, the electronic device includes processors and one or more memories stored thereon a priority-based resource allocation program that, when executed by the processor, performing operations comprising:

receiving a resource application submitted by a job, the resource application including resource demand information and job priority information, the job priority information including a priority level and a priority interval that includes multiple priority levels;

determining, according to the resource demand information of the resource application, whether remaining resources meet the resource application;

sequentially traversing, in an allocated resource application queue, allocated resource applications having priority intervals lower than that of the resource application in response to determining that the remaining resources do not meet the resource application, the allocated resource applications being resource applications that have been completely or partially allocated with resources;

determining not to preempt or traverse an allocated resource application having a priority interval same as that of the resource application and a priority level different from that of the resource application;

forming traversed resource applications by using the traversed allocated resource applications;

using a sum of system resources occupied by the traversed resource applications plus the remaining resources as available resources;

determining, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application;

stopping traversing when the available resources meet the resource application; and

allocating the available resources to the resource application.

The allocated resource application having the priority interval same as that of the resource application and the priority level different from that of the resource application includes:

the allocated resource application having the priority interval same as that of the resource application and the priority level lower than that of the resource application.

the allocated resource application having the priority interval same as that of the resource application and the priority level same as that of the resource application; or

the allocated resource application having the priority interval same as that of the resource application and the priority level higher than that of the resource application.

Compared with the conventional techniques, the present disclosure has the following advantages:

The present disclosure provides a priority-based resource allocation method, apparatus and device. The resource allocation method, after accepting a resource application submitted by a job, determines, according to resource demand information of the resource application, whether remaining resources of a system meet the resource application, and when the remaining resources do not meet the resource application, sequentially traverses, in an allocated resource application queue, allocated resource applications having job priorities lower than that of the resource application; uses the sum of system resources occupied by all traversed allocated resource applications plus the remaining resources as available resources, and allocates the available resources to the resource application having a high job priority. The technical solution provided in the present disclosure solves the defect in the conventional techniques that a resource application having a high job priority must wait in a queue when system resources are exhausted, such that the resource application having a high job priority obtains, to the maximum extent, resources needed for running.

In addition, in an example embodiment of the present disclosure, second remaining resources of the available resources, which remain after the available resources are allocated to the resource application, are allocated to the allocated resource applications whose resources are occupied, according to a sequence of job priorities from high to low. Thus, the system resources are effectively utilized.

Brief Description of the Drawings

In order to describe the technical solutions in the example embodiments of the present disclosure or the conventional techniques more clearly, accompanying drawings needed for describing the example embodiments are introduced briefly in the following. It is apparent that the accompanying drawings in the following description are only example embodiments described in the present disclosure.

FIG. 1 is a schematic diagram of a resource application queue as a comparison art;

FIG. 2 is a flowchart of a priority-based resource allocation method according to an example embodiment of the present disclosure;

FIG. 3 is a schematic diagram of an allocated resource application queue according to an example embodiment of the present disclosure;

FIG. 4 is a schematic diagram of a priority-based resource allocation apparatus according to an example embodiment of the present disclosure;

FIG. 5 is a schematic diagram of the resource preemption unit according to an example embodiment of the present disclosure; and

FIG. 6 is a schematic diagram of an electronic device according to an example embodiment of the present disclosure.

Detailed Description

Many specific details are illustrated in the following descriptions to facilitate full understanding of the present disclosure. However, the present disclosure can be implemented in many other manners different from that described herein, and persons of ordinary skill in the art may perform similar promotion without departing from the essence of the present disclosure; therefore, the present disclosure is not limited by specific implementations disclosed in the following.

Example embodiments of the present disclosure provide a priority-based resource allocation method and a priority-based resource allocation apparatus; the example embodiments of the present disclosure also provide an electronic device. The method, apparatus, and electronic device are described in detail one by one in the following example embodiments.

Before specific steps of this example embodiment are described in detail, a running environment of resource allocation involved in this technical solution is described briefly first.

In a distributed system, there are generally three roles: a resource scheduler, a job manager and subordinate job nodes, and machine nodes. When a user needs to submit a job, the job manager may apply for a certain amount of machine resources, such as a CPU and a memory, from the resource scheduler for the job node to use. After receiving a job resource application, the resource scheduler may calculate, according to remaining available resources in the distributed system, a list of available resources that can be allocated to the job manager, for example, a 1-core CPU and 1 GB memory are available on a machine A, and a 2-core CPU and 2 GB memory are available on a machine B. Then, the resource scheduler may send the resource list to the job manager and a corresponding machine node. After receiving the resource list, the job manager may allocate the job node to the corresponding machine node to execute the job program. When execution of the job node is completed, the job manager may notify the resource manager to return resources of the job node, and the resource manager, after receiving a return application, would notify the machine node to release the resources of the job node. The released resources may be reallocated to another job.

Job priorities of resource applications of the job manager are different. When system resources are sufficient, the resource scheduler may meet resource demands of all resource applications, and record, according to a sequence of job priorities from high to low, the resource applications that have been allocated with resources into an allocated resource application queue. In the case where the overall system resources are exhausted, when the resource scheduler receives a resource application having a high job priority, determining how to allocate resources to the resource application having a high job priority is the problem to be solved in the technical solution provided in the present disclosure.

The present disclosure provides a priority-based resource allocation method, and an example embodiment of the resource allocation method is as follows:

FIG. 2 is a flowchart of a priority-based resource allocation method according to an example embodiment of the present disclosure. Step S202, receiving a resource application submitted by a job, the resource application including resource demand information and job priority information.

As described above, when a user needs to submit a job, the user may submit the resource application to the resource scheduler through the job manager. The resource application includes a resource application name, a resource application volume, a resource application quantity, and a job priority.

5 The resource application name is a label name of the resource application. The resource application name is the same as a label name of the job.

10 The resource application volume is the minimum resource occupation needed for running the job corresponding to the resource application. The minimum resource occupation is the minimum requirement on a processor, a memory, and the like when the job runs normally in a job node in a distributed system.

The resource application quantity is the quantity of the resource application volumes that the resource application applies for. The resource application quantity may be understood as the quantity of job nodes needed for normally running the job in the distributed system.

15 For example, 100 CPUs and 100G memory are needed for running a user job, the user job has 100 job nodes to be executed, and the execution of each job node needs 1 CPU and 1G memory; therefore, 1 CPU and 1G memory are the resource application volume, and 100 is the resource application quantity.

20 The resource application volume and the resource application quantity form the resource demand information of the resource application. A product of the resource application volume and the resource application quantity form a total resource demand amount of the resource application.

25 The job priority is a definite quantity of priority levels defined by the user; if the priority level is higher, it indicates that the corresponding job is more important, and the resource application should be met more preferentially; if the priority levels are the same, time sequences thereof are compared, and the resource application submitted earlier has a higher priority.

Step S204, determining, according to the resource demand information of the resource application, whether remaining resources of the system meet the resource application, and when the remaining resources do not meet the resource application, performing step S206.

30 The resource demand information of the resource application includes the resource application volume and the resource application quantity. The resource scheduler compares

the resource demand with the size of the remaining resources, and if the remaining resources are less than the resource demand, step S206 is performed.

For example, when the resource scheduler compares the resource demand with the size of the remaining resources, there is another situation, that is, the remaining resources are greater than or equal to the resource demand, and in this case, corresponding resources are allocated to the job according to the resource demand information of the resource application, and the resource application is recorded into an allocated resource application queue as an allocated resource application.

The allocated resource queue is a set of allocated resource applications that have been completely or partially allocated with resources and are arranged according to the job priorities. The allocated resource application includes the resource application name, the resource application volume, an allocated resource quantity, allocated resources, and the job priority;

wherein, the resource application name, the resource application volume, and the job priority are consistent with corresponding items of the resource application; the allocated resource quantity is the quantity of the resource application volumes that have been allocated to the resource application; and the allocated resources are a product of the resource application volume and the allocated resource quantity.

FIG. 3 is a diagram of an allocated resource application queue according to an example embodiment of the present disclosure.

As shown in FIG. 3, the allocated resource queue includes 3 allocated resource applications, and the resource application names are respectively: resource application A, resource application B, and resource application C, wherein other items of the resource application A include:

resource application volume: CPU: 1, memory (MEM): 1
allocated resource quantity: 20
allocated resources: CPU: 20, memory (MEM): 20
job priority: 3

Step S206, sequentially traversing, in the allocated resource application queue, all current allocated resource applications having job priorities lower than that of the resource application; forming traversed resource applications by using all the traversed allocated

resource applications, and using the sum of system resources occupied by all the traversed resource applications plus the remaining resources as available resources.

All the current allocated resource applications having job priorities lower than that of the resource application are sequentially traversed according to a sequence of the job priorities from low to high.

An example embodiment of the present disclosure clusters job priorities into different priority intervals according to a set dimension. Each of the priority intervals includes determined priority levels. Moreover, system resources occupied by the allocated resource application cannot be preempted by another resource application in the same priority interval.

In the example embodiment of the present disclosure, when the priority interval is taken into consideration, the sequentially traversing, in the allocated resource application queue, all the current allocated resource applications having job priorities lower than that of the resource application becomes: traversing, in the allocated resource queue according to a sequence of the job priorities from low to high, all the allocated resource applications of which priority intervals are lower than that of the resource application.

In the example embodiment of the present disclosure, the job priorities are classified into 10 levels from low to high, that is, a job priority 1, a job priority 2, ..., a job priority 10. The 10 levels are clustered into 3 priority intervals according to a sequence of the priorities from low to high: a priority interval I including the job priorities 1, 2, 3, and 4; a priority interval II including the job priorities 5, 6, and 7; a priority interval III including the job priorities 8, 9, and 10. If the job priority of the submitted resource application is 9, when the priority interval is taken into consideration, only the priority interval I and the priority interval II can be traversed, and the allocated resource applications having the job priority 8 will not be traversed.

The example embodiment of the present disclosure, during traversing, will not reclaim system resources occupied by the allocated resource applications that have been traversed, and the system resources are reclaimed only after a final resource allocation solution is determined. In some cases, the resource scheduler may, during traversing, simultaneously reclaim system resources occupied by the allocated resource applications that have been traversed.

Step S208, determining, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application; and when the

available resources meet the resource application, stopping traversing, and allocating the available resources to the resource application.

Every time one of the allocated resource applications is traversed, the system may determine whether the available resources meet the resource application:

5 If the available resources meet the resource application, the traversing is stopped, and the available resources are allocated to the resource application.

If the available resources cannot meet the resource application, the traversing is continued until the available resources meet the resource application.

10 The operation of stopping traversing and allocating the available resources to the resource application may have two cases. If the available resources just meet the resource demand of the resource application, the allocation ends. If there are second remaining resources after the available resources are allocated to the resource application, that is, system resources remaining after the available resources are allocated to the resource application, the second remaining resources are continuously allocated to the traversed resource applications.

15 In the example embodiment provided in the present disclosure, the second remaining resources are allocated to the traversed resource applications according to a sequence of the job priorities from high to low.

As an optional solution, there are the following two solutions for processing the available resources:

20 Solution 1: it is further determined whether the second remaining resources meet any one of the traversed resource applications, if yes, system resources occupied by the traversed resource application are reclaimed first, the available resources are then allocated to the resource application, and the second remaining resources are allocated to any one of the traversed resource applications.

25 Solution 2, it is further determined whether the second remaining resources meet any one of the traversed resource applications, and if yes, system resources occupied by any one of the traversed resource applications are reserved, and remaining available resources, excluding the resources occupied by any one of the traversed resource applications, are allocated to the resource application.

30 A difference between the solution 1 and the example embodiment lies in that: in the solution 1, system resources occupied by each traversed allocated resource application are reclaimed first, and then the available resources are allocated; in the example embodiment,

the system resources occupied by each traversed allocated resource application will not be reclaimed in advance; instead, the system resources occupied by all the traversed allocated resource applications are reclaimed after an allocation solution is determined.

5 The solution 2 aims to ensure preferential execution of the traversed resource application that has completely acquired resources needed.

If the available resources still cannot meet the resource application after all the allocated resource applications having low job priorities are traversed, partial resources are allocated to the resource application.

10 The allocating partial resources to the resource application includes: allocating the remaining resources to the resource application, an unmet part of the resource application entering a to-be-allocated resource application queue to wait for next resource allocation; and allocating the remaining part of the available resources after the available resources are allocated to the resource application to the traversed resource applications according to a sequence of the job priorities from high to low.

15 The to-be-allocated resource application queue is a set of to-be-allocated resource applications that are completely or partially not allocated with resources and arranged according to the job priorities; the to-be-allocated resource application includes the resource application name, the resource application volume, the resource application quantity, a to-be-allocated resource quantity, and the job priority.

20 The resource application name, the resource application volume, the resource application quantity, and the job priority are the same as those of the resource application.

The to-be-allocated resource quantity is the quantity of the resource application volumes needed by the unmet part of the resource application.

25 As an optional solution, when the available resources still cannot meet the resource application after all the allocated resource applications having low job priorities are traversed, system resources occupied by the allocated resource applications may keep unchanged, the remaining resources are allocated to the resource application, and an unmet part of the resource application enters a to-be-allocated resource application queue to wait for next resource allocation.

30 In the following, the whole process is illustrated by using a specific example with reference to FIG. 3.

Assume that total resources of the system are CPU: 100, memory: 100. The resource application A 302 occupies CPU: 20, memory: 20, has a resource application volume CPU: 1, memory: 1, and a resource application quantity 20, and has a job priority of 3. The resource application B 304 occupies CPU: 60, memory: 40, has a resource application volume: CPU :3, memory: 2, and a resource application quantity: 20, and has a job priority of 2. The resource application C 306 occupies CPU: 20, memory: 10, has a resource application volume: CPU: 2, memory: 1, and a resource application quantity: 10, and has a job priority of 1. Therefore, remaining resources of the system are CPU: 0, memory: 30.

At this point, a resource application E is received, of which the resource application volume is CPU: 1, memory: 1, the application quantity is 30, and the job priority is 4. As the remaining resources cannot meet the resource application E, the next step is performed to traverse allocated resource applications having job priorities lower than that of the resource application E according to a sequence of the job priorities from low to high in an allocated resource application queue 308. First, the resource application C 306 is traversed, and then, after the traversing, new available resources are CPU: 20 ($=0+20$), memory: 40 ($=30+10$). The available resources still cannot meet the resource application E, and therefore, the resource application B 304 is traversed continuously, such that new available resources are CPU: 80 ($=20+60$), memory: 80 ($=40+40$). The available resources meet the resource application E, and therefore, the available resources are allocated to the resource application E. After the allocation, remaining available resources, that is, second remaining resources, are CPU: 50 ($=80-30$), memory 50 ($=80-30$), and the second remaining resources are reallocated to the traversed resource applications, that is, the resource application B 304 and the resource application C 306, according to a sequence of the job priorities from high to low. The resource application B 304 is checked first, and 16 is obtained after the second remaining resources are divided by a larger resource application volume of all resource application volumes, i.e., 3, of the resource application B 304, that is, system resources that can be allocated to the resource application B 304 are: CPU: 48 ($=3*16$), memory: 32 ($=2*16$). System resources remaining after the second remaining resources are allocated, that is, third remaining resources, become: CPU: 2 ($=50-48$), memory: 18 ($=50-32$). An unmet part of the resource application B 304, that is, CPU: 12 ($=3*(20-16)$), memory: 8 ($=2*(20-16)$), enters the to-be-allocated resource application queue. Then, the resource application C 306 is checked, and the third remaining resources meet one resource application volume of the

resource application C 306. After the third remaining resources are allocated to the one resource application volume of the resource application C, remaining system resources become CPU: 0 ($=2-2*1$), memory: 17($=18-1*1$), and cannot be used for allocation. An unmet part of the resource application C 306, that is, CPU: 18 ($=2* (10-1)$), memory: 9 ($=1*(10-1)$), enters the to-be-allocated resource application queue.

The final result is: the resource application E obtains system resources CPU: 30, memory: 30; the resource application B obtains system resources CPU: 48, memory: 32; the resource application C obtains system resources CPU: 2, memory: 1; and the new remaining resources of the system are CPU: 0, memory: 17.

Finally, the resource scheduler will reclaim the resources occupied by the resource application B and the resource application C, and allocate the system resources according to the new solution above.

The example embodiment of the present disclosure provides a priority-based resource allocation method, and corresponding to the priority-based resource allocation method above, the present disclosure further provides a priority-based resource allocation apparatus.

FIG. 4 is a schematic diagram of a priority-based resource allocation apparatus 400 according to an example embodiment of the present disclosure.

The apparatus 400 includes one or more processor(s) 402 or data processing unit(s) and memory 404. The apparatus 400 may further include one or more input/output interface(s) 406, and network interface(s) 408. The memory 404 is an example of computer readable media.

The memory 404 may store therein a plurality of modules or units including a resource application receiving unit 410, and a resource preemption unit 412.

The resource application receiving unit 410 is configured to accept a resource application submitted by a job, the resource application including resource demand information and job priority information; determine, according to the resource demand information of the resource application, whether remaining resources of a system meet the resource application; and trigger the resource preemption unit 412 if the remaining resource do not meet the resource application.

The resource preemption unit 412 is configured to traverse, in the allocated resource application queue according to a sequence of the job priorities from low to high, allocated

resource applications having job priorities lower than that of the resource application, and reallocate the system resources according to a result after the traversing.

As shown in FIG. 4, the apparatus 400 may further include a resource allocation unit 414 stored on memory 404.

5 When the resource application receiving unit 410 determines, according to the resource demand information of the resource application, whether remaining resources of the system meet the resource application, if the remaining resources meet the resource application, the resource allocation unit 414 is triggered.

10 The resource allocation unit 414 is configured to allocate, according to the resource demand information of the job, corresponding resources to the resource application; and record the resource application into the allocated resource application queue as an allocated resource application.

FIG. 5 is a schematic diagram of the resource preemption unit 412 according to an example embodiment of the present disclosure.

15 As shown in FIG. 5, the resource preemption unit 412 includes a traversing sub-unit 502, and a traversing post-processing sub-unit 504.

20 The traversing sub-unit 502 is configured to traverse, in the allocated resource queue according to a sequence of the job priorities from low to high, the allocated resource applications having job priorities lower than that of the resource application; determine, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application; stop traversing if the available resources meet the resource application, and allocate the available resources to the job; and continue traversing, if the available resources cannot meet the resource application, until the available resources meet the resource application.

25 The traversing post-processing sub-unit 504 is configured to process a situation in which the available resources still cannot meet the resource application after the traversing.

Optionally, in the traversing sub-unit 502, the stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application includes:

30 stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application; and allocating, according to a

sequence of the job priorities from high to low, second remaining resources after the allocation to the traversed resource applications.

Optionally, in the traversing sub-unit 502, the stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application includes:

stopping traversing if the available resources meet the resource application, and further determining whether the second remaining resources meet any one of the traversed resource applications, if yes, allocating the available resources to the resource application, and allocating the second remaining resources to any one of the traversed resource applications.

Optionally, in the traversing sub-unit 502, the stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application includes:

stopping traversing if the available resources meet the resource application, and further determining whether the second remaining resources meet any one of the traversed resource applications, if yes, allocating remaining available resources, excluding resources occupied by any one of the traversed resource applications, to the resource application.

Optionally, the traversing post-processing sub-unit 504 is configured to:

allocate the available resources to the resource application when the available resources still cannot meet the resource application after all the allocated resource applications having low job priorities are traversed, an unmet part of the resource application entering a to-be-allocated resource application queue to wait for next resource allocation; and

allocate the remaining part of the available resources after the available resources are allocated to the resource application to the allocated traversed resource applications, of which resources are preempted, according to a sequence of the job priorities from high to low.

Optionally, the traversing post-processing sub-unit 504 is configured to:

keep system resources occupied by the allocated resource applications unchanged, and allocate the remaining resources to the resource application when the available resources still cannot meet the resource application after all the allocated resource applications having low job priorities are traversed, an unmet part of the resource application entering the to-be-allocated resource application queue to wait for next resource allocation.

The above example embodiments provide a priority-based resource allocation method and a priority-based resource allocation apparatus. Moreover, the present disclosure further provides an electronic device, and an example embodiment of the electronic device is described as follows:

5 FIG. 6 is a schematic diagram of an electronic device 600 according to an example embodiment of the present disclosure.

The electronic device 600 includes: a display 602, one or more processor(s) 604 or data processing unit(s) and memory 606. The electronic device 600 may further include one or more input/output interface(s) and network interface(s), not shown in FIG. 6

10 The memory 606 is configured to store a priority-based resource allocation program, the program, when being read and executed by the processor, performing the following operations:

 accepting a resource application submitted by a job, the resource application including resource demand information and job priority information;

15 determining, according to the resource demand information of the resource application, whether remaining resources of a system meet the resource application, and if yes, allocating corresponding resources to the job according to the resource demand information of the resource application; and recording the resource application into an allocated resource application queue as an allocated resource application;

20 if no, traversing, in the allocated resource application queue according to a sequence of job priorities from low to high, the allocated resource applications having job priorities lower than that of the resource application; and

 forming traversed resource applications by using all the traversed allocated resource applications, and using the sum of system resources occupied by all the traversed resource applications plus the remaining resources as available resources; determining, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application; stopping traversing if the available resources meet the resource application, and allocating the available resources to the resource application; continuing traversing, if the available resources cannot meet the resource application, until the available
30 resources meet the resource application; or continuing traversing till all the allocated resource applications having low job priorities are traversed, and allocating resources to the resource application.

In a typical configuration, a computer device includes one or more processors (CPUs), input/output interfaces, network interfaces, and memories.

The memory may include a volatile memory, a random access memory (RAM) and/or a non-volatile memory or the like in a computer readable medium, for example, a read-only memory (ROM) or a flash RAM. The memory is an example of the computer readable medium.

The computer readable medium includes non-volatile and volatile media as well as movable and non-movable media, and can implement information storage by means of any method or technology. Information may be a computer readable instruction, a data structure, and a module of a program or other data. A storage medium of a computer includes, for example, but is not limited to, a phase change memory (PRAM), a static random access memory (SRAM), a dynamic random access memory (DRAM), other types of RAMs, a ROM, an electrically erasable programmable read-only memory (EEPROM), a flash memory or other memory technologies, a compact disk read-only memory (CD-ROM), a digital versatile disc (DVD) or other optical storages, a cassette tape, a magnetic tape/magnetic disk storage or other magnetic storage devices, or any other non-transmission medium, and can be used to store information accessible to a computation device. According to the definition of this text, the computer readable medium does not include transitory media, such as a modulated data signal and a carrier.

Persons of ordinary skill in the art should understand that, the example embodiments of the present disclosure may be provided as a method, a system, or a computer program product. Therefore, the present disclosure may be implemented as a complete hardware example embodiment, a complete software example embodiment, or an example embodiment combining software and hardware. Moreover, the present disclosure may be a computer program product implemented on one or more computer usable storage media (including, but not limited to, a magnetic disk memory, a CD-ROM, an optical memory, and the like) including computer usable program code.

The present disclosure is disclosed in the example embodiments in the foregoing, but they are not intended to limit the present disclosure, and any persons skilled in the art can make possible variations and modifications without departing from the spirit and scope of the present disclosure. Therefore, the protection scope of the present disclosure should be subject to the scope defined by the claims of the present disclosure.

Claims

What is claimed is:

1. A method comprising:

5 receiving a resource application, the resource application including resource demand information and job priority information of a job;

determining, according to the resource demand information of the resource application, whether remaining resources meet the resource application;

10 sequentially traversing, in an allocated resource application queue, allocated resource applications having job priorities lower than that of the resource application in response to determining that the remaining resources do not meet the resource application, the allocated resource applications being resource applications that have been completely or partially allocated with resources;

15 forming traversed resource applications by using the traversed allocated resource applications;

using a sum of system resources occupied by the traversed resource applications plus the remaining resources as available resources;

determining, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application;

20 stopping traversing when the available resources meet the resource application; and allocating the available resources to the resource application.

2. The method of claim 1, further comprising:

25 allocating corresponding resources to the job according to the resource demand information of the resource application, in response to determining that the remaining resources meet the resource application; and

recording the resource application into the allocated resource application queue as one of the allocated resource applications.

3. The method of claim 1, further comprising continuing traversing, when the available resources do not meet the resource application, until the available resources meet the resource application.

5 4. The method of claim 1, wherein the determining, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application further includes:

allocating partial resources to the resource application in response to determining that the available resources still do not meet the resource application after the allocated resource
10 applications having low job priorities are traversed.

5. The method of claim 4, wherein the allocating the partial resources to the resource application includes:

allocating remaining resources to the resource application, an unmet part of the
15 resource application entering a to-be-allocated resource application queue to wait for a next resource allocation.

6. The method of claim 5, wherein the allocating the partial resources to the resource application further includes:

20 allocating a remaining part of the available resources, after the available resources are allocated to the resource application, to the traversed resource applications according to a sequence of the job priorities from high to low.

7. The method of claim 4, wherein the allocating the partial resources to the resource
25 application includes, when the available resources still do not meet the resource application after the allocated resource applications having low job priorities are traversed:

keeping system resources occupied by the allocated resource applications unchanged;
and

allocating the remaining resources to the resource application, an unmet part of the
30 resource application entering a to-be-allocated resource application queue to wait for a next resource allocation.

8. The method of claim 1, further comprising:

allocating, according to a sequence of the job priorities from high to low, second remaining resources to the traversed resource applications, the second remaining resources being resources that remain after the available resources are allocated to the resource application.

9. The method of claim 8, further comprising:

determining whether the second remaining resources meet a respective traversed resource application of the traversed resource applications; and

allocating the second remaining resources to the respective traversed resource application of the traversed resource applications, in response to determining that the second remaining resources meet the respective traversed resource application of the traversed resource applications.

10. The method of claim 8, further comprising:

determining whether the second remaining resources meet a respective traversed resource application of the traversed resource applications; and

allocating remaining available resources, excluding resources occupied by the traversed resource applications, to the resource application.

11. The method of claim 1, wherein:

the resource application includes a resource application name, a resource application volume, a resource application quantity, and a job priority;

the resource application name is a label name of the resource application;

the resource application volume is a minimum resource occupation needed for running the job corresponding to the resource application;

the resource application quantity is a quantity of the resource application volume that the resource application applies for; and

the job priority is a quantity of priority levels defined by a user.

12. The method of claim 11, wherein:

a higher priority level indicates that a corresponding job is more important and the resource application should be met more preferentially; and

5 when priority levels are the same, the resource application submitted earlier in a time sequence is met more preferentially.

13. The method of claim 1, wherein:

the job priorities are clustered into different priority intervals according to a set dimension; and

10 the sequentially traversing, in the allocated resource application queue, the allocated resource applications having job priorities lower than that of the resource application includes:

traversing in the allocated resource application queue, according to a sequence of the job priorities from low to high, the allocated resource applications of which priority intervals
15 are lower than a priority interval of the resource application.

14. The method of claim 13, wherein the sequentially traversing, in the allocated resource application queue, the allocated resource applications having job priorities lower than that of the resource application further includes:

20 determining not to preempt or traverse an allocated resource application of which priority interval is same as that of the resource application.

15. The method of claim 14, wherein the allocated resource application of which priority interval is same as that of the resource application and priority level is lower than the
25 resource application.

16. The method of claim 1, wherein the sequentially traversing, in the allocated resource application queue, the allocated resource applications having job priorities lower than that of the resource application further includes:

traversing, in the allocated resource queue, according to a sequence of the job priorities from low to high, the allocated resource applications having job priorities lower than that of the resource application.

17. The method of claim 1, wherein the allocated resource queue is a set of allocated resource applications that have been completely or partially allocated with resources and are arranged according to the sequence of the job priorities.

18. An apparatus comprising:

a resource application receiving unit that receives a resource application, the resource application including resource demand information and job priority information of a job, and determines, according to the resource demand information of the resource application, whether remaining resources of a system meet the resource application, and triggers a resource preemption unit in response to determining that the remaining resources do not meet the resource application, and triggers a resource allocation unit in response to determining that the remaining resources meet the resource application;

the resource allocation unit that allocates, according to the resource demand information of the job, corresponding resources to the resource application; and records the resource application into an allocated resource application queue as an allocated resource application; and

the resource preemption unit that traverses, in the allocated resource queue according to a sequence of the job priorities from low to high, the allocated resource applications having job priorities lower than that of the resource application; determines, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application, stops traversing when the available resources meet the resource application, allocates the available resources to the job, and continue traversing when the available resources do not meet the resource application, until the available resources meet the resource application.

19. An electronic device comprising:

one or more processors; and

one or more memories stored thereon a priority-based resource allocation program

5 that, when executed by the processor, performing operations comprising:

receiving a resource application, the resource application including resource demand information and job priority information of a job, the job priority information including a priority level and a priority interval that includes multiple priority levels;

10 determining, according to the resource demand information of the resource application, whether remaining resources meet the resource application;

sequentially traversing, in an allocated resource application queue, allocated resource applications having priority intervals lower than that of the resource application in response to determining that the remaining resources do not meet the resource application, the allocated resource applications being resource applications that have been completely or partially allocated with resources;

15 determining not to traverse an allocated resource application having a priority interval same as that of the resource application and a priority level different from or same as that of the resource application;

20 forming traversed resource applications by using the traversed allocated resource applications;

using a sum of system resources occupied by the traversed resource applications plus the remaining resources as available resources;

determining, every time one of the allocated resource applications is traversed, whether the available resources meet the resource application;

25 stopping traversing when the available resources meet the resource application;

and

allocating the available resources to the resource application.

20. The electronic device of claim 19, wherein the allocated resource application having the priority interval same as that of the resource application and the priority level different from or same as that of the resource application includes:

an allocated resource application having the priority interval same as that of the
5 resource application and the priority level lower than that of the resource application.




FIG. 1 (PRIOR ART)

2/6




FIG. 2

3/6

RESOURCE APPLICATION A

RESOURCE APPLICATION VOLUME: CPU: 1, MEMORY: 1

RESOURCE APPLICATION QUANTITY: 20

ALLOCATED RESOURCES: CPU: 20, MEMORY: 20

JOB PRIORITY: 3

302

RESOURCE APPLICATION B

RESOURCE APPLICATION VOLUME: CPU: 3, MEMORY: 2

RESOURCE APPLICATION QUANTITY: 20

ALLOCATED RESOURCES: CPU: 60, MEMORY: 40

JOB PRIORITY: 2

304

RESOURCE APPLICATION C

RESOURCE APPLICATION VOLUME: CPU: 2, MEMORY: 1

RESOURCE APPLICATION QUANTITY: 10

ALLOCATED RESOURCES: CPU: 20, MEMORY: 10

JOB PRIORITY: 1

306

ALLOCATED RESOURCE APPLICATION QUEUE

308

FIG. 3




FIG. 4




FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2017/036894

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 9/46; G06F 9/50; G06F 15/173; G06F 17/30; H04L 12/24 (2017.01)

CPC - G06F 9/46; G06F 9/50; G06F 9/5005; G06F 9/505; G06F 9/5072; G06F 9/5083; G06F 15/173; G06F 17/30; G06F 17/30474; G06F 17/30545; H04L 12/24 (2017.02)

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)

See Search History document

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

USPC - 709/223.000; 709/224.000; 709/226.000; 718/102.000; 718/103.000; 718/104.000 (keyword delimited)

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

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 9,128,883 B2 (GOKHALE et al) 08 September 2015 (08.09.2015) entire document	18
---		---
Y		1-17, 19, 20
Y	US 7,953,000 B2 (BAKER et al) 31 May 2011 (31.05.2011) entire document	1-17, 19, 20
Y	US 2013/0311999 A1 (FETTERMAN et al) 21 November 2013 (21.11.2013) entire document	12, 14, 15, 19, 20
A	BAKER. A stack-based resource allocation policy for realtime processes. 11th Real-Time Systems Symposium. 191-200, 1990. [retrieved on 27.07.2017]. Retrieved from the Internet. <URL:http://www.cs.cmu.edu/~ssaewong/research/srp.pdf>. entire document	1-20
A	US 9,003,037 B2 (BEVERIDGE) 07 April 2015 (07.04.2015) entire document	1-20
A	US 8,640,137 B1 (BOSTIC et al) 28 January 2014 (28.01.2014) entire document	1-20
A	US 2013/0311686 A1 (FETTERMAN et al) 21 November 2013 (21.11.2013) entire document	1-20
A	US 8,639,772 B2 (GENTILE et al) 28 January 2014 (28.01.2014) entire document	1-20
A	US 7,339,942 B2 (MORGAN et al) 04 March 2008 (04.03.2008) entire document	1-20
A	US 8,773,998 B2 (POLK et al) 08 July 2014 (08.07.2014) entire document	1-20
A	US 2016/0147564 A1 (RAYTHEON COMPANY) 26 May 2016 (26.05.2016) 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

28 July 2017

Date of mailing of the international search report

23 AUG 2017

Name and mailing address of the ISA/US

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

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2017/036894

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	URGAONKAR et al. Dynamic resource allocation and power management in virtualized data centers. 2010 IEEE Network Operations and Management Symposium. 1-8, 2010. [retrieved on 27.07.2017]. Retrieved from the Internet. <URL: https://www.researchgate.net/profile/Ulas_Kozat/publication/220708034_Dynamic_resource_allocation_and_power_management_in_virtualized_data_centers/links/0fcfd50ad7cfd2fc9000000.pdf >. entire document	1-20
A	US 5,963,911 A (WALKER et al) 05 October 1999 (05.10.1999) entire document	1-20