

(19) World Intellectual Property Organization  
International Bureau(43) International Publication Date  
27 December 2007 (27.12.2007)

PCT

(10) International Publication Number  
WO 2007/149188 A1(51) International Patent Classification:  
G06F 9/50 (2006.01)

(74) Agents: POMERENKE, Ronald M. et al.; 2055 Gateway Place, Suite 550, San Jose, California 945110 (US).

(21) International Application Number:  
PCT/US2007/012746

(22) International Filing Date: 29 May 2007 (29.05.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
11/471,466 19 June 2006 (19.06.2006) US(71) Applicant (for all designated States except US):  
DISKEEPER CORPORATION [US/US]; 7590 N. Glenoaks Blvd., Burbank, California 91504 (US).

(72) Inventors; and

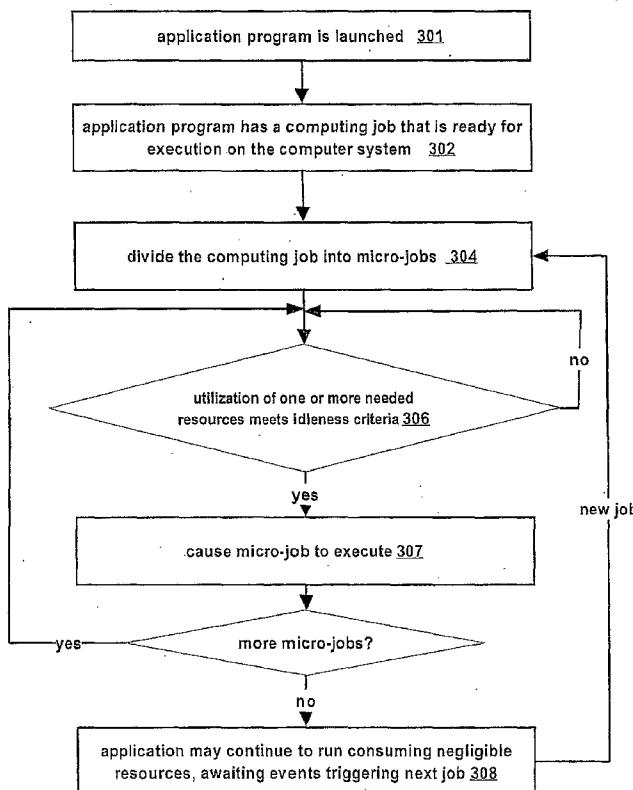
(75) Inventors/Applicants (for US only): JENSEN, Craig [US/US]; 4245 Mesa Vista Drive, La Canada, California 91011 (US). STAFFER, Andrew [CA/US]; 13270 Alta Vista Way, Sylmar, California 91342 (US). THOMAS, Basil [IN/US]; 14537 Willowgreen Lane, Sylmar, California 91342 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: METHOD, SYSTEM, AND APPARATUS FOR SCHEDULING COMPUTER MICRO-JOBS TO EXECUTE AT NON-DISRUPTIVE TIMES



**(57) Abstract:** A method, system, and apparatus to divide a computing job into micro-jobs and allocate the execution of the micro-jobs to times when needed resources comply with one or more idleness criteria is provided. The micro-jobs are executed on an ongoing basis, but only when the resources needed by the micro-jobs are not needed by other jobs. A software program utilizing this methodology may be run at all times while the computer is powered up without impacting the performance of other software programs running on the same computer system.



**Published:**

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

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**METHOD, SYSTEM, AND APPARATUS FOR SCHEDULING COMPUTER MICRO-JOBS  
TO EXECUTE AT NON-DISRUPTIVE TIMES**

**FIELD OF THE INVENTION**

**[0001]** The present invention relates to executing software applications in a computing environment. In particular, embodiments of the present invention relate to dividing a computing or input-output job of an application into micro-jobs and allocating the execution of the micro-jobs to periods of low computer resource utilization, so that execution of the micro-jobs does not significantly impact performance of other applications and jobs.

**BACKGROUND**

**[0002]** In the early days of mainframe computers, the concept of running software programs in batches of jobs was the norm. There were a limited number of computers, so users had to schedule their job(s) to run on the computer when the computer was not being used for some other, more important job. In such systems, each job was scheduled to run to completion without interruption, followed by the next job and then the next. The limited computer time available necessitated running lower-priority jobs “off-hours” so as not to delay higher-priority applications.

**[0003]** More recently, multi-tasking computer systems have allowed the concurrent or interleaved execution of two or more jobs by a single CPU. A multi-tasking computer system allows many applications to execute in the same general time period. Typically, multi-tasking systems have complex internal scheduling algorithms, wherein processes are scheduled in accordance with assigned priorities. However, the applications still contend for computing resources. To alleviate resource contention, an application in a multi-tasking system may be run off “off-hours” on an operator-scheduled basis.

**[0004]** The applications that are run off-hours may include maintenance jobs, such as backup, indexing, software updates, virus and malware scans and defragmentation. Candidates for off-hours processing may also include software applications that run reports, perform financial calculations, etc. However, some applications, such as indexers, should be run during production time. Therefore, not all applications are good candidates for off-hours execution.

**[0005]** Another problem with scheduling a job to run off-hours is that the computer may be turned off at the time the job is scheduled to run. A further problem is that some

machines do not have clearly identified off-hours. For example, many computer systems are used twenty-four hours a day for a computing activity that is considered significant enough that the activity should not be interrupted for a substantial period. Therefore, there are no "off-hours" in which to schedule jobs. A still further problem is that typically a user has to determine when the job should be scheduled for off-hours computing. Thus, setting the schedule takes up a user's time and is subject to user error.

**[0006]** As previously mentioned, running the computing job can interfere with a user's ability to use the computer and can take resources away from other, possibly more pressing applications and jobs. Throttling is a technique for minimizing these negative impacts. Throttling prevents an application or job from using more than an allocated amount of resources. Types of throttling include disk I/O throttling, CPU throttling and network throttling. For example, CPU throttling can involve establishing a target CPU utilization limit for an application and forcing the application to stop working if the application exceeds the target limit. Throttling is sometimes applied to computer resources for maintenance applications or less important computing jobs. While throttling has benefits, the computing job's resource use is not totally transparent to other jobs and applications.

**[0007]** At the same time, it is observable that considerable computing resources go unused, even during the processing of urgent, top-priority jobs. The wide differences in the speeds of CPUs, memory, disk drives and networks typically cause one or more of these components to sit idle while one of the other components is fully consumed. A three-gigahertz CPU, for example, often sits idle while waiting for a disk drive to retrieve data at an average access time measured in milliseconds.

**[0008]** To recover and utilize these otherwise lost resources, what is needed is a technique that allows one or more jobs to execute in a computer system without significantly impacting other jobs or applications. The technique should not consume a user's time in scheduling the job nor should it negatively impact the user's interaction with the computer system when the job is running. The technique should not require scheduling the job to run off-hours. The technique should be utilizable by and beneficial to a computer system that has no off-hours.

**[0009]** The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0010] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0011] FIG. 1 is a diagram of an architecture for executing micro-jobs, in accordance with an embodiment of the present invention.

[0012] FIG. 2 illustrates a comparison of a conventional application memory footprint compared to an application memory footprint in accordance with an embodiment of the present invention.

[0013] FIG. 3 is a flowchart illustrating steps of a process of executing a micro-job scheduler enabled application using micro-jobs, in accordance with an embodiment of the present invention.

[0014] FIG. 4 illustrates steps of a process of defragmenting a storage medium using micro-jobs, in accordance with an embodiment of the present invention.

[0015] FIG. 5 is a block diagram that illustrates a computer system upon which an embodiment of the invention may be implemented.

## DETAILED DESCRIPTION

[0016] In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

## OVERVIEW

[0017] The majority of computers do not utilize all of their resource capacity 100% of the time. This is typically true even of computers that seemingly are in high use twenty-four hours a day, seven days a week; as well as computers that are only turned on for a portion of each day. Therefore, computer time and resources are wasted. For example, over a twenty-four hour period, a computer system that is used quite heavily, and which may have brief spikes in activity, may on average use only about five to twenty percent of its resources.

[0018] A method, system, and apparatus are disclosed herein to utilize these unused computer resources by dividing a computing job into micro-jobs and executing the micro-jobs when the required resources to run the micro-job are idle. As used herein, the term micro-job is a portion of a computer process the execution of which can be delayed without affecting the outcome of the rest of the process. As used herein, the term “idle resource” “idle-time” or the like refers to a time when a resource is less than 100% utilized, that is, a portion of the resource that is not being utilized is considered to be “idle” even though some other portion of the resource is being utilized.

[0019] Thus, the micro-jobs are executed on an ongoing basis, but without an attempt to get the computing job to complete as rapidly as possible. A software program utilizing this methodology may be run at all times while the computer is powered up with negligible impact on the performance of other software programs running at the same time on the same computer, effectively resulting in considerably more useful work getting done per unit of time.

[0020] In contrast to attempting to run the job as quickly as possible given current resource constraints or to scheduling the job to run “off-hours” to avoid significantly impacting other jobs and applications, the job is run on the computer on an ongoing basis, but in such minute pieces that it may be imperceptible to the user or other computer jobs. Thus, the job may be completely transparent to the user and to other jobs and applications. The user does not need to schedule the job; with this method it can be run at any time, including during performance-critical times.

[0021] As used throughout this description, the term “transparent” with respect to other jobs, processes, and applications means that other jobs and applications are able to execute without suffering any significant negative performance impact due to execution of the micro-jobs. A negative impact can be determined by comparing how long a test job takes to run by itself and while a computing job that is divided into micro-jobs is executing. There should be no significant difference in the run time of the test job for these two cases.

[0022] A micro-job scheduler (MJS) determines when the micro-jobs should be executed, in accordance with an embodiment of the present invention. The application program divides its computing jobs into a plurality of micro-jobs, in one embodiment. As used herein, the term plurality means any number greater than one. The MJS determines when the micro-jobs should execute. Applications that are enabled to work with the MJS are referred to herein as MJS-enabled applications.

[0023] In one embodiment, the MJS has a memory manager. An MJS-enabled application requests memory from the MJS memory manager, as opposed to requesting memory from the operating system. In one embodiment, the MJS-enabled application has a very small in-memory footprint (meaning the amount of memory allocated to the MJS-enabled application and from which the MJS-enabled application executes). In order to achieve the small footprint, limits are placed on the amount of memory a program using MJS utilizes.

[0024] In another embodiment, computer resource utilization is monitored and analyzed to determine whether resource utilization complies with one or more idleness criteria. The MJS causes micro-jobs to be executed when the one or more idleness criteria are satisfied. Since the time required for a particular micro-job to execute is smaller than the typical window of idle time for the resource used by the particular micro-job, the resource is relinquished before the resource is needed by another job. Thus, resource utilization by the micro-job may go unnoticed and the micro-job may be invisible to its application's environment.

[0025] An MJS-enabled application sends resource utilization threshold parameters to the MJS in order to define conditions under which its resource utilization should be yielded, in accordance with an embodiment of the present invention. These resources include, but are not limited to, disk I/O, CPU and network utilization. For example the MJS-enabled application can request a micro-job to be executed pending minimum threshold levels for any combination of the above three resources.

[0026] In one embodiment, the application is a storage medium defragmenter. A defragmenter might take 12 minutes of wall clock time to complete a daily scheduled defragmentation job. The MJS-enabled defragmenter divides the defragmentation job into many micro-jobs, which can be executed a few milliseconds at a time, selecting times when the resources needed by the defragmenter are not being used by any other job or application. The MJS-enabled defragmenter can run whenever the computer system is on, dividing the 12-minute job into micro-jobs that execute over a longer period.

[0027] As a particular example, the MJS-enabled defragmenter might execute a first micro-job to determine whether a file is fragmented or not. Execution of the first micro-job consumes only otherwise idle resources. That is, resources whose utilization complies with one or more idleness criteria. As a result, execution of the micro-job is transparent to other applications. Moreover, memory utilization is kept low and memory allocation associated with the defragmenter need not be changed prior to or following execution of the micro-job.

[0028] Prior to executing each micro-job, the MJS-enabled defragmenter makes a determination whether computer resource utilization complies with one or more idleness criteria. Thus, the MJS-enabled defragmenter determined whether computer resource utilization is currently sufficiently low to allow the micro-job to proceed. If resource utilization is too high, execution of the micro-job is deferred. After the first micro-job executes, the MJS-enabled defragmenter does not necessarily immediately execute a second micro-job. Rather, the execution of subsequent micro-jobs may spread out over time if other applications require the same resources.

#### ARCHITECTURE OVERVIEW

[0029] FIG. 1 is a diagram of an architecture 100 for executing micro-jobs, in accordance with an embodiment of the present invention. Each MJS-enabled application 115(1) – 115(n) divides its computing job (or jobs) into micro-jobs 125 to execute. For example, an application programmer can place calls at appropriate locations in the application code that request permission from the MJS 110 to execute a micro-job 125, which, in effect, divides the computing job into micro-jobs 125. The micro-computing jobs are substantially smaller (for example, the smallest) work units that can be completed as a single unit while safely allowing for a pause in execution until the next micro-job 125 executes, in accordance with one embodiment. By safely allowing for a pause in execution, it is meant that the execution of a particular micro-job can be delayed without affecting the outcome which results from execution of the all of the micro-jobs. Keeping the micro-jobs 125 small allows the MJS-enabled application 115 to use only a small amount of computer resources at one time. Thus, execution of a micro-job 125 consumes a sufficiently small amount of resources so as to not significantly impact performance of other applications in the computer system, in accordance with one embodiment of the present invention. As examples, the MJS-enabled applications 115(1) – 115(n) may perform maintenance such as backup, indexing, software updates, virus and malware scans, and defragmentation. However, the MJS-enabled applications 115(1) – 115(n) may also be software other than maintenance.

[0030] The micro-job scheduler (MJS) 110 determines when micro-jobs 125 can be executed. In this embodiment, the MJS 110 has an application program interface (API) 130 to allow a particular MJS-enabled application (e.g., 115(1)) to request that one or more micro-jobs 125 be allowed to execute. The API 130 also allows an MJS-enabled application 115 to specify by how much the micro-jobs 125 may be spread out, as will be

discussed more fully below. An example API is included herein below. However, the architecture 100 is not limited to the example API.

[0031] The micro-job scheduler 110 maintains a micro-job queue, such that the MJS 110 can determine which micro-job 125 should be allowed to execute next. The execution of micro-jobs is timed by the MJS 110 so as to have negligible impact on other applications. In one embodiment, the MJS 110 schedules micro-jobs such that the micro-jobs utilize only idle resources. The MJS 110 determines whether resource utilization complies with one or more idleness criteria to make a scheduling decisions. In other words, the MJS 110 is resource based, in this embodiment. The execution of the micro-jobs can be specified by the MJS-enabled application in an API call or other method of communication between the application and the MJS. When the scheduler 105 determines that the next micro-job 125 may execute without impacting other jobs, the MJS 110 responds to the MJS-enabled application 115 by instructing the MJS-enabled application 115(1) to execute the micro-job 125. The MJS bases its determination on when to schedule jobs based on resource utilization, in one embodiment. As an example, the MJS may analyze disk activity. If an application other than the application with the micro-job is using the disk, then the MJS waits until the other application is done to schedule the micro-job. The MJS continues to monitor the disk I/O utilization, and allows another micro-job to be scheduled if no other application is seeking access to disk I/O. However, if another application seeks utilization of disk I/O, then the MJS will not allow another micro-job to be scheduled, wherein the other application can utilize the disk I/O.

[0032] As another example, the MJS may analyze network activity. If network traffic is too high, the MJS will not schedule any micro-jobs until traffic slows. If network traffic is low enough, then the MJS schedules a micro job for execution. The MJS continues to check to make sure that network traffic stays low enough. If network traffic stays low enough, another micro-job may be scheduled. However, if traffic gets too high, no further micro-jobs are scheduled to execute.

[0033] The MJS may make resource-based scheduling decisions based on any type of computing resource and any combination of resources. In one embodiment, the MJS has multiple queues of micro-jobs awaiting permission to execute. Each queue may correspond to a particular resource. For example, there may be a queue for micro-jobs that need to utilize disk I/O, a queue for micro-jobs that need to utilize a network, a queue for micro-jobs that need to utilize a CPU, etc. There may also be one or more queues for micro-jobs that utilize a combination of resources. The MJS deploys micro-jobs when the particular resource or combination of resources is available. A particular micro-job might

require the use of two resources. For example, the particular micro-job might require use of a network resource and disk resource. However, the particular micro-job does not need CPU resource. Even if the CPU resource utilization is currently high, the particular micro-job can still be scheduled and executed.

[0034] While the MJS 110 is depicted in FIG. 1 as a separate program from the MJS-enabled application 115, the MJS 110 may be integrated into the MJS-enabled application 115. Thus, the MJS 110 can cause the micro-jobs 125 to execute without sending a response to the MJS-enabled application 115(1), in accordance with one embodiment. Therefore, the API 130 is optional.

[0035] In one embodiment, the MJS 110 is part of an operating system. In another embodiment, the MJS 100 executes outside from an operating system. If the MJS executes outside of the operating system, the MJS self-limits in its own resource utilization, in one embodiment. For example, the MJS 110 monitors its own resource utilization and if its own resource utilization gets too high, the MJS 110 makes a request to the operating system to stop scheduling the MJS 110 for a period of time.

[0036] The MJS-enabled application 115 sends parameters to the MJS 110 to control resource utilization, in accordance with an embodiment of the present invention. Control of resource utilization includes, but is not limited to, disk I/O, CPU and network. For example, the MJS-enabled application 115 can request a micro-job be executed pending any combination of threshold levels of the above three resources. Moreover, the MJS-enabled application 115 can specify different resource threshold levels for different micro-jobs 125. For example, the MJS-enabled application 115 specifies different resource threshold level with each micro-job 125, in accordance with one embodiment. Therefore, fine-grained resource management is possible. When the MJS 110 calculates resource utilization, it is the resource utilization of processes other than the MJS-enabled application that is measured, in accordance with one embodiment of the present invention. The following example in which the CPU utilization threshold is set to 20 percent is used to illustrate. If CPU utilization is below 20 percent prior to allowing the MJS-enabled application to execute, CPU utilization may increase to over 20 percent when the micro-job(s) execute. This increase beyond 20 percent is not considered a CPU resource utilization violation, in this example. Similar principles apply to network and disk I/O resources.

[0037] The MJS 110 also has a memory manager 140, in one embodiment. When the MJS 110 initializes, it is allocated memory by the operating system, some of which it uses for its own purposes and some of which it allocates to the MJS-enabled applications 115.

When an MJS-enabled application 115(1) launches, it requests a memory allocation from the MJS 110. The MJS 110 may determine how much memory to allocate to the MJS-enabled application 115(1) based on factors such as current computer system memory utilization by all processes and the needs of the MJS-enabled application 115(1).

Memory requirements may be specific to each MJS-enabled application 115 and may be programmed into the MJS-enabled application 115 by a computer software programmer.

[0038] FIG. 2 illustrates a comparison of a conventional application memory footprint 202 compared to an MJS-enabled application memory footprint 204(1) - 204(n), in accordance with an embodiment of the present invention. Because the MJS-enabled application 115(1) executes micro-jobs 125, the memory allocation 204(1) can be very small. Moreover, because the memory allocation 204(1) is so small, it may not be necessary for the MJS-enabled application 115(1) to ever relinquish its allocated memory 204(1). Thus, the MJS-enabled applications 115 do not cause frequent memory allocations and de-allocations. This reduction in or elimination of memory allocations is one reason why the MJS-enabled applications 115 have no significant impact on other applications and jobs.

#### PROCESS FLOW

[0039] FIG. 3 is a flowchart illustrating steps of a process 300 for executing an MJS-enabled application using micro-jobs, in accordance with an embodiment of the present invention. In step 301, an MJS-enabled application program is launched. The MJS-enabled application program is launched when the computer system is booted, in one embodiment. If the MJS-enabled application does not have a job to perform, the MJS-enabled application sits in an idle state until it has a job to run. In this idle state, the MJS-enabled application may perform some functions such as occasional monitoring. In step 302, the MJS-enabled application has a job to perform, such as defragmenting a storage medium, or scanning for viruses. The job could be to defragment a single disk and the files stored on it, wherein the MJS-enabled application defragments that disk on an ongoing basis.

[0040] A small amount of memory is allocated to the MJS-enabled application when it is launched. The amount that is allocated can be very small because the MJS-enabled application typically only attempts to execute a single micro-job at one time. However, in some cases, the MJS-enabled application might attempt to execute multiple micro-jobs without waiting for other processes to execute. For example, if the MJS determines that the required computer system resources are idle, the MJS may allow the MJS-enabled

application to execute multiple micro-jobs in a row without another process utilizing the resources that are utilized by the micro-jobs.

[0041] In step 304, the computing job is divided into micro-jobs. The micro-jobs are of a size such that execution of the micro-jobs utilizes a sufficiently small amount of resources so as to not significantly impact performance of other jobs in the computer system, in accordance with an embodiment of the present invention. Dividing the computing job into micro-jobs may be accomplished by instructions within an MJS-enabled application. In general, these instructions are decision points in the MJS-enabled application. For example, the instructions can be API calls to the MJS that request permission to execute a micro-job. However, the MJS can be integrated with the MJS-enabled application, in which case the instructions could be calls to a scheduling function within the MJS-enabled application. Other techniques might be used to divide the computing job into micro-jobs.

[0042] In step 306, a determination is made, for each of the micro-jobs, as to whether utilization of one or more resources of the computer system to be used by the particular micro-job satisfies one or more idleness criteria. Thus, the micro-jobs are only executed at times when the resources of the computer system needed by the micro-jobs are sufficiently idle. The idleness criteria is based on resource thresholds, in one embodiment. For example, resource thresholds may be used, wherein a micro-job of the MJS-enabled application only executes if resource utilization by other processes is below a threshold specified by the MJS-enabled application. The example API described below contains one example of some resource threshold parameters. However, process 300 is not limited to these resource threshold parameters. Step 306 spreads the execution of micro-jobs over time subject to availability of the resources needed by each micro-job. Therefore, execution of the micro-jobs does not significantly impact performance of the other jobs and applications in the computer system.

[0043] The idleness criteria for a particular resource may be based on one or more factors. For example, CPU utilization is used as an idleness criteria for CPU resource utilization, in one embodiment.

[0044] In step 307, the micro-jobs are caused to be executed. In one embodiment, the MJS gives permission to the MJS-enabled application that a particular micro job can be executed. If there are more micro-jobs to execute, then control passes to step 306 to determine whether utilization of one or more resources of the computer system to be used by the next micro-job satisfies one or more idleness criteria.

**[0045]** In a maintenance-type embodiment, the MJS-enabled application runs on an on-going basis, continuing to run as long as the computer system remains booted. Thus, even if the MJS-enabled application completes its job, it continues to run, as indicated in step 308. Thus, the MJS-enabled application does not consume additional resources that are typical of launching an application. When the MJS-enabled application determines that it has another job to do, the MJS-enabled application divides the new computing job into micro-jobs in step 304, and the micro-jobs are executed over time in step 306.

#### EXAMPLE API

**[0046]** An embodiment of the present invention is an API for allowing an MJS-enabled application to interface with an MJS. The example API has the following resource threshold parameters for CPU, disk, and network.

- CPU Utilization threshold
- Pending Disk I/O Count threshold
- Network Utilization threshold

**[0047]** The above parameters can be specified for each micro-job. In other words, different micro-jobs can be assigned different resource threshold parameters. For example, for a micro-job that uses the network, a network threshold may be used. However, the network threshold could be zero for micro-jobs that do not use the network. Thus, fine-grained resource management is provided for, in accordance with an embodiment of the present invention.

**[0048]** As a particular example, the MJS-enabled application can request that a particular micro-job be executed only if the CPU utilization is below 50%, and the I/O Disk Utilization is below 40%, and network traffic is below 60%. Any combination of the resource threshold factors can be used, including none at all. The CPU utilization threshold differentiates between MJS's use of the CPU as opposed to that of any other job, in accordance with an embodiment of the present invention.

**[0049]** The following two parameters are used to specify how frequently resource utilization should be measured.

- CPU Utilization Window
- Network Utilization Window

**[0050]** The CPU Utilization Window parameter defines a time window over which CPU utilization is calculated. For example, CPU utilization over the last  $n$  milliseconds is averaged. The network utilization window defines a time window over which network utilization is calculated. These parameters may be internal to the MJS. However, an MJS-enabled application may override these parameters. The pending disk I/O is absolute at any point in time and thus it does not have to be calculated.

**[0051]** A mandatory idle time parameter may be passed from the MJS-enabled application to the MJS engine to control how the micro-jobs are spread out over time. The mandatory idle time parameter is optional. Furthermore, when used, the mandatory idle parameter may have a value of zero.

- Mandatory Idle Time

**[0052]** The MJS keeps track of “idle time,” which is defined as system idle time after all micro-jobs have executed. As previously mentioned, MJS-enabled application(s) can queue up micro-jobs with the MJS. When there are no micro-jobs on the MJS queue, the MJS waits for the specified Mandatory Idle Time and then wakes up and authorizes the MJS-enabled application(s) to perform additional work. For example, an MJS-enabled defragmenter might first execute a number of micro-jobs to defragment a disk drive, then be paused by the MJS micro-job scheduler. After the specified Mandatory Idle Time, the MJS calls the MJS-enabled defragmenter to authorize additional work. For example, the MJS-enabled defragmenter might execute a clean-up job, such as releasing memory. Mandatory Idle Time can be a default parameter that can be adjusted by an MJS-enabled application.

**[0053]** The following parameters relate to waiting to execute a micro-job when resource utilization is above a threshold.

- Wait Time
- Maximum Wait Time

**[0054]** If the MJS determines that resource utilization is currently too high to execute a micro-job, the MJS waits for the specified Wait Time and then re-checks resource utilization. The Wait Time parameter can be increased each time the MJS determines that resource utilization is too high. For example, the MJS can increase the Wait Time parameter until the Max Wait Time is reached. These parameters can be specified by the

MJS-enabled application when it is first started. An MJS-enabled application can adjust these parameters during its run time.

#### DEFRAMGMENTER EMBODIMENT

[0055] The MJS-enabled application is a defragmenter, in accordance with one embodiment. FIG. 4 illustrates steps of a process 400 of an MJS defragmenter using micro-jobs, in accordance with an embodiment of the present invention. Process 400 is an example of one way in which at least a portion of a defragmenter can make use of the micro-job concepts. In this example, the MJS-enabled application interfaces with the MJS via an API. However, as previously mentioned, the MJS may be integrated into the MJS-enabled application, wherein the API is not necessary.

[0056] The process of defragmenting includes scanning a disk drive for fragmented files. The scanning can be broken down into separate micro-jobs of getting a file record and determining whether the file is fragmented. Scanning the disk drive is depicted in steps 402 – 408 of FIG. 4.

[0057] In step 402, the MJS-enabled defragmenter calls the micro-job API to request permission from the MJS to execute a micro-job. The API call can specify resource utilization threshold parameters for this micro-job. Alternatively, previously defined parameters can be applied for this micro-job. In response to receiving the request, the MJS determines when the micro-job can be executed. This determination can be based on the resource utilization threshold parameters.

[0058] In step 404, after the MJS-enabled application receives permission, the MJS defragmenter executes a micro-job, which in this case is to get the next file record. In step 406, the MJS-enabled application again calls the micro-job API. When the MJS-enabled application receives permission to execute from the MJS, the MJS-enabled application determines whether the file from step 404 is fragmented. If the file is not fragmented, process 400 returns to step 402.

[0059] If the file is fragmented, the MJS-enabled defragmenter can defragment the file, as depicted in steps 410 – 416. Step 410 is a call to the micro-job API. Step 412 is to find free disk space for the file and obtain an allocation of the free space.

[0060] Step 414 of calling the micro-job API and step 416 of moving a piece of the file are repeated until the entire file is moved. For example, the file to be moved could be moved in pieces small enough to cause no significant impact on the performance of other applications.

[0061] Thus, the micro-job concept provides for a dynamic defragmenter that monitors fragmentation on an on-going basis and defragments files as soon as fragmentation occurs. This results in an MJS-enabled defragmenter that defragments fragmented files immediately upon the occurrence of fragmentation, without waiting for a scheduled defragmentation run time. Thus, an MJS-enabled defragmenter is a real-time defragmenter, in accordance with an embodiment of the present invention.

#### SHELL MICRO-JOB SCHEDULER

[0062] In one embodiment, the MJS automatically divides a computing job into micro-jobs. For example, the MJS works as a shell wrapped around application programs that are not MJS-enabled. Thus, the shell MJS is a complete software application through which any executable can be run. The shell MJS automatically divides the computing job from the executable into micro-jobs. In other words, the application programmer does not need to divides the application into micro-jobs, in this embodiment.

[0063] The shell MJS divides the computing job from the executable into micro-jobs based on resource utilization, in one embodiment. The shell MJS may analyze the application and how the application runs to see what resources the application uses. For example, the MJS analyzes what resources the application uses and to what degree the application uses the resources. For example, when a disk defragmenter runs, the shell MJS can determine what resources the applications use (e.g., CPU, network, disk I/O). The shell MJS automatically determines how to divide the application in into micro-jobs based on this analysis, in one embodiment. The shell MJS may also determine how to schedule the micro-jobs, based on this analysis.

[0064] The shell MJS may use various parameters to determine how to divide the computing job into micro-jobs and/or determine how to schedule the micro-jobs for execution. These parameters may be based on user input or established by the shell MJS. For example, the user may specify that a particular application has a high priority.

#### HARDWARE OVERVIEW

[0065] FIG. 5 is a block diagram that illustrates a computer system 500 upon which an embodiment of the invention may be implemented. Steps of process 300 and process 400 are stored as instructions one or more of the computer-readable media of system 500 and executed on the processor of computer system 500. Computer system 500 includes a bus 502 or other communication mechanism for communicating information, and a processor 504 coupled with bus 502 for processing information. Computer system 500

also includes a main memory 506, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 502 for storing information and instructions to be executed by processor 504. Main memory 506 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 504. Computer system 500 further includes a read only memory (ROM) 508 or other static storage device coupled to bus 502 for storing static information and instructions for processor 504. A storage device 510, such as a magnetic disk or optical disk, is provided and coupled to bus 502 for storing information and instructions. The computer system 500 can have any number of processors 504. For example, computer system 500 is a multi-processor system, in one embodiment. The processor 504 can have any number of cores. In one embodiment, the processor 504 is a multi-core processor 504. Computer system 500 can be used in a hyper threaded machine.

[0066] Computer system 500 may be coupled via bus 502 to a display 512, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 514, including alphanumeric and other keys, is coupled to bus 502 for communicating information and command selections to processor 504. Another type of user input device is cursor control 516, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 504 and for controlling cursor movement on display 512. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0067] The invention is related to the use of computer system 500 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 500 in response to processor 504 executing one or more sequences of one or more instructions contained in main memory 506. Such instructions may be read into main memory 506 from another machine-readable medium, such as storage device 510. Execution of the sequences of instructions contained in main memory 506 causes processor 504 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0068] The term "machine-readable medium" as used herein refers to any medium that participates in providing data that causes a machine to operate in a specific fashion. In an embodiment implemented using computer system 500, various machine-readable media are involved, for example, in providing instructions to processor 504 for execution.

Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 510. Volatile media includes dynamic memory, such as main memory 506. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 502. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infrared data communications. All such media must be tangible to enable the instructions carried by the media to be detected by a physical mechanism that reads the instructions into a machine.

[0069] Common forms of machine-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0070] Various forms of machine-readable media may be involved in carrying one or more sequences of one or more instructions to processor 504 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 500 can receive the data on the telephone line and use an infrared transmitter to convert the data to an infrared signal. An infrared detector can receive the data carried in the infrared signal and appropriate circuitry can place the data on bus 502. Bus 502 carries the data to main memory 506, from which processor 504 retrieves and executes the instructions. The instructions received by main memory 506 may optionally be stored on storage device 510 either before or after execution by processor 504.

[0071] Computer system 500 also includes a communication interface 518 coupled to bus 502. Communication interface 518 provides a two-way data communication coupling to a network link 520 that is connected to a local network 522. For example, communication interface 518 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 518 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication

interface 518 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0072] Network link 520 typically provides data communication through one or more networks to other data devices. For example, network link 520 may provide a connection through local network 522 to a host computer 524 or to data equipment operated by an Internet Service Provider (ISP) 526. ISP 526 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the “Internet” 528. Local network 522 and Internet 528 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 520 and through communication interface 518, which carry the digital data to and from computer system 500, are exemplary forms of carrier waves transporting the information.

[0073] Computer system 500 can send messages and receive data, including program code, through the network(s), network link 520 and communication interface 518. In the Internet example, a server 530 might transmit a requested code for an application program through Internet 528, ISP 526, local network 522 and communication interface 518.

[0074] The received code may be executed by processor 504 as it is received, and/or stored in storage device 510, or other non-volatile storage for later execution. In this manner, computer system 500 may obtain application code in the form of a carrier wave.

[0075] In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this application, in the specific form in which such claims issue, including any subsequent correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

---

## CLAIMS

What is claimed is:

1. A machine-implemented method comprising the steps:  
for each of a plurality of micro-jobs;  
determining when utilization of one or more resources required to run a particular micro-job complies with one or more idleness criteria;  
and  
causing the particular micro-job to be executed based, at least in part, on a determination that the utilization of one or more resources required to run the particular micro-job complies with one or more idleness criteria.
2. The method of Claim 1, further comprising dividing a computing job into the plurality of micro-jobs.
3. The method of Claim 2, wherein dividing a computing job into the plurality of micro-jobs is performed by an application program that initiates the computing job.
4. The method of Claim 2, wherein dividing a computing job into a plurality of micro-jobs is performed by a software process other than an application program that initiates the computing job.
5. The method of Claim 1, wherein determining when utilization of one or more resources required to run a particular micro-job complies with one or more idleness criteria is performed by an operating system.
6. The method of Claim 1, wherein determining when utilization of one or more resources required to run a particular micro-job complies with one or more idleness criteria is performed by a micro-job scheduler outside of an operating system.
7. The method of Claim 1, further comprising deferring execution of at least one of the micro-jobs based on a wait time that specifies a minimum amount of time to wait between execution of each micro-job.

8. The method of Claim 7, further comprising an application program specifying the wait time.
9. The method of Claim 8, wherein the specifying the wait time comprises specifying a different minimum amount of time for different ones of the micro-jobs.
10. The method of Claim 1, wherein the one or more idleness criteria is based on one or more resource utilization thresholds.
11. The method of Claim 10, further comprising specifying the one or more resource utilization thresholds.
12. The method of Claim 10, further comprising specifying a different resource utilization threshold for different ones of the micro-jobs.
13. The method of Claim 1, wherein causing the particular micro-job to be executed based, at least in part, on a determination that the utilization of one or more resources required to run the particular micro-job complies with one or more idleness criteria is further based on limiting the number of micro-jobs allowed to execute before allowing other jobs to utilize computer resources.
14. The method of Claim 13, wherein limiting the number of micro-jobs allowed to execute is determined by examining computer resource utilization by processes other than the application program.
15. The method of Claim 1, further comprising determining an amount of memory to allocate to an application that initiates the computing job, wherein determining an amount of memory is based on a much smaller amount of memory required for executing micro-jobs of the application rather than on a memory requirement for the application as a whole.
16. The method of Claim 15, wherein determining an amount of memory is further based on the amount of memory available in the computer system.

17. The method of Claim 1, wherein the micro-jobs are portion of a computer process, and wherein the execution of each micro-job can be delayed without affecting the outcome of the rest of the computer process .

18. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, cause the one or more processors to carry out the steps of:

for each of a plurality of micro-jobs;

    determining when utilization of one or more resources required to run a particular micro-job complies with one or more idleness criteria;  
    and

    causing the particular micro-job to be executed based, at least in part, on a determination that the utilization of one or more resources required to run the particular micro-job complies with one or more idleness criteria.

19. A computer-readable medium as recited in Claim 18, further comprising instructions which, when executed by the one or more processors, cause the one or more processors to perform dividing a computing job into the plurality of micro-jobs.

20. A computer-readable medium as recited in Claim 18, further comprising instructions which, when executed by the one or more processors, cause the one or more processors to perform deferring execution of at least one of the micro-jobs based on a wait time that specifies a minimum amount of time to wait between execution of each micro-job.

21. A computer-readable medium as recited in Claim 20, wherein the instructions for deferring execution of at least one of the micro-jobs based on a wait time that specifies a minimum amount of time to wait between execution of each micro-job comprise instructions for performing specifying a different minimum amount of time for different ones of the micro-jobs.

22. A computer-readable medium as recited in Claim 18, wherein the instructions for performing determining when utilization of one or more resources required to run a particular micro-job complies with one or more idleness criteria comprise instructions for performing determining when utilization of one or more resources required to run a particular micro-job complies with one or more idleness criteria based on one or more resource utilization thresholds.

23. A computer-readable medium as recited in Claim 22, further comprising instructions which, when executed by the one or more processors, cause the one or more processors to perform specifying the one or more resource utilization thresholds.

24. A computer-readable medium as recited in Claim 22, further comprising instructions which, when executed by the one or more processors, cause the one or more processors to perform specifying a different resource utilization threshold for different ones of the micro-jobs.

25. A computer-readable medium as recited in Claim 18, wherein the instructions for performing causing the particular micro-job to be executed based, at least in part, on a determination that the utilization of one or more resources required to run the particular micro-job complies with one or more idleness criteria comprise instructions for performing causing the particular micro-job to be executed further based on limiting the number of micro-jobs allowed to execute before allowing other jobs to utilize computer resources.

26. A computer-readable medium as recited in Claim 25, further comprising instructions which, when executed by the one or more processors, cause the one or more processors to perform limiting the number of micro-jobs allowed to execute before allowing other jobs to utilize computer resources by examining computer resource utilization by processes other than the application.

27. A computer-readable medium as recited in Claim 18, further comprising instructions which, when executed by the one or more processors, cause the one or more processors to perform determining an amount of memory to allocate to the application, wherein determining an amount of memory is based on a minimum memory requirement for executing micro-jobs of the application.

28. A computer-readable medium as recited in Claim 27, wherein the instructions for performing the determining of an amount of memory further comprise instructions for performing the determining of an amount of memory based on the amount of memory available in the computer system.

29. A computer-readable medium as recited in Claim 18, wherein the micro-jobs are portion of a computer process, and wherein the execution of each micro-job can be delayed without affecting the outcome of the rest of the computer process .

30. A system, comprising:  
one or more processors; and;  
computer-readable medium communicatively coupled to the one or more processors; wherein the computer-readable medium has stored thereon one or more stored sequences of instructions which, when executed by the one or more processors, cause the one or more processors to perform:  
for each of a plurality of micro-jobs;  
determining when utilization of one or more resources required to run a particular micro-job complies with one or more idleness criteria;  
and  
causing the particular micro-job to be executed based, at least in part, on a determination that the utilization of one or more resources required to run the particular micro-job complies with one or more idleness criteria.

31. A system as recited in Claim 30, wherein the computer readable medium further comprises instructions which, when executed by the one or more processors, cause the one or more processors to perform dividing a computing job into the plurality of micro-jobs.

32. A system as recited in Claim 30, wherein the computer readable medium further comprises instructions which, when executed by the one or more processors, cause the one or more processors to perform deferring execution of at least one of the micro-jobs

based on a wait time that specifies a minimum amount of time to wait between execution of each micro-job.

33. A system as recited in Claim 32, wherein the instructions for deferring execution of at least one of the micro-jobs based on a wait time that specifies a minimum amount of time to wait between execution of each micro-job comprise instructions for performing specifying a different minimum amount of time for different ones of the micro-jobs.

34. A system as recited in Claim 30, wherein the instructions for performing determining when utilization of one or more resources required to run a particular micro-job complies with one or more idleness criteria comprise instructions for performing determining when utilization of one or more resources required to run a particular micro-job complies with one or more idleness criteria based on one or more resource utilization thresholds.

35. A system as recited in Claim 34, wherein the computer readable medium further comprises instructions which, when executed by the one or more processors, cause the one or more processors to perform specifying the one or more resource utilization thresholds.

36. A system as recited in Claim 34, wherein the computer readable medium further comprises instructions which, when executed by the one or more processors, cause the one or more processors to perform specifying a different resource utilization threshold for different ones of the micro-jobs.

37. A system as recited in Claim 30, wherein the instructions for performing causing the particular micro-job to be executed based, at least in part, on a determination that the utilization of one or more resources required to run the particular micro-job complies with one or more idleness criteria comprise instructions for performing causing the particular micro-job to be executed further based on limiting the number of micro-jobs allowed to execute before allowing other jobs to utilize computer resources.

38. A system as recited in Claim 37, wherein the computer readable medium further comprises instructions which, when executed by the one or more processors, cause the

one or more processors to perform limiting the number of micro-jobs allowed to execute before allowing other jobs to utilize computer resources by examining computer resource utilization by processes other than the application.

39. A system as recited in Claim 30, wherein the computer readable medium further comprises instructions which, when executed by the one or more processors, cause the one or more processors to perform determining an amount of memory to allocate to the application, wherein determining an amount of memory is based on a minimum memory requirement for executing micro-jobs of the application.

40. A system as recited in Claim 39, wherein the instructions for performing the determining of an amount of memory further comprise instructions for performing the determining of an amount of memory based on the amount of memory available in the computer system.

41. A system as recited in Claim 30, wherein the micro-jobs are portion of a computer process, and wherein the execution of each micro-job can be delayed without affecting the outcome of the rest of the computer process .

1/5

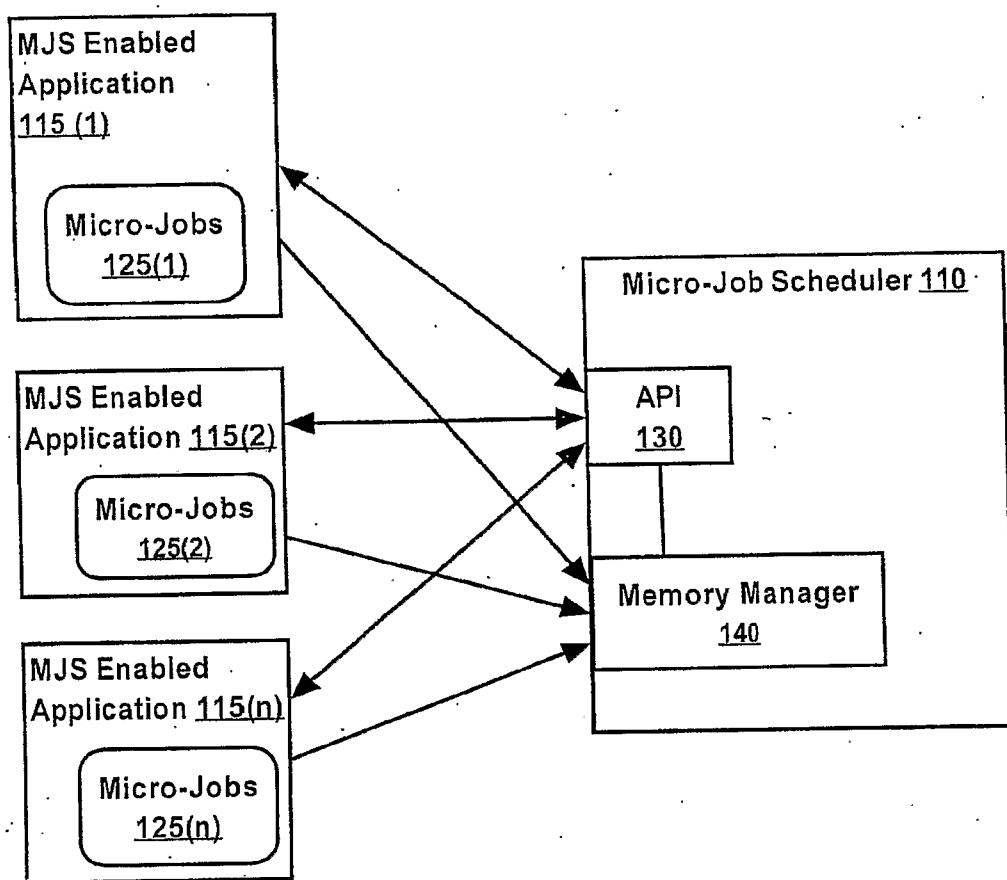
100

FIG. 1

2/5

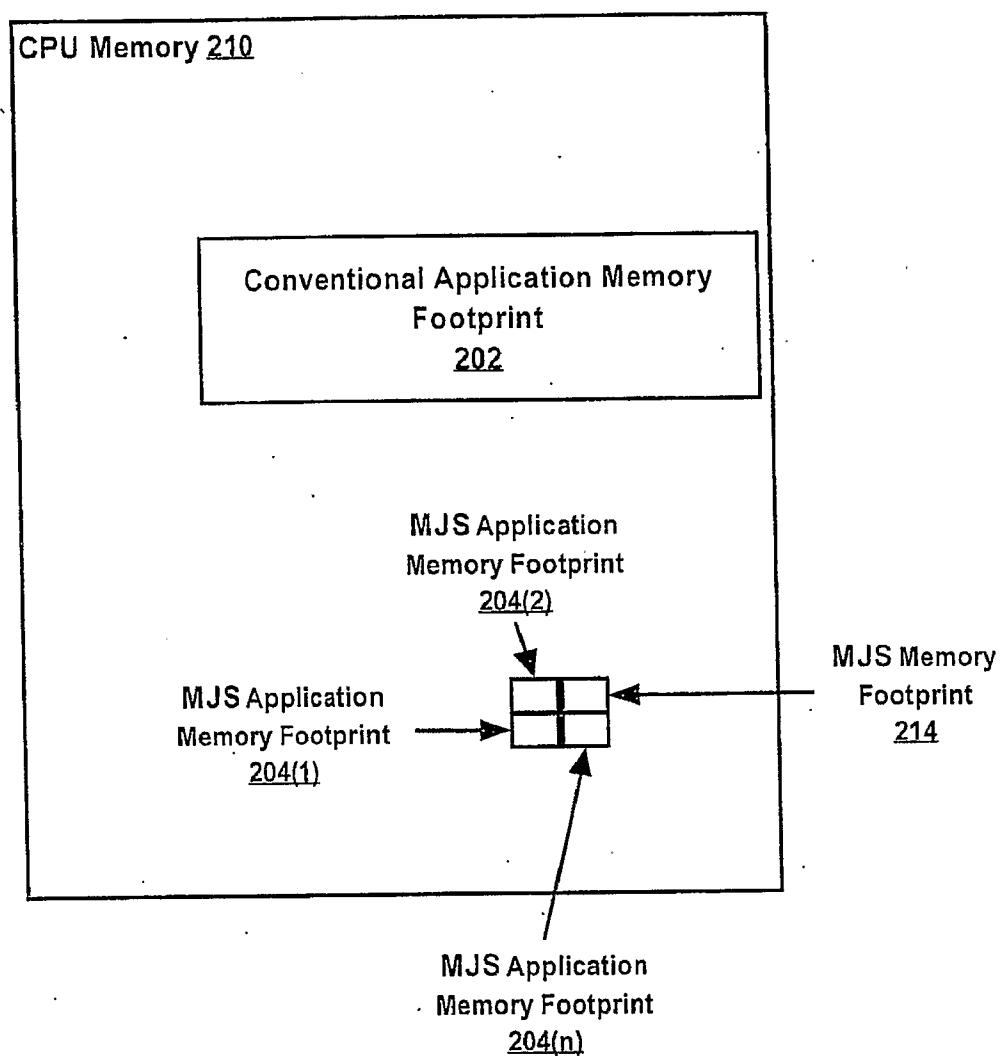
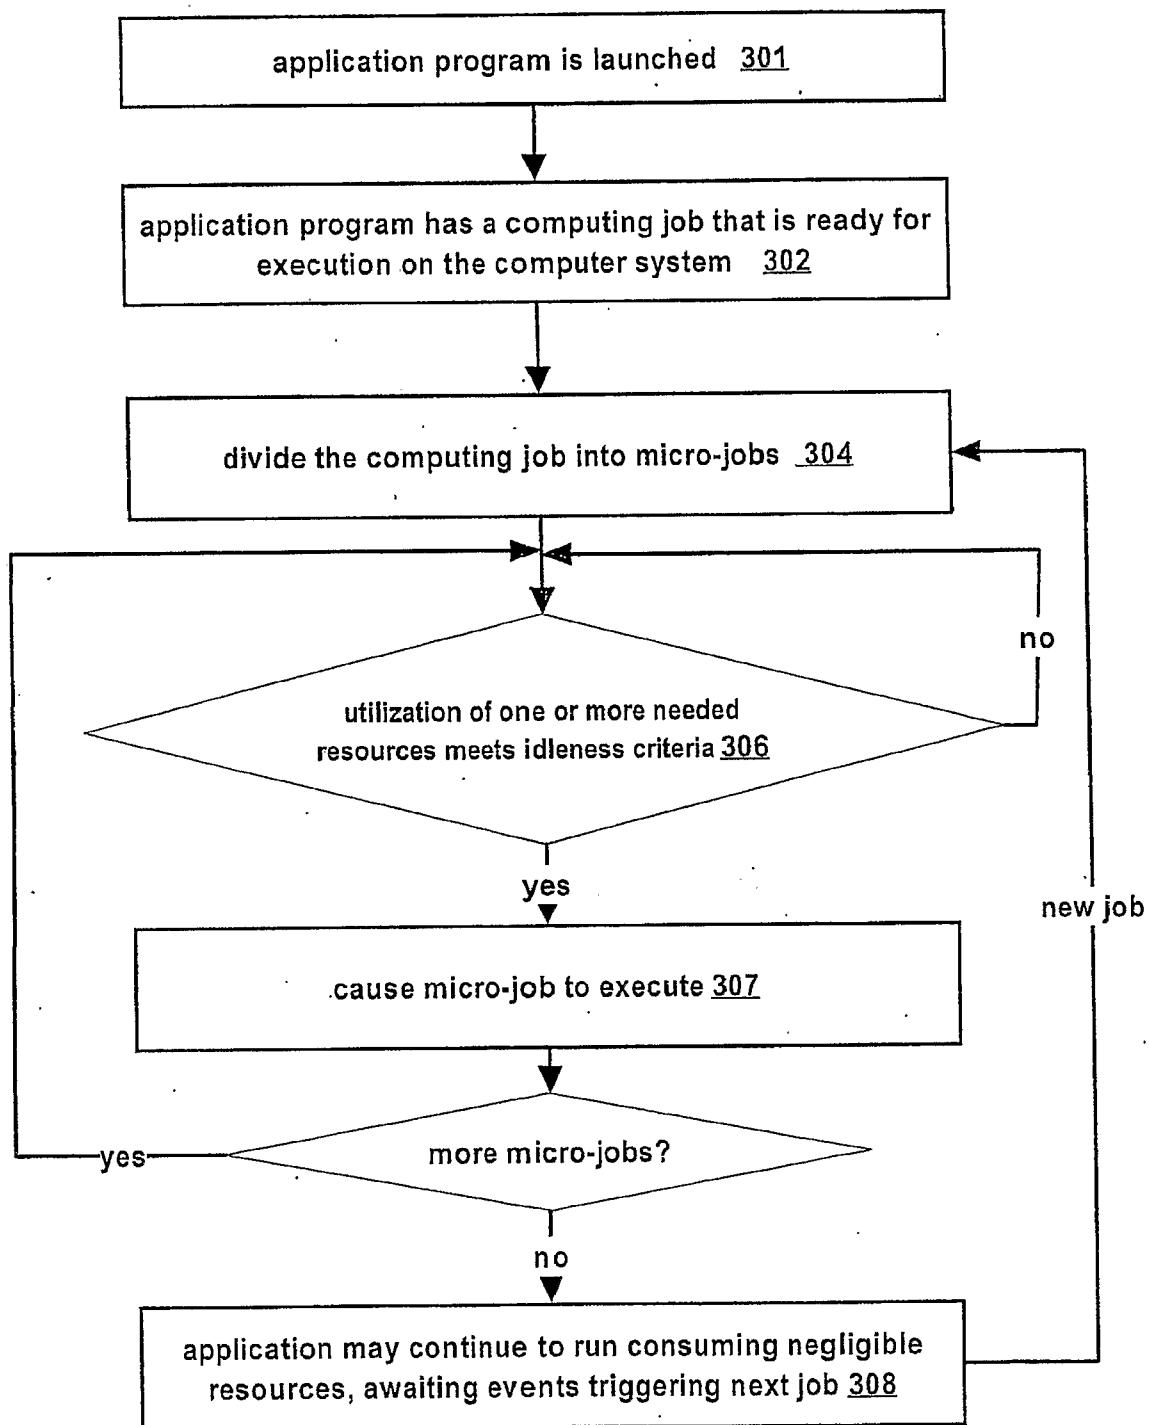


FIG. 2

3/5

300*FIG. 3*

4/5

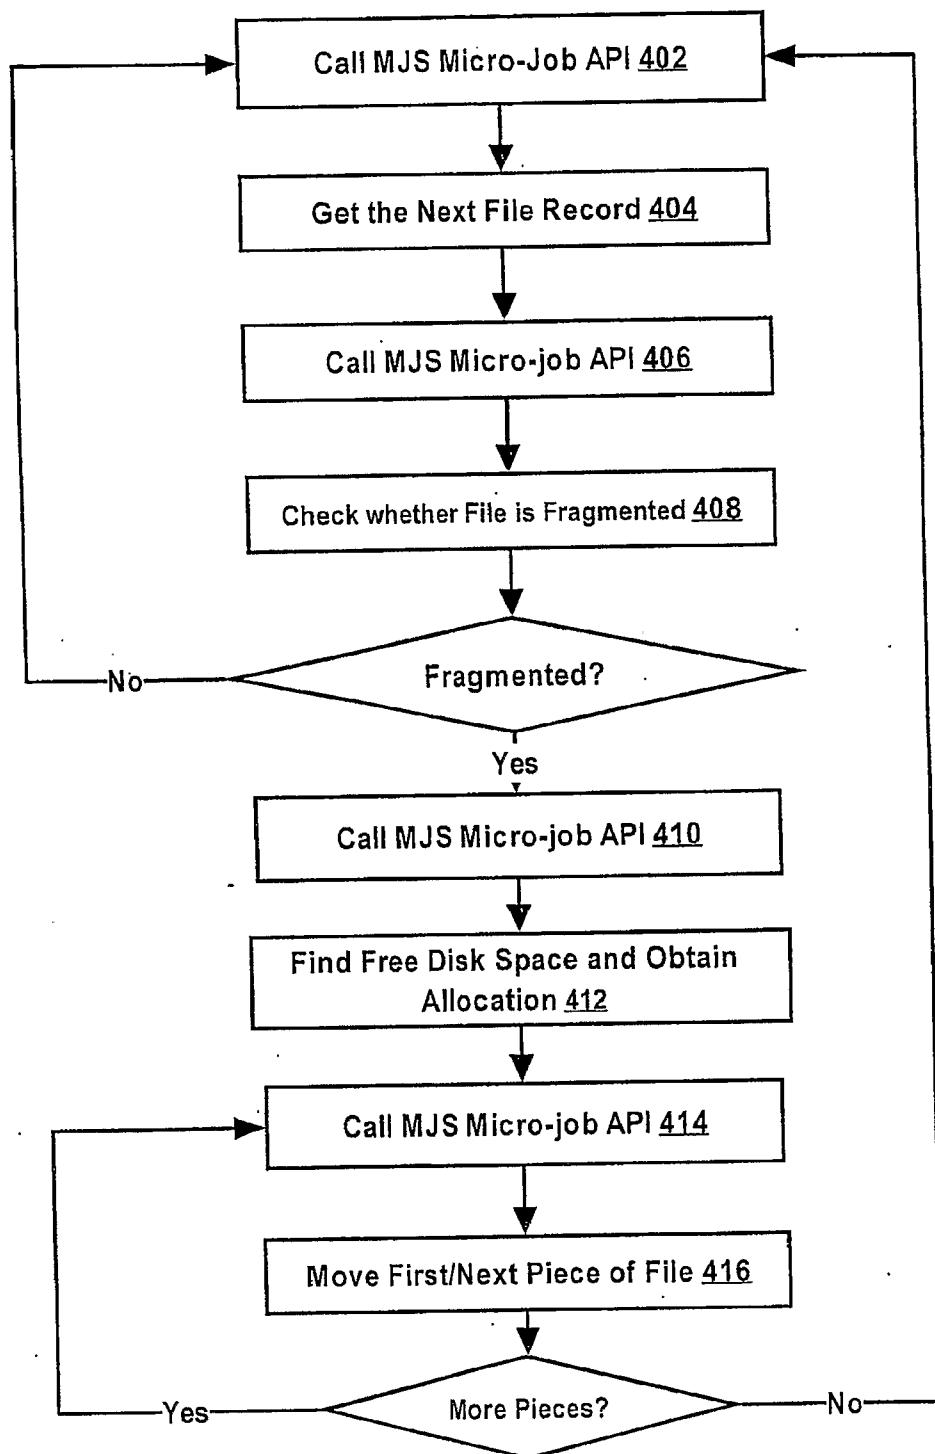
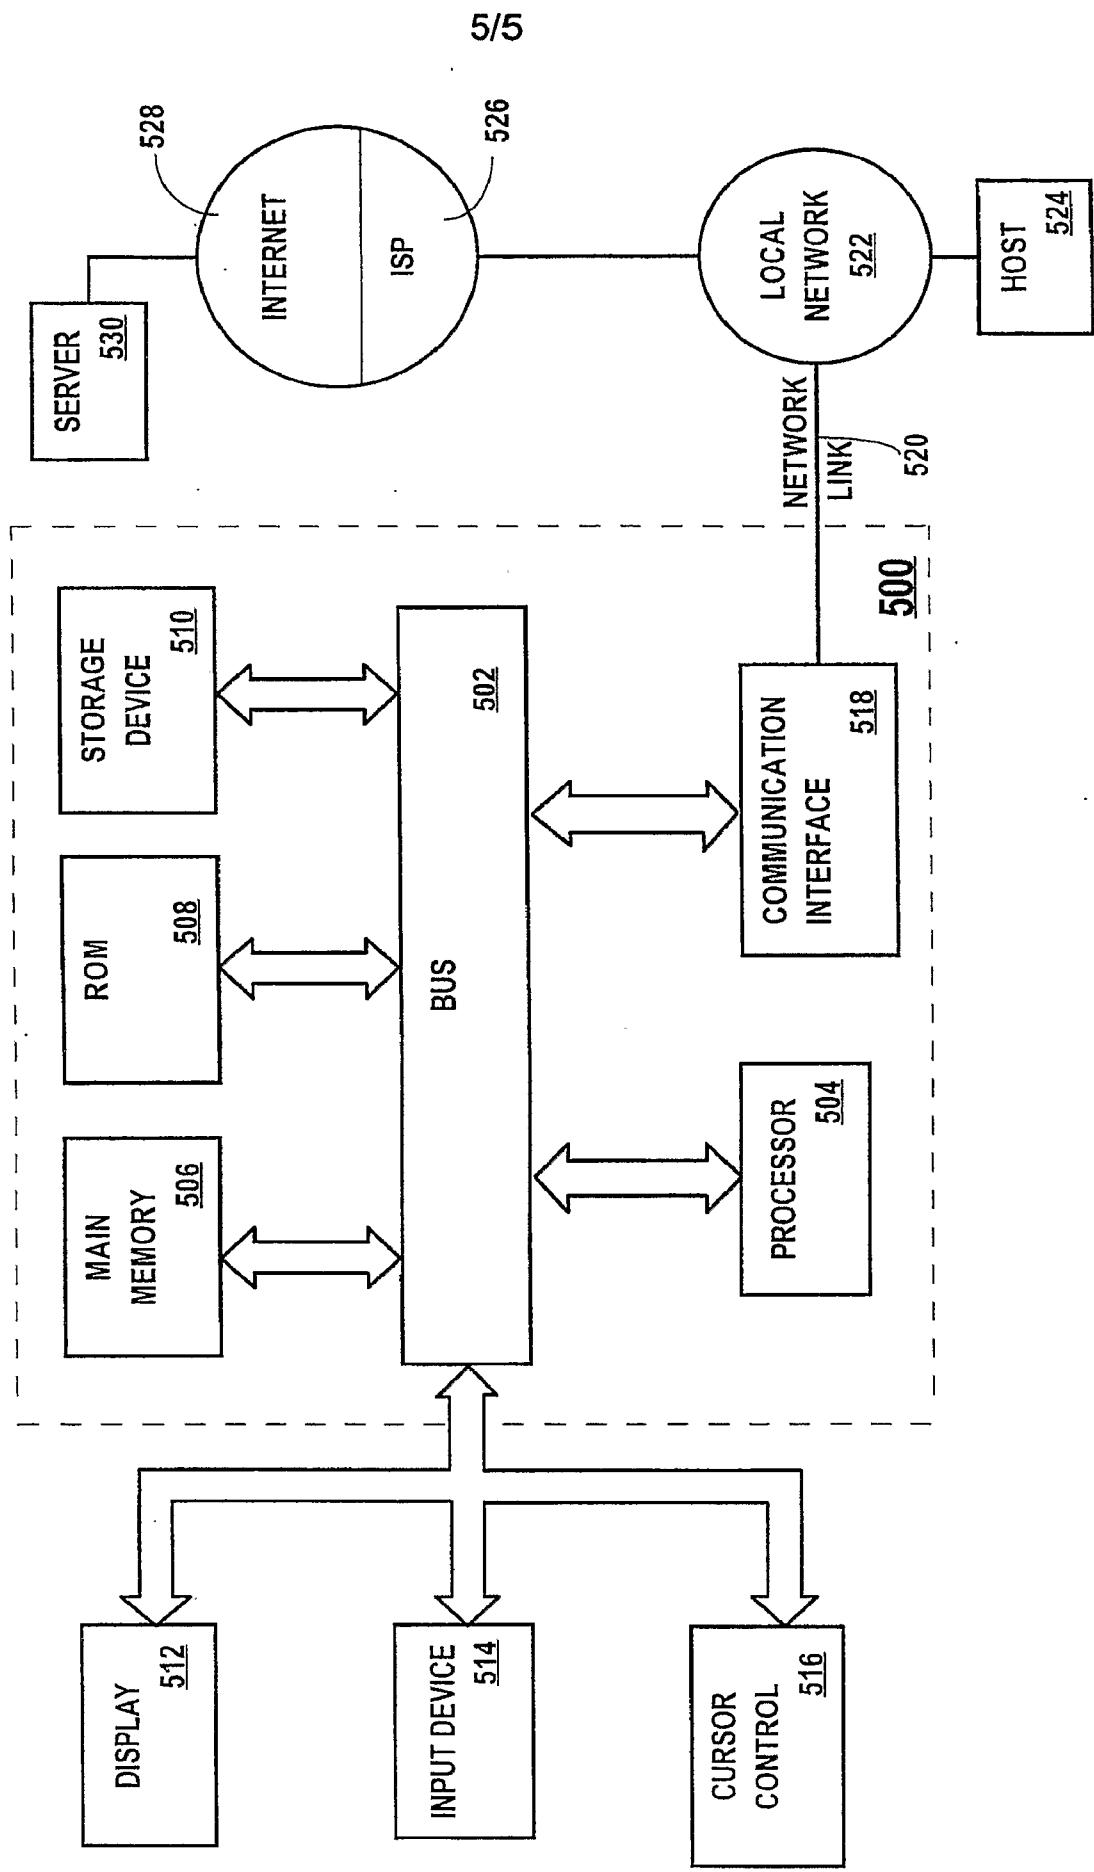
400

FIG. 4

FIG. 5



## INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2007/012746

A. CLASSIFICATION OF SUBJECT MATTER  
INV. G06F9/50

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

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
G06F

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

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

EPO-Internal, INSPEC

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	LIFKA D A: "The ANL/IBM SP scheduling system" JOB SCHEDULING STRATEGIES FOR PARALLEL PROCESSING. IPPS'95 WORKSHOP. PROCEEDINGS SPRINGER-VERLAG BERLIN, GERMANY, 1995, pages 295-303, XP002456555 ISBN: 3-540-60153-8 the whole document	1,18,30
Y		2-17, 19-29, 31-41
	----- -/-	

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 document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

26 October 2007

Date of mailing of the international search report

21/11/2007

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

Skomorowski, Markus

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2007/012746

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	<p>KUHN N ET AL: "Task decomposition in dynamic agent societies" PROCEEDINGS ISADS 93. INTERNATIONAL SYMPOSIUM ON AUTONOMOUS DECENTRALIZED SYSTEMS (CAT. NO.93TH0502-5) IEEE COMPUT. SOC. PRESS LOS ALAMITOS, CA, USA, 1993, pages 165-171, XP002456556 ISBN: 0-8186-3125-2</p> <p>abstract</p> <p>page 165, left-hand column, paragraph 1. -- right-hand column, paragraph 2.</p> <p>page 168, left-hand column, line 1 -- page 170, right-hand column, paragraph 7.</p> <p>-----</p>	2-17, 19-29, 31-41
Y	<p>WO 00/70455 A (HONEYWELL INC [US]; BINNS PAMELA A [US]; VESTAL STEPHEN C [US]) 23 November 2000 (2000-11-23)</p> <p>abstract</p> <p>page 3, line 5 -- page 4, line 16</p> <p>page 7, line 8 -- line 25</p> <p>-----</p>	2-17, 19-29, 31-41
A	<p>DAS R ET AL: "Towards Commercialization of Utility-based Resource Allocation" AUTONOMIC COMPUTING, 2006. ICAC '06. IEEE INTERNATIONAL CONFERENCE ON DUBLIN, IRELAND 13-16 JUNE 2006, PISCATAWAY, NJ, USA, IEEE, 13 June 2006 (2006-06-13), pages 287-290, XP010932299</p> <p>ISBN: 1-4244-0175-5</p> <p>abstract</p> <p>page 287, right-hand column, line 30 --</p> <p>page 288, left-hand column, line 30</p> <p>page 289, left-hand column, paragraph D. --</p> <p>page 290, right-hand column, paragraph V.</p> <p>-----</p>	

**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No

PCT/US2007/012746

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
WO 0070455	A 23-11-2000	AT	253751 T	15-11-2003
		AU	769245 B2	22-01-2004
		AU	4851700 A	05-12-2000
		CA	2371340 A1	23-11-2000
		DE	60006422 D1	11-12-2003
		DE	60006422 T2	09-09-2004
		EP	1244963 A2	02-10-2002
		JP	2002544621 T	24-12-2002
		US	6567840 B1	20-05-2003