

(51) International Patent Classification:
G06F 9/50 (2006.01)

(21) International Application Number:

PCT/US2012/022571

(22) International Filing Date:

25 January 2012 (25.01.2012)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/435,969 25 January 2011 (25.01.2011) US

(71) Applicant (for all designated States except US):
GOOGLE INC. [US/US]; 1600 Amphitheatre Parkway,
Mountain View, CA 94043 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **BORNSTEIN, Daniel, R.** [US/US]; 65 Norfolk Street, Unit 5, San Francisco, CA 94013 (US). **SHAPIRO, Carl, S.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **BROWN, Jeffrey, M.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **TATE, Christopher, Lynn** [US/US]; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US).(74) Agent: **LOOS, Thomas, J.**; McDonnell Boehnen Hulbert & Berghoff LLP, 300 South Wacker Drive, Suite 3200, Chicago, IL 60606 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

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

(54) Title: COMPUTING PLATFORM WITH RESOURCE CONSTRAINT NEGOTIATION




FIG. 3

(57) Abstract: Various techniques are described for resource management on a computing platform. A computing platform can receive a query message that specifies an amount of a resource proposed for allocation. The computing platform can select a selected recommendation level from a plurality of recommendation levels, based on an evaluation of a request for the amount of the resource proposed for allocation. The computing platform can generate a resource allocation recommendation that includes the selected recommendation level with respect to the amount of the resource proposed for allocation. The computing platform can send the resource allocation recommendation.

COMPUTING PLATFORM WITH RESOURCE CONSTRAINT NEGOTIATION

CROSS REFERENCE TO RELATED APPLICATIONS

- 5 [0001] This application claims priority to U.S. Pat. App. No. 61/435,969, entitled “Computer Platform with Resource Constraint Negotiations”, filed January 25, 2011, which is incorporated by reference herein for all purposes.

BACKGROUND

- 10 [0002] Computing platforms, such as personal computers, tablet computers, laptops, cellular telephones, and other mobile or portable devices can include a system monitor to monitor and manage their resources and performance. For instance, the system monitor can monitor and manage resources within the computing platform, such as memory, that can be requested by and allocated to various software processes running on the computing platform. A process
15 can request additional resources from the computing platform, and the system monitor can grant the request and allocate the requested resources or, in some instances, terminate the process based on the effect granting the request would have on overall system resources or performance.

SUMMARY

- 20 [0003] In one aspect, an article of manufacture is provided. The article of manufacture includes a non-transitory computer-readable storage medium with instructions stored thereon that, in response to execution by a processor of a computing platform, cause the computing platform to perform operations. The instructions include: (a) instructions for receiving a
25 query message, where the query message specifies an amount of a resource proposed for allocation, (b) instructions for selecting a selected recommendation level from a plurality of recommendation levels based on an evaluation of a request for the amount of the resource proposed for allocation, (c) instructions for generating a resource allocation recommendation including the selected recommendation level, and (d) instructions for sending the resource
30 allocation recommendation.

[0004] In another aspect, a method is provided. A computing platform receives a query message that specifies an amount of a resource proposed for allocation from the computing

platform. The computing platform selects a selected recommendation level from a plurality of recommendation levels based on an evaluation of a request for the amount of the resource proposed for allocation. Using the computing platform, a resource allocation recommendation including the selected recommendation level is generated. Using the
5 computing platform, the resource allocation recommendation is sent.

[0005] In another aspect, an article of manufacture is provided. The article of manufacture includes a non-transitory computer-readable storage medium with instructions stored thereon that, in response to execution by a processor of a computing platform, cause the computing platform to perform operations. The instructions include: (a) instructions for sending a query
10 message that specifies an amount of a resource proposed for allocation, (b) instructions for, in response to the query message, receiving a resource allocation recommendation related to the amount of the resource proposed for allocation, the resource allocation recommendation comprising a selected recommendation level of a plurality of recommendation levels, (c)
instructions for generating a response to the resource allocation recommendation, where the
15 response is based on the selected recommendation level, and (d) instructions for sending the response to the resource allocation recommendation.

[0006] In another aspect, a method is provided. A computing platform sends a query message that specifies an amount of a resource proposed for allocation from the computing platform. In response to the query message, the computing platform receives a resource
20 allocation recommendation related to the amount of the resource proposed for allocation at the computing platform. The resource allocation recommendation includes a selected recommendation level of a plurality of recommendation levels. Using the computing platform, a response to the resource allocation recommendation based on the selected recommendation level is generated. The computing platform sends the response to the
25 resource allocation recommendation.

BRIEF DESCRIPTION OF DRAWINGS

[0007] FIG. 1 is a block diagram illustrating an example computing platform that can be used to establish a channel of communication between a system monitor and one or more software
30 processes in accordance with one or more aspects of this disclosure.

[0008] FIG. 2 is a block diagram illustrating a system monitor and various processes in accordance with one or more aspects of this disclosure.

[0009] FIG. 3 is an example interaction diagram for an interaction between a computing platform and a process, in accordance with one or more aspects of this disclosure.

[0010] FIGS. 4A and 4B depict an example interaction diagram for an interaction between a computing platform and three processes, in accordance with one or more aspects of this disclosure.

5 [0011] FIG. 5 is a flow diagram illustrating one example method, in accordance with one or more aspects of this disclosure.

[0012] FIG. 6 is a flow diagram illustrating another example method, in accordance with one or more aspects of this disclosure.

DETAILED DESCRIPTION

10 Overview

[0013] In a modern computing platform, an operating system is software that performs various tasks to govern the computing platform. These tasks can be related to execution of application software, management of resources, such as memory, processor(s), and input/output devices associated with the computing platform.

15 [0014] The operating system can include a “system monitor” to manage the resources between the competing requirements of one or more processes or executing computer programs. In some cases, a process can use sub-processes, such as threads, to carry out sub-tasks of the process. The system monitor can allocate resources to the processes to permit successful execution of the processes and/or sub-processes. These resources can include
20 computer memory, computer processor usage, battery power, network bandwidth, access to data provided to and/or received from input/output devices, and other computing-platform-related resources.

[0015] In some cases, allocation of various data items in the operating system can be controlled by the system monitor as well. For example, the operating system can be
25 configured with a fixed number of control blocks required to start processes, and consequently have a fixed number of processes that can be started at one time. In this example, the system monitor can manage the control block resources by requiring a process to specifically request and receive a control block before starting the process. If a control-block request is denied, then the process is unable to execute due to the lack of a control
30 block.

[0016] The system monitor can be configured to provide recommendations and other information about resource requests. For example, suppose a process requires an additional 10 megabytes (MBs) of memory to continue execution. A channel of communication can be

established between the system monitor and the process to enable the process to query if a given request for a particular resource – in this case, 10 MB of memory – would be acceptable. In response to the query, the system monitor can provide one of a plurality of recommendation levels as part of a resource allocation recommendation. Example
5 recommendation levels can range from “approved” to “strongly recommended against.”

[0017] The system monitor can use resource allocation recommendations both to inform processes about resource constraints and indicate a severity level of the resource constraints. That is, a process can explicitly ask the system monitor how good or bad an idea it would be — from the system monitor’s perspective — for the process to receive the requested allocated
10 resources.

[0018] For example, suppose a process is requesting memory in a memory-constrained environment, *e.g.*, one with a relatively large number of simultaneously active processes. In this memory-constrained environment, following the resource allocation recommendations provided by the system monitor can lead the process to be relatively conservative in
15 requesting memory and/or be relatively aggressive in reclaiming memory. In a less memory-constrained environment; *e.g.*, one with a relatively small number of simultaneously active processes, the resource allocation recommendations can inform the process that the computing platform can support requests for additional memory.

[0019] Processes can be expected to follow the resource allocation recommendations of the system monitor by requesting resources when “approved” and to consider delaying
20 requesting resources and/or reclaiming memory when “strongly recommended against.” By providing and following resource allocation recommendations, the computing platform can enable the system monitor and processes to negotiate resource constraints and better allocate resources. Further, the processes operating on the computing platform can be better informed
25 about the status of resources on the computing platform when making resource allocation requests.

An Example Computing Platform

[0020] FIG. 1 is a block diagram illustrating an example computing platform 100 that can be used to establish a channel of communication between a system monitor and one or more
30 software processes in accordance with one or more aspects of this disclosure. Computing platform 100 can be a personal computer, tablet computer, laptop, cellular telephone, personal digital assistant and/or some other mobile/portable device.

[0021] As shown in Figure 1, computing platform 100 includes control unit 120. Control unit 120 can include hardware, software, and/or firmware. In some examples, hardware, such

as one or more processors, can execute software instructions (usually in the form of a computer or software program) to implement one or more aspects of this disclosure. The software instructions can be stored on a non-transitory computer-readable data storage medium. Example non-transitory computer-readable data storage media can include volatile
5 memory, such as random access memory (RAM), including dynamic RAM (DRAM), static RAM (SRAM), and other volatile memories, and/or non-volatile memories, such as read-only memory (ROM), flash memory, magnetic memories such as hard and floppy disks, compact discs, electronic programmable ROM (EEROM), or any other type of non-volatile memory.

[0022] Non-transitory computer-readable data storage media can also include, as another
10 example, storage devices such as compact disk ROMs (CD-ROMs), digital video disk ROMs (DVDs), hard drives, so-called “thumb drives,” and/or any other type of storage device. In some examples, control unit 120 can represent hardware, which, for example, can include one or more digital signal processors (DSPs), general purpose microprocessors, application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), or other
15 equivalent integrated or discrete logic circuitry.

[0023] Control unit 120 executes or implements an operating system 140 that includes a system monitor 160. System monitor 160 is configured to monitor and manage at least resource 180. Resource 180 can include, but is not limited to, physical RAM residing within
20 computing platform 100, usage of part or all of control unit 120, access/control of data between control unit 120 and one or more of devices 170A-170N, and/or resources within operating system 140, such as new threads or processes.

[0024] In addition, system monitor 160 monitors all processes that are active on computing platform 100, noting a process type and/or a process status of each monitor process. A process type can be used to classify a kind of a process, such as a normal process, game, core
25 service, user-visible, high-resource, and background service, among other process. A process status can indicate an execution state of the process, such as running in the foreground, running in the background, blocked, idle and terminated. Other process types and statuses are possible as well.

[0025] System monitor 160 can use the process type and/or process status of a process in
30 determining whether a process should be granted more resources or terminated, for example. One or more applications and services can execute on computing platform 100 and, in particular, on control unit 120. These applications and services, referred to collectively throughout this disclosure as “processes,” are depicted in FIG. 1 as software processes 200A-200N.

[0026] Devices 170A-170N permit interaction with computing platform 100 by enabling input into and/or output from computing platform 100. Example devices 170A-170N include, but are not limited to, touch screens, touch pads, still and/or video cameras, microphones, computer mice, keyboards, buttons, keys, wired and/or wireless network
5 interfaces, communication devices, storage devices, displays, light emitting diodes (LEDs), printers, light bulbs, speakers, and audio/video ports.

[0027] Computing platform 100 can be configured with resource constraint negotiation capabilities, such as resource allocation recommendations, to aid software developers in writing applications that enhance performance of computing platform 100. As an example
10 of the resource constraint negotiation capabilities, a software process, such as Process 200A, running on computing platform 100 can generate a query message configured to request allocation of one or more resources. Once generated, Process 200A can send the query message to system monitor 160.

[0028] System monitor 160 can monitor an amount of the resource that is currently allocated
15 to process 200A and an amount of the resource for computing platform 100. Then, in response to the query message, system monitor 160 can evaluate how satisfying an allocation request for resources specified in the query message would affect computing platform 100. To satisfy an actual allocation request, system monitor 160 can allocate a resource such as memory, control unit usage, and/or bandwidth, create a resource such as a new thread,
20 reserve a resource, such as access to one of devices 170A-170N, and/or initiate a resource, such as starting a new process.

[0029] After determining how satisfying an allocation request for resources specified in the query message affects computing platform 100, system monitor 160 can generate a resource allocation recommendation regarding allocation of the requested resource. As one example,
25 system monitor 160 can respond with an “approved” resource allocation recommendation when computing platform 100 currently has enough resources to readily grant the allocation request. As another example, when computing platform 100 has just enough available resources, then system monitor 160 can respond with a “proceed with caution” resource allocation recommendation. Other recommendations are possible as well.

[0030] System monitor 160 can transmit the resource allocation recommendation to process
30 200A. Then, process 200A can use the resource allocation recommendation to determine whether to make an allocation request requesting the resources specified in the query message.

An Example System Monitor

[0031] FIG. 2 is a block diagram illustrating a system monitor in accordance with one or more aspects of this disclosure. As seen in FIG. 2, system monitor 160 includes several modules including process interface module 220, resource monitoring module 240, resource query module 260, request determination module 280, and resource allocation module 300.

[0032] Process interface module 220 can establish a channel of communication between system monitor 160 and each software process 200A-200N running on computing platform 100 to allow system monitor 160 to receive and transmit messages to and from one or more of processes 200A-200N. Via process interface module 220, each process 200A-200N can query system monitor 160 regarding the acceptability of a given request for a resource. In response to the query, system monitor 160 can inform the requesting process, via process interface module 220, how the additional resource allocation for the process would affect computing platform 100 and transmit a recommendation to the process on how the process should proceed with the request.

[0033] In addition, system monitor 160 can transmit message(s), via process interface module 220, to inform processes 200A-200N of changes to resource constraints. For example, system monitor 160 can send a resource change message to each of processes 200A-200N when memory has been added to computing platform 100 and thus indicate that processes 200A-200N have more memory available. A resource change message need not be in response to a query from a particular process but can instead be sent as an information message to some or all of processes 200A-200N.

[0034] Process interface module 220 can be implemented using various inter-process communication (IPC) techniques. For example, process interface module 220 can be implemented as a local socket, e.g., a Unix-like socket. Numerous other IPC techniques can be used to communicate between system monitor 160 and processes 200A-200N. For example, process interface module 220 can be implemented using remote procedure calls (RPCs).

[0035] Resource monitoring module 240 can monitor the resources of computing platform 100, either continuously or periodically. That is, resource monitoring module 240 can monitor available system-wide memory, bandwidth, processor usage, battery power, and/or other resources.

[0036] Resource query module 260 can receive a query message from one or more processes 200A-200N and evaluate, based on the system-wide status of the resources as monitored by resource monitoring module 240, the overall system-wide effect if system monitor 160

granted a request for additional resources as specified in the query message. To be clear, the query message is not a request for allocating additional resources. Rather, the querying process, e.g., process 200A, is asking system monitor 160 to determine the effects on computing platform 100 if a request were made for the resources listed in the query message.

5 In response to the query message, resource query module 260 generates and sends a resource allocation recommendation with respect to the additional amount of the resource proposed for allocation by the software process.

[0037] For example, suppose process 200A generates a query message to send to system monitor 160 and, in particular, resource query module 260, via process interface module 220.

10 The query message can specify a requesting process, e.g., process 200A, the resource, e.g., memory, and the amount of the resource, e.g., 5 megabytes, that it is considering requesting. In some embodiments, the query message can also specify a process type and/or a process status. In response, system monitor 160 can provide a resource allocation recommendation guiding the querying process on proceeding or not proceeding with a subsequent resource allocation request.

15 [0038] Request determination module 280 receives resource allocation requests for additional resources from requesting processes. As previously discussed, a requesting process, such as process 200A, can send a query message, and in response, receive a resource allocation recommendation. In response to the resource allocation request, requesting process 200A can send a resource allocation request to actually obtain additional resources from system monitor 160.

[0039] Upon receiving the resource allocation request via process interface module 220, request determination module 280 determines, based on the system-wide status of the resources as monitored by resource monitoring module 240, that providing the resources requested by requesting process 200A is either acceptable or unacceptable and correspondingly either grants or denies the resource allocation request. If request determination module 280 grants the resource allocation request, system monitor 160 allocates the requested resources to requesting process 200A via resource allocation module 300 and sends a resource allocation response with the requested resource. If request determination module 280 denies the request from requesting process 200A, system monitor 160 may, in some instances, terminate requesting process 200A.

Example Interactions between Processes and the System Monitor

[0040] FIG. 3 is an example interaction diagram for an interaction between system monitor 160 for computing platform 100 and an example process 200A executing on computing

platform 100, in accordance with one or more aspects of this disclosure. In the non-limiting example of FIG. 3, which is provided for purposes of illustration only, it is assumed that computing platform 100 has the following resource constraints: 100 megabytes (MB) of memory available, 50 MB of which are being used by operating system 140, e.g., a kernel and any privileged system processes. It is further assumed that the system requires that an emergency buffer of 10 MB of free memory be maintained and that a user-initiated process resident in memory, which has recently been used by a user, is currently using 10 MB of memory. In the description below, reference is made to one or more of the modules shown in FIG. 2 for purposes of illustration only. Any combination of one or more modules, including one or more of the modules shown in FIG. 2, can be utilized to perform the acts shown in FIGS. 3, 4A, and 4B.

[0041] FIG. 3 shows block 400 where process 200A can determine a need for additional resources, e.g., more memory. At block 420, process 200A can generate and send a query message, such as query message (QM) 430, to system monitor 160. For example, query message 430 can specify a requesting process e.g., process 200A, the type of resource, e.g., memory, and a resource amount, e.g., 5 MB, to be requested from computing platform 100.

[0042] At block 440, system monitor 160 can receive query message 430 via process interface module 220, and forward the query message 430 to resource query module 260, which monitors the system-wide status of resources available to computing platform 100. In some examples, resource monitoring module 240 can monitor resources of computing platform 100 prior to receiving query message 430. In other examples, resource monitoring module 240 can begin monitoring resources of computing platform 100 in response to query message 430.

[0043] Query message 430 can prompt system monitor 160 to evaluate the effect on computing platform 100 if a resource allocation request were made for resources specified in the query message. Based on the evaluation, system monitor 160 can generate and send a resource allocation recommendation to recommend how process 200A can proceed with a subsequent resource allocation request.

[0044] At block 460, resource query module 260 can generate, based on a status of resources of computing platform 100, resource allocation recommendation (RAR) 490 estimating an overall system-wide effect if system monitor 160 grants a request for the resources specified in the query message. In this example, based on the assumptions indicated above prior to the discussion of block 400, resource query module 260 can evaluate free memory to determine 30MB of non-emergency memory is available within computing platform 100 (100 MB

overall - 50 MB core - 10 MB emergency buffer - 10MB in use = 30 MB available), and that query message 430 proposes allocation of 5 MB, or about 17% of the free 30 MB. Resource query module 260 can also compare the amount of resources allocated to process 200A with that of other processes executing on computing platform 100 to determine whether a request
5 from process 200A for the proposed 5 MB allocation is reasonable. If the request is reasonable in comparison to other processes, then resource query module 260 can recommend allocation of 5 MB to process 200A, as 5 MB is a relatively small amount of the 30 MB of free memory.

[0045] At block 480, resource query module 260 can generate and send resource allocation
10 recommendation 490 to process 200A in response to query message 430. Resource allocation recommendation 490 can include a recommendation level, or indication of approval or disapproval concerning allocation of more resources to process 200A. In this manner, resource query module 260 determines the level of approval for resource allocation recommendation 490 based at least upon the amount of the resource that is currently available
15 on the computing platform.

[0046] In some embodiments, resource query module 260 can additionally or alternatively determine the recommendation level in resource allocation recommendation 490 based upon the amount of the resource that is currently allocated to the querying software process. For example, resource query module 260 can utilize a policy that restricts any process from
20 allocating more than a predetermined amount of the resource. That is, even if there are sufficient resources available, resource query module 260 can determine that the process sending a query message for a resource is already allocated a maximum allowable amount of the resource. Further, under this example policy, resource query module 260 can terminate the process if attempts to allocate more resources beyond the maximum allowable amount.
25 That is, resource query module 260 can determine the resource allocation recommendation based upon (a) the amount of the resource that is currently allocated to the software process, (b) the amount of the resource that is currently available on the computing platform, or (c) both the currently-allocated resource and the currently-available resource amounts.

[0047] At block 500, process 200A receives resource allocation recommendation 490 from
30 system monitor 160 and proceeds on that basis. In this example, upon receiving a favorable recommendation, process 200A can proceed to allocate additional memory using a resource allocation request. As another example, upon receiving an unfavorable recommendation, process 200A can try to free resources via a garbage collection operation, generate an “out of memory” error, or take some other action. In this manner, process 200A negotiates with

system monitor 160 for additional resources based on the resource constraints of computing platform 100.

[0048] As indicated above, system monitor 160 can generate a resource allocation recommendation with a recommendation level indicating a level of approval or disapproval for a proposed resource request. For instance, in the example above, process 200A queried the system monitor about requesting 5 MB of additional memory, with 30 MB of memory available to computing platform 100. In such an example, the system monitor can respond with a recommendation level such as “APPROVED” because there is more additional capacity available (30 MB) than what the process can request (5 MB). Other recommendation levels can include “APPROVED WITH CAUTION,” “MILDLY RECOMMEND AGAINST,” “MODERATELY RECOMMEND AGAINST,” and “STRONGLY RECOMMEND AGAINST.” Numeric values could be used as well; *e.g.*, APPROVED = 0, APPROVED WITH CAUTION = 1, etc. Of course, these are only some of the potential gradations that can indicate recommendation levels. More, different, and/or fewer, recommendation levels are considered within the scope of this disclosure.

[0049] The recommendation levels can form part of a system-wide resource policy, *e.g.*, a memory growth policy. As one example of a memory growth policy, assume that computing platform 100 allocates “normal processes” an initial amount of memory (*e.g.* 5 MB), and allows the normal process to increase allocated memory without complaint from system monitor 160 to a “minimal maximum” amount, such as 25 MB. For example, computing platform 100 can attempt to guarantee that a normal process that uses less than the minimal maximum will not be denied a request for additional memory. However, beyond the “minimal maximum” size of 25 MB of memory, computing platform 100 can increasingly disapprove of memory growth requests as the process grows beyond the 25 MB of memory.

In some embodiments, memory can be allocated as “pages” or fixed-size blocks, such as blocks of size one kilobyte (KB), two KBs, four KBs, or some other fixed size. Then, if 1 MB of memory were allocated to the process, the allocated memory can be provided as 1,024 one-KB pages, as 512 two-KB pages, or as 256 four-KB pages.

[0050] Minimal maximum values and/or resource policies can depend on process type. For example, an “immersive application” type of process can be determined to often require more resources, such as memory, than normal processes. Then, immersive application processes can be allocated more memory initially and have a higher minimal maximum than normal processes. However, as the system can see immersive application processes as potential resource “hogs”, one resource policy to moderate immersive application process resource

growth is to immediately terminate a given immersive application process upon requesting more memory beyond the minimal maximum for the given immersive application process.

More generally, immersive application (or similar resource intensive) processes can be allocated additional bandwidth, processor usage, and/or other resources, but can also be subject to tighter termination controls when resource requests exceed maxima such as a bandwidth maximum, a processor-usage maximum, and/or other-resource maxima.

[0051] In another aspect of the system-wide resource policy, resource query module can track which processes choose to follow (or not follow) resource allocation recommendations. For example, suppose a process P1 sends a query message for X MB of memory and, in response, receives a resource allocation recommendation with a recommendation level of “STRONGLY RECOMMEND AGAINST.” Then, further suppose that process P1 immediately requests X MB of memory, thus ignoring the resource allocation recommendation.

[0052] System monitor 160 could track how many resource allocation recommendations each process, including P1, either followed or ignored. Then, suppose system monitor 160 determined a need for more resources, and also determined a need to purge processes to reclaim resources to satisfy this need. In this scenario, processes such as P1 that ignored resource allocation requests can be purged before processes that had always followed resource allocation requests. The follow/ignore recommendation aspect of a resource allocation policy can be used as a primary factor to determine processes to be purged, or as one factor among many, such as amount of resources held, process age, and other factors, in considering which process(es) are to be purged.

[0053] Table 1 below includes example pseudocode that represents how resource query module 260 of system monitor 160 can respond to a request from a “normal” type of process as discussed above (as opposed to another type of process, for example):

```
// Example responses, in order, are:
// APPROVED, APPROVED_WITH_CAUTION, MILD_RECOMMEND_AGAINST,
// MODERATE_RECOMMEND_AGAINST, STRONG_RECOMMEND_AGAINST
if (requesting process's memory is below the minimal maximum)
{
    if (system has enough free/unreserved memory to satisfy
        the request) {
        respond APPROVED;
    } else if (there is any memory that can safely be freed)
    {
        respond APPROVED_WITH_CAUTION;
    } else {
```

```

        respond MILD_RECOMMEND_AGAINST;
    }
} else {
    // The process is interested in growing beyond the
    minimal maximum.
    if (system has enough free/unreserved memory to satisfy
    the request) {
        respond APPROVED_WITH_CAUTION;
    } else if (there are idle processes that can be evicted
    to make enough room) {
        respond MILD_RECOMMEND_AGAINST;
    } else if (there are larger processes that can be evicted
    to make room) {
        respond MODERATE_RECOMMEND_AGAINST;
    } else {
        respond STRONG_RECOMMEND_AGAINST;
    }
}

```

Table 1

[0054] As seen above, there are five levels of responses, or recommendation levels that can be generated by system monitor 160 via the example pseudocode shown in Table 1 above. Again, there can be more, or less, than five levels of responses in other examples. The recommendation levels can be represented, in one example, by bit patterns in the resource allocation recommendation message sent from system monitor 160 to the process. For example, in order to indicate one of the five recommendation levels listed above, a three bit field ($2^3=8$) in the resource allocation recommendation can represent the five levels.

[0055] In the example pseudocode of Table 1, system monitor 160 can generate and send an “APPROVED” resource allocation recommendation if the querying process, e.g., process 200A, is using an amount of memory that is below the predetermined “minimal maximum” amount of memory, e.g., 32 MB, and if computing platform 100 has enough free or unreserved memory to satisfy a request for the additional memory from the process. In response to receiving the “APPROVED” resource allocation recommendation from system monitor 160, the process, e.g., process 200A can proceed to request allocation of new memory.

[0056] If, however, computing platform 100 does not have enough free or unreserved memory pages to satisfy a request for the additional memory from the process, but there are memory pages that can be safely freed (e.g., memory devoted to buffering input/output or other pages that can be flushed because the data in the page exists elsewhere in computing platform 100), then system monitor 160 can generate and send an “APPROVED WITH CAUTION” resource allocation recommendation to the querying process, e.g., process 200A.

Process 200A can, in response to the “APPROVED WITH CAUTION” resource allocation recommendation, proceed to allocate new memory pages and/or initiate a garbage collection operation to reclaim memory pages no longer used by the process.

[0057] If no pages can be freed to satisfy the additional memory request, then system monitor 160 can generate and send an “MILD RECOMMEND AGAINST” resource allocation recommendation, to process 200A. In response to receiving the “MILD RECOMMEND AGAINST” response from system monitor 160, process 200A can initiate a garbage collection operation to reclaim pages no longer being used and allocate new memory pages if the garbage collection operation failed to free sufficient memory to satisfy the additional memory request.

[0058] In the example pseudocode presented in Table 1 above, the process, e.g., process 200A, can seek to grow beyond the “minimal maximum” size designated by computing platform 100. For example, process 200A can seek to increase its allocated memory beyond the predetermined 32 MB “minimal maximum” size specified by computing platform 100. In the example pseudocode presented above, system monitor 160 generates and sends an “APPROVED WITH CAUTION” response if the querying process, e.g., process 200A, is requesting an amount of memory that would exceed the predetermined “minimal maximum” amount of memory, e.g., 32 MB, and if computing platform 100 has enough free or unreserved memory pages to satisfy a request for the additional memory from the process.

As described above, in response to receiving the “APPROVED WITH CAUTION” response from system monitor 160, the process, e.g., process 200A can proceed to allocate new memory pages, but also initiate a garbage collection operation to reclaim pages that are no longer being used.

[0059] If one or more idle processes could be terminated such that sufficient memory could be freed to satisfy the request once the idle process(es)’ memory had been reclaimed, system monitor 160 can generate and send a “MILD RECOMMEND AGAINST” response. As described above, in response to receiving the “MILD RECOMMEND AGAINST” response from system monitor 160, the process, e.g., process 200A can proceed to initiate a garbage collection operation to reclaim memory no longer being used by the process, and allocate new memory only if the garbage collection operation failed to free sufficient memory.

[0060] If there are no idle processes that could be terminated to free sufficient memory for the request, but larger processes could be terminated that could provide sufficient pages to satisfy the request, system monitor 160 can generate and send a “MODERATE RECOMMEND AGAINST” response.

[0061] In response to receiving the “MODERATE RECOMMEND AGAINST” response from system monitor 160, process 200A can proceed in a manner similar to having received a “MILD RECOMMEND AGAINST” response, but use a more aggressive garbage collection policy. An example aggressive garbage collection policy can include freeing all “softly-referenced” objects. A softly-referenced object is an object designated by a programmer to be optionally freed by a garbage-collected process. By designating an object as a softly-referenced object, the programmer has signaled explicitly that the softly-referenced object can be freed if necessary to satisfy a higher-priority request.

[0062] In the case that the request cannot be satisfied by terminating either idle or larger processes, system monitor 160 can generate and send a “STRONG RECOMMEND AGAINST” response. In response to receiving the “STRONG RECOMMEND AGAINST” response from system monitor 160, the process, e.g., process 200A can generate a suitable error; e.g., an “OUT OF MEMORY” error.

Example Scenario with Resource Allocation Recommendations

[0063] FIGS. 4A and 4B depict an example scenario 600 for an interaction between a computing platform and three processes, in accordance with one or more aspects of this disclosure. In scenario 600, an example resource 180 of computing platform 100 is 100 MB of memory. Other example resources 180 could be used as well or instead. FIG. 4A shows allocated portions of resource 180 in grey, reserved memory using cross-hatching, and unallocated portions in white.

At the beginning of scenario 600 and as shown at the top of FIG. 4A, (a) 50 MB of resource 180 are allocated to core operating system (OS) 140, (b) three user-initiated Processes A, B, and C are each currently allocated 10 MB of memory, each of which having been recently used by the user, and (c) 10 MB is reserved as an emergency buffer (EB). In scenario 600, system monitor 160 attempts to maintain a 10 MB “emergency buffer minimum”, or minimum amount of memory allocated to the emergency buffer. As shown in FIG. 4A, at the beginning of scenario 600, 10 MB of free (F) or unallocated memory is available for use by user-initiated processes such as Processes A, B, and C.

[0064] Scenario 600 continues with Process A determining a need for additional memory. In preparation for requesting more memory, Process A generates and sends query message 610. Query message 610 asks system monitor 160 to determine the effects on computing platform 100 if a request were made for resources specified in query message 610 and to provide an opinion on how Process A should proceed. In scenario 600, query message 610 specifies a

requesting process of Process A, a type of resource of memory, and a resource amount of 5 MB that the requesting process proposed for allocation from computing system 100.

[0065] System monitor 160 receives query message 610 using process interface module 220, e.g., a socket, and forwards query message 610 to resource query module 260. Resource query module 260 makes evaluation 612 to consider the overall system-wide effect of granting a request for resources specified in query message 610, based on the system-wide status of at least resource 180.

[0066] In making evaluation 612, resource query module 260 determines that 10MB of non-emergency memory is available to processes within computing platform 100. Also, granting a request to Process A for an additional 5 MB would leave 5 MB of memory still free in resource 180. Further, resource query module 260 determines that providing process A with an additional 5 MB of memory, for a total of 15 MB, is not excessive compared to the 10 MB used by similar Processes B and C. Thus, system monitor 160 can determine that the request would be approved.

[0067] As shown in FIG. 4A, after evaluation 612, system monitor 160 generates and sends resource allocation recommendation (RAR) 620 indicating approval of a request as specified in query message 610. In scenario 600, resource allocation recommendation 620 indicates the request is "APPROVED" using the recommendation levels described above in the context of Table 1.

[0068] Based on resource allocation recommendation 620 indicating approval, Process A can generate and send allocation request 622, requesting allocation of an additional 5 MB of resource 180. In response to resource allocation request 622, resource allocation module 300 allocates the requested 5 MB of memory to Process A. System monitor 160 can inform Process A of the allocation via resource allocation response 624.

[0069] In some embodiments, when evaluation 612 determines that a resource allocation request based on query message 610 has an APPROVED recommendation level, and perhaps other recommendation levels, system monitor 160 can allocate the requested resources in response rather than responding with a resource allocation recommendation. That is, in response to query message 610, system monitor can send a resource allocation response indicating allocation of the requested resources, such as resource allocation response 624. In some of these embodiments, the resource allocation response can include data representing the recommendation level that would be provided with a resource allocation response.

[0070] Scenario 600 continues with Process B determining a need for an additional 5 MB of resource 180. Process B then generates and sends query message 630 with a requesting process of Process B, a type of resource of memory, and a resource amount of 5 MB.

5 [0071] As shown in FIG. 4A, system monitor 160 receives query message 630 from Process B. Query message 630 is forwarded to resource query module 260 to make evaluation 632 to consider the effect of granting a request for resources specified in query message 630.

[0072] In making evaluation 632, resource query module 260 determines that 5 MB of non-emergency memory is available to processes within computing platform 100. Also, granting a request to Process B for an additional 5 MB would leave no free memory for resource 180, outside of any free memory in the emergency buffer. In addition, resource query module 260 can determine that Process B's resource allocation of 15 MB is the same as that of Process A, and similar to the 10 MB allocated to Process C. Therefore, Process B's allocation request is reasonable in comparison to resources used by similar Processes A and C, and so the request should be approved, but the lack of non-emergency free memory after approval indicates that
10 caution is warranted.

[0073] As shown in FIG. 4A, after evaluation 632, system monitor 160 generates and sends resource allocation recommendation 640 to Process B. In scenario 600, resource allocation recommendation 640 indicates cautionary approval of the request specified in query message 630 with the "APPROVED WITH CAUTION" recommendation level of the
20 recommendation levels described above in the context of Table 1.

[0074] In scenario 600, Process B can request 5 MB of resource 180 in accord with an APPROVED WITH CAUTION recommendation level of resource allocation recommendation 640. Thus, Process B can generate and send allocation request 642 to requesting allocation of an additional 5 MB memory from resource 180. In response to
25 resource allocation request 642, resource allocation module 300 allocates the requested 5 MB of resource 180 to Process B and informs Process A of the 5 MB allocation from resource 180 via resource allocation response 644.

[0075] Turning to FIG. 4B, as the request is approved with caution, Process B can initiate measures to reclaim resource 180, such as invoking a garbage collection operation to reclaim
30 memory pages no longer being used by Process B. The garbage collection operation for a process can be invoked via garbage collection message 646 and/or can involve an internal memory-reclaiming operation of the process. In scenario 600, Process B invokes garbage collection garbage collection message 646 to system monitor 160. Process B can request the additional 5MB of resource 180 before, after, or simultaneously with invocation of the

garbage collection operation. In scenario 600, this garbage collection operation did not reclaim an appreciable amount of memory from Process B.

[0076] Scenario 600 continues with Process A determining a need for an additional 5 MB from resource 180. Process A then generates and sends query message 650 that specifies a requesting process of Process A, a type of resource of memory, and a resource amount of 5 MB that the requesting process proposed for allocation from computing system 100.

[0077] As shown in FIG. 4B, system monitor 160 receives query message 650 from Process B. Query message 650 is forwarded to resource query module 260 to make evaluation 652 to consider the effect of granting a request for additional resources from Process A as specified in query message 650, based on a system-wide status of at least resource 180.

[0078] In making evaluation 652, resource query module 260 determines that, while Process A is active, reclaimed and/or emergency buffer memory would have to be used to satisfy a request corresponding to query message 650. If no memory is reclaimed, system monitor 160 would have to satisfy the request from emergency buffer memory alone, which would take the emergency buffer below the emergency buffer minimum. Therefore, resource query module 260 can strongly recommend against accepting a request based on query message 650.

[0079] After evaluation 652, system monitor 160 generates and sends resource allocation recommendation 660 to Process A. In scenario 600, resource allocation recommendation 660 indicates lack of approval of the request specified in query message 650 with a “STRONG RECOMMEND AGAINST” level of the recommendation levels described above in the context of Table 1.

[0080] Upon receiving resource allocation recommendation 660 with a strong recommendation against allocation of resource 180, FIG. 4B shows that Process A utilizing internal functionality to perform aggressive garbage collection (GC) 662a. An aggressive garbage collection operation for Process A can include a normal garbage collection operation to reclaim memory pages that are no longer being used by Process A (*e.g.*, a garbage collection operation such as performed for Process B as discussed above). An aggressive garbage collection operation can further include reclaiming memory pages that are only softly-referenced by Process A. In scenario 600 and as shown in FIG. 4B, the aggressive garbage collection operation frees up 1 MB of resource 180 for system monitor 160 to reclaim from Process A, leaving Process A with 14 MB of allocated memory and 1 MB of free memory in resource 180. Process A uses Free Memory message 662b to inform system monitor 160 about and/or provide system monitor 160 with the freed 1 MB of memory. In

some embodiments, Process A can send multiple Free Memory messages 662b to inform and/or provide memory to system monitor 160.

[0081] Subsequently, Process A sends resource allocation request 664 to request 5 MB of resource 180. In scenario 600, resource allocation request 664 can be granted using the 1 MB of free memory and 4 MB from the emergency buffer of resource 180. After granting resource allocation request 664, resource 180 has 6 MB remaining in the emergency buffer, and Process A has allocated a total of 19 MB. System monitor 160 sends resource allocation response 666 providing the allocated portion of resource 180 to Process A.

[0082] In some scenarios, resource allocation request 664 can be delayed in response to the recommendation against resource allocation, with the intent of waiting for some other process(es) to free up part or all resource 180. For example, Process A can wait a short length of time *e.g.*, a few nanoseconds, or some longer length of time before sending resource allocation request 664. The length of time to delay sending resource allocation request 664 can depend upon the resource(s) being requests, the amount(s) of resource(s) being requested, a recommendation level received in the resource allocation recommendation 660, recommendation level(s) received in resource allocation recommendation(s) prior to resource allocation recommendation 660, amount(s) of resource(s) allocated to the process; *e.g.*, Process A, and other factors as well or instead.

[0083] As another example, if resource allocation recommendation 660 indicates lack of approval of the request specified in query message 650 with a “STRONG RECOMMEND AGAINST” level, or perhaps a “MODERATE RECOMMEND AGAINST” level of the recommendation levels described above in the context of Table 1, a priority of the application can be determined. If the priority is relatively low, system monitor 160 and/or the application itself can voluntarily postpone some or all operations until a later resource allocation recommendation 660 indicates a higher level of approval of the request. In particular embodiments, the application can free resources or perhaps self-terminate to allow other processes to complete using the now-available resources.

[0084] Subsequently, system monitor 160 determines a need to restore the emergency memory buffer to the emergency buffer minimum of 10 MB. During scenario 600, Process C has been idle and so system monitor 160 can terminate at least one idle software process, such as Process C, when an emergency buffer is below a threshold level, such as an emergency buffer minimum of 10 MB.

[0085] To free up memory for the emergency buffer, system monitor 160 can terminate idle process C via terminate message 670. Upon termination, Process C releases 10 MB of

resource 180, thereby decreasing memory usage of computing platform 100 below the emergency buffer zone of 10 MB. Then, 4 MB of the 10 MB reclaimed from now-terminated Process C can be allocated to the emergency buffer, restoring the emergency buffer to the emergency buffer minimum of 10 MB. As shown in FIG. 4B, the remaining 6
5 MB of reclaimed memory can be designated as free memory.

[0086] System monitor 160 can inform some or all processes of changes to resource constraints. For example, in response to termination from Process C, scenario 600 continues with system monitor 160 generating and transmitting status change message 672 to Processes A and B. Status change message 672 can include a resource indicator for indicating the
10 resource whose status has changed, a free-resource value indicating a free amount of the resource, and perhaps other data, such as a recommendation level. FIG. 4B shows status change message indicating a status change to resource 180 to currently indicating that 6 MB of memory of resource 180 is free, and that resource 180 is associated with an APPROVED recommendation level. In this fashion, system monitor 160 can indicate to processes
15 executing on computing platform 100 that there is no longer an impending memory emergency.

[0087] In scenarios not shown in FIGS. 4A-4B, system monitor 160 can use other criteria along with or instead of the idle process criterion discussed above to determine process(es) for termination. For example, suppose that Process C was active during scenario 600 when
20 memory was allocated from the emergency buffer to take the emergency buffer below the emergency buffer minimum. Then, to restore the emergency buffer, system monitor 160 can use criteria such as terminating a process with a maximum or, in some cases, a minimum amount of resource 180 allocated over all processes. Other example criteria include terminating a process making the most, or in some cases, least recent resource allocation
25 request.

[0088] Additional criteria for process termination can be based on compliance/non-compliance with resource allocation recommendations. That is, system monitor 160 and/or other component(s) of computing platform 100 could track compliance with resource allocation recommendations and terminate process(es) that do not comply with resource
30 allocation recommendations. For example, suppose a Process IG receives a resource allocation recommendation that recommends against making a resource allocation request, and Process IG ignores the recommendation and persists in making a resource allocation request without attempting to reclaim resources or otherwise ameliorate the negative recommendation. Then, system monitor 160 and/or other component(s) of computing

platform 100 can note that Process IG ignored a resource allocation recommendation not to request resources, and therefore increase the probability that Process IG will be terminated when necessary. Many other criteria for terminating processes are possible as well.

Example Operations

5 [0089] FIG. 5 is a flow chart illustrating an example method 700 of resource constraint negotiation on a computing platform, such as computing platform 100, in accordance with one or more aspects of this disclosure.

[0090] Method 700 begins at block 710 where a query message can be received at the computing platform. The query message can specify an amount of a resource proposed for
10 allocation from the computing platform, such as discussed above in the context of at least FIGS. 4A-4B. For example, the resource can be memory, processor usage, bandwidth, and/or other resources.

[0091] At block 720, the computing platform can select a selected recommendation level from a plurality of recommendation levels. The selection of the selected recommendation
15 level can be based on an evaluation of a request for the amount of the resource proposed for allocation, such as discussed above in the context of at least FIGS. 4A-4B. An example plurality of recommendation levels is discussed above at least in the context of Table 1.

[0092] In some embodiments, the evaluation of the request for the amount of the resource is based at least upon an amount of the resource allocated to a software process and an amount
20 of the resource currently available on the computing platform, such as discussed above in the context of at least FIGS. 4A-4B. In other embodiments, the evaluation of the request for the amount of the resource is based at least on usage of the resource on the computing platform and resource usage of a software process associated with the query message compared to at least one other software process of the software processes, such as discussed above in the
25 context of at least FIGS. 4A-4B.

[0093] At block 730, the computing platform can generate a resource allocation recommendation, which can include the selected recommendation level, such as discussed above in the context of at least FIGS. 4A-4B and Table 1.

[0094] At block 740, the computing platform can send a resource allocation
30 recommendation, such as discussed above in the context of at least FIGS. 4A-4B.

[0095] In some embodiments, in response to sending the resource allocation recommendation, a resource allocation request for the amount of the resource can be received, such as discussed above in the context of at least FIGS. 4A-4B. In particular embodiments, in response to receiving the resource allocation request, method 700 can

include (a) allocating the amount of the resource proposed for allocation and (b) generating a resource allocation response that includes the allocated amount of the resource, such as discussed above in the context of at least FIGS. 4A-4B.

[0096] In other embodiments, in response to sending the resource allocation
5 recommendation, a message to reclaim the resource can be received, such as discussed above in the context of at least FIG. 4B.

[0097] In some embodiments, method 700 can additionally include terminating at least one software process executing on the computing platform in response to an available amount of the resource on the computing platform being below a threshold level, such as discussed
10 above in the context of at least FIG. 4B.

[0098] In other embodiments, method 700 can additionally include terminating at least one software process executing on the computing platform based on non-compliance with resource allocation recommendations, such as discussed above in the context of at least FIG. 4B.

15 [0099] In still other embodiments, method 700 can additionally include transmitting a message indicating a change in a status of the resource, such as discussed above in the context of at least FIG. 4B.

[00100] Figure 6 is a flow chart of an example method 800 of resource constraint negotiation on a computing platform, such as computing platform 100, in accordance with
20 one or more aspects of this disclosure.

[00101] Method 800 begins at block 810 where a query message can be send from the computing platform. The query message can specify an amount of a resource proposed for allocation from the computing platform, such as discussed above at least in the context of FIGS. 4A-4B.

25 [00102] At block 820 in response to the query message, the computing platform can receive a resource allocation recommendation related to the amount of the resource proposed for allocation. The resource allocation recommendation can include a selected recommendation level of a plurality of recommendation levels, such as discussed above in the context of at least FIGS. 4A-4B and Table 1.

30 [00103] At block 830, the computing platform can generate a response to the resource allocation recommendation. The response can be based on the selected recommendation level, such as discussed above in the context of at least FIGS. 4A-4B.

[00104] In some embodiments, the selected recommendation level is an approval level and the response includes a resource allocation request for the amount of the resource proposed for allocation, such as discussed above in the context of at least FIGS. 4A-4B.

[00105] In other embodiments, the selected recommendation level is a
5 recommendation-against level, and the response includes a delayed resource allocation request for the amount of the resource proposed for allocation, such as discussed above in the context of at least FIG. 4B.

[00106] In even other embodiments, the selected recommendation level is a recommendation-against level, and the response includes a garbage collection message, such
10 as discussed above in the context of at least FIG. 4B.

[00107] At block 840, the computing platform can send the response to the resource allocation recommendation, such as discussed above in the context of at least FIGS. 4A-4B.

[00108] In some embodiments, method 800 can additionally include receiving a resource allocation response in response to the resource allocation request. The resource
15 allocation response can include an allocated amount of the resource. The allocated amount of the resource can equal or exceed the amount of the resource proposed for allocation, such as discussed above in the context of FIGS. 4A-4B.

Conclusion

[00109] The above detailed description describes various features and functions of the disclosed systems, devices, and methods with reference to the accompanying figures. In the
20 figures, similar symbols typically identify similar components, unless context dictates otherwise. The illustrative embodiments described in the detailed description, figures, and claims are not meant to be limiting. Other embodiments can be utilized, and other changes can be made, without departing from the spirit or scope of the subject matter presented
25 herein. It will be readily understood that the aspects of the present disclosure, as generally described herein, and illustrated in the figures, can be arranged, substituted, combined, separated, and designed in a wide variety of different configurations, all of which are explicitly contemplated herein.

[00110] With respect to any or all of the ladder diagrams, scenarios, and flow charts in
30 the figures and as discussed herein, each block and/or communication may represent a processing of information and/or a transmission of information in accordance with example embodiments. Alternative embodiments are included within the scope of these example embodiments. In these alternative embodiments, for example, functions described as blocks, transmissions, communications, requests, responses, and/or messages may be executed out of

order from that shown or discussed, including substantially concurrent or in reverse order, depending on the functionality involved. Further, more or fewer blocks and/or functions may be used with any of the ladder diagrams, scenarios, and flow charts discussed herein, and these ladder diagrams, scenarios, and flow charts may be combined with one another, in part or in whole.

[00111] A block that represents a processing of information may correspond to circuitry that can be configured to perform the specific logical functions of a herein-described method or technique. Alternatively or additionally, a block that represents a processing of information may correspond to a module, a segment, or a portion of program code (including related data). The program code may include one or more instructions executable by a processor for implementing specific logical functions or actions in the method or technique. The program code and/or related data may be stored on any type of computer readable medium such as a storage device including a disk or hard drive or other storage medium.

[00112] The computer readable medium may also include non-transitory computer readable media such as computer-readable media that stores data for short periods of time like register memory, processor cache, and random access memory (RAM). The computer readable media may also include non-transitory computer readable media that stores program code and/or data for longer periods of time, such as secondary or persistent long term storage, like read only memory (ROM), optical or magnetic disks, compact-disc read only memory (CD-ROM), for example. The computer readable media may also be any other volatile or non-volatile storage systems. A computer readable medium may be considered a computer readable storage medium, for example, or a tangible storage device.

[00113] Moreover, a block that represents one or more information transmissions may correspond to information transmissions between software and/or hardware modules in the same physical device. However, other information transmissions may be between software modules and/or hardware modules in different physical devices.

[00114] While various aspects and embodiments have been disclosed herein, other aspects and embodiments will be apparent to those skilled in the art. The various aspects and embodiments disclosed herein are for purposes of illustration and are not intended to be limiting, with the true scope and spirit being indicated by the following claims.

CLAIMS

1. An article of manufacture including a non-transitory computer-readable storage medium having instructions stored thereon that, in response to execution by a processor of a computing platform, cause the computing platform to perform operations, the instructions comprising:

instructions for determine whether a resource is required for execution of a process;

instructions to, in response to determining that the resource is required, sending a query message from the process to a system monitor, wherein the query message specifies an amount of the resource proposed for allocation;

instructions for, in response to the query message, the system monitor to determine a system effect based on an evaluation of a request for the amount of the resource proposed for allocation;

instructions for the system monitor to select a selected recommendation level from a plurality of recommendation levels based on the evaluation of the request for the amount of the resource proposed for allocation;

instructions for the system monitor to generate a resource allocation recommendation comprising the selected recommendation level;

instructions for sending the resource allocation recommendation from the system monitor to the process; and

instructions to execute the process based on the received resource allocation recommendation.

2. The article of manufacture of claim 1, wherein the evaluation of the request for the amount of the resource is based at least upon an amount of the resource allocated to a software process and an amount of the resource currently available on the computing platform.

3. The article of manufacture of claim 1, wherein the evaluation of the request for the amount of the resource is based at least on usage of the resource on the computing platform and resource usage of a software process associated with the query message compared to at least one other software process.

4. The article of manufacture of claim 1, further comprising instructions to:

terminate at least one software process executing on the computing platform in response to an available amount of the resource on the computing platform being below a threshold level.

5 5. The article of manufacture of claim 1, further comprising instructions to:
terminate at least one software process executing on the computing platform based on non-compliance with resource allocation recommendations.

10 6. The article of manufacture of claim 1, further comprising instructions to:
transmit a message indicating a change in a status of the resource.

7. The article of manufacture of claim 1, wherein the instructions further comprise:
instructions for, in response to sending the resource allocation recommendation,
receiving a resource allocation request for the amount of the resource at the system monitor.

15 8. The article of manufacture of claim 7, wherein the instructions further comprise
instructions for, in response to receiving the resource allocation request, the system monitor
to:

20 allocate the amount of the resource proposed for allocation; and
generate a resource allocation response that comprises the allocated amount of the
resource.

9. The article of manufacture of claim 1, wherein the resource comprises one or more of
volatile memory, bandwidth, and disk storage.

25 10. The article of manufacture of claim 1, wherein the instructions further comprise:
instructions for, in response to sending the resource allocation recommendation,
receiving a message at the system monitor to reclaim the resource.

30 11. A method, comprising:
receiving a query message at a computing platform, wherein the query message
specifies an amount of a resource proposed for allocation from the computing platform;

selecting a selected recommendation level from a plurality of recommendation levels based on an evaluation of a request for the amount of the resource proposed for allocation using the computing platform;

- generating, using the computing platform, a resource allocation recommendation
5 comprising the selected recommendation level; and
sending the resource allocation recommendation using the computing platform.

12. The method of claim 11, wherein the evaluation of the request for the amount of the resource is based at least upon an amount of the resource allocated to a software process and
10 an amount of the resource currently available on the computing platform.

13. The method of claim 11, wherein the evaluation of the request for the amount of the resource is based at least on resource usage on the computing platform and resource usage of a software process associated with the query message compared to at least one other software
15 process.

14. The method of claim 11, further comprising:
terminating at least one software process executing on the computing platform in
response to an available amount of the resource on the computing platform being below a
20 threshold level.

15. The method of claim 11, further comprising:
terminating at least one software process executing on the computing platform based
on non-compliance with resource allocation recommendations.

25

16. The method of claim 11, further comprising:
transmitting a message indicating a change in a status of the resource.

17. The method of claim 11, further comprising:
30 receiving a resource allocation request for the amount of the resource, in response to
sending the resource allocation recommendation.

18. The method of claim 17, further comprising, in response to the receiving the resource allocation request:

allocating the amount of the resource proposed for allocation; and
generating a resource allocation response that comprises the allocated amount of the
resource

5 19. The method of claim 11, further comprising:

receiving a message to reclaim the resource, in response to sending the resource
allocation recommendation.

10 20. The method of claim 11, wherein the resource comprises one or more of volatile
memory, bandwidth, and disk storage.

21. An article of manufacture including a non-transitory computer-readable storage
medium having instructions stored thereon that, in response to execution by a processor of a
computing platform, cause the computing platform to perform operations, the instructions
15 comprising:

instructions for sending a query message that specifies an amount of a resource
proposed for allocation;

instructions for, in response to the query message, receiving a resource allocation
recommendation related to the amount of the resource proposed for allocation, the resource
20 allocation recommendation comprising a selected recommendation level of a plurality of
recommendation levels;

instructions for generating a response to the resource allocation recommendation,
wherein the response is based on the selected recommendation level; and

instructions for sending the response to the resource allocation recommendation.
25

22. The article of manufacture of claim 21, wherein the selected recommendation level is
an approval level, and wherein the response comprises a resource allocation request for the
amount of the resource proposed for allocation.

30 23. The article of manufacture of claim 22, wherein the instructions further comprise:
instructions for, in response to the resource allocation request, receiving a resource
allocation response that comprises an allocated amount of the resource.

24. The article of manufacture of claim 22, wherein the allocated amount of the resource equals or exceeds the amount of the resource proposed for allocation.

25. The article of manufacture of claim 21, wherein the selected recommendation level is a recommendation-against level, and wherein the response comprises a garbage collection message.

26. A method, comprising:

10 sending a query message from a computing platform, wherein the query message specifies an amount of a resource proposed for allocation from the computing platform;
in response to the query message, receiving a resource allocation recommendation related to the amount of the resource proposed for allocation at the computing platform, the resource allocation recommendation comprising a selected recommendation level of a plurality of recommendation levels;
15 generating, using the computing platform, a response to the resource allocation recommendation wherein the response is based on the selected recommendation level; and
sending the response to the resource allocation recommendation from the computing platform.

20 27. The method of claim 26, wherein the selected recommendation level is an approval level, and wherein the response comprises a resource allocation request for the amount of the resource proposed for allocation.

28. The method of claim 26, further comprising:

25 in response to the resource allocation request, receiving a resource allocation response that comprises an allocated amount of the resource, wherein the allocated amount of the resource equals or exceeds the amount of the resource proposed for allocation.

29. The method of claim 26, wherein the selected recommendation level is a recommendation-against level, and wherein the response comprises a delayed resource allocation request for the amount of the resource proposed for allocation.

30. The method of claim 26, wherein the selected recommendation level is a recommendation-against level, and wherein the response comprises a garbage collection message.






FIG. 1

**FIG. 2**

**FIG. 3**




FIG. 4A




FIG. 4B




FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2012/022571

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/50
ADD.

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, IBM-TDB, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2009/183168 A1 (UCHIDA SATOSHI [JP]) 16 July 2009 (2009-07-16) paragraphs [0015] - [0017], [0038] - [0043], [0048] - [0057], [0060] - [0066], [0082], [0239] -----	1-21
X	US 2006/143617 A1 (KNAUERHASE ROBERT C [US] ET AL) 29 June 2006 (2006-06-29) paragraphs [0027], [0028]; figure 3 -----	1-30

☐

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

30 March 2012

Date of mailing of the international search report

11/04/2012

Name and mailing address of the ISA/

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

Authorized officer

Thibaudeau, Jean

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2012/022571

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2009183168 A1	16-07-2009	JP 2009169672 A US 2009183168 A1	30-07-2009 16-07-2009
US 2006143617 A1	29-06-2006	NONE	