



- (51) International Patent Classification:
G06F 1/32 (2006.01) *G06F 1/26* (2006.01)
- (21) International Application Number:
PCT/US2015/054438
- (22) International Filing Date:
7 October 2015 (07.10.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/582,985 24 December 2014 (24.12.2014) US
- (71) Applicant: **INTEL CORPORATION** [US/US]; 2200 Mission College Blvd., Santa Clara, California 95054 (US).
- (72) Inventors: **BODAS, Devadatta V.**; 35224 7th Avenue SW, Federal Way, Washington 98023-8144 (US). **RA-JAPPA, Muralidhar**; 3195 W. Tyson Place, Chandler, Arizona 85226 (US). **SONG, Justin J.**; 2403 Craig Road

SE, Olympia, Washington 98501 (US). **HOFFMAN, Andy**; 231 Clay Court SE, Olympia, Washington 98513 (US). **PATTERSON, Michael K.**; 2800 Center Drive North, Dupont, Washington 98327 (US).

(74) Agent: **FARIES, Floron C.**; International IP Law Group, P.L.L.C., c/o CPA Global, P.O. Box 52050, Minneapolis, Minnesota 55402 (US).

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

[Continued on next page]

(54) Title: CONTROL OF POWER CONSUMPTION




FIG. 6 106

[Continued on next page]

WO 2016/105625 A1

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

Declarations under Rule 4.17:

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

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

(57) Abstract: A system and method for computing at a facility having systems of multiple compute nodes to execute jobs of computing. Power consumption of the facility is managed to within a power band. The power consumption may be adjusted by implementing (e.g., by a power balloon) activities having little or no computational output.

CONTROL OF POWER CONSUMPTION

Cross-Reference to Related Applications

[0001] The present application claims the benefit of the filing date of U.S. Patent Application No. 14/582,985, filed December 24, 2014, which is incorporated herein by reference.

Technical Field

[0002] The present techniques relate generally to supply of electricity to a facility performing distributed computing. More particularly, the techniques relate to management of power consumption of the facility.

Background Art

[0003] A variety of computing configurations and systems may consume relatively large amounts of power. Such systems may include those for cloud computing, Big Data analytics, web services, enterprise services, distributed computing, high performance computing (HPC), and so on. The facility housing such systems can require significant energy.

[0004] Distributed computing, HPC systems, and other aforementioned configurations may facilitate scientists and engineers to solve complex science, engineering, and business problems using applications that benefit from high bandwidth, low latency networking, and very high compute capabilities. These systems may also execute data storage and retrieval, perform more straightforward tasks, and so on. Unfortunately, again, systems such as distributed systems, HPC systems, and others, which may have hundreds or thousands of processors, servers, or compute nodes performing tasks, typically consume significant power. Such may be especially problematic in the "Big Data" era. Further, variations in power consumption and issues of power allocation may also be problematic.

[0005] The competitive business of data and computing services drives manufacturers in the continuous improvement of their processes and products in order to lower production costs and deliver reliable service. Indeed, as technologies advance in services for data, computing, and telecommunications, a competitive

need exists to continuously increase consistency of service and the efficiency of power utilization.

BRIEF DESCRIPTION OF DRAWINGS

- [0006] FIG. 1 is a diagrammatical representation of an exemplary distributed computing system in accordance with embodiments of the present techniques.
- [0007] FIG. 2 is a block diagram of an exemplary facility for distributed computing in accordance with embodiments of the present techniques.
- [0008] FIG. 3 is a diagrammatical representation of an exemplary facility for distributed computing in accordance with embodiments of the present techniques.
- [0009] FIG. 4 is a diagrammatical representation of a system level and node level of an exemplary facility for distributed computing in accordance with embodiments of the present techniques.
- [0010] FIGS. 5A and 5B area flow diagram of an example of a power manager hierarchy and actions in accordance with embodiments of the present techniques.
- [0011] FIG. 6 is block diagram of exemplary managers and functions at the system level and node level in accordance with embodiments of the present techniques.
- [0012] FIG. 7 is a bar chart comparing power consumption of two example jobs of distributed computing in accordance with embodiments of the present techniques.
- [0013] FIG. 8 is a block diagram depicting an example of a tangible non-transitory, computer-readable medium that can facilitate power management of a distributed computing system in accordance with embodiments of the present techniques.
- [0014] The same numbers are used throughout the disclosure and the figures to reference like components and features. Numbers in the 100 series refer to features originally found in Fig. 1; numbers in the 200 series refer to features originally found in Fig. 2; and so on.

DETAILED DESCRIPTION

[0015] Embodiments of the present techniques are directed to a facility having systems of multiple compute nodes configured to execute jobs of computing. Power

consumption is managed or controlled at the facility to within a power band and/or within a desired or specified ramp or rate of change in power. The rate of change may be some magnitude of Watts per minute or per hour, for example. The management or control of power to within band or ramp may be at the facility level, system level, and job/node level.

[0016] Further, it should be noted that controlling power consumption to within a power band is much different than merely only making efforts to use less power or less electricity, as is conventional. Instead, a power band, for example, may be a specified minimum to a specified maximum of power, and in embodiments, the band is a relatively narrow range. For instance, the power band may be related to a negotiated range with a utility provider. A negotiated allocated amount of power from the utility provider to the facility may be a relatively narrow range of power (e.g., an allocation of 10 MW allowed to drop to 8.5 MW but not below). Thus, with the power band, the maximum limit of the band may be an allocated amount of power consumption (e.g., 10 MW). The minimum limit of the band may be the lower amount (e.g., 8.5 MW) that the power consumption is to reach, as negotiated via the allocation. It should be noted that facility use of local energy generation and storage may provide flexibility (e.g., widening the facility control power-band slightly) with respect to power consumption at the facility and satisfying negotiated allocation with the utility provider. Nevertheless, with or without local energy generation/storage, the facility may implement activities to control power consumption at the facility to within power band.

[0017] Various activities may be implemented to adjust power consumption. An example of a facility-level activity may be to waste power by lowering the temperature of a cooling system beyond that needed. Such may be implemented to increase power consumption so that the facility power consumption stays within a power band. Another example of a facility-level activity may be to activate or adjust local energy generation/storage to maintain power consumption within power band or within power ramp limits or targets. Energy generation may be from renewable and/or non-renewable sources disposed local to the facility, for example. Energy storage may involve power storage, (e.g., batteries), thermal storage (e.g., ice, cooled liquid), and so on.

[0018] To manage and control power, activities with low-priority computational output or no computational output may be implemented at the system and job/node levels to affect power consumption. As discussed below, such activities may be implemented by a power balloon module, at the system level and at the job/node level. The activities may be to increase or decrease power consumption, with respect to both band and rate. For example, a power balloon may specify or initiate an activity that wastes energy so to slow a sudden decrease in power. In another example, a manager module may initiate an activity that increases power with no computational output. For instance, the activity may be to power an idle node (without executing a job on that node) to waste/expend power, so to increase power consumption of the system to above a specified minimum target.

[0019] Again, a power balloon or manager module may initiate activities having little or no computation/work output, or no additional computational output, to adjust and control power at the system and/or node levels. These activities may include “green” and “non-green” activities.

[0020] The green techniques are useful but produce no computational output or no new computational job output, or no new job work for the client. In other words in particular embodiments, green actions may produce useful work but no computational output for a user or client. Examples of green techniques (implemented by a power balloon or manager to adjust/control power) are diagnostic action and maintenance action on the node or on the system of nodes. At the system level, green activity by the power balloon or manager may include system diagnostics, calibration of unallocated but powered nodes (per node power), maintenance of storage system (compression, virus scan, error detection and correction), maintenance of system memory (error detection and correction), running of low-priority jobs, and so forth. At the node level, green activities initiated by the power balloon or manager may include node diagnostic, calibration of the nodes, maintenance of local storage (compression, virus scan, error detection and correction), maintenance of system memory (error detection and correction), and so forth.

[0021] In contrast, non-green techniques are useless activity, except for the intended purpose of wasting energy and being energy inefficient to increase power

consumption. Non-green activities implemented via a power balloon or manager produce no useful work, and no output or at least no work output, no computational output (or no additional computational output). The power balloon or manager implements the non-green activity to consume energy to maintain the power consumption at the node or system (and thus at the facility) above lower limits or targets of the power band, or below upper limits or targets of power rate of change.

[0022] At the job/node level, non-green actions initiated by the power balloon or manager may include powering off various power features (e.g., features that reduce power and develop energy consumption proportional to computation load and utilization of components), and running various power virus processes, and so forth. At the system level, non-green actions initiated or implemented by the power balloon to control power include powering-on unallocated powered down nodes, turning on a dummy job (per node power), changing energy policy of jobs (to make them run in less energy efficient fashion), and the like.

[0023] In configuration and operation, the power balloon or appropriate manager module will select which power balloon activity to apply. The selection may be based on the amount of power that would be consumed by the activity, the amount of power desired to be consumed at the facility, system, or node, and based on the priority of the activities (e.g., of the useful activities), and other considerations.

[0024] With respect to computing such as distributed computing, high performance computing (HPC) or Big Data analytics, cloud computing, and so on, the cost of energy to a facility or datacenter, as related to electricity or power, may involve both meter charge and demand charge. Meter charge may be proportional to the amount of energy used over a time period. Demand charge may be a function of the specified demand for allocation of power, the amount of power actually used as a percentage of demand, and the fluctuation or variation in the rate of power consumption, and so on.

[0025] Today, a typical specified demand by such a facility or datacenter may be about 1 megawatt (MW) of power, for example. In the future, demand is expected to grow to as much as 45 MW or higher, for instance, especially for facilities or datacenters that host high performance computing (HPC) or super computers, and/or perform Big Data analytics. At such high demand, the demand charge may

be greater than the meter charge. To manage and control demand charge at datacenters including datacenters having high predetermined demand, the datacenters in examples may benefit from the following: (1) when X MW of power is allocated, the power consumption should generally not fall below X minus Y MW, where X minus Y may be labeled as a power band; and (2) the fluctuation or variation in power should be generally limited over time, including with rates such as power per minute and power per hour.

[0026] At present, capacitors, such as a voltage regulator (VR), power supply, and uninterruptable power supply (UPS), provide protection against relatively small fluctuations of power consumption for small durations. Unfortunately, for a UPS or batteries to support multiple MW consumption would generally require large floor space or real estate, and will be prohibitively expensive in many examples. Conventional solutions will generally not be capable to effectively satisfy controlling rate of power consumption at future datacenters, or at datacenters having relatively large power demand and large actual consumption of power.

[0027] In accordance with embodiments of the present techniques, the managing of the rate of power consumption can be generally achieved and implemented by a combination of controlling power consume by computing systems, using energy storage and energy generation, and so on. From an energy/environmental conservation perspective, controls can be generally categorized as green or non-green. In particular, green mechanisms typically facilitate that energy is spent for useful purpose. On the other hand, non-green mechanisms may result in wasted or inefficient expenditure of energy to for the intended purpose of increasing power or slowing a decrease in power. Green and non-green activities can be implemented at facility level, system level, and job/node level. Green activities at the facility level may include using excess energy to: (1) charge batteries to store electrical energy, (2) cool liquid and store the cooled liquid in a reservoir for later use (e.g., to supplement a cooling system), and (3) produce ice to store thermal energy. Facility-level green activities may also include using local generation of energy involving operating renewable energy sources (e.g., solar, wind, geothermal etc.) and/or traditional energy sources such as diesel generators. Facility-level non-green activities may include operating facility racks and the datacenter at low temperature

(such may be analogous wasting energy by running air-conditioning with windows open).

[0028] Some embodiments herein may employ a hierarchical structure to manage the rate of power consumption. The structure may involve a facility power manager, system power-performance manager, job power-performance manager, and so forth. As discussed below, various mechanisms and structures may be implemented to control rate of power consumption at different levels, e.g., within the manager hierarchy.

[0029] As indicated, capacitors at various power delivery levels, e.g., VR, power supply, power distribution units (PDU), UPS with battery, etc., may provide local storage and discharge of energy. At component and system levels, the capacitors may suppress transients at microseconds and milliseconds. In a particular example, for a 1 MW data center, a UPS may support loss of energy for a few minutes. To address upcoming user requirements, such as with 10 MW to 45 MW datacenters, for instance, the extending of traditional approaches may generally be relatively expensive and inadequate. In another particular example, if an HPC job crashes or finishes (completes), the datacenter could experience a 4-6 MW drop in power consumption in milliseconds, for instance. Conventional mechanisms do not protect against or address such an unfortunate and likely scenario. In contrast, as discussed herein, embodiments of the present techniques protect or address sudden and relatively large changes in power consumption at datacenters and facilities for distributed computing including HPC and Big Data analytics.

[0030] Historically, efforts in managing power consumption focused on reducing power consumption without consideration of power band. Conversely, embodiments herein accommodate the need that power draw from a utility company or provider be kept close to the power allocation of the datacenter. Again, some embodiments may employ hierarchical targets and controls to manage or control rate of power consumption at the datacenter or similar facility.

[0031] A facility or datacenter generally hosts multiple compute systems, energy generators, power storage (e.g., batteries), cooling units, etc. In embodiments, the facility may have a facility-level power manager (FPM) to meet user requirements for variations per minute and per hour, and the FPM to maintain facility consumption

within a min-max level or range based upon allocation by the utility supplier entity, for instance. The FPM may generate and provide targets to system level power-performance manager (SPPM) for system level control. The FPM and SPPM may manage variation in power consumption. The SPPM may work with a system of nodes, as well as a job-level power-performance manager or controller, for example, and provide power consumption targets to node-level power performance managers (NPPM). Both the SPPMs and job power-performance managers (JPPM) may comprehend architectures, power usages, and local energy storage to implement various schemes, as discussed below. Thus, again, hierarchical targets and controllers may be implemented, including a hierarchy for controlling to within power band and for controlling rate of power consumption.

[0032] FIG. 1 is an exemplary system 100 for computing, such as a distributed computing system, datacenter, supercomputer(s), Big Data analytics center, cloud computing, enterprise services, high performance computing (HPC) center, and so on. Embodiments herein related to facility power level and system power level are generally applicable to various computing configurations. In the illustrated embodiment, the system 100 has a facility 102 with a computing architecture 104 including one or more systems 106 of nodes 108. In an example, the facility 102 includes three systems 106 each having 20,000 nodes 108. In alternate embodiments, the system level may be fluid in the sense that the groupings of nodes 108 can change over time, such as per job. The nodes 108 may generally be compute nodes but can also include operating system (OS) nodes, input/output (I/O) nodes, and other nodes. Each node 108 may include one or more processors 110, wherein each processor 110 has one or more processing cores 112.

[0033] Each node 108 may also include memory 114 storing code 116 executable by the processors 108. The nodes 108 may include additional hardware, software, firmware, and so on. The computing architecture 104 may also include one or more computing devices 118, such as servers, hosts, etc., each having at least one processor 120 and memory 122 storing code 124 executable by the processor 120.

[0034] The stored executable code 116 and 124 may include various manager or control modules, including managers at different levels in a hierarchy. As mentioned, exemplary managers include node power-performance manager(s)

(NPPM), system power-performance(s) (SPPM), and a facility power manager (FPM), as represented by dashed boxes 126, 128, and 130, respectively, and other modules. In embodiments, job power performance managers (JPPM) (not shown) are distributed among the nodes 108. The portion of that JPPM that runs on an individual node is referred to as the NPPM. Again, these controllers or managers 126, 128, 130 are executable code (e.g., 116 and/or 126) stored in memory (e.g., 114 and/or 124). The managers 126, 128, 130 may have submodules and be a part of other modules. Moreover, as indicated, the managers 126, 128, and 130, and other controller modules, may provide for a hierarchical management or control with respect to power. Further, respective variants of the managers 126, 128, 130, and/or other controllers and managers, may provide for in-band or out-of-band control.

[0035] As can be appreciated, the computing architecture 104 at the facility 102 may consume power. Additionally, the facility 102 may include other consumers 132 of power, such as a building cooling system, lighting, and so on. The facility 102 may receive electricity or power 134 from one or more utility companies or providers 136, for example. The power 134 received may be subjected to conversion at the facility 102 as needed. Moreover, the facility 102 may also have power generation facilities 138 (e.g., diesel generators, etc.), disposed locally and/or remote. Further, the facility 102 may have power storage 140 such as batteries and other energy storage such as a reservoir for coolant. Lastly, the facility 102 may be in communication (e.g., via a network 142) with other distributed computing systems 144 (e.g., datacenters, super computers, Big Data, etc.), or non-distributed or traditional compute systems and computing devices 146 (e.g., at client locations, utility companies or distributors, and so forth).

[0036] FIG. 2 is a diagrammatical representation of an exemplary facility 102 for distributed computing. As mentioned, in some embodiments, a hierarchy may be implemented to manage or control rate of power consumption for the distributed computing system 100 (e.g., datacenter, super computer, etc.) including the facility 102. At the facility level, the facility power manager (FPM) 130 may accommodate requirements for facility power variations per time (e.g., per minute and per hour), and maintain the power consumption at the facility 102 between a minimum power

consumption and a maximum power consumption. Such a range may be based on the allocation of power by the utility company or provider, or by a facility manager (e.g., human administrator).

[0037] The FPM 130 may generate and provide power targets (e.g., for both consumption and variation) to system power-performance manager (SPPM) 128 for the system level power. The FPM 130 and SPPM 128 may manage variation in power consumption. Based on constraints and targets enunciated by the SPPM 128 and/or other factors, the SPPM 128 may manage the power consumption of its respective system 106 of nodes 108.

[0038] FIG. 3 is an example facility 102 for distributed computing. As discussed, the facility manager FPM 130 gives direction to the multiple system managers SPPM 128 for the systems 106 of nodes. Further, the FPM 130 may communicate with and have policies/rules set by a human administrator 300, utility company or provider 302, and so forth. In examples, the FPM 130 may use or be a demand/response interface to the utility provider 302. Moreover, the FPM 130 may perform accountings of power and cooling at the facility 102, and the like, including to communicate capacity and requirements. The FPM 130 may account for a cooling system 132C, manage co-generation of energy (e.g., at generators 138), allocate power budgets to the one or more SPPM 128, and so forth.

[0039] A FPM 130 may use various mechanisms to meet requirements by the datacenter operator with respect to power consumption variation, such as for delta Watts per minute ($\Delta W/\text{min}$) and $\Delta W/\text{hour}$. A datacenter operator may also have the FPM 130 maintain facility-level power consumption at, or slightly below, power allocation by the utility provider 302. With targets to keep energy efficiency high, the FPM may attempt to use green mechanisms before resorting to non-green mechanisms.

[0040] For a facility 102 with local generators 138, the FPM 130 may use local generators 138 for facility control of power. Depending on the particular facility 102, there may be various types of local generators 138. Examples of local generators 138 are diesel engines, solar power, and so forth. In a particular example, if facility 102 demand for power as negotiated with the utility provider 302 is set at say 12 MW, and the facility 102 has local power generation of 2 MW, for instance, then the

facility 102 may instead renegotiate demand for 10 MW from the utility provider 302. If so, the utility provider 302 may require the facility 102 draw from the utility provider between 8.5 MW to 10 MW (e.g., allowed variation of 15%), for example. In this numerical example, the 2 MW local generation may facilitate datacenter facility 102 level fluctuation between 8.5 to 12MW (~30% variation). Local power generation may therefore increase flexibility or level of power, as well as help meet goal(s) for limited variation (e.g., ΔW /hour) over the long term.

[0041] For a facility 102 with local power storage 140, FPM 130 may use the power storage (e.g., battery, local refrigeration, etc.). When facility 102 level actual power consumption may fall below minimum demand (say 8.5 MW in the above numerical example), the FPM 130 may channel energy to charge batteries 140B part of power storage 140. The resulting battery charge may later then be used when facility 102 actual demand for power grows. Another approach for energy storage is to use excess energy (e.g., when actual power consumption of the facility drops below negotiated demand) to cool liquid or to generate ice. Ice or cool liquid can be stored and used later for cooling to reduce energy by the cooling system 132C, for example, when actual power consumption of the facility 102 approaches or exceeds demand, i.e., the upper limit of the power band.

[0042] Other examples of mechanisms the FPM 130 may employ is for the FPM 130 to designate settings on the cooling systems 132C, such as adjusting the temperature or temperature set point of coolant, air, or water in the cooling system 132C. Such control may affect power consumed by both the cooling system 132C and the computing architecture 104 including the compute systems. For example, to increase facility power consumption, the FPM 130 may lower temperature of outlet the outlet air of an air conditioning unit in the cooling system 132C. Although this may be a non-green mechanism, such lowering of the cooling air temperature may help improve life of components (e.g., in the computing architecture 104) cooled by the air. Therefore, this non-green activity while generally useless could have a benefit other than raising power consumption to within power band.

[0043] FIG. 4 is a diagrammatical representation 400 of a system level and node level of an exemplary facility 102 for distributed computing. A Datacenter Manager 406 and a rack controller or manager 402, for example, monitor shared power supply

404 at the system level, and cooling in racks of nodes 108, and so forth. The rack manager 402 may direct node managers 403 (e.g., out-of-band node managers) with respect to monitoring power for the nodes 108, for example. Both in-band and out-of-band management, monitoring, and control may be accommodated. A datacenter manager 406, for example, may part of the out-of-band management hierarchy. A datacenter manager 406 communicates health and capability of power and cooling system to SPPM 128.

[0044] In examples, the node manager 403, rack manager 402 and datacenter manager 406 are out-of-band and not a direct part of system software. These three managers 402, 403 and 406 may instead be part of independent mechanism that monitors power and cooling infrastructure for failures, for instance. In cases of failure in power supply or cooling, the rack manager 402 may take action to reduce or prevent resulting undesired events, may communicate with a job manager, resource manager, SPPM so that they will operate the system and nodes comprehending lower power/cooling capability (due to a failure).

[0045] As discussed with respect to FIGS. 1-3, in comprehending the facility 102 or datacenter level capability and energy cost of the facility 102, e.g., losses in power and cooling systems, building, etc., the FPM 130 may generate target(s) for power amount or level, and for variation in power, for the respective SPC 128 and systems 106 of nodes 108 (see FIGS. 1-3). These targets may be communicated to the SPC 128 for systems 106 power management and control.

[0046] FIGS. 5A and 5B are a flow diagram of an example management hierarchy 500 having exemplary actions 502. FIG. 5A is the top part of the diagram. FIG. 5B is the bottom part of the diagram. Entities in the management hierarchy 500 may include the aforementioned FPM 130, SPPM 128, JPPM 504 and NPPM 126, or variants thereof. In operation, the FPM 130 may manage facility 102 resources and define system power budgets to the one or more SPPM 128. At the system level, a resource manager (not separately depicted) may include, for example, the SPPM 128 and a job power-performance manager (JPPM) 504, and may direct the NPPMs 126 with respect to power consumption (including variation). This power consumption at the system level and the job level may be for the multiple nodes 108 in the systems 106, such as in implementation or execution of a distributed

computing job on the nodes 108. Additional entities in the management hierarchy may include a facility manager (e.g., an automated or human administrator), an energy utility 508 (e.g., utility company or provider), building management 510, and so forth.

[0047] As discussed with respect to the preceding figures, the FPM 130 may generally communicate via a demand/response interface with an energy utility 508, receive and set rules and policies with a facility manager 506, and manage facility 102 energy resources. The FPM 130 may account for facility 102 power and cooling, including to communicate capacity needs and requirements, manage use of local energy, and allocate budgets to the system (i.e., to the resource manager or SPPM 128).

[0048] In the illustrated embodiment, of FIG. 5A, as noted by the actions 502, the FPM 130 may query and receive forecasted system power. The SPPM 128 may develop the system demand forecast. The FPM 130 may also query and receive current and historical system power information from the SPPM 128. Further, the FPM 130 may query and receive the status of the power and cooling infrastructure from building management 510. The FPM 130 may also query and receive from the building management 510 the state of local energy storage and generation.

[0049] Moreover, as noted in FIG. 5A, the facility manager 506 may receive a policy for power ramp and power band, and in turn facility manager 506 may set policies and rules for control of power ramp and power band, to the FPM 130. The FPM 130 may receive from the utility energy 508, utility expectation on power usage and variations, as well as pricing information and incentives with respect to energy. Of course, the reverse direction may be realized, in that the FPM 130 may also provide similar information, and other data and information, to the facility manager 506 and the utility energy 508.

[0050] Continuing with the diagram to FIG. 5B, the FPM 130 based on the various aforementioned information received, and other information, may develop a facility demand forecast, develop targets and plans for local energy storage/generation and set those targets to building management 510, and develop system level targets for power band and ramp and set those limits to the SPPM 128. In response, the SPPM 128 may develop policies for power sequencing of jobs, develop policies for using

“green” and “non-green” power activities, and set job level power band and ramp limits to the JPPM 502.

[0051] As depicted as actions 502 in FIG. 5A, the JPPM 502 may set node level power band and ramp limits (e.g., provided alongside power budgets) to the NPPM 126. The JPPM 502 may also communicate policies for managing power ramps and bands, including green and non-green approaches, to the NPPM 126. The NPPM 126 may maintain power consumption at respective node(s) within band and ramp targets or limits. Lastly, it should be emphasized that other management hierarchies not depicted in FIG. 5, including labels, levels, and configurations, are applicable to the present techniques.

[0052] FIG. 6 is block diagram of exemplary managers and functions at the system level and node level. In a system 106, there may be mechanisms to maintain system 106 power-performance to give relatively high energy efficiency within the system 106 level power budget. The SPPM 128 may work in conjunction with a resource manager 600, for example, to meet requirements for rate of change in power consumption. The resource manager 600 (e.g., executable code directing a processor) may have elements such as job scheduler (not shown), job launcher (not shown), etc., to manage user submitted jobs and launch the jobs in a prioritized order in such a manner to meet a goal, for example, that over 95% of the nodes are used for computations. The resource manager (RM) 600 may have components, such as a system power-performance manager (SPPM) 128 and a job power-performance manager (JPPM) 604 (which may be analogous to 502 of FIGS. 5A and 5B), to manage power-performance of the job. The RM 600 by selecting various types of jobs and controlling gradual increase in job power can help meet the system 106 level target for power consumption.

[0053] As for selection of jobs, the RM 600 may select jobs 606 from queues and schedule the jobs 606 for execution. In embodiments, the RM 600 can access tools that estimate power needs for jobs. These estimates may help the RM 600 to select combinations of jobs that may run to meet system 106 level power requirements for power control.

[0054] The RM 600 may address power sequencing 608 for a job. In a particular example, if a job starts say on 20,000 nodes 108, there may be thus a sudden

increase in power consumption at the system 106 and the facility 102. In certain embodiments, a RM 600 can employ multiple techniques to make this power increase more gradual. For example, the RM 600 may coordinate with the FPM 130 to address systematic power and to make power consumption gradual for the system 106. The FPM 130 can dampen a sudden or stark increase in power consumption by using local energy storage 140 and local energy generation 138. In another example to make a system power increase more gradual, the RM 600 may control performance of compute nodes 108 that run the job in such a way that the nodes start the job at a low (e.g., lowest) frequency. The RM 600 may then gradually in a controlled manner increase the frequency of the nodes 108 executing the job.

[0055] As for non-critical jobs 610, when a previous job finishes, the RM 600 may designate running or execution of non-critical but useful jobs. Examples of such jobs 610 may be a system health check, diagnostics, security scan, power-performance calibration of nodes 108, defragmentation of storage devices, scanning and correction of errors in system memory and storage drives, and so forth. These may be green mechanisms that improve reliability of the systems 106 and other facility 102 systems.

[0056] A power control job 612 may be a job that generates little or no useful work and consumes energy to maintain power level over minimum. These may clearly be a non-green option. Nevertheless, such power control jobs 612 may help maintain system 106 level power consumption within target. In certain embodiments, these kinds of activities may be implemented via a system power balloon application 614 (e.g., executable code), as discussed in more detail below.

[0057] As for power consumption targets for nodes 108, the RM 600 may control power-performance of the node 106 through a job or workload manager, such a JPPM 604. The RM 600 via a JPPM 604, for example, can provide specific targets in certain embodiments so that node 108 level power consumption does not change drastically and suddenly. Again, to meet ramp or band targets, certain activities directed to avoiding sudden changes in power rate (and also to managing power consumption within band), may be implemented via a power balloon technique or mechanism (module or executable code stored in memory). In certain embodiments, there may be a power balloon at both the system level and the node level. The

illustrated embodiment depicts a system power balloon 614 and node power balloon 622.

[0058] To maintain node 108 level power consumption within target provided by the job manager or JPPM 604, the control of power management (PM) features, as noted in block 618, may be addressed. A reduction in node 108 level power may be the result of power management (PM) or power reduction technologies in a node or a platform. These technologies may reduce component level power when the component's usage or computation load decreases. By disabling PM features in some embodiments, the component level power may remain steady or substantially steady. However, the RM 600 and/or node power balloon 622 may direct the PM for a more gradual approach to reduce power slowly. A more gradual approach may help meet goals for rate of change in power and also may be energy efficient.

[0059] Furthermore, non-critical applications 620 may exist for maintenance. In examples with a power drop at a node 108, the RM 600 via the node power balloon 622 may run non-critical but useful applications 620. Although such non-critical applications 620 may not generate computer output for any user, such applications 620 may perform maintenance and therefore improve reliability of the node 108. Examples of such applications include node health check, diagnostics, security scan, calibration of nodes, defragmentation of storage drives, scanning and correction of errors in system memory and storage drives, and the like.

[0060] The SPPM 128, for instance, may implement a present technique called power balloon (see, e.g., FIG. 7). A power balloon (614 and 622) may manage green and non-green applications to (1) maintain system power within a power band and (2) manage system level power fluctuations, and other results. Examples of green applications, as alluded above as non-critical applications 620, may be system health check, diagnostics, security scan, calibration of nodes, defragmentation of storage drives, scanning and correction of errors in system memory and storage drives, and so forth. Examples of non-green applications, such as those implemented by a power balloon application, may be powering on unallocated nodes (unallocated nodes are generally powered down to save energy), running dummy jobs such as a power virus, asking a job manager (or JPPM 502 or 604) to run jobs without being energy efficient, and so on. As mentioned and also depicted in FIG. 6,

a power balloon 614, 622 may operate on the system level and the node level, respectively. As for the node level, a node power balloon 622 may provide for a local technique at a node 108. In particular examples, the node power balloon 622 may be a non-green technique in that the node may not generate useful compute output in response to the power balloon activity. The node power balloon 622 may also implement green actions. Again, the same or similar actions (green and non-green) may be implemented on the system level via a system power balloon 614.

[0061] FIG. 7 illustrates how a node power balloon 622 can be used to beneficially maintain node 108 power within a band. In some embodiments, this mechanism may be responsive even when there is power consumption change in milliseconds, for instance. The power balloon 622 can be invoked by an operating system, or by an out-of-band mechanism or firmware. In the illustrated embodiment, an out-of-band node manager 616 may invoke the power balloon 622. In general, a power balloon either increases or decrease power to meet requirements for power ramp and for power band control. Further, a power balloon may be implemented at a node level or a system level. In the particular illustrated embodiment of FIG 7, the power balloon is implemented at the node level and increases power to meet power band. However, power balloon implementation may be equally or similarly applicable or adopted at the system level, and also to adjust or control power ramp, i.e., rate of power change (increase or decrease).

[0062] FIG. 7 is a bar chart 700 comparing power consumption of two example compute jobs 702 and 704, or the same compute job at different times, executing on a compute node 108 in a distributed computing system 106. The desired range or band for the node 108 power is from a minimum power 708 to a maximum power 710. The first compute job 702 executes within the desired band. In contrast, the second compute job 704 executes below the desired minimum power 708 for the node 108. However, with implementation of the aforementioned power balloon application 622, the power of node 108 when executing the second job 704 may be increased to at or above the minimum power 708 level, as indicated by the added power 706, to within the desired power band.

[0063] In sum, a node-level power balloon 622 can implement green and non-green actions to manage power band and power fluctuations. Green actions or

methods may include running node calibration, diagnostic, memory or disc scans to detect and correct errors, virus scan, etc. Non-green actions may include disabling or turning off power reduction features, running a dummy program (e.g., power virus) that consumes power but generates no work output, and so forth. Again, a power balloon implementing the aforementioned actions giving the consequence at the node level depicted in FIG. 7, can also be implemented on the system level (see system power balloon 614 of FIG. 6).

[0064] FIG. 8 is a block diagram depicting an example of a tangible non-transitory, computer-readable medium that can facilitate power management of a distributed computing system in accordance with embodiments of the present techniques. The computer-readable medium 800 may be accessed by a processor 802 over a computer interconnect 804. The processor 802 may be one or more compute node processors (e.g., 110), a server processor (e.g., 110, 120), or another processor. The tangible, non-transitory, computer-readable medium 800 may include executable instructions or code to direct the processor 802 to perform the operations of the techniques described herein.

[0065] The various software components discussed herein may be stored on the tangible, non-transitory, computer-readable medium 800, as indicated in FIG. 8. For example, a power management/control module 806 (executable code/instructions) may direct the processor 802 to manage and control power consumption and variation at a distributed computing facility. Indeed, the module 806 executable instructions or code may include a FPM, resource manager, SPPM, JPPM, NPPM, system power balloon, node power balloon, and so on. The code or module 806 may be similar or the same as the aforementioned executable code 116, 124 in FIG. 1. It should be understood that any number of additional software components not shown in FIG. 8 may be included within the tangible, non-transitory, computer-readable medium 800, depending on the application.

[0066] Some embodiments may be implemented in one or a combination of hardware, firmware, and software. Some embodiments may also be implemented as instructions stored on a machine-readable medium, which may be read and executed by a computing platform to perform the operations described herein. A machine-readable medium may include any mechanism for storing or transmitting

information in a form readable by a machine, e.g., a computer. For example, a machine-readable medium may include read only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; or electrical, optical, acoustical or other form of propagated signals, e.g., carrier waves, infrared signals, digital signals, or the interfaces that transmit and/or receive signals, among others.

[0067] An embodiment is an implementation or example. Reference in the specification to “an embodiment”, “one embodiment”, “some embodiments”, “various embodiments”, or “other embodiments” means that a particular feature, structure, or characteristic described in connection with the embodiments is included in at least some embodiments, but not necessarily all embodiments, of the present techniques. The various appearances of “an embodiment”, “one embodiment”, or “some embodiments” are not necessarily all referring to the same embodiments. Elements or aspects from an embodiment can be combined with elements or aspects of another embodiment.

[0068] Not all components, features, structures, characteristics, etc. described and illustrated herein need be included in a particular embodiment or embodiments. If the specification states a component, feature, structure, or characteristic “may”, “might”, “can” or “could” be included, for example, that particular component, feature, structure, or characteristic is not required to be included. If the specification or claim refers to “a” or “an” element, that does not mean there is only one of the element. If the specification or claims refer to “an additional” element, that does not preclude there being more than one of the additional element.

[0069] It is to be noted that, although some embodiments have been described in reference to particular implementations, other implementations are possible according to some embodiments. Additionally, the arrangement and/or order of circuit elements or other features illustrated in the drawings and/or described herein need not be arranged in the particular way illustrated and described. Many other arrangements are possible according to some embodiments.

[0070] In each system shown in a figure, the elements in some cases may each have a same reference number or a different reference number to suggest that the elements represented could be different and/or similar. However, an element may

be flexible enough to have different implementations and work with some or all of the systems shown or described herein. The various elements shown in the figures may be the same or different. Which one is referred to as a first element and which is called a second element is arbitrary.

[0071] Examples are provided. Example 1 is a facility for computing, the facility having systems of multiple compute nodes to execute jobs of computing. The computing facility has memory storing code, wherein the code is executable by the multiple compute nodes or by a processor to manage power consumption of the facility to within a power band, and wherein to manage power consumption comprises to adjust power consumption via a power balloon. The executable code comprises the power balloon.

[0072] Example 2 incorporates the subject matter of Example 1. In this example, to adjust power consumption, via the power balloon, controls rate of the power consumption.

[0073] Example 3 incorporates the subject matter of any combination of Examples 1-2. In this example, the computing comprises distributed computing and the job comprises a distributed computing job. Further, to manage power consumption includes managing power consumption at a facility level, a system level, and a node level, and wherein the power band comprises a range from a minimum power consumption to a maximum power consumption.

[0074] Example 4 incorporates the subject matter of any combination of Examples 1-3. In this example, the power balloon is a node power balloon.

[0075] Example 5 incorporates the subject matter of any combination of Examples 1-4. In this example, the power balloon is a system power balloon.

[0076] Example 6 incorporates the subject matter of any combination of Examples 1-5. In this example, to adjust power consumption via a power balloon comprises to implement a green technique.

[0077] Example 7 incorporates the subject matter of any combination of Examples 1-6. In this example, to adjust power consumption via a power balloon comprises to implement an activity that produces useful work but no computational output for a client.

- [0078]** Example 8 incorporates the subject matter of any combination of Examples 1-7. In this example, to adjust power consumption via a power balloon comprises to implement a non-green technique.
- [0079]** Example 9 incorporates the subject matter of any combination of Examples 1-8. In this example, to adjust power consumption via a power balloon comprises to implement an activity that produces no useful work and no additional computational output, the activity to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.
- [0080]** Example 10 incorporates the subject matter of any combination of Examples 1-9. In this example, to adjust power consumption via a power balloon comprises to run a diagnostic or perform maintenance, or a combination thereof.
- [0081]** Example 11 incorporates the subject matter of any combination of Examples 1-10. In this example, to adjust power consumption via a power balloon comprises to power-on an unallocated powered-down node.
- [0082]** Example 12 incorporates the subject matter of any combination of Examples 1-11. In this example, to adjust power consumption via a power balloon comprises to run a process of power thermal utility to increase energy consumption of a compute node.
- [0083]** Example 13 is a facility for computing, having systems of compute nodes, each system comprising a plurality of the compute nodes to execute jobs of distributed computing. The facility for computing includes memory storing executable code to manage power consumption at the facility, wherein to manage power consumption comprises to control the power consumption to within a power band and to control rate of the power consumption at the facility, by initiating an activity to adjust the power consumption, the activity giving no computational output.
- [0084]** Example 14 incorporates the subject matter of Example 13. In this example, the activity acts on a system of compute nodes or on a compute node, or both, and wherein the facility comprises a cloud computing facility, a distributed computing facility, a high performance computing (HPC) facility, a Big Data analytics facility, or any combination thereof.

- [0085]** Example 15 incorporates the subject matter of any combination of Examples 13-14. In this example, executable code comprises a power balloon to initiate the activity.
- [0086]** Example 16 is a method of computing, including executing jobs of computing on systems of computing at a facility, wherein each system comprises multiple compute nodes. The method includes managing, via a processor, power consumption at the facility to within a power band via a power balloon comprising implementing green activities and non-green activities.
- [0087]** Example 17 incorporates the subject matter of Example 16. In this example, managing power consumption comprises controlling rate of power consumption.
- [0088]** Example 18 incorporates the subject matter of any combination of Examples 16-17. In this example, the power balloon is at a node level.
- [0089]** Example 19 incorporates the subject matter of any combination of Examples 16-18. In this example, the power balloon is at a system level.
- [0090]** Example 20 incorporates the subject matter of any combination of Examples 16-19. In this example, managing power consumption via a power balloon comprises implementing an activity that produces useful work but no computational output for a client.
- [0091]** Example 21 incorporates the subject matter of any combination of Examples 16-20. In this example, managing power consumption via a power balloon comprises implementing an activity that produces no useful work and no additional computational output, wherein the activity is implemented to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.
- [0092]** Example 22 incorporates the subject matter of any combination of Examples 16-21. In this example, the computing comprises distributed computing, the system comprises a distributed computing system, and the facility comprises a distributed computing facility.
- [0093]** Example 23 is a non-transitory, computer-readable medium comprising instructions executable by a processor to manage, via a power balloon implementing green techniques and non-green techniques, power consumption at a facility for

distributed computing to within a power band, the facility comprising systems of multiple compute nodes, and wherein to manage power consumption comprises to control rate of power consumption at a facility level, a system level, and a compute node level. The executable instructions include the power balloon.

[0094] Example 24 incorporates the subject matter of Example 23. In this example, to manage power consumption via the power balloon comprises to implement an activity on a system of compute nodes at the facility to produce useful work but no computational output for a user, and wherein the power band comprises a range from a minimum power consumption to a maximum power consumption.

[0095] Example 25 incorporates the subject matter of any combination of Example 23-24. In this example, to manage power consumption via a power balloon comprises to initiate an activity on a system of compute nodes at the facility to produce no useful work and no additional computational output, wherein the activity is implemented to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.

[0096] Example 26 is an apparatus for managing power consumption at a facility for distributed computing, comprising means for managing, via a power balloon implementing green techniques and non-green techniques, power consumption of the facility to within a power band, the facility comprising systems of multiple compute nodes, and wherein to manage power consumption comprises to control rate of power consumption at a facility level, a system level, and a compute node level.

[0097] Example 27 incorporates the subject matter of Example 26. In this example, the means for managing power consumption via a power balloon comprises means for implementing an activity on a system of compute nodes at the facility to produce useful work but no computational output for a user, and wherein the power band comprises a range from a minimum power consumption to a maximum power consumption.

[0098] Example 28 incorporates the subject matter of any combination of Examples 26-27. In this example, the means for managing power consumption via a power balloon comprises means for initiating an activity on a system of compute nodes at the facility to produce no useful work and no additional computational

output, wherein the activity is implemented to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.

[0099] Example 29 is a facility for computing, the facility having systems of multiple compute nodes to execute jobs of computing. The computing facility having memory storing executable code to manage power consumption of the facility to within a power band, wherein to manage power consumption comprises to control rate of power consumption, and wherein to manage power consumption comprises to adjust power consumption via a power balloon.

[0100] Example 30 incorporates the subject matter of Example 29. In this example, the computing comprises distributed computing and the job comprises a distributed computing job, wherein to manage power consumption comprises to manage power consumption at a facility level, a system level, and a node level, and wherein the power band comprises a range from a minimum power consumption to a maximum power consumption.

[0101] Example 31 incorporates the subject matter of any combination of Examples 29-30. In this example, the power balloon comprises a node power balloon.

[0102] Example 32 incorporates the subject matter of any combination of Examples 29-31. In this example, the power balloon comprises a system power balloon.

[0103] Example 33 incorporates the subject matter of any combination of Examples 29-32. In this example, to adjust power consumption via a power balloon comprises to implement a green technique comprising to implement an activity that produces useful work but no computational output for a client.

[0104] Example 34 incorporates the subject matter of any combination of Examples 29-33. In this example, to adjust power consumption via a power balloon comprises to implement a non-green technique comprising to implement an activity that produces no useful work and no additional computational output, the activity to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.

[0105] Example 35 incorporates the subject matter of any combination of Examples 29-34. In this example, to adjust power consumption via a power balloon comprises to run a diagnostic or perform maintenance, or a combination thereof.

[0106] Example 36 incorporates the subject matter of any combination of Examples 29-35. In this example, to adjust power consumption via a power balloon comprises to run a process of power thermal utility to increase energy consumption of a compute node.

[0107] Example 37 is a method of computing, including executing jobs of computing on systems of computing at a facility, wherein each system comprises multiple compute nodes. The method includes managing, via a power balloon, power consumption at the facility to control rate of power consumption and to control power within a power band, the power balloon implementing green activities and non-green activities.

[0108] Example 38 incorporates the subject matter of Example 37. In this example, managing power consumption via the power balloon comprises implementing an activity that produces useful work but no computational output for a client.

[0109] Example 39 incorporates the subject matter of any combination of Examples 37-38. In this example, in managing power consumption via the power balloon comprises implementing an activity that produces no useful work and no additional computational output, wherein the activity is implemented to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.

[0110] Example 40 incorporates the subject matter of any combination of Examples 37-39. In this example, managing power consumption via the power balloon comprises to power-on an unallocated powered-down node.

[0111] Example 41 is a non-transitory, computer-readable medium comprising instructions, including a power balloon, executable by a processor to manage power consumption of a facility to control rate of power consumption of the facility and to control power consumption of the facility to within a power band, the power balloon implementing green activities and non-green activities, and wherein the facility comprises a distributed computing facility.

[0112] Example 42 incorporates the subject matter of Example 41. In this example, to manage power consumption via the power balloon comprises implementing an activity that produces useful work but no computational output for a client.

[0113] Example 43 incorporates the subject matter of any combination of Examples 41-42. In this example, to manage power consumption via the power balloon comprises implementing an activity that produces no useful work and no additional computational output, wherein the activity is implemented to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.

[0114] Example 44 incorporates the subject matter of any combination of Examples 41-43. In this example, to manage power consumption via the power balloon comprises to power-on an unallocated powered-down node.

[0115] Example 45 incorporates the subject matter of any combination of Examples 41-44. In this example, the power balloon is configured for a node level.

[0116] Example 46 incorporates the subject matter of any combination of Examples 41-45. In this example, the power balloon is configured for a system level.

[0117] Example 47 incorporates the subject matter of any combination of Examples 41-46. In this example, to manage power consumption via the power balloon comprises implementing an activity that produces no useful work and no additional computational output.

[0118] Example 48 incorporates the subject matter of any combination of Examples 41-47. In this example, the distributed computing facility comprises high performance computing (HPC) systems comprising compute nodes.

[0119] Example 49 incorporates the subject matter of any combination of Examples 41-48. In this example, to manage power consumption via the power balloon comprises to implement a green technique.

[0120] Example 50 incorporates the subject matter of any combination of Examples 41-49. In this example, to adjust power consumption via the power balloon comprises to implement a non-green technique.

[0121] Example 51 is a facility for computing, the facility having systems of multiple compute nodes configured to execute jobs of computing, and memory

storing code executable by the multiple compute nodes or by a processor to manage power consumption of the facility to within a power band.

[0122] Example 52 incorporates the subject matter of Example 51. In this example, the computing comprises distributed computing and the job comprises a distributed computing job, wherein to manage power consumption comprises to manage power consumption at a facility level, a system level, and a node level, and to utilize local energy storage or local energy generation, or both.

[0123] Example 53 incorporates the subject matter of any combination of Examples 51-52. In this example, the code includes a power balloon, and wherein to manage the power consumption comprises to adjust the power consumption via the power balloon.

[0124] Example 54 incorporates the subject matter of any combination of Examples 51-53. In this example, to adjust power consumption, via the power balloon, controls rate of the power consumption.

[0125] Example 55 incorporates the subject matter of any combination of Examples 51-54. In this example, the power balloon comprises a node power balloon.

[0126] Example 56 incorporates the subject matter of any combination of Examples 51-54. In this example, the power balloon comprises a system power balloon.

[0127] Example 57 incorporates the subject matter of any combination of Examples 51-56. In this example, to adjust power consumption via the power balloon comprises to implement a green technique.

[0128] Example 58 incorporates the subject matter of any combination of Examples 51-57. In this example, to adjust power consumption via the power balloon comprises to implement an activity that produces useful work but no computational output for a client.

[0129] Example 59 incorporates the subject matter of any combination of Examples 51-58. In this example, to adjust power consumption via the power balloon comprises to implement a non-green technique.

[0130] Example 60 incorporates the subject matter of any combination of Examples 51-59. In this example, to adjust power consumption via the power

balloon comprises to implement an activity that produces no useful work and no additional computational output, the activity to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.

[0131] Example 61 incorporates the subject matter of any combination of Examples 51-60. In this example, to adjust power consumption via the power balloon comprises to run a diagnostic or perform maintenance, or a combination thereof.

[0132] Example 62 incorporates the subject matter of any combination of Examples 51-61. In this example, to adjust power consumption via the power balloon comprises to power-on an unallocated powered-down node.

[0133] Example 63 incorporates the subject matter of any combination of Examples 51-62. In this example, to adjust power consumption via the power balloon comprises to run a process of a power thermal utility to increase energy consumption of a compute node.

[0134] Example 64 incorporates the subject matter of any combination of Examples 51-63. In this example, to manage the power consumption comprises to control rate of the power consumption.

[0135] Example 65 incorporates the subject matter of any combination of Examples 51-64. In this example, to control rate of power consumption comprises to adjust operating frequency of the compute nodes to affect power sequencing in order to reduce initial power consumption caused by start of a job.

[0136] Example 66 incorporates the subject matter of any combination of Examples 51-65. In this example, to manage power consumption comprises to utilize local energy storage.

[0137] Example 67 incorporates the subject matter of any combination of Examples 51-66. In this example, to manage power consumption comprises to utilize local energy generation.

[0138] Example 68 is a method of computing including executing jobs of computing on systems of computing at a facility, wherein each system has multiple compute nodes. The method includes managing, via a processor, power consumption of the facility to within a power band and to control rate of power consumption.

[0139] Example 69 incorporates the subject matter of Example 68. In this example, managing power consumption comprises managing power consumption via a power balloon.

[0140] Example 70 incorporates the subject matter of Example 69. In this example, the power balloon is at a system level or a node level, or both, and managing the power consumption via the power balloon comprises implementing a green activity, which produces useful work but no computational output for a client.

[0141] Example 71 incorporates the subject matter of any combination of Examples 69-70. In this example, managing the power consumption via the power balloon comprises implementing a non-green activity on a system of multiple compute nodes and that produces no useful work and no additional computational output, wherein the activity is implemented to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.

[0142] Example 72 incorporates the subject matter of any combination of Examples 68-71. In this example, managing power consumption comprises controlling rate of power consumption at a facility level, a system level, job level, and a compute node level.

[0143] Example 73 incorporates the subject matter of any combination of Examples 68-71. In this example, managing power consumption comprises controlling rate of power consumption at a job level.

[0144] Example 74 incorporates the subject matter of any combination of Examples 68-73. In this example, the systems comprise high performance computing (HPC) systems.

[0145] Example 75 incorporates the subject matter of any combination of Examples 68-74. In this example, the power balloon is at a system level or a node level, or both.

[0146] Example 76 is a facility for computing, the facility comprising: systems of multiple compute nodes configured to execute jobs of distributed computing; and memory storing code executable by the multiple compute nodes or by a processor to manage power consumption of the facility to within a power band, and to control rate of the power consumption.

[0147] Example 77 incorporates the subject matter of Examples 76. In this example, to manage power consumption comprises to manage power consumption at a facility level, a system level, and a node level, and to utilize local power storage or local power generation, or both.

[0148] Example 78 incorporates the subject matter of any combination of Examples 76-77. In this example, the code comprises a power balloon, and wherein to manage power consumption comprises to adjust power consumption via the power balloon implementing a green technique that produces useful work but no computational output for a client.

[0149] Example 79 incorporates the subject matter of any combination of Examples 76-78. In this example, the code comprises a power balloon, and wherein to manage power consumption comprises to adjust power consumption via the power balloon implementing a non-green activity that produces no useful work and no additional computational output, the non-green activity to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.

[0150] Example 80 incorporates the subject matter of any combination of Examples 76-79. In this example, the non-green activity comprises to power-on an unallocated powered-down node.

[0151] Example 81 incorporates the subject matter of any combination of Examples 76-80. In this example, to manage power consumption comprises to adjust power consumption via a power thermal utility to increase energy consumption of a compute node.

[0152] Example 82 incorporates the subject matter of any combination of Examples 76-81. In this example, to control rate of power consumption comprises to adjust operating frequency of the compute nodes to affect power sequencing in order to reduce initial power consumption caused by start of a distributed computing job.

[0153] Example 83 is a non-transitory, computer-readable medium comprising instructions executable by a processor to manage power consumption by a facility to within a power band and to control rate of power consumption by the facility, the facility having distributed computing systems having multiple compute nodes.

[0154] Example 84 incorporates the subject matter of Example 83. In this example, to manage power consumption comprises to manage power consumption at a facility level, a system level, and a compute node level, and wherein the power band comprises a range from a specified minimum power consumption to a specified maximum power consumption.

[0155] Example 85 incorporates the subject matter of any combination of Examples 83-84. In this example, the instructions executable by the processor comprise a power balloon, and wherein to manage power consumption comprises to adjust power consumption via the power balloon implementing a green activity that produces useful work but no computational output for a user.

[0156] Example 86 incorporates the subject matter of any combination of Examples 83-85. In this example, instructions executable by the processor comprise a power balloon, and wherein to manage power consumption comprises to adjust power consumption via the power balloon implementing a non-green activity that produces no useful work and no additional computational output, wherein the activity is implemented to consume energy to maintain the power consumption at the facility above a specified minimum limit of the power band.

[0157] It is to be understood that specifics in the aforementioned examples may be used anywhere in one or more embodiments. For instance, all optional features of the computing device described above may also be implemented with respect to either of the methods described herein or a computer-readable medium. Furthermore, although flow diagrams and/or state diagrams may have been used herein to describe embodiments, the present techniques are not limited to those diagrams or to corresponding descriptions herein. For example, flow need not move through each illustrated box or state or in exactly the same order as illustrated and described herein.

[0158] The present techniques are not restricted to the particular details listed herein. Indeed, those skilled in the art having the benefit of this disclosure will appreciate that many other variations from the foregoing description and drawings may be made within the scope of the present techniques. Accordingly, it is the following claims including any amendments thereto that define the scope of the present techniques.

CLAIMS

What is claimed is:

1. A facility for computing, the facility comprising:
systems of multiple compute nodes configured to execute jobs of computing;
and
memory storing code executable by the multiple compute nodes or by a processor to manage power consumption of the facility to within a power band.
2. The facility of claim 1, wherein the computing comprises distributed computing and the job comprises a distributed computing job, wherein to manage power consumption comprises to manage power consumption at a facility level, a system level, and a node level, and to utilize local power storage or local power generation, or both.
3. The facility of claim 1, wherein the code comprises a power balloon, and wherein to manage the power consumption comprises to adjust the power consumption via the power balloon.
4. The facility of claim 3, wherein to adjust power consumption, via the power balloon, controls rate of the power consumption.
5. The facility of claims 3 and 4, wherein the power balloon comprises a node power balloon.
6. The facility of claims 3 and 4, wherein the power balloon comprises a system power balloon.
7. The facility of claims 3 and 4, wherein to adjust power consumption via the power balloon comprises to implement a green technique.

8. The facility of claims 3 and 4, wherein to adjust power consumption via the power balloon comprises to implement an activity that produces useful work but no computational output for a client.

9. The facility of claims 3 and 4, wherein to adjust power consumption via the power balloon comprises to implement a non-green technique.

10. The facility of claims 3 and 4, wherein to adjust power consumption via the power balloon comprises to implement an activity that produces no useful work and no additional computational output, the activity to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.

11. The facility of claims 3 and 4, wherein to adjust power consumption via the power balloon comprises to run a diagnostic or perform maintenance, or a combination thereof.

12. The facility of claims 3 and 4, wherein to adjust power consumption via the power balloon comprises to power-on an unallocated powered-down node.

13. The facility of claims 3 and 4, wherein to adjust power consumption via the power balloon comprises to run a process of a power thermal utility to increase energy consumption of a compute node.

14. The facility of claims 1 and 2, wherein to manage the power consumption comprises to control rate of the power consumption.

15. The facility of claim 14, wherein to control rate of power consumption comprises to adjust operating frequency of the compute nodes to affect power sequencing in order to reduce initial power consumption caused by start of a job.

16. A facility for computing, comprising:
systems of compute nodes, each system comprising a plurality of the compute nodes to execute jobs of distributed computing; and
memory storing executable code to manage power consumption at the facility, wherein to manage power consumption comprises to control power consumption of the facility to within a power band and to control rate of the power consumption of the facility, by initiating an activity to adjust the power consumption, the activity giving no computational output.

17. The facility of claim 16, wherein the activity acts on a system of compute nodes or on a compute node, or both, and wherein the facility comprises a cloud computing facility, a distributed computing facility, a high performance computing (HPC) facility, a Big Data analytics facility, or any combination thereof.

18. The facility of claims 16 and 17, wherein executable code comprises a power balloon to initiate the activity.

19. A method of computing, comprising:
executing jobs of computing on systems of computing at a facility, wherein each system comprises multiple compute nodes; and
managing, via a processor, power consumption of the facility to within a power band and to control rate of power consumption.

20. The method of claim 1, wherein managing power consumption comprises managing power consumption via a power balloon.


21. The method of claim 20, wherein the power balloon is at a system level or a node level, or both, and managing the power consumption via the power balloon comprises implementing a green activity, which produces useful work but no computational output for a client.

22. The method of claim 20, wherein managing the power consumption via the power balloon comprises implementing a non-green activity, which produces no useful work and no additional computational output, wherein the activity is implemented to consume energy to maintain the power consumption at the facility above a minimum limit of the power band.


23. A non-transitory, computer-readable medium comprising instructions executable by a processor to manage power consumption by a facility to within a power band and to control rate of power consumption by the facility, the facility having distributed computing systems having multiple compute nodes.

24. The non-transitory, computer-readable medium of claim 23, wherein the instructions executable by the processor comprise a power balloon, and wherein to manage power consumption comprises to adjust power consumption via the power balloon implementing a green activity that produces useful work but no computational output for a user.

25. The non-transitory, computer-readable medium of claims 23 and 24, wherein the instructions executable by the processor comprise a power balloon, and wherein to manage power consumption comprises to adjust power consumption via the power balloon implementing a non-green activity that produces no useful work and no additional computational output, wherein the activity is implemented to consume energy to maintain the power consumption at the facility above a specified minimum limit of the power band.



100
FIG. 1



200
FIG. 2

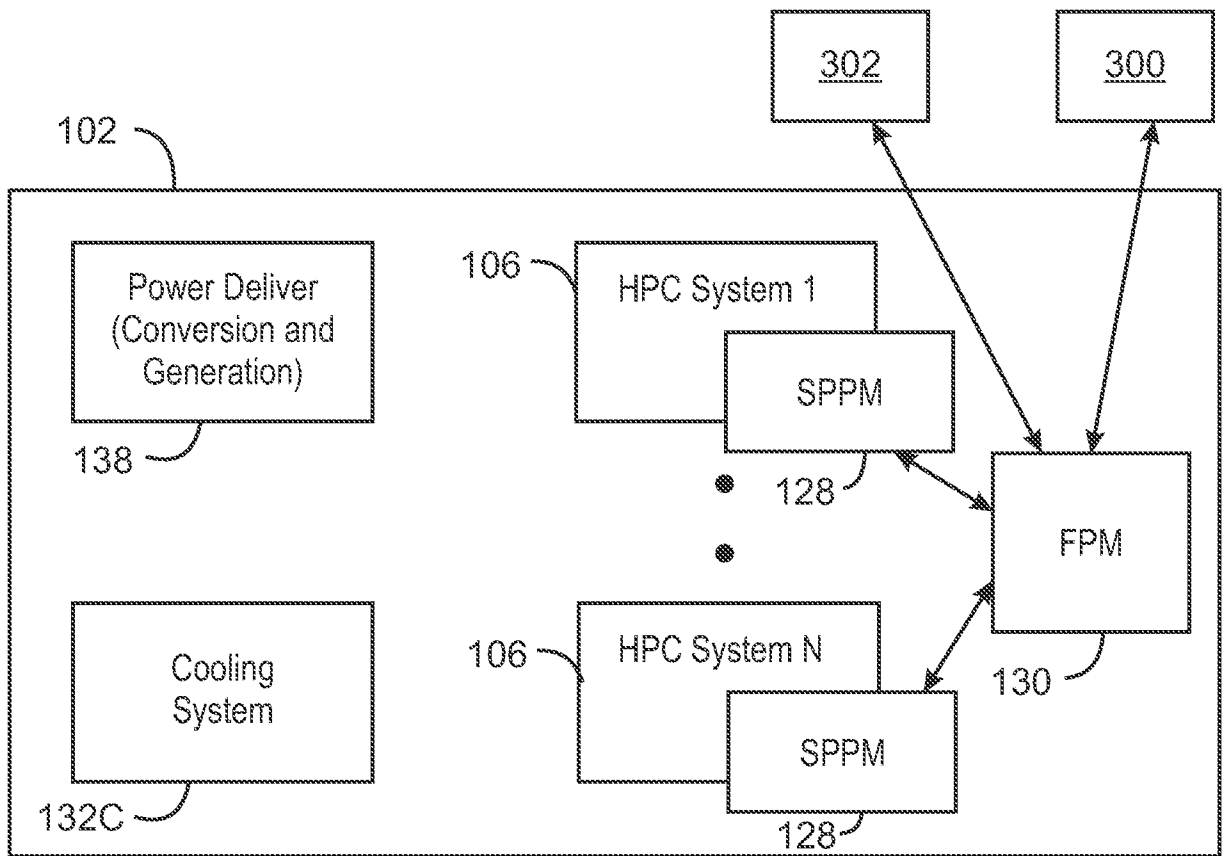
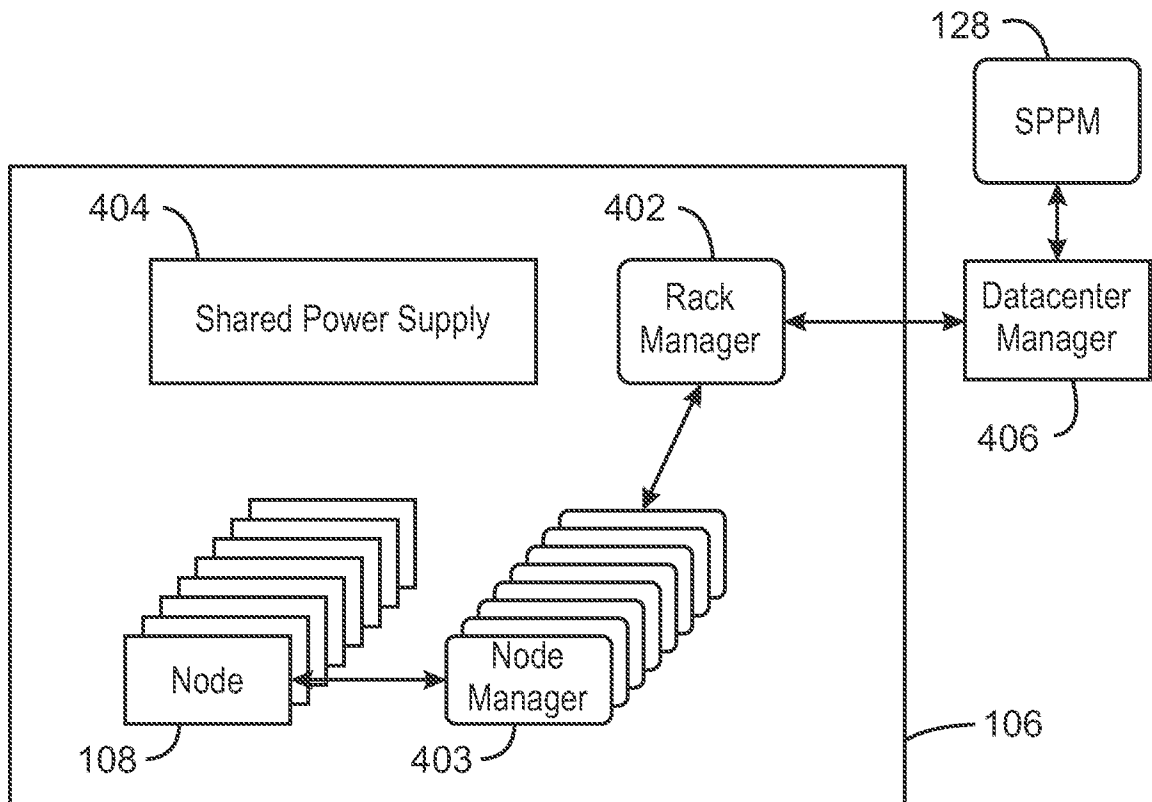
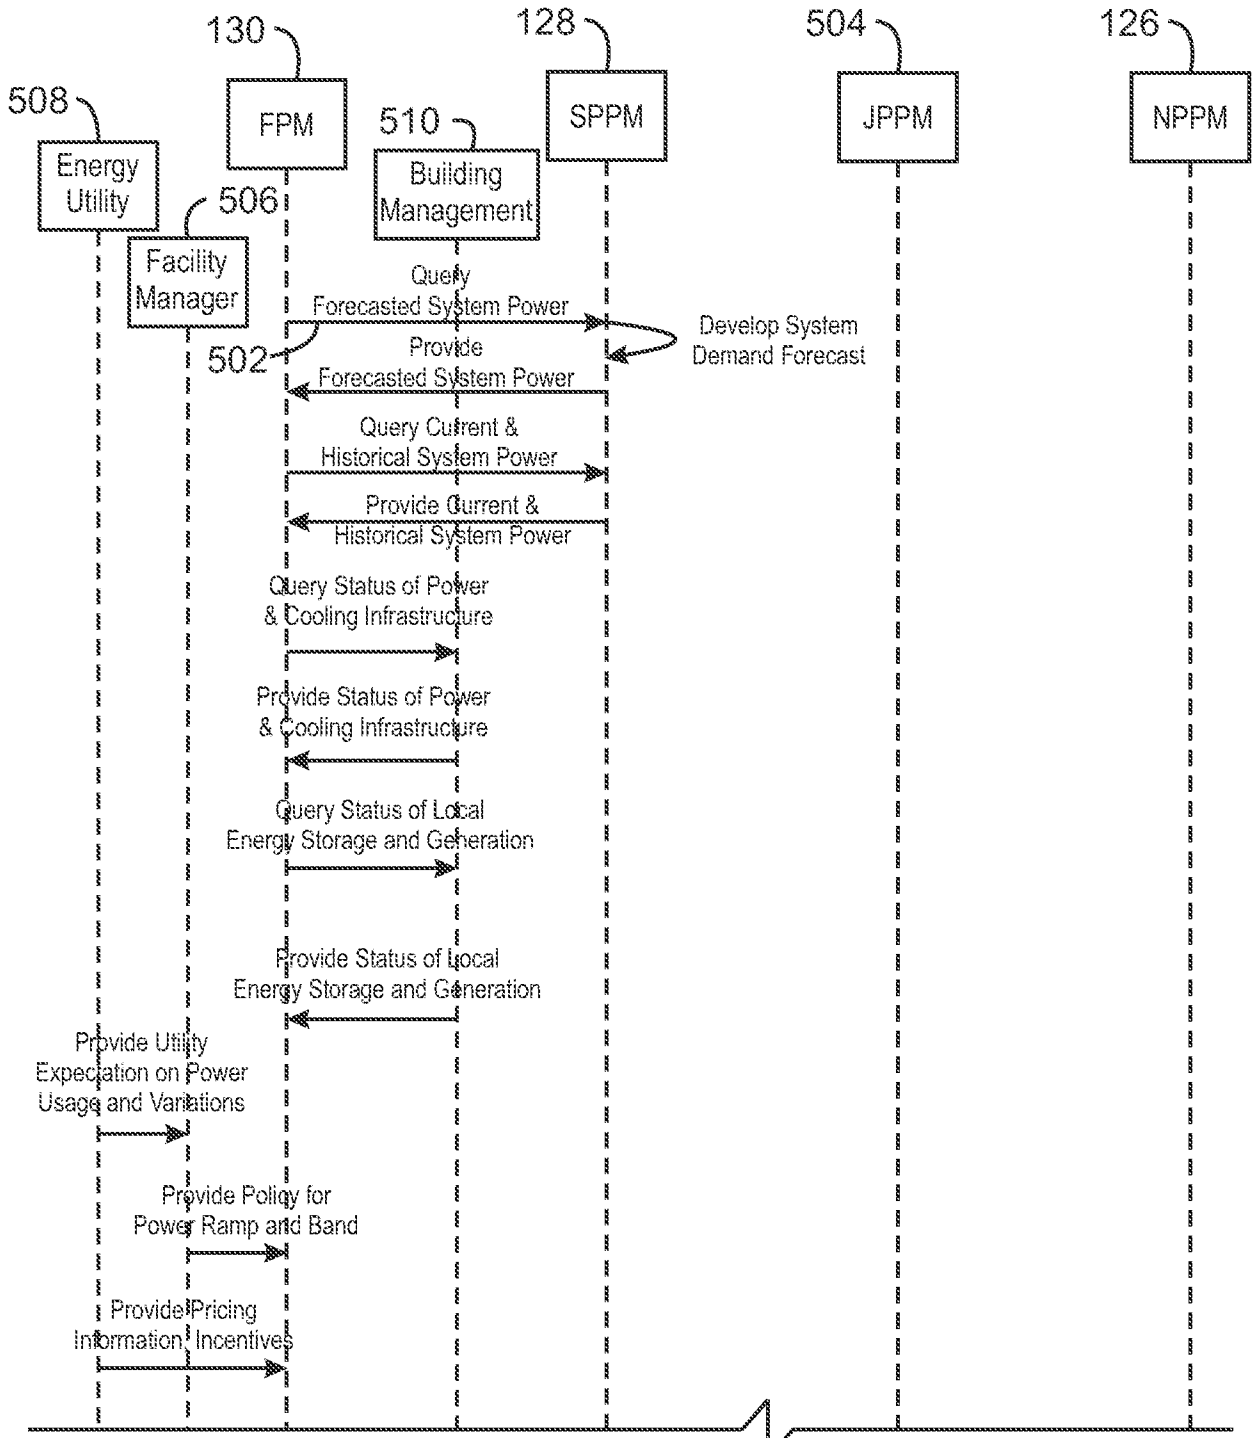


FIG. 3



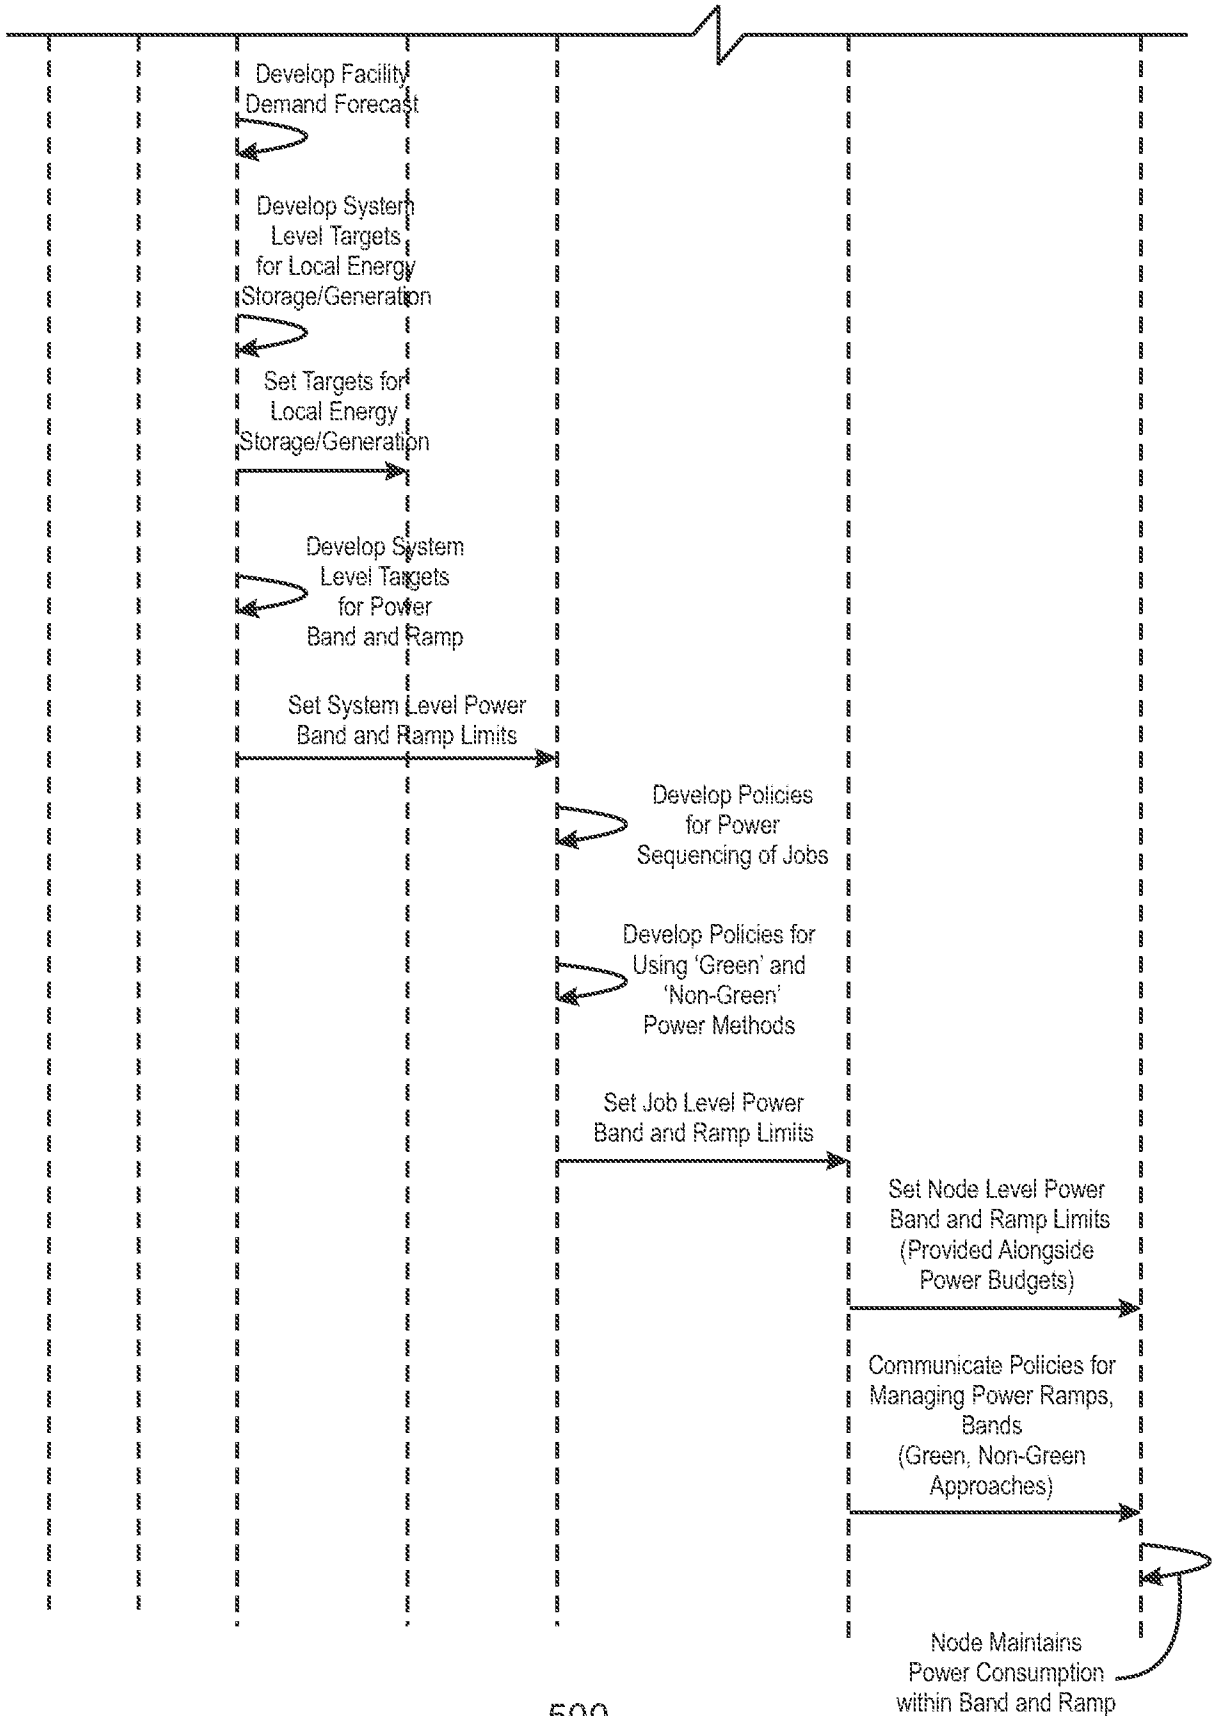
400
FIG. 4




Continued on FIG. 5B

500
FIG. 5A

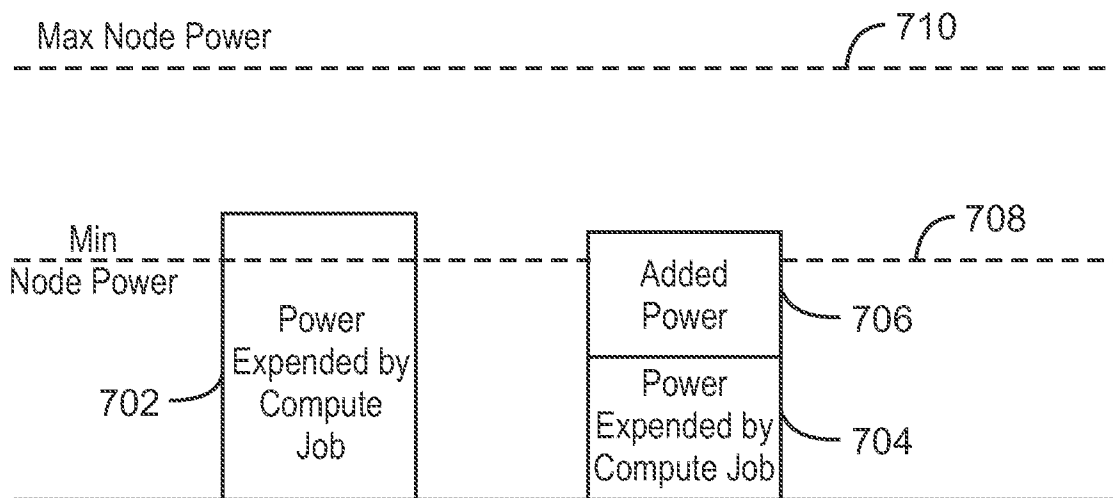
Continued on FIG. 5A




500
FIG. 5B



106
FIG. 6



700
FIG. 7



800
FIG. 8

A. CLASSIFICATION OF SUBJECT MATTER**G06F 1/32(2006.01)i, G06F 1/26(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06F 1/32; G06F 1/26; G06F 1/00Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: power, consumption, band, computing, job, node, manage, rate**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2009-0049313 A1 (THOMAS M GOODING et al.) 19 February 2009 See paragraphs [0041], [0048], [0056]; claim 15; and figures 1, 8.	1-4, 16-17, 19-24
A	US 2009-0070611 A1 (FRED A. BOWER, III et al.) 12 March 2009 See paragraph [0035]; and figure 5.	1-4, 16-17, 19-24
A	US 2013-0054987 A1 (CLEMENS PFEIFFER et al.) 28 February 2013 See paragraph [0015]; and figure 2.	1-4, 16-17, 19-24
A	US 2009-0235097 A1 (JAMES R. HAMILTON et al.) 17 September 2009 See paragraph [0063]; and figure 7.	1-4, 16-17, 19-24
A	US 7210048 B2 (DEVADATTA V. BODAS) 24 April 2007 See column 3, line 45 - column 4, line 6; and figure 3.	1-4, 16-17, 19-24

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

21 January 2016 (21.01.2016)

Date of mailing of the international search report

22 January 2016 (22.01.2016)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

LEE, Dong Yun

Telephone No. +82-42-481-8734



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

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

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.: 15
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
Claim 15 is not clear because it refers to claim 14, which does not comply with PCT Rule 6.4(a).

3. Claims Nos.: 5-14, 18, and 25
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

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

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

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of any additional fees.

3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

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

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/054438

Patent document cited in search report	Publication date	Patent family member(s)	Publication date		
US 2009-0049313 A1	19/02/2009	CN 101779189 A	14/07/2010		
		CN 101779189 B	09/01/2013		
		EP 2179359 A1	28/04/2010		
		EP 2179359 B1	27/02/2013		
		JP 2010-537266 A	02/12/2010		
		JP 4659138 B2	30/03/2011		
		KR 10-1137073 B1	19/04/2012		
		KR 10-2010-0041829 A	22/04/2010		
		TW 200917052 A	16/04/2009		
		US 7941681 B2	10/05/2011		
		WO 2009-024459 A1	26/02/2009		
		US 2009-0070611 A1	12/03/2009	None	
		US 2013-0054987 A1	28/02/2013	AU 2012-302000 A1	20/03/2014
CA 2847258 A1	07/03/2013				
EP 2751632 A1	09/07/2014				
JP 2014-527394 A	09/10/2014				
WO 2013-033217 A1	07/03/2013				
US 2009-0235097 A1	17/09/2009	US 8001403 B2	16/08/2011		
US 7210048 B2	24/04/2007	CN 101346681 A	14/01/2009		
		CN 101346681 B	06/07/2011		
		KR 10-0824480 B1	22/04/2008		
		KR 10-2005-0102113 A	25/10/2005		
		US 2004-0163001 A1	19/08/2004		
		WO 2004-074983 A2	02/09/2004		
		WO 2004-074983 A3	14/08/2008		