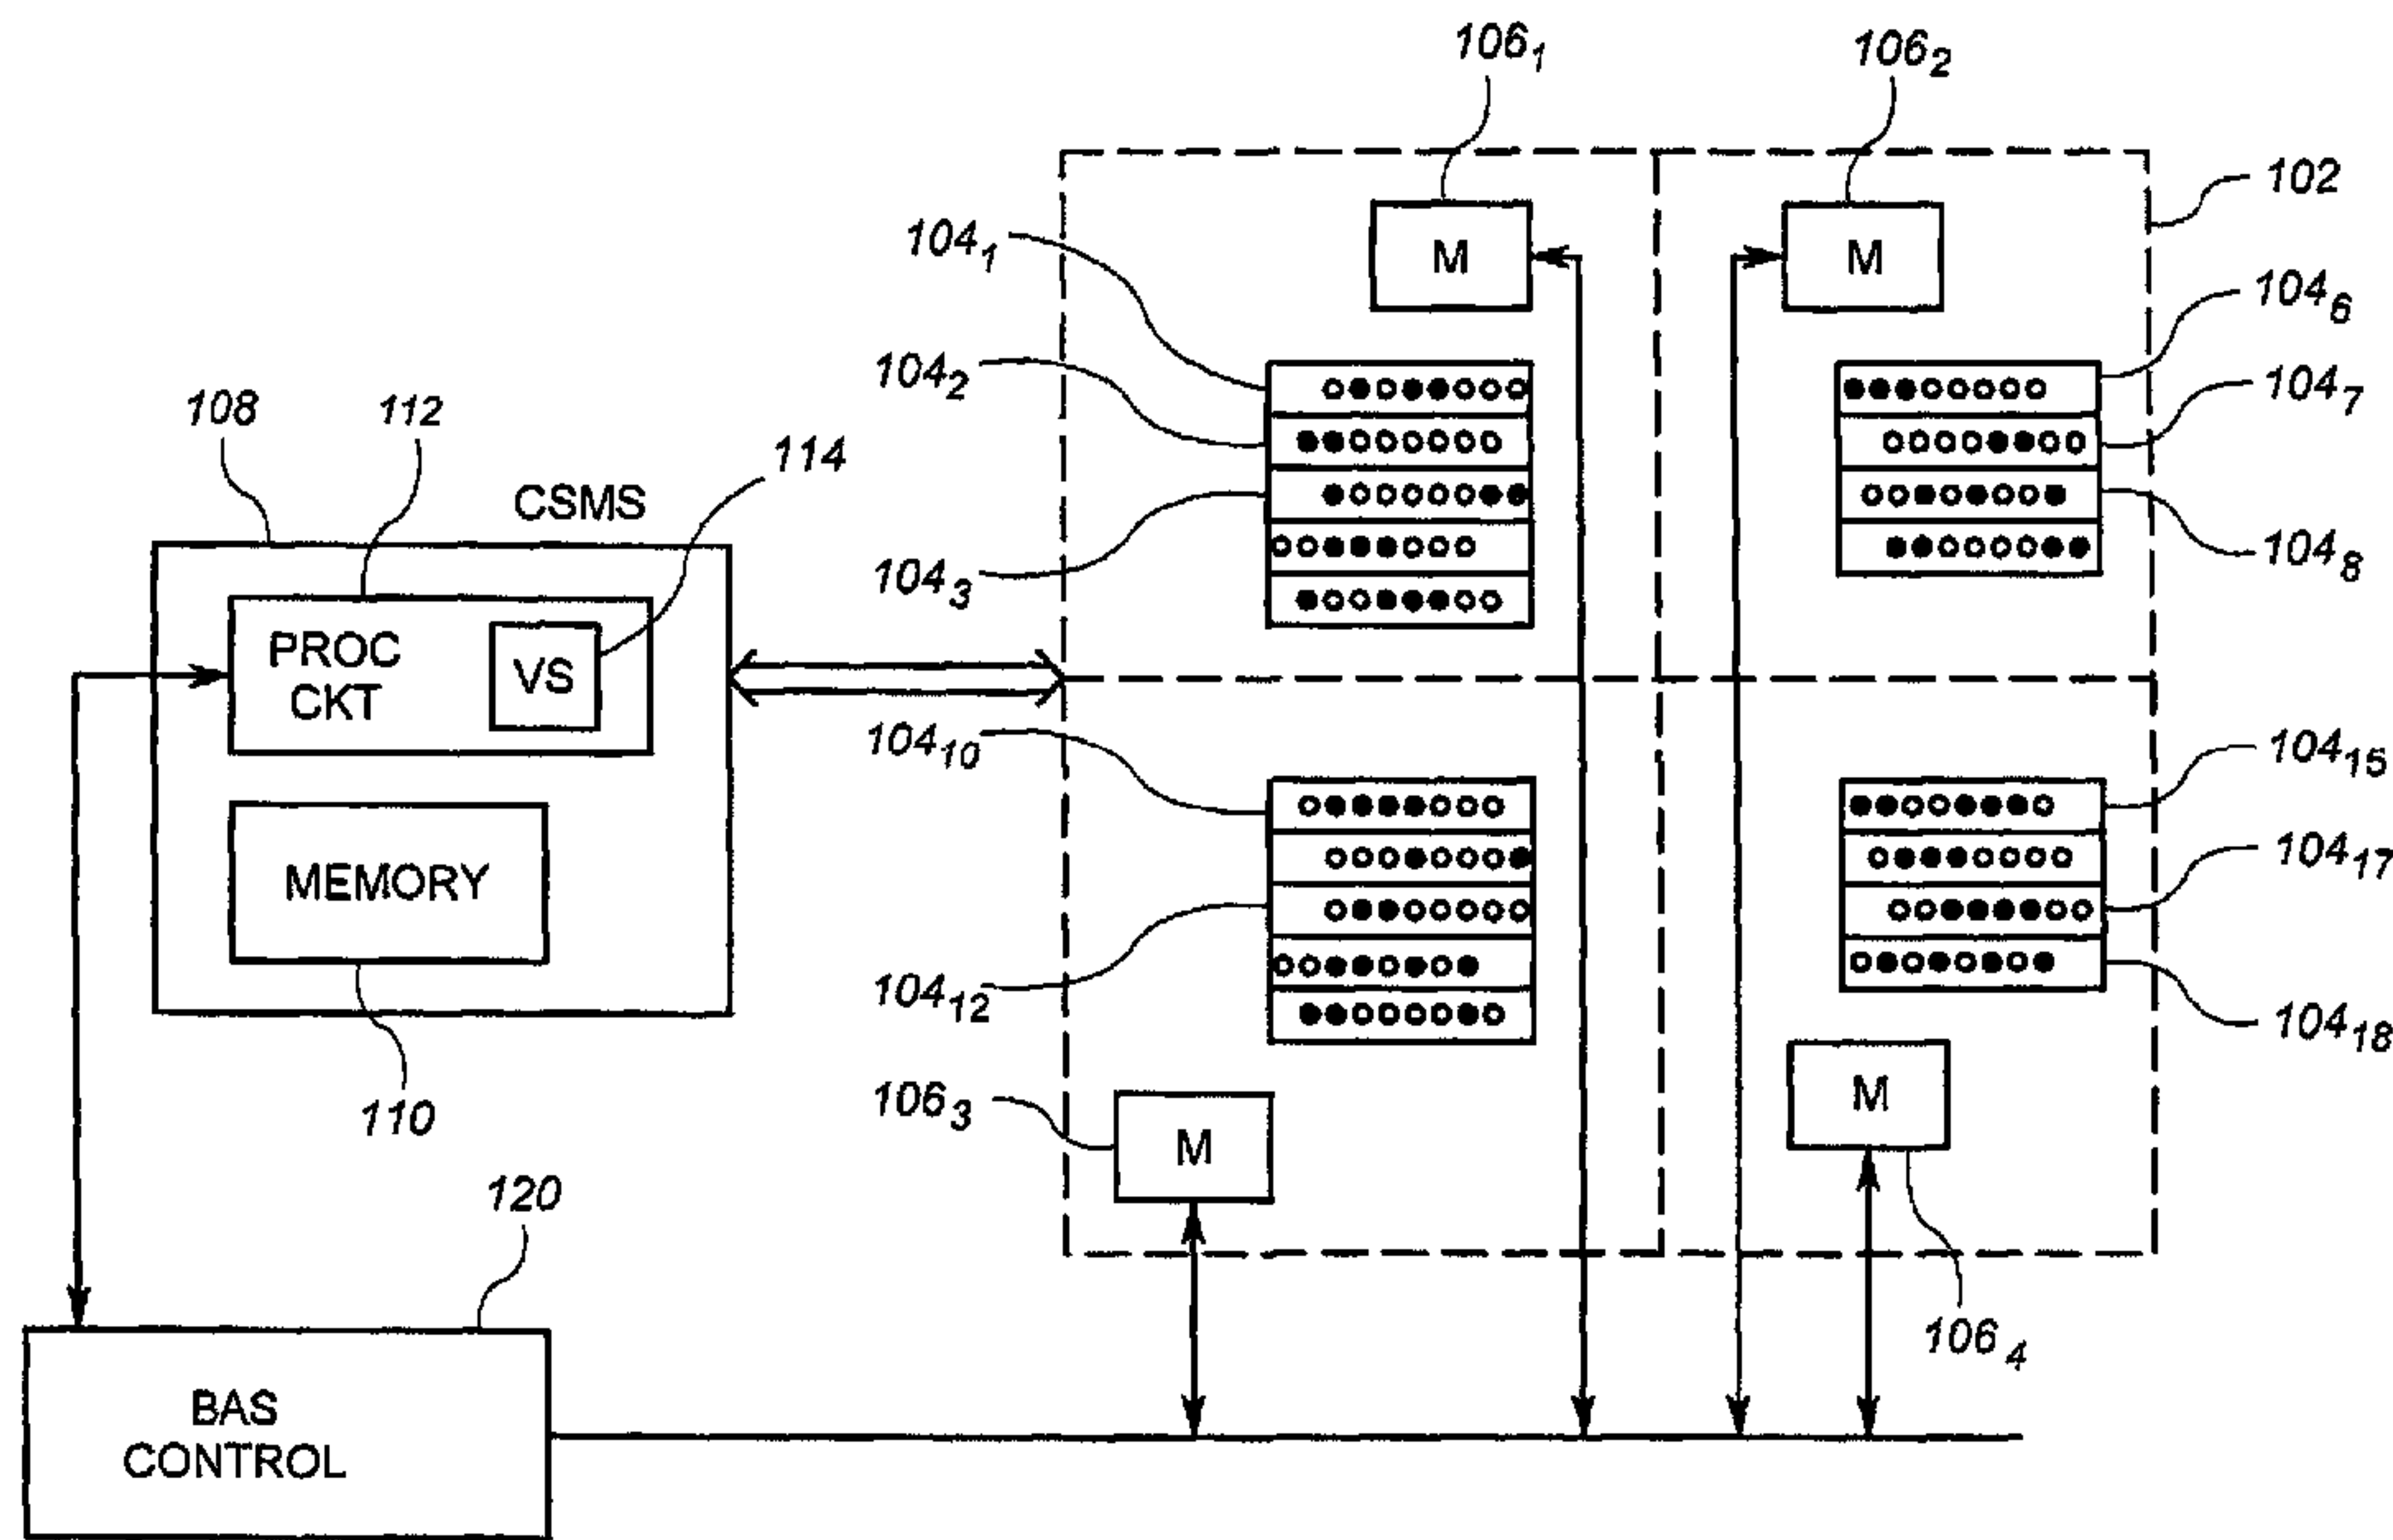




(86) Date de dépôt PCT/PCT Filing Date: 2009/05/05
 (87) Date publication PCT/PCT Publication Date: 2009/11/12
 (45) Date de délivrance/Issue Date: 2019/03/26
 (85) Entrée phase nationale/National Entry: 2010/11/03
 (86) N° demande PCT/PCT Application No.: US 2009/002762
 (87) N° publication PCT/PCT Publication No.: 2009/137026
 (30) Priorités/Priorities: 2008/05/05 (US61/050,420);
 2008/05/05 (US61/050,425); 2008/05/05 (US61/050,429);
 2009/05/04 (US12/435,401)

(51) Cl.Int./Int.Cl. *G06F 1/20* (2006.01)
 (72) Inventeurs/Inventors:
 PIANTA, WILLIAM THOMAS, US;
 SONGKAKUL, PORNSAK, US
 (73) Propriétaire/Owner:
 SIEMENS INDUSTRY, INC., US
 (74) Agent: SMART & BIGGAR

(54) Titre : AGENCEMENT DE GESTION D'OPERATIONS DE CENTRE DE DONNEES DESTINE A AMELIORER L'EFFICACITE DE REFROIDISSEMENT
 L'EFFICACITE DE REFROIDISSEMENT
 (54) Title: ARRANGEMENT FOR MANAGING DATA CENTER OPERATIONS TO INCREASE COOLING EFFICIENCY



(57) Abrégé/Abstract:

A method includes a step of obtaining efficiency characteristic information for each of a plurality of air conditioning units in a location containing a plurality of server computers. The method also includes employing one or more processing units to allocate one or more processing tasks to one of the plurality of server computers based on the efficiency characteristic information.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
12 November 2009 (12.11.2009)(10) International Publication Number
WO 2009/137026 A3(51) International Patent Classification:
G06F 1/20 (2006.01)(21) International Application Number:
PCT/US2009/002762(22) International Filing Date:
5 May 2009 (05.05.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/050,429 5 May 2008 (05.05.2008) US
61/050,425 5 May 2008 (05.05.2008) US
61/050,420 5 May 2008 (05.05.2008) US
12/435,401 4 May 2009 (04.05.2009) US(71) Applicant (for all designated States except US):
SIEMENS INDUSTRY, INC. [US/US]; 3333 Old Milton Parkway, Alpharetta, GA 30005-4437 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **PIENTA, William, Thomas** [US/US]; 414 West Olive Avenue, Prospect Heights, IL 60070 (US). **SONGKAKUL, Pornsak** [US/US]; 2113 West Lilac Lane, Mequon, WI 53092 (US).(74) Agents: **RAJ, Rashmi, S.** et al.; Siemens Corporation-Intellectual, Property Dept., 170 Wood Avenue South, Iselin, NJ 08830 (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, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, 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))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(88) Date of publication of the international search report:
7 January 2010

(54) Title: ARRANGEMENT FOR MANAGING DATA CENTER OPERATIONS TO INCREASE COOLING EFFICIENCY

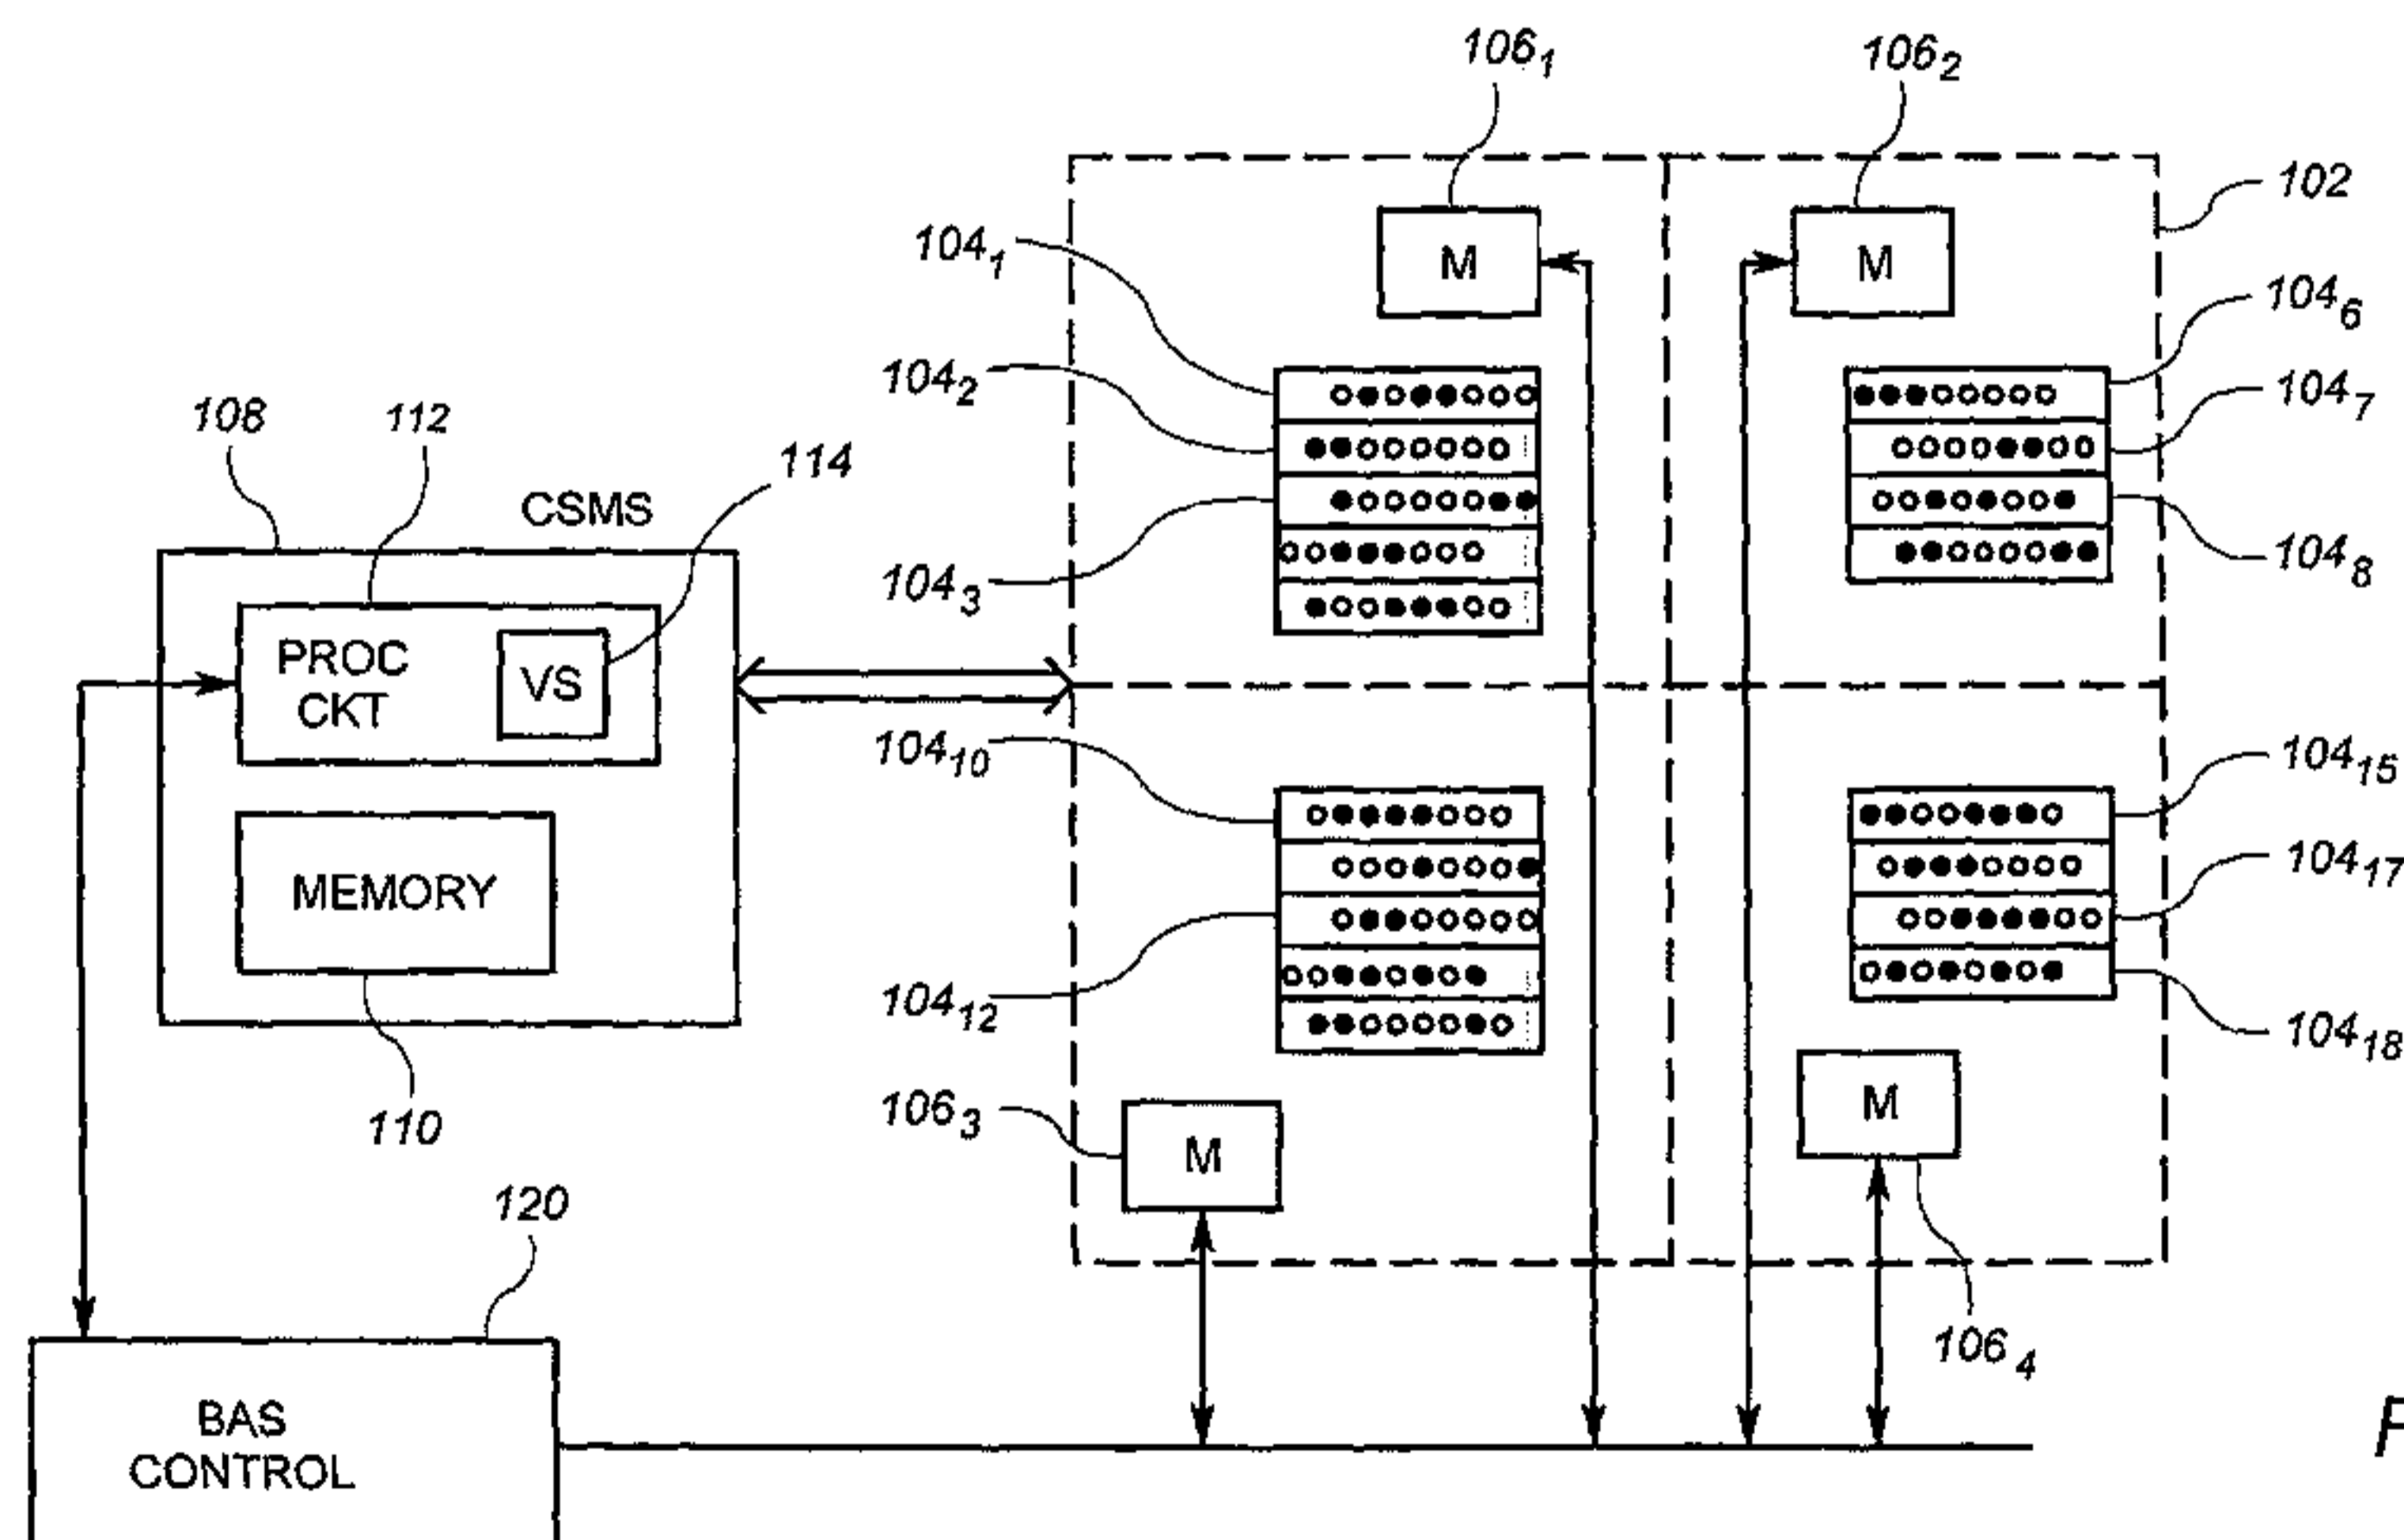


FIG. 1

(57) Abstract: A method includes a step of obtaining efficiency characteristic information for each of a plurality of air conditioning units in a location containing a plurality of server computers. The method also includes employing one or more processing units to allocate one or more processing tasks to one of the plurality of server computers based on the efficiency characteristic information.

54106-624

Arrangement for Managing Data Center Operations to Increase Cooling Efficiency

Field of the Invention

The present invention relates generally to data processing centers,
5 and more particularly, to the management of the operations of a data processing center.

Background of the Invention

Data centers are parts of buildings or facilities in which a large
number of server computers are located. The dense packing of the server
10 computers results in the generation of a large amount of heat in a localized area. The data center must be cooled in a reliable manner in order to avoid the shutting down of, or damage to, the server computer hardware. Shutting down of server computers due to heat overload can cause significant economic loss.

Accordingly, specialized cooling units have been developed for
15 implementation directly in data centers. These specialized cooling units are sometimes known in the art as

computer room air conditioning units (“CRACs”) or computer room air handling units. In this disclosure, air conditioning unit or CRAC will be understood to encompass any device used to effect cooling in a data center. CRACs have been employed as a result of the fact that the ordinary HVAC systems of buildings are not optimally configured to handle the concentrated heat generated with data centers. Thus, CRACs are often used in connection with, but in addition to, the ordinary cooling units of a building employed for human comfort systems.

Many CRACs have simple, embedded controls that adjust the unit output based on factors such as sensed ambient air temperature. In some cases, CRACs have controllers that interact with the building automation system that controls or includes the building HVAC system, among other things.

While CRACs provide a solution to the need for enhanced cooling power within a data center having several server computers, the power consumed by CRACs is significant. It is therefore desirable to reduce the power consumption associated with the adequate cooling of data centers.

Summary

The present invention addresses the above identified needs, as well as others, by providing a method and arrangement that allocates processing tasks to servers based on the location of the servers with respect to air conditioning units. Moreover, embodiments of the invention allocate processing tasks to servers based on efficiency characteristics of air conditioning units within the data center. As a result, heat generated by the execution of processing tasks is advantageously distributed to the air conditioning units based on their

efficiencies.

A first embodiment is a method that includes a step of obtaining efficiency characteristic information for each of a plurality of air conditioning units in a location containing a plurality of server computers. The method also includes employing one or more processing units to allocate one or more processing tasks to one of the plurality server computers based on the efficiency characteristic information.

In some embodiments, a first processing request is allocated to a first of the plurality of computers based on a thermal proximity of the first of the plurality of computers to a select one of the plurality of air conditioning units. In addition, the select one of the plurality of air conditioning units is selected based on the efficiency characteristic information.

A second embodiment is an arrangement that includes a computer server management system having a memory and processing circuit. The processing circuit is configured to coordinate the usage of a plurality of server computers. The memory stores efficiency characteristic information for a plurality of air conditioning units in a location containing a plurality of server computers. The processing circuit is further configured to allocate one or more processing tasks to one of the plurality of server computers based on the efficiency characteristic information.

In some embodiments, there is provided a method, comprising: a) obtaining efficiency characteristic information for each of a plurality of air conditioning units in a location containing a plurality of server computers, wherein the efficiency characteristic information is representative of an efficiency performance curve for a range of variable cooling output; and b) employing one or more processing units to allocate one or more processing tasks to one of the plurality of server computers based on the efficiency characteristic information and to: i) identify for a total thermal load level of each of said plurality of server computers a combination of air conditioning units capable of cooling the total thermal load level using less energy than a plurality of other combinations of air conditioning units capable of cooling the total thermal load level; ii) allocate processing tasks

to server computers based on their thermal proximity to the identified combination of air conditioning units.

In some embodiments, there is provided an arrangement, comprising: a computer server management system having a processing circuit and a memory, the processing circuit configured to coordinate the usage of a plurality of server computers, the processing circuit executing virtualization software configured to manage application processing on the plurality of server computers; the memory storing efficiency characteristic information for a plurality of air conditioning units in a location containing a plurality of server computers, wherein the efficiency characteristic information is representative of an efficiency performance curve for a range of variable cooling output; wherein the processing circuit is further configured to allocate one or more processing tasks to one of the plurality of server computers based on the efficiency characteristic information and to: i) identify for a total thermal load level of each of said plurality of server computers a combination of air conditioning units capable of cooling the total thermal load level using less energy than a plurality of other combinations of air conditioning units capable of cooling the total thermal load; ii) allocate processing tasks to server computers based on their thermal proximity to the identified combination of air conditioning units.

The above described features and advantages, will become more readily apparent to those of ordinary skill in the art by reference to the following detailed description and accompanying drawings.

Brief Description of the Drawings

Fig. 1 shows a schematic block diagram of an exemplary arrangement according to a

first embodiment of the invention implemented to coordinate application processing in an exemplary data center;

Fig. 2 shows a set of exemplary efficiency characteristic curves for air conditioning units in the data center shown in Fig. 1;

Fig. 3 shows an exemplary set of operations that may be carried in accordance with the present invention;

Fig. 4 shows in further detail an exemplary embodiment of at least one of the operations of Fig. 3; and

Fig. 5 shows in further detail a second embodiment of at least one of the operations of Fig. 3.

Detailed Description

Fig. 1 shows an arrangement 100 according to an exemplary embodiment of the invention. The arrangement 100 is shown used in conjunction with a data center 102 that includes a plurality of server computers 104₁, 104₂... 104₁₈ and a plurality of air conditioning units 106₁, 106₂, 106₃, and 106₄. The arrangement 100 includes a computer server management system 108 having, among other things, a memory 110 and a processing circuit 112. In this embodiment, the arrangement 100 further includes a BAS element 120, which is communicatively connected to the processing circuit 112.

Each of the server computers 104₁ to 104₁₈ is part of a set of computers that provide application processing services to at least one, and typically a large number of, client computers, not shown. The server computers 104₁ to 104₁₈ are typically arranged in racks and dispersed throughout the space of the data center 102. For example, as shown in Fig. 1,

the server computers 104₁, 104₂, 104₃, 104₄, and 104₅ may be grouped on a first rack, the server computers 104₆, 104₇, 104₈, and 104₉ may be grouped on a second rack, the server computers 104₁₀, 104₁₁, 104₁₂, 104₁₃, and 104₁₄ may be grouped on a third rack, and the server computers 104₁₅, 104₁₆, 104₁₇, and 104₁₈ may be grouped on a fourth rack.

Each of the air conditioning units 106₁, 106₂, 106₃, and 106₄, is a computer room air conditioner or computer room air handler unit, collectively referred to as CRACs, or any air conditioning unit that may suitably be employed to specifically cool a data center or other area that is a high heat generator, and requires significant localized cooling. Such devices are well known in the art. In this embodiment, each of the air conditioning units (ACUs) 106₁, 106₂, 106₃, and 106₄ is operably coupled to the BAS element 120 such that the BAS element 120 can perform at least some measure of control over the operations of the ACUs 106_n. For example, if an ACU 106_n has self-contained temperature sensing and control, the BAS element 120 may be operably connected to override the on/off local control, and/or to provide a set point to the ACU 106_n. Other ACUs may be configured for more extensive control by an external field controller of the building automation system. In either event, the BAS element 120 preferably is operably connected to provide overall management and/or control of the each of the ACUs 106₁, 106₂, 106₃, and 106₄.

The BAS element 120 is one or more devices that are configured to communicate with, and operate within, a building automation system such as an HVAC system or the like. Such systems are known in the art and may have a general architecture of the APOGEE™ system available from Siemens Building Technologies Inc. The BAS element 120 includes at least one processing circuit 140 and a memory 142. The BAS element 120 may suitably take the form of a supervisory work station in a BAS such as the INSIGHT™ work station

available from Siemens Building Technologies, Inc., of Buffalo Grove, Illinois. In the alternative, the BAS element 120 may suitably be a configurable field controller, such as the PXC Modular field controller, also available from Siemens Building Technologies, Inc. In general, the processing circuit 140 is configured via other circuits to communicate BAS data (such as set points, sensor values, and commands) with other BAS devices such as other controllers, or even with sensors and actuators. The BAS element 120 may further includes special digital or analog I/O devices as may be necessary to communicate with control elements of the ACUs 106₁, 106₂, 106₃, and 106₄. In this embodiment, the BAS element 120 is further operably connected to communicate information with the computer server management system 108, and particularly the processing circuit 112. To this end, a suitable data interface is provided between the BAS element 120, which is configured for a BAS system, and the computer server management system 108, which is typically not set up for communication with a BAS system.

The computer server management system 108 is a computing system that is generally configured to coordinate the usage of the plurality of server computers 104₁, 104₂. Such devices are generally known. To coordinate the server usage, the processing circuit 112 of the computer server management system 108 executes virtualization software 114. Virtualization software 114, as is known in the art, is software that, when executed by a computer processor, manages the allocation of application processes among a plurality of server computers, such as in a data center.

In accordance with this embodiment of the present invention, the processing circuit 112 is further configured to allocate application processes among the server computers 104₁, 104₂, etc. based on the efficiency characteristics of the ACUs 106₁, 106₂, 106₃, 106₄. As will

be discussed below in detail, the processing circuit 112 allocates applications processes among the server computers 104₁ to 104₁₈ such that the resultant heat generated by the processing of the applications is distributed among the ACUs 106₁, 106₂, 106₃, 106₄ in a manner that takes into account the relative efficiencies of the units for different thermal load levels.

To this end, the memory 110 stores efficiency characteristic information for the plurality of ACUs 106₁, 106₂, 106₃, 106₄ in the data center 102. The efficiency characteristic information may suitably comprise part-load performance curves for each of the ACUs 106₁, 106₂, 106₃, 106₄. For example, Fig. 2 shows exemplary graphs 202, 204, 206 and 208, each graph showing efficiency curves 212, 214, 216 and 218 for various loads. The x-axis as represents the cooling output, and the y-axis represents the power consumption.

The efficiency curves 212, 214, 216 and 218 may be generated in a plurality of ways, some depending on the type of air conditioning unit. In general, the curves may be generated by identifying the power consumption for a plurality of cooling output values. The curves 212, 214, 216 and 218 may be generated from the resultant data points via interpolation or curve-fitting techniques.

To obtain the data points, the cooling output is generally measured as the change in air temperature (ΔT) between the input and output of the ACU, multiplied by the air flow through the unit. The associated power consumption can be measured in various ways depending on the type of unit. If the unit is entirely self-contained (e.g. contains a compressor and refrigerant), the electrical power consumption of the ACU provides a good measure of the power consumption. If, however, the device relies on coolant provided and/or pumped from another source, such as chilled water, then the power consumption can be

measured as the flow of coolant multiplied by the change in temperature of the coolant between the input and output of the ACU. The electrical power consumption of the unit's fan must also be added. Obviously, this measure of power consumption can be made more accurate by taking into account coolant transmission loss to and from the ACU, as well as other losses.

In any event, after various data points of cooling output versus power consumption are obtained for each air-conditioning unit 106₁, 106₂, 106₃, 106₄, the efficiency curves 212, 214, 216, and 218 are generated using straight-line interpolation or curve-fitting may be used. In any event, data representative of the curves 212, 214, 216 or 218 is stored in the memory 110.

Referring again to Fig. 1, as discussed above, the computer server management system 108 is configured to allocate (via the virtualization software) one or more processing tasks to one of the plurality computer servers 104₁ - 104₁₈ based in part on the efficiency characteristic information of the air-conditioning units 106₁, 106₂, 106₃, 106₄. In a first exemplary operation, applications are directed to computer servers 104_n, 104_o, etc. based on the thermal proximity of the servers 104_n, 104_o, etc. to the air-conditioning units 106₁ to 106₄ and the stored efficiency characteristics of the air-conditioning units 106₁ to 106₄.

It will be understood that the phrase "thermal proximity" as used herein means an extent to which two items are thermally coupled, e.g. how much heat is transferable between them. For example, if a computer server is adjacent to the input of an ACU such that the unit receives and removes all of the heat generated by the computer server, then they are closely thermally coupled, and thus have a high degree of "thermal proximity". While thermal proximity is often a result of physical proximity, it is possible for a server computer to be

physically proximate to, but not very tightly thermally coupled to, a particular ACU. Relative thermal proximity between cooling devices such as ACUs and heat generating equipment in an environment such as a data center may readily be determined or at least estimated by those of ordinary skill in the art.

In any event, when an application is assigned to a server computer 104_n , the execution of the application causes that server computer 104_n generate heat energy. The processing circuit 112 allocates the processing tasks such that the heat generated by the server computers 104 executing the processing tasks is distributed efficiently among the various ACUs 106_1 , 106_2 , 106_3 and 106_4 .

By way of example, consider a situation in which 100 applications must be allocated to the server computers 104_1 to 104_{18} . In the prior art, one way to allocate the applications may be to simply allocate a substantially an equal number of applications to each of the processors, such that in this example each of the server computers 104_1 to 104_{18} would have five or six of the one hundred applications. Alternatively, the allocation may be based on attempting to keep the busy-ness of each of the server computers 104_1 to 104_{18} roughly equal. Thus, if a particular server computer 104_n has a number of particularly computationally intensive tasks, it may have fewer overall applications. The computing speed and efficiency of the server computers 104_1 to 104_{18} may also be taken into account. In either event, the prior art method can include more or less equal distribution of processing load among the server computers.

However, it may be the case that such an allocation does not result in efficient use of the ACUs 106_1 , 106_2 , 106_3 , 106_4 . To this end, assume that the ACU 106_1 generally is responsible for cooling (i.e. is thermally proximate to) the server computers 104_1 to 104_5 , the

ACU 106₂ generally is responsible for cooling the server computers 104₆ to 104₉, the ACU 106₃ generally is responsible for cooling the server computers 104₁₀ to 104₁₄, and the ACU 106₄ generally is responsible for cooling the server computers 104₁₅ to 104₁₈. It may be the case that with the heat loading cause by a more or less equal allocation of applications among the processors 104₁ to 104₁₈ does not result in the ACUs 106₁, 106₂, 106₃, 106₄ operating at their most efficient levels possible. Indeed, even assuming that the heat generated by server computers 104₁ to 104₁₈ is equally distributed among the ACUs 106₁, 106₂, 106₃, 106₄, the ACUs 106₁, 106₂, 106₃, 106₄ as a group may not be operating at their highest efficiency levels. In particular, each ACU 106_n may have a different level at which they operate most efficiently, and may be more or less efficient than other units at various output levels. This is evidenced by the variability of the curves 212, 214, 216 and 218. Thus, in accordance with the present invention, the processing circuit 112 assigns the applications to the server computers 104₁ to 104₁₈ such that the heat is allocated in accordance with the combined efficient operating levels of the ACUs 106₁, 106₂, 106₃, 106₄.

In a simple illustration, assume that at a certain *average* thermal load level, the first ACU 106₁ may be operating at a much more efficient level than the second ACU 106₂ another. In such a case, it might be advantageous to place more processing load on the server computers 104₁ to 104₅ because they are more thermally proximate to the more efficient first ACU 106₁, and place less processing load on server computers 104₆ to 104₉ because they are more thermally proximate to the less efficient ACU 106₂. However, at the same time, placing too much of the processing load on processors thermally proximate to the first ACU 106₁ and too little on the server computers thermally proximate to the second ACU 106₂ may also result in reduced efficiency. The processing circuit 112 thus uses efficiency information

such as the curves 212, 214, 216 and 218 to distribute the heat load to the ACUs 106₁, 106₂, 106₃, 106₄ in an efficient manner. The heat load is distributed to a particular ACU 106_m by assigning processing tasks to server computers 104_n that are thermally proximate to (i.e. tightly thermally coupled to) that ACU 106_m.

The processing circuit 112 is therefore configured to determine the allocation of at least some processes based on the efficiency characteristics of the four ACUs 106₁, 106₂, 106₃, 106₄. In general, such a determination is further dependent upon the thermally proximity of the server computers 104₁ to 104₁₈ to the ACUs 106₁, 106₂, 106₃, 106₄.

Fig. 3 shows an exemplary embodiment of the general operations of the processing circuit 112 of Fig. 1. In step 305, the processing circuit 112 obtains efficiency characteristics for the ACUs 106₁, 106₂, 106₃, 106₄. As discussed above, the efficiency characteristics may suitably comprise the part-load curves 212, 214, 216 and 218 of cooling output as a function of energy consumption, as illustrated in the graphs 202, 204, 206 and 208 of Fig. 2. Typically, the processing circuit 112 obtains this information via communications, or by user input, as a result of a set-up procedure for the data center 102. The processing circuit 112 stores the efficiency information in the memory 110. In some embodiments, the efficiency characteristics of the ACUs 106₁, 106₂, 106₃, 106₄ may be stored in a memory of the BAS element 120, or another computer.

In step 310, the processing circuit 112 determines how to allocate the processing load based on the efficiency characteristics of the ACUs 106₁, 106₂, 106₃, 106₄. To this end, the processing circuit 112 also has information that correlates the server computers 104₁ to 104₁₈ to the ACUs 106₁, 106₂, 106₃, 106₄. For example, each server computer 104_n is identified as having a thermally closest ACU 106_m. By thermally closest, it is meant that the heat load

generated by the server computer 104_n is borne primarily by that ACU 106_m , or in other words, the server computer 104_n is most thermally proximate to the ACU 106_m .

In a more sophisticated embodiment, each server computer 104_n may have a thermal impact *function* that consists of a weighted polynomial, wherein each term represents the thermal load bearing capacity of one of the ACUs 106_1 , 106_2 , 106_3 , 106_4 . This embodiment acknowledges that the heat load generated by a server computer 104_n may in some cases be handled, at least in part, but more than one of the ACUs 106_1 , 106_2 , 106_3 , 106_4 . Such a thermal impact function for a server computer 104_n would be expressed as:

$$H_{total,n} = I_{1061,n}H_{total,n} + I_{1062,n}H_{total,n} + I_{1063,n}H_{total,n} + I_{1064,n}H_{total,n}$$

wherein $I_{1061,n}$, $I_{1062,n}$, $I_{1063,n}$, $I_{1064,n}$ are weighting factors totaling 1, and $H_{total,n}$ is the total heat generated by the server computer 104_n . The weighting factors represent the proportion of heat generated by the server computer 104_n that is loaded on to each of the ACUs 106_1 , 106_2 , 106_3 , 106_4 . For example, the server computer 104_{15} may have weighting factors $I_{1061,15} = 0.02$; $I_{1062,15} = 0.12$; $I_{1063,15} = 0.01$; and $I_{1064,15} = 0.85$, while the server computer 104_{12} may have weighting factors $I_{1061,12} = 0.02$; $I_{1062,12} = 0.00$; $I_{1063,12} = 0.9$; and $I_{1064,12} = 0.08$.

The processing circuit 112 then uses the computer server-ACU correlation information to distribute the processing tasks such that the heat load will cause the mix of ACUs 106_1 , 106_2 , 106_3 , and 106_4 to be operating at a collectively best (or very good) efficiency load levels.

In a very simple example of step 310, the processing circuit 112 may execute an algorithm that starts with a first proposed task allocation, estimates the heat generated thereby, and then identifies the load on the conditioning units 106_1 , 106_2 , 106_3 , and 106_4 using the server/ACU thermal proximity correlation information. Once the load on the

ACUs 106₁, 106₂, 106₃, and 106₄ is estimated, the processing circuit 112 uses the stored efficiency information (e.g. curves 212, 214, 216 and 218) to calculate the overall efficiency (or power consumption value) for the first proposed task allocation. Thereafter, the processing circuit 112 generates a *second* proposed task allocation that is slightly (or significantly) different from the first proposed task allocation. The processing circuit 112 uses the same procedure to determine the overall efficiency of the second proposed task allocation. The processing circuit 112 repeats this process for additional proposed task allocations. The processing circuit 112 may also use any number of converging algorithms to solve for a best (or very good) proposed task allocation based on the overall efficiencies of the various proposed task allocations.

However, it will be appreciated that the allocation of tasks to achieve efficient heat loading of the ACUs 106₁, 106₂, 106₃, and 106₄ may be carried out in other ways. Some other ways are discussed below in connection with Figs. 4 and 5.

Referring again to Fig. 3, once the processing circuit 112 determines the allocation of computing tasks in step 310, the processing circuit 112 in step 315 causes the allocation of the processing tasks in accordance with the determination. The operations of step 315 may be carried out using the known capabilities of computer server management systems. After step 315, the computer servers 104₁ to 104₁₈ may suitably execute the processing task that have been assigned to them.

In an optional step 320, the processing circuit 112 further communicates load information to the BAS control 120. The load information either identifies the overall processing load on the servers 104₁ to 104₁₈, any change in load on the servers 104₁ to 104₁₈, or any estimate of thermal load on the ACUs 106₁, 106₂, 106₃, and 106₄. The BAS control

120 in some cases may use this information to make proactive decisions on thermal cooling, or to control aspects of the ACUs 106₁, 106₂, 106₃, and 106₄ to accommodate any needed increase or decrease in cooling output.

Fig. 4 shows an example of the processing task allocation determination that may be used as step 310 in Fig. 3. The operations of Fig. 4 describe a processing task allocation determination that occurs incrementally during ongoing operation of the data center. More specifically, the operations of Fig. 4 may be used when a new task or a new set of processing tasks are requested of the server computers 104₁ to 104₁₈. While the operations of Fig. 4 are described with respect to a single additional application or processing request, it will be appreciated that the steps of Fig. 4 may suitably be used with a group of additional processing requests.

In some cases, the processing circuit 112 can proactively determine where the next n applications should be assigned before the requests are actually received.

In step 405, the processing circuit 112 of the computer server management system 108 receives an additional processing request. To this end, from time to time, client computers that access the data center 102 provide requests for application tasks to be executed by one or more of the computer servers 104₁ to 104₁₈. Assuming that the additional processing request is non-trivial in nature, the processing circuit 112 proceeds to step 410 to determine an allocation for the task.

In step 410, the processing circuit 112 determines which of the ACUs 106 _{m} would handle additional heat load with the best efficiency. Such a determination may be carried out by identifying the current thermal load on each of the ACUs 106₁, 106₂, 106₃, 106₄, and using the efficiency characteristics to identify the ACU 106 _{m} that would require the least

amount of incremental additional power consumption to provide the additional cooling.

For example, referring to Fig. 2, assume that an existing heat loading on the 106₁, 106₂, 106₃, 106₄, is shown as point A on each of the graphs 202, 204, 206 and 208. It can be seen that slope increases as a function of *inefficiency* in the curves 212, 214, 216 and 218. In other words, the slope increases because more power is required to achieve incremental gains in cooling. Assuming a current thermal loading as indicated by points A in the graphs 202, 204, 206 and 208, the processing circuit 112 in step 410 of Fig. 4 would determine that the ACU 106₂ would handle the additional heat load resulting from additional processing best, compared to the ACUs 106₁, 106₃, and 106₄. The processing circuit 112 would make that determination because the slope of the curve 214 at point A is the minimum among those of curves 212, 214, 216 and 218.

Referring again to Fig. 4, after step 410, the processing circuit 112 executes step 415. In step 415, the processing circuit 112 causes the additional processing task to be allocated or assigned to one the server computers that is thermally proximate to (i.e. is cooled most by) the ACU 106_m that was identified in step 410. In the exemplary operation of step 410 discussed above in connection with Fig. 2, the processing circuit 112 in step 415 would cause assignment of the additional task to one of the server computers (e.g. servers 104₆ to 104₉) that is thermally proximate or thermally closest to the ACU 106₂.

As discussed above, the processing circuit 112 may perform step 410 (and step 415) proactively, in anticipation of new request. In such a case, the result of steps 410 and 415 is to identify where the next *n* processing tasks *will* be assigned once they are received. The processing circuit 112 stores the generated information and assigns incoming processing task requests accordingly.

In general, the processing task assignment operation of Fig. 4, which is based on incremental efficiency, can result in adding new tasks in an energy efficient manner. However, the operations of Fig. 4 may not result in the overall best efficiency. For example, the operations of Fig. 4 do not necessarily account for terminated processing tasks in a way that maximizes efficiency. Terminated tasks can reduce the thermal load in the data center. While the reduced thermal load due to termination of tasks is at least taken into account in step 410 (by considering the current thermal loading of the ACUs 106₁, 106₂, 106₃, 106₄), the current state of the thermal loading on the ACUs 106₁, 106₂, 106₃, 106₄ can be inefficient. The inefficiency can be due to the fact that terminations of tasks can drastically alter the thermal loading in ways that the mere incremental assignment of additional tasks cannot adequately compensate.

In addition, the operations of Fig. 4 presume that all (or at least a given set) of ACUs are currently on, and can be variably loaded to achieve an optimum result. In many cases, however, operation of less than all of the ACUs may be sufficient to handle the thermal load. The intelligent selection of which ACUs should be activated to handle certain levels of load can increase energy consumption efficiency.

To address the issues regarding the operations of Fig. 4, Fig. 5 shows a set of operations that may be used to allocate and/or reallocate *existing* application tasks to achieve thermal cooling efficiency. To this end, the operations of Fig. 5 determine which of the ACUs 106₁, 106₂, 106₃, 106₄ should be activated for various overall thermal load levels in the data center 102. Once a best or near best ACU configuration is determined, the application processes may be allocated to the server computers 104₁ to 104₁₈ accordingly.

The steps of Fig. 5 may, but need not, be executed on a computer server management

system 108 that also employs the steps of Fig. 4. The operations of Fig. 5 may be used as a part of step 310 of Fig. 3. To this end, the operations described below use load modeling to determine a most (or approximately most) efficient allocation of heat load (and consequently processing load) in the data center 102. The result of the operations of Fig. 5 identifies a best configuration ACUs to be turned on for a given thermal load TL for the overall data center 102. The computer server management system 108 may then assign or re-assign computer processing tasks among the server computers 104₁ to 104₁₈ in accordance with the identified ACU configuration.

Thus, in one example, the processing circuit 112 may first obtain from the BAS element 120 the current thermal load TL in the data center, based on loading of the ACUs 106₁, 106₂, 106₃, 106₄. The processing circuit 112 may then execute the operations of Fig. 5 to determine if a *better* allocation of the existing processing tasks can achieve a more efficient loading of the ACUs 106₁, 106₂, 106₃, 106₄.

Alternatively, the processing unit 112 may predetermine the best ACU configurations for a plurality of given thermal loads, and then store the results. Thus, the processing unit 112 may have stored in the memory 110 a plurality of data points of best (or near best) thermal load allocation among the ACUs 106₁, 106₂, 106₃, 106₄ for a plurality of data center thermal load levels. Then, from time to time, the processing circuit 112 obtains the current thermal load TL and identifies a best thermal load model based on the stored data points.

As a consequence, the operations of Fig. 5 may be generated prior to actual data center operations, such that the results are stored for future use. Alternatively, the operations of Fig. 5 may be generated in quasi-real time, in response to receipt of actual calculated thermal load TL .

In step 505, the processing circuit 112 selects a configuration of ACUs to be activated for a given overall thermal load TL for the data center 102. In step 510, the processing circuit 112 determines a load ratio of the selected ACU configuration, LR_{comb} , by dividing the TL value by the full loading capacity of the selected ACU configuration, $MAXL_{comb}$. In other words,

$$LR_{comb} = TL / MAXL_{comb}$$

In step 515, the processing circuit 112 determines whether selected configuration is a valid combination based on the load ratio value LR_{comb} . For example, if the configuration cannot adequately cool the data center 102 under the conditions of thermal load TL , then selected ACU configuration is not valid. In the current embodiment, the processing circuit 112 carries out step 515 by determining if the calculated load ratio LR_{comb} is between 0.3 and 1.0. If so, then the processing circuit 112 proceeds to step 520. If not, then the ACU configuration is not valid and the processing circuit 112 proceeds to step 525, discussed further below.

To reach step 520, it has been determined that the selected ACU configuration/combination is valid. In step 520, the processing circuit 112 determines the power draw for the selected ACU combination, PD_{comb} . To this end, the processing circuit 112 may first determine, based on the efficiency characteristics for the ACUs 106₁, 106₂, 106₃, 106₄ (e.g. curves 212, 214, 216 and 218 of Fig. 2), the best allocation of the thermal load TL among the selected combination of ACUs. Using this best allocation of the load TL among the selected ACU combination, the processing circuit 112 then calculates the overall power draw for the selected combination, PD_{comb} , by adding the individual predicted power draws for each of the ACUs in the selected ACU combination.

Thereafter, in step 530, the processing circuit 112 determines whether the power draw for the selected combination is lower than the current lowest power draw for all previously evaluated ACU combinations. In other words, the processing circuit 112 determines whether $PD_{comb} < PD_{bestcomb}$ wherein $PD_{bestcomb}$ is the power draw of the best combination of ACUs calculated thus far, for the current TL . If the answer in step 530 is negative, then the processing circuit 112 proceeds to step 525, discussed further below. If the answer in step 530 is positive, however, then the processing circuit 112 in step 535 stores the selected combination of ACUs as the current best combination for the load TL , and sets $PD_{bestcomb} = PD_{comb}$. The processing circuit 112 then proceeds to step 525.

In step 525, the evaluation of the selected ACU configuration/combination for the data center load level TL is complete. In step 525, the processing circuit 112 determines whether there are any other combinations of ACUs to be evaluated. If so, then the processing circuit 112 returns to step 505 to select another configuration of ACUs. If not, then the evaluation is complete, and the processing circuit 112 in step 535 stores the ACU combination associated with $PD_{bestcomb}$ as the appropriate ACU combination for the data center thermal load level TL .

The processing circuit 112 may subsequently allocate application processes such that the thermal load TL will be borne by the combination of ACUs that correspond to the $PD_{bestcomb}$. In addition, as per step 320 of Fig. 3, the processing circuit 112 informs the BAS element 120 of the combination of ACUs that should be activated based on the current thermal load TL .

In another embodiment, a set of operations may be performed from time to time to determine whether a reallocation of application processes among the server computers 104₁

to 104₁₈ may allow for one of the currently operating ACUs 106_n to be shut down. The process involves calculating the maximum cooling capability (at an acceptable or maximum efficiency) of each online ACU unit. The total of the maximum capability from the on-line ACUs is then compared to the current actual thermal load to determine the spare capacity of the on-line ACUs. If the load of the least efficient on-line ACU drops below the spare capacity for a given time period, then the least efficient on-line ACU is shut down, and the processing tasks are reallocated among the server computers 104₁ to 104₁₈ accordingly.

It will be appreciated that the above described embodiments are merely illustrative, and that those of ordinary skill in the art may readily devise their own implementations and modifications that incorporate the principles of the present invention and fall within the spirit and scope thereof.

CLAIMS:

1. A method, comprising:

a) obtaining efficiency characteristic information for each of a plurality of air conditioning units in a location containing a plurality of server computers, wherein the efficiency characteristic information is representative of an efficiency performance curve for a range of variable cooling output; and

b) employing one or more processing units to allocate one or more processing tasks to one of the plurality of server computers based on the efficiency characteristic information and to:

i) identify for a total thermal load level of each of said plurality of server computers a combination of air conditioning units capable of cooling the total thermal load level using less energy than a plurality of other combinations of air conditioning units capable of cooling the total thermal load level;

ii) allocate processing tasks to server computers based on their thermal proximity to the identified combination of air conditioning units.

2. The method of claim 1, wherein step b) further comprises employing the one or more processing units to allocate a first processing request to a first of the plurality of computers based on a thermal proximity of the first of the plurality of computers to a select one of the plurality of air conditioning units.

3. The method of claim 2, wherein the one or more processing units identifies the select one of the plurality of air conditioning units based on the efficiency characteristic information.

4. The method of claim 1, further comprising:

employing the one or more processing units to allocate one or more processing tasks to a first of the plurality server computers based in part on a distance of the first of the

plurality of server computers from a first air conditioning unit, and causing the first air conditioning unit to reduce a cooling output thereof.

5. The method of claim 4, further comprising causing the first air conditioning unit to be turned off.

6. The method of claim 4, wherein the first of the plurality of server computers is more distant than at least one other server computer from the first air conditioning unit.

7. The method of claim 1, further comprising shutting down air conditioning units outside of the identified combination of air conditioning units responsive to the completion of substep ii).

8. An arrangement, comprising:

a computer server management system having a processing circuit and a memory, the processing circuit configured to coordinate the usage of a plurality of server computers, the processing circuit executing virtualization software configured to manage application processing on the plurality of server computers;

the memory storing efficiency characteristic information for a plurality of air conditioning units in a location containing a plurality of server computers, wherein the efficiency characteristic information is representative of an efficiency performance curve for a range of variable cooling output;

wherein the processing circuit is further configured to allocate one or more processing tasks to one of the plurality of server computers based on the efficiency characteristic information and to:

i) identify for a total thermal load level of each of said plurality of server computers a combination of air conditioning units capable of cooling the total thermal load level using less energy than a plurality of other combinations of air conditioning units capable of cooling the total thermal load;

ii) allocate processing tasks to server computers based on their thermal proximity to the identified combination of air conditioning units.

9. The arrangement of claim 8, wherein the processing circuit is further configured to allocate a first processing request to a first of the plurality of computers based on a thermal proximity of the first of the plurality of computers to a select one of the plurality of air conditioning units.

10. The arrangement of claim 9, wherein the processing circuit is further configured to identify the select one of the plurality of air conditioning units based on the efficiency characteristic information.

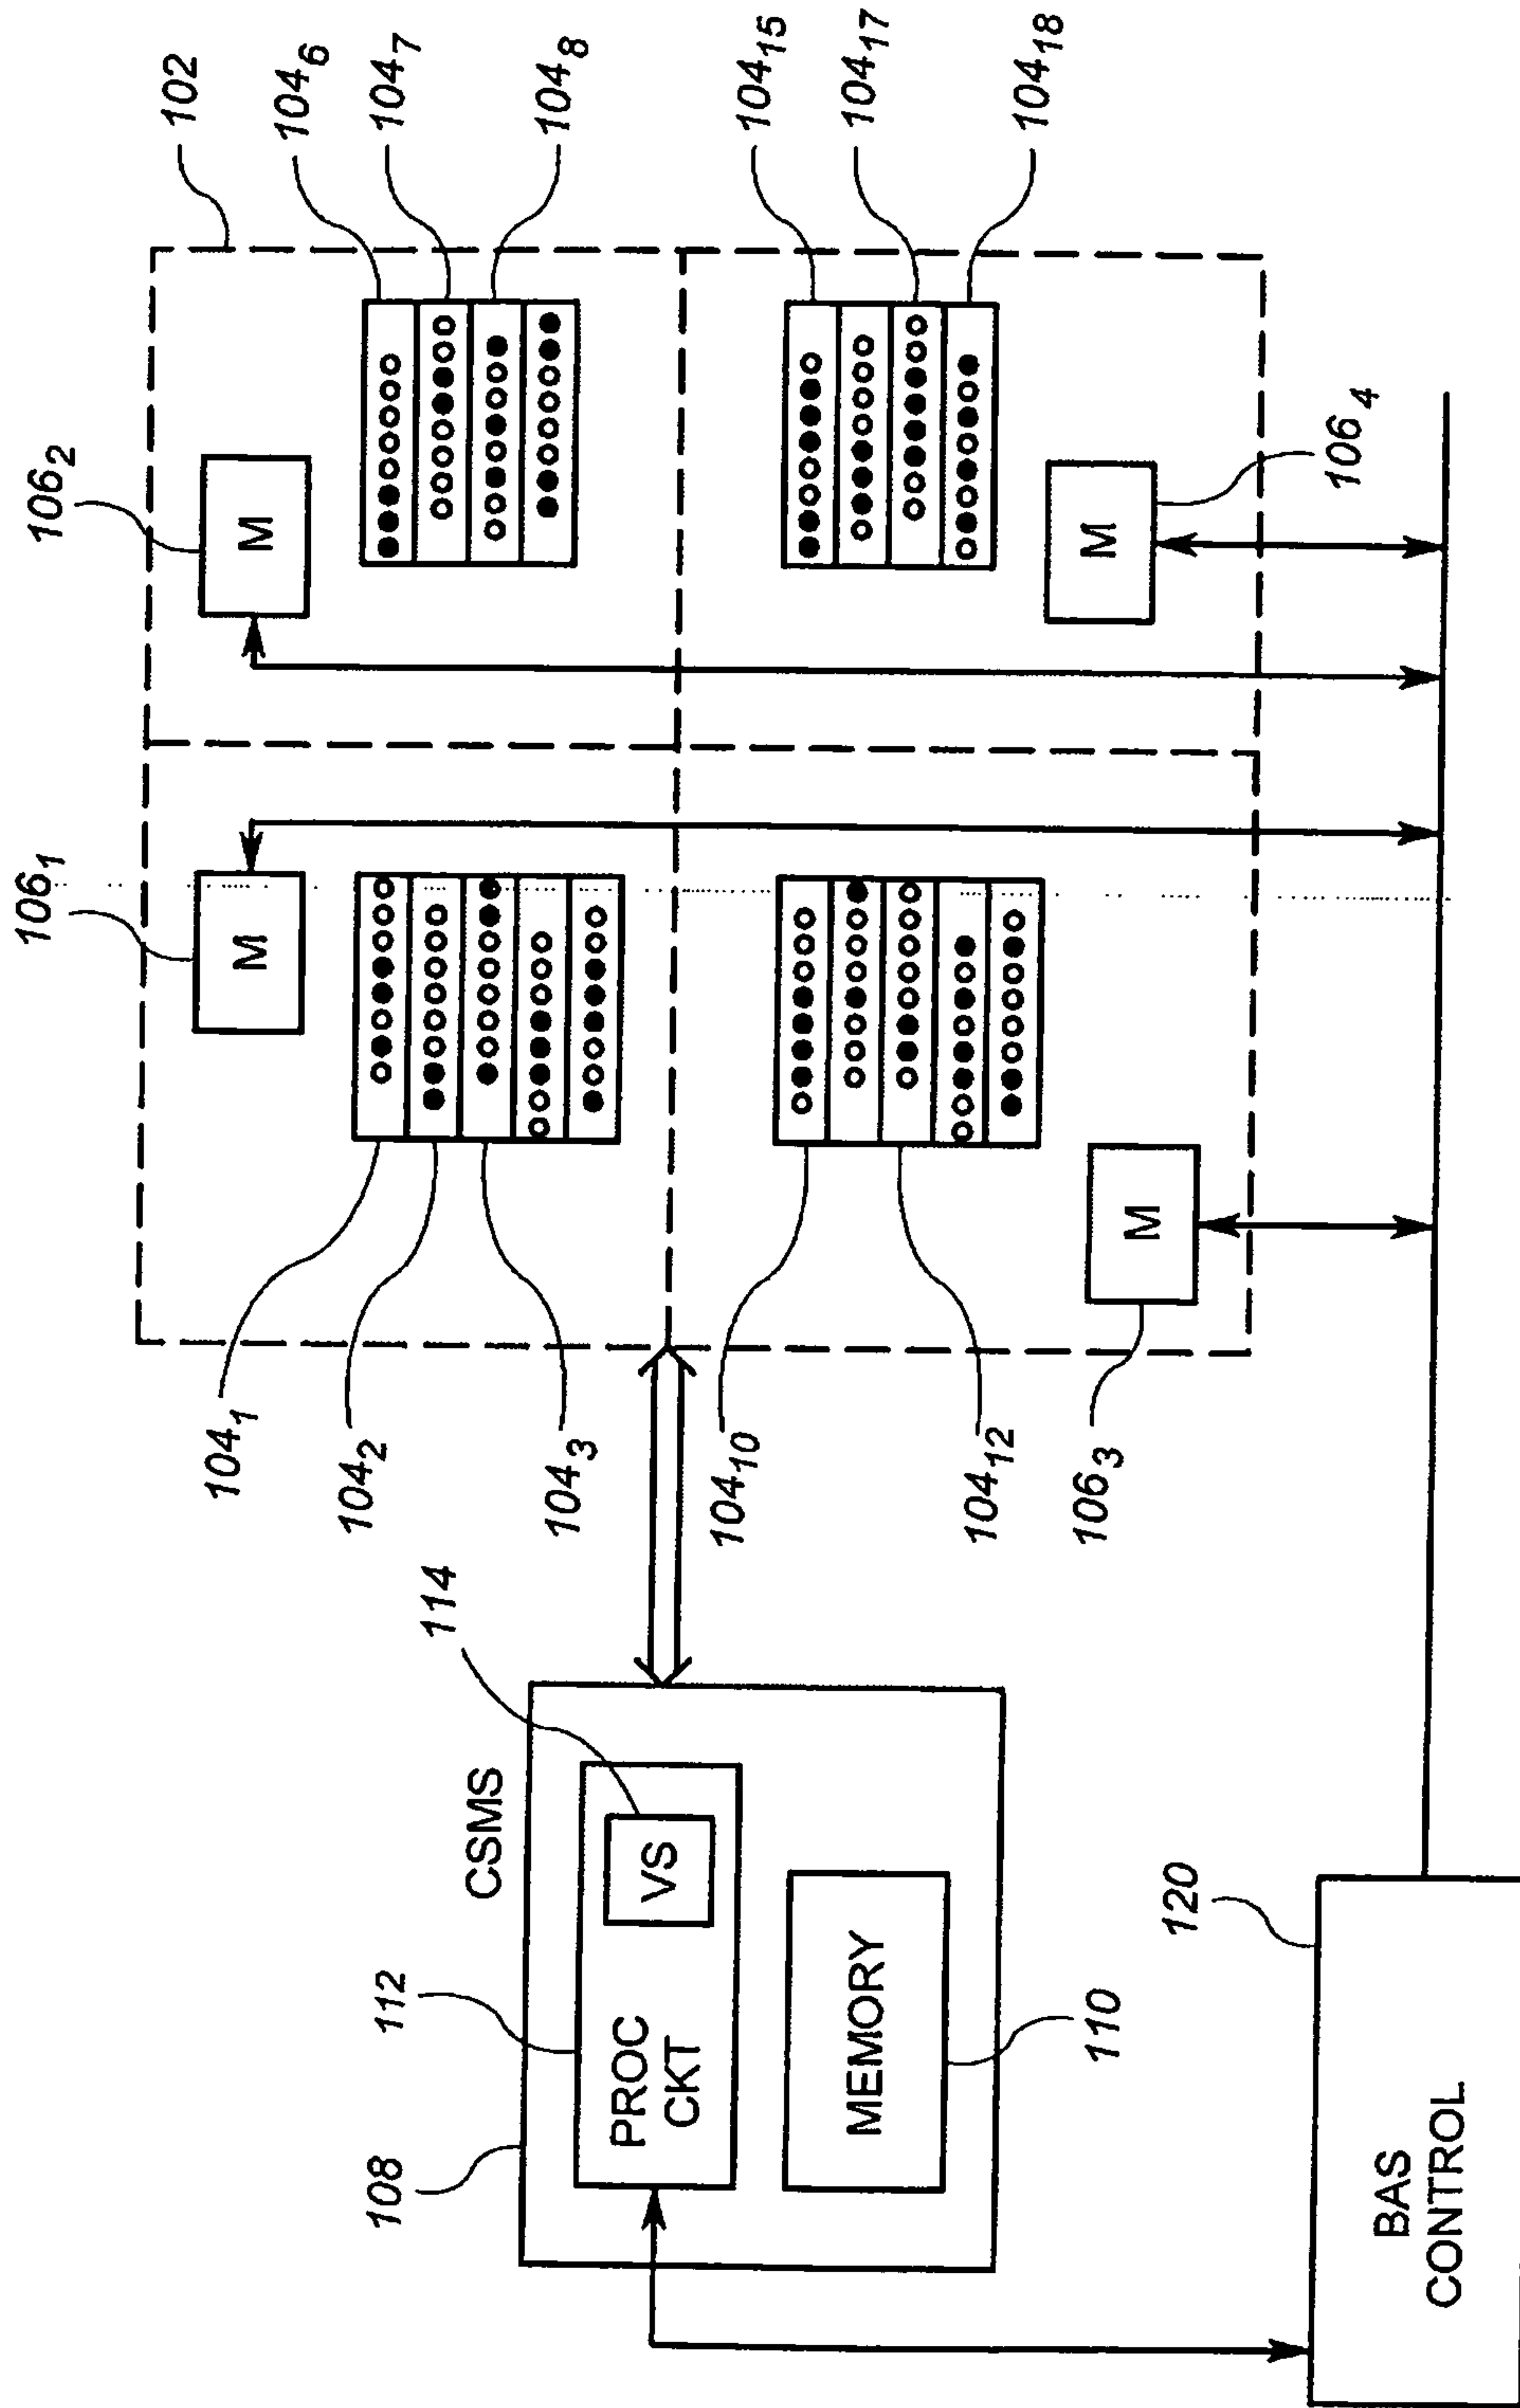


FIG. 1

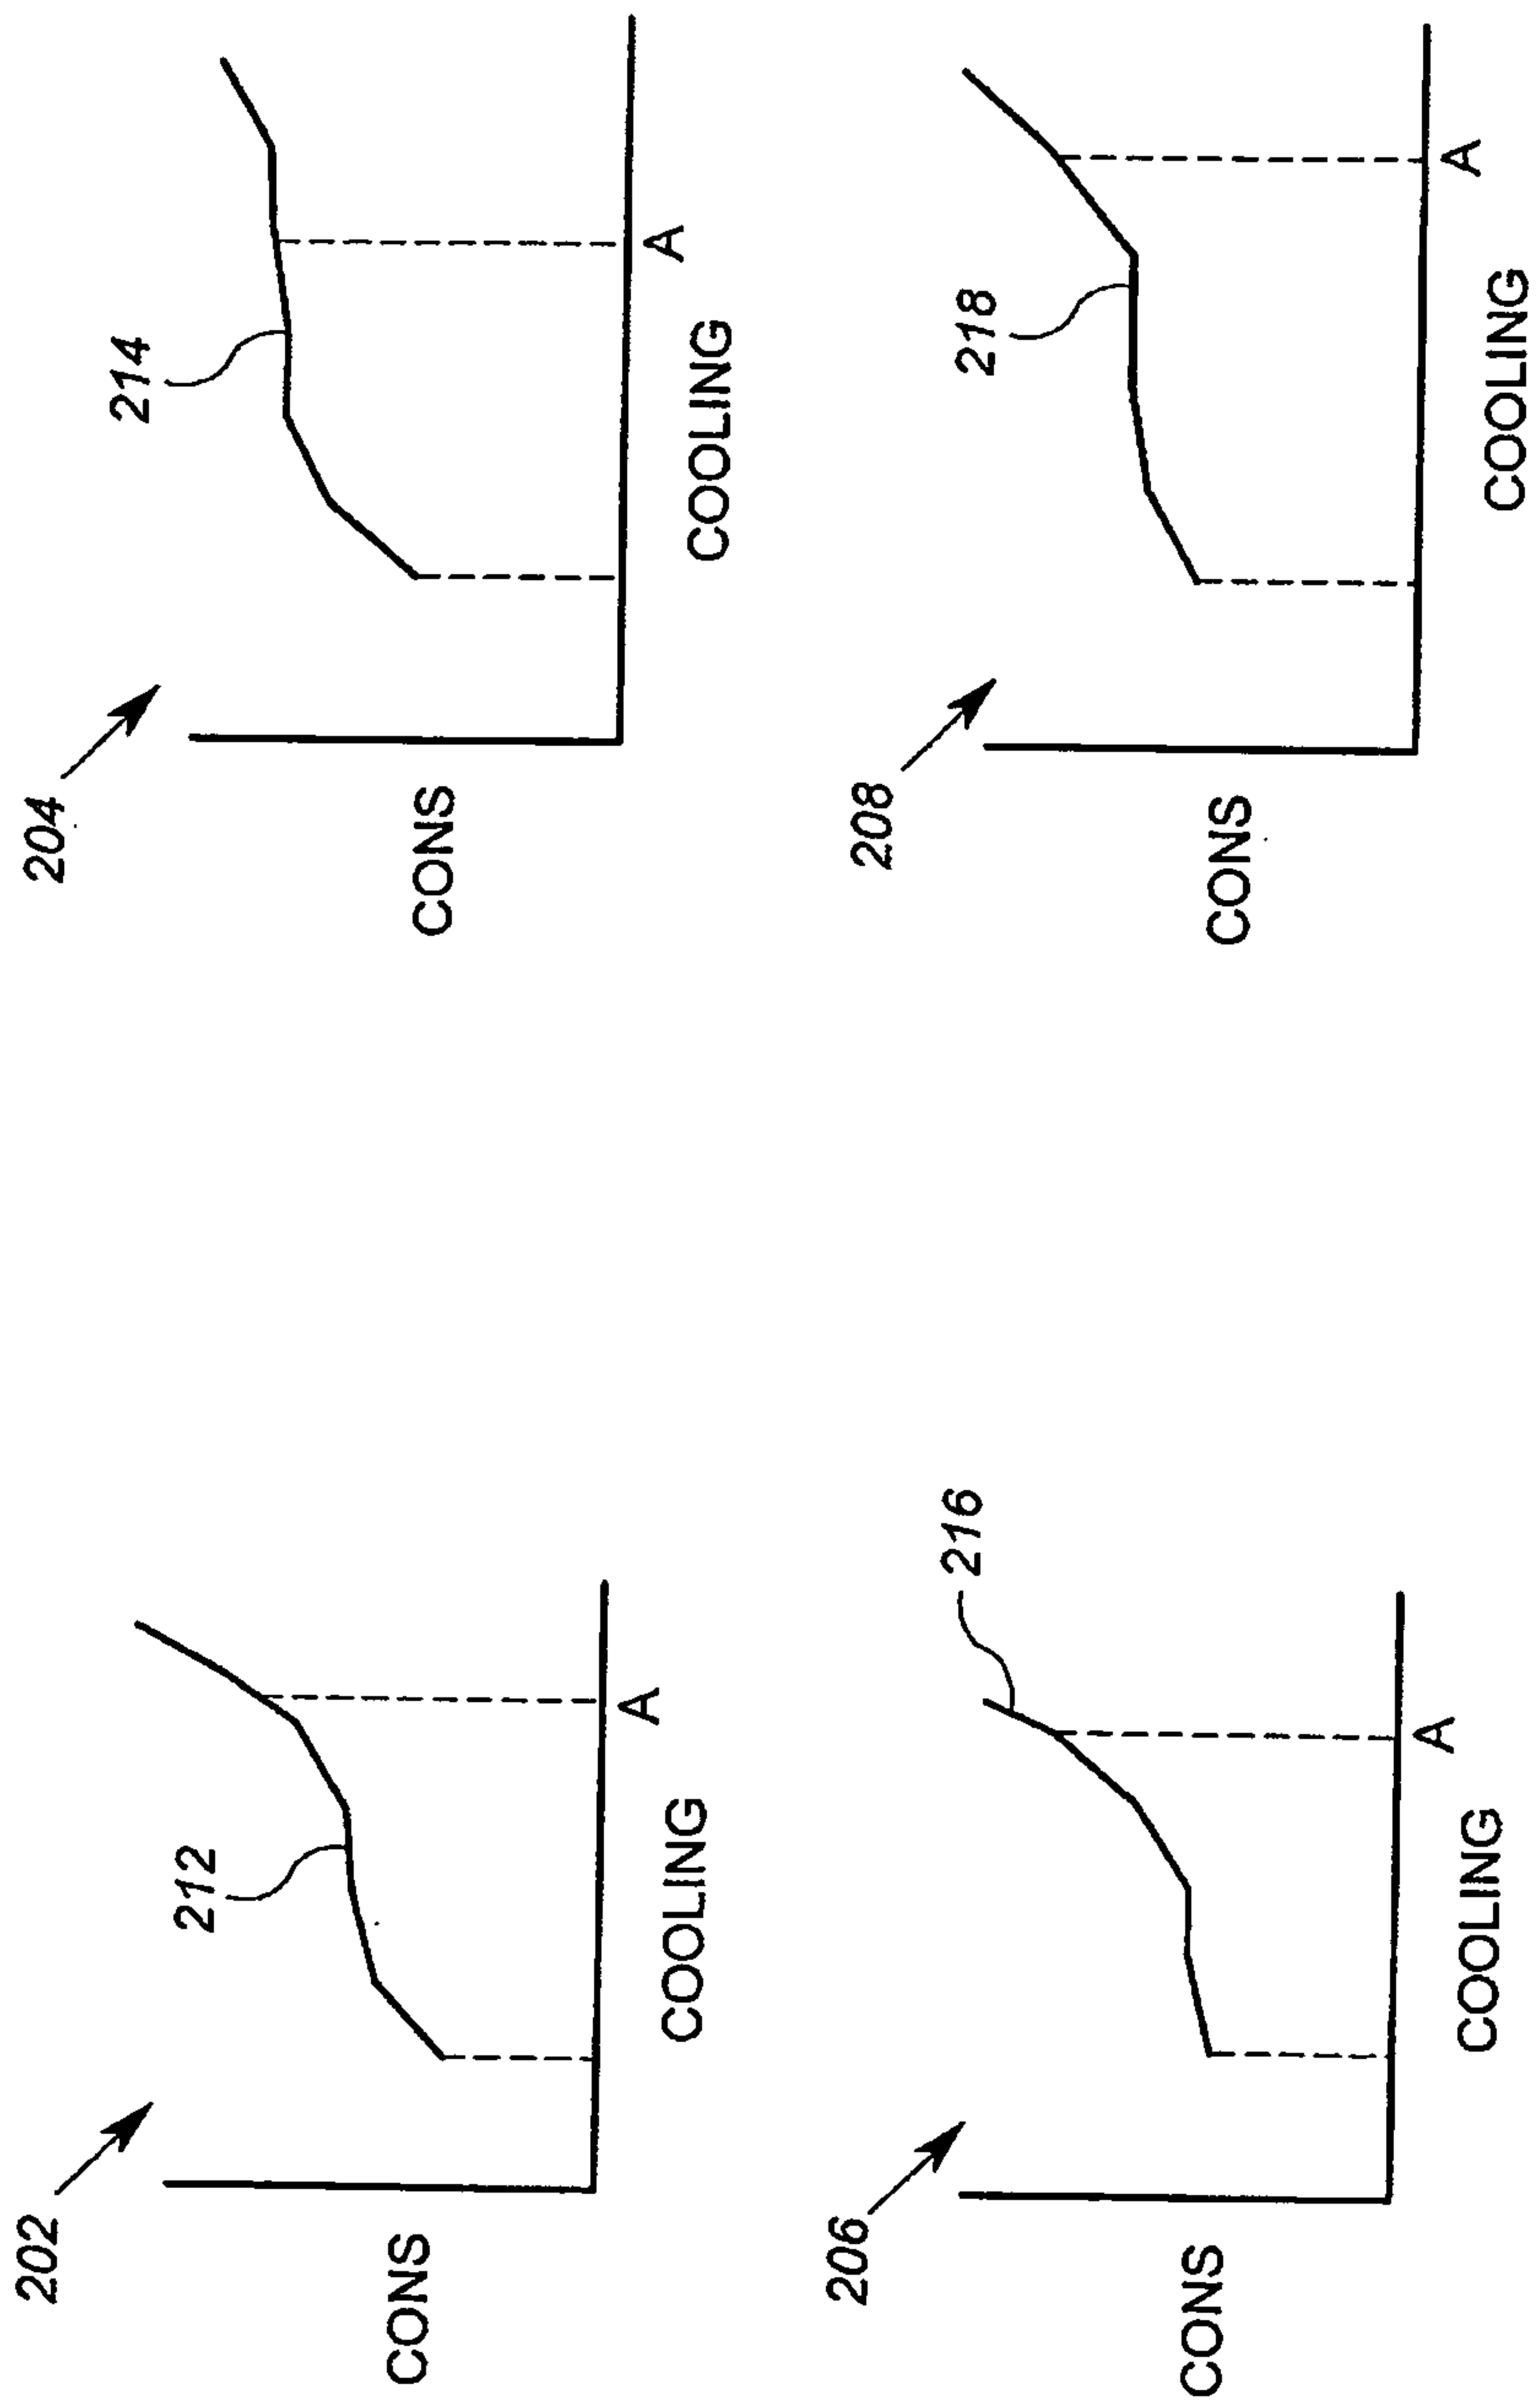


FIG. 2

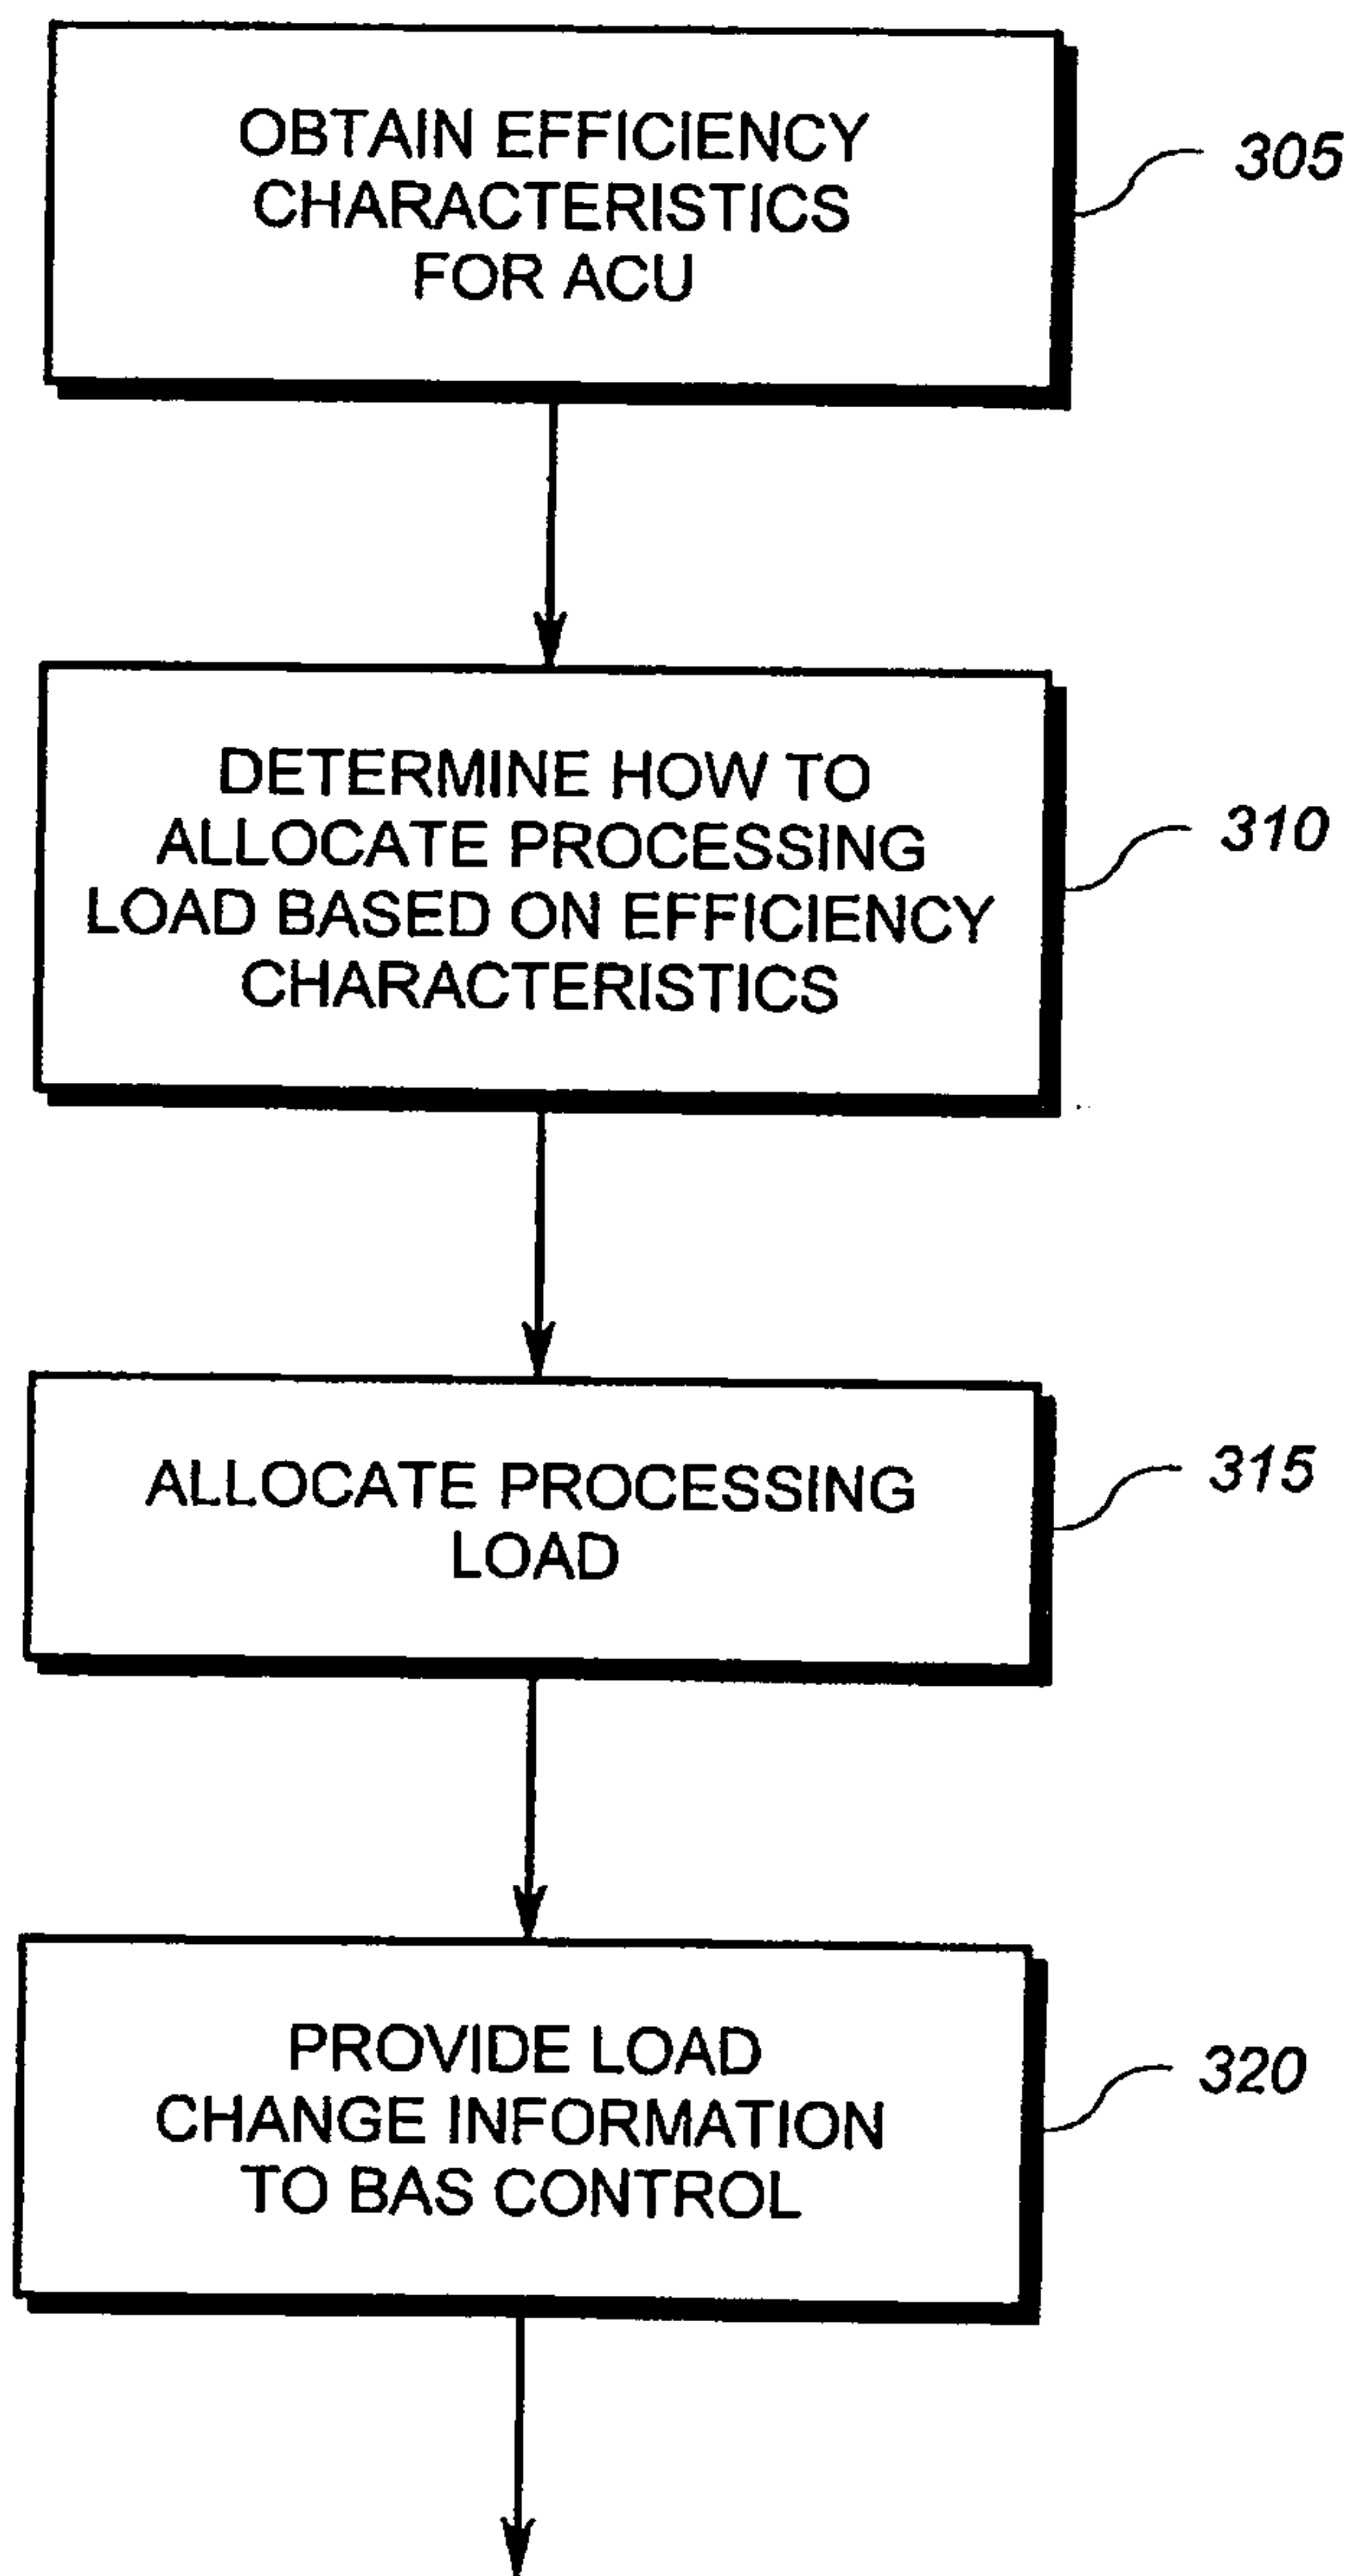


FIG. 3

4 / 5

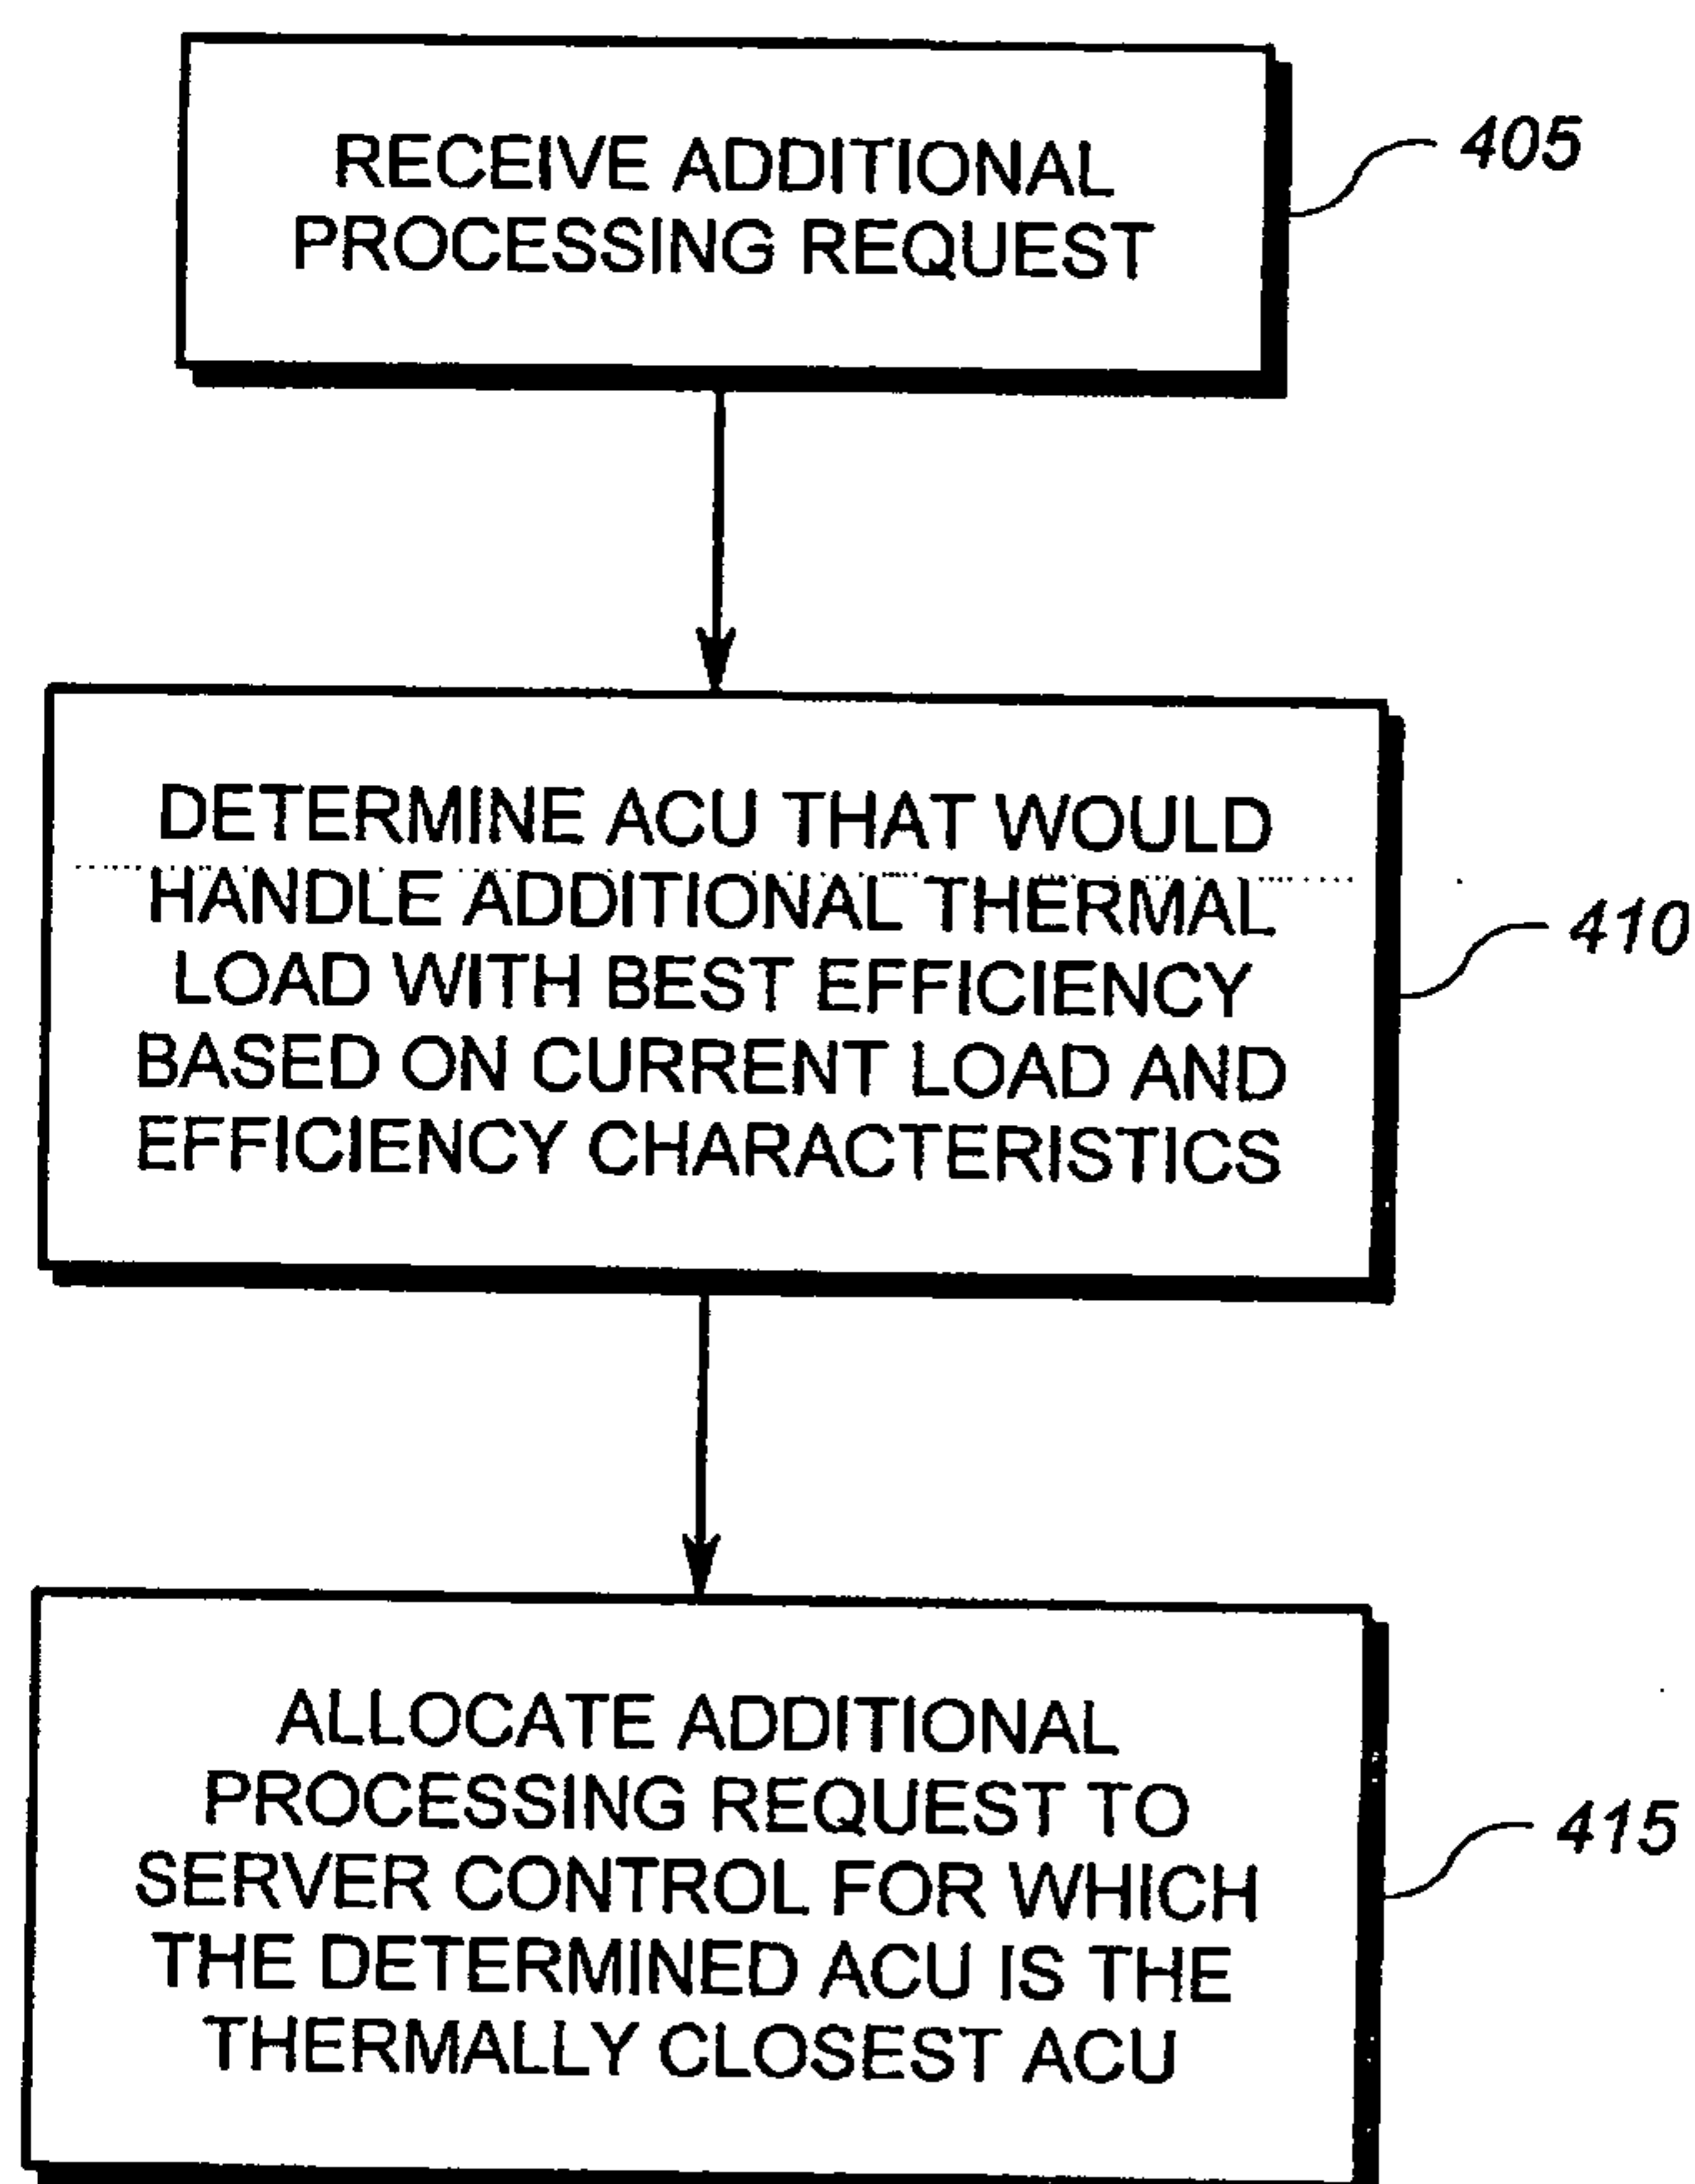


FIG. 4

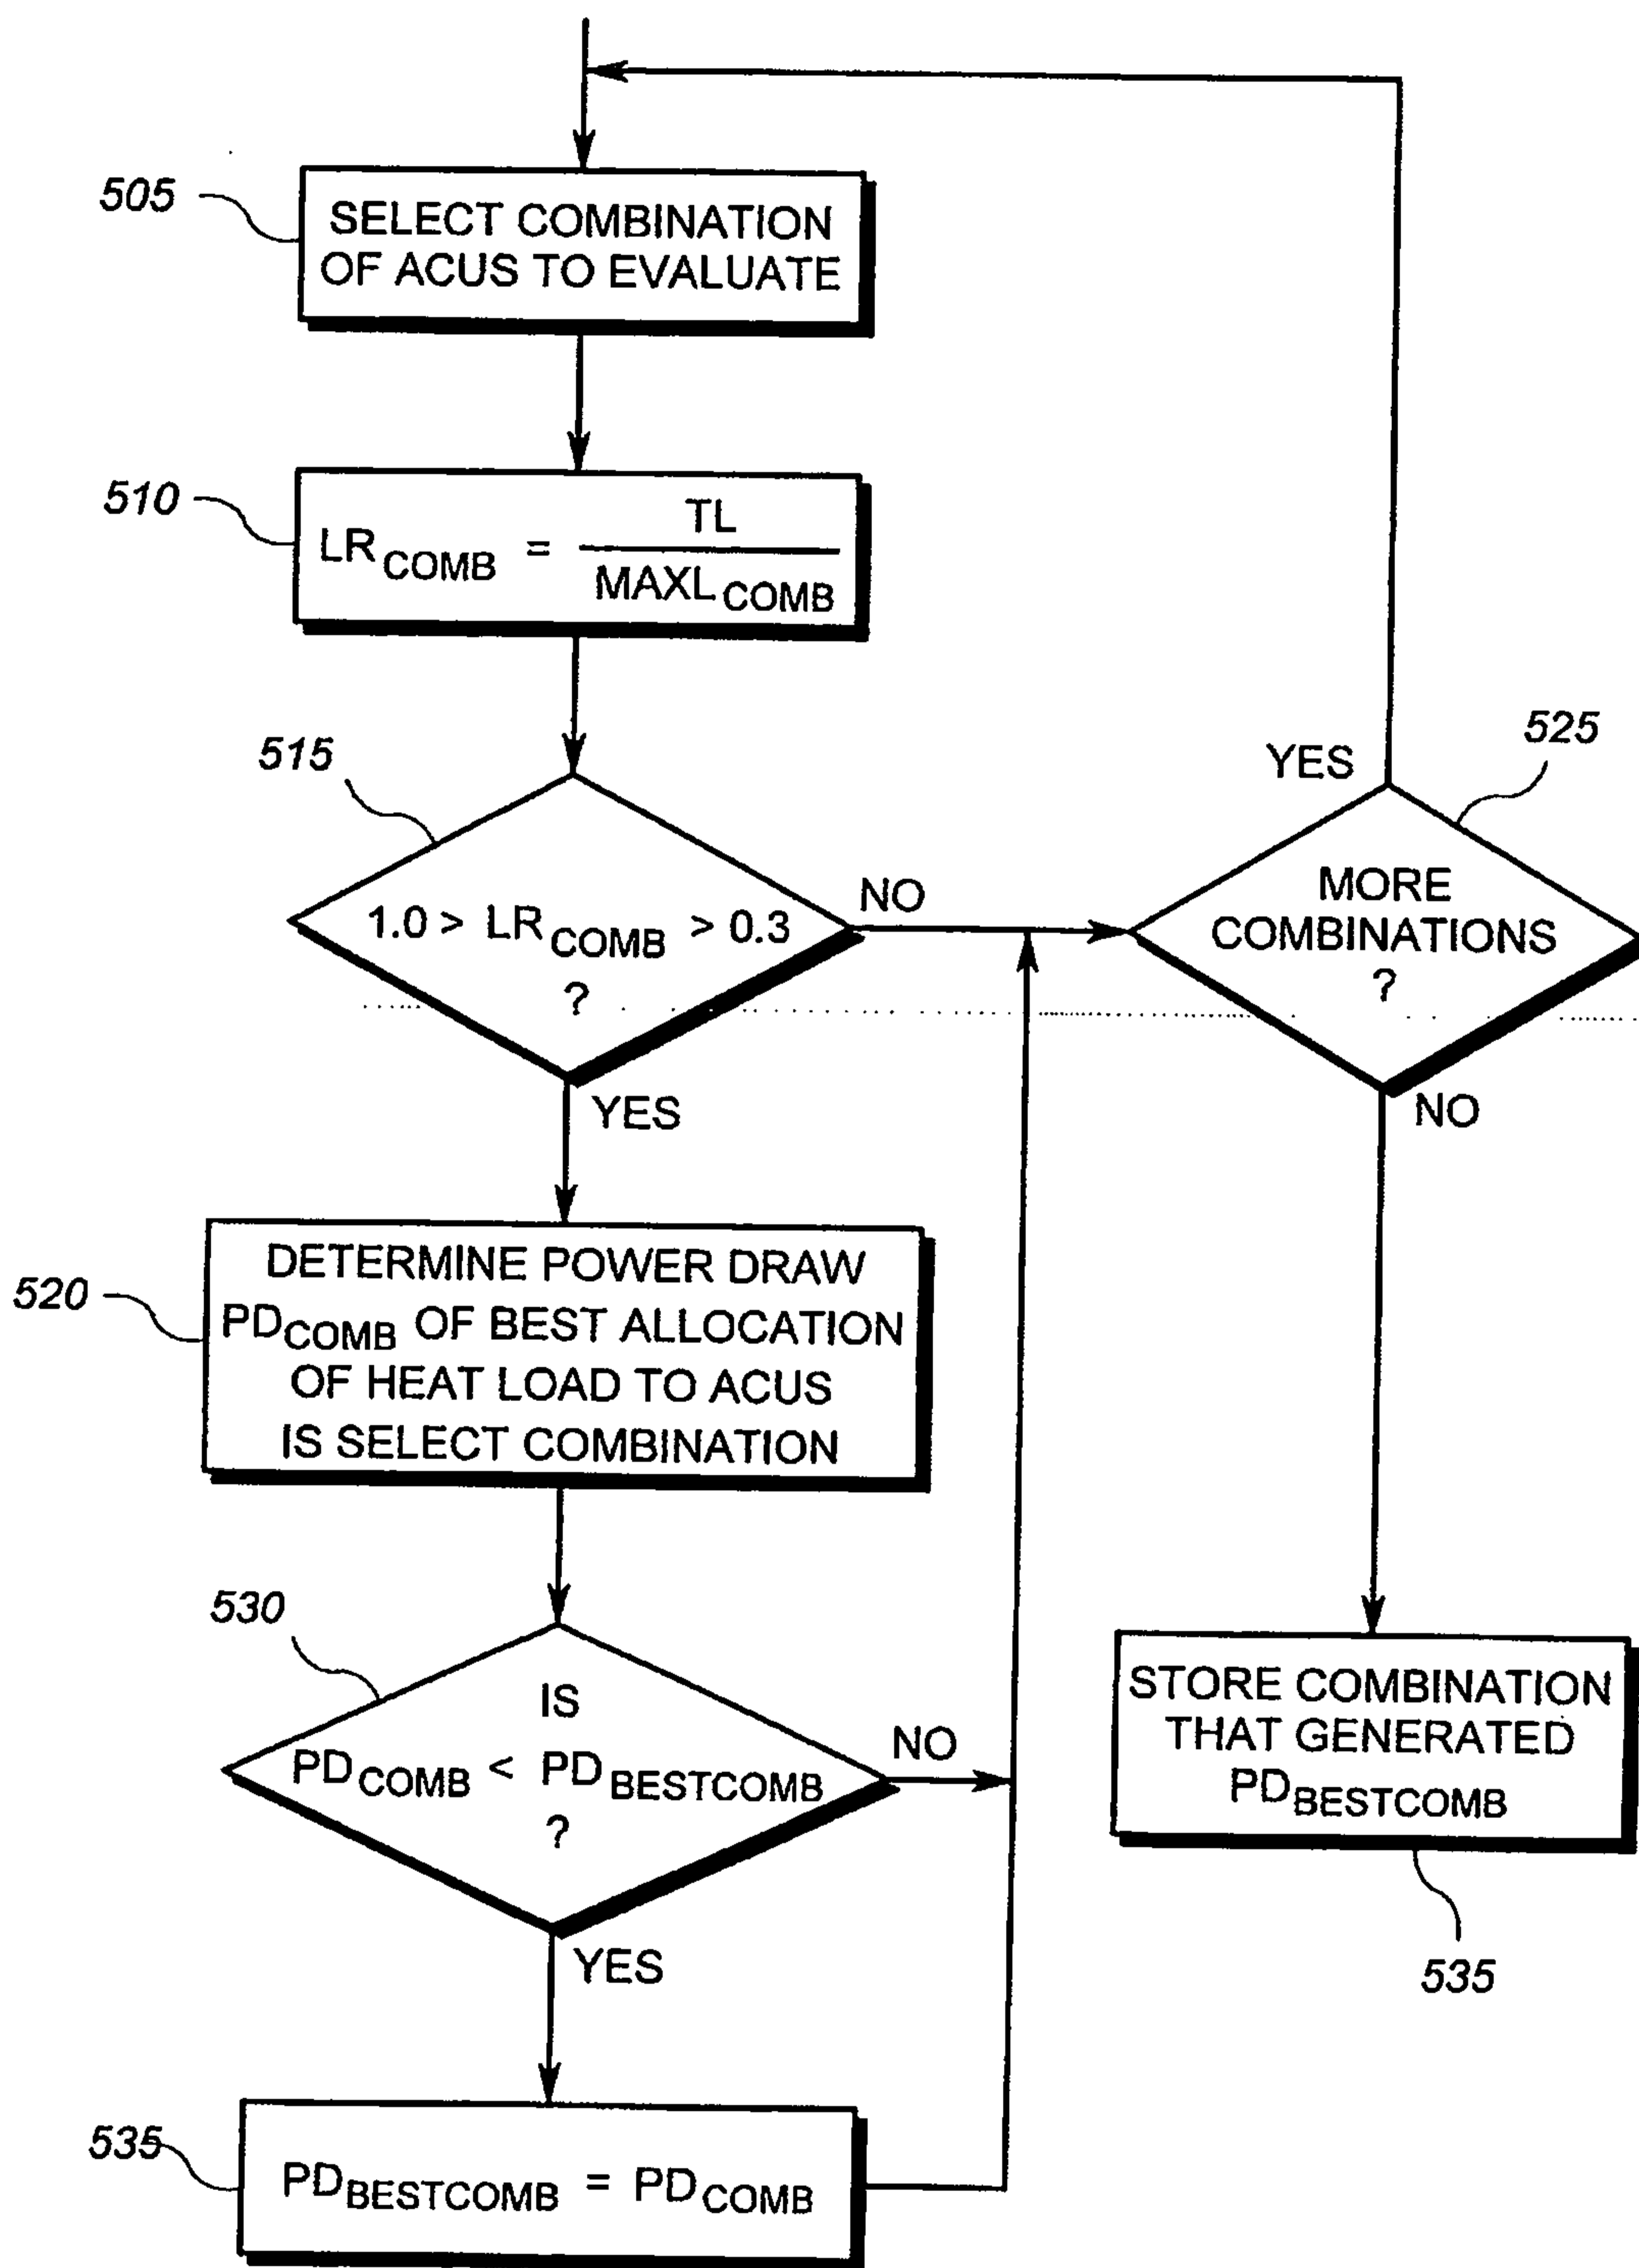


FIG. 5

