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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
9 July 2015 (09.07.2015)

WIPO | PCT

(10) International Publication Number
WO 2015/102797 A1

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

(21) International Application Number:
PCT/US2014/068352

(22) International Filing Date:
3 December 2014 (03.12.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
14/145,915 31 December 2013 (31.12.2013) US

(71) Applicant: **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, WA 98052-6399 (US).

(72) Inventors: **GARDEN, Euan, Peter**; c/o Microsoft Corporation, One Microsoft Way, LCA - International Patents (8/1172), Redmond, WA 98052-6399 (US). **JUSTICE, John, Raymond**; c/o Microsoft Corporation, One Microsoft Way, LCA - International Patents (8/1172), Redmond, WA 98052-6399 (US). **SHARMA, Madhumitra**; c/o Microsoft Corporation, One Microsoft Way, LCA - International Patents (8/1172), Redmond, WA 98052-6399 (US).

(74) Agents: **HOWARD, Jason, O.** et al.; (USOC - Shook, Hardy & Bacon), One Microsoft Way, Microsoft Corporation, LCA - International Patents (8/1172), Redmond, WA 98052-6399 (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.

(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 applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

[Continued on next page]

(54) Title: MULTIMODE GAMING SERVER

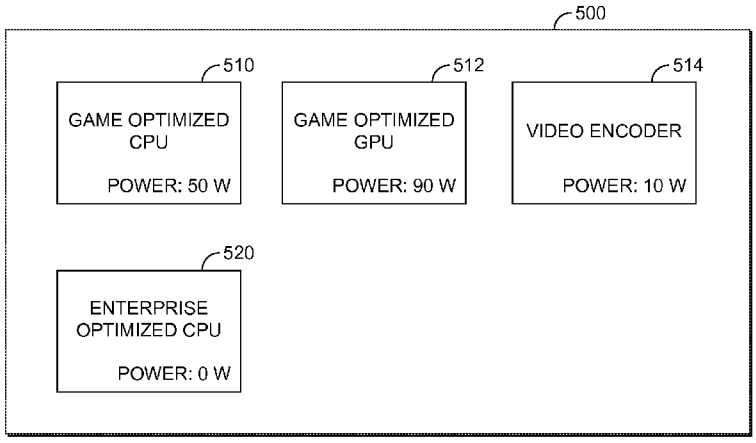

FIG. 5.

(57) Abstract: Aspects of the present invention relate to a multimode gaming server with different types of computing resources provided within the server. The different computing resources can be optimized for different computing tasks. For example, a first type of resource can be optimized for producing high definition graphics and a second type of resource for enterprise computing. Each resource may be activated or deactivated as demand for different computing tasks change throughout the day. In one aspect, the resources are different chip sets in different mother board sockets. In one aspect, provisioning of the other components (e.g., cooling, power supply, network bandwidth) in the multimode server is not adequate for both computing resources to run simultaneously.

WO 2015/102797 A1

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

MULTIMODE GAMING SERVER

BACKGROUND OF THE INVENTION

Generally, the servers selected for deployment in a data center can perform a wide range of computing tasks, but may not perform some specialized computing tasks very efficiently. For example, video intensive compute projects are better performed on a server with powerful GPU and video encoders. An enterprise server may be able to perform some video related work through the CPU, but the work may be inefficient. On the other hand, a server designed for video related work may perform general computing projects inefficiently.

SUMMARY OF THE INVENTION

This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the detailed description. This summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used in isolation as an aid in determining the scope of the claimed subject matter.

Aspects of the present invention relate to a multimode gaming server with different types of computing resources provided within the server. The different computing resources can be optimized for different computing tasks. For example, a first type of resource can be optimized for producing high definition graphics and a second type of resource for enterprise computing. Each resource may be activated or deactivated as demand for different computing tasks change throughout the day. In one aspect, the resources are different chip sets in different mother board sockets. In one aspect, provisioning of the other components (e.g., cooling, power supply, network bandwidth) in the multimode server is not adequate for both computing resources to run simultaneously.

In one aspect, the cooling capacity for the server is intentionally sized to be incapable of providing enough cooling to maintain an acceptable operational temperature for the multimode server, if all of the computing resources in the server are simultaneously in an active processing mode. The data center's control fabric can maintain an acceptable operating temperature within the server by assigning workloads to only one type of computing resources within the multimode server at a given point in time. For example, at any given

time only the gaming optimized compute resource in a server may be assigned a workload and be in an active processing state. The remaining compute resources in the servers are set to a low-power state.

In one aspect, different computing resources within the multimode server have a similar maximum power usage when in an active processing mode. For example, a gaming optimized resource (e.g., a GPU and specialty CPU) having a maximum power usage of 150 W can be deployed in the same multimode server as an enterprise CPU having a 150 W capacity.

In one aspect, computing resources are selected for inclusion in a multimode server according to anticipated peak periods of usage. Resources designed for specialized workload having a peak period of usage that differs from each other can be included in a multimode gaming server. For example, a first type of computing resources associated with a specialized workload with peak hours from 4 PM to 12 PM can be matched with a second type of computing resource having a specialized workload with peak hours from 6 AM to 2 PM. In other words, during a given period either a first type or a second type of computing resource within a server will be in a high demand.

BRIEF DESCRIPTION OF THE DRAWING

Aspects of the invention are described in detail below with reference to the attached drawing figures, wherein:

FIG. 1 is a block diagram of an exemplary computing environment suitable for implementing aspects of the invention;

FIG. 2 is a diagram depicting a gaming environment, in accordance with an aspect of the present invention;

FIG. 3 is a diagram depicting a remote gaming environment having one or more data centers with an nonhomogeneous arrangement of gaming servers and general purpose servers, in accordance with an aspect of the present invention;

FIG. 4 is a diagram depicting an arrangement of multimode gaming servers in various modes, in accordance with an aspect of the present invention;

FIG. 5 is a diagram depicting a mother board with active gaming resources, in accordance with an aspect of the present invention;

FIG. 6 is a diagram depicting a mother board with active general purpose resources, in accordance with an aspect of the present invention; and

FIG. 7 is a diagram depicting a method for managing workloads within a data center, in accordance with an aspect of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

The subject matter of aspects of the invention is described with specificity herein to meet statutory requirements. However, the description itself is not intended to limit the scope of this patent. Rather, the inventors have contemplated that the claimed subject matter might also be embodied in other ways, to include different steps or combinations of steps similar to the ones described in this document, in conjunction with other present or future technologies. Moreover, although the terms "step" and/or "block" may be used herein to connote different elements of methods employed, the terms should not be interpreted as implying any particular order among or between various steps herein disclosed unless and except when the order of individual steps is explicitly described.

Aspects of the present invention relate to a multimode gaming server with different types of computing resources provided within the server. The different computing resources can be optimized for different computing tasks. For example, a first type of resource can be optimized for producing high definition graphics and a second type of resource for enterprise computing. Each resource may be activated or deactivated as demand for different computing tasks change throughout the day. In one aspect, the resources are different chip sets in different mother board sockets. In one aspect, provisioning of the other components (e.g., cooling, power supply, network bandwidth) in the multimode server is not adequate for both computing resources to run simultaneously.

In one aspect, the cooling capacity for the server is intentionally sized to be incapable of providing enough cooling to maintain an acceptable operational temperature for the server, if all of the computing resources in the server are simultaneously in an active processing mode. The data center's control fabric can maintain an acceptable operating temperature within the server by assigning workloads to only one type of computing resources within the multimode server at a given point in time. For example, at any given time only the gaming optimized compute resource in a server may be assigned a workload

and be in an active processing state. The remaining compute resources in the servers are set to a low-power state.

In one aspect, different computing resources within the multimode server have a similar maximum power usage when in an active processing mode. For example, a gaming optimized resource (e.g., a GPU and specialty CPU) having a maximum power usage of 150 W can be deployed in the same multimode server as an enterprise CPU having a 150 W capacity. Even though the total power usage of the two types of computing resources is similar, the distribution of power usage in the resource can differ greatly. For example, the gaming optimized resource may have a graphics processing unit ("GPU") that uses 100 W and a central processing unit ("CPU") that uses 50 W. The enterprise resource may not have a GPU, but could have a more powerful CPU that consumes 150 W.

In one aspect, computing resources are selected for inclusion in a multimode server according to anticipated peak periods of usage. Resources designed for specialized workload having a peak period of usage that differs from each other can be included in a multimode gaming server. For example, a first type of computing resources associated with a specialized workload with peak hours from 4 PM to 12 PM can be matched with a second type of computing resource having a specialized workload with peak hours from 6 AM to 2 PM. In other words, during a given period either a first type or a second type of computing resource within a server will be in a high demand.

As used herein, a "gaming optimized computing resource" is adapted to output a rendered video game image to a client device, such as a game console. The video game image may be rendered as a streaming video communicated to the client. In order to render a high quality video game image, a gaming optimized computing resource can have a graphics processing unit that is more powerful than a graphics processing unit, if any, found in a general-purpose computing resource. The gaming optimized computing resource may also have dedicated video encoding capabilities.

Power consumption can be used as a proxy for a processor's capabilities. In one aspect, a gaming optimized computing resource can be defined by the inclusion of a GPU that consumes greater than a threshold percentage of power used by the game optimized computing resource during peak power consumption. In one aspect, the threshold percentage of power is greater than 40% of peak power, for example greater than 50%, for example greater than 60%, for example greater than 70%, or for example greater than 80%. For

example, a GPU in a gaming optimized computing resource could use 100 W, with a total peak power usage (e.g., GPU and CPU) of 150 W in the gaming optimized server resource.

As used herein, the terms "general-purpose computing resource" or "general processing optimized computing resource" describes a resource designed to emphasize computing process typically associated with a central processing unit. General-purpose computing resources can be capable of performing specialized computing processes, but may not be optimized for that purpose. For example, a CPU can perform graphics processing less efficiently than the same or similar tasks can be performed by a GPU.

Aspects of the present invention may transition various types of computing resources between different power modes or states. As used herein, the term "low-power mode" means a resource is presently operating at less than 20% of the resource's maximum rate of power. As an example, a resource in low-power mode may be shut off at the motherboard socket, but able to respond to a power-on command.

As used herein, the phrase "active processing mode" means a computing resource is actively processing a computing workload. A computing resource in active processing mode can be using greater than 20% of the resource's maximum rated power.

Having briefly described an overview of aspects of the invention, an exemplary operating environment suitable for use in implementing aspects of the invention is described below.

Exemplary Operating Environment

Referring to the drawings in general, and initially to FIG. 1 in particular, an exemplary operating environment for implementing embodiments of the invention is shown and designated generally as computing device 100. Computing device 100 is but one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computing device 100 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated.

The invention may be described in the general context of computer code or machine-useable instructions, including computer-executable instructions such as program components, being executed by a computer or other machine, such as a personal data assistant or other handheld device. Generally, program components, including routines, programs, objects, components, data structures, and the like, refer to code that performs

particular tasks or implements particular abstract data types. Embodiments of the invention may be practiced in a variety of system configurations, including handheld devices, consumer electronics, general-purpose computers, specialty computing devices, etc. Embodiments of the invention may also be practiced in distributed computing environments where tasks are performed by remote-processing devices that are linked through a communications network.

With continued reference to FIG. 1, computing device 100 includes a bus 110 that directly or indirectly couples the following devices: memory 112, one or more processors 114, one or more presentation components 116, input/output (I/O) ports 118, I/O components 120, and an illustrative power supply 122. Bus 110 represents what may be one or more busses (such as an address bus, data bus, or combination thereof). Although the various blocks of FIG. 1 are shown with lines for the sake of clarity, in reality, delineating various components is not so clear, and metaphorically, the lines would more accurately be grey and fuzzy. For example, one may consider a presentation component such as a display device to be an I/O component 120. Also, processors have memory. The inventors hereof recognize that such is the nature of the art, and reiterate that the diagram of FIG. 1 is merely illustrative of an exemplary computing device that can be used in connection with one or more embodiments of the invention. Distinction is not made between such categories as "workstation," "server," "laptop," "handheld device," etc., as all are contemplated within the scope of FIG. 1 and refer to "computer" or "computing device."

Computing device 100 typically includes a variety of computer-readable media. Computer-readable media can be any available media that can be accessed by computing device 100 and includes both volatile and nonvolatile media, removable and non-removable media. By way of example, and not limitation, computer-readable media may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules or other data.

Computer storage media includes RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices. Computer storage media does not comprise a propagated data signal.

Communication media typically embodies computer-readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier

wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer-readable media.

Memory 112 includes computer-storage media in the form of volatile and/or nonvolatile memory. The memory 112 may be removable, nonremovable, or a combination thereof. Exemplary memory includes solid-state memory, hard drives, optical-disc drives, etc. Computing device 100 includes one or more processors 114 that read data from various entities such as bus 110, memory 112 or I/O components 120. Presentation component(s) 116 present data indications to a user or other device. Exemplary presentation components 116 include a display device, speaker, printing component, vibrating component, etc. I/O ports 118 allow computing device 100 to be logically coupled to other devices including I/O components 120, some of which may be built in. Illustrative I/O components 120 include a microphone, joystick, game pad, satellite dish, scanner, printer, wireless device, etc.

Exemplary Online Gaming Environment

Turning now to FIG. 2, an online gaming environment 200 in which a multimode gaming servers may be deployed within a data center is shown, in accordance with an embodiment of the present invention. The online gaming environment 200 comprises various game clients connected through a network 220 to a game service 230. Exemplary game clients include a game console 210, a tablet 212, and a personal computer 214. Use of other game clients, such as smart phones, are also possible. The game console 210 may have one or more game controllers communicatively coupled to it. In one embodiment, the tablet 212 may act as an input device for a game console 210 or a personal computer 214. In another embodiment, the tablet 212 is a stand-alone game client. Network 220 may be a wide area network, such as the Internet.

Game service 230 comprises multiple computing devices communicatively coupled to each other. In one embodiment, the game service 230 is implemented using one or more data centers that comprise multimode gaming servers. The data centers may be spread out across various geographic regions including cities throughout

the world. In this scenario, the game clients may connect to the closest data centers. Embodiments of the present invention are not limited to this setup.

The game service 230 allows the game to be executed within the computing devices provided by the game service 230. A communication session between the game service and game clients carries input traffic to the game service 230 and returns a rendered game image. In this embodiment, a computing device that is part of the game service executes the video game code using a control stream generated by input devices associated with the various game clients. The rendered video game is then communicated over the network to the game client where the rendered game is output for display.

The game service 230 may be provided by a data center that uses a combination of gaming optimized servers to execute the game and render a video game image. The gaming optimized servers may be deployed with multimode gaming servers. When a suitable task is not available for the multimode gaming servers, the gaming CPU and GPU within the multimode gaming servers can be placed in a low-power mode and the non-gaming optimized CPU activated.

Exemplary Game Service

Turning now to FIG. 3, an exemplary remote gaming environment 300 is shown, in accordance with an embodiment of the present invention. The gaming environment 300 includes a game client 310 that is shown communicatively coupled to a game service 340 through a network 330. The gaming service may use one or more multimode gaming servers to accommodate peak gaming demand.

In one embodiment, the network may be the Internet. The game client 310 is connected to a first game input device 312, a second game input device 314, and a display 316. Exemplary game input devices include game pads, keyboards, a mouse, a touch pad, a touch screen, a microphone for receiving voice commands, a depth camera, a video camera, a keyboard, and a trackball. Embodiments of the present invention are not limited to these input devices. The display device 316 is capable of displaying video game content. For example, the display 316 may be a television or computer screen. In another embodiment, the display 316 is a touch screen integrated with the game client 310.

The game client 310 is a computing device that is able to execute video games. The game client 310 could be a tablet or a laptop computer. In another embodiment, the game client 310 is a game console and the display 316 is a remote display

communicatively coupled to the game console. The game client 310 includes an operating environment 320, a game execution environment 322, a game data store 324, a game service client 326, and a player profile data store 328.

The operating environment 320 may be provided by an operating system that manages the hardware and provides services to application running on the game client 310. The operating environment may allocate client resources to different applications as part of the game migration. For example, the operating environment may give control of the display to the game execution environment 322 once game play is migrated to the game client 310.

The game execution environment 322 comprises the gaming resources on the client 310 required to execute instances of a game or a game preview. The game execution environment 322 comprises active memory along with computing and video processing. The game execution environment 322 receives gaming controls and causes the game to be manipulated and progressed according to its programming. In one embodiment, the game execution environment 322 outputs a rendered video stream that is communicated to the display 316.

The game data store 324 stores downloaded games, game previews, and partially downloaded games.

The game service client 326 is a client application that displays rendered video game images received from the game service 340. The game service client 326 may also process game input and change it into an easily uploadable format that is communicated to the game service 340. The game service client 326 may also scale the rendered video game images received from the service 340 to a size optimized for display 316.

The player profile data store 328 stores player profile information for individual games. The player profile information may also save tombstones or game-saved data for individual games, including previews. Both the game-save file and the tombstone record game progress. The game execution environment 322 then reads the game-saved data to start the game where the player left off on the server. The opposite scenario is also possible where the game-saved data and player profile information is uploaded from the game client 310 to the game service 340 to initiate the game.

The game service 340 comprises a connection manager 342, a player profile data store 344, a game execution environment 348, and a game data store 350. Though depicted as a single box, the game service 340 could be implemented in a data center that comprises numerous machines, or even several data centers.

The connection manager 342 builds a connection between the client 310 and the service 340. The connection manager 342 may also provide various authentication mechanisms to make sure that the user is authorized to access the game service 340. The connection manager 342 may also analyze the bandwidth available within a connection and throttle the download of a game during game play to make sure that game play is not degraded.

The player profile data store 344 may work in conjunction with the connection manager 342 to build and store player information. Part of the player profile may comprise demographic and financial information such as a player's name, address and credit card information or other mechanism for paying for or purchasing games and experiences provided by the game service.

In addition, the player profile data store 344 may store a player's progress within an individual game. As a player progresses through a game or game preview, the player's score and access to game levels may be stored. Further, the player profile data store 344 may store information about individual player preferences such as language preferences. Information regarding a player's game client and speed of the network connection may also be stored and utilized to optimize the gaming experience. For example, in one embodiment, when a geographically proximate data center is busy, players with higher latency Internet connections may be preferentially connected to proximate data centers while players with lower latency connections may be connected to data centers that are further away. In this way, the players with the network connections that are best able to handle the additional latency are connected to data centers that create additional latency because of their location.

The player profile data store 344 may also store a usage history for the individual player. A player's history of purchasing games, sampling games, or playing games through a game service that does not require the purchase of the games may be stored. The usage information may be analyzed to suggest games of interest to an individual player. In one embodiment, the purchase history may include games that are not purchased through the game service. For example, the purchase history may be augmented by the player entering in a key from a game purchased in a retail store. In some embodiments, the player may then have access to that game both on their game client 310 and through the game service when they are no longer at their game client.

The game execution environment 348 comprises the gaming resources required to execute instances of a game. These are the resources described previously that

are managed by the game manager 352 and other components. The game execution environment 348 comprises active memory along with computing and video processing. The game execution environment 348 receives gaming controls through an I/O channel and causes the game to be manipulated and progressed according to its programming. In one embodiment, the game execution environment 348 outputs a rendered video stream that is communicated to the game client. In other embodiments, the game execution environment 348 outputs game geometry, or other representations, which may be combined with local objects on the gaming client to render the gaming video.

The game data store 350 stores available games. The games may be retrieved from the data store and activated through an active memory. The game data store 350 may be described as passive or secondary memory. In general, games may not be played off of the game data store 350. However, in some embodiments, the secondary memory may be utilized as virtual memory, in which case portions of the game data store 350 may also serve as active memory. This illustrates that active memory is not necessarily defined by a particular hardware component, but is defined by the ability of the game resources to actively manipulate and access objects within the memory to execute the game.

Turning now to FIG. 4, an arrangement of multimode gaming servers within a data center 400 is shown, in accordance with an aspect of the present invention. The arrangement comprises rack 410, rack 412, rack 414, and rack 416. Four racks are shown for the sake of simplicity; an actual implementation could include tens, hundreds, or thousands of racks deployed within a data center. Each rack can comprise a quantity of servers, power distribution equipment, and networking equipment. In one arrangement, a networking cable is run to a router/switch within the rack. Each server in the rack then connects to the router. Similarly, power may be run to a power distribution station associated with the rack. Each server is then coupled to the power distribution station.

Additionally, each rack can include cooling equipment, such as fans. In one arrangement, a fan wall is provided behind the servers to draw air through the servers. In a vertical cooling arrangement, one or more fans are located above or below the rack to facilitate airflow to the servers within the rack. The cooling equipment can also include thermocouples and other sensors that measure temperature, pressure, and air flow throughout the rack. The rack may include one or more fixed or adjustable baffles to distribute air where needed for cooling.

A control fabric 402 is communicatively coupled to the racks and computing devices within the racks. The control fabric 402 manages the state of each multimode gaming server. For example, the control fabric 402 can transition a multimode server between modes by activating a first type of computing resource and deactivating a second type of computing resource. The control fabric 402 can distribute workloads to computing devices. The control fabric 402 can also manage cooling equipment within the racks. For example, the control fabric 402 can lower fan speed within a rack when the servers within the rack are in a low power mode.

Racks 414 and 416 illustrate multimode servers in a mixture of modes. Aspects of the present invention are not limited to mixing modes within a data center unit, such as a rack chassis. In one aspect, all multimodal gaming servers within a data center unit are operating in the same mode. As FIG. 4 illustrates, a mixture of modes within a rack is also possible.

Illustratively, rack 416 includes multimode server 420 in game mode, multimode server 422 in general processing mode, and multimode server 424 in general processing mode. Illustratively, rack 414 includes multimode server 430 in game mode, multimode server 432 in game mode, and multimode server 434 in game mode. In one aspect, the computing devices within a rack have a homogeneous hardware configuration that allows them to switch between a gaming optimized mode and a non-gaming optimized mode.

The multimode servers can be transitioned between a gaming optimization and general-purpose optimization by activating different computing resources within the server. In one aspect, the general-purpose resources and game-optimized resources have peak usage periods that do not significantly overlap. Racks with an arrangement of multimode gaming servers, such as racks 414 and 416, may be deployed within a data center in combination with racks of single mode servers optimized for a single function, such as gaming. The quantity of single mode servers may be specified to accommodate base demand for the computing service provided by an optimized server. The deployment of single mode servers to meet base demand allows the single mode servers to be active above a threshold amount of time on average. For example, an amount of single mode servers deployed may be limited to those able to be active, on average, 80% of a day. Multimode servers can be used during peak usage periods to accommodate demand in combination with single mode servers optimized for the work load, such as gaming.

Turning now to FIG. 5, different types of resources within a multimode server are shown, in accordance with an aspect of the present invention. The multipurpose server includes motherboard 500. The motherboard 500 includes a first type of computing resource optimized for gaming. The first type of computing resource comprises a game optimized CPU 510, a game-optimized GPU 512, and a video encoder 514. This illustrates that a computing resource, as used herein, can comprise multiple hardware items. Also, though not shown, the computing resources may include memory and other components that support the resource. For example, a CPU can have dedicated DRAM memory. Aspects of the invention are not limited to use with a separate video encoder. A video encoder could be part of a CPU or GPU.

In one aspect, the game optimized CPU 510 and game-optimized GPU 512 are the same chips as found in a commercially available game console. Exemplary game consoles include the Xbox 360, Sony's PlayStation® family, Xbox One, and Nintendo's Wii™, and such. The hardware configuration associated with the game optimized computing resources can be configured to allow games written for a commercially available game console to run on the multimode server without modification to the game code and to interact with the hardware in the same way the games interact with hardware in the game console. For example, a process performed by a game console's GPU can be performed by the game-optimized computing resource's GPU.

The game optimized CPU 510, the game-optimized GPU 512, and the video encoder 514 can all be coupled to sockets within the motherboard. In one aspect, a computing resource is deactivated for transition to a low power mode by turning off the socket to which the resource is attached.

The general-purpose computing resource on motherboard 500 comprises enterprise-optimized CPU 520. In game mode, as shown in FIG. 5, the enterprise-optimized CPU 520 is drawing 0 W. In contrast, the game optimized CPU 510 is drawing 50 W, the game-optimized GPU 512 is drawing 90 W, and the video encoder 514 is drawing 10 W for a total use of 150 W in game mode.

Turning now to FIG. 6, power use in general purpose compute mode is illustrated, in accordance with an aspect of the present invention. Now the game optimized CPU 510, the game-optimized GPU 512 and a video encoder 514 are all drawing 0 W. In contrast, the enterprise-optimized CPU 520 is drawing 150 W. In one aspect, the rated power consumption of computing resources associated with different modes is substantially equal.

In this case, the computing resources associated with the gaming mode draw the same amount of power as the computing resources associated with the enterprise, or general compute mode. In one aspect, the cooling system for a multimode server is only capable of providing a knot cooling for one type of computing resource to be active at a given point in time.

Turning now to FIG. 7, a method 700 for managing workloads within a data center is provided, in accordance with an aspect of the present invention. Method 700 may be performed by a control fabric that manages workloads within a data center.

At step 710, substantially all of a first type of computing resource within a plurality of multimode servers are set to a low powered mode during a first time period. Each multimode server has multiple computing resources optimized for different types of work. The multiple computing resources comprise at least the first type of computing resource and a second type of computing resource. In one aspect, the first period of time corresponds to a low demand period for a workload the first type of computing resource is optimized to process. For example, a low demand period for a gaming workload may occur during the day and gaming optimized servers may be set to a low power mode during this period of time.

At step 720 substantially all of the second type of computing resource within the plurality of multimode servers are set to the low power mode at a second time period. The first and second time periods do not substantially overlap in one aspect. The nonoverlapping time periods allow the first type of computing resource and a second type of computing resource to satisfy peak demand for the computing loads they are optimized to handle.

The type of computing resource that is active within an individual multimode server can be adjusted based on workload demand. As the supply of a first type of workload increases, resources adapted to process the first type of workload can be activated. As demand for different types of workload changes, a mixture of states may be present in the plurality of multimode servers. For example, all the multimode servers could be set to process a first type of workload, a third of the multimode servers could be set to process a first type of workload, half of the multimode servers could be set to process a first type of workload, or none of the multimode servers could be set to process a first type of workload. Multimode servers that are not set to handle the first type of workload could be set to use a different type of workload.

Aspects of the invention have been described to be illustrative rather than restrictive. It will be understood that certain features and subcombinations are of utility and

may be employed without reference to other features and subcombinations. This is contemplated by and is within the scope of the claims.

CLAIMS

What is claimed is:

1. A multimode server comprising a gaming optimized computing resource having a first hardware configuration and a general processing optimized computing resource having a second hardware configuration that is different from the first hardware configuration, wherein the gaming optimized computing resource and the general processing optimized computing resource are power balanced to use a substantially equal amount of power at peak power, and wherein the multimode server comprises a control to activate either the gaming optimized computing resource or the general processing optimized computing resource at a given point in time, but not both simultaneously.

2. The multimode server of claim 1, wherein the gaming optimized computing resource is designed for a first workload with a peak usage during a first time period and the general processing optimized computing resource is designed for a second workload with a peak usage during a second time period that does not substantially overlap with the first time period.

3. The multimode server of claim 1, wherein the gaming optimized computing resource has a graphics processing unit ("GPU") and a central processing unit ("CPU"), and wherein a maximum power usage of the GPU comprises more than 40% of the gaming optimized computing resource's maximum power usage.

4. The multimode server of claim 1, wherein a power source provided to the multimode server does not supply enough power to run the gaming optimized computing resource and the general processing optimized computing resource simultaneously.

5. The multimode server of claim 1, wherein at least part of the gaming optimized computing resource is connected to a first socket on a mother board and at least part of the general processing optimized computing resource is connected to a second socket on the mother board.

6. A method for managing workloads within a data center, the method comprising: during a first time period, setting substantially all of a first type of computing resource within a plurality of multimode servers to a low power mode, each multimode server having multiple computing resources optimized for different types of work, the multiple computing resources comprising at least the first type of computing resource and a second type of computing resource; and during a second time period, setting substantially all of the second type of computing resource within the plurality of multimode servers to the low power mode, wherein the second time period does not substantially overlap with the first time period.

7. The method of claim 6, wherein said setting the first type of computing resource to the low power mode comprises deactivating one or more motherboard sockets to which the first type of computing resource is attached.

8. The method of claim 6, wherein the first type of computing resource is designed for a first workload with a peak usage during the first time period and the second type of computing resource is designed for a second workload with a peak usage during the second time period.

9. The method of claim 6, wherein the first type of computing resource and the second type of computing resource generate a substantially similar amount of heat when in use.

10. The method of claim 6, wherein the first type of computing resource outputs a rendered video game image over a wide area network to a remotely located gaming device.

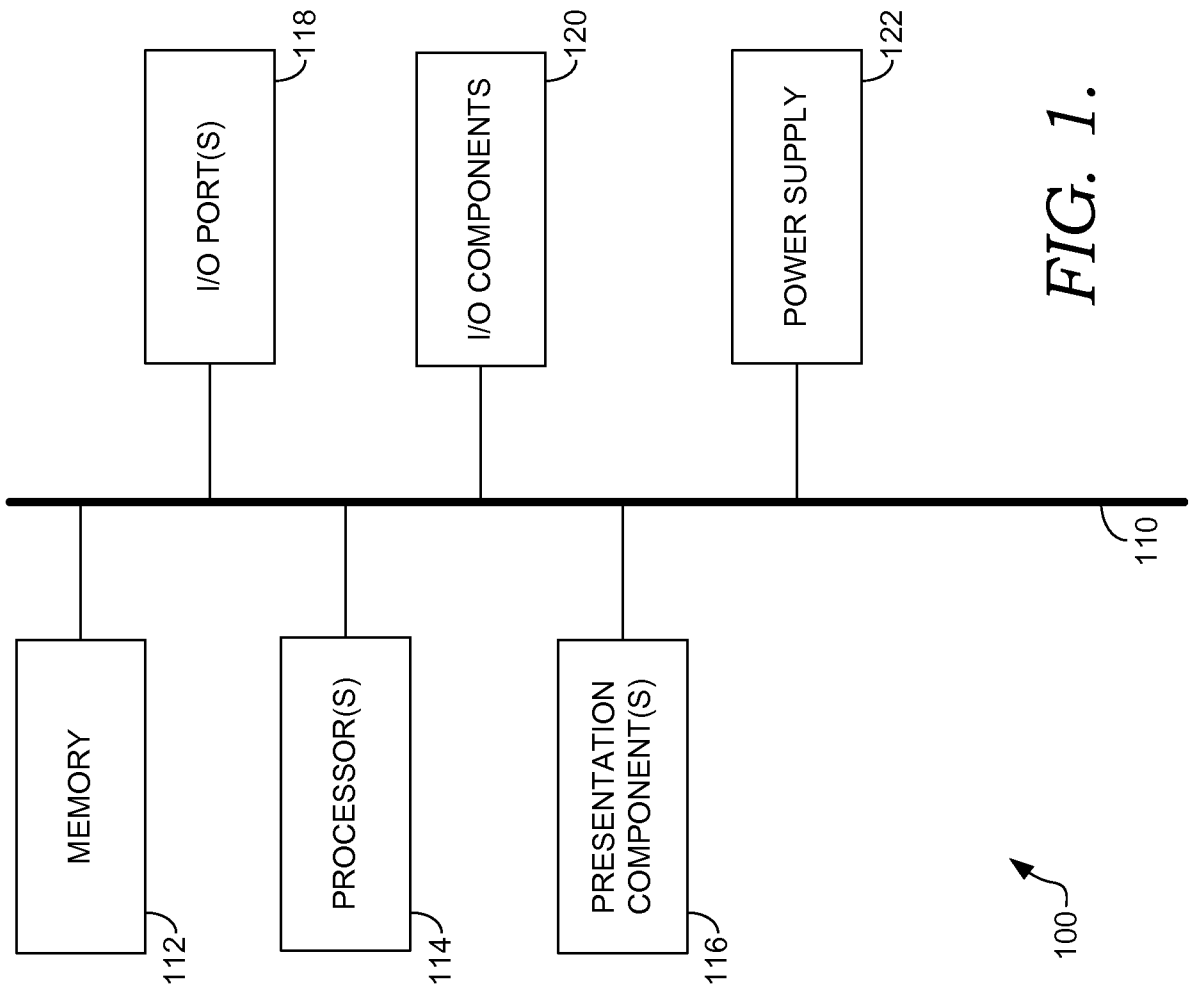
I/O PORT(S)
118
I/O COMPONENTS
120
POWER SUPPLY
122
110
MEMORY
112
PROCESSOR(S)
114
PRESENTATION COMPONENT(S)
116
100

FIG. 1.

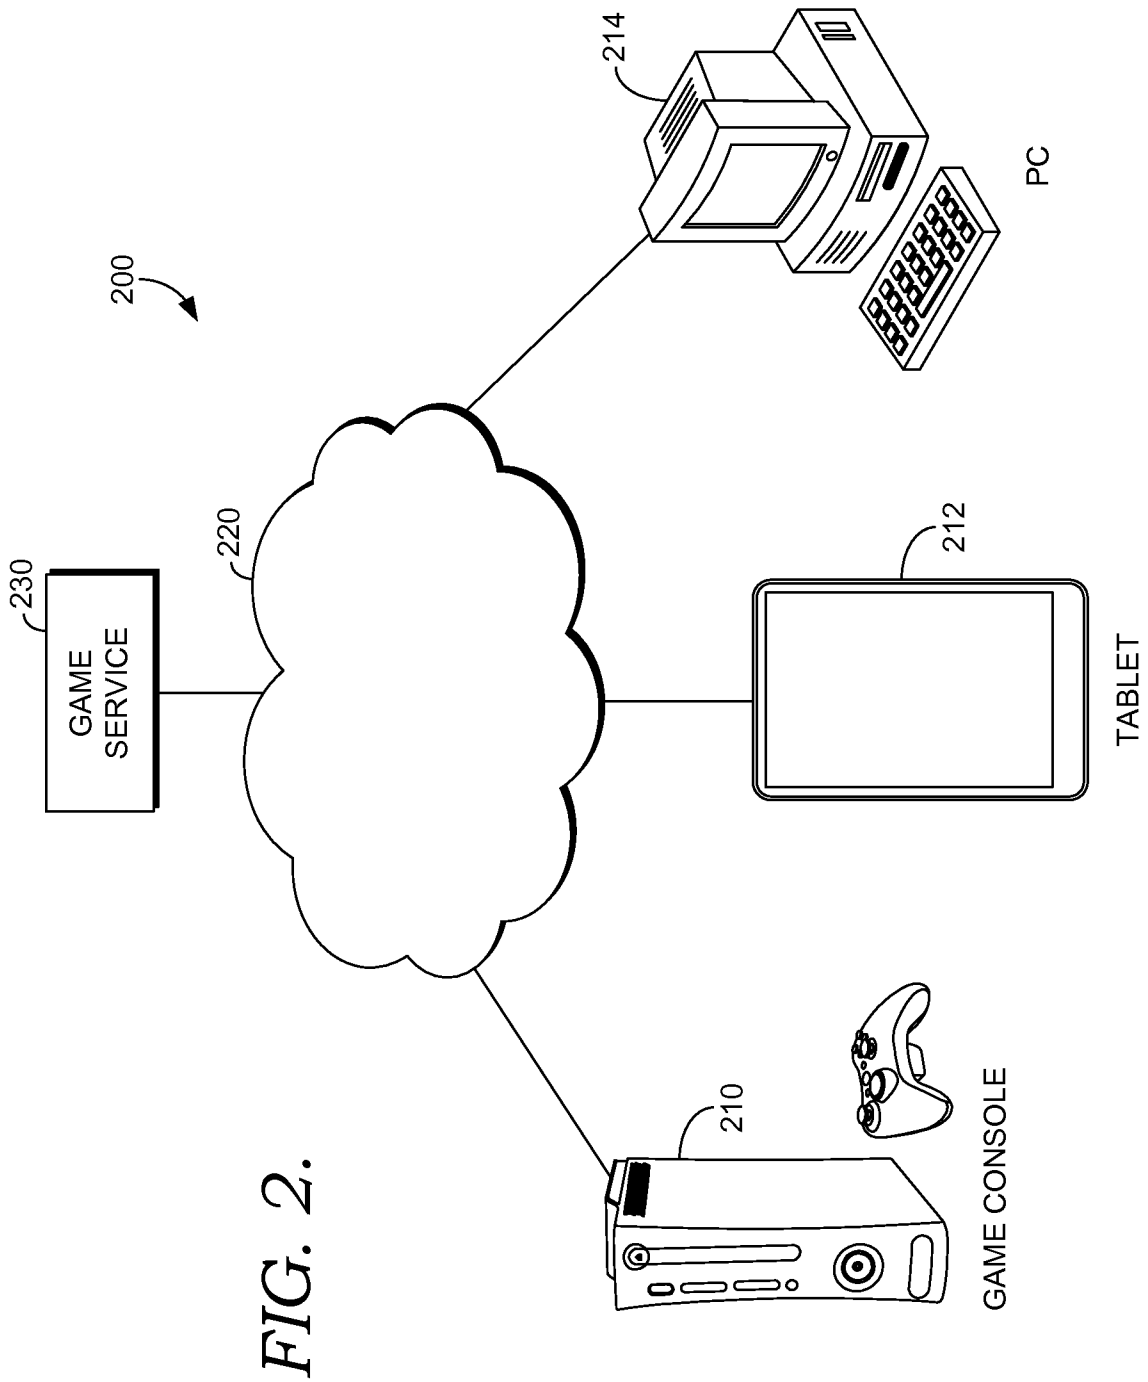

FIG. 2.

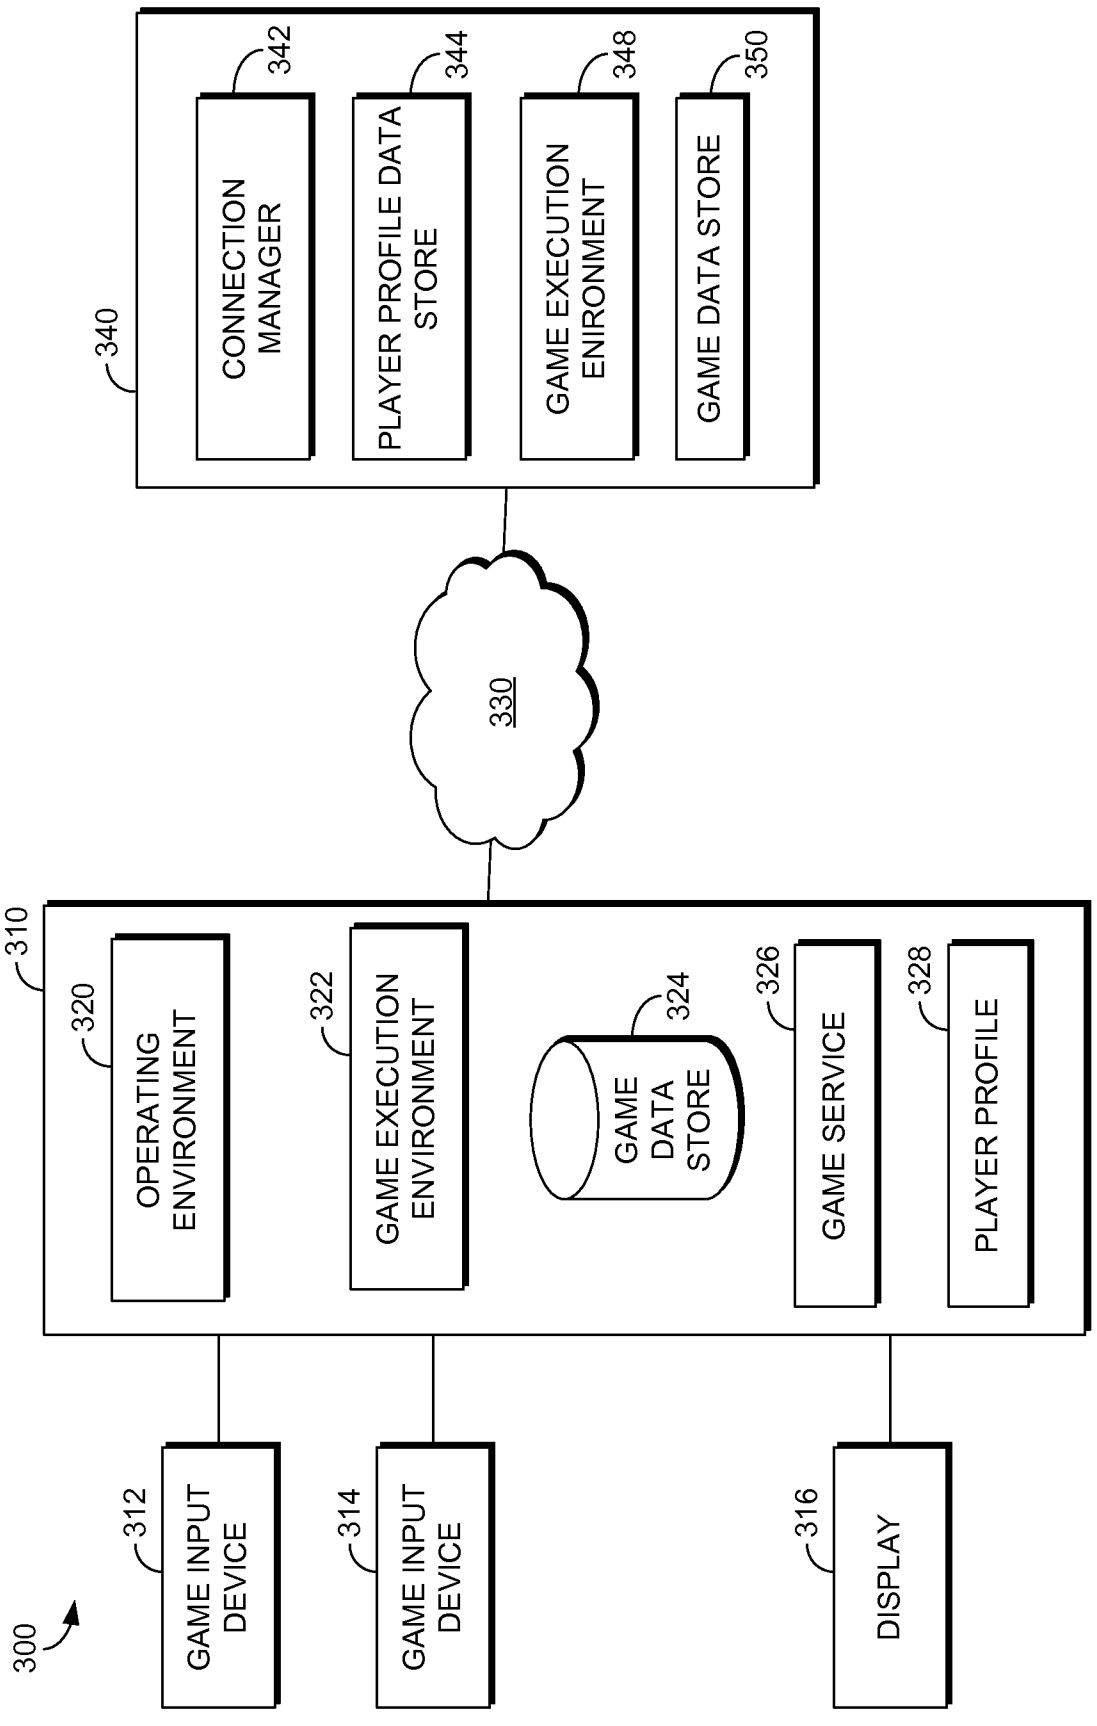
300
310
312
GAME INPUT DEVICE
314
GAME INPUT DEVICE
316
DISPLAY
320
OPERATING ENVIRONMENT
322
GAME EXECUTION ENVIRONMENT
324
GAME DATA STORE
326
GAME SERVICE
328
PLAYER PROFILE
330
340
342
CONNECTION MANAGER
344
PLAYER PROFILE DATA STORE
348
GAME EXECUTION ENIRONMENT
350
GAME DATA STORE

FIG. 3.

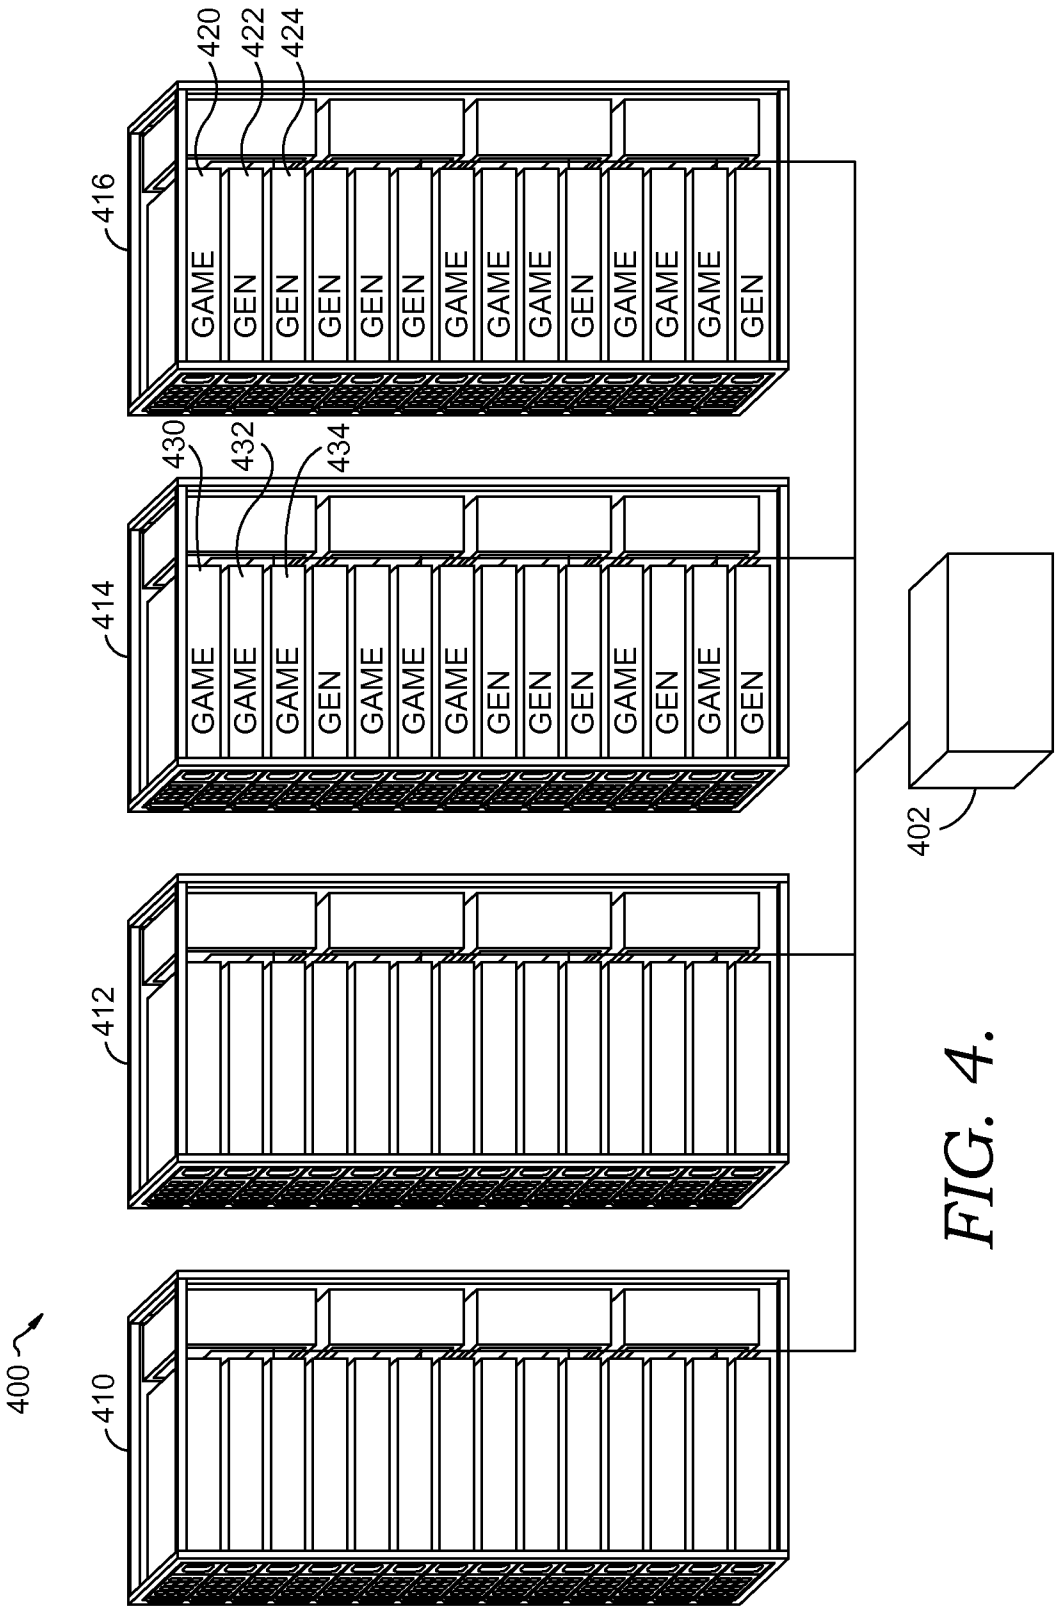

FIG. 4.

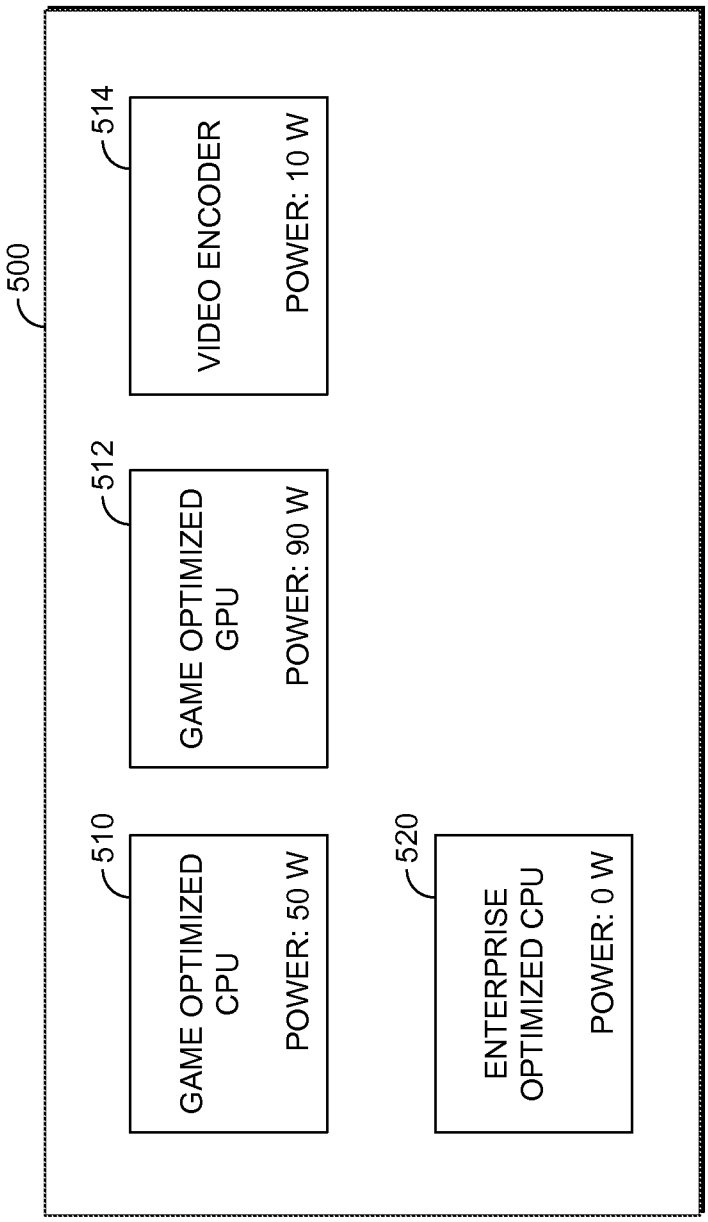
500
514
VIDEO ENCODER
POWER: 10 W
512
GAME OPTIMIZED GPU
POWER: 90 W
510
GAME OPTIMIZED CPU
POWER: 50 W
520
ENTERPRISE OPTIMIZED CPU
POWER: 0 W

FIG. 5.

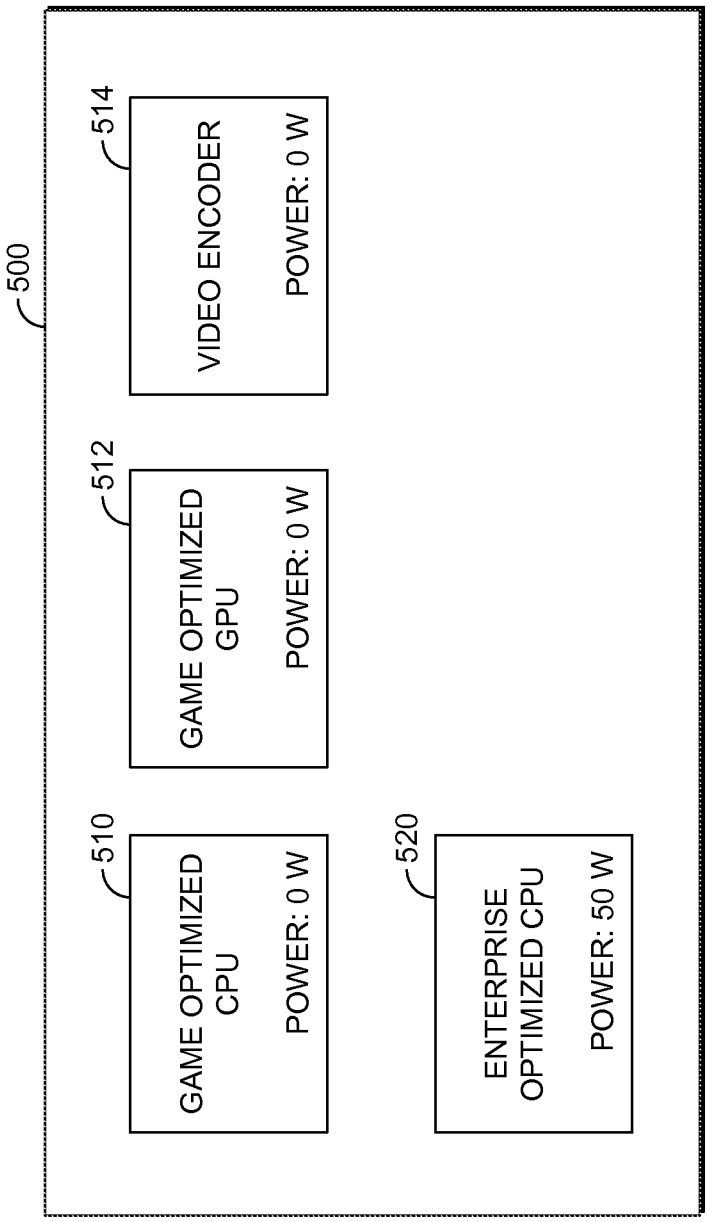
500
514
VIDEO ENCODER
POWER: 0 W
512
GAME OPTIMIZED GPU
POWER: 0 W
510
GAME OPTIMIZED CPU
POWER: 0 W
520
ENTERPRISE OPTIMIZED CPU
POWER: 50 W

FIG. 6.

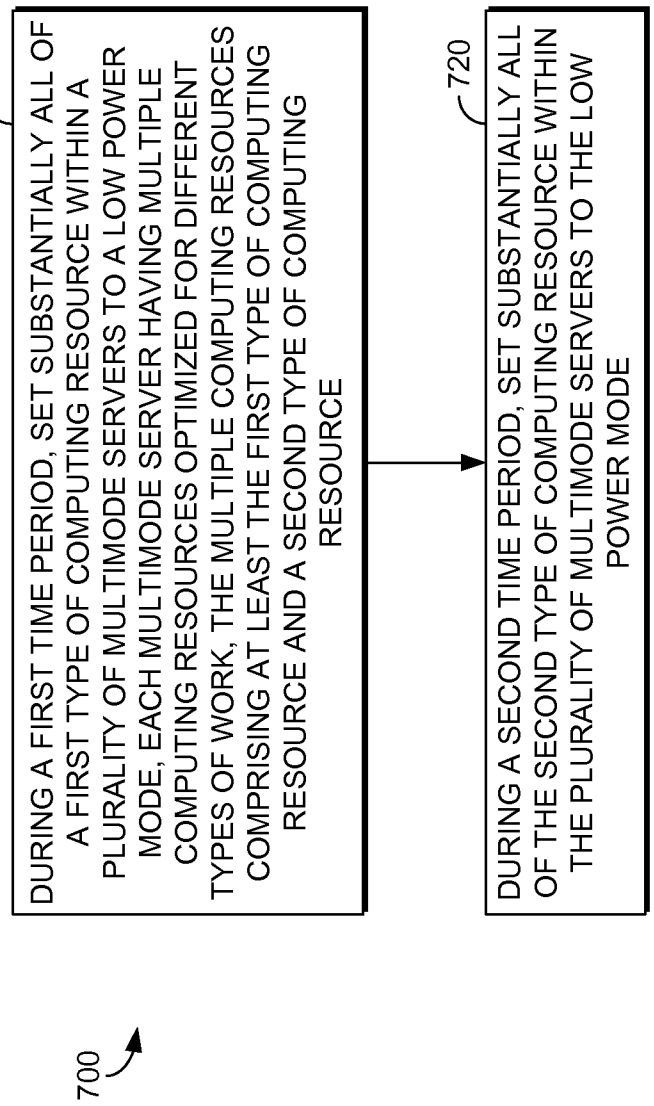
700
710
DURING A FIRST TIME PERIOD, SET SUBSTANTIALLY ALL OF A FIRST TYPE OF COMPUTING RESOURCE WITHIN A PLURALITY OF MULTIMODE SERVERS TO A LOW POWER MODE, EACH MULTIMODE SERVER HAVING MULTIPLE COMPUTING RESOURCES OPTIMIZED FOR DIFFERENT TYPES OF WORK, THE MULTIPLE COMPUTING RESOURCES COMPRISING AT LEAST THE FIRST TYPE OF COMPUTING RESOURCE AND A SECOND TYPE OF COMPUTING RESOURCE
720
DURING A SECOND TIME PERIOD, SET SUBSTANTIALLY ALL OF THE SECOND TYPE OF COMPUTING RESOURCE WITHIN THE PLURALITY OF MULTIMODE SERVERS TO THE LOW POWER MODE

FIG. 7.

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/068352

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

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

B. FIELDS SEARCHED

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

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2010/318818 A1 (MANGIONE-SMITH WILLIAM HENRY [US]) 16 December 2010 (2010-12-16) abstract; figures 1,46,52 paragraphs [0087], [0096], [0212]	1-5
A	US 2012/016528 A1 (RAMAN MADHUSUDAN [US] ET AL) 19 January 2012 (2012-01-19) paragraphs [0001], [0027], [0028], [0030], [0033], [0077], [0079], [0094]	1-5

☐ 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	Date of mailing of the international search report
20 February 2015	27/02/2015

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Kingma, Ype

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2014/068352

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010318818 A1	16-12-2010	NONE	
US 2012016528 A1	19-01-2012	US 2012016528 A1	19-01-2012
		US 2013103221 A1	25-04-2013

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2014/068352

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.: 6-10
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:

see FURTHER INFORMATION sheet PCT/ISA/210

3. ☐ Claims Nos.:
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 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.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

Continuation of Box II.2

Claims Nos.: 6-10

Due to unclarities in the claims the unity between the group of claims 1-5 and
the group of claims 6-10 cannot be established.

The reasons are given under Box VIII and Box III in the separate sheet of the written opinion.

The applicant's attention is drawn to the fact that claims relating to inventions in respect of which no international search report has been established need not be the subject of an international preliminary examination (Rule 66.1(e) PCT). The applicant is advised that the EPO policy when acting as an International Preliminary Examining Authority is normally not to carry out a preliminary examination on matter which has not been searched. This is the case irrespective of whether or not the claims are amended following receipt of the search report or during any Chapter II procedure. If the application proceeds into the regional phase before the EPO, the applicant is reminded that a search may be carried out during examination before the EPO (see EPO Guidelines C-IV, 7.2), should the problems which led to the Article 17(2) declaration be overcome.

(19)中华人民共和国国家知识产权局

(12)发明专利申请

(10)申请公布号 CN 105934745 A
(43)申请公布日 2016.09.07

(21)申请号 201480071809.1

(22)申请日 2014.12.03

(30)优先权数据
14/145915 2013.12.31 US

(85)PCT国际申请进入国家阶段日
2016.06.30

(86)PCT国际申请的申请数据
PCT/US2014/068352 2014.12.03

(87)PCT国际申请的公布数据
WO2015/102797 EN 2015.07.09

(71)申请人 微软技术许可有限责任公司
地址 美国华盛顿州

(72)发明人 E.P.贾登 J.R.朱斯蒂斯 M.沙马

(74)专利代理机构 中国专利代理(香港)有限公司 72001
代理人 刘红 景军平

(51)Int.Cl.
G06F 9/50(2006.01)

权利要求书1页 说明书8页 附图7页

(54)发明名称

多模游戏服务器

(57)摘要

本发明的各方面涉及多模游戏服务器，其中在服务器内提供不同类型的计算资源。不同的计算资源能够针对不同的计算任务进行优化。例如，第一类型的资源能够被优化用于制作高清晰度图形，而第二类型的资源能够被优化用于企业计算。可以针对不同的计算任务在整个一天内改变而根据需求来激活或去激活每个资源。在一个方面，这些资源是不同的母板插座上不同的芯片组。在一个方面，多模服务器中其他组件(例如，冷却，电源，网络宽带)的供应对于两个计算资源同时运行而言不是足够的。

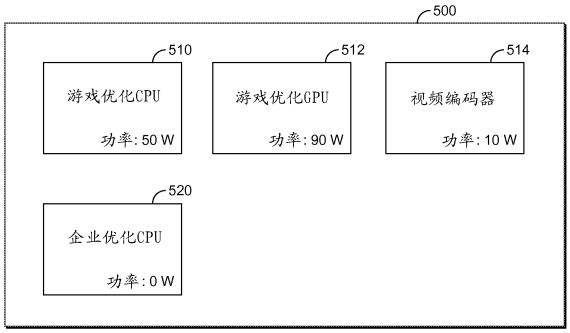

CN 105934745 A

1.一种多模服务器,包括具有第一硬件配置的游戏优化计算资源和具有与第一硬件配置不同的第二硬件配置的一般处理优化计算资源,其中游戏优化计算资源和一般处理优化计算资源被功率平衡,以便在峰值功率上使用实质上等量的功率,以及其中多模服务器包括控制来在给定的时间点激活或游戏优化计算资源或一般处理优化计算资源,但是不是同时激活二者。

2.权利要求1的多模服务器,其中游戏优化计算资源被设计用于在第一时间周期期间具有峰值使用的第一工作负荷,而一般处理优化计算资源被设计用于在与第一时间周期实质上不重叠的第二时间周期期间具有峰值使用的第二工作负荷。

3.权利要求1的多模服务器,其中游戏优化计算资源具有图形处理单元(“GPU”)和中央处理单元(“CPU”),以及其中GPU的最大功率使用包括大于40%的游戏优化计算资源的最大功率使用。

4.权利要求1的多模服务器,其中提供给多模服务器的电源没有供给足够的功率来同时运行游戏优化计算资源和一般处理优化计算资源。

5.权利要求1的多模服务器,其中至少部分的游戏优化计算资源被连接至母板上的第一插座,并且至少部分的一般处理优化计算资源被连接至母板上的第二插座。

6.一种用于在数据中心内管理工作负荷的方法,该方法包括:在第一时间周期期间,实质上将多个多模服务器内所有的第一类型的计算资源设置为低功率模式,每一个多模服务器具有针对不同类型的工作进行优化的多个计算资源,多个计算资源至少包括第一类型的计算资源和第二类型的计算资源;以及在第二时间周期期间,实质上将多个多模服务器内所有的第二类型的计算资源设置为低功率模式,其中第二时间周期实质上与第一时间周期不重叠。

7.权利要求6的方法,其中所述的将第一类型的计算资源设置为低功率模式包括去激活第一类型的计算资源所附着的一个或多个母板插座。

8.权利要求6的方法,其中第一类型的计算资源被设计用于在第一时间周期期间具有峰值使用的第一工作负荷,而第二类型的计算资源被设计用于在第二时间周期期间具有峰值使用的第二工作负荷。

9.权利要求6的方法,其中第一类型的计算资源和第二类型的计算资源在使用的时候生成实质上类似的热量。

10.权利要求6的方法,其中第一类型的计算资源通过广域网将所渲染的视频游戏图像输出至远程定位的游戏设备。

多模游戏服务器

背景技术

[0001] 一般而言，被选择用于在数据中心中部署的服务器能够执行广泛的计算任务，但是可能无法非常有效地执行一些专业计算任务。例如，在具有强大GPU和视频编码器的服务器上较好地执行视频密集计算项目。企业服务器可能能够通过CPU执行一些视频相关的工作，但是该工作可能是效率低的。另一方面，被设计用于视频相关工作的服务器可能效率低地执行一般计算项目。

发明内容

[0002] 提供这个概述部分来以简化形式介绍下面在具体描述部分中进一步描述的概念的选择。这个概述部分并不打算用于识别所请求保护的主题的关键特性或基本特性，也不打算被孤立用作辅助手段来确定所请求保护的主题的范畴。

[0003] 本发明的各方面涉及多模游戏服务器，其中在服务器内提供不同类型的计算资源。不同的计算资源能够针对不同的计算任务进行优化。例如，第一类型的资源能够被优化用于制作高清晰度图形，而第二类型的资源能够被优化用于企业计算。可以针对不同的计算任务在整个一天内改变而根据需求来激活或去激活每个资源。在一个方面，这些资源是不同的母板插座中不同的芯片组(chip set)。在一个方面，多模服务器中其他组件(例如，冷却，电源，网络宽带)的供应对于两个计算资源同时运行而言不是足够的。

[0004] 在一个方面，服务器的冷却容量被故意设计尺寸，以致在服务器中所有的计算资源同时处于活动(active)处理模式中时不能提供足够的冷却来为多模服务器保持可接受的操作温度。数据中心的控制构造(fabric)能够通过在给定的时间点只给多模服务器内一种类型的计算资源分配工作负荷而在服务器内保持可接受的操作温度。例如，在任何给定的时间，只有服务器中的游戏优化计算资源可以被分配工作负荷并且处于活动处理状态中。服务器中剩余的计算资源被设置为低功率状态。

[0005] 在一个方面，多模服务器内不同的计算资源当处于活动处理状态中时具有相类似的最大功率使用。例如，具有150W的最大功率使用的游戏优化资源(例如，GPU和专业CPU)能够被部署在与具有150W容量的企业CPU相同的多模服务器中。

[0006] 在一个方面，根据预期的峰值使用周期来选择计算资源以便包含在多模服务器中。被设计用于具有彼此不同的峰值使用周期的专业工作负荷的资源能够被包括在多模游戏服务器中。例如，与具有从4PM到12PM的高峰小时的专业工作负荷相关联的第一类型的计算资源能够与具有带有从6AM到2PM的高峰小时的专业工作负荷的第二类型的计算资源进行匹配。换句话说，在给定的周期期间，服务器内或第一类型或第二类型的计算资源将处于高需求中。

附图说明

[0007] 下面参考所附的附图来具体描述本发明的各方面，其中：

图1是适于实现本发明的各方面的示例性计算环境的方框图；

图2是根据本发明的一方面、描绘游戏环境的图表；

图3是根据本发明的一方面、描绘具有一个或多个带有游戏服务器和通用服务器的非均质(nonhomogeneous)安排的数据中心的远程游戏环境的图表；

图4是根据本发明的一方面、描绘多模游戏服务器在各种模式中的安排的图表；

图5是根据本发明的一方面、描绘具有活动游戏资源的母板的图表；

图6是根据本发明的一方面、描绘具有活动通用资源的母板的图表;和

图7是根据本发明的一方面、描绘用于在数据中心内管理工作负荷的方法的图表。

具体实施方式

[0008] 在本文利用特异性来描述本发明的各方面的主题以满足法定要求。然而,该描述本身并不打算用于限制这个专利的范畴。相反,发明人已设想:结合其他的目前或未来的技术,所请求保护的主题也可能采用其他的方式来体现,以包括与这个文档中所描述的步骤不同的步骤或相类似的步骤的组合。此外,虽然在本文可以使用术语“步骤”和/或“块”来暗指所采用的方法的不同元素,但是这些术语不应被解释为暗示在本文所披露的各种步骤之中或之间的任何特殊顺序,除非且除了明确地描述个别步骤的顺序之外。

[0009] 本发明的各方面涉及多模游戏服务器,其中在服务器内提供不同类型的计算资源。不同的计算资源能够针对不同的计算任务进行优化。例如,第一类型的资源能够被优化用于制作高清晰度图形,而第二类型的资源能够被优化用于企业计算。可以针对不同的计算任务在整个一天内改变而根据需求来激活或去激活每个资源。在一个方面,这些资源是不同的母板插座中不同的芯片组。在一个方面,多模服务器中其他组件(例如,冷却,电源,网络宽带)的供应对于两个计算资源同时运行而言不是足够的。

[0010] 在一个方面,服务器的冷却容量被故意设计尺寸,以致在服务器中所有的计算资源同时处于活动处理模式中时不能提供足够的冷却来为服务器保持可接受的操作温度。数据中心的控制构造能够通过在给定的时间点只给多模服务器内一种类型的计算资源分配工作负荷而在服务器内保持可接受的操作温度。例如,在任何给定的时间,只有服务器中的游戏优化计算资源可以被分配工作负荷并且处于活动处理状态中。服务器中剩余的计算资源被设置为低功率状态。

[0011] 在一个方面,多模服务器内不同的计算资源当处于活动处理状态中时具有相类似的最大功率使用。例如,具有150W的最大功率使用的游戏优化资源(例如,GPU和专业CPU)能够被部署在与具有150W容量的企业CPU相同的多模服务器中。即使两种类型的计算资源的总功率使用是类似的,资源中功率使用的分布也能够极为不同。例如,游戏优化资源可以具有使用100W的图形处理单元(“GPU”)和使用50W的中央处理单元(“CPU”)。企业资源可以没有GPU,但是能够具有消耗150W的更强大CPU。

[0012] 在一个方面,根据预期的峰值使用周期来选择计算资源以便包含在多模服务器中。被设计用于具有彼此不同的峰值使用周期的专业工作负荷的资源能够被包括在多模游戏服务器中。例如,与具有从4PM到12PM的高峰小时的专业工作负荷相关联的第一类型的计算资源能够与具有带有从6AM到2PM的高峰小时的专业工作负荷的第二类型的计算资源进行匹配。换句话说,在给定的周期期间,服务器内或第一类型或第二类型的计算资源将处于高需求中。

[0013] 如在本文所用的，“游戏优化计算资源(gaming optimized computing resource)”适应于将所渲染的视频游戏图像输出至客户端设备，诸如游戏机(game console)。视频游戏图像可以作为被传送至客户端的流媒体视频来渲染。为了渲染高质量视频游戏图像，游戏优化计算资源能够具有比在通用计算资源中找到的图形处理单元(如果有的话)更强大的图形处理单元。游戏优化计算资源也可以具有专用视频编码能力。

[0014] 功率消耗能够被用作处理器的能力的代理(proxy)。在一个方面，游戏优化计算资源能够利用包含消耗比游戏优化计算资源在峰值功率消耗期间所使用的功率的阈值百分比更大的GPU来定义。在一个方面，功率的阈值百分比大于40%的峰值功率，例如大于50%，例如大于60%，例如大于70%，或例如大于80%。例如，游戏优化计算资源中的GPU能够使用100W，其中在游戏优化服务器资源中具有150W的总峰值功率使用(例如，GPU和CPU)。

[0015] 如在本文所用的，术语“通用计算资源(general-purpose computing resource)”或“一般处理优化计算资源(general processing optimized computing resource)”描述被设计来强调典型地与中央处理单元相关联的计算处理的资源。通用计算资源可能能够执行专业计算处理，但是可能没有针对那个用途进行优化。例如，CPU能够比利用GPU能够执行相同或相类似任务更低效地执行图形处理。

[0016] 本发明的各方面可以在不同的功率模式或状态之间转变各种类型的计算资源。如在本文所用的，术语“低功率模式(low-power mode)”意味着：资源目前正操作在低于20%的资源的最大额定功率(rate of power)上。作为示例，低功率模式中的资源可以在母板插座上被关闭，但是能够响应通电指令。

[0017] 如在本文所用的，短语“活动处理模式(active processing mode)”意味着：计算资源正在活动地处理计算工作负荷。活动处理模式中的计算资源能够正在使用大于20%的资源的最大额定功率。

[0018] 在简要描述本发明的各方面的综述之后，下面描述适于在实现本发明的各方面中使用的示例性操作环境。

[0019] 示例性操作环境

一般参考附图，并且初始地，特别参考图1，用于实现本发明的实施例的示例性操作环境被显示并且一般被标明为计算设备100。计算设备100只是合适的计算环境的一个示例而并不打算针对本发明的使用或功能的范畴建议任何限制。计算设备100 也不应被解释为具有与所举例说明的组件之中的任何一个或组合相关的任何依赖或要求。

[0020] 可以在利用计算机或其他机器诸如个人数据助理或其他手持设备执行的包括计算机可执行指令诸如程序组件的计算机代码或机器可用指令的一般上下文中描述本发明。一般而言，包括例程、程序、对象、组件、数据结构等等的程序组件指的是执行特殊任务或实现特殊抽象数据类型的代码。本发明的实施例可以在包括手持设备、消费电子产品、通用计算机、专业计算设备等等的各种系统配置中进行实践。本发明的实施例也可以在其中利用通过通信网络链接的远程处理设备来执行任务的分布式计算环境中进行实践。

[0021] 继续参考图1，计算设备100 包括直接或间接耦合以下设备的总线110：存储器112，一个或多个处理器114，一个或多个呈现组件116，输入/输出(I/O)端口118，I/O组件120，以及说明性的电源122。总线110代表什么可以是一条或多条总线(诸如地址总线、数据总线或其组合)。虽然图1的各种块为了清晰起见而利用线条来显示，但是事实上，划定各种

组件并不是如此清楚的，并且隐喻地，这些线条将更准确地是灰色和模糊的。例如，可以将呈现组件诸如显示设备视为I/O组件120。处理器也具有存储器。其发明人认识到：这就是本领域的性质，并且发明人重申：图1的图表只是举例说明能够结合本发明的一个或多个实施例来使用的示例性计算设备。在诸如“工作站”、“服务器”、“膝上型计算机”、“手持设备”等等之类的类别之间并不进行区分，这是因为所有的类别被设想在图1的范畴之内并称为“计算机”或“计算设备”。

[0022] 计算设备100典型地包括各种计算机可读媒体。计算机可读媒体能够是任何的可用媒体，其能够利用计算设备100来访问并包括易失性与非易失性媒体、可移动与不可移动媒体二者。通过示例而非限制，计算机可读媒体可以包括计算机存储媒体和通信媒体。计算机存储媒体包括在任何的用于存储信息诸如计算机可读指令、数据结构、程序模块或其他数据的方法或技术中实现的易失性与非易失性、可移动与不可移动媒体二者。

[0023] 计算机存储媒体包括RAM、ROM、EEPROM、闪存或其它的存储技术、CD-ROM、数字多用途盘(DVD)或其它的光盘储存器、磁带盒、磁带、磁盘储存器或其他的磁存储设备。计算机存储媒体不包括传播的数据信号。

[0024] 通信媒体典型地体现计算机可读指令、数据结构、程序模块或调制的数据信号诸如载波或其它传输机制中的其他数据并包括任何的信息递送媒体。术语“调制的数据信号”意味着这样的信号，其特征之中的一个或多个特征以在该信号中编码信息的方式进行设置或改变。通过示例而非限制，通信媒体包括有线媒体诸如有线网络或直接连线连接和无线媒体诸如声学、RF、红外和其他的无线媒体。上述之中任何的组合也应该被包括在计算机可读媒体的范畴之内。

[0025] 存储器112包括采用易失性和/或非易失性存储器的形式的计算机存储媒体。存储器112可以是可移动、不可移动或其组合。示例性存储器包括固态存储器、硬盘驱动器、光盘驱动器等等。计算设备100包括从各种实体诸如总线110、存储器112或I/O组件120中读取数据的一个或多个处理器114。(一个或多个)呈现组件116向用户或其他设备呈现数据指示。示例性呈现组件116包括显示设备、扬声器、打印组件、振动组件等等。I/O端口118允许计算设备100被逻辑耦合至包括I/O组件120的其他设备，其中的一些设备可以被内置。说明性的I/O组件120包括麦克风、操纵杆、游戏垫、卫星天线、扫描仪、打印机、无线设备等等。

[0026] 示例性在线游戏环境

现在转到图2，根据本发明的实施例，显示其中可以在数据中心内部署多模游戏服务器的在线游戏环境200。在线游戏环境200包括通过网络220连接至游戏服务230的各种游戏客户端。示例性游戏客户端包括游戏机210、平板电脑212和个人计算机214。其他的游戏客户端诸如智能电话的使用也是有可能的。游戏机210可以具有一个或多个与之通信耦合的游戏控制器。在一个实施例中，平板电脑212可以充当游戏机210或个人计算机214的输入设备。在另一实施例中，平板电脑212是独立的游戏客户端。网络220可以是广域网，诸如因特网。

[0027] 游戏服务230包括彼此通信耦合的多个计算设备。在一个实施例中，使用包括多模游戏服务器的一个或多个数据中心来实现游戏服务230。数据中心可以跨越包括世界各地的城市的各种地理区域进行展开(spread out)。在这种情形中，游戏客户端可以连接至最近的数据中心。本发明的实施例并不限于这种布置。

[0028] 游戏服务230允许在利用游戏服务230所提供的计算设备内执行游戏。游戏服务和游戏客户端之间的通信会话将输入业务(traffic)运送至游戏服务230并返回所渲染的游戏图像。在这个实施例中，为游戏服务的一部分的计算设备使用利用与各种游戏客户端相关联的输入设备所生成的控制流来执行视频游戏代码。所渲染的视频游戏随后通过网络被传送至游戏客户端，其中在游戏客户端中所渲染的游戏被输出，以便显示。

[0029] 游戏服务230可以利用使用游戏优化服务器的组合来执行游戏和渲染视频游戏图像的数据中心来提供。游戏优化服务器可以与多模游戏服务器一起进行部署。当合适的任务不可用于多模游戏服务器时，多模游戏服务器内的游戏CPU和GPU能够被放置在低功率模式中并且非游戏优化CPU被激活。

[0030] 示例性游戏服务

现在转到图3，根据本发明的实施例，显示示例性远程游戏环境300。游戏环境300包括被显示为通过网络330通信耦合至游戏服务340的游戏客户端310。游戏服务可以使用一个或多个多模游戏服务器来容纳(accommodate)峰值游戏需求。

[0031] 在一个实施例中，网络可以是因特网。游戏客户端310被连接至第一游戏输入设备312、第二游戏输入设备314和显示器316。示例性游戏输入设备包括游戏垫、键盘、鼠标、触摸板、触摸屏、用于接收话音命令的麦克风、深度照相机、摄像机、键盘和轨迹球。本发明的实施例并不限于这些输入设备。显示设备316能够显示视频游戏内容。例如，显示器316可以是电视或计算机屏幕。在另一实施例中，显示器316是与游戏客户端310相集成的触摸屏。

[0032] 游戏客户端310是能够执行视频游戏的计算设备。游戏客户端310能够是平板电脑或膝上型计算机。在另一实施例中，游戏客户端310是游戏机，并且显示器316是通信耦合至游戏机的远程显示器。游戏客户端310包括操作环境320、游戏执行环境322、游戏数据储存器324、游戏服务客户端326和玩家简档(profile)数据储存器328。

[0033] 操作环境320可以利用操作系统来提供，其中操作系统管理硬件并提供服务给运行在游戏客户端310上的应用。操作环境可以将客户端资源分配给不同的应用作为游戏迁移的一部分。例如，一旦玩游戏被迁移至游戏客户端310，操作环境可以将显示器的控制给予游戏执行环境322。

[0034] 游戏执行环境322包括执行游戏或游戏预览的实例所要求的客户端310上的游戏资源。游戏执行环境322连同计算与视频处理一起包括主动存储器(active memory)。游戏执行环境322接收游戏控制并使得游戏根据其编程来操纵和进展。在一个实施例中，游戏执行环境322输出被传送至显示器316的所渲染的视频流。

[0035] 游戏数据储存器324存储所下载的游戏、游戏预览和部分下载的游戏。

[0036] 游戏服务客户端326是显示从游戏服务340接收的所渲染的视频游戏图像的客户端应用。游戏服务客户端326也可以处理游戏输入并将其改变成被传送至游戏服务340的轻松上传格式。游戏服务客户端326也可以将从服务340接收的所渲染的视频游戏图像缩放至针对显示器316进行优化的尺寸。

[0037] 玩家简档数据储存器328存储个别游戏的玩家简档信息。玩家简档信息也可以保存包括预览的个别游戏的墓碑(tombstone)或游戏保存的数据。游戏保存文件和墓碑二者记录游戏进展。游戏执行环境322随后读取游戏保存的数据，以便在玩家在服务器上停止的地方开始游戏。其中从游戏客户端310上传游戏保存的数据和玩家简档信息至游戏服务340

以启动游戏的相反情景也是有可能的。

[0038] 游戏服务340包括连接管理器342、玩家简档数据储存器344、游戏执行环境348和游戏数据储存器350。虽然被描绘为单个块，但是能够在包括许多机器的数据中心或甚至若干数据中心中实现游戏服务340。

[0039] 连接管理器342在客户端310和服务340之间建立连接。连接管理器342也可以提供各种鉴别机制来确保用户被授权访问游戏服务340。连接管理器342也可以分析在连接内可用的带宽并在玩游戏期间节流(throttle)游戏的下载以确保玩游戏不被降级。

[0040] 玩家简档数据储存器344可以结合连接管理器342来工作，以建立和存储玩家信息。玩家简档的一部分可以包括人口统计信息和财务信息诸如玩家的姓名、地址和信用卡信息或其它的用于付款或购买游戏与利用游戏服务所提供的体验的机制。

[0041] 另外，玩家简档数据储存器344可以存储个别游戏内玩家的进展。在玩家通过游戏或游戏预览而进展时，玩家的分数以及针对游戏级别的访问可以被存储。进一步，玩家简档数据储存器344可以存储有关个别玩家偏好诸如语言偏好的信息。涉及玩家的游戏客户端和网络连接的速度的信息也可以被存储并被利用来优化游戏体验。例如，在一个实施例中，当地理上最接近的数据中心忙碌时，具有较高延迟(latency)因特网连接的玩家可以被优选连接至最接近的数据中心，而具有较低延迟连接的玩家可以被连接至进一步远离的数据中心。以这种方式，具有最好能够应付附加延迟的网络连接的玩家被连接至因为其位置而创建附加延迟的数据中心。

[0042] 玩家简档数据储存器344也可以存储个别玩家的使用历史。玩家的购买游戏、采样游戏或通过不要求购买游戏的游戏服务来玩游戏的历史可以被存储。使用信息可以被分析来向个别玩家建议感兴趣的游戏。在一个实施例中，购买历史可以包括没有通过游戏服务购买的游戏。例如，通过玩家从零售店中购买的游戏中输入密钥，可以扩大(augment)购买历史。在一些实施例中，玩家随后可以在其游戏客户端310上并在其不再位于其游戏客户端上的时候通过游戏服务二者来访问那个游戏。

[0043] 游戏执行环境348包括执行游戏的实例所需要的游戏资源。这些是以前描述的利用游戏管理器352和其他组件来管理的资源。游戏执行环境348连同计算与视频处理一起包括主动存储器。游戏执行环境348通过I/O通道接收游戏控制并使得游戏根据其编程来操纵和进展。在一个实施例中，游戏执行环境348输出被传送至游戏客户端的所渲染的视频流。在其他的实施例中，游戏执行环境348输出游戏几何形状或其他表示，其可以与游戏客户端上的本地对象进行组合来渲染游戏视频。

[0044] 游戏数据储存器350存储可用游戏。可以从数据储存器中检索这些游戏并通过主动存储器来激活这些游戏。游戏数据储存器350可以被描述为被动或辅助存储器。一般而言，可能无法玩游戏数据储存器350中的游戏。然而，在一些实施例中，辅助存储器可以被用作虚拟存储器，在这种情况中游戏数据储存器350的各部分也可以充当主动存储器。这举例说明：不一定利用特殊硬件组件来定义主动存储器，但是利用游戏资源主动操纵和访问存储器内的对象以执行游戏的能力来定义主动存储器。

[0045] 现在转到图4，根据本发明的一方面，显示多模游戏服务器在数据中心400内的安排。该安排包括支架410、支架412、支架414和支架416。为了简单起见而显示四个支架；实际的实现方式能够包括在数据中心内部署的数十、数百或数千的支架。每个支架能够包括大

量的服务器、配电设备和联网设备。在一种安排中，连网缆线被运行至支架内的路由器/转换器。支架中的每个服务器随后连接至路由器。类似地，功率可以被运行至与支架相关联的配电站。每个服务器随后被耦合至配电站。

[0046] 此外，每个支架能够包括冷却设备，诸如风扇。在一种安排中，在这些服务器后面提供扇墙(fan wall)，以便通过这些服务器吸入空气。在垂直冷却安排中，一个或多个风扇被定位于支架之上或之下，以促进气流至支架内的服务器。冷却设备也能够包括热耦合和其他的传感器，其测量整个支架的温度、压力和空气流量。支架可以包括一个或多个固定或可调挡板来在需要冷却的地方分布空气。

[0047] 控制构造402被通信耦合至支架和支架内的计算设备。控制构造402管理每个多模游戏服务器的状态。例如，控制构造402能够通过激活第一类型的计算资源和去激活第二类型的计算资源而在模式之间转变多模服务器。控制构造402能够分布工作负荷至计算设备。控制构造402也能够管理支架内的冷却设备。例如，控制构造402能够在支架内的服务器位于低功率模式中时降低支架内的风扇速度。

[0048] 支架414和416举例说明在模式的混合中的多模服务器。本发明的各方面并不限于数据中心单元诸如支架机箱内的混合模式。在一个方面，数据中心单元内所有的多模态游戏服务器正操作在相同的模式中。如图4所举例说明的，支架内模式的混合也是有可能的。

[0049] 说明性地，支架416包括游戏模式中的多模服务器420、一般处理模式中的多模服务器422和一般处理模式中的多模服务器424。说明性地，支架414包括游戏模式中的多模服务器430、游戏模式中的多模服务器432和游戏模式中的多模服务器434。在一个方面，支架内的计算设备具有允许其在游戏优化模式和非游戏优化模式之间转换的均质硬件配置。

[0050] 多模服务器能够通过激活服务器内不同的计算资源而在游戏优化和通用优化之间进行转变。在一个方面，通用资源和游戏优化资源具有没有明显重叠的峰值使用周期。具有多模游戏服务器的安排的支架诸如支架414和416可以与针对单个功能诸如游戏进行优化的单个模式服务器的支架相组合而被部署在数据中心内。单个模式服务器的数量可以被指定，以容纳利用优化服务器所提供的计算服务的基本需求。满足基本需求的单个模式服务器的部署允许单个模式服务器平均超过阈值量的时间是活动的。例如，所部署的单个模式服务器的数量可以限于在一天之中平均80%的时间是活动的。在峰值使用周期期间能够使用多模服务器来与针对工作负荷诸如游戏进行优化的单个模式服务器相组合来容纳需求。

[0051] 现在转到图5，根据本发明的一方面，显示多模服务器内不同类型的资源。多用途服务器包括母板500。母板500包括针对游戏进行优化的第一类型的计算资源。第一类型的计算资源包括游戏优化CPU 510、游戏优化GPU 512和视频编码器514。这举例说明：如在本文所用的，计算资源能够包括多个硬件项。并且，虽然未显示，但是计算资源可以包括存储器和其他的支持资源的组件。例如，CUP能够具有专用DRAM存储器。本发明的各方面不限于与单独的视频编码器一起使用。视频编码器能够是CPU或GPU的一部分。

[0052] 在一个方面，游戏优化CPU 510和游戏优化GPU 512是在市售的游戏机中找到的相同芯片。示例性游戏机包括Xbox 360、Sony(索尼)的PlayStation® 家族、Xbox One和Nintendo(任天堂)的Wii™等等。与游戏优化的计算资源相关联的硬件配置能够被配置，以允许针对市售的游戏机进行编写的游戏运行在多模服务器上而不修改游戏代码并且采用

这些游戏与游戏机中的硬件交互的相同方式来与硬件交互。例如，利用游戏机的GPU执行的过程能够利用游戏优化的计算资源的GPU来执行。

[0053] 游戏优化CPU 510、游戏优化GPU 512和视频编码器514全部能够被耦合至母板内的插座。在一个方面，计算资源被去激活，以便通过关闭该资源所附着的插座而转变至低功率模式。

[0054] 母板500上的通用计算资源包括企业优化CPU 520。在游戏模式中，如图5所示，企业优化CPU 520正汲取0W。相反，对于游戏模式中150W的总使用，游戏优化CPU 510正汲取50W，游戏优化GPU 512正汲取90W，而视频编码器514正汲取10W。

[0055] 现在转到图6，根据本发明的一方面，举例说明通用计算模式中的功率使用。现在，游戏优化CPU 510、游戏优化GPU 512和游戏编码器514都正在汲取0W。相反，企业优化CPU 520正在汲取150W。在一个方面，与不同模式相关联的计算资源的额定功率消耗实质上是相等的。在这种情况中，与游戏模式相关联的计算资源汲取和与企业或一般计算模式相关联的计算资源相同的功率量。在一个方面，多模服务器的冷却系统只能提供节冷却(knot cooling)，以便一种类型的计算资源在给定的时间点是活动的。

[0056] 现在转到图7，根据本发明的一方面，提供用于在数据中心内管理工作负荷的方法700。方法700可以利用在数据中心内管理工作负荷的控制构造来执行。

[0057] 在步骤710，实质上多个多模服务器内所有的第一类型的计算资源在第一时间周期期间被设置为低功率模式。每一个多模服务器具有针对不同类型的工作进行优化的多个计算资源。多个计算资源至少包括第一类型的计算资源和第二类型的计算资源。在一个方面，第一时间周期对于第一类型的计算资源被优化来处理的工作负荷而言对应于低需求周期。例如，针对游戏工作负荷的低需求周期可以发生在白天(during the day)，并且游戏优化服务器可以在这个时间周期期间被设置为低功率模式。

[0058] 在步骤720，实质上多个多模服务器内所有的第二类型的计算资源在第二时间周期被设置为低功率模式。第一时间周期与第二时间周期在一个方面实质上不重叠。不重叠时间周期允许第一类型的计算资源和第二类型的计算资源满足其被优化来应付的计算负荷的峰值需求。

[0059] 在个别多模服务器内是活动的计算资源的类型能够基于工作负荷需求进行调节。当第一类型的工作负荷的供给增加时，适于处理第一类型的工作负荷的资源能够被激活。当针对不同类型的工作负荷的需求改变时，状态的混合可以存在于多个多模服务器中。例如，所有的多模服务器能够被设置来处理第一类型的工作负荷，三分之一的多模服务器能够被设置来处理第一类型的工作负荷，一半的多模服务器能够被设置来处理第一类型的工作负荷，或者没有多模服务器能够被设置来处理第一类型的工作负荷。没有被设置来应付第一类型的工作负荷的多模服务器能够被设置来使用不同类型的工作负荷。

[0060] 本发明的各方面已被描述为说明性的而非限制性的。将明白：某些特性和次组合是实用的并且可以被采用而不参考其他的特性与次组合。这利用权利要求书的范畴来设想并位于权利要求书的范畴之内。

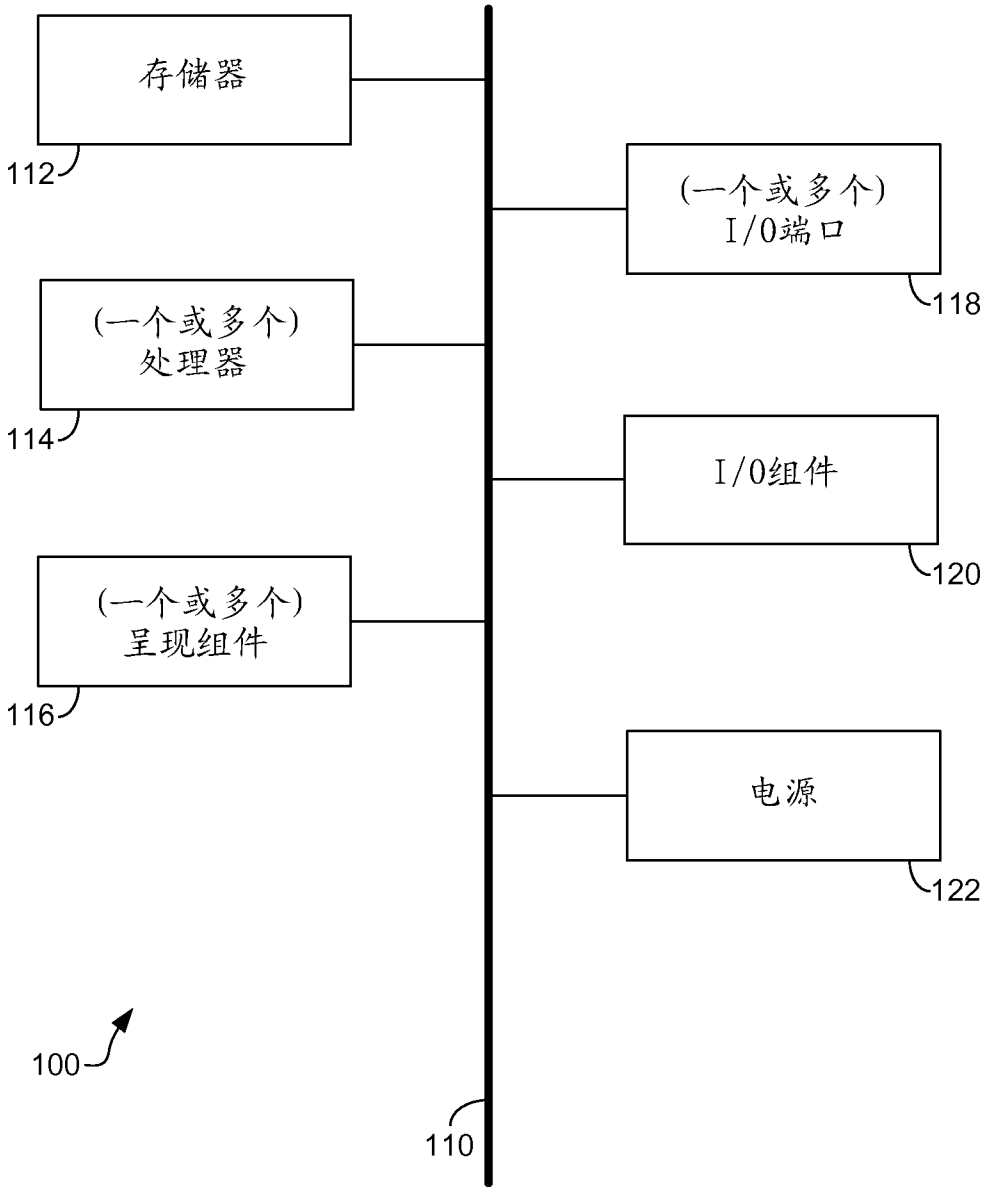
存储器
112
(一个或多个)
处理器
114
(一个或多个)
呈现组件
116
(一个或多个)
I/O端口
118
I/O组件
120
电源
122
100
110

图 1

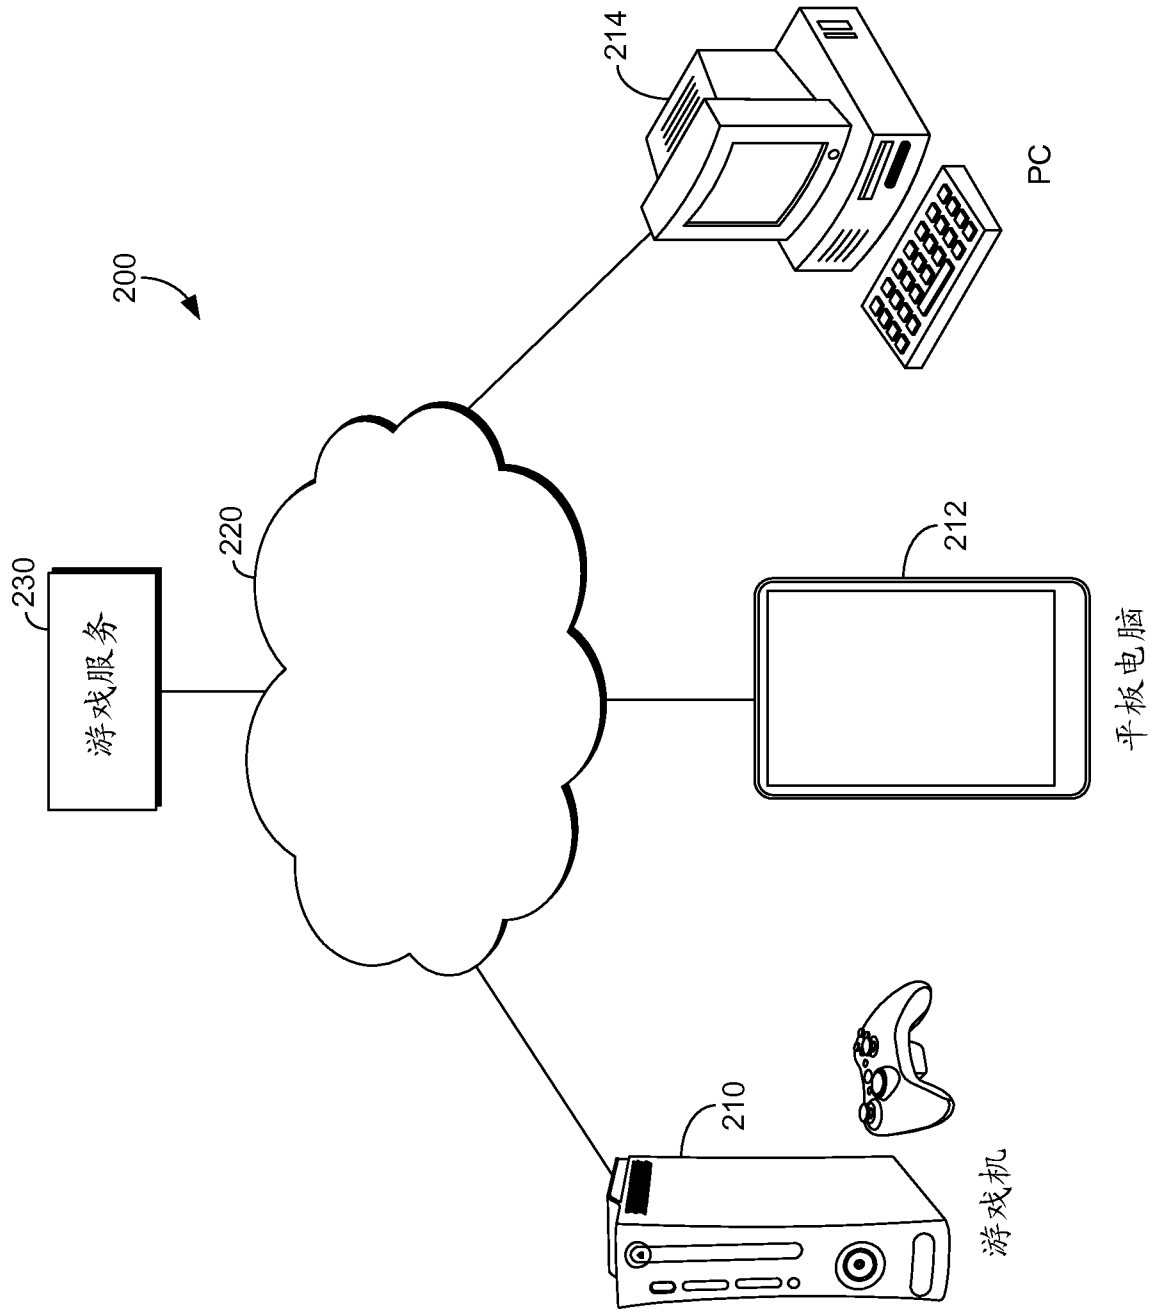

图 2

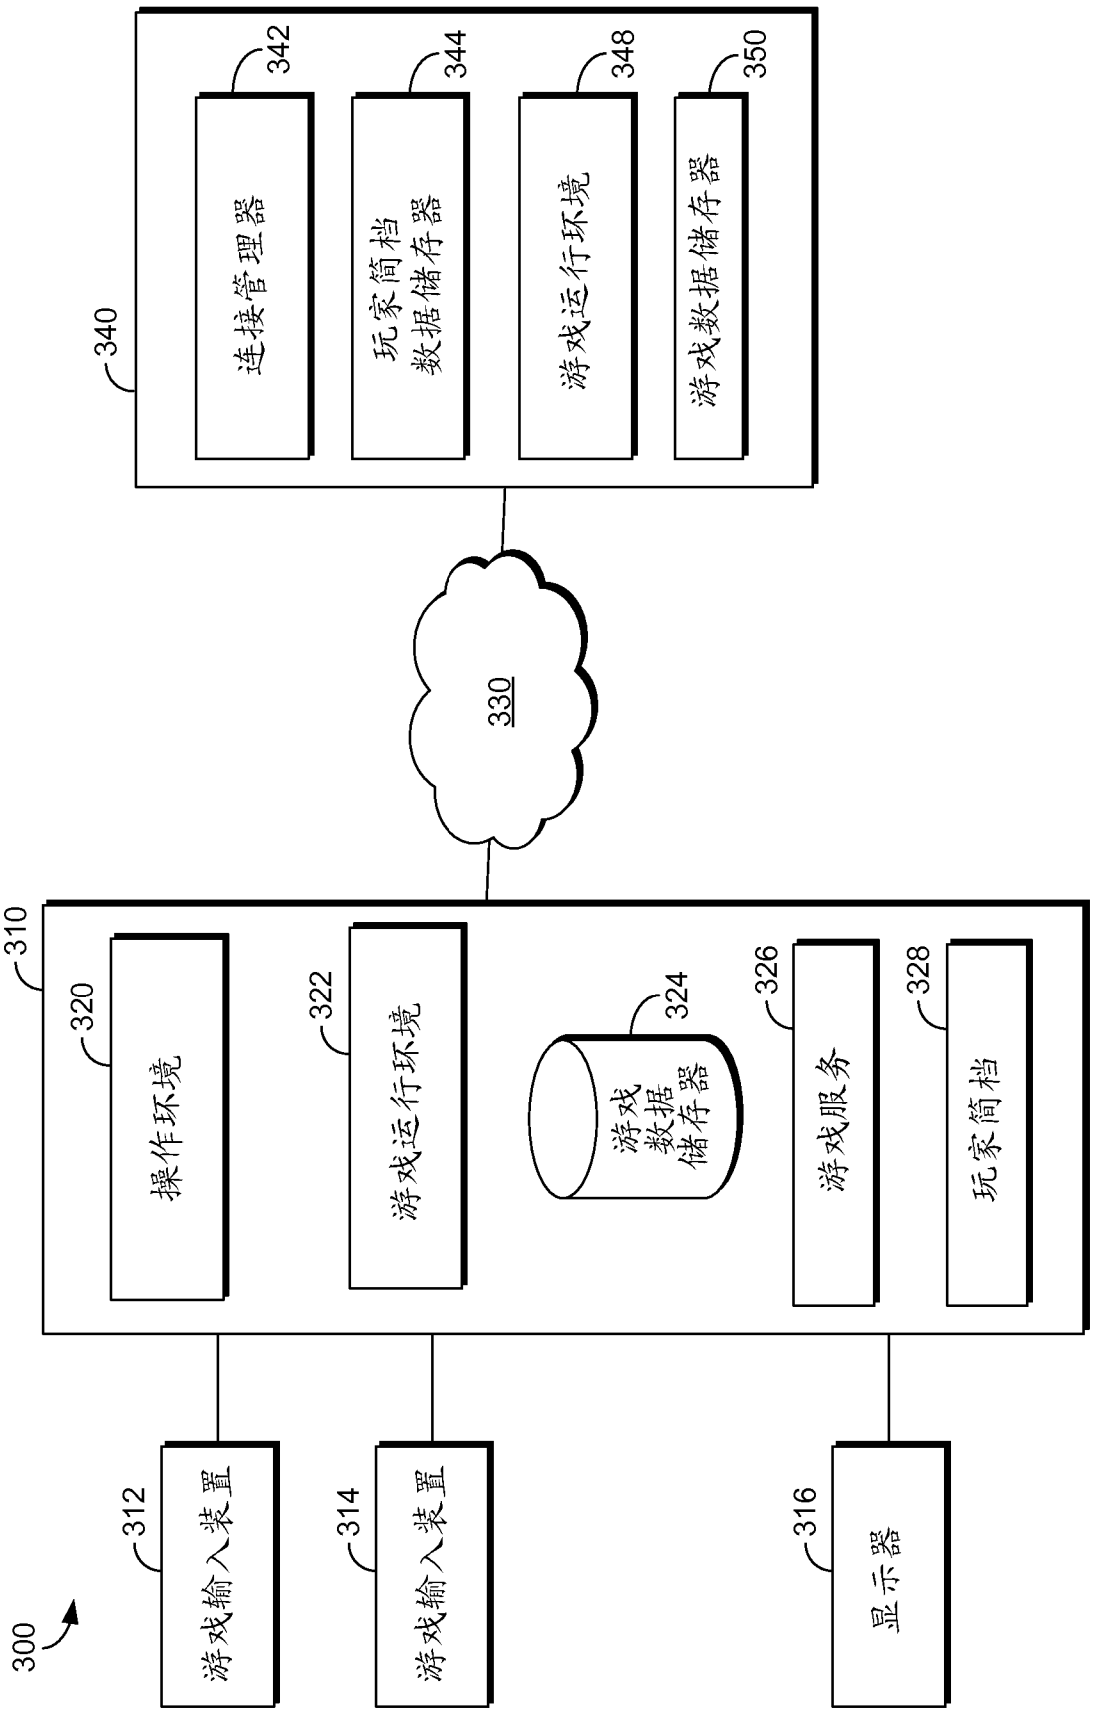
340
342
连接管理器
344
玩家简档
数据储存器
348
游戏运行环境
350
游戏数据储存器
330
310
320
操作环境
322
游戏运行环境
324
游戏
数据
储存器
326
游戏服务
328
玩家简档
312
游戏输入装置
314
游戏输入装置
316
显示器
300

图 3

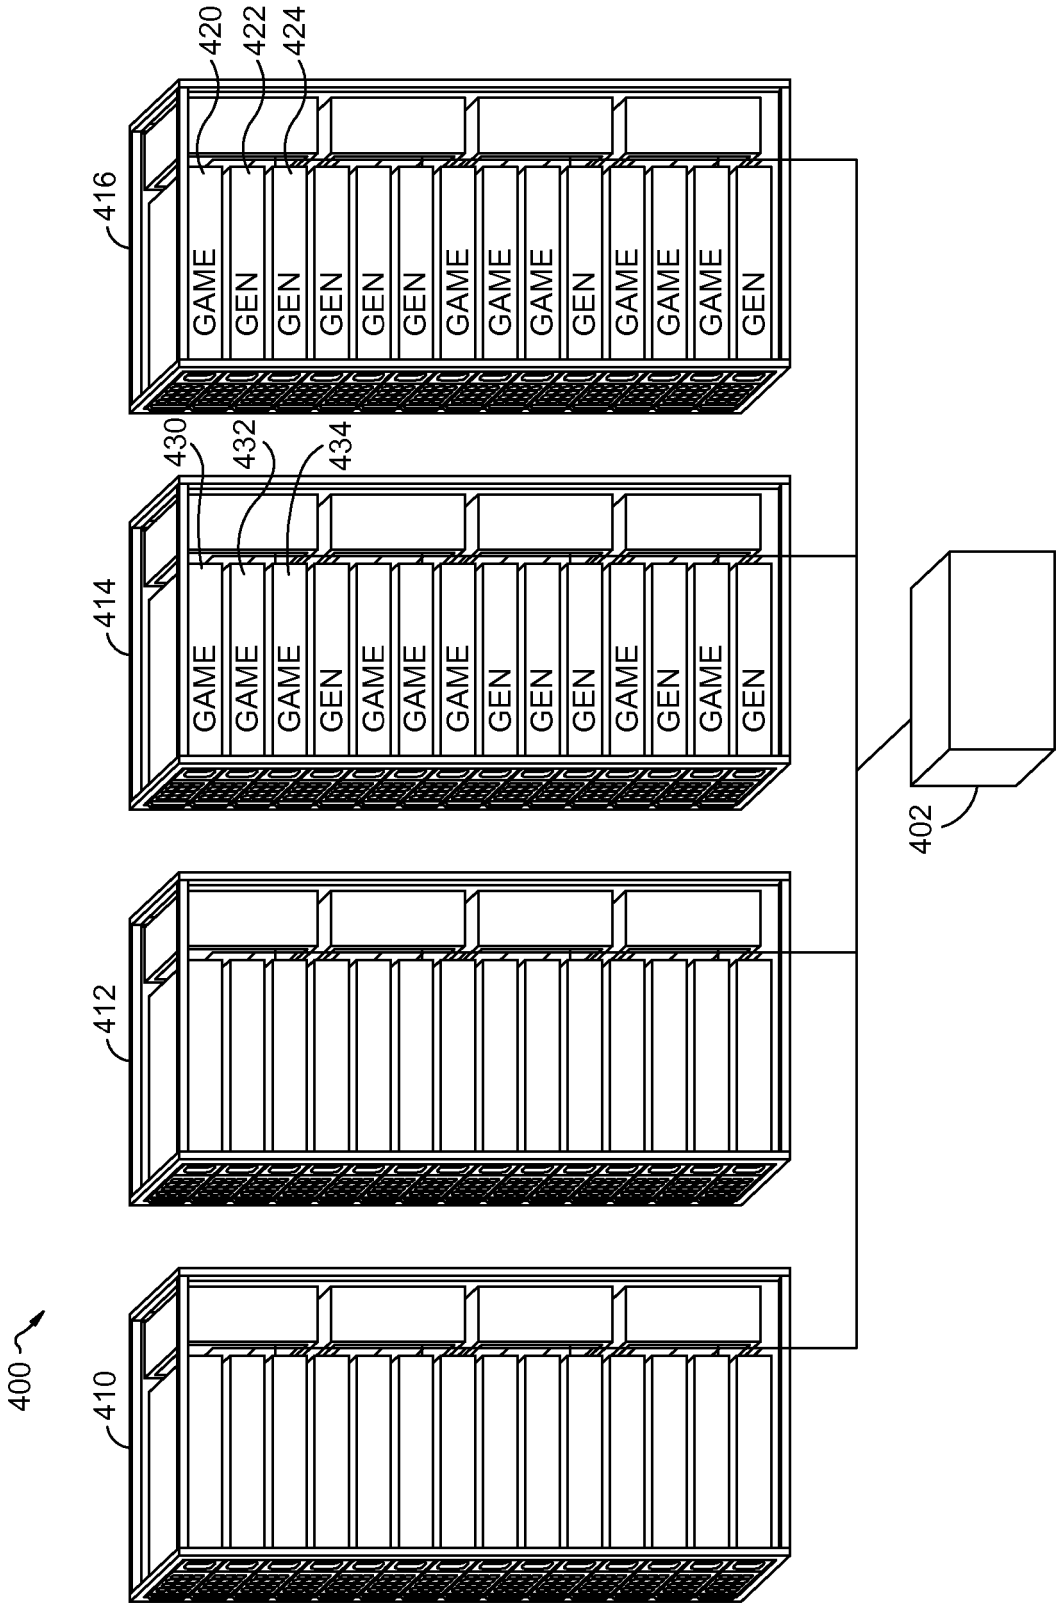
420
422
424
416
GAME
GEN
GEN
GEN
GEN
GEN
GAME
GAME
GAME
GEN
GAME
GAME
GAME
GEN
430
432
434
414
GAME
GAME
GAME
GEN
GAME
GAME
GAME
GEN
GEN
GEN
GAME
GEN
GAME
GEN
402
412
400
410

图 4

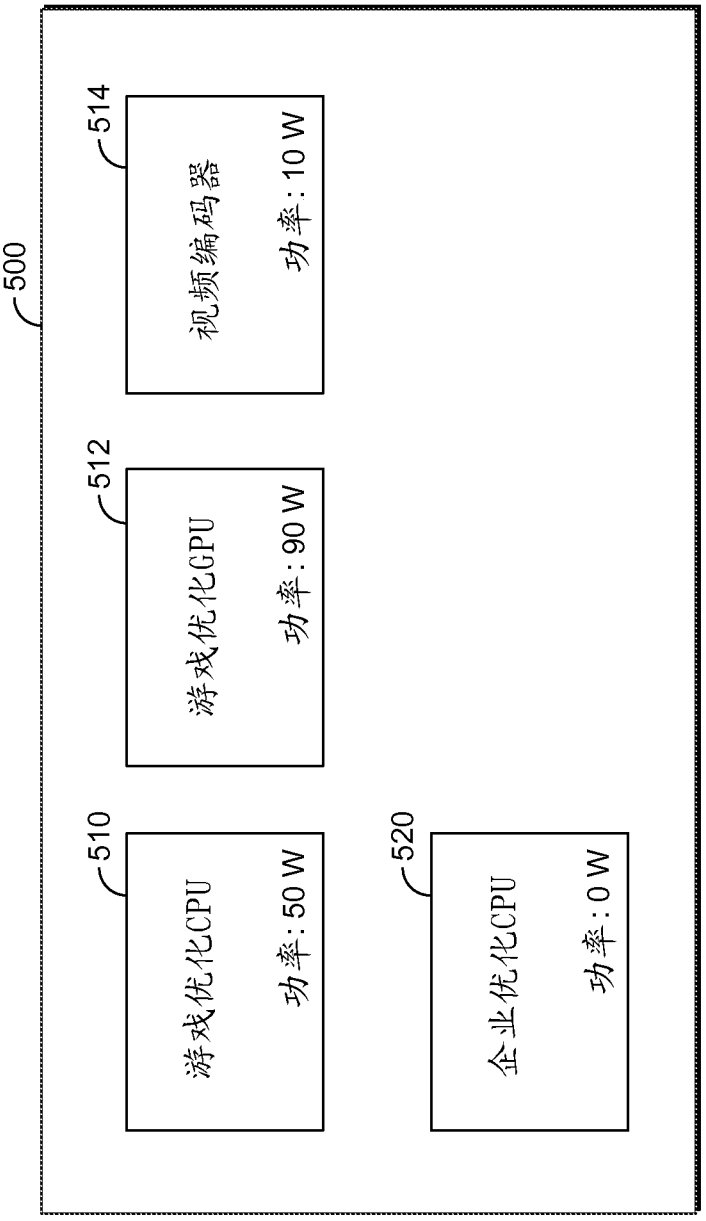
500
514
视频编码器
功率：10 W
512
游戏优化GPU
功率：90 W
510
游戏优化CPU
功率：50 W
520
企业优化CPU
功率：0 W

图 5

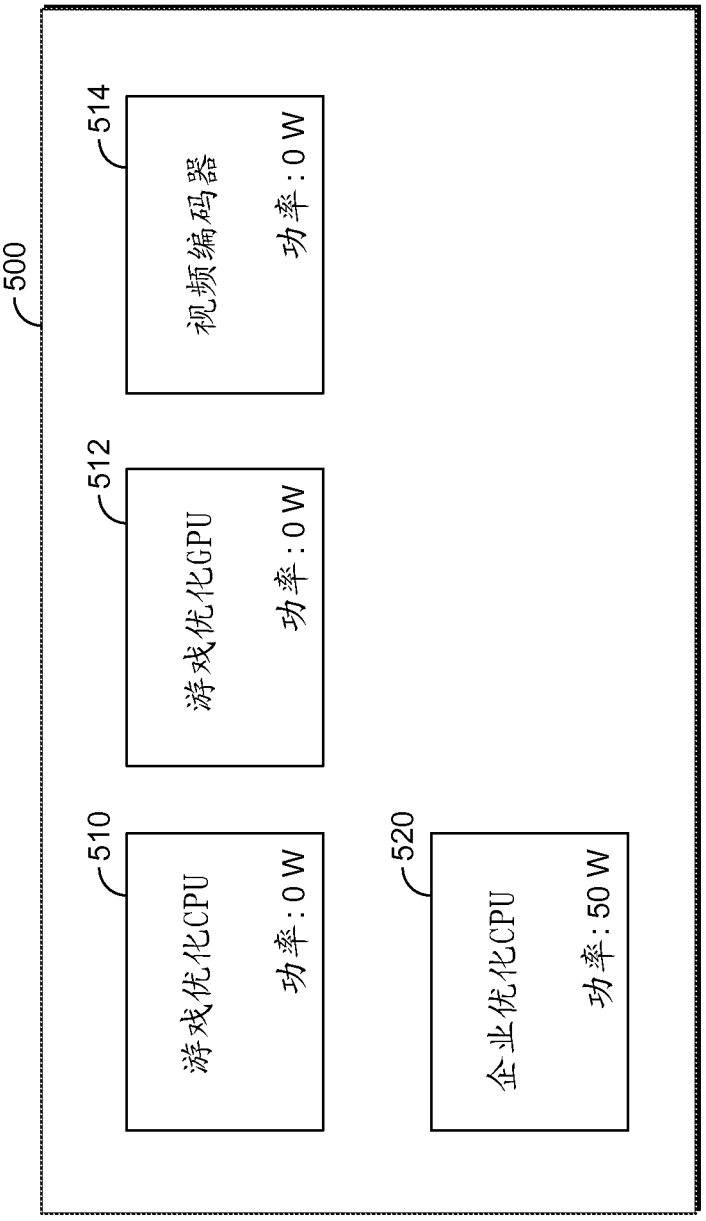
500
514
视频编码器
功率：0 W
512
游戏优化GPU
功率：0 W
510
游戏优化CPU
功率：0 W
520
企业优化CPU
功率：50 W

图 6

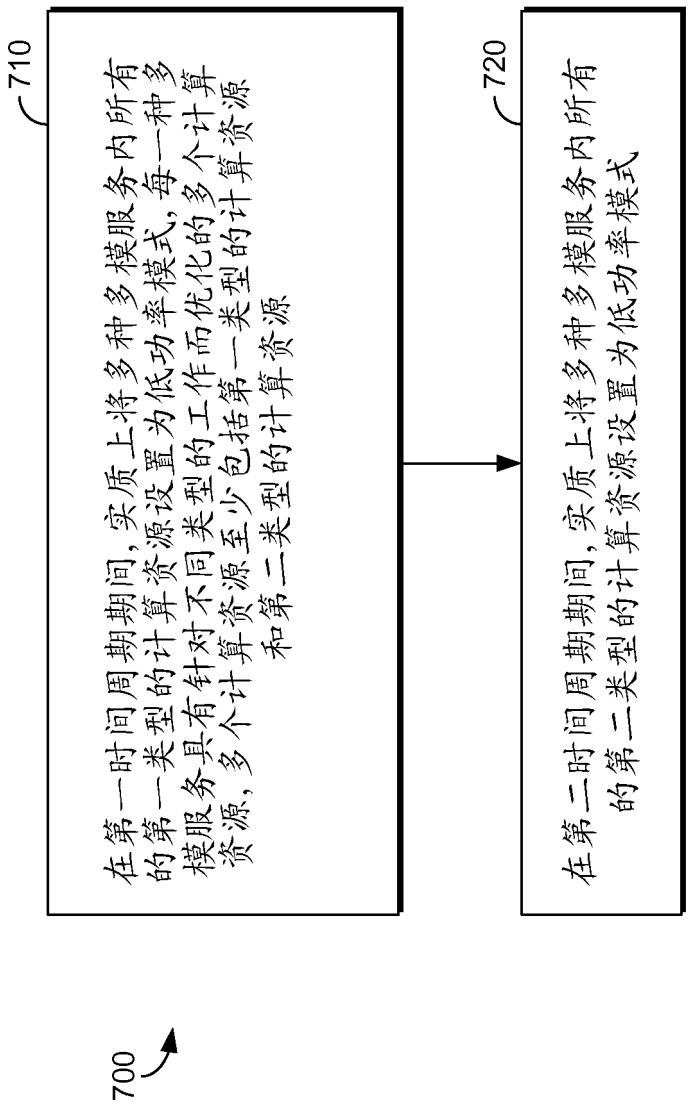
710
在第一时间周期期间，实质上将多种多模服务内所有的第一类型的计算资源设置为低功率模式，每一种多模服务具有针对不同类型的工作而优化的多个计算资源，多个计算资源至少包括第一类型的计算资源和第二类型的计算资源
720
在第二时间周期期间，实质上将多种多模服务内所有的第二类型的计算资源设置为低功率模式
700

图 7