(19) **United States**
(12) **Patent Application Publication**
Chin

(10) **Pub. No.: US 2008/0114997 A1**
(43) **Pub. Date: May 15, 2008**

(54) **SYSTEM AND METHOD OF DYNAMIC POWER MANAGEMENT**

(75) Inventor: **David Kun Chin**, Los Altos, CA (US)

Correspondence Address:
STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.
1100 NEW YORK AVENUE, N.W.
WASHINGTON, DC 20005

(73) Assignee: **Broadcom Corporation**

(21) Appl. No.: **11/979,919**

(22) Filed: **Nov. 9, 2007**

Related U.S. Application Data

(60) Provisional application No. 60/858,937, filed on Nov. 15, 2006.

Publication Classification

(51) **Int. Cl.**
G06F 1/28 (2006.01)

(52) **U.S. Cl.** **713/321**

(57) **ABSTRACT**

In a communications system, such as a Power-Over-Ethernet system, where power supply equipment (PSE) supplies power to powered requiring devices (PRDs), a system and method of dynamic power management is implemented. The system and method monitors the power consumed at each port by the PRDs. Based on this monitoring, the PSE dynamically determines the minimum power which can be allocated to each PRD, and so dynamically maximizes the available reserve power. The PSE maintains a queue or queues wherein PRDs are listed in order of a power allocation priority. When additional power is available, the PSE preferentially allocates power to a PRD or PRDs which have higher priority. The system and method of the present invention minimizes the power allocated to each individual network device, as a result of which the total number of network devices that can be supported with the available power may be maximized.

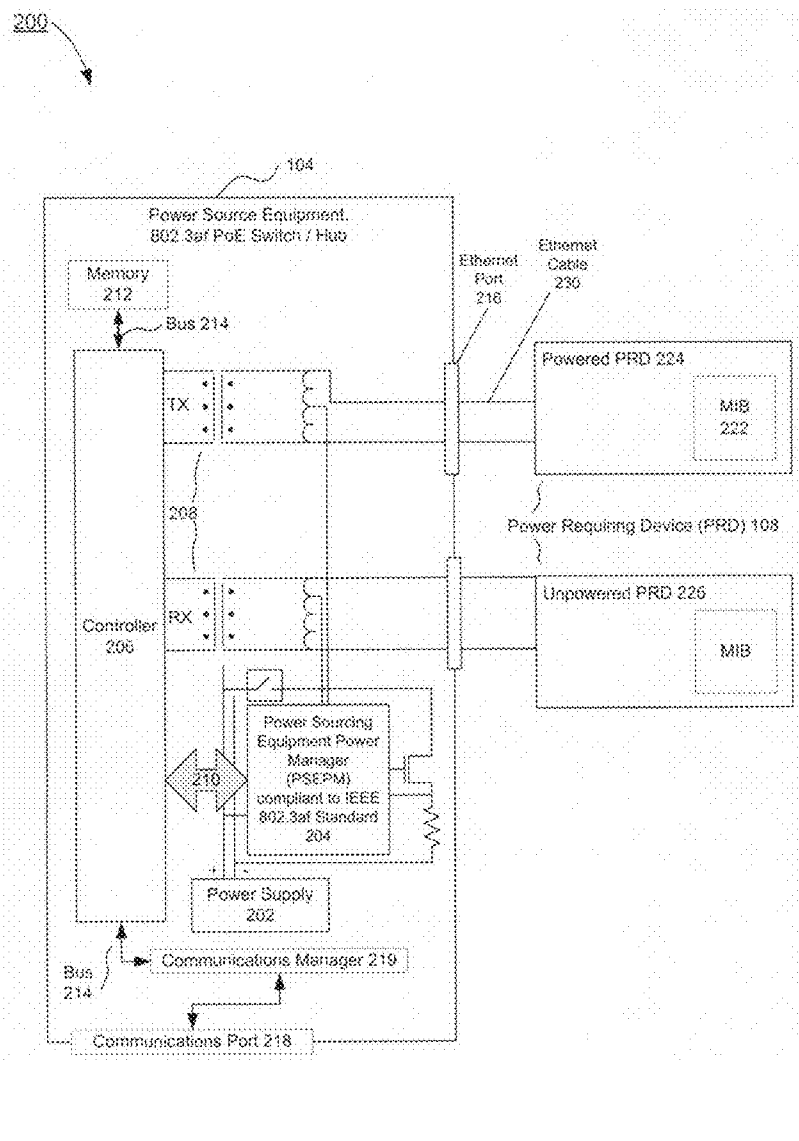

[0055] FIGS. 1A and 1B offer illustrations of Power-Over-Ethernet (PoE) system deployments 100 and 140… Exemplary deployment 100 in FIG. 1A shows an end span configuration. In this configuration, an uninterruptible power supply (UPS) 102 is connected to a communications system 104. The IEEE standard 802.3af is embodied within the technology of the communications system itself, which in this case is a PoE Ethernet switch. That is to say, the Ethernet switch 104 actually contains the power source equipment (PSE), and in fact is considered to be the power source equipment.

[0056] The PoE switch ultimately provides both data and power to Power Requiring Devices (PRDs) 108, such as the PoE VoIP phone, APs, and video camera illustrated.

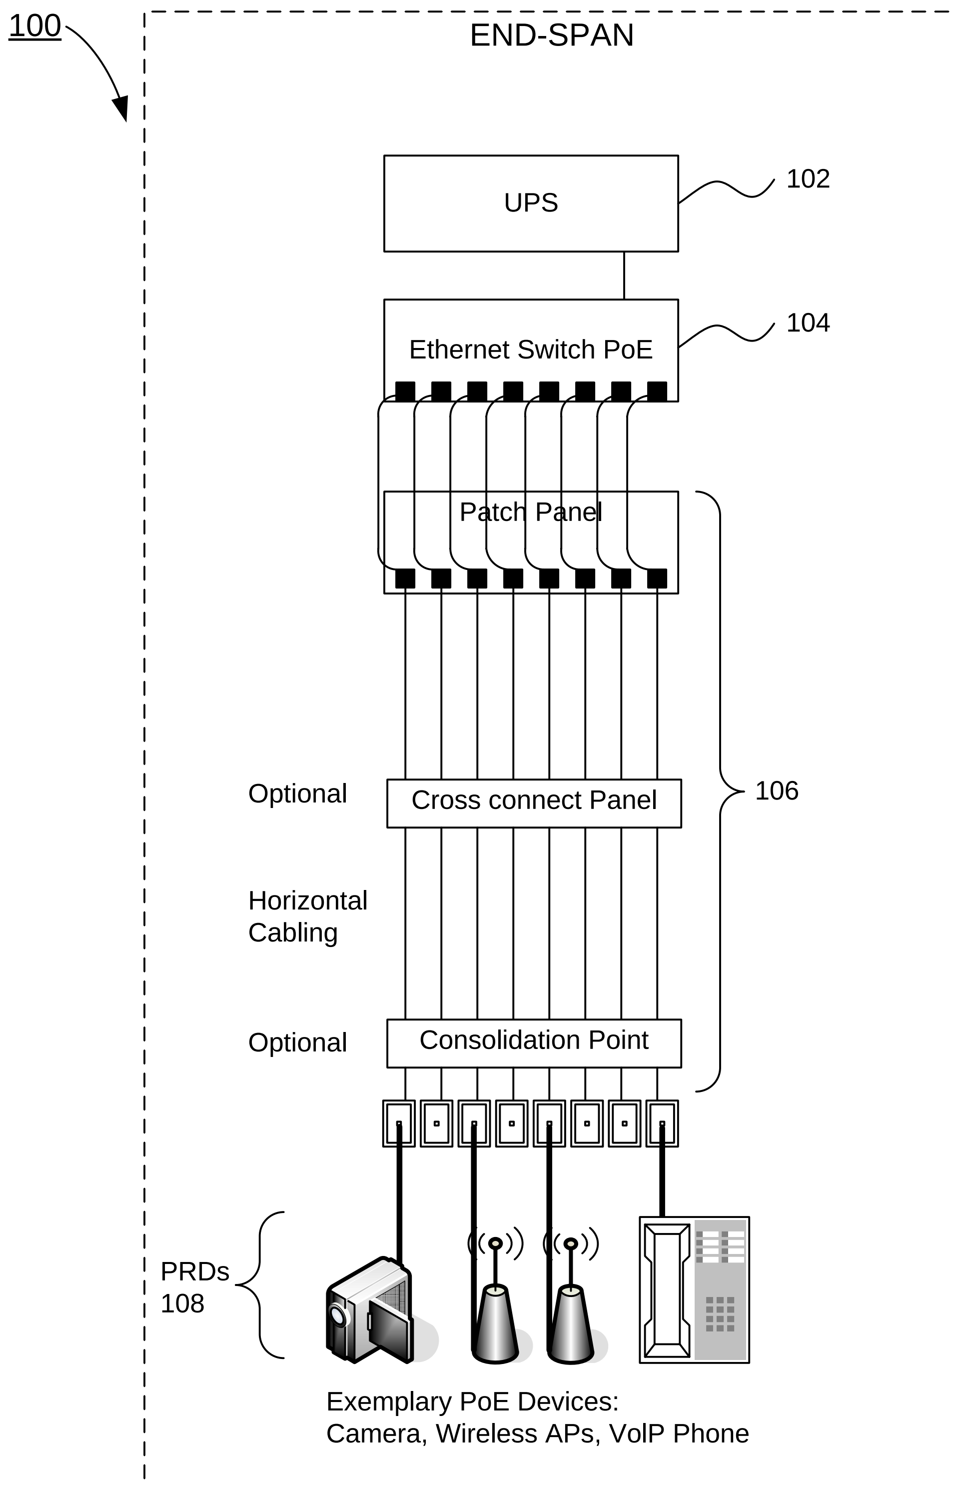

FIG. 1A

[0060] FIG. 2 is a block diagram 200 showing both part of the internal structure of representative communications system 104, such as a PoE switch 104, as well as part of the internal structure of representative PRD 108….

[0061] Note that in the ordinary and routine course of operations of PSE 104 and one or more PRDs 108 which may be attached to PSE 104, a PRD 108 which is an unpowered PRD 226 may become a powered PRD 224, and similarly a PRD 108 which is a powered PRD 224 may become an unpowered PRD 226….

[0063] PoE switch 104 contains an internal power supply 202 which in turn may be fed from external UPS 102. PoE switch 104 also contains power sourcing equipment power manager (PSEPM) 204 which is compliant to the PoE IEEE 802.3af standard.

[0064] PSE 104 also contains a controller 206. Controller 206 may be a general purpose microprocessor or it may be another ASIC….

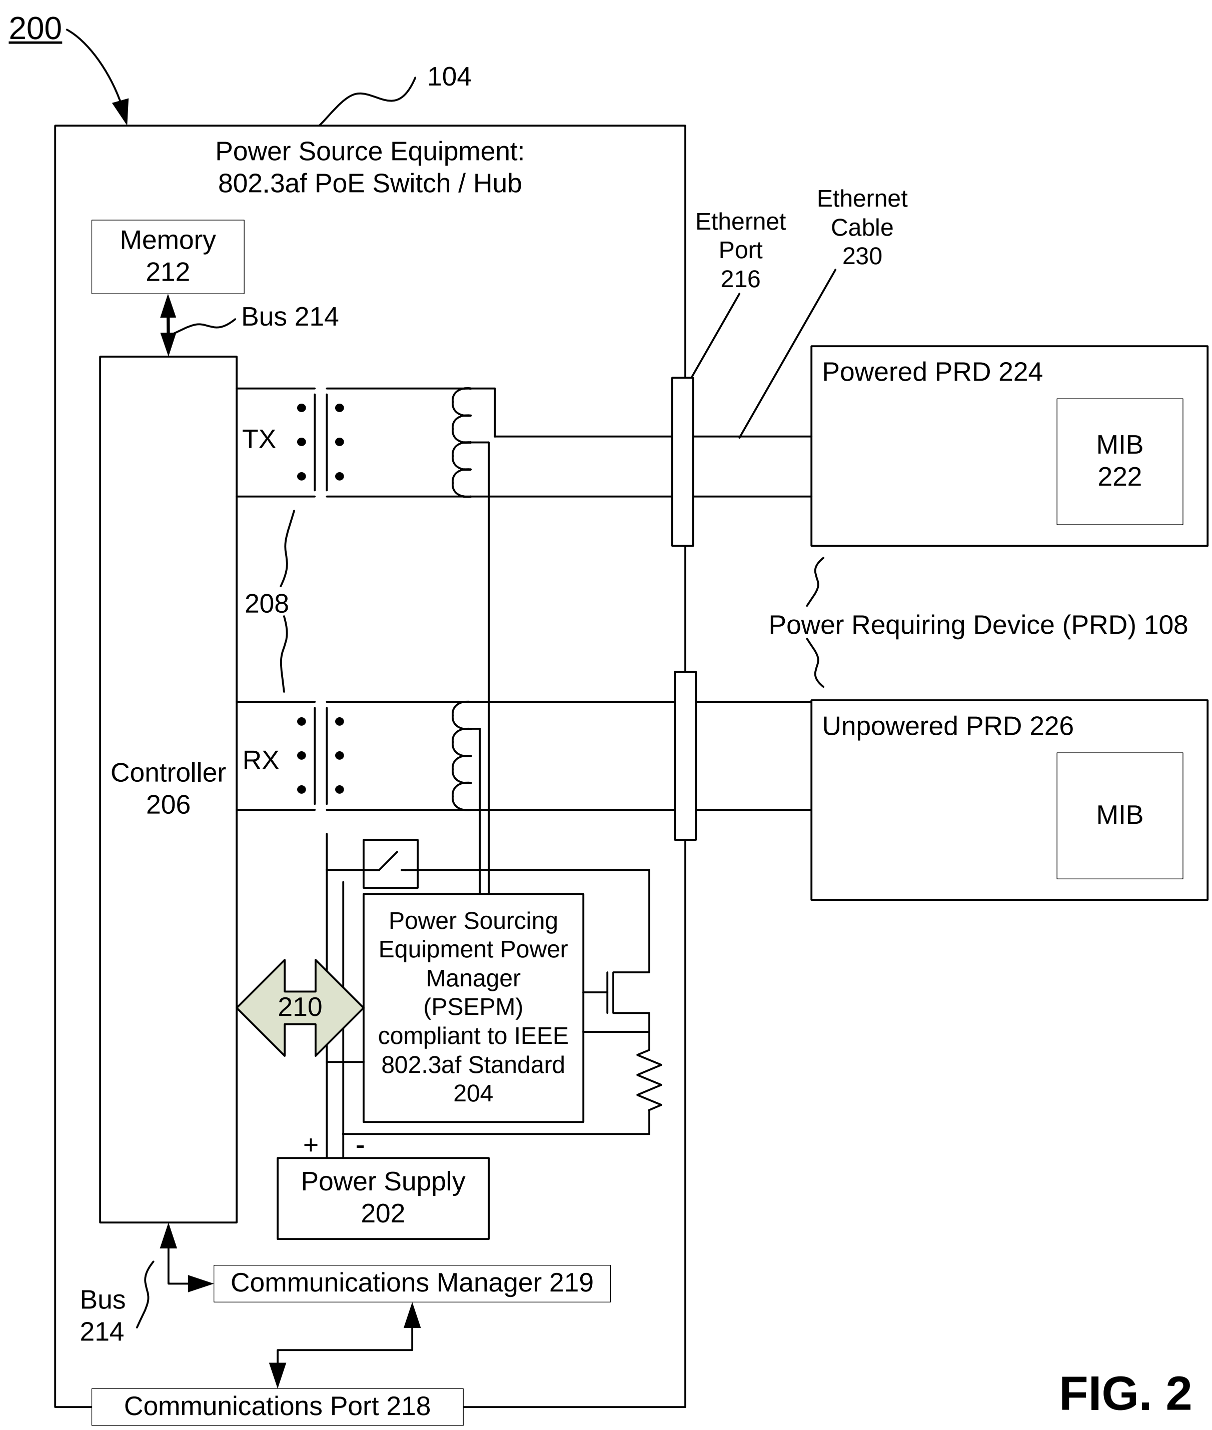

FIG. 2

[0076] FIG. 3A illustrates a method 300 according to one embodiment of the present invention, whereby a PSE 104 and associated PRDs 108 are initially powered on, and wherein the current method of dynamic power management is initiated.

[0077] In initial step 302 the PSE 104 is powered on. In step 304, the initial power requirements of the initially attached PRDs 108 are then determined. Communication is established between PSE 104 and the PRDs 108 according to the 802.3af standard via methods which are well known in the art.

[0078] During the initial handshaking between PSE 104 and PRD 108, PSE 104 determines the initial power allocation to be made to each PRD 108 based on any of several criteria. These criteria may include, for example and without limitation, a power classification which is built into the MIB 222 of the PRD 108, or a maximum power which can be allocated to the communications ports 216 of the PSE 108, or a predefined power allocation which has been defined by a user of the system and which is stored in memory 212 of PSE 104.

[0080] Once power has been allocated to all attached devices, or once the maximum available power has been allocated, the method of dynamic power management 316 begins. The method 316 has two phases, a dynamic power initialization phase 312, following by a second phase of ongoing dynamic power management 314.

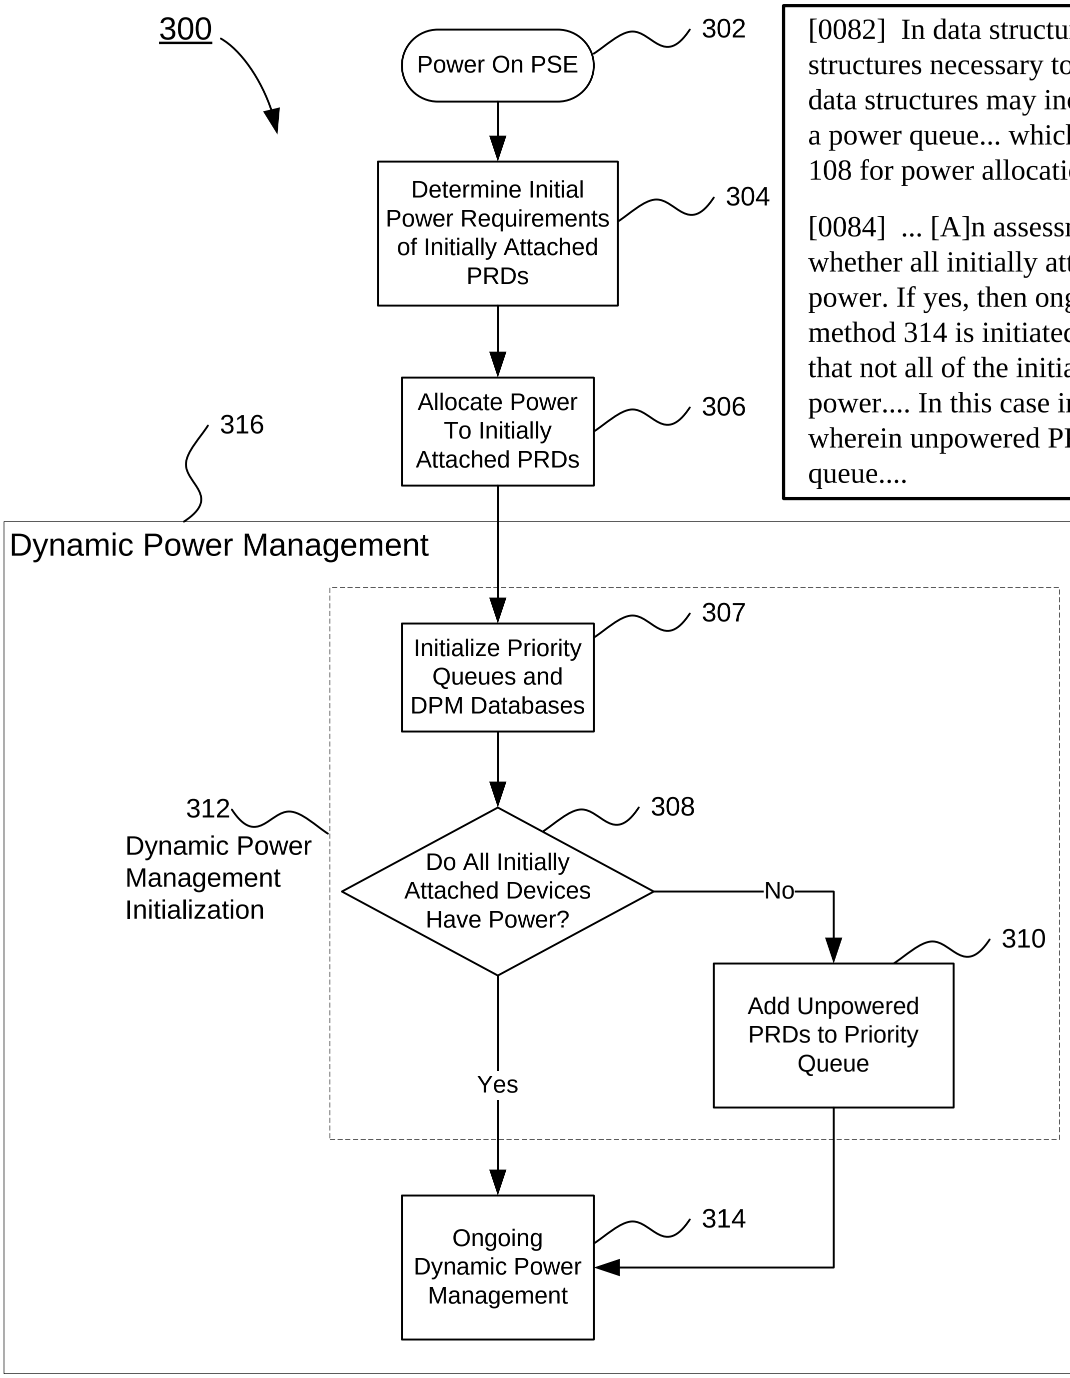

[0082] In data structure initialization step 307, data structures necessary to the method are initialized. These data structures may include a priority queue and possibly a power queue... which may indicate the priority of PRDs 108 for power allocation or power deallocation. ….

[0084] ... [A]n assessment is made in step 308 as to whether all initially attached devices have been allocated power. If yes, then ongoing dynamic power management method 314 is initiated. In step 308, it may be determined that not all of the initially attached devices were allocated power.... In this case initialization step 310 is performed, wherein unpowered PRDs 226 are added to the priority queue....

FIG. 3A

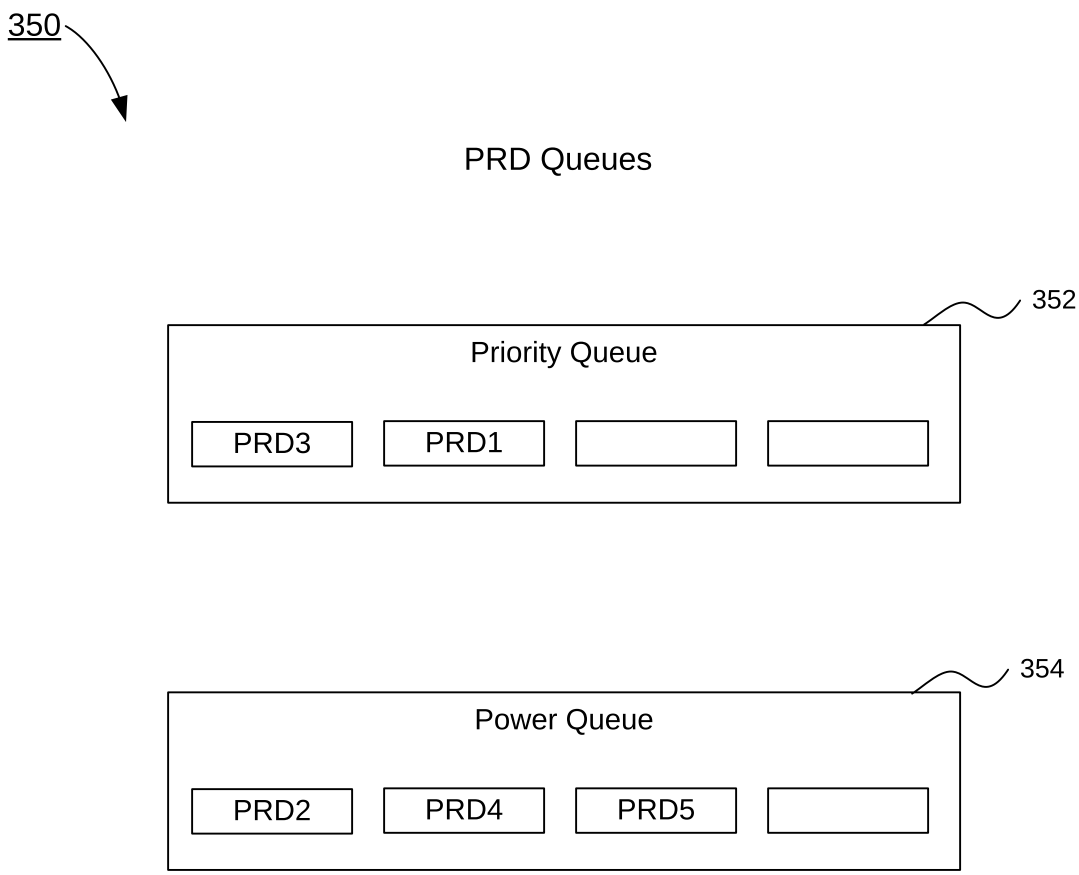

FIG. 3B

[0086] FIG. 3B is a representative depiction of several PRD queues 350. In one embodiment, the queues are used to store or to indicate the priority of PRDs for power allocation. In an alternative embodiment, the priority of PRDs may be stored or indicated by means other than a queue or queues, including, for example and without limitation, a database of PRDs wherein the database has a field or other data structure for indicating the priority of PRDs.

[0089] The PRD queue or PRD queues store the identity of devices using some kind of unique device identifier which may be, for example and without limitation, an IP address of a PRD, the port number of a port to which the PRD is attached, a unique number generated from one or several selected parameters retrieved from a PRD MIB, or any other value which can serve as a unique identifier for a device.

[0090] Priority queue 352 is used to track those PRDs 108 which are connected to PSE 104 but which are not currently powered. Typically, devices are added to priority queue 352 when they are connected to the system and it is determined that there is insufficient power available in order to allocate power to them.

[0091] [A]nother queue is power queue 354. Power queue 354 stores the identity of all PRDs which are currently powered. The purpose of the power queue is to maintain a list of these devices in order of a device priority, wherein the device priority may be assigned by the user.

[0093] FIGS. 4A and 4B together present a flowchart illustrating one embodiment 314A of the method of ongoing dynamic power management 314. The flowchart has been broken into two figures for ease of reading only, and embodies one unified method.

[0094] The steps with shadowed highlighting, namely steps 402, 404, 410, 420, and 430, initiate monitoring of a port 216 and determine the status of a port 216, such as determining if a port 216 has a PRD 108 connected to it or if the port 216 requires power. The steps that do not have shadowed highlighting represent responses made based on the status determinations of steps 404, 410, 420, and 430. The steps described below are performed on a repeated or ongoing basis.

[0095] Step 402 begins with monitoring a port 216 in the communications system.

[0096] In step 404 a determination is made if the port 216 currently being monitored previously required power but no longer requires power. If yes, then in step 406 the method of the present invention stops supplying power to the port 216. In addition, the method in step 408 recalculates a total reserve power (TRP) which is available.

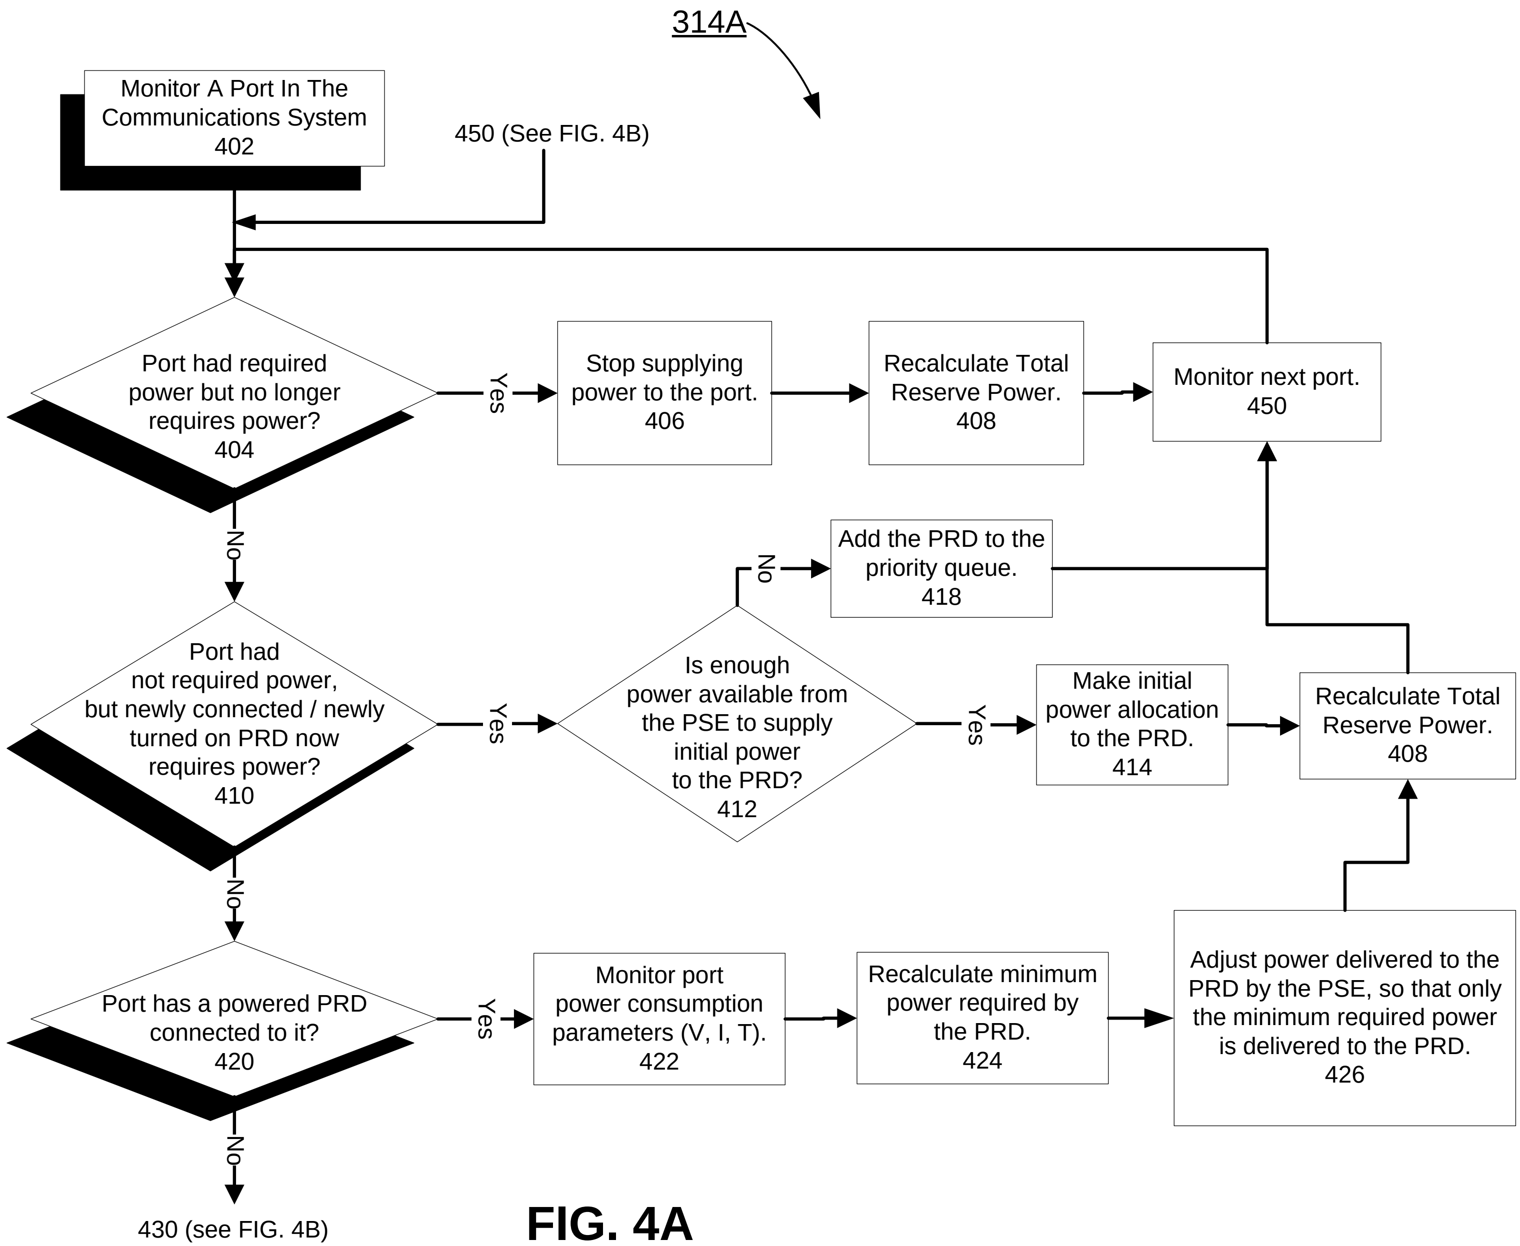

FIG. 4A

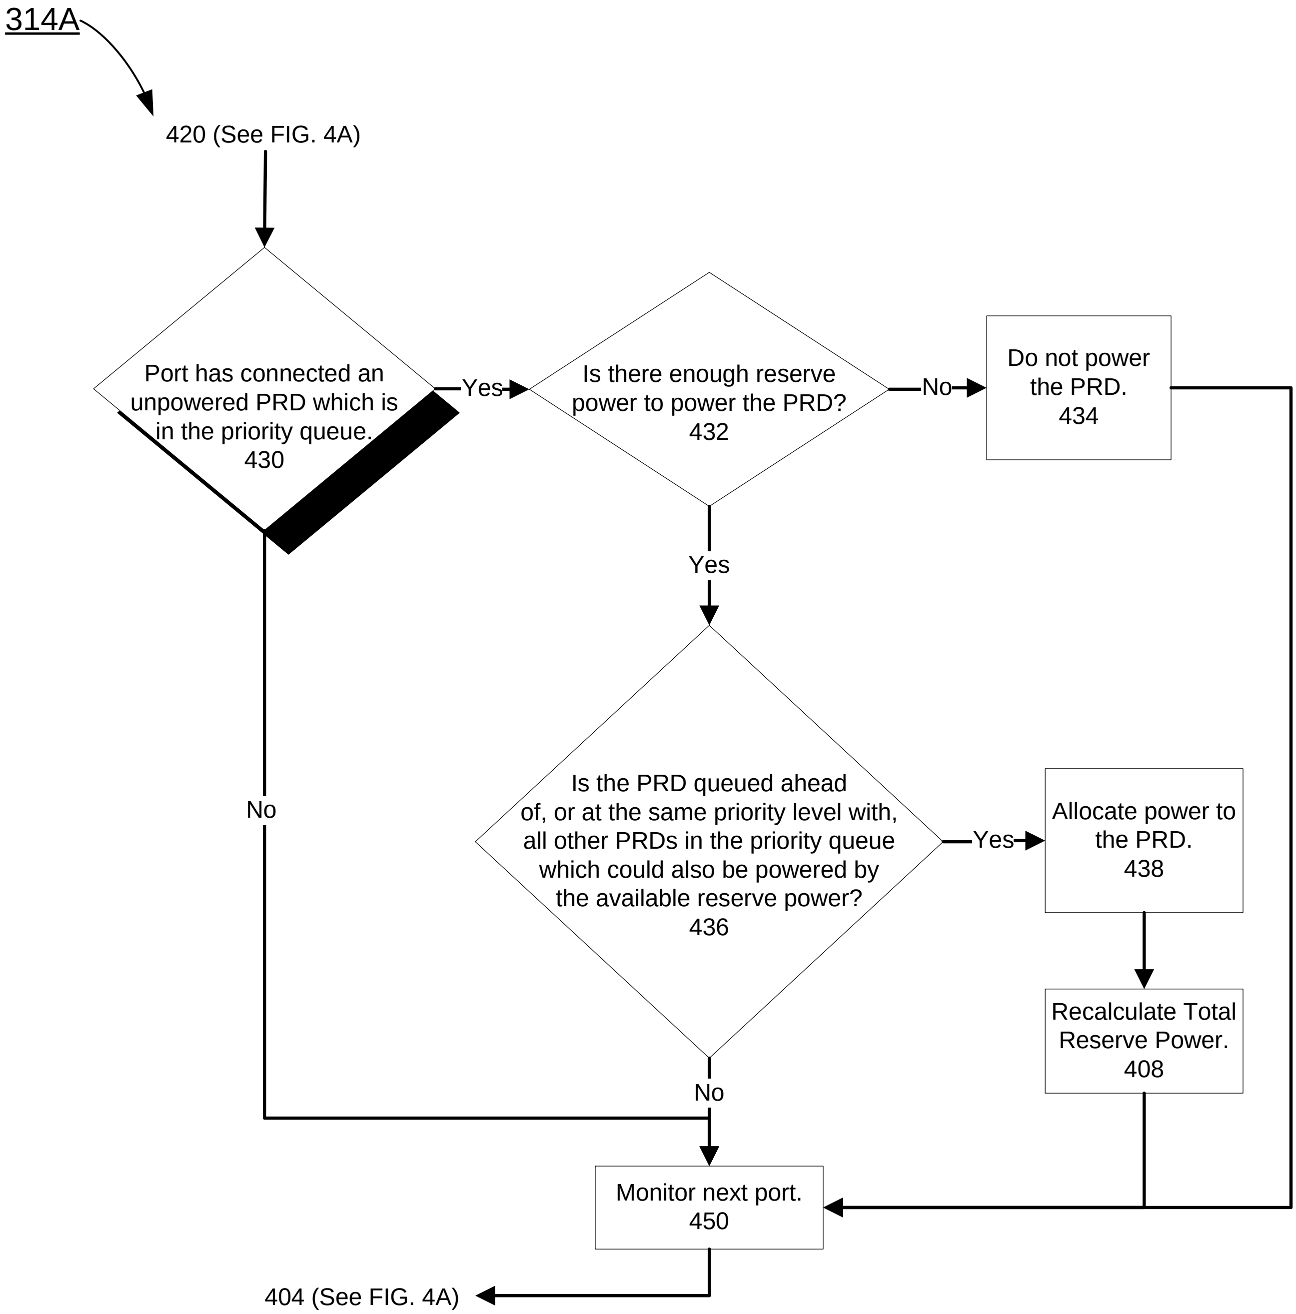

FIG. 4B

[0101] The choice as to whether to calculate the total power used by powered PRDs 224 based on the maximum power consumed or the power being presently consumed; or whether, on the other hand, to calculate the power being consumed based on an average power consumed by the powered PRDs 224, may be based in part on the number of powered PRDs 224 connected to the PSE 104. When only a few PRDs 108 are connected to PSE 104, it is possible that at a given time all the powered PRDs 224 may be drawing their maximum power or near their maximum power. In this case, it may be beneficial to calculate the powered PRD 224 power consumption as the sum of the maximum power which has been observed over time to be drawn by all the powered PRDs 224, or to calculate the power consumption based on the power presently being delivered to the powered PRDs 224.

[0102] However, as more and more powered PRDs 224 are connected to PSE 104, it becomes statistically less likely that all the powered PRDs 224 will be drawing a maximum power at the same time. This is because in typical usage, powered PRDs 224 tend to have peaks and valleys in their levels of power consumption, with the peaks occurring only intermittently. Therefore, as the number of powered PRDs 224 connected to PSE 104 increases, it becomes increasingly reliable to calculate the power used by all powered PRDs 224 as a sum of an average power calculated for each powered PRD 224 over time.

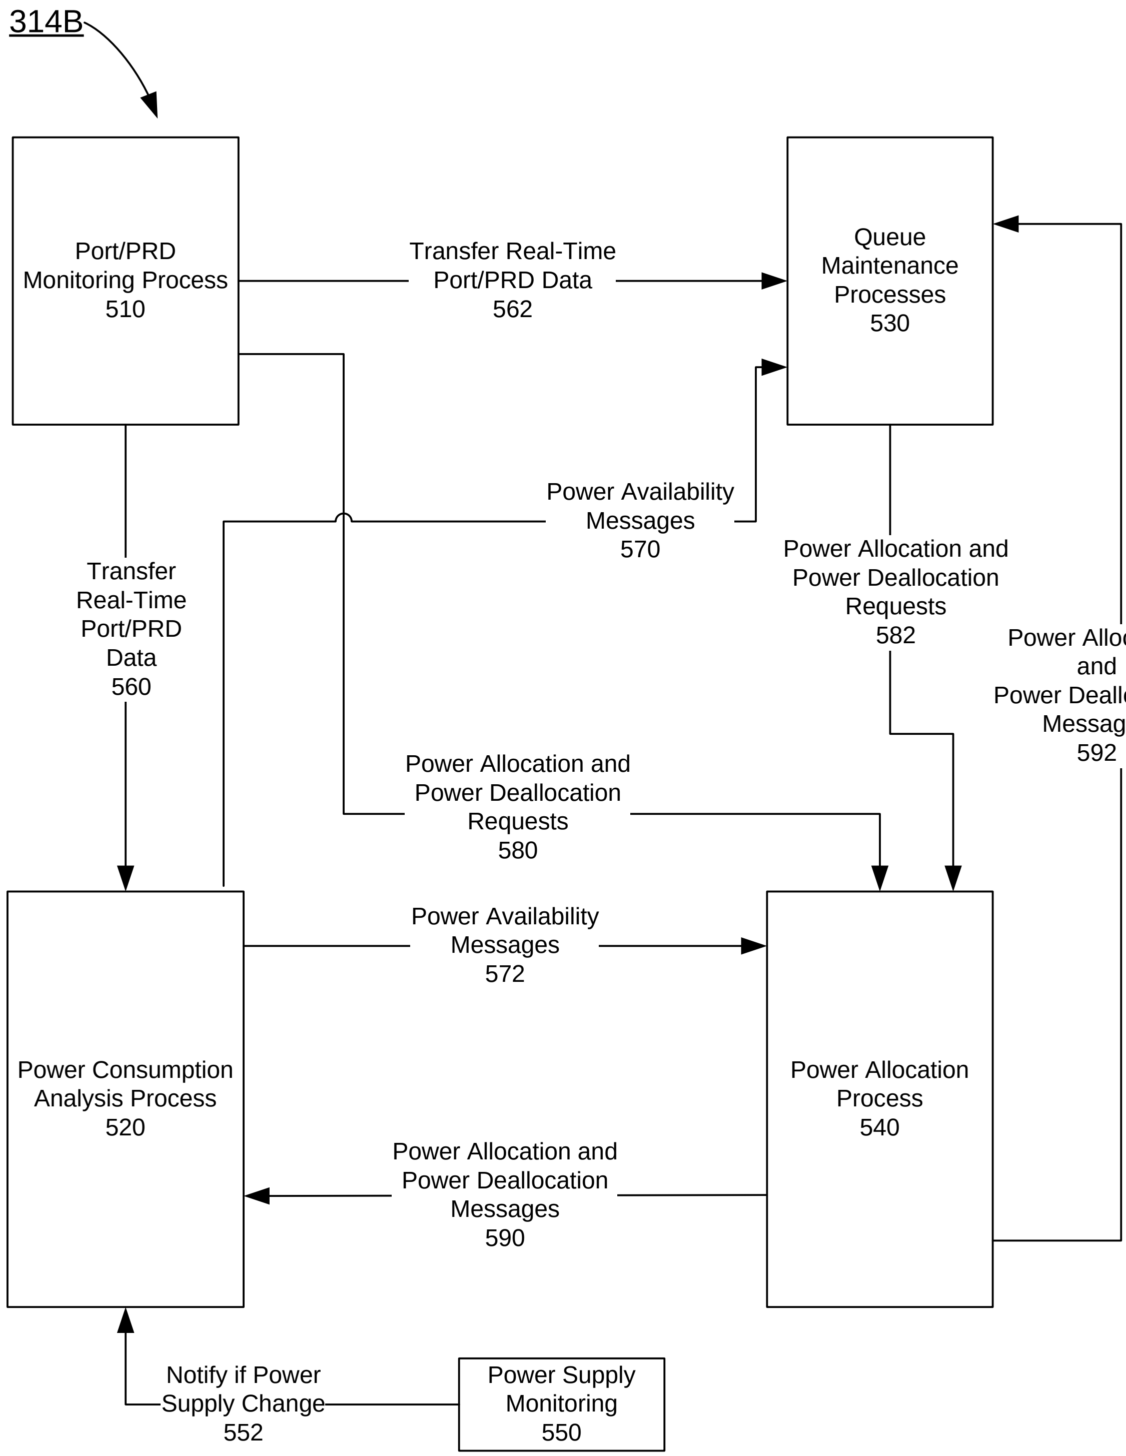

FIG. 5

[0131] FIG. 5 presents another embodiment of the method of ongoing dynamic power management 314 of the present invention, namely, method 314B wherein multiple processes occur in parallel. These processes include port/PRD monitoring process 510, power consumption analysis process 520, queue maintenance process 530, power allocation process 540, and power supply monitoring process 550.

[0132] In turn, these processes communicate with each other through a series of messages. These messages provide necessary information for each process to perform its calculations or to perform the tasks assigned to the process. For example, port/PRD monitoring process 510 provides to power consumption analysis process 520 real-time port/PRD data 560. In particular, port/PRD monitoring process 510 supplies to power consumption analysis process 520 data 560 which may include, for example and without limitation, whether or not a PRD 108 is connected to a port 216, whether the PRD 108 is consuming power, the temperature of the port 216, the current (I) being consumed by the port 216, and the voltage (V) being used to power the port 216.

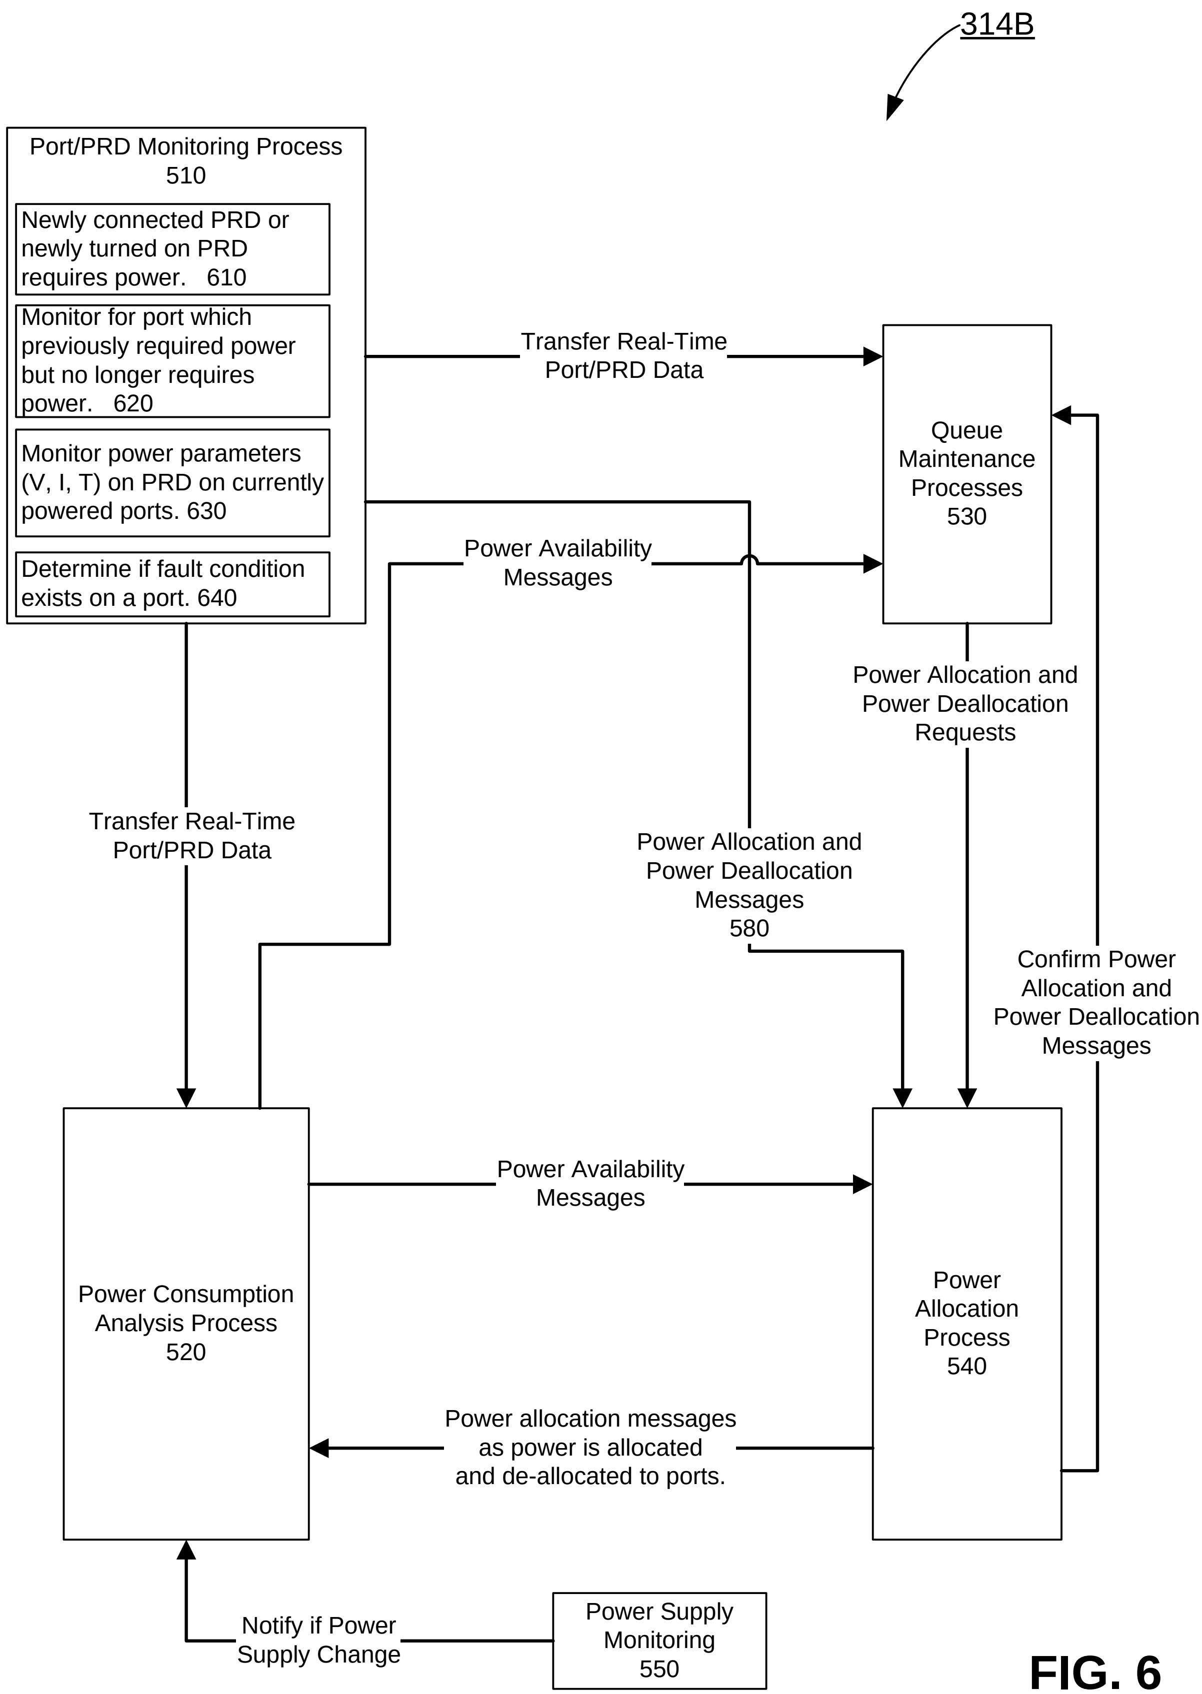

FIG. 6

[0151] FIG. 6 is another view of method 314B. Many features of FIG. 6 previously illustrated in FIG. 5 are repeated here and will not be described again in detail.

[0152]FIG. 6 shows in greater detail the steps involved in port/PRD monitoring process 510.

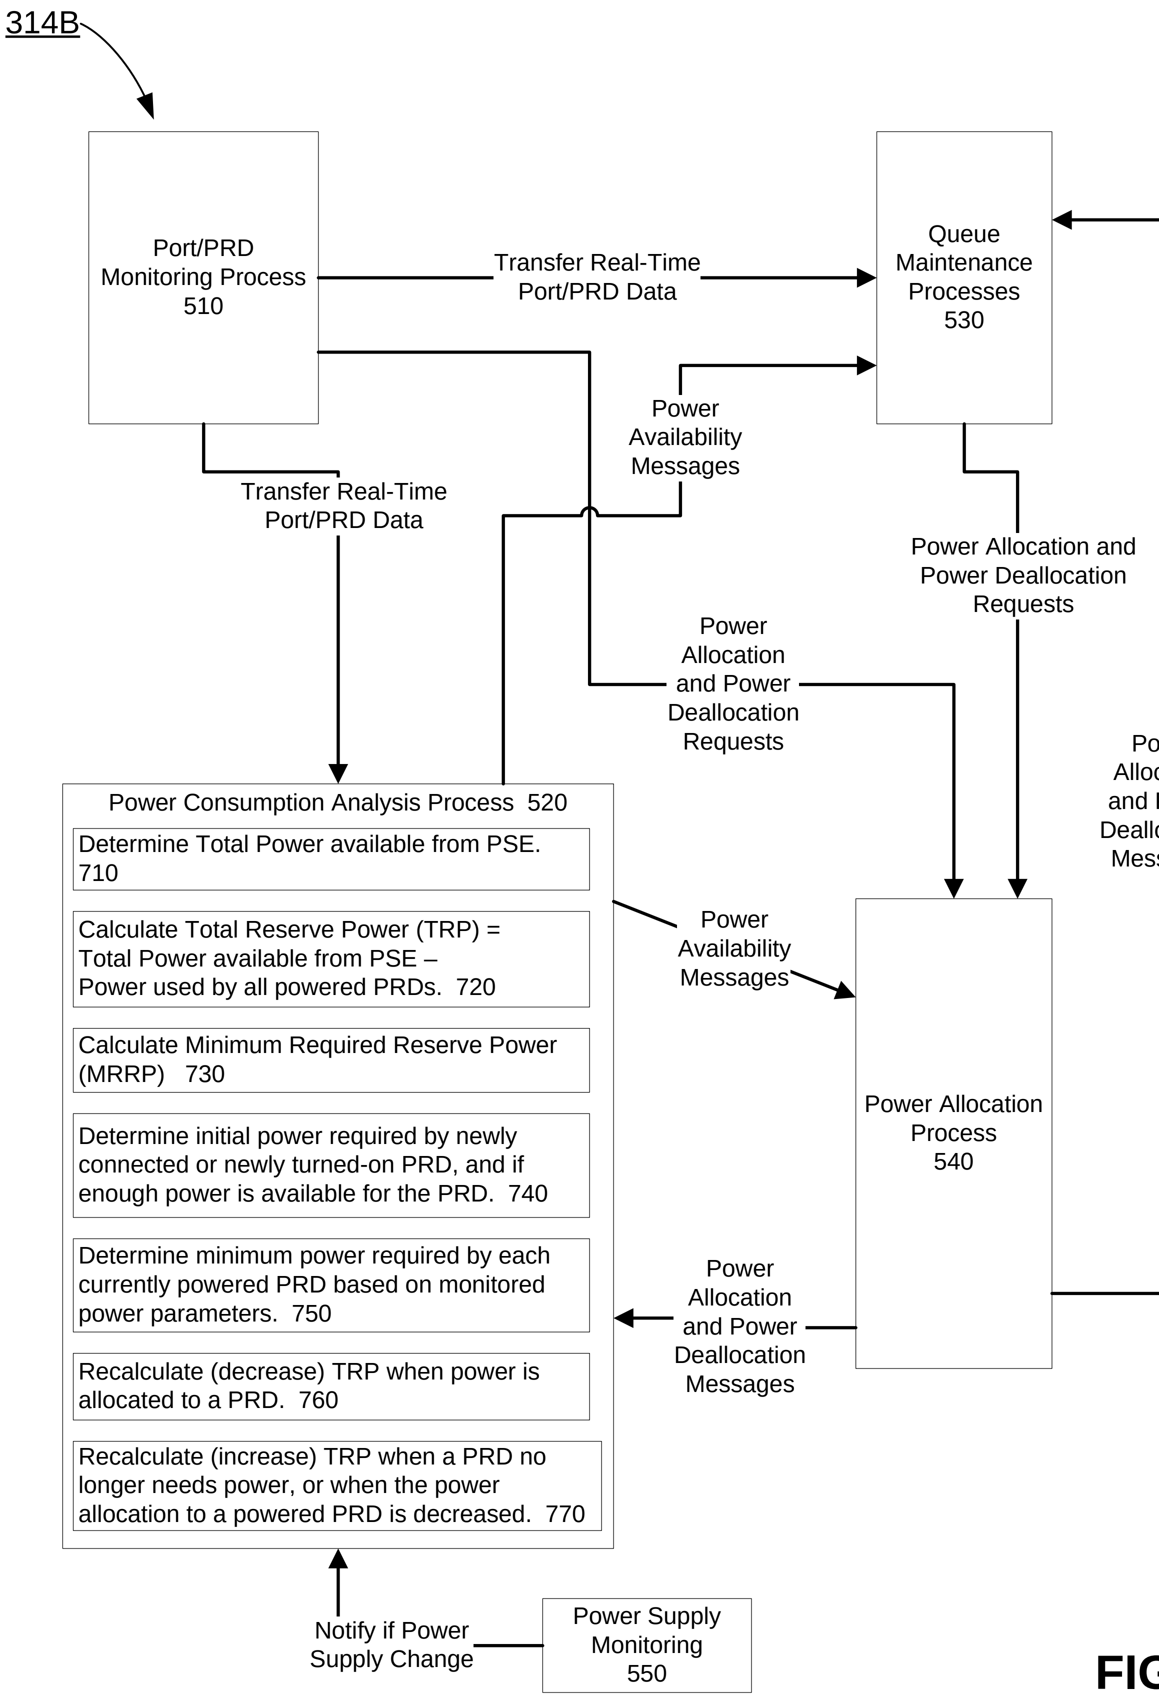

FIG. 7

[0159] FIG. 7 is still another view of method 314B. Many features of FIG. 7 previously illustrated in FIG. 5 are repeated here and will not be described again in detail.

[0160] FIG. 7 shows in greater detail the steps involved in power consumption analysis process 520. In step 710, power consumption analysis process 520 determines the total power available from the PSE. This is based on information supplied by power supply monitoring process 550 via messages 552.

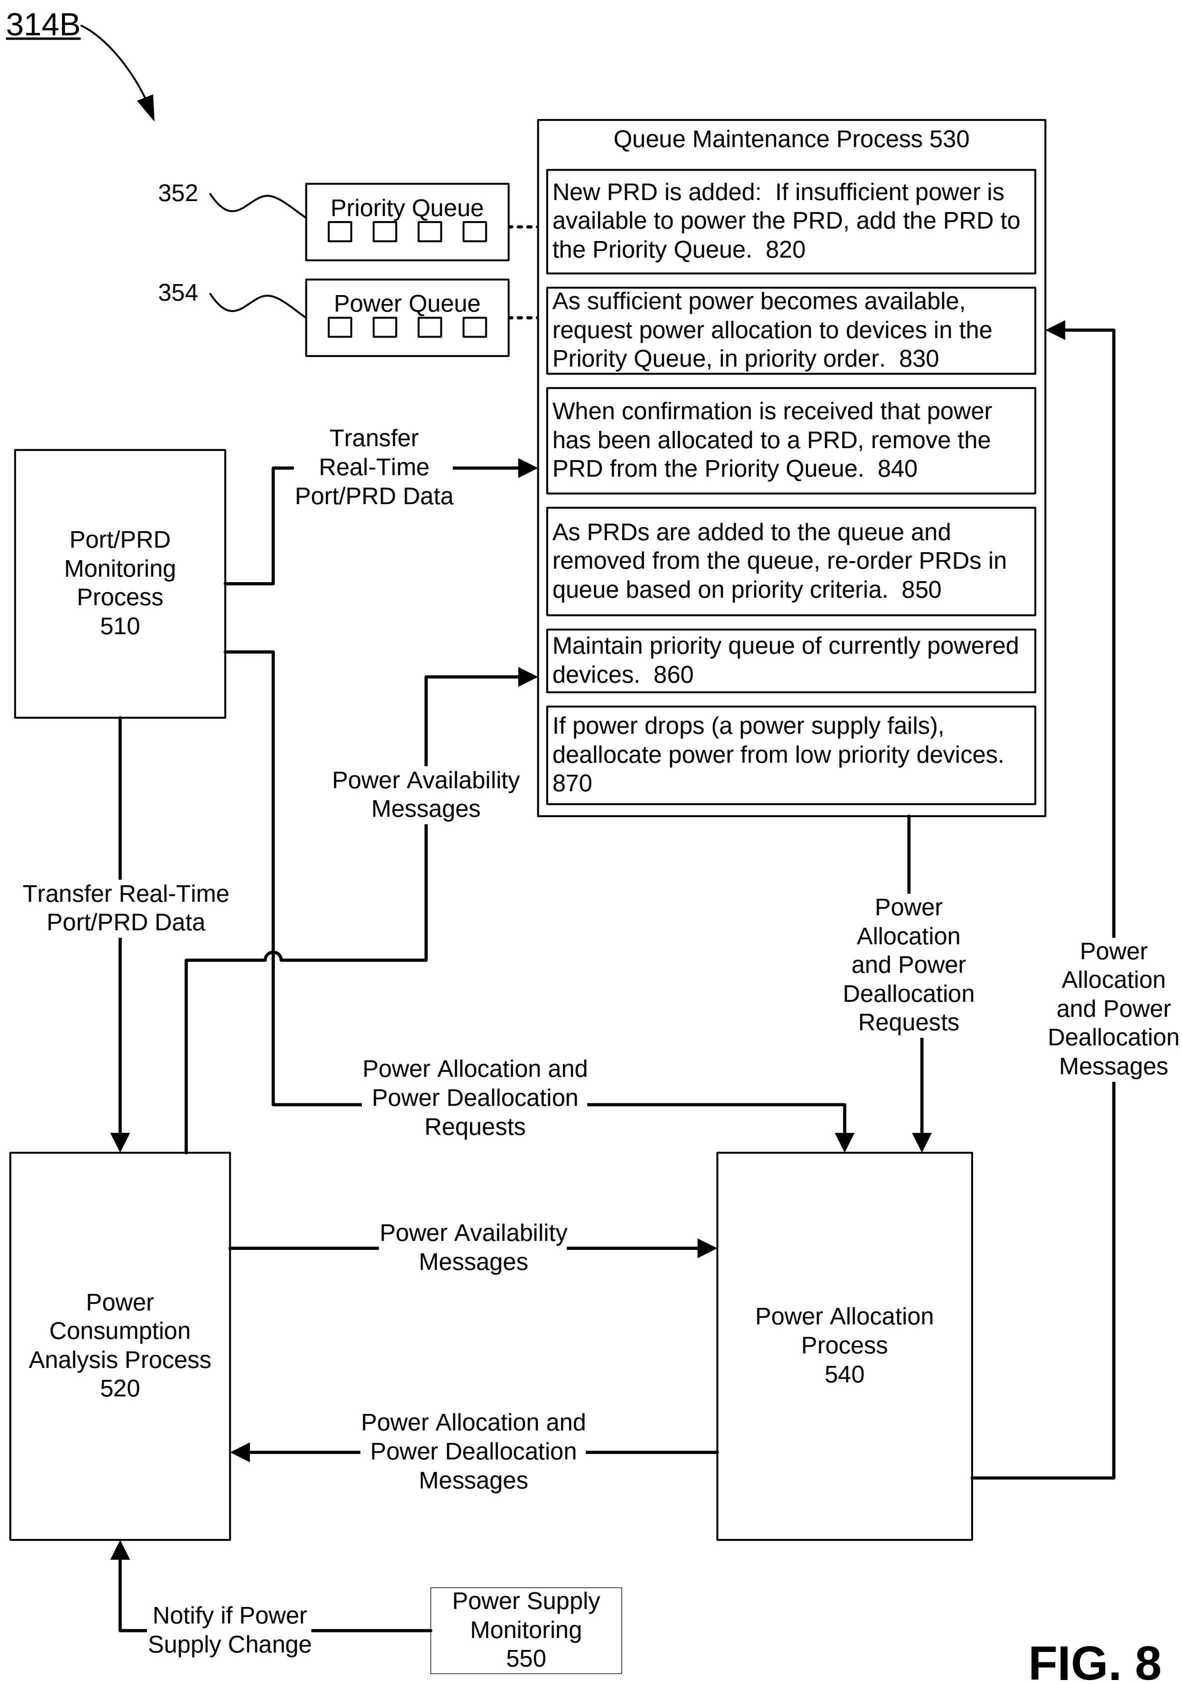

FIG. 8

[0176] FIG. 8 is still another view of method 314B. Many features of FIG. 8 previously illustrated in FIG. 5 are repeated here and will not be described again in detail.

[0177] FIG. 8 shows in greater detail the steps involved in queue maintenance process 530, which is depicted here with its associated priority queue 352 and power queue 354.

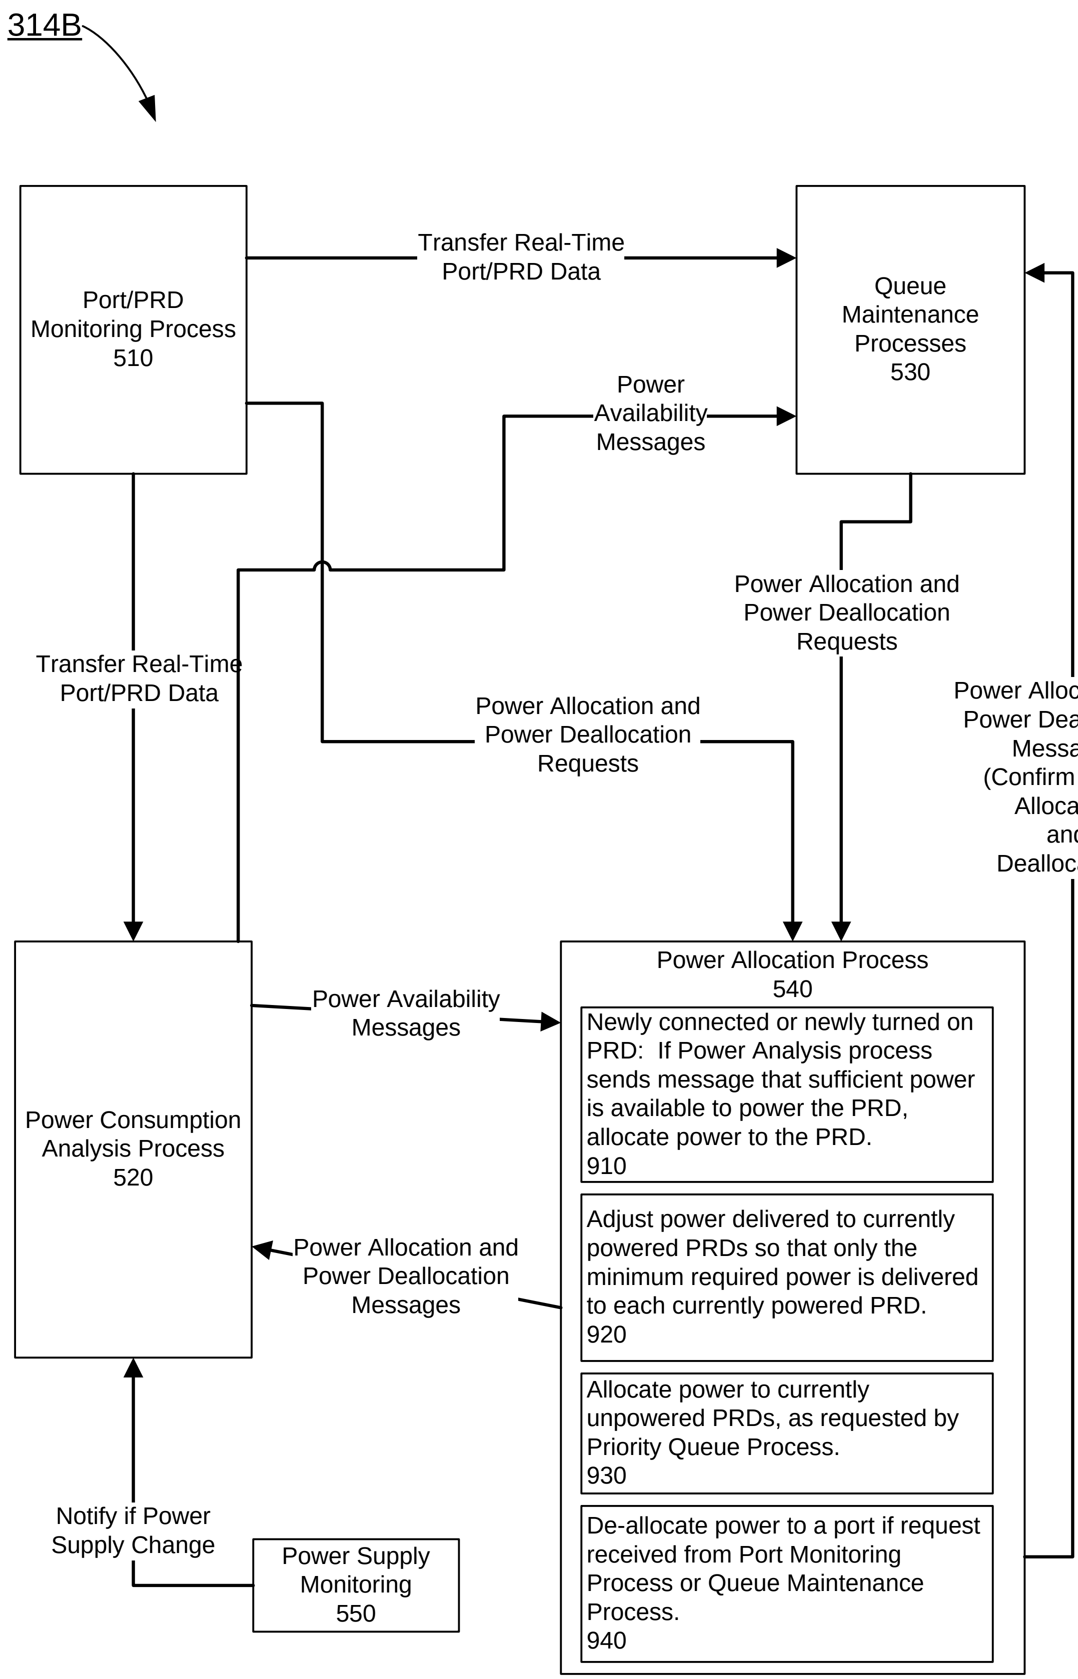

FIG. 9

[0196] FIG. 9 is still another view of method 314B. Many features of FIG. 9 previously illustrated in FIG. 5 are repeated here and will not be described again in detail.

[0197] FIG. 9 shows in greater detail the steps involved in power allocation process 540.

SYSTEM AND METHOD OF DYNAMIC POWER MANAGEMENT

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This patent application claims the benefit of U.S. Provisional Patent Application No. 60/858,937, filed Nov. 15, 2006, entitled "System And Method Of Dynamic Power Management," which is incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

[0002] 1. Field of the Invention

[0003] The present invention pertains to the field of providing power from power source equipment (PSE) to one or more devices which are part of a communications system or communications network, and to managing the power which is supplied to the networked devices.

[0004] 2. Background Art

[0005] Network devices which require power are conventionally referred to in the art as "powered devices", or "PDs" for short. In this document, for reasons explained further below, such devices will be referred to as "power requiring devices", or "PRDs", wherein the terms "power requiring device", "PRD", and the plurals thereof are entirely synonymous with the terms "powered device", "PD", and the plurals thereof, as conventionally employed in the art.

[0006] Conventionally, in a communications system, such as a network environment where cabling is employed to provide power to PRDs, at least two separate cables are connected to each PRD. One PRD cable conveys data to and from the (data communications equipment (DCE), while a second cable supplies each PRD with power from a power supply which is typically separate from the DCE. The disadvantages of this approach have long been obvious. Using two cables increases the physical bulk of cabling which must be wired through the physical environment, demanding more space and making it more difficult to track which cables are serving which functions. If power is supplied to a PRD via a local power outlet, for example, a standard wall outlet, then it may be necessary to locate a separate uninterruptible power supply (UPS) near many or all of the PRDs. As a result, multiple UPS may be required. Some PRDs have required transformers co-located with the PRD, further increasing bulk and complicating deployment efforts. Overall, dual cable solutions typically cost more than a single cable solution and tend to be unwieldy to deploy and maintain.

[0007] In addition, a key limitation of dual cable solutions is that they inherently tend to exclude the transmission of data which is related to the power consumption of a PRD. For example, if a PRD is plugged directly into a wall socket, or connected to a wall socket via a dedicated UPS or transformer, there typically does not exist any means to monitor the power delivered to the PRD, nor to regulate the power delivered to the PRD. In turn, this makes it difficult to monitor and control overall network power consumption.

[0008] A single cable solution, wherein the single cable carries both data and power, has long been known in certain fields. For example, the cabling employed by the plain old telephone system (POTS) network carries both power and data. Other technologies, such as USB communications systems and IEEE **1392** ("Firewire") systems, typically are capable of supporting both data and power transmission on a single cable. The advantages of single cable solutions are obvious: Lower cost, half as many wires to be tracked and untangled, and, if all the cables extend from a single DCE, then a single UPS can provide power support to all the network devices via the DCE. Furthermore, single cable solutions are inherently more friendly to technology which monitors and regulates power consumption by the PRDs, since the power cable is already designed to support data transmission as well.

[0009] A problem has existed, however, in adapting certain widely used legacy technologies to take advantage of single cable power and data solutions. In particular, Ethernet communications systems, which are widely deployed for wired computer networks worldwide, have long relied on one cable (the Ethernet cable) to carry data to PRDs and another separate cable to supply power to the PRDs, with all the disadvantages already noted above.

[0010] Recently a solution has emerged to enable Ethernet networks to take advantage of single cable data and power solutions. Specifically, the IEEE 802.3af protocol defines standards wherein power can be carried to PRDs from a DCE, such as an Ethernet switch, over the same industry-standard Ethernet cabling and Ethernet connectors (i.e., the RJ-45 plug and jack) that formerly was used to carry only data. As a result, where it was formerly necessary to employ a separate source of power to deliver power to PRDs, the power source equipment (PSE) can now be integrated into the DCE itself. Newly designed PRDs, such as wireless access points (WAPs) for 802.11 and Bluetooth devices, web cameras, IP telephones, security access devices, point-of-service (POS) terminals, and similar technologies are designed to accept power directly over the Ethernet cabling, dramatically simplifying deployment of network devices.

[0011] The IEEE 802.3af standard is also referred to as the "Power Over Ethernet" standard, or "PoE", and the terms "802.3af" and "PoE" are used interchangeably in this document, along with such terms as "802.3af-compliant", "PoE-compliant", and related terminology.

[0012] An advantage of PoE is that a PoE-compliant PSE, or PoE PSE for short, can detect the presence of legacy Ethernet devices which are not designed to accept power over the Ethernet cable, and so not deliver power to such devices. This makes it possible to support both PoE-compliant PRDs and legacy PRDs in the same network. Additionally, legacy Ethernet switches, which are not PoE-compliant, can be used with the newer PoE-compliant PRDs. This is done by inserting, between the legacy Ethernet switch and the PRDs, a so-called "mid-span" PSE, also referred to as a PoE mid-span hub. The PoE mid-span hub accepts the Ethernet cabling from the Ethernet switch, and loads the power onto Ethernet cables which extend from the PoE hub to the PRDs. When data is received from the PRDs, the mid-span PoE hub passes the signals on to the Ethernet switch.

[0013] A further advantage of the IEEE 802.3af standard is that PoE-compliant PRDs can have an embedded management information base (MIB), which contains data on the power requirements of the device as determined by the device manufacturer. Specifically, 802.3af-compliant PRDs can be classified at power levels designated as '1', '2', or '3', which respectively indicate increasing power requirements. (A default '0' value encompasses all three power levels. A fifth designation, '4', is currently reserved for future use.)

The PoE PSE can read the power classification from a PoE PRD via the Ethernet connection. This enables the PoE PSE to allocate only the designated amount of power to a PRD, which conserves power for other attached devices and contributes to the general goal of efficient power utilization. Yet another advantage of combining data signaling and power delivery over a single cable is that PRDs can be shut down remotely, via control signals sent from the DCE, without the need for a power switch or reset button.

[0014] In summary, PoE technology, and similar protocols which may be employed in the context of other kinds of communications links such as USB and Firewire, enable power to be carried over the same cable that is used for data transmission. These technologies further enable intelligent management by DCEs of the power consumption by PRDs.

[0015] In spite of the improvements in network power allocation and power management that have come with 802.3af and similar protocols, shortcomings remain. One problem is that, in real-world operation, PRDs often use significantly less power than the maximum power they might request. For example, an 802.3af-compliant PRD might classify itself in Class 2, meaning that the DCE should be able to deliver 7.0 watts to the port to which the PRD is connected, and therefore that at least 7.0 watts should be held in reserve by the DCE to support the Class 2 PRD. However, in actual use, the PRD might never use anywhere near this much power, or possibly the PRD might draw near the maximum power at infrequent intervals, while using dramatically less power most of the time. This might be the case, for example, with a wireless access point that has a light load of wireless devices associated with it.

[0016] In addition, it is not always a requirement of the standards that a device provide any power classification. For example, PoE devices are not required to provide a power classification. As a result, when a PoE-compliant PRD is connected to the DCE, and the PRD does not report any power classification, by default the DCE may allocate or reserve the maximum allowed power under the 802.3af protocol, namely 15.4 watts, to the port to which the PRD is connected. This will be the case even if the PRD never actually consumes anywhere near 15.4 watts.

[0017] As a result of these factors, the actual power allocated to PRDs, or held in reserve by the DCE to support the PRDs, may be well in excess of the amount of power that these devices actually require to support their operation. If many devices are connected to the DCE, a further result may be that insufficient power is available to power all the devices, leaving some devices effectively non-operational. It is desirable, then, that a means be found to ensure that the minimum possible power is allocated to each PRD, while still ensuring full functionality of each PRD. In turn, this will ensure that the number of PRDs which may actually be allocated power tends to be maximized, while still minimizing overall power consumption.

[0018] A further short-coming of existing power technologies is that they provide no way to prioritize attached network devices. That is to say, if insufficient power is available to support the operations of all PRDs attached to the DCE, it may be desirable that some specific PRDs have a higher priority for power allocation than other PRDs. The exact choice of which PRDs should have higher priority will naturally depend on many factors specific to the nature of the network, the usage being made of the network devices, and the specific requirements of a company or other enterprise implementing the network.

[0019] Therefore, it is desirable that means be provided so that network engineers can indicate a priority among network devices, such that if insufficient power is available to power all PRDs connected to the DCE, then those PRDs assigned the highest priority are preferentially allocated power ahead of lower priority PRDs whenever possible. It is further desirable that network engineers be provided the maximum possible flexibility in determining the criteria which are used to assign power allocation priority levels to PRDs.

[0020] A further shortcoming of present technologies is that, in the event of a sudden decrease in available power, no method is provided to deallocate power from the PRDs. It is desirable that, in the event of an unexpected decrease in available power, a means exist to deallocate power from some of the PRDs, preferably in an order which reflects some kind of priority among the attached PRDs, where that priority may or may not be the same priority used to assign power to the PRDs to begin with.

[0021] Given the foregoing, what is needed is a system and method of improved, dynamic power management and allocation for power requiring devices which are part of a communications system or communications network.

[0022] In particular, what is needed is a system and method of dynamic power management and allocation for determining the actual power consumption of power requiring devices which are part of the communications system, and adjusting the power allocation so that the minimum necessary power is allocated to power requiring devices which are part of the communications system. What is further needed is a system and method of dynamic power management and allocation for determining a priority among power requiring devices which are part of the communications system, and for ensuring that whenever possible power requiring devices of a higher prior are preferentially allocated power ahead of power requiring devices of a lower priority.

BRIEF SUMMARY OF THE INVENTION

[0023] The present invention meets the above-identified needs by providing a system and method for dynamic management in a communications system or communications network. In an exemplary embodiment, the present invention may be employed as part of a data communications equipment (DCE) device, such as a network switch, wherein the DCE is also the PSE and wherein multiple network devices are powered via the PSE. The PSE may have a plurality of communications ports and a plurality of PRDs which transmit and receive data over the ports, and receive power via the ports. The PSE monitors the power consumed at each port by those PRDs which are currently powered.

[0024] Based on the power monitoring, the PSE dynamically determines the minimum power which should be allocated to each PRD, and so dynamically determines the amount of reserve power available. The PSE also maintains a listing, possibly in the form of a PRD queue or queues, which indicates the priority of PRDs. When additional power is available to be allocated, the PSE preferentially allocates power to a PRD or PRDs which have higher priority. If a decrease in available power mandates that some currently powered PRDs no longer receive power, the PSE preferentially ensures that PRDs with a higher priority continue to receive power.

[0025] The method of the present invention thus provides dynamic power management, wherein the power allocated to each individual network device may be minimized, and wherein as a result the total number of network devices that can be supported with the available power may be maximized. One result may be a more efficient utilization of the available power.

BRIEF DESCRIPTION OF THE DRAWINGS/FIGURES

[0026] The features and advantages of the present invention will become more apparent from the detailed description set forth below when taken in conjunction with the drawings in which like reference numbers indicate identical or functionally similar elements. Additionally, the left-most digit of a reference number identifies the drawing in which the reference number first appears.

[0027] FIG. **1**A depicts an exemplary communications system deployment, and specifically a Power-Over-Ethernet (PoE) system deployment, in which the present invention might be employed, wherein the PoE system is deployed in an end-span configuration.

[0028] FIG. **1**B depicts an exemplary communications system deployment, and specifically a Power-Over-Ethernet (PoE) system deployment, in which the present invention might be employed, wherein the PoE system is deployed in a mid-span configuration.

[0029] FIG. **2** contains block diagrams showing part of the internal structure of a representative communications system, specifically a PoE switch, which may be used in one embodiment of the present invention, and also showing part of the internal structure of a representative power requiring device. The figure also shows the connection between the PoE switch and the power requiring device.

[0030] FIG. **3**A illustrates a method according to one embodiment of the present invention, whereby a representative communications system and associated power requiring devices are initially powered on, and wherein the current method of dynamic power management of the present invention is initiated.

[0031] FIG. **3**B is a representative depiction of several queues, specifically a priority queue and a power queue, and wherein the queues are employed in one embodiment of the present invention to store priority information on PRDs.

[0032] FIGS. **4**A and **4**B together present a flowchart illustrating one embodiment of a method of ongoing dynamic power management according to the present invention. The single flowchart has been split between two figures for ease of reading only, and embodies a unified method.

[0033] FIG. **5** presents another embodiment of a method of ongoing dynamic power management according to the present invention, wherein multiple monitoring, analysis, and control processes occur in parallel or substantially in parallel, and wherein these processes communicate with each other through a series of messages.

[0034] FIG. **6** is another view of the method presented in FIG. **5**. The figure shows in detail the steps involved in a port/PRD monitoring process.

[0035] FIG. **7** is another view of the method presented in FIG. **5**. The figure shows in detail the steps involved in a power consumption analysis process.

[0036] FIG. **8** is another view of the method presented in FIG. **5**. The figure shows in detail the steps involved in a queue maintenance process, wherein the queues are those illustrated in FIG. **3**B.

[0037] FIG. **9** is another view of the method presented in FIG. **5**. The figure shows in detail the steps involved in a power allocation process.

DETAILED DESCRIPTION OF THE INVENTION

I. Introduction

[0038] The present invention is directed to a system and method for dynamic power management and allocation in a communications system or communications network. In an exemplary embodiment, power is provided from power source equipment (PSE) to one or more power requiring devices (PRDs) via a means of power conveyance, such as network cabling. In one embodiment of the present invention, the PSE is contained in the same hardware unit or device which also serves as the data communications equipment (DCE), that is, the hardware unit or device which transmits data to and/or accepts data from the PRDs. Further, data sent back and forth via the means of power conveyance may include, in addition to any network communications proper, data which is related to power management of the PRD or PRDS.

[0039] It should be understood that this communications system configuration is exemplary only, and other communications system or network configurations are possible and are consistent with the functioning of the present invention. For example, in an alternative embodiment of the present invention, the DCE and the PSE could be in physically separate units or cases, and could send data and power respectively to the PRDs via separate cables or other means. The dynamic power management methods of the present invention might then entail some further connection between the DCE and the PSE, such that power management related data and power management control information could be passed back and forth between the DCE and the PSE.

[0040] Persons skilled in the relevant art(s) will appreciate that while the methods described in the following discussion will be applied in the context of devices which use electrical power and electrical signaling, nothing in the present invention limits the invention to the management of power supplied as electrical power via conventional cables which typically use insulated or shielded metal wires, or other metallic pathways or metallic means of power conductance.

[0041] Rather, the method of the present invention may be applied to other technologies for both power conveyance and signal communication. For example, optical communications are increasingly used as a means to convey data, and in the near future it may be possible to power a communications system using optical or other purely electromagnetic-wave-based power conveyance means including, for example and without limitation, laser light, infrared light, microwaves, or x-rays. The method of the present invention can be employed in these contexts as well to provide dynamic power management, wherein the power allocated to each individual network PRD is substantially minimized, and wherein as a result the total number of network PRDs that can be supported with the available power may be substantially maximized.

[0195] Powered PRDs **226** may be preferably deallocated power in such a manner that those devices of lowest priority are deallocated first while still ensuring that sufficient power is available to power those devices of highest priority. Once a decision has been made by queue maintenance process **530** to deallocate power from a device, messages **582** are sent to power allocation process **540** requesting the deallocation from designated PRDs **108**. Power allocation process **540** responds with messages **592** indicating whether or not power has been successfully deallocated. Queue maintenance process **530** responds by adding and removing PRDs **108** from priority queue **352** and power queue **354** as appropriate, and if necessary reordering devices in those queues as appropriate. Messages are also sent from power allocation process **540** to power consumption analysis process **520** indicating that power has been deallocated.

[0196] FIG. **9** is still another view of method **314**B. Many features of FIG. **9** previously illustrated in FIG. **5** are repeated here and will not be described again in detail.

[0197] FIG. **9** shows in greater detail the steps involved in power allocation process **540**.

[0198] In step **910**, in one embodiment of the present invention, a message **580** may be received by power allocation process **540** from port/PRD monitoring process **510** indicating that a new PRD **108** has been connected to PSE **104**, or that a previously connected PRD **108** that was powered off is now powered on. If a message or messages **572** from power consumption analysis process **520** indicate that sufficient power is available to power PRD **108**, then power allocation process **540** allocates power to the PRD **108**. In an alternative embodiment of the present invention, power allocation process **540** itself makes the determination as to whether there is sufficient power available to power the PRD **108**, using methods analogous to those described above for power consumption analysis process **520**.

[0199] If sufficient power is available, power allocation process **540** allocates the power to PRD **108**. The detailed steps involved in power allocation, including any necessary, associated data communication with PRD **108**, may conform to protocols appropriate to the power delivery technology employed by PSE **104**, such as the power allocation protocols indicated by IEEE standard 802.3af. Having allocated the power, power allocation process **540** sends a message **592** to queue maintenance process **530** and also sends a message **590** to power consumption analysis process **520**.

[0200] In step **920**, power allocation process **540** adjusts the power allocated to currently powered PRDs **224** so that only the minimum required power is delivered to each currently powered PRD. Power allocation process **540** is able to do this based on power availability messages and power analysis messages provided by power consumption analysis process **520**, as already described above.

[0201] In step **930**, power allocation process **540** receives messages **582** from queue maintenance process **530** requesting that power be allocated to currently unpowered PRDs **226**. In one embodiment of the present invention, power allocation process **930** may use data contained in a message or messages **572** from power consumption analysis process **520** to first confirm that sufficient power is available for the allocation. In an alternative embodiment of the present invention, power allocation process **540** itself makes the determination as to whether there is sufficient power available to power the PRD **108**. Upon having made the confirmation, power allocation process **540** allocates power to the currently unpowered PRD **226** wherein the PRD **108** becomes a powered PRD **224**.

[0202] Power allocation process **540** then sends a message **590** to power consumption analysis process **520** and also sends a message **592** to queue maintenance process **530**, wherein both messages indicate that power has been allocated.

[0203] Finally, in one embodiment of the present invention, in step **940** power allocation process **540** monitors for requests **580** from port/PRD monitoring process **510** or queue maintenance process **530** to deallocate power from a port. As discussed above, these requests may come for a variety of reasons including, for example and without limitation, a device being disconnected from a port, a fault condition being detected on a port, a decrease in overall power availability, or for other reasons.

[0204] Having received the message to deallocate power to a port, power allocation process **540** deallocates power from the port. The detailed steps involved in power deallocation, including any necessary, associated data communication with PRD **108**, will conform to protocols appropriate to the power delivery technology employed by PSE **104**, such as the power allocation protocols indicate by IEEE standard 802.3af. Having deallocated the power, power allocation process **540** sends a message **592** to queue maintenance process **530** and also sends a message **590** to power consumption analysis process **520** indicating that power was deallocated from the PRD **108**.

V. Conclusion

[0205] As will be appreciated by persons skilled in the relevant art(s), the processes described here represent only one possible embodiment of the present invention. Many of the steps described could, in alternative embodiments of the present invention, be allocated differently among the different processes described. In addition, additional processes, or a different organization of the various steps of the method into processes of different structure, could still implement the overall effect and intent of the method.

[0206] Finally, it will be understood that the actual process of allocating power or deallocating power is accomplished through at least one of the hardware, software and firmware embodied in the PSE **104**, and in particular in some embodiments of the present invention, in the PSEPM **204** which handles the power management protocols mandated by the 802.3af standard. Therefore, the method of the present invention in an exemplary embodiment is to have the algorithms or processes of the invention send messages to PSEPM **204** for purposes of controlling the actual power at the ports **216** and/or for obtaining power information and related data from the ports **216**. In an alternative embodiment, the protocols and methods to ensure compliance with 802.3af may be handled through controller **206** which controls ports **216**. Controller **206**, in turn, is controlled by the algorithms or processes of the present invention.

What is claimed is:

1. In a communications system having power source equipment (PSE) and a power requiring device (PRD), the PSE providing power to the PRD over the communications system, a method of power management comprising:

(a) monitoring a power consumption of the PRD; and

(b) dynamically adjusting the power allocated to the PRD based on the monitored power consumption of the PRD.

2. The method of claim **1**, further comprising allocating power to the PRD when the communications system is first powered up based on at least one of a predefined power allocation, a PRD power classification, or a maximum power which can be allocated to a communications port of the communications system to which the PRD is connected.

3. The method of claim **1**, further comprising

(c) checking for a PRD that previously required power but no longer requires power, and no longer allocating power to the PRD.

4. The method of claim **3**, wherein step (c) comprises at least one of:

(i) determining that a PRD that was connected to the communications system is no longer connected to the communications system;

(ii) determining that a PRD that is connected to the communications system and was powered on is now powered off; and

(iii) checking for a fault condition in a communications port of the communications system to which the PRD is connected.

5. The method of claim **4**, wherein step (c)(iii) comprises monitoring the communications port for fault conditions including at least one of short circuits, in-rush, and thermal shutdown (TSD).

6. The method of claim **1**, further comprising:

(d) maintaining a priority queue of power requiring devices (PRDs) which are connected to the communications system, wherein the PRDs in the priority queue are not currently powered;

(e) determining a total reserve power currently available from the PSE;

(f) allocating power to an unpowered PRD in the priority queue while the total reserve power is greater than or equal to a minimum required reserve power, wherein the unpowered PRD becomes a powered PRD; and

(g) removing the powered PRD from the priority queue.

7. The method of claim **6**, wherein step (d) comprises listing the PRDs in the priority queue in order of a predetermined PRD priority;

wherein the PRD priority is determined by a set of priority elements, the set of priority elements including at least one of:

a PRD power classification;

a communications port number of a communications port of the communications system to which the PRD is connected;

a preassigned priority value of the communications port to which the PRD is connected; and

a time order in which the PRD is connected to the communications system.

8. The method of claim **6**, wherein step (e) comprises calculating the total reserve power as a maximum power supplied by the PSE minus a sum of the power currently being supplied to all PRDs or calculating the total reserve power as the maximum power supplied by the PSE minus an average of the power supplied over a period of time to all PRDs which are currently powered.

9. The method of claim **6**, wherein step (f) comprises setting the minimum required reserve power equal to a maximum power that can be used by at least one of a highest powered PRD connected to the PSE or a highest powered communications port which is part of the communications system.

10. The method of claim **6**, further comprising:

(h) maintaining a power queue of PRDs which are connected to the communications system, wherein the PRDs in the power queue are currently powered and wherein the PRDs in the power queue are ranked in an order of priority;

(i) determining that there is a decrease in the maximum power supplied by the PSE;

(j) deallocating power from PRDs in the power queue, wherein a PRD which ranks lowest in priority in the power queue is deallocated first, and wherein PRDs from which power is deallocated are deallocated PRDs, and wherein power is deallocated from PRDs in the power queue until the maximum power supplied by the PSE is at least sufficient to supply power to the powered PRDs; and

(k) assigning to the priority queue the deallocated PRDs.

11. The method of claim **10**, wherein PRDs in the power queue are ranked in priority according to a PRD importance, wherein the PRD importance is determined by a set of priority elements, the set of priority elements including at least one of:

a power classification of each powered PRD;

a preassigned priority value of the communications port to which each powered PRD is connected;

a communications port number of the communications port to which each powered PRD is connected;

a time when the powered PRD was connected to the communications port; and

a user-designated PRD importance.

12. The method of claim **1**, wherein step (a) comprises monitoring at least one of a port current, a port voltage, and a port temperature at a communications port where the PRD is connected.

13. The method of claim **1**, wherein step (b) comprises dynamically determining a minimum necessary power required by the PRD, and allocating the minimum necessary power to the PRD.

14. The method of claim **13**, wherein the minimum necessary power is determined by applying a law of averages based on the power consumption of the PRD over a period of time.

15. The method of claim **1**, wherein the communications system is a Power-over-Ethernet system; and

wherein power is delivered from the PSE to the PRD via a communications port of the Power-over-Ethernet system.

16. In a communications system having power source equipment (PSE), a plurality of communications ports, and a plurality of power requiring devices (PRDs), the PSE providing power to each power requiring device (PRD) over a corresponding one of the communications ports,

wherein each PRD which is currently powered via a corresponding communications port is a powered PRD,

wherein each communications port with a powered PRD connected to it has a communications port power consumption, and

wherein the PSE supplies a maximum power,

a method of power management comprising:

(a) allocating power to each PRD when the communications system is first powered up;

(b) monitoring the communications port power consumption for each port which currently has a powered PRD connected;

(c) dynamically adjusting the power allocated to each powered PRD based on the monitored communications port power for the communications port to which the PRD is connected;

(d) maintaining a priority queue of one or more PRDs which are connected to the communication ports in the communications system but which are not currently powered;

(e) determining a total reserve power currently available from the PSE; and

(f) allocating power to one or more currently unpowered PRDs in the priority queue while the total reserve power is greater than a minimum required reserve power, wherein the one or more unpowered PRDs become powered PRDs and are removed from the priority queue.

17. The method of claim **16**, further comprising checking for a communications port that previously required power but no longer requires power, and no longer allocating power to the communications port, wherein checking for the communications port comprises at least one of:

(g) determining that the PRD that was connected to the communications port is no longer connected to the communications port;

(h) determining that the PRD that is connected to the communications port and was powered on is now powered off; and

(i) checking for a fault condition in the communications port, wherein the fault condition includes at least one of short circuits, in-rush, and thermal shutdown (TSD).

18. The method of claim **16**, further comprising:

(g) maintaining a power queue of PRDs which are connected to the communications system, wherein the PRDs in the power queue are currently powered and wherein the PRDs in the power queue are ranked in priority in an order of priority;

(h) determining that there is a decrease in the maximum power supplied by the PSE;

(i) deallocating power from a PRD in the power queue, wherein the PRD which ranks lowest in the power queue is deallocated first, and wherein PRDs from which power is deallocated are deallocated PRDs, and wherein power is deallocated from PRDs in the power queue until the maximum power supplied by the PSE is at least sufficient to supply power to the powered PRDs; and

(j) assigning to the priority queue the deallocated PRD.

19. The method of claim **18**, wherein PRDs in the power queue are ranked in priority according to a PRD importance, wherein the PRD importance is determined by a set of priority elements, the set of priority elements including at least one of:

a power classification of each powered PRD;

a preassigned priority value of the communications port to which each powered PRD is connected;

a communications port number of the communications port to which each powered PRD is connected;

a time when the powered PRD was connected to the communications port; and

a user-designated PRD importance.

20. The method of claim **16**, wherein step (a) comprises allocating power to the PRDs based on at least one of a predefined power allocation, a PRD power classification, or a maximum power which can be allocated to each communications port to which each PRD is connected.

21. The method of claim **16**, wherein step (b) comprises monitoring at least one of a port current, a port voltage, and a port temperature at the port where the PRD is connected and powered.

22. The method of claim **16**, wherein step (c) comprises dynamically determining a minimum necessary power required by each powered PRD, and allocating the minimum necessary power to each powered PRD.

23. The method of claim **16**, wherein step (d) comprises listing the unpowered PRDs in the priority queue in order of a descending PRD priority, wherein the PRD priority is determined by a set of priority elements, the set of priority elements including at least one of:

a power classification of each PRD;

a preassigned priority value of the communications port to which each unpowered PRD is connected; and

a communications port number of the communications port to which each unpowered PRD is connected.

24. The method of claim **23**, wherein a set of unpowered PRDs having the same priority based on the set of priority elements are assigned to the priority queue based on a time order in which each unpowered PRD in the set of unpowered PRDs was connected to its corresponding communications port.

25. The method of claim **16**, wherein step (e) comprises calculating the total reserve power as a maximum power supplied by the PSE minus a sum of the power currently being supplied to all PRDs or calculating the total reserve power as the maximum power supplied by the PSE minus an average of the power supplied to all PRDs over a period of time.

* * * * *