

(19) **DANMARK**

(10) **DK/EP 2084591 T3**



(12) **Oversættelse af
europæisk patentskrift**

Patent- og
Varemærkestyrelsen

-
- (51) Int.Cl.: **G 06 F 8/65 (2018.01)** **G 06 F 1/26 (2006.01)**
- (45) Oversættelsen bekendtgjort den: **2019-12-16**
- (80) Dato for Den Europæiske Patentmyndigheds bekendtgørelse om meddelelse af patentet: **2019-09-11**
- (86) Europæisk ansøgning nr.: **06816708.9**
- (86) Europæisk indleveringsdag: **2006-10-10**
- (87) Den europæiske ansøgnings publiceringsdag: **2009-08-05**
- (86) International ansøgning nr.: **US2006039715**
- (87) Internationalt publikationsnr.: **WO2008045044**
- (84) Designerede stater: **AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IS IT LI LT LU LV MC NL PL PT RO SE SI SK TR**
- (73) Patenthaver: **Google LLC, 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA**
- (72) Opfinder: **KRIEGER, Ken, P.O. Box 11694, Jackson, Wyoming 83002, USA**
BORCHERS, Albert, 60 Nichols Dr., Santa Cruz, California 95060, USA
- (74) Fuldmægtig i Danmark: **AWA Denmark A/S, Strandgade 56, 1401 København K, Danmark**
- (54) Benævnelse: **OPDATERING AF EN STRØMFORSYNINGSMIKROSTYREENHED**
- (56) Fremdragne publikationer:
US-A- 6 009 177
US-A1- 2005 088 201
US-A1- 2006 020 832
US-B1- 6 584 559

DESCRIPTION

TECHNICAL FIELD

[0001] The following disclosure relates to updating power supply microcontrollers, for example, updating firmware used by the microcontroller to control certain functions of a power supply.

BACKGROUND

[0002] To operate, computing systems need a source of power. Certain computing systems include power supplies, which provide power to internal components, such as motherboards, processors, and memory.

[0003] Power supplies may convert one form of electrical power to another desired form and voltage. In certain systems, this involves converting 120 or 240 V alternating current supplied by a utility company to a lower direct current voltage used for the internal components.

[0004] Power supplies can include switched-mode power supplies and linear regulator power supplies. Certain switched-mode power supplies include a transistor and transformer network that switches the load current rapidly on and off in order to stabilize the output voltage.

[0005] Certain linear regulator power supplies use a transformer to convert AC voltage to a lower AC voltage. For example, the linear regulator power supply can rectify the lower AC voltage with one or more diodes. Additionally, these linear regulator power supplies can include a capacitor to smooth out variations in the rectified signal. Next, a linear regulator can be used to decrease the voltage based on the requirements of the circuits powered by the power supply.

[0006] U.S. Pat. Publ. No. 2006/0020832 A1 discloses a circuit capable of updating power supply specifications of a microprocessor.

SUMMARY

[0007] The invention relates to a system as set out in claim 1 and a method for loading instructions into a programmable power supply as set out in claim 9. Preferred embodiments are described in the dependent claims.

[0008] The present specification relates to updating power supplies.

[0009] The systems and techniques described here may provide one or more of the following

advantages. First, cost savings and operational up-time may be increased by upgrading or restoring a power supply's functionality while maintaining the power supply's output. Second, a system may be upgraded conveniently through an internal bus. Third, a system may decrease the probability of executing corrupted or outdated firmware. Fourth, a system may update multiple power supplies at once or in a sequence. Fifth, operational problems related to a system can be repaired by updating firmware. Sixth, a power supply may be upgraded conveniently without removing from the power supply from a system.

[0010] The details of one or more embodiments are set forth in the accompanying drawings and the description below. Other features and advantages will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

[0011]

FIG. 1 is a schematic diagram illustrating an exemplary power distribution architecture 100 for a data center 105.

FIG. 2 shows an example method that a controller of a power supply can use as a start-up sequence for the controller.

FIG. 3 shows an example method performed by a controller of a power supply during a firmware update for the power supply.

FIG. 4 shows a more detailed method of a firmware update step of FIG. 3.

FIG. 5 shows an example of a method that a controller of a power supply can use to restrict access to a portion of memory during the firmware update process.

FIG. 6 is a schematic diagram illustrating an example of a power delivery system that includes a plurality of power supply management units (PSMU).

[0012] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

[0013] FIG. 1 is a schematic diagram illustrating an exemplary power distribution architecture 100 for a data center 105. The data center 105 includes a number of modular rack-mounted bases (which may also be referred to as trays) 110, each of which can include a power supply 115 integrated with a computer motherboard 120. The power supply 115 includes a controller 125 and a non-volatile memory (NVM) 130 (e.g., an EEPROM, flash memory, etc.).

Additionally, although depicted in the schematic as internal, the NVM can be either internal or external to the power supply 115 or the tray 110.

[0014] In this example, the NVM 130 includes two memory portions, where one portion stores application firmware 130a and the other stores a boot loader 130b. During normal operation, the controller 125 is configured to execute program instructions stored in the application firmware 130a to control the power supply 115 (e.g., performing fail-safe operations to protect the motherboard 120 when AC power fails and adjusting cooling mechanisms in the power supply). The controller 125 can execute program instructions stored in the boot loader 130b to maintain a power output to the motherboard 120 and to update the program instructions stored in the application firmware 130a.

[0015] In some implementations, the boot loader 130b includes instructions that when executed prepare the controller 125 for execution of the application firmware 130a. For example, the boot loader 130b instructions may reset variables within the controller and load addresses or points used by the controller to execute the application firmware 130a.

[0016] In the depicted example, the data center 105 receives AC power from an AC power source 135. For example, the AC power source 135 may deliver substantially symmetric three phase voltages suitable for powering substantially balanced three phase loads. In some implementations, the AC power may be received from an electric utility grid, a hydropower plant, a wind turbine system, a photovoltaic power system, or a combination thereof. In other implementations, the AC power source may be a tidal power plant that includes wave-to-electricity converters (e.g., Pelamis commercially available from Ocean Power Delivery Limited of Edinburgh, United Kingdom) to supply power to a water-based data center, for example.

[0017] The data center 105 includes a number of racks 140a, 140b, and 140c. Each of the racks 140a-c includes a number of the trays 110. The racks 140a-140c may be powered by the AC power that is delivered to the data center 105 from the electric utility 135. In some implementations, the racks 140a-c provide structures to facilitate the operations of the trays 110. For example, the racks 140a-c may include power cables to supply AC power to the trays 110 and tubes to supply liquid coolant to the trays 110. Additionally, the racks 140a-c may also include network cables that link the trays 110 in the racks 140a-c so that information can be exchanged between the trays 110 the data center 105.

[0018] The architecture 100 includes a network 145 and a remote computer 150. Each of the racks 140a-c is coupled to a network connection 155. The network connection 155 provides an information channel to the network 145, which may include, for example, a local area network, virtual private network, wide area network (e.g., the Internet), or a combination of such networks, which may be wired, fiber optic, and/or wireless. The computer 150 represents one of many possible devices that could communicate data directly or indirectly with one or more trays to access, store, process, and/or retrieve information using a processor 160 and associated memory 165 on the motherboard 120. In some implementations, additional processors (e.g., servers) may facilitate such communication. For example, the computer 150

may be included in a server, a desktop computer, a laptop computer, and/or a handheld processor-based device. One or more servers may pre- or post-process, supervise, route, and/or balance the flow of data associated with the communication.

[0019] By coupling the trays 110 to the network 145, the data center 105 can provide various services to network users. For example, the data center 105 can store various network information that a user of the computer 150 can search via the network 145. In another example, the user of the computer 150 can use the data center 105 to execute applications (e.g., word processing applications) remotely via the network 145. In another example, the data center 105 can be a security data center that store security data for the computer 150. In another example, the data center 105 can execute security critical operations, such as monetary transactions for clients of a bank.

[0020] In some implementations, the data center 105 can execute processing operations using one or a combination of more than one of the trays 110. As shown, each of the trays 110 provides processing functionality to the data center 105 using the motherboard 120. In various implementations, the motherboard 120 may include two, three, four, or any other practicable number of processors 160. Each of the processors 160 can execute instructions stored in the memory 165, for example. In some embodiments, the motherboard 120 may be replaced with tray of data storage devices (e.g., hard disc drives, flash memory, RAM, or any of these or other types of memory in combination). The memory 165 may include volatile and/or non-volatile memory that may be read and/or written to by the processor 160. The motherboard 120 may further include some or all of a central processor unit(s) (CPU), networking devices, memory (e.g., cache, non-volatile, flash), and/or disk drives, for example, along with various memories, chip sets, and associated support circuitry.

[0021] The motherboard 120 operates using the DC power supplied from the power supply 115. As shown, an AC input voltage signal is delivered to each of the trays 110 to be processed by the power supply 115. In some examples, the AC input voltage signal may be received from the AC power source 135 via the power cables in the racks 140a-c connecting each of the trays 110. The power supply 115 includes an AC-to-DC converter 170 that converts the AC input voltage signal to a regulated DC voltage. The converter 170 outputs the regulated DC voltage onto a DC bus 175. In some implementations, the AC-to-DC converter 170 may regulate the DC voltage to a static voltage set point. In some other implementations, the voltage set point may be dynamically determined. For example, the power supply 115 can determine the voltage set point based on, for example, a present temperature of the motherboard 120.

[0022] Various operations in the power supply 115 are controlled by the controller 125. For example, the controller 125 can execute application firmware 130a to perform power management operations, such as determining the voltage set point of the power output, monitoring signal quality of the AC input voltage signal, monitoring temperature in the power supply 115, and/or other power management functions. In some implementations, the controller 125 can execute the application firmware 130a to control a speed of a cooling fan in

the power supply 115 to regulate the air flow and cooling effects for the power supply 115. For example, when the temperature of the power supply 115 rises, the controller 125 can increase the fan speed to cool down the power supply 115 to avoid damage in the power supply 115 due to excessive heat.

[0023] In some implementations, the controller 125 executes instructions in the application firmware 130a that perform fail-safe operations to keep the motherboard 120 operational during an AC power failure. In some examples, the AC power failure may cause or potentially cause improper operation of the motherboard 120. As an illustrative example, when a fault condition in the AC input voltage signal is detected by the controller 125 (e.g., by detecting a root mean square (r.m.s.) of the input voltage signal to be less than a voltage threshold, or by detecting a frequency of the input voltage signal to be lower than a frequency threshold for a predetermined period of time), the controller 125 can control a switch 185 to connect a battery 190 to the DC bus 175 so that the motherboard 120 can continue to operate without substantial interruption. The battery 190 may continue to provide operating power to the circuits on the motherboard 115 until the battery 190 substantially discharges. In some implementations, the controller 125 can monitor the battery 190. When the battery 190 is substantially discharged, the controller 125 can execute power down operations so as to avoid, for example, data loss from power loss in the motherboard 120. In some implementations, when the AC power is restored, the controller 125 can perform operations to transition from using the battery 190 to the AC/DC converter 170 to supply power to the motherboard 120.

[0024] In some implementations, the controller 125 can also receive commands and data from a management bus 195 via a wired connection (e.g., through a System management bus (SMBus), an inter-integrated circuit (I²C) bus, a power management bus (PMBus), a twisted pair bus, power lines (such as those providing the AC power to the power supply 115), etc.). In other implementations, the controller 125 receives commands and data from the management bus wirelessly (e.g., through a radio frequency (RF) connection, a Bluetooth connection, a wireless LAN connection, infrared, or through other wireless/mobile communication protocols).

[0025] In some implementations, the controller 125 can perform operations to update the program instructions in the NVM 130 while supplying power to the motherboard continuously. For example, the controller 125 may determine whether an update of the program instructions is permitted based on present conditions of the power supply 115 (e.g., the stability and presence of the AC input signal, the presence of a battery, the remaining battery life of the battery, etc.) Some examples of the operations that can be performed by the controller 125 to update program instructions in the NVM 130 are described with reference to FIGS. 2-5.

[0026] Program instructions for performing the fail-safe operations, the power supply condition monitoring operations, and the communication operations between the power supply 115 and the remote host, for example, are stored in the NVM 130. In some implementations, the controller 125 can perform a checksum operation to verify and maintain the integrity of the program instructions stored in the NVM 130. For example, if the checksum operation detects an error in the application firmware 130a, the controller 125 may notify the remote host to

download a replacement of the application firmware.

[0027] In the depicted example, the NVM 130 includes a plurality of memory portions 180, where one portion of the memory includes the application firmware 130a and one portion includes the boot loader 130b. The application firmware portion 130a and the boot loader 130b may be independent sets of program instructions in the NVM 130. In one implementation, each of the independent sets of program instructions can be executed by the controller 125 to maintain power output to the motherboard 120. In another implementation, the power output is independent of the functionality of the controller 125 and the program instructions stored in the non-volatile memory (NVM) 130.

[0028] In some implementations, the application firmware 130a and the boot loader 130b each contain specialized instructions that are independent of instructions stored in other portions of the NVM 130. For example, the application firmware 130a may have instructions that when executed monitor a temperature of fan exhaust and adjust a fan speed based on this temperature or the application firmware can include instructions that outline a power failure protection process for switching over to the battery if the AC power becomes unstable during normal operation, whereas the boot loader 130b may be responsible for initializing the startup sequences for the application firmware and managing a firmware update process used to update, or modify, the firmware application 130a.

[0029] By executing the program instructions in the boot loader 130b, the controller 125 can modify instructions for the application firmware portion 130a while supplying power continuously to the motherboard 120. For example, the firmware updates can upgrade the operations of the power supply 115 (e.g., adding additional power supply control features or correcting errors in operations of the power supply 115). In another example, the firmware updates can change some of the parameters used by the power supply 115 (e.g., the fan speed, the DC output voltage, or other operation parameters).

[0030] During the firmware update process, the controller 125 and the power supply 115 can remain functional so that the power supply supplies a substantially continuous power output to the motherboard 120. In some implementations, the program instructions stored in the boot loader 130b may include operations that are used for replacing, deleting, supplementing, or modifying program instructions in the application firmware 130a without interrupting power output to the motherboard 120.

[0031] FIGS. 2-5 are flow diagrams that illustrate some examples of methods that the controller 125 can use to update program instructions in one of the memory portions 180 (e.g., the application firmware 130a) while continuously supplying power to the motherboard 120. For example, the power distribution architecture 100 may receive updates from an administrator of the system 100 that includes the power supply 115, where the updates include instructions to modify the application firmware of the power supply 115. The controller 125 can use some of the methods described below to replace the application firmware portion 130a (or other memory portions of the NVM).

[0032] FIG. 2 shows an exemplary method 200 that a controller of a power supply can use as a start-up sequence for the controller. For example, the controller 125 may perform the method to restart the power supply 115, or a portion of the power supply (e.g., restarting the controller 125 while maintaining a substantially stable power output to the motherboard 120), after a hardware reset or a software reset. In some cases, a controller of a power supply can use the method 200 to determine which of the memory portions are accessed by the controller 125 for execution.

[0033] For example, the controller 125 uses the method 200 to determine whether the application firmware 130a or the boot loader 130b is to be used to control the power supply 115. During a startup process, a controller of a power supply can also perform the method 200 to check the integrity of the memory portions. For example, performing the method 200 may substantially reduce and/or prevent the controller from executing program instructions that have been corrupted or erroneously modified.

[0034] In general, a controller can use the method 200 when the controller is reset using software or hardware. The method 200 may further be performed under the control, supervision, and/or monitoring of the controller 125 (or other processing modules) in the system 100. Operations may also be supplemented or augmented by other processing and/or control elements that may be in operative communication with the controller through a network connection 155 coupled to the tray 110. Some or all of the operations may be performed by one or more processors executing instructions tangibly embodied in a signal. The processing may be implemented using analog and/or digital hardware or techniques, either alone or in cooperation with one or more processors executing instructions.

[0035] The method 200 begins when the controller receives a software reset 205 or a hardware reset 210. For example, the controller 125 can receive a software reset after the controller 125 executes a reset instruction. In another example, the power supply 115 can be reset when the controller 125 receives a software reset instruction from the remote host, such as after the controller has updated the application firmware 130a. In another example, the controller 125 can receive a hardware reset after the power supply 115 is powered up after being introduced into the system 100.

[0036] Next, the controller determines whether the power supply operates in a firmware update mode in step 215. In some implementations, the power supply 115 may be operating in various operating modes, including, for example, the firmware update mode (e.g., boot loader mode) and a normal operating mode (e.g., application firmware mode). For example, when the power supply 115 is operating in the firmware update mode, the controller 125 can execute operations to update the application firmware 130a.

[0037] In some implementations, the controller 125 can check a status bit or a status register of the controller 125 or in the NVM 130 to determine the present operating mode of the power supply 115. For example, the controller 125 can check a boot flag register in the NVM 130. If

the boot flag register stores a value of zero, then the controller 125 determines that, for example, the power supply 115 is operating in the firmware update mode. Otherwise, the controller 125 can determine that the power supply 115 is operating in the normal operating mode.

[0038] The operating mode of the power supply can be set in a previous session before the software or the hardware reset is executed. For example, the remote host can transmit a reset instruction to the controller and specify an operating mode (e.g., firmware update mode or normal operating mode) to be executed after reset. In another example, the power supply may be restarted in a pre-determined operating mode (e.g., the firmware update mode) when a reset is issued due to certain events (e.g., a sudden shutdown).

[0039] If the controller determines that the power supply is in the firmware update mode, then, in step 220, the controller starts the power supply in firmware update mode and the method 200 ends. If the controller determines that the power supply is not in the firmware update mode, then the controller calculates, in step 225, a checksum of an application image in a memory. For example, the controller 125 can perform the checksum operation to protect the integrity of the program instructions in the NVM 130 from, for example, accidental data modifications due to memory corruptions or system errors. In some implementations, the controller 125 may add up basic components (e.g., each byte) of the stored instructions in the application firmware image in the NVM 130. In other implementations, the controller 125 may perform a cyclic redundancy check (CRC) to generate the checksum or other operations to calculate the checksum.

[0040] In step 230, the controller determines whether the calculated checksum is correct. For example, the controller can compare the calculated checksum to a checksum of the application firmware that was previously calculated and stored in a register in the NVM 130. If the two values match, then the controller 125 can determine that the application firmware image in the NVM 130 is correct. If the two values do not match, then the controller 125 can determine that the application firmware image in the NVM 130 is not correct.

[0041] In another implementation, the controller 125 can transmit the calculated checksum to a host (such as the firmware update host shown in FIG. 6) for verification. For example, the host can store a checksum of the last application firmware update transmitted to the power supply 115. The controller 125 transmits the calculated checksum to the host for verification that the calculated checksum is substantially equivalent to the checksum stored by the host. If the calculated checksum and the previously stored checksum agree, the host can transmit an indicator to the controller specifying that the checksums are substantially equivalent. Otherwise, the host can indicate that the checksums do not match and that the controller should initiate actions to update or modify the application firmware.

[0042] If the controller determines that the calculated checksum is not correct in step 230, then the step 220 is executed as described above. If the controller determines that the calculated checksum is correct, then the controller starts, in step 235, executing instructions from the

application firmware and the method 200 ends.

[0043] FIG. 3 shows an example of a method 300 that a controller of a power supply can perform when the application firmware 130a in a memory of the power supply is being updated. In some implementations, the controller 125 can use the method 300 to avoid power failures during a firmware update process. In general, the method 300 includes operations to prevent power failures to the controller 125 and fail-safe operations to continue the firmware update process if a main power supply, such as the AC power source 135, fails.

[0044] In some implementations, some or all operations in the method 300 may be executed after the step 215 is executed in the method 200 (FIG. 2). In some implementations, the controller 125 can perform operations in the method 300 when the controller 125 receives an update firmware command from the remote host.

[0045] The method 300 begins in step 305 by determining whether AC power is present and stable. For example, the controller 125 can measure the r.m.s. (root mean square) value of the magnitude and the frequency of the input voltage signal to determine whether the AC power is present and stable. If the AC power is not present or stable, then the power supply controller remains in the normal operating mode (e.g. application firmware mode) in step 310. If the AC power is present and, for example, has a relatively stable voltage magnitude and frequency, then the power supply controller can determine, in optional step 315, whether a back up battery is properly connected. For example, the controller 125 can detect the presence of the battery 190 and can determine whether the battery is correctly installed based on a measured output voltage of the battery, charge counting, previous battery tests, or other methods. However, the power supply 115 does not require the inclusion of a back up battery. If the power supply 115 does not include a back up battery then the method 300 can proceed from step 305 to step 320 if AC power is present and stable.

[0046] In optional step 315, if the back up battery is not properly connected, then the step 310 is executed as described above. If the back up battery is properly connected, then firmware update operations begin in step 320. Some examples of operations for updating firmware in a power supply without interrupting power output are described with reference to FIG. 4.

[0047] Next, the controller determines, in step 325, whether the AC power has been interrupted during the execution of the firmware update instructions. If the AC power is not interrupted, then the controller 125 determines whether the firmware update process is finished in step 330. If the firmware update process is not finished, then the controller continues to monitor the AC power by executing step 325.

[0048] In the step 325, if the AC power is interrupted, then the power supply and the controller can switch to the back up battery in optional step 335, and the controller can monitor the progress of the firmware update for completion as described in step 330.

[0049] As described above, the power supply need not include the back up battery. In certain

implementations without a battery, the tray 110 powers down if AC power is interrupted during execution of the firmware update. When the tray is booted after losing power, the controller can start in boot loader mode and initiate the firmware update operations again as described in association with FIG. 2.

[0050] If the firmware update process is finished as determined in step 330, then, in step 340, the controller switches to operate in the normal operation mode and the method 300 ends. Optionally, the controller may execute normal AC power fail sequence if the power supply is presently using battery back-up power in step 345.

[0051] FIG. 4 shows an example of a method 400 that a controller of a power supply can use to update an application firmware of the controller while maintaining a substantially stable power output. In general, while maintaining the power output of the power supply, the controller can use the method 400 to receive an update for a memory portion (e.g., the application firmware 130a) and protect data integrity of the received update portion. For example, the controller 125 can perform one or more checksum operations on the received data to detect errors in the received data.

[0052] The method 400 begins in step 405 when the controller receives a command to set the operation mode to the firmware update mode. For example, the controller may receive a command from the remote host to restart the controller in firmware update mode.

[0053] In step 410, the controller sets a power supply operation mode to the firmware update mode. In some implementations, the controller may check the main power and the back-up power available for the power supply when performing the step 410. For example, the controller can perform the method 300 when performing the step 410.

[0054] Next, the controller may receive a command to start firmware update process in step 415. After setting the power supply operation in the firmware update mode, the controller receives, in step 420, a portion of a firmware update. For example, a complete firmware update can be segmented into portions, which can be separately transmitted from a host to the controller 125. In some implementations, the portion of a firmware update is transmitted in one or more data packets.

[0055] In one implantation, the controller can segment each portion into smaller portions once received at the controller. For example, the larger portion can be segmented to accommodate a maximum block transfer size of a bus used to update the NVM. In a more specific example, the portion of a firmware update can be substantially equivalent to 64 bytes, which may be the size of a memory page in the NVM. Before transmission to the NVM, the portion can be segmented into two 32-byte portions, where one portion represents a high byte portion of the larger 64-byte portion, and the other portion represents a low byte portion of the larger 64-byte portion. Each smaller portion can then be separately transmitted to the NVM using the bus with a maximum block transfer size of 32-bytes.

[0056] In step 425, the controller writes the received portion to a memory, such as the portion 130a of NVM 130. In some implementations, the controller may combine more than one received data packets to form the entire firmware update portion before writing into the memory. For example, the controller can combine the low byte and the high byte of the firmware update portion before writing a full word into the memory.

[0057] Next, the controller calculates a checksum on the received portion in step 430 and transmits, in step 435 the calculated checksum to the remote host. In step 440, the controller determines whether the calculated checksum is correct. For example, the controller 125 may receive, via the management bus 195, a message from the remote host indicating whether the calculated checksum is correct. If the calculated checksum is determined to be not correct, then step 441 is executed, where the controller requests retransmission of the firmware update and the step 420 is repeated so that the received portion is retransmitted from the host to the controller.

[0058] In some implementations, the step 430 can occur before the step 425. For example, the controller can calculate a checksum of the firmware portion (or the entire received firmware update) before the firmware portion is written to the NVM. The controller can compare the checksum to a pre-calculated checksum (e.g., a checksum that was calculated by the firmware update host before transmission of the firmware or firmware portion to the power supply). If the checksums do not match, then the controller may not write the checksum to memory, but instead, the controller may request retransmission of the firmware.

[0059] If the calculated checksum is determined to be correct, then, the controller or host determines, in step 445, whether there are more firmware update portions to write to complete the firmware update. In some implementations, the controller may include a counter and store a total size of the firmware update. For example, the controller can use the counter to count the size of the received firmware update in this session and compare the counter value with the total size of the firmware update (previously transmitted by the host) to determine whether there are more firmware update portions to write.

[0060] If, in step 445, the controller determines that there are more firmware update portions to write, then the controller advances, in step 450, to a next portion of the firmware update and the step 420 is repeated. If there are no more firmware update portions to write, then a checksum command is received in step 455. In step 460, the controller calculates and transmits a checksum for the entire firmware update to, for example, the remote host. In a next step 465, the controller determines whether the transmitted checksum for the entire firmware update is correct.

[0061] If the controller determines that the transmitted checksum is not correct, then the controller may restart the firmware update in step 470 and the steps 420-465 are repeated.

[0062] If the controller determines that the transmitted checksum is correct, then the controller, in step 475, can receive a software restart command. In some implementations, before

receiving the software restart command from the host, the controller receives a query from the host regarding the status of the firmware update process. If the controller transmits, for example, a response indicating that the firmware update is successfully installed, then the host may send the software restart command. Otherwise, the host may, for example, wait for a predetermined time and transmit the query to the controller for the present firmware update status.

[0063] Next, the controller stores the calculated checksum of the updated firmware in step 480. In some implementations, the controller can use the stored checksum to verify that the firmware has not been erroneously modified or corrupted as described in association with step 230 of FIG. 2. In step 485, the controller restarts the power supply in normal operating mode (e.g., application firmware mode) that uses the updated application firmware, and the method 400 ends.

[0064] FIG. 5 shows an example of a method 500 that a controller of a power supply can use to restrict access to a portion of memory during the firmware update process. In the depicted method 500, the portion of memory is an application memory space in which the controller 125 uses to execute commands received from a remote host.

[0065] To maintain data integrity, a memory portion (e.g., the application firmware 130a) may not be accessible unless a unlock command is received from the remote host, and the power supply is operating in the normal operating mode (e.g., in application firmware mode). Accordingly, by switching to an initialization mode, such as boot loader mode, the controller 125 can restrict access to an application firmware. In some implementations, this restriction is used to prevent access to the application firmware 130a while it is being updated.

[0066] Specifically, the controller 125 can set the application firmware 130a of the NVM 130 to be inaccessible unless a unlock command is received from the remote host via the management bus 195 and unless the power supply 115 is operating in the normal operating mode. When the power supply 115 is operating in the firmware update mode, the controller 125 can restrict access to the application firmware portion until the controller 125 switches the operating mode of the power supply 115 back to the normal operating mode. By restricting access to the portion being updated, the controller is prevented from executing instructions stored in the portion.

[0067] The method 500 begins in step 505 when the controller receives an unlock command and other subsequent commands. Next, the controller, in step 510, determines whether the power supply 115 is in the firmware update mode. If the controller determines that the power supply 115 is in the firmware update mode, then the controller transmits an error message to the remote host indicating that presently no application firmware command can be executed in step 515, and the method 500 ends.

[0068] In step 510, if the controller determines that the power supply 115 is not in the firmware update mode, then the controller unlocks, in step 520, the application memory space in

response to the unlock commands.

[0069] Next, the controller checks whether a time threshold is reached in step 525. For example, the unlock command may only unlock the application memory space for a limited period of time, such as one second. The commands subsequent to the unlock command that are received within one second of the unlock command can be executed. If the controller determines that the time threshold is reached (e.g., one second has lapsed), then the method 500 ends. If the controller determines that the time threshold is not reached, then the controller executes the received subsequent commands in step 530.

[0070] In other implementations, the method 500 can be modified so that it substantially prevents malicious or unintentional commands from being executed by the controller. For example, the unlock command may be required before any command is executed by the controller, regardless of whether the controller is updating the firmware or not.

[0071] This requirement may prevent computing devices in communication with the power supply 115 from causing the controller to execute unauthorized commands. In this implementation, the steps 510, 515, and 520 can be eliminated. Unless the computing device supplies an unlock command, subsequent commands are not executed by the controller. If an unlock command is transmitted (step 550), the following commands are considered legitimate and authorized. The method proceeds to step 525 to determine if the time threshold for execution has lapsed, and if not, any subsequent commands are executed by the controller (step 530). If the time threshold has lapsed, the controller no longer executes received commands unless proceeded by another unlock command.

[0072] FIG. 6 is a schematic diagram illustrating an example of a power delivery system 600 that includes a plurality of power supply management units (PSMUs) 605a, 605b, 605c, 605d and a firmware update host 610. In some implementations, each of the PSMUs include all or a part of the power supply 115. In the depicted example, the PSMUs 605a-d are connected to the firmware update host 610 via the management bus 195. The firmware update host 610 can transmit and receive data and commands to and from the PSMUs 605a-d. In one example, the firmware update host 610 can update an application firmware used by the PSMUs 605a-d using the management bus 195.

[0073] Each of the PSMUs 605a-d can control a power supply, such as the power supply 115, that supplies power to a load (e.g., the motherboard 120). In this example, the PSMUs 605a-d includes memories 615a-d, respectively, and controllers 620a-d, respectively. In some implementations, each of the controllers 620a-d can execute program instructions stored in the memory 615a-d on the same PSMU 605a-d to control the power supply. For example, each of the memories 615a-d can include the application firmware 130a and the boot loader 130b.

[0074] The controllers 620a-d can also execute commands received from the firmware update host 610. In some implementations, the controllers 620a-d may by default restrict access to, for example, the application firmware in the memories 615a-d. To execute commands in the

application firmware on the controllers 620a-d, the firmware update host 610 can transmit an unlock command to unlock the application firmware before transmitting subsequent commands.

[0075] Through the management bus 195, the firmware update host 610 can transmit an updated application firmware and also commands to the PSMUs 605a-d. For example, the firmware update host 610 can read data stored in a specific location in the memories 615a-d. In another example, the firmware update host 610 can write data to a specific location in the memories 615a-d. Additionally, the firmware update host 610 can request one of the controllers 620a-d to calculate a checksum value for the application firmware. In another example, the firmware update host 610 can reset one or more of the PSMUs 605a-d to the boot loader mode.

[0076] The firmware update host 610 can also receive messages from the 620a-d. For example, the controllers 620a-d can transmit a request for a firmware update when the controller determines that the application firmware stored in the memories 615a-d is corrupted or outdated. For example, the controllers 620a-d can be configured so that when a checksum value of the application firmware image in the memories 615a-d is different from a previously stored value, the controllers 620a-d requests a replacement for the application firmware to avoid malfunctioning of the power supply.

[0077] The firmware update host 610 can be a local processing unit or a remote processing unit. In some implementations, the firmware update host 610 may be a local processing unit, such as one of the trays 110 in the architecture 100. For example, in the architecture 100, one of the trays 110 may be pre-selected or instructed to be the firmware update host 610 to update the other trays 110. In some implementations, the firmware update host 610 may be located remotely. For example, the firmware update host 610 may be a laptop computer connected to the racks 140a-c via the management bus 195 or a server connected to the data center 105 through the network 145.

[0078] In normal operation (e.g., application firmware mode), the firmware update host 610 can transmit commands to the controllers 620a-d by first transmitting an unlock command followed by subsequent commands for execution. In some implementations, the controllers 620a-d can use the method 500 to execute the received commands.

[0079] During a firmware update operation, the firmware update host 610 can transmit a restart command to the controllers 620a-d to reset the PSMUs 605a-d in the firmware update mode (e.g., boot loader mode). In some examples, the controllers 620a-d may execute operations described in the method 200 after the PSMUs 605a-d are reset. Next, the controllers 620a-d checks whether the available power is stable for completing the application firmware update process. If the available power is stable, then the controllers 620a-d can interact with the firmware update host to update the application firmware. For example, the controllers 620a-d can update the application firmware using operations in the method 300.

[0080] Depending on the implementation of the system 100, the firmware update process may be executed differently. In some implementations, the PSMUs 605a-d can be updated simultaneously. In other implementations, the firmware update host 610 can update a subset of the PSMUs 605a-d at a time. For example, the firmware update host 610 may only update the PSMUs 605 that have the batteries 190 properly connected to the power supplies 150, and the remaining PSMUs may be updated at a later time when the batteries are properly connected.

[0081] In some implementations, the firmware update host 610 can select one of the PSMUs 605a-d to be updated. For example, the firmware update host 610 may query one of the PSMUs 605a-d for the version of the application firmware stored in the PSMUs 605a-d. If the firmware update host 610 determines that the application firmware in the queried PSMU is not updated, then the firmware update host 610 may start the firmware update process in the queried PSMU. Alternatively, each PSMU may generate an update request based on a seed value, such as a serial number for the associated power supply, stored in the memory 615. Because each serial number may be unique, the requests for updating may occur at different times.

[0082] In some implementations, the firmware update host 610 can specify a specific group of the PSMUs 605a-d to be updated. For example, there may be some hot spots (e.g., hardware that is running at an elevated temperature relative to other hardware within the system 100) identified in the data center 105 (FIG. 1). The firmware update host 610 can address a group of the trays 110 (e.g., by using an address pin in the trays 110) at the hot spots in order to transmit an application firmware update to the group so that, for example, cooling fans on these trays 110 are operated at a higher speed than the surrounding trays 110.

[0083] A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made. For example, although the motherboard 120 was described above as the load for the power supply 115, other electrical loads may be used. For example, the load can include network devices, integrated chips, disk-only systems, memory devices, etc.

[0084] In some implementations, more than one power supply, such as the power supply 115 may power a single motherboard, such as the motherboard 120. In other implementations, more than one motherboard may be powered by a single power supply. Various combinations, such as two power supplies powering three motherboards may also be implemented.

[0085] In other implementations, advantageous results may be achieved if the steps of the disclosed techniques were performed in a different sequence, if components in the disclosed systems were combined in a different manner, or if the components were replaced or supplemented by other components. Accordingly, other embodiments are within the scope of the following claims.

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US20060020832A1 [0006]

P A T E N T K R A V

1. System k e n d e t e g n e t v e d:

- et effektkonverteringsmodul (115);
en dataport (195);
- 5 en styreenhed (125) til at modtage data fra dataporten (195); og
et datalager (130), som indeholder instruktioner, som når de udføres af styreenheden
(125) foretager operationer for at styre effektkonverteringsmodulet (115), hvor operatio-
nerne omfatter:
at udføre instruktionerne i en første del (130b) af datalageret (130) under opstart af styre-
10 enheden (125);
at udføre instruktionerne i en anden del (130a) af datalageret (130) efter opstart af styre-
enheden (125), hvor instruktionerne i den første delen (130b) og instruktionerne i den an-
den del (130a) hver kan udføres af styreenheden for at styre effektkonverteringsmodulet
(115) for at opretholde et effektoutput til en belastning;
- 15 at modtage et signal fra dataporten (195), som indikerer at instruktionerne i den anden del
(130a) skal modificeres;
at modtage modificerede instruktioner fra dataporten (195), som skal lagres i den anden
del (130a); og
at udføre, som svar på signalet, instruktionerne i den første delen (130b) til at lagre de
20 modificerede instruktioner i den anden del (130a) og at opretholde et effektoutput til en
belastning under instruktionsindlæsningsoperationer;
hvor systemet yderligere omfatter:
en primær effektkildeindgangsport til at tilvejebringe effekt til styreenheden (125) og en
sekundær effektkildeindgangsport til at tilvejebringe en alternativ effektkilde til styreenhe-
25 den (125), og
hvor operationerne yderligere omfatter:
at monitorere en effekt modtaget fra den primære effektkildeindgangsport for at be-
stemme tilstedeværelse eller stabilitet af vekselstrømseffekt ved den primære effektkilde-
indgangsport, og
30 at skifte til effekt, som er tilgængeligt fra den sekundære effektkildeindgangsport, hvis den
modtagne vekselstrømseffekt er bestemt til at være fraværende eller ustabil.

2. System ifølge krav 1, hvor de modificerede instruktioner erstatter mindst en del af
instruktionerne tidligere lagret i den anden del (130a) af datalageret (130).

3. System følge krav 1, hvor de modificerede instruktioner supplerer instruktionerne
35 tidligere lagret i den anden del (130a) af datalageret.

4. System ifølge krav 1, hvor operationerne yderligere omfatter:
at begrænse adgang til instruktionerne i den anden del (130a) efter modtagelse af signa-
let, som indikerer at instruktionerne i den anden del (130a) skal modificeres; og

at gendanne adgang til de modificerede instruktioner i den anden del (130a) efter instruktionerne i den første del (130b) har lagret de modificerede instruktionerne i den anden del (130a).

5 5. System ifølge krav 1, hvor den primære effektkildeindgangsport er en vekselstrømsindgangsport og hvor den sekundære strømkildeindgangsport er en jævnstrømsindgangsport.

6. System ifølge krav 5, hvor jævnstrømsindgangsporten er en batteriindgangsport.

10 7. System ifølge krav 1, hvor operationerne yderligere omfatter at transmittere en kontrolsum til et værtssystem til validering, hvor kontrolsummen er mindst delvist genereret fra de modificerede instruktioner lagret i den anden del (130a).

15 8. System ifølge krav 1, hvor operationerne yderligere omfatter at anmode om en retransmission af de modificerede instruktioner til dataporten (195), hvis en kontrolsum baseret på mindst en del af de modificerede instruktioner i den anden del (130a) ikke stemmer overens med en lagret kontrolsum tidligere bestemt baseret på den mindst ene del.

20 9. Fremgangsmåde til at indlæse instruktioner ind i en programmerbar strømforsyning (115), som har hukommelse, hvor fremgangsmåden er kendt ved: at opretholde et effektoutput til en belastning for den programmerbare strømforsyning ved at udføre instruktioner i en første del (130b) af hukommelsen under instruktionsindlæsningsoperationer, hvor instruktionsindlæsningsoperationerne omfatter:

at begrænse adgang til en anden del (130a) af den programmerbare strømforsynings hukommelse, hvor hukommelsen er forbundet til en styreenhed (125) af den programmerbare strømforsyning;

25 at modtage instruktioner gennem en intern bus anvendt til kommunikation med strømforsyningen, hvor de modtagne instruktioner er lagret i den anden del (130a) af hukommelsen;

og

at gendanne adgang til den anden del (130a) af hukommelsen, så styreenheden (125) kan eksekvere de modtagne instruktioner,

30 hvor instruktionerne i den første del (130b) og instruktionerne i den anden del (130a) hver kan udføres af styreenheden for at styre effektkonverteringsmodulet (115) for at opretholde et effektoutput til en belastning;

at monitorere en effekt modtaget fra en primær effektkildeindgangsport i den programmerbare strømforsyning for at bestemme tilstedeværelse eller stabilitet af vekselstrømseffekt i den primære effektkildeindgangsport, og

35 at skifte til effekt som er tilgængeligt fra en sekundær effektkildeindgang, hvis den modtagne vekselstrømseffekt er bestemt til at være fraværende eller ustabil.

10. Fremgangsmåden ifølge krav 9, hvor at begrænse adgang til den anden del af hukommelsen omfatter at indlede en batterireservetilstand, hvor et batteri anvendes til at

forsyning et output fra strømforsyningen.

11. Fremgangsmåden ifølge krav 9, hvor at gendanne adgang til delen af hukommelsen omfatter at indlede en vekselstrømsdrevet tilstand, hvor vekselstrømseffektkilden anvendes til at forsyne et output fra strømforsyningen.

5 12. Fremgangsmåden ifølge krav 9, som yderligere omfatter at modtage en forespørgsel gennem den interne bus, hvor forespørgslen omfatter et svar på hvorvidt de modtagne instruktioner blev installeret succesfuldt.

DRAWINGS




FIG. 1

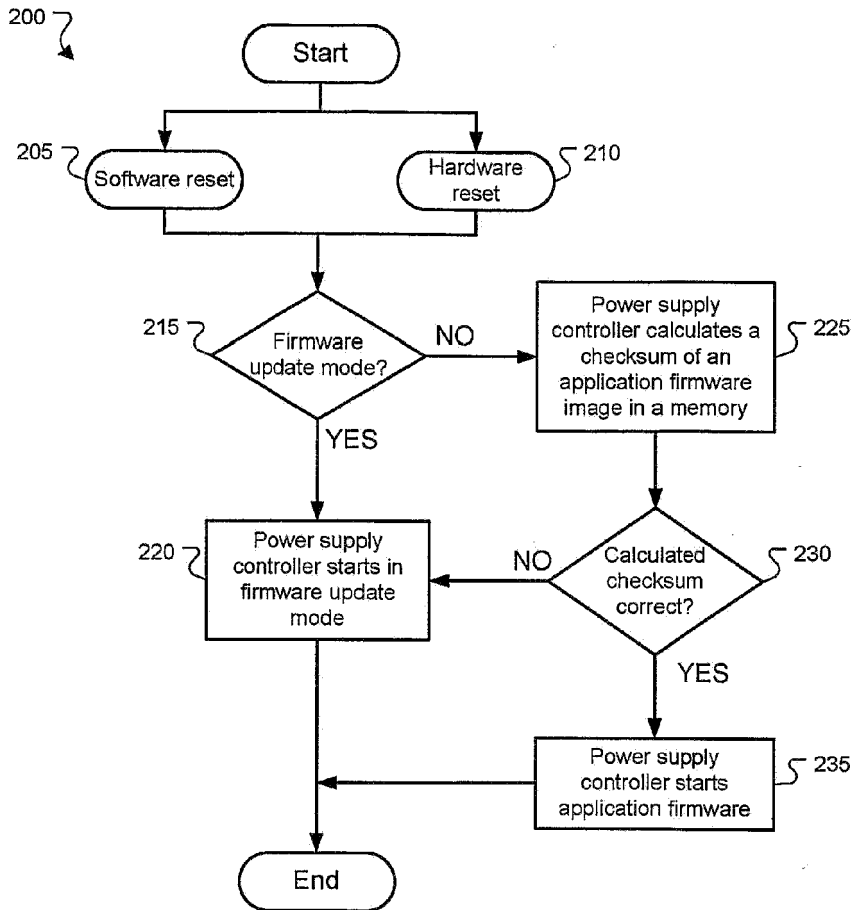


FIG. 2

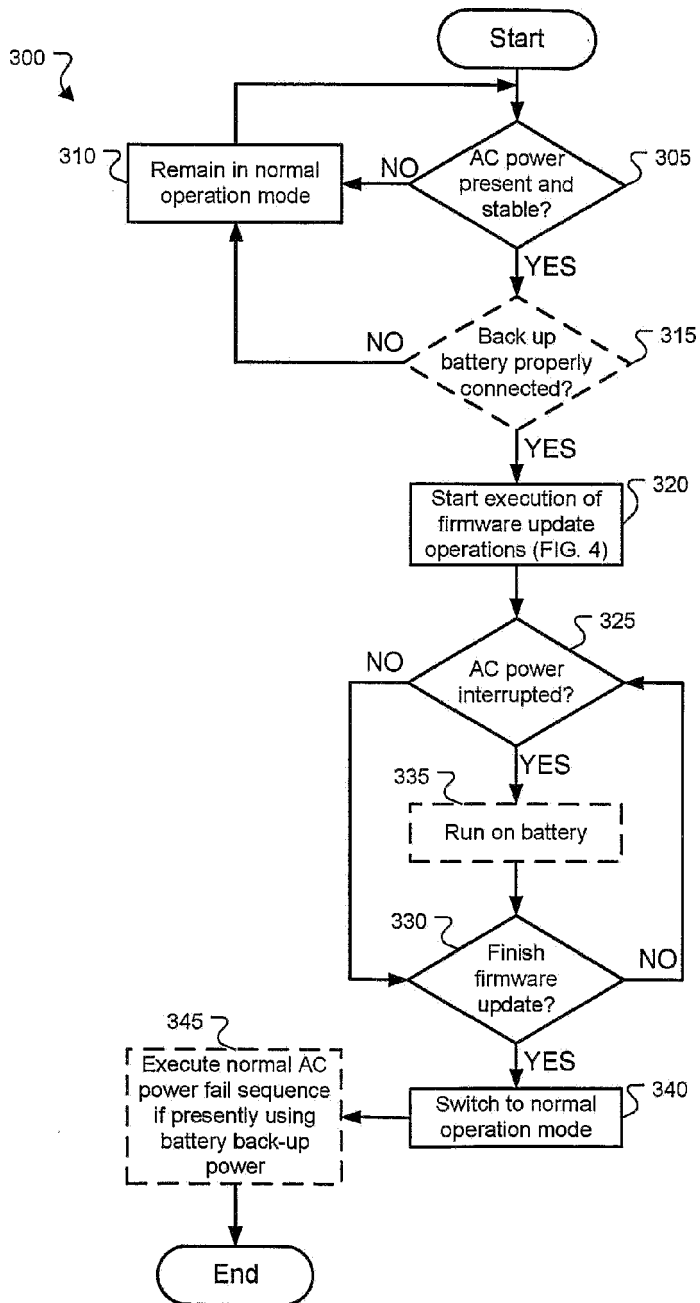


FIG. 3

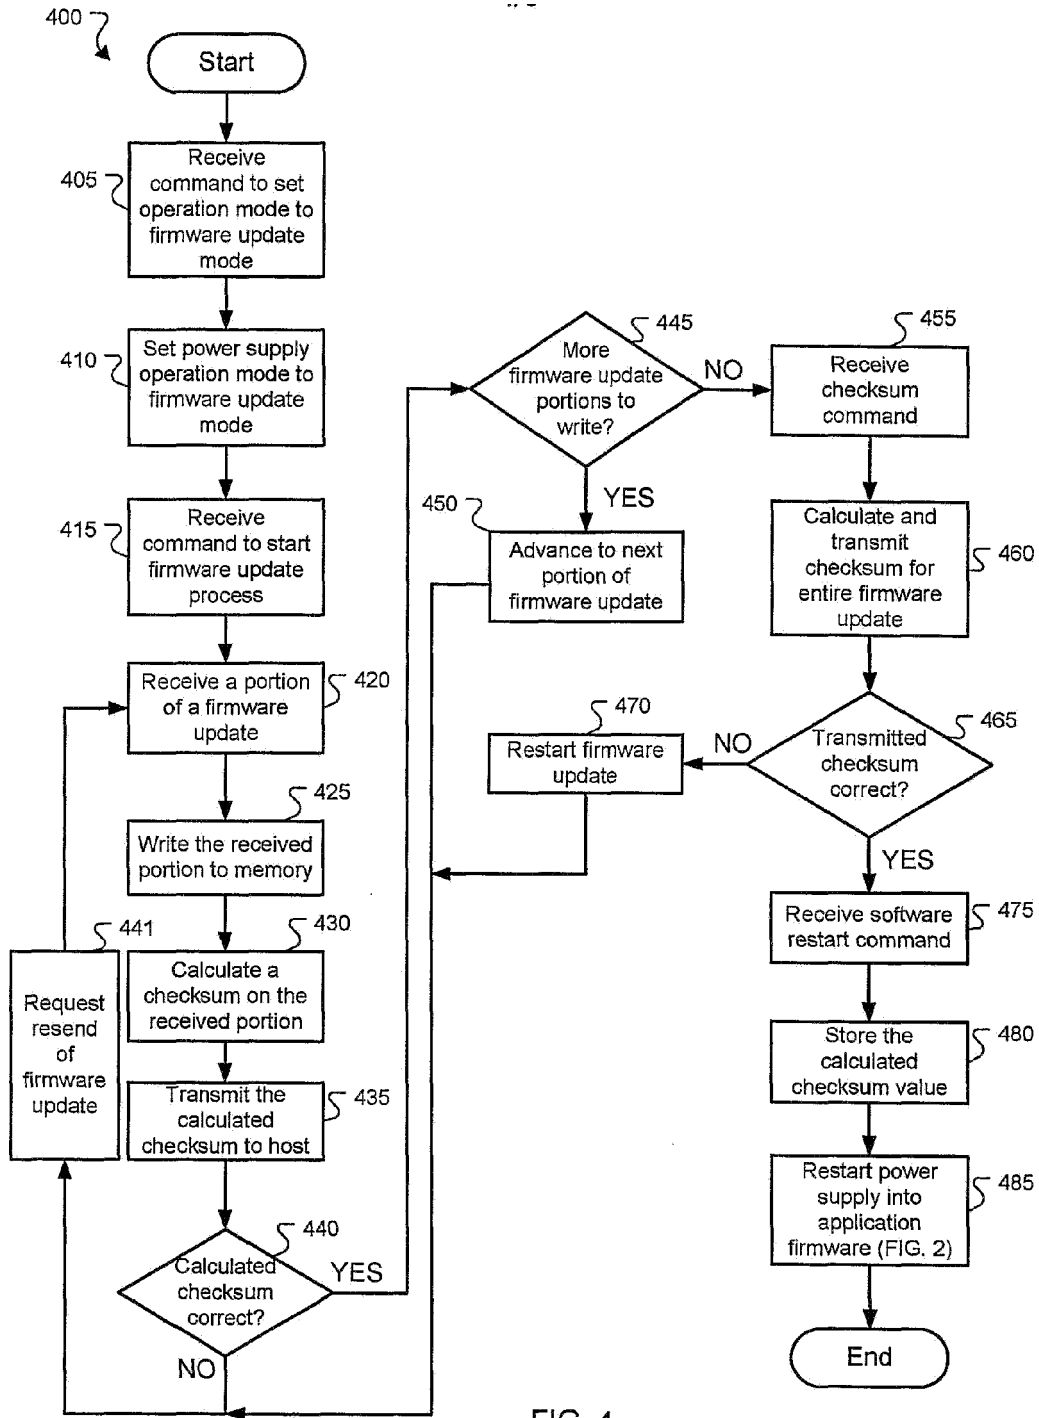


FIG. 4

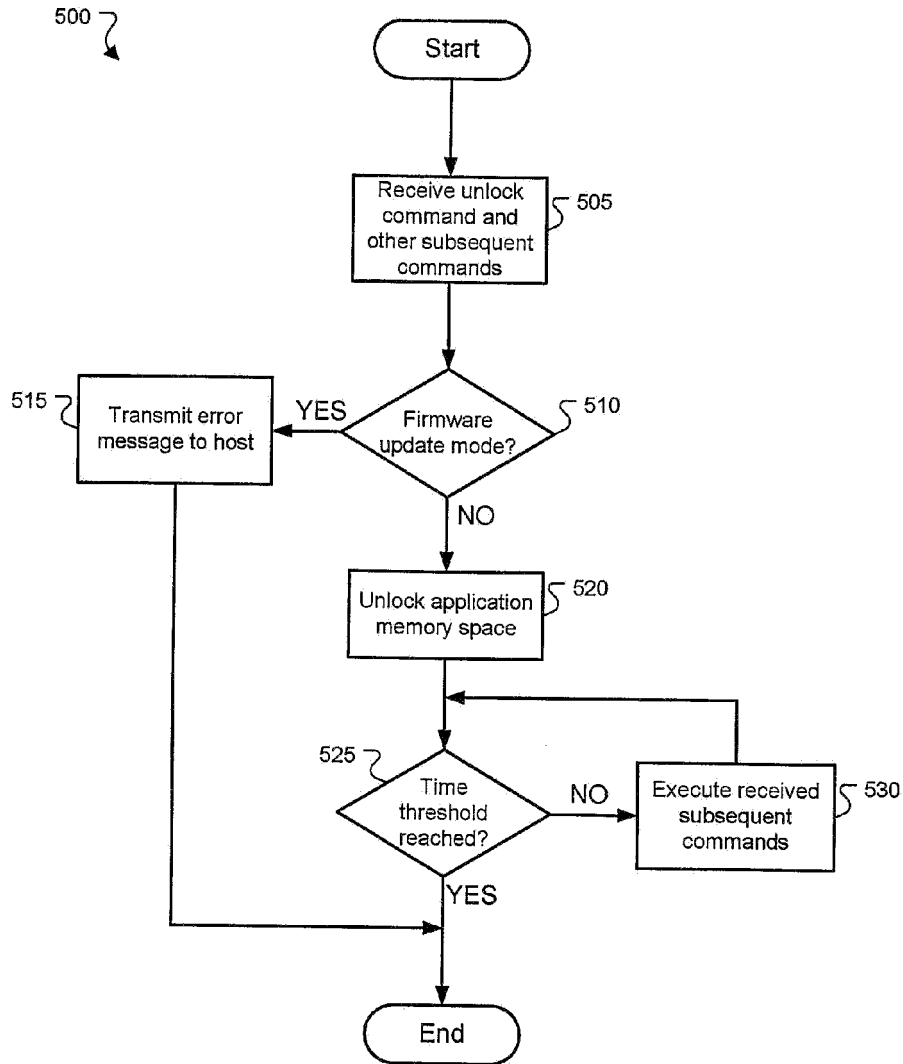


FIG. 5




FIG. 6