

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
17 April 2008 (17.04.2008)

PCT

(10) International Publication Number
WO 2008/045044 A1

(51) International Patent Classification:

G06F 1/26 (2006.01) G06F 9/44 (2006.01)
G06F 9/06 (2006.01) G06F 15/177 (2006.01)
G06F 9/30 (2006.01)

(21) International Application Number:

PCT/US2006/039715

(22) International Filing Date: 10 October 2006 (10.10.2006)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant (for all designated States except US): EX-AFLOP LLC [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): KRIEGER, Ken [US/US]; P.O. Box 11694, Jackson, Wyoming 83002 (US). BORCHERS, Albert [US/US]; 106 Arroyo Court, Santa Cruz, California 95060 (US).

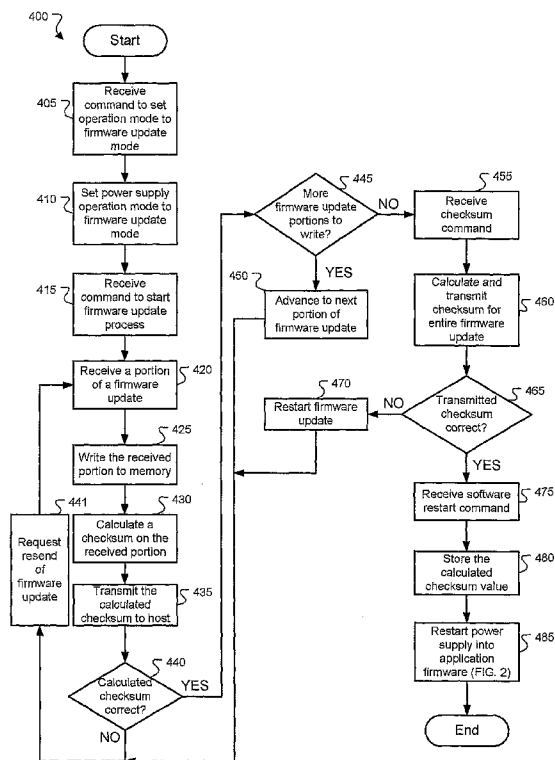
(74) Agents: MCGUIRE, Joshua J. et al.; Fish & Richardson P.C., P.O. Box 1022, Minneapolis, Minnesota 55440-1022 (US).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— with international search report

(54) Title: UPDATING A POWER SUPPLY MICROCONTROLLER



(57) Abstract: A system is described that includes a power conversion module, a data port, a controller to receive data from the data port, and a data store containing instructions that when executed by the controller perform operations to control the power conversion module. The operations include executing instructions in a first portion of the data store during booting of the controller, executing instructions in a second portion of the data store after the controller has booted, receiving a signal from the data port indicating the instructions in the second portion are to be modified, receiving from the data port modified instructions to be stored in the second portion, and executing, in response to the signal, the instructions in the first portion to store the modified instructions in the second portion.

WO 2008/045044 A1

Updating a Power Supply Microcontroller

TECHNICAL FIELD

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

BACKGROUND

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

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.

15 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.

Certain linear regulator power supplies use a transformer to
20 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.

5

SUMMARY

The present specification relates to updating power supplies.

In a first general aspect, a system includes a power conversion module, a data port, a controller to receive data from the data port, and a data store containing instructions that when executed by the controller
10 perform operations to control the power conversion module. The operations include executing instructions in a first portion of the data store during booting of the controller, executing instructions in a second portion of the data store after the controller has booted, receiving a signal from the data port indicating the instructions in the second portion are to
15 be modified, receiving from the data port modified instructions to be stored in the second portion, and executing, in response to the signal, the instructions in the first portion to store the modified instructions in the second portion.

In a second general aspect, a method for loading instructions into
20 a programmable power supply having memory is described. The method includes maintaining a power output for the programmable power supply during instruction loading operations. The instruction loading operations include restricting access to a portion of the programmable power supply's memory, where the memory is coupled to a controller of the

programmable power supply. The operations also include receiving instructions through an internal bus used for communication with the power supply, where the received instructions are stored in the portion of the memory, and restoring access to the portion of the memory so that
5 the controller can execute the received instructions.

In another general aspect, a method is described. The method includes loading instructions into a power supply system, and executing instructions in a first portion of a data store during initialization of a controller, wherein the data store and the controller are included in the
10 power supply system. The method also includes executing instructions in a second portion of the data store after the power supply system has been initialized, receiving a signal indicating the instructions in the second portion are to be modified, and executing the instructions in the first portion in response to the signal to store modified instructions in the
15 second portion.

In yet another general aspect, a system is described. The system includes a power converter, a processor, and means for loading instructions while maintaining a power output from the power converter, the instructions to be executed by the processor to control the power
20 converter. Loading the instructions includes executing a first set of instructions to boot the processor, executing a second set of instructions after the processor has booted, receiving a signal to modify the second set of instructions in response to a received signal, receiving a modified set of instructions, and using the first set of instructions to update the
25 second instructions with the modified set of instructions, wherein

PCT/US2006/039715

execution of the second set of instructions is restricted until the update is complete.

The systems and techniques described here may provide one or more of the following advantages. First, cost savings and operational up-
5 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
10 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.

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

DESCRIPTION OF DRAWINGS

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

20 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
5 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).

Like reference symbols in the various drawings indicate like
10 elements.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

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
15 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
20 power supply 115 or the tray 110.

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
5 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.

In some implementations, the boot loader 130b includes
instructions that when executed prepare the controller 125 for execution
10 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.

In the depicted example, the data center 105 receives AC power
from an AC power source 135. For example, the AC power source 135
15 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
20 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.

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
5 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
10 information can be exchanged between the trays 110 the data center 105.

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
15 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
20 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
25 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.

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.

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.

The motherboard 120 operates using the DC power supplied from
5 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-
10 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
15 determined. For example, the power supply 115 can determine the voltage set point based on, for example, a present temperature of the motherboard 120.

Various operations in the power supply 115 are controlled by the controller 125. For example, the controller 125 can execute application
20 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
25 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.

5 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
10 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
15 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
20 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
25 supply power to the motherboard 120.

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, 5 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 10 communication protocols).

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 15 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 20 to FIGS. 2-5.

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 25 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.

5 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
10 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.

15 In some implementations, the application firmware 130a and the boot loader 103b 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
20 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.

By executing the program instructions in the boot loader 130b, the
controller 125 can modify instructions for the application firmware portion
5 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
10 by the power supply 115 (e.g., the fan speed, the DC output voltage, or
other operation parameters).

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.
15 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.

FIGS. 2-5 are flow diagrams that illustrate some examples of
20 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).

5 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
10 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.

 For example, the controller 125 uses the method 200 to determine
15 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
20 executing program instructions that have been corrupted or erroneously modified.

 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

5 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.

10 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
15 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.

Next, the controller determines whether the power supply operates
20 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.

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.

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).

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.

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.

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 check sum 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.

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.

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.

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.

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 property 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.

In optional step 315, if the back up battery is not property connected, then the step 310 is executed as described above. If the back up battery is property 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.

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.

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
5 step 335, and the controller can monitor the progress of the firmware update for completion as described in step 330.

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
10 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.

If the firmware update process is finished as determined in step 330, then, in step 340, the controller switches to operate in the normal
15 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.

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
20 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.

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.

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.

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.

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
5 transmitted to the NVM using the bus with a maximum block transfer size of 32-bytes.

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
10 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.

Next, the controller calculates a checksum on the received portion
15 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
20 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.

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
5 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.

10 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
15 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.

If, in step 445, the controller determines that there are more
20 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
25 example, the remote host. In a next step 465, the controller determines

whether the transmitted checksum for the entire firmware update is correct.

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

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
10 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
15 present firmware update status.

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.
20 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.

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.

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.

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.

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
5 ends.

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.

Next, the controller checks whether a time threshold is reached in
10 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
15 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.

In other implementations, the method 500 can be modified so that it substantially prevents malicious or unintentional commands from being
20 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.

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
5 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
10 unlock command.

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
15 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
20 firmware used by the PSMUs 605a-d using the management bus 195.

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
25 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.

5 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
10 610 can transmit a unlock command to unlock the application firmware before transmitting subsequent commands.

 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
15 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
20 update host 610 can reset one or more of the PSMUs 605a-d to the boot loader mode.

 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
5 requests a replacement for the application firmware to avoid malfunctioning of the power supply.

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
10 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
15 racks 140a-c via the management bus 195 or a server connected to the data center 105 through the network 145.

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
20 execution. In some implementations, the controllers 620a-d can use the method 500 to execute the received commands.

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
5 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.

Depending on the implementation of the system 100, the firmware
10 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
15 connected to the power supplies 150, and the remaining PSMUs may be updated at a later time when the batteries are properly connected.

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
20 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
25 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.

In some implementations, the firmware update host 610 can specify a specific group of the PSMUs 605a-d to be updated. For
5 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
10 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.

A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For
15 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.

In some implementations, more than one power supply, such as
20 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.

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
5 other components. Accordingly, other embodiments are within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A system comprising:
 - a power conversion module;
 - a data port;
 - 5 a controller to receive data from the data port; and
 - a data store containing instructions that when executed by the controller perform operations to control the power conversion module, the operations comprising:
 - executing instructions in a first portion of the data store
 - 10 during booting of the controller;
 - executing instructions in a second portion of the data store after the controller has booted;
 - receiving a signal from the data port indicating the instructions in the second portion are to be modified;
 - 15 receiving from the data port modified instructions to be stored in the second portion; and
 - executing, in response to the signal, the instructions in the first portion to store the modified instructions in the second portion.

- 20 2. The system of claim 1, wherein the modified instructions replace at least a portion of the instructions previously stored in the second portion of the data store.

3. The system of claim 1, wherein the modified instructions supplement the instructions previously stored in the second portion of the data store.
4. The system of claim 1, wherein the data port receives data from a system management bus (SMBus).
5. The system of claim 1, wherein the operations further comprise restricting access to the instructions in the second portion after receiving the signal indicating the instructions in the second portion are to be modified.
6. The system of claim 5, wherein the operations further comprise restoring access to the modified instructions in the second portion after the instructions in the first portion has stored the modified instructions in the second portion.
7. The system of claim 1, further comprising a primary power source input port to provide power to the controller and a secondary power source input port to provide an alternative power source to the controller.
8. The system of claim 7, wherein the primary power source input port is an alternating current (AC) input port and the secondary power source input port is a DC input port.

9. The system of claim 8, wherein the DC input port is a battery input port.

10. The system of claim 7, wherein the operations further comprise
5 monitoring a power received from the primary power source input port to determine presence or stability of AC power at the primary power source input port.

11. The system of claim 10, wherein the operations further comprise
10 switching to power available from the secondary power source input if the AC power received is determined to be absent or unstable.

12. The system of claim 1, wherein the operations further comprise transmitting a checksum to a host system for validation, the checksum at
15 least partially generated from the modified instructions stored in the second portion.

13. The system of claim 1, wherein the operations further comprise receiving at the data port a retransmission of the modified instructions if
20 the checksum is invalid.

14. The system of claim 1, wherein the operations further comprise requesting a retransmission of the modified instructions to the data port if a checksum based on at least a portion of the modified instructions in the

second portion does not match a stored checksum previously determined based on the at least a portion.

15. The system of claim 1, wherein the data port comprises an infrared
5 receiver.

16. The system of claim 1, wherein the data port comprises a radio frequency (RF) receiver.

10 17. The system of claim 1, wherein the data port comprises a fiber optic receiver.

18. The system of claim 1, further comprising a common bus that couples a plurality of data ports, each data port associated with a power
15 conversion module, a controller, and a data store.

19. The system of claim 1, wherein the common bus is further coupled to a system host configured to transmit modified instructions to the plurality of data ports.

20

20. The system of claim 19, wherein the system host is further configured to transmit individualized modified instructions to at least one of the plurality of data ports.

21. The system of claim 19, wherein the system host is further configured to receive requests to transmit the modified instructions to the plurality of data ports at different times.

5 22. The system of claim 1, wherein the operations further comprise executing the modified instructions.

23. The system of claim 1, wherein the first and second portions of the data store are logical divisions or physical divisions.

10

24. The system of claim 1, wherein the instructions comprises data.

25. A method for loading instructions into a programmable power supply having memory, the method comprising:

15 maintaining a power output for the programmable power supply during instruction loading operations, the instruction loading operations comprising:

restricting access to a portion of the programmable power supply's memory, the memory coupled to a controller of the programmable power
20 supply;

receiving instructions through an internal bus used for communication with the power supply, wherein the received instructions are stored in the portion of the memory; and

restoring access to the portion of the memory so that the controller
25 can execute the received instructions.

26. The method of claim 25, wherein the instruction loading operations comprise modifying instructions previously stored in the portion of the memory.
- 5
27. The method of claim 25, wherein the instructions comprise firmware
28. The method of claim 27, wherein the firmware comprises data and executable instructions.
- 10
29. The method of claim 25, wherein restricting access to the portion comprises preventing the controller from executing instructions stored in the portion.
- 15
30. The method of claim 25, wherein the bus comprises a system management bus.
31. The method of claim 25, wherein limiting access to the portion of the memory comprises initiating a battery backup mode, wherein a battery is
- 20
- used to supply an output of the power supply.
32. The method of claim 25, wherein restoring access to the portion of the memory comprises initiating an alternating current (AC) powered mode, wherein an AC power source is used to supply an output of the
- 25
- power supply.

33. The method of claim 25, further comprising receiving a query through the internal bus, the query comprising a response as to whether the received instructions were installed successfully.

5

34. A method comprising loading instructions into a power supply system:
executing instructions in a first portion of a data store during initialization of a controller, wherein the data store and the controller are included in the power supply system;

10 executing instructions in a second portion of the data store after the power supply system has been initialized;

receiving a signal indicating the instructions in the second portion are to be modified; and

15 executing the instructions in the first portion in response to the signal to store modified instructions in the second portion.

35. A system comprising:

a power converter;

a processor; and

20 means for loading instructions while maintaining a power output from the power converter, the instructions to be executed by the processor to control the power converter, wherein loading the instructions comprises:

executing a first set of instructions to boot the processor;

executing a second set of instructions after the processor has
booted; receiving a signal to modify the second set of
instructions in response to a received signal;
receiving a modified set of instructions; and
5 using the first set of instructions to update the second instructions
with the modified set of instructions, wherein execution of the second set
of instructions is restricted until the update is complete.




FIG. 1

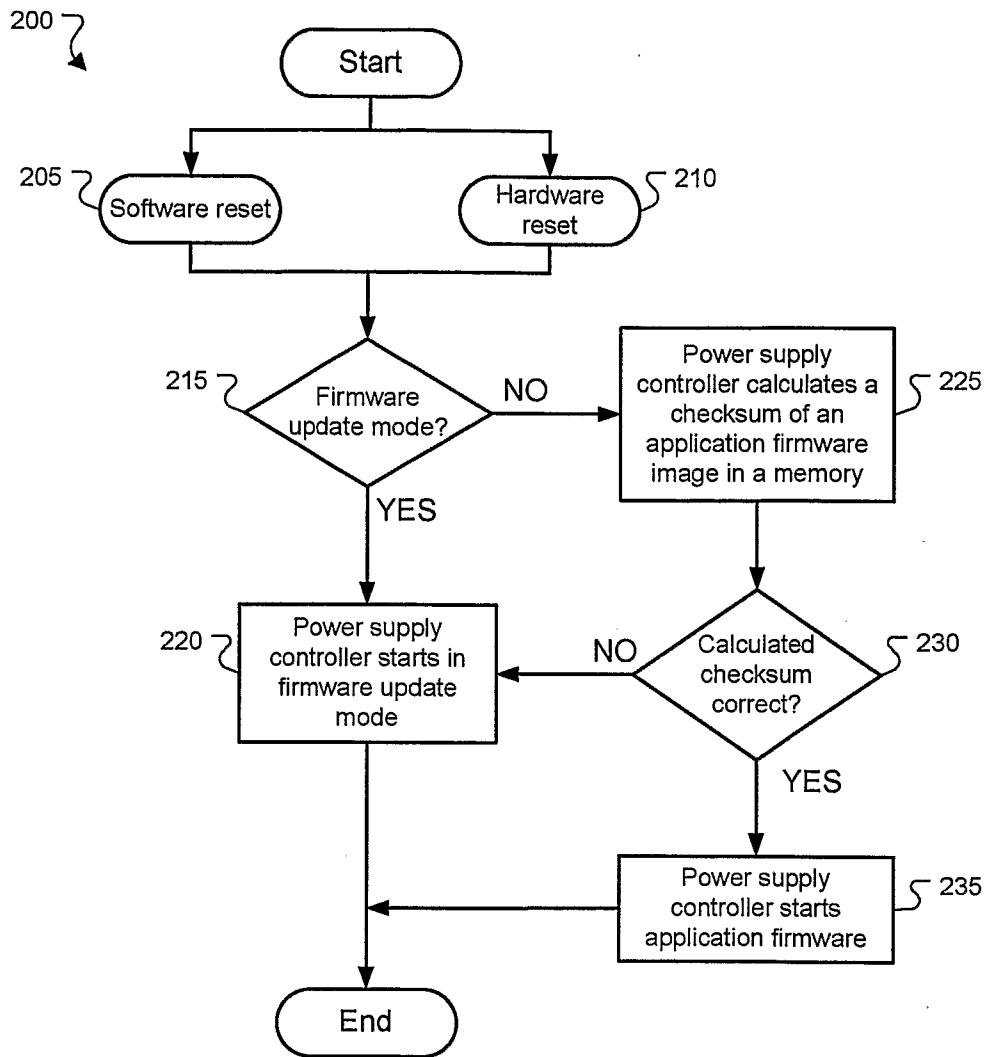


FIG. 2

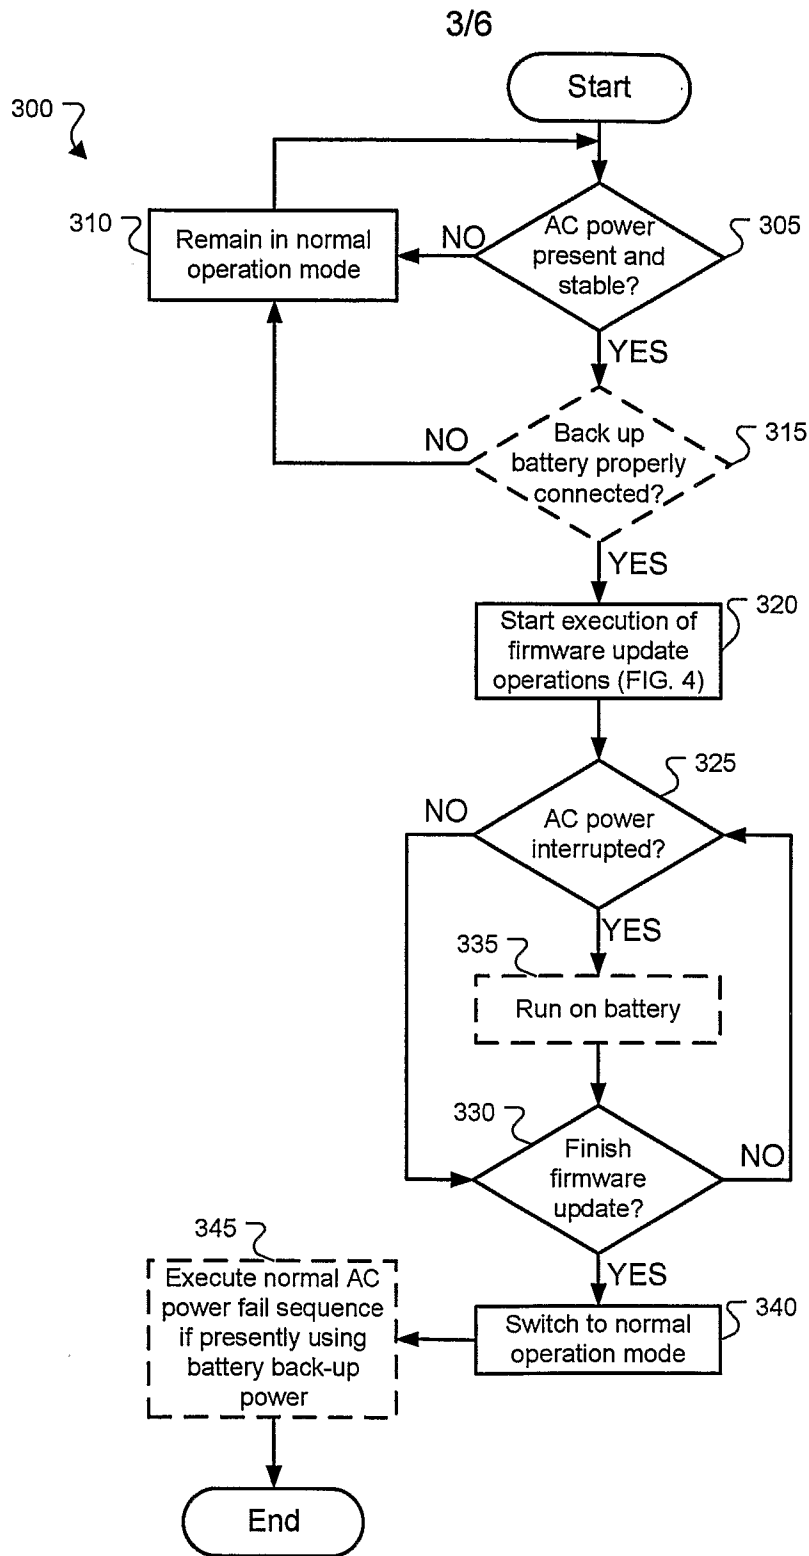


FIG. 3

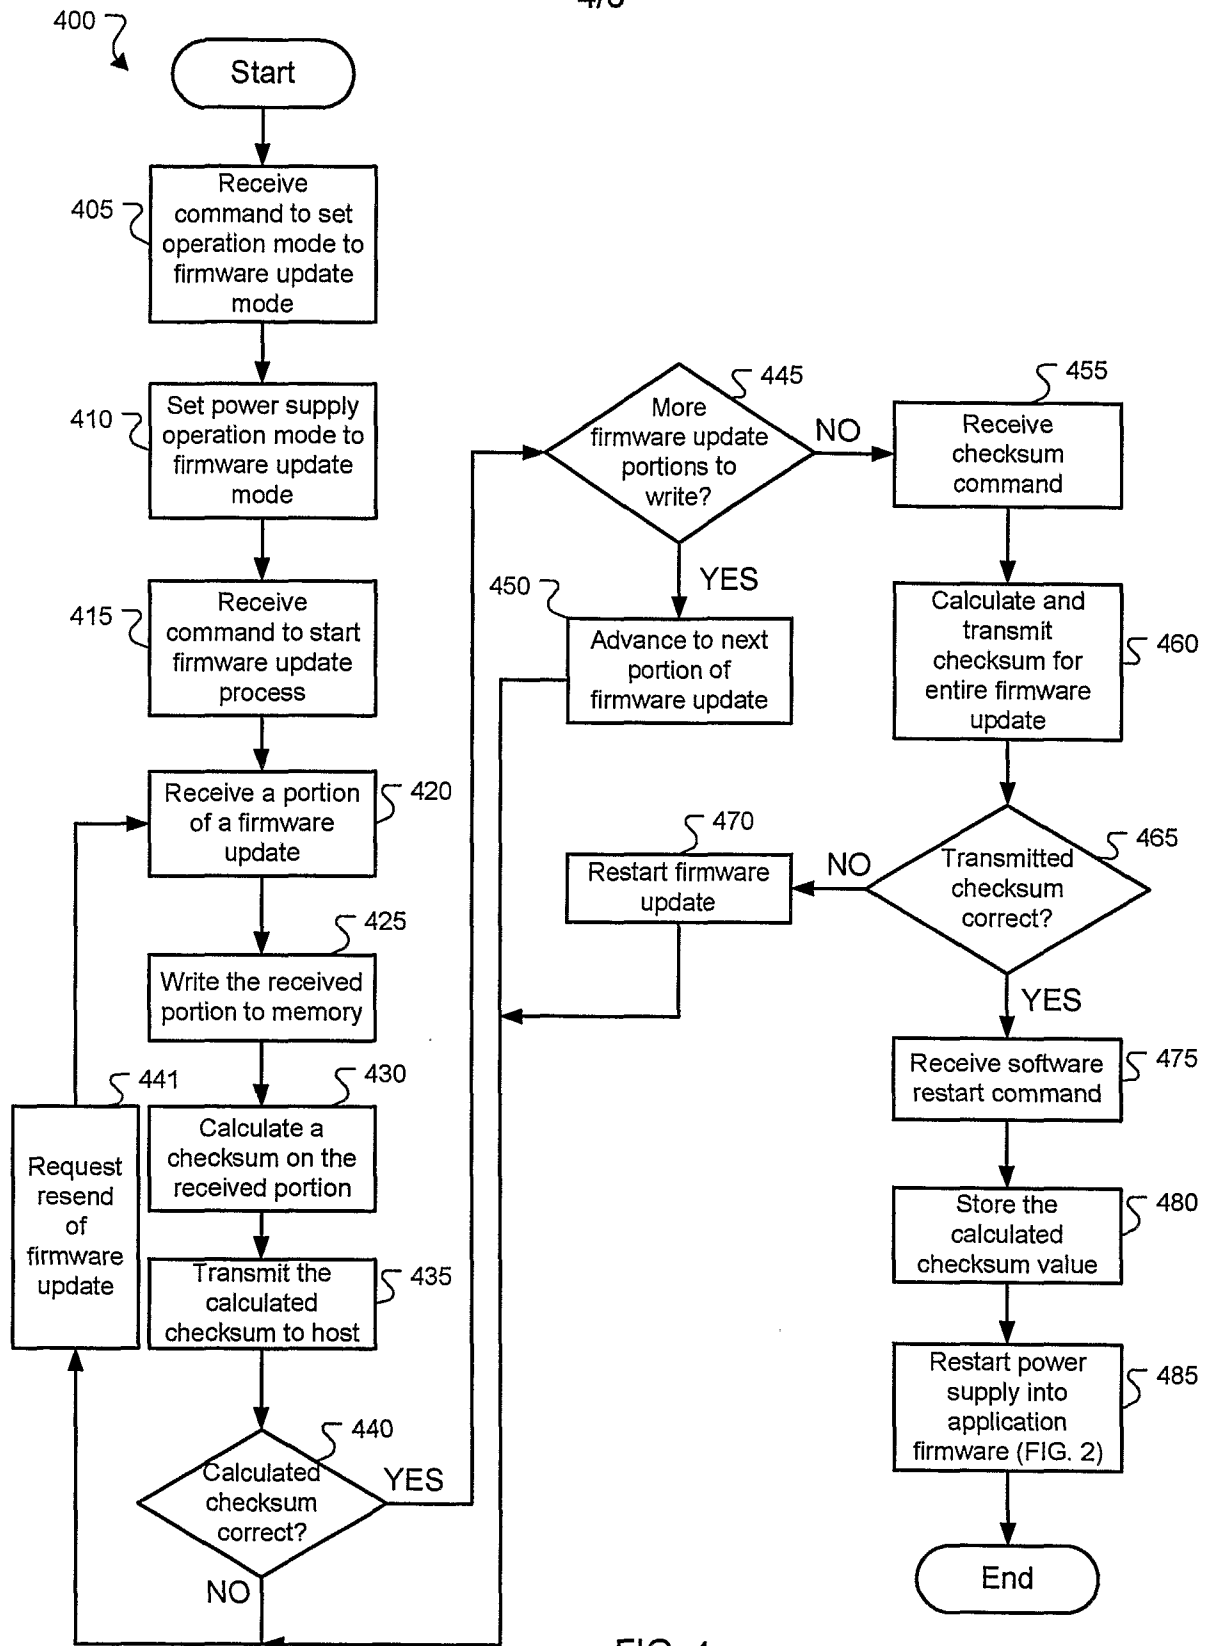


FIG. 4

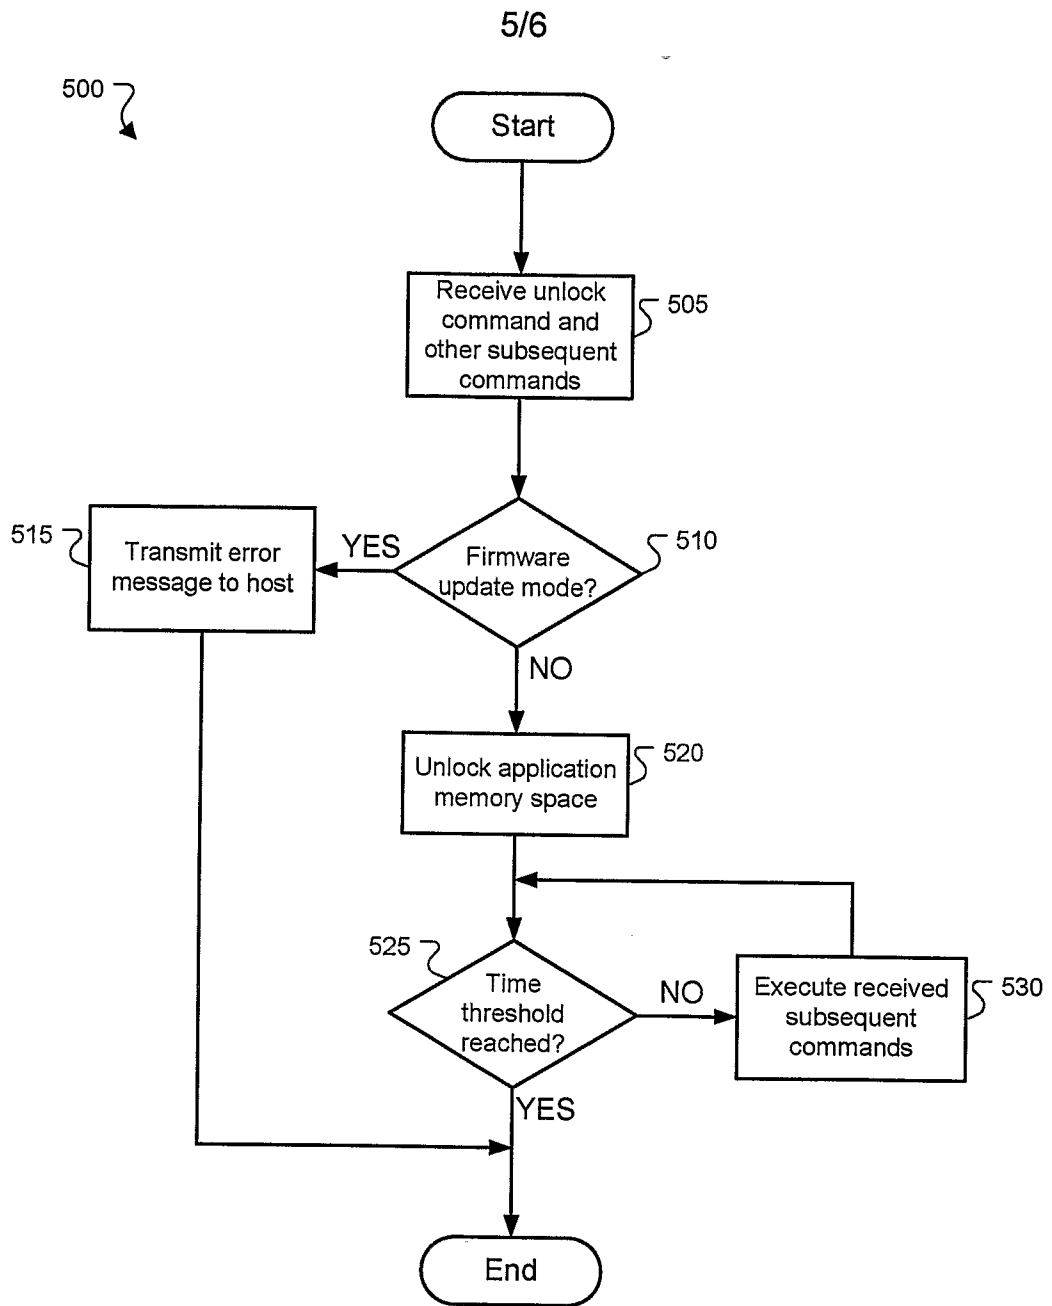


FIG. 5

6/6




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US06/39715

A. CLASSIFICATION OF SUBJECT MATTER
 IPC: **G06F 1/26**(2006.01),**9/06**(2006.01),**9/30**(2006.01),**9/44**(2006.01),**15/177**(2006.01)

 USPC: 713/1,2,100
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 U.S. : 713/1,2,100,300

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 2006/0020832 A1 (HUNG) 26 January 2006 (26.01.2006), paragraphs [0020]-[0022]	1, 2, 4-6, 18-20, 22-30, 34 and 35 ----- 3, 12-17 12-14
Y	US 6,584,559 B1 (HUH et al) 24 June 2003 (24.06.2003), column 3, line 61 through column 4, line 31	3
Y	US 6,009,177(SUDIA) 28 December 1999 (28.12.1999) colum 38, line 64 through col. 39, line 32	3
Y	US 2005/0088201 A1 (DEVLIN et al) 28 April 2005 (28.04.2005), page 4, paragraph [0046].	1, 2, 4-6, 12-20, 22-30, 34 and 35

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

* Special categories of cited documents:		
"A" document defining the general state of the art which is not considered to be of particular relevance	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent published on or after the international filing date	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means		
"P" document published prior to the international filing date but later than the priority date claimed	"&"	document member of the same patent family

Date of the actual completion of the international search 01 October 2007 (01.10.2007)	Date of mailing of the international search report 05 NOV 2007
-------------------------------------------------------------------------------------------	--------------------------------------------------------------------------

Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450 Alexandria, Virginia 22313-1450 Facsimile No. (571) 273-3201	Authorized officer James K. Trujillo Jacqueline A. Whitfield Special Project Asst. Telephone No. (571) 272-2100
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------