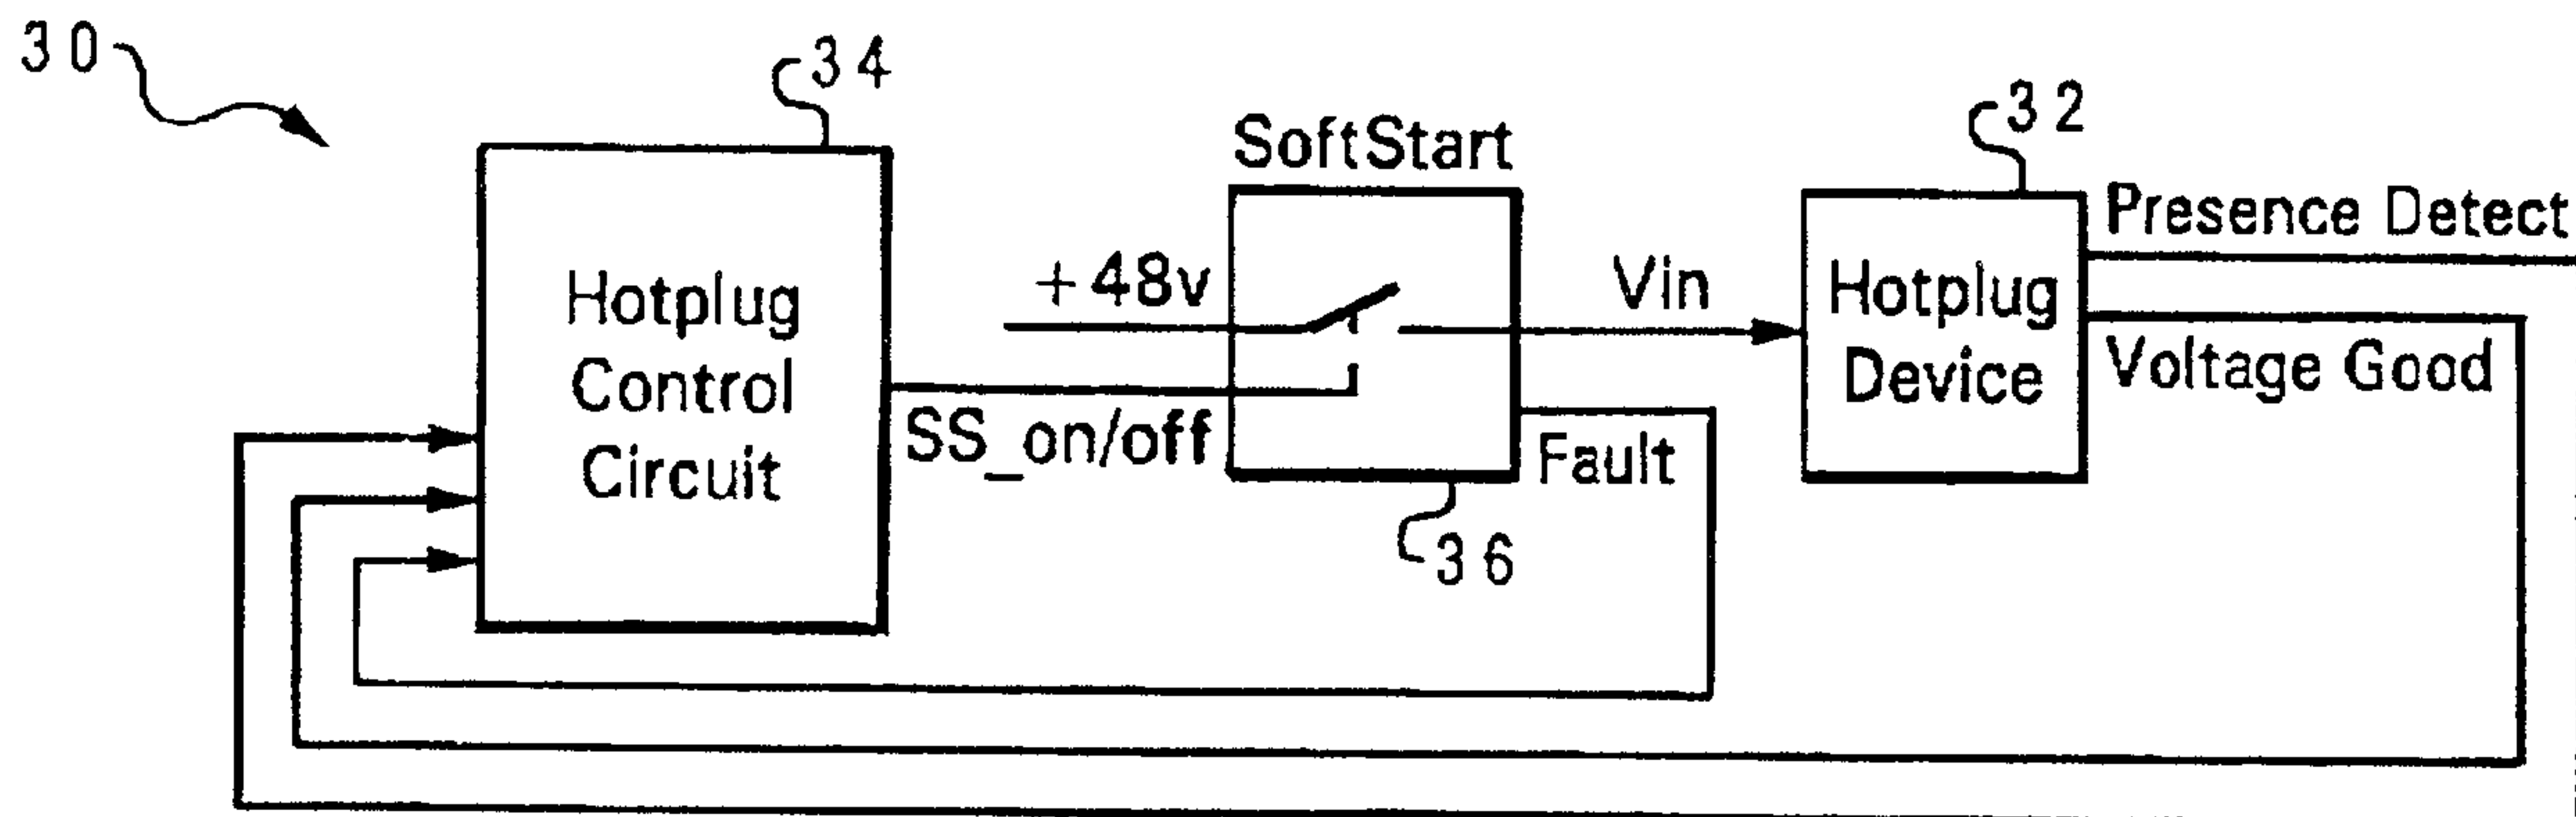




(86) Date de dépôt PCT/PCT Filing Date: 2000/03/24  
 (87) Date publication PCT/PCT Publication Date: 2000/10/05  
 (45) Date de délivrance/Issue Date: 2005/01/11  
 (85) Entrée phase nationale/National Entry: 2001/09/18  
 (86) N° demande PCT/PCT Application No.: GB 2000/001132  
 (87) N° publication PCT/PCT Publication No.: 2000/058846  
 (30) Priorité/Priority: 1999/03/30 (09/281,082) US

(51) Cl.Int.<sup>7</sup>/Int.Cl.<sup>7</sup> G06F 13/40  
 (72) Inventeurs/Inventors:  
 JOHARI, GIRISH CHANDRA, US;  
 THOMSEN, PETER MATTHEW, US;  
 WALTER, LUCINDA MAE, US;  
 MUELLER, MARK WAYNE, US  
 (73) Propriétaire/Owner:  
 INTERNATIONAL BUSINESS MACHINES  
 CORPORATION, US  
 (74) Agent: ROSEN, ARNOLD

(54) Titre : SURVEILLANCE DE BRANCHEMENT A CHAUD DE SYSTEMES INFORMATIQUES A PROCESSEURS MULTIPLES  
 (54) Title: HOT PLUG CONTROL OF MULTIPROCESSOR BASED COMPUTER SYSTEM



(57) **Abrégé/Abstract:**

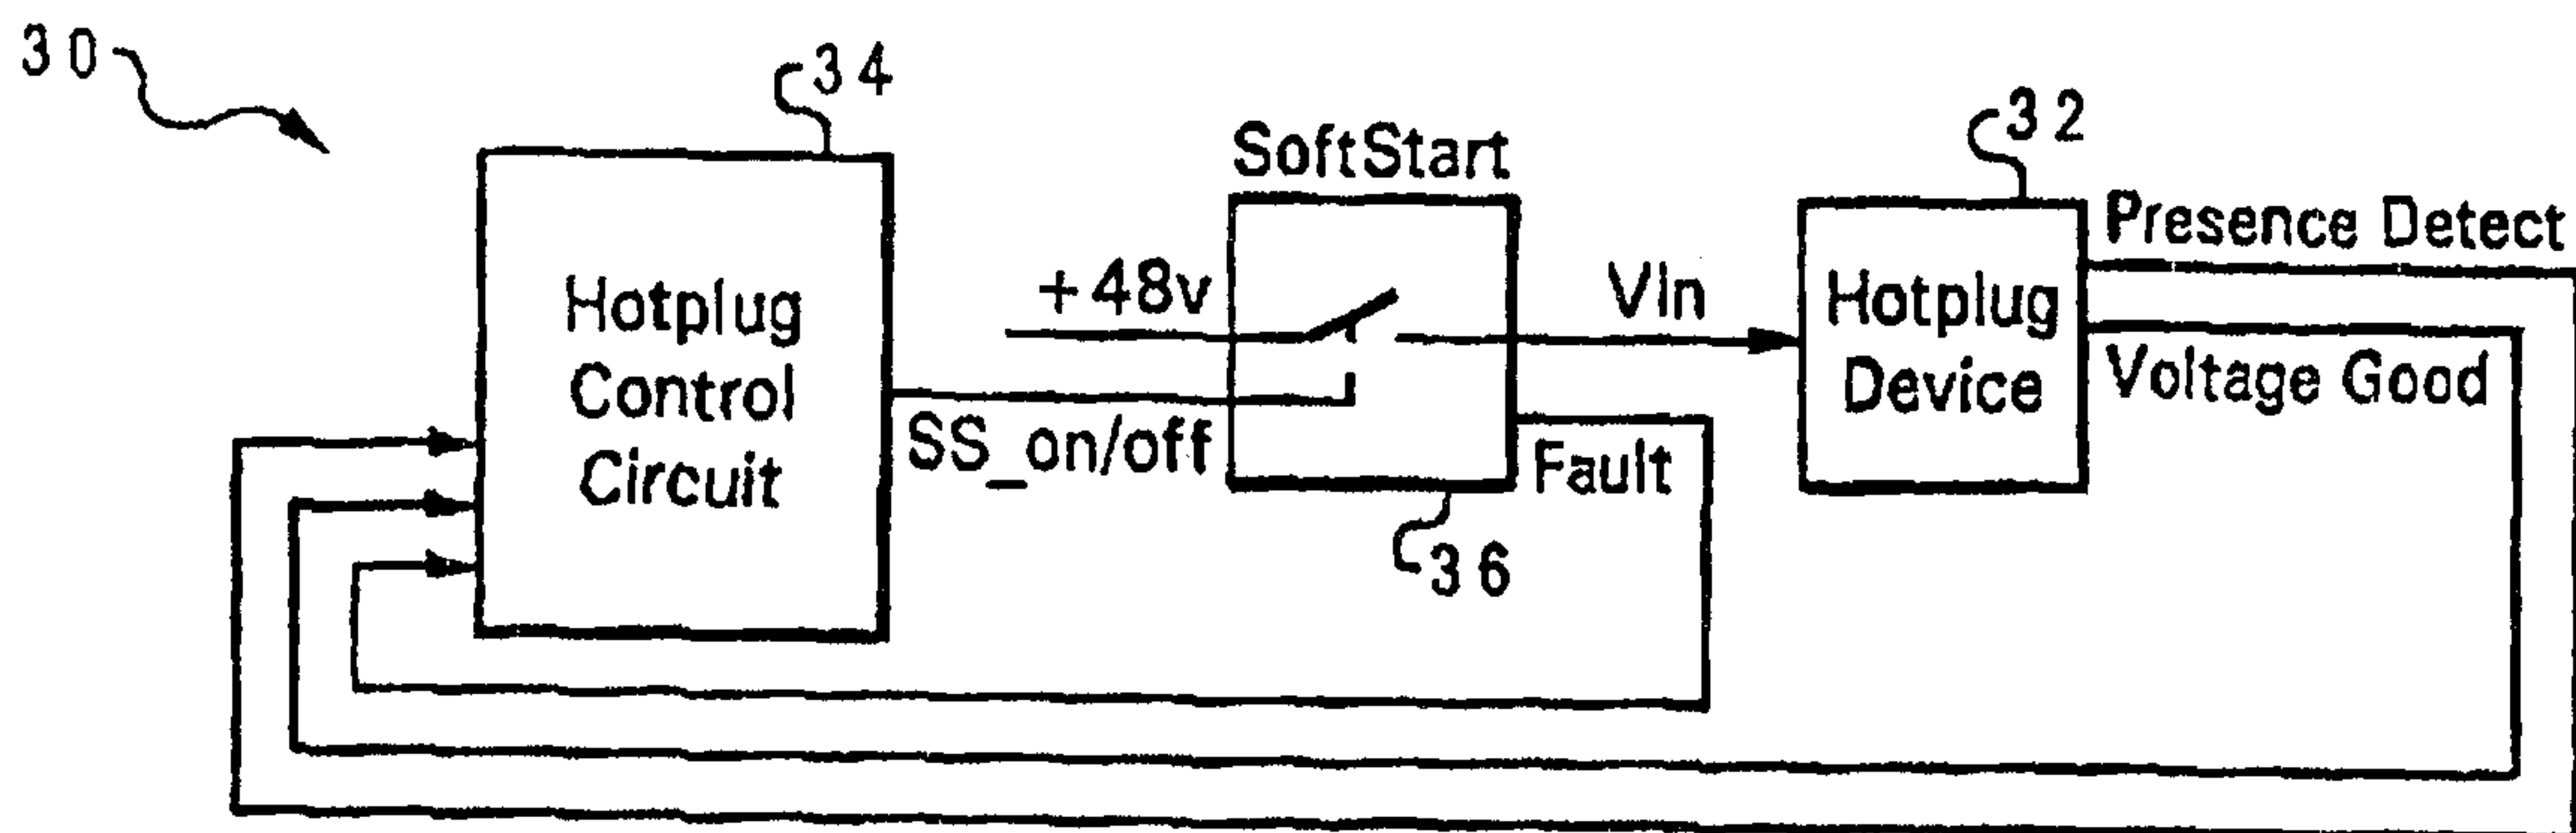
A method of servicing a computer system without interrupting operation of the computer system, by connecting a computer component to a board of the computer system, detecting connection of the computer component to the system board using a control circuit, supplying power to the voltage input of the computer component in response to detecting the connection, and thereafter monitoring the power supplied to the voltage input of the computer component. The method may be used for core computer components such as CPU modules and voltage regulator modules. Power to the voltage input of the computer component is turned off in response to a determination that a current level of the power supplied to the voltage input exceeds a specified level. A fault signal is latched in an active state in response to the determination; the fault signal is reset when the component is removed from the system. The method also applies to a plurality of hot-pluggable components, wherein the power supplied to each component is individually monitored.

**PCT**WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau

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

<p>(51) International Patent Classification <sup>7</sup> : <b>G06F 13/40</b></p>	<p><b>A1</b></p>	<p>(11) International Publication Number: <b>WO 00/58846</b></p> <p>(43) International Publication Date: 5 October 2000 (05.10.00)</p>
<p>(21) International Application Number: PCT/GB00/01132</p> <p>(22) International Filing Date: 24 March 2000 (24.03.00)</p> <p>(30) Priority Data: 09/281,082                      30 March 1999 (30.03.99)                      US</p> <p>(71) Applicant: INTERNATIONAL BUSINESS MACHINES CORPORATION [US/US]; New Orchard Road, Armonk, NY 10504 (US).</p> <p>(71) Applicant (for MC only): IBM UNITED KINGDOM LIMITED [GB/GB]; P.O. Box 41, North Harbour, Portsmouth, Hampshire PO6 3AU (GB).</p> <p>(72) Inventors: JOHARI, Girish, Chandra; 6908 Rimner Cove, Austin, TX 78759 (US). THOMSEN, Peter, Matthew; 12714 Palfrey Drive, Austin, TX 78727 (US). WALTER, Lucinda, Mae; 4600 Secluded Hollow, Austin, TX 78727 (US). MUELLER, Mark, Wayne; 11007 Callanish Park Drive, Austin, TX 78750 (US).</p> <p>(74) Agent: MOSS, Robert, Douglas; IBM United Kingdom Limited, Intellectual Property Department, Hursley Park, Winchester, Hampshire SO21 2JN (GB).</p>		<p>(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p><b>Published</b> <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>

(54) Title: HOT PLUG CONTROL OF MULTIPROCESSOR BASED COMPUTER SYSTEM



## (57) Abstract

A method of servicing a computer system without interrupting operation of the computer system, by connecting a computer component to a board of the computer system, detecting connection of the computer component to the system board using a control circuit, supplying power to the voltage input of the computer component in response to detecting the connection, and thereafter monitoring the power supplied to the voltage input of the computer component. The method may be used for core computer components such as CPU modules and voltage regulator modules. Power to the voltage input of the computer component is turned off in response to a determination that a current level of the power supplied to the voltage input exceeds a specified level. A fault signal is latched in an active state in response to the determination; the fault signal is reset when the component is removed from the system. The method also applies to a plurality of hot-pluggable components, wherein the power supplied to each component is individually monitored.

HOT PLUG CONTROL OF MULTIPROCESSOR BASED COMPUTER SYSTEMField of the Invention

5           The present invention generally relates to computer systems, particularly to a method of upgrading or servicing computer components.

Background of the invention

10           Modern computing systems are often constructed from a number of processing units and a main memory, connected by a generalised interconnect. The basic structure of a conventional multiprocessor computer system 10 is shown in **Figure 1**. Computer system 10 has several processing units (CPUs) 12a, 12b, and 12c which are connected to various peripheral, or input/output (I/O) devices 14 (such as a display monitor, keyboard, and permanent storage device), memory device 16 (random-access  
15           memory or RAM) that is used by the processing units to carry out program instructions, and firmware 18 whose primary purpose is to seek out and load an operating system from one of the peripherals (usually the permanent memory device) whenever the computer is first turned on.

20           Processing units 12a-12c communicate with the peripheral devices, memory and firmware by various means, including a bus 20. Computer system 10 may have many additional components which are not shown, such as serial and parallel ports for connection to, e.g., modems or printers. There are  
25           other components that might be used in conjunction with those shown in the block diagram of **Figure 1**; for example, a display adapter might be used to control a video-display monitor, a memory controller can be used to access memory 16, etc. The computer can also have more than three processing units. In a symmetric multiprocessor (SMP) computer, all of the  
30           processing units 12a-12c are generally identical, that is, they all use a common set or subset of instructions and protocols to operate, and generally have the same architecture.

35           Conventional computer systems often allow the user to add or remove various components after delivery from the factory. For peripheral devices, this can be accomplished using an "expansion" bus, such as the Industry Standard Architecture (ISA) bus or the Peripheral Component Interconnect (PCI) bus. Another component that is commonly added by the user is main memory. This memory is often made up of a plurality of

memory modules that can be added or removed as desired. Even processing units can be added or swapped out, in more recent computer designs.

5 Expansion buses such as the ISA and PCI buses were originally very limited, in that the entire computer system had to be powered down before any peripheral device could be added to or removed from a PCI adaptor slot, and then powered up again (rebooted) to properly initialise the operating system and any new peripheral device. More recently, computer hardware components such as "hot-pluggable" PCI adapters have been devised  
10 that can be added or removed from a computer system while the system is fully operational, without any service interruption. Each PCI adapter slot along the PCI bus has a separate power line, a separate reset line and a switch connecting the slot to the PCI bus, allowing the slot to be electrically isolated from the PCI bus and reactivated after insertion of  
15 a new PCI device into the slot.

This hot-plug capability has never been expanded to core or low-level components such as processors, system memory, or voltage regulator modules (VRMs), which are used to produce the required power  
20 sources/references at precise voltages. While processors and system RAM can be added or swapped out in some conventional systems, these systems must still be powered down for such upgrades or service. Furthermore, components such as VRMs are generally not removable, and any replacement requires field service by a qualified engineer, since the VRM is hardwired  
25 into the system.

Unfortunately, a user may not only want to add another PCI device, but further might want to replace a defective processor, memory bank, or VRM, without service interruption. For many computer systems  
30 (particularly large servers used in a client-server network), there may be hundreds of users connected to it, and the down time required to perform such a service operation can be extremely expensive. Also, in systems which are used in mission-critical applications, it is highly desirable to be able to perform all maintenance or upgrade operations without service  
35 interruption, particularly when it is necessary to replace a defective component.

One problem in providing such hot-pluggable devices relates to control of the voltages and currents involved. Individual control must be  
40 maintained for the power characteristics of each hot-pluggable device, but

presently available power supply designs are incapable of providing such control. Expanded hot-plug capabilities would also necessitate the generation appropriate status signals for other parts of the computer system, e.g., the firmware or operating system which supervises the hot-plug operations, in a manner heretofore not considered. It would, therefore, be desirable to provide a method of controlling the voltage sources to hot-pluggable devices in a computer system, to allow upgrading or servicing of system components without requiring a powering down or interruption of the system. It would be further advantageous if the method could easily handle a large number of hotpluggable devices, and monitor the devices for power faults.

US-A5 875 308 describes a peripheral component interconnect (PCI) architecture for a data processing system including a PCI host bus, a number of PCI local buses and a PCI hot-plug bridge. Each of the local PCI buses has an adapter card slot. The hot-plug bridge, connected between the PCI host bus and the PCI local buses, is utilized for controlling power to each of the PCI local buses, such that a PCI adapter card may be removed from or added to any one of the adapter card slots during power up while there is processing ongoing within adapter cards situated in other adapter card slots.

EP-A-0 772 134 describes a computer system provided with connector slots for receiving feature cards to implement functions such as I/O, memory or the like. A reset control signal issued from an I/O bridge chip is used to initiate the functions of ceasing data processing activity from a card to be removed, decoupling the slot from the bus and causing the power to be gradually decreased. The reset control signal remains active until the original card is removed and the new card is installed in the slot. Once the new card is installed in the connector, power is brought up, the slot is coupled to the bus and the reset signal from the bridge chip is deactivated. An individual slot can be isolated from other slots in the computer system such that particular adapter cards can be changed without the need to power down the entire computer system.

WO 93 15459 describes a method for the insertion or removal of a circuit device into a slot in a backplane of a computer system having powered circuit devices interconnected by a communication bus. The method employs a slot controller to notify the system that a circuit device is to be inserted, identifying the location for the circuit device in the system and inserting the circuit device into the slot in the backplane while the system remains powered. A detect signal is provided by isolating a ground

pin on the backplane and attaching a pull-up resistor such that when a circuit device is inserted the pin will be grounded. Once a circuit has been detected, the slot controller will arbitrate for the bus, wait for the existing bus traffic to subside and then power up and reset the new board.

#### Disclosure of the Invention

Accordingly, the present invention provides a method of servicing a computer system without interrupting operation of the computer system, generally comprising the steps of connecting at least one computer component to a board of the computer system, the computer component having a voltage input, detecting connection of the computer component to the system board using a control circuit of the computer system, supplying power to the voltage input of the computer component in response to said detecting step, and thereafter monitoring the power supplied to the voltage input of the computer component, characterized in that the method includes connecting at least one hotplugged, lower-level non-peripheral component of the computer system into a connector on a system board of the computer system, enabling an individual soft start circuit to smoothly turn on supply voltage to the or each component connected to the system board and to individually monitor the supplied power, and employing a programmable gate array including soft start control logic adapted to receive presence detect signals from the soft start circuits indicating the presence of components connected to the system board and to respond to the presence detect signals by supplying on/off signals to enable the individual soft start circuits, and fault control logic to respond to fault signals from the soft start circuits to supply reset signals to the soft start circuits. The invention may further include the step of turning off power to the voltage input of the computer component in response to a determination that a current level of the power supplied to the voltage input of the computer component exceeds a specified level. A fault signal is latched in an active state in response to the determination; the fault signal is reset when the component is removed from the system.

The method also applies to a plurality of hot-pluggable components, wherein the power supplied to each component is individually monitored. The control circuit can sequence power to the components in any desired order. A plurality of voltage good signals from the computer components are consolidated in the control circuit, and the control circuit generates

4a

a system power good signal based on the plurality of voltage good signals from the computer components.

Viewing the present invention from a second aspect, there is provided a power subsystem for a computer system, comprising a circuit board having at least one connector for receiving a component of the computer system, said computer component having a voltage input, means for detecting connection of the computer component to said circuit board, means for supplying power to the voltage input of the computer component in response to detection of the connection and means for monitoring the power supplied to the voltage input of the computer component characterized in that the subsystem includes; connectors mounted on a system board to receive hotplugged, lower-level non-peripherals components of the computer system, individual soft start circuits adapted to be enabled to smoothly turn on supply voltage to components connected to the system board connectors and to individually monitor the supplied power; and a programmable gate array, the programmable gate array including; soft start control logic having a plurality of presence detect inputs adapted to receive presence detect signals from the soft start circuits indicating the presence of components connected to the system board connectors and to respond to the presence detect signals by supplying on/off signals to enable the individual soft start circuits, and fault control logic adapted to receive fault signals from the soft start circuits and to respond to removal of computer components by supplying reset signals to the soft start circuits.

It is an advantage of the present invention to provide an improved method of upgrading and servicing components of a computer system.

It is another advantage of the present invention to provide such a method that allows a wide variety of computer components to be upgraded or serviced, without interrupting system operation.

Furthermore, it is yet another advantage of the present invention to provide such a method that carefully controls and monitors the power supplied to any such hot-pluggable devices, individually.

AUS980836

5

Brief Description of the Drawings

The present invention will now be described, by way of example only, with reference to preferred embodiments thereof as illustrated in the following drawings:

5

**Figure 1** is a block diagram of a prior-art multiprocessor computer system;

10

**Figure 2** is block diagram of a power subsystem for a computer system, illustrating the control and monitoring of one of a plurality of hot-pluggable devices used by the computer system, in accordance with one embodiment of the present invention;

15

**Figure 3** is a pictorial representation of one implementation of a field programmable gate array used with the hot-plug control circuit of **Figure 2**; and

20

**Figure 4** is a schematic diagram of a soft start circuit used to supply power to a hot-pluggable device, in accordance with one embodiment of the present invention.

Description of Embodiments of the invention

25

With reference now to the figures, and in particular with reference to **Figure 2**, there is depicted one embodiment of a power subsystem **30** constructed in accordance with the present invention, for a computer system having a plurality of hot-pluggable devices. **Figure 2** depicts only one such hot-pluggable device **32**, but it is understood that the following description applies to any number of hot-pluggable devices that may be provided by the overall computer architecture.

30

35

While the present invention may be applied to hot-plug peripheral devices, it is particularly adapted for use with non-peripheral components, such as the central processing units (CPUs or processors), or even lower-level components like a voltage regulator module (VRM). These components may be rendered hot-pluggable as described in U.S. Patent Nos. 6,378,027 and 6,449,676. The CPUs and VRMs may be added or removed using connectors mounted on a system board.

In this embodiment, power subsystem 30 includes a single hotplug control circuit 34, and individual soft start circuits 36 (one for each hot-pluggable device 32). Prior to a device being connected, soft start circuit 36 is off, resulting in no input voltage ( $V_{in}$ ) to device 32. When the device is placed in a corresponding slot or socket, the "Presence Detect" signal output from device 32 becomes active. The Presence Detect line has a pull-down resistor to ground on the hotpluggable device. The signal floats when the device is not present, and is grounded when the device is present. Once hotplug control circuit 34 detects the presence of device 32, it enables soft start circuit 36, which turns on  $V_{in}$  to hotplug device 32. Soft start circuit 36 is described further below.

As noted, only one hotplug control circuit 34 is provided in this embodiment, however it is adapted to handle multiple hot-pluggable devices (17 devices in the example discussed in conjunction with **Figure 3**). Hotplug control circuit 34 can sequence the voltage to the devices in any desired (pre-defined) order. Hotplug control circuit 34 also monitors soft start circuit 36 for faults via a "Fault" signal. If a fault is detected, hotplug control circuit 34 shuts off soft start circuit 36.

In the illustrative embodiment, hotplug control circuit 34 is implemented with a field programmable gate array (FPGA). **Figure 3** shows a detailed plan of a hotplug control FPGA 38 configured in accordance with the present invention.

Hotplug control FPGA 38 is adapted for use with hot-pluggable VRMs, hot-pluggable CPU modules, etc. A plurality of Presence Detect signals are input into the soft start control logic 35, which has as its outputs a plurality of respective soft start on/off lines. A plurality of fault signals are similarly input into the fault control logic 37, which has as its outputs a plurality of respective reset lines. The voltage good signals from each hot-pluggable VRM or CPU quad are respectively consolidated in power good control logic 39, which generates the Power Good signals for the rest of the system.

**Figure 4** depicts one embodiment of soft start circuit 36. The SOFT\_START\_ON/OFF signal is an LVTTTL (low-voltage transistor-transistor logic) level signal that enables a power MOSFET 40 to smoothly bring the HOTPLUG\_INPUT\_VOLTAGE signal up to the input voltage supplied to the circuit. In this example, the input voltage is 48 volts, provided by a

power supply (not shown) connected to an external power source, e.g., a 110 volt AC wall outlet.

5 A logic circuit 42 (Unitrode part# UCC3917) provides a fault output 44 if the voltage across the current sense resistor 41 exceeds a specified level (e.g., 50 mV). Several comparators 46, 48 and 50 latch the fault signal, and keep it high (active) until the SOFT\_START\_RESET signal allows it to be reset. As long as the fault signal is active, hotplug control circuit 34 maintains the SOFT\_START\_ON/OFF signal at the low level to keep 10 the power HOTPLUG\_INPUT\_VOLTAGE 43 turned off. The reset signal can be activated upon, e.g., removal of the device which is also detected via the Presence Detect signal.

15 Pins C1P and C1N of the Unitrode part are connected to an upper charge pump capacitor, while pins C2P and C2N are connected to a lower charge pump capacitor. Pin OUTPUT is the output to the NMOS pass element, and pin SENSE is the sense voltage input from sense resistor 41. The capacitor value on pin CT determines the maximum fault time before 20 retrying; this retry feature is disabled in the illustrated embodiment by the SOFT\_START\_RESET circuitry. The resistance on pin MAXI determines the maximum allowable sourcing current. Pin FLTOUT# is used for fault output indication. The reference signals include pin V<sub>SS</sub> (the negative reference for the device), pin V<sub>OUT</sub> (the ground reference for the chip, which is floating with respect to system ground), and V<sub>REF</sub>/CATFLT# (the output 25 reference for programming MAXI, and catastrophic fault output).

Thus the present invention provides an effective method of individually controlling voltage sources to hot-pluggable devices. It is possible (and convenient) to use components such as CPU modules and VRMs 30 as hot-pluggable devices. The present invention is also scalable to practically any number of hot-pluggable devices since the FPGA is easily modified.

Although the invention has been described with reference to specific 35 embodiments, this description is not meant to be construed in a limiting sense. Various modifications of the disclosed embodiments, as well as alternative embodiments of the invention, will become apparent to persons skilled in the art upon reference to the description of the invention. It is therefore contemplated that such modifications can be made without

departing from the scope of the present invention as defined in the appended claims.

CLAIMS

1. A method of servicing a computer system without interrupting operation of the computer system, comprising the steps of:

connecting at least one computer component (32) to a board of the computer system, the computer component having a voltage input;

detecting connection of the computer component (32) to the board using a control circuit (34) of the computer system;

supplying power to the voltage input (Vin) of the computer component (32) in response to said detecting step; and

monitoring the power supplied to the voltage input (Vin) of the computer component (32);

characterized in that the method includes;

connecting at least one hotplugged, lower-level non-peripheral component (32) of the computer system into a connector on a system board of the computer system,

enabling an individual soft start circuit (36) to smoothly turn on supply voltage to the or each component connected to the system board and to individually monitor the supplied power; and

employing a programmable gate array (38) including soft start control logic (35) adapted to receive presence detect signals from the soft start circuits (36) indicating the presence of components connected to the system board and to respond to the presence detect signals by supplying on/off signals to enable the individual soft start circuits, and fault control logic (37) to respond to fault signals from the soft start circuits to supply reset signals to the soft start circuits.

2. A method as claimed in Claim 1 further comprising the steps of:

determining that a current level of the power supplied to the voltage input (Vin) of a connected computer component (32) exceeds a specified level; and

turning off power to the voltage input of the computer component (32) in response to said determining step.

3. A method as claimed in Claim 2 further comprising the steps of:

latching a fault signal in an active state in response to said determining step;

removing the computer component (32) from the system board;

detecting removal of the computer component (32); and

resetting the fault signal in response to said step of detecting removal of the computer component (32).

4. A method as claimed in any preceding claim wherein said connecting step connects a CPU module to the system board.

5. A method as claimed in any one of claims 1 to 4 wherein said connecting step connects a voltage regulator module to the system board.

6. A method as claimed in Claim 1 wherein the control circuit (34) sequences power to a plurality of components (32) in a pre-defined order.

7. A method as claimed in Claim 6 further comprising the steps of:

consolidating a plurality of voltage good signals from the computer components, in the control circuit (34); and

generating a system power good signal using the control circuit based on the plurality of voltage good signals from the computer components (32).

8. A power subsystem for a computer system, comprising:

a circuit board having at least one connector for receiving a component (32) of the computer system, said computer component (32) having a voltage input (Vin);

a control circuit (34) for detecting connection of the computer component (32) to said circuit board;

means for supplying power to the voltage input (Vin) of the computer component (32) in response to detection of the connection; and

means for monitoring the power supplied to the voltage input (Vin) of the computer component (32);

characterized in that the subsystem includes;

connectors mounted on a system board to receive hotplugged, lower-level non-peripheral components (32) of the computer system,

individual soft start circuits (36) adapted to be enabled to smoothly turn on supply voltage to components (32) connected to the system board connectors and to individually monitor the supplied power; and

a programmable gate array (38), the programmable gate array (38) including;

soft start control logic (35) having a plurality of presence detect inputs adapted to receive presence detect signals from the soft start circuits (36) indicating the presence of components (32) connected to the system board connectors and to respond to the presence detect signals by supplying on/off signals to enable the individual soft start circuits (36), and

fault control logic (37) adapted to receive fault signals from the soft start circuits and to respond to removal of computer components by supplying reset signals to the soft start circuits (36).

9. A power subsystem as claimed in Claim 8 further comprising means for turning off power to the voltage input (Vin) of the computer component (32) in response to a determination that a current level of the power supplied to the voltage input of the computer component (32) exceeds a specified level.

10. A power subsystem as claimed in Claim 9 further comprising:

means (46,48,50) for latching a fault signal in an active state in response to the determination that the current level exceeds a specified level.

11. A power subsystem as claimed in Claim 10 wherein said supplying means sequences power to the components (32) in a pre-defined order.

12. A power subsystem as claimed in Claim 11 further comprising:

means for consolidating a plurality of voltage good signals from the computer components; and

means (39) for generating a system power good signal based on the plurality of voltage good signals from the computer components.




Fig. 1  
Prior Art




Fig. 2




Fig. 3

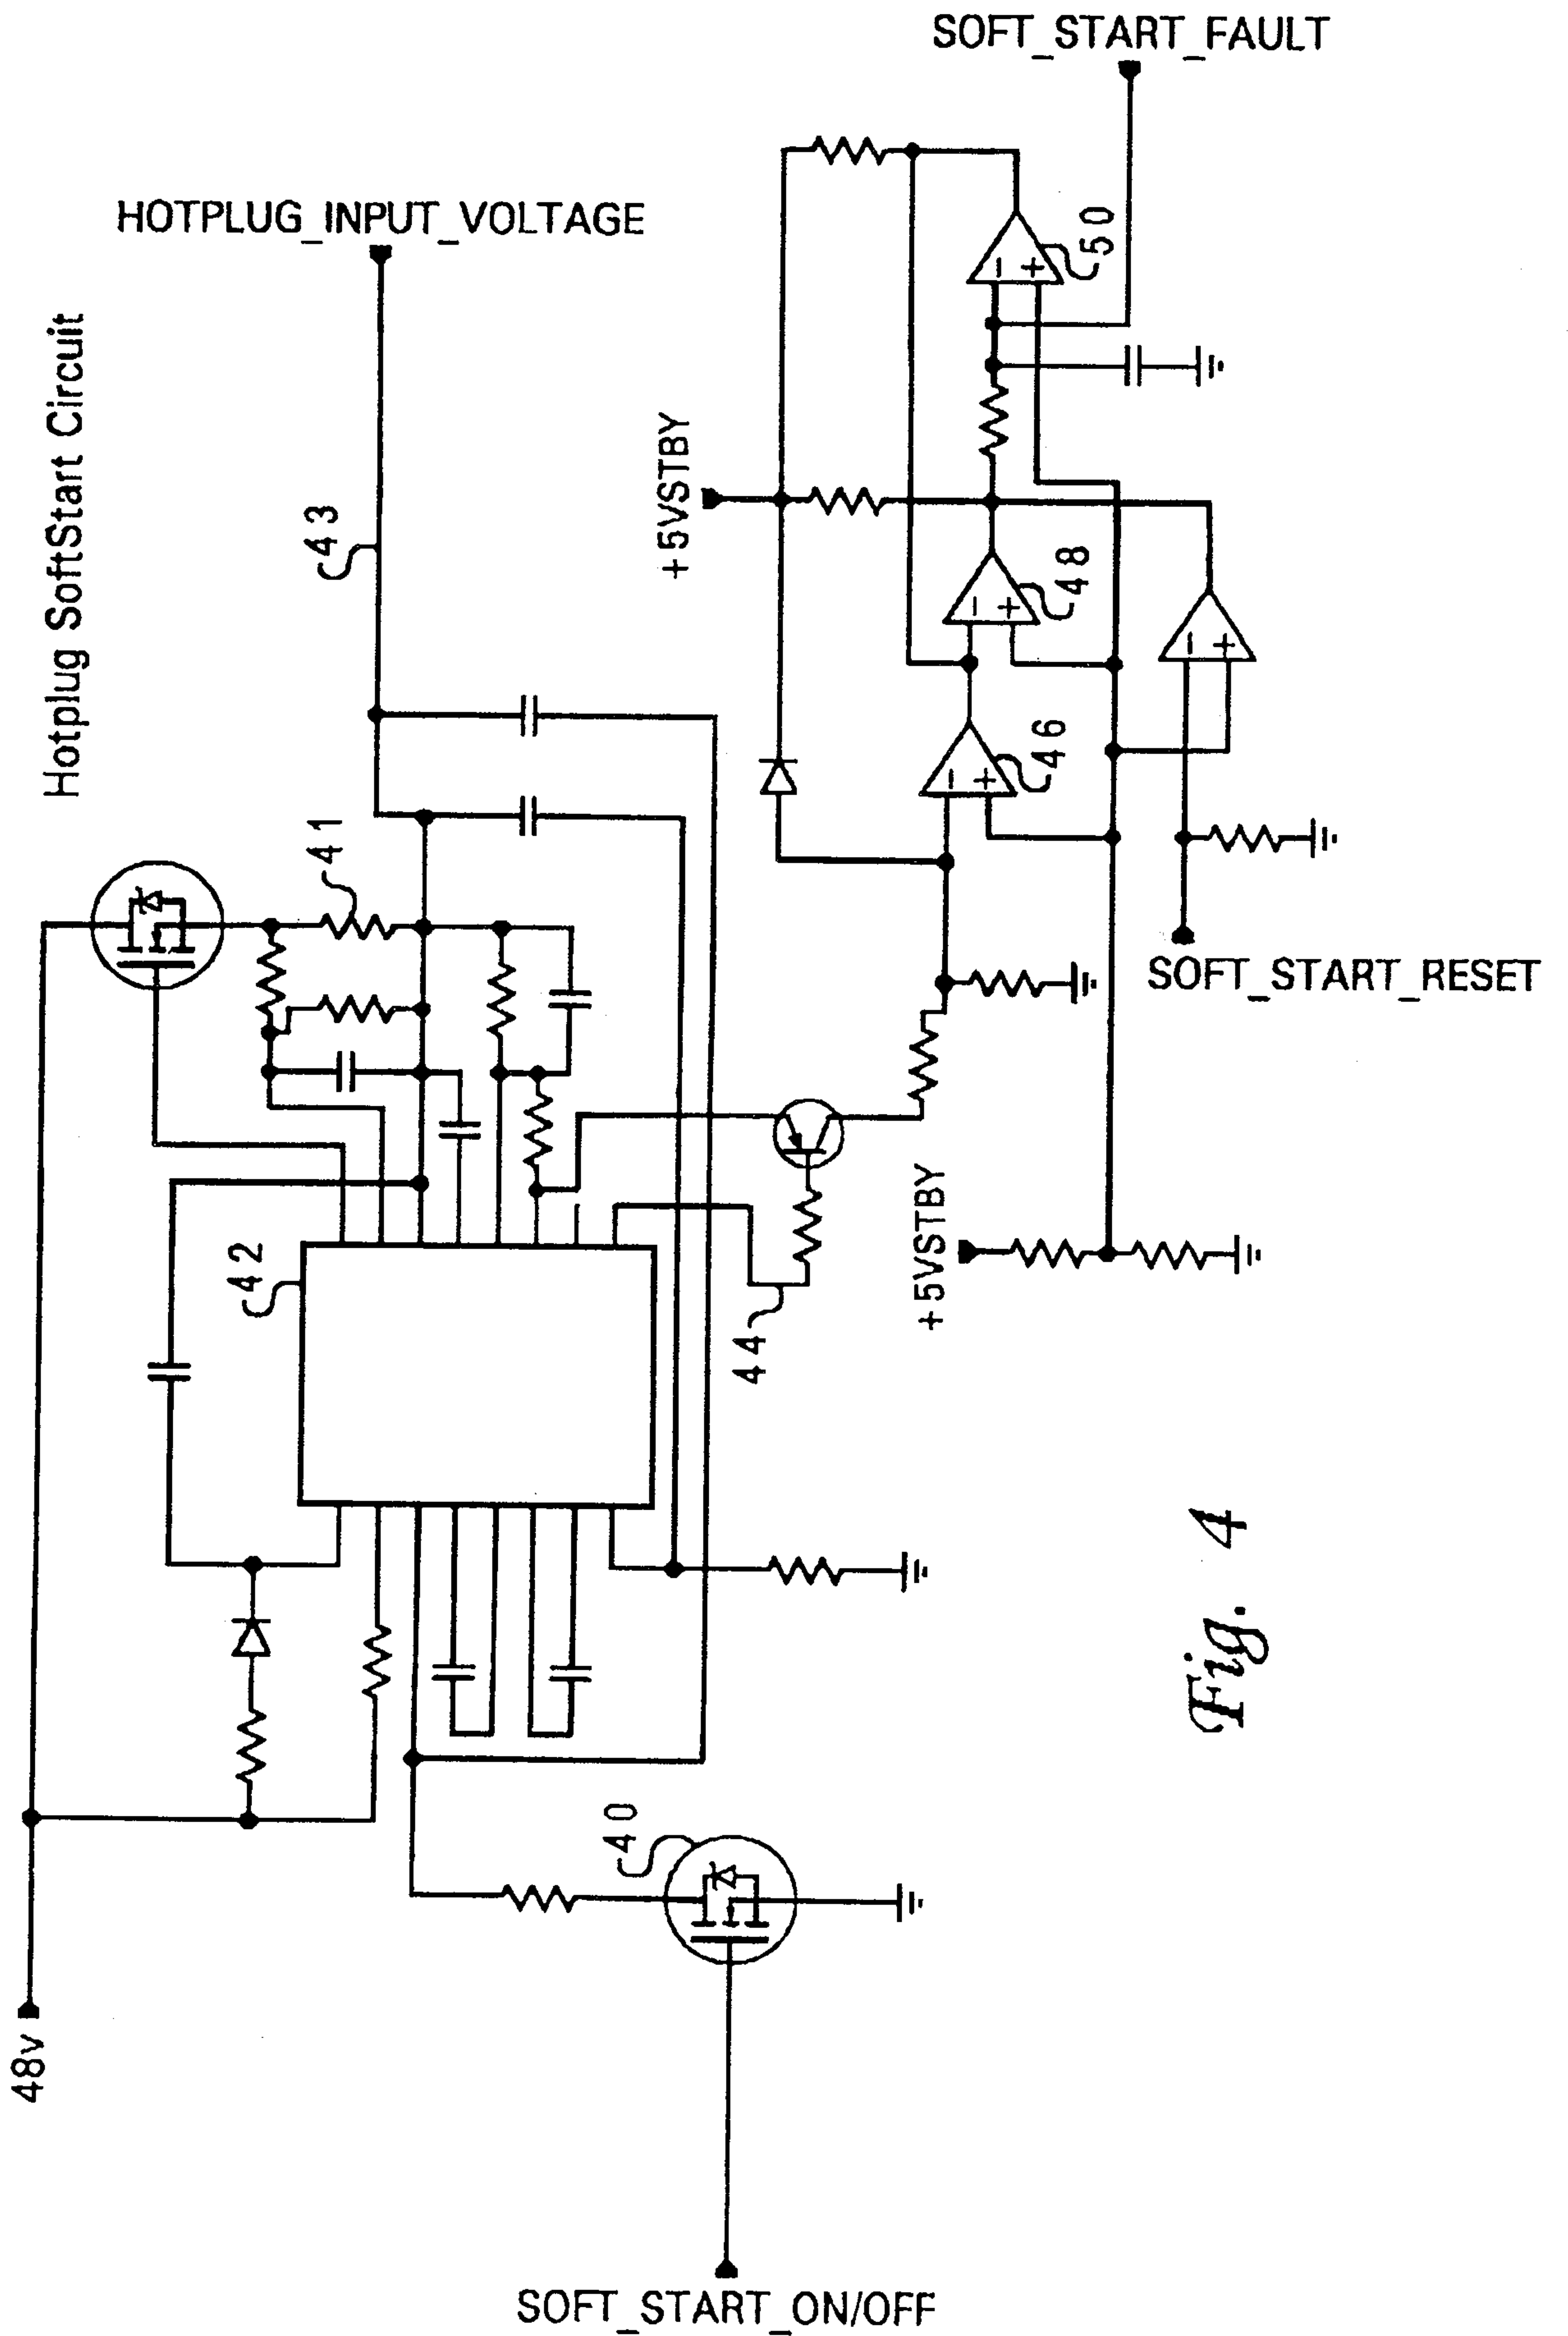


Fig. 4

