

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 October 2010 (21.10.2010)

(10) International Publication Number
WO 2010/119093 A1

- (51) International Patent Classification:
G06F 1/26 (2006.01) *H02J 1/00* (2006.01)
- (21) International Application Number:
PCT/EP2010/054951
- (22) International Filing Date:
15 April 2010 (15.04.2010)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
12/425,275 16 April 2009 (16.04.2009) US
12/425,267 16 April 2009 (16.04.2009) US
- (71) Applicant (for all designated States except US): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).
- (71) Applicant (for MG only): **IBM UNITED KINGDOM LIMITED** [GB/GB]; PO Box 41, North Harbour, Portsmouth Hampshire PO6 3AU (GB).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **ARIMILLI, Ravi, Kumar** [US/US]; IBM Corporation, 11400 Burnet RD, Austin, Texas 78758-3493 (US). **SEMINARO, Edward, Joseph** [US/US]; IBM Corporation, 2455 South RD,

Poughkeepsie, New York 12601-5400 (US). **COVI, Kevin, Robert** [US/US]; IBM Corporation, 2455 South RD, Poughkeepsie, New York 12601-5400 (US). **AHLADAS, Steven, John** [US/US]; IBM Corporation, 2455 South RD, Poughkeepsie, New York 12601-5400 (US). **FAHR, Gerald** [US/US]; IBM Corporation, 2455 South RD, Poughkeepsie, New York 12601-5400 (US). **BARUS, Daniel** [US/US]; IBM Corporation, 2455 South RD, Poughkeepsie, New York 12601-5400 (US).

(74) Agent: **LITHERLAND, David, Peter**; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester Hampshire SO21 2JN (GB).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG,

[Continued on next page]

(54) Title: POWER CONVERSION, CONTROL, AND DISTRIBUTION SYSTEM

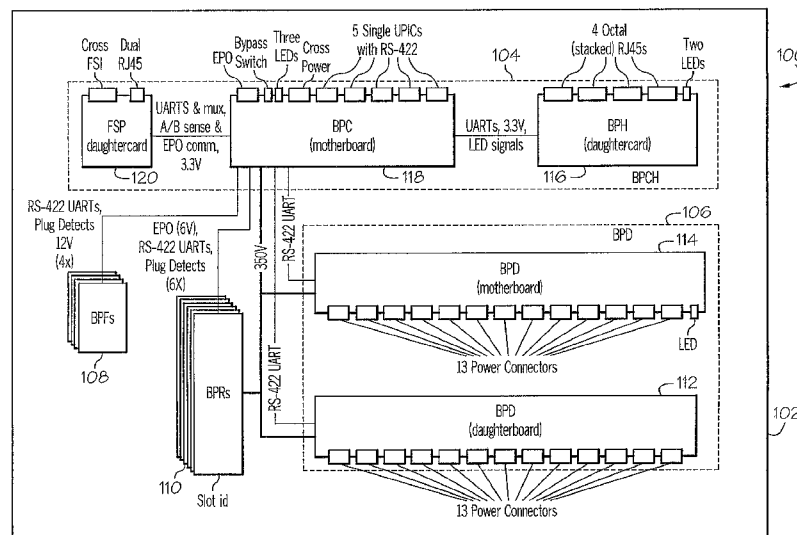


FIG. 1

(57) Abstract: A power conversion, control, and distribution system includes multiple bulk power regulator (BPR) subassemblies, a bulk power distribution (BPD) subassembly, and a bulk power controller and hub (BPCH) subassembly. The BPR subassemblies are each configured to provide regulated DC power from both AC input power and DC input power. The BPD subassembly is configured to distribute the regulated DC power. The BPCH subassembly is coupled to the multiple BPR subassemblies and the BPD subassembly. The BPCH subassembly is configured to monitor and control the BPR assemblies and the BPD assembly.

WO 2010/119093 A1

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, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG). **Published:** — *with international search report (Art. 21(3))*

POWER CONVERSION, CONTROL, AND DISTRIBUTION SYSTEM

Field of the Invention

5 This disclosure relates generally to a power system and, more specifically to a power conversion, control, and distribution system for providing bulk power to a node (or nodes) of a high performance computing cluster or other system requiring high-availability.

Background of the Invention

10

Conventional high-availability systems (e.g., servers) have employed power supply redundancy in an attempt to prevent outages of the systems in the event a power supply fails. Unfortunately, power supply redundancy alone is frequently not sufficient to achieve one-hundred percent availability for a system.

15

Summary of the Invention

According to one aspect of the present disclosure, a power conversion, control, and distribution system includes a power conversion, control, and distribution system,
20 comprising: multiple bulk power regulator (BPR) subassemblies each configured to provide regulated DC power from both AC input power and DC input power; a bulk power distribution (BPD) subassembly configured to the regulated DC power; and a bulk power controller and hub (BPCH) subassembly coupled to the multiple BPR subassemblies and the BPD subassembly, wherein the BPCH subassembly is configured to monitor and control the
25 BPR assemblies and the BPD assembly.

According to another aspect of the invention there is provided a bulk power assembly,
comprising: a bulk power distribution (BPD) subassembly that includes multiple static
30 circuit breakers that are configured to distribute DC bulk power and provide relay and fault protection; and a bulk power controller and hub (BPCH) subassembly coupled to the BPD subassembly, wherein the BPCH subassembly is configured to monitor and control the BPD subassembly, and wherein the BPCH subassembly includes a bulk power control (BPC)

motherboard that includes multiple static circuit breakers that function to route the DC bulk power and a bulk power hub (BPH) daughtercard that is isolated in a safety extra low voltage (SELV) block and is configured to provide Ethernet communication functions.

5 **Brief Description of the Drawings**

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

10 FIG. 1 is a diagram of a relevant portion of a bulk power assembly (BPA) configured according to various aspects of the present disclosure.

FIG. 2 is a diagram of a relevant portion of an example bulk power controller and hub (BPCH) subassembly that may be employed in the BPA of FIG. 1, according to various
15 embodiments of the present disclosure.

FIG. 3 is a diagram of a relevant portion of an example bulk power hub (BPH) daughtercard that may be employed in the BPCH subassembly of FIG. 2, according to various
20 embodiments of the present disclosure.

FIG. 4 is a diagram of a relevant portion of another example BPCH subassembly that may be employed in the BPA of FIG. 1, according to various embodiments of the present disclosure.

FIG. 5 is a diagram of a relevant portion of an example bulk power distribution (BPD) subassembly that may be employed in the BPA of FIG. 1, according to various embodiments
25 of the present disclosure.

FIG. 6 is a diagram of a relevant portion of another example BPD subassembly that may be employed in the BPA of FIG. 1, according to various embodiments of the present disclosure.
30

FIG. 7 is a diagram of a relevant portion of an example bulk power regulator (BPR) subassembly that may be employed in the BPA of FIG. 1, according to various embodiments of the present disclosure.

5 FIG. 8 is a diagram of a relevant portion of an example rectifier front-end that may be employed in the BPR of FIG. 7, according to various embodiments of the present disclosure.

Detailed Description

10 As will be appreciated by one of ordinary skill in the art, various portions of the present invention may be embodied as a method, system, device, or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, microcode, etc.) or an embodiment combining software and hardware aspects that may all
15 generally be referred to herein as a circuit, module, or system. The present invention may, for example, take the form of a computer program product on a computer-usable storage medium having computer-usable program code, e.g., in the form of one or more design files, embodied in the medium.

20 Any suitable computer-usable or computer-readable storage medium may be utilized. The computer-usable or computer-readable storage medium may be, for example, but is not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device. More specific examples (a non-exhaustive list) of the computer-readable storage medium includes: a portable computer diskette, a hard disk, a
25 random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM) or flash memory, a portable compact disc read-only memory (CD-ROM), an optical storage device, or a magnetic storage device. As used herein, the term "coupled" includes a direct electrical connection between elements or blocks and an indirect electrical connection between elements or blocks achieved using one or more
30 intervening elements or blocks.

As used herein the term bulk power enclosure (BPE) refers to a sheet metal enclosure designed to hold modular components included in a bulk power assembly (BPA) or power conversion, control, and distribution system. In one or more embodiments, a BPE may be designed to accommodate the following modular subassemblies: one to six hot-swappable
5 bulk power regulator (BPR) subassemblies that each generate -350 VDC from either AC input power or DC input power; a bulk power controller and hub (BPCH) subassembly that performs selected control/monitoring functions, distributes -350VDC via electronic circuit breakers, provides power control via a 32-port RS-422 multiplexer (MUX), and acts as a 34-port Ethernet switch for system service; a bulk power distribution (BPD) subassembly that is
10 configured to distribute -350 VDC to external power devices via electronic circuit breakers; and four bulk power fan (BPF) subassemblies that remove heat load from air dissipated in the various subassemblies.

To support the subassemblies, the BPE may incorporate: a water manifold designed to
15 accept system water inlet and outlet connections and distribute water as required to the subassemblies included in the BPE; an electromagnetic interference (EMI) filter for each power input to ensure line conducted noise of the BPA meets requirements; a pair of power inputs that attach to input line cords via detachable connectors; a backplane that interconnects signals and provides power for the subassemblies; guides and fasteners for
20 mating all subassemblies into the BPE (and the BPE to a frame in a system application); and mechanical features to prevent air leakage during periods when BPR subassemblies are hot-swapped.

With reference to FIG. 1, an example BPA 100 is illustrated that includes a number of
25 modular subassemblies within a BPE 102. The BPA 100 may, for example, be configured to provide up to 250KW of power in a single rack. In this embodiment, electrical connections between a BPCH subassembly 104 and a BPD subassembly 106 and between the BPCH subassembly 104 and BPR subassemblies 110 are made via a backplane of the BPE 102. Electrical connections to the BPF subassemblies 108 are made from the BPCH subassembly
30 104 to the backplane of the BPE 102 and then from the backplane to each of the BPF subassemblies 108 through a cable assembly. In one embodiment, each pair of input power connectors (not shown in FIG. 1) supply input power to three BPR subassemblies 110. It

should be noted that all output power and signals from the BPE 102 occur on a tailstock of either the BPCH subassembly 104 (signals & some power) or the BPD subassembly 106 (most of the power).

5 In a disclosed embodiment, nominal AC input voltages for the BPA 100 may include 200V, 208V, 220V, 240V, 380V, 400V, 415V, or 480V root-mean square (RMS). Steady-state AC input voltages may range from 176V to 508V. In various embodiments, the BPA 100 is configured for four-wire balanced three-phase system AC inputs. The BPA 100 is configured to connect to WYE (i.e., earthed, but unused neutral) and DELTA (i.e., earthed
10 corner or center tapped phase) sources. Steady-state DC input voltages may range from 231VDC to 600VDC. DC inputs to the BPR subassemblies 110 are two potential (+ and -) lines plus earth. In various embodiments, the + line of the DC input is referenced to earth. In order to connect the two potential DC input conductors to a three-phase AC input, one of the DC potentials is connected to two of the phase inputs. In this case, the polarity of the
15 input voltage is arbitrary (i.e., + and - may be connected to any of the phase inputs). A line filter may be implemented for each of the two power cords in the BPA 100 to ensure a fully populated BPE 102 meets applicable line conducted noise limitations.

BPR subassembly location pins may be provided to allow a BPR subassembly host
20 microcontroller to identify what slot of the BPE 102 a BPR subassembly 110 is plugged into. A '0' on any of the pins is asserted by connecting the pin to the ground reference of the BPR subassembly host microcontroller (i.e., the plus rail of the BPR subassembly output). A '1' on the pin is asserted by leaving the pin open. Each of the two power input connectors may have an associated single LED that indicates sufficient input voltage is available to enable
25 the BPR subassemblies 110 to start. In one embodiment, the single LED is illuminated if any one of the line-to-line input voltages is greater than 180VAC or 231VDC (i.e., the steady-state input operating range of the BPR subassemblies in this example). In one embodiment, the single LED is configured to be off if all of the line-to-line input voltages are less than 150VAC or 185VDC (the low line fault range for the BPR subassemblies in
30 this example).

While the majority of heat dissipated in the subassemblies included in the BPE may be dissipated to water, forced air may be required to remove the remaining heat. In one embodiment, the forced air may be provided by four BPF subassemblies (speed controlled fan modules) located within the BPE. Each fan module may include two counter rotating fans. In general, the BPF subassemblies are configured to provide sufficient air flow to cool all subassemblies under all non-fault operating conditions. Under the condition that one of the fans has failed, the implemented fans should have sufficient capacity to cool all the BPR subassemblies at up to 75% of their rated output. In general, the electrical and mechanical design of the fans permits replacement of a failed fan while the BPA 100 is in service. In various embodiments, the design of the BPE 102 includes appropriate hardware to prevent air leakage and allow recirculation under conditions where one of the BPF subassemblies 108 has failed or has been removed from the BPE 102. Each of the BPF subassemblies 108 includes a microcontroller that has a serial interface that may be used to control fan speed and monitor the health of a respective BPF subassembly.

For example, a BPF microcontroller may include the follow connections: Tach; TachSel; PWMSpeed; Vfan; FanGood#; ID#; Xmit; and Rcv. In this case, Tach provides a pulse train from a fan output tachometer. The pulse train may have a duty cycle between 40 and 60% for non-fault operating conditions. The fan speed may, for example, be calculated as 15 RPMs for every rising edge per second. Fan operating speeds may, for example, range from 0 to 14,000 RPMs. TachSel is used to select which of two fan tachometers is connected to the tach input. PWMSpeed (which may be a 28.8KHz pulse width modulated (PWM) waveform) is used to control the speed of both attached fans. For example, at a PWM duty cycle of '0' the fan speed is zero and at a duty cycle of 100% the fan speed is at a maximum. Vfan may correspond to the fan voltage divided by $4.00 \pm 2\%$. The Vfan signal may be, for example, passed through and low-pass filtered with a single pole at 600Hz. When FanGood# is low, an LED may be lit to indicate a BPF subassembly is operational. When the ID# is low another LED may be driven to indicate when the BPF subassembly is operational in the BPE. The Xmit and Rcv pins may be implemented to form a full-duplex UART interface connection with a bulk power controller (BPC) motherboard 118. The interface may, for example, translate to an RS-422 protocol for connection to the BPC subassembly 104.

Each BPF subassembly may provide a plug detect signal that indicates the fan is fully seated. The plug detect signal may be connected to the ground reference on the fan and run as a separate signal to the BPC motherboard 118. In at least one embodiment, when the BPF subassembly is not plugged, the BPC motherboard 118 pulls the plug detect signal to a high (i.e., not plugged) state. The BPA may be designed to operate at maximum load with an inlet temperature entering the BPA between 15°C and 40°C at an altitude of 0-10,000 ft. In at least one embodiment, the range for relative humidity is 20-80% non-condensing.

In one or more embodiments, none of the components that are included in the BPE 102 proper are water cooled. However, in this embodiment, the BPR subassemblies 110 and BPD subassembly 106 are water cooled. The BPE 102 includes the water manifold and connections to the BPR and BPD subassemblies. Cooling water temperature may range from about 13°C to 24°C. In one embodiment, the water flow delivered to a fully configured BPE (e.g., including six of the BPR subassemblies 110 and one BPD subassembly 106) is about 3.5 gallons per minute at a pressure drop of 10 pounds per square inch (PSI). In one or more embodiments, deionized cooling water with 1000 parts per million (PPM), by weight, Benzotriazole™ is used. In various embodiments, system controls are designed to assure water temperature is non-condensing based on actual psychrometric conditions.

As noted above, the BPCH subassembly 104 is part of the BPA 100. With reference to FIG. 2, the BPC motherboard 118 of the BPCH subassembly 104 includes static circuit breakers (SCBs) 212, which are used to control and distribute DC power. For example, five SCBs may be implemented in the SCBs 212. Bulk DC power (-350V) applied to the BPCH subassembly 104 is switched through different outputs by the SCBs 212, which provide relay and fault protection functions. The BPCH subassembly 104 includes power supplies 210 that provide various voltages required in the BPA 100. In a disclosed embodiment, the BPCH subassembly 104 controls and monitors four BPF subassemblies. The BPCH subassembly 104 also provides communication paths to associated output ports. The BPCH subassembly 104 also includes a flexible support processor (FSP) daughtercard 120 that provides system level processing and Ethernet communication functions. A bulk power hub (BPH) daughtercard 116 provides Ethernet hub functions. The BPCH subassembly 104 has no single points of failure in the circuitry providing BPF subassembly power and control. In

one or more embodiments, the BPCH subassembly 104 also provides UEPO light-emitting diodes (LEDs), sensors, and connections that require 6V power, which is not required for operation of the BPCH subassembly 104.

5 The SCBs 212 perform the function of a circuit breaker, solid state relay, and in-rush control circuit. The SCBs 212 are designed to limit overloads by disabling a solid-state switch when output current exceeds a predefined threshold. A fuse may be attached to each of the SCBs 212 for backup protection. A scenario which would cause the backup fuse to open is
10 abnormal as it requires both an overload on an output of the BPCH subassembly 104 and a failure of one of the SCBs 212 to current-limit the overload. In one or more embodiments, the fuse for each of the SCBs 212 is installed permanently in the BPCH subassembly 104 and is non-replaceable. A 12V bias for analog circuitry of the SCBs 212 is generated on the BPCH subassembly 104 and is referenced to -350VDC. The bus that provides the 12V bias may be protected by a fusible resistor at the SCB bias input. When one of the SCBs 212 is
15 disabled, the rate at which the output voltage falls is dependent on the load impedance. In one or more embodiments, the SCBs 212 are turned on with a logical 'low' (-NBL) signal from the BPCH subassembly microcontroller 204. When the BPCH subassembly microcontroller 204 is reset, all SCB (-NBL) signals are 'high' disabling the SCBs. If one of the SCBs 212 latches off, the (-NBL) line must be toggled off, then on, to re-enable the SCB.
20 The state of each SCB output voltage (-DCVP) is reported to the BPCH subassembly microcontroller 204. The SCB output voltage is 'low' when output voltage is present. Without +12V bias, the SCB output voltage line is 'high'. If an SCB latches off due to over-current, its (+OC) signal is 'high'. The signals are reported to the BPCH subassembly microcontroller 204. Without +12V bias, the (+OC) line is 'high'.

25
In one embodiment, the BPH daughtercard 116 provides 32 10/100 MB/S Ethernet interfaces and 2 GB/s Ethernet interfaces. The BPH daughtercard 116 also provides thirty-two RS-422/RS-485 UART interfaces. Each 10/100 MB/s RJ-45 connector includes both types of interfaces (the 2 GB/S interfaces do not include RS-422/RS-485 UARTs). With reference
30 to FIG. 3, in one embodiment, the BPH daughtercard 116 includes three microcontrollers 302, 304, and 306 (i.e., two for the UART communication and multiplexing and one for

configuring and managing the Ethernet switch modules 308 and 310). In at least one embodiment, the BPH daughtercard 116 is virtual local area network (VLAN) configured.

The BPH daughtercard microcontroller 302 (BPH 0) configures the Ethernet switch modules 308 and 310 by via synchronous serial interfaces of the modules 308 and 310.

Communication to the BPCH subassembly 104 is on serial interfaces from the FSP 202 on TD1 using a universal asynchronous receive/transmit (UART) protocol. In one or more embodiments, the interface levels are TTL compatible. The BPCH subassembly 104 also has an RS-232 serial interface on T4, which is utilized for UEPO communication and BPCH subassembly 104 testing. The interface is compatible with a properly configured personal computer COM (serial) port. The protocol for the TTL interface and the RS-232 interface is identical. In an embodiment, four UARTs are wired from the FSP daughtercard 120. These interfaces use TTL/CMOS (3.3V) levels.

The BPCH subassembly 104 provides a hardware multiplexer 208 for UART communication on one of the UARTs of the FSP 202. The multiplexer 208 allows the FSP 202 to select a port for communication. The FSP interface is TTL/CMOS (3.3V) compatible, but the BPCH subassembly 104 provides RS-422/RS-485 (differential, point-point) driver/receivers on all other ports. The multiplexer 208 has four address bits controlled by the FSP 202. The BPH daughtercard microcontrollers 304 and 306 (respectively labeled 'BPH 1' and 'BPH 2') are used for multiplexing RS-422/RS-485 UART interfaces. The BPH daughtercard microcontrollers 304 and 306 each provide sixteen RS-422/RS-485 UART interfaces, which may be half-duplex interfaces. The BPH daughtercard microcontrollers 304 and 306 provide direction control for the RS-422/RS-485 drivers. The BPC motherboard microcontroller 204 controls and monitors the BPF subassemblies 108 using multiplexed RS-422/RS-485 UART interfaces (via UART multiplexer 206).

The BPCH subassembly 104 includes an RS-422/RS-485 UART serial control interface for each of the four BPF subassemblies 108. These interfaces allow fan speed control and sensing. Each of the BPF subassemblies 108 is also provided protected 12V power from power supply 210. The BPCH subassembly 104 can enable/disable and sense each power

output. The BPCH subassembly 104 includes power supplies that generate the following voltages from 350VDC input: +3.3V; +12V (SCB); +12V (BPF). Each of the four BPF subassemblies 108 provide current limiting and in-rush control for the +12V (BPF) voltage. Any BPF subassemblies 108 can be hot plugged/unplugged without disturbing the +12V (BPF) voltage to the other BPF subassemblies 108. A short on any one BPF subassemblies 108 does not affect the other BPF subassemblies 108 or the BPCH subassembly 104.

The power supplies 210 have the following protection and features: cross-powered 350V input; redundancy in that +12V (BPF) is generated with redundant circuitry for improved system availability; each redundant supply level is sensed by the BPC motherboard microcontroller 204; input current limiting in that a short in any BPCH subassembly bias supply does not disturb the -350VDC bus; the BPCH subassembly bias supplies are efficient and designed to minimize cooling requirements; the BPCH subassembly supplies do not start until the input voltage is over a set voltage; excessive voltage causes the supplies to disable switching as long as the over-voltage persists; and excessive BPCH subassembly board temperature causes the supplies to latch off.

The 6V UEPO power (used for the UEPO loop and panel) from the BPR subassemblies 110 is available independently from the BPCH subassembly bias power supplies. Once the UEPO loop is completed at power-up, the BPR assemblies 110 automatically enable the 350VDC to ensure that the BPCH subassembly power supplies start. If all 350VDC power is disabled, the system requires manual intervention to restart, because the FSP control logic is off. The 6V UEPO voltage from the BPR subassemblies 110 is sent to the UEPO switch connected to T4. The output is not over-current protected on the BPCH subassembly 104, as the BPCH subassembly 104 and FSP 202 are not powered from the 6V UEPO voltage. The BPCH subassembly 104 provides an RS-232 UART serial interface to the UEPO panel. This interface is also used during testing to send commands to the BPCH subassembly 104.

The UEPO panel provides the following sensing information to the BPCH subassembly 104: ambient temperature; humidity; barometric pressure; +UEPO power available; +room EPO complete/bypassed; and +UEPO complete. The BPCH subassembly 104 also supports the following panel functions directly: A/B side sensing; FSP to FSP handshake signals; and a

frame service LED. The BPC subassembly 104 provides a UEPO bypass switch to allow concurrent maintenance of the UEPO panel and/or cable. When the switch is set to “bypass”, the 6V UEPO loop is completed internally to the BPC motherboard 118. The maintenance can then be performed without affecting the 350VDC output from the BPR subassemblies 110. The BPC motherboard microcontroller 204 senses/controls: the -NBL, -DCVP, and +OC signals from each SCB; the -NBL, -DCVP, and -OC signals from each BPF power output; when BPR1-6 and BPF 1-4 are plugged; when UEPO (6V) power is available and UEPO is complete; when the UEPO panel is connected to T4; a bypass switch to concurrently maintain the UEPO panel/cable; and the 350V input from the T5 cross-power connector.

The BPC motherboard microcontroller 204 communicates with the UEPO panel and buffers the sensor information received. The BPC motherboard microcontroller 204 monitors the BPF subassemblies 108 and alters fan speed upon command from the FSP 202. The BPC motherboard microcontroller 204 also provides SCB status and enables/disables SCBs upon command from the FSP 202, provides BPF/BPR plug detect status, and provides bias sensing. The BPC motherboard microcontroller 204 stores code images for the UEPO panel and BPF subassemblies 108. After a successful power-on self-test, the BPC motherboard microcontroller 204 enables the BPC subassembly 104 ‘good’ LED. All SCBs are then disabled upon initial power-on (hardware reset). The BPC motherboard microcontroller 204 does not typically initialize fan speed of the BPF assemblies 108. The status of all BPC motherboard microcontroller 204 outputs is preserved during and immediately after code download. The BPH daughtercard microcontroller 302 continually validates the VLAN configuration of the switch modules 310 and 312 by reading the configuration registers and comparing the values against a table stored in flash associated with the microcontroller. This ensures the security of the network and checks for a reset condition of the switch modules which will have these registers cleared if a reset takes place.

The status of the port links are continually checked. If a link is found to appear or disappear, a 'link status change' flag and corresponding mask are set for a network manager to pick-up. Upon receiving an appropriate FSP command, the BPH daughtercard microcontroller 302 can alter an LED status for any port. The BPH daughtercard microcontroller 302 provides

bias sensing for regulators 312. In one embodiment, upon boot, the BPH daughtercard microcontroller 302 first checks that the switch modules 308 and 310 are being held in reset to ensure that the hub is never in an 'open network' configuration. The default VLAN configuration, which is stored in flash, is loaded into the switch modules 308 and 310. Then port forwarding is enabled on all available ports. The BPH daughtercard microcontroller 302 then enables the control flags for the port link monitor functions.

With reference to FIG. 4, a diagram 400 illustrates that the BPCH subassembly 104 includes an SELV block 402 that includes the BPH daughtercard 116 and the FSP daughtercard 120. The SELV block 402 is coupled to 350V circuitry 404 with transformers and opto-couplers 406.

With reference to FIG. 5, the BPD subassembly 106 includes SCBs 502 and 504, which are used to control and distribute DC power. In a disclosed embodiment, bulk DC power (350V) applied to the BPD subassembly 106 is switched through 26 different outputs (T1-T26) by the SCBs 502 and 504, which provide relay and fault protection functions. In this embodiment, the BPD subassembly 106 executes commands that initiate from the BPC motherboard 118 and does not perform any independent control function. For example, the commands provided to the BPD subassembly 106 include enabling/disabling SCBs 502 and 504 and querying status of the SCBs 502 and 504. In the disclosed embodiment, the BPD subassembly 106 has two independent microcontrollers 506 and 508 that each sense one-half the SCBs 502 and 504. In this manner, any microcontroller failure does not disable all of the BPD subassembly 106. The BPD subassembly 106 has two independent 3.3V/12V bias supplies 510 and 512. Each of the bias supplies 510 and 512 supply power to one of the microcontrollers 506 and 508 and one of the SCBs 502 and 504. In this manner, any bias supply failure does not disable the entire BPD subassembly 106. The 3.3V bias, 5V bias, 12V bias, SCB controls, and microcontroller circuitry are all referenced to 350V(-). The RS-422 interface to the BPC motherboard 118 is designed as a safety extra low voltage (SELV) interface to optically isolate the interface from the BPD subassembly 106.

The SCBs 502 and 504 in the BPD subassembly 106 function as circuit breakers, solid state relays, and in-rush control circuits. The SCBs 502 and 504 are designed to limit overloads

by disabling a solid-state switch when output current exceeds a predefined threshold. Each of the SCBs may have an associated fuse for backup protection. A scenario that would cause the backup fuse to open is abnormal as it requires both an overload on an output of the BPD subassembly 106 and an SCB failure to current-limit the overload. In at least one
5 embodiment, the fuse is installed permanently in the BPD subassembly 106 and is non-replaceable. When an SCB is disabled, the rate at which the output voltage falls is dependent on the load impedance.

The SCBs 502 and 504 are turned on with a logical 'low' (-NBL) signal from an associated
10 one of the microcontrollers 506 and 508. When one of the microcontrollers 506 and 508 is reset, all associated -NBL signals are 'high' disabling the associated SCBs. If an SCB latches off, the -NBL line is toggled off, then on, to re-enable the SCB. The state of each SCB output voltage (-DCVP) is reported to an associated one of the microcontrollers 506 and 508. The SCB output voltage is 'low' when output voltage is present. Without +12V
15 bias, the SCB output voltage line is 'high'. If an SCB latches off due to over-current, its (+OC) signal is 'high'. These signals are reported to an appropriate one of the microcontrollers 506 and 508. Without +12V bias, the (+OC) line is 'high'.

Communication to the BPD subassembly 106 is on serial interfaces from the BPC
20 motherboard 118 on T0A/T0B. In one or more embodiments, the protocol on the interface is a UART protocol and the interface levels are RS-422 compatible. The bit transmission rate on the interface may be, for example, 115.2 KBaud. In one embodiment, the BPD subassembly 106 has a single LED to indicate when the BPD subassembly 106 is good. The single LED is used for identification when replacing a BPD subassembly 106. Each of the
25 microcontrollers 506 and 508 also provides thermal protection for the BPD subassembly 106. In one or more embodiments, the SCBs 502 and 504 are disabled if an associated cold plate is too hot. In general, the bias supplies 510 and 512 are configured to protect themselves from insufficient airflow. After a successful power-on self-test, the microcontroller 506 enables the BPD good LED. All SCBs 502 and 504 are disabled upon
30 initial power on (hardware reset). The microcontrollers 506 and 508 may be the same or different. With reference to FIG. 6, a diagram 600 illustrates that the BPD subassembly 106

includes an SELV block 602 that is coupled to 350V circuitry 604 with transformers and opto-couplers 606.

5 According to various aspects of the present disclosure, the BPR subassembly 110 is configured to: provide a regulated -350 VDC isolated output to the BPCH subassembly 104 and BPD subassembly 106; draw low harmonic low noise power factor corrected input power; provide a line regulated standby voltage of 6V to the BPCH subassembly 104 and system emergency power off (EPO) loop; provide an EPO mechanism via an internal relay; and monitor a number of internal signals with microcontrollers to access the health of itself
10 and provides this information via a serial link. The high power levels of the BPR subassembly 110 are achieved through a modular arrangement of smaller power converter stages.

With reference to FIG. 7, each of the BPR subassemblies 110 include three electrically
15 identical power converter sections denoted by the letters R, S, and T. Each of the sections includes a rectifier front-end 702 and a group of DC/DC converters 704 that are commoned together to form an output power bus. Each of the sections has no dependency on the other two sections. That is, a failure in one section does not prevent the other two sections from operating. A standby housekeeping power supply (labeled 'STBY HSKPR') 706 provides
20 power to the control circuits for rectifier R and DC/DC converter group R when power is applied to an input of the BPR subassembly 110. Once rectifier R is started, the standby housekeeping power supply 706 bootstraps its input power feed to the output of rectifier R. A host microcontroller 708, which is powered from the standby housekeeping power supply 706, may then start the main housekeeping power supply (labeled 'MAIN HSKPR') 710.

25 The main housekeeping power supply 710 receives input power from the output of rectifier R and powers the controls for the remainder of the BPR subassembly 110. A relay is used to allow the -350V output to be enabled by providing needed operating bias to the DC/DC converter control circuitry. The relay also serves as a UEPO loop function. When the relay
30 is not activated all DC/DC converter controls mechanically have their bias removed. This prevents the output of the DC/DC converter from being active for cases where safety power

removal is necessary. The BPR subassembly 110 'BPR Good' signal is turned on when the bias is activated and the output is enabled.

5 A BPR subassembly may be configured with an on/off slide switch with a mechanical interlock to ensure that the BPR subassembly is always installed or removed in a disabled condition. In order to enhance current sharing on the input and output connectors between the three rectifiers and three groups of DC/DC converters, the three-phase input connections and DC output connections for each of the power converter stages may have their own dedicated set of pins. The BPR subassembly 110 includes a host and six ancillary
10 microcontrollers (one for each rectifier and DC/DC converter group).

With reference to FIG. 8, each rectifier 702 includes a single phase AC/DC converter whose control circuits perform resistive emulation (i.e., sinusoidal unity power factor input currents) for AC inputs. The rectifier switching power stage is a buck+boost converter with
15 four interleaved sections, denoted as sections W, X, Y, and Z. Each of the single phase rectifiers is connected to the third phase of the AC input line to provide increased immunity to power line disturbances. Under normal operation, no current flows on connection C. When required, the single phase input diode bridge connected to connections A and B becomes a three-phase diode bridge that includes connection C (by firing a pair of SCRs on
20 connection C). The three rectifier sections are alternately connected to the three-phase input line to provide a nominally balanced three-phase DELTA input, with one rectifier operating between any one line-to-line voltage of the AC input. Since each of the three rectifiers nominally operate from a different pair of line voltages, the (-) output of each rectifier is not common.

25 Input power to the BPR subassemblies 110 is provided from the backplane of the BPE 102. For example, one slightly longer input pin per phase may be implemented to engage AC first when a BPR subassembly 110 is plugged into the backplane. In this case, the pins are then paralleled with the remaining pin of the phase when the BPR subassembly 110 is completely
30 seated. In a typical implementation, an initial contacting pin/socket is configured to handle any arcing damage.

Connected to the output of each of the rectifiers is a group of DC/DC converters (e.g., a group may include four DC/DC converters) that generate a regulated output voltage and also provide galvanic isolation from the input line. The four DC/DC converters in a given group are referred to herein as W, X, Y, and Z. Each output of the DC/DC converters includes an ORing diode, which prevents a fault on one output section from bringing down the parallel converters. In one embodiment, DC/DC converters in the same group share a common voltage error signal, which forces all of the converters to operate at the same output current.

The BPR subassembly worst case maximum in-rush current per any phase at the initial application of AC is composed of input capacitor charging current. In various implementations, the charging current will not exceed 100A peak with a pulse width of 200uS and is dependent on the input source impedance. When the BPR subassembly is turned on to produce output voltage, the input soft start current will be 50/60Hz sinusoidal for AC inputs, or a ramp in the case of DC input. For most input fault conditions, the current causes the internal BPR input protection fuses to open. For most output load faults, the current (after the initial AC application in-rush) will not exceed the maximum steady-state values. This condition is when the BPR is operating in output current limit. Maximum input current with the -350V and 6V at maximum loading are typically less than or equal to 38A RMS balanced three-phase with a balanced AC input voltage of 428 volts. Under these conditions, the input current to each rectifier should be less than 22A RMS single phase. Input current should generally be limited to 58A RMS maximum by the BPR subassembly (33.5A RMS per rectifier).

In the case of DC inputs, the maximum input current with the -350V and 6V at maximum loading should be less than or equal to 66A at 437VDC. Under these conditions the input current to each of the three rectifiers should be less than 22A. System application conditions should generally prevent input currents from exceeding maximum input currents. In at least one embodiment, each rectifier is protected by one 30A, 600V fast-acting semiconductor fuses on each of the AC input phases.

Each rectifier is configured to switch over to processing power from a normally disabled phase if a loss occurs on one of the normally active phases. Loss of any phase will also

result in a phase fault indication in that rectifier. Loss of voltage on the normally disabled phase has no effect on input power processing provided the other active phases are present. Each rectifier is designed to start under a single phase condition provided the line voltage is within the steady-state operating requirements. In one or more embodiments, software is employed to control the operation of connection C. The BPR subassembly 110 is configured to operate error free during a battery discharge condition when operating off DC input.

Heat dissipated in the BPR subassemblies is primarily through a water cooled cold plate. In general, heat load to air should not exceed 250 Watts under non-fault conditions. Due to the BPR subassemblies modular design, a power on of only some of the power stages may be permitted to support lesser loads. Typically, low-power modes should keep the main housekeeping supply 710 off to eliminate the power required for the controls of two of the three rectifiers and their associated DC/DC converters. With a single DC/DC converter turned on and operated at up to 1.5A of output current, the BPR subassembly 110 should generally be able to operate indefinitely without water cooling.

In various embodiments, the output of the BPR subassembly 110 is magnetically isolated from the high voltage input and controls and is floating with respect to the green wire ground. The +6V output is magnetically isolated from the AC input and is considered a SELV output. Each of the DC/DC converters has a diode in its power path before the output connector to prevent internal output fault shorts from affecting the common bus for parallel N+1 redundant applications. Output sharing between up to twelve parallel BPR subassemblies may be accomplished by impedance matching. This results in the output voltage having a negative slope characteristic as output current is increased. In at least one embodiment, voltage regulation ensures that BPR subassemblies operating in parallel share within about ± 7.5 amps of the average current of all regulators in operation when using software adjustments.

DC/DC converters in the same group share a common regulation reference and voltage error signal with a current mode loop for each power stage. As a result DC/DC converters in the same group typically share current to within about ± 0.63 amps of the average current for the group. Each of the three DC/DC groups output voltage may be configured to be individually

adjustable via software. For example, the BPR subassembly 110 may be configured to maintain its output voltage at 350 volts $\pm 15V$ for any combination of the following conditions (voltage regulation is measured at the output terminals of the BPR subassembly 110): static input voltage changes within the steady-state operating range; error free power line disturbances; transient load changes (steps); and variations within converter's operating environment.

'Sparing' refers to the ability of a product to continue to provide partial function in the event portions of the product fail. In the case of the BPR subassembly 110, the system application leverages the modular power trains to perform sparing. In at least one embodiment, the BPR subassembly 110 allows operation at reduced power levels in the event one or more rectifier front-ends or DC/DC converters is defective. While operating under these conditions, the output application current should be reduced (e.g., by 6.3A for each DC/DC converter not operational. In the disclosed embodiment, it should be appreciated that one failed rectifier causes a group of four DC/DC converters to be non-operational. In general, current derating with input voltage may be applied to the lower capacity while operating in this mode.

In at least one embodiment, the +6V floating level (which is used for EPO coil(s) pick voltage) should be present within no more than three seconds after input voltage is applied to the BPR subassembly 110. In a typical design, the +6V level should maintain regulation of 5.9V to 7.8V for a DC load current range of 0 to 0.5A under non-fault conditions. The fundamental voltage ripple on the +6V should generally not exceed 60mV peak-to-peak. +BPR NBL energizes the relay that mechanically applies internal bias to the DC/DC control circuitry. Without this bias voltage, the 350V output cannot be enabled and this feature allows the BPR subassembly output voltage to be mechanically disabled for UEPO safety requirements. Once the pick voltage is applied, the 350V output can be turned on. When pick voltage is removed the BPR output voltage will be turned off within 10mS and then start to decay. The rate of decay is dependent on the load.

The BPR subassembly 110 may include an amber LED labeled as "Identify" and a green LED labeled as "BPR Good" that is readily visible when inserted into a BPA. These LEDs may be managed by the host microcontroller 708. In one or more embodiments, the BPR

subassembly 110 includes a full-duplex RS-422 interface that allows for control and monitoring from the BPC motherboard 118 and connects to the host microcontroller 708. Each rectifier includes a primary control board 712. The primary control board 712 includes circuitry (including a rectifier microcontroller) to operate the rectifier power train and
5 includes an ancillary microcontroller that reports to the host microcontroller 708. Each of the primary control boards 712 are referenced to different voltages on the input voltage mains.

In at least one embodiment, the rectifier microcontroller includes the following signal
10 connections: PrimeVoltsA; PrimeVoltsP; ACamps; Bulk V; IBBW, IBBX, IBBY, and IBBZ; TP; TempDiode; TempB+B; Bulk Adj; Line Fault; AMUX0, AMUX1, AMUX2; T/OnW#, T/OnX#, T/OnY#, T/onZ#; SCRnbl; Crowbar; Xmit; Rcv; and HOV.

PrimeVoltsA corresponds to the voltage of the three-phase input connections connected to a
15 three-phase full-wave diode bridge. PrimeVoltsP corresponds to the peak voltage of the three-phase input connected to a three-phase full-wave diode bridge. ACamps is the scaled line current into the rectifier. Bulk V is the scaled voltage at the output of the rectifier; IBBW, IBBX, IBBY, and IBBZ indicate the relative current through each of the buck+boost inductors. TP is the scaled highest output voltage of the housekeeping supply that powers
20 the secondary control board 714. TempDiode is the temperature of the cooling block for the input diode bridge measured with a thermistor. TempB+B is the higher temperature of the cooling blocks for the buck or boost switches measured with a thermistor. Bulk Adj (which may be is a 2KHz PWM signal) is used to adjust the regulation point of the rectifier output voltage.

25
Line Fault is used to indicate whether a line fault has occurred (e.g., if the instantaneous voltage of a full-wave three-phase bridge off the input is above 212Volts $\pm 5\%$) and is used to determine whether to fire the SCRs 802 and 804; AMUX0, AMUX1, and AMUX2 form the address in standard binary form for the analog multiplexer with AMUX0 as the least
30 significant bit. T/OnW#, T/OnX#, T/OnY#, and T/onZ# indicate whether an associated phase of the input rectifier is turned on. SCRnbl controls whether the SCRs on the input diode bridge are activated. Crowbar, when high, indicates the boost switches in the rectifier

should turn on until the input current protection causes power to be removed from the rectifier input; Xmit and Rcv are used to form a full-duplex UART connection to the host microcontroller 708; and HOV indicates whether a hardware overvoltage has been detected on the output bulk voltage.

5

In the disclosed embodiment, each of the group of DC/DC converters 704 includes a secondary control card 714 that includes circuitry (e.g., a DC/DC group microcontroller) to operate the DC/DC power trains of the converters 704. The card 714 includes one ancillary microcontroller that reports to the host microcontroller 704. The secondary control boards 714 are referenced to the common positive rail of the DC/DC converter output power buses. In at least one embodiment, the DC/DC group microcontroller includes the following signal connections: TempDCW, TempDCX, TempDCY, and TempDCZ; T/OnW, T/OnX, T/OnY, and T/OnZ; VoAdj; IsnsAdj; IlimAdj; TempDCA; ISW, ISX, ISY, and ISZ; HOV; AMUX0, AMUX1, and AMUX2; Xmit; and Rcv.

10

15

TempDCW, TempDCX, TempDCY, and TempDCZ correspond to temperatures of the cooling block for the output diode bridges of the respective DC/DC converters measured with a thermistor. T/OnW, T/OnX, T/OnY, and T/OnZ are controlled by a master turn on bit which determines whether all DC/DC converters are on. VoAdj (which may be a 57KHz PWM signal) is used to adjust the regulation point of all of the DC/DC converters. IsnsAdj (which may be a 57KHz PWM) is used to null the offset of the Iout signal. IlimAdj (which may be a 57KHz PWM signal) is used to adjust the current limit point of all of the DC/DC converters.

20

25

TempDCA measures the local air temperature of the secondary control card 714 via a connection to a temperature sensor powered from AVCC. ISW, ISX, ISY, and ISZ are signals whose values are proportional to the average primary current of the respective DC/DC converters. HOV indicates whether a hardware (i.e. fast) overvoltage has been detected on the output bus. AMUX0, AMUX1, and AMUX2 form the address in standard binary form for the analog multiplexer with AMUX0 as the LSB. Xmit and Rcv form a full-duplex UART connection to the host microcontroller 708. The host microcontroller 708

30

aggregates and acts on the information from the six ancillary microcontrollers (i.e., the rectifier microcontrollers and the DC/DC group microcontrollers).

Accordingly, a number of techniques have been disclosed herein that readily facilitate
5 implementation of a power conversion, control, and distribution system for a high performance computing (HPC) cluster or clusters.

The flowchart and block diagrams in the figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products
10 according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks
15 shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified
20 functions or acts, or combinations of special purpose hardware and computer instructions.

The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used herein, the singular forms "a,"
25 "an," and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising," (and similar terms, such as includes, including, has, having, etc.) are open-ended when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one
30 or more other features, integers, steps, operations, elements, components, and/or groups thereof.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below, if any, are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present invention has been presented for purposes of
5 illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. The embodiment was chosen and described in order to best explain the principles of the invention and the practical application, and to enable others of ordinary skill in the art to understand the
10 invention for various embodiments with various modifications as are suited to the particular use contemplated.

Having thus described the invention of the present application in detail and by reference to preferred embodiments thereof, it will be apparent that modifications and variations are
15 possible without departing from the scope of the invention defined in the appended claims.

CLAIMS

1. A power conversion, control, and distribution system, comprising:
 - multiple bulk power regulator (BPR) subassemblies each configured to provide
 - 5 regulated DC bulk power from both AC input power and DC input power;
 - a bulk power distribution (BPD) subassembly configured to distribute the regulated DC bulk power; and
 - a bulk power controller and hub (BPCH) subassembly coupled to the multiple BPR subassemblies and the BPD subassembly, wherein the BPCH subassembly is configured to
 - 10 monitor and control the BPR subassemblies and the BPD subassembly.

2. The power conversion, control, and distribution system of claim 1, further comprising:
 - multiple bulk power fan (BPF) subassemblies configured to remove heat load from
 - 15 air dissipated by the BPR subassemblies and the BPD subassembly.

3. The power conversion, control, and distribution system of claim 1 or claim 2, further comprising:
 - a bulk power enclosure (BPE) including a backplane configured to couple the BPCH
 - 20 subassembly to the BPR subassemblies and the BPD subassembly.

4. The power conversion, control, and distribution system of any preceding claim, wherein each of the BPR subassemblies includes multiple electrically identical power converter sections.
- 25

5. The power conversion, control, and distribution system of claim 4, wherein each of the multiple electrically identical power converter sections includes a rectifier front-end and a group of DC/DC converters.

- 30

6. The power conversion, control, and distribution system of claim 5, wherein respective outputs of the group of DC/DC converters are electrically connected to form an output power bus.

7. The power conversion, control, and distribution system of claim 5, wherein each of the BPR subassemblies includes a relay that is controlled to disable the group of DC/DC converters.

5 8. The power conversion, control, and distribution system of claim 5, wherein the rectifier front-end includes a single phase AC/DC converter that performs resistive emulation for AC power inputs.

10 9. The power conversion, control, and distribution system of claim 5, wherein the group of DC/DC converters function as a buck plus boost converter with four interleaved sections.

10. The power conversion, control, and distribution system of claim 5, wherein the group of DC/DC converters includes four DC/DC converters that are configured to provide galvanic isolation from an input power line.

15

11. The power conversion, control, and distribution system of claim 10, wherein each output of the four DC/DC converters includes an ORing diode that prevents a fault on one or more of the four DC/DC converters from rendering remaining ones of the four DC/DC converters inoperable.

20

12. The power conversion, control, and distribution system of any preceding claim wherein the BPCH subassembly includes a bulk power control (BPC) motherboard that includes multiple static circuit breakers that function to route the DC bulk power and a bulk power hub (BPH) daughtercard that is isolated in a safety extra low voltage (SELV) block and is configured to provide Ethernet communication functions.

25

13. A high performance computing cluster, comprising a power conversion, control, and distribution system as claimed in any preceding claim, and further comprising:

30

one or more processing nodes coupled to the BPD subassembly and configured to receive operating power derived from the regulated DC bulk power.

14. A method of providing operating power to a high performance computing cluster, comprising:

generating, using multiple bulk power regulator (BPR) subassemblies, regulated DC bulk power from either AC input power or DC input power;

5 distributing, using a bulk power distribution (BPD) subassembly, the regulated DC bulk power to one or more processing nodes of a high performance computing cluster;

monitoring and controlling, using a bulk power controller and hub (BPCH) subassembly, the BPD subassembly and the multiple BPR subassemblies; and

10 coupling, using a backplane of a bulk power enclosure (BPE), the BPCH subassembly to the multiple BPR subassemblies and the BPD subassembly, wherein the one or more processing nodes are configured to receive operating power derived from the regulated DC bulk power.

15 15. The method of claim 14, wherein each of the BPR subassemblies includes multiple electrically identical power converter sections, and wherein each of the multiple electrically identical power converter sections includes a rectifier front-end and a group of DC/DC converters.

20 16. The method of claim 15, wherein respective outputs of the group of DC/DC converters are electrically connected to form an output power bus and each of the BPR subassemblies includes a relay that is controlled to disable the group of DC/DC converters.

25 17. The method of claim 16, wherein the group of DC/DC converters function as a buck plus boost converter with four interleaved sections.

30 18. The method of claim 17, wherein the group of DC/DC converters includes four DC/DC converters that are configured to provide galvanic isolation from an input power line and each output of the four DC/DC converters includes an ORing diode that prevents a fault on one or more of the four DC/DC converters from rendering remaining ones of the four DC/DC converters inoperable.

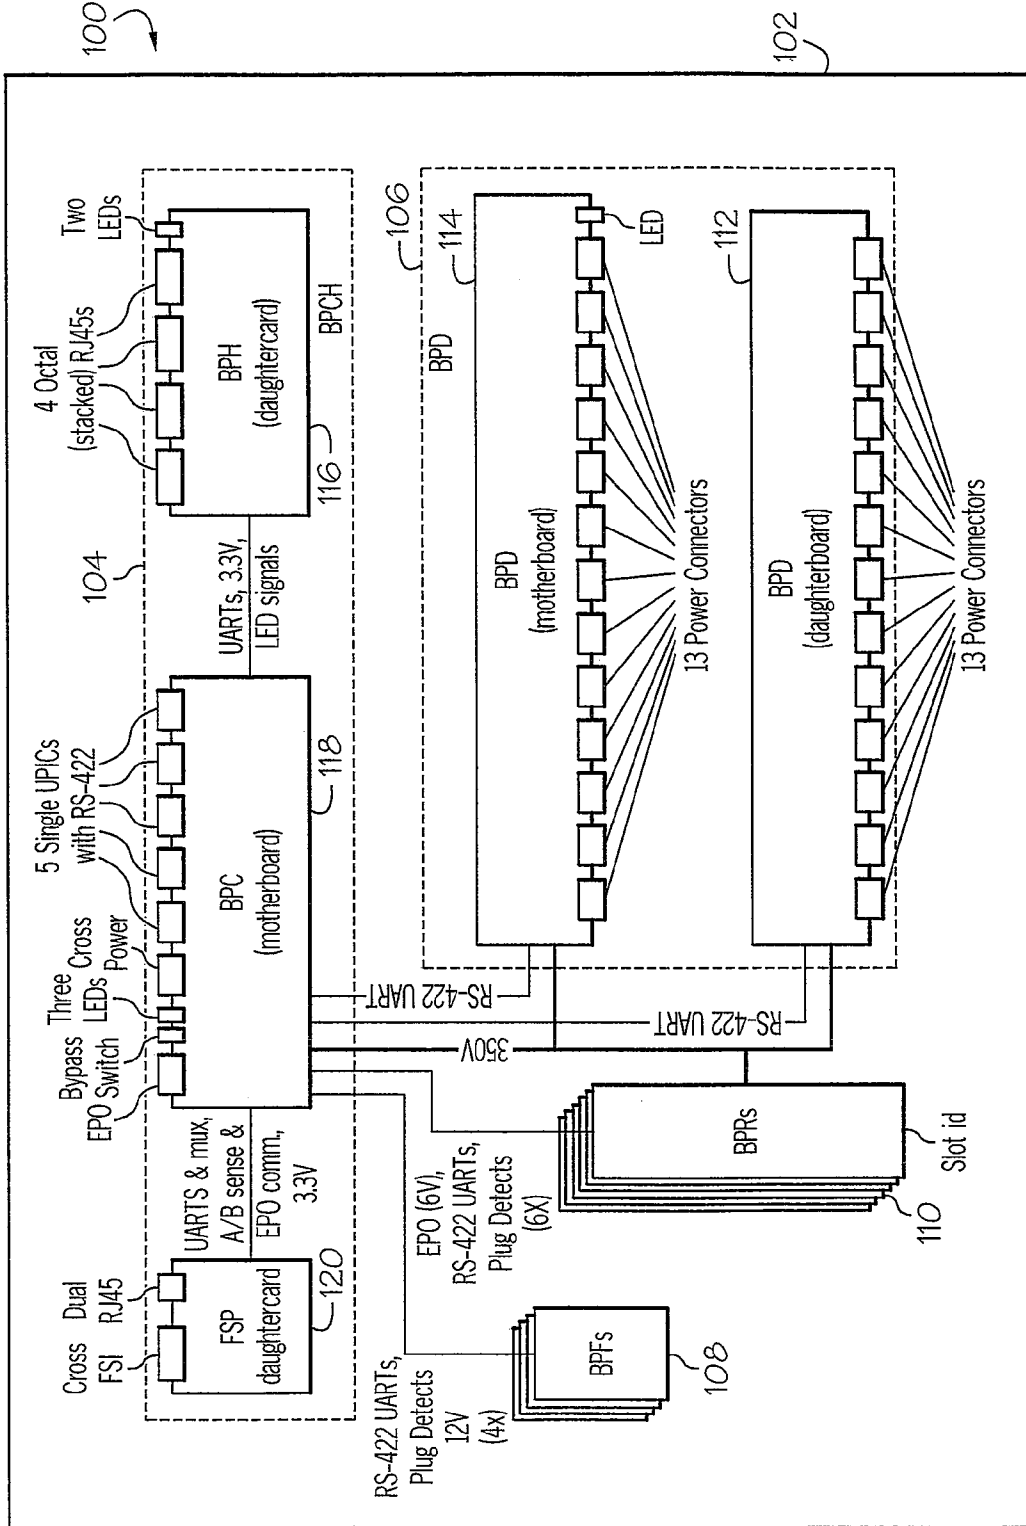


FIG. 1

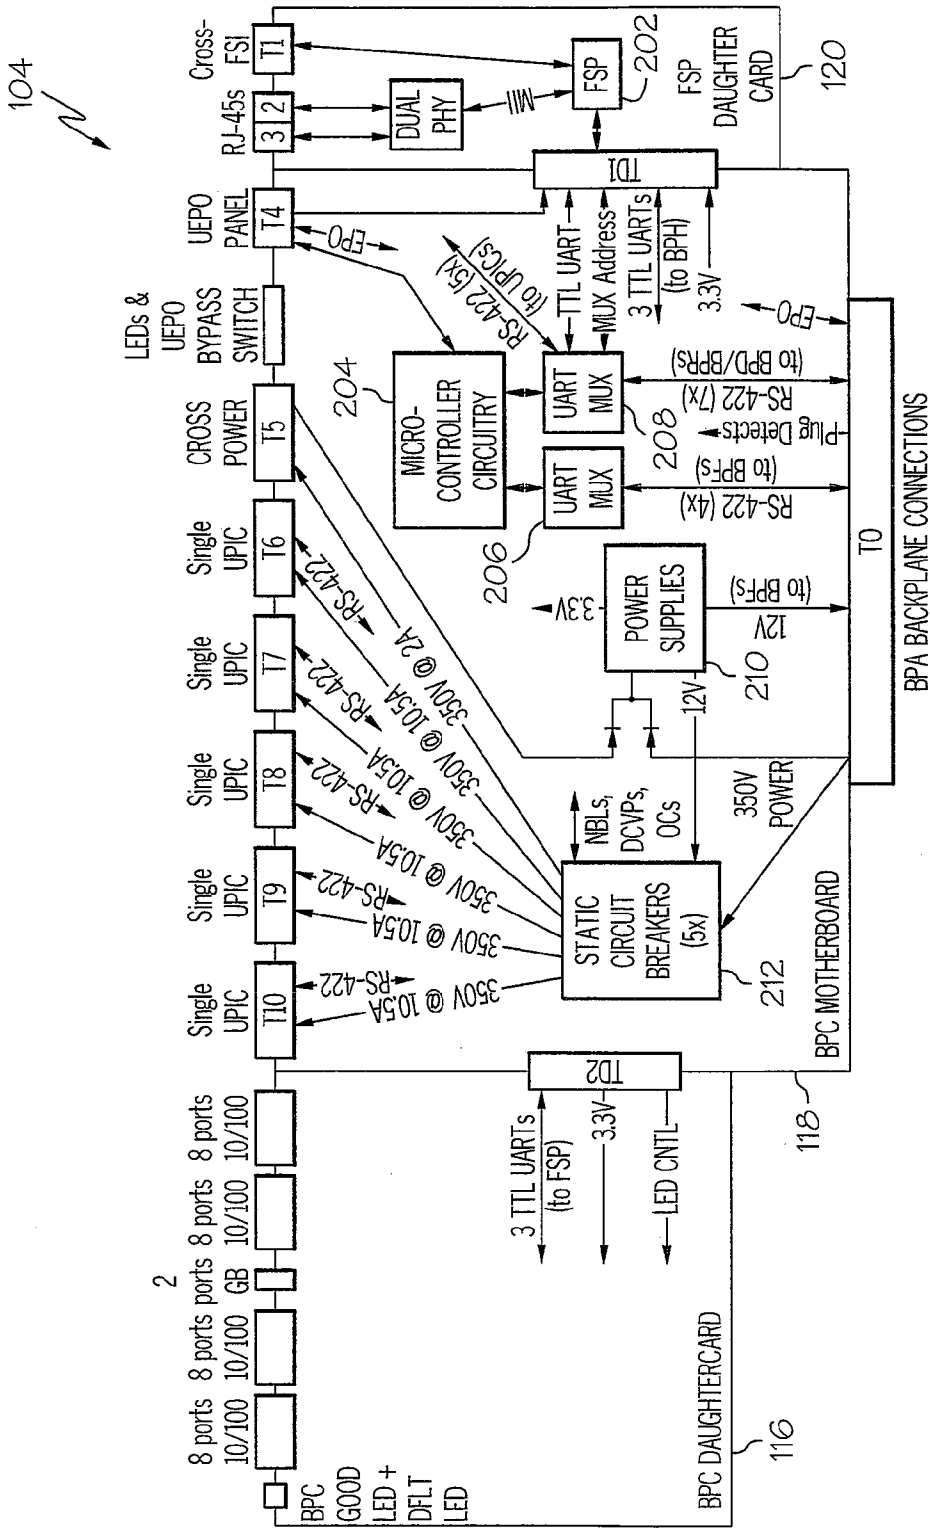


FIG. 2

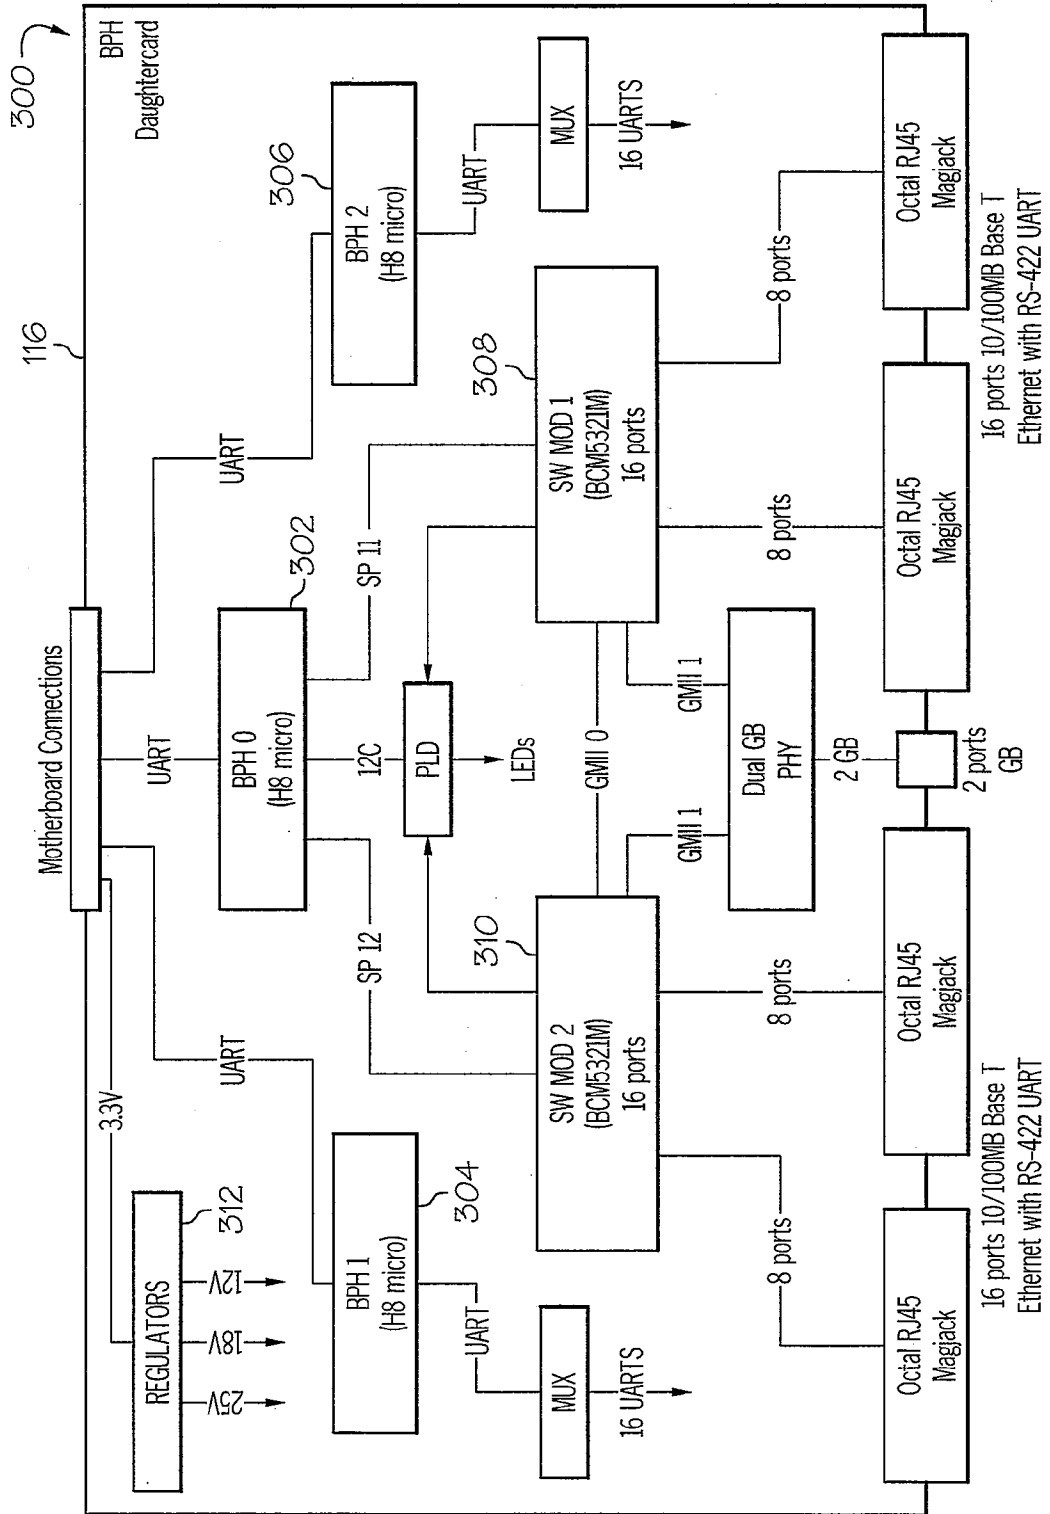


FIG. 3

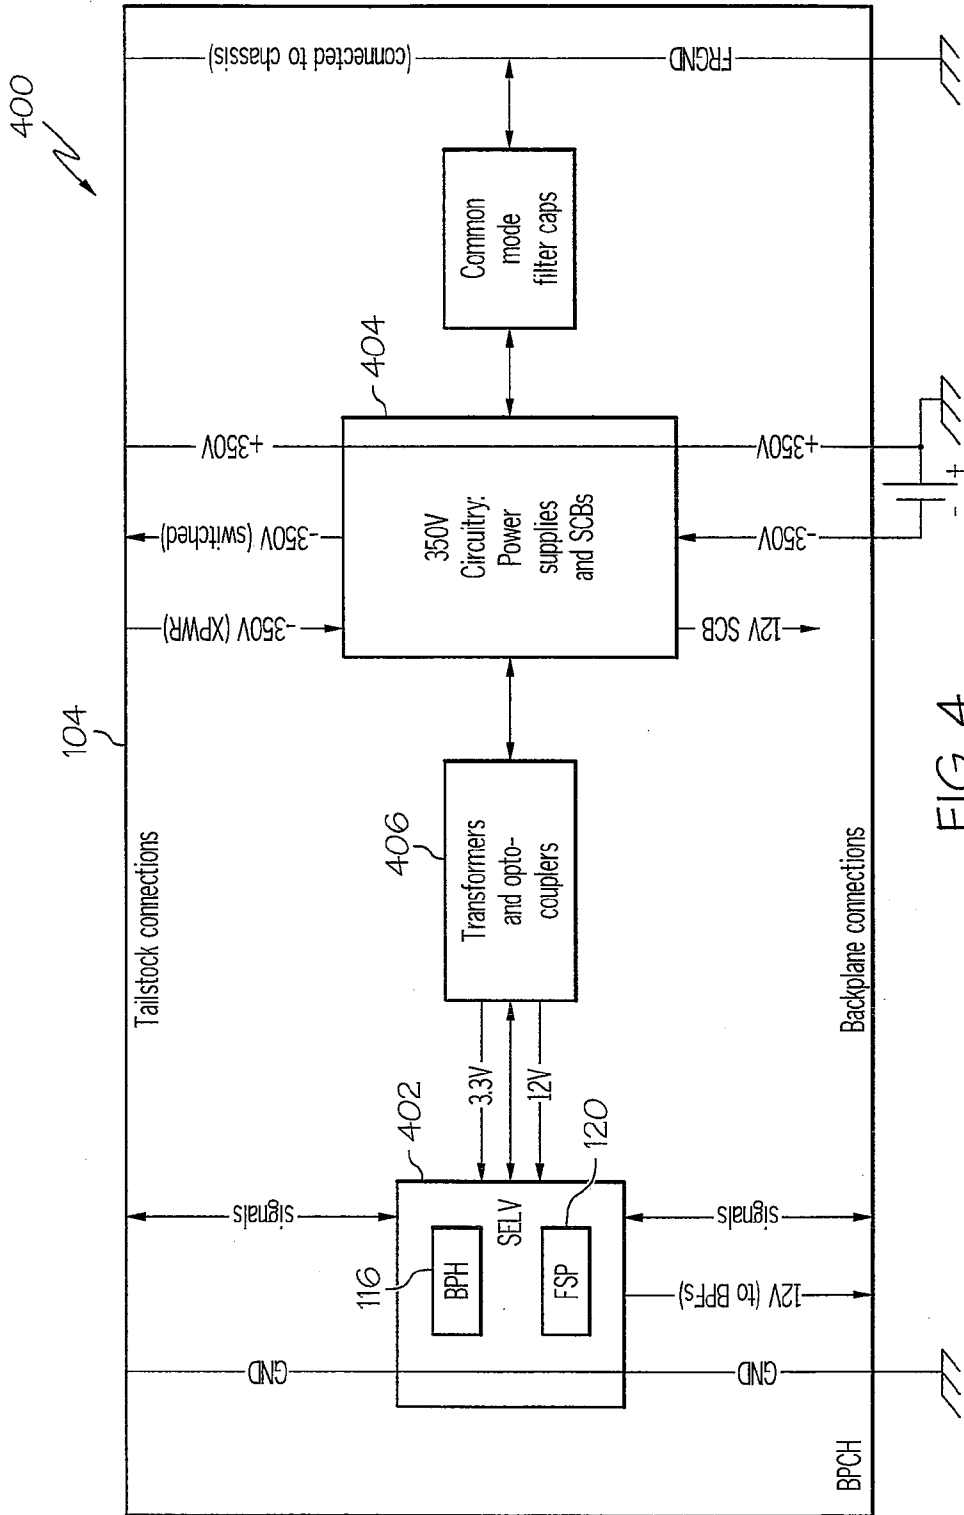


FIG. 4

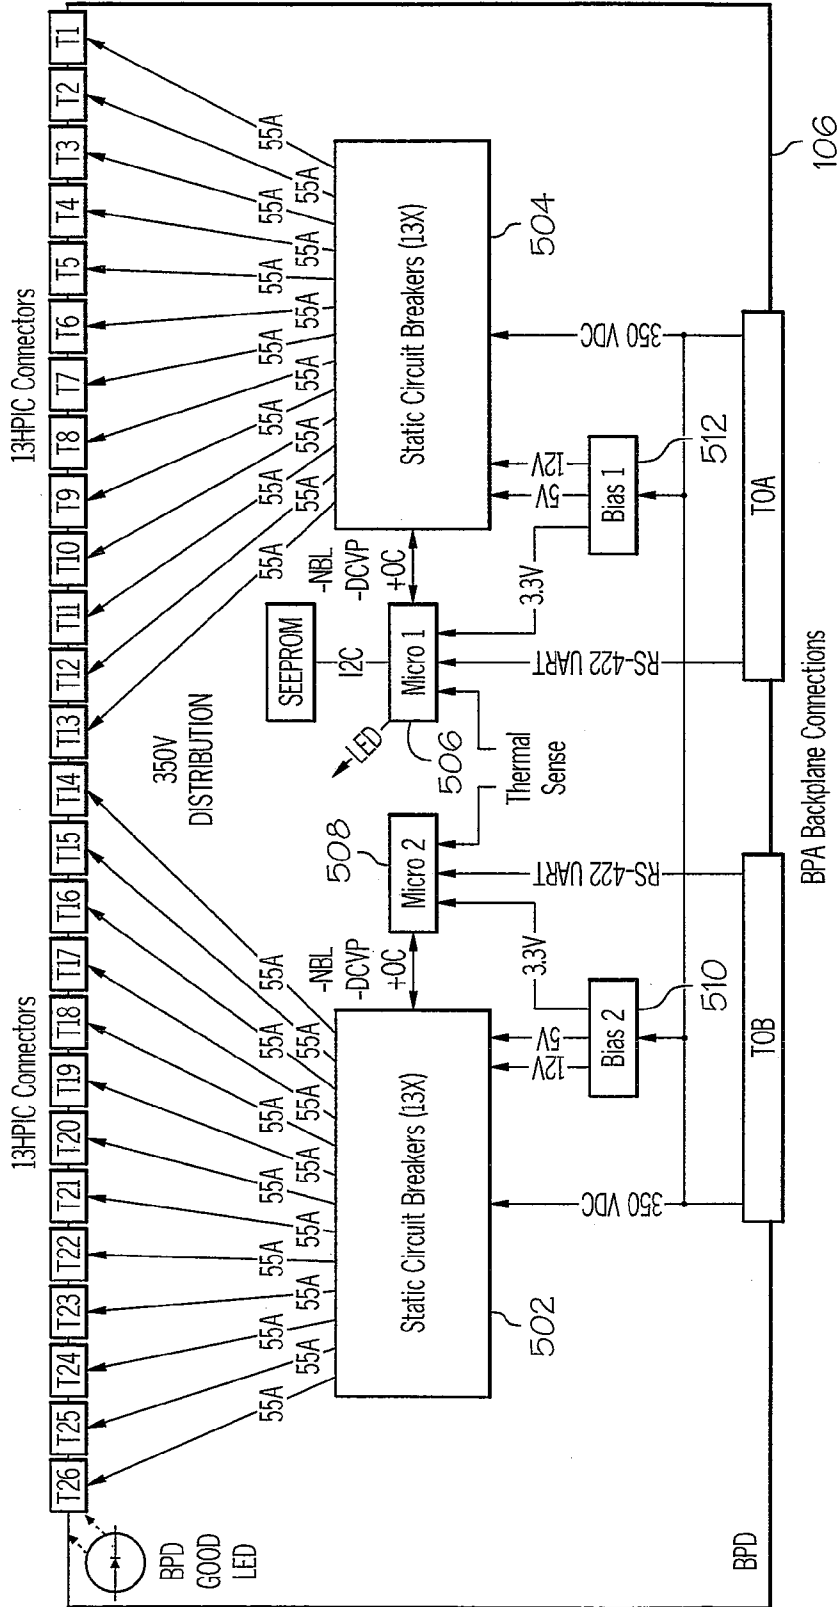


FIG. 5

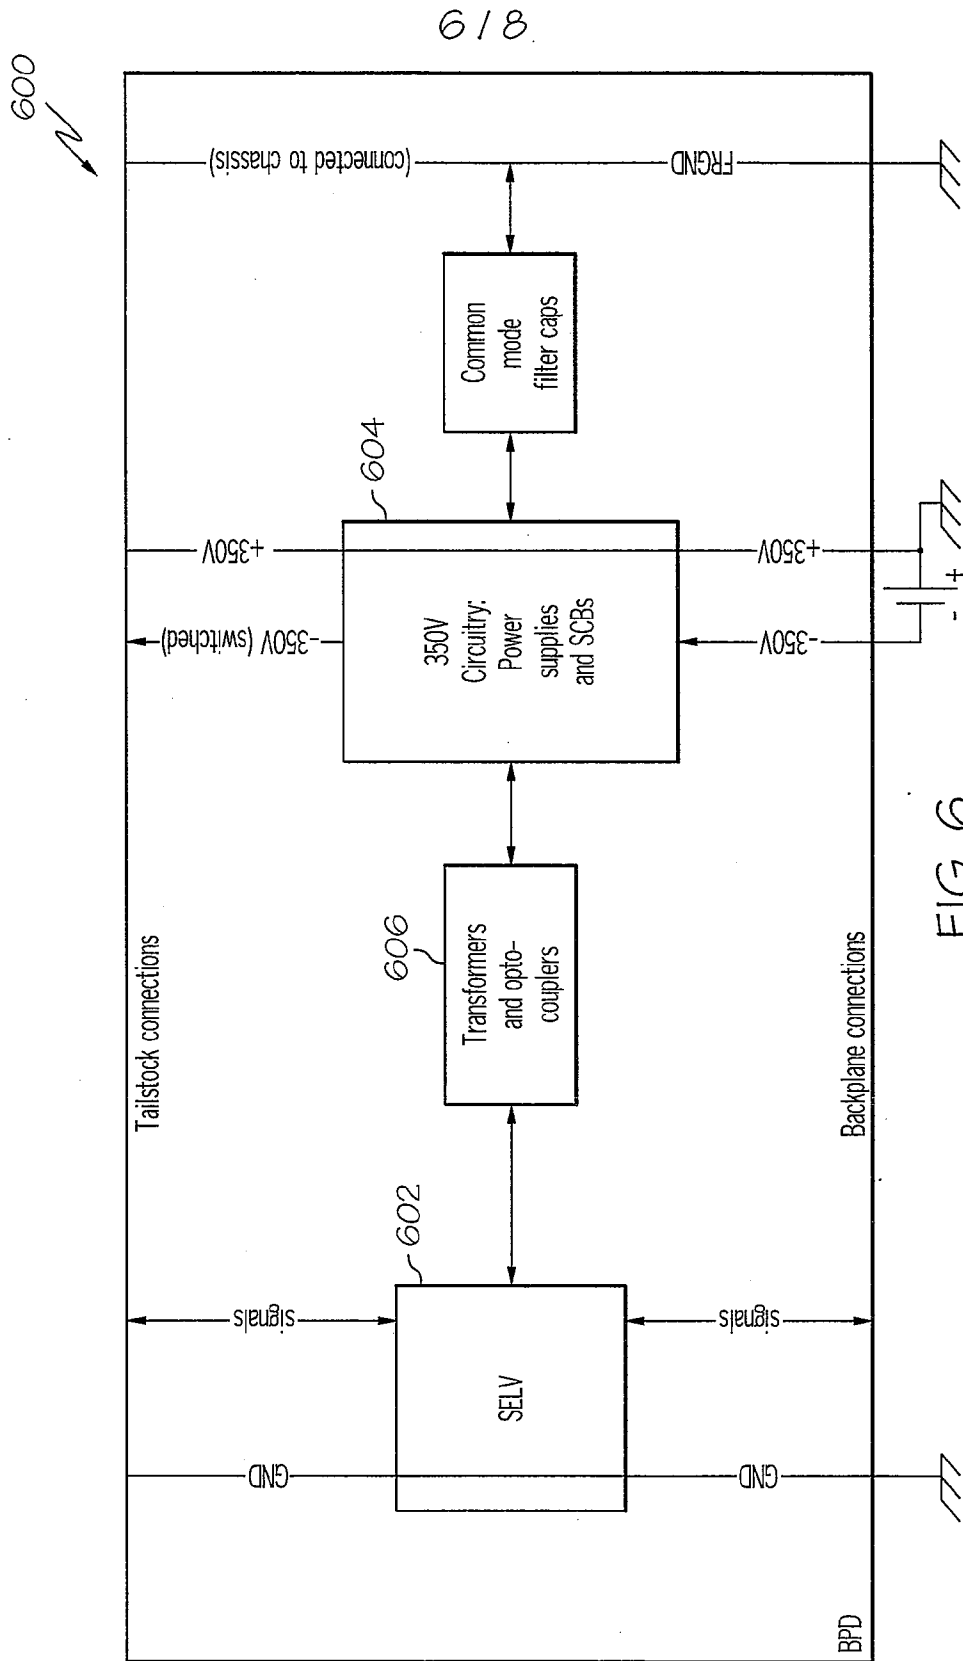


FIG. 6

718

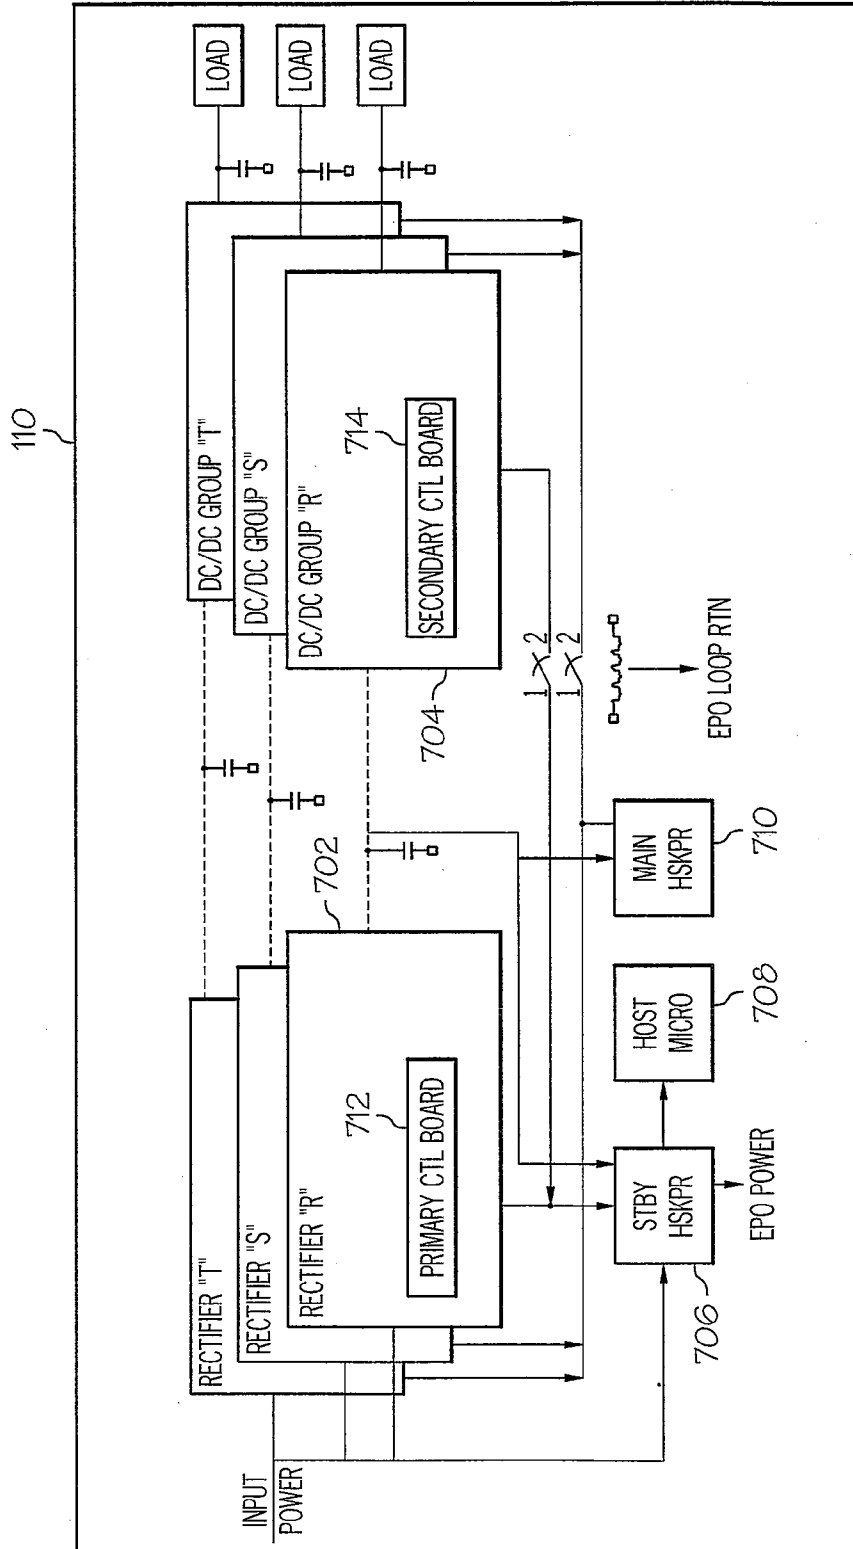


FIG. 7

8 / 8

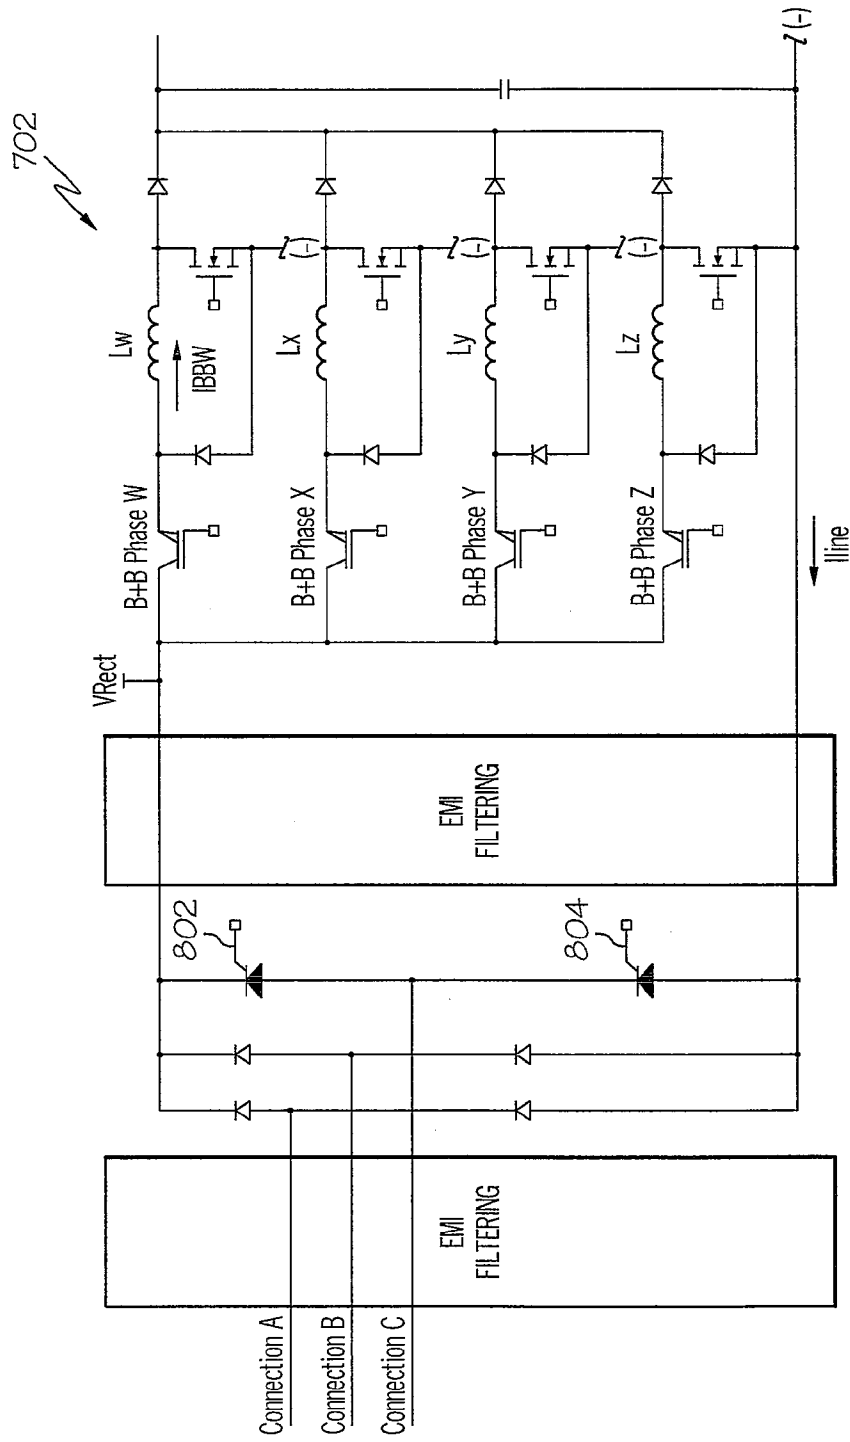


FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2010/054951

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06F1/26 H02J1/00
 ADD.
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
 Minimum documentation searched (classification system followed by classification symbols)
 G06F H02J

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 practical, search terms used)
 EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2002/181249 A1 (COFFEY JOSEPH [US]) 5 December 2002 (2002-12-05) figures 2,3 paragraphs [0001], [0010], [0012], [0032] -----	1-18
A	US 2008/040623 A1 (SONE MASAHIRO [JP]) 14 February 2008 (2008-02-14) * abstract figures 1,6,7,8,9 paragraphs [0039], [0042], [0060], [0063] -----	1-18

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

Date of the actual completion of the international search

16 July 2010

Date of mailing of the international search report

28/07/2010

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

Authorized officer

Zettler, Karl-Rudolf

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2010/054951

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2002181249	A1	05-12-2002	NONE
US 2008040623	A1	14-02-2008	NONE