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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
21 September 2006 (21.09.2006)

PCT

(10) International Publication Number
WO 2006/099613 A2

(51) International Patent Classification:
H04L 12/413 (2006.01)

(21) International Application Number:
PCT/US2006/010023

(22) International Filing Date: 16 March 2006 (16.03.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/082,607 16 March 2005 (16.03.2005) US

(71) Applicant *(for all designated States except US)*: **CISCO TECHNOLOGY, INC.** [US/US]; 170 West Tasman Drive, San Jose, California 95134 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **LUM, Meilissa, Rui Ming** [US/US]; 1399 Bouret Drive, San Jose, California 95118 (US). **KARAM, Roger** [US/US]; 981 Bonita Avenue, #15, Mountain View, California 94040 (US).

(74) Agent: **HANISH, Jonathan, D.**; SIERRA PATENT GROUP, LTD., 1657 Hwy 395, Suite 202, Minden, Nevada 89423 (US).

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

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

Declaration under Rule 4.17:

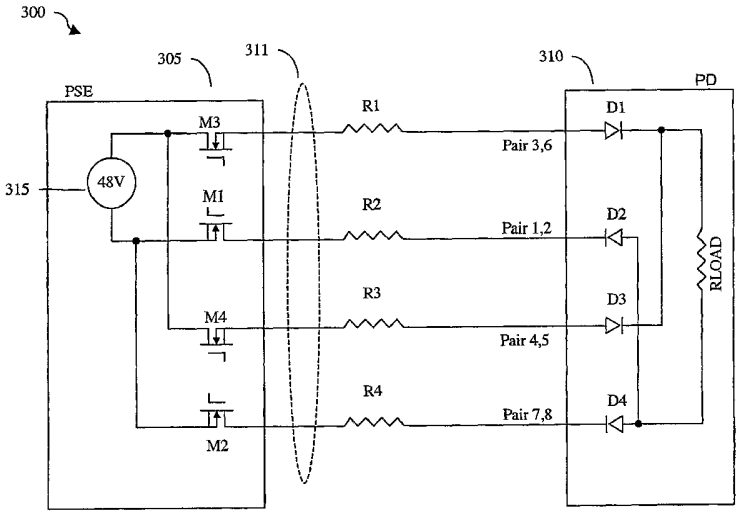
— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD AND APPARATUS FOR CURRENT SHARING ETHERNET POWER ACROSS FOUR CONDUCTOR PAIRS



(57) Abstract: Methods and apparatus for delivering power down all conductors of an Ethernet connection are disclosed. Embodiments include control circuitry coupled to both conductor of an Ethernet cable. The control circuitry is configured to sense the current flowing through power switches coupled inline with the respective conductors and control the switch in order to maintain the current flow through the power switch below a predetermined level. The circuitry prevents Ethernet ports from being shutting down due to conditions caused by imbalances in the effective resistance of Ethernet conductors and circuit paths.

WO 2006/099613 A2

SPECIFICATION

METHOD AND APPARATUS FOR CURRENT SHARING ETHERNET POWER ACROSS FOUR CONDUCTOR PAIRS

FIELD OF THE INVENTION

The present invention relates generally to networking equipment which is powered by and/or powers other networking equipment over wired data telecommunications network connections.

BACKGROUND OF THE INVENTION

Inline Power (also known as Power over Ethernet and PoE) is a technology for providing electrical power over a wired telecommunications network from power source equipment (PSE) to a powered device (PD) over a link section. The power may be injected by an endpoint PSE at one end of the link section or by a midspan PSE along a midspan of a link section that is distinctly separate from and between the medium dependent interfaces (MDIs) to which the ends of the link section are electrically and physically coupled.

PoE is defined in the IEEE (The Institute of Electrical and Electronics Engineers, Inc.) Standard Std 802.3af-2003 published 18 June 2003 and entitled "IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements: Part 3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications: Amendment: Data Terminal Equipment (DTE) Power via Media Dependent Interface (MDI)" (herein referred to as the "IEEE 802.3af standard"). The IEEE 802.3af standard is a globally applicable standard for combining the transmission of Ethernet packets with the transmission of DC-based power over the same set of wires in a single Ethernet cable. It is contemplated that Inline Power will power such PDs as Internet Protocol (IP) telephones, surveillance cameras, switching and hub equipment for the

telecommunications network, biomedical sensor equipment used for identification purposes, other biomedical equipment, radio frequency identification (RFID) card and tag readers, security card readers, various types of sensors and data acquisition equipment, fire and life-safety equipment in buildings, and the like. The power is direct current, 48 Volt power available at a range of power levels from roughly 0.5 watt to about 15.4 watts in accordance with the standard. There are mechanisms within the IEEE 802.3af standard to allocate a requested amount of power. Other proprietary schemes also exist to provide a finer and more sophisticated allocation of power than that provided by the IEEE 802.3af standard while still providing basic compliance with the standard. As the standard evolves, additional power may also become available. Conventional 8-conductor type RG-45 connectors (male or female, as appropriate) are typically used on both ends of all Ethernet connections. They are wired as defined in the IEEE 802.3af standard.

FIGS. 1A, 1B and 1C are electrical schematic diagrams of three different variants of PoE as contemplated by the IEEE 802.3af standard. In FIG. 1A a data telecommunications network 10a comprises a switch or hub 12a with integral power sourcing equipment (PSE) 14a. Power from the PSE 14a is injected on the two data carrying Ethernet twisted pairs 16aa and 16ab via center-tapped transformers 18aa and 18ab. Non-data carrying Ethernet twisted pairs 16ac and 16ad are unused in this variant. The power from data carrying Ethernet twisted pairs 16aa and 16ab is conducted from center-tapped transformers 20aa and 20ab to powered device (PD) 22a for use thereby as shown. In FIG. 1B a data telecommunications network 10b comprises a switch or hub 12b with integral power sourcing equipment (PSE) 14b. Power from the PSE 14b is injected on the two non-data carrying Ethernet twisted pairs 16bc and 16bd. Data carrying Ethernet twisted pairs 16ba and 16bb are unused in this variant for power transfer. The power from non-data carrying Ethernet twisted pairs 16bc and 16bd is conducted to powered device (PD) 22b for use thereby as shown. In FIG. 1C a data telecommunications network 10c comprises a switch or hub 12c without integral power sourcing equipment (PSE). Midspan power insertion equipment 24 simply passes the data signals on the two data carrying Ethernet twisted pairs 16ca-1 and 16cb-1 to corresponding data carrying Ethernet twisted pairs 16ca-2 and 16cb-2. Power from the PSE 14c located in the midspan power insertion equipment 24 is injected on the two non-data carrying Ethernet twisted pairs 16cc-2 and 16cd-2 as shown. The power from non-data carrying Ethernet twisted pairs 16cc-

2 and 16cd-2 is conducted to powered device (PD) 22c for use thereby as shown. Note that powered end stations 26a, 26b and 26c are all the same so that they can achieve compatibility with each of the variants described above.

Turning now to FIGS. 1D and 1E, electrical schematic diagrams illustrate variants of the IEEE 802.3af standard in which 1000 BaseT communication is enabled over a four pair Ethernet cable. Inline Power may be supplied over two pair or four pair. In FIG. 1D the PD accepts power from a pair of diode bridge circuits such as full wave diode bridge rectifier type circuits well known to those of ordinary skill in the art. Power may come from either one or both of the diode bridge circuits, depending upon whether Inline Power is delivered over Pair 1-2, Pair 3-4 or Pair 1-2 + Pair 3-4. In the circuit shown in FIG. 1E a PD associated with Pair 1-2 is powered by Inline Power over Pair 1-2 and a PD associated with Pair 3-4 is similarly powered. The approach used will depend upon the PD to be powered.

Inline Power is also available through techniques that are non-IEEE 802.3 standard compliant as is well known to those of ordinary skill in the art.

In order to provide regular Inline Power to a PD from a PSE it is a general requirement that two processes first be accomplished. First, a "discovery" process must be accomplished to verify that the candidate PD is, in fact, adapted to receive Inline Power. Second, a "classification" process must be accomplished to determine an amount of Inline Power to allocate to the PD, the PSE having a finite amount of Inline Power resources available for allocation to coupled PDs.

The discovery process looks for an "identity network" at the PD. The identity network is one or more electrical components which respond in certain predetermined ways when probed by a signal from the PSE. One of the simplest identity networks is a resistor coupled across the two pairs of common mode power/data conductors. The IEEE 802.3af standard calls for a 25,000 ohm resistor to be presented for discovery by the PD. The resistor may be present at all times or it may be switched into the circuit during the discovery process in response to discovery signals from the PSE.

The PSE applies some Inline Power (not "regular" Inline Power, i.e., reduced voltage and limited current) as the discovery signal to measure resistance across the two pairs of conductors to determine if the 25,000 ohm resistance is present. This is typically implemented as a first voltage for a first period of time and a second voltage for a second period of time, both voltages exceeding a maximum idle voltage (0 – 5 VDC in accordance with the IEEE 802.3af standard) which may be present on the pair

of conductors during an "idle" time while regular Inline Power is not provided. The discovery signals do not enter a classification voltage range (typically about 15 – 20V in accordance with the IEEE 802.3af standard) but have a voltage between that range and the idle voltage range. The return currents responsive to application of the discovery signals are measured and a resistance across the two pairs of conductors is calculated. If that resistance is the identity network resistance, then the classification process may commence, otherwise the system returns to an idle condition.

In accordance with the IEEE 802.3af standard, the classification process involves applying a voltage in a classification range to the PD. The PD may use a current source to send a predetermined classification current signal back to the PSE. This classification current signal corresponds to the "class" of the PD. In the IEEE 802.3af standard as presently constituted, the classes are as set forth in Table I:

Class	PSE Classification Current Range (mA)	Corresponding Inline Power Level (W)
0	0 – 5	15.4
1	8 – 13	4.0
2	16 – 21	7.0
3	25 – 31	15.4
4	35 – 45	15.4

TABLE I

The discovery process is therefore used in order to avoid providing Inline Power (at full voltage of -48VDC) to so-called "legacy" devices which are not particularly adapted to receive or utilize Inline Power.

The classification process is therefore used in order to manage Inline Power resources so that available power resources can be efficiently allocated and utilized.

It may be desired to provide inline power down all four pairs in a cable. As mentioned above, the existing IEEE 802.3af standard specifies that a PD can draw approximately 12.95W (@ 350 ma) down two pairs. However, if a PD requires more power, it may utilize all four pairs to draw approximately 12.95W from each set of

pairs. It is therefore desired to ensure that during such an operation, the PD does not draw more than 12.95W (or 350 ma) from each set of pairs.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated into and constitute a part of this specification, illustrate one or more embodiments of the present invention and, together with the detailed description, serve to explain the principles and implementations of the invention.

In the drawings:

FIGS. 1A, 1B, 1C, 1D and 1E are electrical schematic diagrams of portions of data telecommunications networks in accordance with the prior art.

FIG. 2 is an electrical schematic diagram of a typical Ethernet 10/100 Base T connection in accordance with the prior art.

FIG. 3 is a circuit diagram of an Ethernet system in which power is being delivered down all four pairs of the Ethernet cable.

FIG. 4 is a flow diagram of a method for monitoring current in a conductor in accordance with the teachings of this disclosure.

FIG. 5 is a flow diagram of a method for monitoring voltage in an Ethernet system in accordance with the teachings of this disclosure.

FIG. 6 is a circuit diagram of a system for balancing current delivered down four conductor pairs between Ethernet devices in accordance with the teachings of this disclosure.

FIG. 7 is a circuit diagram of a fully implemented system implemented in the PD for balancing current delivered down four conductor pairs between Ethernet devices in accordance with the teachings of this disclosure.

FIG. 8 is a circuit diagram of a system implemented in the PSE for balancing current delivered down four conductor pairs between Ethernet devices in accordance with the teachings of this disclosure.

FIG. 9 is a conceptual block diagram of a system implemented in a midspan device for delivering current down four conductor pairs to a PD in accordance with the teachings of this disclosure.

DETAILED DESCRIPTION

Embodiments of the present invention described in the following detailed description are directed at power and data redundancy in a single wiring closet. Those of ordinary skill in the art will realize that the detailed description is illustrative only and is not intended to restrict the scope of the claimed inventions in any way. Other embodiments of the present invention, beyond those embodiments described in the detailed description, will readily suggest themselves to those of ordinary skill in the art having the benefit of this disclosure. Reference will now be made in detail to implementations of the present invention as illustrated in the accompanying drawings. Where appropriate, the same reference indicators will be used throughout the drawings and the following detailed description to refer to the same or similar parts.

In the interest of clarity, not all of the routine features of the implementations described herein are shown and described. It will, of course, be appreciated that in the development of any such actual implementation, numerous implementation-specific decisions must be made in order to achieve the developer's specific goals, such as compliance with application- and business-related constraints, and that these specific goals will vary from one implementation to another and from one developer to another. Moreover, it will be appreciated that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking of engineering for those of ordinary skill in the art having the benefit of this disclosure.

Turning now to FIG. 2 a typical 2-pair Ethernet (10 Base T, 100 Base T and 1000BT if 4-pairs were used) connection is illustrated. Box 100 encompasses the Ethernet port as it might exist in a network device such as a switch, hub, router or like device. Within port 100 is a PHY or physical layer device 102 which includes transmit circuitry 104 and receive circuitry 106. The transmit circuitry 104 interfaces to a connector such as an RJ-45 connector (not shown here) and through the connector to a cable 108 which includes at least two pairs of conductors, the Pair 1 - 2 (110) and the Pair 3 - 6 (112). The interface between the transmit circuitry 104 and the cable 108 includes a center-tapped magnetic device such as transformer T1. T1 has a PHY-side including pins 1 and 2 and center tap 6, and a wire side including pins 3 and 5 and center tap 4. The PHY side is also referred to as the primary side; the wire side is also referred to as the secondary side of the magnetic device T1. Termination circuitry 114 provides a Vdd bias (here illustrated as +3.3VDC) to the primary of T1. The secondary of T1 is coupled to cable pair 112 which is, in turn,

coupled in operation to a network device 118 which may be another hub, switch or router or a PD such as a Voice Over Internet Protocol (VOIP) telephone or other network device.

The interface between the receive circuitry 106 and the cable 108 includes a center-tapped magnetic device such as transformer T2. T2 has a PHY-side including pins 1 and 2 and center tap 6, and a wire side including pins 3 and 5 and center tap 4. The PHY side is also referred to as the primary side; the wire side is also referred to as the secondary side of the magnetic device T2. Termination circuitry 116 provides a ground bias to the primary of T2. The secondary of T2 is coupled to cable pair 110 which is, in turn, coupled in operation to a network device 118. If the pairs of conductors shown belonged to a 1000 Base T wired data telecommunications network segment then each pair would transmit and receive at the same time and all four pairs in the cable would be used.

Center tap pins 4 of T1 and T2 are coupled to inline power circuitry including a 48 VDC power supply 120 for providing Inline Power over cable 108, control circuitry 122 and switch circuitry 124.

Figure 3 is a schematic diagram of an Ethernet system 300 configured to deliver power down four conductor pairs without current balancing. FIG. 3 shows a PSE 305 coupled to a PD 310 using all four conductor pairs of an Ethernet cable 311. Typically, power is not delivered down the unused pairs 4,5 and 7,8. However, in FIG. 3, pair 4,5 is coupled to pair 3,6, and pair 7,8 is coupled to pair 1,2. Thus, the pairs formed by conductors 3,6 and 4,5 are coupled from the positive side of 48V power source for PSE 315 to the PD 310 through diodes D1 and D3. The pairs formed by conductors 1,2 and 7,8 are coupled from the negative side of power source 315 of PSE 305 to the PD 310 through diodes D2 and D4.

To switch power to all four pairs, the PSE 305 includes power switches M1-M4 disposed inline with the conductor pairs. Power switches M1 and M2 are configured to close the path formed by the power source 315, the conductor pairs 1,2 and 7,8, and load of the PD 310 (shown as a load resistance RLOAD). Likewise, power switches M3 and M4 are configured to close the return path formed by the power source 315, the conductor pairs 3,6 and 4,5, and RLOAD. The effective resistance of the total path from the power source to the PD load, including that wires forming the cable 311, is shown as R1-R4 for pairs 3,6; 1,2; 4,5; and 7,8, respectively.

The result of coupling conductor pairs as shown in FIG. 3 is that the conductor pairs are electrically connected as parallel resistors, i.e., R1 and R3, and R2 and R4, each respectively form a circuit of parallel resistors.

As is appreciated by those skilled in the art, the IEEE 802.3af standard specifies a maximum 350 mA current level down powered conductor pairs. However, a small imbalance in the effective total path resistance consistent of intrinsic cable resistance, resistance of connectors, resistance of coupling windings of magnetics and the likes of one or more pairs can cause a shift in the voltage drop across the effective resistance of such pairs that would lead to mismatch in the current delivered down each pair. Ideally for a specific load, such current must be equal in all pairs. As can be seen from FIG. 3, because of the diode coupling used in the IEEE spec, a small change in the effective total path resistance of two or more pairs can therefore result in a proportional change in the magnitude of current delivered down each pair. Changes in the total effective path resistance through a conductor pair may be caused by a wide variety of factors, such as tolerances in conductor windings and connector couplings of the conductors, as well as tolerance swings in other components, such as coupling transformers typically used in Ethernet ports.

According to the IEEE 802.3af specification today, the PSE can shut down a port when the current is in the range of 350-400ma. Hence, when providing power down all four pairs, it is desired that the PSE port doesn't shut down because it is fooled into seeing a short circuit condition when in fact the current level exceeding the 350ma threshold is a result of the mismatch of the effective DCR (Direct Current resistance) of the total path.

Thus, a small imbalance of resistance between R1 and R3, or between R2 and R4, may result in the magnitude of current down each pairs to be mismatched. If the imbalance is large enough, the current down one pair may exceed the 350-400mA threshold (typically designed around 375mA) limit, and the system i.e., the PSE, may mistakenly detect a short circuit condition and shut down the device, when in reality the DCR mismatch causes the fault to appear due to effective resistance mismatch.

The present disclosure ensures that the current down any pair does not exceed a cutoff threshold, such as 350 mA, which is used as a soft limit to trigger actions that give the PSE or the PD a chance to initiate a corrective action to such an overcurrent condition pending the evaluation of whether it is a short, or a mismatch. While the 350ma soft limit may be deployed in both the PSE and the PD, the PSE has the option

of using the same threshold typically midway between the 350ma and the 400ma limit as a trigger instead of adding a new 350mA threshold. The present disclosure provides for methods and apparatus for monitoring the system voltage and current levels, either at the PSE or the PD, and responsive to the monitoring, modulating the gate voltage of power FET switches thereby increasing their ON-resistance to maintain a desired current level. Integrated circuits already present in a typical Ethernet device may be utilized to monitor the voltage and/or current levels in a pair. If conditions are detected that are out of range, the control circuitry will decrease the bias voltage applied to the gate of the power FET, causing the FET to conduct less current, limiting the current in the process. This will force a portion of the drawn current to flow down the other pair, thereby balancing the current flow between the conductor pairs. Such balancing does not necessitate the current in all pairs to be equal in magnitude, rather, the goal is to insure that neither hot-swap engine (circuitry known in the art for power management, i.e., detecting shorts on each sets of pairs) present on pairs 1,2 and 3,6 or pairs 4,5 and 7,8 is fooled into shutting the power down due to the effective resistance mismatch.

The present disclosure provides for a sense of such a mismatch, and provides a solution by sensing the current at a threshold below the 350mA to serve as a warning to a correcting loop inside the PSE or PD. Thus, the present disclosure allows for a current sense to make sure that prior to the activation of the short circuit threshold, in the range of 350-400ma, conditions are detected that indicate that the current is going towards the short circuit threshold to initiate an attempt at keeping the power on if such a condition is due to a mismatch. If the condition is caused by an imbalance of DCR, the current may be balanced ahead of such a short-circuit indication. Thus, the goal of this disclosure is keeping the power up while under a mismatch condition up to certain reasonable limit, but if the mismatch is very high, then a failure may exist, such as a broken wire, and power will need to be shut down.

The current sense of this disclosure takes place on each set of pairs independent of the other. Embodiments of this disclosure modulate the gate of the power switches to limit the current in one pair, enabling the other pair to conduct more of the current. Thus, the present disclosure effectively forces a limited mismatch in current so as to not cause an overcurrent or a shutoff condition. This is done while monitoring for true overcurrent conditions, such as short circuits, in parallel.

Figure 4 is a flow diagram of a method for monitoring current in accordance with the teachings of this disclosure. The process 400 begins in query 410, where it is determined whether the current in one conductor pair has exceeded a predetermined level. The process of query 410 may loop until a condition is detected in which the current level exceeds a pre-determined level. In preferred embodiments, a level of 350 mA is chosen as a threshold level, though other levels may be chosen based on particular applications.

If it is determined that the current exceeds a predetermined level, then the process moves to act 420, where the current is decreased in the conductor. The process then moves to query 430, where it is determined whether the current has fallen below a predetermined level. If the current has fallen below a predetermined level, then the process moves to act 440, where the current level is increased.

The process then moves to query 450, where the current level is again sensed to determined whether the current level is below a predetermined level. If the current is still below a predetermined level, then the process loops back to query 410 after setting the current to a default level in act 460. If the current level is not below a predetermined level in query 450, then the process loops back to act 420, where the current is decreased. As will be appreciated, the process of FIG. 4 reacts in a fold-back manner if the current is aggressively exceeding the over-current threshold. If the current is hovering around the 350mA limit, it will soft limit at 350ma due to the loop corrective action – the process is the same. By alternatively lowering and raising the current level in steps in an attempt to adjust the current level to fall within a desired range the PSE or the PD can provide the means to keep the power on.

In the event of a fault such as a short circuit, the disclosed system may be unable to lower the current to within a range that allows the power to be kept on. Such a case will be detected in query 430, where the current has not fallen below a predetermined level even the attempt to lower the current in act 420. In such a case, the process will move to query 470, where the current is sensed to determined whether it is out of a predetermined range. If the current level is out of range, then the system is powered off in act 480, and the process ends. If the current has been determined to be in range in query 470, then the process loops back to act 420, where the current is lowered again.

To further provide for a fail-safe short circuit protection, the voltage across the port may also be sensed. Figure 5 is a flow diagram of a method 500 for determining

whether a short circuit condition exists in an Ethernet system. The process begins in query 510, where the process waits for a detection of a voltage level exceeding a predefined level. If such a detection has been made, the process moves to act 520, where a timer is then set. The process then enters the loop defined by acts 530-550, where the voltage is continued to be sensed until either the timer expires or the voltage is sensed below the predefined value. If the voltage is found to be below the predefined level in query 530 before the timer expires, the process then moves back to query 510.

If the voltage remains above the predefined level, the timer will be incremented in act 550. If the voltage level remains above the predefined level when the timer expires in query 530, then the port will be powered off in act 560.

Embodiments of circuitry for the above described methods will now be disclosed.

Figure 6 is a circuit diagram of a system 600 for balancing current delivered down four conductor pairs between Ethernet devices. The embodiment of FIG. 6 shows an example where the current balancing is performed in the PD. For simplicity, common circuitry such as diode bridges have been eliminated, and the implementation for one rail for each sets of pairs is shown, i.e., MOSFETs M1 and M2 present on pairs 3,6 and 4,5 show only half the supporting circuitry needed, the same circuitry would be used for pairs 1,2 and 7,8. FIG. 6 shows a PSE 610 coupled to a single PD 620 through an Ethernet cable 630. The effective path resistance including the intrinsic resistance of the cable 630 has been shown as a single resistance representing the combined resistances for each conductor pairs. Thus, FIG. 6 shows the cable 630 as comprising Rcable 36, Rcable 12, Rcable45, and Rcable 78.

In a typical PD port, control circuitry is provided in the form of a hot-swap controller that is used to sense overcurrents. In FIG. 6, a pair of controllers X1 and X2 are provided. The overcurrent sensing is done by the controller chip by measuring the voltage across a sense resistor (Rsense 1 and Rsense2). If there is an overcurrent, the voltage across the sense resistor will be too high, and the controller chip will decrease the gate-to-source voltage of a FET coupled is in series with the sense resistor. FIG. 6 shows FETs M1 and M2 coupled inline with Rsense1 and Rsense2, respectively.

By controlling the gate voltage, the controller chip can limit the current to a predetermined value. When in current limit, the FET is held in the linear region, such

that it acts like a variable resistor. When Vgs (gate to source voltage of the FET) decreases, the resistance of the FET increases, thus limiting the amount of current, thereby acting like a constant current source.

In preferred embodiments, the controllers X1 and X2 may comprise an LT1641 IC from Linear Technology or a functional equivalent. A datasheet may be found at:
http://www.lineartech.com/pc/downloadDocument.do?navId=H0,C3,P1781,D1566.

As can be seen from FIG. 6, the controllers X1 and X2 are coupled between their respective send and return pairs. Thus, X1 is coupled between pair 3,6 and 1,2, and X2 is coupled between pair 4,5 and 7,8. A sense input SENSE of each controller is coupled to sense the voltage across its respective FET and a pair of series resistors. A gate controller output GATE is coupled to the gate the controller's respective FET. The outputs of the FETs are then diode or-ed to the PD load RLOAD through the use of additional diodes D5 and D6.

In a typical operation, a capacitor is coupled from the TIMER output to ground. When an over-current situation is detected, a voltage is applied to charge the capacitor to initiate a timer-based current limit routine. When the voltage across the capacitor reaches a predefined level, the port is shut down. For example, if the LT1641 senses an overcurrent, the TIMER pin will source 80uA of current. The normal operation of this pin is to charge the capacitor, such that when the voltage reaches 1.233V, the output will be shut off.

In a preferred embodiment, the controller setup is modified to provide a controller that has the ability to be placed in a current limit state indefinitely. By doing so, the controller chip may remain in a current limit state such that if one pair is in current limit, the controller chip for the other pair can react and balance the current without the first chip shutting down. In this disclosure, the TIMER cap was replaced with a 10k resistor, so that the controller chip would not shut down its output if there is an overcurrent. By using a 10k resistor instead of a capacitor, we can use this pin as an indication if the controller chip is in current limit or not. If the voltage is 0.8V (80uA*10k), the chip is in current limit. If the voltage is 0V, the chip is not in current limit. Thus, through the use of the diodes D5 and D6 of FIG. 6 for example, the controller chip for the other pair can react and balance the current without the first chip shutting down. The second pair increases its delivered current to keep the power up because one of the controllers behaves as a current source under this condition,

i.e., the controller with an active current limit behaves as a current source, thus forcing the voltage at the FET to decrease and the excess current to flow on the other pair.

In preferred embodiments, the controller chips are configured to perform current limiting according to the method of FIG. 4 to comply with the limits as set forth in the IEEE 802.3af specification, or approximately 350-400mA. It is contemplated that any hot-swap controller chip that has an external timer, or that may have the time routine modified in accordance with this disclosure, may be employed in the present disclosure. Such circuitry may reside on a single IC, possibly with integrated FET switches, and will support the same action on all 4-pairs as needed.

FIG. 6 shows that the current sense of this disclosure takes place on each set of pairs independent of the other, that is, controller chip X1 operates on pairs 3,6 and 1,2, and controller chip X2 operates on pairs 4,5 and 7,8. The controller chips are configured to each operate independently in a current limit state to modulate the gate of the power FETs limit the current in one pair, while allowing the other pair to draw more of the current. Thus, the present disclosure effectively forces a pair conducting less current than the current-limited pair to draw more current, thus achieving a forced current balance between the pairs.

Thus, the system of FIG. 6 may be described as a four-pair power delivery system for Ethernet applications, where the current flow in each pair is dynamically monitored and adjusted so as to prevent over-current situations caused by differences in the electrical characteristics of the conducting paths. The dynamic monitoring may be performed by circuitry already present that has been modified from its original state with the addition of relatively few components.

Figure 7 is a circuit diagram of a full implementation with current limiting provided on the negative side as well. As currently specified, the IEEE specification does not mandate where the current limit needs to be in the PSE port. Thus, the switch may have protection residing in the either the positive or negative leg. FIG. 7 thus provides protection in both legs (positive and negative) in the PD so that the PSE port is not mistakenly tripped.

FIG. 7 shows controllers X1 and X2 coupled to pairs 3,6 and 4,5, respectively, as described above, and are shown as being coupled to V-, which is the negative side of the 48V power supply. Particular circuit details have been omitted for clarity.

Controllers X3 and X4 have been provided to control the FETs M3 and M4 that are inline on pairs 1,2 and 7,8, respectively.

The controllers X3 and X4 preferably comprise negative hot-swap controllers, such as the LT4252A, also available from Linear Technologies. M3 and M4 comprise the n-channel FETs that the controller chips control in case of an overcurrent. As mentioned above in relation to the positive controllers X1 and X2, it is preferred that the chips be placed in a current limit state indefinitely. It is contemplated that the negative hot-swap controllers by adding conventional circuitry to the TIMER pin. This circuitry is preferably configured to keep the TIMER pin in a state such that the chip stays in a current limit state indefinitely.

The current limits of the LT4252A are 483mA typical, with a sense resistor value of 0.124ohm (Rsense3 and Rsense4). This threshold can be scaled down to below 350mA by choosing suitable values, thus placing the controller chips in a current limit state well before a shut down condition.

Embodiments of this disclosure may also be implemented in the PSE. Figure 8 is a diagram of such an implementation. The structure is similar to that as described in FIG. 7, expect that the controller chips and control FETs are now disposed in the PSE.

However, if the PSE is configured to perform the current limiting of this disclosure, the PSE must also be able to sense and shut the port down if there is a short circuit condition, as mandated by the current IEEE802.3af specification.

In a preferred embodiment, the PSE is configured to monitor the output voltage of each set of pairs determine if there is a short circuit. FIG. 8 includes voltage control circuitry X5 and X6 for monitoring the overall output voltage of the power supply on each conductor pair set. The circuitry X5 and X6 are preferably coupled in parallel with the voltage source to sense the overall system voltage level.

It is contemplated that the voltage control circuitry of X5 and X6 may be comprised of conventional circuitry for determining whether the voltage has fallen below a predetermined level. It is further contemplated that the control circuitry X5 may be configured to perform the method of FIG. 5 wherein if the output voltage falls below 30V for the timer period, then the PSE knows there is a short circuit condition and may shut down the port. Under normal operating conditions and with a 5.1ohm imbalance, the output PSE port should never see a voltage less than 30V.

Additionally, the PSE may include fault bits associated with each pair set that may be set by the voltage control circuitry X5. If the PSE senses that fault bit have been set for both pairs, then the PSE may then shut off power to both powered pairs of the port.

The benefits of this disclosure may also be applied to a midspan device. Figure 9 shows a block diagram of a PSE providing 15.4W on the Ethernet pairs and a midspan device providing 15.4W on the unused pairs.

A midspan device is typically inserted between the switch port and a device such as an IP phone. The switch port may send 15.4W on its used pairs (wires 3,6 and 1,2), as is typical. The midspan device passes the used pairs straight through with 15.4W from the switch. However, using the circuitry of this disclosure, the midspan device may also put an additional 15.4W on the midspan pairs (wires 4,5 and 7,8). Thus, the PD is supplied with 15.4W from the switch port and 15.4W from the midspan port.

While embodiments and applications of this invention have been shown and described, it will now be apparent to those skilled in the art having the benefit of this disclosure that many more modifications than mentioned above are possible without departing from the inventive concepts disclosed herein. Therefore, the appended claims are intended to encompass within their scope all such modifications as are within the true spirit and scope of this invention.

CLAIMS

What is claimed is:

1. An apparatus for delivering power down all conductors of an Ethernet connection comprising:

an Ethernet conductor having eight conductors arranged as four pairs of conductors for connecting a Power Source Equipment (PSE) to a Powered Device (PD), each of said conductors of said pairs including a power switch disposed inline with said conductor;

said PSE being configured to provide inline power to said PD down all of said conductors;

control circuitry coupled to each of said conductors for controlling said power switch;

the control circuitry configured to sense the current flowing through its respective power switch; and

the control circuitry being configured to control its respective control switch to maintain the current flow through the power switch below a predetermined level.

2. The apparatus of claim 1, wherein each of said conductor pair comprises a pair of conductors electrically connected in parallel between said PSE and said PD.

3. The apparatus of claim 2, wherein said control circuitry senses a potential over-current condition due to an imbalance in the effective resistance between said pair of conductors.

4. The apparatus of claim 3, wherein said control circuitry in one of said conductors of said conductor pairs causes current to flow in the other conductor of said conductor pair to prevent a shut down of the Ethernet port to which said conductor pair is connected.

5. The apparatus of claim 4, wherein said control circuitry comprises hot-swap circuitry that is held in a current limit state indefinitely to maintain current flow below said predetermined level.

6. The apparatus of claim 5, wherein said predetermined current level is approximately 350 – 400 mA.

7. The apparatus of claim 4, wherein said control circuitry is disposed in said PSE.

8. The apparatus of claim 7, further comprising voltage control circuitry disposed in said PSE, the voltage control circuitry configured for detecting a short circuit condition in said Ethernet connection and shutting down said connection responsive to said detection.

9. The apparatus of claim 4, wherein said control circuitry is disposed in said PD.

10. The apparatus of claim 4, wherein said control circuitry is disposed in an Ethernet midspan device coupled between said PSE and said PD.

11. An apparatus for delivering power down all conductors of an Ethernet connection comprising:

Ethernet conductor means having eight conductors arranged as four pairs of conductors for connecting a Power Source Equipment (PSE) to a Powered Device (PD), each of said conductors of said pairs including power switching means disposed inline with said conductor;

said PSE being comprising means for providing inline power to said PD down each of said conductors of said Ethernet conductor means; and

control circuitry means coupled to each of said conductors for maintaining the current flow through the power switching means below a predetermined level.

12. The apparatus of claim 11, wherein each of said conductor pair comprises a pair of conductors electrically connected in parallel between said PSE and said PD.

13. The apparatus of claim 12, wherein said control circuitry means comprises means for sensing a potential over-current condition due to an imbalance in the effective resistance between said pair of conductors.

14. The apparatus of claim 13, wherein said control circuitry means comprises means for causing current to flow in the other conductor of said conductor pair to prevent a shut down of the Ethernet port to which said conductor pair is connected.

15. The apparatus of claim 14, wherein said control circuitry means is disposed in said PSE.

16. The apparatus of claim 15, further comprising voltage control means disposed in said PSE, the voltage control means comprising means for detecting a short circuit condition in said Ethernet connection and shutting down said connection responsive to said detection.

17. The apparatus of claim 14, wherein said control circuitry is disposed in said PD.

18. The apparatus of claim 14, wherein said control circuitry is disposed in an Ethernet midspan device coupled between said PSE and said PD.

19. An Ethernet Power Source Equipment (PSE) device for delivering power down all conductors of an Ethernet connection to an Ethernet Power Device (PD) comprising:

an Ethernet port disposed in the PSE for connecting an Ethernet conductor having eight conductors arranged as four pairs of conductors;

a power switch disposed inline with each of said conductors for providing inline power to a PD down all of said conductors;

control circuitry coupled to each of said conductors within said PSE for controlling a respective said power switch; and

the control circuitry configured to sense the current flowing through said respective power switch and maintain the current flow through the power switch below a predetermined level.

20. The device of claim 19, wherein each of said conductor pair comprises a pair of conductors electrically connected in parallel between said PSE and said PD.

21. The device of claim 20, wherein said control circuitry senses a potential over-current condition due to an imbalance in the effective resistance between said pair of conductors.

22. The device of claim 21, wherein said control circuitry in one of said conductors of said conductor pairs causes current to flow in the other conductor of said conductor pair to prevent a shut down of the Ethernet port to which said conductor pair is connected.

23. The device of claim 22, wherein said control circuitry comprises hot-swap circuitry that is held in a current limit state indefinitely to maintain current flow below said predetermined level.

24. The device of claim 23, wherein said predetermined current level is approximately 350 – 400 mA.

25. The apparatus of claim 22, further comprising voltage control circuitry disposed in said PSE, the voltage control circuitry configured for detecting a short circuit condition in said Ethernet port and shutting down said port responsive to said detection.

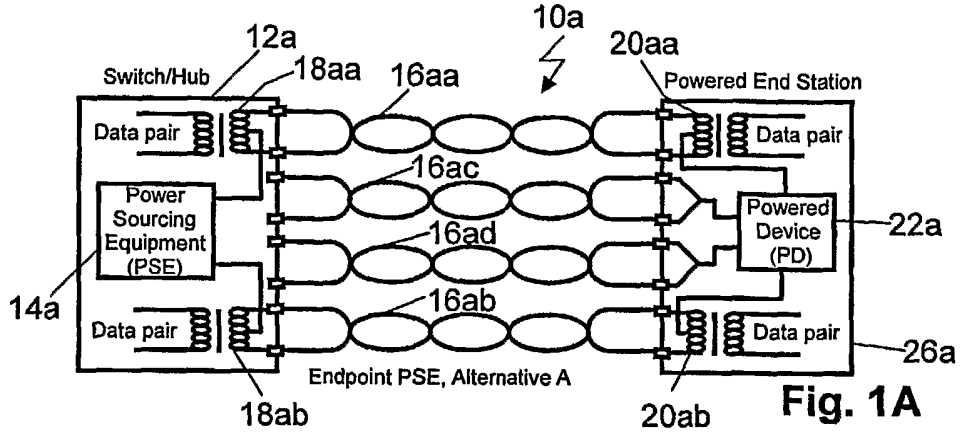


Fig. 1A
Prior Art

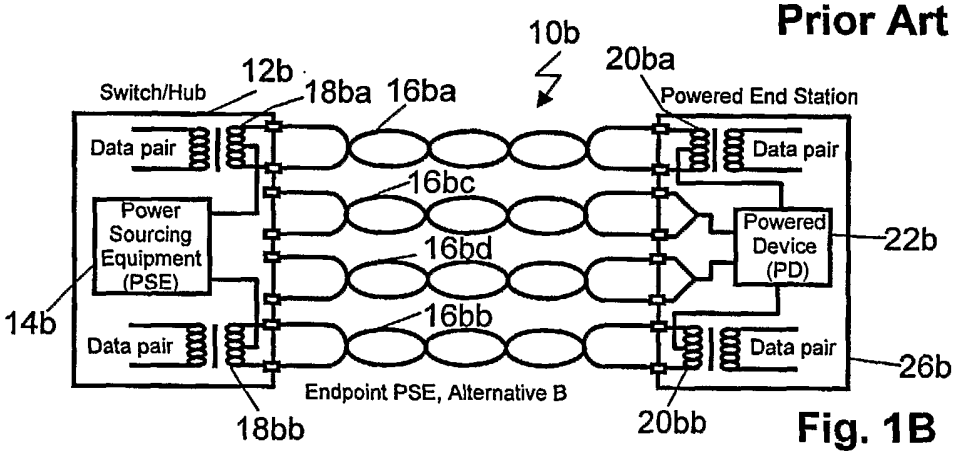


Fig. 1B
Prior Art

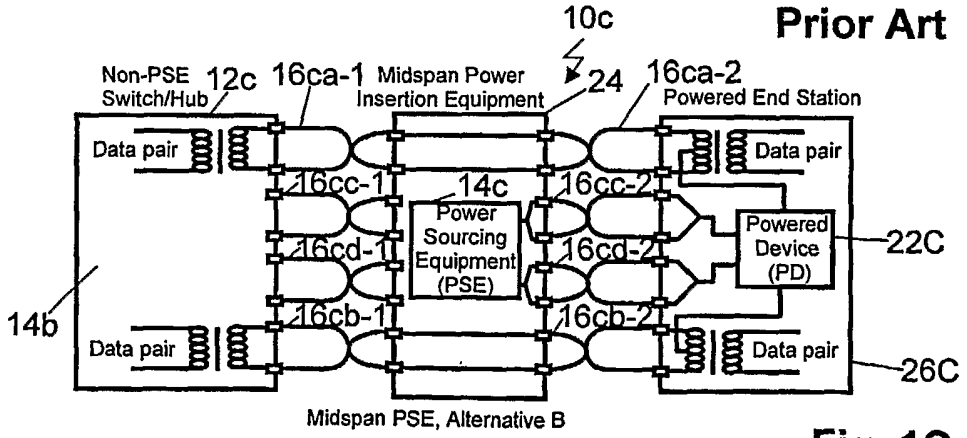


Fig. 1C
Prior Art

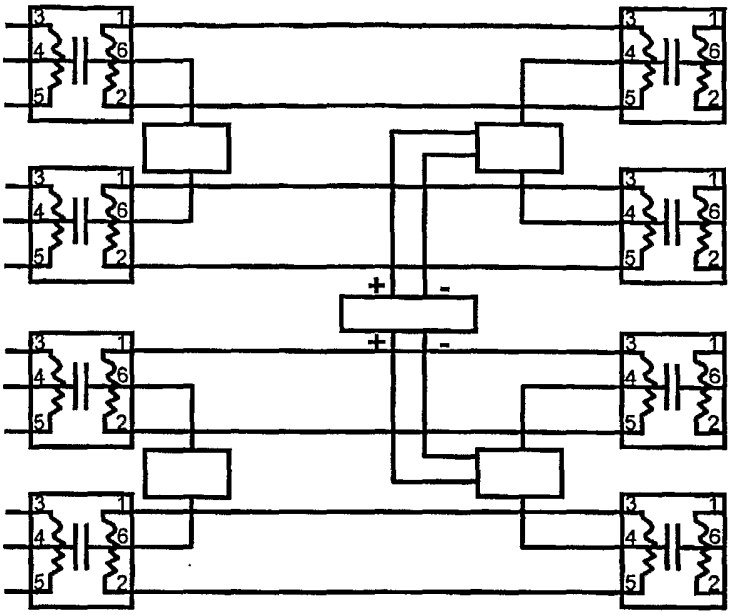


Fig. 1D
Prior Art

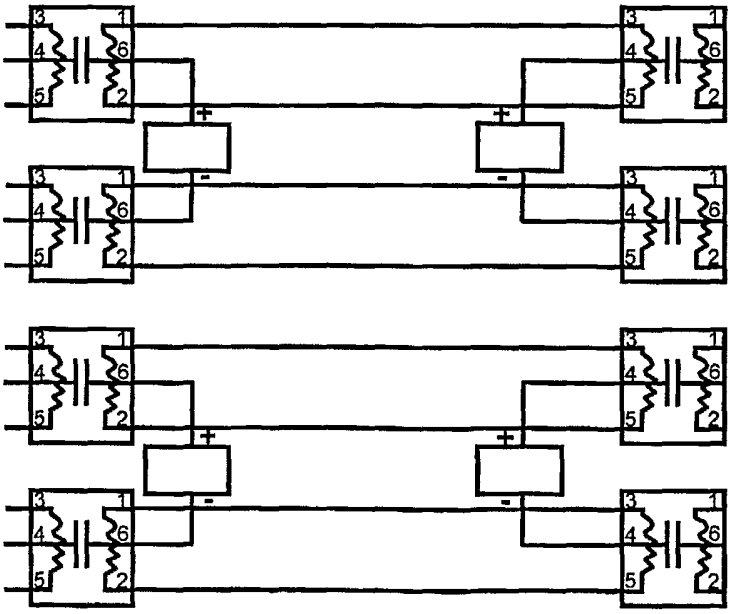


Fig. 1E
Prior Art

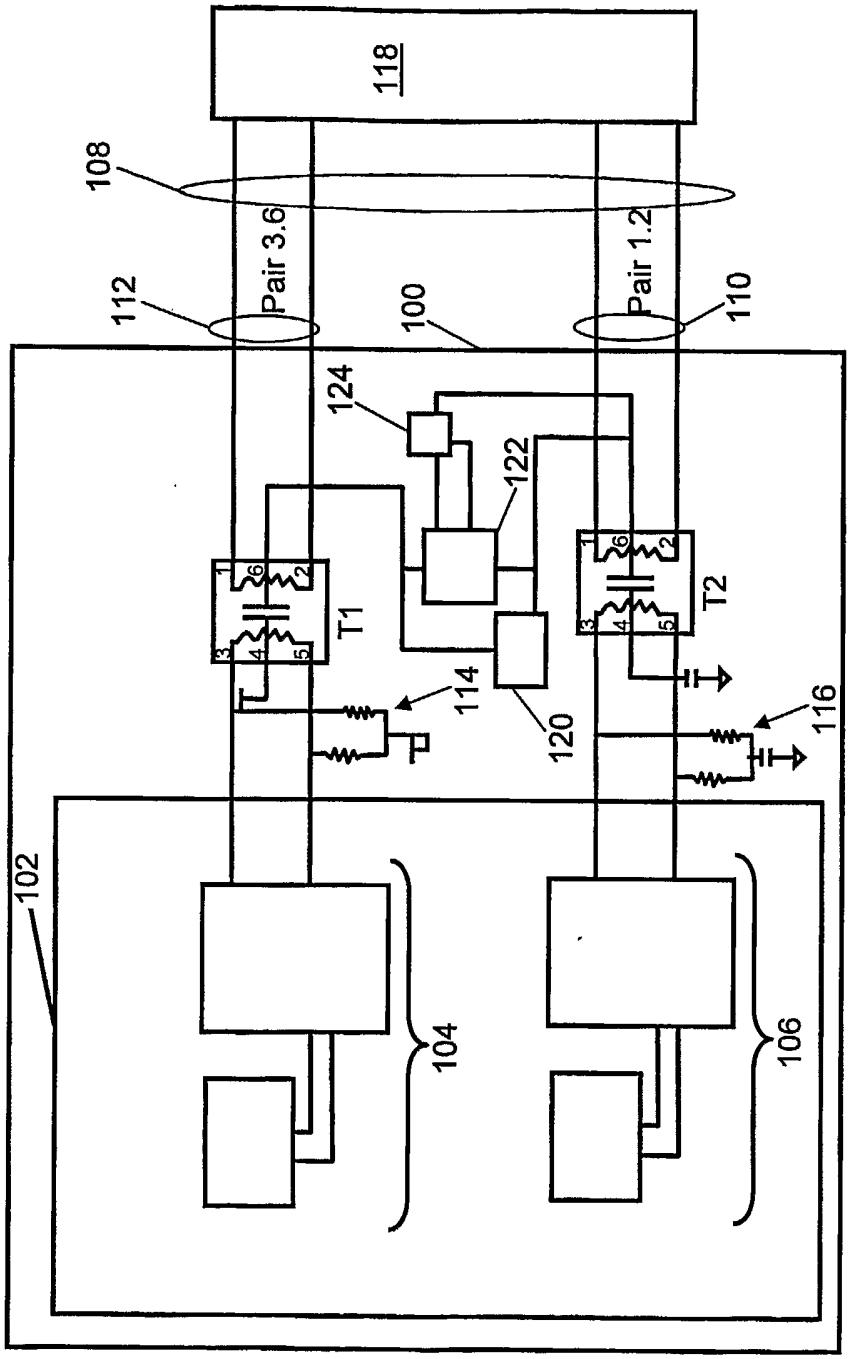


Fig. 2
Prior Art

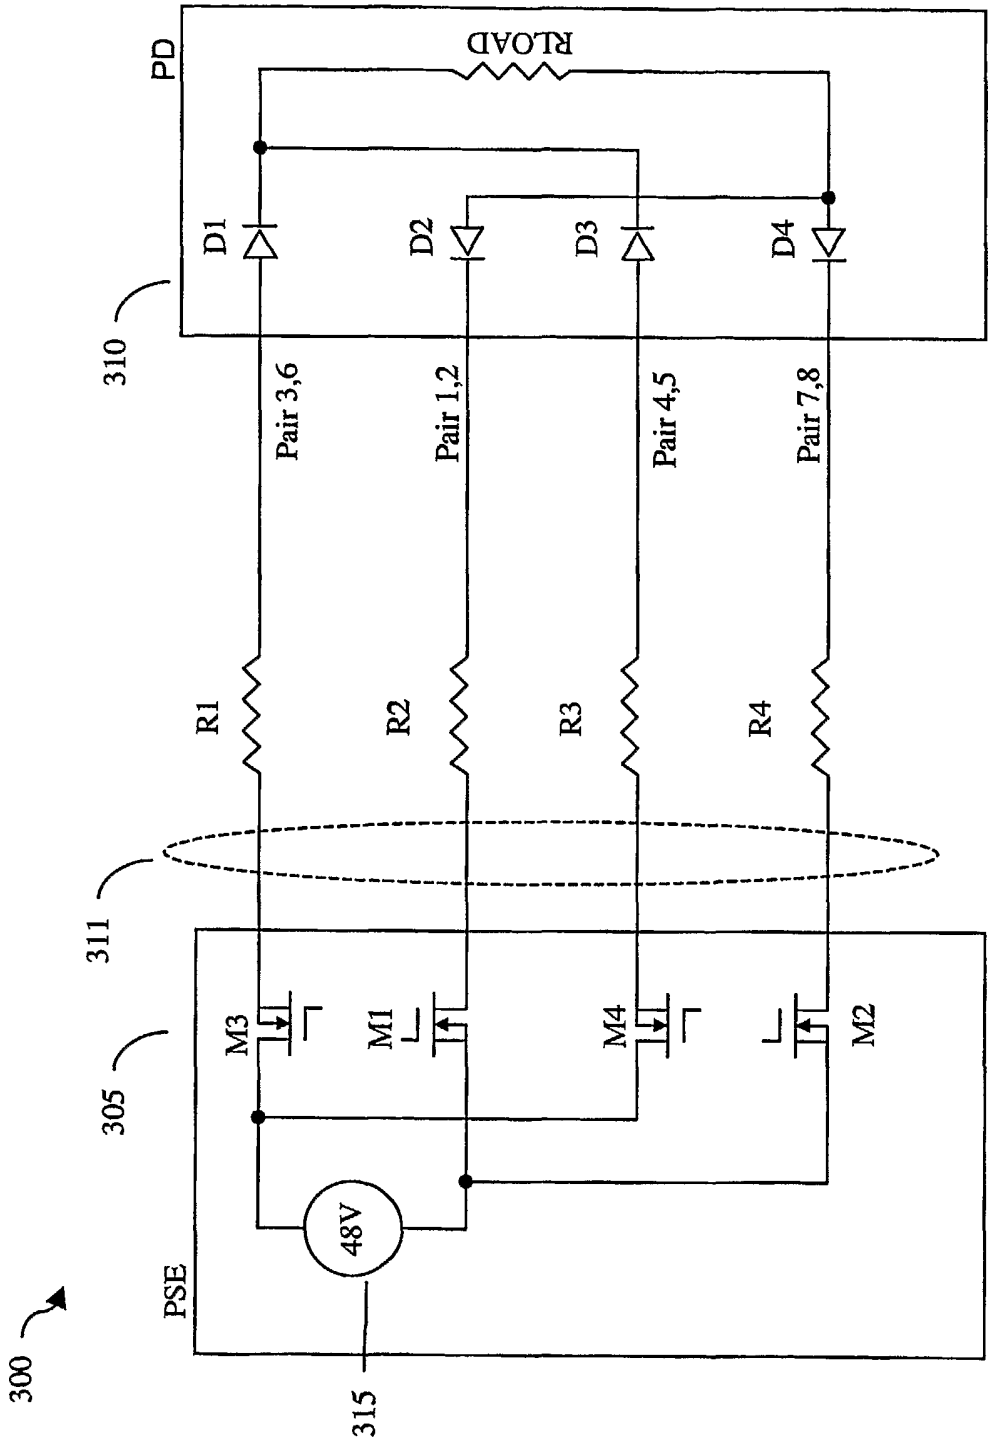

300
PSE
305
M3
M1
M4
M2
48V
315
311
R1
R2
R3
R4
Pair 3,6
Pair 1,2
Pair 4,5
Pair 7,8
310
PD
RLOAD
D1
D2
D3
D4

FIG. 3

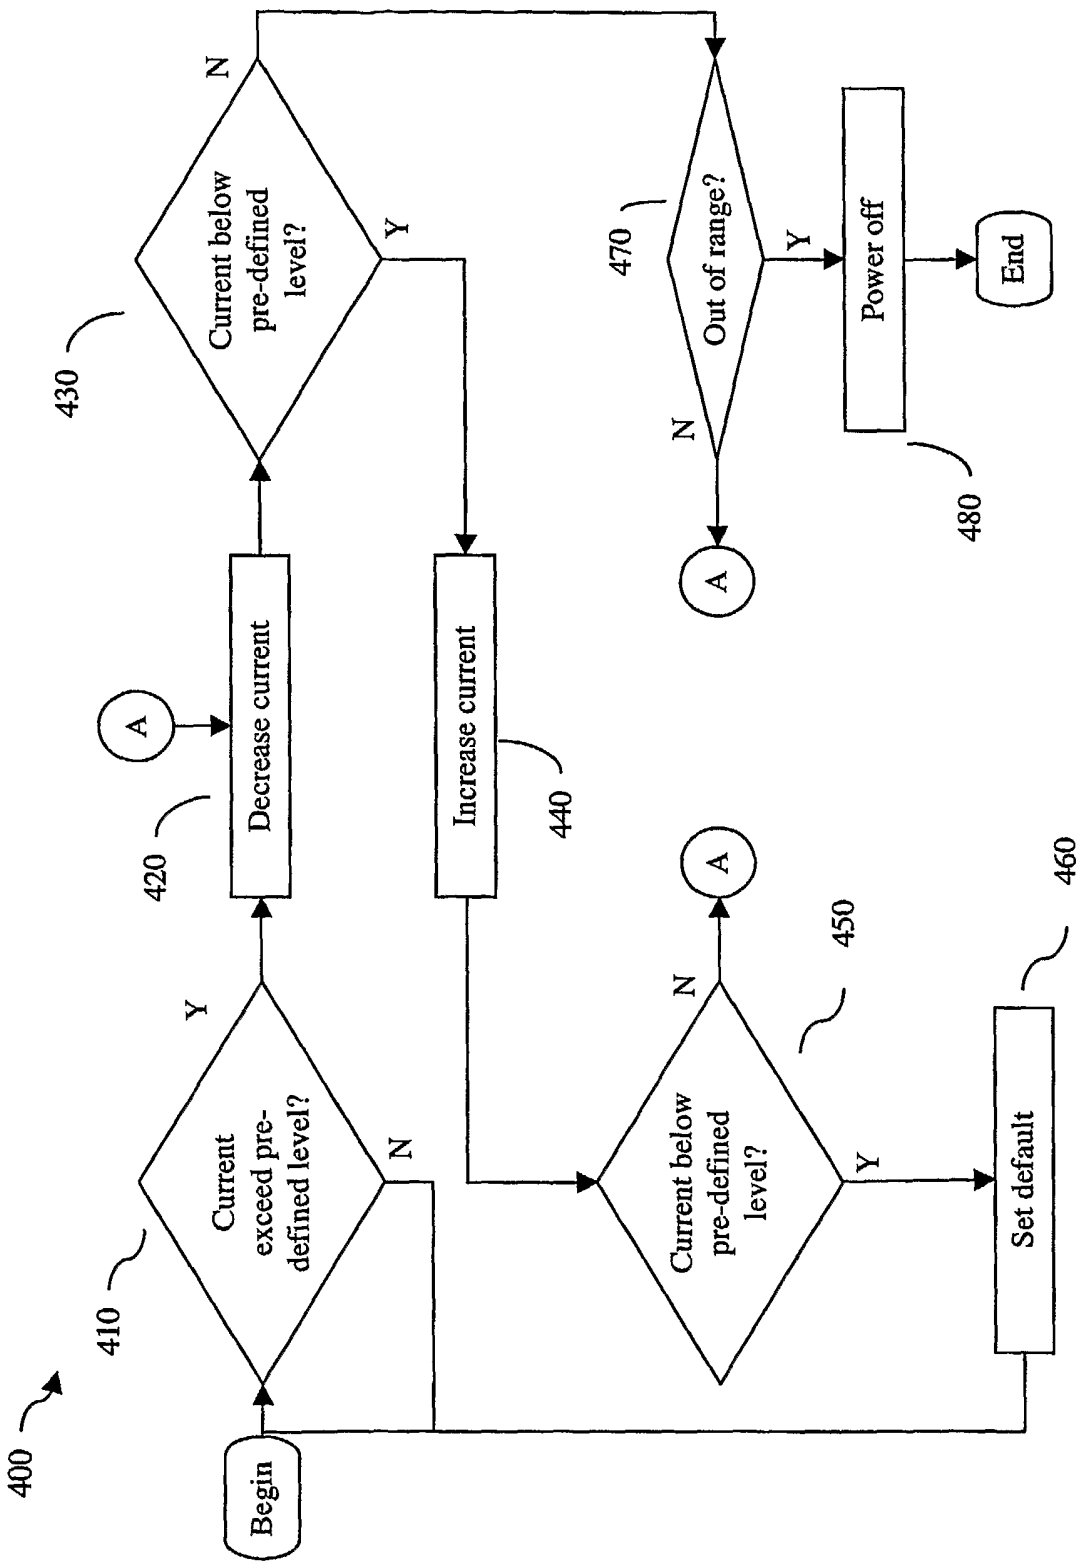

400
Begin
410
Current exceed pre-defined level?
Y
N
420
Decrease current
430
Current below pre-defined level?
N
Y
440
Increase current
450
Current below pre-defined level?
N
Y
460
Set default
470
Out of range?
N
Y
480
Power off
End
A
A
A

FIG. 4

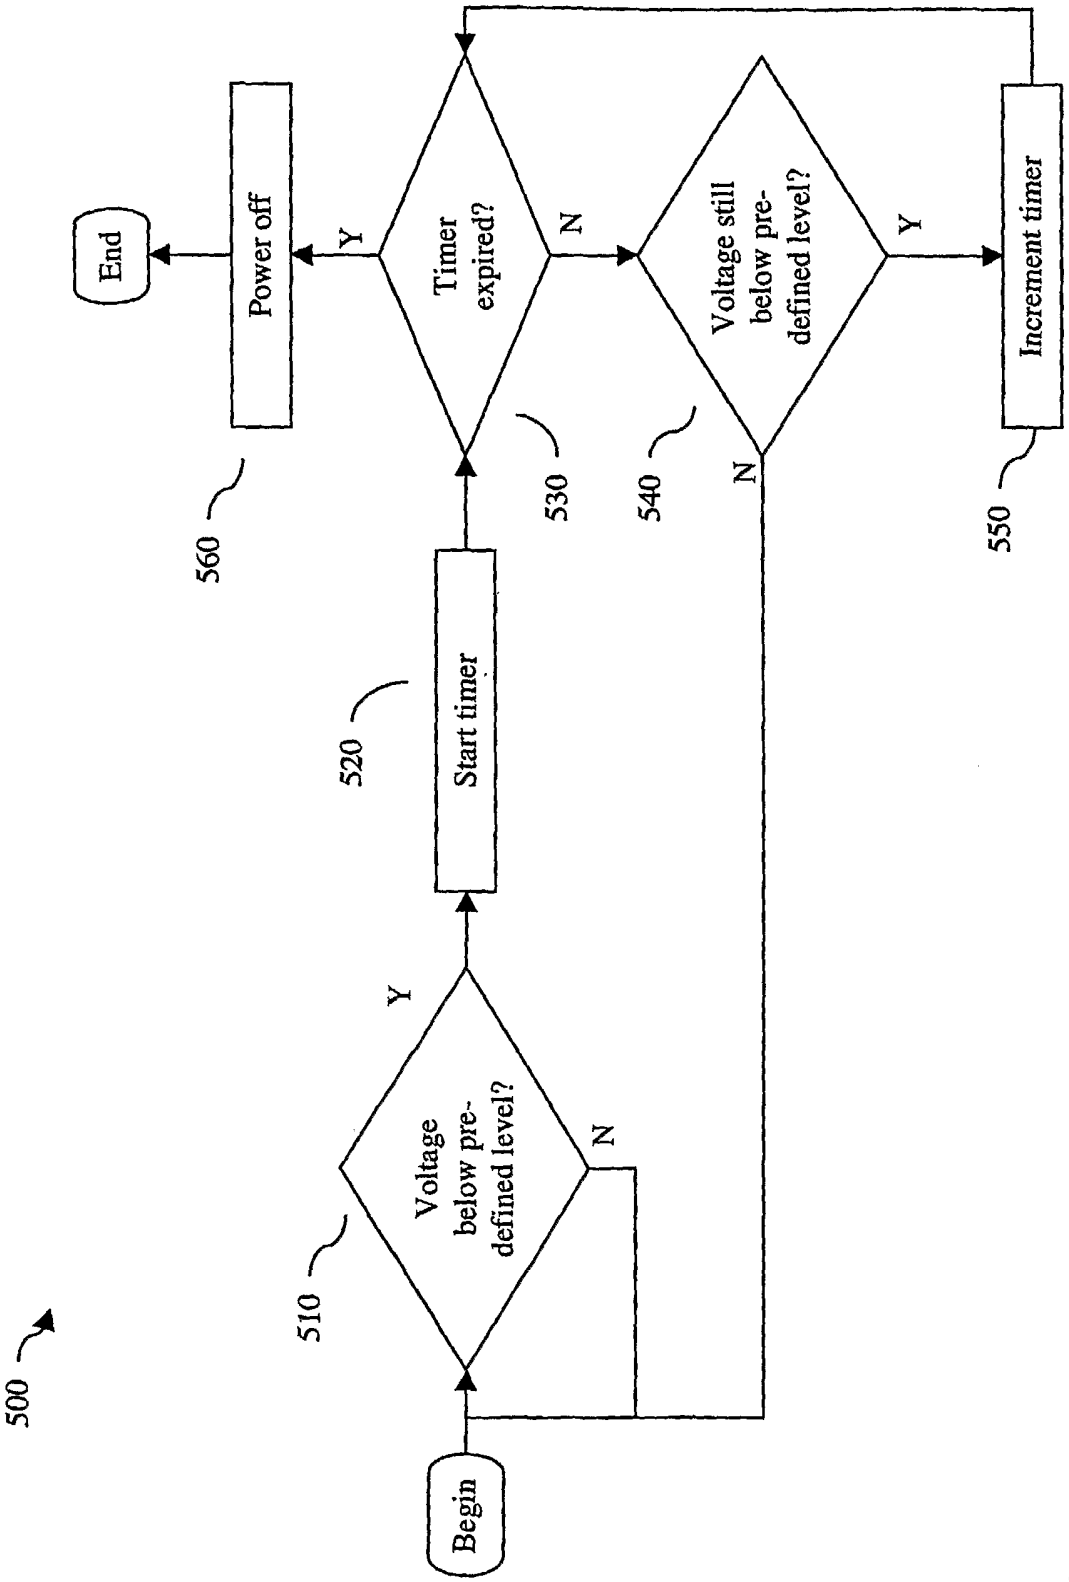

500
Begin
Voltage below pre-defined level?
510
Y
N
Start timer
520
Timer expired?
530
Voltage still below pre-defined level?
540
Increment timer
550
Power off
560
End

FIG. 5

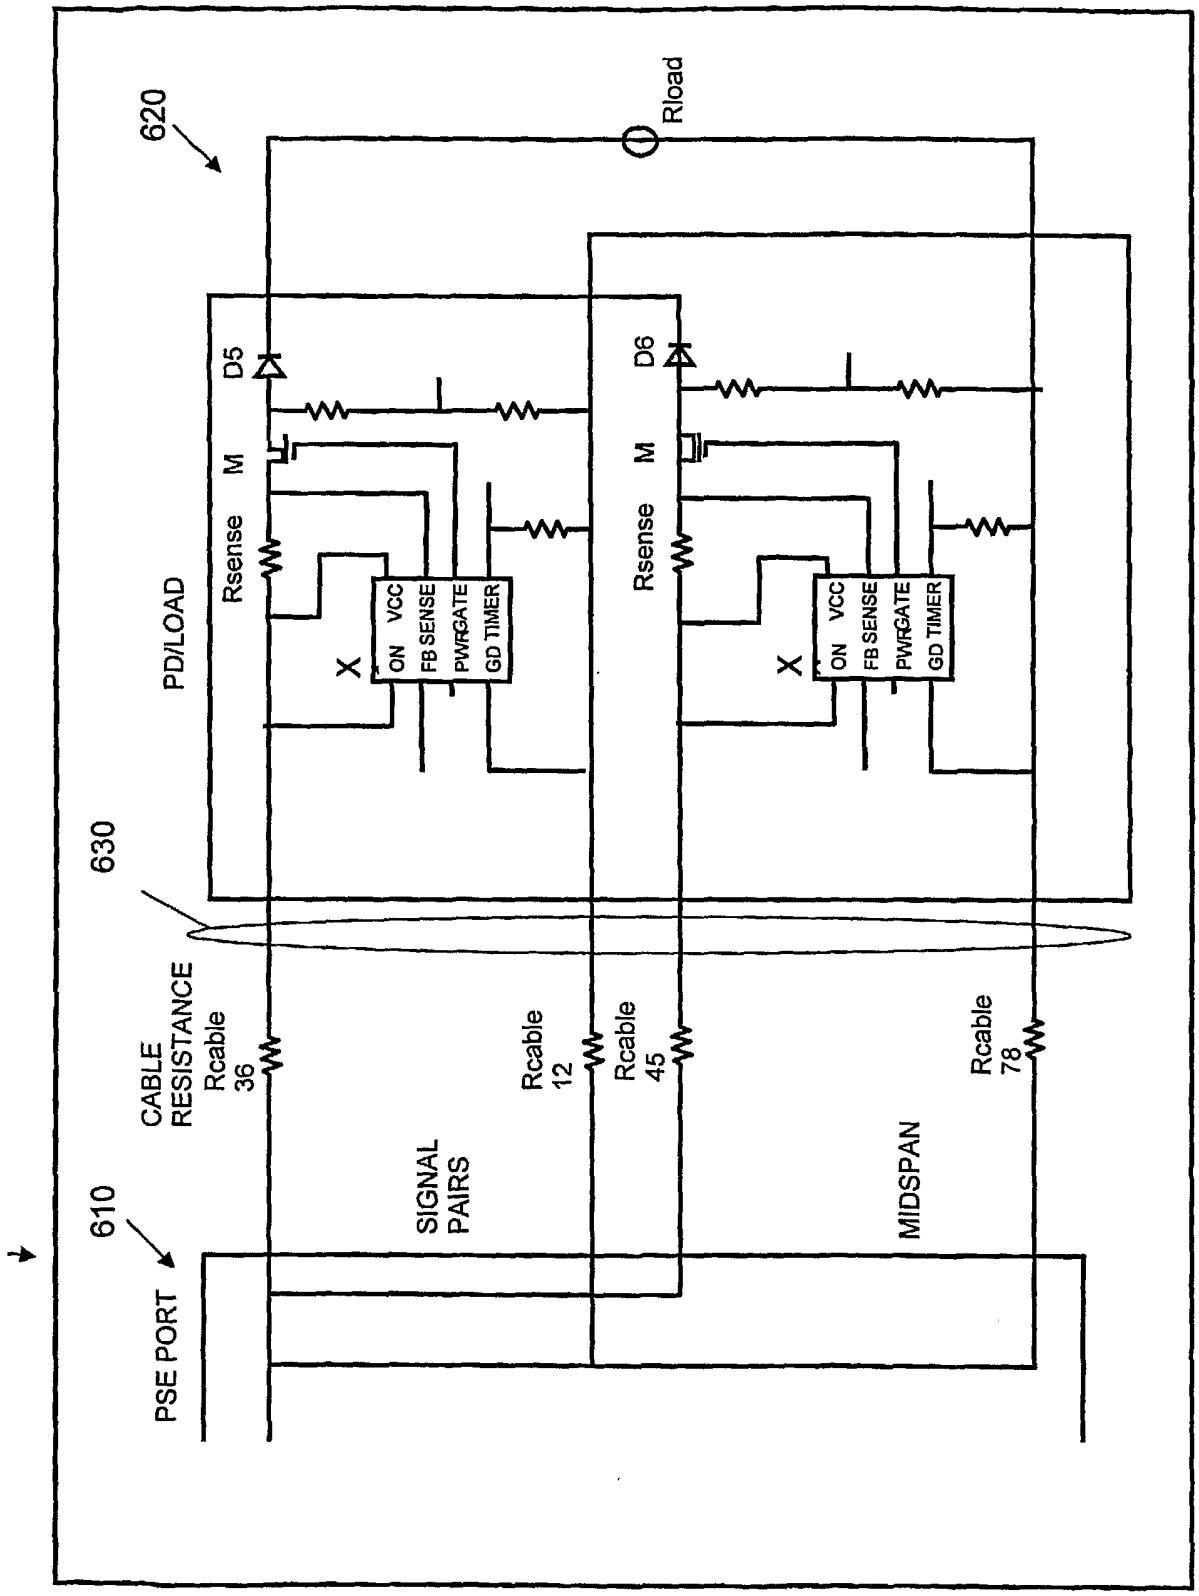

620
Rload
D5
D6
M
M
Rsense
Rsense
PD/LOAD
X
X
VCC
ON
FB SENSE
PWRGATE
GD TIMER
630
CABLE RESISTANCE
Rcable 36
Rcable 12
Rcable 45
Rcable 78
SIGNAL PAIRS
MIDSPAN
610
PSE PORT

FIG. 6

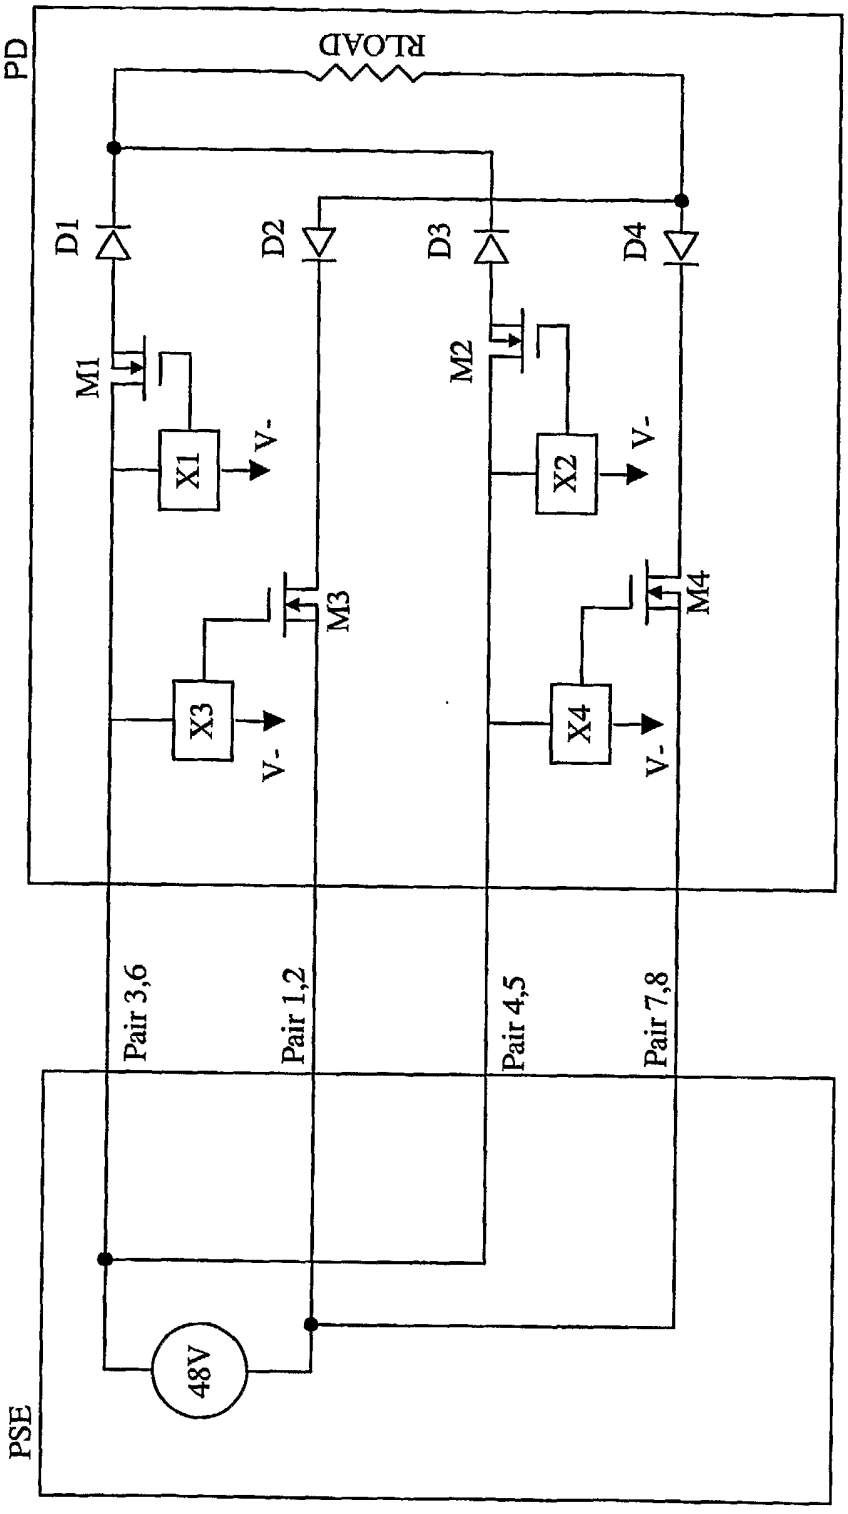

PD
RLOAD
D1
D2
D3
D4
M1
M2
M3
M4
X1
X2
X3
X4
V-
V-
V-
V-
Pair 3,6
Pair 1,2
Pair 4,5
Pair 7,8
48V
PSE

FIG. 7

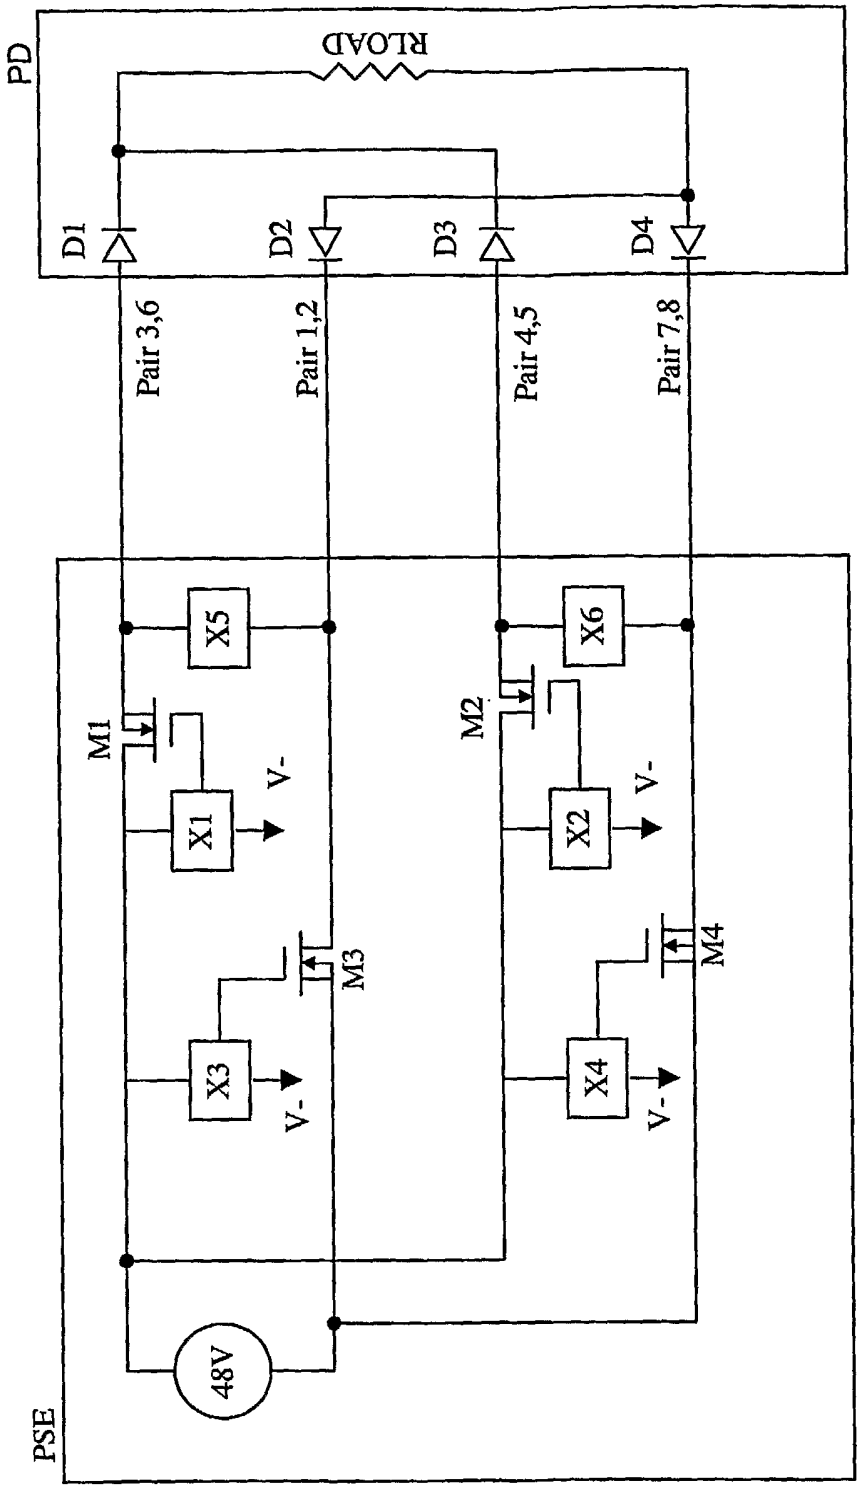

PD
RLOAD
D1
D2
D3
D4
Pair 3,6
Pair 1,2
Pair 4,5
Pair 7,8
X5
X6
M1
M2
V-
X1
X2
M3
M4
X3
X4
48V
PSE

FIG. 8

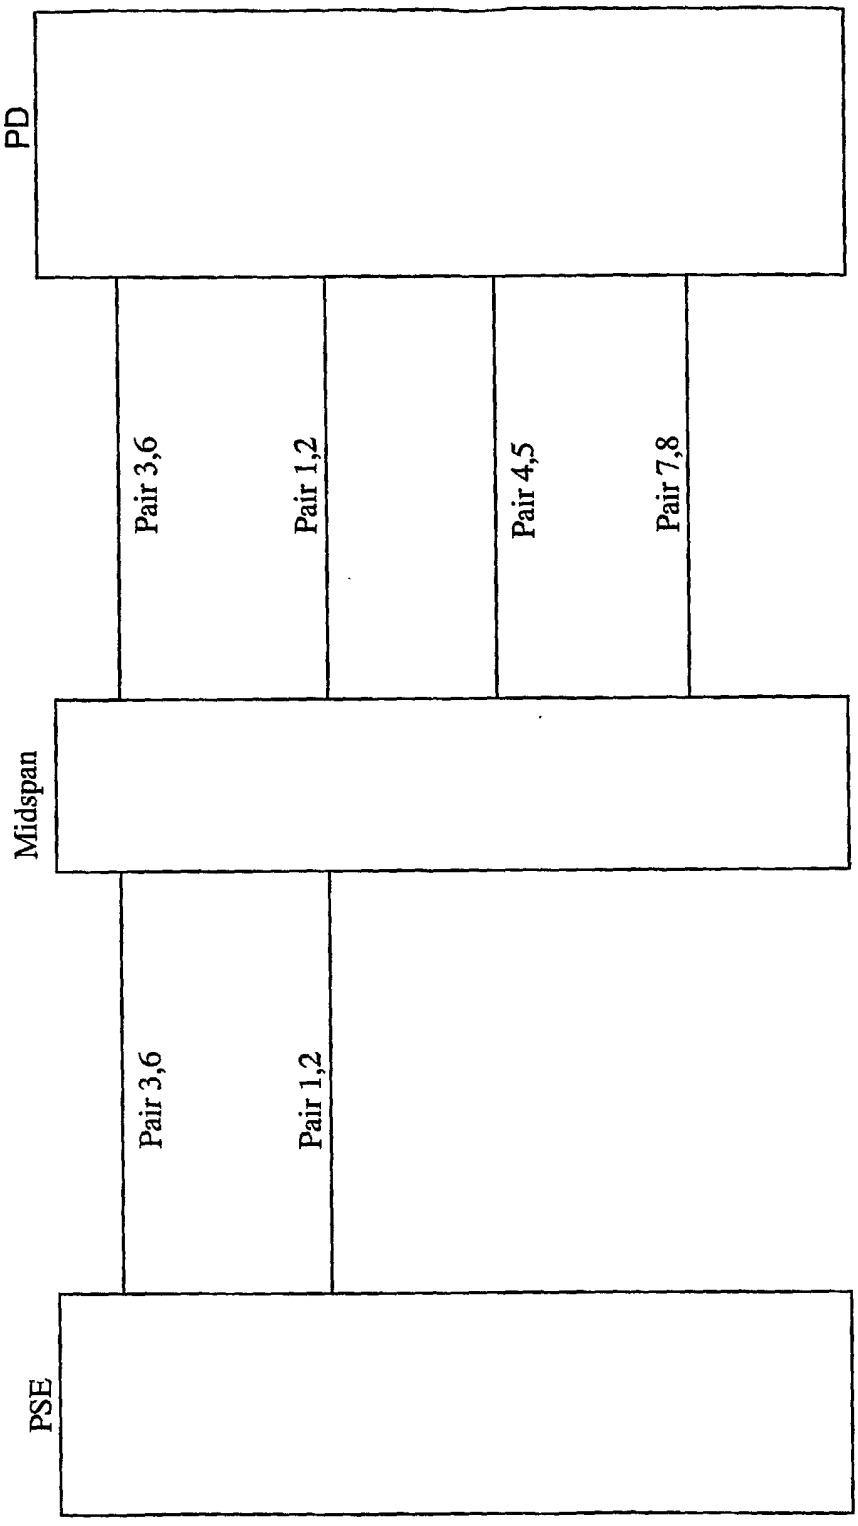


FIG. 9