US 20070044493A1

(19) United States

(12) Patent Application Publication

Kearney et al.

(10) Pub. No.: US 2007/0044493 A1

(43) Pub. Date: Mar. 1, 2007

(54) SYSTEMS AND METHODS FOR COOLING ELECTRONICS COMPONENTS EMPLOYING VAPOR COMPRESSION REFRIGERATION WITH SELECTED PORTIONS OF EXPANSION STRUCTURES COATED WITH POLYTETRAFLUORETHYLENE

(75) Inventors: **Daniel J. Kearney**, Ulster Park, NY (US); **Mark A. Marnell**, Kingston, NY (US); **Donald W. Porter**, Highland, NY (US)

Correspondence Address:
HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 COLUMBIA CIRCLE
ALBANY, NY 12203 (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(21) Appl. No.: 11/209,241

(22) Filed: Aug. 23, 2005

Publication Classification

(51) Int. Cl.
F25B 41/06 (2006.01)
F25D 23/12 (2006.01)

(52) U.S. Cl. .... **62/259.2**; 236/92 B

(57) ABSTRACT

Systems and Methods of cooling heat generating electronics components are provided employing vapor compression refrigeration. In one embodiment, the vapor compression refrigeration system includes a condenser, at least one expansion structure, at least one evaporator, and a compressor coupled in fluid communication to define a refrigerant flow path, and allow the flow of refrigerant therethrough. The at least one evaporator is coupled to the at least one heat generating electronics component to facilitate removal of heat produced by the electronics component. At least a portion of the at least one expansion structure is coated with a polytetrafluorethylene in the refrigerant flow path for inhibiting accumulation of material thereon. The polytetrafluorethylene coating has a thickness sufficient to inhibit accumulation of material in a pressure drop area of the expansion structure without significantly changing a pressure drop characteristic of the pressure drop area.

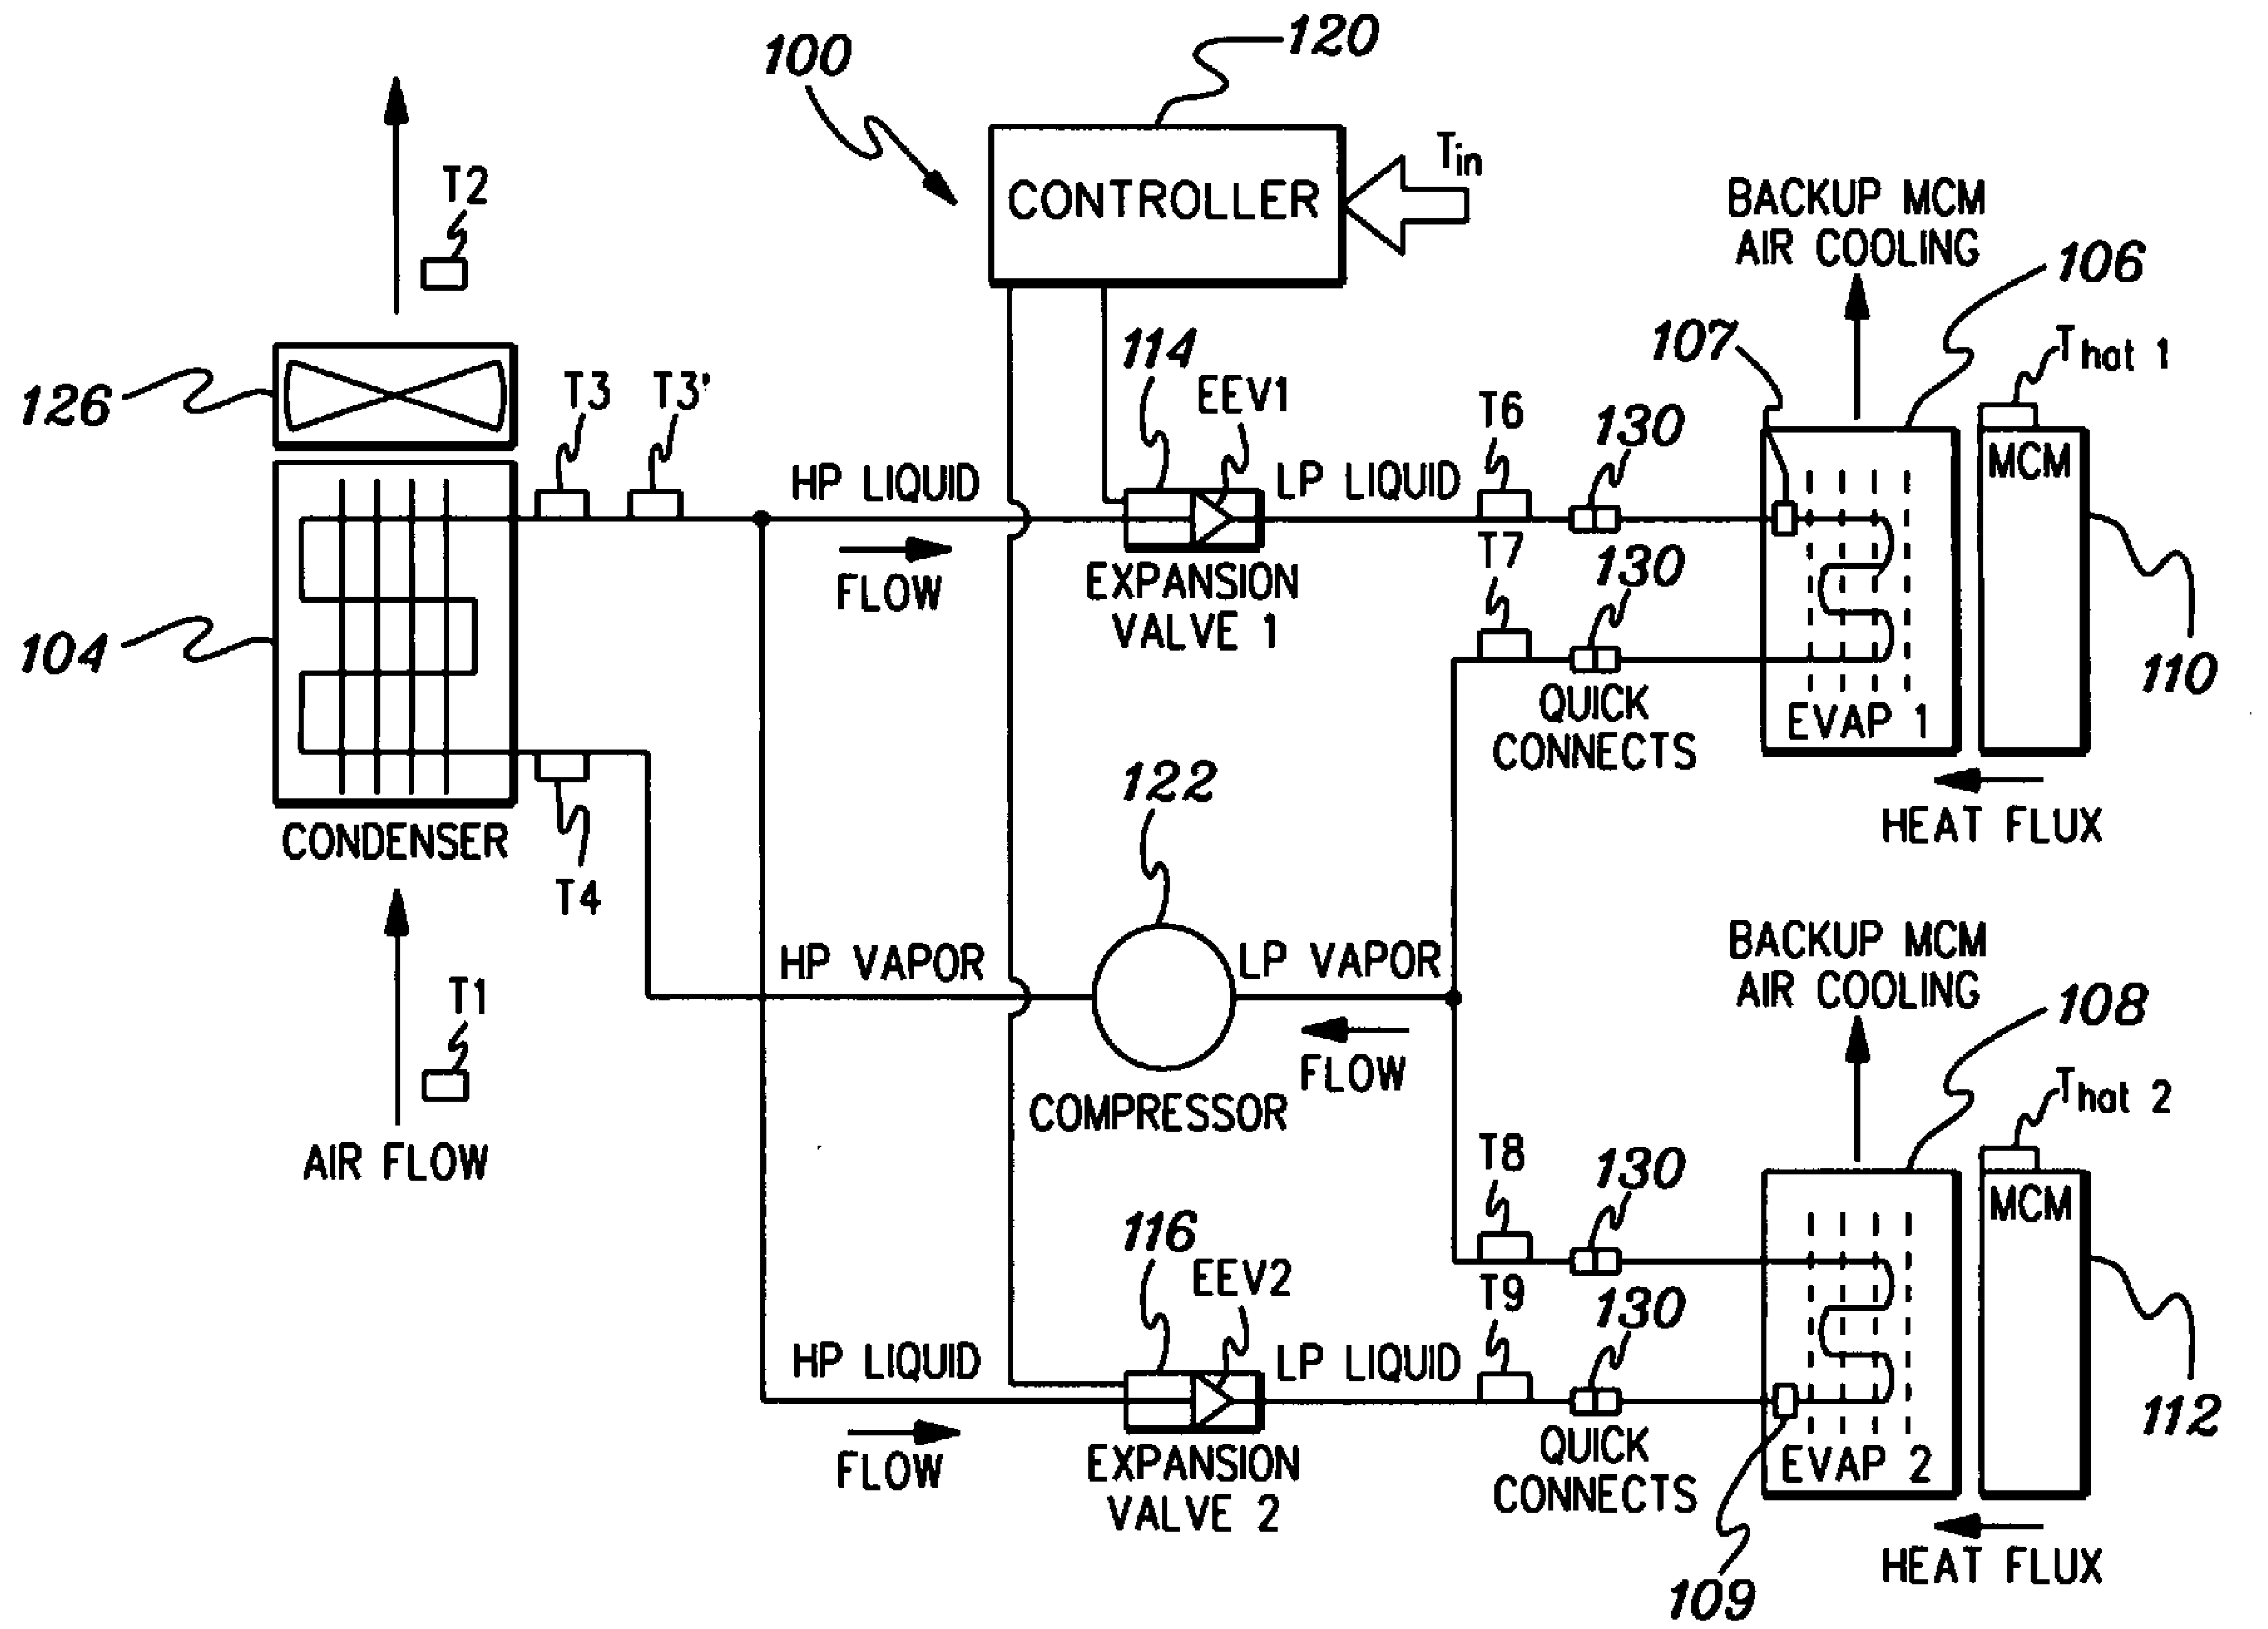

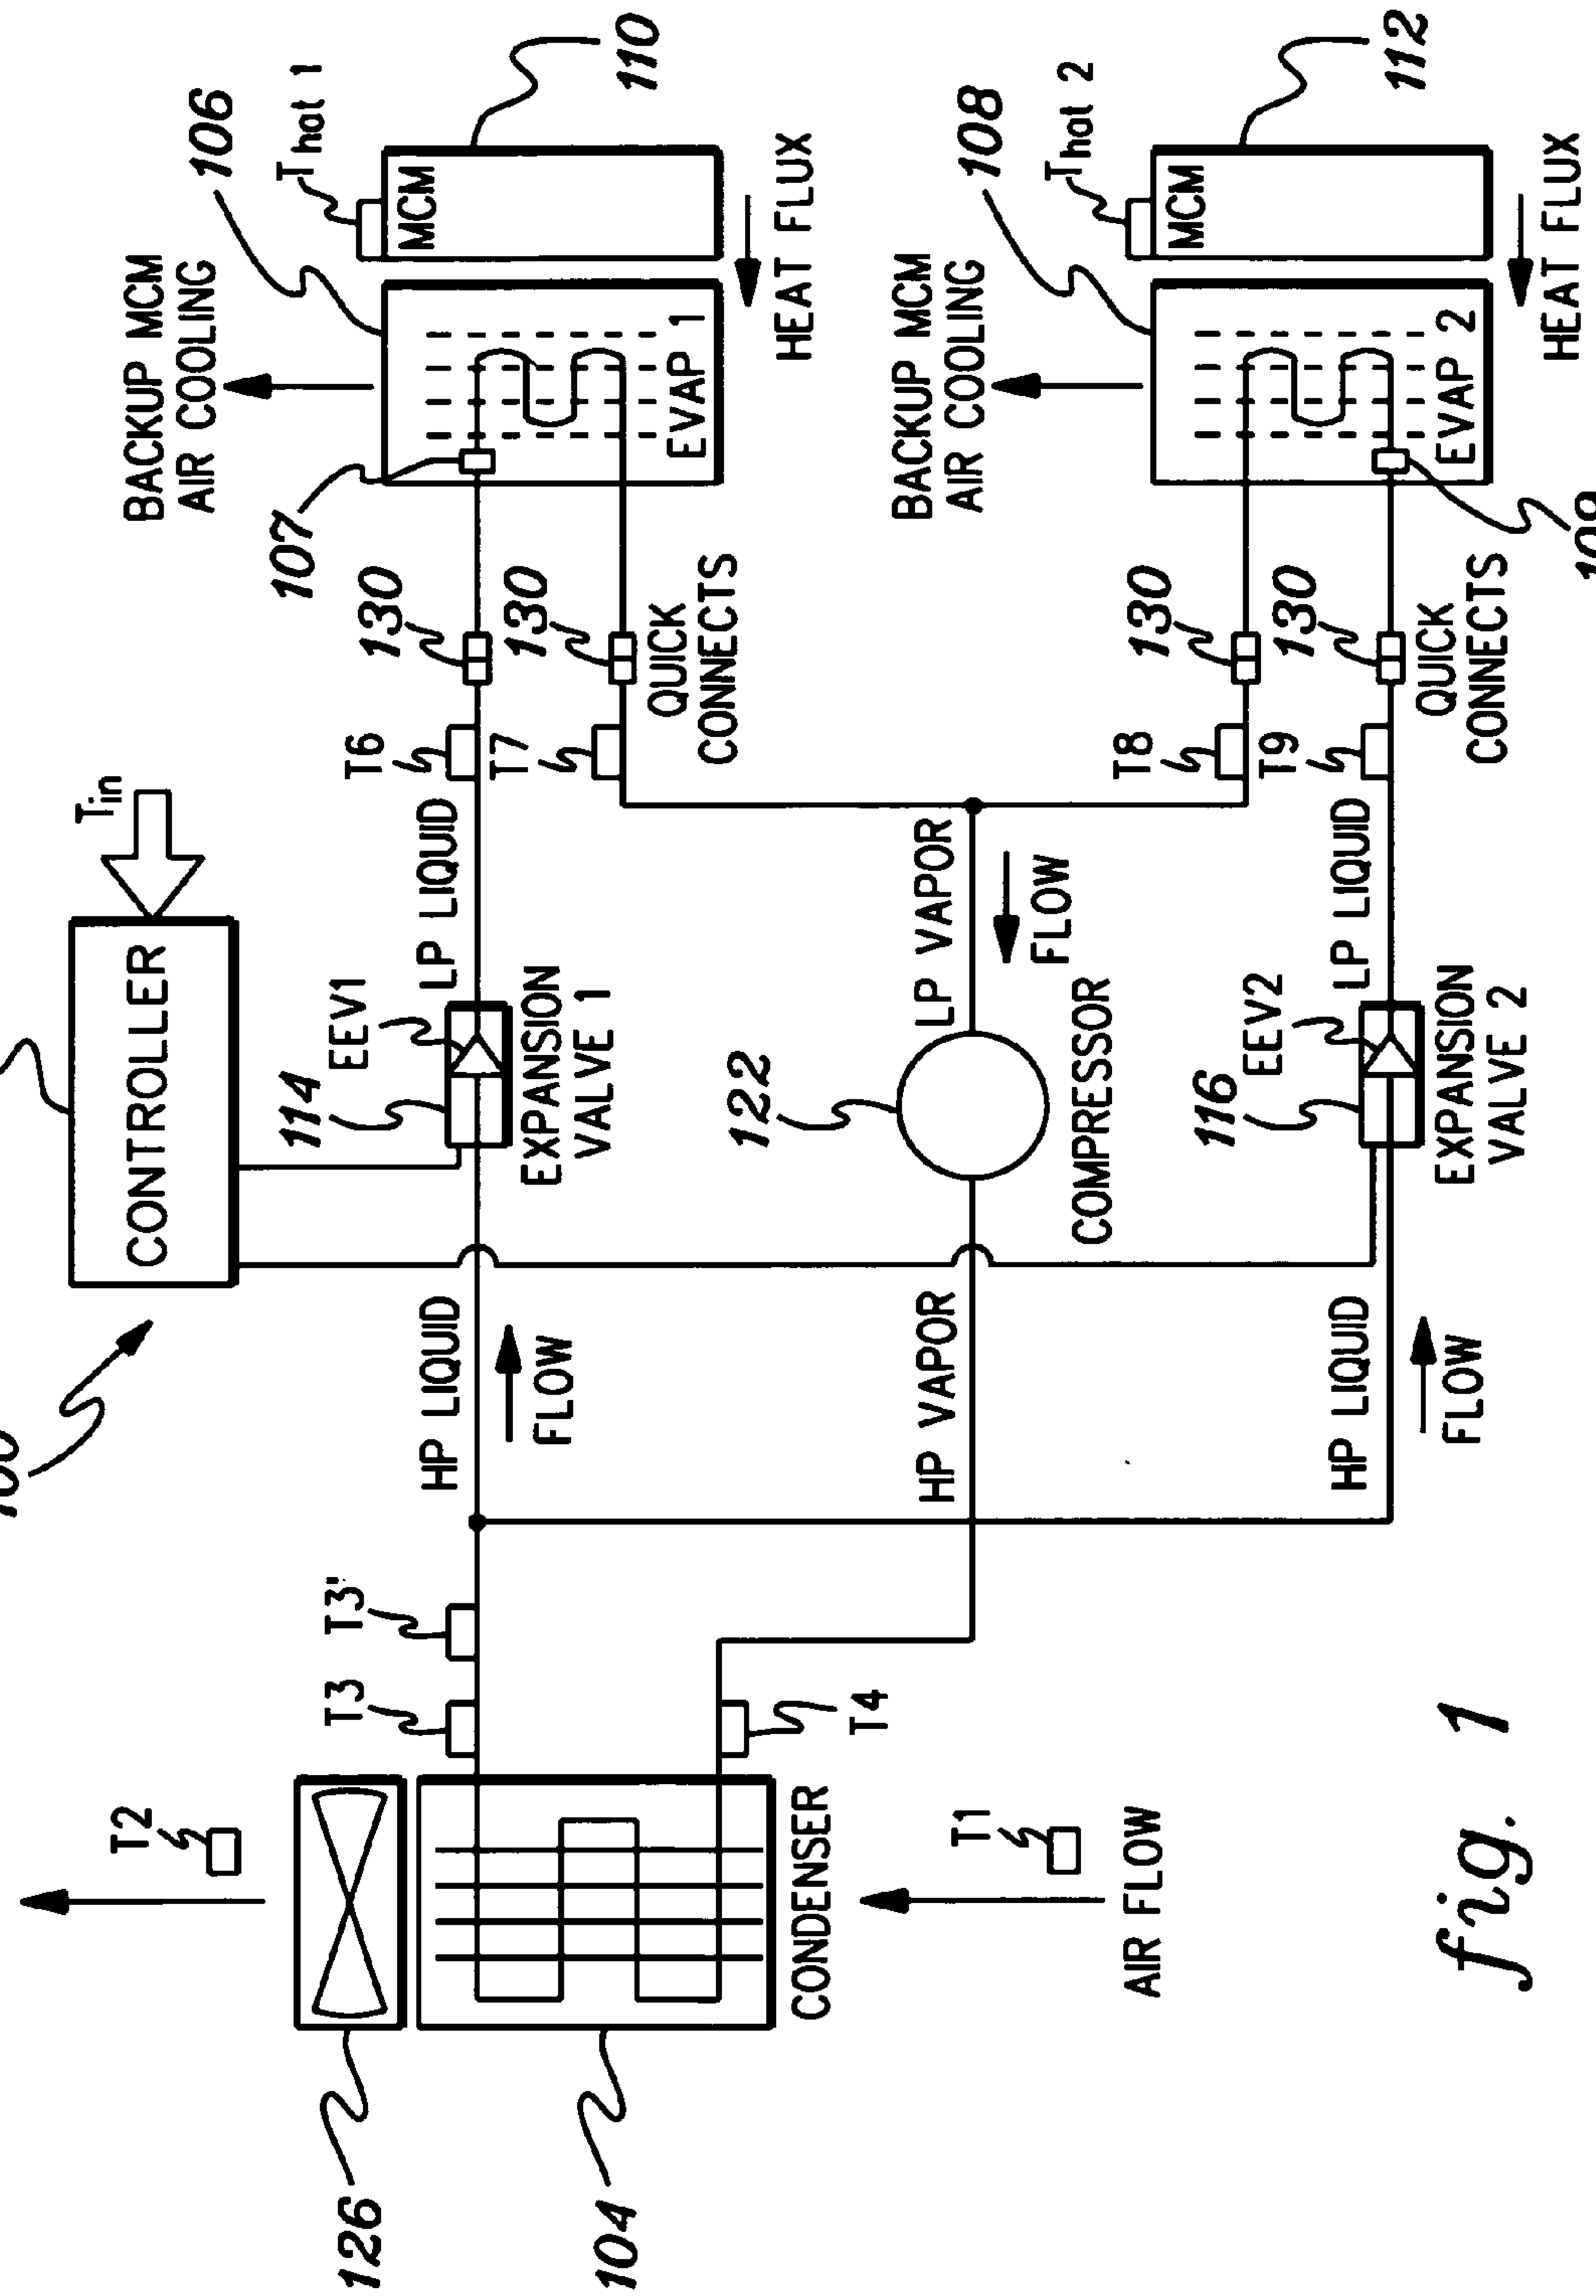

100
120
CONTROLLER
T_{in}
BACKUP MCM AIR COOLING
106
T_{hot} 1
110
MCM
EVAP 1
HEAT FLUX
107
108
T_{hot} 2
112
EVAP 2
109
130
QUICK CONNECTS
T6
T7
T8
T9
LP LIQUID
LP VAPOR
FLOW
EEV1
114
EXPANSION VALVE 1
EEV2
116
EXPANSION VALVE 2
122
COMPRESSOR
HP LIQUID
HP VAPOR
T3
T3'
T4
T2
T1
126
104
CONDENSER
AIR FLOW

fig. 1

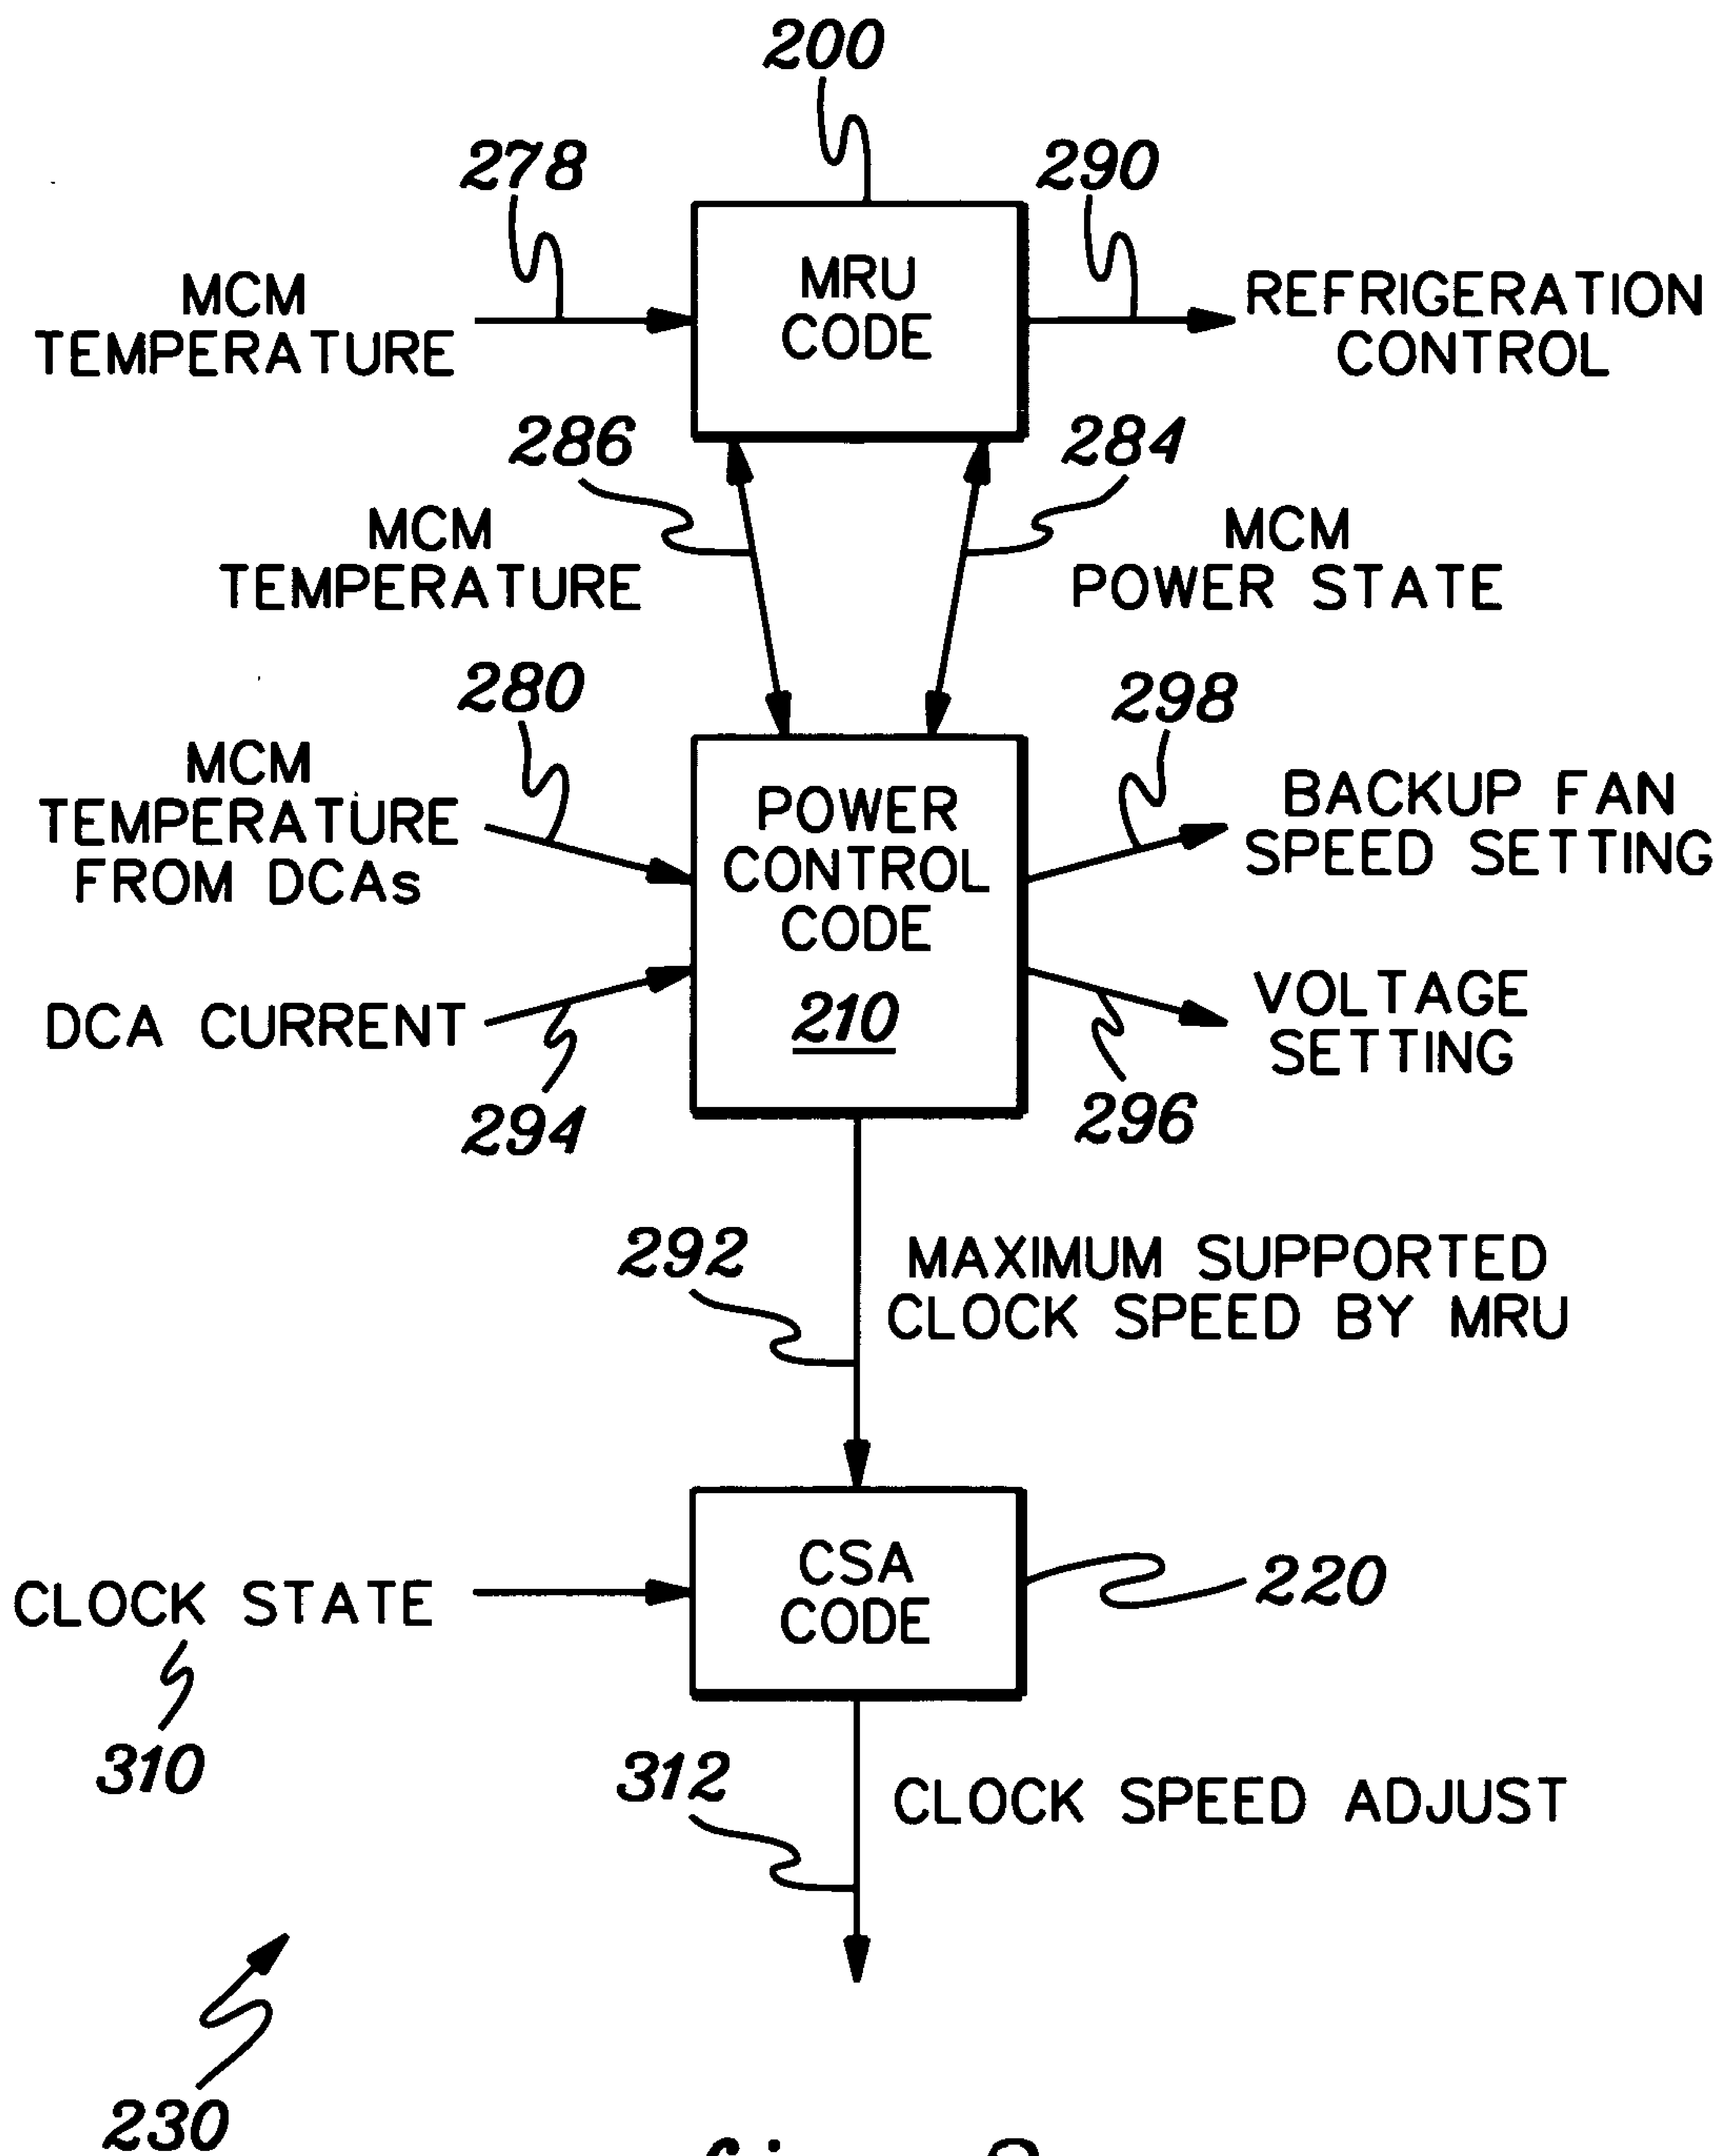
200
278
290
MCM
TEMPERATURE
MRU
CODE
REFRIGERATION
CONTROL
286
284
MCM
TEMPERATURE
MCM
POWER STATE
280
298
MCM
TEMPERATURE
FROM DCAs
POWER
CONTROL
CODE
210
BACKUP FAN
SPEED SETTING
DCA CURRENT
VOLTAGE
SETTING
294
296
292
MAXIMUM SUPPORTED
CLOCK SPEED BY MRU
CLOCK STATE
CSA
CODE
220
310
312
CLOCK SPEED ADJUST
230

fig. 2

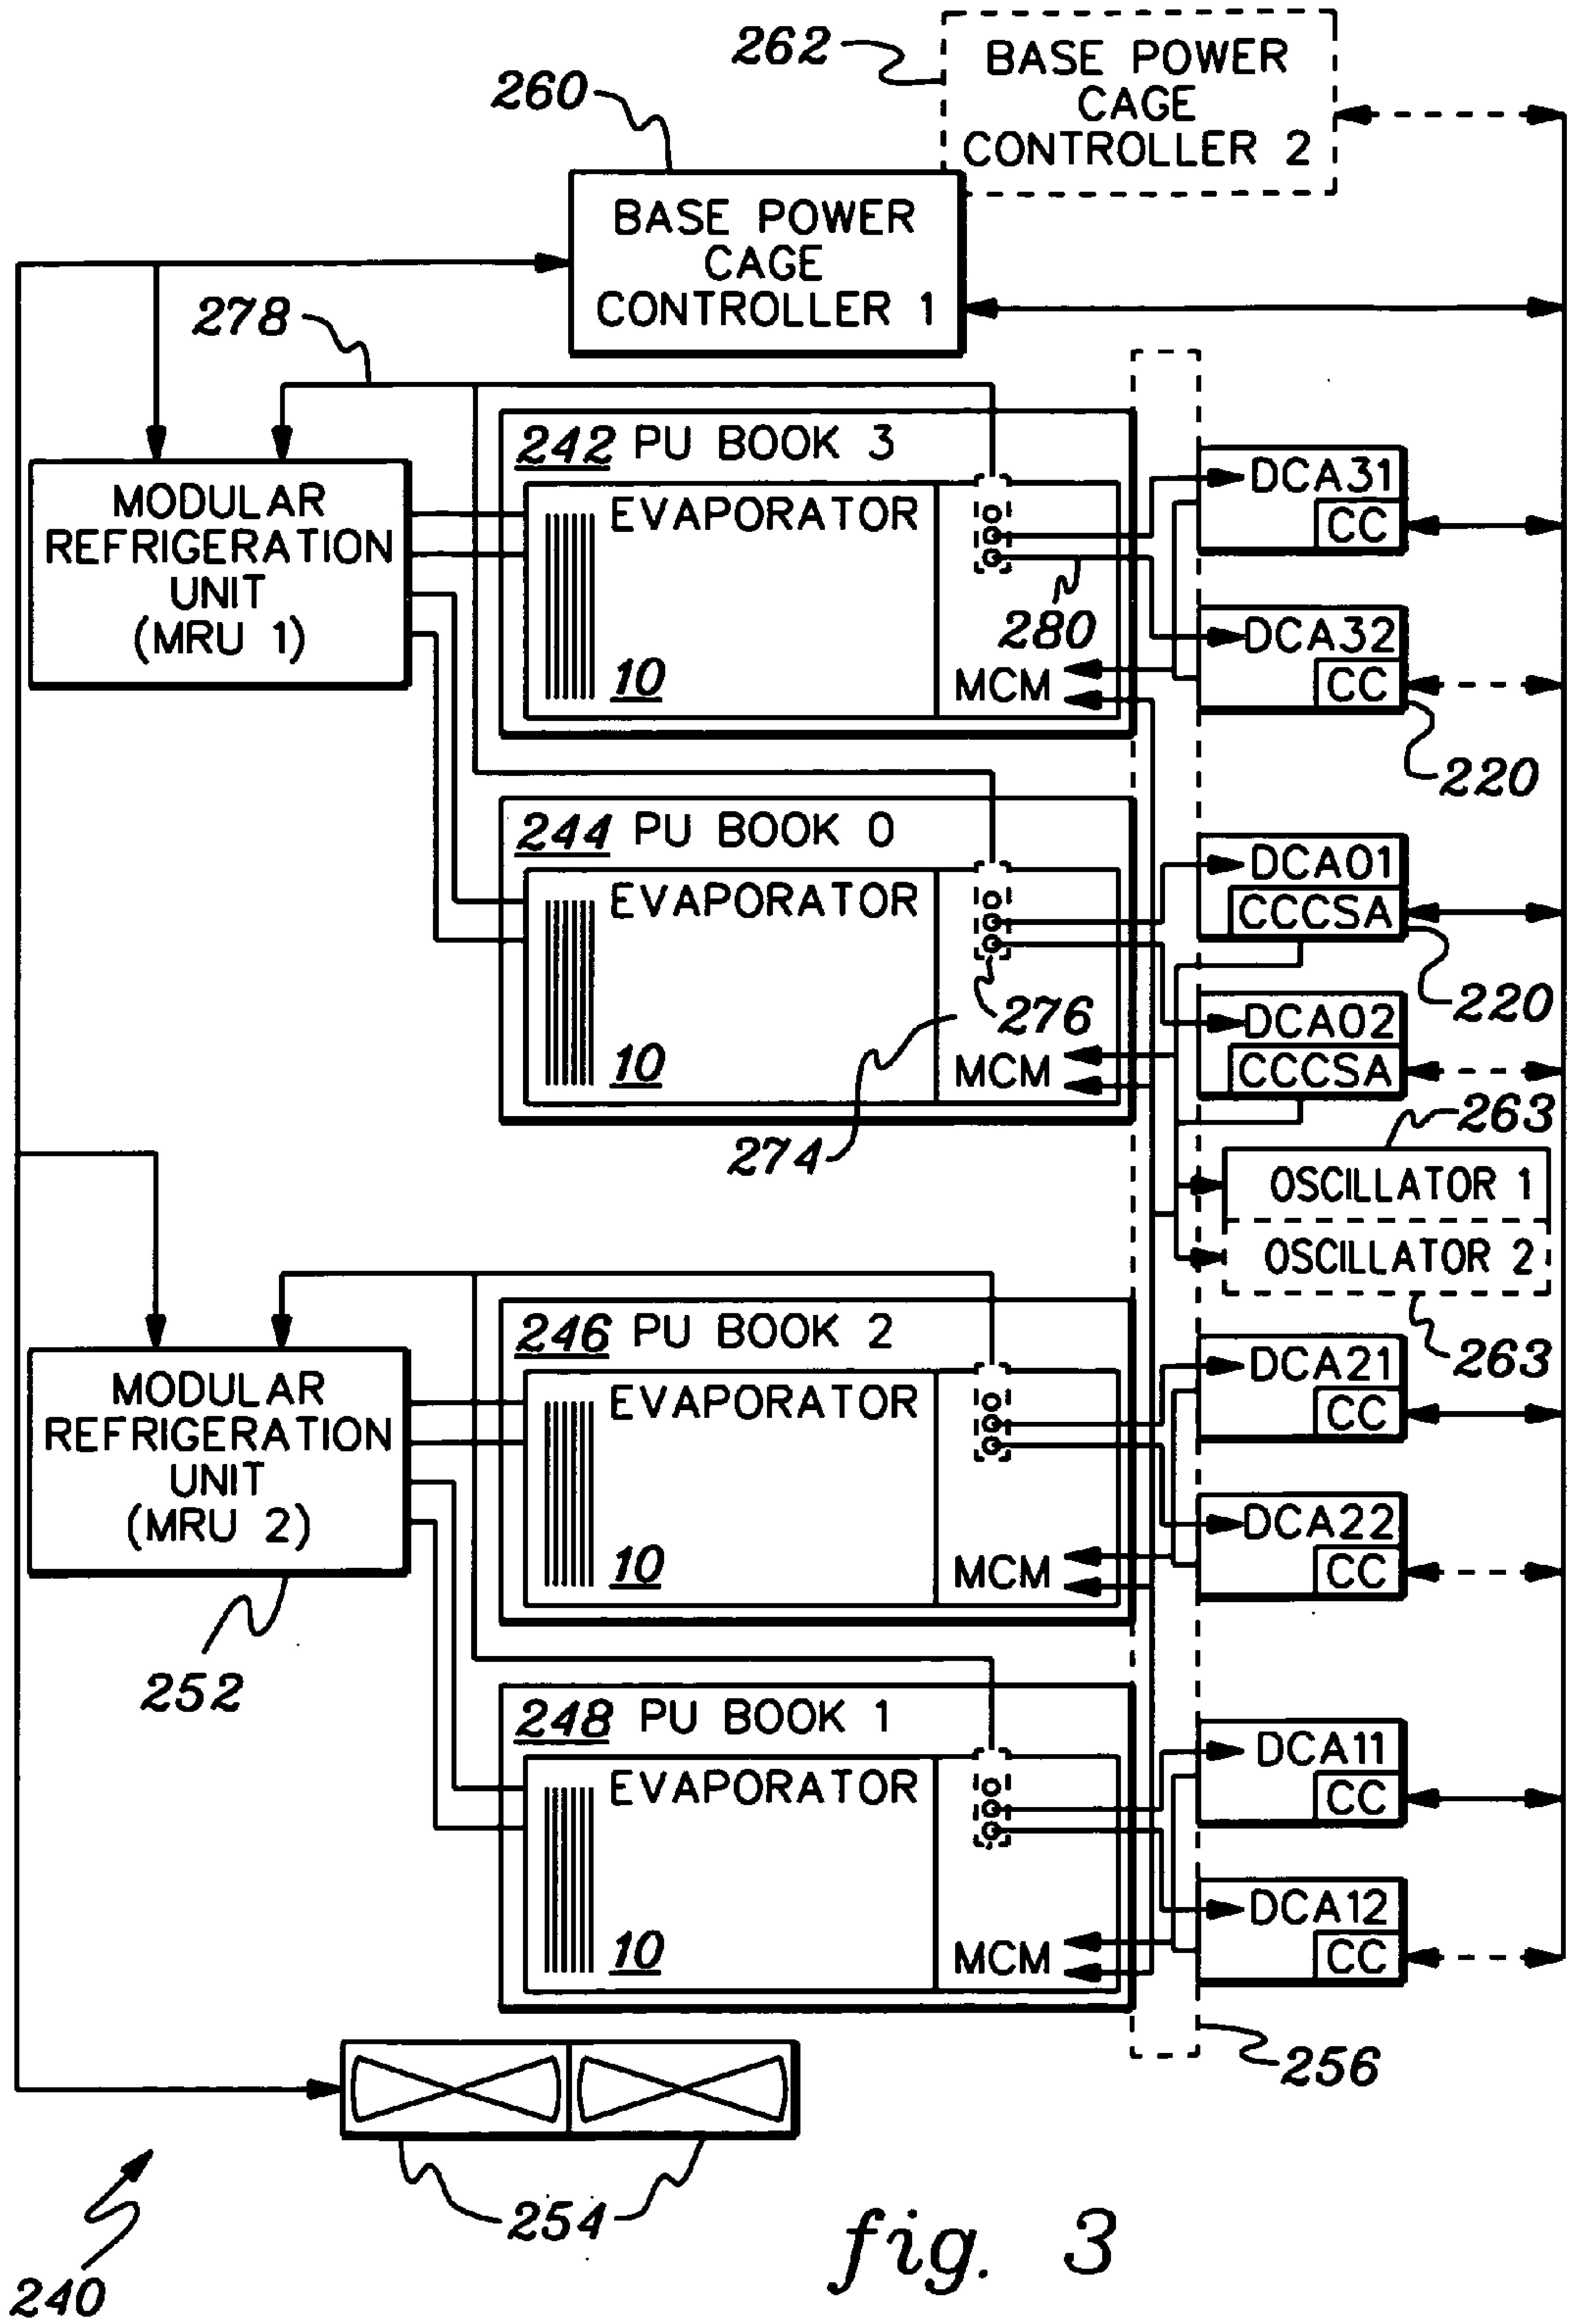
262
BASE POWER CAGE CONTROLLER 2
260
BASE POWER CAGE CONTROLLER 1
278
MODULAR REFRIGERATION UNIT (MRU 1)
242 PU BOOK 3
EVAPORATOR
10
MCM
280
DCA31
CC
DCA32
CC
220
244 PU BOOK 0
EVAPORATOR
10
MCM
276
274
DCA01
CCCSA
DCA02
CCCSA
220
263
OSCILLATOR 1
OSCILLATOR 2
263
MODULAR REFRIGERATION UNIT (MRU 2)
252
246 PU BOOK 2
EVAPORATOR
10
MCM
DCA21
CC
DCA22
CC
248 PU BOOK 1
EVAPORATOR
10
MCM
DCA11
CC
DCA12
CC
256
254
240

fig. 3

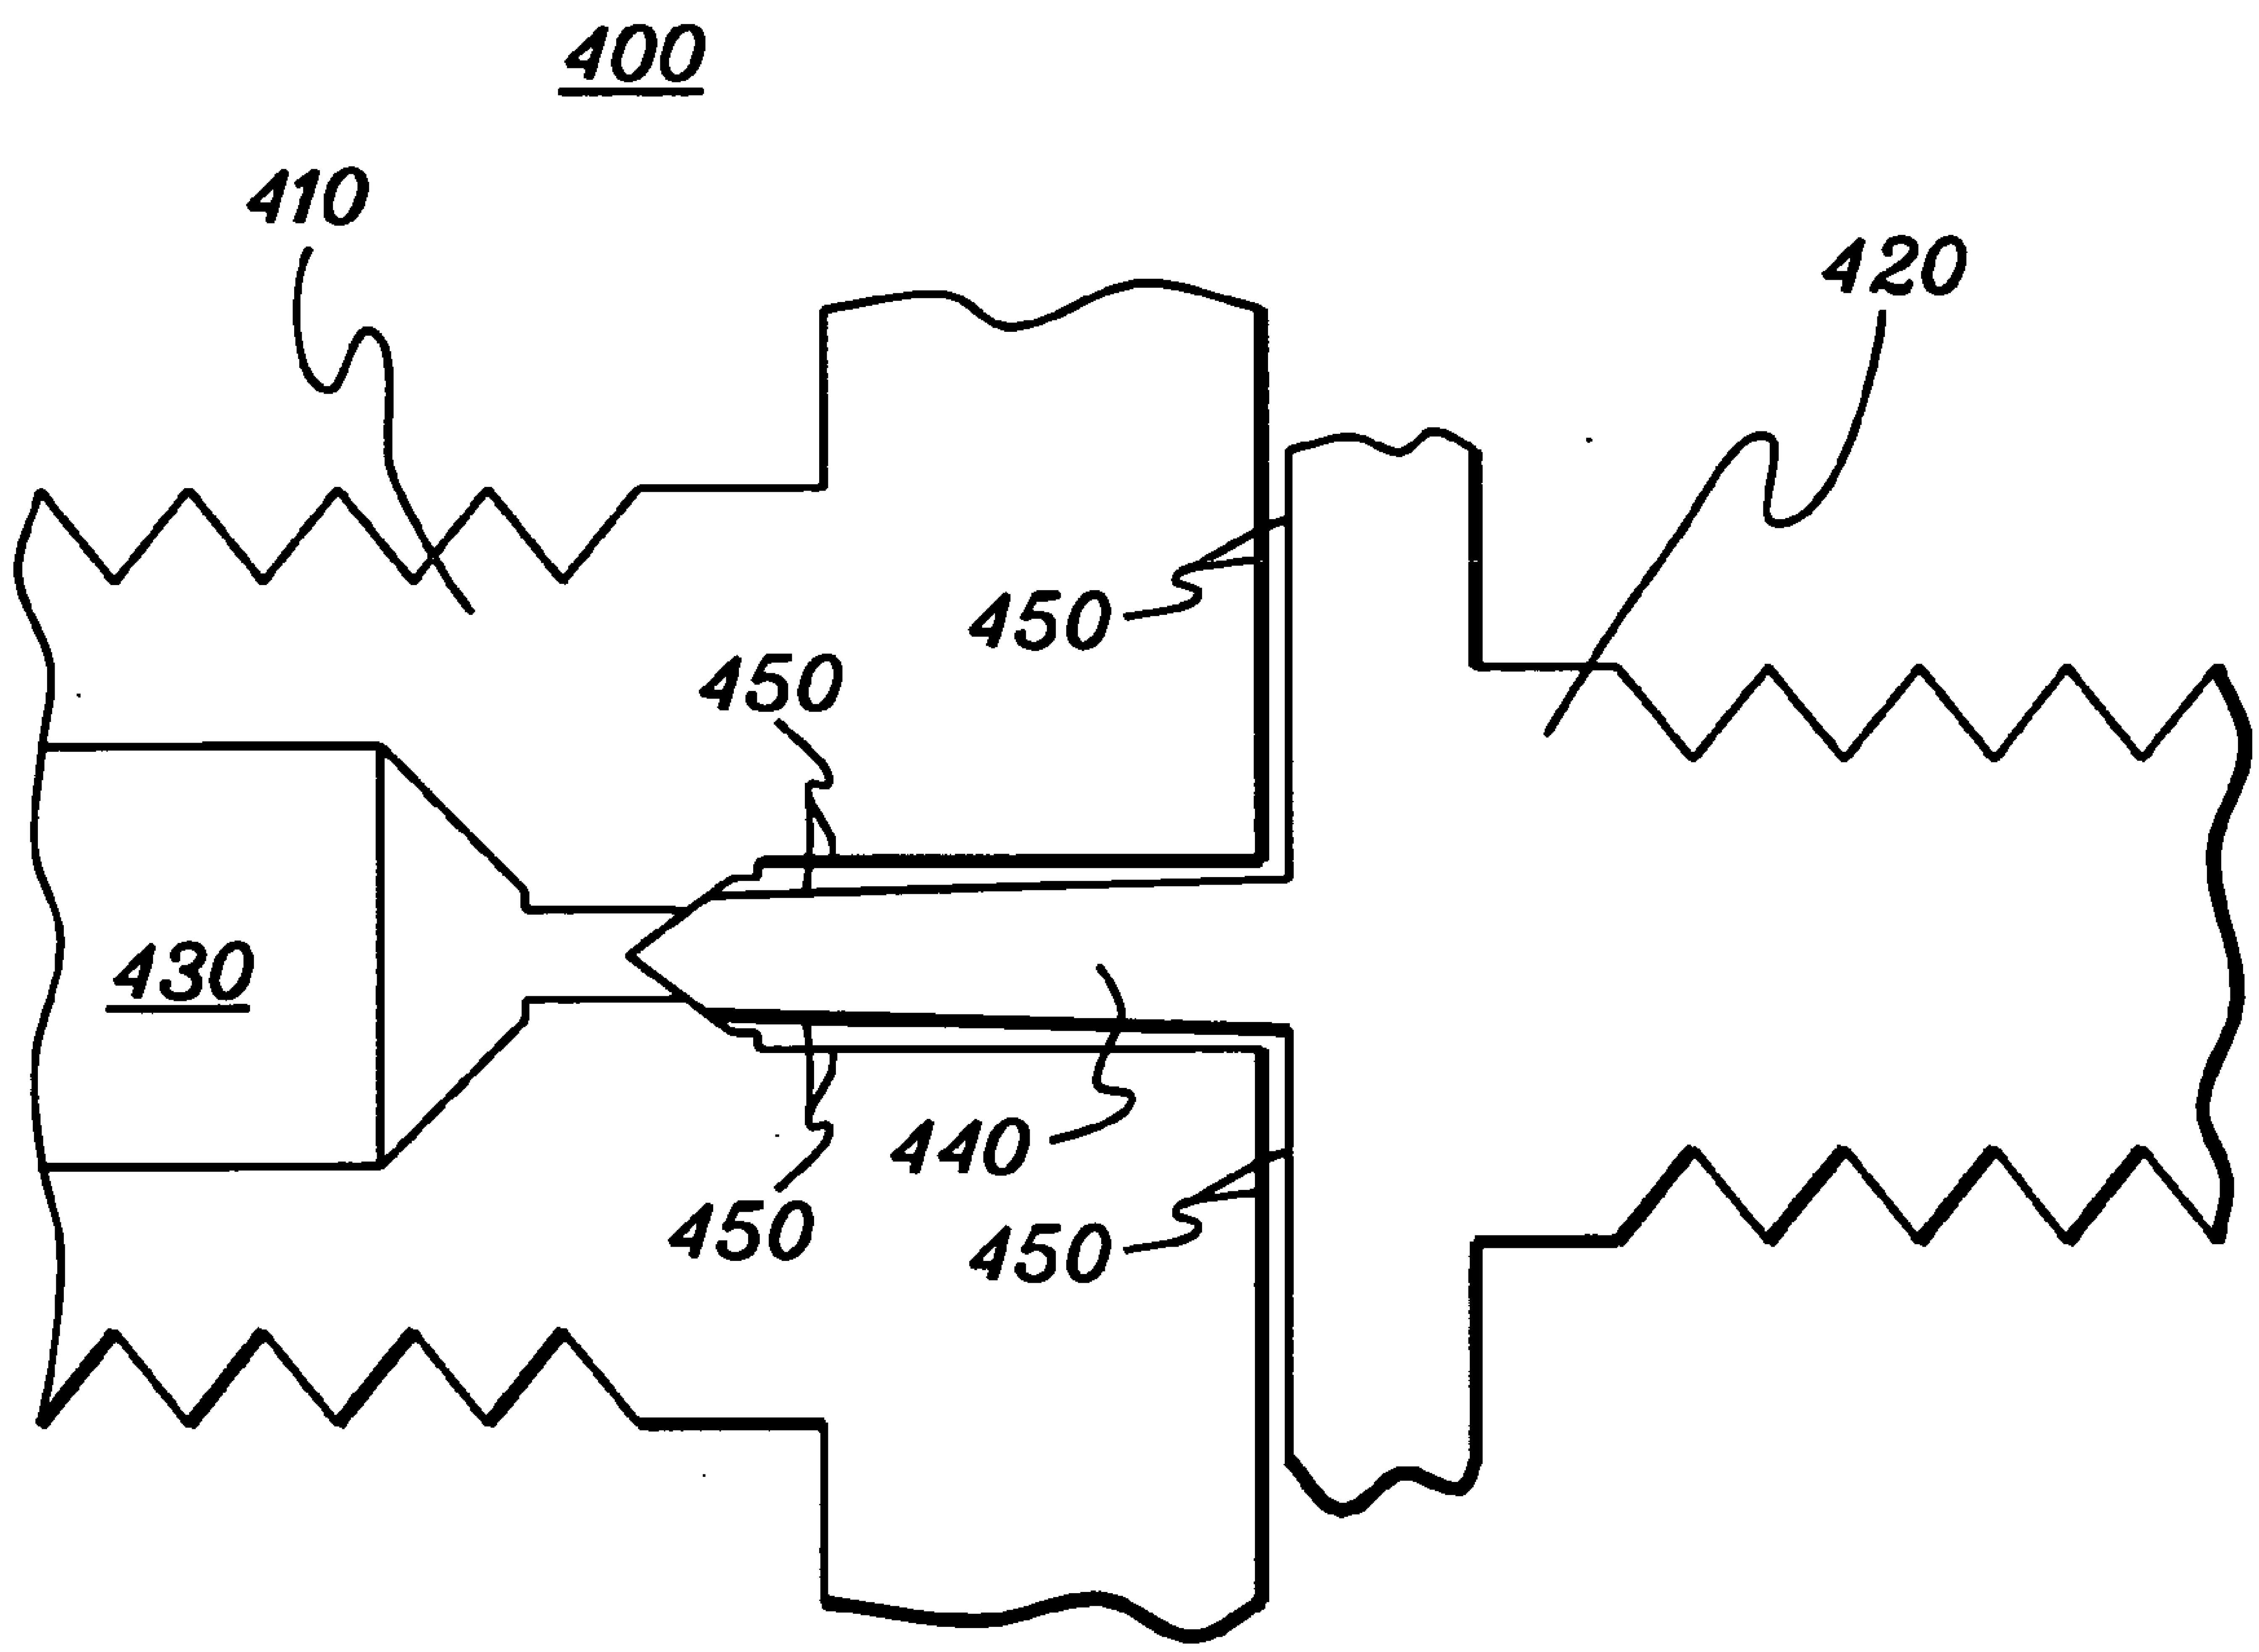

fig. 4

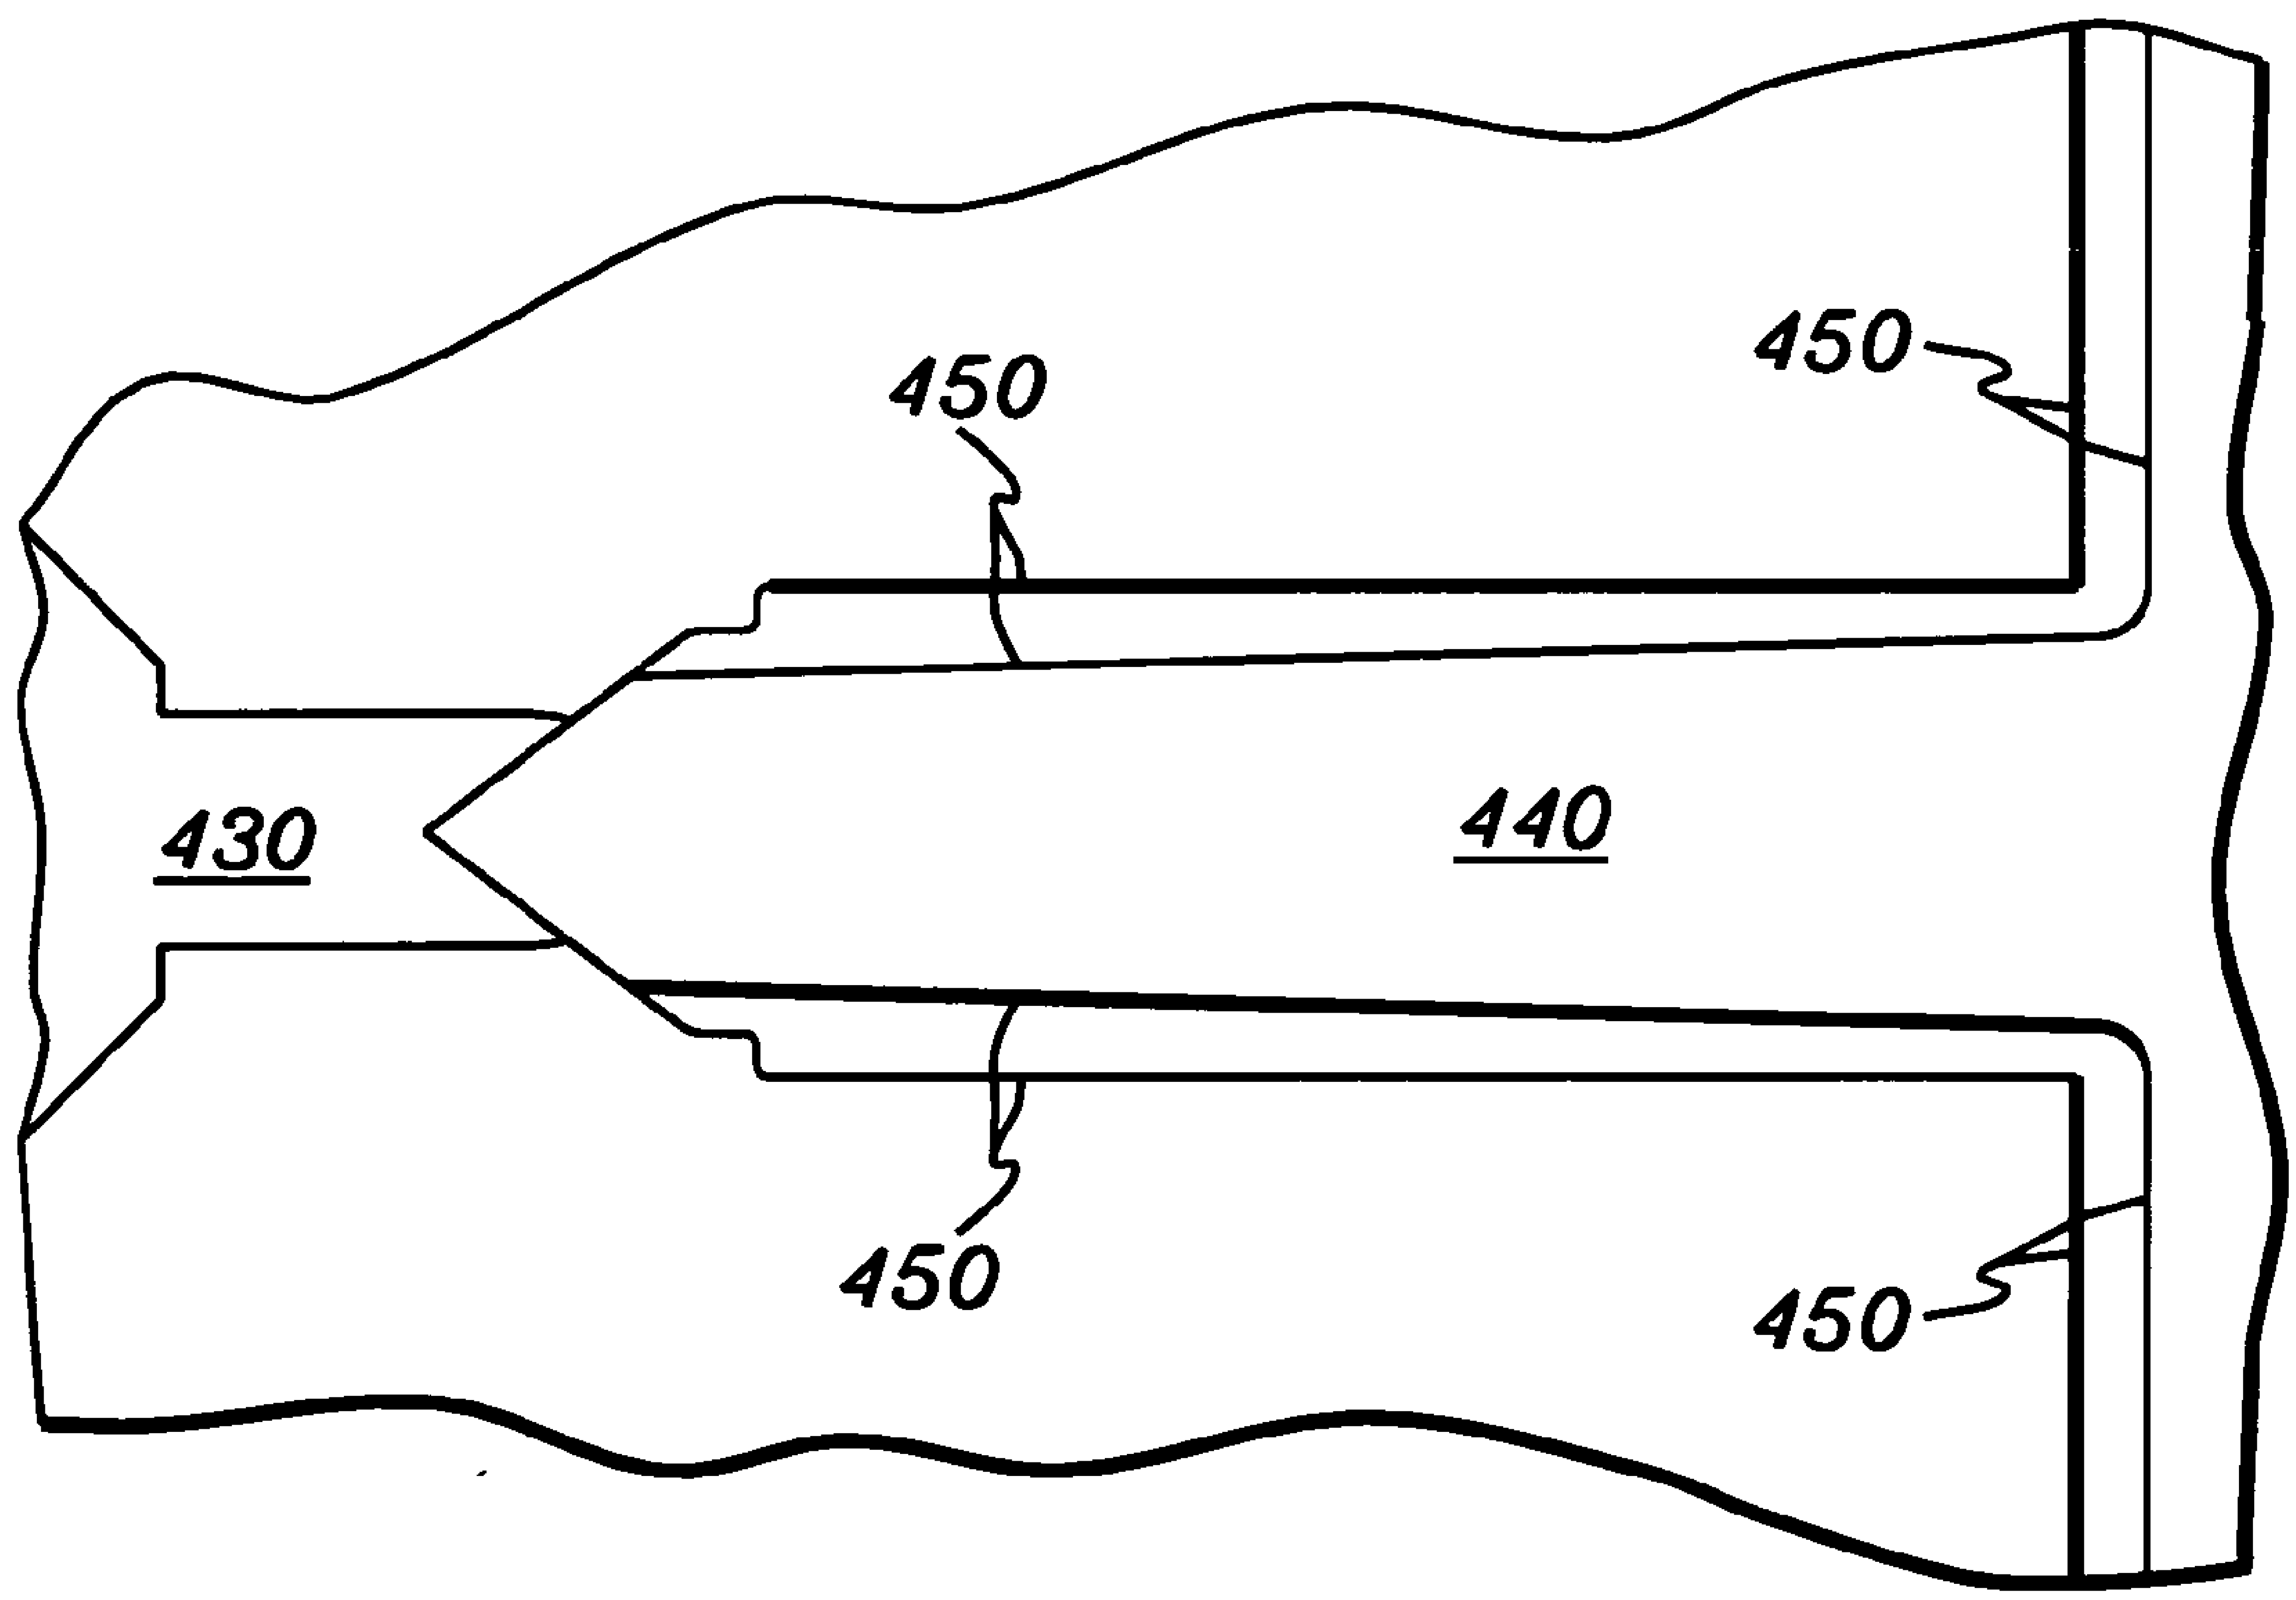

fig. 5

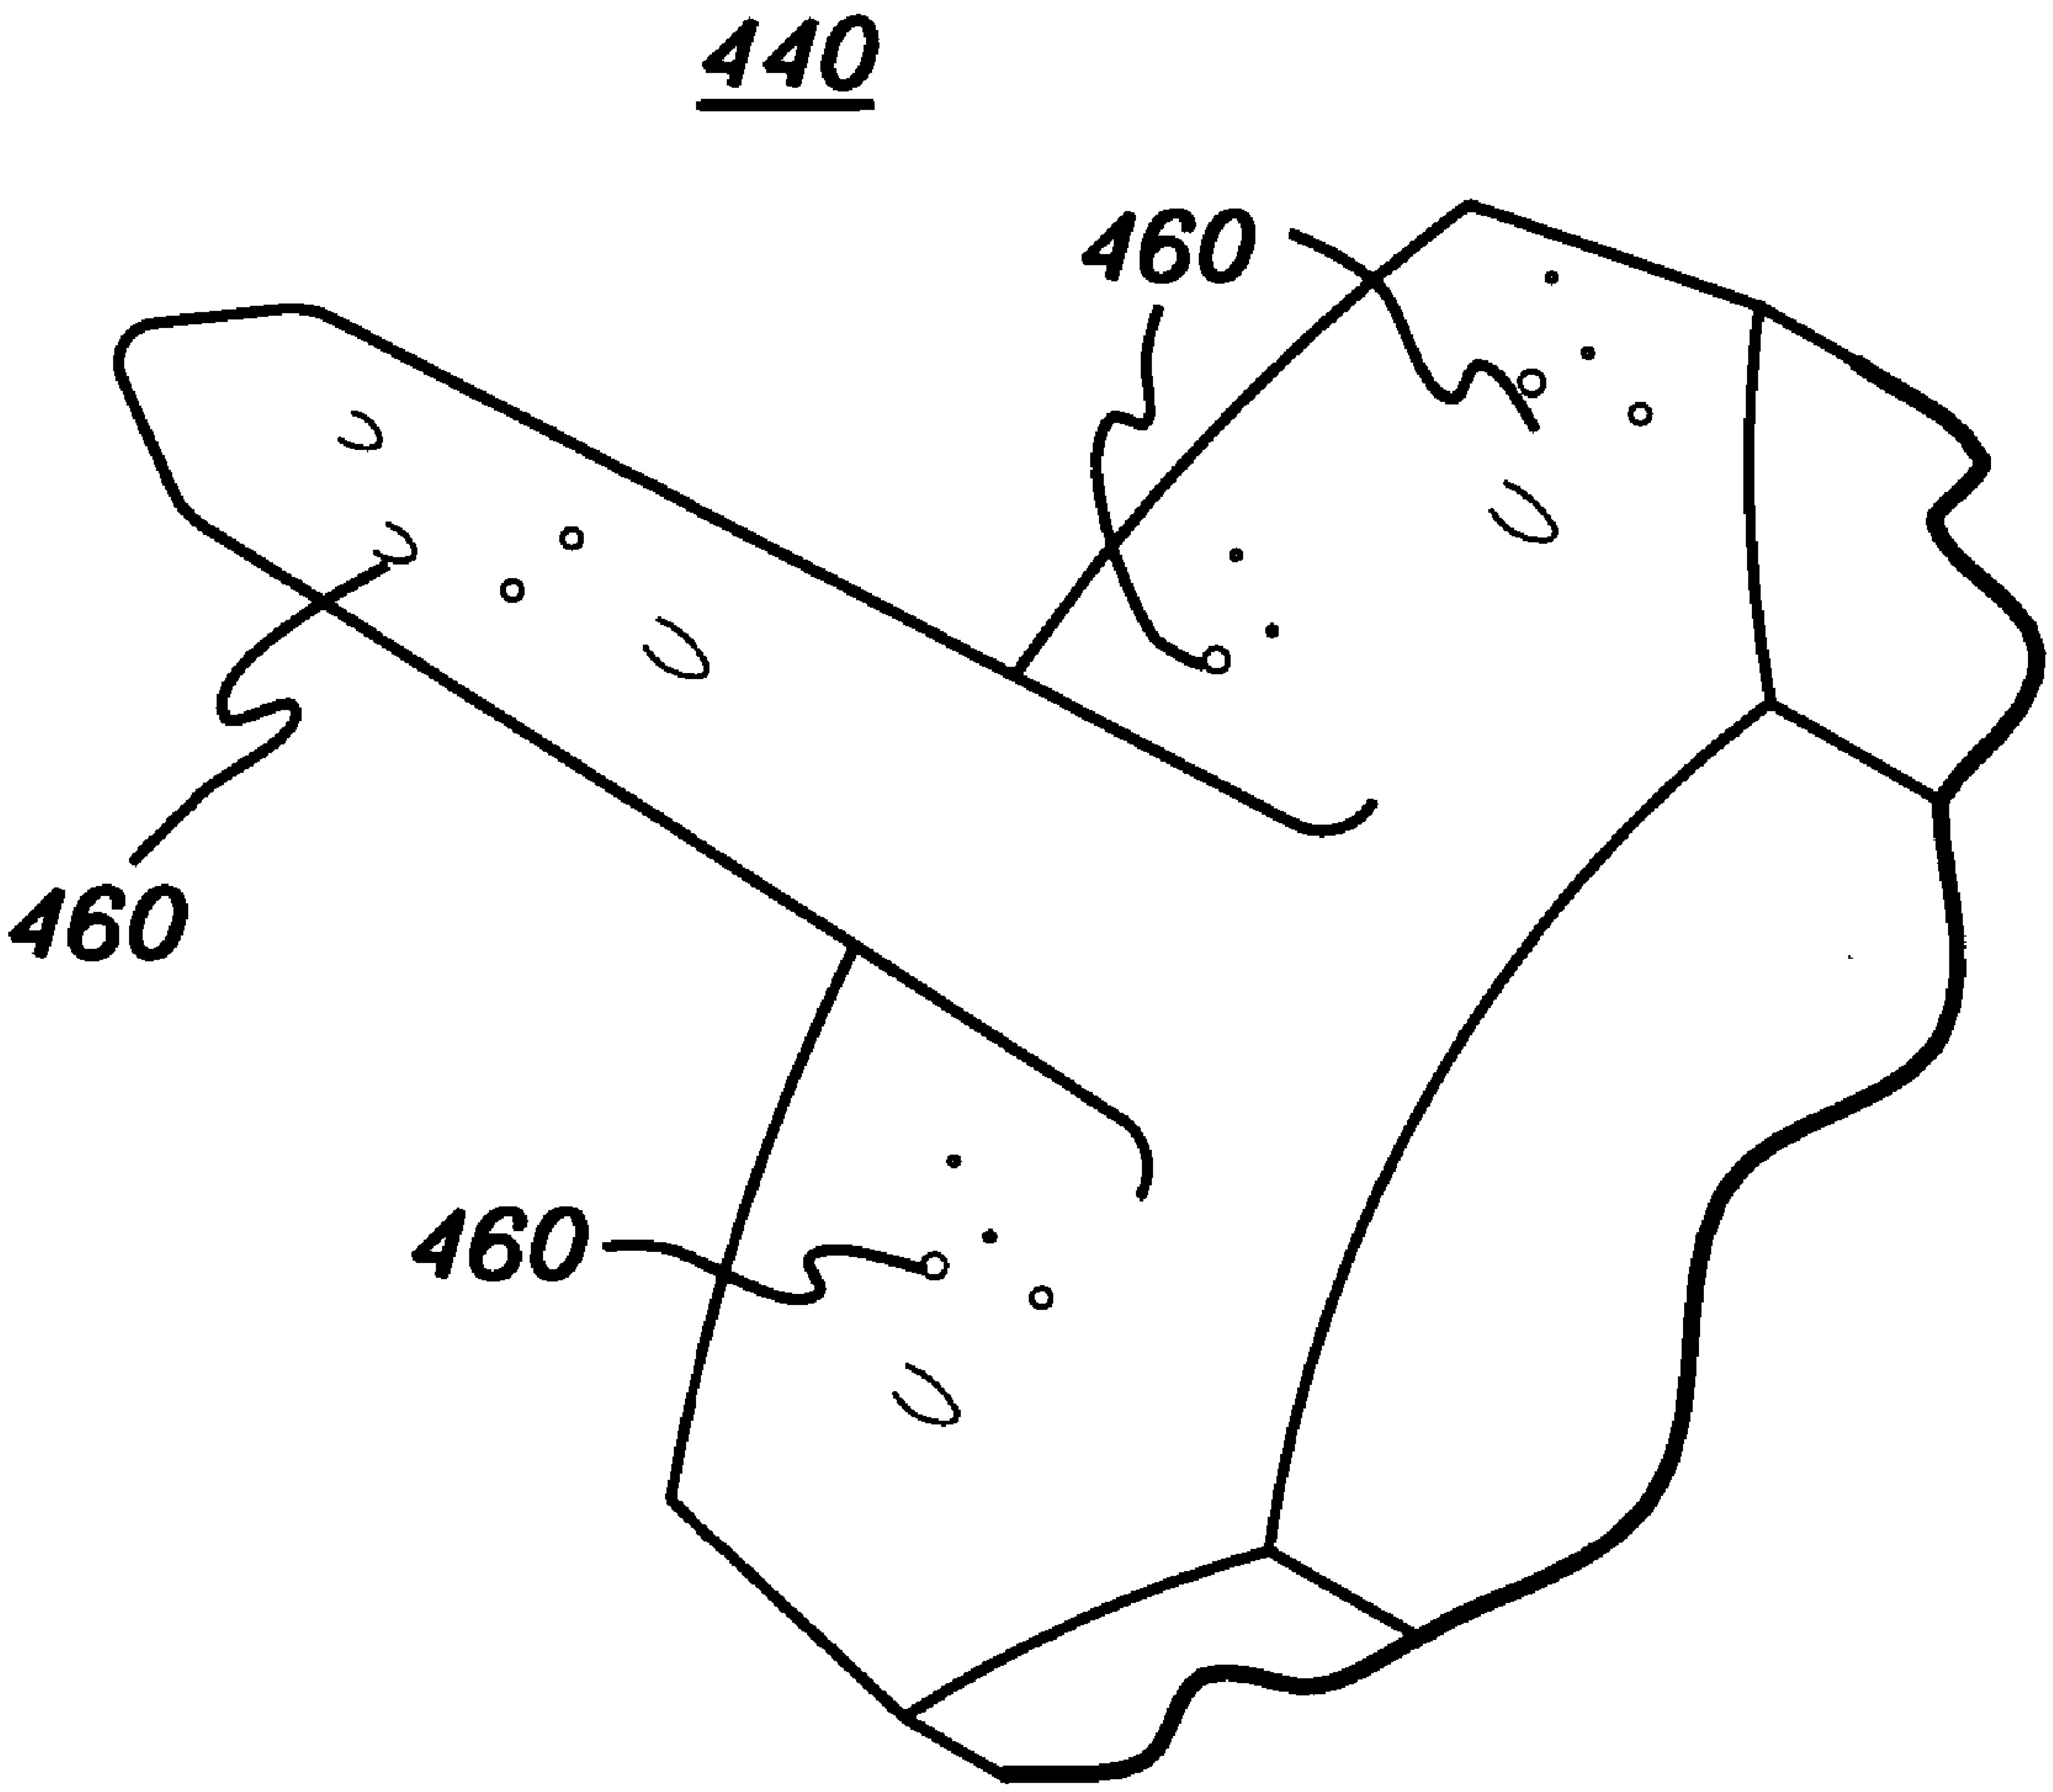

fig. 6

700

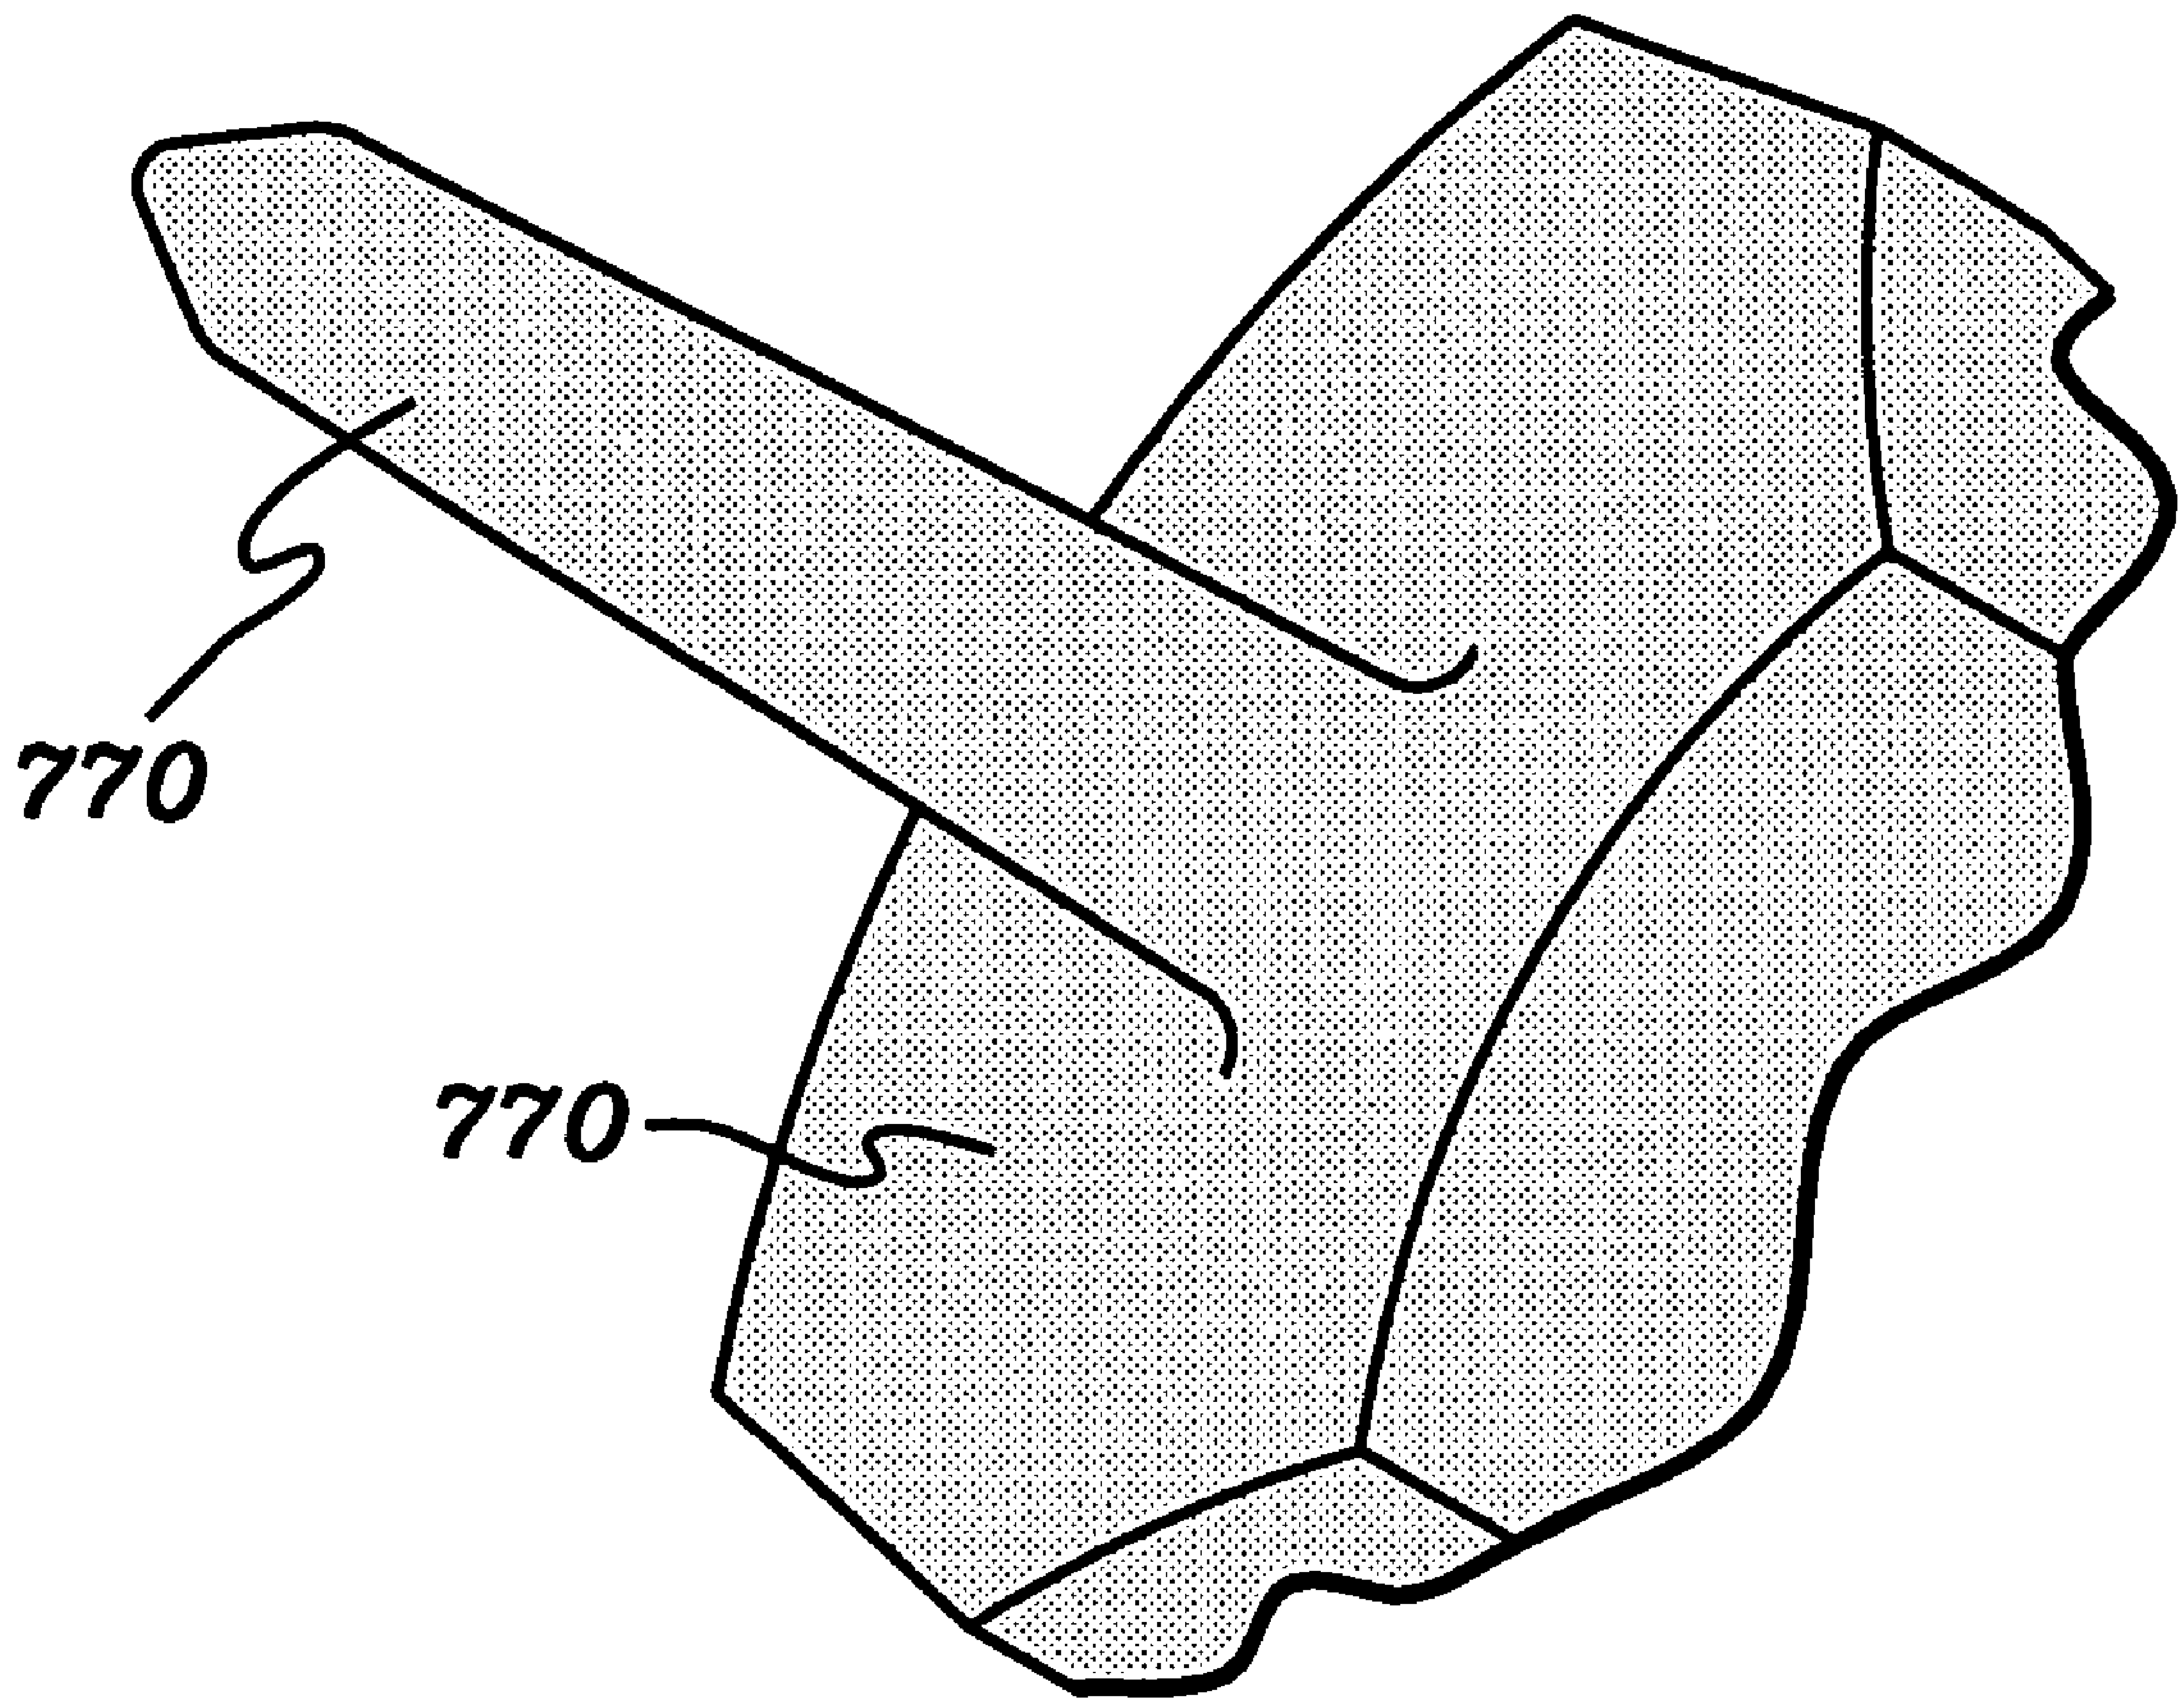

fig. 7

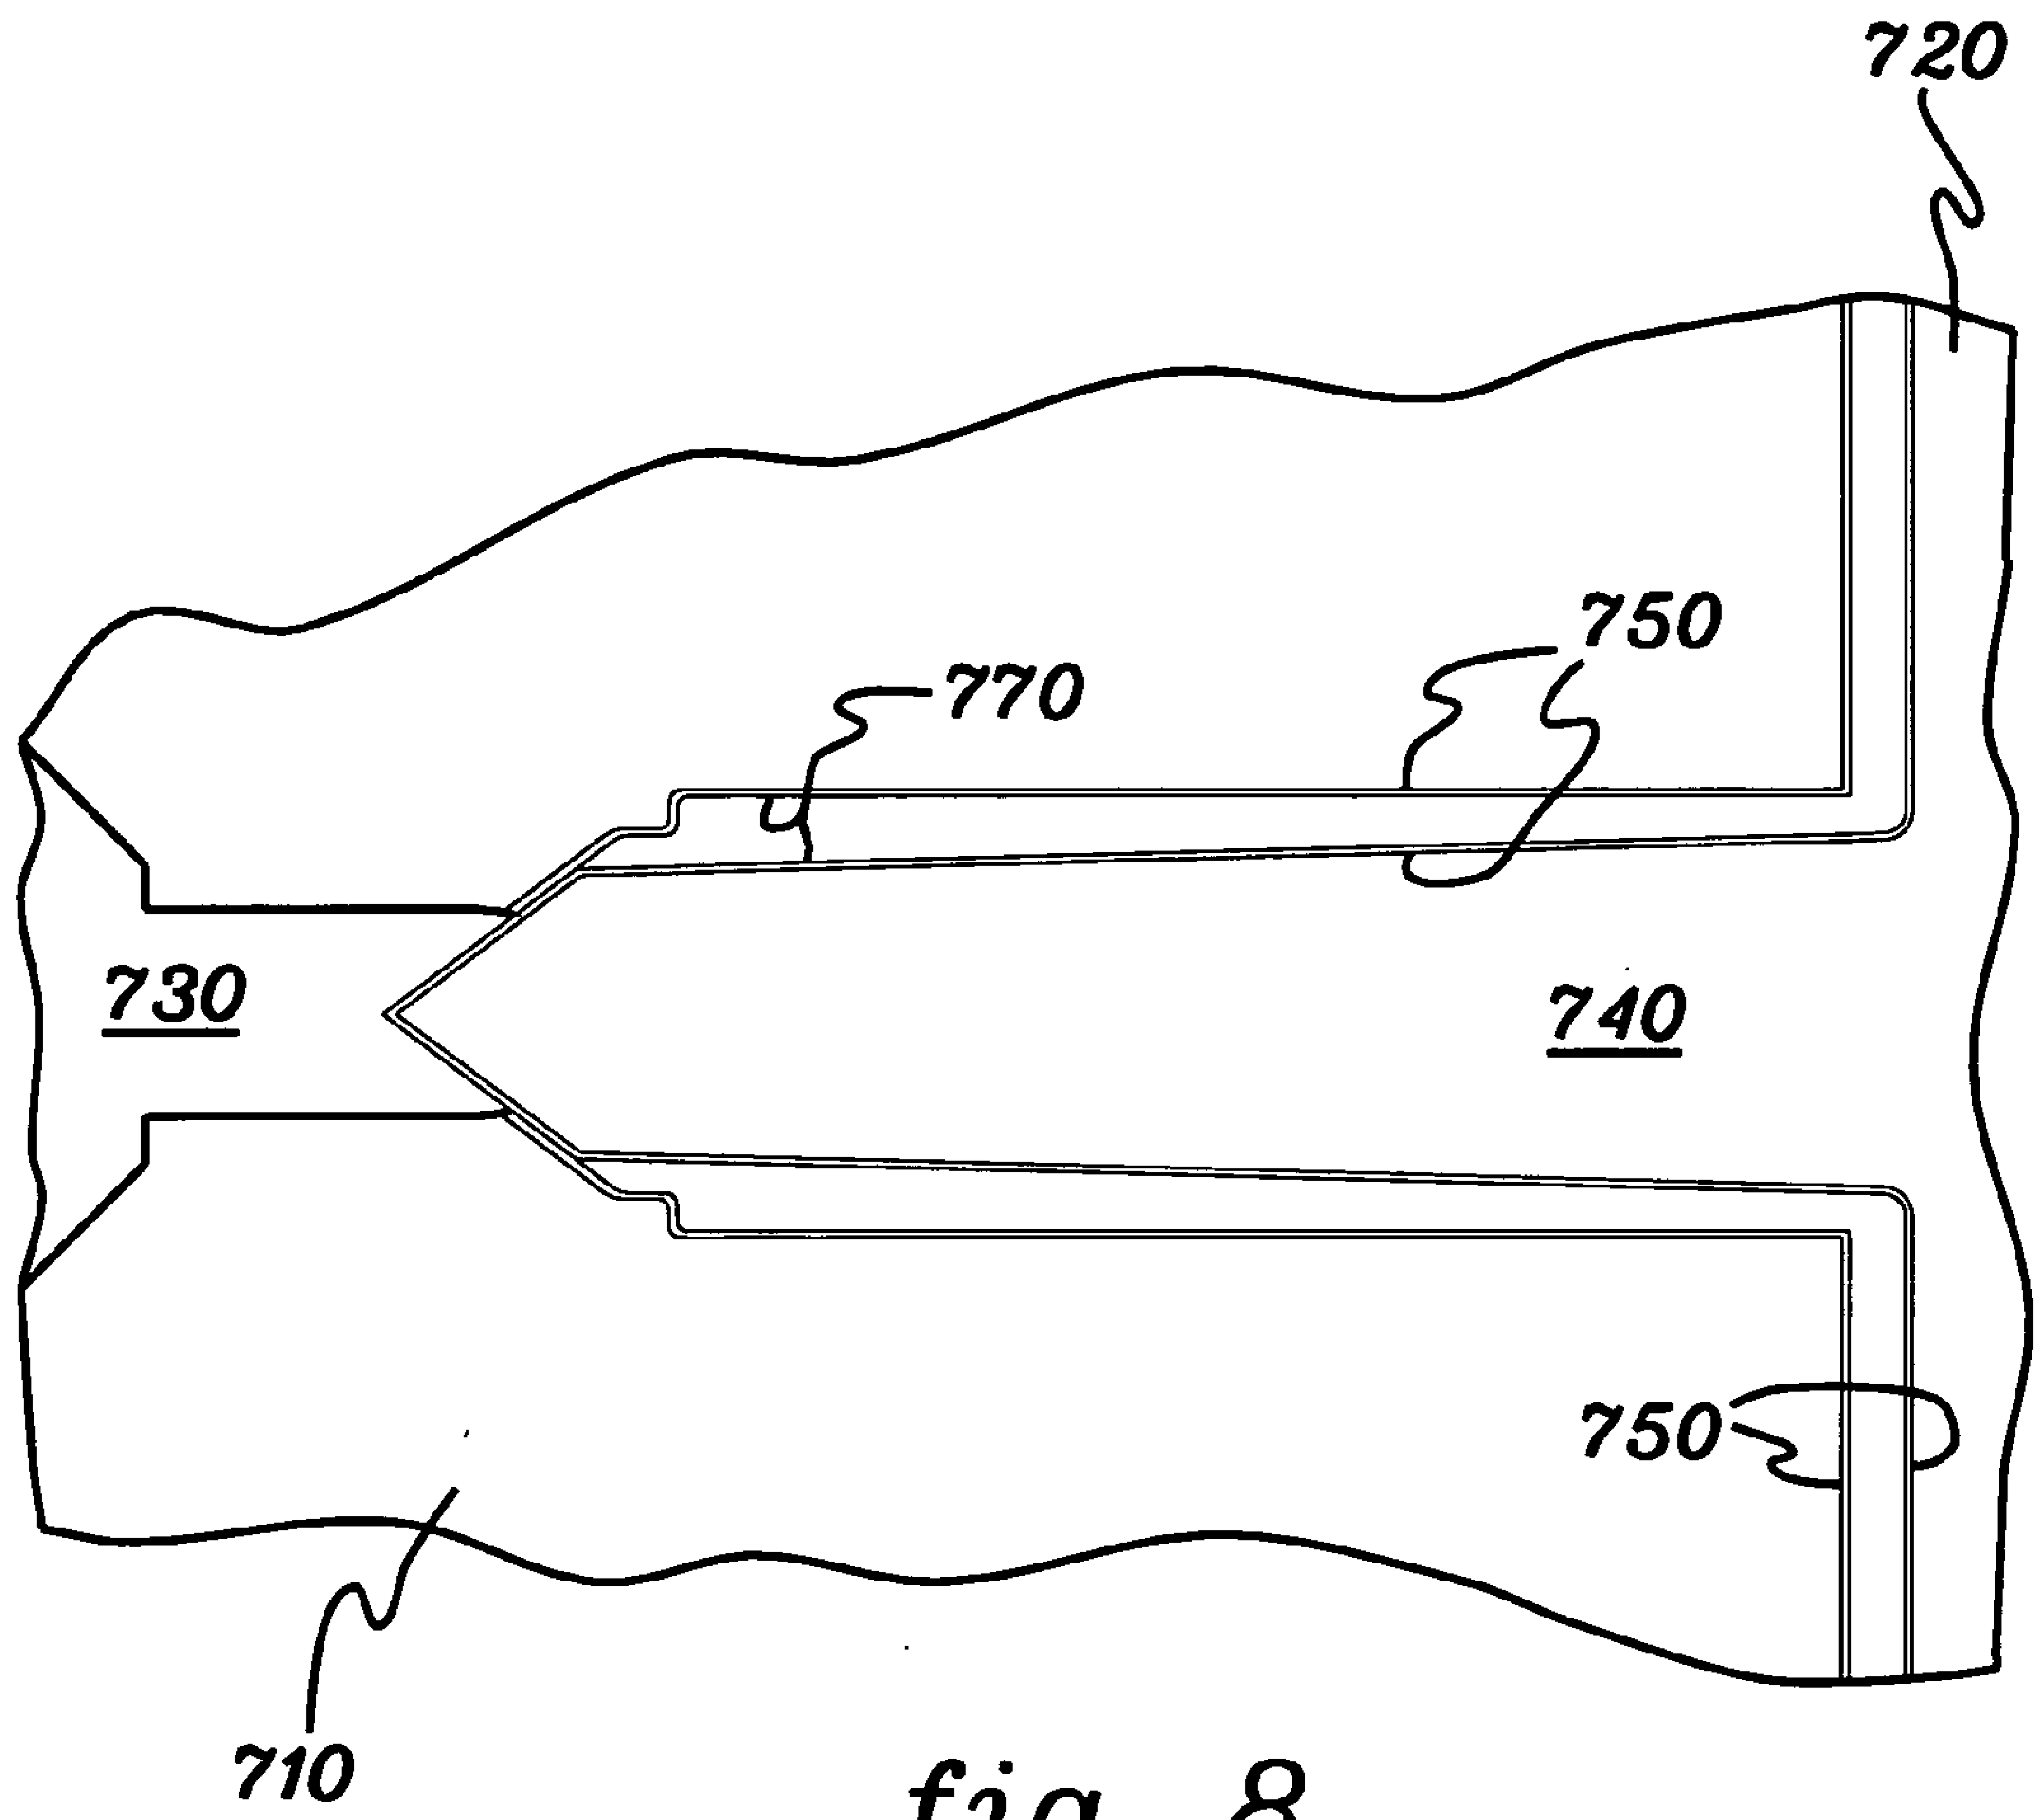

fig. 8

SYSTEMS AND METHODS FOR COOLING ELECTRONICS COMPONENTS EMPLOYING VAPOR COMPRESSION REFRIGERATION WITH SELECTED PORTIONS OF EXPANSION STRUCTURES COATED WITH POLYTETRAFLUORETHYLENE

TECHNICAL FIELD

[0001] The present invention relates generally to heat transfer mechanisms, and more particularly, to cooling systems and methods for removing heat generated by one or more heat generating electronics components. More particularly, the present invention relates to cooling systems and methods employing vapor compression refrigeration.

BACKGROUND OF THE INVENTION

[0002] As is known, operating electronic devices produce heat. This heat should be removed from the devices in order to maintain device junction temperatures within desirable limits. Failure to remove produced heat results in increased device temperatures, potentially leading to thermal runaway conditions. Several trends in the electronics industry have combined to increase the importance of thermal management, including heat removal for electronics devices, particularly in technologies where thermal management has traditionally been less of a concern, such as CMOS. In particular, the need for faster and more densely packed circuits has had a direct impact on the importance of thermal management. First, power dissipation, and therefore heat production, increases as device operating frequencies increase. Second, increased operating frequencies may be possible at lower device junction temperatures. Further, as more and more devices are packed onto a single chip, power density (Watts/cm^2) increases, resulting in the need to remove more power from a given size chip or module. Additionally, a common packaging configuration for many large computer systems today is a multi-drawer rack, with each drawer containing one or more processor modules along with associated electronics, such as memory, power and hard drive devices. These drawers are removable units so that in the event of failure of an individual drawer, the drawer may be removed and replaced in the field. A problem with this configuration is the increase in heat flux at the electronics drawer level. The above-noted trends have combined to create applications where it is no longer desirable to remove heat from modem devices solely by traditional air cooling methods, such as by using traditional air cooled heat sinks. These trends are likely to continue, furthering the need for alternatives to traditional air cooling methods.

[0003] One approach to avoiding the limitations of traditional air cooling is to use a cooling liquid. As is known, different liquids provide different cooling characteristics. For example, refrigerants or other dielectric fluids (e.g., fluorocarbon fluid) may have an advantage in that they may be placed in direct physical contact with electronic devices and interconnects without adverse affects such as corrosion or electrical short circuits. For example, U.S. Pat No. 6,052,284, entitled "Printed Circuit Board with Electronic Devices Mounted Thereon", describes an apparatus in which a dielectric liquid flows over and around several operating electronic devices, thereby removing heat from the devices. Similar approaches are disclosed in U.S. Pat. No. 5,655,290, entitled "Method for Making a Three-Dimensional Multichip Module" and U.S. Pat. No. 4,888,663, entitled "Cooling System for Electronic Assembly".

[0004] Notwithstanding the above, there remains a large and significant need to provide further useful cooling system enhancements for facilitating cooling of heat generating electronics components, such as one or more electronics modules disposed, e.g., in a book of an electronics rack of a computer installation.

SUMMARY OF THE INVENTION

[0005] In vapor compression refrigeration systems employed for cooling one or more heat generating electronics components, it has been discovered that material can agglomerate in certain pressure drop areas of expansion structures within the vapor compression refrigeration system. During refrigerant/oil transport through a hot compressor, any long-chain molecules and other typically non-soluble compounds at room temperature can go into solution in the hot mixture. These, as well as other physically transported impurities, then fall out of solution when the refrigerant/oil cools down. A layer of "waxy" material can build up in the pressure drop areas and act as a sticky substance which then catches other impurities. This material has been found to amass on expansion structures such as expansion valves, and particularly on the pin and orifice control region in the refrigerant flow path of the expansion valve. This amassing of material can interfere with the normal control volumes and interfere with the control of motor steps (due to unpredictable valve characteristic changes). This is particularly true when the vapor compression refrigeration system is employed in a cooling application for removing heat from a heat generating electronics component as described herein since control of the valve in this environment is a very sensitive application and expansion structure geometries are typically very small. To eliminate all contaminants from the vapor compression refrigeration system would be too costly, if not impossible. Thus, presented herein is a solution based on coating only selected pressure drop areas of the vapor compression refrigeration system to eliminate or reduce the clogging effect of debris and impurities in critically tight areas. This application is particularly significant in a cooling system where little of the expansion valve's available valve volume is employed during a vapor compression cycle.

[0006] The shortcomings of the prior art and additional advantages are provided through the provision of a cooling system for cooling at least one heat generating electronics component. The cooling system includes a vapor compression refrigeration system. The vapor compression refrigeration system has a condenser, at least one expansion structure, at least one evaporator and a compressor all coupled in fluid communication to define a refrigerant flow path and allow the flow of refrigerant therethrough. The at least one evaporator facilitates removal of heat produced by the at least one heat generating electronics component, while at least a portion of the at least one expansion structure is coated with a polytetrafluorethylene in the refrigerant flow path. The polytetrafluorethylene coating inhibits accumulation of material on selected pressure drop surfaces of the at least one expansion structure.

[0007] In another embodiment, a vapor compression refrigeration cooling system is provided for cooling at least one heat generating electronics component. This cooling system includes: a condenser, a first electrically controlled expansion valve coupled to the condenser, a first evaporator coupled to the first electrically controlled expansion valve; a second electrically controlled expansion valve coupled to the condenser, a second evaporator coupled to the second electrically controlled expansion valve; a controller providing control signals to the first electrically controlled expansion valve and the second electrically controlled expansion valve to control operation of the first electrically controlled expansion valve and the second electrically controlled expansion valve; and a compressor coupled to the first evaporator, the second evaporator and the condenser. The condenser, the first electrically controlled expansion valve, the first evaporator, the second electrically controlled expansion valve, the second evaporator, and the compressor are coupled in fluid communication to define multiple refrigerant flow paths, each refrigerant flow path allowing flow of refrigerant therethrough. The first evaporator and the second evaporator facilitate removal of heat produced by the at least one heat generating electronics component. At least a portion of the first electrically controlled expansion valve and at least a portion of the second electrically controlled expansion valve are coated with a polytetrafluorethylene in the respective refrigerant flow paths for inhibiting accumulation of material thereon.

[0008] In a further aspect, a method of fabricating a vapor compression refrigeration system for cooling at least one heat generating electronics component is provided. The method includes: (i) providing a condenser, at least one expansion structure, at least one evaporator, and a compressor; (ii) providing a polytetrafluorethylene coating on at least a portion of the at least one expansion structure; (iii) coupling the condenser, at least one expansion structure, at least one evaporator and compressor in fluid communication to define a refrigerant flow path; and (iv) providing refrigerant within the refrigerant flow path of the vapor compression refrigeration system to allow for cooling of the at least one heat generating electronics component employing sequential vapor compression cycles, wherein the polytetrafluorethylene coating is provided on the at least a portion of the at least one expansion structure in the refrigerant flow path for inhibiting the accumulation of material thereon.

[0009] Further, additional features and advantages are realized through the techniques of the present invention. Other embodiments and aspects of the invention are described in detail herein and are considered a part of the claimed invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] The subject matter which is regarded as the invention is particularly pointed out and distinctly claimed in the claims at the conclusion of the specification. The foregoing and other objects, features, and advantages of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

[0011] FIG. **1** depicts one embodiment of a cooling system comprising a vapor compression refrigeration system, in accordance with an aspect of the present invention;

[0012] FIG. **2** illustrates one example of a flowchart that shows how a Modular Refrigeration Unit (MRU) code which contains a method to monitor and regulate multi-chip module (MCM) temperature under primary MRU cooling, a power control code (PCC) which contains a method to determine and communicate the thermal state or range that equates to a specific temperature and voltage condition, and a Cycle Steering Application (CSA) code which contains a method of matching the various logic clocks to the thermal degrade states that exist, may interact in a single temperature-power-logic control system, in accordance with an aspect of the present invention;

[0013] FIG. **3** depicts a system schematic where the MRU code, PCC code, and CSA code are physically located in a server having four processor books or nodes, cooled in primary mode by two MRUs, and in back-up mode by blowers, in accordance with an aspect of the present invention;

[0014] FIG. **4** is a cross-sectional, elevational view of one embodiment of an expansion structure comprising an expansion valve having an expansion pin and an expansion orifice which are part of a refrigerant flow path of a vapor compression refrigeration cooling system, in accordance with an aspect of the present invention;

[0015] FIG. **5** is an enlarged, cross-sectional view of the expansion orifice and expansion pin illustrated in FIG. **4**, in accordance with an aspect of the present invention;

[0016] FIG. **6** is an isometric view of one embodiment of an expansion pin for of an expansion valve of a vapor compression refrigeration system, wherein material/debris is shown amassed on exposed surfaces of the expansion pin, which would be in a pressure drop area of the refrigerant flow path (not shown);

[0017] FIG. **7** is an isometric view of an expansion pin of an expansion valve of a vapor compression refrigeration system, wherein the expansion pin is coated with a layer of polytetrafluorethylene in pressure drop areas the refrigerant flow path, in accordance with an aspect of the present invention; and

[0018] FIG. **8** is a cross-sectional, elevational view of an expansion orifice and expansion pin of an expansion valve of a vapor compression refrigeration system showing selected pressure drop areas of the expansion pin and inner surface of the expansion orifice coated with a polytetrafluorethylene, in accordance with an aspect of the present invention.

BEST MODE FOR CARRYING OUT THE INVENTION

[0019] As used herein, the term "electronics rack" includes any frame, rack, blade server system, etc., having at least one heat generating electronics component of a computer system or electronics system, and may be, for example, a stand alone computer processor having high, mid or low end processing capability. In one embodiment, an electronics rack may comprise multiple books, each book having one or more heat generating electronics components requiring cooling. Each "heat generating electronics component" may comprise an electronic device, an electronics module, an integrated circuit chip, a multi-chip module, etc. An "expansion structure" is any structure or area in a vapor compression refrigeration system where there is a pressure drop, and thus refrigerant expansion during a refrigerant compression/expansion cycle. As used herein, the term "expansion structure" includes any structure of a pressure drop area and adjacent areas where an agglomeration would effect an expansion structure characteristic, including any thermally effected conduction zones and any downstream mass transport zones. Examples of expansion structures include expansion valves, including electronic expansion valves, thermal expansion valves, hot-gas bypass valves, or mechanical expansion valves, as well as other refrigerant expansion structures such as a fixed expansion orifice in an evaporator. As used herein, an "expansion orifice" means any opening defined by a component within the vapor compression refrigeration system, and includes a fixed orifice in an evaporator, as well as an opening defined by an inner surface of an expansion valve. Further, the word "refrigerant" is used herein to refer to any coolant which can be employed in a vapor compression/expansion system.

[0020] One example of refrigerant within a cooling system in accordance with an aspect of the present invention is R-134A coolant (i.e., 1,1,1,2 tetrafluoroethane), however, the concepts disclosed herein are readily applied to other types of refrigerants, other dielectric fluids (e.g., fluorocarbon fluid), or other types of coolants while still maintaining the advantages and unique features of the present invention.

[0021] FIG. **1** depicts a cooling system **100** as an exemplary embodiment of the present invention. Cooling system **100** includes a condenser **104** and two evaporators **106** and **108**. Evaporators **106** and **108** cool heat generating electronics components **110** and **112**, respectively. In this embodiment, components **110** and **112** are multi-chip modules (MCMs), but it is understood that other components (e.g., single processors, memory) may be similarly cooled.

[0022] Both evaporators **106** and **108** are supplied refrigerant from a common condenser **104**. An expansion valve **114** receives high pressure liquid refrigerant from condenser **104** and generates low pressure liquid refrigerant to evaporator **106**. An expansion valve **116** receives high pressure liquid refrigerant from condenser **104** and generates low pressure liquid refrigerant to evaporator **108**. Expansion valves **114** and **116** are electrically controllable. A controller **120** provides control signals to expansion valve **114** and expansion valve **116** to control refrigerant flow and pressure drop across each expansion valve. In an exemplary embodiment, expansion valves **114** and **116** each includes a stepper motor that responds to control signals from the controller **120**. The stepper motor opens or closes an orifice in the expansion valve to regulate refrigerant flow and pressure drop. Controller **120** executes a computer program to control the expansion valves **114** and **116**.

[0023] The low pressure liquid refrigerant exits the expansion valves **114** and **116** and is supplied to evaporators **106** and **108**, respectively. The refrigerant in each evaporator **106** and **108** is converted to low pressure vapor refrigerant, in part, though further fixed expansion structures **107**, **109**, respectively, and provided to a common compressor **122**. High pressure vapor from compressor **122** is supplied to condenser **104**. Fan **126** establishes air flow across condenser **104** to facilitate cooling the high pressure vapor refrigerant to high pressure liquid refrigerant.

[0024] A plurality of temperature sensors are distributed throughout the cooling system **100**. The sensors may be thermistors or other known temperature sensors. Sensor T**1** measures air temperature entering condenser **104**. Sensor T**2** measures air temperature exiting condenser **104**. Sensors T**3** and T**3**' provide redundant measurement of refrigerant temperature exiting condenser **104**. Sensor T**4** measures refrigerant temperature entering condenser **104**. Sensor T**6** measures refrigerant temperature entering evaporator **106** and sensor T**7** measures refrigerant temperature exiting evaporator **106**. Sensor T**8** measures refrigerant temperature entering evaporator **108** and sensor T**9** measures refrigerant temperature exiting evaporator **108**. Sensor T_{hat1} measures temperature at electronics component **110** and sensor T_{hat2} measures temperature at electronics component **112**.

[0025] Each temperature sensor generators a temperature signal which is supplied to controller **120** and shown as T_{in}. The control **120** adjust the expansion valves **114** and/or **116** in response to one or more of the temperature signals to maintain the logic modules **110** and **112** at a predefined temperature. Controller **120** controls expansion valves **114** and/or **116** to obtain desired superheat valves while maintaining each electronics component at a desired temperature. Each component **110** and **112** may be maintained at a different temperature or the same temperature, even if each component has different heat loads.

[0026] Evaporators **106** and **108** may be connected to the refrigerant supply and refrigerant return lines through quick disconnect connectors **130**. The controllable expansion valves **114** and **116** allow an evaporator to be removed for maintenance or upgrade while the other evaporator, condenser and compressor continue to operate. For example, expansion valve **114** can be closed and the refrigerant from evaporator **106** removed by the suction of the compressor **122**. Evaporator **106** can then be removed for service, upgrade, etc.

[0027] Although two evaporators are shown connected to one modular refrigeration unit (MRU) (condenser, compressor, expansion valves and controller), it is understood that more than two evaporators may be coupled to each MRU.

[0028] In an exemplary embodiment in accordance with the present invention, one embodiment joins methods to monitor and control the temperatures of electronics components **110**, **112**, to report the temperature state and to adjust the voltage levels appropriately and to adjust the various clock speeds which govern CMOS circuits that are effected by the change in temperature and/or voltage.

[0029] A detailed description of one method of monitoring and controlling the temperature of a hybrid cooling system **100** is described below with reference to FIGS. **2** and **3**. FIG. **2** illustrates a flowchart that shows how a Modular Refrigerant Unit (MRU) code **200**, which contains a method to monitor and regulate the MCM (i.e., one example of a component) temperature under primary MRU cooling, interfaces with a Power Control Code (PCC) **210**, which contains a method to determine and communicate the thermal state or range that equates to a specific temperature and voltage condition of each MCM, and a Cycle Steering Application (CSA) code **220**, which contains the method of matching the various logic clocks to the thermal degrade state that exist. The MRU code, PCC code and CSA code, all interact into a single temperature-power-logic control system generally indicated as **230**.

[0030] FIG. **3** shows one embodiment of a system schematic wherein the MRU code **200**, the PCC code **210** and the

CSA code **220** are physically located in a server that has four Processor (PU) books or nodes **242**, **244**, **246**, **248**, respectively, each having an electronics component or MCM cooled in primary mode by one of two MRUs **250**, **252** and in backup mode by two blowers **254**. The backup blowers **254** provide air cooling of all PU books **242**, **244**, **246**, **248**, for MRU failures or light logic load state. Each MCM is operably connected to a main system board generally indicated at **256**. The MRU code **200** is in each MRU **250**, **252**. The PCC code **210** is split between Base Power Cage Controllers or Base Power Assembly **260**, **262** and digital converter assemblies (DCA) cage controllers (DCA **01**, **02**, **11**, **12**, **21**, **22**, **31**, **32**). The Base Power Assembly **260**, **262** provides high voltage DC power to the entire server **240** and the DCA converts the high DC power to low DC voltages used by each circuit. The CSA code **220** is located in the first Processor book **244** (labeled PU Book **0**) of multi-node server **240**.

[0031] Each MCM (not shown) in each PU book **242-248** includes a hat **274** in operable communication with a cooling unit **10** and connected to a thermal sensor assembly **276**. Each thermal sensor assembly **276** preferably includes three thermistors configured to sense a temperature of a corresponding MCM.

[0032] The thermal sensors are compared for miscompare properties and for insanity limits to make sure the temperatures measured are accurate. One sensor is directly sensed by the Modular Refrigeration Unit (MRU) indicated generally at **278** and the other two are read by the power supply feeding the MCM power indicated generally at **280** to insure full redundancy and accuracy of this reading. The MRU reads an MCM hat thermistor sensor directly through its drive card to enable continual monitoring and thermal regulation in case of a cage controller (cc) failover. MCM hat thermistors that are read by each DCA power supply as well as by the MRU are compared to each other by the MRU and Power Control Code to identify any faulty sensors and eliminate the faulty sensors from consideration generally indicated at **286** in FIG. **2**. This insures redundancy of control and cooling status function. The power supply thermistor also serves for thermal protection of the MCMs, dropping power if the temperatures are near damage limits.

[0033] The control of the primary cooling system is done by using a Proportional Integral Derivative (PID) control loop of an electronic expansion valve to each evaporator as described with reference to FIG. **1** and generally indicated at **290** in FIG. **2**. The PID control loop regulates the coolant flow to each MCM being cooled. The coolant flow is increased by opening the electronic expansion valve if the MCM is too warm or is higher than targeted and the flow is reduced by closing the valve position if the MCM is too cold or cooler than targeted.

[0034] When the PID control has opened its electronic expansion valve to the fully open position providing maximum coolant to a given MCM, the compressor speed then executes its own PID control loop to deliver additional cooling capacity to the MCM. In other words, a second PID control loop controls the compressor speed if the valve regulating the flow of coolant to a respective evaporator has reached its maximum cooling position.

[0035] Similarly, the blower speed of blower **126** cooling the refrigerant condenser **104** is controlled by the cooling capacity needs from the MRU. More specifically, blower speed controls provide more air for cooling the MRU condenser **104** when the thermistors T**1** and T**2** on the condenser **103** and ambient air indicate that inadequate condensing is taking place. Also, the speed of condenser blower **126** is increased in a warm ambient.

[0036] MCM power data **284**, read by the Power Control Code **210** and provided to the MRU code **200** every 2.5 seconds, determines if a given MCM no longer has its clocks functioning. If the MCM power stays low, indicating a non-functional Processor book, for sufficient time, the refrigerant coolant supply is stopped by completely closing the expansion valve to that MCM only and turning on the backup blowers **254** at a reduced speed. In this manner, other MCMs in the same server can stay refrigerant cooled while the MCM that has check stopped or otherwise ceased to function logically will be air cooled. Refrigerant cooling and MCM without adequate logic power can lead to condensation forming on its external surfaces. For example, when regulating light heat loads to a fixed temperature, the expansion device must significantly close the refrigerant flow rate, which lowers the pressure and hence the refrigerant temperature inside the evaporator cooling the MCM. When the clocks are off, the expansion valve closes so far that the evaporator pressure may be sub-atmospheric, which creates very cold local temperatures. These cold local temperatures with low heat flux and outside regions of the MCM can get cold enough to form condensate after extended operation in this condition.

[0037] The MRU code **200** also provides a function that enables virtually all of the refrigerant to be removed from the evaporator of a corresponding cooling unit before the refrigerant lines are opened for servicing the MCM or cooling hardware, as discussed above with respect to FIG. **1**. This is provided by closing the electronic expansion valves for some period before turning off the compressor, resulting in a partial vacuum that removes the refrigerant from the evaporator and connecting hoses, The benefits include better ecology and consistent refrigerant charge before and after the MRU is reconnected.

[0038] This temperature control code, together with primary and/or secondary cooling hardware, has the ability to program and run the MCMs at different or "biased" conditions to enable the MCM to be tested beyond the normal temperature conditions it sees in actual use. The temperature bias testing may be done while the logic voltage is also biased. In the prior art, these bias cooling functions required special tester cooling hardware and test code which was costly and inefficient compared to combining this stress test thermal function in the actual cooling system. Secondary cooling uses a PID loop also to achieve MCM temperature target that may be outside of the normal operating range.

[0039] Still referring to FIGS. **2** and **3**, a detailed description of the Power Contol Code (PCC) **210** which principally includes a method for monitoring the actual thermal or degrade state and for making suitable power and cooling adjustments, as well as reporting this state to the CSA code **220**, follows below. The thermal states of each MCM are monitored and the state of each MCM is communicated to a function that determines the proper clock cycle time, called the Cycle Steering Application (CSA) code **220**. This function tells the CSA code **220** both which cycle time range of the circuits are now operating in and whether the cause of the failure of the primary cooling means has been repaired or not.

[0040] In particular, PCC **210** continually monitors and posts "cooling state" data to the CSA code **220** indicated generally as **292**. The thermal state is defined by discrete temperature ranges that are associated with a given clock speed as the proper speed to operate. In other words, the full operating temperature range from coldest to ambient to shut-down for thermal protection is subdivided into smaller discrete operating ranges. The coldest steady state temperature range is called the normal state, and is the temperature range kept under normal primary cooling means (e.g., MRUs **250**, **252** and cooling units **10**). When the primary cooling means no longer functions properly, the cooling state, sensed via the MCM sensors **276**, is reported as a specific "degrade state". By way of example, there may be between **2** and **4** degrade states between normal operation and thermal shut-down, but more or less are also contemplated, and hence, these concepts are not limited to between **2** and **4**. Within a given degrade state, there exists one "optimum" set of clock speeds.

[0041] The PCC **210** reads the actual current **294** and voltage **284** being supplied to each MCM as well as its temperature **286**. Based on the leakage characteristics of the CMOS technology, the capacity left in the power supply providing the current to the MCM, and operating temperatures, the PCC **210** may either increase or decrease or leave alone the applied voltage level to each set of circuits indicated generally **296**.

[0042] When the voltage is increased, the increased voltage enables a higher range of operating temperatures before a given degrade state is indicated to the CSA code **220** to slow the clocks. Hence, the higher voltage can delay the need to operate in a slower clock range. This is because CMOS switches faster at higher voltages somewhat offsetting the slowing effects of warmer circuits.

[0043] Normally, it is desirable to increase voltage applied to the circuits to offset some of the slowing effect on circuit switching of warmer circuits. Typically, a 6% increase in voltage will cause circuits to switch about 4% faster, offsetting a 25° C. temperature rise. However, with recent circuit technology, power increases strongly with higher temperature and increased voltage. In some cases it may require the voltage to be dropped when the junction temperature rises significantly, even though this lowering of voltage will increase the amount of slowing of the clock frequency that is needed. This disclosure includes all three voltage responses to loss of normal cooling: doing nothing, increasing voltage, and lowering voltage. A voltage alteration may be done to all components in a system or just to specific electronics components that are exceeding normal cooling limits.

[0044] Under circumstances where additional leakage currents due to hotter CMOS circuit temperatures cause concern of either heating the MCM beyond its safe operating temperature range or requires additional current than the DCAs are able to provide, the PCC **210** lowers the voltage applied to the CMOS circuits when a temperature degradation occurs. The effect on the "cooling degrade state" is to hasten its arrival as the combination of lower voltage and warmer circuits requires faster clock speed adjustments.

[0045] The PCC **210** takes into account both the MCM temperatures and applied voltage when it notifies the CSA code **220** of a change in "cooling state". The PCC **210** continually monitors the MCM thermistors **276** and provides the MRU with information if a sensor value is erroneous as well as the actual good values.

[0046] The PCC **210** sends the message to the CSA code **220** when the first degrade state is reached, indicating that the primary cooling system is not functioning normally. When it has been determined that this degrade state is due to a failure of the cooling hardware, the PCC **210** sets a fault flag for the primary cooling system, which is not removed until the primary cooling system is repaired. The PCC **210** posts this interrupt to the CSA code **220**.

[0047] The PCC **210** automatically turns on the backup cooling blowers or cooling fans **254** if the temperatures are above acceptable levels for the primary cooling system. The fan speeds are controlled in such a manner that the MCM temperature will not oscillate between cooling states unless the room ambient also oscillates.

[0048] The PCC **210** turns on the backup cooling blowers **254** at a speed to provide a temperature sufficiently above the temperature the first degrade state occurred so as to prevent "cooling state oscillation" when the backup blowers **254** are first turned on generally indicated at **298**. Steady state air cooling mode will be in degrade one or a slower degrade state, but if the backup blowers **254** are turned on immediately after the first degrade state is posted, then the additional backup cooling may cause a temporary spike down into the normal range temperature only to be soon followed by revisiting the first degrade state. It will be recognized by one skilled in the pertinent art that it is advantageous to minimize the occurrences of changing degrade states.

[0049] The PCC **210** continually samples the current and voltage being used by each MCM and communicates this power data to the MRU code as MCM powers state **284**. The PCC **210** also suitably adjust the power supply voltage levels at **296** being applied to the circuits. Raising the voltages will offset some of the speed lost by higher operating temperatures for some servers still operating in a safe temperature range and with extra power available from the power supply. For an MCM within server **240** which is operating near its upper temperature limit or for which the power supply has no additional current to supply, the PCC **210** either leaves the voltage unchanged or lowers it to reduce leakage currents in CMOS circuits. Hence, by sensing MCM temperatures and current being used by the MCM, the PCC **210** determines what if any voltage adjustment is suitable.

[0050] At all times, the existing temperatures and voltage conditions together define a suitable "thermal state" or range within which a specific set of clock speeds is optimum. The PCC **210** notifies the CSA code **220** of the proper speed range or "thermal state" that the MCMS are operating in at all times at **292**. This speed range may also be called a degrade state as described above.

[0051] The PCC **210** maintains a cooling state for each MCM available for the CSA code **220** to monitor at any time. The PCC **210** also provides periodic redundancy checks to insure that the backup blowers **254** are operating properly. When a primary cooling source having a fault, such as an MRU, is repaired, the PCC **210** clears defect status registers set which are visible to the CSA code **220**. Likewise, the PCC **210** also sends an interrupt to the CSA code **220** if the primary cooling system, e.g., MRUs **250**, **252**, needs service.

[**0052**] The Cycle Steering Application (CSA) code **220** provides a fail-safe method of adjusting the clock speeds in an optimum manner when the cooling state changes. This method of clock speed adjustment includes determining if a cooling failure has been repaired prior to increasing the clock speeds to prevent oscillating clock speeds. It should be noted that the clock speed follows the temperature and voltage conditions at all times. Further, the time from a change of circuit temperature to a corresponding change in clock speed is slow enough that the temperatures of the circuits change minimally, less than about 1° C., during this process.

[**0053**] The CSA node **220** includes an interrupt handler that reads directly from the PCC **210** the cooling state of each MCM as well as receiving interrupts on these states.

[**0054**] For systems with multiple processor books or nodes, the CSA code **220** determined which MCM has the slowest cooling state. This is the state that governs the safe clock speed of the system indicated generally at **310** in FIG. **2**. The multiple clock boundaries on multiple oscillators with predefined ratios are always maintained.

[**0055**] The CSA code **220** determines if any cooling defective hardware registers are set whenever a cooling state is increased calling for a faster clock speed. If the hardware defect register is set, it means the cause of the cooling degradation has not yet been fixed and the change in cooling state is likely due to transient change in ambient or other transient conditions. Hence, the server clock speeds are not re-adjusted faster until the defective cooling hardware is replaced and the register cleared. This is true even after the machine is re-initial microcode loaded (reIMLed) or rebooted. If there is uncertainty in the cooling state due to communication problems, the slowest, safest cooling state is employed by the CSA code **220**.

[**0056**] When the CSA code **220** determines it is appropriate to make a change in several clock speeds, it alters the phase lock loops (PLL) on the clock synthesizers in a sequence of very small steps until its new targeted clock speed is reached generally indicated as **312**. The phase lock loops are stepwise changes always retaining the optimum operating ratio between the various clocks that may be affected. The steps are sufficiently small to pose no risk to proper operation due to change in clock ratios during this adjustment process.

[**0057**] Every step is performed in a two step commit algorithm, e.g., the current step and the next step PLL values are saved in a persistent storage concept made up by using SEEPROMS residing on the current and backup cage controller **262**, **262**. After the change is written to the PLL and read back for verification, the saved current value is updated. This is done to provide protection in case a speed change is interrupted by a cage controller switchover.

[**0058**] The width of the small steps taken on the phase lock loops is less than the normal jitter of the phase lock loop normal output. This allows the step variation not to be detected by the target clock receiving circuitry. In this manner, all of the affected clocks are stepped in small increments until the targeted clock speed is achieved.

[**0059**] The PLLs are on two oscillator cards **263**, one in charge, one in backup mode. At all times the optimum ratio between clocks is maintained as the phase lock loops are moved in minimal increments or decrements.

[**0060**] Prior to power good time, the CSA code **220** issues a "Pre-Cooling" command to insure that the MCM temperatures are in proper normal state prior to turning on the clocks. This also prevents a sudden surge of power from the CMOS logic beginning to switch. Without pre-cool, this could cause a quick degrade state to occur because the refrigerant system takes some time to get its cooling cycle established. When pre-cooled state is reached the PCC **210** notifies the CSA code **220** of the same and IML is initiated.

[**0061**] The PLLs are initially loaded with a pattern, which is hard wired on the cards and loaded in parallel at power good time. Normally, PLLs are loaded serially, but this is exposed to shift errors which would lead to wrong clock speed settings.

[**0062**] The exact process of initializing clocks includes first verifying the right oscillator card **263**. Then, the pattern matching the actual system speed is loaded into the line drivers and read back to insure that there are no errors or hardware failures. Next, the loaded and verified pattern is read into the phase lock loops, with this pattern again read back to be verified. Now the system clock is started using the phase lock loop output as input. At the completion of IML, the system is degraded to its slowest clock state and upgraded back to its normal state with the required number of small incremental steps to the phase lock loops. This insures that all necessary patterns can be loaded into the phase lock loops without system error. This process takes a fraction of a second to complete on every server that is IMLed.

[**0063**] The pattern to be loaded for speed adjustment purposes such as when going from one cooling state to another is generated by a set of digital I/O lines controlled by the FGAs DIO engines, which is a part of the cage controller (cc) hardware. The FGAs DIO engines are digital I/O lines controlled by cage controller code that interface to the PLLs that control the system oscillators **263**. They are CSA code driven which is running on the PU Book **0** cage controller (cc). Before changing the PLL pattern due to a change in cooling state, the existing pattern is monitored to make sure the adjusting processes were not interrupted, by saving the line settings of the current pattern.

[**0064**] The CSA code **220** issues a warning service reference code (SRC) to the operator whenever the CSA code leaves normal clock speed. When the service is completed, the PCC **210** removes the error states and interrupts the CSA code **220**. The CSA code **220** removes SRC once notified.

[**0065**] The CSA code **220** monitors the actual speeds used for an IML to assure these speeds are never increased in actual operation even though the cooling state later permits the increased speed. The reason for this is that the initialization of "Elastic Interfaces" (EI) done during IML allows only for speed reduction and its clearing, not faster speeds than those present during IML initialization and self-tests.

[0066] Hence, the CSA code **220** notifies the operator that re-ILM should be avoided while a cooling failure service register is flagged so that when the cooling hardware problem is repaired, the server can return to its fast normal speed without needing a subsequent re-IML. Also contemplated is a repair and verify procedure that verifies that the clocks have returned to full speed while a customer engineer is present.

[0067] As a further enhancement on the above-described cooling system, a polytetrafluorethylene coating is employed on selected pressure drop areas of expansion structures within the vapor compression refrigeration system.

[0068] As noted, it has been discovered that material can agglomerate in certain pressure drop areas of the expansion structures within the refrigeration system. During refrigerant/oil transport, certain impurities and chemically reacted byproducts may come out of solution in the pressure drop areas as the refrigerant cools down. By way of example, FIGS. **4** & **5** depict part of an expansion valve, generally denoted **400**, which includes a first element **410** having an expansion orifice **430**, and a second element **420** having a tapered expansion pin **440**. As shown, the expansion pin **440** controls the amount of refrigerant passing through expansion orifice **430**, where refrigerant is assumed to flow left-to-right in the drawings illustrated. For the cooling applications described hereinabove, the expansion pin **440** is stepped open in very small increments to allow controlled flow of refrigerant through expansion orifice **430** into a pressure drop area defined between opposing surfaces **450** of elements **410** & **420**

[0069] During refrigerant/oil transport through a hot compressor, any long-chain molecules and other typically non-soluble compounds at room temperature can go into solution in the hot mixture. These, as well as other physically transported impurities, then fall out of the solution when the refrigerant/oil cools down, for example, in the pressure drop areas of the expansion structure. A layer of "waxy" material can build up in the pressure drop areas and act as a sticky substance which then catches other impurities. FIG. **6** depicts one example of an expansion pin **440** wherein contaminant material **460** has amassed in certain pressure drop areas of surfaces of the pin exposed to the refrigerant flow path. This amassing of material can interfere with the normal control volumes and interfere with the control of motor steps (e.g., due to unpredictable vavle characteristic changes). This is particularly true in a vapor compression refrigeration system employed as described above since the control of the expansion valves in this implementation is very sensitive and refrigerant expansion structure geometries are typically very small. Experimentation has shown that cleaning contaminant material from the pressure drop areas of expansion valves will typically fix any valve control problem resulting therefrom.

[0070] Thus, the solution presented herein is to apply a polytetrafluorethylene coating to at least portions of one or more expansion structures within the vapor compression refrigeration system in the pressure drop areas of the expansion structures. For example, FIG. **7** depicts a polytetrafluorethylene coating **770** over an expansion pin **700** of an expansion valve to be disposed within the vapor compression refrigeration system. In FIG. **8**, the polytetrafluorethylene coating is shown also disposed on the inner surface of element **710** defining expansion orifice **730** in the pressure drop area of the expansion valve defined between the opposing surfaces **750** of element **710** and element **720**, that is, the area which contains the tapered expansion pin **740** as shown. The polytetrafluorethylene coating can be applied to the exposed surfaces of a refrigerant expansion structure in the pressure drop area employing any conventional technique, such as vapor deposition. The polytetrafluorethylene coating has a thickness sufficient to inhibit the accumulation of material in any pressure drop area without changing a pressure drop characteristic of the pressure drop area. For example, if the expansion orifice is 30 mils in diameter, then the thickness of the polytetrafluorethylene coating may be 5 microns or less. Again, the goal of applying a polytetrafluorethylene coating is to make the exposed surfaces sufficiently slippery in the pressure drop areas of the expansion structures to inhibit the agglomeration of material onto those surfaces. This goal is achieved by the combination of refrigerant force through the pressure drop area and the surface energy properties of the polytetrafluorethylene, which together will reduce or eliminate contaminants from agglomerating.

[0071] Although preferred embodiments have been depicted and described in detail herein, it will be apparent to those skilled in the relevant art that various modifications, additions, substitutions and the like can be made without departing from the spirit of the invention and these are therefore considered to be within the scope of the invention as defined in the following claims.

What is claimed is:

1. A cooling system for cooling at least one heat generating electronics component, the cooling system comprising:

a vapor compression refrigeration system, the vapor compression refrigeration system comprising a condenser, at least one expansion structure, at least one evaporator, and a compressor coupled in fluid communication to define a refrigerant flow path and allow the flow of refrigerant therethrough; and

wherein the at least one evaporator facilitates removal of heat produced by the at least one heat generating electronics component, and wherein at least a portion of the at least one expansion structure is coated with a polytetrafluorethylene in the refrigerant flow path for inhibiting accumulation of material thereon.

2. The cooling system of claim 1, wherein the at least a portion of the at least one expansion structure comprises a pressure drop area of the at least one expansion structure.

3. The cooling system of claim 2, wherein the vapor compression refrigeration system comprises multiple expansion structures coupled in the refrigeration path, each expansion structure comprising a pressure drop area coated with a polytetrafluorethylene in the refrigerant flow path.

4. The cooling system of claim 2, wherein the polytetrafluorethylene coating has a thickness sufficient to inhibit accumulation of material in the pressure drop area without changing a pressure drop characteristic of the pressure drop area.

5. The cooling system of claim 1, wherein the at least one expansion structure comprises an expansion valve including an expansion pin and an expansion orifice defining a pressure drop area, and wherein the pressure drop area is coated with a polytetrafluorethylene in the refrigerant flow path.

6. The cooling system of claim 5, wherein the expansion valve is an electronic expansion valve.

7. A vapor compression refrigeration cooling system for cooling at least one heat generating electronics component, the cooling system comprising:

a condenser;

a first electrically controlled expansion valve coupled to the condenser;

a first evaporator coupled to the first electrically controlled expansion valve;

a second electrically controlled expansion valve coupled to the condenser;

a second evaporator coupled to the second electrically controlled expansion valve;

a controller providing control signals to the first electrically controlled expansion valve and the second electrically controlled expansion valve to control operation of the first electrically controlled expansion valve and the second electrically controlled expansion valve;

a compressor coupled to the first evaporator, the second evaporator and the condenser; and

wherein the condenser, the first electrically controlled expansion valve, the first evaporator, the second electrically controlled expansion valve, the second evaporator, and the compressor are coupled in fluid communication to define multiple refrigerant flow paths, each refrigerant flow path allowing the flow of refrigerant therethrough, and wherein the first evaporator and the second evaporator facilitate removal of heat produced by the at least one heat generating electronics component, and wherein at least a portion of the first electrically controlled expansion valve and at least a portion of the second electrically controlled expansion valve are coated with a polytetrafluorethylene in respective refrigerant flow paths for inhibiting accumulation of material thereon.

8. The cooling system of claim 7, wherein the at least a portion of the first electrically controlled expansion valve comprises a pressure drop area of the first electrically controlled expansion valve, and wherein the at least a portion of the second electrically controlled expansion valve comprises a pressure drop area of the second electrically controlled expansion valve.

9. The cooling system of claim 8, wherein the pressure drop areas comprise areas where refrigerant expansion occurs during a vapor compression cycle of the vapor compression refrigeration system.

10. The cooling system of claim 8, wherein the polytetrafluorethylene coating has a thickness sufficient to inhibit accumulation of material in the pressure drop areas without changing pressure drop characteristics of the pressure drop areas.

11. The cooling system of claim 7, wherein the first electrically controlled expansion valve comprises a first expansion pin and a first expansion orifice defining a first pressure drop area, and wherein the second electrically controlled expansion valve comprises a second expansion pin and a second expansion orifice defining a second pressure drop area, and wherein the first pressure drop area and the second pressure drop area are coated with a polytetrafluorethylene in the refrigerant flow path.

12. The cooling system of claim 7, wherein the cooling system is for cooling multiple heat generating electronics components, and wherein the first evaporator facilitates removal of heat produced by a first electronics component of the multiple heat generating electronics components and the second evaporator facilitates removal of heat produced by a second electronics component of the multiple heat generating electronics components.

13. A method of fabricating a vapor compression refrigeration system for cooling at least one heat generating electronics component, the method comprising:

(i) providing a condenser, at least one expansion structure, at least one evaporator, and a compressor;

(ii) providing a polytetrafluorethylene coating on at least a portion of the at least one expansion structure;

(iii) coupling the condenser, at least one expansion structure, at least one evaporator and compressor in fluid communication to define a refrigerant flow path; and

(iv) providing refrigerant within the refrigerant flow path of the vapor compression refrigeration system to allow for cooling of the at least one heat generating electronics component employing sequential vapor compression cycles, wherein the polytetrafluorethylene coating is provided on the at least a portion of the at least one expansion structure in the refrigerant flow path for inhibiting the accumulation of material thereon.

14. The method of claim 13, wherein the providing (ii) comprises providing the polytetrafluorethylene coating on a pressure drop area of the at least one expansion structure.

15. The method of claim 14, wherein the providing (i) comprising providing multiple expansion structures, and wherein the coupling (iii) comprises coupling the multiple expansion structures in the refrigerant flow path, each expansion structure comprising a pressure drop area coated with a polytetrafluorethylene in the refrigerant flow path.

16. The method of claim 14, wherein the providing (ii) comprises providing the polytetrafluorethylene coating with a thickness sufficient to inhibit accumulation of material in the pressure drop area without changing a pressure drop characteristic of the pressure drop area.

17. The method of claim 13, wherein the providing (i) comprises providing an expansion valve as the at least one expansion structure, the expansion valve including an expansion pin and an expansion orifice defining a pressure drop area, and wherein the providing (ii) comprises providing the polytetrafluorethylene coating in the pressure drop area in the refrigerant flow path.

18. The method of claim 17, wherein the providing (i) comprises providing an electronic expansion valve as the expansion valve.

* * * * *