

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
7 August 2008 (07.08.2008)

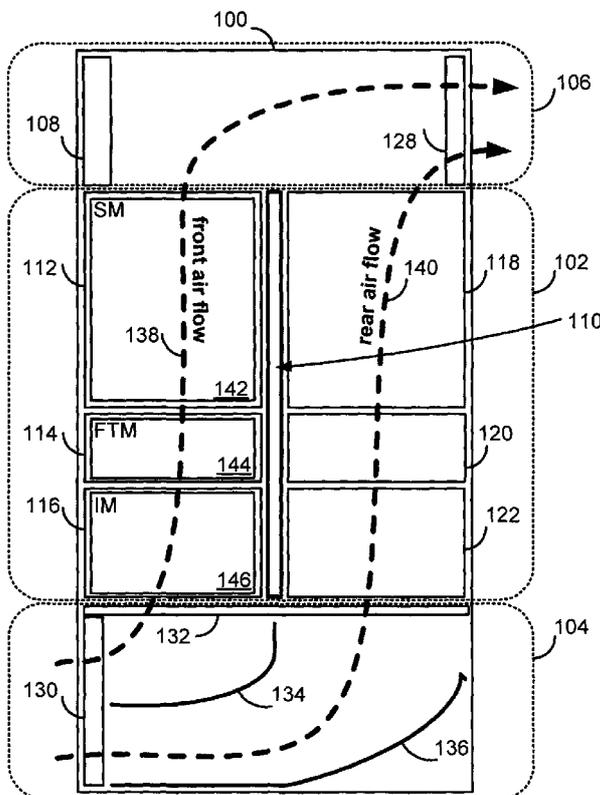
PCT

(10) International Publication Number
WO 2008/092247 A1

- (51) **International Patent Classification:**
G06F 1/20 (2006.01)
- (21) **International Application Number:**
PCT/CA2008/000180
- (22) **International Filing Date:** 29 January 2008 (29.01.2008)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
11/669,897 31 January 2007 (31.01.2007) US
- (71) **Applicant (for all designated States except US):** **LIQUID COMPUTING CORPORATION et al** [CA/CA]; 340 Terry Fox, Suite 300, Ottawa, Ontario K2K 3A2 (CA).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **BISSON, Sylvio** [CA/CA]; 100 du Vieux Rouet Gatineau, Secteur Aylmer,
- (74) **Agents:** **LEDWELL, Kent, M.** et al.; Gowling Lafleur Henderson LLP, 160 Elgin Street, Suite 2600, Ottawa, Ontario K1P 1C3 (CA).
- (81) **Designated States (unless otherwise indicated, for every kind of national protection available):** AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW
- (84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,

[Continued on next page]

(54) **Title:** COOLING HIGH PERFORMANCE COMPUTER SYSTEMS



<Fig. 1a

(57) **Abstract:** A computer system may include a chassis defining a front and a rear. The chassis may include a vertically oriented midplane disposed therein, the midplane including a plurality of front module slots for receiving front electronic modules from the front of the chassis, and a plurality of rear module slots for receiving rear electronic modules from the rear of the chassis. A cooling system may be provided within the chassis and may generate an upwardly-directed front air flow within the chassis directed at selected ones of the front electronic modules and an upwardly-directed rear air flow within the chassis directed at selected ones of the rear electronic modules. The front air flow is separate from and independent of the rear air flow. The selected front and rear electronic modules may be disposed in the chassis so as to separate the front air flow into a plurality of substantially equal front air streams and the rear air flow into a plurality of substantially equal rear air streams, respectively.

WO 2008/092247 A1



ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,
FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL,
NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG,
CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report*

COOLING HIGH PERFORMANCE COMPUTER SYSTEMS**BACKGROUND OF THE INVENTION****Field of the Invention**

[0001] Embodiments of the present invention relate generally to cooling systems for high performance computer systems and components thereof.

Description of the Prior Art and Related Information

[0002] The present invention relates to cooling of high performance electronic systems, and computer systems in particular, such as described in the commonly assigned and co-pending U.S. Patent Application No. 11/530,410, filed September 8, 2006, entitled "Methods and Systems for Scalable Interconnect", which application is hereby incorporated herein by reference in its entirety. Briefly, electronic components of such high performance computer systems are housed in modules that are located in one or more cages (card cages). In turn, these cages are disposed in equipment chassis or cabinets.

[0003] Simple air circulation around the electronic components that generate heat is often not sufficient, and methods of assisted cooling are required. Conventional cooling methods include: equipping heat generating components with heat sinks and fans; directing air by fans and baffles at the heat sinks; heat pipes; and liquid cooling such as described in U.S. Patent No. 4,561,040 in which a cooling liquid is circulated through a heat sink that is attached to the device. A more recent example of liquid cooling available for electronic devices is found in http://www.alienware.com/intro_pages/liquidcooling.aspx.

[0004] In conventional packaging of an electronic system in an equipment chassis or cabinet, with electronic components mounted on modules, air is circulated through the system

using blowers (fans). The fans may be mounted, for example, in the bottom of the chassis (or cabinet) and blow air from an air inlet upward to pass over the modules. The blown air, now heated by the modules, escapes the chassis or cabinet through an air outlet that may be disposed at or near the top of the chassis.

[0005] The chassis may be located side-by-side with other chassis in rows, separated by aisles. Conventionally, cool air may be provided by the building system in "cool" aisles while the warm air, having passed through the chassis, is exhausted into "warm" aisles. The ASHREA (American Society of Heating, Refrigerating and Air-Conditioning Engineers) and NEBS (Network Equipment Building Standards) provide guidelines for the installation and cooling aspects of equipment chassis or cabinets.

[0006] Baffles in the chassis may be used to direct the air flow from the "cool" aisle into the chassis, air filters are generally used in the air inlet, and the chassis may be enclosed to confine the airflow over the components that are to be cooled. Air filters may be used in the air inlet to insure that only filtered air passes over the components to be cooled.

[0007] The electronic modules may be open at the bottom and at the top, which allows air to pass directly over the electronic components and their respective heat sinks.

[0008] While such conventional methods of cooling rack mounted electronic systems are fairly well established in the industry, the increasing power dissipation of high performance electronic components requires more efficient cooling methods.

SUMMARY OF THE INVENTION

[0009] It is therefore an object of the present invention to develop a comprehensive cooling architecture to provide improved cooling of densely packed high power electronic components.

[0010] According to an embodiment thereof, the present invention is a computer system. The computer system may include a chassis defining a front and a rear, the chassis including a vertically (for example) oriented midplane, the midplane including a plurality of front module slots for receiving front electronic modules from the front of the chassis, and a plurality of rear module slots for receiving rear electronic modules from the rear of the chassis; a cooling system, the cooling system being configured to generate an upwardly-directed front air flow within the chassis directed at selected ones of the front electronic modules and being configured to generate an upwardly-directed rear air flow within the chassis directed at selected ones of the rear electronic modules, the front air flow being separate from and independent of the rear air flow.

[0011] The selected front electronic modules may dissipate more power than remaining ones of the front electronic modules and the selected rear electronic modules may dissipate more power than remaining ones of the rear electronic modules. The cooling system may include one or more front fans disposed between the selected front electronic modules and remaining ones of the front electronic modules and one or more rear fans disposed between the selected rear electronic modules and remaining ones of the rear electronic modules. The front fan(s) may be configured to draw air past the remaining front electronic modules and to accelerate the drawn air to form the front air flow. The rear fan(s) may be configured to draw air past the remaining rear electronic modules and to accelerate the drawn air to form the rear air

flow. The cooling system may include front fans disposed below the selected front electronic modules and rear fans disposed below the selected rear electronic modules. The cooling system may include an inlet separator baffle and an exhaust separator baffle. The inlet separator baffle, the midplane and the exhaust separator baffle may be collectively configured in the chassis such that the front air flow and the rear air flow are isolated from one another and independently controlled. The cooling system may include a front fan tray module coupled to the chassis for generating the upwardly-directed front air flow and a rear fan tray module coupled to the chassis for generating the upwardly directed rear air flow. The front fan tray module and the rear fan tray module may be each removably coupled to the chassis. The front electronic modules may include front function modules and front interconnect modules. The rear electronic modules may include rear function modules and rear interconnect modules. The front fan tray modules may include a number of front fans that is less than the number of the front function modules and the rear fan tray module may include a number of rear fans that is less than the number of the rear function modules. The front air flow may generate a front plenum from which the front air flow divides into a plurality of front air streams by the selected front electronic modules. Similarly, the rear air flow may generate a rear plenum from which the rear air flow divides into a plurality of rear air streams by the selected rear electronic modules. The selected front electronic modules may be or may include front function modules and the selected rear electronic modules may be or include rear function modules. The plurality of front air streams may be substantially equal to one another and the plurality of rear air streams may be substantially equal to one another. Each of the front and rear fan modules may include at least a first and a second fan and the front plenum may be disposed between the at least first and second front fan of the front fan module

and the selected front electronic modules, and the rear plenum may be disposed between the at least first and second rear fan of the rear fan module and the selected rear electronic modules.

[0012] The computer system may further include a plurality of sets of louvers, each of the plurality being disposed between an underlying one of the first and second fans of the front and rear fan tray modules and the front and rear plenums, respectively. The louvers of each of the plurality of sets of louvers may be configured to pivot to an open configuration when their respective underlying fan are operating and to pivot to a closed configuration when their respective underlying fan are not operating. The louvers of each of the plurality of sets of louvers may be configured to pivot to the closed configuration by a combination of gravity and the pressure of the air in either the front or rear plenum. The louvers of each set of louvers may be radially aligned with the center of their respective underlying fan. The sets of louvers disposed between the first and second fans of the front fan tray module and the front plenum may be pivotally coupled to a front louver tray configured to fit over the front fan tray module. The sets of louvers disposed between the first and second fans of the rear fan tray module and the rear plenum may be pivotally coupled to a rear louver tray configured to fit over the rear fan tray module.

[0013] The computer system may also include a plurality of front function modules received into selected ones of the front module slots and a plurality of rear function modules received into selected ones of the rear module slots, each of the front and rear function modules including: an electronic component, and a liquid cooling sub system including a cool plate coupled to the electronic component to absorb heat generated by the electronic component; a pump; a fluid-to-air heat exchanger; tubing connecting the cool plate, the pump and the heat exchanger in a closed circuit, and closed circuit cooling fluid that is pumped through the closed

circuit to transport heat from the cool plate to the heat exchanger, and from the closed circuit cooling fluid to the front or rear air flow. Each of the front and rear function modules may have, for example, a six-sided volume shape that may be substantially closed on first to fourth sides, open on a fifth side for admitting the air flow, and terminated by the heat exchanger on a sixth side that is opposite the fifth side, such that most of the air flow that is directed thereto flows through the heat exchanger.

[0014] According to another embodiment, the present invention is a computer system that may include a chassis; a plurality of electronic modules to be cooled mounted within the chassis; at least two axial fans that may be horizontally mounted in a fan tray module within the chassis, the at least two fans being configured to generate an air flow into and pressurize a plenum within the chassis, the plenum being located between the fans and the plurality of electronic modules to be cooled.

[0015] The computer system may also include a set of pivotally coupled passive louvers disposed over an underlying one of each of the at least two axial fans, between the underlying axial fan and the plenum, the louvers of each set of louvers being configured to pivot to an open position when its underlying fan is operating, and to pivot to a closed position when its underlying fan is stopped. The louvers of each set of louvers may be configured to pivot to the open position when sufficient air flow is generated by its underlying fan and to pivot to the closed position when its underlying fan ceases to operate by a combination of gravity and the air pressure in the plenum generated by a remaining operating fan or fans.

[0016] The present invention, according to still another embodiment thereof, is a computer system comprising: a chassis; a plurality of N electronic modules mounted within the chassis; a plurality of M axial fans that are collectively configured to generate a combined air

flow into an air plenum within the chassis for providing individual cooling air streams to a each of the plurality of N of electronic modules, wherein $N > M$ and wherein the plurality of N electronic modules are disposed within the chassis such that the air plenum divides the combined air flow into N substantially equal individual air streams.

[0017] The air plenum within the chassis may be configured to divide the combined air flow into substantially equal individual air streams, even with only M-I of the M axial fans operating. Airflow generated by each the N fans may pass through one of N sets of louvers such that each of the M axial fans that is operating holds the corresponding louver arrangement open, but failure of a fan causes the corresponding louver arrangement to close, to restrict an amount of air of the combined air flow that escapes through the failed fan.

[0018] Another embodiment of the present invention is a method of cooling a computer system, comprising: providing a chassis defining a front and a rear, the chassis including a vertically (for example) oriented midplane, the midplane including a plurality of front module slots for receiving front electronic modules from the front of the chassis, and a plurality of rear module slots for receiving rear electronic modules from the rear of the chassis; generating an upwardly-directed front air flow within the chassis directed at selected ones of the front electronic modules, and generating an upwardly-directed rear air flow within the chassis directed at selected ones of the rear electronic modules, such that the front air flow is separate from and independent of the rear air flow.

[0019] The front cooling flow generating step may include providing front fans below the selected front electronic modules and the rear cooling flow generating step may include providing rear fans disposed below the selected rear electronic modules. The method may further include providing an inlet separator baffle and an exhaust separator baffle such that the inlet

separator baffle, the midplane and the exhaust separator baffle are configured such that the front air flow and the rear air flow are isolated from one another and independently controlled. The front air generating step may include providing a front fan tray module coupled to the chassis and the rear air generating step may include providing a rear fan tray module coupled to the chassis. The providing step may be carried out with the front electronic modules including front higher power dissipating modules and front lower power dissipating modules, with the rear electronic modules including rear higher power dissipating modules and rear lower power dissipating modules and the front air generating step generating step may be carried out with the front fan tray modules including a number of front fans that is less than the number of the front higher power dissipating modules and the rear air generating step may be carried out with the rear fan tray module including a number of rear fans that is less than the number of the rear higher power dissipating modules. The front air flow generating step may generate a front plenum from which the front air flow divides into a plurality of front air streams by selected ones of the front electronic modules and the rear air flow generating step may generate a rear plenum from which the rear air flow divides into a plurality of rear air streams by selected ones of the rear electronic modules. The front air flow generating step may be carried out with the selected front electronic modules being configured to dissipate more power than remaining ones of the front electronic modules and the front air flow generating step may be carried out with the selected rear electronic modules being configured to dissipate more power than remaining ones of the rear electronic modules.

[0020] The front air flow generating step may be carried out with the plurality of front air streams being substantially equal to one another and the rear air flow generating step may be carried out with the plurality of rear air streams being substantially equal to one another. The

front air flow generating step may include providing a first front fan and a second front fan and generating a front plenum between the first and second front fans and selected ones of the front module slots, and the rear air flow generating step may include providing a first rear fan and a second rear fan and generating a rear plenum between the first and second rear fans and selected ones of the rear module slots. The selected front module slots may be configured to receive electronic modules that may dissipate more power than remaining ones of the front electronic modules and the selected rear module slots may be configured to receive electronic modules that may dissipate more power than remaining ones of the rear electronic modules. The method may also include steps of providing first and second sets of front louvers disposed between an underlying one of the first and second front fans, respectively, and the front plenum and providing first and second sets of rear louvers disposed between an underlying one of the first and second rear fans, respectively, and the rear plenum. Constituent louvers of the sets of front and rear louvers may be configured to pivot to an open configuration when their respective underlying fan is/are operating and to pivot to a closed configuration when their respective underlying fan is/are stopped. The louvers of each of the first, second, front and rear sets of louvers may be configured to pivot to the closed configuration by a combination of gravity and the pressure of the air in either the front or rear plenum. The method may also include a step of restricting the amount of air that leaks back through a stopped one of the first, second, front and rear fans. The method may also include a step of providing a plurality of front function modules received into selected ones of the front module slots and a plurality of rear function modules received into selected ones of the rear module slots. Each of the front and rear function modules may include an electronic component, and a liquid cooling sub system including a cool plate coupled to the electronic component to absorb heat generated by the electronic component; a

pump; a fluid-to-air heat exchanger; tubing connecting the cool plate, the pump, and the heat exchanger in a closed circuit. Closed circuit cooling fluid that may be pumped through the closed circuit to transport heat from the cool plate to the heat exchanger, and from the closed circuit cooling fluid to the front or rear air flow. Each of the front and rear function modules may, for example, have a six-sided volume shape that may be substantially closed on first to fourth sides, open on a fifth side for admitting the air flow, and terminated by the heat exchanger on a sixth side that is opposite the fifth side, such that most of the air flow that may be directed thereto flows through the heat exchanger.

[0021] Still another embodiment of the present invention is a method of cooling electronic modules within a computer system, comprising: providing a chassis; providing a backplane having a plurality of modules slots disposed away from one another in a vertical (for example) orientation and mounting the backplane in the chassis; coupling a plurality of electronic modules into corresponding ones of the plurality of module slots; generating an upwardly-directed air flow into an air plenum within the chassis below the electronic modules, and generating substantially equal individual cooling air streams from the air plenum to a space between the plurality of electronic modules.

[0022] The generating step may be carried out by a plurality of fans and the individual cooling air streams providing step may provide substantially equal individual cooling air streams even when fewer than all of the plurality of fans are operating.

BRIEF DESCRIPTION OF THE DRAWINGS

[0023] Embodiments of the present invention will now be described relative to the Figures, in which:

[0024] Fig. 1a is a cross-sectional side view of a chassis 100 of a high performance computer system according to an embodiment of the invention;

[0025] Fig. 1b shows a front view of the chassis 100 of Fig. 1;

[0026] Fig. 1c shows a rear view of the chassis 100 of Fig. 1;

[0027] Fig. 2a is a perspective view of a fan tray module 200 according to an embodiment of the invention;

[0028] Figs. 2b is a perspective view of the fan tray module 200 of Fig. 2a;

[0029] Fig. 2c is another perspective view of the fan tray module 200 of Fig. 2a;

[0030] Fig. 2d is yet another perspective view of the fan tray module 200 of Fig. 2a;

[0031] Fig. 3 illustrates an air flow distribution in a part of the chassis 100 of Fig. 1;

[0032] Fig. 4 is a chart of simulated flow rate data of the air flow 138 of Fig. 3;

[0033] Fig. 5 is a chart of simulated flow rate data of the upper air streams 308 of Fig. 3;

[0034] Fig. 6 shows a physical model of a computing module 600 according to an embodiment of a service module 112 of Fig. 1;

[0035] Fig. 7 is an exemplary thermal diagram 700 of the front part 300 of the electronics sub-system 102 of Fig. 3;

[0036] Fig. 8 illustrates a typical liquid cooled chip assembly 800 according to an embodiment of the invention; and

[0037] Fig. 9 is a graphical illustration of the specification (solid curve) and test results (diamond marks on are near the solid curve) of the overall performance of the liquid cooling subsystem 707 of Fig. 7.

DETAILED DESCRIPTION

[0038] The cooling system of the present invention has been developed for high performance computer systems (such as disclosed in, for example commonly assigned and co-pending U.S. Patent Application No. 11/530,410, incorporated by reference above). Such high performance computer systems may include one or more chassis. Each chassis may include, for example:

- a vertically oriented midplane, having "function" and "interconnect" slots;
- up to twenty (for example) high-power "service" (or "function") modules; and
- up to twenty (for example) lower-powered "interconnect" modules.

[0039] These components are mounted (plugged in) in cages within the chassis that may also include elements of the cooling system.

[0040] Fig. 1a (cut away side view), Fig. 1b (front view), and Fig. 1c (rear view) are simplified diagrams of a chassis 100 of a high performance computer system according to an embodiment of the invention. The chassis 100 may be housed in an equipment frame or a cabinet 101 shown in a dashed outline. The dimensions of the chassis 100 may, for example, be approximately 19" wide (for mounting in a 19 inch rack), 45.5" high, and 40" deep.

[0041] The exemplary chassis 100 includes an electronics sub-system 102 in the middle part of the chassis, an air inlet area 104 in the lower part of the chassis, and an air exhaust

area 106 in the upper part. A breaker interface panel 108 is located near the front of the chassis 100 in the air exhaust area 106.

[0042] Other equipment outside the scope of this invention may be mounted in remaining space above or below the chassis 100 in a standard seven foot (7') equipment frame or cabinet.

[0043] In a first alternate embodiment of the invention, two chassis 100 may be mounted above one another in a single equipment frame or cabinet of sufficient height. In a second alternate embodiment of the invention, two chassis 100 with modified air inlet and outlets areas may be mounted above one another in a single equipment frame or cabinet of standard height (e.g. 7 feet). In either alternate embodiment, the size of the electronics sub-system 102 is unchanged, but in the second alternate embodiment, the heights of the air inlet and outlet areas are reduced.

[0044] As shown in the simplified cut-away side view of Fig. 1a, the electronics sub-system 102 may include a vertically oriented midplane 110 in the approximate center of the chassis 100, and two sets of cages (module cages) in the front and the rear of the midplane 110. The front of the midplane 110 may include, from top to bottom; a front service cage 112, a front fan tray cage 114, and a front interconnect cage 116. Similarly, the rear of the midplane 110 may include, from top to bottom; a rear service cage 118, a rear fan tray cage 120, and a rear interconnect cage 122.

[0045] As shown, the air inlet area 104 may include: a front grille 130; an air filter 132 and an inlet separator baffle 134. The air exhaust area 106 may include a rear grille 128. As shown, the air inlet area 104 may include a front grille 130; an air filter 132; an inlet separator baffle 134; and an inlet deflector baffle 136. Fan tray modules may be located within the front

and rear fan tray cage (114 and 120, respectively), to be described in more detail below, which include fans driving two separate cooling fluid (e.g., air, other gases, liquid) flows through the chassis 100.

[0046] In operation, the cooling system shown in Figs. Ia, Ib and Ic generates a front air flow 138 that enters the chassis near the bottom front of the chassis 100 through the front grille 130, and that travels successively through the air filter 132, the front interconnect cage 116, the front fan tray cage 114, and the front service cage 112, before exiting the chassis 100 at the top rear through the rear grille 128.

[0047] Similarly, a rear air flow 140 is generated and enters the bottom front of the chassis 100 (below the entry of the front air flow 138) through the front grille 130, and travels successively through the air filter 132, the rear interconnect cage 122, the rear fan tray cage 120, and the rear service cage 118, before exiting the chassis 100 near the top rear thereof through the rear grille 128 (below the exit front air flow 138).

[0048] As may be appreciated from Fig. Ia, the front air flow 138 is isolated from the rear air flow 140 by the inlet separator baffle 134 and the midplane 110. Such isolated front and rear air flows enables the front and back modules of the electronics sub-system 102 to be cooled independently of one another.

[0049] In the illustrated exemplary embodiment, the front service cage 112 and the rear service cage 118 provide support for inserting service modules (SM) 142; the front fan tray cage 114 and the rear fan tray cage 120 provide support for inserting fan tray modules (FTM) 144; and the front interconnect cage 116 and the rear interconnect cage 122 provide support for inserting interconnect modules (IM) 146.

[0050] The front view of the chassis 100 within the cabinet is shown in Fig. Ib, using the same reference numerals for corresponding items as in Fig. Ia, i.e. from the top to the bottom:

- the air exhaust area 106 including the breaker interface panel 108 (visible from the front);
- the electronics sub-system 102 including the front service cage 112, the front fan tray cage 114, and the front interconnect cage 116, and
- the air inlet area 104 including and the front grille 130.

[0051] The rear view of the chassis 100, illustrated in Fig. Ic, similarly shows, using the same reference numerals for corresponding items in Fig. Ia, i.e. from the top to the bottom:

- the air exhaust area 106 including the rear grille 128;
- the electronics sub-system 102 including the rear service cage 118, the rear fan tray cage 120, and the rear interconnect cage 122, and
- the air inlet area 104 including the air filter 132.

[0052] In the exemplary embodiment illustrated in the three views (Figs. Ia, Ib, and Ic), the electronics sub-system 102 provides plug-in space for functional modules within the three types of cages, such as, for example,

- ten service modules (SM) 142 in the front service cage 112;
- another ten service modules (SM) 142 in the rear service cage 118;
- two fan tray modules (FTM) 144 in the front fan tray cage 114;
- another two fan tray modules (FTM) 144 in the rear fan tray cage 120;
- ten interconnect modules (IM) 146 in the front interconnect cage 116, and
- another ten interconnect modules (IM) 146 in the rear interconnect cage 122.

[0053] To identify the individual slots in each cage, the modules are labeled (left to right):

- service modules SMO to SM9 in the front service cage 112;
- service modules SM10 to SM19 in the rear service cage 118;
- fan tray modules FTMO to FTM1 in the front fan tray cage 114;
- fan tray modules FTM2 to FTM3 in the rear fan tray cage 120;
- interconnect modules IMO to IM9 in the front interconnect cage 116, and
- interconnect modules IM10 to IM19 in the rear interconnect cage 122.

[0054] The cages provide the mounting support for the modules that are plugged into connectors (not shown) of the midplane 110 (shown in Fig. 1a, and as also shown in Fig. 10 of the U.S. Patent Application No. 11/530,410, incorporated by reference above).

[0055] The electronics sub-system 102 may be designed with forced air cooling to facilitate its thermal dissipation. The air inlet area 104 in the bottom of the chassis 100 below the electronics sub-system 102 allows air to enter from the front of the chassis 100. The air exhaust area 106 mounted in the top of the chassis 100 above the electronics sub-system 102, allows air to be exhausted through the rear of the chassis 100. The air moves in the form of two distinct air flows; namely the front air flow 138 and the rear air flow 140, through the cages of the electronics sub-system 102 from the bottom to the top, as indicated in the cut away side view of Fig. 1a.

[0056] The service modules 142 and the interconnect modules 146 may be configured as plug-in modules that carry electronic components which collectively dissipate sufficiently high power that cooling is required. The cooling system according to embodiments of the present invention may include two forced air flows (e.g., the front air flow 138 and the rear air flow

140), and may be designed to handle a total peak electrical load of about 17 kW of power within the chassis 100. It is to be understood, however, that the present cooling system may be readily designed to handle peak electrical loads that are greater or small than the exemplary amount of 17kW. The forced air flows (138 and 140) may be generated by fans located in the fan tray modules 114 (FTMO to FTM 3), each of which may include, for example, two axial fans (for example, axial fans manufactured by EBM Papst of Farmington, CT, part number EBM6448/2TDHHP).

[0057] Each service module 142 may further comprise a self-contained liquid cooling subsystem (to be described in greater detail below, Fig. 6). Briefly, the purpose of the liquid cooling subsystem is to transfer heat from certain electronic components to a heat exchanger mounted within the module. The heat exchanger is then exposed to the air flow (138 or 140) for cooling.

[0058] A fan tray module 200 according to an embodiment of the present invention is shown in the perspective drawings shown in Figs 2a, 2b, 2c, and 2d. Fig. 2a shows a fan tray module 200 with two axial fans 202 and 204, and a damper plate 206. The view shown in Fig. 2a is from the top rear (plug-in side) of the fan tray module 200. The damper plate 206 is shown lifted away from the axial fans 202 and 204 in order to give a view of the axial fans 202 and 204; fully assembled, the damper plate would be mounted directly adjacent the fans 202 and 204. The damper plate 206 may include, as shown, two circular fan covers 208 and 210. When assembled, the damper plate 206 is mounted so that the circular fan covers 208 and 210 cover the fans 202 and 204 respectively. Each fan cover (208 and 210) may include a set of radially disposed louvers 212, held within a circular opening in the damper plate 206, between a rim 214 of the opening and a hub 216 disposed in the center of the circular opening.

[0059] As shown in the magnified detail view, each louver 212 includes a blade 218 of a light weight solid material (such as sold under the trade name Makrolon or Valox)formed into a narrow truncated pie slice shape (for example). The louver 212 is designed to pivot along an axis 220 formed by one long edge of the blade 218. At the corners of the blade 218, the axis 220 is extended with short knobs (pegs) 222 and 224. Each louver 212 is inserted in a predetermined position between the rim 214 and the hub 216 of a fan cover (208 or 210) by squeezing the knobs 222 and 224 into corresponding predetermined indentations in the rim 214 and the hub 216, respectively.

[0060] Each fan cover (208 or 210) may include a set of, for example, 18 louvers 212 as described, and arranged so that each louver 212 can freely rotate around its axis 220 from a closed position to an open position. In the closed position, each louver 212 is positioned essentially in the plane of the damper plate 206, by virtue of being supported by its knobs 222 and 224, and resting through the force of gravity on the edge of the radially adjacent louver 212. In Fig. 2a, the louvers) 212 are shown in the closed position.

[0061] Fig. 2b shows the fan tray module 200 with the louvers 212 in the open position. As shown, the blade of each louver 212 is turned about 45 to 70 degrees from the closed position. In operation of the present cooling system, the (louvers 212 of a fan cover (208 or 210) may be pivoted into the open position shown in Fig. 2b and may be held in that open position by the force of the air flow generated by the corresponding fan 202 or 204, respectively. For clarity, the illustration in Fig. 2b shows the louvers 212 opening as much as 90 degrees from the closed position; however, the shape of the knobs (222 and 224) and the corresponding indentations in the rim 214 and the hub 216 should be designed to stop pivoting at about 45 to 70

degrees from the closed position to ensure the louvers 212 will close with gravity and not stay open when the fan is stopped.

[0062] Indeed, while the closed position of the louvers 212 is predisposed through the action of gravity, the louvers 212 of a fan may be held open through the upwardly-directed force imposed thereon by the airflow generated by the an operating fan. However, if the fan should stop, it is the intention of the design that the louvers 212 close as soon as possible, not only through the action of gravity but also as a result of the air pressure from above that is generated by the other fan (202 or 204) in the same fan tray module 200. Fans in the adjacent fan tray module in the same fan tray cage 114 are also expected to assist through the air pressure they contribute, in holding the louvers 212 of the stopped fan in the closed position.

[0063] Figs. 2c and 2d show the fan tray module 200 in a bottom front view, with the louvers) 212 in the closed and open position respectively. In these views, the damper plate 206 is shown lifted away from the axial fans 202 and 204 in order to provide a view of the damper plate 206. Fully assembled, however, the damper plate 206 would be mounted directly adjacent the fans 202 and 204.

[0064] Returning now to the views of the chassis 100 illustrated in Figs. 1a to 1c. Each fan tray module 144 (FTMO and FTMI) in the front fan tray cage 114 may be implemented as is the fan tray module 200 shown in and described relative to Figs. 2a, 2b and 2c. Each such fan tray module provides a vertical air flow. Similarly, the fan tray modules 144 (FTMO and FTMI) in the rear fan tray cage 120 may be implemented as shown in Figs. 1a and 1c.

[0065] Thus, the front air flow 138 (as a whole) may be generated by the four fans located in the fan tray modules 144 (FTMO and FTMI) located in the front of the midplane 110. This front air flow 138 is then substantially evenly divided as it passes through the ten

interconnect modules 146 (IMO to IM9) below the fan tray cage 114, and through the ten service modules 142 (SMO to SM9) above the fan tray cage 114.

[0066] Fig. 3 shows an aspect of the front air flow 138 within a front part 300 of the electronics sub-system 102 in more detail. The rear air flow 140 within the rear part of the electronics sub-system 102 is analogous to the front air flow 138, and is not shown in detail.

[0067] The front part 300 of the electronics sub-system 102 may include the front service cage 112, the front fan tray cage 114, and the front interconnect cage 116, each with their full complement of corresponding modules (service modules 142, fan tray modules 144, and interconnect modules 146).

[0068] After passing through the air filter 132 (not shown in Fig. 3), the front air flow 138 enters the front part 300 of the electronics sub-system 102. Confined by the side walls of the chassis 100, and the circuit boards and front plates of each of the ten interconnect modules 146 (IMO to IM9), and further restricted by the midplane 110, the front airflow 138 divides into ten lower air streams 302, each lower air stream 302 passing through the space of one interconnect module 146 in the front interconnect cage 116. Note that Fig. 3 is idealized in order to illustrate the concepts more clearly; it is not to scale and illustrates the air flow approximately only. In a physical system, the air flow is affected by the structural parts of the electronics sub-system 102 as well as the physical presence of components mounted within each of the modules themselves.

[0069] Each of the ten interconnect modules 146 (IMO to IM9) is thus cooled by approximately one tenth of the total front air flow 138. After leaving the interconnect modules 146, the lower air streams 302 combine again into the front air flow 138 in a fan intake space 304 within the lower part of the front fan tray cage 114. The air flow is then propelled upwards by the combined action of the fans in the fan tray modules 144, into a "plenum" 306 in the upper part of

the fan tray modules 144. The plenum 306 is a pressurized zone from which the air flow will again divide and stream into the ten service modules 142 in the front service cage 112, in the form of ten substantially equal upper air streams 308. Air pressure in the plenum 306 also aids in holding the louvers 212 above a failed (stopped) fan closed, thus preventing or reducing air leakage through the stopped fan which would result in an even greater pressure drop in the plenum 306 than the drop already caused by the reduction of active fans from four to three. Simulations and measurements have indicated that the pressure drop due to the failure of one fan is small. Simulated air flow rate data are presented in the written description of Figs. 4 and 5, below.

[0070] Each of the ten service modules 142 (SM0 to SM9) is cooled by approximately one tenth of the total front air flow 138. Each of the ten upper air streams 308 serves to cool some components mounted on the corresponding service module 142 directly, or through heat sinks mounted on those components. Other components on the service module 142 which are connected to the liquid cooling subsystem (to be described below) are cooled indirectly by the upper air stream 308 passing through the heat exchanger. Heat exchangers are shown in Fig. 3 symbolically with reference numeral 310 near the top of each service module 142. However, the heat exchanger 310 need not be mounted at the top of each service module. Indeed, a heat exchanger may be mounted, for example, at or near the bottom of each service module 142 or anywhere between the top and bottom of each service module. Likewise, the position of other components may be changed from that shown and described herein without departing from the scope and spirit of the present inventions.

[0071] It should be noted that the interconnect modules 146 may be designed to dissipate less power and require comparatively less cooling than the service modules 142, the cooling requirements of the service modules 142 thus being dominant.

[0072] Simulations of the combined system of a representative worst case model of the service module 142 have indicated that the overall cooling requirement of the entire chassis 100 (a peak of 17 kW) may be achieved if an individual air flow intensity of at least 70 cubic feet per minute (CFM) through the heat exchanger 310 of each of the service modules 142 is maintained.

[0073] Fig. 4 is a chart of simulated data of the total front air flow 138 of Fig. 3, and the contribution by each of the fan tray modules 144 (FTMO and FTMI). The histogram shows the air flow rate in cubic feet per minute (CFM), for the individual air flows generated by the fan tray modules FTMO and FTMI, and the combined air flow 138 which is the sum of the individual air flows.

[0074] The grey shaded set of columns indicates air flow rates with all four fans in the fan tray modules operating normally. The rate of the air flow 138 is approximately 1245 CFM, of which about 635 CFM is contributed by the FTMO and 610 CFM is contributed by the FTMI. The contribution from each fan tray module (FTMO and FTMI) is designed to be equal, but asymmetries in the geometry of the physical modules (which were modeled fairly accurately in the simulation) may give rise to the slightly different efficiencies of the two fan tray modules.

[0075] The second set of columns (solid black) indicates air flows when one of the two fans in the fan tray module FTMI has failed (stopped). As indicated in the chart, the total air flow 138 is now reduced to approximately 1165 CFM of which about 737 CFM is contributed by the FTMO (two fans operating) and 428 CFM is contributed by the single fan operating in FTMI .

[0076] The simulation indicates that the total air flow generated by three fans is only slightly reduced from the normal air flow generated by four fans, even though the fans operate at the same speed. In effect, the three non-failed fans individually each move more air into the common plenum 306 (Fig. 3), thus partially compensating for the lost air flow of the failed fan. This is due to the fact that the operating characteristics of the fans (at the selected flow rate) are such that a small drop in back pressure in the plenum 306 (due to the failure of one fan) can result in a significant increase in output flow from each of the other three fans.

[0077] As described above, the louvers 218 (Figs. 2a-d) of the circular fan cover above the failed fan (for example the circular fan cover 208 above the fan 202) close and thus restrict, substantially prevent or substantially reduce the escape of air from the plenum 306 downward through the failed fan.

[0078] Similarly, Fig. 5 is a chart of simulated data of the upper air streams 308 of Fig. 3, showing their distribution among the service modules 142 (SMO to SM9). The grey shaded set of columns indicate air flow rates when all four fans in the fan tray modules 144 (FTM1 and FTM2) operate normally. The total air flow 138 of approximately 1245 CFM (Fig. 4) is distributed among the service modules SMO to SM9 more or less evenly. Again a completely even distribution (ideally 124.5 CFM per service module) is not achieved due to geometrical factors.

[0079] When all four fans are operating (grey columns) the air flow rate in the upper air streams 308 ranges from a minimum of about 88 CFM (SM9) to a maximum of about 131 CFM (SM5). In order to maximize the flow rate of the upper air streams 308, leakage of air from the plenum 306 due to various factors may be minimized by well fitting parts and gaskets if necessary.

[0080] When (as shown in Fig. 4) one of the fans in the fan tray module FTMI fails, the overall air flow from the plenum 306 into the service modules is reduced, but as shown in the solid black columns of Fig. 5, the distribution of the flow rates of the upper air streams 138 is still quite good.

[0081] The total reduced air flow rate of about 1165 CFM is distributed among the service modules SMO to SM9 more or less evenly. With one of the four fans failed (as simulated) the air flow rate in the upper air streams 308 ranges from a minimum of about 77.3 CFM (SM9) to a maximum of about 123.6 CFM (SMO). As might be expected with only one fan operating in the right hand fan tray module FTMI, the drop of flow rate is more pronounced in the right hand set of service modules (SM 6 to SM 9) which are located above the right hand fan tray module FTMI. And, perhaps surprisingly, because of the higher output individually of each of the three non-failed fans, the air flow rate in some left hand modules (SMO to SM3) is actually increased.

[0082] The valuable contribution of the plenum 306, interposed between the fan tray modules 144 and the service modules 142, is that even with only three fans operating in a very asymmetrical configuration (two working fans below five service modules, only one working fan below the remaining five service modules), sufficient cooling air flow may be provided to all service modules. The foregoing description, detailing air flow design in the front part of the chassis 100, equally applies to the rear part of the chassis 100, being of equivalent design.

[0083] In order to preserve a horizontally balanced air flow, it is preferred that service filler modules (not shown) be provided in all slots in the service module cage 112 and 118 that do not have functional service modules 142 inserted therein. Similarly, interconnect filler modules should be provided in all slots in the service module cage 116 and 122 that do not have functional interconnect modules 146 inserted therein.

[0084] Additionally, filler modules (such as service filler modules and interconnect filler modules) also balance the air flow distribution by presenting equivalent air impedance to air flow; provide increased product safety, provide Electromagnetic Containment and Immunity Compliance (EMC/EMI) for the product, and serve as a placeholder during shipping and handling of the chassis.

[0085] Filler modules typically do not contain any electrical components and do not make any physical or electrical connection with any mid-plane connectors. There are no fan tray filler modules, as all systems require that all fan tray cages (114, 120) be populated with functional fan tray modules 144.

[0086] An example of a service module is the Compute Module (CM), a functional and system description of which is provided in the commonly assigned and co-pending application incorporated by reference above. The CM provides the core block for multi-processor processing and may include, for example, four processor chips, associated memory and HyperTransport channels to enable both inter-processor communications and networking through external I/O. Because of the large amount of heat generated in this module, special consideration is given to cooling this module. Furthermore, because of the required packaging density (2 x 10 CMs in a 19" cage), conventional air heat sinks cooled simply by forced air, are not feasible.

[0087] Fig. 6 shows a simplified physical model 600 of the service module 142 (based on the example of a compute module CM). Fig. 6 omits mechanical details, both in the interest of clarity and to emphasize the cooling aspects of the service module 142.

[0088] The exemplary embodiment of the service module 142, as shown in the physical model 600, includes a multilayer printed circuit board (PCB) 602. Attached or otherwise coupled to the PCB 602 are:

- a front plate 604;
- a number of connectors 606 for insertion into the midplane 110 (Fig. 1a);
- four liquid cooled chip assemblies (608, 610, 612, and 614);
- two pumps (616 and 618);
- a heat exchanger 620;
- ten tubing segments (622, 624, 626, 628, 630, 632, 634, 636, 638 and 640), and
- two exemplary memory assemblies 642.

[0089] Other components that may also be mounted on the PCB 602 are not shown in this simplified physical model 600 in the interest of clarity. Also, the service modules 142 need not be structured exactly as shown in Fig. 6 and as described above, as variations and modifications are possible, as those of skill in this art will recognize.

[0090] The physical model 600 is shown in the orientation in which it is operated, that is, the PCB 602 is disposed vertically, with the heat exchanger 620 mounted along the top, preferably occupying as much of the top edge (back to front) as possible, and preferably having a thickness sufficient to occupy virtually all available space between the present service module and its neighbor, in order to minimize the amount of air escaping around the heat exchanger 620. Additional measures, such as installing gaskets, may be taken in order to minimize air leakage and force as much of each upper air stream 308 (Fig. 3) to go through the heat exchanger 620 (310 in Fig. 3).

[0091] The two pumps (616 and 618) drive two closed liquid circuits which share the heat exchanger 620 as follows:

[0092] Each closed liquid circuit extends from an outlet 644 of the heat exchanger 620 ultimately back to an inlet 646 thereof.

[0093] The first closed liquid circuit, (passing through pump 616) extends through the tubing segment 622, into the pump 616 and out through the tubing segment 624 to an inlet 648 of the liquid cooled chip assembly 612. At the liquid cooled chip assembly 612, the first closed liquid circuit divides into a first part that passes straight through the inlet 648 and continues through the tubing segment 636 to an inlet 650 of the liquid cooled chip assembly 614. Thereafter, the first closed liquid circuit continues through the liquid cooled chip assembly 614 to an outlet 652 thereof. A second part of the first closed liquid circuit passes through the liquid cooled chip assembly 612 to an outlet 654 thereof, and continues through the tubing segment 638, to merge with the first part at the outlet 652. From the outlet 652 of the liquid cooled chip assembly 614, the remerged first closed liquid circuit then continues through the tubing segment 640 to the inlet 646 of the heat exchanger 620.

[0094] The second closed liquid circuit, (passing through pump 618) extends analogously from the outlet 644 of the heat exchanger 620 to its inlet 646 through the tubing segment 626 to the pump 618. The second closed liquid circuit reemerges from the pump 618 through the tubing segment 628 to an inlet 656 of the liquid cooled chip assembly 608. Thereafter, the second closed liquid circuit divides into a first part and a second part. The first part continues through the tubing segment 632 to an inlet 658 of the liquid cooled chip assembly 610, to the outlet 660 thereof. The second part leaves the liquid cooled chip assembly 608 through the outlet 662 thereof, continues through the tubing segment 630 where it merges with

the first part. The second part continues through the outlet 660 of the liquid cooled chip assembly 610 out through the tubing segment 634 to the inlet 646 of the heat exchanger 620. The arrows in Fig. 6 indicate the direction of flow of the cooling fluid. The simplified physical model 600 of the service module 142 shown in Fig. 6 is merely an exemplary embodiment of the invention. It is within the scope of the invention to include more or fewer components within the first, second or additional closed circuits by additional liquid circuit tubing, pumps, and cold plates, as required.

[0095] Shown in Fig. 7 is an exemplary thermal diagram 700 of the front part 300 of the electronics sub-system 102 (see Fig. 3), combined with details (further simplified from the physical model of Fig. 6) of one service module 142 (see Fig. 6). The same reference numerals as in Figs. 3 and 6 are used for corresponding components. The thermal design of the rear part of the electronics sub-system 102 may be analogous to the front part thereof shown in Fig. 7.

[0096] The thermal diagram 700 shows the front service cage 112 with the ten service modules 142, one of which is shown in the form of a simple thermal model 702; the front fan tray cage 114 with the two fan tray modules 144, one of which is shown in the form of a simple thermal model 704, and the front interconnect cage 116 with the ten interconnect modules 146, one of which is shown in the form of a simple thermal model 706.

[0097] Derived from the physical model 600 of the service module 142 (see Fig. 6) the thermal model 702 of the service module 142 includes a liquid cooling subsystem 707. The liquid cooling subsystem 707 may include, as shown, the heat exchanger 620; a liquid pumping system 708 (which includes the pumps 616 and 618 of Fig. 6); a feeder tubing system 710 (the tubing segments 622 and 626 of Fig. 6) from the heat exchanger 620 to the liquid computing system 710; an inlet tubing system 712 (which includes the tubing segments 624, 628, 632, and

636 of Fig. 6); a return tubing system 714 (the tubing segments 630, 634, 638, and 640 of Fig. 6), and a set of liquid cooled components 716 (the four liquid cooled chip assemblies 608 to 614 of Fig. 6). The set of liquid components 716 together dissipate, in this exemplary embodiment, a total of up to $4 \times 95 \text{ Watts} = 390 \text{ Watts}$. The thermal model 702 of Fig. 7 may also include a set of air cooled components 718 (which may include the memory assemblies 642 of Fig. 6, as well as other heat generating components not shown in Fig. 6) that together may dissipate up to about 385 Watts in the example developed herein.

[0098] The thermal model 704 of the fan tray module 144 may include a set of fans 720, which together may dissipate up to 36 Watts. The thermal model 706 of the interconnect module 146 may include another set of air cooled components 722, which may dissipate up to 100 Watts (1000 Watts collectively for the interconnect cage 116). Power dissipation numbers are exemplary worst case numbers and may deviate from typical observed values which are generally lower)

[0099] Air flows and air streams are shown in Fig. 7 symbolically in the form of white arrows. All temperatures shown in the diagram (in degrees Celsius) were obtained with a simulation of the thermal models.

[0100] The front air flow 138 enters the diagram at the bottom with an exemplary air temperature $T_1 = 35^\circ$, which is the highest temperature expected to be supplied according to the ASHREA standard (American Society of Heating, Refrigerating and Air-Conditioning Engineers). The front air flow 138 passes through the interconnect modules 146, in the form of the lower air streams 302 (see Fig. 3). Each of the lower air streams 302, flowing through its respective interconnect module 146, is heated by the another set of air cooled components 722,

and subsequently leaves the front interconnect cage 116 at an increased air temperature of $T_2 = 36.4^\circ$, whereupon each of the heated air streams 302 enters the fan intake space 304.

[0101] The fans in the front fan tray cage 114 propel the air from the fan intake space 304 into the plenum 306, at an exemplary rate of 1245 cubic feet per minute (CFM), thus in effect creating the front air flow 138. The temperature of the air in the plenum 306, after passing through the fans is $T_3 = 36.4^\circ$. The increase in the air temperature due to the dissipation in the fans is negligible.

[0102] From the plenum 306, the air flow divides into the upper air streams 308 shown in Fig. 3, each entering the space of a representative one of the service modules 142, as described earlier.

[0103] Within the representative service modules 142, the air may be further heated by the set of air cooled components 718 to an air temperature of $T_4 = 43.3^\circ$, when it reaches the heat exchanger 620.

[0104] The set of liquid cooled components 716, dissipating $4 \times 95 = 390$ Watts, transfers generated heat to the liquid circulating in the liquid cooling subsystem 707, as described earlier. The feeder tubing system 710 joins the heat exchanger 620 to the liquid pumping system 708; the inlet tubing system 712 joins the liquid pumping system 708 to the set 716 of liquid cooled components; and the return tubing system 714 joins the set of liquid cooled components 716 back to the heat exchanger 620. A cooling fluid (for example, water mixed with glycol) circulates through this closed liquid cooling circuit. In the present exemplary embodiment, the cooling fluid flows from the heat exchanger 620, through the liquid pumping system, 708 and the set of liquid cooled components 716, before returning to the heat exchanger

620. In another embodiment, the cooling fluid may flow in the opposite direction and/or follow a different path

[0105] In terms of cooling the set of liquid cooled components 716, a large proportion of their generated heat (at a component temperature $T5 = 68^\circ$) is efficiently transferred to the circulating cooling fluid, which may attain a liquid temperature of $T6 = 66.4^\circ$ as indicated in the diagram. A certain proportion of the heat is then transferred in the heat exchanger 620 from the circulating cooling fluid to the upper air stream 308 passing through the heat exchanger 620, thereby heating the air from $T4 = 43.3^\circ$ to an exhaust air temperature of $T7 = 51.7^\circ$ and at the same time cooling the circulating cooling fluid from $T6 = 66.4^\circ$ to a temperature of $T8 = 60.5^\circ$. When the cooled ($T8 = 60.5^\circ$) liquid circulates through the liquid pumping system 708 back to the set of liquid cooled components 716, the cycle repeats.

[0106] Fig. 8 illustrates a typical liquid cooled chip assembly 800 according to an embodiment of the invention (for example, the liquid cooled chip assemblies 608 of Fig. 6). The liquid cooled chip assembly 800 may include, for example, an AMD Opteron CPU chip 802 attached to a cold plate 804 for transferring heat from the CPU chip 802 to the cooling fluid which circulates through the cold plate 804. The cold plate 804 has an inlet 806 and an outlet 808 (corresponding to the inlet 656 and the outlet 662 of the liquid cooled chip assemblies 608) for attaching tubing segments (e.g. the tubing segments 628, 630, and 632 in the case of the liquid cooled chip assemblies 608 as described in Fig. 3 above). The liquid cooling components including the cold plates, tubing material, pumps, and heat exchanger, may be custom made to dimensional and performance specifications by, for example, Thermacore, Inc. of Lancaster PA.

[0107] Fig. 9 shows a graph illustrating the specification (solid curve) and test results (diamond marks on or near the solid curve) of the overall performance of the liquid cooling

subsystem 707 of Fig. 7. Shown in the diagram is the heat resistance R_{c-a} from the chip (the AMD Opteron CPU chip 802, Fig. 8) to the air flowing through the heat exchanger 620 in degrees Celsius per Watt, as a function of the amount of air flow (the upper air stream 308) through the heat exchanger 620 assuming an air intake temperature of $T_4 = 43.3^\circ$ (Fig. 8) at the heat exchanger 620. For example, as shown in the diagram of Fig. 9, the heat resistance has been determined to be 0.215 degrees/Watt at an assumed air flow rate of 90 cubic feet per minute (CFM). Thus, with an assumed chip dissipation of 95 W, and an air flow of 90 CFM, the chip temperature rise above the temperature T_4 of the air flow is expected to be $T_{de.i.ta} = 95 \times 0.215 = 20.4^\circ$. Added to the air temperature of $T_4 = 43.3^\circ$, it is then expected that a chip temperature $T_{ch.i.p} = T_4 + T_{de.i.ta} \approx 43.3^\circ + 20.4^\circ = 63.7^\circ$ would be reached, which is well within the operating margins of the device (the AMD Opteron CPU chip 802). As the amount of airflow modeled in Fig. 7 (the upper air stream 308) is 115 CFM, a margin exists to allow for air leakage and variation in air flow relative to the typical service module 142. Even when the uneven distribution between service module slots, as well as the possibility of fan failure, is taken into account (as depicted in the simulation results of Figs. 4 and 5), sufficient cooling of high-power chips is provided.

[0108] It may be appreciated that the numerical description of the thermal design has been provided as an example of a working design for handling particular heat loads that are largely determined by the electronic circuits in the various modules. Other combinations of heat dissipation and air flow rates may be appropriate in other cases, as will be readily apparent to skilled persons.

[0109] While the foregoing detailed description has described preferred embodiments of the present invention, it is to be understood that the above description is illustrative only and

not limiting of the disclosed invention. Those of skill in this art will recognize other alternative embodiments and all such embodiments are deemed to fall within the scope of the present invention. Thus, the present invention should be limited only by the claims as set forth below.

WHAT IS CLAIMED IS:

1. A computer system, comprising:

a chassis defining a front and a rear, the chassis including a vertically oriented midplane, the midplane including a plurality of front module slots for receiving front electronic modules from the front of the chassis, and a plurality of rear module slots for receiving rear electronic modules from the rear of the chassis;

a cooling system, the cooling system being configured to generate an upwardly-directed front air flow within the chassis directed at selected ones of the front electronic modules and being configured to generate an upwardly-directed rear air flow within the chassis directed at selected ones of the rear electronic modules, the front air flow being separate from and independent of the rear air flow.

2. The computer system of claim 1, wherein the selected front electronic modules dissipate more power than remaining ones of the front electronic modules and wherein the selected rear electronic modules dissipate more power than remaining ones of the rear electronic modules.

3. The computer system of claim 1, wherein the cooling system includes at least one front fan that is disposed between the selected front electronic modules and remaining ones of the front electronic modules and at least one rear fan that is disposed between the selected rear electronic modules and remaining ones of the rear electronic modules.

4. The computer system of claim 3, wherein the at least one front fan is configured to draw air past the remaining front electronic modules and to accelerate the drawn air to form

the front air flow and wherein the at least one rear fan is configured to draw air past the remaining rear electronic modules and to accelerate the drawn air to form the rear air flow.

5. The computer system of claim 1, wherein the cooling system includes front fans disposed below the selected front electronic modules and rear fans disposed below the selected rear electronic modules.

6. The computer system of claim 1, wherein the cooling system includes an inlet separator baffle and an exhaust separator baffle, and wherein the inlet separator baffle, the midplane and the exhaust separator baffle are collectively configured in the chassis such that the front air flow and the rear air flow are isolated from one another and independently controlled.

7. The computer system of claim 1, wherein the cooling system includes a front fan tray module coupled to the chassis for generating the upwardly-directed front air flow and a rear fan tray module coupled to the chassis for generating the upwardly directed rear air flow.

8. The computer system of claim 7, wherein the front fan tray module and the rear fan tray module are each removably coupled to the chassis.

9. The computer system of claim 7, wherein the front electronic modules include front function modules and front interconnect modules, wherein the rear electronic modules include rear function modules and rear interconnect modules and wherein the front fan tray modules includes a number of front fans that is less than a number of the front function modules and wherein the rear fan tray module includes a number of rear fans that is less than a number of the rear function modules.

10. The computer system of claim 1, wherein the front air flow generates a front plenum from which the front air flow divides into a plurality of front air streams by the selected front electronic modules and wherein the rear air flow generates a rear plenum from which the rear air flow divides into a plurality of rear air streams by the selected rear electronic modules.

11. The computer system of claim 10, wherein the selected front electronic modules are front function modules and wherein the selected rear electronic modules are rear function modules.

12. The computer system of claim 10, wherein the plurality of front air streams are substantially equal to one another and wherein the plurality of rear air streams are substantially equal to one another.

13. The computer system of claim 10, wherein each of the front and rear fan modules includes at least a first and a second fan and wherein the front plenum is disposed between the at least first and second front fan of the front fan module and the selected front electronic modules, and wherein the rear plenum is disposed between the at least first and second rear fan of the rear fan module and the selected rear electronic modules.

14. The computer system of claim 13, further including a plurality of sets of louvers, each of the plurality being disposed between an underlying one of the first and second fans of the front and rear fan tray modules and the front and rear plenums, respectively, and wherein the louvers of each of the plurality of sets of louvers are configured to pivot to an open configuration when their respective underlying fan is operating and to pivot to a closed configuration when their respective underlying fan is not operating.

15. The computer system of claim 14, wherein the louvers of each of the plurality of sets of louvers are configured to pivot to the closed configuration by a combination of gravity and a pressure of the air in either the front or rear plenum.

16. The computer system of claim 14, wherein the louvers of each set of louvers are radially aligned with a center of their respective underlying fan.

17. The computer system of claim 14, wherein the sets of louvers disposed between the first and second fans of the front fan tray module and the front plenum are pivotally coupled to a front louver tray configured to fit over the front fan tray module and wherein the sets of louvers disposed between the first and second fans of the rear fan tray module and the rear plenum are pivotally coupled to a rear louver tray configured to fit over the rear fan tray module.

18. The computer system of claim 1, further comprising a plurality of front function modules received into selected ones of the front module slots and a plurality of rear function modules received into selected ones of the rear module slots, each of the front and rear function modules including:

an electronic component, and

a liquid cooling sub system including a cool plate coupled to the electronic component to absorb heat generated by the electronic component; a pump; a fluid-to-air heat exchanger; tubing connecting the cool plate, the pump and the heat exchanger in a closed circuit, and closed circuit cooling fluid that is pumped through the closed circuit to transport heat from the cool plate to the heat exchanger, and from the closed circuit cooling fluid to the front or rear air flow.

19. The computer system of claim 18, wherein each of the front and rear function modules has a six-sided volume shape that is substantially closed on first to fourth sides, open on

a fifth side for admitting the air flow, and terminated by the heat exchanger on a sixth side that is opposite the fifth side, such that most of the air flow that is directed thereto flows through the heat exchanger.

20. A computer system, comprising:

a chassis;

a plurality of electronic modules to be cooled mounted within the chassis;

at least two axial fans that are horizontally mounted in a fan tray module within the chassis, the at least two fans being configured to generate an air flow into and pressurize a plenum within the chassis, the plenum being located between the fans and the plurality of electronic modules to be cooled.

21. The computer system of claim 20, further comprising a set of pivotally coupled passive louvers disposed over an underlying one of each of the at least two axial fans, between the underlying axial fan and the plenum, the louvers of each set of louvers being configured to pivot to an open position when its underlying fan is operating, and to pivot to a closed position when its underlying fan is stopped.

22. The computer system of claim 21, wherein the louvers of each set of louvers are configured to pivot to the open position when sufficient air flow is generated by its underlying fan and to pivot to the closed position when its underlying fan ceases to operate by a combination of gravity and an air pressure in the plenum generated by a remaining operating fan or fans.

23. A computer system comprising:

a chassis;

a plurality of N electronic modules mounted within the chassis;

a plurality of M axial fans that are collectively configured to generate a combined air flow into an air plenum within the chassis for providing individual cooling air streams to a each of the plurality of N of electronic modules, wherein $N > M$ and wherein the plurality of N electronic modules are disposed within the chassis such that the air plenum divides the combined air flow into N substantially equal individual air streams.

24. The computer system of claim 23, wherein the air plenum within the chassis is configured to divide the combined air flow into substantially equal individual air streams, even with only M-I of the M axial fans operating.

25. The computer system of claim 23, wherein airflow generated by each the N fans passes through one of N sets of louvers such that each of the M axial fans that is operating holds the corresponding louver arrangement open, but failure of a fan causes the corresponding louver arrangement to close, to restrict an amount of air of the combined air flow that escapes through the failed fan.

26. A method of cooling a computer system, comprising:

providing a chassis defining a front and a rear, the chassis including a vertically oriented midplane, the midplane including a plurality of front module slots for receiving front electronic modules from the front of the chassis, and a plurality of rear module slots for receiving rear electronic modules from the rear of the chassis;

generating an upwardly-directed front air flow within the chassis directed at selected ones of the front electronic modules, and

generating an upwardly-directed rear air flow within the chassis directed at selected ones of the rear electronic modules, such that the front air flow is separate from and independent of the rear air flow.

27. The method of claim 26, wherein the front cooling flow generating step includes providing front fans below the selected front electronic modules and wherein the rear cooling flow generating step includes providing rear fans disposed below the selected rear electronic modules.

28. The method of claim 26, further including providing an inlet separator baffle and an exhaust separator baffle such that the inlet separator baffle, the midplane and the exhaust separator baffle are configured such that the front air flow and the rear air flow are isolated from one another and independently controlled.

29. The method of claim 26, wherein the front air generating step includes providing a front fan tray module coupled to the chassis and wherein the rear air generating step includes providing a rear fan tray module coupled to the chassis.

30. The method of claim 29, wherein the providing step is carried out with the front electronic modules including front higher power dissipating modules and front lower power dissipating modules, with the rear electronic modules including rear higher power dissipating modules and rear lower power dissipating modules and wherein the front air generating step generating step is carried out with the front fan tray modules including a number of front fans that is less than a number of the front higher power dissipating modules and wherein the rear air generating step is carried out with the rear fan tray module including a number of rear fans that is less than a number of the rear higher power dissipating modules.

31. The method of claim 26, wherein the front air flow generating step generates a front plenum from which the front air flow divides into a plurality of front air streams by selected ones of the front electronic modules and wherein the rear air flow generating step generates a rear plenum from which the rear air flow divides into a plurality of rear air streams by selected ones of the rear electronic modules.

32. The method of claim 31, wherein the front air flow generating step is carried out with the selected front electronic modules being configured to dissipate more power than remaining ones of the front electronic modules and wherein the front air flow generating step is carried out with the selected rear electronic modules being configured to dissipate more power than remaining ones of the rear electronic modules.

33. The method of claim 31, wherein the front air flow generating step is carried out with the plurality of front air streams being substantially equal to one another and wherein the rear air flow generating step is carried out with the plurality of rear air streams being substantially equal to one another.

34. The method of claim 26, wherein the front air flow generating step includes providing a first front fan and a second front fan and generating a front plenum between the first and second front fans and selected ones of the front module slots, and wherein the rear air flow generating step includes providing a first rear fan and a second rear fan and generating a rear plenum between the first and second rear fans and selected ones of the rear module slots.

35. The method of claim 34, wherein the selected front module slots are configured to receive electronic modules that dissipate more power than remaining ones of the front electronic

modules and wherein the selected rear module slots are configured to receive electronic modules that dissipate more power than remaining ones of the rear electronic modules.

36. The method of claim 34, further including steps of providing a first and second sets of front louvers disposed between an underlying one of the first and second front fans, respectively, and the front plenum and providing a first and second sets of rear louvers disposed between an underlying one of the first and second rear fans, respectively, and the rear plenum, and wherein constituent louvers of the sets of front and rear louvers are configured to pivot to an open configuration when their respective underlying fan is operating and to pivot to a closed configuration when their respective underlying fan is stopped.

37. The method of claim 36, wherein the louvers of each of the first, second, front and rear sets of louvers are configured to pivot to the closed configuration by a combination of gravity and a pressure of the air in either the front or rear plenum.

38. The method of claim 34, further comprising a step of restricting an amount of air that leaks back through a stopped one of the first, second, front and rear fans.

39. The method of claim 26, further comprising a step of providing a plurality of front function modules received into selected ones of the front module slots and a plurality of rear function modules received into selected ones of the rear module slots, each of the front and rear function modules including:

an electronic component, and

a liquid cooling sub system including a cool plate coupled to the electronic component to absorb heat generated by the electronic component; a pump; a fluid-to-air heat exchanger; tubing connecting the cool plate, the pump, and the heat exchanger in a closed circuit, and closed circuit

cooling fluid that is pumped through the closed circuit to transport heat from the cool plate to the heat exchanger, and from the closed circuit cooling fluid to the front or rear air flow.

40. The method of claim 39, wherein each of the front and rear function modules has a six-sided volume shape that is substantially closed on first to fourth sides, open on a fifth side for admitting the air flow, and terminated by the heat exchanger on a sixth side that is opposite the fifth side, such that most of the air flow that is directed thereto flows through the heat exchanger.

41. A method of cooling electronic modules within a computer system, comprising:

- providing a chassis;
- providing a backplane having a plurality of modules slots disposed away from one another in a vertical orientation and mounting the backplane in the chassis;
- coupling a plurality of electronic modules into corresponding ones of the plurality of module slots;
- generating an upwardly-directed air flow into an air plenum within the chassis below the electronic modules, and
- generating substantially equal individual cooling air streams from the air plenum to a space between the plurality of electronic modules.

42. The method of claim 41, wherein the generating step is carried out by a plurality of fans and wherein the individual cooling air streams providing step provides substantially equal individual cooling air streams even when fewer than all of the plurality of fans are operating.

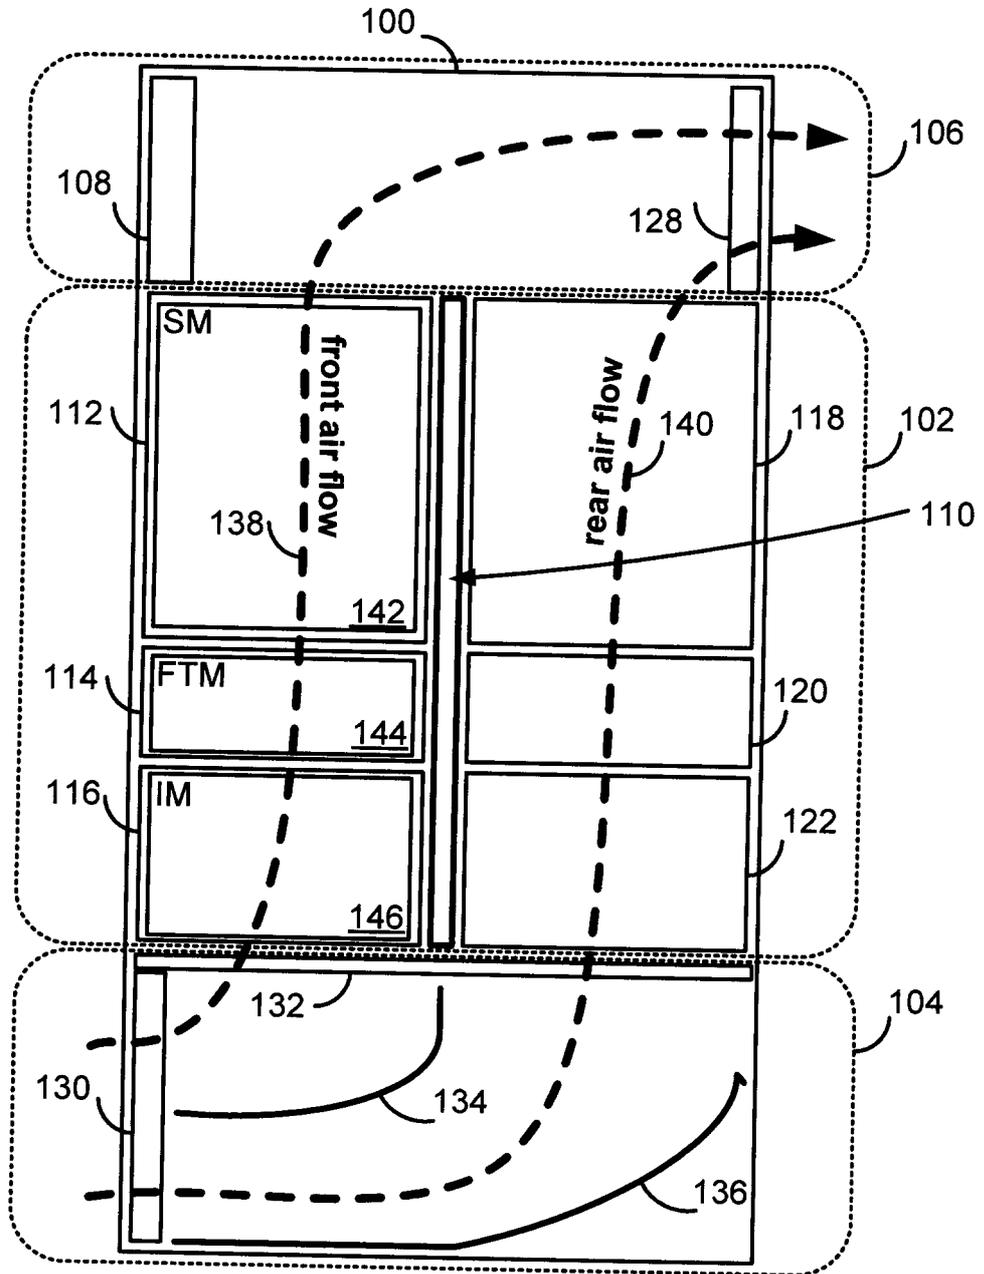


FIG. 1a

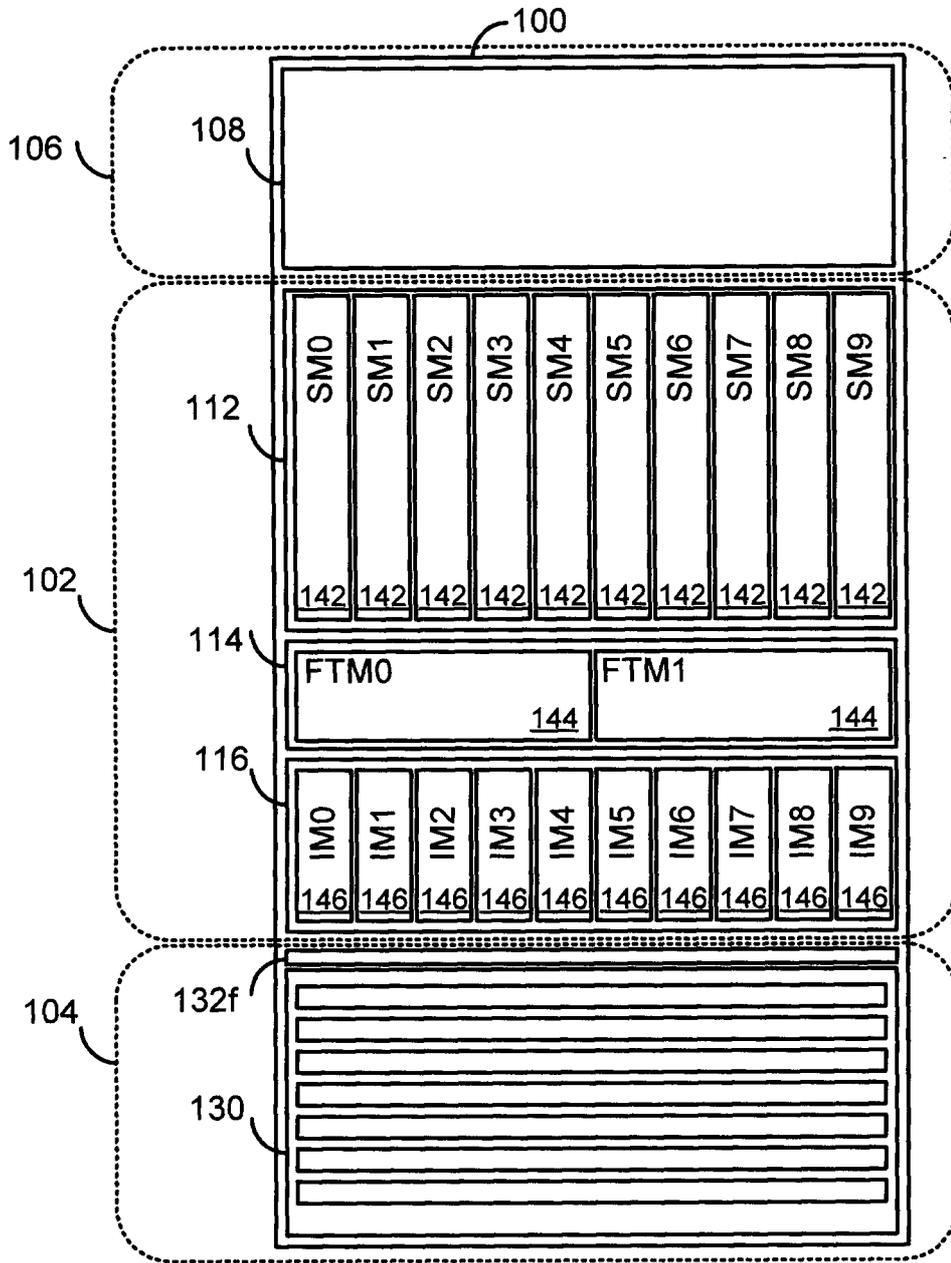


Fig. 16

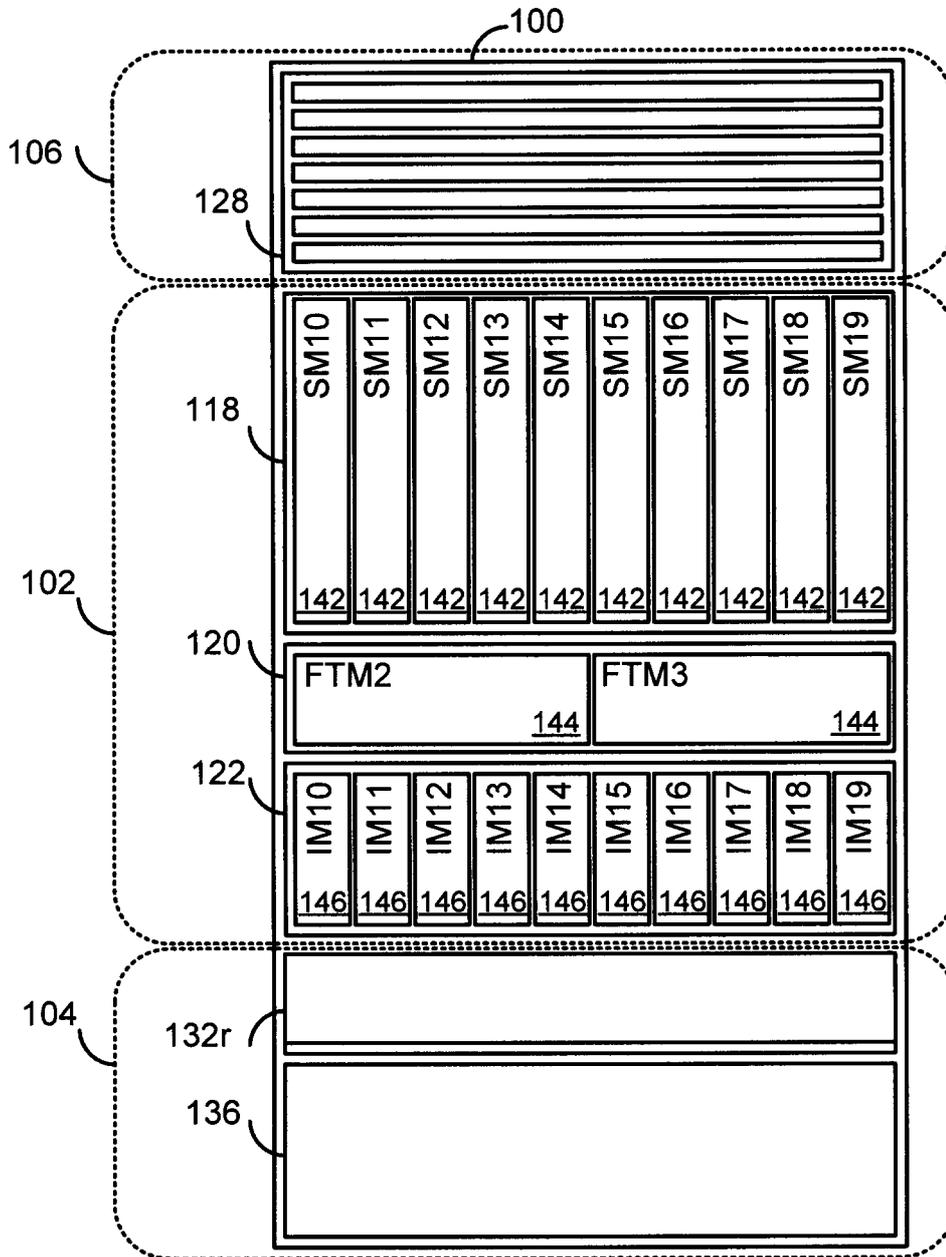


Fig. 1c

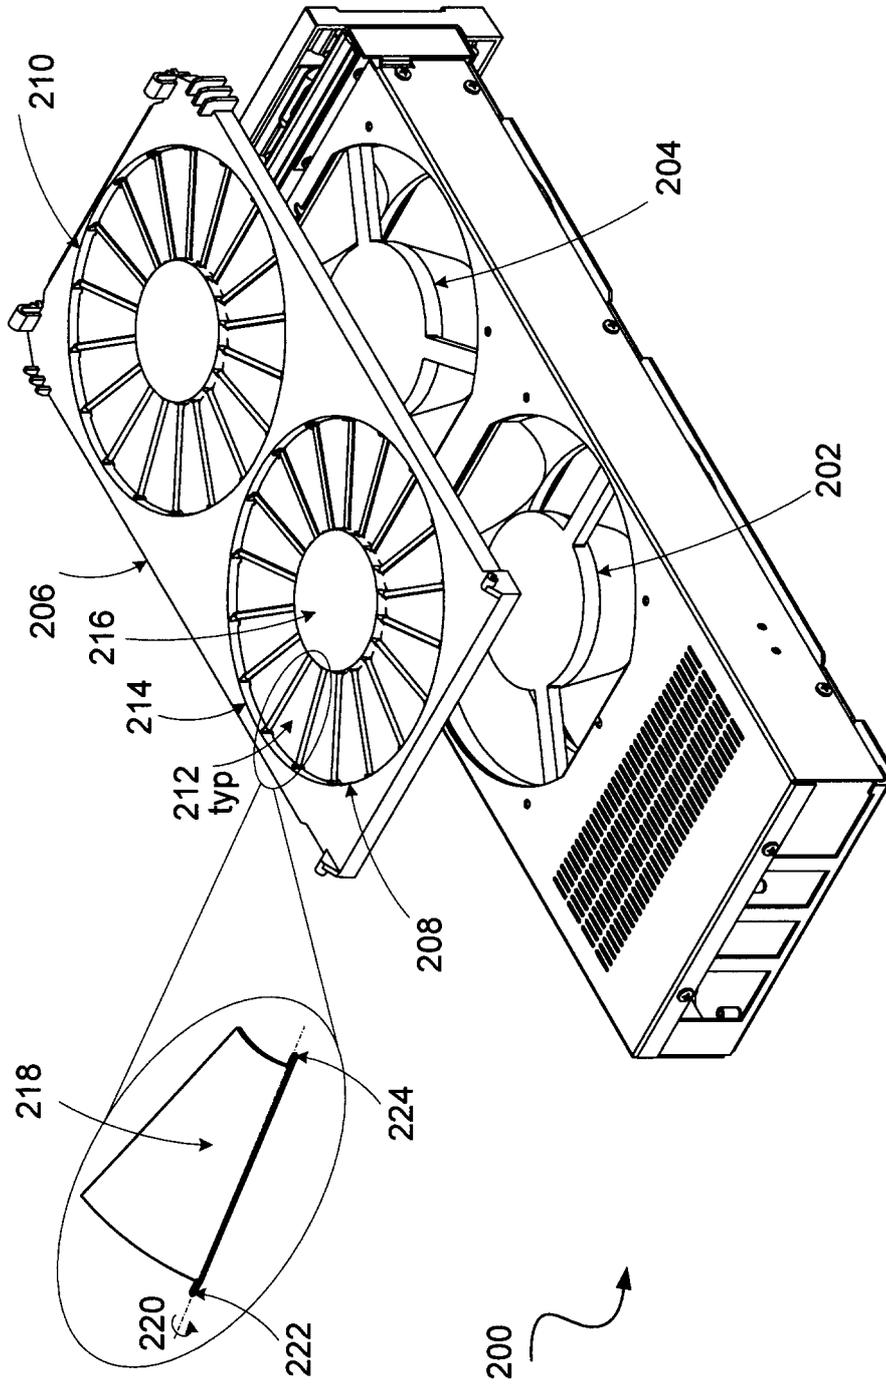


FIG. 2a

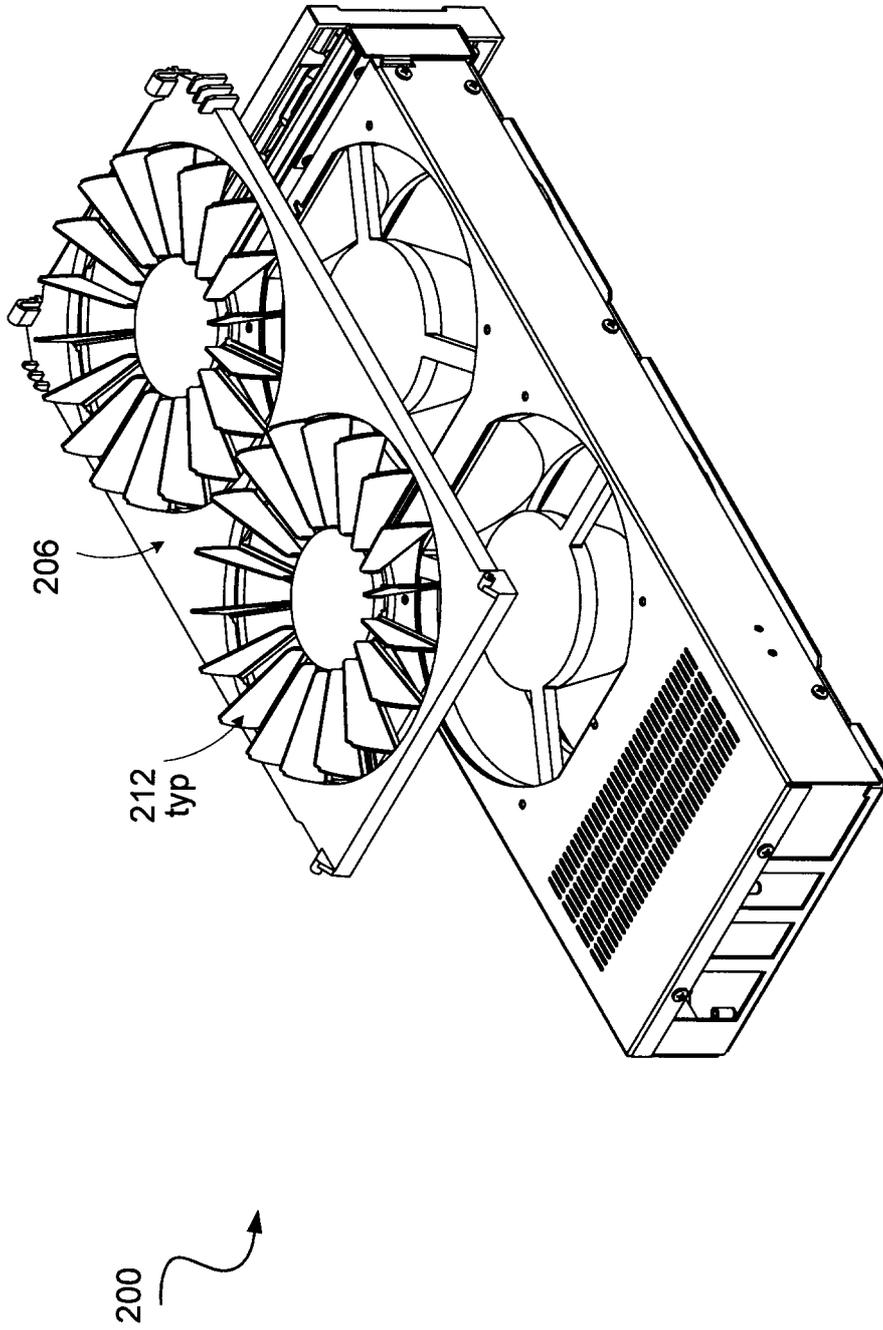


FIG. 26

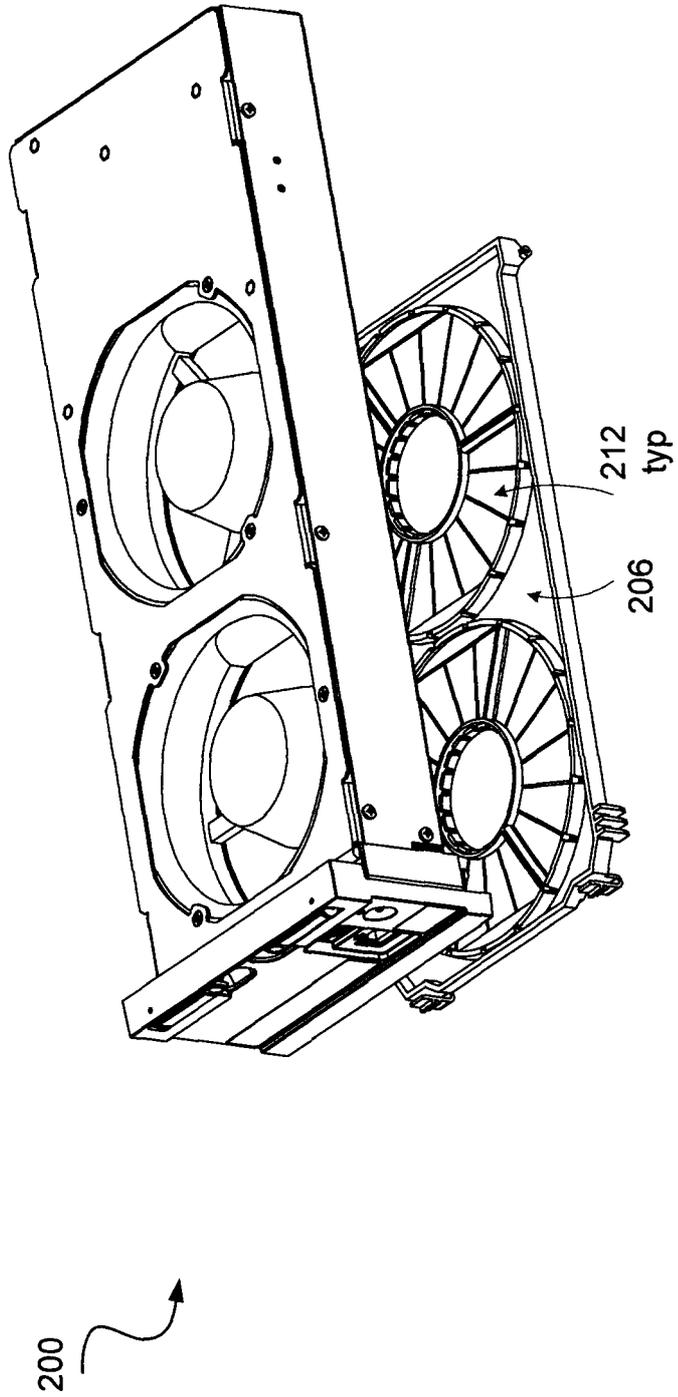
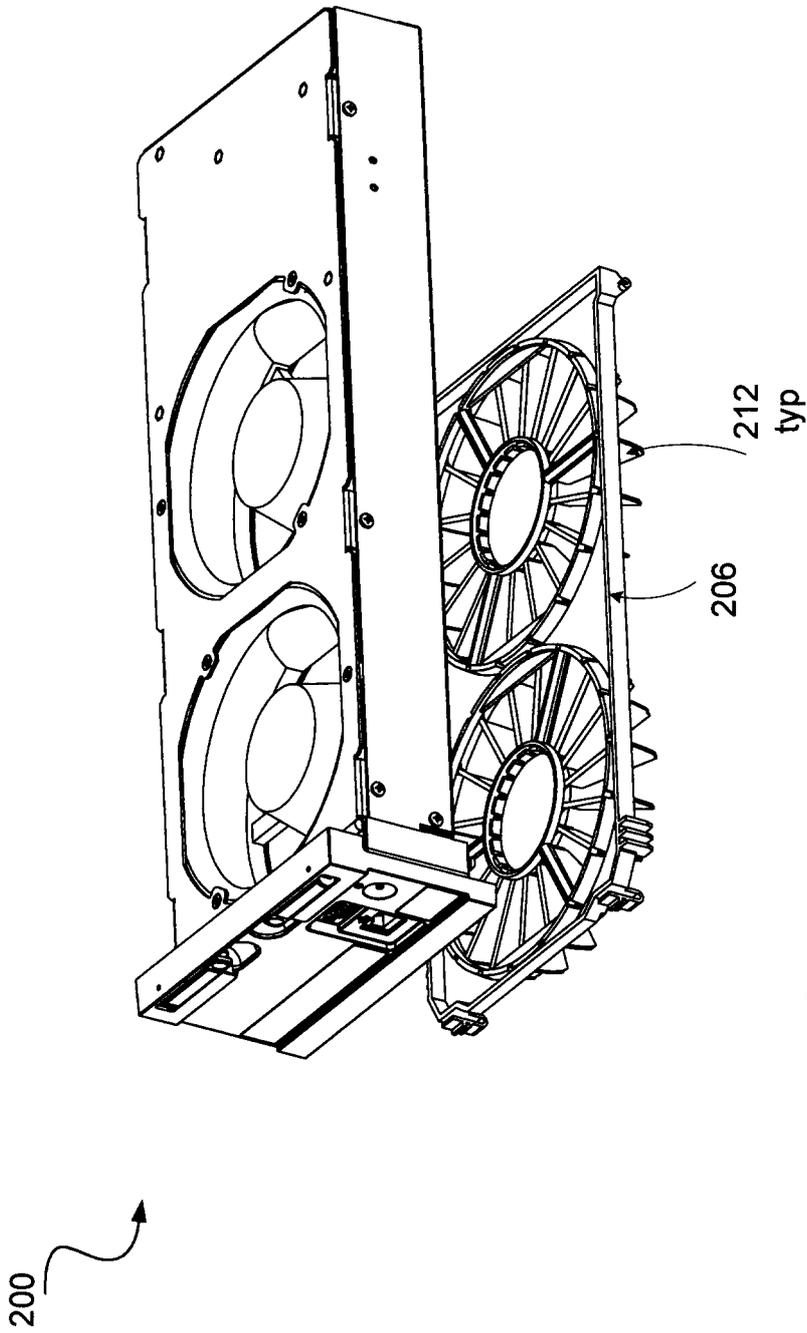


FIG. 2c



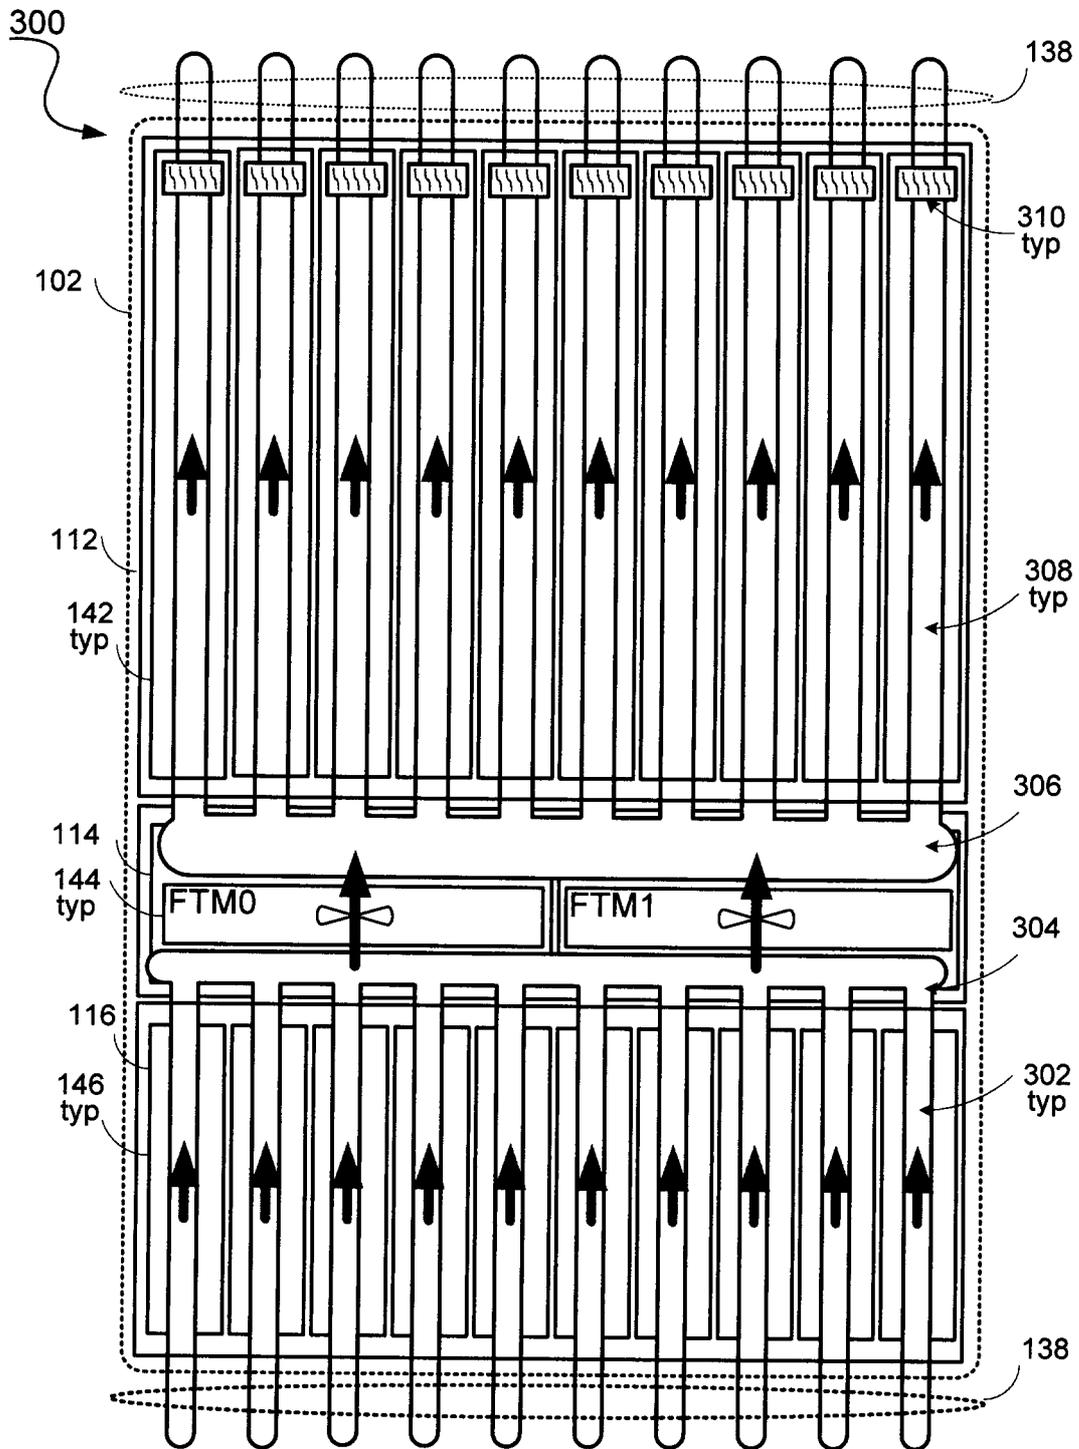


FIG. 3

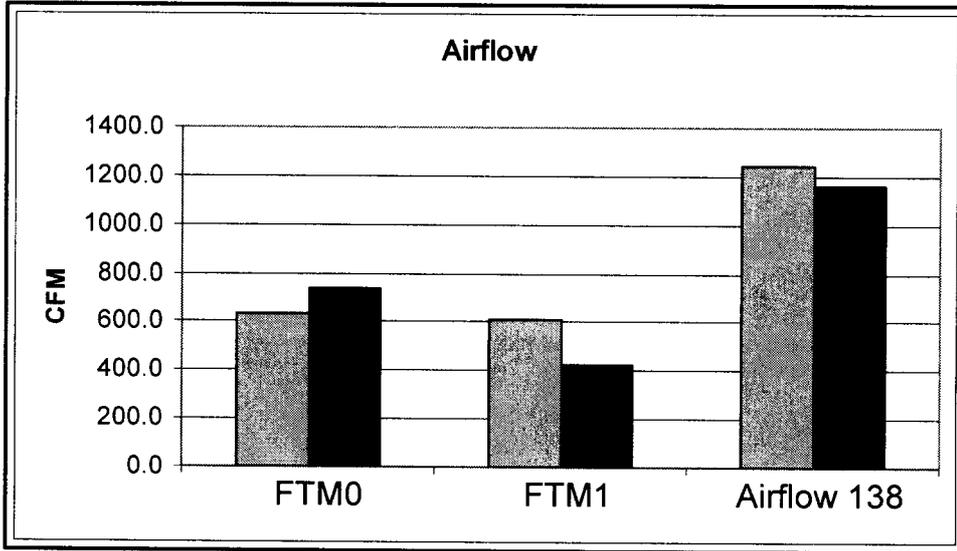


FIG. 4

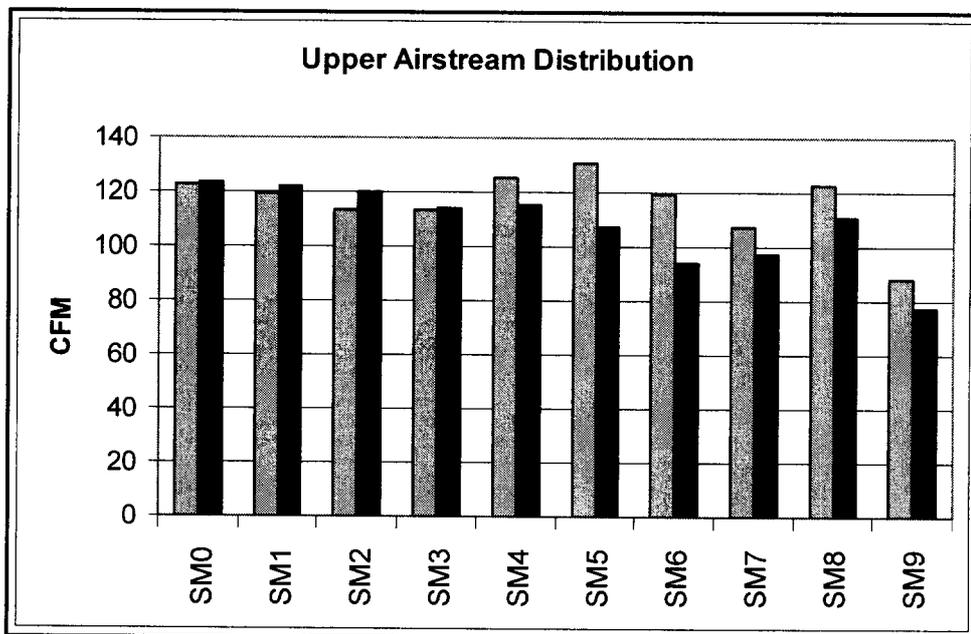


FIG. 5

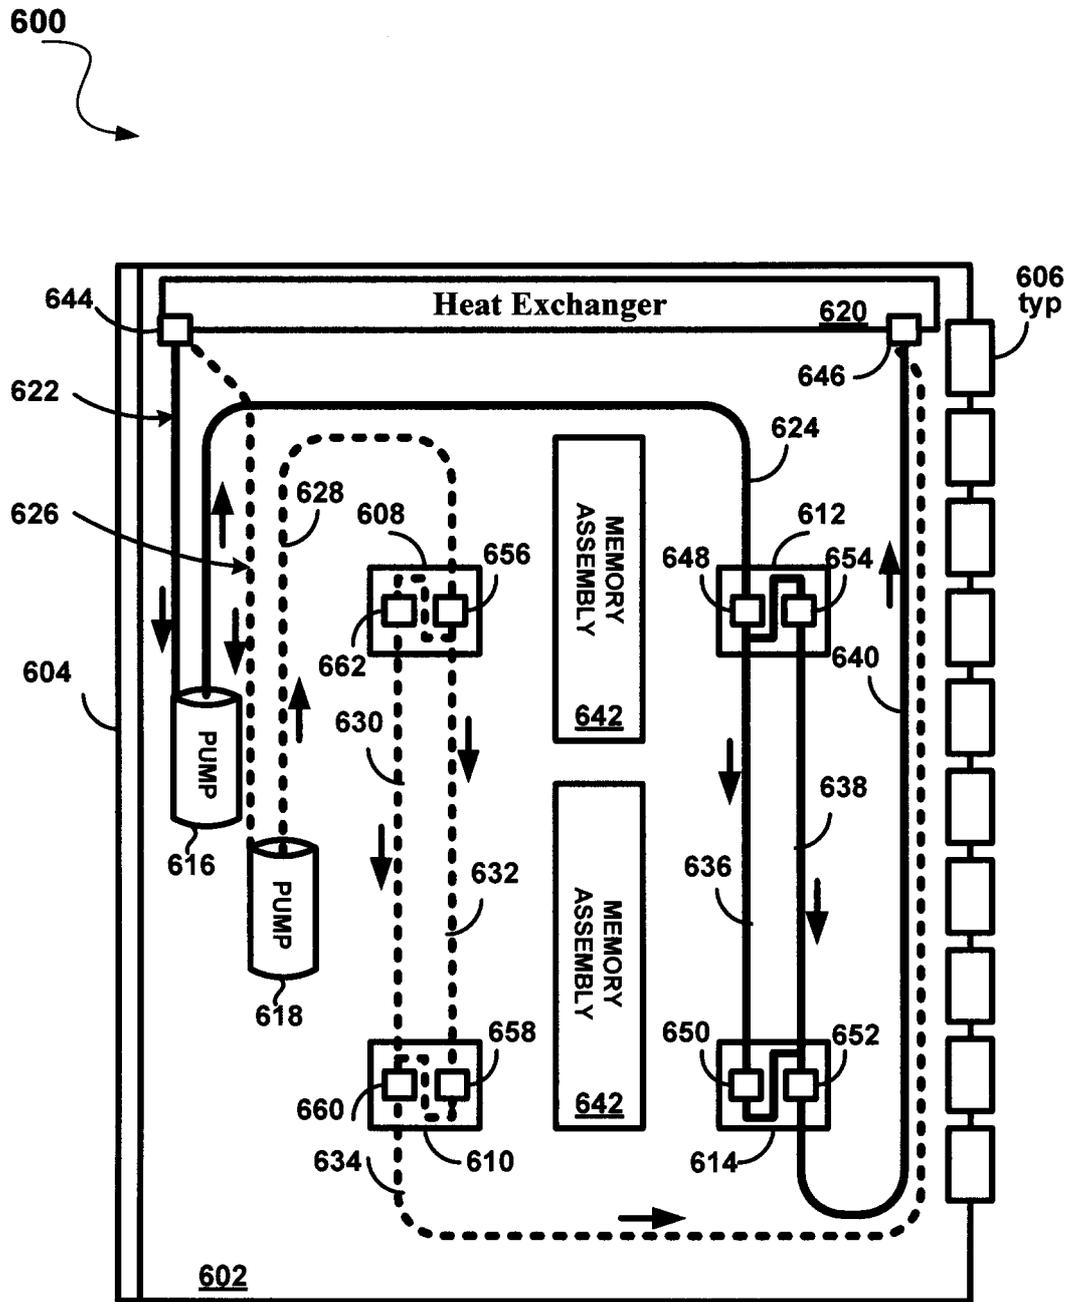


FIG. 6

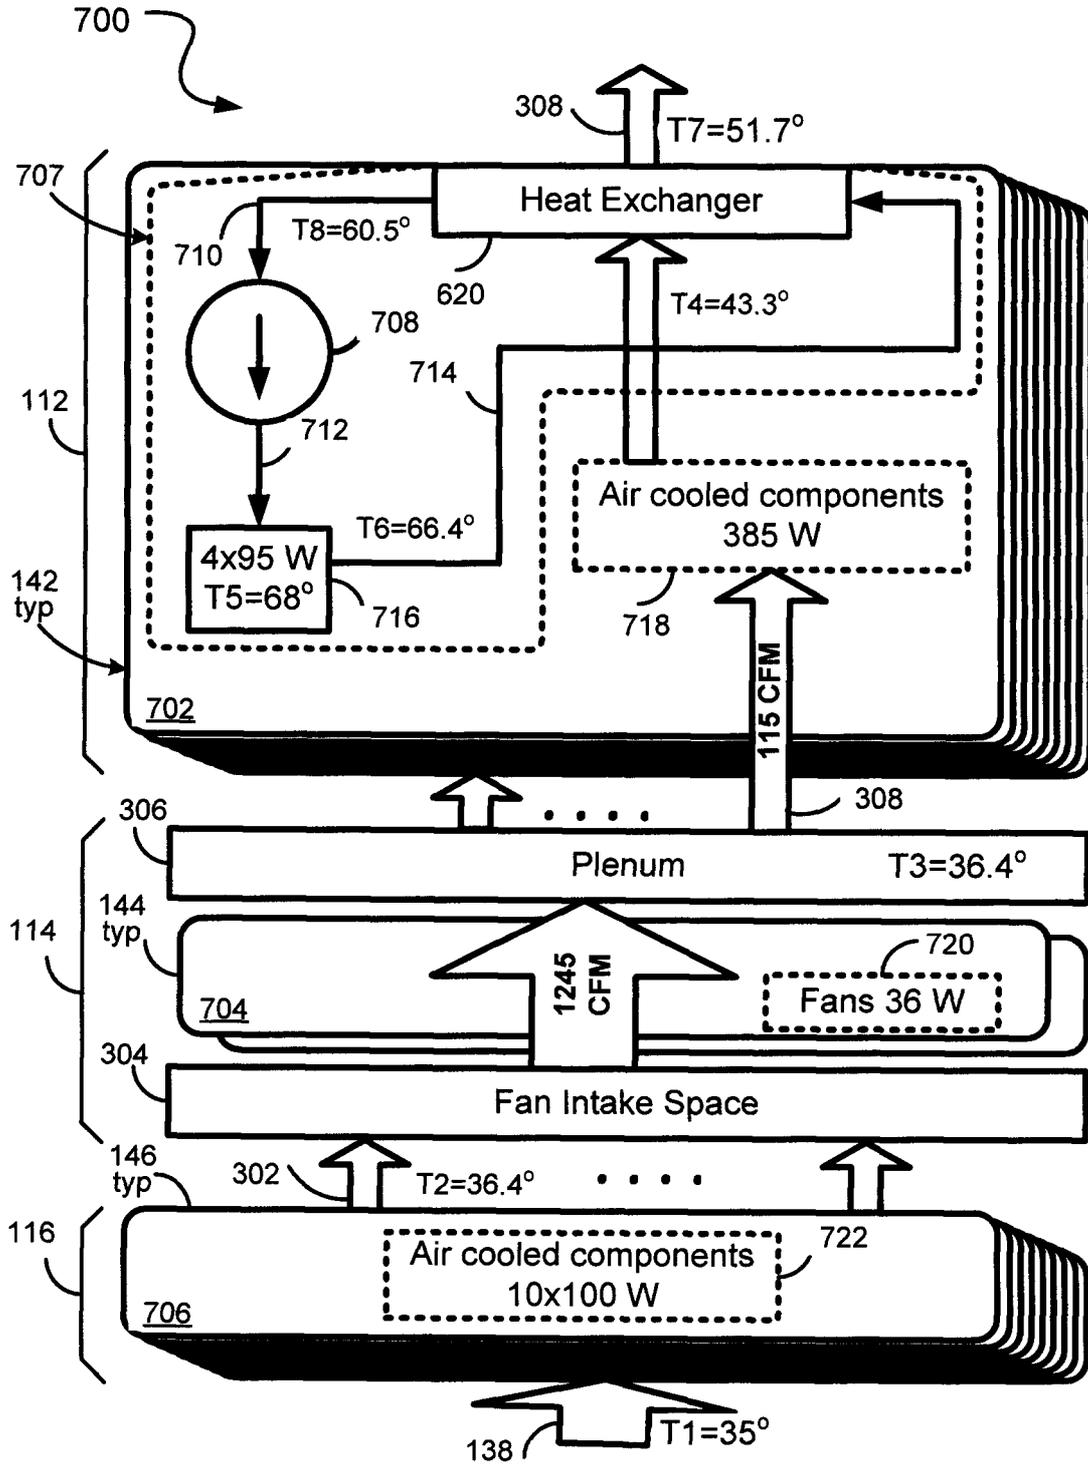


FIG. 7

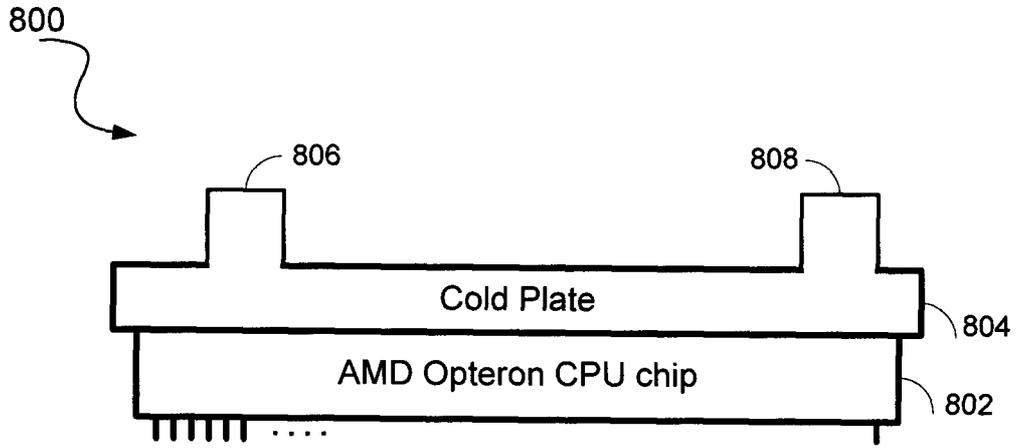


FIG. 8

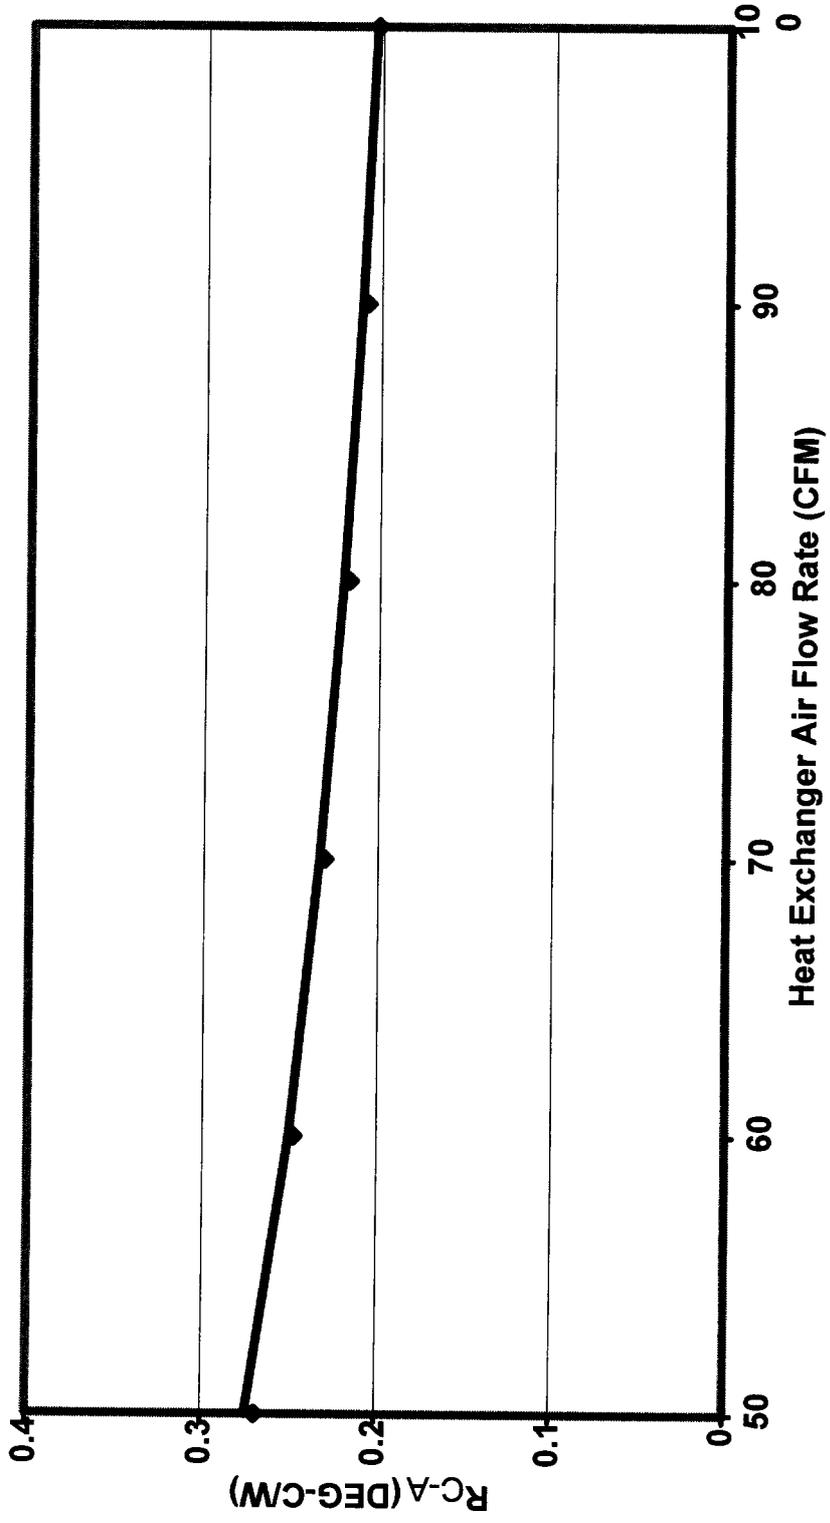


FIG. 9

A CLASSIFICATION OF SUBJECT MATTER IPC G06F 1/20 (2006 01) According to International Patent Classification (FPC) or to both national classification and IPC		
B FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) IPC G06F1/20 (2006 01) ,H05K (2006 01) Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic database(s) consulted during the international search (name of database(s) and, where practicable, search terms used) QPAT, Delphion, Esp@cenet, IEEEExplore, Google (cooling, computer, chassis, plenum, air, flow, streams, separate, independent, baffle, louvers, front, rear, upward, vertical)		
C DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
X	US 6,388,880 B1 (El-Ghobashy et al) 14 May 2002 (14-05-2002)	1-9, 20, 26-30, 34, 35
Y	* whole document * ---	10-17, 21-25, 31-33, 36-38, 41, 42
Y	US 2002/0126449 A1 (Casebolt) 12 September 2002 (12-09-2002) * whole document * ---	10-17, 23-25, 31-33, 41, 42
Y	US 6,445,581 B1 (Carbonneau et al) 3 September 2002 (03-09-2002) * whole document * ---	14-17, 21, 22, 24, 25, 36-38, 42
A	US 2004/0190243 A1 (DiStefano) 30 September 2004 (30-09-2004) * whole document *	20, 23, 25, 26, 31, 33, 41
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C <input checked="" type="checkbox"/> See patent family annex		
* Special categories of cited documents	T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention	
A document defining the general state of the art which is not considered to be of particular relevance	X document of particular relevance the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone	
E earlier application or patent but published on or after the international filing date	Y document of particular relevance the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents such combination being obvious to a person skilled in the art	
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	& document member of the same patent family	
O document referring to an oral disclosure use exhibition or other means		
P document published prior to the international filing date but later than the priority date claimed		
Date of the actual completion of the international search	Date of mailing of the international search report	
08 May 2008 (08-05-2008)	14 May 2008 (14-05-2008)	
Name and mailing address of the ISA/CA Canadian Intellectual Property Office Place du Portage I, C1 14 - 1st Floor, Box PCT 50 Victoria Street Gatineau, Quebec K1A 0C9 Facsimile No 001-819-953-2476	Authorized officer Jeffrey Orser 819- 934-2669	

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of the first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons

- 1 Claim Nos
because they relate to subject matter not required to be searched by this Authority, namely

- 2 Claim Nos
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically

- 3 Claim Nos
because they are dependant claims and are not drafted in accordance with the second and third sentences of Rule 6 4(a)

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows

Group A Claims 1-19 and 26-40 are directed to a system and method of cooling a computer system wherein a cooling system generates a front air flow and a rear air flow that are separate and independent of each other

Group B Claims 20-25, 41 and 42 are directed to a system and method of cooling a computer system wherein a plurality of axial fans generates an air flow into and pressurizes a plenum within the chassis such that the plenum divides the combined air flow into a plurality of substantially equal individual air streams

- 1 As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims
- 2 As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees
- 3 As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claim Nos
- 4 No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims, it is covered by claim Nos

- Remark on Protest** The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation
- No protest accompanied the payment of additional search fees

C (Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
P, A	US 2007/0178822 A1 (Lanus et al.) 2 August 2007 (02-08-2007) * abstract * * paragraphs [0028]-[0032] * * Figure 1 * ---	1, 20, 23, 26, 41
A	US 2005/0168945 A1 (Coglitore) 4 August 2005 (04-08-2005) * whole document * ---	1, 20, 23, 26, 41
A	US 6,272,007 B1 (Kitlas et al.) 7 August 2001 (07-08-2001) * abstract * * column 2, line 58 to column 3, line 27 * * column 5, lines 27-30 * * column 7, line 38 to column 8, line 12 * * column 17, lines 8-58 * ---	1, 26
A	US 5,706,668 A (Hilpert) 13 January 1998 (13-01-1998) * whole document * ---	18, 19, 39, 40
A	US 6,000,623 A (Blatti et al.) 14 December 1999 (14-12-1999) * whole document * ---	20, 23, 41

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No
PCT/CA2008/000180

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
US6388880	14-05-2002	NONE	
US2002126449	12-09-2002	US6525935 B2 WO02073383 A 1	25-02-2003 19-09-2002
US6445581	03-09-2002	NONE	
US20041 90243	30-09-2004	NONE	
US20071 78822	02-08-2007	NONE	
US20051 68945	04-08-2005	NONE	
US6272007	07-08-2001	AT248396 T AU5899500 A DE60004851 D 1 DE60004851 T2 EP1 194830 A 1 US6373697 B 1 US6542362 B2 WO0101229 A 1	15-09-2003 31-01-2001 02-10-2003 15-07-2004 10-04-2002 16-04-2002 01-04-2003 04-01-2001
US5706668	13-01-1998	DE4445818 A 1	14-06-1995
US6000623	14-12-1999	NONE	