

Hybrid cooling with cycle steering in the IBM eServer z990

G. F. Goth
D. J. Kearney
U. Meyer
D. W. Porter

The IBM eServer™ z990 introduces a new mode for cooling multichip processor modules that enables significantly more processors to be refrigerant-cooled than previously. In recent IBM zSeries® offerings, including G4, G5, G6, and z900, chip junctions in a single multichip module (MCM) located in a central electronic complex (CEC) frame were cooled for reliability and performance benefits, using refrigerant technology, to temperatures lower than those achievable with air cooling. In the z990 system, a hybrid cooling approach is used, allowing refrigeration to be extended to four MCMs in a single CEC, which makes possible denser systems and greater power efficiency compared with prior modular refrigeration technologies used. In the event of a malfunction of the primary refrigeration cooling system, a backup air-cooling system is automatically engaged until the refrigeration problem is fixed. System sensors monitor the cooling state at all times. When air cooling is required, the chip circuit temperatures increase and the logic clocks are optimally adjusted to match the new junction temperatures. When refrigeration cooling is restored, the clocks are adjusted back to their fast speed. This technique allows the z990 system to match the processor density of direct-air-cooled systems while retaining a system performance and reliability benefit from refrigeration.

Introduction

The z990 hardware architecture was introduced to answer the difficult challenge of retaining the low junction temperatures enabled by IBM modular refrigeration units (MRUs) while cooling four times as many multichip modules (MCMs) as prior systems [1] without increasing MRU space or power consumption. All prior MRUs had chilled only a single MCM, so cooling four MCMs required a reinvention of the technology.


When many cooling options had been considered, the one that best met density, efficiency, and performance goals was to integrate refrigeration technology as the primary means of cooling, with direct air cooling as the backup cooling solution in the event of a refrigeration failure. In this manner, the cooler silicon temperatures of refrigeration would support very high circuit reliability and

faster clock speeds without the space, power consumption, and cost of prior MRUs.

To successfully cool the MCMs, a robust, error-free method was used to match the logic clock cycle times to the temperatures of the circuits they govern. In industry, techniques such as cutting the clock speeds in half have resulted in large degradations of performance. For the z990 system, a much different approach was taken by IBM. The change in clock speed was tailored to maximize performance levels, whether the system was MRU-cooled or in a slightly degraded air-cooling mode based on circuit and wiring temperature-speed relationships. The air-cooled clock speeds are normally reduced just 4% unless in a hot environment, when an 8% cycle time degradation is employed. At all times the clocks involved in this adjustment maintain optimum performance levels.

©Copyright 2004 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/04/\$5.00 © 2004 IBM



(a)



(b)

Figure 1

(a) z990 and (b) z900 systems. (PU: processor unit.)

The hardware that enables this technology is quite different from that used in prior MRU systems [1]. The




Figure 2

AMA and z990 separate blowers/plenums.

base of the evaporator heat sink incorporates parallel refrigeration channels with copper fins soldered to the side opposite the MCM. A new insulation system was developed with an electromagnetic compatibility (EMC) protective coating and with improved quick connects to ease assembly. One of the major cooling differences in the z990 system compared with prior MRU generations is the method by which refrigerant expansion occurs and how the temperatures of multichip modules are regulated. Prior MRUs used a standard expansion device with a pressure bulb in the MRU, coupled with a hot-gas-bypass valve for control of the MCM temperature [1]. Such an approach is ideal for controlling the temperature of a single MCM, but it is ineffectual in cooling multiple MCMs with unequal heat loads. The z990 system architecture has one MRU cooling two MCMs at vastly different heat loads under certain conditions. The refrigerant flow to each evaporator is tightly regulated by its own electronic expansion valve controlled by MRU code. To meet packaging density needs, the final refrigerant expansion is done at an orifice integral to the evaporator.

In summary, the z990 system represents a radical improvement in cooling efficiency. Spatially this is evident in **Figure 1**, which compares a z990 system with a z900. Both utilize 29.5-inch-wide frames, with the z900 measuring two EIA¹ units (3.5 in.) taller than the z990. Power efficiency improvements allow the four MCMs in the z990 to be cooled with 60% of the compressor power used to cool the single MCM in the z900. A comparison of the

¹ EIA: Electronic Industry Association. One EIA unit = 1.75 in.

z990 and z900 attributes is shown in **Table 1**. This paper discusses the processor cooling hardware, cooling controls, and clock speed optimization technology IBM employs to accomplish this.

MCM cooling technology

An overview of the z990 central electronic complex (CEC) is given in [2]. The CEC cooling is achieved through a combination of air-cooling and refrigerant-cooling techniques. **Figure 2** shows the air-moving assembly (AMA) that houses the two primary blowers used to redundantly cool the memory cards, the riser card that converts CEC I/O data for communication to the I/O function within the cargo cages via the self-timed interface (STI) connectors, and distributed converter assembly (DCA) power supplies. Also shown in this figure are the backup blowers that cool the MCM heat sinks.

Normally the MCMs are cooled by refrigeration, with air cooling turned on only if the MRU fails; this enables a denser, more efficient CEC than if redundant refrigeration is used. **Figure 3(a)** shows a 3D model of a single processor node book with the evaporator/heat sink attached to the MCM. **Figure 3(b)** is a closeup photograph of the MCM area with the heat sink and module hat partially cut away, revealing the processor silicon beneath. Up to four of these nodes may be populated in one CEC.

Evaporator/heat sink design details are shown in **Figure 4**. The oxygen-free copper (CDA101) fins shown in green in **Figure 4(a)** are 0.7 mm thick. When the backup blowers are off, the gray cover ensures that the airflow cooling the memory and riser cards does not circulate through the MCM heat sink, as this would add unnecessary heat load for the MRU to remove. In the evaporator base beneath the fins is a thin, low pressure-drop evaporator shown in gray in **Figure 4(a)** and split to reveal the refrigerant channels in **Figure 4(b)**. The refrigerant (R134A), commonly used in automobile air conditioners, flows from the MRU to the evaporator in the upper right yellow tube and into the evaporator near its center. As the refrigerant enters the evaporator, it expands through a fixed orifice to its lowest temperature

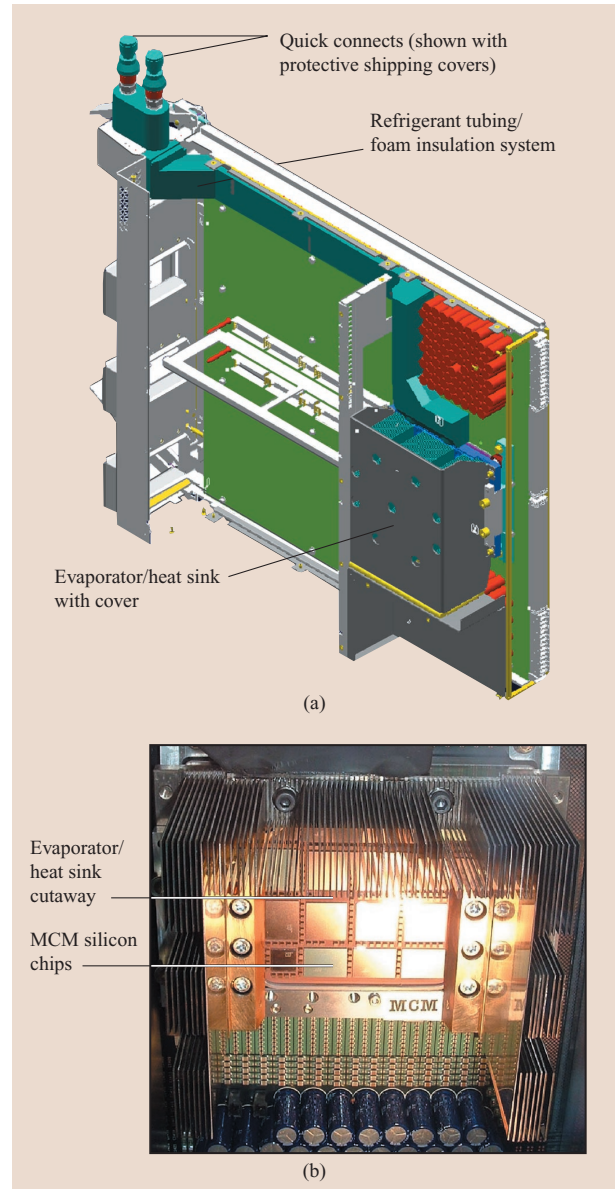


Figure 3

(a) Single node (PU book). (b) Closeup of evaporator/heat sink cutaway.

Table 1 Comparison of MCM cooling solutions in z990 and z900 systems.

Attribute	z990 – 16-way MCM cooling	z900 – 20-way MCM cooling
Compressor power/MCM power	1 : 3	3 : 1
No. of MRUs per MCM	1/2	2
MRU weight (lb)	60	120
Total no. of EIAs for MRU	7 for 4 MCMs	8 for 1 MCM
Evaporator/heat sink volume (in. ³)	100	800+
Special hardware for MCM condensation protection	None	Heaters, O-rings and C-rings, pressurized dry air, seal test equipment, humidity sensors




Figure 4

(a) Evaporator arm, evaporator, and heat sink assembly. (b) Evaporator details. (c) MRU hoses connected to evaporator arm.

and pressure state in the system. This keeps the inlet refrigerant tube much warmer than if all refrigerant expansion occurred in the MRU, enabling thinner tube insulation. Internal spacers maintain a minimum consistent insulation thickness of rigid polyurethane foam with a density between 35 and 40 lb per cubic foot around both tubes. The natural skin of this insulation was tested and shown to be an effective water barrier. Additionally, the flow impedance generated by the fixed orifice aids thermal cross-regulation control between the two refrigerant loops.

The quick-connects are shown at the ends of the yellow refrigerant tubes in Figure 4(a) and inside the opened insulation "clamshells" in Figure 4(c). To disconnect a processor node, the clamshell insulation and the quick-connects are opened. In contrast to prior MRUs, the z990 uses true quick-connect couplings, eliminating torquing of these connections in the factory and field. The hinged clamshell insulation system seals individually around the supply and return lines with an interference fit between the flexible hose insulation and the molded evaporator arm insulation.

Two backup blowers utilize 146-mm forward-curved scrolls and operate at 2800 rpm when cooling the four MCM heat sinks. About 260 cubic feet per minute of airflow per blower is generated at a pressure drop of slightly more than one inch of water. For optimum thermal performance under these airflow conditions, the copper fins are on a 1.95-mm pitch. The gray cover [Figure 3(a)] seals to a plenum beneath the heat sink, which causes all of the backup blower airflow to travel through the fins. The serpentine end-milled refrigerant path within the evaporator base of the heat sink is only 6 mm tall to minimize the thermal resistance this base creates for the heat sink function while providing adequate evaporation surface area. At sea level the air-cooled heat sink thermal resistance is 0.036°C/W measured to the MCM hat thermal surface under the airflow described above. The interface between the evaporator and the MCM hat uses a synthetic oil (PAO100, with a conductivity of 0.18°C/W-m) to reduce its thermal resistance when joined by the nine captive screws shown in white in Figure 4(b). Under air-cooling conditions, the junction temperatures may rise to 90°C.

The MRU is shown in **Figure 5**. It reuses the blower and condenser employed in the z900 system. The compressor shown under the hoses is similar to the compressor used to cool the 12-way z900 MCM but is more power-efficient. New to this system are the electronic expansion valves (EEVs) shown with blue insulation located adjacent to the compressor. The temperature of each MCM is regulated by its own valve. Component placement was optimized

for one- and two-processor-unit (PU) book operation, such that a disconnected line would not collect compressor oil during one PU book operation. Not shown are the blower located above the condenser and the power drive card (MDA-RG) located beneath the condenser. Both the blower and the drive card are field-replaceable units (FRUs) and are accessible from the rear of the MRU.

It is instructive to follow a pressure vs. enthalpy diagram of typical refrigerant cycles when the MRU is functioning, as shown in **Figure 6** for a high heat load (above 800 W) and a low heat load (below 400 W). Point A (Figure 6) is the MRU tube between the condenser and the EEV, with point B representing the evaporator inlet tube. As the refrigerant expands through the orifice in the evaporator block inlet [shown in Figure 4(b)], its temperature and pressure drop to their lowest values, point C. After this expansion the cold liquid refrigerant flows in three parallel serpentine paths designed with a thermal bias toward having the superheated refrigerant gas warm both the perimeter of the evaporator and the attached MCM hat before it exits the evaporator block, as represented by point D in the figure. This prevents condensation on the evaporator and MCM perimeters. The superheated refrigerant returns to the MRU through the lower tube of Figure 4(a) and enters the compressor, leaving at point E. Passage through the air-cooled condenser returns the refrigerant to its state at point A. Note that it is in the lower-heat-load case that the evaporator temperatures are lowest and require the most attention from a condensation perspective.

Figures 7(a) and **7(b)** respectively show the very different refrigerant schematics of the z990 and the z900 systems. Shown are the electronic expansion valves (EEV1 and EEV2) for regulation of the z990 MCM temperature. All prior MRUs used the standard thermal expansion valve (TEV) and hot-gas bypass valve (HGBV for the z900) with the code regulating the hot-gas flow for temperature control. This methodology, while successful for a single MCM, proved inadequate to regulate two MCMs at different heat loads. The hot gas introduced into the lower-powered, cooler side quickly increases the higher-powered MCM temperature as well, which is unacceptable. In contrast, the EEVs allow each MCM to be precisely controlled independently of its heat load, with cross-regulation effects readily managed. In addition to controlling the temperature of two MCMs, the EEVs also enable complete removal of the refrigerant in the lines between the expansion valve and the compressor prior to opening the quick-disconnects. When the EEVs are fully closed while the compressor is running, a partial vacuum removes




Figure 5

Isometric view of z990 MRU without covers.




Figure 6

Pressure (P) vs. enthalpy in z990 refrigeration cycles.

all refrigerant. Also absent in the z990 system are the humidity sensor, the heater, the vacuum test port, and other complexities of the z900 system.

For reasons of efficiency and package density, we selected the MCM hat temperature setpoint at 25°C, well above the room dew point. This eliminates condensation and any need for the insulation around the MCM that was required in prior systems. After accounting for the conduction resistances and variations in hat temperature due to the changing state and temperature of the refrigerant flow, our circuits are cooled to the nominal

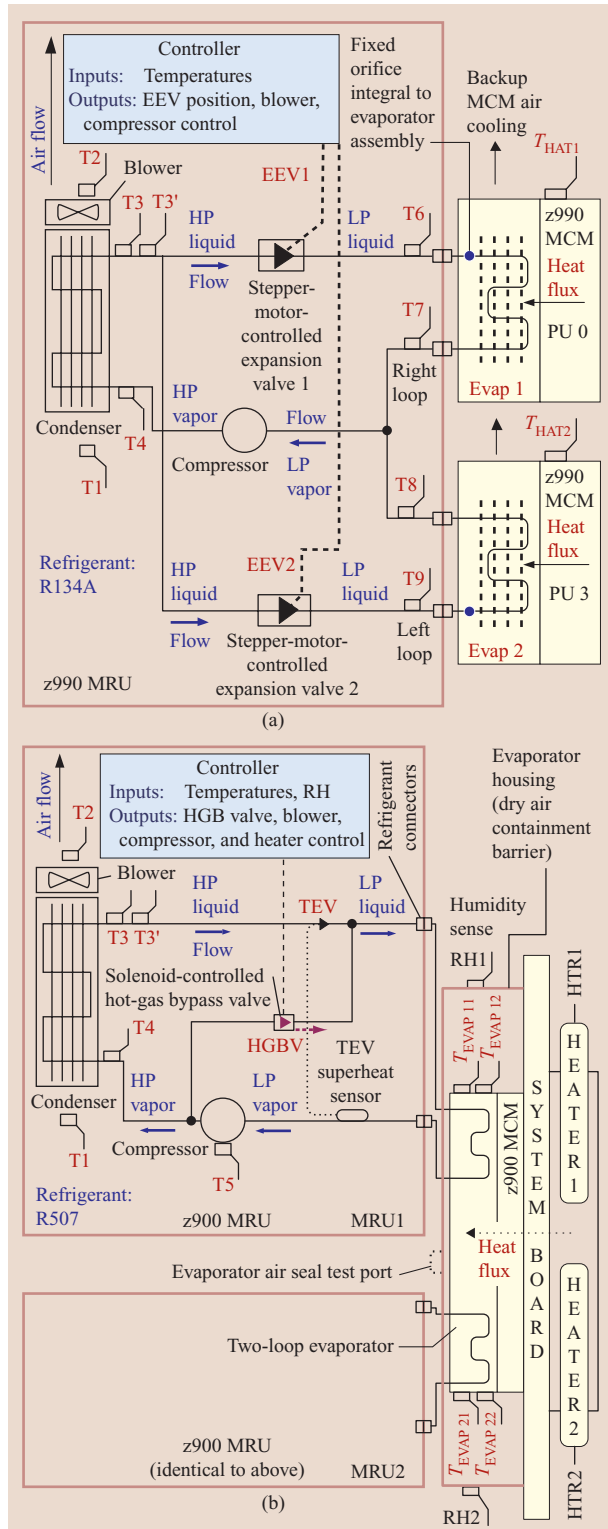


Figure 7

IBM zSeries refrigeration schematics: (a) z990—Two-loop MRU cooling two MCMs with air-cooling backup. (b) z900—Two MRUs cooling one MCM.

junction temperatures shown in **Figure 8** for an MCM fabricated by IBM. z990 MCMs manufactured by Hitachi employ a different thermal path from the silicon to the hat thermal surface, and the processors operate about 1°C cooler than temperatures shown in Figure 8, with the other chips at roughly the same temperature.

Figure 9 is a flowchart of the interaction of the cooling hardware, MRU code, power control code, and cycle steering application (CSA) code. The MRU code continually monitors three thermistors located in the center of each MCM hat. One is monitored directly by the power drive card, with the other sensors read by the DCA power supplies and sent to the MRU via the power control code. The power supplies also use these temperatures for thermal protection. In the process of monitoring and controlling the MCM hat temperatures, the MRU code detects whether an MRU is functioning improperly. The power control code automatically starts the backup blowers if a hat temperature threshold is crossed, ensuring cooling even if the power drive card (MDA-RG) fails. After comparing the actual MCM thermistor reading to the targeted setpoint, the MRU code opens or closes the electronic expansion valve and regulates the exact refrigerant flow to each evaporator, as discussed later. In the case of a thermistor reading above the maximum normal range, after the thermistor value is verified by comparing it with two other thermistor values, the MRU code sets a flag which triggers the power control code to notify the CSA to alter the clock speed by a 4% increment. This complete cycle is done within two seconds, during which a rise of less than 1°C occurs.

Cooling system integration—Code objectives and functions

The MRU and power control code designs are used to monitor and control the hybrid cooling system of the IBM z990 server and communicate to the cycle steering application the thermal condition of each module. We describe the design goals and functions in each code section. All code functions are designed to enhance the availability of the server.

Code design objectives

The code design for the z990 hybrid cooling system has the following priorities.

First, the code design operating the intelligent FRUs must provide the proper base function. That is, the code must operate the MRUs and backup blowers properly to support the cooling requirements of the MCMs while operating in various power-system or logic states. These operating states include fault conditions and subsequent service states for both power system and logic FRUs. For instance, the MRU code detects the low-power

nonfunctioning operation of a single MCM and compensates for that MCM power state to prevent condensation while maintaining other MCMs at normal MRU temperatures for optimal performance.

Second, the code and hardware are designed to allow system-level “self-test” operation in a manufacturing test environment. The hybrid cooling system code is capable of driving the various hardware FRUs to their specified limits; i.e., the cooling system provides manufacturing test modes that cause logic or other power FRUs to operate at or beyond their “corner” values. For example, the MRU FRU code design provides modes that allow the MCM temperature setpoints to be biased while a system undergoes verification tests. This enables all z990 MCMs to be tested at worst-case thermal conditions prior to shipping to a customer. Other parameters such as clock speed and voltage bias may concurrently be varied while these thermal conditions exist. Furthermore, FRUs can be turned off and on to prove proper service reporting and recovery while the system operates seamlessly.

Third, the code design is able to create real error conditions. This enables true fault testing and verifies the high availability of the power thermal system design. A low-level, low-overhead, secure microcode design was created that allows software control to inject apparent hardware errors. For example, the microcode running in an MRU drive assembly has software capabilities to cause the electronic expansion valve to “get stuck” at any arbitrary valve position with all other downstream FRU and higher-level system code oblivious of the method of the injected error. By causing the valve to apparently become stuck in various positions, both over-temperature and under-temperature faults can be created. By injecting a particular fault condition, downstream FRU code, as well as cage controller and service code, can be tested for proper fault recovery and reporting operations.

Finally, the microcode design also supports testing at the subassembly vendors—the same code used in a z990 intelligent FRU is also used at the subassembly vendors in order to test the MRU. This provides early awareness of any defect mechanisms in the subassemblies of the hybrid cooling system code.

Code design overview

The power thermal control code resides in the base power cage controller, the DCA cage controller, and the MRU drive card. This power-control infrastructure is based on the z900 power system control network (PSCN) concept [1] modified from a one- to a four-PU-book, “hybrid-cooled” z990 configuration. Each processor unit (PU) book is supported by a pair of DCA power supplies, with each DCA containing a cage controller operating as master or slave on the power system control network. A cage controller is also present in each base power

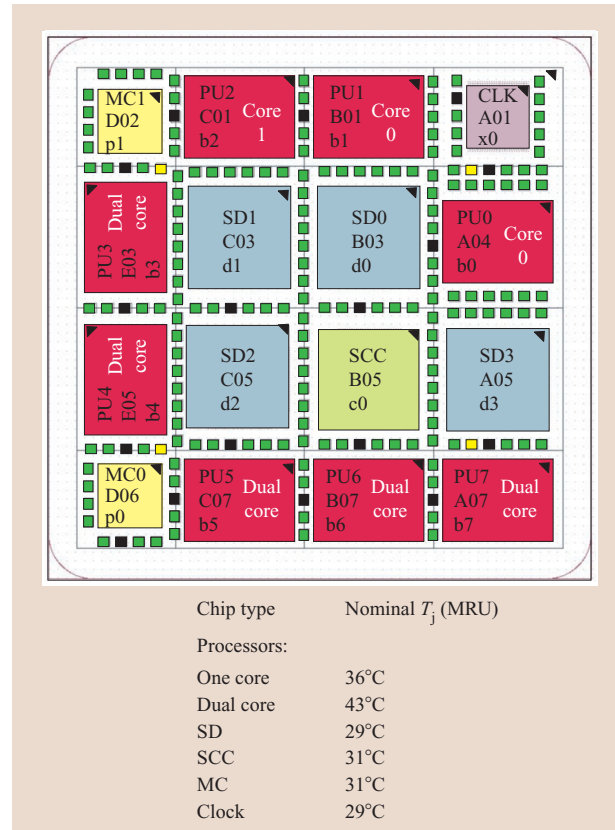


Figure 8

MCM chip locations and temperatures.

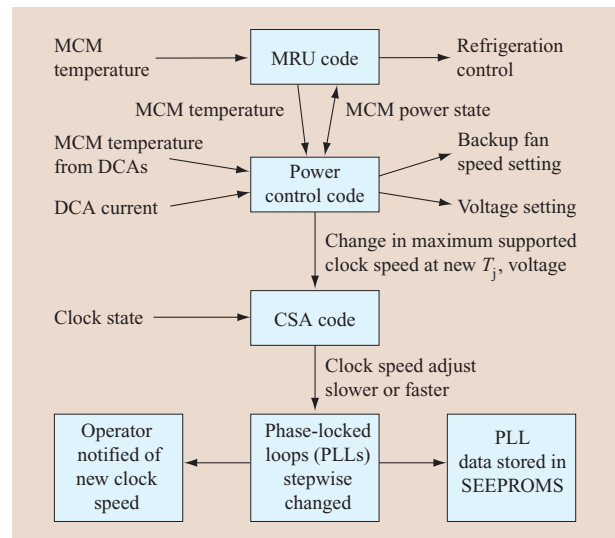


Figure 9

Flowchart of code-layer interactions.

assembly (BPA), operating likewise in master/slave mode. Power thermal system control and state data provided to and from the DCAs and BPA are via a 100-MB Ethernet communications data bus. The base power cage controllers interface with the MRU drive card via RS485 connections. The system elements (SEs) connect to the Ethernet bus and provide higher-level instructions to, or receive state data from, the power thermal system, as well as performing other logic-state monitoring and control. An MCM temperature sense thermistor assembly, comprising three thermistors in a single probe, is connected, one thermistor each to the appropriate MRU loop and each DCA providing power to that PU book.

Cage controller support functions specific to hybrid cooling

Within the cage controllers reside some specific functions needed to support hybrid cooling, including the following:

1. *Configuration*

The cage controller determines what the particular system hardware configuration is by analyzing which intelligent FRUs are installed and how they are cabled. The cage controller then correlates the proper PU books, DCAs, and MRU loops accordingly. Errors in configuration, such as misplugged cables, are sensed and reported for repair.

2. *Interface for normal power thermal control*

The cage controller receives high-level commands such as “power-on” as a result of customer input from the SE, and in response issues commands, in proper sequence, to individual intelligent FRUs. A power-on (PON) of a one-PU-book system, for example, consists of turning on the proper DCAs and MRU loop for PU book 0. Other control commands provided by the cage controller functions include engineering commands for special manufacturing test conditions not used during normal operation. Another command, precooling, is addressed later.

3. *Power-state sampling and communications*

Another cage controller function is to transfer power-state information to and from the intelligent FRUs, particularly the MDA-RG/MRU, so that proper temperature control is maintained. The heat flux value as indicated by the voltage and current output of each DCA is provided to the MRU by the cage controllers. This is similar to the method used in the z900 server [1]. The power-state data is used by the MRU microcode in order to prevent condensation from forming should the heat load be too low for too long. Low heat flux occurs either when there is an MCM checkstop² or when the MCM is not IMLed and that

PU book is therefore not functioning. In this non-operating extended low-heat-flux state, the MRU refrigeration flow to that PU book is turned off to prevent condensation; backup blowers are turned on to low speed. Another improvement compared with the z900 system is that the DCA-sensed thermistors are located at the same MCM thermal location as the MRU-sensed thermistor, and the data is posted to the MRU and used for backup temperature control.

4. *Cooling status*

The cage controller independently monitors the MCM temperature and provides appropriate interrupts to the cycle steering application (CSA) code should the temperature exceed trigger limits. The interrupts cause the CSA to take necessary actions for proper logic clock control for a particular MCM temperature range, as discussed in the next section. This cage controller function also controls the backup blower state. Backup blower speed is set according to the particular MCM temperature and its power state. If the MCM power is low when a “degrade range 1” CSA trigger is reached, the blower speed is set to “low.” If an MCM power state is high, i.e., functional, and its temperature reaches the first degrade range, the blower is set to high speed by the cooling status function. Regardless of the power state, the blower speed is set to high speed if the second degrade range is reached. The CSA trigger limits are contained in MRU FRU code for flexibility.

5. *Redundancy checking for service procedures*

The cage controller provides functions in response to service procedure commands to ensure that the system can function properly during maintenance procedures. This is referred to as *redundancy checking*. For example, a command to query whether it is acceptable to power off an MRU while the system is operating causes the backup blowers to be turned on and checked for errors, including speed verification. This ensures that backup air cooling will work prior to powering off an MRU should the need arise.

6. *Cyclic monitoring and error detection and fault isolation (EDFI)*

The cage controller provides cyclical functions to detect error conditions throughout the power system. Intelligent FRUs post errors when possible, and these errors are detected by a polling EDFI function. Errors detected then trigger service procedures to replace the suspect FRUs. Furthermore, the cyclical cage controller functions periodically cause the backup blowers to turn on in order to verify full functionality.

Functions of MRU microcode on the MDA-RG drive assembly

Here we turn our attention from the power control functions provided by the cage controllers to those

² The term *checkstop* indicates the halting of all processing.

functions executed in the MRUs (MDA-RG drive assemblies). The MRU microcode provides the following functions:

1. *Autonomous control and thermistor rationalization*

The MCU FRU code provides fault tolerance during cage controller failovers by operating autonomously based on data within the MDA-RG. An MRU loop will regulate using its direct sense thermistor channel before relying on data from cage controllers. The thermistor rationalizing function defaults to the direct sense thermistor unless it is detected as faulty, either “insane” or containing a “miscompare” error. A thermistor is classified as insane when it is detected to be open or shorted by simply comparing the thermistor reading to fixed high or low thresholds. A thermistor contains a miscompare error when its value exceeds the other two thermistor values by some smaller error threshold, a function continuously monitored by the MDA-RG code.

2. *Proportional integral derivative (PID) [3] MCM temperature control for each of the two MRU loops*

Each loop corresponds to a respective PU book; MRU 1 left loop/right loop supports books 3 and 0, respectively, and similarly MRU 2 supports books 1 and 2. The independent loop FRU code

- Controls stepper valve position for refrigerant flow of each loop and is calibrated by “homing closed” at power on and off.
- Provides independent per-loop cooling regulation. A PID controller is implemented in the MRU FRU code which is optimized for simultaneous maximum heat flux conditions on each loop and sufficient cross-regulation. The PID coefficients were chosen such that perturbations in one loop do not affect the other loop, but individual loop response is capable of compensating. Therefore, each cooling loop is essentially independent of the operating state of the other loop.
- Utilizes power and thermal data posted periodically (every 2.5 s) from the DCAs through the master BPA cage controller. This data path is used for backup temperature regulation and for monitoring the logic heat flux operating state.
- Executes various commands received from power control code, such as loop on/off, precool, and temperature bias during manufacturing tests.
- Provides function to pump out the refrigerant between compressor and EEV prior to disconnecting refrigerant lines.

3. *Other MRU code functions*

- Compressor speed has its own PID control when either stepper valve has reached its maximum open value.

- Blower speed is determined by the number of loops operating, the power state of each loop, and the rationalized condenser air inlet temperature, as well as acoustic considerations.
- Thermal sensors are rationalized to be fault-tolerant.

Cycle steering application goal: A fail-safe design

This “fail-safe” design principle is used for the design and implementation of the z990 cycle steering. It is a physical characteristic that the electrical resistance increases with temperature. The higher resistance leads to longer electrical delays for signal propagation in the wiring. Also, warmer circuit temperatures slow their switching rates. The cycle time determines when all signals must have arrived at their next gates so that they can be processed in the next cycle. The slowdown is about 1.6% per 10°C. Should the logic circuits experience a defect in the cooling system such that the cooling capacity in connection with the ambient conditions leads to increasing chip operating temperatures, this condition is detected, leading to an incremental speed reduction of some percent. This speed reduction compensates for the increased on-chip signal propagating delays. The MCM temperatures are monitored, and when boundaries are crossed, the cycle steering application (CSA) is notified to make an appropriate adjustment.

To make the design fail-safe, this degradation takes place whenever the communication path has a problem and the actual temperature is not known by CSA. In the case of a communication problem, with or without a real cooling problem, our fail-safe design will detect the messaging problem and reduce the speed regardless of the actual temperatures at that moment.

The CSA code runs on the DCA cage controller (CC) of PU book 0. The interrupt-handling part of CSA, as well as other program-issued requests, retrieves the current cooling states of all books. In the case of retrieval problems or implausible results, CSA degrades the machine speed to a safe setting. Should the MCM temperature still rise beyond an operating limit of this safe slowed-down speed, the machine is powered down in order to prevent hardware damage. The relation is illustrated in **Figure 10**. Here the horizontal axis is MCM hat temperature and the vertical axis is processor frequency in gigahertz (GHz). The slope of the line shows the inherent speed vs. temperature dependency of the circuits and steps, indicating the temperature at which the clocks are adjusted.

Cycle steering application code functions

The system speed follows the temperature. When the temperature that caused a speed degradation is corrected, CSA increases speed to return to the appropriate speed, corresponding to the requirement of the actual

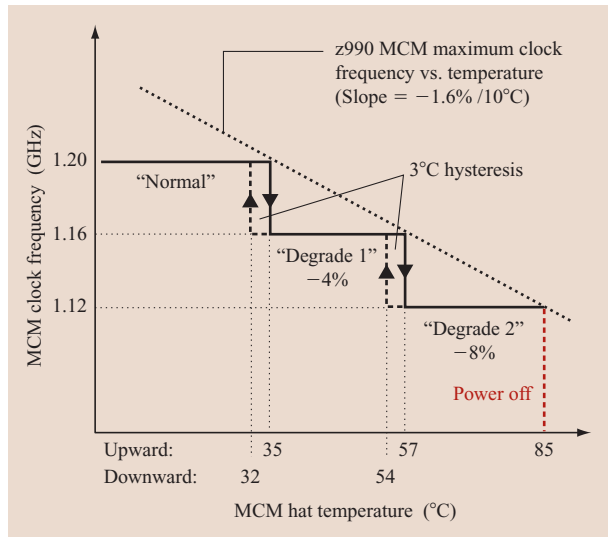


Figure 10

z990 cycle steering behavior.

temperature. Care and effort have been taken to successfully eliminate the risks of changing system speed while operating. As an example, fully loaded systems were tested for several days by downgrading, then upgrading the machine speed twice a second, error-free, to demonstrate reliability. Furthermore, at every completion of a system initialization or IML, the ability of the system to withstand a full downgrade and upgrade speed change is tested to ensure that it works with the physical factors and speed settings of the actual machine.

Setting the phase-locked loops (PLLs)

The clock sources are derived from Motorola MC12430 high-frequency PLL clock synthesizers. Their PLL is controlled by means of a variable value that is written into the chip, internally kept, and can be reloaded. If the reload is done in minimal increments or decrements, the phase is maintained and the pulse width is changed. The minimal step causes a width change that is smaller than the normal jitter on the PLL output. Thus, the variation cannot be detected by or affect the target clock receiving circuitry.

CSA makes use of the above effect and always moves the PLLs by a sequence of minimal steps until the target speed is reached. Every step is performed in a two-phase commit algorithm; i.e., the PLL values of the current and the next step are saved in a persistent storage concept. After the change has been written to the PLL and read back for verification, the saved current value is updated. This procedure provides the necessary precautions to complete a pending speed change even if this process is interrupted and a CC takeover is performed.

The PLLs sit on the two oscillator cards, one in charge of the current power-on, one as the backup card. The PLLs are initially loaded with a pattern which is hard-wired on the cards and loaded in parallel at power-on time. PLLs are usually loaded serially, but this loading exposes them to shift errors which can lead to wrong speed settings, since this may or may not be the correct pattern used later on. The pattern to be loaded for speed-adjustment purposes is generated from a set of digital I/O (DIO) lines, controlled by the FRU gate array (FGA) DIO engines. These lines can be read as input lines or driven high or low as output lines. One extra line serves as the gate for loading the pattern into the addressed PLL. The line settings are monitored to maintain stable levels. An interrupt is issued if any level should change and identifies a problem so that the failure can be repaired as soon as possible. The settings which make up the current patterns are kept, so that a glitch on their gates will not harm the PLL settings. Also, a failure to maintain the pattern will not cause a problem until the PLL must be changed, which in turn is not done without a previous check.

Initialization and hardware testing

These means are used to carefully check and initialize the speed setting of the system. After power-on and during the initialization phase, the initial PLL settings are done in the following sequence:

1. Read the default resistors of the oscillator card and compare with the card-level requirements. A mismatch points to a defective or wrong oscillator card. A repair is/can be initiated.
2. The pattern matching the actual system speed is loaded into the line drivers and read back. Any shortages or other failures corrupting this pattern are detected and a repair is initiated.
3. The loaded and verified pattern is gated into the PLL, and the pattern is again read back and verified.
4. The system clock is started, now using the PLL output as input.
5. At the completion of IML, the system is degraded 8% and upgraded again to nominal speed with the required number of minimal changes to the PLL. This ensures that all necessary patterns can be loaded into the PLL and the system does not execute an error. This process takes less than half a second.

Cycle-steering-caused speed degradations usually take place in scenarios in which either the cooling unit cannot maintain the chip temperatures below 35°C due to failures or extreme ambient conditions, or the messaging interface is temporarily suspended because of CC takeovers.

MRU and cage controller code updates

There are two exceptions to our general rules; both are related to concurrent code maintenance. One feature of the z990 system is the use of a digital signal processor (DSP) in the motor drive cards for blowers and the MRU. The DSP takes the 350 volts dc from the BPA and converts it to the appropriate voltage waveform for the compressor, blower, and EEVs. This DSP allows a change in this waveform by programming, a useful feature in the early development of the code and hardware. In the unlikely event that it is advantageous to change the DSP code in a machine that is already running, the machine goes into air-cooling mode for one to two minutes before returning to MRU cooling. Hence, this type of code download may create a brief speed degradation even though the cooling hardware is functional.

The second scenario is loading new cage controller code. Here, the exception is leaving the fail-safe design for a few seconds. When the PU book 0 cage controller or the BPA cage controller microcode is updated and rebooted, CSA or power control cooling status is not running for a few seconds during this code load. The probability of a cooling problem raising the MCM temperature into the next temperature range while the CSA or cooling status is not running is estimated to be so low that the risk of executing an error is negligible.

Running system in air-cooling mode

Besides the general CSA tasks, special handling has been implemented to support outstanding repair actions (RAs). Depending upon the severity of the cooling malfunction, the system may or may not require speed degradation by either 4% or 8%. If it is required, a special indication of a “defective MRU” is sent from the cooling system to CSA. To prevent a permanent oscillation between normal and degraded modes of operation, the normal speed is resumed only after a successful concurrent repair of the fault (described previously).

Furthermore, lack of cooling capacity during the IML is monitored and remembered at the next IML, which is then performed with the required speed degradation as a precaution in order not to expose the success of the IML to failures unless the MRU has been replaced.

One task of CSA is to monitor the speeds used for an IML and to ensure that these never exceed their IML values even though future temperature changes would permit an increase. The reason is that the initialization of elastic interfaces (EIs)³ that is done during an IML allows only speed reduction and its clearing, not faster

³ An elastic interface compensates, at startup time, for the latency of the logic circuits and packaging so that the controls and data lines are calibrated to the topology of the interface. Once calibrated, the interface protocols operate on the basis of the calibrated latencies [4].

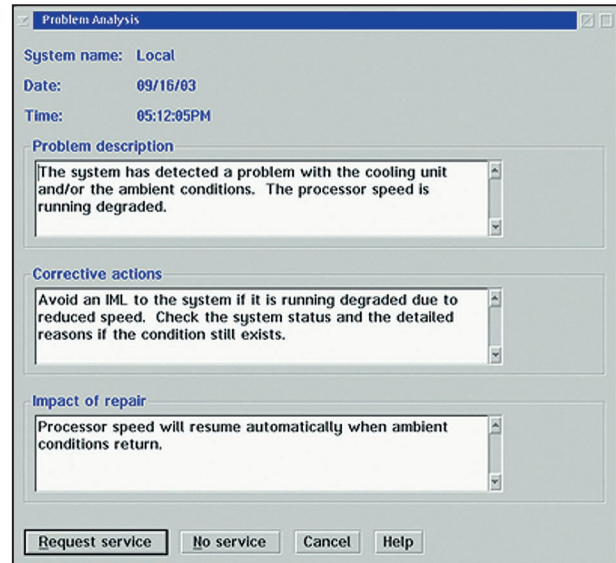


Figure 11

Problem analysis message.

speed than is present at initialization time. As a result, even with a marginal cooling unit the system continues to operate in a safe environment, possibly with speed reduced by either 4% or 8%.

Display of cycle steering state

There are two means of notifying the system operator; they are used to serve different purposes. Whenever CSA leaves normal speed, a CSA warning service reference code (SRC) is issued. This SRC is accompanied by a “system warning” message that uses the “hardware messages” interface box to alert the operator (flashes blue). Double-clicking the icon opens the hardware messages box, showing the system warning. Displaying the system warning then reveals the actual problem analysis message, as shown in **Figure 11**.

This warning is given to prevent a re-IML before a concurrent repair of the MRU can be performed; the re-IML would regain full clock speed automatically. As previously mentioned, performing an IML while the MCM hat temperature is degraded causes a recalibration of the elastic interface. The elastic interface must be calibrated at normal temperature to guarantee future increases in clock speed. Therefore, an IML at a degraded temperature requires an additional IML after the MCM temperature has been returned to the normal range. It is thus desirable to repair the cause of the degraded temperature before a new IML is performed.

Additional messages are displayed if an IML is performed while the system is running in a degraded temperature

condition. A message is displayed to indicate that another IML will be required after repair of the element causing the degraded state. A second message is presented when a repair has been completed successfully and the temperature returns to normal range. This message indicates that the processor speed remains degraded and an additional IML is desirable to return the processor speed to normal. All CSA-related SRCs do not result in a service call. Repair service reference calls are generated by the underlying defective hardware that induced the degraded temperatures and CSA warnings.

The system cycle steering state can be inspected at any time. This is important in providing information as to whether the system is running degraded, whether an IML should be planned or postponed, etc. The system CPC icon shows any degraded state, and the degraded details panel reveals the reasons, similar to those shown in Figure 11, with similar text.

Pre-cooling the MCMs

Finally, CSA routes commands from the SE to the cooling unit and vice versa. Starting the clocks, either for self-test or final IML, causes sudden heat dissipation. The cooling unit requires some time to react, because in idle mode the temperatures are not held extremely low in order to prevent condensation inside the machine. It may take longer than desired to return the temperature range to normal. Therefore, the cooling unit is notified about the clock start request. The cooling unit now starts maximum cooling, and when an acceptable temperature is reached, the clocks can be started. This is indicated in the response of the cooling unit to the pre-cool command.

Anatomy of an over-temperature event in a nominal condition

To illustrate how the cooling hardware and sensors, the MRU code, the power control code, and the CSA code functions work together in the event of a cooling failure, we describe what happens if an electronic expansion valve (EEV) were to become stuck in a low-flow position. This scenario and others were artificially induced using software inject capabilities.

1. The EEV controlling the refrigerant flow to PU book 1 becomes stuck at a position of 50 (a position in which the valve is almost closed) and provides insufficient refrigerant to cool typical power in the MCM.
2. The MCM thermistor sensed directly by the MRU code temperature value increases above its target, and the control code attempts to open the valve farther (but the valve is still stuck at 50). The MCM hat temperature continues to rise.

3. All MCM thermistor temperatures are compared, and all values indicate that a problem is not detected with direct sense value.
4. The MCM hat temperature now exceeds the over-temperature (OT) limit of 34°C. The MRU FRU code posts an error to its internal error registers.
5. Error detection fault isolation (EDFI) code detects the OT error flag in MRU error registers. Logs are captured of the state of the MRU and all of its sensors, and a “service reference code” is posted indicating a defective MRU due to an OT condition. The MCM hat temperature is still rising slowly.
6. The MCM temperature now exceeds the cooling status “degrade 1” threshold (35°C), while the power state remains high. An interrupt is posted to CSA.
7. The CSA code reads the “degrade 1” temperature range and a defective MRU.
8. The CSA code sets PLLs to 4% slow cycle time to all nodes.
9. The MCM hat temperature now exceeds the backup blower monitor function threshold (38°C)—blowers are turned on to 2800 rpm. The blowers turn on at 3°C above the “degrade 1” state; this prevents clock speed oscillations when cooling starts.
10. The MCM hat temperature stabilizes to air-cooling levels, typically within the upper half of “degrade 1” temperatures.
11. The MRU service actions are started. A redundancy check of backup blowers indicates that the MRU can be replaced concurrently.
12. MRU 2 is deactivated. The EEVs to both loops (PU book 1 with bad valve and PU book 2) are turned to a closed position, while the compressor is left on briefly to remove refrigerant before power is removed.
13. MRU 2 cables are disconnected, refrigerant lines are disconnected, and MRU 2 is removed from the frame.
14. A new MRU 2 is installed and activated, cooling loops are turned on, and error registers are cleared.
15. MRU cooling now returns to PU books 1 and 2. The MCM hat temperature drops because of refrigeration flow. When the hat temperature drops to the “blower turn-off” threshold (35°C = 38°C – 3°C hysteresis), backup blowers are turned off, and the hat temperature continues to drop.
16. Temperature is now below the cooling status “degrade 1” threshold (32°C = 35°C – 3°C hysteresis). An interrupt indicating normal temperature range is posted. The CSA checks the temperature range and the MRU defective state before returning cycle time to normal.
17. The MRU is regulating properly at a 25°C setpoint; no defects are present, and normal operations are resumed.

Hardware reliability and system verification

The fail-safe philosophy previously described for our control code is extended to the design and testing of the cooling hardware in numerous areas, including those discussed in the following sections.

Condensation control

When a refrigerant cooling system is employed in a server, it is essential to ensure that no condensation forms. In prior systems, IBM utilized an extensive seal system around the MCMs, complete with desiccants and tests of the seal and humidity sensors. In the z990 system, we have not employed complex seal systems, but instead have followed several discrete steps to ensure that no moisture forms:

1. Established 25°C hat target, which is above the ambient dew point.
2. Designed the evaporator such that the perimeter is warmed by the superheated refrigerant.
3. Tested MCM exterior-surface and board-surface temperatures in an environmental chamber at temperature, altitude, and humidity extremes.
4. Implemented a control code to detect power states that may cause condensation, or a non-IMLed PU book, and take mitigating action via backup air cooling.
5. Used double-redundant thermistors with miscompare and insanity checking.
6. Caused the final refrigerant expansion to occur inside the evaporator in order to localize the coldest regions.
7. Used highly effective refrigerant tube insulation and “clamshells” over the quick-connects.

Hardware qualification

In-depth testing was performed on critical components or components not previously used in IBM refrigerated products to ensure their reliability and life expectancy:

- *Hoses* – Successful application testing was performed on the four bronze flexible refrigerant hoses which make up the dual path supply and return lines shown in Figure 5. This coupling of technology and material has proven over time to be a reliable method of delivering refrigerant to evaporator assemblies, but the range of movement and handling aspects of the design must be demonstrated each time it is introduced into a product.
- *Couplings* – A major change was implemented on the z990 system with regard to the type and manufacturer of the refrigerant couplings used in this design. Prior IBM zSeries* product lines employ refrigerant cooling using a threaded, spring-loaded, double-sealing coupling design which is soldered or brazed onto the copper tubes/bronze hoses. These have proven, over time, to have higher-than-desired problems with reliability, both

in refrigerant escape and threading and installation damage. To improve this design point, a different-style coupling was selected. Specifically, a nonthreaded quick-connect coupling with mechanical attachment to the tubes and hoses was chosen and tested for leak rate, loss of refrigerant during plug/unplug, and performance attributes in the z990 MRU application. Test results were very positive, and hardware experience over the two years of product development activity has shown a very low rate of failure.

- *Electronic expansion valves (EEVs)* – New to our refrigerant design point, these stepper-motor-controlled valves replace traditional thermal expansion valves and provide the needed control between the dual sides of the MRU. Samples of these parts were tested for wear characteristics, since the MRU code constantly moves and repositions the valves throughout the life of the product. The minimum criteria for iterations of open/close commands based on life expectancy of the overall cooling system were met, and testing to failure continues.
- *Braze joints and general materials* – Consistent with prior refrigeration designs, the integrity and life of copper tube joints, connections to new refrigeration components, and insulating materials were tested and verified to be at acceptable levels.
- *MRU life test* – Samples of the overall z990 refrigeration unit were purchased from the assembly supplier, run through initial structural and stress tests, and then placed under long-term test in the application environment. Units from this sample “pool” have been and are being pulled at periodic intervals and put through destructive analysis to evaluate wear characteristics, condition of oils and refrigerant, condition of insulation materials, and fatigue of metal tubing or components. Data compiled to date on tested and analyzed units has been positive and has enabled release of the cooling system for production and product general availability.

EMI prevention

For the first time since the introduction of refrigerant-based cooling into IBM zSeries products, high-frequency electromagnetic interference (EMI) was found to be a problem with cooling hardware. It was determined that high-frequency electronics in the z990 MCM were using the evaporator coldplate and copper supply and return tubes as antennae for transmitting interference outside the PU book. The rigid insulation which fills the exit from the book does not shield the book from EMI radiation. A method was needed to provide a complete three-dimensional metal seal around the book opening and insulation, with intimate attachment to the book sheet metal and coupling bodies. Traditional EMI grounding




Figure 12

Arc-sprayed locations for prevention of EMI.

methods proved ineffective. A solution was found by coating the completed evaporator arm, at the coupling end, with a metal spray (arc spray). This is shown in **Figure 12**, with the gray area indicating where the metal spray is applied. The robust spray adheres around each coupling; when assembled into the PU book, it contacts an EMC gasket material attached to the book sheet metal to complete the grounding contact to the metal coating of the evaporator arm. This allows the removal and reinsertion of the evaporator assembly, if required, without affecting the integrity of the EMI shield. Additional UL testing was successfully completed to meet approval requirements for metal arc sprays.

Testing MRU unit at supplier and in IBM systems tests

After every MRU is assembled, leak-tested, and charged, it must pass a stringent and comprehensive 24-hour test at the supplier, where the capacity is verified at 850 W from both MCMs, as well as its ability to control light-load MCMs and the operation of its internal sensors. Because the EEVs are all slightly different, they are tested to ensure that they perform well in a given MRU. Each MRU is tested at specific valve settings and heat loads; the temperature of the MCM hat must be in a tight range to ensure successful control in all field environments and at all heat loads. Since this test is more stringent than commercial EEV tolerances guarantee, an occasional EEV is removed, even though it is within the tolerance range of the supplier. Before and after this 24-hour functional

test, as well as tests performed within IBM, the MRUs are precisely weighed to verify that there has been no loss of refrigerant charge, which is the principal cause of refrigerant system failures. Similarly, each evaporator is X-ray tested as well as helium-leak tested to ensure that there are no marginal braze joints that may induce subsequent fails.

When the MRUs are received by IBM from the manufacturer, they again must pass a series of system tests in an environmental chamber where their functionality is verified using actual z990 systems operating in extreme environments. Part of this testing is verifying error-free switchovers between refrigerant- and air-cooling modes.

Processor book functional testing at all cooling conditions

Each complete PU book is tested under and beyond all possible cooling conditions to which it may be subjected in the field using the evaporator/heat sink, MRU, and backup air system. In one of the early book tests called system run-in, the backup air system alone is utilized to drive the junction temperatures to 90°C, controlled by adjusting the backup blower rpm based on the hat thermistor readings. This test is intended to detect early life failures that may have escaped prior chip- and module-level stress tests. Each PU book is tested in hot and cold environments with MRU cooling, with air cooling alone, and in transitions between the two cooling modes. No book leaves IBM either in a system or as a field replacement or upgrade without having been completely tested under both cooling modes and in transition between the cooling modes. In this manner, we ensure that cycle steering works flawlessly on all books shipped.

Concluding remarks

Market-driven forces require the IBM eServer* z990 to introduce a four-PU-book system with each MCM cooled by refrigeration. The normal industry solution to heat removal is forced-air cooling, but the resulting warm junction temperatures would create unacceptable reliability and slow clock speeds in the zSeries. The full-clock-speed benefit of modular refrigeration unit cooling is determined by adding the circuit speed increase with lower temperature to the higher voltages enabled by these cooler junction temperatures. It was estimated that as an air-cooled system, z990 logic voltages would have to be lowered between 5% and 10% as a function of MRU-cooled voltage levels to mitigate the harmful reliability effects of warmer circuit temperatures. By both lowering the junction temperatures and allowing higher voltages, this equates to a performance gain between 10% and 14% when compared with a completely air-cooled z990 system. The few hours' time a hybrid system spends in air-cooled

backup does not materially add to the MCM failure rate.

While we reuse the z900 acronym MRU to describe this z990 approach, it differs markedly from the prior MRU applications. Coupling the cooling attributes of a z900 with a robust, fail-safe design approach, the z990 system offers state-of-the-art cooling.

*Trademark or registered trademark of International Business Machines Corporation.

References

1. P. Singh, S. J. Ahladas, W. D. Becker, F. E. Bosco, J. P. Corrado, G. F. Goth, S. Iruvanti, M. A. Nobile, B. D. Notohardjono, J. H. Quick, E. J. Seminario, K. M. Soohoo, and C. Wu, "A Power, Packaging, and Cooling Overview of the IBM eServer z900," *IBM J. Res. & Dev.* **46**, No. 6, 711–738 (November 2002).
2. T. M. Winkel, W. D. Becker, H. Harrer, H. Pross, D. Kaller, B. Garben, B. J. Chamberlin, and S. A. Kuppinger, "First- and Second-Level Packaging of the z990 Processor Cage," *IBM J. Res. & Dev.* **48**, No. 3/4, 379–394 (May/July 2004, this issue).
3. K. Astrom and T. Hagglund, *PID Controller: Theory, Design and Tuning*, Second Edition, Instrument Society of America, Research Triangle Park, NC, 1995.
4. M. L. Fair, C. R. Conklin, S. B. Swaney, P. J. Meaney, W. J. Clarke, L. C. Alves, I. N. Modi, F. Freier, W. Fischer, and N. E. Weber, "Reliability, Availability, and Serviceability (RAS) of the IBM eServer z990," *IBM J. Res. & Dev.* **48**, No. 3/4, 519–534 (May/July 2004, this issue).

Received October 9, 2003; accepted for publication March 30, 2004; Internet publication June 1, 2004

Gary F. Goth *IBM Systems and Technology Group, 2455 South Road, Poughkeepsie, New York 12601 (gfgoth@us.ibm.com)*. Mr. Goth received his B.S.E. degree from Princeton University in 1971, his M.S. degree in mechanical engineering from Union College in 1975, and his M.S. degree in statistics in 1978 from Rensselaer Polytechnic Institute. After seven years with the General Electric Company, he joined IBM in 1979. In IBM he has held managerial and technical positions in server group development. Currently a Senior Technical Staff Member, Mr. Goth is the thermal team leader for high-end servers. He holds 24 patents, has coauthored several technical publications, and has received IBM Outstanding Innovation and Outstanding Technical Achievement Awards in thermal technology.

Daniel J. Kearney *IBM Systems and Technology Group, 2455 South Road, Poughkeepsie, New York 12601 (yenraek@us.ibm.com)*. Mr. Kearney is an Advisory Engineer in the Power/Thermal Systems Development and Integration Department in the IBM Systems and Technology Group. He received his B.S. degree in electrical engineering from the L. C. Smith College of Engineering at Syracuse University in 1983, subsequently joining IBM. Mr. Kearney has been involved with power-supply design, power and thermal systems integration, and code development and testing. In 1988, he was a development manager for the power systems design group responsible for the ES/9000 mainframe. Mr. Kearney has contributed to the development of the IBM 9076 SP2 bulk power system, and the S390 G3 through G6 power systems integration and testing. He was responsible for power/thermal system code integration for the refrigeration components of the zSeries z900, and is currently responsible for power/thermal system code integration for the hybrid cooling system of the z990 server. Mr. Kearney has received two IBM Outstanding Technical Achievement Awards, has been granted ten U.S. patents, and has achieved his third IBM Invention Achievement Award plateau. He is a registered Professional Engineer in the state of New York.

Udo Meyer *IBM Systems and Technology Group, IBM Deutschland Entwicklung GmbH, Schoenaicherstrasse 220, 71032 Boeblingen, Germany (umeyer@de.ibm.com)*. Mr. Meyer received the graduate engineer (Dipl. Ing. grad.) degree in precision engineering from the Gauss-Schule in Berlin. He joined IBM in 1967, initially as a Customer Engineer, and worked as a region-designated specialist on IBM high-end computers. Since 1973 Mr. Meyer has worked in the IBM Boeblingen Laboratory, developing HW-related firmware and supporting control programs. He has received an IBM Team Award for his technical contributions to the IBM S/390 G5. He holds two patents.

Donald W. Porter *IBM Systems and Technology Group, 2455 South Road, Poughkeepsie, New York 12601 (porter@us.ibm.com)*. Mr. Porter received his B.S. and M.S. degrees in mechanical engineering from Rensselaer Polytechnic Institute in 1981 and 1994, respectively, and became a licensed Professional Engineer in the state of New York in 1997. He joined IBM in 1981 as a Manufacturing Engineer, large systems, and has held various management and non-management positions in both the manufacturing and development organizations. Mr. Porter is currently an Advisory Engineer; he is team leader for the zSeries modular cooling unit (MCU), responsible for development and release of refrigeration technology. He specializes in applying unique cooling solutions to IBM products and developing those designs through the conceptual, release, shipment, and field-support stages.