

IPACK2007-33870

CLOSED LOOP LIQUID COOLING FOR HIGH PERFORMANCE COMPUTER SYSTEMS

Sukhvinder Kang
kang@aavid.com

David Miller

John Cennamo

**Aavid Thermalloy LLC,
Concord, NH, 03301**

ABSTRACT

The power dissipation levels in high performance personal computers continue to increase rapidly while the silicon die temperature requirements remain unchanged or have been lowered. Advanced air cooling solutions for the major heat sources such as CPU and GPU modules use heat pipes and high flow rate fans to manage the heat load at the expense of significant increases in the sound power emitted by the computer system. Closed loop liquid cooling systems offer an excellent means to efficiently meet the combined challenges of high heat loads, low thermal resistance, and low noise while easily managing die level heat fluxes in excess of 500 W/cm².

This paper describes the design and attributes of an advanced liquid cooling system that can cool single or multiple heat sources within the computer system. The cooling system described use copper cold plates with meso scale channels to pick up heat from CPU and GPU type heat sources and highly efficient liquid-to-air heat exchangers with flat copper tubes and plain fins to transfer the heat to air by forced convection. A water based coolant is used for high thermal performance and additives are used to provide burst protection to the cooling system at temperatures down to -40 °C and corrosion protection to critical components. A highly reliable compact pump is used to circulate the fluid in a closed loop. The overall system is integrated using assembly methods and materials that enable very low fluid permeation for long life.

INTRODUCTION

The power dissipation levels in high performance electronics systems such as computers continue to increase rapidly while the silicon die temperature requirements remain unchanged or have been lowered. These electronics systems are typically cooled by forced air flow using axial fans or

centrifugal blowers. When the heat loads are small and relatively diffuse, thermal conduction through aluminum plates is sufficient to spread the heat into finned heat sinks for convection from the fins into the air flow stream. In recent years, as heat loads have increased and become more concentrated, better heat conductors such as copper plates and heat pipes have been used to improve the spreading of heat from heat sources such as the central processing unit (CPU) and the graphics processing unit (GPU) modules into the heat sink fins. These heat spreaders have made it possible to extend air cooling by enabling efficient convection from finned heat sinks that can be situated within a larger volume of space around the heat sources. However, the usable space for the finned heat sinks remains limited by, among other constraints, the heat collection and transport limitations of heat pipes [1] so higher flow rate system fans or additional local fans have to be used resulting in significant increases in the sound power emitted by the electronics system.

As the trend towards higher power dissipation and more concentrated heat sources continues, a more attractive solution is to use closed liquid cooling loops to efficiently spread heat to finned surfaces that can be situated almost anywhere within the electronics system. These closed loops may be two-phase capillary pumped systems such as loop heat pipes [2] or single or two-phase mechanically pumped systems such as the single-phase liquid cooling system used in the Apple G5 Power Mac computer [3]. The cooling loops may even be gravity aided thermosiphon systems using dielectric fluids [4] or vapor compression refrigeration systems such as that used in the IBM Z-series mainframe computers [5].

It is important to recognize that the closed loop cooling systems function only as efficient heat spreaders internally within the

electronics system enclosure but the electronics system is still “air cooled” from an external perspective. No additional limitation is placed on the type of environment and orientation that the electronics system can function in versus more conventional air cooled systems. With the efficient heat spreading within the electronics system enclosure made possible by the closed loop liquid cooling systems, the need for high speed fans is mitigated and the noise levels can be significantly lowered. At the same time it becomes possible to air-cool significantly greater and more concentrated heat loads within the electronics enclosure.

This paper describes the design and attributes of a high-performance single-phase liquid cooling system that can cool single or multiple heat sources within a computer system. The cooling system utilizes copper cold plates with meso scale channels to pick up heat from concentrated heat sources and a highly efficient liquid-to-air heat exchanger with flat copper tubes and plain fins to transfer the heat to the air flow. Several water based coolants with suitable corrosion inhibitors and antifreeze additives were qualified to provide burst protection to the cooling system at temperatures down to $-40\text{ }^{\circ}\text{C}$ and provide corrosion protection. A highly reliable centrifugal pump was used to circulate the fluid in the closed loop. The overall system was integrated using assembly methods and flexible tubes that enable very low fluid permeation for long operating life.

NOMENCLATURE

Q	Heat Load (W)
R_{fa}	Heat Exchanger fluid exit to air inlet thermal resistance ($^{\circ}\text{C}/\text{W}$)
R_{sf}	Cold Plate outer surface to fluid inlet thermal resistance ($^{\circ}\text{C}/\text{W}$)
R_{tim}	Thermal interface resistance ($^{\circ}\text{C}/\text{W}$)
T	Temperature ($^{\circ}\text{C}$)
ΔP	Pressure Difference (Pa or psi)

CLOSED LOOP COOLING SYSTEM DESIGN

The cooling system described here and pictured schematically in Figure 1 was designed for cooling two heat sources, the CPU and the GPU, in a Dell Dimension 9100 commercial PC. The computer had a dual core Intel Pentium D 840 CPU running at 3.2 GHz with a target heat dissipation of 130 Watt and a NVIDIA GeForce 7800 GTX graphics card with a target GPU heat dissipation of 80 Watts.

The overall liquid cooling system consists of the following key components:

- Pump
- Cold plates
- Heat Exchanger
- Flexible Tubing

- Coolant




Figure 1 Liquid Cooling System for a PC

Pump

The pump used in the cooling system is a well known highly reliable DC pump with a spherical bearing system that is designed to perform reliably for more than 50,000 hours of continuous operation [6].

Cold Plates

A cold plate is the component within the liquid cooling system that physically contacts the electronic module and enables the transfer of heat from the electronic module into the liquid coolant. The present cold plate embodies a parallel plate design with meso-scale fluid channels like the cold plates used in the Apple G5 Power Mac computer [3]. Thermal performance estimates for the cold plate were made using well known heat transfer correlations for laminar flow in rectangular ducts summarized in [7]. Figure 2 provides an exploded view of the present cold plate design showing the base plate, the fins, and the cover. An elastomeric o-ring is provided between the base plate and cover to prevent coolant leakage from the cold plate.

In the CPU cold plate, a uniform array of 48 copper fins each 0.3 mm thick and 6.35 mm tall are attached to a 2.5 mm thick copper base plate over an area that is 22 mm wide and 12 mm long. The base plate is circular with a diameter of 52 mm. A cover with barbed inlet and outlet fittings is provided over the base plate such that liquid flow enters into an inlet plenum on one side of the fins and flows transversely through the channels between the fins before exiting through the outlet. The fluid cavity in the cover is designed with minimum clearance around the fin array to efficiently direct the liquid flow through the cooling channels between the fins without bypass around the fin array.

The GPU cold plate design is identical to the CPU cold plate except that the area covered by the fins is 22 mm wide by

18 mm long to account for the different die size and only 38 fins are used at a larger fin pitch so that the pressure drop is similar to the CPU cold plate. The fluid cavity in the cover of the GPU cold plate is modified accordingly to accept the altered finned area.




Figure 2 Copper Cold Plate

Heat Exchanger

The heat exchanger component of a liquid cooling system enables the transfer of heat from the liquid coolant into the stream of air that flows through it. Its function is similar to a radiator in an automobile. In the present design pictured in Figure 1, copper alloy tubes and liquid manifolds are used to provide excellent chemical compatibility with the aqueous liquid coolant. The only metals that the liquid coolant comes in contact with are copper and copper alloys to avoid any possibility of galvanic corrosion. Low mass aluminum plate fins are attached to the outer surface of the copper alloy tubes for heat transfer to the air. Plate fins provide the most attractive balance between the air pressure drop and air side thermal performance. Thermal and pressure drop estimates for the heat exchanger were made using well known heat transfer correlations for rectangular ducts and heat exchangers summarized respectively in [7] and [8].

The size of the present heat exchanger is 124 mm wide by 146 mm tall by 25 mm deep. It consists of four copper alloy flat tubes, each 19 mm wide by 2.34 mm high and 124 mm long, connected to the two manifolds. The face area of the fins, 124 mm wide by 122 mm tall, is populated with a uniformly spaced array of 0.254 mm thick and 25 mm deep aluminum fins on a 1.254 mm pitch along the length of the flat tubes.

Liquid flows in and out of the heat exchanger through two barbed fittings provided in the manifolds. Another quick disconnect fitting is provided on the manifold for charging the

cooling system with coolant. Air flow through the heat exchanger is provided by a 120 x 120 mm system fan in the computer. A steel sheet metal shroud is provided around the heat exchanger to secure it to the computer chassis and to direct the air flow from the fan through the heat exchanger. The steel shroud also provides a platform for mounting the liquid pump downstream of the fluid outlet on the heat exchanger manifold.

Flexible Tubing

The interconnect tubing provides the means to transfer cooling fluid between the liquid cooling system components. A corrugated Fluorinated Ethylene Propylene (FEP) tube that was specifically developed for long life liquid cooling systems was chosen [9]. The tubing is made sufficiently flexible because of the corrugations that it transmits very low mechanical forces to the cold plates while providing a very low rate of coolant loss by permeation through its walls. These attributes are important to prevent mechanical damage to the electronics components that the cold plates are mounted on while enabling factory sealed liquid cooling systems that do not need to be refilled during a typical computer system lifetime of 5 to 7 years.

Experimental measurements show that the permeation rate of water through the tube at a temperature of 45 °C is 5.6 mg/m-day. For the 0.5 m length of tubing used in the system, the anticipated loss of coolant is less than 10 g over seven years of operation. A 15 cm³ flexible accumulator is provided within the pump to accommodate up to 15 g of coolant loss.

Coolant

Three aqueous coolants were studied for use in single-phase liquid cooling systems. While the exact compositions evaluated and selected are proprietary, the coolants consist of distilled water, an antifreeze additive such as ethylene glycol, and a suitable corrosion inhibitor. The coolant composition is selected to provide high thermal performance, protection from damage in freezing environments encountered in shipping, and long-term protection against corrosion during system operation.

From a thermal performance perspective alone, pure water would be the best choice for the coolant but the cooling system would burst due to the large volumetric expansion of water (~10%) when it freezes. A suitable corrosion inhibitor is necessary to prevent electrochemical interaction between the ionic aqueous coolant and metal components in the cooling system. Dielectric fluids such as fluorocarbons (FCs) and hydro-fluoro-ethers (HFEs) offer protection against freezing and corrosion but their single-phase heat transfer performance is markedly low compared to aqueous coolants.

The aqueous coolant composition was developed to provide good thermal performance at typical room temperature operating environments (10 °C to 40 °C) and good protection against bursting during shipping environments down to -40 °C.

Commercial corrosion inhibitor packages were selected based on the metals systems they are designed to protect. For completeness, simulated system corrosion tests [10] were conducted with cooling systems that contained (1) only copper, copper alloys, and joining materials such as soldering and brazing alloys and (2) the same copper and joining alloys as well as aluminum alloys that are commonly used to make brazed aluminum heat exchangers.

In the corrosion tests, electrical contact was allowed between like metals such as copper, copper alloys, solders and brazing alloys but metals with large differences in galvanic potential were separated at least 6 mm apart using electrically isolative spacers. This mimics real systems since copper and copper alloy components could be soldered or brazed to each other whereas copper and aluminum components would be connected with an electrically insulating tube or spacer. The simulated systems included high liquid velocity over the copper test surfaces to mimic potential erosion corrosion in the narrow fluid passages in the cold plates. Tests were carried out at high fluid temperatures to enhance chemical activity. The coolant was aerated to simulate dissolved air in the initial coolant charge and air and coolant permeation through the flexible FEP tubing over time.

Tests were conducted for a period of 1064 hours at 88 °C for each antifreeze and corrosion inhibitor combination. The mass loss of the metal coupons was subsequently measured to determine the corrosion rate. The surface of the metal coupons was also examined under an optical microscope to check for localized corrosion like erosion and pitting. Sample corrosion test data presented in Figure 3 shows that the copper and copper alloy materials are much more robust to corrosion when used with aqueous coolants compared to aluminum alloys.




Figure 3 Sample Corrosion Test Data

Cooling System Integration

The liquid cooling system is intended as a fully integrated and tested system for easy assembly into the computer system for which it was designed. During cooling system integration,

a special fixture was used to hold the pump, the two cold plates, and the heat exchanger in the correct relative positions. Four separate pieces of FEP tubing were made in the precise lengths needed to interconnect each leg of the cooling system. Custom tooling was used to expand each end of a FEP tube and force it over the barbed fitting on the cooling system component – the tubing then contracted around the barbed fitting and formed a leak tight joint. The fully assembled cooling system was evacuated and checked for leaks using a helium leak detector. The measured helium leak rate was 4E-7 cc-atm/s. For the demonstration cooling system, the corrosion tests had not been completed so a commercial automotive coolant composed of 50% ethylene glycol by volume in water was charged into the cooling system to a positive system pressure of ~0.6 bars.

When the assembly and coolant charging were completed, the pump was attached to the steel housing of the heat exchanger. The cooling system was then placed inside the computer system and the CPU and GPU cold plates were clamped to the surfaces of the respective components while the steel housing was attached to the computer frame. The computer with the liquid cooling system installation is pictured in Figure 4.

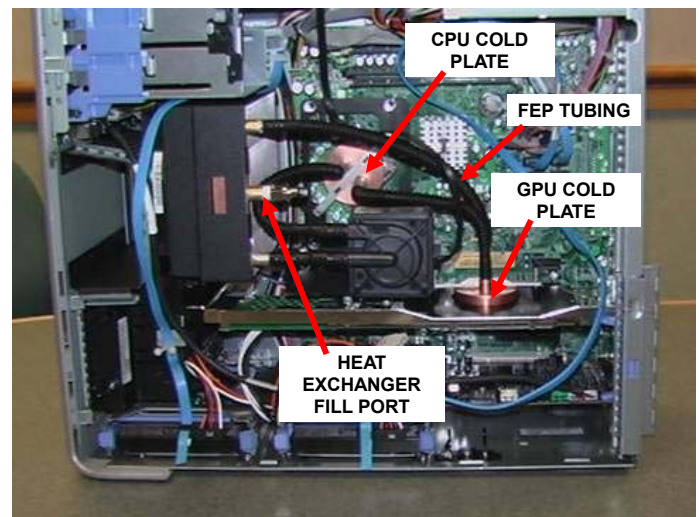


Figure 4 Liquid Cooling System installed in the PC

THERMAL PERFORMANCE TESTING

Lab Bench Testing

The thermal performance of the individual cold plates was measured using a laboratory liquid flow bench with water as the coolant. In these tests, each cold plate was mounted on custom heat sources corresponding to each application (12x21 mm for the CPU and 18x18 mm for the GPU) while water was supplied to the cold plate at a known flow rate and temperature. Because of the small heat capacities and the low thermal

resistance of the cold plates steady state temperatures were achieved within one minute.

Thermocouples were used to measure the temperature at the center of the cold plate base plate where it contacted the heat source and in the liquid inlet line of the cold plate. Liquid flow rate was measured using a coriolis mass flow meter (Micro Motion R025) with an accuracy of +/-0.5% and the pressure drop across the cold plate was measured using a differential pressure sensor (Sensotec FDW) with a 15 psid range and an accuracy of +/- 0.075 psi. After the water flow was established, electrical power was switched on to dissipate 150 watts in the heat source. The thermocouple data was recorded using an Agilent 34970 data logger while the heat dissipation, mass flow rate, and pressure drop were recorded by hand from the instrument displays.

Individual experimental data sets consisted of the liquid flow rate, the pressure drop across the cold plate, and the thermal resistance of the cold plate. Cold plate thermal resistance (R_{sf}) was calculated by dividing the heat dissipated in the heat source by the temperature difference from the base plate outer surface to the liquid inlet. No correction was made for extraneous heat losses from the heat source. The extraneous losses can be estimated by noting that the thermal resistance of the extraneous heat loss paths is about 8 °C/W. Pressure drop readings were corrected by subtracting the pressure drop separately measured across the same fittings and length of tubing that was used to connect the device under test to the liquid flow bench.

The liquid-to-air heat exchanger was tested by attaching the liquid flow path of the heat exchanger to the same laboratory liquid flow bench and attaching the heat exchanger to a duct on the suction side of a laboratory air flow bench. In this setup, the liquid flow rate and the air flow rate through the heat exchanger could be independently controlled. Air flow rate was measured using the discharge coefficient of standard ASME long radius nozzles along with the barometric pressure and air temperature and the air side pressure drop was measured with an accuracy of +/- 0.3 Pa using a differential pressure transducer (Setra 264 with a range of 125 Pa) connected across a pressure tap in ambient air and a wall pressure tap in the duct downstream of the heat exchanger. An electrical heater was used to heat the liquid entering the heat exchanger about 30 °C above the air inlet temperature. Thermistors, YSI 46044 series, were used to read the liquid inlet and exit temperatures with an accuracy of +/- 0.05 °C. Thermocouples were used to read the air inlet and exit temperatures. The net rate of heat transfer from the water to the air was calculated by averaging the energy balance on the liquid and air streams. With the water flow rate set to 2 liters per minute (LPM), at the lowest air flow rate of 20 cubic feet per minute (CFM), the energy balance in the water and air

streams differed by 16%. The difference dropped to 7% at 40 CFM and 2% at 80CFM.

Individual experimental data sets consisted of the liquid flow rate, the pressure drop across the heat exchanger, and the thermal resistance of the heat exchanger. The thermal resistance was calculated by dividing the heat loss from the heat exchanger by the temperature difference from the liquid exit to air inlet temperature. This definition makes it easy to calculate the liquid temperature at the inlet of the cold plates by simply multiplying the heat exchanger thermal resistance by the total heat load on the liquid cooling loop.

In-System Testing

The in-system tests were more qualitative in nature. The main objective of these tests was to study the maximum device temperature reported by a case temperature thermocouple on the CPU while using a synthetic computational load on the CPU and simultaneously executing an automated video game to maximize the heat dissipation in the GPU. The well-known CPU Burn program [11] was used to maximize the CPU thermal load and the “Medal of Honor” Video Game was used to simultaneously load the GPU.

A second objective was to determine the fan speed response to the computational loading of the CPU and GPU to assess the potential acoustic benefits of using the liquid cooling system. For comparison, the computer system, with the as received air-cooled heat sinks on the CPU and GPU was observed to increase the system fan speed to the maximum setting under similar computational loads.

RESULTS

CPU Cold Plate Performance: Figure 5 shows the thermal performance of the CPU cold plate measured in the bench tests using water as the coolant. The thermal resistance R_{sf} data was collected using a 21 mm wide by 12 mm long heat source aligned with the 22 mm wide by 12 mm long finned area inside the cold plate. The thermal resistance of the cold plate agreed well with the design intent. The pressure drop, shown here for completeness, was used in the total system flow impedance to estimate the operating point of the completed liquid cooling system. The thermal performance and pressure drop of the GPU cold plate was similar to that of the CPU cold plate.

Given the good agreement between the test data and the design thermal resistance calculations for water, we predicted the thermal resistance for the 50% ethylene glycol coolant used in the unit built for cooling the PC computer. This prediction is displayed on Figure 5. With the 50% ethylene glycol coolant the liquid flow rate is expected to be only 0.4 GPM and the cold plate performance is predicted to be 0.066 C/W.

Heat Exchanger Performance: Figure 6 displays the performance of the liquid to air heat exchanger measured in the

bench tests using 0.53 GPM of water coolant flow rate over a range of air flow rates. The figure presents the fluid-to-air thermal resistance R_{fa} of the heat exchanger and the air side pressure drop versus the air flow rate. The design thermal resistance, included on the graph, compares well with the test data and provides a good basis to predict the cold plate performance for other coolants and flow rates. The performance with the 50% ethylene glycol coolant at a flow rate of 0.4 GPM is included in the graph. Separate tests were done to determine the air flow rate through the heat exchanger in combination with the system fan as the fan speed was changed by varying the DC voltage supplied to the fan. At the lowest or idling fan speed, the expected air flow rate is ~48 CFM. At this air flow rate, the thermal resistance of the heat exchanger is measured at 0.047 °C/W with water and predicted to be 0.084 °C/W with a 0.4 GPM flow rate of the 50% ethylene glycol coolant. At the highest fan speed, the CFM through the heat exchanger is expected to be ~95 CFM, and the thermal resistance is predicted to be 0.073 °C/W with the 50% ethylene glycol coolant.




Figure 5 CPU Cold Plate Performance

Cooling System performance in the computer: Based on the bench test results with water and scaled predictions for the 50% ethylene glycol coolant it is possible to predict the thermal performance of the liquid cooling system in the computer system. Under a full computational and video graphics load, the CPU and GPU could dissipate as much as 130 watts and 80 watts respectively. In addition, we assume that the pump electrical power of 10 Watts is transferred to the liquid coolant

being pumped. Using this information the CPU case temperature (T_c) can be predicted as follows:

$$T_c = T_a + R_{fa} \cdot (Q_{cpu} + Q_{gpu} + Q_{pump}) + (R_{sf} + R_{int}) \cdot Q_{cpu}$$

where the relevant thermal resistances are:

- $R_{fa} = 0.084 \text{ C/W}$
- $R_{sf} = 0.066 \text{ C/W}$
- $R_{int} = 0.026 \text{ C/W}$ between the CPU and cold plate
- $T_a = 25 \text{ C}$ air temperature at heat exchanger inlet

Substituting the above resistances and power levels we calculate a case temperature of:

$$T_c = 25 + 0.084 \times 220 + (0.066 + 0.026) \times 130 = 55.4^\circ \text{C}$$




Figure 6 Heat Exchanger Performance

In system tests: The liquid cooling system with the 50% ethylene glycol aqueous coolant was installed in the computer and tested for qualitative performance. The computer system BIOS provided automatic fan speed control in response to temperature sensors on the CPU die. We independently monitored the case temperature of the CPU module using a thermocouple placed in a slot cut into the top surface of the integrated heat spreader that formed the CPU case. A second thermocouple was located in the air inlet grille ahead of the computer system fan.

As noted earlier, during these tests both the CPU and GPU were fully loaded with computational tasks using a CPU

exerciser and a video game. The CPU process monitor showed 100% activity in both CPU cores. In long duration tests lasting many hours under these conditions, the CPU case temperature reading was observed to peak at 57 °C in a room temperature environment of 25 °C. This is consistent with the expected performance based on the bench test results. The GPU temperature was not monitored but video performance issues were never encountered during the tests.

An important additional observation is that the computer did not speed up the system fan above its lowest speed over many months of testing under synthetic and actual computational loads. This resulted in a very low noise experience for the user compared to the high fan speeds and noisy operation that was observed with the original air cooled heat sinks.

CONCLUSIONS

As the power dissipation in high performance personal computers continues to increase, closed loop liquid cooling systems offer an excellent means to provide efficient reliable cooling with high cooling capacity and low noise. The design of all the key components of a liquid cooling system for PC's was described. System integration for long life operation using low permeation flexible FPE tubing, vacuum leak testing and coolant filling to a positive system pressure were described.

The wetted components of the cooling system were made from copper and copper alloys which provide very good reliability with aqueous coolants in single-phase liquid cooling systems. Aluminum plate fins were used on the air side of the heat exchanger for low mass. Corrosion tests were performed to identify attractive coolant chemistries that provide long term corrosion protection of the wetted components as well as freeze protection down to -40 °C.

Bench tests were performed with water as the liquid coolant to measure the performance of the CPU cold plate and the liquid to air heat exchanger to validate the design methods for the respective components. The validated design methods were used to scale the performance to the 50% ethylene glycol coolant used in the demonstration cooling system for a PC.

The performance of the liquid cooling system in the computer was in accord with predictions based on bench test data using water as the coolant. Low noise operation of the computer was demonstrated under high computational loads over several months as a significant benefit of the liquid cooling system.

ACKNOWLEDGMENTS

The authors would like to recognize Tom Seow for designing the cold plates and David Gailus for assistance and advice on materials selection and for the corrosion testing.

REFERENCES

1. Peterson, G.P., *An Introduction to Heat Pipes: Modeling, Testing and Applications*, Fourth Ed., Wiley, New York, 1994.
2. Zimbeck, W., Yun, J., Cennamo, J., Kroliczek, E., "Multi-evaporator hybrid loop heat pipe for cooling high performance servers", Proc. THERMES 2007, Millpress, Rotterdam, 2007, pp. 249-256.
3. Eric Bangeman, "Dual 2.5GHz Power Mac G5", ARS Technica Review, Sept. 29, 2004, <http://arstechnica.com/reviews/hardware/G5.ars/6>
4. Tuma, P., "Evaporator/Boiler design for thermosyphons utilizing segregated hydrofluoroether working fluids," Proc. 22nd IEEE SEMI-THERM, 2006
5. Schmidt R. R. and Notohardjono, B. D., "High end server low temperature cooling," IBM J. Res. & Dev., Vol. 46, No. 6, Nov 2002, pp. 739-751.
6. Laing DDC-1 Pump, <http://www.lainginc.com/DDC-1.htm>
7. Kakac, S., Shah, R.K., Aung, W., *Handbook of Single-Phase Convective Heat Transfer*, 1987, Second Edition, Wiley & Sons, New York, 1987, chap. 3 & 4.
8. Kays, W.M., London, A.L., *Compact Heat Exchangers*, Third Edition, McGraw-Hill, Inc., New York, 1984.
9. Valdez, F. A., "Tubing Selection for Sealed-for-Life Liquid Cooling Units for Electronics," Proc. IEEE IOTHERM 2006, pp. 167-174.
10. *Standard Test Method for Simulated Service Corrosion Testing of Engine Coolants*, ASTM D2570-96(2002)e1
11. CPU Burn-in program, <http://users.bigpond.net.au/cpuburn/>