Proceedings of IPACK'03
International Electronic Packaging Technical Conference and Exhibition
July 6–11, 2003, Maui, Hawaii, USA

INTERPACK2003-35338

DESIGN OPTIMIZATION OF A ROUTER BOARD USING COMPUTATIONAL FLUID DYNAMICS

Timothy J. Dake and Joseph Majdalani[1]
Department of Mechanical and Industrial Engineering
Marquette University, Milwaukee, WI 53233
Phone: (414) 288-6877; Email: maji@mu.edu

Susheela Narasimhan
Cisco Systems
250 West Tasman Drive, San Jose, CA 95134
Phone: (408) 527-6334; Email: snn@cisco.com

The cooling of boards and components in telecommunication chassis is becoming progressively more challenging as the power dissipation of the components increases. The internet market is demanding higher speeds and the internet gear makers try to accommodate the increasing demand of the consumer. Hence there are situations where the power requirement for one rack unit of the chassis is reaching 500-600 W currently and expected to reach about 800-1000 W in 1-2 years. The power dissipations on components are currently in the vicinity of approximately 60-80 W and expected to reach roughly 100-120 W in the near future. This paper presents a design effort wherein a chassis has been analyzed and designed for a cooling capacity on the order of 600 W for the board to the extent that one is able to cool about 80 W for the ASICs and other components on the board.

1 INTRODUCTION

The design limitations of telecom chassis require that the length, width and height of the unit be severely restricted to standardized dimensions in order to accommodate several boards within a single router. This constraint poses certain barriers in the planning for the airflow and in the optimization of the thermal management of the board chips. To successfully address the problems of supplying sufficient airflow and maximizing the heat transfer from the chips to the airflow, it becomes necessary to employ the latest tools such as CFD and a mixture of new technologies in forced convection cooling.

The processor chips utilized in the new generation of routers have a maximum case service temperature specified at 75°C. The typical inlet temperature of the cooling airstream is 55°C. This temperature of inlet air is high because the internet gear manufacturer must adhere to Network Equipment Building System (NEBS) and other industry requirements in designing the chassis. The current generation of production model routers have a processor power dissipation of about 80 W. This places a firm demand on the thermal engineers employed in the electronics cooling industry to have a design that will give a thermal resistance of about 0.18 °C/W. The maximum amount of airflow possible is constrained by several factors that include noise, EMI issues, characteristics of commercially available fans, board obstructions and layout, and utilization of good design practices that limit the maximum velocity of the airflow. The thermal resistance of the heat sink coupled with the pressure drop in the airflow, can create an upper limitation on the forced convection heat transfer. The latest generation of heat sinks employs a large mass with a complex geometry in the fin arrangement.

Heat sinks are frequently used to enhance the rate of heat dissipation from printed circuit boards and other power sources within computers and electronic enclosures. Recently, the study of heat sinks under both forced and natural convection has received much attention in view of the consistent trends in microelectronic design leading towards higher packaging densities and higher power dissipation rates. Due to their

[1] To whom correspondence should be addressed.

inherent simplicity, reliability, and low long-term costs, forced convection heat sinks have proven to be instrumental in cooling single or multiple-chip circuit boards with high dissipative heat loads [1]. As a result of the modern trends of increased chip size and performance, there is a mounting number of cases that arises in industrial design where a multiple array of heat sinks need to be modeled using computational fluid dynamics (CFD).

In the case of the new router, several chips of varying wattage are employed on a common board. The flow of the air over and around the chips is difficult to predict and analyze. The interaction of the airflow with the chips is complex leading to a varied temperature and velocity distribution as the airflow changes course within the router chassis. The heat transfer and thermodynamic analysis is complicated requiring the use of a computational package. Studies of heat sink performance are prevalent in the literature and most have been focused on design optimization and selection details of heat sink components. For a more elaborate literature survey, the reader may wish to consult with the works of Lee [2] or Patel and Belady [3] on design optimization and selection.

As the limitation of using forced convection cooling methods is approached, the emphasis must shift from an optimization of the heat sink design to an assurance of the optimal management of the airflow through the heat sink. Much of the efficacy of the ability of the airflow to remove heat is lost in the bypass of air over and between the heat sinks. This air does not make contact with the heat sink surface although power is used to move the air. Recapturing the work potential of this bypass air provides an opportunity to enhance the heat transfer from the chips. The airflow direction and volume are carefully controlled to make as full a use of the heat sink potential as is possible. Careful analysis of these critical facets of the airflow leads to a maximization of the heat transfer from the chip.

This paper consists of one such study that is carried out using a commercial CFD package to analyze the robustness of the airflow model under a forced convection flow regime. The router is modeled as a system to account for all of the airflow and the restrictions within the total system.

2 METHODOLOGY

The chassis used in this study is a 4RU box with three boards. The rack space available per board is 1.333 Rack Units (2.31"). The three dimensions of the chassis are given by 17.25"(L)×21.67"(W)×6.92"(H). Two large system exhaust fans are used to draw the flow out of the rear of the chassis. The two fans used in this set-up have the specifications given in Table 1.

Figure 1 shows the chassis. The board layout is as shown in Fig. 2. The two high power dissipating components (U1 and U2) on the board are the processors which have a power dissipation of generally 80 W each. The temperature of the incoming air is specified as 55°C. Even though NEBS specifies a maximum air inlet temperature of 50°C for short times, Telecom Gear

Table 1. Fan specifications

Static pressure at zero flow	1.083" of water
Flow at zero static pressure	325 CFM
Noise	59 DBA
Voltage	24 V
Current	1.1 A
Power	26.4 W
RPM	4300

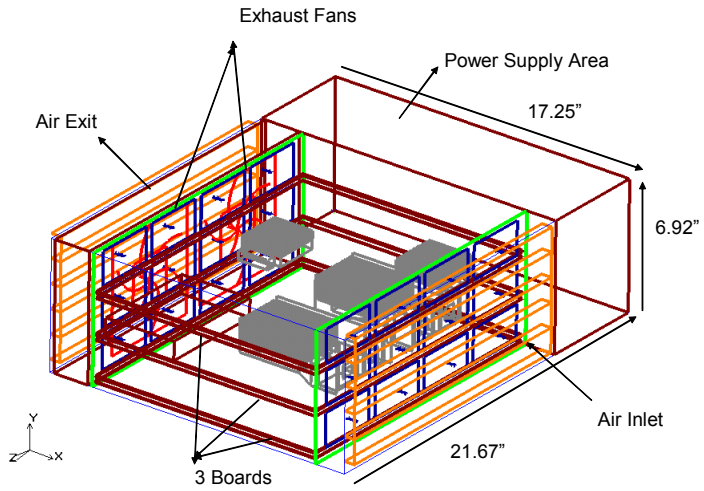

Fig. 1 Computational model of the chassis used in the CFD simulations.

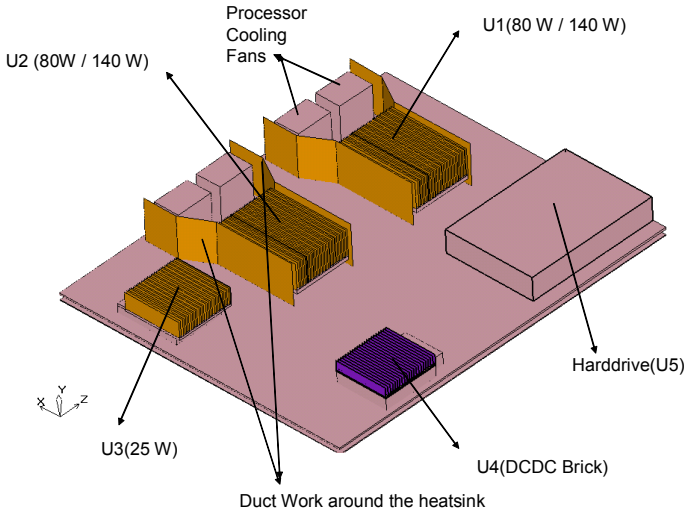

Fig. 2 Component detail of the middle board.

Manufacturer’s Internal MDVT (Mechanical Design and Verification Test) specifications call for 55°C in order to achieve an extra 5°C margin. The maximum allowable case temperature on the processors is about 75°C. This specification places a demand on the thermal resistance of the cooling device to be in

the region of 0.18-0.20 °C/W. The third processor (U3) has a 25 W power dissipation. The other components on the board include a harddrive (U5) and a DCDC converter (U4).

The challenge is to cool the two high power dissipating processors and maintain the processors at or below the maximum 75°C specified by the processor manufacturer. The methodology to cool the processors consists of employing (1.) two 40×40×28 mm fans to blow the air through the heatsink cooling the processor, (2.) duct work to guide the airflow into the heatsink and (3.) the use of a standard footprint heatsink of dimensions 2.5"×3.5"×1.0". The heatsink base is 4 mm thick using a total of 36 fins on the heatsink with each fin having a thickness of 0.4 mm. The heatsink base thickness of 4 mm was selected in such a way to allow for maximized heat spreading across the base. The fin thickness and number of fins were selected to achieve an optimal performance of the heatsink. The board under consideration is in the middle slot with an inplane conductivity of 18 W/m-K and a normal conductivity of 0.35 W/m-K. For socketed processors, the majority of the heat travels to the heatsink and hence an effort to increase the board conductivity by increasing the number of copper layers will not yield any increase in heatsink performance.

A phase change interface material (with a thermal impedance of 0.08 °C-in^2/W) was used between the processor and the heatsink in the simulation. The upper and lower slots are modeled using flow resistance objects which have appropriate loss coefficient values to take into account the flow impedance of a single board.

3 CFD SIMULATIONS

The CFD simulations are carried out using Icepak, an electronics cooling software package [4] that has been previously used in analyzing the heat transfer character in electronics and chip cooling applications. According to the code algorithm, the flow and temperature fields in the computational domain are obtained by solving the three dimensional Navier-Stokes equations along with the energy conservation equation using a finite volume solver. The finite volume solver also employs multigrid acceleration for the solution of the non-linear differential equations. The multigrid acceleration helps to reduce the error by iterating on a series of coarse and fine grid meshes alternatively. The convergence criterion used is based on a tolerance that ensures that the residual error remains under 10^{-3} in the momentum balance, and less than 10^{-7} in the energy balance [4].

Two different simulations were carried out for the chassis. The first simulation was conducted with a power of 80 W on each of the processors whereas the second simulation was performed with a power of 140 W for each of the processors (specifically, U1 and U2 shown in Fig. 2).

4 RESULTS AND DISCUSSION

A total flow in the vicinity of 350 CFM was obtained for the whole chassis. The flow for each slot was more or less 116 CFM. Using the equation,

$$\dot{Q} = \dot{m} C_p \Delta T \qquad (1)$$

we can see that each slot can cool in the region of 640 W per slot. Here $\dot{m}$ is the equivalent mass flow rate of air (for 116 CFM), C_p is the specific heat of air (1005 J/kg K) and ΔT is 10°C. Each system fan delivered generally 175 CFM of air through the chassis. Figure 3 shows the airflow vectors through the chassis and Fig. 4 shows the speed profiles for the chassis.

It can be seen that the speed of the airflow approaching the heatsink is about 4-5 m/s for the processors. The presence of the local processor fans and also the duct work around the processor heatsink help to provide high speed and a large quantity of

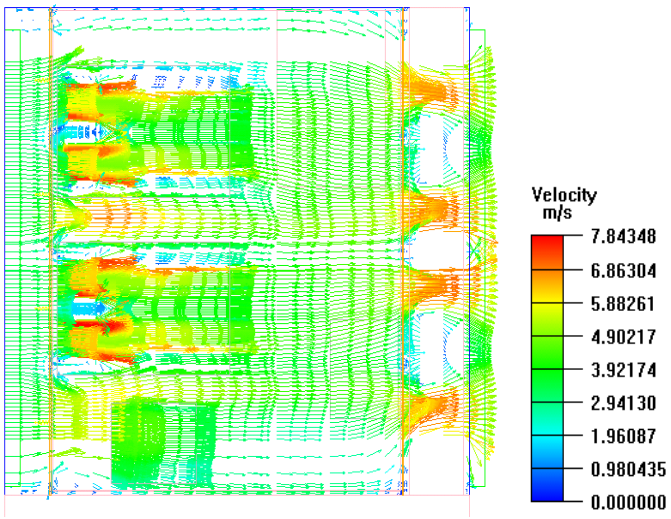

Fig. 3 Vector plot of air flowing through the middle slot.

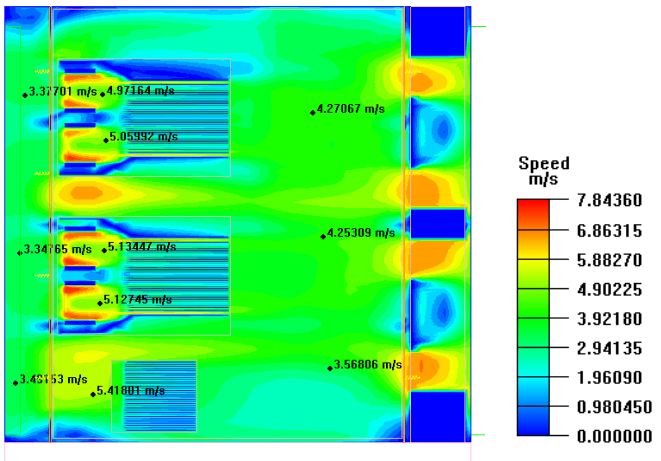

Fig. 4 Speed contours above the middle board.

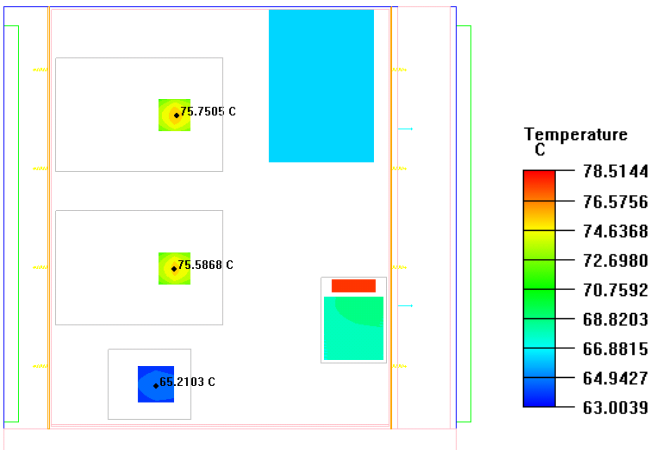

Fig. 5 Component temperatures (case 1: 80 W per processor).

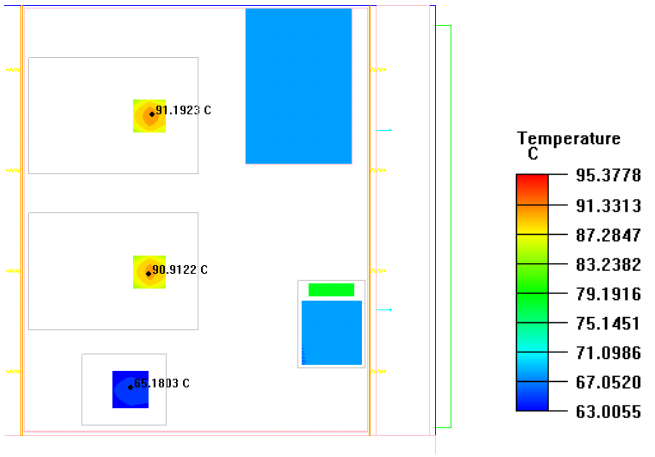

Fig. 6 Component temperatures (case 2: 140 W per processor).

volumetric airflow to the processor heatsinks. Figure 5 shows the temperatures of the different components on the board for Case 1 in which simulations were run with 80 W on each of the processors.

From Fig. 5, it can be observed that the processors are at a temperature approaching 75°C, which is close to the maximum temperature specification for the processors. The heatsinks used on the processor have dimensions of 2.5"×3.5"×1" standard footprint with the dimension of 3.5 inches being in the airflow direction. The number of fins used in the heatsink is 36, which gives a fin pitch of about 0.07 inches. This fin pitch is good for the high speed flow approaching the heatsinks. The airflow is augmented by the presence of two processor fans in the front of the processors. In standard board configurations, (with the absence of processor fans to accelerate the flow into the processor heatsink) the number of fins cannot be made as dense as the one proposed in this study. For most telecom applications, the flow approaching the heatsink will be around 200-300 LFM (1-1.5 m/sec). With such a velocity, a high fin density like the one used in this study might not be feasible. Hence, the use of on-board fans with ducting to channel air into the processor heatsinks provides the forced aircooling option for high power dissipating components (with 80 W) with a very low maximum case temperature specification.

Figure 6 shows the temperature profiles on the processor for the second simulation where 140 W were applied to the processors. The figure shows that the maximum temperature on the processors is about 95°C. This is suitable for cooling ASICs, which have a maximum junction temperature specification of about 125°C. Therefore, if we have ASICs with very extreme power dissipations of 140 W, it can be seen that these processors can be cooled sufficiently and effectively.

The heatsink temperature profiles are depicted in Fig. 7 for the Case 2 simulation wherein a power of 140 W was used on the processors. Figure 8 shows the temperature profiles on the base of the heatsink. It should be noted that the base of the heatsink used in this simulation was made out of aluminum.

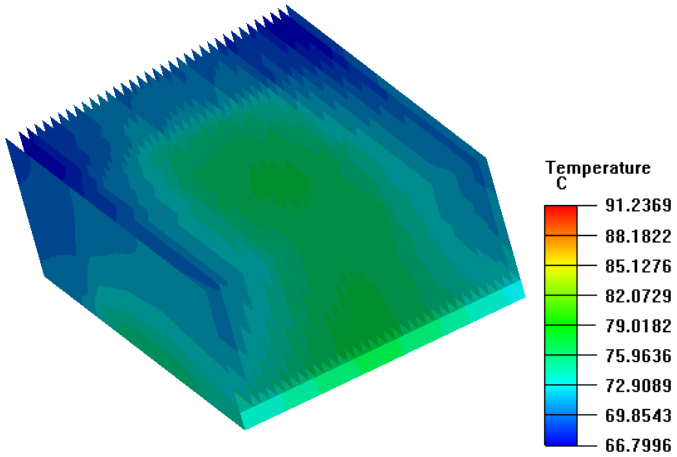

Fig. 7 Temperature profiles along heat sink (case 2: 140 W per processor).

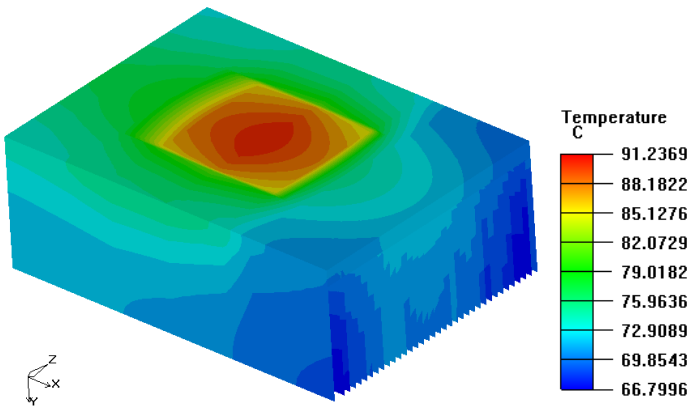

Fig. 8 Base temperature (case 2: 140 W per processor).

Consequently it can be determined from Fig. 8 that there is a significant temperature gradient on the base of the heatsink due to the spreading resistance of the heatsink. Some significant improvement in performance and the cooling capacity can be obtained when the heatsink uses a copper base, aluminum base with heatpipes, copper base with heatpipes or a heatsink base holding a vapor chamber.

5 CONCLUSIONS

In this study, CFD analysis was carried out on a 3 slot, 4 RU chassis with large system fans. The total CFM obtained for the whole system was about 350 and roughly 116 CFM per slot. The total cooling capacity of each slot was approximately 640 W. One of the boards analyzed in the system had two high power dissipating components (processors). Simulations were carried out for these two critical components at both 80 W and 140 W per processor. It was observed that in both cases, with the processor at 80 W or at 140 W each, the usage of on-board fans, a very dense fin pitch heatsink and duct work to guide the airflow dramatically influenced the heat transfer to achieve the objective.

By a well planned manipulation of the airflow and a thorough, careful design of the heatsink, a maximum heat transfer from the chip to the airstream was accomplished without the need for a complex, high cost, and space consuming liquid, jet impingement, or compact refrigeration system. The efficacy of CFD simulations in lending support to experimentation trials was confirmed by the success of our numerical predictions. Though the use of a computational code, the time to complete the thermal analysis was drastically shortened. In addition to the savings at the virtual design stage, a drastic reduction in the number of required test trials was achieved. Most notably, this design effort resulted in an extension of forced convection to a high power dissipation application that would other require alternative cooling technologies.

6 ACKNOWLEDGEMENTS

This work was partially supported by NASA and the Wisconsin Space Grant Consortium. Particular thanks are owed to M. Michael Yovanovich, R. Aileen Yingst, and Steven I. Dutch.

REFERENCES

[1] Lee, S., Culham, J. R., and Yovanovich, M. M., 1991, "Effect of Common Design Parameters on the Thermal Performance of Microelectronic Equipment. Part 1 – Natural Convection," Heat Transfer in Electronic Equipment, pp. 47-54.

[2] Lee, S., 1995, "Optimum Design and Selection of Heat Sinks," Eleventh IEEE SEMI-THERM Symposium, pp. 48-54.

[3] Patel, C. D., and Belady, C. L., Modeling and Metrology in High Performance Heat Sink Design, Hewlett Packard Laboratories, Palo Alto, California, 1997.

[4] Icepak Reference Manual, 3.1 ed., Fluent Inc., Palo Alto, California, 1999.