AIAA 2003-4186

Optimization of a New Telecom Board Using CFD and Enhanced Forced Convection

Timothy J. Dake and Joseph Majdalani
Marquette University
Milwaukee, WI 53233

Susheela Narasimhan
Cisco System Incorporated
San Jose, CA 95134

36th AIAA Thermophsics Conference
23–26 June 2003
Orlando, FL

36th AIAA Thermophysics Conference
23-26 June 2003, Orlando, FL

AIAA 2003-4186

Optimization of a New Telecom Board Using CFD and Enhanced Forced Convection

Timothy J. Dake* and Joseph Majdalani†
Marquette University, Milwaukee, WI 53233
and
Susheela Narasimhan‡
Cisco Systems, San Jose, CA 95134

Following a series of numerical simulations using a lightweight aluminum base/copper fin heatsink, stable cooling at the board level has been achieved for the new telecom motherboard corresponding to a one-rack unit intended for the next generation of telecoms. The new system relies on standard forced convection cooling, thus precluding the need for the more advanced concepts derived from the use of heatpipes, refrigerants or liquid based thermal management systems. The board under consideration operates with multiple chips that vary in thermal loading up to 70 Watts for a total cooling capacity per board of 200 W, given an inlet air temperature of 55°C. This paper addresses the methodology followed in developing the thermal management solution for this board level structure. A description of the design and the approaches taken in optimizing the thermal outcome is detailed with an emphasis on the measures undertaken to limit the chip temperature rise and to maximize the airflow. The study illustrates the manner in which bimetallic heat sinks combined with carefully designed ducting for optimal airflow management can be conveniently used to reduce the needed computational cost, time and resources while producing sufficiently reliable results.

I. Introduction

THE DESIGN limitations of network telecoms require that the depth of the individual rack unit be severely restricted in order to facilitate placing several rack units within a single telecom housing of a standardized height.[1-20] 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 of maximizing the heat transfer from the chips to the airflow, it becomes necessary to employ the latest tools such as computational fluid dynamics (CFD) and a mixture of new technologies in forced convection cooling.

Numerous innovative technologies offer the potential for increasing the heat transfer from the chip thereby allowing an increase in the permissible heat generation in the chip.[21,22] While these nascent technologies hold the promise of extending the viability of forced convection cooling in the near future, they do so at a price. Some heat transfer gains are made at the expense of increased complexity, weight, cost, space requirement and power consumption coupled with reduced reliability, design flexibility and thermal tolerance. In recognition of the importance of these constraints, the design effort concentrated in the thermal management of the telecom under consideration has been focused on optimizing the utility of forced convection cooling through the use of non-powered heat transfer methods. The enhancements employed were considered on the basis of a zero gain in power requirement,[23] minimal loss in cabinet interior space, the avoidance of liquids in the cabinet,[24-26] and careful cost containment. Adoption of these limitations has forced the emphasis in design optimization to be concentrated on the maximization of the thermal absoption potential of the airflow and on overcoming the heat spreading resistance of the heat sink.

The potential of the airflow to absorb and remove heat is governed by the extremes of temperature in the chip and the airstream. The chips utilized in the new telecom have a maximum case service temperature specified at 70°C. The typical inlet temperature of the cooling airstream is 55°C. This inlet temperature is in accordance with the NEBS (Network Equipment Building Systems) standards governing chassis design. The maximum amount of airflow is constrained by several factors including noise, commercially available fans, board obstructions and good design practices that limit the maximum velocity of the airflow. The thermal resistance of the heat sinks and the pressure drop in the airflow create an upper limit on the forced convection

heat transfer. The latest generation of heat sinks employs a large mass with a complex geometry in the fin design and arrangement. The new board cools the chip with a standard heat sink footprint of dimensions 3.5" × 3.5"× 0.87". In fact, several of the standard footprints must be engaged to achieve the thermal stability of the board system.

To accommodate the ability of the airflow to remove heat, as great a percentage as possible of the airstream needs to be brought into contact with the heat sink surfaces. Additionally, the heat sink must be consistently maintained at a uniform wall temperature. Combining both of the aforementioned conditions dictates that the temperature difference between the heat sink and the airstream be maximized over as large an area as is feasible.

To this end, the design effort consisted of two goals: (1) the optimization of the heat sink with respect to the "evening out" of the spreading resistance, and (2) the control of the airflow to maximize the amount of air impinging on the heat sink surface. The first goal was achieved through an analysis of the material used to construct the heat sink and a study of the flow of heat to the heat sink. The second goal was attained by analyzing the amount and route of the air that bypassed the heat sink surfaces and by correcting for the bypass in the airstream. Working to achieve these stated goals within the given constraints, a system was successfully developed to reach the previously unattainable thermal load of 200 W per board.

II. Description of the Models

In all cases examined, the physical model of the chip and heat sink assembly remained geometrically constant. The differences between the models lay in the materials used to model the heat sink, in the inclusion of the theoretical vapor chamber placed between the chip and heat sink, and in the power requirements for the two chips considered. Aside from the increase in the height of the assembly due to the height contribution of the vapor chamber, the dimensions of the assembly generally remained fixed.

The ducting employed to channel the airflow defines the geometry of the system separate from the chip and heat sink. The vent and fan locations are functions of the design of the managed airflow. The locations of the chips and memory modules drive the ducting design. Clearly, the physical model dictates the reaction of the airflow and heat pathways to the imposed constraints. The discussion of the model begins with an examination of the effect of the material choice on the heat sink. Once material optimization is accomplished, attention will be turned to the addition of a vapor chamber and an optimal array of variable speed fans to effectuate the desired cooling.

A. Physical Description of the System

The new telecom utilizes several processors per board of varying size; two chips were examined for use on the primary board. A medium power processor generating 59.8 W with a maximum case service temperature of 71°C and a high power processor with a maximum case service temperature of 75°C generating 69.7 W were studied. The secondary board utilized a processor with a heat output of 38 W. In both cases, a maximum noise level specification of 48 dBA could be tolerated according to established standards.[27] The memory modules were powered at 10 W total. The board configuration is depicted in Fig. 1.

Each rack unit (RU) is equipped with one line card using both a primary and a secondary board. There are a total of four fans used; two on each of the primary and secondary boards. The tubeaxial fans used measured 40×40×28 mm with the properties given in Table 1. The fans are coupled in pairs to deliver approximately 60 CFM total draw-through airflow at a maximum free delivery of 9,500 RPM. The cooling capacity per board was determined to be 200 W based on an air temperature difference of 10°C. To achieve the requisite cooling, variable speed fans are required.

Table 1 Fan properties

Fan characteristic	value	unit
Static pressure	0.429	in H_2O
Volumetric flowrate	15.8	CFM
Noise	42	dBA
Speed	9,500	RPM

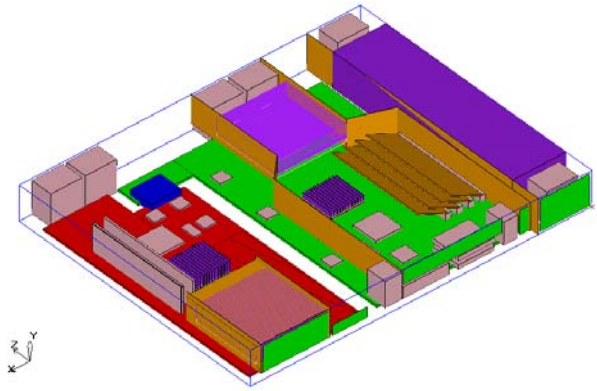

Fig. 1 Isometric view of existing board layout showing the ducting, processor, vent and fan locations.

The board layout separating the primary and secondary boards into individual airflow regions permits each board to experience an influx of ambient temperature air and precludes the necessity of compensating for wake effects. This separation prohibits the warmer air from the higher power chip on the primary board from being funneled over the processor on the secondary board. The four memory sticks on the secondary board are placed alongside the processors out of the warmed airstream. The draw through fan configuration assists in supplying a consistently smooth air flow.

B. Numerical Analysis

For simulation purposes, the mock-up enclosure was constructed with strip heaters to mimic the chip package. The test rig was powered at 38 W and 59.8 W. The tests were conducted with two fans that proved insufficient for the task. The fan RPM was increased to provide the necessary airflow raising the total airflow for both boards to 76 CFM. The new fan specifications are given in Table 2.

A careful examination of the problem leads to the determination that the airflow through the telecom was decelerating at an unacceptable rate due to the number of protuberances in the airstream and the excessive pressure drop created by the variation in the size of the channels. These predominantly rectangular channels are formed by the board, line cards, baffles, chips and telecom case. The variation in the geometry due to the placement and size of the chips and other protuberances generates a cross sectional area that changes rapidly and radically.

Chips and heat sinks are laid out on the boards in such a manner that the severely constricted sections of the channels accelerate the airflow between the protuberances while in other places within the chassis the lack of structures creates empty spaces of large cross sectional area that suddenly slow the airflow and sap energy from it. These sudden decelerations in airflow velocity may be inferred from Fig. 2. It is this slowing effect that negates the cooling efficacy of the airstream by reducing the velocity and therefore its potential to transport large chunks of heated fluid away from the board. This conclusion was verified by numerical simulation of the airflow profile.

To sustain the velocity, baffles (depicted in orange in Fig. 1) were installed in an effort to partition the large, open spaces and redirect the airflow towards the chips and heat sinks. The cooling fans were located at the downstream ends of the channels to pull the air at a constant velocity over the chips. The vent sizes (shown in green in Fig. 1) were designed to permit the largest amount of airflow to enter through maximized inlet

Table 2 Revised fan properties

Fan characteristic	value	unit
Static pressure	0.8	in H_2O
Volumetric flowrate	19	CFM
Noise	45	dBA
Speed	11,300	RPM

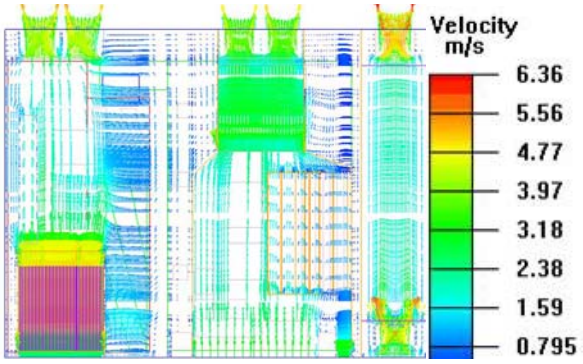

Fig. 2 Top view of the velocity vectors across the board.

vents while regulating the outlet velocity to compensate for the drop in pressure through the chassis by way of the minimized outlet vents. The ducting was added progressively to determine the optimal configuration for capturing the maximum airflow. A critical step was the avoidance of as much air bypass as possible and the prevention of an excessive airside pressure drop exceeding fan capability.

An additional concern was the amount of noise generated by the fans and power supply. In accordance with NEBS standards, the total decibel level must not exceed 48 dBA at room temperature. To achieve the proper design with this acoustical constraint, it was necessary to chart the noise versus speed of the selected fan for the cases of the power supply, the fans at 2.5 W per module alone, and the combined fans and power supply. Also, the speed at various voltages was studied and fan curves drawn to find the optimal point at which the fan speed and temperature difference would meet at less than the maximum noise level.

C. Computational Simulations

Computational simulations allowed for a closely controlled set of conditions. The use of a commercial software package, designed for simulating electronics enclosures, was engaged to model the system. A power supply of a steady 58 W on the chip was specified at varying heat sink temperatures dependent on the air temperature difference.

The initial objective of the simulations was to recreate the results of the experiments and confirm the baseline. This step was followed by the introduction of the variance in heat sink materials and the subsequent inclusion of the vapor chamber in the model. The flow vectors, velocities, and localized board temperatures were computed and displayed.

A. Heat Sink Material Optimization

In all cases, the heat sink assembly under consideration consists of thirty-five 0.015" thick fins attached to a base with the dimensions of 3.5"×3.5"×0.87". The parameter differing between the chosen cases is the material used for each heat sink. For the first case under consideration, the parallel-plate heat sink is made entirely of copper. In the initial trials, heat generation is simulated by means of a steady planar source dissipating 70 watts. The source is placed at the center of the base under the heat sink. Across the inlet section to the heat sink, a uniform velocity of approximately 2.0 m/s is imposed.

For the second case studied, the material is changed to an all aluminum heat sink. All other parameters remain unchanged. The third case focuses on a heat sink with an aluminum base and copper fins. At issue are two important facets of the material selection, the thermal resistance of the material and the cost. The all aluminum heat sink is the less costly alternative and is also lighter. The all copper version has a higher cost although it also has the lower thermal resistance and the higher thermal conductivity.

The third alternative of the composite design offers the lower cost of the aluminum in the base and the superior thermal resistance of the copper in the fins. The aspect ratio of the base thickness to the fin height dictates the choice of material for each section. The objective of the material study is to identify a superior material or composite of materials that provides a minimal thermal resistance and an optimal heat transfer to the airstream.

Our study suggests that the aluminum fin-copper base composite heat sink provides the most cost effective alternative. The lighter weight of the composite coupled with the ability of the copper base to minimize the spreading resistance between the processor and the heat sink fins has helped to maintain the chip temperature at less than the specified 71°C maximum.

B. Ducting Design

The design of the layout for the new boards permits a "scoop" or baffle configuration in the sheet metal ducting that leads from the inlet to the heat sink. By compartmentalizing the two boards, the airflow volume is partitioned to permit a better match of volume to temperature difference than if a single channel is employed.

The use of ducting increases both the volume and the velocity of airflow through the heat sink. Careful attention is given to making certain that the airflow from the inlet is not completely ducted through the heat sinks. A portion of the airflow is allowed to bypass the ducts and flow unobstructed over the remaining sections of the board that do not lie in line with the heat sink paths. The correct duct design is determined by CFD analysis of the chip temperatures. No amount of air is routed through the heat sinks than is necessary to control the chip temperatures properly. The power supply is also segregated from the two boards and is given its own airflow channel.

Figure 3 depicts the board layout with the ducting shown in brown. The purple hatched areas are the processors. The long, rectangular, orange outlined areas are the memory modules.

C. Vapor Chamber Incorporation

Similarly, the placement of a vapor chamber between the chip and heat sink is considered a success as evidenced by the temperature distribution through the heat sink. Comparison to the temperature distribution without the vapor chamber indicates that the chamber permits and additional temperature drop on the processor of approximately 3°C. If the vapor chamber is effective in reducing the spreading resistance, the temperature distribution is expected to display a narrower difference in extremes. In lieu of the parabolic shape of the isotherms characteristic of a chip and heat sink assembly, the inclusion of the vapor chamber must generate isotherms that more closely follow the rectangular shape of the heat sink.[21]

To better facilitate the flow of the heat from the processor to the fins, the copper base of the heat sink

Fig. 3 Board layout with ducting.

incorporated the vapor chamber, thus eliminating one material interface. The vapor chamber accounts for a decrease in the processor temperature of 2-3 degrees over the model as designed without the chamber. This design feature also negates the loss in fin height and fin gap flow area required by placing the vapor chamber between the chip and heat sink.

D. Fan Selection

The optimization of the fan has necessitated a more extensive study to determine the proper fan speed that would provide sufficient cooling airflow and still not exceed the noise level limit. The study focused first on the noise level and compared the voltage and current inputs with the noise produced. This comparison is presented in Table 3.

As the speed of the fan changes with room temperature, the noise is dependent on the temperature of the incoming air. For the purpose of analysis, four speed levels are selected and fan curves drawn relating static pressure to CFM. These results are depicted in Fig. 4.

The results of the matching of the four fan speeds to the inlet air temperatures with respect to the processor temperatures are given in Table 4. In the most restrictive case, the temperatures remain below the maximum service temperature recommended by the manufacturer.

Table 3 Comparison of noise to power input

Voltage	Current	Speed	Noise System and PS	Noise System	Noise PS
V	A	RPM	dBA	dBA	dBA
12	0.69	9017	59.6	58.0	54.8
11	0.61	8296	57.8	56.2	53.6
10	0.55	7565	54.9	53.2	50.3
9	0.48	6850	51.4	49.8	47.2
8	0.42	6010	48.2	47.0	43.3
7	0.36	5156	45.2	44.2	39.8

Table 4 Effect of variable fan speed on temperature control

Case	Speed	T_{in} air	Primary board Processor temperature	Secondary board Processor temperature
	RPM	°C	°C	°C
A	5,100	30	59.3	56.3
B	7,200	37	58.5	56.1
C	9,200	45	62.3	60.2
D	11,300	55	70.0	68

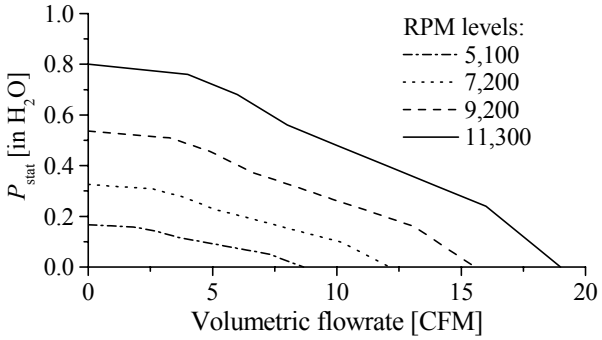

Fig. 4 Fan curves describing the static pressure loss versus air flowrate at four different fan speeds.

E. Numerical Simulations

The CFD simulations were undertaken using a software package dedicated to the analysis of flow and heat transfer in electronics and chip cooling applications. The flow and temperature fields in the computational domain are obtained by solving the three dimensional Navier-Stokes equations along with the energy equation based on a finite volume solver. The finite volume solver used by the code employs multigrid acceleration. The multigrid acceleration serves in reducing the error by iterating in an alternative fashion on a series of coarse and fine grids. The convergence criterion used is based on a tolerance that ensures that the residual error remains less than 10^{-3} in the momentum balance, and less than 10^{-7} in the energy balance. In this study, the vapor chamber was created in a commercial CAD package independent of the code used for the simulations and then imported into the CFD program.

III. Results and Discussion

The robustness of the computational model is tested by comparing its results to experimentally acquired data. A series of comparisons indicates that, for the computational model, an appreciable increase in heat transfer is gained at the expense of a small, tolerable loss in accuracy.

The worst case scenario is that of the primary board equipped with the high power processor operating at a 55°C air inlet temperature. In this instance, the power input to the chip is 69.7 W. Testing and computational predictions verify that this case will experience problems with the memory module temperatures. Also, the primary processor chip temperature exceeds the manufacturer's limitation. Utilization of a variable speed fan enables us to avoid excessive noise while achieving the highest possible cooling.

Use of the vapor chamber incorporated in to the heat sink ensures that the processors and memory modules are thermally safe. The temperature distribution for the 55°C air inlet and 2.5 W per DIMM case is demonstrated in Fig. 5. Coupling the vapor chamber with the careful ducting design, the new telecom board system has made possible the use of a 70 W processor without the employment of refrigerant or liquid cooling.

The inclusion of the vapor chamber in the base of the heat sink combined with the material selection in the components of the heat sink generates a solution that meets the more stringent NEBS requirements. The processor temperature corresponding to the worst case has measured 71.3°C; close to the maximum temperature permissible.

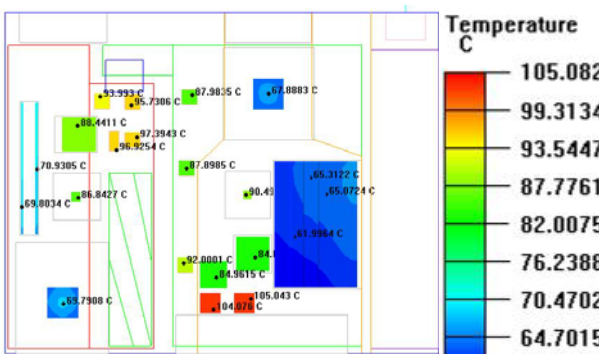

Fig. 5 Key temperatures along the telecom board corresponding to the stringent NEBS specification requiring an inlet air temperature of 55 °C.

IV. Concluding Remarks

The results obtained from numerical simulations appear quite promising. The simplicity of the components used, in conjunction with the associated low costs, have enabled us to extend the viability of forced convection cooling for another generation of telecoms. The implementation of this study is currently awaiting the fabrication of these units.

Acknowledgments

The first two authors gratefully acknowledge the support received from NASA and the Wisconsin Space Grant Consortium. We are especially thankful for the direction and management of R. Aileen Yingst, Sharon D. Brandt, Steven I. Dutch, and Thomas H. Achtor.

References

[1]Andrews, J. A., "Package Thermal Resistance Model Dependency on Equipment Design," *IEEE Transactions on Components Hybrids and Manufacturing Technology,* Vol. 11, No. 4, 1988, pp. 528-537.

[2]Asako, Y., and Faghri, M., "Three-Dimensional Heat Transfer and Fluid Flow Analysis of Arrays of Square Blocks Encountered in Electronic Equipment," *Numerical Heat Transfer,* Vol. 13, 1988, pp. 481-498.

[3]Azar, K., McLeod, R. S., and Caron, R. E., "Narrow Channel Heat Sink for Cooling of High Powered Electronic Components," Eighth IEEE SEMI-THERM Symposium IEEE Service Center 1992.

[4]Bar-Cohen, A., Elperin, T., and Eliasi, R., "Θ_{jc} Characterization of Chip Packages –Justification, Limitations, and Future," *IEEE Transactions on Components Hybrids and Manufacturing Technology,* Vol. 12, No. 4, 1989, pp. 724-731.

[5]Bar-Cohen, A., and Kraus, A. D., *Advances in Thermal Modeling of Electronic Components and Systems*, Vol. 2, ASME Press, New York, NY, 1990.

[6]Bar-Cohen, A., "Air-Cooled Heat Sinks -Trends and Future Directions," *Advances in Electronic Packaging,* Vol. 19, No. 2, 1997.

[7]Butterbaugh, M. A., and Kang, S. S., "Effect of Airflow Bypass on the Performance of Heat Sinks in Electronic Cooling," *Advances in Electronic Packaging,* Vol. 10, No. 2, 1995.

[8]Culham, J. R., Lee, S., and Yovanovich, M. M., "Effect of Common Design Parameters on the Thermal Performance of Microelectronic Equipment. Part II. Forced Convection," *Heat Transfer in Electronic Equipment*, 1991, pp. 55-62.

[9]Culham, J. R., Lemczyk, T. F., Lee, S., and Yovanovich, M. M., "Meta - a Conjugate Heat Transfer Model for Air Cooling of Circuit Boards with Arbitrarily Located Heat Sources," *Heat Transfer in Electronic Equipment*, 1991, pp. 117-126.

[10]Culham, J. R., Lee, S., Jeakins, W. D., and Yovanovich, M. M., "Thermal Simulation of Electronic Systems with Non-Uniform Inlet Velocities," *Computer Aided Design in Electronic Packaging,* Vol. 3, 1992, pp. 33-40.

[11]Culham, J. R., Lee, S., and Yovanovich, M. M., "Conjugate Heat Transfer from a Raised Isothermal Heat Source Attached to a Vertical Board," Ninth IEEE SEMI-THERM Symposium IEEE Service Center 1993.

[12]Culham, J. R., Teertstra, P., and Yovanovich, M. M., "Thermal Wake Effects in Printed Circuit Boards," *Journal of Electronics Manufacturing,* Vol. 9, 1999, pp. 99-106.

[13]Dake, T. J., Majdalani, J., and Narasimhan, S., "Design Optimization of a Router Board Using Computational Fluid Dynamics," The ASME International Electronic Packaging Conference and Exhibition IPACK Paper 2003-35338, July 2003.

[14]Gauche, P., Coetzer, C. B., and Visser, J. A., "Characteristics of Heat Sink Flow Bypass for Thermal Modelling," Proceedings of the Fifth International Conference for Advanced Computational Methods in Heat Transfer 1998.

[15]Gautier, "Construction and Validation of Thermal Models of Packages," Seventh IEEE SEMI-THERM Symposium IEEE Service Center 1991.

[16]Gavali, S., Patankar, S. V., Karki, K. C., and Miura, K., "Effect of Heat Sink on Forced Convection Cooling of Electronic Components: A Numerical Study," *Advances in Electronic Packaging,* Vol. 4, No. 2, 1993.

[17]Goldberg, N., "Narrow Channel Forced Air Heat Sink," *IEEE Transactions on Components Hybrids and Manufacturing Technology,* Vol. 1, 1984, pp. 154-159.

[18]Gopalakrishna, S., "Numerical and Experimental Study of Forced Convection over Power Supply Heat Sinks," ASME Winter Annual Meeting ASME 1991.

[19]Knight, R. W., Goodling, J. S., and Hall, D. J., "Optimal Design of Forced Convection Heat Sinks – Analytical," *ASME Journal of Electronic Packaging,* Vol. 113, 1991, pp. 313-321.

[20]Kusha, B., Rosenblat, S., and Lee, S., Private communication, Thermal Modeling of Heat Sinks, 2000.

[21]Mehl, D., Dussinger, P., and Grubb, K., "Use of Vapor Chambers for Thermal Management," *Proceedings of the Technical Program National Electronic Packaging & Production Conference,* Vol. 3, 1999, pp. 1358-1366.

[22]Toth, J., DeHoff, R., and Grubb, K., "Heat Pipes: The Silent Way to Manage Desktop Thermal Problems," *Thermomechanical Phenomena in Electronic Systems Proceedings of the Intersociety Conference*, 1998, pp. 449-455.

[23]Beitelmal, A. H., Saad, M. A., and Patel, C. D., "Effect of Inclination on the Heat Transfer between a Flat Surface and an Impinging Two-Dimensional Air Jet," *International Journal of Heat and Fluid Flow,* Vol. 21, 2000, pp. 156-163.

[24]Garner, S. D., "Heat Pipes for Electronics Cooling Applications," *Electronics Cooling*, 1996.

[25]Whitaker, S., "Forced Convection Heat Transfer Correlations for Flow in Pipes, Past Flat Plates, Single Cylinders, Single Spheres, and for Flow in Packed Beds and Tube Bundles," *AIChE Journal,* Vol. 18, 1972, pp. 361-371.

[26]Zuo, Z. J., North, M. T., and Ray, L., "Combined Pulsating and Capillary Heat Pipe Mechanism for Cooling of High Heat Flux Electronics,", Vol. HTD 364, American Society of Mechanical Engineers, New York, 1999, pp. 237-243.

[27]Peterson, A. P. G., and Ervin E. Gross, J., *Handbook of Noise Measurement*, 5th ed., General Radio Co., West Concord, MA, 1963.