CFD and EDA tools
The Interoperability of FLOTHERM® and Board Station®/AutoTherm®: Concurrent Design of a Freescale PowerPC™ RISC Microprocessor-based Microcomputer

Gary Kromann (1), Vincent Pimont (2) and Steve Addison (3)

(1) FREESCALT
Advanced Packaging Technology
3501 Ed Bluestein Blvd.
mailstop: TX11, K1
Austin, Texas, USA 78721
Email: R14103@email.sps.mot.com
TEL: (512) 933-3565
FAX: (512) 933-6344

(2) Mentor Graphics
Systems on Board Division
8005 SW Boeckman Road
Wilsonville, Oregon, USA 97070-7777
Email : vincent_pimont@mentorg.com
TEL : (503) 685 1015
FAX : (503) 685 7991

(3) Flomerics Inc.
Western Region
2880 Zanker Road
Suite 203
San Jose, CA, USA 95134
Tel: (408) 954 7332
Fax: (408) 954 7335
Email: steve@ca.flomerics.com

Summary

This paper discusses an attempt to bring thermal analysis early in the printed-circuit board design process, when designing Motorola's PowerPC 603™ and PowerPC 604™ microprocessor-based desktop system. The goal was to assess a methodology that should help to define a real concurrent design process for future projects. We emphasize here the thermal aspects of this concurrent process that required the use of a board-level (AutoTherm from Mentor Graphics) and system-level thermal analysis tool (FLOTHERM from Flomerics).

After describing the project, and the dataflow currently available between AutoTherm and FLOTHERM, we describe the practical steps that were carried out in this project, and how thermal design has finally been included as one of the constraint during the component placement phase on the printed-circuit board design.

Overall the experience gained through this project on multi-level thermal analysis, as well as, working in a cross-functional team environment is presented. Also presented are the steps for implementing such a concurrent design flow.

© Freescale Semiconductor, Inc., 2004. All rights reserved.

freescale™ semiconductor

INTRODUCTION

Computer system performance has dramatically improved over the past three decades. Much of this improvement is a result of the increased integration of components at the semiconductor level made possible by reduced feature sizes. The level of integration now possible at the chip level has resulted in several semiconductor integrated-circuit (IC) trends, all of which are increasing: gate count, chip inputs/outputs, chip size, operating frequency, and power consumption. All these trends have resulted in increasing thermal flux at the chip level. Furthermore at the package level, there is a continuing trend of packaging chips in an ever decreasing footprint and volume. Together with the increasing use of surface-mount technology and expanding use of "chip scale" packages, these trends have resulted in higher power dissipation at the board-level. On the system-level, as seen with recent development of desktop, laptop and palmtop systems, more features are being placed into a shrinking enclosure with less and less air and space available for designing a thermal management scheme. Thermal control of microelectronic devices is required for proper operation and acceptable reliability and is becoming an increasingly critical part of the design of microelectronic systems [Lasance,1995].

Today's electronic systems require a very diverse set of requirements that must be met by the electrical, thermal, mechanical, and packaging engineers. Various computer simulation tools are to be used to predict the physics of these class of problems. To reduce the design- and analysis-cycle time, the interoperability of such computer simulation tools offer a potential reduction in the product development cycle times, while reducing potential pre-processing errors. Here we will discuss the interoperability between an electronic-design automation (EDA) tool and an computational-fluid dynamics (CFD) tool to complete the thermal analysis of a microcomputer system. This paper discusses the concurrent design process methodology within the printed-circuit board (PCB) layout and the thermal system simulation for a desktop microcomputer. More specifically we will describe data flow and the model simulation results within the Mentor Graphics' Board Station/AutoTherm and Flomerics' FLOTHERM environment for a PowerPC 603 and PowerPC 604 microprocessor-based desktop system.

BRINGING THERMAL CONSTRAINTS INTO THE PRINTED-CIRCUIT BOARD

Assessing Concurrent Design and Analysis

For any electronic system design, an electrical designer generates a set of schematics. From these schematics, layout designer creates a physical design, and then faces the challenge of meeting many potential contradictory constraints, such as: component selection, electrical performance, thermal performance, manufacturability, cost, geometry, etc. [Seaton, 1996]. At this point, components are selected from company libraries, and the PCB layout database can be generated, giving the basis for the layout and routing of the printed-circuit board (PCB).

Today the success of the project depends upon the entire team working together; that is, the electrical, mechanical, and packaging engineers working hand-in-hand with the layout team to translate and integrate the primary physical constraints of the design. At this point, rules-of-thumb may be efficiently replaced by more sophisticated numerical thermal analysis of the "virtual physical designs" to reveal to the entire team the quality and the flaws of the on-going design. These computer simulation tools take into account many complex boundary conditions, that neither intuition nor global metrics like the junction-to-ambient thermal resistance can do [Andrews, 1988], [Lin et al, 1993], [Rosten, 1996]. This process is what was attempted in the Motorola's YellowKnife project described below, where electrical and packaging engineers worked together with the layout designers to decide upon the best possible trade-offs for the motherboard. We will focus in this paper on the thermal aspects only, knowing that the key role of the mechanical packaging engineer includes the solutions of such problems as : thermal control (e.g. heat sink, heat sink attach, air mover and location geometry), electrical cabling, connectors, enclosure selection, mechanical shock and vibration.

Advantages of Concurrent Design and Analysis

With thermal analysis performed during the design cycle, the microprocessor-based system design is expected to get a better chance of being designed right the first time through. It is then important to build good thermal management into a product design at the earliest possible stage, or else you run the risk of thermal problems later delaying the entire project. Using thermal analysis tools such as AutoTherm and FLOTHERM, thermal engineers can ensure that system's thermal requirements are being adequately addressed before the PCB is fabricated. Some other key advantages are:

- Reduce Product Development Cycle time by performing the analysis during the design phase, not after.
- Improved Thermal Prediction Techniques.
- Using a EDA database does not require the thermal analyst to pre-process component geometries, powers and locations; thereby, improve the accuracy of the model and reduce cycle time.
- Numerically parametric studies of thermal enhancement features, from PCB conductivity to fan performance.
- Board-level simulation; interaction, thermal loading of adjacent components.
- System-level simulation can show the impact of card placement, fan location.

Thus, in the long run, up front thermal analysis will produce higher reliability designs that don't require the expense and time needed for an extra PCB turn to get thermal issues resolved. This results in shorter design cycles without compromising product quality which speed time-to-market while reducing product costs.

Linking Printed-circuit Board-Level and System-Level Thermal Analysis

AutoTherm is a typical PCB level analysis tool provided by an EDA company. While giving the flexibility of performing numerous what-ifs analyses, impacting the component placement and the PCB structure, according to basic boundary conditions around a single PCB, it is integrated within Mentor Graphics' Board Station design environment. Any change made in AutoTherm to placement or properties, as well as additional thermal data may be read directly back into the current layout design. AutoTherm enables engineers to analyze all the details of a PCB, including the thermal behavior of capacitors and resistors. On-line tools of AutoTherm can help to create the component models

With the shrinking of typical PCB enclosures, the air flow around the PCB and components is more likely to be influenced by the system level geometry. AutoTherm provides basic flow capabilities, based on a perfect fluid approach, which can be used for initial estimates before the enclosure has been fully specified. But, as more accurate information on the system becomes available, it is important to take this into account within the EDA environment where it may impact the PCB layout.

Therefore, Flomerics and Mentor Graphics have developed an interface between AutoTherm and FLOTHERM, allowing to bring FLOTHERM's system-level examination of thermal into the EDA environment.

The dataflow between the applications is presented Figure 1. In addition to the 2D data exchange possible through FLOGATE, tridimensional translation of AutoTherm models into FLOTHERM is now available, allowing a deeper integration for critical designs. In this latter case, the PCB and selected components (i.e. significant for the power they dissipate and/or for their geometry) are automatically extruded as cuboids in the enclosure. As an option, the cuboids representing the components can support AutoTherm's two resistors component model, instead of uniform conductivity cuboids. A flow analyzer "FLOMENTOR" has been developed to extract local heat transfer conditions from the 3D velocity and temperature fields around the PCB and the components. AutoTherm will then use this local heat transfer data to perform a detailed analysis of the entire PCB.

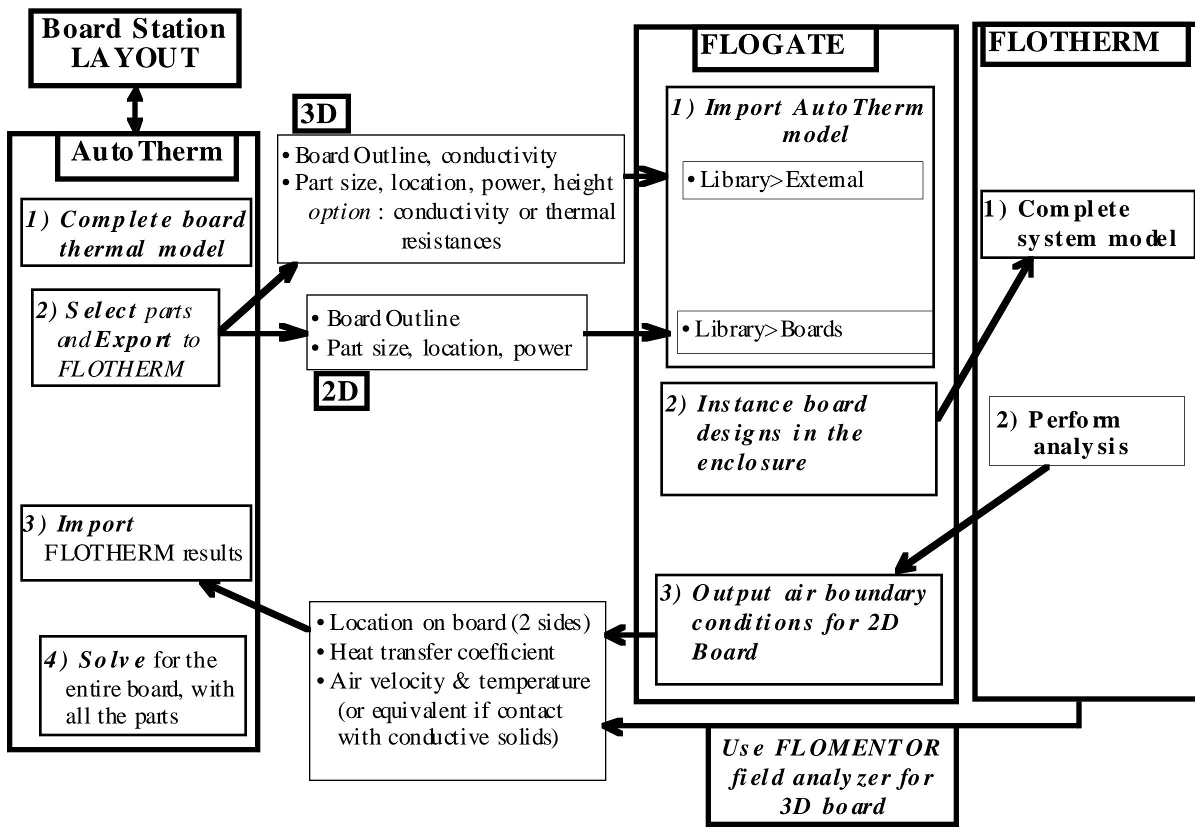

Figure 1 - Dataflow between Mentor's PCB design environment and FLOTHERM

MOTOROLA'S YELLOWKNIFE CASE STUDY

PowerPC™ RISC Microprocessor Family

The scaleable PowerPC™ Reduced-Instruction-Set-Computer (RISC) architecture microprocessor family jointly developed by Apple, IBM, and Freescale, is being designed into high-performance cost-effective computers (including notebooks, desktops, workstations, and servers). The PowerPC microprocessor family includes: the PowerPC 601™, PowerPC 602™, PowerPC 603™, PowerPC 603e™, PowerPC 604™, PowerPC 604e™, and the PowerPC 620™ microprocessor [Kromann et al, 1995], [Gerke et al, 1995], [Kromann, 1996] [Freescale related documentation cited in references]. Each microprocessor is designed to meet the needs of a different segment of the marketplace.

PowerPC Architecture reference system Microcomputer Hardware

The microcomputer system investigated in this study is a PowerPC Common Hardware Reference Platform (CHRP) compliant reference system (code named "Yellowknife"). The major system features are given appendix A. The Yellowknife design is a single processor system that accepts either the PowerPC 603 or the PowerPC 604 microprocessor. Yellowknife supports all 2.5V and 3.3V PowerPC 603 or the PowerPC 604 microprocessor operating in modes which result in external processor bus speeds up to 66 MHz. The maximum power dissipation of the PowerPC 604 microprocessor is 22 watts [Freescale, 1995a]. Voltage regulator will be put on PCB to provide different voltage for different processor. There is one "socket 3" PGA zero-insertion-force (ZIF) socket on PCB. Socket 3 is a 17 X 17 fully-populated footprint. A ball-grid-array(BGA)-to-pin-grid-array (PGA) interposer will be used to convert the BGA footprint to the PGA footprint.

The Yellowknife Chassis uses the Intel's ATX form-factor chassis described Figure 2. The Yellowknife chassis supports a total of five I/O slots for add-in cards. Three of the slots support PCI cards and the remaining two support ISA cards. The memory subsystem comprises in particular a DRAM expandable 64-bit memory bank , allowing to support up to 132 MB. The system is made of up to 4 SIMM cards (Single-In-line Memory Module) each of them dissipating approximately 5 watts.

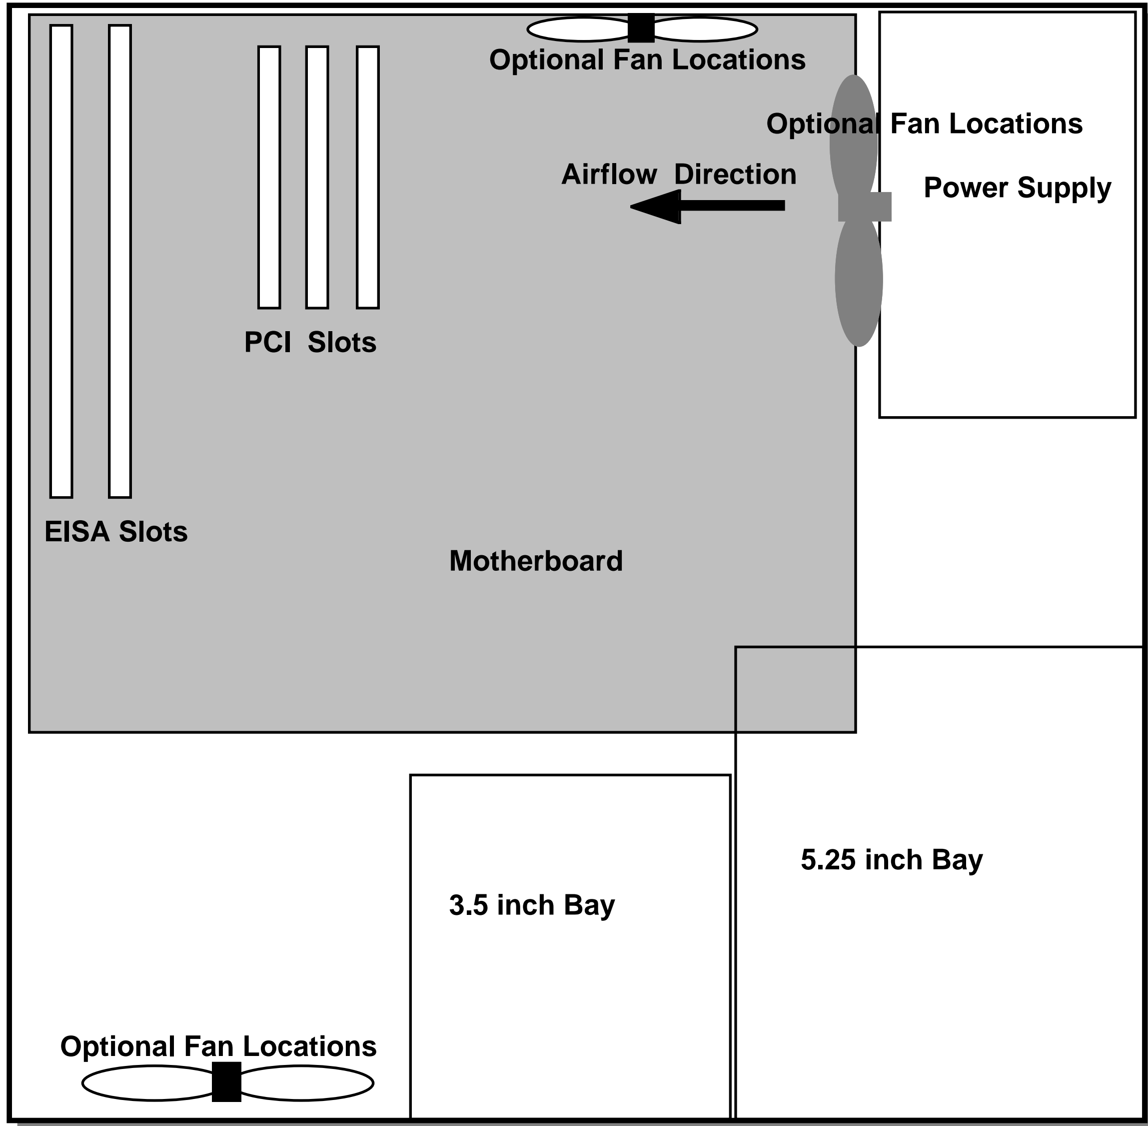

Figure 2 - Salient Features of the ATX chassis (top view)

The Yellowknife chassis Drive Bays has two external 5.25" drive bays and two internal 3.5" drive bays. This combination of available drive bays allows systems to be built in configurations such as the following that will be used for the initial systems assembled: 1) one IDE hard drive, 2) one SCSI hard drive, 3) one 3.5" floppy drive, and 4) one CD-ROM drive. Other configurations are clearly possible, including support for an external tape drive and/or a third hard drive.

PRINTED-CIRCUIT BOARD-LEVEL THERMAL SIMULATION

A first packaging of the design, taking into account some timing and electrical constraints, was made available by the layout team to the packaging engineer. The translation from Board Station LAYOUT into AutoTherm is being automated for the geometrical description of the PCB. The automation of the PCB thermal model creation depends on the status of the company components libraries that may (or may not) contains the adequate component thermal models and properties. For those components that have no complete information, the first step was to gather and check the necessary information, in particular :

1) Verify the component package geometry
2) Identify the power consumption of each device (core plus I/O)
3) Identify mechanical geometry interference issue (e.g. heat sinking interference with option cards)
4) Option cards: cache, memory, LISA, CPI,
5) Possible mass storage units: disc and tape drive units

A preliminary component placement must be determined that meets the design requirements of a cross-functional design team, concurrent with the product requirements. Then the next step of the thermal analyst is to complete the board-level and system-level thermal simulation before PCB routing. That is taking the initial component placement from Mentor Graphics' Board Station to AutoTherm, then to FLOTHERM and back to AutoTherm. The component libraries within the EDA tool, include the component: geometry, location, and power dissipation (maximum and/or typical) The populated mother-board and add-in cards will be placed. Upon completion of initial thermal analysis, any component placement changes can be modeled to improve the thermal performance. Next, the PCB board conductivity may be increased locally or globally (e.g., thermal vias, adding planes, adding additional metal in a signal plane etc.). Finally, the thermal enhancements such as heat sink sizing and fan sizing might be considered.

Preliminary Motherboard Thermal Analysis

This section presents the key steps of the analysis interations completed on the main motherboard only, from the initial placement to the final preliminary placement. In this preliminary thermal analysis was done very early in the design stage prior to the selection of the system enclosure. The key objectives and strategies are to :

- Identify key components that grossly exceed the manufactures recommended operating temperatures. These components are identified as requiring additional heat sinking solutions; however, its also very illustrative to the electrical engineers. In addition, one might recheck the power consumption and the component package thermal resistance data contained in the library database. Next the thermal analyst can work with the PCB layout designer to generates features that can be activated within Board Station to allow for geometry constraints. Finally, as needed move components to locations where spacing will not result in a geometric interference.
- Identify key components that are much lower than the manufactures recommended operating temperatures. To reduce the solution solver times, these components may be deleted.
- Identify key components that are slightly higher than the manufactures recommended operating temperatures. Here we might consider simple local thermal enhancements such as, spreading components out, rearranging components, improving board-level conduction coupling (e.g., adding thermal vias, underfill compounds) and/or improving the PCB conductivity (locally and/or globally).

In the early stages of the Yellowknife program, the system enclosure had not yet been identified. Therefore, an enclosure was modeled within Autotherm with an inlet and an outlet of approximately 1 m/s airflow. At this stage adding heat sinks was not considered. A variety of analysis were completed and the key results were :

1) In concert with the electrical engineering team an initial placement was considered for the key components. Then PCB placement began.

2) This first simulation run exposed the components that grossly exceed the manufactures recommended operating temperatures. For instance, as not heat sinks were considered, the microprocessor showed a die-junction temperature of approximately three times the manufactures specification; while, two other components exceeded die-junction temperature of approximately two times the manufactures specification These components are identified as requiring additional heat sinking solutions and/or relocation candidates (however, relocation decisions were made as a team, as there may be other electrical issues (e.g. timing, signal intregity)). In addition, its also very illustrative to the design team that without a doubt, these components will require heat sinking solutions.

3) From the initial component placement, the following were the key components that were identified as potentially requiring thermal enhancements: microprocessor, 4 FETS, SCSI controller, an ASIC, Cache Tag and Clock Driver (Figure 3). The following is a summary of the components simulation run this were the steps taken.

Microprocessor : will require an active or passive heat and lets try to place in an area that will not get shadowed. Also, within the layout tool geometry constraints were activated in an area around the CPU.

FETs : the four FETS initially were placed very close together in a quad arrangement. A second and third simulation showed that by spreading these components out or placing them in an in-line arrangement would reduce the temperature rise by approximately 20%. In addition, another simulation run showed the local enhancement of PCB conductivity could further reduce the die-junction temperatures by approximately 5 to 10%. Therefore, here were greatly reduced a potential thermal problem by spreading the components out and improving the local PCB conductivity.

Cache Tag : its location was critically placed near the microprocessor, for routing and signal integrity/crosswalk issues. However, it was place far enough away from the microprocessor to add a heat sink as needed.

ASIC : Initial analysis showed this component may require a heat sink. Originally this was placed between the ISA and PCI slots, so it was moved to in front of the PCI slot where spacing was indented for heat sinking and better unobstructed airflow.

SCSI controller : same strategy as the ASIC component.

Clock Driver was shown to have an die-junction temperature rise that exceeded the manufactures specifications by approximately 30%. Therefore, the clock driver chip as moved away from the PCI-controller chip and the Cache Tag chip

4) After all components were placed (figure 3), another simulation was conducted on all the components. As expected many of the smaller components (e.g. passives, inductors, capacitors, resistors) may be deleted from the database for the thermal simulation, which greatly improves solution run times.

5) At this point the PCB will be imported to FLOTHERM for system-level simulation.

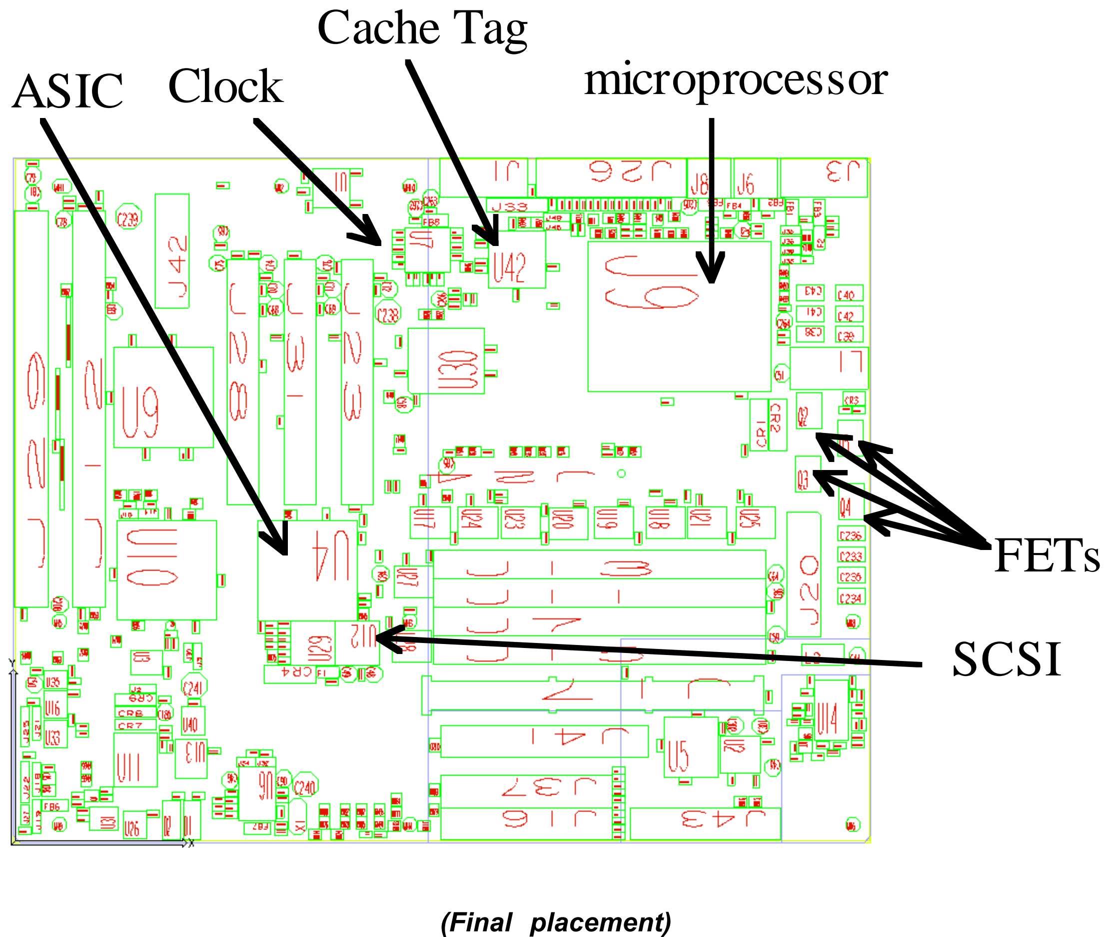

(Final placement)

Figure 3 - Fully Populated PCB in AutoTherm after transfer from LAYOUT

System Level Modeling in FLOTHERM

ATX Chassis Overview

The FLOTHERM model of the chassis is composed of the following main elements:

- the Outer Case and Internal Structure;
- the Power Supply Unit (PSU);
- the PSU Fan; and
- the Disk Drives & Ancillaries.

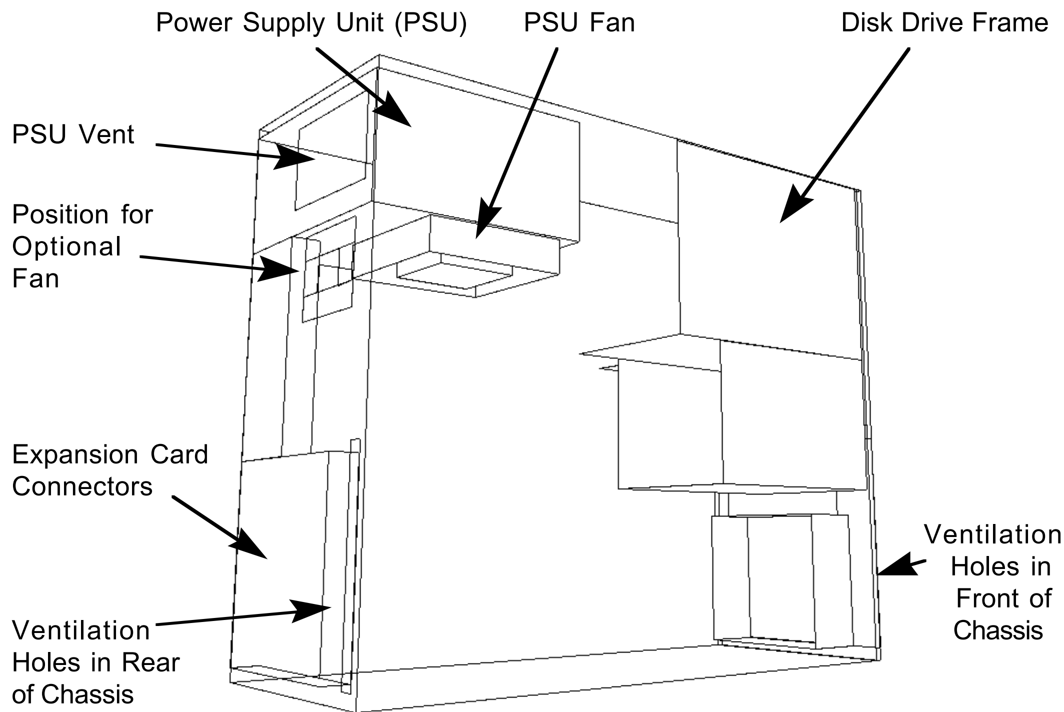

Figure 4 - FLOTHERM Model of the System Chassis

Outer Case and Internal Structure

The outer case of the computer consists of a series of plastic sheets lined with thin metal plates. In this arrangement there is a thin layer of air trapped between the metal and the plastic cutting and, therefore, the case is almost adiabatic. It is modeled in FLOTHERM using *External Walls* with a conductivity derived from the thermal resistances of the metal sheets, the air gap and the plastic wall in series (Figure 5).

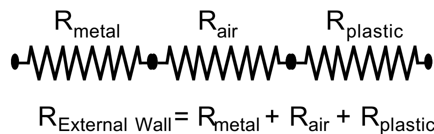

$$R_{External\ Wall} = R_{metal} + R_{air} + R_{plastic}$$

Figure 5 - Thermal Resistance of the Case Walls

In this simulation, radiation from the outside of the case has been omitted in order to arrive at a conservative answer. The perforations and slots in the case are modeled as *Vents* with the Loss Factor set to 1.7 (based on Device - [Flomerics,1], [Flomerics,2]). This estimate of loss factor is expected to be adequate for a first analysis although, should the thermal design prove to be marginal, a more detailed assessment of the pressure loss through the vents would be carried out.

The internal metal structure used to support the disk drives and ancillaries (for example, tape drive and CDROM) is modeled using *Internal Plates*. In doing this, we are assuming that the heat conduction within the fabric of the structure is not significant and that the heat dissipated by the disk drives and ancillaries is convected away in the cooling flow. This is, again, expected to give a conservative answer for the motherboard thermal analysis. It should, however, be noted that for some heavily equipped configurations (typically PC's acting as servers), the cooling of the drives

and the drive bays may in fact be a more pressing problem than the motherboard and processor(s).

Power Supply Unit (PSU)

The modeling of the PSU and its associated fan (see next section) presents the most difficult part of creating and validating this chassis model. There are a number of factors which need to be addressed:

1. the pressure drop characteristics of the PSU;
2. the effect of component proximity on the fan characteristic; and
3. power dissipation of the PSU and PSU fan.

The main problem for a designer addressing these problems are practical issues because he/she has little control over which PSU is used. In particular:

- manufacturing departments may change the installed unit to save cost;
- the PSU manufacturer might change the internal layout of components; or
- the end user may install his/her own PSU as a replacement part.

And, whilst the mechanical form factor and electrical specification may be the same, there is usually no attention paid to thermal specification. With the PSU being an integral part of the thermal design of this type of chassis, thermo/fluid compatibility becomes increasingly important.

There are two approaches that can be taken:

1) At the preliminary design stage, the designer has little choice but to use correlations based on the expected component packing density and test for the sensitivity of the design to the assumed values. Correlations that might be useful can be found in texts such as [Perry & Chiltern], [Fried & Idelchick], [Flomerics, 3].

2) Once a PSU has been specified (or if a module of similar configuration is available), a bench test can be carried out to determine the flow rate for the unit under conditions with no back pressure. The PSU manufacturer may also be able to supply this data. Unfortunately, this cannot simply be assumed to be the flow rate through the system since the unit will also see a back pressure due to other elements in the system. The procedure for creating a compact model of the PSU from this data is described in more detail in Appendix B.

In the analysis reported in this paper where we have no data on the performance of the PSU, we have chosen to apply a pressure loss coefficient of $40m^{-1}$ (i.e. 40 dynamic heads per meter of length) which is based on a rough assessment of the likely component packing density of the PSU. When assessing the final results, considerable attention would be paid to the effect of varying this parameter within practical limits and this might result in design constraints being placed on the PSU vendor.

PSU Fan

The physical structure of the PSU fan was modeled using a combination of *Internal Plates* and a *Cuboid Block* to create an annular duct. The fan characteristic is taken as a straight line approximation to the supplied fan curve in the region of the operating point. This level of detail is more than adequate for a fan drawing configuration but might need additional detail for the fan blowing configurations (see for example reference [Flomerics, 4]).

Assessing the effect of swirl in blowing configuration is a further difficulty. Modeling swirl in FLOTHERM is not very difficult, simply involving the use of a number of *Volume Sources*. However, there are considerable problems in obtaining the data from fan manufacturers in the first place. This parameter is rarely (if ever) measured and, furthermore, varies with the operating point of the fan. Some background information on modeling fans with swirl is contained in Appendix C.

Disk Drives & Ancillaries

For the designer of a chassis such as this, the modeling of disk drives and similar ancillaries (for example, CDROM and tape drives) represents a further complication. Once again, the issue is one of control. Although the initial design may assume a certain configuration, the end-user of the system (particularly for SOHO systems) may choose a totally different configuration and/or the drive manufacturer may change the specification.

Fortunately, the coupling between the thermal behavior of the drives and the motherboard is far weaker than the coupling between the motherboard and the PSU. Therefore, in the absence of better data, modeling the ancillaries as *Cuboid Blocks* of appropriate conductivity with a pessimistic estimate of power dissipation represents an acceptable solution for an engineer struggling to meet design deadlines. In the analysis described here, blocks of conductivity 50 W/mK have been used. The value of 50, although seemingly arbitrary, is a reasonable guess given the high metal content of a typical disk drive assembly. We are also assuming a worst case scenario where little heat is conducted into the fabric of the chassis (through the plates). Under these circumstances, the air side heat transfer coefficient dominates and the actual value of conductivity chosen is secondary.

Motherboard

After simplifying the motherboard in AutoTherm, it was translated into FLOTHERM as a 3D assembly of cuboids. During the first analyses, of the remaining components, the components dissipating at least 1 W were sent to FLOTHERM. In that case, the remaining power is distributed in the PCB, ensuring that the overall heat load is consistent between the two models. At this point, the components were transferred with the default conductivity of 10 W/mK, while the PCB had a conductivity of around 19 W/mK.

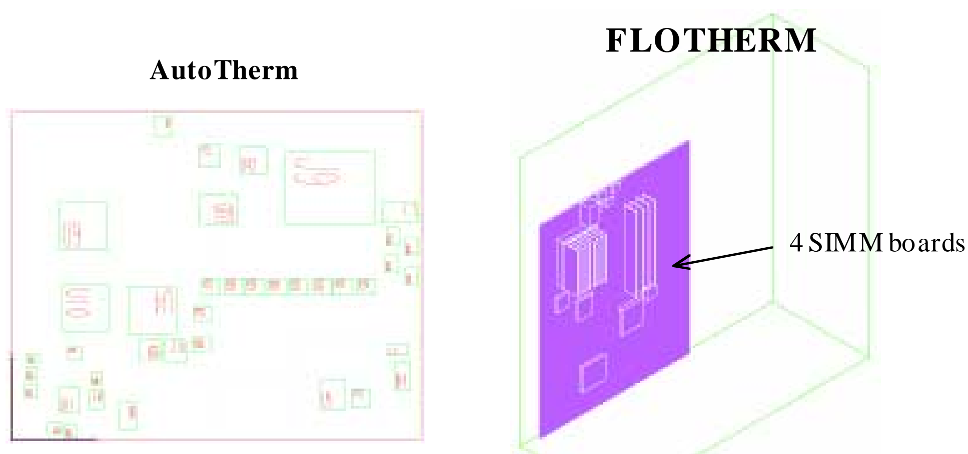

Figure 6 - Simplification in AutoTherm and transfer of the PCB and selected componen into FLOTHERM

(The SIMM boards as well as the heat sink on the processor have been added)

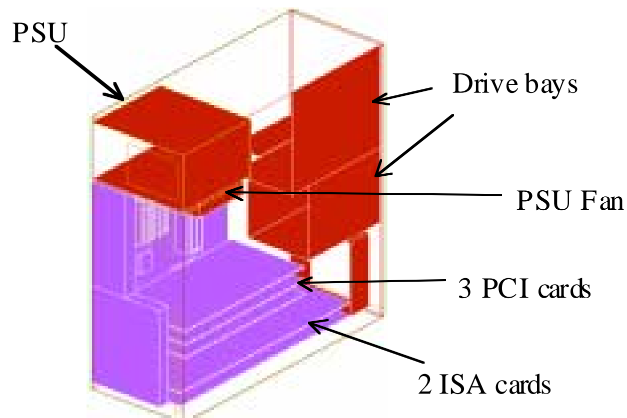

Figure 7 - Final system in FLOTHERM

Using AutoTherm and FLOTHERM Together

While performing AutoTherm board analyses using system-level FLOTHERM results, the question of the consistency of the models between AutoTherm and FLOTHERM came up several time. Being in an on-going design process on the board side, we have to use the same boundary conditions from FLOTHERM for several options for the board. In that case, it may be important to know when it is necessary to go back to FLOTHERM with an updated board models, and which risks of important inconsistencies may be taken during this "creative" step of the board design process. Note, the following results were obtained with the fan sucking the air from the enclosure through the power supply unit (PSU).

Technical aspects of the 3D AutoTherm-FLOTHERM interface

When using the FLOMENTOR 3D field analyzer around a 3D PCB that was designed in AutoTherm, the heat flux going out a surface element of the board or of a component is used to compute a heat transfer coefficient value h, according to the ambient air local temperature. The values computed around a component will be then combined in a single h value that reflects the local 3D flow effects around the component.

This process takes into account the possibility of having external blocks on the PCB, in contact with the board or the any other components. Therefore, the full modeling of a heat sink on the processor with conductive blocks, as well as, the conduction heat flux coming from the SIMMs into the board, may be modeled in this manner. Also with this modeling feature, any flux through any additional blocks put on the board is taken into account. We'll make use of this extended capability below.

Blocking effect of the inductor at the exit of the heatsink

In a first step, a inductor close to the processor and that dissipates no power, was not transferred to FLOTHERM. The figure 8 shows that the inductor affects the flow at the exit of the heat sink, and creates a recirculation zone. Despite the relatively small blockage (the inductor is 10 mm high

while the fins’ base is approximately 6 mm), there is influence on the flow field at the base of the heat sink, with a potential impact on the heatsink performance.

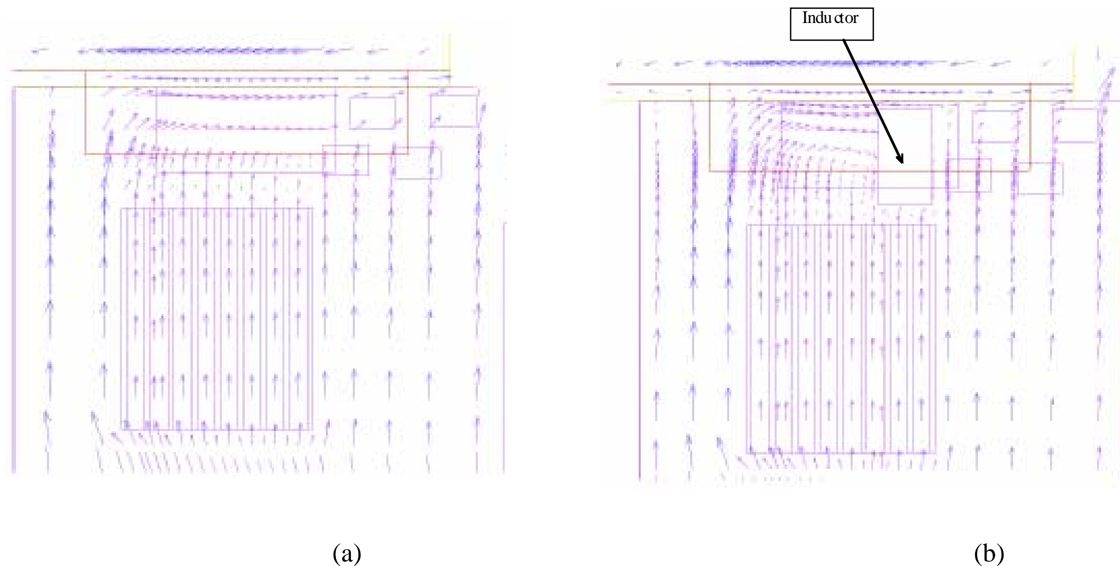

Figure 8 - Flow field in the heatsink : (a) no inductor ; (b) inductor

Surprisingly, the temperature results remain quite close, within 2%. One of the reason may be that the inductor is acting as a fin for a region of the board that is pretty hot, phenomenon enhanced due the default conductivity of 10 w/mK for all the components.

Taking into account local enhancement of the mother board thermal conductivity

The present AutoTherm-FLOTHERM interface transfers an average board conductivity from AutoTherm to FLOTHERM. We then investigated the significance in modeling the local board conductivity enhancement resulting from the addition of a copper pad in the region of the FETs, both in AutoTherm and FLOTHERM. It appears that the average junction temperature drop computed in AutoTherm for the FETs was differing from less than 1.5% (relative to the average temperature rise above the ambient) when using FLOTHERM results obtained with the original mother board or with the enhanced board.

But the major effect of the copper pad is missed when it is not modeled into FLOTHERM. The temperature maps around the FETs are similar when using the original FLOTHERM results (no copper pad) : the diffusion effect of the copper pad in AutoTherm is finally somewhat overriden by the convective heat flux information coming from FLOTHERM which does not take into account the modified heat flux distribution in this region of the board, that has also the highest heat load.

When modeling the copper pad also in FLOTHERM, the AutoTherm simulation shows that the resulting diffusion of the heat from the hottest parts (close to the processor) reduces the maximum temperature difference between the FETs, from more than 10 C in the initial model to less than 6.3 C. As a consequence, the coldest FET sees its temperature rising a little, but overall, the temperature level for the FETs has been improved.

Impact of the SIMMs power dissipation on the mother board

The first analyses performed on the mother board were made with an estimate power dissipated on the SIMMs of 1 w. The final power appears to be around 5 w per SIMM. The four SIMMs are modeled as dissipating blocks (1 block for the FR4 substrate and 1 block representing the 8 SOJs mounted on the substrate) that are in contact with the mother board. Thanks to the ability of the interface to take into account the conduction heat flux coming from a block into the board or any

other component, it was possible to translate this information into AutoTherm. That is, the conduction heat flux coming from the block representing the SIMMs that takes into account the convection ocurring around them (figure 9) . The impact is in the immediate vicinity of the SIMMs. The processor die-junction temperature remains in the same range (110 C), while the FETs are now all approximately the same, yet with a slight increase in temperature to 100 ∞C.

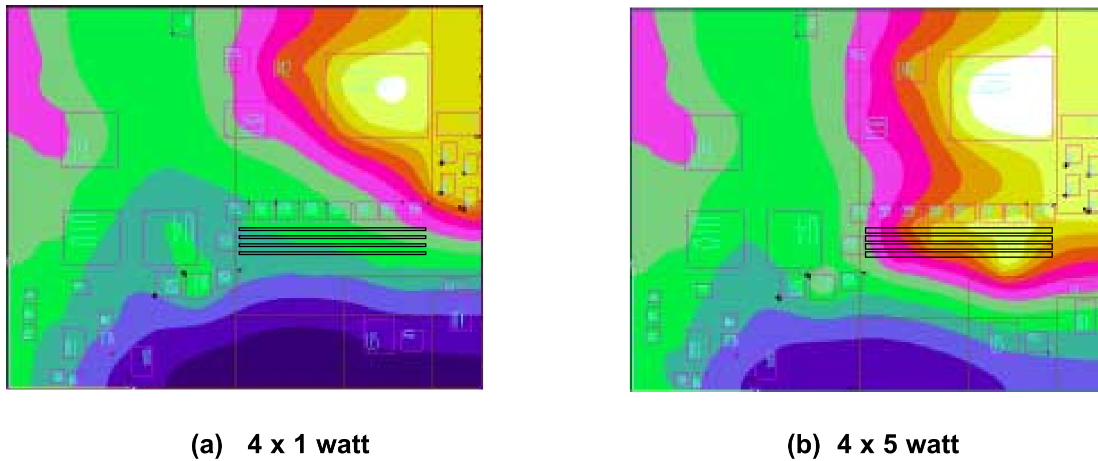

(a) 4 x 1 watt **(b) 4 x 5 watt**

Figure 9 - Influence of the power of the SIMMs on the mother board.

The major impact is on the components that are the closest to the SIMMs (just above them in figure 9). They see their die-junction temperature rising from a range of 54 to 64 C to a range of 65 - 92 C. Thus, this clearly demonstrates the local component interaction effects on an otherwise low-power component. That is, here we see these “low power” components (around 17 mW) being “heated” by their “hot” neighbors.

Detailed System Level Results

Although the final design was not complete at the time of writing and, furthermore, many of the design parameters are pre-defined by the chassis specification [INTEL, 1996], some of the initial results are valuable in identifying design features which could be changed to optimize the thermal performance. Indeed, the earlier that this kind of analysis can be done in the design cycle, the more value can be derived.

Two sample runs for the system model are shown here. In the first, the PSU fan is blowing air across the processor site (as in the ATX specification). The peak temperature in the processor is about 96.5∞C. But it is also notable that the temperature in the SIMMS modules is very high and this is a cause for concern.

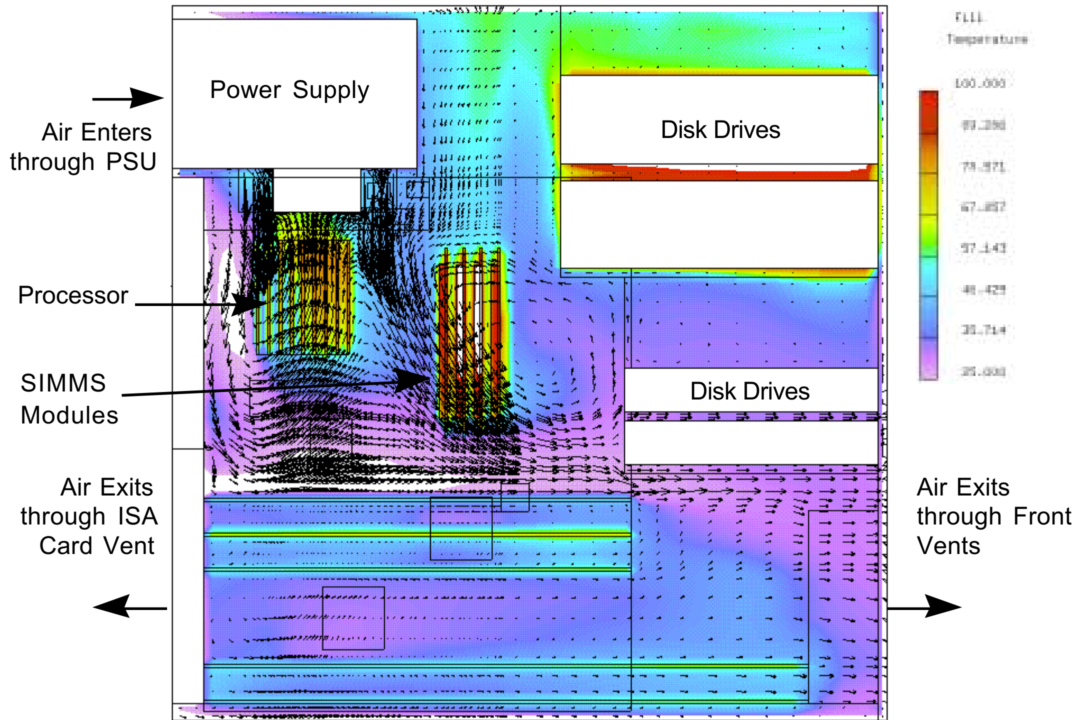

Figure 10 - Results for the Fan Blowing Configuration
Temperature Fill on a plane through the Processor Heatsink
Velocity Vectors on a plane through center of the PSU Fan

Figure 11 shows that the cooling of the processor heatsink is heavily influenced by the spreading of the air from the PSU fan. This tells us that:

∑ Firstly, modeling the swirl from the fan is important in this case since it influences the spreading through centrifugal action; and

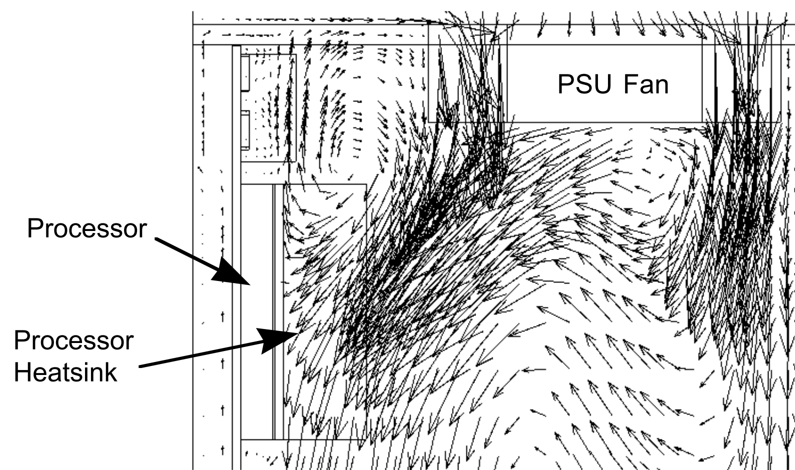

Figure 11 - Results for the Fan Blowing Configuration
Close-up of Flows Through the Processor Heatsink

- Secondly, we would be wise to allow a considerable margin in the design since the swirl from a fan is very poorly defined quantity and may vary considerably from fan to fan.

Indeed, it might also be wise to look more closely at the effect of using longer fins on the processor heatsink to bring it further into the main cooling airstream.

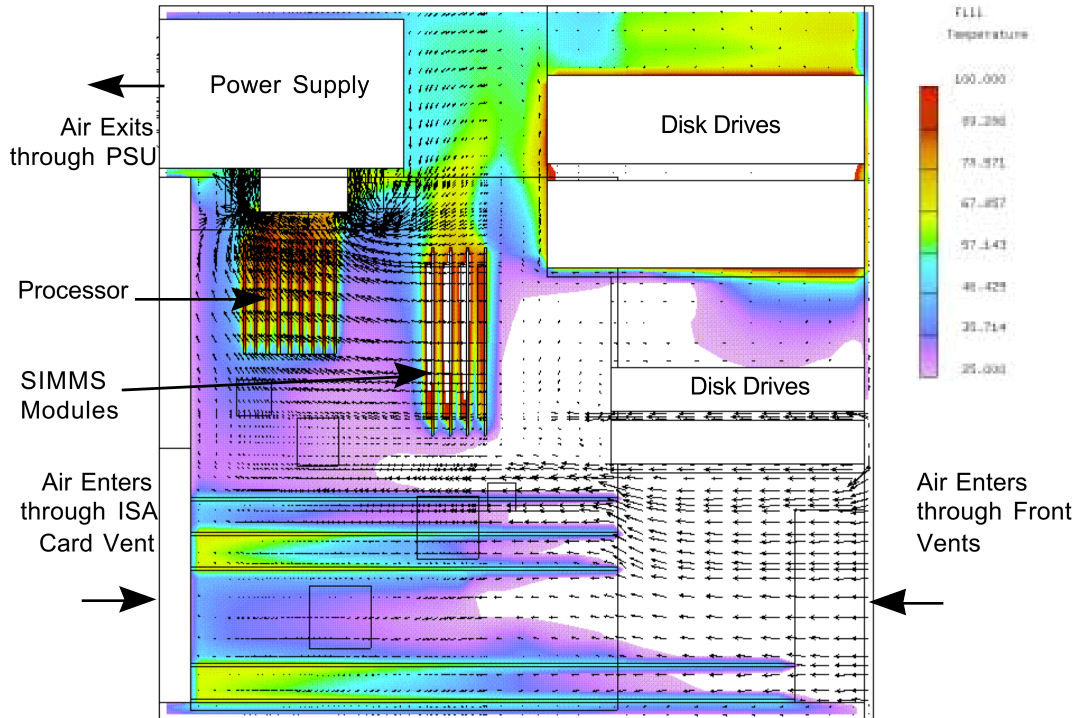

Figur
e 12 - Results for the Fan Exhausting Configuration
Temperature Fill on a plane through the Processor Heatsink
Velocity Vectors on a plane through center of the PSU Fan

Figure 12 shows the results for the fan exhausting. The peak temperature in the processor is about 102.5∞C and, once again, temperature in the SIMMS modules is very high.

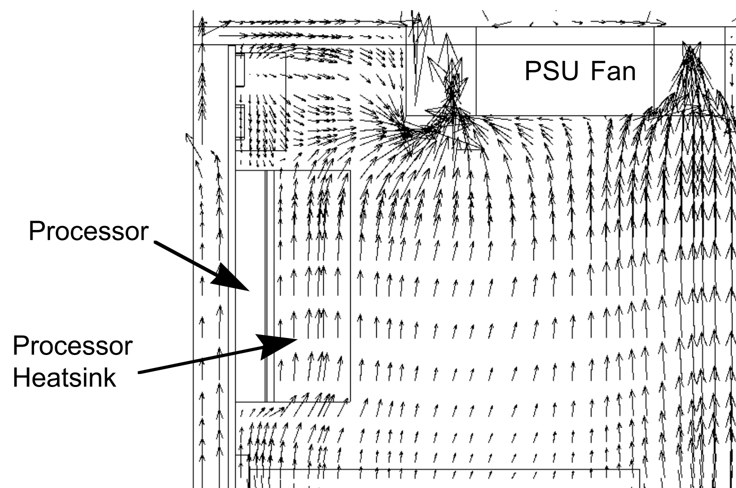

Figure 13 - Results for the Fan Exhausting Configuration
Close-up of Flows Through the Processor Heatsink

Figure 13 shows how changing the direction of flow of the fan results in a flow that is far more even through the processor heatsink. Thus, even though the overall performance of the processor is degraded for this particular case (by about 6∞C), a design based on an exhaust fan configuration would be expected to be much more tolerant of variations in PSU, PSU fan or processor heatsink configuration.

Finally, one of the benefits that is often sought when using an exhaust fan configuration is improved airflow through vents in the casing. An exhaust fan creates a negative pressure (relative to the atmosphere) inside the case which tends to result in more even flow through the case vents (see, for example, Reference [FLOMERICS, 5]). In this system, we might expect to find improved cooling of the expansion cards. However, as comparison of Figures 10 and 12 shows, this does not seem to be the case here.

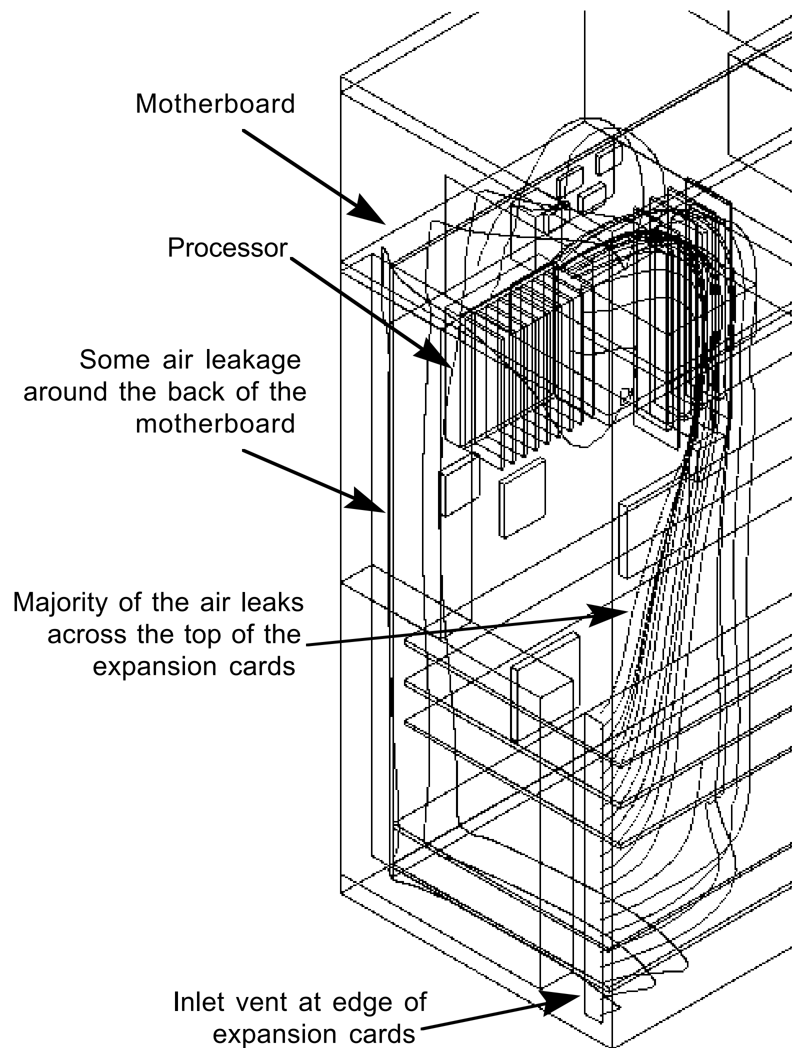

Figure 14 - Results for the Fan Exhausting Configuration
Particle Tracks for Flows Through the Expansion Card Vent

Figure 14 shows one reason why. Particle tracks for flow through the expansion card vent at the back of the computer clearly show how most of the air that enters bypasses the expansion cards either around the back of the motherboard or over the top of the cards themselves. This may, in

part be due to the opposing air flow from the vents in the front of the case and changes in vent positioning would be worth investigating.

CONCLUSION

With the interface between AutoTherm and FLOTHERM now available, the electronic companies have the tools needed to redefine their overall PCB design process, bringing thermal analysis into this process. With the increasing complexity of boards, board-level thermal analysis has to linked tighter with the layout tools, to answer for example questions like : “if I need to implement thermal enhancements to cool this component, is the board still routable?” But with the shrinking of the enclosure in which these boards have to fit, it is imperative to bring system-level thermal constraints, computed with FLOTHERM, into the EDA environment.

In the future, it is expected that the various analyses needed for evaluating PCBs will be part of the front-end of the design process (Seaton, 1996). These tools will help the designers to define the product performance requirements, and also to translate these requirements into constraints for further automation of the design itself. They will also be used to interactively assist the designer in making design decision on the fly, or get back to the specialist if required. But the implementation of these tools will still rely on well established design process flows that can be tested and implemented with the tools available today.

Bringing successfully analyses tools in the PCB design process means that these tools may be used on-the-fly by a non-specialist. It may be achieved for example for thermal analysis if the model creation can be automated, using in particular components model libraries and if reliable boundary conditions obtained with FLOTHERM are available (keeping in mind that, as we have seen above, the synchronization of board-level and system-level modeling need still to be worked out). The creation and the validation of this data remains the domain of the specialist, working in parallel on different projects, within cross-functional teams.

Appendix A - Major Features of the Yellowknife system

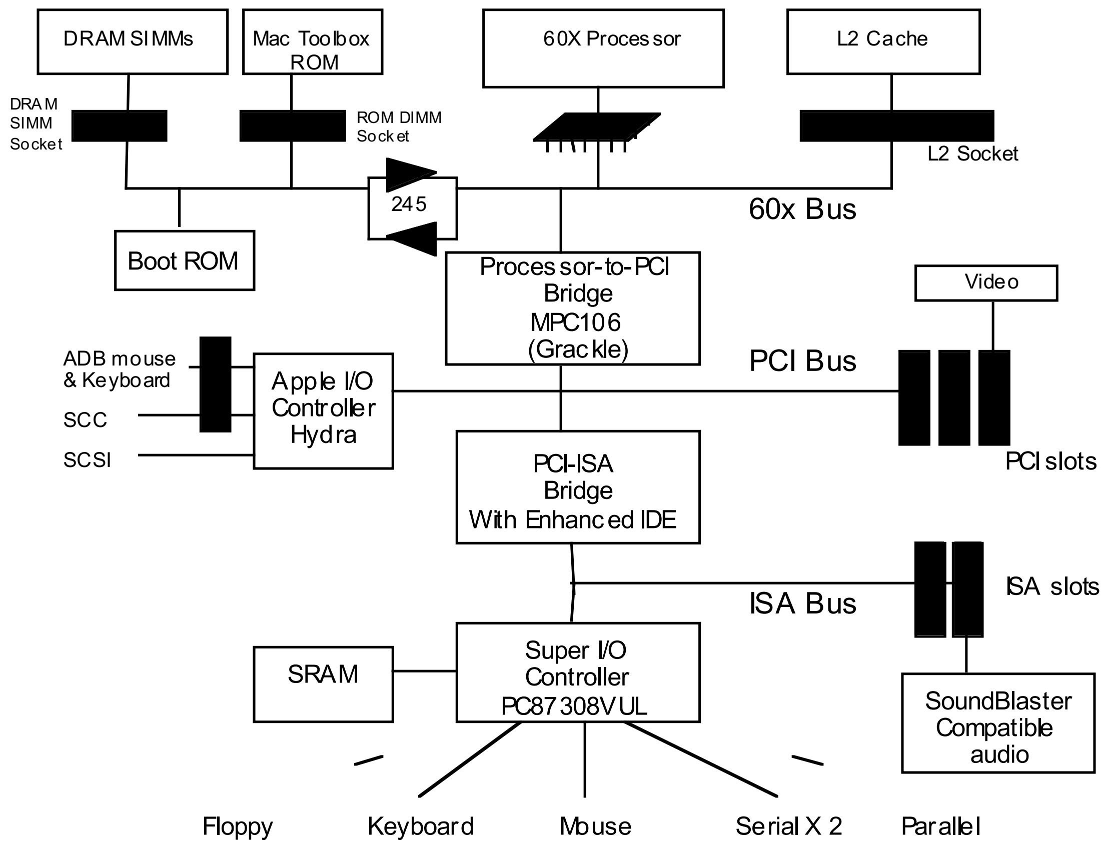

Figure A.1 - Block Diagram of the Yellowknife Evaluation System

Table A.1 - Summary of the Major Features of Yellowknife

Processor and Chipset support	CHRP - compliant
	Processors supported, when available: all 2.5V and 3.3V 603s and 604s with external bus frequencies up to 66MHz
	One PGA (Socket 3) connector on board
	MPC106 Processor to PCI bridge
	Winbond 565 or VLSI Tollgate as PCI to ISA bridge
	Mac I/O chip (Hydra) from TI
Memory support	
	Four 36-bit (72-pin) SIMM sockets
	8 MB minimum, 128MB maximum DRAM
	Support for 4MB, 8MB, 16MB and 32MB SIMMs
	Support for Page mode and EDO DRAMs
	160-pin COAST socket on board
	Second-Level cache selectable between 0KB, 256KB, 512KB
	Flash EPROM for Open firmware
	ROM SIMM socket on motherboard for MacToolbox ROM

Peripheral support

Three PCI slots (one used for graphics card)
Two ISA slots
High performance PCI-bus graphics card
Two serial ports (buffered, 16550-compatible) and two Apple SCC serial ports
One DB25 parallel port
Enhanced IDE interface support for CD-ROM and Hard drive
MESH SCSI interface support
Support PS/2 type 101-key keyboard and Apple ADB keyboard
PS/2 mouse port and Apple ADB mouse port support
Floppy interface support two drives
SoundBlaster compatible audio support on add-on card

Chassis

ATX size chassis with ATX power supply
Two external 5.25" drive bays (one with IDE CD-ROM drive, one unused)
Three external 3.5” drive bays (one with floppy drive and two unused)

Software Target to run Mac OS and Windows NT in 96

Appendix B - Constructing a Compact Model of the PSU

Given the flow rate for a PSU with no back pressure (say, from a bench test or manufacturers specifications) and the characteristics for the PSU fan, the following procedure can be used to construct a FLOTHERM compact model.

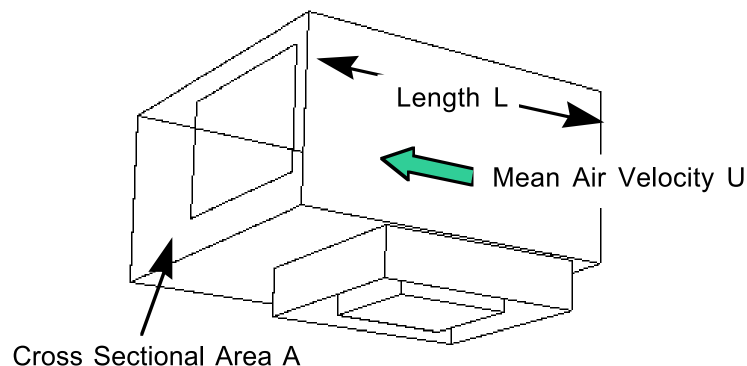

Figure B.1 - PSU Parameters

Estimating the Parameters

The measured or quoted zero back pressure flow rate (mass flow rate m_0 or volumetric flow rate Q_0) can be used to estimate the PSU pressure drop (Dp) from the fan characteristic curve.

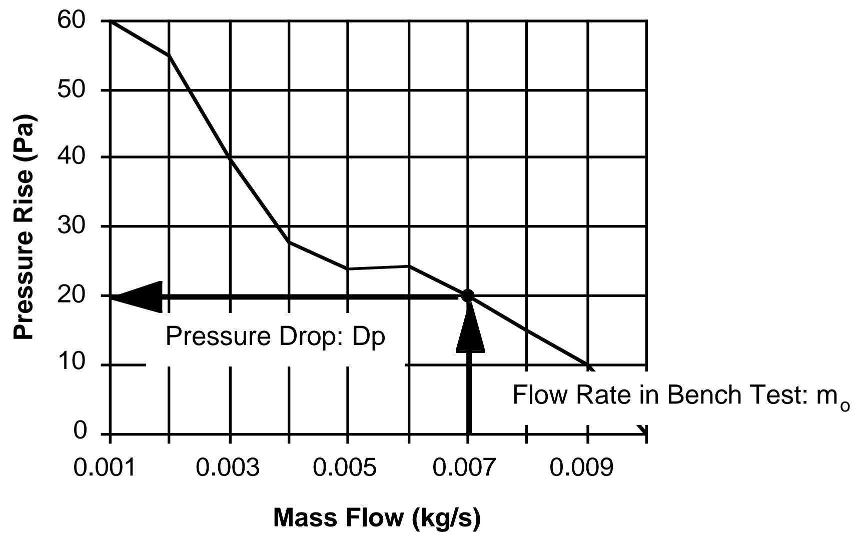

Figure B.2 - PSU Pressure Drop (Dp) Estimated from the Zero Back Pressure Flow Rate (Q_0)

The flow rate (m_0 or Q_0), density (r) and approximate cross sectional area (A) of the PSU give an estimate of the air flow velocity (U) through the unit:

$$U \approx Q_0 / A = \dot{m}_0 / \rho A$$

We now take the length of the unit (L). In the case of an “L”-shaped configuration as we have in this chassis, an approximate length scale is used. Then the loss factor per unit length is given by:

$$F \approx \frac{\Delta p}{\frac{1}{2} \rho U^2} \frac{1}{L}$$

Creating the Compact Model

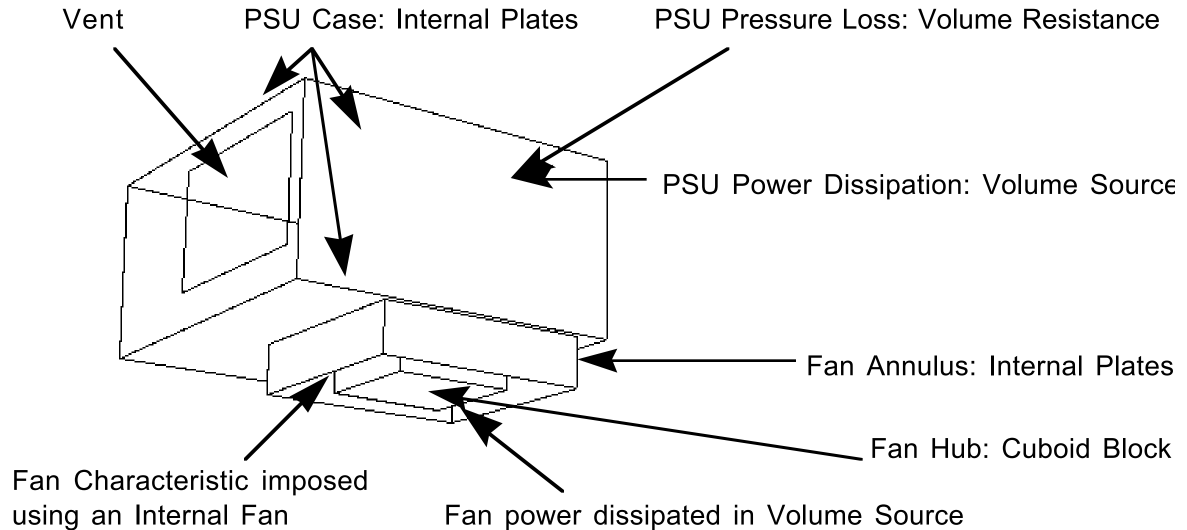

Figure B.3 - FLOTHERM Compact Model of the PSU

The FLOTHERM compact model can now be constructed as follows:

1. The case is modeled as a series of *Internal Plates* and a *Vent*.
2. The pressure loss through the unit is modeled using a *Volume Resistance* over the whole volume of the PSU. The Loss Factor in each direction is set to F (as calculated above) based on Approach Velocity. The Free Area Ratio is not set (leave it as 1 - it will be ignored by FLOTHERM).
3. The power dissipation of the PSU is modeled using a *Volume Source* over the whole volume of the PSU with a Total Source of Temperature set to the dissipated power.
4. The PSU fan is modeled as:
   - a series of *Internal Plates* to represent the annulus;
   - an *Internal Fan* object (remembering to use the “Punch Hole” option to cut a hole in the Internal Plates of the PSU case); and
   - a *Cuboid Block* representing the hub of the fan.

   The power dissipation of the fan is represented by a *Volume Source*. Modeling fans with swirl is discussed in Appendix C.

Calibration of the Compact Model

Whilst this could give acceptable results, it would be prudent to “calibrate” the compact model by running it in a “computational wind tunnel” and adjusting the volume resistance to match the experimentally-determined or manufacturer-supplied data.

Reservations

There are a couple of reservations about the way the PSU is being modeled:

1. We are making a fairly radical assumption that the fan characteristic quoted applies to the fan bolted to the PSU. However, as is well known (see, for example, [HARMESEN,1] and [Flomerics, 6]) fan characteristics are degraded by the proximity of objects. In an ideal world, we would be able to obtain manufacturers data measured in this configuration. But, in the

real world, we have to make do with the information that we have and make due allowance for inaccuracies in our assessment of the final results.

2. We are assuming that the principal mechanism for cooling the PSU is through convection and that the PSU designer has not used the case of the PSU as a heat sink. Should the designer be using the case as a heatsink by bolting components directly to it, then a more detailed assessment of the PSU needs to be undertaken using cuboid blocks for the structural elements. In this case, more detailed information from the PSU vendor will be required (ideally, a validated FLOTHERM model which you can include directly into your chassis model).

Appendix C - Modeling Fans with Swirl

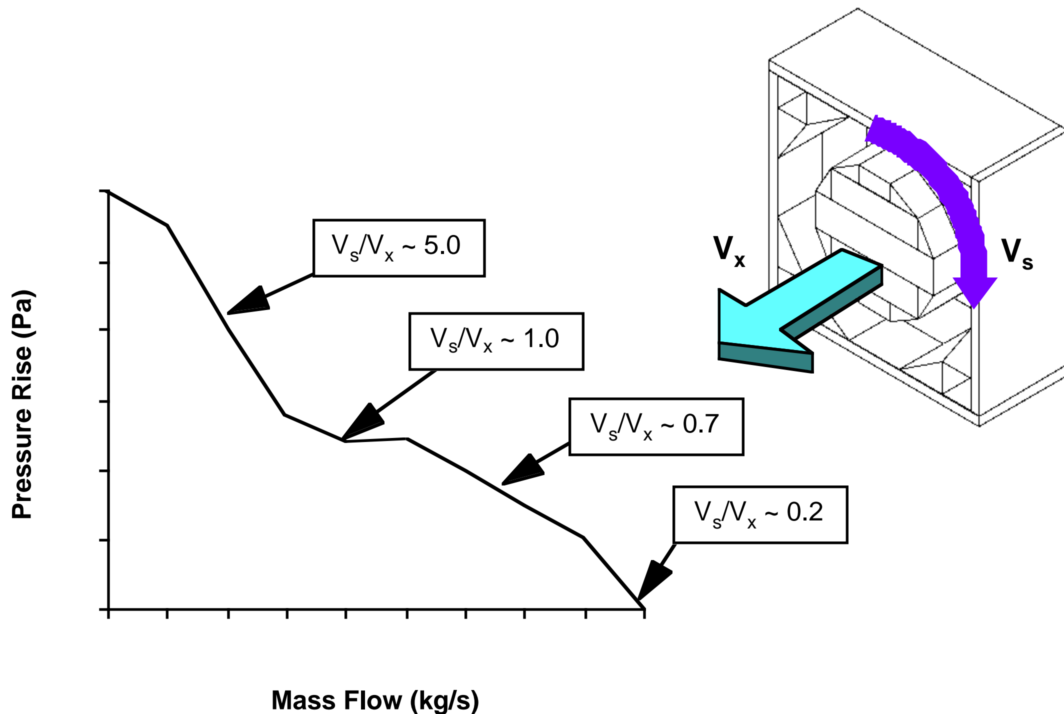

Figure C.1 - Representative Values of Swirl Coefficient V_s/V_x [HARMESEN,2]

Data regarding fan swirl in fans of the size used in electronics cooling applications is hard to come by. The above graph is based on data from a major electronics fan manufacturer and shows typical values for the ratio of the swirl to axial velocity V_s/V_x where the swirl velocity is defined as the average tangential velocity across the length of the fan blades. In the range of interest to most users, the ratio is about 0.7.

To incorporate swirl in a fan in FLOTHERM in Version 1.4, we estimate the axial flow velocity for the fan at it's operating point and then use a series of volume sources to impose a pre-defined tangential velocity component around the circumference as described in Reference [Flomerics, 7] and shown in Figure C.2.

Axial Velocity $U = V_x$ set by Internal Fan at all Points

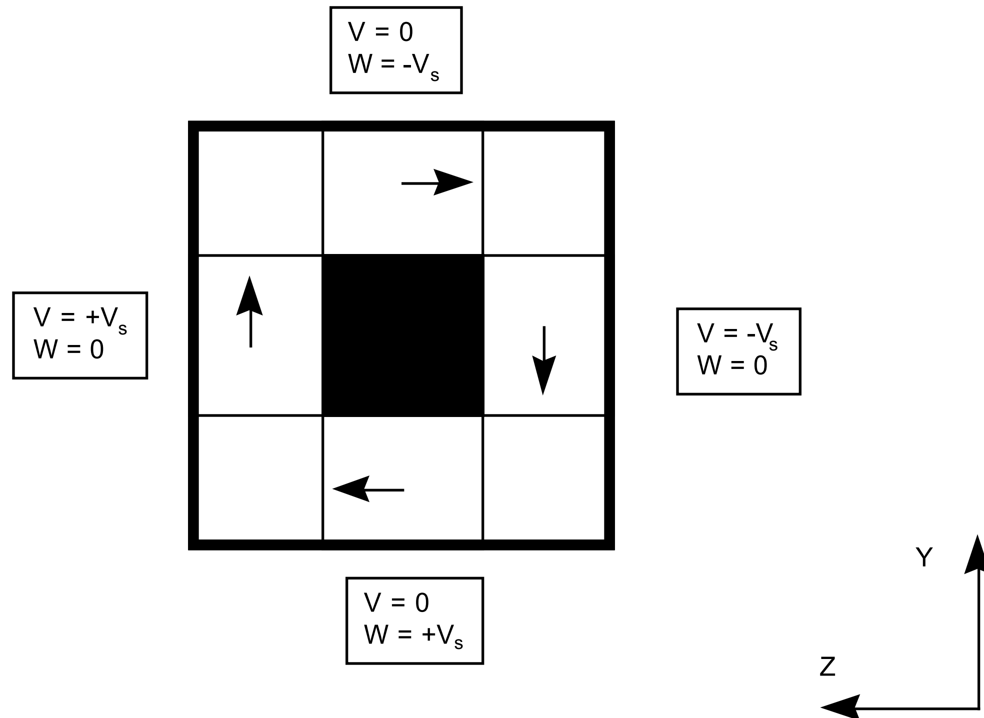

Figure C.2 - Setting Volume Sources to Create Swirl for an Axial Fan

ACKNOWLEDGMENTS

The authors wish to acknowledge the contributions of the Yellowknife RISC product team and the support of Grant King of Mentor Graphics.

REFERENCES

ANDREWS, J., 1988, "Package Thermal Resistance Model Dependency on Equipment Design", IEEE Components, Hybrids, and Manufacturing Technology, vol. 11, No. 4, pp. 528-537.

FLOMERICS,1, http://www.flomerics.com/tips/flow_res/flow_res.htm

FLOMERICS, 2, “Specifying Flow Resistances”; FLOTHERM User News; March 1995.

FLOMERICS, 3, FLOTHERM Advanced Training Course: Lecture II - Uses of Resistances; 28 April 1993.

FLOMERICS, 4, http://www.flomerics.com/desktop/desktop.htm

FLOMERICS, 5, http://www.flomerics.com/ft_about/whatkinds/intel/intel.htm

FLOMERICS, 6, “Don’t Be So Attached to Your Fan Curve”, FLOTHERM User News; August 1995.

FLOMERICS, 7, FLOTHERM Lecture Course Notes; pV-6 Fans with Swirl; Ref : FLOTHERM/LC/0993/1/0.

FRIED & IDELCHICK, “Flow Resistance - A Design Guide for Engineers” ; publ. Hemisphere.

GERKE, R.D., and KROMANN, Gary, 1995, The Effect of Solder-joint Temperature Rise on Ceramic-Ball-Grid Array to Board Interconnection Reliability: Freescale PowerPC 603™ and PowerPC 604™ Microprocessors and MPC 105 Bridge/Memory Controller,ASME/JSME INTERpack-95, March 1995. (Also available from Freescale's PowerPC website: http://www.mot.com/powerpc/)

HARMESEN, S., 1, "Equipment Fans for Electronic Cooling", p55, Publ. Verlag Moderne Industrie.

HARMESEN, S., 2, Private communication from Dr. S. Harmesen of PAPST-Motoren, 7 July 1994.

INTEL, 1996, ATX Specification - Release 1.1 Febr. 1996 : http://www.intel.com/pc-supp/motherbd/atx.htm

KROMANN, Gary, et al, 1995, Motorola's PowerPC 603™ and PowerPC 604™ RISC Microprocessor: the C4/Ceramic-ball-grid Array Interconnect Technology, Proceedings of the 45th Electronic Components & Technology Conf., May 1995. (Also available from Freescale's PowerPC website: http://www.mot.com/powerpc/)

KROMANN, Gary, 1996, Thermal Management of a C4/Ceramic-Ball-Grid Array: The Freescale PowerPC 603™ and PowerPC 604™ RISC Microprocessors, Semitherm-96, March 1996. (Also available from Freescale's PowerPC website: http://www.freescale.com/powerpc/)

LASANCE, C., 1995, The need for a Change in Thermal Design Philosophy, Electronics Cooling, Volume 1, no.2, pp24-26

LIN, P.C. and NQUYEN, L.I., 1993, Thermal Characteristics of IC Packages, EUROTHERM Seminar 29, Delft, The Netherlands, 14-16 June 1993.

Freescale, 1994, PowerPC 603 RISC Microprocessor User's Manual, IBM/Freescale.

Freescale, 1995a, PowerPC 604 RISC Microprocessor Hardware Specifications, Freescale. (Also available from Motorola's PowerPC website: http://www.mot.com/powerpc/)

Freescale, 1996a, PowerPC 603e RISC Microprocessor Hardware Specifications, Freescale. (Also available from Motorola's PowerPC website: http://www.mot.com/powerpc/)

Freescale, 1996b, PowerPC 604e RISC Microprocessor Hardware Specifications, Freescale. (Also available from Motorola's PowerPC website: http://www.mot.com/powerpc/)

PERRY & CHILTERN, Eds, "Chemical Engineer's Handbook".

ROSTEN, H., 1996, "DELPHI - A Status Report on the European-funded Project for the Development of Libraries and Physical Models for an Integrated Design Environment", Presented at 46th Electronic Components & Technology Conference, Orlando, May 28-31 1996.

SEATON, 1996, Getting Systems Onto Printed Circuit Assemblies, Printed Circuit Design, 7/96, pp 23-29.

Trademarks

FLOTHERM is a registered trademark of Flomerics Ltd.

AutoTherm and Board Station are registered trademarks of Mentor Graphics Corp.

PowerPC, PowerPC 601, Power 602, PowerPC 603, PowerPC 604, and PowerPC 620, are trademarks of International Business Machines Corporation and are used by Freescale, Inc. under license from International Business Machines Corporation.