

Novel Simulation Techniques for Design of Air-cooled Electronics

David A. Johnson, Mark J. Welch, Brent A. Cullimore
C&R Technologies, Inc.
9 Red Fox Lane
Littleton Colorado, 80127-5710
303 971 0292, 303 742 1540 (FAX), dave@crtech.com

ABSTRACT

This paper describes a new means of analyzing the thermal response of air-cooled and liquid-cooled electronics that overcomes limitations in available tools and current design methods. It also shows how these new tools and methods can extend the reach of such thermal/fluid analyses by helping to size and locate components as well as dealing with both pre-test uncertainties and post-deployment variations in manufacturing, environment, and usage.

As the time lag from design to market diminishes, the pressure to abandon "build and test" approaches to electronics thermal cooling has created a wide variety of design analysis methods ranging from simple hand calculations of energy balances to detailed three-dimensional CFD (Computational Fluid Dynamics) approaches. Surprisingly few options are available between these two extremes, leaving most designers feeling that they face an "all or nothing" choice. Hand calculations and other simple software approaches, while contributing to an engineer's intuition, cannot be relied upon for the entire design cycle, especially with the reduced emphasis on hardware prototyping that is necessary to speed up product development time. Fluid network modeling (FNM) approaches offer more analytic power but lack strong connectivity to geometric thermal models, and are therefore cumbersome to use. CFD approaches include limited geometric thermal modeling, but are relatively inflexible because they focus on detailed point design evaluations, and therefore contribute little to design knowledge.

This paper will describe a new approach using multidimensional heat transfer modeling in combination with

ducted or quasi-multidimensional flow solutions for fast and easily modifiable models of electronics packaging that lends itself to high-level operations such as sizing and reliability estimation.

SUMMARY OF CURRENT METHODS

In order to create a background for understanding the new approach presenting here, a brief summary of the current methods will be presented along with their strong and weak points.

1D, 2D, 3D Thermal Modeling

A variety of network-style thermal conduction/capacitance modeling tools exist, including Thermal Solution's Sauna®, Network Analysis' SINDA/G®, Thermal Associates' TAK, and the SINDA side of C&R's SINDA/FLUINT. Usually these codes are erroneously considered "finite difference" when in fact they are geometry-less thermal network (circuit) solution engines that can be used to solve not only finite difference problems and 1D lumped parameter problems, but also finite element problems (with proper input preparation). They usually feature concurrently executed user logic and/or other equation-style inputs. Increasingly, thermal network analyzers are used with graphical user interfaces (usually geometry-based) that help prepare inputs, although most can still be accessed at the "thermal circuit level." Such access is important for high-level lumped parameter modeling in which a complex component such as a battery might be represented using effective thermal mass, conduction, surface area, etc.

Similarly, there is no shortage of software tools for modeling steady or transient conduction within shells or solids, usually

using finite elements (e.g., Harvard Thermal's TAS), occasionally using finite differences (e.g., SDRC's TMG®), and in at least one instance (C&R's Thermal Desktop®, Ref Panczak) both finite elements and finite differences can be used in a mix-and-match fashion. Indeed, almost every finite element method (FEM) structural program offers such "heat transfer modeling" as an option. With varying degrees of flexibility, most of these thermal analysis codes also supply means of generating models from CAD data or at least structural FEM models themselves generated from CAD representations.

Because of the larger emphasis on structural analysis, few of the available surface and solid (2D/3D) codes are specifically designed for thermal management tasks. Only those that *are* so oriented tend to support analysis of higher level assemblies critical to product-level heat transfer, including effects such as contact conductance and efficient radiation calculations. Few provide any fluid flow capabilities, excepting those that use full CFD (e.g., SDRC's ESC®), an approach which will be summarized later. A few others provide simple fluid flow networks, although in many such codes *answers* such as flow rate and heat transfer coefficients must be supplied as *inputs*, and interconnections with the 2D/3D thermal geometry is not automated.

Notwithstanding these difficulties, the main point of this subsection is to emphasize that most thermal engineers have access to or can relatively easily generate 2D/3D thermal conduction models, and some can generate models with thermal features such as contact conductance and radiation, but few can link air flow modeling into these models without resorting to a full 3D CFD solution.

1D Flow Networks

Once again, there is no shortage of 1D fluid network or "piping system" codes, almost all of which have friendly circuit-style sketchpad (2D but nongeometric) graphical user interfaces. However, very few of these are applicable to thermal control coolant loops much less air flow within boxes since most of them lack heat transfer altogether, and often even energy conservation. Transient capabilities are usually absent, or when present are focussed on hydrodynamic transients (i.e., waterhammer).

A very few flow network (FNM) codes, however, *are* intended for heat transfer applications, examples including Innovative Research's Macroflow® and C&R's FLUINT side of SINDA/FLUINT (Cullimore 2001a, perhaps using the SinapsPlus® sketchpad graphical interface). Prior to the developments described later in this paper, however, these thermal-oriented FNM codes did not provide direct access to and interconnection with 2D/3D thermal models nor were they able to extract data from CAD drawings for faster model building.

3D CFD

One of the most important recent enhancements of an engineer's ability to predict the performance of air-cooled electronics enclosures is the advent of CFD tools such as Flomerics's Flotherm®, Fluent's IcePak®, and SDRC's ESC®. It is therefore important to emphasize that while the quasi-3D methods that will be described in this paper propose alternatives to full CFD solutions, they can only replace *some* types of current CFD applications. No single method is the solution for all thermal/fluid analytical requirements. However, it is the opinion of the authors that, in leaping to full CFD solutions, an important "middle ground" has been overlooked.

Therefore, the immediate question is why CFD solutions aren't globally applicable to all problems, especially with the focus on enclosure-level modeling of heat transfer and fluid flow paths that is taken in this paper. There are three reasons.

First, CFD codes can be comparatively expensive both to acquire and maintain, but also to become proficient in and to retain that proficiency. High-level models of enclosures can be time-consuming both to create and to solve. Many organizations only perform such analyses occasionally, and smaller organizations tend to shy away from *any* analytic solution because of such hurdles, or they use simple in-house codes or hand calculations.

Second, CFD codes are comparatively inflexible with respect to rapid model changes, including "what if" style parametric analyses, sensitivity/bottleneck studies, model verification, component placement or sizing studies, or other tasks important in preliminary design. They are also not oriented toward automatic model calibration to test data (Cullimore, 2001b) that is important in later design phases.

Third, CFD codes struggle with producing accurate film coefficients with reasonable meshes. In 1D fluid network codes, frictional losses and convection heat transfer are treated empirically, using Nusselt-based duct flow correlations and often making rough assumptions such as fully developed flows. CFD codes usually make no such assumptions, relying more on "first principles" approaches and the actual geometry. However, closure of the momentum and energy equations at the wall requires some trade-off between a very fine mesh and an empiricism. The error terms for energy solutions are larger than those for momentum solutions because the former is based upon the latter: *heat transfer coefficients are highly derived and can have relatively large uncertainties*. Because a large increase in mesh size results in only a modest decrease in this uncertainty, a point of diminishing returns is quickly reached.

These uncertainties in CFD results are rarely greater than those which occurs due to the misapplication of correlations in 1D codes, however. If that statement is true, then it might seem that a CFD approach would always be superior. The key becomes how those uncertainties are overcome in both

approaches, as will be discussed later (see also Cullimore, 2001b).

GEOMETRY-BASED FLOW NETWORKS

The principal innovation described in this paper can be summarized succinctly: a complete fluid network modeling (FNM, sometimes referred to as *thermohydraulic*) code has been deployed within a CAD-based 3D thermal modeling package specifically to address the aforementioned shortfalls in thermal design of electronics packaging.

The modeling of arbitrary networks of fans, pumps, ducts, valves, filters, and other miscellaneous loss elements may be accomplished as with almost any flow network modeler. Fluids may be user-specified, but common choices for air-cooled electronics include dry air and moist air (including condensation and other psychrometric effects), and for coolant loops fluids choices include water, water-glycol, PAO, etc. Even two-phase systems such as vapor compression cycle refrigeration systems can be modeled.

Convection: Simple Air Passages

The benefits of quickly laying out an air “circuit” within a 3D thermal model framework can be seen in Figure 1. While the geometry provides a quick basis for the selection and input of 1D flow parameters such as hydraulic diameter, flow area, and the “wetted” perimeter for heat transfer calculations, the real productivity improvement is the automatic generation of convection heat transfer between 1D fluid network elements and the 2D heat transfer surfaces (boards, chassis, etc.). This can be seen by comparing the top part of Figure 1, in which fluid-to-wall heat transfer connections have been suppressed, with the bottom part of Figure 1, in which they have been activated.

The thermal surface to which the connections are made might be a discretized shell element (rectangle, cone, disk, etc.), perhaps with its own radiation, insulation, heat generation, internal conduction, contact conduction to other elements, etc. Or, the surface might be a surface coated on an arbitrary 3D finite element model, perhaps imported from another source.

Appropriate interconnections to the nearest fluid element appear automatically, apportioned according to the surface area of each thermal node (whether based on finite differences or finite elements). If the resolution of the fluid network is altered or if the network is moved relative to the surface, interconnections are automatically regenerated. Similarly, if the resolution of the thermal model changes (such as the increase demonstrated in the inset of Figure 1), convection connections are again updated automatically. Automatic connections can be extended to more complicated circuit board geometries, as depicted in Figure 2.

The heat transfer coefficients are generated empirically, and therefore will not always be appropriate for the exact situation. Also, flow distributions between parallel legs (including by-pass ratios over heat sinks) will similarly be approximate. However, the entire model can be generated parametrically (algebraically, symbolically) using scaling factors on each heat transfer surface or loss factors and by-pass ratios on each flow passage. These scaling factors can then be assigned as uncertainties for a statistical analysis or as unknowns to be back-calculated automatically from test data. These advanced features, enabled by a fast-solving parametric model, are described later.

Convection: Coldplates and Coolant Loops

In some applications such as aircraft avionics, coolant loops are used in addition to or instead of air cooling. A flow network approach is clearly warranted in this case; a CFD analysis of a piping loop would represent significant “overkill.”

However, without the intimate interconnections with 2D/3D thermal geometry, model generation and management is extremely cumbersome and error prone. Figure 3 illustrates a liquid cooled coldplate: the attachment of 1D fluid flow elements to a 2D thermal plate.

Again, the surface to which the coolant channel is attached can be a coating on a solid, or can be a thermal finite element or finite difference shell of any resolution. The axial resolution in the fluid element can also change, with connections established automatically. The connection between the two automatically includes convection based on built-in correlations, but can also include an linear or areal contact conductance as well as an extruded container/fin cross section (which itself can be discretized if needed).

MODEL CALIBRATION AND ADVANCED DESIGN: BETTER USE OF MOORE’S LAW?

Moore’s Law (stated in 1965 by G. Moore, one of the founders of Intel) implies that computer processor speed has increased by a factor of from 1.5 to 2 every year for three decades, and will continue to do so for the foreseeable future. It is as if the computer-aided engineering (CAE) community has been blessed with a geometrically increasing “budget.” But how is that budget being spent in the thermal/fluid analysis arena?

The answer is that most of the budget is absorbed by increasingly larger (more detailed) models, with added phenomenological modeling absorbing most of the remaining budget. Improved graphics, user interfaces, and interconnection/interchange between software has also taken advantage of faster machines. Nonetheless, the basic approach of *point design evaluation* remains largely unchanged: using a fixed design in a specific environment, predict the steady-state and/or transient performance.

The above represents not what an engineer needs to accomplish, but rather what is convenient to solve numerically assuming inputs are known precisely. Specifically, point design evaluation is merely a *subprocess* of what an engineer must do to produce a useful and efficient design. Sizing and locating components and coping with uncertainties and variations are the

real tasks. Simulations alone cannot produce effective designs, they can only verify deterministic instances of them.

Why perform expensive and detailed point design simulations when the heat transfer coefficients are only accurate to within 20 to 50%? Even if further improvements in CFD technology




Figure 1. Example of Air Flow Model without Convection Connections (top), with Convection Connections Automatically Placed (bottom), and Additional Detail in the Thermal Model (inset)

were to completely eliminate this uncertainty and yet maintain reasonable run times, there remain variations in environment, usage, fabrication, installation, etc.

Therefore, the real power of a simplified and parametric flow analysis is to trade accuracy within a *single* relatively slow evaluation for the ability to run multiple faster analyses and




Figure 2. Postprocessed Example of a Circuit Board Requiring Complex Convection




Figure 3. Postprocessed Example of a Cold Plate: 1D Coolant Elements Attached to a 2D Plate

thereby:

1. automatically calibrate a model to available test data, effectively eliminating or at least reducing uncertainties such as heat transfer coefficient and contact conduction. This approach uses modeling not to attempt to replace testing with more and more detailed analyses, but rather to extend (and therefore significantly reduce the need for) testing with fast-to-generate and fast-to-solve analyses;
2. allow remaining variations and unknowns to be evaluated together statistically to determine tolerancing, to focus on critical (bottle-neck) uncertainties, and to avoid margin stack-up; and
3. get help producing a sensible, robust, and efficient design in the first place using automated design synthesis and optimization techniques. Using multidisciplinary optimization (MDO) techniques, this design synthesis can transcend thermal/fluid considerations, including cost and reliability models, etc., although the need for fast solutions becomes even more critical in those cases.

These advanced techniques are elaborated in Cullimore (2001b). The purpose of this subsection is simply to compare a detailed point design evaluation technique using a CFD approach with a faster albeit more approximate technique that provides alternate means of dealing with inaccuracies that are intrinsic in *both* approaches, while also lending itself better to supporting high-level design decisions in a timely fashion.

LIMITATIONS OF THE GEOMETRY-BASED FLOW NETWORK APPROACH

As was mentioned before, this paper describes an alternative to full CFD solutions for *certain* applications, including box-level design and analysis of air-cooled and liquid-cooled electronics. This section lists a few example applications where the geometry-based flow network (quasi-3D) approach is clearly inappropriate.

In a nutshell, a quasi-3D approach cannot be used in any application where 2D or 3D flow fields occur *and* their behavior cannot be adequately characterized empirically (allowing for scaling and correction factors as discussed above).

One example is natural convection *within* an open cell or cavity. Quasi-3D approaches can handle natural convection, but are limited to ducted or contained flows (such as flow rising between a pair of cards and descending between the last card and the enclosure wall).

Another example is detailed flows *within* a heat sink, perhaps during the design of the heat sink itself. Otherwise, a correlation-based approach would be applicable once vendor data for the particular heat sink is available for use as an input to a higher-level modeling approach.

CONCLUSIONS

Thermal design analysis of high-level assembly such as electronics enclosures has been accomplished with either nongeometric 1D flow network codes, or with geometrically faithful 3D CFD codes. The former are quick to generate and solve, but cumbersome to integrate with surface and solid thermal models. The latter are slow to generate and solve, and retain uncertainties in critical parameters such as heat transfer coefficients.

An intermediate approach has been outlined in this paper: 1D flow circuits connected directly to 2D/3D thermal models derived from CAD or FEM models. This approach eliminates the problems with stand-alone flow network codes, and is much faster to generate and solve than CFD codes. However, it introduces approximation. These approximations are overcome exploiting the fast-to-solve nature of this approach and the spreadsheet-like input flexibility of the tools: by statistically treating the criticality of each approximation, and by providing automated means of calibrating models with available test data. For the same reasons, this quasi-3D approach is also more amenable to sizing, selection, parametric, sensitivity, and “what if” analyses than are CFD methods.

REFERENCES

- Panczak et al, Thermal Desktop Version 3.3 User’s Manual and tutorials, PDF available at www.crtech.com, December 2000
- Cullimore et al, SINDA/FLUINT Version 4.4 User’s Manual, PDF available at www.crtech.com, May 2001
- Cullimore, B., “Dealing with Uncertainties and Variations in Thermal Design,” IPACK2001-15516, July 2001
- Moore, G., “The Continuing Silicon Technology Evolution Inside the PC Platform,” developer.intel.com/update/archive/issue2/feature.htm, May 2000