CAD-based Methods for Thermal Modeling of Coolant Loops and Heat Pipes

David A. Johnson, Jane Baumann, Brent Cullimore
C&R Technologies
Littleton, Colorado
303-971-0292

Abstract

As air cooling of electronics reaches the limits of its applicability, the next generation of cooling technology is likely to involve heat pipes and single- or two-phase coolant loops (including perhaps loop thermosyphons, spray cooling, vapor compression refrigeration cycles, and loop heat pipes). These technologies are not suitable for analysis using 2D/3D computational fluid dynamics (CFD) software, and yet the geometric complexities of the thermal/structural models make network-style schematic modeling methods cumbersome.

This paper describes techniques whereby CAD line-drawing methods can be used to quickly generate 1D fluid models of heat pipes and coolant loops within a 3D thermal model. These arcs and lines can be attached intimately or via lineal contact or saddle resistances to plates and other surfaces, whether those surfaces are modeled using thermal finite difference methods (FDM) or finite element methods (FEM) or combinations of both. The fluid lines can also be manifolded and customized as needed to represent complex heat exchangers and plumbing arrangements.

To demonstrate these concepts, two distinct examples are developed: a copper-water heat pipe, and an aluminum-ammonia loop heat pipe (LHP) with a serpentined condenser. A summary of the numerical requirements for system-level modeling of these devices is also provided.

Beyond Air Cooling

Forced air cooling is arguably the most common method in use today for cooling electronics. However, air velocities much higher than 3 to 4 m/s are difficult to obtain. Therefore, due to increasingly high heat fluxes and/or inaccesible packaging (such as multichip modules or MCMs), air cooling is rapidly reaching the limits of its usefulness. Such limits have in fact already been exceeded in high power density applications. In many other situations, air cooling no longer represents the best engineering solution, but nonetheless continues to be used along with extreme design measures so as to avoid the political and infrastructural hurdles (but *not* technological hurdles) of moving on to the next step.

What is that next step? What lies beyond air cooling? Single-phase liquid cooling arguably represents the smallest technological step (Ref 1-2), although hermetically sealed heat pipes are also gaining favor as means of extending air cooling. Some organizations have reasoned that if the hurdles associated with adding liquids to the system must be overcome, then two-phase systems should leap-frog single-phase systems to exploit their lower flowrates and higher heat transfer efficiencies. Such two phase systems include "passive" technologies (no pump or compressor) such as heat pipes, loop heat pipes and loop thermosyphons, as well as "active" pumped two-phase coolant loops including evaporative spray coolers (Ref 3). Others (Ref 4 for example) have recognized that, having bothered with the introduction of two-phase systems, one might as well exploit the potential for vapor-cooled refrigeration systems and thereby eliminate the ultimate limit in the rejection path: the temperature difference between semiconductor junction and the ambient.

Whether the answer is single-phase coolant loops (perhaps including ducted air), heat pipes, two-phase coolant loops, or refrigeration cycles, *a change in modeling technology will be required*: the growing emphasis on CFD modeling of air systems will not suffice. But before describing alternatives to CFD, the status of structural (conduction/capacitance/ radiation) thermal modeling will be briefly reviewed.

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-independent 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., or where incorporation of compact models is required.

Similarly, there is no shortage of software tools for modeling steady or transient conduction within shells or solids, usually using finite elements (e.g., MSC/NASTRAN®), occasionally using finite differences (e.g., SDRC's TMG®), and in at least one instance (Ref 5) 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. Most of these thermal analysis codes also supply means of generating models from CAD data, albeit with varying degrees of flexibility.

At the very least, structural FEM models can be generated from CAD representations using a wide variety of software. Unfortunately, such models, being based on structural meshes, are rarely appropriate for direct use as thermal models. 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., Fluent's IcePak®, SDRC's ESC®).

A few other 2D/3D codes provide fluid flow networks. With one exception (Ref 6), most of this class of software require that *answers* such as flow rates and heat transfer coefficients must be supplied as *inputs*. Worse, interconnections with the 2D/3D thermal geometry are not automated. Alternative graphical user interfaces for flow network solvers are based on schematics with the surfaces and solids associated with the thermal model either absent or oversimplified: the emphasis is placed on *either* the 1D fluid modeling, *or* on the 3D thermal modeling, but not both in the same package.

In summary, 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 or ducted coolant flow modeling into these models without resorting to a full 3D CFD solution.

Applicability of 3D CFD Modeling

Ducted single-phase flow with heat transfer may be modeled using a variety of 2D/3D CFD methods. Such models are used in the automotive industry, for example, to determine branching and splitting of flows in complex air ducts, and to determine the register exit velocity profiles as needed to verify even flow distribution into the passenger compartment.

However, very small CFD elements or volumes are required within the boundary layers of objects in freestream (unducted) flow, and computational resource requirements usually increase geometrically with increased discretization. In adiabatic *ducted* flow, CFD elements must be small throughout the model. In ducted flow *with heat transfer*, most CFD codes require even smaller elements to avoid large error terms in estimating conjugate heat transfer at the wall. The cost of solving these models is very high for realistically complex systems such as an entire coolant loop, thereby making transient analyses essentially untenable. Even making parametric or iterative steady-state runs can be too time consuming, especially since few CFD codes offer full parametric modeling capabilities: model and mesh changes are difficult to make *between* runs much less *within* runs.

Two-phase flow with phase change, such as occurs in heat pipes (including loop heat pipes), thermosyphons, spray coolers, and vapor compression cycles is currently beyond the realm of practical commercial CFD modeling for system-level modeling, although it is applied at university-level research.

For these reasons, some CFD providers have recently begun to offer 1D flow modeling alternatives, recognizing that the above limitations are likely to remain intractable for many years to come.

Applicability of 1D Flow Modeling

One dimensional flow models might still be called "computational fluid dynamics" by some engineers, but 1D models are distinguished by the complete elimination of the mesh in the nonaxial dimensions. Instead, well-established empirical correlations are used for both heat transfer and pressure drop. In other words, the boundary layers in 1D duct flow are not solved from "first principles" as in a CFD approach, but rather using computationally efficient assumptions based on copius testing. Because the radial and circumferential dimensions do not need to be discretized, even the axial dimension does not usually require as much subdivision as it would in a CFD approach. Thus, 1D flow models are many orders of magnitude faster to solve than are 3D flow model for ducted systems.

In the 1D approach, momentum conservation is applied axially, with wall friction applied to the axial flow momentum equation using correlations appropriate for the duct shape, fluid, current flowrate, etc. In other words, the only "velocity field" is a single vector in the axial direction (at any point along the flow stream).

Energy and mass (and species etc.) can be conserved at axial points along the flow direction. Heat transfer coefficients can vary around the circumference in a quasi-2D fashion, again using an empirical approach. There is no subdivision of the fluid control volumes in the radial or circumferential directions, resulting in simple fast-solving network schematics.

For single-phase flow, the speed enhancements over CFD methods are dramatic. For two-phase flow, the 1D approach is "enabling" since such problems are essentially intractible using 3D CFD approaches, which must resolve and track each phase and must handle both the sharp gradients and the intense coupling with thermodynamics and heat transfer that is required in two-phase flows.

A "first principals" CFD approach (i.e., eliminating Reynolds- and Nusselt-based correlations) is considered by some engineers to be more accurate. While this opinion is difficult to defend for ducted flows, there *are* some circumstances where an empirical 1D approach is strained. One example is two-phase flow, where 20% error in predicted friction or heat transfer coefficients would be considered "excellent" in the emprical correlations underlying a 1D flow model. Fortunately, the fast solution speed of 1D methods enables higher-level methods for dealing with such uncertainties (Ref 7, 8).

1D solution speeds also allow detailed transient analyses to be made, along with rapid model changes (including parametric sweeps during a single solution run). Such parametric model changes are important precursors for higher-level analyses and design activities such as automated sizing, selection, and location of components (Ref 9).

In summary, the "loss" of the extra mesh dimensions yields an enormous gain in solution speed, and this gain can be applied to higher-level engineering tasks rather than to single "point design simulation" (i.e, predicting how a single design point responds to a single scenario). 1D flow solutions are clearly superior to 2D/3D CFD solutions for ducted flow problems such as those encountered in electronics cooling applications.

However, one problem has existed with the 1D flow network modeling approach for thermal modeling: the lack of integration with 3D thermal models.

1D Flow Modeling within 3D Thermal Models

Reference 6 introduced a methodology for building 1D flow models within 3D (i.e., FDM and/or FEM) thermal models. Selecting 1D flow methods requires that simplifying assumptions be made for modeling air-cooled electronics. While such simplications are not always appropriate for modeling air flows, they *are* appropriate for ducted air or coolant flows, as was discussed above.

However, significant expansions of the methods detailed in that reference were required in order to apply them to ducted flow systems such as coolant loops, heat pipes, and refrigeration systems. Specifically:

- means had to be supplied of drawing free-form lines and arcs using CAD tools, and then enabling these 1D lines elements to be considered as either pipes or ducts (for coolant loops, loop heat pipes, loop thermosyphons, vapor compression cycles, etc.) or as fixed or variable conductance heat pipes
- these fluid lines, whether ducts or heat pipes, had to be able to include the pipe wall or container, if applicable, without violating the 1D assumption: 1D thermal conductive/capacitance network elements were required
- the fluid lines had to be attachable to thermal solids and surfaces with appropriate models for fins, saddles, bonds, contact conductance, etc.
- the fluid lines had to have variable axial resolution, and yet be able to be subdivided as needed to form tees, manifolds, etc.
- the axial discretization (both number and method) of the fluid lines needed to be specifiable independent of the spatial discretization (again, both number and method) of the surface or solid to which the fluid line was to be attached

These improvements have been completed successfully, yielding a methodology uniquely suited to electronics cooling applications requiring ducted air or coolant flow networks.

Two brief applications will be described to illustrate these ideas. First, modeling of constant (or "fixed") conductance heat pipes (CCHPs, FCHPs) will be presented and applied to an example scenario. Second, the replacement of the heat pipe with a loop heat pipe (LHP) will be used to illustrate both LHP modeling techniques as well as the more general case of modeling one- or two-phase coolant loops.

System-Level "Compact" Heat Pipe Modeling

Heat pipe modeling is plagued by two misconceptions. The first is that full two-phase thermohydraulic modeling is required because the devices are "two-phase." While full fluidic solutions *are* applicable to LHPs (see below), they represent "overkill" with respect to heat pipe modeling at the system level. Even during the design of the heat pipes themselves (versus their implementation into a design), simple methods are used by most manufacturers.

The second misconception is that heat pipes can be represented by solid bars or rods with an artifically high thermal conductivity, which is not only disruptive to the numerical solution (especially in transient analyses), but is also not an equivalent representation. Unlike a highly conductive bar, a heat pipe's conductance or resistance is independent of transport length, provided that its internal limits (such as boiling, wicking, entrainment, viscosity, and sonic limits) have not been exceeded. Furthermore, some types of heat pipes can exhibit up to a two-fold difference in convection coefficients between evaporation and condensation, and in realistically complex geometries the analyst shouldn't assume a priori which sections will absorb heat and which will reject it: the resulting temperature profiles should instead govern such decisions during the solution itself.

It is also important to be able to track power throughputs in a heatpipe in a format comparable with the vendor-supplied rating: the integrated power-length product ($Q*L_{eff}$). Given a safety margin, this comparison is all that is usually needed to ensure that the heat pipe has not exceeded its operational limits. The power-length product is also important when designing arrays of parallel (and perhaps redundant) heat pipes to make sure that each is carrying an appropriate load.

Fortunately, a relatively simple network-based heat pipe modeling method is available that has been used for years in the aerospace industry, which has about 3 decades of experience using heat pipes. To explain this method, first consider a simple one-dimensional finite difference wall model with only axial gradients considered, as presented in

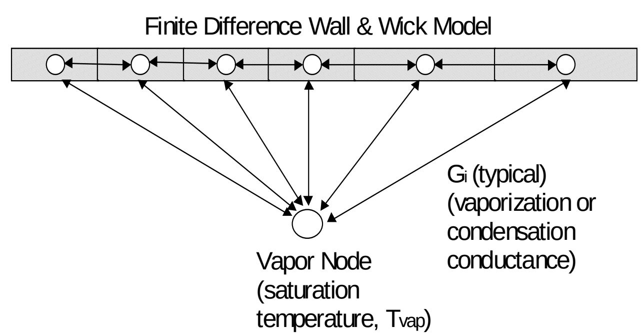

Figure 1.

Figure 1: System-level Network Model of a Heat Pipe

The key to this approach is the addition of a massless node representing the vapor saturation temperature (T_{vap}). All wall nodes are then attached to this node with a conductive "fan" where the conductance of the i^{th} leg (whose temperature is T_i, whose internal surface area is A_i , whose volume is V_i) is equal to:

$$G_i = H_e*A_i \qquad (T_i > T_{vap})$$
$$\text{or} \quad G_i = H_c*A_i \qquad (T_i < T_{vap})$$

where H_e is the coefficient of heat transfer for vaporization, and H_c is the corresponding coefficient for condensation. These values are normally provided by the heat pipe vendor.

This method can be easily extended to a two-dimensional heat pipe wall, and even to arbitrarily shaped vapor chamber fins. Consider, for example, Figure 2, which depicts an Intel Xeon™ CPU chip cooler that employs embedded heat pipes (Ref 10). In this case, the size of the heat pipe diameter compared to the lateral fins presents problems with a completely 1D approach to modeling the heat pipe. Therefore, a 2D cylindrical shell has been used instead, permitting temperature gradients to exist around the circumference of the pipe. Nontheless, the algorithms presented in this section are still applicable.

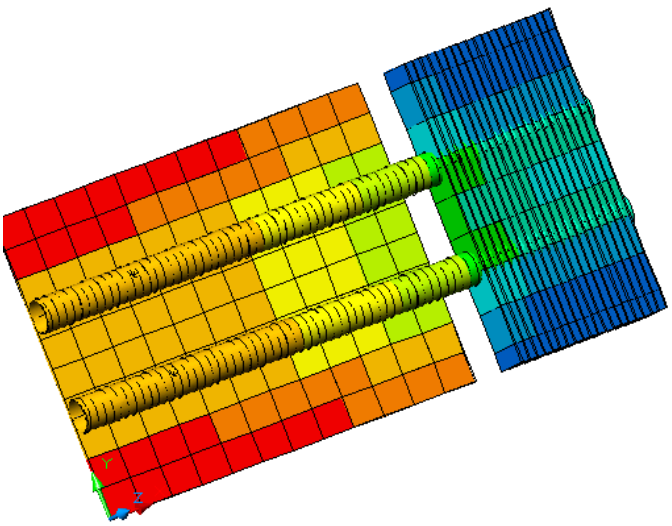

Figure 2: Chip-to-Fin Heat Pipes Modeled as a 2D Cylindrical Shell Attached to Finite Difference Plates

Variable conductance heat pipes (VCHPs) employ noncondensible gas (NCG) reservoirs to limit overall conductance (and therefore power throughput) in order to reduce or eliminate the need for make-up heaters under cold environmental conditions. Gas generation in aging constant conductance heat pipes (CCHPs), which are the most common type used in electronic cooling applications, represents a degradation mechanism for the same reasons: it blocks the flow of the working fluid vapor to the cold wall by forming a barrier through which the vapor must diffuse, and therefore inhibits condensation.

Blockage by noncondensible gases can also be modeled in the network-style approach, but it cannot be accommodated in a "conductive bar" approach. A common assumption is that the gas forms a flat front across the width of the pipe, and that any portion of the condenser covered by the gas is inactive in proportion to that blockage.

For a known amount of gas (usually specified in gm-mole or lb-mole for a degraded heat pipe since the constituents of the NCG are unknown), the length of the blocked portion is calculated using the current saturation pressure corresponding to the temperature of the vapor node: $P_{sat}(T_{vap})$, This pressure to calculate the current mass of the NCG:

$$M_{gas} = \Sigma_i (M_i) = \Sigma_i (V_i * \rho_i) \qquad \text{for all i axial segments}$$

where* $\rho_i = [P_{sat}(T_{vap}) - P_{sat}(T_i)]/RT_i$

This is an iterative algorithm because the current size of the blockage affects the wall conductances G_i, which in turn affect the saturated vapor temperature T_{vap}, which is used to update the pipe pressure and hence size of the gas blockage. The algorithm is complicated by the fact that the gas introduces a nonuniform temperature field, and so the partial pressure of the local working fluid in each blocked or partially blocked section must be taken into account per the above equation. In other words, the warmer the liquid in each blocked section, the less gas will exist in that section. This leads to a requirement for adequate resolution (mesh, discretization) in the anticipated cold (gas-blocked) sections of the pipe.

Despite the apparent complexity, such algorithms are not difficult to write, and have been used for years for modeling both variable conductance heat pipes and gas-degraded constant conductance pipes. The real difficulty lies is in the estimation of the amount of NCG generation that can be expected over the life of a CCHP. This value varies with materials, manufacturing techniques (especially cleaning procedures), and even installation techniques (bending, brazing, etc.). The application engineer is advised to request vendor data, and then to apply healthy conservatism to the date provided given the large uncertainties involved.

The next section provides a specific demonstration of both this modeling technique and the effects of NCG generation, using 1D finite difference elements to represent the heat pipe.

Sample Heat Pipe Application

To illustrate both the application of the heat pipe modeling techniques described above, and to demonstrate the utility of the hybrid 1D fluid – 3D thermal technique, consider the cooling of a 8cm x 12cm PCB board with five dissipative components. Each of the components dissipates 5W, but the only sink available is via natural convection to the air within the compartment. A 8cm x 15cm x 1.27 mm thick alumimum housing wall can be used to double the convective area available, but it is located 8cm away from the PCB board.

To solve the problem without introducing fans, a 1cm diameter copper-water heat pipe is placed between the board and the wall. It is laid underneath a row of chips which represent the hot spots on the board, and makes two 90 degree bends to maximize contact length with each plate. The total length of the pipe is just over 36cm.

Figure 3 shows a parametric study on the affects of gas blockage, from no blockage at beginning-of-life (BOL) to about 8.5e-9 kg-mole of NCG at the end-of-life (EOL). The progression of the gas blockage through the pipe as it ages can be best seen on the lower plate, although the temperature

---

* The perfect gas law has been used for clarity, although real gases densities are equally easy to calculate provided an adequate PVT surface is available for the NCG.

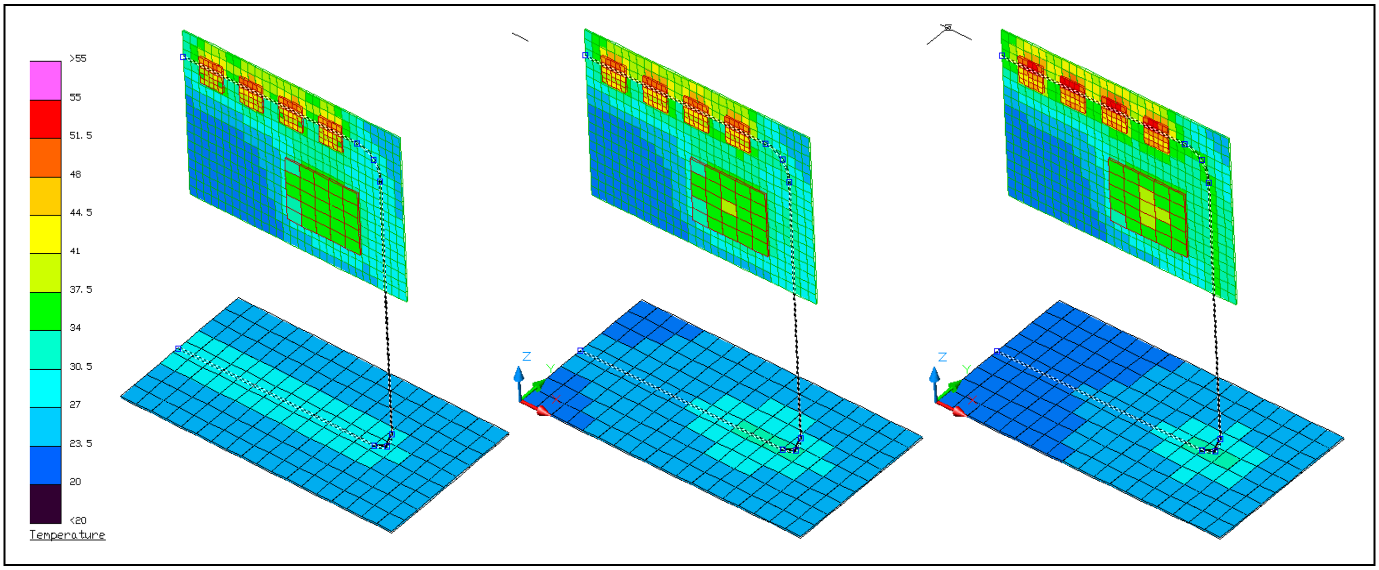

Figure 3: Parametric Study of Heat Pipe Degradation from Zero NCG (left) to 8.5e-9 kg-mole (right)

of the components can be seen to increase as well. Note that the pipe itself is barely visible, and is evident in the Figure only because it is in a highlighted "select" state. This is perhaps a disadvantage of simplified 1D modeling: less physically representative CAD drawings since an abstraction has been made.

Figure 4 depicts the size of the blockage and the corresponding increase in the component temperatures as a function of NCG amount.

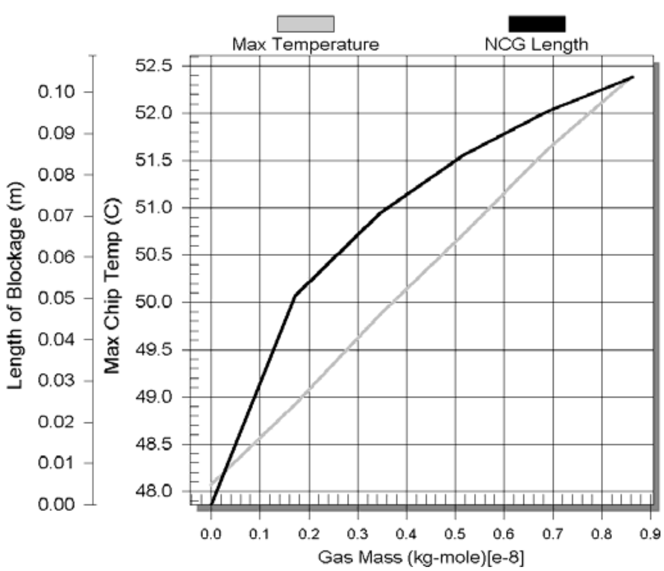

Figure 4: Effect of Degradation via NCG Generation

System-Level Loop Heat Pipe Modeling

Although increasingly of interest to the electrical packaging community, loop heat pipes (LHPs) were chosen as a topic for this paper strictly as a vehicle for discussion of 1D fluid modeling techniques. Single-phase loops and other types of two-phase loops (including vapor compression cycles) could similarly have been chosen for elaboration.

Despite the similarity in their names, LHPs are actually quite different from traditional heat pipes, and the distinctions include modeling techniques, which are completely different. As was shown above, a full fluidic solution is not necessary to simulate the performance of traditional heat pipes, but a more complete thermohydraulic solution *is* necessary to simulate LHPs, even under steady-state conditions.

LHPs operate under the same physical principals as heat pipes, but the separation of vapor and liquid flows into simple, small diameter tubing has significant repurcusions on both their operation and their applications. The isolation of the pumping into a concentrated zone (the evaporator) means not only that flexible, routable lines can be used to form the loop and the especially the condenser, it also means that a smaller pore size wick can be used, effectively eliminating many gravitational and orientation constraints that are otherwise imposed on heat pipes. Unlike thermosyphons, for example, an LHP can operate with the source above the sink.

Simple modeling of LHPs using thermal (resistance/capacitance) networks is inappropriate because two-phase flow and condensation processes exist whose accurate simulation is critical to successful LHP performance predictions. It is very important in loop heat pipes to accurately predict not only the condenser performance (specifically, the subcooling production) but also to track seemingly minor heat gains or losses in the liquid line and the compensation chamber (the large volume colocated with the evaporator), especially at low powers. A "transistor" effect occurs with LHPs: a difference of 1W heating on the liquid line or compensation chamber can easily halve or double the overall loop thermal resistance (which is usually on the order of 0.01 to 0.05 K/W for small devices). A 1W difference in subcooling prediction ($Q_{subcool} = m*C_{p,liq}*\Delta T_{subcool}$ where m is the loop mass flow rate) has similar consequences.

Similarly, tracking pressure drops through the loop is important for the same reason: the sensitivity to heating or cooling of the liquid side of the loop. Perhaps nonintuitively, the overall loop thermal resistance can change as a function of its orientation in gravity because of this effect.

This sensitivity is caused by the fact that there are two saturation conditions on each side of the wick, which is usually metalic (and therefore conductive). An increase in pressure difference across the wick generates an equivalent temperature difference per the Clausius-Clapeyron equation:

$$\Delta T_{wick} = \Delta P_{wick} * v_{fg} / h_{fg} * T_{vap}$$

causing some heat to conduct "backwards" into the core instead of being vaporized:

$$Q_{back} = \Delta T_{wick} / R_{wick} \approx Q_{subcool}$$

Any such "back conduction" plus any liquid line heat leaks must be counterbalanced by increased subcooling production, and increased subcooling production means an decreased active* (two-phase) zone in the condenser, which translates into a increase in the overall loop thermal resistance.

Reference 11 provides a good review of LHP operating principals, and Reference 12 provides a good summary of LHP modeling techniques, so these descriptions will not be repeated here.

Fortunately, despite the apparent complexities of LHP operation, they are not difficult to simulate provided the engineer

1. has acquired relevant performance metrics from the LHP vendor (including critical information such as the total wet wick thermal resistance)
2. has access to a sufficiently detailed two-phase thermohydraulic code that includes at least rudimentary capillary modeling components, and
3. has created a sufficiently detailed thermal/fluid model of the condenser, return lines, and compensation chamber

The focus of this paper is on the last item: the ability to lay out a condenser and route pipes, and to integrate those lines with the thermal model of the structure. This will be the subject of the subsequent example, in which the prior heat pipe-based design is revisited.

Fortunately, LHP performance is relatively insensitive to NCG generation (excepting perhaps start-up considerations, per Reference 13). While start-up (short time-scale) transients can be quite complicated (Ref 14), normal thermally-dominated transients can be easily accommodated provided the two-phase analyzer permits quasi-steady two-phase hydraulics to be combined with transient thermal/structural responses.

* Although heat transfer does occur in the single-phase zone, comparatively little overall heat is rejected in that zone compared to the two-phase (condensing) zone with its orders-of-magnitude higher heat transfer coefficients. The single-phase zone is therefore often refered to as the *blocked* or *inactive* zone.

Sample Loop Heat Pipe Application

To illustrate both a typical LHP modeling application and to illustrate a different use of 1D fluid modeling within 3D FDM/FEM thermal models, the previous heat pipe example will be revisited using an LHP instead.

LHPs cannot completely eliminate the use of the heat pipe, however, unless the conductivity of the PCB were somehow dramatically increased. In other words, a heat pipe is still needed to collect the heat from the dissipating components and transmit that heat to an LHP evaporator. LHPs are not suitable for isothermalizing components, nor can they acquire heat over a large footprint.

LHPs *can*, however, reject heat over an arbitrarily large footprint, and need not be constrained in one plane as must heat pipes. In other words, an LHP can better exploit the available area on the aluminum wall, and this makes up for the introduction of the additional thermal interface resistance between the heat pipe and LHP.

For this design, a single serpentine condenser was used in order to make the task of hermetic sealing easier. Manifolded, parallel passages could have alternatively been used, in which case the two-phase thermohydraulic analyzer must be able to model distribution in parallel legs with very low pressure drops, and must be able to track liquid-vapor interfaces because of the strong effects of gravity on such distributions.

Ammonia was chosen for the working fluid both because of the design maturity of ammonia systems but also because the low vapor pressure of water at these temperatures (3000 Pa absolute at about 25°C) makes it somewhat less suitable for LHPs than for heat pipes. Given the selection of ammonia, copper is no longer available as housing and piping material, so aluminum and stainless steel are used instead, along with a sintered nickel wick. A single continuous run of ASTM B307 4mm (nominal) aluminum tubing (1.9mm ID, 3.2 mm OD) is used for both the transport lines and the serpentine condenser.

Figure 5 depicts the performance of the system, which also includes the final results of the gas-free heatpipe system (described above) for comparison. The evaporator and the compensation chamber are visible as 2D shell elements in the lower right section of the PCB. The evaporator (but specifically *not* the compensation chamber) connect to the PCB isothermalizing heat pipe, which is still present within the board (though not visible in the figure for clarity). This heat pipe no longer serves as the transport device, so it no longer extends past the circuit board.

The saturation temperature for the LHP is approximately 26°C, which is a few degrees cooler than that of the heat pipe design (30°C), but the chip temperatures are approximately the same in both cases since the power transported to the aluminum plate was about the same: a little over 11W. As was expected, the extra thermal interface between the collection heat pipe and the LHP evaporator was compensated by the better exploitation of the aluminum wall plate as a sink. In other words, the serpentine condenser essentially eliminates gradients in that plate (see Figure 5). Such a configuration is

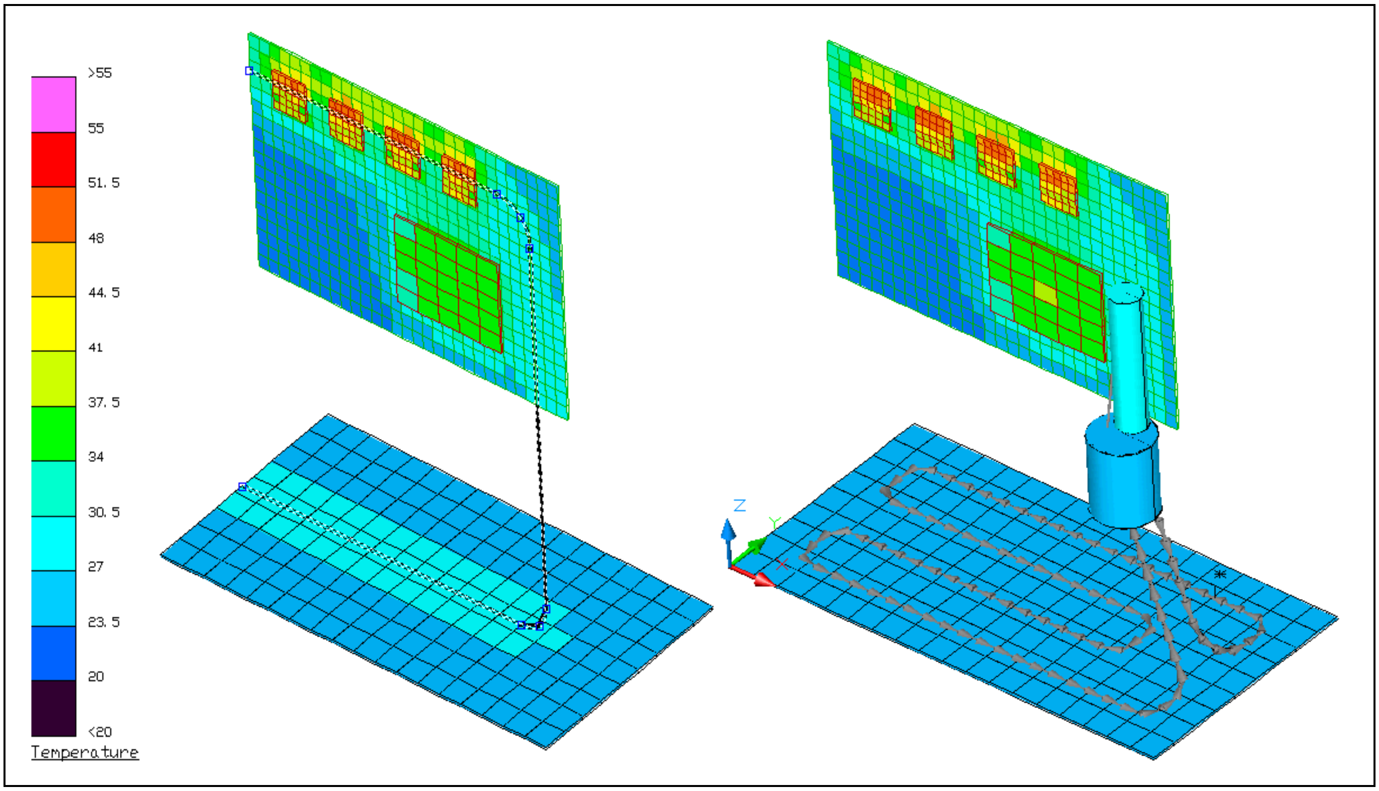

not feasible in a heat pipe because of static pressure differences caused by being out-of-plane.

Figure 5: LHP Replacement System with Serpentine Condenser (Prior Heat Pipe Solution Shown at Left)

The above example is *not* to be misconstrued as a comparison between heat pipes and LHPs, since neither design was optimized against a fix set of requirements. Rather, it was intended to show the utility of including diverse 1D objects within 3D thermal geometry, and to illustrate two specific modeling techniques as examples.

Nonetheless, some pros and cons of heat pipes versus LHPs were introduced, so a brief discussion is warranted. A heat pipe is a simpler and less expensive device than a LHP, and should therefore be selected preferentially, everything else being equal. However, heat pipes are limited in their rejection footprint, and often must be oriented in single planes and with restricted orientations with respect to gravity. Loop heat pipes, on the other hand, can use an arbitrarily complex, small flow area pipe or network of flow passages as the condenser, along with thin, flexible transport lines. LHPs have few gravitational or orientation restrictions. However, they are not as robust with respect to starting up (Ref 13, 14), and the compensation chamber can present an awkward packaging problem because of its intolerance of heating exacerbated by its necessary proximity to the evaporator. LHPs also have restricted heat acquisition footprints because large evaporator sizes and noncompact evaporator shapes represent performance degradations to an LHP due to the previously discussed back-conduction term, which also affects start-up reliability.

Conclusions

Air-cooling of electronics is reaching its limits for all but low-power applications. The successor technologies include heat pipes, vapor chamber fins, loop heat pipes, loop thermosyphons, pumped single-phase coolant loops, spray cooling, and vapor compression cycle refrigeration loops. *All* of these successor technologies are difficult to simulate using 2D/3D CFD techniques: 1D flow modeling techniques are much more appropriate. However, 1D flow modeling techniques were not previously compatible with the widespread used of 2D/3D thermal (conduction/ radiation/capacitance) modeling software.

This paper has introduced a 1D flow modeling tool specifically intended to redress this gap in simulation technology, and has used heat pipe and loop heat pipe examples to demonstrate the concepts involved. The speed of the resulting simulations enables higher-level tasks such as optimization, worst-case scenario seeking, automated calibration to test data, and reliability/sensitivity assessments via statistical design methods.

References

1. J. Wei et al, “Thermal Management of Multiple MCMs with Low-temperature Liquid Cooling,” Proceedings of InterPack ’01 Pacific Rim International Electronic Packaging Conference, IPACK2001-15523, July 2001.
2. S. Downing, “An Integrated Cooler for High Heat Flux Electronics,” IPACK2001-15783, July 2001.
3. G.W. Pautsch, “Overview on the System Packaging of the Cray SV2 Supercomputer,” IPACK2001-15513, July 2001.
4. C.E. Bash, “Analysis of Refrigerated Loops for Electronics Cooling” IPACK2001-15619, July 2001.

5. T. Panczak et al, Thermal Desktop User's Manual, available in PDF at www.crtech.com
6. D. Johnson et al, "Novel Simulation Techniques for Design of Air-cooled Electronics," IPACK2001-15523, July 2001.
7. J. Baumann et al, "Nonlinear Programming Applied to Calibrating Thermal and Fluid Models to Test Data," SEMI-THERM, March 2002.
8. B. Cullimore, "Dealing with Uncertainty and Variations in Thermal Design," IPACK2001-15516, July 2001.
9. B. Cullimore, "Nonlinear Programming Applied to Thermal and Fluid Design Optimization," ITHERM, May 2002.
10. Thermacore's Xeon™ Processor Cooler, http://www.thermacore.com/therma_sink.htm
11. J. Ku, "Operating Characteristics of Loop Heat Pipes," SAE 1999-01-2007, July 1999.
12. J Baumann et al, "Steady State and Transient Loop Heat Pipe Modeling," SAE 2000-ICES-105, July 2000.
13. J. Baumann et al: "Noncondensible Gas, Mass, and Adverse Tilt Effects on the Start-up of Loop Heat Pipes," SAE 1999-01-2048.
14. J. Baumann et al, "An Analytical Methodology for Evaluating Start-up of Loop Heat Pipes," AIAA 2000-2285.