The thermal demands of electronic design

The dense packaging, fast switching of ICs today create so much heat that designers must bone up on old remedies, find out about new ones

by Stephen E. Grossman, *Packaging & Production Editor*

☐ Heat—the inescapable byproduct of every device and circuit—is looming larger and larger as an electronic design problem. With large-scale integration packing ever more active devices on chips, and with circuit boards carrying IC packages in the hundreds, the pressure is on circuit designers to mind their thermal manners as never before. Still another dimension to the problem is added by each advance in device performance: higher switching speeds in digital ICs, for example, or higher output powers in analog circuits.

The prudent designer will regard heat as his implacable enemy. Poor thermal design at the least reduces lifetime, undermines reliability, and degrades system performance. At worst, it can cause catastrophic failure, unsafe products, and costly, even fatal damage.

Thermal management, then, is too important to be left only to specialists. It should be the concern of all electronic engineers right from the circuit concept stage, through the selection of components, materials, into layout and final packaging. In the end, a heat specialist may be required, but his services may be costly if preliminary attention hasn't been paid to heat fundamentals.

David Hegarty is well acquainted with their importance, for as chief applications engineer at Wakefield Engineering Inc., Wakefield, Mass., a company which specializes in thermal management, he often receives pleas for help which arrive too late. Says Hegarty: "Frequently, the circuit designer lays out the whole board—then he comes to us for help in coping with the heat. Too often it is just too late for an efficient and economical design. So the thermal solution may end up being a kluge."

At the other end of the scale are firms like IBM, which, recognizing the critical nature of thermal management to its business, involves teams of heat transfer specialists and thermal designers in the development process. Richard Chu, who manages IBM's thermal development


1. Taking the heat off. Liquid cooling is an increasingly popular way to remove heat from semiconductor devices with high dissipation levels. This cold plate, made by Wakefield Engineering, can dissipate over 600 watts. Note that the coolant lines pass directly under the devices.

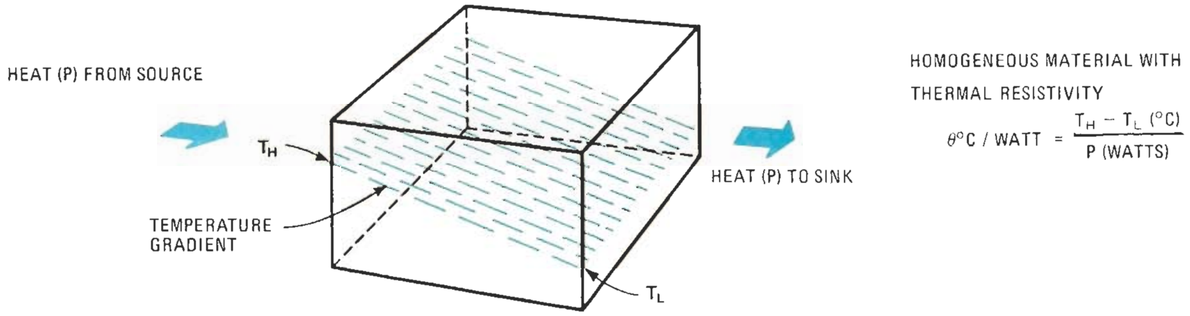

2. Thermal Ohm's law. Components with arbitrary geometries can be assigned a thermal resistance (θ) by determining the heat flow and the temperature drop parallel to the heat flow path and calculating the quotient of temperature to power in degrees centigrade per watt. Thermal resistance plays an important role in solving many thermal problems in electronics.

group at the Poughkeepsie, N.Y., development laboratories, describes his company's attitude: "Here at IBM we look at heat transfer from the very beginning of a design. We also make a point of considering the impact of thermal design on both product performance and reliability. It is not simply a matter of saying 'the lower the temperature, the better.' "

So it behooves the electronic designer to get some feel for the terminology and the units of measurement of this important domain, and to keep up to date with its emerging techniques (Fig. 1). Such knowledge will enable him to include more thermal management in his

How heat hits semiconductors

Because the electrical and mechanical properties of electronic devices are temperature-dependent, it comes as no surprise that both excessive heat and erratic, uncontrolled temperature excursions accelerate semiconductor device failure.

There are two prominent failure mechanisms: thermal mismatch, and hot spots.

Thermal mismatch causes a stress known as shear to develop along a bonded interface between two dissimilar materials when the junction is subjected to temperature variation. It occurs as the result of unequal coefficients of thermal expansion in the two mated materials. As an example, if silicon is bonded to a ceramic substrate at a high temperature, the silicon will shrink less than the ceramic when they cool, and the shear stress that develops may rupture the bond. This accounts for the popularity of epoxy adhesive to bond large silicon dice to substrates because bonding temperatures are low—typically 200°C. (Eutectic die-attach temperatures rise at least twice that high.)

Current crowding is a phenomenon that develops in bipolar devices causing localized device heating. It leads to hot spots that raise temperatures to destructive levels and cause premature device failure. Such crowding across the emitter-base junction occurs because an unavoidable voltage gradient develops across the base, adjacent to the emitter.

One technique that device designers have employed to combat this effect is to configure the base-to-emitter geometry as interdigitated fingers. This maximizes the parameter-to-surface area ratio and counteracts the crowding effect.

A second technique is to ballast the emitter by increasing the resistivity of the paths between the emitter contact metalization and the emitter-base junction as shown. It is accomplished by raising the bulk resistivity of the emitter region. A self-regulating effect results because the bulk resistance response to rising current is to develop a reverse bias which prevents hot spots from developing.

Yet another technique is to enlarge the silicon pellet so that the interface area is larger and provides a broader path over which the heat can flow to the sink. But enlarging the interface of the pellet with the heat sink increases the likelihood of a bond failure because of the thermal match problem. Copper is a good sink material because it has a low thermal resistance, but it is a poor thermal match to silicon. Its coefficient of thermal expansion is about 17.5 × 10^{-6} in./in./°C, much larger than that of silicon at 2.3 × 10^{-6} in./in./°C. Some manufacturers turn to molybdenum, because its coefficient matches silicon's well (3 × 10^{-6} in./in./°C. Then they clad it with copper to reduce the thermal resistivity of the interface with silicon.

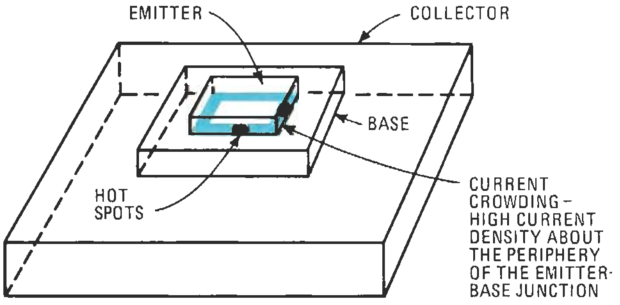

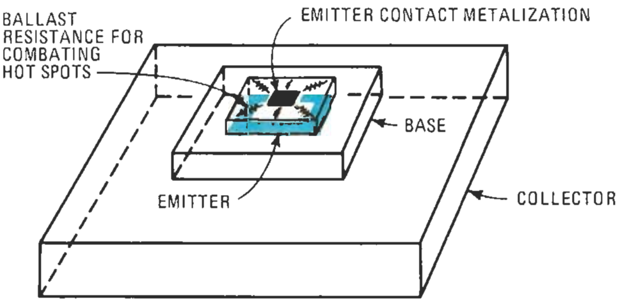

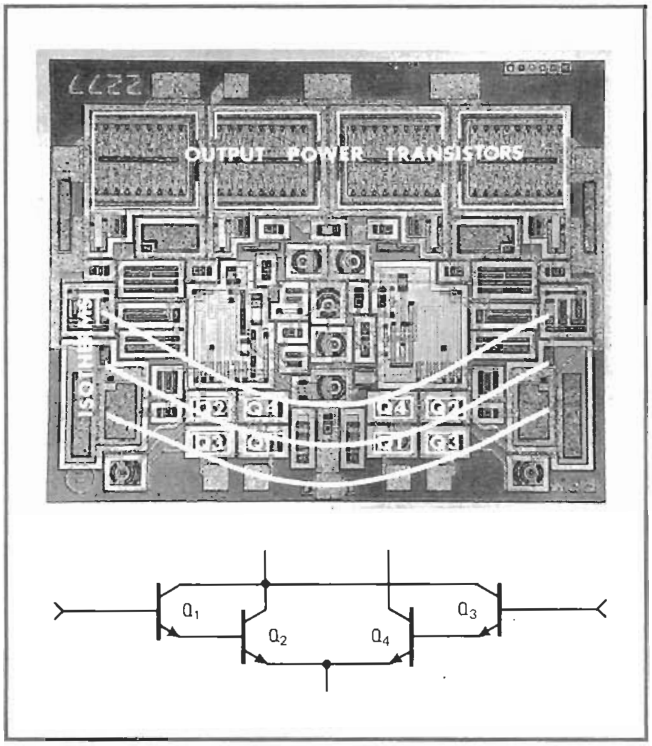

3. Hot lines and cold. Isotherms, shown on this dual power IC chip, are lines that connect points having the same temperature. By judicious placement of the driver devices (Q1-Q4) relative to the isotherms, power gain remains virtually independent of temperature.

future design plans, and will also provide him with the vernacular to communicate with heat transfer specialists should the need arise.

Heat always follows a thermal path that is downhill, so to speak, traveling from the relatively high temperature of the source to the cooler temperature of the sink.

For simplicity, it can be assumed that this thermal path is homogeneous with a uniform cross section (Fig. 2). Under these conditions, the thermal path can be thought of as a thermal resistor, which has the dimensions of degrees per watt instead of ohms. One degree per watt also represents a thermal gradient through which 1 watt of heat may be transferred through the material, and thus also connotes a potential gradient which serves the same role as a circuit voltage. That is, temperature gradient is a potential that determines the power flow, in watts, through a given component.

The concept of thermal resistance then conveniently reduces the analysis of thermal problems to a simple relationship resembling Ohm's law. Thus, in order to raise the heat transfer capability from one point to another, without changing the temperatures of the source and sink, it's only necessary to lower the thermal resistance between the two points as far as possible.

The problem can become complicated, however, when the thermal resistance path is not homogeneous. Then the over-all gradient becomes a function of the composite thermal paths.

The levels

It is commonplace in packaging to speak of separate packaging levels: the device, the surface which supports the device—the printed-circuit board or the backplane—and the enclosure: This division also lends itself to analysis of the thermal aspects of packaging.

IC devices come into consideration because device designers are packaging more and more active circuits into the device package. Though power levels for the commonly-used families of logic gates are in the 10- to 100-milliwatt range, the problem is simply that putting thousands of these gates into a small IC package translates into extraordinarily high power densities never before encountered in electronic packaging. Since silicon cannot withstand operating temperatures much above 150°C, the device designer cannot allow the junction temperatures to rise above the device limit. So he must lower the thermal resistance of the path out of the package. This means either altering the geometry, or selecting packaging materials with lower thermal resistances, or doing both.

Geometry is also the key to assuring stable power-IC performance independent of temperature, for it enables device designers to take into account the fact that gain is dependent on temperature. Figure 3 illustrates the device geometry of a Sprague parallel power amplifier chip. The output power transistors are identified near the top of the chip. Isotherms, which are lines connecting points having the same temperature, have been drawn through the region at the bottom of the chip, where the Darlington-connected driver pairs are located. Note that transistor Q_4 lies on the isotherm nearest to the power transistors. This transistor is the hottest. Transistor Q_3 is a degree or so cooler, while paired-transistors Q_1 and Q_2 are at about the same temperature. However, in both cases the transistor pairs are equidistant from the center isotherm. Consequently, the gain of the Q_1-Q_2 pair is about equal to the gain of the Q_3-Q_4 pair, and over-all gain of each amplifier section remains essentially independent of temperature variation.

The next level of thermal concern is the printed-circuit boards and the mechanically wired backplanes. Here the predominant mode of cooling is convective air.

In free convection the boards are positioned in a vertical attitude so that warm air rises along the board surfaces carrying the heat away. Free convection is economical, because no air-moving equipment is required, and it is also reliable, which explains its popularity where low maintenance costs are crucial—in telephone companies, for example.

But the higher density of modern electronics needs a faster rate of air flow than 0.5 linear feet per second—which is about the highest that can be expected in free convection. Consequently, forced air cooling has become commonplace, and today fans and blowers are performing two principal functions: cooling localized hot spots such as high-power transistors, and flushing hot air from an equipment enclosure, so that the ambient temperature is held below a specified design limit.

But even forced air cooling is reaching its limit for localized cooling as digital designers drive devices at higher switching speeds and raise the power they dissipate. As a result, liquid cooling is being used more and more.

Bill Allen, a manager of large computer circuit design, at Burroughs Corp., Paoli, Pa., points out some of the arguments in its favor: "Temperature variation

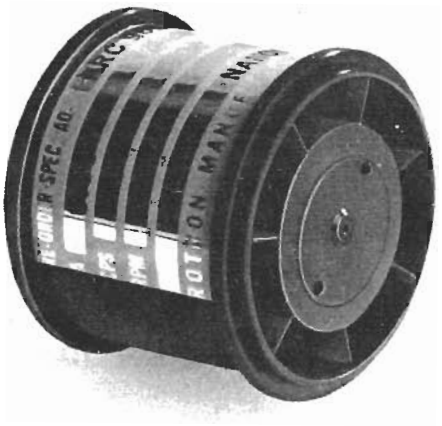

4. Breezes to order. This fan weighs only an ounce, fits in a 1-inch cube, and is well suited for microelectronic applications. Rotron, the manufacturer, rates life expectancy in excess of 20,000 hours. Device delivers up to 8 cubic feet per minute.

within a mainframe enclosure plays havoc with both noise margin and propagation delay. But by switching to liquid cooling such variations are easily minimized. Liquid also enables the size of the power supplies, a major source of heat in most large systems, to be reduced. Also, air-moving systems often create a lot of acoustic noise, and liquid cooling rids the enclosure of blowers and fans."

Cold plates are heat-exchanging devices which transfer heat from an electronic device to a moving liquid. The unit shown in Fig. 1 dissipates a lot of power—600 w developed by 21 TO-3 devices and two TO-66 devices. Thermal resistivity from device case to coolant is about 0.3°C/w per transistor, and the coolant flow rate is 1½ gallons per minute.

Cold plates can be operated as either an open-loop system, which would be a simple tap-to-drain arrangement with flow provided by the local water supply, or a closed-loop system. The closed-loop system, though more costly, enables the coolant to remain free of contamination. Frequently, a deionization system is employed which removes impurities from the coolant and also prevents foreign materials from blocking the coolant lines.

Closed-loop systems normally employ a heat exchanger to cool the fluid. The heat exchanger is a device which functions very like an automobile coolant system. A pump maintains the coolant flow. The hot coolant from the cold plate is passed through a radiator which is cooled by a forced air flow. The coolant is then returned to the cold plate.

Tom Coe, president of Wakefield Engineering points out: "Computer designers are rediscovering the heat exchanger, a familiar friend to engineers who have worked in induction heating and high power rf."

Heat pipes

Still waiting in the wings is the heat pipe, which has yet to take hold as a major participant in commercial heat transfer applications even though the technique dates back to a 1942 patent. Heat at one end of a sealed pipe causes liquid inside that end of the pipe to vaporize and travel to the opposite, cooler end, where it condenses. The liquid makes the return trip to the hot end by being drawn by capillary action along a wick lining the pipe [*see* "Heat Pipes—A Cool Way to Cool Circuitry," *Electronics*, Feb. 16, 1970, p. 94].

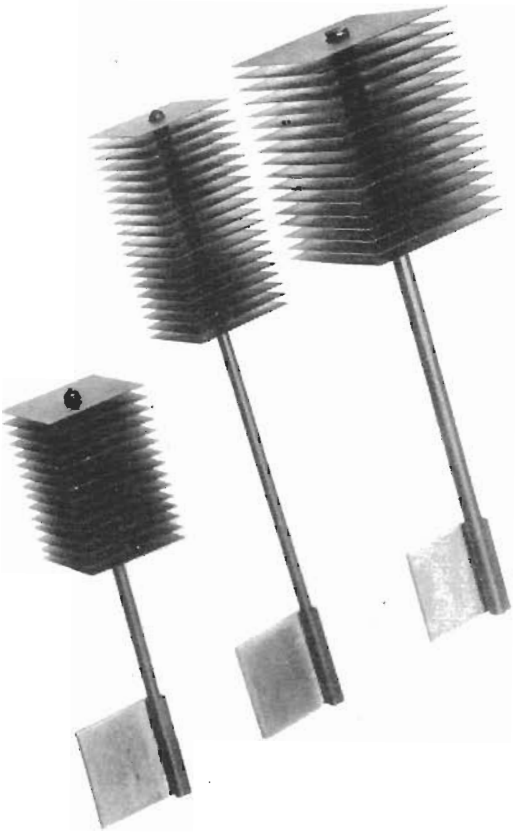

5. Heat pipes. These units are designed to cool discrete devices which mount on the tabs at bottom and transport heat to the finned radiator at top. Manufactured by Jermyn, they are among the first off-the-shelf products to employ the heat-pipe principle. The advantage of heat pipes is that they provide thermal resistivities several orders of magnitude lower than the best thermal conductors.

Thermal experts are enthusiastic about heat pipes because these devices function with a very small temperature drop—several hundred times less than the best thermal conductors. Negligible temperature drop permits a designer to position a sink in a remote location and then use the heat pipe to transport the heat from a pad to a finned radiator (Fig. 5). The pads on these assemblies are drilled by the user and the semiconductor devices are secured in place.

First of a series

The article which begins on page 102 is the first in a series on thermal design. It discusses techniques for raising the power capability of plastic-packaged power ICs.

In the next article in the series, Forest Golden of General Electric will discuss the thermal problems encountered by engineers who design and use discrete power semiconductor devices. Subsequent articles will explore thermal topics of concern to device, circuit and package designers. □