Liquid cooling safeguards high-power semiconductors

Liquid coolants require less heat-sink volume than forced-air systems to carry away heat from kilowatt-level circuits

by John A. Gardner Jr., *Wakefield Engineering Inc., Wakefield, Mass.*

□ Liquid systems have seldom been used to cool semiconductors in the past because convective-air cooling has usually proven adequate at low power levels. But as semiconductor rectifiers and thyristors grew in power capability, design engineers were confronted with the problem of cooling devices that dissipate hundreds and sometimes thousands of watts.

When semiconductors were first used, it seemed that large heat sinks, together with high-capacity fans and blowers could fulfill the most extreme cooling requirements. However, as power levels and device sizes were increased, space requirements for cooling grew exponentially and demanded precious space in the electronic packages; designers, consequently, turned to liquids.

Liquid provides a larger margin of reserve cooling power than other cooling techniques to cope safely with peak loads and transient conditions because thermal inertia enables fluid to absorb momentary heat pulses with only a slight temperature rise. Liquid cooling also minimizes acoustic interference, a persistent problem when cabinets are air-cooled. Noise can be readily abated by locating the heat-exchanger and pumping equipment at a distance from the electronic components being cooled.

If a designer had his choice, he would select natural convection cooling to reap the benefits of cost and reliability that must be sacrificed by adding the electromechanical components required for forced-air and fluid systems. However, costs and complexity of liquid-cooling systems cannot be denied. And many engineers

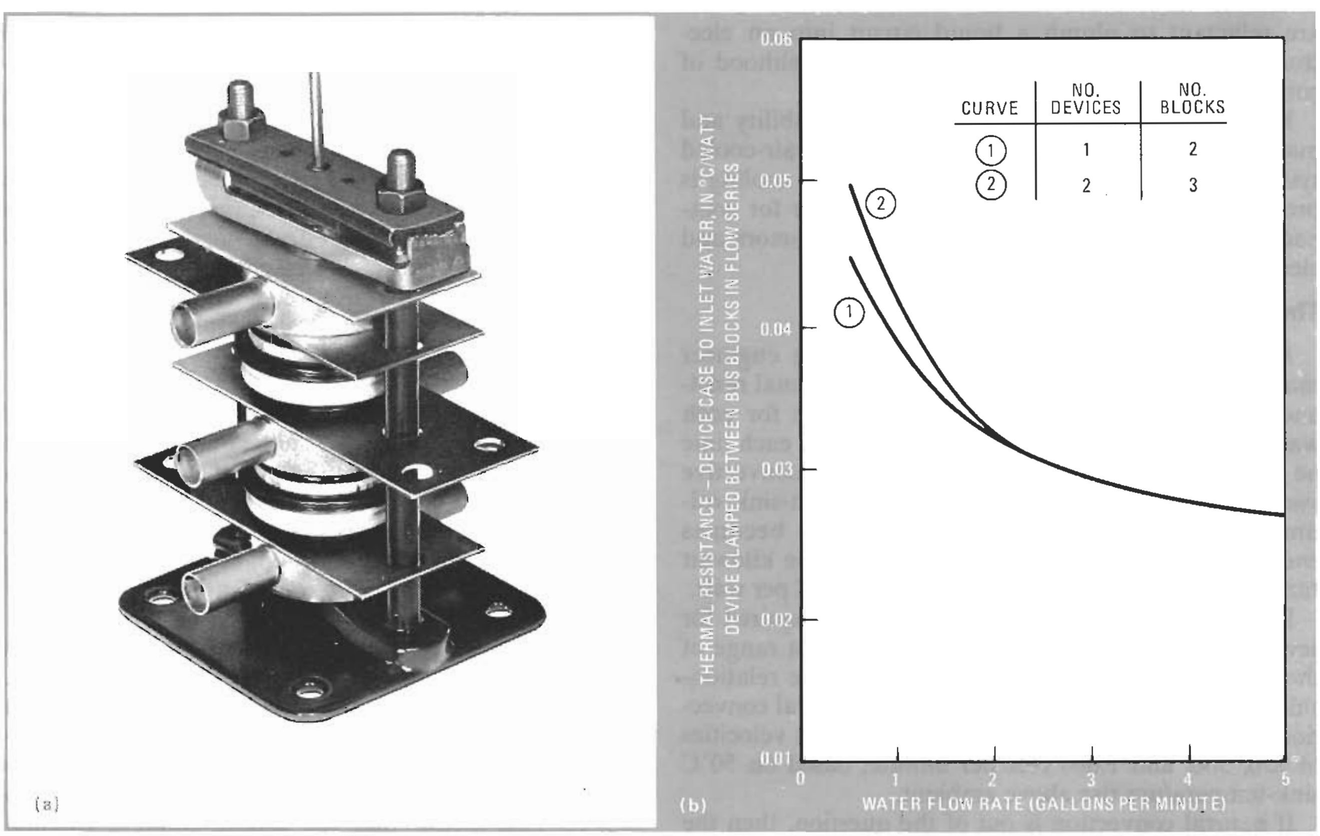

Two stories. Stacking two pressure-mounted devices, such as semiconductor rectifiers and thyristors, in a liquid-cooled assembly (a) is a compact technique for cooling kilowatt-level devices. Thermal resistance from the case to the coolant is low at low flow rates, and it diminishes little at flow rates above two gallons per minute (b). Adding the second device (curve 2) raises the thermal resistance a trifle (curve 1).

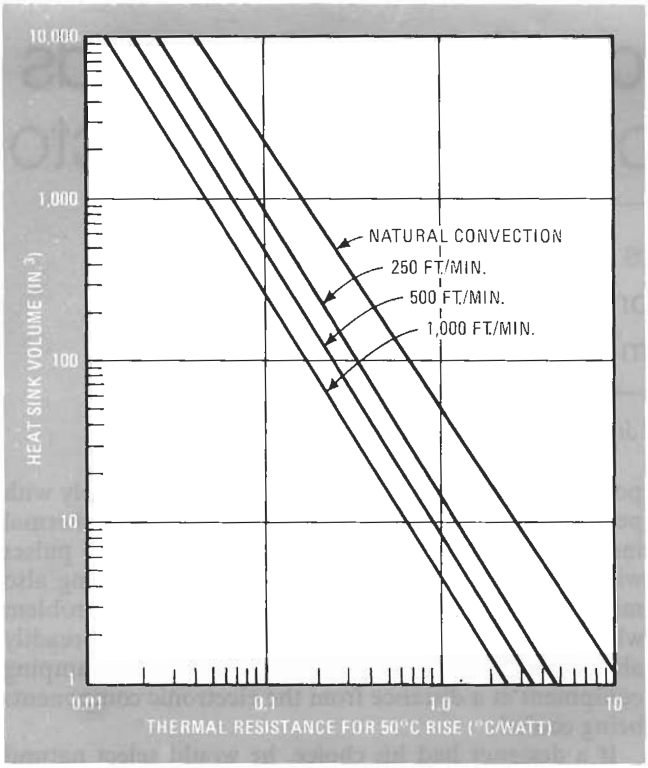

1. Gobbling space. The volume of an air-cooled sink grows enormously as a designer lowers the thermal resistance. These curves indicate the various sink volumes for cooling by natural convection and forced air. However, liquid-cooling requires much smaller sink volumes—desirable for compact electronic packaging.

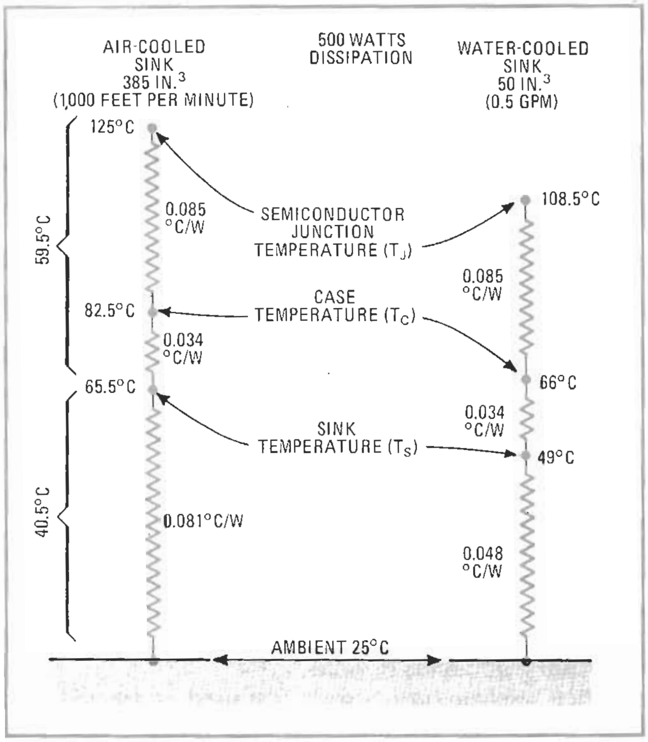

2. Cool the junction. Holding the semiconductor junction below a limiting temperature is crucial to thermal cooling—a key factor in long-term reliability. Here, circuit equivalents have been drawn for both an air- and a water-cooled sink dissipating 500 watts. The water-cooled system keeps the junction almost 17°C cooler and requires about one seventh the volume of the air-cooled sink.

are reluctant to plumb a liquid circuit into an electronic-equipment cabinet because of the likelihood of corrosion, leakage, and condensation.

Moreover, factors such as component reliability and maintenance demands all weigh in favor of air-cooled systems. But despite these drawbacks, liquid cooling is proving to be a highly satisfactory technique for compact, silent cooling of high-power semiconductors and electronic systems.

Thermal resistance

As power dissipation rises, the packaging engineer must whittle away at the case-to-ambient thermal resistance (Θ_{C-A})—the temperature rise in degrees for each watt of power transferred. As a rule of thumb, each time he halves the thermal resistance in a natural convective system, the designer must quadruple the heat-sink volume. Obviously, the demand for space becomes enormous when dissipation levels rise into the kilowatt range and thermal resistance falls below 0.1°C per watt.

Figure 1 illustrates the envelope volume required for several widely used shapes of heat sinks for a range of thermal resistances. The four curves show the relationship of volume to thermal resistance for natural convection cooling and forced-convection cooling at velocities of 250, 500, and 1,000 feet per minute, based on 50°C sink-temperature rise above ambient.

If natural convection is out of the question, then the designer must turn either to forced air or liquid cooling, and both are feasible in the range of 500 to 1,000 w. However, when volume is a crucial consideration, liquid has the edge. A forced-air cooling system requiring 500 cubic inches of heat-sink volume can't compete with liquid, which can deliver the same cooling capability in a sink volume from 60 to 120 cubic inches—an improvement of a full order of magnitude.

Whereas air-cooled systems require careful analysis of localized ambients within equipment cabinets to avoid interactive heating effects, analysis for liquid-cooling systems is relatively straightforward.

The arguments favoring cooling with liquids at higher-power levels emerge more clearly in an example.

A high-power example

Consider a 500-w pressure-mounted semiconductor rectifier with a maximum junction temperature of 125°C and junction-to-case thermal resistance, as shown on the manufacturer's data sheet, of 0.085°C/w. The case-to-sink thermal resistance is determined from experimental data to be an additional 0.034°C/w. Adding the two thermal resistances and multiplying by 500 w yields a rise of 59.5°C between the sink and junction.

These thermal resistances and the temperature rises across the resistances are shown schematically in Fig. 2. Such sketches help the designer to visualize how each portion of the thermal path contributes to the rise above the ambient temperature. If the ambient is assumed to be 25°C, the rise from ambient to sink is limited to 40.5°C. This means that the sink-to-ambient thermal resistance can be 40.5°C/500 w, or 0.081°C/w, at most. Figure 1 discloses that this thermal resistance requires a sink volume of about 385 cubic inches if air flow is to be

1,000 feet per minute. This requirement can be satisfied by two sinks of about 3 inches by 7 in. by 7.5 in.

By contrast, consider the requirements for liquid-cooling the same semiconductor rectifier. Figure 3(a) shows the device clamped between liquid-cooled blocks. Thermal resistance values for blocks of this type are shown in Fig. 3(b). At a flow rate of 0.5 gallon per minute, the sink-to-inlet water thermal resistance of the water system is 0.048°C/w—about half the value of the previously calculated forced-air system.

As for volume, the sinks occupy about 50 cubic inches, a mere 13% of the sink size required in the moving-air system. Liquid flows from the lower coolant block to the upper one in flow series, resulting in only a slight warmup of the water passing through the lower block. This heating is of little consequence.

Series flow

Stacking of devices, as shown on page 103, adds only slightly to the volume and causes little degradation of the case-to-inlet-water thermal resistance for the water system. At the same rate of 0.5 gallon per minute, cooling a second device in series degrades the thermal resistance only 0.005°C/w—the difference between plots 1 and 2 in the performance graph on page 103. Plot 1 is the case-to-inlet-water thermal resistance of a single pressure-mounted semiconductor. Plot 2 is the thermal resistance for each of two devices mounted in a stack so that both devices share a common pole piece.

Plots 1 and 2 are virtually coincident because the conductive path from the devices to the coolant-pole pieces is far more efficient than the thermal path between adjacent semiconductor devices. However, this is seldom the case in air-cooling systems.

Liquid cooling is also attractive for cooling groups of lower-power devices, such as semiconductor devices in TO-3 cases that have power levels in the range of 50 to 150 w, as shown in Fig. 4. The channel-plate cooler (Fig. 5) is designed as an inexpensive arrangement for cooling both stud-mounted and bolted semiconductor devices. Sink thermal resistances range from about 0.6 to 0.25°C/w, depending on the center-to-center spacing of the devices and coolant-flow rate.

On the other hand, the high-density cooler shown in Fig. 4 is virtually unaffected by thermal interaction because the devices are located directly over the coolant-flow lines, thereby optimizing the thermal path to the fluid. This arrangement is well suited for cooling large numbers of such smaller devices as those mounted in the popular TO-3s. This cooler, which measures 6 in. by 7 in. by 1 in. and occupies a volume of 85 cubic inches, is capable of dissipating 2 kw if the coolant-inlet temperature is 40°C or below.

Table 1, which lists the thermal resistivities of various liquid-cooler geometries, shows that the pressure-mounted assemblies offer thermal resistances an order of magnitude lower than the channel-plate mountings and the high-density cooler. However, the latter are more than adequate for clusters of lower-power, smaller devices. The pressure-mounted coolers offer an additional advantage. Since the bus plates and the bolts are cooled along with the device, the bus-current rating can be higher than in an air-cooled configuration. In general, thermal resistances diminish substantially as flow rates are raised to 2 gallons per minute, but not beyond.

When a semiconductor must be electrically isolated from a liquid-cooled sink, an interface material, such as beryllium oxide, offers high thermal conductivity and excellent electrical isolation. The penalty is a rise in over-all thermal resistance, caused by an increased interface thermal resistance, as indicated in the second column of Table 1.

Electrical isolation between pressure-mount cooling

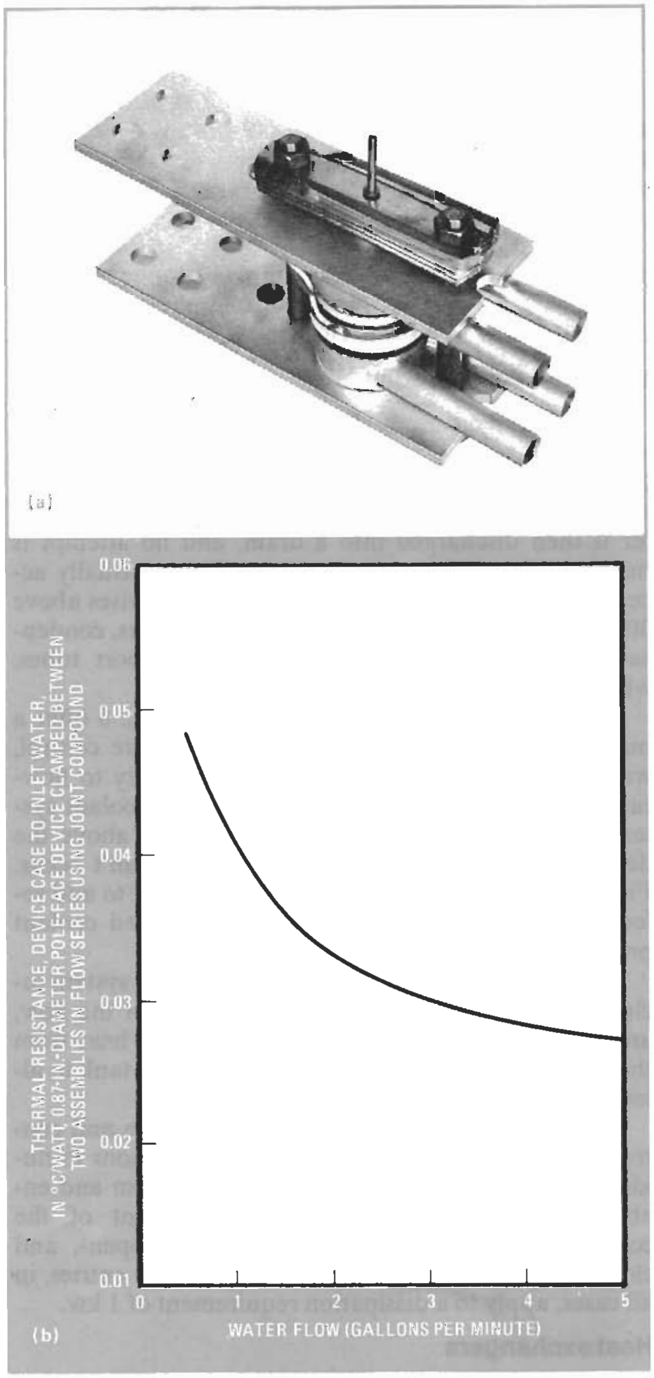

3. Cool cooler. The liquid-cooled solid-state copper blocks sandwiching the semiconductor device (a) provide two thermal paths (top and bottom) and dissipate 2 kilowatts. Clamp puts an 800- to 2,000-pound bite on the semiconductor device to assure a low thermal resistance—on the order of 0.034°C per watt at the interface (b).

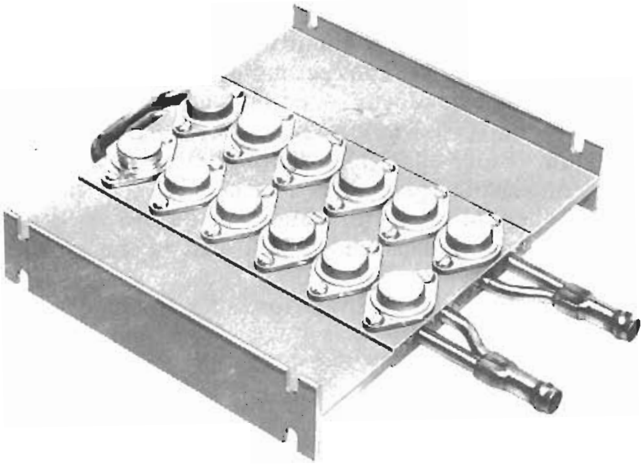

4. **Dense package.** This configuration is well suited for cooling TO-3 packages because the devices are placed directly over the coolant lines, ensuring short, efficient thermal paths to the coolant. Unit can dissipate 2 kilowatts, but occupies only 85 cubic inches.

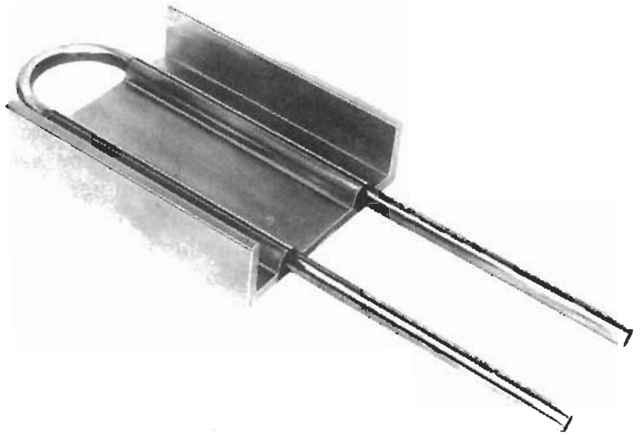

5. **Channel chiller.** This economical U-shaped aluminum plate, priced at $16, can cool a small number of relatively high-power devices. Thermal resistance varies from about 0.25°C/W to 0.6°C/W, depending on spacing between devices and coolant-flow rate.

blocks can be achieved with rubber liquid-transport tubing. A good rule of thumb is to employ one foot of tubing for each 1,000 volts of potential difference.

Open and closed loops

Liquid systems are commonly designed in an open-loop configuration in which tap water is fed to the cooler or cold plates through a pressure-reducing station, which ensures a constant flow rate. The heated water is then discharged into a drain, and no attempt is made to control water temperature. This is usually acceptable if the inlet water temperature never rises above 30°C. However, during humid summer months, condensation can form on cold plates and transport tubes, which may be troublesome.

By contrast, the closed-loop systems of Fig. 6 offer a number of advantages, including temperature control, water conservation, and reduced susceptibility to flow-rate variation. Moreover, by operating the coolant system so that the water temperature remains above the dew point, condensation on cold surfaces can't occur. Finally, a closed-loop system enables the user to add selected solutions to the water to attain desired coolant properites.

The hardware components of a closed-loop system include a cooler, a circulating pump to sustain the flow, an air-liquid heat exchanger to transfer the heat from the liquid to the surrounding air, and storage tank to allow for expansion.

The storage tank permits normal expansion and contractions that accompany temperature variations in fluids. It is also a deaeration point for the system and enables periodic sampling and replenishment of the coolant. The relative costs of forced-air, open-, and closed-loop systems are listed in Table 2. The entries, in all cases, apply to a dissipation requirement of 1 kw.

Heat exchangers

The air-liquid, high-efficiency heat exchanger with an attached fan, shown in Fig. 7, is typical of a compact series of exchangers suitable for closed-loop cooling of electronic components. Copper and brass lines are commonly used to carry the coolant and provide long-term high performance with most heat-transfer fluids.

This type of cooler is available with either single or double-pass flow on the water side. With double-pass flow—this means that the water makes a round trip through the exchanger region—thermal performance is enhanced. Also, inlet and outlet fittings mount on the same side of the exchanger, which is frequently a convenience. However, a larger-capacity pump is required to cope with the increased pressure drop that is characteristic of the double-pass system.

The fan of the illustrated exchanger draws the air through the exchanger core before the air passes through the fan itself, which produces even heat distribution across the core. Thus, the operating temperature of the fan assembly, including the bearings, will be elevated above the ambient temperature. And since the life of fan motors depends on their operating temperature, the temperature of the air leaving the core is a critical parameter in length of fan life.

Designing a liquid-cooled system

Generally, design requirements of a liquid-cooled system are less complex to compute than those of an air-cooled system because the string of thermal resistances from the device case to the coolant loop is less critical. That is because the thermal capacity of the liquid-cooling loop is large enough that the interacting secondary resistances among devices play a negligible role. This is not true when air-cooled sinks are employed.

Once the designer knows the power dissipation required in a cooler or cold plate and selects a flow rate, he can readily determine the rise in the cooler's water temperature by using the alignment chart in Fig. 8. As an example, for 1,000 w and eight gallons per minute, the rise is less than 1°C.

If the eight gallons per minute were to be split equally among four cold plates, each dissipating a kilowatt, the water temperature rise would only be 2°C. However, a careful analysis is mandatory because, in some cases, the temperature rise may be substantial.

Here are the parameters required to determine the thermal resistance (Θ) of a heat exchanger:

■ Total power dissipated by the components that need cooling (P).
■ Temperature of the water entering the heat exchanger ($T_{water\ in}$). The temperature drop from the cold plates to the heat exchanger should be subtracted if it is not negligible.
■ Ambient air temperature ($T_{air\ in}$).

These values enable the designer to calculate the thermal resistance of the heat exchanger:

$$\Theta = (T_{water\ in} - T_{air\ in})/P \text{ in } °C/W$$

Once this thermal resistance is determined, the designer should check performance curves for various heat exchangers. These curves show that the flow rates of both the water and the air govern the performance of the exchanger.

It is likely that more than one type of heat exchanger will fulfill the cooling requirement. Selection can be narrowed by examining such factors as available space and position of inlet and outlet fittings. Finally, the pressure drop of the exchanger and all coolers, lines, and fittings must not exceed the pump capacity.

Selecting a pump

Once the heat exchanger and the coolers are selected, the drop in pressure through the cooler plates, the heat exchanger, and all interconnecting tubing and fittings is summed to determine the total head that must be deliv-

TABLE 1. LIQUID COOLERS – SOME REPRESENTATIVE THERMAL RESISTANCES

Type	Thermal resistance per device case-to-inlet water at 2 gallons per minute (°C/watt)
Liquid-cooled bus, 0.87-in. device interface, cooled on both ends	0.033
Same as above, but electrically isolated from the cooler	0.078
Channel plate 1 in.-diameter-devices, 2 in. center-to-center spacing, 0.038°C/watt case-to-sink impedance	0.38
High density cooler, TO-3 case style, thermal grease on device interface	0.30

TABLE 2. COMPARISON OF COOLING SYSTEM COSTS

System	Devices	Approximate cost
Forced-air convection	Two, pressure-mounted	$75
Liquid-cooled bus blocks, open-loop system	Two, pressure-mounted	$58
Liquid-cooled channel plate, open-loop system	Four, stud-mounted	$16
Liquid-cooled, high-density cooler, open-loop system	12, bolted	$21
Closed-loop system, cost to be added to liquid systems above	—	$138

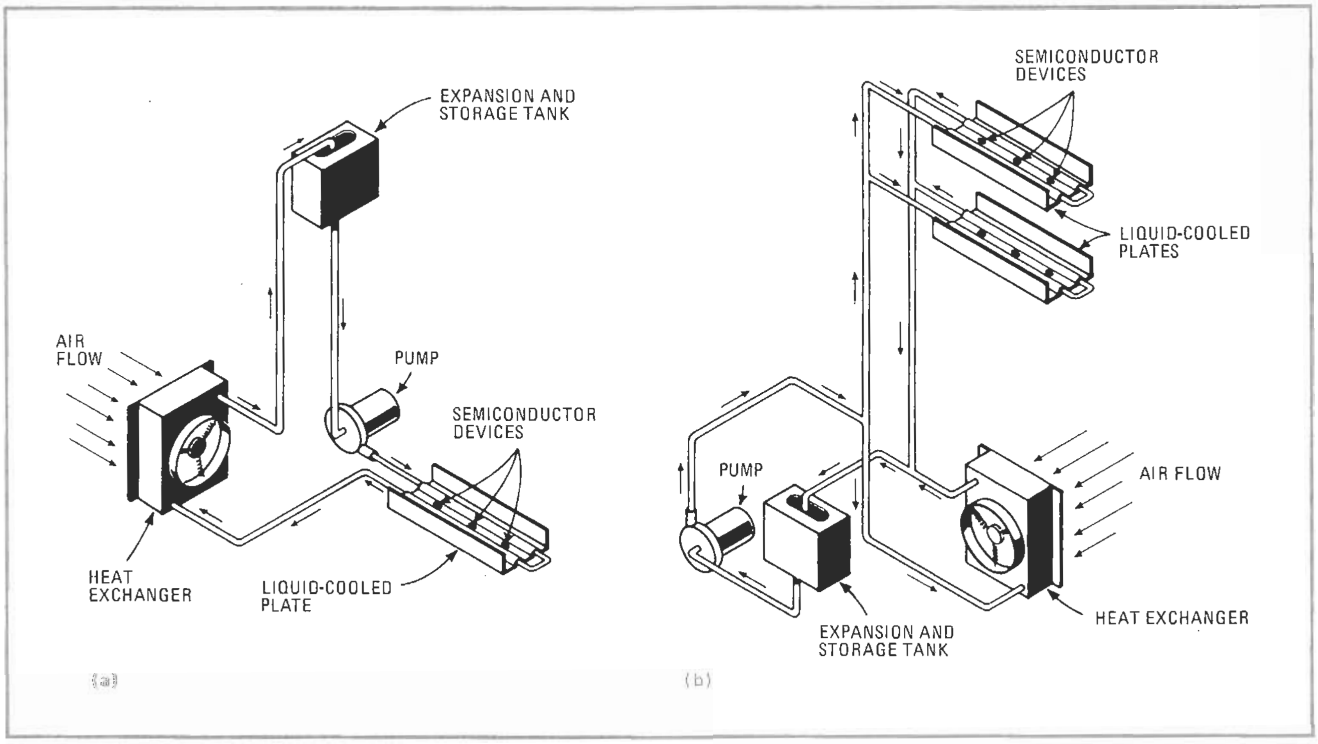

6. Flow system. A simple series-flow system is well suited for cooling a single cold plate (a). Connecting two or more liquid-cooled plates in a parallel-flow system (b) reduces the pressure drop so that a large-capacity pump isn't needed.

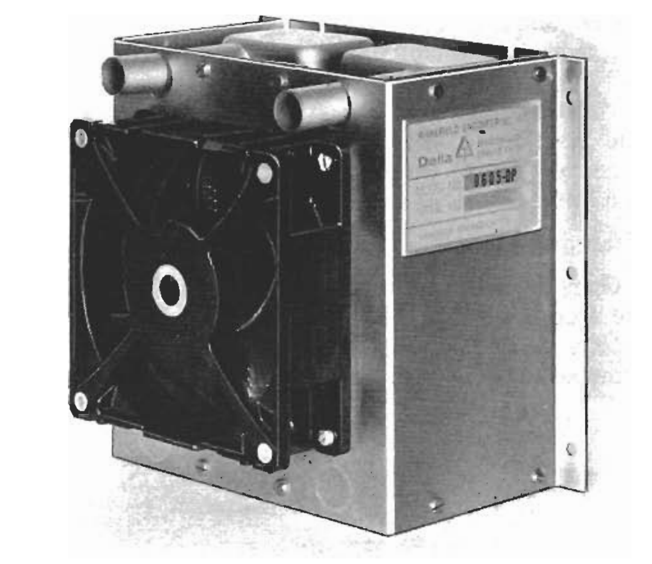

7. Cool exchanger. This double-pass heat exchanger transfers heat from the entering liquid to an air stream that is driven through the exchanger by a fan. The liquid enters and exits through the fittings above the fan. Raising the air and liquid-flow rates improves thermal performance. However, to minimize erosion and corrosion, liquid-flow rates through coolant lines should not exceed 10 feet per second.

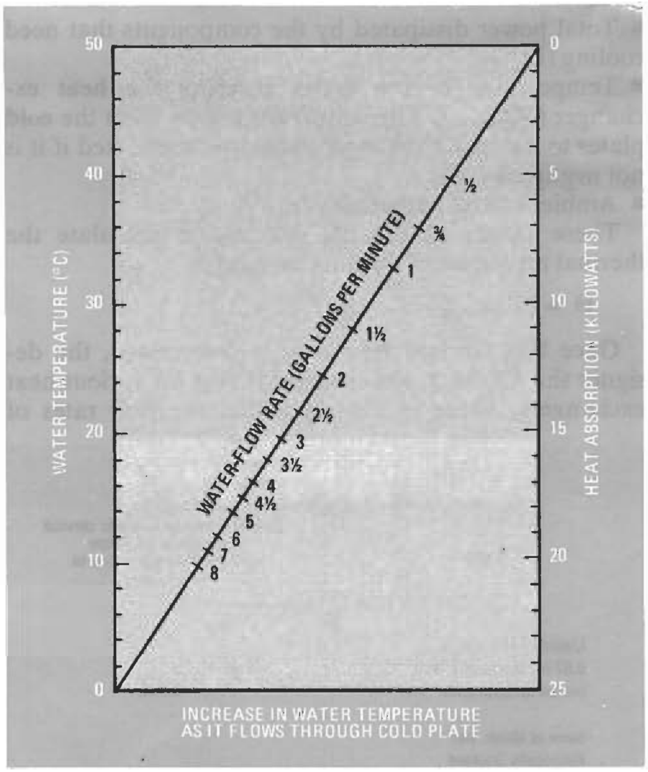

8. Warmup. By aligning a straight edge with values of the heat absorption and the water-flow rate in a cold plate, the temperature rise can be read from the scale at the left. A similar chart, found in heat-exchanger manuals, enables the designer to determine the temperature rise through a heat exchanger.

ered by the circulating pump. Head is the pressure, in pounds per square inch, delivered by a pump at a specified flow rate. Heat-exchanger manufacturers usually supply this necessary data. During the past several years, a number of centrifugal pumps have been developed to operate without rotating shaft seals so that long-term leakproof operation is assured. Flow capacities are sufficient for most electronic-package cooling.

If the pressure drop through a closed-loop system appears to be excessive, the parallel-flow system of Fig. 6(b) may be suitable. The advantage is that, unlike a series-flow system, the flow rate through each cold plate is not necessarily the same as the flow rate through the heat exchanger. And since drop in pressure through cold plates is usually much higher than it is through a heat exchanger, parallel connection limits pressure drop without significantly degrading cooling performance. Moreover, the flow rate of the heat exchanger can remain at a relatively high value, assuring high performance as a result of low thermal resistance.

Reliable transport of a liquid demands careful attention to both flow velocities and the materials contacting the fluid. Although copper tubing is relatively expensive, it offers the best envelope because the smooth wall surface resists corrosion in most environments. Copper also conducts heat well and resists the mechanical erosion and the chemical corrosion which are most severe at such points of high turbulence as sharp bends. Copper tubing, which is also easy to install, offers a good electrochemical and thermal match with other materials commonly employed in heat-exchanger and cold-plate construction.

Selecting the fluid

Water offers the best over-all coolant characteristics in terms of density, viscosity, thermal conductivity, and heat capacity. In closed-loop operation, where control of the content of the circulating fluid is possible, additions to compensate for losses of fluid can be made from time to time. Water that has been distilled, deionized, and demineralized provides the most efficient long-term performance. When both aluminum and copper-brass metals are present in a fluid circuit, specially—inhibited ethylene-glycol solutions can prevent deterioration of the fluid passages. However, because of their lower thermal conductivity, they do degrade thermal performance. Solutions of this type are mandatory where the ambient temperature can drop below the freezing point of water, or where surface temperatures exceed the boiling point of water.

Exotic dielectric oils are employed where severe electrical-insulation requirements team with freezing temperatures. Unfortunately, many of these oils, especially the chlorinated series, place severe demands on pump seals and plumbing joints in the fluid circuits. Again, even the best dielectric oils, as well as the series of silicone oils, require higher-performance heat exchangers than do water-cooled systems.

There are a number of variations in liquid-cooling systems, and one is the cold-sump system. This technique employs a refrigerant loop to cool a reservoir of refrigerated water, which is then circulated through a closed-loop cooling system. Cold sumps usually have a large cooling capability and may serve a number of heat loads simultaneously at remote locations. They are frequently selected for large complexes like computer installations. □