5

COOLING

Heatsink Cooling

All electronic circuits discharge heat generated by the movement of electrons. As operating frequencies increase, thermal loads go up (assuming die sizes and fabrication techniques remain constant). Given these facts, a successful overclocker must pay attention to cooling systems and sometimes devise external solutions dedicated to dissipating heat and regulating the processor's temperature.

Figure 5-1: Forced-air heatsink cooler

Most common thermal regulation solutions are built around the concept of forced-air cooling. The standard processor cooler includes a massive metal heatsink and a high revolutions-per-minute (RPM) fan to dissipate heat through convection. Forced-air coolers represent the most cost-effective solution for the widest range of system platforms; the parts are simple in design and readily available. Many solutions are possible, but a heatsink cooler is still the best choice for most overclocking scenarios.

Size, density, shape, and material influence a heatsink's ability to dissipate heat. The best coolers available are often fabricated entirely of copper, with its excellent thermal transfer properties. Copper is also the most expensive option, due to high procurement and manufacturing costs. Aluminum is the most widely used heatsink material. It provides acceptable thermal conductivity at a relatively low cost. At the least expensive end of the heatsink spectrum is cast iron. Iron is a poor thermal conductor, and these coolers should be avoided.

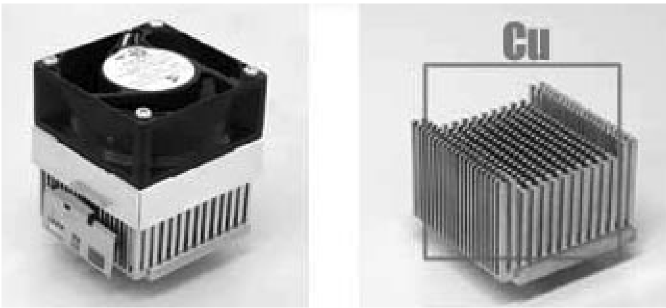

Figure 5-2: All-copper heatsink

Many of today's high-quality heatsink coolers combine two or more metals to maximize thermal conductivity while minimizing cost. The most cost-effective designs use copper as the core material (due to its superior thermal efficiency) and aluminum for the radiating fins (due to its stress-handling properties). This combination provides a good cooling system with improved ruggedness. Copper is not as rigid as aluminum, so you must exercise great care when handling and installing copper coolers.

Cooling fan efficiency is directly proportional to size, RPM, and blade design. The most common fans are 60 millimeters across, with average turning rates in the range of 4500 to 6000 RPM. Fans that operate at high speeds move more air for their size than those with low RPM, though high-speed fans also create more noise. For example, 60-millimeter high-speed fans exceeding 7000 RPM are available for midrange coolers, but prove too noisy for most users.

Figure 5-3: Cooling fan

Large fans 70 to 90 millimeters in size usually move more air than small ones (assuming that the RPM is the same). An adapter mechanism is often required to take advantage of a large fan, however, since most heatsink designs are built around a 60-millimeter fan.

Heatsink Lapping

One popular way to achieve maximum thermal transfer is by lapping the heatsink's base contact surface. The term *lapping* describes the process of sanding the heatsink-to-processor contact surface to eliminate microscopic air pockets caused by machining or extruding in the heatsink fabrication process. In addition, external factors such as rough handling or poor shipping practices can lead to scratches in the contact surface. By smoothing out the surface of the heatsink, you can maximize the surface area that will be in contact with the processor. Even the smallest imperfection can create the opportunity for an air pocket to interfere with contact. This will mean a loss in cooling efficiency, because air is a poor thermal conductor.

The process of lapping can be tedious, especially with poorly machined heatsink coolers. The base surface must be sanded with fine-grit sandpaper to obtain a mirrorlike finish and eliminate as many imperfections as possible. Disassemble the fan cooler from the heatsink before you begin lapping to prevent damage to the fan or electrical connections during the procedure. Clean the heatsink thoroughly to remove any deposits or filings introduced during the lapping process. Alcohol is the best cleaning solution and it won't damage finishes on the heatsink surface.

The best way to begin lapping a heatsink is to start with a low-grit paper (400 to 600), then slowly progress to higher grades until you reach the desired results, usually at the 1200- to 1600-grit level.

The most accurate method for sanding is to move the heatsink's base contact surface across the sandpaper in an alternating circular motion, with the paper firmly affixed to a perfectly flat surface. A small piece of extruded glass, such as a windowpane or mirror, works well as the flat surface. Glass will not scratch the heatsink surface if you make an error. Regular scotch tape or even a mild spray adhesive works best for attaching the sandpaper to the glass. Lapping requires great patience. Some heatsinks can require up to an hour of fine sanding before the desired surface consistency is reached. The procedure is described below.

1. With the sandpaper firmly attached to the glass, move the heatsink across the paper in a figure-eight motion.
2. Apply even pressure throughout the length of the motion to guard against introducing deformed flat spots or upraised regions.
3. Alternate the pattern of movement at regular intervals.
4. Polish and clean the heatsink once its base is sanded to a uniform finish. Cleaning with alcohol will preserve any finishes applied to the surface.
5. Apply a quality metal polish evenly and buff it as specified on the package. A rotary tool with a cloth-polishing wheel operating at a low speed works well for many materials.
6. Once the heatsink is polished, clean it with alcohol again. As with air gaps caused by scratches, polishing materials left on the surface reduce thermal conductivity.

A properly lapped heatsink will result in a processor that runs cooler and with more stability, thus improving your likelihood of achieving overclocking success.

Heat Transfer Compounds

Applying a thermal interface material between the processor core and the heatsink surface will optimize heat transfer, even with a perfectly flat cooler that has undergone the lapping process. Most coolers included with preconfigured OEM systems feature a rubbery synthetic thermal pad in between the heatsink and the processor. Several aftermarket retail coolers also use this thermal pad because it is less expensive than the better solutions.

Thermal paste is preferable to a thermal pad. Remove any existing heatsink pad before you apply the paste. Removal of the rubbery substance should be attempted with a plastic scraper only. The edge of an old plastic credit card usually does the trick, as the card offers a good combination of flexibility and rigidity. Harder scraping implements, such as a razor blade, might scratch the heatsink's base contact surface, no matter how careful you are. Alcohol will remove any remaining deposits.

Figure 5-4: Thermal paste

Thermal paste is the preferred thermal interface material because this thin liquidlike substance works well to fill small voids between the processor and heatsink surfaces. Silicon-based and zinc-oxide-based pastes are the most commonly available materials, offering good heat transfer at low cost. Silver-bearing pastes, such as Arctic Silver™ (http://www.arcticsilver.com), are rapidly emerging as the best choice for overclocking because they are efficient for the widest range of processors.

Good-quality silver paste can offer thermal conductivity upwards of 9 watts per meter Kelvin, compared to 0.7 watts per meter Kelvin for traditional pastes. Silver has two negative properties. It can become electrically conductive under extreme pressure, such as that between the processor and heatsink. While this is not normally a concern, you must take great care when applying the paste. Silver-based pastes also cost two or three times as much as standard thermal paste products. Considering that only a small amount of paste is needed for each cooler installation, the cost factor is not a serious barrier.

Application of the thermal paste is a relatively straightforward process, assuming you observe a few simple guidelines:

1. The paste should be applied to the processor core, not the heatsink's base, in a fresh installation.
2. Apply only enough paste to barely obscure the color of the processor core. A thick layer of paste acts as a thermal insulator instead of a thermal capacitor. The paste should be spread as evenly as possible to ensure proper surface contact.
3. Use an edge of plastic or other synthetic material to apply the paste; oils found in human skin can break down or disrupt the polymers in thermal paste.

These instructions can be disregarded if the heatsink comes with the paste already applied.

Case Cooling

Proper case cooling is essential for maximizing the efficiency of the processor cooler. The average cooler fan ranges from 80 to 120 millimeters in size. As with processor cooler fans, large-diameter case fans allow slow blade rotation speeds for low sound-to-noise ratios. Due to increased power current demands, only standard 4-pin Molex pass-through adapters should power large case fans. The common three-wire fan header connectors found on most motherboards cannot supply enough amperage to power most case fan designs. Simple 4-wire to 3-wire adapters are available, at low cost, for case fans with three-prong connecters.

Figure 5-5: ATX power supply

Many modern power supply units are moving to a two-fan design to ensure adequate removal of high-temperature ambient air from both the power transformer and internal case devices. At least one case fan is recommended in addition to any power supply fans already in place. The additional case fan should be mounted to take in air, since most ATX-format power supply fans are mounted to exhaust air from the case.

The goal is to maintain a balance of exhaust air to intake air for consistent flow throughout the case interior. Most cases include predesignated mounting points for fans to enable the best possible airflow regulation. The most common mounting point on a tower is the bottom front side of the case. This placement allows cool air to circulate directly through the devices located in the expansion slots area.

Figure 5-6: Proper fan placement

Air filtration is another concern when installing any additional case cooling fans. Dust deposits that accumulate on system devices can become a serious problem, even in the cleanest operating environments. Microscopic dust particles quickly accumulate on surfaces due to increased airflow inside the case. Basic filtering solutions work best in most scenarios. A piece of low-density foam, sandwiched between two coarse mesh layers and secured between the fan and mounting surface, offers excellent filtration for all but the tiniest of dust particles. The foam filter should be installed in front of the fan's air intake side for maximum efficacy, though either side should work.

Increased noise from multiple case fans can be a real concern, especially for users in a small office or home setting. The best trick to reduce fan noise is to eliminate or dampen vibrations caused by the fans' rotation, which are transferred to the case chassis. A simple rubber washer, installed between each fan and case mount point, usually provides enough cushioning. A thin layer of silicon or nonadhesive RTV sealant can be applied along the fans' edges if you are concerned about possible air leaks in true sealed-case architectures.

Additional cooling components are available to boost the effectiveness of case cooling designs. *Slot coolers* can be installed in an expansion slot to exhaust hot air from warm cards, such as the powerful AGP video cards. *Bay coolers* can be installed to regulate hard-drive temperatures. The intake-style varieties will often direct airflow onto internal system components. *Nidec coolers* are a type of squirrel-cage fan that moves air through an internal slotted wheel. Nidec coolers can move a large amount of air for their low profile. These coolers are best for directing air into the case, and can be secured easily to nearly any internal surface with screws or even double-faced adhesive tape.

Alternative Cooling Technologies

Alternative cooling systems (not based on forced-air technologies) are available, though the associated costs are prohibitive for most enthusiasts. The best of these systems can cool to subzero temperatures, at which electrical circuits can operate quite efficiently. The latest experiments in super-cooling techniques involving liquid helium or nitrogen have revealed unlimited overclocking possibilities.

The predominant concern with subzero cooling is condensation due to the difference between processor temperature and air temperature. Various solutions have been adopted over the years, but most have failed to prove viable. Enthusiasts look toward silicon sealants, Styrofoam blocks, or rubber barriers to seal the processor socket region from outside humidity. These tricks work well for the short term, but tend to break down or fail during long-term use.

The best condensation-blocking solutions are usually the most expensive. Some high-end manufacturers have literally submerged the whole processor, motherboard and all, into a nonconductive liquid to isolate vital system components from variability in air temperature or humidity. Others have opted for a complex processor-to-motherboard socket connector that uses a custom interface to seal the processor completely. While either method works efficiently, implementing such designs can be expensive.

Boot-time problems are also a significant concern. Most technologies lack a real-time response, so the cooling device must reach a certain temperature before the computing system can be booted into operation. Professional products use timing devices and a complex series of electrical relays to shorten the wait. Enthusiasts must often develop homegrown techniques to accomplish the same goal. Responding to the demand for simplified installation and use of alternative cooling technologies, commercial retailers are beginning to stock a number of helpful devices.

Peltier Cooling

Peltier thermoelectric devices are gaining popularity due to their low cost, though many users do not anticipate the secondary costs involved. A Peltier circuit is a thin ceramiclike disk that acts like a heat pump once an electrical current has been applied across its substrate. A nontechnical explanation is that one side of the Peltier disk becomes cooler while the other side becomes warmer.

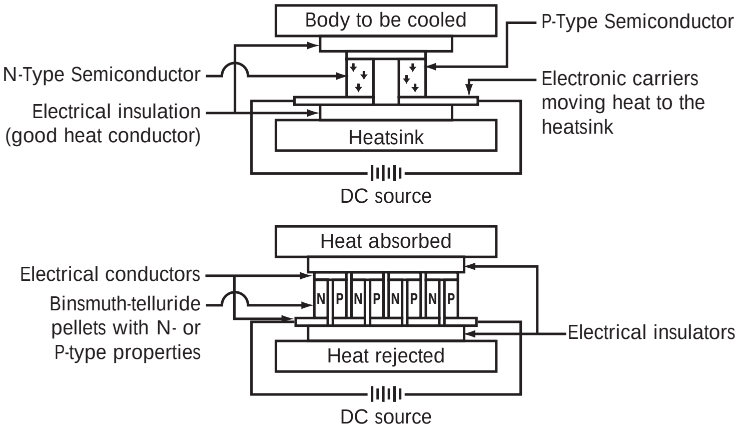

Figure 5-7: Peltier diagram

Cooling the warm side enhances the temperature variance gradient, increasing the circuit's thermal efficiency. This warm-side cooling is usually accomplished by means of a forced-air heatsink, though some use water or vapor-phase cooling to further increase the thermal load coefficient between the warm and cool sides of a Peltier system. A properly designed system can often cool even the hottest of today's processors to near zero temperatures.

Peltier devices require substantial amperage loads to operate at peak efficiency. Average 300-watt power supply units found in today's better computers cannot deal with the demands of a quality Peltier that is 72 watts or higher. An upgrade to a server-grade unit that is 400 watts or higher or adding a second power supply is often required. Wiring an external 12- to 13.8-volt power supply to the Peltier circuit is simple enough, but a relay circuit may be needed to initialize the Peltier at boot time.

Vapor-Phase Cooling

Vapor-phase cooling is active refrigeration technology, similar to the processes involved in heat pumps or air conditioners. A gaseous substance, like freon, is compressed, condensed, and forced through an exchanger device to provide a cooling effect. Vapor-phase coolers are specialized systems, usually tailored to the requirements of a specific computing system.

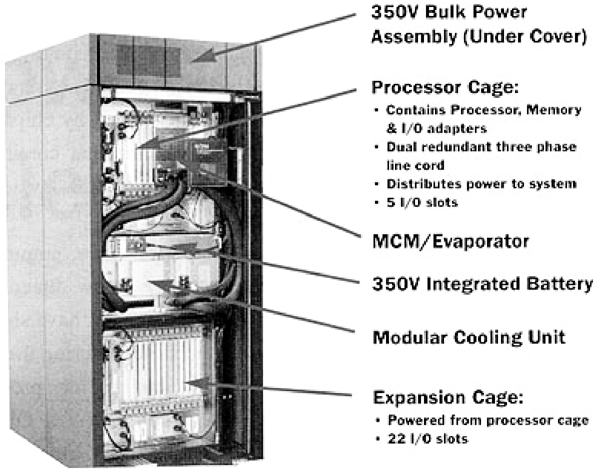

Figure 5-8: KyroTech architecture

Kryotech (http://www.kryotech.com) is a popular manufacturer of vapor-phase refrigerated computing systems. The company provides both pre-configured systems and custom cooling solutions for later system integration. Another popular product is the VapoChill cooler. The VapoChill (http://www.vapochill.com) is a highly coveted aftermarket cooling solution for overclockers. Though its price was previously very high, the VapoChill's North America price was reduced to $469 in October 2002 for their standard edition. The VapoChill is capable of cooling an overclocked processor to less than zero degrees centigrade.

Liquid Cooling

Liquid cooling was once the most popular alternative cooling solution, though the efficiency of today's forced-air systems has decreased demand. In liquid coolers a pump circulates coolant through a holding tank or radiator element and then throughout the system. Fans cool the radiator for maximum thermal efficiency.

Liquid is denser than air and therefore offers greater thermal transference. Adding other cooling technologies, such as a Peltier circuit or vapor-phase system, can further extend the thermal dissipation efficiency of a liquid cooling technique. These options can be introduced, as required, at either the processor or radiator stage of the process.

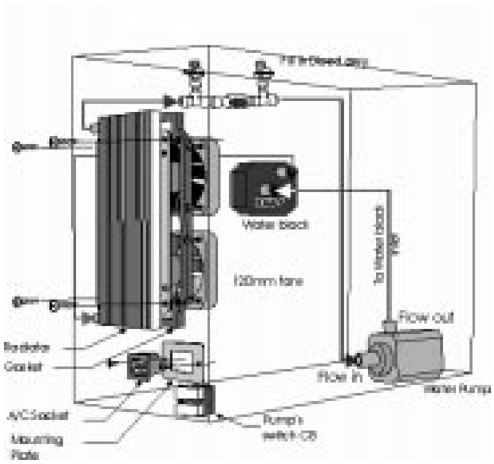

Figure 5-9: Liquid cooler architecture

Many enthusiasts build their own custom liquid cooling systems, but the decreasing cost of aftermarket products tends to encourage purchase over construction. Many quality systems now cost less than $200, making the expenditure of effort and money for a home-built system seem pointless. Moreover, the efficiency of forced-air coolers can reach or even exceed that of most generic liquid coolers. High-end forced-air solutions can cost more than $100, but that is only half the cost of a quality liquid cooler.

Figure 5-10: SwiftTech liquid cooler

Submersion Cooling

Submersion cooling is rapidly gaining favor in hardcore overclocking circles. Submersion implies literally sinking the bulk of a computing system directly into a liquid substance. The liquid can then be chilled to subzero temperatures and with it the computer. The primary advantage of this solution is that no condensation can form on sensitive electronic components. Submersion is a straightforward process: just dip the motherboard into a vat or reservoir of nonconducting liquid. However, you must avoid submerging any drives or power supply units, as these devices will fail, causing system damage.

Submersion cooling is expensive. Fluorinert liquid from the Electronics Manufacturing Division of 3M is the most attractive submersion cooling fluid. However, the cost of acquiring it can be staggering, especially because permits are required to purchase and handle it. Just a few gallons of this material can cost several thousand dollars. (And Fluorinert can evaporate, necessitating the design of a cooling vat system, which prevents evaporation of the liquid, along with your money.)

Enthusiasts have tried other liquids, like pure alcohol, but results vary wildly. Alcohol should never be used, as this material has a very low flash point. The goal is to cool the computer, not set it on fire, or worse, cause it to explode. One alternative is pure mineral oil. *Pure* is the key concept, as most bulk mineral oils on retail shelves have added water and possibly scenting agents. The presence of either impurity will cause the mineral oil to conduct electricity.