

Computer Hardware Cooling Technologies



Olinda Braden

First Edition, 2012

ISBN 978-81-323-4112-3



© All rights reserved.

Published by:

White Word Publications

4735/22 Prakashdeep Bldg,

Ansari Road, Darya Ganj,

Delhi - 110002

Email: info@wtbooks.com

Table of Contents

Chapter 1 - Computer Cooling

Chapter 2 - Computer Fan

Chapter 3 - Heat Sink

Chapter 4 - Laptop Cooler

Chapter 5 - Thermal Management of Electronic Devices and Systems

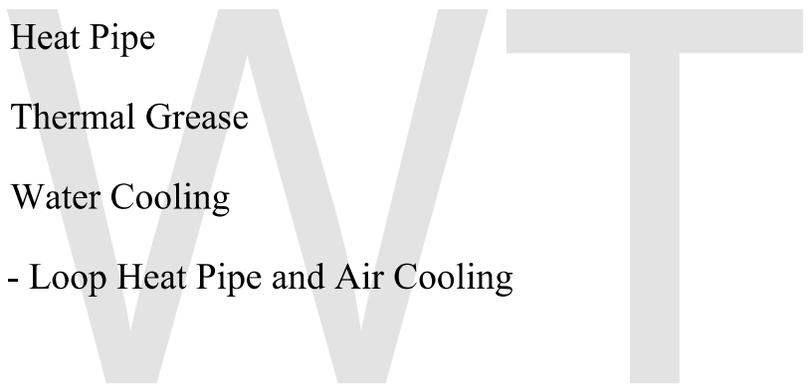
Chapter 6 - Fan Control

Chapter 7 - Heat Pipe

Chapter 8 - Thermal Grease

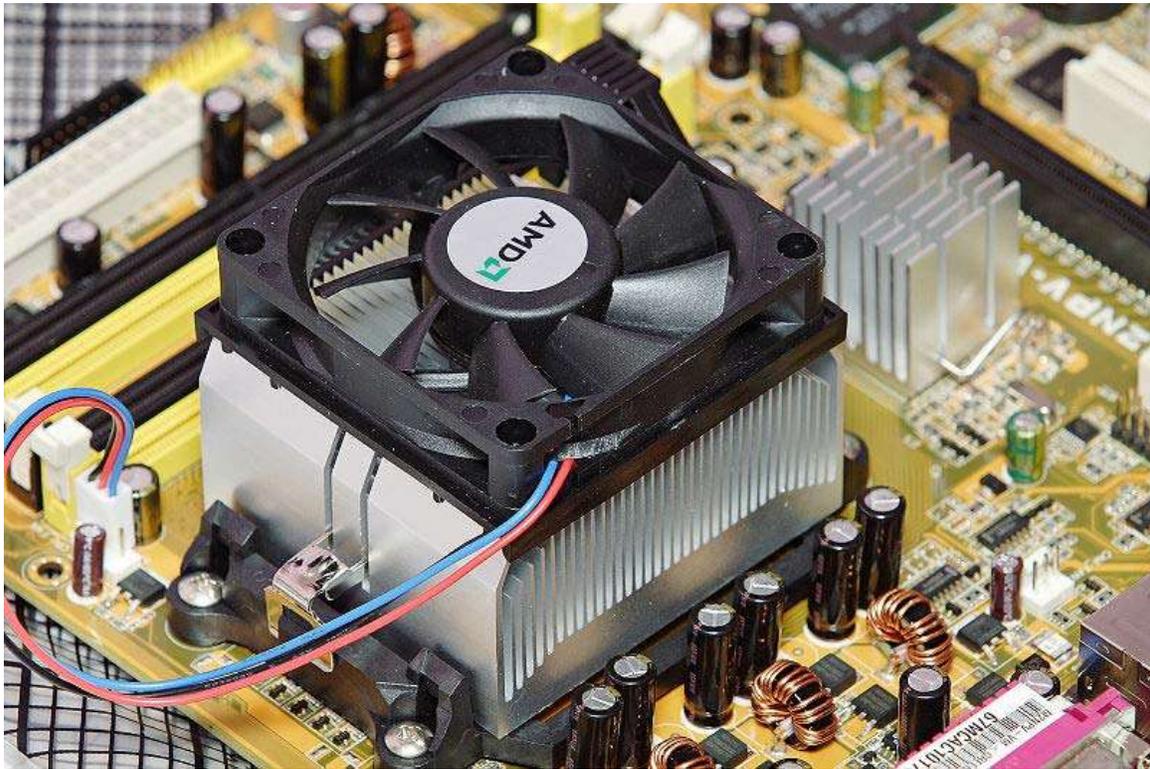
Chapter 9 - Water Cooling

Chapter 10 - Loop Heat Pipe and Air Cooling



Chapter 1

Computer Cooling



An OEM AMD heatsink mounted onto a motherboard.

Computer cooling is required to remove the waste heat produced by computer components, to keep components within their safe operating temperature limits. Various cooling methods help to improve processor performance or reduce the noise of cooling fans.

Components which produce heat and are susceptible to performance loss and damage include integrated circuits such as CPUs, chipset and graphics cards, along with hard drives (though excessive cooling of hard drives has been found to have negative effects).

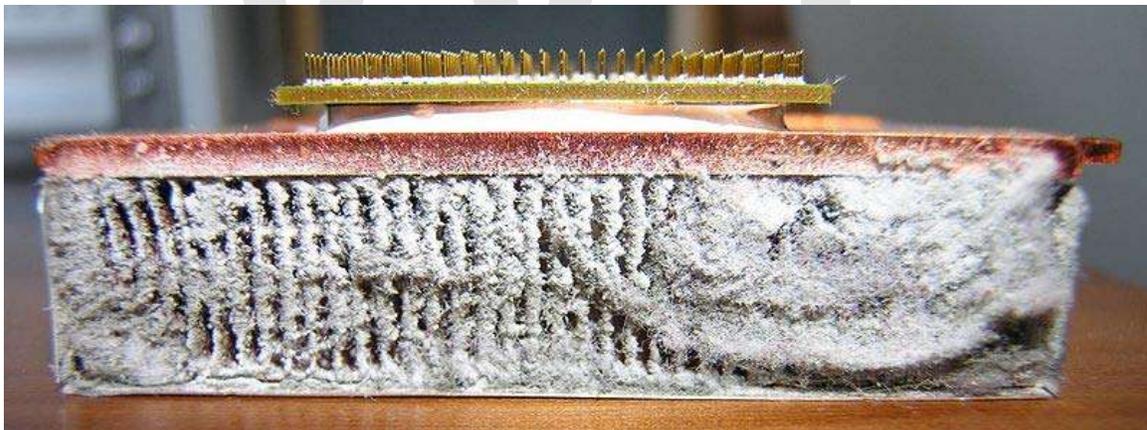
Overheated parts fail early and may give sporadic problems resulting in system freezes or crashes.

Both integral and peripheral means are used to keep the temperature of each component at a safe level. With regard to integral means, CPU and GPUs are designed with energy efficiency, including heat dissipation, in mind; though improved efficiency may only allow increased performance instead of reduced heat. Peripheral means include heat sinks to increase the surface area which dissipates heat, fans to speed up the exchange of air heated by the computer parts for cooler ambient air, and in some cases softcooling, the throttling of computer parts in order to decrease heat generation.

As a safety measure, many computers are designed to turn themselves off if the internal temperature exceeds a certain point. Alternatively, some have an option in their BIOS that allows the user to determine if the system emits an alarm beep or shuts itself down when the core temperature reaches the level set by the user. However, setting this incorrectly can result in hardware damage or erratic system behaviour.

Causes of heat build up

The amount of heat generated by an integrated circuit (e.g., a CPU or GPU), the prime cause of heat build up in modern computers, is a function of the efficiency of its design, the technology used in its construction and the frequency and voltage at which it operates.



The dust buildup on this laptop CPU heat sink after three years of use has made the laptop unusable due to frequent thermal shutdowns.

In operation, the temperature of a computer's components will rise until the heat lost to the surroundings is equal to the heat produced by the component, and thus the temperature of the component reaches equilibrium. For reliable operation, the equilibrium temperature must be sufficiently low for the structure of the computer's circuits to survive.

Cooling can be hindered by:

- **Dust** acting as a thermal insulator and impeding airflow, thereby reducing heat sink and fan performance.
- **Poor airflow** including turbulence due to friction against impeding components such as ribbon cables, or improper orientation of fans, can reduce the amount of air flowing through a case and even create localized whirlpools of hot air in the case.
- **Poor heat transfer** due to a lack of, or poor application of thermal compounds and sufficient surface area of heat sinks to radiate off the heat.

Damage prevention

Thermal sensors in some CPUs and GPUs can shut down the computer when high temperatures are detected. However, reliance on such measures may not prevent repeated incidents from permanently damaging the integrated circuit.

An integrated circuit may also shut down parts of the circuit when it is idling, or to scale back the clock speed under low workloads or high temperatures, with the goal of reducing both power use and heat generation.

Air cooling

Fans are most commonly used for air cooling. A *computer fan* may be attached to the computer case, or attached to a CPU, GPU, chipset, PSU, hard drive or PCI slot. Common fan sizes include 40, 60, 80, 92, 120, and 140 mm. Recently, 200mm fans have begun to creep into the performance market, as well as even larger sizes such as 230 and 240mm.

In desktops



Typical airflow through a desktop ATX case.

Desktop computers typically use one or more fans for cooling. Almost all desktop power supplies have at least one fan to exhaust air from the case. Most manufacturers recommend bringing cool, fresh air in at the bottom front of the case, and exhausting warm air from the top rear.

If there is more air being forced into the system than is being pumped out (due to an imbalance in the number or strength of fans), this is referred to as a "positive" airflow, as the pressure inside the unit would be higher than outside. A balanced or neutral airflow is the most efficient, although a slightly positive airflow results in less dust build up if dust filters are used. Negative pressure inside the case can create problems such as clogged optical drives due to sucking in air (and dust).

In high density computing

Data centers typically contain many racks of flat 1U servers. Air is drawn in at the front of the rack and exhausted at the rear. Because data centers typically contain such large numbers of computers and other power-consuming devices, they risk overheating of the various components if no additional measures are taken. Thus, extensive HVAC systems are used. Often a raised floor is used so the area under the floor may be used as a large plenum for cooled air and power cabling.

Another way of accommodating large numbers of systems in a small space are blade chassis. In contrast to the horizontal orientation of flat servers, blade chassis are often oriented vertically. This vertical orientation facilitates convection. When the air is heated by the hot components, it tends to flow to the top on its own, creating a natural air flow along the boards. This stack effect can help to achieve the desired air flow and cooling. Some manufacturers expressly take advantage of this effect.

In laptop computing

Most laptops use air cooling in order to keep the CPU and other components within their operating temperature range. Because the fan's air is forced through a small port, the fan and heatsinks can be clogged by dust or be obstructed by objects placed near the port. This can cause overheating, and can be a cause of component failure in laptops. The severity of this problem varies with laptop design, its use and power dissipation. With recent reductions in CPU power dissipation, this problem can be anticipated to reduce in severity.

Liquid submersion cooling

An uncommon practice is to submerge the computer's components in a thermally conductive liquid. Personal computers that are cooled in this manner do not generally require any fans or pumps, and may be cooled exclusively by passive heat exchange between the computer's parts, the cooling fluid and the ambient air. Extreme component density supercomputers such as the Cray-2 and Cray T90 used additional liquid to chilled liquid heat exchangers in order to facilitate heat removal.

The liquid used must have sufficiently low electrical conductivity in order for it not to interfere with the normal operation of the computer's components. If the liquid is somewhat electrically conductive, it may be necessary to insulate certain parts of components susceptible to electromagnetic interference, such as the CPU. For these reasons, it is preferred that the liquid be dielectric.

Liquids commonly used in this manner include various liquids invented and manufactured for this purpose by 3M, such as Fluorinert. Various oils, including but not limited to cooking, motor and silicone oils have all been successfully used for cooling personal computers.

Evaporation can pose a problem, and the liquid may require either to be regularly refilled or sealed inside the computer's enclosure. Liquid may also slowly seep into and damage components, particularly capacitors, causing an initially functional computer to fail after hours or days immersed.

Waste heat reduction

Where full-power, full-featured modern computers are not required, some companies opt to use less powerful computers or computers with fewer features. For example: in an office setting, the IT department may choose a thin client or a diskless workstation thus cutting out the heat-laden components such as hard drives and optical disks. These devices are also often powered with direct current from an external power supply brick which still wastes heat, but not inside the computer itself.

The components used can greatly affect the power consumption and hence waste heat. A VIA EPIA motherboard with CPU typically generates approximately 25 watts of heat whereas a Pentium 4 motherboard typically generates around 140 watts. While the former has considerably less computing power, both types are adequate and responsive for tasks such as word processing and spreadsheets. Choosing a LCD monitor rather than a CRT can also reduce power consumption and excess room heat, as well as the added benefit of increasing available physical desk space.

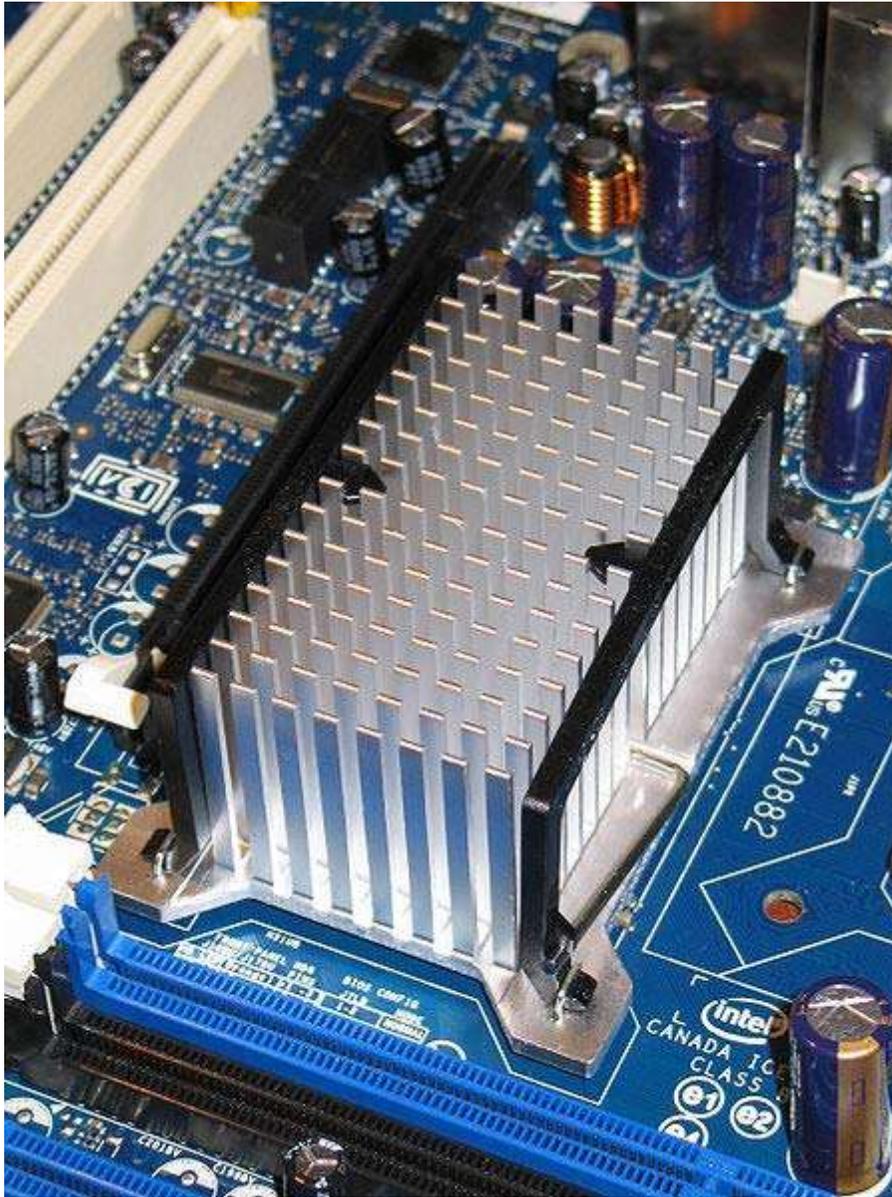
Conductive and radiative cooling

Some laptop components, such as hard drives and optical drives, are commonly cooled by having them make contact with the computer's frame, increasing the surface area which can radiate and otherwise exchange heat.

Spot cooling

In addition to system cooling, various individual components usually have their own cooling systems in place. Components which are individually cooled include, but are not limited to, the CPU, GPU and the Northbridge chip. Some cooling solutions employ one or more methods of cooling, and may also utilize logic and/or temperature sensors in order to vary the power used in active cooling components.

Passive heat-sink cooling



Passive heatsink on an Intel GMA graphics chip.

Passive heat-sink cooling involves attaching a block of machined or extruded metal to the part that needs cooling. A thermal adhesive may be used. More commonly for a personal-computer CPU, a clamp holds the heat sink directly over the chip, with a thermal grease or thermal pad spread between. This block usually has fins and ridges to increase its surface area. The heat conductivity of metal is much better than that of air, and it radiates heat better than does the component that it is protecting (usually an integrated circuit or CPU). Until recently, fan-cooled aluminium heat sinks were the norm for desktop computers. Today, many heat sinks feature copper base-plates or are entirely made of copper.

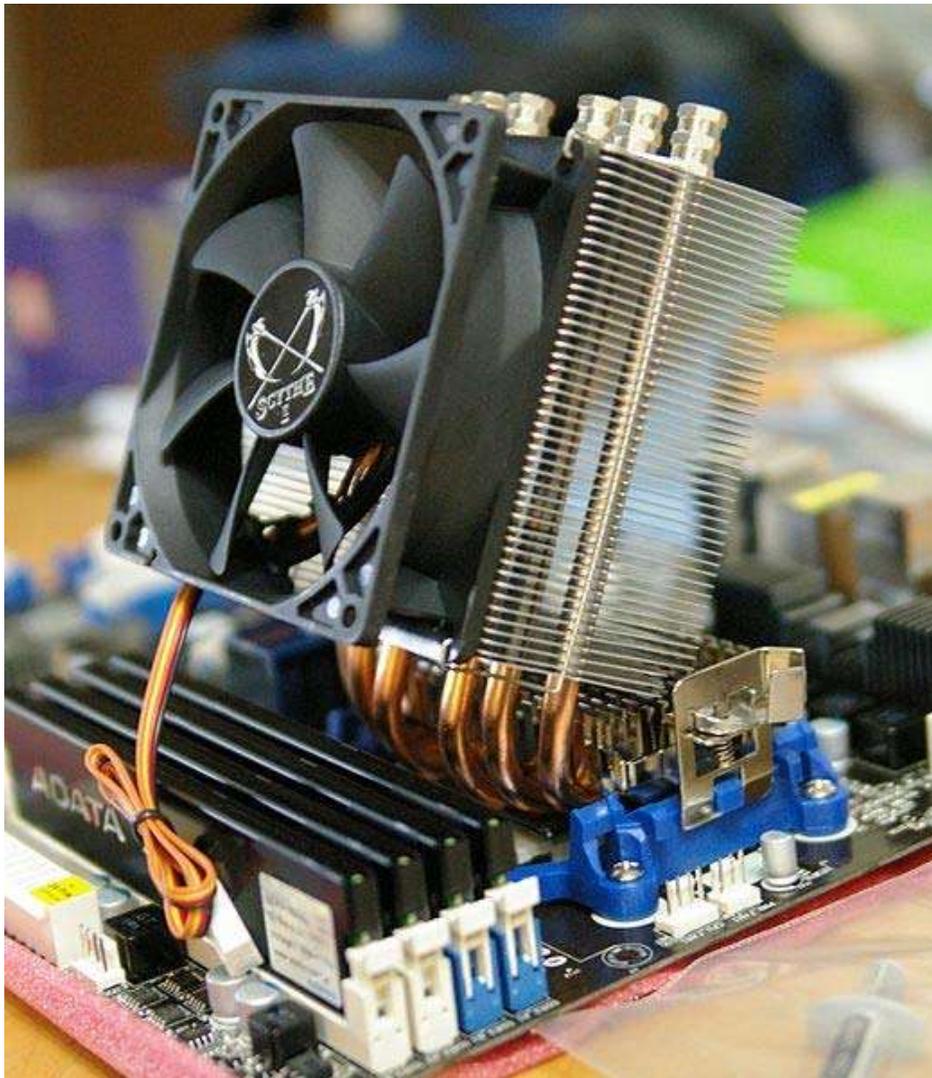
Dust buildup between the metal fins of a heat sink gradually reduces efficiency, but can be countered with a gas duster by blowing away the dust along with any other unwanted excess material.

Passive heat sinks are commonly found on older CPUs, parts that do not get very hot (such as the chipset), and low-power computers.

Usually a heat-sink is attached to the integrated heat spreader (IHS), essentially a large, flat plate attached to the CPU, with conduction paste layered between. This dissipates or spreads the heat locally. Unlike a heat sink, a spreader is meant to redistribute heat, not to remove it. In addition, the IHS protects the fragile CPU.

Passive cooling involves no fan noise.

Active heat-sink cooling



Active heat sink with a fan and heat pipes.

Active heat-sink cooling uses the same principle as passive, with the addition of a fan that blows over or through the heat sink. The air movement increases the rate at which the heat sink can exchange heat with the ambient air. Active heat sinks are the primary method of cooling modern processors and graphics cards.

The buildup of dust is greatly increased with active heat-sink cooling, because the fan continually takes in the dust present in the surrounding air.

Peltier cooling or thermoelectric cooling

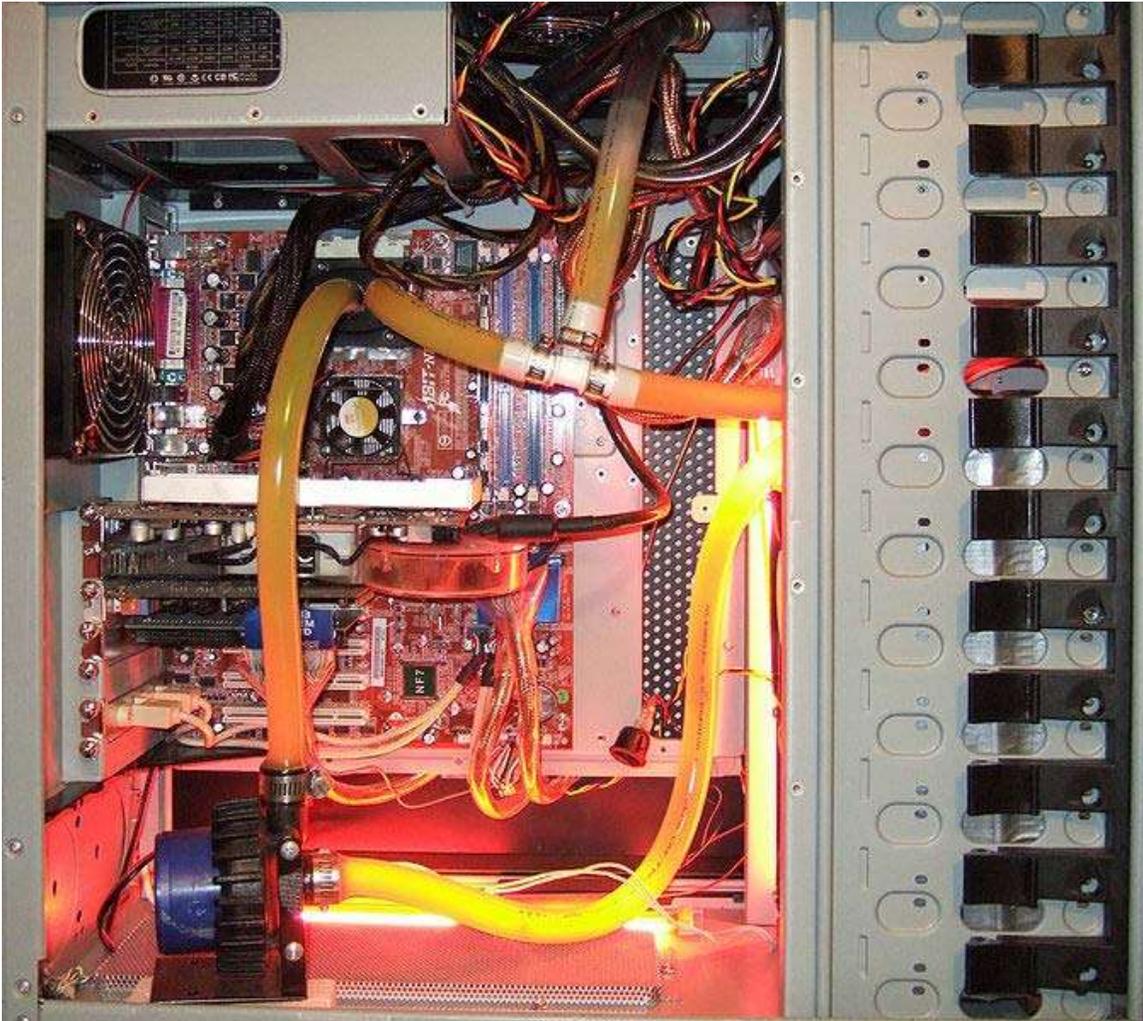
In 1821 T. J. Seebeck discovered that different metals, connected at two different junctions, will develop a micro-voltage if the two junctions are held at different temperatures. This effect is known as the "Seebeck effect"; it is the basic theory behind the TEC (thermoelectric cooling).

In 1834 Jean Peltier discovered the inverse of the Seebeck effect, now known as the "Peltier effect". He found that applying a voltage to a thermocouple creates a temperature differential between two sides. This results in an effective, albeit extremely inefficient heat pump.

Modern TECs use several stacked units each composed of dozens or hundreds of thermocouples laid out next to each other, which allows for a substantial amount of heat transfer. A combination of bismuth and tellurium is most commonly used for thermocouples.

As active heat pumps, TECs can cool the surface of components below ambient temperatures. This is impossible with common radiator cooled water cooling systems and heatpipe HSFs.

Water cooling



DIY Water cooling setup showing 12v pump, CPU Waterblock and the typical application of a T-Line

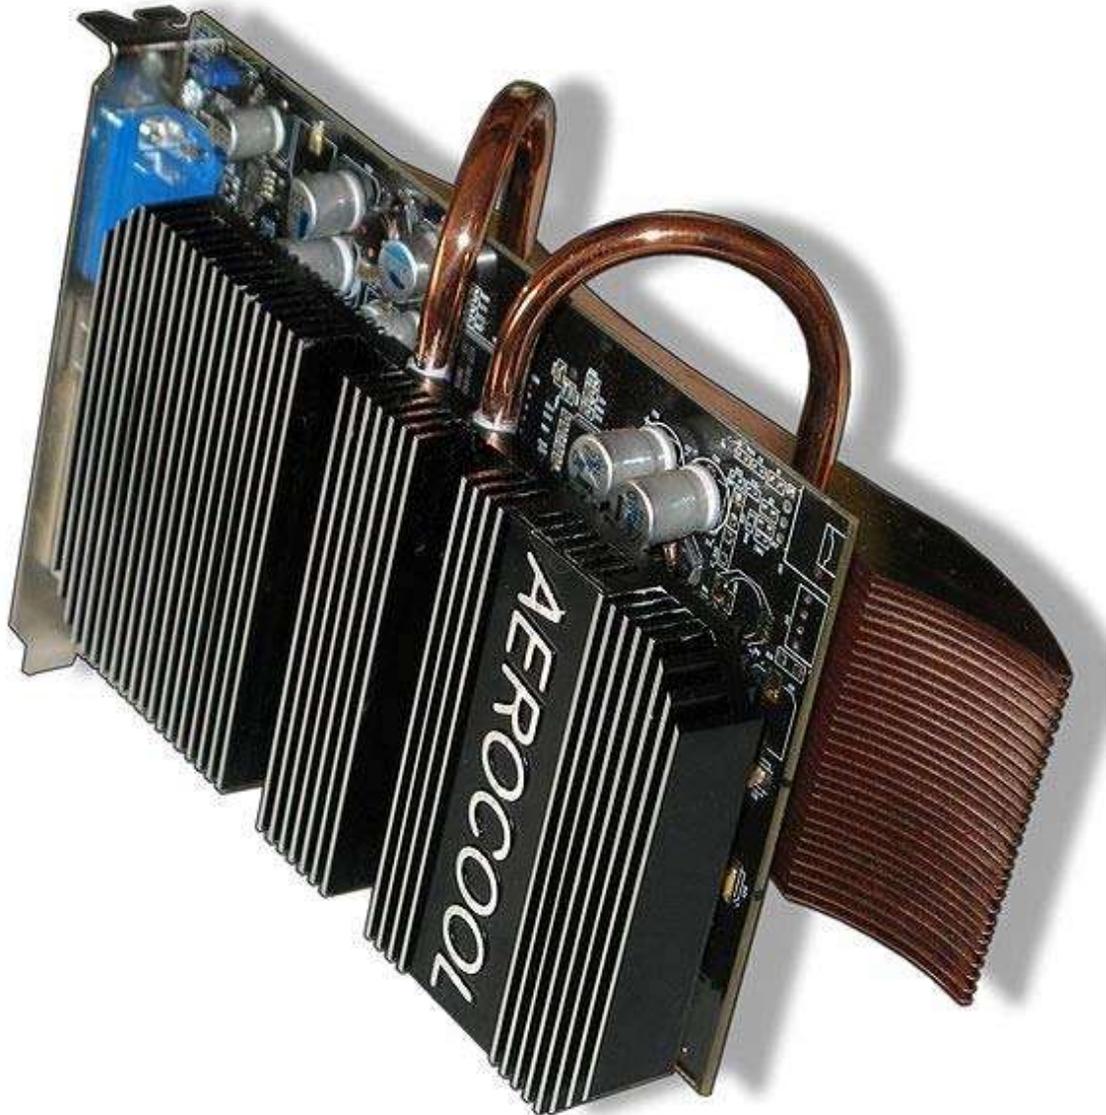
While originally limited to mainframe computers, water cooling has become a practice largely associated with overclocking in the form of either manufactured kits, or in the form of do-it-yourself setups assembled from individually gathered parts. The past few years has seen water cooling increasing its popularity with pre-assembled, moderate to high performance, desktop computers. Water has the ability to dissipate more heat from the parts being cooled than the various types of metals used in heatsinks, making it suitable for overclocking and high performance computer applications.

Advantages to water cooling include the fact that a system is not limited to cooling one component, but can be set up to cool the central processing unit, graphics processing unit, and/or other components at the same time with the same system. As opposed to air cooling, water cooling is also influenced less by the ambient temperature. Water cooling's comparatively low noise-level compares favorably to that of active cooling, which can

become quite noisy. One disadvantage to water cooling is the potential for a coolant leak. Leaked coolant can damage any electronic components it comes in contact with. Another drawback to water cooling is the complexity of the system; an active heat sink is much simpler to build, install, and maintain than a water cooling solution.

Computing folklore holds that users of the Sinclair ZX81, one of the first home computers, had to balance a carton of milk on top of the case to cool it down – perhaps an early form of water cooling.

Heat pipe



A graphics card with a heatpipe cooler design.

A heat pipe is a hollow tube containing a heat transfer liquid. As the liquid evaporates, it carries heat to the cool end, where it condenses and then returns to the hot end (under capillary action, or, in earlier implementations, under gravitation). Heat pipes thus have a much higher effective thermal conductivity than solid materials. For use in computers, the heat sink on the CPU is attached to a larger radiator heat sink. Both heat sinks are hollow as is the attachment between them, creating one large heat pipe that transfers heat from the CPU to the radiator, which is then cooled using some conventional method. This method is expensive and usually used when space is tight (as in small form-factor PCs and laptops), or absolute quiet is needed (such as in computers used in audio production studios during live recording). Because of the efficiency of this method of cooling, many desktop CPUs and GPUs, as well as high end chipsets, use heat pipes in addition to active fan-based cooling to remain within safe operating temperatures.

A new design wrinkle is known as HDT for Heatpipe Direct Touch. In this usage, the heat pipe is in direct contact with the CPU chip skin. Heatpipe Direct Touch was first introduced in the Zaward ZikaRay ZIKA-01 heatsink on February 2007, using a patent obtained from Golden Sun News Techniques Corporation in Taiwan.

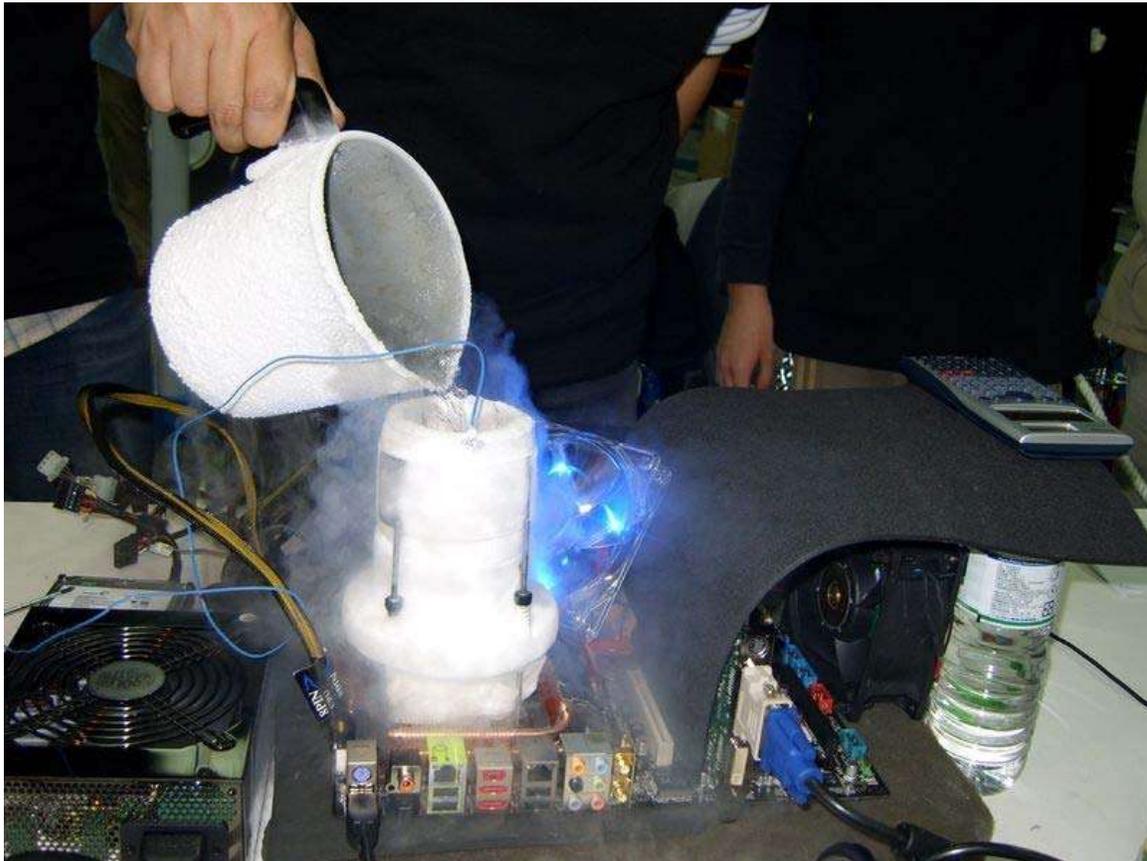
Phase-change cooling

Phase-change cooling is an extremely effective way to cool the processor. A vapor compression phase-change cooler is a unit which usually sits underneath the PC, with a tube leading to the processor. Inside the unit is a compressor of the same type as in a window air conditioner. The compressor compresses a gas (or mixture of gases) which condenses it into a liquid. Then, the liquid is pumped up to the processor, where it passes through an expansion device, this can be from a simple capillary tube to a more elaborate thermal expansion valve. The liquid evaporates (changing phase), absorbing the heat from the processor as it draws extra energy from its environment to accommodate this change. The evaporation can produce temperatures reaching around -15 to -150 degrees Celsius. The gas flows down to the compressor and the cycle begins over again. This way, the processor can be cooled to temperatures ranging from -15 to -150 degrees Celsius, depending on the load, wattage of the processor, the refrigeration system and the gas mixture used. This type of system suffers from a number of issues but mainly one must be concerned with dew point and the proper insulation of all sub-ambient surfaces that must be done (the pipes will sweat, dripping water on sensitive electronics).

Alternately a new breed of cooling system is being developed inserting a pump into the thermo siphon loop. This adds another degree of flexibility for the design engineer as the heat can now be effectively transported away from the heat source and either reclaimed or dissipated to ambient. Junction temperature can be tuned by adjusting the system pressure; higher pressure equals higher fluid saturation temperatures. This allows for smaller condensers, smaller fans and/or the effective dissipation of heat in a high ambient environment. These systems are in essence the next generation liquid cooling paradigm as they are approximately 10 times more efficient than single phase water. Since the system uses a dielectric as the heat transport media, leaks do not cause a catastrophic failure of the electric system.

This type of cooling is seen as a more extreme way to cool components, since the units are relatively expensive compared to the average desktop. They also generate a significant amount of noise, since they are essentially refrigerators, however the compressor choice and air cooling system is the main determinant of this, allowing for flexibility for noise reduction based on the parts chosen.

Liquid nitrogen



Liquid nitrogen may be used to cool an overclocked PC.

As liquid nitrogen boils at -196°C , far below the freezing point of water, it is valuable as an extreme coolant for short overclocking sessions.

In a typical installation of liquid nitrogen cooling, a copper or aluminum pipe is mounted on top of the processor or graphics card. After being heavily insulated against condensation, the liquid nitrogen is poured into the pipe, resulting in temperatures well below -100°C .

Evaporation devices ranging from cut out heat sinks with pipes attached to custom milled copper containers are used to hold the nitrogen as well as to prevent large temperature changes. However, after the nitrogen evaporates, it has to be refilled. In the realm of personal computers, this method of cooling is seldom used in contexts other than overclocking trial-runs and record-setting attempts, as the CPU will usually expire within

a relatively short period of time due to temperature stress caused by changes in internal temperature.

Although liquid nitrogen is non-flammable, it can condense oxygen directly from air. Mixtures of liquid oxygen and flammable materials can be dangerously explosive.

Liquid helium

Liquid helium, colder than liquid nitrogen, has also been used for cooling. Liquid helium evaporates at $-269\text{ }^{\circ}\text{C}$, and temperatures ranging from -230 to $-240\text{ }^{\circ}\text{C}$ have been measured from the heatsink.

Soft cooling

Softcooling is the practice of utilizing software to take advantage of CPU power saving technologies to minimize energy use. This is done using halt instructions to turn off or put in standby state CPU subparts that aren't being used or by underclocking the CPU.

Undervolting

Undervolting is a practice of running the CPU or any other component with voltages below the device specifications. An undervolted component draws less power and thus produces less heat. The ability to do this varies by manufacturer, product line, and even different production runs of the same exact product (as well as that of other components in the system), but modern processors are typically shipped with voltages higher than strictly necessary. This provides a buffer zone so that the processor will have a higher chance of performing correctly under sub-optimal conditions, such as a lower quality mainboard (motherboard). However, too low a voltage will not allow the processor to function correctly, producing errors, system freezes or crashes, or the inability to turn the system on. Undervolting too far does not typically lead to hardware damage, though in worst-case scenarios, program or system files can be corrupted.

This technique was generally employed by those seeking low-noise systems, as less cooling is needed because of the reduction of heat production. Since the popularity of hand-held or remote computers (Unmanned vehicles, mobile & cordless phones/camera/viewers, etc), undervolting is used to prolong battery endurance.

Integrated chip cooling techniques

Conventional cooling techniques all attach their “cooling” component to the outside of the computer chip, or via IHS and/or heat sinks. This “attaching” technique will always exhibit some thermal resistance, reducing its effectiveness. The heat can be more efficiently and quickly removed by directly cooling the local hot spots. At these locations, power dissipation of over $300\text{W}/\text{cm}^2$ (typical CPU are less than $100\text{W}/\text{cm}^2$, although future systems are expected to exceed $1000\text{W}/\text{cm}^2$) can occur. This form of local cooling is essential to developing high power density chips. This ideology has led to

the investigation of integrating cooling elements into the computer chip. Currently there are two techniques: micro-channel heat sinks, and jet impingement cooling.

In micro-channel heat sinks, channels are fabricated into the silicon chip (CPU), and coolant is pumped through them. The channels are designed with very large surface area which results in large heat transfers. Heat dissipation of $3000\text{W}/\text{cm}^2$ has been reported with this technique. In comparison to the Sun power density of around $7400\text{W}/\text{cm}^2$. The heat dissipation can be further increased if two-phase flow cooling is applied. Unfortunately the system requires large pressure drops, due to the small channels, and the heat flux is lower with dielectric coolants used in electronic cooling. Another local chip cooling technique is jet impingement cooling. In this technique, a coolant is flown through a small orifice to form a jet. The jet is directed toward the surface of the CPU chip, and can effectively remove large heat fluxes. Heat dissipation of over $1000\text{W}/\text{cm}^2$ has been reported. The system can be operated at lower pressure in comparison to the micro-channel method. The heat transfer can be further increased using two-phase flow cooling and by integrating return flow channels (hybrid between micro-channel heat sinks and jet impingement cooling).

Cooling and overclocking

Extra cooling is usually required by those who run parts of their computer (such as the CPU and GPU) at higher voltages and frequencies than manufacturer specifications call for, called overclocking. Increasing performance by this modification of settings results in a greater amount of heat generated and thus increasing the risk of damage to components and/or premature failure.

The installation of higher performance, non-stock cooling may also be considered modding. Many overclockers simply buy more efficient, and often, more expensive fan and heat sink combinations, while others resort to more exotic ways of computer cooling, such as liquid cooling, Peltier effect heatpumps, heat pipe or phase change cooling.

There are also some related practices that have a positive impact in reducing system temperatures:

Heat sink lapping

Heat sink lapping is the smoothing and polishing of the contact (bottom) part of a heat sink to increase its heat transfer efficiency. The desired result is a contact area which has a more even surface, as a less even contact surface creates a larger amount of insulating air between the heat sink and the computer part it is attached to. Polishing the surface using a combination of fine sandpaper and abrasive polishing liquids can produce a mirror-like shine, an indicator of a very smooth metal surface. Even a curved surface can become extremely reflective, yet not particularly flat, as is the case with curved mirrors; thus heat sink quality is based on *overall flatness*, more than optical properties. Lapping a high quality heat sink can damage it, because, although the heat sink may become shiny,

it is likely that more material will be removed from the edges, making the heat sink less effective overall.

If attempted, a piece of float glass should be used, as it self-levels as it cools and offers the most economical solution to producing a perfectly flat surface.

Use of exotic thermal conductive compounds

Some overclockers use special thermal compounds whose manufacturers claim to have a much higher efficiency than stock thermal pads. Heat sinks clean of any grease or other thermal transfer compounds have a very thin layer of these products applied, and then are placed normally over the CPU. Many of these compounds have a high proportion of silver as their main ingredient due to its high thermal conductivity. The resulting difference in the temperature of the CPU is measurable (several celsius degrees), so the heat transfer does appear to be superior to stock compounds. Some people experience negligible gains and have called to question the advantages of these exotic compounds, calling the style of application more important than the compound itself. Also note that there may be a 'setting' or 'curing' period and negligible gains may improve over time as the compound reaches its optimum thermal conductivity.

Use of rounded cables

Most older PCs use flat ribbon cables to connect storage drives (IDE or SCSI). These large flat cables greatly impede airflow by causing drag and turbulence. Overclockers and modders often replace these with rounded cables, with the conductive wires bunched together tightly to reduce surface area. Theoretically, the parallel strands of conductors in a ribbon cable serve to reduce crosstalk (signal carrying conductors inducing signals in nearby conductors), but there is no empirical evidence of rounding cables reducing performance. This may be because the length of the cable is short enough so that the effect of crosstalk is negligible. Problems usually arise when the cable is not electromagnetically protected and the length is considerable, a more frequent occurrence with older network cables.

These computer cables can then be cable tied to the chassis or other cables to further increase airflow.

This is less of a problem with new computers that use Serial ATA which has a much narrower cable.

Airflow optimization

The colder the cooling medium (the air), the more effective the cooling. Cooling air temperature can be improved with these guidelines:

- Supply cool air to the hot components as directly as possible. Examples are air snorkels and tunnels that feed outside air directly and exclusively to the CPU or GPU cooler. For example, the BTX case design prescribes a CPU air tunnel.
- Expel warm air as directly as possible. Examples are: Conventional PC (ATX) power supplies blow the warm air out the back of the case. Many dual-slot graphics card designs blow the warm air through the cover of the adjacent slot. There are also some aftermarket coolers that do this. Some CPU cooling designs blow the warm air directly towards the back of the case, where it can be ejected by a case fan.
- Air that has already been used to spot-cool a component should not be reused to spot-cool a different component (this follows from the previous items). The ATX case design can be said to violate this rule, since the power supply gets its "cool" air from the inside of the case, where it has been warmed up already. The BTX case design also violates this rule, since it uses the CPU cooler's exhaust to cool the chipset and often the graphics card.
- Prefer cool intake air, avoid inhaling exhaust air (outside air above or near the exhausts). For example, a CPU cooling air duct at the back of a tower case would inhale warm air from a graphics card exhaust. Moving all exhausts to one side of the case, conventionally the back, helps to keep the intake air cool.
- Hiding cables behind motherboard tray or simply apply ziptie and tucking cables away to provide unhindered airflow.

Fewer fans strategically placed will improve the airflow internally within the PC and thus lower the overall internal case temperature in relation to ambient conditions. The use of larger fans also improves efficiency and lowers the amount of waste heat along with the amount of noise generated by the fans while in operation.

There is little agreement on the effectiveness of different fan placement configurations, and little in the way of systematic testing has been done. For a rectangular PC (ATX) case, a fan in the front with a fan in the rear and one in the top has been found to be a suitable configuration. However, AMD's (somewhat outdated) system cooling guidelines notes that "A front cooling fan does not seem to be essential. In fact, in some extreme situations, testing showed these fans to be recirculating hot air rather than introducing cool air." It may be that fans in the side panels could have a similar detrimental effect—possibly through disrupting the normal air flow through the case. However, this is unconfirmed and probably varies with the configuration.

Chapter 2

Computer Fan



A 3D illustration of four 80 mm fans, a type of fan commonly used in personal computers (sometimes as a set, or mixed with other fan sizes).

A **computer fan** is any fan inside, or attached to, a computer case used for cooling purposes, and may refer to fans that draw cooler air into the case from the outside, expel warm air from inside, or move air across a heatsink to cool a particular component. The use of fans to cool a computer is an example of active cooling.

Usage

As processors, graphics cards, RAM and other components in computers that have increased in speed and power consumption, the amount of heat produced by these components as a side-effect of normal operation has also increased. These components need to be kept within a specified temperature range to prevent overheating, instability, malfunction and damage leading to a shortened component lifespan.

While in earlier personal computers it was possible to cool most components using natural convection (passive cooling), many modern components require more effective active cooling. To cool these components, fans are used to move heated air away from the components and draw cooler air over them. Fans attached to components are usually used in combination with a heatsink to increase the area of heated surface in contact with the air, thereby improving the efficiency of cooling.

In the IBM compatible PC market, the computer's power supply unit (PSU) almost always uses an exhaust fan to expel warm air from the PSU. Active cooling on CPUs started to appear on the Intel 80486, and by 1997 was standard on all desktop processors. Chassis or case fans, usually one exhaust fan to expel heated air from the rear and optionally an intake fan to draw cooler air in through the front, became common with the arrival of the Pentium 4 in late 2000. A third vent fan in the side of the PC, often located over the CPU, is also common. The graphics processing unit (GPU) on many modern graphics cards also requires a heatsink and fan. In some cases, the northbridge chip on the motherboard has another fan and heatsink. Other components such as the hard drives and RAM may also be actively cooled, though as of 2007 this remains relatively unusual. It is not uncommon to find five or more fans in a modern PC.

Cooling fan applications



An 80x80x25mm computer fan

Case mount

Used to aerate the case of the computer. The components inside the case cannot dissipate heat efficiently if the surrounding air is too hot. Case fans move air through the case, usually drawing cooler outside air in through the front (where it may also be drawn over the internal hard drive racks) and expelling it through the rear. There may be a third fan in the side or top of the case to draw outside air into the vicinity of the CPU, which is usually the largest single heat source. Standard case fans are 80 mm, 92 mm or 120 mm in width and length. As case fans are often the most readily visible form of cooling on a PC, decorative fans are widely available and may be lit with LEDs, made of UV-reactive plastic, and covered with decorative grilles. Decorative fans and accessories are popular with case modders. Air filters are often used over intake fans, to prevent dust from entering the case.

A power supply (PSU) fan often plays a double role, not only keeping the PSU itself from overheating, but also removing warm air from inside the case. PSUs with two fans are also available, which typically have a fan on the inside to supply case air into the PSU and a second fan on the back to expel the heated air.

CPU fan

Used to cool the CPU (central processing unit) heatsink.

Graphics card fan

Used to cool the graphics processing unit or the memory on graphics cards. These fans were not necessary on older cards because of their low power dissipation, but most modern graphics cards, especially those designed for 3D graphics and gaming, need their own dedicated cooling fans. Some of the higher powered cards can produce more heat than the CPU (up to 289 watts), so effective cooling is especially important. Passive coolers for new video cards, however, are not unheard of, such as the Thermalright HR-03.

Chipset fan

Used to cool the northbridge of a motherboard's chipset, which may be necessary for system bus overclocking.

Other types of fans

Other less commonly encountered fans may include:

- **PCI slot fan:** A fan mounted in one of the PCI slots, usually to supply additional cooling to the PCI and/or graphics cards.
- **Hard disk fan:** A fan mounted next to or on a hard disk drive. This may be desirable on faster-spinning (e.g. 10,000 RPM) hard disks with greater heat production.
- **CD burner fan:** Some internal CD and/or DVD burners included cooling fans.

Physical characteristics

The width and height of these usually square fans are measured in millimeters; common sizes include 60 mm, 80 mm, 92 mm and 120 mm. Fans with a round frame are also available; these are usually designed so that one may use a larger fan than the mounting holes would otherwise allow (i.e., a 120 mm fan with 90 mm holes). The amount of airflow which fans generate is typically measured in cubic feet per minute (CFM), and the speed of rotation is measured in revolutions per minute (RPM). Often, computer enthusiasts choose fans which have a higher CFM rating, but produce less noise (measured in decibels, or dB), and some fans come with an adjustable RPM rating to produce less noise when the computer does not require additional airflow. Fan speeds may be controlled manually (a simple potentiometer control, for example), thermally, or by the computer hardware or by software. It is also possible to run many 12V fans from the 5 V supply, at the expense of airflow, but with reduced noise levels.

The other consideration when choosing a computer fan is static pressure. A fan with high static pressure is more effective at forcing air through restricted spaces, such as the gaps between a radiator or heatsink. Therefore, enthusiasts often prioritize static pressure over CFM when choosing a fan for use with a heatsink. The relative importance of static pressure depends on the degree to which the airflow is restricted by geometry (i.e. static pressure becomes more important as the spacing between heatsink blades decreases). Static pressure is usually measured in either mm Hg or mm H₂O.

The type of bearing used in a fan can affect its performance and noise output. Most computer fans use one of the following bearing types:

- **Sleeve bearing** fans use two surfaces lubricated with oil or grease as a friction contact. Sleeve bearings are less durable as the contact surfaces can become rough and/or the lubricant dry up, eventually leading to failure. Sleeve bearings may be more likely to fail at higher temperatures, and may perform poorly when mounted in any orientation other than vertical. The lifespan of a sleeve bearing fan may be around 40,000 hours at 50 °C. Fans that use sleeve bearings are generally cheaper than fans that use ball bearings, and are quieter at lower speeds early in their life, but can grow considerably noisier as they age.
- **Rifle bearing** fans are similar to sleeve bearing, but are quieter and have almost as much lifespan as ball bearings. The bearing has a spiral groove in it that pumps fluid from a reservoir. This allows them to be safely mounted horizontally (unlike sleeve bearings), since the fluid being pumped lubricates the top of the shaft. The pumping also ensures sufficient lubricant on the shaft, reducing noise, and increasing lifespan.
- **Ball bearing** fans use ball bearings. Though generally more expensive, ball bearing fans do not suffer the same orientation limitations as sleeve bearing fans, are more durable especially at higher temperatures, and quieter than sleeve bearing fans at higher rotation speeds. The lifespan of a ball bearing fan may be around 63,000 hours at 50 °C.
- **Fluid bearing** fans have the advantages of near-silent operation and high life expectancy (comparable to ball bearing fans). However, these fans tend to be the most expensive. The **enter bearing** fan is a variation of the fluid bearing fan, developed by Everflow.
- **Magnetic bearing** or **maglev** fans, in which the fan is repelled from the bearing by magnetism.

Fan Sizing

Fans are available in wide variety of sizes and capacities. In general the faster the fan, the more noise it produces. Within a given physical size capacity is roughly proportional to current draw. For a given flow a larger fan will be quieter than a smaller fan.

Fan connector

The standard connectors for computer fans are

3-pin Molex connector KK Family

This connector is used when connecting a fan to the motherboard or other circuit board. It is a small thick rectangular in-line female connector with two tabs on the outer-most edge of one long side. The size and spacing of the pin sockets is identical to a standard 3-pin female IC connector. The three pins are used for ground, +12 V power, and a tachometer signal. Molex Part number of receptacle is 22-01-3037. Molex Part number of individual crimp contacts is 08-55-0101.

4-pin Molex connector KK Family

This is a special variant of the Molex KK connector with four pins but with the locking/polarisation features of a 3-pin connector. The additional pin is used for a pulse-width modulation signal to provide variable speed control. These can be plugged into 3-pin headers, but will lose their fan speed control. Molex Part number of receptacle is 22-01-3047. Molex Part number of individual crimp contacts is 08-55-0101.

4-pin Molex connector

This connector is used when connecting the fan directly to the power supply. It consists of two wires (red/12V and black/ground) leading to and splicing into a large in-line 4-pin male-to-female Molex connector.

Dell, Inc. proprietary

This connector is an expansion of a simple 3-pin female IC connector by adding two tabs to the middle of the connector on one side and a lock-tab on the other side. The size and spacing of the pin sockets is identical to a standard 3-pin female IC connector and 3-pin Molex connector. Some models have the wiring of the white wire (speed sensor) in the middle, whereas the standard 3-pin Molex requires the white wire as pin #3, thus compatibility issues may exist.

Alternatives

If a fan is not desirable, because of noise, reliability, or environmental concerns, there are some alternatives:

- Very Rarely, such as ultra silent home theatre machines, can rely on passive cooling alone and do not require a case fan to keep computer components at ordinary operating temperatures. More commonly (such as in simple business and home machines) a power supply fan alone is sufficient to cool the machine.
- Undervolting and/or underclocking to reduce power dissipation
- Larger heatsinks (for example, some motherboards have northbridge fans; others have larger, more costly heatsinks)
- Natural convection cooling: carefully designed, correctly oriented, and sufficiently large CPU coolers can dissipate up to 100 W by natural convection alone
- More unusual solutions, *e.g.* heatpipes bonded to the metal case, water cooling, or refrigeration
- Motherboards sunk in liquid oil provides excellent convection cooling and protects from humidity and water without the need for heatsinks or fans. Special care must be taken to ensure compatibility with adhesives and sealants used on

the motherboard and ICs. This solution is used in some external environments like wireless equipments located in the wild.

- Ionic wind cooling is being researched, whereby air is moved by ionizing air between 2 electrodes. this replaces the fan and has the advantage of no moving parts.

WWT

Chapter 3

Heat Sink

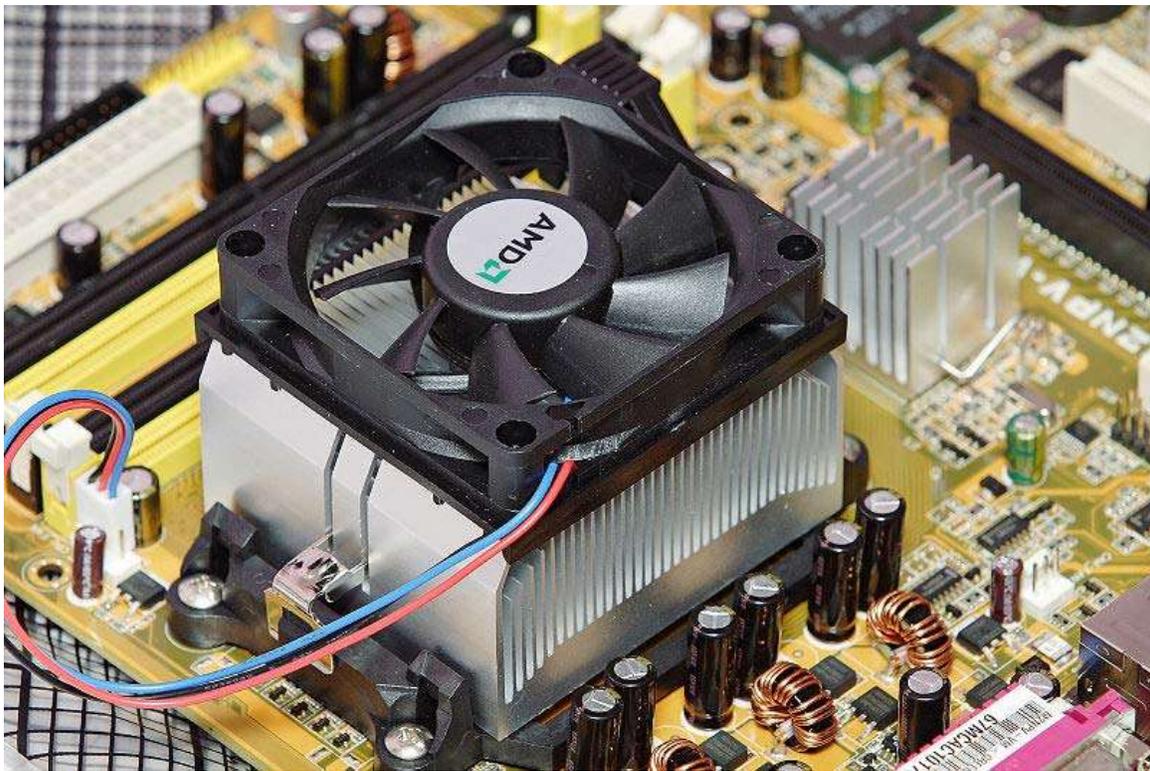


Figure 1: An active (fan cooled) heat sink used for the processor cooling on a PC motherboard. To its right is a smaller pin fin heat sink used to cool the Northbridge of the motherboard.

A **heat sink** is a term for a component or assembly that transfers heat generated within a solid material to a fluid medium, such as air or a liquid. Examples of heat sinks are the heat exchangers used in refrigeration and air conditioning systems and the radiator (also a heat exchanger) in a car. Heat sinks also help to cool electronic and optoelectronic devices, such as higher-power lasers and light emitting diodes (LEDs).

A heat sink is physically designed to increase the surface area in contact with the cooling fluid surrounding it, such as the air. Approach air velocity, choice of material, fin (or other protrusion) design and surface treatment are some of the design factors which influence the thermal resistance, i.e. thermal performance, of a heat sink. One engineering application of heat sinks is in the thermal management of electronics, often computer CPU or graphics processors. For these, heat sink attachment methods and thermal interface materials also influence the eventual junction or die temperature of the processor(s). Thermal adhesive (also known as thermal grease) is added to the base of the heatsink to help its thermal performance. Theoretical, experimental and numerical methods can be used to determine a heat sink's thermal performance.

Basic heat sink heat transfer principle

A heat sink is an object that transfers thermal energy from a higher temperature to a lower temperature *fluid medium*. The fluid medium is frequently air, but can also be water or in the case of heat exchangers, refrigerants and oil. If the fluid medium is water, the 'heat sink' is frequently called a cold plate.

To understand the principle of a heat sink, consider Fourier's law of heat conduction. Joseph Fourier was a French mathematician who made important contributions to the analytical treatment of heat conduction. Fourier's law of heat conduction, simplified to a one-dimensional form in the x -direction, shows that when there is a temperature gradient in a body, heat will be transferred from the higher temperature region to the lower temperature region. The rate at which heat is transferred by conduction, q_k , is proportional to the product of the temperature gradient and the cross-sectional area through which heat is transferred.

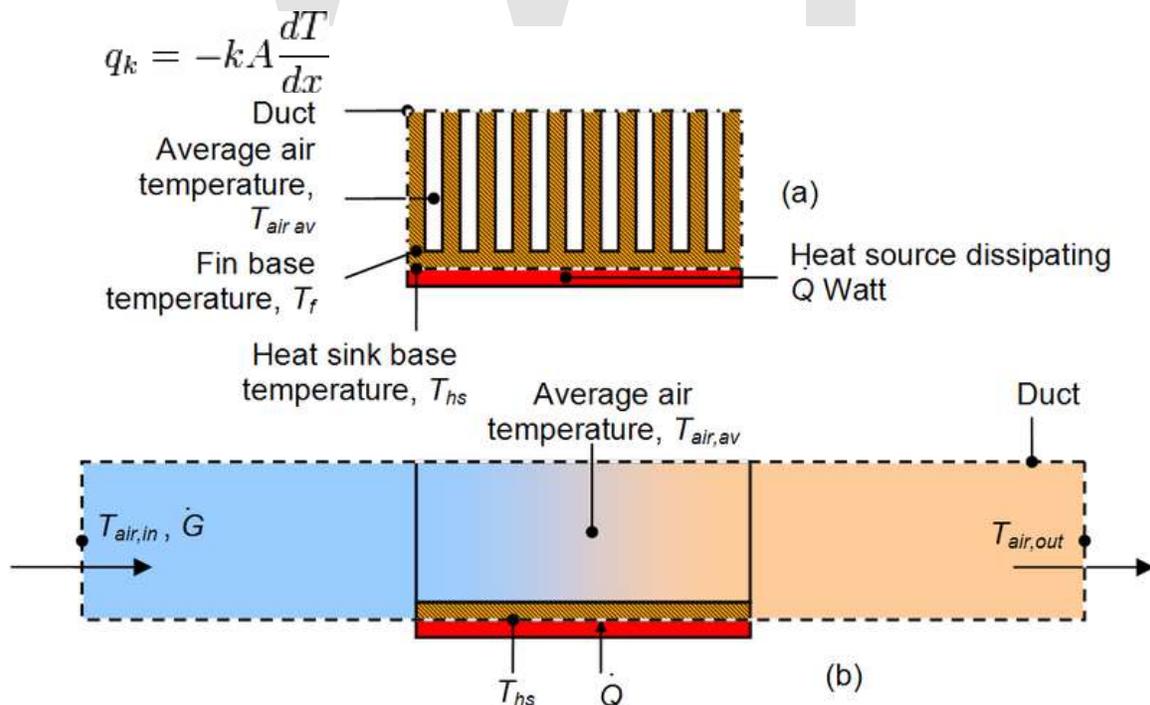


Figure 2: Sketch of a heat sink in a duct used to calculate the governing equations from conservation of energy and Newton's law of cooling.

Consider a heat sink in a duct, where air flows through the duct, as shown in Figure 2. It is assumed that the heat sink base is higher in temperature than the air. Applying the conservation of energy, for steady-state conditions, and Newton's law of cooling to the temperature nodes shown in Figure 2 gives the following set of equations.

$$\dot{Q} = \dot{m}c_{p,in}(T_{air,out} - T_{air,in}) \quad (1)$$

$$\dot{Q} = \frac{T_{hs} - T_{air,av}}{R_{hs}} \quad (2)$$

where $T_{air,av} = \frac{T_{air,out} + T_{air,in}}{2} \quad (3)$

Using the mean air temperature is an assumption that is valid for relatively short heat sinks. When compact heat exchangers are calculated, the logarithmic mean air temperature is used. \dot{m} is the air mass flow rate in kg/s.

The above equations show that

- When the air flow through the heat sink decreases, this results in an increase in the average air temperature. This in turn increases the heat sink base temperature. And additionally, the thermal resistance of the heat sink will also increase. The net result is a higher heat sink base temperature.
 - The increase in heat sink thermal resistance with decrease in flow rate will be shown in later here.
- The inlet air temperature relates strongly with the heat sink base temperature. For example, if there is recirculation of air in a product, the inlet air temperature is not the ambient air temperature. The inlet air temperature of the heat sink is therefore higher, which also results in a higher heat sink base temperature.
- Therefore, if there is no air or fluid flow around the *heat sink*, the energy dissipated to the air can not be transferred to the ambient air. Therefore, the heat sink functions poorly.
- Furthermore, a heat sink is not a device with the "magical ability to absorb heat like a sponge and send it off to a parallel universe".

Other examples of situations in which a heat sink has impaired efficiency:

- Pin fins have a lot of surface area, but the pins are so close together that air has a hard time flowing through them.
- Aligning a heat sink so that the fins are not in the direction of flow.
- Aligning the fins horizontally for a natural convection heat sink. Whilst a heat sink is stationary and there are no centrifugal forces and artificial gravity, air that

is warmer than the ambient temperature *always* flows upward, given essentially-still-air surroundings; this is convective cooling.

Design factors which influence the thermal performance of a heat sink

Material

The most common heat sink material is aluminium. Chemically pure aluminium is not used in the manufacture of heat sinks, but rather aluminium alloys. Aluminium alloy 1050A has one of the higher thermal conductivity values at 229 W/m•K. However, it is not recommended for machining, since it is a relatively soft material. Aluminium alloys 6061 and 6063 are the more commonly used aluminium alloys, with thermal conductivity values of 166 and 201 W/m•K, respectively. The aforementioned values are dependent on the temper of the alloy.

Copper is also used since it has around twice the conductivity of aluminium, but is three times as heavy as aluminium. Copper is also around four to six times more expensive than aluminium, but this is market dependent. Copper and aluminium prices can be compared in figures 3 and 4, or on Internet websites, such as the London Metal Exchange. Aluminium has the added advantage that it is able to be extruded, while copper can not. Copper heat sinks are machined and skived. Another method of manufacture is to solder the fins into the heat sink base.

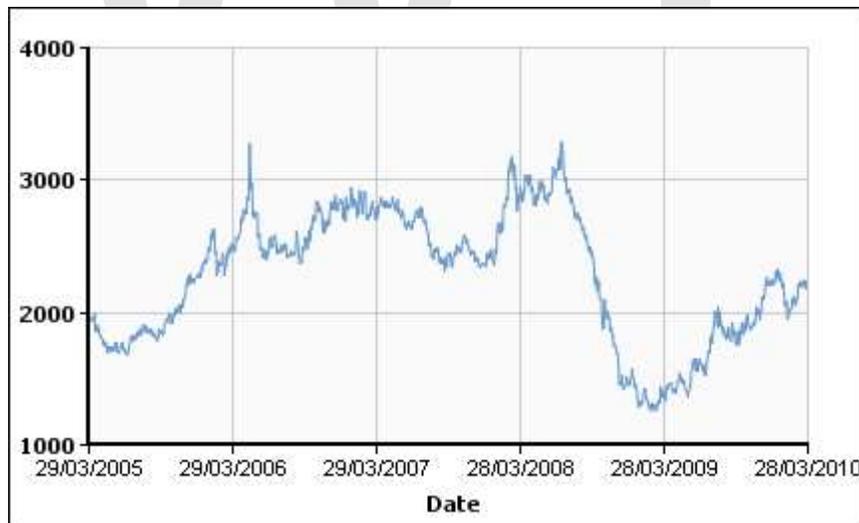


Figure 3: Aluminium cash buyer prices listed in US Dollar per metric tonne. Prices are shown for the period 29 March 2005 to the 29 March 2010.

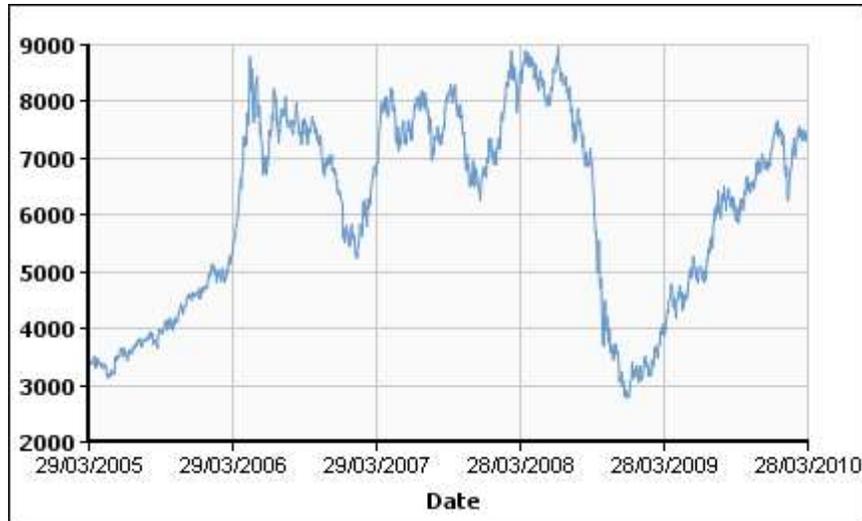


Figure 4: Copper cash buyer prices listed in US Dollar per metric tonne. Prices are shown for the period 29 March 2005 to the 29 March 2010.

Another heat sink material that can be used is diamond. With a value of 2000 W/mK it exceeds that of copper by a factor of five. In contrast to metals, where heat is conducted by delocalized electrons, lattice vibrations are responsible for diamond's very high thermal conductivity. For thermal management applications, the outstanding thermal conductivity and diffusivity of diamond is an essential. Nowadays CVD diamond is used as submounts for high-power integrated circuits and laser diodes.

Composite materials can be used. Examples are a copper-tungsten pseudoalloy, AlSiC (silicon carbide in aluminium matrix), Dymalloy (diamond in copper-silver alloy matrix), and E-Material (beryllium oxide in beryllium matrix). Such materials are often used as substrates for chips, as their thermal expansion coefficient can be matched to ceramics and semiconductors.

Fin efficiency

Fin efficiency is one of the parameters which makes a higher thermal conductivity material important. A fin of a heat sink may be considered to be a flat plate with heat flowing in one end and being dissipated into the surrounding fluid as it travels to the other. As heat flows through the fin, the combination of the thermal resistance of the heat sink impeding the flow and the heat lost due to convection, the temperature of the fin and, therefore, the heat transfer to the fluid, will decrease from the base to the end of the fin. This factor is called the fin efficiency and is defined as the actual heat transferred by the fin, divided by the heat transfer were the fin to be isothermal (hypothetically the fin having infinite thermal conductivity). Equations 6 and 7 are applicable for straight fins.

$$\eta_f = \frac{\tanh(mL_c)}{mL_c} \quad (6)$$

$$mL_c = \sqrt{\frac{2h_f}{kt_f}} L_f \quad (7)$$

Where:

- h_f is the convection coefficient of the fin
 - Air: 10 to 100 W/(m²K)
 - Water: 500 to 10,000 W/(m²K)
- k is the thermal conductivity of the fin material
 - Aluminum: 120 to 240 W/(m·K)
- L_f is the fin height (m)
- t_f is the fin thickness (m)

To increase the fin efficiency of fins:

- Decrease the fin aspect ratio, by:
 - Increasing the fin thickness, or
 - Decreasing the fin length
- Increase the thermal conductivity of the fins, (e.g. by using copper instead of aluminum)

Spreading resistance

Another parameter that concerns the thermal conductivity of the heat sink material is spreading resistance. Spreading resistance occurs when thermal energy is transferred from a small area to a larger area in a substance with finite thermal conductivity. In a heat sink, this means that heat does not distribute uniformly through the heat sink base. The spreading resistance phenomenon is shown by how the heat travels from the heat source location and causes a large temperature gradient between the heat source and the edges of the heat sink. This means that some fins are at a lower temperature than if the heat source were uniform across the base of the heat sink. This nonuniformity increases the heat sink's effective thermal resistance.

To decrease the spreading resistance in the base of a heat sink:

- Increase the base thickness
- Choose a different material with better thermal conductivity
- Use a vapour chamber or heat pipe in the heat sink base.

Fin arrangements



Figure 5: A pin-, straight- and flared fin heat sink types

A pin fin heat sink is a heat sink that has pins that extend from its base. The pins can be cylindrical, elliptical or square. A pin is by far one of the more common heat sink types available on the market. A second type of heat sink fin arrangement is the straight fin. These run the entire length of the heat sink. A variation on the straight fin heat sink is a cross cut heat sink. A straight fin heat sink is cut at regular intervals but at a coarser pitch than a pin fin type.

In general, the more surface area a heat sink has, the better it works. However, this is not always true. The concept of a pin fin heat sink is to try to pack as much surface area into a given volume as possible. As well, it works well in any orientation. Kordyan has compared the performance of a pin fin and a straight fin heat sink of similar dimensions. Although the pin fin has 194 cm² surface area while the straight fin has 58 cm², the temperature difference between the heat sink base and the ambient air for the pin fin is 50 °C. For the straight fin it was 44 °C or 6 °C better than the pin fin. Pin fin heat sink performance is significantly better than straight fins when used in their intended application where the fluid flows axially along the pins rather than only tangentially across the pins.

Comparison of a pin fin and straight fin heat sink of similar dimensions. Adapted from data of						
Heat sink fin type	Width [cm]	Length [cm]	Height [cm]	Surface area [cm ²]	Volume [cm ³]	Temperature difference, $T_{\text{case}} - T_{\text{air}}$ [°C]
Straight	2.5	2.5	3.2	58	20	44
Pin	3.8	3.8	1.7	194	24	51

Another configuration is the flared fin heat sink; its fins are not parallel to each other, as shown in figure 5. Flaring the fins decreases flow resistance and makes more air go through the heat sink fin channel; otherwise, more air would bypass the fins. Slanting them keeps the overall dimensions the same, but offers longer fins. Forghan, et al. have published data on tests conducted on pin fin, straight fin and flared fin heat sinks. They found that for low approach air velocity, typically around 1 m/s, the thermal performance is at least 20% better than straight fin heat sinks. Lasance and Eggink also found that for

the bypass configurations that they tested, the flared heat sink performed better than the other heat sinks tested.

Surface colour

The heat transfer from the heatsink is mediated by two effects: conduction via the coolant, and thermal radiation. The surface of the heatsink influences its emissivity; shiny metal absorbs and radiates only a small amount of heat, while matte black radiates highly.

In coolant-mediated heat transfer, the contribution of radiation is generally small. A layer of coating on the heatsink can then be counterproductive, as its thermal resistance can impair heat flow from the fins to the coolant. Finned heatsinks with convective or forced flow will not benefit significantly from being colored.

In situations with significant contribution of radiative cooling, e.g. in case of a flat non-finned panel acting as a heatsink with low airflow, the heatsink surface finish can play an important role. Matte-black surfaces will radiate much more efficiently than shiny bare metal.

The importance of radiative vs coolant-mediated heat transfer increases in situations with low ambient air pressure (e.g. high-altitude operations) or in vacuum (e.g. satellites in space).

Engineering applications

Processor/Microprocessor cooling

Heat dissipation is an unavoidable by-product of all but micropower electronic devices and circuits. In general, the temperature of the device or component will depend on the thermal resistance from the component to the environment, and the heat dissipated by the component. To ensure that the component temperature does not overheat, a thermal engineer seeks to find an efficient heat transfer path from the device to the environment. The heat transfer path may be from the component to a printed circuit board (PCB), to a heat sink, to air flow provided by a fan, but in all instances, eventually to the environment.

Two additional design factors also influence the thermal/mechanical performance of the thermal design:

1. The method by which the heat sink is mounted on a component or processor. This will be discussed under the section *attachment methods*.
2. For each interface between two objects in contact with each other, there will be a temperature drop across the interface. For such composite systems, the temperature drop across the interface may be appreciable. This temperature change may be attributed to what is known as the thermal contact resistance. *Thermal interface materials* (TIM) decrease the thermal contact resistance.

Attachment methods for microprocessors and similar ICs

As power dissipation of components increases and component package size decreases, thermal engineers must innovate to ensure components won't overheat. Devices that run cooler last longer. A heat sink design must fulfill both its thermal as well as its mechanical requirements. Concerning the latter, the component must remain in thermal contact with its heat sink with reasonable shock and vibration. The heat sink could be the copper foil of a circuit board, or else a separate heat sink mounted onto the component or circuit board. Attachment methods include thermally conductive tape or epoxy, wire-form z clips, flat spring clips, standoff spacers, and push pins with ends that expand after installing.

- **Thermally conductive tape**



Figure 6: Roll of thermally conductive tape.

Thermally conductive tape is one of the most cost-effective heat sink attachment materials. It is suitable for low-mass heat sinks and for components with low power dissipation. It consists of a thermally conductive carrier material with a pressure-sensitive adhesive on each side.

This tape is applied to the base of the heat sink, which is then attached to the component. Following are factors that influence the performance of thermal tape:

1. Surfaces of both the component and heat sink must be clean, with no residue such as a film of silicone grease.
2. Preload pressure is essential to ensure good contact. Insufficient pressure results in areas of non-contact with trapped air, and results in higher-than-expected interface thermal resistance.
3. Thicker tapes tend to provide better "wettability" with uneven component surfaces. "Wettability" is a term used to describe the percentage area of contact of a tape on a component. Thicker tapes, however, have a higher thermal resistance than thinner tapes. From a design standpoint, it is best to strike a balance by selecting a tape thickness that provides maximum "wettability" with minimum thermal resistance.

- **Epoxy**

Epoxy is more expensive than tape, but provides a greater mechanical bond between the heat sink and component, as well as improved thermal conductivity. The epoxy chosen must be formulated for this purpose. Most epoxies are two-part liquid formulations that must be thoroughly mixed before being applied to the heat sink, and before the heat sink is placed on the component. The epoxy is then cured for a specified time, which can vary from 2 hours to 48 hours. Faster cure time can be achieved at higher temperatures. The surfaces to which the epoxy is applied must be clean and free of any residue.

The epoxy bond between the heat sink and component is semi-permanent/permanent. This makes re-work very difficult and at times impossible. The most typical damage caused by rework is the separation of the component die heat spreader from its package.

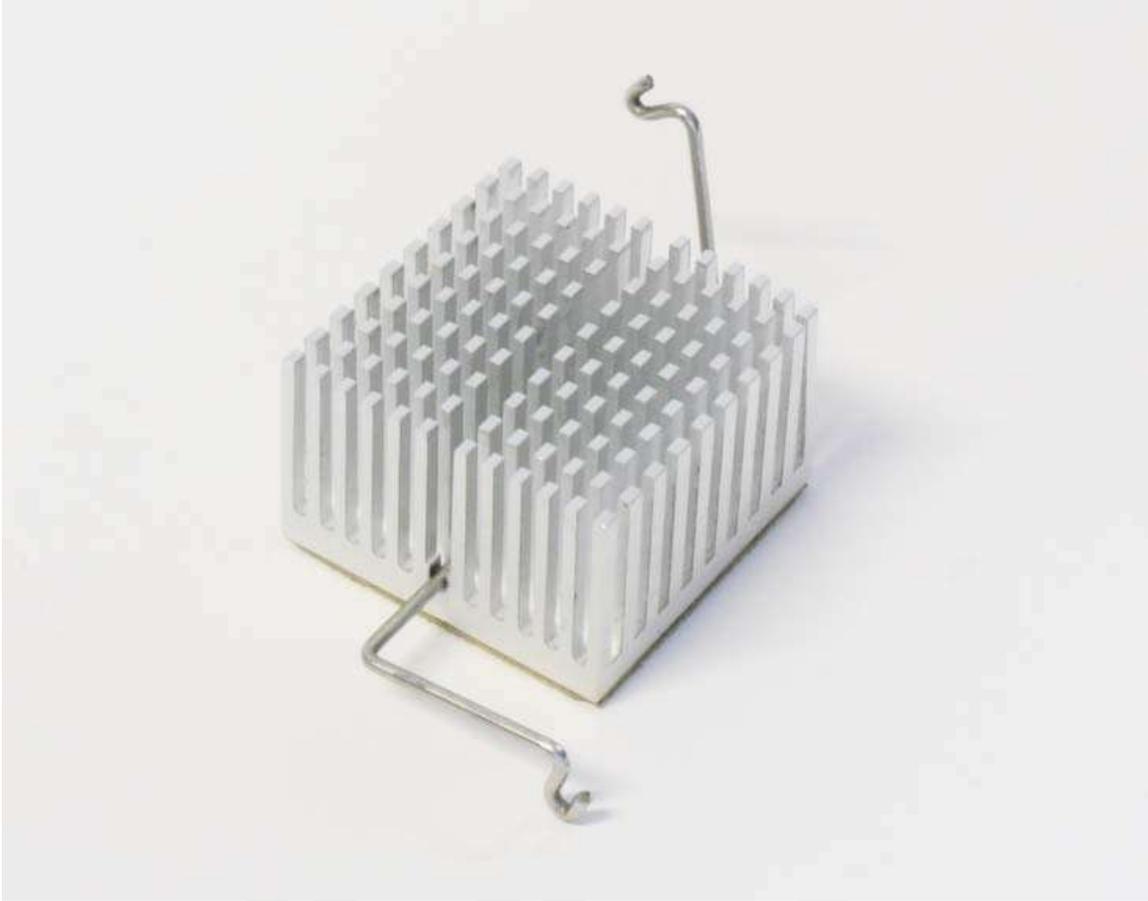


Figure 7: A pin fin heat sink with a z-clip retainer.

- **Wire form Z-clips**

More expensive than tape and epoxy, wire form z-clips attach heat sinks mechanically. To use the z-clips, the printed circuit board must have anchors. Anchors can be either soldered onto the board, or pushed through. Either type requires holes to be designed into the board. The use of RoHS solder must be allowed for because such solder is mechanically weaker than traditional Pb/Sn solder.

To assemble with a z-clip, attach one side of it to one of the anchors. Deflect the spring until the other side of the clip can be placed in the other anchor. The deflection develops a spring load on the component, which maintains very good contact. In addition to the mechanical attachment that the z-clip provides, it also permits using higher-performance thermal interface materials, such as phase change types.



Figure 8: Two heat sink attachment methods, namely the maxiGRIP (left) and Talon Clip

- **Clips**

Available for processors and ball grid array (BGA) components, clips allow the attachment of a BGA heat sink directly to the component. The clips make use of the gap created by the ball grid array (BGA) between the component underside and PCB top surface. The clips therefore require no holes in the PCB. They also allow for easy rework of components. Examples of commercially available clips are the maxiGRIP™ and superGRIP™ range from Advanced Thermal Solutions (ATS) and the Talon Clip™ from Malico. The three aforementioned clipping methods use plastic frames for the clips, but the ATS designs uses metal spring clips to provide the compression force. The Malico design uses the plastic "arm" to provide a mechanical load on the component. Depending on the product requirement, the clipping methods will have to meet shock and vibration standards, such as Telecordia GR-63-CORE, ETSI 300 019 and MIL-STD-810.



Figure 9: Push pins.

- **Push pins with compression springs**

For larger heat sinks and higher preloads, push pins with compression springs are very effective. The push pins, typically made of brass or plastic, have a flexible barb at the end that engages with a hole in the PCB; once installed, the barb retains the pin. The compression spring holds the assembly together and maintains contact between the heat sink and component. Care is needed in selection of push pin size. Too great an insertion force can result in the die cracking and consequent component failure.

- **Threaded standoffs with compression springs**

For very large heat sinks, there is no substitute for the threaded standoff and compression spring attachment method. A threaded standoff is essentially a hollow metal tube with internal threads. One end is secured with a screw through a hole in the PCB. The other end accepts a screw which compresses the spring, completing the assembly. A typical heat sink assembly uses two to four standoffs, which tends to make this the most costly heat sink attachment design. Another disadvantage is the need for holes in the PCB.

Summary of heat sink attachment methods			
Method	Pros	Cons	Cost
Thermal tape	Easy to attach. Inexpensive.	Cannot provide mechanical attachment for heavier heat sinks or for high vibration environments. Surface must be cleaned for optimal adhesion. Moderate to low thermal conductivity.	\$
Epoxy	Strong mechanical adhesion. Relatively inexpensive.	Makes board rework difficult since it can damage component. Surface must be cleaned for optimal adhesion.	\$\$
Wire form Z-clips	Strong mechanical attachment. Easy removal/rework. Applies a preload to the thermal interface material, improving thermal performance.	Requires holes in the board or solder anchors. More expensive than tape or epoxy. Custom designs.	\$\$\$
Clip-on	Applies a preload to the thermal interface material, improving thermal performance. Requires no holes or anchors. Easy removal/rework.	Must have "keep out" zone around the BGA for the clip. Extra assembly steps.	\$\$\$
Push pin with compression springs	Strong mechanical attachment. Highest thermal interface material preload. Easy removal and installation.	Requires holes in the board which increases complexity of traces in PCB.	\$\$\$\$
Stand-offs with compression springs	Strongest mechanical attachment. Highest preload for the thermal interface material. Ideal for large heat sinks.	Requires holes in the board which increases complexity of trace layout. Complicated assembly.	\$\$\$\$\$

Thermal interface materials

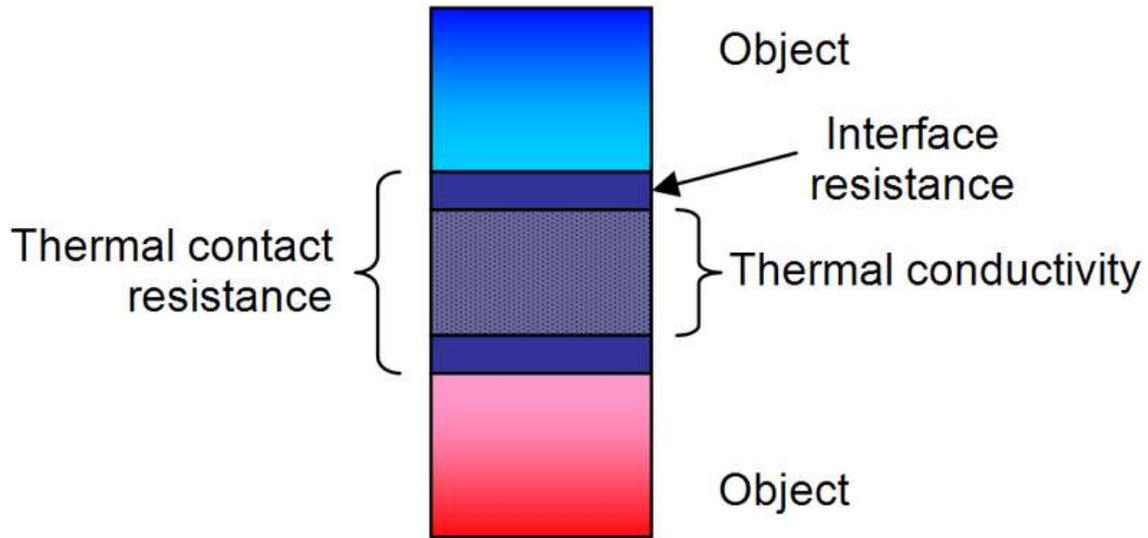


Figure 10: Thermal conductivity and the interface resistance form part of the thermal interface resistance of an thermal interface material.

Thermal contact resistance occurs due to the voids created by surface roughness effects, defects and misalignment of the interface. The voids present in the interface are filled with air. Heat transfer is therefore due to conduction across the actual contact area and to conduction (or natural convection) and radiation across the gaps. If the contact area is small, as it is for rough surfaces, the major contribution to the resistance is made by the gaps. To decrease the thermal contact resistance, the surface roughness can be decreased while the interface pressure is increased. However, these improving methods are not always practical or possible for electronic equipment. Thermal interface materials (TIM) are a common way to overcome these limitations,

Properly applied thermal interface materials displace the air that is present in the gaps between the two objects with a material that has a much-higher thermal conductivity. Air has a thermal conductivity of $0.022 \text{ W/m}\cdot\text{K}$ while TIMs have conductivities of $0.3 \text{ W/m}\cdot\text{K}$ and higher.

When selecting a TIM, care must be taken with the values supplied by the manufacturer. Most manufacturers give a value for the thermal conductivity of a material. However, the thermal conductivity does not take into account the interface resistances. Therefore, if a TIM has a high thermal conductivity, it does not necessarily mean that the interface resistance will be low.

Selection of a TIM is based on three parameters: the interface gap which the TIM must fill, the contact pressure, and the electrical resistivity of the TIM. The contact pressure is

the pressure applied to the interface between the two materials. The selection does not include the cost of the material. Electrical resistivity may, or may not, be important, depending upon electrical design details.

Selection Based on Interface Gap	
Interface gap values	Products types available
< 2 mil	Thermal grease, epoxy, phase change materials
2 – 5 mil	Phase change materials, polyimide, graphite or aluminium tapes
5 – 18 mil	Silicone coated fabrics
> 18 mil	Gap fillers

Selection Based on Contact Pressure		
Contact pressure scale	Typical pressure ranges	Product types available
Very low	< 70 kPa	Gap fillers
Low	< 140 kPa	Thermal grease, epoxy, polyimide, graphite or aluminium tapes
High	2 MPa	Silicone coated fabrics

Selection Based on Dielectric Strength			
Electrical insulation	Dielectric strength	Typical values	Product types available
Not required	N/A	N/A	Thermal grease, epoxy, phase change materials, graphite or aluminium tapes.
Required	Low	< 300 V/mil	Silicone coated fabrics, gap fillers
Required	High	> 1500 V/mil	Polyimide tape

TIM Application Notes Based on Product Type		
Product type	Application notes	Thermal performance
Thermal paste	Messy. Labour intensive. Relatively long assembly time.	++++
Epoxy	Creates 'permanent' interface bond.	++++
Phase change	Allows for pre-attachment. Softens and conforms to	++++

	interface defects at operational temperatures. Can be repositioned in field.	
Thermal tapes, including graphite, polyimide, and aluminium tapes	Easy to apply. Some mechanical strength.	+++
Silicone coated fabrics	Provide cushioning and sealing while still allowing heat transfer.	+
Gap filler	Can be used to thermally couple differing-height components to a heat spreader or heat sink. Naturally tacky.	++

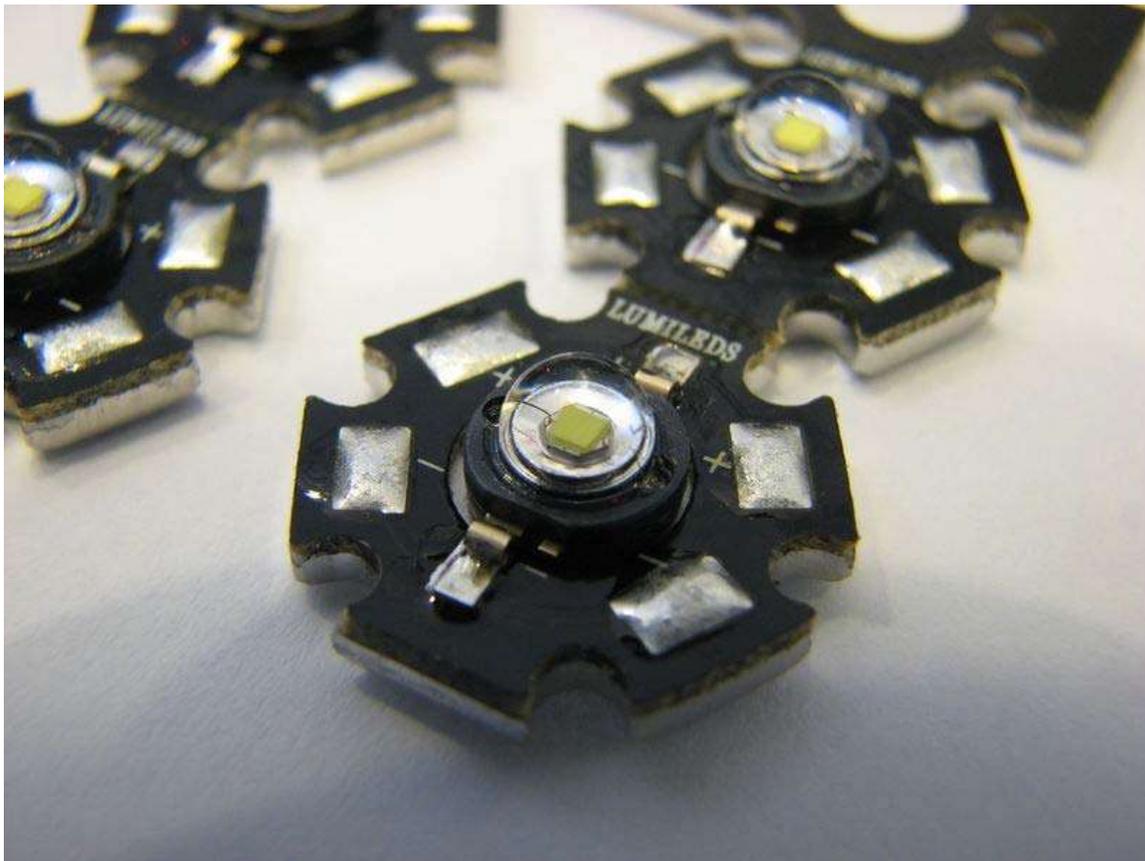


Figure 11: High power LEDs from Philips Lumileds Lighting Company mounted on 21 mm star shaped metal-core PCBs

Light emitting diode lamps

Light emitting diode (LED) performance and lifetime are strong functions of their temperature. Effective cooling is therefore essential. A case study of a LED based downlighter shows an example of the calculations done in order to calculate the required heat sink necessary for the effective cooling of lighting system. The article also shows that in order to get confidence in the results, multiple independent solutions are required that give similar results. Specifically, results of the experimental, numerical and theoretical methods should all be within 10% of each other to give high confidence in the results.

Firestopping and fireproofing

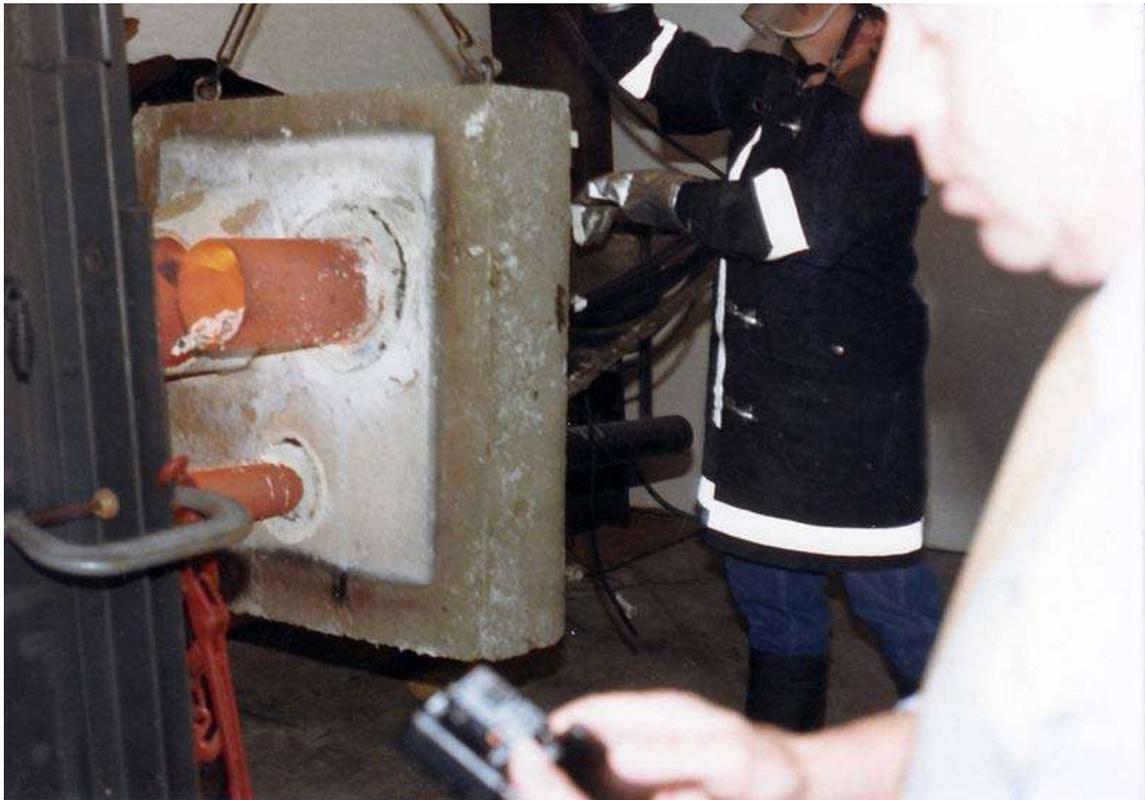


Figure 12: Fire test where the steel pipe penetrant to absorb and conduct heat from the furnace, through to the unexposed side.

A heat sink is rarely a desired thing in passive fire protection. Rather, it is usually a problem that must be overcome to maintain fire-resistance ratings. The proven ability to *overcome* heat sinks in construction is subject to building code and fire code regulations.

Firestopping

- **Problem** – Metallic penetrants and sleeves, at a density of 7.9 kg/L are denser than common firestops or concrete. Consequently, during a fire, they will absorb

more heat and conduct it to the unexposed side of a fire barrier (thus "*cooling*" the exposed side at the expense of the unexposed side), such as the cold side of a firewall. This is undesirable. Even if the fire is stopped by the barrier, one must keep the unexposed side cool to prevent autoignition of combustibles on the unexposed side of a fire barrier. The unexposed side may very well be an area of refuge, which must be safeguarded to comply with the building code. Greater penetrant and sleeve conductivity leads to lower T-ratings. Higher density firestops, such as firestop mortars act as heat sinks to absorb heat away from small penetrants, such as cables, thus increasing T-ratings.

- **Benefit** – a rare exception where heat sinks are *beneficial* in firestops is where intumescent must be activated, such as in a firestop containing a plastic pipe. Heat sinks such as wire mesh and extra metallic sleeving may be used to carry heat to intumescent to activate expansion which should choke off a melting plastic pipe or melting pipe covering, such as foamed plastic or fibreglass.

Fireproofing

In fireproofing of structural steel as well as providing circuit integrity to cables, cable trays, junction boxes and electrical conduit, the metallic items that are protected by the fireproofing measures act as heat sinks. Fireproofing methods are used to *defeat* the heat sink properties of the items they protect. In the case of circuit integrity measures, electrical services will fuse and short circuit above 140°C.

In soldering

Temporary heat sinks were sometimes used while soldering circuit boards, preventing excessive heat from damaging sensitive nearby electronics. In the simplest case, this means partially gripping a component using a heavy metal crocodile clip, hemostat or similar clamp. Modern semiconductor devices, which are designed to be assembled by reflow soldering, can usually tolerate soldering temperatures without damage. On the other hand, electrical components such as magnetic reed switches can malfunction if exposed to hotter soldering irons, so this practice is still very much in use.

Methods to determine heat sink thermal performance

In general, a heat sink performance is a function of material thermal conductivity, dimensions, fin type, heat transfer coefficient, air flow rate, duct size. To determine the thermal performance of a heat sink, a theoretical model can be made. Alternatively, the thermal performance can be measured experimentally. Due to the complex nature of the highly 3D flow in present in applications, numerical methods or CFD can also be used. Here we will discuss the aforementioned methods for the determination of the heat sink thermal performance.

A heat transfer theoretical model

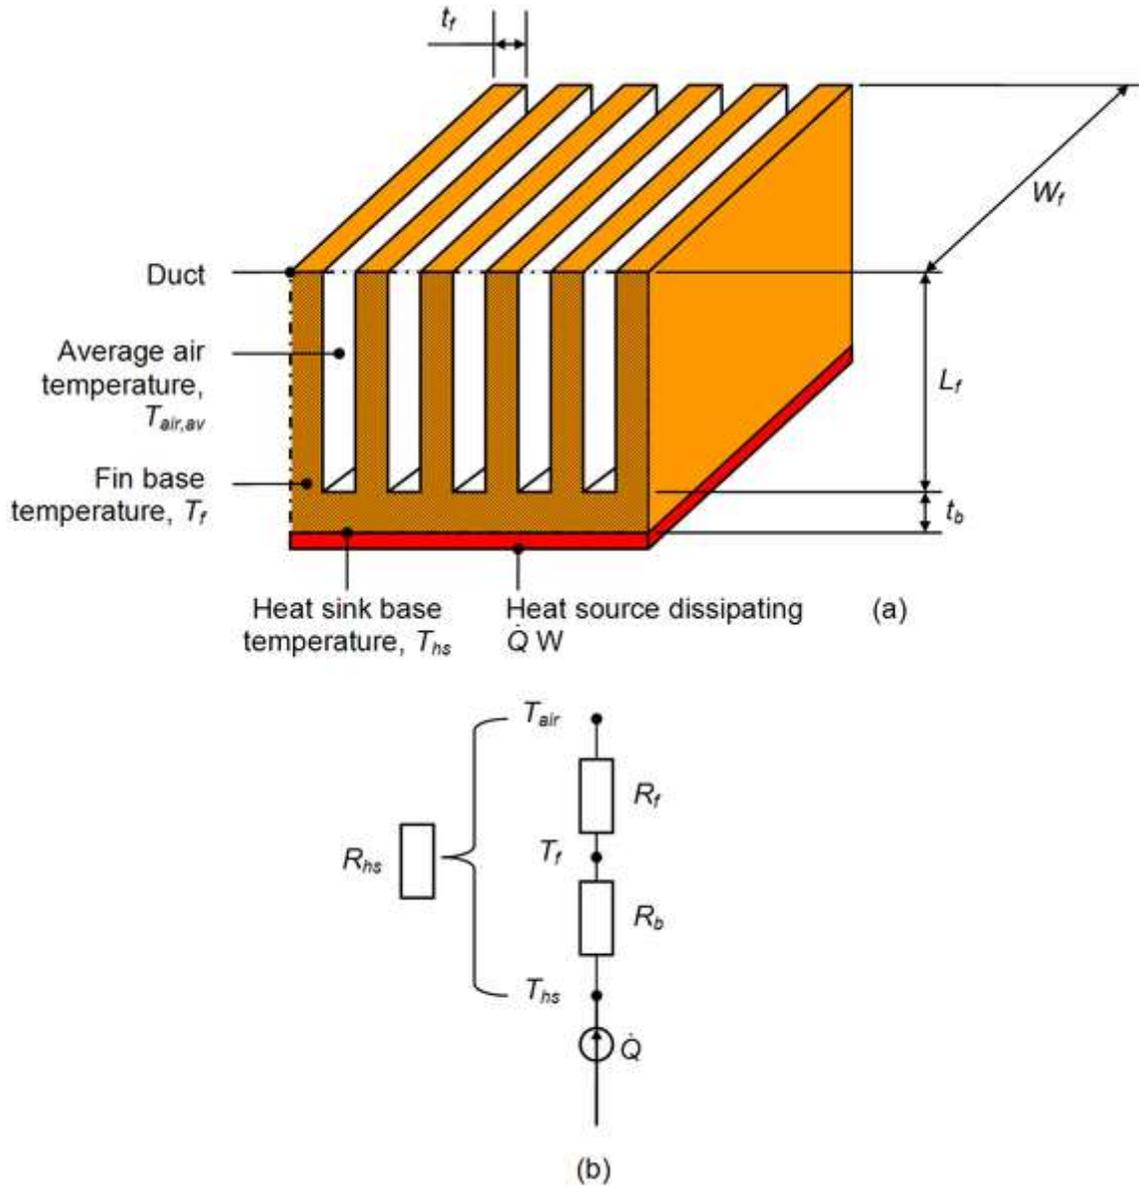


Figure 13: Sketch of a heat sink with equivalent thermal resistances.

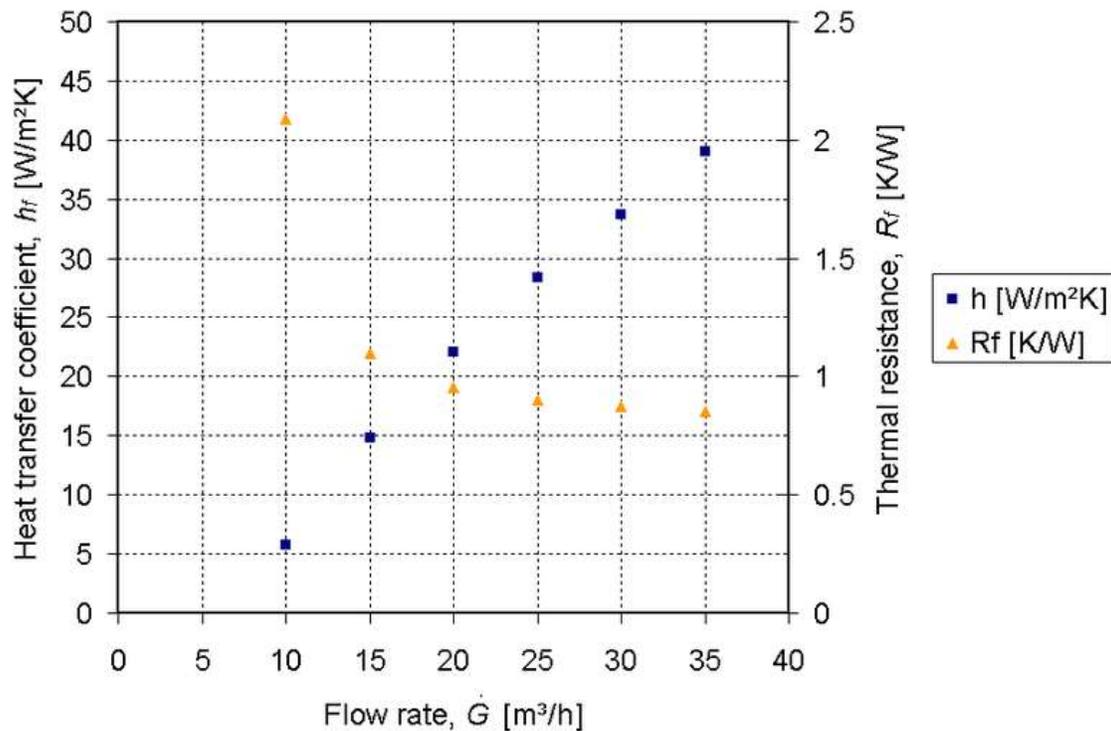


Figure 14: Thermal resistance and heat transfer coefficient plotted against flow rate for the specific heat sink design used in. The data was generated using the equations provided here. The data shows that for an increasing air flow rate, the thermal resistance of the heat sink decreases.

One of the methods to determine the performance of a heat sink is to use heat transfer and fluid dynamics theory. One such method has been published by Jeggels, et al., though this work is limited to ducted flow. Ducted flow is where the air is forced to flow through a channel which fits tightly over the heat sink. This makes sure that all the air goes through the channels formed by the fins of the heat sink. When the air flow is not ducted, a certain percentage of air flow will bypass the heat sink. Flow bypass was found to increase with increasing fin density and clearance, while remaining relatively insensitive to inlet duct velocity.

The heat sink thermal resistance model consists of two resistances, namely the resistance in the heat sink base, R_b , and the resistance in the fins, R_f . The heat sink base thermal resistance, R_b , can be written as follows if the source is a uniformly applied the heat sink base. If it is not, then the base resistance is primarily spreading resistance:

$$R_b = \frac{t_b}{kA_b(4)}$$

where t_b is the heat sink base thickness, k is the heat sink material thermal conductivity and A_b is the area of the heat sink base.

The thermal resistance from the base of the fins to the air, R_f , can be calculated by the following formulas.

$$R_f = \frac{1}{nh_f W_f (t_f + 2\eta_f L_f)} \quad (5)$$

$$\eta_f = \frac{\tanh mL_c}{mL_c} \quad (6)$$

$$mL_c = \sqrt{\frac{2h_f}{kt_f}} L_f \quad (7)$$

$$D_h = \frac{4A_{ch}}{P_{ch}} \quad (8)$$

$$Re = \frac{4\dot{G}\rho}{n\pi D_h \mu} \quad (9)$$

$$f = (0.79 \ln Re - 1.64)^{-2} \quad (10)$$

$$Nu = \frac{(f/8)(Re - 1000)Pr}{1 + 12.7(f/8)^{0.5}(Pr^{\frac{2}{3}} - 1)} \quad (11)$$

$$h_f = \frac{Nuk_{air}}{D_h} \quad (12)$$

$$\rho = \frac{P_{atm}}{R_a T_{in}} \quad (13)$$

The flow rate can be determined by the intersection of the heat sink system curve and the fan curve. The heat sink system curve can be calculated by the flow resistance of the channels and inlet and outlet losses as done in standard fluid mechanics text books, such as Potter, et al. and White.

Once the heat sink base and fin resistances are known, then the heat sink thermal resistance, R_{hs} can be calculated as:

$$R_{hs} = R_b + R_f \quad (14)$$

Using the equations 5 to 13 and the dimensional data in, the thermal resistance for the fins was calculated for various air flow rates. The data for the thermal resistance and heat transfer coefficient are shown in Figure 14. It shows that for an increasing air flow rate, the thermal resistance of the heat sink decreases.

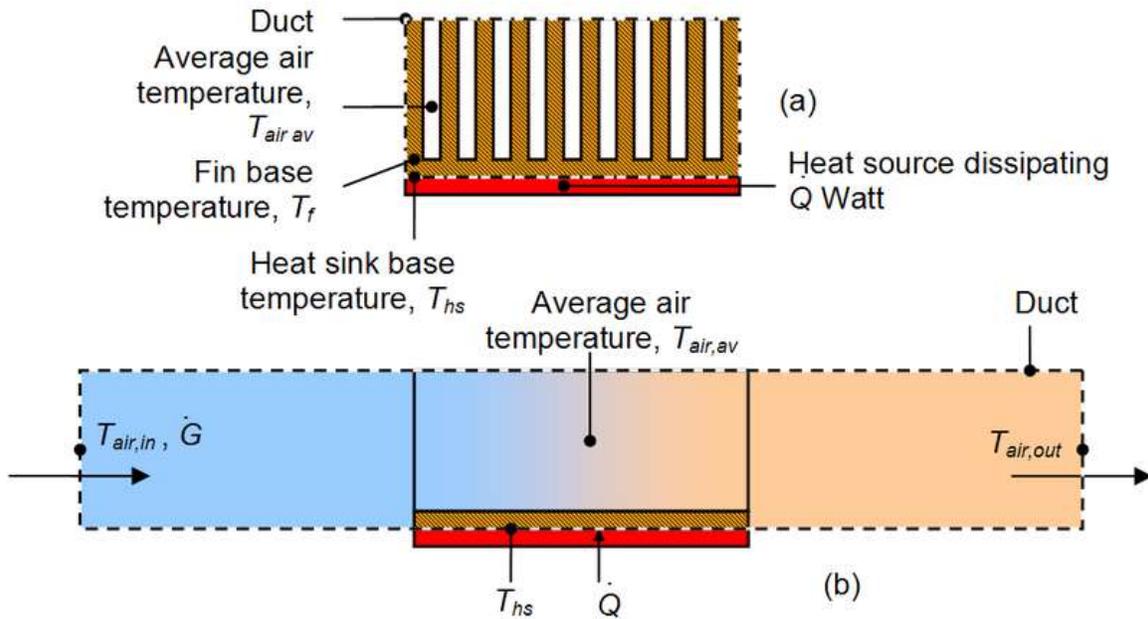


Figure 15: Sketch of a heat sink in a duct used to calculate the governing equations from conservation of energy and Newton's law of cooling.

Experimental methods

Experimental tests are one of the more popular ways to determine the heat sink thermal performance. In order to determine the heat sink thermal resistance, the flow rate, input power, inlet air temperature and heat sink base temperature need to be known, as shown in figure 15. However, figure 15 shows a test setup for a ducted flow heat sink application. Vendor-supplied data is commonly provided for ducted test results. However, the results are optimistic and can give misleading data when heat sinks are used in an unducted application. Another problem with the set up, be it unducted or ducted, is the losses into the board. These must be taken into account. More details on heat sink testing methods and common oversights can be found in Azar, et al.

Numerical methods

In industry, thermal analyses are often ignored in the design process or performed too late — when design changes are limited and become too costly. Of the three methods mentioned here, theoretical and numerical methods can be used to determine an estimate of the heat sink or component temperatures of products before a physical model has been made. A theoretical model is normally used as a first order estimate. Numerical methods or computational fluid dynamics (CFD) provide a qualitative (and sometimes even quantitative) prediction of fluid flows. What this means is that it will give a visual or post-processed result of a simulation, like the images in figures 16 and 17, but the quantitative or absolute accuracy of the result is not guaranteed.

CFD can give an insight into flow patterns that are difficult, expensive or impossible to study using experimental methods. Experiments can give a quantitative description of flow phenomena using measurements for one quantity at a time, at a limited number of points and time instances. If a full scale model is not available or not practical, scale models or dummy models can be used. The experiments can have a limited range of problems and operating conditions. Simulations can give a prediction of flow phenomena using CFD software for all desired quantities, with high resolution in space and time and virtually any problem and realistic operating conditions. However, the results still need to be validated. Another problem with CFD is that the inputs need to be correct. It is the classic case of "Garbage in, garbage out."

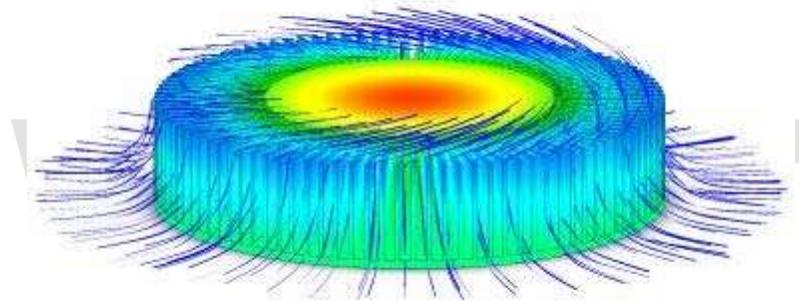


Figure 16: Radial heat sink with thermal profile and swirling forced convection flow trajectories predicted using a CFD analysis package

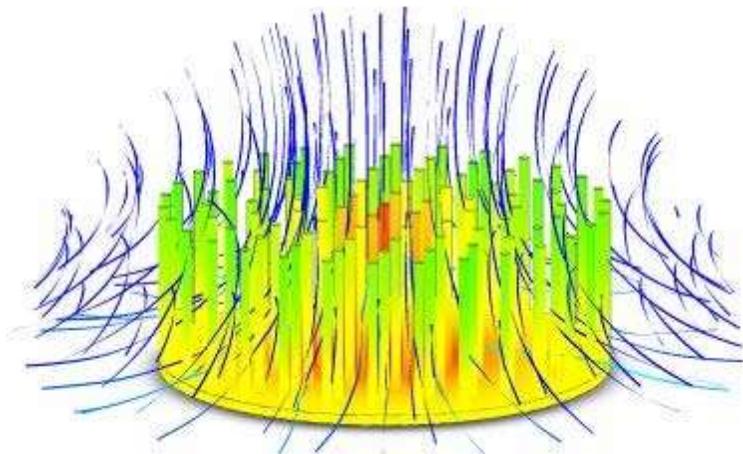


Figure 17: Pin fin heat sink with thermal profile and diode convection flow trajectories predicted using a CFD analysis package

Chapter 4

Laptop Cooler



An Active Laptop cooler



Passive Laptop cooler

A **laptop/notebook cooler, cooler pad or chill mat** is an accessory for laptop computers that helps reduce their operating temperature. Normally used when the laptop's fan device is unable to sufficiently cool the laptop, a cooling pad may house active or passive cooling methods and rests beneath the laptop. Active coolers move air or liquid to direct heat away from the laptop quickly, while passive methods may rely on thermally conductive materials or increasing passive airflow.

Active coolers

Many cooler pads support the use of a power adapter. They typically run on power drawn through one of the laptop's USB ports. Many cooler pads therefore compensate with a built-in USB hub.

Some of these "powered" coolers are designed to draw heat from the underside of the computer; others work in the opposite way by blowing cool air towards the machine. The

fan speed is adjusted manually or automatically on certain models and on others stays at a fixed speed.

The contribution of heat from the energy requirement of the cooler is very minimal impact as USB ports are limited to 2.5 watts of output power, whereas a laptop can easily consume 40 W of electrical power.

Just as there are advantages of having fans in a cooler - there are disadvantages as well. The power drawn by the cooler increases with the wear of the bearings or bushings, or when the vents of fan are blocked by dust, and as friction rises. When the friction becomes significant the power drawn by the cooler increases as well and there have been instances of USB port failures due to overdraw. In a situation like this, although the cooler may appear to work but it may be ineffective or counterproductive by leading to overheating of key components of the mainboard, in rare cases - leading to premature failure of the mainboard.

Passive coolers

Typically, a conductive cooling pad allows for the cooling of a laptop without using any power. These "Pads" are normally filled with an organic salt compound that allows them to absorb the heat from the laptop. They are good for a limited amount of time from around 6–8 hours of cooling. Other designs are simply a pad that elevates the laptop so that the fans in the laptop are allowed greater airflow.

The conductive cooling pads are not advisable for laptops that have fan vents built into the bottom as the cooling pad blocks the vents leading to overheating or premature system failure. The best way to determine if a cooling pad would be suitable for your laptop would be to take a look at the bottom of the laptop and look for air vents or fan vents. If they are on the side and not on the bottom, it is usually safe to use the cooler pad otherwise not.

The other variety that can be used simply as a hard resting surface that provides gap between the cooler and the laptop is normally safer to use. The disadvantages of using these are - they are usually bulky, have to be carried around and sometimes they are too big or too small for a laptop so a bit cumbersome to use.

Multi-Surface Cooler

A Type of Passive Cooler that allows both airflow between the laptop base and cooler, as well as, between the base of the Cooler and Users's lap. These laptop coolers are well suited to laptops that have vents on its base because it prevents these vents from being blocked regardless of what ever surface the laptop is used on. Therefore, these Multi-Surface Coolers are suitable for use on desk, lap and uneven/soft surfaces (couch, bed/duvet, carpet) and the outdoors.

Multipurpose Coolers

Recent advancements have brought forward coolers that are multipurpose. Features include card readers for various forms of media such as keydrives, memory cards, and 2.5" Laptop hard disk drives.

In addition to the above coolers that are a combination of mini work desk with fans are a convenient addition to users that want to use the laptop on a bed or a couch - although they tend to be too heavy and bulky to be carried conveniently everywhere, limiting mobility.

One more variant of this is a cooler with writing pad having an area meant to be used for placing a book or a writing pad – designed with students in mind although the bigger size limits its mobility and the weight usually results in tired legs for the user when used for a prolonged period of time.

One recent addition to the above is an attachable laptop cooler and a comfort pad built into one. It offers combination of cooling, mobility, and comfort as our lives integrate more with mobile computers.

As the laptop gets more powerful and faster it generates more heat. Manufacturers are trying to embed newer technologies to reduce heat generation, improve efficiency, and lower power consumption, which will one day eliminate the use of these devices.

But if Laptop Manufacturer's make smaller and thinner laptops that house higher speed components, then laptops may only be comfortable to use with cooler pads to protect from toasted skin syndrome.

Chapter 5

Thermal Management of Electronic Devices and Systems

Heat generated by electronic devices and circuitry must be dissipated to improve reliability and prevent premature failure. Techniques for heat dissipation can include heatsinks and fans for air cooling, and other forms of computer cooling such as liquid cooling.

In cases of extreme low environmental temperatures, it may actually be necessary to heat the electronic components to achieve satisfactory operation.

Overview

Thermal resistance of devices

This is usually quoted as the thermal resistance from junction to case of the semiconductor device. The units are °C/W. For example, a heatsink rated at 10 °C/W will get 10°C hotter than the surrounding air when it dissipates 1 Watt of heat. Thus, a heatsink with a low °C/W value is more efficient than a heatsink with a high °C/W value.

Thermal time constants

A heatsink's thermal mass can be considered as a capacitor (storing heat instead of charge) and the thermal resistance as an electrical resistance (giving a measure of how fast stored heat can be dissipated). Together, these two components form a thermal RC circuit with an associated time constant given by the product of R and C. This quantity can be used to calculate the dynamic heat dissipation capability of a device, in an analogous way to the electrical case. A specific type of thermal interface material is put between the heat sink and the heat source to increase thermal throughput, such as a microprocessor chip or other power handling semiconductor to stabilise its temperature through increased thermal mass and heat dissipation (primarily by conduction and convection and to a lesser extent by radiation).

Thermal interface material

A **Thermal Interface Material or Mastic** (aka **TIM**) is used to fill the gaps between thermal transfer surfaces, such as between microprocessors and heatsinks, in order to increase thermal transfer efficiency. These gaps are normally filled with air which is a very poor conductor.

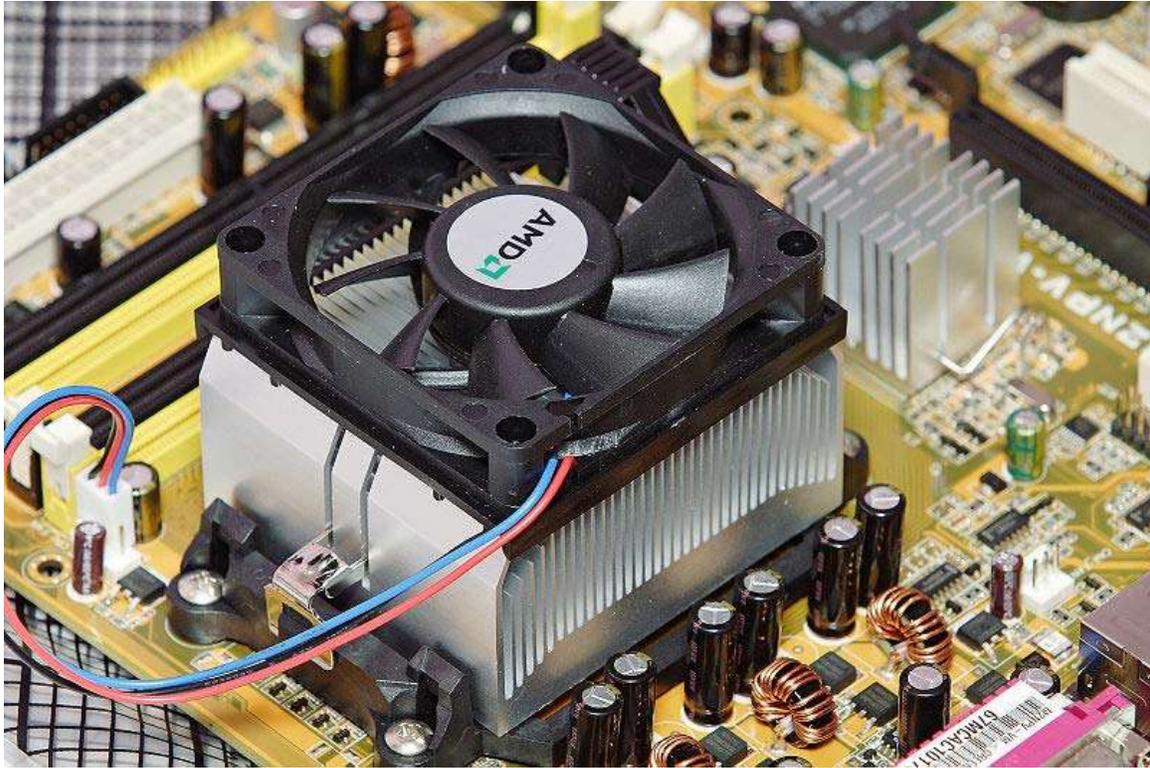
Device heatsinking

Heat sinks

Heat sinks are widely used in electronics, and have become almost essential to modern central processing units. In common use, it is a metal object brought into contact with an electronic component's hot surface — though in most cases, a thin thermal interface material mediates between the two surfaces. Microprocessors and power handling semiconductors are examples of electronics that need a heat sink to reduce their temperature through increased thermal mass and heat dissipation (primarily by conduction and convection and to a lesser extent by radiation). Heat sinks have become almost essential to modern integrated circuits like microprocessors, DSPs, GPUs, and more.

A heat sink usually consists of a metal structure with one or more flat surfaces to ensure good thermal contact with the components to be cooled, and an array of comb or fin like protrusions to increase the surface contact with the air, and thus the rate of heat dissipation.

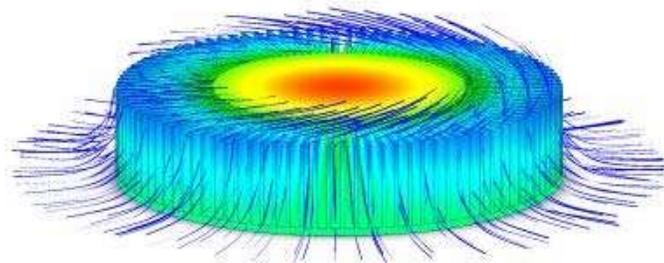
A heat sink is sometimes used in conjunction with a fan to increase the rate of airflow over the heat sink. This maintains a larger temperature gradient by replacing warmed air faster than convection would. This is known as a forced air system.



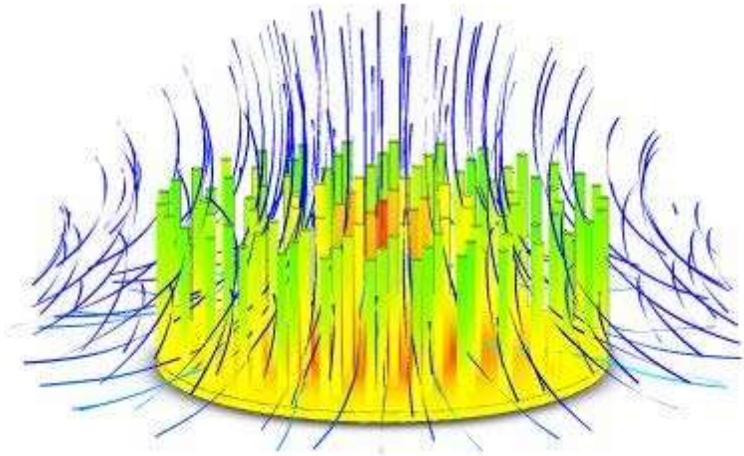
CPU heat sink with fan attached

A **heat sink** (or heatsink) is an environment or object that absorbs and dissipates heat from another object using thermal contact (either direct or radiant). Heat sinks are used in a wide range of applications wherever efficient heat dissipation is required; major examples include refrigeration, heat engines and cooling electronic devices.

Principle



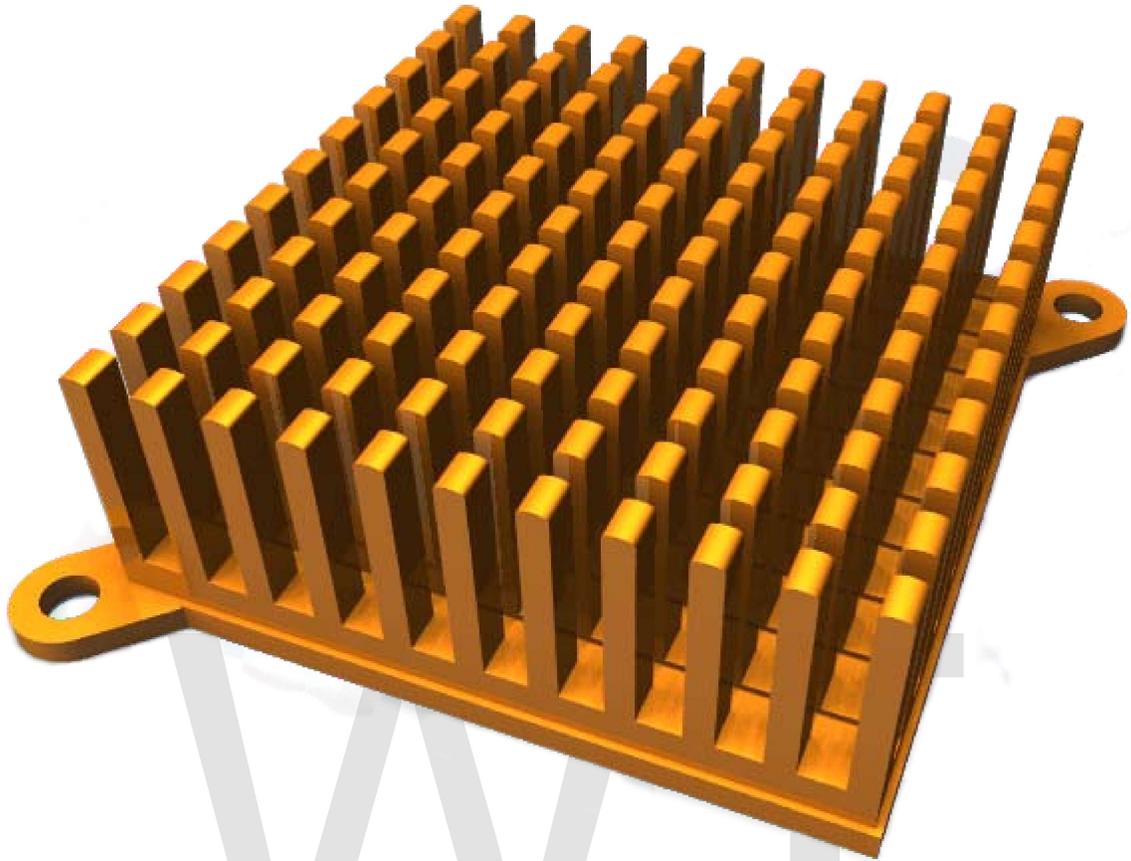
Radial Heat Sink with Thermal Profile and Swirling Forced Convection Flow Trajectories (using CFD analysis)



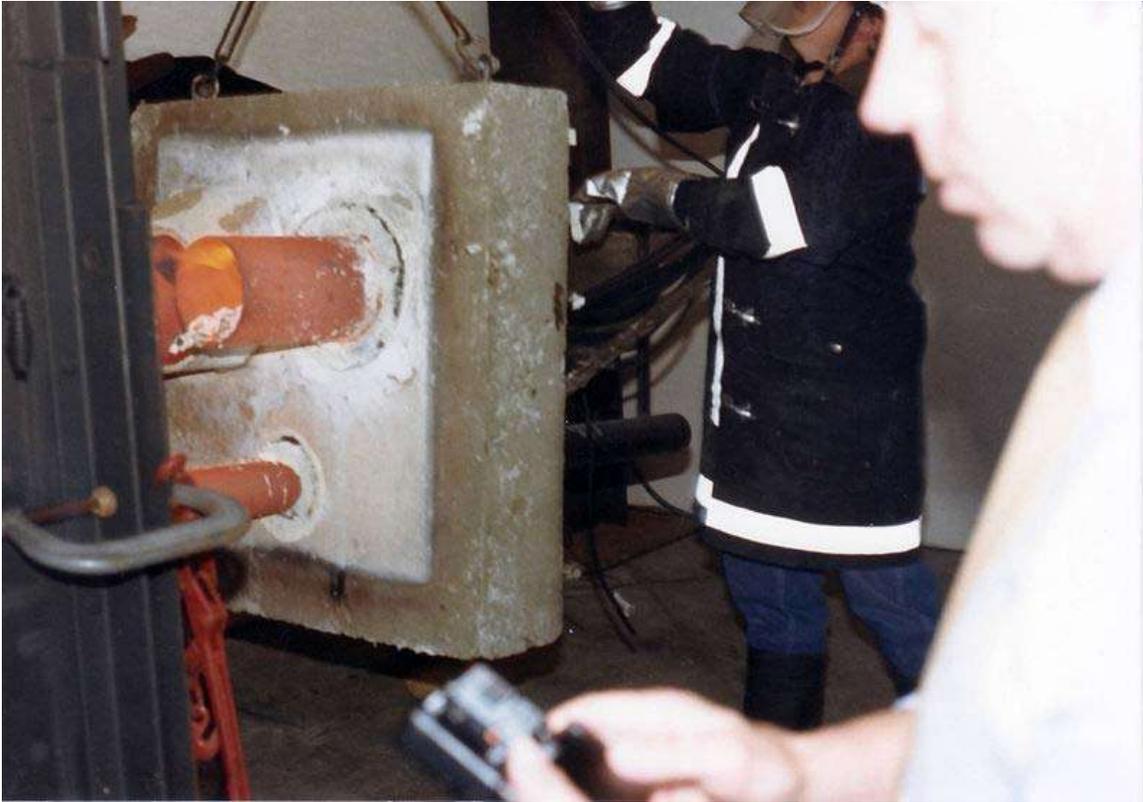
Pin Fin Heat Sink with Thermal Profile and Dione Convection Flow Trajectories (using CFD analysis)



Heat sink in a workstation computer



A motherboard heat sink



Fire test where the steel pipe penetrants clearly act to absorb and conduct heat from the furnace, through to the unexposed side.

Heat sinks function by efficiently transferring thermal energy ("heat") from an object at high temperature to a second object at a lower temperature with a much greater heat capacity. This rapid transfer of thermal energy quickly brings the first object into thermal equilibrium with the second, lowering the temperature of the first object, fulfilling the heat sink's role as a cooling device. Efficient function of a heat sink relies on rapid transfer of thermal energy from the first object to the heat sink, and the heat sink to the second object.

The most common design of a heat sink is a metal device with many fins. The high thermal conductivity of the metal combined with its large surface area result in the rapid transfer of thermal energy to the surrounding, cooler, air. This cools the heat sink and whatever it is in direct thermal contact with. Use of fluids (for example coolants in refrigeration) and thermal interface material (in cooling electronic devices) ensures good transfer of thermal energy to the heat sink. Similarly, a fan may improve the transfer of thermal energy from the heat sink to the air.

Construction and materials

A heat sink usually consists of a base with one or more flat surfaces and an array of comb or fin-like protrusions to increase the heat sink's surface area contacting the air, and thus

increasing the heat dissipation rate. While a heat sink is a static object, a fan often aids a heat sink by providing increased airflow over the heat sink — thus maintaining a larger temperature gradient by replacing the warmed air more quickly than passive convection achieves alone — this is known as a forced-air system.

Ideally, heat sinks are made from a good thermal conductor such as silver, gold, copper, or aluminum alloy. Copper and aluminum are among the most-frequently used materials for this purpose within electronic devices. Copper (401 W/(m·K) at 300 K) is significantly more expensive than aluminum (237 W/(m·K) at 300 K) but is also roughly twice as efficient as a thermal conductor. Aluminum has the significant advantage that it can be easily formed by extrusion, thus making complex cross-sections possible. Aluminum is also much lighter than copper, offering less mechanical stress on delicate electronic components. Some heat sinks made from aluminum have a copper core as a trade off. The heat sink's contact surface (the base) must be flat and smooth to ensure the best thermal contact with the object needing cooling. Frequently a thermally conductive grease is used to ensure optimal thermal contact; such compounds often contain colloidal silver. Further, a clamping mechanism, screws, or thermal adhesive hold the heat sink tightly onto the component, but specifically without pressure that would crush the component.

Performance

Heat sink performance (including free convection, forced convection, liquid cooled, and any combination thereof) is a function of material, geometry, and overall surface heat transfer coefficient. Generally, forced convection heat sink thermal performance is improved by increasing the thermal conductivity of the heat sink materials, increasing the surface area (usually by adding extended surfaces, such as fins or foam metal) and by increasing the overall area heat transfer coefficient (usually by increase fluid velocity, such as adding fans, pumps, etc.).

Online heat sink calculators from companies such as Novel Concepts, Inc., can accurately estimate forced convection heat sink performance. For more complex heat sink geometries, and/or heat sinks with multiple materials, and/or heat sinks with multiple fluids, computation fluid dynamics (CFD) analysis is recommended.

Use in electronics

PC marketplace

Due to recent technological developments and public interest, the retail heat sink market has reached an all time high. In the early 2000s, CPUs were produced that emitted more and more heat than earlier, escalating requirements for quality cooling systems.

Overclocking has always meant greater cooling needs, and the inherently hotter chips meant more concerns for the enthusiast. Efficient heat sinks are vital to overclocked computer systems because the higher a microprocessor's cooling rate, the faster the

computer can operate without instability; generally, faster operation leads to higher performance. Many companies now compete to offer the best heat sink for PC overclocking enthusiasts. Prominent aftermarket heat sink manufacturers include: Aero Cool, Foxconn, Thermalright, Thermaltake, Swiftech, and Zalman.

In soldering

Temporary heat sinks were sometimes used while soldering circuit boards, preventing excessive heat from damaging sensitive nearby electronics. In the simplest case, this means partially gripping a component using a heavy metal crocodile clip or similar clamp. Modern semiconductor devices, which are designed to be assembled by reflow soldering, can usually tolerate soldering temperatures without damage. On the other hand, electrical components such as magnetic reed switches can malfunction if exposed to higher powered soldering irons, so this practice is still very much in use.

In Batteries

In the battery used for electric vehicles, Nominal battery performance is usually specified for working temperatures somewhere in the + 20°C to +30°C range however the actual performance can deviate substantially from this if the battery is operated at higher or in particular lower temperatures, so some electric cars have heating and cooling for their batteries

Recent developments

More recently, synthetic diamond cooling sinks are being researched to provide better cooling. Also, some heat sinks are constructed of multiple materials with desirable characteristics, such as phase change materials, which can store a great deal of energy due to their heat of fusion.

Convective air cooling

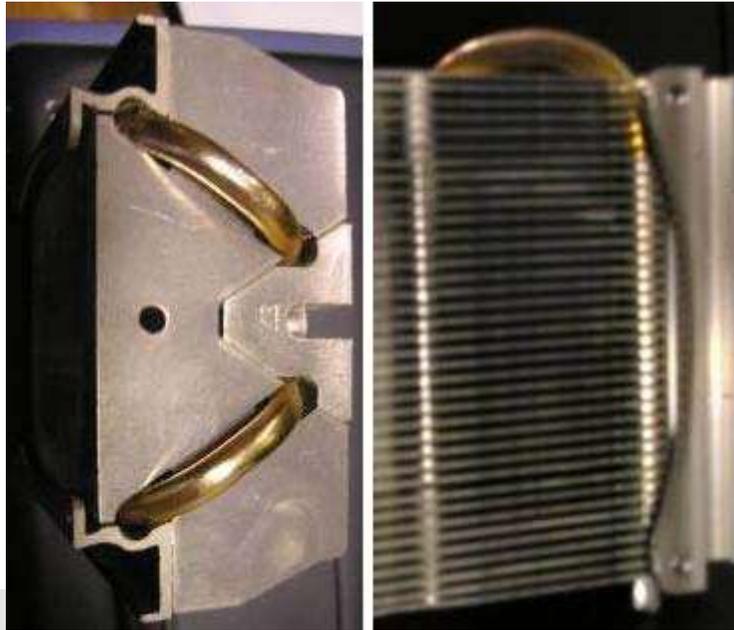
This term describes device cooling by the convection currents of the warm air being allowed to escape the confines of the component to be replaced by cooler air. Since warm air normally rises, this method usually requires venting at the top or sides of the casing to be effective.

Forced air cooling

If there is more air being forced into a system than being pumped out (due to an imbalance in the number of fans), this is referred to as a 'positive' airflow, as the pressure inside the unit is higher than outside.

A balanced or neutral airflow is the most efficient, although a slightly positive airflow can result in less dust build up if filtered properly.

Heat pipes



A heat sink (aluminium) incorporating a heat pipe (copper)

A **heat pipe** is a heat transfer mechanism that can transport large quantities of heat with a very small difference in temperature between the hot and cold interfaces. A typical heat pipe consists of sealed hollow tube made of a thermoconductive metal such as copper or aluminium. The pipe contains a relatively small quantity of a "working fluid" or coolant (such as water, ethanol or mercury) with the remainder of the pipe being filled with the vapour phase of the working fluid, all other gases being excluded. The advantage of heat pipes is their great efficiency in transferring heat. They are actually more "conductive" than a copper bar of equivalent cross-section.

Peltier cooling plates

Peltier cooling plates take advantage of what is known as the Peltier effect to create a heat flux between the junction of two different types of materials. This effect is commonly used for cooling electronic components and small instruments.

There are no moving parts and such a device is maintenance free. Due to the relatively low efficiency, thermoelectric cooling is generally only used in environments where the solid state nature outweighs the poor efficiency. Thermoelectric junctions are generally only around 10% as efficient as the ideal refrigerator (Carnot cycle), compared with 40% achieved by conventional compression cycle systems.

Synthetic Jet Air Cooling

A Synthetic Jet is produced by a continual flow of vortices that are formed by alternating brief ejection and suction of air across an opening such that the net mass flux is zero. A

unique feature of these jets is that they are formed entirely from the working fluid of the flow system in which they are deployed can produce a net momentum to the flow of a system without net mass injection to the system.

Synthetic jet air movers have no moving parts and are thus maintenance free. Due to the high heat transfer coefficients, high reliability but lower overall flow rates, Synthetic jet air movers are usually used at the chip level and not at the system level for cooling. However depending on the size and complexity of the systems they can be used for both at times.

WWT

Chapter 6

Fan Control

Fan control is the management of the rotational speed of an electric fan. There are many types of electric fans and many types of fan controls.

Need for fan control

As modern PCs grow more powerful so do their requirements for electrical power. Computers convert most of this electrical power into heat generated by all major components.

Some early generation PCs did not need active ventilation. Power supplies eventually needed forced cooling, and soon took up the duty of cooling the rest of the PC with the ATX standard. The byproduct of increased heat generation is that the fan(s) need to move increasing amounts air and thus, need to be more powerful. Since they must move more air through the same area of space, fans will naturally become more noisy.

In fact, if one installs extra fans in a PC case, the noise levels can reach 70 dB. Since fan noise increases with the fifth power of the fan rotation speed, reducing rotations per minute (RPM) by a small amount potentially means a reduction in fan noise. This must be done cautiously, as excessive reduction in speed may cause components to overheat and be damaged. If done properly fan noise can be drastically reduced.

Fan types

Common cooling fans have at least two pins on the connector. These two pins connect to the fan motor, usually a brushless DC type. Extra pins may be present, with either or both of the following features:

- A tachometer that measures the actual speed of the fan as a pulse train, frequency being proportional to speed.
- Control of fan speed using PWM (Pulse-width modulation.)

Two-pin fans operate either as an on/off fan, or can be controlled by varying the voltage. will have either two, three, or four pins on the connector.

Types of control

No control

The simplest method of fan control is simply to leave the fan on all the time. This creates quite a bit of noise and increases power requirements, but keeps the system the coolest.

Thermostatic

In this style of fan control, the fan is either on or off. A system thermistor checks the temperature inside the chassis, and if it detects a temperature outside of range, it spins the fans up to maximum. When the temperature drops below a threshold again, the fans are turned back off. This control method reduces power requirements during periods of low usage, but when the system is operating at capacity, the fan noise can become a problem again.

Linear voltage regulation

A standard cooling fan is essentially a bladed DC motor. By varying the voltage input across the acceptable range for a fan, the speed of the fan will increase (to added voltage) and decrease (to reduced voltage). A faster fan, obviously, means more air moved, and thus a higher heat exchange rate. There are a few ways to perform this regulation.

Resistors

Resistors are the simplest method of reducing fan noise, but they add to the heat generated inside the computer case. Since the voltage drop is proportional to the current, the fan may not start. They need to be of the appropriate power rating (i.e. higher than the fan). For variable fan control, potentiometers could be used along with a transistor such as a MOSFET whose output voltage is controlled by the potentiometer. It is possible to use a rheostat instead.

Diodes

A diode in series with the fan will reduce the voltage being output to the fan. You can use a zener diode (select one for the desired voltage drop) or a silicon diode (Produce the required voltage drop by connecting multiple diodes in series. Each diode reduces the voltage by approximately 0.75 volts.)

Voltmodding

Voltmodding describes the practice of varying the voltage fed to a component; in this case, a computer fan. This can be achieved by connecting the ground wire of the fan to the +5V rail and the positive wire to the +12V rail of a typical PC power supply to achieve a theoretical +7V (positive seven volts). However, this is a potentially risky method: the parts on the +5V power line might be exposed to overvoltage in case of a short in the fan. A less common variation is to increase the voltage to the fan by connecting the ground wire to a -12V rail (located on the motherboard ATX connector) and the positive wire to the +12V rail, producing a total of +24V. 24V can destroy the fans, other options are 17V and 15.3V.

Pulse-width modulation

Pulse-width modulation (PWM) is a common method of controlling fans. Modern computer motherboard PWM control when used with multi-core CPUs reads data from Digital Temperature Sensors on each core of the CPU.

Unlike the linear methods above that are based on voltage loss, PWM switches the input voltage between (nearly) fully on and fully off. This means there is practically no voltage or power loss and associated heat output. PWM controller can be a relatively small, low-power and cool-running, albeit complex, component that doesn't require heavy duty resistors, diodes or transistors and associated heatsinking.

Fan speed controllers



A fan controller with LEDs to indicate fan status and potentiometer and switch to control fan speed.

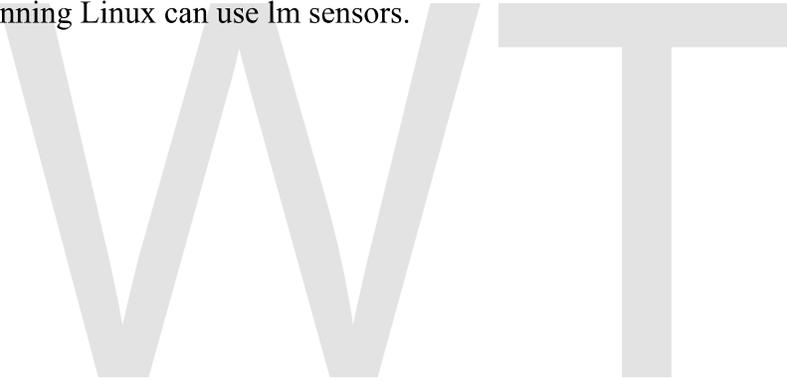
Another method, popular with gamers, is the manual fan speed controller. They can be mounted in an expansion slot, a 5.25" or 3.5" drive bay or come built into the computer's case. Using switches or knobs, attached fans can have their speeds adjusted by one of the above methods.

Software

The method by which the software physically controls the fan is usually PWM. Many companies now provide software to control fan speeds on their motherboards under Microsoft Windows.

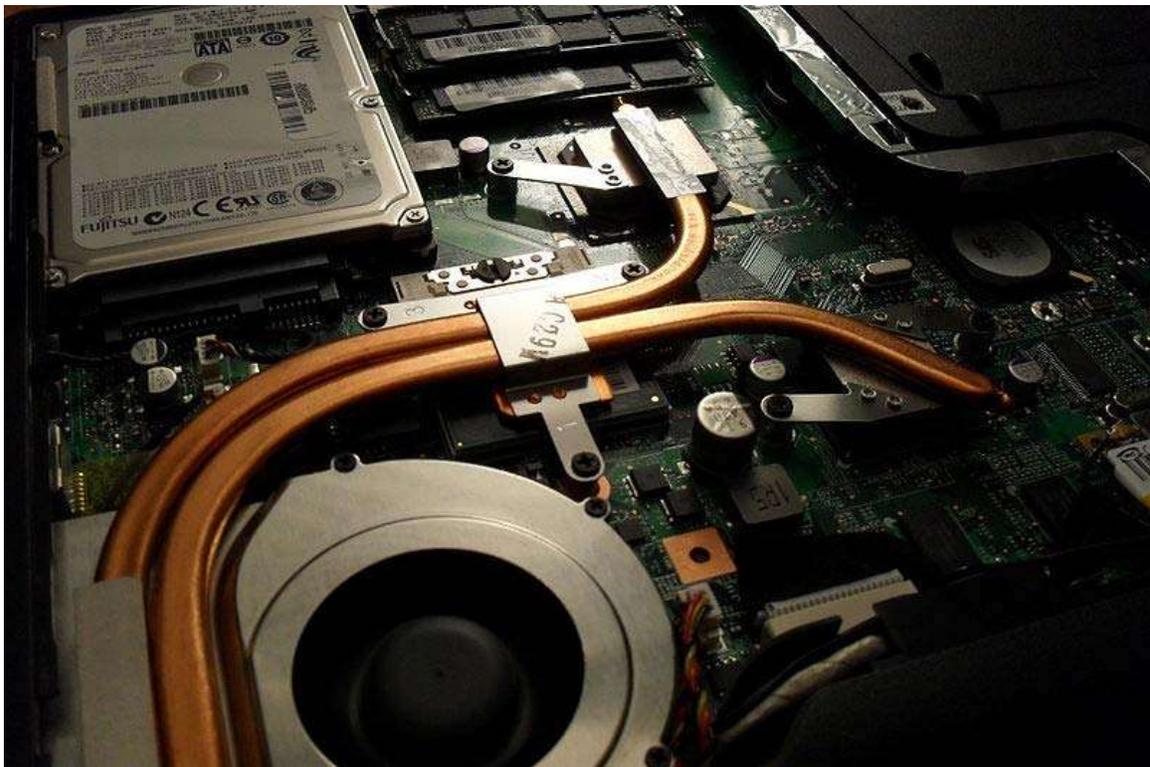
- AOpen motherboards can use "SilentTEK".
- ASUS motherboards can use "Q-Fan".
- MSI motherboards can use "Core Center".
- Universal abit motherboards can use "μGuru".
- Gigabyte motherboards can use "EasyTune 6".
- Intel desktop boards (older socket 478 etc) use "Active Monitor" and "Desktop Control Centre".
- Intel desktop boards (newer socket 775 etc) use "Desktop Utilities".
- Dell Inspiron/Latitude/Precision computers can use "I8kfanGUI".
- Various computers can use the freeware "SpeedFan".

Computers running Linux can use lm sensors.



Chapter 7

Heat Pipe



A laptop heat pipe system

A **heat pipe** is a heat transfer mechanism that combines the principles of both thermal conductivity and phase transition to efficiently manage the transfer of heat between two solid interfaces.

At the hot interface within a heat pipe, which is typically at a very low pressure, a liquid in contact with a thermally conductive solid surface turns into a vapor by absorbing heat from that surface. The vapor condenses back into a liquid at the cold interface, releasing the latent heat. The liquid then returns to the hot interface through either capillary action or gravity action where it evaporates once more and repeats the cycle. In addition, the

internal pressure of the heat pipe can be set or adjusted to facilitate the phase change depending on the demands of the working conditions of the thermally managed system.

Structure, design and construction

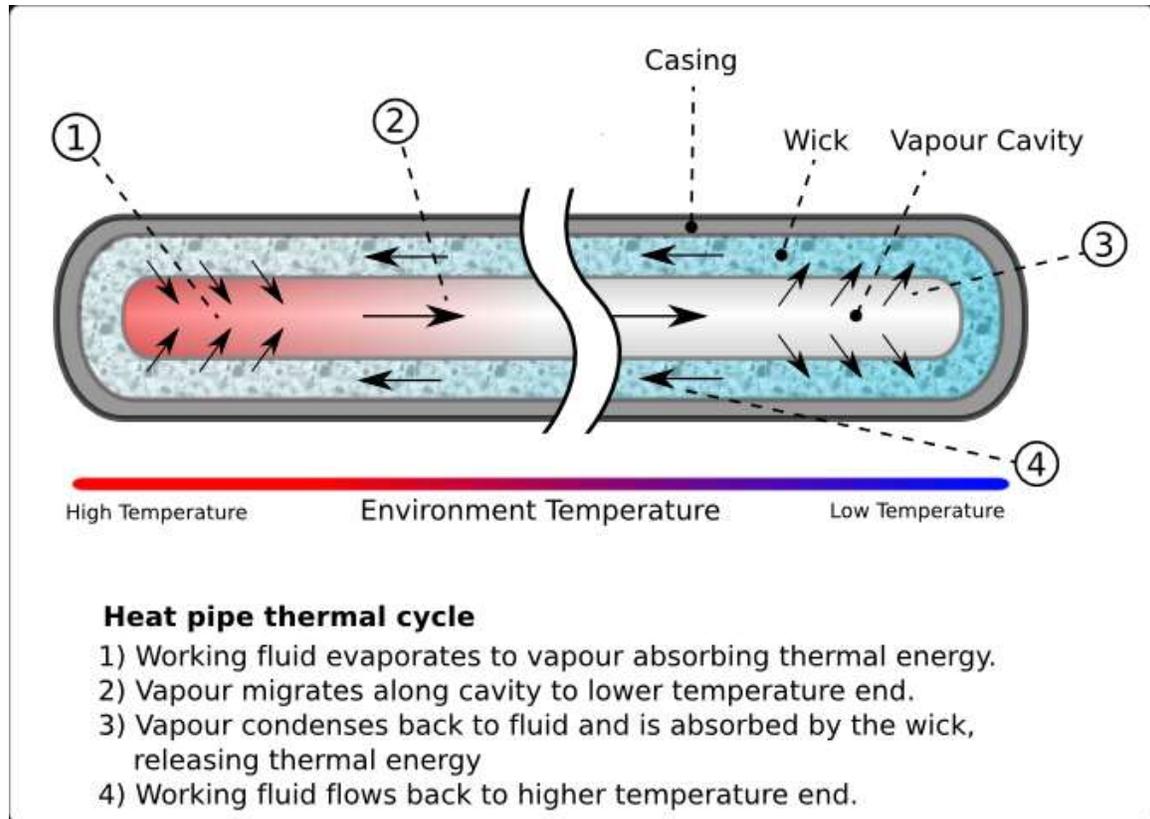
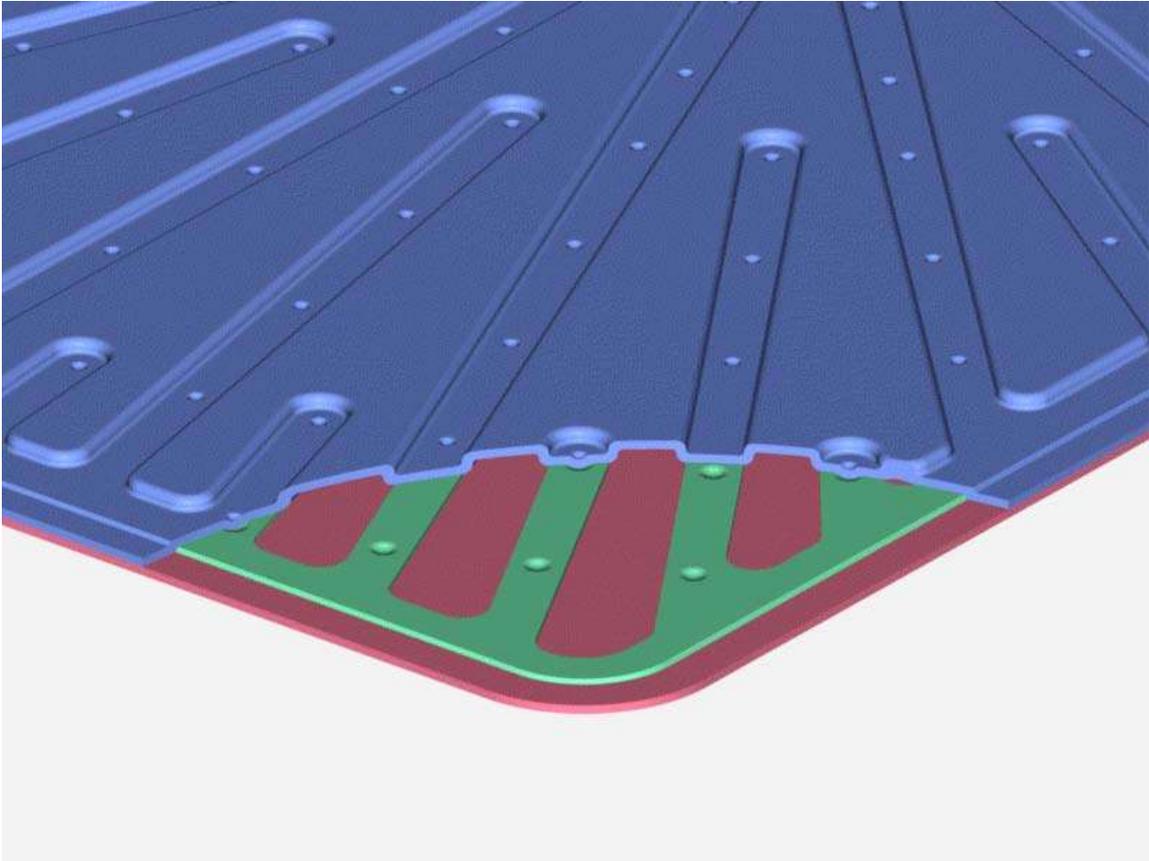
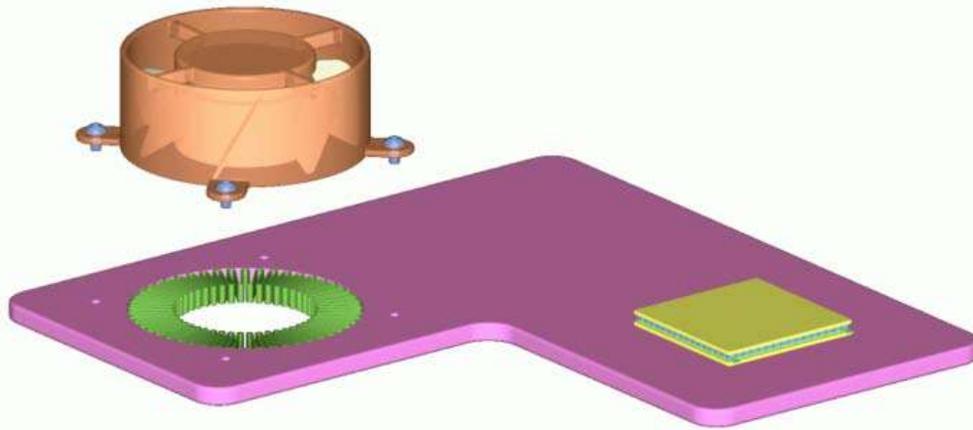


Diagram showing components and mechanism for a heat pipe containing a wick



Cut-away view of a 500 μm thick flat heat pipe, with a thin planar capillary (aqua colored)



Thin flat heat pipe (heat spreader) with remote heat sink and fan

A typical heat pipe consists of a sealed pipe or tube made of a material with high thermal conductivity such as copper or aluminium at both hot and cold ends. A vacuum pump is used to remove all air from the empty heat pipe, and then the pipe is filled with a fraction of a percent by volume of *working fluid* (or coolant) chosen to match the operating temperature. Examples of such fluids include water, ethanol, acetone, sodium, or mercury. Due to the partial vacuum that is near or below the vapor pressure of the fluid, some of the fluid will be in the liquid phase and some will be in the gas phase. The use of a vacuum eliminates the need for the working gas to diffuse through any other gas and so the bulk transfer of the vapor to the cold end of the heat pipe is at the speed of the moving molecules. In this sense, the only practical limit to the rate of heat transfer is the speed with which the gas can be condensed to a liquid at the cold end.

Inside the pipe's walls, an optional wick structure exerts a capillary pressure on the liquid phase of the working fluid. This is typically a sintered metal powder or a series of grooves parallel to the pipe axis, but it may be any material capable of exerting capillary pressure on the condensed liquid to wick it back to the heated end. The heat pipe may not need a wick structure if gravity or some other source of acceleration is sufficient to overcome surface tension and cause the condensed liquid to flow back to the heated end.

A heat pipe is not a thermosiphon, because there is no siphon. Thermosiphons transfer heat by single-phase convection.

Heat pipes contain no mechanical moving parts and typically require no maintenance, though non-condensing gases (that diffuse through the pipe's walls, result from breakdown of the working fluid, or exist as impurities in the materials) may eventually reduce the pipe's effectiveness at transferring heat. This is significant when the working fluid's vapour pressure is low.

The materials chosen depend on the temperature conditions in which the heat pipe must operate, with coolants ranging from liquid helium for extremely low temperature applications (2–4 K) to mercury (523–923 K) & sodium (873–1473 K) and even indium (2000–3000 K) for extremely high temperatures. The vast majority of heat pipes for low temperature applications use some combination of ammonia (213–373 K), alcohol (methanol (283–403 K) or ethanol (273–403 K)) or water (303–473 K) as working fluid. Since the heat pipe contains a vacuum, the working fluid will boil and hence take up latent heat at well below its boiling point at atmospheric pressure. Water, for instance, will boil at just above 273 K (0 degrees Celsius) and so can start to effectively transfer latent heat at this low temperature.

The advantage of heat pipes over many other heat-dissipation mechanisms is their great efficiency in transferring heat. They are a fundamentally better heat conductor than an equivalent cross-section of solid copper (a heat sink alone, though simpler in design and construction, does not take advantage of the principle of matter phase transition). Some heat pipes have demonstrated a heat flux of more than 230 MW/m², nearly four times the heat flux at the surface of the sun.

Active control of heat flux can be effected by adding a variable volume liquid reservoir to the evaporator section. Variable conductance heat pipes employ a large reservoir of inert immiscible gas attached to the condensing section. Varying the gas reservoir pressure changes the volume of gas charged to the condenser which in turn limits the area available for vapor condensation. Thus a wider range of heat fluxes and temperature gradients can be accommodated with a single design.

A modified heat pipe with a reservoir having no capillary connection to the heat pipe wick at the evaporator end can also be used as a thermal diode. This heat pipe will transfer heat in one direction, acting as an insulator in the other.

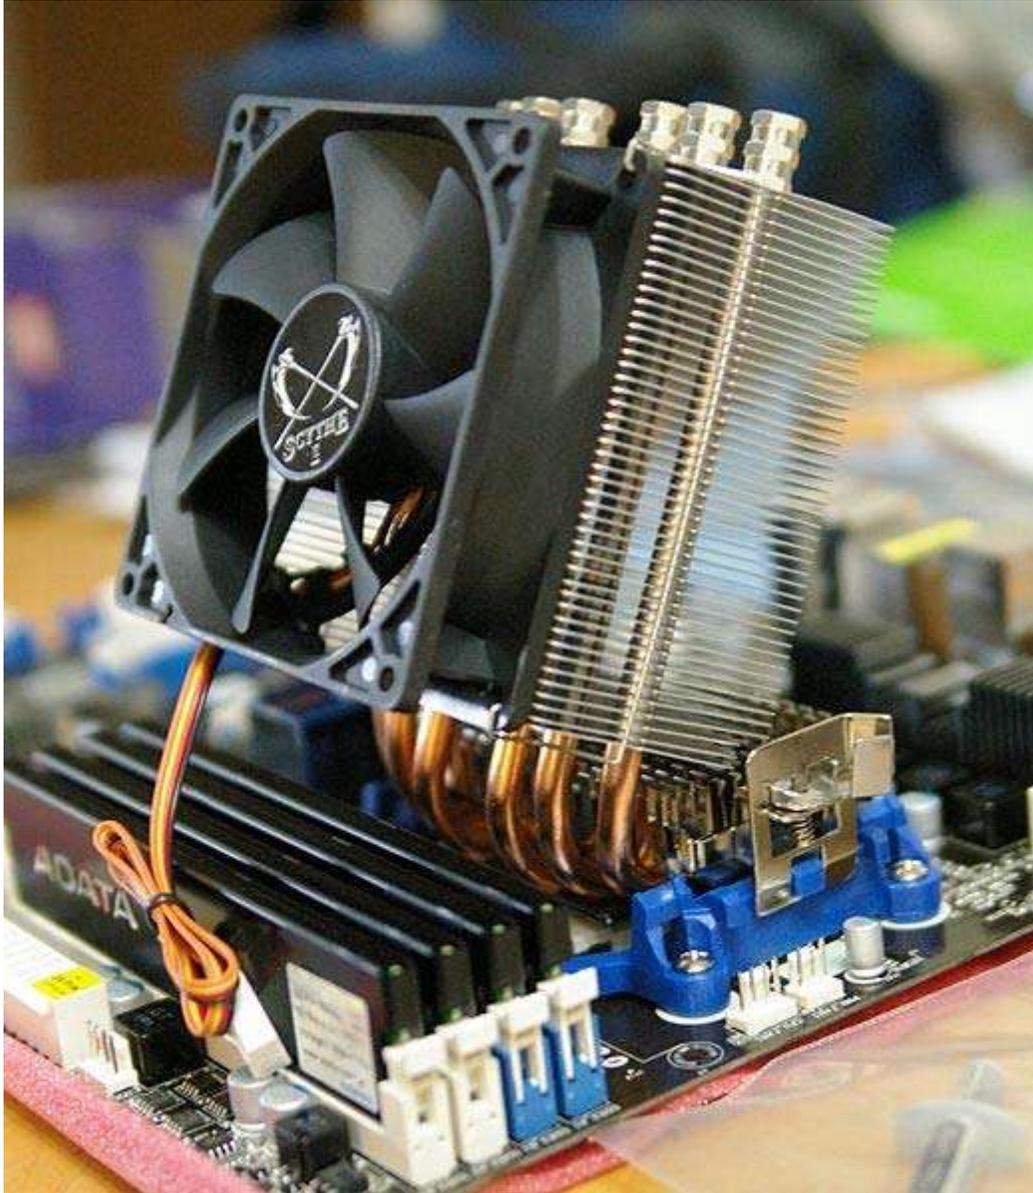
Vapor Chamber or Flat heat pipes

Thin planar heat pipes (heat spreaders) have the same primary components as tubular heat pipes. These components are a hermetically sealed hollow vessel, a working fluid, and a closed-loop capillary recirculation system.

Compared to a one-dimensional tubular heat pipe, the width of a two-dimensional heat pipe allows an adequate cross section for heat flow even with a very thin device. These

thin planar heat pipes are finding their way into “height sensitive” applications, such as notebook computers, and surface mount circuit board cores. It is possible to produce flat heat pipes as thin as 0.5 mm (thinner than a credit card).

Heat transfer



A heat sink (aluminium) with heat pipe (copper)

Heat pipes employ evaporative cooling to transfer thermal energy from one point to another by the evaporation and condensation of a working fluid or coolant. Heat pipes rely on a temperature difference between the ends of the pipe, and cannot lower temperatures at either end beyond the ambient temperature (hence they tend to equalise the temperature within the pipe).

When one end of the heat pipe is heated the working fluid inside the pipe at that end evaporates and increases the vapour pressure inside the cavity of the heat pipe. The latent heat of evaporation absorbed by the vaporisation of the working fluid reduces the temperature at the hot end of the pipe.

The vapour pressure over the hot liquid working fluid at the hot end of the pipe is higher than the equilibrium vapour pressure over condensing working fluid at the cooler end of the pipe, and this pressure difference drives a rapid mass transfer to the condensing end where the excess vapour condenses, releases its latent heat, and warms the cool end of the pipe. Non-condensing gases (caused by contamination for instance) in the vapour impede the gas flow and reduce the effectiveness of the heat pipe, particularly at low temperatures, where vapour pressures are low. The velocity of molecules in a gas is approximately the speed of sound and in the absence of non condensing gases, this is the upper velocity with which they could travel in the heat pipe. In practice, the speed of the vapour through the heat pipe is dependent on the rate of condensation at the cold end.

The condensed working fluid then flows back to the hot end of the pipe. In the case of vertically-oriented heat pipes the fluid may be moved by the force of gravity. In the case of heat pipes containing wicks, the fluid is returned by capillary action.

When making heat pipes, there is no need to create a vacuum in the pipe. One simply boils the working fluid in the heat pipe until the resulting vapour has purged the non condensing gases from the pipe and then seals the end.

An interesting property of heat pipes is the temperature over which they are effective. Initially, it might be suspected that a water charged heat pipe would only work when the hot end reached the boiling point (100 °C) and steam was transferred to the cold end. However, the boiling point of water is dependent on absolute pressure inside the pipe. In an evacuated pipe, water will boil just slightly above its melting point (0 °C). The heat pipe will operate, therefore, when the hot end is just slightly warmer than the melting point of the working fluid. Similarly, a heat pipe with water as a working fluid can work well above the boiling point (100 °C), if the cold end is low enough in temperature to condense the fluid.

The main reason for the effectiveness of heat pipes is the evaporation and condensation of the working fluid. The heat of vaporization greatly exceeds the sensible heat capacity. Using water as an example, the energy needed to evaporate one gram of water is equivalent to the amount of energy needed to raise the temperature of that same gram of water by 540 °C (hypothetically, if the water was under extremely high pressure so it didn't vaporize or freeze over this temperature range). Almost all of that energy is rapidly transferred to the "cold" end when the fluid condenses there, making a very effective heat transfer system with no moving parts.

Origins and research in the United States

The general principle of heat pipes using gravity (commonly classified as two phase thermosiphons) dates back to the steam age. The modern concept for a capillary driven heat pipe was first suggested by R.S. Gaugler of General Motors in 1942 who patented the idea. The benefits of employing capillary action were independently developed and first demonstrated by George Grover at Los Alamos National Laboratory in 1963 and subsequently published in the Journal of Applied Physics in 1964. Grover noted in his notebook:

"Heat transfer via capillary movement of fluids. The "pumping" action of surface tension forces may be sufficient to move liquids from a cold temperature zone to a high temperature zone (with subsequent return in vapor form using as the driving force, the difference in vapor pressure at the two temperatures) to be of interest in transferring heat from the hot to the cold zone. Such a closed system, requiring no external pumps, may be of particular interest in space reactors in moving heat from the reactor core to a radiating system. In the absence of gravity, the forces must only be such as to overcome the capillary and the drag of the returning vapor through its channels."

Between 1964 and 1966, RCA was the first corporation to undertake research and development of heat pipes for commercial applications (though their work was mostly funded by the US government). During the late 1960s NASA played a large role in heat pipe development by funding a significant amount of research on their applications and reliability in space flight following from Grover's suggestion. NASA's attraction to heat pipe cooling systems was understandable given their low weight, high heat flux, and zero power draw. Their primary interest however was based on the fact that the system wouldn't be adversely affected by operating in a zero gravity environment. The first application of heat pipes in the space program was in thermal equilibration of satellite transponders. As satellites orbit, one side is exposed to the direct radiation of the sun while the opposite side is completely dark and exposed to the deep cold of outer space. This causes severe discrepancies in the temperature (and thus reliability and accuracy) of the transponders. The heat pipe cooling system designed for this purpose managed the high heat fluxes and demonstrated flawless operation with and without the influence of gravity. The developed cooling system was the first description and usage of variable conductance heat pipes to actively regulate heat flow or evaporator temperature.

Corporate R&D

Publications in 1967 and 1968 by Feldman, Eastman, & Katzoff first discussed applications of heat pipes to areas outside of government concern and that did not fall under the high temperature classification such as: air conditioning, engine cooling, and electronics cooling. These papers also made the first mentions of flexible, arterial, and flat plate heat pipes. 1969 publications introduced the concepts of the rotational heat pipe with its applications to turbine blade cooling and the first discussions of heat pipe applications to cryogenic processes.

Starting in the 1980s Sony began incorporating heat pipes into the cooling schemes for some of its commercial electronic products in place of both forced convection and passive finned heat sinks. Initially they were used in tuners & amplifiers, soon spreading to other high heat flux electronics applications. During the late 1990s increasingly hot microcomputer CPUs spurred a threefold increase in the number of U.S. heat pipe patent applications. As heat pipes transferred from a specialized industrial heat transfer component to a consumer commodity most development and production moved from the U.S. to Asia. Modern CPU heat pipes are typically made from copper and use water as the working fluid.

Applications

Grover and his colleagues were working on cooling systems for nuclear power cells for space craft, where extreme thermal conditions are found. Heat pipes have since been used extensively in spacecraft as a means for managing internal temperature conditions.

Heat pipes are extensively used in many modern computer systems, where increased power requirements and subsequent increases in heat emission have resulted in greater demands on cooling systems. Heat pipes are typically used to move heat away from components such as CPUs and GPUs to heat sinks where thermal energy may be dissipated into the environment.

Solar Thermal

Heat pipes are also being widely used in solar thermal water heating applications in combination with evacuated tube solar collector arrays. In these applications, distilled water is commonly used as the heat transfer fluid inside a sealed length of copper tubing that is located within an evacuated glass tube and oriented towards the sun.

In solar thermal water heating applications, an evacuated tube collector can deliver up to 40% more efficiency compared to more traditional "flat plate" solar water heaters. Evacuated tube collectors eliminate the need for anti-freeze additives to be added as the vacuum helps prevent heat loss. These types of solar thermal water heaters are frost protected down to more than -3 °C and are being used in Antarctica to heat water.

Pipelines over permafrost

Heat pipes are used to dissipate heat on the Trans-Alaska Pipeline System. Without them residual ground heat remaining in the oil, as well as that produced by friction and turbulence in the moving oil would conduct down the pipe's support legs. This would likely melt the permafrost on which the supports are anchored. This would cause the pipeline to sink and possibly sustain damage. To prevent this each vertical support member has been mounted with 4 vertical heat pipes.

Cooking

Heat pipes have been designed to speed the cooking of roasts. The pipe is poked through the roast. One end of the pipe extends into the oven where it draws heat to the middle of the roast.

Ventilation heat recovery

In heating, ventilation and air-conditioning systems, HVAC, heat pipes are positioned within the supply and exhaust air streams of an air handling system, or in the exhaust gases of an industrial process, in order to recover the heat energy.

The device consists of a battery of multi-row finned heat pipe tubes located within both the supply and exhaust air streams. Within the exhaust air side of the heat pipe, the refrigerant evaporates, taking its heat from the extract air. The refrigerant vapour moves towards the cooler end of the tube, within the supply air side of the device, where it condenses and gives up its heat. The condensed refrigerant returns by a combination of gravity and capillary action in the wick. Thus heat is transferred from the exhaust air stream through the tube wall to the refrigerant, and then from the refrigerant through the tube wall to the supply air stream.

Because of the characteristics of the device, better efficiencies are obtained when the unit is positioned upright with the supply air side mounted over the exhaust air side, this allows the liquid refrigerant to flow quickly under gravity back to the evaporator. Generally, gross heat transfer efficiencies of up to 75% are claimed by manufacturers.

Limitations

Heat pipes must be tuned to particular cooling conditions. The choice of pipe material, size and coolant all have an effect on the optimal temperatures in which heat pipes work.

When heated above a certain temperature, all of the working fluid in the heat pipe will vaporize and the condensation process will cease to occur; in such conditions, the heat pipe's thermal conductivity is effectively reduced to the heat conduction properties of its solid metal casing alone. As most heat pipes are constructed of copper (a metal with high heat conductivity), an overheated heatpipe will generally continue to conduct heat at around 1/80 of the original conductivity.

In addition, below a certain temperature, the working fluid will not undergo phase change, and the thermal conductivity will be reduced to that of the solid metal casing. One of the key criteria for the selection of a working fluid is the desired operational temperature range of the application. The lower temperature limit typically occurs a few degrees above the freezing point of the working fluid.

Most manufacturers cannot make a traditional heat pipe smaller than 3mm in diameter due to material limitations (though 1.6mm thin sheets can be fabricated). Experiments

have been conducted with micro heat pipes, which use piping with sharp edges, such as triangular or rhombus-like tubing. In these cases, the sharp edges transfer the fluid through capillary action, and no wick is necessary.

WWT

Chapter 8

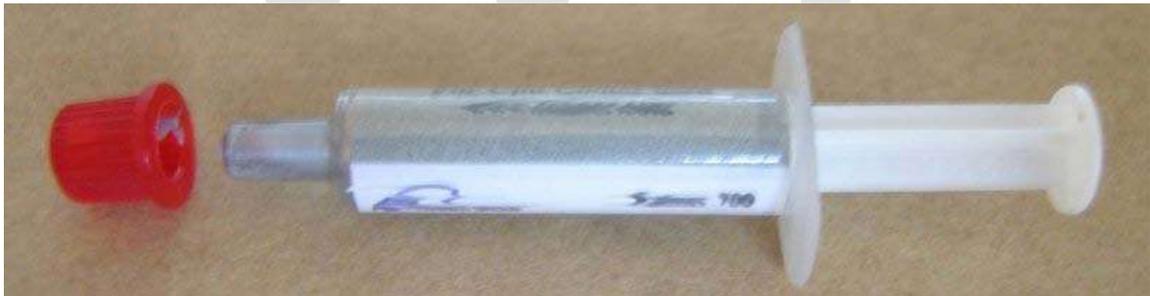
Thermal Grease



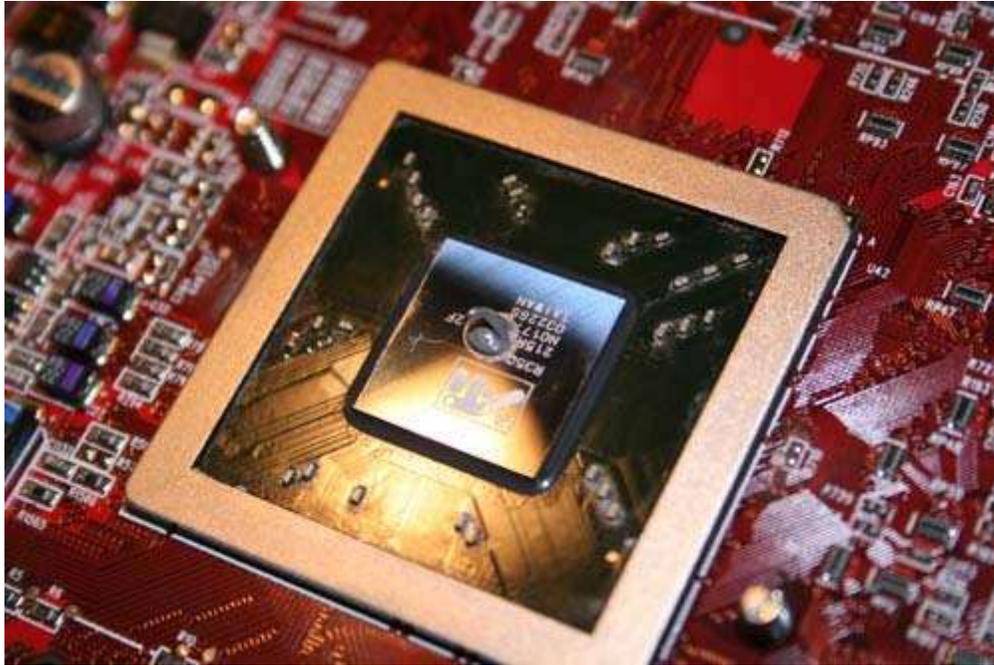
From left to right: Arctic Cooling MX-2 and MX-3, Tuniq TX-3, Cool Laboratory Liquid Metal Pro(Liquid Metal based), Shin-Etsu MicroSi G751, Arctic Silver 5, Powdered Diamond. In background Arctic Silver grease remover



Silicone thermal compound



Metal (silver) thermal compound



Metal thermal grease applied to a chip



Surface imperfections

Thermal grease (also called **thermal gel**, **thermal compound**, **thermal paste**, **heat paste**, **heat sink paste**, **heat transfer compound**, **heat transfer paste (HTP)** or **heat sink compound**) is a fluid substance, originally with properties akin to grease, which increases the thermal conductivity of a thermal interface by filling microscopic air-gaps

present due to the imperfectly flat and smooth surfaces of the components; the compound has far greater thermal conductivity than air (but far less than metal). In electronics, it is often used to aid a component's thermal dissipation via a heat sink.

Thermal conductor types

Thermal greases use one or more different thermally conductive substances:

- Ceramic-based thermal grease has generally good thermal conductivity and is usually composed of a ceramic powder suspended in a liquid or gelatinous silicone compound, which may be described as 'silicone paste' or 'silicone thermal compound'. The most commonly used ceramics and their thermal conductivities (in units of $W/(m \cdot K)$) are: beryllium oxide (218), aluminum nitride (170), aluminum oxide (39), zinc oxide (21), and silicon dioxide (1). Thermal grease is usually white in colour since these ceramics are all white in powder form.
- Metal-based thermal grease contain solid metal particles (usually silver or aluminum). It has a better thermal conductivity (and is more expensive) than ceramic-based grease.
- Carbon based. There are products based on with carbon-based conductors, using diamond powder, or short carbon fibers, they have the best thermal conductivity and are generally more expensive than metal-based thermal grease.
- Liquid metal based. Some thermal pastes are made of liquid metal alloys of gallium. These are rare and expensive.

All but the last classification of compound usually use silicone grease as a medium, a heat conductor in itself, though some manufacturers prefer use of fractions of mineral oil.

All these compounds conduct heat far better than air, but far worse than metal. They are intended to fill gaps that would otherwise hold air, not to create a layer between component and heatsink—this will decrease the effectiveness of the heatsink. Ideally perfectly smooth and flat metallic surfaces would not need heatsink compound.

Purpose

Thermal grease is primarily used in the electronics and computer industries to assist a heat sink to draw heat away from a semiconductor component such as an integrated circuit or transistor.

Thermally conductive paste improves the efficiency of a heatsink by filling air gaps that occur when the imperfectly flat and smooth surface of a heat generating component is pressed against the similar surface of a heatsink, air being approximately 8000 times less efficient at conducting heat than, for example, aluminum, a good heatsink material. Surface imperfections and departure from perfect flatness inherently arise from

limitations in manufacturing technology and range in size from visible and tactile flaws such as machining marks or casting irregularities to sub-microscopic ones not visible to the naked eye.

Thermal conductivity and "conformability" (i.e., the ability of the material to conform to irregular surfaces) are the important characteristics of thermal grease.

Both high-power handling transistors, such as those in an audio amplifier, and high-speed integrated circuits, such as the central processing unit (CPU) of a personal computer, generate sufficient heat to benefit from the use of thermal grease to improve the effectiveness of a heatsink.

The need for heatsink compound can be minimised or removed by lapping the surfaces of the hot component and the matching heatsink face so that they are virtually perfectly flat and mirror-smooth. Computer overclockers, who increase computer speed by measures which increase heat production, resort to lapping and other extreme cooling methods such as water-cooling.

Properties

The metal oxide and nitride particles suspended in silicone thermal compounds have thermal conductivities of up to 220 W/(m·K). (In comparison, the thermal conductivity of metals used particle additions, copper is 380 W/(m·K), silver 429 and aluminum 237.) The typical thermal conductivities of the silicone compounds are 0.7 to 3 W/(m·K). Silver thermal compounds may have a conductivity of 3 to 8 W/(m·K) or more.

In compounds containing suspended particles, the properties of the fluid may well be the most important. As seen by the thermal conductivity measures above, the conductivity is closer to that of the fluid components rather than the ceramic or metal components. Other properties of fluid components that are important for thermal grease might be:

1. How well it fills the gaps and conforms to both the component's and the heat sink's uneven surfaces.
2. How well it adheres to those surfaces
3. How well it maintains its consistency over the required temperature range
4. How well it resists drying out or flaking over time
5. Whether it degrades with oxidation or breaks down over time

The compound must have a suitable consistency to apply easily and remove all excess to leave only the minimum needed.

Application and removal

Computer processor heatsinks utilize a variety of designs to promote better thermal transfer between components. Some thermal greases have a durability up to at least 8

years. Flat and smooth surfaces may use a small line method to apply material, and exposed heat-pipe surfaces will be best prepared with multiple lines.

Excess grease separating the metal surfaces more than the minimum necessary to exclude air gaps will only degrade conductivity, increasing the risk of overheating. Silver-based thermal grease can also be either slightly electrically conductive or capacitive; if some flows onto the circuits it can cause malfunctioning and damage.

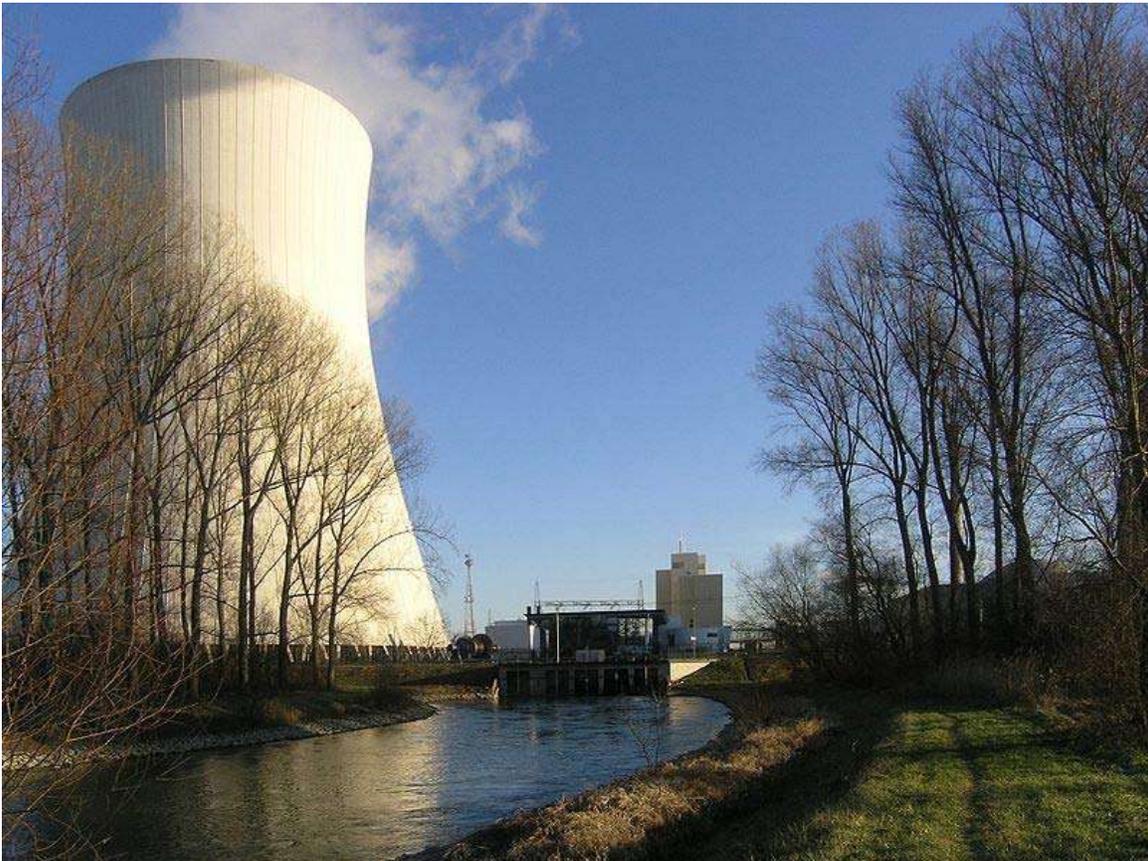
Over time, some thermal greases may dry out, have reduced heat transferring capabilities, or set like glue and make it difficult to remove the heat sink. If too much force is applied the processor may be damaged. Heating the grease by turning the processor on for a short period often softens the adhesion. It is recommended that thermal grease be re-applied with each removal of the heatsink.

Silicone oil-based thermal grease can be removed from a component or heatsink with an alcohol (such as rubbing alcohol) or acetone. Special-purpose cleaners are made for removing heatsink grease and cleaning the surfaces.



Chapter 9

Water Cooling



Cooling tower and water discharge of a nuclear power plant

Water cooling is a method of heat removal from components and industrial equipment. As opposed to air cooling, water is used as the heat conductor. Water cooling is commonly used for cooling automobile internal combustion engines and large industrial facilities such as steam electric power plants, hydroelectric generators, petroleum

refineries and chemical plants. Other uses include cooling the barrels of machine guns, cooling of lubricant oil in pumps; for cooling purposes in heat exchangers; cooling products from tanks or columns, and recently, cooling of various major components inside high-end personal computers. The main mechanism for water cooling is convective heat transfer.

Advantages

The advantages of using water cooling over air cooling include water's higher specific heat capacity, density, and thermal conductivity. This allows water to transmit heat over greater distances with much less volumetric flow and reduced temperature difference.

For cooling CPU cores its primary advantage is that its tremendously increased ability to transport heat away from source to a secondary cooling surface allows for large, more optimally designed radiators rather than small, inefficient fins mounted directly on the heat source.

The water jacket around an engine is also very effective at deadening mechanical noises, which makes the engine quieter. However, the primary disadvantage is that it costs significantly more than an air cooled engine system.

Open method

An open water cooling system makes use of evaporative cooling, lowering the temperature of the remaining (unevaporated) water. This method was common in early internal combustion engines.

Automotive usage

Pressurization

Modern automotive cooling systems are slightly pressurized. This raises the boiling-point of the coolant and reduces evaporation.

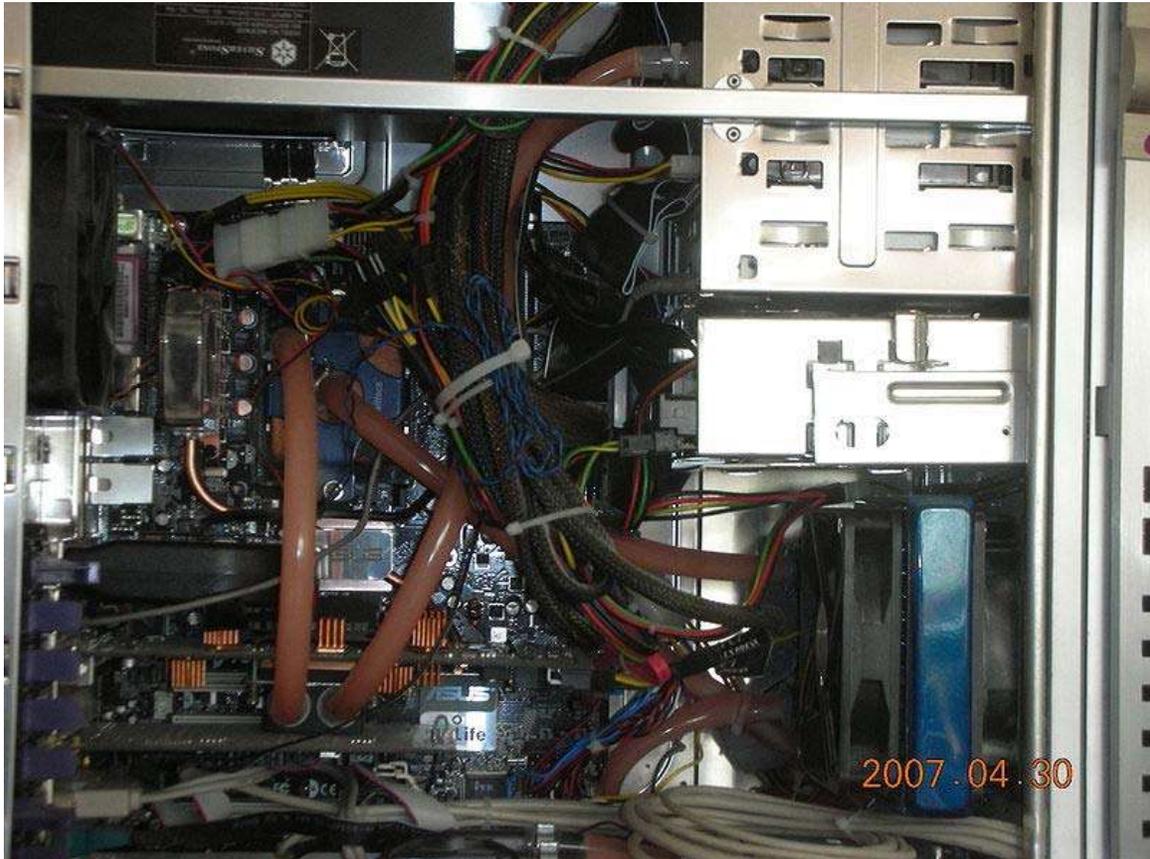
Antifreeze

The use of water cooling carries the risk of damage from freezing. Automotive and many other engine cooling applications require the use of a water and antifreeze mixture to lower the freezing point to a temperature unlikely to be experienced. Antifreeze also inhibits corrosion from dissimilar metals and can increase the boiling point, allowing a wider range of water cooling temperatures. Its distinctive odor also alerts operators to cooling system leaks and problems that would go unnoticed in a water-only cooling system. The heated water can also be used to warm the air conditioning system inside the car, if so desired.

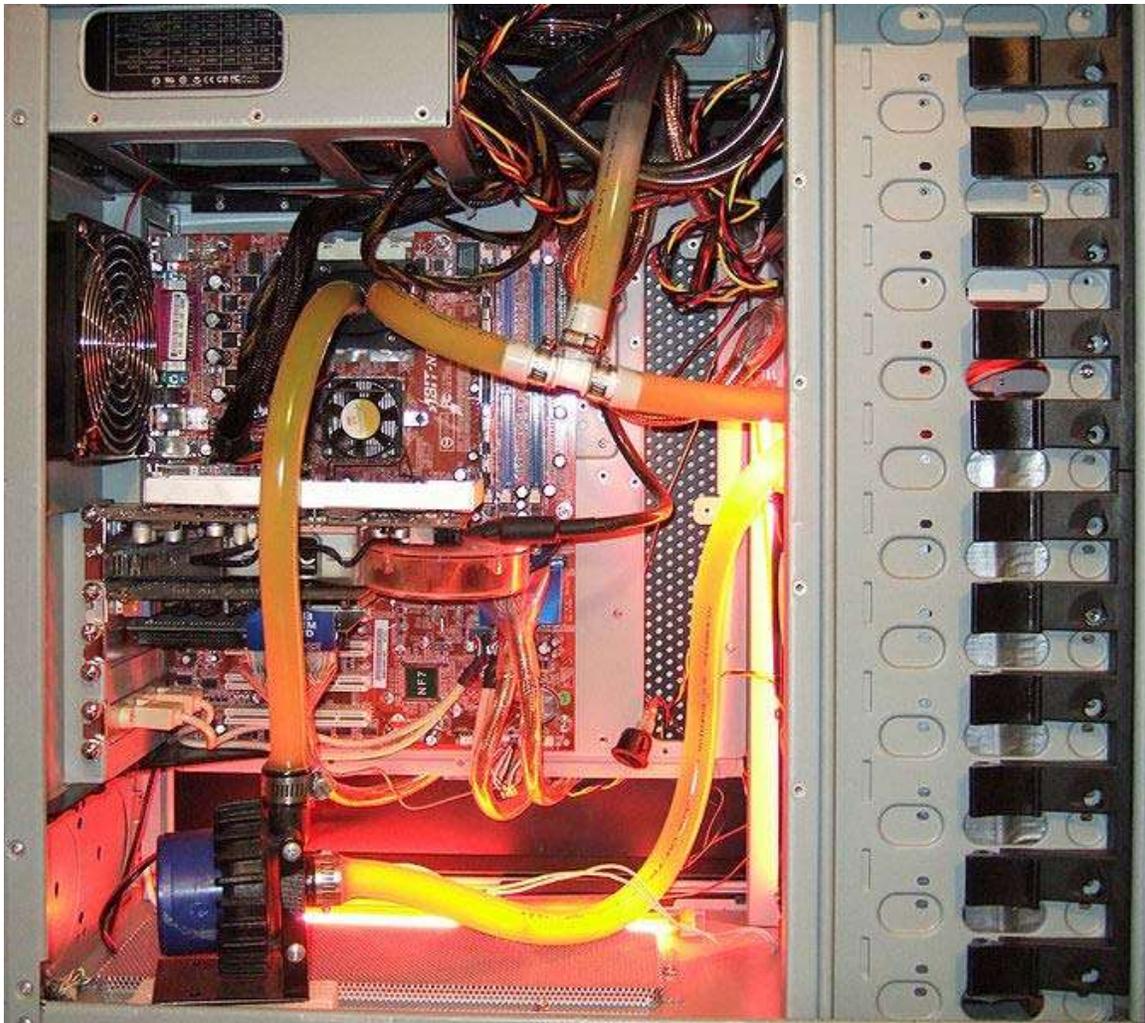
Other additives

Other less common chemical additives are products to reduce surface tension. These additives are meant to increase the efficiency of automotive cooling systems. Such products are used to enhance the cooling of underperforming or undersized cooling systems or in racing where the weight of a larger cooling system could be a disadvantage.

Computer usage



Interior of a water cooled computer, showing CPU/GPU waterblock, tubing and pump.



DIY Watercooling setup showing Laing Thermotech D4 12v pump, Swiftech STORM CPU Waterblock and the typical application of a T-Line.

Cooling hot computer components with various fluids has been in use since at least as far back as the development of Cray-2 in 1982, using Fluorinert. Through the 1990s, water cooling for home PCs slowly gained recognition amongst enthusiasts, but it started to become noticeably more prevalent after the introduction of AMD's hot-running Athlon processor in mid 2000. As of 2010, there are several manufacturers of water cooling components and kits, and some custom computer retailers include various setups of water cooling for their high performance systems.

Water cooling can be used to cool many computer components, but especially the CPU. Water cooling usually uses a CPU water block, a water pump, and a heat exchanger (usually a radiator with a fan attached). Water cooling can allow quieter (potentially fanless) operation, or improved processor speeds (overclocking), or a balance of both. Less commonly, GPUs, Northbridges, hard disk drives, memory, Voltage regulator modules (VRMs), and even power supplies can be water-cooled.

Water coolers for desktop computers were, until the end of the '90s, homemade. They were made from car radiators (or more commonly, a car's heater core), aquarium pumps and home-made water blocks, laboratory-grade PVC and silicone tubing and various reservoirs (home made using plastic bottles, or constructed using cylindrical acrylic or sheets of acrylic, usually clear) and or a T-Line. More recently a growing number of companies are manufacturing water-cooling components compact enough to fit inside a computer case. This, and the trend to CPUs of higher power dissipation, has greatly increased the popularity of water cooling, although only a very small minority of computers are water-cooled.

Dedicated overclockers occasionally use vapor-compression refrigeration or thermoelectric coolers in place of more common standard heat exchangers. Water cooling systems in which water is cooled directly by the evaporator coil of a phase change system are able to chill the circulating coolant below the ambient air temperature (impossible with a standard heat exchanger) and, as a result, generally provide superior cooling of the computer's heat-generating components. The downside of phase-change or thermoelectric cooling is that it uses much more electricity, and antifreeze must be added due to the low temperature. Additionally, insulation, usually in the form of lagging around water pipes and neoprene pads around the components to be cooled, must be used in order to prevent damage caused by condensation of water vapour from the air on chilled surfaces. Common places from which to borrow the required phase change systems are a household dehumidifier or air conditioner.

An alternative cooling system, which enables components to be cooled below the ambient temperature, but which obviates the requirement for antifreeze and lagged pipes, is to place a thermoelectric device (commonly referred to as a 'Peltier junction' or 'pelt' after Jean Peltier, who documented the effect) between the heat-generating component and the water block. Because the only sub-ambient temperature zone now is at the interface with the heat-generating component itself, insulation is required only in that localized area. The disadvantage to such a system is that pelts typically dissipate a lot of power, and the water cooling system must remove this heat in addition to that generated by the component.

Another possible danger is damaging condensation, due to the ambient air right around the pelt being cold. A proper installation requires that the Peltier be "potted" with silicone epoxy. The epoxy is applied around the edges of the device, preventing air from entering or leaving the interior.

Apple's Power Mac G5 was the first mainstream desktop computer to have water cooling as standard. Dell followed suit by shipping their XPS computers with liquid cooling, using thermoelectric cooling to help cool the liquid. Currently, Dell's only computers to offer liquid cooling are their Alienware desktops.

Industrial usage



A Marley mechanical induced draft cooling tower.

Industrial cooling towers may use river water, coastal water (seawater) or well water as their source of fresh cooling water. The large mechanical induced-draft or forced-draft cooling towers in industrial plants continuously circulate cooling water through heat exchangers and other equipment where the water absorbs heat. That heat is then rejected to the atmosphere by the partial evaporation of the water in cooling towers where upflowing air is contacted with the circulating downflow of water. The loss of evaporated water into the air exhausted to the atmosphere is replaced by "make-up" fresh river water or fresh cooling water. Since the evaporation of pure water is replaced by make-up water containing carbonates and other dissolved salts, a portion of the circulating water is also continuously discarded as "blowdown" water to prevent the excessive build-up of salts in the circulating water.

High grade industrial water (produced by reverse osmosis) and potable water is sometimes used in industrial plants requiring high-purity cooling water.

Some nuclear reactors use heavy water as cooling. Heavy water is employed in nuclear reactors because it is a weaker neutron absorber. This allows for the use of less enriched fuel. For the main cooling system, normal water is preferably employed through the use of a heat exchanger as heavy water is much more expensive. Reactors that use other materials for moderation (graphite) may also use normal water for cooling.

Environmental impacts



Cooling water intake of a nuclear power plant.

On very large rivers, but more often at coastal and estuarine sites, "direct cooled" systems are often used instead. These industrial plants do not use cooling towers and the atmosphere as a heat sink but put the waste heat to the river or coastal water instead. These **once-through cooling** (OTC) systems thus rely upon a good supply of river water or seawater for their cooling needs. The warmed water is returned directly to the aquatic environment, often at temperatures significantly above the ambient receiving water. Thermal pollution of rivers, estuaries and coastal waters is an issue which needs to be addressed when considering the siting of such plants.

Other impacts include "impingement" (the capture of larger organisms such as fish and shrimp on screens protecting the small bore tubes of the heat exchangers from blockage) and "entrainment" (the combined effects of temperature, pressure, biocide residual and turbulence/shear on smaller organisms entrained with the cooling water and then expelled back to the aquatic environment in the effluent). The cooling water in such heat exchange cycles is often treated with a biocide to prevent fouling in heat exchangers like condensers and other equipment, but in some instances such control can be exercised instead through frequent cleaning, antifouling paints (both toxic-release and non-toxic), or heat treatment.

Industrial cooling water regulations

The U.S. Clean Water Act requires the Environmental Protection Agency (EPA) to issue regulations on industrial cooling water intake structures. EPA issued final regulations for

new facilities in 2001 (amended 2003). Other EPA regulations for existing facilities were challenged in litigation and EPA issued new proposed regulations in March 2011.

WWT

Chapter 10

Loop Heat Pipe and Air Cooling

Loop heat pipe

A **loop heat pipe** (LHP) is a two-phase heat transfer device that uses capillary action to remove heat from a source and passively move it to a condenser or radiator. LHPs are similar to heat pipes but have the advantage of being able to provide reliable operation over long distance and the ability to operate against gravity. Different designs of LHPs ranging from powerful, large size LHPs to miniature LHPs (micro loop heat pipe) have been developed and successfully employed in a wide sphere of applications both ground based as well as space applications.

Construction

The most common coolants used in LHPs are anhydrous ammonia and propylene.

Mechanism

Limitations of heat pipes

Heat pipes are excellent heat transfer devices but their sphere of application is mainly confined to transferring relatively small heat loads over relatively short distances when the evaporator and condenser are at same horizontal level. This limitation on the part of the heat pipes is mainly related to the major pressure losses associated with the liquid flow through the porous structure, present along the entire length of the heat pipe and viscous interaction between the vapor and liquid phases, also called entrainment losses. For the applications involving transfer of large heat loads over long distances, the thermal performance of the heat pipes is badly affected by increase in these losses. For the same reason conventional heat pipes are very sensitive to the change in orientation in gravitational field. For the unfavorable slopes in evaporator-above-condenser configuration, the pressure losses due to the mass forces in gravity field adds to the total pressure losses and further affect the efficiency of the heat transfer process.

As a result of these limitations, different solutions involving structural modifications to the conventional heat pipe have been proposed. Some of these modifications incorporate arterial tubes with considerably low hydraulic resistance for liquid return to the heat source (arterial heat pipes), while others provide spatial separation of the vapor and liquid phases of the working fluid at the transportation section (separated line heat pipes).

Though these new forms of heat pipes are able to transfer significant heat flows and can increase heat transport length, they remain very sensitive to spatial orientation relative to gravity. To extend functional possibilities of two-phase systems towards applications involving otherwise inoperable slopes in gravity, the advantages provided by the spatial separation of the transportation line and the usage of non-capillary arteries are combined in the loop scheme. This scheme allows heat pipes to be created with higher heat transfer characteristics while maintaining normal operation in any directional orientation. The loop scheme forms the basis of the physical concept of Two-Phase Loops (TPLs).

Origins

Loop heat pipes were patented in USSR in 1979 by Valery M. Kiseev, Jury F. Maidanik, Jury F. Gerasimov, all of the former Soviet Union. The patent for LHPs was filed in the USA in 1982 (Patent 4,467,861).

Applications

The first space application occurred aboard a Russian spacecraft in 1989. LHPs are now commonly used in space aboard satellites including; Russian Granat, Obzor spacecraft, Boeing's (Hughes) HS 702 communication satellites, Chinese FY-1C meteorological satellite, NASA's ICESat .

Air cooling



A Cooler Master V8 has a lot of heatpipes.

Air cooling is a method of dissipating heat. It works by making the object to be cooled have a larger surface area or have an increased flow of air over its surface, or both. An example of the former is to add fins to the surface of the object, either by making them integral or by attaching them tightly to the object's surface (to ensure efficient heat transfer). In the case of the latter it is done by using a fan blowing air into or onto the object one wants to cool. In many cases the addition of fins adds to the total surface area making a heatsink that makes for greater efficiency in cooling.

In all cases, the air has to be cooler than the object or surface from which it is expected to remove heat. This is due to the second law of thermodynamics, which states that heat will only move spontaneously from a hot reservoir (the heat sink) to a cold reservoir (the air).

Examples



Air-cooled *boxer* engine on a 1954 BMW motorcycle

Vehicles

Air is mainly used for Air-cooled engine cooling internal combustion engines(ICE),particularly those powering aircraft, because it is a readily available fluid and is often at a suitable temperature to be used efficiently. While many such ICE are called "liquid cooled" the cooling liquid is usually cooled by air passing through a radiator or heat exchanger. Examples of direct air cooling in automobiles are rarer. The most common example is the Flat engine| flat or *Boxer* engine,once used extensively by porsche and still in use on BMW Motorrad| BMW motorcycles.Notable past models include the Volkswagen Beetle and related models, the Citroën 2CV,the Chevrolet Corvair and the Porsche 911 until 1998.

Turbines

Gas turbine engines (e.g. turbojets, turbofans, etc) incorporate turbines, which are exposed to the hot gases exiting the combustion chamber. Where necessary, relatively cold air is bled from the compression system and used to cool the turbine blades and vanes, to prevent them from melting.

Electronics

Its use is widespread in computers and CPU cooling, where the computer processors produce large quantities of heat that, if not dissipated, could damage the CPU and other electronic components. In this case air has the advantage of being a good insulator too. However, in the future, new processors might generate too much heat to be dissipated through direct air cooling and it would follow that such direct cooling for computers and their components will become obsolete. Water cooling is somewhat popular in very high-power situations, such as large servers or heavily overclocked amateur systems.

Industries

A very large number of industrial processes use air as a cooling medium, either directly or indirectly. Air conditioning is a very common process in which the air in a room, or a whole building, is cooled in order to maintain a comfortable environment for its occupants. Often the air has been cooled by chilled water or brine and the heat transferred to that medium is transported outside the building where, often, fan-driven water-to-air heat exchanging is again effected to reject the heat into the atmosphere. A common sight around, for example, power stations are the large waisted concrete towers that emit steam more or less constantly. These are, in part, using air cooling on a grand scale.