Cooling Electronic Components

Some really cool developments

By Reg Miles

regmiles@lycos.co.uk

One thing has characterised the electronics industry since the change from valves to solid-state devices — both components and products have got smaller. The inevitable consequence of this is that components are running much hotter, but have less space for cooling to take place. This article looks at some of the recent developments and innovations in cooling techniques.

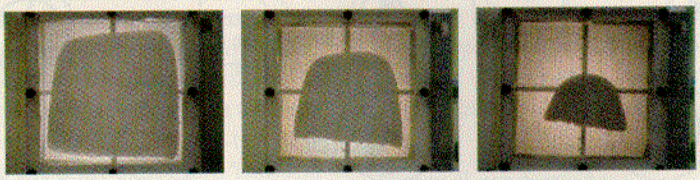

t = 15 min t = 30 min t = 45 min

Figure 1. Visualisation of phase change energy storage (courtesy Cooling Technologies Research Consortium).

The heatsink is no longer the complete solution that it once was, but it is still widely used. The efficiency of any design depends on the surface area of the fins, and the heat transfer coefficient (how effectively heat can actually be removed from its surface). The trouble with using air to do so is that it is a rather good insulator. But it does have obvious advantages — particularly for consumer uses.

Conventionally, heatsink fins are thin and flat. Recently, however, folded fins have been introduced made of corrugated metal sheet bonded, brazed or soldered to the heat spreader base.

This design was originally introduced for military and aerospace use, where the large surface area and light weight were advantageous, and it has spread rapidly. For omnidirectional airflows, the cross-cutting of flat fins into short, peg-like sections gives a c. 20% improvement by comparison with bidirectional types. Augmented fin techniques either add a curvature to the leading and trailing vertical edges of the divided fins (bent fin) or split and bend the fins to give the shape of a tuning fork to create a degree of turbulence; this improves heat transfer by breaking up the slow moving layer of air at the surface of the fins caused by friction between them and the air particles - a region known as the boundary layer. Pin fins (round, elliptical or hexagonal columns) achieve the same effect - and can still be used effectively with natural convection and slow moving air.

Most heatsinks are made of aluminium. Copper is also used, when the superior conduction outweighs the increased weight and cost. A recent alternative is aluminium silicon carbide (AlSiC); which is light, strong and conductive, although expensive at present.

Natural graphite is another newcomer: this has the conductivity of copper, at less than a quarter of its weight. Pyrolitic graphite and graphite fibre-based materials have still better conductive properties; but to achieve them requires temperatures of over 3000 °C, which is costly and limits them to aerospace applications at present. Graphite materials, incidentally, exhibit strong anisotropic properties, with different conductivity along axes in different directions — which limits the range of applications.

Heatsink compounds

In order to maximise the conduction from component(s) to the base of the heatsink it is usual to apply some form of thermally conducting compound that fills in any unevenness in the materials, so that there is no insulating air trapped between them. The most common is a silicone-based thermal grease; but there are a variety of other compounds available. One recent development is the use of a phase change material (PCM). This is solid up to a predetermined temperature, and then liquefies as the temperature exceeds that — spreading out and filling the interface.

PCM is also used to prevent components from exceeding a particular operating temperature, or to cool components that operate transiently. Its expansion as it changes phase and liquefies absorbs the excess heat (see **Figure 1**); and when the temperature drops the PCM will solidify again in readiness for the next temperature rise. Heatsinks filled with PCM are available, and PCM reservoirs that can be attached to PCBs and individual components.

Fans and jets

When components generate a lot of heat, then a fan is probably the answer. Or, a more recent development, the jet actuator. A device that consists of a plenum chamber and a small, electromagnetically driven diaphragm that sucks air in and blows it out again. This pulsed delivery helps to overcome the boundary layer problems described before, because the layer will thin between blasts. A small version of this, known as a Micro-Jet Array, has recently been developed which can be placed immediately below the component to blow air over it at a velocity of around 70 km/h.

Thermosiphon et al.

An even more effective method of avoiding hot spots is to conduct the heat away from the component to where it can be cooled near the outlet (or even outside the enclosure, if that is practicable); and this can be achieved by using a thermosiphon or a heat pipe — without any power being consumed. A thermosiphon is an evacuated sealed tube, typically made of copper, containing a small amount of fluid —· usually water. Being in a vacuum the fluid readily vaporises when heat is applied, enabling it to absorb a large amount of heat, the vapour rises to the other end of the pipe where the heat is conducted away and the vapour condenses on the inner surface of the tube and runs back down for the cycle to be repeated. It only requires a slightly lower temperature for the vapour to condense, so it is very efficient.

A heat pipe adds a porous wick lining on the inner surface of the tube that recirculates the liquid by capillary action to the hot end. The finer the pore structure of the wick the more the capillary action can overcome gravity: grooved and fine screen wicks will only cope with the evaporator slightly above the condenser, whereas sintered metal powder wicks can cope with any orientation.

Thermosiphons and heat pipes are becoming increasingly attractive because of their efficiency and their passive operation. They are also small: typically 3-4mm in diameter, with thinner, micro, versions now coming into use; and in lengths to suit the applications. A mobile phone, for example, would require only a short pipe; and would probably use the antenna as the condenser.

In a notebook computer it would be somewhat longer; and use either an aluminium plate under the keyboard or a heatsink as the condenser. Multiple heat pipes can also be used to carry heat from several components to the one heatsink.

A recent variation on the thermosiphon and heat pipe is to make them into a loop. The vapour travels from the evaporator to the condenser and the condensed liquid travels through a continuation of the pipe back to the evaporator. The advantages of the loop arrangement are that the vapour and liquid lines can be flexible, and can be a lot longer than a conventional pipe — more than a metre. A further variation is the pulsating heat pipe, with an internal serpentine channel in which expansion and contraction due to vaporisation and condensation set up a pulsating motion that pushes the vapour to the cold end and the liquid back to the hot end.

Temperature vs. size

The problem of cooling electronics components with increasingly high operating temperatures is being exacerbated in many cases by their smaller size. This is particularly true in applications such as telecommunications, where the heatsinks needed to cool devices like RF amplifiers are larger than what they cool. The result is a hot spot directly over the device due to 'thermal spreading resistance' — only that part of the heatsink is doing its job. The fan size and speed can be increased to deal with it, or aluminium can be replaced by copper or some more exotic material (chemical vapour deposited (CVD) diamond is being used for heat spreaders where the temperature is critical). But these alternatives have their disadvantages in noise, weight and cost.

An attractive solution is to embed heat pipes in the base of an aluminium heatsink. This overcomes much of the resistance, and spreads the heat fairly evenly. An even more effective solution is to use a vapour chamber, which is a vacuum vessel with a wick lining that works on exactly the same principle as the heat pipe: wherever heat is applied the fluid in the wick at that point is vaporised; and wherever the vapour comes into contact with a cooler part its latent heat is released and it condenses back into the wick. A DARPA (US Defence Advanced Research Projects Agency) project, led by Florida International University, is developing a similar heat spreader; but containing a piezo-driven micro pump to transfer the fluid to a heat exchanger, eventually to be integrated into a single module.

When there is no practical or aesthetic objections, a cold plate (or water block) makes a very effective alternative to air cooling, heat pipes and vapour chambers. This is moving into the realms of computer overclocking, where the motivation is not to prolong the reliable life of the component, the CPU, but to push it to its limits. Although, of course, it is a perfectly valid means of prolonging the reliable life of any components that do get very hot.

The cold plate may be just a container for the liquid (probably water) or have a liquid carrying pipe embedded in it. Depending on its design, electronic components may be attached to one or both sides. The system, at its simplest, would consist of cold water going into the cold plate through a tube from a supply and out through another tube. A more likely arrangement is the recirculating one, comprising a cold plate, a reservoir or expansion tank, a pump, and a heat exchanger (probably a tubed radiator — a small version of the type used in cars) with a fan to provide air cooling (a liquid to air system). Warm liquid from the cold plate goes to the reservoir, then to the pump, on to the radiator where it is cooled, and back again to the cold plate. If the cold plate has pipes embedded in it these can be fed in series or in parallel: in the former case starting at one end and snaking along the plate from side to side; in the latter case having a number of pipes going straight across from a main supply pipe on one side to a main return pipe on the other. Kits have already become available: Maplin has one (www.maplin.co.uk); and Koolance has a range of kits, and water cooled PCs (www.koolance.com). Figure 2 shows their CPU-200 Cooler installed.

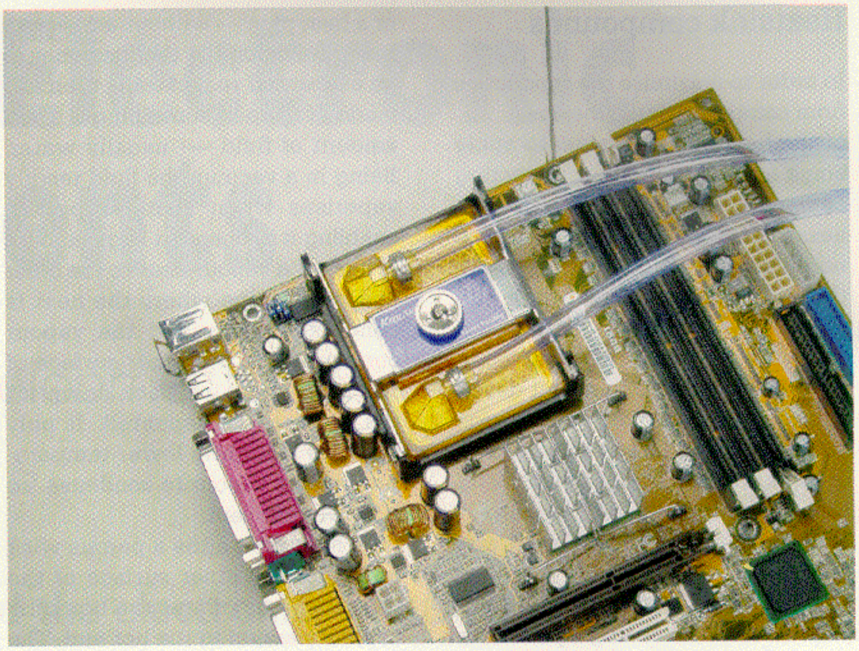

Figure 2. Koolance CPU cooler installed (courtesy Koolance).

Liquid cooling — the future?

Liquid cooling is very suitable for combining with a thermoelectric cooler (TEC) — also known as a Peltier element. This consists of two semiconductor elements, primarily bismuth telluride, heavily doped to create n-type and p-type couples (see **Figure 3**). At the cold junction, heat from the component is absorbed by electrons as an electric current raises them from a low energy level in the p-type element to a higher energy level in the n-type element; at the hot junction, electrons move from n-type back to low energy p-type and the heat is expelled. The flow of electrons is maintained by the DC current, which goes up one couple and down the other, but the carrier current, and heat, go down both. Normally, the n-type and p-type couples are combined in various numbers to produce a module, where they are connected electrically in series and thermally in parallel. A module can be small, but still produce a typical temperature difference of 70 degrees Celsius. For still greater cooling the modules can be cascaded.

Superlattice block

A new development from the Research Triangle Institute promises to give the TEC a considerable boost. It uses stacks of thin films of two alternating semiconducting materials to control the transport of phonons and electrons in the superlattices: the p-type Bi2Te3/Sb2Te3 superlattices block the former (to prevent heat being conducted back) and transmit the latter. It is claimed to be 2.4 times more efficient than a conventional TEC; and responds 23,000 times faster, because of its thinness.

Apparently, dots of the material applied to just the hot spots on an electronic component would be more effective than cooling the whole device — and save on power. Meanwhile DARPA is funding research into integrated thermo-electric-fluidic refrigeration plates, to replace the more expensive and less efficient separate TEC and coolant — a project led by CFD Research Corporation.

A variation on the TEC is the thermionic cooling device. This employs two materials separated by either a barrier layer of around one micrometre thick or a vacuum gap of a few nanometres: as the electrons absorb heat and become more energetic they tunnel from the cold side (emitter) to the hot side (collector), aided by a voltage bias — and the barrier layer or, especially, the vacuum gap prevent phonons from returning. The low voltage means that cooling can be achieved without unwanted heating. The main problem with the development of the type using a vacuum gap is that of obtaining a gap of consistent size across an area measured in square centimetres — in the laboratory, and then in production.

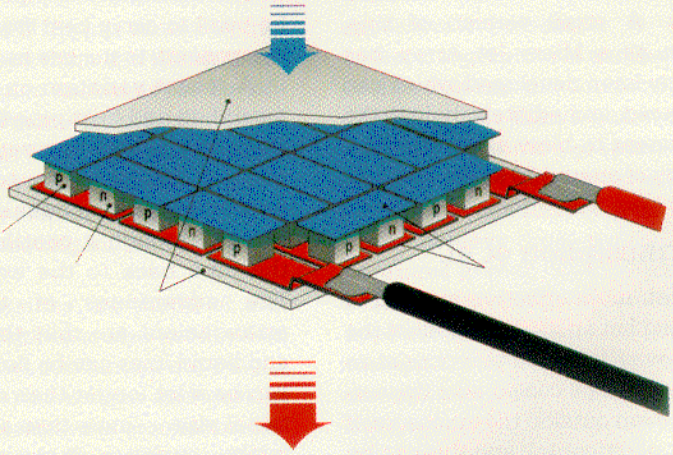

Figure 3. Basic structure of a thermoelectric cooler (TEC) (courtesy Melcor).

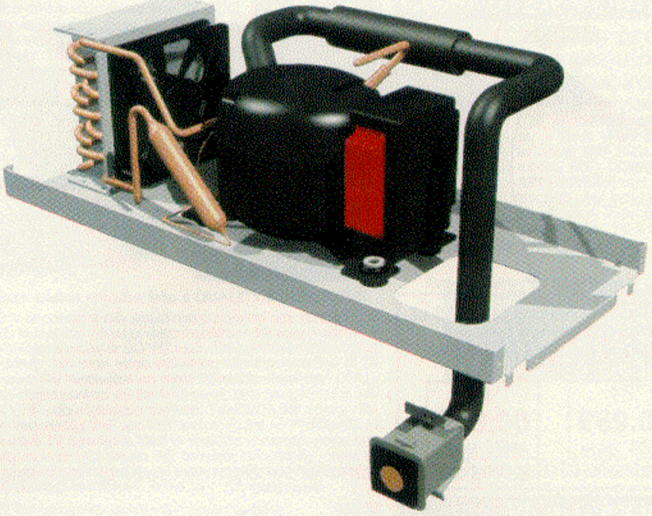

Figure 4. Vapour compression cooling system for home constructors (courtesy Vapochill/Asetek)

The fridge approach

Returning to liquid cooling, such a system can make use of conventional refrigeration to cool the liquid. In this chilled liquid (or liquid to liquid) system the refrigerator will replace the heat exchanger. Again, water is likely to be the circulating fluid. If it is required to cool at subzero temperatures it can be mixed with antifreeze — which does reduce the excellent heat transfer properties of water and increases the viscosity, but not significantly so. As an alternative, the system can be a refrigerator in its own right, which is normally referred to as vapour compression cooling. VapoChill (www.vapochill.com) has a range of workstations, power PCs, and a kit for home constructors (see **Figure 4**). To save space a capillary tube is used to reduce the pressure prior to the evaporator mounted on the CPU, otherwise it is the same as a normal refrigerator. Such systems are intended to be used at low temperatures, in the range from 20 °C down to –40 °C. Because it is the refrigerant itself that is being circulated, the flow rate can be less than with chilled liquid cooling. A Stanford University led project for DARPA is investigating the use of a micro cooler, using an electrokinetic pump and a micromachined evaporator incorporating a temperature sensor; with the long term goal of developing refrigeration devices based on novel compressors for the vapour phase using electrokinetic technology.

Operating CMOS devices at low temperatures gives a significant increase in performance. This is mainly because of an exponential reduction in leakage currents and an increase in transistor switching speed. The latter results from an improvement in mobility, which increases as the temperature falls because thermal vibrations in the semiconductor crystal lattice are reduced and, with it, electron-phonon scattering that slows the carrier velocity (transistor switching speed being proportional to the mean carrier velocity in the device). However, devices can fail due to mismatches in the thermal expansion of their materials, or electronic failure due to hot carrier effects — they must be suitable for use at those low temperatures.

Another barrier to using chilled cooling is the condensation that will form when warm, moist ambient air meets the cold parts of the system. All parts of the cooling system that are inside the case must therefore be insulated to prevent condensation forming.

One refrigeration technology that looks particularly promising for cooling electronics (and there are several alternatives to the familiar vapour compression system) is the thermoacoustic refrigeration device (TRD), or thermoacoustic cooler (TAC). This uses standing acoustic waves in an enclosed cavity to generate the mechanical compression and expansion of a pressurised gas (normally helium), in conjunction with a thermoacoustic core consisting of a porous stack of plates. Its operation is in the manner of a pump: expanding the gas enables it to absorb heat, while compressing it expels that heat; as the gas particles oscillate back and forth they cool down the stack by absorbing heat as they pass through it in one direction and heat it up as they travel back in the opposite direction — resulting in temperature gradient across the stack. This is exploited by having a cold and a hot heat exchanger at either end through which circulating fluid passes (water, with or without antifreeze).

Because the technique has the potential for high efficiency operation without the need for cooling liquids or mechanical moving parts (except for the acoustic driver), it is suitable for miniaturisation. DARPA has a project under way, led by Rockwell Science Centre, to develop a microelectromechanical system (MEMS) version of it; utilising piezoelectric materials and transducers for high efficiency acoustic generation. And a related project, led by the University of Utah, to integrate it with microelectronic circuits. An additional advantage of the TRD is that the degree of cooling can be controlled by adjusting the amplitude, rather than by switching it on and off like a normal refrigerator.

Watch your calories

In addition to refrigeration systems using liquid or gas to achieve cooling, research continues into systems that employ the magnetocaloric effect. In this a magnetic field is cyclically applied to a paramagnetic solid that makes it expel absorbed heat when the field is on, because all the electron spins are aligned and entropy (and its capacity to hold heat) is reduced; and lets it absorb heat again when it is off, because the electron spins are randomised again — increasing its heat absorbing capacity.

When the refrigerant (which can be water with antifreeze, or any other suitable liquid) is pumped into contact with the paramagnetic material while the magnetic field is off, its heat is absorbed, and the cold liquid goes on to cool the electronic component(s); the warmed liquid then goes back to the material (with the magnetic field on) and carries away the previously absorbed heat, and then on to a fan cooled heat exchanger. From where the cycle is repeated.

Until recently, it required superconducting magnets to work. Then researchers at the Ames Laboratory working in partnership with the Astronautics Corporation developed a refrigerator using a permanent magnet, and a rotating disc containing an alloy of gadolinium, silicon and germanium (GdSiGe) as the paramagnetic material. Nanocomposites have also been proposed to further increase the magnetocaloric effect.

A variation on magnetocaloric refrigeration is electrocaloric refrigeration, which uses electric fields to achieve the same effect. This was originally developed by researchers in Russia.

It has the advantage that electrical fields are easier and cheaper to develop than magnetic fields. Research into paraelectric materials continues, to find the optimum combination with which to achieve realistic temperature changes.

Other developments

Just for completeness, there are a couple of other methods of liquid cooling that are used for specialised purposes. One is spray cooling, where, as the name implies, components and circuit boards are sprayed. A refinement on this has recently been announced by researchers at UCLA (University of California Los Angeles): they have developed micro sprays that cool individual components, using a matrix of nozzles 35 micrometres in diameter. These are made by reactive-ion etching; a process that gives a particularly smooth internal wall to minimise clogging by trapped contaminants. The sprayed water cools by both thermal convection and evaporation. The nozzle matrix was tailored to the distributed heat of the component, which was coated with Parylene-C, a conformal polymer with excellent dielectric properties. DARPA is funding similar research, led by Carnegie Mellon University. In this case the atomised droplets (of dielectric fluids) will be created by swirlers and vibrating piezoelectric transducers, as well as by micro nozzles. It is intended to use integrated software on the chip to control droplet sizes, spray frequencies and spray locations, based on sensing the temperature, thermal gradients and film thickness.

NASA has also been experimenting with spray cooling; and liquid immersion, with the components in a liquid bath. This will enable NASA to make increased use of commercially available components in space: by allowing them to be decoupled from the chassis, and providing radiation shielding, in addition to good thermal management.

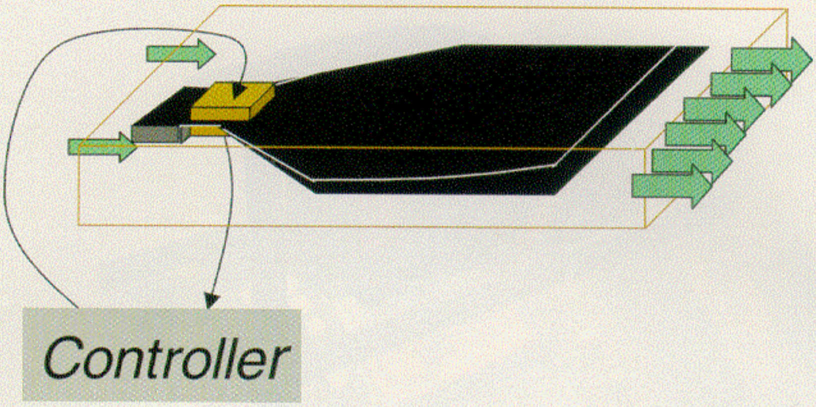

Figure 5. Principle of 'flapping' piezoelectric fans (courtesy Cooling Technologies Research Consortium).

The Cooling Technologies Research Consortium at Purdue University in Indiana was originally founded in 1999 by Suresh Garimella, a professor of mechanical engineering at Purdue, and now has ten consortium members: Aavid Thermalloy, Apple, Delphi-Delco Electronics, Eaton Corporation, General Electric, Intel, Modine Manufacturing, Nokia, Rockwell Automation and Sandia National Laboratories. In addition, there are two Consortium Supporters: Johnson Matthey and Philips Research.

"Industry comes to us with a technical problem, and we conduct research to help solve those problems.", Garimella said.

A number of projects are under way, one of which is investigating micro channel heatsinks. These have micrometre-sized channels that carry a cooling liquid; and have many times the heat transfer coefficient of conventional liquid cooled heatsinks.

The CTRC is investigating the characteristics of heat transfer and fluid mechanics in micro channels, which differ in many respects from conventional designs because of the small scale.

Another project deals with piezoelectric fans — developed by Garimella and Arvind Raman, an assistant professor of mechanical engineering at Purdue. These have piezoceramic patches bonded onto thin, low frequency flexible blades to drive the fan at its resonant frequency — alternating current causes the ceramic to expand and contract, giving a flapping motion (**Figure 5** shows the principle). The technology is attractive because the fans can be small: small enough to fit into a mobile phone or portable computer; and smaller still to fit onto a chip and cool it directly, with blades only 100 micrometres long. They also consume only 2 mW of electricity, compared to about 300 mW for a conventional fan. And, without magnets, there is no electromagnetic noise to interfere with signals. The Consortium is also investigating and improving flat heat pipes and phase change materials.

Conclusion

Cooling used to be almost an afterthought, and something of an art when it was applied. But the latest hot-running components have put an end to that casual approach. It is now essential to consider the needs at an early stage in the design process, and research is turning its application into a science.

(020411-1)

Addendum:

Maplin, in addition to a water cooling kit, has a TEC, and a range of fans, heatsinks and thermal compounds, for the home constructor.

JUNE 2003
£3.45

ELEKTOR ELECTRONICS

THE ELECTRONICS & COMPUTER MAGAZINE

www.elektor-electronics.co.uk

AVR TV TENNIS

back to the seventies

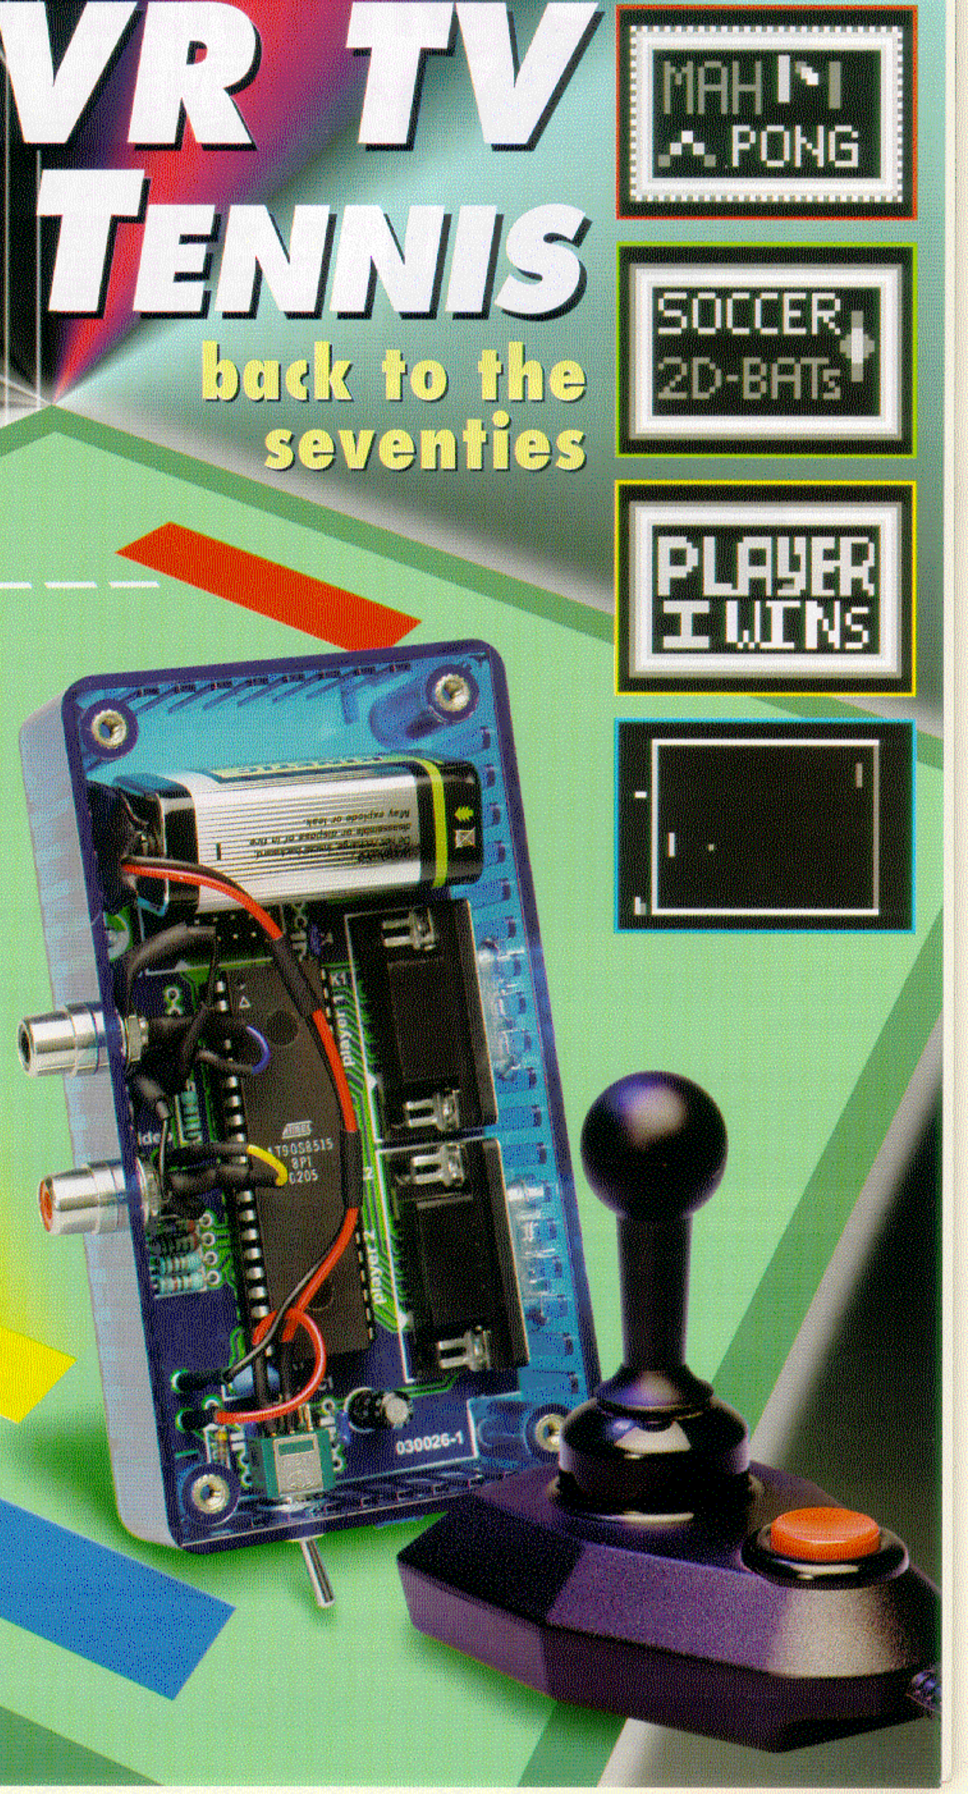

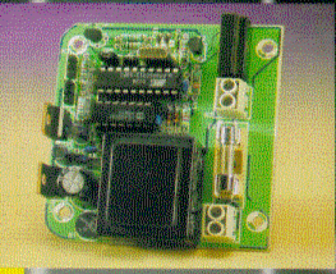

Night Light Control

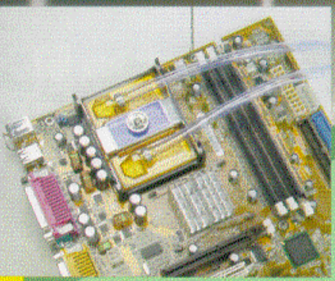

Cooling Electronic Components

Electronic Knotted Handkerchief

Modifying PC AT Power Supplies

Lightning Intensity Detector

C Compilers for the PIC Micro

06
9 770268 451104