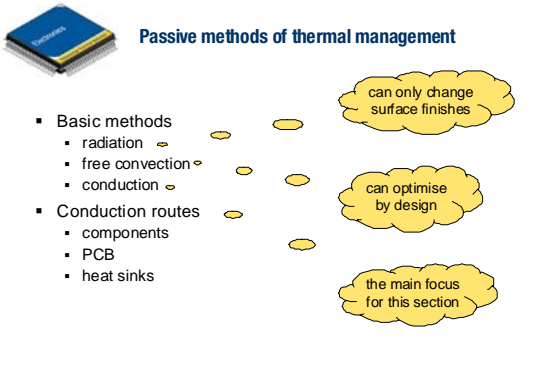
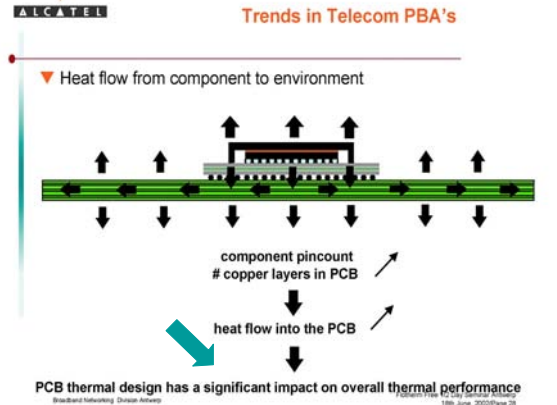
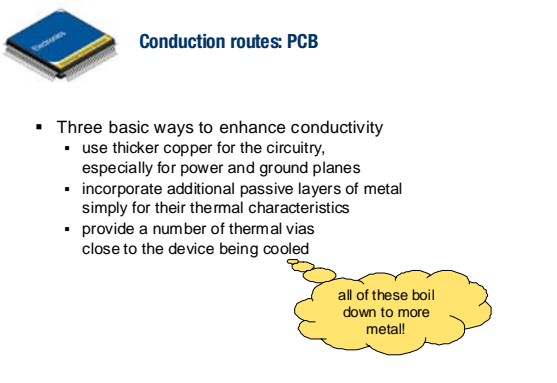
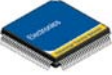

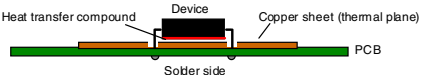

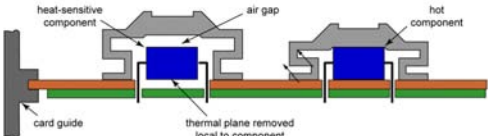


Passive methods of thermal management

 <p>Passive methods of thermal management</p> <ul style="list-style-type: none"> Basic methods <ul style="list-style-type: none"> radiation free convection conduction Conduction routes <ul style="list-style-type: none"> components PCB heat sinks <p>can only change surface finishes</p> <p>can optimise by design</p> <p>the main focus for this section</p> <p>Electronics KTN – Knowledge For Growth</p>	<p>By passive methods of thermal management, we mean ways of removing heat without putting in extra energy or using moving parts, and there are the same three basic methods of transfer that we met in Section 2: radiation, convection and conduction. Now with radiation we can't do much to influence the amount of heat transferred, except possibly change the surface finish to improve its emissivity. We can't just call up an environment with colder surfaces! It's much easier to put foil around something and keep the radiation in!</p> <p>With convection, the only passive method available to us is free or natural convection, and this we can influence, by choosing designs that maximise the airflow.</p> <p>So the main focus for this section on passive methods will inevitably be on conduction, because conduction routes are available through components, first to the board and out through heat sinks to the surroundings, and we can influence all of these routes by appropriate design choices.</p>
 <p>Trends in Telecom PBA's</p> <p>Heat flow from component to environment</p> <p>component pincount # copper layers in PCB</p> <p>heat flow into the PCB</p> <p>PCB thermal design has a significant impact on overall thermal performance</p> <p>Broadband Networking, Design Academy PCB Design File #2 2007 Seminar Agenda 18th June, 2007 Page 28</p>	<p>To remind you of what happens at a component level, this is a view of a face-down BGA. Internally this has a flip-chip construction, and the reverse of the die is stuck to the top cover. So heat is being conducted both upwards to the case and downwards through the interposer and the balls to the PCB. A significant proportion of the heat is being removed via the substrate, and the effectiveness of the transfer will depend on the number of balls, and the copper content of the board. So the thermal design of the board itself will have a significant impact on overall thermal performance.</p>
 <p>Conduction routes: PCB</p> <ul style="list-style-type: none"> Three basic ways to enhance conductivity <ul style="list-style-type: none"> use thicker copper for the circuitry, especially for power and ground planes incorporate additional passive layers of metal simply for their thermal characteristics provide a number of thermal vias close to the device being cooled <p>all of these boil down to more metal!</p> <p>Electronics KTN – Knowledge For Growth</p>	<p>Boards will normally be designed with power and ground plane and signal layers, and the design will depend primarily on the electrical requirements. But there are things that we can do to enhance conductivity:</p> <ul style="list-style-type: none"> we can use thicker copper within the board, especially for power and ground planes which have a high percentage of area coverage, we can include additional passive layers of metal simply for their thermal characteristics, and we can provide thermal vias close to the device being cooled. <p>We will look at these in turn, and see how they might be helpful to thermal management. Note that all of them boil down to having more metal within the board.</p>



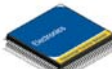
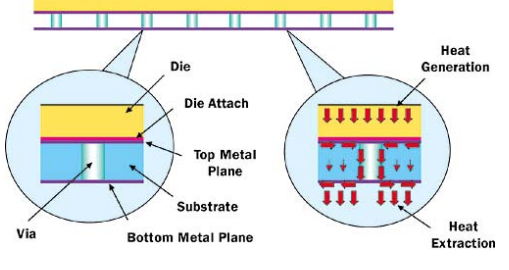

Passive methods of thermal management

 <h3>Using heavier copper/internal thermal planes</h3> <p>more likely to be used for CTE control</p> <ul style="list-style-type: none"> Simplest solution <ul style="list-style-type: none"> heavier copper has no extra process steps internal layers involve some extra processes Four potential problem areas <ul style="list-style-type: none"> can't etch very fine tracks in thick copper stress problems that can lead to board warp delamination potential with thick embedded planes more difficult to solder because of greater thermal mass First two have implications for design Last two will have an impact on the process settings, and may require changes to the laminate material <p>Electronics KTN – Knowledge For Growth</p>	<p>The simplest solution is to use thicker copper for the conductive layers, which needs no extra process steps. Whilst one can include within the structure copper layers that serve no function apart from thermal conduction, adding internal layers involves extra processing, and boards with this type of construction are more likely to be used when the aim is to control the coefficient of thermal expansion by using invar or similar.</p> <p>However, there are four potential problem areas with adding metal:</p> <ul style="list-style-type: none"> if the copper is being used for signal layers, it is not possible to etch fine tracks in thick copper; heavy copper can produce stress problems that lead to board warp, unless care is taken to balance the internal construction; thick embedded planes are a similar challenge, and also have the potential for delamination; (finally) any solution with additional copper becomes more difficult to solder because of the greater thermal mass. <p>Of these, the first two problem areas have implications for design, and the last two have an impact on the process settings, both for fabrication and assembly, and may require changes to the laminate material.</p>
 <h3>External thermal planes</h3> <p>extra metal can be indirect</p> <ul style="list-style-type: none"> Simple solution for less complex boards <ul style="list-style-type: none"> metal layer bonded to component side of board must include holes for component mounting Transfer heat from device to thermally conductive sheet with greater surface area Usually made of copper or aluminium  <p>Electronics KTN – Knowledge For Growth</p>	<p>Rather than putting the heavy copper inside the board, one solution that is particularly appropriate with simpler through-hole boards, and used to be quite common, is to put the additional metal in the form of an external thermal plane. This is bonded to the component side of the board and has holes for component mounting.</p> <p>The aim is to transfer heat from the device to a thermally-conductive sheet with greater surface area, and these thermal planes, often as fingers of metal underneath devices, are typically made of copper or aluminium. In the past nickel has also been used, but it is better to employ a metal that conducts heat as well as possible.</p>
 <h3>Thermal bridge</h3> <ul style="list-style-type: none"> Extruded/stamped metal  <ul style="list-style-type: none"> To protect heat-sensitive components by transferring heat around the device <ul style="list-style-type: none"> thermal plane removed from area below heat-sensitive part To conduct heat from top of device to thermal plane on PCB <p>Electronics KTN – Knowledge For Growth</p>	<p>While we are considering external metal, bonded to the board, we will divert to consider thermal bridges. These metal sections are made by extrusion or stamping and can be used in one of two ways:</p> <ul style="list-style-type: none"> to protect sensitive components by transferring heat around the device, for which the thermal plane is removed from the area below the heat-sensitive part; or alternatively, as shown on the right of this diagram, to conduct heat from the top of the device to a thermal plane on the board.

Passive methods of thermal management

 <h3>Thermal bridge</h3> <ul style="list-style-type: none">More efficient than relying on conduction to air and through component leads alone  <p>Electronics KTN – Knowledge For Growth</p>	<p>A more common form of thermal bridge is shown here, soldered to copper areas on the board, and providing local heat-sinking to take heat out of the board. This style of bridge can also be used, with an appropriate thermal interface material, to remove heat from the top of the component.</p>
 <h3>Alternative materials and constructions</h3> <ul style="list-style-type: none">Most PCBs made of glass-fibre reinforced epoxy (often “FR-4”, a generic description!)<ul style="list-style-type: none">high thermal resistanceanisotropicfillers make only a marginal improvement <p>Electronics KTN – Knowledge For Growth</p>	<p>Persuading heat out of the board would be so much easier if boards were not made of laminates that are relatively poor conductors! The most common material is a glass-fibre reinforced epoxy often referred to as “FR-4”, but this is a generic description for a material that occurs in enormous variety. Given the range available, with different resins and reinforcements ranging from cost-conscious to high-specification/high cost, it’s a bit like saying “car”, when you are including everything from a Mini to a Mercedes!</p> <p>But, whatever the laminate, the base laminate will be made of a poorly-conducting resin, with a poorly-conducting reinforcement. Not only does a laminate not conduct heat very well, but it also conducts to different extents in different directions: in other words, it is anisotropic. As with any resin, it’s possible to add fillers that improve the thermal conductivity, but there are mechanical side effects, and the result is still not a good conductor of heat. Within the board, it’s the copper that conducts heat, which is why we focused on it.</p>
 <h3>Alternative materials and constructions</h3> <ul style="list-style-type: none">Most PCBs made of glass-fibre reinforced epoxy (often “FR-4”, a generic description!)<ul style="list-style-type: none">high thermal resistanceLow thermal resistance materials<ul style="list-style-type: none">include alumina, berylliaAlternative constructions<ul style="list-style-type: none">T-LAM™ — thin dielectric layer between copper foil tracks and a metal base plateIMS (“insulated metal substrate”) — interconnect on thin insulating layer on metal substrateProblems of cost and industry acceptance <p>Electronics KTN – Knowledge For Growth</p>	<p>There are of course alternative ways of providing an insulating substrate on which to mount the interconnection required by the circuit. There are ceramic materials with high thermal conductivity, which include alumina, beryllia and aluminium nitride, but these are brittle, rigid and require specialised metallisation.</p> <p>As well as changing the base material, we can look at different constructions, of which there are many. One is T-LAM™, which has a thin dielectric layer between copper foil tracks and a metal base plate. Then there is the “Insulated Metal Substrate”, where the interconnect is built up on a thin insulating layer on a metal substrate. A different construction, but a similar principle, seeking to combine the electrical isolation given by polymeric substances with the excellent heat conduction of metals. Whilst such constructions have been developed for specialist purposes, there are problems both of cost and industry acceptance, so for most purposes the thermal engineer will have to live with FR-4.</p>

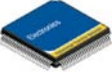
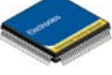
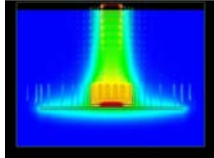
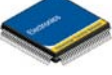
Passive methods of thermal management

 <p>Thermal vias</p> <ul style="list-style-type: none"> Conductivity of each via will depend on <ul style="list-style-type: none"> the dimensions of the via the thickness of the walls whether or not it is filled the nature of any insulation layer on either side of the via  <p><small>Electronics KTN – Knowledge For Growth</small></p> <p><small>http://www.pcblibrary.com</small></p>	<p>We've already seen that you can improve the lateral conductivity of a board by increasing the thickness of copper, and one way which finds favour for increasing the thermal conductivity through the board is the thermal via. From the point of view of construction and materials, there is no difference between a thermal via and one that is intended for electrical connection: the thermal conductivity will depend on its dimensions and on the thickness of the walls – a typical process will plate 25-30µm of copper round the inside of the via hole, though this may be thicker near the surface.</p> <p>The effectiveness of the via as a thermal connection will also depend on whether or not the via is filled, and if so whether with resin or copper, and on the nature of any insulation layer on either side of the via, which would be part of the thermal path.</p>
 <p>Heat flow in a thermal via</p>  <p><small>Electronics KTN – Knowledge For Growth</small></p> <p><small>Bruce Guerin, Electronics Cooling, August 2004</small></p>	<p>As indicated here, when it comes to moving heat, only a little goes through the laminate, but a great deal goes down the coating of the wall. Thermal vias are thus a very good strategy for taking heat out of BGAs, with vias in the pads to which they are connected, and it's even possible to extract heat from the opposite side of the board using a Thermal Interface Material and some kind of extended surface.</p>
 <p>Thermal vias</p> <ul style="list-style-type: none"> Conductivity of via will depend on <ul style="list-style-type: none"> the dimensions of the via the distance between vias the thickness of the walls whether or not it is filled the nature of any insulation layer on either side of the via Multiple small vias can be more effective <ul style="list-style-type: none"> choice will depend on interconnect density and board fabrication constraints <p><small>Electronics KTN – Knowledge For Growth</small></p>	<p>Typically, because what you are trying to do is to get lots of copper, you want lots of small vias, rather than a few big ones, but the choice will be dictated by the board fabricator – after all, holes cost money – and also by the general interconnect density.</p>

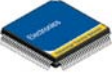
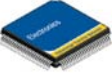

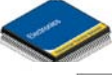
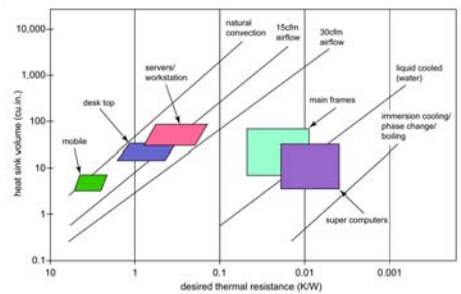
Passive methods of thermal management

<div data-bbox="151 212 268 286" data-label="Image"> </div> <h3>Thermal vias</h3> <ul style="list-style-type: none"> Diminishing returns from additional vias <ul style="list-style-type: none"> each via adds the cost of drilling not worthwhile using vias any smaller than can be drilled economically at high yield <div data-bbox="434 347 705 577" data-label="Figure"> <table border="1"> <caption>Approximate data from the graph</caption> <thead> <tr> <th>Number of vias</th> <th>Thermal resistance (°C/W)</th> </tr> </thead> <tbody> <tr><td>1</td><td>16</td></tr> <tr><td>5</td><td>10</td></tr> <tr><td>10</td><td>7</td></tr> <tr><td>25</td><td>4</td></tr> <tr><td>50</td><td>3</td></tr> <tr><td>75</td><td>2.5</td></tr> <tr><td>100</td><td>2.2</td></tr> <tr><td>125</td><td>2</td></tr> </tbody> </table> </div> <p data-bbox="146 604 319 618">Bruce Guerin, Electronics Cooling, August 2004</p> <p data-bbox="491 589 702 604">Electronics KTN – Knowledge For Growth</p>	Number of vias	Thermal resistance (°C/W)	1	16	5	10	10	7	25	4	50	3	75	2.5	100	2.2	125	2	<p>But it is not worth having too many vias, as each one adds to the cost of drilling, and it is not worthwhile using vias any smaller than can be drilled economically at high yield. This graph shows the diminishing returns from having additional vias.</p>
Number of vias	Thermal resistance (°C/W)																		
1	16																		
5	10																		
10	7																		
25	4																		
50	3																		
75	2.5																		
100	2.2																		
125	2																		
<div data-bbox="151 672 268 745" data-label="Image"> </div> <h3>Impact on the device thermal footprint</h3> <ul style="list-style-type: none"> Local heat removal has substantial impact on “thermal footprint” of device Size of thermal footprint set by competition between heat spreading and heat extraction <div data-bbox="438 757 710 1041" data-label="Diagram"> </div> <p data-bbox="491 1052 702 1068">Electronics KTN – Knowledge For Growth</p>	<p>Local heat removal, for example using thermal vias, is effectively sucking heat out where the hot component is, so not only are you removing the heat, but you are also reducing the size of the thermal footprint of the device. The heat is going down rather than sideways, and that can be quite significant if you have many components, where otherwise components that are running hot will be adding to the background temperature of nearby parts.</p>																		
<div data-bbox="151 1135 268 1209" data-label="Image"> </div> <h3>Heat sinks</h3> <ul style="list-style-type: none"> Categories of heat sink Heat sink technologies <ul style="list-style-type: none"> material properties and selection manufacturing methods matched fan/heat sink assemblies Heat sink attachment <ul style="list-style-type: none"> Thermal Interface Materials (TIMs) Factors affecting heat sink performance <ul style="list-style-type: none"> why heat sinks don't always work! Choosing a heat sink <p data-bbox="491 1516 702 1532">Electronics KTN – Knowledge For Growth</p>	<p>We have looked at ways in which board materials and design can be modified to improve thermal transfer, but there are only so many things you can do with the board. So typically a passive thermal management approach is going to involve a heat sink.</p> <p>Over the next slides we will be looking in more detail at different types of heat sink. Not only do they fall into four different categories, but we need to think about the technologies used to make them, and about how they will be fitted, in terms of thermal interface materials and so on. And we will finally look at some factors that affect heat sink performance – why heat sinks don't always work! – before we consider how to choose a heat sink.</p>																		
<div data-bbox="151 1599 268 1673" data-label="Image"> </div> <h3>Wisdom from Tony Kordyban</h3> <p>“The term heat sink must have been coined by an electrical engineer. It is too close in concept to the current sink used in circuit theory classes. The term sink is not to be found in heat transfer text books. When they mention those aluminium combs, they call them extended surfaces . . .</p> <p>“ . . . A sink is a place we can dump something and then pretend it no longer exists, such as the kitchen sink. We dump our greasy water down the drain and forget about it. But like that dirty water that shows up on somebody's beach, heat that flows into a heat sink doesn't just disappear, and we ultimately can't forget about it.</p> <p>“ . . . Heat can't flow forever into a chunk of aluminium and just disappear or migrate into a parallel universe. If you put heat into a chunk of anything, either the heat will flow out into something else or the chunk will get hotter.</p> <p>“ . . . So you attach your extended surface to the top of your component. The reason it looks like a thick comb is that all those fins pack a lot of extra surface area into a small volume. Heat energy now flows out of your component into the extended surface device, which has much more surface in contact with the air than the original component package.</p> <p>“ . . . The heat energy has still not disappeared. It is in the air. But keeping the room cool is somebody else's problem . . .”</p> <p data-bbox="459 1955 702 1971">Tony Kordyban, <i>Hot air rises and heat sinks</i></p> <p data-bbox="491 1977 702 1993">Electronics KTN – Knowledge For Growth</p>	<p>Who thought up the term “heat sink”? As Tony Kordyban said, it must have been coined by somebody who had been to circuit theory classes! It's worth pausing the presentation here, and reading this extract from his book. He rightly points out that heat that flows into a heat sink just doesn't disappear, and we ultimately can't forget about it. And he explains in his inimitable fashion what a heat sink looks like, and how it works. If you want more quotations like it, as they say, “buy the book”.</p>																		

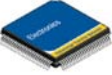
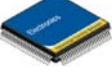
Passive methods of thermal management

 <h3>Categories of heat sink — 1</h3> <ul style="list-style-type: none">Four categories<ul style="list-style-type: none">by primary cooling methodPassive heat sinks<ul style="list-style-type: none">used where natural convection alone is sufficientdo not rely on specified local air velocityno moving parts and therefore highly reliablecan only remove small amounts of heatrequire careful positioning of enclosure vents <p>Electronics KTN – Knowledge For Growth</p>	<p>While the extended surface itself is definitely a passive method of thermal management for a device, the categories into which we put heat sinks relate to the primary method of cooling them, and that isn't always passive.</p> <p>We can divide heat sinks into four separate categories, the first of which are passive heat sinks. Just plain metal, and no other additions. If you remember Newton's Law of Cooling, the equation depends on the area, so what a heat sink tries to do is to get more area in a given space. Which is where the idea of "extended surface" comes from! But we will still call them heat sinks here because that's the more usual term.</p> <p>A passive heat sink is used where natural convection alone is sufficient, so we are not looking at having a specified air velocity local to the heat sink. There are no moving parts, and we have a very reliable product, but it can only remove small amounts of heat, and requires careful positioning of the vents in the enclosure, in order to maximise the free airflow.</p>
 <h3>Modelling natural convection</h3> <ul style="list-style-type: none">Natural convection<ul style="list-style-type: none">easier to model than forced convection?wrong! – with forced convection you are in control of what is happeningPlume of heat/movement of air<ul style="list-style-type: none">depends on hot air risingis much more difficult to simulate than forced convection case  <p>Electronics KTN – Knowledge For Growth</p>	<p>If you think back to one of the case studies in Section 3, the first case study was about natural convection and you would have thought that it would be easier to model than forced convection. In fact, with forced convection, where you have a fan or blower, you're in control of what is happening, whereas with natural convection the plume of heat and the movement of air that you get depend solely on the fact that hot air rises, and that is much more difficult to simulate than the forced convection case.</p>
 <h3>Categories of heat sink — 2</h3> <ul style="list-style-type: none">Passive heat sinksSemi-active heat sinks<ul style="list-style-type: none">hardware as for passive heat sinkuses system fans to give predictable air velocityfans<ul style="list-style-type: none">close enough to the heat sink to be effectivenot attached to itremoves more heat than natural convection <p>Electronics KTN – Knowledge For Growth</p>	<p>Our next category of heat sink is what is sometimes referred to as a semi-active heat sink. This uses the same form as the passive heat sink, but provides predictable airflow by using fans within the enclosure. These fans are close enough to the heat sink to be effective, but they are not attached to it, and not dedicated to specific devices. They remove more heat than is possible with only natural convection, and offer more design choices.</p>



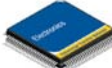
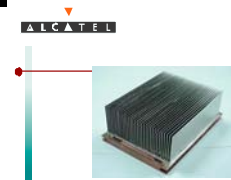
Passive methods of thermal management

 <p>Categories of heat sink — 3</p> <ul style="list-style-type: none"> ▪ Passive heat sinks ▪ Semi-active heat sinks ▪ Active heat sinks <ul style="list-style-type: none"> ▪ dedicated fan or blower <ul style="list-style-type: none"> ▪ fitted to specific heat sink or group of heat sinks ▪ fan and heat sink(s) carefully matched ▪ faster, localised air movement than semi-active ▪ removes more heat from smaller area <p style="text-align: right; font-size: small;">Electronics KTN – Knowledge For Growth</p>	<p>The third category is the active heat sink, which is the sort of thing you are likely to find fitted to the processor on a conventional tower PC. The active heat sink has a dedicated fan or blower which is fitted to a specific heat sink or group of heat sinks, and the fan and heat sinks are carefully matched for optimal performance. This gives faster localised air movement than a semi-active heat sink, and is able to remove more heat from a smaller area. Quite often the dedicated fan on the active heat sink will be supplemented by a larger air mover fitted to the enclosure.</p>
 <p>Categories of heat sink — 4</p> <ul style="list-style-type: none"> ▪ Passive heat sinks ▪ Semi-active heat sinks ▪ Active heat sinks ▪ Liquid-cooled heat sinks <ul style="list-style-type: none"> ▪ liquids have much higher heat transfer coefficients (typically $>1,000\text{W}\cdot\text{m}^{-2}\cdot\text{K}^{-1}$) than gases (typically $10\text{--}50\text{W}\cdot\text{m}^{-2}\cdot\text{K}^{-1}$) ▪ provide a 'cold plate' ▪ pump to circulate fluid past heat source ▪ external heat exchanger usually needed <p style="text-align: right; font-size: small;">Electronics KTN – Knowledge For Growth</p>	<p>The fourth category of heat sink is more unusual, and we will be looking at liquid-cooled heat sinks in later sections. They are much more effective, because liquids have much higher heat transfer coefficients than gases. So it's easy to provide a "cold plate", but you need external pressure or a pump to circulate the fluid past the source of heat.</p> <p>Also, unless you are taking water in from the main and just discharging it to waste, which isn't environmentally friendly and costs money, you need some kind of external heat exchanger. This will itself be an extended surface, often with fan assistance. So the liquid-cooled heat sink is really a speciality product for demanding applications.</p>
 <p>Categories of heat sink — 5</p> <p>Passive ➔ Semi-active ➔ Active ➔ Liquid-cooled</p> <p>➔ additional complication and expense</p> <ul style="list-style-type: none"> ▪ but higher effectiveness ▪ For extreme use <ul style="list-style-type: none"> ▪ thermoelectric cooling ▪ phase-change systems <ul style="list-style-type: none"> ▪ refrigerator ▪ direct boiling of liquid ("ebullient cooling") ▪ closed-loop heat pipe ▪ Choice depends on application <p style="text-align: right; font-size: small;">Electronics KTN – Knowledge For Growth</p>	<p>So we have four categories of heat sink. Going from one type to the next means more complication, and more expense, but higher effectiveness.</p> <p>As we will see later, that's not the end of the road, because there are other approaches for extreme uses, which you can examine when one of the standard categories doesn't work. These include thermoelectric cooling and phase-change systems. The choice depends totally on the application and on the severity of the thermal challenge.</p>
 <p>Thermal resistance limits for different types of cooling method</p>  <p style="text-align: right; font-size: small;">Electronics KTN – Knowledge For Growth</p> <p style="font-size: x-small;">Intel Technology Journal 2000 Q3</p>	<p>This is one of a number of diagrams in this material that show similar information, but in slightly different ways. In this diagram we see the thermal resistance ranges for different types of cooling method. The passive heat sink is at the extreme left, and we move progressively from active heat sinks to more challenging applications.</p> <p>What this diagram indicates is that, the more power you have, the less you can rely on convection and heat sinks, and the more you have to look at special-purpose ways of pulling heat out. Put in a different way, the closer you get to the state-of-the-art, the more effort you have to put into it.</p>

Passive methods of thermal management

 <h3>Heat sink technologies — 1</h3> <ul style="list-style-type: none">Contributors to effective thermal resistance<ul style="list-style-type: none">mounting areathermal pathpinsMaterial properties<ul style="list-style-type: none">aluminium and copperweight may be significant<ul style="list-style-type: none">copper the better conductor, but more densemany designs only use copper as insertsceramic composites and plasticsmild steel at lower-cost end of market <p>Electronics KTN – Knowledge For Growth</p>	<p>Most of our extended surfaces will be air-cooled, whether or not that air is moved artificially, and we are looking to make our extended surface effective. Major contributors to effective thermal resistance – that means low thermal resistance – are a bigger mounting area, a shorter thermal path, and lots of fins, to give us a large surface area.</p> <p>We have some choices of material, of which aluminium is probably the most commonly seen. It's not as good a conductor as copper, but it's a lot lighter, and when you are buying a heat sink, you are buying volume rather than weight, and you get much more aluminium for your buck.</p> <p>Copper is the better conductor, but it is also more dense, and the weight of the heat sink may be significant in some applications. As a result, a number of designs use copper as an insert, rather than as the whole heat sink. Typically we will be using metal rather than ceramic composites or loaded plastics, but this does depend on the requirement.</p> <p>At the lower-cost end of the market, coated mild steel will be perfectly adequate, and this is cheap to fabricate by pressing. If it does the job, there is no benefit in putting in more cooling than you need, because more cooling than you need costs money, and this is where simulation and verification by measurement help you to be certain that you are providing enough cooling, but not too much.</p>
 <h3>Heat sink technologies — 2</h3> <ul style="list-style-type: none">Material selection<ul style="list-style-type: none">wide range of aluminium alloys<ul style="list-style-type: none">6061 and 6063 for extrusionA380 for die-castingfar from a single material (as with "FR-4")choice related to<ul style="list-style-type: none">manufacturing methodassembly processdegree of post-fabrication finishing needed <p>Electronics KTN – Knowledge For Growth</p>	<p>Looking in a bit more detail at the material, we will find that our "aluminium heat sink" is not a single material, but has at least three common alloys, depending on the method of manufacture. This is a bit like the FR-4 generic term used for laminate. Which particular material is used depends on the manufacturing method, on any assembly processes that are involved in manufacturing and use, and the degree of post-fabrication finishing that is needed to meet customer expectations.</p>


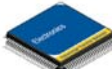
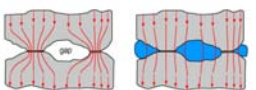
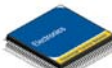
Passive methods of thermal management

 <h3>Heat sink technologies — 3</h3> <ul style="list-style-type: none"> Main manufacturing methods <ul style="list-style-type: none"> stamping — suitable for high-volume: clips/interface materials pre-applied extrusion — more flexible technique: can create a complex fin structure machining — for “pin fin” designs Alternatives for complex, high-density, pin fin heat sinks in aluminium or copper/bronze <ul style="list-style-type: none"> sand casting, lost core casting and die casting forging — can embed copper heat spreaders injection moulding  <p style="text-align: right; font-size: small;">Electronics KTN – Knowledge For Growth</p>	<p>The main manufacturing methods used for heat sinks are stamping, extrusion and machining. Stamping is a good technique for high-volume mild steel parts, where the heat sink is clipped into position, and interface materials can be pre-applied during manufacture.</p> <p>Extrusion is a more flexible technique that can create a complex fin structure, and aluminium is easily extruded, though there are limitations on complexity and feature size. Or one can machine from solid, which is a way of making parts such as pin fin designs.</p> <p>For those who don't like subtractive processes like machining, there are alternatives for complex high-density pin fin heat sinks, in the form of different kinds of casting, forging and even injection moulding. With some of these processes, it is possible to embed copper heat spreaders into the structure. So heat sinks can be made in many different ways, with different look and performance, and significantly different in cost.</p>																					
 <h3>Heat sink technologies — 4</h3> <ul style="list-style-type: none"> Higher-specification (higher-cost) heat sinks <ul style="list-style-type: none"> options extended by a range of fabrication processes much greater fin height-to-gap aspect ratio (between 20 and 40 is achievable) Bonded/fabricated fin heat sinks <ul style="list-style-type: none"> expose more surface area to air stream thermal performance improved significantly Typical manufacturing methods <ul style="list-style-type: none"> resin-bond flat fins onto grooved extruded base brazing to assemble fins to heat sink body <p style="text-align: right; font-size: small;">Electronics KTN – Knowledge For Growth</p>	<p>For higher-specification heat sinks the options are extended by a range of different fabrication processes that allow us to make fins that are thinner, more closely-spaced, and with a greater height-to-gap ratio. A bonded or fabricated fin heat sink, for example, which will expose more surface area to the air stream and improve thermal performance significantly, can be made by resin-bonding flat fins onto a grooved extruded base, or assembled by brazing fins onto the heat sink body.</p>																					
 <h3>Advanced Alcatel Heatsink Datasheet</h3> <p>□ Solder attachment of aluminium fins to copper base</p> <table border="1" data-bbox="215 1456 630 1601"> <tbody> <tr> <td>Length</td> <td>[mm]</td> <td>50</td> </tr> <tr> <td>Width</td> <td>[mm]</td> <td>70</td> </tr> <tr> <td>Height</td> <td>[mm]</td> <td>13</td> </tr> <tr> <td>Fin thickness</td> <td>[mm]</td> <td>0.4</td> </tr> <tr> <td>Min Fin Gap</td> <td>[mm]</td> <td>1</td> </tr> <tr> <td>Base Material</td> <td></td> <td>Copper - k=390 W/mk</td> </tr> <tr> <td>Fin Material</td> <td></td> <td>Aluminium - k=180 W/mk</td> </tr> </tbody> </table>	Length	[mm]	50	Width	[mm]	70	Height	[mm]	13	Fin thickness	[mm]	0.4	Min Fin Gap	[mm]	1	Base Material		Copper - k=390 W/mk	Fin Material		Aluminium - k=180 W/mk	<p>This is an example of an advanced heat sink, where aluminium fins are soldered to a copper base. Bear in mind that it is not worth buying an expensive part like this and then just blowing air on top of it, but you need to direct the airflow so that as much as possible of the extended surface is in contact with cooling air.</p>
Length	[mm]	50																				
Width	[mm]	70																				
Height	[mm]	13																				
Fin thickness	[mm]	0.4																				
Min Fin Gap	[mm]	1																				
Base Material		Copper - k=390 W/mk																				
Fin Material		Aluminium - k=180 W/mk																				

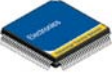
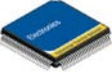


Passive methods of thermal management

 <h3>Heat sink technologies — 5</h3> <ul style="list-style-type: none">▪ Folded fin construction<ul style="list-style-type: none">▪ corrugated sheet of aluminium/copper▪ bonded with resin or braze to base plate or direct to the hot surface.▪ can be relatively light-gauge material yet still create a robust structure▪ More complex structures<ul style="list-style-type: none">▪ larger surface area▪ higher impedance to airflow▪ frequently used in combination with matched fan   <p>Electronics KTN – Knowledge For Growth</p>	<p>Another construction you will occasionally see is the “folded fin”, which uses a corrugated sheet of aluminium or copper that is bonded with resin or brazed to a base plate, or occasionally direct to the hot surface. Because the sheet is bonded to a support, it can be made of relatively thin-gauge material, yet still create a robust structure.</p> <p>As with all these more complex structures that offer a larger surface area, there can be a higher impedance to airflow, so heat sinks of this type are frequently used in combination with a matched fan.</p>
 <h3>Views of matched fan-heat sink assembly</h3>  <p>Electronics KTN – Knowledge For Growth</p>	<p>The idea of matching the fan to the heat sink is something that we will commonly see in heat sink designs for microprocessors. Just because the solution has to be cheap doesn't mean that good technology is not involved. In this case, the pairs of fins have been curved and angled, not only to fit in a greater surface area, but also to interact as much as possible with the air from the matched fan.</p>
 <h3>Heat sink attachment – 1</h3> <ul style="list-style-type: none">▪ Two stages in the assembly<ul style="list-style-type: none">▪ mounting heat sink to component or module▪ fixing heat sink within overall equipment▪ may be separate or combined operation▪ Key factors<ul style="list-style-type: none">▪ thermal quality of connection▪ mechanical security depends on<ul style="list-style-type: none">• mass of heat sink and other components• the mechanical environment <p>Electronics KTN – Knowledge For Growth</p>	<p>We haven't yet chosen our heat sink, but we need to think about how it will be fixed, preferably without involving too many bits and pieces and a lot of effort. This is important, because you are trying both to get a good thermal connection, and to make sure that the whole thing doesn't fall apart!</p> <p>There are two stages in the assembly, mounting the heat sink to the component or module, and fixing the heat sink within the overall equipment. The first stage might be carried out by the component or module manufacturer, or performed separately by the assembly house. Or both stages can be combined in a single operation – things that clip into place are an example of that.</p> <p>The corresponding key factors are the thermal quality of the connection, which we can enhance with suitable TIMs, and its mechanical security. How big a challenge this second factor is will depend on the mass of the heat sink and any other components, and also on the severity of the mechanical environment.</p>

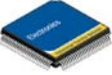

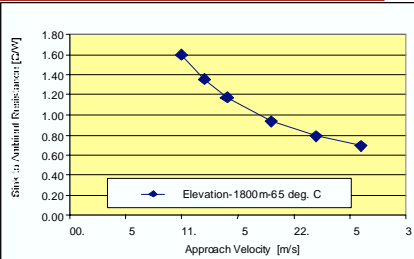

Passive methods of thermal management

 <h3>Heat sink attachment – 2</h3> <ul style="list-style-type: none"> Primary methods of assembly <ul style="list-style-type: none"> nuts/washers/bolts self-tapping screws rivets spring clips adhesives Choice of method will depend on <ul style="list-style-type: none"> the requirement for rework the application life the preferences of the assembler the equipment available <p style="text-align: right; font-size: small;">Electronics KTN – Knowledge For Growth</p>	<p>Primary methods of assembly “include but are not limited to”, as they say, nuts/washers/bolts, self-tapping screws, rivets, spring clips and adhesives. All of these vary in their cost, in their ease of assembly, in whether or not they can be made foolproof, and in whether they can be easily disassembled.</p> <p>The choice of method will depend on whether there is any requirement for rework – would you be able to take this to pieces in the field? And it also depends on the intended application, as this affects both the technical requirements and the target cost. The choice will also depend on the preferences of the assembler, and on the equipment and materials available, an aspect that is becoming increasingly important with the trend for assembly partners to be 6,000 miles away!</p>
 <h3>Need to manage the interface</h3> <ul style="list-style-type: none"> Contact only between high points <ul style="list-style-type: none"> always microscopically a rough surface Lack of intimate contact results <ul style="list-style-type: none"> increased thermal resistance often unstable in value Range of different interface materials <ul style="list-style-type: none"> oils and greases compliant filled rubbers shims of soft metal  <p style="text-align: right; font-size: small;">Electronics KTN – Knowledge For Growth</p>	<p>Typically we will be doing more than just pressing hot parts against the heat sink, but we will need to use some kind of thermal interface material. In an ideal world, surfaces would be microscopically flat, and touch at all points. In practice, they are far from flat and on the microscopic scale they have rough surfaces.</p> <p>The result is that mating surfaces make contact only between high points, and the lack of intimate contact that results produces an unexpectedly high thermal resistance, and one that is often unstable in value. In order to improve the intimacy of contact, we normally use materials that are referred to as TIMs, Thermal Interface Materials. The original materials were loaded oils and greases, but silicone grease is not favoured by most assembly houses, as it impairs solderability. Alternatives include compliant filled rubbers and shims of soft metal.</p>
 <h3>TIM materials</h3> <ul style="list-style-type: none"> There are a lot of TIMs! <ul style="list-style-type: none"> PCMs, gap fillers, putties, greases, insulators, solders Kordyban's Grilled Cheese Sandwich (GCS) theory: <ul style="list-style-type: none"> more to it than high bulk k flatness, pressure, contact TIM performance a function of <ul style="list-style-type: none"> elastomer, filler, filler shape, size and loading Innovations <ul style="list-style-type: none"> metal TIMs, nanoparticle fillers, carbon nanotubes <p style="text-align: right; font-size: small;">Electronics KTN – Knowledge For Growth</p>	<p>TIMs are a bit of a minefield and guidance from Tony Kordyban is always welcome. He refers to the TIM as a grilled cheese sandwich, which is actually more what would be referred to in the UK as a “toastie”. In other words, the quality of the result depends on intimate connection, the right material in the sandwich, flatness, pressure and intimate contact. As with the toastie, if you are trying to produce the right result you need enough filling, sufficient contact with the surfaces, and the right combination of pressure and temperature.</p> <p>The performance of particular TIMs will depend on the size and shape of the filler, the amount of filler in the structure, and whether this is elastomer or grease. And it's an area where there are many innovations, some coming out of nanotechnology.</p>


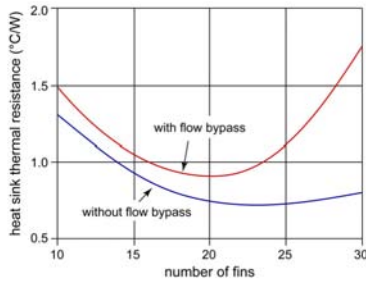
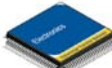

Passive methods of thermal management

 <p>TIM materials</p> <p>Kordyban's Korner -- GCS Theory Applied to Thermal Interface Materials</p> <p>Tony Kordyban, Thermal Guru, tonykordyban@gmail.com</p> <p><i>There are 819 kinds of Thermal Interface Material (TIM) - those greases, putties, pads and phase change compounds that go between hot electronic components and their heat sinks - and just as many ways to misuse them. This article covers all of them by stating: If a TIM disappoints, it is all your fault for making one or more of these silly mistakes.</i></p> <p style="text-align: right;">Electronics KTN – Knowledge For Growth</p>	<p>There is more to it than having something which is a really good thermal conductor, and there are many TIMs, and their performance is exceptionally variable. From the convenience point of view, probably the best ones are demountable, in order to be able to take the assembly apart. Of course, if you put it back together again, almost certainly your silicone rubber shim will have been squashed, and remain a bit squashed, so it may not fit together correctly, but if you compare that with the problems of repairing a soldered or brazed part it's a lot easier to deal with.</p>
 <p>TIM materials</p> <ul style="list-style-type: none"> ▪ Which material to choose? ▪ Decision points <ul style="list-style-type: none"> ▪ TIM an adhesive? <ul style="list-style-type: none"> ▪ needs maintained pressure? <p>Top 20 considerations:</p> <ul style="list-style-type: none"> • Adjacent materials • Q dissipated • Clamping arrangement • Footprint dimensions • Clamping pressure • Wapage • Surface finish • Electrical isolation • Reuse / rework • Horizontal / vertical • Vibration / shock • Vacuum application • Belcore reqs. • Flammability rating • Shape (preformed) • Adhesive <ul style="list-style-type: none"> • Non-stick • Sliding joint • TIM packaging <p style="text-align: right;">Electronics KTN – Knowledge For Growth</p>	<p>As always, which material you choose will depend on the application. This is Tony Kordyban's list of 20 things that you might like to think about in relation to the contents of your sandwich, you know, the cheesy bit! In particular you have to make the decision as to whether the TIM will itself be an adhesive, and bond the surfaces together, or whether it will be a style that needs maintained pressure in order to perform.</p>
 <p>How not to do it!</p>  <p style="text-align: right;">Electronics KTN – Knowledge For Growth</p>	<p>This is a real example of how not to do it. The package is a Column Grid Array, with a ceramic base and a metal lid. So not too much heat will be going down the columns, and most of the heat needs to leave the package at the top surface. A large heat sink has been chosen, and applied with what appears to be an uncontrolled amount of thermal adhesive.</p> <p>While the heat sink was removed deliberately (and didn't just fall off!), one wonders how effective the heat sink was. For example, what was the variability in thermal contact resistance between the heat sink and the device? Was the thickness of the TIM controlled? And how effective would the peripheral areas of the heat sink be, given the length of their thermal path from the device?</p>


Passive methods of thermal management

 <p>Factors affecting heat sink performance – 1</p> <ul style="list-style-type: none"> Major effects on heat transfer <ul style="list-style-type: none"> fluid velocity through heat sink type of flow (turbulent flow removes heat from a heat sink surface more readily than laminar flow) <p>the significance of laminar flow</p> <p>"Despite the multitude of materials and interfaces within an electronic package, the largest thermal resistance, and consequently the controlling resistance in the path between the source and the sink, is generally the boundary layer or film resistance at the solid/fluid interface.</p> <p>Culham, Teertstra and Yovanovich, <i>Natural convection modeling of heat sinks using web-based tools</i></p> <p>Electronics KTN – Knowledge For Growth</p>	<p>Whilst we need to get heat to the heat sink, the major factor that affects its performance is heat transfer from the extended surface. This is limited by the fluid velocity through the heat sink and is significantly affected by the type of flow. Turbulent flow removes heat from a heat sink surface much more readily than laminar flow.</p> <p>The idea of laminar flow is something that always reminds me of times on a British beach when the wind howled. You got down as close as you could to the sand, and found that the wind speed close to the ground was less than it was when you were standing up. Much later I realised that this was because of the drag exerted by the ground on the wind, and it's a very similar situation with a heat sink. If the flow is laminar, that is the airflow is moving across the surface, it's not making very good contact with the surface.</p> <p>For maximum effect, the airflow should actually impinge on the surface, as well as being turbulent. It's interesting that Culham and his colleagues, based on their modelling, made the statement that it's the boundary layer or film resistance at the solid/fluid interface that is the largest thermal resistance in the path.</p>
 <p>Advanced Alcatel Heatsink Performance</p>  <p>Note: No side or top flow bypass</p>	<p>It's always helpful to have more air, as this graph of sink-to-ambient thermal resistance against approach velocity indicates. But note the diminishing returns – this is not a linear graph.</p>
 <p>Factors affecting heat sink performance – 2</p> <ul style="list-style-type: none"> Fin air flow velocity affected by <ul style="list-style-type: none"> fin thickness fin shape number of fins and fin spacing The problem of "flow bypass" <p>Electronics KTN – Knowledge For Growth</p>	<p>What is actually important is the airflow around the surface of the fins, and this may be different from the approach velocity, and is often much lower. The local airspeed will be affected by the thickness of the fins, their shape, and especially their height and spacing. If we have many tall fins packed closely together, the air may fail to remove heat from the bottom of the deep channels, a problem that is given the name of "flow bypass".</p>

Passive methods of thermal management

 <h3>Effect of flow bypass</h3>  <p style="text-align: center;">Electronics KTN – Knowledge For Growth</p>	<p>Here we can see graphically what happens when air doesn't get down into the crevices. Beyond a certain number of fins (and this of course does depend on the spacing and other characteristics), adding extra fins becomes counter-productive, so you are using a heat sink technology that's costing you too much.</p>
 <h3>Factors affecting heat sink performance – 2</h3> <ul style="list-style-type: none"> ▪ Fin air flow velocity affected by <ul style="list-style-type: none"> ▪ fin thickness ▪ fin shape ▪ number of fins and fin spacing ▪ The problem of "flow bypass" ▪ Fin heat distribution, affected by <ul style="list-style-type: none"> ▪ fin thickness ▪ fin height <p style="text-align: center;">Electronics KTN – Knowledge For Growth</p>	<p>The amount of heat that is transferred from an extended surface is a function of its temperature, and fins will inevitably be cooler at the top than the bottom, so that the parts of the fin furthest away from the heat source will be moving less heat per unit area than those close to the source.</p> <p>This actually compounds the flow bypass problem, because it is the hottest parts that get the least airflow. It's absolutely obvious, and it comes straight out of the equations, but it is amazing how many people like big heavy heat sinks, where a small one might do just as well for less money!</p>
 <h3>Choosing a heat sink</h3> <ul style="list-style-type: none"> ▪ Luxeon's suggested approach <ul style="list-style-type: none"> ▪ Step 1: Determine reliable Θ_{JA} from $\Theta_{JA} = \frac{(T_J - T_A)}{P}$ with absolute max. value of T_J and worst-case operating ambient condition T_A ▪ Step 2: Subtract the Θ_{JS} figure for the device, in order to obtain the target Θ_{SA} ▪ Step 3: Determine the best heat sink configuration for the application from the target value of Θ_{SA} <p style="text-align: center;">Electronics KTN – Knowledge For Growth</p>	<p>As with choosing a car, people have different ways of choosing heat sinks. It is tempting to go for high quality of manufacture, a fine glossy surface, and a top-end specification, but that isn't an engineering approach! What Luxeon suggested as a sequence was to start with a reliable figure of the required thermal resistance from junction to ambient, based on the worst-case power, the absolute maximum value of junction temperature, and the worst case of operating ambient. The junction-to-heat-sink value for the device is subtracted, giving a target thermal resistance for the heat sink, from which one can then select the best heat sink configuration.</p> <p>While a step-by-step approach commends itself, we have to be aware of the many uncertainties associated both with the internal thermal resistance from junction to heat sink, and the characteristics of the heat sink, which are highly dependent on local airflow. As an approach, it is fine, but only as far as it goes, and you will need to build in some margin, and preferably include some sensitivity analysis. If you are operating close to the margin, then you really need to simulate the thermal environment, in order to get more reliable and accurate results.</p>

Passive methods of thermal management

 <p>Passive methods of thermal management</p> <ul style="list-style-type: none">▪ Basic methods<ul style="list-style-type: none">▪ radiation▪ free convection▪ conduction▪ Conduction routes<ul style="list-style-type: none">▪ components▪ PCB▪ heat sinks <p>can only choose the right parts</p> <p>can optimise by design</p> <p>considerable scope when combined with forced convection</p> <p><small>Electronics KTN – Knowledge For Growth</small></p>	<p>In this section we have looked at various passive methods of thermal management, and seen our reliance on conduction routes through components, boards and heat sinks. For conduction within the component, the only design choice is to choose appropriate parts, and of course to make sure that you buy the part that has been thermally modelled, and not something that looks the same but doesn't have the same construction!</p> <p>For the board, the designer can choose appropriate materials, and can employ features such as thermal vias to aid in heat distribution. But the major conduction route will be through the heat sink, and here there is considerable scope for choosing cost-effective solutions, though we will probably need to rely on more than just free air movement. It's this topic of forced convection that we will see in our next section.</p>
---	--