The rationale for taking action

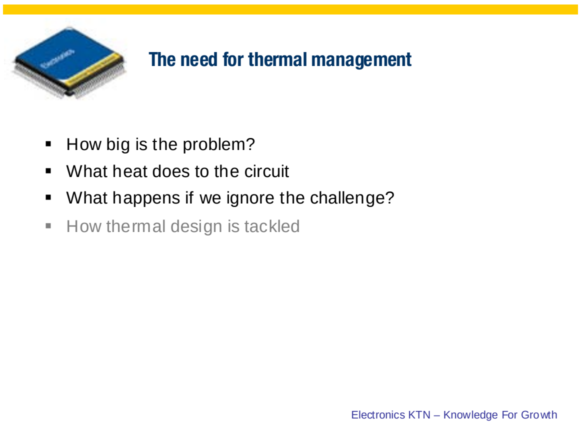

So far we have taken a look at the size of the problem, and at some of the technology trends. We are now looking at the rationale for taking action, in terms of what heat does to the circuit and what happens if we ignore the challenge, and in the section that follows we will be outlining how thermal design is tackled.

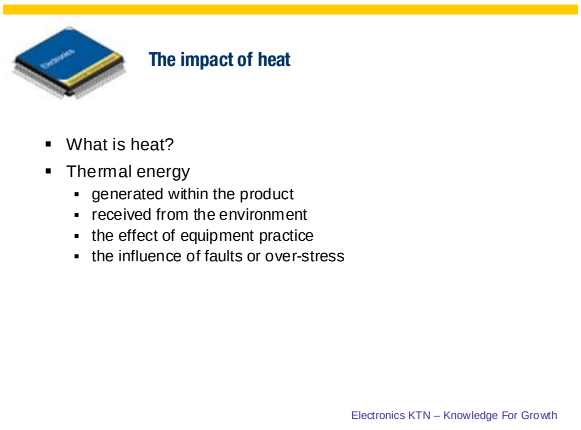

Before looking at what heat does to the circuit, we will be reviewing some basics, and in particular where the heat we have to manage is coming from.

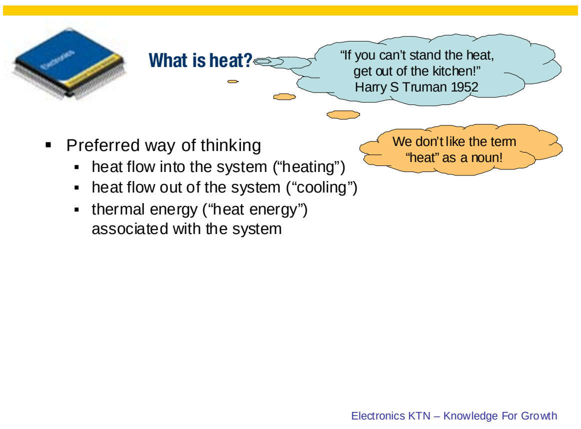

I am starting off by asking the question: What is heat? I know Harry Truman said "If you can't stand the heat, get out of the kitchen", but thermal engineers don't actually like the word "heat" used as a noun.

The reason is that we are focused on temperature, on the heat flow into the system and the heat flow out of the system, and on the thermal energy that is associated with the system. In other words, we're talking about "heating" and "cooling", and it's the heat energy, or the thermal energy, that we are trying to manage.

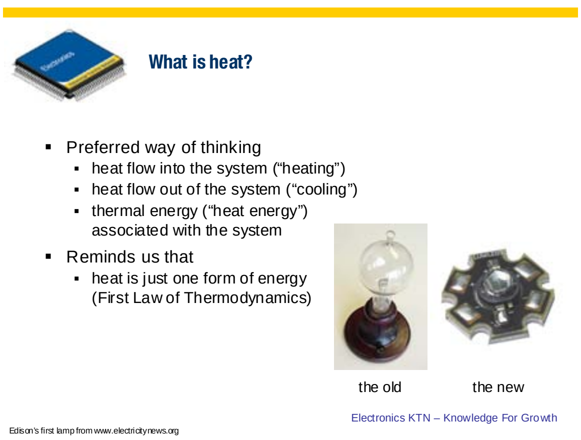

The term "heat energy" also reminds us that heat is just one form of energy: generally we are converting electrical energy into heat, plus some useful work.

As an aside, one of the problems with moving from lamps to LEDs is that, with a conventional tungsten lamp, while you're getting light out of it, most of the heat is being radiated at long wavelengths, and the amount of heat energy that stays around the envelope is relatively small. With an LED, though it is vastly more efficient, but the spare heat that you are putting in has to be taken out of the package, so that heat management of LEDs actually presents more problems than with ordinary filament lamps.

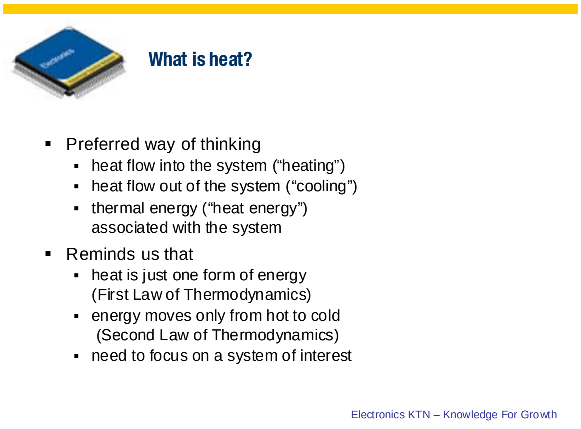

Using the expression "heat energy" also reminds us that thermal energy only moves from hot to cold – that's the only way it will flow. As Flanders and Swann put it: "You can't get heat from a cooler to a hotter, and that's the Second Law". And that is what a thermal management solution provides, a cooler place to go.

The final benefit of thinking of heat as energy is that it reminds us that we need to focus on a system of interest. We are looking at heat flow into the system and heat flow out of the system, and for that we need to define what our system is. We might be looking at the individual package; we might be looking at the board; or we might be looking at the total system. And our system might be the whole world because, at the end of the day, the heat that we generate is going out into the environment. So it's a system view that the thermal manager takes.

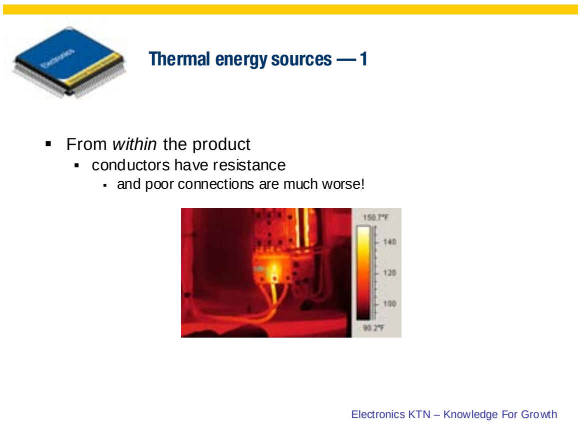

So where does this heat come from? Or rather, what are the thermal energy sources? Some of them are within the product. For example all conductors have some resistance, and poor connections are much worse, as you will see in this infrared photograph of an connector, showing a significant temperature rise probably due to a loose or dirty connection.

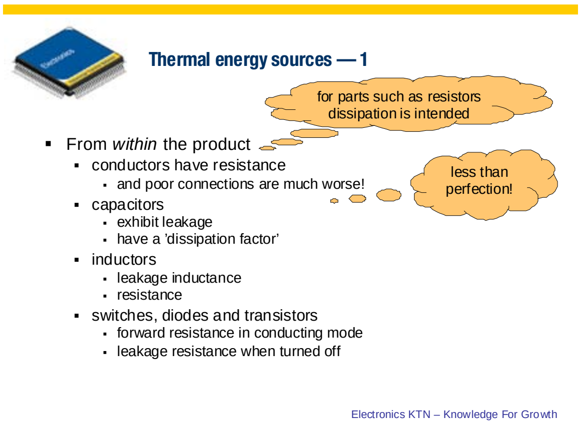

Components such as resistors are intended to dissipate heat – I squared R heating makes this inevitable. But most other components dissipate heat because they are less than perfect. You have seen with interconnections that they may have resistance, and this is particularly important at higher current. Capacitors and inductors are also less than perfect, so there will be some kind of dissipation in them. Certainly if you have any kind of semiconductor this is going to be less than 100% efficient, devices exhibiting a forward resistance when conducting, and leakage when turned off.

Thermal energy sources — 2 ▪ From the environment ▪ external solar radiation Electronics KTN – Knowledge For Growth	But we can't be totally inward-looking in our thermal design. Not only is energy converted into heat within the product, but heat will be gained from the surroundings, and this can be particularly important in applications exposed to an adverse environment, such as automotive and aerospace.
The effects of some sun!  Electronics KTN – Knowledge For Growth Doug Klassen on 40on2.blogspot.com 14 August 2006	If you've ever asked the question how hot your bike gets while merely sitting in the sun, Doug Klassen has posted this picture. Admittedly Doug lives in Arizona, but it indicates why designers of vehicle electronics have to consider some highly adverse conditions!

Thermal energy sources — 2
▪ From the environment
  ▪ external solar radiation
  ▪ heat inputs from local environment
    ▪ other electronics
    ▪ external sources of heat
  ▪ critical applications in automotive/aerospace
    ▪ high maximum temperature
    ▪ wide temperature excursion
    ▪ simultaneously subjected to vibration and shock
Electronics KTN – Knowledge For Growth
	Even when there is no significant heat input from the external environment, there will of course be heat inputs from the local environment, whether these are other pieces of electronic equipment or external sources of heat, like the vehicle engine or exhaust system, which points up the challenges within automotive and aerospace applications.
	
	Not only can maximum temperatures be high, but environments with a high daytime temperature are often very cold at night, and it can also be very cold in the upper atmosphere, so that the diurnal temperature excursions seen by the electronic assemblies are very wide. And it's those changes which tend to impair reliability. At the same time, products are subjected to vibration and shock, which can be a lethal combination – joints that have been weakened by temperature excursion can fail in shock conditions.

The rationale for taking action

Thermal energy effects

- Effect of "equipment practice"
  - thermal energy gain from
    - adjacent components
    - nearby circuits
  - *distribution* of heat will depend on
    - board layout
    - board arrangement within racking
    - way equipment is arranged within room

Electronics KTN – Knowledge For Growth

The performance of an electronic product depends not only on the elements that have been brought together and interconnected, but is impacted by the enclosure and by the way in which the elements are arranged within it. Not everything is designed from scratch, and usually there will be some elements of standard "equipment practice" which will determine dimensions, spacing, and the way that boards and circuits are interconnected.

Equipment practice choices therefore will affect the thermal energy gained from adjacent components and nearby circuits, and the distribution of heat will depend on the board layout, the board arrangement within the racking and, at a more macro level, the way that the whole equipment is arranged within the room.

Thermal energy effects

- Effect of "equipment practice"
  - heat from
    - adjacent components
    - nearby circuits
  - *distribution* of heat will depend on
    - board layout
    - board arrangement within racking
    - way equipment is arranged within room
- Influence of faults or over-stress conditions
  - the challenge of "over-clocking"
  - blocked filters or fan failure
  - component failure

Electronics KTN – Knowledge For Growth

The thermal energy distribution will also depend on faults and overstress conditions, and that is where modelling is particularly useful in telling us the potential effect of making changes. For example, what happens if the user decides to overclock the processor? Or if something goes wrong, like a door being left open, filters becoming blocked or a fan failing? After all, this happens in real life, and likely failure modes are something that we can simulate.

Hopefully though you won't get this kind of extreme thermal problem which was actually a fire due to battery failure!

The impact of heat

- What is heat?
- Thermal energy
- Temperature effects on the circuit
  - changes with temperature
  - changes over life
    - drift
    - failure

Electronics KTN – Knowledge For Growth

The thermal energy distribution within the system will result in temperature effects on the circuit, and here we can distinguish two sorts of change – properties that change with temperature, and those that also change over life.

Changes with temperature: passives

- Resistors: TCR
  - positive TCR for metals

Material	Resistivity (μΩ-cm)	TCR (ppm/°C)
Alloy 42	66.5	1,000
aluminium	2.83	3,400
copper	1.72	3,900
kovar	48.9	3,700
nickel	7.80	6,000
silver	1.63	3,800

Electronics KTN – Knowledge For Growth

Passive components will generally have well-defined variation with temperature. For example, metals have a high positive Temperature Coefficient of Resistance …

The rationale for taking action

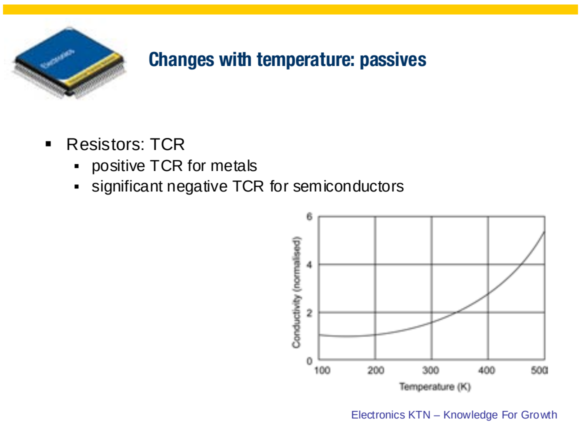

… whereas semiconductors become considerably more conductive with temperature, so that their TCR is highly negative.

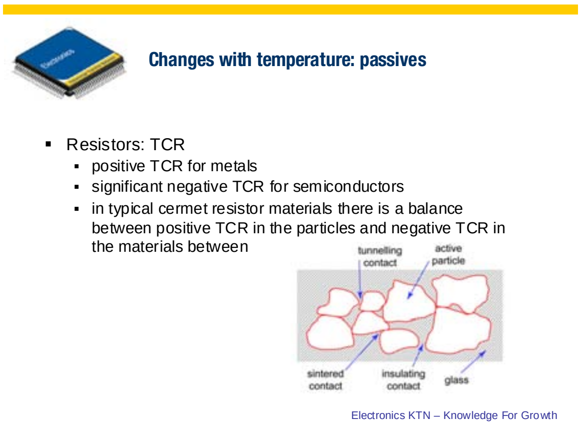

Resistors need to be more stable for most electronics applications and are controlled in the range around zero in the complex cermets and similar that are typically used, by balancing metallic and semiconductor properties.

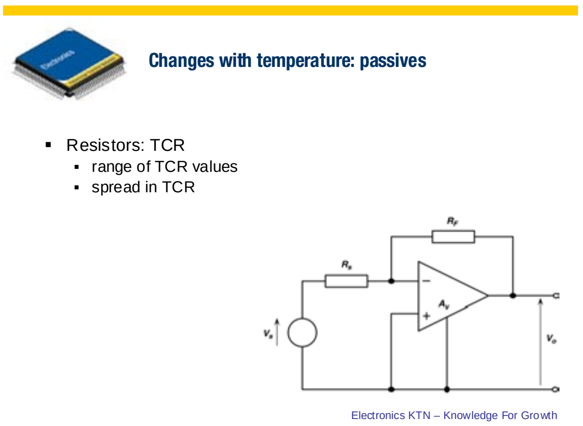

As always, there will be a spread in characteristics that has to be allowed for by the electronic designer. If we take the case of a basic inverting amplifier: if the two resistors “track” perfectly, the only effect on performance will be caused by any temperature changes to the open-loop gain of the operational amplifier. However, if the match is less than perfect, and this can be for temperature reasons, there will be an error.

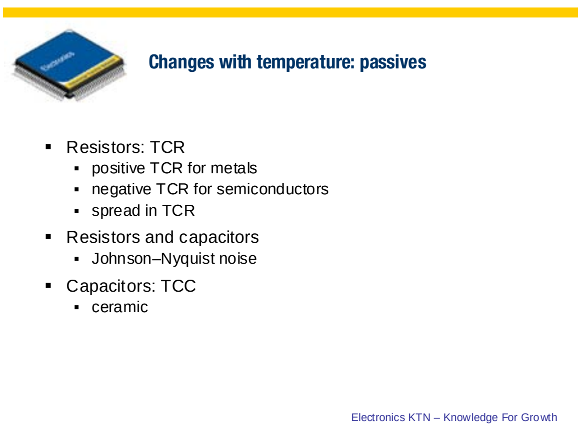

Another parameter that is temperature-dependent is Johnson-Nyquist noise, otherwise known as thermal noise, which is generated by thermal movement of charge carriers inside any conductor. Also referred to as white noise, the power spectral density is nearly constant through the frequency spectrum, and the power is a function of the absolute temperature in degrees Kelvin.

With capacitors, we have a Temperature Coefficient of Capacitance that is related to the basic structure of the material.

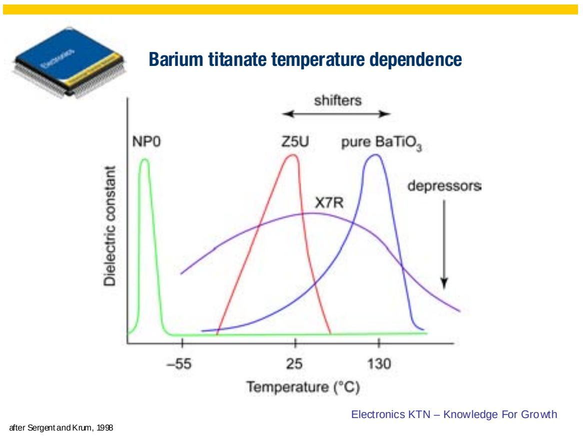

For ceramic dielectric types, there is a trade-off between stability of value and the permittivity of the basic dielectric. If you want a stable part, this needs to have an NP0 characteristic, which has a low permittivity at normal working temperatures, and in consequence is only available in modest values.

Higher values are obtained only at the expense of some stability, but it's worth adding to the mix to increase the dielectric constant over a relatively narrow temperature range, if all you need is a decoupling capacitor that is as small as possible.

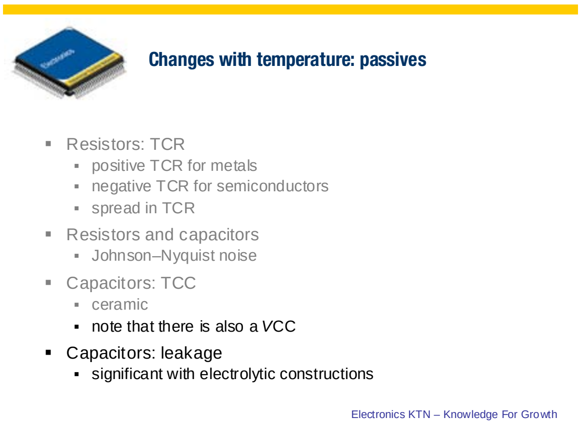

Of course, there is also a voltage coefficient, and there are also changes in leakage with temperature, but the last of those apply much more to electrolytic types than to ceramic capacitors.

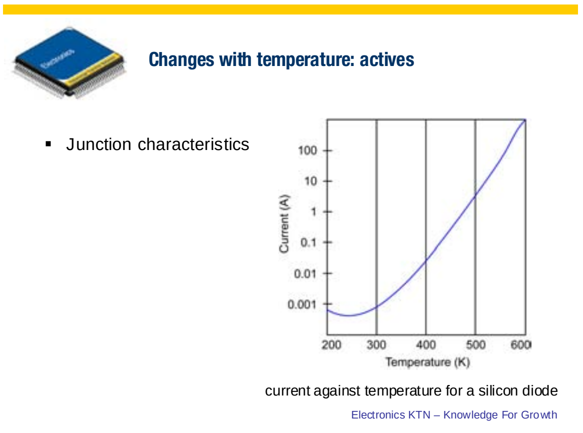

The characteristics of semiconductors are significantly more temperature-sensitive than most passive components. With diodes, for example, the reverse saturation current at room temperature doubles with each increase in temperature of approximately 10°C, giving the resultant curve shown here.

Under constant current conditions, the voltage across the diode will decrease linearly with temperature, having a temperature coefficient in the range -1 to -3 mV/K, a feature often used in temperature-sensing applications.

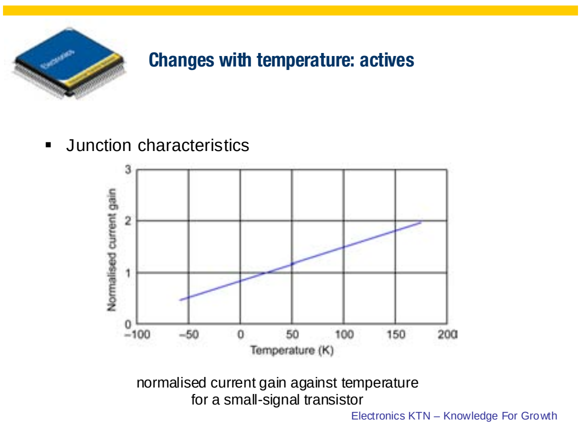

Transistor characteristics also change with temperature, and this illustrates the typical change with temperature of the common-emitter current gain for a small-signal transistor.

<table>
<tr>
<td>
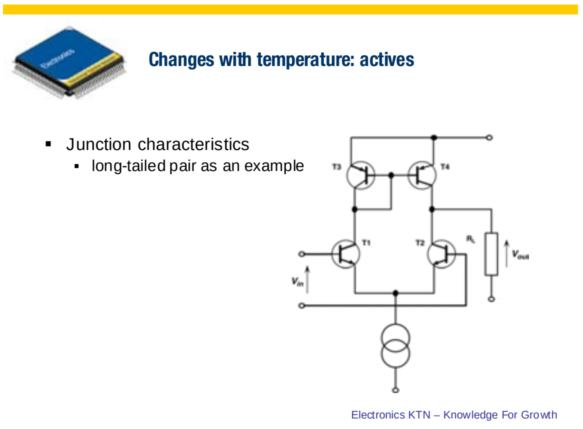

</td>
<td>
Even when the power dissipation and consequent internal heat generation is minimal, differences in characteristics between matched devices may impact on overall circuit performance. A simple example of this is the long-tailed pair, where the balance between the two halves of the circuit will be impaired if the transistors forming the long-tailed pair are not matched.

For this they have to be in intimate thermal contact: when such configurations used to employ discrete transistors, it wasn’t uncommon for the parts to be coupled physically by a common heat sink. This is obviously less of a problem when the circuit is implemented within an integrated circuit, but bear in mind that there can be substantial gradients even within an IC die.
</td>
</tr>
<tr>
<td>
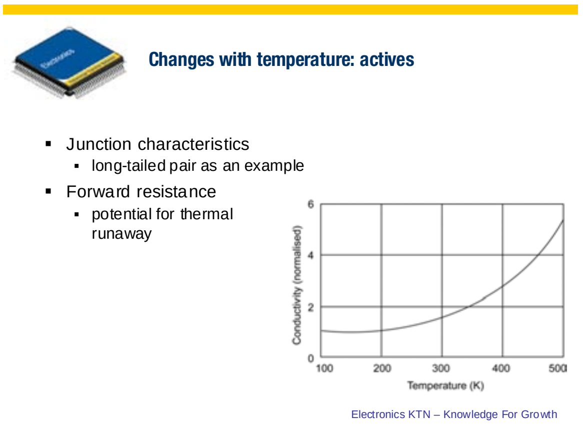

</td>
<td>
We’ve already seen that the conductivity of the basic silicon increases with temperature, and related effects can have serious effects for bipolar power devices. For example, in a circuit where a constant voltage is maintained across a base-emitter junction, elevated temperatures can produce a phenomenon called “thermal runaway”, in which the increased temperature causes the collector current to rise; in turn, this causes the temperature to increase further, raising the current still more. The end result can be failure as a result of excessive heat, excessive current, or both.

The problems are usually minimised by placing a resistor in the emitter circuit. As the current increases, the voltage across the emitter resistor increases, lowering the effective base-emitter voltage, which acts to reduce the current, thus providing protective negative feedback.
</td>
</tr>
<tr>
<td>
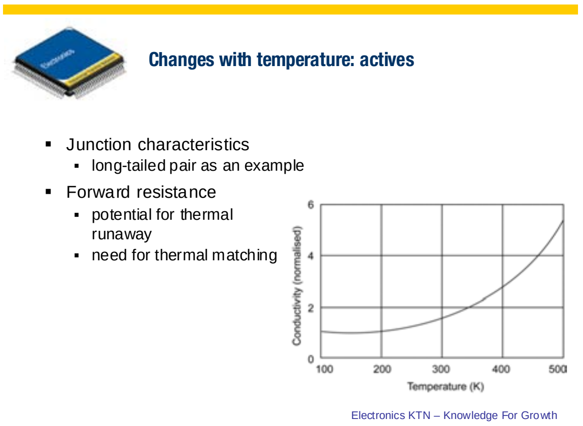

</td>
<td>
In power devices, it is common to operate devices in parallel to increase the available current. If two discrete devices are not perfectly matched, both electrically and thermally, one will tend to carry more current than the other, resulting in a rise of temperature that further increases the imbalance. To avoid one device carrying the bulk of the current, resulting in failure, negative feedback is usually fitted in the form of low-ohm resistors in the individual circuits, although obviously this lowers the overall efficiency.
</td>
</tr>
</table>

The rationale for taking action

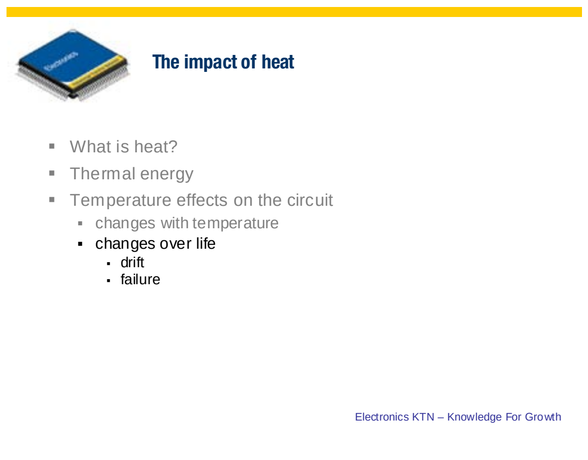

We've seen that at any one time the parameters of a component will change with temperature, and this may have a corresponding effect on circuit performance, and the consequent need to make modifications to the electronic and thermal design of the product.

But temperature effects on components also have a time dimension, and we will see changes over life, whether these are merely parametric drift or result in device failure.

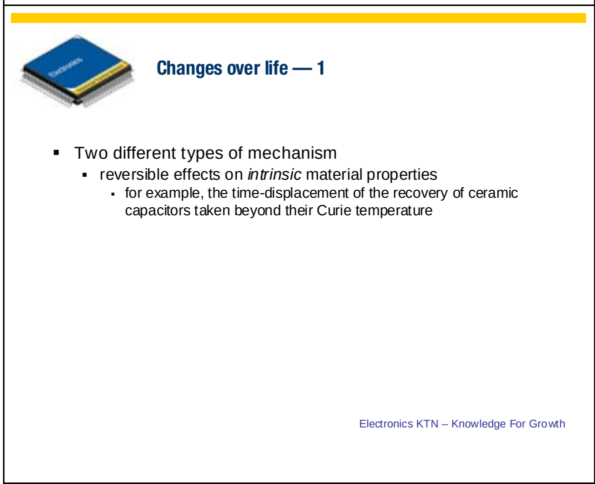

There are two different types of mechanism, the first of which is a reversible effect on the intrinsic properties of the material. An example of this is a high-permittivity capacitor. When this is being soldered into position it is taken above its Curie point, and will have at that time the highest value of capacitance it will ever possess! If you measure it only an hour later, you will find it has reduced in value by a percent or so, and this it continues to do throughout its life. Every decade hour, the value of capacitance will reduce by the same amount – the same percentage in the first hour, the next 10 hours, the next 100 hours, the next 1,000 hours, and so on. But this effect is reversible – if it goes out of the bottom end of the tolerance band, you can take the part back above its Curie point and restore its initial value.

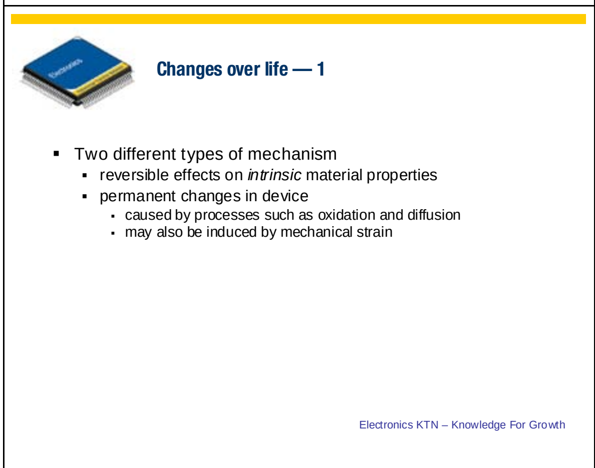

Much more common are permanent changes to the device that may be caused by processes such as oxidation and diffusion, or may also be induced by mechanical strain. If you have a structure that is being heated up and then cooled down, during the process it will be subject to strain because of differences in the coefficient of temperature expansion of the materials. Such strain, when repeated many times, may have an adverse impact on performance.

The rationale for taking action

Changes over life — 1

- Two different types of mechanism
  - reversible effects on *intrinsic* material properties
  - permanent changes in device
    - caused by processes such as oxidation and diffusion
    - may also be induced by mechanical strain
- Timescale for changes
  - from extremely short-term to very long-term
  - power pulses produce transient changes
  - most components will drift after an extended period, especially at high temperature
    - changes are predictable, so can be allowed for

Electronics KTN – Knowledge For Growth

The timescale for all these changes varies enormously. Some are extremely short term – for example, if the system is subjected to power pulses, we will be able to see transient changes that are reversible.

But after an extended period most components will drift permanently. For example, resistors will tend to increase in resistance, capacitors will tend to reduce in capacitance value but become higher in leakage. Such changes are usually faster at higher temperatures, but are predictable, so can be allowed for in design.

Changes over life — 2

- Drift
  - can be allowed for in design phase
  - mostly no adverse implication for reliability

Electronics KTN – Knowledge For Growth

Changes due to temperature or to parametric drift are predictable, within limits, and can be allowed for at the design phase by proper tolerancing. We can also allow for the effect of parts heating up when powered, although for that we need information on their temperature rise.

Such temporary temperature-related changes and permanent parameter drift mostly have no adverse implication for reliability.

Changes over life — 2

- Drift
  - can be allowed for in design phase
  - mostly no adverse implication for reliability
- *But*, temperature cycling induces strains
  - caused by expansion mismatches in structure
  - level determined by
    - absolute temperature
    - temperature differences within the structure

Electronics KTN – Knowledge For Growth

However, temperature excursions induce strains that are caused by mismatches in the CTE (coefficient of thermal expansion) within the structure, the strain level depending both on the actual temperature and on the temperature differences. And in many cases it is the temperature differences within the structure that are more important than the temperature itself.

Changes over life — 3

- Strain-induced change
  - only totally reversible if strain takes place within a linear part of the stress-strain curve
  - once materials start to yield, permanent distortion will occur
  - eventual catastrophic failure of some sort likely

Electronics KTN – Knowledge For Growth

Systems will recover from strain-induced change, but only if the materials are within their elastic range. If the strain takes place within a linear part of the stress-strain curve, the changes induced by strain will be totally reversible, and the structure will revert to its original condition when the stress is removed.

However, once materials start to yield, permanent distortion will occur, and eventual catastrophic failure of some sort is likely.

The rationale for taking action

Changes over life — 3

- Strain-induced change
  - only totally reversible if strain takes place within a linear part of the stress-strain curve
  - once materials start to yield, permanent distortion will occur
  - eventual catastrophic failure of some sort likely
- Time to failure will depend on
  - the number of stress cycles
  - their amplitude
  - rate of change within cycle

Electronics KTN – Knowledge For Growth

The time to failure, and the resistance to failure of the component or structure, will depend both on the number of stress cycles, in this case the number of thermal cycles, and also on the severity of the stress. This is a function of the amplitude of the thermal cycle and of the rate of change – moving a part very quickly from hot to cold or vice versa is significantly more stressful than subjecting it to a gradual change, primarily as a result of the temperature variations within the device.

Repetitive thermal cycling

A computer that is turned on twice a day, every day for 15 years, will accumulate about 11,000 thermal fatigue cycles. A television that is turned on 10 times a day, every day for 15 years, will accumulate about 55,000 thermal fatigue cycles. An automobile that is started 10 times a day, every day for 20 years, would accumulate 73,000 thermal stress cycles. A satellite in orbit around the Earth experiences a thermal cycle about every 90 minutes. In 20 years it can accumulate about 117,000 thermal cycles.

Dave Steinberg, *Preventing thermal cycling and vibration failures in electronic equipment*

Electronics KTN – Knowledge For Growth

It's a great temptation to underestimate the number and severity of the thermal cycling that takes place under normal operation. Dave Steinberg points out here that electronic products experience tens of thousands of temperature cycles, even under benign conditions.

Changes over life — 4

- Parametric failures
  - generally associated with functional cores of the electronic components themselves
- Mechanical failure
  - can be more dramatic
  - many of the causes relate to the joints, or to assembly features of the internal component

Electronics KTN – Knowledge For Growth

Many permanent changes in a device, whether caused by oxidation or diffusion or induced by mechanical strain, can eventually result in parametric failure, depending on the tolerancing of design. Such failures are generally associated with the functional core of the electronic components themselves.

Mechanical failure, on the other hand, can be more dramatic, and many of the causes relate to the joints, or to assembly features of the internals of the components.

Changes over life — 4

- Parametric failures
  - generally associated with functional cores of the electronic components themselves
- Mechanical failure
  - can be more dramatic
  - many of the causes relate to the joints, or to assembly features of the internal component
- Eventually, joints will inevitably fail
  - this is actually unimportant!
  - what matters is that the expected number of cycles to failure should be well in excess of the design life of the product

Electronics KTN – Knowledge For Growth

Eventually, joints will inevitably fail, but this is actually unimportant! What matters is that the expected number of cycles to failure should be well in excess of the design life of the product.

Changes over life — 5

- Expected time-to-failure a complex function
  - the combination of materials used
  - design of components, assembly and joint itself
  - quality of the assembly process
  - any creep there is in the material
  - whether any stress raisers are present
  - the strain in the joint that results from the local temperature and temperature gradients

Electronics KTN – Knowledge For Growth

The expected time-to-failure is a complex function, involving:

- the combination of materials used
- the design of the components, the assembly and the joint itself
- the quality of the assembly process
- any creep that there is in the material
- whether any stress raisers are present, like cracks or scratches
- and, of course, the strain in the joint that results from the local temperature and temperature gradients.

It is the last of these, the strain in the joint, that is the most significant, and it’s also the one which the packaging engineer can most influence by appropriate thermal design.

The impact of heat

- What is heat?
- Thermal energy
- Temperature effects on the circuit
  - changes with temperature
  - changes over life
- Reliability implications
  - the parts we buy

Electronics KTN – Knowledge For Growth

So temperature can have effects on the circuit that are not reversible. This has implications for reliability, and also affects a number of choices that we make as designers. The first of which concerns the parts that we procure.

Maximum operating temperature

- Can we run hot?
  - usually some failure mechanism beyond a safe maximum

Electronics KTN – Knowledge For Growth

Typically we will be selecting parts on the basis of their maximum operating temperature, or conversely accepting the specification as given by the manufacturer of the part we need to use for electronic reasons.

Almost every one of the parts will have some kind of maximum operating temperature quoted. So can we run hot? Well, probably not, because there will usually be some failure mechanism beyond that safe maximum. This might be increased leakage which is an intrinsic consequence of the technology, as in the case of a tantalum capacitor, or it may be that the physical structure itself will start to deform, as with connectors with a thermoplastic body.

Maximum operating temperature

- Can we run hot?
  - usually some failure mechanism beyond a safe maximum
- The LED and power resistor examples
  - maximum temperature + de-rating curves

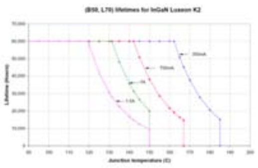

Philips Lumileds Lighting Company

Electronics KTN – Knowledge For Growth

We have already seen the case of the LED, where life and maximum temperature were related.

The rationale for taking action

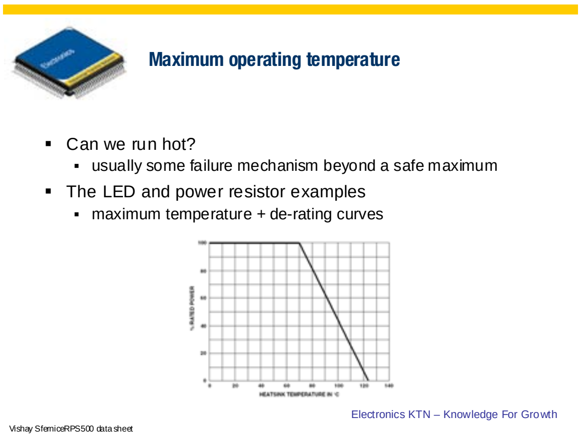

Another example is the power resistor. A data sheet for a power resistor will combine temperature and power information to give a de-rating curve that is based on the maximum allowable core temperature.

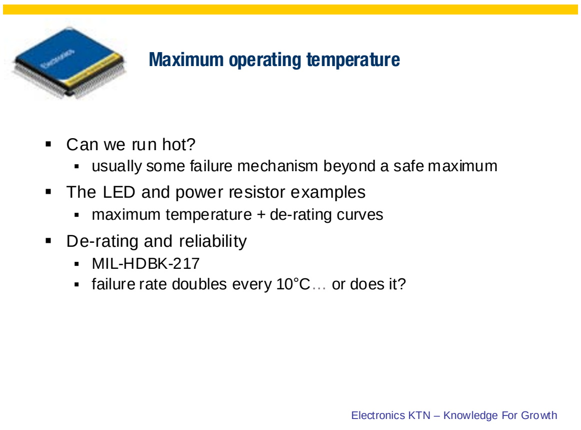

Applications requiring high reliability have for many years been based on the reasonable assumption that components stressed at less than their maximum rating will enjoy a longer life, and MIL-HDBK-217 is probably the best-known example of a repository of de-rating information that can be used for the life prediction of electronic assemblies. And there are equivalents in automotive and aerospace.

Embedded within them is the concept that the failure rate will double for each 10°C increase in temperature. However, whilst that rule of thumb works well for biological reactions, where it was first reported upon by Arrhenius, …

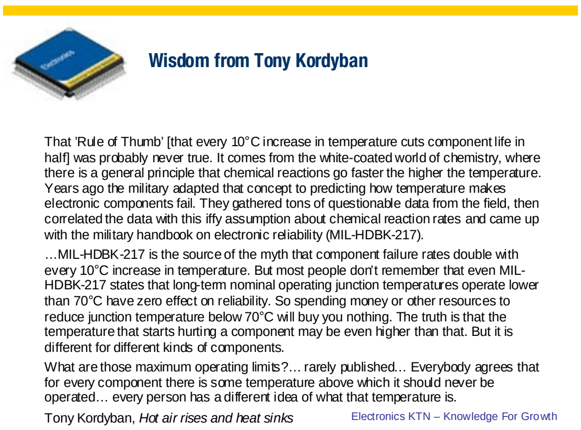

… as Tony Kordyban says, it's questionable whether the rule applies to electronic reliability. I would always recommend people to have this book on their shelves, or at least pause the presentation and read the words. Possibly the key point that Tony is making is that reducing the junction temperature below 70°C buys you nothing, and the real turning point for components, which will be different for different kinds of components, is likely to be higher than 70°C.

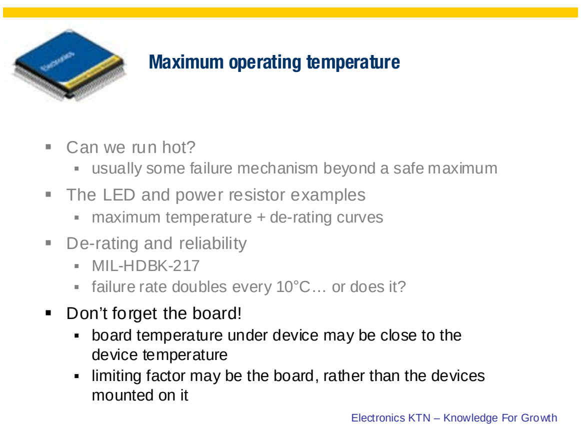

When it comes to maximum operating temperature, the tendency is to focus on components, but please don't forget the board! The trends towards surface mount components mean that the board is often used as a heat sink, and as a result the temperature under the device may be close to the silicon temperature. With semiconductors rated at 150°C or even 175°C, the limiting factor in the application may actually be the glass transition temperature of the board laminate, rather than the devices mounted on it. High T_g boards are common, but they're frequently associated with processing and reliability issues of other kinds.

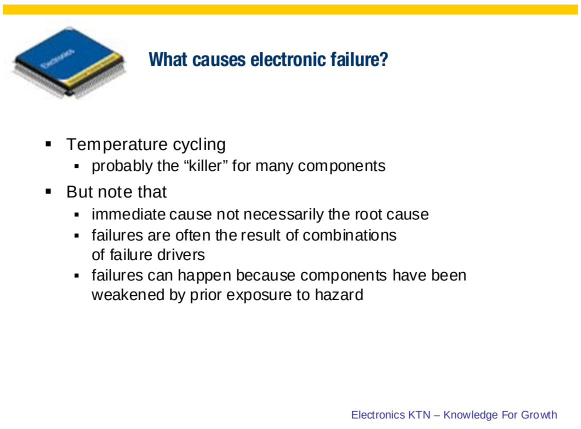

Whilst keeping maximum operating temperature in mind, we should also bear in mind that it is actually temperature cycling which is the killer for most components, with premature failure caused by the cumulative effect of hot and cold transitions and the associated strains.

It is also worth remembering that thermal cycling is unlikely to be the only cause, and failures are often the result of a combination of failure drivers. It is easy to point the finger at the straw that broke the camel's back, when the real effect was something different and earlier. The most obvious example here is of joints that have been weakened by temperature cycling and succumb to the shock of a drop test.

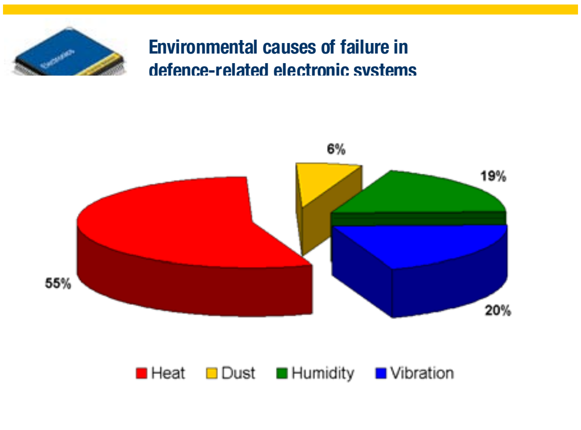

Even given the likelihood that most failures will have more than a single cause, it is interesting to see the degree to which thermal problems are perceived as being at the root of failure. This is 1990s work carried out as part of the US Avionics Integrity Program, and shows over half of electronic failures being caused by thermal issues. And, as we have already said, given large increases in power and reduction in size since then, we can't expect that this problem will have gone away.

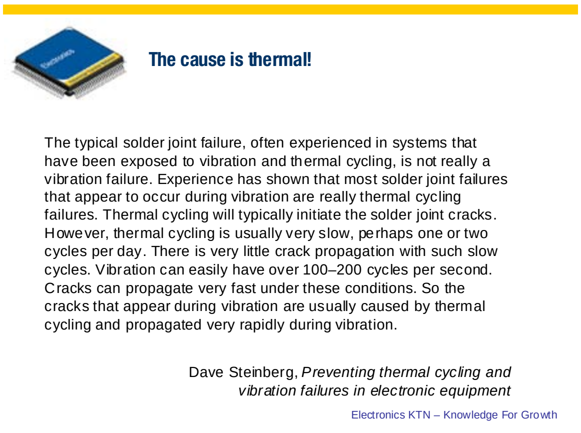

Dave Steinberg would certainly support this view. Dave has written a number of fairly approachable books on the subject of thermal and mechanical failure, and he makes the comment that most solder joint failures that appear to occur during vibration are really thermal cycling failures. In other words, the root cause of the problem is thermally-induced stress, not the way that we have put the structure together.

Thermal cycling is relatively slow, but it can initiate the problem which is then exposed either by vibration or during life events such as inadvertent drop test. So the immediate cause of your phone's failure might be the fact that you dropped it, but the root cause might be solder embrittlement. And this embrittlement might be a process or metallurgy problem, but may also be the result of thermal and power cycling during life.

The rationale for taking action

Cooler is more reliable?

"This is generally true, but not because many chip-level failure mechanisms are accelerated by absolute temperature.

"Within the manufacturer's specified operating temperature range, most of the reported failure mechanisms are not due to high steady-state temperatures. Rather, they depend on temperature gradients within the electronics assembly, the temperature cycle magnitude, and how fast the equipment heats up and cools down.

"Generally, reducing the operating temperature also reduces these accelerants, so in a general sense it is true to state that cooler is more reliable."

John Parry, Flomerics

Electronics KTN – Knowledge For Growth

Having listened to Tony Kordyban, of course we wouldn't try and cool products below 70°C, but we generally try and make things cooler under the belief that they will become more reliable in consequence.

John Parry supports this when he says that the cause of failure lies not so much in the absolute temperature but in the temperature gradients, the magnitude of the temperature cycle, and how fast the changes are. In other words, it is a question of severity rather than absolute temperature: the faster and the more uncontrolled the temperature excursion, the greater the chance of inducing failure. He makes the point though, that by clamping down the temperature you automatically clamp down those other factors also. It is not directly the case that cooler is necessarily more reliable, but it can have an impact.

The impact of heat

- What is heat?
- Thermal energy
- Temperature effects on the circuit
  - changes with temperature
  - changes over life
- Reliability implications
  - the parts we buy
  - the electronic and thermal design
  - the testing we carry out
- What happens if we ignore the challenge?

Electronics KTN – Knowledge For Growth

Thermal considerations have reliability implications, not only for the parts that we buy and for aspects of the electronic and thermal design, but also for the testing we carry out on the finished product. For example, we can use soak testing, that is high temperature endurance, as a reliability screen. Or we can temperature cycle, or power cycle, or carry out similar screening procedures either at a component or a system level to try and weed out early life failures. Of course, when doing this, you have to be careful that you don't actually reduce the life expectancy of the items that pass the screening!

But what happens if we ignore the thermal challenge. Our next slide is a very widely reported case…

The X-box story…

- The red ring of death…
  "…what happens when Microsoft goes too fast"

Electronics KTN – Knowledge For Growth

…which is the Xbox problem referred to in the media as "the red ring of death", which happened to an embarrassingly large number of Xbox360 games consoles when they were first introduced. When the green light turns red in this way, the system has failed, and may not recover when power is removed and the unit left to cool.

The rationale for taking action

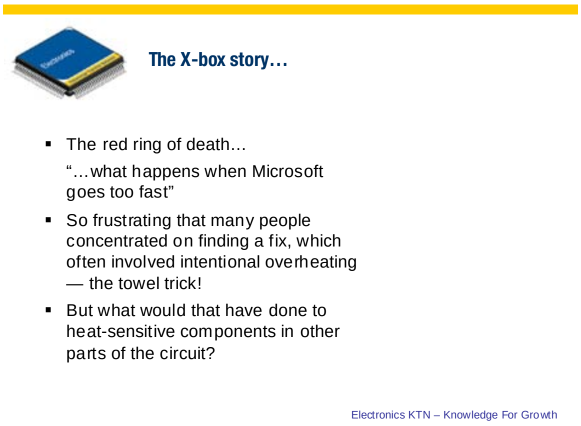

It was so frustrating that people looked for a fix, and some found that if you put a towel over the device and allowed it to "cook", on many occasions the problem went away. The supposition was that joint failure had occurred when the internal structure had been distorted, probably by heat, and that these joints could be re-melted.

One wonders what would have happened to heat-sensitive components in a situation where enough heat was being generated inside the case in order to re-melt the solder!

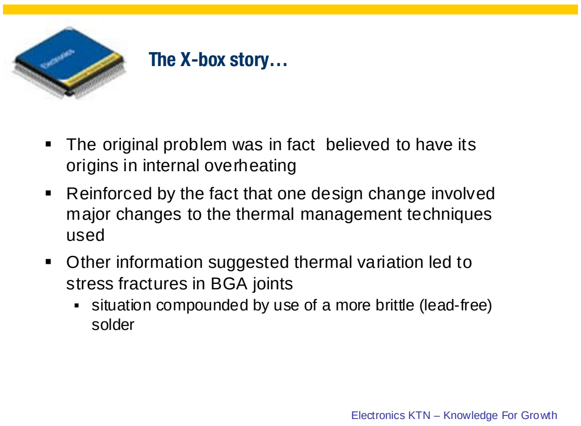

Now companies dealing with $1bn warranty problems are understandably reluctant to make detailed information available, and the story almost has an urban myth status. However, the problem was originally believed to have its origins in internal overheating, a supposition that was reinforced by the fact that one rectification change had involved major modifications to thermal management techniques.

Yet other information suggested that thermal cycling had led to stress fractures in BGA joints, a situation compounded by the use of a more brittle lead-free solder. In other words, thermal cycling was the trigger, but the root cause was the materials.

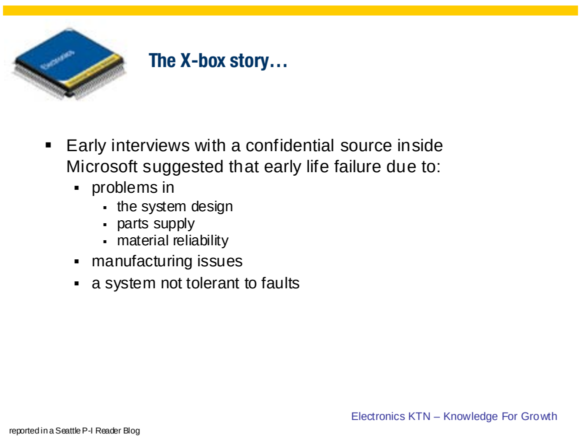

So what of the thermal problem? The immediate fix was certainly a thermal solution, to put in better cooling, although thermal problems weren't explicitly defined at that stage as the cause. One source suggested that there were actually a number of different problems, many due to other people, except perhaps the last point made, that the system wasn't tolerant to faults.

A year after the problem, and probably with the benefit of some hindsight, a conference speaker revealed a vital piece of information, which was that Microsoft had decided to take some design work in-house, with the result that the graphics driver could have been designed to dissipate 40% less power, but wasn't.

Which indicates that the original design was wrong, in part possibly because the thermal challenge wasn't dealt with formally, and in part because the situation was more challenging than might have been expected.

The rationale for taking action

There are certainly some learning points here, not the least of which is that from every perspective we should look to reduce dissipation as our primary aim, in the same way as in environmental engineering "Reduce" gets a higher priority than "Re-use" or "Re-cycle". And then, having kept the thermal problem to a minimum, we need to manage it properly.

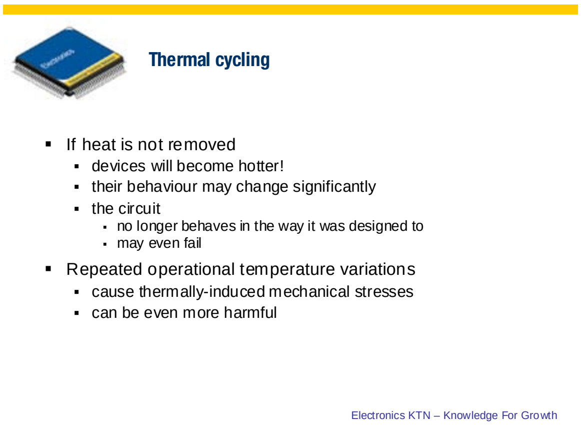

We've seen that, if heat is not removed, devices will become hotter, and their behaviour may change significantly, so that the circuit no longer behaves in the way it was designed to, and may even fail.

Also, whilst high temperatures themselves can be damaging, it is repeated operational temperature variations causing thermally-induced mechanical stresses that can be even more harmful.

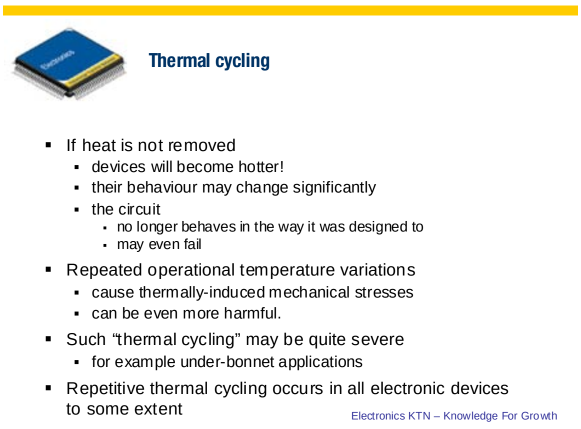

Such "thermal cycling" can be quite severe, even in everyday applications. For example, circuitry housed near the engine of a car experiences both the high temperatures developed by the engine when the car is being driven, and rapid cooling when the engine is turned off.

But repetitive thermal cycling occurs in all electronic devices to some extent. Even a desktop PC is stressed every time it is switched on for a period and then turned off.

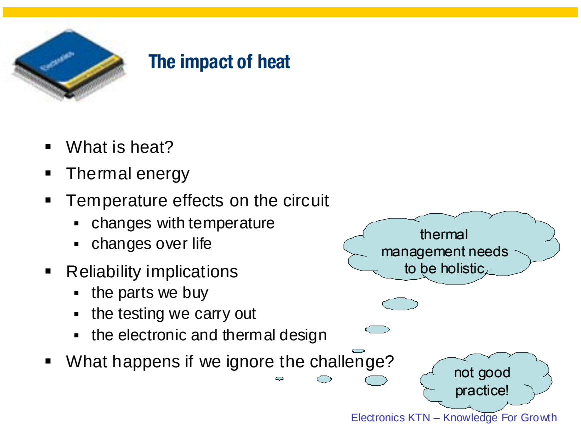

So we need to manage the impact that thermal energy will have on our product and on its performance and reliability and it's not good practice to ignore that challenge.

And we also need to keep in mind that thermal management has to be holistic, and to take account of every aspect of design, manufacture and test, including areas such as procurement, in order to make sure that our solution is both technically and commercially fit for purpose.