PCB Thermal Simulation - The State of the Art

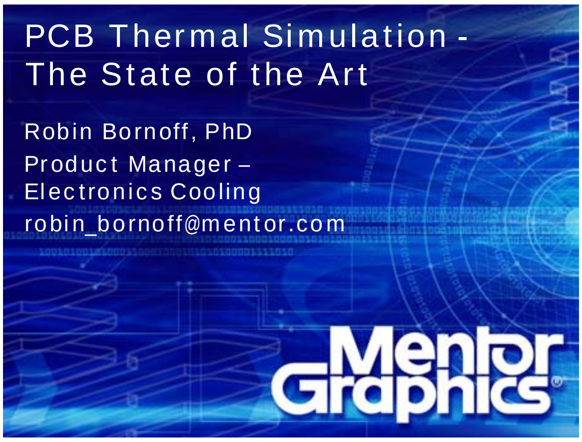

PCB Thermal Simulation - The State of the Art

Robin Bornoff, PhD
Product Manager – Electronics Cooling
robin_bornoff@mentor.com

Okay, good morning or afternoon everybody, thank you very much for attending today's webinar. The title of the webinar today is PCB thermal simulation, the state of the art.

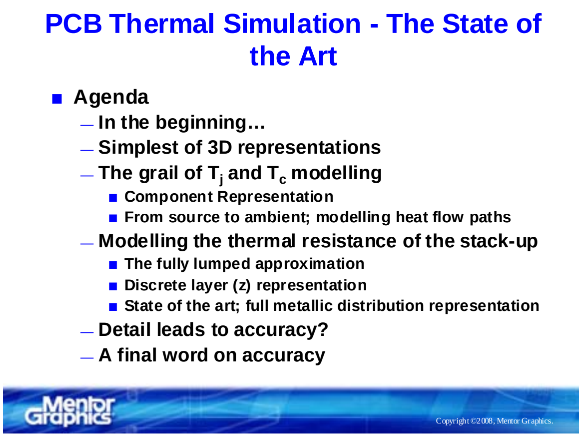

PCB Thermal Simulation - The State of the Art

- Agenda
  - In the beginning…
  - Simplest of 3D representations
  - The grail of T_j and T_c modelling
    - Component Representation
    - From source to ambient; modelling heat flow paths
  - Modelling the thermal resistance of the stack-up
    - The fully lumped approximation
    - Discrete layer (z) representation
    - State of the art; full metallic distribution representation
  - Detail leads to accuracy?
  - A final word on accuracy

Okay, in terms of today's agenda, I'll start off by talking a little bit about some of the first models, CSD models, containing representations of PCBs, that were being constructed a couple of decades ago now. I'll then move up and talk about the tentative steps that were taken to create 3D representations of components and PCBs.

I'll then deviate slightly and talk about the grail of junction case temperature modelling, very much associated with the representation of PCBs. To do that, I will outline various component representation options and methodologies and then talk a little bit about modelling heat flow paths from the heat source right up to the ambient.

The main bulk of the presentation will be focused on the theory behind all the modelling methodologies available, in terms of thermal resistance representation of the PCB stack-up, everything from a single block fully lumped approximation right up to the forthcoming state of the art capability, which will include a full metallic distribution representation of a PCB.

I'll then talk a little bit about the relationship between PCB model detail and its relationship to accuracy, then finally I'll focus a little more on the concept of accuracy, setting expectations as to what you should and shouldn't expect with regard to predictive accuracy of CFD models containing PCBs.

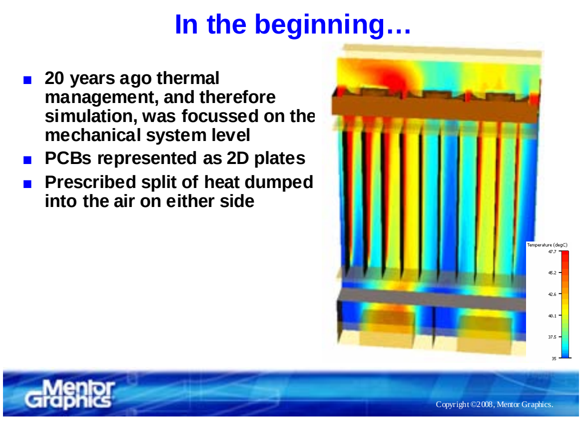

In the beginning…

- 20 years ago thermal management, and therefore simulation, was focussed on the mechanical system level
- PCBs represented as 2D plates
- Prescribed split of heat dumped into the air on either side

Okay, in the beginning. 20 years ago, which was the point at which Flotherm was released, the first CFD tool focused specifically on electromagnetic thermal predictions, thermal management, and therefore the capabilities of the simulation tool, was focused on the mechanical system level. In terms of PCB representation – the example I'm showing here is a compact model of a sub-rack – PCBs are modelled very simply as simple, two-dimensional plates, sufficient just to deflect and channel air flow.

In terms of the thermal effects that PCBs had on the air, such initial models had a prescribed split of heat that was dumped directly into the air on either side of the PCB.

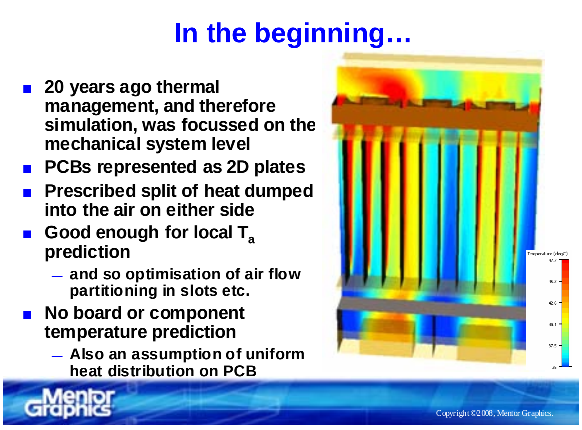

In the beginning…

- 20 years ago thermal management, and therefore simulation, was focussed on the mechanical system level
- PCBs represented as 2D plates
- Prescribed split of heat dumped into the air on either side
- Good enough for local T_a prediction
  - and so optimisation of air flow partitioning in slots etc.
- No board or component temperature prediction
  - Also an assumption of uniform heat distribution on PCB

This was good enough for local ambient air temperature prediction, and that was the intention of such a model. Having a prediction of the air flow distribution and the temperature distribution within that air, allowed design to investigate effects such as fan selection, fan placement, venting configuration optimisation of air flow partitioning within slots, for example, so a sort of air flow management to achieve thermal management at the mechanical system level.

In these initial models there was no direct prediction of board or component temperature and there was also an assumption of a uniform heat distribution across the surface of the PCB. Despite the fact that these initial models were very simple, they did provide value and we still see examples of such a modelling level being used even today.

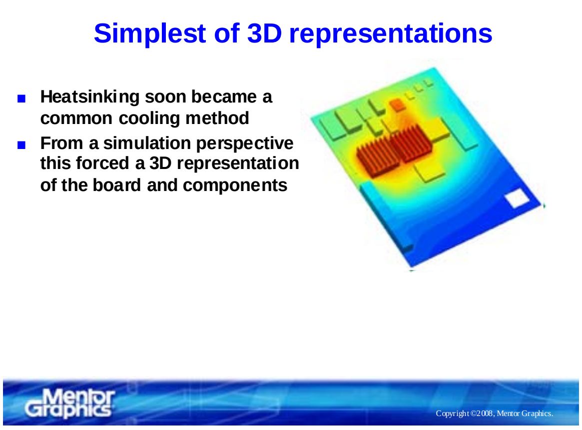

Simplest of 3D representations

- Heatsinking soon became a common cooling method
- From a simulation perspective this forced a 3D representation of the board and components

The simplest of three-dimensional representations, although there was never a single cause motivating people to go for a more detailed three-dimensional representation PCB, we believe that it was the advent of heat sinking becoming a very common cooling methodology that motivated a more accurate three-dimensional representation of the PCB and also the layout of components on the board and of course the heat sink sitting on top of those components.

From a simulation perspective, to be able to investigate the effectiveness of heat sinks, of course, an accurate three-dimensional prediction of the air flow in and around the heat sinks was required, and it was this that motivated three-dimensional representations to ensure that the three-dimensional air flow field was predicted accurately.

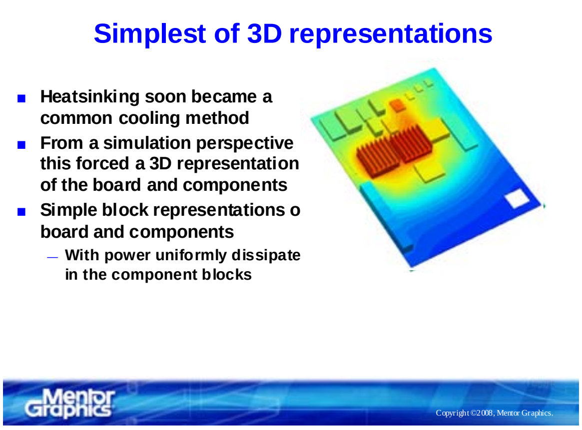

Simplest of 3D representations

- Heatsinking soon became a common cooling method
- From a simulation perspective this forced a 3D representation of the board and components
- Simple block representations o board and components
  - With power uniformly dissipate in the component blocks

In terms of the level of detail, at this stage the PCB was represented as a single cuboid, the components were represented as simple blocks and the heat sinks were defined as simple blocks as well. In terms of the thermal prediction, heat was assumed to be dumped uniformly in the single blocks representing the components themselves.

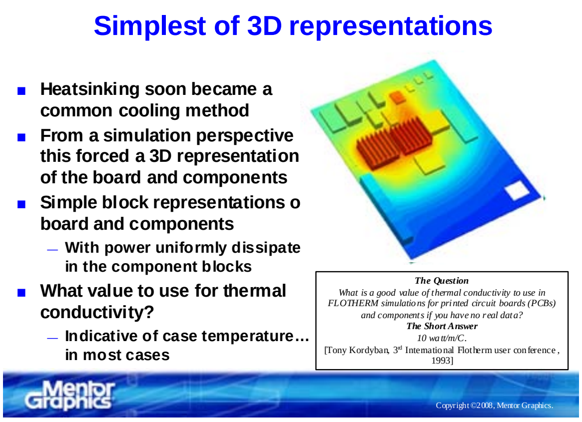

Simplest of 3D representations

- Heatsinking soon became a common cooling method
- From a simulation perspective this forced a 3D representation of the board and components
- Simple block representations o board and components
    - With power uniformly dissipate in the component blocks
- What value to use for thermal conductivity?
    - Indicative of case temperature… in most cases

The Question
What is a good value of thermal conductivity to use in FLOTHERM simulations for printed circuit boards (PCBs) and components if you have no real data?
The Short Answer
10 watt/m/C.
[Tony Kordyban, 3rd International Flotherm user conference, 1993]

Now, inherent in any CFD simulation that includes three-dimensional solid objects, to obtain an prediction of the temperature within those solid objects, the primary input to the CFD model is a required value of thermal conductivity for the individual cuboids. So one of the first questions is: what value of thermal conductivity should I use for such simple representations of PCBs? Now, this was exactly the question posed by a seminal paper published by Tony Kordyban at the third International Flotherm User Conference way back in 1993. He posed the question: What is a good value of thermal conductivity to use in flow simulations to PCBs and components if you have no real data? – or maybe, at that stage, no real experience as to what conductivity you should use. He summarised that, the short answer being 10W per metre Kelvin, which was obviously better than no value at all and has proved its worth time and time again.

In terms of setting expectations as to what you can expect by using this basic approach with a single isotropic thermal conductivity value, you will obtain temperatures indicative of component case temperatures, no better, and then not in all cases, but in most cases. So a very simple model, a very simple prescription of thermal conductivity, and an indication roughly as to what the case temperature is but good enough for heat sink design and optimisation in situ.

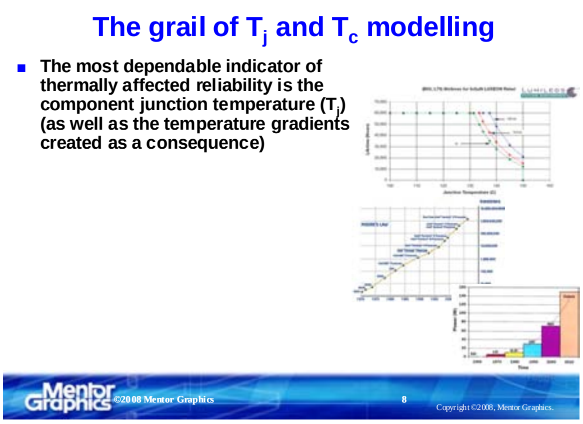

The grail of T_j and T_c modelling

- The most dependable indicator of thermally affected reliability is the component junction temperature (T_j) (as well as the temperature gradients created as a consequence)

Quite quickly it became the focus of the electronics cooling community to focus on the goal or the grail of accurate junction and case temperature modelling. The motivation for this is clear and well known. The most dependable indicator of thermally affected reliability is the component junction temperature, which is not explicitly just the junction temperature itself, it's the temperature gradients that are created as a consequence, both gradients in time and also gradients in space as well. These are the main driving factors that result in thermally mechanically induced failure mechanisms.

A little chart I have here just shows a Lumileds LED, a Luxeon Rebel LED, and off the spec sheet a graph showing the relationship between light time in hours and the junction temperature, where if maximum temperatures are exceeded, there is a noticeable drop in the expected lifetime. Such relationships, or similar relationships, exist for all powered electronic parts.

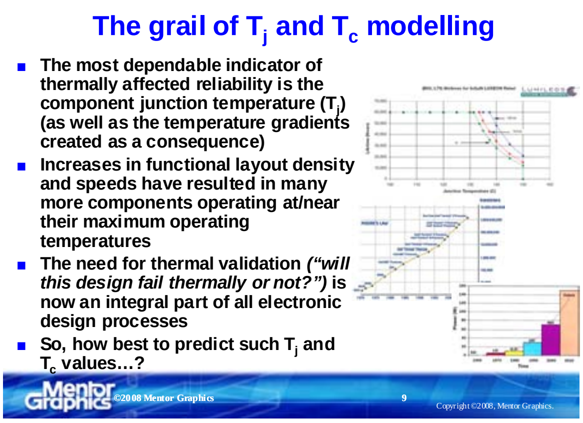

The grail of T_j and T_c modelling

- The most dependable indicator of thermally affected reliability is the component junction temperature (T_j) (as well as the temperature gradients created as a consequence)
- Increases in functional layout density and speeds have resulted in many more components operating at/near their maximum operating temperatures
- The need for thermal validation *("will this design fail thermally or not?")* is now an integral part of all electronic design processes
- So, how best to predict such T_j and T_c values…?

In addition to this, over the years – and a trend that continues – is increases in functional layout density, packing more actives onto a PCB, the drive for electronic miniaturisation forcing this as well, in conjunction with processing speeds increasing, have resulted in many more components operating at or near their maximum operating temperatures. The reason for this is that the thermal density, the watts per inch squared on PCBs, is reaching such a level so as to push many, many more components right up to their maximum operating temperature.

As a consequence of this, the need for thermal validation (i.e. will this design fail due to thermal effects or not?) is now an integral part of all electronic design processes, or maybe I should say all good electronic design processes.

The main focus, I think, of the rest of the presentation is really centred around how best to predict such junction and case temperature values.

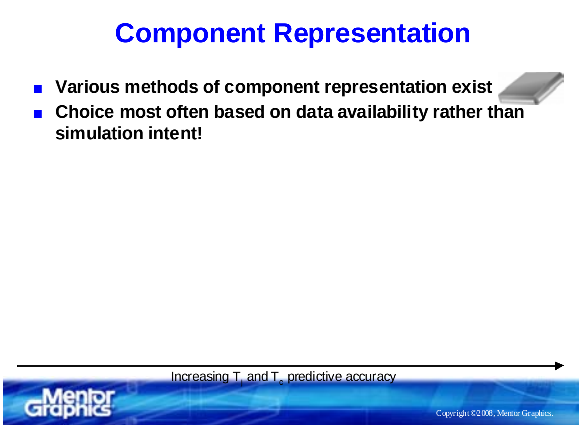

Component Representation

- Various methods of component representation exist
- Choice most often based on data availability rather than simulation intent!

Increasing T_j and T_c predictive accuracy

To answer this, it's important to focus firstly on various options for component representation. There are various methods of component representation, various options you have available to you. The choice of component representation is more often based on data availability rather than your simulation intent of trying to obtain most accurate data.

I have here a little figure that shows various component representation methods and their relationship to increasing junction and case temperature predictive accuracy.

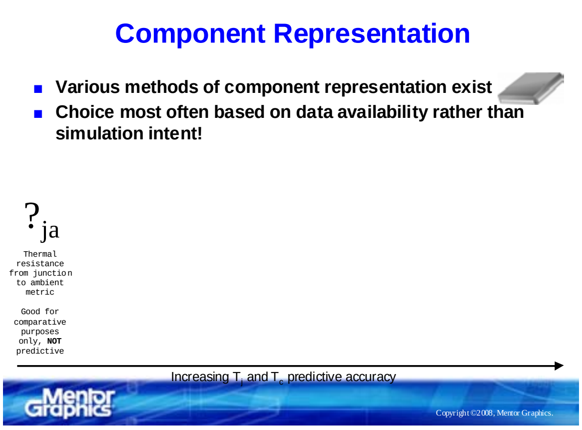

Component Representation

- Various methods of component representation exist
- Choice most often based on data availability rather than simulation intent!

$?_{ja}$

Thermal resistance from junction to ambient metric

Good for comparative purposes only, **NOT** predictive

Increasing T_j and T_c predictive accuracy

The first is not really a model itself, it's a metric, Θ_{JA}, thermal resistance between junction and the ambient, so this is a metric that is nearly always included on component spec sheets, despite the fact that, in theory, if you know the thermal resistance between junction and ambient, and you know what the ambient temperature is going to be and you know what the component power dissipation is, then in theory you could back out what the expected junction temperature of the component is. However, as JEDEC indicated in their description of such a metric, the intention of this metric is used for comparative purposes when selecting components, it should not be used for predictive purposes. So, despite the attractiveness of its simplicity, it should not be used for predictive purposes.

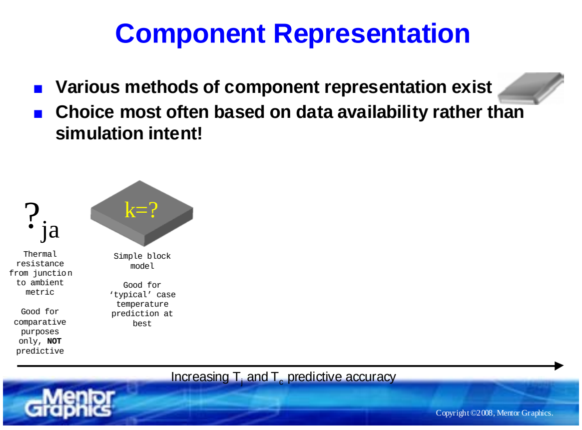

Component Representation

- Various methods of component representation exist
- Choice most often based on data availability rather than simulation intent!

$?_{ja}$

Thermal resistance from junction to ambient metric

Good for comparative purposes only, **NOT** predictive

k=?

Simple block model

Good for 'typical' case temperature prediction at best

Increasing T_j and T_c predictive accuracy

The next level up, as alluded to already, is the simplest of all three-dimensional representations: a single block representing the component, whereby a value of thermal conductivity is required to be defined. Based upon Tony Kordyban's work: initial value of 10 is better than nothing. In recent releases of our software we've come up with a library of more refined thermal conductivity values, appropriate for different package styles, but even so, expectations should be that the most you can expect from such a representation is predictions of typical case temperatures at best, certainly no more than that.

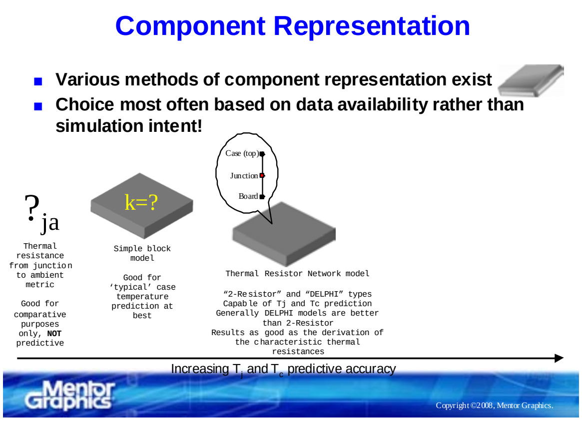

Component Representation

- Various methods of component representation exist
- Choice most often based on data availability rather than simulation intent!

Thermal resistance from junction to ambient metric

Good for comparative purposes only, **NOT** predictive

Simple block model

Good for 'typical' case temperature prediction at best

Thermal Resistor Network model

"2-Resistor" and "DELPHI" types
Capable of Tj and Tc prediction
Generally DELPHI models are better than 2-Resistor
Results as good as the derivation of the characteristic thermal resistances

Increasing T_j and T_c predictive accuracy

Copyright ©2008, Mentor Graphics.

The next two component representations are both capable of predicting junction and case temperatures. The first is what is known as a thermal resistor network model representation of a component. I'll talk a little more about the concept of thermal resistor networks and how they can be used in various modelling methodologies for the PCB itself, but as far as component definition is concerned, there are two types of thermal resistor network, there's a two-resistor network and a DELPHI type of network.

These networks are sort of abstract approximations of the conduction paths the heat experiences as it is dissipated in the junction in the silicon within the package and conducts its way to either different points within or areas on the periphery of the package itself. As far as the two-resistor model is concerned, this is characterised by two resistances, Θ_{JC} and Θ_{JB}, sometimes available from component manufacturers.

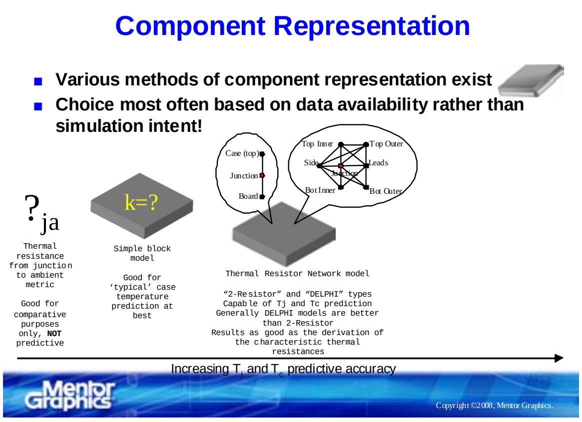

Component Representation

- Various methods of component representation exist
- Choice most often based on data availability rather than simulation intent!

Thermal resistance from junction to ambient metric

Good for comparative purposes only, **NOT** predictive

Simple block model

Good for 'typical' case temperature prediction at best

Thermal Resistor Network model

"2-Resistor" and "DELPHI" types
Capable of Tj and Tc prediction
Generally DELPHI models are better than 2-Resistor
Results as good as the derivation of the characteristic thermal resistances

Increasing T_j and T_c predictive accuracy

Copyright ©2008, Mentor Graphics.

The DELPHI type of resistor network has a more advanced topology that in theory is capable of representing the more advanced conduction paths the heat experiences as it makes its way through more complicated packages. Generally it's a rule that DELPHI models are going to be better in terms of accuracy than the two-resistor models. I have a note here that the results are as good as a derivation of the characteristic thermal resistances.

Sometimes Θ_{JC} and Θ_{JB} are available in spec sheets, as I said; sometimes you're lucky enough that your component suppliers will supply such information. We released a tool many years ago called FloPack that was developed specifically to be able to create such representations from parametric definitions of various components.

The real advantage of such a representation of a component is, yes, it's capable of predicting junction and case temperature but it also does not contain any proprietary information about the construction of the package and as such, in theory, it should be a preferable method of characterising and disseminating thermal models of components from suppliers to end users.

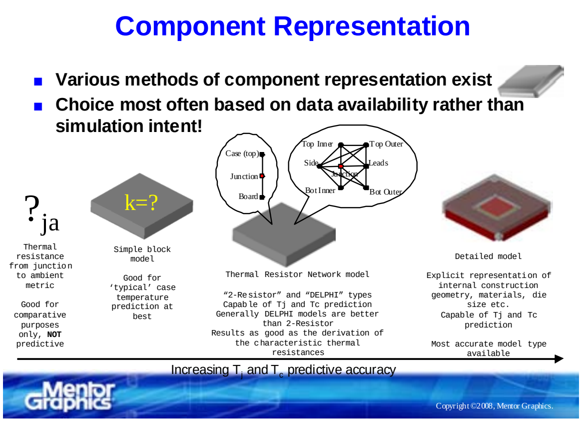

Component Representation

- Various methods of component representation exist
- Choice most often based on data availability rather than simulation intent!

Thermal resistance from junction to ambient metric

Good for comparative purposes only, **NOT** predictive

Simple block model

Good for 'typical' case temperature prediction at best

Thermal Resistor Network model

"2-Resistor" and "DELPHI" types
Capable of Tj and Tc prediction
Generally DELPHI models are better than 2-Resistor
Results as good as the derivation of the characteristic thermal resistances

Detailed model

Explicit representation of internal construction geometry, materials, die size etc.
Capable of Tj and Tc prediction

Most accurate model type available

Increasing T_j and T_c predictive accuracy

Copyright ©2008, Mentor Graphics.

Finally, the most accurate type of representation is a detailed model. By "detailed", I mean it is an explicit representation of the internal construction geometry, including the materials, the die size, various other constituents of the package itself. This is not only capable of predicting junction and case temperature values: by its very nature it will enable observation of temperatures throughout the package itself, and it's the most accurate model type available.

Again, some component suppliers will supply such detailed models. Alternatively, FloPack as a tool can be used to generate such models itself.

Accuracy comes with a price – having to specify explicitly the internal construction of the package. By its very nature there has to be a definition of the die size and other proprietary type information, which is a downside from the component supplier's point of view, of such model definition types.

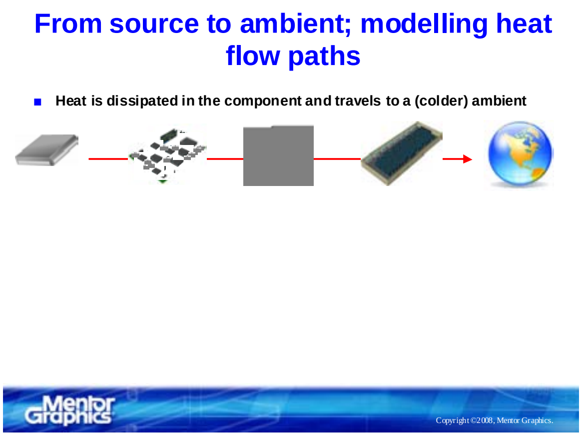

From source to ambient; modelling heat flow paths

- Heat is dissipated in the component and travels to a (colder) ambient

From source to ambient, modelling heat flow paths. I think it's important to appreciate this concept. Any heat that is dissipated in the component has to travel or travels naturally to a colder ambient. In extreme cases, heat dissipated in a component makes its way, through various tortuous routes, to heat up the atmosphere around our planet. It does so by passing through various stages. The first stage, obviously, is heat passing from the component to its local environment, which is PCB and the air around the PCB. Beyond that, the heat is managed to be extracted from the PCB to the local environment in which that sits, which is the chassis, the sub rack, the handheld units, or whatever.

From that electronic product, that is normally placed in some sort of built environment – in this case, for example, a data centre – and heat is extracted at that level and then passed on to heat up the atmosphere.

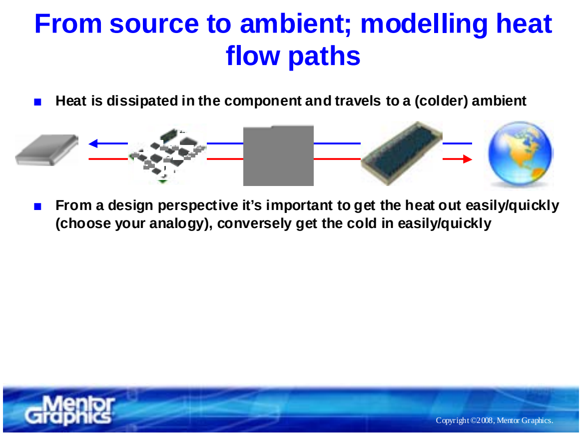

From source to ambient; modelling heat flow paths

- Heat is dissipated in the component and travels to a (colder) ambient
- From a design perspective it's important to get the heat out easily/quickly (choose your analogy), conversely get the cold in easily/quickly

So from a design perspective, it's important, as a goal, to get the heat out from the package to the ambient as easily and quickly as possible, depending on which analogy you like to use.

Conversely, another way of looking at it, which is just as useful, I think, is to make it as easy as possible in your design to get the cold from the ambient down to where the heat is being dissipated to quench that heat dissipation and to reduce the resulting source temperature.

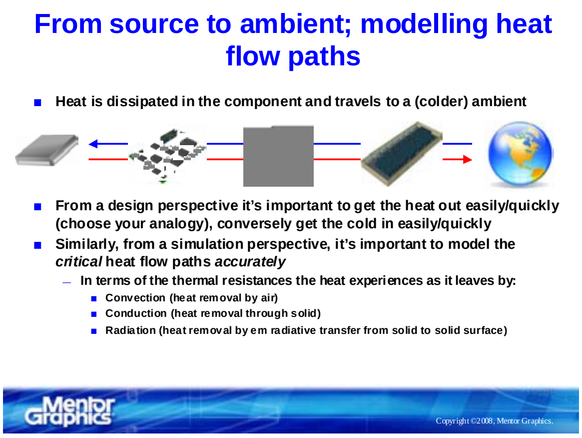

From source to ambient; modelling heat flow paths

- Heat is dissipated in the component and travels to a (colder) ambient
- From a design perspective it's important to get the heat out easily/quickly (choose your analogy), conversely get the cold in easily/quickly
- Similarly, from a simulation perspective, it's important to model the *critical* heat flow paths *accurately*
  - In terms of the thermal resistances the heat experiences as it leaves by:
    - Convection (heat removal by air)
    - Conduction (heat removal through solid)
    - Radiation (heat removal by em radiative transfer from solid to solid surface)

Similarly, from a simulation perspective, in terms of the need or the desire for accurate prediction, it's important to model the critical heat flow paths accurately. These are heat flow paths where the heat experiences, the majority of the heat experiences, as it passes from the package, through the various levels, out.

Then this is in terms, as I said, of the thermal resistance the heat experiences as it leaves by the three classical modes of heat transfer, so heat convected away, so it's heat removal by air; by conduction, which is heat removal from solid; and through radiation, which is heat removal by electromagnetic radiative transfer from solid to solid surfaces.

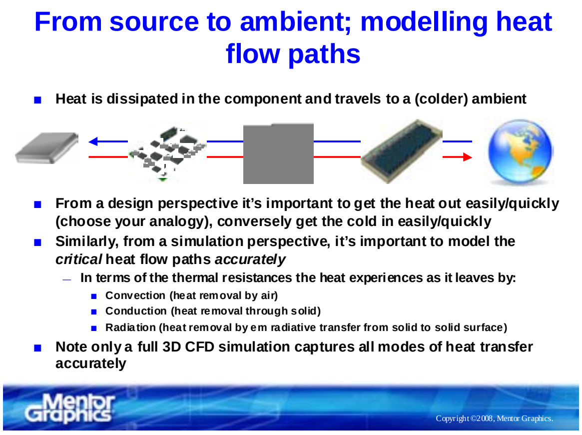

From source to ambient; modelling heat flow paths

- Heat is dissipated in the component and travels to a (colder) ambient
- From a design perspective it's important to get the heat out easily/quickly (choose your analogy), conversely get the cold in easily/quickly
- Similarly, from a simulation perspective, it's important to model the *critical* heat flow paths *accurately*
  - In terms of the thermal resistances the heat experiences as it leaves by:
    - Convection (heat removal by air)
    - Conduction (heat removal through solid)
    - Radiation (heat removal by em radiative transfer from solid to solid surface)
- Note only a full 3D CFD simulation captures all modes of heat transfer accurately

I have a note in here: only a full 3D CFD simulation will capture all modes of heat transfer accurately. So a CFD simulation that has explicit representation or accommodation of convective, conductive and radiative heat transfer approaches. There are other simulation tools that are attractive in terms of their speed that have conduction-only simulations and prescribe the other modes of heat transfer, specifically prescribe convective losses by specified heat transfer coefficients and local ambient temperatures. As I said, the advantage is they're very fast but great care has to be taken that you don't prescribe incorrectly the other modes of heat transfer that the simulation tool is not accommodating. Not a problem with a CFD simulation.

Modelling the thermal resistance of the stack-up

- Arguably the most critical thermal resistance the heat experiences (outside of the package) is the conductive resistance of the PCB itself
- Various methods exist for the representation of the PCB stack-up

Let's focus now on modelling the thermal resistance of the PCB stack-up itself.

Arguably the most critical thermal resistance the heat experiences, that is outside of the package, is the conductive resistance of the PCB itself. I say "arguably", I mean, in theory the thermal resistance of the heat as it passes through into the air through the boundary layer is also a dominant resistance in the full heat flow path. For the sake of this presentation, though, I will focus on the thermal resistive effects, the conductive resistive effects of the PCB substrate underneath the component.

In terms of the PCB representation, again various methods exist. Just for component representation there are various options, there again are various modelling methodology options for the PCB stack-up.

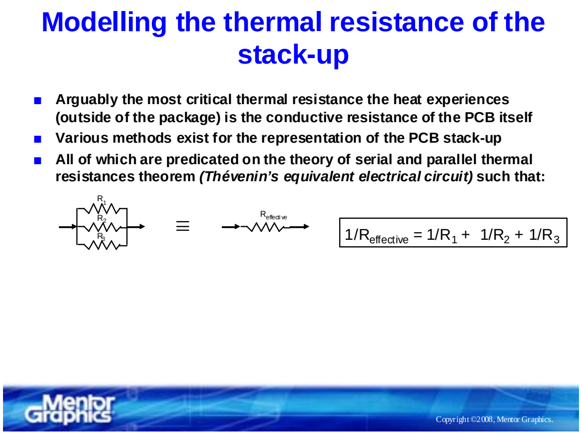

Modelling the thermal resistance of the stack-up

- Arguably the most critical thermal resistance the heat experiences (outside of the package) is the conductive resistance of the PCB itself
- Various methods exist for the representation of the PCB stack-up
- All of which are predicated on the theory of serial and parallel thermal resistances theorem *(Thévenin's equivalent electrical circuit)* such that:

$1/R_{effective} = 1/R_1 + 1/R_2 + 1/R_3$

All the methods of PCB representation are predicated on the theory of serial and parallel thermal resistance theorem and there is a direct analogy here to the electrical equivalent, which is a Thévenin's equivalent electrical circuit theorem. Such that, if heat passes in parallel through three resistances – when I say three resistances, think of it in terms of three different types of material, three different types of solid material – then this can be collapsed or sort of composited down to passing through a single effective thermal resistance, where the effective thermal resistance is related to the individual thermal resistances using the equation that you see on the screen.

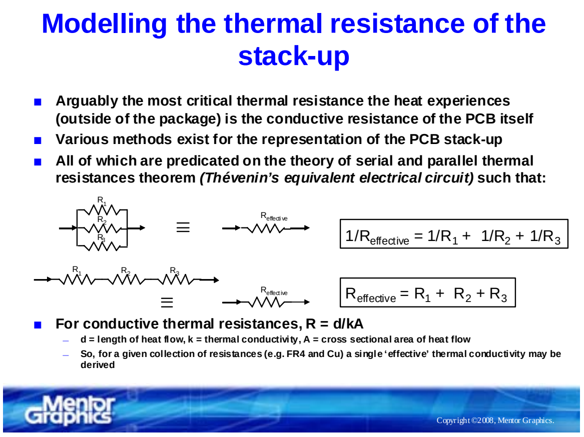

Modelling the thermal resistance of the stack-up

- **Arguably the most critical thermal resistance the heat experiences (outside of the package) is the conductive resistance of the PCB itself**
- **Various methods exist for the representation of the PCB stack-up**
- **All of which are predicated on the theory of serial and parallel thermal resistances theorem *(Thévenin's equivalent electrical circuit)* such that:**

$1/R_{effective} = 1/R_1 + 1/R_2 + 1/R_3$

$R_{effective} = R_1 + R_2 + R_3$

- **For conductive thermal resistances, R = d/kA**
  - d = length of heat flow, k = thermal conductivity, A = cross sectional area of heat flow
  - So, for a given collection of resistances (e.g. FR4 and Cu) a single 'effective' thermal conductivity may be derived

This side of this is heat that passes through thermal resistances or different materials sequentially in theories. These too can be represented as a single effective thermal resistance, which is simply the sum of the individual thermal resistances of the three different materials.

For conductive thermal resistances, the thermal resistance is defined by ∂ over kA, where ∂is the length of heat flow – it's the length of the solid through which the heat is flowing, in the flow direction – k is the thermal conductivity of that material and A is the cross-sectional area normal to the direction of heat flow.

So for a given collection of resistances – and as far as PCBs are concerned we nearly always talk about FR-4 and copper – a single effective thermal conductivity may be derived from the single effective thermal resistance, which itself is derived from the contribution of the various materials that constitute the PCB. As mentioned before, it is thermal conductivity that is the primary input to a solid object representation in a CFD simulation that includes conduction.

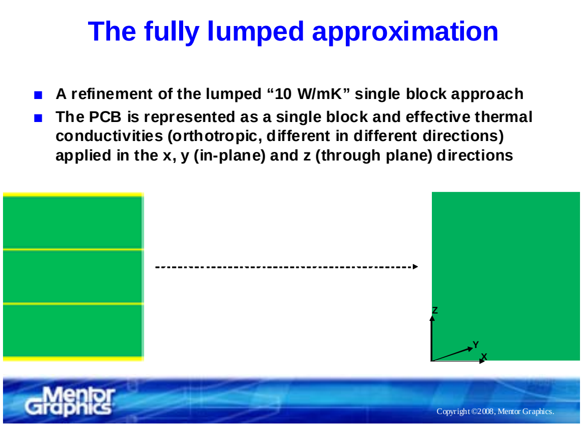

The fully lumped approximation

- **A refinement of the lumped "10 W/mK" single block approach**
- **The PCB is represented as a single block and effective thermal conductivities (orthotropic, different in different directions) applied in the x, y (in-plane) and z (through plane) directions**

Based upon this theory, let's have a look at various approximations of a representation of PCBs. Despite the fact that we're talking about some mathematical theory, suffice it to note that Flotherm and other associated CFD tools automate to various degrees the calculation of these conductivities, using the theories that are being described, so the following slides just talk about the methodology, they don't describe what you have to do to input data into the tool itself.

The first fully lumped approximation, this is a refinement of the lumped isotropic 10W/m·K single block approach. The PCB is still represented as a single block, but effective thermal conductivities, multiple conductivities, so orthotropic, being different in different directions, are applied in the X/Y in-plane directions and the Z, the through-plane directions, of the single block representation.

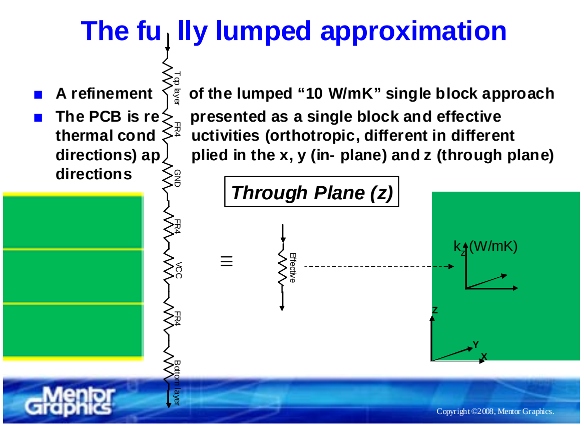

The fu lly lumped approximation

- **A refinement of the lumped "10 W/mK" single block approach**
- **The PCB is re presented as a single block and effective thermal cond uctivities (orthotropic, different in different directions) ap plied in the x, y (in- plane) and z (through plane) directions**

Through Plane (z)

If we look at the through plane thermal resistance of the PCB – and it's got a long resistance here – any heat passing vertically up and down the PCB in the Z direction is going to experience thermal resistances due to sequentially, top layer, bit of FR4 dielectric to the ground, FR4 power, FR4, and then the bottom layer, in this 2S2P basic board example.

This is made equivalent to a single effective thermal resistance, and from that single effective thermal resistance an effective conductivity in the Z direction can be derived and applied in the Z direction to the single block representation.

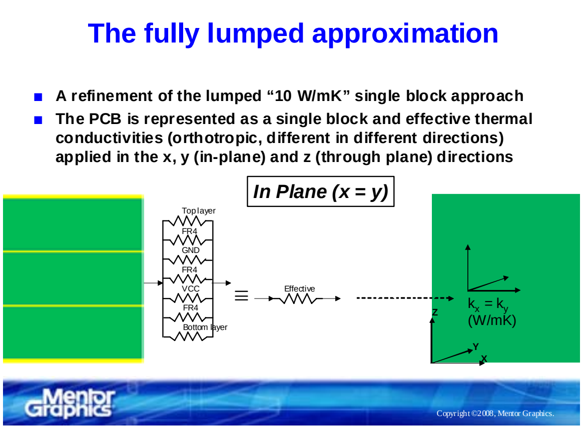

The fully lumped approximation

- A refinement of the lumped "10 W/mK" single block approach
- The PCB is represented as a single block and effective thermal conductivities (orthotropic, different in different directions) applied in the x, y (in-plane) and z (through plane) directions

In Plane (x = y)

Similarly, if you consider heat moving in the plane of the board, then it passes through all seven different thermal resistances in parallel, again whereupon an effective thermal resistance can be calculated, and then effective thermal conductivity set the same for the X and Y direction. This is because whether heat travels in the X direction or the Y direction, the thermal resistor parallel topology remains the same, so it's orthotropic but conductivities in X and Y are the same and different in Z.

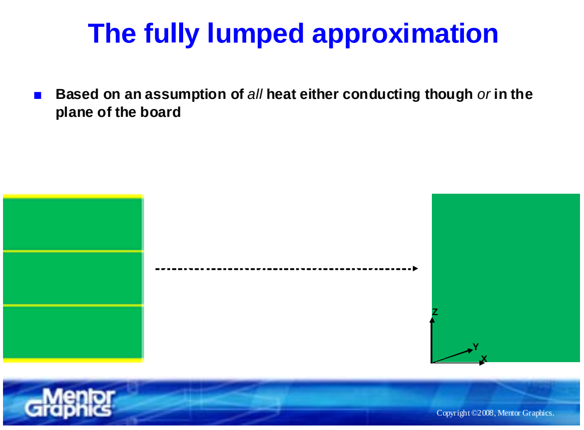

The fully lumped approximation

- Based on an assumption of *all* heat either conducting though *or* in the plane of the board

I think it's very important to be aware of the assumptions on which those effective orthotropic conductivities were derived. Basic assumption that either heat conducts through or in the plane of the board, so if all heat does do that, then the orthotropic conductivity approximation in the single block is going to be relatively accurate.

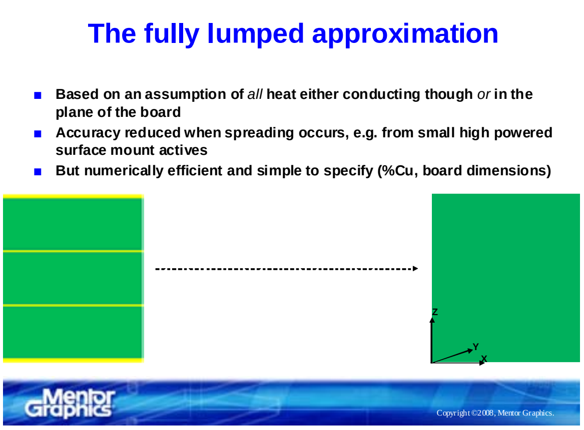

The fully lumped approximation

- Based on an assumption of *all* heat either conducting though *or* in the plane of the board
- Accuracy reduced when spreading occurs, e.g. from small high powered surface mount actives
- But numerically efficient and simple to specify (%Cu, board dimensions)

Unfortunately heat has a tendency, especially if released from a relatively small area, to spread, not just down and sideways but a combination of the two. So, if spreading effects are important, then the rather simple, single block orthotropic conductivity approach will not accurately resolve that heat spreading resistance. If, however, all of your heat dissipation is uniformly distributed, on either the top or bottom surface, then this single block orthotropic approach produces some usefully accurate results.

The main advantage, however, is the fact that such a single block representation is numerically very efficient to represent and simple to specify: you simply need to know the percentage copper in the board and the board dimensions.

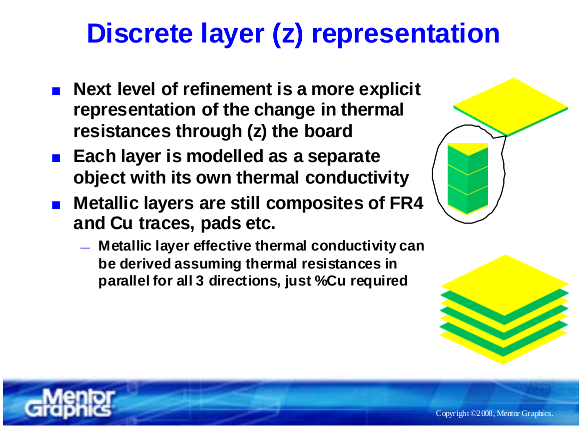

Discrete layer (z) representation

- Next level of refinement is a more explicit representation of the change in thermal resistances through (z) the board
- Each layer is modelled as a separate object with its own thermal conductivity
- Metallic layers are still composites of FR4 and Cu traces, pads etc.
  - Metallic layer effective thermal conductivity can be derived assuming thermal resistances in parallel for all 3 directions, just %Cu required

The next level of detail, if you like, is a discrete layer representation in the Z direction. This intends to represent the change in thermal resistance as heat passes through the Z direction of the board. Here the PCB model is represented using individual cuboids with their own materials to represent each individual layer within the PCB.

Suffice it to note that, even for the metallic layers – the metallic layers shown in yellow, here – they themselves in fact are still composites of a certain amount of FR4 and copper, in terms of the traces and pads, et cetera. The approach most often used here is, in terms of calculating the thermal conductivity for the metallic layers, you assume that the copper and FR4 is distributed on those metallic layers such that heat will pass through them in parallel, regardless of which direction the heat flows through that metallic layer, and though not indicated here, you could make a similar argument for the dielectric layers in terms of their containing thermal or electrical vias.

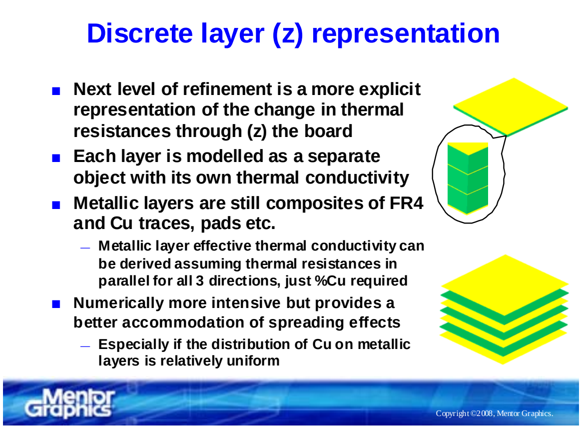

Discrete layer (z) representation

- Next level of refinement is a more explicit representation of the change in thermal resistances through (z) the board
- Each layer is modelled as a separate object with its own thermal conductivity
- Metallic layers are still composites of FR4 and Cu traces, pads etc.
  - Metallic layer effective thermal conductivity can be derived assuming thermal resistances in parallel for all 3 directions, just %Cu required
- Numerically more intensive but provides a better accommodation of spreading effects
  - Especially if the distribution of Cu on metallic layers is relatively uniform

This is numerically more intensive, you need at least one grid line for each of the cuboids that constitute the stack-up, but it does provide better accommodation of spreading effects especially if – and this is the key here – the distribution of copper on the metallic layers is relatively uniform, then the isotropic assumption of conductivity on that layer is valid. If there are large variances on the copper density, if you like, as you go across the XY plane, this will necessitate moving up to the next, . . .

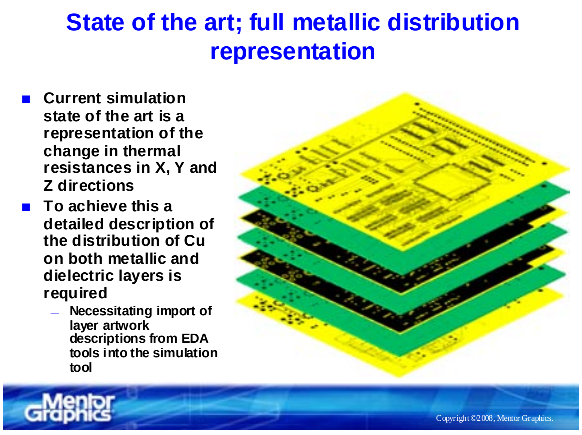

State of the art; full metallic distribution representation

- Current simulation state of the art is a representation of the change in thermal resistances in X, Y and Z directions
- To achieve this a detailed description of the distribution of Cu on both metallic and dielectric layers is required
  - Necessitating import of layer artwork descriptions from EDA tools into the simulation tool

. . . what we're calling the state of the art representation of PCBs. The current simulation state of the art is a representation for change in thermal resistances not only in the Z direction but also in the X and Y direction as well.

To achieve this as an input to the model it is required to have a detailed description of the distribution of copper on both the metallic – so the yellow layers here – and also the dielectric layers, so we're looking at, or needing to define the distribution of traces and pads, et cetera, on metallic layers and then the distribution of electrical and thermal vias passing through the dielectric layers. In terms of model input, this necessitates import of layer artwork descriptions from EDA tools into the simulation tool.

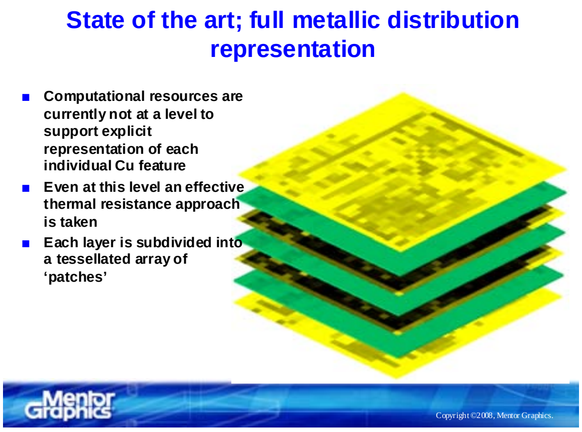

State of the art; full metallic distribution representation

- Computational resources are currently not at a level to support explicit representation of each individual Cu feature
- Even at this level an effective thermal resistance approach is taken
- Each layer is subdivided into a tessellated array of 'patches'

First thing, in terms of expectations. Today, computational resources are currently not at a level to support an explicit representation of each individual copper feature. These copper features can be incredibly small; to have a grid cell capturing each and every copper feature would result in intractable model sizes that would exceed the capabilities, the memory and speed capabilities of current simulation hardware. So even at this level, an effective thermal resistance approach is taken.

Each layer is subdivided into a tessellated array of patches.

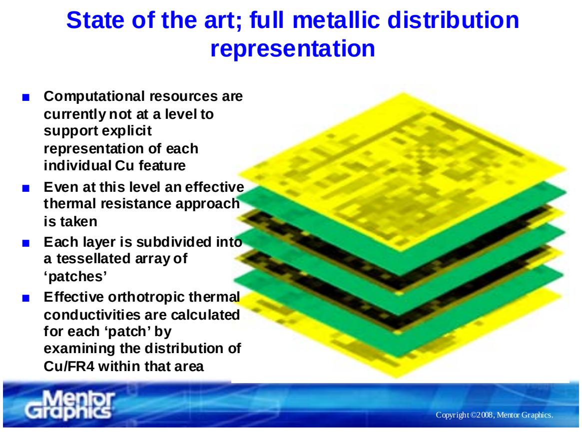

State of the art; full metallic distribution representation

- Computational resources are currently not at a level to support explicit representation of each individual Cu feature
- Even at this level an effective thermal resistance approach is taken
- Each layer is subdivided into a tessellated array of 'patches'
- Effective orthotropic thermal conductivities are calculated for each 'patch' by examining the distribution of Cu/FR4 within that area

Let me go page up again to show the original distribution. If I just go page down, you should see it merging into a tessellated approximation, or pixellised approximation, of the true metallic distribution. For each of these individual patches on each layer, calculations are conducted to derive effective orthotropic thermal conductivities for each of these little patches, by examining the distribution of copper and FR4 within that specific local, rectangular area.

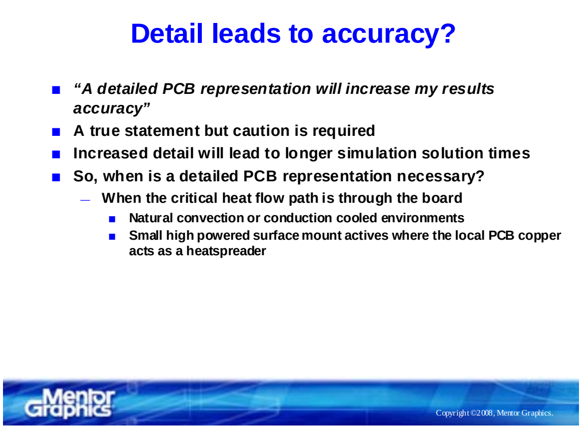

Detail leads to accuracy?

- *"A detailed PCB representation will increase my results accuracy"*
- A true statement but caution is required
- Increased detail will lead to longer simulation solution times
- So, when is a detailed PCB representation necessary?
  - When the critical heat flow path is through the board
    - Natural convection or conduction cooled environments
    - Small high powered surface mount actives where the local PCB copper acts as a heatspreader

Detail leads to accuracy. A detailed PCB representation will increase my results accuracy. That is – yes, it's a true statement but I think it's important to apply some caution here.

Increased detail, it's true for any simulation modelling approach, will lead to longer solution times, heavier models in terms of how you manipulate them, et cetera. I think it's important to appreciate or to pose the question: when is a detailed PCB representation necessary? The simple answer to this is when the critical heat flow path is going through the board itself. Generally true of natural convection, certainly true of conduction-cooled environments, so when a PCB is wedge locked, for example, into the side of a sealed chassis, where heat or copper is actually designed in to suck the heat down into the board and out through the edges.

Also, when there are very small high-powered surface mount actives on the PCB where the local PCB copper is of a length scale comparable to the actual active and therefore acts as a spreader, and, just as you would want to model an air side heat sink in detail, you're going to want to model that local copper in detail, to model the local spreading effects that it provides.

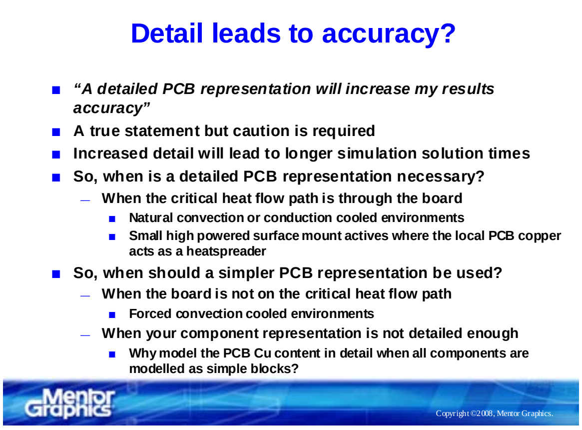

Detail leads to accuracy?

- *"A detailed PCB representation will increase my results accuracy"*
- A true statement but caution is required
- Increased detail will lead to longer simulation solution times
- So, when is a detailed PCB representation necessary?
  - When the critical heat flow path is through the board
    - Natural convection or conduction cooled environments
    - Small high powered surface mount actives where the local PCB copper acts as a heatspreader
- So, when should a simpler PCB representation be used?
  - When the board is not on the critical heat flow path
    - Forced convection cooled environments
  - When your component representation is not detailed enough
    - Why model the PCB Cu content in detail when all components are modelled as simple blocks?

Copyright ©2008, Mentor Graphics.

Conversely, when should a simpler PCB representation be used, or rather, when can you get away with a much simpler PCB representation? Well, exactly the opposite is true: when the board is not on the critical heat flow path. Primarily, this is true for forced convection cooled environments. If you blow loads of air over a board, you airwash it, you're designing it to manage the heat to be extracted from the top surface of the components, possibly through a heat sink, into the air and out. The board is not on the critical heat flow path and therefore the representation of the board would not unduly affect the predictive accuracy, and therefore you can and should represent the PCB as simply as possible.

Another point I make here is when should a simple PCB representation be used? When your component representation is not detailed enough. As I said, your component selection is probably going to be dictated by availability of data. If all you have are simple block representations of components then it really doesn't make any modelling sense to model the PCB thermal resistance in absolute detail, when your model is going to be swamped by issues associated with your single block representation. Having a single block for a component and going for a lot of detail in the board is still only going to give you component predictions that are going to be indicative of case temperature, it can never be more than that.

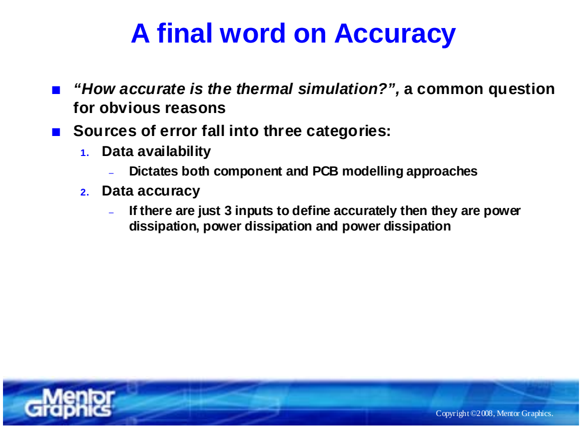

A final word on Accuracy

- *"How accurate is the thermal simulation?"*, a common question for obvious reasons
- Sources of error fall into three categories:
  1. Data availability
     - Dictates both component and PCB modelling approaches
  2. Data accuracy
     - If there are just 3 inputs to define accurately then they are power dissipation, power dissipation and power dissipation

Copyright ©2008, Mentor Graphics.

My final slide, a final word on accuracy. How accurate is the thermal simulation? A very common question, for obvious reasons – I was going to say it's the second most common question, the first being, "How expensive is your software?" Okay, sources of error falls into three categories, in our experience. First, as I've said right at the beginning, is data availability, so this dictates both component and PCB modelling approaches.

Secondly, data input accuracy. I've said here that if there are just three inputs to define accurately, then they are power dissipation, power dissipation and power dissipation. I cannot stress this enough. In all the experiences we have had in benchmarking PCB level thermal simulations with experimental data, by far and away the most common cause of error has been inaccurate definitions of the power dissipation. If you are striving for accuracy, then put a lot of effort into obtaining as accurate a description of the power dissipation as possible.

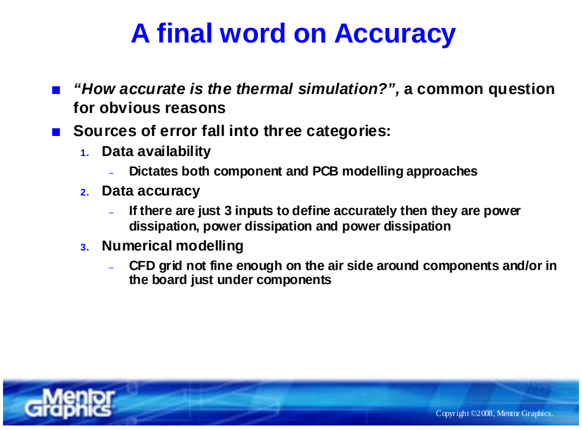

A final word on Accuracy

- *"How accurate is the thermal simulation?"*, a common question for obvious reasons
- Sources of error fall into three categories:
  1. Data availability
     - Dictates both component and PCB modelling approaches
  2. Data accuracy
     - If there are just 3 inputs to define accurately then they are power dissipation, power dissipation and power dissipation
  3. Numerical modelling
     - CFD grid not fine enough on the air side around components and/or in the board just under components

Copyright ©2008, Mentor Graphics.

Finally, there are risks or errors associated with numerical modelling approaches, the most common being that the CFD grid is not fine enough on the air side around the component and/or in the board just under the component. The temperature gradients, both in the air side and the board side, are very, very high, especially around thermally-critical components, it's very important you have a fine enough CFD solution grid to be able to capture those gradients. With our FloPCB products and the coming EDA interface window with Flotherm, FloEDA, we've put a lot of effort into automated gridding algorithms to automatically put fine grid in and around components that are capable of predicting junction and case temperature.

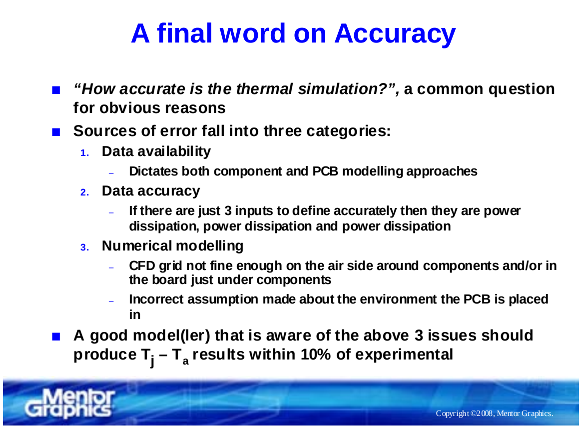

A final word on Accuracy

- *"How accurate is the thermal simulation?"*, a common question for obvious reasons
- Sources of error fall into three categories:
  1. Data availability
     - Dictates both component and PCB modelling approaches
  2. Data accuracy
     - If there are just 3 inputs to define accurately then they are power dissipation, power dissipation and power dissipation
  3. Numerical modelling
     - CFD grid not fine enough on the air side around components and/or in the board just under components
     - Incorrect assumption made about the environment the PCB is placed in
- A good model(ler) that is aware of the above 3 issues should produce $T_j - T_a$ results within 10% of experimental

The second source of error from a numerical modelling perspective is an incorrect assumption made about the environment that the PCB is placed in. I alluded to this before. For example, you can model the component and the board in loads of detail but if you don't get your contact resistance right for the chassis which it's connected to, and that is on the critical heat flow path, then of course that's going to swamp back errors right back to the junction temperature of the source of the heat.

I'll conclude by saying a good model, or modeller, that is aware of the above three issues, should produce T_J minus T_A results, temperature rise results that are within 10 per cent of experimental. If you find your results to be sort of 15 to 20 per cent plus over experimental, I would advise going back to one of these three possible sources of error and confirming whether they're responsible. If you have results with less than 5 per cent or a less than 3 per cent, then, in our experience, that is a very happy coincidence.

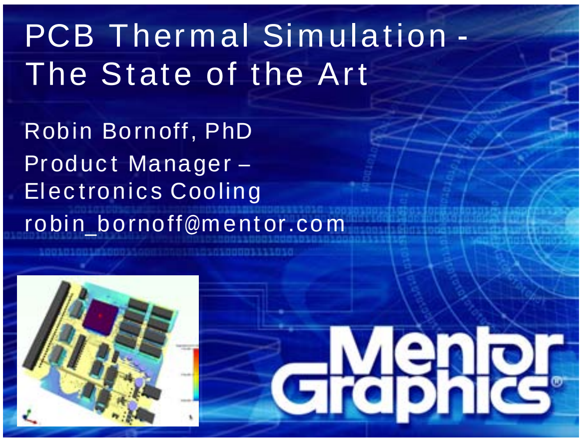

PCB Thermal Simulation - The State of the Art

Robin Bornoff, PhD
Product Manager – Electronics Cooling
robin_bornoff@mentor.com

That concludes my presentation.