[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: Why don't we get started. So before going back to talking about more kinds of power electronic circuits, I wanted to spend one lecture talking about heat transfer and how you model and design thermal systems. And this is principles of power electronics, chapter 25. Now, to be clear, you can spend an entire course-- or for that matter, multiple entire courses studying heat transfer and the ways to do that.

And clearly, we don't have time to do that here, so what I'm going to attempt to do in this lecture is give you just enough information to be dangerous. In that a lot of the things we need to do as a power electronics designer, let us offset some of the detailed computational or deep analytical modeling onto heat sink designers and other people. And what we need to know is enough information about how to model these things so that we can predict, things like temperature rises and select the kind of correct components. And that's what we're going to try to get to today.

Now, if we think about heat transfer, there's sort of-- and first of all, why am I obsessed about this? It's because despite our best attempts, we're going to dissipate some power in our devices and our components. And if I want to make something really high power density, I've really, first of all, got to minimize how much heat I dissipate. But in the end, I've also got to take whatever heat I dissipate and get it out of the system. And how good a job I do of getting it away from the components and keeping their temperature rise low says a lot about how high performance a power converter I can build.

Now, let's assume we're going to dissipate a certain amount of power and ask, how do I get the heat out of my system? And there's really three fundamental ways you can do that. You can do it through conduction. Conduction is just simply, if I think of heat as being vibrations of particles, I can let the heat conduct through some heat conductor and be carried away. So some physical element, like a piece of metal, will carry heat along it.

I can do it by convection. I transfer heat to some moving fluid-- that could be air or water-- and I let the fluid carry the heat away. Or I can do it by radiation. If I make something hot, it emits photons at whatever wavelength. It emits electromagnetic radiation that carries the heat away from my device.

And depending upon the application, all three of these can be important. However, radiation, if I looked at the rate of heat transfer-- so this is heat power being transferred-- tends to be proportional to the source thing that's delivering the heat away to the fourth power minus the environment around it to the fourth power.

So this high to the fourth power means that when things get really hot, radiation really dominates heat transfer. But for most terrestrial applications, it's not the highest-- it's not the most important thing. So if I have an incandescent light bulb and it's really, really hot, then yeah, a large part of the energy may get out of it by just radiating away from it as blackbody radiation.

If I'm in a satellite, and there's no way for heat to conduct out of the thing, and there's no air or anything for the heat to go into, well, then eventually, all the heat is going to go away from that thing by radiation. So in those kind of applications, yeah, I really may care very heavily about radiation. But in most cases, it's only a small-ish fraction of the heat I'm getting rid of, because we don't get crazy high temperatures. We don't have thousands of degrees Kelvin to remove the heat.

So let's go back and look at convection and conduction, and I'll start with conduction. Suppose I took-- actually, I'll just go over here. Suppose I took some-- I'll keep everything 1D, and I'll get some thing that I'll put at temperature T1. And then I'll have some block of material, and I will have a second surface that I will keep a temperature T2.

And in between them, I have some material through which heat is going to flow. It's going to conduct. And this material, whatever it is, might have some cross-sectional area A, and it might be transferring heat through some length l. And in this kind of situation, with 1D heat transfer, for many materials, what you get is the following relationship.

Q-- that's the amount of heat transferred-- is proportional to the temperature difference T1 minus T2 times the cross-sectional area divided by the thermal resistance, RTh and divided by the length. So if I have more cross-sectional area, more heat flows. If I have to transfer it through a thicker piece of material, the heat flow goes down.

And this proportionality constant is called the thermal resistivity of the material. I could rewrite this as T1 minus T2 divided by RTh times l divided by A. That would be another way to write this. So this would reflect how heat-- this would be reflect conduction of heat in steady state.

This relationship-- and this is sort of a bulk relationship-- I could create an incremental relationship, which would be Q is equal to minus A over rho Th dT dx. That would be if I treated everything as a little increment in space. And if I want to start doing Fourier heat transfer, I get into these kind of differential equations descriptions. But for a bulk piece of element, we can write it like this.

This relationship, where Q is my heat flow-- so Q is whatever heat is flowing this way-- suggests a circuit analog where-- here's my circuit model. I have some temperature source T1 and another temperature source or temperature sink T2.

So temperature now means-- a voltage source here means constant temperature. And then there's between them some thermal resistance RTh which is equal to Rho Th times l over A. And that has the same form as an electrical resistor. And then my heat flow Q is the current flowing around this loop.

And here, I have some point to which these temperatures are referenced. Ideally speaking, this should be absolute 0. In practice, people often just reference it to ambient and pretend ambient is the equivalent of absolute 0. It'd be kind of like having an offset from absolute 0.

Any questions about that structure? You get-- if you analyze this circuit, you get exactly this relationship. The reason this works is temperature is conserved going around any loop. If I measure temperature differences around a loop, I ought to get 0. And if I consider heat inputs and heat transfers, conservation of power is the equivalent of KCL in this circuit model.

So I'm showing you a simplified circuit model for a very particular kind of heat transfer, but you can use this model to represent heat transfers in a variety of conditions. Any questions about that? So I should say, by the way, Q is in watts. That's the heat flow in watts.

T is in Kelvin, or often, people use it in degrees Celsius. And RTh ends up having units of degree C per watt or Kelvin per watt. In this regard, thermal resistance is very like regular resistance, because if I think about a regular resistor, if I make the resistor proportionally longer, the thermal resistance goes up. If I increase the cross-sectional area, the thermal resistance goes down.

And here, all I have is thermal resistivity instead of electrical resistivity. So the form of the relationships are exactly the same. So basically, we can assume all of our devices in steady state have some thermal resistances. We can put power into the system. If we know some fixed temperatures, like some ambient temperatures, we can create a circuit and then analyze what our heat flows will look like.

This is conduction. What might my other means of transferring heat energy look like? Well, my other means of heat transferring energy might be something like this. I've got some structure which has heat going into it and then some fluid going over it. This could be a flow of air, or a flow of water, or something like that.

And what we expect to happen is the heat is going to come out of whatever the item is up into the air, and then the air is going to carry it away. So we're going to transfer the heat into the moving fluid, and the moving fluid kind of walks away with it. What does this look like? Heat transfer for this case looks like Q is the heat transfer is equal to hA times the temperature of the surface minus the temperature of the fluid.

So if I have some fluid at T fluid, and I have some temperature at the surface T surface, the heat transfer again in some range is going to be proportional to the temperature difference between the surface of this thing and the fluid that's carrying the heat away. And this is the so-called wetted area.

This is the area at the interface that the fluid is flowing over to grab the heat and carry it away. And then this is the heat-transfer coefficient. Now, here's where it's-- right here is essentially where I'm burying all the ugliness. This is where I'm massively oversimplifying it, and the reason you can take whole courses in this-- cool courses in this topic is because what is this heat transfer coefficient?

Well, it depends upon is the flow laminar or turbulent? It depends upon all kinds of stuff. And it can even be a function of temperatures in the general case. So I have to tell you-- I'm telling you this magic number. How you find that magic number is tricky in a lot of cases. But what we're going to see is, as designers, we often don't have to worry about that.

So in this case, I could have the same model where I would just say, RTh is simply equal to hA inverse. If I told you the area through which I was transferring heat and I knew that heat transfer coefficient, great. I've got my answer, and I go on my merry way.

So putting aside that it's difficult to find this number, or in some cases, this equivalent number, we have a means of predicting what it is, and I can just model that heat transfer as a thermal resistor again. Any questions about that? So how might we use this? Well, in the typical case, I'm going to dissipate some power in my device. I've got to connect that device through some heat transfer means, and the heat transfer is going to carry it away.

So here's a typical structure. Maybe I have some power device. This could be a MOSFET. And basically, somewhere inside this device, I'm dissipating a bunch of power. So that's a source of heat. And then what I'm going to do is I'm going to say, I need to connect that down onto a heat sink.

And here, when I connect this device, I'm going to have some thermal-- I'm going to have some interface here. What is that interface? I'll talk about that in a second, but that's going to be between the device, and then maybe a heat sink. And maybe the heat sink has a bunch of fins that's designed to help me cool this thing.

And so the heat is going to come from the device through this interface into the heat sink, and then maybe the heat sinks, transferring it off to-- by convection past the heat sink out into space around it, and then I forget about the heat. It's just going out somewhere.

Well, what is this interface? This device here-- and let me just illustrate just as an example. Here is a typical MOSFET. This is in a TO-220 FET package, which is maybe not the most modern of packages, but they're still used in a lot of cases. And so the device is inside this plastic package. These leads, gate drain and source are connected to the device, and then one side of the device is connected down onto this pad through which you're going to transfer heat.

So that's the case of the device. So inside this package somewhere, I'm dissipating heat, and it goes to the case of the device. Then I'm going to mount the case of the device down to some heat sink. But there is some interface of the heat sink. Now, in a practical device, it depends on the device type. In this particular device-- this is a power MOSFET-- that heat sink-- or that terminal of the device through which all the heat is going to go is actually connected to the drain of the device-- to the high side, the top side of the MOSFET.

Why? Because it's a vertical MOSFET, and that's the way they could connect it. And so if that drain is, say, for a boost converter, that may be flapping up and down in voltage. So it's kind of unfortunate that you're trying to transfer heat out of something that's flying around in voltage.

So in that case, this thing might be a silicone insulation pad. What does that thing do? Essentially, it's electrically insulating, but reasonably thermally conductive, so it lets heat transfer through, but not electrons. And by the way, that's not a very common thing. Most things that are good thermal conductors are also good electrical conductors, but you can find some materials that will help you out there.

So this thing might be some electrical insulation pad to let you connect this thing to a heat sink thermally but not electrically. You can find-- a typical brand name for this is like a Sil-Pad. That's Bergquist's trade name for it. There's a lot of companies that make it, and they come in varying degrees of quality and electrical insulation capability and so forth.

In some applications, maybe you don't need electrical insulation. If you happen to have a package that was already electrically insulated-- and they do make some FETs that way-- or the package happened to be connected to ground, maybe you could connect that directly down to the heatsink. But often, you will have some other interface between your device and your heatsink.

Why? Because if you think about surface roughness, maybe here's the back package of the device and here's the surface of the heatsink. You don't want any air in between it. So in that case, what you might do is use thermal grease. So this could also be instead thermal grease.

Why would I use thermal grease? Because thermal grease has a lot better heat transfer characteristics than air does. So you wouldn't want any scratches or open areas in the back of your device limiting heat flow, so you fill that with thermal grease. Of course, thermal grease has way worse thermal conductivity than metal.

So what you would do is you just put enough thermal grease on it. And sometimes, if you've ever mounted, say, a CPU down to the CPU cooler, they'll give you this material to put in there. And the idea is you put it on, and then you basically try to squeegee it all off.

It's just to fill in cracks. You don't want like a thick layer of this stuff, because it will add thermal resistance to this path. The Sil-Pad adds thermal resistance. So does thermal grease. So you put as little as on as you can get away with it just to make it essentially flat, and then mount that device down to the heatsink.

Any questions about that? So it depends whether you need electrical insulation, or whether you just use thermal grease, or you use an insulating pad, or you use some kind of non-insulating thermal pad that'll be compliant. And that's the physical structure. What does this thing look like as a circuit model?

Well, here, I will have some Pdiss. That's the power I'm dissipating in the device. And ideally, that power dissipation is connected to absolute 0. But maybe I'll just call this T ambient and recognize that, oh, maybe this is connected to really 0 Kelvin. So I can either treat this as 0 Kelvin or call it whatever my ambient temperature is and work from there.

That's the power dissipation. And we often say this temperature is at T junction from-- so it's either-- this is T junction to ambient or T junction absolute depending upon how I decide to write this. The junction is just my notion of wherever the power is being dissipated in the device. It may be an actual physical semiconductor junction. It may just be wherever in this device is where the heat is being dissipated.

Then I've got to transfer the heat to the surface of the device, because I said the heat's actually being dissipated inside, and there's some thermal resistance junction between where the heat is being dissipated inside the semiconductor to the actual case. So there's going to be some RTh junction to case.

And then there's going to be some other thermal resistance, and so this temperature here would be the T case. Then there's going to be some other thermal resistance, RTh case to sink. That depends upon is this grease, is it a thermal pad, what is it? And I care because that's going to contribute to temperature rise.

And then over here, I have my heat sink. And that heat sink, if I include what's going to transfer off it by conduction and convection, maybe I say there is some sink to ambient. And I connect this back to ambient, and that's my thermal model, and this is Tsink. So this is my model for the whole thing.

Any questions about that? And I should say this is in thermal steady state. Now, fortunately, if you get a lot of typical devices, and you look carefully in the device datasheet, here we go. What they tell me is they've done some measurements.

Somebody with a lot of thermal modeling experience has said-- done some measurements to figure out how hot the junctions are going to get from above the case for a given amount of dissipation. And they give me a thermal resistance junction to case. And they say here, it's 2.5 degrees C per watt.

And then they tell me, oh, if you go get just the right amount of thermal grease and you do a really good job of mounting this down to a heat sink so we're not worrying about any additional insulation pad here, it's another half a degree C per watt between the case and the sink. That's if I don't need electrical isolation. If I do, I have to look up the thermal resistance of my Sil-Pad or whatever else I'm putting in there.

And then I'm-- usually, as the designer, I'm trying to pick this. But if I went out and picked a particular heat sink, I can look up its value of RTh sink to ambient. So then I can find delta T junction to ambient is simply going to be Pdiss, which is this power dissipation, times RTh junction to case plus RTh case to sink plus RTh sink to ambient, and then I'm done.

If, in my example-- typical example numbers that I might use-- I said T ambient might be 40 degrees C. I don't know. That's not-- that's a very that's actually a pretty cool ambient temperature. Some industrial stuff you design, they assume the ambient is going to be 100 degrees C.

But 40 degrees C is hot but not crazy hot if you're in a room. And maybe I would say this kind of device, Pdiss might be equal to 10 watts, This is something I got from my circuit design. So here, I could say, that's 10 watts times-- junction to case I said was 2.5 degrees C per watt, plus 0.5 degrees C per watt was the-- if I just had thermal degrees in there.

And if I went out and I bought a Redpoint thermal km 51 heatsink-- that's just one of these convection-cooled heat sinks to which I might mount this device-- that is 4.8 degrees C per watt. Which if I say this is plus T ambient-- so this would be just T junction-- would be T ambient plus this stuff. This would give me 118 degrees C, and maybe I'd be a happy guy.

Typical devices that you can buy in plastic packages, 150 to 175 is the maximum temperature you're allowed to get the junction. You can get devices that are specially packaged that will go higher, or if they're not going to-- bad stuff won't happen to the device. It's often limited by the packaging. But typically, 150 or 175 is what you can get, and this just gives me some margin. Any--

AUDIENCE: Is that on the data sheet?

PROFESSOR: Yes, usually, a data sheet will tell you what the maximum allowable temperature-- here we go. Operating and storage temperature range-- operating junction and storage temperature range. 150 degrees C for this part. A lot of more modern devices might be 175. It just depends what you're buying.

So what would you be-- really job? Usually, what you do is-- you know all this stuff. You know your ambience of specification, you know your dissipation because you calculated it, you know these things. And your job is to go pick the heat sink or maybe the thermal interface material. And you have to pick this so that this number doesn't get too high.

So I don't know this stuff in advance. This is the thing I'm choosing as the heat sink designer. What happens if you can't do it? Well, you either got to go reduce your dissipation, or find a better package with a lower junction to case, or some other combination of things. And you're often pushed up against the wall if you're-- especially if you're trying to go for high power density. Questions?

AUDIENCE: So it seems like a lot to do for one component in your system. What do you do for a workflow when you have a lot of components in the circuit [INAUDIBLE]?

PROFESSOR: So if you're a power converter designer, usually, you only have a few sources of thermal dissipation that are the most important, and they're often your power devices. At a broader scale, if you have to consider the whole thermal design, well, either you get a bigger network, or you start to pull out computational finite element tools or something like that.

But very often, for a power converter design, you think-- you've got your-- in a buck converter, you might have your MOSFET and your diode as being big sources of loss, and then maybe your inductor is next. And you just model those into the system and figure out how you're going to heatsink it.

I should say, by the way, that a very typical thing that you would do is-- a very typical thing to do is-- I'm showing you one picture here, but it may be true that you come up and you might do the following. You might have a heat sink, and then your diode and your MOSFET are right next to each other and mounted to the same heatsink.

So what would you do? You'd say, here's PFET and here's TJ-FET. And that would be RTh. You'd have junction to case, case to sink. So maybe that would RTh-- junction to sink would be the total. And you'd do the same thing for RTh junction to sink for your diode. And here's Tj diode.

And then you'd go out and get one RTh sink to ambient, and you'd go calculate what the junction temperature of the diode is and what the junction temperature of the FET is. And you'd have to pick these things low enough, and you can expand out from there. Questions?

So that's one way to think about it. This is the steady state thermal model. This is, I plug my converter in, I let it run for a long time. Everything heats up to its maximum temperature, and I find out if it blows up or not. Did the thing get too hot when everything settled out?

But this isn't really a transient thermal model. I mean, if I put my tea kettle on, and I turn on the power, which is basically dumping heat into the water from my tea kettle, it doesn't instantly get up to maximum temperature. Eventually, it'll get up to-- hopefully, if I have-- pouring enough heat in to get up to boiling, but it doesn't do it instantly.

This model does not tell you at all about instantaneous behavior or transient behavior. And sometimes that's quite important, because we sometimes don't even wait for things to get to steady state, or we want to know, how quickly is my BIEST going to heat up? Well, it turns out you can measure that as well, or you can model that as well.

And what we do with it is to model things with a thermal capacitance which has units of joules per degree C or joules per Kelvin. How many joules do I have to put into this structure to raise it one degree Kelvin just by heating up the mass of this thing? And so if I know the heat capacitance for some material per volume or per mass, I can multiply that and get the heat capacitance of the whole thing.

So maybe what I would do is I would say, suppose I have some surface here, which I know is going to be at some T ambient-- or let me call this T2. I guess I did call it T ambient. So suppose it's T ambient. And here, I have some block of material-- piece of element of my system.

And let me characterize this by its average temperature T1. And now, I'm going to go pour heat into this thing, Q. This is the power dissipation. Maybe it's at this surface. This is going in and heating this thing up, and this thing is transferring energy to ground.

How would I model that? Well, the steady state model would look like this. So this is Pdiss-- or this is Q, I guess. What did I call it in my note? I guess I called it Q. Then I have some-- if this is T ambient, this is RTh and this is T1.

So this is the thermal model in steady state for this thing. How do I account for the fact that this thing has some thermal mass that takes time to heat up? Well, I find out what is thermal capacitance is and I put it right here. CTh.

And now, this gives me a model so that if at some time-- suppose here's T. If I come in here and I say, here is Q of T. It steps. This is Q of T, and at T equals 0, it turns on. What I'm going to get is a T1 with respect to ambient that's going to do this.

Eventually, I'm going to get the final value is going to be-- T1 minus t ambient is going to be equal to Q times RTh. That's this final value, and it's going to rise with a time constant tau that's equal to Rh Th-- Rh-- RTh CTh.

So it'll have a time constant. It'll rise up, so it won't immediately get hot. It'll just rise over time. Now, I should say, fundamentally, a thermal capacitance should technically always be referenced to absolute 0. If I'm just dealing with some ambient, and I want to call it an offset, I can reference it to a fixed temperature.

But thermal capacitance is in terms of modeling them heating up. I really have to-- in terms of modeling the physics, I always have to have that thermal capacitance to my reference point, which would ideally be 0 degrees Kelvin. And I'll come back to that in a moment. Any questions about that, though?

AUDIENCE: So in this case where we have multiple devices, how would we insert the thermal capacitances?

PROFESSOR: That's right. So in this case, if I was talking-- and you've got to say the thermal capacitance of what now-- the heat sink, the device? If I was talking about the thermal capacitance of the device, I would put that right here. So this device package might have a thermal capacitance, and this device might have a thermal capacitance, and I ought to place it to the reference potential. The heat sink might have a thermal capacitance, and then I get a higher order thing.

Now, when I get to T equals infinity and I have some step response here, it all settles out to the same steady state model. Along those lines, what happens if I'm not too happy with the resolution I have here? I'm basically treating T1 of this whole block as being the same thing.

But in reality, I recognize if the heat's really coming in over here-- this ends hotter, this ends cooler. It's not got-- spatially, it doesn't have the same temperature everywhere, and maybe even temporally, it doesn't have the same temperature everywhere. If I'm not happy with this model, I can break it up into more lumps.

So I can come over and say here's my BIEST. Let me break this up into T1, T2. Here's T ambient, here's Q. And now, I have a different thermal model. I have Q coming in.

I could say, here's CTh1 for just this region. Here's RTh 1 to 2, which is the thermal resistance here. Here's CTh for this block, and here's the thermal resistance coming back to ambient here. And this is going to be T2, this is going to be T1. And now I have something that's more spatially accurate.

If all things were equal, RTh, CTh would be a factor of 4 smaller. I'd get better time resolution and better spatial resolution as to how I'm modeling my system. If I carry this to an extreme, I get-- basically, I'm back to Fourier heat transfer and partial differential equations. So this is your best lumped model approximation to Fourier heat transfer.

Any other questions? How do people model this in, say, the thermal capacitance aspect? Sometimes you can look up the thermal-- if you have a material, you can go look up its heat capacity, and then multiply by the mass, and you get a thermal capacitance, and you can do it from fundamentals this way.

That's not always the way they give it to you for a device, because if I think about it, a device is kind of a complicated thing, right. There's the silicon, and then maybe inside there's some header, and there's different stuff going on. Sometimes, what they will do is they will give you a lumped model. And unfortunately, this is often called a transient thermal impedance, or ZTh.

And it's a-- I say unfortunately because it's not an impedance in the way you think of. An impedance is a frequency domain thing, but this is just the language that has stuck. It's unfortunate, but that's what they call it. But it's not an impedance at all.

Here is an example. This is the transient-- or here, they actually they call it the maximum effective transient thermal impedance junction to case. So what does this thing mean? This thing means if I took the device case, and I mounted it against a really mega good heat sink that's just going to hold the temperature exactly at ambient-- a perfect heat sink-- then I stepped the power into it.

And when they say a single pulse, what this really is, is a plot of the step response. It's basically this plot for the junction temperature above the case, where the case is held exactly at constant. So what this is saying is if I put if I put one pulse into this device, and I waited 10 to the minus 4 seconds, I would have a 0.1 degree C per watt of that step rise.

So if I had a 10 watt step, I'd have a 1 degree C rise after 100 microseconds. After a millisecond, I'd have a higher value, et cetera. So this curve-- oh, you-- yeah, if you can see my cursor, this curve is really a log log plot of the step response.

If it has one time constant, it would look like one thing. If it has multiple time constants, the shape of this curve changes. So a transient thermal impedance, what they're really telling you is essentially, this step response for the device assuming it's against the perfect heatsink.

What they also sometimes tell you is if I now have little pulses. I step it up, and I step it down, and I wait a while. I step it up, and I step it down, and I wait a while. They will tell you the maximum temperature it will get during the step for different duty ratios of pulses.

Why do they tell you this? Because look, if I put 100 watts into this device continuously, and I put the case at 0 degrees C and I held it there continuously, what this would say is 100 times 2.5 would give me 250 degrees C rise. Well, that would blow the thing up, because this said 150-degree C device.

But if I only did it for a millisecond, this thing would tell me that I'd have only 0.1, 0.2-- I'd be right here on the curve. Like 0.3 degrees C per watt, and maybe it would only go up 30 degrees C. So I can have a 1 millisecond pulse that's really big, and it won't kill the device because of the thermal capacitance of the device in its package. So thermal-- transient thermal impedance is a step response plotted on a log log plot.

The other thing I should tell you is that you could somehow extract that model and come up with a thermal model for the device. Maybe it's only one RC that you approximate it with. But in a complicated big module, it might be a more complicated model, because there might be the thermal capacitance of the silicon, the thermal capacitance of the header and so forth.

You will often see thermal models-- this is called a Cauer form network. You will often see thermal models in datasheets, or app notes, or something that are something like this and that are done like this. The thermal resistances and thermal capacitances, that's sort of a foster form of network.

That's just wrong, unfortunately. You can kind of try to take this kind of model and map it over to that kind of model, but any real thermal capacitance, if you're going to represent Fourier heat transfer, thermal capacitances ought to always be down to-- we can use ambient, but it should always technically be absolute 0. These are referencing the thermal capacitance to other places in the thermal path where the temperature is changing.

So it's basically just not right, but people do it because it's convenient. And I suppose if the thermal capacitance here or the thermal resistance here were super low, maybe it's almost like this was at Earth, so they can get away with it sometimes. But it's just a warning. If you want to stick with the physics, this is the correct kind of model to use.

So that's just giving you a heads up. I want to just tell you about two more things. What I've told you so far is if I have a big package, then I'm going to mount it against the heat sink, and that's a very typical thing you have to do.

In some cases-- and in a lot of designs, it doesn't work that way. You might have something like this. Here's my device that I might mount-- instead of mounting my device directly to a heat sink, you will often have a circuit board. Here's my circuit board.

And I'll have some device, and maybe this device will have a heat pad. Here's my package. It'll have a bunch of electrical connections off the device onto the circuit board. So the first thing I do is I mount my device to the circuit board. And my electrical connections would come off of that.

I'd run my gate and my source. I'd have a bunch of wires on the circuit board going different places. So I've made my electrical connection to device, but now I also need to get the heat out of it. So what do I do? Often, what you do is, instead of trying to mount the device directly against the heatsink, you mount the device to the board, carry the heat through the board, and then to a heatsink on the other side.

So maybe what I would do is, on the other side of the board, maybe I would put a thermal interface pad and then a heat sink. The heatsinks mounted to the other side of the board. And I'm going to try to get heat out of the device through the package, through the board, through the interface material into the heat sink. And this is necessary because you want to connect the device to the board for the electrical connection-- is the most important thing.

So one of the things you often do is you will often put thermal vias in the board. So what you'll do is, you'll go to the board-- and if I looked at this from the side, you'll drill a bunch of holes through the board. You'll put holes in the board like this, and you'll put what's known as a via farm.

If I look straight down on top of it, you'd see a bunch of holes drilled in the board like this. Why? Because those holes will then be plated with copper, which is a good thermal conductor. And I-- it'll help me transfer heat through down through the board to the interface material to the heatsink. Sometimes, you fill those holes with copper in a fancy board, or you fill it with thermally conductive epoxy to help the heat transfer through the board. Sometimes, you'll try to use a thin board so that you have lower thermal resistance.

This app note that you'll be provided with in the handouts talks about, if I have thermal vias of a certain length, how many degrees C per watt for a 12 mil via going through a 63 mil board and so forth? You can figure out your thermal model transferring your heat through the circuit board. So that's just a heads up about how you might do things for PC board-mounted components.

The last thing I'd like to do for you today-- any questions about that? The last thing I wanted to do for you today is just show you a thermal transfer example. So here we go. Here's the thing I'm going to show you.

This is for a particular matching network design that we did, and this thing has these four transistors that are mounted between the circuit and ground. Each of these transistors can dissipate something like 30 watts in the worst case. It's a high power design. So four devices times 30 watts is like 100 and-- well, it's 33 watts, so roughly speaking, it's like 140 watts-- something like that. 130 watts.

That's a lot of heat to pull out of these little devices. Here's the devices. These are the package. One side of this has these electrical connections to the devices, and the other side of this has a big thermal pad out of which the heat is supposed to come. So this is a two-sided device, where you have electrical connections on one side and a thermal pad for heat transfer on the other side of the device.

So one side of this device, the side that's showing flat here, is going to be connected to the circuit board. On the other side of the device, we're going to want to connect to the heatsink. So what do we do? Here you can see one side. This is the top side of the circuit board.

This is not the side of the circuit board where the devices are on. The devices are mounted to the bottom side of the circuit board, and part of the bottom side of the circuit board is this over ground plane you see here. These are the devices. The side of the device you cannot see is what's got the electrical connections for the drain and the gate and the source and the drivers. And the drivers are all on this side of the board.

So what's exposed here on the bottom side of the board is the heat pad through which I'm going to try to transfer heat. So what do we do? We created a heat spreader, because we're going to flip this board over and [SMACKS FIST] bang, down on the heat spreader. You notice that the heat spreader has some area milled out.

So if you look at this, there's a space for the devices to sit in this space in these four spaces for the four devices so that the circuit board will sit flat with the heat spreader, and then the devices will sit down and sit flat against the little milled out area. These little circles here are just areas for extra solder to flow out of the device. So what we're going to do is we have the device soldered to the board, and then we're going to solder it to this heat spreader.

How do we make that work? Well, the first thing we do is we solder the circuit-- all the chips to the board with high temperature solder. So then I've mounted my device to the board. Now, I put thermal-- I put solder ball-- I put solder paste on it. Here, you can see the dots, the solder paste.

I flip this thing over under this heat sink, and then I reflow it again. And this solder is low temperature solder. So I don't unsolder the device to the board while I do solder it to the heat spreader. Make sense to everybody?

So we got this hierarchy of temperatures of my solder. How do I cool this device? Now I got this heat spreader, which has got, really low thermal resistance now. So I've gone from my chip directly to the heat spreader. And by the way, I said these-- this thermal pad on the device is actually connected to source in this particular device, and source is ground in the particular circuit.

So I can basically make this whole heat spreader and the whole thing electrical ground. That's kind of nice. Here's my-- so I'm going to connect the heat spreader onto this Dynatron copper heatsink. And here's what you can see about this heat sink.

Here's the thermal resistance to ambient of the heat sink versus air flow. So this is where I said, yeah, to really know-- this heat-transfer coefficient's a lot of work, and to design the heat sink to get a low heat transfer coefficient or a high heat transfer coefficient is hard. But somebody else did that for me.

All I got to know is how much air I'm going to blow across this, and I can get down to crazy low thermal resistances. If this wasn't low enough, then I'd have to go to liquid cooling and get even lower ones. But we were able to do it with air. So all we did was then we basically mounted the device to the heat spreader-- the heat spreader of the heat sink and then blew air across it.

Here, you can see another view of it. Here's the heat sink, the device mounted to it, and we're to blow a lot of air to it. Here's the actual BIEST So you can see here's the circuit mounted to the heat spreader. The heat sink is now all silver because we-- it's silver-colored because we silver-plated it. And here, you can see the fan blowing air through the fins.

The result of that-- as I said, we had like 100-and-something watts of heat dissipation in the worst case. But if you look at the temperature rises of the thing, the heat sink surface only goes up 20 C. Why? Because we got crazy heat flow.

So this is the kind of thing you do. And we use exactly these models to essentially figure out what's the thermal resistance from the junction to the case. Then we have the case going down to the heat spreader. And it was soldered, so that's pretty low.

And then the heat spreader to the heat sink, and the heat sink to ambient we looked up from the data sheet. And boom, we have our model, and we have our nice thing where we can pour lots of dissipated power into it and still keep it cool. So I'm out of time. Are there any final questions.

AUDIENCE: Will these slides be posted?

PROFESSOR: Yeah, we can put up these slides. Yep. So this-- yeah?

AUDIENCE: I'm just wondering what the device was.

PROFESSOR: What the dev-- it was a GaN system. There were four GaN systems. GaN on silicon 650-volt transistors. So yeah, that's a lightning introduction to quick and dirty heat sinking. I still encourage you to actually take much more sophisticated classes on heat transfer.

I'm not saying this replaces that, but it's enough to get you to build some pretty cool stuff. So next time, we'll go back to talking more about circuits. Have a great day.