AND8215/D

Semiconductor Package Thermal Characterization

Prepared by: Roger Paul Stout, PE
ON Semiconductor

ON Semiconductor®

http://onsemi.com

APPLICATION NOTE

Glossary of Symbols

$R(t)_{JA}$, $R_{\theta JA}$, $R_{\theta JMA}$, θ_{JA}, θ_{JC}	various notations indicating transient and steady state thermal resistance to ambient, case, normalized to actual power along path of interest
Ψ_{JT}, Ψ_{JL}, R_{JS}	("junction to top", "junction to lead", "junction to solder") thermal resistance parameter normalized to total package power
TSP	Temperature Sensitive Parameter
DUT	Device Under Test
TC	thermocouple

GENERAL INFORMATION

In order to measure thermal resistance of packaged semiconductors, some basic information needs to be provided. Die size, thickness and active area are used to calculate certain thermal transient characteristics of the device. Certain material properties are also necessary, specifically density, specific heat, and thermal conductivity of the primary materials in the package (encapsulant, silicon, die attach, leadframe, etc.), and from these, derived thermal transient properties of diffusivity and effusivity.

The package type is also important, surface mount or through−hole, in order to determine mounting requirements. Surface mount devices are tested on FR4 boards with minimum or 1″ pad areas. TO−220 and larger power devices are tested on a cold plate.

When measuring temperature of any power device, it is basically impossible to put a physical thermometer onto a device's junction while under power. Instead, we must utilize some temperature "sensing" method internal to the device. For instance, in power MOSFET's we ordinarily use the device's inherent "body diode." The forward−biased voltage drop of this pn junction has a very linear relationship with temperature, so, when properly calibrated, we can use it to tell us what junction temperature results from any power condition.

Thermal Parameter Test Procedure

The Temperature Sensitive Parameter (TSP)

To thermally characterize a semiconductor package, it is necessary to have a temperature sensitive parameter available (such as a diode or a resistor) within the device being tested, which can be used to measure the die surface temperature. The voltage of this TSP (in theory, at a fixed current) is measured in a calibration oven at temperatures of 25, 50, 75, 100, and 125°C. The current used is very low (typically 1.0 mA) to prevent significant self heating of the device. (In practice, a constant current supply is approximated by using a constant voltage supply and a large resistor in series with the TSP, as shown in Figure 1).

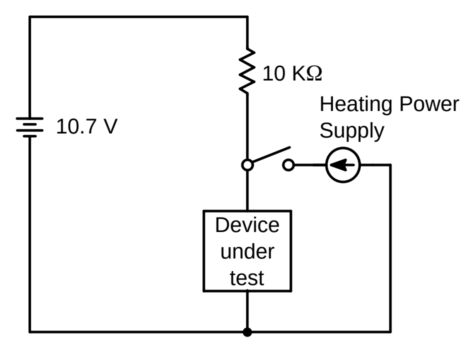

Figure 1. Basic Thermal Test Circuit

Calibration

Before calibration, a surface mount device is assembled onto a thermal test board such as the min−pad board illustrated in Figure 2, or the 1″ pad board as in Figure 3. Tab and “back of board” thermocouples are typically attached for external package temperature measurements.

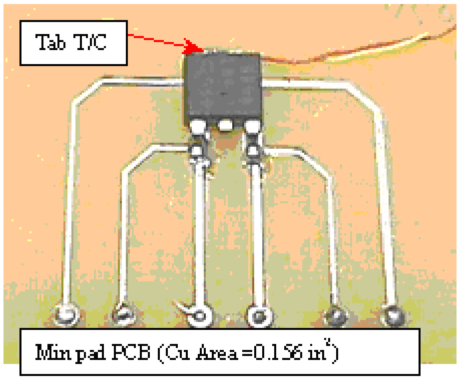

Figure 2. D²PAK on Min Pad Board

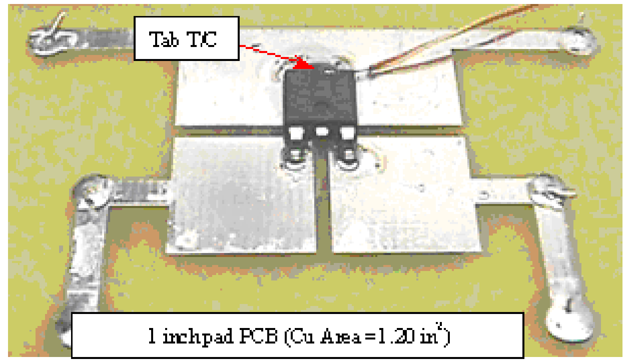

Figure 3. D²PAK on 1″ Pad Board

Devices to be characterized in a non−surface mount condition (such as socketed TO−220’s) are calibrated directly in a socket such as will be used in subsequent testing.

Thermal Test Configurations and Fixturing

After calibration, a surface mount device is put into a one cubic foot still air test chamber with the board in a horizontal position, as shown in Figure 4. An ambient thermocouple is positioned within the chamber approximately 1″ below the test board and off to the side. Generally, still air characterization is limited to about 2.0 W maximum power dissipation (and may be much lower).

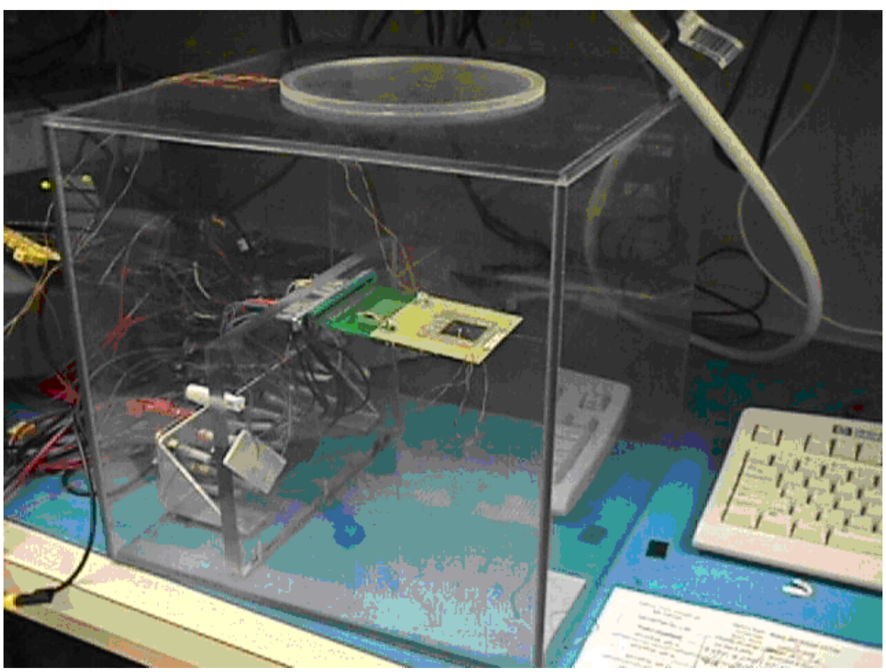

Figure 4. One Cubic Foot Still Air Chamber

For power devices requiring large amounts of power for characterization (for instance, heat−sinked TO−220’s), a coldplate (Figures 5, 6 and 7) is used with the device mechanically or hydraulically clamped to the surface. A thin layer of commercial “thermal grease” is applied at the interface between the device and the coldplate to minimize thermal contact resistance. Again, external package temperature measurements are generally made with a K−type thermocouple glued or soldered to the exposed heatsink tab of the package. Additional thermocouples are embedded within the coldplate at various locations to monitor the coldplate conditions.

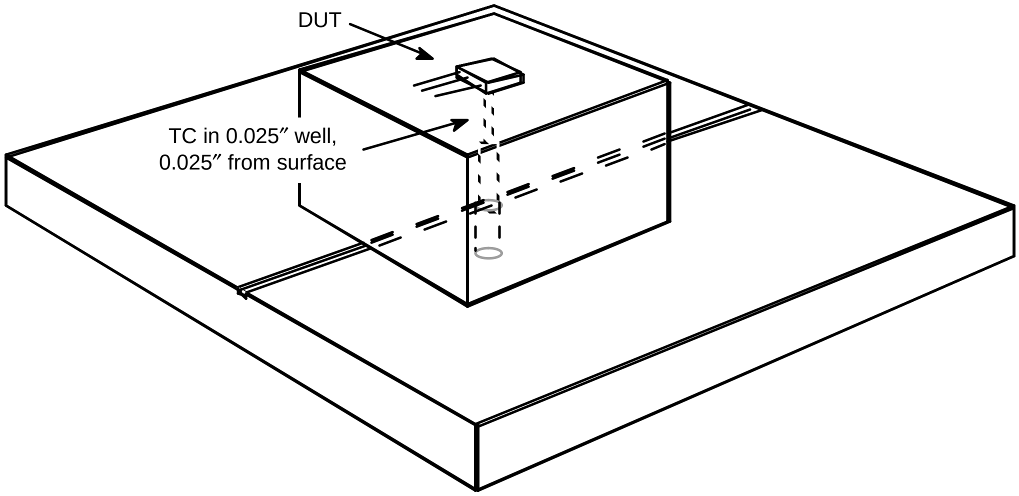

Figure 5. Oblique Schematic of Coldplate

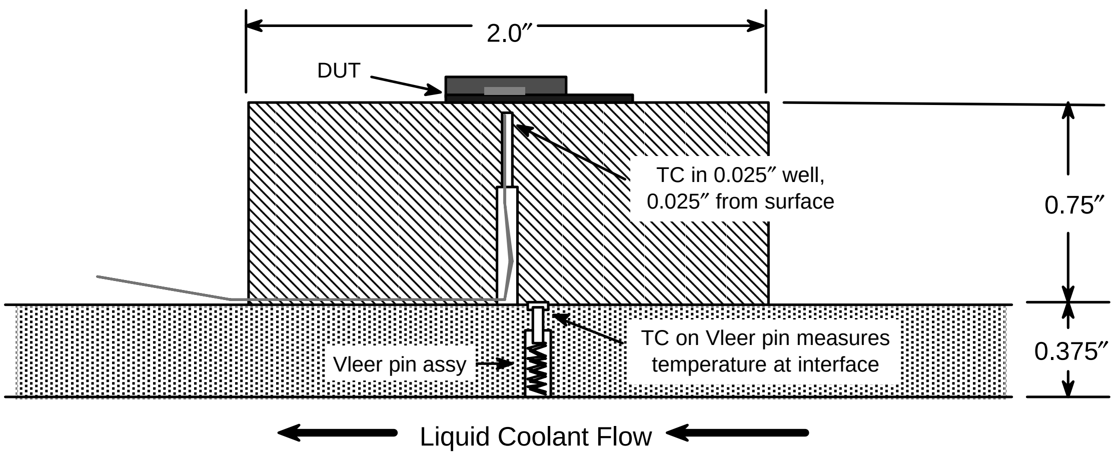

Figure 6. Cutaway Section View of Coldplate

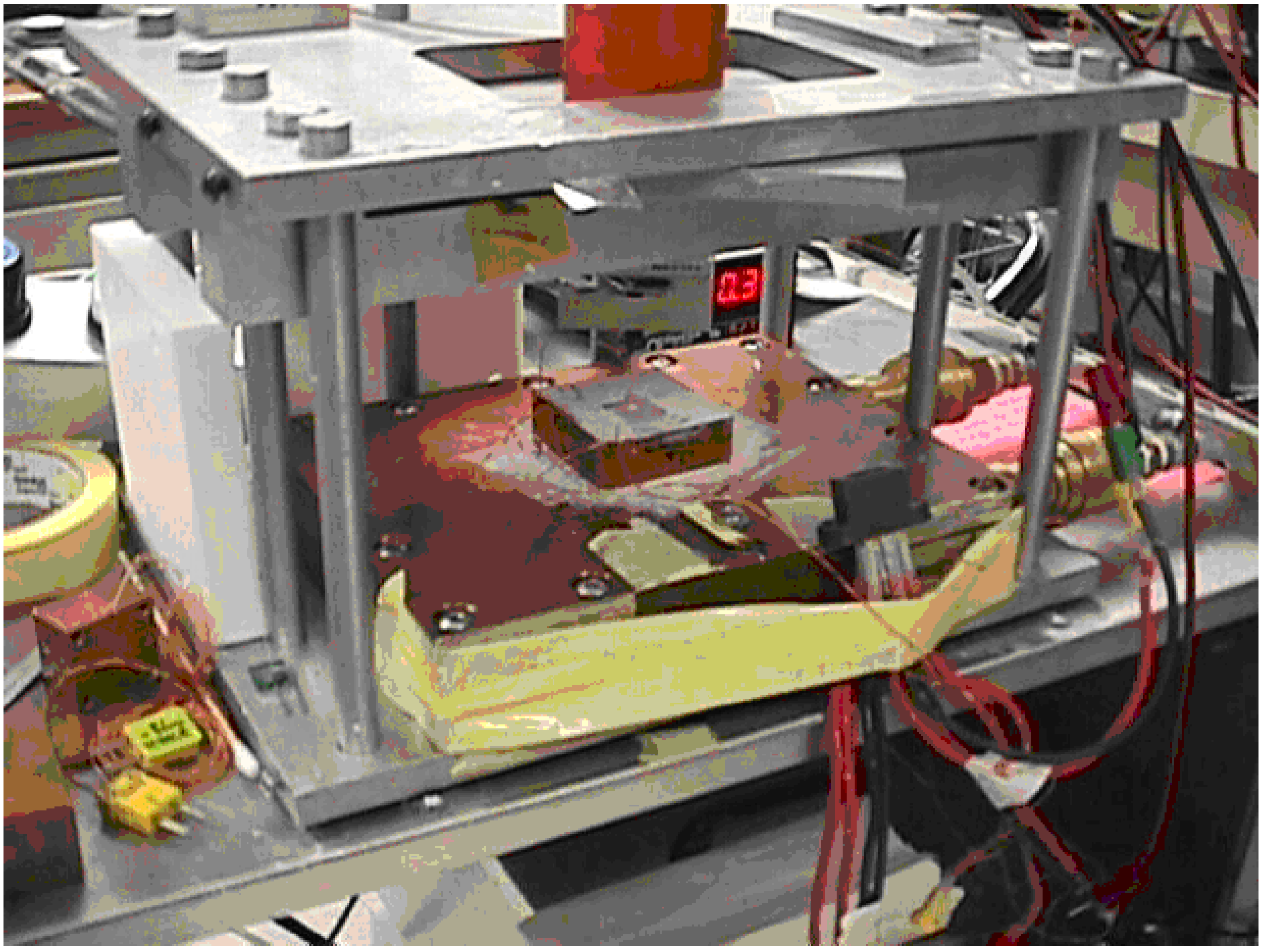

Figure 7. Photograph of Coldplate with Device on 1″ Pad Test Board

For coldplate tests, the semiconductor package is clamped directly onto a copper block 2″ square by 0.75″ thick (Figure 8). Instead of a thermocouple being in direct contact with the back side of the device under test (DUT), the “case” TC is mounted in a narrow well drilled up the center from directly underneath the DUT (see Figure 6). Over the final quarter−inch of depth, the TC well is only about 0.025″ diameter, and comes within about 0.025″ of the surface without breaking through. In other words, this copper block is a close approximation to a pristine semi−infinite heatsink design, though to get there we have sacrificed our ability to measure the case temperature by direct contact. Instead, we must rely on the embedded TC to get as “close as possible” to the case temperature, assuming negligible temperature gradient through the thermal grease and intervening 0.025″ copper. This (almost) solid copper block is in turn clamped against the upper surface of the cold plate (see Figure 6). The Vleer−pin now presses the thermocouple against the underside of the copper block. The TC embedded within the copper block is an approximation of the “case” temperature, while the Vleer−pin TC serves as a “distant” temperature reference. Indeed, measured values of “thermal resistance” for these reference TC’s were only 0.23°C/W, “case” TC to coolant (with standard deviation 0.02°C/W), and 0.03°C/W, Vleer to coolant (with standard deviation 0.01°C/W).

Most high power devices tested on the coldplate are also tested (at much lower power) in natural convection conditions where the device is not specifically heatsinked at all, except by virtue of the high−power test socket in which it is held, typically as shown in Figure 9.

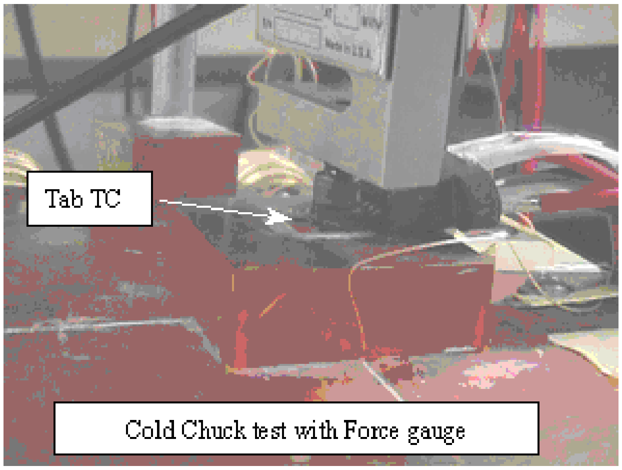

Figure 8. TO−220 Clamped to Coldplate

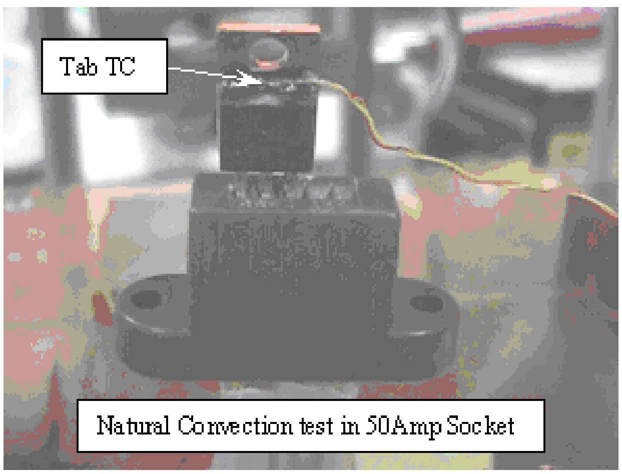

Figure 9. TO−220 in Socket for Free Convection Test

Thermal Test Procedure − Basic Test Method

The basic test method is to heat the device with constant power input, and allow it to reach thermal equilibrium. Heating is accomplished preferably by using the primary heat dissipating structure as occurs in normal device operation. (Likewise, preferably, the same structure serves the purpose of the TSP, so that temperature information is available wherever the heating happened.) This may be the collector−emitter channel in a bipolar device, the source−drain channel in a MOS device (often reverse biased to utilize the so−called “body diode”), the anode−cathode channel in a thyristor, an ESD diode on an output driver of a logic device, etc.

Once heated steady state is reached, “hot” values are recorded for all thermocouples, power is switched off, and the device junction temperature is recorded as the entire system relaxes back to ambient. Power switching is synchronized between the data−logging computer (which controls the test) and a custom−designed switching box. The switching box switches from fully on to fully off in about 10 μs. The first reliable temperature readings (that is, stable voltage readings not obviously contaminated with electrical switching noise) on most junctions are generally not seen until about 250 μsec after the switching box switch is triggered; however, this is somewhat device dependent. These voltage readings are converted to temperatures using the calibration data previously recorded. (As suggested in Figure 1, in many cases the “measurement” circuit is never actually disconnected during heating phase; but in any case, during the measurement phase, the same circuit as was used during calibration is once again used. There are also rare situations when a TSP completely electrically isolated from the heating circuit is available, in which case temperature measurements may indeed be made simultaneous to heating. Obviously, however, this precludes the possibility of measuring the temperature exactly *at* the heat source.)

Transient Results

General Background

The temperature of the junction at the moment of switching off the heating supply is the steady state junction temperature. It is an unfortunate reality that this temperature can never actually be measured (that is, immediately and instantaneously), due to the electrical transients which must settle out before a trustworthy temperature signal can be detected. Much of the art in thermal characterization is therefore the method used in deducing how much the temperature may have dropped during this initial unmeasurable period. In any case, temperature data is collected from the earliest possible moment, as the device cools down to a final steady−state (unpowered) at ambient. The cooling data can be converted to a transient heating curve and an RC network can be determined.

Transient results require a considerable amount of post−processing in order to get the most useful information. First of all, there is the normal amount of experimental variability encountered in any laboratory setting: intrinsic instrument resolution, etc. In semiconductor package measurements, there are also experimental variations attributable to specific aspects of the typical device and environment. For instance, short−time scatter due to die attach irregularities; intermediate time scatter due to board−level soldering, gap between package and board, etc.; long−time scatter due to test board variation (trace thickness) and random convection noise.

Second, as previously stated, the very nature of the measurements precludes ever knowing a “true” initial temperature at the heat source itself. Since junction temperatures in general cannot be measured (at low current) simultaneous to heating the device (at high current), the switching time required to change current opens the door for an unknown amount of cooling to occur before the first reliable temperature measurement can be made. (There is an alternative “heating curve” test method that at least permits the experimentalist to know the starting ambient temperature to any desired degree of precision. An equivalent problem exists, however, because still the junction cools by an unknown amount, each and every time the heating pulse is discontinued for the eventually required measurement.)

Minimizing Scatter

In the “cooling curve” method, the starting steady−state temperature is fundamentally a function of the entire system thermal resistance (i.e. test board variations, thermal grease variations, clamping pressure variations, etc.), let alone individual package manufacturing variations. So even though each of five samples of a certain device might be expected to have very similar transient performance for the first several milliseconds, in practice all their transient cooling curves will begin at different temperatures. (In contrast, the “heating curve” method, by definition, guarantees a common, well defined ambient starting condition which is not at all a function of the test fixturing.) Nevertheless, for the sake of analysis, we’d like to artificially bring each set of these cooling curves together to begin at a common value, and allow the intrinsic system influences to add up to differences in the data towards the “cold,” or steady state, end of the curve. Because our primary interest is in the thermal performance of the package itself, the best approach is to shift the curves (of a single data set representing a single configuration) relative to each other so as to minimize the scatter between them over some designated time range. Typically the range selected is from the earliest possible instant (end of switching noise) out through the end of “die level” influences (based on theoretical expectations arising from die properties). Figures 10 and 11 illustrate this process.

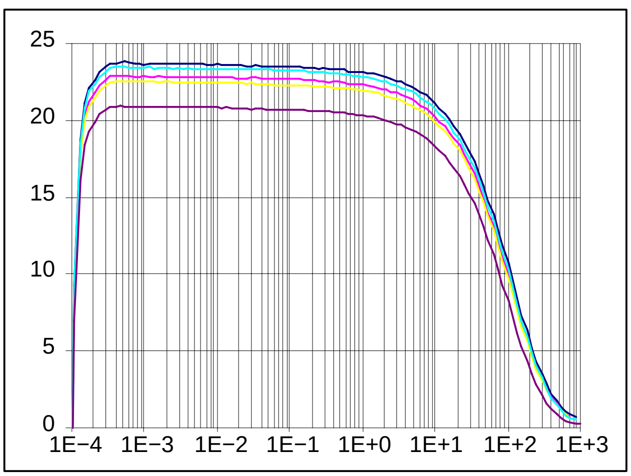

Figure 10. Raw Cooling Curves

For a more detailed explanation, see AND8216/D.

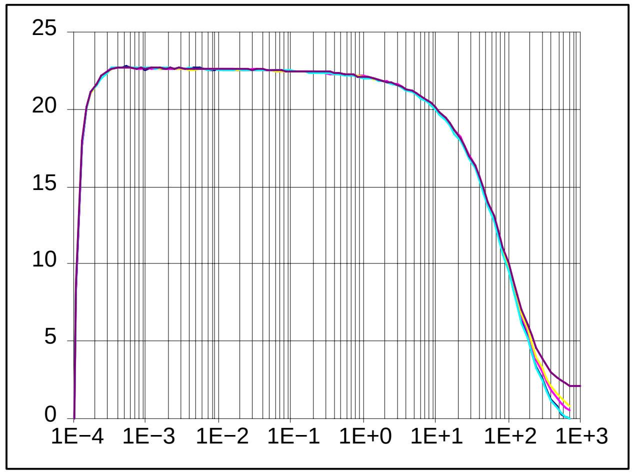

Figure 11. Scatter Minimized from 0.5−2.0 msec

Unmeasurable Short−time Characteristics

To handle the initial “unmeasurable” period in the transient cooling curve, MIL standards suggest that a sqrt(time) assumption be used to back−extrapolate from the first reliable measurements to time zero. (This arises from the closed−form mathematical solution to a uniform heat flux applied to a semi−infinite domain.) We find that if we use a linear regression technique to examine the form of our data between, say, 250 μsec and 1.0 msec, the “correlation coefficient” is only so−so for an assumed sqrt(t) function (say 0.94, where 1.00 is a perfect fit). In fact, over this timescale, the correlation is no worse for our real data if we assume a perfect exponential behavior rather than a sqrt(t), or even a straight−line linear extrapolation. On the one hand, this may be simply because our measurement resolution is not good enough to discriminate the form with this degree of accuracy. On the other hand, it could be because the actual form is neither sqrt(t), nor exponential, nor linear. Consider that these MIL standards were written several decades ago when silicon die thicknesses were undoubtedly much larger, and the sqrt(t) approximation was good out to several milliseconds (assuming uniform surface heating). A 10−mil thick die, however, can only be assumed to exhibit sqrt(t) behavior for about 400 μsec − departing from this idealized behavior at about the same timescale at which the first reliable temperatures can be measured. Compounding matters is that the heat may not be fairly approximated as a surface flux (especially as the die gets thinner) − in fact certain silicon technologies minimize their electrical resistance properties by dissipating large fractions of their thermal energy through the entire thickness of the die. Last, but certainly not least, is the question of whether the “temperature sensitive parameter” being calibrated in the first place is actually at the surface of the die at all, or whether it covers so large an area (or even volume of material) that it is measuring some sort of average temperature, rather than the idealized “peak” junction temperature! So for all these reasons, it is really not clear what should be done to extrapolate back to time zero. Ordinarily we go ahead and use the sqrt(t) fit, recognizing its limitations, but also that it tends to be conservative with respect to power distributions which are not actually surface heating. One can calculate the sqrt(t) constant based on a simple, closed form solution to one−dimensional surface heating of a semi−infinite solid:

$$\theta(t) = \frac{2}{A\sqrt{\pi k \rho c_p}}\sqrt{t} = b\sqrt{t} \quad \text{(eq. 1)}$$

where $$b = \frac{2}{\sqrt{\pi}\,A\eta} \quad \text{(eq. 2)}$$

the quantity $\eta = \sqrt{k\rho c_p}$ is known as the thermal effusivity. (eq. 3)

A = active area, k = thermal conductivity, ρ = density, c_p = specific heat.

This sqrt(t) behavior is followed, within a couple of percent, for about 40% of the "characteristic time" for the continuous medium in question (in this case, through the thickness, L, of the silicon die). The characteristic time is given by:

$$\tau = \frac{L^2}{\alpha} \quad \text{(eq. 4)}$$

where the quantity $\alpha = \frac{k}{\rho c_p}$ is known as the thermal diffusivity. (eq. 5)

As one may surmise from the preceding discussion, the decision making process for handling the "unmeasurable" time period of an experimentally generated thermal transient response curve, is somewhat subjective. It is worthwhile, therefore, to present some additional guidelines and rules of thumb which are considered during this process.

1. Calculate the characteristic time[1] (Equation 4) of a hypothetical piece of silicon of the thickness of the die. For the temperature ranges of interest to us, our standard value for the thermal diffusivity of silicon is 0.0122 msec/mil^2 (in units convenient for the sake of this discussion). A 10−mil thick die thus has a 1.22 ms characteristic time.
2. Calculate the sqrt(t) proportionality (Equation 2) for the full die area, and also for the active die area. Using our standard silicon properties, the thermal effusivity (Equation 3) is 0.0138 W√s /mm^2/°C. Thus, for example, using a full die area of 3.0 x 3.0 mms, you obtain a value for b (Equation 1 and 2) of 9.1°C/W/√s. If the active area is only 2.0 x 2.0 mils, b would be 20.6°C/W/√s.
3. Using the 40% rule of thumb, identify the latest time at which the sqrt(t) proportionality might be expected to be valid. For a 10−mil thick die, that would be 0.5 ms. Compare the die−thickness time scales with the time at which the experimental data is believed to be good, to see whether the actual die properties should have anything to do with it or not. (Obviously "good" is somewhat circular reasoning, but for instance we keep clear of obvious electrical noise in the data, like when it appears that the junction is getting hotter after we kill the power, rather than getting cooler! Clearly the data can't be correct until it at least looks like it's getting cooler, and probably not even then.)

(a) If the data looks clean all the way back into the range of 40% of tau, then we hopefully can reconcile the actual slope of the curve with a theoretical sqrt(t) curve which is based on active die area (if it's in the 40% range, die thickness hasn't come into play yet). For 15−mil thick die this is not unusual.

(b) On the other hand, if it looks like the data is clean back into somewhere between 40% and 100% of tau, we want to suppose that actual die−level thermal characteristics are driving the data, but it isn't so clear whether the theoretical sqrt(t) curve means as much. In this case, we'll play a visual curve matching game trying to convince ourselves that we can make a smooth blending of the sqrt(t) curve into the real data, etc., but see item (4).

(c) Worst case is (like for particularly thin die) when our data doesn't look clean even as early as 100% of tau, or maybe it looks clean but doesn't very much look like the theoretical sqrt(t) at all.

[1]It should be noted in passing that this is not a "time constant" in the strict mathematical sense of an exponential function decay time (because it isn't), but it does relate to how long it takes the heat to pass through a slab of material of the indicated thickness, under the following conditions: surface heating, one−dimensional heat flow through the slab, thermal ground on the "destination" side of the slab.

4. Active area and other factors affecting the 40% rule − the theoretical sqrt(t) curve is based on the assumption of uniform surface heating, and 1−D heat flow. In the real world, obviously that never happens. However, silicon having a much higher thermal conductivity than the mold compound which typically surrounds it on all but the back side, the 1−D approximation is reasonable, at least if 100% of the surface area of the die was heated. But that's rarely the case. Among other things, you've got wirebond pads that may have active silicon underneath ("bond over active" technology), so the heat loss *upwards* is not uniformly small; second, wirebonds may be over *inactive* areas, and there's usually some dead zone around the perimeter to allow for minor variations in saw−cutting the die apart, etc. One can also look at how far apart active zones are (sometimes there are multiple heat sources, especially in large−scale analog devices), and how far away these active regions are from the edges of the die. Finally, one must also look at the width of the heated region as compared to the die thickness. Following a 45−degree spreading angle (for the lack of any better rule), see if the heat gets to the full die dimensions before it gets through the die, etc. The competing effects are that if the heated area is large compared to the thickness, one expects the sqrt(t) proportionality based on active area to resemble that of the actual curve; but if the distance to the edges is small, one expects to reach something more like a sqrt(t) based on the total area by a "tau". On the other hand, if the percent area is small, and if the width of the heated area is small compared to the die thickness, the situation begins to more closely resemble a point source of heat, and the sqrt(t) behavior will never occur in the first place. (The same thing goes for volumetrically distributed heat sources, which never follow a sqrt(t) either.) For small die with relatively large pads and dead zones, it's very subjective. Even so, the short−time problem can be bounded with a sqrt(t) based on active area as the upper bound, and a sqrt(t) based on 100% area as a lower bound. And the earlier in time the data is "clean," the closer together come these two bounds.
5. One must consider the presence of thermally significant materials on the top side of the die. Indeed, we've already glossed over the typical situation, where mold compound encapsulates the die. Though a moderately poor thermal conductor, the effusivity of a typical compound reduces the effective sqrt(t) constant of the surface heating model by about 10%. So the numbers quoted earlier are, strictly speaking, applicable to a "bare" die situation more so than to an encapsulated die. In bond−over−active technology, the wire itself (particularly the ball bond resting on the bond pad) may need to be considered, especially when the die is very small. Likewise, many power devices have a "clip" of some sort (usually copper) in lieu of wires, soldered directly over the active area of the silicon. These materials change the theoretical sqrt(t) proportionality; however, the same % area considerations (meaning, now, the % coverage of the other material) described above still come into play. If the clip is thinner than the die, however, it affects the initial sqrt(t), yet its effect dies out within its own "tau" (which is based on its thickness and alpha), probably long before the data ever became "good." So there still may be a brief period where the die−only sqrt(t) shows up.
6. Finally, in special cases, the heat source may be better modeled as a volumetric heat distribution rather than surface heating, in which case the initial transient response is proportional to linear time rather than sqrt (t).

Finite Element Models (FEM)

If any of these particularly difficult experimental situations arises, whether due to complex geometries, multiple material combinations, or simply when the experimentally unmeasurable period extends well past the expected sqrt(t) behavior based on silicon thickness, we typically supplement the measured data with a finite−element model (FEM) at the die level. Such a model generally includes the die attach material and leadframe immediately adjacent to the silicon. For computational reasons, mold compound above the die is generally omitted from such model, but more thermally conductive materials will be included (especially "clips"), if present.

Last but not least, a die−level FEM is also very useful for adjusting/interpreting short−time data when the TSP is known to be separate from the heated region, or when the practical exigencies of measurement preclude heating of the desired region at all. (In this latter case, you heat what you can in order to measure the overall package characteristics, and rely on the model to fill in what you want to know at the die level.) The following figure illustrates the situation of a very small active area as compared to overall die dimensions, and asymmetrically located on the die surface. No simple short−time model will match the peak temperature transient profile in such a device, so the FEM gives us the best means of predicting the initial transient heating with which to augment the measurements.

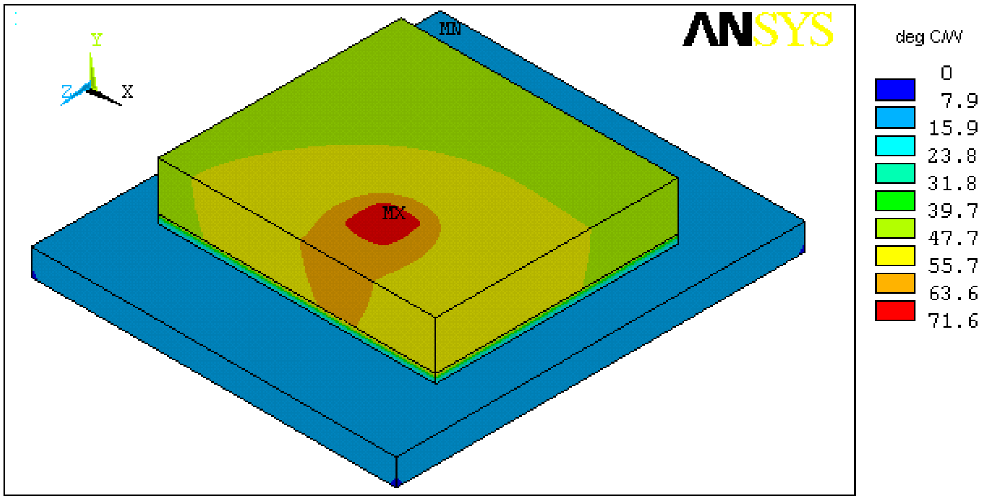

Figure 12. Finite Element Model of Quad−Op−Amp Die on Leadframe, One Output Heated

Steady State Values

Once the initial "unmeasurable" temperature change at the die level has been determined, the overall steady state junction−to−ambient (or junction−to−coldplate) values come from adding the deduced short−time temperature change to the averaged net value from the cooling curves. Additional steady state values for junction−to−board, junction−to−tab, etc. are also calculated from available thermocouple measurements. Clearly, as the "thermal distance" between the junction and the thermocouple in question increases, the overall variability in the steady state result will increase as it includes progressively more of the sources of experimental variability (die attach, board mounting, board/convection variations). In many cases, certain points in the cooling curve may be identified with physical location in the test configuration. For instance, a pronounced "bend" in the cooling curve at about five seconds in a 1″ pad test (Figure 13) corresponds to the heat "arrival" at the edge of the 1″ pad. A bend in the coldplate curve at about one second (Figure 14) corresponds to heat "arrival" at the perimeter of the 2″ copper block. Knowing these characteristics of the test fixturing helps correlate steady state results with the transient results.

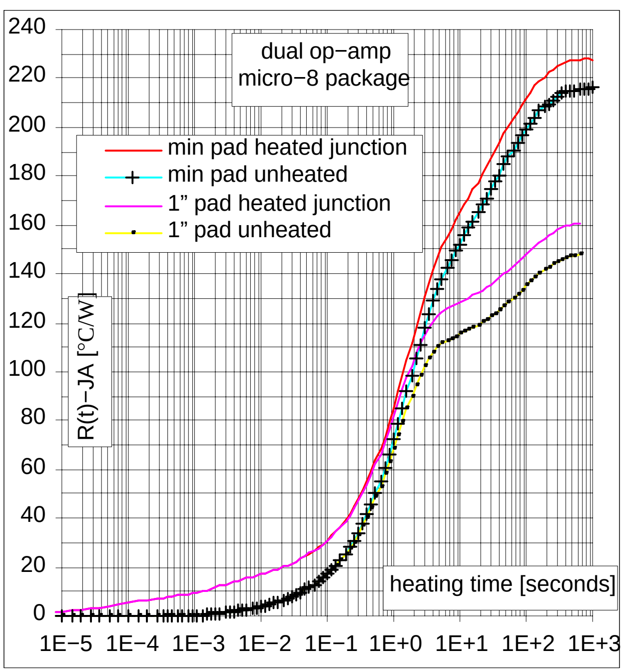

Figure 13. Min Pad vs. 1″ Pad Heating Curves

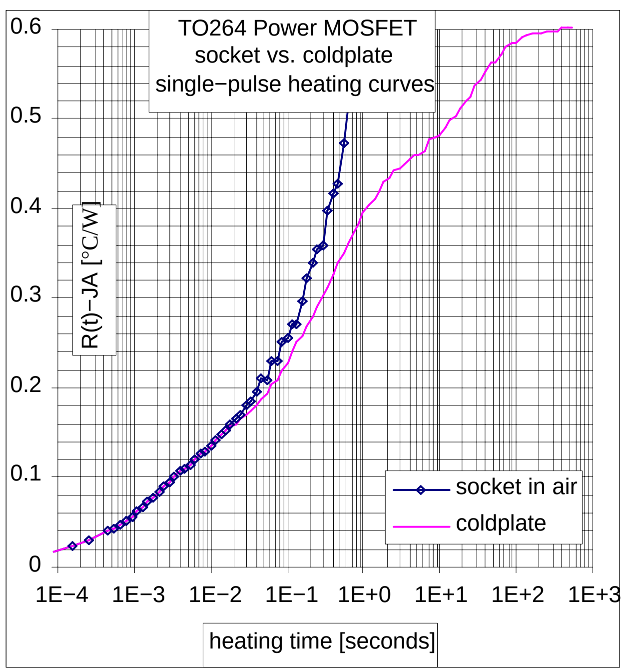

Figure 14. Socket vs. Coldplate Heating Curves

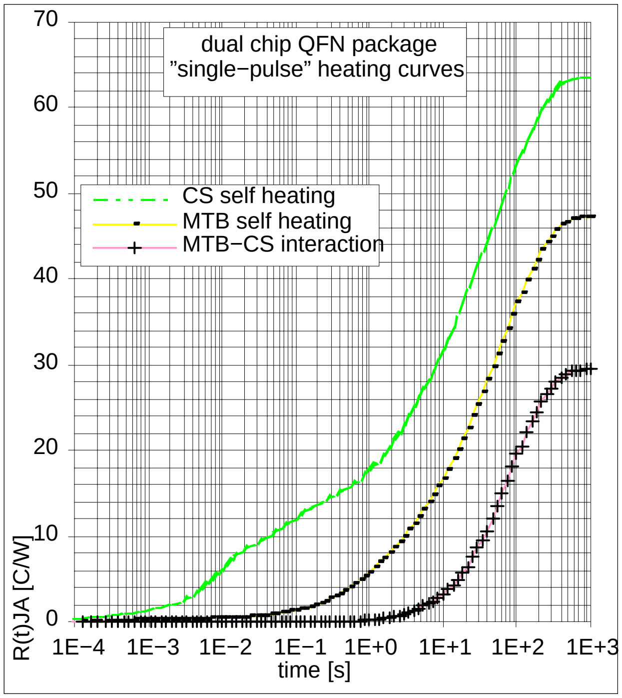

Figure 15. Heating Curves for Two−Chip Device

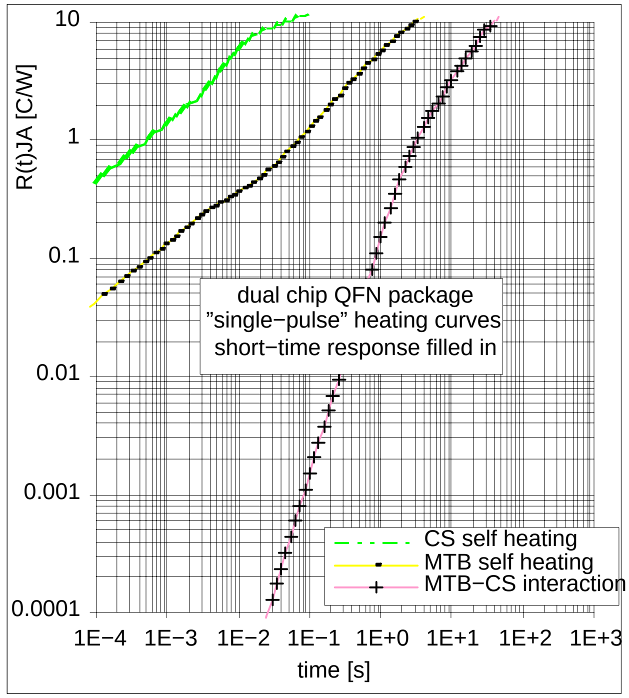

Figure 16. Short−Time Response Filled In

Thermal RC Network Analysis

The complete transient curve for a device (model or theory blended with measurements) is often used for a thermal RC network analysis. Proprietary software provides a best fit model to the data (usually within a fraction of a percent of mean error between RC network and input curve). Any network topology may be selected, but usually a simple "ladder" network of resistors with capacitors tied to ground is chosen (such as in Figure 17). Multiple−junction devices are typically modeled with branched ladder networks, and the input transient curves (examples in Figures 13, 15) must include "self heating" and "interaction heating" for all the junctions represented in the model. Figure 18 depicts a four−input RC thermal network, where only the nodes and connecting resistors are indicated, the (grounded) capacitors being implied.

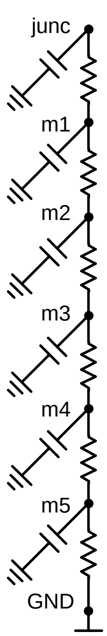

Figure 17. Simple RC "Ladder"

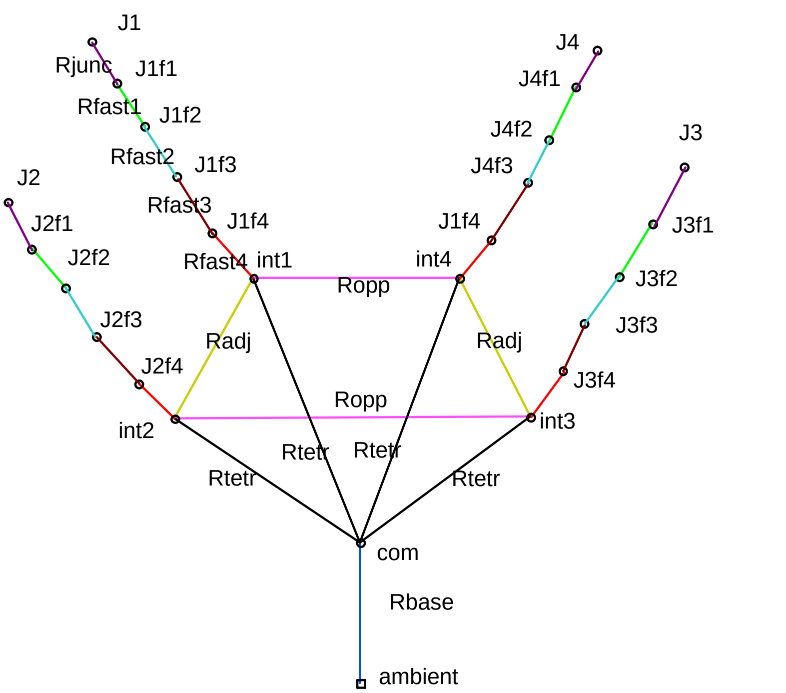

Figure 18. Four Input Thermal RC Network, Capacitors Implied

For the specific non−symmetric dual−chip device whose heating curves were shown in Figures 15 and 16, we illustrate in Figure 19 the RC network used to fit the data; here the capacitors are implied. Figure 20 shows the same network again, only with the capacitors explicitly included. Note that each variable−temperature network node connects to ground through a capacitor.

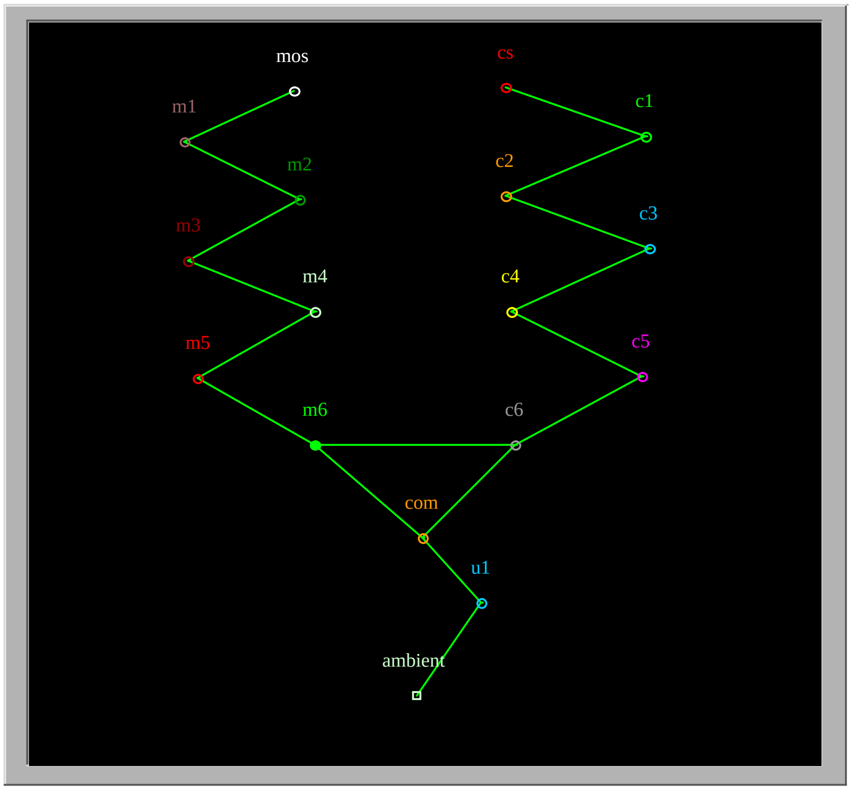

Figure 19. Two Input Thermal RC Network, Grounded Capacitors Implied

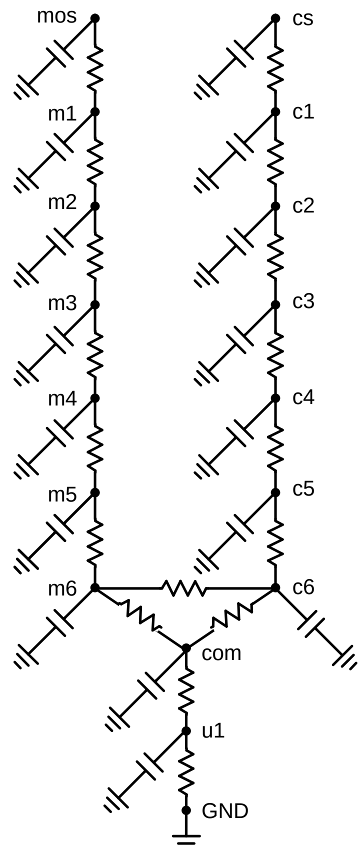

Figure 20. Explicit Capacitor Version

Grounded vs. Non−Grounded Capacitor Networks

It is important to understand that Thermal RC models which have any presumed physical significance at all, that is, where nodes in the model bear some correlation with physical locations in the thermal system, must utilize grounded capacitors. Even so, the thermal behavior of the heat−input junction of the network can be described by a mathematical equation (or set of equations) which resemble an RC network whose capacitors are not tied to ground, but are connected between nodes in parallel with the resistive links. In this purely "mathematical" version of the network, amplitudes of the individual circuit "elements" correspond to resistances, and the time constants of the mathematical terms may be interpreted as RC products; thus the capacitances can be "deduced" from the time constants, though neither they, nor the resistors, have any particular correlation with physical locations in the actual thermal system.

To try to clarify this point (which is *not* what Figures 19 and 20 were illustrating), the following table (Table 1) provides the R and C values corresponding to the data in Figures 15 and 16 and the network shown in Figures 19 and 20. Table 2 provides the mathematical time constants and amplitudes corresponding to the set of equations which describe the thermal performance of the two junction nodes, both their isolated "self heating" and their "interaction heating." Note that the R's in Table 1 and the amplitudes of Table 2 bear little resemblance to each other.

Table 1. A SPICE Compatible Thermal RC Model

Element Name	Node 1	Node 2	Element Value
C_C1	gnd	cs	9.21909E−6
C_C2	gnd	c1	4.36252E−5
C_C3	gnd	c2	1.30876E−4
C_C4	gnd	c3	1.75727E−4
C_C5	gnd	c4	8.87879E−4
C_C6	gnd	c5	2.06324E−2
C_C7	gnd	c6	3.25284E−1
C_C8	gnd	mos	9.59163E−5
C_C9	gnd	m1	4.53881E−4
C_C10	gnd	m2	1.36164E−3
C_C11	gnd	m3	4.08493E−3
C_C12	gnd	m4	1.66362E−2
C_C13	gnd	m5	6.00462E−2
C_C14	gnd	m6	4.60781E−1
C_C15	gnd	com	3.03315E+0
C_C16	gnd	u1	2.51716E+1
R_R1	cs	c1	3.97030E−2
R_R2	c1	c2	1.19109E−1
R_R3	c2	c3	3.57327E−1
R_R4	c3	c4	2.59622E−1
R_R5	c4	c5	8.51408E+0
R_R6	c5	c6	6.55035E+0
R_R7	c6	com	3.92892E+1
R_R8	mos	m1	3.81609E−3
R_R9	m1	m2	1.14483E−2
R_R10	m2	m3	3.43448E−2
R_R11	m3	m4	1.03035E−1
R_R12	m4	m5	2.57613E−1
R_R13	m5	m6	5.90459E+0
R_R14	m6	com	2.40044E+1
R_R15	c6	m6	5.31662E+1
R_R16	com	u1	1.72692E+1
R_R17	gnd	u1	4.35939E+0

Table 2. Mathematical Model − Roots and Amplitudes for $R(t) = \sum_{i=1}^{n} \left(\sum_{j=0}^{m} A_{ij} t^j \right) e^{r_i t}$

		MTB40N10E Self Heating		CS5342 Self Heating		Interaction	
i	r_i	A_{i0}	A_{i1}	A_{i0}	A_{i1}	A_{i0}	A_{i1}
1	−3.34538E+6	−2.51600E−3	2.88500E−3	−2.61750E−2	−3.00200E−2	−3.76038E−46	−9.55148E−41
2	−2.27628E+5	−6.15500E−3	−1.12100E+0	−6.36000E−2	1.16007E+1	−4.70539E−32	−9.49738E−28
3	−4.66437E+4	−5.60013E−17		−3.34200E−2		7.92060E−26	
4	−2.25273E+4	−1.83010E−2		1.00724E−15		−6.31528E−23	
5	−9.77316E+3	2.79361E−17		−4.13157E−1		3.43633E−20	
6	−2.18916E+3	−6.57940E−2		−4.03874E−16		−6.66750E−16	
7	−2.29490E+2	−1.77952E−1		−6.91961E−14		6.02476E−10	
8	−9.99632E+1	2.51105E−15		−7.68731E+0		−3.73658E−8	
9	−7.42153E+0	−1.50758E−7		−6.73210E+0		1.00700E−3	
10	−2.42893E+0	−4.24632E+0		−1.89000E−4		−2.83180E−2	
11	−1.62443E−1	−5.32689E+0		−1.09936E+1		7.65256E+0	
12	−1.02279E−1	−7.89576E+0		−8.32867E+0		−8.10933E+0	
13	−1.83258E−2	−1.31004E+1		−1.30787E+1		−1.30896E+1	
14	−6.96478E−3	−1.61600E+1		−1.61463E+1		−1.61531E+1	
15	0.00000E+0	4.70001E+1		6.35032E+1		2.97268E+1	

Notes for Table 2: In this specific case, $n = 15$ (number of “roots”, including the constant) and $m = 1$ (maximum root multiplicity less one). The “roots” here are the coefficients in the exponents of the exponential function, so called due to the method of solution by which they are obtained (being the roots of a polynomial, the determinant of the system of equations representing the transient response of the RC network). As such, they are related to the “time constants” of the exponential responses of the network, through the expression $\tau_i = -\frac{1}{r_i}$, hence the equation can be rewritten

$$R(t) = \sum_{i=1}^{n} \left(\sum_{j=0}^{m} A_{ij} t^j \right) e^{-\frac{t}{\tau_i}}.$$

All nodes of a particular network share the same set of time constants, but the amplitudes of the terms corresponding to each time constant differ from node to node. In certain cases, roots are repeated, from which arise the higher order ($j > 0$) terms t^j. Note that most of the higher order A_{i1} coefficients are exactly zero (hence absent from the table); some others are effectively zero and the associated term may therefore be omitted if convenient. (For instance, terms $i = 1$ through $i = 8$ for the “interaction” equation are certainly negligible.) Also, in determining which higher order terms are negligible, it may be useful to note that in general, the maximum of any term $y_{ij} = A_{ij} t^j e^{r_i t}$, is $y_{max} = A_{ij} \tau_i^j e^{-j}$, and it occurs at time $t = j\tau_i$.

Short−time Limits of RC Networks

A significant limitation of RC networks is that for times shorter than the approximated RC time constant of the junction node (which, for a well defined network, will be the fastest node), transient response falls off linearly (going back in time) rather than according to the previously described sqrt(t) constant. This means that if the RC network model is exercised for power fluctuations on the order of (or faster than) this minimum response time, a thermal RC network will vastly underestimate the sqrt(t) response demanded by surface heating theory. (Interestingly, if the RC model is extended to such short times and implied physical dimensions as correspond to locally volumetric heating, *volumetric* heating theory demands *linear* temperature increases with time. In other words, carried to times too short for the *surface* heating model to apply, the RC network model actually becomes “correct” again. However, this requires that one have an input transient curve which is known to exhibit the correct behavior over the entire time scale of interest, whether this be linear at the shortest times, or sqrt(t) at the shortest times. (Observe, for instance, the 1:2 slope in the 1E−4 to 1E−3 s range of the responses of the heated junctions in Figure 16.) The RC response will not magically become linearly correct just because it *fails* to follow the sqrt(t) where that is appropriate. The point, here, is that one must be cautious in the use of an RC network model which is not designed to provide sqrt(t) behavior, over a time scale in which sqrt(t) behavior is expected. If one must work in this time scale without a suitable RC network, Equation 1 should be used directly in lieu of the RC model. For more information, see also AND8218/D.

For further information on Thermal Resistance Measurements:

1. MIL−STD−883E−1012.1, Test Method Standard, Microcircuits, Thermal Characteristics.
2. EIA/JEDEC, JESD24−3, 24−4, 51, 51−1 through 51−8.
3. R.R. Tumala, E.J. Ramaszewski, A.G. Klopfenstein, *Microelectronics Packaging Handbook*, 2nd Ed., Chapman & Hall, New York, 1997.
4. “Basic Semiconductor Thermal Measurements,” Application Note #AN1570, ON Semiconductor, <www.onsemi.com>.
5. Y.L. Xu, R.P. Stout, D.T. Billings, “Electrical Package Thermal Response Prediction to Power Surge” ITHERM, May 2000.
6. R.P. Stout, D.T. Billings, “How to Extend a Thermal−RC−Network Model to Respond to an Arbitrarily Fast Input,” IEEE CPMT−A Journal, Dec. 1998.
7. Michael Pecht, *Handbook of Electronic Package Design*, Marcel Dekker, New York, 1991.
8. Tony Kordyban, *Hot Air Rises and Heat Sinks (Everything You Know About Cooling Electronics is Wrong)*, ASME Press, 1998 (ISBN 0−7918−0074−1).

PUBLICATION ORDERING INFORMATION

LITERATURE FULFILLMENT:
Literature Distribution Center for ON Semiconductor
P.O. Box 61312, Phoenix, Arizona 85082−1312 USA
Phone: 480−829−7710 or 800−344−3860 Toll Free USA/Canada
Fax: 480−829−7709 or 800−344−3867 Toll Free USA/Canada
Email: orderlit@onsemi.com

N. American Technical Support: 800−282−9855 Toll Free USA/Canada

Japan: ON Semiconductor, Japan Customer Focus Center
2−9−1 Kamimeguro, Meguro−ku, Tokyo, Japan 153−0051
Phone: 81−3−5773−3850

ON Semiconductor Website: http://onsemi.com

Order Literature: http://www.onsemi.com/litorder

For additional information, please contact your local Sales Representative.