

How To Use Thermal Data Found in Data Sheets

Packaging Technology Development

Prepared by: Roger Paul Stout, PE
ON Semiconductor



ON Semiconductor®

<http://onsemi.com>

APPLICATION NOTE

Table of Contents

Definitions	2
Steady State Data	2
Theta (θ) and Psi (Ψ) Numbers	2
Theta (θ) values	2
Psi (Ψ) values	4
Multiple Junction Devices and Matrix Formulations	5
Theta–JA vs. Copper Area	6
Transient Data	7
Heating Curves	7
Duty Cycle Curves	10
Thermal RC Network Models	11
Grounded Capacitor (Cauer) Networks	12
Non–Grounded Capacitor Networks	12
Comparison and Contrast of Cauer and Foster Ladders	12
Multiple–Junction Devices and Transient Response	13
Circuit Simulators	13
Spreadsheet Models	15
RC Models and Short–Time Transient Response	15
Periodic Waveforms Using Foster RC Models	17
Surface Heating, Square–Root of Time, and Short–Time Transient Response	18

Definitions

ambient, T_A = where all the heat ends up, the environment thermally “far” from the device (not to be confused with the case temperature, which may be vastly different from the air temperature just millimeters away).

case temperature, T_C = representative point on the external “case” of the device, location must be well defined along with use of any parameters based on this value.

junction, T_J = the hottest point inside the semiconductor device.

psi-JT, $\Psi_{JT} = \frac{T_J - T_T}{P_d}$ thermal characterization parameter, measured junction to case top (T_T)

psi-Jx, $\Psi_{Jx} = \frac{T_J - T_x}{P_d}$ thermal characterization parameter, measured junction to location defined (T_x)

psi-xA, $\Psi_{xA} = \frac{T_x - T_A}{P_d}$ thermal characterization parameter, measured from location defined (T_x) to ambient

theta-JA, $\theta_{JA} = \frac{T_J - T_A}{P_d}$ overall thermal resistance of device plus external system

theta-JC, $\theta_{JC} = \frac{T_J - T_C}{P_d}$ ideally, thermal resistance of just the device as measured to the case

P_d = total power dissipation of device

min-pad = in reference to a thermal test board, a board having only the minimum amount of metal pads and traces required to mount the device and carry power and signals to and from the device; the traces may actually have significantly more area than the mounting pad itself, and the total size of the board (typically 3” square), and its thickness (0.062”) may be significantly different than what will be used in an actual application. These variables are only some of those that render min-pad data sheet values of limited use in a real application environment.

1” pad = in reference to a thermal test board, a board having a nominally 1” square area of copper plating, at the center of which is mounted the package; typically the additional trace area required to carry power and signals to and from the device will be a small fraction (<10%) of the pad; but for larger devices, e.g. a D2pak, the actual heatsink itself may be a significant fraction of 1-sq-in., so the difference between the min-pad and 1” pad values will not be as large for large devices as it is for tiny devices. Copper spreader thickness (typically 1-oz, meaning 1-sq.-ft of the material in that thickness would weigh 1-oz), and overall board size (typically 3” square) and thickness (0.062”) will potentially make the value significantly different than what would be experienced in an actual application. These variables are only some of those that render 1” pad data sheet values of limited use in a real application environment.

Steady State Data

By “steady state”, we mean operating conditions wherein power dissipation in each relevant device, has been constant for a long enough period of time that temperature changes are no longer occurring. Starting from zero power, with all temperatures initially at ambient, the sudden application of constant non-zero power will result in monotonically increasing temperatures. The highest temperatures will eventually be reached, therefore, at steady state. Steady state thermal data is often provided in the form of specific values of thermal resistance or impedance. In addition, other charts may be presented showing how steady-state thermal characteristics typically depend on certain external conditions, such as the amount of heat spreading metal that has been provided on the application board for the specific device in question. In the case of multiple-junction devices, there may be a matrix form of the steady state thermal characteristics as well.

Theta (θ) and Psi (Ψ) Numbers

Theta (θ), sometimes denoted R_θ , values are true “thermal resistances.” That is, they tell you that if you know the temperatures at two points (connected by the thermal resistance) then the amount of heat that flows from one point to the other is completely determined by that resistance. Conversely, if you know the heat flow along the path, and you know its resistance, then you can predict the temperature difference that will result due to this heat flow. If there are other heat paths in the system, they have their own characteristics, and they are independent of what happens along the particular path of interest. Typical units would be °C/W.

In the context of semiconductor packages and devices, there are usually going to be, *at most*, two “true” thermal resistances, theta-JA, and theta-JC, and these must be defined carefully. But the single most important fact about these values (what, indeed, makes them “theta”s), is that the total power dissipated by the device flows between the two “points” being described (the junction being one “point,” and either ambient or the case temperature being the other “point”). That is, there are no extraneous, parallel thermal paths in the system allowing some of the heat to “leak” away. All heat leaving the junction, the first point, eventually arrives at or passes through the other point – either ambient or case, respectively.

Mathematically, we define these two quantities as follows:

$$\theta_{JA} = \frac{T_J - T_A}{P_d} \quad (\text{eq. 1})$$

$$\theta_{JC} = \frac{T_J - T_C}{P_d} \quad (\text{eq. 2})$$

Knowing the appropriate values, actual operating junction temperature therefore could be predicted according to:

$$T_J = \theta_{JA} \cdot P_d + T_A \quad (\text{eq. 3})$$

or

$$T_J = \theta_{JC} \cdot P_d + T_C \quad (\text{eq. 4})$$

For theta-JA, obviously, by definition, all the power leaving the junction eventually arrives at ambient, hence the ratio of this temperature difference over the total package power, is a true system thermal resistance. A data sheet may provide one or more theta-JA values for different representative mounting situations. For instance, min-pad board values, and 1" pad board values; or perhaps a chart of theta-JA vs. copper spreader area. But these values may not apply to a real application, even if the copper area is "correct." Other variables, such as the presence of other power dissipating devices, the nature of the air flow conditions, the thickness and detailed layout of the spreader itself, will all affect the value.

For theta-JC, the two points are the junction (J), and the "case" (C) – it's the definition or selection of the "case" temperature that makes the challenge here. If we are justified in making the assumption that 100% of the dissipated power actually flows past the "C" point we've defined, then again, the ratio of this temperature difference to the total package power, is a true thermal resistance. Typically the only test situation that results in a reasonable approximation of this 100%-of-heat condition is a coldplate test, and that being for power packages clamped directly onto the coldplate. (Even in such a test, from 1-10% of the power may "leak" out through other paths, depending on the particular package design and test fixturing.) One must still carefully designate the location of the "case" measurement. For a good theta-JC measurement, the "case" is usually defined as the center point of the heatsink at the heatsink/coldplate interface, which will be the hottest point on the coldplate, but not necessarily the coldest point on the "case" of the device at all. One practical difficulty in making a good theta-JC measurement, is doing so without disturbing the heat flow. A groove in the surface of the coldplate may reduce the interface area significantly; holes bored through the coldplate disturb the heat flow as well, though if small enough, perhaps not as much. Another issue is the question of how well the case-measuring thermocouple makes contact with the case, when it can't be seen. As an alternative, measuring the temperature at the exposed edge of the heatsink (say on the tab of a Dpak or TO220) may circumvent both these difficulties, but it may yield a significantly different result (perhaps 20-40% higher than the "true" theta-JC value). Clearly, to use a theta-JC value successfully in a real application environment (i.e. to predict the actual junction temperature), the application must assure that virtually 100% of the device power flows through the case. Finally, a theta-JC value as measured on a coldplate, may be vastly different than a corresponding "psi-JC" value, even through the thermocouple location may be

identical for both measurements (see following discussion on "psi" values). This arises because the fraction of the heat passing through the "case" point will very likely be substantially less in a non-coldplate test setup, as compared to the coldplate setup. Indeed, if the theta-JC value was derived from a 100%-of-heat condition, it is axiomatic that any corresponding psi-JC value will be lower. For instance, if only 10% of the heat flows past the "case" in a non-coldplate mounting situation, then psi-JC will be one tenth of the theta-JC value!

Clearly, the difference between theta-JA and theta-JC is that theta-JA necessarily includes the entire system, not just the package, whereas theta-JC is idealized as a "package only" property. It is not at all unusual for theta-JC to be a small fraction of theta-JA, which is a way of saying that the thermal design of the external environment is more important in determining the operating junction temperature of a device, than is the thermal design of the device itself.

The problem with these values appearing on data sheets, is that theta-JA will very likely not apply to a particular customer application (because there will be differences in the system external to the package, such as air flow conditions, amount of metal thickness, area, and layout in the board, proximity and power dissipation of neighboring devices, and so forth). So theta-JA may *seem* convenient, because all you need to know is the ambient temperature. In reality, unless the application is exactly like the thermal test situation, a different theta-JA will apply, and the difference may be substantial. Data sheets may present either or both "min pad" values and "1" pad values", but the real theta-JA in an application may be better than the "1" pad" value, or it may be worse than the "min-pad" value. In any event, if the system is different, the data sheet's theta-JA is not really a useful value.

Theta-JC *may* be more useful, in the sense that it *may* actually describe the device's characteristics in the real application (including the requirement that essentially 100% of the heat pass through the identified case location). Even so, it is truly useful only if the external system is fully defined as well. The problem here is that one cannot simply assume that the case temperature can be controlled to an arbitrarily chosen value; rather, the design of the external thermal system must ensure that this is true for the amount of power being dissipated by the device. For example, consider a particular TO264 power transistor, with a theta-JC of 0.4°C/W. If max Tj is 150°C and the case could be held at 25°C, then one could *in principle* dissipate 312.5 W [$P_d = (T_J - T_C) / \theta_{JC}$].

However, what sort of external system can "hold" the case at 25°C? How about a water-cooled coldplate capable of sinking 400 W, whose thermal resistance is approximately 0.2°C/W, measured from the mounting point on the coldplate to the "infinite" coolant supply? To sink 312.5 W implies that the coolant itself would have to be held at (0.2°C/W * 312.5 W, or) 62.5°C *below* case temperature, or -37.5°C! In fact, a realistic "real world" application of this

TO264 device might utilize a forced-air-cooled heatsink, with a hundred square inches of surface area, and a net thermal resistance of 0.2°C/W (much similar in capacity to the water-cooled coldplate just illustrated). But in this real-world system, it is *ambient*, not the device case, that is limited to 25°C. Since the total system resistance (theta-JA) is 0.6°C/W (0.4 for the device, plus 0.2 for the heatsink), actual maximum power dissipation is really 208 W, and the case temperature at steady-state equilibrium will be about 66°C.

Psi (Ψ) values, as contrasted with theta values, are not really thermal resistances, though they have the same units. JEDEC defines the term as a “thermal characterization parameter.”¹ It is nothing more than the ratio of the temperature difference between two selected points in a system, and the total power dissipation of the device in question. The equation defining it is essentially identical to that for theta, that is:

$$\Psi_{Jx} = \frac{T_J - T_x}{P_d} \quad (\text{eq. 5})$$

and

$$\Psi_{xA} = \frac{T_x - T_A}{P_d} \quad (\text{eq. 6})$$

observe

$$\Psi_{Jx} + \Psi_{xA} = \frac{T_J - T_x}{P_d} + \frac{T_x - T_A}{P_d} = \theta_{JA} \quad (\text{eq. 7})$$

Note that we have defined two variations, one referring the junction to some arbitrary package location *x*; the other referring the arbitrary package location *x* to ambient. This misleadingly suggests that the former is mainly a “package” characteristic, and the latter mainly an “environment” characteristic. The reality is that the chosen package location merely *arbitrarily* divides the overall system theta-JA into two pieces that are guaranteed to add up to the correct total (Equation 7). It does *not* follow that location *x* will have a predictable temperature between the two endpoints as the environment changes around it; it will only be predictable (from *either* endpoint) if the environment does not change. (Contrast this with theta-JC, which, at least ideally, will always yield a predictable temperature relative to the junction, regardless of what happens to the environment beyond the case.) Be that as it may, just as with theta values, knowing the appropriate inputs, operating temperatures might be predicted according to:

$$T_J = \Psi_{Jx} \cdot P_d + T_x \quad (\text{eq. 8})$$

or

$$T_x = \Psi_{xA} \cdot P_d + T_A \quad (\text{eq. 9})$$

In the context of semiconductor devices and packages, common data sheet psi values include psi-JL_n (where a particular lead, *n*, is designated), psi-JT (where T represents the case top), psi-J-tab (where the tab is the exposed heatsink tab on a suitable power device), and psi-J-board (where perhaps the board directly underneath the center of the package is designated, for instance for a BGA style package).

It is usually possible to know the total power dissipation of the device in question, but it is far more difficult to know what fraction of the heat flows out through the case top, vs. through the leads, vs. through the air gap under the package, and so forth. Though it may be feasible to make temperature measurements at myriad locations all over a package, it is difficult or impossible to make actual heat flow measurements along selected paths. Further, those paths may vary dramatically in their actual thermal resistances as connected to the junction, and in their individual sensitivity to external changes. Therefore, as mounting conditions vary, the relative heat flow along the various possible paths may shift significantly. Psi values reported on data sheets, therefore, must only be used for application temperature estimates, when it is known that the heat flow is similarly distributed. The *minimum* stipulation for valid application of a psi value is that the same fraction of heat flows along the *particular* psi-path as occurred during the lab measurement (regardless of how much difference there might be along other paths); generally this will be difficult or impossible to ascertain.

Consider a particular example of a 2-leaded axial device. It was measured on a thermal test board having a symmetric layout with equal amounts of trace metal to each lead. In this test scenario, theta-JA was 45°C/W, and psi-JL was 15°C/W (same value to each lead, due to symmetry).² Now in a particular application, the device is mounted on a board with a 1” square pad allocated to just one of the two leads; the other lead has minimal traces (identical to the original test board). Measurements made on this application board now yield theta-JA of 31°C/W, and extremely non-symmetric values of psi - JL1 = 21°C/W and psi - JL2 = 9°C/W. Depending on which lead temperature were to be used as the reference point, it should be clear that if the data sheet value for psi-JL was used to predict junction temperature, it would either result in extremely high or extremely low values. (Obviously, if psi values were measured for the actual application, there would be no reason to use the data sheet value. The emphasis here is on the extreme difference that the actual psi values have as compared to the data sheet, even though the package is the same, and one of its leads is even mounted exactly as during the data sheet measurements.) If one returns to the concept of true thermal resistance, it may be seen that in the original

² It is unfortunate in this particular example that the data sheet was pre-1995, and referred to the value as “Thermal resistance, junction-to-lead,” with no explanation of how to apply the value. So is it a *psi*, or a *theta*? Nothing like a small two-to-one discrepancy in resulting calculations!

¹ *Guidelines for Reporting and Using Electronic Package Thermal Information, EIA/JESD51-12*, Electronic Industries Association, 2005.

lab measurements, the true θ_{JA} value would have been 30°C/W, as each lead carried exactly one half the total power.

$$\text{that is, } \theta_{JL} = \frac{T_J - T_L}{\left(\frac{P_d}{2}\right)} = 2\Psi_{JL} \quad (\text{eq. 10})$$

In fact, knowing that θ_{JA} value would allow one to predict junction temperatures precisely, in any application environment, however non-symmetrically the heat spreading metal was arranged, so long as both lead temperatures are measured instead of relying on just one, specifically:

$$T_J = \frac{\theta_{JL}P_d + T_{L1} + T_{L2}}{2} \quad (\text{eq. 11})$$

Using Equation 10, it may be of interest to rearrange this expression into:

$$T_J = \Psi_{JL}P_d + \text{avg}(T_{L1}, T_{L2}) \quad (\text{eq. 12})$$

Thus in this particular two-lead device example, it is seen that even ψ_{JA} may be used to advantage if both lead temperatures are measured in the actual application. In more complex package situations, however, there may be multiple leads, some of which have direct connections to the internal “flag” (where the silicon is mounted) and some don’t, or there may be other significantly non-symmetric thermal paths. Then, changing the amount of metal trace area associated with individual leads, or adding an external heatsink to the exposed case, may drastically modify the relative heat flows as compared to those present when the data sheet values were determined, and may thus completely invalidate the use of the published ψ_{JA} values. As a rule, when it comes to choosing which lead to characterize for a data sheet, it is preferred to pick the one that carries the largest fraction of the power, if such a lead can in fact be identified. This may well be the one with the largest ψ_{JA} , especially if it has a direct internal connection to the flag and has been given a disproportionately large heat spreader on the outside. Such a value will be the least sensitive to modest changes in the actual application environment. On the other hand, if a data sheet specifies a lead which is known *not* to have the most heat flow, then one cannot be sure in a particular application, whether the actual ψ_{JA} value will be higher or lower than that provided in the data sheet (the two-lead example above bearing witness).

Multiple Junction Devices and Matrix Formulations

In referring to multiple junction devices, we are generally referring to devices that contain relatively independent electrical components, for which the ratio of power dissipation between the various possible “junctions” may vary widely. This might be an analog device with two different regulated voltage outputs, driving vastly different loads (that from one application to another may differ substantially). It might be a dual-rectifier package, which in one application utilizes both channels fairly equally, but in another application perhaps utilizes one in preference over the other. It might be a single application that from time to

time moves between two vastly different operating points (different power distributions); which condition is “worst case” needs to be determined. So long as constant power conditions are considered, a matrix approach utilizing steady state values is a concise method of describing the system’s thermal characteristics. It relies on the principle of linear superposition, which states that the temperature rise at any given point in the system is the sum of the independently derived temperature increases attributable to each heat source in the system. Stated in matrix form, we would say it this way:

$$\begin{Bmatrix} \Delta T_{J1} \\ \Delta T_{J2} \end{Bmatrix} = \begin{pmatrix} \theta_{11} & \Psi_{12} \\ \Psi_{12} & \theta_{22} \end{pmatrix} \begin{Bmatrix} P_{d1} \\ P_{d2} \end{Bmatrix} \quad (\text{eq. 13})$$

Note that the matrix notation is simply a shorthand version of the following pair of equations:

$$\begin{aligned} \Delta T_{J1} &= \theta_{11} \cdot P_{d1} + \Psi_{12} \cdot P_{d2} \\ \Delta T_{J2} &= \Psi_{12} \cdot P_{d1} + \theta_{22} \cdot P_{d2} \end{aligned} \quad (\text{eq. 14})$$

(Note that if every matrix element is referenced to ambient, for instance, then each temperature rise so computed should be added to ambient, in order to predict the actual temperatures resulting from the two applied power dissipation values). For a two junction device, this means that the temperature rise of the first junction is the sum of its “self heating” characteristic, times its own power dissipation, added to the “interaction” characteristic (how much each junction heats the other) times the power dissipation of the other junction. Another principle applicable to linear systems is known as *reciprocity*, which states that the amount by which one junction *heats* another (in terms of temperature rise per unit power input) is equal to the amount that it will be *heated by* the other, hence the symmetry of the matrix across the main diagonal.

Observe that here we have used ψ_{JA} to represent the interaction terms. This is strictly correct because in general we cannot say that any particular fraction of the heat dissipated at either junction is passing “through” the other junction; *some* of it certainly flows that way, causing the temperature to rise there (indeed, if it did not, the interaction term would be zero and the two junctions would be completely thermally independent). On the other hand, we have used θ_{JA} for the self heating terms (along the main diagonal). This is *not* strictly correct unless we have the assurance stipulated in the preceding discussion of θ_{JA} values, that is, the reference point to which θ_{JA} refers is the ultimate destination of 100% of the power dissipated by (in this case) both junctions. However, even the basic matrix formulation itself, as illustrated above, is an oversimplification of the required mathematical description if multiple reference temperatures are needed.

Suffice it to say, for idealized simple situations, a matrix formulation typically will consist of θ_{JA} self heating and ψ_{JA} interaction heating values, or possibly θ_{JB} and ψ_{JB} values, where a board location central to the package is identified and *presumed* to represent the major

heat flow path and a common reference temperature. One way of describing the situations in which this simple matrix approach is valid, would be to say that as long as there is a single temperature *boundary condition* (for instance, ambient, or board), the formulation is applicable. (Then that temperature reference, or boundary condition, becomes the common value to which all temperature increases are related.) But in typical situations where lead, board, or case-top temperatures have been characterized using psi parameters, these additional temperature locations are not true boundary conditions, rather auxiliary reference points where the temperature is measured but not controlled.

Even with this restriction, it should be noted that a matrix description can be readily extended to any number of junctions and auxiliary temperatures, where every heat source of interest yields a self heating theta, and a set of interaction psi's, for every other junction and point of interest. For example, if we have three heat sources, and one lead and one board temperature reference location, we would have:

$$\begin{Bmatrix} \Delta T_{J1} \\ \Delta T_{J2} \\ \Delta T_{J3} \\ \Delta T_{L1} \\ \Delta T_B \end{Bmatrix} = \begin{bmatrix} \theta_{J1A} & \Psi_{12} & \Psi_{13} \\ & \theta_{J2A} & \Psi_{23} \\ & \Psi_{13} & \theta_{J3A} \\ \Psi_{L1-1A} & \Psi_{L1-2A} & \Psi_{L1-3A} \\ \Psi_{B1A} & \Psi_{B2A} & \Psi_{B3A} \end{bmatrix} \begin{Bmatrix} P_{d1} \\ P_{d2} \\ P_{d3} \end{Bmatrix} \quad (\text{eq. 15})$$

Here we do not have a square matrix, because we are using psi characteristics to describe the temperature at some points in the system that are not themselves heat sources. (If they *were* heat sources, we could expand the matrix back to a larger square with two additional heat input variables in the heating vector, but this would necessarily require two additional columns of psi characteristics that need to be measured before computations may be completed. Clearly it is much simpler, if those points really are unheated, to avoid the extra work.) Perhaps obviously, the symmetry consequent of the reciprocity theorem applies only to the square *sub*-matrix representing the temperatures at only the heat sources. Once again, the preceding matrix notation is simply a shorthand version of the following system of equations:

$$\begin{aligned} \Delta T_{J1} &= \theta_{J1A} \cdot P_{d1} + \Psi_{12} \cdot P_{d2} + \Psi_{13} \cdot P_{d3} \\ \Delta T_{J2} &= \Psi_{12} \cdot P_{d1} + \theta_{J2A} \cdot P_{d2} + \Psi_{23} \cdot P_{d3} \\ \Delta T_{J3} &= \Psi_{13} \cdot P_{d1} + \Psi_{23} \cdot P_{d2} + \theta_{J3A} \cdot P_{d3} \\ \Delta T_{L1} &= \Psi_{L1-1A} \cdot P_{d1} + \Psi_{L1-2A} \cdot P_{d2} + \Psi_{L1-3A} \cdot P_{d3} \\ \Delta T_B &= \Psi_{B1A} \cdot P_{d1} + \Psi_{B2A} \cdot P_{d2} + \Psi_{B3A} \cdot P_{d3} \end{aligned} \quad (\text{eq. 16})$$

to all of which, ambient must be added.

In real life, you would measure three power levels and *three* temperatures (lead, board, and ambient). The matrix method then allows you to calculate *five* temperatures: the three junctions as well as the lead and board temperatures. If there is agreement between the calculated temperatures and the measured values (two opportunities: lead and

board), you would then have some confidence that all the characterization inputs (theta's and psi's) were valid for the system under consideration. If there was a significant discrepancy between the calculated and measured values for lead and board, you would have an indication that heat flow distribution in the actual application differs significantly from that during which the psi's originally had been deduced, casting into doubt the calculated junction temperatures as well.

For completeness, we shall demonstrate the additional complexity required for a true multiple heat source, multiple temperature boundary condition model. Suppose we have a six-lead package with two independent silicon devices (i.e., two junctions, meaning two heat sources) inside. If we can make temperature measurements at all six leads, under the assumption that virtually 100% of the heat generated internally, at either junction, must exit the package along one of the leads, then we could write the following matrix equation, where each lead is treated as a separate boundary condition:

$$\begin{Bmatrix} T_{J1} \\ T_{J2} \end{Bmatrix} = \begin{pmatrix} \Psi_{11} & \Psi_{12} \\ \Psi_{12} & \Psi_{22} \end{pmatrix} \begin{Bmatrix} P_{d1} + \frac{T_{L1}}{\Psi_{1L1}} + \frac{T_{L2}}{\Psi_{1L2}} + \frac{T_{L3}}{\Psi_{1L3}} + \frac{T_{L4}}{\Psi_{1L4}} + \frac{T_{L5}}{\Psi_{1L5}} + \frac{T_{L6}}{\Psi_{1L6}} \\ P_{d2} + \frac{T_{L1}}{\Psi_{2L1}} + \frac{T_{L2}}{\Psi_{2L2}} + \frac{T_{L3}}{\Psi_{2L3}} + \frac{T_{L4}}{\Psi_{2L4}} + \frac{T_{L5}}{\Psi_{2L5}} + \frac{T_{L6}}{\Psi_{2L6}} \end{Bmatrix} \quad (\text{eq. 17})$$

Two important differences exist between this complete model and the previous, simpler matrix formulation. First, the left hand side of the equation is the actual junction temperature prediction, *not* a temperature rise over some common reference temperature. Second, the temperature boundary conditions (of which there are six in this "complete" model) show up as quasi-heat-inputs. Each has its own associated weighting, expressed here as yet another psi value. So even though there are only two heat sources, there are actually 16 (15 independent) distinct parameters characterizing this model. Obviously it is one thing to state that such a model may exist; it is quite another thing to experimentally (or otherwise) derive all these coefficients. This is the domain of "compact models," for which a tremendous amount of literature and research has been generated over the past decade. The minimum number of independent external boundaries necessary for a certain desired level of accuracy, whether (and how) temperature nonlinearities must be considered, and even the internal structure of the model, are all significant issues that must be addressed for a thorough development of such a model.

Theta-JA vs. Copper Area

One of the most problematic issues in providing manufacturer's thermal data lies in choosing what data is helpful for purposes of illustration, but which may unfortunately grossly misrepresent how the package will actually perform in a customer's application. An excellent example of this is in the variation of Theta-JA with copper area, as shown in the next figure.

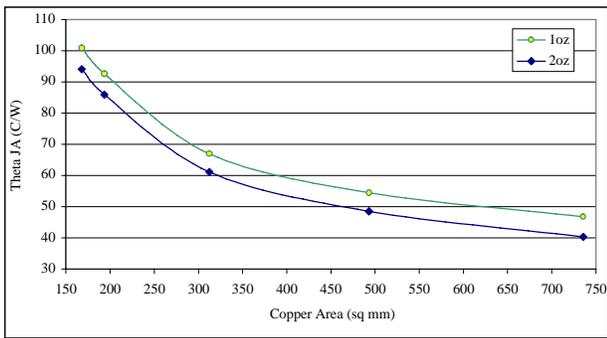


Figure 1. Theta-JA vs. Copper Area for a Dpak

The two parameters explicit on this chart are the thickness of the heat-spreading copper (one curve for each of two representative thicknesses), and the copper area (the x-axis). Even considering just these two items, many questions arise. For instance, what should be done if the copper thickness is outside the range provided? Indeed (though it would probably be considered ridiculous to try to be this precise in the first place), should one interpret linearly even within the range provided? What happens if the copper area falls below the left end or above the right end of the curves provided? Is a linear extrapolation appropriate? Should a curve (of what type?) be fit to the points known, and extended (and if so, how far is appropriate)? (Certainly, any polynomial extrapolation to the right will eventually pass through or end in physically impossible values, such as zero, or negative values. Worse, it may even *increase* as copper area increases.) Does the metal area include the traces directly associated with the device of interest? If so, clearly at some point toward the left ends of these curves, one transitions from large “blocky” areas to long skinny areas. Surely the heat spreading ability will shift markedly at this transition. So if in a particular application, a different amount of purely “trace” metal is utilized than that upon which this chart was derived, significant departures from an “obvious” extrapolation will occur.

However, beyond those explicit parameters, there are many very important factors *not* explicit on this chart affecting the actual theta-JA value that will be experienced. (And even if they were known, perhaps through conscientious provision of footnotes by the manufacturer, and diligent reference to the same by the customer, the same questions would remain regarding how to handle deviations from the stipulated conditions.) For instance, how large is the board beyond the copper area? How much airflow is there? Is the airflow the same on both top and bottom of the board? If the air is “still” (also known as “free convection,” being driven by the buoyancy due to the temperature difference in the air as it is heated by the device in question), what is the orientation of the board with respect to gravity, and how does it change with orientation (significantly)?

Finally, of course, how near are the nearest neighboring heat sources? Do they share common metal (i.e. ground

planes)? How much separation is there between metal areas not electrically connected?

The moral here is that these charts may be used *only* to gain a very, *very* rough idea of just how much leverage one *might* have, for the device of interest, in adjusting overall system thermal characteristics. Obviously the starting point is to subtract out the “intrinsic” device characteristics (psi-J-lead, for instance), and see if what’s left over might possibly give enough margin for a required compensation. If there is, a comprehensive analysis of the external thermal system is mandatory, taking into account at least all the variables that have been highlighted in this discussion.

Transient Data

Transient results may be presented in the form of heating curves, duty cycle curves, thermal-RC-network models, or thermal-RC mathematical models. In conjunction with any of these descriptions may also be found a “sqrt(t)” surface heating model, suitable for particularly short duration heating situations outside the scope of the other approaches.

Heating Curves

A heating curve (also known as a transient response curve, or a single-pulse heating curve), shows how the junction temperature of a device increases with time, given a constant power input at the junction, in some particular environment. In the following chart, it is clear that for times shorter than about 0.2 s, the device in question has the same thermal transient response independent of what sort of board it is mounted on. Somewhere between 0.1 and 1 s the effects of the environment (the board) start to be felt, and clearly by 1000 s the difference *between* those two particular example environments (about 50°C/W) is larger than the original contribution due to the device itself at 0.2 s (5–6°C/W). If two different environments are depicted on the same heating curve chart, therefore, one gets an idea as to when the environment begins to enter into play, and how much is due to the package alone.

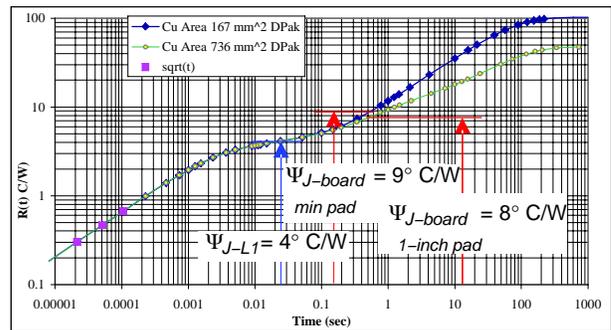


Figure 2. Typical Heating Curves

It may also be useful to note the point(s) on the transient heating curve where some of the other “steady state” theta or psi values occur. For instance, a psi-JL or psi-JB value may be found to occur roughly at the point where two

different environment curves separate. Possibly two different psi-J-tab values are given, one for each of the mounting conditions. In that case, it will be seen that they occur at roughly the same time along the two curves. One may also, therefore, use the psi-JL values, if provided, to help ascertain the time scales at which package effects “end” and environment effects “begin,” even when only one environment’s curve may be presented (for instance, if a 1” pad curve is given, but no min-pad board curve).

The use of a heating curve is straightforward, in the situation where constant power is applied to a device. If it is desired to know how hot the junction has become at a certain moment, the R(t) value is simply looked up and used just as if it was a theta value:

$$T_J(t) = R(t) \cdot P_D + T_A \quad (\text{eq. 18})$$

Sometimes heating curves are referenced to lead temperature or board temperature, rather than ambient, so one must pay attention to this detail in making use of the curve. It should be evident that if the curve ends up at the steady state theta value for the particular application and environment, then this computation yields the steady state junction temperature. The heating curve simply provides a time-variable generalization of the concepts previously discussed in the steady state context.

Heating curves may be used to estimate more complex power situations. For complete generality, a non periodic power input with ramps or smooth curves, rather than square edges, can be modeled using the single-pulse curve. Figure 3, illustrates the basic steps involved in such an analysis.

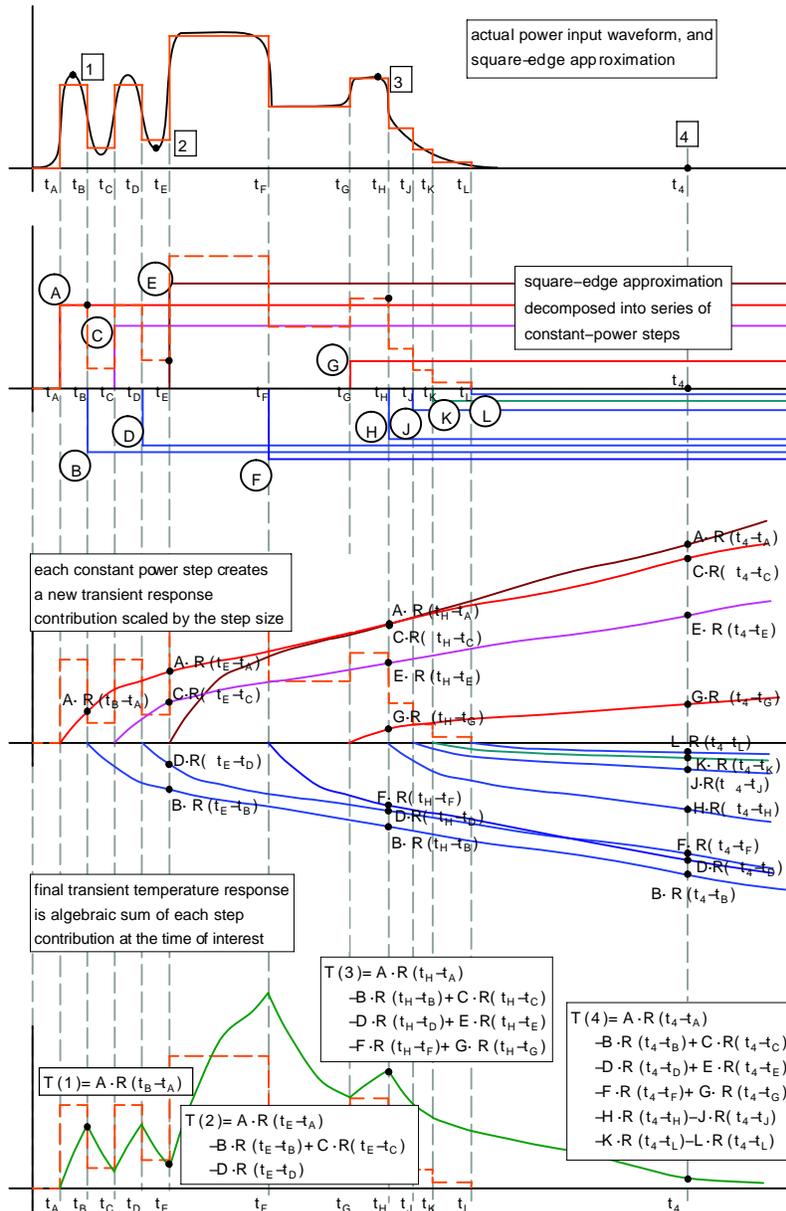


Figure 3. A Complicated Non-Periodic Power Analysis

Beginning with the upper diagram, the black trace represents the actual power input. The orange squared-off trace in that same diagram is an approximation of the actual power input, based on rectangular blocks of constant power. Usually an attempt is made to keep the total energy input correct, i.e. to make sure the area under the approximating curve is the same as that under the true power curve. The points labeled 1–4 are points at which a temperature estimate is desired. Points 1 and 3 are intended to capture the local peak temperatures, and point 2 is intended to capture the local minimum temperature. This means that the choices of the precise times for the square-edged approximation will dictate the times at which these local extrema will occur, and will not necessarily be at the exact times dictated by “real” power input waveform. In the second diagram, the approximated power input (now shown as a dashed orange line), is broken down into a sequence of steps of constant power, labeled A–L. The relative amplitudes of the power steps are the changes in power as each new step is introduced. Thus note that A, C, E, and G are positive going steps, and all the others are negative going steps. In the third figure, the transient responses due to each step of power are indicated. Each response is the basic single-pulse response scaled by the amplitude of the power step in question. In the bottom figure, the resulting temperature profile is illustrated, and the net contribution of the various pieces needed to compute the temperature (at the four points of interest) is indicated. The response at point 4 is shown mainly to illustrate that the temperature may be computed at *any* time of interest, whether or not it corresponds to a change in power; all that is needed is to include every step (in this instance, all of them) that has begun prior to the moment of interest, computing its contribution based on the elapsed time since each initiation. Clearly this could be done anywhere during the power input profile, for instance, anywhere between t_F and t_G .

One difficulty that sometimes arises, is that an $R(t)$ curve may need to be read with great precision when no such precision is apparently available. The preceding example, for instance, requires several pairs of nearby $R(t)$ values to be read accurately, so that a small difference between them may be used to calculate a temperature change. The recommended procedure is to assume that over small intervals, the transient curve can be represented by a power law (which will be a straight line segment on a log–log graph), of the form:

$$R(t) = a \cdot t^n \quad (\text{eq. 19})$$

Given two points far enough apart to span any particular range of interest (yet close enough together to be connected by an effectively straight line segment), the power law exponent n may be calculated using:

$$n = \frac{\log[R(t_2)/R(t_1)]}{\log[t_2/t_1]} \quad (\text{eq. 20})$$

Then the value a small distance ϵ from a nearby value at t may be calculated from:

$$R(t + \epsilon) = R(t) \left(1 + \frac{\epsilon}{t}\right)^n \quad (\text{eq. 21})$$

Yet a different approach to using the single-pulse heating curve may be possible for estimating peak temperatures of short wave trains with long gaps between the trains, but repeated periodically. For instance (refer to Figures 4 and 5), consider ten, 100 W pulses, with a 1-millisecond period and a 5% duty cycle (i.e. 0.05 ms on, 0.95 ms off); then every 100 milliseconds, repeat the same pulse train. Do this for 45 seconds. The question is, what is the highest temperature reached during this scenario?

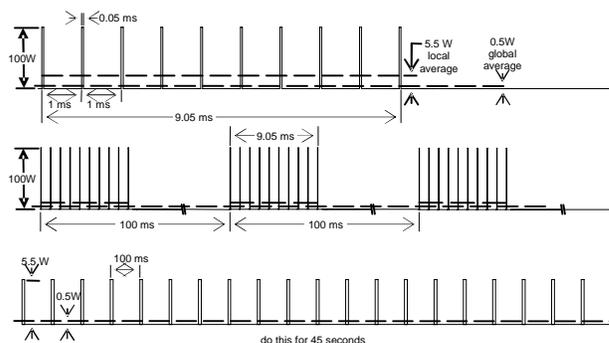


Figure 4. Periodic Short Pulse Trains Problem

Before we can answer that, we must first determine *when* this will occur, which is straightforward in this example. It will occur at the end of the 10th pulse at the end of the last pulse train at the end of the 45-second point (i.e. the end of the very, very last pulse to be applied). So now that we know *when* to find the temperature rise, we may find the *amount* of the temperature rise, which we can break down into three pieces.

Suppose we’re working with the device whose transient response is given in Figure 2, on the 736 mm² Cu area board. First, as seen in the upper portion of Figure 4, there is an overall average power dissipation to be concerned with. Out of every 100 ms, there are 10x0.05 msec of 100 W power applied – and the power is zero the rest of the time. So we have an average power of $(10 \cdot 100 \cdot 0.05 / 100) = 0.5$ W, when looking at the overall waveform. This average power adds a “background” temperature rise to the system, which we thus obtain by pretending there is a constant 0.5 W of power applied for 45 s (last portion of Figure 4). Clearly in the environment we’ve chosen for this example, the overall system hasn’t reached steady state by 45 s. But from the single-pulse curve, we *can* say that at 45 s, $R(t)$ is about 30°C/W, hence the background temperature rise will be $0.5 \text{ W} \cdot 30^\circ\text{C/W}$, or 15°C.

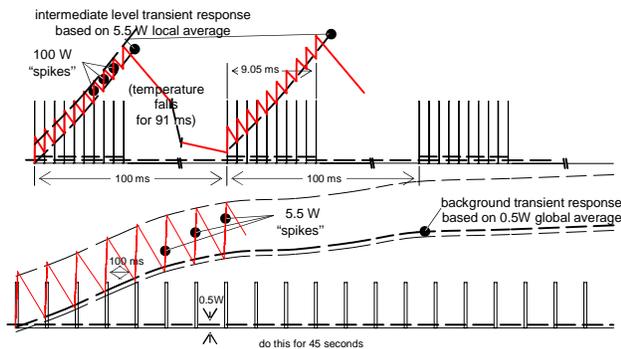


Figure 5.

Similarly, each pulse train of 10 fast pulses may be viewed as a short block of “constant” power that rides on top of this background average. Over a block of 10 fast pulses, we have an “average” power of $(10 \cdot 100 \cdot 0.05 / 9.05) = 5.5 \text{ W}$ (note that from the beginning of the first of ten pulses, to the time the 10th pulse turns off, is 9.05 ms). Between each of the 10–pulse trains, the power is off for so long (90.95 ms) that the temperature will basically fall all the way back to the “background” temperature rise, whatever that may have been at the beginning of the current train (which, again, is what ends up being about 15°C by the 45 s point in the problem). So again, reading the single–pulse curve at 9.05 ms, we find $R(t)$ of about 3.6°C/W, so at the end of the 10th pulse within each brief train, the average temperature rise will be $(5.5 - 0.5) \text{ W} \cdot 3.6^\circ\text{C/W}$, or about 18°C. (Note that we have subtracted out the 0.5 W of “background” average power, so as not to double–count it in computing the temperature rise of the average pulse train.)

Finally, we have to ask what the temperature rise of each individual 100 W pulse is around the quasi–average junction temperature of this question. We’ve already determined that by the end of 45 ms, 0.5 W will have translated into a 15°C rise, and another 5 W will result in a net 18°C rise between the first and last pulse of the ten, every time a 10–pulse train fires off. So what happens on that very last pulse of each ten? 94.5 W of the 100 W (having already accounted for the effect of 5.5 W), goes into “spiking” the temperature on top of whatever it was just before that pulse started. Again, from the single–pulse heating curves, $R(0.05 \text{ ms})$ is about 0.45°C/W, so a final 94.5 W pulse adds another $(94.5 \cdot 0.45 =) 43^\circ\text{C}$ to the cumulative temperature rise. Our conclusion, therefore, is that the peak temperature at the end of the final pulse at the 45 s point in the example, will be $15 + 18 + 43 = 76^\circ\text{C}$ above ambient. Clearly, due to the lack of precision in reading the single–pulse response chart, there is some inherent uncertainty in this result. If more accuracy than this is desired, thermal RC models may be used. These will be discussed subsequently. Before this topic is addressed, however, another application or extension of the use of transient heating curves needs to be covered.

Duty Cycle Curves

A transient heating curve may also be known as a “single pulse heating” curve. This is because it is derived from either an experiment or a model that produces the junction temperature rise in response to a sudden application of constant power – the longer the heating power, or “pulse,” is applied, the hotter the junction. Clearly, the heating power could be turned off at any moment, and the temperature rise at that instant would be known – so whether the pulse is in reality turned off, or simply continued as more data is collected at longer and longer times, the resulting plot may be interpreted as being the result of a single pulse whose “width” is indicated on the x–axis. This leads us naturally to the question of what happens if a pulse is repeated periodically, rather than applied once and never again (or in any other way non periodically, as in the previous examples).

If a pulse train of square pulses, of equal width and regularly spaced, is applied to a device, it turns out that the single pulse heating curve (just described) may be transformed into a family of curves, each of which represents the peak junction temperature that will be eventually reached once the pulse train has been applied for a long enough period of time (see AND8219/D). These curves are generally called “duty cycle” curves, and are parameterized by the percent of “on” time. In Figure 6 following, the single pulse curve for the 1” pad thermal test board, shown previously in Figure 2, has been so transformed, and the resulting family of duty cycle curves is presented.

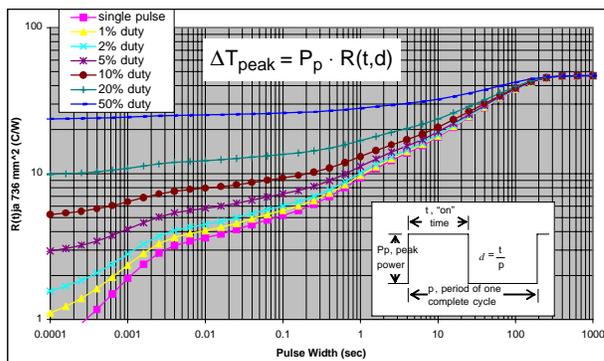


Figure 6. Typical Duty Cycle Curves

The x–axis of the chart is the individual pulse width, that is, the “on” time. So if the pulse width is “t” and the total length of a cycle (on time plus off time) is “p”, then the duty cycle, d, will be the ratio of t/p (and is usually expressed in terms of percent, as shown in the figure). To read the chart, then, one figures out the percentage of “on” time for the pulse train of interest, and then looks up the appropriate transient response value for that specific “on” time (on the corresponding % duty cycle curve). This gives a value, often denoted $R(t,d)$, in units of °C/W. Since these are square

waves, the peak power is the height of the individual pulses (note also that the average power will be $d \cdot P_p$, the duty cycle times the peak power). Note that to calculate *peak* temperature rise, one must use the *peak* power for the pulse, *not* the average power.³

The most common mistake made in using duty-cycle curves, is that they may only be applied for situations where the original single-pulse curve was the “correct” curve for the application environment under consideration. This means that it must end up at the correct steady state value. If the single pulse curve was for a min-pad board, for example, then *none* of the resulting duty cycle curves may be used for a 1” pad application, regardless of how short the pulses or what the duty cycle. To understand this, it may be helpful to consider the mathematical expression typically used to derive these curves:

$$R(t, d) = (1-d) \cdot R(t) + d \cdot R(\infty) \quad (\text{eq. 22})$$

If d is vanishingly small, then the result is the original curve (which is clearly valid only until environmental effects come into play). For any finite d , however, *regardless* of how short the pulses of interest (i.e. “on” time), the duty cycle curve carries along a contribution from the steady state end of the original curve, i.e. $R(\infty)$.

Given the appropriate single-pulse curve, if the pulse train is periodic (even if not square), the square-wave duty cycle curves may provide a time-saving approximation. For instance, pulses that are trapezoidal or triangular in shape, partial sinusoids, etc., may be approximated by square pulses with the same total energy (i.e. area under the pulse), where the height and width of the equivalent square pulse are adjusted such that the end of the pulse coincides with the moment of peak temperature – though this itself may require some experience to judge when that is likely to occur.

Thermal RC Network Models

Thermal RC network models are an alternative way of describing the same transient thermal response previously discussed (see AND8214/D and AND8221/D). An entire transient response curve can usually be represented in just a handful of resistor and capacitor elements. If the correct computational tools are readily available, RC networks may therefore be a convenient and compact representation. Two general forms of RC networks are possible, those with grounded capacitors, and those in which the capacitors are not grounded. These will be discussed in turn.

³ (Think *peak power* for *peak temperature*, if that helps. But to see why this must be so, think about that single-pulse curve being the equivalent of the 0% duty cycle case. For a given pulse width, if the only thing that changed was the period, you’d be staying at the same position on the x-axis while you moved from one curve to another. As the period went to infinity, you’d end up on the 0%, or single-pulse, curve; but if the power you multiplied by was the *average* power, you’d also be moving toward a zero-power average, hence the temperature rise would approach zero for a fixed pulse width. Obviously this would be incorrect – because the whole point of the single-pulse curve is to give you the actual temperature rise based on the power level of the pulse *while it’s on*, so clearly you’re supposed to be using the instantaneous power level, not the average power level.)

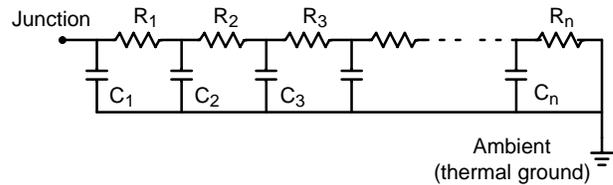


Figure 7. Grounded Capacitor Thermal Network (“Cauer” Ladder)

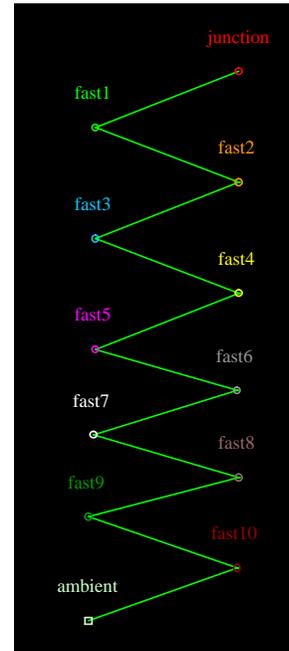


Figure 8. RC Network Schematic
Capacitors (not shown) tie each node to ground

Figures 7 and 8 illustrate a typical “grounded-capacitor” thermal ladder network. In fact, any network topology of resistors might be chosen to represent a physical thermal system (i.e. not just a linear string of resistors, but just as well a star, a bridge, or whatever). The main advantage of a grounded-capacitor network is that it derives from the fundamental heat-transfer physics. Every node in the network is connected to thermal ground through a capacitor. If a simple chain of nodes is used, it is convenient to draw the network as shown in Figure 7 because it resembles a ladder, though because the lower edge of each rung attaches directly to ground, the connections between the rungs are essentially through the resistors. Often for clarity, the capacitors are omitted entirely, in which case Figure 8 is an equivalent model. A grounded capacitor network such as shown in Figure 7 or 8 is known as a *Cauer* ladder.

Because this network derives from the real physics, there is at least a chance that experimental data from various points within the physical system can be correlated with specific individual nodes of the network model. As we move from junction to ambient, for instance, we might find physical locations corresponding to nodes in this order: silicon junction, back of silicon chip, edge of leadframe, lead (at package boundary), lead (at board interface), board (at some distance from package), and finally ambient. Of course, we may not *have* any intermediate location data to correlate with, or the intermediate data we have might not happen to land “on” a node of the model (rather, somewhere in between nodes). Also, the physical system might *not* be well represented by such a simple chain of resistors, so no correlation might be possible except at the junction itself. (This is actually quite typical, for in many environments, the heat flow follows at least two separate and distinct paths from junction to ambient, e.g. upward through the case, outward through the leads into the board, and downward through the air gap and thus directly to ambient on the back side of the board. When the heat flow is believed or known to flow along multiple parallel paths, it clearly would be better to model the system with a more complex network. Likewise, one would not expect to find a nodal correlation with physical locations if the network was willfully chosen as a simple ladder, when multiple significant parallel paths were present.) Only in the case where a single path to ground dominates heavily, would such a simple linear resistor topology be expected to yield good correlations at the intermediate nodes. Nevertheless, the point is that there *could* be such a correlation.

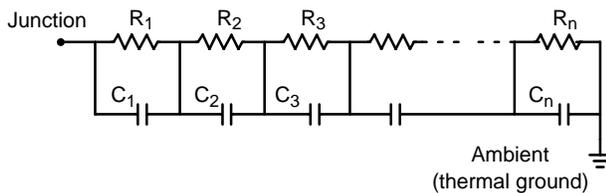


Figure 9. Non-Grounded Capacitor Thermal Ladder (“Foster” Ladder)

Non-Grounded Capacitor Networks

Contrast the grounded network of Figure 7, with the non-grounded-capacitor network of Figure 9. Figure 9 is a true “ladder” of resistors and capacitors, and is sometimes known as a *Foster* ladder. Each rung is joined to the next rung (and only to the next rung) through both the resistor and the capacitor; only the final capacitor is directly connected to thermal ground.

Difficult though it may be to grasp at first, this network has no physical basis. In the thermal/electrical analogy, a thermal capacitor is simply an element that stores energy

based only on *one* temperature, that is, the temperature of the node whose thermal mass it represents. (In contrast, an electrical capacitor stores energy based on the *difference* in voltage between its two terminals.) Hence a thermal capacitor whose energy storage is based on a difference in temperature between two ungrounded nodes in a network (as is the case of most of the capacitors in Figure 9) has no physical meaning. However, there is a mathematical simplicity underlying Foster ladders. In their mathematical description, one finds that each resistor-capacitor pair contributes an “amplitude” to the overall system response, and a unique time constant associated with that amplitude. Indeed, a Foster ladder may be viewed as nothing more than a schematic of the mathematical fit to a real transient response curve. Given a transient response curve of junction temperature vs. time, a series of exponential terms consisting of amplitudes and time constants may be fit to the curve to whatever degree of accuracy is desired. (More terms usually implies a better fit.) Once done, the terms may be interpreted as an RC ladder (i.e. the Foster ladder) where each amplitude is a resistor, and each time constant is the product of its associated resistor and a capacitor in parallel with it.

Comparison and Contrast of Cauer and Foster Ladders

Clearly, the mathematical terms representing (or represented by) a Foster ladder may be added together in any order to achieve the same sum. Thus the rungs (RC pairs) of the schematic may be listed (or diagrammed) in any order and still represent the same response! Because the overall response (from junction to ambient) is immune to reordering of the individual rungs (as long as each RC pair remains a pair), the temperature that might be calculated at any other node *between* any two rungs is physically meaningless. By contrast, though a Cauer network must necessarily have a mathematical representation comprising amplitudes and time constants, one finds that *every* amplitude and *every* time constant depends on *every* resistor and *every* capacitor, in a highly complicated and algebraically intractable tangle intimately dependent on the physical location of the elements in the network.

Even so, Foster networks are typically drawn with the rungs placed in order from junction to ambient with the smallest values (i.e. fastest responding rung) at the junction end, and the largest values (i.e. slowest responding rung) at the ambient end. This is superficially similar in appearance to a typical Cauer ladder, which almost always (and necessarily) has the smallest elements nearest the junction, and the largest elements nearest thermal ground. But in the Cauer ladder, the choice is not arbitrary; rather it is imposed by the intrinsic relationship between time response and location in the model.

So where does this leave us? A grounded–capacitor model is most useful when a physically meaningful model is required, for instance, to separate the package from the environment in order to replace the environment portion of the network with a different network, representing a different environment. However, to best use a grounded–capacitor model, a circuit simulating tool is required. Of course, if a circuit simulating tool is being used for thermal calculations, any complicated, time varying, power input may be imposed on the circuit without particular cost. Finally, multiple heat source models can be built with equal facility, and again, arbitrarily complex asynchronous power inputs at any number of nodes may be managed without difficulty.

On the other hand, non–grounded–capacitor models are mathematically very simple, and quite detailed thermal calculations may be performed with spreadsheet based tools. Though only the junction has the designed correlation with physical reality in a non–grounded capacitor model, the fact remains that if the model is capable of producing the known transient response, applicable to the particular environment in question, then it may be used successfully for temperature predictions with arbitrarily complex power inputs. Even multiple input thermal models may be constructed using equivalent non–grounded networks.

It should also be mentioned that every Cauer ladder has a Foster equivalent, and vice versa. The conversion from one to the other is a non–trivial operation, but algorithms do exist for that purpose. In ON Semiconductor data sheets, generally the two equivalent networks are provided, enabling a knowledgeable customer to take advantage of the strengths of each.

Multiple–Junction Devices and Transient Response

In the preceding section, allusions have been made to multiple–input transient models. Just as with steady state descriptions of a thermal system, transient descriptions of multiple–junction devices may be constructed. If a matrix method is followed, the only difference is that every element of the matrix is a function of time. For every heat source in the device, there will be a “self heating” transient response curve, and for every other point of interest in the system (whether another heat source or a passive temperature monitoring location) an “interaction” transient response curve will exist.

Bounded by the same limiting assumptions, the principles of linear superposition and reciprocity continue to apply. That is to say, the time–varying response at any point in the system may be treated as the linear superposition of its response to each independent heat source, as if each heat source was powered individually and independently of the

others. Further, the less intuitive truth of the reciprocity theorem applies in the time domain: namely, the transient response of point “A” in the network due to (constant) heat input at some other point ‘B’, will be identical to the transient response at point ‘B’, if the same amount of heat input is applied at point ‘A.’ So in the matrix description, symmetry across the main diagonal will still exist. Perhaps the most powerful implication of the reciprocity theorem is experimental: in effect, only half of the total possible interaction thermal transient responses need be measured.

Circuit Simulators

Cauer models are of little use without a circuit simulator, due to the messy algebra required to describe their mathematical response. Thus, if a Cauer model is all that is available (at least, if it consists of more than about two nodes), a circuit simulator is a must. Of course, if a circuit simulator is available, a circuit is a circuit, and it should therefore be evident that both Cauer and Foster ladders may be analyzed with equal ease. Indeed, for single–input networks, there will be no difference in the overall approach, only in the details of the network connectivity and element values.

For multiple–input networks, the Cauer network will be straightforward (see Figure 10). Recalling that a Cauer network will have been derived on some premise of having physical significance, the interactions between the various possible heat sources (and possibly passive “boundary” nodes) will be built into the topology of the network itself. Resistors and capacitors will exist that “automatically” provide the correct interaction responses due to heat inputs at every source; reciprocity and superposition are necessary consequences of the method. One simply enters the grounded–capacitor Cauer model into the simulator, with all nodes and interconnections explicit in the schematic, and the task is done.

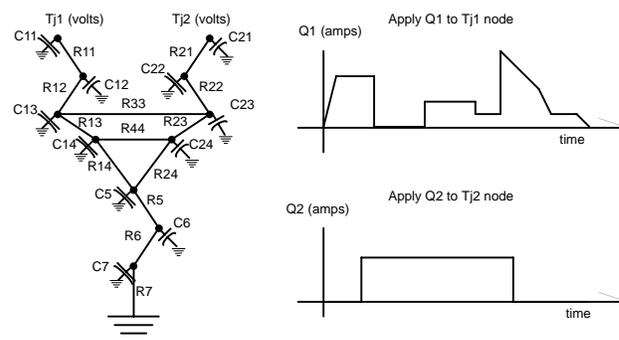


Figure 10. Implementing a Multiple Input Cauer Network in a Circuit Simulator

A multiple-source Foster model is more complicated to implement in a circuit simulator, and exactly how it is done will depend on the features of the simulator available. Since a Foster model is nothing more (in essence) than a schematic of the mathematical fit to a transient response curve, the resistors and capacitors in a particular “self heating” Foster ladder will not correlate with the resistors and capacitors in any of the interaction networks; nor will the self heating Foster ladder elements of any two heat sources have any correlation with each other, even though we may know that there is much underlying common thermal path between the two sources. Moreover, depending on how the Foster ladders were derived, even the time constants between various curves in the model may not match! (It may be observed that if “closed form” Foster ladders are derived from a Cauer model directly, at least the time constants will be shared across all the self heating and interaction curves; but if the Foster ladders are simply independent mathematical fits to response curves, from whatever source, there is no particular reason the time constants will match, unless by intentional choice.) Again, depending on how the

Foster ladders are derived, there may even be “negative” amplitudes (this is, in fact, guaranteed to occur in “closed form” solutions to the interaction responses in a physically significant grounded-capacitor models.) Clearly, if negative amplitudes arise in the Foster representation, a circuit simulator must permit negative resistances (and obviously negative capacitors, since a positive time constant can only result for a negative resistor, from its product with a negative capacitor). Alternatively, the simulator must provide a programmatic method of subtracting the response of one node from another, so a negative contribution can be constructed from positive sub-circuits. Similarly, to implement a multiple input Foster model in a circuit simulator, care must be taken to intentionally create “summing” nodes that implement the principle of linear superposition between the various otherwise independent self heating and interaction heating portions of the overall model. If the circuit simulating tool does not provide features adequate for these tasks, a spreadsheet implementation will be the best alternative. Figure 11 illustrates the possible steps.

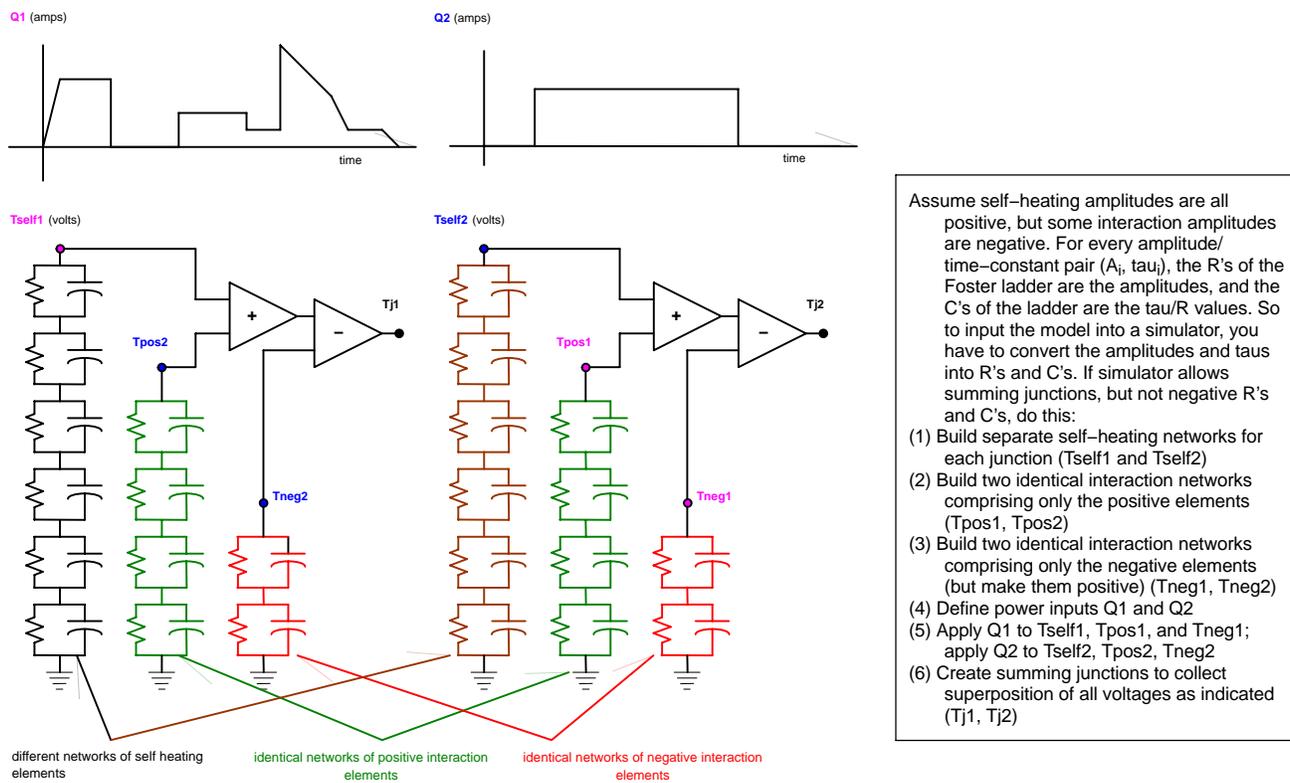


Figure 11. Implementing a Multiple Input Foster Network in a Circuit Simulator

Spreadsheet Models

As previously described, Cauer models basically require a circuit simulator, even for single-input models. For Foster ladders, however, spreadsheet tools are convenient for implementing both single input and multiple input models. This is the consequence of the mathematical simplicity of Foster models, and the triviality with which spreadsheets permit superposition to be introduced. For instance, consider the ease with which a constant power transient response for a single-input Foster ladder may be written in Microsoft® Excel. Suppose the following meanings are given to certain cells in a spreadsheet:

cell A1 is the power level
cells B1:B10 are the amplitudes
cells C1:C10 are the time constants (where C1 is the time constant for the B1 amplitude, and so forth)
cell D1 is the time after a constant power step begins

Then the Excel formula to calculate the temperature rise at time **D1** is:

$$\{=A1*SUM(B1:B10*(1-EXP(-D1/C1:C10)))\}$$

Though by no means necessary, it may also be noted that by using Excel's *Name* capability and judicious use of *absolute* vs. *relative* reference notation, we can make this formula more mnemonic, and easy to copy to different locations for computing results at many different times. Amending the previous example; the mathematical expression of single-pulse heating curve in terms of Foster-type amplitudes and taus:

$$R(t) = \sum_{i=1}^n R_i(1-e^{-t/\tau_i}) \quad (\text{eq. 23})$$

Define Names

power	\$A\$1
amplitudes	\$B\$1:\$B\$10
taus	\$C\$1:\$C\$10
time	D1

We are now permitted the more readable formula:
 $\{=power*SUM(amplitudes*(1-EXP(-time/taus)))\}$

If this formula was entered into cell **E1**, it could be copied down into cells **E2** through **E100**, for instance, resulting in the time response at each of the times found in cells **D2** through **D100**. The *Table* feature of Excel may also be used to advantage to create a table of many values from a single formula.⁴

As time varying power inputs are introduced, and multiple heat sources are introduced, it obviously gets more complicated, but remains quite manageable for a relatively limited number of inputs and time steps. The method is that illustrated earlier (example given in Figure 3), with the following embellishments: (1) the temperature at any point of interest is the superposition of the response at that point due to *all* heat sources (so just as in the steady state case, if a junction temperature is being calculated, there will typically be one self heating contribution, and multiple interaction terms); (2) a new "step" in time must be made whenever the power input changes at *any* heat source, even if there is no change in power at that instant at the point in question (because the interactions will change, even when the self heating contribution does not change).

RC Models and Short-Time Transient Response

It may be shown mathematically that when the time scale is shorter than its fastest time constant, an RC model's transient response becomes proportional to time (i.e. *linear* in time). This will not be a problem if (1) the time scale of interest is somewhat greater than the fastest time constant, or (2) if it is known that the linear response with time is appropriate for the system under consideration. However, as will be discussed subsequently, for many semiconductor devices there is a range of time (typically between about 1 microsecond and 1 millisecond) when the concept of "surface heating" closely approximates the real thermal physics. In surface heating, device transient response is

⁴To those unfamiliar with "array" formulas in Excel, the preceding example accomplishes some very powerful operations in a compact notation. First, the use of the array syntax itself (the colon as part of a cell reference) tells Excel to execute the same computation for each cell of the range in turn; thus, since there are ten cells in each of the two arrays identified, ten parallel computations result. This means that ten different terms representing ten amplitudes and time constants, all evaluated for the same time (**D1**), are calculated at once. Second, the squiggly braces {} surrounding the formula indicate that (in this case) the formula was actually entered into the spreadsheet with the *Ctrl-Shift-Enter* keystroke, rather than the ordinary *Enter* keystroke. (In other words, simply typing in squiggly braces does not accomplish the same thing at all!) This tells Excel that we wish it to actually return all the available array results in however many cells are assigned to the formula. (More typically, this "array formula" entry method is used to distribute the array results over an array of cells; for instance, if we had selected the ten cells **E1:E10** for the formula entry, typed the formula into the formula entry blank, and hit *Ctrl-Shift-Enter*, one of each of the ten individual array results would be listed in each of the ten selected cells.) Here, however, we have no need to see all ten results individually, yet we still wish to access them all even though only a single cell is the target for the formula's result. Thus, finally, we include the **SUM** function to tell Excel to add up those ten individual results, rather than report just the first one in the single cell we've picked for the location of the formula.

AND8220/D

proportional to the square-root of time, rather than linear in time. Now a properly constructed RC model (i.e. one with time constants and amplitudes so designed) is capable of following this square-root behavior with excellent accuracy, but only for time scales greater than the model's shortest time constant. It is therefore important, whenever using an RC model, to consider whether the shortest time constant is fast enough for the needs of the analysis. For a Foster ladder, the fastest time constant is known exactly (the smallest RC product in the model). For a Cauer ladder, a good (although not exact) estimate of the fastest time constant is similarly obtained as the product of the RC pair closest to the junction. (In fact, ill-formed Cauer ladders, depending on exactly how the R's and C's appear in order, may yield actual fastest time constants several orders of magnitude faster than that of the RC pair nearest the junction; alternatively, ill-formed Cauer ladders may result in Foster amplitudes that are infinitesimal, yet with finite time constants. However, these are aberrations that do not materially affect the ability of the network to respond to

square-root-of-time needs on the order of the RC pair nearest the junction.) In any case, if the shortest legitimate time constant is not smaller than the shortest time scales of interest, and especially in the microsecond to millisecond time scale, extreme caution should be taken in interpreting RC model results. If a linear model is used when a square-root model is appropriate, temperature changes as predicted by the model will occur much too slowly, and *significant* underestimates of maximum junction temperatures may result.

With that caution in mind, the following table presents RC models for the same D2pak device on two different thermal test boards (basically min-pad and 1" pad, with some extra trace area included in the total area quantity). For each board, both Cauer and Foster networks are given. It should be emphasized that these Foster networks are in fact the exact mathematical equivalents of the corresponding Cauer networks. A number of the concepts addressed in the preceding discussions may be illustrated.

Table 1. RC Networks ("R" values are °C/W; "C" values are J/°C; "tau's" in seconds)

Drain Copper area (1 oz thick)			241 mm ²	788 mm ²		241 mm ²	788 mm ²	
(SPICE deck format)			Cauer network			Foster network		
			241 mm ²	653 mm ²	units	Tau	Tau	units
C_C1	Junction	Gnd	6.3269E-6	6.3269E-6	W-s/C	2.9892E-7	2.9892E-7	sec
C_C2	node1	Gnd	2.9939E-5	2.9939E-5	W-s/C	4.3949E-6	4.3949E-6	sec
C_C3	node2	Gnd	8.9817E-5	8.9817E-5	W-s/C	3.8122E-5	3.8122E-5	sec
C_C4	node3	Gnd	1.9877E-4	1.9877E-4	W-s/C	2.9542E-4	2.9542E-4	sec
C_C5	node4	Gnd	1.3388E-3	1.3388E-3	W-s/C	2.3055E-3	2.3055E-3	sec
C_C6	node5	Gnd	2.5099E-2	2.5099E-2	W-s/C	1.2749E-2	1.2766E-2	sec
C_C7	node6	Gnd	3.1191E-1	3.1815E-1	W-s/C	3.3747E-1	4.1823E-1	sec
C_C8	node7	Gnd	2.2054E-1	4.7830E-1	W-s/C	3.3611E+0	2.7622E+0	sec
C_C9	node8	Gnd	8.8815E-1	1.9594E+0	W-s/C	2.1614E+1	3.0643E+1	sec
C_C10	node9	Gnd	1.8889E+0	6.0036E+0	W-s/C	1.1357E+2	1.2328E+2	sec
			241 mm ²	653 mm ²		R's	R's	
R_R1	Junction	node1	0.0578524	0.0578524	C/W	0.03814	0.03814	C/W
R_R2	node1	node2	0.173557	0.173557	C/W	0.093163	0.093163	C/W
R_R3	node2	node3	0.520671	0.520671	C/W	0.201565	0.201565	C/W
R_R4	node3	node4	1.07638	1.07638	C/W	0.936692	0.936690	C/W
R_R5	node4	node5	1.44732	1.44732	C/W	1.730444	1.730479	C/W
R_R6	node5	node6	0.510799	0.510799	C/W	0.690301	0.691548	C/W
R_R7	node6	node7	2.84846	2.31584	C/W	0.333827	0.60289	C/W
R_R8	node7	node8	9.11661	4.38504	C/W	4.196175	3.230389	C/W
R_R9	node8	node9	34.2576	20.0524	C/W	6.059695	5.266272	C/W
R_R10	node9	gnd	24.9485	11.0277	C/W	60.677683	28.776447	C/W

NOTE: The boldface elements represent the part of the network most closely associated with the package; the remaining elements represent the environment. Listing the Foster rungs in ascending order of time constant provides a rough, though imperfect, equivalent, as the fast response rungs necessarily result in the most significant contributions to the short-time (hence the package) portion of the curve. As emphasized previously, however, the exact location of nodes within the Foster rungs has no direct physical significance, and any apparent correlation with the Cauer resistors is purely coincidental.

First, the fastest time constant for these networks is $2.98E-7$ seconds (given exactly in the Foster *Tau* column). An approximation of this value is found by the RC product in the Cauer network closest to the junction, i.e. C_{C1} times R_{R1} , yielding $3.66E-7$ s. Second, for convenience (remember, the choice is arbitrary and does not affect the junction response whatsoever), the rungs of the Foster ladder are listed in ascending order of time constant, but it is clear that the R's do not correlate very well with the R's of the "corresponding" rungs of the Cauer networks. Third, beginning at the short time ends of the ladders (both Cauer and Foster), the models are identical between the two boards. That is to say, for single pulse heating response, only the package (which is the same) matters at first, and only after some time has passed and the heat has begun to cross from package into the board, does the environment influence the response.

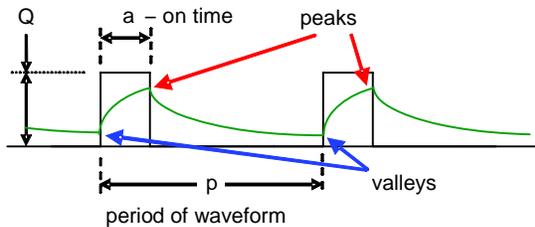


Figure 12. Basic Square Wave

Periodic Waveforms Using Foster RC Models

Square-wave duty-cycle curves have been discussed, and they are often derived from the simple formula expressed previously in Equation 22. However, given an RC model (in particular, the amplitude/time-constant Foster expression) of a single-pulse transient curve, exact closed form solutions for an infinite train of equal square pulses may be derived. We will simply present several of these solutions and illustrate how they may be applied (see AND8219/D). Given the single-pulse heating curve formulation of an n-rung RC model as indicated in Equation 23, we have the following:

$$R(a, d) = \sum_{i=1}^n R_i \frac{1 - e^{-\frac{a}{\tau_i}}}{1 - e^{-\frac{p}{d\tau_i}}} \quad \begin{array}{l} \text{peaks of simple square} \\ \text{wave train of duty cycle} \\ \text{d, on-time a} \end{array} \quad \text{(eq. 24)}$$

$$Y(a, d) = \sum_{i=1}^n R_i e \left(1 - \frac{1}{d}\right) \frac{a}{\tau_i} \frac{1 - e^{-\frac{a}{\tau_i}}}{1 - e^{-\frac{p}{d\tau_i}}} \quad \begin{array}{l} \text{valleys of simple} \\ \text{periodic square} \\ \text{wave train} \end{array} \quad \text{(eq. 25)}$$

Observe that the on-time, period, and duty cycle of the waveform are related through the equality $a = p \cdot d$. When the on-time is plotted on the x-axis, and the duty cycle is used as a curve parameter, Equation 24 gives us the family of duty cycle curves previously encountered in Figure 6, based on the Foster RC resistor model as fit to the original

$R(t)$ single-pulse heating curve. Indeed, if the RC model is a good fit, the duty cycle curves derived from Equation 24 will be more accurate than if they are derived from the more approximate formula of Equation 22 (with the possible exception that for very short duty cycle values, and on-times smaller than the smallest RC time constant, we may have the same limitation related to square-root of time previously discussed).

When a single pulse is repeated, (Figure 12), obviously the peaks occur at the ends of the "on" times, and the valleys occur at the ends of the "off" times (i.e., the beginning of each "on" time). Further, when only a single square pulse is repeated, it doesn't matter where the pulse is positioned within the period, if all one is concerned with is the peaks and valleys. In fact, for convenience these preceding formulas were derived under the assumption that the "on" time of each pulse begins at the beginning of each cycle.

However, if we generalize the problem slightly and allow a single square pulse to be positioned at an arbitrary point within the cycle, some more powerful formulas may be derived. For the formulas that follow, Figure 13 defines the parameters for a generalized square pulse within a period of length p. All times are relative to the beginning of a cycle.

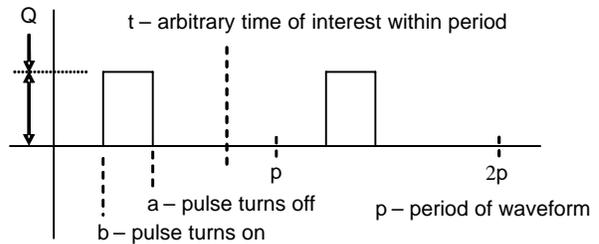


Figure 13. Generalized Square Wave

After an infinite number of identical cycles, the following three formulas describe the shape of the temperature response for the ranges indicated:

$$F(a, b, p, t) = \sum_{i=1}^n R_i \left(\frac{e^{-\frac{a-t-p}{\tau_i}} - e^{-\frac{b-t-p}{\tau_i}}}{1 - e^{-\frac{p}{\tau_i}}} \right) \quad \begin{array}{l} \text{good} \\ \text{(computable)} \\ \text{only for} \\ 0 \leq t < b \end{array} \quad \text{(eq. 26)}$$

$$F(a, b, p, t) = \sum_{i=1}^n R_i \left(1 + \frac{e^{-\frac{a-t-p}{\tau_i}} - e^{-\frac{b-t}{\tau_i}}}{1 - e^{-\frac{p}{\tau_i}}} \right) \quad \begin{array}{l} \text{good} \\ \text{(computable)} \\ \text{only for} \\ b \leq t < a \end{array} \quad \text{(eq. 27)}$$

Note: if $t = 0$ and $b = 0$, Equation 25 results

$$F(a, b, p, t) = \sum_{i=1}^n R_i \left(\frac{e^{-\frac{a-t}{\tau_i}} - e^{-\frac{b-t}{\tau_i}}}{1 - e^{-\frac{p}{\tau_i}}} \right) \quad \begin{array}{l} \text{good for} \\ 0 \leq t \leq p \\ \text{(computable)} \\ \text{only for } t > a \end{array} \quad \text{(eq. 28)}$$

Note: if $t = a$ and $b = 0$, Equation 24 results

For these formulas, the “computability” restriction is a practical matter arising when positive arguments appear in the exponential terms of the various numerators. Note also that these formulas describe the response curve, but the power level of the applied pulse has not yet been considered. We defer consideration of pulse power to the following formula, which now expresses the completely general superposition of any number of square pulses occurring with the same frequency, all positioned within the same cycle of period p :

$$G(t) = \sum_k Q_k F(a_k, b_k, p, t) \quad (\text{eq. 29})$$

Equation 29 now permits us to predict the “steady–state” transient behavior of any complexity of periodic power, assuming we break down the cycle into a series of square–edged pulses – a process we illustrated (see Figure 3) in an earlier example of a non–periodic waveform. By “steady state” transient response, we mean the shape of the temperature response curve for a typical cycle after an infinite number of such identical cycles has occurred. An important observation now must be made. Whereas the “peak” and “valley” temperatures for an infinitely repetitive *single* pulse can be predicted (i.e. Equation 24, 25) without knowing the details of the profile, this is not possible for a general periodic waveform, even when the waveform is a relatively simple combination of just a few square sub–pulses. Consider the following example, with the periodic power input of Figure 14 applied to the RC model given in Table 4.

Table 2. RC Model for 3–Pulse Example

tau (sec)	R [°C/W]	tau (sec)	R [°C/W]
1.00E–6	0.01104	1.00E–1	1.128566
1.00E–5	0.012806	1.00E+0	3.539626
1.00E–4	0.069941	1.00E+1	5.423616
1.00E–3	0.275489	1.00E+2	12.08694
1.00E–2	0.019806	1.00E+3	16.2933

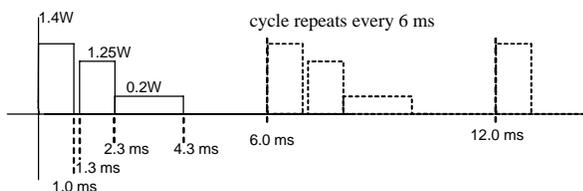


Figure 14. 3–Pulse Periodic Input

Applying Equation 26, 27, and 28 to the respective portions of each of three separate square pulses comprising the repeated pattern, and Equation 29 to superimpose their effects, we find the following temperature response:

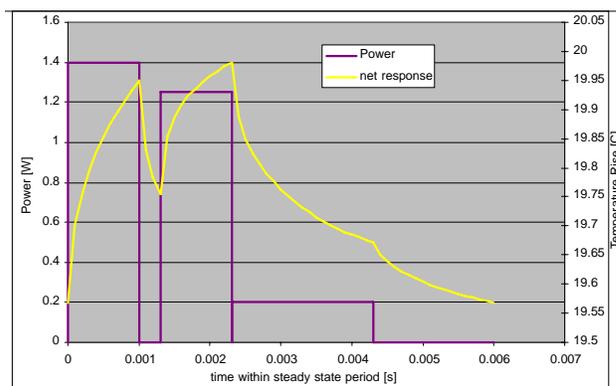


Figure 15. 3–Pulse Periodic Example Steady–State Transient Response

What makes this example particularly interesting is that the peak temperature (during a steady state cycle) occurs at the end of the *second* pulse, which has a lower power and even a small gap of zero power, between it and the higher power pulse immediately preceding it in the cycle. Knowing that the single pulse response is proportional to power, and that the peak temperature always occurs at the end of a square pulse (even when infinitely repeated), one might easily overlook the possibility demonstrated here. In other words, for a generalized periodic waveform, even when constructed of (or perhaps approximated by) just a small number of square sub–components, one does well to compute the response throughout the entire range of a cycle, not at just the “obvious” points.

Surface Heating, Square–Root of Time, and Short–Time Transient Response

In most thermal transient tests, experimental data is acquired as early as 1E–5 s (10 microseconds). However, in most cases, due to electrical switching transients, it is inconsistent between test devices out to as late as 1E–3 s. Even when measurement consistency occurs at earlier times, it is rarely reliable (meaning, consistent with expected theoretical behavior) at times earlier than 1E–4 s. Usually, in fact, measured signals corresponding to expected theoretical behavior do not occur until somewhere between 3E–4 s and 1E–3 s. The two major contributors to this correlation (or lack thereof) are electrical transient effects in the devices (from the measurement perspective), and die geometry effects (from the theoretical perspective).

More specifically, die thickness, and actual active heated area as compared to overall die size, affect the theoretical behavior. The simplest, commonly used theory for short–time thermal transient behavior is the surface heating model. This assumes constant power, 1–dimensional heat flow (i.e., flow direction is exactly perpendicular to the heated surface, with no spreading effects), and results in a surface temperature rise that is proportional to the square–root of heating time.

Because of this, it is often referred to as “sqrt(t)” heating. An important aspect of sqrt(t) heating is that on a log–log plot (refer back to Figure 2), such a heating “curve” is a straight line that rises one decade of temperature (or thermal resistance) for every two decades of time increase (i.e. a factor of 10 in temperature for every factor of 100 in time – hence the sqrt(t)). On the log–log plot this appears, therefore, as a 1:2 slope. The vertical position of this theoretical line (the intercept on the log–log plot, or the proportionality factor on a linear vertical scale) is determined by the area being heated, and the material properties of the die and whatever material adjoins the heated surface of the die. (Typically, this adjoining material is mold compound, which effectively lowers the heating rate of a silicon free surface by about 10%. However, if the silicon is covered by a copper “clip,” the short–time heating rate may be lowered by as much as 70%.) Also consequent to the sqrt(t) theory, the thinner the die, the sooner the heat reaches the back side of the silicon and thus ceases to follow the sqrt(t) model; a die of half the thickness will thus end its sqrt(t) behavior in one fourth the time. Typically, we consider that theoretical behavior should persist until about 1E–3 s for a 15 mil (380 micron) thick die, but when the thickness is as small as 10 mil (250 micron), theoretical behavior will last only 4E–4 s; for a 7–mil (180 micron) thick die, sqrt(t) will last for only 2E–4 s. Die thickness is also directly related to the other “extreme” of transient behavior, that is, how long it will be until local steady state (meaning the final temperature gradient, or maximum temperature difference between heated surface and back of silicon) is reached. All else being equal, this should take no longer than 2.5E–3 s for a 15 mil die, and 5E–4 s for a 7 mil die (as before, thickness increasing in proportion to the square root of the time required).

On the other hand, a lumped parameter RC model, due to the exponential nature of the equations describing its behavior, will always become linear in time as the shortest times are approached. Therefore an RC model will always fail to approximate a sqrt(t) behavior if taken to times smaller than its shortest time constant. Hence, as already discussed, if it is known that sqrt(t) behavior is a reasonable approximation to the actual behavior (typically in the 1E–6 to 1E–3 s range), but the RC time constants do not begin lower than this range, the sqrt(t) model should be used directly for short–pulse temperature estimates, otherwise severe underestimates of temperature changes will result. (A final caution, however, is that extremely short repetitive pulses require a more thorough analysis – one approach being to extend an RC model’s short–time response down into the required time scale where it accurately approximates the sqrt(t) behavior. (See AND8218/D for details.)

The following tables provide definitions and formulas useful for 1D surface heating estimates, and some typical material property values found in semiconductor packaging.

Table 3. 1D Surface Heating Formulas and Definitions

	Thermal Property
$\theta(t) = b\sqrt{t}$	Temperature rise at surface [°C/W]
$b = \frac{\theta}{A} \frac{2}{\sqrt{\pi}} \frac{1}{\eta}$	Sqrt(t) proportionality [°C/W/√s]
$\eta = \sqrt{k\rho c_p}$	Thermal effusivity [W√s/mm²/°C]
$\eta_e = \eta_1 + \eta_2$	Effective effusivity for heating at planar interface between two materials
$\tau = \frac{L^2}{\alpha}$	Characteristic time through thickness L [s]
$\alpha = \frac{k}{\rho c_p}$	Thermal diffusivity [mm²/s]

where:

Q	Heat input (W)
A	Heated area (mm²)
k	Thermal conductivity (W mm⁻¹ C⁻¹)
ρ	Density (kg mm⁻³)
c _p	Specific heat (J kg⁻¹ mm⁻¹)

Table 4. Material Properties for Short–Time Thermal Response

Material	Thermal Property	
	Diffusivity, α [mm²/s]	Effusivity, η [W√s/mm²/°C]
Silicon	52.7	0.0138
Typical mold compound	0.31	0.00126
Copper	111	0.0360
Gold	128	0.0281
Air	24.9	6.0E–6

ON Semiconductor and  are registered trademarks of Semiconductor Components Industries, LLC (SCILLC). SCILLC reserves the right to make changes without further notice to any products herein. SCILLC makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does SCILLC assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation special, consequential or incidental damages. "Typical" parameters which may be provided in SCILLC data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. SCILLC does not convey any license under its patent rights nor the rights of others. SCILLC products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the SCILLC product could create a situation where personal injury or death may occur. Should Buyer purchase or use SCILLC products for any such unintended or unauthorized application, Buyer shall indemnify and hold SCILLC and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that SCILLC was negligent regarding the design or manufacture of the part. SCILLC is an Equal Opportunity/Affirmative Action Employer. This literature is subject to all applicable copyright laws and is not for resale in any manner.

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.