



WP258 (v1.0) February 8, 2007

Considerations for Heatsink Selection Xilinx Thermal Data Application

By: Abu Eghan

The traditional approach to selecting a heatsink with a given junction-to-case thermal resistance (θ_{JC}) assumes a one-dimensional heat flow path in the package. Similarly, when the case temperature of a component is known, and a user wants to predict the junction temperature (T_j), the data sheet θ_{JC} is deployed in one dimension. In this simplified approach, the user assumes that all generated heat flows serially from the junction to the case, then across the interface into the heatsink, and is finally dissipated from the heatsink to the air stream. Unless the component is floating or is on an insulating board, this assumption cannot be true. Any appreciable heat flow into the board will distort the T_j prediction or lead to a more conservative, and perhaps expensive, heatsink choice. For most systems, such over-design might be tolerable; however, this may not always be the case. This review will point out some of the pitfalls of this one-dimensional approach and offer an example to illustrate a more complete approach.

© 2007 Xilinx, Inc. All rights reserved. All Xilinx trademarks, registered trademarks, patents, and further disclaimers are as listed at <http://www.xilinx.com/legal.htm>. All other trademarks and registered trademarks are the property of their respective owners. All specifications are subject to change without notice.

NOTICE OF DISCLAIMER: Xilinx is providing this design, code, or information "as is." By providing the design, code, or information as one possible implementation of this feature, application, or standard, Xilinx makes no representation that this implementation is free from any claims of infringement. You are responsible for obtaining any rights you may require for your implementation. Xilinx expressly disclaims any warranty whatsoever with respect to the adequacy of the implementation, including but not limited to any warranties or representations that this implementation is free from claims of infringement and any implied warranties of merchantability or fitness for a particular purpose.

Traditional Approach Explored


The illustration in [Figure 1](#) shows the basic thermal circuit used in heatsink selection. The temperatures can be represented as follows:

- T_a - Ambient Temperature
- T_j - The device junction temperature
- T_c - Case temperature of the device
- T_s - Temperature of the heatsink

With P being the Power to be dissipated, the resistances involved can be defined as follows:

- $\Theta_{JA} = (T_j - T_a)/P$; this figure of merit for components is supplied by component supplier*.
- $\Theta_{JC} = (T_j - T_c)/P$; this is supplied by component supplier*.
- $\Theta_{CS} = (T_c - T_s)/P$; this comes from the attachment material supplier.
- $\Theta_{SA} = (T_s - T_a)/P$; this is supplied by the heatsink supplier in the form of performance curves and data tabulation for various airflows, etc.

The following relationship exists between the resistances: $\Theta_{JA} = \Theta_{JC} + \Theta_{CS} + \Theta_{SA} = (T_j - T_a)/P$.



WP258_01_020607

Figure 1: Traditional Approach Illustrated

So, given the datasheet Θ_{JC} , and a known heatsink Θ_{SA} with known heatsink glue Θ_{CS} , you can predict T_j . Alternately, given a thermal budget $(T_j - T_a)$, it is possible to use the above expression to determine a heatsink with Θ_{SA} to meet a specified thermal requirement. For the most part, this linear approach works well. It works with the assumption that almost all Power P goes through the top path. This, however, may not be true on all boards or applications since a sizable amount of heat can be dissipated to the board. This is where the predictions using this approach can be inaccurate.

In general, when the board contribution is on the same order as the contribution from the top path, the prediction needs to be examined. Fortunately, in low Θ_{JC} cases, the efficiency of the paths is skewed towards the top, and therefore, the one-dimensional approach provides reasonable predictions.

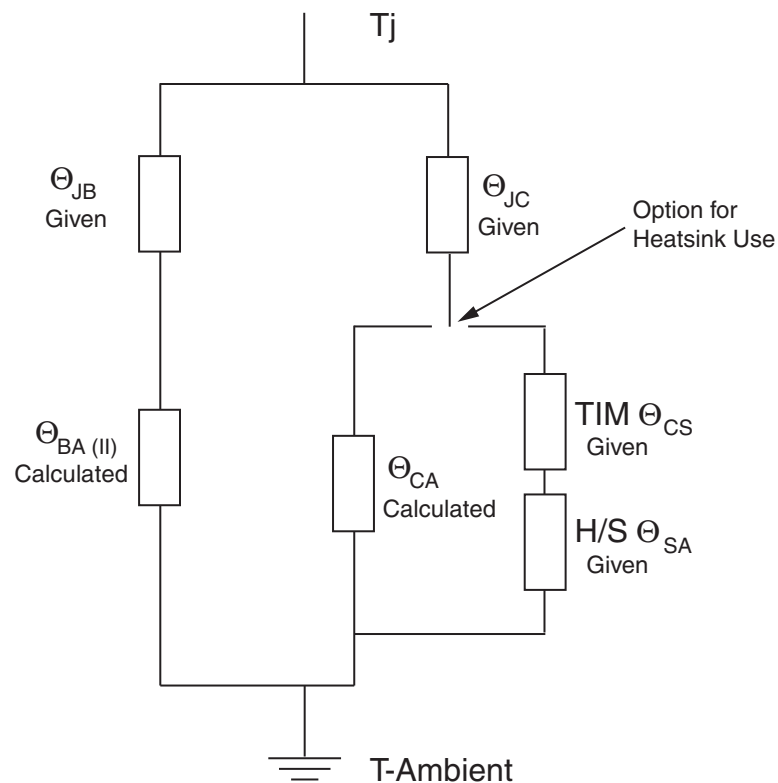
Simplified Parallel Resistor Network Model

In most applications, the package is mounted on a board. Indeed, for most high-performance FPGA components, the package substrates tend to be very efficient heat spreaders with low junction-to-board thermal resistance - Θ_{JB} . With an efficient board providing an alternate heat path, a better model might be needed to account for the power that is not through the top of the package. Obviously, if cost and over-design are not an issue, the board contribution can be ignored.

Packages can be represented by multi-resistor networks. An example of such a network is the DELPHI compact model, which Xilinx provides for high-performance FPGAs-such as Xilinx Virtex™-4 and Virtex-5 FPGAs (see [Download Center](#)). However, for this discussion, we will use a simpler 2-resistor model (2-R model) where as the supplier, we will provide both Θ_{JC} and Θ_{JB} . Given those two resistances, the package and the board can be represented by the simplified network, shown in [Figure 2](#).

Note that any resistance contributions from the sides of the package have been ignored. These side resistances are usually higher and can be ignored for simple cases such as this. Adding them in, as Delphi models do, will further improve prediction accuracy.

In the model below, the traditional path through the top (right side) has been considered. In this case, there are two options: one using heatsink, and the other without heatsink. The path through the board is also provided-left branch.



WP258_02_020507

Figure 2: The 2-Parallel Path Approach;
TIM - Thermal Interface Material; H/S - Heatsink

The challenge in this approach is to obtain the Θ_{BA} (Board-to-Ambient thermal resistance). This resistance is not supplied by component vendors, but rather depends on the environment, the board size, its in-plane conductivity, etc.

If you are using a calculating tool, the boundary conditions you set will calculate Θ_{BA} for you. Without a tool, you can either make inferences from previous experiences, or attempt some approximations with an analytical approach. On a typical 2S2P test board for area array surface mount package thermal measurement (EIA/JESD51-9, see [Ref 1]), Θ_{BA} can range from 1 to 10°C/watt depending on the component footprint, air-speed in use, and, for natural convection, the temperature gradient-to-ambient. The determination of Θ_{BA} and Θ_{CA} can be made with known heat transfer coefficients and the surface area of the component, the details of which are beyond the scope of this application note.

The advantage of the parallel path model is that it takes the board path and, therefore, the user-condition into consideration. This leads to better T_j predictions over the simpler one-dimensional approach. However, unlike the one-dimensional approach, where all the variables can be obtained from the supplier datasheets, the parallel path model requires extra effort to obtain the board-to-ambient resistance (Θ_{BA}) associated with the use conditions. Obtaining the extra data might not be that trivial and obvious for the casual user.

Case Study

Let us review how this parallel network works for four cases on the same horizontally mounted, 3 mm thick, medium size square board; 300 mm on a side with effective aggregate copper thickness of 0.575 mm. The heatsink glue (TIM) used is 0.1°C/watt (Θ_{CS}) for all heatsink cases.

- Case 1 is molded BGA type package - Xilinx device XC2V2000-FG676.
- Case 2 is a Flip-Chip - Xilinx device XC4VLX60-FF1148 with a heatsink at 400 LFM.
- Case 3 - Xilinx component XC5VLX110-FF1760 with an efficient heatsink ($\Theta_{SA} = 1.6$ °C/watt).
- Case 4 - Same component in Case 3 with reasonable, but not so efficient heatsink ($\Theta_{SA} = 3.0$ °C/watt).

Table 1 compares the predicted junction temperatures, given $T_a = 45^\circ\text{C}$ in all cases. The T_j for the traditional approach, where board contribution is ignored, is compared with the T_j using the 2-path approach (last 2 columns).

Table 1: Comparing Tj Predictions

Device		Θ_{JC} °C/watt	Θ_{JB} °C/watt	Θ_{SA} °C/watt	Watts	Tj	Tc	% Flux - Top	% Flux - Board	Mod-Tj	Θ_{JC}
						Traditional Ta = 45°C				Board Considered, Ta = 45°C	
Case 1	XC2V2000- FG676	3.2	6.7	4.6 (200LFM)	4	76.6°C	63.8°C	56%	44%	62.8°C	55.6°C
Case 2	XC4VLX60- FF1148	0.2	2.6	2.6 (400LFM)	8	68.2°C	66.6°C	64%	36%	59.9°C	58.9°C
Case 3	XC5VLX110- FF1760	0.12	2.0	1.0 (400LFM)	8	59.6°C	58.2°C	78%	22%	56.4°C	55.3°C
Case 4	XC5VLX110- FF1760	0.12	2.0	3.0 (200LFM)	12	83.6°C	82.2°C	59%	41%	67.6°C	66.8°C

Nowhere in the above cases do we have the condition where 100% heat flux is through the top. The closest, or best case (Case 3), still shows an 8 to 2 ratio of heat flux. So only 80% of the power flows through the top path. The ratio favors the heatsink path when using a more efficient heatsink and a lower Θ_{JC} . As the top flux dominates at 8:2, the traditional Tj prediction is not significantly different from the parallel path approach at 12 watts. This is why that approach works and has survived.

Case 4 is a repeat of Case 3 with reduced heatsink efficiency. Notice that the flux ratio changes from 8:2 to 6:4, and that this change impacts the Tj predictions. As the resistances of the two paths (top and bottom) approach each other, as is typical of molded BGAs (e.g., Xilinx FG, and FGG series packages) or medium priced heatsinks, the same component can exhibit significant Tj differences between the traditional approach and the parallel path approach. This is also evident in Case 1, where the Θ_{JC} of the molded package is not that great.

Predicting Tj from Monitored Tc (T-top)

In most practical cases, Tc (or T-top, temperature at the top of the package) is easily monitored and the user then tries to predict the Tj. Again, due to the fraction of power in the path of interest, there is a chance that the Tj will be overestimated. In the worst case, 100% of the power can be used with the published Θ_{JC} to make the prediction. Users need to be aware that the contribution from the 2nd will discount the power by some finite amount-so that the power will be a fraction of the P used, and subsequently lead to a lower Tj. On this board in review, the top flux can be assumed to be in the 60 to 80% range. This will change for other boards, but not significantly.

Summary

In the above examples, we have demonstrated how the two different heatsink application models can lead to differences in T_j prediction. The traditional one-dimensional approach is simple to implement, but it might lack accuracy because it ignores other paths in the system. In some cases, this methodology can lead to over-design of the heatsink. As a user implementing surface mount devices on efficient boards, you need to be aware that a parallel path through the board exists that can influence your predictions. Whether that extra path is important to consider depends on the relative contributions from that path and the effort needed to estimate the board-to-ambient resistance of your board.

References

1. EIA/JESD51-9 - Test Boards for Area Array Surface Mount Package Thermal Measurements:
<http://www.jedec.org/download/search/jesd51-9.pdf>

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
2/8/07	1.0	Initial Xilinx release.