

Thermal Characterization of

a Multi-Chip Module

Multi-chip modules (MCMs) look like single chip components, but as the name describes, they actually contain multiple chips in one package. One example of an MCM is Intel's Core i7 CPU series. These are still called CPUs, but they are actually modules that house the CPU itself, along with other components such as the memory controller and graphics processor [1]. The advantages of MCM construction are many, including reduced size vs discrete components, faster data rates and better signal quality (with less EMI) between the components on an MCM, and simpler motherboard design, since many connections are now made within the MCM [2].

MCMs do create challenges for thermal engineers however, not only because of the possible increase in power density, but because the traditional JEDEC JESD51 standard only holds true when the junction-to-case thermal resistance, Θ_{j-c} , describes the thermal path from one heat source to the case (or other reference point). In the case of an MCM with multiple junctions, a single Θ_{j-c} is not enough to determine the individual temperature of each junction. In fact, an illustration of an MCM with 6 chips (see Figure 2) shows just how complex the thermal interactions are between all of the dies and the external environment. In effect, MCMs are like miniature PCBs with the difference that the chips are not exposed to the coolant directly.



Figure 1. Intel Clarkdale with CPU and GPU [1]

One relatively quick way to approximate the junction temperatures in an MCM is to imagine an internal point that all of the heat from the individual dies flows through, at T_{INT} , and then from there, it flows through a reference point T_{REF} , then dissipated to the environment [4]. This reference point can be thought of as the case temperature of the MCM. From each individual junction to the internal point, there is a thermal resistance that is specific to that junction. This idea is illustrated

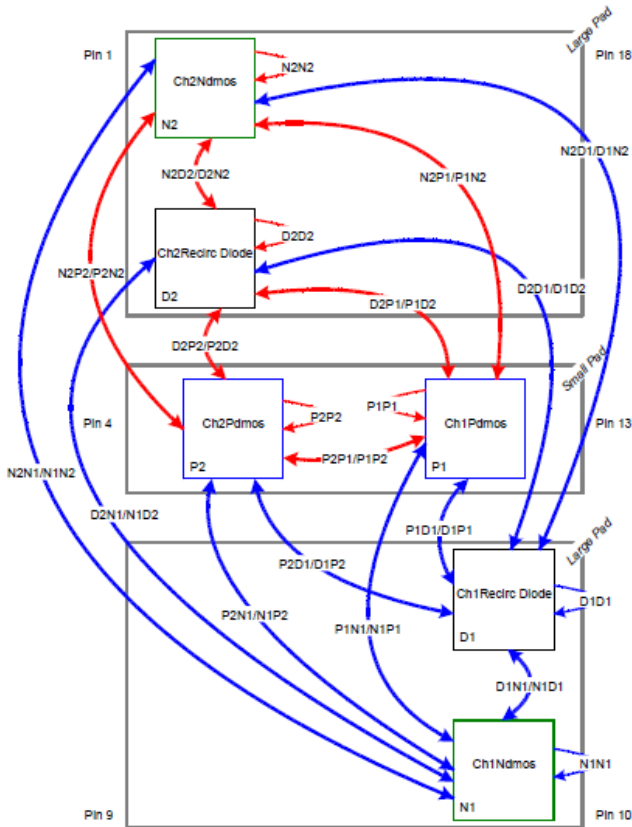


Figure 2. Thermal Interaction Between Dies in an MCM [3]

in Figure 3 below. Note that resistors between the junctions themselves are not shown, which greatly simplifies this approach, but also affects its accuracy.

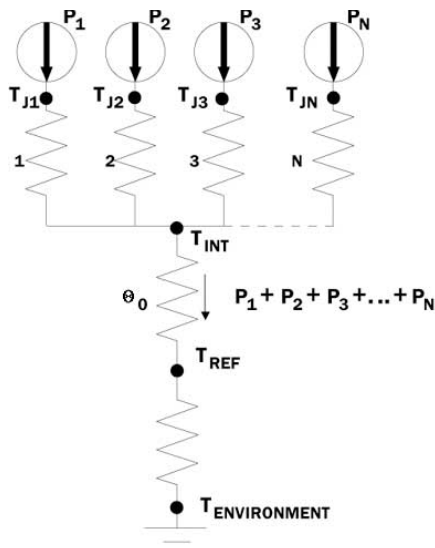


Figure 3. Simplified Thermal Resistor Network for MCM [4]

The equation for junction temperature rise over the reference point can then be written as equation (1) below. The accuracy of this equation depends largely on the thermal conductivity of the heat spreader (package lid) of the MCM. The closer this surface is to being isothermal, then the better the approximation will be.

$$\Delta T_{j-REF} = P_i \times \theta_i + P_{TOTAL} \times \theta_0 \quad (1)$$

In equation (1), the values for θ can be found using analytical analysis with simple conduction equations, when chip size and thermal conductivity values are known. Guenin compared results of this method with FEA simulations, and found errors of less than 5%, but warns that errors as large as 15% would not be unusual [4]. This is due to ambiguity associated with material properties of the package and the heat transfer coefficient of the coolant.

In order to create a more accurate method to calculate the junction temperatures in MCMs, the interactions between chips must be taken into account. The heat generated in a particular chip does not simply travel to the package case, but also co-heats the other components within the MCM. Galloway and Sutton explore the use of a superposition method to take this into account [5]. The superposition method states that the total effect of multiple heat sources can be determined by adding up the effects of turning on each heat source one at a time. One way to order these thermal interactions is to arrange these equations in a linear algebra format that results in a standard set of matrices. A_{ij} below (2) is the influence coefficient matrix, where each coefficient A_{ij} represents the effect of an individual power source P_i on the temperature rise over ambient at node j .

$$\begin{bmatrix} A_{11} & A_{12} & \dots & A_{1n} \\ A_{21} & A_{22} & \dots & A_{2n} \\ \vdots & \vdots & \dots & \vdots \\ A_{n1} & A_{n2} & \dots & A_{nn} \end{bmatrix} \begin{bmatrix} P_1 \\ P_2 \\ \vdots \\ P_n \end{bmatrix} = \begin{bmatrix} \Delta T_1 \\ \Delta T_2 \\ \vdots \\ \Delta T_n \end{bmatrix} \quad (2)$$




Figure 4. MCM Thermal Test Vehicle (copper lid removed) [5]

Galloway and Sutton applied the superposition method to an MCM with five components: An ASIC (Application Specific Integrated Circuit) surrounded by four memory modules. The MCM had a copper lid for heat spreading, and is shown in Figure 4 below.

The MCM test vehicle was operated at eight different power combinations and with three different heat sinks. Non-linear optimization was used to determine the best fit to the experimental data, and one example of the results is shown in equation (3) below. Note that this equation allows the calculation of ΔT_{j-a} for a given power dissipation only with a specific heat sink at a specific air velocity. In this case, the heat sink had a resistance of approximately $\Theta_{c-a} = 1.88^\circ\text{C/W}$ at 1.0 m/s.

$$\begin{bmatrix} 2.11 & 1.98 & 2.29 & 1.51 & 1.89 \\ 1.75 & 8.48 & 2.08 & 1.34 & 1.44 \\ 1.84 & 2.05 & 8.04 & 1.33 & 1.75 \\ 1.76 & 1.83 & 1.86 & 8.07 & 1.70 \\ 1.84 & 1.82 & 2.14 & 1.57 & 8.34 \end{bmatrix} \begin{bmatrix} P_A \\ P_{M1} \\ P_{M2} \\ P_{M3} \\ P_{M4} \end{bmatrix} = \begin{bmatrix} \Delta T_A \\ \Delta T_{M1} \\ \Delta T_{M2} \\ \Delta T_{M3} \\ \Delta T_{M4} \end{bmatrix} \quad (3)$$

From an application point of view, it is more useful to know what the Θ_{j-c} value is for a component; that way, any heat sink may be specified that will meet the thermal requirements. To achieve this, equation

(3) can be thought of as the sum of two effects, the ΔT_{j-c} which includes thermal interactions within the MCM itself, and the ΔT_{c-a} , which depends on the heat sink used. This may be written as equation (4), where $[A_{ij}][P_j]$ represents ΔT_{j-c} , and the second term represents ΔT_{c-a} :

$$T_j = [A_{ij}][P_j] + \theta_{ca} \times \sum P_j + T_a \quad (4)$$

Using the junction and case temperature data gathered previously, non-linear optimization was again used to approximate $[A_{ij}]$, which gave the coefficients shown below in equation (5).

$$[A_{ij}] = \begin{bmatrix} 0.0793 & 0.0074 & 0.0603 & 0.0350 & 0.0704 \\ -0.173 & 6.459 & -0.0080 & -0.0248 & -0.0830 \\ -0.156 & 0.0469 & 5.675 & -0.00204 & 0.0275 \\ -0.174 & -0.0455 & -0.0871 & 6.503 & 0.0214 \\ -0.157 & -0.0354 & 0.0188 & 0.0882 & 6.471 \end{bmatrix} \quad (5)$$


Note that in equation (5), the coefficients on the diagonal are much higher. This is because these values represent Θ_{j-c} for the individual components in the MCM. The other values are lower because they represent thermal paths between the components. The average temperature error was found to be less than 5%, but some particular configurations had higher error values. In particular, with one of the heat sinks, the estimated T_j of the memory modules could be as much as 29% different compared to experimental data.

This was due to the fact that one heat sink was made of copper, compared to the other two heat sinks, which were made of aluminum. Despite efforts to calculate values in Equation (5) independent of the heat sink used, it was impossible to eliminate the effects of the heat sink on heat spreading between the components. For the most accurate results, the influence matrix should be based on a heat sink with similar thermal performance as the heat sink in the end application.

Multiple heat sources on the same carrier as seen in MCMs and PCBs pose a unique thermal calculation challenge. It can be seen that even though the superposition method has the potential for very good accuracy, it is still dependent on specific boundary conditions. If these conditions are known, and they are not changing, then junction temperatures may be predicted with good degree of accuracy. The method shown in equation (1) seems convenient, as Θ_j are calculated independently of the heat sink, but its accuracy is limited because it does not take co-heating of components into account. Which method is more suitable to an application depends on how much is known about the boundary conditions, since the superposition method requires more detail. In the end, it should be remembered that both methods are still approximations, and should be validated by laboratory testing.


References:

1. Worrel, J., "Intel migrates to desktop Multi-chip Modules (MCMs) with 14nm Broadwell," <http://www.fudzilla.com/home/item/26786-intel-migrates-to-desktop-multi-chip-module-mcm-with-14nm-broadwell>, [Mar. 26, 2014].
2. "Multichip Module Design Handbook," Edition 3, Europractice MCM Service, 1996.
3. "Thermal Characterization of Multi-Die IC Products," Fairchild Semiconductor Corporation, 2012.
4. Guenin, B., "Thermal Calculations for Multi-chip Modules," <http://www.electronics-cooling.com/2002/11/thermal-calculations-for-multi-chip-modules/>, [Mar. 25, 2014].
5. Galloway, J., Sutton, M., "Extracting Thermal Data from High Power MCM Packages," Amkor Technology, Chandler, AZ, 2007.



CHILLERS

THERMOELECTRIC ASSEMBLIES




HAND HELD ANEMOMETER

HEAT PIPES

LEAK DETECTORS

FLOW METERS



THERMOELECTRIC MODULES

HEAT EXCHANGERS



ATS ADVANCED THERMAL SOLUTIONS, INC.
Innovations in Thermal Management®