

A Multi-stage Thermal Management Strategy for 3D Multicores

Dipika Suresh, Amit Kumar Singh and Akash Kumar

Department of Electrical and Computer Engineering,

National University of Singapore, Singapore

Email: akash@nus.edu.sg

Abstract - 3D integration technology has the potential to enhance IC performance, improve functionality and lessen wiring of ICs. However, it poses several challenges, where the key challenge is heat generation from internal active layers due to power dissipation. To mitigate this challenge, thermal aware design has become a necessity. Towards thermal aware design, this paper proposes a two stage design technique. In the first stage, a temperature-power thermal model is created to calculate power dissipated by an IC at an input temperature. The proposed model calculates power dissipated by 2D and 3D ICs with an average error of 0.37% and 25% respectively. Power calculation helps in process variation, validation of power models and minimization of temperature gradients. In the second stage, thermal aware mapping is performed for the ICs. For thermal aware mapping, three mapping algorithms are proposed to account for different resource (processor) availability scenarios. Each algorithm utilizes temperature-power thermal model (from the first design stage) to map applications to processing elements in a 3D IC. The proposed two stage design technique performs faster temperature to power calculations than existing techniques. It provides a simplified approach to mapping compared to existing techniques by utilizing power dissipated by processing elements to map applications.

I. INTRODUCTION

3D integration technology entails the design and manufacture of ICs with multiple layers of active devices. This technology has the potential to enhance chip performance, improve functionality and lessen wiring. Additionally, 3D integration technology reduces IC manufacturing cost, improves device packing density and decreases the distances travelled by signals in the IC [2].

The benefits of 3D integration technology can be availed after handling the numerous challenges it poses, the key challenge being heat generation from internal active layers due to power dissipation. Increase in power density per unit volume creates internal hotspots and large temperature gradients in the chip. To mitigate this challenge, thermal aware design has become a necessity. Thermal aware design is a design methodology that uses temperature as a guideline throughout the design flow.

There are several existing thermal aware design techniques. One such design technique calculates the temperatures of processors in an IC by creating a compact thermal model and taking power dissipated by the IC as input [10]. A drawback of this technique is that input power is difficult to measure and can only be roughly approximated using power simulators. Another design technique is the inverse

heat conduction process (IHCP) which utilizes input temperatures of an IC to calculate power dissipated by it. IHCP eliminates the need to physically measure power and relies instead on temperature measurement using infra-red sensors. However, the current approaches to IHCP, which include image processing algorithms [11] and spatially resolved imaging of microprocessor power [6], take a considerably long time to calculate power. Thermal aware mapping is a popular thermal aware design technique used to map tasks to IC processors while accounting for thermal factors such as peak temperature of processors, task load etc. Some techniques evaluate thermal distribution in an IC and perform mapping at run-time while others do the same at design time. Popular mapping techniques include clock gating and dynamic voltage frequency scaling which incorporate protection mechanisms (stalling, decreasing V/f etc) when the temperature of a core exceeds the acceptable limit. The primary drawback of such mapping techniques is that they do not account for task deadlines and throughput while mapping.

This paper proposes a two stage thermal aware design technique. In the first stage, by utilizing temperature-power relations, power dissipated by each processor in an IC is calculated for an input temperature. This can be utilized to calculate power dissipated by each processor when temperature gradients are minimized (by setting similar input temperature for all processors). The power values computed can be used to perform thermal aware mapping of tasks. In the second stage, thermal aware mapping is performed by utilizing the power values computed in the first stage. Three mapping algorithms are proposed, which account for different resource (processor) constraints. Algorithm one performs mapping when the number of tasks to be mapped is less than or equal to the number of processors in the IC. This is done by comparing power consumption of each task with power dissipated by each processor and mapping task with lowest power consumption to processor with minimum power dissipation. Algorithm two performs mapping when the number of tasks to be mapped is greater than the number of processors in the IC. Since number of tasks is greater than the number of processors, tasks with the lowest power consumption are combined to form a super task. Mapping is performed between task/super task and processors in the same way as algorithm one. Algorithm three derives a simplified relation between temperature and power. The simplified temperature-power relation is derived by observing how power calculated varies with input temperature. This relation along with algorithm one or two (depending on number of processors and tasks) is utilized to perform mapping.

II. RELATED WORK

There are several existing thermal aware design techniques. One such design technique relies on thermal models to simulate thermal aspects of an IC architecture [10]. HotSpot is a thermal modelling tool which utilizes this design technique. HotSpot generates compact thermal models using the principle of thermal-electric duality [10]. It can be used in early architecture design stages where a detailed IC layout is not available. With Hotspot, reasonably accurate spatial and temporal temperature variations of silicon die can be obtained to help in efficient design decisions during early design stages. Moreover, HotSpot makes it possible to study thermal evolution over long periods of full-length applications [10].

Another consideration for thermal aware design is the inverse heat conduction process (IHCP). In IHCP, power maps are solved from thermal maps, hence removing the need to directly make power measurements. Calculation of power maps assists in validation of power models, and supports new techniques to manage runtime power dissipation on a per-chip basis. One suggested approach to IHCP is the utilization of image processing algorithms to generate power maps from temperature maps [11]. One of the main drawbacks of this approach is that treating temperature and power maps as images loses micro architectural information such as functional blocks. Another approach is the SIMP (spatially-resolved imaging of microprocessor power) methodology from IBM which captures thermal maps from IR cameras and solves power maps through LS (least squared) [6]. However, LS problem formulation and applied constraints have not been discussed at length.

Thermal aware mapping and scheduling for 3D IC architecture is another well researched thermal aware design technique. Thermal-aware mapping and scheduling can be roughly split into dynamic (runtime) techniques and static (design time) techniques. Dynamic techniques measure or estimate the current thermal distribution in the IC and take actions to minimize temperature gradients. Some popular dynamic thermal management techniques are described below.

- *Clock gating*: In clock gating, each processor in the IC runs at its default highest frequency and voltage setting until the processor reaches its threshold temperature. When this occurs, the processor is stalled and its clock is gated to reduce its power consumption [4].
- *Dynamic voltage frequency scaling with temperature trigger (DVFS - TT)*: In DVFS-TT, the V/f setting of a processor is reduced to the next lower V/f value when the temperature of the processor exceeds its threshold [4]. A primary drawback of DVFS-TT and clock gating is that these techniques can potentially have an unexpected impact on software execution, for example cause tasks to miss deadlines.
- *Temperature balancing of cores*: In this technique, at each scheduling point, the scheduler sorts the power consumption of all tasks and the current temperature of each core. It then assigns the task with the highest power to the coolest core [12].

Static mapping approaches aim at finding a thermal-aware mapping at design time, by using a model of the physical chip, or by using general knowledge about the thermal behaviour of 3D ICs. Some popular static mapping techniques are described below.

- *Temperature aware floorplanning*: This technique attempts to explore how changing the layout / floorplan of an architecture impacts its thermal properties [7].

- *Genetic algorithms*: [1] suggests a static mapping technique based on genetic algorithms. This technique pays attention to thermal and communication considerations while designing the architecture. However, throughput requirements are not taken into account.
 - *Thermal aware mapping for streaming applications*: This method provides an integrated thermal-aware approach for mapping streaming applications. Thermal characteristics of the 3D IC are extracted from and supplied to a resource allocation algorithm. This approach takes communication and throughput considerations into account during mapping [5].
- The above mapping techniques tend to alter IC specifications. They often do not account for communication delays and task deadlines. The two stage thermal aware design technique proposed in this paper proposes a design time mapping (static mapping) which does not alter IC properties. As it is a design time approach, it accounts for task deadlines during mapping (tasks are assumed to be independent).

III. TEMPERATURE-POWER MODEL

This section covers the proposed temperature-power model for 2D and 3D IC architecture.

A. 2D IC architecture

A. Steady state analysis: By thermal-electric duality, electric current flow through a resistance is analogous to heat flow through a thermal resistance. The circuits used to define heat flow are equivalent to those describing current flow. According to [9], based on thermal-electric duality, power and temperature are related as follows:

$$A_r \cdot P = T \quad (1)$$

where A_r is the thermal resistance matrix, P is the power map matrix, T is the temperature map matrix.

The above temperature-power relation for steady state has been implemented in Hotspot thermal modelling tool. An important consideration is the initialization of temperatures of "package" layers such as thermal interface layer, heat spreader and heat sink. Power dissipated by silicon die depends on the temperatures of these layers because the flow of heat in the IC is from silicon die to interface layer to heat spreader and heat sink. The initial temperatures for these layers that produce most accurate power is discussed in the experimental results section.

B. Transient analysis: Similar to steady state, thermal-electric duality can be used to generate an equivalent circuit to represent heat flow at a particular point in time t (transient). According to [9], based on thermal-electric duality, power and temperature are related as follows:

$$\left(\frac{dT}{dt}\right) \cdot C + T(t) \cdot A_c = P \quad (2)$$

where $\frac{dT}{dt}$ is the rate of change of temperature, C is the capacitance matrix, A_c is the thermal conductance matrix, P is the power map matrix, T is the temperature map matrix.

Another unknown in transient analysis is the rate of change of temperature $\left(\frac{dT}{dt}\right)$ [9]. The value of $\frac{dT}{dt}$ can be approximated by the Central Difference Method. As per this method, the first derivative of a function $T(t)$ can be approximated as:

$$\frac{dT}{dt} = \frac{T_{+1} - T_{-1}}{2h} \quad (3)$$

where $T_{+1} = T(t + h)$, $T_{-1} = T(t - h)$, h is the sampling interval (time interval between two measured temperatures). The central difference method is preferred over other methods such as the forward difference method due to greater accuracy in result [9]. The proposed temperature-power relation for transient analysis has been implemented in HotSpot. As in the case of steady state, temperatures of thermal interface layer, heat spreader and heat sink need to be initialized here as well. Sampling interval is another factor that impacts accuracy of power values calculated. Therefore, it should be chosen appropriately, which is discussed in the experimental results section.

B. 3D IC Architecture

A. Steady state analysis: Thermal-electric duality can be used to obtain temperature-power relation for 3D architecture as well. Thermal conductance can be calculated for every layer in an IC and used to generate power maps. To ease the calculation of thermal conductance of layers in the IC, every layer can be divided into multiple uniform grid cells. The thermal conductance of each grid cell is taken as the sum of the thermal conductance between the cell and its neighbours. Apart from thermal conductance, the initial temperatures of surrounding layers impact power dissipation of a layer. By utilizing the grid approach, the impact of temperature of neighbouring layers can be examined using the weighted temperature of a grid cell. The weighted temperature of a grid cell is the sum of the temperatures of its neighbours. By thermal-electric duality, the following represents relation between power dissipated and temperature of a grid cell:

$$P = T * csum - wsum \quad (4)$$

where T is the temperature of the grid cell, P is the power dissipated by the grid cell, $csum$ is the thermal conductance of the grid cell and $wsum$ is the weighted temperature of the grid cell. Power dissipated by an individual component is calculated by summing up the powers of the grid cells within that component. The proposed temperature-power relation for 3D architecture has been implemented in HotSpot, which supports grid level analysis of ICs. As in 2D architecture, temperatures of thermal interface layer, heat spreader and heat sink need to be initialized. The temperature of a thermal interface layer lying between two layers of Si die is initialized as the average of the temperatures of the Si die layers above and below it. Similarly, the temperature of a thermal interface layer lying between a layer of Si die and the heat spreader / heat sink is initialized as the average of the temperatures of the Si die layer above it and the ambient temperature.

B. Transient analysis: A time dependent relation between temperature, power and thermal conductance is used to calculate the transient power of each Si die layer. Similar to steady state analysis, the impact of surrounding layers on Si die is accounted for using weighted temperatures. The details of the implementation of transient analysis in HotSpot is omitted due to limitations on the number of permissible pages for the paper.

IV. MAPPING TECHNIQUES

A powerful consequence of utilizing the temperature-power relations discussed above to calculate power maps from temperature maps is that the power maps generated can assist in mapping of tasks on the processing elements (PEs) of an IC. All PEs in an IC

can be set to uniform temperatures to minimize temperature gradient. The power dissipated by PEs at these uniform temperatures can be calculated using the implementation in HotSpot discussed previously. Thereafter, tasks can be mapped to processing elements based on their power dissipation and the task power consumption.

Three mapping algorithms are proposed, which can be applied based on the number of PEs and tasks to be mapped.

One-on-one mapping (OM): This algorithm can be applied when the number of tasks to be mapped is less than or equal to the number PEs in the IC. The algorithm considers the following:

- *Temperature setting:* A uniform temperature is chosen for all PEs in the IC.
- *Power calculation by IHCP:* With the chosen temperature as input for all PEs of the IC, the power dissipated by each PE is calculated using temperature-power relations discussed earlier.
- *Sorting of power dissipated:* The power dissipated by PEs is sorted in ascending order.
- *Sorting of power consumption of tasks:* The power dissipated by tasks is sorted in ascending order.
- *Mapping tasks on PEs:* Task with smallest power consumption is mapped to PE with minimum power dissipation. Similarly, each task is mapped to a PE in the IC.

Many-to-one mapping (MM): This technique can be applied when number of tasks is greater than the number of PEs in the IC. The first 4 steps of this algorithm are the same as that of OM. Additionally, this algorithm includes the following:

- *Task combination:* The two tasks with the smallest power consumption are combined together to form a "super task". In a similar manner, tasks are combined to form super tasks until the summation of the number of tasks and super tasks equals the number of PEs.
- *Mapping tasks on PEs:* Task/super task with smallest power consumption is mapped to PE with minimum power dissipation. Similarly, each task/super task is mapped to a PE in the IC. Each PE has a queue of tasks to be executed. When a super task is mapped to a PE, the task within the super task which has a closer deadline will be executed first and the task which has a deadline that is further away will be put in the task queue and executed only when the previous task's execution is complete.

Direct mapping (DM): To obtain a more direct relation between power and temperature, power values for a particular IC architecture are calculated in HotSpot (using temperature-power relations discussed above) for a set of uniform temperatures. A relation between power and temperature is drawn based on how power calculated changes with a change in input temperature. During mapping, a uniform temperature can be set and using the relation derived, power can be easily calculated from temperature. With calculation of power dissipated simplified, tasks can be mapped to processors using either of algorithms described above.

V. EXPERIMENTAL RESULTS

A. Validation of temperature-power model

2D Architecture

A. Experimental setup

To evaluate the accuracy of power calculated using the temperature-power relations discussed previously, HotSpot thermal modelling tool

Parameter	Value
Initial temperature	333.15
Ambient temperature	318.15
Si thermal conductance	150
Si specific heat	1.75×10^6
Heat sink side	0.06
Heat sink thickness	0.0069
Heat sink conductivity	400
Heat sink specific heat	3.55×10^6
Interface thickness	2×10^{-6}
Interface conductivity	4
Interface specific heat	4×10^6

TABLE I: Physical Properties and HotSpot Parameters (all parameters in standard units)

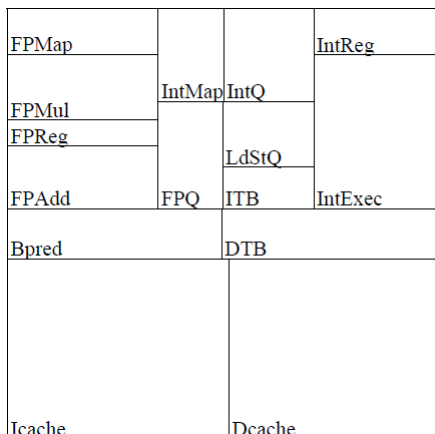


Fig. 1: Rough floorplan of Alpha 21364 microprocessor

(with the temperature-power relation implemented) is utilized. Table 1 lists the important physical properties and HotSpot parameters used for the HotSpot runs. The input temperature used for each run is obtained from HotSpot using compact thermal models (calculation of temperature from power) [10]. The 2D IC used for the HotSpot runs is the Compaq Alpha 21364 microprocessor, the layout is shown in Figure 1 [10]. The microprocessor consists of numerous functional units such as cache, floating point registers, floating point adders, load-store queue etc. The power dissipated by each functional unit is calculated during a HotSpot run. The power calculated is compared with power simulated by Wattch power simulator [3] for the Alpha 21364 microprocessor. The power simulated by Wattch simulator shall be referred to as expected power in the upcoming sections of the paper.

B. Steady state analysis

Figure 2 shows the power difference between calculated power and expected power in each functional unit for the Alpha 31264 microprocessor. It is observed in Figure 2 that the difference is close to zero. Thus, calculated power is close to expected power, which validates temperature to power calculation in steady state.

As mentioned in the previous section, HotSpot considers an IC to have "package" layers consisting of a thermal interface layer (iface), heat spreader and heat sink layer in addition to a Silicon (Si) die layer. If the temperatures of the package layers are not provided as input, they need to be approximated. Iface layer has the same




Fig. 2: Expected power vs Calculated power




Fig. 3: Expected power vs Calculated power - different iface temperatures

layout as the Si die layer. Its temperature is approximated using the temperature of Si layer above it. Figure 3 depicts the difference between calculated power and expected power in each functional unit for various approximations to iface temperature. Each case in Figure 3 represents the difference between calculated and expected power for a particular iface temperature approximation. For example, case 2 approximates iface temperature as temperature of Si die layer minus 0.5. From Figure 3 it is observed that an approximation of iface temperature as temperature of Si minus 1.5 produces the least error in calculated power.

Therefore, based on Figure 2 and Figure 3, the following conclusions are drawn:

- Temperature to power calculation for steady state analysis is quite accurate.
- Temperature of thermal interface has a significant impact on calculated power.
- If thermal interface temperature is not available, it can be approximated as temperature of Silicon die minus 1.5.

C. Transient analysis

The accuracy of the temperature-power relation has been validated for transient analysis as well. Similar parameters as those of steady state are considered. Additionally, another parameter needs to be set, namely the sampling interval. The sampling interval for this run is set to $3.33 \mu s$. Figure 4 shows the power difference between

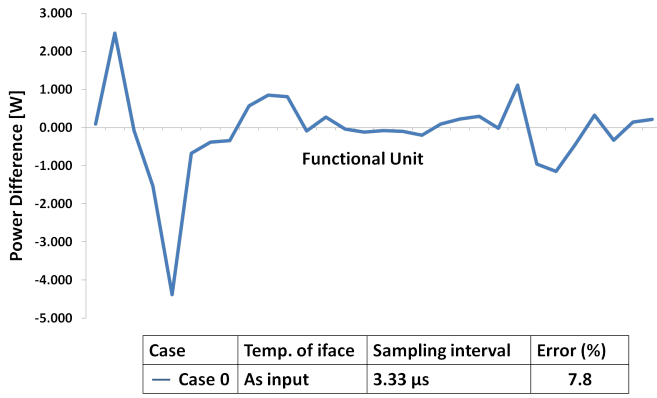


Fig. 4: Expected power vs Calculated power

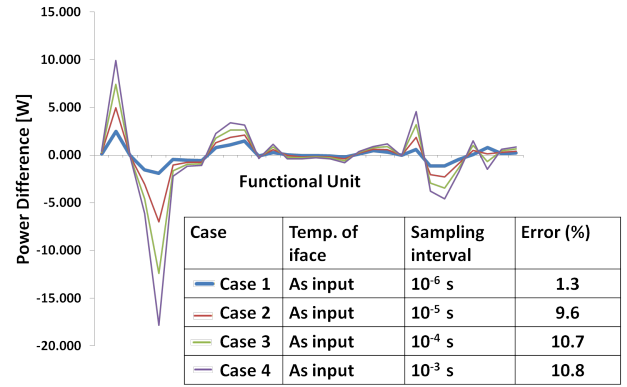


Fig. 6: Expected power vs Calculated power - different sampling intervals

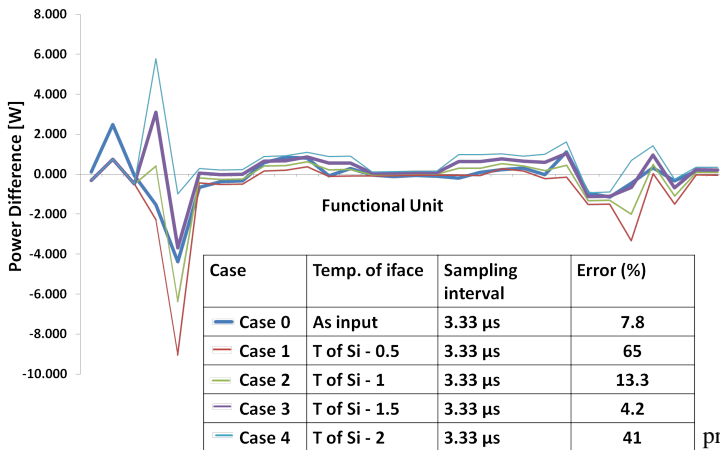


Fig. 5: Expected power vs Calculated power - different iface temperatures




Fig. 7: 3x3x3 ARM multicore

calculated power and expected power in each functional unit for the Alpha 31264 microprocessor. It is observed in Figure 4 that the difference is close to zero. Thus, calculated power is close to expected power, which validates temperature to power calculation for transient analysis. Figure 5 shows the difference between calculated power and expected power in each functional unit for different approximations of iface temperature (sampling interval set to 3.33 μ s). As seen from Figure 5, approximating iface temperature as temperature of Si minus 1.5 calculates most accurate power. Figure 6 shows the difference between calculated power and expected power in each functional unit for different sampling intervals (temperature of iface is taken as input). It is observed in Figure 6 that error decreases as duration of sampling interval decreases. From Figure 4, Figure 5 and Figure 6, in addition to previous conclusions drawn, the following conclusion is drawn:

- Sampling interval has an impact on power calculated. As sampling interval increases, accuracy of power calculated decreases.

3D Architecture

A. Experimental setup

HotSpot is used to evaluate the accuracy of temperature-power relations for 3D architecture. Table 1 lists the important physical

properties and HotSpot parameters used for the HotSpot runs. The input temperature used for each run is obtained from HotSpot using compact thermal models (calculation of temperature from power) [10]. The 3D IC used for the HotSpot runs is 3x3x3 ARM multicore, the layout is shown in Figure 7 [5]. The IC consists of three layers of Si die. Each Si die layer consists of 9 cores arranged as a 3x3 grid. The power dissipated by each core is calculated during a HotSpot run. The power calculated is compared with power simulated by Watch power simulator for 3x3x3 ARM multicore (expected power).

B. Steady state analysis

Figure 8 shows the power difference between calculated power and expected power in each core of the 3x3x3 ARM multicore. Figure 8 indicates that calculated power is not significantly close to expected power. The primary reason for this is that interface layers, heat spreader and heat sink layers have a greater impact on Silicon layers power for 3D architecture. The greater the temperature difference between the Si die layer and the interface layers surrounding it and / or the heat sink (if it is the bottomest layer in a 3D IC), greater is the heat flow and power dissipated.

To improve accuracy of power calculated, temperatures of thermal interface layers (iface) and heat spreader (hsp)/ heat sink (hs) are approximated. The temperature of iface is approximated as the average of temperatures of Si layers above and below it. The temperature of hsp and hs is approximated as the average of temperature of Si layer above it and ambient temperature. Figure 9 indicates which approximation of iface, hsp and hs produces most accurate calculated power values. As indicated in the graph and percentage error calculated, if temperature of thermal interface layer




Fig. 8: Expected power vs Calculated power

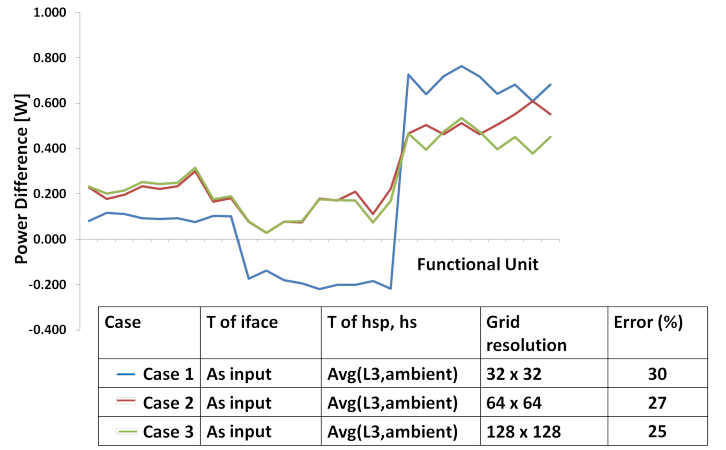


Fig. 10: Expected power vs Calculated power - different grid resolutions




Fig. 9: Expected power vs Calculated power - different iface, hsp, hs temperatures




Fig. 11: 2x2x2 ARM multicore

is not taken in as an input, it can be approximated as the average of the temperatures of the Si die layers above and below it. Moreover, if temperature of heat spreader/heat sink is not taken as input, it can be approximated as average of temperature of Si die layer above it and ambient temperature.

Figure 10 shows the difference between calculated power and expected power in each core for different grid resolutions (grid cell size). It is observed in Figure 10 that error decreases as grid size decreases. Based on Figure 8, Figure 9 and Figure 10, the following conclusions are drawn:

- Temperature to power calculation for steady state analysis for 3D architecture is quite accurate.
- Temperature of thermal interface has an impact on power calculated.
- If thermal interface temperature is not available, it can be approximated as average of the temperatures of the Si die layers above and below it.
- Temperature of heat spreader and heat sink has an impact on power calculated.
- If heat spreader/heat sink temperature is not available, it can be approximated as average of the temperature of the Si die layer above it and ambient temperature.
- Grid size has an impact on power calculated. As grid size increases, accuracy of power calculated decreases.

B. Validation of mapping techniques

A. Experimental setup

The proposed mapping techniques are validated by considering the mapping of eight independent tasks onto a 2x2x2 ARM multicore. This IC has two active layers, each layer consisting of four cores. The layout of the IC is given in Figure 11 [12]. Layer 1 is further away from the heat sink and layer 2 is closer to the heat sink. Input temperatures, as those depicted in Figure 11, are given as input to HotSpot tool to calculate power dissipated at each core of the IC. The 8 tasks to be mapped have loads of 10 W, 20 W, 20 W, 30 W, 35 W, 40 W, 50 W and 60 W. The proposed mapping will be validated against mapping techniques provided in [12] and [5].

B. One-on-one mapping

Power dissipated by each core is calculated by HotSpot. The power dissipated is sorted in ascending order. Table 2 indicates the rough mapping performed by OM. Based on OM, about 27% of the power is dissipated in layer 1 (further away from heat sink) and 73% of the

Core	Approximate Power (W)	Mapped task power (W)
A	1.31	20
B	3.245	30
C	1.725	20
D	0.69	10
E	24	35
F	28	60
G	24.8	40
H	27.2	55

TABLE II: Mapping by OM




Fig. 12: Power vs. Temperature

power dissipated in layer 2 (closer to the heat sink). This aligns with the mapping technique proposed in [5] which assigns 25% of power in layer 1 and 75% of power in layer 2. Moreover, as per [12], tasks 10 W, 20 W, 20 W and 30 W are mapped to layer 1 and tasks 35 W, 40 W, 55 W and 60 W are mapped to layer 2. Hence, OM performs mapping in accordance with these two mapping techniques.

C. Mapping by DM

A more direct relation between temperatures and powers, which can be utilized for mapping, is derived by calculating power dissipated by a particular architecture at different temperatures. The example below illustrates how a relation between temperature and power can be drawn for a 2x2x2 ARM multicore IC.

Example

The temperature-power relation implemented in HotSpot for 3D architecture is run 5 times. The following set of conditions remain the same for all runs:

- Ambient temperature = 300 K.
- Initial temperature of thermal iface 1 = Average of {(average temperature of Si layer 1) , (average temperature of Si layer 2)}
- Initial temperature of hsp & hs = Average of {(average temperature of Si layer 2) , ambient temperature}
(The temperatures of iface and hsp/hs are set as above because these approximations calculate most accurate power as deduced in previous sections)
- Grid resolution = 128x128 (high grid resolution decreases error in calculated power)

The only parameter that varies across the 5 runs is the input temperature. The following input temperatures are provided for the 5 cases:

Case 1: Input temperature = 330 K,

Case 1: Input temperature = 335 K,

Case 2: Input temperature = 340 K,

Case 3: Input temperature = 345 K,

Case 4: Input temperature = 350 K,

Figure 12 depicts power calculated for a processor in the IC varies with the input temperature. A rough linear relation can be deduced between power and temperature for the 2x2x2 ARM multicore IC. This relation eases the calculation of temperature from power and power from temperature. This relation can be applied to either mapping algorithm described previously.

Similar to OM, in DM (as indicated in Table 3), 25% of the power is dissipated in layer 1 and 75% of the power is dissipated in layer 2. This is in accordance with [5] and [12]. Hence, DM is verified.

Core	Approximate Power (W)	Mapped task power (W)
A	25.8	10
B	27.3	20
C	28.15	30
D	26.23	20
E	28.16	35
F	29.7	55
G	30.3	60
H	28.5	40

TABLE III: Mapping by DM

VI. CONCLUSIONS AND FUTURE WORK

This paper proposes a two stage thermal aware design technique. The first stage is a temperature-power model that calculates power maps from temperature maps. The proposed model calculates power dissipated by 2D and 3D ICs with an average error of 0.37% and 25% respectively. In the second stage, thermal aware mapping is explored. Three mapping algorithms are proposed which provide a direct and simplified approach to thermal aware mapping for 3D IC architecture.

In the future, improvements will be made to the accuracy of the temperature-power model. A possible way of achieving higher accuracy is by implementing least square optimization as suggested by [9]. The model could be extended to transient analysis of 3D architecture. With regards to mapping techniques, emphasis will be placed on mapping with constraints on performance. Moreover, inter-task communications will be accounted for in the future.

REFERENCES

- [1] Addo-Quaye, C. (2005). Thermal-aware mapping and placement for 3-D NoC designs . SOC Conference, 2005. Proceedings. IEEE International, 45894759, 25 - 28 .
- [2] Awasthi, M., Venkatesan, V., & Balasubramonian, R. (2007). Understanding the Impact of 3D Stacked Layouts on ILP. Journal of Instruction-Level Parallelism, 9, 1-27.
- [3] Brooks, D., Tiwari, V., & Martonosi, M. (2000). Watch: a framework for architectural-level power analysis and optimizations (Vol. 28, No. 2, pp. 83-94). ACM.
- [4] Coskun, A., Ayala, J., Atienza, D., Rosing, T., & Leblebici, Y. (2009). Dynamic thermal management in 3D multicore architectures. Design, Automation & Test in Europe Conference & Exhibition, 1410 - 1415 .
- [5] Cox, M., Singh, A., Kumar, A., & Corporaal, H. (2012). Thermal Aware Mapping of Streaming Applications on 3D Multi-Processor Systems.
- [6] Hamann, H., Weger, A., Lacey, J., Hu, Z., Bose, P., Cohen, E., et al. (2007). Hotspot-limited microprocessors: Direct temperature and power distribution measurements. IEEE J. Solid-State Circuits, 1, 56-65.
- [7] Han, Y., Koren, I., & Moritz, C. A. (2005). Temperature Aware Floorplanning. Workshop on Temperature Aware Computer Systems.
- [8] Huang, W., Brown, C. L., Ghosh, S., Veluswamy, S., Sankaranarayanan, K., Skadron, K., et al. (2006). HotSpot: a compact thermal modeling methodology for early-stage VLSI design . IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 14, 501 - 513.
- [9] Qi, Z., Meyer, B., Huang, W., Ribando, R., Skadron, K., & Stan, M. (2010). Temperature-to-power mapping. IEEE International Conference on Computer Design (ICCD), 384 - 389.
- [10] Skadron, K., Stan, M. R., Sankaranarayanan, K., Huang, W., Velusamy, S., & Tarjan, D. (2004). Temperature-aware microarchitecture: Modeling and implementation. ACM Transactions on Architecture and Code Optimization, 1, 94 - 125.
- [11] Wang, X., Shakouri, A., Farsiu, S., & Milanfar, P. (2007). Extraction of power dissipation profile in an IC chip from temperature map. Annual Thermal Measurement, Modeling and Management Symposium, 51-56.
- [12] Zhou, X., Yang, J., Xu, Y., Zhang, Y., & Zhao, J. (2010). Thermal-Aware Task Scheduling for 3D Multicore Processors. IEEE Transactions on Parallel and Distributed Systems, 21(1), 60-71.