Physical Characterization of Steady-State Temperature Profiles in Three-Dimensional Integrated Circuits

Sumeet S. Kumar, Amir Zjajo, Rene van Leuken
Circuits and Systems Group
Delft University of Technology
The Netherlands
Email: {*s.s.kumar, a.zjajo, t.g.r.m.vanleuken*}*@tudelft.nl*

***Abstract*—The thermal performance of three-dimensional integrated circuits is influenced by a number of design and technology parameters. However, the relationship between these parameters and thermal behaviour of die stacks is complex and not well understood. In this paper, we perform a detailed evaluation of the influence of stack composition and depth, thickness of dies, physical location of power dissipating elements and stack power density on steady-state temperature profiles. We examine how each of these parameters affects heat spread within 3D ICs, and highlight the causes for hotspot formation. The results of our analysis illustrate the implications of effective thermal conductivity on temperature sensing zones on dies, and the significant impact of stack power density on overall operating temperature.**

I. Introduction

Three-dimensional integrated circuits (3D IC) enable the large scale integration of heterogeneous devices into a single system without increasing its area footprint. However, the consequent increase in power density raises the significant challenge of thermal management [1]. In conventional single-die ICs, the heat dissipated by circuit elements is conducted to the heatsink through a relatively small number of intermediate layers. In 3D ICs on the other hand, each die can add up to 12 layers of varying thickness and conductivity to the heat flow path between dissipating elements and the heatsink [2]. This impedes the flow of heat away from power dissipation sites, and results in aggravated operating temperatures. A number of studies in literature examine specific thermal characteristics of die stacks. Oprins et al. [3] investigated the thermal coupling between memory and logic dies in a two-tier stack in order to determine operating temperatures and thermal profiles in the memory. Their study used a thermal model calibrated with a $130nm$ stacked-die silicon test chip, and examined the influence of microbumps and underfill thermal conductivity on peak temperatures. Matsumoto et al. [4] performed an experimental measurement of thermal resistance of multi-layer die stacks interconnected with C4 microbumps. Their work revealed the dependence of thermal resistance on microbump pitch and size. Most recently, Vaisband et al. [5][6] investigated thermal conduction paths within a two-tier die stack in order to evaluate the magnitude of heat transfer between dies, and to determine the temperature dependence of thermal conductivity.

This research was supported in part by the CATRENE programme under the Computing Fabric for High Performance Applications (COBRA) project CA104. The authors would like to thank Yash Joshi of the Delft University of Technology for his help with the thermal conductivity estimation.

The design of the *Through Silicon Via (TSV)* based vertical interconnect significantly influences thermal conductivity, and thereby impacts system performance, as we previously showed in [7]. However, in addition to TSVs, the composition and depth of die stacks, die thickness, location of power dissipating elements, and stack power density also form critical design parameters that influence thermal behaviour. In this paper, we characterize the influence of these parameters on the steady-state temperature profiles of dies in 3D ICs using a high-level flow. The characterization is performed using a physical model of a real 3D stacked test chip [2], and provides insights into the formation of hotspots, and the implications of effective thermal conductivity on the placement of temperature sensors for accurate tracking of hotspots. Furthermore, this paper highlights the complexity of heat flow within 3D ICs, and the impact of stack power density on the operating temperature of dies.

II. Thermal Characterization Flow

Three-dimensional integrated circuits are composed of multiple stacked dies interconnected using TSVs. These enable the high-density integration of devices without increasing the area footprint of the system. However, the number of devices that can be integrated within the system is limited by its thermal behaviour. Nagata [8] determined this limit as:

$$\frac{\alpha N_G E}{t_{pd}} \leq g \cdot \Delta T \tag{1}$$

where N_G is the number of gates that can be integrated within a system with a clock period t_{pd}. Essentially, the relation dictates that the maximum number of gates that can be integrated within a single chip is limited by the average thermal conductance g and temperature gradient ΔT between the dissipating elements and the ambient air. In order to increase integration density, either energy dissipation E of the gates, or their activity rate α must be decreased. Alternatively, the thermal conductance of the system must be improved so as to efficiently dissipate the larger amount of generated heat (Q). Heat flow between dissipating elements and sink surfaces in 3D ICs follows the Fourier heat transfer equation:

$$c_v \frac{\delta T}{\delta t} = \nabla \cdot g(\nabla T)^T + Q \tag{2}$$

$$\nabla T = [\frac{\delta T}{\delta x}, \frac{\delta T}{\delta y}, \frac{\delta T}{\delta z}] \quad and \quad Q \propto \frac{\alpha N_G E}{t_{pd}} \tag{3}$$

978-1-4799-8391-9/15/$31.00 ©2015 IEEE

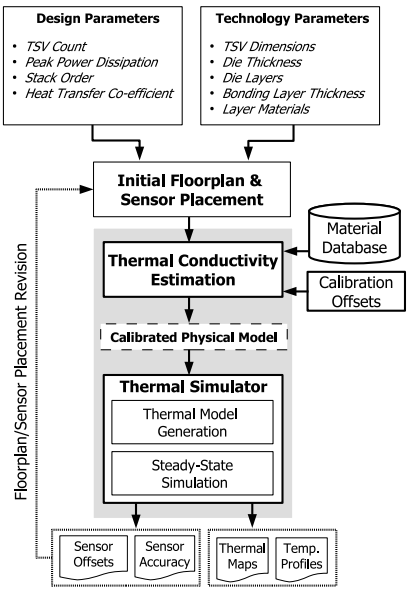

Fig. 1. Thermal characterization flow

where Q is the heat flowing from a power dissipating element towards multiple sink surfaces, at a rate $\delta T/\delta t$ through a material of volumetric heat capacity c_v. The conductance matrix g defines the thermal conductance along the orthogonal x, y and z axes of the material as a function of effective thermal conductivity (κ_{eff}) of the materials encountered along those respective axes. The equation (2) can be rewritten in the one-dimensional steady-state form as:

$$Q = g \cdot \Delta T \quad where \quad g = \kappa_{eff} \frac{A}{l_{x,y,z}} \tag{4}$$

$$\kappa_{eff} = A_{tsv}\kappa_{tsv} + A_{mat}\kappa_{mat} \tag{5}$$

This equation indicates the relationship between thermal conductance, and effective thermal conductivity of the material. The conductance g across a die layer is thus a function of its κ_{eff}, its area (A), and its material thickness ($l_{x,y,z}$). The κ_{eff} of a layer is not only dependent on its material, but also on structures such as TSVs which act as high conductance thermal pathways, improving the overall conductivity of the layer itself. For a material layer with TSVs, the κ_{eff} may be computed as the weighted average of the thermal conductivities of the TSV material (κ_{tsv}) and the layer material (κ_{mat}) respectively, as given in (5). A_{mat} is the area of the material layer in a die with length h_{die} and width w_{die}, and A_{tsv} represents the total area of n TSVs of radius r_{tsv} on this layer.

$$A_{tsv} = n(\pi r_{tsv}^2) \quad and \quad A_{mat} = (h_{die}w_{die}) - A_{tsv} \tag{6}$$

From (4), the magnitude of thermal gradients observed at a given point across any layer of the stack is determined by effective thermal conductivity (κ_{eff}) of that layer, its thickness, and the power density at that point in the die stack. Lateral thermal gradients on stacked dies are predominantly influenced by the effective thermal conductivity of the bulk silicon. As κ_{eff} increases more heat is conducted away from the power dissipation site resulting in a better spread of temperatures. Low κ_{eff} on the other hand results in the stagnation of heat and thus the formation of hotspots. As the thickest die layer, the bulk silicon's κ_{eff} also determines the heat flow to other dies in the stack, thereby affecting vertical thermal gradients.

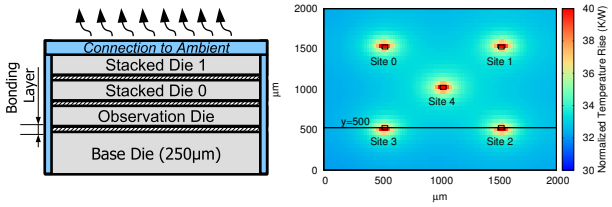

Fig. 2. Composition of the die stack (left) and normalized map of temperature rise across the die with all heaters dissipating 25mW (right)

Improving conductivity increases heat flow towards the sink surfaces, yielding lower peak temperatures in stacked dies.

The spread of temperatures is also dependent on die thickness ($l_{x,y,z}$). The use of die thinning to improve integration density causes a decrease in the thickness of bulk material, which hampers the lateral spreading of heat, and leads to the formation of high temperature hotspots. The spread of temperatures is however limited by lateral thermal resistance beyond a certain die thicknesses. Therefore, dies thinner than this threshold can be expected to experience hotspots of greater magnitude than thicker dies. Power density influences the concentration of generated heat (Q) in die stacks, and can increase the peak temperatures of hotspots. However, since vertical heatflow is determined by κ_{eff}, the temperature gradient of a die is predominantly affected only by local power dissipation as well as dissipation in the tiers directly above and below it. Heat from other tiers in the stack, in the steady-state condition, raises the overall operating temperature of all the dies. Consequently, increasing power density causes the temperature profiles to shift to higher ranges, while temperature gradients are affected only if the dissipated power is in a tier with sufficient κ_{eff} to the observation die.

A high-level exploration flow is used to characterize the influence of these parameters on the thermal behaviour of die stacks. The flow comprises of two stages: thermal conductivity estimation, and thermal simulation, as illustrated in Figure 1. In the first stage, input design and technology parameters are translated into a physical model of the die stack. This describes the dies and their constituent layers, materials and dimensions, floorplans, peak power dissipation values for components, locations of temperature sensors, and the number and dimensions of TSVs used in the design. The thermal conductivity of each layer in the stack is determined using a material database. The resulting physical model contains all the requisite data for accurate thermal modelling of the die stack. In the second stage of the flow, the thermal behaviour of this model is simulated in order to obtain fine-grained steady-state temperature maps corresponding to the peak power dissipation of components. The obtained maps are homogenized into zones based on their temperature difference compared to the hotspot. These zones indicate regions across which hotspot temperature can be measured with similar accuracy. This information can be used to improve calibration of temperature sensors, and in refining their placement in the floorplans.

III. Experimental Results

The calibrated and validated model is used to characterize the influence of the design and technology parameters on thermal behaviour.

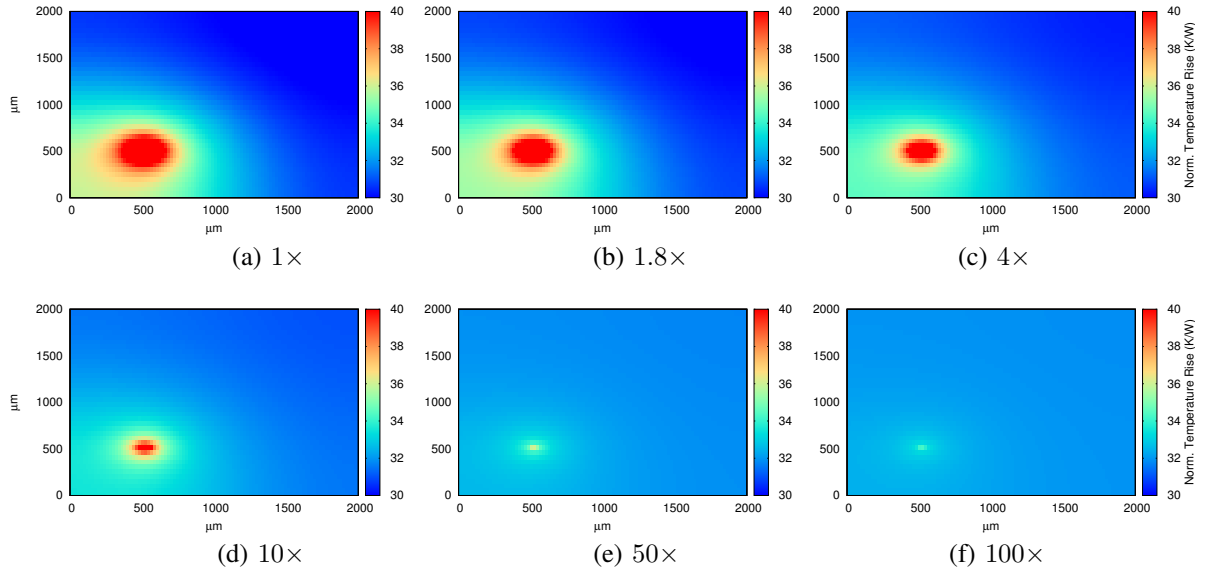

Fig. 3. Thermal map depicting normalized temperature rise across the observation die with increasing effective thermal conductivity relative to conventional silicon (1×).

A. Experimental Setup and Validation

In order to validate our flow, we created a physical model of the 3D stacked test chip described in [2] and carefully calibrated the heat transfer co-efficient of the connection to ambient using available data. Thereafter, we carried out a series of thermal simulations using test floorplans containing heating elements of size and power dissipation identical to those in [2]. Our setup accurately reproduced the same temperature profiles, with a maximum temperature deviation of 6% at hotspot peaks.

The experimental setup utilizes a 3D stack consisting of a $250\mu m$ base die with multiple thinner dies bonded above using $1\mu m$ thick bonding layers, as illustrated in Figure 2. Each die in the stack contains five configurable power dissipating elements of $50\mu m \times 50\mu m$ to generate heat. The complete stack is connected to ambient air through an interface with the previously calibrated heat transfer co-efficient. The material composition and layer dimensions for dies is identical to the 3D test chip described in [2]. The dies used in the characterization have a size of $2000\mu m \times 2000\mu m$, and the thermal simulations use a grid width of $10\mu m$, which provides good accuracy with acceptable simulation times. Simulations are carried out using the 3D-ICE thermal simulator [9] and are run until steady state temperatures are reached. To quantify the variation in temperatures with changing parameter values, the normalized temperature profile along the horizontal line at $y = 500\mu m$ is plotted. For clarity of presentation, we only apply power at Site 3, keeping all other heaters on the die idle. This allows us to clearly observe the temperature gradients that result from power dissipation at a single site under different physical conditions.

B. Experimental Results

1) Thermal Conductivity: The formation of hotspots is critically influenced by the effective thermal conductivity of the bulk silicon, as observed in Figures 3*(a)-(f)* which depict the temperature rise across the dissipating die with increasing

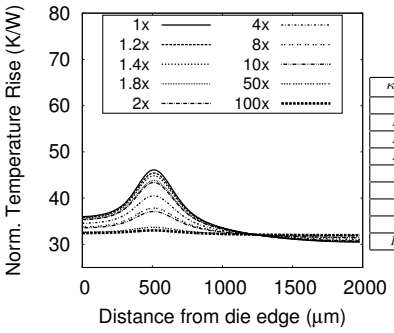

κ_{eff}	ACCURACY-LINKED SENSING ZONE RADIUS (μm)			
	0K	*1K*	*2K*	*3K*
1×	0	80	360	1460
1.2×	0	80	480	1460
1.4×	0	100	620	1460
1.8×	0	120	1140	1460
2×	20	120	1460	1460
4×	20	280	1460	1460
8×	20	1460	1460	1460
10×	20	1460	1460	1460
100×	300	1460	1460	1460

Fig. 4. Influence of effective thermal conductivity on normalized temperature rise in the observation die (left).Variation in the size of sensing zones for different measurement accuracies with changing thermal conductivity (right)

κ_{eff}. This die forms the middle layer of a five-die stack mounted on a $250\mu m$ base, and dissipates $120mW$ of power at the Site 3 heater. As effective conductivity increases, the heat produced due to this localized power dissipation is conducted more efficiently to adjacent dies in the stack, and in turn to the sink surfaces. The generated heat is thus prevented from stagnating in the die, thereby resulting in a decreased hotspot temperature, and smaller thermal gradients. Figure 4 shows the change in temperature profiles observed with varying κ_{eff} values. Note that κ_{eff} is specified relative to the thermal conductivity of pure silicon.

The Figures 3 and 4 also illustrate the improved heat spread due to decreased lateral thermal resistance resulting from the higher κ_{eff}. The lateral spread of heat holds a number of implications for the placement of temperature sensors. The accuracy of temperature sensors is influenced by their distance from the hotspot [10], and owing to their size, sensors cannot always be placed at close proximity to the region of interest. This necessitates the use of calibration techniques to offset induced measurement errors [10][11]. We observe that the radius of the zone within which hotspot temperatures can be tracked with 100% accuracy increases with κ_{eff}. The

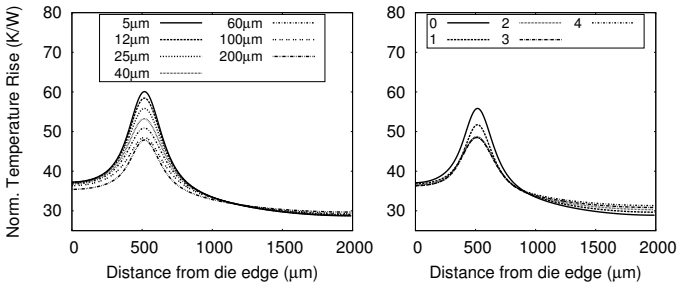

Fig. 5. (a) Influence of varying die thickness on normalized temperature rise. (Thickness of base die: $250\mu m$) (b) Influence of stack depth on normalized temperature rise on an active observation die mounted on a $250\mu m$ thick base die. (Thickness of observation/stacked dies: $25\mu m$)

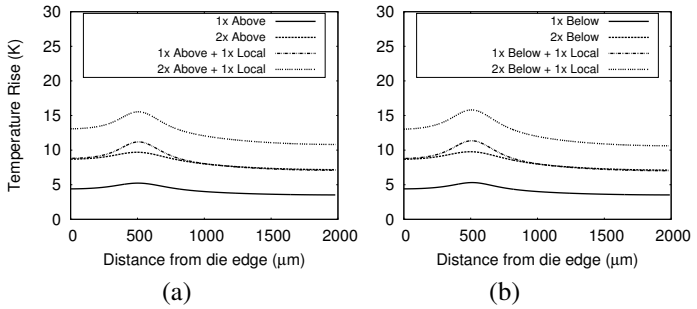

Fig. 6. Influence of stack power density on the actual temperature rise in an observation die. Power sources (a) above observation die (b) below observation die. Note that the case marked *+Local* uses an active Site 3 heater on the observation die. Each heater dissipates $120mW$ of power.

table in Figure 4 lists the maximum radius of sensing zones within which temperatures can be measured to within $0K$, $1K$, $2K$ and $3K$ accuracy of the hotspot. The radius is observed to increase with κ_{eff}, indicating that temperature sensor placement becomes less restrictive as effective thermal conductivity increases.

2) Die Thickness and Stack Depth: Figure 5(a) shows the temperature profile on a single die of varying thickness, stacked above a $250\mu m$ base die. The thickness of the die is observed to influence the peak temperatures of hotspots. The relatively shallow depth of bulk silicon in thin dies inhibits the spread of heat away from the hotspot, and thus causes the highest peak temperatures. While increasing the thickness improves lateral spreading, this benefit diminishes as die thickness exceeds $100\mu m$. The influence of die thickness can also be emulated by a stack of dies. Figure 5(b) reports the temperature profiles resulting from the stacking of multiple $25\mu m$ thick dies on top of the dissipating die used in the previous case. Peak temperatures are seen to decrease as stack depth increases. The additional dies stacked above the dissipating die in this case serve as bulk material and facilitate the lateral spreading of generated heat, as evidenced by the increased temperature rise registered at the eastern edge of the die. The heat spreading effect is observed for stacks with up to two-dies above the dissipating die, and translates to an effective thickness ceiling of $75\mu m$ (excluding the base die), which is lower than the $100\mu m$ ceiling observed in the previous case. This observation is explained by the fact that die stacks contain a range of materials such as *silicon dioxide*, which have a thermal conductivity much lower than that of silicon. Consequently, the thermal characteristics of a die stack with effective thickness of $100\mu m$ are certain to differ from those of a single die of the same thickness.

3) Power Density: Figure 6 shows the temperature profile of an observation die located in the middle of a five-die stack mounted on a $250\mu m$ base die. The graph reports the temperature rise experienced in the observation die due to power dissipated at Site 3 heaters on other tiers of the stack, i.e. changing stack power density. Since the primary sink in the stack is located at the surface of the topmost die, generated heat must flow through all intermediate dies. This is evidenced by the higher temperature rise noted when the dissipating dies are located below the observation die. However, it is interesting to note that similar behaviour is observed even when the dissipating dies are placed above the observation die. This leads us to conclude that in addition to the heat flowing towards the primary sink, significant conduction also occurs towards other sink surfaces in the stack, resulting in complex heat flow patterns. The overall steady-state temperature, on the other hand, is observed to be notably influenced by the stack power density and the heat transfer co-efficient of the sink. It is important thus to take the maximum power density of the stack into consideration when determining the specifications of the heatsink and heatspreader.

IV. Conclusions

In this paper, we evaluated the influence of critical design and technology parameters on the thermal behaviour of three-dimensional die stacks. Using a thermal model based on state-of-the-art measurement data from a 3D test chip, we examined the influence of die thickness, stack power density and effective thermal conductivity on temperature profiles of stacked dies. Importantly, the results of our evaluation uncovered the influence of effective thermal conductivity on the sizing of sensing zones for accurate temperature measurement.

References

[1] K. Puttaswamy and G.H. Loh. Thermal analysis of a 3d die-stacked high-performance microprocessor. In *Proceedings of the ACM Great Lakes Symposium on VLSI*, pages 19–24, 2006.

[2] H. Oprins and et al. Fine grain thermal modeling and experimental validation of 3d-ics. *Microelectronics Journal*, 42(4):572–578, April 2011.

[3] H. Oprins and et al. Numerical and experimental characterization of the thermal behavior of a packaged dram-on-logic stack. In *Proceedings of the Electronic Components and Technology Conference*, pages 1081–1088, May 2012.

[4] K. Matsumoto and Y. Taira. Thermal resistance measurements of interconnections, for the investigation of the thermal resistance of a three-dimensional (3d) chip stack. In *Proceedings of the Semiconductor Thermal Measurement and Management Symposium, 2009*, pages 321–328, March 2009.

[5] B. Vaisband, I. Savidis, and E.G. Friedman. Thermal conduction path analysis in 3-d ics. In *Proceedings of the International Symposium on Circuits and Systems*, pages 594–597, June 2014.

[6] I. Savidis and E. G. Friedman. Thermal coupling in tsv-based 3-d integrated circuits. In *Proceedings of the Workshop on 3D Integration - Design, Automation and Test in Europe*, Mar 2014.

[7] S.S. Kumar, A. Aggarwal, R.S. Jagtap, A. Zjajo, and R. van Leuken. System level methodology for interconnect aware and temperature constrained power management of 3-d mp-socs. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 22(7):1606–1619, July 2014.

[8] M. Nagata. Limitations, innovations, and challenges of circuits and devices into a half micrometer and beyond. *IEEE Journal of Solid-State Circuits*, 27(4):465–472, Apr 1992.

[9] A Sridhar et. al. 3d-ice: Fast compact transient thermal modeling for 3d ics with inter-tier liquid cooling. In *International Conference on Computer-Aided Design*, pages 463–470, Nov 2010.

[10] A. Zjajo, N. van der Meijs, and R. van Leuken. Dynamic thermal estimation methodology for high-performance 3-d mpsoc. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 22(9):1920–1933, Sept 2014.

[11] S.O. Memik, R. Mukherjee, M. Ni, and J. Long. Optimizing thermal sensor allocation for microprocessors. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27(3):516–527, March 2008.