Integration of Thermal Management and Floorplanning Based on Three-Dimensional Layout Representations

Puskar Budhathoki*, Johann Knechtel†, Andreas Henschel* and Ibrahim (Abe) M. Elfadel‡
*Institute Center for Smart Infrastructure (iSmart), Masdar Institute of Science and Technology, Abu Dhabi, UAE
†Institute of Electromechanical and Electronic Design, Dresden University of Technology, Dresden, Germany
‡Institute Center for Microsystems (iMicro), Masdar Institute of Science and Technology, Abu Dhabi, UAE
pbudhathoki@masdar.ac.ae, johann.knechtel@ifte.de, ahenschel@masdar.ac.ae, ielfadel@masdar.ac.ae

***Abstract*—Three-dimensional integrated circuits hold great promise for performance improvement and power savings, by reducing space and wiring. However, the technology comes with new challenges already during early stages such as the development of adequate floorplan representations. Moreover, it is required to deal with increased thermal stress in 3D integrated circuits already during physical design. As opposed to the application of conventional layerwise 2D floorplanning approaches, genuine 3D layout representations naturally integrate vertical relations between layers. We present a design flow that integrates 3D floorplanning with thermal management. The main contribution of our approach is the flexible integration of floorplanners that use 3D layout representations and thermal simulations. Thereby, floorplanning (traditionally) accounts for wirelength and chip area, and the placement of thermal through silicon vias (TSVs) is optimized by subsequent TSV-density adaptation. The impact of thermal TSV insertion strongly increases with the number of chip layers. We show that already 0.5% average TSV density can decrease the maximum chip temperature by 15, 45 and 90 degrees for two, three and four layer configurations of exemplary 3D designs, respectively. The presented approach paves the way for thermal-driven design methods that profit from the advantages of 3D representations for floorplanning techniques.**

I. Introduction

Traditional integrated circuit consists of one active device layer covered with several metal layers. Metal layers are used to realize interconnection between transistors that are integrated in the active layer. In the past several years, continuous scaling of CMOS technology into deep submicron regime causes larger delay, higher power consumption and signal integrity problems. One solution to this problem is the implementation of 3D circuits based on 3D integration technology [1], [2]. 3D integrated circuits are a promising technique to boost the performance by stacking active devices rather than shrinking an existing 2D plane. 3D IC technology greatly reduces wire length [3], has the potential to reduce interconnect delay, and can improve system performance. The reduction of wire length also helps to lessen the power consumption [4]. Furthermore, 3D IC technology helps to integrate disparate technologies onto different layers and supports heterogeneous system-on-chips as well. Stacking of the active functional blocks and reduced die footprint causes on chip temperature to increase which in turn degrades the chip reliability, as given by the Arrhenius equation:

$$k = A * exp^{(\frac{-E_a}{R*T})}$$

where k is the rate coefficient, A is a constant, E_a is the activation energy, R is the universal gas constant, and T is the temperature (in Kelvin). Increased power density has made thermal management a critical issue in 3D IC design [5].
During floorplanning, data structures are abstract models–often topology-preserving maps–to store information of layout elements and their properties, such as connectivity to neighbors. The solution space should be non-redundant, as small as possible and include the best solutions. It is desirable to include vertical dependencies between functional modules. We therefore focus on genuine 3D data structures, which perform best on comparable benchmarks with respect to the above listed criteria [6].

II. Problem Formulation

The input to the 3D floorplanning algorithm are set of n blocks, connectivity between blocks specified by a netlist and a maximum number of layers. Each block B_i can be represented by a tuple (h_i, w_i) for a given set of blocks B = $\{B_1, B_2, B_3, ..., B_n\}$. h_i and w_i are the height and the width of respective block B_i. Moreover, the number of maximal employed layers is represented as *k*. A 3D floorplan is a lower-left corner coordinate (x_i, y_i, l_i) assignment for each block B_i such that no blocks overlap and which satisfies $1 \leq l_i \leq k, l_i \in N$. We formulate the 3D floorplanning task with TSVs as a weighted multi-objective optimization problem, minimizing the occupied area of the floorplan (A), total wirelength (w) and number of signal TSVs (n_{vias}):

$$\alpha \cdot A + \beta \cdot w + \gamma \cdot n_{vias} \quad (1)$$

where α, β and γ represent weights that assign different priorities to the three objectives. Wirelength is calculated using the traditional Manhattan distance for every connection specified in the netlist. The total number of inserted signal TSVs needs to be minimized due to limited routing resources and fabrication cost. We consider Face-to-Back die-die bonding. The thermal TSVs pass only through the passive, not the active silicon layer; however, the difference in thermal impact is expected to be quite low. We therefore neglect TSV space requirements in active layers in the current stage.

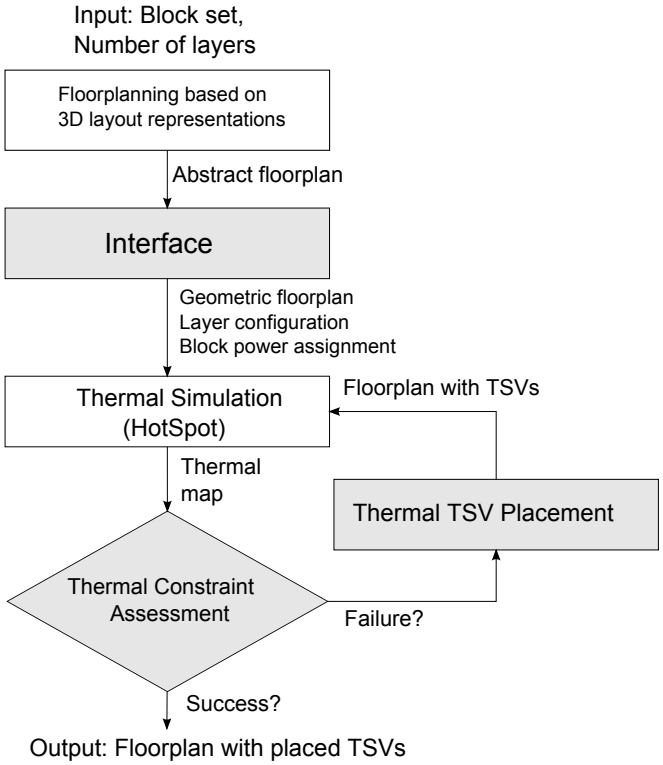

Fig. 1: Workflow diagram for integration of floorplanning and thermal management

III. Approach

A. Overview of the Algorithm

The workflow is shown in Fig. 1. Here we present an area/wirelength-driven floorplanning algorithm based on 3D layout representations with subsequent TSV placement.

The main contributions of our work are also highlighted in Fig. 1. Subsequently, we describe the individual modules employed by our design flow and motivate their choice over alternatives.

B. Floorplanning

As argued initially, 3D layout representations provide a number of advantages over layer-wise 2D floorplanning methods. In particular, we investigate Corner Block Lists (CBLs) and its 3D extension [7], [8], which we deemed suitable for our purposes for several reasons: the runtime complexity of 3D Corner block lists ($O(n)$) is favorable in comparison to other 3D layout representations as shown by [6]. Moreover, the solution space ($O(n!3^{n-1}2^{4n-4})$) is comparably small and CBLs can represent non-slicing structures (such as "wheels").

As floorplanning module, we employ Corblivar [9], a very recent CBL-based tool. Corblivar is tailored for planning of massively-parallel interconnect structures in 3D ICs. Besides this, Corblivar accounts for optimization of area, wirelength, and thermal distribution. Corblivar is further able to restrict the number of dies as well as their outlines, which is often required for manufacturing purposes.

C. Floorplanning/Thermal Management Interface

This module converts topological maps into Cartesian coordinates and generates an output format readable for the thermal simulation module. Layer configurations containing material characteristics and TSV densities are generated.

D. Thermal Simulation

HotSpot is a thermal modeling tool which provides temperature estimation of a functional module by employing the principle of thermal-electrical duality [10]. HotSpot has been successfully used for many floorplanning tasks, predominantly in the realm of 2D thermal management. A circuit-solving technique is applied for an RC network of thermal capacitances and resistances (derived from functional modules and their power consumption). Then, temperatures at the center of functional modules are calculated by solving the RC network's equation. Vertical layers along with their physical properties are provided in the layer configuration. Recently, various extensions enable multi-layer simulations, assuming that floorplans for each layer are calculated and inter-layer material properties are provided. The HotSpot extension tool presented in [11] allows modeling of TSVs with related thermal resistivity and specific heat capacity values in specified (delimited) regions. Accurate thermal simulation for larger scale benchmarks and realistic applications are computationally expensive. By modifying grid size during simulation, we trade off accuracy with speed.

E. Thermal Constraint Assessment and TSV Placement

The layered heat maps from the simulation output are parsed and scrutinized for temperature constraint violations. In the case of violation, we iteratively increase thermal TSV density. This process is repeated until the maximal temperature is below a certain threshold. Note that currently TSV density is increased globally. A more differentiated approach, similar to the TSV placement algorithm presented in [12], is currently in progress. During every thermal-simulation and TSV-placement cycle, the conductivity k_{new} is updated for each grid cell using the formula provided by [12]:

$$k_{new} = k_{old}\frac{t_{curr}}{t_{target}} \tag{2}$$

where k_{old} is the previous conductivity of the grid cell, t_{curr} is the current temperature and t_{target} is the target temperature of the whole package chip. Subsequently, the required TSV density in a grid cell is calculated with the formula from [12]

$$v = min\left(v_{max}, c \cdot \frac{k_{new} - k_{old}}{k_{via} - k_{old}}\right) \tag{3}$$

where v_{max} is the maximum allowable TSV density in each grid cell, c is a user specified constant and k_{via} is the thermal conductivity of TSVs. We then calculate the compound thermal resistivity which is based on the parallel joint resistance of silicon and copper. Here, the TSV density impacts the area fraction for joint-resistance calculation. The HotSpot extension tool assumes that specific heat capacity has lower impact in comparison to thermal resistivity and does not consider its effect for temperature estimation. However, compound specific heat capacity based on mass weighted mean as mentioned by Tonpheng et al. [13] is considered during TSV placement in our model.

IV. Experimental Results

The algorithm for our design flow that integrates floorplanning with thermal management and thermal TSVs placement is implemented in C++. Our experiments are performed on the GSRC benchmark set [14]. The specification of the chip package along with material properties and dimensions [15], [16] are listed in Fig. 2.

	Area (mm^2)	Thickness (μm)	Thermal Conductivity (W/m-K)	Specific Heat Capacity (J/kg-K)
Active SI layer	5X5-15X15	2	117.5	700
Passive Si layer	5X5-15X15	98	117.5	700
BEOL	5X5-15X15	12	2.25	517
Bonding layer	5X5-15X15	20	0.2	2187
Heat Spreader	30X30	1000	400	397
Heat Sink	60X60	6900	400	397

Fig. 2: Dimension and material properties of package model

Fig. 3 shows the thermal profile of the critical die (i.e., the die furthest away from the heatsink) for the n300 benchmark when blocks are stacked in two dies. Note that the thermal profile without and with thermal TSVs are similar since thermal TSVs are distributed uniformly in all layers. However, the comparison also clearly shows the maximum temperature reduction by 21.54K, since thermal TSVs efficiently transfer heat between different layers.

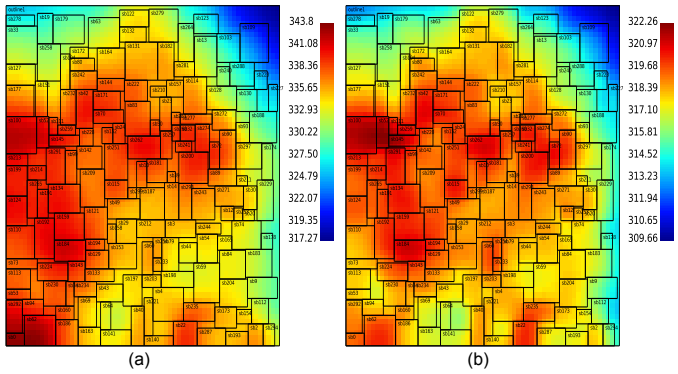

Fig. 3: (a) Thermal profile of critical die without TSVs (b) Thermal profile of critical die with TSV density of 0.01 for n300 benchmark

Table I shows the comparative study of wirelength, chip area and maximum temperature for different GSRC benchmarks between our approach and area/wirelength driven floorplanning (AWF) [12]. The reduction of chip area and compact packing density in our case causes the maximum temperature of the block to increase without thermal TSVs. The area/wirelength driven floorplanning in [12] did not mention 3D bonding technology and TSV-related implications on wirelength, thus is too optimistic. In our case, wirelength estimation exhibits notably larger numbers, this is due to the more realistic model of [9] with (i) the consideration of actual TSV lengths for each net, (ii) the consideration of TSV landing pads and related, additional wiring, as well as (iii) the wires to connect to package pins. On the other hand, our approach shows promising results in temperature reduction with the same average thermal TSV density.

TABLE I: Comparison between our approach and [12] on temperature without TSVs and with TSVs on GSRC benchmarks

Metric	Our Approach, 4 dies			Approach of [12], 4 dies		
	n100	n200	n300	n100	n200	n300
Chip Area(μm^2 x 10^5)	0.53	0.52	0.81	0.61	0.52	0.83
Wirelength(μm x 10^5)	5.13	7.97	10.89	1.38	2.18	3.16
Temp w/o TSVs(^{o}C)	161.57	163.95	174.97	123.6	135.6	186.9
Temp w/ TSVs(^{o}C)	35.01	36.22	40.08	92.3	105.2	146.9
Average TSV density	0.046	0.036	0.026	0.046	0.036	0.026

Fig. 4 shows the relationship between maximal temperature of the critical die and the average density of thermal TSVs used for the GSRC n100 benchmark. It shows that without including thermal TSVs, the temperature of the critical die grows rapidly as the number of stacked dies increases in 3D ICs. We can conclude from Fig. 4 that even with 0.5% of average thermal TSV density for an n100-benchmark 4-die setup, we can lower the temperature of the critical die by 90K while the temperature reduction in 3-die and 2-die setups are 45K and 15K, respectively. However, the reduction of maximal temperatures on a critical die for applying more than 1% average thermal TSV density decreases strongly.

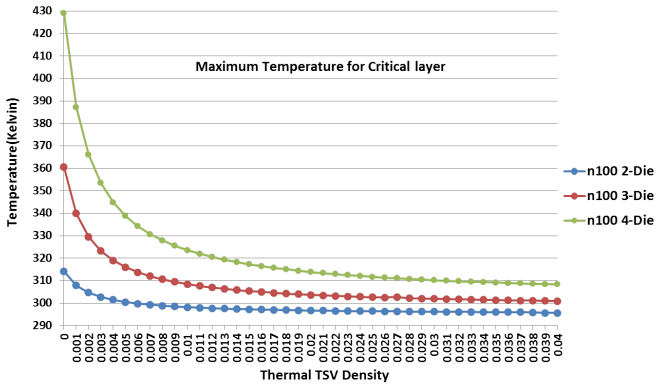

Fig. 4: Characteristic curve for Thermal TSV density vs. Maximum Temperature for the critical die of n100 benchmark

V. SUMMARY

We conceptualized the integration of 3D layout representation based floorplanning with thermal management. Floorplanning results were subjected to a loop of thermal simulations and TSV placements. This way, we first determine the necessary TSV density to avoid thermal constraint violations and second, we also unravel the thermal characteristics of the layout. Finally, we showed that the impact of thermal TSVs insertions strongly increases with number of layers.

VI. Discussion and Future Work

The presented work follows the approach conceived by Wong et al. [12], with the crucial extension that our design flow is modular in 3D layout representations rather than layered 2D sequence pairs. Moreover, we implemented a workflow that utilizes a full thermal simulation cycle. The presented work enables tight integration between floorplanning and thermal management for three-dimensional circuits. At present, TSV density is changed uniformly across the whole layer. The next obvious step is to differentiate and place thermal TSVs only around hotspots, thus minimizing number of TSVs and hence, manufacturing costs.

We have implemented thermal TSV placement as a post processing step to floorplanning. In order to accommodate fully thermal-driven design, i.e. to simultaneously minimize area, wirelength and maximal chip temperature, information from thermal simulations must be fed back to the floorplan optimization loop. With 3D layout representations, it is possible to create heuristics that can (conservatively) estimate whether a proposed floorplan perturbation (potentially) causes a thermal constraint violation. This way, full thermal simulations would be omitted whenever possible, e.g., on exchange or rotation of comparably cold blocks. Besides, note that the employed modules for thermal simulation are fast approximations and hence—depending on grid size granularity—can be suited to the computational challenge of integrating thermal simulations into the optimization loop of the floorplanner.

Another possible extension to our algorithm is the consideration of spatial constraints introduced by thermal TSV insertion, as for example implemented in [17].

Acknowledgment

The authors would like to thank Advanced Technology Investment Company (ATIC, Abu Dhabi, United Arab Emirates) for funding support of the TwinLab collaboration between Masdar Institute, UAE and TU Dresden, Germany, Ref. 372/002/6754/102d/146/64947.

References

[1] E. Beyne, "The rise of the 3rd dimension for system intergration," in *Interconnect Technology Conference, 2006 International*. IEEE, 2006, pp. 1–5.

[2] W. R. Davis, J. Wilson, S. Mick, J. Xu, H. Hua, C. Mineo, A. M. Sule, M. Steer, and P. D. Franzon, "Demystifying 3D ICs: the pros and cons of going vertical," *Design & Test of Computers, IEEE*, vol. 22, no. 6, pp. 498–510, 2005.

[3] J. W. Joyner, R. Venkatesan, P. Zarkesh-Ha, J. A. Davis, and J. D. Meindl, "Impact of three-dimensional architectures on interconnects in gigascale integration," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 9, no. 6, pp. 922–928, 2001.

[4] J. W. Joyner and J. D. Meindl, "Opportunities for reduced power dissipation using three-dimensional integration," in *Interconnect Technology Conference, 2002. Proceedings of the IEEE 2002 International*. IEEE, 2002, pp. 148–150.

[5] S. Das, A. Chandrakasan, and R. Reif, "Timing, energy, and thermal performance of three-dimensional integrated circuits," in *Great Lakes Symposium on VLSI: Proceedings of the 14 th ACM Great Lakes symposium on VLSI*, vol. 26, no. 28, 2004, pp. 338–343.

[6] R. Fischbach, J. Lienig, and J. Knechtel, "Investigating modern layout representations for improved 3d design automation," in *Proceedings of the 21st edition of the great lakes symposium on Great lakes symposium on VLSI*, ser. GLSVLSI '11. New York, NY, USA: ACM, 2011, pp. 337–342. [Online]. Available: http://doi.acm.org/10.1145/1973009.1973076

[7] X. Hong, G. Huang, Y. Cai, J. Gu, S. Dong, C.-K. Cheng, and J. Gu, "Corner block list: an effective and efficient topological representation of non-slicing floorplan," in *Computer Aided Design, 2000. ICCAD-2000. IEEE/ACM International Conference on*. IEEE, 2000, pp. 8–12.

[8] Y. Ma, X. Hong, S. Dong, and C. Cheng, "3D CBL: An efficient algorithm for general 3D packing problems," in *Circuits and Systems, 2005. 48th Midwest Symposium on*. IEEE, 2005, pp. 1079–1082.

[9] J. Knechtel, E. F. Y. Young, and J. Lienig, "Structural Planning of 3D-IC Interconnects by Block Alignment," in *ASP DAC*, 2014, accepted for publication.

[10] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. R. Stan, "HotSpot: A compact thermal modeling methodology for early-stage VLSI design," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 14, no. 5, pp. 501–513, 2006.

[11] J. Meng, K. Kawakami, and A. K. Coskun, "Optimizing energy efficiency of 3-D multicore systems with stacked DRAM under power and thermal constraints," in *Proceedings of the 49th Annual Design Automation Conference*, ser. DAC '12. New York, NY, USA: ACM, 2012, pp. 648–655. [Online]. Available: http://doi.acm.org/10.1145/2228360.2228477

[12] E. Wong and S. K. Lim, "3D floorplanning with thermal vias," in *Design, Automation and Test in Europe, 2006. DATE'06. Proceedings*. IEEE, 2006, pp. 1–6.

[13] B. Tonpheng, J. Yu, and O. Andersson, "Thermal conductivity, heat capacity, and cross-linking of polyisoprene/single-wall carbon nanotube composites under high pressure," *Macromolecules*, vol. 42, no. 23, pp. 9295–9301, 2009. [Online]. Available: http://pubs.acs.org/doi/abs/10.1021/ma902122u

[14] "GSRC benchmarks," 2000. [Online]. Available: http://vlsicad.eecs.umich.deu/BK/GSRCbench/

[15] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunschwiler, and D. Atienza, "3D-ICE: fast compact transient thermal modeling for 3D ICs with inter-tier liquid cooling," in *Proceedings of the International Conference on Computer-Aided Design*, ser. ICCAD '10. Piscataway, NJ, USA: IEEE Press, 2010, pp. 463–470. [Online]. Available: http://dl.acm.org/citation.cfm?id=2133429.2133527

[16] J.-H. Park, A. Shakouri, and S.-M. Kang, "Fast thermal analysis of vertically integrated circuits (3d ics) using power blurring method," in *InterPACK*, vol. 9, 2009, pp. 19–23.

[17] J. Knechtel, I. L. Markov, J. Lienig, and M. Thiele, "Multiobjective optimization of deadspace, a critical resource for 3D-IC integration," in *Proceedings of the International Conference on Computer-Aided Design*. ACM, 2012, pp. 705–712.