A 3D Physical Design Flow Based on OpenAccess

Jason Cong[1,2], *Fellow, IEEE*, and Guojie Luo[1], *Student Member, IEEE*
[1]Computer Science Department, University of California, Los Angeles, CA 90095, USA
[2]California NanoSystems Institute, Los Angeles, CA 90095, USA
{cong, gluo}@cs.ucla.edu

***Abstract*—3D IC technologies have recently attracted great attention due to the potential performance improvement, power consumption reduction and heterogeneous integration. In this paper we present a 3D physical design flow based on OpenAccess (named 3D-Craft) to facilitate the rapid adoption of 3D IC technologies. The OpenAccess extension for 3D-Craft is discussed, and the key components including the 3D placer mPL-3D and the 3D router TMARS are presented. We also demonstrate the application of 3D-Craft for the 3D physical design of an open-source processor, and show that the 3D implementation can reduce both the half-perimeter wirelength and the routed wirelength by about 30% compared to the 2D implementation.**

I. INTRODUCTION

The 3D IC technologies promise to further increase integration density, beyond Moore's Law, and offer the potential to significantly reduce interconnect delays and improve system performance. Furthermore, the shortened wirelength, especially that of the clock net, also lessens the power consumption of the circuit. 3D IC technologies also provide a flexible way to carry out the heterogeneous system-on-chip (SoC) design by integrating disparate technologies, such as memory and logic circuits, radio frequency (RF) and mixed signal components, optoelectronic devices, etc., onto different layers of a 3D IC.

Device layers in a 3D IC are connected using through-silicon vias (TS via). However, TS vias are usually etched or drilled through device layers by special techniques and are costly to fabricate. A large number of the TS vias will degrade the yield of the final chip. Also, under current technologies, TS via pitches, usually around 5-10μm, are very large compared to the sizes of regular metal wires. In 3D IC structures, TS vias are usually placed at the whitespace between the macro blocks or cells, so the TS vias affect both the routing resources and the overall chip areas. Therefore, the number and distribution of TS vias in a 3D IC needs to be considered, not only at the routing stage, but also at the floorplanning and placement stages.

Another critical challenge of 3D IC design is heat dissipation, which already posed a serious problem even for 2D IC designs. The thermal problem is exacerbated in the 3D cases mainly for two reasons: (1) The vertically stacked multiple layers of active devices cause a rapid increase in power density; (2) The thermal conductivity of the dielectric layers between the device layers is very low compared to silicon and metal. For instance, the thermal conductivity of SiO_2 at room temperature (300K) is 1.4 W/mK [12], which is very much smaller than the thermal conductivity of silicon (150 W/mK) and copper (401 W/mK). Therefore, the thermal issue needs to be considered during every stage of 3D physical design.

In recent years, 3D IC physical design has attracted an increasing amount of attention. There is a significant amount of work on the floorplanning [1][2][14][15][19][21], placement [8][10][11][5][7][18] and routing [3][4][20] for 3D ICs. However, all these tools developed by different groups, using different formats to represent the design data, create barriers for researchers who need to make use of the existing design automation tools to conduct further studies on 3D IC. This problem motivates us to develop an infrastructure for the 3D design data representation and assist in the interoperation of physical design tools.

In this paper we present an OpenAccess [22] extension for 3D physical design automation. The main difficulty in developing the OpenAccess extension is to make it applicable, not only to a specific 3D IC technology available today, but also applicable to other possible 3D IC technologies in the future. The issues in detail include how to represent the multiple-tier structure for a 3D IC technology, and how to represent through-silicon vias (TS via) in a 3D design. To solve these issues, we have made the following contributions in this paper:

- We define a database architecture for 3D physical design based on OpenAccess. This architecture is capable of representing the multi-tier structure in a general way, and it is also capable of representing the structure of TS vias and their occurrence in a 3D design.
- We implement 3D-Craft based on the OpenAccess extension. This is a 3D physical design flow including 3D placement, 3D routing and thermal TS via insertion, and the interface with commercial detailed routers.
- The ability of 3D-Craft is demonstrated through the 3D physical design of an open-source microprocessor LEON3 [23]. The physical design results show that the 3-tier 3D implementation can reduce both the half-perimeter wirelength and routed wirelength by about 30% compared to the 2D implementation.

The remainder of this paper is organized as follows. Section II illustrates an overview of 3D-Craft, introduces OpenAccess and discusses the issues and solutions in the extension for 3D ICs. Section III presents the key components in 3D-Craft. The

Manuscript received May 19, 2009. This research was partially supported by IBM under a DARPA subcontract, and supported by National Science Foundation under grants CCF-0430077 and CCF-0528583.

application of 3D-Craft on an open-source core is demonstrated in Section IV. Finally, Section V concludes the work and proposes future work.

II. Overview of 3D-Craft and OpenAccess Extension for 3D Physical Design

The 3D-Craft's physical design flow is illustrated in Fig. 1. Several physical design tools are integrated based on the 3D OpenAccess, including a 3D placer, a 3D global router with thermal TS via planner, and a commercial 2D detailed router. A thermal resistive network model is also integrated for thermal evaluations. The detailed information for this collection of physical design tools will be presented in Section III. Before we present the components in 3D-Craft, we shall first introduce OpenAccess and discuss the issues and solutions in the extension for 3D physical design.

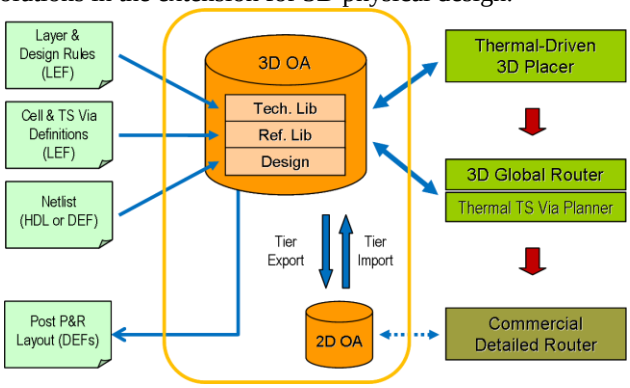

Fig. 1. Overview of 3D-Craft

A. Preliminaries on OpenAccess

OpenAccess is an infrastructure designed and maintained by Si2 [22] to achieve the interoperability of EDA applications and design data. The design data management is through a C++ API that defines classes and member functions to create, access and manipulate the databases. The OpenAccess API consists of a set of packages, where the technology database package and the design database package are the most important packages related to place and route (P&R) tools.

The technology library holds data that is generally applied across all the designs developed from a specific technology. For example, the minimum wire width of a metal layer is specified in the technology library as a layer constraint.

The design database refers to design-specific data, which includes cell/macro libraries and the top design that consists of instances of cells and macros. The data in the design library is represented in a hierarchical way, an example of which is illustrated in Fig. 2. In this example, the design library contains the cells XOR/AND/OR, the "macro" HalfAdder, and the "top design" FullAdder. The top design consists of instances of designs in the cell library and the macro library.

The *oaAppDef* mechanism is a way to add extension values to existing database objects, through which we attach additional information to OpenAccess objects to realize the 3D design representation. An example is shown in Fig. 3, where the attributes "name," "cell name," and "origin" are native attributes of the *oaInst* objects. And we use *oaIntAppDef* to add an integer attribute to these objects for the extended information in 3D designs. So every instantiation of *oaInst* includes the native attributes of "name," "cell name," and "origin," as well as the extended attribute of "tier." The detailed use of *oaAppDef* for OpenAccess extension will be discussed in Section II.C.

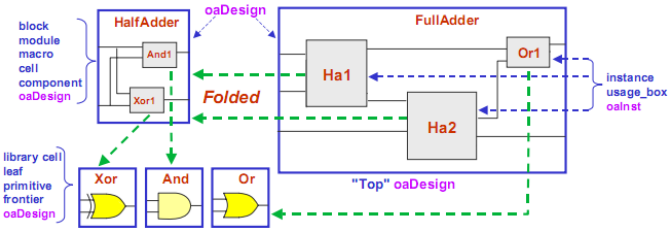

Fig. 2. An example of the hierarchical design database [22]

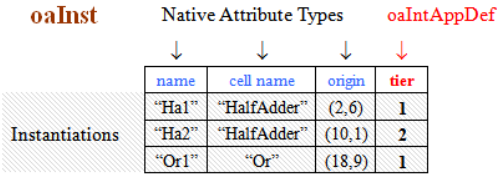

Fig. 3. An example of the *oaAppDef* mechanism

B. Issues in Extending OpenAccess for 3D Physical Design

Although OpenAccess is general enough for data representation of traditional 2D designs, it lacks some necessary features for 3D IC technologies; therefore, there are two very important issues to be considered in the extension.

The first issue is the multiple-tier structure in a 3D design. Physically, a 3D design can be viewed as a stack of multiple 2D designs, where a single 2D design is called a *tier* in terms of 3D IC technologies. There are face-to-back, face-to-face and back-to-back bonding of tiers to form a 3D IC. Therefore, the OpenAccess extension should be able to represent the multiple-tier structure and also the bonding method.

The other issue is representation of through-silicon vias (TS via). TS vias are used for the interconnection between different tiers. In current 3D IC technologies, the size of a TS via is comparable to the size of an inverter, and it also has electrical and mechanical characteristics that are different from traditional vias; thus, TS vias should be represented explicitly and acknowledged by physical design tools.

C. Solutions for Extending OpenAccess for 3D Designs

The database architecture for 3D designs consists of a technology library, a reference library and a design library.

The technology library stores the technology information of each tier, where each tier has a structure similar to a single 2D design. The design rules for the metal layers and the traditional via layers are described in the technology library, as well as the RC characteristics of these layers. To represent all this information, the metal layers and normal via layers are listed in order from bottom to top. A string called *structure* is added to the *oaTech* object by the *oaAppDef* mechanism to

describe the bonding order of tiers and the order of the metal layers and normal via layers. For example, the string "f5b5b5" tells us that there are three tiers, "f5," "b5," and "b5," in a 3D chip. The first tier in a face-on-top ("f") direction starts with a silicon layer, followed by five ("5") interconnect layers from metal 1, via 12, metal 2, via 23 to metal 3. The second and third tiers in a back-on-top ("b") direction start with five ("5") interconnect layers first, from metal 3, via 23, metal 2, via 12 to metal 1, followed by the silicon layer. In such a way, the multi-tier structure is represented in the OpenAccess database, which is the solution of the first issue referred to Section II.B.

The reference library contains a TS via library and a cell/macro library. Under the assumption that the size of a TS via is no larger than the standard cell height, we capsulate the TS via as a pseudo cell in the TS via library. The metal layer consumptions, as well as the silicon layer consumption which a TS via drills through, are captured by the pseudo cell. Different TS vias connecting different pairs of tiers may have different structures; thus there are multiple pseudo cells in the TS via library. This solves the second issue referred to in Section II.B. The cell/macro library contains the abstract (bounding box and I/O pins) of standard cells and intellectual property (IP) blocks that are available for a design instance.

The design library stores the netlist and its P&R information. The original netlist representation before placement and routing is the same as the netlist representation for 2D designs. After placement, an integer called *tier* is added to every *oaInst*, as shown in Fig. 3, to represent which tier a cell instance is placed on. After global routing, the TS via locations are determined, and instances of the pseudo cells in the TS via library are created to represent these TS vias. When the TS via locations are determined, it is straightforward to partition the 3D design to multiple 2D designs physically and run commercial detailed routers.

III. Components in 3D-Craft

A. Resistive Thermal Model

In this work we use the thermal resistive model proposed in [17]. Compared with an accurate simulation tool, the error of this resistive network model is smaller than 2% [17]. A tile structure is imposed on the circuit stack, as shown in Fig. 4(a). Each tile stack contains an array of tiles, one from each device layer, as shown in Fig. 4(b). A tile stack is modeled as a resistive chain as shown in Fig. 4(c).

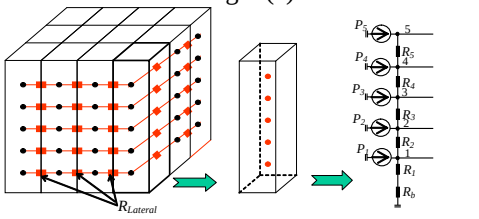

(a) Tiles Stack Array (b) Single Tile (c) Tile Stack Analys

Fig. 4. Resistive Thermal Model

The tile stacks are connected by lateral resistances. A voltage source is used for the isothermal base of heat sink temperature, and current sources are present in every tile to represent heat sources.

B. Multilevel Placement for 3D ICs (mPL-3D)

Given a circuit represented as a hypergraph $H=(V,E)$, the placement region R, and the number device layers K, the task of 3D placement problem is to assign every cell $v_i \in V$ a triple (x_i, y_i, z_i), which indicates that this cell is placed on the device layer $z_i \in \{1,2,\mathrm{K},K\}$ with its center at $(x_i, y_i) \in R$. The objective is to minimize the weighted sum of total wirelength and TS via number, under the non-overlap constraints and the TSV density constraints in each tier.

$$\begin{aligned} \text{minimize} \quad & OBJ(x,y,z) = \sum\nolimits_{e\in E}\big(l(e)+\alpha\cdot v(e)\big) \\ \text{subject to} \quad & \big(\text{non-overlap constraints}\big) \\ & \text{(TSV density constriants)} \end{aligned} \tag{1}$$

The objective function is a weighted sum of wirelength and TS via number. We estimate the wirelength $l(e)$ by the half-perimeter model, and estimate the TS via number $v(e)$ by assuming the cross-tier net is routed as a vertical-trunk tree. The weighting factor α is used to achieve the tradeoffs between the wirelength $l(e)$ and the TS via number $v(e)$.

The multilevel analytical placement engine [7] is used for 3D global placement, where it solves the following problem in a multilevel framework:

$$\begin{cases} \text{minimize } OBJ(x,y,z)+\mu\cdot Penalty(x,y,z) \\ \text{increase } \mu \text{ until converge} \end{cases} \tag{2}$$

To apply an analytical solver, the range of the discrete variable z is relaxed from the set $\{1,2,...,K-1\}$ to a continuous interval $[1,K]$. Then the max operation in the function $OBJ(x,y,z)$ is smoothed by the log-sum-exp function [13], and the non-overlap constraints are replaced by the projected smoothed density constraints on tiers and pseudo tiers [7] and added to the objective as a penalty function.

The intermediate solution obtained from global placement roughly satisfies the density constraints, thus the tier assignment can be done by simply snapping the cells or the macros to the nearest tier. After tier assignment, the remaining 2D legalization and detailed placement are done tier-by-tier as traditional 2D detailed placement. A two-step approach [6] is used for the mixed-size legalization, where the macros are legalized first, followed by standard-cell legalization.

The experimental results in [7] indicate that a 3D standard cell placement with four tiers can reduce the wirelength by about 50% compared to a 2D placement, if the sizes of TS vias are ignored. The results also show that the analytical placement engine is able to achieve good-quality tradeoffs between wirelength and TS via number, so it is adaptive for different 3D IC technologies.

C. 3D Multilevel Routing with Thermal TS Via Planning (TMARS)

The thermal-driven 3D routing with a TS via planning problem can be described as follows, given these three inputs:

1) The target 3D IC technology, including design rule, height and thermal conductivity of each material layer.

2) A 3D circuit placement or floorplan result with whitespace reserved between blocks for interlayer interconnects.

3) A given maximum temperature T_0, e.g., 80°C.

Given these three inputs, the circuit must be routed according to the connecting rules and design rules, so that the weighted cost of wirelength and the total TS via number is minimized.

The size and thermal conductivity disparity between TS vias and the regular signal wires and vias make it difficult to handle them together. An individual step of TS via planning also gives us more control over the temperature of the circuit, since the TS vias are planned directly instead of planned through shortest-path searching. Therefore, our TMARS [3][4] is implemented as a multilevel 3D routing system with a novel thermal TS via planning algorithm. With a more global view and the planning power of a multilevel planning scheme, the TS via planning step can effectively optimize temperature and wirelength through direct planning of the TS vias.

The experimental results [3][4] show that this method of simultaneous global routing and thermal TS via planning can reduce the number of TS vias by 67% with the same temperature constraint, compared to the method of thermal TS via insertion by post-processing.

D. 3D Detailed Routing

After the 3D global routing with signal and thermal TS via planning, the 3D design can be decomposed into several 2D designs, and there is no difference between 3D detailed routing and 2D detailed routing. Therefore, it is best to make use of a well-developed commercial detailed router.

The decomposed 2D designs can be routed by any router, including the commercial routers in the Cadence Encounter [24] or the Magma Talus [25]. The routers read in the decomposed circuit and complete the global and detailed routing.

Magma Talus supports reading in existing global routing using Magma-TCL script. So, we implemented a tool to exchange the routing data with Magma Talus. The global routing paths are represented as boxes in the Magma data model. Thus the result of 3D global routing is imported to Talus by creating boxes using Magma-TCL script. The tool we implemented in the flow converts the global routing into a Magma-TCL script to create those boxes. After loading the script and running the detailed router on these 2D designs, the tool again parses the detailed routing information and writes back to the 3D design library. The routed 3D design is obtained after this step.

IV. 3D Physical Design Evaluation

We evaluate the benefit on wirelength reduction using 3D IC technology and the capability of our 3D-Craft tool with the open-source processor LEON3 [23]. It is a synthesizable and configurable 32-bit processor compliant with the SPARC V8 architecture. We synthesize a single-core LEON3 processor based on the NCSU standard cell kit [9] for the MITLL 0.18μm FD-SOI technology. The size of a TS via is 3×3 μm^2 with an extra 1.5 μm spacing on each side, thus a TS via consumes a area of 36 μm^2. We perform physical design and evaluate the wirelengths for the 2D and the 3-tier 3D implementations.

The synthesized netlist consists of 95061 standard cells, 97880 nets, and 150 I/O ports. The total cell area is 11.05 mm^2, We define the 2D and 3D placement regions such that there are 20% white spaces in total, excluding the area consumed by TS vias. In detail, we define a 3.72×3.72 mm^2 placement region for the 2D implementation, and a 2.15×2.15 mm^2 placement region for the 3-tier 3D implementation. Since 2D physical design can be treated as a special case of 3D physical design, we use 3D-Craft to generate a placement and global routing for both implementations. The placement is generated by mPL-3D, and the global routing is generated by Cadence Encounter.

We list the statistics of the physical design results in TABLE I. The half-perimeter wirelength (HPWL) and the wirelength estimated after global routing (routed WL) are computed by Cadence Encounter. The TS vias are inserted by mPL-3D assuming there is only one TS via for every inter-tier net, and for each tier we only count the TS vias (#TSV) that consumes silicon area. The area utilization is also reported in the last column. In this example, the HPWL of the 3-tier 3D implementation is 29% shorter than the 2D implementation, and the routed wirelength is 37% shorter, which demonstrates the benefit of the 3D implementation.

TABLE I Statistics of the 2D and 3D implementations

		HPWL (mm)	routed WL (mm)	#TSV	utilization
2D		11.23	17.11	N/A	0.80
3D	Bottom tier	2.26	3.33	0	0.80
	Middle tier	3.57	4.5	729	0.81
	Top tier	2.2	2.96	1191	0.81
	Total	**8.03**	**10.79**	**1920**	**0.81**

The placement, TS distribution and routing congestion analysis are shown in Fig. 5. The placements and the TS via distribution are shown in the first two pictures of each tier, respectively. The routing congestions on the right are analyzed by the Cadence Encounter, where the shading spots represent congested regions. Although there are congested region in the middle tier, we observe similar congestions in the 2D implementation, thus the comparison is fair. Since the MITLL 3D IC technology provides only 3 metal layers for each tier, the routing congestions can be probably solved by adopting 3D IC technologies that provide more routing resources, or by integrating a routability-driven 3D placer in 3D-Craft.

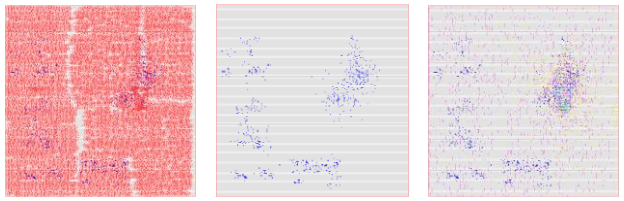

(a) Bottom tier

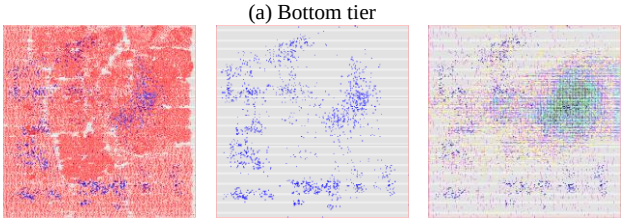

(b) Middle tier

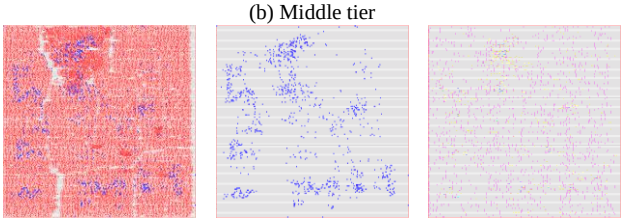

(c) Top tier

Fig. 5. The 3D placement by mPL-3D, the TS via distribution, and the routing congestion analysis by Cadence Encounter

V. Conclusions and Future Work

In this paper we propose our extension of OpenAccess for 3D designs and present 3D-Craft as a referenced 3D physical design flow. The database architecture for 3D designs is described, and the implementation details based on an OpenAccess extension through the *oaAppDef* mechanism is introduced. Several physical design tools with 3D awareness are integrated in the 3D-Craft, including the 3D placer mPL-3D and the 3D global router TMARS with thermal TS via insertion. The application of 3D-Craft on the open-source processor LEON3 shows that the 3-tier 3D implementation can reduce both half-perimeter wirelength and routed wirelength by about 30% compared to the 2D implementation.

Future work includes the integration of the following physical design tools: 3D cube packing tools which support planning real 3D modules, the 3D power/ground network optimization tools, and the 3D clock tree routing tools.

Acknowledgment

This research was partially supported by IBM under a DARPA subcontract, and supported by National Science Foundation under grants CCF-0430077 and CCF-0528583. The authors would like to thank Dr. Philip Chong for his help on the extension mechanisms of OpenAccess, and wish to thank Dr. Marek Turowski and Mr. Patrick Wilkerson from CFD Research Corporation (CFDRC) for the compact resistive network thermal model.

References

[1] L. Cheng, L. Deng, and M. D. Wong, "Floorplanning for 3D VLSI design," Proceedings of IEEE/ACM ASP-DAC 2005, pp. 405-411, 2005.

[2] J. Cong, J. Wei, and Y. Zhang, "A Thermal-Driven Floorplanning Algorithm for 3D ICs," Proceedings of ICCAD 2004, pp. 306–313, 2004.

[3] J. Cong and Y. Zhang, "Thermal-Driven Multilevel Routing for 3-D ICs," Proceedings of the 2005 Conference on Asia South Pacific Design Automation, pp. 121-126, 2005.

[4] J. Cong and Y. Zhang, "Thermal Via Planning for 3-D ICs," Proceedings of the 2005 IEEE/ACM International Conference on Computer-Aided Design, pp. 745-752, 2005.

[5] J. Cong, G. Luo, J. Wei, and Y. Zhang, "Thermal-Aware 3D IC Placement Via Transformation," Proceedings of the 2007 Conference on Asia South Pacific Design Automation, pp. 780-785, 2007.

[6] J. Cong and M. Xie, "A Robust Mixed-Size Legalization and Detailed Placement Algorithm," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 27, no. 8, pp. 1349-1362, August 2008.

[7] J. Cong and G. Luo, "A Multilevel Analytical Placement for 3D ICs," Proceedings of the 2009 Conference on Asia and South Pacific Design Automation, pp. 361-366, 2009.

[8] S. Das, "Design Automation and Analysis of Three-Dimensional Integrated Circuits," PhD Dissertation, Massachusetts Institute of Technology, Cambridge, MA, 2004.

[9] 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," IEEE Design & Test of Computers, vol. 22, no. 6, pp. 498-510, 2005.

[10] B. Goplen and S. Sapatnekar, "Efficient Thermal Placement of Standard Cells in 3D ICs using a Force Directed Approach," Proceedings of the 2003 IEEE/ACM International Conference on Computer-Aided Design, pp. 86-89, 2003.

[11] B. Goplen and S. Sapatnekar, "Placement of 3D ICs with Thermal and Interlayer Via Considerations," Proceedings of the 44th Annual Conference on Design Automation, pp. 626-631, 2007.

[12] A.S. Grove, "Physics and Technology of Semiconductor Devices," John Wiley & Sons, Inc., Hoboken, NJ, 1967.

[13] D. Hill, "Method and System for High Speed Detailed Placement of Cells within an Integrated Circuit Design," US Patent 6370673, April 9, 2002.

[14] Z. Li, X. Hong, Q. Zhou, Y. Cai, J. Bian, H. H. Yang, V. Pitchumani, and C.-K. Cheng, "Hierarchical 3-D Floorplanning Algorithm for Wirelength Optimization," IEEE Trans. Circuits and Systems I, Vol.53, issue 12, pp. 2637-2646, 2007.

[15] Y. Ma, X. Hong, S. Dong and C.K.Cheng, "3D CBL: an Efficient Algorithm for General 3-Dimensional Packing Problems," Proceedings of the 48th MWS-CAS 2005, Vol. 2, pp. 1079- 1082, 2005.

[16] J.-M. Lin and Y.-W. Chang, "TCG: A Transitive Closure Graph Based Representation for Non-Slicing Floorplans," Proceedings of ACM/IEEE Design Automation Conference (DAC-2001), pp. 764-769, Las Vegas, NV, June 2001.

[17] P. Wilkerson, A. Raman, and M. Turowski, "Fast, Automated Thermal Simulation of Three-Dimensional Integrated Circuits," ITherm 2004, Las Vegas, Nevada, 1-4 June 2004.

[18] H. Yan, Q. Zhou, and X. Hong, "Thermal Aware Placement in 3D ICs Using Quadratic Uniformity Modeling Approach," Integration, the VLSI Journal, vol. 42, no. 2, pp. 175-180, 2009.

[19] T. Yan, Q. Dong, Y. Takashima, and Y. Kajitani, "How Does Partitioning Matter for 3D Floorplanning?," Proceedings of the 16th ACM GLS-VLSI , pp.73-78, 2006.

[20] T. Zhang, Y. Zhan, and S.S. Sapatnekar, "Temperature-Aware Routing in 3D ICs," Proceedings of the 2006 Conference on Asia South Pacific Design Automation, pp.309-314, 2006.

[21] P. Zhou, Y. Ma, Z. Li, R.P. Dick, L. Shang, H. Zhou, X. Hong, and Q. Zhou, "3D-STAF: scalable temperature and leakage aware floorplanning for three-dimensional integrated circuits," Proceedings of ICCAD 2007, pp. 590-597, 2007.

[22] http://OpenEDA.Si2.org

[23] http://www.gaisler.com

[24] http://www.cadence.com

[25] http://www.magma-da.com