Physical Design Automation for 3D Chip Stacks – Challenges and Solutions

Johann Knechtel
Masdar Institute of Science and Technology
Abu Dhabi, UAE
knechtel@ieee.org

Jens Lienig
Dresden University of Technology
Dresden, Germany
jens@ieee.org

ABSTRACT

The concept of 3D chip stacks has been advocated by both industry and academia for many years, and hailed as one of the most promising approaches to meet ever-increasing demands for performance, functionality and power consumption going forward. However, a multitude of challenges has thus far obstructed large-scale transition from “classical” 2D chips to stacked 3D chips. We survey major design challenges for 3D chip stacks with particular focus on their implications for physical design. We also derive requirements for advances in design automation, such as the need for a unified workflow. Finally, we outline current promising solutions as well as areas needing further research and development.

Keywords

Electronic design automation; physical design; through-silicon via-based 3D ICs; interposer stacks; monolithic 3D ICs; high-level design challenges; power delivery; thermal management; clock delivery; partitioning; floorplanning; placement; routing; testing; path-finding; chip-package co-design; multi-physics simulation; assembly design kit; design standards

1. INTRODUCTION

3D chip stacks offer the potential to meet current and future requirements of electronic circuits, such as for performance, functionality, and power consumption. Two design paradigms, namely “More Moore” (shrinking device nodes) and “More-than-Moore” (heterogeneous integration), are advocating 3D chip stacks in particular [1] (Fig. 1). Despite significant projected benefits over 2D chips, the overall adoption of 3D stacks still lags behind expectations. What are reasons for the current lack of industrial adoption, and what are specific implications for physical design?

Successful adoption of 3D chip stacks requires addressing different classical and novel challenges which simultaneously affect the manufacturing processes, design practices and physical design tools. If not properly addressed, the fairly complex challenges may render 3D chip stacks commercially unviable, despite recent advances in manufacturing yield and cost reduction. Physical design automation partially meets these challenges at present, but further efforts are needed to fully exploit the potential of 3D chip stacks and to facilitate their commercial breakthrough.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ISPD’16 April 03-06, 2016, Santa Rosa, CA, USA

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4039-7/16/04.

DOI: <http://dx.doi.org/10.1145/2872334.2872335>

This work is licensed under a Creative Commons

Attribution-NonCommercial-ShareAlike International 4.0 License.

Figure 1: The well-known “More Moore” trend is slowly but surely reaching its limits for CMOS technology. “More than Moore”, i.e., exploiting heterogeneous integration, has been identified as novel, important direction. 3D chip stacks offer the potential to meet both trends at the same time.

In this paper, we will first review the diversity of stacking options for 3D chips and formulate related high-level design challenges. Then, along with further needs for design automation, we discuss the following classical challenges and related promising solutions:

- Power delivery and thermal management;
- Clock delivery;
- Partitioning and floorplanning;
- Placement;
- Routability estimation and routing; and
- Testing.

It is important to understand that these challenges are much more complex for 3D chip stacking, and that existing design tools for 2D chips cannot be directly applied.

We also elaborate on novel challenges that are currently (at least partially) addressed, but which require further efforts, such as:

- Pathfinding and design-space exploration;
- Chip-package co-design; and
- Multi-physics simulation and verification.

Our paper concludes with promising trends for physical design of 3D stacks, which meet the identified open challenges, and whose application can ease the adoption of 3D stacks.

2. 3D-CHIP STACKING OPTIONS AND HIGH-LEVEL DESIGN CHALLENGES

3D-chip stacking falls into four categories (Fig. 2): (i) package stacking, (ii) interposer stacks, (iii) through-silicon via (TSV)-based 3D ICs, and (iv) monolithic 3D ICs. Each option has its

Figure 2: Implementation options for 3D chip stacks. Originating with package stacking, 3D integration has evolved through interposer stacks (also known as “2.5D integration”) towards TSV-based and monolithic 3D ICs. Advances in design automation are typically called for the latter two options but are also required for modern, large-scale interposer stacks.

scope of application, with distinctive benefits and drawbacks as well as requirements for design and manufacturing processes [1, 2]. In the following, related key aspects are reviewed and design challenges are outlined.

2.1 Options for 3D Chip Stacking

Package stacking, based on wire bonding and/or flip-chip bonding, has been widely adopted and thus is not reviewed here.

TSV-based 3D ICs have initially attracted the most attention and efforts; many academic/industrial prototypes/products nowadays are based on TSV technology [3–7]. TSVs are metal plugs (copper or tungsten) that penetrate whole dies—which are stacked and bonded—to interconnect those dies. Different stacking scenarios are applicable, such as face-to-back die-to-wafer stacking [2].

Depending on the TSV type, different challenges arise: via-first and via-middle TSVs obstruct the device layer and result in placement obstacles; via-last TSVs obstruct the device layer and the metal layers, resulting in placement and routing obstacles. Due to their relatively large size and intrusive character, TSVs cannot be deployed excessively and/or arbitrarily but have to be optimized in count and arrangement [8, 9]. Note that TSVs do not scale at the same rate as transistors, thus the TSV-to-cell mismatch will remain for future nodes and may even increase [10].

Interposer stacks are a cost-efficient but only seemingly straightforward option for 2.5D or 3D integration [11–13]. Mainly pre-designed dies are stacked in lateral and/or vertical fashion on interposer where redistribution layers and TSVs realize the connectivity. Interposer are typically fabricated as passive carriers, but may also embed components like decaps or even “glue logic” [11]. The concept of interposer stacks supports various applications and is thus widely acknowledged in the industry. Notable products include the Xilinx *Virtex-7* FPGA family [14] and a GLOBALFOUNDRIES prototype with two *ARM Cortex-A9* chips [15].

The design of interposer stacks is obstructed by the lack of dedicated design tools [12]. For example, routing highly-interconnected interposer and the related design of large-scale NoCs requires further research [16]. Other challenges such as simulation and verification of signal integrity across an interposer stack have been recently addressed [17], but require further integration efforts [18].

Monolithic 3D ICs have recently gained more attention [19], thanks to advances of the processing technology. The key feature of monolithic integration is that active layers are sequentially manufactured into one stack rather than bonded using separate dies. Due to their relatively small vias, monolithic 3D ICs enable fine-grain transistor-level integration, which is especially sought-after for high-density logic integration [19].

As for design challenges, routing becomes notably more complex due to high congestion along with increased delays [19]. Novel approaches covering partitioning, placement and routing for monolithic integration are required; such a holistic approach has recently been proposed in [20]. Furthermore, thermal properties differ from “classical” TSV-based 3D ICs: on the one hand, the small vias are not as effective as TSVs for conducting heat towards the heatsink; on the other hand, monolithic stacks do not exhibit “thermal barriers” in form of bonding layers. Thus, the thermal coupling within monolithic stacks is larger and more uniform than for TSV-based 3D ICs. This, in turn, calls for dedicated thermal management [21].

2.2 High-Level Design Challenges

Exploring and selecting the most suitable chip-stacking options for a particular design is much more complex than handling the decisions required for classical 2D-IC design. Given an abstract design description in early planning phases, one has to consider the following challenges, among others:

- How can intellectual property (IP) blocks, pre-designed modules or even whole dies be effectively reused in the stack?
- Into how many dies/layers should the design be split up, and how is the design performing after being spread across multiple dies? What is an appropriate partitioning strategy?
- What are the bandwidth, power and integrity requirements for global interconnects? What is an appropriate system-level interconnect structure for separated components/dies?
- Which bonding and packaging technologies are applicable, and what technology/device node should be considered for each die? Are constraints (such as thermal design, power and signal integrity) satisfiable with the chosen technologies?
- How are heterogeneous components such as MEMS or HF modules designed and properly integrated along with digital modules? (This challenge is especially relevant for 3D chip stacks targeting widely-anticipated and multifunctional applications such as Internet of Things, IoT.)
- How to test separate components/dies and the overall stack?

Most of these challenges are interacting, and any respective decision does impact the overall performance, reliability and cost of the 3D chip stack. It is apparent that solving such a convoluted set of challenges requires sophisticated design know-how, EDA capabilities and well-defined project structures. Specifically, methodologies and tools are required that enable (i) an accurate (yet fast) exploration of the design space and (ii) rapid prototyping/evaluation of different stacking configurations (see Sec. 4).

3. CLASSICAL DESIGN CHALLENGES: AGGRAVATED BUT SOLVABLE

Classical challenges such as power delivery, floorplanning and placement become much more complex when designing 3D chip stacks instead of 2D chips. Existing methodologies and tools for 2D-chip design cannot be directly applied and are only to some degree extendible. During the recent years, however, many efforts in academia and industry have been undertaken which render those classical challenges still aggravated but solvable. In the following, key challenges and promising solutions are discussed.

3.1 Power Delivery and Thermal Management

High integration density—a key benefit of 3D stacking—notably complicates both power delivery and thermal management. A stack of d dies/layers potentially exhibits a d -fold power density compared to 2D chips. This implies that approximately $d \times$ power has to be delivered through the stack, without excessive static and dynamic voltage drop, despite the notable parasitics of 3D power-distribution networks (PDNs) [22]. At the same time, approximately $d \times$ of heat has to be removed from the stack. Solutions for power delivery and thermal management are discussed next.

Arrangement of TSVs: For power/ground (PG) TSVs, a distributed, irregular topology is superior to both regularly placed (single) TSVs and clustered TSVs in order to limit power noise [22,23]. Proper planning of TSVs may also serve power delivery and thermal management at the same time. For example, aligned PG-TSV stacks may simultaneously increase heat dissipation and reduce power noise [23, 24]. Regular signal TSVs can be similarly arranged into (possibly aligned) TSV cluster for improved heat conduction without excessive routing overheads [25–27].

Design of PDN and PG grids: A promising (yet expensive) PDN architecture is the “multi-story power delivery” [28], where several power supplies (e.g., one per die) distribute the load more effectively compared to a single power supply. As for decaps, their allocation has to be carefully investigated to manage their complex impact on the power-noise distribution in 3D chip stacks [29].

The design of PG grids should generally account for the overall 3D PDN, not only for the dies each grid is attached to [24, 30].

Low-power design: Consumed power and, in turn, generated heat can be reduced by deployment of low-power circuitry. In this context, sub-/near-threshold circuitry for 3D chip stacks has gained attention; detailed studies and prototypes are available [31,32]. Arrangement of multiple voltages domains in order to reduce power and temperature has been recently proposed as well [33].

Thermal-aware design: Spreading high-power modules in an orchestrated manner, e.g., during thermal-aware floorplanning [27, 34–36] or placement [37], is a simple but effective measure. Along with high temperature comes a strong impact on transistor behaviour; thermal-aware design is thus also important for reliability.

Technological approaches: In order to limit package impedance and external current supply, one can bring voltage converters closer to the logic by “sandwiching” dedicated converter dies into the 3D stack [38]. To reduce power noise, classical CMOS decaps and/or metal-insulator-metal (MIM) decaps can be deployed within each die [39] or even in dedicated “decap dies” as well [30]. As for thermal paths, internal paths (i.e., paths across and within dies) as well as external paths (i.e., paths to the heatsink and the package) can be improved by introducing micro-fluidic channels [40].

3.2 Clock Delivery

The key challenge when designing 3D clock networks is to ensure reliable, uniform and high-speed delivery of clock signals to all sinks which are spread across different dies. Besides thermal management, this is one of the main challenges for “real logic-on-logic stacking” and has many implications for physical design. These concerns, along with promising solutions, are discussed next.

Redundancy and arrangement of TSVs: Since the advent of TSV technology, their reliability has been a major concern—failure of any clock or signal TSV would render the whole stack defective. Thus, different redundancy architectures have been proposed in general [41, 42] and for clock networks in particular [43, 44].

The optimized assignment of clock signals to TSV arrays—which are generally more practical than separate TSVs [8]—has also been proposed for design of reliable and low-power clock networks [45]. Another flow [46] accounts for PG- and signal-TSV obstacles during placement of clock TSVs, but neglects their co-optimization while arranging those different types of TSVs.

Holistic and variation-aware design: The design of clock networks is highly coupled to most physical-design stages like placement, routing and timing optimization. All stages should account for the 3D clock network’s properties, e.g., arrangement of clock TSVs and clock-trees on individual dies. Capabilities for tuning the network and to provide feedback to other stages are also needed.

Another requirement for clock delivery is the analysis of timing variations and parameter variations, which are a general concern for 3D chip stacks with their inter-die variations. These variations can be mitigated by implementation of multiple clock domains [47]. Note that variations are exacerbated in interposer stacks and TSV-based 3D ICs, mainly due to thermo-mechanical stress [48].

Clock-network synthesis: Minimizing the clock skew is a key objective for clock-network synthesis. It is considered by most previous studies such as [43–45, 49–51]. Variation-aware timing simulation [49, 52], statistical clock-skew models [53] or post-silicon techniques such as body biasing [50] have been applied as well.

As for network topologies, it has been shown that H-trees are superior to classical grid networks in terms of power and mean skew but may suffer from larger maximum skew due to variations [52]. Other proven approaches such as the method of means and medians (MMM) and deferred-merge embedding (DME) have been successfully extended as well [51]. The related approach in [51] allows to effectively trade-off TSV count and power consumption, to manage slew variations and to insert buffers, all within low runtime.

Testability and yield management: Stacking only known-good dies is essential in order to manage the overall yield of 3D chips. This stacking paradigm requires the clock networks to be tailored for pre-bond testability which can be achieved by encapsulating separate, fully functional clock networks into individual dies.¹ These separate networks should be subsequently linked by multiple TSVs to improve wirelength and clocking performance [43].

3.3 Partitioning and Floorplanning

Traditionally, the main objectives for partitioning are to (i) minimize the cuts/connections between partitions and to (ii) reduce the overall design complexity by means of divide-and-conquer. For 3D stacks, however, the objective (i) has to be carefully adapted. For example, min-cut partitioning is not practical since it cannot account for the stacking order of partitioned layers/dies which, in turn, impacts the overall number of vertical interconnects [56]. Furthermore, depending on other, more pressing criteria such as thermal management, the minimization of cuts/connections may not be the primary concern.

In general, partitioning tools need to become more technology-aware in order to properly account for different stacking scenarios and their implications. For example, monolithic stacks are vertically interconnected by very small vias which can (and should) be employed in large numbers, to enable high-performance and low-power transistor-level integration [19]. Nevertheless, related partitioning must account for routing congestion [20]. For TSV-based 3D ICs, it depends on the particular dies’ stacking interface whether interconnects should be limited: a face-to-back (or back-to-back) interface requires TSVs which may need to be restricted, while a face-to-face interface allows for a large number of interconnects, comparable to monolithic stacks. Practical approaches for 3D partitioning must also address other design concerns such as power-density constraints [57] or performance-aware logic splitting [6].

Floorplanning determines the module arrangement within (nowadays mainly pre-fixed) die outlines. For 3D chip stacks, the module arrangement is interdependent across the whole stack, rendering 3D floorplanning much more complex than 2D floorplanning.²

Existing floorplanning methodologies already address key challenges such as thermal management [27, 34–36], co-arrangement of modules and TSVs [59] or planning of system-level buses [27, 36, 60]. Still, 3D floorplanning is a highly technology-dependent and iterative process—fast, accurate and configurable design evaluation is currently targeted but still to be enhanced [61, 62]. So far, only few tools offer high-level, multi-objective exploration, which is needed for microarchitecture-focused 3D design [55, 63]. Furthermore, splitting and folding modules across multiple dies—to potentially reduce power and wirelength even more—has to be also

¹Modules are preferably (re)used as is—splitting them across dies is only warranted for particular scenarios [8, 54, 55]. Thus, timing paths remain mostly within dies which implies that embedding of functional clock networks into separate dies is practical.

²Wang *et al.* [58] have shown the NP-complete nature of both 2.5D floorplans (zero-height modules; emulated 3D arrangements) with more than two dies and for 3D floorplans (non-zero height modules; native 3D arrangements) in general.

explored during floorplanning (and during placement, see below). Related tools/flows are available [54, 55, 64, 65], but they tend to neglect properties of stacking interfaces and/or apply stochastic optimization where results are largely irreproducible and unsteady.

3.4 Placement

As with floorplanning, placement for 3D chips becomes much more challenging. The main objectives are (i) wirelength minimization, (ii) thermal management, and (iii) vertical-interconnect-aware placement or co-placement of cells and vertical interconnects (applies for TSV-based 3D ICs). Placement may further account for specific properties of 3D stacks, e.g., thermo-mechanical stress induced by TSVs (which can even be exploited for timing [66]). There are three categories for placement approaches, discussed next.

Folding-based placement reuses 2D placement solutions/tools and derives a 3D placement by folding modules and applying local refinements [67]. The main benefit is that proven 2D placers can be applied; the disadvantage is the limited consideration of implications arising from die stacking, e.g., increased routing congestion. Recall that folding has been recently advocated for 3D floorplanning as well; the potential of orchestrated folding for floorplanning and placement, however, has not been addressed yet.

Analytic placement is based on numerical analysis of non-linear equations which encode objectives such as minimization of wirelength and cell overlap. In the context of 3D chips, analytic placement is either tailored as 2.5D or 3D placement: the 2.5D approach encodes cell-die assignments as fixed variables, but still accounts for any interdependencies of cells across different dies [37, 66, 68–70]; the 3D approach sets cell-die assignments as flexible variables to be optimized as part of the overall placement solution [37, 71, 72].

Analytic placement suggests itself for placement of 3D chips—cells are distributed optimally (“only” with respect to the tailored equations) throughout the 2.5D/3D domain. However, it has some drawbacks: its complexity calls for techniques such as clustering and coarsening [72] which may lower the final design quality; some formulations (e.g., Huber functions) offer only local smoothness which requires further efforts for global smoothing across the third dimension [37]; any issue/objective not readily accounted within the placement equations (such as thermo-mechanical stress induced by TSVs) may be notably exacerbated and difficult to post-process. Thus, analytic placement is usually not applied as stand-alone but rather as integrated stage (for global placement), as described next.

Hierarchical placement typically applies three stages: global placement, legalization and detailed placement. For 3D chip stacks, most hierarchical placers invoke partitioning and/or floorplanning before actual placement [37, 66, 68–70, 73]. This way, related efforts for module arrangement are acknowledged and cells are pre-assigned to particular dies (but not necessarily fixed to them). This, in turn, helps to improve the convergence of 3D placement. Note that hierarchical placement is not restrictive in terms of applicable placement techniques. For example, Kim *et al.* [69] apply force-directed (2.5D/3D) global placement, Goplen *et al.* [73] promote 3D recursive-bisection placement, and Lu *et al.* [71] propose (globally smooth) electrostatics-based 3D placement.

For placement of TSV-based 3D ICs, different thermal- and TSV-aware 3D placers have been proposed [68, 70, 73]. In this context, the co-placement of cells and TSVs has proven beneficial for thermal management [68, 70] as well as for reducing wirelength [69] and critical delays [66]. In order to avoid potentially opposing and/or unsteady TSV arrangement across partitioning, floorplanning and placement, a closer integration of those stages with dedicated feedback loops to orchestrate TSV (co-)placement is required.

3.5 Routability Estimation and Routing

Routability estimation provides fast but limited local and global predictions of routing supply and demand, without invoking actual routing. This is particularly relevant when designing 3D stacks in order to fully exploit their potential for massive interconnectivity

and to enable high-performance and low-power 3D chips. The key challenge for both routability estimation and routing of 3D stacks arises from the 3D arrangement of nets, i.e., nets are spanning more than one die, which requires to account for vertical interconnects. Related EDA techniques are reviewed next.

Analytic estimation is mainly based on simple geometric and/or global estimates and, thus, not effective for controlling physical-design stages, e.g., to avoid local routing congestion. Nevertheless, analytic estimates are helpful for fast design evaluation; the well-known half-perimeter wirelength (HPWL) is such an estimate. However, note that the HPWL should be determined step-wise in 3D chips to achieve reasonable accuracy; all partial nets (i.e., net segments encapsulated in separate dies/layers) are to be independently estimated and subsequently summed up [8]. Kim *et al.* [74] proposed an interconnect model which is TSV-aware and emulates optimal buffer insertion. Thus, this model helps to efficiently evaluate wirelength, delay and power consumption of 3D chips.

Heuristic estimation applies computationally-inexpensive probabilistic models to capture local routing demand or congestion. Models for 3D chips have to account for vertical interconnects, along with their capacities, distribution and potential blockages. Fischbach *et al.* [75] extended the well-known concept of routing graphs which may serve as generic data structure or “backbone” for any model of choice. Depending on the desired accuracy-runtime trade-off, their extended graph encodes each die’s multiple metal layers as one merged graph layer or as separate layers. In any case, the stacking and interconnects technology defines the graph’s vertical capacities. Panth *et al.* [20] use such a 3D routing graph along with simplified construction of 3D rectilinear Steiner trees for efficient routability estimation in monolithic 3D chips.

Global routing for 3D chips has already been implicitly addressed by conventional methodologies due to the routing layers’ “3D arrangement” in 2D chips. Nevertheless, dedicated global routers have been proposed for TSV-based 3D ICs to account for routing obstacles (induced by TSVs) and to facilitate thermal management [26, 76–79]. Most of these routers construct Steiner or minimum spanning trees, use thermal-RC-network analysis, and (re)arrange signal TSVs, possibly along with dummy thermal TSVs and/or thermal wires. In general, accounting for signal, PG, clock and thermal TSVs along with their networks—all competing for placement and routing area—poses a major design challenge [80]. Thus, global routing (along with partitioning, floorplanning and placement) has to implement an orchestrated (re)arrangement of TSVs to enable design closure.

As for global routing of interposer stacks, early work by Minz *et al.* [81] focused on delay-, congestion- and crosstalk-aware redistribution of pins and nets to properly access and utilize the interposer’s routing channels. In order to reduce TSV usage, Wang *et al.* [82] developed a 3D NoC with shared vertical links, and Foroutan *et al.* [83] optimized the assignment of such shared links in “vertically-partially-connected 3D NoCs”. Loh *et al.* [16] outlined open challenges for large-scale and heterogeneous 3D NoCs, such as the synchronization of routers across sub-networks with potentially different topologies.

Detailed routing is not fundamentally different for 3D stacks; its task remains routing the net segments within metal layers (of individual dies). Assuming that multiple types of interconnects and obstacles may be modeled, available detailed routers can be reused.

3.6 Testing

As indicated when discussing testability of clock networks (see Sec. 3.2), stacking known-good dies is essential for 3D chips. Further testing approaches of 3D chip stacks are described next.

Fault models and at-speed testing has been proposed for both interposer stacks [84–86] and TSV-based 3D ICs [43, 87, 88]. These studies are typically focused on specific scenarios. For example, Taouil *et al.* [87] developed a methodology tailored for testing of memory-on-logic stacks and Deutsch *et al.* [88] applied thermo-

mechanical-stress-aware generation of test patterns for TSV-based 3D ICs. Offering a more holistic approach, Wang *et al.* [84] enabled unified testing of wires, microbumps and TSVs for interposer stacks, with low overhead and compliance to the *IEEE 1149.1* standard. Agrawal *et al.* [89] proposed an efficient heuristic for test-flow selection which can be applied for different configurations of 3D chip stacks; it is notably more efficient and qualitatively competitive to an optimal approach when stacking up to three dies.

Since Design-for-Test (DfT) architectures require access to all modules within 3D chip stacks, such architectures should be based on well-defined components and testing interfaces (see Sec. 5).

3.7 Needs for Advances in Design Automation

We discussed how classical challenges become much more complex when designing 3D chip stacks, and we reviewed promising and available solutions. There is nevertheless a need to further improve “classical” design automation, with particular focus on:

- Synchronized arrangement of all types of TSVs during floorplanning, placement, design of clock networks and routing;
- Perpetual consideration of variations—mainly induced by chip stacking in general and/or by thermo-mechanical stress around TSVs in particular;
- Technology- and stacking-aware design planning and evaluation, in particular for partitioning and floorplanning;
- Measures for 3D design-space exploration, e.g., by exploring options for module folding during partitioning and/or floorplanning, or by system-level arrangement of components and dies within the 3D stack; and
- System-level design of global interconnects. This requires to optimize buses and vertical interconnects (mainly relevant for TSVs) along with design objectives/constraints for, e.g., bandwidth, power consumption, delay and signal integrity.

Note that some of these needs (e.g., 3D design-space exploration) have recently been addressed to some degree by high-level design automation (see Sec. 4).

As outlined in Fig. 3, EDA tools should be unified into a workflow in order to meet (3D-specific) classical as well as novel design challenges. This calls for, among others, extensive code modularization, the design/use of internal/external APIs and for file formats tailored for exchanging 3D-chip designs (see Sec. 5 for the latter).

4. NOVEL DESIGN CHALLENGES AND EMERGING SOLUTIONS

With the aforementioned classical design challenges of 3D chip stacks being addressed at least to some degree by today’s tools, the focus for EDA R&D is shifting towards more high- and system-level challenges such as simulation of multi-physics effects. These challenges are especially pressing for heterogeneous stacks, comprising components such as MEMS or photonics along with “regular” digital modules. We outline promising approaches and trends suitable to address such novel challenges next.

4.1 Pathfinding and Design Exploration

Due to the complex and diversified nature of 3D chip stacks, the automation of their high-level design is extremely challenging. Novel tools are required to explore (at an early stage) the complex impact arising from any decision/step taken in the design flow on criteria such as functionality and power consumption. *Pathfinding* is addressing this need by stepwise passing down high-level descriptions within customized EDA suites. Feedback loops are essential for this purpose, to pass down specifications (e.g., physical constraints) as well as to annotate findings back (e.g., simulation results). Pathfinding flows should cover the following three stages:

1. **System-Level Design:** A high-level description (e.g., in SystemC) is generated and/or iteratively adapted. Technology

Figure 3: Workflow for designing 3D chip stacks. Pathfinding and prototyping link system-level and IC-level design; this link eases the design closure for the highly iterative process. Modeling and simulation as well as management of TSVs must interact with multiple design stages. Most stages require thermal- and variation-aware optimization; feedback loops are essential.

- and stacking configurations have to be already considered, e.g., stacking impacts the options for high-level partitioning.
2. **Component Design:** Components are abstract design modules such as RTL modules—they serve as “bridging” parts between system design and physical design. Components must be modularized in order to encapsulate the high-level design and to enable design reuse for different 3D stacks.
3. **Physical-Design Prototyping:** The components are then fed to (customized) physical-design stages. In contrast to classical physical design, the focus lies initially on fast design exploration, i.e., to obtain “coarse layouts”. The components are also to be annotated (e.g., with their power consumption) to facilitate guidance of subsequent pathfinding iterations.

Pathfinding for 3D chip stacks has been initially proposed in 2009 by Milojevic *et al.* [90]. They applied automated synthesis of RTL modules and physical-design prototyping using feedback loops. More recently, Martin *et al.* [18] proposed to encapsulate TSV arrays using parameterized multi-port modules. They do so in order to feed 3D chip stacks containing such and other components to electromagnetic solvers for analysis of crosstalk across the whole stack. Yazdani and Park [91] automated and optimized the arrangement of buffers, Cu pillars and package bumps for *Wide I/O* memory integration [92] in interposer stacks. Priyadarshi *et al.* [93] proposed transaction-level-based and thermal-aware pathfinding, complementing previous RTL-based approaches.

Besides these mainly academic efforts, commercial tool chains have been becoming available as well, for example *Sphinx 3D Path Finder* (used in [18]) or *Mentor Graphics Xpedition Path Finder Suite* [94].

4.2 Chip-Package Co-Design

The diverse nature of 3D chip stacking, along with its manifold technology configurations, calls for co-design of chip and package. Ideally, the whole 3D stack should be designed within one EDA suite. However, this is not necessarily practical due to several reasons such as:

1. The design of complex 3D chip stacks is typically the effort of a large team, with several engineers (potentially working

in separate companies and/or different locations) being responsible for different parts of the whole stack.

- Existing design know-how and EDA tools are mainly focused on individual parts or technologies, not whole 3D stacks.
- The multitude of interfaces in 3D stacks require—both in terms of physical and functional design—measures for design orchestration and/or a unified data base; this is difficult to establish for heterogeneous design environments.

As illustrated in Fig. 3, it is more practical to introduce and agree on a common workflow, possibly with dedicated points for data handover and feedback iterations between different design parties.

An elaborate example for co-designing a mixed-signal 3D stack is given by Cederström in [95]. Multiple EDA tools have been leveraged, applied by different engineers, linked via custom scripts, and configured to exchange design data. A representation of the whole 3D stack, usable in such a heterogeneous design environment, is called for (see Sec. 5 for a related, novel solution). It has been stressed that TSVs are physically part of the interposer but functionally part of the transceivers (in the ICs). This required multiple iterations for design closure, along with notable manual effort. Cederström hence proposes that the top metal layers, pins and bumps of the ICs are managed by packaging engineers. This way, the iterations for designing the package and the transceiver-TSV links are eased and IC designers are involved more effectively.

Analysing TSVs more efficiently for their impact on whole 3D chip stacks has been successfully pursued for some time; the related modeling and simulation techniques (see “medium” design phases, Sec. 4.3) can also be extended for chip-package co-design [96].

Pathfinding tools may be helpful for co-design as well; for example, *Mentor Graphics Xpedition Path Finder* [94] is tailored for chip-package design with features such as a “virtual die model” (to encapsulate details of IC design) or the generation of pin arrays.

4.3 Multi-Physics Simulation and Verification

Simulation is an integral part of design automation while verification is traditionally separated from the actual physical design. Because of strong coupling of the different physical domains—especially but not exclusively occurring in complex, heterogeneous stacks—simulation and verification for 3D chip stacks should account for multi-physics effects (Fig. 4). Such a framework has been proposed by Schneider *et al.* [97]. They advocate tailoring simulation techniques for each design level; they further promote the (re)use of parameterizable models.

Next, we review measures for simulation and verification of 3D chip stacks across different design levels.

At transistor level, i.e., the lowest level of abstraction, very detailed simulation models are required—they must capture all active and passive devices along with their physical properties. For the thermal and mechanical domains, models are implemented as fine-grain meshes which are fed to finite element/difference analysis, e.g., using tools from ANSYS or COMSOL. These models are also used to verify the reliability and performance of 3D interconnects in detail, e.g., for TSV arrays [98–100]. For the electrical domain, well-known techniques such as SPICE simulations, LVS or ERC verification are applied as well, but with additional consideration of 3D interconnects and multi-physics effects [101]. Some tools (e.g., *Mentor Calibre*) have also been tailored for this need.

For “medium” design phases, such as placement and routing, models are more abstract and capture the system behaviour. Here, the main challenge for simulation is to survey the complex, multi-physics coupling between individual components and to evaluate the overall system behaviour, all in relatively short runtime. Models are often derived from detailed simulations of small components, e.g., single IP blocks or few TSVs. These “characterizing simulations for base components” are independent from the actual design process and can be conducted in advance by experienced engineers, who then encapsulate their findings, typically into parameterizable modules or analytical models. A prominent example for this is

Figure 4: Coupling of various domains in 3D chips, along with potential implications for reliability and performance, require simulation and verification to account for multi-physics effects.

the simulation of TSV arrays—they have been successfully modeled using manifold approaches: as multi-port components along with S-parameters for simulation of signal coupling [17, 18, 102]; as MIMO channels to regulate equalization of such coupling [103]; via superposition of TSV’s stress components to capture their impact on power and/or timing [48, 96]; and via superposition of stress components to determine thermo-mechanical stress itself [9, 96, 100]. The automated generation of such models, also for other components than TSV arrays, remains an open challenge [104].

High-level design renders simulation challenging due to its abstract nature; 3D chips are typically described as a chip-package stack in SystemC, VHDL or Verilog, or even in custom languages. As with pathfinding (Sec. 4.1), simulations are initially approximative (and rely on design experience); more accurate simulations will become feasible after subsequent design/simulation stages fed their findings back. In general, models have to be scalable to support 3D stacks with diverse, multi-scale geometries [102]. Regarding needs for EDA tools, a unified data handling of high-level descriptions is sought after, to ease streaming the 3D-chip configuration to and from different simulation and design tools [105]. A flexible configuration or even the generation of PDKs is required to streamline verification at this chip-package system level [105]; Cibrario *et al.* [106] released such a PDK generator for the (currently widely adopted) 3D stacking technologies from CEA-LETI.

5. PROMOTING 3D CHIPS VIA DESIGN STANDARDS AND FILE FORMATS

Design standards and file formats for seamless data exchange are important measures to achieve the advances required in design automation for 3D chip stacks more efficiently. As for standardization, recent progress covers memory integration and testing:

- The JESD229 standard [92], commonly known as *Wide I/O*, considers one up to four memory dies stacked on top of a controller die, with four up to eight 128-bit-wide memory channels. The standard further covers details on AC and DC characteristics, packages, and bump assignments. First *Wide I/O* stacks and design flows are available [12, 91, 96, 107].
- JESD235A [108], better known as *High Bandwidth Memory (HBM)*, is an alternative standard for memory integration and adopted by industry, e.g., by Hynix [109] and Nvidia [110].
- Yet another standard, the *Hybrid Memory Cube (HMC)* [3], has recently gained more attention as well. It is backed by Altera, Micron, Open-Silicon, Samsung and Xilinx.
- As for testing, standards are currently under development; *IEEE P1838* evolves as the most promising option [111]. The therein proposed wrapper circuitry enables controllability and observability at the die boundaries. Besides the novel wrapper, existing test facilities are considered for reuse whenever possible: *IEEE 1149.x* for test access, *IEEE 1500* for die test, and *IEEE P1687* for internal debugging.

These standardization efforts are crucial for adoption of 3D chip stacks since they facilitate first commercial products (i.e., memory stacks) and are streamlining DfT, which is a pressing concern from industrial perspective. However, further aspects such as thermal management and power delivery (along with related analysis/verification) also call for standardization efforts [102].

As indicated before, another essential need is the seamless data exchange between different parties. Heinig and Fischbach [112] proposed an *assembly design kit (ADK)* for this purpose, which is analogous to the well-known PDK but tailored for 3D chip stacks. An ADK serves the exchange of both design data as well as manufacturing requirements. Using input from all parties, it encodes details of manufacturing steps and procedures, material properties, geometrical descriptions of the stack along with its components and interconnects, as well as assembly rules such as minimal pitches. Besides, an ADK contains customized data converters to “translate” any 3D-stack design into appropriate settings/commands for the manufacturing equipment. Ferguson and Ramadan (*Mentor Graphics*) also advocate the concept of ADKs [113]. They highlight its potential for simplified verification and sign-off procedures, especially for interfaces between dies and/or the package.

6. SUMMARY

In this paper, we discuss the most relevant aspects of automating the physical-design process for 3D chip stacks. We highlight how this process becomes increasingly difficult and demanding as compared to well-engineered design automation for 2D chips, and we review the state-of-art for such 3D-chip design automation.

Classical challenges such as thermal management, clock delivery and placement also became notably more complex, and have seen much academic and industrial R&D in the past. Novel, 3D-specific challenges, such as pathfinding or multi-physics simulation, have undergone similar progress recently. However, further problems remain such as the system-level design of global interconnects and stacking-aware partitioning. Other challenges to be addressed are more technology-dependent, e.g., the synchronized arrangement of all different types of TSVs. Finally, the automation and data management for high-level design is of particular importance—such measures are essential for initial design exploration, collaborative and distributed design teamwork, and final design closure.

We conclude that design automation for 3D chip stacks is challenging and not yet fully solved, though more and more sophisticated EDA solutions are being proposed. Recent efforts such as ADKs and pathfinding tools, among others, help to streamline the design process and, thus, render 3D chip stacks more applicable.

7. ACKNOWLEDGMENTS

This work was supported by *Mubadala*, Abu Dhabi, UAE, under the *TwinLab 3D Stacked Chips* collaboration between the Masdar Institute of Science and Technology, UAE, and the Dresden University of Technology, Germany, Ref. 372/002/6754/102d/146/64947. The authors are thankful to Prof. Ibrahim (Abe) M. Elfadel (Masdar Institute) for fostering the *TwinLab* collaboration and for his feedback on this paper.

8. REFERENCES

- [1] W. Arden *et al.*, “More-than-Moore” White Paper,” ITRS, Tech. Rep., 2010.
- [2] R. R. Tummala, *System on Package: Miniaturization of the Entire System*. McGraw-Hill Professional, 2008.
- [3] (2014, Feb) Hybrid Memory Cube specification 2.0. [Online]: http://hybridmemorycube.org/files/SiteDownloads/20141119_HMCC_Spec2.0Release.pdf
- [4] S. Borkar, “3D integration for energy efficient system design,” in *Proc. DAC*, 2011, pp. 214–219.
- [5] D. H. Kim *et al.*, “3D-MAPS: 3D massively parallel processor with stacked memory,” in *Proc. ISSCC*, 2012, pp. 188–190.
- [6] G. Neela and J. Draper, “Logic-on-logic partitioning techniques for 3-dimensional integrated circuits,” in *Proc. ISCAS*, 2013, pp. 789–792.
- [7] T. Thorolfsson, S. Lipa, and P. D. Franzon, “A 10.35 mW/GFlop stacked SAR DSP unit using fine-grain partitioned 3D integration,” in *Proc. CICC*, 2012, pp. 1–4.
- [8] J. Knechtel, I. L. Markov, and J. Lienig, “Assembling 2-D blocks into 3-D chips,” *Trans. CAD*, vol. 31, no. 2, pp. 228–241, 2012.
- [9] C. Zhang and L. Li, “Characterization and design of through-silicon via arrays in three-dimensional ICs based on thermomechanical modeling,” *Trans. ED*, vol. 58, no. 2, pp. 279–287, 2011.
- [10] V. S. Nandakumar and M. Marek-Sadowska, “Layout effects in fine-grain 3-D integrated regular microprocessor blocks,” in *Proc. DAC*, 2011, pp. 639–644.
- [11] J. H. Lau, “TSV interposer: The most cost-effective integrator for 3D IC integration,” SEMATECH Symposium Taiwan, 2011.
- [12] D. Milojevic *et al.*, “Design issues in heterogeneous 3D/2.5D integration,” in *Proc. ASPDAC*, 2013, pp. 403–410.
- [13] C. Zhang and G. Sun, “Fabrication cost analysis for 2D, 2.5D, and 3D IC designs,” in *Proc. 3DIC*, 2012, pp. 1–4.
- [14] P. Dorsey, “Xilinx stacked silicon interconnect technology delivers breakthrough FPGA capacity, bandwidth, and power efficiency,” Xilinx, Inc., Tech. Rep., 2010.
- [15] (2013, Nov) Open-Silicon and GLOBALFOUNDRIES demonstrate custom 28nm SoC using 2.5D technology. [Online]: <http://www.open-silicon.com/2013/11/open-silicon-globalfoundries-demonstrate-custom-28nm-soc-2-5d-technology/>
- [16] G. H. Loh *et al.*, “Interconnect-memory challenges for multi-chip, silicon interposer systems,” in *Proc. MEMSYS*, 2015, pp. 3–10.
- [17] W. Yao *et al.*, “Modeling and application of multi-port TSV networks in 3-D IC,” *Trans. CAD*, vol. 32, no. 4, pp. 487–496, 2013.
- [18] B. Martin, K. Han, and M. Swaminathan, “A path finding based SI design methodology for 3D integration,” in *Proc. ECTC*, 2014, pp. 2124–2130.
- [19] Y.-J. Lee and S. K. Lim, “Ultrahigh density logic designs using monolithic 3-D integration,” *Trans. CAD*, vol. 32, no. 12, pp. 1892–1905, 2013.
- [20] S. Panth *et al.*, “Placement-driven partitioning for congestion mitigation in monolithic 3D IC designs,” *Trans. CAD*, vol. 34, no. 4, pp. 540–553, 2015.
- [21] S. K. Samal *et al.*, “Fast and accurate thermal modeling and optimization for monolithic 3D ICs,” in *Proc. DAC*, 2014, pp. 1–6.
- [22] M. B. Healy and S. K. Lim, “Distributed TSV topology for 3-D power-supply networks,” *Trans. VLSI*, vol. 20, no. 11, pp. 2066–2079, 2012.
- [23] H.-T. Chen *et al.*, “A new architecture for power network in 3D IC,” in *Proc. DATE*, 2011, pp. 1–6.
- [24] J. Knechtel *et al.*, “Multiobjective optimization of deadspace, a critical resource for 3D-IC integration,” in *Proc. ICCAD*, 2012, pp. 705–712.
- [25] Y. Chen *et al.*, “Through silicon via aware design planning for thermally efficient 3-D integrated circuits,” *Trans. CAD*, vol. 32, no. 9, pp. 1335–1346, 2013.
- [26] P. Hsu, H. Chen, and T. Hwang, “Stacking signal TSV for thermal dissipation in global routing for 3-D IC,” *Trans. CAD*, vol. 33, no. 7, pp. 1031–1042, 2014.
- [27] J. Knechtel, E. Young, and J. Lienig, “Planning massive interconnects in 3D chips,” *Trans. CAD*, vol. 34, no. 11, pp. 1808–1821, 2015.
- [28] P. Jain *et al.*, “A multi-story power delivery technique for 3D integrated circuits,” in *Proc. ISLPED*, 2008, pp. 57–62.
- [29] K. Kim *et al.*, “Effects of on-chip decoupling capacitors and silicon substrate on power distribution networks in TSV-based 3D-ICs,” in *Proc. ECTC*, 2012, pp. 690–697.
- [30] G. Huang *et al.*, “Power delivery for 3D chip stacks: Physical modeling and design implication,” in *Proc. EPEPS*, 2007, pp. 205–208.
- [31] S. Samal *et al.*, “Ultralow power circuit design with subthreshold/near-threshold 3-D IC technologies,” *Trans. CPMT*, vol. 5, no. 7, pp. 980–990, 2015.
- [32] D. Fick *et al.*, “Centip3De: A cluster-based NTC architecture with 64 ARM Cortex-M3 cores in 3D stacked 130 nm CMOS,” *J. Solid-State Circ.*, vol. 48, no. 1, pp. 104–117, 2013.
- [33] S. Jin *et al.*, “On optimizing system energy of voltage-frequency island based 3-D multi-core SoCs under thermal constraints,” *Integration*, vol. 48, pp. 36–45, 2015.
- [34] J. Cong and Y. Ma, “Thermal-aware 3D floorplan,” in *Three Dimensional Integrated Circuit Design*, Y. Xie, J. Cong, and S. Sapatnekar, Eds. Springer US, 2010, ch. 4, pp. 63–102.
- [35] P. Budhathoki *et al.*, “Integrating 3D floorplanning and optimization of thermal through-silicon vias,” in *3D Stacked Chips – From Emerging Processes to Heterogeneous Systems*, I. A. M. Elfadel and G. Fettweis, Eds. Springer, 2016, ch. 10.
- [36] J. Knechtel, E. F. Y. Young, and J. Lienig, “Structural planning of 3D-IC interconnects by block alignment,” in *Proc. ASPDAC*, 2014, pp. 53–60.
- [37] G. Luo, Y. Shi, and J. Cong, “An analytical placement framework for 3-D ICs and its extension on thermal awareness,” *Trans. CAD*, vol. 32, no. 4, pp. 510–523, 2013.
- [38] J. Sun *et al.*, “3D power delivery for microprocessors and high-performance ASICs,” in *Proc. Appl. Power Electr. Conf.*, 2007, pp. 127–133.
- [39] P. Zhou, K. Sridharan, and S. S. Sapatnekar, “Congestion-aware power grid optimization for 3D circuits using MIM and CMOS decoupling capacitors,” in *Proc. ASPDAC*, 2009, pp. 179–184.
- [40] D. Sekar *et al.*, “A 3D-IC technology with integrated microchannel cooling,” in *Proc. IITC*, 2008, pp. 13–15.
- [41] A.-C. Hsieh and T. Hwang, “TSV redundancy: Architecture and design issues in 3-D IC,” *Trans. VLSI*, vol. 20, no. 4, pp. 711–722, 2012.

- [42] I. Loi *et al.*, "A low-overhead fault tolerance scheme for TSV-based 3D network on chip links," in *Proc. ICCAD*, 2008, pp. 598–602.
- [43] C.-L. Lung *et al.*, "Through-silicon via fault-tolerant clock networks for 3-D ICs," *Trans. CAD*, vol. 32, no. 7, pp. 1100–1109, 2013.
- [44] H. Park and T. Kim, "Synthesis of TSV fault-tolerant 3-D clock trees," *Trans. CAD*, vol. 34, no. 2, pp. 266–279, Feb 2015.
- [45] X. Zhao and S. K. Lim, "TSV array utilization in low-power 3D clock network design," in *Proc. ISLPED*, 2012, pp. 21–26.
- [46] X. Zhao and S. K. Lim, "Through-silicon-via-induced obstacle-aware clock tree synthesis for 3D ICs," in *Proc. ASPDAC*, 2012, pp. 347–352.
- [47] S. Garg and D. Marculescu, "Mitigating the impact of process variation on the performance of 3-D integrated circuits," *Trans. VLSI*, vol. 21, no. 10, pp. 1903–1914, 2013.
- [48] S. Marella and S. Sapatnekar, "A holistic analysis of circuit performance variations in 3-D ICs with thermal and TSV-induced stress considerations," *Trans. VLSI*, vol. 23, no. 7, pp. 1308–1321, 2015.
- [49] J.-S. Yang *et al.*, "Robust clock tree synthesis with timing yield optimization for 3D-ICs," in *Proc. ASPDAC*, 2011, pp. 621–626.
- [50] T.-Y. Kim and T. Kim, "Post silicon management of on-package variation induced 3D clock skew," *J. Semicond. Technol. Scie.*, vol. 12, no. 2, pp. 139–149, 2012.
- [51] X. Zhao, J. Minz, and S. K. Lim, "Low-power and reliable clock network design for through-silicon via (TSV) based 3D ICs," *Trans. CPMT*, vol. 1, no. 2, pp. 247–259, 2011.
- [52] H. Xu, V. F. Pavlidis, and G. De Micheli, "Effect of process variations in 3D global clock distribution networks," *J. Emerg. Technol. Comput. Syst.*, vol. 8, no. 3, pp. 20:1–20:25, 2012.
- [53] X. Zhao, S. Mukhopadhyay, and S. K. Lim, "Variation-tolerant and low-power clock network design for 3D ICs," in *Proc. ECTC*, 2011, pp. 2007–2014.
- [54] M. Jung *et al.*, "On enhancing power benefits in 3D ICs: Block folding and bonding styles perspective," in *Proc. DAC*, 2014, pp. 1–6.
- [55] Y. Ma *et al.*, "Investigating the effects of fine-grain three-dimensional integration on microarchitecture design," *J. Emerg. Technol. Comput. Syst.*, vol. 4, no. 4, pp. 17:1–17:30, 2008.
- [56] Y.-S. Huang, Y.-H. Liu, and J.-D. Huang, "Layer-aware design partitioning for vertical interconnect minimization," in *ISVLSI*, 2011, pp. 144–149.
- [57] H.-L. Chang *et al.*, "A 3D IC designs partitioning algorithm with power consideration," in *Proc. ISQED*, 2012, pp. 137–142.
- [58] R. Wang, E. F. Y. Young, and C.-K. Cheng, "Complexity of 3-D floorplans by analysis of graph cuboidal dual hardness," *Trans. DAES*, vol. 15, no. 4, pp. 33:1–33:22, 2010.
- [59] C.-R. Li, W.-K. Mak, and T.-C. Wang, "Fast fixed-outline 3-D IC floorplanning with TSV co-placement," *Trans. VLSI*, vol. 21, no. 3, pp. 523–532, 2013.
- [60] A. Quiring, M. Olbrich, and E. Barke, "Fast global interconnect driven 3D floorplanning," in *Proc. VLSISOC*, 2015, pp. 313–318.
- [61] R. Fischbach, J. Lienig, and J. Knechtel, "Investigating modern layout representations for improved 3D design automation," in *Proc. GLSVLSI*, 2011, pp. 337–342.
- [62] S. K. Lim, "Research needs for TSV-based 3D IC architectural floorplanning," *J. Inf. Commun. Converg. Eng.*, vol. 12, no. 1, pp. 46–52, 2014.
- [63] M. Healy *et al.*, "Multiobjective microarchitectural floorplanning for 2-D and 3D ICs," *Trans. CAD*, vol. 26, no. 1, pp. 38–52, 2007.
- [64] R. K. Nain and M. Chranowska-Jeske, "Fast placement-aware 3-D floorplanning using vertical constraints on sequence pairs," *Trans. VLSI*, vol. 19, no. 9, pp. 1667–1680, 2011.
- [65] Y. Liu *et al.*, "Fine grain 3D integration for microarchitecture design through cube packing exploration," in *Proc. ICCD*, 2007, pp. 259–266.
- [66] K. Athikulwongse *et al.*, "Stress-driven 3D-IC placement with TSV keep-out zone and regularity study," in *Proc. ICCAD*, 2010, pp. 669–674.
- [67] J. Cong *et al.*, "Thermal-aware 3D IC placement via transformation," in *Proc. ASPDAC*, 2007, pp. 780–785.
- [68] J. Cong, G. Luo, and Y. Shi, "Thermal-aware cell and through-silicon-via co-placement for 3D ICs," in *Proc. DAC*, 2011, pp. 670–675.
- [69] D. H. Kim, K. Athikulwongse, and S. K. Lim, "Study of through-silicon-via impact on the 3-D stacked IC layout," *Trans. VLSI*, vol. 21, no. 5, pp. 862–874, 2013.
- [70] K. Athikulwongse, M. Pathak, and S. K. Lim, "Exploiting die-to-die thermal coupling in 3D IC placement," in *Proc. DAC*, 2012, pp. 741–746.
- [71] J. Lu *et al.*, "ePlace-3D: electrostatics based placement for 3D-ICs," *ArXiv e-prints*, Dec. 2015.
- [72] M.-K. Hsu, V. Balabanov, and Y.-W. Chang, "TSV-aware analytical placement for 3-D IC designs based on a novel weighted-average wirelength model," *Trans. CAD*, vol. 32, no. 4, pp. 497–509, 2013.
- [73] B. Goplen and S. Sapatnekar, "Placement of 3D ICs with thermal and interlayer via considerations," in *Proc. DAC*, 2007, pp. 626–631.
- [74] D. Kim, S. Mukhopadhyay, and S. Lim, "TSV-aware interconnect distribution models for prediction of delay and power consumption of 3-D stacked ICs," *Trans. CAD*, vol. 33, no. 9, pp. 1384–1395, 2014.
- [75] R. Fischbach, J. Lienig, and T. Meister, "From 3D circuit technologies and data structures to interconnect prediction," in *Proc. SLIP*, 2009, pp. 77–84.
- [76] J. Cong and Y. Zhang, "Thermal-driven multilevel routing for 3-D ICs," in *Proc. ASPDAC*, 2005, pp. 121–126.
- [77] T. Zhang, Y. Zhan, and S. S. Sapatnekar, "Temperature-aware routing in 3D ICs," in *Proc. ASPDAC*, 2006, pp. 1–6.
- [78] M. Pathak and S. K. Lim, "Performance and thermal-aware Steiner routing for 3-D stacked ICs," *Trans. CAD*, vol. 28, no. 9, pp. 1373–1386, 2009.
- [79] P.-Y. Hsu, H.-T. Chen, and T. Hwang, "Stacking signal TSV for thermal dissipation in global routing for 3D IC," in *Proc. ASPDAC*, 2013, pp. 699–704.
- [80] Y.-J. Lee and S. K. Lim, "Co-optimization and analysis of signal, power, and thermal interconnects in 3-D ICs," *Trans. CAD*, vol. 30, no. 11, pp. 1635–1648, 2011.
- [81] J. R. Minz and S. K. Lim, "Block-level 3-D global routing with an application to 3-D packaging," *Trans. CAD*, vol. 25, no. 10, pp. 2248–2257, 2006.
- [82] Y. Wang *et al.*, "Economizing TSV resources in 3-D network-on-chip design," *Trans. VLSI*, vol. 23, no. 3, pp. 493–506, 2015.
- [83] S. Foroutan, A. Sheibanyrad, and F. Petrot, "Assignment of vertical-links to routers in vertically-partially-connected 3-D-NoCs," *Trans. CAD*, vol. 33, no. 8, pp. 1208–1218, 2014.
- [84] R. Wang, K. Chakrabarty, and S. Bhawmik, "Interconnect testing and test-path scheduling for interposer-based 2.5-D ICs," *Trans. CAD*, vol. 34, no. 1, pp. 136–149, 2015.
- [85] L.-R. Huang *et al.*, "Parametric fault testing and performance characterization of post-bond interposer wires in 2.5-D ICs," *Trans. CAD*, vol. 33, no. 3, pp. 476–488, 2014.
- [86] S. Huang *et al.*, "Pulse-vanishing test for interposers wires in 2.5-D IC," *Trans. CAD*, vol. 33, no. 8, pp. 1258–1268, 2014.
- [87] M. Taouil *et al.*, "Post-bond interconnect test and diagnosis for 3-D memory stacked on logic," *Trans. CAD*, vol. 34, no. 11, pp. 1860–1872, 2015.
- [88] S. Deutsch *et al.*, "TSV stress-aware ATPG for 3D stacked ICs," in *Proc. ATS*, 2012, pp. 31–36.
- [89] M. Agrawal, K. Chakrabarty, and B. Eklow, "Test-cost optimization and test-flow selection for 3D-stacked ICs," Duke University, Tech. Rep., 2012.
- [90] D. Milojevic *et al.*, "Automated pathfinding tool chain for 3D-stacked integrated circuits: Practical case study," in *Proc. 3DIC*, 2009, pp. 1–6.
- [91] F. Yazdani and J. Park, "Pathfinding methodology for optimal design and integration of 2.5D/3D interconnects," in *Proc. ECTC*, 2014, pp. 1667–1672.
- [92] (2011, Dec) JEDEC standard: JESD229 Wide I/O. [Online]: <http://www.jedec.org/standards-documents/results/jesd229>
- [93] S. Priyadarshi *et al.*, "Thermal pathfinding for 3-D ICs," *Trans. CPMT*, vol. 4, no. 7, pp. 1159–1168, 2014.
- [94] (2014, June) Mentor Graphics launches Xpedition Path Finder suite for efficient IC/Package/PCB design optimization, assembly, and visualization. Mentor Graphics. [Online]: <https://www.mentor.com/company/news/mentor-xpedition-path-finder>
- [95] L. Cederström, "EDA environments for 3D chip stacks," in *3D Stacked Chips – From Emerging Processes to Heterogeneous Systems*, I. A. M. Elfadel and G. Fettweis, Eds. Springer, 2016, ch. 9.
- [96] M. Jung, D. Z. Pan, and S. K. Lim, "Chip/package mechanical stress impact on 3-D IC reliability and mobility variations," *Trans. CAD*, vol. 32, no. 11, pp. 1694–1707, 2013.
- [97] P. Schneider *et al.*, "Towards a methodology for analysis of interconnect structures for 3D-integration of micro systems," in *DTIP*, 2007, pp. 162–168.
- [98] K. H. Lu *et al.*, "Thermo-mechanical reliability of 3-D ICs containing through silicon vias," in *Proc. ECTC*, 2009, pp. 630–634.
- [99] H. M. Lee *et al.*, "Impact of TSV induced thermo-mechanical stress on semiconductor device performance," in *Proc. EDAPS*, 2012, pp. 189–192.
- [100] Y. S. Chan, H. Y. Li, and X. Zhang, "Thermo-mechanical design rules for the fabrication of TSV interposers," *Trans. CPMT*, vol. 3, no. 4, pp. 633–640, 2013.
- [101] R. Fischbach, A. Heinig, and P. Schneider, "Design rule check and layout versus schematic for 3D integration and advanced packaging," in *Proc. 3DIC*, 2014, pp. 1–7.
- [102] M. Swaminathan, "Electrical design and modeling challenges for 3D system integration," in *DesignCon Tutorial*, 2012, pp. 1–31.
- [103] T. Seifert, F. Pauls, and G. Fettweis, "Multi-TSV crosstalk channel equalization with non-uniform quantization," in *3D Stacked Chips – From Emerging Processes to Heterogeneous Systems*, I. A. M. Elfadel and G. Fettweis, Eds. Springer, 2016, ch. 4.
- [104] P. Schneider *et al.*, "Integration of multi physics modeling of 3D stacks into modern 3D data structures," in *Proc. 3DIC*, 2010, pp. 1–6.
- [105] A. Heinig and R. Fischbach, "Overview of 3D CAD design tools," in *Physical Design for 3D Integrated Circuits*, A. Todri-Sanial and C. S. Tan, Eds. CRC Press, Taylor & Francis, 2016, ch. 14, pp. 311–320.
- [106] G. Cibrario *et al.*, "A 3D process design kit generator based on customizable 3D layout design environment," in *Proc. 3DIC*, 2013, pp. 1–5.
- [107] A. Heinig, R. Fischbach, and M. Dittrich, "Thermal analysis and optimization of 2.5D and 3D integrated systems with Wide I/O memory," in *Proc. ITherm*, 2014, pp. 86–91.
- [108] (2015, Nov) JEDEC standard: JESD235A high bandwidth memory (HBM). [Online]: <http://www.jedec.org/standards-documents/docs/jesd235a>
- [109] SK HYNIX INC. (2014, Dec) SK Hynix HBM graphics memory. [Online]: http://www.skhynix.com/inc/pdf/Download.jsp?path=/datasheet/Databook/Databook_Q4'2014_Graphics.pdf
- [110] H. Reiter. (2014, Dec) 3D ASIP 2014: All aboard the 3D IC train! [Online]: <http://www.3dincites.com/2014/12/3d-asip-2014-addresses-3d-benefits-challenges-solutions/>
- [111] E. J. Marinissen, "Status update of IEEE Std P1838," in *Proc. Int. Workshop Testing 3D Stack Integr. Circ.*, 2014.
- [112] A. Heinig and R. Fischbach, "Enabling automatic system design optimization through assembly design kits," in *Proc. 3DIC*, 2015, pp. TS8.31.1–TS8.31.5.
- [113] J. Ferguson and T. Ramadan. (2015, Nov) Why do we need assembly design kits for packages? [Online]: <http://www.3dincites.com/2015/11/why-do-we-need-assembly-design-kits-for-packages/>