Invited: EDA for Heterogeneous Integration

Emad Haque*, Pragnya Nalla†, Chetal Choppali Sudarshan*, Divya Yogi†, Hangyu Zhang†,
Chaitali Chakrabarti*, Vidya A. Chhabria*, Ramesh Harjani†, Jeff Zhang*, and Sachin S. Sapatnekar†
* Arizona State University, Tempe, AZ, USA † University of Minnesota, Minneapolis, MN, USA

***Abstract*—The advent of heterogeneous integration (HI) places new demands on EDA tooling. Building large systems requires (1) methods for chiplet disaggregation that map the system to smaller chiplets, working in conjunction with system-technology co-optimization to determine the right design decisions that optimize computation and communication, together with the choice of substrate and chiplet technologies; (2) multiphysics and multiscale analyses that incorporate thermomechanical aspects into performance analysis, ranging from fast machine-learning-driven analyses in early stages to signoff-quality multiphysics-based analysis; (3) physical design techniques for placing and routing chiplets and embedded active/passive elements on and within the substrate, including the design of thermal and power delivery solutions; and (4) underlying infrastructure required to facilitate HI-based design, including the design and characterization of chiplet libraries and the establishment of data formats and standards. This paper overviews these issues and lays out a set of EDA needs for HI designs.**

***Index Terms*—Heterogeneous integration, chiplets, design automation, disaggregation, power delivery, thermal management.**

I. Introduction

The push towards greater integration is driven by the increasing demands of high-performance computing, artificial intelligence, and big-data applications that require large amounts of data movement. For cloud applications, conventional monolithic computing systems with limited die sizes cannot service the performance needs of modern AI applications: for instance, it takes 20 NVIDIA GPUs to hold one copy of the GPT4 model parameters [1], and data movement costs are very high. The slowdown in Moore's law leaves a gap in satisfying the exponentially-growing hardware needs of advanced chips. For edge applications, bandwidth limitations from sensor/memory to compute modules pose fundamental challenges in the big data era: a 2-bit DVS camera with 4M pixels, uses a 1TBps data rate at 1 MHz [2]; graph neural nets (GNNs) on the Nell dataset with 64K graph nodes require up to 2.7 TB of data to be communicated between processing elements [3]; large language models (LLMs) demand over TBps-level data rates [4].

Advanced packaging offers avenues to mitigate these problems. Heterogeneous integration (HI) technologies place multiple chiplets on a substrate and connect them with dense interconnects, thus providing a viable alternative for integrating large systems at low cost and high yield. Recent years have seen the emergence of EMIB [5], CoWoS [6], and Foveros [7], as well as roadmaps towards finer-pitch inter-chiplet interconnects [8], as technologies based on through-silicon vias (TSVs)/through-glass vias (TGVs), microbumps, hybrid bonding, Cu-Cu bonding, and Cu pillars, have expanded the space of solutions on HI substrates. Recent designs from industry and research groups have demonstrated the viability of HI [9]–[15].

Advances in HI have explored a variety of architectural solutions to address these challenges [6], [9], [11], [12], [14]. One such approach, described in [16], employs hybrid bonding to vertically connect chiplets, with the upper-tier chiplets interfacing directly with the package through through-dielectric vias (TDVs). Alternative architectures leverage silicon interposers to enable high-density connectivity between chiplets. Ongoing research has explored multiple interposer materials and configurations, including passive [6], active [14], glass [17], organic [11], and inorganic [18] substrates. These interposers support a variety of vertical interconnect technologies, such as through-silicon vias (TSVs), TDVs, through-glass vias (TGVs), through-silicon interposers (TSIs), microbumps, copper pillars, and both copper-to-copper (Cu-Cu) and hybrid bonding techniques.

Moreover, HI raises new opportunities for integrating disparate technologies in the form of 2D or 3D chiplets on a 2.5D substrate. For example, digital subsystems may use cutting-edge, advanced CMOS nodes while analog/mixed-signal (AMS) or radio frequency (RF) components may use older CMOS nodes that provide better and more robust performance, or non-CMOS nodes such as GaAs and SiGe; emerging memory technologies may be easily integrated; integrated photonics can provide high-bandwidth communication; etc. An example integrated system may contain sensing (in the form of image sensors, phased arrays, and/or RF antennas), analog preprocessing, and a digital baseband: such a system can provide the horsepower to perform on-chip sensing, rapid computational processing, and on-chip and off-chip communication.

However, even as HI packs greater amounts of computation per unit footprint, it brings forth new challenges at the design phase, which will be further exacerbated as complex systems, with hundreds to thousands of chiplets, are built. Chief among these are:

- How can the complexity of scale posed by these systems, with significantly more transistors and/or more diverse technologies be efficiently handled?
- How can large systems be disaggregated into smaller chiplets?
- How should individual chiplets be designed?
- How can the power delivery solution service the needs of increased amount of power per unit footprint?
- How can the heat generated by this increased power be removed?
- How should local thermal hotspots be avoided and mitigated?
- How can system reliability be assured over the system lifetime?
- How should the chiplets be placed and routed in 3D space, accounting for multiphysics effects arising from electrical, thermal, and mechanical considerations, incorporating cooling, power delivery, test, and security considerations?

All of these issues indicate that a substantial EDA effort is an essential component to build systems with multi-trillion transistors and multi-kW power dissipation. Design techniques must be enhanced to capture multiphysics interactions [19] between the electrical, thermal, and mechanical domains. Analyzers must be able to capture multiscale effects at different length scales, from the package level to the chiplet level to the function block and gate levels. While some ideas from six decades of chip design can be useful, they must be supplemented with new techniques and algorithms, as well as new thinking. In this paper, we attempt to list some key challenges in this domain, and suggest how existing EDA expertise can be adapted and enhanced to meet the substantial challenges of building HI systems.

II. System Disaggregation

Disaggregation refers to the decomposition of the overall system into smaller chiplets interconnected by a communication fabric.

Each chiplet may be fabricated in a different technology, and these choices have ramifications on cost and performance, requiring system-technology co-optimization [20]; chiplets may be built from scratch (incurring higher costs) or from a chiplet library (enabling design reuse, reducing development time, and improving the carbon footprint); communication overheads associated with partitioning the system into multiple chiplets must be factored into the analysis of system performance; new strategies must be developed for these complex systems to overcome verification, test, and security challenges.

Given the number of design choices, disaggregation inherently involves optimization over a complex multidimensional space. An exploration of the cost/benefit tradeoffs must address both the technology choices (choice of design substrate, e.g., silicon, organic, or glass; substrate area; ability to integrate active and passive devices) and system design choices, under cost functions related to power, performance, system throughput, and dollar cost. In this context, the use of chiplet libraries to reduce design turnaround time has attracted a great deal of attention. In [21], a method for building AI hardware accelerators from a chiplet library was demonstrated to give up to 4$\times$ reduction in nonrecurring engineering (NRE) costs with high chiplet utilization. In [22], disaggregated systems have been shown to provide benefits in reducing embodied carbon emissions by 30% over monolithic systems, for example, due to design reuse from the chiplet library, (thus providing benefits beyond NRE costs) and the use of smaller die sizes that provide better yields than a large monolithic die. A methodology for analyzing the optimal range of sizes for disaggregated chips, including tradeoffs between cost, yield, communication, assembly costs, and NRE costs, has been presented in [23]. Communication fabrics using networks-on-package [24] or photonic interconnects for communication [25] within and outside the package form a vital part of the solution. Disaggregation must be supported by efficient modeling that comprehends multiphysics interactions on circuits, using tools such as HISIM [19].

III. Power Delivery

To deliver power to multi-kW HI systems with low losses, a strong focus must be placed on developing robust power delivery solutions that assure power integrity. HI systems can experience sustained power densities of 1–10 W/mm^2 [8]; traditional core power densities are in the range of 0.1–0.5 W/mm^2. This requires a power delivery architecture that operates at the printed circuit board (PCB) level, the HI substrate level, and the chiplet level. At the PCB level, off-chip multistage step-down power converters are used to convert high-voltage (e.g., 48V [26]) low-current power to high-current power at CMOS-compatible voltages, through multiple stages. An intermediate bus voltage (IBV) is used to distribute power on the PCB (e.g., 1.8V [27]) and feed the package, where the IBV is converted to a CMOS-compatible voltage using voltage regulators. At the chiplet level, on-chip voltage regulators, such FIVRs [28], and low-dropout regulators (LDOs) are well-established solutions. Solutions at several levels have been proposed in [14], [26], [29], [30].

At the chiplet level, voltage regulation can be performed using switched-capacitor voltage regulators (SCVRs) [31], [32] or LDOs [33]–[35]. A large load causes the regulator output to degrade, leading to significant efficiency losses: the DC power loss across the power delivery network (PDN) [36] increases quadratically as the output voltage falls. An HI substrate may employ capacitive or inductive converters. In a silicon interposer, it is difficult to build sufficiently high-Q inductors, and SCVRs can be used to convert an input IBV to the output level(s) required by the chiplet. Other types of substrates may integrate inductors and facilitate higher power, particularly for large-area designs. It is well known that the use of distributed smaller voltage regulators, placed close to the load, provides significant advantages over a single large regulator [33], [37]. The problem of selecting converters over the space of design choices, and optimally placing them to meet power integrity constraints, is a significant issue in HI systems. Moreover, these converters must be codesigned with the PDN for maximum effectiveness of the power delivery solution. Methods for automatic PDN construction such as [38] may be useful when adapted to the HI context, and can be coupled with algorithms for decoupling capacitor placement within the HI substrate [39], [40].

IV. Thermal Considerations

Thermal management and power delivery are two sides of the same coin: the power delivery scheme transmits energy into the system so that it can be dissipated during computation, but the resulting heat must be removed from the system lest it should excessively raise on-chip temperature levels. Elevated temperatures not only cause performance degradation, but can also result in accelerated aging [41]. This calls for thermal analysis techniques that can be used to identify hot-spots and support electrothermal and reliability analysis. For early design, ML-based approaches that have shown promise for early design at the full-chip level can be extended to HI systems. For example, in [42], the thermal problem is viewed as translating an input power map image, with knowledge of the heat removal paths, to an output thermal map using a U-Net-based approach. In HI systems, the thermal solution may be more complex, e.g., involving cold plates or active microchannel-based cooling within the HI stack. Such fast models can be used to drive electrothermal analysis, as has been demonstrated at the chip level for a power amplifier structure [43].

Thermal solutions require the use of a mix of passive and active cooling strategies. Passive cooling involves the use of thermal vias [44], while active cooling may involve cold plates [45], [46] near the package, or microchannels within the package [47].

V. Electrothermomechanical Reliability

At elevated temperatures, reliability problems become significant. Most device reliability mechanisms accelerate at higher temperature [48], particularly bias temperature instability and hot carrier injection, which degrade the threshold voltage and drive current of a transistor over time. Hence, reliability must be a first-class objective during HI system design, with cooling solutions designed to mitigate reliability hot spots. Electromigration (EM) is also significantly accelerated by temperature, and therefore it will be essential to build thermally-aware EM design solutions, accounting for both EM in wires [49]–[53] that distribute power and signals, and in bumps [54], [55] that connect the chiplets, substrate, package, and PCB.

Mechanical reliability also plays a very strong role in HI systems, in several ways. First, substrate warping due to mismatches in the coefficients of thermal expansion (CTEs) can cause a large and thin substrate to warp. Today's techniques capture these effects using first-order models or detailed numerical simulations, but it will be essential to develop methods that present tradeoffs between the simplicity of the former and the accuracy of the latter, possibly driven by ML models. Second, issues related to CTE mismatches between bumps/bonding pads and the chiplets, or TSVs and silicon, are known to cause circuit performance drifts [56]: with the increasing integration of sensitive AMS/RF blocks, such analyses will become important. Third, stress migration [57] due to residual stress in the chip structure can exacerbate the impact of EM, and must be factored in during EM analyses.

VI. Physical Design

Almost every performance parameter in an HI system is integrally connected to physical design. Decisions that are made during layout impact the spatiotemporal distribution of power dissipation in the system, and hence the power delivery and thermal solution. In turn, the thermal map of the system affects the performance and reliability of the chiplets, interconnects, and active devices. The complexity of physical design is compounded by several factors:

(1) Placement considerations must incorporate the requirements of *power delivery* solutions. For instance, the use of distributed regulators is more efficient than centralized regulators (as mentioned earlier); PDNs must share interconnect resources with signal/clock nets. These issues have been addressed at the chip(let) level [37], [38], and new approaches must incorporate HI-specific considerations.

(2) Embedded *thermal solutions* for *in situ* cooling remove thermal flux close to the heat-generating elements, but also constitute routing blockages: therefore, the task of thermal management, placement, and routing must be closely intermeshed. Physical design must operate in tandem with fast multiphysics models that capture electrothermal effects. For early-stage design, ML predictors show great promise.

(3) The incorporation of *AMS/RF components* in the system requires methods for distributing clean power supplies, providing appropriate isolation against electromagnetic interference (EMI) for sensitive analog components (chiplets as well as embedded passives and active devices), and shielding/isolation for critical routes.

(4) *Test* considerations drive the need to develop solutions for large systems. The area and communication of test access and test overheads, such as built-in self-test, and their impact on the performance of the primary system, must be factored in during physical design.

Physical design must also account for the need to manage the cost and overhead of communication between chiplets. To a first order, this cost is strongly related to the distance between the chiplets, but also depends on the choice of communication protocol. In future, as substrate interconnect pitches grow finer, this may be supplemented with direct signaling between neighboring chiplets, potentially paving the way for simpler interfaces without strong ESD protections. However, long-distance communication is inevitable and will require the use of networks-on-package [19].

Traditionally, digital chips have largely been built to be synchronous. For large-area systems with dimensions in the hundreds of mm on a side, global synchronization is practically infeasible due to the complexities of distributing a synchronized clock signal over long distances. Therefore, design automation solutions and physical design must consider clock distribution over locally synchronous regions, and synchronization across these regions.

VII. Standards and Representations

In addition to tools and methodologies, data formats and standards play a strong role in the widespread deployment of HI systems. Similar to process design kits (PDKs), there has been a move towards defining package characteristics through assembly design kits (ADKs) [58] that capture design rules and geometries as well as electrical, thermal, and mechanical properties. There have been significant efforts to build standards for inter-chiplet communication, starting from the open-source advanced interface bus (AIB) [59] protocol, and with several newer contenders including BoW [60], and UCIe [61]. For physical representations of 3D design objects, from chiplets to interconnects to PDN structures, the 3Dblox format is moving towards standardization [62], [63]. To scale up to large systems, such efforts are essential and provide structured methods for representing objects and relationships within a 3D HI system.

VIII. Conclusion

This paper attempts to provide an overview of the EDA needs of future HI systems – system disaggregation (including system-technology cooptimization), multiscale multiphysics analyses (interaction of electrical, thermal, mechanical, and reliability considerations), power delivery, thermal management, physical design, verification, and test. The challenge before the community is to build upon past achievements to deliver integrated EDA solutions that solve these problems for HI systems by delivering a new generation of EDA algorithms, tools, and methodologies.

IX. Acknowledgments

This effort is supported in part by the NSF under awards 2212345, 2324945, 2324946, 2403408, 2403409, and 2437795.

References

[1] B. Dally, "Hardware for deep learning," in *Proc. Hot Chips*, 2023.

[2] G. Krishnan *et al.*, "3D-ISC: A 65nm 3D compatible in-sensor computing accelerator with reconfigurable tile architecture for real-time DVS data compression," in *Proc. ASSCC*, 2023, 3 pages.

[3] S. K. Mandal, G. Krishnan, A. A. Goksoy, G. R. Nair, Y. Cao, and U. Y. Ogras, "COIN: Communication-aware in-memory acceleration for graph convolutional networks," *IEEE J. Emerging Sel. Top. Circuits Syst.*, vol. 12, no. 2, pp. 472–485, 2022.

[4] W. Wang and M. Ghobadi, "Spine-free networks for LLM training," *IEEE Micro*, 2025.

[5] R. Mahajan *et al.*, "Embedded multi-die interconnect bridge (EMIB) – a high density, high bandwidth packaging interconnect," in *Proc. ECTC*, 2016, pp. 557–565.

[6] M.-S. Lin *et al.*, "A 7nm 4GHz Arm-core-based CoWoS chiplet design for high-performance computing," *IEEE J. Solid-State Circuits*, vol. 55, no. 4, pp. 956–966, 2020.

[7] D. B. Ingerly *et al.*, "Foveros: 3D integration and the use of face-to-face chip stacking for logic devices," in *Proc. IEDM*, 2019, pp. 19.6.1–19.6.4.

[8] "Heterogeneous Integration Roadmap, Chapter 1," 2023. [Online]. Available: http://eps.ieee.org/hir

[9] W. Gomes *et al.*, "Lakefield and mobility compute: A 3D stacked 10nm and 22FFL hybrid processor system in 12×12mm^2, 1mm package-on-package," in *Proc. ISSCC*, 2020, pp. 144–146.

[10] ——, "Ponte Vecchio: A multi-tile 3D stacked processor for exascale computing," in *Proc. ISSCC*, vol. 65, 2022, pp. 42–44.

[11] T. Burd, N. Beck, S. White, M. Paraschou, N. Kalyanasundharam, G. Donley, A. Smith, L. Hewitt, and S. Naffziger, ""Zeppelin": An SoC for multichip architectures," *IEEE J. Solid-St. Circ.*, vol. 54, no. 1, pp. 133–143, 2019.

[12] S. Naffziger, K. Lepak, M. Paraschou, and M. Subramony, "AMD chiplet architecture for high-performance server and desktop products," in *Proc. ISSCC*, 2020, pp. 44–45.

[13] R. Agarwal, P. Cheng, P. Shah, B. Wilkerson, R. Swaminathan, J. Wuu, and C. Mandalapu, "3D packaging for heterogeneous integration," in *Proc. ECTC*, 2022, pp. 1103–1107.

[14] P. Vivet *et al.*, "IntAct: A 96-core processor with six chiplets 3D-stacked on an active interposer with distributed interconnects and integrated power management," *IEEE J. Solid-State Circuits*, vol. 56, no. 1, pp. 79–97, 2021.

[15] S. Pal, J. Liu, I. Alam, N. Cebry, H. Suhail, S. Bu, S. S. Iyer, S. Pamarti, R. Kumar, and P. Gupta, "Designing a 2048-chiplet, 14336-core waferscale processor," in *Proc. DAC*, 2021, pp. 1183–1188.

[16] A. Elsherbini *et al.*, "Enabling next generation 3D heterogeneous integration architectures on Intel process," in *Proc. IEDM*, 2022, pp. 27.3.1–27.3.4.

[17] B. M. D. Sawyer, Y. Suzuki, R. Furuya, C. Nair, T.-C. Huang, V. Smet, K. Panayappan, V. Sundaram, and R. Tummala, "Design and demonstration of a 2.5-D glass interposer BGA package for high bandwidth and low cost," *IEEE Trans. Compon. Packag. Technol.*, vol. 7, no. 4, pp. 552–562, 2017.

[18] T. Iwai, T. Sakai, D. Mizutani, S. Sakuyama, K. Iida, T. Inaba, H. Fujisaki, A. Tamura, and Y. Miyazawa, "Multilayer glass substrate with high density via structure for all inorganic multi-chip module," in *Proc. ECTC*, 2019, pp. 1952–1957.

[19] Z. Wang *et al.*, "HISIM: Analytical performance modeling and design space exploration of 2.5D/3D integration for AI computing," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, 2025, (early access).

[20] S. Jeloka, B. Cline, S. Das, B. Labbe, A. Rico, R. Herberholz, J. DeLaCruz, R. Mathur, and S. Hung, "System technology co-optimization and design challenges for 3D IC," in *Proc. CICC*, 2022, 6 pages.

[21] P. S. Nalla, E. Haque, Y. Liu, S. S. Sapatnekar, J. Zhang, C. Chakrabarti, and Y. Cao, "CLAIRE: Composable chiplet libraries for AI inference," in *Proc. DATE*, 2025, 6 pages.

[22] C. C. Sudarshan, N. Matkar, S. Vrudhula, S. S. Sapatnekar, and V. A. Chhabria, "ECO-CHIP: Estimation of carbon footprint of chiplet-based architectures for sustainable VLSI," in *Proc. HPCA*, 2024, pp. 671–685.

[23] A. Graening, S. Pal, and P. Gupta, "Chiplets: How small is too small?" in *Proc. DAC*, 2023, 6 pages.

[24] Z. Wang *et al.*, "Exploiting 2.5D/3D heterogeneous integration for AI computing," in *Proc. ASP-DAC*, 2024, pp. 758–764.

[25] R. Mahajan *et al.*, "Co-packaged photonics for high performance computing: Status, challenges and opportunities," *J. of Lightwave Technol.*, vol. 40, no. 2, pp. 379–392, 2022.

[26] K. Radhakrishnan, M. Swaminathan, and B. K. Bhattacharyya, "Power delivery for high-performance microprocessors – Challenges, solutions, and future trends," *IEEE Trans. Compon. Packag. Technol.*, vol. 11, no. 4, pp. 655–671, 2021.

[27] P. R. Prakash, A. Nabih, Y. Liang, S. Kudva, M. Mosa, C. T. Gray, and Q. Li, "A 2400 W/in^3 1.8 V bus converter enabling vertical power delivery for next-generation processors," in *Proc. APEC*, 2024, pp. 910–917.

[28] E. A. Burton, G. Schrom, F. Paillet, J. Douglas, W. J. Lambert, K. Radhakrishnan, and M. J. Hill, "FIVR – fully integrated voltage regulators on 4th generation Intel® Core™ SoCs," in *Proc. APEC*, 2014, pp. 432–439.

[29] Y. Safari and B. Vaisband, "Integrated power delivery methodology for 3D ICs," in *Proc. ISQED*, 2022, pp. 114–119.

[30] H.-P. Le, C. Hardy, H. Pham, M. M. Jatlaoui, F. Voiron, P. Mercier, P.-H. Chen, and S. Jha, "Vertical power delivery and heterogeneous integration for high-performance computing," in *Proc. BCICTS*, 2023, pp. 32–35.

[31] H.-P. Le, J. Crossley, S. R. Sanders, and E. Alon, "A sub-ns response fully integrated battery-connected switched-capacitor voltage regulator delivering 0.19W/mm^2 at 73% efficiency," in *Proc. ISSCC*, 2013, pp. 372–373.

[32] R. Harjani and S. Chaubey, "A unified framework for capacitive series-parallel DC-DC converter design," in *Proc. CICC*, 2014, 8 pages.

[33] T. Yu and M. D. Wong, "Efficient simulation-based optimization of power grid with on-chip voltage regulator," in *Proc. ASP-DAC*, 2014, pp. 531–536.

[34] Z. Zeng, X. Ye, Z. Feng, and P. Li, "Tradeoff analysis and optimization of power delivery networks with on-chip voltage regulation," in *Proc. DAC*, 2010, pp. 831–836.

[35] S. A. Sadat, M. Canbolat, and S. Köse, "Optimal allocation of LDOs and decoupling capacitors within a distributed on-chip power grid," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 23, no. 4, 2018, 15 pages.

[36] S. Krishnakumar and I. Partin-Vaisband, "Vertical power delivery for emerging packaging and integration platforms – power conversion and distribution," in *Proc. SOCC*, 2023, 6 pages.

[37] P. Zhou, A. Paul, C. H. Kim, and S. S. Sapatnekar, "Distributed on-chip switched-capacitor DC–DC converters supporting DVFS in multicore systems," *IEEE Trans. VLSI Syst.*, vol. 22, no. 9, pp. 1954–1967, 2014.

[38] V. A. Chhabria and S. S. Sapatnekar, "OpeNPDN: A neural-network-based framework for power delivery network synthesis," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 41, no. 10, pp. 3515–3528, 2022.

[39] H. Su, S. Sapatnekar, and S. Nassif, "Optimal decoupling capacitor sizing and placement for standard-cell layout designs," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 22, no. 4, pp. 428–436, 2003.

[40] M. Popovich, M. Sotman, A. Kolodny, and E. G. Friedman, "Effective radii of on-chip decoupling capacitors," *IEEE Trans. VLSI Syst.*, vol. 16, no. 7, pp. 894–907, 2008.

[41] Y. Zhan, S. V. Kumar, and S. S. Sapatnekar, "Thermally-aware design," *Found. Trends in Electron. Des. Autom.*, vol. 2, no. 3, pp. 255–370, March 2008.

[42] V. A. Chhabria, V. Ahuja, A. Prabhu, N. Patil, P. Jain, and S. S. Sapatnekar, "Encoder-decoder networks for analyzing thermal and power delivery networks," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 28, no. 1, Dec. 2022, 27 pages.

[43] N. Karmokar, S.-W. Tam, T. V. D. Dinh, V. A. Chhabria, R. Harjani, and S. S. Sapatnekar, "Analyzing the impact of FinFET self-heating on the performance of RF power amplifiers," in *Proc. ICCAD*, 2024, 9 pages.

[44] B. Goplen and S. Sapatnekar, "Thermal via placement in 3D ICs," in *Proc. ISPD*, 2005, pp. 167–174.

[45] D. Lorenzini, C. Green, T. E. Sarvey, X. Zhang, Y. Hu, A. G. Fedorov, M. S. Bakir, and Y. Joshi, "Embedded single phase microfluidic thermal management for non-uniform heating and hotspots using microgaps with variable pin fin clustering," *Int. J. Heat Mass Transfer*, vol. 103, pp. 1359–1370, 2016.

[46] C. H. Hoang, A. Azizi, N. Fallahtafti, S. Rangarajan, V. Radmard, C. Arvin, K. Sikka, S. Schiffres, and B. Sammakia, "Design and thermal analysis of a 3-D printed impingement pin fin cold plate for heterogeneous integration application," *IEEE Trans. Compon. Packag. Technol.*, vol. 12, no. 7, pp. 1091–1099, 2022.

[47] G. Yan *et al.*, "Toward TSV-compatible microfluidic cooling for 3D ICs," *IEEE Trans. Compon. Packag. Technol.*, vol. 15, no. 1, pp. 104–112, 2025.

[48] J. Fang, S. Gupta, S. V. Kumar, S. K. Marella, V. Mishra, P. Zhou, and S. S. Sapatnekar, "Circuit reliability: From physics to architectures: Embedded tutorial paper," in *Proc. ICCAD*, 2012, pp. 243–246.

[49] F. Najm and V. Sukharev, "Electromigration simulation and design considerations for integrated circuit power grids," *J. Vac. Sci. Technol., B*, vol. 38, no. 6, pp. 063 204:1–063 204:12, Nov/Dec 2020.

[50] F. N. Najm, "Equivalent circuits for electromigration," *Microelectron. Rel.*, pp. 114 200:1–114 200:16, Aug. 2021.

[51] M. A. A. Shohel, V. A. Chhabria, N. Evmorfopoulos, and S. S. Sapatnekar, "Analytical modeling of transient electromigration stress based on boundary reflections," *Proc. ICCAD*, 2021.

[52] M. A. Al Shohel, V. A. Chhabria, N. Evmorfopoulos, and S. S. Sapatnekar, "Frequency-domain transient electromigration analysis using circuit theory," in *Proc. ICCAD*, 2023, 8 pages.

[53] V. Mishra and S. S. Sapatnekar, "The impact of electromigration in copper interconnects on power grid integrity," in *Proc. DAC*, 2013, 6 pages.

[54] H. Xiong, Z. Huang, and P. Conway, "Effects of stress and electromigration on microstructural evolution in microbumps of three-dimensional integrated circuits," *IEEE Trans. Device Mater. Reliab.*, vol. 14, no. 4, pp. 995–1004, 2014.

[55] Z. Shen, S. Jing, Y. Heng, Y. Yao, K. N. Tu, and Y. Liu, "Electromigration in three-dimensional integrated circuits," *Appl. Phys. Rev.*, vol. 10, no. 2, pp. 021 309:1–021 309:30, 05 2023.

[56] S. K. Marella and S. S. Sapatnekar, "A holistic analysis of circuit performance variations in 3-D ICs with thermal and TSV-induced stress considerations," *IEEE Trans. VLSI Syst.*, vol. 23, no. 7, pp. 1308–1321, 2015.

[57] A. Kteyan and V. Sukharev, "Physics-based simulation of stress-induced and electromigration-induced voiding and their interactions in on-chip interconnects," *Microelectron. Eng.*, vol. 247, pp. 111 585:1–111 585:7, 2021.

[58] Open Compute Project, "Assembly design kits (ADK), version 1.0," Jan. 2025.

[59] D. Kehlet, "Accelerating innovation through a standard chiplet interface: The advanced interface bus (AIB)," Intel, https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/accelerating-innovation-through-aib-whitepaper.pdf.

[60] R. Farjadrad, M. Kuemerle, and B. Vinnakota, "A bunch-of-wires (BoW) interface for interchiplet communication," *IEEE Micro*, vol. 40, no. 1, pp. 15–24, 2020.

[61] D. Das Sharma, G. Pasdast, Z. Qian, and K. Aygun, "Universal chiplet interconnect express (UCIe): An open industry standard for innovations with chiplets at package level," *IEEE Trans. Compon. Packag. Technol.*, vol. 12, no. 9, pp. 1423–1431, 2022.

[62] J. Chang, "3Dblox: Unleash the ultimate 3DIC design productivity," 2024, p. 215.

[63] "3Dblox – Chiplet connectivity and physical properties description language (ieee p3537)," https://sagroups.ieee.org/3537/.