

Network on Interconnect Fabric

Boris Vaisband, Adeel Bajwa, and Subramanian S. Iyer

Center for Heterogeneous Integration and Performance Scaling (CHIPS)

Department of Electrical and Computer Engineering

University of California, Los Angeles, CA 90095 USA

[vaisband,abajwa,s.s.iyer]@ucla.edu

Abstract—Silicon interconnect fabric (Si-IF) supports integration of bare dies using thermal compression bonding on a Si wafer substrate. Fine pitch (2 to 10 μm) horizontal and vertical interconnects are feasible within the Si-IF using standard Si processing techniques. A network on interconnect fabric (NoIF) is proposed in this paper. The NoIF enables integration of ultra large scale heterogeneous systems within the technologically mature Si-IF platform. NoIF is based on utility dies which serve as intelligent nodes within the network. NoIF enables global communication, power conversion and management, synchronization, processing and memory capabilities, redundancy allocation, and test of the Si-IF, and the utility and functional dies.

I. INTRODUCTION

Modern applications include a variety of heterogeneous circuit blocks. Diverse technologies, substrate and interconnect materials, and processes are required to coexist within a single system [1]. In addition to heterogeneity, ultra large scale integration is necessary for a variety of applications, *e.g.*, neuromorphic systems [2]. Silicon interconnect fabric (Si-IF) is a compatible platform to satisfy the needs of modern systems [3].

On-chip dimensions have been aggressively scaling for several decades. The scaling of packages and printed circuit boards (PCBs) however, has been almost stagnant during this time. Due to the difference in scaling pace, on chip I/Os exhibit a significantly smaller pitch (few μm) than package-level C4 bumps (tens of μm) and PCB-level ball grid arrays (hundreds of μm). Effectively, the package is used to space-out the on-chip interconnect to match the PCB. In addition, the package is large and the distance between packages can be of the order of tens of millimeters. An interposer platform has been used in recent years to alleviate the issues associated with long board-level interconnects, reducing the distance between chips on the interposer to approximately ten millimeters. Si-IF provides an opportunity to overcome both, the distance between chips and the large vertical interconnect pitch issues.

Bare (unpacked) dies are placed directly on the passive Si-IF platform that effectively replaces the PCB [4]. Many dies, each independently fabricated using a dedicated technology and process, are integrated within the Si-IF creating an ultra large scale heterogeneous system within a single platform. The dies are connected using fine pitch vertical interconnects (pillars) designed directly on the Si-IF. Ideally, an entire




Fig. 1. An illustration of a 300 mm diameter Si-IF completely populated with heterogeneous dies.

300 mm diameter Si-IF can be populated with disparate dies of different size and aspect ratio, as illustrated in Figure 1. The pitch of the vertical pillars that are used to bond dies to the Si-IF is in the order of a few μm , and the minimal distance between adjacent dies on the Si-IF is approximately 50 to 100 μm .

To enable the Si-IF as a practical platform for ultra large scale heterogeneous integration, system-level issues, similar to a large system on chip (SoC), must be addressed. A network on interconnect fabric (NoIF) is proposed in this paper to support global communication, power conversion and management, synchronization, and to facilitate testing within the Si-IF. NoIF, the main contribution of this paper, borrows from the network on chip (NoC) concepts [5]. Unlike NoCs that are designed within SoCs, the Si-IF is a passive platform for interconnect and other passive components (*i.e.*, capacitors and inductors) only. Utility dies (UDs) therefore, serve as intelligent nodes within the NoIF. The purpose of this paper is to describe the NoIF and the UD as components within the network.

The rest of the paper is composed of the following sections. The Silicon interconnect fabric technology is reviewed in Section II. An overview of the proposed NoIF is provided in Section III. Architecture, design, and test aspects related to the NoIF approach are discussed in Section IV. Some conclusions are offered in Section V.

II. SILICON INTERCONNECT FABRIC

The key advantage of Si-IF technology is that it is based on mature silicon fabrication processes. The Si-IF is a silicon wafer of typical size (*e.g.*, diameter of 300 mm) with up to four Cu metal layers with a pitch of 2 to 10 μm , fabricated using conventional damascene process. Si-IF wires are terminated with Cu pillars (capped with Au) with height and diameter dimensions of 2 to 5 μm . The Si-IF platform can accommodate heterogeneous dies of various size (edge length of 0.5 to

This research is supported in part by the UCLA CHIPS Consortium, the Defense Advanced Research Projects Agency under Grant No. FA8650-16-1-7648, and the Office of Naval Research under Grant No. N00014-16-1-2639.




Fig. 2. SEM cross section of the fine pitch interconnects. Bond interface is shown in inset.




Fig. 3. Fabricated 100 mm Si-IF with a large number of dies of different sizes (4, 9, 16, and 25 mm²), attached with 10 μm pitch.

5 mm), aspect ratio, substrate and interconnect material, and technology. Each die is bonded to the Cu pillars on the Si-IF using thermal compression bonding (TCB) process, as demonstrated in Figure 2. Packages and PCBs become obsolete when utilizing the Si-IF technology, effectively, the Si-IF replaces the PCB and is directly connected to an interface socket using connectors. A fabricated Si-IF with a large number of dies of different sizes (4, 9, 16, and 25 mm²) is demonstrated in Figure 3. The dies are attached at 10 μm pitch, as compared to 400 μm pitch on a PCB.

To further emphasize the advantages of Si-IF technology over C4 and μ-bumps, a comparison of geometric and electrical properties of these technologies is summarized in Table I. The Cu pillars, utilized by the Si-IF technology, exhibit superior characteristics as compared to the other technologies. The small diameter of the Cu pillars enables a large number of I/Os between the die and the Si-IF, whereas the low effective contact resistance of the pillars alleviates signal degradation issues and reduces power.

In addition to the favorable geometric and electrical properties, the Si-IF technology exhibits improved thermal characteristics. The Si substrate has a significantly lower thermal conductivity as compared to other substrates, such as FR4 [3], making the Si-IF an effective heat spreader. Coefficient of thermal expansion (CTE) mismatch between the Si-IF

TABLE I
COMPARISON OF GEOMETRIC AND ELECTRICAL PROPERTIES OF CU PILLARS (CAPPED WITH AU), C4 AND μ-BUMPS

Interconnect type	Diameter [μm]	Contact pad area [μm ²]	Specific contact resistance [Ω·μm ²]
C4 bump [6]	100	~7,800	78
C4 bump [6]	50	~1,950	48.7
μ-bump [7]	23	~415	19.5
μ-bump [7]	16	~201	8.64
Cu pillar [4]	5	~19.5	0.82

and typical Si chips is also low leading to reduced thermal-mechanical stress.

III. OVERVIEW OF NETWORK ON INTERCONNECT FABRIC

Although the technology that supports ultra large scale heterogeneous integration on Si-IF is available, as described in Section II, the system is expected to suffer from system-level design and test challenges, similar to SoCs. For example, global routing congestion and excessive power dissipation in the interconnects are typical concerns in large scale systems [8].

NoIF borrows architectural concepts from the mature NoC approach. Since the Si-IF is a passive substrate where each component is a die, dedicated utility dies are used as nodes within the NoIF. Each UD includes circuit components to manage all service aspects of the nearby dies, including global communication within the Si-IF. A system-level schematic of a NoIF is depicted in Figure 4. Unlike in NoCs where the location of each routing node is arbitrarily defined (the actual physical placement can be anywhere on the chip), the nodes of the NoIF are dies that are placed in the appropriate physical locations using the same TCB process as any other die within the system. UDs are connected using wide lines that serve as routing paths for global communication of data and other management signals.

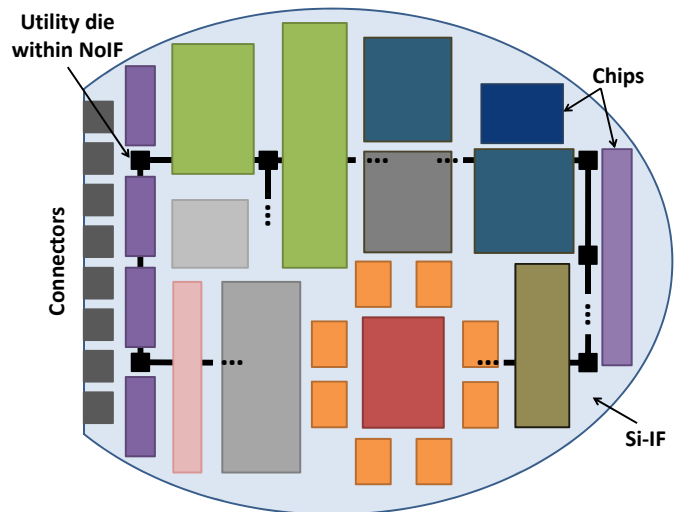


Fig. 4. System-level schematic of the NoIF including utility dies, serving as nodes within the NoIF.

NoIF supports a wide variety of services due to the unique technology of the Si-IF. Unlike NoCs, where the nodes are typically dedicated to global routing, UD within the NoIF are small dies (dielets in the order of 1 mm^2) that include intelligent components to perform multiple tasks (a detailed discussion of the tasks is provided in Section IV):

- Global routing - similar to NoC, NoIF supports global communication between remotely placed dies.
- Power conversion, regulation, and management - each UD contains power conversion and regulation circuits, for example a buck converter and linear drop out (LDO) regulators. In addition to conversion and regulation, UD support power management across the entire Si-IF [9].
- Synchronization - NoIF supports clock management through circuits for multi-clock domains within the UD.
- Processing capability - a processing core controls the main functions of each UD.
- Memory and queuing - embedded memory within each UD enables queue based communication and global signaling, as well as supports local memory requirements of the processing core.
- Test - NoIF enables unique testing opportunities for the Si-IF, UD, and functional dies (dies that are part of the functionality intended for the designed system).
- Redundancy allocation - UD contain additional circuit components to be used in case of failure of adjacent dies, increasing the overall reliability of the system.

To ensure feasibility and facilitate reuse of the NoIF, UD must be identical general purpose dies. The size, number, and floorplan of UD within the Si-IF platform are critical parameters that determine the usefulness of the NoIF versus the overhead that the network incurs. These parameters have significant effect on the performance of the network for all of the aforementioned tasks. A conceptual block-level schematic of a utility die is shown in Figure 5. Each UD is connected by global links to adjacent UD, and by local interconnects

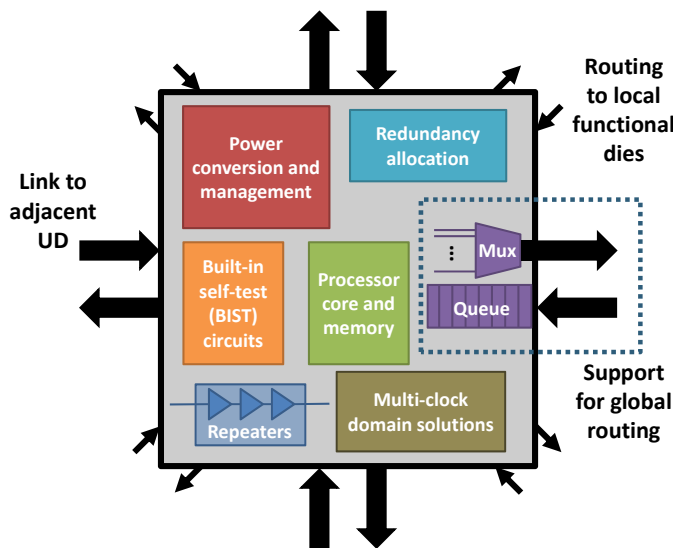


Fig. 5. Block-level schematic of an intelligent utility die within the NoIF.

to surrounding functional dies. All incoming data into the UD (global and local) is queued to enable intelligent priority based communication. The order of all data at the output of the UD (global and local) is determined according to priorities and routed using multiplexers (Muxs). Each UD has a minimum area limitation of $\sim 0.25 \text{ mm}^2$ defined by the ability of the pick-and-place tool to pick and bond dies using TCB. UD are, therefore, able to support the components described in Figure 5.

IV. ARCHITECTURE, DESIGN, AND TEST OF NOIF

A detailed discussion of the architecture, design, and test methodologies is provided in this section. The concepts of communication within the NoIF are similar to NoC, however, the NoIF approach requires significant expansion on the other system-level issues.

A. NoIF Architecture

The basic architecture of NoIF is based on UD and the wide global links that enable communication between adjacent UD. Existing NoC communication algorithms and architectures [10], [11] are exploited to enhance the basic architecture. Each die within the Si-IF that needs to communicate with another remotely placed die (not nearest neighbor) will utilize the NoIF. The minimal communication distance that will ensure that the utilization of NoIF is useful, depends on the delay and power associated with the communication of the signal. However, global routing congestion may require utilization of the NoIF even when the benefit is not obvious. According to the SuperCHIPS protocol developed for communication between dies within the Si-IF [12], significant benefits are exhibited for Si-IF links of length $\leq 500 \mu\text{m}$, as compared to PCB and interposer technologies. Although, these links are useful to connect dies to nearest neighbors since dies can be placed as close as 50 to $100 \mu\text{m}$ within the Si-IF, farther communication requires utilization of the NoIF.

B. Global and Semi-Global Routing

Large links are utilized to communicate between adjacent UD. Incoming data will be placed in input queues and distributed to other UD according to assigned priority. The priority of incoming data is determined by the processing core according to arrival time, urgency, and relevant link availability. In some cases, urgent data can be directed to a physically longer path that will provide a shorter temporal path. The management of the global routing is executed by the local processing cores within each UD. The processing cores can communicate over the network to inquire about link availability and use this information to generate the shortest physical or temporal paths. A classical tradeoff between speed and power is considered by the processing core while generating the optimal path for the data.

In addition to communication with other UD, repeater circuits are available within each UD to support semi-global communication between dies that are not nearest neighbors. If a UD is located between two dies and is identified as the

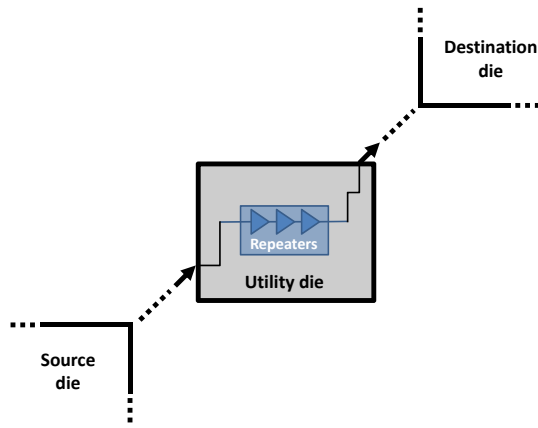


Fig. 6. Schematic of semi-global communication utilizing repeaters within a utility die.

local NoIF node for both dies, the communication between the dies will be according the SuperCHIPS protocol. The communication signal is then routed through the repeaters that are available within the UD to enhance the signal. A schematic of semi-global communication is depicted in Figure 6.

Si-IF supports heterogeneous integration, therefore compatible with optical interconnect (both waveguide-based and free-space). The power, area, and speed tradeoffs between electrical and optical communication should be considered when deciding to utilize optical interconnects. It is likely that optical interconnects are beneficial for global communication across the entire Si-IF (> 100 mm).

C. Power Conversion and Management

Since each UD is an integrated circuit (IC), power conversion circuits can be directly designed on the UD. The power to the Si-IF is expected to be supplied via socket connectors (shown in Figure 4). The power is distributed across the entire wafer via high voltage power planes (to minimize IR power losses) and then locally converted by power circuits within the distributed UDs. Power conversion within each UD is performed according to the voltage requirements of adjacent dies (local voltage domains). The power is regulated using LDO circuits to ensure high quality of power.

Delivering high quality power to all devices within an ultra large scale system is a significant challenge. Moreover, heterogeneous integration imposes additional constraints on the delivered power - a variety of voltage domains, supply currents, and quality of power demands. To satisfy the power constraints of an ultra large scale heterogeneous system within Si-IF, a power management system is necessary.

Existing NoC power management concepts [9] are adopted for the NoIF. Each node communicates unique power requirements to other nodes and the NoIF dynamically adapts the power conversion and delivery schemes to ensure high quality of power for each device, stability of the power system [13], and security.

Power gating is also managed by the NoIF, as part of the power management scheme. Each UD is locally responsible

to power-gate adjacent dies to conserve power when possible, by utilizing power gating circuitry within the UD. Decoupling capacitors (decaps) are another important power component within UDs that can be fabricated using Si processes such as deep trench technology. Multiple size decaps ensure that the Ldi/dt noise is within specifications for various frequencies of operation.

Power management constraints will drive system level floor-planning and Si-IF resource allocation. A schematic of a UD within NoIF surrounded by multiple voltage domains is shown in Figure 7.

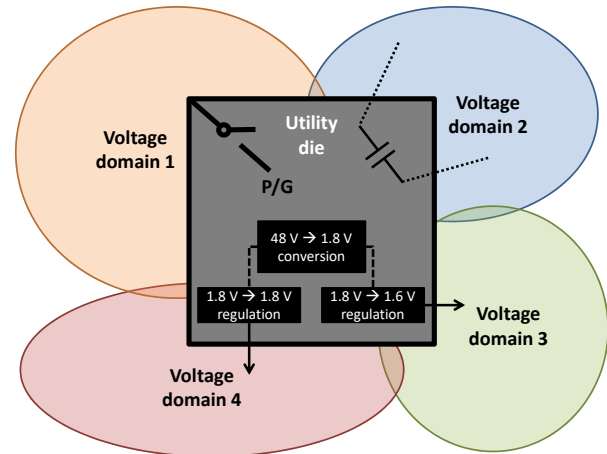


Fig. 7. Schematic of a UD surrounded by multi-voltage domains. The utility die dynamically manages the conversion, regulation, and gating of power delivered to the adjacent dies. Decoupling capacitors are also available for passive frequency dependent power regulation.

D. Synchronization

Clock distribution is a fundamental design challenge in synchronous integrated circuits. Various synchronization methodologies have been developed for SoCs and board level systems that can be adopted by the NoIF. For ultra large scale systems, a globally asynchronous locally synchronous (GALS) approach is typically used [8], [14]. In this methodology, the communication between dies within the Si-IF is not clocked (*e.g.*, using hand shake protocol), alternatively, local on-chip circuits are clocked. Adjacent dies within Si-IF can also communicate synchronously using the SuperCHIPS communication protocol. GALS approach is especially beneficial for heterogeneous integration where different dies can utilize various clocks (*i.e.*, multi-clock domains).

The proposed NoIF supports GALS by enabling asynchronous communication between UDs and local clock distribution, managed by clock circuits on each UD. Similarly to power gating, clock gating is also supported by the nodes within the NoIF.

E. Processing Capability

The processing cores provide the intelligence to the UDs. These controllers are responsible for the functionality of the UD in terms of communication, power and clock management, and testing. Each processing core manages the traffic of data

between local dies and adjacent UD's by controlling the queues and determining the priorities of all incoming signals. Bypass routing is available for high priority data packets.

NoIF is effectively a many-core system that manages the communication within the Si-IF. In addition to communication, the processing cores support test of the Si-IF, UD's, and functional dies, as described in Subsection IV-H.

F. Memory and Queuing

Queues are required to support access to a mutual resources - links between UD's. The incoming data from local dies is queued and arranged according to priorities determined by the processing core. Access to global links is managed by multiplexers, controlled by the processing core. Standard queue management architecture can be exploited.

A small embedded memory is required to support the processing core and the NoIF. Network parameters and data is stored in this high performance memory.

G. Floorplanning

Since UD's are actual dies placed within the Si-IF, it is necessary to optimize the physical location of each UD. The floorplan optimization problem within the Si-IF is similar to floorplanning SoCs. A simulated annealing engine can be used to iteratively converge to an optimal solution considering a weight function. In addition to standard floorplanning parameters such area and interconnect length, the floorplanning weight function of NoIF consists of parameters that are unique to the Si-IF technology. The parameters that should be considered within the NoIF floorplanning weight function are:

- Number of UD's - a key parameter within the NoIF is the number of utility dies. Each utility die is in fact an overhead, since the main function of the system is within the functional dies, however it is a necessary overhead. The tradeoff between the overhead that each UD incurs and the benefit that it provides is quantified in terms of the area of each UD versus the area (or number of functional dies) it can effectively support.
- Voltage and clock domains - the number of voltage and clock domains will dictate both the number and placement of UD's. It is preferable to group functional dies with similar voltage or clock requirements to simplify the power/clock management problem for the UD.
- Heterogeneity of the system within the Si-IF - this parameter can significantly affect the floorplan of the NoIF. A homogeneous system with regular dies will require a similarly regular NoIF. Alternatively, a heterogeneous system consistent of disparate size and aspect ratio dies, will require a unique NoIF floorplan. A hybrid NoIF is also possible, where part of the system is homogeneous and the other part is heterogeneous. The floorplan of the UD's will exhibit similar properties.
- Area - since the area and aspect ratio of UD's is identical, the area of the entire floorplan (including functional dies) should be considered while iterating the placement of the

UD's. The locations of the UD's will affect placement of functional dies.

- Interconnect length - two types of interconnects should be considered (1) global links between UD's, and (2) local interconnects between functional dies and closest UD. Note functional dies can be assigned (connected) to multiple UD's and the processing cores will determine the optimal routing path in real time.
- Heat dissipation - this issue is significantly alleviated by the Si-IF technology as compared to, for example, three-dimensional (3-D) IC's [1], due to the high thermal conductance of the Si substrate. The Si-IF spreads the heat effectively across the entire wafer. Nonetheless, additional thermal aware floorplanning techniques can be required to separate thermally aggressive from thermally sensitive dies [15].

The NoIF can exhibit an irregular floorplan based on the specific requirements of the system within the Si-IF. An example NoIF floorplan adapted for heterogeneous and homogeneous systems within a single Si-IF is shown in Figure 8.




Fig. 8. An example NoIF floorplan compatible with both homogeneous and heterogeneous systems that coexist within a single Si-IF platform.

H. Test

Although Si-IF employs mature Si processing technologies, failure is bound to happen either during fabrication of the Si-IF or during the TCB process. In addition, the functional bare dies can also exhibit failure during fabrication and handling. Testing of the overall system and individual components within the system is necessary to minimize failure and maximize yield.

Built-in self-test (BIST) algorithms and compatible circuits are designed into the UD's (as depicted in Figure 5). The BIST approach enables a system to perform self testing according to predefined inputs and expected outputs. Similar to other previously mentioned topics, NoIF can borrow from previous NoC-related work [16].

The NoIF supports testing on both the system and component level. After placing the UD's on the Si-IF, communication signals are used to test the bonding (pillars), interconnect fabric, and the UD itself. The network will effectively gather information regarding itself and the Si-IF platform on which it is designed. UD failures are recorded by all other UD's to update the constraints of the routing algorithms. In addition, defective interconnect fabric is identified and bypass routes

are determined. Once the Si-IF and UDs have been tested, the functional dies are placed. Note, the test results of the Si-IF and UDs can affect the placement of functional dies. For example, if a certain UD is defective, functional dies will be preferably placed closer to other functional UDs. After the functional dies are bonded, BIST algorithms are executed to identify defects within the functional dies and RA is performed.

I. Redundancy Allocation

The BIST engine (described in IV-H) is capable to identify failures, however, redundancy allocation (RA) is required to correct for failures. Together, BIST and RA provide a built-in self-test and repair (BISTR) capability. Two levels of RA are supported by the NoIF. At the circuit level, additional circuits are available within the UD to accommodate circuit failure. These additional circuits can be used as backup for circuits on the UD or on functional dies. However bypassing local circuitry on a functional die by communicating with a UD should only be done in extreme cases. At the system level, The UDs can store information regarding the functionality of other UDs and functional dies in embedded nonvolatile memory, such as OTPM/MTPM [17]. A certain die can fail either during processing (fabrication and bonding) or during operation. In either case, failures are detected by the UD and the information is stored to adapt routing algorithms and other functionality of the UD, such as cut off power to the malfunctioning die to save power.

RA is a key capability within the Si-IF since rework of dies, *i.e.*, replacement of bonded dies after identified failure, is not possible in this technology. Rigorous testing should, therefore, be performed prior to bonding dies to the Si-IF to avoid significant reduction in performance of the overall system, as described in the following subsection.

V. CONCLUSIONS

Feasibility of the Si-IF substrate technology has been proven and the next step is to enable the Si-IF platform for integration of ultra large scale heterogeneous systems. An NoIF is proposed in this paper. The network exhibits similarities to the mature NoC methodology, but also requires a unique approach to ensure high performance of the integrated systems.

The NoIF is based on intelligent utility dies that incorporate active circuits to provide solutions for a wide variety of architecture, design, and test aspects. The UDs are robust in terms of the services they are able to support within the NoIF. The network will enable basic global communication requirements, as well as power and clock management, processing and embedded memory, redundancy allocation, and BIST capabilities.

Additional research is necessary for the NoIF architecture to mature. Nonetheless, the NoIF approach is expected to become feasible in the near future (at least in a basic form) due to a variety of important heterogeneous and ultra large scale applications that significantly benefit from the Si-IF technology, or are not otherwise possible.

REFERENCES

- [1] Boris Vaisband, *3-D ICs as a Platform for Heterogeneous Systems Integration*, Ph.D. Thesis, 2017.
- [2] Z. Wan and S. S. Iyer, "Three-Dimensional Wafer Scale Integration for Ultra Large Scale Cognitive Systems," *Proceedings of the IEEE SOI-3D-Subthreshold Microelectronics Technology Unified Conference*, October 2017.
- [3] S. S. Iyer, "Heterogeneous Integration for Performance and Scaling," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, Vol. 6, No. 7, pp. 973 – 982, July 2016.
- [4] A. A. Bajwa, S. Jangam, S. Pal, N. Marathe, T. Bai, T. Fukushima, M. Goorsky, and S. S. Iyer, "Heterogeneous Integration at Fine Pitch ($\leq 10 \mu\text{m}$) Using Thermal Compression Bonding," *Proceedings of the IEEE International Electronic Components and Technology Conference*, pp. 1276 – 1284, May 2017.
- [5] S. Kumar, A. Jantsch, J.-P. Soininen, M. Forsell, M. Millberg, J. Öberg, K. Tiensyrjä, and A. Hemani, "A Network on Chip Architecture and Design Methodology," *Proceedings of the IEEE Computer Society Annual Symposium on VLSI*, pp. 117 – 124, April 2002.
- [6] S. L. Wright, R. Polastre, H. Gan, L. P. Buchwalter, R. Horton, P. S. Andry, E. Sprogis, C. Patel, C. Tsang, J. Knickerbocker, J. R. Lloyd, A. Sharma, and M. S. Sri-Jayantha, "Characterization of Micro-Bump C4 Interconnects for Si-Carrier SOP Applications," *Proceedings of the IEEE International Electronic Components and Technology Conference*, pp. 633 – 640, June 2006.
- [7] B. Dang, S. L. Wright, P. S. Andry, C. K. Tsang, C. Patel, R. Polastre, R. Horton, K. Sakuma, B. C. Webb, E. Sprogis, G. Zhang, A. Sharma, and J. U. Knickerbocker, "Assembly, Characterization, and Reworkability of Pb-Free Ultra-Fine Pitch C4s for System-on-Package," *Proceedings of the IEEE International Electronic Components and Technology Conference*, pp. 42 – 48, May 2007.
- [8] E. Salman and E. G. Friedman, *High Performance Integrated Circuit Design*, McGraw-Hill Publishers, 2012.
- [9] I. P. Vaisband, R. Jakushokas, M. Popovich, A. V. Mezhiba, S. Kose, and E. G. Friedman, *On-Chip Power Delivery and Management, 4th Edition*, Springer, 2016.
- [10] J. Liu, J. Harkin, L. P. Maguire, L. J. McDaid, J. J. Wade, and G. Martin, "Scalable Networks-on-Chip Interconnected Architecture for Astrocyte-Neuron Networks," *IEEE Transactions on Circuits and Systems I: Regular Papers*, Vol. 63, No. 12, pp. 2290 – 2303, December 2016.
- [11] P. Vivet, Y. Thonnart, R. Lemaire, C. Santos, E. Beigné, C. Bernard, F. Darve, D. Lattard, I. Miro-Panadès, D. Dutoit, F. Clermidy, S. Chermamy, A. Sheibanyrad, F. Pétrot, E. Flamand, J. Michailos, A. Arriordaz, L. Wang, and J. Schloeffel, "A $4 \times 4 \times 2$ Homogeneous Scalable 3D Network-on-Chip Circuit With 326 MFlit/s 0.66 pJ/b Robust and Fault Tolerant Asynchronous 3D Links," *IEEE Journal of Solid-State Circuits*, Vol. 52, No. 1, pp. 33 – 49, January 2017.
- [12] S. Jangam, S. Pal, A. Bajwa, S. Pamarti, P. Gupta, and S. S. Iyer, "Latency, Bandwidth and Power Benefits of the SuperCHIPS Integration Scheme," *Proceedings of the IEEE International Electronic Components and Technology Conference*, pp. 86 – 94, May 2017.
- [13] I. Vaisband and E. G. Friedman, "Stability of Distributed Power Delivery Systems With Multiple Parallel On-Chip LDO Regulators," *IEEE Transactions on Power Electronics*, Vol. 31, No. 8, pp. 5625 – 5634, August 2016.
- [14] E. Kasapaki, M. Schoeberl, R. B. Sørensen, C. Müller, K. Goossens, and J. Sparsø, "Argo: A Real-Time Network-on-Chip Architecture With an Efficient GALS Implementation," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 24, No. 2, pp. 479 – 492, February 2016.
- [15] B. Vaisband and E. G. Friedman, "3-D Floorplanning Algorithm to Minimize Thermal Interactions," *Proceedings of the IEEE International Symposium on Circuits and Systems*, pp. 2133 – 2136, May 2015.
- [16] C. Grecu, A. Ivanov, R. Saleh, and P. P. Pande, "Testing Network-on-Chip Communication Fabrics," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 26, No. 12, pp. 2201 – 2214, December.
- [17] J. Viraraghavan, D. Leu, B. Jayaraman, A. Cestero, R. Kilker, Ming Yin, J. Golz, R. R. Tummuru, R. Raghavan, D. Moy, T. Kempanna, F. Khan, T. Kirihata, and S. Iyer, "80Kb 10ns Read Cycle Logic Embedded High-K Charge Trap Multi-Time-Programmable Memory Scalable to 14nm FIN with No Added Process Complexity," *Proceedings of the IEEE International Symposium on VLSI Circuits*, pp. 1 – 2, June 2016.