Global and Semi-Global Communication on Si-IF

Special Session Paper

Boris Vaisband*

Department of Electrical and Computer Engineering
University of California, Los Angeles
Los Angeles, CA, USA
vaisband@ucla.edu

Subramanian S. Iyer

Department of Electrical and Computer Engineering
University of California, Los Angeles
Los Angeles, CA, USA
s.s.iyer@ucla.edu

ABSTRACT

On-chip scaling continues to pose significant technological and design challenges. Nonetheless, the key obstacle in on-chip scaling is the high fabrication cost of the state-of-the-art technology nodes. An opportunity exists however, to continue scaling at the system level. Silicon interconnect fabric (Si-IF) is a platform that aims to replace both the package and printed circuit board to enable heterogeneous integration and high inter-chip performance. Bare dies are attached directly to the Si-IF at fine vertical interconnect pitch (2 to 10 μm) and small inter-die spacing ($\leq 100 \mu\text{m}$). The Si-IF is a single-hierarchy integration construct that supports dies of any process, technology, and dimensions. In addition to development of the fabrication and integration processes, system-level challenges need to be addressed to enable integration of heterogeneous systems on the Si-IF. Communication is a fundamental challenge on large Si-IF platforms (up to 300 mm diameter wafers). Different technological and design approaches for global and semi-global communication are discussed in this paper. The area overhead associated with global communication on the Si-IF is determined.

CCS CONCEPTS

- **Hardware** $\rightarrow$ Buses and high-speed links; Metallic interconnect; Die and wafer stacking; Package-level interconnect;
- **Networks** $\rightarrow$ Routers.

KEYWORDS

silicon interconnect fabric, communication, heterogeneous integration, global interconnect

ACM Reference Format:

Boris Vaisband and Subramanian S. Iyer. 2019. Global and Semi-Global Communication on Si-IF *Special Session Paper*. In *International Symposium on Networks-on-Chip (NOCS '19)*, October 17–18, 2019, New York, NY, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3313231.3352379>

*The author is currently with McGill University, 3480 Rue University, 547 McConnell Engineering Building, Montreal, QC, Canada H3A 0E9. boris.vaisband@mcgill.ca

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

NOCS '19, October 17–18, 2019, New York, NY, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6700-4/19/10...\$15.00

<https://doi.org/10.1145/3313231.3352379>

1 INTRODUCTION

On-chip features have scaled approximately 1,000X in the past fifty years, whereas the package features have only scaled approximately 4X within the same period of time, as depicted in Figure 1. Nonetheless, further on-chip scaling is associated with significant challenges, most importantly, high cost per transistor. The high cost of on-chip scaling, combined with the need for heterogeneous integration and increased system performance (higher bandwidth and lower energy) [8, 18] are the main drivers for the "More than Moore" approach.

Figure 1: Scaling trends of on-chip and package features.

The silicon interconnect fabric (Si-IF) is a platform that enables scaling off-chip features and is effectively a system-on-wafer (SoW) exhibiting the performance of a system-on-chip (SoC). In the Si-IF approach, bare dies are integrated using thermal compression bonding (TCB) [7] on a Si wafer, eliminating the need for both package and printed circuit board (PCB). The dies are attached using fine pitch (2 to 10 μm) vertical pillars at high proximity ($< 100 \mu\text{m}$) [2]. Since no solder is used with TCB, no intermetallics are present at the interface between the Si-IF pillars and the die pads, significantly reducing the interface resistance [10]. An example of over a hundred indium phosphide dies integrated on Si-IF is shown in Figure 2 [16]. Similarly, any type of die can be integrated on the Si-IF, enabling a true heterogeneous system [8, 18]. The Si-IF fabrication process borrows heavily from the highly mature Si technology at relatively large features (minimum of 2 μm), leading to low cost of fabrication. In addition, the Si-IF is a passive platform (no active devices), allowing to maintain a low fabrication cost.

The Si-IF alleviates bottleneck challenges in modern integrated circuits [13]. Specifically, the number of inputs/outputs (I/Os) on a chip, governed by Rent's rule [12], can be increased by up to

Figure 2: Over a hundred indium phosphide dies integrated on the Si-IF [16]. The dies are of various size and aspect ratio.

two orders of magnitude if integrated at fine pitch on the Si-IF. For example, the aggregate bandwidth of the IBM POWER9 chip is 1,206 GB/s [5] with 2,359 C4 pads dedicated to differential signaling [6], resulting in a data rate per pad of 8.2 Gb/s. Whereas, if integrated using a fine pitch of 3.5 μm , the number of pads dedicated to signaling could be increased by a factor of ~ 35 (assuming a C4 pad pitch on the package of 125 μm). Moreover, integration of an IBM POWER9 on Si-IF will result in significant reduction in communication energy.

Since the Si-IF is agnostic to the type of dies that are integrated on the wafer, diverse technologies, materials, and functionalities within a single system are supported by the platform. The Si-IF is, therefore, compatible with a wide range of applications. For example, intelligent IoT edge devices that are expected to support both sensing and actuation, energy harvesting and storage, processing, memory, and communication [20]. Moreover, traditionally homogeneous applications, such as CPUs, are also shifting towards heterogeneity to alleviate the bottlenecks of classical von-Neumann architectures. Various emerging technologies, *e.g.*, memristors, that require integration of disparate materials are explored for those applications. CPU-memory or GPU-memory systems are also fundamentally heterogeneous. Ideally each of the components (processor and memory) is fabricated using an optimal technology node and process for the function that the component is expected to perform. All of these applications can benefit from integration on the Si-IF platform.

A wide variety of system-level challenges need to be addressed to enable the full potential of the Si-IF platform. Specifically, power delivery and thermal management [1, 15], testing, communication [19], etc. In this paper, global and semi-global communication on the Si-IF is discussed.

The rest of the paper is composed of the following sections. The Si-IF technology is reviewed in Section 2. Global communication aspects including a comparison to classical package-PCB constructs are provided in Section 3. Semi-global communication and the differentiation between global and semi-global communication is discussed in Section 4. A scheme for external, *i.e.*, off-Si-IF, communication is presented in Section 5, followed by some conclusions in Section 6.

2 SI-IF TECHNOLOGY

The Si-IF technology is relatively mature [17]. A comparison is drawn in this section between the Si-IF platform and classical packaging constructs (package and PCB). The advantages and compatibility of the Si-IF as a platform for a large space of applications are discussed.

The Si-IF is a single hierarchy integration scheme, where dies are attached directly to a Si wafer that serves as both the package and the PCB. The Si-IF includes wiring levels that are similar to on-chip top-level metals (fat wires) that support seamless on-chip-like communication with neighboring dies. Unlike the Si-IF, the classical package-PCB structures are meant mostly to fan out the pitch of the vertical interconnect from the fine on-chip pitch (2 to 10 μm) to the large PCB-level pitch (500 to 1,000 μm). Signals then traverse on the PCB towards the neighboring packaged chip and follow a reverse, fan-in, path to reduce the vertical interconnect pitch to the on-chip pitch of the second chip. Packages serve as an intermediate stage (between the chip and the PCB) with a vertical interconnect pitch of (100 to 200 μm). The pitch on the PCB cannot be significantly reduced due to material limitation (FR4 is relatively flexible and therefore prohibits fabrication of fine features). The difference between the Si-IF and a package-PCB structure is depicted in Figure 3.

Interposers have been used in recent years to bring dies closer together, however, interposers add an additional layer of hierarchy to classical packaging, making the structure further complex and significantly increasing the cost. In addition, interposers can only accommodate a small number of dies and the interconnect pitch is still about an order of magnitude larger than the pitch on the Si-IF.

The Si-IF addresses additional key challenges of classical packaging, for example, heat dissipation. In the package-PCB structure, the thermal path from the die towards the PCB is blocked (high thermal resistance) and a heat sink is typically placed on top of the die. Alternatively the Si-IF is a Si wafer that exhibits low thermal resistance and acts as a heat spreader. Moreover, heat sinking can be applied to both sides of the structure, on top of the dies and on the back of the Si-IF [1, 15].

A key advantage of the Si-IF, as mentioned in Section 1, is the large number of I/Os the platform supports. Rent's rule dictates that the number of signal terminals T is related to the number of internal chip components g (*e.g.*, gates, blocks) as follows

$$T = t \cdot g^p, \quad (1)$$

where t and p are constants that represent the technology and circuit complexity. Typically, the exponent p is bounded $0 < p < 1$ and approaches unity as the complexity of the chip decreases. In microprocessors, for example, the values of t and p are, respectively, 0.8 and 0.45 [4]. Current packaging technologies cannot support the number of terminals that Rent's rule dictates, therefore the utilization of serializer/deserializer (SerDes) circuits – a "bypass" of Rent's rule. To eliminate the use of power hungry SerDes, a fine I/O pitch is required similar to the pitch that the Si-IF supports. From (1), across various technology nodes (from 45 to 14 nm), the required vertical I/O pitch to enable elimination of SerDes, is approximately 3 to 7 μm . This range corresponds to the Si-IF vertical interconnect pitch that supports a massive number of parallel short interconnects to communicate among neighboring dies.

Figure 3: A comparison between the Si-IF and a classical package-PCB construct [3, 8, 11].

3 GLOBAL COMMUNICATION ON SI-IF

Communicating to neighboring dies on the Si-IF (up to 500 μm) is a solved problem based on the approach described in Section 2 (also known as the SuperCHIPS protocol [10, 11]). Global communication however, requires additional attention.

Since the Si-IF is an SoC-like SoW, the global communication approach of SoCs, *i.e.*, networks-on-chip (NoCs) can be adapted for the wafer-scale integration. The network on interconnect fabric (NoIF) [19] borrows conceptually from NoCs. The NoIF however, performs additional services to support the system that is integrated on the Si-IF, including power management, synchronization, testing, and more [19]. Similarly to NoCs, the NoIF includes routing nodes, however, the nodes within the NoIF are dies, similar to the functional dies (FDs) that are integrated on the wafer. The node dies are called utility dies (UDs). UD include various circuits to support the services that the NoIF provides to the system. A schematic of the NoIF is depicted in figure 4.

The Si-IF facilitates IP reuse since the integration on the platform is backwards compatible to larger pitches and also solder-based integration. IP reuse is a major advantage of the Si-IF since a large portion of modern integrated circuits (ICs) is unnecessarily redesigned at advanced technology nodes. For example, SerDes circuits can provide high performance at older technology nodes (65 and 45 nm), however, since they are located within modern chips that do not support heterogeneous fabrication or integration, the SerDes circuitry must be redesigned. The Si-IF supports integrating off-the-shelf IPs such as SerDes components without the need to redesign them. This advantage can be utilized for the NoIF to support

global communication on the wafer. μ -SerDes IPs can be integrated as a sub-component of a hierarchical UD (HUD), next to a UD control die (also a sub-component of the HUD). Additional examples of sub-component dies within the HUD are antenna die for wireless communication between Si-IFs and clock distribution across the Si-IF platform, and linear dropout (LDO) IP for power regulation. The HUD concept is illustrated in Figure 5. Communication among the sub-components of the HUD is realized using SuperCHIPS. In addition to μ -SerDes circuits, global communication across the Si-IF can be realized using optical interconnect. This approach however, requires development of photonic waveguides within the Si-IF platform as well as conversion circuits (electrical-photonic-electrical). Existing communication protocols and all NoC-related routing algorithms can be utilized for the NoIF to ensure high performance of global communication.

Similarly to NoCs, NoIF is a necessary overhead for the system integrated on the Si-IF. The area overhead that the NoIF imposes is evaluated based on the parameters listed in Table 1 and plotted as a function of the number of UD and the area of each UD in Figure 6. A typical minimal overhead of NoCs for SoCs at 32 nm is $\sim 12\%$ [14], marked for reference in Figure 6. Note that the columns for 4,096 UD at 16 and 25 mm^2 are missing since the values are above 100%. The NoIF exhibits a wide range of area overhead depending on the number and size of UD. Nonetheless, majority of the UD number-area combinations exhibit an area overhead that is smaller than typical NoCs.

Figure 4: A NoIF based on UD integrated alongside FDs on the Si-IF.

Figure 5: Hierarchical utility die consisting of sub-component dies integrated alongside functional dies.

Table 1: Si-IF and UD parameters used to determine the area overhead of the NoIF.

Parameter	Value
Si-IF diameter (mm)	300
Si-IF area (mm ²)	~70,686
Effective area percentage (%)	90
Effective area (mm ²)	~63,617
Number of UDs	{4, 16, 64, 128, 256, 1,024, 4,096}
Area of each UD (mm ²)	{1, 4, 9, 16, 25}

Figure 6: Area overhead of NoIF as a function of the number of UDs and the area of each UD. Note, large UDs refer to HUDs.

4 SEMI-GLOBAL COMMUNICATION ON SI-IF

Utilizing the NoIF incurs a certain cost. Data needs to be packetized and routed within the network until it reaches the destination FD. For communicating with distant FDs, it is the only reasonable approach. However, at shorter distance, the cost of propagating data within the NoIF may be higher than direct wire communication. Semi-global communication targets a range of destination FDs that are located outside of the range of the SuperCHIPS approach (effective up to 500 μm), but also placed too close to justify the cost of utilizing the NoIF. Semi-global communication does not load the NoIF. The communication range that is the cutoff point between global and semi-global communication in terms of energy and delay, is found to be approximately 10 mm [21]. Since an optimal die area for integration on the Si-IF is from 1 to 100 mm² [9], this 10 mm distance may span one to ten dies. Repeaters, placed on both FDs and UDs, are required to propagate signals across 10 mm.

The design space of the NoIF shown in Figure 6 can be narrowed down by addressing only the relevant number and size of UDs. The distance between neighboring UDs as a function of the number of UDs and the area of each UD is shown in Figure 7 (evaluation is based on the parameters from Table 1. Note that the columns for 4,096 UDs at 16 and 25 mm² are missing since the values are smaller than zero, i.e., overlapping UDs).

A zoom-in on the number and size of UDs that result in a distance of 10 mm between UDs is shown in Figure 8. For clarity, the UD area axis in Figure 8 is reversed (as compared to Figures 6 and 7). Figure 8 defines the right number and size of UDs (around the 10 mm cutoff

Figure 7: Distance between neighboring UDs as a function of the number of UDs and the area of each UD. Note, large UDs refer to HUDs. **Figure 8: Zoom-in on the distance between neighboring UDs (10 mm) that is the cutoff point between global and semi-global communication, as a function of the number of UDs and the area of each UD. For clarity, the UD area axis in Figure 8 is reversed (as compared to Figures 6 and 7).**

point) that are required to ensure efficient communication on the Si-IF.

It can be concluded from Figure 8 that, depending on the size of each UD, the number of UDs on the Si-IF should be between 300 and 500. According to Figure 6, the area overhead associated with this number of UDs, depending on the size of each UD, is approximately between 0.5 and 19.5%.

5 EXTERNAL COMMUNICATION

In addition to communication on the Si-IF, a scheme is required to communicate between the system on the IF and external components, such as various interface I/Os, panels, additional Si-IFs, etc. To support external communication for the Si-IF, an FR4 ring connector along the periphery of the Si wafer is developed, as demonstrated in Figure 9. Pads are fabricated on the ring connector and compatible pads are fabricated on the periphery of the Si-IF. The ring connector attaches to the pads on the Si-IF using solder paste and enables communication to/from the integrated system. Any circuitry required for the realization of various communication protocols can be placed on the ring connector board and signals

Figure 9: Ring connector attached to an Si-IF wafer for external communication.

are routed on the FR4 as needed (e.g., into a plug-and-play socket connector).

The Si-IF side in Figure 9 includes 144 peripheral pads, where each pad is $7 \times 1 \text{ mm}^2$. The FR4 side includes complimentary pads and two levels of traces to accommodate any required wiring. The board side also includes peripheral screws for attachment to PowerTherm [1, 9] (a high power delivery component of the Si-IF that doubles as a heat removal system), that is necessary for systems that are expected to operate at extremely high power.

6 CONCLUSIONS

The Si-IF is a platform for integration of heterogeneous systems that eliminates the need for packaging and PCBs. Dies are integrated on the Si-IF at fine pitch of the vertical interconnect (2 to 10 μm) and small inter-die spacing ($<100 \mu\text{m}$). Integration on the Si-IF enables a massive number of I/O/s to be utilized for communication, thus significantly reducing the energy and increasing the bandwidth of communication. Local (up to 500 μm) communication is realized using the SuperCHIPS approach. Global communication on the Si-IF requires the development of a network on IF that, in addition to communication, provides a wide variety of services to the system on the Si-IF. The NoIF is based on HUDs as nodes within the network. Global communication is realized using $\mu\text{-SerDes}$ IPs integrated as sub-components of the HUDs. Since global communication is only effective starting a communication distance of 10 mm, semi-global communication is realized using repeaters that are included in both UD and FDs. The area overhead of the NoIF is determined according to a relevant number and size of UD. The number and size of UD is extracted from the minimal distance between UD that makes utilization of the NoIF cost-effective.

ACKNOWLEDGMENTS

This research is supported in part by ASCENT, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) sponsored by the Defense Advanced Research Projects Agency, 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.

REFERENCES

- [1] Pranav Ambhore, Umesha Mogera, Boris Vaisband, Mark Goorsky, and Subramanian S. Iyer. 2019. PowerTherm Attachment Process for Power Delivery and Heat Extraction in the Silicon-Interconnect Fabric. In *Proceedings of the IEEE International Electronic Components and Technology Conference*.
- [2] Adeel Bajwa, SivaChandra Jangam, Saptadeep Pal, Boris Vaisband, Randall Irwin, Mark Goorsky, and Subramanian S. Iyer. 2018. Demonstration of a Heterogeneously Integrated System-on-Wafer (SoW) Assembly. In *Proceedings of the IEEE International Electronic Components and Technology Conference*. 1926–1930.
- [3] Adeel A. Bajwa, SivaChandra Jangam, Saptadeep Pal, Niteesh Marathe, Tingyu Bai, Takafumi Fukushima, Mark Goorsky, and Subramanian S. Iyer. 2017. Heterogeneous Integration at Fine Pitch ($\leq 10 \mu\text{m}$) Using Thermal Compression Bonding. In *Proceedings of the IEEE International Electronic Components and Technology Conference*. 1276–1284.
- [4] William D. Brown. 1999. *Advanced Electronic Packaging: With Emphasis on Multichip Modules*. IEEE, Chapter 1: Introduction and Overview of Microelectronics Packaging.
- [5] Sungjun Chun, Wiren D. Becker, Jon Casey, Steve Ostrander, Daniel Dreps, Jose A. Hejase, Ryan M. Nett, Brian Beaman, and Jason R. Eagle. 2018. IBM POWER9 Package Technology and Design. *IBM Journal of Research and Development* 62, 4/5 (July–September 2018), 12:1–12:10.
- [6] Christopher Gonzalez, Eric Fluhr, Daniel Dreps, David Hogenmiller, Rahul Rao, Jose Paredes, Michael Floyd, Michael Sperling, Ryan Kruse, Vinod Ramadurai, Ryan Nett, Saiful Islam, Juergen Pille, and Donald Plass. 2017. 3.1 POWER9™: A Processor Family Optimized for Cognitive Computing with 25Gb/s Accelerator Links and 16Gb/s PCIe Gen4. In *Proceedings of the IEEE International Solid-State Circuits Conference*. 50–51.
- [7] Mark S. Goorsky, Kari Schjolberg-Henriksen, Brett Beekley, Tingyu Bail, Karthick Mani, Pranav Ambhore, Adeel Bajwa, Nishant Malik, and Subramanian S. Iyer. 2018. Characterization of Interfacial Morphology of Low Temperature, Low Pressure Au-Au Thermocompression Bonding. *Japanese Journal of Applied Physics* 57, 251 (2018), 02BC03.
- [8] Subramanian S. Iyer. 2016. Heterogeneous Integration for Performance and Scaling. *IEEE Transactions on Components, Packaging, and Manufacturing Technology* 6, 7 (July 2016), 973–982.
- [9] Subramanian S. Iyer, SivaChandra Jangam, and Boris Vaisband. 2019. From Homogeneous SoCs to Heterogeneous SoWs. In *Proceedings of the Government Microcircuit Applications & Critical Technology Conference (GOMACTech)*.
- [10] SivaChandra Jangam, Adeel Bajwa, Kannan Kalappurakai Thankappan, Prem-sagar Kittur, and Subramanian S. Iyer. 2018. Electrical Characterization of High Performance Fine Pitch Interconnects in Silicon-Interconnect Fabric. In *Proceedings of the IEEE International Electronic Components and Technology Conference*. 1823–1828.
- [11] SivaChandra Jangam, Saptadeep Pal, Adeel Bajwa, Sudhakar Parmarti, Puneet Gupta, and Subramanian S. Iyer. 2017. Latency, Bandwidth and Power Benefits of the SuperCHIPS Integration Scheme. In *Proceedings of the IEEE International Electronic Components and Technology Conference*. 86–94.
- [12] Mary Y. Lanzerotti, Giovanni Fiorenza, and Rick A. Rand. 2005. Microminiature Packaging and Integrated Circuitry: The Work of E. F. Rent, with an Application to On-Chip Interconnection Requirements. *IBM Journal of Research and Development* 49, 4.5 (July 2005), 777–803.
- [13] Saptadeep Pal, Daniel Petrisko, Matthew Tomei, Puneet Gupta, Subramanian S. Iyer, and Rakesh Kumar. 2019. Architecting Waferscale Processors: A GPU Case Study. In *Proceedings of the IEEE International Symposium on High-Performance Computer Architecture*.
- [14] Partha P. Pratim, Cristian Grecu, Michael Jones, André Ivanov, and Resve Saleh. 2005. Performance Evaluation and Design Trade-offs for Network-on-Chip Interconnect Architectures. *IEEE Trans. Comput.* 54, 8 (August 2005), 1025–1040.
- [15] Ujjash Shah, Umesha Mogera, Pranav Ambore, Boris Vaisband, Subramanian S. Iyer, and Timothy S. Fisher. 2019. Dynamic Thermal Management for Silicon Interconnect Fabric using Flash Cooling. In *Proceedings of the IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*.
- [16] Eric Sorensen, Boris Vaisband, SivaChandra Jangam, Tim Shirley, and Subramanian S. Iyer. 2019. Integration and Characterization of InP Dies on Silicon Interconnect Fabric. In *Proceedings of the IEEE International Electronic Components and Technology Conference*.
- [17] UCLA. 2019. *Si-IF Design Manual*. (available under NDA).
- [18] Boris Vaisband. 2017. *3-D ICs as a Platform for Heterogeneous Systems Integration*. Ph.D. Dissertation. University of Rochester, Rochester, NY.
- [19] Boris Vaisband, Adeel Bajwa, and Subramanian S. Iyer. 2018. Network on Interconnect Fabric. In *Proceedings of the IEEE International Symposium on Quality Electronic Design*.
- [20] Boris Vaisband and Eby G. Friedman. 2018. Heterogeneous 3-D ICs as a Platform for Hybrid Energy Harvesting in IoT Systems. *Future Generation Computer Systems* 87 (October 2018), 152–158.
- [21] Boris Vaisband and Subramanian S. Iyer. 2019. Communication Considerations for Silicon Interconnect Fabric. In *Proceedings of the ACM/IEEE International Workshop on System Level Interconnect Prediction*.