Design Flow for Active Interposer-Based 2.5D ICs and Study of RISC-V Architecture with Secure NoC

Heechun Park, Jinwoo Kim, Venkata Chaitanya Krishna Chekuri, Majid Ahadi Dolatsara, Mohammed Nabeel, Alabi Bojesomo, Satwik Patnaik, Ozgur Sinanoglu, Madhavan Swaminathan, Saibal Mukhopadhyay, Johann Knechtel, and Sung Kyu Lim

***Abstract*—Interposer-based 2.5D integrated circuits (ICs) enable the chip-level reuse of hard intellectual properties (IPs), also known as *chiplets*. Such system-level integration shortens the design cycle considerably for large-scale and heterogeneous chips. Besides traditional interposers, which only provide passive elements and routing, *active interposers* are furthermore comprised of logic components. When implemented carefully using a dedicated electronic design automation (EDA) flow, an active interposer can significantly improve the design quality and flexibility for 2.5D ICs. In this paper, we present a complete EDA flow and design strategies targeting such active interposer-based 2.5D ICs. Our key contributions include the co-analysis of power, performance, signal and power integrity, and the related co-optimization of chiplets and the active interposer. Our benchmark is a 64-core RISC-V architecture, organized into multiple chiplets and interconnected by a system-level network-on-chip (NoC). For efficiency, we embed the NoC routers and integrated voltage regulators (IVRs) into the active interposer. Moreover, we integrate security monitors into the interposer-based NoC to protect the system and its shared memories against adversarial traffic. The simple yet powerful benefit of this implementation is to offer *security by construction*, as it is based on a clear physical separation between critical and trusted components (the system-level NoC) versus commodity components (the chiplets). We contrast our active, secured design to a passive, unsecured design baseline of the same RISC-V benchmark, and we find that the active design reduces the silicon area by 18.5%, power by 3.2%, and IR-drop by 73.7%, respectively.**

***Index Terms*—2.5D IC, chiplet, active interposer, EDA flow, signal integrity, power integrity, network-on-chip (NoC), hardware security.**

I. Introduction

INTERPOSER-BASED 2.5D integrated circuit (IC) technology enables chip-level reuse of heterogeneous intellectual property (IP) components. The underlying functional blocks are designed by different vendors and fabricated as separate chips, called *chiplets*, which are then integrated onto an interconnect carrier, the *interposer*. Compared to the traditional 2D IC technology, the 2.5D IC technology reduces the time to market drastically as it allows designers to choose commodity, off-the-shelf chiplets and to integrate them directly at the system level. Thus, 2.5D integration also renders the system-level design iterations more flexible and less complex, as such would only require to exchange the concerned chiplet(s) and re-evaluate the overall system, while for traditional 2D ICs, one is required to re-design and re-implement the whole system from scratch.

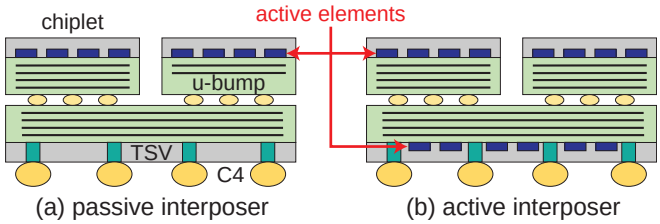

Fig. 1: The two flavors for interposer-based 2.5D ICs.

A. Passive Versus Active Interposer-Based 2.5D Designs

Fig. 1 shows 2.5D ICs in abstraction based on different types of interposers. In both cases, the chiplets are flipped and mounted onto the interposer using microbumps (μ-bumps). The passive interposer-based 2.5D IC (Fig. 1(a)), called *passive 2.5D design* for simplicity in the remainder, is the classical approach that allows only for passive elements to be used in the interposer. Although such passive 2.5D designs allow for heterogeneous reuse of chiplet IP, refraining from active elements (i.e., standard cells, repeaters) in the interposer requires extra efforts for the 2.5D design flow, such as properly driving signals through long distances, clock distribution without clock buffers, and a delicate implementation of system-level interconnects like a network-on-chip (NoC).

An active interposer-based 2.5D IC (Fig. 1(b)), called *active 2.5D design* for simplicity, represents an emerging approach for better utilization of the interposer silicon. In addition to the inter-chiplet connections, an active interposer also allows for active elements. This approach helps to solve essential challenges arising from passive 2.5D designs by allowing for, e.g., repeaters for long-distance wiring in the interposer,

Heechun Park, Jinwoo Kim, Venkata Chaitanya Krishna Chekuri, Majid Ahadi Dolatsara, Madhavan Swaminathan, Saibal Mukhopadhyay, and Sung Kyu Lim are with the School of Electrical and Computer Enginering, Georgia Institute of Technology, Atlanta, GA, 30332, USA (e-mail: heechun@gatech.edu; jinwookim@gatech.edu; vchekuri3@gatech.edu; majidahadi@gatech.edu; madhavan.swaminathan@ece.gatech.edu; saibal.mukhopadhyay@ece.gatech.edu; limsk@ece.gatech.edu).

Satwik Patnaik was with the Department of Electrical and Computer Engineering, Tandon School of Engineering, New York University, Brooklyn, NY, 11201, USA. He is currently with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843 USA (e-mail: satwik.patnaik@tamu.edu).

Mohammed Nabeel, Alabi Bojesomo, Ozgur Sinanoglu, and Johann Knechtel are with the Division of Engineering, New York University Abu Dhabi (NYU AD), Saadiyat Island, 129188, UAE (e-mail: mtn2@nyu.edu; asbojesomo@gmail.com; ozgursin@nyu.edu; johann@nyu.edu).

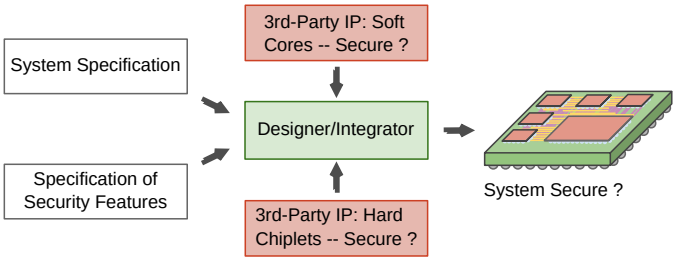

Fig. 2: In modern IC supply chains, potentially malicious third-party IP components play a major role, requiring tailored security features as well as trusted design and manufacturing of those features.

regular clock-tree synthesis with clock buffers, flexible NoC design [1], and improved signal and power integrity [2], [3].

Given their large size, a practical concern for active interposers are yield rates. However, when designing a low-utilized active interposer (e.g., below 5% utilization), which is also known as *minimally-active interposer* [4], the yield losses are expected to be comparable to passive interposers. Since process nodes affect defect rates and yield, and also manufacturing cost, older and mature nodes are preferred for active interposers [4]. For example, for the recent prototype of an active 2.5D design [5], the active interposer is implemented using a mature $65nm$ technology, while the chiplets are realized using a $28nm$ technology. Moreover, with a utilization of only 0.08 transistors/μm^2, the interposer of [5] can be considered as minimally-active as well.

B. Hardware Security

The notion of hardware security has become ever-more pressing to ensure confidentiality, integrity, and availability of data processing—also known as the CIA triad—in electronic systems [6]. The CIA triad also applies for the hardware itself [7]. Prominent security features to hinder software- and/or hardware-based threats at runtime include enclaves for trusted execution and memory management (e.g., *Intel SGX*, *ARM TrustZone*, or the academic *MIT Sanctum*; all three are reviewed in [6]), wrappers for secure design-time integration of third-party IP modules [8], verifiers of computation and external communication [9], or secure NoC architectures [10].

It is essential to note that all these features require a fully trusted design and manufacturing flow, a requirement that is difficult to realize for outsourced IC supply chains relying on third-party IP components (Fig. 2). This concern is even more true for large-scale and heterogeneous systems like 2.5D ICs, where efforts by the designer/integrator to incorporate security features can be hampered as follows. For hard-IP chiplets, security features may not be present at all, or undermined by untrustworthy design and manufacturing procedures initiated by the chiplet vendor. For soft-IP cores, security features have to be tailored for interaction with the cores as well as for system-level compatibility. Furthermore, security features targeting soft-IP cores still necessitate trustworthy fabrication.

C. Contributions

Prior works focused on trade-off analysis for active versus passive interposers in terms of high-level cost [11] or studied essential components for active 2.5D designs, such as the power delivery network (PDN) [3] or NoCs [1]. Recently, a sophisticated prototype of an active 2.5D design was presented in [5]. [1] However, to our best knowledge, there is no prior art proposing an electronic design automation (EDA) flow tailored for active interposers and studying passive versus active 2.5D designs in detail. Besides, prior studies considered only passive interposer for IP protection [12] or protection against Trojan insertion [13], but no prior work considered active interposers to enforce system-level security at runtime.

In this paper, we make the following contributions:

- We present an end-to-end EDA flow for active 2.5D designs, based on an industrial-grade tool setup.
- We showcase the flow by implementing a 64-core RISC-V architecture [14] using chiplets and an active interposer, and we discuss relevant design strategies. We leverage the $28nm$ process node for the commodity chiplets, and the mature $65nm$ node for the minimally-active interposer.
- We propose and implement security features for the NoC embedded in the active interposer. This way, the active interposer becomes the trustworthy backbone for system-level communication, with a clear physical separation from the commodity chiplets to be monitored at runtime.
- We conduct a thorough design study for that benchmark architecture, where we also compare the active 2.5D design with a passive 2.5D baseline counterpart. According to our experiments, the active 2.5D design, even with security features, fares better in terms of power, performance, and area (PPA) as well as signal/power integrity when compared to the passive 2.5D design without security features.

II. RISC-V Architecture with Secure NoC

A. Overview

We use a 64-core design of the proven *RISC-V Rocket Core* architecture [14]. The passive 2.5D baseline of our 64-core design has been introduced as *ROCKET-64* in [15]. Fig. 3 illustrates the architecture in overview: it contains 8 *Rocket* tiles, a central NoC, a memory controller (MC), and 4 integrated voltage regulators (IVRs). The NoC and its communication protocol are introduced in Sec. II-C, while the security features for the NoC are detailed in Sec. II-D.

Each of the 8 *Rocket* tiles contains 8 *Rocket* cores, with a private $16KB$ L1 cache in each core, and a $1MB$ L2 cache shared across the 8 cores. As explained in [15], to construct a *Rocket* tile, we generate an 8-core *Rocket* system using the

[1]The prototype described in [5] was realized by CEA-LETI researchers in collaboration with STMicroelectronics. Their active 2.5D design comprises 96 MIPS cores with caches divided into 6 chiplets using a $28nm$ technology, and a $65nm$ active interposer which embraces dynamic voltage-frequency scaling (DVFS) with switched capacitor voltage regulators (SCVR), NoC links for interconnects, memory-IO controller, and other regular infrastructures such as clock network, thermal and stress sensors, and a design-for-test (DfT) scheme. This demonstrator is a remarkable achievement and could serve well as point of inflection for the broader acceptance of active interposers. Unfortunately, in terms of design automation and design flow, the reference [5] does not provide many insights. In any case, given that the architectures as well as physical implementations are quite different, we have to refrain from a direct comparison of our work with [5].

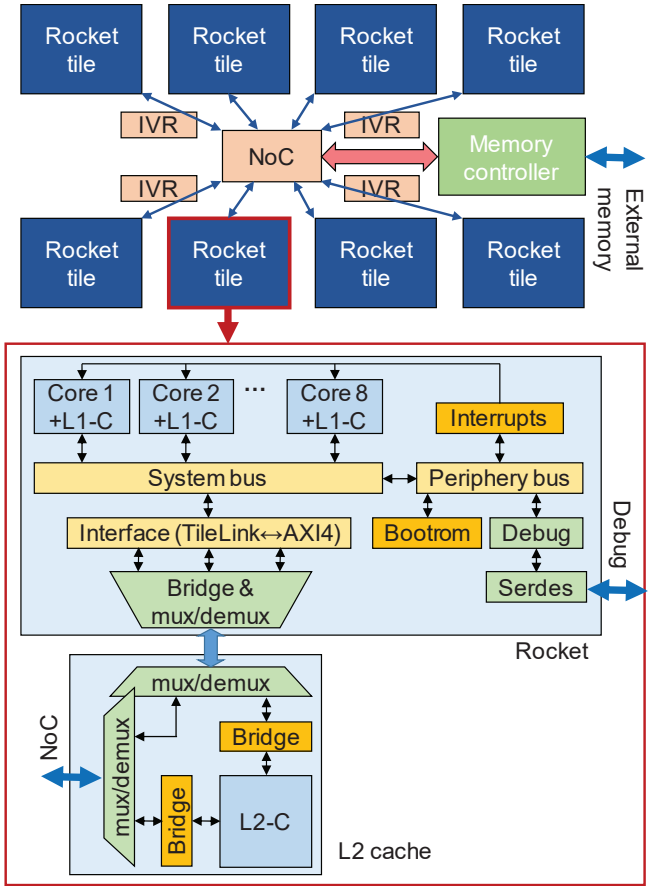

Fig. 3: The *ROCKET-64* architecture, extended from [15].

Rocket chip generator provided in [14]. For the L1 caches, we attach 8 copies of a $16KB$ memory block generated using a commercial $28nm$ memory compiler. We use the same memory compiler for the $1MB$ L2 cache, which consists of 8 copies of a $128KB$ memory block.

B. Organization and Strategies for Active 2.5D Design

There are 30 chiplets in the original passive 2.5D design of [15], summarized in Table I. Note that the chiplets are implemented using different, dedicated nodes and the metal layers for the passive interposer follow the $65nm$ mature node. The support of such a heterogeneous technology setup is one of the main benefits of interposer-based system design. Also note that, to balance the area required for the signal/power μ-bump arrays in the chiplets with the chiplets' design utilization, each *Rocket* tile is separated into a *Rocket* chiplet and an L2 cache chiplet. Besides the functional chiplets, there are four inductor and four capacitor chiplets which work along with the IVRs for power delivery. We use IVRs instead of off-chip voltage regulators, to reduce the power settling time and the interposer PDN impedance. Instead of using one IVR and mounting a digital low-dropout (DLDO) module into each Rocket tile, as done in [15], here we use 4 IVR modules for improved power integrity for both the passive and the active 2.5D design (more details are discussed in Sec. IV-D).

TABLE I: Chiplets in the 2.5D *ROCKET-64* Designs

Chiplet (#)	Footprint ($mm \times mm$)	#I/O Bumps (Signal / Power)	Tech Node
Passive & Active 2.5D Designs			
Rocket (8)	1.70×1.70	58 / 383	$28nm$
L2 Cache (8)	1.46×1.46	92 / 104	$28nm$
Mem. Ctrl. (1)	0.80×1.40	588 / 112	$28nm$
Inductor (4)	1.60×3.40	-	-
Capacitor (4)	0.70×1.20	-	-
Passive 2.5D Design Only			
NoC (1)	0.68×1.56	655 / 108	$28nm$
IVR (4)	0.48×1.20	12 / 240	$130nm$

For our active 2.5D design, we integrate the NoC and the IVR modules directly into the $65nm$ minimally-active interposer, instead of implementing them as chiplets. Doing so has been suggested by prior studies—an NoC within an active interposer allows for more flexible NoC design and better chiplet connectivity [1][11], whereas IVRs within an active interposer achieve a more reliable power supply [2][3]. Embedding the NoC and IVR modules in the interposer also naturally saves chiplet cost along with all the related silicon. While embedding modules into the interposer may also pronounce other concerns like thermal management or area blockage by additional through-silicon vias (TSVs) required to link the embedded modules to the package, we argue a minimally-active interposer design helps to keep such concerns in bound. Finally, the other chiplets in the active 2.5D design remain the same as for the passive 2.5D design [15].

Fig. 4 illustrates the cross-section of both passive and active 2.5D designs. Note that the passive 2.5D design (Fig. 4(a)) derived from [15] already comprise an advanced interconnect structure when compared to classical passive 2.5D designs, in the sense that bi-directional Advanced Interface Bus (AIB) drivers [16] are included here to reduce signal bump count and to mitigate signal distortion in the long and unbuffered interposer wires. Still, each AIB driver has to be designed to support a particular range of interconnect length during chiplet implementation and, thus, the placement of chiplets is restricted during subsequent 2.5D integration. Therefore, relying on AIB drivers alone does not provide a fundamental solution to advance large-scale and complex 2.5D ICs.

Contrary to the passive 2.5D design, our active 2.5D design (Fig. 4(b)) is completely free from any long unbuffered wires, as we embed repeaters inside the active interposer.[2] These repeaters serve to reduce signal distortion and net switching power (or energy per bit transmitted). Note that the chiplets for the active 2.5D design still contain AIB drivers; these drivers are bi-directional and hence essential for reducing the original signal bump count by half.

Embedding the NoC into the active interposer serves another benefit; it enables the assembly of potentially untrusted commodity chiplets into a system that can yet remain secure. The key difference between our approach and prior security studies is a clear physical separation between commodity and security components; we do *not* require any security features within commodity chiplets, nor any trust assurance regarding their de-

[2]This insertion of repeaters is performed automatically using the EDA flow described in Section III.

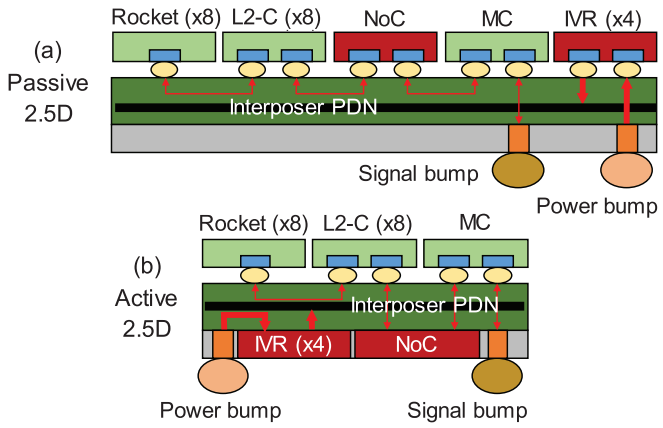

Fig. 4: Cross-section of both *ROCKET-64* design implementations.

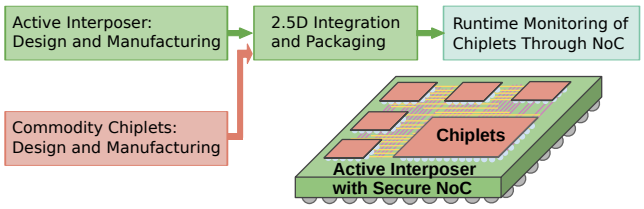

Fig. 5: Active 2.5D design and manufacturing stages (trusted; green), chiplets design and manufacturing (untrusted; red), and security features in action (turquoise).

TABLE II: Interposer Design Rules

	Passive	Active
Metal Layer #	4	4 (M2–M5)
Wire Width	$0.4\mu m$	$0.4\mu m$ (NDR)
Wire Min. Pitch	$0.8\mu m$	$0.8\mu m$ (NDR)
Via Diameter	$0.7\mu m$	$0.7\mu m$ (NDR)
μ-Bump Diameter	$20\mu m$	
μ-Bump Min. Pitch	$40\mu m$	
μ-Bump Array Pitch (Avg.)	$80\mu m$	
C4-Bump Diameter	$90\mu m$	
C4-Bump Min. Pitch	$180\mu m$	
C4-Bump Array Pitch (Avg.)	$400\mu m$	$360\mu m$

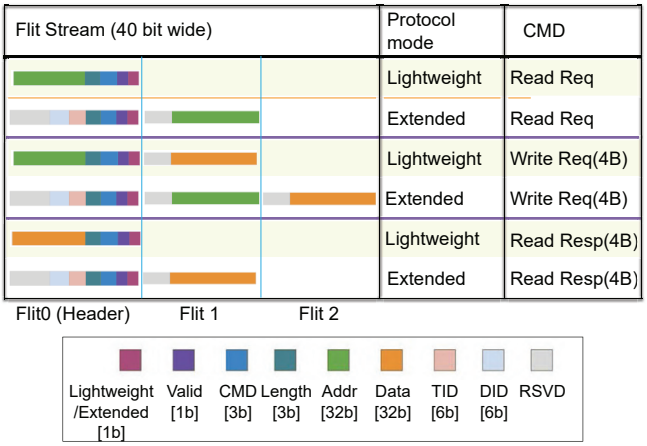

Fig. 6: Flit representation of the Hybrid-Link protocol [15].

sign and manufacturing. That is an essential aspect for modern supply chains, where enforcing security and trust assurances across all components and involved parties is impractical. Thus, using our approach, one can maintain full flexibility for 2.5D integration using various off-the-shelf chiplets and still provide a secure platform. Toward this end, we require some security features for the NoC (detailed in Sec. II-D), and we require a trustworthy design and manufacturing flow for the active interposer containing this secure NoC as well as for the final 2.5D IC assembly (Fig. 5). As active interposers are manufactured preferably using mature technologies, we argue that it is realistic to be able to commission fully trusted design and manufacturing facilities toward that end.

As just motivated in the context of security, as well as for minimally-active interposers and high yield rates (Sec. I), we leverage a mature, commercial-grade $65nm$ technology node for the active interposer. We tuned the process design kit (PDK) rules to match those of the passive 2.5D design rules in [15] as closely as possible using non-default rules (NDRs); see Table II. Only the average pitch of C4-bump arrays is reduced for the active 2.5D design, to fit all C4-bumps within the smaller die outline of the active-interposer floorplan, but the reduced pitch has still significant margin over the given minimum pitch of $180\mu m$. Regarding metal layer usage, all standard cells in the active interposer have their pins located regularly in M1 (along with some cell-level routing), but the pins are connected directly to M2 using vias. In other words, we block M1 for signal routing, which is required to enable a fair comparison for interposer routing across the passive and the active 2.5D designs.

C. Communication Protocol and Overview on NoC

A standardized communication protocol is indispensable for any system to accommodate IP blocks from different vendors, and the Hybrid-Link (HL) protocol proposed in [15] is suitable for inter-chiplet communication in 2.5D ICs. Fig. 6 shows the functional unit (flit) representation of the HL protocol. The protocol has two different modes: *lightweight* for simple point-to-point connections, and *extended* for transaction-based communication. The selection of mode depends on the needs of the applications running on the system. In comparison to other protocols used within the *Rocket* chiplet, like *AXI4* and *TileLink*, the HL protocol is a simple serializing protocol with a default flit width of only 40 bits; this is essential to reduce the number of signal μ-bumps per chiplet and, thus, limit their die outlines and silicon area.

The NoC is generated using OpenSMART [17], and all its connected components communicate using the HL protocol. That is, all chiplets have protocol translators attached at their AIB interfaces to communicate with other chiplets using the HL protocol. The NoC consists of 12 routers in a 4×3 mesh topology (Fig. 7). The 4 routers in the middle row connect to the 4 four channels of the MC chiplet, linking the system to some shared, external DDR4 memory. The remaining 8 routers are connected to the L2 chiplets using a secure network interface (SNI), which is described in Sec. II-D. Note that each *Rocket* chiplet is connected via direct interposer links to its L2 cache chiplet, bypassing the NoC for access to this chiplet-level shared cache. The L2 chiplets, in turn, pass only requests to the external memory onto the NoC.

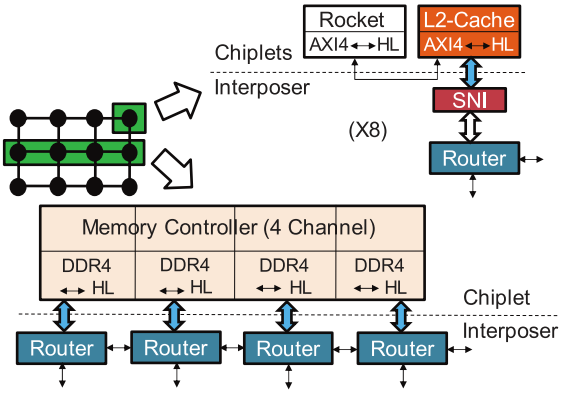

Fig. 7: NoC architecture.

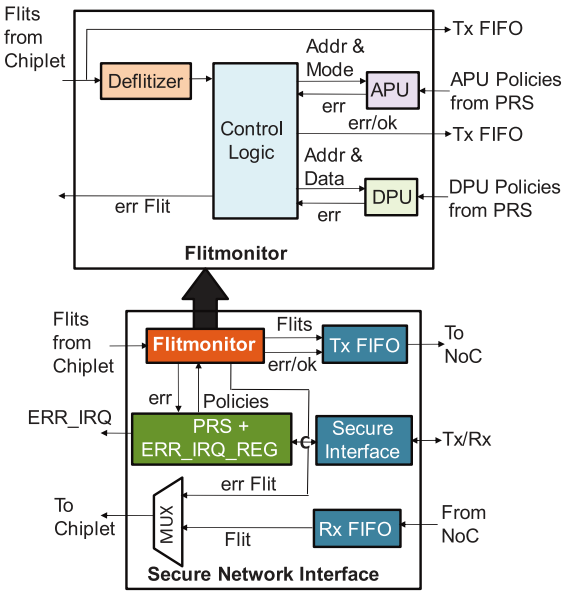

Fig. 8: Micro-architecture of NoC security features.

D. Threat Model and Security Features for NoC

Since *ROCKET-64* is a multi-core architecture interconnected by an NoC and using an external shared memory, we have to consider the following security threats [8][10]: (i) snooping of NoC communication, (ii) spoofing of communication/transaction identifiers, (iii) modification or diversion of NoC communication, (iv) malicious access or modification of data in the shared memory, and (v) leaking of sensitive data from cores/caches into the shared memory.

We assume that any of these threats can be introduced by (a) the cores themselves, either unintentionally via "hardware bugs" or intentionally via Trojans, or (b) malicious software running on the cores. However, we assume a trusted runtime environment; any physical attacks conducted by end-users are out of scope. We also assume that attacks are exercised explicitly by NoC communication targeting at the shared memory; any adversarial side-channel activities arising within chiplets and their cores (e.g., as in [18]) or across cores and caches (e.g., as in [19]) are also out of scope.

Recall that the NoC, including all SNIs, is residing exclusively in the trusted active interposer. This way, our architecture rules out—by construction—the threats (i) to (iii). Consider for example spoofing: each SNI assigns hard-coded core identifiers to any request and, thus, a malicious core cannot masquerade itself. For the remaining threats, i.e., malicious activities targeting at the external shared memory, we enforce *policy-driven monitoring* of all memory requests with the help of SNIs. At its heart, our security scheme is inspired by the seminal work of Fiorin *et al.* [10]. However, the notable difference for ours is the clear physical separation enabled by 2.5D integration—we ensure security by construction, even in the presence of untrustworthy chiplets, whereas the work in [10] requires a fully trusted supply chain for all components.

As shown in Fig. 8, each SNI comprises an address protection unit (APU), a data protection unit (DPU), a flit handler, and control logic—all encapsulated into a flitmonitor. Further, each SNI comprises a policy register space (PRS), a secure interface (UART in our implementation), and some glue logic. The PRS is managed via the secure interface by an external trusted system, which also schedules and controls the workload of the 2.5D *ROCKET-64* "workhorse" system.

Next, we provide some details on the working of the SNI. After decomposing incoming Hybrid-Link flits, the control logic delegates the relevant parameters to the APU and, for write requests, also to the DPU. Then, APU and DPU leverage the PRS to check whether to approve or reject the transaction. For burst transfers, policies are checked for each beat in the burst until final approval/rejection. Each APU policy describes the type of transaction, allowed for a particular address range and core. Each DPU policy describes sensitive data that can only be written in a particular address range by a specific core. If no matching policies are found, the transaction is rejected, an interrupt request (ERR_IRQ in Fig. 8) is delegated to the external trusted system, and an error response is sent to the respective core. Note that all flits undergoing APU and DPU checks are also queued in a Tx FIFO until approval/rejection; in case of rejection, the FIFO is reset, preventing all the flits of the entire transaction from being pushed to the NoC.

III. EDA Flow for Active 2.5D Designs

The interposer in a passive 2.5D design usually contains only chiplet-to-chiplet interconnects and some additional passive elements (e.g., decoupling capacitors), which renders such an interposer analogous to conventional packaging substrates like printed circuit boards (PCBs). Therefore, in the passive 2.5D design flow proposed in [15], a commercial package-centric design tool (Cadence SiP Layout) is used for automated routing of the passive interposer, and the design quality is measured by RLGC-based SPICE modeling and related simulation tools (Ansys HFSS, Synopsys HSPICE). For an active 2.5D design, however, such a flow is not applicable. In addition to the inter-chiplet connections, an active interposer also contains active components (i.e., functional blocks with standard cells). Thus, the active interposer is more like a separate chip. Even a low-utilized active interposer can include a large number of active elements (e.g., our active interposer with 2.68% utilization contains more than 480K standard cells,

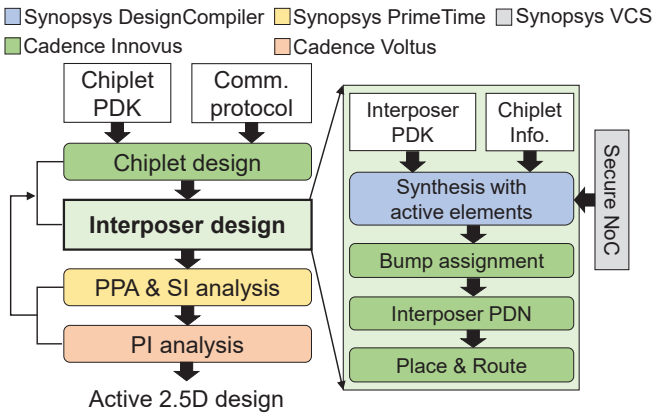

Fig. 9: Our EDA flow for active 2.5D designs.

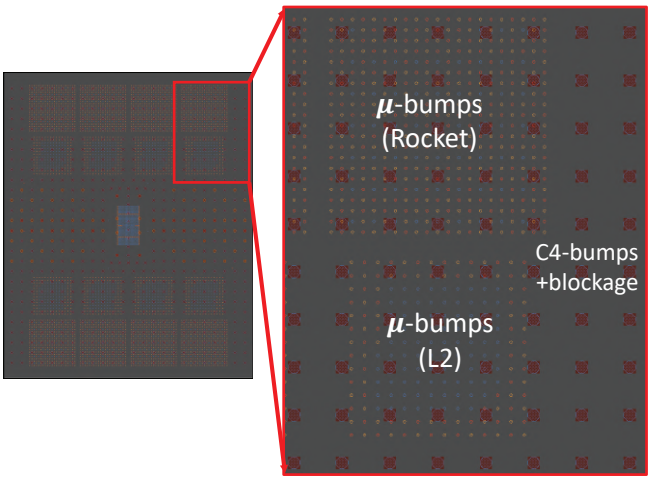

Fig. 10: Bump arrangement and representations as I/O pins for the active 2.5D design of the *ROCKET-64* architecture.

see Table IV in Sec. IV-A). Thus, analyzing all the related interconnects using the modeling and simulation methods proposed for in the passive 2.5D design flow is not practical. Moreover, the passive 2.5D design flow does not support the placement of active components, to begin with.

Therefore, we propose a novel flow for active 2.5D designs (Fig. 9). This flow utilizes the classical chip-centric design and analysis tools (e.g., Cadence Innovus, Synopsys PrimeTime, Cadence Voltus), which support verified place-and-route (P&R) algorithms and timing/power evaluation for the active elements in the interposer within a reasonable runtime. Next, we describe the stages of the flow. More details are also discussed throughout the study in Sec. IV.

First, all chiplets are designed and implemented separately, including any additional blocks required for 2.5D integration (e.g., I/O drivers, protocol translators). Note that this step is largely independent of the actual active 2.5D design and may, especially in commercial setups, also be conducted by external vendors, as long as these vendors incorporate a standardized chiplet communication protocols, like Hybrid-Link.

Second, we leverage the interposer PDK and all the chiplets' physical constraints (e.g., pin locations) to synthesize the interposer RTL with all its active elements, including the secure NoC, using a commercial synthesis tool (Synopsys Design Compiler). We also verify and simulate the security features using Synopsys VCS.

Third, all I/O bumps, including μ-bumps and C4-bumps, are assigned to the interposer floorplan based on a GUI-guided arrangement of chiplets on the interposer (and considering the flip-chip attachment of chiplets). Each chiplet's μ-bump array is transformed into a corresponding array of I/O pins, and the C4-bumps are represented by I/O pins with placement blockages reserved for the pads of TSVs required to connect the C4-bumps with the metal layers of the active interposer. For example, Fig. 10 shows the bump arrangement for the active 2.5D design of this work.

Fourth, we perform the power delivery network (PDN) design, clock-tree synthesis (CTS), and conduct P&R, all using a commercial design tool (Cadence Innovus), to generate the final interposer design. That design now includes an optimized placement of all modules and all clock trees, along with all the inter-chiplet links. The prior flow for passive 2.5D design could not support those heterogeneous requirements.

Fifth, we perform design analyses (PPA as well as signal and power integrity; SI and PI), again using commercial tools (Synopsys PrimeTime and Cadence Voltus). This stage provides more accurate results when compared to the selective, modeling-based analysis for the passive 2.5D design flow, and it does so within reasonable runtime.

Lastly, in case the PPA metrics are not satisfactory, or some SI/PI violations are noted, the interposer design stages are revisited. While doing so, the concerning designs constraints, such as the chiplet floorplan(s) or the bump assignments, are revised. If co-optimization of chiplets and the interposer are possible (i.e., in case chiplets are not procured as physical hard IP from other vendors), the chiplet designs are also revisited, e.g., to reconfigure the I/O driver strengths or power/ground networks as needed. In general, while we propose a notion of iterative co-analysis and co-optimization, we note that current commercial tools would not allow for this to be fully automated. Given the advanced nature of an active, heterogeneous interposer design, there is still manual supervision and decision making required during co-analysis and co-optimization, such as the revision of design constraints indicated above.

IV. Study on Active 2.5D Design of RISC-V Architecture with Secure NoC

Next, we provide a thorough study of our active 2.5D design. The passive 2.5D design represents the baseline to quantify the design benefits of the active 2.5D design. For a fair comparison, both 2.5D designs have most functional chiplets in common as hard IP, with their respective PPA results reported in Table III. Fig. 11 shows the layouts of all chiplets except the passive inductor/capacitor chiplets.

Fig. 12 shows the layouts of both the 2.5D designs. For the passive 2.5D design (Fig. 12(a)(b)), all inter-chiplet connections are routed using the *Automatic Router* of Cadence SiP Layout. For the active 2.5D design (Fig. 12(c)(d)(e)), the 12 routers of the NoC are placed underneath their corresponding chiplets, to avoid long wiring between chiplets and their NoC

TABLE III: PPA Results of Chiplets in Both 2.5D Designs

Chiplet	*Rocket*	L2 Cache	Mem. Ctrl.	NoC (Passive 2.5D Design Only)
# Cells	923,764	3,670	80,986	52,074
WL (m)	18.25	0.26	1.70	1.07
Util. (%)	60.73	79.99	13.10	27.30
Max Freq. (GHz)	1.10	1.14	1.13	1.11
Total Power (W)	1.035	0.019	0.068	0.045

The IVR chiplet/module as well as the inductor and capacitor chiplets do not constitute functional/logic components, hence they are not listed here; related details are explained in Sec. IV-D.

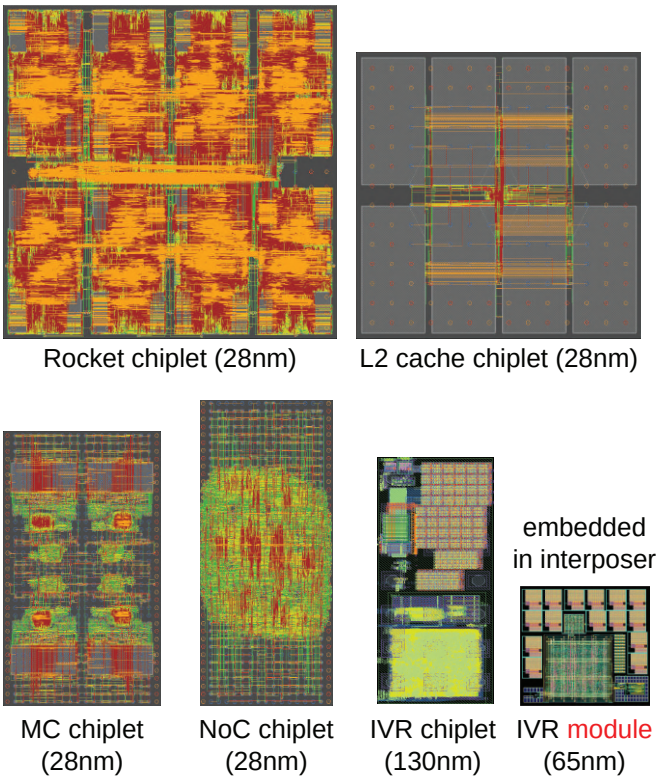

Fig. 11: GDSII layouts of the chiplets used in our 2.5D designs. Note that the NoC and IVR chiplets are used only in the passive 2.5D design, whereas the NoC and IVRs are implemented as modules in the active interposer (see also Fig. 12 for the latter).

interface. Recall that the *Rocket* chiplets are using direct links to connect to their L2 cache chiplets, while only the L2 cache chiplets are connected to the NoC (Fig. 7). Therefore, the related routers are placed underneath the L2 cache chiplets. The 4 IVR modules are placed at the center of the active interposer. After GUI-guided floorplanning, regular P&R is performed for the interposer using Cadence Innovus. More details for the EDA flow are given in Sec. III.

A. PPA Evaluation and Comparison

Table IV summarizes our active 2.5D design in comparison with the passive counterpart. Note that this comparison considers an unsecured NoC chiplet for the passive 2.5D design, and a secured NoC module for the active 2.5D design.

The active interposer incurs 18.5% less silicon area, by embedding the NoC and the IVRs and thereby reducing the footprint required for placement of all chiplets along with their μ-bump arrays. The core utilization of the active interposer is

TABLE IV: PPA Comparison of Passive Versus Active 2.5D Design

	Passive	Active
Chiplet Technology	$28nm$, $130nm$	$28nm$
Interposer Technology	$65nm$	$65nm$
Security	Unsecured	Secured
Footprint (mm)	10.8×10.8	8.8×10.8 (-18.5%)
Interposer Cell #	0	480,709
(Repeater #)	(0)	(98,639)
Interposer Utilization	0%	2.68%
# Metal Layers	4	4
Interposer Net #	1,420	481,485
Interposer WL (m)	5.068	30.134 (5.95×)
Avg. Net WL (mm)	3.582 (57.2×)	0.063
Interposer Power	172.8 mW	167.2 mW (-3.2%)
Net Power	172.8 mW	80.2 mW (-53.6%)
Cell Power	-	86.6 mW
Leakage Power	-	0.4 mW

only 2.68%,[3] which is considered as minimally-active [4] and thus free from additional yield losses.

The active interposer is reported with 5.95× total wirelength of the passive interposer. That is because the active interposer includes all the wiring for the secure NoC, which is scattered over the interposer as the routers are attached underneath their corresponding chiplets, whereas the passive interposer only contains a much smaller number of simple inter-chiplet connections. Still, each connection in the passive interposer is long and unbuffered; these unbuffered wires are on average 57.2× longer than the buffered wires in the active interposer. These long unbuffered wires consume a large amount of power during signal transfers, such that the nets' power consumption for the passive interposer is reported as more than 2× that of the active interposer. Regarding the total power consumption including all active elements, the active 2.5D design still consumes less power than the passive counterparts, even with the additional security features in the NoC.

B. Clock Networks for the Active 2.5D Design

While each chiplet has its own clock network built in, the traditional passive 2.5D design requires additional efforts to deliver the clock signal to all chiplets as it is not possible to use clock buffers within the interposer, and there is *no* clock-network synthesis in the passive 2.5D design flow [15].

The active 2.5D design, on the contrary, allows for clock buffers to be inserted in the active interposer and, thus, a reliable clock distribution can be achieved. Our active 2.5D design flow makes the interposer clock generation straightforward, by utilizing the CTS (clock-tree synthesis) engine of the commercial design tool (i.e., Cadence Innovus).

There are two different clock domains in the active interposer of our active 2.5D design: the one driving all chiplets (*chiplet clock*), and the one driving the registers in the interposer (*interposer clock*). Due to the different technology nodes leveraged for the chiplets and the interposer ($28nm$ and $65nm$), we apply different target frequencies for the two clock domains ($1GHz$ and $200MHz$), and we use two different

[3]Note that all cells are placed sparsely, due to the NDRs leveraged from the passive 2.5D design imposing a 4× larger wire width and pitch than the regular rules for the $65nm$ PDK. Thus, the cell utilization may visually appear higher than 2.68% in Fig. 12(d), but it is not.

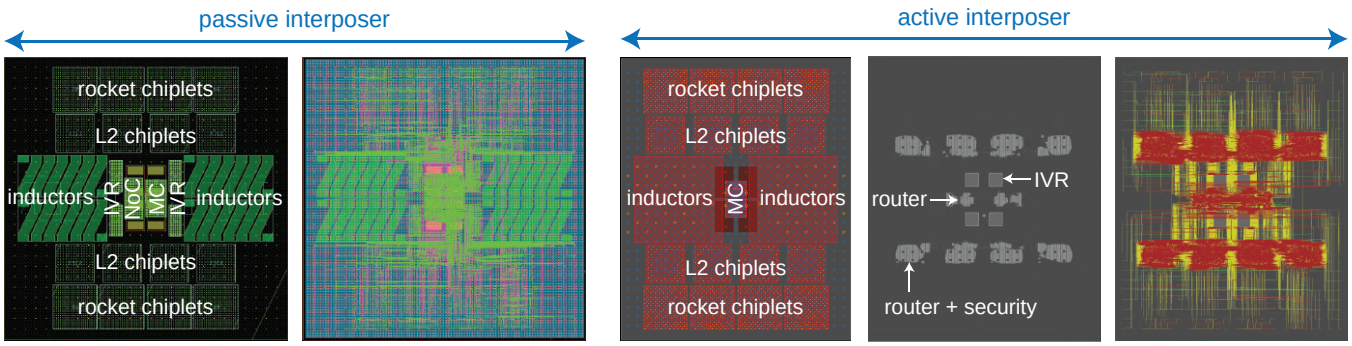

Fig. 12: GDSII layouts of our passive and active 2.5D *ROCKET-64* designs.

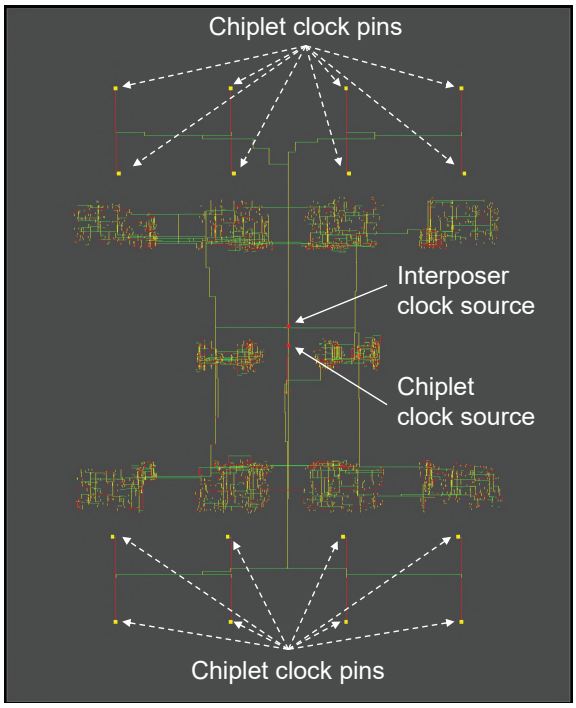

Fig. 13: Layout of clock trees in the active 2.5D *ROCKET-64* design.

clock sources from two dedicated external clock ports. Both clock ports are located in the middle of the interposer, to allow for better clock skew control. During CTS, both clock networks are generated and buffered separately.

Fig. 13 shows both clock networks in the active interposer layout; Table V provides the metrics for both networks. The chiplet clock tree has only 17 target leafs which are the clock source ports for the functional chiplets (8 *Rocket* chiplets, 8 L2 cache chiplets, and the MC chiplet), and it is constructed as a simple H-tree with the leafs connected to the respective μ-bumps of the chiplets. The interposer clock network, however, has a more complex structure since the active interposer holds 61,617 registers which are spread out. In any case, the CTS engine of Cadence Innovus can manage this structure well.

Given that the active interposer is implemented in the more mature $65nm$ node, the performance of the interposer lacks behind the $28nm$-based chiplets. Still, recall that such heterogeneous setup represents one of the key benefits of interposer-based systems; one can procure and integrate commodity chiplets based on advanced nodes, while a mature node is preferred for the active interposer to manage yield. Moreover, our flow is able to consider particular cases of high-performance signals, e.g., the link to the external DDR4 memory, simply by assigning higher priority on those particular nets and adding NDRs, to have them routed with optimized transmission delay.

TABLE V: Metrics of the Active Interposer Clock Networks

	Chiplet Clock	Interposer Clock
Target Freq.	$1GHz$	$200MHz$
Clock Skew	$0.23ns$	$1.40ns$
Clock Leaf #	17	61,617
Clock Buf/Inv #	52	2,092
Clock WL	$34.37mm$	$958.54mm$
Clock Power	$3.97mW$	$41.69mW$

Besides the IVR modules, which are not performance-critical, the active interposer comprises the secure NoC. However, the NoC handles only a fraction of the system's overall traffic (i.e., the requests to the shared external memory), whereas the majority of signal transfers stay directly within the *Rocket* chiplets or across the *Rocket* chiplets and their L2 caches, with the latter using direct interposer links.[4] Still, to allow for synchronization between chiplets and the NoC, we require FIFO structures at all NoC interfaces. Overall, the active 2.5D design is realized without any "bottlenecks" arising from the heterogeneous technology setup.

C. Evaluation and Comparison of Signal Integrity

For the signal integrity (SI) analysis of the passive interposer, recall that we employ RLGC modeling of transmission line samples (Fig. 14(a)) and Monte-Carlo-driven SPICE simulations. As doing so requires considerable runtime, only a few critical interconnects with worst-case crosstalk can be studied.

For the active 2.5D design, there are many transistors and a significant number of wires in the interposer—RLGC modeling and simulation would become even more complex and more time-consuming. As indicated, the active 2.5D design is fully compatible with commercial SI tools (e.g., Synopsys

[4]These direct links are driven by AIB interfaces designed within the chiplets themselves. Thus, these links are not subject to the interposer target frequency and rather enable high-performance transmission with delays of only $\approx 50ps$.

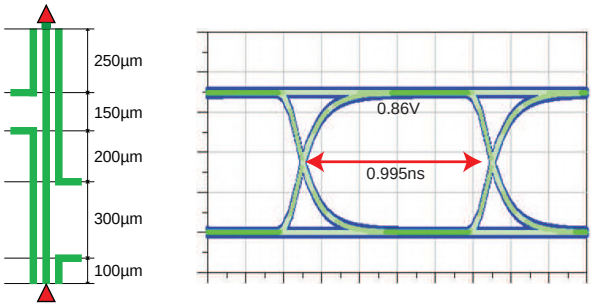

(a) interposer wire sample (b) eye diagram

Fig. 14: Example for transmission line sample and eye diagram simulation for the passive 2.5D *ROCKET-64* design.

TABLE VI: Signal Integrity Results for Passive and Active Interposer

Passive Interposer		(Sample, Runtime: 30 min)	
Data Rate	I/O Driver Imp.	Eye Width	Eye Height
1 $Gbps$	50 Ω	0.995 ns	0.86 V
Active Interposer		**(Full design, Runtime: 8 min)**	
Max Voltage Bump		Delta Delay	
Bump Width	Bump Height	Max Net Δ-Delay	Δ-WNS
0.701 ns	0.187 V	66.34 ps	85.01 ps

PrimeTime SI), by simply extracting the standard parasitic exchange format (SPEF) file from the active interposer layout and importing that file into the SI tool. This way, we can obtain a comprehensive SI analysis for the entire interposer, and the runtime for analysis is even much shorter than the simulation-based analysis for a small selection of worst-case wires in the passive design. Note that package-centric design tools (e.g., Cadence SiP Layout) commonly do not support RC parasitic extraction and SPEF files but only allow for extraction of, e.g., S-parameters and related models. Hence, it is not practical to use chip-centric SI tools like Synopsys PrimeTime SI to analyze the passive interposer as well.

Given that the SI analysis involves different steps and covers substantially different components for the passive versus active interposer, the results and their interpretation are not directly comparable in a consistent format. More specifically, the simulation-based analysis of the passive interposer refers to eye diagrams (Fig. 14(b)), where a larger width and a larger height of the eye imply better SI. For the active interposer, we consider the following two metrics: voltage bump and delta delay. Voltage bump is an unexpected pulse generated from crosstalk by neighboring aggressors, and delta delay is the additional delay a net experiences due to crosstalk. The lower both metrics, the better the SI. Note that the eye width of the passive interposer and the delta delay of the active interposer both describe the "delay" for SI, while the eye height of the passive interposer and the voltage bump of the active interposer represent the "signal robustness" for SI.

Table VI shows the SI analysis results for both interposers. SI results for the passive interposer cover average cases, whereas the SI results for the active interposer cover the worst-case for each metric. While the passive interposer incurs a relatively low voltage drop, recall that the related SI analysis covers only a small fraction of the interconnects. The active interposer holds more complex routing structures, which naturally result in somewhat larger crosstalk effects (voltage bump, delta delay). However, the SI analysis is more accurate and covers worst-case scenarios; thus, the findings are more robust than those for the passive interposer. Finally, also note that our SI analysis accurately covers the active interposer, and represents a solid approximation for a comprehensive 2.5D system-wide SI analysis. We have accounted for the external I/O signals of the interposer and all chiplets, using custom constraints within the design flow, to satisfy timing closure of the entire active 2.5D design (i.e., also across all the bumps). Once accurate models for the particular parasitics of the μ-bumps and C4-bumps are available, a more detailed 2.5D system-wide SI analysis can be easily conducted as well, given that both the interposer and chiplets RC parasitics are already expressed using SPEF.

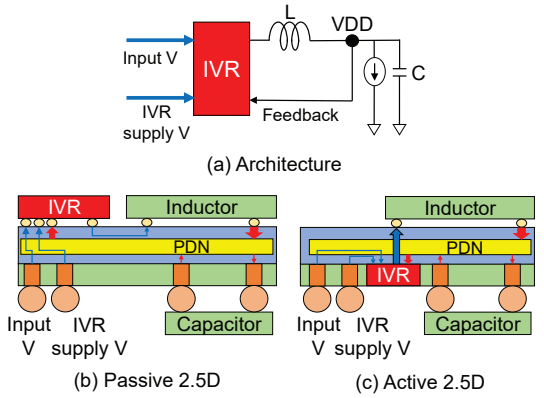

Fig. 15: Conceptual view of our power delivery scheme and its implementation for both 2.5D design designs. Note that both designs contain 4 instances of that circuitry.

D. Evaluation and Comparison of Power Integrity

Power integrity (PI) for the 2.5D designs concerns the power delivery to all chiplets with as little voltage drop as practically possible. As indicated, we use IVR modules for both 2.5D designs instead of choosing off-chip voltage regulators, to reduce the power settling time and the interposer PDN impedance. Instead of using one IVR module and mounting a digital low-dropout (DLDO) module into each *Rocket* chiplet, as done in [15], here we use 4 IVR modules to improve PI.

Fig. 15(a) illustrates the power delivery circuitry, comprising an IVR along with coupled inductor and capacitor. The IVR receives the input voltage ($1.2V$) and converts it to the $0.9V$ supply voltage for the chiplets and active elements. This supply voltage is charged using the coupled inductor/capacitor, and is fed to the interposer PDN mesh. The IVR module then monitors the supply voltage level via its feedback loop, and compensates the voltage consistently as needed. Each of the 4 IVR modules is designed and placed as separate chiplet onto the interposer for the passive 2.5D design (Fig. 15(b)), or embedded directly within the interposer for the active 2.5D design (Fig. 15(c)).

Table VII lists the power and currents for all chiplets. To satisfy the power requirements of all chiplets and active elements, we set the target current load of both 2.5D designs

TABLE VII: Power Consumption for Chiplets and Interposer in the 2.5D *ROCKET-64* Designs, VDD is 0.9V

Chiplet	Power (W)	Count	Current (A)
Rocket	1.035	8	9.200
L2 Cache	0.019	8	0.169
Mem. Ctrl.	0.045	1	0.050
Common Chiplets			9.419
NoC	0.068	1	0.076
(Passive Int.)	0.173	1	0.192
Total (Passive 2.5D)			9.687
(Active Int.)	0.167	1	0.186
Total (Active 2.5D)			9.605
Target Supply Current			10.0

TABLE VIII: PDN Metrics for Passive Versus Active Interposer

		Passive	Active
Mesh	Interposer Metal #	2 (M1–M2)	2 (M2–M3)
	Width / Space (μm)	40 / 100	12 / 12
	Occupancy (%)	62	30
IVR	Implementation	Chiplet	Module Block
	Tech. Node	$130nm$	$65nm$
	Size (mm)	0.48×1.20	0.5×0.5
	Conversion (V)	$3.6 \rightarrow 0.9$	$1.2 \rightarrow 0.9$
	Settling Time (ns)	125	223
	Efficiency (%)	76.0	78.3
	Consumed Power (W)	0.70	0.62
Max IR-Drop (mV)		106.30	27.94

to $10A$. Therefore, in both designs, each power supply circuit (IVR module) is designed for a $2.5A$ target load.

Table VIII shows the PDN metrics for both interposers. For the passive interposer, PDN meshes are placed manually using the package design tool (Cadence SiP Layout), whereas for the active interposer, we construct the power ring and stripes using the automated power planner of the chip design tool (Cadence Innovus). The wire width of the active interposer PDN is set to the maximum allowed in the commercial $65nm$ PDK. The PDN occupancy of the metal layers is less for the active interposer compared to passive interposer; that is because routing resources in the active interposer are also required for signal routing. The IVR metrics also differ, since the IVR is re-designed for the active 2.5D design using the commercial $65nm$ PDK.

The PI analysis flow is different for passive versus active interposers. While we can only use RLGC modeling and SPICE simulation (as with the SI analysis) for the passive interposer, the active interposer is in principle compatible with the commercial tool (Cadence Voltus). However, the tool by itself can only consider 2D designs. Thus, we model the active 2.5D design as follows. For the power-rail analysis of the active interposer, we describe the four μ-bump positions of the inductor chiplets as voltage sources, and we add external current regions, which indicate the specific areas that draw the currents, to all chiplets' VDD μ-bump locations. This way, we can effectively conduct a 2.5D sytem-wide PI analysis of all components. Fig. 16 shows the resulting IR-drop map of our active 2.5D design, with the active interposer at the heart of the PI analysis. The maximum IR-drop of 27.94 mV is just 3.1% of the supply voltage; this drop is 73.7% lower than the drop observed for the analysis of the passive 2.5D design.

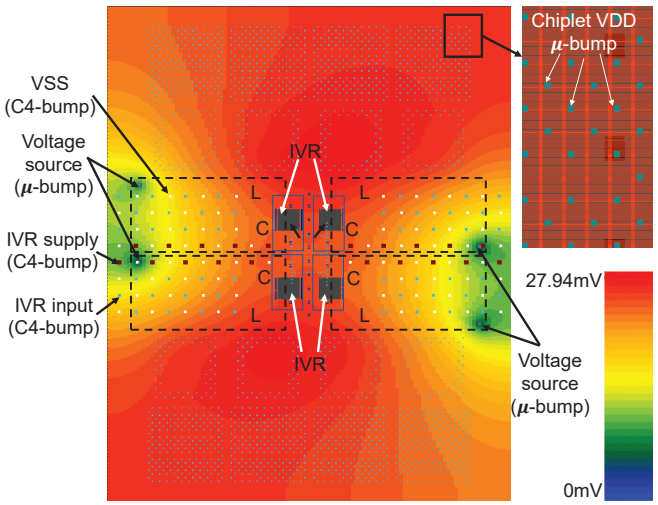

Fig. 16: IR-drop map for the active interposer. All chiplet's VDD μ-bump locations are described as external current regions, to capture their currents drawn from the power rails. IVRs are power supply units; their IR-drop is omitted.

Overall, for the same target current load, the active 2.5D design exhibits better PI (i.e., less IR-drop) because there are 4 IVR modules embedded in the active interposer itself, instead of using 4 IVR chiplets as is the case for the passive 2.5D design. Embedded IVR modules help to reduce the length of power paths and thus to better maintain the level of supply voltage for all the chiplets. Moreover, the chip-centric design tool (Cadence Innovus) tends to provide more optimized PDN structures when compared to the manual approach leveraged for the passive interposer.

E. Evaluation of Secure NoC

The security features of the NoC have been verified and simulated using Synopsys VCS. We implement a testbench to run multiple testcases, in the form of various requests to the external memory originating from different cores running different applications. We compile the related, exemplary APU policies describing protected shared and private memory regions in the external memory, and exemplary DPU policies describing protected data assets (e.g., private crypto keys) from leaking into the external memory. For simplicity of simulation, we assume that the external trusted system provides the APU/DPU policies during bootstrapping and initializes the SNIs accordingly. We also consider only the essential component, i.e., the SNI as security-enforcing interface right before the NoC router, and we neither consider the NoC itself (with all its routers working in tandem) nor all chiplets. Next, we describe a representative simulation run in some detail.

Fig. 17 shows the VCS simulation waveform for two memory requests, both originating from the core with ID 02. The first request is a simple Hybrid-Link read request (lightweight mode), which is approved. The related incoming flit (with packet $*$_5022) is checked against all the policies relevant for this core (only one policy is illustrated in the waveform). As none of the policies is disapproving this request, the flit is passed onto the NoC. The second request is a more complex

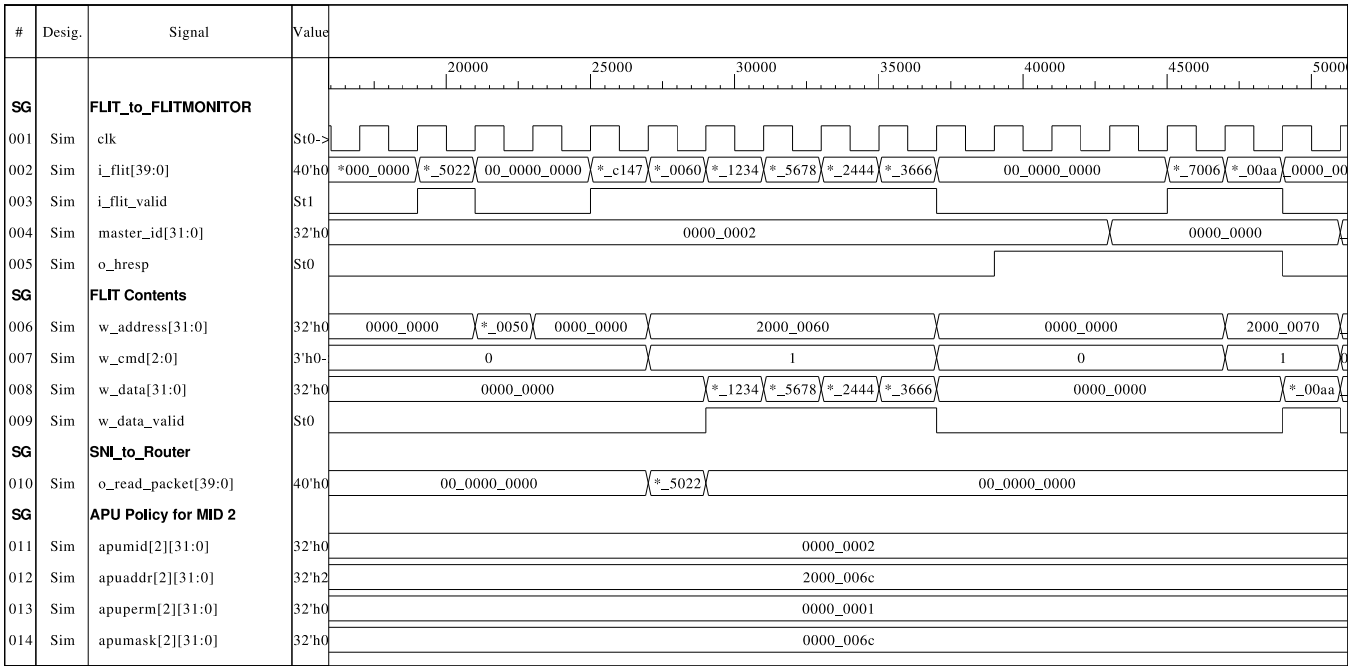

Fig. 17: Exemplary Hybrid-Link requests being checked by the secure network interface (SNI) residing before the related NoC router. Not all but only one APU policy is illustrated, there are zero-valued flits injected between requests, and the VCS waveform is stripped via graphical post-processing; this is all done for clarity of the simulation result.

Hybrid-Link write request (extended mode), which is rejected. The related incoming 6 flits (with packets $*\_c147, \ldots, *\_3666$ in burst mode) are checked one by one. Here, the policy illustrated in the waveform leads to rejection of this request, indicated by the signal o_hresp. The rejection is because the concerned address 2000_0060 is within a memory region accessible only for reading by that core, but not for writing.

As for latency, handling of the first read request incurs 4 cycles in total, whereas handling of the second write request incurs only 1 additional cycle. That is because the second request can be decided upon already early on during its burst transmission of data, effectively realizing pipelining. Also, note that the other incoming flits are set to all zero for clarity of the waveform; in normal operation, the SNI can handle a continuous stream of requests through pipelining.

As for the physical design implementation, the overheads incurred by the additional security features are listed in Table IX, with respect to the NoC itself as well as with respect to the overall system. While some costs are expected, it is important to note that the core utilization remains within the range of minimally-active interposers [4]. Also, note that costs are marginal in the context of the entire 2.5D *ROCKET-64* design: these system-level costs are 1.17% more power, 4.56% more gates, and 11.83% more wirelength.

Recall that we place an SNI in the source router of each of the L2 cache chiplets. This decision incurs a trade-off: given that our NoC has 8 L2 cache chiplets/source routers, but only 4 target routers (for the 4 MC channels), placing SNIs at the source routers instead of the target routers incurs twice the minimal cost required for securing the NoC. However, when securing the source routers, we ensure that adversarial traffic never passes onto the NoC, preventing any malicious traffic to begin with, and potentially reducing the load and energy consumption of the NoC accordingly at runtime.

TABLE IX: Design Cost for NoC Security Features in the Active 2.5D Design

	NoC w/o Security	NoC w/ Security	Cost w.r.t. 2.5D Design
Cell #	132,945	480,709 (3.62×)	+4.56%
Utilization (%)	1.58	2.68	-
Wirelength (m)	11.108	30.134 (2.71×)	+11.83%
Total Power (mW)	66.7	167.2 (2.51×)	+1.17%

V. Conclusion

In this paper, we propose an end-to-end EDA flow for active 2.5D designs, based on commercial chip design and analysis tools. Among other steps, this flow enables signal and power integrity analysis with high accuracy, which provides robust insights for the benefits of active interposers over passive interposers. We demonstrate our flow for a large-scale *RISC-V ROCKET-64* benchmark architecture and discuss related design strategies. Utilizing the notion of 2.5D integration, we also provide a scheme for clear physical separation between commodity chiplets and a trustworthy communication backbone, namely a secured NoC residing in the active interposer. This scheme hinders critical threats for heterogeneous chiplets integration by construction, and further offers runtime monitoring of system-level memory requests. Our active 2.5D design with secure NoC achieves better layout results (smaller footprint, less and more stable power consumption) compared to its unsecured passive 2.5D counterpart. For future work, we plan for system-level functional simulation using *gem5*, among other aspects.

REFERENCES

[1] J. Yin *et al.*, "Modular Routing Design for Chiplet-Based Systems," in *ACM/IEEE ISCA*, 2018, pp. 726–738.

[2] J. Kim, "Active Si interposer for 3D IC integrations," in *IEEE 3DIC*, 2015, pp. TS11.1.1–TS11.1.3.

[3] S. Kim, Y. Kim, K. Cho, J. Song, and J. Kim, "Design and Measurement of a Novel On-Interposer Active Power Distribution Network for Efficient Simultaneous Switching Noise Suppression in 2.5-D/3-D IC," *Trans. Compon., Pack., Manuf. Tech.*, vol. 9, no. 2, pp. 317–328, 2019.

[4] A. Kannan, N. E. Jerger, and G. H. Loh, "Enabling interposer-based disintegration of multi-core processors," in *IEEE/ACM MICRO*, 2015, pp. 546–558.

[5] P. Vivet *et al.*, "A 220GOPS 96-core processor with 6 chiplets 3D-stacked on an active interposer offering 0.6ns/mm latency, 3Tb/s/mm2 inter-chiplet interconnects and 156mW/mm2@ 82%-peak-efficiency DC-DC converters," in *IEEE ISSCC*, 2020, pp. 46–48.

[6] P. Maene, J. Gtzfried, R. de Clercq, T. Mller, F. Freiling, and I. Verbauwhede, "Hardware-based trusted computing architectures for isolation and attestation," *IEEE TC*, vol. 67, no. 3, pp. 361–374, 2018.

[7] J. Knechtel, S. Patnaik, and O. Sinanoglu, "Protect your chip design intellectual property: An overview," in *Proc. Conf. Omni-Layer Intell. Syst.*, 2019, pp. 211–216.

[8] A. Basak, S. Bhunia, T. Tkacik, and S. Ray, "Security Assurance for System-on-Chip Designs with Untrusted IPs," *IEEE TIFS*, vol. 12, no. 7, pp. 1515–1528, 2017.

[9] H. Zhang *et al.*, "Architectural support for containment-based security," in *Proc. Arch. Supp. Programm. Lang. Op. Sys.*, 2019, pp. 361–377.

[10] L. Fiorin, G. Palermo, S. Lukovic, V. Catalano, and C. Silvano, "Secure Memory Accesses on Networks-on-Chip," *IEEE TC*, vol. 57, no. 9, pp. 1216–1229, 2008.

[11] D. Stow, Y. Xie, T. Siddiqua, and G. H. Loh, "Cost-effective Design of Scalable High-performance Systems Using Active and Passive Interposers," in *IEEE/ACM ICCAD*, 2017, pp. 728–735.

[12] Y. Xie, C. Bao, and A. Srivastava, "Security-aware 2.5D integrated circuit design flow against hardware IP piracy," *Computer*, vol. 50, no. 5, pp. 62–71, 2017.

[13] F. Imeson, A. Emtenan, S. Garg, and M. V. Tripunitara, "Securing computer hardware using 3D integrated circuit (IC) technology and split manufacturing for obfuscation," in *Proc. USENIX Sec. Symp.*, 2013, pp. 495–510.

[14] K. Asanovic *et al.*, "The Rocket Chip Generator," *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2016-17*, 2016. [Online]. Available: https://github.com/chipsalliance/rocket-chip

[15] J. Kim *et al.*, "Architecture, Chip, and Package Co-design Flow for 2.5D IC Design Enabling Heterogeneous IP Reuse," in *ACM/IEEE DAC*, 2019, pp. 1–6.

[16] D. Kehlet *et al.*, "Accelerating Innovation Through a Standard Chiplet Interface: The Advanced Interface Bus (AIB)," https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/accelerating-innovation-through-aib-whitepaper.pdf.

[17] H. Kwon and T. Krishna, "OpenSMART: Single-cycle multi-hop NoC generator in BSV and Chisel," in *IEEE ISPASS*, 2017, pp. 195–204.

[18] R. J. Masti, D. Rai, A. Ranganathan, C. Müller, L. Thiele, and S. Capkun, "Thermal Covert Channels on Multi-core Platforms," in *Proc. USENIX Sec. Symp.*, 2015, pp. 865–880.

[19] D. A. Osvik, A. Shamir, and E. Tromer, "Cache attacks and Countermeasures: the Case of AES," in *IACR Crypt. ePrint Arch.*, 2005.

Heechun Park received the B.S. degree in Electrical and Computer Engineering from Seoul National University, Seoul, South Korea in 2011, and the integrated M.S. and Ph.D. degree from the same department in 2018. Dr. Park worked as a Post-Doctoral Fellow with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, from 2018 to 2020. He is about to work as a Senior Researcher in Inter-university Semiconductor Research Center, Seoul National University. Dr. Park's research interests cover physical design of digital integrated circuits (ICs) including vertically stacked 3D and 2.5D ICs. He also stretches his interests to system-level hardware architecture design, and supporting computer-aided design (CAD) with deep learning.

Jinwoo Kim received the B.S. degree in Electrical and Computer Engineering and the M.S. degree in Electrical Engineering and Computer Science from Seoul National University, Seoul, South Korea, in 2011 and 2013, respectively. He is currently pursuing the Ph.D. degree with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA. His current research interests include interposer-based 2.5D IC design and co-analysis, and 3D IC design and methodology.

Venkata Chaitanya Krishna Chekuri (S'18) received the B.Tech degree in Electronics and Communication Engineering from Manipal Institute of Technology, Manipal, Karnataka, India, in 2015. He is currently pursuing the M.S. and Ph.D. degree in Electrical and Computer Engineering with the Georgia Institute of Technology, Atlanta, GA, USA, under the supervision of Prof. S. Mukhopdhyay. He was an Intern with the VLSI Methodology Group, NVIDIA Corp., Santa Clara, CA, USA in 2017. His current research interests include low-power design, and power management in digital circuits.

Majid Ahadi Dolatsara received his B.Sc. degree in Electrical Engineering from K.N.Toosi University of Technology, Tehran, Iran, in 2013, and M.Sc. degree in Electrical Engineering from Colorado State University, Fort Collins, Colorado in 2016. He is currently a Ph.D. candidate in Electrical Engineering at Georgia Institute of Technology, Atlanta, Georgia. His research interest includes electronic design automation and machine learning for signal and power integrity and packaging. Mr. Ahadi Dolatsara was a recipient of the Best Poster Paper Award at the 23rd IEEE International Conference on Electrical Performance of Electronic Packaging and Systems in 2014.

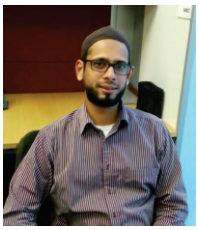

Mohammed Nabeel received his B.tech degree in Electrical and Electronics Engineering from the National Institute of TechnologyCalicut, India. He is currently working as a Research Engineer at the Center for Cyber Security at New York University Abu Dhabi.

Alabi Bojesomo obtained the M.Sc. in Microsystems Engineering from Masdar Institute of Science and Technology Abu Dhabi in 2016, and a B.Sc. in Electrical Engineering from Obafemi Awolowo University (OAU) Ile-Ife in 2011. He is currently a Ph.D. Student at Khalifa University Abu Dhabi, UAE. His research interests include MEMS, Deep learning and hardware security.

Satwik Patnaik received the Ph.D. degree in Electrical Engineering from Tandon School of Engineering, New York University, Brooklyn, NY, USA. He is currently a Post-Doctoral Researcher with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX, USA. Dr. Patnaik received the Bronze Medal in the Graduate Category at the ACM/SIGDA Student Research Competition held at ICCAD 2018, and the Best Paper Award at the Applied Research Competition held in conjunction with Cyber Security Awareness Week in 2017.

Ozgur Sinanoglu is a professor of Electrical and Computer Engineering at New York University Abu Dhabi. He obtained his PhD in Computer Science and Engineering from University of California San Diego. He has industry experience at TI, IBM and Qualcomm, and has been with NYU Abu Dhabi since 2010. During his PhD, he won the IBM PhD fellowship award twice. He is also the recipient of the best paper awards at IEEE VLSI Test Symposium 2011 and ACM Conference on Computer and Communication Security 2013.

Prof. Sinanoglus research interests include design-for-test, design-for-security and design-for-trust for VLSI circuits, where he has more than 200 conference and journal papers, and 20 issued and pending US Patents. Sinanoglu has given more than a dozen tutorials on hardware security and trust in leading CAD and test conferences, such as DAC, DATE, ITC, VTS, ETS, ICCD, ISQED, etc. He is serving as track/topic chair or technical program committee member in about 15 conferences, and as (guest) associate editor for IEEE TIFS, IEEE TCAD, ACM JETC, IEEE TETC, Elsevier MEJ, JETTA, and IET CDT journals.

Prof. Sinanoglu is the director of the Design-for-Excellence Lab at NYU Abu Dhabi. His recent research in hardware security and trust is being funded by US National Science Foundation, US Department of Defense, Semiconductor Research Corporation, Intel Corp and Mubadala Technology.

Madhavan Swaminathan is the John Pippin Chair in Microsystems Packaging & Electromagnetics in the School of Electrical and Computer Engineering (ECE), Professor ECE with a joint appointment in the School of Materials Science and Engineering (MSE), and Director of the 3D Systems Packaging Research Center (PRC), Georgia Tech (GT) (http://www.prc.gatech.edu) . He also serves as the Site Director for the NSF Center for Advanced Electronics through Machine Learning (CAEML) and Theme Leader for Heterogeneous Integration, SRC JUMP ASCENT Center. He formerly held the position of Founding Director, Center for Co-Design of Chip, Package, System (C3PS), Joseph M. Pettit Professor in Electronics in ECE and Deputy Director of the Packaging Research Center (NSF ERC), GT. Prior to joining GT, he was with IBM working on packaging for supercomputers. He is the author of 500+ refereed technical publications, holds 31 patents, primary author and co-editor of 3 books, founder and co-founder of two start-up companies, and founder of the IEEE Conference Electrical Design of Advanced Packaging and Systems (EDAPS), a premier conference sponsored by the EPS society. He is an IEEE Fellow and has served as the Distinguished Lecturer for the IEEE EMC society. He received his BE degree from Regional Engineering College, Tiruchirapalli (now NITT) in 1985 and MS/PhD degrees in Electrical Engineering from Syracuse University in 1989 and 1991, respectively.

Saibal Mukhopadhyay received the B.E. degree in Electronics and Telecommunication Engineering from Jadavpur University, Kolkata, India, in 2000, and the Ph.D. degree in Electrical and Computer Engineering from Purdue University, West Lafayette, IN, USA, in 2006. He is currently the Joseph M. Pettit Professor with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA. He has authored or co-authored more than 200 papers in refereed journals and conferences. He holds five U.S. patents. His current research interests include the design of energy efficient, intelligent, and secure systems. His research explores a cross-cutting approach to design spanning algorithm, architecture, circuits, and emerging technologies. He was a recipient of the Office of Naval Research Young Investigator Award in 2012, the National Science Foundation CAREER Award in 2011, the IBM Faculty Partnership Award in 2009 and 2010, the SRC Inventor Recognition Award in 2008, the SRC Technical Excellence Award in 2005, the IBM Ph.D. Fellowship Award from 2004 to 2005, the IEEE Transactions on Very Large Scale Integration Best Paper Award in 2014, the IEEE Transactions on Components, Packaging and Manufacturing Technology Best Paper Award in 2014, and multiple best paper awards in International Symposium on Low-Power Electronics and Design in 2014, 2015, and 2016. Dr. Mukhopadhyay is a Fellow of IEEE.

Johann Knechtel (Member, IEEE) received the M.Sc. degree in Information Systems Engineering (Dipl.-Ing.) in 2010 and the Ph.D. degree in Computer Engineering (Dr.-Ing., summa cum laude) in 2014, both from TU Dresden, Germany. He is a Research Scientist with New York University Abu Dhabi, UAE. From 2015 to 2016, he was a Postdoctoral Researcher with the Masdar Institute of Science and Technology, Abu Dhabi; from 2010 to 2014, he was a Ph.D. Scholar and Member with the DFG Graduate School on "Nano- and Biotechnologies for Packaging of Electronic Systems" hosted at TU Dresden; in 2012, he was a Research Assistant with the Chinese University of Hong Kong, Hong Kong; and in 2010, he was a Visiting Research Student with the University of Michigan at Ann Arbor, MI, USA. His research interests cover VLSI physical design automation, with particular focus on emerging technologies and hardware security. He has (co-)authored around 50 publications.

Sung Kyu Lim is a Professor of the School of Electrical and Computer Engineering, Georgia Institute of Technology. He received the B.S., M.S., and Ph.D. degrees from the Computer Science Department, University of California, Los Angeles (UCLA), in 1994, 1997, and 2000, respectively. He joined the School in 2001. His research focus is on the architecture, design, test, and electronic design automation (EDA) solutions for 2.5-D and 3-D ICs. His research is featured as Research Highlight in the Communication of the ACM in January, 2014. He is the author of Practical Problems in VLSI Physical Design Automation (Springer, 2008) and Design for High Performance, Low Power, and Reliable 3D Integrated Circuits (Springer, 2013).

Dr. Lim received the National Science Foundation Faculty Early Career Development (CAREER) Award in 2006. He received the ACM SIGDA Distinguished Service Award in 2008. He was an Associate Editor of the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION SYSTEMS during 2007-9 and the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS during 2013-18. He received the Best Paper Award from several conferences in EDA. His works have been nominated for the Best Paper Award at several top venues in EDA and circuit/packaging design. Dr. Lim has published more than 350 papers on 2.5-D and 3-D ICs.