Resonant Clock Synchronization with Active Silicon Interposer for Multi-Die Systems

Ragh Kuttappa, *Student Member, IEEE,* Baris Taskin, *Senior Member, IEEE*,
Scott Lerner, *Student Member, IEEE,* and Vasil Pano, *Member, IEEE*

***Abstract*—This paper presents the integration of resonant clocking to multi-die architectures to synchronize individual chiplets connected through an active silicon interposer. The proposed inter-chiplet synchronization through the active silicon interposer rotary oscillator array (ASI-ROA) provides a unitary clock domain to the multiple die (i.e. multiple chiplets) in the package with a very low design overhead. System performance analysis is performed with parasitics-extracted, post-layout simulation models of two different sizes of representative heterogeneous multi-die architectures, each with varying number of RISC-V cores per die. Each RISC-V core of the multi-die package belongs to the unitary clock domain, designed with ASI-ROA to operate at a frequency of $2\,\mathrm{GHz}$. The proposed architecture is investigated for robustness in frequency and skew across the multi-die system (MDS) with SPICE based simulations of post layout models, demonstrating variations of only $80\,\mathrm{MHz}$ for a $2\,\mathrm{GHz}$ target frequency. The power savings are upto 41% for the overall MDS, compared to an equivalent implementation with a contemporary ADPLL used to synchronize the multiple chiplets over the active interposer. The average clock skew of the completely resonant architecture presented in this work is $8.2\,\mathrm{ps}$.**

***Index Terms*—Resonant rotary clock, synchronization, Multi-Die Systems, low power, VLSI.**

I. Introduction

SILICON interposer based systems allow for integration of heterogeneous dies capitalizing on the yield and cost benefits [1–3]. Heterogeneous integration of multi-die system (MDS) leverages stacking multiple components of the system, including I/O, memory, CPU cores, and GPU cores. Multi-die systems have stimulated the innovation of interconnects between different dies, such as AMD's infinity fabric, Intel's Foveros interposer, and Marvell's Mochi Interconnect [4, 5]. In addition to standard measures of performance, these interconnects differ in their active footprint on the interposer. The footprint on the interposer is important because passive interposers demonstrate superior yield with cost reduction, while active interposers demonstrate superior performance while sacrificing yield [1–10].

Designing robust, high-speed, low-skew, low-jitter, and low-power synchronization (i.e. clocks or asynchronous systems) across multiple die and the interposer(s) is challenging over a large area [3, 6]. To address these challenges, heterogeneous architectures in MDS are designed with clock/data forwarding or asynchronous clocks that require additional circuits and clock domain crossing consideration. The Ground Referenced Signaling (GRS) solution by NVIDIA, for instance, uses high-speed interconnects between dies for clock forwarding from on-chip phase locked loops (PLL). To that end, authors from NVIDIA proposed a globally-asynchronous locally-synchronous (GALS) solution in literature [11, 12]. Clock forwarding solution is the currently adopted industry solution, with the advantages and disadvantages of (costly) high speed interconnects with PLLs delivering the synchronicity of chiplets. The GALS solution benefits from the lack of the costly and sensitive PLLs, however, is subject to the same design overhead and verification challenges that have distanced designers from asynchronous solutions in general. The proposed solution in this paper combines the benefits of high speed interconnects with the benefit of lacking a PLL, while enabling synchronous operation across the chiplets on the package.

Leveraging the high quality interconnect performance offered by the 2.5D system, cross-die synchronization is proposed through resonant clocking in this work. In particular, multi-die systems are synchronized with rotary traveling wave oscillators (RTWOs) combined in a Rotary Oscillatory Array (ROA) topology (of connected RTWOs) over the active silicon interposer for heterogeneous integration, as shown in Fig. 1(a). The RTWOs are designed with transmission lines and inverter pairs that power and amplify the signals adiabatically. The transmission lines are akin to the those high speed interconnects in clock forwarding, however, are different in being routed in the mobius-shaped [13] topology of ROAs. The preservation of energy within the mobius-shaped ROAs enable pseudo-adiabatic operation of the synchronizing clock signal, leading to substantial power savings (>50% reported in literature [13]) compared to traditional clock distribution networks (e.g. tree, spine, mesh). The RTWOs are distributed over the active silicon interposer as an active-silicon interposer rotary oscillatory array (ASI-ROA) (Fig. 1(b)) providing a lightweight synchronized clocking source with minimal footprint. Package-wide solutions are subject to variations without a global feedback between domains of chiplets. The proposed active silicon interposer rotary oscillator array (ASI-ROA) is the feedback between the chiplets, providing robustness to variations exacerbated by the large interposer sizes.

ASI-ROA provides the generation and distribution network of the unitary clocking domain of the MDS. The feasibility of such resonantly synchronous multi-die systems were pre-

The authors are with the Department of Electrical and Computer Engineering, Drexel University, Philadelphia, PA 19104 USA (e-mail: fr67@drexel.edu;taskin@coe.drexel.edu;spl29@drexel.edu;vp93@drexel.edu). Corresponding author: Ragh Kuttappa. This material is based upon work supported by the National Science Foundation under Grant 1816857 and Grant 2008629.

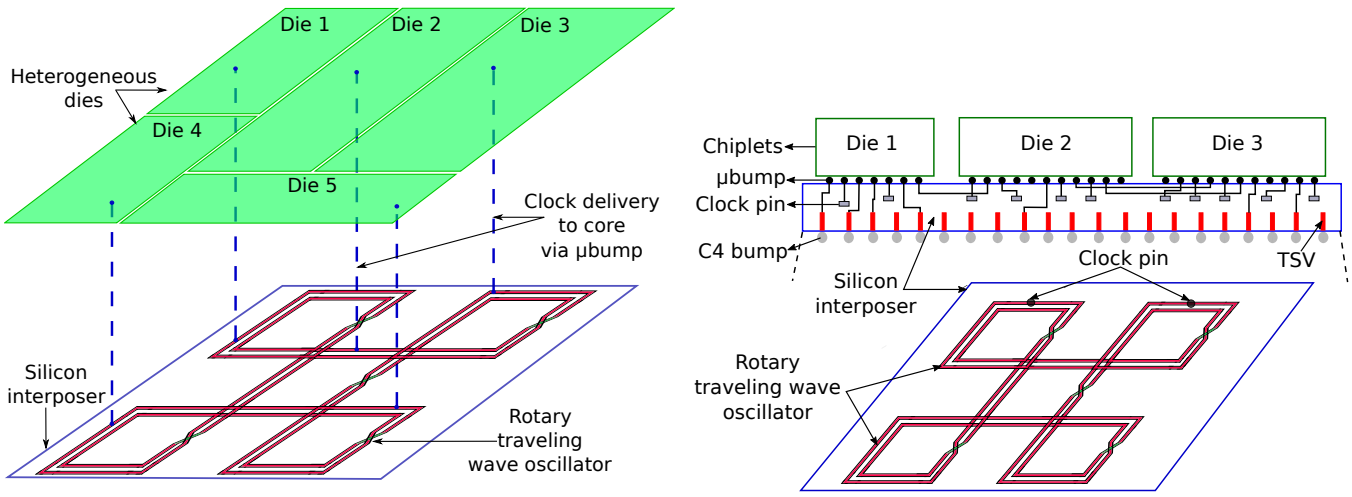

(a) Heterogeneous synchronization with rotary traveling wave oscillators (b) Cross-sectional view of 2.5/3D system with rotary traveling wave oscillators

Fig. 1: Silicon interposer based synchronization for multi-die systems.

viously proposed in [14], for homogenous multi-die systems. This paper builds on the proof-of-concept in [14] to demonstrate the innovations in the circuit and systems integration of resonant multi-die synchronization, with goals of integrating to the emerging need for heterogeneous integration with varying load currents and for intellectual property (IP) reuse and cross-die synchronization for hard IPs. There is heterogeneity in architecture and heterogeneity in technology. In this work, the focus is on a heterogeneous architecture i.e., each chiplet can be a different architecture (accelerator, memory, etc.). The underlying theme for the work in this manuscript is for chiplets that are in the "same clock domain", regardless of homogeneous v.s. heterogeneous integration.

A. Novelty and Contributions

Heterogeneous multi-die systems leverage different fabrication technologies for the dies and the interposer(s), subject to varying sources and magnitudes of variation. The RTWOs that are the clock oscillators within the ASI-ROA architecture in the interposer have self-tuning attributes, in terms of jitter and phase correction through self-resonance, and phase locking between the RTWOs over a large area, that provide robustness to process variations. To evaluate the efficacy of the proposed robust synchronization architecture of ASI-ROA, multiple dies with varying number of the RISC-V cores per die is used to design and simulate a heterogeneous MDS. It is also shown that the methodology with ASI-ROA can accommodate hard IPs at the chiplet level, where performance and variation analysis are performed for a heterogeneous MDS with one representative hard IP and the entire MDS with hard IPs. The proposed ASI-ROA architecture presents a loading balancing cost function that considers the clock pin location, clock pin load, μ_{bump} location, and clock source location that is applicable for the methodology presented in this work and other industrial methodologies. The contributions of this work include:

1) Robust clock generation on the active silicon interposer with an active footprint of 3.1% providing a unitary clocking domain across the MDS,
2) Synchronization of multi-die systems, with or without hard IPs,
3) A redefined load balancing process utilizing the interposer space to accommodate MOSCAPs, and
4) Power analysis comparison to standard all digital PLL (ADPLL) based clock source with buffered clock tree.

B. Paper Outline

The rest of this paper is organized as follows. Background material on resonant rotary clocking and clock network design are reviewed in Section II. The proposed ASI-ROA architecture and methodology for synchronized clock generation and distribution for active interposer based multi-die systems are presented in Section III. The experimental setup and results are presented in Section IV. Concluding remarks are provided in Section V.

II. Resonant Rotary Clocking Background

Rotary traveling wave oscillators (RTWO) are a type of resonant clocking that produces varying phases of clock signals with low-power, constant magnitude, and low jitter [13]. An RTWO is modeled as an LC oscillator,

$$f_{osc} \approx \frac{1}{2\sqrt{L_T C_T}}. \quad (1)$$

The inductance L_T depends on the geometry of the rotary ring and C_T is the total capacitance of the ring, interconnects and devices [13]. For large floorplans, multiple RTWO rings are connected across corner tapping points in an array topology as shown in Fig. 1(b), called a rotary oscillatory array (ROA) [13].

Capacitive load and inductance affect the frequency of oscillations of an RTWO as given by (1). The RTWO rings in

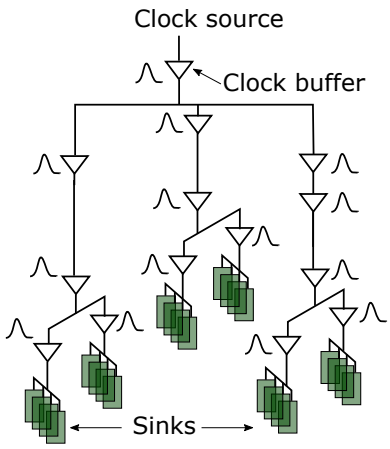

Fig. 2: Bounded skew clock tree illustration showing the clock buffers having variations in process, voltage, and temperature.

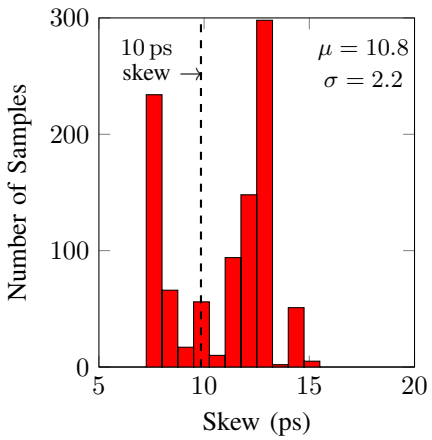

Fig. 3: Histogram for skew variation for buffered clock tree considering PVT, target skew=10 ps.

an ROA topology [13]) are designed with the same physical parameters including the perimeter and structure. Therefore, the inductance L_T of the rotary rings is identical. The total capacitance from (1) is estimated as

$$C_T = \Sigma\ C_{tra} + \Sigma\ C_{inv} + \Sigma\ C_{reg} + \Sigma\ C_{wire}, \qquad (2)$$

where C_{tra}, C_{inv}, C_{reg}, and C_{wire} are the capacitances contributed by the transmission line, inverter pairs, registers, and the register tapping wires, respectively. The capacitance of the transmission lines and inverter pairs are extracted from the geometries of the ROA topology. The capacitance of the registers C_{reg} and the register tapping wires C_{wire} depend on the number of registers connected to each ring and the tapping wire length of each connection. In order to maintain the self-oscillating operation of an RTWO (and the ROA) with L_T and C_T, a balanced distribution of the parasitic components across multiple tapping points on an ROA is necessary.

A. Capacitance imbalance

An increase in the load imbalance across the clock network leads to 1) the clock signal not reaching full-V_{dd} and 2) oscillations not sustaining. To investigate the tolerance of the clock network, the cns02 of the ISPD'10 benchmark suite is designed to operate at 2 GHz with an industrial 28 nm technology node, similar to [15]. The load across the clock network is varied and it is observed that when the capacitive load variation is greater than 25% across the tapping points, the clock signal starts to deteriorate. The resonant clock design process should satisfy this capacitive imbalance tolerance in design (i.e. by balancing capacitance of IPs, chiplets etc. within the tolerance limit). In case where such balancing is not feasible, additional capacitances can be used, minimally, for capacitance balancing purposes.

B. Clock network design

Standard clock network synthesis techniques employ either buffered trees, mesh networks, or hybrid (tree + mesh) topologies [16–21]. Designing clock distribution networks over a large die area incur variations in skew, that require de-skew circuits or significant design overhead. In Fig. 2, an illustration of a bounded-skew tree with clock buffers and sinks is shown. Variations in process, voltage, and temperature affect the delay of each clock buffer. Each clock buffer has its own variation that will affect the overall clock skew at the last stage of the clock tree. Further, variations in skew affect the combinational logic that lead to timing violations. To present the variations in skew, cns02 of the international symposium on physical design 2010 (ISPD'10) benchmark suite is designed with an industrial tool synthesizing a bounded skew tree distributing a 2 GHz clock signal at a commercially available industrial 28 nm technology node. The cns02 benchmark has a die size of 13 mm × 7 mm with 2249 sinks. In Fig. 3, the variation in skew for the 2 GHz bounded skew clock tree with a skew budget of 10 ps is presented. Variations of a pessimistic uniformly random distributed ±10% in threshold voltage, supply voltage, and temperature with fixed load are considered. The mean skew and variation observed with 1000 Monte Carlo runs is 10.8 ps and 2.2 ps, respectively. Overall, the variations cause 61% of the runs to violate the 10 ps skew budget. The variations in skew are primarily due to the variations in the clock buffers in the clock tree. It is expected that the increasing interposer sizes will exacerbate the skew, as insertion delays will increase.

III. Active Silicon Interposer ROA (ASI-ROA)

The proposed heterogeneous multi-die system (MDS) synchronization architecture is illustrated in Fig. 1(a). An ROA is placed on the active silicon interposer (ASI-ROA) in order to provide the clock source and distribution to the heterogeneous dies via the micro-bumps. The ROA structure of RTWOs in the proposed ASI-ROA limits jitter and skew variations since the RTWO rings on the interposer are locked in phase through self-resonance. That is, any affect of variation in lagging or leading the clock signal will be compensated by the resonance, and the system will stabilize back to its resonance. The self-resonance, on the other hand, requires algorithmic solutions that produce a balanced (L,C) load across the ROA that would

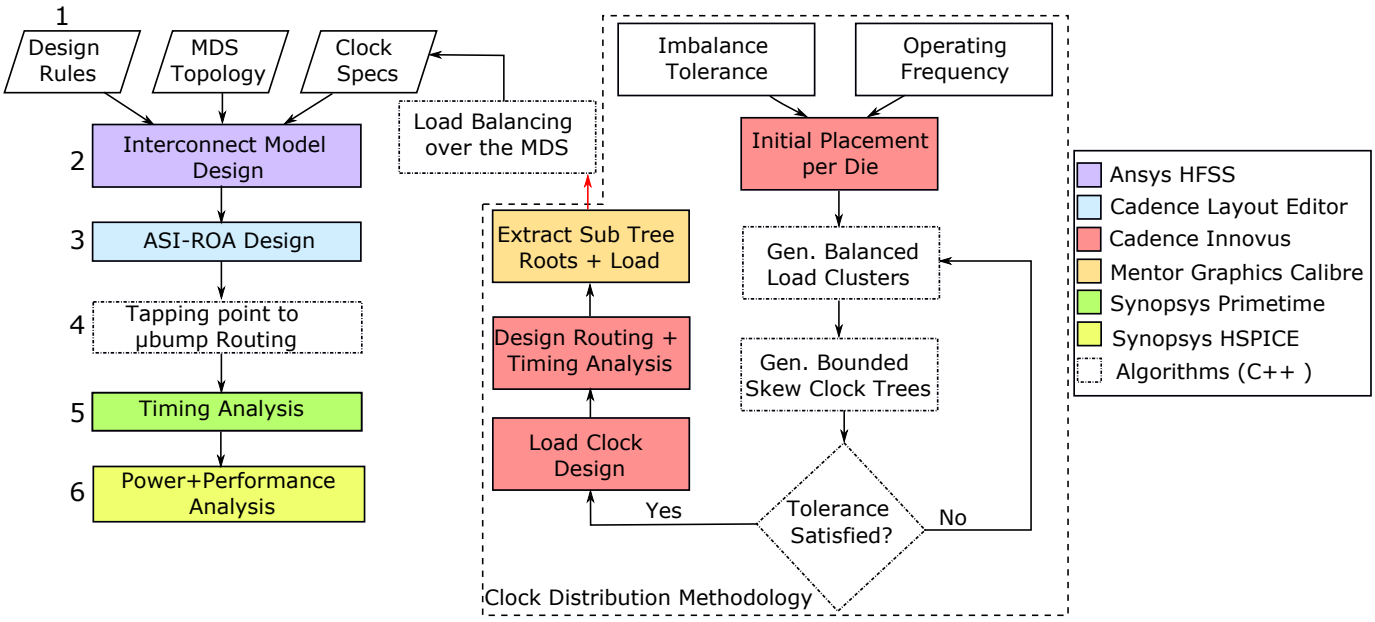

Fig. 4: Overall synchronization methodology for MDS.

establish such a stable resonant operation. The sensitivity of the ASI-ROA to capacitive imbalance is mitigated by applying ROA synthesis algorithms such as [22]. Designing custom clock distribution network provides the desired balanced load, in addition to savings in power (resonant clock source + unbuffered steiner tree), superior skew matching, and robustness to variations thanks to eliminating clock buffers. In this work, unbuffered custom clock trees are designed per die with balanced load across the ASI-ROA. In addition, dies that are hard IP and do not allow for alteration of the clock network are accommodated by delivering ASI-ROA to the external clock pins of the hard IP die, while monitoring the capacitive balance across the MDS. The following sections detail the clock synchronization architecture proposed in this work: In Section III-A, the overall synchronization methodology is presented. In Section III-B, the active silicon interposer based rotary oscillator array design is presented.

A. Overall Methodology

The methodology to synchronize the MDS with ASI-ROA is illustrated in Fig. 4. The overall methodology is annotated with the following 6 stages, and the proposed innovation is marked as the "Clock Distribution Methodology" on Fig. 4: 1) The design rules for the interposer technology, the topology of the MDS (die placement), and clock specifications (subtree roots + load) are the inputs. 2) The transmission line (TL) interconnect models for the RTWOs of ASI-ROA are modeled with a full-wave electromagnetic (EM) solver, such as Ansys High Frequency Structural Simulator (HFSS) in this work, based on the operating frequency. The TL models are used to generate SPICE compatible models. 3) The layout for the ASI-ROA are designed with a custom layout editor such as Cadence layout editor in this work. 4) The tapping points of ASI-ROA are algorithmically matched with the micro-bumps of the heterogeneous package and then, the layout is finalized with the layout editor. 5) Timing analysis for the ASI-ROA clock network of the entire MDS, with extracted models, is performed considering the skew at each node for the dies in the MDS and ASI-ROA. In experimentation, Synopsys PrimeTime is used for the static timing analysis (STA). 6) Power and performance analysis of the entire MDS are performed with a SPICE tool such as HSPICE in this work. The clock distribution methodology for each die, one of the inputs to Stage 1 above, is illustrated on the right in Fig. 4.

The clock specifications (Stage 1) requires clock frequency, micro-bump locations (*i.e.* clock source) and the load being driven by each micro-bump. To generate balanced load across the synchronized clock network for the MDS, a two step load balancing approach is proposed. Step 1: Balanced load subtree networks are designed for each die of the MDS. For hard-IP dies where clock design/distribution is not permitted or necessary, the capacitive load at the external clock pin(s) is formulated and propagated to step 2 for balancing. Step 2: Load across the ASI-ROA is balanced considering the loads of the unbalanced clock trees designed for the dies and the load(s) of the hard-IP block(s). This process is to address the capacitive balancing requirement for resonance discussed in Section II-A.

Step 1: Subtree Network Generation: Each die undergoes initial placement after which the floorplan specifications along with load information are extracted to balance the load across each core in the die. The inputs are the imbalance tolerance of the RTWOs and the target frequency. To generate a balanced load, the k-means method [23] is applied. The inputs to the k-means method are the register locations, number of subnetwork trees m per core, and the capacitance tolerance ratio τ. Each die in the MDS is modeled with m tapping points, determined heuristically at design time. The value of τ is set to 0.25 as discussed in Section II-B. The outputs are the budgeted skew and the capacitively balanced clusters

TABLE I: Balanced Assignment Problem (BAP).

min	$\max_j(\sum_i x_{ij}C_{ij}) - \min_j(\sum_i x_{ij}C_{ij})$
$s.t.$	$\sum_i x_{ij} = 1, \forall j$
	$\sum_j x_{ij} \geq 1, \forall i$
	$x_{ij} \in \{0,1\}$

TABLE II: Linear Bottleneck Assignment Problem (LBAP).

min	$\max_j(\sum_{\forall i} x_{ij}C_{ij})$
$s.t.$	$\sum_i x_{ij} = 1, \forall j$
	$\sum_j x_{ij} \geq 1, \forall i$
	$x_{ij} \in \{0,1\}$

$K_1, K_2, ..., K_m$. The total capacitance across each cluster i is estimated by

$$C_i^{tot} = \sum_{j=1}^{n_i}(C_{i,j} + C_{i,j}^{wire}). \tag{3}$$

In (3), $C_{i,j}$ is the capacitance of the register $r_{i,j}$ in the cluster K_i, where $j = 1, ..., n_i$ and n_i is the total number of registers in the cluster K_i. $C_{i,j}^{wire}$ represents the parasitic capacitance of a wire with a Manhattan distance between the register $r_{i,j}$ and the centroid cen_i of the cluster K_i. The interconnect length between the centroid and each register is used to approximate the total wire length of the cluster. C_{tot} is the total capacitance of the registers and wires.

The capacitance across all the clusters are balanced heuristically by comparing the difference between the maximum and minimum capacitance tolerance ratio τ. To avoid convergence problems, once a register is moved from a cluster, the register is not permitted to be moved again. After the clusters are generated, skew budgeted subnetwork trees are synthesized. The bounded-skew tree based on the deferred-merge and embedding (BST/DME) [24] algorithm is used to generate m unbuffered steiner trees. The local trees are unbuffered to maintain the adiabatic switching of the ASI-ROA, where capacitively balanced topologies of local unbuffered trees become part of the resonating structure. After generating the local trees, all the subtree networks in the die are compared to ensure the loads satisfy the capacitance tolerance set by the user. In cases with clusters not satisfying the condition, the clock trees are redesigned with a different tolerance ratio to obtain optimal balanced load across the die. The clock designs are then loaded into the place and route tool (P&R) to perform placement and routing and initial timing analysis for the dies. The final designs specifications required for the clock specification step in the methodology is obtained after extraction, utilizing an industrial extraction tool.

Step 2: Load Balancing over the MDS: The ASI-ROA requires balanced load across all the tapping points. The number of subtree networks and load driven by each sub-tree network in a heterogeneous MDS are different for each die. The extracted load information from Step 1 is used in this step. For hard IPs, the clock pin information and load driven by each pin is considered at this stage (detailed later in this section). The load balancing problem is formulated as balanced assignment problem of subtree roots to the micro-bump where the maximum and minimum capacitive load is minimized. The clock is delivered from the ASI-ROA to the core in each die via the micro-bump. The micro-bumps are modeled as lumped *RC* elements [25]. For a micro-bump with a pitch of 40µm and height/width of 10µm/25µm the delay incurred is $15\,\mathrm{fs}$ [25]. To balance the load, a bottleneck assignment problem as listed in Table I is considered. The binary variable x_{ij} indicates whether the sub-tree i and the micro-bump j are connected. The constraints in Table I permit the subtree root to connect to one micro-bump while each micro-bump can have one or more subtrees connected to satisfy the design objective. The balanced assignment problem (BAP) is adopted from [26].

To obtain an optimal solution for Table I, first a linear bottleneck assignment problem (LBAP) is solved as listed in Table II. The cost matrix C_{ij} is defined as the total capacitance of each cluster (sub-tree + wires). The Hungarian algorithm is modified and applied to find an assignment in the cost matrix which produces an optimal assignment (x_{ij}) of the sub-tree roots to the micro-bumps. The optimal assignment and the corresponding capacitance values are then iterated to decrease the difference between the maximum and minimum capacitance values while simultaneously checking for an assignment between the maximum and minimum value. The algorithm returns the optimal balanced assignment such that the difference between the max and min capacitance values is minimized. Therefore, the optimal capacitive load balanced solution of connecting the sub-tree roots to the tapping points is obtained that also satisfies the skew requirement without the excessive use of tapping wires. The overall time complexity of the clock distribution methodology presented in this work that includes balancing load clusters and clock tree synthesis (CTS) is $O(N^4 + |S|^3)$ where N is the number of rings and S is the set of registers. Even though the complexity approaches $O(N^4)$, the number of rings, N, is small, resulting in an average runtime that is closer to $O(|S|^2)$ due to memory optimizations lookups in the CTS algorithm.

IP reuse: 2.5D systems, by design, heavily leverage IP reuse and off-the-shelf IPs. The proposed clock synchronization methodology can be leveraged for IPs (reuse) at the load balancing stage where the clock pin and load information from the hard IPs can be used to balance the load across the ASI-ROA. An IP pre-designed with a traditional clock distribution network and implemented within the standard ASIC flow is considered as a hard IP. At the load balancing stage (Step 2), the load driven by the hard IP is considered along with the loads driven by the dies designed with the proposed methodology. Redesigning the hard IP would not be feasible if the tolerance set by the bottleneck assignment cannot be achieved. To balance the load, on-chip capacitors, only if necessary, are proposed on the RTWO rings in the active silicon interposer for additional load. The on-chip capacitor value is chosen based on the load imbalance in the cost function. Active silicon interposers traditionally have large die sizes and placing additional capacitors will not adversely affect the MDS operation ($10\,\mathrm{pF}$ MOSCAP estimated die area is $0.1\,\mathrm{mm}^2$) or increase the active portion of the interposer.

Algorithm 1: Generation of tapping point set

Input: Tapping point set S and micro-bump locations Bu
Output: Optimized tapping point set S_{opt}, optimal matching record F_{opt}, best transformation vector t_{opt}

Initialize $cost = \infty$, $cost_{new} = 0$;
while $|cost - cost_{new}| > \triangle$ **do**
  $[F, cost_{new}] = best\_matching(S, Bu)$;
  $\mathbf{t} = best\_move(S, Bu, F)$;
  $cost = cost_{new}$
  $cost_{new} = COST(F, S, Bu, t)$
  $S = S + \mathbf{t}$;
end while
$S_{opt} = S, F_{opt} = F, \mathbf{t}_{opt} = \mathbf{t}$;

B. ASI-ROA Design

The performance metrics of the proposed ASI-ROA of RTWOs are characterized using a full-wave EM solver, Ansys High Frequency Structural Simulator (HFSS) [27] and SPICE. Depending on the frequency of operation, transmission line (TL) models are obtained using S-parameters and converted to SPICE compatible models. After obtaining the TL models, layout for the RTWO are designed in the active silicon interposer technology node. The tapping point locations for the same phase of the ASI-ROA is then determined after simulating the post layout models, simulated with SPICE.

The resonant clock delivery from the interposer to the MDS requires capacitance and delay matching between the clock source and the micro-bumps. The optimized micro-bump locations and ASI-ROA tapping point set S are the inputs for Algorithm 1. The delay is matched during the load balancing step, thanks to which the micro-bump locations along with the load information is sufficient at this stage. The objective of the optimization process is to minimize the total delay between the micro-bump and the tapping. To ensure the total delay is matched across the ASI-ROA, a process based on the earth movers distance approach (EMD), which is a measure of distance between discrete and finite distributions, is used [15]. The $EMD_{\mathbf{t}}$ approach [15] is modified to account for the locations and parasitics of the micro-bump in this work. The $EMD_{\mathbf{t}}$ shifts the square shaped ASI-ROA to generate a one-to-one matching between the micro-bump and the tapping points, which results in a balanced delay among the subnetwork trees connected through the $bumps$. For any transformation $\mathbf{t} \in \mathbf{R}^2$, the tapping points on the RTWO rings of the ASI-ROA are shifted by the same transformation $\mathbf{t}$:

$$S + \mathbf{t} = \{s_1 + \mathbf{t}, s_2 + \mathbf{t}, ..., s_m + \mathbf{t}\}. \quad (4)$$

The function $COST$ in Algorithm 1 is to calculate the delay from $s_i + \mathbf{t}$ to the load of a micro-bump, d_j. The problem is formulated as $F = (f_{ij}) \in d_j$, where f_{ij} = 1 indicates a one-to-one matching from $s_i + \mathbf{t}$ to micro-bump d_j. The Elmore delay model for the cost function is:

$$t(s_i + \mathbf{t}, d_j) \approx r \cdot Cap(d_j). \quad (5)$$

The total summed delay represents the delay from the tapping points to each micro-bump (d_j). Algorithm 1 is employed to find the best match and perform best move of the ASI-ROA to minimize the cost function. Considering the routing wires are either horizontal or vertical, Manhattan routing is used to calculate the total minimum delay. Algorithm 1 terminates when the predefined tolerance set by $\triangle$ is achieved resulting in the optimal transformation $\mathbf{t}_{opt}$ and the best matching result F_{opt}. The EMD_t converges quickly for the problems investigated in this work, as the number of tapping points and rings are small.

The extracted layout views of the ASI-ROA along with the MDS load are first simulated to obtain the skew across the entire MDS. SPICE based simulations are employed to evaluate the power and performance metrics for the MDS. The skew values are then translated into Synopsys PrimeTime to perform timing analysis for the MDS.

C. Pre-bond Testing of the Proposed Architecture

Prior to bonding the dies to the interposer it is important to ensure the clock signals have the desired performance. In [28], authors utilize multiple through silicon vias (TSV) to reduce the overall wirelength between the clock source locations per sub-tree such that a bounded skew with minimum wirelength is achieved across the 3D clock tree. The multiple TSVs are used as redundant clock tapping points to enable pre-bond testing in [28] where the entire clock tree is not connected. In [29], redundant trees with transmission gates are designed to enable pre-bond testing and then to be turned off to synchronize the clock trees.

In this work, the 2D clock tree designed per die and the clock sub-tree locations are considered during the load balancing stage to generate one-to-one matching between the clock pins (tapping points) and the sub-tree root. Unlike CTS for 3D ICs, the RTWOs designed in this work are partitioned at zero skew points, between the regional and local trees. The clock signals can be probed at the μ_{bump} locations to ensure the desired performance can be achieved. The proposed architecture can be implemented for MDS with hard and soft IPs. For the hard IP based solution, the clock pin for each die can be tested prior to bonding. The clock source on the interposer can be probed at the μ_{bump} location to ensure a reliable clock signal is available. The challenges associated with pre-bond testing in terms of emulating the resonance between the clock trees and the interposer rings, need to be addressed carefully.

IV. Architecture Evaluation

In the following section the evaluation specifications and setup, variation analysis and power consumption results are presented for the proposed architecture.

A. Evaluation Setup

The evaluation is performed over two arbitrarily selected interposer die size as listed in Table III: 1) A $10\,\mathrm{mm} \times 6\,\mathrm{mm}$ MDS with 4 dies, and 2) a $10\,\mathrm{mm} \times 12\,\mathrm{mm}$ MDS with 5 dies. Each MDS has non-uniform RISC-V cores per die, listed in Table III. This non-uniformity of number of cores is selected

TABLE III: Design parameters for the analysis.

Category	10 mm×6 mm	10 mm×12 mm
# die	4	5
# cores in each die	4, 8, 4, 6	4, 8, 8, 6, 8
RISC-V dimension	1.5 mm×1.2 mm	1.5 mm×1.2 mm
V_{dd}	1.2 V	1.2 V
Nom. Temp.	25°C	25°C
μ_{bump} pitch	40 µm	40 µm
μ_{bump} height/width	10 µm/5 µm	10 µm/5 µm
RTWO W_{tline}	20 µm	20 µm
RTWO Sep_{tline}	25 µm	25 µm

TABLE IV: Parameters for PVT variation.

Category	Var. range (28 nm)	Var. range (65 nm)	Nom. RTWO (65 nm)
V_{dd}	± 10%	± 10%	1.2 V
W_{tline}	–	± 10%	20 µm
Sep_{tline}	–	± 10%	25 µm
Temp.	– 40°C to 125°C	– 40°C to 125°C	25°C
σ_{V_t}	60 mV	60 mV	–

in order to demonstrate that the proposed method is not limited to uniform orientation of MDS cores, for which regular topologies of clock networks would suffice, but the proposed method is applicable to non-regular topologies, as well. The active silicon interposer is designed with a 65 nm industrial technology node with 7 metal layers and a 28 nm industrial technology node with 7 metal layers is used for the multiple die. The skew budget for the analysis is aggressively set to 10 ps. The transmission lines for the RTWOs are modeled with Ansys High Frequency Structural Simulator (HFSS) [27] with the top two metal layers in the 65 nm technology node. The die and interposer are placed and routed with Cadence Innovus. Mentor Graphics Calibre is utilized to extract the parasitic impedances of the circuits after place and route. The micro-bumps have a pitch of 40 µm and are modeled as lumped *RC* elements [25]. SPICE based analysis is performed on post-layout models to more accurately characterize the ASI-ROA. The design parameters for the SPICE analysis are listed in Table III.

An ASI-ROA operating at 2 GHz is designed with a perimeter length of 1200 µm per individual RTWO ring. The ASI-ROA is designed with 24 inverter pairs for each RTWO ring to generate the operating frequency of 2 GHz. For the interposer die size of 10 mm×12 mm, the active area occupied by the ASI-ROA is 3.1% of the total interposer area. The MDS and the active silicon interposer operate at the same voltage ($V_{dd} = 1.2$ V). The results are presented to investigate the three main contributions of this work: 1) the robustness of ASI-ROA and 2) synchronization of the MDS, and 3) power consumption of the MDS designed with the ASI-ROA.

B. Variation Analysis

Three different variation profiles are considered 1) within die (WID), 2) die-to-die (D2D), and 3) die-to-interposer (DIP) for the clock generation and delivery to the multi-die system. The parameters included for the variation analysis in the die and interposer are listed in Table IV, similar to [14]. The design parameters listed in Table IV are varied for 500 Monte-Carlo simulations on post-layout models of the multi-die system. The σ_{V_t} is varied by 60 mV along with ±10%

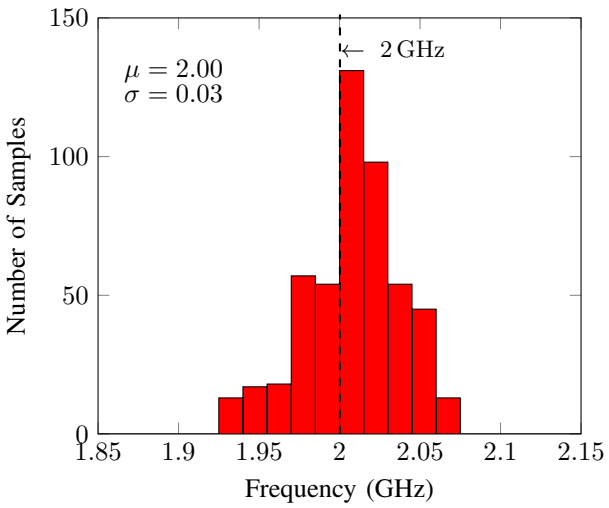

Fig. 5: Post layout simulated frequency of the ASI-ROAs distributed over 10 mm×12 mm synchronizing the 5 dies in a MDS. ASI-ROA frequency design target is 2 GHz.

TABLE V: RTWO frequency mismatch between fastest and slowest RTWO rings distributed in an ROA topology across PVT corners for a target RTWO frequency of 2 GHz.

Die dimension	SF	SS	TT	FS	FF
10 mm × 12 mm	0.04%	0.09%	0.001%	0.07%	0.01%

variations in the geometries of the RTWOs and supply voltage. The variations of ±10% is chosen based on the foundry rules provided to us based on all the possible corners for the process nodes. The temperature is varied from −40 °C to 125 °C for the analysis. The frequency variation of the ASI-ROA designed for a 10 mm×12 mm die size across the MDS consisting of 5 dies is shown in Fig. 5. The largest frequency variation observed in the ASI-ROA is 80 MHz which is a 4% variation from the target frequency of 2 GHz. This variation is for all the RTWOs on an ASI-ROA with respect to the nominal 2 GHz frequency, e.g. an ASI-ROA with 24 RTWOs would have all 24 RTWOs at 1.920 MHz for the worst case. The variation between the individual RTWOs for an overall ASI-ROA frequency (e.g. @2 GHz nominal) is discussed next.

The transmission lines of the RTWO structures distributed over the active silicon interposer are interlocked at the corners which mitigate the variations in σ_{V_t} and transmission line geometries thanks to the resonating structure of the ASI-ROA and the algorithmic load balancing across the MDS. In Table V, the frequency mismatch between the slowest and fastest RTWO rings distributed in an ASI-ROA topology over the larger interposer die dimensions is listed across PVT corners. The post-layout models are simulated across the different corners and it can be observed that the worst case variation between the RTWOs is ≈0.1% (*i.e.* ≈2 MHz) at the SS corner. The ASI-ROA at the TT corner has a frequency mismatch of 0.001% (*i.e* 0.02 MHz). RTWOs, through self-resonance, self-tuning, and in-phase locking across rotary rings in an ROA topology ensure the frequency is matched across the entire network. A local variation affecting a local RTWO ring, therefore, gets compensated across the ROA, as each local RTWO ring is phase locked.

TABLE VI: Power consumption of ADPLL clocked circuit as compared to the proposed ASI-ROA distributed on an active silicon circuit operating at a frequency of 2 GHz, a V_{dd} of 1.2 V, and temperature of 25 °C.

Die dimension	# dies	ADPLL clocked design			This Work: ASI-ROA				
		ADPLL power (mW)	Total power (mW)	Skew (ps)	ASI-ROA power (mW)	Total power (mW)	Skew (ps)	# RTWO rings	Active area
10 mm × 6 mm	4	443.9	1400.3	7.2	150.7 (-66%)	870.3 (-38%)	8.4	132	2%
10 mm × 12 mm	5	1025.9	2536.1	7.2	282.9 (-72%)	1395.4 (-45%)	8.0	204	3.1%
Average	-	-	-	-	-69%	-41%	8.2	-	-

TABLE VII: Power consumption of ADPLL clocked circuit as compared to the proposed ASI-ROA distributed on an active silicon circuit operating at a frequency of 2 GHz, a V_{dd} of 1.2 V, and temperature of 25 °C. One die designed with buffered clock tree with Cadence Innovus.

Die dimension	# dies	ADPLL clocked design			This Work: ASI-ROA					
		ADPLL power (mW)	Total power (mW)	Skew (ps)	ASI-ROA power (mW)	Total power (mW)	Skew (ps)	# RTWO rings	Active area	Balancing capacitor
10 mm × 6 mm	4	443.9	1400.3	7.2	300.5 (-32%)	1020.3 (-27%)	8.4	84	1.2%	-
10 mm × 12 mm	5	1025.9	2536.1	7.2	506.9 (-50%)	1619.1 (-36%)	8.0	156	2.4%	2×10 pF
Average	-	-	-	-	-41%	-31%	8.2	-	-	-

C. Power Consumption and Clock Skew

The power consumption and skew of the ASI-ROA as compared to an ADPLL operating at 2 GHz are listed in Table VI, Table VII, and Table VIII. In Table VI, the power consumption of the ASI-ROA designed with the methodology proposed in this work is listed *i.e.,* all the dies in the MDS have unbuffered steiner clock trees. In Table VII, the power consumption of the ASI-ROA designed with the methodology proposed in this work along with the largest die designed with a standard ASIC flow (hard IP) is listed. In Table VIII, the power consumption of the ASI-ROA designed with the methodology proposed in this work with all dies designed with a standard ASIC flow (hard IP) is listed. The results presented in Table VIII are for a MDS with ASI-ROA serving as a synchronization source only *i.e.,* without the proposed clock distribution network. Standard cell library for the 28 nm technology node is used to implement the ADPLL based designs and placed and routed with Cadence Innovus. The ADPLL is reimplemented from [30] to operate at 2 GHz and occupies an area of $0.005\,\mathrm{mm}^2$ while consuming 7.5 mW power with a supply voltage of 1.2 V. Each die in the multi-die module includes one ADPLL to synchronize the RISC-V cores.

A total power savings of approximately 45% is achieved in post-layout simulation for the multi-die system when compared against an ADPLL clocked multi-die system for the 10 mm×12 mm MDS. The total power savings are achieved thanks to the clock power savings of 72% when compared to the ADPLL based designs for the same die size. The active area footprint of the inverter pairs is 3.1% of the total silicon interposer area for the 204 RTWO rings in the ASI-ROA topology. The adiabatic nature of the ASI-ROA along with the unbuffered steiner trees that deliver the clock to the dies maintain the adiabatic property of the clock source which provide significant savings in power. The average clock skew is 8.2 ps when designed with a clock skew constraint of 10 ps.

The slew of the clock buffers at the last stage of the clock tree have a high impact. The clock tree designed with the buffered clock trees + ADPLL clock source has a slew of 50 ps. This slew is 10% of the clock period considered in this work which is the norm. The clock slew of the unbuffered steiner clock tree is 60ps, which is 12% of the clock period. Larger clock slews have detrimental effects resulting in an increase in the clk-to-q delay and power consumption. Slew-aware methodologies, such as those in [19, 20, 31] for buffered trees (i.e. slew-aware steiner trees), can be developed for unbuffered trees as well with continued research. Unlike those in [19, 20, 31], the solution will need to rely on innovations in ROA routing, rather than merging and buffering innovations for slew-awareness. The ASI-ROA provides an average of 41% total power savings when compared to the ADPLL based design operating at 2 GHz, thanks to the 69% reduction in the power consumption of the clock for the two die sizes listed in Table VI. The proposed architecture delivers the clock signals across the MDS with a 2% increase in slew while providing ≈41% overall power savings for the MDS with an entirely resonant (clock source + tree) clock generation and distribution solution.

ASI-ROA for IP reuse: In Table VII, the power consumption of the ASI-ROA as compared to ADPLL is presented for designs with the largest die considered to be a hard IP. One largest die is chosen, and placed and routed with a standard ASIC flow utilizing Cadence Innovus (bounded skew buffered clock trees). The clock pins for largest die with the load information is considered at the load balancing stage (Step 2 in Section III) along with the sub-tree roots and load information for the other dies designed with the methodology proposed in this work. The overall power savings drop to 27% (11% reduction) for the die size of 10 mm×6 mm since the largest die is designed with buffered clock trees and the clock distribution network does not maintain the adiabatic property of the clock source. The clock power savings drop to 32% (34% reduction) when compared to the ASI-ROA design with the entire clock network resonant for a 10 mm×6 mm die size. The buffered portion of the clock networks do not resonate with the rest of the ASI-ROA, due to which savings in power diminish.

For the larger die size 10 mm×12 mm, with the reported experimental setup of technologies and cores on dies, the load balancing stage does not converge. As explained in Section III,

TABLE VIII: Power consumption of ADPLL clocked circuit as compared to the proposed ASI-ROA distributed on an active silicon circuit operating at a frequency of 2 GHz, a V_{dd} of 1.2 V, and temperature of 25 °C. All dies implemented with buffered clock trees with Cadence Innovus.

Die dimension	# dies	ADPLL clocked design			This Work: ASI-ROA					
		ADPLL power (mW)	Total power (mW)	Skew (ps)	ASI-ROA power (mW)	Total power (mW)	Skew (ps)	# RTWO rings	Active area	Balancing capacitor
10 mm × 6 mm	4	443.9	1400.3	7.2	415.9 (-6%)	1387.8 (-0.9%)	7.2	4	<0.01%	1×5 pF
10 mm × 12 mm	5	1025.9	2536.1	7.2	986.9 (-4%)	2521.9 (-0.6%)	7.2	5	<0.01%	2×30 pF
Average	-	-	-	-	-5%	-0.75%	7.2	-	-	-

this is due to the hard IP die not enabling the design of the clock network on the die, and instead, the die clock network being provided as an input to the load balancing process. With some numerical values, load balancing within the expected balance margin (<25%) cannot be achieved. To address this, capacitors are placed on the rotary rings which require load balancing. The load imbalance percentage is determined and then capacitors are placed on the active silicon interposer to balance the load. For the larger die size of $10\,\mathrm{mm}\times12\,\mathrm{mm}$ two $10\,\mathrm{pF}$ capacitors are placed to balance the load, determined in Step 2 (Section III). A total power saving of 36% is achieved thanks to 50% savings in the clock power. Designing the heterogeneous MDS with dies designed with the proposed methodology and using a hard IP achieves an average power saving of 31% when compared to the ADPLL based design. The savings in power is achieved with an average skew of $8.2\,\mathrm{ps}$. The active silicon interposer has a large die area ($120\,\mathrm{mm}^2$) with minimal footprint of the ASI-ROA (occupies 2.4% active area of the active silicon interposer), which makes it tolerable to place a capacitor to balance the load to save design time.

ASI-ROA for hard IP: It must be noted that, all dies in a MDS can be treated as hard IPs. In such a scenario of 100% IP reuse, ASI-ROA can still be used to synchronize the MDS, however the savings in power will reduce as the conventional clock trees are designed with clock buffers making the distribution network non-resonant. State of the art industrial designs employ buffered or mesh based clock distribution networks that are readily implementable with industrial tools. However, buffered clock trees consume significant power while providing reliable clocks signals for the designs. Designing unbuffered steiner trees require dedicated algorithms and tool flows to enable the power savings that can be achieved with resonant rotary clocks. In Table VIII, the power consumption of the ASI-ROA as compared to ADPLL is presented for all the dies considered to be a hard IP (buffered clock trees + resonant clock source). All the dies are placed and router with a standard ASIC flow utilizing Cadence Innovus (bounded skew buffered trees). The clock pins for all the dies with the load information is considered at the load balancing stage (Step 2 in Section III). In a completely hard IP solution, the top level clock buffer is the load that needs to be driven by each tapping point of the ASI-ROA. Each die in the MDS is driven by a one RTWO ring in the ASI-ROA topology, due to this there is a significant reduction in the active area footprint. For the largest MDS of size $10\,\mathrm{mm}\times12\,\mathrm{mm}$ the total power saving of 0.6% is achieved when compared to the ADPLL based solution. The clock power savings thanks to the resonant clock source in the active silicon interposer is 4% when compared to the ADPLL based designs operating at 2 GHz. The $10\,\mathrm{mm}\times12\,\mathrm{mm}$ MDS requires two $30\,\mathrm{pF}$ capacitors to balance the load. The clock signals delivered to the hard IPs in the MDS have the same rise and fall times across the die.

V. Conclusion

Synchronization of multi-die system with rotary traveling wave oscillators is presented in this work. The proposed architecture is implemented within ASIC physical design flows on two different die dimensions and evaluated using post-layout extracted view SPICE simulations for robustness and power consumption. The proposed methodology is leveraged for IP reuse providing a feasible solution for load imbalance mitigation. For the largest die area the proposed methodology provides a total power savings of approximately 41% for the MDS implemented with the proposed methodology when compared to an ADPLL based design. The average clock skew of the completely resonant architecture presented in this work is $8.2\,\mathrm{ps}$. For the largest die area the proposed methodology along with the hard IP provides 31% power savings when compared to an ADPLL based design. Furthermore, the ASI-ROA has an active area footprint of only 3.1% over a $10\,\mathrm{mm}\times12\,\mathrm{mm}$ active silicon interposer die size for a completely resonant solution. In addition, the ASI-ROA is leveraged to provide a synchronous clocking solution for the MDS with all the dies considered as hard IPs.

References

[1] A. Arunkumar *et al.*, "Mcm-gpu: Multi-chip-module gpus for continued performance scalability," in *Proc. ACM Int. Symposium on Computer Architecture (ISCA)*, 2017, pp. 320–332.

[2] R. Mahajan *et al.*, "Embedded multi-die interconnect bridge (EMIB) – a high density, high bandwidth packaging interconnect," in *Proc. IEEE Electronic Components and Technology Conference (ECTC)*, May 2016, pp. 557–565.

[3] N. E. Jerger *et al.*, "NoC architectures for silicon interposer systems: Why pay for more wires when you can get them (from your interposer) for free?" in *Proc. IEEE/ACM Int. Symp. on Microarchitecture (MICRO)*, Dec. 2014, pp. 458–470.

[4] W. Gomes *et al.*, "8.1 lakefield and mobility compute: A 3D stacked 10nm and 22FFL hybrid processor system in 12x12mm^2, 1mm package-on-package," in *IEEE Int. Solid- State Circuits Conference - (ISSCC)*, 2020, pp. 144–146.

[5] Marvell Corp., https://www.marvell.com/architecture/mochi/, 2018.

[6] A. Kannan, N. E. Jerger, and G. H. Loh, "Enabling interposer-based disintegration of multi-core processors," in *IEEE/ACM Int. Symp. on Microarchitecture (MICRO)*, Dec. 2015, pp. 546–558.

[7] D. Stow *et al.*, "Cost-effective design of scalable high-performance systems using active and passive interposers," in *Proc. IEEE/ACM Int. Conf. on Computer-Aided Design (ICCAD)*, Nov. 2017, pp. 728–735.

[8] J. Kim, "Active Si interposer for 3D IC integrations," in *Proc. Int. 3D Systems Integration Conf. (3DIC)*, Aug. 2015, pp. TS11.1.1–TS11.1.3.

[9] P. Vivet *et al.*, "8.1 a 4x4x2 homogeneous scalable 3d network-on-chip circuit with 326mflit/s 0.66pj/b robust and fault-tolerant asynchronous 3d links," in *IEEE Int. Solid-State Circuits Conf. (ISSCC)*, Jan. 2016, pp. 146–147.

[10] ——, "IntAct: A 96-core processor with six chiplets 3D-stacked on an active interposer with distributed interconnects and integrated power management," *IEEE Journ. of Solid-State Circuits (JSSC)*, vol. 56, no. 1, pp. 79–97, 2021.

[11] Y. S. Shao *et al.*, "Simba: Scaling deep-learning inference with multi-chip-module-based architecture," in *Proc. IEEE/ACM Int. Symp. on Microarchitecture (MICRO)*, 2019.

[12] W. J. Turner *et al.*, "Ground-referenced signaling for intra-chip and short-reach chip-to-chip interconnects," in *Proc. IEEE Custom Integrated Circuits Conference (CICC)*, April 2018, pp. 1–8.

[13] J. Wood, T. C. Edwards, and S. Lipa, "Rotary traveling-wave oscillator arrays: a new clock technology," *IEEE Journ. of Solid-State Circuits (JSSC)*, vol. 36, no. 11, pp. 1654–1665, Nov. 2001.

[14] R. Kuttappa *et al.*, "Robust low power clock synchronization for multi-die systems," in *Proc. IEEE/ACM Int. Symp. on Low Power Electronics and Design (ISLPED)*, July 2019, pp. 1–6.

[15] Y. Teng and B. Taskin, "Sparse-rotary oscillator array (SROA) design for power and skew reduction," in *Proc. Design, Automation and Test in Europe Conf. Exhibition (DATE)*, Mar. 2013, pp. 1229–1234.

[16] H. Jang, D. Joo, and T. Kim, "Buffer sizing and polarity assignment in clock tree synthesis for power/ground noise minimization," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 30, no. 1, pp. 96–109, Jan 2011.

[17] J. Lu, X. Mao, and B. Taskin, "Clock mesh synthesis with gated local trees and activity driven register clustering," in *Proc. Int. Conference on Computer-Aided Design (ICCAD)*, Nov. 2012, pp. 691–697.

[18] J. Lu, Y. Teng, and B. Taskin, "A reconfigurable clock polarity assignment flow for clock gated designs," *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, vol. 20, no. 6, pp. 1002–1011, June 2012.

[19] C. Sitik *et al.*, "Design methodology for voltage-scaled clock distribution networks," *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, vol. 24, no. 10, pp. 3080–3093, Oct 2016.

[20] W. Liu *et al.*, "Slects: Slew-driven clock tree synthesis," *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 4, pp. 864–874, April 2019.

[21] R. Islam and M. R. Guthaus, "Hcdn: Hybrid-mode clock distribution networks," *IEEE Trans. on Circuits and Systems I: Regular Papers*, vol. 66, no. 1, pp. 251–262, Jan 2019.

[22] R. Kuttappa *et al.*, "RotaSYN: Rotary Traveling Wave Oscillator SYNthesizer," *IEEE Trans. on Circuits and Systems I: Regular Papers*, vol. 66, no. 7, pp. 2685–2698, July 2019.

[23] R. He *et al.*, "Balanced k-means algorithm for partitioning areas in large-scale vehicle routing problem," in *Proc. of the Int. Symp. on Intelligent Information Technology Application (IITA)*, vol. 3, Nov 2009, pp. 87–90.

[24] K. D. Boese and A. B. Kahng, "Zero-skew clock routing trees with minimum wirelength," in *Proc. IEEE Int. ASIC Conf. and Exhibit (ASIC)*, Sep. 1992, pp. 17–21.

[25] S. M. Alam *et al.*, "Inter-strata connection characteristics and signal transmission in three-dimensional (3D) integration technology," in *Proc. Int. Symp. on Quality Electronic Design (ISQED)*, Mar. 2007, pp. 580–585.

[26] R. E. Burkard and E. Cela, *Linear Assignment Problems and Extensions*. Boston, MA: Springer US, 1999, pp. 75–149.

[27] "High frequency structural simulator (HFSS): User's guide, ansoft corporation," 2018.

[28] X. Zhao *et al.*, "Pre-bond testable low-power clock tree design for 3d stacked ics," in *Proc. IEEE/ACM Int. Conference on Computer-Aided Design (ICCAD)*, 2009, pp. 184–190.

[29] T. Kim and T. Kim, "Clock tree synthesis with pre-bond testability for 3d stacked ic designs," in *Proc. Design Automation Conference (DAC)*, 2010, pp. 723–728.

[30] M. Lee *et al.*, "A 0.3-to-1.2V frequency-scalable fractional-N ADPLL with a speculative dual-referenced interpolating TDC," in *IEEE Int. Solid - State Circuits Conference - (ISSCC)*, 2018, pp. 122–124.

[31] C. Sitik *et al.*, "Finfet-based low-swing clocking," *ACM Journ. of Emerging Technologies in Computing Systems (JETC).*, vol. 12, no. 2, pp. 13:1–13:20, Sept. 2015.

Ragh Kuttappa (S'15) received the Bachelor's degree from Visvesvaraya Technological University, India, in 2012 and the Master's degree from San Francisco State University, San Francisco, CA, in 2015. He is currently pursuing the Ph.D. degree from the Department of Electrical and Computer Engineering, Drexel University, Philadelphia, PA. His current research interests include electronic design automation for VLSI, low-power circuits and resonant clocking. He has held an intern position at Samsung Austin Research Center (SARC) in 2017.

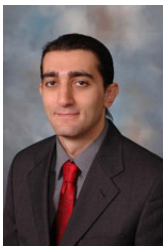

Baris Taskin (S'01-M'05-SM'12) received the B.S. degree in electrical and electronics engineering from Middle East Technical University, Ankara, Turkey, in 2000, and the M.S. and Ph.D. degrees in electrical engineering from the University of Pittsburgh, Pittsburgh, PA, in 2003 and 2005, respectively.

Dr. Taskin is currently a Professor of Electrical and Computer Engineering at Drexel University, Philadelphia, PA. His current research interests include electronic design automation for VLSI, low-power circuits, resonant clocking, clock network synthesis, wireless IC interconnects, and networks-on-chip for chip multiprocessors.

Dr. Taskin is a recipient of a number of awards for his research and professional contributions, including the National Science Foundation Faculty Early Career Development (NSF CAREER) Award in 2009, an ACM SIGDA Distinguished Service Award in 2012 and the Delaware Valley Young Electrical Engineer of the Year Award from the IEEE Philadelphia Section in 2013 and the Drexel ECE Department's Outstanding Research Award in 2015. Dr. Taskin was the General Chair for ACM Great Lakes Symposium on VLSI (GLSVLSI) 2019 and was the Chair of the IEEE Circuits and Systems Society's Technical Committee on VLSI Systems and Applications (IEEE CAS VSA-TC) 2018-2020.

Scott Lerner (S'09) received the B.S. degrees in electrical and computer engineering from Drexel University, Philadelphia, PA, USA in 2015. He is currently a PhD student in the electrical engineering department at Drexel University.

He has a varied background of work experience from working at Advanced Micro Devices in Austin, TX on low-power cl ock tree synthesis, to DRAM product engineer at Micron Technologies in Boise, ID, to RF spectrum denial at Lockheed Martin in Cherry Hill, NJ. He was a recipient of the National Science Foundation Graduate Research Fellowship Program (GRFP) and National Defense Science and Engineering Graduate Fellowship (NDSEG) in 2015. His current research interests include low-power design, clock tree synthesis, and design for reliability.

Vasil Pano (S'16) received his B.S degree and PhD degrees from Drexel University, Philadelphia, PA, in Computer (B.S degree) and Electrical (PhD degree) Engineering in 2014 and 2019, respectively.

He is currently a Post Doctoral Researcher at Drexel University, Philadelphia, PA. Dr. Pano has held an graduate intern position at Intel Corporation, Hillsboro, OR in their Data Center Group working towards a scalable network interconnect for exa-scale architectures. His research focuses on network-on-chips, wireless on-chip interconnects, and multi-die systems.