

Scalable Power-Efficient Kilo-Core Photonic-Wireless NoC Architectures

Avinash Kodi, Kyle Shiflett, Savas Kaya and Soumyasanta Laha
Department of Electrical Engineering and Computer Science
Ohio University, Athens, Ohio 45701
Email: kodi@ohio.edu, kaya@ohio.edu, laha@ohio.edu

Ahmed Louri
Department of Electrical and Computer Engineering
George Washington University, Washington DC 20052
Email: louri@gwu.edu

Abstract—As technology scales, hundreds and thousands of cores are being integrated on a single-chip. Since metallic interconnects may not scale effectively to support thousands of cores, architects have proposed emerging technologies such as photonics and wireless for intra-chip communication. While photonics technology is limited by the complexity and thermal effects, wireless technology for on-chip communication is limited by the available bandwidth. In this paper, we combine the benefits of both technologies into novel architecture that takes advantage of the communication benefits of both technologies while circumventing their limits. We discuss the scalability of the proposed architecture to kilo-core system using wireless technology. We evaluate the power consumption, throughput and latency for 256 and 1024 core architectures when compared to photonics-only, wireless-wired, wireless-photonics and wired-only architectures on synthetic traffic traces. Our simulation results indicate that the proposed architecture and design methodology can have significant impact on the overall network power and performance.

Keywords—network-on-chip, emerging technology, wireless, photonics, performance analysis

I. INTRODUCTION

Technology scaling has enabled integrating hundreds of homogeneous and heterogeneous cores within a single chip. Several commercial and academic chips have integrated hundreds and even thousands of cores such as Kalray-256 MPPA, kilocore from UC Davis [1], NVIDIA GTX1080 and several others. Aggressive scaling the number of cores has continued to disrupt the design of energy-efficient on-chip communication fabric since data movement between the processing cores and the memory hierarchy becomes critical. According to the International Technology Roadmap for Semiconductors (ITRS), the development of traditional metallic interconnects would not be sufficient to support the growing number of multicores as metallic interconnects do not scale due to the increased energy and multi-hop requirements [2]. Emerging interconnects technologies such as photonics and wireless are under serious consideration to overcome the challenges stated above.

Photonic interconnects offers several advantages over metallic interconnects such as distance-independent energy consumption particularly for short intra-chip distances, higher bandwidth-density due to wavelength-division multiplexing (WDM) and CMOS compatibility [3], [4]. While several work have proposed photonic technology for on-chip network, there are several hurdles for implementing

such architectures. First, mitigating thermal and parametric variations with exceedingly large number of components for kilo-core architectures is difficult. For example, a 64×64 crossbar using photonics will require 448 modulators, 7 waveguides and 28224 photodetectors using single-writer multiple-reader (SWMR). If we scale to 1024×1024 , then we will need approximately 7168 modulators, 112 waveguides, and 7.3 million photodetectors which is prohibitive and not easily scalable to mitigate thermal variations. Second, network latency and insertion losses tend to increase with either a long snake-like waveguide (single crossbar) or with a multi-hop network (decomposed crossbar). Therefore, while photonic networks are extremely energy-efficient, design and implementation of photonic interconnect layers are much more complex for scalable multicore architectures.

Wireless technology offers several advantages over the metallic technology such as (1) distance independent one-hop communication, (2) lower energy requirement compared to a long metallic link, (3) multicasting and broadcasting with omnidirectionality, and (4) absence of any physical channels. However, on-chip wireless technology has limited bandwidth at 60 GHz center frequency and is not energy efficient at shorter distances. Many of the current efforts for chip-to-chip communications have focused on the millimeter wave bands, and the initial results have exploited the growing technology base in the 30-100 GHz range [5], [6], [7]. Hence, to overcome limited bandwidths, metallic interconnects are used for short distance communications whereas wireless interconnects are used for long distance communications using frequency division multiplexing (FDM), time division multiplexing (TDM), and space division multiplexing (SDM).

In our prior work, we evaluated OWN (Optical-Wireless NoC) architecture that combined the best of photonics and wireless technologies by overcoming the complexity of photonics and limited bandwidth of wireless [8]. While the prior work focused on OWN architecture (connectivity, routing), transceiver design and power-efficiency to achieve the high wireless bandwidth were not considered. In prior work, we did not identify how the wireless bandwidth will be achieved and what technology will be used to achieve the high wireless bandwidth. Moreover, optimistic energy/bit across the entire wireless spectrum was assumed which was unrealistic. Further, wireless channel allocation and

implementation for 256 and 1024 cores were not completely analyzed with different technologies.

In this paper, we extend the energy-efficiency analysis by projecting ideal and conservative wireless energy-efficiency for 256 and 1024 core architectures. First, we discuss the architecture and on-chip communication for combining two diverse technologies into an integrated platform that can scale to large number of cores. While prior work considered 256 and 1024 architectures, in this work, we clearly show how to scale the architecture and what channels to be allocated such that the same transceivers can be used for kilo-core architecture. Second, we discuss the advances and breakthroughs needed by the wireless technology to meet the bandwidth demands of on-chip communication. With detailed analysis, we project the scaling of power-efficiency with different link distance and link efficiency factor for various wireless technologies. Using the wireless power-efficiency, we propose four different architecture configurations where wireless channels can be implemented with different power-efficiency. We simulate the design for wireless technology centered at 100 GHz with CMOS to validate the wireless designs. Third, we simulate the proposed wireless-photonic hybrid architecture for 256 and 1024 cores with synthetic traffic traces and compare against state of the art electronic-wireless, photonics-only and electronic-only architectures. Our simulation results indicate the technology used to design wireless architectures can have significant impact on the overall network power and performance. The major contributions of this work are as follows:

- **OWN Architecture:** We refine and clarify the connectivity, routing and communication using both wireless and photonic technologies in OWN architecture that can be seamlessly scaled from 256 to 1024 cores.
- **Wireless Channel Allocation:** We consider the on-chip distances between wireless transceivers to allocate wireless channels according to energy/bit from CMOS and beyond-CMOS technologies to provide the best energy-efficiency for OWN-256 and OWN-1024 architectures. To validate, we design transceiver circuits for CMOS-only technology and speculate on beyond-CMOS technologies.
- **Performance:** We simulate the OWN architecture for synthetic traffic traces and compare against electronic-wireless, photonics-only and electronic-only architectures. OWN-256 and OWN-1024 improves power savings over a pure-electrical CMESH network in excess of 30% while improving the throughput by 3-5% and latency by 50%.

II. RELATED WORK

Traditional metallic interconnects are designed in 2-D Mesh, Concentrated Mesh (CMesh), or Torus topologies. Since metallic interconnects may not scale for large core counts, architectures employing emerging technologies such

as wireless or photonics are proposed. One such architecture that employs wireless technology is WCube [6]. WCube extends the CMesh architecture by inserting micro-wireless router for a subnet or group of routers. The inter-subnet communication uses wireless technology whereas the intra-subnet communication uses wired technology. Similarly, WiNoC [5] and iWISE [7] uses wireless technology, and uses both wired and wireless technology for inter-subnet communication. More recently, WiSync has been proposed to implement fine-grain synchronization using wireless communication with each core having a transceiver and an antenna to communicate with other cores [9].

Optical NoCs are drawing considerable interest due to their inherent energy and bandwidth advantages. Corona [10] proposes an optical ring-crossbar network using the broadcasting capability of the optical links. Single-writer-multiple-reader (SWMR) technique is used for arbitration, and off-chip laser source and dense wavelength division multiplexing (DWDM) is used for data communication. However, Corona requires a very high number of ring resonators and consumes high power as a portion of the wavelength is peeled off by every router on the path. Firefly [11] reduces the optical crossbar costs by utilizing electrical mesh while 3D-NoC [12] reduces the cost utilizing decomposed crossbars. Similar to 3D-NoC, OWN [8] proposes to use smaller crossbars to reduce the cost but uses wireless technology to connect the crossbars. While OWN showed the architecture design of scaling the nodes using optimistic energy-efficiency values, in this work, we extend the design space by evaluating different wireless technologies for enabling wireless communication. We evaluate different wireless configurations to determine the best scenario for implementing wireless routers for on-chip communication and show the scalability of the OWN architecture to 1024 nodes using the proposed wireless channel allocation.

III. ARCHITECTURE

In this section, we first describe OWN architecture for 256 cores and 1024 cores. We then describe the inter-router communication, wireless channel allocation and antenna placement.

A. OWN for 256 cores

Figure 1(a) shows the proposed OWN architecture for 256 cores. Each core is identified as a quadruple (g, c, t, p) where g identifies the group, c identifies the cluster, t identifies the tile and p identifies the processing element. There are a total of G groups, C clusters per group, T tiles per cluster and P processors per tile with $0 \leq g \leq G - 1$, $0 \leq c \leq C - 1$, $0 \leq t \leq T - 1$, and $0 \leq p \leq P - 1$. For 256 cores design shown in Figure 1(a), $G = 0$, $C = 4$, $T = 16$ and $P = 4$. Each cluster is interconnected by a photonic crossbar that snakes through all the 16 tiles. The photonic waveguide (shown as a ring) is in reality a bus that connects all the tiles in

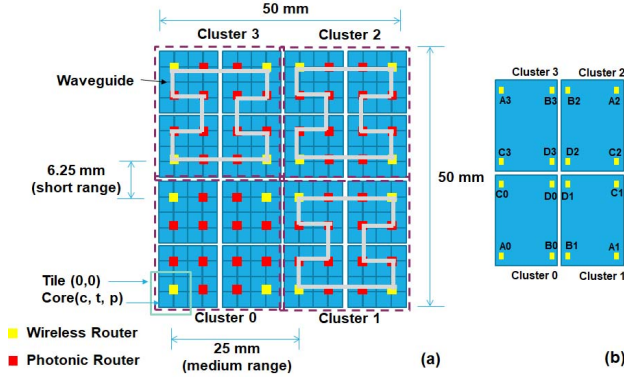


Figure 1. (a) Proposed OWN architecture for 256 cores consisting of 4 clusters with each cluster consisting of 64 cores grouped into 16 tiles with each tile consisting of 4 cores connected to either a wireless or photonic router. (b) Wireless antenna placement within each cluster, for example, A0, B0, C0 and D0 are wireless antenna within cluster 0.

a multiple-writer-single-reader (MWSR) fashion where one tile reads from the waveguide with different tiles writing to it. The bus originates and terminates at the home tile i.e. tile where the multiplexed signal will be dropped. To avoid contention, token arbitration is used such that only one tile can write to it. To enable effective communication, we need 16 waveguides with one home waveguide per tile and 16 tokens that circulate among the tiles. Similar to prior work, we assume off-chip laser source that can generate 64 wavelengths which is pumped into the chip using a separate *power* waveguide and the signal is split across 16 tiles using a star splitter [12]. Each router that connects only to the photonic interconnect is shown in red while those that connect to both photonics and wireless are shown in yellow 1(a).

In Figure 1(b), we show the placement of wireless antennas for inter-cluster communication. We assume that we have 16 wireless channels each with a bandwidth of 32 Gbps. More details on the wireless bandwidth will be discussed in Section 4. We place four wireless transceivers on the four corners of the cluster to facilitate inter-cluster communication such that each of the four routers have wireless antennas. These routers will also connect to the photonic interconnect, therefore, the radix of these routers are 20 (15 to photonic interconnect, 1 wireless and 4 cores). If all the wireless transceivers were located in close proximity (center of the cluster), then all inter-cluster traffic will be directed to the center which could lead to load and thermal imbalance. Therefore, by isolating the four transceivers to the four corners, we balance the load imbalance as well as thermal impact within the cluster. We assume that each individual cluster has a dimension of $25 \times 25 \text{ mm}^2$. This is similar to 61-core, Xeon Phi processor built in 22 nm technology node with a die area of 720 mm^2 which is close to a chip

Table I
VARIOUS WIRELESS CONNECTIONS PROPOSED IN OWN ARCHITECTURE. DIAGONAL OR CORNER-TO-CORNER (C2C), EDGE-TO-EDGE (E2E) AND SHORT-RANGE (SR) ARE DIFFERENT WIRELESS DISTANCES CONSIDERED IN OWN.

Wireless Channels	Naming Convention	Link Distance	Link Factor	Wireless Connections
Diagonal Links	C2C	$\sim 60 \text{ mm}$	1	A3-B1, A0-B2, B1-A3, B2-A0
Edge Links	E2E	$\sim 30 \text{ mm}$	0.5	A2-B3, A1-B0, B3-A2, B0-A1
Short Range	SR	$\sim 10 \text{ mm}$	0.15	C3-C0, C2-C1, C0-C3, C1-C2

dimension of $26 \times 26 \text{ mm}^2$. We assume that we can put 4 such individual chips together and connect via wireless interconnects with 2.5D integration such that each chip is powered separately and is connected to memory via photonic interconnects such as photonic DRAM (PIDRAM) [13]. Prior work such as the design in Galaxy [14] have assumed that multi-chip modules can be designed with photonic interconnects. Here we assume that the individual clusters are photonics interconnects, however they are connected via wireless interconnects.

Table I shows three distances - diagonal links (C2C), edge links (E2E) and short range (SR) - under consideration within the OWN architecture. With four clusters, we need 12 wireless channels that are used to connect all clusters together. For example, cluster 3 communicates with cluster 1 on two wireless channels (A3-B1, B1-A3) and cluster 0 communicates with cluster 2 (A0-B2, B2-A0) using diagonal links which are the longest distance ($\sim 60 \text{ mm}$). Clusters 3 and 2 communicate using two wireless channels (A2-B3, B3-A2) and clusters 0 and 1 communicate using two wireless channels (A1-B0, B0-A1) using the edge links which are medium range distance ($\sim 30 \text{ mm}$). Finally, clusters 0 and 3 communicate using two wireless channels (C0-C3, C3-C0) and clusters 1 and 2 communicate using two wireless channels (C1-C2, C2-C1) using short range links with distances ($\sim 10 \text{ mm}$). The associated distances contribute to the link factor which can be reduced due to shorter distances leading to improved energy-efficiency (in Section 4). There could be different assignments for inter-cluster connections, however they will typically fall within the three distances mentioned. The antennas (D0-D3) will be used for intra-cluster communication as explained next.

B. OWN for 1024 cores

Figure 2 shows the proposed 1024 core architecture with $G = 4$, $C = 4$, $T = 16$ and $P = 4$. This architecture uses the 256-core OWN designed previously as the building block (now called a *group*) and combines four such groups together. Within each cluster inside a group, we still have photonic interconnects as before and the wireless routers (A-D) are located at the same locations. However, we also

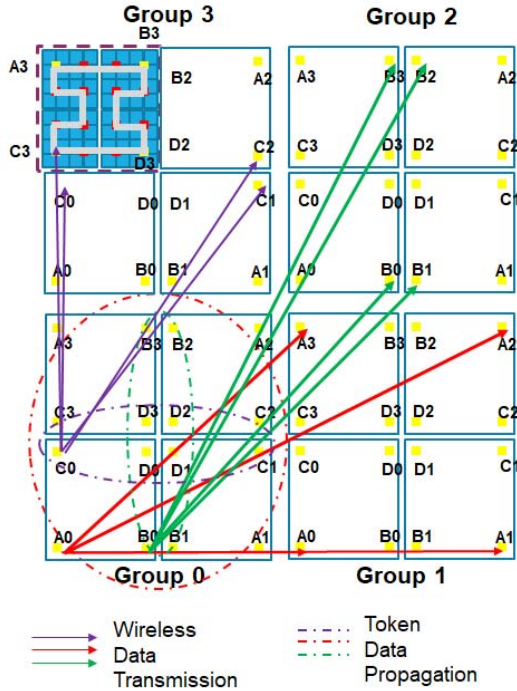


Figure 2. Proposed OWN architecture for 1024 cores consisting of 4 groups with each cluster group consisting of 256 cores. Each cluster has transceivers located as in 256-core OWN architecture, however one additional wireless channel is used for intra-group communication. Communication for group 0 is shown with data paths and token paths.

Table II
WIRELESS CHANNELS ARE SHOWN FOR INTRA-GROUP AND INTER-GROUP COMMUNICATION WITH GROUP 0 AS THE SOURCE GROUP AND GROUPS 1-3 AS THE DESTINATION GROUPS.

Source	Destination	Wireless Channel	Token Sharing	Direction
Group 0	Group 1	A0(G0) – A0(G1) A0(G0) – A1(G1) A0(G0) – A2(G1) A0(G0) – A3(G1)	A0(G0), A1(G0) A2(G0), A3(G0)	Horizontal Inter-Group
Group 0	Group 2	B0(G0) – B0(G2) B0(G0) – B1(G2) B0(G0) – B2(G2) B0(G0) – B3(G2)	B0(G0), B1(G0) B2(G0), B3(G0)	Diagonal Inter-Group
Group 0	Group 3	C0(G0) – C0(G3) C0(G0) – C1(G3) C0(G0) – C2(G3) C0(G0) – C3(G3)	C0(G0), C1(G0) C2(G0), C3(G0)	Vertical Inter-Group
Group 0	Group 0	D0(G0) – D1(G0) D0(G0) – D2(G0) D0(G0) – D3(G0)	D0(G0), D1(G0) D2(G0), D3(G0)	Intra-Group

need to ensure that intra-group communication along with inter-group communication across clusters, therefore, the previously proposed MWSR approach may not be sufficient. Instead we adopt the single-writer-multiple-reader (SWMR) approach where we multicast the request to several wireless transceivers in different clusters. Figure 2 shows the wireless communication proposed for 1024 nodes. In this design, the same wireless channel is used for inter-group communication with different clusters receiving the same signal; the in-

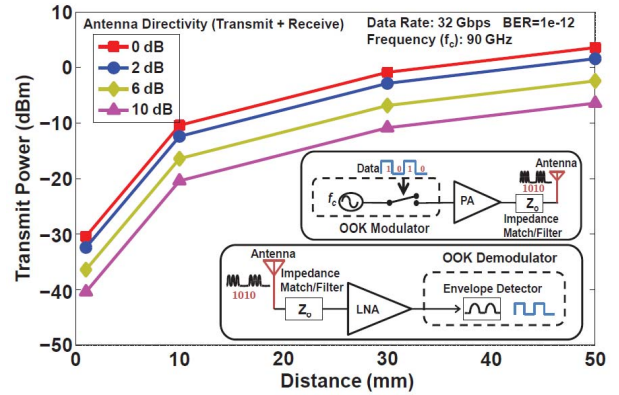


Figure 3. The link budget estimation at the data rate of 32 Gbps and the center frequency of 90 GHz for different antenna directivities. Right Inset: The OOK Transmitter (Top) and Receiver (Bottom).

tended destination cluster will simply forward the signal and the rest will discard it. For example, A0 in group 0 transmits the same signal to A0, A1, A2, and A3 in group 1 at the same time. This ensures that all four wireless transceivers receive the signal, and then the intended receiver will forward the packet on the photonic interconnect. The remaining receivers will discard the data since it is not intended for the receiving group. Table II shows the wireless channel assigned between group 0 and group 1-3. Similar allocation is made for other inter-group communication. Now, since only one cluster with group 0 can transmit at any time, we ensure that token is propagated across different transmitters within the group to enable the communication (this is shown by the dotted line). In traditional SWMR consumes more power since the signal needs to separately reach all the receivers; however using wireless simplifies the design since the signal is multicast and there is no additional transmitter power required. However, receiver power is consumed since the data has to be analyzed before discarding it. If the actual clusters are set in 2D design, then the prior distances (from Table I) will not be applicable; however each group can be integrated in a 3D layout enabling similar distances from before.

IV. WIRELESS TECHNOLOGY

In this section, we explore the feasibility of integrating electrical, wireless and optical interconnects in OWN architecture and the best strategies to reach the ambitious targets. While there has been several studies to integrate photonic interconnects [3], [4], in this work, we focus on the challenges of integrating wireless transceivers. First, we introduce circuit building blocks for wireless transceivers in 65-nm CMOS relevant for implementation of OWN wireless links at 100 GHz. Then, we discuss alternative pathways for wireless transceiver beyond CMOS in BiCMOS and SiGe technologies. CMOS and BiCMOS technology represents

the current state-of-the-art in the wireless design, with pure SiGe HBT design being a more speculative solution that is likely to shape Si integration above ~ 500 GHz.

A. Wireless Transceivers in CMOS

In order to design a very efficient wireless communication channel, we first study the link budget and introduce the wireless transceiver design to be employed. The modulation scheme proposed is the non-coherent On-Off keying (OOK) because of its design simplicity as well as power and area efficiency [15]. The OOK modulator and demodulator are depicted in the inset of Figure 3. It requires an oscillator and modulated power amplifier (PA) driving the antenna on the transmitter side and a low-noise amplifier (LNA) followed by an envelope detector on the receiver end. For efficiency, it is important to tune the oscillator signal and the PA gain for short distances involved, limited to around 50 mm. The RF output power of the transmitter for various distances and antenna gains can be obtained from Figure 3. For a data rate of 32 Gbps at the center frequency of 90 GHz and isotropic antenna (0 dB directivity), the maximum power required for an OOK transmitter is ≥ 4 dBm for a maximum distance of 50 mm in OWN-256 design.

The carrier signal may be generated via a power-efficient Colpitt oscillator at 90 GHz, as shown in the right lower inset of Figure 4(a). To achieve higher operating frequency, and reduce non-linear effects, no external capacitors have been used in the design. The gate-source and gate-drain capacitances of M_1 , which is inherent to the device, is substituted for the external capacitors. These resonate with the inductor, L , to produce the oscillation. The PSD at 1 V supply has been plotted and can be observed in the left upper inset of the figure Figure 4(a). The phase noise at 1 MHz offset is observed to be around -86 dBc/Hz.

The PA in our design is a one-stage class-AB amplifier (inset of Figure 4(b)) with a DC power dissipation of 14 mW at 1 V supply. It can be biased to produce a sufficient RF power (P_{RF}) of 7 dBm (≥ 4 mW required) with sufficiently low-distortion as verified from the 1-dB compression point of ~ 5 dBm. The PA achieves a peak gain of 3.5 dB centered around 90 GHz with a bandwidth of around 20 GHz considering a gain of 2 dB, as seen in Figure 4(b). The PA reflection loss $\geq 10\%$ indicates that there is sufficient output matching for a bandwidth of 16 Gbps transmission. Clearly a wider bandwidth design is necessary for 32 Gbps operation, which can be achieved by higher-order matching circuits and higher transconductance or using SiGe Heterojunction Bipolar Transistor (HBT). In the receiver end, a wideband common-source degeneration cascade-cascode LNA is designed, which has a gain of 10 dB, as can be seen in Figure 4(c). The LNA gain is sufficient for 50 mm operation and can be further lowered depending on the performance of the envelope detector to be implemented by a diode connected transistor.

The above CMOA designs illustrate that basic building blocks of the OOK transmitter operating at 100 GHz bands is already achievable. To achieve wireless communication at 500 GHz or beyond, the design of the transceiver needs to accommodate different device technologies with higher transition frequency (f_T) such as SiGe HBT in BiCMOS platforms. Access to both CMOS and HBT transistors on the same BiCMOS framework is especially welcome as LNA & PA will require the use of HBT to boost the gain while all other elements can be built using low-power MOSFET's. Depending on the sub-32nm RF CMOS technologies being developed using 22/16 nm FinFET, high-efficiency oscillator and PAs with back-gate tunability are also expected, and are very suitable for compact OOK designs.

B. Wireless Transceivers Beyond CMOS

Due to limited gain and increasing parasitics, a CMOS-only RF solution will be limiting PA and LNA designs in sub-32nm technology [16]. Thus, SiGe BiCMOS technology is the only feasible semiconductor process that has the unique potential to address all device, circuit and integration requirements for the proposed OWN-256 architecture. Combining the best of advances in ultra-low power CMOS devices [17], [18], THz SiGe HBT transistor technology and high-performance passives, the BiCMOS technology platforms rival III-V semiconductors in performance [19]. Indeed, such SiGe HBT devices are routinely used today to drive state-of-the-art fibre-optic networks where BiCMOS integration can reduce cost and size [20]. SiGe HBT can perform similar tasks, including signal drive, modulation and low-noise transimpedance amplifiers in the optical links layer of OWN architecture. However, they can also provide a unique opportunity to efficiently implement OWN wireless networks, since both CMOS and HBT transistors can be selectively utilized in the same process, leaving it to designers to decide if or when to recourse to higher-gain power-hungry SiGe HBT devices for wireless routers. Thus, utilization of SiGe BiCMOS process for OWN essentially becomes a strategic optimization between the use of low-power but performance- and band-limited CMOS transceivers versus more capable yet less-efficient SiGe HBT devices. The most realistic case is to adapt a *hybrid* scheme and utilize CMOS in all active circuits where possible, limiting the use of HBTs only to few critical elements critical for operation, notably in PA and LNAs. Such optimization is further complicated by the fact that mm-wave capable BiCMOS technologies and back-end RF components typically lag several generations behind the digital CMOS processes. Thus, some of the critical power and bandwidth performance figures for both CMOS and HBT devices are not yet available, making the precise OWN design pathway unclear. As a result, we develop two possible scenarios for the implementation of OWN-256 design, as presented in Table III, which differ in terms of available power efficiency and bandwidth. Although




Figure 4. (a) The power spectrum density (PSD) of the oscillation at the frequency of 90 GHz. Left upper inset: Phase noise of the oscillator. Right upper and lower inset: 90 GHz oscillation in time domain and Colpitt Oscillator circuit respectively. (b) The linearity of the PA in terms of 1-dB compression point. This verifies the PA can achieve the required power level estimation of the link budget. (c) The wideband LNA circuit and its gain around 90 GHz.

Table III
COMPARISON OF POWER EFFICIENCY OF WIRELESS NETWORK-ON-CHIP (WiNoC) IMPLEMENTATION USING CMOS, BiCMOS AND SiGe TECHNOLOGIES.

Scenario #1: BW=32 GHz (IDEAL): OOK ; Efficiency ramps: +0.05pJ/bit (CMOS) +0.1pJ/bit (SiGe) +0.07pJ/bit (Hybrid)

Technology	CMOS-only: 0.1pJ/bit/s				BiCMOS Hybrid: 0.3pJ/bit/s				SiGe HBT Only: 0.5pJ/bit/s							
Link #	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
BandF _c (GHz)	20	60	100	140	180	220	260	300	340	380	420	460	500	540	580	620
Bandwidth (GHz)	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32
Function	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx
Link pairs	A1:B4	A1:B4	B2:A3	B2:A3	B1:A2	B1:A2	B3:A4	B3:A4	C1:C3	C1:C3	C2:C4	C2:C4	n/a	n/a	n/a	n/a
Link Distance (LD)	C2C	C2C	C2C	C2C	E2E	E2E	E2E	E2E	SR	SR	SR	SR	SR	SR	SR	SR
LD Factor	1	1	1	1	0.5	0.5	0.5	0.5	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.15
Effic. (pJ/bit)	0.10	0.15	0.20	0.25	0.58	0.65	0.72	0.79	1.30	1.40	1.50	1.60	1.70	1.80	1.90	2.00
Total Power (mW)	3.20	4.80	6.40	8.00	9.28	10.40	11.52	12.64	6.24	6.72	7.20	7.68	8.16	8.64	9.12	9.60

Scenario #2: BW=16 GHz (CONSERVATIVE) OOK ; Efficiency ramps: +0.05pJ/bit (CMOS) +0.07pJ/bit (SiGe) +0.06pJ/bit (Hybrid)

Technology	CMOS-only: 1pJ/bit/s						BiCMOS Hybrid: 1.5pJ/bit/s						SiGe HBT Only: 2pJ/bit/s			
Link #	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Band f _c (GHz)	40	60	80	100	120	140	160	180	200	220	240	260	280	300	320	340
Bandwidth (GHz)	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16
Function	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx
Link pairs	A1:B4	A1:B4	B2:A3	B2:A3	B1:A2	B1:A2	B3:A4	B3:A4	C1:C3	C1:C3	C2:C4	C2:C4	n/a	n/a	n/a	n/a
Link Distance (LD)	C2C	C2C	C2C	C2C	E2E	E2E	E2E	E2E	SR	SR	SR	SR	SR	SR	SR	SR
LD Factor	1	1	1	1	0.5	0.5	0.5	0.5	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.15
Effic. (pJ/bit/efficiency)	1.00	1.05	1.10	1.15	1.20	1.25	1.86	1.92	1.98	2.04	2.10	2.16	2.84	2.91	2.98	3.05
Total Power (mW)	16.00	16.80	17.60	18.40	9.60	10.00	14.88	15.36	4.75	4.90	5.04	5.18	6.82	6.98	7.15	7.32

speculative for $f > 500$ GHz, these scenarios will allow us to explore the limits and most efficient use of BiCMOS technology for OWN architecture at different spectral and power limitations.

Technology Choices: The two (ideal and conservative) scenarios summarized in Table III is built on the assumption that both BiCMOS device technologies and the following RF back-end auxiliaries (LC passives, transmission lines, isolation structures, and vias) will continue advancing in terms of raw performance (higher g_m and f_t/f_{max}), leakage reduction, integration and size reduction. This is conceivable because of the aforementioned lag between digital and RF CMOS technology nodes, continuing advances in HBT optimization and recent advances in materials such as graphene, ferroelectric polymer composites and magnetic nanostructures in particular [16], [21]. Hence, base efficiencies of 0.1pJ/bit and 0.5pJ/bit is assumed for transceivers built using CMOS and HBT devices, respectively, in the BiCMOS technology. Additionally, we also consider that

these performance limits will deteriorate as the frequency of operation (link frequency) gets higher, since silicon is not an optimal substrate for THz integration and parasitics/losses increase at higher frequencies. In the table, these limits are expressed as *efficiency ramps* of +0.05pJ/bit (CMOS) +0.07pJ/bit (BiCMOS) and +0.1pJ/bit (HBT) devices in the ideal case and +0.05pJ/bit (CMOS) +0.06pJ/bit (BiCMOS) and +0.07pJ/bit (HBT) for the conservative case. Since BW is twice smaller in the conservative case (16 vs. 32 GHz) and link frequencies are lower, this leads to a greater increase of losses in the ideal scenario. From Table III, links 1-12 are used for inter-cluster communication whereas links 13-16 are reserved for reconfiguration channels that could adaptively be utilized to improve performance.

Bandwidth Allocations: The second important assumption in Table III is the BW of the resulting transceivers and their allocation to 16 bands for the two scenarios. For the ideal case, we assume bandwidth of 32 GHz for all bands, which will be more challenging for lower frequency

links utilizing only CMOS. In the conservative outlook, the assumption is to allocate only 16 GHz BW per channel, which would save some power by minimizing SiGe HBT usage. It is worth noting that in both scenarios, link frequencies are chosen such that there is at least 4 GHz or 8 GHz isolation between the adjacent bands in the conservative or ideal cases, respectively. This is to ensure that there is no significant intermodulation between them, thereby saving significant power or area that would have been committed to inefficient passive/active filters at such elevated frequencies. Moreover, we also made specific assumptions in the frequency-technology pairings shown in the table. For instance, we consider ~ 300 GHz as a limit beyond which to use SiGe HBT-only circuitry in the wireless routers except its digital infrastructure, which can always be re-visited as will be discussed later in the results section.

Distance Scaling: Another important assumption critical for OWN implementation is the scaling of transceiver radiated power according to the location of routers in the OWN-256 floor-plan. Since the chip is large (~ 50 mm) and routers are positioned at different locations, some fairly close to one another, such power optimization will be highly desirable to ensure that OWN-256 design not waste excess power over shorter distances. This is noted in the Table III as *link distance* (LD) factor, which changes from 1 for C2C, (corner-to-corner) links, 0.5 for E2E, (edge-to-edge) links to 0.15 for SR (short-range, 10 mm) links. LD factor is the result of power changes as a function of distance as indicated in the link budget calculations of Figure 3.

V. PERFORMANCE EVALUATION

To evaluate the performance of the proposed NoC architecture, we compare the 256-core and 1024-core OWN with CMESH, wireless-CMESH [6], optical crossbar (OptXB) [10] and photonic-Clos (p-Clos) [22] architectures. We used Dsent v. 0.91 [23] to calculate the area and power of the wired links and routers for a bulk 45 nm LVT technology. To simulate network performance for different types of synthetic traffic patterns such as uniform (UN), bit-reversal (BR), matrix transpose (MT), perfect shuffle (PS), and neighbor (NBR), we have used a cycle accurate simulator [24] keeping the router and core frequency same for all the networks. Since we are simulating large network sizes (beyond 64), we have simulated the proposed designs with synthetic traffic only. In the future, we will evaluate with real workloads.

A. Simulation Methodology

In order for a fair comparison between different topologies, we have kept the bisection bandwidth same for all the architectures by adding appropriate delay into the network. We assume 4 virtual channels per input port with a regular 5-stage pipelined router (routing computation (RC), virtual channel allocation (VCA), switch allocation (SA), switch

traversal (ST) and link traversal (LT)) for each of the architecture. For OWN-256 architecture, the maximum radix is 20 (1 wireless transceiver, 15 optical transceiver and 4 cores) for wireless routers and 19 for photonic routers. Under worst case scenario, a packet will take three hops to reach the destination (one photonic to wireless router within the cluster, inter-cluster wireless hop and finally photonic hop to reach the destination tile). In order to avoid deadlocks, we allocate 2 VCs for data packet communication over the photonic link and 2 VCs for wireless link. This 50% allocation ensures that both intra- and inter-cluster has the same priority within the router. CMESH is designed with 4 cores per router with a maximum radix of 8 and XY dimension-order routing (DOR) to prevent deadlocks. The maximum diameter is $2(\sqrt{n} - 1)$ where n is the number of routers. For the photonic crossbar (OptXB), we assume the 4 cores are concentrated together and the maximum diameter is one. For the p-Clos architecture, we assumed that the maximum number of hops is two i.e. all concentrated nodes are connected to one level of switches before they are connected back to the router. We implement MWSR with token arbitration with a router radix of 67 (63 for the crossbar and 4 cores). Wireless CMESH also has a core concentration of 4 and a total of 64 routers. Each wireless cluster has 4 routers connected by an electrical crossbar, and one router is a wireless router and 16 of the wireless clusters make up the 256-core chip. Wireless routing is implemented as XY DOR to prevent deadlocks and the maximum hop count is \sqrt{n} where n is the number of routers. The radix of the wireless-CMESH is 11 (3 electrical, 4 wireless x-y and 4 cores).

For 1024-core architecture, the maximum number of hops is still three as before since we implement SWMR along with MWSR (one photonic hop within the cluster, one inter-group wireless multicast and one intra-cluster photonic hop). The maximum radix is 22 (15 photonic, 3 wireless and 4 cores). To avoid deadlocks, the VC allocation is restricted as follows: VC0 for intra-group communication, VC1 for inter-group vertical, VC2 for inter-group horizontal and VC3 for inter-group diagonal. The OptXB, p-Clos, CMESH and wireless-CMESH are scaled to 1024 cores by increasing the radix and the hop count.

B. Power and Performance for 256 cores

Table IV shows the different configurations that we tested in our simulation. Configuration 1 assumes SiGe for long range, CMOS for medium range and short range, Configuration 2 assumes CMOS for long range, BiCMOS for medium range and SiGe for short range, Configuration 3 assumes SiGe for long range, BiCMOS for medium range and CMOS for short range and finally Configuration 4 assumes CMOS for long and medium range and BiCMOS for short range. These are different cases with scenarios picked from Table III. Figure 5 shows the average wireless link power

Configuration	Wireless Technology
Configuration 1	(Long Range - SiGe)(Mid Range - CMOS)(Short Range - CMOS)
Configuration 2	(Long Range - CMOS)(Mid Range - BiCMOS)(Short Range - SiGe)
Configuration 3	(Long Range - SiGe)(Mid Range - BiCMOS)(Short Range - CMOS)
Configuration 4	(Long Range - CMOS)(Mid Range - CMOS)(Short Range - BiCMOS)

Table IV
DIFFERENT WIRELESS NETWORK-ON-CHIP (WiNoC) IMPLEMENTATION USING CMOS, BiCMOS AND SiGe TECHNOLOGIES.




Figure 5. Average wireless link power consumed for different scenarios for random traffic.

considering the two scenarios for different configurations under evaluation for random traffic pattern. We measured the total number of packets sent and received to evaluate the percentage of traffic that uses the wireless channels. From 5, it is clear that configurations 1 and 3 that use SiGe for long range consume significantly more power under both scenarios (32 GHz and 16 GHz wireless bandwidth). Configuration 2 and 4 reduce the power consumption significantly as they rely on CMOS technology. For example, under scenario 1, configuration 1 power is reduced by 60% and 80% by configuration 2 and configuration 4. Similarly, under scenario 2, configuration 1 power is reduced by 47% and 57% by configuration 2 and configuration 4 respectively. Clearly, 32 GHz channel bandwidth relying on CMOS technology with BiCMOS would appear to be a promising approach. However, III shows only four channels with CMOS and we would need at least 8 channels to be designed with CMOS technology. One approach is to implement space-division multiplexing such that the same channel frequency is used on different non-intersecting areas. From Figure 1(b), we could assign B3-A2 and B0-A1 the same channel frequency since the signals do not intersect. Similarly, we can allocate C0-C3 and C1-C2 the same wireless channel, and thereby implement CMOS at multiple locations. While this is a promising approach, care must be taken to ensure that the transmission power is kept at a minimum to limit interference.

It is important to emphasize that the present simulation study is a first attempt to indicate the optimization required to utilize of SiGe BiCMOS technology for kilocore OWN




Figure 6. Power consumed for different configurations including wireless-CMESH, all-photonic crossbar, photonic-Clos and CMESH architectures.

architectures. Clearly, depending on the eventual process parameters, and the quality of RF back-end components, it is possible to come up with additional scenarios to optimize the use SiGe BiCMOS for wireless NoCs. For instance, avoiding SiGe-HBT only transceiver designs all together could save significant power, if performance of SiGe BiCMOS is adequate up to 500GHz regime. Similarly, one can also consider an additional scenario between the two-extreme (best or worst) cases, which may correspond to actual process conditions in reality. Such additional studies will be the subject of our subsequent investigations as the SiGe BiCMOS technology develops further.

Figure 6 shows the power consumed for different configurations as well for different topologies under uniform random traffic. We have considered the power consumed by the photonic link, wireless link, electrical link and the router microarchitecture. The OptXB consumes the least power since the energy-efficiency of photonic links is extremely high (1-2 pJ/bit) and therefore, the photonic power is minimal. The radix of the router microarchitecture contributes to the power consumption, but it is not significant. The OWN in configuration 4 consumes the next least power (almost 2X of OptXB). It must be noted that designing optical snake-like waveguide interconnecting 64 routers with 64 wavelengths will require more than a million ring resonators alone [10]. Therefore, while OptXB consumes the least power, it is quite challenging to integrate all photonic components while mitigating thermal and process variations for more than a million components. The p-Clos architecture consumes slightly more than a crossbar since it has more hops and router power adds up. The wireless-CMESH consumes 7% more power than OWN since there are more wireless hops to navigate when compared OWN. However, the router radix is almost half of OWN and therefore, the router does not consume as much power as OWN. OWN Configurations 1-3 consume power proportional to the wireless link power as shown in Figure IV and perform accordingly. CMESH consumes the most power among all the topologies. When compared to OWN (Configuration 4), CMESH requires 30% in excess power and the majority of the power is dissipated in the routers.

Figure 7(a) shows the throughput for different synthetic traffic traces for all topologies under evaluation. As OWN-256 Configuration 4 showed the best power results,

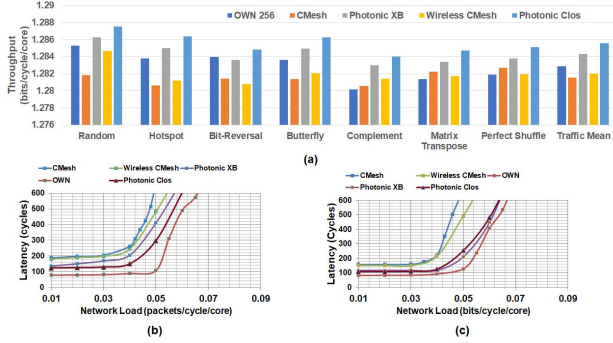


Figure 7. (a) Throughput for several synthetic traffic patterns and average packet latency at saturation for (b) random and (c) bit reversal traffic patterns for CMESH, OWN-256 (with configuration 4), photonic crossbar, photonic Clos and wireless CMESH architectures.

we have assume configuration 4 for 256 and 1024 core throughput, latency and power results. OWN-256 shows 1-2% higher throughput when compared to CMESH and wireless-CMESH architecture. The photonic architectures are marginally better than the OWN design. Since the bisection bandwidths are similar, and topologies have similar throughput result. Figure 7(b,c) show the network latency for different architectures for random and bit reversal traffic patterns. From the result, we observe that OWN saturates at the highest network load. The next best performing network is the p-Clos which saturates 10% earlier than OWN. CMESH, wireless-CMESH and photonic crossbar saturate 20% earlier than OWN. OptXB shows a slight decrease in throughput since token transfer consumes a few extra cycles. OWN reduces the hop count, but has higher link count which allows OWN to handle more packets than other networks.

C. Power and Throughput for 1024 cores

Figure 8(a) and (b) show the throughput and power consumed for 1024-core architecture. We compare the result on a select few synthetic traces for different architectures. The throughput variation is not significant across different architectures. From the power result, we observe that the high radix of OptXB adds considerable power to the total power consumed. Similarly, p-Clos also adds power due to the increase in the number of routers. In this case, the OWN architecture consumes 30% more power compared to OptXB; however the design complexity and scalability of OptXB is challenging. It must be noted that in the 1024-core case, we need 16 wireless channels and not 12 as in 256-core case. Therefore, we require all channels described in Table III. In 1024 case, the major component of power consumed in wireless-CMESH is the wireless link since extra hops needs to be navigated as we implement XY DOR routing algorithm. However, since the router radix is constant, the router power is lesser in this case as well. For the 1024-OWN, the router power is significant since the radix is

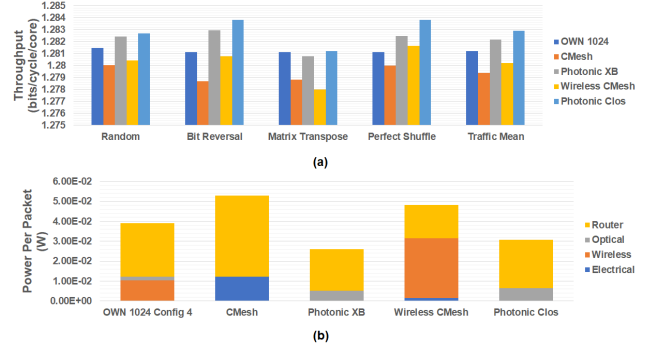


Figure 8. (a) Throughput for different synthetic traffic for CMESH, OWN-1024 (with configuration 4), photonic crossbar, photonic Clos and wireless CMESH and (b) average power consumed per packet for different architectures.

twice of wireless-CMESH architecture and consumes 3% lesser power than wireless-CMESH architecture. Therefore, reducing the radix can enable building more power-efficient architectures, however the latency may increase due to multiple hops.

VI. CONCLUSIONS

In this paper, we analyzed the impact of wireless technology on power-efficiency for wireless-photonic hybrid NoC architectures. We discussed the scaling trends of using CMOS, BiCMOS, and SiGe technologies for implementing 256 and 1024 OWN architectures. On the architecture side, we analyze the wireless channel allocation, distances between transceivers and routing techniques to enable inter-group and inter-cluster communication within the limits of wireless bandwidth. Relying on CMOS and BiCMOS technologies and utilizing SDM techniques can significantly improve the power-efficiency of wireless technologies for future multicores. OWN-256 and OWN-1024 improves power savings over a pure-electrical CMESH network in excess of 30% while improving the throughput by 3-5% and latency by 20%.

VII. ACKNOWLEDGEMENT

This research was partially supported by NSF grants CCF-1054339 (CAREER), CCF-1420718, CCF-1318981, CCF-1513606, CCF-1703013, CCF-1547034, CCF-1547035, CCF-1540736, CCF-1702980 and and by the David and Marilyn Karlgaard Endowment.

REFERENCES

- [1] B. Bohnenstiehl, A. Stillmaker, J. Pimentel, T. Andreas, B. Liu, A. Tran, E. Adeagbo, and B. Baas, "A 5.8 pj/op 115 billion ops/sec, to 1.78 trillion ops/sec 32nm 1000-processor array," in *IEEE Symposium on VLSI Circuits*, 2016.
- [2] A. Mammela and A. Anttonen, "Why will computing power need particular attention in future wireless devices?" *IEEE Circuits and Systems Magazine*, vol. 17, no. 1, pp. 12–26, Firstquarter 2017.

- [3] J. S. Orcutt, B. Moss, C. Sun, J. Leu, M. Georgas, J. Shainline, E. Zraggen, H. Li, J. Sun, M. Weaver, S. Urošević, M. Popović, R. J. Ram, and V. Stojanović, "Open foundry platform for high-performance electronic-photonics integration," *Opt. Express*, vol. 20, no. 11, pp. 12 222–12 232, May 2012.
- [4] M. Hochberg and T. Baehr-Jones, "Towards fabless silicon photonics," *Nature photonics*, vol. 4, no. 8, pp. 492–494, 2010.
- [5] A. Ganguly, K. Chang, S. Deb, P. P. Pande, B. Belzer, and C. Teuscher, "Scalable hybrid wireless network-on-chip architectures for multicore systems," *Computers, IEEE Transactions on*, vol. 60, no. 10, pp. 1485–1502, 2011.
- [6] S.-B. Lee, S.-W. Tam, I. Pefkianakis, S. Lu, M. F. Chang, C. Guo, G. Reinman, C. Peng, M. Naik, L. Zhang *et al.*, "A scalable micro wireless interconnect structure for cmpps," in *Proceedings of the 15th annual international conference on Mobile computing and networking*. ACM, 2009, pp. 217–228.
- [7] D. DiTomaso, A. Kodi, D. Matolak, S. Kaya, S. Laha, and W. Rayess, "A-winoc: Adaptive wireless network-on-chip architecture for chip multiprocessors," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 26, no. 12, pp. 3289–3302, Dec 2015.
- [8] M. A. I. Sikder, A. K. Kodi, M. Kennedy, S. Kaya, and A. Louri, "Own: Optical and wireless network-on-chip for kilo-core architectures," in *High-Performance Interconnects (HOTI), 2015 IEEE 23rd Annual Symposium on*. IEEE, 2015, pp. 44–51.
- [9] S. Abadal, A. Cabellos-Aparicio, E. Alarcn, and J. Torrellas, *WiSync: An architecture for fast synchronization through on-chip wireless communication*. Association for Computing Machinery, 3 2016, vol. 02-06-April-2016, pp. 3–17.
- [10] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System implications of emerging nanophotonic technology," in *ACM SIGARCH Computer Architecture News*, vol. 36, no. 3. IEEE Computer Society, 2008, pp. 153–164.
- [11] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: illuminating future network-on-chip with nanophotonics," in *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3. ACM, 2009, pp. 429–440.
- [12] R. Morris, A. Kodi, and A. Louri, "Dynamic reconfiguration of 3d photonic networks-on-chip for maximizing performance and improving fault tolerance," in *Microarchitecture (MICRO), 2012 45th Annual IEEE/ACM International Symposium on*, Dec 2012, pp. 282–293.
- [13] S. Beamer, C. Sun, Y.-J. Kwon, A. Joshi, C. Batten, V. Stojanović, and K. Asanović, "Re-architecting dram memory systems with monolithically integrated silicon photonics," *SIGARCH Comput. Archit. News*, vol. 38, no. 3, pp. 129–140, Jun. 2010.
- [14] Y. Demir, Y. Pan, S. Song, N. Hardavellas, J. Kim, and G. Memik, "Galaxy: A high-performance energy-efficient multi-chip architecture using photonic interconnects," in *Proceedings of the 28th ACM International Conference on Supercomputing*, ser. ICS '14. New York, NY, USA: ACM, 2014, pp. 303–312.
- [15] S. Laha, S. Kaya, D. W. Matolak, W. Rayess, D. DiTomaso, and A. Kodi, "A new frontier in ultralow power wireless links: Network-on-chip and chip-to-chip interconnects," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 2, pp. 186–198, Feb 2015.
- [16] S. P. Voinigescu, S. Shopov, J. Hoffman, and K. Vasilakopoulos, "Analog and mixed-signal millimeter-wave sige bimos circuits: State of the art and future scaling," in *2016 IEEE Compound Semiconductor Integrated Circuit Symposium (CSICS)*, Oct 2016, pp. 1–4.
- [17] A. Balteanu, S. Shopov, and S. P. Voinigescu, "A 2 44gb/s 110-GHz Wireless Transmitter with Direct Amplitude and Phase modulation in 45-nm soi cmos," in *IEEE Compound Semiconductor Integrated Circuit Symposium (CSICS)*, Oct 2013, pp. 1 – 4.
- [18] K. Nakajima, A. Maruyama, T. Murakami, M. Kohtani, T. Sugiura, E. Otobe, J. Lee, S. Cho, K. Kwak, J. Lee, M. Fujishima, and T. Yoshimasu, "A low-power 71ghz-band cmos transceiver module with on-board antenna for multi-gbps wireless interconnect," in *Microwave Conference Proceedings (APMC), 2013 Asia-Pacific*, Nov 2013, pp. 357–359.
- [19] E. Seok, D. Shim, C. Mo, R. Han, S. Sankaran, W. K. C. Cao, and K. K. O, "Progress and challenges towards Terahertz CMOS integrated circuits," *IEEE JSSC*, vol. 45, no. 8, pp. 1554–1564, 2010.
- [20] S. P. Voinigescu, S. Shopov, J. Hoffman, and K. Vasilakopoulos, "Analog and mixed-signal millimeter-wave sige bimos circuits: State of the art and future scaling," in *2016 IEEE Compound Semiconductor Integrated Circuit Symposium (CSICS)*, Oct 2016, pp. 1–4.
- [21] A. Pan and C. O. Chui, "Rf performance limits of ballistic si field-effect transistors," in *Silicon Monolithic Integrated Circuits in Rf Systems (SiRF), 2014 IEEE 14th Topical Meeting on*. IEEE, 2014, pp. 68–70.
- [22] A. Joshi, C. Batten, Y. J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic, "Silicon-photonics cmos networks for global on-chip communication," in *2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*, May 2009, pp. 124–133.
- [23] C. Sun, C.-H. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic, "Dsnt-a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling," in *Networks on Chip (NoCS), 2012 Sixth IEEE/ACM International Symposium on*. IEEE, 2012, pp. 201–210.
- [24] A. Kodi and A. Louri, "A system simulation methodology of optical interconnects for high-performance computing systems," *J. Opt. Netw.*, vol. 6, no. 12, pp. 1282–1300, 2007.