

Opportunistic Beamforming in Wireless Network-on-Chip

Sergi Abadal*, Adrián Marruedo*, Antonio Franques[‡], Hamidreza Taghvaei*,
Albert Cabellos-Aparicio*, Jin Zhou[†], Josep Torrellas[‡], Eduard Alarcón*

*NaNoNetworking Center in Catalunya (N3Cat), Universitat Politècnica de Catalunya (UPC), Barcelona, Spain

[†]Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign (UIUC), Illinois, USA

[‡]Department of Computer Science, University of Illinois at Urbana-Champaign (UIUC), Illinois, USA

Email: abadal@ac.upc.edu

Abstract—Wireless Network-on-Chip (WNoC) has emerged as a promising alternative to conventional interconnect fabrics at the chip scale. Since WNoCs may imply the close integration of antennas, one of the salient challenges in this scenario is the management of coupling and interferences. This paper, instead of combating coupling, aims to take advantage of close integration to create arrays within a WNoC. The proposed solution is opportunistic as it attempts to exploit the existing infrastructure to build a simple reconfigurable beamforming scheme. Full-wave simulations show that, despite the effects of lossy silicon and nearby antennas, within-package arrays achieve moderate gains and beamwidths below 90° , a figure which is already relevant in the multiprocessor context.

I. INTRODUCTION

Network-on-Chip (NoC) has become the *de facto* standard for the interconnection of cores in multicore processors. However, as we enter the manycore era, the communication requirements increase up to a point where conventional NoCs alone may not suffice [1]. Their limited scalability is in fact turning communication into the performance bottleneck of manycore systems, thus calling for new solutions at the interconnect level [2].

Advances in integrated antennas [3], [4] and transceivers [5], [6] have led to the proposal of Wireless Network-on-Chip (WNoC) as a complement of or alternative to existing NoCs [7]. As shown in Figure 1, WNoCs basically consist of the co-integration of RF front-ends with cores or clusters of cores. Information can be thus modulated and radiated, and radiated signals then propagate through the computing package until reaching the intended destinations. The main advantage of this approach is that distant cores can communicate with low latency thanks to the speed-of-light propagation. In fact, communication is naturally broadcast as long as antennas are roughly omnidirectional. Further, the lack of additional wires between cores provides system-level flexibility not achievable with other technologies.

From the network architecture perspective, one can distinguish between a large set of WNoC proposals that deploy multiple point-to-point wireless links over a wired NoC [8]–[10] and, to a lesser extent, broadcast-based WNoCs [11], [12]. On-chip antennas used in these proposals are generally variants of printed dipoles [13], [14] or vertical monopoles using through-silicon vias (TSVs) [15]–[17], with rather omnidirectional




Fig. 1. Cross-section and planar view of a multicore processor with a wireless on-chip network. Thanks to the proposed architecture, antennas can operate in isolation (blue, omnidirectional) or form small arrays (green, directional).

radiation patterns. As a result, MAC protocols or multiplexing methods are required to avoid collisions and interference in the WNoC [18]–[20]. However, this approach has important limitations because the number of non-overlapping frequency, code, or time slotted channels achievable in this resource-constrained scenario is relatively small.

An alternative or complement to the multiplexing schemes mentioned above would be *spatial multiplexing* as proposed in some works [21]–[25]. By using directional antennas, several wireless point-to-point links can coexist in the same frequency-time window and increase the overall available throughput. The main downturn of this approach, however, is that the antennas need to be carefully aligned and that the established links cannot be reconfigured, thereby losing the system-level flexibility and inherently broadcast appeal of the WNoC paradigm. This issue could be partially overcome by means of dynamic beamforming, but this would require the use of antenna arrays in each wireless interface as proposed in [26], which is clearly unaffordable given the evident area limitations of the manycore scenario.

This paper proposes an opportunistic solution to this problem, seeking to implement spatial multiplexing in a flexible and affordable way for wireless on-chip networks. The main idea is to leverage the known channel characteristics [28],

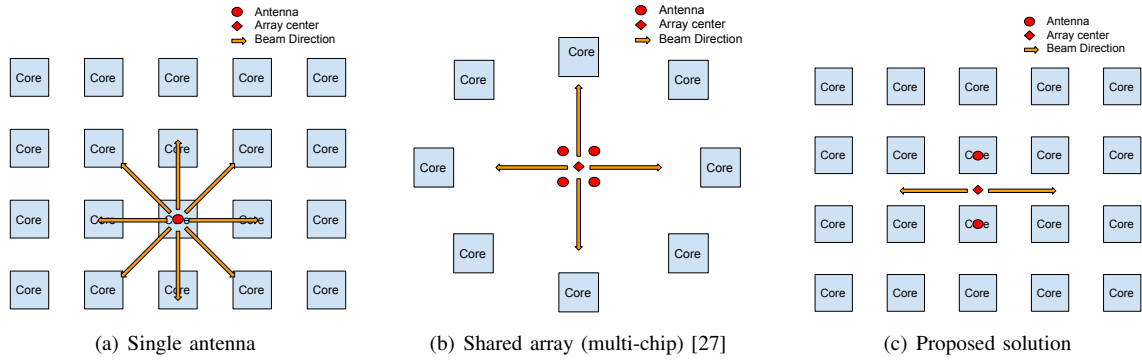


Fig. 2. Summary of antenna configurations in a wireless on-chip network.

the already existing high density of on-chip antennas, and the already existing tight synchronization among cores to create small antenna arrays as shown in Fig. 1. The proposed solution incurs into small overhead as it only adds a very simple phase shifter to each antenna and an array controller per each group of n antennas. With our scheme, the system can create directional arrays and modify their structure on demand driven by the communication needs of the particular application being run, or simply remain omnidirectional.

The remainder of this paper is organized as follows. Section II presents an overview of the idea and details a potential implementation. Section III analyzes the theoretically formable patterns, which are later evaluated via full-wave simulations in Section IV. Finally, Section V concludes the paper.

II. OPPORTUNISTIC BEAMFORMING WITHIN A CHIP PACKAGE

The great majority of WNoC works consider the collocation of antennas and transceivers either to (groups of) cores [12] or to selected routers [8]–[11]. In these cases, schematically represented in Fig. 2(a), each antenna operates in isolation with a rather broad beam and must be carefully integrated to avoid undesired coupling effects. This, however, restricts the number of wireless interfaces and limits the potential of WNoC in manycore processors.

Very few works have explored the possibility of actually leveraging coupling to create arrays in chip-scale environments. Only Baniya *et al.* have proposed the integration of small arrays for beam switching in chip-to-chip communication. Their scheme, shown in Fig. 2(b), considers that groups of cores share a four-element array that can switch between different radiation directions depending on the location of the receiver. The arrays are built by design rather than opportunistically taking advantage of existing antennas, which complicates the layout and reduces the overall flexibility. Moreover, their work assumes an unconventional chip package.

Next, we provide an overview of our proposal in Section II-A, to then discuss the architecture in Section II-B.

A. Overview of the idea

We propose to take advantage of the already existing high density of antennas and, with minor changes, provide means for beamforming within a chip package. The scheme assumes

that each core (or group of cores) has its own antenna. By default, each antenna operates in isolation and can be tuned to radiate omnidirectionally to create a *broadcast channel* as depicted in Fig. 2(a). When needed, two or more antennas are activated simultaneously and form a small array that delivers a *multicast channel* through directional radiation as shown in Fig. 2(c). A controller synchronizes the transceivers to ensure that the constructive interference among antennas results into the desired directional radiation.

The solution is opportunistic and may be even regarded as partially distributed as:

- It exploits already existing antennas.
- Cores are, by definition, tightly synchronized by means of a global clock common to the whole processor.
- Data may be already present in several cores either due to existing architectural mechanisms [29], [30] or enforced by software.
- It admits a few (architecturally relevant) radiation directions, easy to derive given the destination address.

In support of this last argument, it is worth noting that parallel programming libraries include collective primitives that are used in a variety of fundamental algorithms and that generate all-to-all communication patterns [31]. In a conventional mesh NoC, collectives are generally performed within all cores of the same row first, and then within all cores of each column (or vice versa) [32], [33]. Therefore, row/column communication patterns can be architecturally relevant for WNoCs in manycore systems.

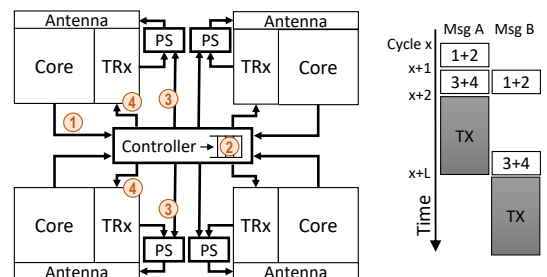


Fig. 3. Proposed architecture in a 2×2 cluster of cores (not to scale). Thick lines refer to added components. TRx stands for transceiver, whereas PS stands for phase shifter. Right plot illustrates the timing of the different steps.

B. Architecture

Figure 3 shows a schematic representation of our proposed solution, exemplified in a group of 2×2 cores. Each transceiver




Fig. 4. Theoretical gain patterns in the chip plane for two omnidirectional antennas separated by $d = \lambda/4$ with different phase shifts $\Delta\Phi$.

(TRx) is augmented with a phase shifter (PS) with a very limited number of states (initially set to 0°), whereas a controller is added to each cluster of cores. By default, cores use their omnidirectional antennas in isolation and the controller does not intervene. However, when the creation of a directional channel is required, cores communicate with the controller as depicted in Fig. 3:

- 1) The core sends the wireless packet to the controller, which places it into the queue.
- 2) Upon arrival, the controller checks the destination address and evaluates the best beam direction.
- 3) Based on the chosen direction, the controller notifies the relevant phase shifters.
- 4) While setting the phase shifters, the controller sends the wireless packet to the relevant transceivers, which modulate the information and radiate it.

These steps are analogue to the pipeline stages of a NoC router and, thus, we can use similar timings [34]. As shown in Fig. 3, steps 1-2 and 3-4 can be performed in the first and second cycle, respectively, in a pipelined fashion. In any case, wireless transmissions are typically longer than two cycles and therefore the controller does not become a bottleneck.

The beamforming policy enforced by the controller can take different forms. In a first approach, a greedy algorithm could direct the beam as close as possible to the destination, which is feasible since the positions of the controller and the destination are known and static. This, however, could increase the likelihood of collisions if not used judiciously. An alternative is to co-design the controller with the link-layer or network-layer protocols to create non-overlapping spatial channels. In the latter case, beams could be reconfigured every R cycles according to past communication demands [8].

To drive the phase shifters, the controller includes a *beam table* with {beam direction, phase shift vector} pairs. The row/column of the controller and the destination are compared, thus determining the beam direction. The number of beam directions is assumed minimal, which implies very simple phase shifters and a small beam table. As we have seen in Section II-A, such coarse-grained configuration is already architecturally relevant.

The present architecture can be scaled to larger arrays, and thus sharper beams, to increase the number of spatial channels in chips with many cores and antennas. A hierarchical controller structure or existing architectural/software mechanisms can be exploited to ensure proper data replication and antenna control in dynamic beamforming schemes for WNoC.

III. ARRAY FORMATION ANALYSIS

To analyze which are the patterns that can be formed, we resort to fundamental antenna array theory [35]. We first assume omnidirectional antennas, which in the chip scenario could be achieved with vertical monopoles. It is then considered that antennas are deployed homogeneously with a fixed distance between them, but also that the frequency of operation is a design choice, leaving antenna spacing as a parameter in terms of λ . We simplify the design space by focusing on short spacings, with the aim of (1) favoring close integration of antennas and (2) avoiding grating lobes appearing when spacing becomes larger than λ , which could create undesired interferences and complicate the architecture.

For simplicity, we start with simple two-element arrays and explore several configurations with phase shifts $\Delta\Phi$ of 0° , 90° , 180° , and 270° . The conventional choice is $d = \lambda/2$, which delivers broadside and end-fire patterns with 6 dB and 4.62 dB of peak gain, as well as beamwidths of 60° and 120° for the shifts of $\Delta\Phi = 0$ and $\Delta\Phi = 180$, respectively (patterns not shown for the sake of brevity). It is therefore a good option for row/column communications, although the flexibility is a bit limited: it does not allow to obtain single-sided patterns.

As a feasible alternative, we considered $d = \lambda/4$ which yields the patterns shown in Figure 4. Such scheme offers remarkable single-sided beams for $\Delta\Phi = 90$ and $\Delta\Phi = 270$, with 4.91 dB of peak gain, 166° beamwidth, and a front-to-back ratio of 4.67 dB. With $\Delta\Phi = 180$, the end-fire pattern reduces the beamwidth to 90° and increases the peak gain by 1.11 dB with respect to $d = \lambda/2$ because it matches with the Hansen-Woodyard condition ($\Delta\Phi \approx 2\pi d/\lambda + \pi/n$ with $n = 2$) used to optimize end-fire radiation. In the diagonal directions, where $d = \lambda\sqrt{2}/4$, both $\Delta\Phi = 0$ and $\Delta\Phi = 180$ provide interesting beams with width 92° and 104° and peak gain 4.22 dB and 5.36 dB, respectively. Should the architect need a diagonal one-sided beam, the frequency can be adjusted accordingly to achieve $d = \lambda/4$ in the diagonal direction.

IV. FAR-FIELD CHARACTERIZATION

The analysis of Section III provides interesting, but entirely theoretical design points. We confirm the results through full-wave electromagnetic simulations with CST MWS [36]. The chip package scheme shown in Fig. 1 is modeled in CST, including a 11- μm layer of silicon dioxide ($\epsilon = 3.9$, lossless) as insulator, a 700- μm layer of bulk silicon ($\epsilon = 11.9$ and resistivity $\rho = 10 \Omega\cdot\text{cm}$) as substrate, and a 200- μm layer of thermal interface material ($\epsilon = 8.6$, lossless). The top and

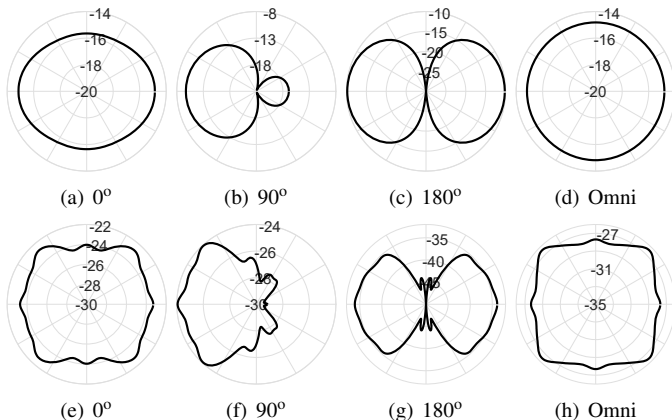


Fig. 5. Radiation patterns of a two-antenna array with separation $\lambda/4$ within a realistic chip package and surrounded by interfering antennas. Patterns (a-c) are evaluated in the near field, whereas patterns (e-g) are evaluated at a distance of 5λ . Plots (d) and (h) illustrate single-antenna patterns.

bottom boundaries (heat sink and micro-bumps, respectively) are modeled as perfect electrical conductors, whereas lateral boundaries are considered as perfect matched layers.

For this study, we choose monopole antennas because most of the power is radiated laterally towards the chip edges. Monopoles can be implemented with TSVs and their length controlled thanks to existing electroplating techniques [17]. In CST, the monopole is modeled as a thin vertical cylinder through the silicon and the length is optimized to minimize the return loss at 60 GHz. Monopole arrays are placed in the center of a 20×20 chip and are surrounded by more antennas at $\lambda/4$ distance to recreate a high-density WNoC.

Radiation patterns: we simulate two-antenna array to verify that the patterns analyzed in Sec. III are possible within the chip environment. We obtain the gain (IEEE) in the azimuthal plane, within the silicon, both near to and far from the array. Results in Fig. 5(d) and 5(h) show that when a single antenna is excited, the pattern remains roughly omnidirectional even with the presence of interfering antennas around. The main reason is that the coupling between nearby antennas is low given the presence of the lossy silicon between them. This also explains how the theoretical directional patterns can be replicated with reasonable accuracy, as shown in Fig. 5(a-c) and 5(d-f). Near the array ($\sim \lambda/2$), the lossy silicon leads to a reduction of the gain by ~ 15 dB in average. Far from the array ($\sim 5\lambda$), we observe that the gain decreases sharply for $\Delta\Phi = 180$, to the point of discouraging the use of this radiation mode. We speculate that this is due to the presence of reflections coming from the ground plane or the heat sink.

Scaling trends: we simulate linear arrays with three and four antennas to evaluate the potential of the proposed approach when scaled. Table I compares several alternatives with phase shifts $\Delta\Phi$ of 0° , 90° , or following the Hansen-Woodyard condition. The gain is measured at distance of 5λ . The omnidirectional mode improves in terms of gain, whereas the end-fire mode with $\Delta\Phi = 90$ improves in terms of beamwidth. The end-fire mode with the theoretical Hansen-Woodyard condition does not follow a clear trend. It is worth noting that other $\Delta\Phi$ values might provide better performance, but we restrict our

TABLE I
PERFORMANCE OF SCALED ARRAYS WITHIN A REALISTIC CHIP PACKAGE

Size	Phases	Radiation Type	Max Gain	Beamwidth
1	0	Omnidirectional	-26.2 dB	360°
2	0 0	Omnidirectional	-22.9 dB	360°
3	0 0 0	Omnidirectional	-21.3 dB	360°
4	0 0 0 0	Omnidirectional	-20.2 dB	360°
2	0 90	End-fire one-sided	-23.7 dB	138.7°
3	0 90 180	End-fire one-sided	-23.5 dB	122.4°
4	0 90 180 270	End-fire one-sided	-23.3 dB	109.9°
2	0 180	End-fire two-sided	-32.9 dB	112.8°
3	0 150 300	End-fire one-sided	-30 dB	120.3°
4	0 135 270 45	End-fire one-sided	-32.6 dB	193.1°




Fig. 6. Example of two parallel channels created with the proposed approach.

exploration to cases with simple phase shifters.

Overhead: for a first overhead estimation, we note that phase shifters at 60 GHz as small as 0.034 mm^2 are available in 65-nm CMOS [37]. The memory required at the controller is negligible compared to the large caches present in current multiprocessors. As justified in II-B, we assume a 2-cycle delay and no impact on the network throughput. We leave a more thorough analysis for future work.

Spatial multiplexing: the creation of directional beams allows to create multiple concurrent row/column channels that do not interfere each other. The study herein can be applied to develop a signal-to-interference model within the chip, through which a set of simple clustering and spatial multiplexing rules can be derived. Figure 6 shows a simple example where two independent channels can be created with directional radiation in two different columns. In future work, we plan to systematically analyze the possibilities in this respect.

V. CONCLUSIONS

We have presented an opportunistic scheme that leverages existing antennas in WNoC environments to create reconfigurable arrays. Albeit limited in number of beams, the proposed scheme is architecturally relevant as it can be used to implement row/column communication patterns. We simulated the feasible array configurations within a realistic chip package and found that their patterns are in close agreement with theory –although with a significantly lower efficiency due to the effects of lossy silicon and nearby antennas.

ACKNOWLEDGMENT

This work was supported by ICREA under the ICREA Academia programme, the Spanish MINECO (PCIN-2015-012), the EU's H2020 FET-OPEN program (grant 736876), and the NSF (CCF 16-29431).

REFERENCES

- [1] N. Enright Jerger, T. Krishna, and L.-S. Peh, *On-Chip Networks*, 2nd ed., 2017. [Online]. Available: <http://dx.doi.org/10.2200/S00209ED1V01Y200907CAC008>
- [2] D. Bertozzi, G. Dimitrakopoulos, J. Flich, and S. Sonntag, "The fast evolving landscape of on-chip communication," *Design Automation for Embedded Systems*, vol. 19, no. 1, pp. 59–76, 2015.
- [3] O. Markish, B. Sheinman, O. Katz, D. Corcos, and D. Elad, "On-chip mmWave Antennas and Transceivers," in *Proceedings of the NoCS '15*, 2015, p. Art. 11.
- [4] H. M. Cheema and A. Shamim, "The last barrier: On-chip antennas," *IEEE Microwave Magazine*, vol. 14, no. 1, pp. 79–91, 2013.
- [5] 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, 2015.
- [6] T. Shinde, S. Subramaniam, P. Deshmukh, M. M. Ahmed, M. Indovina, and A. Ganguly, "A 0.24 pJ/bit, 16 Gbps OOK Transmitter Circuit in 45-nm CMOS for Inter and Intra-Chip Wireless Interconnects," in *Proceedings of the GLSVLSI '18*, 2018, pp. 69–74.
- [7] D. Matolak, A. Kodi, S. Kaya, D. DiTomaso, S. Laha, and W. Rayess, "Wireless networks-on-chips: architecture, wireless channel, and devices," *IEEE Wireless Communications*, vol. 19, no. 5, 2012.
- [8] D. DiTomaso, A. Kodi, D. Matolak, S. Kaya, S. Laha, and W. Rayess, "A-WiNoC: Adaptive Wireless Network-on-Chip Architecture for Chip Multiprocessors," *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 12, pp. 3289–3302, 2015.
- [9] N. Mansoor, P. J. S. Iruthayaraj, and A. Ganguly, "Design Methodology for a Robust and Energy-Efficient Millimeter-Wave Wireless Network-on-Chip," *IEEE Transactions on Multi-Scale Computing Systems*, vol. 1, no. 1, pp. 33–45, 2015.
- [10] A. Rezaei, M. Daneshmand, M. Palesi, and D. Zhao, "Efficient Congestion-Aware Scheme for Wireless on-Chip Networks," *Proceedings of the PDP '16*, pp. 742–749, 2016.
- [11] S. Deb, K. Chang, X. Yu, S. P. Sah, M. Cosic, P. P. Pande, B. Belzer, and D. Heo, "Design of an Energy Efficient CMOS Compatible NoC Architecture with Millimeter-Wave Wireless Interconnects," *IEEE Transactions on Computers*, vol. 62, no. 12, pp. 2382–2396, 2013.
- [12] S. Abadal, J. Torrellas, E. Alarcón, and A. Cabellos-Aparicio, "OrthoNoC: A Broadcast-Oriented Dual-Plane Wireless Network-on-Chip Architecture," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 3, pp. 628–641, 2018.
- [13] Y. P. Zhang, Z. M. Chen, and M. Sun, "Propagation Mechanisms of Radio Waves Over Intra-Chip Channels With Integrated Antennas: Frequency-Domain Measurements and Time-Domain Analysis," *IEEE Transactions on Antennas and Propagation*, vol. 55, no. 10, pp. 2900–2906, 2007.
- [14] R. S. Narde, N. Mansoor, A. Ganguly, and J. Venkataraman, "On-Chip Antennas for Inter-Chip Wireless Interconnections: Challenges and Opportunities," in *Proceedings of the EuCAP '18*, 2018.
- [15] J. Wu, A. Kodi, S. Kaya, A. Louri, and H. Xin, "Monopoles Loaded with 3-D-Printed Dielectrics for Future Wireless Intra-Chip Communications," *IEEE Transactions on Antennas and Propagation*, vol. 65, no. 12, pp. 6838–6846, 2017.
- [16] W. Rayess, D. W. Matolak, S. Kaya, and A. K. Kodi, "Antennas and Channel Characteristics for Wireless Networks on Chips," *Wireless Personal Communications*, vol. 95, no. 4, pp. 5039–5056, 2017.
- [17] X. Timoneda, A. Cabellos-Aparicio, D. Manassis, E. Alarcón, and S. Abadal, "Channel Characterization for Chip-scale Wireless Communications within Computing Packages," in *Proceedings of the NOCS '18*, 2018.
- [18] S. Abadal, A. Mestres, J. Torrellas, E. Alarcón, and A. Cabellos-Aparicio, "Medium Access Control in Wireless Network-on-Chip: A Context Analysis," *IEEE Communications Magazine*, vol. 56, no. 6, pp. 172–178, 2018.
- [19] D. DiTomaso, A. Kodi, S. Kaya, and D. Matolak, "iWISE: Inter-router Wireless Scalable Express Channels for Network-on-Chips (NoCs) Architecture," in *Proceedings of the HOTI-19*. IEEE, 2011, pp. 11–18.
- [20] V. Vijayakumaran, M. P. Yuvaraj, N. Mansoor, N. Nerurkar, A. Ganguly, and A. Kwasinski, "CDMA Enabled Wireless Network-on-Chip," *ACM Journal on Emerging Technologies in Computing Systems*, vol. 10, no. 4, p. Art. 28, 2014.
- [21] D. Zhao and Y. Wang, "SD-MAC: Design and Synthesis of a Hardware-Efficient Collision-Free QoS-Aware MAC Protocol for Wireless Network-on-Chip," *IEEE Transactions on Computers*, vol. 57, no. 9, pp. 1230–1245, 2008.
- [22] H. Mondal, S. Gade, M. Shamim, S. Deb, and A. Ganguly, "Interference-Aware Wireless Network-on-Chip Architecture using Directional Antennas," *IEEE Transactions on Multi-Scale Computing Systems*, vol. 3, no. 3, pp. 193–205, 2017.
- [23] A. Mineo, M. Palesi, G. Ascia, and V. Catania, "Exploiting antenna directivity in wireless NoC architectures," *Microprocessors and Microsystems*, vol. 43, no. 6, pp. 59–66, 2016.
- [24] V. Pano, Y. Liu, I. Yilmaz, A. More, B. Taskin, and K. Dandekar, "Wireless NoCs Using Directional and Substrate Propagation Antennas," in *Proceedings of the ISVLSI '17*, 2017, pp. 188–193.
- [25] S. H. Gade, S. S. Rout, and S. Deb, "On-Chip Wireless Channel Propagation : Impact of Antenna Directionality and Placement on Channel Performance," *Proceedings of the NOCS '18*, 2018.
- [26] Z. Liu, Y. Liang, N. Li, G. Feng, H. Yu, and S. Chen, "An Energy-efficient Adaptive Sub-THz Wireless Interconnect with MIMO-Beamforming between Cores and DRAMs," in *Proceedings of the NANOCOM '16*, 2016, pp. 1–6.
- [27] P. Baniya, S. Yoo, K. L. Melde, A. Bisognin, and C. Luxey, "Switched-Beam 60-GHz Four-Element Array for Multichip Multicore System," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 8, no. 2, pp. 251–260, 2018.
- [28] X. Timoneda, S. Abadal, A. Franques, D. Manassis, J. Zhou, J. Torrellas, E. Alarcón, and A. Cabellos-Aparicio, "Engineer the Channel and Adapt to it: Enabling Wireless Intra-Chip Communication," *arXiv preprint arXiv:1901.04291*, 2018. [Online]. Available: <https://arxiv.org/pdf/1901.04291.pdf>
- [29] S. Abadal, E. Alarcón, A. Cabellos-Aparicio, and J. Torrellas, "WiSync: An Architecture for Fast Synchronization through On-Chip Wireless Communication," in *Proceedings of the ASPLOS '16*, 2016, pp. 3–17.
- [30] V. Fernando, A. Franques, S. Abadal, S. Misailovic, and J. Torrellas, "Replica: A Wireless Manycore for Communication-Intensive and Approximate Data," in *Proceedings of the ASPLOS '19*, 2019.
- [31] A. Grama, V. Kumar, A. Gupta, and G. Karypis, *Introduction to parallel computing*. Pearson Education, 2003.
- [32] N. Enright Jerger, L.-S. Peh, and M. Lipasti, "Virtual Circuit Tree Multicasting: A Case for On-Chip Hardware Multicast Support," in *Proceedings of the ISCA-35*, 2008, pp. 229–240.
- [33] T. Krishna, L.-S. Peh, B. Beckmann, and S. K. Reinhardt, "Towards the ideal on-chip fabric for 1-to-many and many-to-1 communication," in *Proceedings of the MICRO-44*, 2011, pp. 71–82.
- [34] S. Park, T. Krishna, C.-H. Chen, B. Daya, A. Chandrakasan, and L.-S. Peh, "Approaching the theoretical limits of a mesh NoC with a 16-node chip prototype in 45nm SOI," in *Proceedings of the DAC-49*, 2012, pp. 398–405.
- [35] C. A. Balanis, *Antenna Theory: Analysis and Design*, 3rd ed., Wiley, Ed., 2005.
- [36] "CST Microwave Studio." [Online]. Available: <http://www.cst.com>
- [37] F. Meng, K. Ma, K. S. Yeo, S. Xu, C. Chye Byoon, and W. Meng Lin, "Miniaturized 3-bit Phase Shifter for 60-GHz Phased-Array in 65-nm CMOS Technology," *IEEE Microwave and Wireless Components Letters*, vol. 24, no. 1, pp. 50–52, 2013.