

Engineer the Channel and Adapt to It: Enabling Wireless Intra-Chip Communication

Xavier Timoneda¹, Sergi Abadal¹, *Member, IEEE*, Antonio Franques, Dionysios Manassis, Jin Zhou², Josep Torrellas, *Fellow, IEEE*, Eduard Alarcón, and Albert Cabellos-Aparicio¹

Abstract—Ubiquitous multicore processors nowadays rely on an integrated packet-switched network for cores to exchange and share data. The performance of these intra-chip networks is a key determinant of the processor speed and, at high core counts, becomes an important bottleneck due to scalability issues. To address this, several works propose the use of mm-wave wireless interconnects for intra-chip communication and demonstrate that, thanks to their low-latency broadcast and system-level flexibility, this new paradigm could break the scalability barriers of current multicore architectures. However, these same works assume 10+ Gb/s speeds and efficiencies close to 1 pJ/bit without a proper understanding of the wireless intra-chip channel. This paper first demonstrates that such assumptions do not hold in the context of commercial chips by evaluating losses and dispersion in them. Then, we leverage the system’s monolithic nature to *engineer the channel*, this is, to optimize its frequency response by carefully choosing the chip package dimensions. Finally, we exploit the static nature of the channel to *adapt to it*, pushing efficiency-speed limits with simple tweaks at the physical layer. Our methods reduce the path loss and delay spread of a simulated commercial chip by 47 dB and 7.3×, respectively, enabling intra-chip wireless communications over 10 Gb/s and only 3.1 dB away from the dispersion-free case.

Index Terms—Dispersive channels, millimeter wave propagation, multipath interference, multiprocessor interconnection, transceivers.

I. INTRODUCTION

MULTICORE processors are present in virtually every computing domain nowadays. They integrate a number

Manuscript received August 8, 2019; revised December 8, 2019 and January 26, 2020; accepted February 8, 2020. Date of publication February 14, 2020; date of current version May 15, 2020. The authors gratefully acknowledge support from the Spanish MINECO under grant PCIN-2015-012, from the EU’s H2020 FET-OPEN program under grants No. 736876 and No. 863337, and by the Catalan Institution for Research and Advanced Studies (ICREA). The associate editor coordinating the review of this article and approving it for publication was B. Shim. (*Corresponding author: Xavier Timoneda.*)

Xavier Timoneda, Sergi Abadal, Eduard Alarcón, and Albert Cabellos-Aparicio are with the NaNoNetworking Center in Catalunya (N3Cat), Universitat Politècnica de Catalunya (UPC), 08034 Barcelona, Spain (e-mail: xavier.timoneda@upc.edu; abadal@ac.upc.edu; eduard.alarcon@upc.edu; acabello@ac.upc.edu).

Antonio Franques and Josep Torrellas are with the Department of Computer Science, University of Illinois at Urbana-Champaign (UIUC), Champaign, IL 61820 USA (e-mail: franque2@illinois.edu; torrella@illinois.edu).

Dionysios Manassis is with the System Integration & Interconnection Technologies, Fraunhofer Institute for Reliability and Microintegration (IZM), 82234 Berlin, Germany (e-mail: dionysios.manassis@izm.fraunhofer.de).

Jin Zhou is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign (UIUC), Champaign, IL 61820 USA (e-mail: jinzhou@illinois.edu).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCOMM.2020.2973988

of processor cores within the same chip and, in the past few years, manufacturers have been consistently increasing the core count seeking higher execution speeds. However, in order to translate this potential into effective performance, the on-chip communication problem must be solved: cores need an integrated interconnect to exchange or share data and, for densely populated chips, traditional interconnects are burdensome and slow down the processor. Communication, not computation, thus becomes the main performance bottleneck in multicore systems [1].

In the past, most chips did not contain more than a handful of cores and on-chip communication was easily performed through a bus. Since buses do not scale well with the number of cores, a completely different approach was soon required. The adopted solution, called Network-on-Chip (NoC), consists of a packet-switched network of routers that are co-integrated with the cores as represented in Figure 1. Since then, NoCs have been widely applied not only in research works [2]–[5], but also in commercial chips such as Tiler’s TILE-GX [6] or Intel’s Xeon Phi [7]. Nevertheless, with the arrival of extreme scaling and massive multicore architectures, standard NoCs start to show performance and efficiency issues [8]. New paradigms are thus required in the manycore era.

The scalability problems of NoCs are mainly the network diameter and overprovisioning. As further elaborated in Sec. II, these cause the communication latency and power to increase, especially for chip-wide transactions. Therefore, any new candidate to improve existing NoCs should address them and, among a few alternatives [9], Wireless Network-on-Chip (WNoC) shows great promise in this regard. In short, WNoC basically consists in overlaying a set of wireless intra-chip links over a backbone wired NoC. This reduces the latency of chip-wide transfers, including broadcasts, by virtue of the omnidirectional speed-of-light propagation of radio waves, and also combats overprovisioning thanks to its global reconfigurability [10]. As shown in the literature, these unique features become key enablers of new multicore architectures capable of pushing current scalability limits [11]–[13].

The WNoC paradigm builds on the foundations of wide-spread millimeter-wave (mm-wave) technology. A wide variety of on-chip antennas is already available [14]–[16] and wireless intra-chip communication with such antennas has been experimentally confirmed in multiple works [17]–[19]. Additionally, 60/90 GHz integrated transceivers specifically designed for WNoC have been tested [20]–[23]. On top of

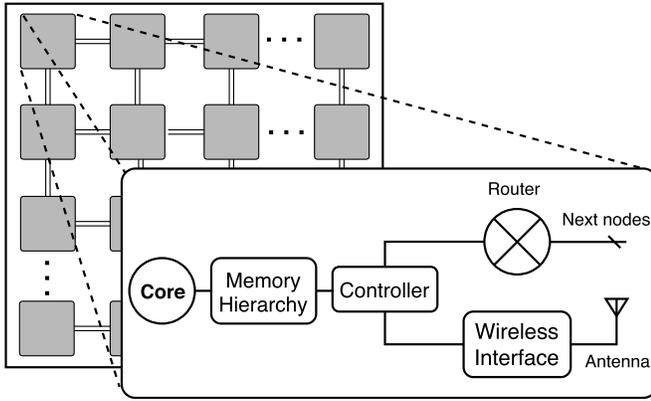


Fig. 1. Sketch of a Wireless Network-on-Chip architecture.

this, a great variety of works have evaluated new topologies and routing protocols [24]–[29] in an attempt to exploit the potential of WNoC at the network level.

The main caveat of the majority of WNoC research is that it lays on incorrect channel models. Many works [18], [30]–[37] either neglect the influence of the chip package, which introduce losses and dispersion, or directly neglect dispersion whatsoever. This does not invalidate the potential of the WNoC paradigm, but leads to erroneous assumptions on the achievable speed and power. For instance, many WNoC architectures assume bandwidths well over 10 GHz [12], [27], [28], [38], [39], which may not be achievable due to multipath effects. Other works obtain power consumption estimates by assuming path losses below 30 dB [40]–[43]. In the present study, we show that these assumptions are false for standard chip packages.

This paper aims to fill this gap and restate the potential of WNoC by proposing, as the main contribution, a novel co-design methodology that (i) properly characterizes the wireless intra-chip channel, and (ii) identifies and exploits its uniquenesses. It can be summarized in three pillars:

- **Channel characterization:** we study the propagation within a realistic computing package, which has been often overlooked. Frequency and time domain analyses are performed to extract attenuation and dispersion scaling trends. With this, we prove that the assumptions made in most WNoC works may not hold true, and that path loss and delay spread often follow contradicting trends.
- **Channel engineering:** the intra-chip channel is unique in that it can be engineered. Therefore, we propose an optimization methodology that explores the package design space to jointly minimize attenuation (path loss) and dispersion (delay spread). We illustrate the methodology by applying it in a particular chip package design and reduce the path loss and delay spread by 30 dB and $3.52\times$ together, respectively, or by 47 dB and $7.32\times$ in separated extreme cases.
- **Static transceiver optimization:** the intra-chip channel is also unique in that it is quasi-deterministic. Based on this, we propose to combat dispersion by predicting the multipath effects and adapting the transceiver

back-end to them. We easily accommodate 10 Gb/s and reach beyond the coherence bandwidth limit, figures that would be unattainable with conventional coding.

Although the static and monolithic nature of the WNoC scenario were already discussed in [24], [44], this is the first work that, to the best of the authors' knowledge, systematically exploits the unique traits of the wireless intra-chip channel. The proposed methodology could potentially lead to the conditions to operate at 10–20 Gb/s with 1–2 pJ/bit, figures that are widely assumed in the literature but that would be otherwise unattainable. It is worth noting that very few other wireless communication scenarios, if any, allow to *engineer the channel* to enhance propagation.

The remainder of this paper is organized as follows. Sec. II provides some background. Sec. III details the proposed methodology, which is then evaluated in Sec. IV. Finally, Sec. V discusses the results and Sec. VI concludes the paper.

II. BACKGROUND

Network-on-Chip: NoCs generally implement a 2-D mesh topology wherein every router is connected to a core and to its four neighbors (Fig. 2). The choice is driven by the regularity of the topology and the short path lengths, which simplifies the routers and the links. Topologies requiring long links are in fact discouraged as their energy and delay scale exponentially with length and technology [45]. Short links, however, come at the cost of a network diameter that scales as $2(k-1)$ in a $k \times k$ mesh. Thus, 64-core chips, which are commercially available [6], [7], have a network diameter of 14 hops with a chip-wide latency of several tens of nanoseconds without contention. This delay would be incurred by transmissions among far-apart cores or, even worse, broadcasts that would also increase contention as they flood the mesh. Alternatively, carefully designed WNoCs can reduce this delay to a few nanoseconds regardless of the location of data and number of destinations. This difference in performance is crucial because communications are often on the critical path of the program and any added delay can slow down execution [11].

Wireless Network-on-Chip: WNoC broadly refers to the implementation of wireless intra-chip links on top of a wired NoC. A packet arriving to a wireless interface is serialized, modulated and radiated by the antenna with a given pattern as we show in Figure 2. Radio waves propagate through the package at nearly the speed of light until reaching the intended destinations, also located within the same package, where they are demodulated and deserialized. Since intermediate router hops are avoided, WNoC reduces the latency of long-range and broadcast communications by an order of magnitude. On the downside, wireless bandwidth is limited and needs to be shared among the cores.

The physical layer of WNoC adapts to chip resource constraints. The use of mm-wave bands allows antennas to be commensurate with cores, whereas simple modulations such as On-Off Keying (OOK) are adopted to avoid bulky or power-hungry components at the transceiver. With such low order modulations, high symbol rates are needed to reach the 10+ Gb/s speeds expected for WNoC. This, together with the

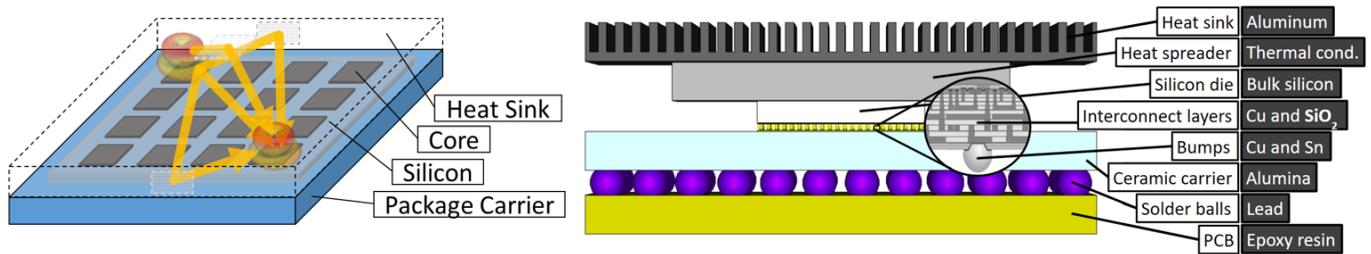


Fig. 2. Wireless propagation within a computer package (left) and typical cross-section of a flip-chip package (right).

stringent Bit Error Rate (BER) requirements of the scenario (10^{-15} to be comparable to that of a wire), makes signals particularly vulnerable to Inter-Symbol Interference (ISI). Fortunately, multipath effects can be mitigated through package-transceiver co-design, as we propose in this work.

At the MAC and network layers, WNoCs are constrained by the resource limitations and latency requirements of the scenario. Directive antennas are prohibitive, with few exceptions, rendering spatial multiplexing impractical. Thus, MAC protocols generally rely on some variant of low-latency collision-avoiding token passing to share one or a few frequency channels, thus leaving ISI as the main source of interference [25], [29]. Moreover, by design, wireless intra-chip networks are generally one-hop in an attempt satisfy the strong latency demands of multiprocessors [12], [26], [38], [39]. As a result, to support the WNoC functionalities, a routing algorithm only needs add the logic to decide when a packet should enter the wireless plane, i.e., for broadcast or long-range transmissions [24]. The simplicity at the network layer makes the wireless option more scalable than other emerging technologies, as briefly discussed in Section V.

Chip Structure and Antenna Placement: The typical cross-section of a standard chip consists of a metal stack with 5–10 layers, separated by an insulator and placed over a lossy silicon substrate [14]. Chips are then generally covered by a package that provides mechanical support and facilitates its interfacing with the rest of components. Flip-chip packages, wherein the chip is flipped over and connected to the PCB board through solder bumps, are currently widespread and preferred over wired bonding. As shown in Figure 2, the chip ends up surrounded by (i) a metallic heat sink contacted by a heat spreader and (ii) the package carrier, with several metal layers on top the PCB.

The flip-chip package does not leave much space for the antennas. Due to the presence of solder bumps, antennas cannot be implemented in the first metal layer anymore [46]. Alternatively, designers have to use the metal layers closer to the silicon or, as proposed recently, drill Through-Silicon Via (TSV) to implement vertical monopoles [47], [48]. Due to the very stringent area constraints of the scenario, directional antennas and MIMO arrays are generally prohibitive, with few exceptions [49].

Chip-scale Channel Characterization: At the chip scale, most channel characterization works have been based on full-wave simulation due to manufacturing costs and the complexity of probing in highly integrated packages [34], [35], [50]. In open packages, however, experimental works have been

more common and have shown a reasonable agreement between measurement and simulation [34], [46], [51]. Several of those works described two propagation aspects worth considering. First, the low resistivity silicon used to facilitate transistor operation introduces significant losses and, therefore, shall be avoided [14]. Second, materials used as heat spreader like Aluminum Nitride (AlN) introduce low electrical losses and, thus, would enhance propagation [31]. This opens interesting perspectives to the manufacturer, which can now take chip design decisions based on the potential for wireless intra-chip communication.

Being enclosed in a metallic package, electromagnetic propagation is confined within the limits of the package. Such field confinement has positive implications on security as eavesdropping or jamming are physically avoided, but also leads to strong multipath effects. This has been formulated by Matolak *et al.* through micro-reverberation theory [44], yet without detailing the package structure. In fact, very few studies include the chip package in their simulations or measurements and, those that do it, are limited to low frequencies or lack proper justifications on the antenna type and placement [34], [46], [51]. Others simply assume free space over the insulator layer [18], [32]–[35].

To find analogous results, we need to refer to works at the data center cabinet scale [52], or at the motherboard scale in desktops or laptops [19], [53], [54], which have structural resemblances. However, the results are not directly applicable to the chip scale due to substantial differences in dimensions, materials, and antenna placement restrictions.

Remind that, without proper understanding of the wireless channel within the package, the impact of the wireless chip-scale paradigm cannot be really assessed. In the next section, we propose a methodology to bridge this gap.

III. SYSTEM DESIGN

Our methodology provides a way to systematically co-design the chip package and the transceiver exploiting the static and monolithic nature of the system. This way, the methodology (i) validates the WNoC concept, (ii) increases the achievable data rate, and (iii) reduces the power consumed by the transceiver circuitry. Here, we first overview our proposal and then detail its design.

A. System Overview

The wireless intra-chip channel is largely unknown and prevents architects from assessing the true potential of WNoC.

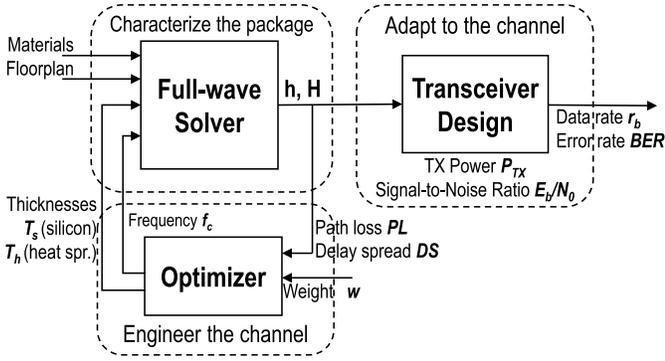


Fig. 3. Proposed methodology: characterize the package, engineer the channel, and adapt to it.

The proposed methodology, summarized in Figure 3, solves the problem in three steps.

First, a comprehensive characterization of the wireless channel within a chip package is performed. Through modeling and full-wave solving, we obtain the response of the wireless channel as a function of the position of the transmitting and receiving antennas within a 4×4 grid, as well as two parameters that chip makers can modify at design time: the frequency band and the dimensions of the package. As further elaborated in Section III-B, the results are processed to evaluate path loss and dispersion over the transmission distance.

The next step in the methodology is referred to as *channel engineering* and is uniquely suited to this monolithic system. Its main goal is to find the combination of package dimensions and frequency band that jointly minimizes path loss and dispersion. To this end, we define a figure of merit that takes both aspects into account with adjustable weights, allowing manufacturers to model the importance of power and performance in the system. This figure of merit drives an optimizer that, thanks to heuristics derived from the previous characterization process, navigates through the package design tradeoffs efficiently. The exploration is possible thanks to the use of full-wave electromagnetic simulations, which avoid the need for building multiple expensive test vehicles. More details on the methodology are given in Section III-C.

Once we have found the best package and frequency band for our purposes, we optimize the transceiver by leveraging the static nature of the channel. As shown in Section III-D, simple but effective modifications are carried out at both sides of the communication: the transmitter uses Return-to-Zero (RZ) to mitigate the ISI level, whereas the receiver uses a small and fixed set of decision thresholds to decode the current symbol based on previous bits. Both modifications are static and allow pushing the data rate beyond the typical limits imposed by the ISI (i.e., the coherence bandwidth).

B. Channel Characterization

Simulation Setup. The structure shown in Figure 2 is modeled in CST Microwave Studio [55], a full-wave solver. The silicon die has a resistivity of $10 \Omega\text{-cm}$, with $\epsilon_{r,Si} = 11.9$ whereas the heat spreader is AlN with $\epsilon_{r,AlN} = 9$ and negligible losses. To reduce the computational burden, the interconnect layers and micro-bump array are approximated as

a solid metallic element. This assumption has been validated in previous simulation works [48], [56] and is justified by the small pitch of the interconnect layers ($< 10 \mu\text{m}$) and bump array ($< 0.1 \text{ mm}$) as compared to the excitation wavelength ($\sim 1 \text{ mm}$). The antenna used for the simulations is a broadband omnidirectional aperture, which allows to focus the study on the channel effects. Unless noted, we consider a homogeneous distribution of 4×4 antennas within a $20 \times 20 \text{ mm}^2$ chip and a central frequency of 60 GHz. The minimum distance among antennas at 60 GHz is 3.57λ , where the λ is the wavelength within silicon. This distance and the high loss of silicon guarantee that there is no near-field coupling among neighboring antennas.

Frequency Domain Analysis. The full-wave solver uses the Finite Elements Method (FEM) to obtain the field distribution, the antenna gain, and the coupling between antennas in the frequency domain. Then, the channel frequency response $H_{ij}(f)$ is evaluated for each antenna pair as

$$G_i G_j |H_{ij}(f)|^2 = \frac{|S_{ji}(f)|^2}{(1 - |S_{ii}(f)|^2) \cdot (1 - |S_{jj}(f)|^2)}, \quad (1)$$

where G_i and G_j are the transmitter and receiver antenna gains, S_{ji} is the coupling between transmitter i and receiver j , whereas S_{ii} and S_{jj} are the reflection coefficients at both ends [57]. Once the whole matrix of frequency responses \mathbf{H} is obtained, a path loss analysis can be performed by fitting the attenuation L over distance d to

$$L = 10n \cdot \log_{10}(d/d_0) + L_0, \quad (2)$$

where L_0 is the path loss at the reference distance d_0 and n is the path loss exponent [18]. The path loss exponent is around 2 in free space, below 2 in guided or enclosed structures, and above 2 in lossy environments. Since losses at the channel are crucial to determine the power consumption at the transceiver (see Section V) we will report improvements in terms of worst-case L_{max} , average L_{avg} , and path loss exponent n .

Time Domain Analysis. In the time domain, we define an input excitation $x_i(t)$ at the input of the transmitting antenna i . Then, CST employs the Finite-Difference Time-Domain (FDTD) method to calculate the output signal $y_j(t)$ at the receiving antenna j . Hence, the impulse response $h_{ij}(t)$ between transmitter i and receiver j can be derived with the classical formulation

$$y_j(t) = x_i(t) \star h_{ij}(t), \quad (3)$$

where \star denotes the convolution operator. Once calculated, it is straightforward to evaluate the Power Delay Profile (PDP) in the channel between transmitter i and receiver j as

$$P_{ij}(\tau) = |h_{ij}(t, \tau)|^2, \quad (4)$$

therefore obtaining a matrix of PDP functions \mathbf{P} for all transmitters and receivers within the chip. To characterize the multipath richness of the channel, we obtain the delay spread τ_{rms} using the PDP of each channel as

$$\tau_{rms}^{(i,j)} = \sqrt{\frac{\int (\tau - \bar{\tau}_{ij})^2 P_{ij}(\tau) d\tau}{\int P_{ij}(\tau) d\tau}}, \quad (5)$$

where $\bar{\tau}_{ij} = \frac{\int \tau P_{ij}(\tau) d\tau}{\int P_{ij}(\tau) d\tau}$ is the mean delay of the channel.

In this work we will assume that all wireless channels are broadcast and, therefore, they should be operated at the lowest speed ensuring correct decoding at all nodes. As a result, we will take the worst delay spread across all pairs of transmitters-receivers (i.e., across all distances) as limiting case and use it to evaluate the coherence bandwidth B_c , as follows

$$\tau_{rms} = \max_{i,j \neq i} \tau_{rms}^{(i,j)} \Rightarrow B_c \propto \frac{1}{\tau_{rms}}. \quad (6)$$

For simplicity, we will take $B_c = \frac{1}{\tau_{rms}}$.

C. Channel Engineering

Our methodology takes path loss and delay spread as two metrics to be optimized. Since both aspects are dependent on multiple inputs, the channel engineering can be formally treated as a Multi-Objective Optimization (MOO) problem. These problems can be solved using algorithms amenable to MOO, such as evolutionary algorithms [58].

Another way to tackle the channel engineering is by reducing it to a single-objective problem using weights. In particular, our methodology defines a single custom figure of merit ϕ_w that we will attempt to maximize. Since the aim is to mitigate the path loss and the delay spread, the figure of merit takes the form

$$\phi_w = \frac{1}{PL^w DS^{(1-w)}} \quad (7)$$

where PL is the path loss metric, DS is the delay spread metric, and $w \in [0, 1]$ models the importance of power or speed in different designs. In other words, w is fixed by the architect: small values will be used in high performance devices where speed needs to be optimized over power, whereas large values imply minimization of the path loss oriented to low-power embedded systems. In this paper, our metrics are $PL = L_{avg}$ and $DS = \tau_{rms}$, with L_{avg} defined as the average path loss across all distances and τ_{rms} as defined in Equation (6). Moreover, we normalize both metrics so that they have the same dynamic range between 0 and 1.

The package engineering process as defined in this work considers three variables that can be modified at design time: the silicon thickness T_s , the heat spreader thickness T_h , and the carrier frequency f_c . Then, the objective is to maximize the figure of merit

$$\max_{T_s, T_h, f_c} \phi_w, \quad (8)$$

this is, to find the T_s , T_h , and f_c values that maximize the figure of merit for a given w and within the bounds given by the manufacturer or the architect¹ [18]. We conservatively assume $T_s \in [0.1, 0.7]$ mm and $T_h \in [0, 0.8]$ mm, which are ranges easily achievable with current silicon thinning and packaging techniques for 3D ICs [59].

¹Although this work considers three key parameters, the optimization can be extended to other decisions such as antenna placement, lateral chip dimensions, or additional material choices. For instance, the resistivity of silicon could be considered as an additional optimization knob as it largely determines the losses within silicon.

To solve the optimization problem, it is first worth noting that the full-wave simulations required to obtain ϕ_w for each $\{T_s, T_h, f_c\}$ combination are very computationally intensive, especially as f_c increases, which renders exhaustive searching impractical. Also, path loss and dispersion are related to $\{T_s, T_h, f_c\}$ in non-monotonic ways and often showing opposed trends. This creates local peaks in the ϕ_w function, thus discarding methods such as the gradient-based *hill climbing*, which tends to get stuck into local maxima.

Among the pool of optimization techniques, one alternative amenable to this problem would be Simulated Annealing (SA), which uses a probabilistic method to avoid local peaks and progressively approach a global optimum. Although SA can be modified to solve MOOs [58], we treat our problem as a single-objective optimization and use conventional SA variants. Since SA has been used in other electromagnetic problems [60], [61] and is widely known, we will not detail its implementation for the sake of brevity. We just note that the results of the channel characterization described in this work can help deriving the appropriate heuristics (e.g., candidate generation, cooling schedule) for SA to converge fast to the global optima.

D. Static Transceiver Optimization

Once the channel is engineered to minimize path loss and delay spread, we leverage the static nature of the channel to perform simple yet effective optimizations in the RF back-end. The idea is to push the symbol rates while resorting to the known, deterministic channel response to keep complexity at a minimum.

Figure 4 shows the block diagram of a typical wireless intra-chip link. As pointed out in Section II, OOK modulation is generally considered. Assuming a bit-energy of $E_b = P_{rx}/r_b$, where P_{rx} is the received power and r_b is the symbol rate, the BER of OOK is lower bounded by

$$BER_{OOK} = \frac{1}{2} \operatorname{erfc} \left(\sqrt{\frac{E_b}{4N_0}} \right) \quad (9)$$

where erfc is the complementary error function and E_b/N_0 is the signal-to-noise ratio. This bound assumes coherent detection with optimal threshold calculation and no ISI. In our case, however, ISI manifests when pushing the data rate beyond the Nyquist rate. To mitigate its effects, we propose two techniques: threshold adaptation and RZ modulation.

Threshold adaptation: The main issue in conventional wireless environments is that multipath effects are space- and time-dependent. Therefore, its impact on the Euclidean distance between the OOK symbols and on the optimal decision threshold cannot be predicted. In the worst case, ISI is modeled as added noise, reducing the noise margin and leading to an approximate BER of

$$BER_{OOK}^{isi} \approx \frac{1}{2} \operatorname{erfc} \left(\sqrt{\frac{E_b}{4(N_0 + I)}} \right) > BER_{OOK} \quad (10)$$

where I is the interference energy.

In WNoC, the channel is time-invariant and we can calculate the exact position of each symbol at all times. This means that

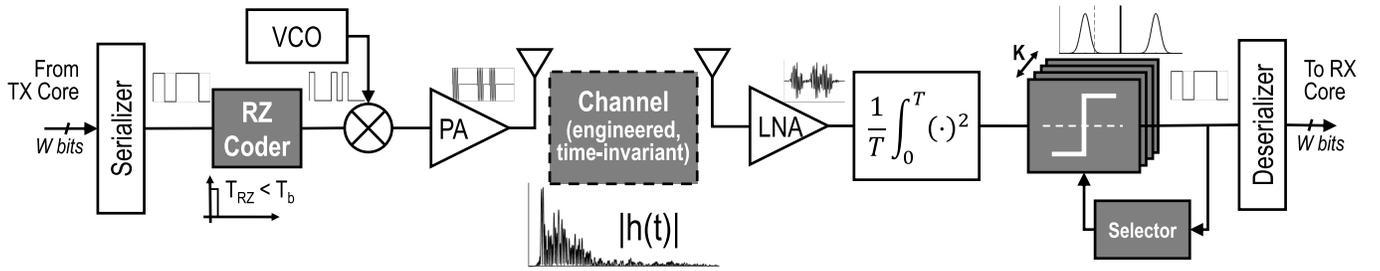


Fig. 4. Physical layer in a wireless intra-chip link with OOK and non-coherent detection. Shaded blocks identify the improvements proposed in this work.

we can find the Euclidean distance between symbols and the optimal decision threshold for any combination of previous symbols even in the presence of ISI. This information can be used to design a receiver composed by K parallel deciders, each with its own threshold, and a register that selects the appropriate leg. Assuming that with K deciders we address all ISI effects, we can approximate the BER as

$$\overline{BER}_{OOK}^{adapt} \approx \frac{1}{K} \sum_{k=1}^K \frac{1}{2} \operatorname{erfc} \left(\sqrt{\frac{\alpha_k E_b}{4N_0}} \right) \quad (11)$$

where α_k models the effect of a given past symbol combination to the Euclidean distance between current symbols. The number of required deciders scales as $K \sim \tau_{rms}/T_b$ where $T_b = 1/r_b$ is the symbol period assuming a binary modulation. In any case, the associated overheads are small compared to the cost of the RF front-end.

Return to zero: a classical way to mitigate ISI effects is by using RZ techniques, which reduce the length of the symbol through duty cycling. On the one hand, this shortens the length of the current symbol as seen by the receiver, which implies lower spillage into the next symbols. On the other hand, the lower ISI comes at the cost of a drop in the received energy, which may offset the gains of reduced ISI if RZ is not designed properly. However, since the channel is time-invariant, we can infer the duty cycle that maximizes the signal-to-interference ratio and, thus, minimizes the BER for any symbol combination. In Equation (11), this would be equivalent to increasing α_k for all k .

IV. EVALUATION

The three pillars of the proposed methodology are evaluated separately. Section IV-A discusses channel scaling trends, Section IV-B shows the gains of the channel engineering process, and Section IV-C illustrates the transceiver improvements.

A. Channel Characterization

Here, we quantify the impact of the silicon thickness T_s , the heat spreader thickness T_h and the central frequency f_c on the path loss and delay spread. Unless noted, we assume a homogeneous distribution of 4×4 antennas and take $f_c = 60$ GHz and the dimensions of a standard chip ($T_s = 0.7$ mm and $T_h = 0.2$ mm) as default values. We obtain the path loss and delay spread for all antenna pairs and perform a linear regression to obtain the dependence with distance.

Figure 5 shows the scaling trends with respect to the silicon thickness. This layer is highly lossy, as mentioned in Sec. II, and we observe that the benefits of thinning it down are significant. A 100- μm chip has a maximum path loss of $L_{max} = 36.29$ dB and a maximum delay spread of $\tau_{rms} = 0.19$ ns. Compared to a standard chip, the thinned alternative is $2.1\times$ better in terms of path loss (39 dB difference) and $2.73\times$ better in terms of worst-case delay spread (0.33 ns difference). Additionally, the path loss exponent is reduced from $n = 4.32$ to $n = 1.32$, confirming the transition from a lossy environment ($n > 2$) to a guided medium ($n < 2$). The performance also scales better in terms of delay spread, reducing the slope from 25.05 to 5.83 ps/mm.

Figure 6 repeats the analysis by varying the heat spreader thickness T_h . Given its low electrical losses, this layer can aid propagation and its inclusion is thereby highly recommended. The delay spread improves up to $3\times$ (from 0.6 to 0.2 ns) due to the presence of a stronger reflection cluster coming from the heat spreader. As for the path loss, the case here presented shows a limited impact in terms of path loss (~ 10 dB improvement in average) because most of the energy is dissipated in the 0.7-mm silicon layer before reaching the heat spreader. Although not shown due to space constraints, the effect of AIN on path loss is much more evident for thinned down silicon as the exponent drops from $n = 4.01$ (no AIN) to close to 1.1 (0.8 mm). In that case, the delay spread also oscillates between 0.2 and 0.6 ns, sometimes contradicting the path loss tendency.

Finally, Figure 7 presents the results of the frequency scaling analysis, which we limit to the 60–120 GHz span due to computational constraints. Additionally, we fix the silicon and heat spreader thicknesses to small and large values, respectively, following the design recommendations justified above. We chose this particular ($T_s = 0.3$ mm and $T_h = 0.8$ mm) because it is close to an optimal point with respect to dispersion. We find that $f_c = 110$ GHz leads to a minimum in terms of delay spread, although the improvement is limited with respect to the other frequencies. The impact on path loss, on the other hand, is substantial yet counter-intuitive at times as the average path loss drops first oscillates around 40–50 dB when shifting the frequency between 60 GHz and 90 GHz, to then increase substantially towards 90 dB at 120 GHz.

B. Engineering the Channel

Here, we show the potential of channel engineering through a partial exploration of the $\{T_h, T_s, f_c\}$ design space. Our

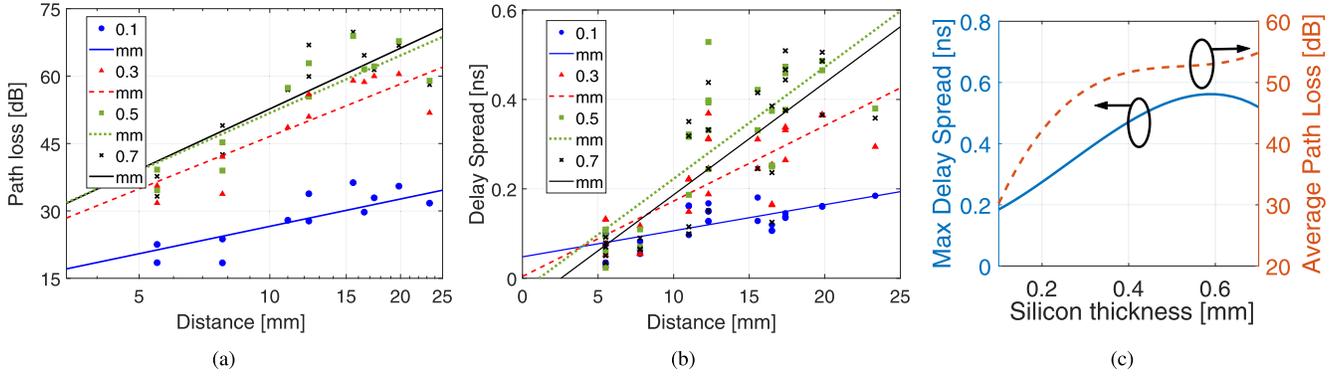


Fig. 5. Scaling analysis of the wireless channel with respect to the silicon thickness T_s for $T_h = 0.2$ mm and $f = 60$ GHz, detailing (a) path loss over distance, (b) delay spread over distance, and (c) maximum delay spread τ_{rms} and average path loss L_{avg} as functions of T_s .

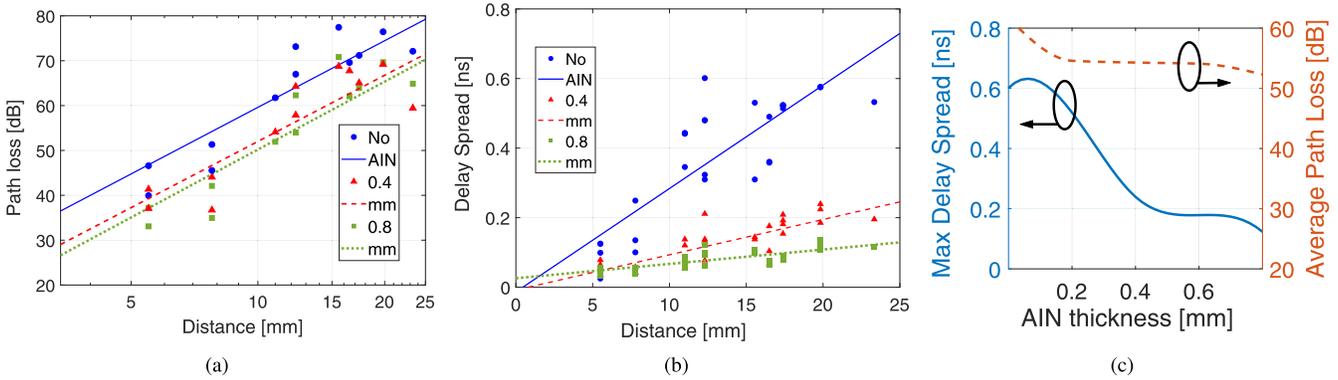


Fig. 6. Scaling analysis of the wireless channel with respect to the heat spreader thickness T_h for $T_s = 0.7$ mm and $f = 60$ GHz, detailing (a) path loss over distance, (b) delay spread over distance, and (c) maximum delay spread τ_{rms} and average path loss L_{avg} as functions of T_h .

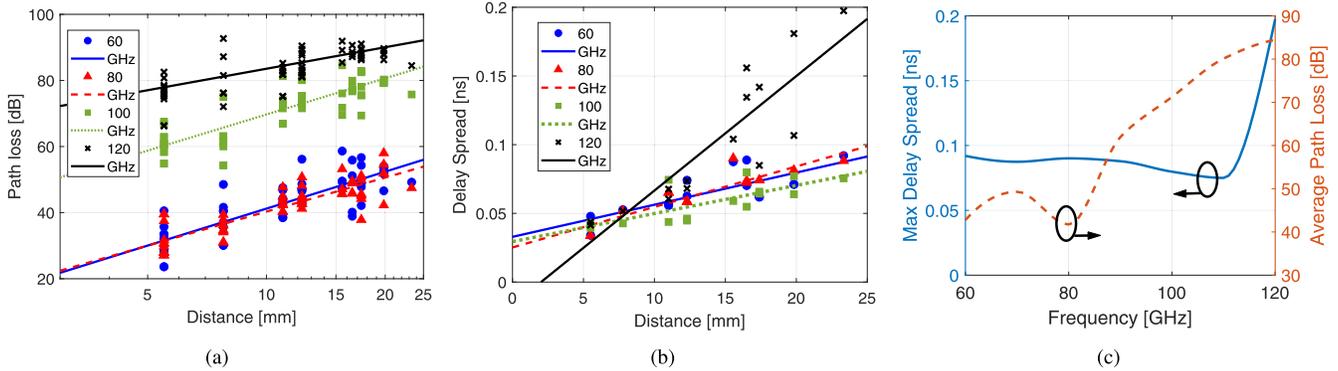


Fig. 7. Scaling analysis of the wireless channel with respect to the frequency f for $T_s = 0.3$ mm and $T_h = 0.8$ mm, detailing (a) path loss over distance, (b) delay spread over distance, and (c) maximum delay spread τ_{rms} and average path loss L_{avg} as functions of f .

aim is not to fully implement the optimizer, but rather to validate the potential of the approach by confirming both the complex interactions between inputs and the presence of local optima, as well as by giving good approximations of the path loss and delay spread improvements that we can expect.

We first plot the figure of merit ϕ_w as function of each exploration parameter while leaving the others fixed. The results, summarized in Figure 8, confirm the main lessons learned in Section IV-A: thin silicon is generally preferable

(left plot), it is hard to obtain clear tendencies with respect to the heat spreader (middle plot), and performance may plateau close to local optima (right plot). The choice of w also plays an important role in the optimization and Figure 8 also confirms it. Since path loss and delay spread often show opposed trends, the shape of ϕ changes in unexpected ways and causes wild variations in the optimal design points. Take, for instance, the frequency scaling trend. The optimal point is clearly at 110 GHz for $w = 0$, but that peak dilutes progressively and disappears around $w = 0.6$. At that point, the optimal

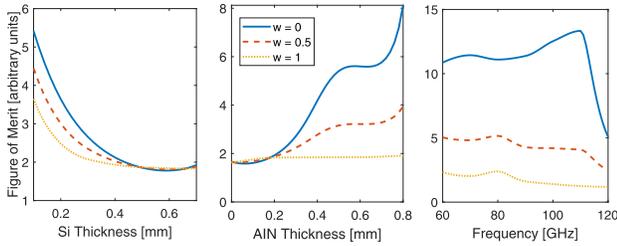


Fig. 8. Figure of merit ϕ_w as function of $\{T_s, T_h, f_c\}$ for different priority weights. Unless noted, $T_s = 0.2$ mm, $T_h = 0.7$ mm, and $f_c = 60$ GHz.

frequency becomes 60 GHz or 80 GHz due to the better path loss behavior.

In order to estimate the maximum gains that we can achieve through channel engineering, we further explored the design space in the quest for points close to a hypothetical global optima. We chose three representative values of w and compared the results with those of a standard chip ($T_s = 0.7$ mm, $T_h = 0.2$ mm, $f_c = 60$ GHz). Figure 9 and Table I illustrate the outcome of this process. There L_{max} and L_{avg} refer to the maximum and average path loss across all measured transmitter-receiver pairs within the 4×4 homogeneous grid of antennas.

We first set $w = 0$ to simulate the extreme of high performance, thereby pushing the limits on the delay spread. The peak has been found around $\{T_s = 0.3$ mm, $T_h = 0.8$ mm, $f_c = 110$ GHz} and yields a worst-case delay spread of $\tau_{rms} = 71.32$ ps for a coherence bandwidth of $B_c = 14.02$ GHz. This is roughly one order of magnitude better than the standard chip case (0.52 ns for 1.92 GHz) and confirms that the speeds assumed in the WNoC literature are feasible. In terms of path loss, this design point is also 10–15 dB better than the standard.

A second representative case would be $w = 1$, which pushes the limits on the path loss. The peak has been found by thinning the silicon down to our lower limit and using a thick spreader: $\{T_s = 0.1$ mm, $T_h = 0.8$ mm, $f_c = 60$ GHz}. This case achieves an outstanding path loss reduction of 47.07 dB for L_{max} and 32.69 dB for L_{avg} ($n = 1.32$). Further, this confirms that the path loss figures assumed in the literature, around 25–35 dB, are indeed achievable even in the presence of a chip package. However, the delay spread is maintained at the levels of the standard chip in this case.

Finally, let $w = 0.5$ to model a channel engineering process searching a balance between power and performance. In this case, a local peak has been found around the point $\{T_s = 0.1$ mm, $T_h = 0.38$ mm, $f_c = 70$ GHz}. With respect to the standard chip, this design allows to improve the coherence bandwidth B_c by $3.52\times$ and the average path loss L_{avg} by over $1.5\times$. Although this may not be a global optimum, it illustrates the potential of the methodology.

C. Static Transceiver Optimization

Since we are interested in pushing the limits of performance, this section evaluates the transceiver improvements in the package engineered for high performance. Thus, we take the

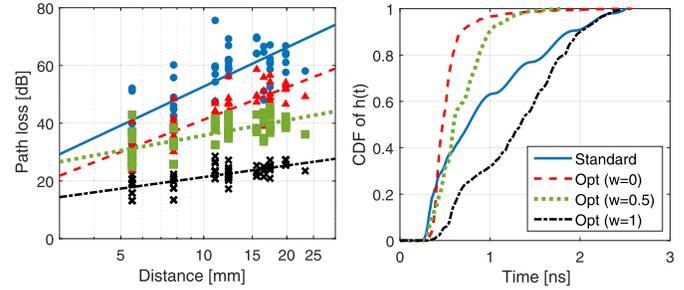


Fig. 9. Comparison between standard package ($T_s = 0.7$ mm, $T_h = 0.2$ mm, $f_c = 60$ GHz) and optimal points for three different power-weighted perspectives from the path loss (left) and delay spread perspectives (right).

TABLE I
SUMMARY OF THE OPTIMIZED PACKAGE DESIGNS

	τ_{rms} (ns)	B_c (GHz)	L_{max} (dB)	L_{avg} (dB)	n
$w = 0$	0.07	14.02	58.62	42.76	3.28
$w = 0.5$	0.15	6.76	45.49	36.48	1.74
$w = 1$	0.59	1.69	28.55	21.88	1.32
<i>Std.</i>	0.52	1.92	75.62	54.57	4.61

worst-case transient response of the $\{T_s = 0.3$ mm, $T_h = 0.8$ mm, $f_c = 110$ GHz} design point with a delay spread of $\tau_{rms} = 71.32$ ps. In all the studied cases, OOK-modulated waveforms are convoluted with the transient response at the channel and fed to the receiver, which determines the hypothetical position of the next ‘0’ or ‘1’ symbol. The BER is calculated assuming independent and equiprobable symbols.

Threshold adaptation: We simulate our proposed receiver with different number of decision thresholds K . We first obtain the threshold values by looking at the previous $\log_2(K)$ symbols and then use conventional $erfc$ formulation to derive the error probability. Figure 10(a) plots the resulting BER for a fixed r_b of 10 Gb/s, assumed in numerous WNoC works, and as a function of E_b/N_0 . Although we are below the coherence bandwidth, ISI effects disable the use of *a priori* thresholds based on steady state measurements alone. The performance for $K = 4$ is far from ideal, but starts to improve significantly. At $K = 8$, the receiver performs close to a coherent receiver in an ISI-free environment. In fact, it only needs to be 24.1 dB above the noise floor achieve the stringent BER required for WNoC (10^{-15}). This is only 3.1 dB over the ideal case.

To further evaluate the potential of the proposed scheme, we fix the received power and push the data rate way beyond the coherence bandwidth. The results, shown in Figure 10(b), reveal that the receiver by default stops working upon reaching the ISI wall at around 5 Gb/s. With as few as $K = 2$ thresholds, our proposed scheme improves the achievable data rate between 20% and 40%. Again, increasing the number of decision thresholds allows to further mitigate ISI (the bitrate increases from 7.32 up to 10.56 Gb/s at $BER = 10^{-9}$), to the point of becoming indispensable as we keep pushing the data rate. These results illustrate the tradeoff between performance and receiver complexity, although the overhead of our proposed scheme is arguably small.

Return-to-zero: One of the conclusions that can be extracted from Figures 10(a) and 10(b) is that we can

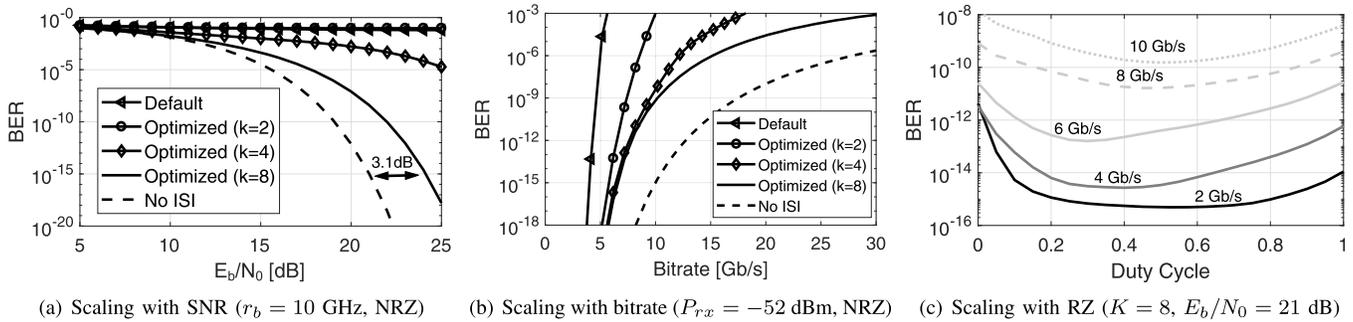


Fig. 10. Impact of transceiver optimizations on the Bit Error Rate (BER) assuming OOK modulation. NRZ stands for Non-Return-to-Zero.

minimize ISI, but we cannot get rid of it completely. The adaptive threshold moves along with the average received energy, but cannot eliminate the case where the ‘0’ and ‘1’ symbols move closer. This is precisely the case targeted by RZ. To evaluate it, we assume a receiver with $K = 8$ and set the E_b/N_0 for all transmission speeds. The results, plotted in Figure 10(c), demonstrate that there is indeed a duty cycle value that minimizes the error rate. The optimal point depends on the transmission speed and yields an improvement of up to two orders of magnitude with respect to non-RZ. The E_b/N_0 scaling analysis, not shown here in the interest of space, also revealed that RZ brings our scheme 1.2 dB closer to the ideal receiver for $BER = 10^{-15}$.

V. DISCUSSION

Impact on transmission speed. The channel engineering process, by means of substantial delay spread cuts, increases the ISI-free speed by an order of magnitude with respect to in a standard chip. Further, the transceiver optimizations have demonstrated that (i) achieving a BER of 10^{-15} at 10 GHz is affordable, and that (ii) it would be otherwise impossible. This thereby proves that our methodology enables the speeds generally assumed in the WNoC literature.

Impact on power consumption. By reducing the path loss by up to 47 dB, we achieve attenuation levels close to those assumed in recent transceiver proposals (26.5 dB in [40], [41] and 26 dB calculated with data in [42], [43]). Meeting such assumptions would lead to a bit energies of 1.95 pJ/bit for [40], [41] or 0.54 pJ/bit for [42], [43], along the lines of what is assumed in the WNoC literature. On top of that, our transceiver only needs an extra 3.1 dB of SNR to compensate for the ISI effects at 10 Gb/s and $BER = 10^{-15}$.

To make an explicit connection between channel losses and efficiency, we note that power amplifiers are the most consuming components of current transceivers, e.g., 70.8% in [42], [43]. Compensating for extra losses, noise figures, or circuit limitations would make these figures to increase even further. In fact, each amplifier has a limit P_{sat} on the output power it can provide. Going beyond that limit would require a re-design of the amplifier and, according to long-time experimentally validated scaling tendencies, the extra effort is generally paid with a reduction of the amplifier efficiency in 2.5% per each extra dBm of P_{sat} [62]. On this same direction,

it is worth noting that increasing the frequency may impact not only on the channel, but also on the efficiency of the transceiver. Although there is no technological impediment preventing the use of the 60–120 GHz band (current technologies offer f_T and f_{max} values around 300 GHz and above), pushing the frequency may initially lead to a loss of efficiency. This difference, however, levels out as technology matures and its use is extended.

Generality of results: we note that the specific results contained in this paper are, by definition, valid for a particular chip arrangement and cannot be generalized to any chip package. The key takeaways of the present work are, however, that the wireless intra-chip channel can be optimized and that the proposed methodology is applicable to any chip package. Such optimization process is unique to this wireless communications scenario.

Research directions: although this work has mitigated the intra-chip channel impairments significantly, we do not consider to have reached a lower bound. Besides the application of simulated annealing techniques to find global optima, we could improve propagation further by (i) directing certain rays via reflectors or leveraging the multiple antennas already in place to perform beamforming, (ii) thinning silicon down to the manufacturing limits [59], or (iii) exploring frequencies up to the terahertz band [63]. Additionally, factors such as the chip’s lateral dimensions, the antenna placement, or the resistivity of the silicon substrate [18] could be brought into the optimization process as long as the computational cost is affordable. At the transceiver side, low-weight coding would help minimizing the impact of ISI at very high speeds [64]. Further, compact and efficient Forward Error Correction (FEC) techniques could allow to reach the required BER without placing a large burden on the amplifiers [65], [66].

Alternative technologies: Transmission of optical signals through integrated nanophotonic waveguides [67]–[69] or of RF signals through transmission lines (TLs) [70], [71] can provide low latency and broadcast. Compared to wireless intra-chip communication, both nanophotonics and TLs are more energy efficient and provide higher bandwidth, because energy is guided rather than radiated. In this respect, the present work aims to reduce the performance and efficiency gap with respect to its alternatives. Beyond that, the main downturn of nanophotonics and TLs is the need of a physical infrastructure

to interconnect the nodes, which complicates the network design. Further, nanophotonics are less scalable due to laser power needs. Light is modulated by the transmitter and then guided to all the receivers. Each receiver extracts a fraction of the light, causing losses, and requiring high laser power for large destinations sets [72]. On the other hand, TLs are less scalable due to the need for amplifiers along the transmission line and centralized arbitration, issues that are exacerbated if the fan-out is large.

VI. CONCLUSION

Wireless intra-chip communication has been proposed as a potential solution to the scalability problems of current multicore processors. However, we have demonstrated that most works on this field are overly optimistic with regards to the channel, assuming figures one or two orders of magnitude better than what we found for a standard chip package. To further address this fundamental issue and restate the potential of WNoC, we proposed a methodology that exploits two unique traits of this new wireless scenario: its monolithic and static nature. The first allows us to engineer the channel, this is, to modify the chip package to enhance propagation in manufacturer-friendly ways. This process is applicable to any chip package and, here, we have illustrated its potential by showing improvements of 47 dB of path loss or more than 10 GHz in coherence bandwidth for a particular system. The second allows us to optimize the transceiver to mitigate multipath effects beyond the Nyquist limit. We demonstrated that we can decode OOK signals at 10 Gb/s with a BER of 10^{-15} with a signal-to-noise ratio only 3.1 dB greater than in a dispersion-free environment.

REFERENCES

- [1] R. Marculescu, U. Y. Ogras, L.-S. Peh, N. E. Jerger, and Y. Hoskote, "Outstanding research problems in NoC design: System, microarchitecture, and circuit perspectives," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 28, no. 1, pp. 3–21, Jan. 2009.
- [2] S. R. Vangal *et al.*, "An 80-tile Sub-100-W tera FLOPS processor in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 29–41, Jan. 2008.
- [3] G. P. Nychis, C. Fallin, T. Moscibroda, O. Mutlu, and S. Seshan, "On-chip networks from a networking perspective: Congestion and scalability in many-core interconnects," in *Proc. Conf. Appl., Technol., Architectures, Protocols Comput. Commun.*, 2012, pp. 407–418.
- [4] 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 *Proc. 49th Annu. Design Autom. Conf. (DAC)*, 2012, pp. 398–405.
- [5] G. Chen *et al.*, "A 340 mV-to-0.9 V 20.2 Tb/s source-synchronous hybrid packet/circuit-switched 16×16 network-on-chip in 22 nm Tri-Gate CMOS," *IEEE J. Solid-State Circuits*, vol. 50, no. 1, pp. 59–67, Jan. 2015.
- [6] D. Wentzlaff *et al.*, "On-chip interconnection architecture of the tile processor," *IEEE Micro*, vol. 27, no. 5, pp. 15–31, Sep. 2007.
- [7] G. Chrysos, "Intel Xeon phi coprocessor-the architecture," *Intel White Paper*, vol. 176, Nov. 2014. [Online]. Available: <https://software.intel.com/en-us/articles/intel-xeon-phi-coprocessor-codename-knights-corner>
- [8] D. Bertozzi, G. Dimitrakopoulos, J. Flich, and S. Sonntag, "The fast evolving landscape of on-chip communication," *Des Automat. Embed. Syst.*, vol. 19, nos. 1–2, pp. 59–76, Mar. 2015.
- [9] J. Kim, K. Choi, and G. Loh, "Exploiting new interconnect technologies in on-chip communication," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 2, pp. 124–136, Jun. 2012.
- [10] D. Matolak, A. Kodi, S. Kaya, D. DiTomaso, S. Laha, and W. Rayess, "Wireless networks-on-chips: Architecture, wireless channel, and devices," *IEEE Wireless Commun.*, vol. 19, no. 5, pp. 58–65, Oct. 2012.
- [11] S. Abadal *et al.*, "Broadcast-enabled massive multicore architectures: A wireless RF approach," *IEEE Micro*, vol. 35, no. 5, pp. 52–61, Sep. 2015.
- [12] R. G. Kim, W. Choi, Z. Chen, P. P. Pande, D. Marculescu, and R. Marculescu, "Wireless NoC and dynamic VFI codesign: Energy efficiency without performance penalty," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 24, no. 7, pp. 2488–2501, Jul. 2016.
- [13] M. A. I. Sikder, A. Kodi, W. Rayess, D. DiTomaso, D. Matolak, and S. Kaya, "Exploring wireless technology for off-chip memory access," in *Proc. IEEE 24th Annu. Symp. High-Perform. Interconnects (HOTI)*, Aug. 2016, pp. 92–99.
- [14] O. Markish, B. Sheinman, O. Katz, D. Corcos, and D. Elad, "On-chip mmWave antennas and transceivers," in *Proc. NoCS*, 2015, Art. no. 11.
- [15] H. M. Cheema and A. Shamim, "The last barrier: On-chip antennas," *IEEE Microw.*, vol. 14, no. 1, pp. 79–91, Jan. 2013.
- [16] J. Wu, A. K. Kodi, S. Kaya, A. Louri, and H. Xin, "Monopoles loaded with 3-D-printed dielectrics for future wireless intrachip communications," *IEEE Trans. Antennas Propag.*, vol. 65, no. 12, pp. 6838–6846, Dec. 2017.
- [17] B. A. Floyd, C.-M. Hung, and K. K. O, "Intra-chip wireless interconnect for clock distribution implemented with integrated antennas, receivers, and transmitters," *IEEE J. Solid-State Circuits*, vol. 37, no. 5, pp. 543–552, May 2002.
- [18] 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 Trans. Antennas Propag.*, vol. 55, no. 10, pp. 2900–2906, Oct. 2007.
- [19] H.-T. Wu, J.-J. Lin, and K. K. O, "Inter-chip wireless communication," in *Proc. EuCAP*, 2013, pp. 3647–3649.
- [20] X. Yu, J. Baylon, P. Wettin, D. Heo, P. Pratim Pande, and S. Mirabasi, "Architecture and design of multichannel millimeter-wave wireless NoC," *IEEE Des. Test*, vol. 31, no. 6, pp. 19–28, Dec. 2014.
- [21] 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 Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 34, no. 2, pp. 186–198, Feb. 2015.
- [22] S. Abadal, M. Iannazzo, M. Nemirovsky, A. Cabellos-Aparicio, H. Lee, and E. Alarcon, "On the area and energy scalability of wireless network-on-chip: A model-based benchmarked design space exploration," *IEEE/ACM Trans. Netw.*, vol. 23, no. 5, pp. 1501–1513, Oct. 2015.
- [23] A. Mineo, M. Palesi, G. Ascia, and V. Catania, "Runtime tunable transmitting power technique in mm-Wave WiNoC architectures," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 24, no. 4, pp. 1535–1545, Apr. 2016.
- [24] S. Abadal, J. Torrellas, E. Alarcon, and A. Cabellos-Aparicio, "OrthoNoC: A broadcast-oriented dual-plane wireless network-on-chip architecture," *IEEE Trans. Parallel Distrib. Syst.*, vol. 29, no. 3, pp. 628–641, Mar. 2018.
- [25] S. Deb, A. Ganguly, P. P. Pande, B. Belzer, and D. Heo, "Wireless NoC as interconnection backbone for multicore chips: Promises and challenges," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 2, pp. 228–239, Jun. 2012.
- [26] S. H. Gade and S. Deb, "HyWin: Hybrid wireless NoC with Sandboxed sub-networks for CPU/GPU architectures," *IEEE Trans. Comput.*, vol. 66, no. 7, pp. 1145–1158, Jul. 2017.
- [27] 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 Trans. Parallel Distrib. Syst.*, vol. 26, no. 12, pp. 3289–3302, Dec. 2015.
- [28] W. Choi *et al.*, "On-chip communication network for efficient training of deep convolutional networks on heterogeneous manycore systems," *IEEE Trans. Comput.*, vol. 67, no. 5, pp. 672–686, May 2018.
- [29] S. Abadal, A. Mestres, J. Torrellas, E. Alarcon, and A. Cabellos-Aparicio, "Medium access control in wireless network-on-chip: A context analysis," *IEEE Commun. Mag.*, vol. 56, no. 6, pp. 172–178, Jun. 2018.
- [30] S. Abadal, C. Han, and J. M. Jornet, "Wave propagation and channel modeling in chip-scale wireless communications: A survey from millimeter-wave to Terahertz and optics," *IEEE Access*, vol. 8, pp. 278–293, 2020.
- [31] L. Yan and G. W. Hanson, "Wave propagation mechanisms for intra-chip communications," *IEEE Trans. Antennas Propag.*, vol. 57, no. 9, pp. 2715–2724, Sep. 2009.
- [32] W.-H. Chen *et al.*, "A 6-Gb/s wireless inter-chip data link using 43-GHz transceivers and bond-wire antennas," *IEEE J. Solid-State Circuits*, vol. 44, no. 10, pp. 2711–2721, Oct. 2009.

- [33] H.-H. Yeh, N. Hiramatsu, and K. L. Melde, "The design of broadband 60 GHz AMC antenna in multi-chip RF data transmission," *IEEE Trans. Antennas Propag.*, vol. 61, no. 4, pp. 1623–1630, Apr. 2013.
- [34] R. S. Narde, J. Venkataraman, A. Ganguly, and I. Puchades, "Intra- and inter-chip transmission of millimeter-wave interconnects in NoC-based multi-chip systems," *IEEE Access*, vol. 7, pp. 112200–112215, 2019.
- [35] S. H. Gade, S. Garg, and S. Deb, "OFDM based high data rate, fading resilient transceiver for wireless networks-on-chip," in *Proc. IEEE Comput. Soc. Annu. Symp. VLSI (ISVLSI)*, Jul. 2017, pp. 483–488.
- [36] W. Rayess, D. W. Matolak, S. Kaya, and A. K. Kodi, "Antennas and channel characteristics for wireless networks on chips," *Wireless Personal Commun.*, vol. 95, no. 4, pp. 5039–5056, Aug. 2017.
- [37] I. El Masri, T. Le Gouguec, P.-M. Martin, R. Allanic, and C. Quendo, "Electromagnetic characterization of the intrachip propagation channel in Ka- and V-bands," *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 9, no. 10, pp. 1931–1941, Oct. 2019.
- [38] S. Abadal, A. Cabellos-Aparicio, E. Alarcon, and J. Torrellas, "WiSync: An architecture for fast synchronization through on-chip wireless communication," *SIGARCH Comput. Archit. News*, vol. 44, no. 2, pp. 3–17, Mar. 2016.
- [39] V. Fernando, A. Franques, S. Abadal, S. Misailovic, and J. Torrellas, "Replica: A wireless manycore for communication-intensive and approximate data," in *Proc. 24th Int. Conf. Architectural Support Program. Lang. Operating Syst. (ASPLOS)*, 2019, pp. 849–863.
- [40] X. Yu, S. P. Sah, H. Rashtian, S. Mirabbasi, P. P. Pande, and D. Heo, "A 1.2-pJ/bit 16-Gb/s 60-GHz OOK transmitter in 65-nm CMOS for wireless network-on-chip," *IEEE Trans. Microw. Theory Techn.*, vol. 62, no. 10, pp. 2357–2369, Oct. 2014.
- [41] X. Yu, H. Rashtian, S. Mirabbasi, P. P. Pande, and D. Heo, "An 18.7-Gb/s 60-GHz OOK demodulator in 65-nm CMOS for wireless network-on-chip," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 62, no. 3, pp. 799–806, Mar. 2015.
- [42] S. Subramaniam, T. Shinde, P. Deshmukh, M. S. Shamim, M. Indovina, and A. Ganguly, "A 0.36pJ/bit, 17Gbps OOK receiver in 45-nm CMOS for inter and intra-chip wireless interconnects," in *Proc. 30th IEEE Int. Syst.-Chip Conf. (SOCC)*, Sep. 2017, pp. 132–137.
- [43] T. Shinde, S. Subramaniam, P. Deshmukh, M. M. Ahmed, M. Indovina, and A. Ganguly, "A 0.24pJ/bit, 16Gbps OOK transmitter circuit in 45-nm CMOS for inter and intra-chip wireless interconnects," in *Proc. Great Lakes Symp. VLSI*, 2018, pp. 69–74.
- [44] D. W. Matolak, S. Kaya, and A. Kodi, "Channel modeling for wireless networks-on-chips," *IEEE Commun. Mag.*, vol. 51, no. 6, pp. 180–186, Jun. 2013.
- [45] *ITRS: International Technology Roadmap for Semiconductors*. Accessed: Nov. 30, 2019. [Online]. Available: <http://www.itrs2.net>
- [46] J. Branch, X. Guo, L. Gao, A. Sugavanam, J.-J. Lin, and K. K. O, "Wireless communication in a flip-chip package using integrated antennas on silicon substrates," *IEEE Electron Device Lett.*, vol. 26, no. 2, pp. 115–117, Feb. 2005.
- [47] V. Pano, I. Tekin, Y. Liu, K. R. Dandekar, and B. Taskin, "In-package wireless communication with TSV-based antenna," in *Proc. ISCAS*, 2019, pp. 19–21.
- [48] X. Timoneda *et al.*, "Millimeter-wave propagation within a computer chip package," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, May 2018, pp. 1–5.
- [49] H. K. Mondal, S. H. Gade, M. S. Shamim, S. Deb, and A. Ganguly, "Interference-aware wireless network-on-chip architecture using directional antennas," *IEEE Trans. Multi-Scale Comp. Syst.*, vol. 3, no. 3, pp. 193–205, Jul. 2017.
- [50] W. Wang *et al.*, "Wireless inter/intra-chip communication using an innovative PCB channel bounded by a metamaterial absorber," *IEEE Antennas Wireless Propag. Lett.*, vol. 15, pp. 1634–1637, 2016.
- [51] K. Kim, W. Bomstad, and K. O. Kenneth, "A plane wave model approach to understanding propagation in an intra-chip communication system," in *Proc. IEEE Antennas Propag. Soc. Int. Symp.*, Nov. 2002, pp. 166–169.
- [52] S. Khademi, S. P. Chepuri, Z. Irahauten, G. J. M. Janssen, and A.-J. van der Veen, "Channel measurements and modeling for a 60 GHz wireless link within a metal cabinet," *IEEE Trans. Wireless Commun.*, vol. 14, no. 9, pp. 5098–5110, Sep. 2015.
- [53] P. Chiang *et al.*, "Short-range, wireless interconnect within a computing chassis: Design challenges," *IEEE Des. Test. Comput.*, vol. 27, no. 4, pp. 32–43, Jul. 2010.
- [54] A. Zajic and P. Juyal, "Modeling of THz chip-to-chip wireless channels in metal enclosures," in *Proc. 12th Eur. Conf. Antennas Propag. (EuCAP)*, 2018, p. 5.
- [55] *CST Microwave Studio*. Accessed: Nov. 30, 2019. [Online]. Available: <http://www.cst.com>
- [56] X. Timoneda, A. Cabellos-Aparicio, D. Manassis, E. Alarcon, and S. Abadal, "Channel characterization for chip-scale wireless communications within computing packages," in *Proc. 20th IEEE/ACM Int. Symp. Netw.-Chip (NOCS)*, Oct. 2018, pp. 1–8.
- [57] J.-J. Lin *et al.*, "Communication using antennas fabricated in silicon integrated circuits," *IEEE J. Solid-State Circuits*, vol. 42, no. 8, pp. 1678–1687, Aug. 2007.
- [58] D. Nam and C. H. Park, "Multiobjective simulated annealing: A comparative study to evolutionary algorithms," *Int. J. Fuzzy Syst.*, vol. 2, no. 2, pp. 87–97, 2000.
- [59] F. Bieck *et al.*, "Carrierless design for handling and processing of ultrathin wafers," in *Proc. 60th Electron. Compon. Technol. Conf. (ECTC)*, Jun. 2010, pp. 316–322.
- [60] J. Simkin and C. W. Trowbridge, "Optimizing electromagnetic devices combining direct search methods with simulated annealing," *IEEE Trans. Magn.*, vol. 28, no. 2, pp. 1545–1548, Mar. 1992.
- [61] L.-S. Shu, S.-Y. Ho, and S.-J. Ho, "A novel orthogonal simulated annealing algorithm for optimization of electromagnetic problems," *IEEE Trans. Magn.*, vol. 40, no. 4, pp. 1791–1795, Jul. 2004.
- [62] H. Wang, F. Wang, H. T. Nguyen, and S. Li, *Power Amplifiers Performance Survey 2000-Present*. Accessed: Nov. 30, 2019. [Online]. Available: https://gems.ece.gatech.edu/PA_survey.html
- [63] Y. Chen, X. Cai, and C. Han, "Wave propagation modeling for mmWave and Terahertz wireless networks-on-chip communications," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2019, pp. 1–6.
- [64] J. M. Jorner and I. F. Akyildiz, "Low-weight channel coding for interference mitigation in electromagnetic nanonetworks in the Terahertz band," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2011, pp. 1–6.
- [65] F. Chang, K. Onohara, and T. Mizuochi, "Forward error correction for 100 G transport networks," *IEEE Commun. Mag.*, vol. 48, no. 3, pp. S48–S55, Mar. 2010.
- [66] A. Ganguly, P. Pande, B. Belzer, and A. Nojeh, "A unified error control coding scheme to enhance the reliability of a hybrid wireless network-on-chip," in *Proc. IEEE Int. Symp. Defect Fault Tolerance VLSI Nanotechnol. Syst.*, Oct. 2011, pp. 277–285.
- [67] C. Batten, A. Joshi, V. Stojanovic, and K. Asanovic, "Designing chip-level nanophotonic interconnection networks," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 2, pp. 137–153, Jun. 2012.
- [68] C. Sun *et al.*, "Single-chip microprocessor that communicates directly using light," *Nature*, vol. 528, no. 7583, pp. 534–538, Dec. 2015.
- [69] G. Kurian *et al.*, "ATAC: A 1000-core cache-coherent processor with on-chip optical network," in *Proc. 19th Int. Conf. Parallel Architectures Compilation Techn. (PACT)*, 2010, pp. 477–488.
- [70] A. Carpenter, J. Hu, J. Xu, M. Huang, and H. Wu, "A case for globally shared-medium on-chip interconnect," *SIGARCH Comput. Archit. News*, vol. 39, no. 3, p. 271, Jul. 2011, pp. 271–282.
- [71] W. Ding, J. Liu, M. Kandemir, and M. J. Irwin, "Traffic steering between a low-latency unswitched TL ring and a high-throughput switched on-chip interconnect," in *Proc. 22nd Int. Conf. Parallel Archit. Compilation Techn.*, Oct. 2013, pp. 309–318.
- [72] S. Abadal, A. Cabellos-aparicio, J. A. Lázaro, M. Nemirovsky, E. Alarcón, and J. Solé-Pareta, "Area and laser power scalability analysis in photonic networks-on-chip," in *Proc. 17th Int. Conf. Opt. Netw. Design Modeling (ONDM)*, Apr. 2013, pp. 131–136.



Xavier Timoneda received the degree in telecommunications systems engineering from the Universitat Politècnica de Catalunya in 2018. He performed his final degree Thesis at the University of Illinois at Urbana-Champaign (UIUC). He is currently a Research Assistant with the Universitat Politècnica de Catalunya. He has authored four scientific publications during his first year as a research assistant, and has recently coauthored a book chapter. His research interests include artificial intelligence and chip-scale communications, currently focusing his research in the development of artificial neural networks for the assembly of software-driven functional metasurfaces.



Sergi Abadal (Member, IEEE) received the B.Sc. and M.Sc. degrees in telecommunication engineering from the Universitat Politècnica de Catalunya (UPC), Barcelona, Spain, in 2010 and 2011, and the Ph.D. degree in computer architecture from UPC in 2016. During his Ph.D., he was awarded by INTEL within its Doctoral Student Honor Program. From 2009 to 2010, he was a Visiting Researcher with the Broadband Wireless Networking Laboratory, Georgia Institute of Technology, Atlanta, USA. He was a Visiting Researcher with the School of

Computer Science, University of Illinois Urbana–Champaign, in 2015 and 2018. Since October 2019, he has been a Principal Investigator of the EU H2020 Project WiPLASH. He currently works as a Post-Doctoral Researcher with the NaNonetworking Center in Catalunya (N3Cat). He has coauthored more than 60 research articles and 7 book chapters. His current research interests are ultra-high-speed on-chip wireless networks and broadcast-enabled manycore processor architectures. Since 2018, he has been an Associate Editor of *Nano Communication Networks* (Elsevier) Journal, where he was appointed as an Editor in 2019.



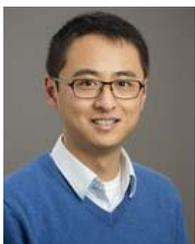
Antonio Franques received the bachelor's degree in telecommunications engineering from the Polytechnic University of Valencia (UPV), Spain. He also performed two years of research in the area of computational mathematics under the supervision of Prof. Juan Ramon Torregrosa and Prof. Alicia Cordero, as a member of the DAMRES Group. He is currently pursuing the Ph.D. degree in computer science with the University of Illinois at Urbana–Champaign (UIUC), and a member of the i-acoma Group. While pursuing the Ph.D., he also interned

twice at AMD Research, working on prototype communication hardware for exascale computing. His research focuses on the application of high-frequency wireless on-chip communication in manycore architectures. Specifically, his goal is to design new shared-memory architectures that reduce the large cost of core-to-core communication in parallel computing.



Dionysios Manassis received the M.Sc. and Ph.D. degrees in materials science & engineering from the Stevens Institute of Technology, NJ, USA. He received the Project Leadership Certificate degrees from Cornell University, NY, USA. He worked as a Technologist with Universal Instruments Corporation, NY, USA. Since 2001, he has been a Senior Technology Scientist with Fraunhofer IZM, Berlin. His main research interests lie on fine pitch flip chip and wafer level CSP bumping, solder balling, materials selection for advanced packaging technologies,

embedding processes for heterogeneous integration of components in PCBs and optical PCBs, and large scale prototype manufacturing. In the above technical fields, he has published extensively in international conferences and peer-reviewed journals.



Jin Zhou received the B.S. degree in electronics science and technology from Wuhan University, Wuhan, China, in 2008, the M.S. degree in microelectronics from Fudan University, Shanghai, China, in 2011, and the Ph.D. degree in electrical engineering from Columbia University, New York, NY, USA, in 2017. From 2011 to 2012, he also worked as an RF Integrated Circuits Design Engineer with MediaTek, Singapore. He is currently an Assistant Professor with the Electrical and Computer Engineering Department, University of Illinois at

Urbana–Champaign (UIUC). He was a recipient of the 2015–2016 Qualcomm Innovation Fellowship and the 2015–2016 IEEE Solid-State Circuits Society Predoctoral Achievement Award. He received the Eli Jury Award from the Department of Electrical Engineering, Columbia University, in 2016, for his outstanding achievement in the areas of systems, communications, signal processing, or circuits.



Josep Torrellas (Fellow, IEEE) received the Ph.D. degree from Stanford University. He is currently the Saburo Muroga Professor of computer science with the University of Illinois at Urbana–Champaign (UIUC). He is also the Director of the Center for Programmable Extreme Scale Computing and the past Director of the Illinois-Intel Parallelism Center (I2PC). He is a fellow of ACM, in 2010, and AAAS, in 2016. He was a Willett Faculty Scholar at UIUC from 2002 to 2009. As of 2016, he has graduated 36 Ph.D. students, who are now leaders

in academia and industry. He is a member of the Computing Research Association (CRA) Board of Directors. He received the IEEE Computer Society 2015 Technical Achievement Award, for pioneering contributions to shared-memory multiprocessor architectures and thread-level speculation, and the 2017 UIUC Campus Award for Excellence in Graduate Student Mentoring. He has served as the Chair for the IEEE Technical Committee on Computer Architecture (TCCA) from 2005 to 2010 and as a Council Member of CRA's Computing Community Consortium (CCC) from 2011 to 2014.



Eduard Alarcón received the Ph.D. degree in electrical engineering from the Universitat Politècnica de Catalunya in 2000. He is currently an Associate Professor with the Universitat Politècnica de Catalunya. He has coauthored more than 400 scientific publications, 8 book chapters, and 12 patents. His research interests include nanocommunications and wireless energy transfer. He was elected as an IEEE CAS Society Distinguished Lecturer, a member of the IEEE CAS Board of Governors from 2010 to 2013, an Associate Editor of the IEEE TRANSACTIONS ON

CIRCUITS AND SYSTEMS I. FUNDAMENTAL THEORY AND APPLICATIONS, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II. ANALOG AND DIGITAL SIGNAL PROCESSING, and Journal of Low Power Electronics, and an Editor-in-Chief of IEEE JOURNAL ON EMERGING AND SELECTED TOPICS IN CIRCUITS AND SYSTEMS.



Albert Cabellos-Aparicio received the Ph.D. degree in computer science engineering from the Universitat Politècnica de Catalunya in 2008. He has been a Visiting Researcher with Cisco Systems and Agilent Technologies and a Visiting Professor with KTH, Sweden, and MIT, USA. He is currently an Assistant Professor with the Universitat Politècnica de Catalunya. He is also a Co-Founder and the Scientific Director of the NaNoNetworking Center in Catalunya. He has coauthored more than 80 research articles. His research interests include

nanocommunications and software-defined networking.