

Received February 20, 2017, accepted March 4, 2017, date of publication March 31, 2017, date of current version May 17, 2017.

Digital Object Identifier 10.1109/ACCESS.2017.2689077

Terahertz Band Intra-Chip Communications: Can Wireless Links Scale Modern x86 CPUs?

VITALY PETROV¹, DMITRI MOLTCHANOV¹, MARIA KOMAR¹, ALEXANDER ANTONOV²,
PAVEL KUSTAREV², SHALOO RAKHEJA³, AND YEVGENI KOUCHERYAVY¹

¹Tampere University of Technology, 33720 Tampere, Finland

²ITMO University, 199034 St. Petersburg, Russia

³Tandon School of Engineering, New York University, New York City, NY 10003 USA

Corresponding author: V. Petrov (vitaly.petrov@tut.fi)

This work was supported by the Academy of Finland FiDiPro Program Nanocommunication Networks 2012–2016. The work of V. Petrov was supported by the Nokia Foundation. The work of S. Rakheja was supported by the National Science Foundation under Grant CCF-1565656.

ABSTRACT Massive multi-core processing has recently attracted significant attention from the research community as one of the feasible solutions to satisfy constantly growing performance demands. However, this evolution path is nowadays hampered by the complexity and limited scalability of bus-oriented intra-chip communications infrastructure. The latest advantages of terahertz (THz) band wireless communications providing extraordinary capacity at the air interface offer a promising alternative to conventional wired solutions for intra-chip communications. Still, to invest resources in this field manufacturers need a clear vision of what are the performance and scalability gains of wireless intra-chip communications. Using the comprehensive hybrid methodology combining THz ray-tracing, direct CPU traffic measurements, and cycle-accurate CPU simulations, we perform the scalability study of x86 CPU design that is backward compatible with the current x86 architecture. We show that preserving the current cache coherence protocols mapped into the star wireless communications topology that allows for tight centralized medium access control a few hundreds of active cores can be efficiently supported without any notable changes in the x86 CPU logic. This important outcome allows for incremental development, where THz-assisted x86 CPU with a few dozens of cores can serve as an intermediate solution, while the truly massive multi-core system with broadcast-enabled medium access and enhanced cache coherence protocols can be an ultimate goal.

INDEX TERMS CPU design, wireless intra-chip communications, terahertz band, massive multi-core, medium access control.

I. INTRODUCTION

Throughout the history of computer architecture evolution, the spatial distribution of computational workload has been the field of special interest. Due to the natural problems of parallel computing, such as the requirement for the workload to be parallelizable and associated increased the complexity of resource management and programming, clock frequency increase has been a straightforward way to achieve the desired speedup. Nowadays, technological issues prevent from the further increase in the clock frequency, making multiprocessing a dominant trend for personal computers. Major CPU manufacturers, Intel and AMD, presented their dual-core CPUs back in 2005, spawning the era of multi-core CPUs. Starting from a simple integration of two computing nodes on a single chip, they have evolved to truly multi-core systems featuring up to 64 computing nodes with deep integration between the components and dynamic threads redistribution between the cores [1].

In spite of the technological process evolution allowing to integrate more cores to the chip in the coming future, providing effective memory synchronization between the chip components becomes challenging with the number of cores growing. Multi-level memory hierarchies, along with multi-core architectures, form a distributed memory subsystem maintaining consistency and coherence across the memory units to make it usable for software developers. Although the properties of a memory architecture depend on its implementation, two basic trends are observed: (i) the sizes of the caches grow, making them compete with other units for the chip area [2] and (ii) communications infrastructure turns from straightforward bus-based interconnects, that suffer from scalability issues, to more complex ones. The recently proposed 3D chip design paradigm [3] can address the first issue by spreading the cores and cache memory between different layers. Meanwhile, the problem of efficient communications still remains open.

The envisioned solution for this problem is the Network-on-Chip (NoC) concept, proposing to replace the communication bus with a network of more complex topology. Although the use of efficient topologies such as meshes or small-worlds may keep the delay at the acceptable level and reduce the area used for communications infrastructure, it introduces additional challenges, such as routing and energy efficiency. As an alternative, allowing to get rid of wiring, the recently proposed wireless NoC (WNoC) paradigm promises to replace the wired connections with a wireless analog [4]. As a result, all the chip components will have only a single wireless transceiver device instead of several wired connections. The WNoC concept simplifies positioning of components, may support the existing memory coherence protocols and allows for broadcast messages. The recent advances in short-range wireless communications, especially, in terahertz (THz) technology, promise substantial amount of resources to handle bandwidth-greedy delay-critical intra-chip communications [5].

The majority of proposed WNoC architectures require a redesign of modern CPU architecture including numerous internal mechanisms. The backward-compatible WNoC-inspired CPU architecture, affordable for major vendors is still missing. To invest resources in this field, CPU vendors need a proof that the wireless links can support the constantly growing intra-chip traffic. On the other hand, the communication society needs a vision on how the prospective communication systems for wireless intra-chip communications may look like to ally their strategies with the major CPU manufacturers.

In this paper, we match the two needs together showing that the THz communications can be applied to intra-chip interfaces without dramatic changes in the modern CPU architecture. Enabling THz intra-chip communications we first propose an improved 3D CPU architecture scalable to massive core solutions that is backward compatible with the current CPU logic. We then construct the system model of the proposed architecture and use a combination of measurement- and simulation-based methodologies to parameterize it. Finally, we assess the scalability of the described solution, estimating the number of active cores that can be effectively supported without drastic changes in the CPU logic.

The major contributions of this study are as follows:

- channel modeling methodology combining ray-tracing simulations and exhaustive search to evaluate the throughput and bit error rate in intra-chip THz communications;
- traffic modeling methodology combining measurements and cycle-accurate CPU simulations to model the volume and statistical characteristics of intra-chip traffic;
- scalability study methodology of THz-assisted multi-core CPUs allowing to estimate the number of supported cores in terms of both throughput for a given traffic volume and tolerable access delay under a given MAC protocol.

The rest of the paper is organized as follows. In Section II we address the state-of-the-art and major trends in CPU design, WNoCs and THz technologies putting emphasis on THz electronics and integration. In Section III, we introduce a WNoC-inspired CPU architecture that is backward compatible with the existing CPU architectures and scalable to a massive multi-core scenario. For this architecture we carry out a detailed performance and scalability study, consecutively addressing the THz propagation in Section IV, traffic characterization in Section V, and medium access control (MAC) in Section VI. Conclusions are drawn in the last section.

II. STATE-OF-THE-ART IN RELATED TECHNOLOGIES

A. CPU ARCHITECTURE

Following the CPU evolution over the last decades, we identify a number of trends, see Fig. 1. Starting from the 1970s the metrics grew linearly allowing for a gradual increase in the CPU performance. In spite of a number of transistors per chip still growing due to advances in the manufacturing process, other characteristics nowadays face a number of challenges.

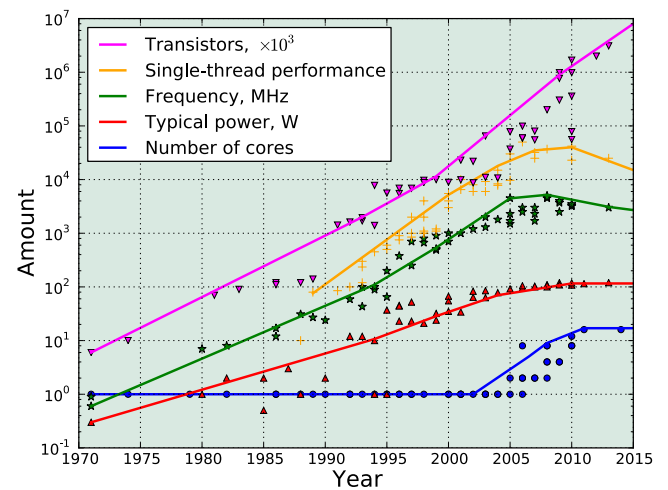


FIGURE 1. Evolution of CPU characteristics.

The first challenge is related to the maximum CPU power consumption limited by: (i) the amount of heat to be dissipated from the CPU, (ii) cooling systems efficiency. One of the design aspects affecting the amount of dissipated heat is the operational frequency. The maximum single-core CPU performance is a function of CPU clock frequency, which cannot grow much higher than 4GHz at room temperature without the principal change of the material used for transistors. Numerically, single-thread processors have generally reached their maximum performance in around 2008. The values of power consumption, clock frequency and single thread performance achieved by 2010 are the highest feasible with the current state of silicon electronics providing the trade-off between the complexity of the cooling systems design and CPU performance. Any further increase in CPU power consumption will require significant advances in these areas.

TABLE 1. Comparison of communications technologies for WNoC design.

Communication technology	Spectrum	Topology	Integration	Theoretical link capacity
Optics [6], [7]	light (430–770THz)	ring/grid/mesh/torus/tree	2D/3D	10–100Gbps
RF wireless (mmWaves) [8], [9]	30–100GHz	mesh/grid	2D	1–10Gbps
Terahertz wireless [4], [10]	0.1–10THz	star/mesh/grid	2D	0.1–2Tbps

Another challenge faced by the CPU industry is the insufficient performance of random access memory (RAM). The key problem is that the regular CPU performance is several times ahead of the RAM memory and CPU-to-RAM bus even when multiple banks are used. This issue has been addressed by the introduction of multilayer caching mechanisms, where intermediate results of CPU computation are temporarily stored in the cache memory – a small storage that has significantly higher speed and lower latency compared to RAM. Integrating the high-speed cache memory on the chip in Intel 80486 CPU in 1989 drastically boosted the speed of regular operations.

To address the abovementioned design constraints, major CPU manufacturers, Intel and AMD, switched to the multi-core design by presenting dual-core Intel Pentium D and Athlon 64 X2, respectively, in 2005. The multi-core CPUs allowed to keep pace with growing computational demands over the last decade while still preserving power/frequency/performance trade-offs achieved for a single core. The multi-core CPU design brought its own challenges that were not evident until the number of cores was kept to just a few. One of the major challenge faced by the industry nowadays is an efficient implementation of the communications infrastructure connecting the growing number of cores and allowing for their effective use.

B. NoC/WNoC DESIGN CONCEPT

Over the last decade, the computing community has been trying to address the issue of implementing massive-core CPUs by interconnecting computational elements via complex networks topologies such as grid, mesh or small-world. This approach is known as Network-on-Chip (NoC) paradigm. The major objectives in NoC research are fast data routing through the NoC, energy efficiency, and integration to existing CPU architectures. In particular, conventional lookup table-based routing algorithms are to be replaced with the faster solutions, implemented in hardware logic. The complexity of the wired switches and data buffers should be kept as low as possible to decrease the power consumption, heating, and space usage.

There are a number of examples of successful applications of NoC concept with graphical processing units (GPU) being possibly the most widely known to the large audience. The reason for an extraordinary increase in performance of such systems is mainly due to the nature of tasks allowing for their perfect parallelization leading to simple processing elements and well-defined traffic patterns. The application of the NoC concept to GP-CPU is more complex as the tasks to be performed greatly vary in their specifics, the

level of parallelization and, thus, may require an intensive exchange of information between computational elements placing additional requirements on the design of the intra-CPU communications infrastructure.

In the recent years, several solutions have been prototyped, including Sony/Toshiba/IBM Cell processor (12 cores), Tiler TILE64 chip (64 cores), and Intel Ter-aFLOPS prototype (80 cores). However, the proposed wired connectivity with a number of switches is applicable to principally new systems, designed having low-performance cores in mind. When applied to the existing general-purpose CPUs, NoC paradigm requires drastic changes in architecture, operating systems, and low-level software design. In particular, the cache structure and associated coherence protocols must be redesigned to operate over multi-hop connectivity. This places high risks on CPU manufacturers as all principally new designs have their own “maturity issues”. Extra wiring required for topology takes valuable area on a chip lowering the number of transistors available for other components. Finally, the delay requirements for information exchange inside a CPU are strict and may prohibit any kind of multi-hop communications.

The WNoC concept, where wired interfaces are replaced by miniaturized wireless transceivers, addresses some of the issues of wired NoCs [4]. Wireless communications between internal elements alleviate the problem of the area taken by wiring. The additional advantages of these systems are miniaturized transceivers [11], broadcast nature of the air interface allowing to simplify one-to-many communications [12], [13] and low power consumption [14]. In its general form, WNoC still relies on multi-hop networking paradigm requiring modifications of the CPU internal logic and bringing new challenges related to the MAC, frequency reuse, and interference mitigation [15]. However, when applied to selected interfaces only, such as the interface between private core caches and shared last level cache, this approach allows for backward compatibility with the current CPU design.

The comparison of communications technologies for NoCs is shown in Table 1. Although the use of optical interconnects has been deeply investigated so far, there are a number of open questions related to the implementation of miniaturized transceivers, multiplexers/demultiplexers, and electrical-to-optical converters. To get rid of additional infrastructure elements, the use of radio frequency (RF) has been proposed.

The modern wireless WNoC proposals mostly concentrate on the millimeter-wave (mmWave) systems in the range of 30–300GHz offering link capacity of up to few Gbps. However, the claimed rates are only achieved when complex modulation and coding schemes are used that may not be

feasible for complexity constrained transceivers. To alleviate this problem it is logical to use even higher frequencies offering the same or higher capacity with simpler of modulation and coding schemes and still benefiting from wireless nature of communications and further miniaturization of antennas.

The natural next step is the THz band, 0.1–10THz. By occupying this ultra-wide band, antennas of hundreds of micrometers in size can be used to transmit and receive data at the rate of up to few Tbps [4], [10]. As a result of recent research efforts, we have a detailed understanding of the THz propagation specifics, prototypes of micro-antennas and integrated on-chip transceivers designs. Using 0.1–0.54THz subband, 0.1aJ per symbol the results of [16] predicts the capacity of 2Tbps at the distances of up to 3 centimeters. THz technology could also be energy efficient spending just 10e–4W for communications. However, there is still no understanding which rates are sufficient for bandwidth greedy intra-chip communications.

C. THz ELECTRONIC AND PHOTONIC TECHNOLOGIES

To enable communication in the THz band, the feasibility and implementation challenges associated with the basic building blocks (amplifiers, mixers, digital-to-analog converters) of the THz communication system must be understood. From a device perspective, the required transceiver features include high output power, high sensitivity, linearity, and low noise. These transceivers will also need to overcome the high path loss of the THz band. Here, we attempt to highlight the potential, limits, and challenges of electronic and nanophotonic technologies for THz communication. For a detailed review of THz technology prospects, readers are referred to [17] and [18].

The continued growth in the performance of digital CMOS technology due to dimensional scaling has pushed the cut-off frequency (f_T) and the maximum oscillation frequency (f_{max}) of the device to several hundreds of GHz. Traditionally, the design of mixed-signal circuits at high frequencies could directly exploit the improvement in the underlying device technology to achieve higher throughput and reduced circuit footprint by simply scaling the size and values of on-chip components (transistors, capacitors, and inductors). While this approach is efficient when the operating frequency is much below the cut-off frequency, at THz frequencies this approach has diminishing improvements, which are further limited by the loss encountered in on-chip metal structures [19]. Therefore, there are two viable paths to overcome the challenges associated with building highly efficient, low-cost and low-energy on-chip THz components. The first approach relies on finding a new device technology that offers better scaling prospects and a higher intrinsic transit frequency. The alternative approach is to exploit the non-linearities of the device for efficient power generation at higher-order harmonics. Approaches that not only rely on the performance of individual devices but use novel circuits concepts to generate, radiate, and control THz frequencies have now been widely adopted in academia and industry [20].

For example, in [21] a CMOS-based oscillator with a cut-off frequency of 220 GHz and a fourth harmonic signal boosted to 870 GHz was demonstrated. In addition, THz components must also meet the area and power constraints for on-chip applications. The integration of the THz technology with existing digital CMOS technology on the same die can use the economies of scale to provide a fully integrated and cost-effective system-on-chip solution. Several innovations in device technologies and circuit design that have happened in the last decade in THz communication are briefly noted.

In [22], 65-nm digital CMOS technology was used to successfully demonstrate a 260-GHz fully integrated transceiver with an equivalent isotropically radiated power (EIRP) of 5 dBm for wireless chip-to-chip communication. The novelty of the work is to exploit quadruplers and spatial power combining with an on-chip antenna array to operate the transceiver beyond the cut-off frequency of the CMOS transistors. The authors reported total power consumption of the transmitter chain to be 688 mW, while the receiver chain consumed 465 mW. In [23], the potential of using advanced CMOS technologies for intra- and inter-chip communication over the THz band is addressed, where the authors demonstrated a fully integrated OOK transceiver that operates at 210 GHz with a power consumption of 421 mW. This work also implements an enhancement layout technique to minimize the internal parasitics. Various critical circuit functionalities for the next-generation THz microsystems, such as phase locking, beam steering, and short-pulse generation, in silicon have been demonstrated in [24]–[26]. For example, in [27] beam steering above 300 GHz in 65-nm CMOS technology with a radiated power of 0.8 mW and a maximum steering angle of 50° has been experimentally demonstrated. Silicon germanium (SiGe) technologies superior high-frequency performance and integration capability with CMOS technology and are, therefore, attractive for implementing on-chip THz components. In [28], an OOK modulator with a power consumption of 30 mW using 130-nm SiGe BiCMOS technology was demonstrated at 240 GHz. An output power of 6dBm and data rate of 13.3 Gbps were measured for their design. A fully integrated transceiver using 130-nm SiGe BiCMOS technology was demonstrated in [29]. The transceiver is operational over a distance of 10 cm with a measured EIRP of -11dBm and has a total power consumption of 380 mW.

Recently, a novel technique to generate THz radiation using Si and SiGe technologies was demonstrated in [27]. This technique is shown to radiate high power levels at a signal purity that cannot be achieved using conventional sources and phased arrays. In their technique, researchers implement a 2D phased array that operates at 338 GHz with a power dissipation of 1.54W using 65 nm bulk CMOS process. Their unique technique exploits the collective performance of a large number of synchronized or coherent sources to achieve the superior phase noise performance (-93dBc/Hz) not achievable in any other prior work. Since their central idea is based on delay-coupled oscillators, it

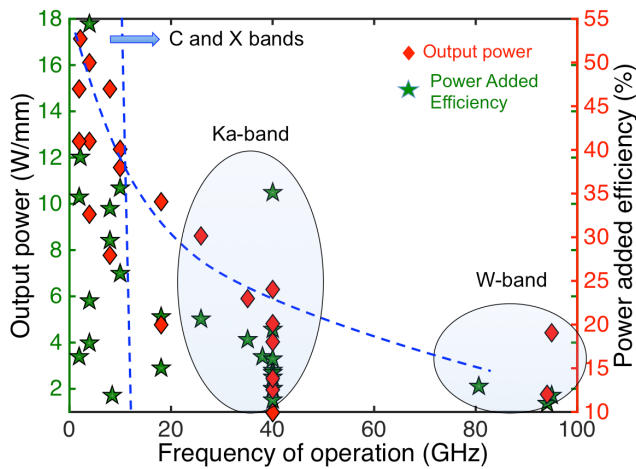


FIGURE 2. Advances in GaN technology over the last decade.

can be extended to compound semiconductor technologies to achieve an even higher power level at higher operating frequencies. In particular, III-V technologies, such as GaAs, InGaAs, GaN, and InP, have higher electron mobility and, therefore, exhibit higher cut-off frequencies. As such, compound semiconductor technologies can offer higher operating frequency and superior power output that is necessary for THz components. Recently, the microwave monolithic integrated circuits (MMICs) using 50-nm composite InGaAs devices were built to achieve 25Gbps wireless data transmission at a frequency of 220 GHz and RF radiated power of -3.4 to -1.4 dBm was reported [30]. GaN technology has recently emerged as the preferred technology to implement high-frequency power amplifiers for various commercial and military products at high frequencies [31]. GaN has a large material breakdown voltage (3MV/cm) and high electron mobility allowing GaN to achieve a high Johnson figure-of-merit of 10THz.V. The progress achieved over the last decade in the output power and power added efficiency of GaN HEMTs is illustrated in Fig. 2. GaN monolithic integrated power amplifiers operational from 75-100 GHz operation with an output power ranging from 350 mW to 840 mW and with power added efficiencies in excess of 15% have been demonstrated utilizing GaN high electron mobility transistors (HEMTs) with a channel length of 120nm [32]–[34].

Above 200 GHz, InP power amplifiers with RF power densities in the 25100-mW/mm range for HEMTs have been reported [35]. As demonstrated in [36], a THz monolithic integrated circuit (TMIC) amplifier using InP technology is able to provide a peak output power of 3mW at 650 THz. The power-added efficiencies of GaN HEMTs are comparable to those of InP HEMTs in the W-band ranging from 15-20%.

Nanoplasmonic technologies rely on the manipulation of the flow of electromagnetic radiation at the scales smaller than the wavelength of light, thereby overcoming the diffraction limit. By virtue of the very high propagation speed of electromagnetic radiation in matter, truly THz devices

that operate above 1 THz can be built. In particular, the two-dimensional carbon-based material graphene supports the propagation of surface plasmon polaritons (SPPs) over a broad frequency spectrum ranging from the microwave to the infrared. SPPs are coupled electron-light oscillations at the interface between a dielectric and a metal that can propagate at the speed of light.

Unlike metals, SPPs in graphene are electrically and chemically tunable. In graphene, the propagation length of plasmons can be several micrometers and their propagation velocity has a lower bound of $v_f/2$, where $v_f = 8 \times 10^5$ m/s is the Fermi velocity of the Dirac fermions in graphene. The long plasmon lifetime and their very high propagation velocity make graphene an ideal platform for implementing plasmonic waveguides for on-chip communication and ultra-broadband antennas for wireless communication. Moreover, such nanoantennas and transceivers with graphene SPPs offering a hundred times reduction in size compared to conventional microstrip antennas, while retaining or even exceeding the figures of merit in terms of bandwidth and gain [37]. Experimental results in SPP excitation and propagation in graphene nanostructures on various dielectric materials such as SiO₂, boron nitride, and diamond-like carbon have already been reported. Challenges related to the electrical excitation of plasma waves in graphene at room temperature must be addressed before the nanoplasmonic technology can be deployed.

D. THz INTEGRATION ISSUES

The co-integration of compound semiconductors and silicon technologies has been heavily researched in the last decade [38] primarily to develop optical devices in conjunction with CMOS digital circuits. Several techniques for the integration of III-V compounds with silicon have now been developed. For example, aspect ratio trapping (ART) technique uses trenches with a high aspect ratio to trap threading dislocations of lattice-mismatched material, yielding high-quality device layers [39]. Another approach involves the transfer of III-V device layer on silicon covered by a thin dielectric [40]. This approach is similar to silicon on insulator technology, which is a well-known technology in the industry. In [41], by using a 1nm thick InP/InAlAs composite buffer layer, InGaAs transistors on silicon have been fabricated. It is expected that the promising results of the co-integration of heterogeneous technologies will be advantageous for the proposed next-generation wireless intra-chip communication.

The design and implementation of all graphene THz front-end components, such as interconnects, phase shifters, filters, and matching networks, has been discussed in [42]. All graphene THz components can be combined with arrays of THz antennas, also implemented in graphene, to achieve dynamic beam forming and steering. Further, the even symmetry of electron and hole transport branches in graphene affords efficient and compact even harmonic generation [43]. The impact of device-level parasitics on the efficiency of




FIGURE 3. The existing and proposed CPU architectures.

frequency multiplication in graphene field-effect transistors was theoretically examined in [44]. By using antenna-coupled graphene field-effect transistors in [45], researchers demonstrate room temperature THz detectors at 0.3THz. More recently, graphene has also been successfully used to build a magnetic-field tunable THz to IR detector with the frequency of operation ranging from 0.7THz to 33THz [46]. Even though graphene is a relatively new technology, its compatibility with silicon CMOS technology [47] and unique two-dimensional physics makes it a viable device technology for the next-generation wireless communication.

III. WIRELESS 3D STACKED x86 CPU ARCHITECTURE

Analyzing the trends in CPU architecture evolution one could notice a substantial gap between academic and industrial efforts. Attempting to minimize the time-to-the-market the major vendors are conservative and will likely continue to develop x86 architecture further with hierarchical caches used for shared data exchange inside the CPU. On the other hand, the academic and long-term industrial projects are strongly pushing towards WNoC design with direct core-to-core communications as a feasible solution to enable processors with hundreds of computing elements. We suggest to address this gap proposing a hybrid CPU architecture with wireless core-to-cache interface that, on one hand, is fully compatible with existing architecture including the cache coherence mechanisms and does not require drastic changes to CPU manufacturing process and lower level software design, and, on the other hand, is scalable up to few tens or even hundreds of computing elements on a single chip. We believe that such an evolutionary architecture will motivate both academia and industry to invest more resources in this field and come up with an operational prototype within the next few years.

The existing and proposed CPU architecture are shown in Fig. 3. We suggest to separate cores and shared LLC cache (L3 in our case) into two different layers one on top of another and replace the wired communication bus connecting major components of existing CPU with a wireless ultra-wideband communication channel. This approach allows getting rid of the bus topology between cores and shared cache that may soon become bottleneck preventing the increase in the number of computing elements. At the same time, the space at the bottom layer is now free to store more L3 cache memory that nowadays takes around 40% of the chip space [2] and needs

to be scaled appropriately with the increase in the number of cores. Thus, even with a single extra layer, the amount of L3 cache can be doubled. Notice that the proposed architecture does not change the way cache memory is organized allowing to reuse the existing cache coherence protocols.

As a carrier technology for core-to-cache traffic, we propose wireless communications in the THz band, as the most capacious technology nowadays allowing for antennas of hundreds of micrometers in size and integrated transceiver electronics. To alleviate propagation losses in the inner space between two chips we suggest the surface of chips be made of soft metal with high reflection coefficient, such as copper, while the inner space could be filled with the noble gas having no absorption lines in the THz band. The use of good heat conductors such as noble gas allows for conventional fans or cooling pipes placed on top of the CPU. We envision THz technology to be sufficient for intra-chip traffic between the cores and shared L3 cache. We support this claim by performing the applicability assessment of the THz band for the proposed architecture.

IV. THz CHANNEL PROPAGATION AND PERFORMANCE

The propagation characteristics of the intra-chip communications environment are expected to be drastically different from the free-space propagation. Particularly, numerous reflections and scatterings from all the sides affect not only the amount of the energy received but also cause distortions of the received signal as a result of inter-symbol interference. The existing multi-ray channel models available in the literature, such as [48] and [49], are mostly large-scale ones characterizing channel performance at larger separation distances compared to those of intra-chip communications. On the other hand, the popular existing commercial multi-ray simulation frameworks do not take into account the specifics of the THz band, particularly, the molecular absorption. Therefore, in this study, we have applied the in-house built multi-ray simulation framework specifically tailored towards THz frequencies.

The tool is developed in C++ and Python, where the former is used for the majority of the modules to ensure the computational efficiency, while the later is used to collect and post-process the obtained results. The tool exploits the ray-tracing methodology [50] and is based on the surface tessellation to miniature segments. The size of each segment is comparable to the wavelength, thus, the segments can be




FIGURE 4. Deployment scenario for the physical layer performance study.

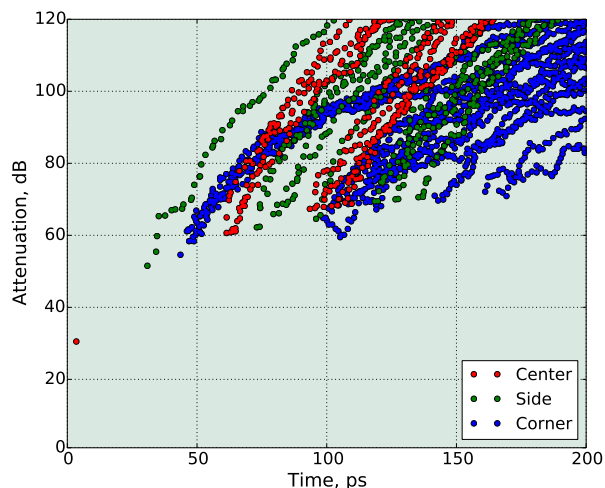


FIGURE 5. Impulse response of a channel in intra-chip THz communications.

considered as point transceivers, receiving some energy from the transmitter and reflecting/scattering a part of it to the receiver or to another point transceiver on a different surface. This approach allows evaluating all the possible reflecting and scattering paths, without the need for extra filtering on the post-processing stage. In addition, the unique feature of the tool is the ability to take into account scattering of the order greater than two within the reasonable time, which is typically not the case for the state-of-the-art commercial solutions.

The THz communications propagation environment of interest is represented by the parallelepiped of 20 mm × 20 mm × 2 mm with the target transmitter (representing the L3 cache interface transceiver) placed in the center of the top surface, as illustrated in Fig. 4. Three options for the target receiver have been considered, “center”, “side”, and “corner”, representing potential locations of the L2 cache interface transceivers. One THz bandwidth channel, 0.5–1.5THz has been considered. Due to the high complexity of the prospective THz communications equipment, a simple on-off keying (OOK) digital modulation has been assumed [51].

We start our analysis by presenting the impulse response of a channel in Fig. 5, where the dots represent the attenuation of the notable components arriving at the receivers. Expectedly,

“center” option is characterized by the best channel conditions as it has a solid and clear Line-of-Sight (LoS) component with the attenuation of around 30dB only. The other two options, “side” and “corner”, have similar attenuation of the LoS component (albeit slightly worse for “corner” due to the greater separation distance). At the same time, “corner” location is, generally, the worst due to the presence of many secondary components, reflected and scattered from the walls.

We then apply the obtained characteristics of the channel to observe the shape of the THz pulse at the receivers. Fig. 6 presents the received signal amplitude as a function of time for “center”, “side”, and “corner” options, respectively. The signal corresponding to “corner” receiver is the least attenuated, while the attenuation of the other two receiver positions is almost similar. Nevertheless, as one may observe, the “corner” receiver experiences worse propagation conditions as the reflected and scattered components remain noticeable even 100ps after the main component.

We now use the obtained received signals evolving in time to estimate the bit error rate (BER) of the OOK modulation scheme. Since the multi-path interference is highly correlated with the transmitted signal, the secondary components cannot be treated as noise, as in [52]–[54]. Therefore, the conventional approach to estimate the BER as the function of the selected modulation scheme and average signal-to-noise ratio (SNR) cannot be applied for the considered channel [51]. Below, we follow the approach originally proposed in [55] estimating BER as a function of the transmission rate by explicitly analyzing all the possible combinations of the transmitted channel symbols. We focus on 10 previous symbols, as a reasonable trade-off between the accuracy and computational complexity. For a given combination of bits, e.g., “0111010101”, and the repetition rate, τ , we explicitly construct the resulting signal at the receiver and estimate the probability of its false demodulation.

Fig. 7(a) presents the obtained results, where BER is shown as a function of inter-symbol duration and the instant value of BER is accompanied by the approximation. Expectedly, the best BER across all inter-symbol durations is observed for “center” position. However, the relative difference between “side” and “corner” positions is much less, compared to the difference between “center” and “side” options. This effect is explained by the fact that the BER for “corner” position is severely affected by the reflected and scattered components, see Fig. 6. Nevertheless, the BER for “corner” position still becomes smaller than 10^{-9} after 36ps.

We conclude the physical layer performance study by presenting the average link layer throughput, T , as a function of intra-symbol duration. We estimate throughput as $T = [1 - BER(\tau)]^l / \tau$, where τ is the inter-symbol duration and l corresponds to the block length, set to 64 bits in the present study. The results are demonstrated in Fig. 7(b). As one may observe, the throughput starts from zero for all the considered positions, as the BER value is close to 1. With the increase of inter-symbol duration the throughput grows, as the decrease




FIGURE 6. Time-domain representation of the received signals for different locations of the L2 cache interface. (a) "Center" location. (b) "Side" location. (c) "Corner" location.

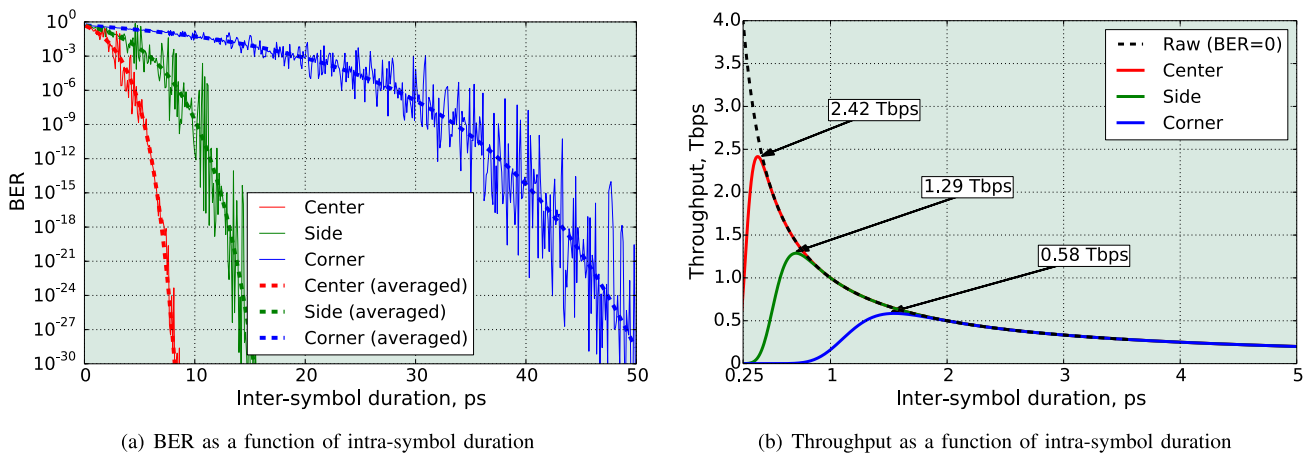


FIGURE 7. BER and throughput characteristics THz communications at the L2-L3 cache interface. (a) BER as a function of intra-symbol duration. (b) Throughput as a function of intra-symbol duration.

in the raw channel capacity is compensated by much faster decrease in BER. However, after a certain point, roughly corresponding to $BER \approx 10^{-3}$, throughput reaches its maximum value and starts decreasing to zero as BER is already negligible and the intra-symbol duration now dominates the trade-off.

The presented physical layer performance study illustrates that the throughput of a link between L2 and L3 cache interfaces ranges from $\approx 580\text{Gbps}$ to $2 \approx 2420\text{Gbps}$, depending on the L2 cache transceiver location. We will later use the former value as the conservative estimation of the link level throughput in intra-chip THz communications.

V. TRAFFIC AND INTERFACE CHARACTERISTICS

To make conclusions about the maximum number of supported cores by assessing MAC performance of the architecture we need to specify input traffic and interface characteristics. In this section, we first characterize the traffic volume using direct measurements and then characterize the principal traffic nature. Finally, we report on the results of the interface benchmarking tests providing the latency characteristics.

The closed nature of x86 CPU development does not allow to completely rely on a single approach for identifying

intra-CPU traffic properties and formulating accurate traffic models. All potential approaches need to be considered to formulate the overall understanding we can build upon to come up with a detailed description of traffic properties.

The microarchitecture-level analysis using publicly available documentation provides the first step towards a traffic model. Fixing the set of algorithms and architectural decisions helps to specify the tools needed at the later stages. First, the analysis of the functionality of a CPU and, particularly, of the cache coherence protocol, allows to understand the effect of different subsystems and make the decision about the level of detail for simulations models. The cycle-accurate system level simulations, when performed by taking into account all major mechanisms implemented in modern CPU, allows understanding the small-scale nature of the traffic at different intra-CPU interfaces. The absolute values of the traffic patterns obtained using this approach may, however, deviate from the reality due to simulation abstractions and undisclosed "know-how" in implementations. Real measurements performed by, e.g., Intel PCM [56] or a similar tool, adds to the understanding of exact values of the traffic volume at internal interfaces. The measurements are also needed to determine the delays between different cache levels and

parameterize the simulations. Using the knowledge provided by all three approaches one can come up with a detailed traffic model for modern CPUs.

A. TRAFFIC VOLUME

We performed measurements of L2-L3 traffic using Intel PCM [56]. The test bench system we used was the eight-cores Intel Core i7-5960X with Haswell architecture featuring 20Mb of L3 cache. The use of 8 cores CPU facilitates the process of calibration for traffic extrapolation. PCM allows reading the values of built-in counters in Intel CPUs including the number of L2 cache misses. Since every L2 miss is followed by the read request to L3 cache, we convert the measured values to the estimated traffic on L2-L3 interface, C , using $C = L2_M L_D / \tau$, where $L2_M$ is the number of L2 misses per measurement round, $L_D = 64$ bytes is the size of the cacheline and τ is the measurement time set to 10 minutes.

For testing purposes, we selected typical applications, a multi-player game, data encryption using AES-128, and H.264 video decoding. To simulate the highest and lowest possible loads we developed two synthetic tests performing data reading from an infinite array located in RAM, see Fig. 8. The “1B” test, emulates the “good programming style”, where the data are read sequentially byte-by-byte. In this case CPU exploits caching hierarchy, where the data are stored in 64 bytes cachelines. The second test, “64B”, reads every 65th byte making caching of data ineffective and generating much more traffic at the L2-L3 interface. This test is considered as “bad programming style”.

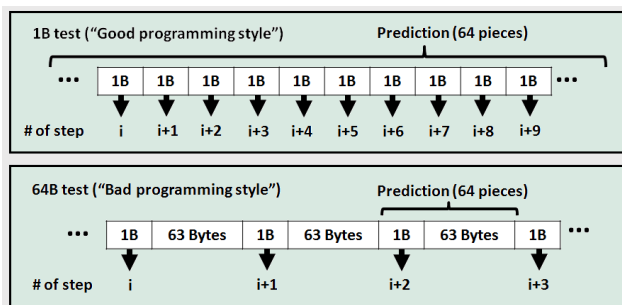


FIGURE 8. “Good” (1B) and “Bad” (64B) programming style tests.

Fig. 9 shows the obtained traffic estimates at the L2-L3 interface. As expected, “64B” test results in the highest load. We used measurements for 1, 2, 4 cores to construct extrapolations in the form An^B for 8 cores, where A and B are constants, n is the number of cores. These curves are shown by dashed lines. As one may observe, the extrapolated traffic only slightly deviates from empirical measurements for 8 cores implying that we can be fairly evident in the extrapolated results. The extrapolation of the data to the massive multi-core scenario and comparison with the achievable throughput of the wireless interface is performed in Section VI.

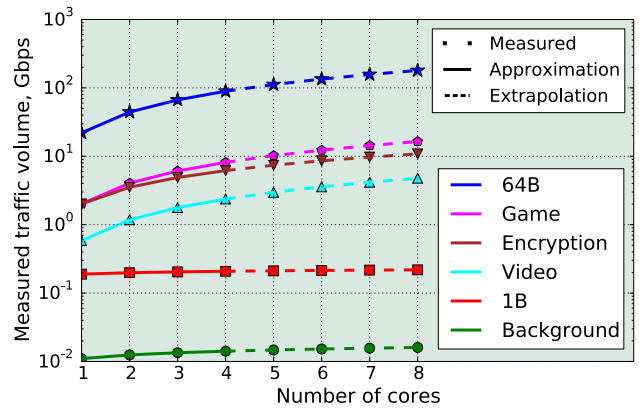


FIGURE 9. Measured L2-L3 traffic volume in Intel Core i7 CPU.

B. STOCHASTIC TRAFFIC PROPERTIES

System performance, particularly, MAC layer characteristics, are sensitive not only to the average traffic load but to the stochastic characteristics of the traffic arrival patterns. Unfortunately, neither the microarchitecture-level analysis nor direct measurements are capable of providing the detailed traffic structure at the transactional level, as both of them bring errors comparable with the measured value. One of the approaches allowing to get high-enough resolution is cycle-accurate CPU simulations. Today, there are a number of simulators supporting x86 architecture, MARSS [57], Gem5 [58], zSim [59], and SST [60]. All of them are flexible tools allowing for detailed time-stamping of events making them suitable for our task. zSim and SST are tailored at simulations of extremely large systems featuring hundreds of cores and, compared to MARSS and Gem5, lack detailed control functionality.

In this study, we rely on the methodology, originally proposed in [61], and apply it to the abovementioned set of tests. Particularly, we implemented a typical Intel x86 architecture in Gem5 including all major features and components of Intel architecture. The chosen cache size and other parameters are typical for modern general-purpose CPUs, however, slight deviations from the real values of particular systems should not change the resulting time series qualitatively. The cache subsystem was assumed to be inclusive with 64KB/2ns, 2MB/12ns and 20MB/30ns size/latency at L1, L2, and L3 caches. The model explicitly takes into account delays associated with information retrieval and emulates the pipelining capability. Systems with 1, 2, 3, 4, 8 and 16 cores have been simulated. The clock frequency was set to 3.0 GHz. Note that the clock produces a quantitative effect only and the obtained results can be scaled to almost any operational frequency. In multi-core configurations, the number of simultaneously run tests were set equal to the number of operational cores. In overall, more than 70 tests have been performed.

To model a typical load we have chosen two tests including simple reading routine (“1B” test) and more comprehensive AES encryption test involving divisions and multiplications. Fig. 10 shows the time series of the traffic at the

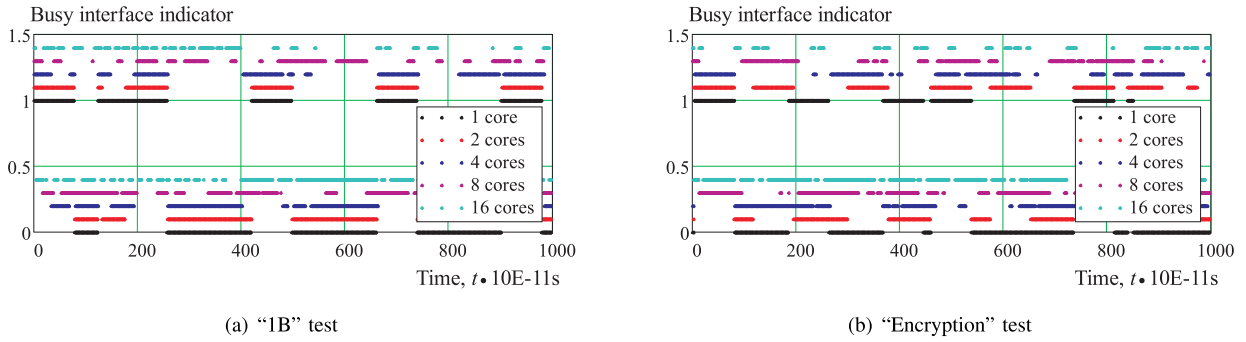


FIGURE 10. The stochastic traffic structure at L2-L3 cache interface for “1B” and “Encryption” tests. (a) “1B” test. (b) “Encryption” test.

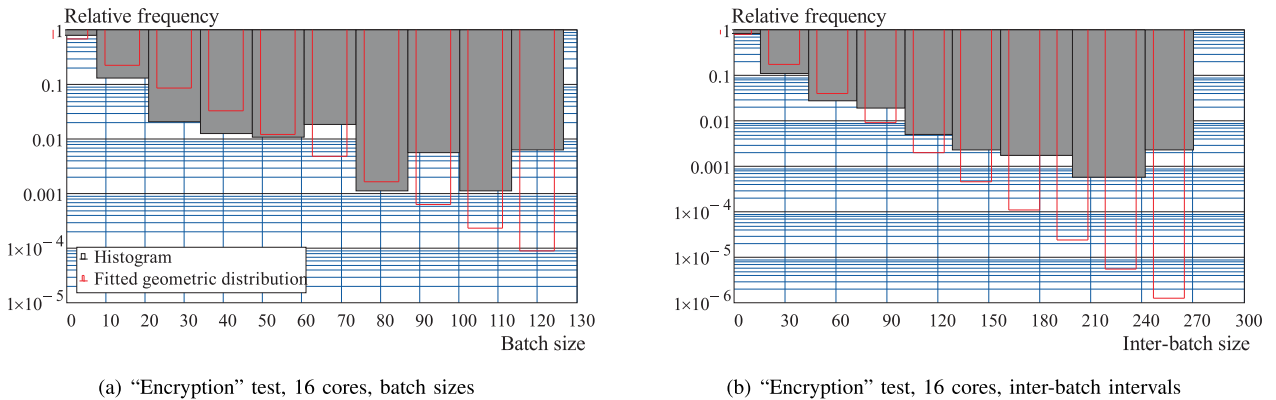


FIGURE 11. The histograms of batch sizes and inter-batch intervals and their Poisson approximations for “Encryption” test. (a) “Encryption” test, 16 cores, batch sizes. (b) “Encryption” test, 16 cores, inter-batch intervals.

L2-L3 interface for two tests, “write” and “Encryption”, for different number of cores by showing busy interface indicators $I_A + b$, where A is the event of busy interface, b is the constant added to distinguish between traces for different number of cores. Observing these data, we can make two qualitative conclusions: (i) the traffic at both interfaces has a stochastic structure and (ii) the traffic has clearly identifiable batches and gaps between them.

Fig. 11 presents the histograms of relative frequencies for batch and inter-batch time intervals, taking “Encryption” test and 16 cores as an example. As one may observe, for 16 cores the histograms have clearly observable geometrically decaying behavior. The corresponding approximations highlight that geometric distribution may provide an accurate first-order approximation for both batch sizes and inter-batch intervals. As a result, one could study the performance of the shared cache interface in prospective systems using Geo^{Geo} arrival process, that is, geometrically distributed inter-batch times batch sizes. These findings will be later used in Section VI, where the system MAC layer characteristics are studied for realistic traffic pattern.

C. DELAY PARAMETERIZATION

To complete parameterization of the model for MAC performance assessment we need to provide the tolerable delay at

the L2-L3 interface. Since Intel PCM tool, we used before, cannot report accurate estimations of the delay at the certain interface, we have developed a principally different methodology for the delay analysis. Our methodology still uses the performance counters embedded into Intel CPUs but is based on the sequential “looped” memory accesses, accompanied by the fine-grained time measurements. Particularly, delay measurements are carried out by the walk-through over the linked list structure. Cache misses are achieved using the access strides.

The conceptually similar approaches have been reported in, e.g., [62]–[64]. However, several important enhancements targeting reported data accuracy have been implemented in the applied tool. Similarly to [65], time is measured using embedded high-precision timer. In each experiment, we achieved interactions with a specific cache level only. Four programmable counters, available in the considered CPU, allow counting every cache level hit to L1, L2, L3 layers. In spite of all the measures taken to ensure the correctness of the measurements there could still be cases when the number of access attempts was not equal to the number of hits to the desired cache level. This happens due to advanced cache prediction algorithms implemented in Intel CPUs, e.g., the needed data have been loaded from lower-level cache to high-level cache, or the part of working array has been invalidated

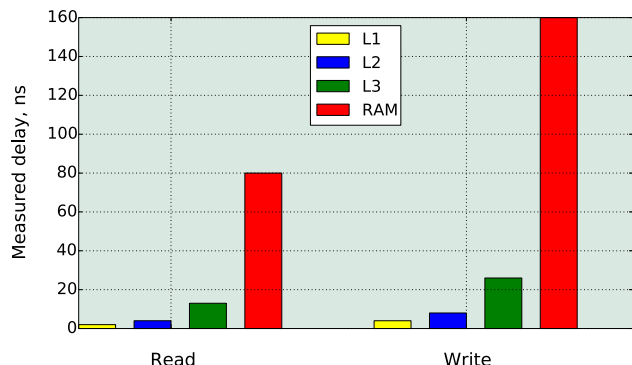


FIGURE 12. Measured access delay at the L2-L3 cache interface.

to save the free space for some new data being potentially requested. Thus, when estimating the statistics only those measurement cycles, where the number of access attempts was equal to the number of hits were taken into account.

Delay measurement are, in general, very sensitive to any background processes running in the operating system. Therefore, certain actions have been taken to ensure the accuracy of our results. First, the effects of background processes were mostly eliminated by implementing the testing program as a module of the GNU GRUB 2.0 bootloader. So, the test starts even before the operation system is loaded. Further, the entire test was explicitly written to work in a single thread, thus, avoiding ambiguities, caused by the context switching. Finally, the compiler effects have been mitigated by mostly relying on the embedded assembler, while C language was used for input and output only.

The results of delay tests for both read and write operations are presented in Fig 12. Note that the measured delay is mostly caused by the cacheline search mechanism. The requirements for the exchange of information between the core and L3 cache are extremely strict: 13ns for reading and 25ns for writing. The time to transmit a single cacheline of 64 bytes over the channel 0.5 – 1.5THz with pulse duration of 800fs is just 0.4ns, so the extra delay caused by the wireless interface (coding, propagation, decoding) should not have any significant impact on the overall delay budget, while the number of active cores in low. On the contrast, once the number of active cores is large, the medium access control protocol may drastically affect the delay performance, since the time, data fragment waits to be sent, may be substantial. We study the scaling of the described system to the massive multi-core scenario and the impact it causes on the delay values in the next section.

VI. CPU SCALING AND MAC PERFORMANCE

In this section, we apply the obtained channel, traffic, and delay findings to study the scalability of the considered CPU architecture with THz intra-chip communications. We, particularly, aim to estimate the maximum amount of cores that can be supported in terms of both throughput delay boundaries.

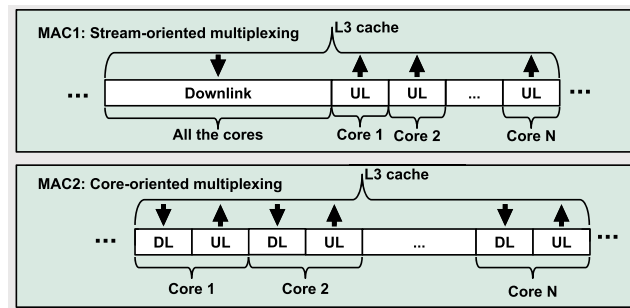


FIGURE 13. Considered MAC protocols for intra-chip communications.

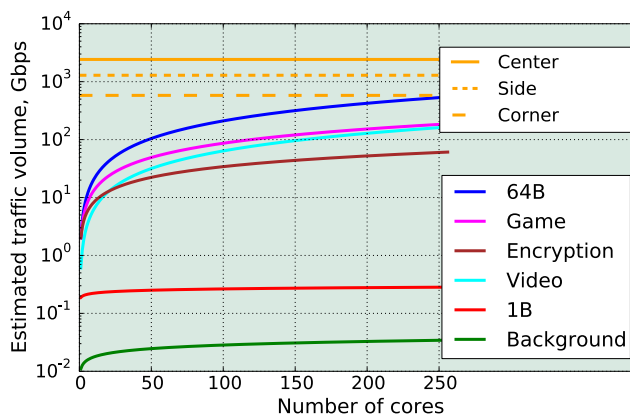


FIGURE 14. Scalability of THz-assisted multi-core CPU in terms of traffic volume.

In order to analyze the system-level performance, a concrete medium access control protocol has to be specified. The random access MAC protocols are highly unlikely to be used in intra-chip communications due to their well-known performance limitations under heavy traffic conditions [66] and the need for extremely fast random number generation for collision resolution. Furthermore, the specifics of the multi-core CPU environment with a well-defined number of communicating entities makes it, in general, more suitable for centralized or hybrid MAC protocols.

Not aiming to compare MAC protocols for intra-chip wireless communications, we specify two simplified time-division multiple access (TDMA) solutions with similar signaling (see Fig. 13). Both of the solutions realistically assume sufficiently accurate time synchronization between the nodes. In the proposed solutions, the entire time is divided into fixed size intervals or frames. Two categories of traffic (‘uplink’, from core to L3 cache and ‘downlink’, from L3 cache to core) are assumed and modeled following the traffic patterns, established in Section V. The major difference between the solutions is the particular approach for data multiplexing in the channel. The first solution suggests so-called ‘stream-oriented’ multiplexing, where every frame is first separated into two equal subframes (one for uplink, another for downlink) and the uplink subframe is then further divided into *N* slots, where *N* is the considered number of cores.




FIGURE 15. The mean wireless access delay comparison for L2-L3 read and write operations. (a) L3 read operation. (b) L3 write operation.

Meanwhile, the second or so-called “core-oriented” multiplexing, first divides the frame into N subframes, where N is the considered number of cores. On its turn, every subframe is further divided into two equal slots (one for uplink, another for downlink).

We study the worst-case scenario, with the maximum traffic load corresponding to the “64B” test in Section V. The slot duration is set equal to the time, required to transmit a single cacheline of size 64 bytes to/from the core. The slot durations are determined separately for “center”, “side”, and “corner” positions specified in Section IV.

We first analyze the CPU scalability in terms of the total throughput provided at the MAC layer. Fig. 14, on one hand, illustrates the estimated traffic volume and, on the other, provides the throughput bounds for three options of core locations inside the chip (582Gbps, 1287Gbps, and 2422Gbps for “corner”, “side”, and “center” location, respectively). Based on the observed results, we conclude that in terms of the total throughput the proposed architecture could be scaled up to 250 active cores for the corner case. Since in real massive multi-core CPUs only a small fraction of cores will have the channel conditions similar to the corner ones, this result is a pessimistic estimate, implying that the actual value could be much greater.

The throughput-based assessment may provide the optimistic upper bound on the number of supported cores even when all the cores are in pessimistic “corner” positions. The reason is that as the number of active cores grows, the total duration of the MAC frame increases resulting in additional access delay. Furthermore, the stochastic nature of traffic may further affect the service performance of cores. To address this question, we now report the simulation results of CPU scaling by taking into account specifics of the channel propagation and traffic arrival process revealed in Section IV and Section V, respectively, as well as the details of the MAC protocols.

Fig. 15 shows the MAC share of the wireless access delay with respect to the L2-L3 access delay estimated in

Section V (see Fig. 12) for both read (downlink) and write (uplink) operations. We, first, introduce the boundaries for the L2-L3 wireless access delay, by setting them to 10%, 33%, and 100% of the L2-L3 access delay for so-called “negligible”, “tolerable”, and “unacceptable” regimes, respectively. The reason for this taxonomy is that additional interface should not bring any significant impact to the total delay budget. We analyze four cases: “center”, “side”, and “corner” core location options introduced earlier, as well as balanced, “hybrid” case, where the semi-square deployment of N cores is assumed. The system setup in “hybrid” case consists of 4 “corner” cores, $4\lceil\sqrt{N}\rceil - 4$ “side” cores, while the rest cores are at the “center” location.

Analyzing the obtained results, we observe that “corner” deployment cannot scale well to massive multi-core scenarios, as the wireless link delay becomes greater than the entire measured delay for L3 access already after 65 and 87 cores for read and write operations, respectively. In contrast, “center” channel conditions allow the system to scale up to 130/200 cores for read/write operations without violating “tolerable” level. It is important to note, the “hybrid” system scales well, reaching 101/134 cores for read/write operations while keeping the delay lower than 1/3 of the L2-L3 access delay reported in Fig. 12. As one may observe, the delay scalability limits are much stricter than throughput-base ones (i.e., 101 cores versus 250 cores), implying that the delay requirements have to be taken into account in CPU scalability studies.

VII. CONCLUSIONS

THz wireless communications are envisioned as a promising technology to satisfy the growing capacity demands of a broad range of prospective applications: from ultra-dense network deployments up to board-to-board and intra-chip communications. Particularly, the massive multi-core CPUs can become reality when facilitated by the THz band intra-chip communication interface. At the same time, the introduction of THz communications inside CPU and following adaptation

of state-of-the-art mechanisms to fully benefit from this new interface may potentially cause the entire system re-design, from broadcast- and multicast-aware data exchange methods to enhanced cache coherence protocols. Thus, the potential need for revolutionary changes cause concerns by the major market players, whose input is of crucial importance to leverage the progress in this area.

Keeping this issue in mind, in this paper, we aimed to provide a high-level evaluation of the system, where the principal architecture and the most of the logic are directly taken from the recent commercial CPUs, while the THz wireless network is implemented “as it is”, without any further integration or adaptation. Further, simplest candidate solutions for physical and medium access control layers have been assumed.

Using a comprehensive hybrid methodology combining direct measurements and cycle-accurate CPU traffic simulations along with ray-tracing of the intra-CPU propagation environment, we performed a first-order performance evaluation of the described system, showing that even under such pessimistic assumptions about the propagation conditions, modulation scheme, and traffic load dynamics, THz wireless network can effectively support the traffic from up to 100 active cores.

Many more parameters and characteristics, such as materials and energy constraints, have to be taken into account in more accurate performance evaluation. At the same time, we believe that our study provides a clear and illustrative example of THz band advantages for intra-chip communications, while the proposed evolutionary approach and preliminary performance insights can boost further investigations in both THz communications and massive multi-core CPUs design.

REFERENCES

- [1] G. Chrysos, “Intel Xeon Phi coprocessor—The architecture,” Intel, Santa Clara, CA, USA, White Paper 176, 2014.
- [2] S. Mittal, J. S. Vetter, and D. Li, “A survey of architectural approaches for managing embedded DRAM and non-volatile on-chip caches,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 26, no. 6, pp. 1524–1537, Jun. 2015.
- [3] W. R. Davis et al., “Demystifying 3D ICs: The pros and cons of going vertical,” *IEEE Des. Test Comput.*, vol. 22, no. 6, pp. 498–510, Nov./Dec. 2014.
- [4] D. Matolak, A. Kodi, S. Kaya, D. Di Tomaso, 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.
- [5] V. Petrov, A. Pyattaev, D. Moltchanov, and Y. Koucheryavy, “Terahertz band communications: Applications, research challenges, and standardization activities,” in *Proc. 8th Int. Congr. Ultra Modern Telecommun. Control Syst. Workshops (ICUMT)*, Oct. 2016, pp. 183–190.
- [6] L. H. K. Duong et al., “A case study of signal-to-noise ratio in ring-based optical networks-on-chip,” *IEEE Des. Test*, vol. 31, no. 5, pp. 55–65, Oct. 2014.
- [7] L. Zhang, M. Yang, Y. Jiang, and E. Regentova, “Architectures and routing schemes for optical network-on-chips,” *Comput. Elect. Eng.*, vol. 35, no. 6, pp. 856–877, Nov. 2009.
- [8] M. F. Chang, V. P. Roychowdhury, L. Zhang, H. Shin, and Y. Qian, “RF/wireless interconnect for inter- and intra-chip communications,” *Proc. IEEE*, vol. 89, no. 4, pp. 456–466, Apr. 2001.
- [9] M. F. Chang et al., “CMP network-on-chip overlaid with multi-band RF-interconnect,” in *Proc. IEEE 14th Int. Symp. High Perform. Comput. Archit. (IEEE HPCA)*, Feb. 2008, pp. 191–202.
- [10] S. Abadal, E. Alarcón, A. Cabellos-Aparicio, M. C. Lemme, and M. Nemirovsky, “Graphene-enabled wireless communication for massive multicore architectures,” *IEEE Commun. Mag.*, vol. 51, no. 11, pp. 137–143, Nov. 2013.
- [11] S. Abadal et al., “Broadcast-enabled massive multicore architectures: A wireless RF approach,” *IEEE Micro*, vol. 35, no. 5, pp. 52–61, Sep./Oct. 2015.
- [12] S. Abadal et al., “Scalability of broadcast performance in wireless network-on-chip,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 27, no. 12, pp. 3631–3645, Dec. 2016.
- [13] S. Abadal, R. Martínez, J. Solé-Pareta, E. Alarcón, and A. Cabellos-Aparicio, “Characterization and modeling of multicast communication in cache-coherent manycore processors,” *Comput. Elect. Eng.*, vol. 51, pp. 168–183, Apr. 2016.
- [14] S. Abadal, M. Iannazzo, M. Nemirovsky, A. Cabellos-Aparicio, H. Lee, and E. Alarcón, “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.
- [15] V. Petrov, M. Komarov, D. Moltchanov, J. M. Jornet, and Y. Koucheryavy, “Interference and sinr in millimeter wave and terahertz communication systems with blocking and directional antennas,” *IEEE Trans. Wireless Commun.*, vol. 16, no. 3, pp. 1791–1808, Mar. 2017.
- [16] P. Boronin, V. Petrov, D. Moltchanov, Y. Koucheryavy, and J. M. Jornet, “Capacity and throughput analysis of nanoscale machine communication through transparency windows in the terahertz band,” *Nano Commun. Netw.*, vol. 5, pp. 72–82, Sep. 2014.
- [17] Y.-S. Lee, *Principles of Terahertz Science and Technology*, vol. 170. Berlin, Germany: Springer, 2009.
- [18] D. Saeedkia, Ed., *Handbook of Terahertz Technology for Imaging, Sensing and Communications*. Amsterdam, The Netherlands: Elsevier, 2013.
- [19] R. Han and E. Afshari, “Filling the terahertz gap with sand: High-power terahertz radiators in silicon,” in *Proc. IEEE Bipolar/BiCMOS Circuits Technol. Meeting (BCTM)*, Oct. 2015, pp. 172–177.
- [20] E. Seok et al., “Progress and challenges towards terahertz CMOS integrated circuits,” *IEEE J. Solid-State Circuits*, vol. 45, no. 8, pp. 1554–1564, Aug. 2010.
- [21] Q. J. Gu et al., “CMOS THz generator with frequency selective negative resistance tank,” *IEEE Trans. THz Sci. Technol.*, vol. 2, no. 2, pp. 193–202, Mar. 2012.
- [22] J.-D. Park, S. Kang, S. V. Thyagarajan, E. Alon, and A. M. Niknejad, “A 260 GHz fully integrated CMOS transceiver for wireless chip-to-chip communication,” in *Proc. Symp. VLSI Circuits (VLSIC)*, Jun. 2012, pp. 48–49.
- [23] S. Moghadami, F. Hajilou, P. Agrawal, and S. Ardalani, “A 210 GHz fully-integrated OOK transceiver for short-range wireless chip-to-chip communication in 40 nm CMOS technology,” *IEEE Trans. THz Sci. Technol.*, vol. 5, no. 5, pp. 737–741, Sep. 2015.
- [24] D. Shim et al., “Paths to terahertz CMOS integrated circuits,” in *Proc. IEEE Custom Integr. Circuits Conf. (IEEE CICC)*, Sep. 2009, pp. 657–664.
- [25] U. R. Pfeiffer, “Silicon CMOS/SiGe transceiver circuits for THz applications,” in *Proc. 12th IEEE Topical Meeting Silicon Monolithic Integr. Circuits RF Syst. (IEEE SiRF)*, Jan. 2012, pp. 159–162.
- [26] R. Han and E. Afshari, “A 260 GHz broadband source with 1.1 mW continuous-wave radiated power and EIRP of 15.7 dBm in 65 nm CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 138–139.
- [27] Y. Touse and E. Afshari, “A scalable THz 2D phased array with +17 dBm of EIRP at 338 GHz in 65 nm bulk CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2014, pp. 258–259.
- [28] P. Rodríguez-Vázquez, N. Sarmah, K. Aufinger, and U. R. Pfeiffer, “An OOK-modulator at 240 GHz with 20 GHz bandwidth,” in *Proc. German Microw. Conf. (GeMic)*, Mar. 2016, pp. 345–348.
- [29] J.-D. Park, S. Kang, and A. M. Niknejad, “A 0.38 THz fully integrated transceiver utilizing a quadrature push-push harmonic circuitry in SiGe BiCMOS,” *IEEE J. Solid-State Circuits*, vol. 47, no. 10, pp. 2344–2354, Oct. 2012.
- [30] I. Kallfass et al., “All active MMIC-based wireless communication at 220 GHz,” *IEEE Trans. THz Sci. Technol.*, vol. 1, no. 2, pp. 477–487, Nov. 2011.
- [31] J. D. Albrecht, T.-H. Chang, A. S. Kane, and M. J. Rosker, “DARPA’s nitride electronic next generation technology program,” in *Proc. IEEE Compound Semiconductor Integr. Circuit Symp. (IEEE CSICS)*, Oct. 2010, pp. 1–4.

- [32] S. Masuda et al., "GaN MMIC amplifiers for W-band transceivers," in *Proc. Eur. Microw. Integr. Circuits Conf. (EuMIC)*, Sep. 2009, pp. 443–446.
- [33] M. Micovic et al., "W-band GaN MMIC with 842 mW output power at 88 GHz," in *IEEE MTT-S Int. Microw. Symp. Dig.*, May 2010, p. 1.
- [34] M. Micovic et al., "GaN MMIC PAs for E-band (71 GHz–95 GHz) radio," in *Proc. IEEE Compound Semiconductor Integr. Circuits Symp.*, Oct. 2008, pp. 1–4.
- [35] L. A. Samoska, "An overview of solid-state integrated circuit amplifiers in the submillimeter-wave and THz regime," *IEEE Trans. THz Sci. Technol.*, vol. 1, no. 1, pp. 9–24, Aug. 2011.
- [36] V. Radisic, K. M. K. H. Leong, X. Mei, S. Sarkozy, W. Yoshida, and W. R. Deal, "Power amplification at 0.65 THz using InP HEMTs," *IEEE Trans. Microw. Theory Techn.*, vol. 60, no. 3, pp. 724–729, Mar. 2012.
- [37] S. Rakheja and P. Sengupta. (2015). "Gate-voltage tunability of plasmons in single and multi-layer graphene structures: Analytical description and concepts for terahertz devices." [Online]. Available: <https://arxiv.org/abs/1508.00658>
- [38] J. A. Del Alamo, "Nanometre-scale electronics with III–V compound semiconductors," *Nature*, vol. 479, no. 7373, pp. 317–323, 2011.
- [39] J. G. Fiorenza et al., "(Invited) aspect ratio trapping: A unique technology for integrating Ge and III–Vs with silicon CMOS," *ECS Trans.*, vol. 33, no. 6, pp. 963–976, 2010.
- [40] H. Ko et al., "Ultrathin compound semiconductor on insulator layers for high-performance nanoscale transistors," *Nature*, vol. 468, no. 7321, pp. 286–289, 2010.
- [41] M. Hudait et al., "Heterogeneous integration of enhancement mode In_{0.7}Ga_{0.3}As quantum well transistor on silicon substrate using thin ($\leq 2 \mu\text{m}$) composite buffer architecture for high-speed and low-voltage (0.5 V) logic applications," in *IEDM Tech. Dig.*, Dec. 2007, pp. 625–628.
- [42] P.-Y. Chen and A. Alú, "All-graphene terahertz analog nanodevices and nanocircuits," in *Proc. 7th Eur. Conf. Antennas Propag. (EuCAP)*, Apr. 2013, pp. 697–698.
- [43] S. Rakheja, H. Wang, T. Palacios, I. Meric, K. L. Shepard, and D. A. Antoniadis, "A unified charge-current compact model for ambipolar operation in quasi-ballistic graphene transistors: Experimental verification and circuit-analysis demonstration," in *IEDM Tech. Dig.*, Dec. 2013, pp. 5.5.1–5.5.4.
- [44] K. N. Parrish and D. Akinwande, "Even-odd symmetry and the conversion efficiency of ideal and practical graphene transistor frequency multipliers," *Appl. Phys. Lett.*, vol. 99, no. 22, p. 223512, 2011.
- [45] L. Vicarelli, M. S. Vitiello, D. Coquillat, A. Lombardo, A. C. Ferrari, W. Knap, M. Polini, V. Pellegrini, and A. Tredicucci, "Graphene field-effect transistors as room-temperature terahertz detectors," *Nature Mater.*, vol. 11, no. 10, pp. 865–871, Sep. 2012.
- [46] Y. Kawano, "Wide-band frequency-tunable terahertz and infrared detection with graphene," *Nanotechnology*, vol. 24, no. 21, p. 214004, 2013.
- [47] K. S. Novoselov, V. I. Fal'ko, L. Colombo, P. R. Gellert, M. G. Schwab, and K. Kim, "A roadmap for graphene," *Nature*, vol. 490, no. 7419, pp. 192–200, 2012.
- [48] C. Han, A. O. Bicen, and I. F. Akyildiz, "Multi-ray channel modeling and wideband characterization for wireless communications in the terahertz band," *IEEE Trans. Wireless Commun.*, vol. 14, no. 5, pp. 2402–2412, May 2015.
- [49] S. Priebe, M. Kannicht, M. Jacob, and T. Kürner, "Ultra broadband indoor channel measurements and calibrated ray tracing propagation modeling at THz frequencies," *J. Commun. Netw.*, vol. 15, no. 6, pp. 547–558, Dec. 2013.
- [50] A. S. Glassner, Ed., *An Introduction to Ray Tracing*. Amsterdam, The Netherlands: Elsevier, 1989.
- [51] B. Sklar, *Digital Communications*, vol. 2. Upper Saddle River, NJ, USA: Prentice-Hall, 2001.
- [52] V. Petrov, D. Moltchanov, and Y. Koucheryavy, "Interference and SINR in dense terahertz networks," in *Proc. IEEE 82nd Veh. Technol. Conf. (VTC Fall)*, Sep. 2015, pp. 1–5.
- [53] V. Petrov, D. Moltchanov, and Y. Koucheryavy, "On the efficiency of spatial channel reuse in ultra-dense THz networks," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2015, pp. 1–7.
- [54] V. Petrov, M. Komarov, D. Moltchanov, J. M. Jornet, and Y. Koucheryavy, "Interference analysis of EHF/THF communications systems with blocking and directional antennas," in *Proc. IEEE Global Commun. Conf. (IEEE GLOBECOM)*, Dec. 2016, pp. 1–7.
- [55] V. Petrov, J. Kokkonen, D. Moltchanov, J. Lehtomäki, and Y. Koucheryavy, "Enabling simultaneous cooling and data transmission in the terahertz band for board-to-board communications," *Phys. Commun.*, vol. 22, pp. 9–18, Mar. 2017.
- [56] Intel PCM, *Specifications*, accessed on Jul. 6, 2015. [Online]. Available: <https://software.intel.com/enus/articles/intel-performance-counter-monitor>
- [57] GNU Lic. *MARSSx86—Micro-Architectural and System Simulator for x86-Based Systems Software Package*, accessed on Feb. 9, 2015. [Online]. Available: <http://marss86.org/~marss86/>
- [58] N. Binkert et al., "The gem5 simulator," *ACM SIGARCH Comput. Archit. News*, vol. 39, no. 2, pp. 1–7, 2011.
- [59] D. Sanchez and C. Kozyrakis, "ZSim: Fast and accurate microarchitectural simulation of thousand-core systems," in *Proc. ACM ISCA*, 2013, pp. 475–486.
- [60] *The Structural Simulation Kit (SST)*, accessed on Apr. 18, 2017. [Online]. Available: <http://sst-simulator.org/>
- [61] D. Moltchanov et al., "Statistical traffic properties and model inference for shared cache interface in multi-core CPUs," *IEEE Access*, vol. 4, pp. 4829–4839, Aug. 2016.
- [62] K. Yotov, K. Pingali, and P. Stodghill, "Automatic measurement of memory hierarchy parameters," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 33, no. 1, pp. 181–192, 2005.
- [63] D. Molka, D. Hackenberg, R. Schone, and M. S. Muller, "Memory performance and cache coherency effects on an Intel Nehalem multiprocessor system," in *Proc. 18th Int. Conf. Parallel Archit. Compilation Techn. (PACT)*, Sep. 2009, pp. 261–270.
- [64] V. Babka and P. Tüma, "Investigating cache parameters of x86 family processors," in *Proc. SPEC Benchmark Workshop Comput. Perform. Eval. Benchmarking*, Jan. 2009, pp. 77–96.
- [65] Z. Fang et al., "Measuring microarchitectural details of multi- and many-core memory systems through microbenchmarking," *ACM Trans. Archit. Code Optim.*, vol. 11, no. 4, p. 55, Jan. 2015.
- [66] R. Rom and M. Sidi, *Multiple Access Protocols: Performance and Analysis*. Springer, 2012.



VITALY PETROV received the M.Sc. degree in information systems security from the Saint Petersburg State University of Aerospace Instrumentation, St. Petersburg, Russia, in 2011, and the M.Sc. degree in communications engineering from the Tampere University of Technology, Tampere, Finland, in 2014, where he is currently pursuing the Ph.D. degree with a focus on enabling terahertz band communications for Beyond 5G wireless networks. He was a Visiting Scholar with the Georgia Institute of Technology, Atlanta, USA, in 2014, and a Strategic Intern with the Nokia Research Center, Helsinki, Finland, in 2012. His current research interests are in Internet-of-Things, terahertz band communications, nanonetworks, cryptology, and network security.



DMITRI MOLTCHANOV received the M.Sc. and Cand.Sc. degrees from the Saint-Petersburg State University of Telecommunications, Russia, in 2000 and 2002, respectively, and the Ph.D. degree from the Tampere University of Technology in 2006. He is currently a Senior Research Scientist with the Department of Electronics and Communications Engineering, Tampere University of Technology, Finland. He authored over 50 publications. His research interests include performance evaluation and optimization issues of wired and wireless IP networks, Internet traffic dynamics, quality of user experience of real-time applications, and traffic localization P2P networks. He serves as a TPC Member in a number of international conferences.



MARIA KOMAR received the B.Sc. and M.Sc. degrees from PG Demidov Yaroslavl State University in 2013 and 2015, respectively. She is currently pursuing the M.Sc. degree with the Department of Electronics and Communications Engineering, Tampere University of Technology, Tampere, Finland. Her research interests are in wireless communications, embedded systems, networks-on-chips, and cryptography.



ALEXANDER ANTONOV received the B.Sc. and M.Sc. degrees from ITMO University, Russia, in 2013 and 2015, respectively, where he is currently pursuing the Ph.D. degree with the Computer Science Department. He has authored several publications indexed in Scopus and Web of Science, and has taken part in a number of international conferences. His research interests include methods and technologies of embedded systems and systems-on-chip design, computer-aided design,

and FPGA-based design.



PAVEL KUSTAREV received the M.Sc. and Cand.Sc. degrees from ITMO University, Russia, in 1995 and 1999, respectively. He was the Head of the over ten research and development projects. He is currently an Associate Professor with the Computer Science Department and the Head of the Program Engineering and Computers Systems Faculty, Saint Petersburg National Research University of Information Technologies, Mechanics and Optics, ITMO University. He has authored over 25 publications. He is also a Supervisor for graduate and postgraduate students. His research interests include architectural design and analysis of systems-on-a-chip, networked embedded systems, and cyber-physical systems-based on wired and wireless networks.



SHALOO RAKHEJA received the B.Tech. degree in electrical engineering from IIT Kanpur, Kanpur, India, in 2005, and the M.S. and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, in 2009 and 2012, respectively. She is currently an Assistant Professor with the Department of Electrical and Computer Engineering, New York University. Her research interests are in modeling and simulating electron transport in nanoscale transistors.

She is also interested in novel, beyond-CMOS devices using alternate state variables, energy harvesting for sensor networks and other mobile devices, and flexible and transparent electronics.



YEVGENI KOUCHERYAVY received the Ph.D. degree from the Tampere University of Technology, Finland, in 2004. He is currently a Professor and the Laboratory Director with the Department of Electronics and Communications Engineering, Tampere University of Technology. He has authored numerous publications in the field of advanced wired and wireless networking and communications. His current research interests include various aspects in heterogeneous wireless communication networks and systems, the Internet of Things and its standardization, and nanocommunications. He is an Associate Technical Editor of the *IEEE Communications Magazine* and an Editor of the *IEEE COMMUNICATIONS SURVEYS AND TUTORIALS*.

...