Graphene-enabled Wireless Communication for Massive Multicore Architectures

Sergi Abadal*, Eduard Alarcón*, Max C. Lemme[†‡], Mario Nemirovsky[§] and Albert Cabellos-Aparicio*

*NaNoNetworking Center in Catalonia (N3Cat)
Universitat Politècnica de Catalunya, Barcelona, Spain
Email: {abadal,acabello}@ac.upc.edu, eduard.alarcon@upc.edu
[†]School of Information and Communication Technology
KTH Royal Institute of Technology, Kista, Sweden
Email: lemme@kth.se
[‡]University of Siegen, Hölderlinstr. 3, 57076 Siegen, Germany
Email: max.lemme@uni-siegen.de
[§]ICREA Senior Research Professor at
Barcelona Supercomputing Center (BSC), Barcelona, Spain
Email: mario.nemirovsky@bsc.es

Abstract—Current trends in microprocessor architecture design are leading towards a dramatic increase of core-level parallelization, wherein a given number of independent processors or cores are interconnected. Since the main bottleneck is foreseen to migrate from computation to communication, efficient and scalable means of inter-core communication are crucial for guaranteeing steady performance improvements in many-core processors. As the number of cores grows, it remains unclear whether initial proposals, such as the Network-on-Chip (NoC) paradigm, will meet the stringent requirements of this scenario. This position paper presents a new research area where massive multicore architectures have wireless communication capabilities at the core level. This goal is feasible by using graphene-based planar antennas, which can radiate signals at the Terahertz band while utilizing lower chip area than its metallic counterparts. To the best of our knowledge, this is the first work that discusses the utilization of graphene-enabled wireless communication for massive multicore processors. Such wireless systems enable broadcasting, multicasting, all-to-all communication, as well as significantly reduce many of the issues present in massively multicore environments, such as data coherency, consistency, synchronization and communication problems. Several open research challenges are pointed out related to implementation, communications and multicore architectures, which pave the way for future research in this multidisciplinary area.

Index Terms—Network-on-Chip, Multicore Processor, Wireless, Graphene, Antennas, Terahertz, Architecture, Manycore, Emerging

I. INTRODUCTION

Throughout the decades, technology advancements in digital circuits, i.e., precise manufacturing techniques, have enabled a steady reduction in the size of transistors. Such tendency has allowed the integration of more transistors on the same chip and resulted in a very high performance increase and cost decrease per transistor. As the level of integration approaches Ultra-Large-Scale Integration (ULSI), the intra-chip communication latency and power consumption become major barriers that prevent the continuation of the trend set by the Moore's Law.

Indeed, the main reasons for the diminishing performance returns of such downscaling trend are as follows. By reducing the width of the on-chip wires, their resistance and therefore the resistive-capacitive (RC) delay are significantly increased. Also, by taking into account the increased clock frequencies imposing reduced symbol times, the charging and discharging the wire within the allotted time becomes a very challenging problem. Finally, the dynamic power demand of a Complementary Metal Oxide Semiconductor (CMOS) transistor grows proportionally to its operation frequency and quadratically to the circuit voltage, justifying the need for low-voltage and frequency-limited designs. Graphene, thanks to its extremely promising properties, could enable the devising of transistors with higher speed and lower energy consumption than traditional CMOS devices. However, such transistors are, thus far, projected for its application in RF circuits rather than in digital computation, due to the intrinsic absence of band gap in graphene [1].

Since better performance is no longer achievable through an increase in clock frequency due to the reasons pointed out above, the natural trend in microprocessor architecture design is to improve the performance by means of parallel architectures. Parallelization is achieved by interconnecting several independent processors forming a Chip Multiprocessor (CMP), and has led towards the recent emergence of multicore and manycore, i.e., more than 16 cores, processors. The main performance bottleneck in these systems is currently defined by the intra-chip communication requirements set by coherency or synchronization, among other common and necessary operations in multicore environments. In this context, the Network-on-Chip (NoC) paradigm was proposed to increase the performance of CMP systems by providing scalable and efficient inter-core communication through wireline routed interconnections. This approach arose as opposed to the traditional bus-based architectures, which scale poorly in terms of delay and energy efficiency due to its time division

multiplexing nature, when the number of cores is increased.

However, as the technology downscaling allows the integration of more cores in the same chip, initial wired NoC solutions pose several challenges in terms of delay, power requirements and chip area utilization, also referred to as area overhead. Consequently and since the inception of the NoC paradigm, many research efforts have been directed towards coping with the increasing demands of the ever-changing field of microprocessor architecture [2]. In this paper, we review the state-of-the-art of this research area and introduce Graphene-enabled Wireless Network-on-Chip (GWNoC), a novel approach that relies in graphene-based nano-antennas (see [3]) to implement wireless communication between the cores of a multiprocessor. Such proposal is expected to considerably outperform other initial Wireless Network-on-Chip (WNoC) designs and aims to provide improved scalability, flexibility and area overhead for multicore systems with hundreds or even thousands of cores. To the best of our knowledge, this is the first work that discusses the employment of nanoscale wireless communication for on-chip networks. Our main contributions are:

- We review the state-of-the-art of the field by briefly surveying different research activities, such as the application of photonics or wireless solutions to traditional NoCs, identifying their potential advantages as well as their open issues.
- We also propose the use of graphene in the WNoC context due to its excellent properties. The resulting GWNoC will provide core-level communication in the THz band by means of graphene-based nano-antennas, achieving a superior performance in terms of bandwidth and area overhead.
- We outline and analyze the open issues and research challenges regarding the combination of graphene and wireless on-chip networks, serving as a roadmap for future investigations.

The remainder of this paper is organized as follows. In Section II, we detail the different emerging solutions that have been proposed to address the shortcomings of traditional NoCs. In Section III, we present GWNoC and analyze its unique features. In Section IV, we analyze the different research challenges of GWNoC from both implementation and network perspectives. Finally, conclusions are drawn in Section V.

II. State-of-the-art of On-chip Networking

Multicore processor architectures rely on complex memory systems in order to provide fast and efficient means for data access and sharing between cores. The characteristics of such systems define the communication demands that the NoC must satisfy. Indeed, memory consistency, i.e., how the memory operations are ordered in execution, and cache coherence operations, i.e., maintaining a single memory image accessible to all processors, are especially critical in terms of latency. As the number of cores on a chip increases, traditional wireline topologies become insufficient for guaranteeing such latency

conditions without significantly affecting other performance metrics.

In this context, disruptive solutions are required in order to alleviate the limitations of NoC in terms of latency while providing high bandwidth and maintaining affordable power and area overheads. Several approaches have been investigated and will be briefly introduced in the following subsections.

A. 3D Network-on-Chip

The creation of tridimensional integrated circuits, wherein layers of active devices are vertically interconnected, has shown to imply significant benefits such as improved noise immunity or higher packing density, as well as several advantages related to the NoC design [4]. For instance, the average wire propagation delay is considerably reduced due to the short distance between layers, i.e., tens of micrometers. Moreover, 3D stacking enables the use of topologies which would be unfeasible in the 2D design space, potentially yielding reduced multihop latency results [4]. Since such advantages are mainly at the network level, the potential improvements are compatible with and practically independent of the underlying interconnection technology. Moreover, it is an effective way to intuitively interface different technologies in hybrid approaches, facilitating modularity by avoiding the integration of different technologies in the same layer.

It is also important to note that 3D stacking presents considerable challenges. The superposition of active layers produces an increase in the heat density that must be circumvented in order to avoid thermal effects. Also, refined techniques are needed for the manufacture of such tridimensional integrated circuits and networks, in particular, alignment methodologies for the precise positioning of the vertical interconnects.

B. RF Interconnects

As the technology downscaling further narrows the performance bottleneck imposed by the delay and bandwidth figures of traditional wireline on-chip networks, research efforts are recently focused on finding scalable alternatives in the physical layer design. In this regard, transmitting modulated RF signals over on-chip transmission lines is proposed in order to complement traditional wireline schemes [5]. Since signals propagate at nearly the speed of light, the propagation delay is significantly reduced and becomes independent of the link length. Moreover, the overall available data rate can be increased by simultaneously transmitting several frequency or code-multiplexed signals through a shared transmission line (see Figure 1). Such approach allows the interconnection of multiple cores using the same transmission line by assigning each core a channel and thus reducing the number of on-chip wires. The bandwidth could be dynamically allocated depending on the needs of each core.

However, the utilization of RF interconnects entails several open challenges. The circuital implementation of frequency or code multiplexing transceivers produces both an area and power overhead that must be controlled as the size of the network increases. Also, the physical topology must be carefully

Fig. 1. Example of a multichannel communication scheme. Depending on the interconnection paradigm, the cores will communicate through a RF transmission line or an optical waveguide, with its respective convention (frequency or wavelength multiplexing)


Fig. 2. Schematic diagram of a hybrid Wireless Network-on-Chip, wherein wireless links connect distant core clusters on top of a wireline mesh. Wireline inter-router links are omitted for simplicity.

designed as impedance mismatch reflections at the terminations of the transmission line may generate interferences.

C. Photonic Network-on-Chip

The advent of nanoscale silicon photonics has enabled the creation of *photonic NoCs* by means of the integration of CMOS-compatible optical building blocks. Such networks maintain the main advantages of RF interconnects in terms of delay and reconfigurability, while yielding extremely high bandwidth and low power. In fact, the power consumption of a photonic interconnect is almost independent of both the transmission bitrate and distance [6]. However, the difficulty of implementing all-optical buffering or header processing turns the design of a photonic NoC into a challenging task. Several approaches have been proposed in order to circumvent such feature, like setting up optical virtual channels by means of a parallel electric NoC control channel [6], or completely avoiding routing by means of shared waveguides and employing one wavelength per core [7].

D. Wireless Network-on-Chip

While technology downscaling is identified as the main culprit of rendering traditional NoCs insufficient, it might also be part of the solution. Indeed, advancements in CMOS integration have enabled the implementation of millimetric on-chip antennas able to radiate in the GHz band, as well as of adequate high frequency transceivers. The Wireless Network-on-Chip (WNoC) paradigm, in which on-chip antennas enable the creation of inter-core wireless links, aims to take advantage of such option.

The advantages of employing wireless communication for intra-chip networks are mainly threefold. First and likewise to the RF/photonic options, the *propagation delay* is significantly reduced and is practically independent of the transmission distance at the chip-level, as messages propagate at nearly the speed of light. Second, the inherent flexibility of wireless communication adds *reconfigurability* options to the design process. Since no path infrastructure is needed to convey information to the receiver, WNoCs can modify the logical topology or other transmission parameters without the need of any physical modification. Finally and due to the aforementioned advantages, inter-core wireless communication offers potential

for improved *scalability* in terms of latency, throughput and energy consumption, as shown in [8], [9].

Such benefits have favored the proposal of several preliminary and diverse WNoC designs. In principle, WNoC is usually regarded as a valid option when it is deployed in order to enhance the wired principal data NoC. In this hybrid approach, the unique advantages of wireless communication over existing wired solutions can be used to create a wireless control system [10]. Also, the insertion of wireless long-range point-to-point links is proposed in order to significantly decrease the average hop count of traditional NoC topologies, either performing fixed and regular positioning ([8], explained in Figure 2) or following the principles of small-world networks [9].

III. GRAPHENE-ENABLED WIRELESS NETWORK-ON-CHIP

The prospects of the WNoC paradigm are certainly promising. However, current implementations show several shortcomings related to the main enabler of WNoC: the on-chip antennas. For instance, the complete replacement of wired links for their wireless counterparts is proposed in [11] for the data plane, but it remains unclear how the area overhead figures will scale as the number of cores increases.

Indeed, the size of future metallic on-chip antennas, i.e., hundreds of micrometers [8], might render unfeasible the approach of integrating at least one antenna per core, as the core sizes continue to shrink with each CMOS technology generation and reach sizes of a few hundreds of micrometers. Also, since the available bandwidth is generally inversely proportional to the antenna size, metallic antennas may not be able to provide enough bandwidth in such a data intensive scenario. Such issues cannot be solved by further reducing the size of a metallic antenna, as this would impose the use of very high resonant frequencies, from the near infrared to the optical ranges. Due to the low mobility of electrons in metals when nanometer scale structures are considered, and the challenges in implementing a transceiver which will be able to operate at this extremely high frequency, the feasibility of inter-core wireless communications would be compromised if this approach would be followed.

Alternatively, we propose the employment of nanoscale wireless communication by means of graphene-based nano-antennas in order to unleash the full potential of the WNoC

Fig. 3. Schematic diagram of a 144-core Graphene-enabled Wireless Network-on-Chip (GWNoC) with one antenna, of width W and length L, per core.

paradigm. Graphene-based nano-antennas just a few micrometers in size, i.e. two orders of magnitude below the dimensions of future metallic on-chip antennas, could provide inter-core communication in the Terahertz (0.1-10 THz) band. These characteristics will both enable size compatibility with each processor core and offer enough bandwidth in massively parallel multiprocessors [12]. We refer to this new nanoscale wireless communication technique as *Graphene-enabled Wireless Network-on-Chip* (GWNoC) and we propose it as the basis of future on-chip network architectures.

Figure 3 shows a simple conceptual implementation of a GWNoC. It is important to note that all the processor cores are equipped with a graphene-based nano-antenna and a nano-transceiver, the latter of which prepares the information for outgoing transmissions and demodulates incoming transmissions. While maintaining the advantages of WNoC, the main benefits of the GWNoC approach are as follows.

Bandwidth and Area Limitations

Graphene-based nano-antennas support the propagation of tightly confined Surface Plasmon Polariton (SPP) waves. Due to their high effective mode index, the propagation speed of SPP waves can be up to two orders of magnitude below the electromagnetic wave propagation speed in vacuum [3]. In other words, graphene-based nano-antennas are expected to be two orders of magnitude smaller than metallic quarter wave antennas for the same resonant frequency. According to our preliminary results [3], a few micrometer wide and long nano-patch antenna could effectively radiate in the Terahertz band (see Figure 4). On the one hand, such reduced dimensions are comparable with future core sizes, i.e. few hundreds of micrometers, enabling the integration of one or multiple antennas per core and giving birth to the concept of *wireless core* in future generation multiprocessors. On the other hand, the Terahertz band may offer enough bandwidth to accommodate the exponentially increasing requirements of multiprocessors.

Additionally, the employment of the Terahertz signals virtually eliminates the near-field effects between neighboring antennas since the near-field region at these frequencies is around one hundred micrometers.

Multicast and All-to-All communication

In multicore environments, a critical part of the on-chip traffic are short, generally multicast, control messages employed for cache coherence, data consistency or synchronization purposes. Moreover, some parallel applications require the transmission of large amounts of data in an all-to-all manner, e.g., 3D FFT calculation [13]. The majority of NoC solutions do not perform well under such conditions due to the difficulty of implementing efficient all-to-all communication in general, and multicast or broadcast schemes in particular. Conversely, since the information is radiated and can be potentially received by any receiver in the transmission range, GWNoC provides native broadcasting and multicasting capabilities, and makes data transmission transparent with respect to the location of data within the chip. To the best of our knowledge, no other interconnect technology inherently offers such options, which ensure the scalability of traditional multicore architectures and open a vast design space at the architectural level. Efficiently transmitting simultaneously from multiple sources to multiple destinations heavily alleviates the constraints of parallel architecture design, therefore reducing the complexity of parallel programming.

Modularity

Replicability and modularity are two properties that can be achieved with the creation of the *wireless core*, in which a graphene nano-antenna and a transceiver are integrated within a core processor. A library of general-purpose or specific wireless cores could be created, allowing the building of custom multicore processors by the integration and initial configuration of a set of such predesigned cores. The cores of such multiprocessors will be wirelessly interconnected, potentially eliminating the necessity of a specific wireline on-chip network. Moreover, the modularity of a GWNoC also enables communication between circuit layers in a 3D approach without including additional wiring.

IV. OPEN ISSUES AND RESEARCH CHALLENGES

The adoption of graphene as the basis of a new generation of WNoC brings up a wide variety of research challenges, covering areas from physical implementation up to computer architecture design. Such challenges need to be addressed in order to pave the way for the development of a new breed of multiprocessors.

Antenna Characterization

The further characterization of the graphene nano-antenna is central in order to assess the feasibility of the GWNoC approach and poses a grand challenge by itself. One important aspect to evaluate is the radiation efficiency of these antennas. Recent works report a radiation efficiency of up to 25%

Fig. 4. Dependence of the first resonant frequency of a nano-patch antenna as a function of its width, determined by the normalized absorption cross section. The antenna length is $L = 5$ μm and the plots correspond, from left to right, to 1 μm, 2 μm, 5 μm, 10 μm and infinitely wide patches.

[14], which is reduced to a total efficiency of 4% when the impedance mismatch between the antenna and a hypothetical load is considered. However, the development of impedance matching circuits for graphene nano-antennas may help improve this low efficiency. The impact of other factors, such as the chemical doping of graphene, in the antenna performance also deserves further attention.

Another issue that remains unclear is how the silicon substrate or the circuits surrounding the antenna could affect the propagation of SPP waves inside the antenna and, therefore, its radiation properties. The work in [3] reports a shift of the resonant frequency in the presence of a silicon substrate of different permittivity and thickness values. However, how the substrate and neighboring circuits will affect other parameters such as the radiation efficiency still need to be investigated.

Implementation

In order to enable on-chip wireless communication, it is necessary to develop circuits to drive the nano-antenna. These circuits need to operate at the same frequency as the antenna itself, which is the Terahertz band according to our preliminary results [3]. The progress in the development of graphene-based components shows that they are excellent candidates for ultra-high-frequency applications [1], aiming at eventually obtaining graphene-based circuits operating in the Terahertz band. For instance, impressive cut-off frequencies (f_T) of 350 GHz have been obtained in Graphene Field-Effect Transistors (GFETs) due to the high carrier mobility in the nanomaterial [1]. Graphene is also uniquely suited for Low Noise Amplifiers (LNAs) as it theoretically offers high frequency and low noise.

Precise, efficient and replicable production of graphene nanoribbons (GNR) or bilayer GFETs is essential in order to provide high quality building blocks for the creation of the nano-antenna and the transceiver. The dimensions of GNRs required for such RF components are quite relaxed compared to GNRs for logic devices [15] and hence it seems quite feasible to achieve high-yield manufacturing of graphene nano-ribbons (GNRs).

Another important issue to investigate is how to integrate the antenna with the transceiver. The fundamental challenge here is to identify suitable heterogeneous integration techniques that make possible the integration of graphene into a semiconductor circuit environment. In general, graphene device technology is compatible with silicon technology [16]. However, graphene-dielectric interfaces and metal-graphene contacts need to be optimized as they limit the overall RF performance by reducing the carrier mobility [16] and introducing parasitic resistances [17]. Ultimately, another option may be to co-integrate other technologies such as InP devices to be able to drive graphene antennas at the required fequencies [18].

Communication and Networking

Being the radiation frequency known, a channel model is fundamental in order to evaluate the available on-chip communication bandwidth. Such models must analyze the presence of *molecular absorption*, which is due to the internal excitation of certain molecules at high frequencies. While preliminary results point out the possibility to utilize very large transmission bandwidths for short-range Terahertz communications [12], the multipath effects of within-package reflections are not taken into account.

The study of the propagation and reflection of the EM waves inside the chip package is also necessary due to the following. Generally, the radiation efficiency of planar antennas is extremely low in the coplanar direction and substantially higher in the transversal direction. Line-of-sight communication becomes very challenging, while communication through reflected EM waves might be feasible since the distance between the antenna and the chip package is in the millimeter range. This and other peculiarities of the on-chip wireless scenario, e.g., the constraints in terms of energy and chip area and energy, require us to rethink the entire protocol stack with respect to classical wireless networks, including the following aspects:

- Special coding and modulation schemes need to be investigated in order to achieve efficient communication through reflected waves while avoiding, or even taking advantage, of multipath propagation.
- Addressing strategies for GWNoC must be carefully decided, as multicast and broadcast are envisaged to be essential for next generation multicore architectures.
- The flexibility offered by wireless communication allows to significantly reduce the multihop latency, at the expense of adding interference and medium access contention due to the existence of simultaneous transmissions in the same medium. A Medium Access Control (MAC) protocol could be designed taking into account the design tradeoff between contention delay and multihop latency when the transmission range of the antennas is adaptively modified. The MAC protocol should also consider possible high-contention phases present in parallel programs.

Multicore Architecture

Given the specifications of the physical implementation and protocols which enable inter-core wireless communication and, on its turn, broadcast and multicast capabilities, a radically new paradigm in multicore architecture can be envisaged. Such all-to-all communication capabilities open a vast range of possibilities in terms of architecture design. For instance, at the memory level, cache coherency protocols could be modified and in a few cases even eliminated due to the possibility of implementing all-to-all on-chip communication. In this regard, we aim to explore the new multicore architecture design space opened by GWNoC in our future work.

V. Conclusions

In this position paper, we present the vision of a novel and multidisciplinary research area in which nanoscale techniques enable core-level wireless communication for massive multicore processors. Although the concept of wireless on-chip networks is discussed in the literature, the size of the proposed on-chip antennas prevents such paradigm from exploiting its potential. We propose the employment of graphene-based nano-antennas for the design and development of flexible and scalable wireless on-chip networks. The high bandwidth, inherent multicast and broadcast capabilities and extremely low area overhead offered by this novel approach could deliver a major breakthrough in massive multicore architectures. Several implementation, networking and architecture research challenges need to be addressed in order to achieve such ambitious goals, as pointed out in the last section of this work.

Acknowledgment

References

[1] Y. Wu, K. A. Jenkins, A. Valdes-garcia, D. B. Farmer, Y. Zhu, A. A. Bol, C. Dimitrakopoulos, W. Zhu, F. Xia, P. Avouris, and Y.-m. Lin, "State-of-the-Art Graphene High-Frequency Electronics," *Nano Letters*, vol. 12, no. 6, pp. 3062–3067, 2012.

[2] T. Bjerregaard and S. Mahadevan, "A survey of research and practices of Network-on-chip," *ACM Computing Surveys*, vol. 38, no. 1, pp. 1–51, Jun. 2006.

[3] I. Llatser, C. Kremers, A. Cabellos-Aparicio, J. M. Jornet, E. Alarcón, and D. N. Chigrin, "Graphene-based nano-patch antenna for terahertz radiation," *Photonics and Nanostructures - Fundamentals and Applications*, vol. 10, no. 4, pp. 353–358, 2012.

[4] B. S. Feero and P. P. Pande, "Networks-on-Chip in a Three-Dimensional Environment: A Performance Evaluation," *IEEE Transactions on Computers*, vol. 58, no. 1, pp. 32–45, Jan. 2009.

[5] E. Socher and M.-C. F. Chang, "Can RF Help CMOS Processors?" *IEEE Communications Magazine*, vol. 45, no. 8, pp. 104–111, Aug. 2007.

[6] A. Shacham, K. Bergman, and L. P. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Transactions on Computers*, vol. 57, no. 9, pp. 1246–1260, Sep. 2008.

[7] G. Kurian, J. Miller, J. Psota, J. Eastep, J. Liu, J. Michel, L. Kimerling, and A. Agarwal, "ATAC: A 1000-Core Cache-Coherent Processor with On-Chip Optical Network," in *Proceedings of the 19th international conference on Parallel architectures and compilation techniques*. ACM, 2010, pp. 477–488.

[8] S.-B. Lee, S.-W. Tam, I. Pefkianakis, S. Lu, M.-C. F. Chang, C. Guo, G. Reinman, C. Peng, M. Naik, L. Zhang, and J. Cong, "A scalable micro wireless interconnect structure for CMPs," in *Proceedings of the 15th annual international conference on Mobile computing and networking - MobiCom '09*. New York, New York, USA: ACM Press, 2009, p. 217.

[9] A. Ganguly, K. Chang, S. Deb, P. P. Pande, B. Belzer, and C. Teuscher, "Scalable Hybrid Wireless Network-on-Chip Architectures for Multi-Core Systems," *IEEE Transactions on Computers*, vol. 60, no. 10, pp. 1485–1502, 2010.

[10] P. Y. Chiang, S. Woracheewan, C. Hu, L. Guo, H. Liu, R. Khanna, and J. Nejedlo, "Short-Range, Wireless Interconnect within a Computing Chassis: Design Challenges," *IEEE Design & Test of Computers*, vol. 27, no. 4, pp. 32–43, Jul. 2010.

[11] D. Zhao, Y. Wang, J. Li, and T. Kikkawa, "Design of multi-channel wireless NoC to improve on-chip communication capacity," in *Fifth IEEE/ACM International Symposium on Networks on Chip (NoCS)*. IEEE, 2011, pp. 177–184.

[12] J. M. Jornet and I. F. Akyildiz, "Channel Modeling and Capacity Analysis for Electromagnetic Wireless Nanonetworks in the Terahertz Band," *IEEE Transactions on Wireless Communications*, vol. 10, no. 10, pp. 3211–3221, 2011.

[13] H. Jagode and J. Hein, "Custom assignment of MPI ranks for parallel multi-dimensional FFTs: Evaluation of BG/P versus BG/L," in *International Symposium on Parallel and Distributed Processing with Applications (ISPA)*. IEEE, 2008, pp. 271–283.

[14] M. Tamagnone, J. S. Gomez-Diaz, J. R. Mosig, and J. Perruisseau-Carrier, "Analysis and design of terahertz antennas based on plasmonic resonant graphene sheets," *Journal of Applied Physics*, vol. 112, p. 114915, 2012.

[15] S. Das and J. Appenzeller, "On the Importance of Bandgap Formation in Graphene for Analog Device Applications," *IEEE Transactions on Nanotechnology*, vol. 10, no. 5, pp. 1093–1098, 2011.

[16] M. C. Lemme, T. J. Echtermeyer, M. Baus, and H. Kurz, "A Graphene Field Effect Device," *IEEE Electron Device Letters*, vol. 28, no. 4, pp. 282–284, 2007.

[17] S. Vaziri, M. Ostling, and M. C. Lemme, "A Hysteresis-Free High-k Dielectric and Contact Resistance Considerations for Graphene Field Effect Transistors," *ECS Transactions*, vol. 41, pp. 165–171, 2011.

[18] M. Kim, J.-S. Rieh, and S. Jeon, "Recent Progress in Terahertz Monolithic Integrated Circuits," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2012, pp. 746–749.

Sergi Abadal received the BSc and MSc degree in Telecommunication Engineering from Universitat Politcnica de Catalunya (UPC), Barcelona, Spain, in 2010 and 2011, respectively. From September 2009 to May 2010, he was visiting researcher at the Broadband Wireless Networking Lab, Georgia Institute of Technology, Atlanta. Since 2011, he is pursuing his PhD at the NaNoNetworking Center in Catalunya (N3Cat, http://www.n3cat.upc.edu) at UPC. His current research interests are graphene-based wireless and nanophotonic communications for on-chip networks.

Eduard Alarcón (S'96, M'01), received M. Sc. (national award) and Ph.D. degrees in EE from UPC Barcelona, Spain, in 1995 and 2000, respectively, where he became Associate Professor in 2001, and has been visiting Professor at University of Colorado at Boulder, USA (2003) and KTH Stockholm (2011). He has coauthored more than 250 scientific publications, 4 book chapters and 4 patents, and has been involved in different national, EU and US R&D projects. Research interests include the areas of on-chip energy management circuits, energy harvesting and wireless energy transfer, and nanocommunications. He was elected IEEE CAS society distinguished lecturer, elected member of the IEEE CAS Board of Governors (2010-2013), recipient of Best paper award at IEEEMWSCAS98, co-editor of 4 journals special issues, 5 conference special sessions, TPC co-chair and TPC member of 15 IEEE conferences, and Associate Editor for IEEE TCAS-I, TCAS-II, JETCAS, JOLPE and Nano Communication Networks.

Albert Cabellos-Aparicio received a BSc (2001), MSc (2005) and PhD (2008) degree in Computer Science Engineering from the Technical University of Catalonia (www.upc.edu). In 2004 he was awarded with a full scholarship to carry out PhD studies at the Department of Computer Architecture, Technical University of Catalonia (UPC), Spain. In september 2005 he became an assistant professor of the Computer Architecture Department and as a researcher in the Broadband Communications Group (http://cba.upc.edu/). In 2010 he joined the NaNoNetworking Center in Catalunya (http://www.n3cat.upc.edu) where he is the Scientific Director. He is an editor of the Elsevier Journal on Nano Computer Network and member of the Project Management Committee of the LISPmob opensource initiative (http://lispmob.org). His main research interests are future architectures for the Internet and Nano-networks.

Max C. Lemme received the Dipl.-Ing. (MSc) and Dr.-Ing. (PhD) degree in Electrical Engineering from RWTH Aachen University in Germany. He is currently Heisenberg-Professor for Graphene-based Nanotechnology at the University of Siegen, Germany and Guest-Professor at KTH Royal Institute of Technology, Sweden since 2010. Before joining KTH, he was a research fellow at Harvard University from 2008-2010 and worked for 10 years at nanotechnology start-up AMO GmbH, Germany, as Head of Technology Department. His research interests include non-conventional nano-CMOS devices, novel high-k materials for gate stacks and, since 2005, graphene technology, devices and circuits. He received a NanoFutur Award by the German Federal Ministry for Education and Research (BMBF) in 2006, a Humboldt-Fellowship in 2007 and an ERC Starting Grant in 2012. He is a Senior Member of the IEEE.

Mario Nemirovsky received the Telecommunications Engineering degree from the National University of La Plata, Argentina, in 1980, and his PhD in Electrical and Computer Engineering from the University of California, Santa Barbara, in 1990, where he was an adjunct professor from 1991 to 1998. After being chief architect in companies such as Apple Inc., National Semiconductors or General Motors (GM), he founded several renowned start-ups including FlowStorm Networks, Xstream Logic, ConSentry Networks or Miraveo. In 2007, he became an ICREA Senior Research Professor at the Barcelona Supercomputing Center (BSC). Mario holds more than 60 issued patents: he pioneered the concepts of Massively Multithreading (MMT) processing for the high performance processor and the by now well-established Simultaneous Multithreding architecture (SMT). He also architected the GM engine control being used in all GM cars for over 20 years. His current research interests include multithreaded multicore systems, high performance systems, network processors and Big Data.