

Received March 20, 2017, accepted April 1, 2017, date of publication April 12, 2017, date of current version May 17, 2017.

Digital Object Identifier 10.1109/ACCESS.2017.2693267

Computing and Communications for the Software-Defined Metamaterial Paradigm: A Context Analysis

SERGI ABADAL¹, CHRISTOS LIASKOS², AGELIKI TSIOLIARIDOU², SOTIRIS IOANNIDIS²,
ANDREAS PITSILLIDES³, JOSEP SOLÉ-PARETA¹, EDUARD ALARCÓN¹,
AND ALBERT CABELLOS-APARICIO¹

¹NaNoNetworking Center in Catalunya, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain

²Institute of Computer Science, Foundation of Research and Technology-Hellas (FORTH), GR-700 13 Heraklion, Greece

³Department of Computer Science, University of Cyprus, Nicosia 2109, Cyprus

Corresponding author: Sergi Abadal (abadal@ac.upc.edu)

This work was supported by the European Union via the Horizon 2020: Future Emerging Topics call (FET Open) through the Project VISORSURF under Grant EU736876.

ABSTRACT Metamaterials are artificial structures that have recently enabled the realization of novel electromagnetic components with engineered and even unnatural functionalities. Existing metamaterials are specifically designed for a single application working under preset conditions (e.g., electromagnetic cloaking for a fixed angle of incidence) and cannot be reused. Software-defined metamaterials (SDMs) are a much sought-after paradigm shift, exhibiting electromagnetic properties that can be reconfigured at runtime using a set of software primitives. To enable this new technology, SDMs require the integration of a network of controllers within the structure of the metamaterial, where each controller interacts locally and communicates globally to obtain the programmed behavior. The design approach for such controllers and the interconnection network, however, remains unclear due to the unique combination of constraints and requirements of the scenario. To bridge this gap, this paper aims to provide a context analysis from the computation and communication perspectives. Then, analogies are drawn between the SDM scenario and other applications both at the micro and nano scales, identifying possible candidates for the implementation of the controllers and the intra-SDM network. Finally, the main challenges of SDMs related to computing and communications are outlined.

INDEX TERMS Metamaterials, software-defined metamaterials, manycores, approximate computing, network-on-chip, nanonetworks.

I. INTRODUCTION

Metamaterials have recently enabled the realization of a wealth of novel electromagnetic (EM) and optical components with engineered functionalities [1]. These include EM invisibility of objects (cloaking), total radiation absorption, filtering and steering of light and sound, as well as ultra-efficient, miniaturized antennas for sensors and implantable communication devices [2], [3]. These applications are possible due to the unnatural physical properties of the metamaterials, which stem from their unique structure generally composed of a pattern of conductive material repeated over a 3D volume. If the pattern is replicated over a 2D surface, we obtain a metasurface instead [4], [5].

Despite its outstanding properties, the adoption of metamaterials and metasurfaces is currently limited due to their non-adaptivity and non-reusability. These properties restrict their applicability to a single functionality per structure (e.g. steering light towards a fixed direction) and to static structures only. Moreover, designing a metamaterial remains a task for specialized researchers, limiting their accessibility from the broad engineering field.

Achieving reconfigurability in metamaterials has been a topic under intense research over the past decade [6]. On the one hand, since the metamaterial properties mostly depend on its conductive pattern, first proposals tried to modulate it using tunable devices or mechanical parts [7]. On the other




FIGURE 1. Schematic representation of a Software-Defined Metamaterial (SDM). External devices drive a network of controllers, whose local decisions determine the global behavior of the metamaterial.

hand, more advanced techniques rely on the use of phase-change media, graphene, or liquid crystals [8]. The main downturn of these techniques is that the reconfigurability boils down to the tunability of a given static property as there is no actual programmatic control over the functionality. Thus, the accessibility issues are not solved either.

Recently, Liaskos *et al.* proposed the concept of Software-Defined Metamaterials (SDMs), a hardware platform that can host metamaterial functionalities described in software [9]. The main idea is to integrate a network of miniaturized controllers within the metamaterial structure. The controllers receive programmatic directives and perform simple alterations on the metasurface structure, adjusting its EM behavior globally, locally, upon request or depending on the environment. In the specific example of Figure 1, the controllers activate or deactivate their associated switch to determine the metamaterial pattern. The required functionality is described in well-defined, reusable software modules, which are disseminated to the controllers from an external interface also shown in Figure 1. This has several advantages. First, the SDM can host multiple functionalities concurrently and adaptively. Second, the SDM can be connected to external devices or even other SDMs to better adapt to the surroundings or increase the operative range. Last but not least, the SDM concept reduces the knowledge required to design a metamaterial for a given purpose.

As mentioned earlier, a network of controllers lies at the heart of an SDM. Both the controllers and their interconnections would ideally be simple, ultra-efficient, yet powerful enough to enable real-time adaptivity and support multiple ways of interacting locally, globally, and with external entities. However, this combination of constraints and requirements poses important challenges, thus requiring a careful definition of the computation and communication mechanisms that will drive the operation of SDMs.

This position paper aims to provide a context analysis of the SDM paradigm from the computing and communication perspectives. We build on the observation that existing approaches may be amenable to this new application if

adapted properly. As the main contribution, this work does not aim to deliver a working solution, but rather:

- To provide a broad analysis of the application context, detailing its particularities regarding the physical implementation, workload characteristics and performance requirements.
- To present an overview of existing computing and networking approaches that could be amenable to SDMs.
- To enumerate the outstanding challenges of this new research area, paving the way for future investigations.

The remainder of this paper is organized as follows. Section II provides background on the reconfigurable metamaterial paradigm and analyzes its main particularities. Then, Section III debates the applicability of current computing techniques to the SDM scenario. Sections IV and V extend the discussion to the networking domain in general and the Network-on-Chip (NoC) paradigm in particular. Finally, Section VI lays out the main computation and communication challenges of SDMs and Section VII concludes the paper.

II. SOFTWARE-DEFINED METASURFACES

For simplicity, let us focus on a particular 2D metasurface case shown in Figure 1. In this case, the dimensions of the rectangular Split Ring Resonators (SRRs) define the refraction angle of an impinging EM wave. Each controller is associated to a switch (or a set of switches) that can be set on conductive or resistive state, therefore shaping the SRRs used as building blocks. Changes of state in each switch can be prescribed via the metasurface interface either because the user desires to change the refraction angle or because external sensing devices detect changes in the EM source. The scale of the controllers and the switches defines the granularity of the formable patterns, eventually determining the number of possible configurations and the frequency at which the metasurface can operate. We refer to the interested reader to [9] for more details.

A. GENERAL STRUCTURE OF AN SDM

The particular example of Figure 1 represents one of the different potential approaches that can be used to attain reconfigurability in an SDM. Other schemes may involve the use of tunable resistors or capacitors, the value of which determines the behavior of the SDM and is dictated by the controller. With the use of graphene, which is inherently tunable, an SDM can be created by allowing controllers to change the electrostatic bias applied to the different areas of the graphene sheet. In any case, and regardless of its physical characteristics, a generic instance of an SDM would have the logical structure shown in Figure 2, with the following set of planes:

- **Metamaterial Plane:** which delivers the desired EM behavior through a reconfigurable pattern. The metamaterial plane can be implemented, for instance, with CMOS switches as illustrated in Figure 1 or materials such as graphene, which can be tuned by simply changing an electrostatic bias [8].




FIGURE 2. Sketch of the logical structure of an SDM. It may include actuators (A), sensors (S), controllers (C) and routers (R).

- **Sensing and Actuation Plane:** which modifies the behavior of the metamaterial plane. Successive SDM generations may integrate sensors within the metasurface, so that state changes can be determined internally without the need to reach an external controller, thereby providing a truly autonomous and adaptive operation.
- **Shielding Plane:** which attempts to decouple the EM behavior of the top and bottom planes, aiming to avoid mutual interferences. A simple metallic layer could be used to this end, as metals mainly reflect EM waves.
- **Computing Plane:** which executes external commands from the interface and internal commands from the rest of controllers or sensors to effectively change the EM profile of the metamaterial plane. Note that one controller can drive the operation of one or several actuators. Possible design approaches are discussed in Section III.
- **Communications Plane:** which coordinates the actions of the computing plane and keeps in touch with external entities via the SDM interface. It may be wired or wireless. Possible design approaches are discussed in Sections IV and V.

At this point, it is important to stress that the programmability of SDMs refers to their EM properties only. This differentiates SDMs from the Claytronics project, which aims to program changes in the physical shape of matter [10]. In any case, we will later see that advances in that application context can be meaningful to the SDM paradigm as they share some basic traits.

B. CURRENT PERSPECTIVES AND VISION

The potential of the SDM concept is vast given the plethora of potential applications in the microwave range and above. However, their feasibility is currently restricted to the development of proof-of-concept devices maintaining a simple architecture and moderate performance. As shown in the left part of Figure 3, those initial SDMs would be limited reactive systems in the microwave range with external sensing and power supply. In the longer term, the SDM vision could incorporate new components such as embedded nanosensors, a full integrated network, or an energy harvesting system, and exploit smaller and faster controllers to create devices capable of reacting microwave or terahertz signals in a truly autonomous manner, without having to rely on the constant intervention of an external controller.

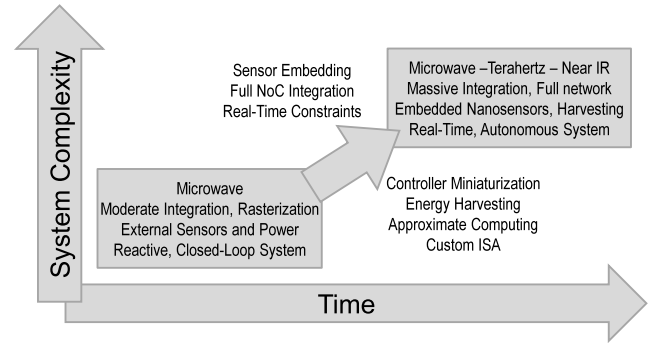


FIGURE 3. Current perspectives and expected evolution of the SDM research activities.

C. CONTEXT ANALYSIS

In the following, the main characteristics of the SDM application are analyzed considering both the current state of things and the full potential of the SDM vision. The main insights are summarized in Table 1.

1) PHYSICAL LANDSCAPE

Computing and communications occur within a constrained environment. The lateral dimensions of the metamaterial building blocks are generally $\lambda/4$ or less, where λ is the wavelength of the EM waves impinging on the metamaterial. This, for the example of Fig. 1, means that a reasonable target of $f = 6$ GHz would require the deployment of an SRR every ~ 1 cm. Assuming that each SRR is composed by dozens of switches, controllers would be placed every ~ 1 mm approximately. Note that such density requirements can be relaxed if concentration is applied, i.e., each controller is shared by a few switches. It is also worth noting that the controllers will operate at a frequency generally much lower than that of the manipulated EM waves.

Since the granularity of the actuation scales inversely to frequency, SDMs in the microwave range and above will account for a considerably dense and highly integrated network of as-small-as-possible controllers. Due to this density and to minimize heat and potential interferences, both the controllers and the network should have a strict power budget also related to the frequency of the impinging EM waves. Link energy figures in NoC, currently in the pJ/bit range and below, can serve as a first reference. In future systems where the SDM is meant to be autonomous and powered by the same EM source than that the controlled by the metamaterial, the energy budget should comply with the limitations of the energy harvester.

The computing and communications devices will be laid out in a planar environment, probably in a chip-like configuration, if we consider the metasurface case; whereas this may not be necessarily so in the broader sense of the SDM paradigm. In both cases, however, the topology of actuators reconfiguring the pattern will be static, controlled, and known beforehand (most likely fairly periodic). As we will see, this offers important optimization opportunities.

TABLE 1. Communications and Computing in the SDM scenario.

Perspective	Traits	Analogous Techniques
Physical Landscape	Planar, dense, constrained, static, controlled (spatially periodic)	Embedded manycores, GPGPUs, NoCs, WNoCs, Nanonetworks, Energy harvesting
Workload Characteristics	Light, highly correlated (dissemination, reductions, sensing), monolithic nature	Multicore processors, NoCs, WNoCs, Sensor networks
Application Requirements	Depends on granularity, from latency-insensitive to real-time, tolerant to errors	Mission-critical systems, Nanonetworking, Approximate computing

2) WORKLOAD CHARACTERISTICS

Although the SDM paradigm opens the door to a large wealth of possibilities at the metamaterial plane, the computing and communication planes only need to perform three distinct actions, summarized in Figure 4:

- 1) Receive and execute external directives. This basically implies the dissemination of data from the interface to all the controllers and the execution of (preferably state-independent) instructions for the initial configuration of the metasurface and the subsequent function updates. After receiving feedback from external sensors or the metasurface itself, the interface may also need to convey messages containing parameter adjustments required to maintain the desired behavior.
- 2) Process and send internal information to the interface. For debugging or SDM interconnectivity purposes, controllers may need to individually or collectively communicate with the interface, therefore generating a *reduction* operation with temporally correlated many-to-one traffic. In the debugging case, the metasurface will send periodic state reports or sporadic failure notifications. In the interconnectivity case, the interface will receive control signals from the different metasurfaces in order to coordinate their joint operation.
- 3) Coordinate their execution strictly within the SDM. To maintain the correct behavior of the SDM, integrated sensors may need to communicate with the controllers and drive their execution. These events generate point-to-point or multicast communication with potentially high spatial correlation. Controllers may also need to locally notify errors and perform flow control within the network.

On top of these considerations, it is important to note that the communication and computation intensity will eventually depend on the desired spatial and temporal granularity, as well as on the variability of the EM waves impinging on the SDM. In any case, given the nature of the application and of the energy constraints of the controllers, the load should be moderate.

Another interesting point is that the SDM will be a monolithic system, meaning that designers will have control over the entire architecture, from the physical implementation up to the compilers. This may have little impact on the computing side since multiprocessors are generally monolithic as well. However, it represents a big departure from traditional networks where the nodes, protocols, and applications are developed by different teams. This implies that protocols can be streamlined by entering into the design loop of the whole architecture as in NoCs.

3) APPLICATION REQUIREMENTS

The requirements set by the application mostly depend on the desired spatiotemporal granularity. In the first SDM generations, where the main objective is to attain reconfigurability via software, latency requirements are expected to be relaxed, probably between a few milliseconds and a few seconds. In a longer term, where SDM applications may demand fast adaptivity, stronger timing requirements on the order of microseconds may be imposed to the controllers and the network. Designs will favor simplicity against performance in the former case, while real-time constraints will suggest the use of mission-critical solutions in the latter case.

An interesting feature stemming from the fundamentals of the SDM application concerns the reliability requirements.

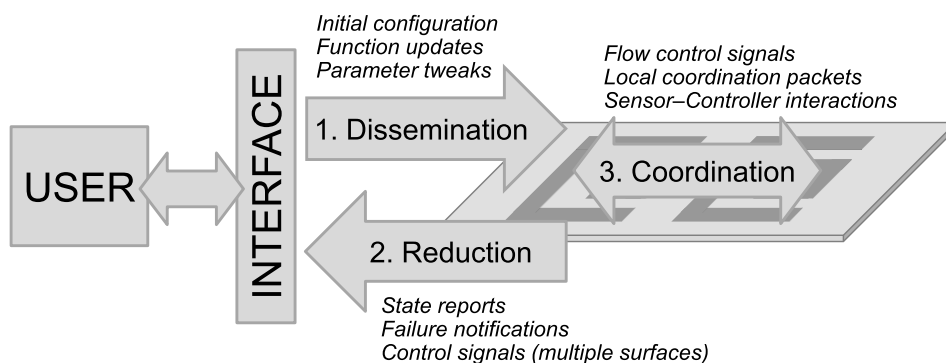


FIGURE 4. Computation and communication flows in SDMs.

Depending on the particular design of the metamaterial pattern, the task of the controller may be, for instance, the choice of a discrete set of voltage levels. The failure of a few controllers or the choice of an incorrect voltage level may not be noticed at the macroscopic level, still obtaining the desired EM behavior. This situation can be quantified and used to improve the efficiency of the controllers and the network.

III. APPLICABILITY OF CURRENT COMPUTING TRENDS

The analysis of the SDM context has clarified that the computing plane will be massive, composed by a potentially huge amount of tiny controllers deployed within a single monolithic system. As a result, simplicity will most likely drive the development of controllers and lead to custom solutions. Each controller will have to handle commands from external entities or from internal controllers or switches, to compute the new state of its associated switches or actuators. This operation is required to obtain the desired feature (e.g. a given pattern, impedance, bias) in the pathway to obtain the target macroscopic EM behavior.

In the following, we revisit how the SDM community can benefit from existing knowledge in the area of computing. The main conclusions are summarized in the left part of Figure 5.

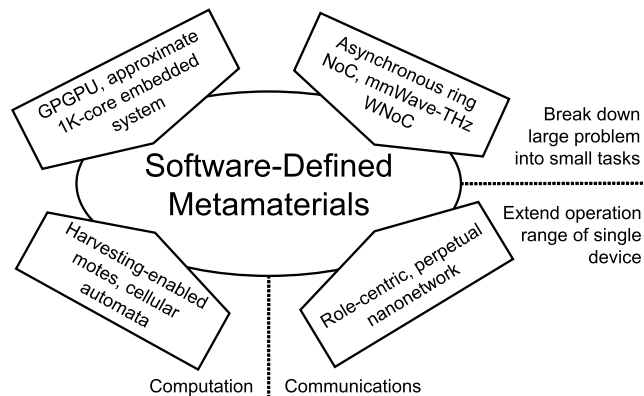


FIGURE 5. The SDM scenario seen from the perspective of possible computation and communications solutions.

A. MASSIVE MANYCORES

Taking into consideration their density and *a priori* monolithic nature, the network of controllers within SDMs can be seen, at large, as a massive manycore processor. Such processors already exist in the research domain, reaching the thousand-core count within a single chip not only in theoretical discussions [11], but also built and demonstrated with CMOS technology [12]. Strictly speaking, however, an SDM does not include a general multiprocessor, but rather an embedded manycore as it can be described as a *computing system with a dedicated function within a larger mechanical or electrical system, often with real-time constraints*. Development of SDM controllers can therefore inherit experience of past custom architectures or software with

real-time constraints for embedded multicores, especially considering that they are already used in other software-defined paradigms [13].

In their work, Liaskos *et al.* [9] discuss the suitability of massively parallel computing architectures mostly due to their node density and the fact that all controllers perform a small set of identical functions. General-Purpose Graphic Processing Units (GPGPUs) such as CUDA-enabled video cards are mentioned as they can handle thousands of threads, conveniently organized in sets and executing simple operations [14]. The possible use of GPGPUs-like computing organizations, at least for proof-of-concept explorations, may be backed up by the vast amount of applied research and knowledge gained through the widespread adoption of these devices in the scientific domain.

B. TOWARDS INFINITESIMAL COMPUTING

The top-down view of an SDM implicitly assumes that a large task is divided into multiple and possibly identical subtasks to reach a common goal. This matches well with the process of reconfiguring the SDM via the software interface. More prospectively, if we envision allowing SDMs to internally sense and adapt to different EM conditions, a bottom-up perspective might be more appropriate.

In strict terms, the controllers and the associated integrated sensors (if any) form a sensor and actuator network [15]. One controller is not significant by itself as it can only impact on one or a few building blocks of the metamaterial, and therefore needs to be connected to other controllers to obtain a desired macroscopic behavior.

Regarding node density and size limitations of controllers, SDMs are conceptually close to paradigms such as smart dust [16], Claytronics [10], or Wireless NanoSensor Network (WNSN) [17]. The potentially infinitesimal motes or nanorobots forming these networks account for tiny computing capabilities and may need energy harvesting modules to operate. Thus, existing knowledge on how to develop and program these systems, e.g. using an event-centric approach, may be highly relevant to the SDM community [18].

Finally, it is worth noting that the periodic layout and simplicity requirements of SDMs allows us to draw a very strong analogy to the cellular automata approach [19]. Cellular automata can achieve very complex emergent behaviors by simply using a few simple rules and communication with the immediate neighbours, therefore becoming an interesting candidate for the implementation of controllers.

C. APPROXIMATE COMPUTING

Approximate and probabilistic computing have been recently proposed to increase energy-efficiency in fields where inexact results are tolerable [20]. As discussed in Section II-F, SDMs may fall into this category depending on the actual implementation of the metasurface pattern. This opens the door to a reduction of the voltage applied to the controller or the use of circuits providing approximate results in exchange for lower power. As long as the error probability remains

bounded along the execution of the controller routines, this approach can reduce power consumption without noticeably degrading the performance of the SDM.

The metamaterial community can leverage existing knowledge in these areas, which have been applied across the computing stack: building approximate circuits, bounding the error probability throughout execution, debugging approximate devices, or combining the approach with energy harvesting, to name a few examples [21]–[24].

IV. APPLICABILITY OF ON-CHIP COMMUNICATION TECHNIQUES

The system-level resemblance between multiprocessors and reconfigurable metamaterials suggest that on-chip communication techniques may be a valid approach for SDM. As such, we next review a set of NoC methods that could be applicable here. We make a distinction between wireline and wireless designs as it remains unclear which option is preferable *a priori*: the wireless option avoids the use of conductive wiring which may interfere with the metamaterial plane, but comes at the expense of a higher complexity, i.e. the design and integration of tiny antennas and transceivers.

A. NETWORK-ON-CHIP

The NoC paradigm essentially refers to packet-switched networks of integrated routers and links. In broad terms, research in this field has been mostly directed to scale designs while obtaining high performance and reasonable efficiency. For high performance, objectives have been to minimize and bound latency in Chip Multiprocessors (CMPs) [25], [26], as well as to make better use of bandwidth in GPGPUs [27]. The main issue is that these proposals generally require fairly complex routers and wide links to implement their improvements and meet manycore requirements. Thus, they are not directly portable to the SDM scenario.

SDMs are much less sensitive to latency than CMPs, which automatically turns proposals seeking simplicity and low power into much better NoC-based candidates for our target scenario. Next, we review several of these techniques.

1) ASYNCHRONOUS NoC

By default, most NoC designs are synchronous. This requires the distribution of a clock signal throughout the chip, which takes precious area and power. To avoid it, one can adopt the Globally Asynchronous Locally Synchronous (GALS) approach consisting in the use of asynchronous links to communicate synchronous cores [28]. In a synchronous controller design, an interface is required to connect with the clockless network; whereas in an asynchronous or event-based approach, no further adaptation will be required.

2) TOPOLOGY AND ROUTER MICROARCHITECTURE

As in CMPs, a bidimensional mesh seems a natural fit for SDMs due to its ease of layout and performance. Yet still, even simpler topologies such as a ring [29] are an intelligent choice since they allow the use of minimalistic router

microarchitectures. In particular, the proposal by Kim [30] eliminates the need for both costly buffers to avoid losses and virtual channels to guarantee deadlock-freedom. Another interesting point to consider here is whether clustering, i.e. serving groups of controllers via the same router, can help reduce the footprint.

3) APPROXIMATE COMMUNICATION

The main idea behind approximate computing has been also applied to NoCs. Li *et al.* [31] proposed to use a lightweight lossy network to carry messages in program sections tolerant to errors. Another approach would be to drop the supply voltage close to near-threshold levels, even if that results into occasional bit flips.

B. WIRELESS NETWORK-ON-CHIP

The Wireless Network-on-Chip (WNoC) paradigm consists in the integration of antennas and transceiver circuits close to the computing cores, introducing higher flexibility at the network level [32], [33]. Driven by the latency sensitivity and moderate throughput of CMPs, WNoCs are designed seeking high data rates and reasonable area. To this end, most proposals employ simple modulations such as On-Off Keying (OOK) and frequencies in the millimeter-Wave (mmWave) range to obtain high bandwidth.

Again, the stringent constraints of SDM suggest to sacrifice performance to reduce footprint. Since communication in SDMs is expected to be occasional and much less latency-sensitive than in NoCs, one can reduce the available bandwidth. This relaxes the requirements cast upon the antenna and transceiver and therefore enables the use of more compact circuits. Another technique that could be leveraged to reduce the footprint would be that of approximate computing: the main idea would be to reduce the gain of the power amplifier to save power even if that increases the bit error rate, as long as this error probability remains bounded within a safe margin. The use of electrically small antennas is another example of this footprint–performance tradeoff.

Although works assuming a large density of antennas within the same chip have been published [34], [35], WNoCs generally complement a wireline NoC and do not need many antennas to achieve meaningful results. The case for SDM, however, is fundamentally different as the objective is to minimize wiring. This will probably require pushing the frequency used for communication up and beyond the mmWave bands for two reasons: (1) to avoid coupling and interferences with the metamaterial plane, and (2) to achieve the target network density and efficiency, as both area and power scale inversely to frequency in on-chip environments (see Fig. 6).

The use of graphene-based antennas in the terahertz band can be a valid option for this particular purpose due to their outstanding properties [36]. The use of graphene as resonant sheets has been widely investigated, showing that patch or dipole antennas a few micrometers long and wide resonate in the terahertz band (0.1–10 THz), this is, between




FIGURE 6. Area and energy scalability of chip-scale wireless transceivers (data and fitting from [41]).

one and two orders of magnitude lower than their metallic counterparts [37]–[40]. Additionally, the unique tunability properties given by the relation between voltage bias and resonant frequency open the door to a set of new opportunistic communication protocols.

V. APPLICABILITY OF NANONETWORKING TECHNIQUES

As discussed in Section III-B, the SDM paradigm representative similarities with sensor and actuator networks. This suggests that ad hoc communication and networking mechanisms generally employed in such networks may be a candidate for the implementation of SDMs. In fact, the expected node density and huge physical constraints of intra-SDM networks lead to considering extreme ad hoc solutions, which mostly lie in the nanonetworking domain [42].

Striving to maintain complexity at a minimum, most nanonetworking research finds consensus on the use of simplistic modulations such as the Time-Spread On-Off Keying (TS-OOK) [43]. In TS-OOK, a logical 0 (1) is represented by means of a silence (short pulse), respectively, with a relatively long time between transmissions. This simplifies the receiver and reduces the probability of collisions. Moreover, this approach can be opportunistically combined with low weight coding [44] and rate division multiple access [45] to maximize its efficiency.

Energy harvesting is another pillar of nanonetworking as it may enable the concept of perpetual networks. Its impact on the design of the protocol stack of nanonetworks has been under intense research over the last years, covering aspects such as the energy consumption policy [46] or the Medium Access Control (MAC) protocol [47] and assessing the potential network performance of perpetual networks [48]. The metamaterial community could benefit from these contributions since an important milestone is to make SDMs reconfigurable without compromising their autonomy. In particular, the work by Cid-Fuentes *et al.* [49], which explores the design of energy harvesting systems in scenarios with high spatiotemporal traffic correlation, would be directly applicable

to SDMs given the high expected correlation of traffic and potential harvesting sources in SDMs.

On top of all this, Liaskos *et al.* [9] provided a view of the main networking challenges of SDM and preliminary potential solutions from the nanonetworking point of view. The authors first discuss the problem of addressing in such dense networks and how it can be simplified taking into consideration the periodic, controlled, and monolithic nature of the system. As in NoCs, nodes can be unambiguously and statically identified with an internal id, leading to a major simplification of routing protocols [50] and a simplification or even complete elimination of addressing in particular case scenarios [51]. Finally, the authors propose the use of role-centric networking techniques, this is, defining custom roles in substitution of the conventional layered approach [52]. Preliminary evaluations were made for the data dissemination case (from interface to controllers), achieving a similar performance and an energy efficiency three times higher than with a generic protocol stack.

VI. OPEN ISSUES AND RESEARCH CHALLENGES

The SDM design and optimization process poses new challenges for the various planes that comprise it. Envisioned milestones are detailed in the next subsections.

A. WIRELESS CHANNEL CHARACTERIZATION

The communications plane constitutes the heart of the SDM. The network of controllers is responsible for receiving external commands and finally altering the SDM structure to meet a given objective. To this end, the efficiency of this network is critical: highly lossy communications may translate to redundant retransmissions of programmatic commands, resulting into higher SDM setup times and reduced adaptivity potential. Thus, understanding and modeling the controllers' communication channel is critical for optimizing their communication accordingly.

The SDM communications plane exhibits some unique attributes that affect the channel modeling. Specifically, the placement of the controllers is expected to exhibit a periodic layout, which is known to yield a well-defined chirality in the communication channel [53]. Additionally, the efficiency of the shielding plane is not a given, and may be subject to metamaterial plane restrictions. For instance, the presence of a highly conductive shielding layer underneath the metamaterial plane may result into a strong and unwanted reflection coefficient. Thus, a non-perfect shielding plane must be taken into account when studying the channel model, factoring for the interference from the metamaterial plane. This cross-talk can yield a highly non-linear channel, given that the programmatic commands exchanged by the controllers alter the metamaterial plane, in turn affecting the interference to the wireless channel. Note that most of these impairments are present in the physically similar WNoC environment, for which comprehensive propagation models have not been developed yet [54], [55].

B. ABSTRACTING THE PHYSICS

SDMs are intended to be usable by non-physicists, which constitutes an attractive and challenging trait. In essence, an SDM user should be able to define the required, high-level SDM functionality without having to specify the low-level actions required to obtain it. Moreover, a user should be able to combine and multiplex SDM functionalities, creating novel SDM *applications*. To these ends, the following SDM software components need to be implemented:

- An SDM compiler, responsible for translating basic SDM functionalities to the corresponding patterns that should be formed over its surface. These basic functionalities are those offered by metasurfaces in general, i.e., EM absorption, steering, polarization, non-linear response [5]. The compiler essentially defines the low-level actions required to form these patterns, such as the state of switches in Fig. 1.
- An SDM standard software library, offering the tools for monitoring, debugging, multiplexing, and abstracting the basic SDM functionalities towards higher-level objectives. For instance, an energy-harvesting high-level objective may be broken down to different EM absorption commands per SDM area unit. Monitoring software tools are required for establishing two-way communication with the SDM, enabling for adaptive behavior and interconnectivity within smart control loops. Finally, debugging tools are necessary for pinpointing both physical flaws (such as SDM malfunction) and programming logic errors.

From another point of view, these components constitute a software form of the physical laws governing the SDM behavior. Three complimentary approaches are envisioned for accomplishing this transformation:

- SDMs can be treated as white-boxes, using existing analytical models of high-level objectives from the metamaterial world [3]. However, very few such models exist and their generality is limited.
- SDMs can be treated as black-boxes, and learning algorithms can be employed for correlating a high-level objective to a low-level SDM internal state. Such algorithms examine multiple random SDM configurations, converging to an understanding of their behavior. Nonetheless, this process can be computationally expensive and of limited efficiency.
- SDMs can be treated as gray-boxes, empowering the learning algorithms with analytical insights to improve their efficiency.

Heuristics optimizers, such as genetic algorithms, may be used for yielding the optimal control plane state that best fits a sought EM behavior [56].

C. MULTI-PHYSICS SIMULATION

Optimizing the design of an SDM via simulations pertains to its metamaterial and controller communication aspects. From the physics point of view, simulations are required for

defining and optimizing the materials, dimensions, geometry and operating spectrum of the SDM, and deducing the supported range of end-functionalities. From the communications point of view, the operational frequency and transmission power of nodes, their topology, allowed dimensions and materials need to be optimized, balancing minimal cross-talk with the metamaterial plane, communication robustness and overall practicality. Additionally, joint physical/networking simulations are required for developing the SDM software components outlined in Section VI-B.

Due to the aforementioned reasons, simulating SDMs is a necessary step in their design. However, it also constitutes a challenge on its own due to the dissimilarity of the two involved disciplines.

The aspect of physics simulations commonly employs diverse computational and analytical methods (effective medium theories, FDTD, FEM, transfer matrix methods, heuristic algorithms, etc.) to study the EM properties of the metamaterial plane. These techniques are known for their vast requirements in computational resources. The aspect of networking commonly operates at more abstract layers using discrete event simulators. Data packet-level propagation is considered sufficient for many networking systems, while statistical channel models simplify the simulation of the physical propagation medium.

Joining these two different aspects into one uniform simulator is an open challenge. Two possible resolutions are envisioned:

- Both aspects can be joined by a simulation at the physical layer. The periodicity in the SDM geometry can be exploited for reducing the required computational resources to a tractable level. Specifically, node-pairs in identical or similar surroundings can be simulated once and then be cached and re-used for the duration of the simulation.
- The two aspects are kept separate, with the more abstract networking events driving the low-level physical layer simulation. The network communication channel is treated statistically, as described in Section VI-A.

In both cases, it is noted that latest computational methods taking advantage of multiple CPUs and GPUs have exhibited several orders of magnitude shorter simulation times for the physics aspect of this challenge [57].

VII. CONCLUSIONS

SDMs are expected to overcome the main limitations of conventional metamaterials in terms of reusability, adaptivity, and accessibility to the engineering community. The materialization of this vision requires embedding a network of tiny controllers within the metamaterial structure, which represents an important challenge due to the particularities of the application context. On the one hand, we have identified the planar, integrated, and monolithic nature of SDM as characteristics suggesting to treat this application as a scaled version of a manycore embedded system with a NoC,

either wired or wireless. On the other hand, its constrained and ultra-dense landscape, as well as the event-based and correlated nature of the workload, brings SDMs closer to the nanosensor network scenario. A graceful combination of both top-down and bottom-up design approaches may lead to a unique, custom solution meeting the demands of this new disruptive paradigm.

REFERENCES

- [1] N. Engheta and R. W. Ziolkowski, *Metamaterials: Physics and Engineering Applications*. Hoboken, NJ, USA: Wiley, 2006.
- [2] J. B. Pendry, D. Schurig, and D. R. Smith, "Controlling electromagnetic fields," *Science*, vol. 312, pp. 1780–1782, Jun. 2006.
- [3] Y. Ra'idi, C. R. Simovski, and S. A. Tretyakov, "Thin perfect absorbers for electromagnetic waves: Theory, design, and realizations," *Phys. Rev. Appl.*, vol. 3, p. 037001, Sep. 2015.
- [4] S. B. Glybovski, S. A. Tretyakov, P. A. Belov, Y. S. Kivshar, and C. R. Simovski, "Metasurfaces: From microwaves to visible," *Phys. Rep.*, vol. 634, pp. 1–72, Apr. 2016.
- [5] H.-T. Chen, A. J. Taylor, and N. Yu, "A review of metasurfaces: Physics and applications," *Rep. Progr. Phys.*, vol. 79, no. 7, p. 076401, 2016.
- [6] G. Oliveri, D. Werner, and A. Massa, "Reconfigurable electromagnetics through metamaterials—A review," *Proc. IEEE*, vol. 103, no. 7, pp. 1034–1056, Jul. 2015.
- [7] T. Hand and S. Cummer, "Characterization of tunable metamaterial elements using MEMS switches," *IEEE Antennas Wireless Propag. Lett.*, vol. 6, pp. 401–404, Apr. 2007.
- [8] X. He, "Tunable terahertz graphene metamaterials," *Carbon*, vol. 82, pp. 229–237, Feb. 2015.
- [9] C. Liaskos et al., "Design and development of software defined metamaterials for nanonetworks," *IEEE Circuits Syst. Mag.*, vol. 15, no. 4, pp. 12–25, Apr. 2015.
- [10] S. C. Goldstein, J. D. Campbell, and T. C. Mowry, "Programmable matter," *Computer*, vol. 38, no. 6, pp. 99–101, May 2005.
- [11] S. Borkar, "Thousand core chips—A technology perspective," in *Proc. DAC-44*, 2007, pp. 746–749.
- [12] B. Bohnenstiehl et al., "A 5.8 pJ/Op 115 billion ops/sec, to 1.78 trillion ops/sec 32nm 1000-processor array," in *Proc. VLSIC*, 2016, pp. 1–2.
- [13] Y. Choi, Y. Lin, N. Chong, S. Mahlke, and T. Mudge, "Stream compilation for real-time embedded multicore systems," in *Proc. GCO*, 2009, pp. 210–220.
- [14] J. D. Owens et al., "A survey of general purpose computation on graphics hardware," *Comput. Graph. Forum*, vol. 26, no. 1, pp. 80–113, 2007.
- [15] T. Melodia, D. Pompili, V. C. Gungor, and I. F. Akyildiz, "Communication and coordination in wireless sensor and actor networks," *IEEE Trans. Mobile Comput.*, vol. 6, no. 10, pp. 1116–1129, Oct. 2007.
- [16] B. Warneke, M. Last, B. Liebowitz, and K. S. J. Pister, "Smart dust: Communicating with a cubic-millimeter computer," *Computer*, vol. 34, no. 1, pp. 44–51, Jan. 2001.
- [17] I. F. Akyildiz and J. M. Jornet, "Electromagnetic wireless nanosensor networks," *Nano Commun. Netw.*, vol. 1, no. 1, pp. 3–19, Mar. 2010.
- [18] P. Levis et al., "TinyOS: An operating system for sensor networks," *Ambient Intelligence*. Berlin, Germany: Springer, 2005, pp. 115–148.
- [19] P. Sarkar, "A brief history of cellular automata," *ACM Comput. Surv. (CSUR)*, vol. 32, no. 1, pp. 80–107, 2000.
- [20] S. Khasanvis et al., "Architecting for causal intelligence at nanoscale," *Computer*, vol. 48, no. 12, pp. 54–64, 2015.
- [21] J. Han and M. Orshansky, "Approximate computing: An emerging paradigm for energy-efficient design," in *Proc. ETS*, 2013, pp. 21–27.
- [22] R. Venkatagiri, A. Mahmoud, S. K. S. Kumar, and S. V. Adve, "Approxilyzer: Towards a systematic framework for instruction-level approximate computing and its application to hardware resiliency," in *Proc. MICRO*, Sep. 2016, pp. 1–14.
- [23] Q. Xu, T. Mytkowicz, and N. S. Kim, "Approximate computing: A survey," *IEEE Design Test*, vol. 33, no. 1, pp. 8–22, Jan. 2016.
- [24] S. Mittal, "A survey of techniques for approximate computing," *ACM Comput. Surv. (CSUR)*, vol. 48, no. 4, 2016, Art. no. 62.
- [25] B. Grot, J. Hestness, S. W. Keckler, and O. Mutlu, "Kilo-NOC: A heterogeneous network-on-chip architecture for scalability and service guarantees," in *Proc. ISCA*, 2011, pp. 401–412.
- [26] T. Krishna, C.-H. O. Chen, W.-C. Kwon, and L.-S. Peh, "Smart: Single-cycle multihop traversals over a shared network on chip," *IEEE Micro*, vol. 34, no. 3, pp. 43–56, May/June. 2014.
- [27] A. Bakhoda, J. Kim, and T. M. Aamodt, "Throughput-effective on-chip networks for manycore accelerators," in *Proc. MICRO*, 2010, pp. 421–432.
- [28] T. Bjerregaard and J. Sparso, "A router architecture for connection-oriented service guarantees in the MANGO clockless network-on-chip," in *Proc. DATE*, 2005, pp. 1226–1231.
- [29] J. Kim and H. K. H. Kim, "Router microarchitecture and scalability of ring topology in on-chip networks," in *Proc. NoCArc*, 2009, pp. 5–10.
- [30] J. Kim, "Low-cost router microarchitecture for on-chip networks," in *Proc. MICRO*, 2009, pp. 255–266.
- [31] Z. Li, J. S. Miguel, and N. E. Jerger, "The runahead network-on-chip," in *Proc. HPCA*, 2016, pp. 333–344.
- [32] S. Deb, A. Ganguly, P. P. Pande, B. Belzer, and D. Heo, "Wireless NoC as interconnection backbone for multicore chips: Promises and challenges," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 2, pp. 228–239, Jun. 2012.
- [33] S. Laha, S. Kaya, D. W. Matolak, W. Rayess, D. DiTomaso, and A. Kodi, "A new frontier in ultralow power wireless links: Network-on-chip and chip-to-chip interconnects," *IEEE Trans. Comput.-Aided Design Integr.*, vol. 34, no. 2, pp. 186–198, Feb. 2015.
- [34] D. Zhao and Y. Wang, "SD-MAC: Design and synthesis of a hardware-efficient collision-free QoS-aware MAC protocol for wireless network-on-chip," *IEEE Trans. Comput.*, vol. 57, no. 9, pp. 1230–1245, Sep. 2008.
- [35] S. Abadal et al., "Broadcast-enabled massive multicore architectures: A wireless RF approach," *IEEE MICRO*, vol. 35, no. 5, pp. 52–61, Jan. 2015.
- [36] S. Abadal, I. Llataer, A. Mestres, H. Lee, and E. Alarcón, and A. Cabellos-Aparicio, "Time-domain analysis of graphene-based miniaturized antennas for ultra-short-range impulse radio communications," *IEEE Trans. Commun.*, vol. 63, no. 4, pp. 1470–1482, Apr. 2015.
- [37] M. Tamagnone, J. S. Gómez-Díaz, J. R. Mosig, and J. Perruisseau-Carrier, "Analysis and design of terahertz antennas based on plasmonic resonant graphene sheets," *J. Appl. Phys.*, vol. 112, p. 114915, Sep. 2012.
- [38] J. M. Jornet and I. F. Akyildiz, "Graphene-based plasmonic nano-antenna for terahertz band communication in nanonetworks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 12, pp. 685–694, Dec. 2013.
- [39] A. Cabellos-Aparicio, I. Llataer, and E. Alarcón, A. Hsu, and T. Palacios, "Use of THz photoconductive sources to characterize tunable graphene RF plasmonic antennas," *IEEE Trans. Nanotechnol.*, vol. 14, no. 2, pp. 390–396, Feb. 2015.
- [40] S. E. Hosseinienejad et al., "Study of hybrid and pure plasmonic terahertz antennas based on graphene guided-wave structures," *Nano Commun. Netw.*, vol. 12, pp. 34–42, Jun. 2017.
- [41] 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, May 2015.
- [42] I. F. Akyildiz, J. M. Jornet, and C. Han, "Terahertz band: Next frontier for wireless communications," *Phys. Commun.*, vol. 12, pp. 16–32, Oct. 2014.
- [43] J. M. Jornet and I. F. Akyildiz, "Femtosecond-long pulse-based modulation for terahertz band communication in nanonetworks," *IEEE Trans. Commun.*, vol. 62, no. 5, pp. 1742–1754, May 2014.
- [44] J. M. Jornet, "Low-weight error-prevention codes for electromagnetic nanonetworks in the terahertz band," *Nano Commun. Netw.*, vol. 5, nos. 1–2, pp. 35–44, 2014.
- [45] J. M. Jornet, J. Capdevila-Pujol, and J. Solé-Pareta, "PHLAME: A physical layer aware MAC protocol for electromagnetic nanonetworks in the terahertz band," *Nano Commun. Netw.*, vol. 3, no. 1, pp. 74–81, Feb. 2012.
- [46] S. Mohrehkesh and M. C. Weigle, "Optimizing energy consumption in terahertz band nanonetworks," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 12, pp. 2432–2441, Dec. 2014.
- [47] S. Mohrehkesh, M. C. Weigle, and S. K. Das, "DRIH-MAC: A distributed receiver-initiated harvesting-aware MAC for nanonetworks," *IEEE Trans. Mol. Biol. Multi-Scale Commun.*, vol. 1, no. 1, pp. 97–110, Jan. 2015.
- [48] J. M. Jornet and I. F. Akyildiz, "Joint energy harvesting and communication analysis for perpetual wireless nanosensor networks in the terahertz band," *IEEE Trans. Nanotechnol.*, vol. 11, no. 3, pp. 570–580, Mar. 2012.
- [49] R. G. Cid-Fuentes, A. Cabellos-Aparicio, and E. Alarcón, "Energy buffer dimensioning through sensor networks in spatio-temporal-correlated energy-harvesting-enabled wireless sensor networks," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 4, no. 3, pp. 301–312, Apr. 2014.

- [50] A. Tsioliariidou, C. Liaskos, E. Dedu, and S. Ioannidis, "Packet routing in 3D nanonetworks: A lightweight, linear-path scheme," *Nano Commun. Netw.*, to be published, doi: <http://doi.org/10.1016/j.nancom.2017.01.001> and <http://www.sciencedirect.com/science/article/pii/S1878778916300898>
- [51] A. Tsioliariidou, C. Liaskos, S. Ioannidis, and A. Pitsillides, "Lightweight, self-tuning data dissemination for dense nanonetworks," *Nano Commun. Netw.*, vol. 8, pp. 2–15, Oct. 2016.
- [52] C. Liaskos and A. Tsioliariidou, "A promise of realizable, ultra-scalable communications at nano-scale: A multi-modal nano-machine architecture," *IEEE Trans. Comput.*, vol. 64, no. 5, pp. 1282–1295, May 2015.
- [53] A. M. Vegni and V. Loscri, "Performance of a chirality-affected channel exhibiting giant optical activity for terahertz communications," in *Proc. NANOCOM*, 2016, Art. no. 2.
- [54] D. Matolak, S. Kaya, and A. Kodi, "Channel modeling for wireless networks-on-chips," *IEEE Commun. Mag.*, vol. 51, no. 6, pp. 180–186, Jun. 2013.
- [55] S. Abadal, I. Llatser, A. Mestres, and J. Solé-Pareta, E. Alarcón, and A. Cabellos-Aparicio, "Fundamentals of graphene-enabled wireless on-chip networking," in *Modeling, Methodologies and Tools for Molecular and Nano-scale Communications*. Cham, Germany: Springer, 2017, pp. 293–317.
- [56] S. Luke, *Essentials of Metaheuristics*, 2nd ed. Lulu, 2013. [Online]. Available: <http://cs.gmu.edu/sean/book/metaheuristics/>
- [57] N. V. Kantartzis, T. T. Zygiroidis, C. S. Antonopoulos, Y. Kanai, and T. D. Tsiiboukis, "A generalized domain-decomposition stochastic FDTD technique for complex nanomaterial and graphene structures," *IEEE Trans. Magn.*, vol. 52, no. 3, pp. 1–4, Mar. 2016.



SERGI ABADAL received the B.Sc. and M.Sc. degrees in telecommunication engineering and the Ph.D. degree in computer architecture from the Universitat Politècnica de Catalunya, Barcelona, Spain, in 2010, 2011, and 2016, respectively. From 2009 to 2010, he was a Visiting Researcher with the Broadband Wireless Networking Laboratory, Georgia Institute of Technology, Atlanta, USA. In 2015, he was a Visiting Researcher with the i-acoma Group, School of Computer Science, Uni-

versity of Illinois at Urbana-Champaign. Since then, he has been a Post-Doctoral Researcher with the NaNoNetworking Center in Catalunya, where he is the Scientific Director of the VISORSURF Project. His current research interests are ultra-high-speed on-chip wireless networks and broadcast-enabled manycore processor architectures. He was awarded by INTEL within its Doctoral Student Honor Program in 2013.



CHRISTOS LIASKOS received the Diploma degree in electrical and computer engineering from the Aristotle University of Thessaloniki (AUTH), Greece, in 2004, the M.Sc. degree in medical informatics from the Medical School, AUTH, in 2008, and the Ph.D. degree in computer networking from the Department of Informatics, AUTH, in 2014. He is currently a Researcher with the Foundation of Research and Technology-Hellas (FORTH). He has published work in several venues, such as the IEEE TRANSACTIONS ON NETWORKING, the IEEE TRANSACTIONS ON COMPUTERS, the IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY, the IEEE TRANSACTIONS ON BROADCASTING, the IEEE TRANSACTIONS ON SYSTEMS MAN & CYBERNETICS, the IEEE TRANSACTIONS ON NETWORKS AND SERVICE MANAGEMENT, the IEEE TRANSACTIONS ON COMMUNICATIONS, and IEEE INFOCOM. His research interests include computer networks and nanotechnology, with a focus on developing nanonetwork architectures and communication protocols for future applications.



AGELIKI TSIOLIARIDOU received the Diploma and Ph.D. degrees in electrical and computer engineering from the Democritus University of Thrace, Greece, in 2004 and 2010, respectively. She is currently a Researcher with the Foundation of Research and Technology-Hellas (FORTH). Her research is mainly in the field of computer networks and specific focus on the problem of quality of service. Additionally, her recent research interests lie in the area of nanonetworks, with specific focus on architecture, protocols, and security/authorization issues. She has contributed to a number of EU, ESA, and National research projects.



SOTIRIS IOANNIDIS received the B.Sc. degree in mathematics and the M.Sc. degree in computer science from the University of Crete in 1994 and 1996, respectively, the M.Sc. degree in computer science from the University of Rochester in 1998, and the Ph.D. degree from the University of Pennsylvania in 2005. He held a research scholar position at the Stevens Institute of Technology until 2007, and since then, he has been the Research Director with the Institute of Computer Science, Foundation for Research and Technology-Hellas. He has coordinated a number of European and National projects (e.g., PASS, EU-INCOOP, and GANDALF) and is currently the Project Coordinator of the SHARCS and CYBERSURE H2020 European projects. He has authored over 100 publications in international conferences and journals, and book chapters, including ACM CCS, ACM/IEEE ToN, USENIX ATC, and NDSS. His research interests are in the area of systems, networks, and security. He is a Marie-Curie Fellow and has participated in numerous international and European projects. He has both chaired and served in numerous program committees in prestigious international conferences.



ANDREAS PITSILLIDES has been a Visiting Professor with the University of Johannesburg Since 2014. He is currently a Professor with the Department of Computer Science, University of Cyprus, and the Head of the Networks Research Laboratory, which he founded in 2002. He has authored over 250 referred papers in flagship journals (e.g., IEEE, Elsevier, IFAC, and Springer), international conferences, and books. His research interests include communication networks, the Internet and Web of Things, smart spaces (Home, Grid, and City), and nanonetworking. He has participated in over 35 EU and locally funded research projects. He served on several international conferences as the General Chair, an International Co-Chair, the Technical Program Chair, and on executive committees, technical committees, a Guest Co-Editor, and an Invited Speaker. He serves on the Editorial Board of the *Journal of Computer Networks*.



JOSEP SOLÉ-PARETA received the M.Sc. degree in and telecom engineering the Ph.D. degree in computer science from the Universitat Politècnica de Catalunya (UPC) in 1984 and 1991, respectively. In 1984, he joined the Computer Architecture Department, UPC, where he is currently a Full Professor. He did a post-doctoral stage (summers of 1993 and 1994) with the Georgia Institute of Technology. He is a Co-Founder of the UPCC-CABA and UPCN3cat. His publications include several book chapters and over 300 papers in relevant research journals (>80), and refereed international conferences. He has participated in many European projects in the computer networking field. His current research interests are in nanonetworking communications, traffic monitoring and analysis, high speed and optical networking, and energy efficient transport networks, with emphasis on traffic engineering, traffic characterization, MAC protocols, and QoS provisioning. He was the Local Chairman of the 25th Conference on Computer Communications (IEEE INFOCOM 2006) held in Barcelona, in 2006, and the General Chairman of the 7th International Conference on Transparent Optical Networks (ICTON 2005), in 2005, and the 5th Workshop on Quality of Future Internet Services (QoFIS04, 2004).



ALBERT CABELLOS-APARICIO received the B.Sc., M.Sc., and Ph.D. degrees in computer science engineering from the Universitat Politècnica de Catalunya (UPC), Barcelona, Spain, in 2001, 2005, and 2008, respectively. He has also been an Assistant Professor with the Computer Architecture Department and a Researcher with the Broadband Communications Group, Technical University of Catalunya, since 2005. In 2010, he joined the NaNoNetworking Center in Catalunya, where he is the Scientific Director. He has also founded the LISPmob open-source initiative along with Cisco. He has been a Visiting Researcher with Cisco Systems, San Jose, CA, USA, and Agilent Technologies, Santa Clara, CA, USA, and a Visiting Professor with the Royal Institute of Technology, Stockholm, Sweden, and the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA. He has given more than ten invited talks (MIT, Cisco, INTEL, MIET, and Northeastern University) and co-authored over 15 journal and 40 conference papers. His main research interests are future architectures for the Internet and nanoscale communications. He is an Editor of *Nano Communication Networks* and the Founder of the ACM NANOCOM conference, the IEEE MONACOM workshop, and the N3Summit.

...



EDUARD ALARCÓN received the M.Sc. (National Award) and Ph.D. degrees in EE from the Universitat Politècnica de Catalunya (UPC), Spain, in 1995 and 2000, respectively. He became an Associate Professor with UPC in 2001. He was a Visiting Professor with the University of Colorado at Boulder, USA, in 2003, 2006, and 2008, respectively, and KTH Stockholm in 2011. He has co-authored over 300 scientific publications, eight book chapters and eight patents, and has been involved in different national, EU, and US R&D projects. His research interests include the areas of on-chip energy management circuits, energy harvesting and wireless energy transfer, nanocommunications, and small satellites. He has been funded and awarded several research projects by companies, including Google, Samsung, and Intel. He has given 30 invited lectures and tutorials worldwide. He is the Vice President of the IEEE CAS Society. He was an elected member of the IEEE CAS Board of Governors from 2010 to 2013, an IEEE CAS Society Distinguished Lecturer, a recipient of the Best paper award at the IEEE MWSCAS98, a Co-Editor of six journal special issues and eight conference special sessions, a TPC Co-Chair and a TPC Member of the 30 IEEE conferences, and an Associate Editor of the IEEE TCAS-I, TCAS-II, JETCAS, JOLPE, and *Nano Communication Networks*.