Self-Awareness in Systems on Chip—A Survey

Axel Jantsch
TU Wien

Nikil Dutt and Amir M. Rahmani
University of California at Irvine

Editor's note:
Self-awareness is desirable feature of emerging computing systems. It helps systems to understand, manage, and report on their own system behavior. This paper presents an overview centered around the paradigm of self-awareness in computing systems.
—Partha Pratim Pande, Washington State University

In addition to its roots in psychology, the notion of self-awareness has been used in computing in a variety of different domains such as autonomic computing, organic computing, adaptive systems, and self-organizing systems, often with different, implicitly given definitions and objectives. Thus, a complete survey with all relevant work is impossible in the limited space of this article and a consistent treatment of this concept across all domains is a challenge. Consequently, our survey is incomplete in that it does not cover all interesting work. Rather it tries to

- explain the motivation of researchers and their interest in this topic,
- show its benefits,
- provide a paradigmatic reference frame that relates to all key ingredients of self-awareness,
- give a cursory historical account, and a representational and fair exposition of the concepts of self-awareness in the various domains with systems on chip as the main focus.

Digital Object Identifier 10.1109/MDAT.2017.2757143
Date of publication: 27 September 2017; date of current version: 14 November 2017.

First, we briefly introduce self-awareness and provide a reference definition of the term. Then, in "Benefits in Systems on Chip," we list benefits and motivate why researchers have so extensively studied and used the concept. In "Self-Awareness Paradigm," we introduce a paradigmatic architecture of a self-aware system (Figure 1) which, in our opinion, is complete in the sense that it contains all crucial elements of self-awareness. Since in the literature the term self-awareness is often used to encompass only parts, frequently different parts, of the elements in our paradigmatic architecture, it serves as reference to which previous work can be easily related. In "Related Research Directions," we discuss the main domains where self-awareness has been more or less extensively used, namely, autonomic computing, self-adaptive systems organic computing, and control theory. In a way, the first three sections can be considered an introduction to our main topic, self-awareness in SoCs. This lengthy introduction is necessary because of the diverse use of the concept in various domains, but we hope to provide a solid basis and understanding for the discussion of work on on-chip self-awareness and to appreciate the use and utility of the involved concepts. Finally, in "Challenges," we list and briefly discuss the most important challenges of further research in this field.

What is self-awareness?

When engineers contemplate a concept they instinctively ask how it can be made useful and the

Copublished by the IEEE CEDA, IEEE CASS, IEEE SSCS, and TTTC

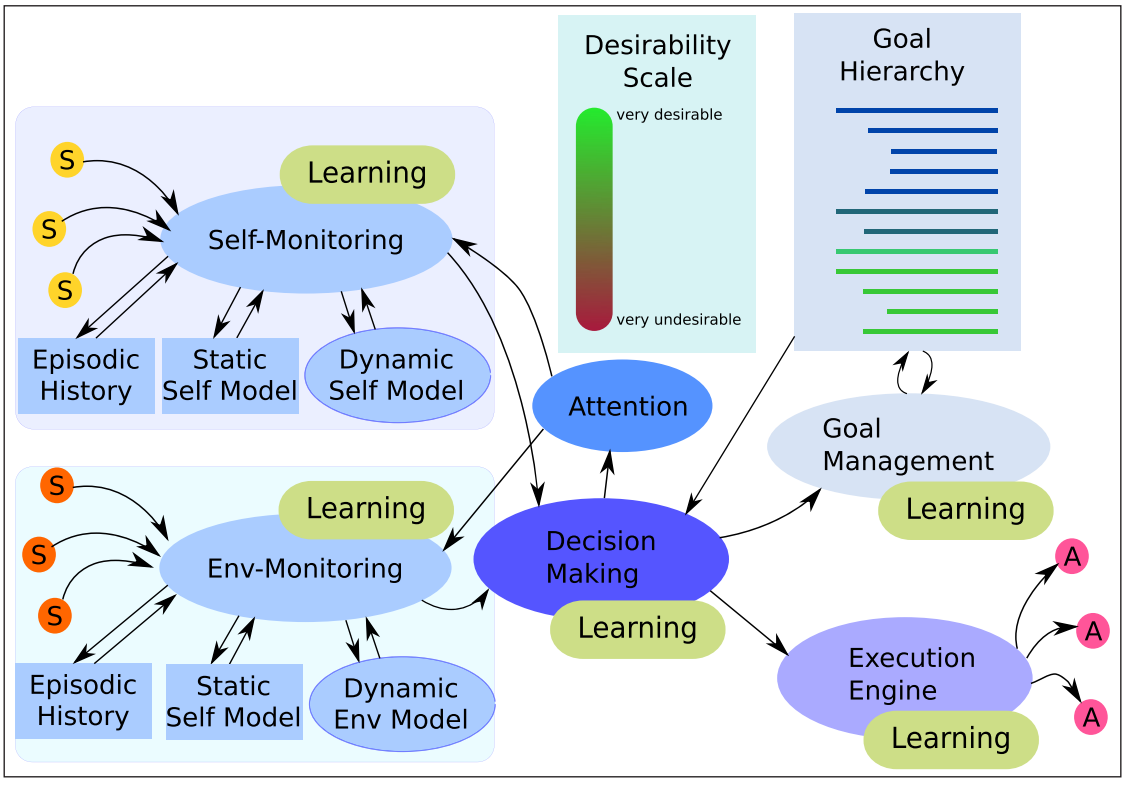

Figure 1. Paradigmatic architecture of a self-aware system. "S" nodes are sensors and "A" nodes are actuators. A proper assessment of the self and the environment (self-monitoring and environmental-monitoring) are the basis for active goal management and effective decision making. The desirability scale is the currency of assessment. All assessments that are considered "good" or "bad" are mapped onto this scale. This allows the comparison of otherwise unrelated properties, e.g., the quality of a signal and the load-level of the battery. Machine learning algorithms are useful for all activities. Both monitoring functions can continuously improve to identify the normality and categorize aberrations. Goal management can learn to dynamically prioritize subgoals in a way to optimize the accomplishments for high-level goals. The decision making can optimize its algorithm based on the effect of its decisions on the system's performance. The execution engine can learn to generate control commands for the best possible effect. Note that many connections are not drawn for the sake of clarity.

desire to fully understand all the details and implications seem to be less important than to find a way to utilize it for a practical end. This has been the fate of self-awareness as a subject of study in the context of computing during the last 20 years. The pyramid of self-* properties (Figure 2), originally proposed in the IBM initiative on autonomic computing in 2003 [1], [2], illustrates this point. Motivated by the objective to make software systems more flexible, and truly self-adaptive researchers have identified self-configuration, self-healing, and other self-* features as essential properties with self-awareness and context-awareness located at the primitive level.

Although adaptive systems with no self-awareness exist, it has been argued that a sophisticated self-model is a prerequisite for sensible adaptive behavior when the environment and the system itself are sufficiently complex and there exists a causal relation between self-* properties and high quality in complex software systems [3].

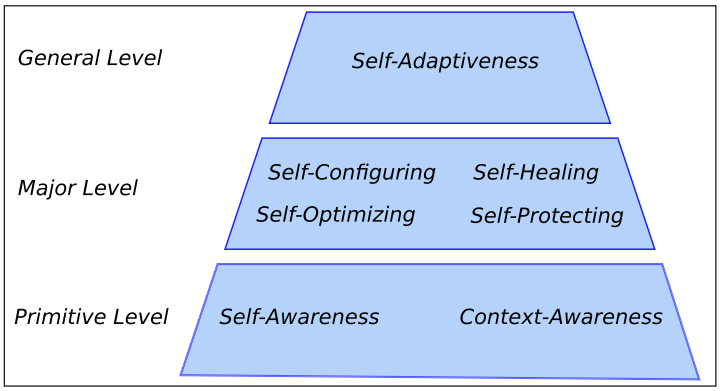

Figure 2. The hierarchy of self-* properties, first proposed in 2001 but cited here from Salehie and Tahvildari [2] in 2009.

As a consequence of this goal oriented approach in the study of self-awareness in computing systems, the research agenda has been dominated by a quest for utility. While this approach leads more directly to applicable results, it also makes a deeper understanding of the concepts appear less desirable. As an example, we quote fully a definition of self-awareness offered by Kounev [4] in 2011 and also in a modified form more recently in 2015 by Kounev et al. [5]:

Self-awareness, in this context, is defined by the combination of three properties that IT systems and services should possess:

1) *Self-reflective*: i) Aware of their software architecture, execution environment, and the hardware infrastructure on which they are running, ii) aware of their operational goals in terms of QoS requirements, service-level agreements (SLAs) and cost- and energy-efficiency targets, and iii) aware of dynamic changes in the above during operation.
2) *Self-predictive*: Able to predict the effect of dynamic changes (e.g., changing service workloads or QoS requirements) as well as predict the effect of possible adaptation actions (e.g., changing service deployment and/or resource allocations).
3) *Self-adaptive*: Proactively adapting as the environment evolves in order to ensure that their QoS requirements and respective SLAs are continuously satisfied while at the same time operating costs and energy-efficiency are optimized.

It is an excellent definition, which we will use as reference in this survey. However, there are two interesting observations to note. First, self-awareness has moved up in prominence. It has been at the bottom of the self-* pyramid in 2001 (Figure 2) as a supporting feature for more advanced adaptive behavior. In the 2011 definition, the term is used to encompass all relevant self-* properties including self-adaptiveness. In a way the pyramid has been turned upside down because the community has realized that self-awareness is not a simple collection of state variables that describe the state of the system (item 1 in the above definition), but it has also to include the operational goals of the system and it has to properly reflect the effects of its own actions and of environmental changes on these state variables. Essentially, it has to include a complete, even though abstracted, model of the static and dynamic system properties. Once the dynamic properties are also properly represented, a prediction, perhaps by simulation, of the system and its interaction with the environment becomes possible and it is self-predictive according to item 2 in Kounev's definition. Based on this capacity of prediction, it can proactively adapt its own actions earlier than it would otherwise be possible (item 3). Hence, rather than being a simple elementary property to be used by a self-adapting system controller, self-awareness, once fully accomplished, makes proactive, adaptive behavior almost straight forward.

The second point to note about the definition is that it is purely utilitarian in that it does not try to capture the essence of the rich concept of self-awareness. It assumes that self-awareness as defined, i.e., self-reflective + self-predictive + self-adaptive, is useful for accomplishing the QoS requirements and the service-level agreements of the system, and thus, it should be implemented. It does not address the question, why it deserves to be a separate concept and what it adds to self-reflective, self-predictive, and self-adaptive. Although we have no definite answer to these questions, it may well worth to ponder them in order to identify additional aspects that are essential for self-awareness but not yet fully accounted for in the state of the art. There is some indication that learning, keeping track of history, and dynamic goal management are such essential aspects. For instance, Chandra et al. [6] argue that a system has to acquire a substantial part of the

self-model during operation by some kind of learning process to be considered self-aware. Externally built and implanted knowledge does not suffice. However, their argument is not of principal nature but claims that it will be very difficult to develop a sufficiently accurate self-model without a dynamic tuning and optimization process. While this may be debatable, it illustrates the complexity of the concept even if the only concern is its usefulness in an engineering endeavor.

As we discuss in "Related Research Directions," autonomic computing is not the only branch of research that has struggled with the concept of self-awareness. Organic computing, bioinspired computing, and self-organization are other prominent lines of research that have approached the topic from different angles and contributed with specific insights and solutions. Before we discuss them in the above-mentioned section, we summarize the potential benefits ("Benefits in Systems on Chip") and sketch a self-aware computing paradigm in "Self-Awareness Paradigm" that serves as a reference and reflects to some degree all major proposals for self-aware systems and architectures.

Benefits in systems on chip

Researchers of self-awareness generally argue that it allows a system to deal better with complexity. The complexity comes from the system itself (its structure and its state space), from the environment, and from the exceedingly diverse goals and objectives it has to meet.

Hardware state assessment and management. Self-awareness in hardware systems often facilitates the management of temperature [7], power/energy [7]–[9], and real-time performance [9]–[11]. Also, aging effects are addressed with increasingly complex models to assess the progress of aging and select counter measures [12].

Resource allocation. Given often the high number of processor, memory, and interconnect resources available, resource allocation is a common target of self-aware enabled management schemes in both hardware and software. Task allocation and MPSoC configuration with higher energy efficiency based on cross-layer self-awareness on chip is proposed by Sarma and Dutt [13]. Because the approach extends across individual components (cores, routers) and layers (HW, NoC, OS middleware, and application), local assessments have to be integrated into a comprehensive system-level assessment. Based on a large number of sensors and comprehensive self-assessment of the system load balancing, task allocation, scheduling, and migration for MPSoCs result in significant improvements [14], [15].

In the SElf-awarE Computing (SEEC) framework, dynamic adaptation through a smart interface between platform and application is achieved [9], [16]. The application registers its performance goals and the platform is responsible for meeting those goals by continuously monitoring the system's performance and appropriately adjusting the resource allocation.

An appropriate self-model allows to allocate scarce communication resources efficiently. Happe and Trammel-Keller [17] use flexible protocol stacks to allow for dynamic rearrangements and optimization of the communication protocols based on needs, requirements, and constraints.

In general, IT systems meeting quality of service requirements and service-level agreements under energy-cost constraints is a tremendously complex task [4]. This challenge has driven the field of autonomic computing during the last two decades and with the growing size and complexity of the applications and the IT systems, the self-models and the self-awareness concepts have grown in complexity and sophistication as well [1], [5].

Reaction to changes in the environment. Many systems that interact with the physical environment by means of sensors and actuators have to adapt to changing conditions. Adaptive systems, as for example surveyed by Krupitzer et al. [18] and further elaborated in "Self-Adaptive Systems," have been studied in various applications. By providing a comprehensive assessment of the state of the system and its environment, self-awareness offers a solid foundation for adaptation decisions and consequently can increase the quality of adaption. Indeed, we expect a direct dependence of the quality of adaption on the quality of self-assessment.

In summary, we conclude that self-awareness leads to more sensible behavior based on more detailed and often more explicit representation of the system's goals, its own state (available resources, faults, etc.), and the environment.

Moreover, it leads to more efficiency due to better and adequate usage of resources. It can be used to detect aberrations of the system's behavior (faults, aging, malicious attacks, design errors, etc.) and of the environment. However, many of these potential benefits are only superficially studied and it remains to be seen what solutions can be found and how effective they are.

Self-awareness paradigm

Figure 1 shows a paradigmatic architecture of a self-aware system that allows to cover and relate most of the systems proposed in the literature. However, this is not the only or the best way to illustrate the structure of self-awareness. Being a loosely defined umbrella concept, there are many options regarding what to include and what to exclude, what to highlight and what to deemphasize, and what to make explicit and what to make implicit. Other researchers have made different choices, e.g., Lewis et al. [19], presumably due to different preferences in their research. For us, attention, goal management, and a central desirability scale are key elements not found in other architectures of self-awareness. Hence, we use Figure 1 as a basis of discussion in this article but in the absence of either theoretical arguments or empirical evidence that clearly favors one architecture over another, we suggest to pragmatically use whatever is more suitable in a given context.

Kounev's [4] definition cited above in "What is Self-Awareness?" is represented in this figure as follows. The *self-reflective* part is located in (i) the static self-model, (ii) the goal management and goal hierarchy, and (iii) the dynamic self-model. The *self-predictive* part is located in the dynamic self-model, and the *self-adaptive* part is located in the decision making, the goal management, and the goal hierarchy. More recently, Kounev et al. [5] have revised this notion and have formulated the following definition.

Self-aware computing systems are computing systems that:

- *learn models* capturing *knowledge* about themselves and their environment (such as their structure, design, state, possible actions, and run-time behavior) on an ongoing basis
- *reason* using the models (for example, predict, analyze, consider, and plan) enabling them to act based on their knowledge and reasoning (for example, explore, explain, report, suggest, self-adapt, or impact their environment) in accordance with *higher-level goals*, which may also be subject to change.

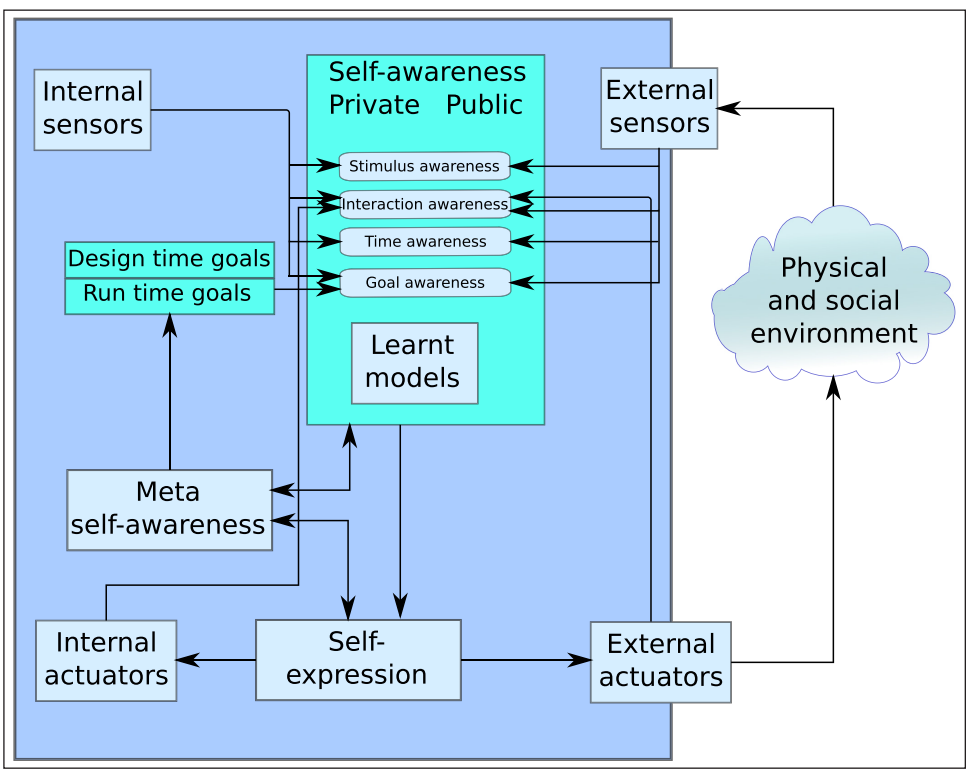

Figure 3. Reference architecture for self-aware computing systems proposed by Lewis et al. [19].

Item 1 can be found in the green learning boxes of Figure 1 and item 2 in the boxes dynamic self-model, goal management, and decision making.

A self-awareness reference architecture has been proposed by Lewis et al. [19] as shown in Figure 3. All its important elements can be mapped to the paradigmatic architecture of Figure 1 but a few points are worth noting.

Meta self-awareness refers to the ability to be aware of and reason about its own self-awareness. It allows to control and dynamically change

the level of self-awareness and thus the resources expended on the self-awareness processes themselves. In some situations, or perhaps most of the time, it is unnecessary to keep these processes active because other tasks have higher priority. In Figure 1, this is not made explicit but can be considered part of the goal management strategy and decision procedure. We have chosen to keep it implicit because meta self-awareness is a rather specialized feature and expected to be present in only few, high-end self-aware systems.

Similarly, *time awareness* is made explicit in Figure 3 and refers to the capability to explicitly reason about the state changes of the system and its environment over time. In the paradigmatic sketch of Figure 1 this is not explicit but the means for it are provided by keeping track of the history, by representing dynamic changes, and by the goals and decision routines. In the same way, stimulus awareness, interaction awareness, and goal awareness refer to abilities to reason about specific aspects of the system. They are not made explicit in Figure 1, but may be included as part of specific goals, decision procedures, and the dynamic internal models.

Figure 3 distinguishes between private and public self-awareness. Public aspects can be inspected from the outside like physical size, battery load level, and initiated actions. Private aspects are not directly visible outside and may refer to internal sensors, counters, and registers. Both are part of the static self-model in Figure 1 but not distinguished.

On the other hand, Figure 3 does not make explicit history mechanisms, the distinction between self-model and environment models, attention, desirability, goal management, and the various places where learning contributes to continuous tuning and optimization.

In Figure 1, the processes of self-monitoring and environment-monitoring are fairly separated and only their results are only combined for decision making. In contrast, Figure 3 treats both as one integrated process. We consider awareness to be the result of a hierarchical process where in the first-level data from individual sensors are preprocessed and filtered individually after which more and more sensory information is gradually fused to establish increasingly abstract concepts. The integration of information from internal and external sensors occurs in most cases rather late in the hierarchy. Thus, it is justified to represent the processes responsible for self- and environment-monitoring as separate activities. However, they may exchange information at every step and for some systems a more integrated solution may be preferable as depicted in Figure 3. In fact, complete isolation and complete integration of these two processes should be considered as extreme points in a continuous design space with practical solutions will almost always fall somewhere in between.

Dutt et al. [20]–[22] have listed relevant features in self-aware systems and have defined them as follows:

- **Semantic interpretation** includes an appropriate abstraction of the primary input data and a disambiguation of possible interpretations.
- **Desirability scale** provides a uniform goodness-scale for the assessment of all observed properties.
- **Semantic attribution** maps properties into the desirability scale suggesting how good or bad an observation is for the system.
- **History of a property:** Awareness of a property implies awareness of its change over time.
- **Goals** provide the context in which interpretation and semantic attribution is meaningful.
- **The purpose** of a smart embedded systems is to achieve all its goals.
- **Expectation on environment:** The system expects a specific environment and detects if the environment deviates significantly from expectations.
- **Expectation on subject:** Similarly, the system's own state and condition are continuously assessed to detect deviations, degradation, performance, and malfunctions.
- **Inspection engine:** Continuously monitoring and assessing the situation requires a specific machinery that integrates all observations into a single, consistent world.

All these processes can be identified in Figure 1. Semantic interpretation and attribution are not shown in the figure and are performed in the monitoring blocks and influenced by the goals and their priorities. A dynamically changing goal hierarchy will also modify the semantic attribution and attention. The inspection engine is not explicit in Figure 1, but is part of the self-monitoring block with the help of several other blocks. An interesting point in this

list of features is the emphasis on data abstraction and the semantic interpretation in the context of goals and an application. The importance of these processes has been elaborated by Taherinejad et al. [23] and they are part of the self and environmental monitoring tasks in Figure 1.

Related research directions

Autonomic computing

After the formulation of its vision in 2003 [1], the field has quickly grown and flourished. In a keynote at the International Conference on Autonomic Computing, Jeff Kephart [24] counted overall 8000 papers published, 200 patents issued, and 200 conferences soliciting papers on the topic of autonomic computing. Since then, the field has continued to be active but has diversified and overlapped with control, machine learning, cloud computing, and web services. A main feature in much of the work on autonomic computing is a variation of the MAPE-K control loop [25] that illustrates the monitor-analyze-plan-execute cycle and is based on knowledge, which often means some kind of model, as shown in Figure 4. Interestingly, the generality of this model has not increased, but in many approaches the models have been customized for a more specific purpose like resource management or maintaining a specific QoS level. An example of work against this trend is the approach of Sanz et al. [26], which proposes a secondary control loop on top of the inner control loop resembling an explicit self-model. This outer loop is derived from the design time model but used during operation.

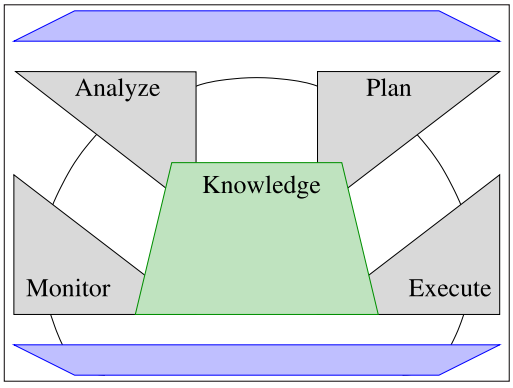

Figure 4. MAPE-K Loop illustrating a monitor-analyze-plan-execute cycle based on Knowledge [25].

However, most of the work in the field has adopted less general and more specialized self-models. Even today, central themes are still self-adaptation, self-optimization, self-configuration, and self-healing [27], [28], but in industrial practice its original vision has not fully materialized. There, trigger-based approaches are still dominant [29], [30], which means that triggering rules fire when a metric such as resource utilization or load imbalance exceeds a threshold value. In academia, a number of systems with model-based performance and resource management have been developed to assure quality of service levels, for instance DiVA [31], MADAM [32], MUSIC [33], and SASSY [34]. They typically use formalisms like Petri nets [35], queuing networks [34], stochastic process algebras [36], statistical regression [37], or kriging models [38] for performance modeling. However, from our perspective, their self-models are limited because they all do not take the software architecture and the execution environment of the system into a detailed account. A survey from Becker et al. [39] confirms this impression. Hence, these systems have limited self-awareness. On the other hand, approaches that do take the software system and execution environment into account are mostly used at design time and not part of the system during operation [40].

It seems that the more sophisticated aspects of the autonomic computing vision has had limited impact and practical solutions based on traditional performance models, and heuristic rule-based approaches have so far been sufficient to address the industry's need. This can on one hand be attributed to the conservative instinct of managers that prefer practically well proven and understood solutions and, on the other hand, to the availability of inexpensive computing, memory, and communication resources that provide little incentives to find the most optimal or efficient solution. We have still limited understanding of the implications at the system level when advanced techniques from the machine learning, the control theory, and optimization domains are integrated with complex models. This has also been concluded by Kounev et al. [5] at the 2015 Dagstuhl Seminar:

> Another finding was that much work remains to be done at the system level. In particular, while there has been considerable success in using machine learning and feedback

control techniques to create adaptive autonomic elements, few authors have successfully built autonomic computing systems containing a variety of interacting adaptive elements. Several authors have observed that interactions among multiple machine learners or feedback loops can produce interesting unanticipated and sometimes destructive emergent behaviors; such phenomena are well known in the multiagent systems realm as well, but insufficiently understood from a theoretical and practical perspective.

Self-adaptive systems

Work on self-adaptive systems naturally emphasize the task of adaptation and considers self-awareness properties only so far as they help to accomplish adaptation. It turns out that for more sophisticated adaptation models of the system and its environment become crucial, leading to *model-based* approaches [18]. Typically, three types of models are distinguished: system models to represent the system state, goal models to represent policies and rules, and environmental models to capture the context [18], [41]. Most work on model-based self-adaptation has been reported for general software systems. Based on the Software Engineering Institute's notion of software product lines [42], a number of approaches model different system features as a basis for selecting dynamically the most appropriate configuration in a particular situation (see [32], [33], [43]–[46]; more examples are discussed in the surveys by Huebscher and McCann [41] and by Krupitzer et al. [18]). However, it should be noted that a model of different software configurations does not constitute self-awareness. Self-awareness, as we understand the term in this survey, is based on a process of dynamically, if not continuously, acquiring data about the system itself and its environment to infer the current state and condition. Thus, most work on self-adaptive software systems do not cover self-awareness as defined in "What is Self-Awareness?" and "Self-Awareness Paradigm" above, even when using various models of the system extensively.

Work on self-adaptive resource constrained cyber-physical systems is more limited but comes closer to our notion of self-awareness. For the domain of smart cities and buildings, Gürgen et al. [47] propose a self-aware cyber-physical architecture that manages the data collection from sensors, the analysis, the planning, and the adaptation of the controlled object (e.g., a building). Smart camera networks have to deal with quickly changing, diverse, and complex environments. Esterle et al. [10] argue that fixed configurations are infeasible and the benefits of self-awareness are due to its coordinating effect on a distributed assessment and decision making, flexible rearrangements of the network under performance, cost, and real-time constraints. In both examples, important features of self-awareness are included but aspects like learning, goal management, attention, and a central desirability scale are only rudimentary present or not at all.

Organic computing

In the early 2000s, the increasing complexity of computing systems led people to conclude that unexpected emergent behavior is unavoidable once a certain complexity has been reached. Consequently, the design of desired and the control of undesired emergent behavior were identified as main challenges. In 2005, Schmeck [48] formulated the vision of Organic Computing as a response to the threatening view of being surrounded by interacting and self-organizing systems, which may become unmanageable, showing undesired emergent behavior. In the following years, the paradigm of Organic Computing was explored in a series of research projects, and in 2011, this work has been nicely summarized in the book *Organic Computing*, edited by Müller-Schloer et al. [49]. Kramer et al. [50] proposed a two level monitoring approach to self-awareness. The low-level monitoring is based on counting events such as cache misses, fault occurrences, or performance counters. In principle, any event that can be counted can be subject to this mechanism. The monitor can be programmed during operation to associate any type of event with event-IDs allowing for flexibility with respect to the kind of events under observation. High-level monitoring uses the event counts for state classification to reflect relevant information about the systems performance and state. Since event grouping and limited event abstraction is possible, the resulting system can be considered rudimentarily self-aware. In a similar spirit, learning classifier systems [51] and eXtended Classifier Systems (XCS) [52] have been used to assess a systems state as a base for decision making such as load management and task allocation [53].

The decisions are coded in rules. A rule consists of a condition, an action, and a predicted reward value. If the condition of a rule matches, i.e., if the system is in the state described by the classifier, and the expected reward is sufficiently high, the action part of the rule is triggered. The rules are optimized by heuristics such as genetic algorithms and reinforcement learning. In the DodOrg project, these and other ideas have been integrated to provide self-awareness in a many-core architecture, which in turn is used for power and thermal management [54].

In summary, the organic computing community has developed a number of innovative approaches to monitoring, adaptation, self-organization, distributed control, and particularly contributed to a better understanding of phenomena of emergent behaviors such as emergent control [55]. However, similar to the autonomic computing endeavors, it has focused more on the decide and act parts of the observe-decide-act (ODA) cycle. For instance, Kramer et al. [50] have observed that in order to enable required self-organization capabilities, a monitoring infrastructure has to provide self-awareness, but have not well defined what is meant with self-awareness and have used a rather limited and static scope of the concept. Interesting aspects of awareness such as abstraction, attention, awareness of the historic changes in its own behavior and in the environment have been hardly touched upon.

Control theory

All the problems mentioned in "Benefits in Systems on Chip" have been successfully addressed without the explicit label of self-awareness. Numerous algorithms for scheduling and task allocation have been developed and deployed in real, demanding large scale systems and temperature and power managers are routinely built into each and every chip on the market. For instance, on-chip dynamic power management [55], [56] has been accomplished by control loops like more or less complex proportional, differential, integral (PID) controllers, where measured or estimated temperature, current flows, and energy levels in batteries are used to tune voltage, frequency, and application load in order to meet given constraints and optimization objectives. Power management is a case in point how exceedingly complex internal models have been used as the problems become increasingly challenging and sophisticated over time. In simple processors of the 1980s and 1990s hardcoded and simple algorithmic solutions have dominated [56], while recent many-core SoCs operating at the edge of thermal stability require advanced power management based on detailed information reflecting the state and objectives of the hardware and the applications [55]. In many core heterogeneous SoCs with several applications concurrently sharing the platform the application behavior regarding the computational load, memory access, and communication can vary over orders of magnitude in short time periods and are often highly unpredictable. P. Bogdan and colleagues have shown how accurate, statistical modeling of workload can significantly improve the efficacy power management [57]–[59].

There have also been efforts to hierarchically manage complex many-core systems by leveraging different structures of feedback control loops. For instance, in [60], a number of nested feedback control loops with different knobs and actuation epochs have been hierarchically deployed for power management with the objective of maximizing the performance while respecting the thermal design power budget. A centralized power management approach with the same objective is presented in [61], which considers both communication and computation characteristics of many-core systems in the power management policy. In a similar fashion, in [62], a coordinated power management approach with multiple scopes of actuation (virtual machine, cluster, server, and core) is presented to cap the power consumption of the system and balance the utilization of the blades. Even though these approaches have proved to be effective to manage complexity, they focus on a single objective which is the main reason why they use several simple single-input single-output PID controllers to form a larger manager for the respective problem at hand, which is often *maximizing performance under a power cap* [63].

There have been recently some contributions to leverage more advanced control theory approaches such as linear-quadratic-Gaussian controller [64] to implement multiple input, multiple output (MIMO) formal control for maximizing resource efficiency. For instance, in [65], Pothukuchi et al. utilize a MIMO controller to track throughput (billions of instructions committed per second) and power consumption for an out-of-order single-core processor in a coordinated manner. Even though MIMO

controllers have the advantage of tracking different references with different priorities, they cannot efficiently be applied to complex systems as obtaining state-space models for complex systems is impractical, if not infeasible.

In general, tracking a single or multiple reference values, which is also called regulatory control, is the main application of control theory. As can be observed from the aforementioned examples, this property is essentially useful for problems where minimizing the tracking error of a parameter is the main goal. For instance, providing a certain quality of service (e.g., frame rate) for a real-time application, capping power consumption of a system, or controlling the thermal behavior of the chip are among popular use-cases for control theory. However, effects of actuations that cannot be modeled using difference functions (e.g., task migration) or problems that need optimization (e.g., minimize an objective function under constraints) cannot be properly addressed using classic control theory. From another perspective, thanks to the feedback-based structure, control theory-based approaches are the best fit for problems such as disturbance rejection (e.g., disturbance due to workload variations when applying dynamic voltage and frequency scaling to control power) or handling noise/uncertainty in measurements (e.g., noisy sensors and virtual sensing), however, it is ineffective to adapt or react to anomalies (e.g., faults), surprises, or radical changes in high-level goals.

In summary, control theory provides guarantees, has the ability to learn from feedback, and has the luxury of formal reasoning and methodology. However, restrictions such as the difficulty to obtain control theoretic models (e.g., transfer functions) in the form of difference equation and lack of a straightforward process to specify reference values limits its efficacy to be solely used for managing complex computing systems. On the other hand, while this approach works for a limited set of parameters and objectives, it does not scale well when the complexity of modeling the system dynamics increases. Modeling complexity escalates with the number of control inputs (i.e., knobs), measured outputs (i.e., sensors), and subsystems (e.g., cores) in multi- and many-core systems. On top of that, heterogeneity of subsystems (e.g., in big.LITTLE style processors) makes the system modeling/identification even more complex. For instance, when aging effects, hard and soft real-time constraints, and transient and permanent faults have to be considered in addition to temperature limits and battery life time, and tuning knobs are available at circuit, architecture, operating system, and application level, the control loops become too complex to be used. It should be noted that the first step to design a control-theoretic approach is to have an accurate-enough model in hand.

Self-awareness offers the promise to be a scalable heuristic in that it can integrate any number of parameters and still provide workable solutions in real time and with sufficient quality. As in control theory, the set points need to be specified by a higher entity, and the integration of self-awareness with control theory can provide a layer of cognition for controllers to coordinate them toward the current goal of the system. So far this claim is still largely a promise but recent work and also the articles in this special issue show encouraging progress.

Self-awareness on chip

Features of self-awareness have found their way in many SoC resource management solutions. The vast majority follow a classic control loop approach, opportunistically extending and customizing them in ad-hoc ways as needed. In the following, we discuss four examples that stand out in that they have self-awareness built into their architectures from the very start.

ASoC. The autonomic SoC platform (ASoC) [66], [67] is based on the organic computing paradigm and aimed at many-core architectures. Functional processing units, which are traditional cores, accelerators, memories, and other functional hardware units, are monitored and controlled by units in a parallel layer, called the *autonomic layer*. For each core or similar components in the functional layer, there is a corresponding element in the autonomic layer, named the *autonomic element*, that consists of a monitor, an evaluator, an actuator, and a communicator, as illustrated in Figure 5. For instance, the autonomic element may monitor the load level in the functional element and update the frequency accordingly. The communicators allow the autonomic elements to communicate with each other. Since each functional element is shadowed by an autonomic element, we have a distributed control system.

The evaluators are rule based. Each rule consists of a matching pattern, an action, and a reward value.

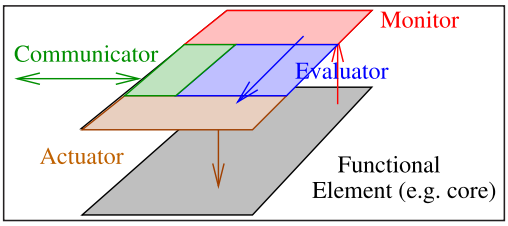

Figure 5. ASoC has two layers: a functional layer consisting of cores and the like, and an autonomic layer that controls the functional elements via a monitor-evaluator-actuator loop [67].

The patterns match the monitored values and determine which rule can apply in a given situation. For example, a pattern could encode "too high load." The action encodes the change of frequency and the reward value estimates how much the action will improve the situation. This reward value is then updated based on the actual improvement as observed by the monitor.

ASoC exhibits some of the features of our paradigmatic architecture in Figure 1. There are self-monitoring, decision making, and execution components. Learning is present in a limited form. The desirability scale and the goals are implicitly coded in the rules. Attention, environmental monitoring, goal management, and the more sophisticated elements of self-monitoring, such as the assessment of the reliability of the measured data, are missing. However, it is conceivable to extend the ASoC architecture to include those elements of self-awareness as well.

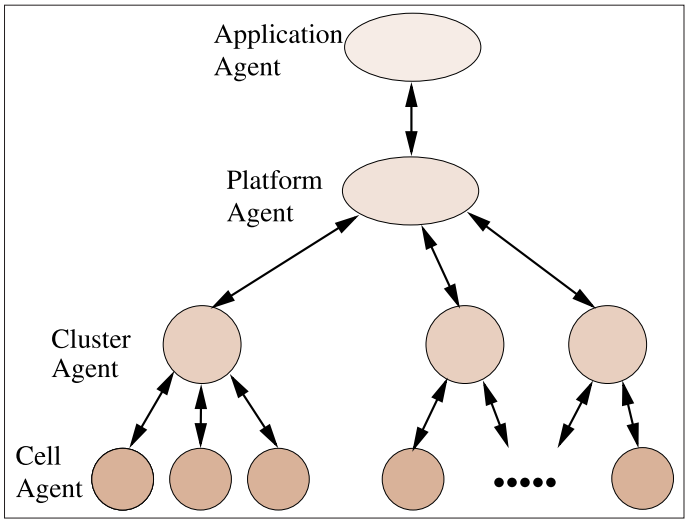

Figure 6. The four-level control structure in HAMSoC [8].

SEEC. Hoffmann and coworkers at MIT have developed SEEC [68], a general framework for self-aware computing using an ODA paradigm. The system cyclically monitors key features, applies a control and decision algorithm, and deploys appropriate actions to adapt to changes in the environment and its own state. It is based on the heartbeats API library [69], which defines a cyclic event called a heartbeat. Through API functions, the application can register rate and latency performance goals in terms of the heartbeat period. Hence, the heartbeats API is a standardized means to monitor the performance of an application. The SEEC controller adapts and optimizes the system's behavior, for instance, by allocating and scheduling resources appropriately. The approach has been evaluated in several applications for performance optimization [68], power management [70], [71], and managing multiple objectives [9]. Also, the concept of *knobs* has been introduced [71] to expose steering facilities such as processor speed or power modes. SEEC allows to adopt different decision making strategies and algorithms that have been studied extensively [72], [73].

Relating the self-awareness features of SEEC to the paradigmatic architecture of Figure 1, we note that the monitoring-deciding-execution loop is thoroughly elaborated in SEEC while learning, history, and attention mechanisms are not emphasized or not used at all. An interesting aspect of SEEC is that the goal formulation and management is assigned to the application. The SEEC platform provides knobs, control algorithms, and measurements to the application, which in turn is responsible to formulate and adjust its goals. Similarly, the desirability scale, as it is related to and dependent on goals, is not part of the SEEC framework.

HAMSoC. With hierarchical agent monitoring SoC (HAMSoC), Guang et al. [74], [75] have proposed a four-level hierarchical control structure, as illustrated in Figure 6. Each cell agent monitors and controls a core, an accelerator, or another functional hardware block, which is similar to the functional elements and autonomic elements in ASoC. Cell agents have only local knowledge but are in turn monitored and steered by cluster agents, which

pursue optimizations for each respective local cluster. The platform agent is responsible for the entire SoC platform and can pursue platform-global optimizations. It interacts with the application agent that provides application specific goals and requirements, based on the heartbeat concept of SEEC [69]. In contrast to ASoC, which has only one level of controllers communicating with each other, HAMSoC proposes a hierarchy of controllers that each is host to different objectives from the local to the application level, as exemplified in a power and resource management scenario [76].

Even though the HAMSoC framework of hierarchical control has the potential to accommodate most of the self-awareness properties of our paradigmatic architecture of Figure 1, it has not been fully elaborated and exploited. The HAMSoC controller hierarchy would be an appealing match to a hierarchical goal management system. However, neither goal management nor attention, history, or learning mechanisms have been explored.

CPSoC. Cyber Physical SoC (CPSoC) [77] is a self-aware embedded system paradigm that enhances traditional MPSoCs with a sensor-actuator-rich platform deploying a closed loop paradigm emulating large-scale cyber-physical systems, enhanced with smartness through adaptivity and limited self-awareness [78]. CPSoC was developed primarily in the context of managing and exploiting hardware variability using the under-designed and opportunistic computing paradigm [79].

The high-level system architecture of CPSoC is shown in Figure 7. The middle of this figure shows various levels of abstraction for the CPSoC platform, from the lowest (device/circuit level) layering up to the highest (application level). At each abstraction level, the CPSoC platform gathers information (through sensor fusion) using virtual and physical sensors, and in turn actuates (through actuator fusion) via virtual and physical actuators. The CPSoC architecture supports two classes of feedback loops: adaptive control (red box in Figure 7) and self-aware supervisory management that generates supervisory policies (tan box in Figure 7). These feedback loops are embedded within the adaptive, reflective middleware that orchestrates cross-layer sensing and actuation.

Figure 8 shows a more detailed view of the CPSoC architecture. On the top right of the figure is a template of an individual CPSoC computational Core, comprised of the computational units, memories, interfaces, and the on-chip sensing and actuations (OCSA) block that allows ubiquitous sensing and actuation at the CPSoC-Core level. These CPSoC Cores are tiled into a (homogeneous or heterogeneous) CPSoC computational fabric (lower right of Figure 8), using a network-on-chip (NoC) interconnect. Note that each router box in the NoC is also equipped with a sensing-and-actuation block (colored green) that enables monitoring and actuation at each NoC router. The left side of Figure 8 expands the abstraction layers of Figure 7, showing the CPSoC tiled hardware fabric at the lowest layer,

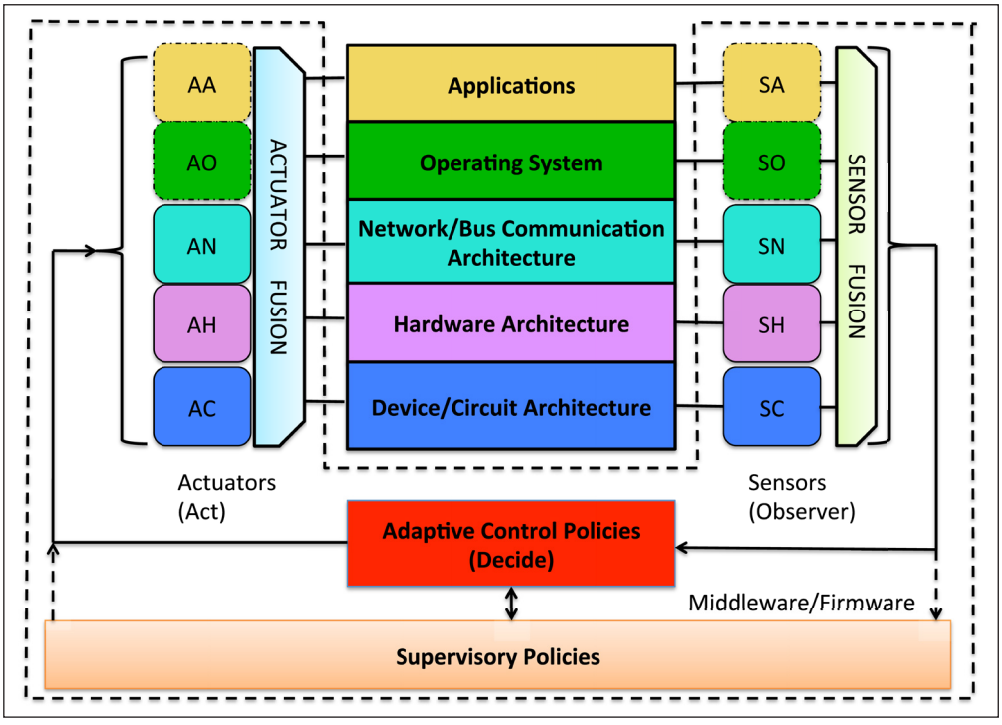

Figure 7. Cross-layer virtual sensing and actuation at different layers of CPSoC [88].

and the applications executing on this platform at the highest layer. The adaptive, reflective middleware layer (yellow box on the left side of Figure 8) orchestrates the distributed sensing and actuation approach, where each component and core can make local decisions to manage the fabric.

The CPSoC architecture achieves self-awareness through three key ideas:

- **Cross-layer virtual and physical sensing & actuation.** CPSoCs are sensor-actuator-rich MPSoCs that include several on-chip physical sensors (e.g., for aging, oxide breakdown, leakage, reliability, and temperature) on the lower three layers as shown by the OCSA block and the introspective sensing units in Figure 8. Virtual sensing and actuation [80] is accomplished across the abstraction stack. For instance, virtual actuations such as application duty cycling and checkpointing are software/hardware interventions that can predictively influence system design objectives. Virtual actuation can be combined with physical actuation mechanisms commonly adopted in modern chips [81].
- **Simple and self-aware adaptations.** Two key attributes of the self-aware CPSoC are adaptation of each layer and multiple cooperative ODA loops. As an example, the unification of an adaptive computing platform (with combined dynamic voltage and frequency scaling, adaptive body biasing, and other actuation means) along with a bandwidth adaptive NoC offers extra dimensions of control and solutions in comparison with traditional MPSoC architecture.
- **Predictive models and on-line learning.** Predictive modeling and on-line learning abilities enhance self-modeling abilities in the CPSoC paradigm. The system behavior and states can be built using on-line or off-line linear or nonlinear models in time or frequency domains [82]. CPSoC's predictive and learning abilities improve autonomy for managing system resources and assisting proactive resource utilization [77].

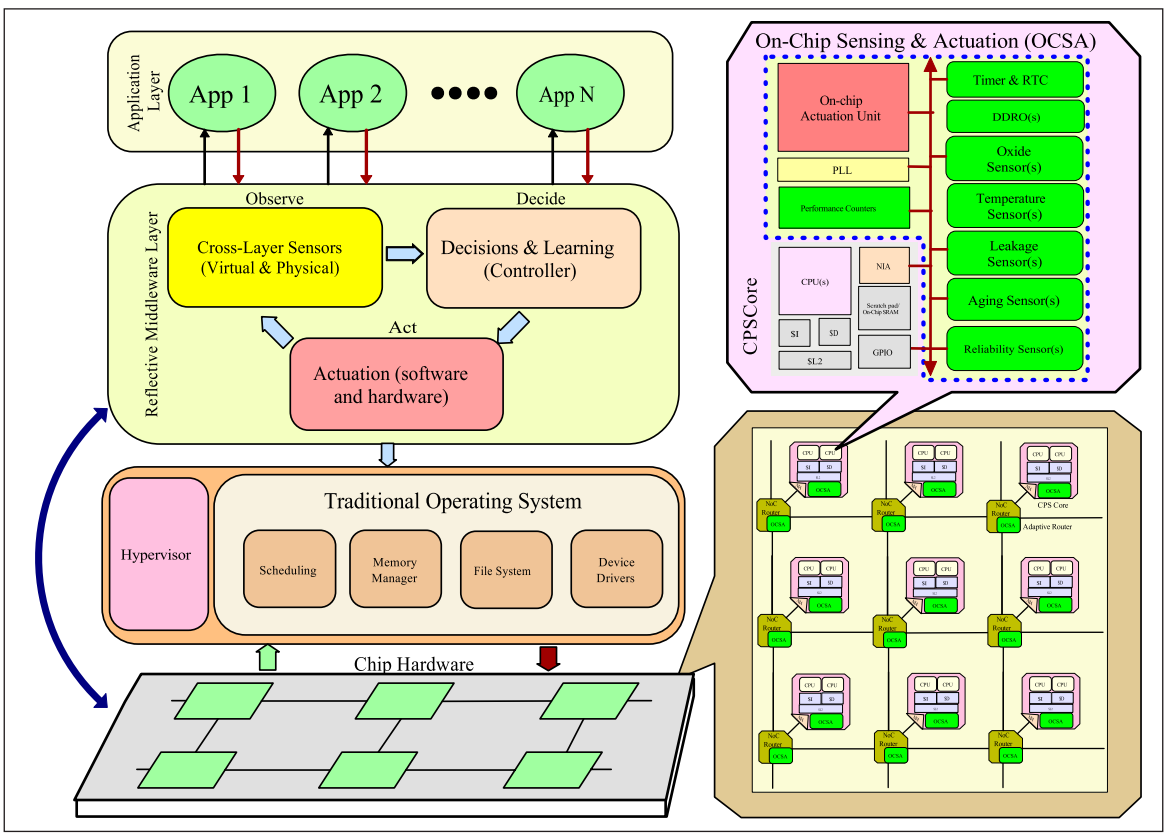

Figure 8. CPSoC architecture with adaptive Core, NoC, and the ODA Loop as Middleware [88].

While CPSoC is a good initial exemplar for a self-aware SoC platform, it handles to a limited extent the self-awareness shown in the paradigmatic architecture of Figure 1. The monitoring-deciding-execution loop is intrinsically part of CPSoC, coupled with some limited learning and history mechanisms. Attention mechanisms have not been considered in CPSoC, and the goal hierarchy and goal management is in a very primitive form. The desirability scale is implicitly encoded within the goals, and has not been explicitly modeled within CPSoC. Furthermore, the self-awareness models in CPSoC did not consider malicious attacks, functional design errors, and nonfunctional aberrations, and show a lot of room for growth in its self-awareness capabilities. Thanks to its modular cross-layer architecture, CPSoC has the potentials to cope with the discussed limitation by providing access to a rich set of cross-layer virtual and physical sensors and actuators, and the capacity to become self-aware in all respects.

Challenges

The term self-awareness encompasses a host of concepts and techniques that together offer great promises to tackle the design, maintenance, and operation of complex, heterogeneous systems that are supposed to be adaptive, autonomous, highly efficient, and always sensible. Even though a significant effort has already been spent in exploring this promise, the more intricate challenges still lie ahead. So far we have focused on picking low hanging fruits by incrementally extending existing architectures and methodologies. However, the research community will make faster progress when we do not exclusively focus on incremental development where each additional feature has to be thoroughly and quantitatively justified by the added value it gives. As this survey shows, self-awareness encompasses a host of concepts and techniques that together facilitate a comprehensive understanding of the system's state and its situation in the world. Picking out individual elements may only result in small gains or none at all. Thus, we recommend to take a step back, try to comprehensively understand self-awareness, what it is, what it consists of, what it is good for, and, based on this understanding, realize it as a whole in cyber-physical computing systems. This approach would be inspired and informed by the widespread presence of self-awareness in animals given that survival of an expensive feature under relentless evolutionary pressure is a strong evidence for its benefit. A case in point is the maintenance of a history. It has hardly been studied in self-aware computing systems and, consequently, there is no strong experimental evidence for its benefits. We still believe it is indispensable for comprehensive self-assessment based on its importance in psychology [83], [84], and the intuitive argument that a comprehensive understanding of the current situation includes the sequence of events and states that historically led to the current situation. Moreover, historical data is required for future learning that involves a reassessment of past situations. Thus, including it in research on self-aware computing systems is justified by the expectation that it will turn out to be beneficial.

Apart from considerations of research strategy, we identify five urgent technical challenges that have to be addressed in order to fully honor the promise of self-awareness: learning, formulation of goals, scalability, ensuring correctness, and an appropriate design methodology.

Learning. For truly self-aware systems, continuous, dynamic learning is indispensable. A major reason for the amazing feats of animals and plants is the relentless learning that goes on all levels from the subcellular organelles to the individual and the community. As Figure 1 indicates, learning is an integral part of many components and functions. Hence, it must be integrated in the sensor and monitoring nodes, in the attention mechanism, the decision making, the goal management, the execution and actuation, and in virtually every part of the system. Learning is only possible when feedback signals are available. Thus, the system must be pervaded by information flows providing feedback to all the learning elements. Many of our machine learning algorithms are not sufficiently efficient and optimized for the requirements of on-chip learning. Hence, we need both adapted machine learning algorithms and a system architecture that lends itself to continuous, pervasive learning, and optimization.

Formulation of goals. We need to be able to formulate quantitative goals for the design and the operation phases and we need to study the involved tradeoffs. The traditional metrics of performance,

power, energy, cost, and fault tolerance are well understood. But quantitative metrics for adaptability, resilience, autonomy, self-assessment, and situation assessment do not exist and are controversial or are limited in scope. However, we need to quantify these properties to explore the tradeoff space spanned by the traditional and the self-awareness metrics. This has to be done for the design phase, but also and even more challenging, for the operation of the system since the system itself has to understand and decide on these tradeoffs in real time. Research on goal formulation and management has been done in the context of artificial agents [85]–[87], which mainly focus on providing the capability to nominate top-level goals and managing the nominated goals by prioritizing them. However, SoC's requirements and restrictions necessitate customized, light-weight, and minimally conflicting approaches which consider the priority, significance, objectives, and requirements of each application, while holistically coupling the overlapping and/or contradicting objectives of different applications to satisfy the system constraints.

Scalable self-awareness. Most applications do not require and cannot afford all features of a full blown self-aware system. To apply self-awareness to a wide range of systems, from resource constrained sensor nodes to multiprocessor platforms, designers should be able to easily select the level of capabilities. To this end, a design space exploration method has to provide the means to trade-off functions and resources in a well-defined self-awareness design space.

Ensuring correctness. Validating a fixed, well defined functionality has been proven difficult enough due to the vast state spaces involved. Validating an adaptive system that, by definition, changes its behavior in ways unpredictable at design time seems to be hopeless. Still, if we cannot guarantee that certain bad behavior can never happen, the appeal of self-aware and autonomous systems will be limited to tiny application domains. Interestingly, self-awareness may be part of the solution because it can comprise a safety monitor that checks for and prohibits all unsafe and bad behavior. For this to work, the space of unsafe and bad behavior has to be specifiable in unambiguous and efficient terms leaving the system to freely explore the vast, unlimited space of safe and good behavior.

Design methodology. Traditional design methodologies rely on the assumption that we can specify, validate, and test the desired and acceptable behavior of the system. When we allow the adaptive, autonomous system to explore behavior that has not been specified at design time, this assumption breaks down. Hence, we have to consider alternative methodologies. For instance, the designers could use a general purpose, self-aware, autonomous machine, that in principle can meet a broad range of goals in any environment. Then the designers "fill" the system with a specific set of goals for a specific application and leave it to the system to find ways to accomplish these goals. Although this vision seems remote, we will be forced to contemplate such options as the pain of designing more and more adaptive systems with traditional methodologies grows.

WHILE THESE CHALLENGES seem formidable, researchers can draw from a range of disciplines with long history and large knowledge. Hence, given the state of the art, as summarized in this survey, we can certainly be confident that the development of fully self-aware SoCs is within the reach of the community in the coming years. ■

Acknowledgments

We acknowledge the financial support by the Marie Curie Actions of the European Union's H2020 Programme, as well as the National Science Foundation under Grant CCF-1704859.

References

[1] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, 2003.

[2] M. Salehie and L. Tahvildari, "Self-adaptive software: Landscape and research challenges," *ACM Trans. Auton. Adaptive Syst.*, vol. 4, no. 2, p. 14, 2009.

[3] M. Salehie and L. Tahvildari, "Autonomic computing: Emerging trends and open problems," in *Proc. ACM SIGSOFT Software Engineering Notes*, ACM, vol. 30, 2005, pp. 1–7.

[4] S. Kounev, "Engineering of self-aware IT systems and services: State-of-the-art and research challenges," in *Proc. 8th European Performance Engineering Workshop*, vol. 6977 of LNCS, Springer, 2011, pp. 10–13.

[5] S. Kounev, X. Zhu, J. O. Kephart, and M. Kwiatkowska, "Model-driven algorithms and architectures for self-aware computing systems (Dagstuhl Seminar 15041)," *Dagstuhl Reports*, vol. 5, no. 1, pp. 164–196, 2015.

[6] A. Chandra, P. R. Lewis, K. Glette, and S. C. Stilkerich, "Reference architecture for self-aware and self-expressive computing systems," *Self-Aware Computing Systems: An Engineering Approach*, P. R. Lewis, M. Platzner, B. Rinner, J. Torresen, and X. Yao, Eds. Springer, 2016, pp. 37–49.

[7] S. Sarma, N. Dutt, P. Gupta, A. Nicolau, and N. Venkatasubramanian, "Cyberphysical-system-on-chip (CPSoC): A self-aware MPSoC paradigm with cross-layer virtual sensing and actuation," in *Proc. Des., Autom. Test in Europe Conf. Exhibition,* Grenoble, France, Mar. 2015.

[8] L. Guang, E. Nigussie, P. Rantala, J. Isoaho, and H. Tenhunen, "Hierarchical agent monitoring design approach towards self-aware parallel systems-on-chip," *ACM Trans. Embed. Comput. Syst.*, vol. 9, no. 3, pp. 1–24, 2010.

[9] H. Hoffmann, "Coadapt: Predictable behavior for accuracy-aware applications running on power-aware systems," in *Proc. 2014 26th Euromicro Conf. Real-Time Systems*, Jul. 2014, pp. 223–232.

[10] L. Esterle, J. Simonjan, G. Nebehay, R. Pflugfelder, G. F. Dominguez, and B. Rinner, "Self-aware object tracking in multi-camera networks," *Self-Aware Computing Systems: An Engineering Approach*, P. R. Lewis, M. Platzner, B. Rinner, J. Torresen, and X. Yao, Eds. Springer, 2016, pp. 261–277.

[11] M. Kurek et al., "Self-aware hardware acceleration of financial applications on a heterogeneous cluster," *Self-Aware Computing Systems: An Engineering Approach*, P. R. Lewis, M. Platzner, B. Rinner, J. Torresen, and X. Yao, Eds. Springer, 2016, pp. 241–260.

[12] H. Giesen, B. Gojman, R. Rubin, J. Kim, and A. De-Hon, "Continuous online self-monitoring introspection circuitry for timing repair by incremental partial-reconfiguration (cosmic trip)," in *Proc. 2016 IEEE 2nd Annu. Int. Symp. Field-Program. Custom Comput. Machines*, 2016, pp. 111–118.

[13] S. Sarma and N. Dutt, "Cross-layer exploration of heterogeneous multicore processor configurations," in *Proc. VLSI Des. Conf.*, 2015.

[14] S. Sarma, N. Dutt, P. Gupta, A. Nicolau, and N. Venkatasubramanian, "On-chip self-awareness using cyberphysical-systems-on-chip (CPSoC)," in *Proc. 12th Int. Conf. Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, New Delhi, India, Oct. 2014.

[15] S. Sarma, T. Muck, L. A. D. Bathen, N. Dutt, and A. Nicolau, "Smart-balance: A sensing-driven Linux load balancer for energy efficiency of heterogeneous MPSoCs, in *Proc. Des. Autom. Conf.*, Jul. 2015.

[16] M. D. Santambrogio, H. Hoffmann, J. Eastep, and A. Agarwal, "Enabling technologies for self-aware adaptive systems," in *Proc. NASA/ESA Conf. Adaptive Hardware and Systems*, 2010, pp. 149–156.

[17] M. Happe and A. Trammel-Keller, "Flexible protocol stacks," *Self-Aware Computing Systems: An Engineering Approach*, P. R. Lewis, M. Platzner, B. Rinner, J. Torresen, and X. Yao, Eds. Springer, 2016, pp. 193–214.

[18] C. Krupitzer, F. M. Roth, S. VanSyckel, G. Schiele, and C. Becker, "A survey on engineering approaches for self-adaptive systems," *Pervasive and Mobile Comput.*, vol. 17, pp. 184–206, 2015.

[19] P. R. Lewis et al., "Architectural aspects of self-aware and self-expressive computing systems," *IEEE Comp.*, vol. 48, no. 8, pp. 62–70, Aug. 2015.

[20] N. Dutt, A. Jantsch, and S. Sarma, "Self-aware cyber-physical systems-on-chip," in *Proc. Int. Conf. Comp. Aided Des.*, Austin, TX, USA, Nov. 2015.

[21] N. Dutt, A. Jantsch, and S. Sarma, "Towards smart embedded systems: A self-aware system-on-chip perspective," *ACM Trans. Embedded Comput. Syst.*, vol. 15, no. 2, pp. 22:1–22:27, May 2016.

[22] A. Jantsch and K. Tammemae, "A framework of awareness for artificial subjects," in *Proc. 2014 Int. Conf. Hardware/Software Codesign and System Synthesis*, New York, NY, USA, 2014, pp. 20:1–20:3.

[23] N. TaheriNejad, A. Jantsch, and D. Pollreisz, "Comprehensive observation and its role in self-awareness – An emotion recognition system example," in *Proc. Federated Conf. Comp. Science Inf. Syst.*, Gdansk, Poland, Sept. 2016.

[24] J. Kephart, "Autonomic computing: The first decade," in *Proc. Int. Conf. Auton. Comput.*, 2011.

[25] IBM Corporation, "An architectural blueprint for autonomic computing," IBM White Paper, 2006.

[26] R. Sanz, I. Lopez, M. Rdorlguez, and C. Hernandez, "Principles for consciousness in integrated cognitive control," *Neural Networks*, vol. 20, no. 9, p. 11, 2007.

[27] D. Ghosh, R. Sharman, H. Raghav Rao, and S. Upadhyaya, "Self-healing systems – survey and

synthesis," *Decis. Support Syst.*, vol. 42, no. 4, pp. 2164–2185, Jan. 2007.

[28] H. Psaier and S. Dustdar, "A survey on self-healing systems: Approaches and systems," *Computing*, vol. 91, no. 1, pp. 43–73, 2011.

[29] Amazon Web Services, *Amazon Auto Scaling*. Accessed: February 9, 2017. [Online]. Available: http://aws.amazon.com/documentation/autoscaling

[30] Microsoft Developer Network, *Autoscaling and Windows Azure*. Accessed: February 9, 2017. [Online]. Available: http://msdn.microsoft.com/en-us/library/hh680945(v=pandp.50).aspx

[31] B. Morin, O. Barais, J. Jezequel, F. Fleurey, and A. Solberg, "Models runtime to support dynamic adaptation," *Computer*, vol. 42, pp. 44–51, 2009.

[32] K. Geihs et al., "A comprehensive solution for application-level adaptation," *Software Prac. Exp.*, vol. 39, pp. 385–422, 2009.

[33] S. Hallsteinsen et al., "A development framework and methodology for self-adapting applications in ubiquitous computing environments," *J. Syst. Software*, vol. 85, pp. 2840–2859, 2012.

[34] D. Menasce, H. Gomaa, S. Malek, and J. Sousa, "SASSY: A framework for self-architecting service-oriented systems," *IEEE Software*, vol. 28, no. 6, pp. 78–85, Nov. 2011.

[35] S. Kounev, "Performance modeling and evaluation of distributed component-based systems using queueing petri nets," *IEEE Trans. Software Engineering*, vol. 32, no. 7, pp. 486–502, 2006.

[36] S. Gilmore, V. Haenel, L. Kloul, and M. Maidl, "Choreographing security and performance analysis for web services," in *Proc. Formal Techniques for Comp. Syst. Business Processes*, 2005, pp. 200–214.

[37] E. Eskenazi, A. Fioukov, and D. Hammer, "Performance prediction for component compositions," in *Proc. Int. Symp. Component-Based Software Engineering*, 2004, pp. 208–293.

[38] A. Gambi, G. Toffetti, C. Pautasso, and M. Pezze, "Kriging controllers for cloud applications," *IEEE Internet Comput.*, vol. 17, no. 4, pp. 40–47, 2013.

[39] M. Becker, M. Luckey, and S. Becker, "Model-driven performance engineering of self-adaptive systems: A survey," in *Proc. ACM SIGSOFT Int. Conf. Quality Software Architectures*, 2012, pp. 117–122.

[40] H. Koziolek, "Performance evaluation of component-based software systems: A survey," *Performance Evaluation*, vol. 67, no. 8, pp. 434–458, 2010.

[41] M. C. Huebscher and J. A. McCann, "A survey of autonomic computting—Degrees, models, and applications," *ACM Comput. Surveys*, vol. 40, no. 3, pp. 7:1–7:28, Aug. 2008.

[42] Software Engineering Institute (Carnegie Mellon University). *Software Product Lines*. Accessed: July 14, 2017. [Online]. Available: http://www.sei.cmu.edu/productlines/

[43] K. Geihs, R. Reichle, M. Wagner, and M. U. Khan, *Modeling of Context-Aware Self-Adaptive Applications in Ubiquitous and Service-Oriented Environments*. Berlin, Germany: Springer, 2009, pp. 146–163.

[44] H. Gomaa and M. Hussein, *Dynamic Software Reconfiguration in Software Product Families*. Berlin, Germany: Springer, 2004, pp. 435–444.

[45] S. Hallsteinsen, M. Hinchey, S. Park, and K. Schmid, "Dynamic software product lines," *Computer*, vol. 41, no. 4, pp. 93–95, Apr. 2008.

[46] B. Morin, O. Barais, G. Nain, and J. M. Jezequel, "Taming dynamically adaptive systems using models and aspects," in *Proc. 2009 IEEE 31st Int. Conf. Software Engineering*, May 2009, pp. 122–132.

[47] L. Gurgen, O. Gunalp, Y. Benazzouz, and M. Gallissot, "Self-aware cyber-physical systems and applications in smart buildings and cities," in *Proc. Conf. Des. Autom. Test in Europe*, EDA Consortium, San Jose, CA, USA, 2013, pp. 1149–1154.

[48] H. Schmeck, "Organic computing – A new vision for distributed embedded systems," in *Proc. 8th IEEE Int. Symp. Object-Oriented Real-Time Distributed Comput.*, May 2005, pp. 201–203.

[49] C. Müller-Schloer, H. Schmeck, and T. Ungerer, Eds. *Organic Computing: A Paradigm Shift for Complex Systems*. Basel, Switzerland: Birkhäuser, 2011.

[50] D. Kramer, R. Buchty, and W. Karl, "Monitoring and self-awareness for heterogeneous, adaptive computing systems," *Organic Computing – A Paradigm Shift for Complex Systems, Autonomic Systems*, C. Müller-Schloer, H. Schmeck, and T. Ungerer, Eds. Basel, Switzerland: Birkhäuser, 2011, pp. 163–178.

[51] J. H. Holland, "Adaptation," *Progress in Theoretical Biology*, R. Rosen and F. M. Snell, Eds. New York, NY, USA: Academic Press, 1976, pp. 263–293.

[52] M. Butz and S. W. Wilson, "An algorithmic description of XCS," *IWLCS'00, Lecture Notes in Artifical Intelligence*, vol. 2321, P. L. Lanzi, W. Stolzmann, and S. W. Wilson, Eds. Springer, 2001, pp. 253–272.

[53] A. Bernauer, J. Zeppenfeld, O. Bringmann, A. Herkersdorf, and W. Rosenstiel, "Combining software and hardware LS for lightweight on-chip learning," *Organic Computing – A Paradigm Shift for Complex Systems*, *Autonomic Systems*, C. Müller-Schloer, H. Schmeck, and T. Ungerer, Eds. Basel, Switzerland: Birkhäuser, 2011, pp. 253–266.

[54] T. Ebi et al., "DodOrg – A self-adaptive organic many-core architecture," *Organic Computing – A Paradigm Shift for Complex Systems, Autonomic Systems*, C. Müller-Schloer, H. Schmeck, and T. Ungerer, Eds. Basel, Switzerland: Birkhäuser, 2011, pp. 353–368.

[55] A. M. Rahmani, P. Liljeberg, A. Hemani, A. Jantsch, and H. Tenhunen, *The Dark Side of Silicon*, 1st ed. Cham, Switzerland: Springer, 2016.

[56] L. Benini and G. DeMicheli, *Dynamic Power Management: Design Techniques and CAD Tools.* Springer Verlag, 1998.

[57] P. Bogdan, "Mathematical modeling and control of multifractal workloads for data-center-on-a-chip optimization," in *Proc. 9th Int. Symp. Networks-on-Chip*, NOCS, ACM, New York, NY, USA, 2015, pp. 21:1–21:8.

[58] P. Bogdan, R. Marculescu, and S. Jain, "Dynamic power management for multidomain system-on-chip platforms: An optimal control approach," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 18, no. 4, pp. 46:1–46:20, Oct. 2013.

[59] R. David, P. Bogdan, and R. Marculescu, "Dynamic power management for multicores: Case study using the intel SCC," in *Proc. 2012 IEEE/IFIP 20th Int. Conf. VLSI System-on-Chip*, Oct. 2012, pp. 147–152.

[60] T. S. Muthukaruppan, M. Pricopi, V. Venkataramani, T. Mitra, and S. Vishin, "Hierarchical power management for asymmetric multi-core in dark silicon era," in *Proc. 50th Annu. Des. Autom. Conf.*, 2013, pp. 174:1–174:9.

[61] A. M. Rahmani et al., "Dynamic power management for many-core platforms in the dark silicon era: A multi-objective control approach," in *Proc. IEEE/ACM Int. Symp. Low-Power Electron. Des.*, 2015, pp. 219–224.

[62] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, and X. Zhu, "No "Power" struggles: Coordinated multi-level power management for the data center," in *ASPLOS*, 2008, pp. 48–59.

[63] H. Zhang and H. Hoffmann, "Maximizing performance under a power cap: A comparison of hardware, software, and hybrid techniques," in *Proc. Int. Conf. Architectural Support for Programming Languages and Operating Syst.*, 2016, pp. 545–559.

[64] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design.* John Wiley & Sons, 2005.

[65] R. P. Pothukuchi, A. Ansari, P. Voulgaris, and J. Torrellas, "Using multiple input, multiple output formal control to maximize resource efficiency in architectures," in *Proc. ISCA*, 2016, pp. 658–670.

[66] A. Bouajila et al., "Organic computing at the system on chip level," in *Proc. 2006 IFIP Int. Conf. Very Large Scale Integr.*, Oct. 2006, pp. 338–341.

[67] A. Bouajila et al., "Autonomic system on chip platform," *Organic Computing – A Paradigm Shift for Complex Systems, Autonomic Systems*, C. Müller-Schloer, H. Schmeck, and T. Ungerer, Eds. Basel, Switzerland: Birkhäuser, 2011, pp. 413–425.

[68] H. Hoffmann, M. Maggio, M. D. Santambrogio, A. Leva, and A. Agarwal, *SEEC: A Framework for Self-Aware Computing*, Tech. Rep. MIT-CSAIL-TR-2010-049, MIT, Cambrige, MA, USA, Oct. 2010.

[69] H. Hoffmann, J. Eastep, M. D. Santambrogio, J. E. Miller, and A. Agarwal, "Application heartbeats: A generic interface for specifying program performance and goals in autonomous computing environments," in *Proc. 7th Int. Conf. Autonomic Comput.*, ACM, 2010, pp. 79–88.

[70] H. Hoffmann, M. Maggio, M. D. Santambrogio, A. Leva, and A. Agarwal, "A generalized software framework for accurate and efficient management of performance goals," in *2013 Proc. Int. Conf. Embedded Software*, Sept. 2013, pp. 1–10.

[71] H. Hoffmann, S. Sidiroglou, M. Carbin, S. Misailovic, A. Agarwal, and M. Rinard, "Dynamic knobs for responsive power-aware computing," *ACM SIGPLAN Notices*, vol. 46, no. 3, pp. 199–212, 2011.

[72] M. Maggio, H. Hoffmann, M. D. Santambrogio, A. Agarwal, and A. Leva, "Decision making in autonomic computing systems: Comparison of approaches and techniques," in *Proc. 8th ACM Int. Conf. Auton. Comput.*, ACM, New York, NY, USA, 2011, pp. 201–204.73.

[73] M. D. Santambrogio, H. Hoffmann, J. Eastep, and A. Agarwal, "Enabling technologies for self-aware adaptive systems," in *Proc. 2010 NASA/ESA Conf. Adaptive Hardware and Systems*, IEEE, 2010, pp. 149–156.

[74] L. Guang, G. Plosila, J. Isoaho, and H. Tenhunen, "HAMSoC: A monitoring-centric design approach for adaptive parallel computing," *Autonomic Networking-on-Chip: Bio-Inspired Specification, Development, and Verification*, P. Cong-Vinh, Ed., CRC Press, Dec. 2011, pp. 135–164.

[75] L. Guang, J. Plosila, J. Isoaho, and H. Tenhunen, "Hierarchical agent monitored parallel on-chip system: A novel design paradigm and its formal specification," *Int. J. Embedded Real-Time Commun. Syst.*, vol. 1, no. 2, pp. 86–105, 2010.

[76] S. M. A. H. Jafri et al., "Self-adaptive NoC power management with dual-level agents: Architecture and implementation," in *Proc. Conf. Self-Adaptive Networked Embedded Systems*, Rome, Italy, Feb. 2012.

[77] S. Sarma, N. Dutt, P. Gupta, N. Venkatasubramanian, and A. Nicolau, "Cyberphysical-system-on-chip (CPSoC): A self-aware MPSoC paradigm with cross-layer virtual sensing and actuation," in *Proc. Des. Autom. Test in Europe Conf.*, 2015, pp. 625–628.

[78] N. Dutt, A. Jantsch, and S. Sarma, "Toward smart embedded systems: A self-aware system-on-chip (SoC) perspective," *ACM Trans. Embed. Comput. Syst.*, vol. 15, no. 2, pp. 22:1–22:27, 2016.

[79] P. Gupta et al., "Underdesigned and opportunistic computing in presence of hardware variability," *IEEE Trans. Comp.-Aided Des. Integr. Circuits Syst.*, vol. 32, no. 1, pp. 8–23, 2013.

[80] S. Sarma, N. Dutt, and N. Venkatasubramanian, "Cross-layer virtual observers for embedded multiprocessor system-on-chip (mpsoc)," in *Proc. 11th Int. Workshop Adaptive and Reflective Middleware*, ARM, New York, NY, USA, 2012, pp. 4:1–4:7.

[81] S. Sarma, N. Dutt, and P. Gupta, *Strength of Diversity: Heterogeneous Noisy Sensor Allocation and Placement for Optimal Reconstruction and Accurate Fullchip Thermal Estimation*, Tech. Rep., University of California Irvine, Irvine, CA, USA, 2014.

[82] L. Ljung, Ed., *System Identification (2nd Ed.): Theory for the User*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1999.

[83] B. J. Baars, *A Cognitive Theory of Consciousness.* Cambridge University Press, 1989.

[84] A. D. Baddeley and G. Hitch, "Working memory," *The Psychology of Learning and Motivation*, vol. 8. Elsevier, 1974, pp. 48–89.

[85] D. Choi, "Reactive goal management in a cognitive architecture," *Cogn. Syst. Res.*, vol. 12, no. 3–4, pp. 293–308, 2011.

[86] U. Jaidee, H. M. Avila, and D. W. Aha, "Integrated learning for goal-driven autonomy," in *Proc. Twenty-Second Int. Joint Conf. Artificial Intelligence – Volume Volume Three*, 2011, pp. 2450–2455.

[87] M. Wilson, M. Molineaux, and D. W. Aha, "Domain-independent heuristics for goal formulation," in *Proc. 26th Artificial Intelligence Research Society Conf.*, 2013, pp. 160–165.

[88] S. Sarma, N. Dutt, N. Venkatasubramanian, A. Nicolau, and P. Gupta, *Cyber Physical-System-On-Chip (CPSoC): Sensor-Actuator Rich Self-Aware Computational Platform*, Tech. Rep., University of California Irvine, Irvine, CA, USA, 2013.

Axel Jantsch is a professor in the Institute of Computer Technology at TU Wien, Austria. His research interests include dependability, robustness, and self-awareness in SoCs and embedded systems. Jantsch has a PhD in computer science from Vienna University of Technology. He is a member of the IEEE.

Nikil Dutt is a Chancellor's Professor of CS, Cognitive Sciences, and EECS at the University of California, Irvine. His research interests include embedded systems, EDA, computer systems architecture and software, and brain-inspired computing. He is an ACM Fellow, IEEE Fellow, and IFIP Silver Core awardee.

Amir M. Rahmani is a Marie Curie Global Fellow at UC Irvine and TU Wien, and a docent at the University of Turku. His research interests span embedded parallel and distributed computing, Internet-of-Things, and e-health. Rahmani has a PhD in ICT from University of Turku. He is a senior member of the IEEE.

■ Direct questions and comments about this article to Axel Jantsch, TU Wien, 1040 Wien, Austria; e-mail: axel.jantsch@tuwien.ac.at.