Editorial: Special Issue on Innovative Design Methods for Smart Embedded Systems

Nowadays, smart systems are considered the solution to a wide range of challenges, from globalization to climate change. As a result, they are becoming more relevant in a large number of critical sectors, like energy management in public spaces, healthcare, automotive, safety and security.

Compared to classical embedded systems, a distinctive aspect for these systems is their *smartness*, which is the ability to learn from the previous experience and to seemingly react to the surrounding environment. However, this tight interaction with the physical environment implies a high level of heterogeneity in the hardware architecture. At the same time, application scenarios are becoming more complex, since an increasing amount of computation is constrained by tight performance, cost and safety requirements. As a result, it becomes mandatory to properly co-optimize architectures, learning methods, and their specific application with respect to the problem under analysis, especially when the systems are used in safety-critical environments or when integrating emerging technologies.

Novel and comprehensive methodologies are thus required to ease the development of next-generation smart systems, with the goal of reducing design costs and time-to-market. On one hand, it becomes necessary to investigate innovative machine-learning and artificial intelligence techniques to improve the *smartness* of the architectures. On the other hand, architectures must be revised to include infrastructural innovations to efficiently and effectively support these techniques. Finally, the integration of these aspects has to be carefully validated as, in this context, it is increasingly necessary to analyze the dynamic behavior of the resulting systems and their reactivity to unpredictable events (e.g., component failures or unexpected environmental changes).

This special issue, entitled "Innovative Design Methods for Smart Embedded Systems," tackles such challenges by providing both machine-learning techniques and application-specific optimization solutions that guarantee that the application of smart innovations meets the imposed requirements and constraints. The high number of submissions allowed to select seven high-quality articles. We hope you will enjoy them and find them as interesting as we did.

The first article, "Toward Smart Embedded Systems: A Self-Aware System-on-Chip (SoC) Perspective," by Nikil Dutt, Axel Jantsch, and Santanu Sarma, provides open challenges and research directions in the context of the ability of smart systems to be self-aware, to evolve, and to achieve a high level of resilience in the face of highly dynamic and unpredictable environments. To this extent, the article reviews the notions of self-awareness, self-adaptivity, and autonomic systems, and it provides a taxonomy of levels of awareness of typical embedded systems. The article then adopts a system-on-chip perspective to show how cyber-physical system-on-chip platforms may achieve self-awareness through a combination of cross-layer sensing, actuation, self-aware adaptations, and online learning.

The second article, "Adaptive and Hierarchical Runtime Manager for Energy-Aware Thermal Management of Embedded Systems," by Anup Das, Bashir M. Al-Hashimi, and Geoff V. Merrett, proposes a learning-based runtime manager for energy-aware thermal management, with the goal of reducing device wear-out, energy consumption

1539-9087/2016/02-ART22e $15.00
DOI: http://dx.doi.org/10.1145/2885505

caused by leakage power, which is critical especially for battery-operated smart systems. To confine the learning space and to minimize the learning cost, the proposed runtime manager combines a heuristic-based thread allocation at a longer time interval (to improve thermal cycling) and a learning-based hardware frequency selection at a much finer interval (to improve average temperature, peak temperature, and energy consumption).

The third article, “Learning Hardware Friendly Classifiers through Algorithmic Stability,” by Luca Oneto, Sandro Ridella, and Davide Anguita, proposes a new learning framework, i.e., the algorithmic-risk-minimization, which enriches the learning process by considering computational and resource constraints, such as limited depth of the arithmetic unit, memory availability, or battery capacity. The proposed framework trains advanced resource-sparing machine-learning models and efficiently deploys them on smart embedded systems, with a case study analysis on a smartphone-based human activity recognition application.

The fourth article, “Real-Time Reachability for Verified Simplex Design,” by Taylor T. Johnson, Stanley Bak, Marco Caccamo, and Lui Sha, builds systems with formally verified behavior, despite the use of unverified and complex control logic, based on machine-learning techniques. The proposed approach combines offline linear matrix inequality and reachability methods along with a real-time reachability computation. This allows to maintain safety but with significantly less conservatism, thus allowing smart controllers to be used more frequently.

The fifth article, “A Real-Time FPGA-Based Accelerator for ECG Analysis and Diagnosis Using Association Rules Mining,” by Xiaoqi Gu, Yongxin Zhu, Shengyan Zhou, Chaojun Wang, Meikang Qui, and Guoxing Wang, presents a streaming architecture implemented on field-programmable gate array (FPGAs), designed for accelerating real-time ECG signal analysis and diagnosis in a pipelining and parallel way. The FPGA-based implementation exploits pipelining operations in hardware to accelerate a quantitative mining algorithm and to achieve better performance and scalability as well as less hardware cost.

The sixth article, “RQNoC: A Resilient QoS NoC with Service Redirection” by Alirad Malek, Ioannis Sourdis, Stavros Tzilis, Yifan He, and Gerard Rauwerda, describes a service-oriented network-on-chip resilient to permanent faults. The authors categorize network resources based on the service they support and provide mechanisms to bypass them when faulty, allowing a traffic class to be redirected. Traffic redirection is performed either by modifying its routing (Service Detour) or by using a router data path dedicated to a different service (SMerge).

The seventh article, “Near-Static Shading Exploration of PV Modules”, by Marioa-Iro Baka, Francky Catthoor, and Dimitrios Soudris, moves the focus to smart building and smart energy environments by proposing an alternative configurable topology for intramodule cell interconnection of photovoltaic panels. In the proposed solution, cell connections are adaptive to reduce the impact of variable shading and to cluster cells with similar I-V characteristics over time. Experimental evidence proves that the proposed module architecture outperforms significantly a conventional 10×6 module under heavy shade.

Sara Vinco
Politecnico di Torino, Italy
Christian Pilato
Columbia University, NY, USA

Guest Editors