SPECIAL ISSUE ON HARDWARE IMPLEMENTATIONS OF SOFT COMPUTING TECHNIQUES

Anguita, D.; Baturone, I.; Miller, J.

Soft computing (SC) techniques, which include methods from fuzzy logic, neural networks, evolutionary computing, and other similar areas, have proven in the last few years how tolerance for imprecision, uncertainty, or approximation can be exploited to achieve tractable, robust, and low cost solutions. Most of SC methods and algorithms have been implemented in software on common computer systems (PCs or workstations); however, the recent interest on pervasive, ubiquitous, and real-time computing, smart adaptive solutions, embedded intelligent systems and devices, has raised sharply the need for efficient hardware implementations to meet restrictive requirements concerning miniaturization, low power consumption, and/or high speed operation.

This special issue targets new developments and state-of-the-art case studies of hardware implementations of soft computing techniques, including neural networks, fuzzy systems and evolutionary computation or their combination. The seven selected papers cover a wide range of approaches: at the higher level they target current-generation microprocessors or digital signal processors (DSPs), while at the lower level they deal with the design complexities of microelectronics and special-purpose VLSI devices.
Following this virtual path, the first article, by Frias-Martinez et al., proposes a fuzzy compiler, which targets single instruction multiple data (SIMD) architectures, like the ones found in modern microprocessors or DSPs. The objective of the fuzzy compiler is to translate the fuzzy system syntax to embedded software so as to obtain the maximum performance from standard architectures. The Intel Pentium III and the Texas Instruments DSP C6x are the analyzed platforms.

The two papers from Cabrera et al. and Reyneri focus on HW/SW co-design issues as a way to exploit the best of hardware and software approaches. They both employ hardware platforms based on processors and field programmable gate arrays (FPGAs), whose performance/cost ratio has been highly increased in the last few years, and adequate CAD tools to automate the design process. The first paper addresses fuzzy control systems analyzing two implementation approaches: an off-the-shelf microcontroller with a medium complex FPGA, and the whole system implemented in a more powerful FPGA as a system on a programmable chip (SoPC). The second paper describes neuro-fuzzy systems and shows how the combination of software and hardware can be exploited for targeting special-purpose digital architectures and for obtaining more cost-effective implementations than software approaches.

Scheuermann et al. present a FPGA implementation of an ant colony optimization algorithm (ACO), which is able to find good solutions of combinatorial optimization problems. The standard ACO algorithm is modified appropriately to suit the resources available in current FPGA architectures. Test results show that the FPGA approach achieves a significant speed-up respect to software implementations.
Wojtasik et al. describe a Fuzzy controller for an implantable medical device: a rate-adaptive heart pacemaker, whose main constraint is an extremely low power consumption. The paper discusses how it can be implemented in hardware, using several approaches: from general-purpose microprocessors (of the ARM family) to dedicated devices based on mixed-mode and fully digital custom VLSI chips.

Salmeri et al. exploit the capability of fuzzy systems to approximate functions, in particular sinusoidal functions, which have been addressed by several classical paradigms. The objective is to implement a digital direct synthesizer, which is used in modern communication and measurement devices. The paper describes the algorithm employed to optimize the design, taking into account the effects of the fixed point implementation, and the requirements in both time and frequency domains.

Finally, the paper from Valle et al. is a step towards implementing smart adaptive systems on silicon. It surveys the current state-of-the-art of hardware technologies at the microelectronics level to implement Neural Networks with on-chip learning algorithms. The work describes the implementation of a Weight Perturbation on-chip learning network using analog VLSI architectures to achieve very low power consumption and high computational density.

We thank all the authors for allowing us to bring together so many and such diverse visions on this subject subject and all the reviewers for their timely and accurate work.

We believe that this special issue will be of interest to a wide range of readers, who are willing to solve real applications with soft computing techniques.