Editorial

Design and Architectures for Signal and Image Processing

Markus Rupp (EURASIP Member),[1] Dragomir Milojevic,[2] and Guy Gogniat[3]

[1] *Institute of Communications and Radio-Frequency Engineering (INTHFT), Technical University of Vienna, 1040 Vienna, Austria*
[2] *BEAMS, Université Libre de Bruxelles, CP165/56, 1050 Bruxelles, Belgium*
[3] *Lab-STICC Laboratory, University of South Brittany, CNRS, UMR 3192, 56321 Lorient, France*

Correspondence should be addressed to Markus Rupp, mrupp@nt.tuwien.ac.at

The development of complex applications involving signal, image, and control processing is classically divided into three consecutive steps: a theoretical study of the algorithms, a study of the target architecture, and finally an implementation.

Today, such sequential design flow is reaching its limits for the following reasons.

(i) The complexity of today's systems designed with the emerging submicron technologies for integrated circuit manufacturing.

(ii) The intense pressure on the design cycle time in order to reach shorter time-to-market, and to reduce development and production costs.

(iii) The strict performance constraints that have to be reached in the end, typically low and/or guaranteed application execution time, integrated circuit area, overall system power dissipation.

An alternative approach to a traditional design flow, called algorithm-architecture matching, aims to leverage the design flow by a simultaneous study of both algorithmic and architectural issues, taking into account multiple design constraints, as well as algorithm and architecture optimizations, not only in the beginning but all the way throughout the design process.

Introducing such design methodology is also necessary when facing the new emerging applications such as high-performance, low-power, low-cost mobile communication systems and/or smart sensors-based systems.

This design methodology will have to face also future architectures based on multiple processor cores and dedicated coprocessors to achieve the required efficiency. NoC-based communications will become also mandatory for many applications to enable parallel interconnections and communication throughputs. Adaptive and reconfigurable architectures represent a new computation paradigm whose trend is clearly increasing.

This forms a driving force for the future evolution of embedded system design methodologies.

This special issue of the EURASIP Journal of Embedded Systems is intended to present innovative methods, tools, design methodologies, and frameworks for algorithm-architecture-matching approach in the design flow including system level design and hardware/software codesign, real-time operating system, system modelling and rapid prototyping, system synthesis, design verification, as well as performance analysis and estimation.

We received 24 submissions for this special issue of which we finally selected 11 for publication.

In the paper entitled "Flexible Hardware-Based Stereo Matching" Kristian Ambrosch et al. propose a novel technique for implementing a flexible block size, disparity range, and frame rate for hardware-based embedded adaptive stereo-vision systems. By reusing existing resources of a static architecture, rather than dynamic reconfiguration, the proposed technique allows both ASIC and FPGA implementations. Using the proposed architecture, the authors show the impact of the flexible stereo matching on the generated disparity maps for the sum of absolute differences (SADs), rank, and census transform algorithms. Finally, the authors quantify the resource usage and achievable performance when synthesized for an Altera Stratix II FPGA.

Back-projection (BP) is a costly computational step in tomography image reconstruction such as positron emission tomography (PET). To reduce the computation time, the paper entitled "High Speed 3D Tomography on CPU, GPU and FPGA" by Nicolas Gac et al. proposes pipelined,

prefetch, and parallelized architecture for PET BP (3PA-PET). The key feature of the proposed architecture is in the original memory access strategy, masking the high latency of the external memory by an efficient use of the intrinsic temporal and spatial locality of the BP algorithm. Proposed architecture is prototyped on a System on Programmable Chip (SoPC) to validate the system and to measure its performances. Time performances are then compared with a desktop PC, a workstation, and a graphic processor unit (GPU).

The paper entitled "A SIMD Programmable Vision Chip with High Speed Focal Plane Image Processing" by Dominique Ginhac et al. describes a high-speed analogue VLSI image acquisition and low-level image processing system based on dynamically reconfigurable SIMD processor array. The chip features a massively parallel architecture enabling the computation of programmable mask-based image processing for each pixel. A 64×64 pixel proof-of-concept chip, built in $35\,\mu$m standard CMOS process, with a pixel size of $35 \times 35\,\mu$m, is shown. A dedicated embedded platform including FPGA and ADCs has also been designed to evaluate the vision chip. The chip can capture up to 10 000 and process up to 5000 images per second.

The paper entitled "Design of a Real-time Face Detection Parallel Architecture Using High-Level Synthesis" by Nicolas Farrugia et al. describes an architecture specified using C language and synthesized using a high-level synthesis tool. Such approach allowed exploration of several implementation alternatives in order to find tradeoffs between processing speed and area of the PE. An instance of 25 PE running at 80 MHz is able to process 127 QVGA or 35 VGA images per second.

The paper entitled "Smart Camera Based on Embedded HW/SW Co-processor" by Romuald Mosqueron et al. describes an image acquisition and a processing system based on a new coprocessor architecture designed for CMOS sensor imaging. The system exploits the full potential CMOS selective access imaging technology because the coprocessor unit is integrated into the image acquisition loop. The acquisition and coprocessing architecture enables the dynamic selection of a wide variety of acquisition modes as well as the reconfiguration and implementation of high-performance image preprocessing algorithms. The experimental results show a large increase of the achievable performances. For instance, the new platform can successfully acquire and fully process up to 50 image-codes per second when applied to the detection and reading of bar codes in the case of a postal sorting application.

For low-volume applications like in professional electronics applications, FPGA are used in combination with DSP and GPP in order to reach required performances. Nevertheless, FPGA designs are static, which raises a flexibility issue with new complex applications. In this scope, dynamic partial reconfiguration (DPR) is used to bring a virtualization layer upon the static hardware of the FPGA. The contribution of the paper "An Evaluation of Dynamic Partial Reconfiguration for Signal and Image Processing in Professional Electronics Applications" by Philippe Manet et al. is to evaluate the interest and limitations when using DPR in real professional electronics applications, and to provide guidelines to improve its applicability. It makes an evaluation of DPR based on experiments made on a set of seven signal and image processing applications carried out in real conditions. It also identifies the missing elements and set of advantages for its use in professional electronic applications. Research directions are also proposed in order to improve its usage.

The paper entitled "Using High-Level RTOS Models for HW/SW Embedded Architecture Exploration: Case Study on Mobile Robotic Vision" by François Verdier et al. deals with the design of a System-on-Chip implementing the vision system of a mobile robot. Specific mechanisms necessary to build a high-level model of an embedded custom operating system able to manage real-time application are described. An executable RTOS model written in SystemC allowing an early simulation of the mobile robotic vision application is also detailed. Based on this model, a methodology is discussed and results are given on the exploration and validation of a distributed platform adapted to the vision system.

Signal processing algorithms become more and more performing and efficient as a result of new developments or at the release of new standards. Textual specifications have been substituted by reference software packages which have become the starting point of any design flow leading to the implementation of the algorithm. Therefore, designing an embedded application has become equivalent to port a generic SW on a, possibly heterogeneous, embedded platform. The paper entitled "A platform for the Development and the Validation of HW IP Components Starting from Reference Software Specifications" by Christophe Lucarz et al. describes a new platform aiming at supporting a step-by-step mapping of a reference software into SW and HW implementations. The platform provides a seamless interface between the software and hardware environments and in addition is supported by profiling capabilities able to analyze the transfers of data between HW and SW parts of the algorithm which help designers to optimize their design.

The paper entitled "A Priori Implementation Effort Estimation for HW Design Based on Independent-Path Analysis" by Rasmus Abildgren et al. presents a metric-based approach for estimating the hardware implementation effort (in terms of time) for an application in relation to the number of linear independent paths of its algorithms. By exploiting the relation between the number of edges and linear independent paths in an algorithm, the authors estimate implementation effort. This approach is implemented in an already existing design framework called "Design-Trotter" and offers a new type of tool to reduce the time-to-market.

The paper entitled "Multiple Word-length High-Level Synthesis" by Philippe Coussy et al. offers tradeoffs between the usage of uniform bit-length designs allowing for traditional automated design flows and nonuniform designs resulting in smaller circuits but to the extent of design complexity. The design flow, based on high-level synthesis (HLS) techniques, automatically generates a potentially pipeline RTL architecture described in VHDL. Both bit-accurate integer and fixed-point data types can be used

in the input specification. The generated architecture uses components (operator, register, etc.) that have different widths. The design constraints are the clock period and the throughput of the application. The proposed approach considers data word-length information in all the synthesis steps by using dedicated algorithms. While providing the same computing dynamic, the total area reduction ranges from 27% up to 80% compared to traditional (none or partial bit-width aware) high-level synthesis flows.

The paper entitled "Accuracy Constraint Determination in Fixed-Point System Design" by Daniel Menard et al. also deals with word length effects. Here, fixed-point system is modelled with an infinite precision version of the system and a single noise source located at the system output. Then, an iterative approach for optimizing the fixed-point specification under the application performance constraint is defined. Finally, the efficiency of this approach is demonstrated by experiments on an MP3 encoder.

Markus Rupp
Dragomir Milojevic
Guy Gogniat