Guest Editors' Introduction

Jarmo Takala · Shuvra S. Bhattacharyya ·
Wonyong Sung

Published online: 24 August 2012
© Springer Science+Business Media, LLC 2012

This special issue contains a selection of recent papers on design and implementation of signal processing systems. The papers in this issue span the areas of image and video processing, multiprocessor implementation techniques for signal processing systems, communication systems, and biomedical signal processing.

The first two papers are in the area of image and video processing. In *Quality-Aware Techniques for Reducing Power of JPEG Codecs*, Emre and Chakrabarti propose bit truncation and voltage overscaling schemes to greatly reduce the power consumption of JPEG codes. To compensate errors introduced by bit truncation, they employ unbiased estimation of the truncation noise. The errors due to aggressive voltage scaling are also compensated by using an algorithm-specific technique that exploits the characteristics of the quantized coefficients after zig-zag scan. A combination of these two techniques results in very high power savings with only a modest increase in PSNR.

J. Takala (✉)
Department of Computer Systems,
Tampere University of Technology, Tampere, Finland
e-mail: jarmo.takala@tut.fi

S. S. Bhattacharyya
Department of Electrical and Computer Engineering,
University of Maryland, College Park, MD, USA
e-mail: ssb@umd.edu

S. S. Bhattacharyya
Institute for Advanced Computer Studies,
University of Maryland, College Park, MD, USA

W. Sung
Department of Electrical Engineering,
Seoul National University, Gwanak-gu, Seoul, Korea
e-mail: wysung@snu.ac.kr

In *Joint Algorithm-Architecture Optimization of CABAC*, Sze and Chandrakasan present several optimizations that can be performed on Context Adaptive Binary Arithmetic Coding (CABAC) to achieve the throughput necessary for real-time, low power, high definition video coding. The combination of syntax element partitions and interleaved entropy slices, referred to as Massively Parallel CABAC, increases the bins per cycle by $2.7\text{--}32.8\times$ at a cost of $0.25\text{--}6.84\%$ coding loss compared with sequential single slice H.264/AVC CABAC. It also provides a $2\times$ reduction in area cost, and lowers the required memory bandwidth. Subinterval reordering is also applied and reduces the critical path by $14\text{--}22\%$, while modifications to context selection decrease memory requirements by 67% .

The next three papers address multiprocessor implementation of signal processing systems. In *Parallel Computation of Adaptive Filtering Algorithms on Multi-core Systems*, Lee, Ahn, and Sung develop methods for implementation of adaptive filters, including adaptive transversal least mean square, gradient adaptive lattice, and QR-decomposition based least-square lattice filters. Their methods exploit both multi-core and SIMD parallelism, and are demonstrated on a personal computer that employs both quad-core and 4-way SIMD structures. They demonstrate speedups ranging from 174 to 355 % for the different types of adaptive filters that they investigate.

In *Optimization and Parallelization of Monaural Source Separation Algorithms in the openBliSSART Toolkit*, Weninger and Schuller develop efficient implementations of monaural audio source separation algorithms that are targeted to graphics processing units (GPUs). They integrate methods for efficient parallel processing and numeric optimizations to reduce both

computation time and memory usage. They demonstrate significant speedups on a desktop personal computer platform equipped with multi-core CPU and GPU devices.

In *Communication-Aware Heterogeneous Multiprocessor Mapping for Real-time Streaming Systems*, Lin, Gerstlauer, and Evans develop methods for mapping real-time streaming systems specified as synchronous dataflow graphs onto bus-based multiprocessor systems. The methods carefully incorporate inter-processor communication costs, and allow designers to optimize multiprocessor implementations in terms of latency, throughput, and processor cost. A variety of mapping approaches is provided to support different trade-offs among solution quality, complexity, and run-time.

The next four papers are in the area of communication systems. In *A Reconfigurable TDMP Decoder for Raptor Codes*, Zeineddine and Mansour consider the challenging task of designing hardware-efficient decoders. The proposed decoder uses architecture-aware Raptor code construction, a turbo-decoding message passing algorithm, and various code enhancements and architectural optimizations. The efficiency is shown with area and power estimates from a synthesized design.

In *Hardware Implementation of Successive-Cancellation Decoders for Polar Codes*, Leroux, Raymond, Sarkis, Tal, Vardy, and Gross propose two successive cancellation polar decoder architectures with linear complexity. The decoding is implemented in the logarithmic domain, thus multiplications and divisions are avoided. The linear complexity has been verified by synthesizing the logic on standard cell technology.

In *DART—A High Level Software Defined Radio Platform Model for Developing the Run-time Controller*, Palkovic, Declerck, Avasare, Glassee, Dewilde, Raghavan, Dejonghe, and Van der Perre describe a platform model framework, which is used to develop reactive controllers for cognitive radio applications requiring rapid switching of modulation and coding schemes when channel conditions change. The applicability of the method is verified by comparing against transaction level modeling simulations.

In *Multi-rate Polyphase DSP and LMS Calibration Schemes for Oversampled ADCs*, Gupta, Tang, Paramesh, and Allstot propose adaptive techniques to estimate and correct errors of low-performance analog circuits in the digital domain. This allows relaxing the design constraints for the analog front-end, resulting in lower power consumption. The experiments show significant power savings compared to sigma-delta A/D converters designed with conventional methods.

We conclude the special issue with two papers on biomedical signal processing systems. In *Low-energy Formulations of Support Vector Machine Kernel Functions for Biomedical Sensor Applications*, Lee, Kung, and Verma develop a formulation for the kernel function of a support-vector machine classifier that can substantially reduce the real-time computations involved in some biomedical applications. Using clinical patient data for EEG-based seizure detection and ECG-based arrhythmia detection, they show that the polynomial transformation yields accuracy performance comparable to the most powerful available transformation (i.e., the radial-basis function). The proposed formulation reduces the energy by over $2,500\times$ in the arrhythmia detector and $9.3\text{--}198\times$ in the seizure detector.

In *A Fully Implantable, Programmable and Multimodal Neuroprocessor for Wireless, Cortically Controlled Brain-Machine Interface Applications*, Zhang, Aghagolzadeh, and Oweiss present the design and implementation of a neuroprocessor for conditioning raw extracellular neural signals recorded through microelectrode arrays chronically implanted in the brain of awake behaving rats. The neuroprocessor design exploits a sparse representation of the neural signals to combat the limited wireless telemetry bandwidth. They demonstrate a multimodal processing capability inherent in the neuroprocessor to support a wide range of scenarios in real experimental conditions. The neuroprocessor supporting 32 channels has been fully implemented on a 5×5 mm nano-FPGA, and the prototyping resulted in 5.19 mW power consumption.

We would like to thank all of the authors of this special issue for their contributions. We would also like to thank the anonymous reviewers for their efforts in ensuring the quality of the papers.