

Editorial

Special Issue: Algorithm/Architecture Co-Exploration of Visual Computing on Emerging Platforms

I. INTRODUCTION

CONCURRENT exploration of both algorithmic and architectural optimizations is an increasingly popular design paradigm. This special issue focuses on the latest research and development of video coding, processing, and computing algorithms on emerging platforms with multiple cores or reconfigurable architectures, including multiprocessor system-on-chips (MPSoCs), field-programmable gate arrays (FPGAs), multicore digital signal processors (DSPs), multicore central processing units (CPUs), and general-purpose computing on graphics processor units (GPGPUs).

Traditional design methodologies that involve sequential design exploration or one-way mapping of a fully specified algorithm to a selected architecture are not adequate for coping with future challenges. The algorithms in forthcoming visual systems are more complex than ever. Since we will witness continuous enhancement of visual quality, future platforms must feature higher performance. Moreover, since many systems require deployment with multiple applications, each with a different performance expectation, increased platform flexibility is much needed. Furthermore, visual computing applications are ubiquitous, available even in many energy-constrained devices; thus, it is critical to further increase power efficiency. The simultaneous optimization of algorithm and architecture starting from early design stages is critical to achieving these objectives.

Systems with multiple cores or reconfigurable architectures open new possibilities for visual system designers in implementing highly complex visual computing algorithms. Advances in semiconductor technology mean emerging platforms will possess an ever-increasing number of processing units and better reconfigurability. For example, in recent years, we have seen the introduction of high-end GPUs with tens of cores, video game consoles consisting of eight-core processors, and some of the latest netbooks even coming with dual-core processors.

This special issue consists of 12 papers that address theoretical as well as practical issues related to the following topics of interest:

- 1) concurrent exploration of both algorithmic and architectural optimization;

- 2) emerging and visually enriched applications or algorithms on multicore or reconfigurable platforms;
- 3) innovative architectures with multiple processors and reconfigurability (including efficient caches, memory subsystems, and on-chip interconnects) for video coding and processing applications;
- 4) dataflow representation used in algorithm/architecture co-exploration for multicore and/or reconfigurable architectures;
- 5) characterization of algorithmic complexity, potential parallelism, memory/data transfer;
- 6) design examples.

II. ORGANIZATION AND OVERVIEW

This special issue starts with “Algorithm/Architecture Co-Exploration of Visual Computing on Emerging Platforms: Overview and Future Prospects.” Due to the unique architectural characteristics of different classes of multicore processors, each visual computing algorithm may prefer one class of processor to the others. No single architecture is a clear winner across the broad spectrum of visual computing algorithms due to a wide variety of algorithm and application characteristics. The new design paradigm in which algorithms and architectures are concurrently explored is surveyed by the guest editors in this paper.

Second, the special issue includes two papers on emerging and visually enriched applications that are enabled by multicore platforms. The algorithms in visual systems are becoming more and more complex. GPUs are popular for many computationally intensive applications. “Fast JND-Based Video Carving with GPU Acceleration for Real-Time Video Retargeting” by Chiang *et al.* proposes new algorithms that can be easily parallelized for GPUs in order to achieve real-time video retargeting. In “Stream-Centric Stereo Matching and View Synthesis: A High-Speed Image-Based Rendering Paradigm on GPUs,” Lu *et al.* specifically steer the proposed image-based rendering system design with two high-level ideas. The first idea involves cost-effective image-based rendering. As long as synthesized views look visually plausible, the estimated disparity and occlusion need not be correct. Hence, the authors jointly optimize stereo matching and view synthesis for favorable end-to-end performance. The second idea is on real-time acceleration with GPUs, in which all functional modules are shaped at an early design stage to fit the massively parallel streaming architecture of GPUs.

First version published October 13, 2009; current version published October 30, 2009.

Digital Object Identifier 10.1109/TCSVT.2009.2034438

1051-8215/\$26.00 © 2009 IEEE

Third, the special issue includes three papers on innovative architectures with multiple processors and reconfigurability for video coding and processing applications. Many video applications involve both low-level and high-level tasks. In low-level processing, simple arithmetic operations are often performed on the pixels in a regular fashion. The high-level processing, however, is often more irregular. Thus, heterogeneous architectures are suitable for such applications. “A Configurable Heterogeneous Multicore Architecture with Cellular Neural Network for Real-Time Object Recognition” by Kim *et al.* provides a good example of a complete algorithm-specific architecture design flow, starting from an algorithm study and workload characterization, and working all the way down to a prototype chip implementation. In “A Self-Reconfigurable Platform for Scalable DCT Computation using Compressed Partial Bitstreams and BlockRAM Prefetching,” Huang and Lee start from the discrete cosine transform (DCT) algorithm and map it onto FPGAs with the intent to utilize the FPGA architecture as efficiently as possible. Furthermore, beyond the normal usage of reconfigurable hardware for fast prototyping, this paper presents a run-time adaptive architecture for DCT. It controls the number of coefficients encoded to reduce power or bandwidth consumption with some loss of quality. “VisoMT: A Collaborative Multithreading Multicore Processor with Fast Data Switching Mechanism for Multimedia Applications” by Ku *et al.* presents a complete design of a multicore architecture. The work includes multithreading cores, a fast data switching mechanism between different levels of storage, and a programming model. Finally, a case study on an AVC/H.264 encoder is also discussed.

Fourth, the special issue includes two papers on dataflow representations providing good models for the co-exploration of algorithms and architectures. Using the example of a reconfigurable video coding (RVC) decoder, “Exploring the Concurrency of an MPEG RVC Decoder Based on Dataflow Program Analysis” by Gu *et al.* discusses techniques and tools that can be used to implement the RVC decoder on parallel embedded computing platforms. This paper explores the concurrency in the specification of the decoder system by dataflow analysis, and presents tools and techniques for implementing it on a parallel computing platform. “A Framework for Heuristic Scheduling for Parallel Processing on Multicore Architecture—A Case Study with Multiview Video Coding (MVC)” by Pang *et al.* present a framework for performing analysis, simulation, and evaluation of dynamic scheduling schemes with exploration at different data granularities for implementing algorithms on different multicore processors.

Fifth, two papers are included on the characterization of algorithmic complexity in early design stages. This is essential to facilitate concurrent exploration of both algorithmic and architectural optimization. “Development of a High-Level Simulation Approach and its Application to Multicore Video Decoding” by Seitner *et al.* introduces a high-level simulation methodology for design space exploration. The method provides an opportunity to quickly analyze the performance of an implementation without performing a “complete” implementation on an FPGA or an application-specific integrated circuit (ASIC). It can save time, labor, and cost compared

to conventional design methods. The authors demonstrate that the proposed infrastructure can be used to implement video decoding on multicore processors. Next, “Profiling-Based Hardware/Software Co-Exploration for the Design of Video Coding Architectures” by Hübner and Stabernack describes a methodology to understand the computational characteristics of each component in a system. This methodology is based on coarse-grain profiling of high-level codes which is more suitable for processor-oriented platforms.

Finally, the special issue concludes with two papers on concrete design examples of mapping visual computing applications onto multicore platforms. The first example is on GPUs. While GPUs are capable of performing massive amounts of computation in parallel, GPU architectures pose some challenges to achieving efficient use of their provided computational capabilities. In particular, rate-distortion (RD) optimization possesses inherent data dependencies and conditional branches. “Highly Parallel Rate-Distortion Optimized Intra Mode Decision on Multicore Graphics Processors” by Cheung *et al.* presents an algorithm that computes approximate RD costs to find the RD-cost optimized intra mode, and a greedy-based block encoding order that accounts for the data dependencies in AVC/H.264 video encoding. The second example is “Multicore Processing and Efficient On-Chip Caching for H.264 and Future Video Decoders.” Finchelstein, Sze, and Chandrakasan demonstrate that multicore processors, combined with some algorithm changes to increase parallelism, are a power-efficient way to obtain the described performance. Nonetheless, external memory accesses consume a lot of power. This paper also shows an efficient on-chip caching scheme to reduce external memory access.

III. SELECTION PROCESS OF THE SPECIAL ISSUE

The goal of this special is to capture state-of-the-art in visual computing algorithms on emerging platforms with projection to the scope of future developmental trend based on algorithm architecture co-exploration. In order to choose the best papers which fit into the scope of the special issue, we used the following principles: The first criterion for evaluation was based on sufficient contributions in innovative algorithm or application within the scope of visual computing. The second factor considered was on novelties for either multicore and/or reconfigurable architecture. Hardware design using existing FPGA in a straightforward manner was not considered as having sufficient contributions in architecture. Finally, innovations were also sought in design methodologies, which map the algorithms onto the emerging platforms, namely the co-exploration of algorithm and architecture. Many papers were not selected due to insufficient content in the manuscripts to fit into the scope of the special issue or IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY even if some of them are excellent papers for the regular issues or other transactions.

This special issue is highly selective and competitive. We received 60 full manuscripts out of 77 abstract submissions in January. After the first round of reviews, 17 manuscripts were asked to be revised. After the second round of reviews,

to reassure that reviewers' concerns are fully addressed, we only accepted seven manuscripts and asked six manuscripts to be revised again. To satisfy the high-quality requirements of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, we chose only three out of those six papers. That is, only ten papers were accepted after multiple rounds of revision. The acceptance rate is one of the lowest within the past few years.

In the end, 12 papers were included in this special issue. We included one paper which was reviewed and accepted from the regular submission. This is because it is highly related to the theme of this special issue. The survey paper by the guest editors was handled independently by another associate editor and reviewed by anonymous reviewers.

We would like to thank everyone who submitted papers to the special issue for their efforts, and express our regret that due to limited space and the need for balanced coverage, not all high-quality submissions could be included. We also thank the authors for their valuable contributions, and the anonymous reviewers for their help in ensuring the quality of the special issue.

We sincerely hope that you enjoy this special issue and find its contents informative and useful.

ACKNOWLEDGMENT

We would like to express our most faithful gratitude to C. W. Chen, the Editor-in-Chief of IEEE TRANSACTIONS

ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, for his encouragement and support on this special issue. We also would like to thank the following Associate Editors of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY for their most valued advices and help in this special issue: I. Ahmad, O. Au, H. Chen, L.-G. Chen, S. Chien, W. Gao, L. Guan, G. Lafruit, S. Li, H. Sun, C. Taylor, G. Wen, and W. Zhu.

YEN-KUANG CHEN, *Guest-Editor*
Intel Corporation, Santa Clara
CA 95054-1549 USA

GWO GIUN (CHRIS) LEE, *Guest-Editor*
National Cheng Kung University
Tainan City 701
Taiwan

MARCO MATTAVELLI, *Guest-Editor*
École Polytechnique Fédérale de Lausanne
CH-1015, Lausanne, Switzerland

EUEE S. JANG, *Guest-Editor*
Hanyang University
Seoul 133-791, Korea