Journal of VLSI Signal Processing 41, 5–7, 2005

Guest Editorial: System-on-a-Chip for Multimedia Systems

The advance in VLSI technology allows the integration of more transistors on a single die. A wide range of functions—processors/controllers, application-specific modules, data storages, and mixed-signal circuits—can be combined on a single chip. Thus, a single chip can become a high-performing and inexpensive system for various multimedia applications—including image, video, graphics, and audio—which take an important role in our daily life. This Special Issue is devoted to discuss the new era of Multimedia SoC.

We would like to address some of the critical challenges in the deployment of System-on-a-Chip for Multimedia Systems. For example, we often optimize our architecture to the computational characteristics of the applications. However, in order to meet the requirements of continuous improvements in multimedia algorithms, the SoC must provide flexibility—in addition to a huge amount of arithmetic processing capability. To provide high computational capability and flexibility, how are the transistors deployed on the chip? For another instance, while many solutions can provide ultimate computational capability and flexibility, many solutions experience the problem of energy inefficiencies. Consuming the lowest level of power to deliver required performance is a key in multimedia SoC. How to deliver required performance with the lowest level of energy? Additionally, because digital, mixed-signal, radio-frequency, and memory blocks are integrated on a single chip, the power and substrate noise may cause sensitive blocks to suffer from functional failures. Moreover, the operating environments, which cannot be predicted by the designers of silicon-IP, vary from chip to chip across applications. How to tolerate noises to assure the yield of SoC design? To address these challenges, this issue contains 3 major parts. We present (1) system architectures that provide flexibility, (2) power-efficient and noise-tolerable designs, and (3) two application case studies.

This Special Issue opens with three system architectures of multimedia SoC. To meet the diverse computing requirements while achieving high computational and energy efficiencies, the system architectures presented here provide desired features via three different approaches: (1) combining multiple heterogeneous processing cores, (2) extending existing programmable processor cores with multimedia enhancements, and (3) reconfiguring computational circuits.

- "HiBRID-SoC: A Multi-Core SoC Architecture for Multimedia Signal Processing" by Stolberg et al. describes an architecture that combines high processing performance for multimedia schemes with high flexibility. Their design integrates multiple programmable cores, which are (1) individually optimized to a particular characteristic of different application fields in order to deliver high performance, and (2) complementing each other with flexibility at reduced system cost.
- "Accelerating Mobile Video: A 64-Bit SIMD Architecture for Handheld Applications" by Paver et al. discusses how to enhance computational competence of existing general-purpose processors using Single Instruction Multiple Data (SIMD) technologies. The paper explains how SIMD can exploit and accelerate the data parallelism inherent in multimedia processing. The technology can be applied to not only the general purpose processors, but also any programmable cores in SoCs.
- "Reconfigurable Discrete Wavelet Transform Processor for Heterogeneous Reconfigurable Multimedia Systems" by Tseng et al. presents an architecture that embraces the reconfigurable computing technology into multimedia SoC, and thus the computational resources can be flexibly adopted to target applications. While the paper reports how the proposed technique can be applied to wavelet filters and wavelet decomposition structures, the proposed approach can also be applied to universal and flexible computing engines for heterogeneous programmable/reconfigurable systems.

In addition to high performance and flexibility, power consumption is yet another key design trade-off. The design challenge is not only to build the system with the highest performance, but also to deliver the required performance with the lowest cost. Low-power design is critical to battery-powered devices, because we want our handheld or portable devices to operate longer. Moreover, energy-efficient design is also important to line-powered equipments, because power dissipation strongly influences the packaging/cooling cost and the reliability of the chip. Furthermore, noise-tolerate embedded memories have become the bottlenecks for the success of SoCs when memory blocks are integrated with digital, mixed-signal, and radio-frequency circuits on a single substrate. Because of the unique characteristics of the memory circuit and layout, the power and substrate noise may cause functional failures. The following three papers address the design of power-efficient and noise-tolerable architectures.

- Beric et al. present an algorithm/architecture co-design methodology to keep the power dissipation low while meeting the performance specification. Their paper "An Efficient Picture-Rate Up-converter" includes a specific design example. They first analyze and optimize the motion estimation algorithm based on the performance specification. Furthermore, they exploit data compression and reference locality to reduce the bottlenecks in memory capacity, data bus bandwidth, and power dissipation.
- Robinson et al. present and analyze the area, time, and energy efficiency of a monolithic integration of photo-detectors, analog-to-digital converters, data storage, and digital processing in their paper "Efficiency Analysis for a Mixed-Signal Focal Plane Processing Architecture." It is shown that the SIMD focal plan architecture, which incorporates analog and digital components to process pixels at the detection site, can improve both the performance and the efficiency of future portable image products.
- Chang et al. present a systematic approach dealing with power and substrate noise for embedded memories in multimedia SoCs. Their paper "Power and Substrate Noise Tolerance of Configurable Embedded Memories in SoC" first investigates the effects of such failures on memory cells, arrays, and circuit design. Then, a noise track-and-filter architecture is presented to provide the required timing relaxation and to minimize the speed degradation, making it very suitable for memory-dominated SoCs.

This Special Issue concluded with two application case studies to illustrate the software and hardware co-design used in multimedia applications. As video compression and audio compression are often used to improve the delivery of multimedia contents over limited channel bandwidth, we selected two case-study papers:

- A video decoder SoC is presented in "A Software-Hardware Co-Implementation of MPEG-4 Advanced Video Coding (AVC) Decoder with Block Level Pipelining" by Wang et al., and
- An audio decoder SoC is presented in "A Hardware/Software Co-Design of MP3 Audio Decoder" by Tsai et al.

Both papers illustrate a common co-design methodology, which includes (1) software and hardware partitioning, (2) software and hardware synchronization, (3) software optimization, (4) fast algorithms for hardware optimization, and (5) dedicated coprocessor design. Both papers reveal that many applications can be characterized into two portions: (1) complex data-dependent and decision-making procedures, and (2) computationally intensive and regular tasks. While we can use a programmable controller to implement the control-intensive tasks, we can use fast dedicated hardware modules to perform regular computation-intensive tasks. The methodology is applicable to not only video and audio decoders, but also other multimedia applications.

In short, we hope that this Special Issue provides enlightening information to the community in a timely fashion on the multimedia SoC. By understanding how the current SoCs are designed, readers can use this issue as a reference when they develop their own SoC solutions. We would like to thank the authors for their excellent contribution. We also appreciate the reviewers for their constructive comments.

Finally, we would like to thank Professor Sun-Yuan Kung, Editor-in-Chief, for suggesting this Special Issue, and for his patience, encouragement, and support for pursuing it. We also deeply appreciate Tzu-Ming Liu and Ching-Che Chung's work on the web system of paper submission and paper review for this Special Issue.

Yen-Kuang Chen
Corporate Technology Group, Intel Corporation, USA

Stella Kuei-Ann Wen
Department of Electronics Engineering, National Chiao Tung University, Taiwan

Chen-Yi Lee
Department of Electronics Engineering, National Chiao Tung University, Taiwan