Master Courses

M1 Design of Digital VLSI Signal Processing Chips

Organizer: Tobias Noll, RWTH Aachen, DE
Speakers: Tobias Noll, RWTH Aachen, DE
Lajos Gaszi, Infineon, DE
Matthias Schöbinger, Infineon, DE
Jos Huisken, Philips Research Lab, NL

The enabling key technologies and driving forces for today's and future Mobile, Internet and Broadband Multimedia systems are Digital Signal Processing and Optoelectronics. At the boundary between the higher layers and the physical transmission medium (wireless, optical, cable, storage, etc.) more and more sophisticated DSP technologies are required. Although the performance of embedded DSP kernels continuously increases, the performance required for the implementation of those advanced DSP systems increases even faster. Adding high performance co-processor blocks dedicated to standard functionalities (like filtering, convolution and error decoding) and applying an optimised HW/SW split significantly reduces DSP load, releasing flexible DSP resources for added value applications. Even more important in many applications is the resulting dramatic reduction of power dissipation and silicon area. An additional implementation alternative becoming more and more attractive is the use of re-configurable FPGA-like blocks. This leads to heterogeneous System-on-Chip architectures, where the most important challenges are to find the ideal HW/SW and analogue/digital partitioning for minimising system costs as silicon area, power dissipation, design effort, testability, etc. Keys to the success of this approach are 1) an early interaction between system level requirements and implementation issues, and 2) short development times. Parameterised, fairly accurate cost models here enable the evaluation of the high-dimensional design space with fast iteration cycles on algorithmic system level and are the basis for tradeoffs between system requirements and implementation cost already in an early phase of the system development. Techniques like minimisation of application-specific weighted complexity measures or Pareto techniques can be applied to support this evaluation process.

In designing the dedicated high-performance/low-power hardware blocks to be used in those heterogeneous architectures a careful optimisation on all levels of deep-sub-micron CMOS design is mandatory. On the algorithmic system and architectural level, equivalence transformations, pipelining, re-timing, look-ahead, as well as pre-compute-and-select techniques together with the use of redundant number representations are applied in order to find attractive solutions coping with the performance requirements as well as the I/O and memory bandwidth problem. Parallelisation and timesharing techniques are applied to match the intrinsic throughput of the optimised building blocks to the specified one. On logic and circuit level, proper clocking schemes and logic building blocks have to be selected, suiting well to the technology and supply voltage to be used, the noise margin requirements and fitting to the switching activity at the according instance. Finally, on the layout level only a careful dimensioning, placement, and routing strategy allows competitive solutions by avoiding instead of coping with the dramatically increasing interconnect delay issues. Preserving the high degree of locality inherent in most DSP algorithms at arithmetic level during the whole mapping process ensures highest throughput, small silicon area and even more important, lowest possible power dissipation.

The proper use of these strategies and techniques will be explained on exemplary challenging, actually industrial relevant systems and architectures. Benchmark comparisons demonstrate the benefit of this approach.

M2 Systems and Networks on Chip

Organizer: Ahmed Jerraya, TIMA, Grenoble, FR
Speakers: Giovanni De Micheli, Stanford U, US
Luciano Lavagno, Politecnico di Torino, IT
Sungjoo Yoo, Inter-University Semiconductor Research Centre, Seoul, KR
Ahmed Jerraya, TIMA, Grenoble, FR

The ITRS roadmap predicts that in 2005, 70 percent of ASICs will include at least one embedded instruction set processor. In this case, most ASICs will be SoCs (Systems-on-Chip). This prediction is not only confirmed but strengthened: SoCs will include several instruction-set processors in the case of applications such as mobile terminals (e.g. GSM), set-top boxes (e.g. pnx 8500 from Philips), game processors (e.g. PlayStation 2 from Sony) and network processors. All the above designs correspond to mass market products and are (or will be) integrated on a single chip for production cost reasons. It is even expected that these applications act as the main drivers for the semiconductor industry. Most system and semiconductor houses are working on platforms allowing the integration of several cores (CPU, DSP, MCU, co-processors) and sophisticated communication networks (hierarchical bus, TDMA-based bus, point-to-point connection and packet routing switch) on a single chip. The trend is then to build large designs as a networked System-on-Chip. The game is now to interconnect standard components as we used to do for boards a few years ago. This evolution is creating several breaking points in the design process.

This course will address the four main challenges for the design community:

- We consider systems on chips (SoCs) that will be designed and produced in five to ten years from today, with gate lengths in the range 50-100nm. We address the distinguishing features of a design methodology that aims at achieving reliable designs under the limitations of the interconnect technology. Specifically, we consider energy consumption reduction, under guaranteed quality of service (QoS), as a main objective in system design. We show that the unreliability of the physical layer is a potential show-stopper for SoC design. We argue that network technology can be used to provide a framework for designing on-chip interconnect. We visit different layers of a micro-network stack abstraction and show new directions toward designing on-chip communication.
- We will overview the main on-chip communication architecture schemes. The on-chip communication architecture implements the communication part (e.g. client-server relation, fifo's, etc..) of the designer's specification. The architecture consists of software and hardware parts. Software parts include middleware, operating system, device driver and hardware part DMA, bus/network adapter, communication network consisting of dynamic/static routers, and memory. We will focus on networks.
- We will introduce a design methodology for architectural exploration of networks on chip based on decoupling functionality from architecture, and computation from communication. A simple example of various mapping and refinement options for a single function-to-function token-based communication will be used to practically illustrate the basic concepts. A larger realistic on-chip network will then be used to describe what sort of information can be obtained by various mapping and performance analysis experiments. Although some specific tools will be used to make the explanation more concrete, the methodology is fully tool-independent, and can be implemented on top of several publicly available design frameworks.
- We will explain the different approaches used to build hardware/software interfaces for networked Systems-on-Chip. This includes building hardware wrappers to interface components with communication networks and software wrappers including OS to isolate software application from architecture.