

06141 Abstracts Collection
Dynamically Reconfigurable Architectures
— Dagstuhl Seminar —

Peter M. Athanas¹, Jürgen Becker², Gordon Brebner³ and Jürgen Teich⁴

¹ Virginia Polytechnic Institute, US
athanas@vt.edu

² Universität Karlsruhe, DE
becker@itiv.uni-karlsruhe.de

³ Xilinx - San José, US
Gordon.Brebner@xilinx.com

⁴ Universität Erlangen, DE
teich@informatik.uni-erlangen.de

Abstract. From 02.04.06 to 07.04.06, the Dagstuhl Seminar 06141 “Dynamically Reconfigurable Architectures” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Dynamically run-time reconfigurable computing architectures, adaptive systems, computational models, circuit technologies, system architecture, CAD tool support

06141 Executive Summary – Dynamically Reconfigurable Architectures

Dynamic and partial reconfiguration of hardware architectures such as FPGAs and XPPs brings an additional level of flexibility in the design of electronic systems by exploiting the possibility of configuring functions on-demand during run-time. This has led to many new ways of approaching existing research topics in the area of hardware design and optimization techniques. For example, the possibility of performing adaptation during run-time raises questions in the areas of dynamic control, real-time response, on-line power management and design complexity, since the reconfigurability increases the design space towards infinity. This Dagstuhl Seminar on Reconfigurable Architectures has aimed at raising a few of these topics e.g. on-line placement, pre-routing/on-line routing trade-off, power minimization etc., and also at presenting novel ideas on how to overcome the difficulties introduced in dynamic reconfigurable systems.

Keywords: Reconfigurable Computing, Reconfigurable Supercomputing, Organic Computing, Dynamic Reconfiguration, Reconfigurable Hardware

Joint work of: Becker, Jürgen; Teich, Jürgen; Brebner, Gordon; Athanas, Peter M.

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/837>

IP Cores Protection in FPGA Environment, Cryptographic Identification Primitives

Wael Adi (TU Braunschweig, D)

With the advent of multi-million gate chips, Field Programmable Gate Arrays (FPGAs) have achieved high usability for design verification, exchange, test and even production. Adding to this is the possibility of reusing readily available licensed IP to shorten the design cycle. A major concern for IP owners is the possible over-deployment of the IP into more devices than originally licensed. In this presentation, two system based on both public and secret-key cryptography embedded in a secured design exchange protocol for protecting the rights of the IP owner are introduced. The systems consist of hardware-supported design encryption and secured device authentication protocols. Design encryption based on secured device identification ensures that the IP can only be deployed into explicitly identified and agreed upon devices. The system is devised for an uncomplicated trustable design exchange scenario. The public-key functions use modular squaring (Rabin Lock) on the FPGA chip instead of exponentiation to reduce the hardware complexity.

Keywords: IP core protection, FPGA design protection, combined secret and Public-Key Identification, Electronic mutation, Write-only-memory, provable secret identity

See also: Wael Adi, R. Ernst, Bassel Soudan, A.Hanoun,, "VLSI Design Exchange with Intellectual Property Protection in FPGA Environment Using both Secret and Public-Key Cryptography". IEEE Computer Society Annual Symposium on VLSI, ISVLSI 2006, Karlsruhe, Germany, March. 2006

The (empty?) Promise of FPGA Supercomputing

Peter M. Athanas (Virginia Polytechnic Institute, USA)

There have been some notable success stories in the past that give merit to the viability of the creation of an FPGA-based supercomputer. When examining the computing potential of these devices, they appear to offer competitive computational characteristics that are highly competitive to contemporary high-performance processors. Recently, there have been supercomputer-class processing blades offered by the leading high-performance computing specialist, yet the

sales of these nodes have been less than spectacular. This talk examines why this may be the case, and explores the viability and cost-performance of FPGA-based supercomputers.

Keywords: FPGA supercomputing

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/732>

The ADRES coarse-grained reconfigurable Array Processor

Mladen Berekovic (Delft University of Technology, NL)

ADRES is a new coarse grain reconfigurable processor device, that is fully programmable in C. ADRES supports design-time and runtime reconfigurability for a broad range of multi-mode embedded applications, such as such as MPEG-2, MPEG-4, AVC/H.264 or Scalable Video Coding.

To address the challenges of multimode operations, ADRES provides improved power efficiency and performance within acceptable area constraints. ADRES targets a power efficiency of at least 40 MOPS/mW while being able to handle a peak performance of 20 GOPS, requirements that clearly go beyond the specs of any state-of-the-art core.

The ADRES array processor is a flexible template instead of a concrete instance. An architecture description language is developed to specify different ADRES instances with full compiler support. A script-based technique allows a designer to easily generate different instances by specifying different values for the communication topology, supported operation set, resource allocation and timing of the target architecture. Together with a retargetable simulator and compiler, this toolchain allows for architecture exploration and development of application domain specific processors.

ADRES supports a VLIW-like programming model with a pure VLIW mode for legacy code, and an array mode with very high dataflow-parallelism for the processing of compute intensive loops. In parallel to ADRES, an C-compiler that can create seamless code for both modes using modulo scheduling techniques and that is based on IMPACT was developed. The compiler takes automatically care of all necessary data transfers between the two modes. Several applications from the wireless and from the multimedia domain have been mapped on ADRES. As an example, an ADRES based system can perform AVC decoding in CIF resolution with less than 50 MHz on a 4x4 array on compiled C-Code. MPEG-2 CIF decoding needs only 27 MHz. Several key benchmark kernel loops have been mapped on ADRES and the results show that ADRES can extend the performance of state-of-the Art VLIW DSPs by a factor of 7, which makes ADRES an attractive alternative to multi-core DSP solutions. Synthesis results show, that such an ADRES core consumes less than 2 mm² in 90nm technology and can run with more than 500 MHz. Considering that the array size can be

further increased to 8x8 or beyond, ADRES offers significant room for further speedup.

ADRES significantly beats a state-of-the art DSP, like the TI C64x, with fewer resources and at comparable power consumption. At the same time ADRES offers performance and power scalability by using more resources and larger array sizes. Therefore, ADRES based designs will reduce the need for multi-processor solutions.

Keywords: Reconfigurable arrays, Tensilica, IMPACT, VLIW, DSP

HW/SW Codesign for Reconfigurable System-on-Chip using a Process Model

Neil W. Bergmann (The University of Queensland, AU)

Abstract: Reconfigurable System-on-Chip is a powerful method to harness the power of FPGA technology. However, there is a very limited pool of designers who can build hardware-software designs. This paper describes a way to simplify rSoC design, by making hardware coprocessors appear like software processes in a Linux development environment. We explain both the ways that hardware processes are controlled by software ghost processes, and also how hardware and software processes communicate.

Keywords: Reconfigurable System-on-Chip, FPGA, Embedded System

Full Paper:

<http://www.itee.uq.edu.au/~bergmann>

Adaptive On-Chip Multiprocessing

Christophe Bobda (TU Kaiserslautern, D)

The last decades have experienced a continuous growth in the performance of multiprocessor according to the Moore law. This increase in performance is mainly due to two main reasons: the clock frequency that keep growing and efficient use von Instruction Level Parallelism (ILP). With the difficulty to continuously improve the clock frequency as well as the ILP has investigation, Chip multiprocessor have begun to appear as an alternative to increase performance through processor level parallelisms. Most of the solutions proposed and developed are mainly SMP-based. Furthermore the architecture is fixed, thus limiting the architectural flexibility. With the growing capacity of FPGAs, it is more and more possible to place several soft or hardcore processors working in parallel on a given FPGA. Moreover, the flexible logic allows for the runtime adaptivity of applications by exchanging hardware accelerator.

In this talk, we present the on going work in adaptive on-chip multiprocessor at the University of Kaiserslautern.

Keywords: Multiprocessing on Chip, FPGA

Full Paper:

<http://soes.informatik.uni-kl.de/people/bobda>

Dynamically Adaptable Behaviours (?)

Gordon Brebner (Xilinx - San José, USA)

This talk introduced the general topic of new thought models for programmability and configurability. A main suggestion was to focus on application behaviours rather than architectures, and thus on adaptability rather than programmability. This would build on the success of research into reconfigurable architectures, to move to a focus on the needs of particular problem instances within broad application domains. Using hyper-programmability, or some alternative mechanism, each instance can be mapped to its own tailored architecture. Downstream from this, further future research areas were discussed. These included: making a big push towards a "third way" of system design that is neither hardware nor software like; devising apt computational models and theories of programming; systematised application and environment specificity; and harnessing programmability within future technologies in a natural manner. A major means of facilitating this future vision will be through education, especially of young people as yet unversed in first and/or second way thinking.

Reconfigurable Architectures and Instruction Sets: Programmability, Code Generation, and Program Execution

Rainer Buchty (Universität Karlsruhe, D)

Within Self-reconfiguring systems two basic problems arise: firstly, on instruction level reconfigurable instruction sets make program generation and execution inherently difficult. Secondly, reconfiguration must not violate certain restrictions vital for the running application.

We describe a combined low-overhead approach which targets both problems by instrumenting an attributed low-overhead run-time environment which is able to dynamically map application-specific instructions to a variety of implementation alternatives while strictly adhering to given application demands. Our approach can be used application-independent and is suitable for use within the adaptive planning stage of a Self-X system as demonstrated by a reference implementation.

Keywords: Self-X, Instruction Set Reconfiguration, Run-time Environment, Code Generation, Programming

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/733>

Enabling RTR for industry

Oliver Diessel (Univ. of New South Wales, AU)

This talk explores the promise of run time reconfigurable (RTR) technology and makes an attempt to identify critical support elements that need to be put in place in order to overcome barriers to enhanced RTR uptake in industry.

We outline a research project underway at the University of New South Wales to develop a positioning satellite receiver that exploits the diversity in satellite signals to mitigate the effects of interference. This project is examined as a case study to motivate the discovery of challenges an industrial organisation faces engineering a dynamically reconfigurable product.

Our progress towards the development of a methodology for providing communications infrastructure for module-based applications illustrates one of the efforts necessary to develop useful synthesis tools for RTR applications development.

We conclude with suggestions for how the academic community can better assist the commercial development of real applications.

Keywords: Run-time reconfiguration, industry support, design tools, module-based design, communications

Joint work of: Diessel, Oliver; Koh, Shannon

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/734>

Reconfiguration Time Aware Processing on FPGAs

Florian Dittmann (Universität Paderborn, D)

The possibility of partial reconfiguration of FPGAs during run-time can be used to implement systems that adapt their execution area over time. Two things are presented in this context:

1) For detailed investigations of partial reconfiguration, the two topics modeling and practical realization of reconfigurable systems must be rooted in the design process. We have developed a tool that meets this requirement. It eases the design of partial bitstreams for Xilinx FPGAs for research purpose. The tool wraps the obstacles of partial bitstream generation, motivating people new to this field. Moreover, the backend of the tool, a single UML class diagram that represents the whole characteristics of the reconfigurable system under development abstractly, allows to model reconfigurable systems in a comprehensive manner on a high level of abstraction. The UML diagram is filled during the design process until enough information for the generation of bitstreams is available.

2) In the single machine environment, several scheduling algorithms exist that allow to quantify schedules with respect to feasibility, optimality, etc. In contrast,

reconfigurable devices execute tasks in parallel, which intentionally collides with the single machine principle and seems to require new methods and evaluation strategies for scheduling. However, the reconfiguration phases of adaptable architectures usually take place sequentially. Run-time adaptation is realized using an exclusive port, which is occupied for some reasonable time during reconfiguration. Thus, we can find an analogy to the single machine environment. We investigate the appliance of single processor scheduling algorithms to task reconfiguration on reconfigurable systems. We determine necessary adaptations and propose methods to evaluate the scheduling algorithms.

Keywords: Real-Time, Partial Reconfiguration, Reconfiguration Time Scheduling

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/735>

Technological Aspects for 3D FPGAs architectures with optoelectronic interconnects

Dietmar Fey (Universität Jena, D)

In the talk a review of proposals concerning the use of optoelectronic interconnects for FPGAs is given. The perspectives and the benefits of using optics for dynamically optical interconnects as well as optically reconfiguration of architectures are critically compared to the possibility of electronic 3D chip stack technology. The author sees a chance for optics in the optically reconfiguration of context memories in smart optical sensors since in this application an optical interface is anyway available or rather necessary for the detection of images. This interface could also be used for a dynamic reconfiguration in combination with modern 3D chip mounting and assembly technologies to realise smart and very compact CMOS camera chip stacks for embedded systems.

Keywords: Optoelectronic interconnects, 3D FPGAs, multi-context architectures, smart optical sensors

Bridging the Gap between Relocatability and Available Technology: The Erlangen Slot Machine

Diana Göhringer (Universität Erlangen, D)

We present an FPGA-based reconfigurable platform called Erlangen Slot Machine (ESM). The main advantages of this platform are: First, the possibility for each module to access peripherals independent from its location through a programmable crossbar, and local SRAM banks for individual modules. This physical design eases the implementation of run-time reconfigurable partial modules and enables an unrestricted relocation of modules on the device. We present our two-board ESM implementation and demonstrate a partially reconfigurable video filter application as well as a relocatable computer game including a dedicated inter-module communication scheme.

Keywords: FPGA-based reconfigurable platform, inter-module communication, crossbar, video filter demo

Joint work of: Göhringer, Diana; Majer, Mateusz; Teich, Jürgen

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/736>

Reconfigurable Supercomputing: What are the Problems? What are the Solutions?

Reiner Hartenstein (TU Kaiserslautern, D)

The dominant paradox of supercomputing is the contrast between decades of excellent technology development (e. g. see the Gordon Moore curve) versus the almost stalled progress in sustained performance in many application areas. The counterparts to this supercomputing paradox are the 3 paradoxes of Reconfigurable Computing: the low power paradox, the high performance paradox, as well as the education paradox. Despite to the fact, that are very power-hungry, software to configware migrations have been reported which reduce the tens or hundreds of thousand dollars electricity bill by an order of magnitude.

Despite of the really awful technological parameters of FPGAs brilliant results with speedups by up to 4 orders of magnitude have been reported from software to FPGA migrations.

The presentation discusses the reasons of these paradoxes. A very important aspect is the way, how data and processors are brought together. From this point of view, supercomputing has used for decades the wrong road map, based on the wrong machine paradigm being extremely memory-cycle-hungry. The presentation illustrates, why the not instruction-stream-centered basic machine paradigm of Reconfigurable Computing provides the right road map to new horizons of supercomputing.

The Reconfigurable Computing education paradox shows, that its extremely high pervasiveness to applications in practically all disciplines of embedded system as well as scientific computing is possible, although computing-related curriculum recommendations completely ignore these subject areas. This leads to the conclusion, that these achievements have been implemented primarily from experts with backgrounds different from computing sciences. CS departments are the best possible institutions to overcome the methodology fragmentation between many FPGA application disciplines using their own domain-specific tool trick boxes, and, to develop models covering all aspects which the application disciplines have in common. Refusing to take this responsibility the undergraduate CS-related curricula also miss the most important job markets for their graduates.

Keywords: FPGA, reconfigurable computing, supercomputing, reconfigurable computing education, curricula, low power, high performance, configware

FlexFilm - an Image Processor for Digital Film Processing

Sven Heithecker (TU Braunschweig, D)

Digital film processing is characterized by a resolution of at least 2K (2048x1536 pixels per frame at 30 bit/pixel and 24 pictures/s, data rate of 2.2 GBit/s); higher resolutions of 4K (8.8 GBit/s) and even 8K (35.2 GBit/s) are on their way. Real-time processing at this data rate is beyond the scope of today's standard and DSP processors, and ASICs are not economically viable due to the small market volume. Therefore, an FPGA-based approach was followed in the FlexFilm project. Different applications are supported on a single hardware platform by using different FPGA configurations.

The multi-board, multi-FPGA hardware/software architecture is based on Xilinx Virtex-II Pro FPGAs which contain the reconfigurable image stream processing data path, large SDRAM memories for multiple frame storage and a PCI express communication backbone network. The FPGA-embedded CPU is used for control and less computation intensive tasks.

This paper will focus on three key aspects: a) the used design methodology which combines macro component configuration and macro-level floorplanning with weak programmability using distributed microcoding, b) the global communication framework with communication scheduling and c) the configurable, multi-stream scheduling SDRAM controller with QoS support by access prioritization and traffic shaping.

As an example, a complex noise reduction algorithm including a 2.5 dimensions DWT and a full 16x16 motion estimation at 24 fps requiring a total of 203 Gops/s net computing performance and a total of 28 Gbit/s DDR-SDRAM frame memory bandwidth will be shown.

Keywords: Digital film, FPGA, reconfigurable, stream-based architecture, weak programming, SDRAM-controller, QoS, communication centric, communication scheduling, PCI-Express

Joint work of: Heithecker, Sven; do Carmo Lucas, Amilcar; Ernst, Rolf

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/737>

Reconfigurable Processing Units vs. Reconfigurable Interconnects

Andreas Herkersdorf (TU München, D)

The question we proposed to explore with the seminar participants is whether the dynamic reconfigurable computing community is paying sufficient attention to the subject of dynamic reconfigurable SoC interconnects. By SoC interconnect, we refer to architecture- or system-level building blocks such as on-chip buses, crossbars, add-drop rings or meshed NoCs.

Our motivation to systematically investigate this question originates from conceptual and architectural challenges in the FlexPath project. FlexPath is a new Network Processor architecture that flexibly maps networking functions onto both SW programmable CPU resources and (re-)configurable HW building blocks in a way that different packet flows are forwarded via different, optimized processing paths. Packets with well defined processing requirements may even bypass the central CPU complex (AutoRoute). In consequence, CPU processing resources are more effectively used and the overall NP throughput is improved compared to conventional NPU architectures.

The following requirements apply with respect to the dynamic adaptation of the processing paths: The rule basis for NPU-internal processing path lookup is updated in the order of 100us, packet inter-arrival time is in the order of 100ns. Partial reconfiguration of the rule basis (and/or interconnect structure) with state of the art techniques would take several ms resulting in a continuously blocked system. However, performing path selection with conventional lookup table search and updates (and a statically configured on-chip bus) takes considerably less than 100ns. Hence, is there a need for new conceptual approaches with respect to dynamic SoC interconnect reconfiguration, or is this a "no issue" as conventional techniques are sufficient?

Keywords: Reconfigurable SoC interconnect

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/779>

AMIDAR: A new Modell for Adaptive Processors

Christian Hochberger (TU Dresden, D)

The AMIDAR model is a new model for processor construction. It borrows concepts from microprogramming, dataflow computers and other well known techniques from the computer architecture. The central aim of the AMIDAR model is to provide a processor that can be adapted easily to the running application. Adaptation can occur in different ways: The communication infrastructure can be adapted, the type and number of functional units can be adapted and finally, specialized functional units can be synthesized and integrated into the processor.

The talk explains the basic ideas of AMIDAR, the achievements we have already made (simulator, various profiling techniques), ongoing work (hardware implementation, synthesis) and will finally show the open challenges (synthesis algorithms, hardware architectures).

Keywords: Adaptive Processor, Java Bytecode Processor

Full Paper:

www.amidar.de

Physical 2D Morphware and Power Reduction Methods for Everyone

Michael Hübner (Universität Karlsruhe, D)

Required flexibility for future embedded systems including dynamic and partial reconfigurable hardware in order to optimize system status is the challenge for actual and future research. Exploiting adaptive hardware for introducing online-placement and routing provides the degree of freedom for run-time multi-adaptive system integration. The presentation introduces novel methods and tools for today and future technology.

Keywords: 2D Placement, Run-time Reconfiguration

Joint work of: Becker, Jürgen; Paulsson, Katarina; Hübner, Michael

Physical 2D Morphware and Power Reduction Methods for Everyone

Michael Hübner (Universität Karlsruhe, D)

Dynamic and partial reconfiguration discovers more and more the focus in academic and industrial research. Modern systems in e.g. avionic and automotive applications exploit the parallelism of hardware in order to reduce power consumption and to increase performance. State of the art reconfigurable FPGA devices allows reconfiguring parts of their architecture while the other configured architecture stays undisturbed in operation. This dynamic and partial reconfiguration allows therefore adapting the architecture to the requirements of the application while run-time. The difference to the traditional term of software and its related sequential architecture is the possibility to change the paradigm of bringing the data to the respective processing elements. Dynamic and partial reconfiguration enables to bring the processing elements to the data and is therefore a new paradigm. The shift from the traditional microprocessor approaches with sequential processing of data to parallel processing reconfigurable architectures forces to introduce new paradigms with the focus on computing in time and space.

Keywords: 2D online placement and routing, Reconfigurable Computing

Joint work of: Becker, Jürgen; Paulsson, Katarina; Hübner, Michael

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/739>

Managing power amongst a group of networked embedded fpgas using dynamic reconfiguration and task migration

David Kearney (Univ. of South Australia, AU)

Small unpiloted aircraft (UAVs) each have limited power budgets. If a group (swarm) of small UAVs is organized to perform a common task such as geolocation then it is possible to share the total power across the group by introducing task mobility inside the group supported by an ad hoc wireless network (where the communication encoding/decoding is also done on fpgas). In this presentation I will describe research into the construction of a distributed operating system where partial dynamic reconfiguration and network mobility are combined so that fpga tasks can be moved to make the best use of the total power available in a swarm of UAVs.

Keywords: Dynamic reconfiguration unpiloted aircraft operating system

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/740>

Superscalar Technology for Reconfigurable Processors

Bernd Klauer (Helmut-Schmidt-Universität - Hamburg, D)

The first superscalar architectures have been designed to allow instruction level parallelism without involving the programmer or the compiler into the parallelization process. Within an instruction window the issue logic schedules instructions by the dataflow principle and maps them on unemployed functional units which are suitable to perform the calculation as indicated by the opcode. The maximum total of concurrently executable instructions is then only restricted by the size of the instruction window and by the total of functional units. With this execution scheme it is possible to build processors with identical instruction sets but differently equipped execution stages.

As the execution stages can be differently equipped with functional units without affecting the executability of the code they are an interesting subject for reconfiguration.

In the presentation a superscalar processor is proposed that contains such a reconfigurable execution stage together with an extended issue logic and a configuration management unit. This unit controls the reconfiguration of the execution stage. It decides whether the current configuration of the execution stage is suitable for the upcoming computations or not. If it rates another configurations better than the current configuration of the execution stage, it can trigger the reconfiguration procedure. Together with the processor architecture some simulation and benchmarking results are shown.

Keywords: Superscalar reconfigurable architecture

Joint work of: Klauer, Bernd; Niyonkuru, Adronis

Back-End Issues in Hardware/Software-Compilation

Andreas Koch (TU Darmstadt, D)

This presentation will give an overview over some of the issues that need to be addressed in the design and implementation of the hardware-generating back-end of a hardware/software compiler. COMRADE is a compile flow that aims to map programs formulated in standard ANSI C to an adaptive computer, such that compute-intensive parts are accelerated on a reconfigurable compute unit, while less critical or unsuitable parts are executed on a standard processor. The talk shows some of the techniques that are used in the COMRADE back-end.

Keywords: Reconfigurable, hardware/software, compilation

Enabling RTR for industry

Shannon Koh (Univ. of New South Wales, AU)

This talk explores the promise of run time reconfigurable (RTR) technology and makes an attempt to identify critical support elements that need to be put in place in order to overcome barriers to enhanced RTR uptake in industry.

We outline a research project underway at the University of New South Wales to develop a positioning satellite receiver that exploits the diversity in satellite signals to mitigate the effects of interference. This project is examined as a case study to motivate the discovery of challenges an industrial organisation faces engineering a dynamically reconfigurable product.

Our progress towards the development of a methodology for providing communications infrastructure for module-based applications illustrates one of the efforts necessary to develop useful synthesis tools for RTR applications development. We conclude with suggestions for how the academic community can better assist the commercial development of real applications.

Floating Point FPGAs

Philip Leong (Imperial College London, GB)

In this talk, a case for developing an FPGA specifically optimised for floating point applications is given. Applications include signal processing, embedded systems and high performance computing, and such a device is likely to have speed and power consumption advantages over conventional FPGA and micro-processor technology. The performance improvement obtained by adding floating point units (FPUs) to an existing fine grain FPGA is estimated to be 2-10x on a number of benchmark applications. Different architectures involving flash based dynamic reconfiguration and the sharing of configuration bits are also discussed.

Keywords: FPGAs, floating point, embedded blocks

FlexFilm - an Image Processor for Digital Film Processing

Amilcar Lucas (TU Braunschweig, D)

On this presentation a multi-board, multi-FPGA hardware/software architecture, for computation intensive, high resolution (2048x2048 pixels), real-time (24 frames per second) digital Film processing. It is based on Xilinx Virtex-II Pro FPGAs, large SDRAM memories for multiple frame storage and a PCI express communication network. The architecture reaches record performance running a complex noise reduction algorithm including a 2.5 dimensions DWT and a full 16x16 motion estimation at 24 fps requiring a total of 203 Gops/s net computing performance and a total of 28 Gbit/s DDR-SDRAM frame memory bandwidth.

To increase design productivity and yet achieve high clock rates (125MHz), the architecture combines macro component configuration and macro level floorplanning with weak programmability using distributed microcoding.

Keywords: Weak-programming, stream-based architecture, digital film, reconfigurable, FPGA, SDRAM-controller, QoS

Joint work of: Lucas, Amilcar; Heithecker, Sven

Full Paper:

www.flexfilm.org

Bridging the Gap between Relocatability and Available Technology: The Erlangen Slot Machine

Mateusz Majer (Universität Erlangen, D)

We present a new concept as well as the implementation of an FPGA-based reconfigurable platform, the Erlangen Slot Machine (ESM). The main advantages of this platform are: first, the possibility for each module to access its peripherals independent from its location through a programmable crossbar, and local SRAMs banks for individual modules. This support eases the design and implementation of run-time reconfigurable partial modules and enables an unrestricted relocation of modules on the device.

We present our two board ESM implementation and demonstrate a partially reconfigurable video filter application.

Keywords: FPGA-based reconfigurable platform, inter-module communication, crossbar, video filter demo

Joint work of: Majer, Mateusz; Göhringer, Diana; Teich, Jürgen

Pre-Routed FPGA Cores for Rapid System Construction in a Dynamic Reconfigurable System

Douglas Maskell (Nanyang Technological University - SGP)

We present a new concept as well as the implementation of an FPGA-based reconfigurable platform, the Erlangen Slot Machine (ESM). The main advantages of this platform are: first, the possibility for each module to access its peripherals independent from its location through a programmable crossbar, and local SRAMs banks for individual modules. This support eases the design and implementation of run-time reconfigurable partial modules and enables an unrestricted relocation of modules on the device.

We present our two board ESM implementation and demonstrate a partially reconfigurable video filter application.

Keywords: FPGA-based reconfigurable platform, inter-module communication, crossbar, video filter demo

Joint work of: Maskell, Douglas ; Oliver, Timothy F.

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/741>

Multi-level Reconfigurable Architectures - The Switch Model

Martin Middendorf (Universität Leipzig, D)

Reconfigurable hardware has been successfully deployed for accelerating computationally demanding applications. While providing enormous flexibility dynamically reconfiguring applications suffer from the large reconfiguration overhead inflicted by contemporary architectures. We propose to design multi-level reconfigurable architectures that can help to reduce this reconfiguration overhead by introducing different levels of reconfiguration, each streamlining the capabilities of its lower levels.

In this talk we present formal models for multi-level reconfiguration. For the switch model where reconfigurable units are seen a cost model that measures the reconfiguration costs are defined. Based on this cost model we study the complexity of several optimization problems. One problem is to find for a given algorithm the optimal time steps when reconfiguration operations should be done of the different levels. Other problems are to find the optimal number of reconfiguration levels and the best granularity for reconfiguration of on the different levels. We present algorithms to solve this problems and present results for several applications.

Low Level Compiler for XMonarch

Vincent J. Mooney (Georgia Institute of Technology, USA)

Ongoing research at the Center for Research on Embedded Systems and Technology (CREST) at Georgia Tech presents a low level compiler for the XMonarch chip which includes a novel Field Programmable Compute Array (FPCA). The FPCA is a course-grained reconfigurable logic device, and the compiler infrastructure used is a modified version of Trimaran.

Keywords: Compiler, FPGA

Design of a Hardware/Software RTOS for FPGAs with Processors

Vincent J. Mooney (Georgia Institute of Technology, USA)

Moore's prediction – commonly known as Moore's "Law" but not a scientific law in the strict sense – indicates that in the next few years we will have digital circuits with ten billion transistors (we already have a billion+ transistor processor chip from Intel).

Clearly, a portion of the billion-transistor integrated circuit market will consist of traditional ASICs, e.g., for super-high volume devices such as cell phones. Another portion of the billion-transistor integrated circuit market will be dominated by processor designs such as Intel's Merced/Itanium architecture.

The rest of the picture is less clear; however, some percentage will likely be dominated by customizable heterogeneous multiprocessor chips with a reasonable (say, 30-60) percent of the chip consisting of reconfigurable and custom digital logic. For lack of a better term, we will refer to such Customizable Heterogeneous Multiprocessor chips as CHM chips. One example of a CHM chip is the Virtex-4. Standard argumentation in favor of RISC indicates that a processor's compiler and architecture must be designed together or codesigned. Similarly, we will argue that CHM chips require codesign of the architecture and the RTOS to run on the architecture.

The Hardware/Software Codesign Group at Georgia Tech is working on some ideas in this domain. Specifically, we will give a brief overview of three recent projects:

- (i) design of a System-on-a-Chip Lock Cache (SoCLC) where lock variables are placed in a special lock cache in a CHM chip – a client-server example using SoCLC shows a reduction in lock latency by a factor of up to 3.65X resulting in an overall speedup of 31% for the application,
- (ii) a specialized hardware structure and associated algorithm which speeds up deadlock detection by two to three orders of magnitude in reconfigurable logic when compared with software algorithms, resulting in a 38% overall speedup in a practical deadlock scenario, and

- (iii) a SoC Dynamic Memory Management Unit (SoCDMMU) integrated with a software RTOS and able to provide worst-case second-level memory allocation in 16 cycles in a four-processor SoC example, resulting in an example where average case application transition time is 4.4X faster and worst case application transition time is over 10X faster using the SoCDMMU versus the traditional software approach.

The talk will end with a brief description of a hardware/software RTOS generation framework able to integrate any mix of the three hardware RTOS units (i, ii or iii above) together with a software RTOS. This talk was given as the keynote at opening of the FPGAWorld 2004 Conference.

Keywords: Hardware/software codesign, RTOS, real-time

Full Paper:

<http://codesign.ece.gatech.edu/papers/papers.html>

See also: FPGAWorld Conference 2004, Vasteras, Sweden

Center for Research on Embedded Systems and Technology

Vincent J. Mooney (Georgia Institute of Technology, USA)

The Center for Research on Embedded Systems and Technology (CREST) at Georgia Tech is a small research center within the School of Electrical and Computer Engineering. A brief overview of CREST's major goals in education, research and commercialization is provided.

Keywords: Embedded

Hardware/Software Codesign of a Real-Time Operating System

Vincent J. Mooney (Georgia Institute of Technology, USA)

The hardware/software codesign group at Georgia Tech has carried out a number of projects in the broad area of codesign of a Real-Timed Operating System. A brief description of some interesting results – plus additional comments – are provided.

Keywords: Codesign

DynaCORE: An Adaptive System-on-Chip Architecture for Deep Packet Processing in Network Applications

Thilo Pionteck (Universität Lübeck, D)

Current network devices have to keep up with increasing bandwidth, growing complexity and rapid changes in network protocols and applications. Conventional hardware systems cannot meet with the required flexibility. Software-based solutions or even hybrid systems such as network processors that combine hardware and software solutions do not achieve the performance requirements. Hence, other architectural solutions have to be found in order to cope with the increasing data rates of network applications.

We address this problem with DynaCORE, an application specific coprocessor for offloading computationally intensive tasks from a network processor. The system-on-chip architecture is based on an adaptable network-on-chip which allows the dynamic replacement of hardware modules as well as the adaptation of the on-chip communication structure. The exploitation of partial dynamic reconfiguration allows a rapid adaptation towards modified system behaviors. The design of DynaCORE, its performance requirements as well as its hardware structure are introduced in this presentation.

Keywords: Network-on-Chip, Network Processor

Mapping Periodic Realtime Tasks to Reconfigurable Hardware

Marco Platzner (Universität Paderborn, D)

The increasing densities of FPGAs and the availability of dynamic reconfiguration modes enable hardware multitasking. Circuits are turned into hardware tasks that are scheduled, loaded, and executed on the reconfigurable resource during runtime. During the last years, the feasibility of hardware multitasking has been demonstrated by several prototypes. For application domains that combine high performance demands with dynamic task sets, a multitasking environment is even essential.

In this talk, we concentrate on the execution of periodic real-time tasks in a hardware multitasking environment, a problem that has not yet received sufficient attention. We first present three scheduling approaches: global EDF, MSDL, and partitioned EDF. Global and partitioned EDF are techniques adopted from multiprocessor scheduling, MSDL is a server-based scheduling technique trying to minimize the FPGA reconfiguration overhead. We discuss the construction of the schedules, efficient schedulability tests, and evaluate the scheduling performance by means of a simulation experiment. Then we turn to implementation-oriented issues and compare the scheduling approaches with respect to suitable FPGA execution models and the number of required device reconfigurations.

Dynamically Reconfigurable Processor-Like Architectures

Wolfgang Rosenstiel (Universität Tübingen, D)

Traditionally, FPGAs are deployed due to their flexibility to change the application over time. Newly developed architectures can be reconfigured within one clock cycle so that components of a device can re-used within a single application promising a better price-performance ratio. The reconfiguration keeping pace with the execution yields an additional degree of freedom that constitutes a new principle of reconfiguration. We name this principle processor-like reconfiguration. A silicon-proven processor-like reconfigurable architecture is NEC's Dynamically Reconfigurable Processor Architecture (DRP) which we use to validate parts of our research.

In our current work, it is evaluated how processor-like reconfiguration can be exploited by a high-level compiler and which architectural resources are needed for an efficient mapping of applications. To accomplish this, the CRC model (Configurable Reconfigurable Core) was developed as a general model for processor-like reconfigurable architectures. The features of the CRC model are modified according to the requirements imposed by mapping applications onto it. For this application mapping, well known techniques from C-based hardware synthesis and from compilers for VLIW processors are deployed. Instances of the CRC model can be synthesized and analyzed at the gate-level for a detailed assessment including a comparison to FPGAs. First results on mapping a real-world example from visual computing have shown considerable advantages of processor-like reconfigurable architectures compared to FPGAs.

Besides the fast reconfiguration mechanism for the functionality, we extend the concept of processor-like reconfiguration for voltage sources. The power dissipation in each time step, the total energy consumption as well as the energy-delay product can be reduced enormously by temporal-spatial voltage assignment. In contrast to other voltage scaling approaches no adaptation of the clock frequency is required.

In particular for coarse-grained reconfigurable architectures, a designer must consider how each application makes use of the provided architectural resources. Therefore, application-domain specific architectures are developed taking into account various techniques that may be used by the compiler.

Keywords: Reconfigurable computing, synthesis, compiler, power optimization

Implications of Organic Computing for Reconfigurable Computing

Hartmut Schmeck (Universität Karlsruhe, D)

The new research program of Organic Computing leads to a number of challenges for system design and architecture. Major requirements are properties like robustness, adaptivity, and flexibility. Adaptivity can be seen as a prerequisite for achieving robust behaviour in spite of disturbing external influences, flexibility refers to the capability of showing different types of behaviour invoked by dynamically changing requirements of the execution environment. Runtime Reconfiguration should be a promising technique for providing this robust, adaptive, and flexible behaviour. The key challenge, though, is the adequate design of observer-controller architectures which are suggested for achieving the necessary control of self-organized behaviour of networked collections of intelligent items. Self-organization will be indispensable because of the infeasibility to manage these systems explicitly and individually, but control is also necessary to prevent undesired behaviour of the system, either locally or globally. Reconfigurable Computing has the potential to provide key components of these architectures and to allow for true organic behaviour.

QUKU: A Coarse Grained Paradigm for FPGAs

Sunil Shukla (Universität Karlsruhe, D)

To fill the gap between increasing demand for reconfigurability and performance efficiency, CGRAs are seen to be an emerging platform. The advantage lies in quick dynamic reconfiguration and power efficiency. Despite having these advantages they have failed to show their mark. This paper describes the QUKU architecture, which uses a coarse-grained dynamically reconfigurable PE array (CGRA) overlaid on an FPGA. The low-speed reconfigurability of the FPGA is used to optimize the CGRA for different applications, whilst the high-speed CGRA reconfiguration is used within an application for operator re-use.

Keywords: FPGA, CGRA, Reconfiguration

Joint work of: Shukla, Sunil; Bergmann, Neil W.; Becker, Jürgen

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/742>

Full Paper:

<http://www.itee.uq.edu.au/~sunil/research.htm>

Efficient architectures for streaming applications

Gerard Smit (University of Twente, NL)

This presentation will focus on algorithms and reconfigurable tiled architectures for streaming DSP applications. The tile concept will not only be applied on chip level but also on board-level and system-level. The tile concept has a number of advantages:

1. depending on the requirements more or less tiles can be switched on/off,
2. the tile structure fits well to future IC process technologies, more tiles will be available in advanced process technologies, but the complexity per tile stays the same,
3. the tile concept is fault tolerant, faulty tiles can be discarded and
4. tiles can be configured in parallel.

Because processing and memory is combined in the tiles, tasks can be executed efficiently on tiles (locality of reference).

There are a number of application domains that can be considered as streaming DSP applications: for example wireless baseband processing (for Hiper-LAN/2, WiMax, DAB, DRM, DVB), multimedia processing (e.g. MPEG, MP3 coding/decoding), medical image processing, color image processing, sensor processing (e.g. remote surveillance cameras) and phased array radar systems. In this presentation the key characteristics of streaming DSP applications are highlighted, and the characteristics of the processing architectures to efficiently support these types of applications are addressed.

Keywords: Reconfigurable streaming efficient

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/743>

Optoelectronic methods to beat Moore's Law

John Snowdon (Heriot-Watt University Edinburgh, GB)

It has been obvious for sometime that a new technology will be required if Moore's law is to be supported. It can readily be seen that optoelectronics offers a bandwidth higher than copper and indeed that latency can be traded off against bandwidth when interfacing is involved. One possibility is to build the compute structure around the optics.

Dynamically Reconfigurable Systems-on-Chip

Walter Stechele (TU München, D)

The design space for dynamically reconfigurable SoCs can be seen in three dimensions:

1. the system architecture for computation and communication, ranging from dataflow-oriented dedicated logic blocks to instruction flow-oriented micro-processor cores, from dedicated point-to-point connections to Networks-on-Chip.
2. the granularity of reconfigurable elements, ranging from simple logic Look-Up-Tables to complex hardware accelerator engines and reconfigurable interconnect structures.
3. the configuration life cycle, ranging from application changes (in the order of seconds) to instruction-based reconfiguration (in the order of nanoseconds).

We propose to use dynamically reconfigurable computing for video processing in driver assistance applications. In future automotive systems, video-based driver assistance will improve security. Video processing for driver assistance requires real time implementation of complex algorithms. A pure software implementation, based on low cost embedded CPUs in automotive environments, does not offer the required real time processing. Therefore hardware acceleration is necessary. Dedicated hardware circuits (ASICs) can offer the required real time processing, but they do not offer the necessary flexibility. Specific driving conditions, e.g. highway, country side, urban traffic, tunnel, require specific optimized algorithms. Reconfigurable hardware offers high potential for real time video processing and adaptability to various driving conditions.

Our system architecture consists of embedded CPU cores for high-level application code, dedicated hardware accelerator engines for low level pixel processing, and an application-specific memory system. The hardware accelerators and the memory system are dynamically reconfigurable, i.e. hardware accelerator engines can be exchanged during runtime, controlled by the application code on the CPU. The life cycle of a configuration depends on the change of driving conditions. A requirement on the reconfiguration time is given by the frame rate of the video signal, e.g. 40 msec for the exchange and relocation of new engines.

Keywords: Dynamic reconfiguration, design space, video processing

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/744>

Libraries for Reconfigurable Computing

Jürgen Teich (Universität Erlangen, D)

We would like to present ideas for the subsequent Monday evening breakout session on the concept of libraries for dynamically reconfigurable computers.

In particular, we would like to address the following questions:

- (1) What are the application domains where such libraries could be useful?
- (2) What are the module types that could be collected?

Finally, also the questions of exchange formats and on the moderation of such libraries will be addressed.

Keywords: Library of Modules for Dynamically Reconfigurable Computers

Joint work of: Teich, Jürgen; van der Veen, Jan

Mojette Transform Implementation on Reconfigurable hHrdware

József Vásárhelyi (University of Miskolc, H)

Inscribing invisible marks (watermarking) into an image has different applications such as copyright, steganography or data integrity checking. Many different techniques have been employed for the last years in different spaces (Fourier, wavelet, Mojette domains, etc.). The presentation outline the development work related to create a functional block scheme of Mojette transform and Inverse Mojette transform using reconfigurable hardware.

Keywords: Image processing, Mojette transformation

Joint work of: Vásárhelyi, József; Serfözö, Péter

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/746>

PISC: Polymorphic Instruction Set Computers

Stamatis Vassiliadis (Delft University of Technology, NL)

We introduce a new paradigm in the computer architecture referred to as Polymorphic Instruction Set Computers (PISC). This new paradigm, in difference to RISC/CISC, introduces hardware extended functionality on demand without the need of ISA extensions. We motivate the necessity of PISCs through an example, which arises several research problems unsolvable by traditional architectures and fixed hardware designs. More specifically, we address a new framework for tools, supporting reconfigurability; new architectural and microarchitectural concepts; new programming paradigm allowing hardware and software to coexist in a program; and new spacial compilation techniques. The paper illustrates the theoretical performance boundaries and efficiency of the proposed paradigm utilizing established evaluation metrics such as potential zero execution (PZE) and the Amdahl's law. Overall, the PISC paradigm allows designers to ride the Amdahl's curve easily by considering the specific features of the reconfigurable technology and the general purpose processors in the context of application specific execution scenarios.

Keywords: Polymorphic Processors, Polymorphic Instruction Set, Reconfigurable Computing, microcode, PISC

Joint work of: Vassiliadis, Stamatis; Kuzmanov, Georgi; Wong, Stephan; Moscu-Panainte, Elena; Gaydadjiev, Georgi; Bertels, Koen; Cheresiz, Dmitry

Full Paper:

<http://ce.et.tudelft.nl/publications.php>

See also: S. Vassiliadis, G.K. Kuzmanov, S. Wong, E. Moscu Panainte, G. N. Gaydadjiev, K. Bertels, D. Cheresiz, PISC: Polymorphic Instruction Set Computers, Proceedings of the International Workshop on Applied Reconfigurable Computing (ARC 2006), pp. 274-286, Delft, The Netherlands, March 2006, LNCS 3985

Reliability-Aware Power Management of Multi-Systems (CMPSoCs)

Klaus Waldschmidt (Universität Frankfurt, D)

Long-term reliability of processors in SoCs and NoCs is experiencing growing attention lately, since decreasing feature sizes and increasing temperatures have a negative influence on the lifespan. Recent work suggests an interplay between power management and reliability, since power management strategies affect the temperature of processors. Power management strategies are examined, which target to actively influence the long-term reliability of a multi-core processor, in integrated systems as e.g. for SoCs and NoCs.

The approach shows that dynamic parallelism can improve the reliability of multi-core systems significantly. First results were achieved by simulating a multi-processor using the Self Distributing Virtual Machine (SDVM) as a basis.

Keywords: Organic Computing, Adaptivity, Power Management, Power Consumption, Reliability, SDVM

Joint work of: Waldschmidt, Klaus; Haase, Jan; Hofmann, Andreas; Damm, Markus; Hauser, Dennis

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/745>

A Reconfigurable Outer Modem Platform for Future Communications Systems

Norbert Wehn (TU Kaiserslautern, D)

Future mobile and wireless communications networks require flexible modem architectures with high performance.

Efficient utilization of application specific flexibility is key to fulfill these requirements.

For high throughput a single processor can not provide the necessary computational power. Hence multi-processor architectures become necessary. This

paper presents a multi-processor platform based on a new dynamically reconfigurable application specific instruction set processor (dr-ASIP) for the application domain of channel decoding. Inherently parallel decoding tasks can be mapped onto individual processing nodes. The implied challenging inter-processor communication is efficiently handled by a Network-on-Chip (NoC) such that the throughput of each node is not degraded. The dr-ASIP features Viterbi and Log-MAP decoding for support of convolutional and turbo codes of more than 10 currently specified mobile and wireless standards.

Furthermore, its flexibility allows for adaptation to future systems.

Keywords: Domain-specific reconfigurable platform, channel coding, outer-modem

Joint work of: Wehn, Norbert; Vogt, Timo; Neeb, Christian

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/730>

Implementing High Performance DSP Systems on Heterogeneous Programmable Platforms

Roger Woods (Queen's University of Belfast, GB)

The talk will look at the issues of implementing complex DSP systems on heterogeneous platforms comprising GPPs, DSP processors and FPGAs. The emphasis is to develop a high level design flow that allows optimisation to be carried out in a top-down manner but which can efficiently exploit IP cores that will have pre-determined features e.g. latency. The talk will describe how dataflow has been used and then modified to allow this to happen. A design example of a normalised lattice filter will be presented. Some conclusions and future work will be outlined along with application to reconfigurable systems.

Keywords: DSP systems, high level, design , synchronous dataflow , IP cores , system level design

Joint work of: Woods, Roger; McAllister, John

Towards an Automated Design of Application-specific Reconfigurable Logic

Peter Zipf (TU Darmstadt, D)

Reconfigurable logic is known to have the potential to provide better solutions than direct ASIC implementations or processors in some situations.

A necessary prerequisite for area advantages compared to ASICs or a better energy efficiency than processors is an application specific design of the reconfigurable unit. Adapting it to the specific requirements of an application helps to

compensate for the area and speed penalty introduced by reconfigurability. The data paths of reconfigurable units are best suited for data flow oriented tasks, but for many applications, both control flow and data flow must be handled, so a integration of the reconfigurable unit into a processor environment is an appropriate choice. By analysing the existing design flow and integration possibilities for reconfigurable units, a basis for discussing possible automation schemes and a standardised interface is defined.

Possible future research could investigate an automated design support for the building blocks of reconfigurable units and the definition of a standard processor interface for some classes of reconfigurable units.

Keywords: Application-specific reconfigurable units, processor integration, design automation

Joint work of: Zipf, Peter; Glesner, Manfred

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/731>