

CONFIGURABILITY FOR SYSTEMS ON SILICON: REQUIREMENT AND PERSPECTIVE FOR FUTURE VLSI SOLUTIONS

Jürgen Becker

Universität Karlsruhe (TH), Institut für Technik der Informationsverarbeitung (ITIV),
Karlsruhe, Germany

INVITED PAPER
MIDEM 2003 CONFERENCE
01.10.2003 - 03.10.2003, Grad Ptuj

Abstract: Systems-on-Chip (SoC) has become reality now, driven by fast development of CMOS VLSI technologies. Complex system integration onto one single die introduce a set of various challenges and perspectives for industrial and academic institutions. Important issues to be addressed here are cost-effective technologies, efficient and application-tailored hardware/software architectures, as well as corresponding IP-based EDA methods. This contribution will provide an overview on recent academic and commercial developments in Configurable Systems-on-Chip (CSoC) architectures, technologies and perspectives in different application fields, e.g. mobile communication and multimedia systems. Due to exponential increasing CMOS mask costs, essential aspects for the industry are adaptivity of SoCs, which can be realized by integrating reconfigurable re-usable hardware parts on different granularities into Configurable Systems-on-Chip (CSoCs).

Konfiguracijski sistemi na siliciju : Zahteve in vidiki za bodoča VLSI vezja

Izvilleček: Zaradi hitrega razvoja CMOS VLSI tehnologij so dandanes sistemi na čipu (SoC) že realnost. Zapletene systemske integracije na eno samo silicijevo tabletko predstavljajo vrsto izzivov za industrijske in akademske ustanove. V mislih imamo poceni tehnologije, učinkovite in uporabniško naravnane programske in strojne rešitve, kakor tudi odgovarjajoče metode elektronskega načrtovanja na osnovi intelektualne lastnine. Prispevek podaja pregled nad akademskim in komercialnim razvojem arhitektur konfiguracijskih sistemov na čipu (SoC) ter pregled nad pričakovanji in razvojem tehnologij na različnih področjih uporabe kot so mobilna telefonija in multimedijaki sistemi. Zaradi visokih cen mask za CMOS tehnologije je prilagodljivost SoC sistemov bistvenega pomena za uporabo v industriji. To dosežemo z integracijo rekonfiguracijskih celic različne granularije v konfiguracijski sistem na čipu (CSoC).

1. Introduction

Due to today's CMOS integration dimensions several designs and implementations of complex systems on silicon, so-called Systems-on-Chip (SoC), have been realized successfully. The term SoC is still not clearly defined and used with various interpretations in different situations. From my point of view, a SoC consists of at least two or more micro-electronic macro-components of complexities previously integrated separately into different single dies. Thus, such components, also often called IP-cores (Intellectual Property), can be distinguished by one or more of the following criteria, characterizing also the major aspects of SoC-level integration decisions (see figure 1):

- integration technology, e.g. different MOS-/Bipolar transistors and materials (Si, SiGe, GaAs, etc.), electronic/mechanical systems (MEMS), etc.
- signal domain, e.g. digital, analog design style, e.g. full-custom, semi-custom, pre-diffused, pre-wired + non-MOS styles
- computing domain, e.g. processor (time domain), dedicated ASIC-based (space domain), dynamically reconfigurable (time / space domain) + various memory-cores and technologies
- specification and programming method, e.g. high-level language HLL (C, C++, SystemC, Matlab, Java, etc.), Assembler language (μ C-specific), hardware description language HDL (Verilog /38/, VHDL /37/, Ella /41/, KARL /39/ /40/).

Thus, SoC-technologies are the consequent continuation of the ASIC technology, whereas complex functionalities, that previously required heterogeneous components to be merged onto a printed circuit board, are integrated within one single silicon chip. The first SoCs appeared in the early 1990s and consisted almost exclusively of digital logic constructions. Today SoCs are often mixed-technology designs, including such diverse combinations as embedded DRAM, high-performance or low-power logic, analog, RF, and even more unusual technologies like Micro-Electro-Mechanical Systems (MEMS) and optical input/output. But this development also raises its problems, e. g. it takes

an enormous amount of time and effort (-> cost) to design and integrate a chip. The cornerstone of the required change in design methodologies will be the augmented use of parts from previous designs and by making use of parts designed by third parties, which is called IP- or Core-based design /10/ /15/ /16/. Dependent on application constraints, important aspects for SoC solutions are:

- time-to-market constraints have to be fulfilled,
- SoC architecture flexibility, e.g. risk minimization by adaptivity for application implementation, e.g. in cases of late specification changes,
- long product life cycles, due to multi-standard/multi-product implementation perspectives, and multi-purpose usage to fabricate high volumes of the same SoC (-> cost decrease per chip).

Recently, in addition to ASIC-based, one new promising type of SoC architecture template is recognized by several academic /4/ /31/ /32/ /28/ /29/ /30/ and first commercial versions /17/ /18/ /19/ /21/ /23/ /24/ /25/: Configurable SoCs (CSoCs), consisting of processor-, memory-, probably ASIC-cores, and on-chip reconfigurable hardware parts for customization to applications. CSoCs combine the advantages of both: ASIC-based SoCs and multichip-board development using standard components, e.g. they require only minimal NRE costs, because they don't need expensive ASIC-tools for developing always different and in the future very expensive mask sets, every time the functionality or standards are changing. Thus, besides other advantages, an enormous cost and risk minimization perspective is obvious for industrial CSoCs.




Fig. 1: SoC-level Integration Design Space

In the following, recent fine- and coarse-grain reconfigurable technologies as well as corresponding academic and commercial developments in architectures and applications are discussed. Reconfigurable hardware architectures have been proven in different application areas /11/ /12/ /32/ /17/ /18/ to produce at least one order of magnitude in power reduction and increase in performance. The focus of this contribution will describe the actual status and results of an industrial/academic CSoC integration, consist-

ing of a SPARC-compatible LEON processor-core, a promising commercial coarse-grain XPP-array of suitable size from PACT XPP Technologies AG (Muenchen, Germany), and application-tailored global/local memory topology with efficient multi-layer Amba-based communication interfaces. The XPP architecture is regular structured for arbitrarily sized implementations, including regularity in combination with locality of data processing, e.g. for reducing power consumption. The complete adaptive SoC architecture is synthesized onto 0.18 and 0.13 μm UMC CMOS technologies at University of Karlsruhe (TH). Due to exponential increasing CMOS mask costs, the essential aspects for the industry are now risk-minimizing adaptivity and low cost of SoCs, which can be realized by integrating reconfigurable re-usable hardware parts on different granularities into CSoCs. In the last years ASIC/SoC markets for computer and communication applications had explosive revenue increases, compared to industrial and automotive areas. Relative to GSM, UMTS and IS-95 will require intensive layer 1 operations, which cannot be performed on today's processors /26/ /27/. Thus, optimized Hw/Sw partitioning of such computation-intensive tasks is necessary, whereas the flexibility to adapt to changing standards and different operation modes has to be considered. Based thereupon and future market demands, now several industrial and academic CSoC approaches arise /17/ /18/ /19/ /21/ /22/ /23/ /25/ /28/ /29/ /30/ /31/ /32/.

2. Reconfigurable Technologies and Power/Cost Trade-offs

Today's processing requirements are rapidly increasing as well as changing for embedded electronic systems, e.g. in emerging applications like mobile communications, multimedia, automotive infotainment, telemetry and others, performance demands are growing rapidly. With the growth rate recently slowing down, the integration density of microprocessors is more and more falling back behind Moore's law. Accelerators occupy most of the silicon chip area. Compared to hardwired accelerators more flexibility is provided by (dynamically) reconfigurable hardware parts, which will be explained later.

The low power optimization requirements are becoming more and more critical, either in the processor and especially in the embedded system world. The capacity of batteries is growing extremely slow (doubling every 30 years), especially compared to the increasing algorithm complexity and performance requirements, e.g. in future wireless algorithms (see figure 2). On the other side, the estimated processor performance and power figures cannot fulfill these requirements as well as the memory throughput demands, e.g. only every 10 years the growth of memory communication bandwidth is doubled. Because of the von Neumann bottleneck, memory bandwidth is an important issue. Avoiding this memory bottleneck not only by using accelerators, but also by innovative computing architectures, or even by breaking the dominance of the von Neu-





Fig. 2: Future Wireless Applications: Algorithm Complexity vs. Performance vs. Power Trade-offs




Source: Gilder Technology Report (Nick Tredennick, USA, 2003)

Fig. 3: Rising Costs per Wafer and the Amortization for Buildings and Equipment /33/


mann machine paradigm is a promising goal of new trends in embedded system development and CSE education.

Another, maybe most important aspect, is the exponential increase of CMOS mask costs, which results in an essential risk and cost factors for all development and production lines, e.g. smaller does in the future not necessarily mean better and cheaper. Moore's law meant doubling the number of transistors per die every eighteen months, which results in more transistors on the same silicon area at equivalent costs, or in the same number of transistors at lower costs. This theory assumes the downscaling of transistor dimensions by $\sqrt{2}$ every eighteen months and that the cost to process a wafer depends mainly on its size. This "law" was fulfilled by the corresponding semiconductor industry for a long time this way and returned the expected efficiency. Unfortunately, we have to deal now with a different situation, because the fixed costs for a semiconductor plant and to process a wafer have been increased exponentially in the last years, e.g. the lithography equipment and the cost per wafer mask set. This results in very high fixed cost factors for each wafer compared to the relative small variable costs to process a wafer through a fab line. The corresponding process and cost interrelations were evaluated and quantized by Nick Tredennick in his Gilder Technology Report /33/. In figure 3 a) the exponentially rising wafer fixed costs and the variable wafer processing costs are illustrated dependent on the transistor technologies, and figure 3 b) shows the cheapest transistors to be fabricated by fully amortized 250 nm fabrication lines. The assumptions in figure 3 do not consider the tremendous and even more increasing mask set costs, so that smaller transistors will be even more expensive. For more details about actual changes in semiconductors, especially about the detailed quantization formulas and assumptions and finally resulting process adoption rates, please see /33/. The former dominance of the procedural von Neumann microprocessor paradigm has been due to its RAM-based flexibility and that in many cases no application-specific silicon is needed.



Source:
Gartner Dataquest 2002

Fig. 4: ASIC/ASSP Semiconductor Consumption of different Application Areas



Source: T. Chassen (ISSCC '99)
and R. Hartenstein (ICECS 2002)

Fig. 5: Energy/Flexibility Conflict of different Hardware Architectures and Circuits

Throughput is the only limitation because of its sequential nature of operation. But now a second RAM-based computing paradigm is heading for mainstream: the application of multi-grain (dynamically) reconfigurable hardware architectures. Such kind of structural programming in space - in contrast to von Neumann based programming in time - provides massive parallelism at logic, operator and arithmetic level, often more efficient than vN-based process level parallelism. As a consequence of all facts and views described above we have to target new ways in exploiting the available silicon and technologies, e.g. not always the newest and most highly integrated versions, in more effective way. To fulfill the cost, power as well as performance

requirements of today's and future algorithm complexities new computing architectures and circuits with more efficiency, flexibility and operation cleverness have to be developed and applied. Thus, today's fine-grain and especially coarse- as well as multi-grain (dynamically) reconfigurable architectures will realize better performance / energy trade-offs than comparable mp, DSP or μ Controller platforms (see figure 5). Moreover, their (online) flexibility and silicon re-use features will result in essential cost and risk minimization effects necessary for future processor, VLSI and System-on-Chip solutions. The application fields and with corresponding complex algorithms and estimated ASIC/ASSP consumption are illustrated in figure 4.

The following section gives an overview on some selected industrial and academic architectures and System-on-Chip solutions applying fine- and coarse-grain (dynamically) reconfigurable hardware datapaths for several of the above mentioned algorithm fields.

3. Academic and Industrial System-on-Chip Solutions

Today's fine-grain and early coarse-grain reconfigurable hardware architectures are very useful in several application fields and are alternatives to specialized (multi-) processor solutions /4/ /7/ /8/ /10/ /11/ /12/ /13/ /17/ /18/ /28/ /29/ /30/ /32/ /34/. But, a minor part of the fine-grain area is used by CLBs (configurable logic blocks), which are the logic resources. Major part of the area is covered by a reconfigurable interconnect fabrics, provid-

Source: R. Hartenstein





Fig. 6: Illustration of fine-grain reconfigurable hardware resources (FPGA: 1 configured "wire" shown) /34/

ing wire pieces, switch boxes, and connect boxes to connect a pin of a CLB, with a pin of another CLB by programming a "soft wire" (an example shown in figure 6). The state of each switching transistor is controlled by a Flip-flop (FF) which is part of "hidden" configuration RAM (not shown in figure 6), also used to program the CLBs to select the particular logic function of each. By downloading new configuration code all this can be re-programmed anywhere and at any time. In the following an efficient fine-grain System-on-Chip solution tailored to baseband voice coding algorithm will be sketched. Within the MAIA CSoC a fine-grain FPGA-core realizes the reconfigurable hardware part. In general, the MAIA architecture consists of one control processor and other satellite units (can be processors, FPGAs or other units such as MAC, see figure 7). During computation and reconfiguration sequential threads are instantiated on the control processor, which configures the satellite processors and the on-chip reconfigurable communication network and manages the overall control flow of applications, either in a static compiled order, or through a dynamic real-time kernel. Thus, the architecture is reconfigurable in two respects - inter-satellite communication configurations and the fine-grain FPGA hardware part. The MAIA processor consists of a microprocessor core (ARM8) and 21 satellite processors: two MACs, two ALUs, eight address generators, eight embedded memories (4 512x16bit, 4 1kx16bit) and an embedded low-energy FPGA. Connections between satellites are accomplished through 2-level hierarchical mesh-structured reconfigura-

ble interconnect network. The ARM8 uses an interface control unit to configure and communicate data with satellites. The address generators and embedded memories are distributed to supply multiple parallel data streams to the computational elements. The MAIA chip was implemented using 0.25U 6-level metal CMOS process with a supply voltage of 1V and additional voltages of 0.4V and 1.5V. The die size of the implementation was 5.2mm x 6.7mm with 1.2 million transistors at 40 MHz with an average power dissipation of 1.5-2 mW. The Maia CSoC is optimized for selected mobile communication application parts, e. g. a full-rate VSELP voice coder algorithm was implemented at 30 MHz with 5.7 GOPS/Watt /31/.

Fine grain morphware lacks area/power-efficiency (figure 6). The physical integration density (transistors per chip) of FPGAs is roughly 2 orders of magnitude worse than the Gordon Moore Curve. Due to reconfigurability overhead roughly about only one percent of these transistors deserve the real application, so that the logical integration density is about 4 orders of magnitude behind Gordon Moore. For high throughput requirements coarse-grain reconfigurable hardware is the much more powerful and more area-efficient, also providing a massive reduction of embedded memory and time needed for configuration /34/. Coarse grain morphware is also about one order of magnitude more energy-efficient than fine-grain solutions (figure 5 and /1/ /2/ /3/). Whereas fine-grain FPGAs are using single bit wide CLBs (figure 6), coarse-grain reconfigurable Computing uses RPU (reconfigurable processing units), which, similar to ALUs, have major path widths, like 32 bits, for instance. Important applications stem from the performance limits of the "general purpose" processor, creating a demand for accelerators. Especially in application areas like multimedia, wireless telecommunication, data communication and others, the throughput requirements are growing faster than Moore's law (growth of required bandwidth: figure 2), along with growing flexibility requirements due to unstable standards and multi-standard operation /4/. Currently the requirements can be met only by coarse-grain hardware arrays from a provider like PACT (figure 9 and /5/).



Source: J. Rabaey, UC Berkeley

Application Example: FIR Filter




Fig. 7: MAIA CSoC and FIR Application /31/ /32/

First, a second selected academic CSoC example will be sketched here. This is an application-tailored architecture called DReAM /4/ /14/, a coarse-grain Dynamically Reconfigurable Architecture for Mobile communication systems. It was designed at the Darmstadt University of Technology for the requirements of future mobile communications systems. Especially the application area of mobile communication requires an adaptable SoC solution. The total system view of such a CSoC is shown in figure 8 /4/. The datapath oriented DReAM array can be seen in figure 8. It consists of an array of coarse-grained, dynamically Reconfigurable Processing Units (RPUs), which are connected with a local and a global communication network. The RPU is the major hardware component of the DReAM, which executes mainly arithmetic data manipulations for signal processing parts. In addition, dual-port




Fig. 8: DReAM CSoC Architecture Datapath and RAKE Application Results /4/

RAMs are used as Look-Up Tables when performing multiplications and the application-specific units are used for PN-code correlation operations. The DReAM architecture provides efficient and fast dynamic reconfiguration possibilities, e.g. only partly and during runtime. Further details to implemented examples and mapping techniques as well as performance results, e.g. a RAKE-Receiver specification for a data rate of 1.5 Mb/s based on a 0.35 μ m CMOS-process, can be found in /4/ /14/.

Next, two commercial CSoC solutions will be described:

- the A7 architecture from Triscend with fine-grain on-chip reconfigurable hardware /19/ /20/
- the dynamically reconfigurable XPP Architecture from PACT /23/, /24/, /6/, /7/

The A7 Configurable System-on-Chip (CSoC) device /19/, /20/ is a complete, high-performance user-programmable system, which contains an embedded 32-bit ARM7TDMI RISC processor and an embedded programmable logic architecture, optimized for processor and bus interface, a high-performance 32-bit internal bus supporting up to 455M-bytes per second peak transfer rates, and 16K-bytes of internal scratchpad SRAM memory and a separate 8K-byte cache. The ARM7TDMI is a general-purpose 32-bit RISC microprocessor that supports the complete ARM 32-bit instruction set and the reduced 16-bit instruction set. The ARM processor is integrated with other system components and the Configurable System Logic (CSL) matrix to provide a complete CSoC system. The embedded SRAM-based Configurable System Logic (CSL)

matrix provides full, easy-to-use system customization. The high-performance programmable logic architecture consists of a highly interconnected matrix of CSL cells. Resources within the matrix provide seamless access to and from the internal high-performance Configurable System Interconnect (CSI) bus, interconnecting the embedded processor, its peripherals, and the CSL matrix at a maximum speed of 60MHz. Each CSL cell performs various potential functions, including combinatorial and sequential logic and the output blocks (PIOs) provide a highly flexible interface between external functions and the internal system bus.

A very interesting and promising approach for CSoC integration is the eXtreme Processing Platform (XPP) /23/ /24/, /6/ /7/ (see figure 9), realizing a new runtime reconfigurable data processing technology that replaces the concept of instruction sequencing by configuration sequencing with high performance application areas envisioned from embedded signal processing to co-processing in different DSP-like application environments. The adaptive reconfigurable data processing architecture consist of following components:

- Processing Array Elements (PAEs), organized as Processing Arrays (PAs),
- a packet oriented communication network,
- a hierarchical Configuration Manager (CM) tree, and
- a set of I/O modules.

This supports the execution of multiple data flow applications running in parallel. A PA together with one low level




Fig. 9: PACT/Leon-based CSoC Architecture and Layout synthesized at Universitaet Karlsruhe (TH)

CM is referred as PAC (Processing Array Cluster). The low level CM is responsible for writing configuration data into the configurable objects of the PA. Typically, more than one PAC is used to build a complete XPP device. Doing so, additional CMs are introduced for configuration data handling. With an increasing number of PACs on a device, the configuration hardware assumes the structure of a tree of CMs. The root CM of the tree is called the supervising CM or SCM. This unit is usually connected to an external or global RAM. The basic concept consists of replacing the Von-Neumann instruction stream by automatic configuration sequencing and by processing data streams instead of single machine words, similar to /12/. Due to the XPP's high regularity, a high level compiler can extract instruction level parallelism and pipelining that is implicitly contained in algorithms /6/. The XPP can be used in several fields, e.g. as image/video processing, encryption, and baseband processing of next generation wireless standards, e.g. to realize also Software Radio approaches. 3G systems, i.e. based on the UMTS standard, will be defined to provide a transmission scheme which is highly flexible and adaptable to new services. Relative to GSM, UMTS and IS-95 will require intensive layer 1 related operations, which cannot be performed on today's processors /26/ /27/. Thus, an optimized HW/SW partitioning of these computation-intensive tasks is necessary, whereas the flexibility to adapt to changing standards and different operation modes (different services, QoS, BER, etc.) has to be considered. Therefore, selected computation-intensive signal processing tasks have to be migrated from software

to hardware implementation, e.g. to ASIC or coarse-grain reconfigurable hardware parts, like the XPP architecture. Within the application area of future mobile phones desired and important functionalities are gaming, video compression for multimedia messaging, polyphone sound (MIDI), etc. Therefore, a flexible, low cost hardware platform with low power consumption is needed for realizing necessary computation-intensive algorithms parts. Thus, PACT implemented several of these functionalities onto the cost-efficient 4x4 XPP array size, e.g. a 256-point FFT, a real 16 tap FIR filter, and a video 2d DCT (8x8) for MPEG-4 systems. Their newest commercial CSoC is called SMExPP and consists of an ARM-7 EJS and an 4x4 XPP array with efficient RAM-topologies promising a high boost in performance and flexibility. The technical and commercial trade-offs of this SMExPP solution is described in /7/ and /8/. First digital TV application performance results were obtained by evaluating corresponding MPEG-4 algorithm mappings onto the introduced ARM/XPP CSoC and based on the 0.13 μm CMOS technology synthesis results. Based on this coarse-grain CSoC version, performance/cost results of an MPEG-4 application is currently under implementation, whereas the Inverse DCT applied to 8x8 pixel blocks can be performed by an 4x4 XPP-Array in 74 clock cycles. Since the IDCT is one of the most complex operations in MPEG-4 algorithms, the preliminary clock frequency of 100 MHz based on 0.13 μm CMOS technology integration is sufficient for this real-time digital TV application scenario.

Another class of resources for reconfigurable computing is called multi-grain reconfigurable hardware, where several fine-grain pathwidth slices (2-/4-bits, for instance) with slice bundling capability including carry signal propagation can be configured to be merged into RPU's with a pathwidth of multiples of the slice path width (e. g. 16, 20, or 24 bits). Moreover, dependent on the targeted algorithm classes, bitlevel data operations, wordlevel arithmetic instructions, or even control-driven FSMs should be supported. These new hybrid architectures, combining the advantages of fine- and coarse-grain circuits into novel generic datapath approaches, are currently under development in different specialized research programs, e.g. funded by the German DFG and other institutions /35/.

4. SoC Education Aspects

The challenges in the development of application-tailored SoCs influences and changes the traditional design flow for chip, and thus today's engineering education. This should have some impact to the way how students in electronic engineering departments are taught, e.g. courses which fully cover all required skills for a SoC designer. The traditional education for students enables them to design stand-alone hardware components such as ASICs, instruction-set processor, memory, FPGA, analog and even RF CMOS chips /36/. Specially educated engineers are responsible for combining these components to a system. With the upcoming of SoCs these till now completely separate categories of design will merge to one design flow. A chip will no longer be assembled at the gate level but at the IP block level and IP interfaces /36/. Multidisciplinary system thinking is required for future designs, e.g. a vertical integration of system and application know-how with CAD and technology knowledge has to be realized in vertical education projects and labs (see figure 10). This education goal could be achieved successfully by the co-working of students and faculty within real system design projects, formalizing and encapsulating application-specific

techniques into reusable methods, libraries and tools shared by the entire educational community. Students and universities need access to the latest technical and industrial developments, and education has to be focused also on techniques and theories which are fundamental and time invariant. Such system architects /36/ should be able to operate efficiently in interdisciplinary teams with highly soft skilled members, required urgently by today's embedded systems divisions.

5. Conclusions and Outlook

The paper has given an introduction and overview on reconfigurable hardware systems and their VLSI integration. It also has pointed out future trends driven by technology progress and EDA innovations. Many system-level integrated future products without reconfigurability will not be competitive. Instead of continuous technology progress and deep-submicron integration more efficient and clever architectures by (dynamically) reconfigurable platform usage will often be the key to keep up the current innovation speed beyond the technology limits of silicon. It is time to revisit the available scientific results from reconfigurable-related R&D to derive promising commercial solutions and corresponding curricular updates in EE and CS education. Exponentially increasing CMOS mask costs demand adaptive and re-usable silicon, which can be efficiently realized by integrating reconfigurable circuits of different granularities into CSoCs, providing a potential for short time-to-market and post-fabrication error/functionality corrections (risk minimization!), multi-purpose/-standard features including comfortable application updates within product life cycles (volume increase: cost decrease). This results in the fact that several major industry players are currently integrating (dynamically) reconfigurable cores/datapaths into their processor architectures and system-on-chip solutions.




Fig. 10: CAD/VLSI Education Challenges and SoC Cost Aspects

6. References

- /1/ R. Hartenstein (invited): The Microprocessor is no more General Purpose; Proc. ISIS 1997
- /2/ R. Hartenstein (invited): Trends in Reconfigurable Logic and Reconfigurable Computing; ICECS 2002
- /3/ A. DeHon: The Density Advantage of Configurable Computing; IEEE Computer, April 2000
- /4/ J. Becker, T. Pionteck, M. Glesner: An Application-tailored Dynamically Reconfigurable Hardware Architecture for Digital Baseband Processing; SBCCI 2000
- /5/ <http://pactcorp.com>
- /6/ V. Baumgarte, et al.: PACT XPP - A Self-Reconfigurable Data Processing Architecture; ERSA 2001
- /7/ J. Becker, M. Vorbach: Architecture, Memory and Interface Technology Integration of an Industrial/Academic Configurable System-on-Chip (CSoC); IEEE Computer Society Annual Workshop on VLSI (WVLSI 2003), Tampa, Florida, USA, February, 2003
- /8/ M. Vorbach, J. Becker: Reconfigurable Processor Architectures for Mobile Phones; Reconfigurable Architectures Workshop (RAW 2003), Nice, France, April, 2003
- /9/ "ASIC Sstem-on-a-Chip", Integrated Circuit Engineering (ICE), <http://www.ice-corp.com>
- /10/ M. Glesner, J. Becker, T. Pionteck: Future Research, Application and Education Perspectives of Complex Systems-on-Chip (SoC); Proc. of Baltic Electronic Conference (BEC 2000), Oct. 2000, Tallinn, Estonia
- /11/ P. Athanas, A. Abbot: Real-Time Image Processing on a Custom Computing Platform, IEEE Computer, vol. 28, no. 2, Feb. 1995.
- /12/ R. W. Hartenstein, J. Becker et al.: A Novel Machine Paradigm to Accelerate Scientific Computing; Special issue on Scientific Computing of Computer Science and Informatics Journal, Computer Society of India, 1996.
- /13/ J. Rabaey, "Reconfigurable Processing: The Solution to Low-Power Programmable DSP", Proceedings ICASSP 1997, Munich, April 1997.
- /14/ J. Becker, N. Liebau, T. Pionteck, M. Glesner: Efficient Mapping of pre-synthesized IP-Cores onto Dynamically Reconfigurable Array Architectures; Proc. 11th Int'l Conference on Field Programmable Logic and Applications, Belfast, Ireland, 2001.
- /15/ Y. Zorian, R. K. Gupta, Design and Test of Core-Based Systems on Chips, in IEEE Design & Test of Computers, pp. 14-25, Oct.-Dec. 1997.
- /16/ B. Tuck, Integrating IP blocks to create a system-on-a-chip, in Computer Design, pp. 49-62, Nov. 1997.
- /17/ Xilinx Corp.: <http://www.xilinx.com/products/virtex.htm>.
- /18/ Altera Corp.: <http://www.altera.com>
- /19/ Triscend Inc.: <http://www.triscend.com>
- /20/ Triscend A7 Configurable System-on-Chip Platform - Data Sheet http://www.triscend.com/products/dsa7csoc_summary.pdf
- /21/ LucentWeb/ <http://www.lucent.com/micro/fpga/>
- /22/ Atmel Corp.: <http://www.atmel.com>
- /23/ PACT Corporation: <http://www.pactcorp.com>
- /24/ The XPP Communication System, PACT Corporation, Technical Report 15, 2000
- /25/ Hitachi Semic.: <http://semiconductor.hitachi.com/news/triscend.html>
- /26/ Peter Jung, Joerg Plechinger, "M-GOLD: a multimode basband platform for future mobile terminals", CTMC'99, IEEE International Conference on Communications, Vancouver, June 1999.
- /27/ Jan M. Rabaey: System Design at Universities: Experiences and Challenges; IEEE Computer Society International Conference on Microelectronic Systems Education (MSE '99), July 19-21, Arlington VA, USA
- /28/ S. Copen Goldstein, H. Schmit, M. Moe, M. Budiu, S. Cadambi, R. R. Taylor, R. Laufer "PipeRench: a Coprocessor for Streaming Multimedia Acceleration" in ISCA 1999. <http://www.ece.cmu.edu/research/piperench/>
- /29/ MIT Reinventing Computing: http://www.ai.mit.edu/projects/transit/dpga_prototype_documents.html
- /30/ N. Bagherzadeh, F. J. Kurdahi, H. Singh, G. Lu, M. Lee: "Design and Implementation of the MorphoSys Reconfigurable Computing Processor "; J. of VLSI and Signal Processing-Systems for Signal, Image and Video Technology, 3/ 2000
- /31/ Hui Zhang, Vandana Prabhu, Varghese George, Marlene Wan, Martin Benes, Arthur Abnous, "A 1V Heterogeneous Reconfigurable Processor IC for Baseband Wireless Applications", Proc. of ISSCC2000.
- /32/ Pleiades Group: http://bwrc.eecs.berkeley.edu/Research/Configurable_Architectures/
- /33/ Nick Tredennick: Gilder Technology Report, vol. IX no. 4, April 2003, USA
- /34/ J. Becker, R. Hartenstein: Configware and Morphware going Mainstream; Journal of Systems Architecture JSA (Special Issue on Reconfigurable Systems), June 2003
- /35/ Deutsche Forschungsgemeinschaft (DFG): Specialized Research Program 1148 „Reconfigurable Computing Systems": <http://www12.informatik.uni-erlangen.de/spprr/>
- /36/ H. De Man: System-on-Chip Design: Impact on Education and Research, IEEE Design & Test of Computers, publ. July-Sept. 1999, Volume 16 3, Page(s) 11-19
- /37/ IEEE Standard VHDL Language Reference Manual, Institute of Electrical and Electronics Engineers Inc., 1994, ISBN 1-55937-376-
- /38/ <http://www.verilog.com>
- /39/ Hauck, R.: KARL-4 - A hardware description language for the design and synthesis of digital hardware; Proc. 2nd ABAKUS workshop, Innsbruck, Austria, Sept. 1988
- /40/ R. Hartenstein: Hardware Description Languages; Elsevier, Amsterdam, 1987.
- /41/ J D Morison, A S Clarke: ELLA 2000, A Language for Electronic System Design; McGraw-Hill Book Company, ISBN 0-07-707821-7

Jürgen Becker
Universität Karlsruhe (TH)
Institut für Technik der Informationsverarbeitung (ITIV)
D-76128 Karlsruhe, Germany
[http://www.itiv.uni-karlsruhe.de/
becker@itiv.uni-karlsruhe.de](http://www.itiv.uni-karlsruhe.de/becker@itiv.uni-karlsruhe.de)

Prispelo (Arrived): 15.09.2003

Sprejeto (Accepted): 03.10.2003