

A Decade of Reconfigurable Computing: a Visionary Retrospective

Reiner Hartenstein (embedded tutorial)

CS Dept. (Informatik), University of Kaiserslautern, Germany
http://www.fpl.uni-kl.de hartenst@rhrk.uni-kl.de

Abstract. *The paper surveys a decade of R&D on coarse grain reconfigurable hardware and related CAD, points out why this emerging discipline is heading toward a dichotomy of computing science, and advocates the introduction of a new soft machine paradigm to replace CAD by compilation.*

1. Introduction

Rapidly increasing number and attendance [1] of conferences on re-configurable computing and the adoption of this topic area by congresses like ASP-DAC, DAC, DATE, ISCAS, SPIE, and others indicate, that reconfigurable platforms are heading from niche to mainstream, bridging the gap between ASICs and microprocessors (fig. 2). It's time to revisit R&D results: the goal of this paper. (On the implementation status of some projects we have only incomplete information.)

2. Coarse-Grained Reconfigurable Architectures

Using FPGAs as accelerator platforms is not subject of this paper. In contrast to FPGA use (fine grain reconfigurable) the area of *Reconfigurable Computing* mostly stresses the use of coarse grain reconfigurable arrays (RAs) with pathwidths greater than 1 bit, because fine-grained architectures are much less efficient, due to a huge routing area overhead and poor routability [2] [3]. Since computational datapaths have regular structure potential, full custom designs of *reconfigurable datapath units* (rDPUs) can be drastically more area-efficient, than by assembling the FPGA way from single-bit CLBs. Coarse-grained architectures provide operator level CFBs, word level datapaths, as well as powerful and very area-efficient datapath routing switches.

A major benefit is the massive reduction of configuration memory and configuration time, as well as drastic complexity reduction of the P&R (placement and routing) problem. Several architectures will be briefly outlined (also see figure 1). Some of them introduce *multi-granular* solutions, where more coarse granularity can be achieved by bundling of resources, such as e. g. 4 ALUs of 4 bits each to obtain a 16 bit ALU.

2.1 Primarily Mesh-Based Architectures

Mesh-based architectures arrange their PEs mainly as a rectangular 2-D array with horizontal and vertical connections

which supports rich communication resources for efficient parallelism. and encourages nearest neighbour (NN) links between adjacent PEs (NN or 4NN: links to 4 sides {east, west, north, south}, or, 8NN: NN-links to 8 sides {east, north-east, north, north-west, west, south-west, south, south-east} like w. CHESS array: fig. 5). Typically, longer lines are added with different lengths for connections over distances larger than 1.

DP-FPGA (Datapath FPGA) [4] has been introduced to implement regularly structured datapaths. It is a FPGA-like mixed fine and coarse grained architecture with 1 and 4 bit paths. Its fabric includes 3 component types: control logic, the datapath, and memory. The datapath block consists of 4 bit-slices: each bit-slice with a lookup table, a carry chain and a 4 bit register. DP-FPGA provides separate routing resources for data (horizontal, 4 bits wide) and control signals (vertical, single bit). A third resource is the shift block to support single-bit or multi bit shifts and irregularities.

The KressArray is primarily a mesh of rDPUs physically connected through wiring by abutment: no extra routing areas needed. In 1995 it has been published [5] as "rDPA" (reconfigurable DataPath Array). "KressArray" has been coined later. The KressArray is a super-systolic array (generalization of the systolic array: fig. 4) which is achieved by DPSS (see § „DPSS“). Its interconnect fabric distinguishes 3 physical levels: multiple unidirectional and/or bidirectional NN links (fig. 3), full length or segmented column and/or row backbuses, a single global bus reaching all rDPUs (also for configuration). Each rDPU can serve for routing only, as an operator, or, an operator with extra routing paths. All connect levels are layouted over the cell, so that wiring by abutment capability is not affected.

A first 32 bit KressArray included an additional control unit for the MoM-3 [6] Xputer [7] [8] [9] [10] with rDPUs supporting all C language operators. With the new Xplorer environment [11] rDPUs also support any other operator repertoires including




Fig. 2: Bridging the gap.

style	project	first publ.	source	architecture	granularity	fabrics	mapping	intended target application
mesh	DP-FPGA	1994	[4]	2-D array	1 & 4 bit, multi-granular	inhomogenous routing channels	switchbox routing	regular datapaths
	KressArray	1995	[5] [11]	2-D mesh	family: select pathwidth	multiple NN & bus segments	(co-)compilation	(adaptable)
	Colt	1996	[12]	2-D array	1 & 16 bit inhomogenous	(sophisticated)	run time reconfiguration	highly dynamic reconfig.
	Matrix	1996	[15]	2-D mesh	8 bit, multi-granular	8NN, length 4 & global lines	multi-length	general purpose
	RAW	1997	[17]	2-D mesh	8 bit, multi-granular	8NN switched connections	switchbox rout	experimental
	Garp	1997	[16]	2-D mesh	2 bit	global & semi-global lines	heuristic routing	loop acceleration
	REMARC	1998	[18]	2-D mesh	16 bit	NN & full length buses	(information not available)	multimedia
	MorphoSys	1999	[19]	2-D mesh	16 bit	NN, length 2 & 3 global lines	manual P&R	(not disclosed)
	CHESS	1999	[20]	hexagon mesh	4 bit, multi-granular	8NN and buses	JHDL compilation	multimedia
	DReAM	2000	[21]	2-D array	8 & 16 bit	NN, segmented buses	co-compilation	next generation wireless
	CS2000 family	2000	[23]	2-D array	16 & 32 bit	inhomogenous array	(not disclosed)	communication
	MECA family	2000	[24]	2-D array	multi-granular	(not disclosed)	(not disclosed)	tele- & datacommunication
	CALISTO	2000	[25]	2-D array	16 bit multi-granular	(not disclosed)	(not disclosed)	tele- & datacommunication
	FIPSOC	2000	[26]	2-D array	4 bit multi-granular	(not disclosed)	(not disclosed)	tele- & datacommunication
linear	RaPID	1996	[27]	1-D array	16 bit	segmented buses	channel routing	pipelining
	PipeRench	1998	[29]	1-D array	128 bit	(sophisticated)	scheduling	pipelining
crossbar	PADDI	1990	[30]	crossbar	16 bit	central crossbar	routing	DSP
	PADDI-2	1993	[32]	crossbar	16 bit	multiple crossbar	routing	DSP and others
	Pleiades	1997	[33]	mesh / crossbar	multi-granular	multiple segmented crossbar	switchbox routing	multimedia

Fig. 1: Summary of the technical details of the different coarse-grained reconfigurable architectures; Note: NN stands for "nearest neighbour".

branching, while loops and do-while loops. I/O data streams from and to the array can be transferred by global bus, array edge ports, or ports of other rDPUs (addressed individually by an address generator). Although KressArrays are dynamically partially configurable, applications tried out so far did not make use of it.

The KressArray Family. Supported by DPSS application development tool and platform architecture space explorer (PSE) environment the basic principles of the KressArray define an entire family of KressArrays covering a wide but generic variety of interconnect resources and functional resources. A later version of this PDE environment (see paragraph “Xplorer.”), supports the rapid creation of RA and rPDU architectures optimized for a particular application domain (like e. g. image processing, multimedia, or others), and rapid mapping of applications onto any RA of this family.

Colt [12] combines concepts from FPGAs and data flow computing [13]. It’s a 16 bit pipenet [14] and relies highly on runtime reconfiguration using wormhole routing. Hereby, the data stream headers hold configuration data for routing and the functionality of all PEs encountered. Colt has a mesh of 16 bit IFUs (Interconnected Functional Units), a crossbar switch, an integer multiplier, and six data ports. Each IFU features an ALU, a barrel shifter to support multiplication and floating point, a decision unit for flow branching, and optional delay units for pipeline synchronization.




Fig. 3: KressArray NN ports examples.

MATRIX [15] (Multiple Alu archiTecture with Reconfigurable Interconnect eXperiment) is a multi-granular array of 8-bit BFUs (Basic Functional Units) with procedurally programmable microprocessor core including ALU, multiplier, 256 word data and instruction memory and a controller which can generate local control signals from ALU output by a pattern matcher, a reduction network, or, half a NOR PLA. With these features, a BFU can serve as instruction memory, data memory, register-file and ALU, or independent ALU function. Instructions may be routed over the array to several ALUs. The routing fabric provides 3 levels of 8-bit buses: 8 nearest neighbour (8NN) and 4 second-nearest neighbour connections, bypass connections of length 4, and global lines.

The Garp Architecture [16] resembles an FPGA and comes with a MIPS-II-like host and, for acceleration of specific loops or subroutines, a 32 by 24 RA of LUT-based 2 bit PEs. Basic unit of its primarily mesh-based architecture is a row of 32 PEs, a reconfigurable ALU. The host has instruction set extensions to configure and control the RA. Array execution is initialized by the number of clock cycles to do. Host and RA share the same memory hierarchy. Memory accesses can be initiated by the RA, but only through the central 16 columns. The blocks in the leftmost column are dedicated controllers for interfacing. For fast reconfigurations, the RA features a distributed cache with depth 4, which stores the least recently used configurations. The routing architecture includes 2 bit horizontal and vertical lines of different length, segmented in a non-uniform way: short horizontal segments spanning 11 blocks, long horizontals spanning the whole array, and different length vertical segments.

RAW (Reconfigurable Architecture Workstation) [17] provides a RISC multi processor architecture composed of NN-connected 32-bit modified MIPS R2000 microprocessor tiles with ALU, 6-stage pipeline, floating point unit, controller, register file of 32 general purpose and 16 floating point registers, program counter, local cached data memory, and 32 Kilobyte SRAM instruction memory. The prototype chip features 16 tiles arranged in a 4 by 4 array. Early concepts [17] having been abandoned included also configurable logic to allow customized instructions. RAW provides both a static (determined at compile-time) and a dynamic network (determined at run-time: wormhole routing for data forwarding). Since processors lack hardware for register renaming, dynamic instruction issuing or caching (like in superscalar processors), statically scheduled instruction streams are generated by compiler, moving the responsibility for dynamic issues to the

array	applications	pipeline properties		mapping	scheduling (data stream formation)
		shape	resources		
systolic array	regular data dependencies	linear only	uniform only	linear projection or algebraic synthesis	
super-systolic RA	no restrictions			simulated annealing, genetic morphing, or other P&R algorithm	
				(e.g. force-directed) scheduling algorithm	

Fig. 4: Pipelined datapath arrays (pipe networks).

development software. But RAW provides possible flow control as backup dynamic support, if the compiler fails to find a static schedule.

REMARC (Reconfigurable Multimedia Array Coprocessor) [18], a reconfigurable accelerator, tightly coupled to a MIPS-II RISC processor, consists of an 8 by 8 array of 16 bit “nanoprocessors” with memory, attached to a global control unit. The communication resources consist of nanoprocessors NN connections and additional 32 bit horizontal and vertical buses which also allow broadcast to processors in the same row or column, respectively, or, to broadcast a global program counter value each cycle to all nanoprocessors, also to support SIMD operations.

MorphoSys (Morphoing System [19]) has a MIPS-like “TinyRISC” processor with extended instruction set, a mesh-connected 8 by 8 RA, a frame buffer for intermediate data, context memory, and DMA controller. The RA is divided into four quadrants of 4 by 4 16 bit RCs each, featuring ALU, multiplier, shifter, register file, and a 32 bit context register for storing the configuration word. The interconnect network features 3 layers: 4 NN ports, links of distance 2, and, inter-quadrant buses spanning the whole array. TinyRISC extra DMA instructions initiate data transfers between the main memory and the “frame buffer” internal data memory for blocks of intermediate results, 128 by 16 bytes in total.

The CHESS Array. The CHESS hexagonal array [20] features a chessboard-like floorplan with interleaved rows of alternating ALU / switchbox sequence (figure 5). Embedded RAM areas support high memory requirements. Switchboxes can be converted to 16 word by 4 bit RAMs if needed. RAMs within switchboxes can also be used as a 4-input, 4-output LUT. The interconnect fabrics of CHESS has segmented four-bit buses of different length. There are 16 buses in each row and column, 4 buses for local connections spanning one switchbox, 4 buses of length 2, and 2 buses of length 4, 8 and 16 respectively. To avoid routing congestion, the array features also embedded 256 word by 8 bit block RAMs. An ALU data output may feed the configuration input of another ALU, so that its functionality can be changed on a cycle-per-cycle basis at runtime without uploading. However, partial configuration by uploading is not possible.

The DReAM Array (Dynamically Reconfigurable Architecture for Mobile Systems [21]) for next generation wireless communication, is a 0.35 μm CMOS standard cell design fabricated by Mietec/Alcatel. Each RPU consists of: 2 dynamically reconfigurable 8-bit Reconfigurable Arithmetic Processing (RAP) units, 2 barrel shifters, a controller, two 16 by 8-bit dual port RAMs (used as LUT or FIFO), and, a




Fig. 5: CHESS array hexagonal floor plan.

Communication Protocol Controller. The RPU array fabric uses NN ports and global buses segmentable by switching boxes.

CS2000 family. Chameleon Systems [22] offers the CS2000 family multi-protocol multi-application reconfigurable platform RCP (reconfigurable communication processor) for telecommunications and data communications [23], with a 32 bit RISC core as a host, licensed from ARC Cores, UK, with full memory controller and PCI controller, connected to a RA of 6, 9, or 12 reconfigurable tiles, where each tile has 7 32-bit rDPUs (each including an 8 word instruction memory), 4 local memory blocks of 128 x 32 bits, 2 16x24-bit multipliers. Every 3 tiles are grouped as a “slice” also including 8 kBytes of local memory. The RA allows multiple independent data streams and its reconfigurable fabric can be changed within a single clock cycle. The CS1200 family aims at initial markets in communication infrastructure and is intended to cope with the chaotic world of evolving standards, protocols and algorithms with application areas as 2nd and 3rd generation wireless basestations, fixed point wireless local loop (WLL), smart antennas, voice over IP (VoIP), very high speed digital subscriber loop (DSL), and, for instance, supports 50 channels of CDMA2000.

The MECA family by Malleable [24], intended to be a family of DSPs (digital signal processors) optimized for VoIP, by compressing voice into ATM or IP packets etc., aims at next generation VoIP and VoATM. The MECA family claims - compared to conventional DSPs- a speed-up factor of 10.

CALISTO (Configurable ALgorithm-adaptive Instruction Set TOpology) by Silicon-Spice [25] intends to be an innovative adaptive instruction set architecture for internet protocols (IP) and ATM packet-based networks with flexibility for Any-Service-Any-Port (ASAP) to deliver voice and data simultaneously over a unified data network. CALISTO is a single-chip communications processor for carrier-class voice gateways, soft switches, and remote access concentrators/remote access servers (RAC/RAS), which aims at applications like echo cancellation, voice/fax/data modems, packetization, cellification, delay equalization.

FIPSOC (Field-programmable System-on-Chip) by SIDA [26] has an 8051 controller, a RA and a RAA (reconfigurable analog array). The 8x12 (8x16 or 16x16) RA is an array of “digital macro cells” (DMC) including a 4-input LUT and 4 latches, programmable to be a 4 bit up/down counter, shift register, 4 bit adder, 16x4 bit RAM etc. The RAA has “configurable analog blocks” (CAB) usable as differential amplifiers, comparators, converters etc. The RA is a multi-context RA featuring 2 extra configuration memories. Device booting can be done from external parallel ROM (as the 8051 controller does), external serial PROM, or RS 232 serial port. FIPSOC is intended to be used as an ASIC emulator for rapid prototyping.

2.2 Architectures Based on Linear Arrays

Some RAs are based on one or several linear arrays, typically also with NN connect, aiming at mapping pipelines onto it. If the pipes have forks, which otherwise would require a 2-D realization, additional routing resources are needed, like longer lines spanning the whole or a part of the array, often being segmented. Two RAs have linear array structure. RaPiD [27] provides different computing resources, like ALUs, RAMs, multipliers, and registers, but irregularly distributed. While RaPiD uses mostly static reconfiguration, PipeRench relies on dynamic reconfiguration, allowing the reconfiguration of a PE in each execution cycle. Besides the mostly unidirectional NN connects, it provides also a global bus.

RaPiD (Reconfigurable Pipelined Datapath) [27] aims at speed-up of highly regular, computation-intensive tasks by deep pipelines on its 1-D RA. RaPiD-1 features 15 DPUs of 8 bit with integer multiplier (32 bit output), 3 integer ALUs, 6 general-purpose datapath registers and 3 local 32 word memories, all 16 bits wide. ALUs can be chained. Each memory has a special datapath register with an incrementing feedback path. To implement I/O streams RaPiD includes a stream generator with address generators, optimized for nested loop structures, associated with

FIFOs. The address sequences for the generators are determined at compile-time. RaPiD’s routing and configuration architecture consists of several parallel segmented 16 bit buses, which span the whole array. The length of the bus segments varies by tracks. In some tracks, adjacent bus segments can be merged. The sophisticated interconnect fabric cannot be detailed here.

PipeRench [29], an accelerator for pipelined applications, provides several reconfigurable pipeline stages (“stripes”) and relies on fast partial dynamic pipeline reconfiguration and run time scheduling of configuration streams and data streams. It has a 256 by 1024 bit configuration

Paradigm	Platform	Programming source
“von Neumann”	Hardware	Software
Xputer	coarse grain Flexware	high level Configware
RL (FPGA etc.)	fine grain Flexware	netlist level Configware

Fig. 6: About terminology.

memory, a state memory (used to save current register contents of a stripe), an address translation table (ATT), four data controllers, a memory bus controller and a configuration controller. The reconfigurable fabric allows the configuration of a pipeline stage in every cycle, while concurrently executing all other stages. The fabric consists of (horizontal) stripes composed of interconnect and PEs with registers and ALUs, implemented as 3-input lookup tables. The ALU includes a barrel shifter, carry chain circuitry, etc. A stripe provides 32 ALUs with 4 bits each. The whole fabric has 28 stripes. The interconnect scheme features local interconnect inside a stripe as well as local and global interconnect between stripes and four global buses.

2.3 Crossbar-Based Architectures

A full crossbar switch, a most powerful communication network is easily to rout. But 2 of the RAs of this category use only reduced crossbars. PADDI for the fast prototyping of DSP datapaths features eight PEs, all connected by a multilayer crossbar. PADDI-2 has 48 PEs, but saves area by restricted crossbar with a hierarchical interconnect structure for linear arrays of PEs forming clusters. The sophistication of these fabrics has again an impact on routing.

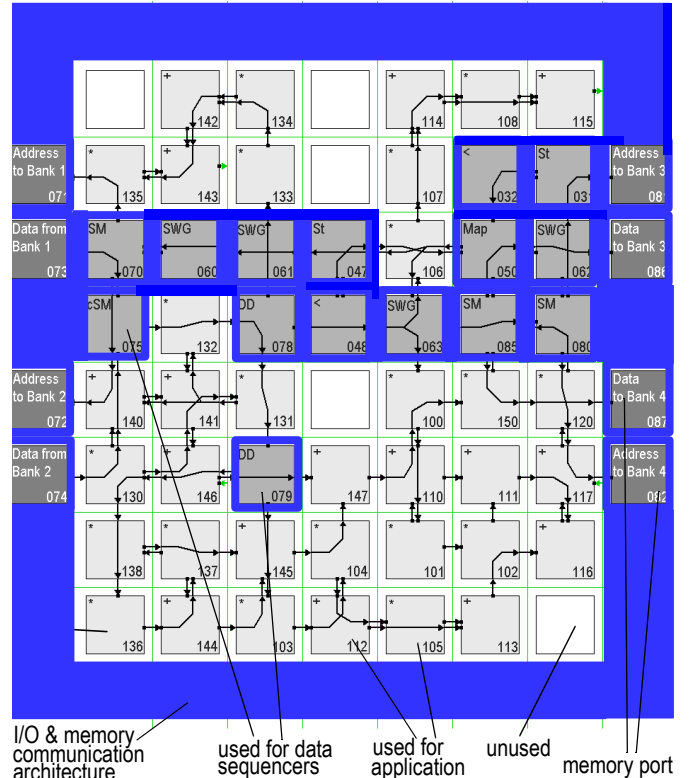


Fig. 7: Mapping application (linear filter) and memory communication architecture (dark background) onto the same KressArray, including the address ports and the data ports to 4 different memory banks (5 of 8 memory port connects are routed through application DPUs).

PADDI-1 (Programmable Arithmetic Device for DSP) [30] [31], for rapid prototyping of computation-intensive DSP data paths, consists of clusters of 8 arithmetic execution units (EXUs) 16 bits wide, including 8 word SRAM (which may be concatenated for 32 bits) and connected to a central crossbar switch box. Interconnect is organized in 2 levels: a 3 bit global bus to broadcast global instructions to each EXU's controller CTL, decoded locally at second level into a 53 bit instruction word.

The PADDI-2 Architecture [32] features a data-driven execution mechanism. Although the basic architectural principles of the original PADDI were kept, the PADDI-2 has several differences. It has 48 EXUs. Each PE features a 16 bit ALU also including Booth multiply and select, multiplication instructions taking 8 cycles on a single processor. Alternatively, eight EXUs can be pipelined, resulting in a single cycle multiplication. PADDI-2 EXUs are packed in 12 clusters of four elements each two level interconnect: six 16 bit intra-cluster data buses (plus one bit control buses, and, inter-cluster 16 data buses, which can be broken up into shorter segments).

The Pleiades Architecture [33] is a generalized low power "PADDI-3" with programmable microprocessor and heterogeneous RA of EXUs, which allows to integrate both fine and coarse grained EXUs, and, memories in place of EXUs. For each algorithm domain (communication, speech coding, video coding), an architecture instance can be created (with known EXU types and numbers). Communication between EXUs is dataflow driven. The control means available to the programmer are basic EXU configurations to specify its operation, and interconnect configurations to build EXU clusters. All configuration registers are part of the processor's memory map and configuration codes are processor's memory writes.

2.4 Memory Bandwidth Problems

RAM cycle time as throughput bottleneck of von-Neumann-style computing, known as "memory bandwidth problem", a communication gap which spans up to 2 orders of magnitude, can be even wider in RA use as accelerator in data-intensive applications - a challenge to RA architectures and mappers (see § "DTSE"). Interfaces like RAMbus etc. may reduce the gap a little bit for both, host and accelerator. But the use of caches, mainly based on instruction loops (write once - multiple read) does not help in RA usage: compare fig. 14. That's why RA architectures should be capable to interface multiple memory banks, only known from the KressArray (example: fig. 7, where 8 ports communicate with 4 memory banks [34]). Also see chapter 4 and 5.

2.5 Future Reconfigurable Architectures

A universal RA obviously is an illusion. The way to go is toward ASPPs (application-specific programmable products) like sufficiently flexible RAs, optimized for a particular application domain like e. g. wireless communication, image processing or multimedia etc. There is a need for tools supporting such dedicated RA architecture development. But architectures have an immense impact on implementability of good mapping tools. "Clever" fabrics are too sophisticated to find good tools. Best are simple generic fabrics architecture principles, or, a mapping tool which generically creates by itself the architectures it can manage easily [11], or, a combination of both approaches like the platform space exploration (s. „The KressArray Family“ and "Xplorer").

3. Programming Coarse Grain RAs

Programming frameworks for RAs (also see fig. 1) are highly dependent on structure and granularity, and differ by language level. For MorphoSys, MATRIX, PADDI-2 and REMARC it's assembler level. Some support the designer by a graphical tool for manual P&R. Others feature automatic design flow from HDL or high-level programming language. Environments differ by the approach used for technology mapping, placement, routing. Using only a simple script for technology mapping [35] DP-FPGA [4] is not considered.

Technology mapping is mostly simpler for coarse grain architectures than for FPGAs. Approaches are: direct mapping, where the operators are mapped straight forward

onto PEs, with one PE for one operator, or, using an additional library of functions not directly implementable by one PE, or, more sophisticated tree matching also capable to merge several operators into one PE by a modified FPGA tool kit. An exception is the RAW compiler doing partitioning instead of technology mapping, since RAW has RISC cores as PEs accepting blocks from program input.

For operator placement, the architecture has an impact. An approach often used for FPGAs synthesis is placement by simulated annealing or a genetic algorithm. Garp uses a tree matching algorithm instead, where placement is done together with technology mapping. The use of greedy algorithms is feasible only for linear arrays (PipeRench), or with a high level communication network (RAW). PADDI is an exception by using a scheduling algorithm for resource allocation.

Routing also features quite different approaches. In two cases, the routing is not done in an extra phase but integrated into the placement and done on the fly. One approach (KressArray) uses a simple algorithm restricted to connects with neighbours and targets with at most the distance of one. The other (RaPiD) employs the pathfinder algorithm [36], which has been developed for FPGA routing. Greedy routing would be not satisfying. General exceptions to the routing approaches is the RAW architecture, which features only one high-level communication resource, so no selection of routing resources is needed, and the PADDI architecture, which features a crossbar switch having the same effect. Greedy routing algorithms are only used for 1-D RAs, or architectures capable to cure routing congestion by other mechanisms, like Colt with wormhole run-time reconfiguration.

3.1 Assembler Programming

Assembler level code for coarse grain architectures can be compared to configuration code for FPGAs. In the case of systems comprising a microprocessor / RA symbiosis, only the reconfigurable part is considered for the classification. Programming is done mainly at a kind of assembler level for PADDI-2, MATRIX, and, RAs of REMARC and MorphoSys.

For Programming PADDI-2 [32], a tool box has been developed which includes software libraries, a graphical interface for signal flow graphs, routing tools, simulation tools, compilation tools and tools for board access and board debugging. Major parts of this process are done manually. The input specifies assembly code for each function in the signal flow graph. The programmer manually partitions the signal flow graph with a graphical tool, which also aids in manual placement and routing. As an alternative to manual placement and routing, an automated tool is provided, which guarantees to find a mapping, if one exists, by exhaustive methods which need much computation time.

For Programming MATRIX [15] an assembly level macro language has been developed. Some work on P&R pointed out the original MATRIX's weak points [37].

REMARC tools [18] allow concurrent programming of the RISC processor (by C using the GCC compiler) and the RA by adding REMARC assembler instructions. The compiler then generates assembly code for the RISC processor with the REMARC assembler instructions embedded which are further processed by a special REMARC assembler generating binary code for the REMARC instructions. Finally, the GCC compiler is used again to generate RISC instruction code to invoke REMARC and its instructions embedded as binary data.

Programming MorphoSys is supported by a SUIF-based compiler [19] for host, and development tools for RA. Host / RA partitioning is done manually by adding a prefix to functions to be mapped onto RA. The compiler generates TinyRISC code for RA activation. Configuration code is generated via a graphical user interface or manually from an assembler level source also usable to simulate the architecture from VHDL.

3.2 Frameworks with FPGA-Style Mapping

Computation-intensive algorithms for mapping onto FPGAs are well-known and can often be used directly for

mapping	Kress DPSS	CHESS	RaPiD	Colt
placement	simulated annealing	simulated annealing		genetic algorithm
routing		Pathfinder		greedy algorithm

Fig. 8: FPGA-Style Mapping for coarse grain reconfigurable arrays.

coarse grain architectures like for CHESS, Colt, KressArray, RaPiD (see fig. 8). All four use simulated annealing or other genetics for placement, and two use pathfinder for routing [36]. The KressArray DPSS (Datapath Synthesis System) accepts a C-like language source. The compilation framework for the RaPiD system works similar, but relies on relatively complex algorithms. Colt tools use a structural description of the dataflow. CHESS has been programmed from a hardware description language (JHDL) source. P&R quality has a massive impact on application performance. But, due to the low number of PEs, P&R is much less complex than for FPGAs and computational requirements are drastically reduced.

DPSS. (DataPath Synthesis System) [5] generates configuration code for KressArrays from ALE-X high-level language sources [5] [38] supporting datapaths with assignments, local variables and loops. After classical optimizations it generates an expression tree. Next processing steps include a front end, logic optimization, technology mapping creating a netlist, simultaneous P&R by simulated annealing, and I/O scheduling (incl. loop folding, memory cycle optimization, register file usage). The result is the application's KressArray mapping and array I/O schedule. Finally configuration binaries are assembled. Routing is restricted to direct NN connect and rout-through of length 1. Other connect is routed to buses or segmented buses. DPSS has also been part of the MoM-3 Xputer compiler accepting and partitioning a subset of C subset into sequential MoM code and structural KressArray code. The more general CoDe-X approach [39] uses this MoM compiler as part of a partitioning co-compiler accepting a C language superset and partitioning the application onto the host and one or several Xputer-based accelerators.

Tools for Colt. [12] accept a dataflow description (below C level) for placement by a genetic algorithm and routing by a greedy algorithm (routing congestion is cured at run-time by wormhole reconfiguration). Data stream headers hold configuration data for routing and the functionality of all PEs encountered.

Programming RaPiD [27] is done in RaPiD-C, a C-like language with extensions (like synchronization mechanisms and conditionals to identify first or last loop iteration) to explicitly specify parallelism, data movement and partitioning. RaPiD-C programs may consist of several nested loops describing pipelines. Outer loops are transformed into sequential code for address generators, inner loops into structural code for the RA. The compilation steps are: netlist generation from structural code, extraction of dynamic control, generation of controller code instruction streams for dynamic control and generation of I/O configuration data for the stream units. The netlist is mapped onto RaPiD by pipelining, retiming, and P&R. Placement is done by simulated annealing, with routing (by pathfinder [36]) done on the fly to measure placement quality [40].

For Programming the CHESS array [20] a compiler [41] has been implemented accepting JHDL [42] sources and generating CHESS netlists. Placement is done by simulated annealing and routing by Pathfinder's negotiated congestion algorithm [36]. Part of the work is not disclosed.

machine category	Computer ("v. Neumann")	Xputer [8] (no transputer!)
machine paradigm	procedural sequencing: deterministic	(no dataflow [13])
driven by:	control flow	data stream(s)
RA support	no	yes
engine principles	instruction sequencing	data sequencing
state register	program counter	(multiple) data counter(s)
communication path set-up	at run time	at load time
data path	resource operation	single ALU sequential
		array of ALUs parallel

Fig. 9: Machine Paradigms.

3.3 Other mapping approaches

Greedy algorithms are poor in mapping to FPGAs. But, although Garp is mesh-based, mapping treats it like a linear array which allows mapping in one step by a simple greedy routing algorithm. RAW features only one communication resource, removing the wire selection problem from routing. Instead, the compiler schedules time multiplexed NN connections. CPU cores inside RAW PEs simplify mapping by loading entire source code blocks. PipeRench resembling a linear array and interconnect fabrics restrictions keep placement simple for a greedy algorithm. PADDI uses a standard P&R approach.

Garp tools [16] use a SUIF-based C compiler [43] to generate code for the MIPS host with embedded RA configuration code to accelerate (only non-nested) loops. At next basic blocks are generated and converted into hyperblocks containing a contiguous group of basic blocks, also from alternative control paths. Control flow inside a hyperblock is converted for predicated execution. Blocks, which cannot be mapped, are removed from hyperblocks. The resulting reduced hyperblock is then the basis for mapping. The next step generates interfacing instructions for the host, and transforms the hyperblock into a DFG (data flow graph). The proprietary Gamma tool [44] maps the DFG onto Garp using a tree covering algorithm which preserves the datapath structure, supports features like the carry chains. Gamma first splits the DFG into subtrees and then matches subtrees with module patterns which fit in one Garp row. During tree covering, the modules are also placed in the array. After some optimizations the configuration code is generated (incl. outing [46]), assembled into binary form, and, linked with the hosts C object code.

RAW tools [47] [48] include a SUIF-based C compiler and a run-time system managing dynamic mechanisms like branch prediction, data caching [49], speculative execution, dynamic code scheduling. The compiler handles resource allocation, parallelism exploitation, communication scheduling, code generation (for host and switch processors in each tile), and, divides execution into coarse-grain parallel regions internally communicating by static network, whereas intra-region communication uses messages. The phases are: pointer analysis [47] with data dependency analysis and memory access generation (if known at compile-time), partitioning application data for distributed memory; space-time scheduling [48] for parallelization; address translation for caching by software. RAW binary is generated by the MIPS compiler back end. The RAW project aims more at parallel processing rather than reconfigurable computing and failed in finding a good automatic mapping algorithm [50].

PipeRench tools [29] [51] use the DIL single-assignment language (SAL) for design entry and as an intermediate form. First, the compiler inlines all modules, unrolls loops and generates a straight-line SA program (SAP). After optimizations and breaking the SAP into pieces fitting on one stripe, a greedy P&R algorithm is run which tries to add nodes to stripes. Once placed, a node is routed and never moved again. P&R is fast by crossbar switch usage, coarse granularity, and, restriction to unidirectional pipelines.

CADDI [52], assembler and simulator, has been implemented for PADDI. First a silage [53] specification is compiled into a CDFG (control /data flow graph), used for estimations of critical path, minimum and maximum bounds for hardware for a given time allocation, minimum bounds of execution time, and for transformations like pipelining, retiming, algebraic transformations, loop unrolling and




Fig. 10: Machine paradigms: a) v. Neumann, b) Xputer, c) parallel Xputer

Explorer System	year	source	inter-active	status evaluation	status generation
DPE	1991	[66]	no	abstract models	rule-based
Clio	1992	[67]	yes	prediction models	advice generator
DIA	1998	[68]	yes	prediction fr. library	rule-based
DSE for RAW	1998	[49]	no	analytical models	analytical
ICOS	1998	[76]	no	fuzzy logic	greedy search
DSE f. Multimedia	1999	[77]	no	simulation	branch and bound
Xplorer	1999	[11][50]	yes	fuzzy rule-based	simulated annealing

Fig. 11: Design Space Exploration Systems.

operation chaining, CDFG to architecture technology mapping is straight-forward since all components are fixed and routing through crossbars is efficient. If several PADDI clusters are involved, a partitioning step comes before resource allocation, assignment, and scheduling. The assignment phase maps operations to EXUs by a rejectionless antivoter algorithm [54].

Compilation for Pleiades. Because of the complex design space created by the heterogeneous architecture, the application mapping problem is not yet solved completely.

C~SIDE development tools for the Chameleon CS2000 family include a GNU C compiler for the RISC host, a HDL synthesizer for the reconfigurable fabric, a simulator, a C-style debugger and a verifier. eBIOS (eConfigurable Basic I/O Services), a kind of operating system, interfaces the RISC processor with the reconfigurable fabric. C~SIDE is a tool box, rather than a co-compiler. About mapping or compilation for the MECA family [24], FIPSOC [26] and MorphICs RA no information is available.

IDE. (Integrated Development Environment) with a C-compiler, debugger, simulator and “Evaluation Module” (EVM) serves for CALISTO from Silicon Spice [25], where a Real-time operating system (RTOS) supports “Any Service Any Port” (ASAP) configurations for up to 240 channels of carrier class G.711 VoIP (voice over IP). IDE is a tool box, but no co-compiler.

3.4 Run-time Mapping

The VIRTEX FPGA family from Xilinx, the RAs being part of the CS2000 series systems from Chameleon and others are run-time reconfigurable. Programming a host/RA combination is a kind of H/S Co-design. However using such devices changes many of the basic assumptions in the HW/SW co-design process: host / RL interaction is dynamic and needs a kind of tiny operating system like eBIOS, also to organize RL reconfiguration under host control. A typical goal is minimization of reconfiguration latency (especially important in communication processors), to hide configuration loading latency, and, list scheduling to find ‘best’ schedule for a series of eBIOS calls (also see § “C~SIDE”). For more about typical aspects of run-time reconfiguration see [55] [56].

3.5 Retrospective

The history of silicon synthesis and application distinguishes three phases [57] (or, 3 Makimoto waves [58]): hardware design (fig. 12 a), introduction of the microcontroller (fig. 12 b), and

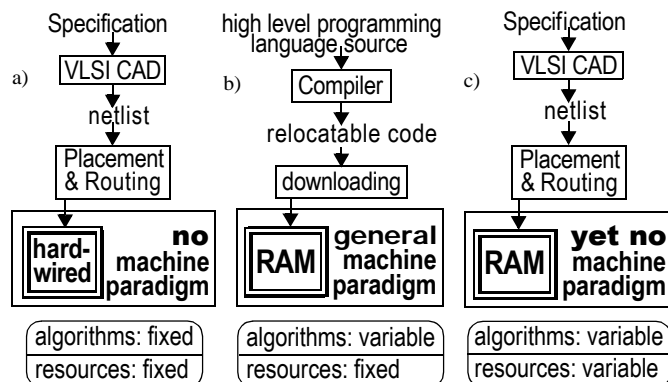


Fig. 12: Synthesis a) hardwired, b) “von Neumann”, c) reconfigurable.

RA usage (fig. 12 c). The transition from phase 1 (structural synthesis) to phase 2 brought a shift from net-list-based CAD (fixed algorithms, no machine paradigm, infinite design space) to RAM-based (procedural) synthesis by compilation, based on a machine paradigm, which reduces the design space by guidance. Note: RAM-based means flexibility and fast turn-around. Accelerator synthesis (fig. 15 b) still uses phase 1 methods (CAD). The third phase introduces reconfigurable hardware, i. e. RAM-based structural synthesis: resources have become variable. But still phase 1 synthesis tools are mainly used: versions of CAD tools. Since structural synthesis has become RAM-based it is time to switch to real compilation techniques, based on a soft machine paradigm. But the R&D scene ignore, that we now have a dichotomy of RAM-based programming: procedural programming versus structural programming, integrating two worlds of computing (fig. 16)

4. Compilation Techniques

“von Neumann” and the classical compiler are obsolete (fig. 15 a). Today, host/accelerator(s) symbiosis is dominant (fig. 15 b) and most of the platforms summarized above make use of it. Newer commercial platforms include all on a single chip. E. g. Altera’s EXCALIBUR combines a core processor (ARM, or MIPS), embedded memory and RL. Sequential code is downloaded to the host’s RAM. But accelerators are still implemented by CAD, a C compiler is only an isolated minor tool, and, *software / configure partitioning is still done manually*. [44] [51] [59] [60] [62], so that massive hardware expertise is needed to implement accelerators.

4.1 Co-Compilation

Using RAs as accelerators again changes this scenario: now implementations onto both, host and RA(s) are RAM-based, which allows turn-around times of minutes for the entire system, instead of months needed for hardwired accelerators. This means a change of market structure by migration of accelerator implementation from IC vendor to customer, demanding automatic compilation from high level programming language sources onto both, host and RA: *co-compilation* including automatic software / configure partitioning. (fig. 15 c). Since compilers are based on a machine paradigm and “v. Neumann” does not support soft datapaths (because “instruction fetch” is not done at run time: fig. 14) we need a new paradigm (Xputer [62]) for the RA side, where the program counter (fig. 10 a) is replaced by a data counter (*data sequencer* [63]: fig. 10 b). Figure 9 compares both paradigms. With multiple data sequencers (fig. 10 c) a single Xputer may even handle several parallel data streams (example in fig. 7).

CoDe-X is the first such co-compilation environment having been implemented ([39] fig. 13), which partitions mainly by identifying loops suitable for parallelizing transformation [3] [39] [61] into code downloadable to the MoM accelerator Xputer.

C~SIDE. Chameleon Systems reports for its CS2000 series co-compilation [23] techniques, combining compiler optimization, multithreading to hide configuration loading latency, and, list scheduling to find ‘best’ schedule for eBIOS calls (see § “C~SIDE”). Whether automatic partitioning is used is undisclosed.

The Xputer Machine Paradigm for soft hardware (fig. 9) [7] [8] [9] [10]. is the diametral counterpart of the von Neumann paradigm: the new “Xputer” machine paradigm. Instead of a “control flow” sublanguage a “data stream” sublanguage like MoPL [64] recursively defines *data goto*, *data jumps*, *data loops*, *nested data loops*, and *parallel data loops*.




Fig. 13: CoDe-X Co-Compiler




Fig. 14: “Instruction Fetch”.




Fig. 15: Computing Platforms: a) “v. Neumann”, b) current, c) emerging.

5. Design Space Explorers (DSEs)

Some development environments aim beyond compilation. DSEs (fig. 11 [50]) select one of many alternative solutions to meet a design goal meeting constraints or desired properties to optimize a *design* or a (by PSE) *programmable platform*. Guidance systems or design assistants are interactive DSEs giving advice during the design flow. Some DSEs avoid the status generation and provide only predictions etc. from a knowledge data base. Advanced DSEs provide status generation, e.g. by expert system, and present advice like a choice of proposals. Non-interactive DSEs automatically generate a solution status from rule-based knowledge or fuzzy learning.

5.1 Design Space Exploration

Interactive design assistants are DPE and Clio (both for VLSI) and DIA. Including effect predictors and proposal generators DPE (Design Planning Environment) [66] using an expert system, Clio [67] using a hierarchy of templates, and DIA (Datapath-Intensive ASICs) [68], targeting semi-custom ASIC behavioural level and based on encapsulated expert knowledge, generate a design flow by creating a schematic, a data flow graph, or a layout from a specification and area, cycle time, power, e.a. constraints and to improve area, power, throughput etc.

DTSE. For data-dominated systems the usual local optimization techniques lead to performance-degrading runtime solutions of access conflicts. A cost overhead of 10 - 100% (in power) for hardware and around 35% (in clock rate) for software has been estimated [69] [70]. For global exploration the use of conflict-directed ordering (CDO) [71] as an extension of force-directed scheduling (FDS) [72] has been proposed [73]. Instead of working on a signal access flow graph (SAFG) [71] a multi-dimensional conflict graph (MD-CG) is used for a generalized CDO (G-CDO) algorithm for the ATOMIUM / ACROPOLIS data transfer and storage exploration (DTSE) system [74] [75].

5.2 Platform Space Explorers (PSEs)

A PSE serves to find an optimum RA or processor array (PA) platform for an application domain by optimizing array size, path width, processor's MIPS, number of ALUs and branch units, local SRAM size, data and instruction cache sizes, local bandwidth, interconnect latency etc. from requirements like chip area, total computation, memory size, buffer size, communication etc. Software or configware application programming is finally not part of exploration, but may serve platform evaluation. All three being non-interactive, the *DSE [49] for RAW* [17] featuring an analytical model, *ICOS* (Intelligent Concurrent Object-oriented Synthesis) [76] featuring object-oriented fuzzy techniques, and “*DSE for Multimedia Processors*” [77] (DSEMMP) aim at automatic synthesis of a multiprocessor platform from system descriptions, performance constraints, and a cost bound and generate an architecture.. DSEMMP aims at shared memory with intel Strong-ARM SA-110 as a starting point.

5.3 Compiler / PSE symbiosis

Since to map an application onto a coarse grain RA may take only minutes, retargettable mappers or compilers may be also used for platform exploration. By profiling the results of the same application or benchmark on different platforms may be compared. Such a compiler / PSE symbiosis like in *Xplorer* provides direct verification and yields more realistic and more precise results than explorers using abstract models for estimation and gives better support for manual tuning.

Xplorer, an interactive PSE framework [11] [50] has been implemented around a modified DPSS mapper [5]. This universal design space exploration environment supports both, optimum architecture selection (e.g. domain-specific) and application development onto it and includes several tools: *architecture editor* (to edit communication resources and annealing parameters), *mapping editor* (to change I/O port type, freeze locations of edge port, cell or cell group etc.), *instruction mapper* to edit and define the operator repertoire, *architecture suggestion generator* [78], *HDL generator* for cell simulation, *retargettable cell layout generator* (planned, similar to [79]), *power estimator* (planned [80], using methods from [82]). A cycle through an exploration loop usually takes only minutes, so that a number of alternative architectures may be checked in a reasonable time. By mapping the application onto it verification is provided directly.

Memory bandwidth. The Xplorer also yields solutions to the memory bandwidth problem [34] by supporting mixed rDPU types, so that both, data sequencers and rDPUs dedicated to the application can be mapped onto the same KressArray what is illustrated by the example in fig. 7 (also see section 2.4). These Xplorer capabilities provide a straight-forward approach to support architectural implementations of the Xputer soft machine paradigm.

5.4 Parallel Computing vs. Reconfigurable

RISC core IP cells are available (e.g. from ARM) so small, that 32 (soon 64 or more) of them would fit onto a single chip to form a massively parallel computing system. But this is not a general remedy for the parallel computing crisis [83], indicated by rapidly shrinking supercomputing conferences. For many application areas process level parallelism yields only poor speed-up improvement per processor added. Amdahls law explains just one of several reasons of inefficient resource utilization. A dominating problem is the instruction-driven late binding of communication paths (fig. 14), which often leads to massive communication switching overhead at run-time. R&D in the past has largely ignored, that the so-called “von Neumann” paradigm is not a communication paradigm. However, some methods from parallel computing and parallelizing compiler R&D scenes may be adapted to be used for lower level parallelism on RA platforms (compare § “Co-Compilation”).

6. Conclusions

Exploding design cost and shrinking product life cycles of ASICs create a demand on RA usage for product longevity. Performance is only one part of the story. The time has come fully exploit their flexibility to support turn-around times of minutes instead of months for real time in-system debugging, profiling, verification, tuning, field-maintenance, and field-upgrades. A new “soft machine” paradigm and language framework is available for novel compilation techniques to cope with the new market structures transferring synthesis from vendor to customer.




Fig. 16: Computing Science Dichotomy.

Reconfigurable platforms and their applications are heading from niche to main-stream, bridging the gap between ASICs and micro-processors. Many system-level integrated future products without reconfigurability will not be competitive. Instead of technology progress better architectures by RA usage will be the key to keep up the current innovation speed beyond the limits of silicon. It is time to revisit past decade R&D results to derive commercial solutions: at least one promising approach is available. It is time for you to get involved. Theory and backgrounds are ready for creation of a dichotomy of computing science (fig. 16) for curricular innovations urgently needed.

7. Literature

1. R. Hartenstein, H. Grünbacher (Editors): The Roadmap to Reconfigurable computing - Proc. FPL2000, Aug. 27-30, 2000; LNCS, Springer-Verlag 2000
2. A. DeHon: Reconfigurable Architectures for General Purpose Computing; report no. ATR 1586, MIT AI Lab, 1996
3. R. Hartenstein: The Microprocessor is no more General Purpose (invited paper), Proc. ISIS'97, Austin, Texas, USA, Oct. 8-10, 1997.
4. D. Cherepacha and D. Lewis: A Datapath Oriented Architecture for FPGAs; Proc. FPGA '94, Monterey, CA, USA, February 1994.
5. R. Kress et al.: A Datapath Synthesis System for the Reconfigurable Datapath Architecture; ASP-DAC'95, Chiba, Japan, Aug. 29 - Sept. 1, 1995
6. H. Reinig: A Scalable Architecture for Custom Computing; Ph.D. Thesis, Univ. of Kaiserslautern, Germany, July 1999.
7. R. Hartenstein, A. Hirschbiel, M. Weber: MoM - a partly custom-design architecture compared to standard hardware; IEEE CompEuro 1989
8. R. Hartenstein et al.: A Novel Paradigm of Parallel Computation and its Use to Implement Simple High Performance Hardware; InfoJapan'90, 30th Anniversary of Computer Society of Japan, Tokyo, Japan, 1990.
9. R. Hartenstein et al.: A Novel ASIC Design Approach Based on a New Machine Paradigm; IEEE J.SSC, Volume 26, No. 7, July 1991.
10. (invited reprint of [8]) Future Generation Computer Systems 7 91/92, p. 181-198, (Elsevier Scientific).
11. U. Nageldinger et al.: KressArray Explorer: A New CAD Environment to Optimize Reconfigurable Datapath Array Architectures; ASP-DAC, Yokohama, Japan, Jan. 25-28, 2000.
12. R. A. Bittner et al.: Colt: An Experiment in Wormhole Run-time Reconfiguration; SPIE Photonics East '96, Boston, MA, USA, Nov. 1996.
13. D. Gajski et al.: A second opinion on datapath machines; Computer, Feb '82
14. K. Hwang: Advanced Computer Architecture; McGraw-Hill, 1993.
15. E. Mirsky, A. DeHon: MATRIX: A Reconfigurable Computing Architecture with Configurable Instruction Distribution and Deployable Resources; Proc. IEEE FCCM'96, Napa, CA, USA, April 17-19, 1996.
16. J. Hauser and J. Wawrzyniek: Garp: A MIPS Processor with a Reconfigurable Coprocessor; Proc. IEEE FCCM'97, Napa, April 16-18, 1997.
17. E. Waingold et al.: Baring it all to Software: RAW Machines; IEEE Computer, September 1997, pp. 86-93.
18. T. Miyamori and K. Olukotun: REMARC: Reconfigurable Multimedia Array Coprocessor; Proc. ACM/SIGDA FPGA'98, Monterey, Feb. 1998.
19. H. Singh, et al.: MorphoSys: An Integrated Re-configurable Architecture; Proceedings of the NATO RTO Symp. on System Concepts and Integration, Monterey, CA, USA, April 20-22, 1998.
20. A. Marshall et al.: A Reconfigurable Arithmetic Array for Multimedia Applications; Proc. ACM/SIGDA FPGA'99, Monterey, Feb. 21-23, 1999
21. J. Becker et al.: Architecture and Application of a Dynamically Reconfigurable Hardware Array for Future Mobile Communication Systems; Proc. FCCM'00, Napa, CA, USA, April 17-19, 2000.
22. <http://www.chameleonsystems.com>
23. X. Tang, et al.: A Compiler Directed Approach to Hiding Configuration Loading Latency in Chameleon Reconfigurable Chips; in [1]
24. <http://www.malleable.com>
25. <http://www.broadcom.com/siliconspice.html>
26. <http://www.sidsa.com>
27. C. Ebeling et al.: „RaPiD: Reconfigurable Pipelined Datapath“, in [28]
28. M. Glesner, R. Hartenstein (Editors): Proc. FPL'96, Darmstadt, Germany, Sept. 23-25, 1996, LNCS 1142, Springer Verlag 1996
29. S. C. Goldstein et al.: PipeRench: A Coprocessor for Streaming Multimedia Acceleration; Proc. ISCA'99, Atlanta, May 2-4, 1999
30. D. Chen and J. Rabaey: PADDI: Programmable arithmetic devices for digital signal processing; VLSI Signal Processing IV, IEEE Press 1990.
31. D. C. Chen, J. M. Rabaey: A Reconfigurable Multiprocessor IC for Rapid Prototyping of Algorithmic-Specific High-Speed DSP Data Paths; IEEE J. Solid-State Circuits, Vol. 27, No. 12, Dec. 1992.
32. A. K. W. Yeung, J.M. Rabaey: A Reconfigurable Data-driven Multiprocessor Architecture for Rapid Prototyping of High Throughput DSP Algorithms; Proc. HICSS-26, Kauai, Hawaii, Jan. 1993.
33. J. Rabaey: Reconfigurable Computing: The Solution to Low Power Programmable DSP; Proc. ICASP'97 Munich, Germany, April 1997.
34. M. Herz: High Performance Memory Communication Architectures for Coarse-Grained Reconfigurable Computing Architectures; Ph. D. Dissertation, Universität Kaiserslautern, January 2001
35. D. Lewis: Personal Communication, April 2000.
36. C. Ebeling et al.: Placement and Routing Tools for the Trypthich FPGA; IEEE Trans VLSI Systems 3, No. 4, December 1995.
37. A. DeHon: Personal Communication, February 2000.
38. T. Molketin: Analyse, Transformation und Verteilung arithmetischer und logischer Ausdrücke; Diplomarbeit, Univ. Kaiserslautern, 1995
39. J. Becker: A Partitioning Compiler for Computers with Xputer-based Accelerators; Ph. D. dissertation, Kaiserslautern University, 1997.
40. C. Ebeling: Personal Communication, March 2000.
41. A. Marshall: Personal Communication; February 2000.
42. B. Hutchings, B. Nelson: Using General-Purpose Programming Languages for FPGA Design; Proc. DAC 2000, Los Angeles, June 2000
43. M. W. Hall et al.: Maximizing Multiprocessor Performance with the SUIF Compiler; IEEE Computer, Dec. 1996
44. T. J. Callahan and J. Wawrzyniek: Instruction-Level Parallelism for Reconfigurable Computing; in [45] pp. 248-257.
45. R. Hartenstein, A. Keevallik (Editors): Proc. FPL'98, Tallinn, Estonia, Aug. 31- Sept. 3, 1998, LNCS, Springer Verlag, 1998
46. J. Hauser: Personal Communication, March 2000.
47. R. Barua et al.: Maps: A Compiler-Managed Memory System for RAW Machines; Proc. ISCA'99, Atlanta, USA, June, 1999.
48. W. Lee et al.: Space-Time Scheduling of Instruction-Level Parallelism on a RAW Machine; Proc. ASPLOS'98, San Jose, Oct. 4-7, 1998.
49. C. A. Moritz et al.: Hot Pages: Software Caching for RAW Microprocessors; MIT, LCS-TM-599, Cambridge, MA, Aug. 1999.
50. U. Nageldinger: Design-Space Exploration for Coarse Grained Reconfigurable Architectures; Dissertation, Universität Kaiserslautern, Febr. 2001
51. M. Budiu and S. C. Goldstein: Fast Compilation for Pipelined Reconfigurable Fabrics; Proc. FPGA '99, Monterey, Feb. 1999, pp. 135-143.
52. D. Chen et al.: An Integrated System for Rapid Prototyping of High Performance Data Paths; Proc. ASAP'92, Los Alamitos, Aug. 4-7, 1992
53. P. H. Hilfinger: A High-Level Language and Silicon Compiler for Digital Signal Processing; Proc. 1985 IEEE CICC., Portland, May 20-23, 1985.
54. M. Potkonjak, J. Rabaey: A Scheduling and Resource Allocation Algorithm for Hierarchical Signal Flow Graphs; Proc. DAC'89, Las Vegas, June 1989
55. J. Noguera, R. Badia: A HW/SW Partitioning Algorithm for Dynamically Reconfigurable Architectures; Proc. DATE 2001
56. J. Noguera, R. Badia: Run-time HW/SW Codesign for Discrete Event Systems using Dynamically Reconfigurable Architectures; Proc. ISSS 2000 (Int'l Symp. System Synthesis), Madrid, Spain, Sept. 20 - 22, 2000
57. N. Tredennick: Technology and Business: Forces Driving Microprocessor Evolution; Proc. IEEE 83, 12 (Dec. 1995)
58. T. Makimoto: The Rising Wave of Field-Programmability; in: [1]
59. M. Weinhardt, W. Luk: Pipeline Vectorization for Reconfigurable Systems; Proc. IEEE FCCM, April 1999
60. M. Gokhale, J. Stone: NAPA C: Compiling for a hybrid RISC / FPGA architecture; Proc. IEEE FCCM April 1998
61. K. Schmidt: A Program Partitioning, Restructuring, and Mapping Method for Xputers; Ph.D. Thesis, University of Kaiserslautern 1994
62. J. Becker et al.: A General Approach in System Design Integrating Reconfigurable Accelerators; Proc. IEEE ISIS'96; Austin, TX, Oct. 9-11, 1996
63. M. Herz, et al.: A Novel Sequencer Hardware for Application Specific Computing; Proc. ASAP'97, Zurich, Switzerland, July 14-16, 1997
64. A. Ast, J. Becker, R. Hartenstein, R. Kress, H. Reinig, K. Schmidt: Data-procedural Languages for FPL-based Machines; in [65]
65. R. Hartenstein, M. Servit (Editors): Proc. FPL'94, Prague, Czech Republic, Sept. 7-10, 1994, LNCS, Springer Verlag, 1994
66. D. Knapp & al.: The ADAM Design Planning Engine, IEEE Trans CAD, July 1991
67. J. Lopez et al.: Design Assistance for CAD Frameworks; Proc. EURO-DAC'62, Hamburg, Germany, Sept. 7-10, 1992
68. L. Guerra et al.: A Methodology for Guided Behavioural Level Optimization; Proc. DAC'98, San Francisco, June 15-19, 1998
69. A. Vandecapelle et al.: Global Multimedia Design Exploration using Accurate Memory organization Feedback; Proc. DAC 1999
70. T. Omnès et al.: Dynamic Algorithms for Minimizing Memory Bandwidth in High throughput Telecom and Multimedia; in: Techniques de Parallelization Automatique, TSI, Editions Hermès, 1999
71. S. Wuytack et al.: Minimizing the required Memory Bandwidth in VLSI System Realizations; IEEE Trans. VLSI Systems, Dec. 1999
72. P. Paulin et al.: Algorithms for High-Level Synthesis; IEEE Design & Test, Dec'89
73. F. Cathoor et al.: Interactive Co-design of High Throughput Embedded Multimedia; DAC 2000
74. L. Nachtergaele et al.: Optimization of Memory Organization and Partitioning for Decreased Size and Power in Video and Image Processing Systems; Proc. IEEE Workshop on Memory Technology, Aug 1995
75. F. Cathoor et al.: Custom Memory Management Methodology - Exploration of Memory Organization f. Embedded Multimedia Systems; Kluwer 1998
76. P.-A. Hsiung et al.: PSM: An Object-oriented Synthesis Approach to Multiprocessor Design; IEEE Trans VLSI Systems 4/1, March 1999
77. J. Kin et al.: Power Efficient Media Processor Design Space Exploration; Proc. DAC'99, New Orleans, June 21-25, 1999
78. R. Hartenstein, M. Herz, T. Hoffmann, U. Nageldinger: Generation of Design Suggestions for Coarse-Grain Reconfigurable Architectures; in [1]
79. V. Moshnyaga, H. Yasuura: A Data-Path Modules Design from Algorithmic Representations; IFIP WG 10.5 Worksh. on Synthesis, Generation and Portability of Library Blocks for ASIC Design, Grenoble, France, Mar 1992
80. U. Nageldinger et al.: Design-Space Exploration of Low Power Coarse Grained Reconfigurable Datapath Array Architectures; in [81]
81. D. Soudris, P. Pirsch, E. Barke (Editors): Proc. PATMOS 2000; Göttingen, Germany Sept. 13 - 15, 2000; LNCS, Springer Verlag, 2000
82. L. Kruse et al.: Lower Bounds on the Power Consumption in Scheduled Data Flow Graphs with Resource Constraints; Proc. DATE, Mrch 2000.
83. R. Hartenstein (invited paper): High-Performance Computing: über Szenen und Krisen; GI/ITG Worksh. Custom Computing, Dagstuhl, June 1996