

A High Performance Machine Paradigm Based on Auto-Sequencing Data Memory

R. W. Hartenstein, A. G. Hirschbiel, M. Riedmüller, K. Schmidt, M. Weber

Universitaet Kaiserslautern, F.B. Informatik, Bau 12,
Postfach 3049, D - 6750 Kaiserslautern, F. R. G.
phone: (+49-631) 205-2606 or: (+49-7251) 3575

Abstract.

This paper introduces a novel (non-von Neumann) programming paradigm of parallel computation featuring a much more efficient implementation of parallel algorithms, as well as a novel (hardware) machine paradigm efficiently supporting such implementations. Acceleration factors of up to more than 2000 have been obtained experimentally on an example architecture for a number of important applications - although using a hardware being more simple than that of a single RISC microprocessor. Due to its auto-sequencing data memory the machine principles are partly related to the organization of associative memories or systems. The machine organization and its most important hardware features are briefly introduced. The programming paradigm and its flexibility is illustrated by a few application examples.

1. Introduction

For a number of real-time applications extremely high throughput (up to several kiloMIPS) is needed at very low hardware cost. For at least another decade this mostly will be possible only with dedicated hardware, but not with programmable von-Neumann-type universal hardware. Even technologically advanced processors very often will be still too slow and/or too expensive. Also parallel or concurrent computers do not meet the requirements, or, are by far too expensive. Their sustained average performance is by orders of magnitude lower, than the peak rate. The reason is, that communication mechanisms offered by this hardware are not sufficiently powerful and/or too inflexible: the hardware is compiler-hostile, since most of the dense data dependencies of parallel algorithms cannot be mapped onto it. Next section gives more details about the reasons.

This paper introduces and illustrates a new machine paradigm being much more efficient in implementation of parallel algorithms by avoiding most of these problems. In contrast to the von Neumann paradigm this paradigm accepts an extraordinarily wide variety of optimization methods, since being supported by a number of innovative hardwired machine features, such as:

- auto-sequencing data memory (minimizing overhead)
- auto-sequencing register file organization (*scan cache*)
- reconfigurable rALU (featuring ultra micro parallelism)
- variable word length memory with innovative interface
- minimized access data trace by ultra micro scheduling

After a brief discussion of the difficulties to achieving massive parallelism on contemporary computer systems and also on VLSI solutions, some fundamental requirements for better hardware efficiency are identified. Also contemporary throughput measures are briefly discussed. After outlining basic xputer principles the MoM xputer architecture is introduced. Simple algorithm examples are used to illustrate xputer operation and programming. After summarizing a number of performance-relevant architectural features contributing to the superiority of xputers more examples are introduced to illustrate semi-associative xputer applications. Finally technology aspects and possible embeddings are discussed.

1.1 From Contemporary Hardware towards what ?

Communication mechanisms within concurrent computer systems are extremely hostile to optimizing compilers. Also vector machines have fundamental performance bottle necks [Vec, Ve3] and their sustained average performance is by several orders of magnitude lower, than their peak rate [Vec], even when creative coding techniques help the compiler [Ve2]. VLIW (Very Long Instruction Word) architectures [Eli, Ced] are much more optimizer-friendly by lower level of parallelism (at instruction level) [Bul, Para, Ga2] and relatively good optimization results have been reported for systolizable algorithms [Bul], but only for algorithms with only locally regular data dependencies (systolic algorithms or systolizable algorithms). VLIW architectures still have substantial drawbacks.

Also data flow machines are optimizer-hostile, since indeterministic operation does not permit optimization at compile time. Also data flow machines throughput is also affected by a number of other drawbacks: several new kinds of bottlenecks have been introduced. Code causes an enormous addressing overhead and data accessing conflicts [Gaj].

A higher degree of parallelism may be achieved by Application-specific Array Processors (ASAPs). Even ASAPs have substantial draw-backs: extensive I/O overhead is caused by scrambling and unscrambling of data streams, expensive design of special hardware is required. A more important drawback is, that only algorithms with **locally** regular data dependencies (systolic or systolizable algorithms, see [SK] and others) are supported. This drawback also holds for parallel computer architectures for systolic emulation [Wp].

1.2 Data-driven Ultra Micro Parallelism

A more detailed comparative analysis has been published elsewhere [Mch]. We strongly believe in the following fundamental requirements to avoid most of these problems, to obtain sufficiently optimizer-friendly hardware, to avoid most of the massive overhead caused (within software and hardware) by von Neumann principles. To obtain sufficiently flexible communication mechanisms parallelism should be implemented at a level much lower than usual: (1) below instruction level (*ultra micro parallelism*). Optimization (parallelization) should be based on very fine granularity resource allocation and scheduling (2) - determined at compile time to a much larger extent than usual (3). The paradigm should be deterministically **data-driven** (4).



Notice: This document has been provided by the contributing authors as a means to ensure timely dissemination of scholarly and technical work on a noncommercial basis. Copyright and all rights therein are maintained by the authors or by other copyright holders, notwithstanding that they have offered their works here electronically. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may not be reposted without the explicit permission of the copyright holder.

The non-von Neumann *xputer paradigm* being introduced in this paper is an approach into this direction. Its novel processor organization is based on *data sequencing* (in contrast to the *control flow sequencing* paradigm of von Neumann machines), so that also optimization methods based on **data-dependencies** are efficiently supported. Its architectural implementation is supported by an auto-sequencing data memory, which is partly related to functional memory [Fl, Gar] and other associative memory applications. But we are not aware of any other development of programmable hardware machine principles, which yields such a good utilization of hardware. Xputer use can avoid the need for special DSP processors or expensive image processing computers.

application area	algorithm example	acceleration factor	r-ALU size **	r-ALU size *
VLSI design automation	CMOS design rule check (pattern matching)	>2000	45	4
	electrical rules check	>300	5	0.5
	Lee routing	>160	5	0.5
digital signal processing and image processing	vector-matrix multiplication	>9	7	0.5
	• single scan cache	150	7	0.5
	• double scan cache	>300	< 0.5	0.02

*) number (or fraction) of PLD chip(s) needed with Plessey ERA 60400, **) with Altera MAX EPM 5128

Fig. 1. Acceleration factors by single processor MoM xputer - compared to a VAX-11/750.

1.3 Technology-independent Cost / Performance Evaluation

Encouraging performance results (fig. 1) have been obtained experimentally on the MoM xputer architecture at Kaiserslautern [MoM, CE] showing, that in several important applications a single xputer processor may even outperform larger parallel computer systems.

A computer-to-xputer performance comparison seems to be the best possible way to evaluate the merits of these results. Since an xputer does not have a hardwired "instruction set", it does not make sense to use MIPS, normally used for computer-to-computer comparison to indicate the progress of technology and physical design, rather than the efficiency of machine principles. But also other computational devices benefit from progress of technology. That's why we prefer the technology-independent measure of *acceleration factor* obtained experimentally (fig. 1) from two equivalent implementations of the same algorithm: one from a computer (VAX-11/750) and one from the technologically comparable MoM xputer [MoM]). We also have found out, that for computed acceleration factor estimates a good model is obtained from comparing the total number or duration of primary memory cycles.

Another important measure is the r-ALU size depending on the computation needed for a particular application and on the number of applications resident simultaneously (its role will be explained later). A rough measure of expense is the number of PLDs (programmable logic devices) needed of a particular type. Fig. 1 shows some such expense figures obtained experimentally on the MoM [MoM] xputer with code from an optimizing compiler having been implemented and tested at Kaiserslautern [CE].

2. Xputer Machine Organization

For clarification xputers are compared to **computers**. The ALU of computers is a very narrow bandwidth device: it can carry out only a single simple operation at a time. Xputers, however, use a PLD-based *r-ALU* ([MoM] fig. 2 b), reconfigurable such, that several highly parallel data paths form also powerful *compound operators*, which need only a few nanoseconds per execution, due to highly parallel dedicated intra-chip read / modify / write interconnect between register files (called *scan caches*, see below) and r-ALU (fig. 4 a). The r-ALU is configured only during loading, not at run time, so that PLD set-up slowness does not affect performance: dedicated wires are fast and avoid buses' multiplexing overhead [Bus]. Although 2 ns gate delay PLDs are available commercially, PLDs might be slower than traditional ALU technologies. This is more than compensated by its micro parallelism and other xputer features. Soft instruction set computers are not new: well known are dynamically microprogrammable architectures, where, however, flexibility is based on sequential programs. Also combinationally programmable repertoires of functions have been proposed [Fl, Gar], however, based on function tables mixed with data within data memory. But the rALU used for xputers is kept strictly separate from data.

In computers control flow is the primary activator (fig. 2 c): the instruction counter is the control state register. The rate of control flow is very high (*control flow overhead*): for each single data manipulation action at least one preceding control action is needed, which requires at least one memory cycle each. If no emit address nor emit data is used, additional control flow and even data operations are needed for address computation (*addressing overhead*).

Driven by a auto-sequencing data memory xputers are deterministically data-driven (fig. 2 d). For auto-sequencing the data memory interface includes a data sequencer, a hardwired data address generator (fig. 2 b, instead of computers' instruction sequencer: fig. 2 a). This hardwired data sequencer provides a repertory of generic data address sequences without any addressing overhead. Such an address sequence makes a scan cache move through data memory space, step by step, scanning a predefined segment of primary memory space along a path, which we call a *scan pattern*.



Notice: This document has been provided by the contributing authors as a means to ensure timely dissemination of scholarly and technical work on a noncommercial basis. Copyright and all rights therein are maintained by the authors or by other copyright holders, notwithstanding that they have offered their works here electronically. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may not be reposted without the explicit permission of the copyright holder.

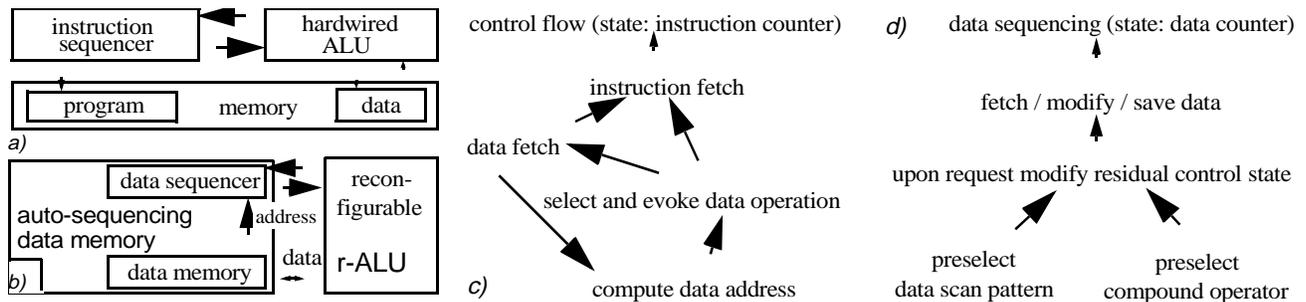


Fig. 2. Computers vs. xputers. basic structure: a) computers, b) xputers; causality in c) computers, d) xputers.

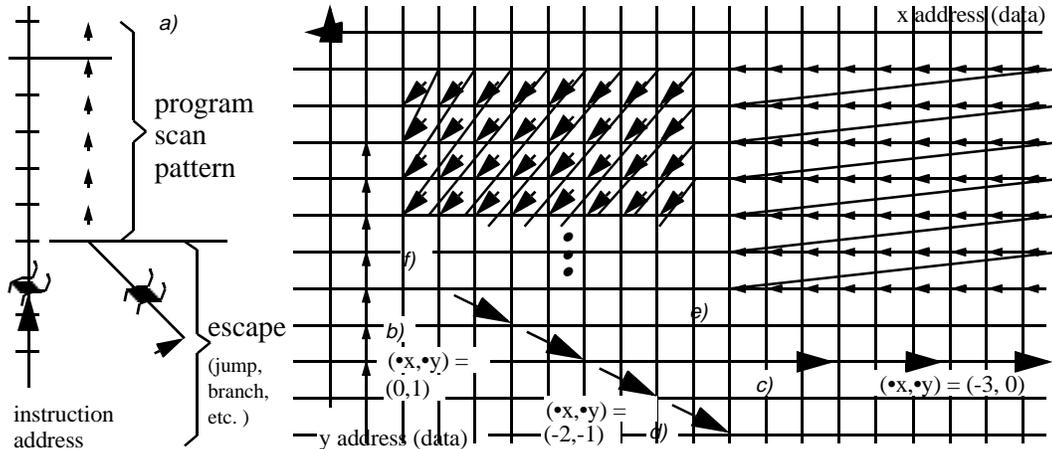


Fig. 3. Scan pattern(s): a) von Neumann control "scan" pattern, b) through f): some xputer auto-sequencing data scan pattern examples: b+c) 1-dimensional, scan, d-f) 2-dimensional scan.

2.1 The MoM (Map-oriented Machine)

Let's illustrate the role of this data sequencer by the MoM xputer architecture example. The MoM having been implemented at Kaiserslautern [MoM] features several peculiarities: (I:) 2-dimensional data address space (fig. 3 b - f), and (II:) a special optimizer-friendly register file organization. The MoM has 4 such register files which we call *scan windows* or *scan caches* (fig. 4 a), because each such cache operates like a size-adjustable window on the data memory (fig. 4 b).

Figures 3 b - d show linear scan pattern examples. Fig. 3 e shows a *video scan* pattern, e. g. for 2-D filtering. Other scan examples are: reflect, shift, shuffle, butterfly etc. Also special scan patterns to emulate systolic arrays, as well as data-dependent scan patterns (e.g. for image preprocessing etc.), are available in hardwired form (i. e. free of any overhead) [CE, MoM]. For more about particular scan patterns and their applications see [MoM, CE]. For a textual scan pattern language see [PP, We]. For stack-based hardware support of nested scan patterns see [Hir].

Computer control flow has only a single "scan pattern" (fig. 3 a, compare 3 b) scanning instructions one by one (as long as no branch nor jump is encountered, which we consider to be an *escape* from the scan). In contrast to those of xputers this scan pattern is not free of overhead: each step requires its own instruction fetch. Each instruction fetch requires a memory access cycle. This especially makes iterative operations inefficient, since the same instruction is fetched again and again. Looping instructions cause additional *control overhead* and thus additional memory access cycles. It is obvious, that the computer paradigm is extremely overhead-prone, whereas the xputer paradigm avoids most kinds of overhead.

3. The Data Sequencing Paradigm

For high level programming of xputers we use a simple model which is supported by the *auto-sequencing* data memory, and, which we call *data sequencing*. This paradigm will be illustrated here by 2 simple algorithm examples. The first example (a systolic algorithm: fig. 5) is not a good one to demonstrate the merits of xputers over vector machines. It has been selected for easy illustration of the data sequencing paradigm.

Fig. 5 a shows it textually and fig 5 b its *data dependence graph* (DG). From this DG the compiler derives: a *data map* (fig. 5 c + 6 c), from this map in fig 6 a) a *cache format spec* (middle of fig. 6 b) and r-ALU *subnet spec* including wiring (left side in fig. 6 b, derived from a single iteration in fig. 6 a), and finally a *scan pattern* (arrows in fig. 6 c). At each step of a scan the r-ALU subnet currently activated applies a read / modify / write cycle to the cache(s) currently active. In our example 8 steps (width = 2) are carried out (fig. 6 c shows initial and final cache locations).

3.1 Fine Granularity Scheduling.

This first example has illustrated the task of the innovative kind of compilers needed for xputer [CE, We]: a kind of fine granularity



Notice: This document has been provided by the contributing authors as a means to ensure timely dissemination of scholarly and technical work on a noncommercial basis. Copyright and all rights therein are maintained by the authors or by other copyright holders, notwithstanding that they have offered their works here electronically. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may not be reposted without the explicit permission of the copyright holder.

scheduling (or: *ultra micro scheduling*) of caches and rALU subnets, and, of data words, ready to be auto-sequenced. This is fundamentally different from sequentially piling up sequential code like conventional compilers do it for computers. Later in a section on xputer high performance features a more detailed impression on this scheduling task will be given.

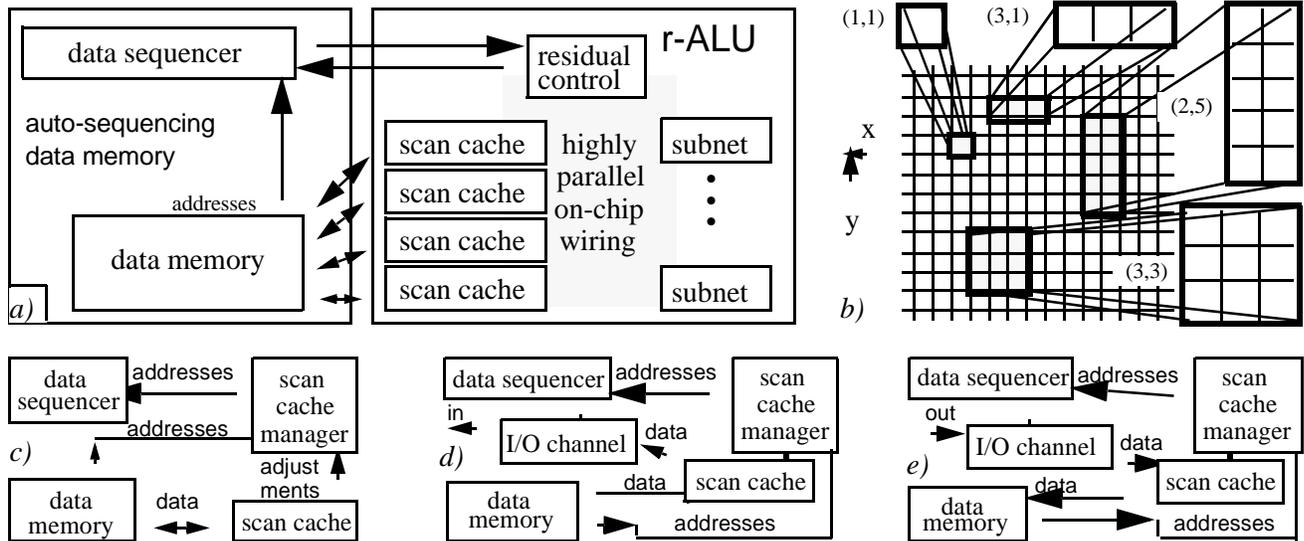


Fig. 4. The MoM xputer architecture: a) basic hardware structure, b) scan cache size adjustment examples c) details of data memory interface, d) scan cache input sequencing, e) scan cache output sequencing.

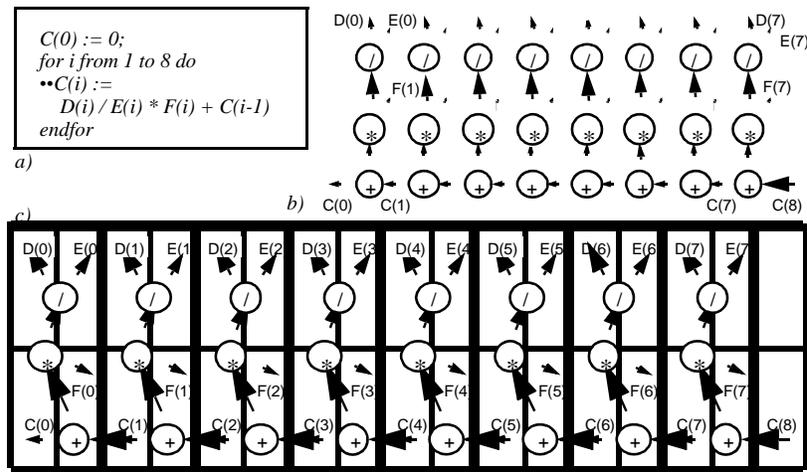


Fig. 5. Systolic algorithm example: a) procedural, b) data dependency graph (DG), c) derive data map from DG

3.2 Organization of Residual Control

At the end of the above data sequence example the cache finds a *tagged control word* (TCW: fig. 6 c) which then is decoded (right side in fig. 6 b) to change the state of the *residual control logic* (fig. 4 a) to select further actions of the xputer. This sparse TCW insertion into data maps we call *sparse control*. Note, that the control state changes only after many data operations (driven by the data sequencer). That's why we use the term *residual control* or *sparse control* for this philosophy. Note, that xputer operation is data-driven so that TCWs may be encountered only from within a data sequence. A TCW decoder is defined at compile time and configured as a subnet within the r-ALU. Fig. 7 a illustrates distribution of the residual control state between a scan state register (holding scan pattern select code and parameters), an ALU state register (holding subnet select code) and residual control state register.

We define, that only conditional branching, operator select and scan pattern select, but not data addressing, are control actions. Thus during a scan there is **no** control action: the data counter is not a state register. But **escape** from a scan **is** a control action (like in computers, see fig. 3 a). Escapes are (fig. 7 a): *normal escape* (by *end of scan flag* from data sequencer), *delimiter escape* (on TCW encounter), *off-limits escape* (address exceeds memory segment limits), *conditional branch escape* (by decision data from r-ALU), and, *event escape* (by external event flag). Upon off-limits escape, branch escape, or event escape a *remote control word* (RCW) or *remote address word* (RAW) is fetched from a remote memory segment via an *escape cache*. A second decision mechanism (*implicit branching*, because



residual control state is not affected) is activated only **within data-dependent scans** (i. e. without escape: curve following etc. [CE, MoM]). Such a data-dependent scan may be exited by conditional branch escape or off-limits escape.

To achieve xputer universality also non-generic scans and individual data accessing are needed, implemented by list-directed scan: next data address is read from a TAW (tagged address word) within the data map or from a RAW (in case of an escape). This list mode can be entered directly during a scan upon TAW encounter. If no TCW is found, a TAW does not activate residual control. Reading addresses from primary memory means addressing overhead, so that list-driven sequencing is slower than hardwired scan patterns. But also in this mode of operation the xputer paradigm is still superior to the computer paradigm.

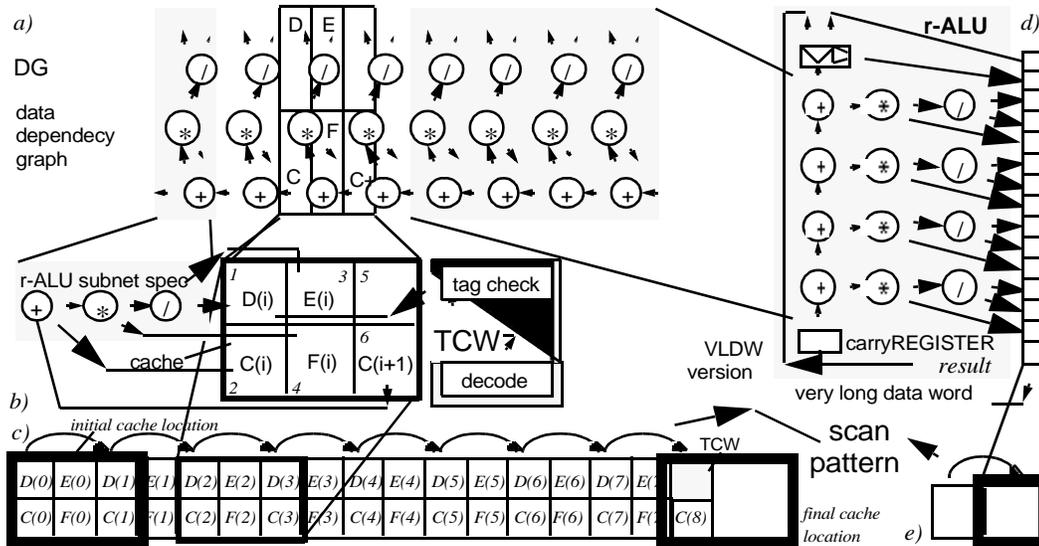


Fig. 6. Mapping a parallel algorithm (fig. 5) onto xputer hardware: a) dependency graph, b) a r-ALU subnet spec + cache size spec (3 by 2 words), c) its data map and scan pattern: 8 steps of width=2; d) VLDW version r-ALU subnet spec (scan cache size: 1 by 1), e) its data map and scan pattern: single step of width=1.

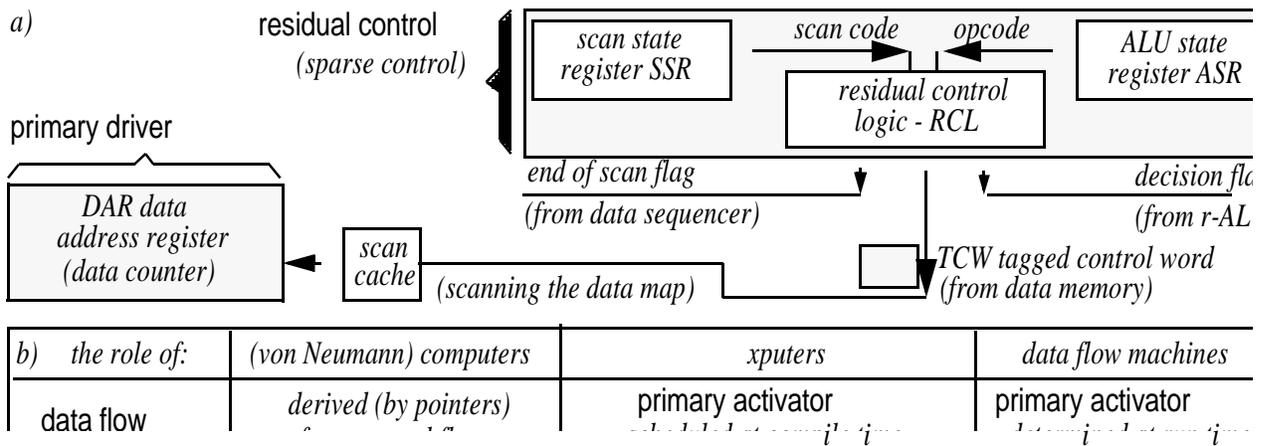


Fig. 7. Xputer control principles; a) state distribution of residual control, b) comparison to other devices.

3.3 I/O Data Sequencing

Xputer I/O is simple: the scan-cache-based data sequencing hardware (more details in fig. 4 c) is linked to an I/O channel (fig. 4 d/e), which is more powerful than DMA known from computers. The data streaming in are not just downloaded into a memory segment. Via a suitable scan pattern selection, along with proper scan cache adjustment, the data sequencer sets up a structured data map already during input operation. Also during output (fig. 4 e) data may be picked (by the data sequencer) from memory in a structured way.



Notice: This document has been provided by the contributing authors as a means to ensure timely dissemination of scholarly and technical work on a noncommercial basis. Copyright and all rights therein are maintained by the authors or by other copyright holders, notwithstanding that they have offered their works here electronically. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may not be reposted without the explicit permission of the copyright holder.

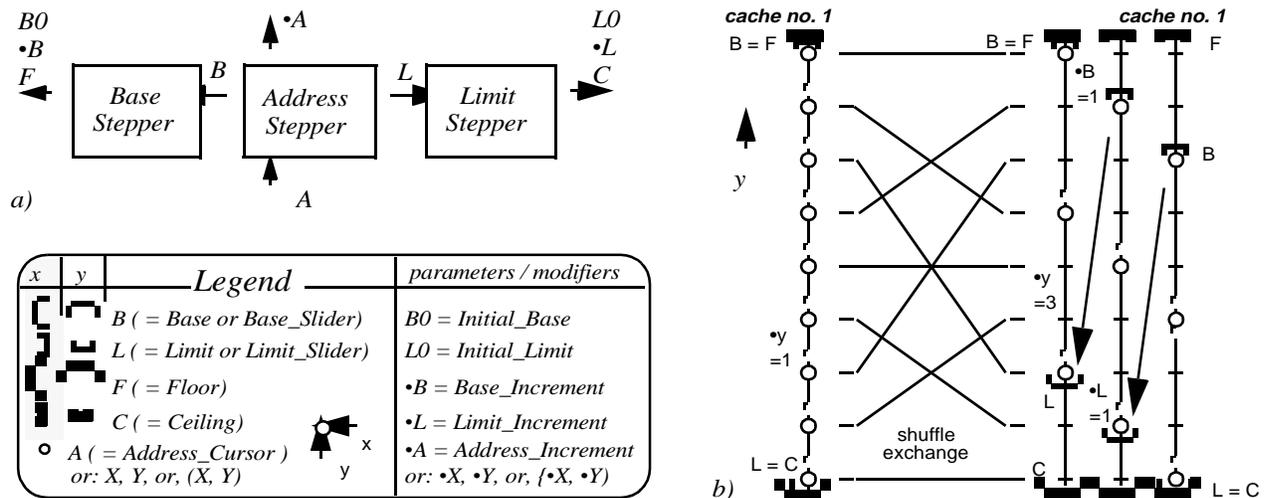
3.4 Highly Flexible Cost / Performance Ratio

Xputer word lengths are compiler-defined: data path, cache, and control words. Thus extensible xputer architectures are feasible, upgradable by inserting more PLDs into free r-ALU sockets and more boards into free memory slots. E. g. it is easy to design a VWL memory (Variable Word Length Memory), where data word length could be changed under software control to support VLDW (very large data word) strategies for more parallelism. Figure 6 d/e shows the VLDW version of fig. 6 a-c. Instead of only a single iteration of the DG this time the r-ALU subnet spec is derived from 4 iterations: the new version of the compound operator (fig. 6 d) is 4 times as powerful. Its use requires a vertical cache format (VLDW cache at right side in fig. 6 d), where a single word holds 14 operands. The scan pattern is very short, so that the 1-by-1 cache visits only 2 locations (fig. 6 e). In total the number of primary memory semi-cycles has been reduced from 41 to 2, so that a speed-up by about a factor of 20 has been obtained. This illustrates the extremely high flexibility of the xputer paradigm with respect to cost/performance trade-off.

3.5 Data Address Generator Hardware

This section illustrates the address generator operation. Fig. 8 a shows the structure of a stepper unit. The MoM uses a twin stepper for (x, y) addresses, providing a separate twin for each cache. The address stepper generates linear address sequences (A, A+yA, A+2*yA, ...) between limits B and L (e.g. fig. 3 b,c,d). The B stepper generates linear sequences B = (B0, B0+yB, ...) up to limit L and within F...C, to group bursts of A sequences. The L stepper operation is independent of that of the B stepper. Fig. 8 b shows address generation for a one-dimensional scan pattern example. It is a 2 cache example for a shuffle exchange scan, where 2 steppers cooperate synchronously on the y address only. A simple linear scan is applied to cache no. 1, a 3-by-3 warped scan to cache no. 2, where shuffling is controlled by equidistant B and L slides within boundaries defined by F and C.

Next example shows the address generation for a two-dimensional scan sequence. Figures 8 c through j illustrate the address generation for a systolic array emulation example. Fig. 8 c shows the systolic array to be emulated. Fig. 8 d shows the array operation by a snapshot sequence, where fig. 8 e illustrates the data access sequence. Figures f through j explain the address generation. Fig. 8 f illustrates the xputer data map, and fig. 8 g shows the snapshot sequence of the scan within the 4-by-4 data map. Fig. 8 h/j show the twin stepper operation to generate the coordinated x address and y address sequences.



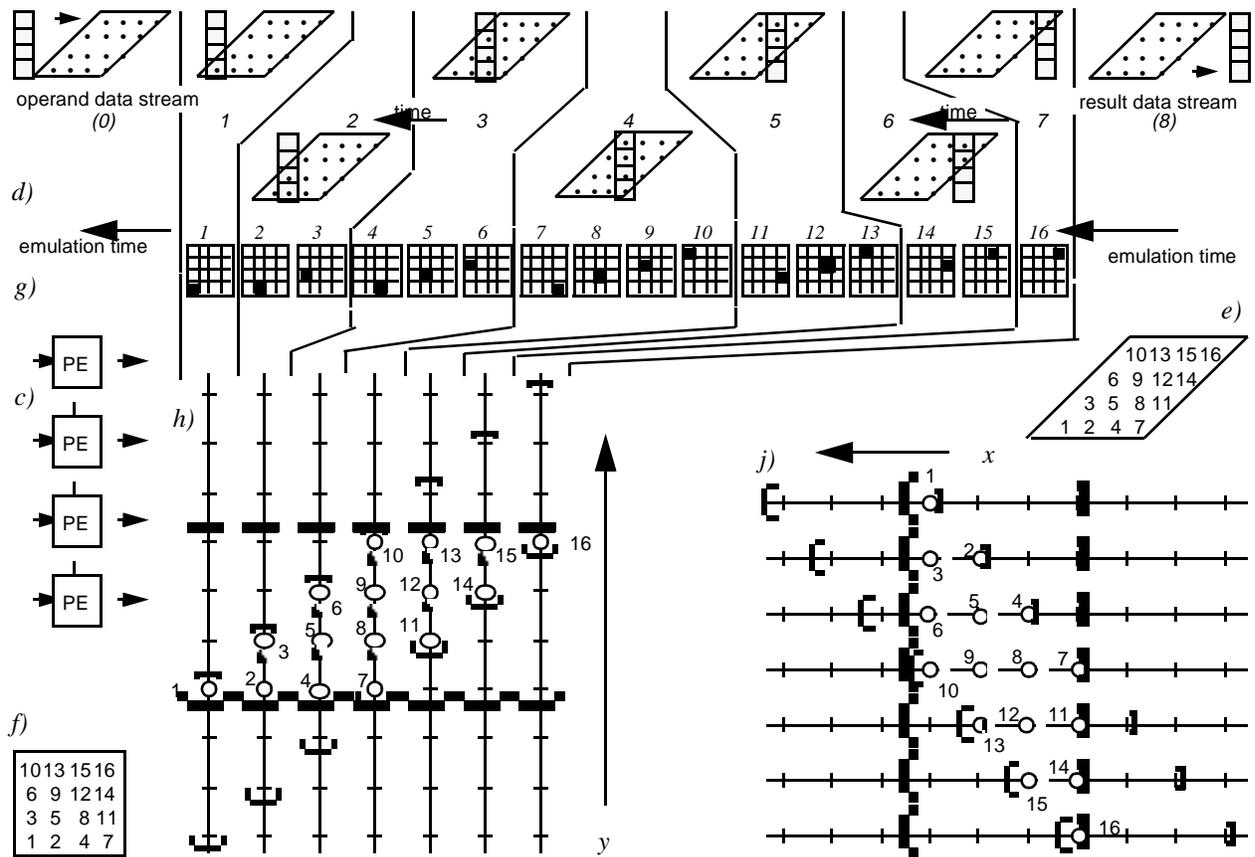


Fig. 8. Address Generator: a) Block Structure, b) Snapshots of Shuffle Exchange Addressing - Systolic Array emulation example: c) systolic array example, d) snapshots of its data stream, e) its data access sequence f) data map for its xputer emulation g) emulation snapshots h) illustration of y address sequence generation, j) illustration of x address sequence generation.

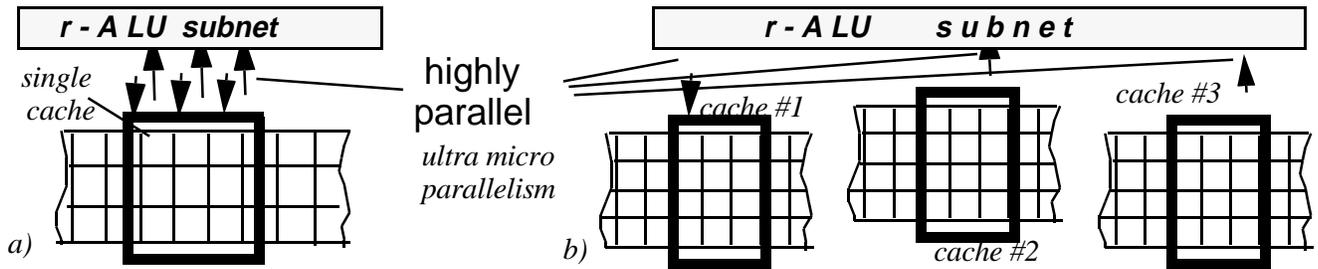


Fig.9. Scan cache: the xputer's basic communication mechanism. a) systolic example, b) non-systolic example



feature, reducing memory bandwidth requirement	memory access cycles saved by:	comment
compound functions ¹ (highly parallel r-ALU data paths)	• 1 complex expressions computed combinationaly	no storage of intermediate variables
ultra micro parallelism ¹ (very fine granularity parallelism)	each on-chip wire connected individually ¹	run time performance not affected ¹
no multiplexing overhead	no buses within CPU	(buses cause overhead [Bus])
data sequencing ²	sparse control ²	avoids control overhead
generic scan patterns ²	avoids addressing overhead	with parallel algorithms
non-generic scan patterns ²	reduces addressing overhead	with 'glue' software
minimized cache updating ¹	access mode tag-controlled ²	
high hit rate memory interleaving ²	by optimized data map ¹	
VLDW - very long data word ²	by VWL (variable word length) memory ^{1, 2}	xputers don't impose format constraints
xputer programming paradigm	by far less overhead-prone ⁴	in all levels of software ³

1) set-up or adjusted at compile time or loading time
2) hardwired feature

3) experimental results: reasons not yet well understood
4) - than von-Neumann-based organization, or, model, respectively

Fig. 10. Influence factors contributing to Xputer Efficiency - compared to those of Computers

4. Non-systolizable Algorithms

The introductory application example in fig. 5 / 6 has been a systolic algorithm, being easy to convert into a data sequencing scheme because of the locality of data communication. In digital signal processing and in other important application areas, however, also non-systolizable algorithms are very important. In contrast to parallel computer systems and VLSI arrays, xputers smoothly accept also non-systolic data sequencing schemes.

The implementation of non systolizable on the Xputer has been shown in [Bil, PP] for the constant geometry FFT. The Xputer FFT implementation uses 3 caches to support processing of non local data dependencies. This method is shown in figure 9.b.

5. Xputer High Performance Features

Having explained introductory sequencing examples we may obtain deeper insight into xputer performance issues more easily. Xputer performance stems from a number of different phenomena and concepts. Fig. 10 surveys the most important mechanisms contributing to the efficiency of parallel algorithm implementations running on xputers, which will be discussed throughout this chapter. Important roots of xputer efficiency are: the r-ALU's ultra micro parallelism, the data sequencing paradigm, and, the high flexibility of xputer memory interface architecture.

Much wider varieties of optimization strategies than possible with computers can be efficiently mapped onto this innovative methodology. Compound operators' ultra micro parallelism reduces memory access by substantially minimizing the number of stored intermediate variables. Often the r-ALU's flexible data path width facilitates better utilization of r-ALU space (e. g. see 2-D filtering example in next chapter). Dedicated intra-r-ALU interconnect avoids using buses being slow and causing *multiplexing overhead* [Bus].

The data sequencing paradigm obviously is by far less overhead-prone, than the von Neumann control flow paradigm. Control flow overhead is almost completely avoided (also no instruction fetch cycles are needed). The above examples have demonstrated, that addressing overhead is substantially reduced not only by hardwired address generator (also see the pattern matching example in next chapter). Not yet all mechanisms of overhead reduction in xputer programs are well understood: we propose basic research also covering overhead mechanisms of the von Neumann paradigm.

Now let's look at memory bandwidth. We may distinguish two kinds of factors: reduced memory bandwidth requirements due to the xputer paradigm, the r-ALU concept, and optimizing compilers (see above), and, providing higher memory bandwidth. Interface flexibility offers an extremely wide variety of strategies (optimum data maps) to meet the bandwidth requirements having been left over, where the xputer scan cache model is an important concept in finding such strategies. Important means are wide memory data paths (VLDW approaches, see above) supported by VWL memories (Variable Word Length memories, see above).

More hardware features having been developed for the MoM xputer [Aq, Hir, MoM] support further reduction of memory access time. Special *access mode tags* per cache word reduce the number of memory semi cycles needed for cache updating. For demonstration let's see the cache configuration in fig. 6 b: using a *read-only* tag for words no.1-4, *write-only* tag for word no.6, and an *ignore* tag for word no. 5, reduces the number of semi cycles per scan step from 12 to 5. The MoM cache mechanism also makes possible very high hit rates in interleaved memory access utilization. See example in fig. 6 c: if in a 4-phase interleaving scheme the groups of all C[i], all D[i], all E[i], and of all F[i] would be stored in separate memory banks, the number of semi cycles for cache update (see cache in fig. 6 b) would be further reduced from 5 to 2 (total speed-up factor: 6).

In unit step sequencing of large caches memory bandwidth bottlenecks can be reduced (due to optimizing compiler strategies) by another cache feature reducing repetitive access to memory locations. The MoM 2-D cache hardware also provides a multidirectional shift path, separately for each dimension, such that, for e. g. a 4-by-4 cache in a video scan the number of semi cycles is reduced from 32 to 8 [MoM]. By combination of this feature with interleaving the memory access rate may be further reduced to 2 (total speed-up factor: 12).



Notice: This document has been provided by the contributing authors as a means to ensure timely dissemination of scholarly and technical work on a noncommercial basis. Copyright and all rights therein are maintained by the authors or by other copyright holders, notwithstanding that they have offered their works here electronically. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may not be reposted without the explicit permission of the copyright holder.

Thus several relatively cheap hardware features supporting optimization may total up another order of magnitude of acceleration.

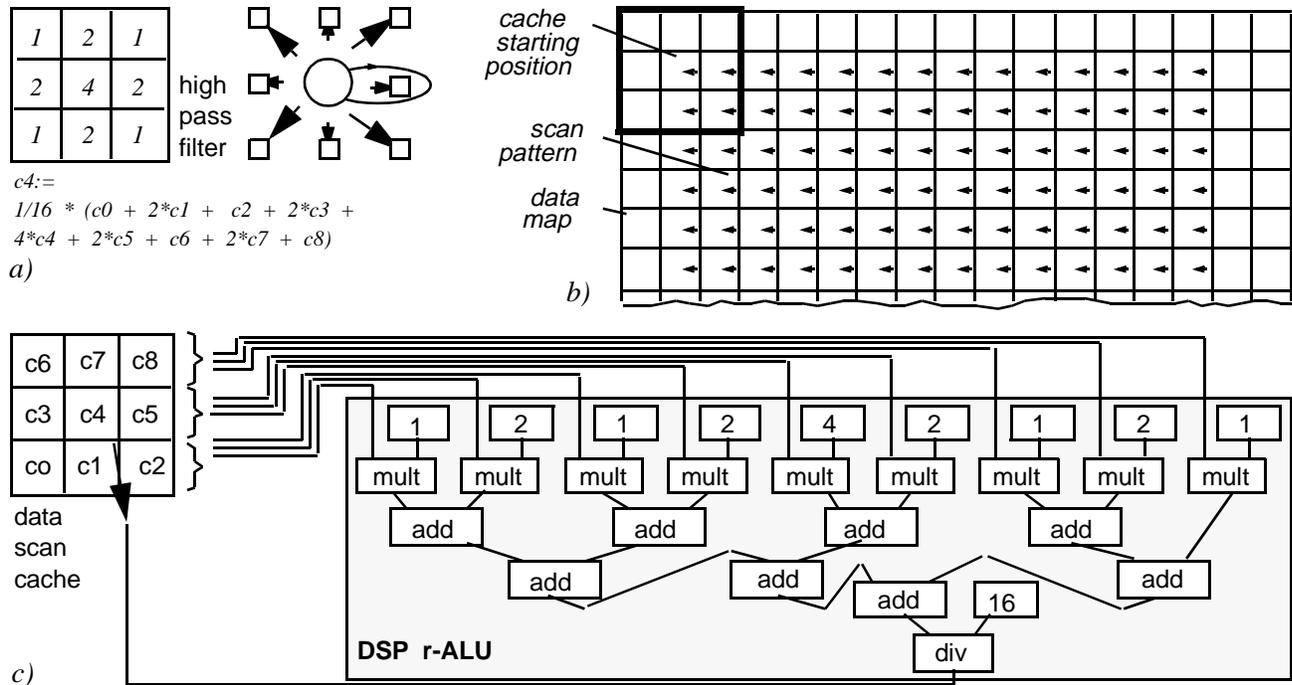


Fig. 11. Xputers in image preprocessing: 2-dimensional filtering example. a) filtering expression and cache map of weights, b) scan pattern example, c) r-ALU subnet derived from DG in fig. a.

Let's revisit the implementation of parallel algorithms by a second look at xputer communication mechanisms - from a higher level point of view. The highly parallel dedicated combinational path between cache(s) and r-ALU (fig. 9) is the basic communication mechanism of xputers (also compare fig. 2 a). Single cache use only supports local communication (systolizable algorithms only), such as e. g. between a data subarray [i] and subarray [i+1] (fig. 9 a) by overlapping cache positioning. Multiple cache use may also support global communication between different data arrays (fig. 9 b) or between distant subarrays. Also comparing acceleration factors in lines no. 4 and 5 within the table in fig. 1 shows, that here multiple-cache solutions tend to be much more efficient. Like cache memories of computers, scan caches in xputers help to reduce performance degradation due to the memory access bottleneck. It is obvious that xputer cache use is fully deterministic, due to a data scheduling strategy being completely compiler-driven.

That's why a much larger variety of optimization strategies may be applied, in contrast to computers permitting only probabilistic strategies which yield only low hit rates. By xputer cache use, however, extraordinarily high hit rates may be achieved, since cache traffic can be scheduled very precisely in detail to the optimum, tailored to any particular sequencing problem. This is because xputer hardware accepts almost any optimized schedule which always provides the right data at the right location at the right time. Thus compilation for xputers is a kind of *very high level synthesis*, where the number of visits to data locations in memory is minimized. This has similarities to the travelling salesman problem, where space-to-time mapping derived from systolic array synthesis methodology is an important method [Kil, Lem].

6. Xputers in Image Processing

Xputers are especially well suitable for image preprocessing, so that no specialized and much more expensive image processing computers are needed. Due to its universality also other kinds of parallel algorithms may be accelerated by the same xputer, and, in mass product applications stand-alone xputer use substantially reduces the total chip count. In image preprocessing systolizable algorithms (mainly using simple scan patterns, see fig. 3 e, f) and methods using data-dependent scan patterns are dominating. This section illustrates xputer use here by electronics design automation examples having been implemented at Kaiserslautern, where integrated circuit layout uses data structures being quite similar to those, well known from image preprocessing.

6.1 Two-dimensional digital filtering

Fig. 11 shows a 2-D digital filtering example implemented at Kaiserslautern: a systolic algorithm example in image preprocessing. A video scan pattern (fig. 11 b) is used to move a 3-by-3-sized single scan cache, which at each location recomputes the center pixel c4, by an expression shown in fig. 11 a. The cache map in fig. 11 a shows integer weight distribution. The r-ALU subnet (fig. 11 c) is derived from the local DG in fig. 11 a. Although including 18 arithmetic functions this compound function is purely combinational and fits on a small fraction of a single 5128 chip (last line in fig. 1) - due to the extraordinarily efficient minimization made possible by the high flexibility of xputer r-ALUs. Since xputer data path width is not hardwired a low path width (e. g. 8 bits for the adders in fig. 11) may save PLD space. Multiplication by 1 saves a multiplier entirely. In case of binary coded integers multiplication by 2 or 4 (see fig. 11) may be replaced by a shift left by 1 bit, or, by 2 bits, respectively. All this demonstrates xputers' high acceptance of a wide variety of optimization strategies. Further minimization yields from memory accessing strategies, possible with xputers only. On-cache shift paths



minimize the number of memory access cycles needed to 1 per word and video scan per line. Combined with suitable memory interleaving this may total up to an order of magnitude (see section 5 for explanations).

6.2 Pattern Matching Applications on Xputers

We use pattern matching examples to illustrate image preprocessing capabilities of xputers, such as applicable also to integrated circuit layout verification and routing using grid-based design rules. A DRC may be carried out by a finite state machine or combinational logic [San]. Such algorithms run very fast on ASIC hardware which, however, have to be reimplemented for changed design rules and for portation. Due to very large primary memories with modern workstations also conventional software implementation is feasible which, however, is very inefficient because of sequential processing of the very large number of reference patterns. But to measure acceleration factors such implementations are needed. The MoM-DE environment with tools like a reference pattern generator and the PISA [San] package facilitate comparative performance measurement by convenient generation of such pattern matching algorithms.

In contrast to computers, here the performance of xputers is competitive to ASIC solutions. E. g. for a grid-based design rule check (DRC) the MoM xputer has been programmed such, that a single video scan over the layout is sufficient. Substantial acceleration is obtained also for other kinds of grid-based layout processing, such as Lee routing [Aq, CE], ERC (electrical rules check), compaction, fault extraction, etc. Reference patterns are configured combinatorially into the r-ALU as a single very powerful compound function linked with a video scan sequence within a 2-dimensional bit map memory segment. A single read-modify-write data loop is performed per cache location without using decision data. Experimental results in grid-based DRC with 4-by-4 cache are acceleration factors of up to 2000 (CMOS design rules).

The extremely high acceleration factor is due to mainly two reasons: all (hundreds of) reference patterns are bundled by a huge compound Boolean operator (massive ultra micro parallelism) and caching completely avoids addressing overhead (an analysis of the VAX version of this algorithm has shown about 90% CPU time for addressing). Also MoM on-cache shift and access mode flag features, also see section 5) contribute to the high performance by minimized storage access time.

6.3 Application Development Support

MoM-DE, the MoM application development environment is running on a host (a μ VAX [Mch]) featuring a self-explanatory syntax-driven editor for a high level language MoPL (Map-oriented Programming Language), roughly a Pascal extension. MoPL sources are accepted by the MoMpiler the "code generator" of which includes a commercial PLD programming tool needed for r-ALU personalization. MoPL includes a sublanguage PaDL, which efficiently supports pattern matching applications in general. An optimizing reference pattern generator has been implemented [San], which accepts VLSI layout design rules [Aq]. For inclusion of pattern matching applications an interactive graphic pattern editor has been implemented [We] for easy editing, modification, inspection and surveying of sets of reference patterns.

7. Embedding and Technology Issues

The most common PLD application is hardware prototyping. But recently an innovative kind of PLD use has been commercialized: ASIC emulation from netlist sources [Qu1, Qu2]: replacing simulation since being a more efficient way of ASIC verification. In contrast to xputers, however, ASIC emulation does not provide a new design paradigm: the netlist is imported: the result of a separate (conventional) hardware design process. Xputers have a programming paradigm: a very high level model of parallel algorithms. Running an implementation on an xputer is execution - but not emulation. Since for some PLDs also compatible gate arrays are available commercially (e. g. by Plessey): xputer machine code may be directly submitted for fabrication. That's why the xputer paradigm may be considered to be an alternative high level synthesis approach to ASIC design [CE] - more precisely: *very high level synthesis*. ASIC emulation nor simulation is needed, since direct execution is available for design verification.

7.1 Partitioning large r-ALUs

To avoid communication bandwidth problems, cache(s) and r-ALU should be on the same chip. If for "large" applications or VLDW use more than a single PLD chip is needed for the r-ALU, also more expensive **inter**-chip wiring is needed in addition to the very efficient **intra**-chip wiring. This is rather a packaging issue of than a performance issue, since still primary memory access remains the only critical bottleneck. In implementing several such "large algorithms" we have always found heuristically a clever partitioning scheme, by slicing caches into multi-bit slices and distributing the compound operator such, that only loose coupling is required between chips.

8. Conclusions

With xputers an innovative computational machine paradigm has been introduced and implemented which achieves for parallel algorithms (also non-systolizable ones) drastically better performance and hardware utilization and drastically more (compiler-) optimizer-friendliness than the von Neumann paradigm (comparative summary: fig. 12). Acceleration factors up to more than 2000 have been obtained experimentally with a simple monoprocessor. For many applications xputers may outperform large parallel computer systems or ASIC solutions. Due to convenient conversion into a gate array the xputer also provides an alternative ASIC design methodology.

Xputers fit well to image preprocessing and digital signal processing, so that often special DSP processors or expensive special image processing computers are not needed. Due to xputer universality also other kinds of parallel algorithms and glue software may run on the same xputer, and, in mass product applications a stand-alone use is possible, which substantially reduces the total chip count. For xputer architectures an extremely low amount of specific hardware is needed, not being performance-critical, so that it's easy to keep up with technology.

An exciting new R&D scene has been opened up: immature, thus promising and challenging. Not really a new theory, but a new mix of backgrounds is needed, derived from languages, compilation, algorithms and applications. Not yet all phenomena are well understood which contribute to the high acceleration factors found experimentally. We need a new direction of (*very*) *high level synthesis*, a new direction in hardware / software performance evaluation redefining the notion of *overhead*, and, a data-sequencing-oriented new direction of research in programming languages.



Notice: This document has been provided by the contributing authors as a means to ensure timely dissemination of scholarship and technical work on a noncommercial basis. Copyright and all rights therein are maintained by the authors or by other copyright holders, notwithstanding that they have offered their works here electronically. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may not be reposted without the explicit permission of the copyright holder.

	(von Neumann) computers	xputers
instructions	hardwired repertory of simple operations	tailored at compile time: powerful compound functors
machine code	sequential: scanned from a program store	combinational: for PLD configuration (at loading or compile time)
sequencer	instruction sequencer	data sequencer
parallelism at	process level or (VLIW:) instruction level	gate level and / or data path level
compilation techniques	traditional	data-driven high level + logic synthesis from algorithm specification
data format	dependant of instruction format	highly flexible: variable word length memory is feasible

Fig. 12. Computer features versus Xputer features: Summary and Conclusions

References

- [Aq] R. W. Hartenstein, A. G. Hirschbiel, M. Weber: MoM - Map Oriented Machine, in: Chiricozzi, D'Amico: Parallel Processing and Applications, North Holland, Amsterdam / New York 1988.
- [Bil] A. Ast, et al.: Using Xputers as Inexpensive Universal Accelerators in Digital Signal Processing; Int'l Conf. on New Trends in Signal Processing, Communication and Control, Ankara, Turkey, July 1990, North Holland 1990
- [Bul] J.Fisher, J.Ellis, J. Ruttenberg, A. Nicolau: A Parallel Compiler and a Dumb Machine; Proc. ACM SIGPLAN '84 Symposium on Compiler Correctness; SIGPLAN Notices 19,6 (June 1984)
- [Bus] R.W.Hartenstein, G.Koch: The Universal Bus Considered Harmful; in: R.Hartenstein, R.Zaks: The Microarchitecture of Computing Systems; North Holland, Amsterdam/New York 1975;
- [CE] R. Hartenstein, A. Hirschbiel, M. Weber: *MoM - a partly custom-designed architecture compared to standard hardware*; Proc. IEEE Comp Euro '89, Hamburg, FRG, IEEE Press, 1989
- [Ced] D. D. Gajski, et al.: CEDAR; COMPCON Spring 1984
- [Eli] J. A. Fisher: Very Long Instruction Word Architectures and the ELI-512; Proc. 10th ISCA, 1983
- [Fl] M. Flinders, P. L. Gardner, J. F. Minshull, R. J. Llewelyn: Functional Memory as a General Purpose Systems Technology; Proc. 1970 IEEE Computer Group Conference, June 1970
- [Gaj] D. D. Gajski, D. A. Padua, D. J. Kuck, R. H. Kuhn: A Second Opinion on Data Flow Machines; Computer, 15, 2 (Febr. 1982)
- [Ga2] D.D. Gajski et al.: Dependence Driven Computation; Proc.COMPCON Spring 1981, IEEE Press 1981
- [Gar] P. L. Gardner: Functional Memory and Its Microprogramming Implications; IEEE TC-20,7 (July'71)
- [Hir] A. G. Hirschbiel: (Ph. D. thesis), Univ. Kaiserslautern, 1990
- [Kil] R. W. Hartenstein, A. G. Hirschbiel, M. Weber: Mapping Systolic Arrays onto the Map-Oriented Machine (MoM), Proc. 3rd Int'l Conf. on Systolic Arrays, Kilarney, Ireland, May 1989.
- [Lee] C. Y. Lee: An Algorithm For Path Connections And Its Applications. IEEE TrC-10 (Sept. 1961)
- [Lem] R. Hartenstein, K. Lemmert, SYS3 - A CHDL-Based CAD System for the Synthesis of Systolic Architectures, Proceedings IFIP CHDL '89, North Holland, Amsterdam / New York 1989.
- [MoM] R. W. Hartenstein, A. G. Hirschbiel, M. Weber: *MoM - Map Oriented Machine*; in: Ambler et al.: (Prepr. Int'l Worksh. on) Hardware Accelerators, Oxford 1987, Adam Hilger, Bristol 1988.
- [Mch] R. Hartenstein et al.: Xputers: an Open Family of non-von-Neumann Architectures; Proc. GI/ITG Conf. on the Architecture of Computing Systems, Munich, 1990; VDE-Verlag Berlin 1990
- [Para] D.Padua et al.: High Speed Multiprocessors and Compilation Techniques; IEEE TC-29, 9 (Sept. '80)
- [PP] R. W. Hartenstein, A. G. Hirschbiel, M. Weber: The Machine Paradigm of Xputers: and its Application to Digital Signal Processing Acceleration; ICPP-90 Int'l Conf. on Parallel Processing, 1990; IEEE Press, Wash., D.C., 1990
- [Qu1] M. D'Amour, et al.: ASIC Emulation cuts Design Risk; High Performance Systems, Oct. 1989
- [Qu2] P. A. Kaufman: Wanted: Tools for Validation, Iteration; Computer Design, Dec. 1989
- [San] R. W. Hartenstein, R. Hauck, A. G. Hirschbiel, W. Nebel, M. Weber: PISA - A CAD package and special hardware for pixel oriented layout analysis, Proc. ICCAD 1984, Santa Clara 1984.
- [SK] S. Y. Kung: VLSI Array Processors; Prentice-Hall, 1988
- [Vec] J. J. Hack: Peak vs. Sustained Performance in Highly Concurrent Vector Machines; Computer 19, 9 (Sept. 1989)
- [Ve2] J. J. Dongarra, E. Eisenstat: Squeezing the Most out of an Algorithm in Cray Fortran; report, ANL/MCS-TM-9, 1983
- [Ve3] J. J. Dongarra: A Survey of High Performance Computers; COMPCON Spring 1987
- [We] M. Weber: (Ph. D. thesis), Univ. Kaiserslautern, October 1990
- [Wp] M. Annaratone et al.: The Warp Computer: Architecture, Implementation, and Performance; IEEE TC-36(12), (Dec. 1987)

