

Pipelining the Connex Arraytm

Dominique Thiébaud¹, and Mihaela Malita²

¹thiebaut@cs.smith.edu, Dept. of Computer Science, Smith College, Northampton, MA.

²mmalita@anselm.edu, Dept. Computer Science, St Anselm College, Manchester, NH.

Abstract

We study the costs and benefits of implementing pipelining computation on the Connex Array circuit, a data-parallel PIM processor currently manufactured by ConnexArray Technologies. We find that the existing SIMD configuration lends itself well to the enhancement, with most of the required hardware modifications localized in the PIM cells, allowing standard design techniques and offering good potentials for frequency increase.

1 Introduction

The growth of performance of current uniprocessor systems is slowing down due in part to two combined effects. One is the exponentially increasing speed gap between processor and semiconductor memory. The second is the von Neumann bottleneck which turns the processor as a sink for all memory words [4]. Moreover, in some areas such as search engines, the sheer size of the infrastructure required is such that companies are faced with major power management issues, requiring them to locate their data centers near cheap power sources[5]. Under these conditions it is paramount to explore new architectural solutions to increase the performance of computer systems, reduce power dissipation, and ensure that processing performance stays on its course.

The Connex Array[3, 1, 2] concept¹ is a 130nm technology circuit. It is a programmable, scalable alternative for low-power high computation. In its current implementation, the Connex Array is a hierarchical in-memory device that permits fast *arithmetic*, *logic*, and *string operations* in a data parallel fashion over 1024 intelligent memory cells each associated its own storage of 256 RAM cells.

Because it is a programmable and scalable device, the Connex Array provides good performance improvement for a wide range of algorithms such as image processing, the alignment of DNA sequences [6], the real-time filtering of network packets for detection of illegal activities [7], or the massive computation of polynomials.

¹The Connex Array is manufactured by ConnexArray Technologies, which recently changed its name to BrightScale.

2 The Connex Array Circuit

The 256 by 1024 memory array can be conceptualized by the programmer as 256 vectors of 1024 words that can be manipulated as single entities, and engaged in simple vector manipulations such as basic vector arithmetic, searching for the occurrences of a pattern in a vector, or inserting or deleting subsections of a vector.

The 1024 intelligent cells are individually addressable as in a regular random-access memory, but can also receive broadcast instructions or data on which they operate in parallel at a peak rate of 1 operation per cycle. This general concept fits the *Processor-In-Memory* paradigm. The intelligent cells are connected by a linear chain network, allowing fast shifting of data between the cells, as well as the insertion or deletion of data from cells while maintaining the relative order of all the data. All these operations are performed in one memory cycle only.

All these operations are overseen by an on-chip controller which issues operations at a rate of one per cycle. The cycle length is determined by the current memory technology, and with today's technology is equal to 5 ns.

The Connex Array is packaged for scalability, with a chip design allowing several Connex Array chips to coexist on the same board and interface with each other to extend the length of the vectors, while receiving instructions and data from only one controller.

3 Pipelining The Connex Array Circuit

The current architecture of the Connex encompasses the PIM and data-parallel paradigms. In this paper we investigate enhancing the architecture by adding MIMD behavior to the cells. Because the PIM cells are currently connected by a simple chain network, this MIMD behavior leads directly to the time parallelism offered by pipelines.

At the hardware level, adding MIMD behavior affects mostly the individual cells, since the network is already implemented.

The new design must address several requirements: 1) the cells must be able to switch between SIMD and MIMD modes of operations, and 2) in MIMD mode, the cells must receive their instructions not from the controller over the broadcast network, but from their local memory, currently holding 256 words.

For the first requirement, we propose to add a status bit accessible by the controller, that forces the cells to switch from SIMD to MIMD and vice versa. The instruction set of the controller is augmented with one instruction that is broadcast to all the cells and switch them from one mode to another.

The second requirement is the most extensive, but has the advantage of being local to the design of the PIM cell. This means that standard automated design techniques can be used to modify the design of the cells.

Because in MIMD mode the instructions must come from the local memory, and instructions contain memory operands, the 256-word memory associated with each PIM cell must be split into two banks accessible in parallel while in MIMD mode, while still providing a uniformly accessible 256-word data memory in SIMD mode. The optimum size of the bank holding the MIMD code is still a research question, and an evaluation of several applications should provide good insight in this matter. In addition each cell must hold its own program counter of at most 8 bits. Finally the control units of each PIM must be modified to allow the execution of new control instructions including unconditional and conditional jumps. Note that the current PIM cells already contain comparator and status bits hardware.

A preliminary analysis reveals that the proposed modifications would increase the chip area which is currently of 64 mm², by 20% for the memory area, and by 2% for the discrete logic in the control unit, for a total area increase of no more than 25%.

4 Performance

The raw performance of the ConnexArray is simply 1024 parallel instructions performed at frequency of 200 MHz, or 200 Gops for an average 2.5 Watt dissipation (5 Watt max). The performance of the pipelined Connex is similar, but its range of application is expanded by the addition of the pipeline in several ways. The first obvious benefit is when the computation requires one subset of cells to perform one task and the complementary set another task can now overlap in MIMD mode.

The second benefit is that data often arrives in the computer serially, and putting the Connex into pipeline mode allows for preprocessing of the data as it enters the array before it can be processed in SIMD mode. Applications in the video, images, and network areas are good targets for this type of operation. Also, because operation of the PIM cells in MIMD mode is local, problems associated with clock skew are eliminated, and it is thus possible to clock the circuit at higher frequency when working in MIMD mode.

On the negative side, we note that the current average power dissipation of 2.5 Watts will be increased in MIMD mode because of the need of the PIM cells to access memory every cycle. Based on measurements performed on prototypes of the original Connex chip, we expect that the increase in power dissipation would at most double the av-

erage power dissipation. Also, while we expect some overhead in the switching of SIMD to MIMD mode and conversally, the potential boost in clock frequency in MIMD mode might compensate for this overhead.

5 Conclusion

The Connex Array is a circuit with great potentials, mainly because of its high level of parallelism, its scalability, its programmability, and its low power dissipation. It has been shown to offer good performance in areas where data-parallelism is large, such as processing DNA structures, processing pixels, checking large tables of IP addresses, and for computing massive numbers of polynomial equations. For applications where the size of the data is also very large, and where pipeline-style computation is required, such as the processing of network packets, we expect that the pipelined Connex Array circuit could provide good performance/power dissipation support such applications and provide significant power and computational performance.

6 Acknowledgements

We wish to thank George Ștefan of ConnexArray Technologies for his insights and careful reading of this document.

References

- [1] G. Ștefan and D. Thiébaud. Memory engine for the inspection and manipulation of data. U.S. Patent No. 6,760,821, July 2004.
- [2] G. Ștefan, D. Tomescu, and D. Thiébaud. Associative memory device. U.S. Patent No. 7,069,386, June 2006.
- [3] G. Ștefan *et al.* The ca1024: A fully programable system-on-chip for cost-effective hdtv media processing. In *Hot Chips: Symposium on High Performance Chips*, August 2006.
- [4] S. Borkar *et al.* Platform 2015: Intel processor and platform evolution for the next decade. *Intel Executive Summary*, 2005.
- [5] George Gilder. The information factories. *Wired Magazine*, 14(10):1–5, October 2006.
- [6] D. Thiébaud, G. Ștefan, and M. Malita. Local alignment of dna sequences with the connex array. In *Proceedings of the Int'l Conf. on Computing in the Global Information Technology (ICCGI06)*, Bucharest, Romania, August 2006.
- [7] D. Thiébaud and M. Malita. Real-time packet filtering with the connex array. In *Proceedings of the Int'l conference Anchor/ISCA06*, Boston, MA, June 2006.