# A Programmable Processor with 4096 Processing Units for Media Applications

A. Krikelis, I. P. Jalowiecki, D. Bean, R. Bishop, M. Facey, D. Boughton, S. Murphy, and M. Whitaker
Aspex Technology Ltd.
Brunel Science Park
Kingston Lane
Uxbridge, UB8 3PH
United Kingdom
argy.krikelis@aspex.co.uk

## Abstract

*Over the past few years, technology drivers for processor designs have changed significantly. Media data delivery and processing – such as telecommunications, networking, video processing, speech recognition and 3D graphics – is increasing in importance and will soon dominate the processing cycles consumed in computer-based systems. This paper describes a processo, called Linedancer, that provides high media performance with low energy consumption by integrating associative SIMD parallel processing with embedded microprocessor technology. The major innovations in the Linedancer is the integration of thousands of processing units in a single chip that are capable to support software programmable high-performance mathematical functions as well as abstract data processing. In addition to 4096 processing units, Linedancer integrates on a single chip a RISC controller that is an implementation of the SPARC architecture, 128 Kbytes of Data Memory, and I/O interfaces. The SIMD processing in Linedancer implements the ASProCore architecture, which is a proprietary implementation of SIMD processing, operates at 266 MHz with program instructions issued by the RISC controller. The device also integrates a 64-bit synchronous main memory interface operating at 133 MHZ (double-data rate, DDR), and a 64-bit 66 MHz PCI interface.*

## Introduction

Over the past few years, technology drivers for microprocessors have changed significantly. High-end systems for technical and scientific applications used to direct the evolution of processor architecture. Now, consumer-level systems drive technology, due to their large volume and attendant profits. Within this environment, important application and technology trends have evolved. Media data delivery and processing – such as telecommunications, networking, video processing, speech recognition and 3D graphics – is increasing in importance and will soon dominate the processing cycles consumed in computer-based systems [1]. SIMD extensions to existing processor architectures [2, 3] for supporting DSP type of operations are essentially narrow vector designs without support for vector memory operations. They have limited scalability because each instruction specifies a fixed number of operations. Most extensions do not support SIMD memory operations, therefore exposing data alignment to user software [4]. Certain instructions, such as random permutations, will not scale well due to interconnect delay scalability problems.

This paper presents the architecture of the Linedancer DSP processor. The Linedancer provides high multimedia performance with low energy consumption by integrating associative SIMD parallel processing with embedded microprocessor technology. The major innovations in the Linedancer is the integration of thousands of processing units in a single chip that are capable to support software programmable high-performance mathematical functions as well as abstract data processing.

## Linedancer Processor

Linedancer is a DSP processor for media processing that integrates 4096 processing units. With 15.6M transistors, the processor is implemented in a 0.18um 5-metal-layer CMOS process. It is packaged in a 560-pin EBGA pack. The core of the device operates at 1.8 Volts, while the I/Os are 3.3 Volts. The device, which will be sampling in the first quarter of 2001, at its peak, dissipates less than 5 Watts at the 266 MHZ typical operating frequency. Table I details the Linedancer device specifications.

| | |
|---|---|
| Technology | 0.18 μm, 5-layer metal CMOS |
| Clock | 266 MHz |
| Power Supply | 1.8 V Core, and 3.3 V I/O |
| Transistor Count | 15.6 millions |
| Power Dissipation | 4 Watts (peak) |
| Die Size | 111 $mm^2$ |
| Package | 680 EBGA |
| Testability | Full-scan |
| Data Transfer – Synchronous Memory Interface | Up to 2,100 Mbytes/sec |
| PCI Interface | 8 bytes @ 66 MHz |
| Integer Performance (8-bit add/subtract) | 126,000 MOPS |
| Integer Performance (8-bit x 8-bit multiplies) | 13,500 MOPS |
| Integer Performance (16-bit x 16-bit multiplies) | 3,800 MOPS |

| | |
|---|---|
| Floating Point Performance (IEEE Single Precision) | 1,700 MFLOPS |

**Table I: Linedancer device specifications**

The Linedancer processor is one of the devices that incorporates ASProCore (**A**ssociative **S**tring **Pro**cessor **Core**), a software programmable associative SIMD structure that provides extremely efficient support for the parallel processing required for media applications.

Figure 1 depicts the block diagram of the processor called the Linedancer. The Linedancer major functional units include the ASProCore, a RISC controller that is an implementation of the SPARC architecture, 128 Kbytes of Data Memory, and I/O interfaces. The ASProCore, which is a proprietary implementation of SIMD processing, is designed to operate at 266 MHz with program instructions issued by the RISC controller - the controller block includes a 128 Kbytes program memory. The device also integrates a 64-bit synchronous main memory interface operating at 133 MHZ (double-data rate, DDR), and a 64-bit 66 MHz PCI interface.

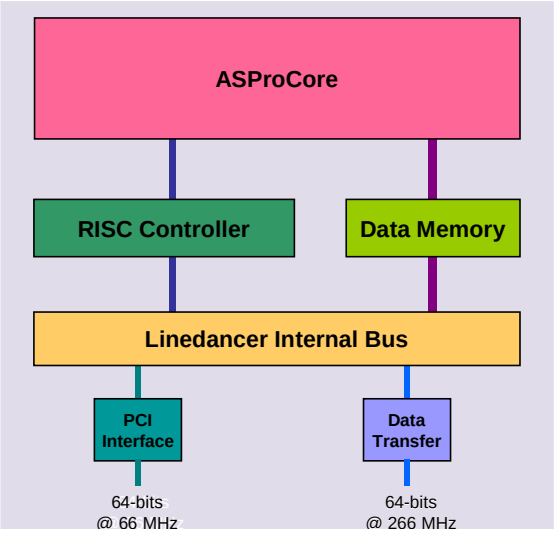


**Figure 1: Linedancer Block Diagram**

The Data Transfer block provides the application programmer with the support required to develop application software that overlaps application processing with data transfers. In inner loops operating over large sets of data, the programmer can program it to move strips, blocks or patches of data on and off the chip. The overall design of the processor allows SIMD processing in ASProCore to fully overlap with data movement.

## ASProCore

The ASProCore (Associative String Processor Core) part of the Linedancer is a programmable, homogeneous and fault-tolerant SIMD parallel processor core incorporating a string of identical processing units, a reconfigurable Intercommunication Network, and a Vector Data Buffer for fully overlapped data input-output as indicated in Figure 2.

As shown in Figure 2, each processing unit incorporates a Data Register and a bit-serial ALU. The size of the Data Register is 200 bits. The Data Register, in addition to storing data for arithmetic operations involving the local ALU, can support associative processing operation (i.e. to direct support for logical and relational operations).

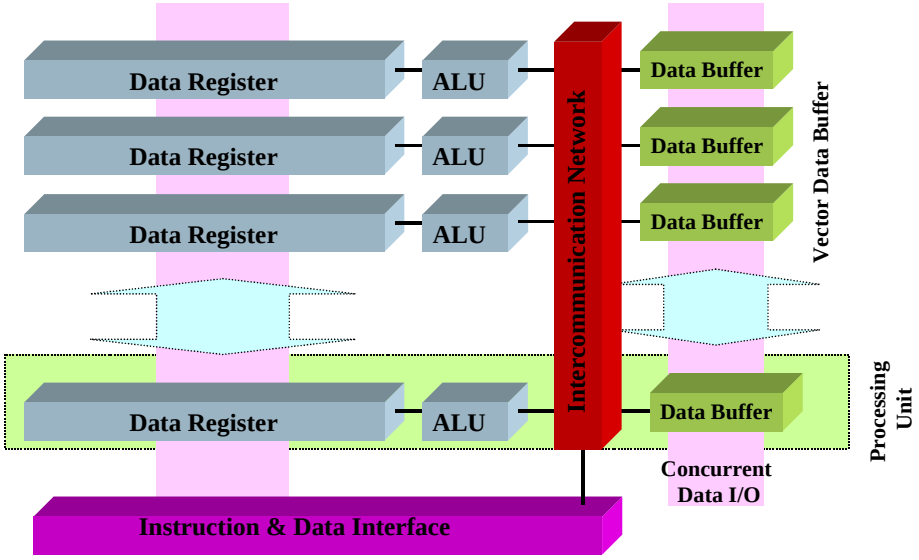


**Figure 2: ASProCore Architecture**

The processing units are connected via the Intercommunication Network. The Intercommunication Network is a flexible network that supports data transfers and navigation of data structures. It can be dynamically reconfigured, in a programmable and user-transparent way, thus providing a cost-effective emulation of common network topologies. The interconnection strategy supports 2 modes of inter-processing communication:

- *asynchronous* bi-directional single-bit communication to connect processing units sources and corresponding processing units destinations of high-speed activation signals, implementing a fully-connected dynamically-configured (programmer-transparently) permutation and broadcast network for processing element selection and inter-processing element routing functions;
- *synchronous* bi-directional multi-bit communication, via a high-speed bit-serial shift register for data/message transfer between processing unit groups.

While being served with control and sequential data via the Instruction and Data Interface, the ASProCore can support parallel data I/O via the Vector Data Buffer. Data is loaded, overlapped with SIMD parallel processing, word-sequentially, bit-parallel into the Vector Data Buffer, It can subsequently be exchanged with the data stored in the Data Register of the local processing unit in a word-parallel, bit-sequentially manner. For data-parallel operations, data are distributed over the processing units and stored in the local Data Register. Successive computational tasks are performed on the stored data and the results are dumped. The ASProCore supports a form of set associative processing, in which a sub-set of active processing units (i.e. those which associatively match broadcast scalar information) support scalar-vector (i.e. between a scalar and Data Registers) and vector-vector (i.e. within Data Registers) operations. Matching processing units are either directly activated or source inter-processing element communications to indirectly activate other processing units. The control interface provides feedback on whether none or some processing units match. The instruction set for the ASProCore is based on 4 basic operations, *match*, *add*, *read* and *write*. More complicated functionality can be performed by combining these operations. A more detailed discussion of the ASProCore architecture can be found in [5].

### Software Programmable

Conceptually the Linedancer processor can be viewed as a general purpose RISC processor with a tightly coupled data parallel co-processor (ASProCore) and a DMA unit for moving data between memory and the co-processor. In inner loops operating over large sets of data, the programmer can program it to move strips, blocks or patches of data on and off the chip. Applications consist of a single program where instructions for the RISC and co-processor can be freely intermixed.

The Linedancer, may be programmed using one or both of the following methods:

- *Application Programming Interface (API)*
  An API provides a suite of specialised functions that can be used to build solutions for particular applications. The functions are callable by a standard C or C++ program without the need for the programmer to have specialised hardware or computer architecture knowledge.

- *Using an extended version of C*
  Linedancer processors can also be programmed using an extended version of C, which has additional language statement that allow operations to be performed in parallel on all the data elements stored in the processors.

As indicated in Figure 3, which depicts the software and hardware layers of Linedancer, the overall programming environment includes: software libraries that provide functions supporting mathematical operation like square root, raise to a power etc. - standard C functions for I/O, memory management etc. - and hardware interface functions for caches control, DMA programming, interrupt handling etc.

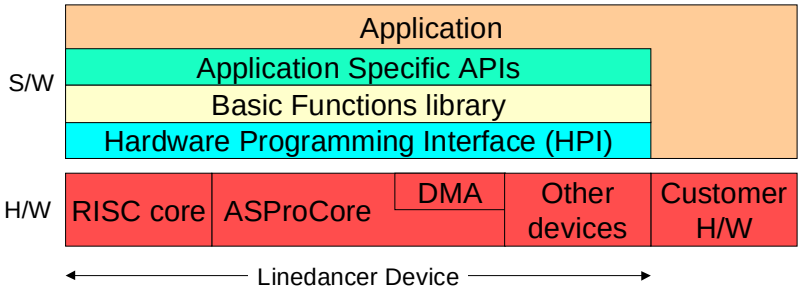


**Figure 3: Linedancer software**

## Performance

Table II presents the sustained performance of a Linedancer device for a number of media related processing; i.e. signal processing, image processing, graphics, etc.

| Media-related Task | Linedancer Performance |
|---|---|
| FIR, (8-taps, 16-bits) | 284.4 Msamples/sec |
| FIR, (16-taps, 16-bits) | 142.4 Msamples/sec |
| Convolution, (3x3 kernel, 8-bits) | 1400 Gpixels/sec |
| Median filtering , (3x3 kernel, 8-bits) | 600 Mpixels/sec |
| DCT/IDCT | 1100 Mpixels/sec |
| Motion Estimation, (vector for *8x8 block over 32x32 area*) | 400 Kvectors/sec |
| Line Resize, (*1024 pixel line or column, 12-tap filter, 8 sets of filters*) | 1400 Mpixels/sec |
| 3D Graphics, (geometry & light transformations) | 33 Mvertices/ec |
| 3D Graphics, (*Gouraund rendering 25 pixels/triangle*) | 25.7 Mtriangles/sec |
| 3D Graphics, (*Gouraund rendering 50 pixels/triangle*) | 12.8 Mtriangles/sec |
| 3D Volumetric Visualisation, (transformation & rendering) | 450 Mvoxels/sec |
| Colour transformation, (*RGB to CMYK*) | 104 Mpixels/sec |

**Table II: Linedancer processor performance for a number of media processing tasks**

## Summary

The Linedancer digital signal processor with it software programmable very high performance is a revolutionary component to solve the computing demands of media applications in areas such as digital TV, digital imaging office products, digital wired and wireless communications and networking. The extremely high degree of data parallelism available in a single device together with associative processing of data information position the Linedancer in a position to exploit the ever increasing demand for sophisticated high-performance devices – instead of the dual device solutions employed with competing DSP devices. With full programmability in high level languages and application programming interfaces, the Linedancer creates a flexible solution for the constantly evolving media standards and applications. The Linedancer's *scalable* architecture and its *portable* media processing code enables a wide range of products for handheld, consumer, and professional environments.

## References

[1] Diefendorff, K., and Dubey, P., "How Multimedia Workloads Will Change Processor Design", *IEEE Computer*, Vol. 30, No. 9, pp:43–45, September 1997.

[2] Peleg, A. and Weiser, U., "MMX Technology Extension to the Intel Architecture", *IEEE Micro*, Vol. 16, No. 4, pp: 42–50, August 1996.

[3] Phillip, M., "A Second Generation SIMD Microprocessor Architecture", *The Proceedings of Hot Chips X Symposium*, August 1998.

[4] Ranganathan, P., S. Adve, S. and N. Jouppi, N., "Performance of Image and Video Processing with General-Purpose Processors and Media ISA Extensions", *The Proceeding of the 26th International Symposium on Computer Architecture*, May 1999.

[5] Krikelis, A., "A Modular Massively Parallel Computing Approach to Image-related Processing", Proceedings of the IEEE, Vol. 84, No. 7, pages 988-1004, July 1996.