Game Processors to Speedup Number Crunching

Francisco Oliveira[1]

[1] Curso de Especialização em Informática
Universidade do Minho, 2006-2007
{Francisco.Oliveira.ext}@gmail.com

Abstract. With a fast growing market value, the video game industry is strongly becoming one of the big technological trend setters. Very large resources are spent with research and development. Current software and hardware are some of the most innovative technologies available. One of the must promising is the new Cell Broadband Engine developed together by Sony, Toshiba and IBM. The main goal of this $400M project is to deliver a processor architecture for the upcoming Sony Playstation3. Meanwhile, processor makers are struggling to find new technologies to increase performance, the scientific community faces slow gains in performance and increasing power demands of traditional processor architectures used in high performance computing. There is an increasing need to search for alternative architectures that enhances today's limitations. In this work, we will show some of the outstanding features of the new Cell processor for use in high performance computing and what new challenges emerges from adopting an heterogeneous processor environment.

Keywords: High Performance Computing, Cell BE, Hybrid Computing, IBM Roadrunner, Heterogeneous Multi Computer System, Game Processor

1 Introduction

1.1 The Videogame Industry

According to reports from the NPD Group and DFC Intelligence, last year, the video game industry had $12.5B in sales in the US. Previewing a 15% annual growth for the upcoming years, it is expected a market value of $17.2B in 2010.

High performance computing had $10B sales during last year, in the entire world. Projections show a market value of $14.3B for 2010.

This clear shows that today's video game market clearly outperforms high performance computing. Standing has one of the fastest growing industries in the US we cannot be surprised to see an increasing technology transfer from video game to other IT domains.

Having a steady annual growth from the very beginning (mid 70's), this business had stamina enough to not allow a clear dominator (opposite to the Microsoft/Intel dominance in the desktop/server business).

Additional key aspects that fit in current processors business demands are the outstanding price/performance ratio and power conscientious of the video game products. We have to remember that one of the most important markets segments of the video game business is portable consoles. For instance, Nintendo portable devices clearly outperform Nintendo's GameCube console sales.

1.2 Processors: Major Players and Market Trends

Facing sustained market dominance from just one manufacture over the past twenty years, the business is struggling to find ways to reverse the increasing vertical aggregation. The biggest market segments are dominated by only one company:

- General-purpose processors: among the best successful business in the IT industry, Intel is unbeatable. AMD and IBM seems to have a technology advantage but they fail to convert that into market share;
- Game processors: clearly dominated by IBM. Next generation video consoles from Microsoft, Sony and Nintendo, all of them, have processors design by IBM and build around Power technology. Although, this apparent monopoly, each of the three main competitors have his own product designed for their specific platforms;
- Graphics processors: ATI and NVIDIA have both very popular products. Most sales for PC gaming;
- Processors for embedded systems: one of the most competitive. We can see architectures from MIPS, Intel, ARM, etc, in a very wide range of products.

Over the past years, some strong establish laws and trends began to decline. One of the strongest was that for each eighteen months we could double the number of transistors in a processor's chip. The last months seem to show that Moore's Law is no longer valid.
Also, processor's clock speed is slowly reaching physical limits.
Cache-based memory architectures are driving us, has W. McKee et al [2] previewed back in 1994, with memory access time increasing in a 7% rate each year, and processor's speed 80%, we are very likely to reach the "Memory Wall" in less than a decade. After hitting that "wall", is pointless to increase processor speed. Systems performance will be dominated by memory access time.
Very recent improvements like multi-threading and deep pipelining were dropped or used with care, because they have a small impact in performance and have a very poor power performance.

2 The Cell Broadband Engine Processor

Around year 2000, Sony Computer Entertainment Incorporated (SCEI) realized that traditional processor architectures could not deliver the computational power required for future entertainment products. The goal was to have a system with 100 times the performance of Sony's Playstation2 video game console.
Preliminary studies showed the need to include aspects from entertainment systems, broadband communications and supercomputer structures.

2.1 Design Goals

One of the key elements that drove early design of the Cell processor was to incorporate reconfigurable I/O and vector elements. This was one of the must crucial aspects of the project: to design a system suitable to high volume production but reconfigurable enough to fulfill a wide range of platforms.
According to the first project draft, the system must deliver outstanding performance on game and multimedia applications. For this to happen, extra efforts were made to design micro-architectures that decrease the pipeline depth and enhanced pipeline slots available.
It must also have real-time responsiveness to the user and the network. Real-time applications and real-time operating systems needs vary slightly from traditional batch processing. In batch processing you only need to keep the processor busy. With real-time environments you need to continuously update contextual data (visual, sound, etc) and user experience requirements.

2.2   Design Concept and Architecture

The keys elements of a Cell BE processor are:
- One or more PPE (Power Processor Element). This is a PowerPC 64-bit processor. The primary function of a PPE is to perform system management of all the Cell BE components.
- One or more SPE (Synergistic Processor Element). The SPE has a simpler design than the PPE. Basically, the SPE is a SIMD (Single Instruction Multiple Data) unit to perform high computational density tasks. The PPE also have one vector/SIMD extension.

The SPE unit concept pretends to fill the gap between a general processor and special purpose hardware. General purpose processors are design to have an average performance with every task an application will need. Special purpose hardware, on the other hand, they are very fast to perform some specific tasks but very difficult to change basic implementation aspects. With a SPE, we have high performance and application/programming level control of all system functionality.

To overcome power inefficiency of cache-based memory hierarchy computer models, Cell implements software-controlled memory architecture. The Memory Flow Controller (MFC) performs data transfer, provides protection and synchronization between main storage and associated local storage using dedicated DMA engines. MFC can execute a sequence of DMA commands, autonomously and asynchronously. With this software controlled memory, we can drastically improve off-chip memory access using application level information by scheduling main memory requests.

MFC works closely with EIB (Element Interconnect Bus). It has four unidirectional data/command rings, associated with twelve ports, one for each element. Is responsible for delivering about 25.6 GB/s (@3.2GHz) of memory bandwidth, and IO bandwidth of 35 GB/s inbound and 40 GB/s outbound.

Local storage is implemented in each SPE, and each SPE only have access to his own local storage. We may consider it a two level register file. One with a 128 x 128 bit single cycle register, and the second with a 16K x 128 bit six cycle register. SPE only have direct access to the 128 x 128 bit registers.

By adopting this memory model and with very simple pipeline-scheduling rules, we can easily calculate and improve code performance.
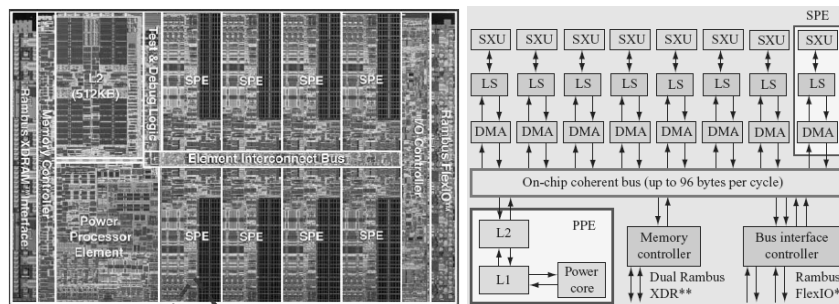

Figure 1. Cell BE photo and block diagram. Source: [11].

2.3   Cell Technology Road Map

Future upgrades are planned for the actual processor. Current version has 1 PPE + 8 SPE, 90 nm CMOS SOI[1]. During 2008, a new SPE unit will be introduced with an enhanced double precision unit (~100 gigaflop/s) and smaller SOI (65 nm). It is expected that around 2010, the Cell BE will have ~1 teraflop/s, 2 PPE + 32 SPE and 45 nm SOI.
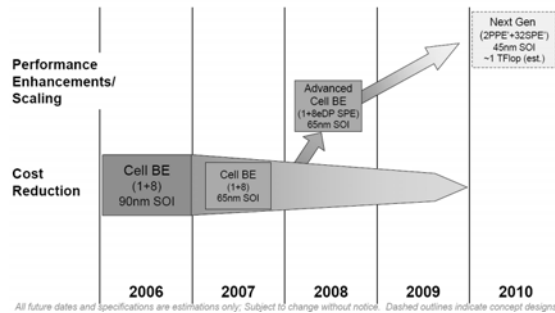

Figure 2. Cell BE Roadmap Version 5.1. Source: [9]

2.4   Cell BE and Scientific Computing

According to Samuel Davis et al [6], Cell processor outperforms clearly, both in performance and power consumption, most common commodity processors available. Using standard scientific computational algorithms, like GEMM (Dense Matrix-Matrix Multiplication), SpVM (Sparse Matrix Vector Multiply, Stencil computations and 1D and 2D FFT (Fast Fourier Transform), against Cray X1E vector processor, super-scalar AMD Opteron and VLIW Intel Itanium2, we can clearly see in the following table, the performance advantages of Cell technology.

Cell w/	Speedup vs.			Power Efficiency vs.		
eDP	X1E	AMD64	IA64	X1E	AMD64	IA64
GEMM	3x	12.7x	9.5x	9x	28.3x	30.9x
SpMV	>2.7x	>8.4x	>8.4x	>8.0x	>18.7x	>27.3x
Stencil	5.4x	37.0x	17.7x	16.2x	82.4x	57.5x
1D FFT	2.3x	10.6x	7.6x	6.9x	23.6x	24.7x
2D FFT	2.3x	13.4x	30.6x	6.9x	29.8x	99.5x

Table 1. Double precision performance and power efficiency of Cell processor with enhanced double precision design vs. main competition. Source: [6].

---

[1] SOI - Silicone on Insulator. Silicone-insulator-silicone substrate in place of conventional substrates in semiconductor manufacturing. Both AMD and IBM use this technology in their processors. INTEL processors use a conventional silicon substrate.
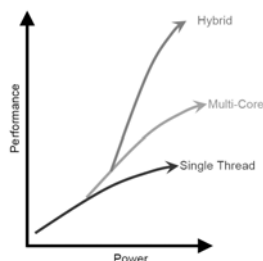
3 Hybrid Computing

Figure 3. Microprocessor Trends. Souce: [10].

One of the most important issues faced when implementing Cell technology in HPC, is the increased heterogonous processor environment.

Historically, hybrid computing was used to classify systems with the ability to make digital and analog calculations. A new kind of hybrid systems or heterogeneous multi-computer systems (HMCS) emerges from specialization of system architectures.

Today, we have the following processors families:

- General purpose processors, used in desktop and sever computers, usually with a very wide instruction set, performing most common computational tasks. Often called scalar processors, although some have vector instructions;
- Embedded system processors, executing one or few pre-defined tasks, usually with very precise algorithms.

A way to surpass today's single thread and multi-core limitations is to bring specialized processors to the same computational entity.

Although this seems to be a trivial solution, we have to take into account, the increase complexity of software and compiler technologies. Processor tasks assignments and memory management need to be managed in the application level.

But, a new world of sophisticated algorithms and specialized instructions may emerge, predicting the best way to optimize such complex environment.

4 The IBM Roadrunner Project

IBM targets more then 1.7 petaflop/s peak performance for the upcoming Roadrunner supercomputer that will be installed in The Los Alamos National Laboratory (LANL) - Department of Energy, and making it almost six times faster that today's front runner: The IBM BlueGene/L – eServer Blue Gene Solution (280.6 teraflop/s)[2] running in The Lawrence Livermore International Laboratory.

The system will be build around IBM System x3755 servers with 16,000 AMD Opteron 64 processors and IBM BladeCenter H systems with more than 16,000 Cell BE processors running the Linux operating system.

IBM System x3755 servers and AMD Opteron 64 processor will form the Base System Clusters connect with an InfiniBand network. All AMD Opteron will deliver 76 teraflop/s.

---

[2] Source: The 28th TOP500 List, www.top500.org, Nov-2006.

Each rack will spend 16 kW, with 1kW for each x3755 and 5kW for each BladeCenter.
The system will be responsible for LANL weapon simulations.

4.1   Project Milestones

The system will go live during 2006 and will be prepared to a major upgrade in 2007 and 2008.
The first version will have Base System Cluster plus one cluster with 7 Cell nodes for development and testing.
During 2007, there will be 6 additional nodes with improved Cell software and blades.
During 2008, all Cell blades will be implemented with the latest Cell specification. The final version will have to deliver a sustained 1 petaflop/s.

References

1. Randal E. Bryant, David O'Hallaron. Computer Systems, a Programmer's Perspective. Pearson Education, Prentice Hall;
2. W. Wulf and S. McKee, Hitting the Memory Wall. Implications of the Obvious. ACM Computer Architecture, News 23, No. 1, 20–24 (March 1995).
3. http://www-03.ibm.com/press/us/en/pressrelease/20210.wss
4. Cell Broadband Engine Architecture v1.0. Sony, Toshiba and IBM, 2005.
5. J. A. Kahle, M. N. Day, H. P. Hofstee, C. R. Johns, T. R. Maeurer, D. Shippy. Introduction to the Cell multiprocessor. IBM Journal of Research and Development, July/September 2005.
6. Samuel Williams, John Shalf, Leonid Oliker, Shoaib Kamil, Parry Husbands, Katherine Yelick.: The Potential of the Cell Processor for Scientific Computing. Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720.
7. Ken Koch & Paul Henning. Beyond a Single Cell. Cell Workshop, University of Tennessee, October 25, 2006, Los Alamos National Laboratory.
8. Robert W. Crandall, J. Gregory Sidak. Video Games: Serious Business For America's Economy. Entertainment Software Association, 2006.
9. Cell BE Roadmap Version 5.1, STI Design Center, 7-Aug-2006.
10. H. Peter Hofstee, Ph. D., Real-time Supercomputing and Technology for Games and Entertainment. IBM Systems and Technology Group, SCEI/Sony Toshiba IBM (STI) Design Center, Austin, Texas, 2006.
11. Cell Broadband Engine Architecture, v1.01, IBM, Out-2006.