Architectural Underpinnings of the Cell Broadband Engine

Nabeel Ali Memon [1], Dr. Imran Amin [2]

[1]MSCS SZABIST
Karachi Pakistan
nabeelalimemon@gmail.com

Abstract: *The unique architecture of the Sony's PlayStation 3 gaming console provided us some extreme improvements and introduced some new techniques in game and graphics development. This gaming console is no ordinary. It is powered by the Cell Broadband Engine. The Cell based processor design of IBM Corporation helped them shape a modern-day supercomputer (codename: Road-Runner) which is the most capable one yet. In this report the author will try to explore some pointers to new possibilities opened up by this unique computing architecture and how it can be harnessed, based on the revision & research of some modern technologies.*

Keywords: *Cell Broadband Engine, PowerPC Element, Synergistic Processing Element, Linux kernel, game development, terrestrial research, protein-folding, cluster computing, super-computing.*

1. CELL IN A NUTSHELL

It was the year 2000 when the Sony Computer Entertainment collaborated with IBM and Toshiba Electronics to seed a research on a new design of multi-core chip. Dr. Peter Hoftsee of IBM research division was made the chief architect of the project Cell Broadband Engine Architecture. The Processor was ready for development work in next couple of years and the very first commercial product which had been powered by the Cell processor was the Sony's PlayStation 3 gaming console. The console was marketed well to keep up with the major market-share of Sony in gaming consoles.

The PlayStation 3 console was never meant to be a scientific number crunching machine but the design of the Cell Broadband Engine was so marvelous that the IBM and Mercury soon put their next production-line of servers and hardware chips based on the Cell. IBM's QS20 and QS21 and Mercury's dual-Cell series of servers are well known in the market. Cell, basically, consists of 9cores in total. 1 is the PowerPC 970 Instruction Set Architecture (ISA) compatible dual-threaded core known as the PowerPC Processing Element (PPE) and the rest of the 8 cores are the Synergistic Processing Elements (SPEs).

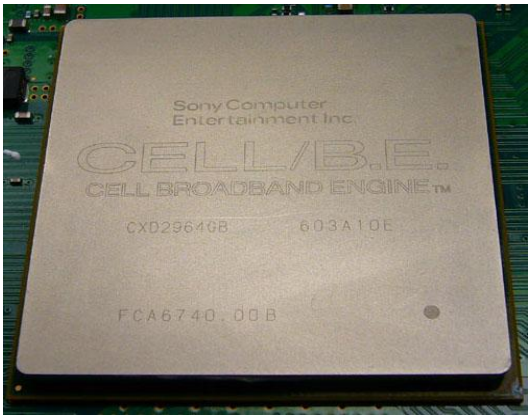

Fig. 1: A packaged Cell Broadband Engine

1.1. Cell in the Super Computer

It was in early 2008 when IBM announced their new Super Computer named Roadrunner which was designed for the total processing power of 1 petaflops and peak processing power of 1.7 petaflops[1]. Roadrunner tops the chart of world's top 500 super computers and is only one of it's kind. It is also efficient at power consumption. Basically, Roadrunner is a hybrid design based on 122400 cores (6120 Opteron (2 cores) + 12240 PowerXCell 8i (9 cores) = 122400 cores)[2] in total (the count used for benchmark provided by the Linpack[3] top500 list). The Cell processors being used in it are slightly modified, but not the whole lot, to be named the IBM's PowerXcell 8i. So, if we count the total number of cores that a Roadrunner has, the result is ought to be:

> 13824 Opteron cores + 116640 Cell cores = 130464 cores for both the computing nodes and the operation nodes.

The AMD uses it's general purpose Instruction Set Architecture (ISA) for the ease of programming while the Cell processors provide acceleration. They are not easy to program though. Thus the whole Roadrunner could not be built with simply using an assembly of Cell processors. Roadrunner is the heaviest super computer built yet.

1.2. Playstation 3 Cluster as the Super Computer

If assembled to make a cluster, playstation 3 consoles may provide the performance that is better than the Blade servers and yet at very low cost. Recent usage of the playstation 3 consoles' cluster to serve as a number crunching machine equivalent to super computers has emerged. A recent paper on playstation 3 based cluster performance and constraints has also been published at the IEEE[c]. The paper concludes playstation 3 as the reliable and the cheapest super-computer-performance cluster[4].

2. HARDWARE ADVANCEMENTS

Microprocessor architects have, over the last couple of decades, been driven by increasing the clock rate of CPU chip for higher performances but according to some architects, considering the constraints, we are gaining approximately only 40% speedup every time we double the number of transistors on a chip[5] means there are not much hopes left. But, recent time has seen the industry jumping to the Multi-core architecture which, by calculation, doubles the speedup with same clock rate of CPU chips used. Still, Cell Broadband Engine is quite different multi-core as it is heterogeneous. It has one dual-threaded Power Processing Element (PPE) which is compatible with the PowerPC 970 Instruction Set Architecture (ISA), has 32 KiB L1 cache and 512 Kib L2 cache and has a clock rate of 3.2 GHz. The other units of a Cell are 8 Synergistic Processing Elements (SPEs) optimized for the SIMD calculations which individually gives out, theoretically, 25.6 GFlops at 3.2 GHz[6] when doing single-precision calculations. The PPE has the potential to load a 32/64 bit operating system and is good at task-switching while SPEs have an entirely different model as opposed to the PPE. Every SPE has a Local Store (LS)[7], which is 256 KiB, to which PPE transfers code and data to operate or work on. LS addresses have an alias in the PPE address map.

Basically a pointer variable to the data structure, in main memory, is passed by PPE to the SPE which in return uses that pointer to issue Direct Memory Access (DMA) commands to bring that data structure in it's LS. After working on that data, SPE again issues a DMA command and the transfer is again coherent in the system in accordance with the PowerPC memory ordering rules and uses the Element Interconnect Bus (EIB). Equivalents of the Power processor locking instructions, as well as a memory-mapped mailbox per SPE, are used for synchronization and mutual exclusion. Thus enhancing the multithreaded application development on this platform.

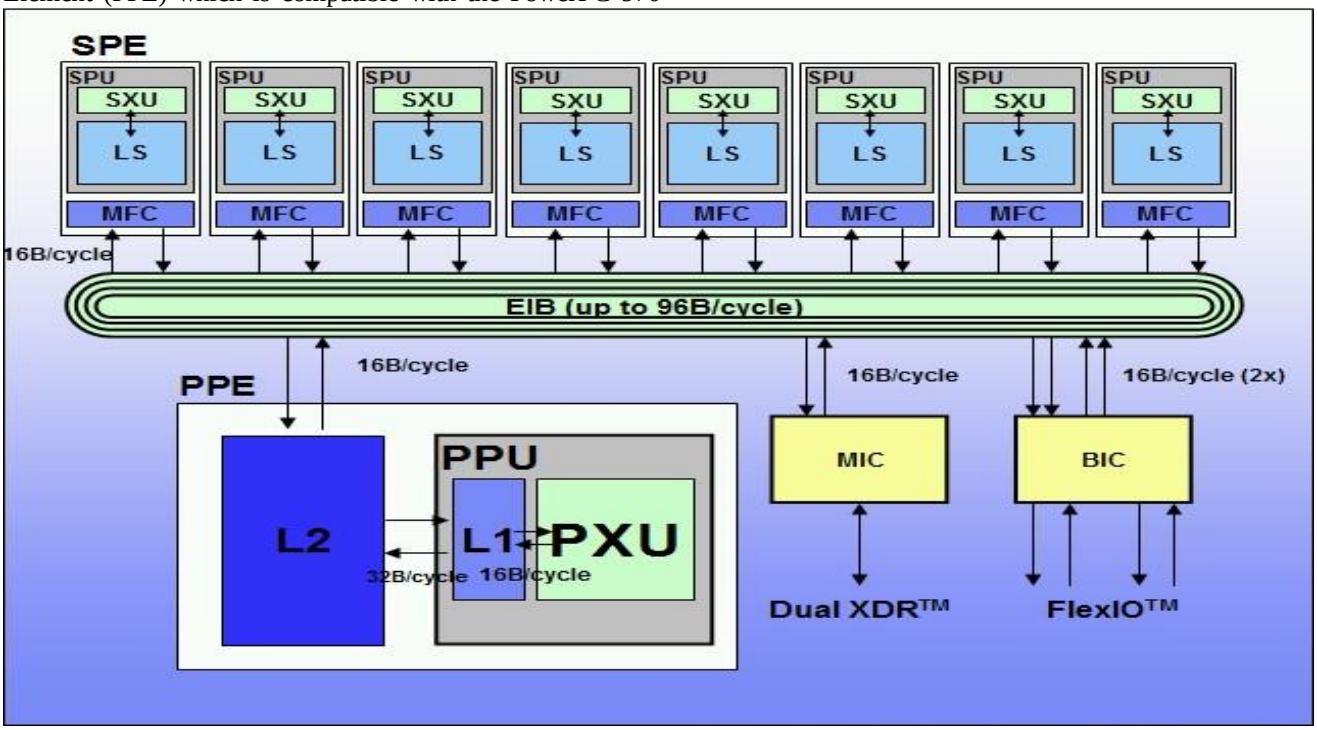

Figure 2: Block diagram of a Cell Broadband Engine

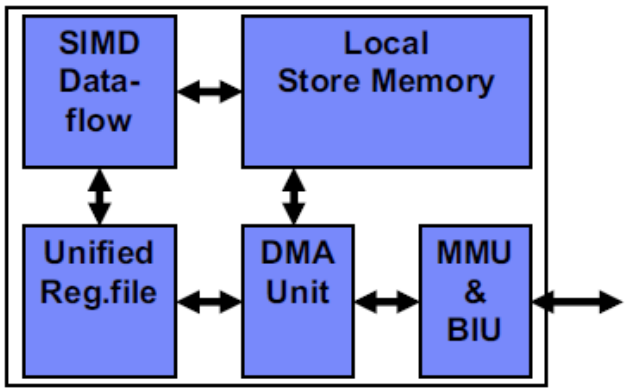

Figure 3: SPU internals

3. PROGRAMMING ENVIRONMENT

To program applications for the Cell platform, IBM has been providing it's Software Development Kit (SDK) for Multi-core Acceleration. Currently, it is version 3.0 while 3.1 is in making. It comes as a bundle of multicore compilers (for SPU and PPU) and profiler along with the core libraries (like Math and SIMD[lib]) and Accelerated Library Framework (ALF). It also includes many other tools like Cell System Simulator [8].

3.1. Linux based Development

Linux has been a very famous open source Operating System as well as a complete development platform. Many big open source projects got courage to begin after the enormous success of Linux like Apache and programming Languages like Ruby. Similarly, Linux has been powering the Cell software development platform too. IBM's SDK[9] for multi-core acceleration works only on Linux OS as it is the most flexible, scalable and open architecture.

3.1.1. Cell Integrated Development Environment

IBM's SDK provides a complete Integrated Development Environment (IDE)[10] for the Cell based software programming. This IDE has been developed on the Eclipse IDE Plugin Development Environment and is only compatible with the Eclipse IDE version 3.2 as of yet. Working on this IDE eases the task of repetition and organizing the software code helping in the development of big code-base. When working on the Cell System Simulator (another software component developed by IBM providing the simulation of Cell hardware for software development on x86 platform) all the other software components can be directly handled from within the Cell IDE with the click of a button.

3.1.2. IBM Software Development kit

Since the release of the Cell, commercially, in playstation 3 game consoles, IBM has let users to use this game console as a completely programmable development environment through the release of it's Cell Software Development Kit (SDK) along with some extra developer tools released for the Cell programming eco-system. Every SDK release is compatible with some PowerPC release of the Fedora Operating System (Red Hat's open source Linux distribution) that is announced by the IBM. Currently, many Research and Development Centres host the SDK and other Open Source tools such as Barcelona Super Computing Facility[11].

3.2. IBM XL C/C++ compilers

IBM XL C/C++ for Multicore Acceleration for Linux is an advanced, high-performance cross-compiler that is tuned for the Cell Broadband Engine Architecture (CBEA). The XL C/C++ compiler, which is hosted on an x86, IBM PowerPC technology-based system, or an IBM BladeCenter QS21, generates code for the PPU or SPU. The compiler requires the GCC toolchain for the CBEA, which provides tools for cross-assembling and cross-linking applications for both the PowerPC Processor Element (PPE) and Synergistic Processor Element (SPE).

4. CURRENT MAINSTREAM APPLICATIONS

Currently, Playstation 3 game console is the only commercially available product having the Cell microprocessor in it which is ready to be harnessed by the usage of Linux and the programming SDK provided by IBM. Round the world, many applications, other than games, have been developed and are being research at all levels.

4.1. High-end Gaming

Many big game titles like Rage by ID Software (having the very famous John Carmack on board) and Fifa 2009 and Metal Gear Solid 4, to name a very few, have been released exhibiting the usage of advanced 3D graphics rendering and the technology used behind it to make it almost real. The heavy number crunching for physics engine's calculations and simultaneous heavy graphics rendering that has been used in these big titles, and many to come, could only be possible with the power of the Cell processor's number crunching and 3D hardware acceleration. Many new game libraries have been developed and some have been published too on the playstation 3 game development.

4.2. Folding@Home

Playstation 3 has not just been used for gaming but it has also been serving the humanity by having any user be able to participate in the medical research by just donating their free processor clocks. That's due to the Distributed Computing Software developed by the Stanford University and named as "Folding @ Home" project[12] which aims at understanding the Protein's folding, misfolding and the diseases related to it. Since these calculations require heavy numerical computing, only super computers were considered to have ability to help out, but the release of the Cell processor, commercially in the playstation 3, led the project to use Cell's computational powers to feed the project. All a

participant has to do is to download a small piece of software, fire it up in their playstation 3 consoles and that is it.

4.3. Astrophysics

Playstation 3 is also serving Astrophysicist Dr. Gaurav Khanna who once used to rely on his funds for his research on the Gravitational Waves' existance and use super computers of National Science Foundation (NSF) deployed throughout the United States. Normally the processing time of those super computers was so expensive that all his grant-money was spent on it. But last summer, after the release of playstation 3, he asked sony for the grant of simply 8 consoles and claimed that he could use them like a super computer. He tweaked his code for the Cell Processor and finally started having all that he wanted just from a cluster of 8 playstation 3 consoles[13]. According to him, he was greatly satisfied with it as it helped him pursue his research without requiring any heavy expenditure of grants and never required any other super computer.

CONCLUSION

With all the programming experience, gained from working on the Cell Broadband Engine Architecture, author concluded that although currently the Cell processor is a specialized environment targeted towards gaming and special-purpose computing, but it is surely the next general-purpose computing platform of the near-future. Recently Toshiba Electronics has launched a laptop having a cell processor, for the first time in a PC environment, that assists the main conventional core-architecture based Intel's processor in that machine. But this product promises to be really good at media-specific encoding/decoding tasks that are faced in everyday computing. All in all, GPU computing tends to bear the capacity that it may replace conventional CPU on the desktop as the future prospects more and more graphics and physics engines changing the way humans interact with the machines.

Currently, programming for the Cell processor may be quite difficult but it is already serving many big projects. For the ease of use as a development environment, projects like Octopiler from IBM and countable other ideas seem promising. IBM is very much interested in developing the Cell programming eco-system and is investing a lot of budget for this purpose. At this moment there is an open room for compiler and application writers to port several favorite languages and applications to the Cell hardware. Upto this point in time, all the major development on the Cell platform have been done through using open source tools like Linux Kernel, GNU toolchain and the IBM SDK which may attract potential developers worldwide.

This research may be concluded in the light of John Carmack's (famous as author of Doom and Quake series) conjecture that the future is for floating-point kernels.

ACKNOWLEDGMENTS

I would like to give thanks to Dr. Imran Amin for his helpful suggestions, proper guidance and supervision.

REFERENCES

[1] "Roadrunner: Science, Cell and a Petaflop/s", page 3, Los Alamos National Laboratory and IBM (2008-06-18). COMPLETE INFO NOT GIVEN

[2] Wikipedia. "Cell (microprocessor). Internet: http://en.wikipedia.org/wiki/Cell_(microprocessor). April 3rd, 2009. [April 15th, 2009].

[3] LINPACK. Internet: http://www.netlib.org/linpack/. [April 15th, 2009].

[4] Blagojevic, F.; Curtis-Maury, M.; Jae-Seung Yeom; Schneider, S.; Nikolopoulos, "Scheduling Asymmetric Parallelism on Playstation 3 Cluster", D.S. Cluster Computing and the Grid, 2008. CCGRID apos;08. 8th IEEE International, 19-22 May 2008, page 1, Digital Object Identifier 10.1109/CCGRID.2008.64.

[5] Dr. Peter Hoftsee, "Power Efficient Processor Architecture and The Cell Processor", 2005. HPCA-11. 11th International Symposium, 12-16 Feb. 2005, page 1, Digital Object Identifier 10.1109/HPCA.2005.26.

[6] Kurzak, J. Buttari, A. Luszczek, P. Dongarra, "Playstation 3 for high performance Scientific Computing", In Computing in Science & Engineering, Volume:10, Issue:3, May-June 2008, page 1, ISSN: 1521-9615.

[7] IBM RedBook,"Cell Broadband Engine Programming Handbook", Version 1.1, page 44, IBM Systems and Technology Group, 2070 Route 52, bldg: 330, Hopewell Junction, NY 12533-6351, April 2007.

[8] IBM Redbook, "Software Development Kit for Multicore Acceleration 3.0 Programming Tutorial", IBM Corporation, 2007, page 124.

[9] "Cell Broadband Engine Resource Center" Internet: http://www.ibm.com/developerworks/power/cell/downloads.html. [April 15th, 2009].

[10] IBM RedBook, "Software Development Kit for Multicore Acceleration 3.0 Programmer's Guide", IBM Corp, page 14, SC33-8325-02, 2007

[11] "Cell Superscalar". Internet: http://www.bsc.es/plantillaG.php?cat_id=179. [April 15th, 2009].

[12] Folding@home. Internet: http://folding.stanford.edu/. [April 15th, 2009].

[13] Bryan Gardiner. "Astrophysicist Replaces Supercomputer with Eight PlayStation 3s." Internet: http://www.wired.com/techbiz/it/news/2007/10/ps3_supercomputer. October 17th, 2007. [April 15th, 2009].