



Cell Broadband Engine™
processor – based systems
White Paper

September, 2006

The Cell Broadband Engine™ (Cell B.E.) processor

Background:

Modern processor architectures have many similarities. A central processing unit (CPU) interprets and executes instructions. A memory subsystem stores programs and data. An Input/Output (I/O) subsystem handles the movement of data and programs between the CPU and the external world, be it local or networked. Special functional units to provide additional functions, like security or special graphics output, are sometimes added to the CPU or the I/O subsystem. To increase effective memory bandwidth and lower latency, memory subsystems become more complex with multiple levels of cache.

Processor speeds improve over time, as technological improvements in manufacturing allow more transistors per unit of chip real estate, closer packing of components, etc. These improvements generate more heat and consume more power, so strategies must be implemented to address these increases

For years, many companies have relied on performance gains from increasing clock speeds of ‘traditional’ microprocessor architectures to continually deliver enhanced value and new, more competitive solutions. However, this approach has been challenged by the physical characteristics of shrinking semiconductors and by traditional processor architecture implementations.

The challenge is to keep a balanced system. Memory speed (i.e. latency and bandwidth) improvements do not keep up with processor speed, so we introduce larger, multiple level caches with exotic management strategies. Pipelines become deeper. Some of those new transistors that Moore promised are used to address these issues, and processor efficiency improves slower than transistor density.

Improved manufacturing techniques, like 7 Level Copper Silicon On Insulator (SOI) or Low K (impedance) dielectrics help, but the inefficiencies remain. As components like gate dielectrics approach real physical limits (i.e. a few layers of atoms), a radical improvement in microprocessor design is needed.

Cell B.E. Processor History:

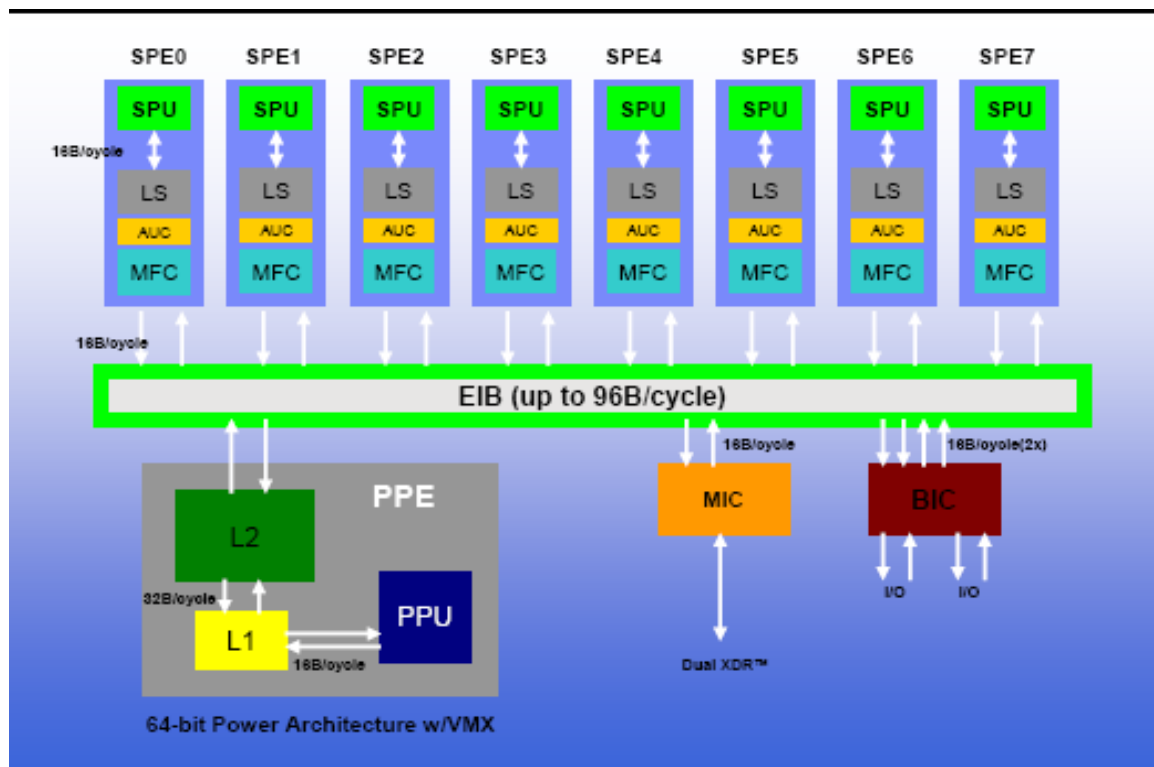
The Cell B.E. architecture is an outgrowth of a strategic alliance between Sony Computer Entertainment Inc, Toshiba Corporation, and IBM Corporation, formed in 2001. The intent was to develop a Cell B.E. processor, based on IBM® Power Architecture™, that would provide very significantly better performance in the games, media and digital content market. The STI (Sony, Toshiba, IBM) Design Center was opened in 2001 in Austin, and development began. The first technical disclosures were in February 2005 at the International Solid-State Circuits Conference.

The Cell B.E. processor is a multi-core broadband processor based on IBM's Power Architecture. A single chip contains a Power Processing Element (PPE) and eight Synergistic Processing Elements (SPEs). Each SPE has an on-chip local store. The PPE has an on-chip L1/L2 cache. Also on the chip are memory and I/O controllers.


The PPE is a general purpose 64-bit RISC processor with 2-way hardware multithreading. L1 cache for the PPE is 32KB instruction and 32KB data. Each of the eight SPEs is a special purpose RISC processor with 128-bit SIMD (Single Instruction Multiple Data) capability and 256KB local memory.

The PPE and the eight SPEs are connected with a high bandwidth Element Interface Bus.

Conceptually it is organized like this:



And looks like this:



Each SPE has a dedicated DMA controller which can support movement of data and instructions between local storage and system memory or between the local storage of different SPEs. SPEs have instructions specifically for scatter/gather operations. Coupled with the ability to have multiple I/O operations outstanding, any algorithms heavily dependent on scatter/gather should demonstrate significant speedup. Life sciences applications like BLAST may expect significant speedups from this architecture.

The very robust communication structure between SPEs is optimal for algorithms where data and/or instructions can be streamed between different levels of the memory hierarchy. For example, computation is performed on a node, the results passed to a nearest neighbor, more computation is performed in a pipelined fashion.. Any application that uses Fast Fourier Transforms successively is a natural fit for this architecture. Image processing (like seismic) and computational fluid dynamics applications should benefit significantly from the Cell B.E..

Less traditional applications, like video surveillance, is also a good fit for Cell B.E. processors. The video encoding capabilities, coupled with an Artificial Intelligence engine, might recognize problems and generate appropriate alerts.

In short, the unique architecture of Cell B.E. makes it almost ideal for compute-intensive tasks like digital content creation (rendering, image processing, simulation, modeling), physics, life sciences, and related workloads such as:

Data Manipulation

Digital media, Image processing, Video processing, Visualization of output, Compression / decompression, Encryption / decryption, DSP.

-

Graphics

Transformation from different domains (time vs. space; 2D vs. 3D, viewpoint transformation), Lighting, Ray-tracing / ray-casting.

Floating Point Intensive applications (Single Precision)

Pattern Matching

Bioinformatics, String manipulation (search engine), Parsing, transformation, and translation (XSLT), Filtering and pruning.

Offload Engines

TCP/IP, Compiler for gaming applications, XML, Network security and intrusion.

The applications list above is not exhaustive, but is meant to give a feel for the breadth of applications that will perform better on Cell B.E. architecture systems.

The Cell B.E. processor-based systems

Hardware:

The IBM BladeCenter® QS20™ is the first in a line of planned IBM offerings using Cell B.E. technology.

The blade form factor was selected because IBM BladeCenter is an industry leading technology that allows dense packaging of processors, provides I/O capability within and outside the BladeCenter chassis, and provides infrastructure like power and cooling. It simplifies management of both the hardware and software.


The mid-plane of the chassis provides communication within the chassis. External communications (to other chassis) is provided by standard I/O adapters on the blades themselves.

The BladeCenter itself is 7U high, and fits in a standard 19" rack. Six BladeCenter chassis will fit in a standard 42U high rack.

BladeCenter has been commercially available since 2004. Blades for the BladeCenter are available with Intel® processors, AMD™ processors, POWER™ processors, and now Cell B.E. processors.

The QS20 has two Cell B.E. processors (a 2-way SMP) operating at 3.2 GHz. Dual GbE controllers provide connection to the chassis mid-plane and to BladeCenter GbE switches, and a total of four 4X InfiniBand connections.

The QS20 blade is a "doublewide" form factor, so a BladeCenter chassis will hold 7 QS20s, allowing from 2 to 14 Cell B.E. processors. Communication between blades within a chassis and between blades in different chassis will be by Infiniband or Ethernet. Communications with the outside world will be the same. Here is a QS20 shown in specific detail:



Software:

The Cell B.E. Processor Based Blade Server will run the Linux™ operating system with extensions to utilize the data parallel SPEs. Each blade will have one OS image. Each blade is a 2-way SMP. Given that the base software is Linux, we can expect that middleware and applications that run in the Linux environment will run properly here. Applications compiled for the Power Architecture or the PowerPC® ISA are binary compatible with the PPE.

Performance:

Preliminary tests have been shown to accelerate some algorithms/workloads/applications to many times the speed of a traditional microprocessor. Since the compilers supplied are enhanced to provide access to the SPEs where appropriate, recompilation and code restructuring are necessary to see improvement in execution performance.

Summary

The Cell B.E. processor ushers in a new era of capabilities for Digital Media, and by extension other High Performance Computing applications. It is a multi-core multi thread design that promises to be faster than traditional microprocessors available in the market today. The core CPU is based on IBM's Power Architecture, in this design called a Power Processor Element, and it also includes 8 Synergistic Processor Elements for each PPE, operating in a data parallel mode.

Recompilation and perhaps restructuring will be needed to take advantage of the SIMD SPEs, and the resultant bandwidth and computational power may yield significant performance for some algorithms and workloads.

Two (2) Cell B.E. Processors enable the dense computing power of the IBM BladeCenter QS20 which is the first execution of a long term Cell B.E. processor-based systems roadmap. It can bring very high computational resources for many workloads such as Digital Media, Medical Imaging, Seismic, DSP ... etc.

The BladeCenter packaging furthers the goal of very high processing power with much lower power and heat.



© IBM Corporation 2006
IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
September 2006
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, BladeCenter, IntelliStation, POWER, Power Architecture, and Power PC are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of U.S. trademarks owned by IBM may be found at:
<http://www.ibm.com/legal/copytrade.shtml>

Cell Broadband Engine™ is a trademark of Sony Computer Entertainment Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

AMD is trademark or registered trademark of Advanced Micro Devices, Inc., in the United States, other countries, or both.

Intel is trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply.

Photographs show engineering and design models. Changes may be incorporated in production models. Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. The IBM home page on the Internet can be found at: <http://www.ibm.com>