

The Basics of Cell Computing Technology

Introduction

Sony Corporation develops technologies for products used in graphics operations and scientific calculation using the Cell Broadband Engine™ (Cell/B.E.) and the RSX employed in the PLAYSTATION®3, which are developed by Sony Computer Entertainment Incorporated (SCEI). This document explains the principal technologies.

Trend toward multi-core processors

Drastic improvements in performance in recent years have enabled microprocessors to handle large amounts of data at high speed. The target applications are being extended, and the amount of computation required in each field is increasing. In addition to conventional general-purpose processors, recently developed application-specific processors exploit the full potential of their performance for a given application. For example, the Graphics Processing Unit (GPU) is a special processor for processing 2D/3D graphics at high speed. This provides very high computational performance that cannot be achieved by conventional general-purpose processors.

The trend in the microprocessor market is moving to asymmetric multi-core processors, where different types of processors optimized for a specific use coexist. Considering this technology trend, Sony Corporation has developed the “Cell Computing Unit”, which uses the high computational performance provided by the Cell/B.E. and the RSX processors.

This technology provides solutions to multimedia computing, which requires large number of computations, such as those encountered in image-processing applications, like computer graphics, and scientific computations for non-image-processing applications, such as encryption technology for security (Fig. 1).

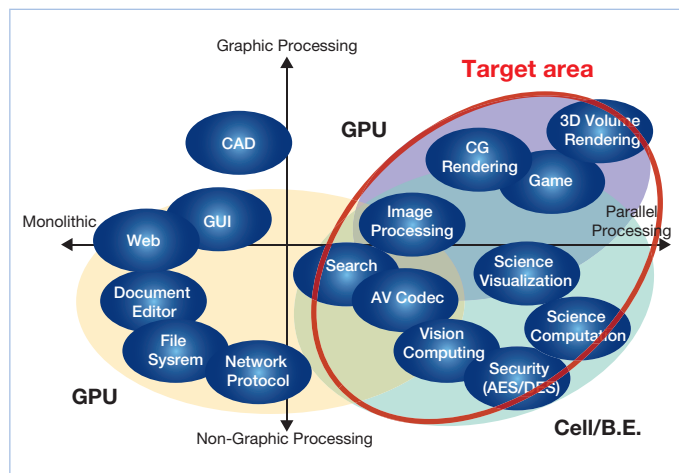


Fig. 1 Target area of Cell Computing Technology

Multi-core processor with powerful computational performance - Cell/B.E. -

The Cell/B.E. is an asymmetric multi-core processor jointly developed by Sony/SCEI, Toshiba Corporation, and IBM Corporation. This processor alone provides high floating-point arithmetic operations at 230 GFLOPS. In anticipation of real-time operation on 4K x 2K or High Definition images, the Cell processor is optimally designed for multimedia processing and usage in distributed computing environments.

Architecture of the Cell/B.E.

Two types of processor cores

Two types of processor cores are mounted in the Cell/B.E.: PowerPC Processor Element (PPE), a control-intensive processor core for general-purpose processing; and Synergistic Processor Element (SPE), a processor core for high-speed data processing. The coordinated operation of these processor cores give rise to the high level of computational performance of the Cell/B.E..

High-speed internal bus

Four ring-shaped broadband buses called Element Interconnect Bus (EIB) connect the cores of the Cell/B.E. and peripheral chips. PPE, SPE, main memory (XDR™ DRAM), and external I/O devices (the RSX, SouthBridge, etc.) are also connected to the EIB (Fig. 2).

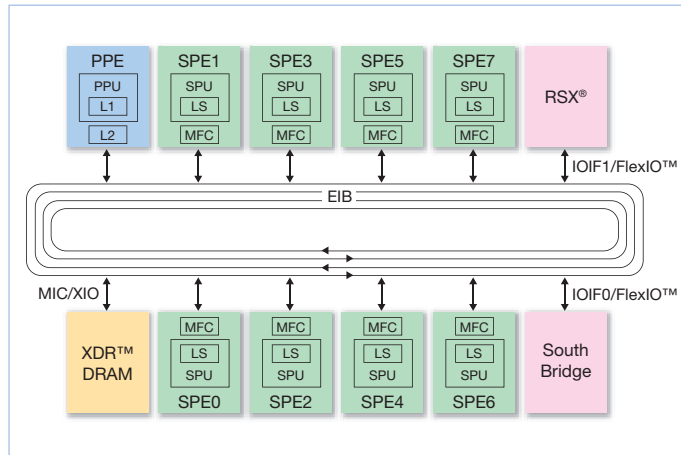


Fig. 2 Overview of the Cell/B.E. and Peripheral Chips

General-purpose processor core - PPE -

PPE is a general-purpose processor core with a 64-bit Power Architecture™ and can run existing software compatible with the PowerPC chip. One PPE is mounted in the Cell/B.E.. The PPE not only executes the operating system (input/output control to main memory, external devices, etc.) but also controls the SPEs.

L1/L2 caches

The PPE is equipped with 32 KB of L1 cache for instructions and data, respectively, and 512 KB of unified L2 cache for both instructions and data.

Hardware multithreading

Simultaneous multithreading technology is implemented in the PPE to provide two-way hardware multithreading.

Vector operation unit

A 128-bit vector operation unit supports the Single Instruction Multiple Data (SIMD) operations by Vector/SIMD Multimedia Extensions (VMX) technology.

Computation-intensive processor core - SPE -

SPE is a computationally intensive processor core that excels at multimedia processing, especially for repeated calculations. The PPE, on the other hand is optimized for complicated program control. Eight SPEs are employed in the Cell/B.E., Each SPE operates independently of the PPE and other SPEs and provides high computational performance.

Vector oriented RISC processor

The SPU instruction set architecture (ISA) supports 128-bit SIMD operations. All basic data processing is performed by this operation. Because of its very simple architecture, the SPE can realize high-speed processing where calculations can be predicted. It is very suitable for execution of an application where real time processing is required.

Abundant register files

As many as 128 of the 128-bit registers are on the SPE because of the large amount of data processing is likely to be performed by this device. By taking advantage of the abundant register files, high-speed arithmetic operations can be performed on a large volume of data.

SPE-dedicated local memory

Each SPE is equipped with 256 KB of scratch pad type memory called Local Store (LS). With no cache mounted, an SPE accesses LS directly at high speed. DMA transfer is used for data transfer between main memory and LS.

Memory flow controller

Access from SPE to main memory, PPE and other external I/O devices is performed through a unit called the Memory Flow Controller (MFC). The MFC is equipped with a DMA Controller to exchange data with the main memory. Because the MFC operates independently of program execution in an SPE, it can transfer data to external devices in parallel.

Graphics engine with high-speed I/O function - RSX -

The RSX is a graphics engine jointly developed by NVIDIA Corporation and SCEI. This device can be connected to the Cell/B.E. using a 128-bit memory interface called FlexIO™.

Programmable shader engine

The RSX processes graphics with vertex shader programs created by the user and computes fragment colors with fragment shader programs. The RSX has multiple fragment shader engines for precise graphics operation without overloading the Cell/B.E., which is difficult in fixed pipelined graphics engines.

Rendering pipeline

The rendering pipeline in the RSX consists of the host/frontend, geometry, raster, fragmentation shader/texture, L2 texture cache, 2D raster, and Raster Operation (ROP). It also includes a frame buffer to control memory access and a display block to scan out from local memory (Fig. 3).

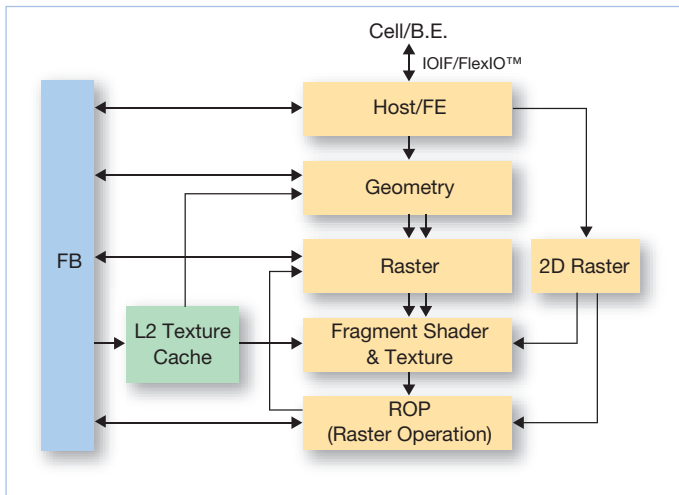


Fig. 3 Major blocks in rendering pipeline

High-speed I/O function

The Cell/B.E. and the RSX are connected via a high-speed bus called FlexIO for efficient graphics operations. The RSX also features an external I/O function with performance exceeding 4 GB/s. Therefore, it can easily handle the I/O of high-resolution video signals, such as 4K x 2K and High Definition. It can also be used in I/F for professional use by taking advantage of this high-speed I/O performance.

**Sony's New Development
-"Cell Computing Unit"**

Features of Cell Computing Unit

The newly developed "Cell Computing Unit" is equipped with the high-performance Cell/B.E. processor with high floating-point arithmetic operation of 230 GFLOPS. Additional use of the RSX graphics engine enables much faster operation. An open system is offered by adopting the Linux operating system. Its eight SPEs can be freely used to support high-level of calculations in multimedia applications.

Software configuration

Cell Computing Unit is equipped with two systems: the main system for general applications and the mini system for maintenance. Linux is the operating system so that they can be switched at system startup. Under normal operation, the main system executes user applications. The mini system for maintenance is mainly used to update firmware and drivers.

Graphics API

Graphics programs for the RSX support the subset of OpenGL 1.5, which is an industry standard for graphics APIs. They also support the Cg shader language developed by NVIDIA Corporation for RSX-native graphics control. Therefore, the RSX acceleration can be programmed at a low level.

Software development

Applications running on the main system of the Cell Computing Unit can be developed under Linux on Intel architecture (x86). A cross-platform toolchain (binutils, compilers, etc.), debuggers, performance analysis tools, and integrated development environments are used.

The SPE runtime management library for the Cell/B.E. and the graphics library for RSX hardware acceleration are currently being prepared.

Porting and optimization of the existing programs

In optimized porting of existing programs, it is important to understand how multiple SPEs are used. The basic optimization process of existing programs with the Cell/B.E. consist of three processes: (1) porting to PPE, (2) performance measurement and profiling, and (3) offloading to SPEs (Fig. 4). SPE offloading ports part of the processing in a program to the various SPEs. In the optimization process, a spiral type process is generally used, where profiling is performed again after optimization. Then, further optimization in performance is possible.

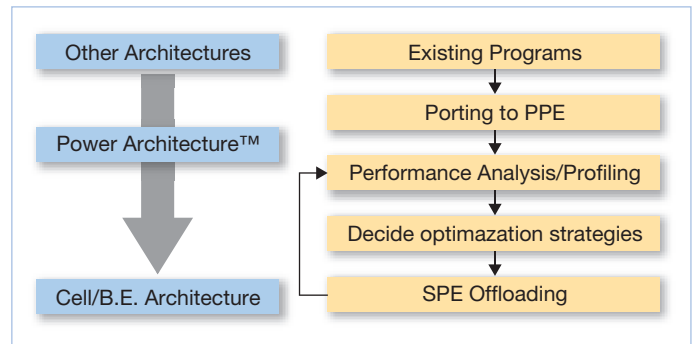


Fig. 4 Basic process of optimization

Faster programs

One can drastically improve the processing performance of programs with the Cell/B.E. by using a variety of optimization techniques. The most important optimization technique is how parallel they are. Programs are designed and implemented so that the following three parallelisms may be improved in every aspect.

Parallelism at the thread level

The PPE can execute two threads in parallel with simultaneous multithreading. When eight SPEs are used, it can execute up to ten threads in parallel. This is because, in general, the PPE thread can directly access external I/O devices and handle such interruptions as hardware exceptions, it controls the entire system, while the SPE thread mainly handles data calculations. What kind of processing is performed in each thread must be considered to distribute the load of the PPE and SPE threads.

Parallelism at the data level

The streaming data operation repeats the same processing for multiple data. In the operation, one SIMD instruction can process multiple data in parallel (Fig. 5). All instructions in the SPE, including arithmetic, comparison, and bit operation instructions, support SIMD operation. The SIMD operation is also available in the PPE by using VMX instructions extended from AltiVec technology.




Fig. 5 SIMD operation (ADD instruction)

Parallelism at the instruction level

An SPE has two long pipelines (Even/Odd) with 26 stages. The instructions executed in these two pipelines are predetermined: arithmetic operations are executed in the Even pipeline, and load/store instructions are executed in the Odd pipeline. By optimizing the instruction sequence of a program, two instructions can be issued simultaneously per cycle (Dual Issue).

Conclusion

Sony Corporation will continue to develop technologies that enable large-scale calculation at high speed, which could not be obtained from traditional servers and workstations. We will strive for development of state-of-the-art technologies in high performance computing, such as CG rendering and scientific calculations by utilizing the maximum performance of the high-end Cell/B.E. processor and the high-speed GPU RSX.

Trademarks

Sony is a registered trademark of Sony Corporation.

PLAYSTATION is a registered trademark of Sony Computer Entertainment Inc.

mental images, mental ray and mental mill are registered trademarks and

MetaSL is a trademark of mental images GmbH.

All other trademarks are the properties of their respective owners.

Disclaimer

© 2008 Sony Corporation, All rights reserved.

Reproduction in whole or in part without written permission of Sony

Corporation is prohibited.

Features and specifications are subject to change without notice.