

Cell Architecture

Disruptive Technology for Processing C4ISR Data

Mercury is preparing to deploy this technology in the defense market.

Warfighting actions are frequently constrained by the lengthy time it takes to process and analyze intelligence imagery data. By employing new multicore chip technology, recently announced by IBM, processing performance is expected to improve by a factor of 10 or better, enabling warfighters on the ground to conduct operations with current intelligence, thereby saving time and lives. The Cell architecture is a breakthrough technology that offers a major leap in processing performance and density. This paper describes battlefield imagery exploitation applications running on Cell Broadband Engine™ (BE) processors. Applications such as situational awareness, change detection, change tracking, hyperspectral image exploitation, and geo-registration are all candidates for this technology.

WHAT IS THE CELL ARCHITECTURE?

The Cell architecture grew from a challenge to provide power-efficient and cost-effective, high-performance processing for a wide range of applications, including the most demanding consumer application: game consoles. The Cell BE processor is an innovative solution designed to significantly accelerate processing for cryptography, graphics transform and lighting, physics, fast Fourier transforms (FFT), matrix operations, and scientific workloads.

A team from IBM Research joined colleagues from IBM Systems Technology Group, Sony, and Toshiba, to lead the development of a novel architecture that represents a quantum leap in performance for consumer applications.

Based on the analysis of available die area, cost and power budgets, and achievable performance, the best

approach to achieving the performance target was to exploit parallelism through a high number of nodes on a multiprocessor chip. To further reduce power, the team opted for a heterogeneous configuration with a novel SIMD-centered architecture. This configuration combines the flexibility of an IBM 64-bit Power Architecture™ core with the functionality of performance-optimized synergistic processing element (SPE) SIMD cores.

IBM Research grew its partnership in the development of the broadband processor architecture beyond its initial definition. During the course of this partnership, members of the original Cell architecture team developed the first SPE compiler, which was a guiding force for the definition of the SPE architecture and the SPE programming environment, and sample code to exploit the strengths of the broadband processor architecture. The extended partnership led to further contributions, including the development of an advanced parallelizing compiler with auto-SIMDization features based on IBM XL compiler technology, the design of the high-frequency Power Architecture core at the center of the Cell architecture, and a full-system simulation infrastructure.

The Cell architecture is not limited to game systems. Other future uses may include HDTV sets, home servers, game servers, and supercomputers. The Cell architecture is also not limited to a single chip, but is a scalable system.

Mercury has announced the Dual Cell-Based Blade, which leverages the investment in the high-performance Cell architecture. Mercury is actively developing other products that incorporate Cell technology and that are targeted at the defense C4ISR market. Specific product designs, capabilities, and availabilities will be made public in the near future.

Cell-Based Processing Can Accelerate Battlefield Applications

Applications for Cell-based processing systems can be either real-time or non-real-time. As multiple UAVs are concurrently deployed, the capability to capture and exploit their sensor data (such as video) at the tactical level becomes a real-time processing challenge. A Cell-based tactical ground station could ingest multiple (uncompressed) video streams, performing functions on each independent stream or a fusion of streams such as video stabilization, image mosaicing (tiling), geo-registration, motion analysis, target tracking, change detection, fusion with other sensors, and video compression.

Based on Mercury's initial analysis, all these functions could take advantage of the Cell architecture to deliver significant performance gains over existing microprocessors. This processing capability of the Cell BE processor would allow the development of a tactical ground station with significantly lower space, weight, and power (SWAP) than other systems, potentially opening up new opportunities and CONOPs for both existing and future sensor assets.

Cell-based systems for applications that are non-real-time, but have both high throughput and low-latency requirements, could also benefit. Computationally intensive applications such as sensor fusion could take advantage of the Cell BE processor. With the proliferation of an increased number in both type and quantity of high bandwidth sensors, image and intelligence analysts are overwhelmed with data. Through the use of high-performance Cell-based processing to perform functions such as sensor fusion and intelligence fusion, a clearer picture of the battlefield would emerge, ultimately resulting in actionable intelligence for the war fighter.

The Cell BE processor delivers unprecedented absolute processing performance in a microprocessor, but is also the highest density processing in terms of space and weight. Mercury is currently developing a Cell-based system* for delivery in 2006 that will be over 60 times faster in terms of raw GFLOPS per cubic foot than today's Intel® Itanium® 2-based SGI® Altix® 350 cluster.

* In a 5U rack-mount configuration, this Mercury system code-named "Turismo" can support 16 Cell BE processors (each operating at 3.0 GHz with a raw GFLOPS rating of 200) with a total of 25 TFLOPS of computational power in a 6-foot rack. This compares to an SGI Altix 350 system, which in a 2U rack-mounted module supports two Intel Itanium 2 processors (each operating at 1.6 GHz with a raw GFLOPS rating of 6.4). Therefore, the SGI Altix 350 system in a 6U (3 dual processor modules) rack supports a total of 38.4 GFLOPS. Normalizing, the SGI system is 6.4 GFLOPS per 1U rack unit, while the Mercury system is 460 GFLOPS per 1U rack unit, which is an astounding 72 times higher per 1U rack unit.

Cell BE Processor: A (Very) Disruptive Technology

The Cell BE processor, at about 200 GFLOPS, is one order of magnitude higher than existing microprocessors in single-precision floating-point performance (see Table 1).

Although processor performance has been steadily improving over time, such a dramatic leap in processing capability is unprecedented. The Cell BE processor's performance level can enable the realization of a new class of computationally intensive applications in intelligence, surveillance, and reconnaissance (ISR). As significant new developments in sensor technology deliver more data and sensor types, and as new algorithms such as data fusion mature, they can take advantage of the Cell BE processor's increased processing power. Sensor types with high data rates combined with unchanged bandwidth restrictions call for the exploitation of data to occur closer to the sensor. Cell-based systems that can access and process this data can enable an accelerated processing timeline with lower latency.

Table 1. Single-Precision Floating-Point Performance

Year	Processor	MHz	Single-Processor GFLOPS
1987	Weitek XL	40	0.015
1990	Intel i860XR	40	0.080
1996	Motorola PowerPC 603e	100	0.200
1996	Analog Devices SHARC®	40	0.120
2000	Motorola 7410 PowerPC	400	3.200
2002	Intel Pentium 4 Xeon®	2200	4.400
2003	Freescale™ 7447 PPC	1000	8.000
2004	Intel Itanium 2	1500	6.000
2004	IBM 970 PowerPC	2000	16.000
2006	IBM Cell BE Processor	3000	200.000

The Cell BE processor is a heterogeneous multicore chip that consists of an IBM 64-bit Power Architecture™ core called the Power™ processing element (PPE), augmented with eight specialized co-processors called SPEs (synergistic processing elements). It is based on a novel single-instruction, multiple-data (SIMD) architecture, which is designed for data-intensive processing like that found in cryptography, media, and scientific applications. The system is integrated by a coherent on-chip element interconnect bus (EIB).

In this architecture, the SPE accelerators operate from a local storage that contains instruction and data for a single SPE. This local storage is the only memory directly addressable by the SPE.

The SPE architecture:

- Provides a large register file
- Simplifies code generation
- Reduces size and power consumption
- Simplifies decoding and dispatch

These goals are achieved by the novel SIMD-based architecture with 32-bit wide instructions encoding a 3-operand instruction format. By designing a new instruction set architecture (ISA), the developers have streamlined the instruction side, and provided 7-bit register operand specifiers to directly address 128 registers from all instructions, using a single pervasive SIMD computation approach for both scalar and vector data. In this approach, a unified 128 entry 128-bit SIMD register file provides scalar, condition, and address operands, such as for conditional operations, branches, and memory accesses.

While the SPE is a novel architecture, the operations selected for the SPE are closely aligned with the functionality of the PPE. This facilitates and simplifies code portability between the PPE and the SIMD-based SPEs.

Floating-point data types automatically support a wide dynamic data range and saturation, so no additional data conditioning is required. To reduce area and power requirements, floating-point arithmetic is restricted to the most common and useful modes. As a result, denormalized numbers are automatically flushed to 0 when presented as input

and when a denormalized result is generated. A single rounding mode is also supported. Single-precision floating-point computation is geared for throughput of media and 3-dimensional graphics objects.

Memory access is performed via a DMA-based interface called the memory flow controller (MFC) using copy-in/copy-out semantics, and data transfers can be initiated by either the PPE or an SPE. The DMA-based interface uses the Power Architecture page protection model, giving a consistent interface to the system storage map for all processor structures despite its heterogeneous instruction-set architecture. The high-performance on-chip EIB interconnects the SPE and Power Architecture computing elements.

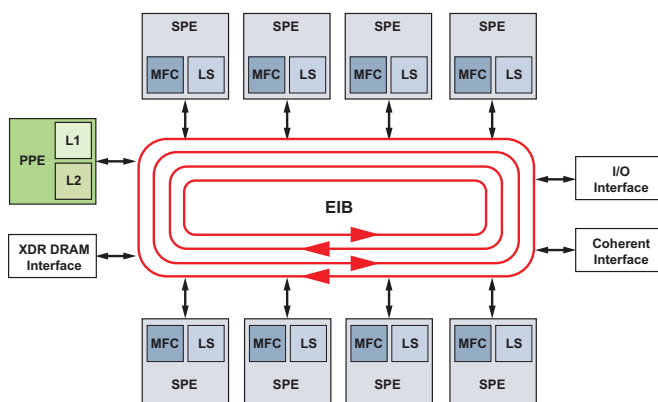
The SPE is an in-order, dual-issue, statically scheduled architecture. Two SIMD instructions can be issued per cycle: one compute instruction and one memory operation. The SPE branch architecture does not include dynamic branch prediction, but relies on compiler-generated branch prediction using prepare-to-branch instructions to redirect instruction prefetches to branch targets.

The SPE was designed with a compiled code focus from the beginning. The early availability of SIMD-optimized compilers has enabled the development of high-performance graphics and media libraries for the broadband architecture entirely in the C programming language.

The number of attached SPEs can be varied to achieve different power/performance and price/performance points, and the Cell architecture is a modular, extendable system, so that multiple Cell BE processors, each with a PPE and attached SPEs, can form a symmetric multiprocessor system.

Summary

Cell technology from IBM represents an enormous, long-term investment and a quantum leap forward in performance, originally focused on the commercial entertainment industry. Mercury is preparing to deploy this technology in the defense market, enabling the next generation in C4ISR processing. The capabilities of applications such as sensor image formation, situational awareness, target recognition and tracking, change detection, and data fusion can be fully realized. This new technology can provide unparalleled performance for rapid processing and dissemination of sensor data.



Cell BE Processor Architecture

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. Challenges Drive Innovation is a trademark of Mercury Computer Systems, Inc. Other products mentioned may be trademarks or registered trademarks of their respective holders. Mercury believes this information is accurate as of its publication date and is not responsible for any inadvertent errors. The information contained herein is subject to change without notice.

This document, including layout and graphics, is a copyrighted work and is protected by worldwide copyright laws and treaties. You may not reproduce this work, prepare derivative works based on this work, distribute copies of this work to the public by sale or other transfer of ownership, or by rental, lease or lending, perform or display this work publicly, or otherwise commercially exploit this work.

Copyright © 2005 Mercury Computer Systems, Inc.

For more information, go to www.mc.com

413-1205-c4isrcellwp



Corporate Headquarters

199 Riverneck Road
Chelmsford, MA 01824-2820 USA
+1 (978) 967-1401 • +1 (866) 627-6951
Fax +1 (978) 256-3599
www.mc.com

Worldwide Locations

Mercury Computer Systems has R&D, support and sales locations in France, Germany, Japan, the United Kingdom and the United States.

For office locations and contact information, please call the corporate headquarters or visit our Web site at www.mc.com.