

Cell Broadband Engine Technology

Bringing impressive performance improvements to systems running compute-, data- and memory-intensive applications



Executive summary

Companies around the world have researchers, engineers, scientists and business professionals who must wait for the results of computerized data analysis before they can take action. This *idle time* directly impacts business in various ways, such as slowing the introduction of new products, inhibiting research breakthroughs and delaying decisions that are critical to the company's long-term success. The obvious impact of these processing delays is cost wasted in people hours, squelched innovation and postponement of key decisions, which can lead to missed opportunities.

IBM®, in collaboration with Sony Computer Entertainment and Toshiba Corporation, established a relationship with the primary objective of developing a new processor with dramatically increased performance, responsiveness and security. Their aggressive target was to far outpace Moore's Law¹ by improving the processing performance by 100 times. The group understood that conventional microprocessors have performance limitations and traditional improvements were not going to meet current and future demands for greater processor performance. Another equally important goal was to overcome the increasing power consumption, cooling resources and floor space needed by systems based on conventional microprocessors that were trying to meet the ever-increasing processing demands of companies.

The results of the relationship is Cell Broadband Engine™ (Cell/B.E.™) technology, an advanced IBM Power Architecture™ technology-based microprocessor. It helps companies to overcome the limitations of today's microprocessors and targets current and future compute-, data- and memory-intensive application workloads. Currently, Cell/B.E. technology is available from IBM on select IBM BladeCenter® products and from IBM Business Partners in industry-standard form factors.

This paper gives examples of how various industries can use systems that are based on this new microprocessor to aid a company's knowledge workers² by providing vital information in a timely manner. The paper also explores how systems that are based on Cell/B.E. technology can overcome the increasing power consumption and cooling resources required by conventional microprocessors.

Addressing processing needs

Knowledge workers are being asked to interpret increasing volumes of data and information in real time, in order to make decisions and take action.

How can you help your knowledge workers perform more effectively?

You can start by providing them with solutions that run on systems that deliver important information and results in real time. By reducing processing time, you can enable knowledge workers to make decisions at the point of need and take action with greater assurance that their decisions are correct and based on the latest information available to them.

How does a Cell/B.E. technology-based system help industries perform at their full potential?

At the heart of the system is Cell Broadband Engine technology, an advanced IBM Power Architecture based microprocessor. Cell/B.E. is designed to provide power-efficient and cost-effective, high-performance processing for compute-, data- and memory-intensive

¹ *Moore's Law* is the "empirical observation made in 1965 that the number of transistors on an integrated circuit for minimum component cost doubles every 24 months." Source: Wikipedia

² A *knowledge worker* is someone who works with information or develops and uses it in the workplace.

workloads. In many instances, systems based on Cell Broadband Engine technology have already shown that they deliver vastly improved real-time responses for these workloads and in particular for broadband media applications. These intensive workloads span a broad range of both commercial and scientific fields and apply to these workload categories including:

► Real-time analytics

In order to be agile, organizations must effectively measure, evaluate and extract meaning from changes in data and information in real time. *Real-time analytics* are processes (often guided by policies, or rules, and procedures) that analyze, draw conclusions and act upon new or updated data and information as soon as an event occurs. Effective use of real-time analytics helps businesses to make decisions more quickly, keep pace with market developments and improve their bottom line.

► Media and graphics creation and management

Media, in this paper, refers to the form and technology used to communicate information, to entertain and to inform the viewer. A prime example is multimedia, which is composed of multiple forms of information, such as text, audio, graphics, animation and video. Another example is digital media, which refers to electronic media based on digital codes.

Graphics is an important way to represent and summarize information and to communicate a message. Computers are being pressed to create, display and manipulate pictures and images. These pictures and images are being used on an increasing scale to represent information to knowledge workers in various fields.

Visualization, a key aspect of this category, is applicable in science, engineering and product visualization and has proven valuable to multimedia and medical environments. For example, scientific visualization uses computer graphics to present huge quantities of laboratory or simulation data to knowledge workers in order for them to extract meaning and draw conclusions. Also, product visualization is used to manufacture components and product assemblies.

Imaging, another aspect of this category, involves the capturing, storing, registering, handling and fetching of images. Two commonly used forms of imaging are document imaging, which involves the management of images of text documents, and graphical imaging, which involves the creation and manipulation of graphics.

High performance compute-, data- and memory-intensive systems are required to handle the sophisticated software, tools and technology used to create, present, manage and maintain various forms of media and graphics.

► Extensive or varied types of data manipulation

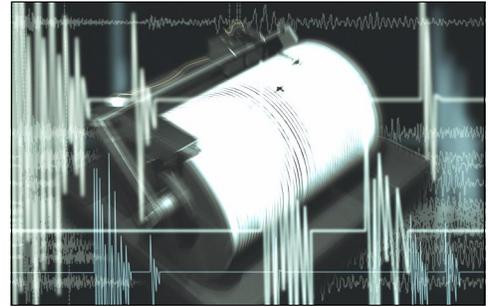
Companies are trying to handle data in numerous formats including unstructured data, which is not organized in a clearly defined and consistent manner. Processing can include:

- *Audio signal processing*: The processing of auditory signals or sound representations
- *Compression and decompression of data*: The encoding of information into a form that uses less space than its original form, making it easier to transport or store in limited space, ready for decoding back to a usable form when needed
- *Encryption and decryption of data*: The process of transforming (computer readable) information or data into an unreadable form, such that the encrypted data can be made readable (decrypted) only by using specific knowledge of the encrypted data
- *Secure transformation of data*: The process of taking data in one format and converting it into another format, while complying with specific security procedures

All of these workloads can be compute-, data- or memory-intensive and may effectively be used with Cell/B.E. technology-based systems. We now look at some examples:

► **Seismic data**

With increasing energy demands, key players in the chemical and petroleum industry are looking for ways to quickly and accurately process massive amounts of data that is essential to identifying likely locations for natural resources. They know that well-used seismic data is invaluable to this quest. These companies are looking for ways to maximize their compute power at a reasonable cost. They want solutions that provide more accurate seismic data processing faster, enabling their engineers and scientists to identify highly likely sites for petroleum reserves helping to maximize the return on investment (ROI).



Cell/B.E. technology-based systems are a good fit for these workloads and for processing massive amounts of seismic data. Tests of applications that manage and display seismic data running on a Cell/B.E. technology-based system have shown an improvement of up to 20 times in performance. On a Cell/B.E. technology-based system, these applications run faster and can process the volumes of data that are involved. In doing so, the applications make results available to scientists and engineers sooner so they may be more effective at drawing conclusions with greater accuracy, reducing the cost of drilling and increasing the potential for profits.

► **Financial information**

Financial firms are providing more robust products and analytics, demanding increasingly complex compute-intensive solutions. To effectively provide information to their clients, financial firms must accurately analyze massive amounts of information, take action quickly and understand and manage any associated risk. To stay at the front of the pack, financial experts must make trading decisions faster, with increased levels of accuracy. For example, financial experts need systems that can reduce the time it takes to value or calculate risk for portfolios that facilitate same-day decision making.



A Cell/B.E. technology-based system can help address these challenges by accelerating key workloads, optimizing the processing of massive amounts of data and returning results to the financial experts at the point of need. These systems provide value by increasing the speed at which volumes of raw data are turned into actionable information. For example, firms that do enterprise risk management and portfolio analysis are benefitting today from using Cell/B.E. technology-based systems.

► **Information-based medicine**

Advanced, computer-based technologies are helping doctors and clinicians provide the quality of care their patients deserve. Medical imaging is an area that is progressing rapidly and creating a new more demanding workload with increasing complexity. New medical imaging capabilities, such as three-dimensional (3-D) imaging, real-time imaging and analysis, and knowledge extraction and simulation, are

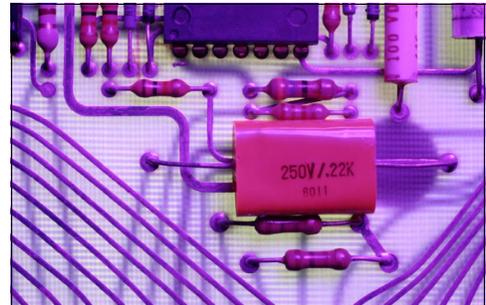


providing critical patient and research information. They are also demanding higher levels of processing performance. Even current image processing is straining existing systems and creating long processing delays.

Through the development of parallel computer architecture and major memory bandwidth improvements, Cell/B.E. technology-based systems are dramatically speeding up the processing of medical images. For example, these processing advances significantly aid *image registration*, which is the computer-enhanced alignment of two medical images obtained at two different dates. By aligning these images over each other, a radiologist is able to more easily detect structural changes such as the growth or shrinkage of tumors. Medical imaging is helping medical professionals to make timely, critical decisions that can save lives.

► **Electronic design automation**

Today, industries, such as aerospace and defense, automotive, and manufacturing, and companies that provide design services in semiconductor development are using electronic design automation (EDA) to help design semiconductor technology for increased manufacturability. In particular, chip manufacturers use EDA to design chips. Currently, EDA is costly and can impact time-to-market because many conventional computers must be harnessed together to perform the necessary computations.



Optical Proximity Correction (OPC), OPC verification and lithographic simulation are compute-intensive EDA functions and are critical to this process. These processes are particularly well-suited for Cell/B.E. technology-based systems. Initial indications are that these EDA systems could provide performance improvements when implemented by an EDA solution provider. In addition, users may more easily identify and eliminate lithography hot spots. This could optimize their designs thus improving manufacturing yield while retaining design and electrical intent.

Systems based on Cell Broadband Engine technology also are a good fit for a variety of signal processing applications. These include applications that process sound, images, biological signals (such as electrocardiograms) and radar signals. Processing of such signals includes storage, reconstruction, separation of information from noise, compression and feature extraction.

Processing limitations with conventional microprocessors

Conventional microprocessors with their single-thread performance are often not adequate for today's new, compute-intensive workloads. Key limitations are related to:

- Performance
- Power consumption

Performance

Performance of conventional microprocessors is limited by memory latency and bandwidth.

Memory latency is the time it takes to retrieve a requested piece of data or code for use by a microprocessor. Processors are over 1000 times faster than in the 1980s, but memory speed has only slightly improved, leading to significant latency issues. In most cases, application performance is limited by memory latency more than any other reason. Today, main memory

operates at about one-third the speed of the processor. Therefore, for up to two-thirds of the total time, the processor could be waiting on memory for data or code to process if the data or code was not pre-fetched by the cache.

Bandwidth is the amount of data that can travel a communication path in a specific amount of time. If an executing application needs a large amount of data, the processor and the application it is executing can sit idle because they are waiting for the delivery of all the data needed to accomplish the task at hand.

Performance of conventional microprocessors is limited by memory latency, bandwidth and the number of microprocessor cores that are available to perform work.

Power consumption

The power consumed by microprocessors is increasing as more transistors are placed in the same space in an effort to gain more compute power. However, placing more transistors in a given area causes an increase in the power and cooling that is needed to operate the microprocessor. Companies are becoming more sensitive to the increasing cost of energy used by their computer centers. Using more power and cooling is causing energy bills to take a larger piece of the IT budget.

Cell Broadband Engine design basics

Cell Broadband Engine technology provides a significant performance advantage over conventional microprocessors for those applications that take advantage of parallelism and software controlled memory management. Parallelism provides an environment where independent subtasks (for a given task) are executed concurrently, dramatically increasing throughput. Software controlled memory management enables the developer to enable the application to make more efficient use of memory.

At the heart of Cell Broadband Engine architecture are two different types of processors:

- ▶ A single Power Processing Element (PPE)
- ▶ Eight optimized and specialized co-processors called *Synergistic Processing Elements* (SPEs)

Conventional processors have one to several processor elements (cores) and no SPEs. SPEs are specialized so they are smaller than conventional cores, allowing the Cell Broadband Engine processor to fit more SPEs into a smaller area, which saves power while boosting performance. The dual thread, 64-bit PPE is the processor that has primary processing control and runs the operating system, manages system resources and manages work shipped to the SPEs. The SPEs are then dedicated only to the compute task that needs to be performed on the data. By separating these activities, the Cell Broadband Engine architecture ensures the SPEs are well utilized and work is spread across the SPEs, allowing an application structured for parallelism to execute much more efficiently. This architecture also allows the system to run a broader range of applications than many types of specialized processors.

The eight SPEs are computational workhorses. They are completely focused on processing data or code as quickly as possible, making them well suited for Single Instruction Multiple Data (SIMD)-coded applications and for applications with numerous floating point calculations, which are usually problematic for conventional microprocessors.

The SPEs have features that make them distinctive, including:

- ▶ *User-mode architecture*, which means the application has complete control over what the processor is doing

No cycles or local memory are stolen by the operating system or daemon tasks. This means that programmers can make more efficient use of each SPE that is under their complete control, rather than under shared control, which is the case with a traditional microprocessor.

The processor core is designed with a reduced instruction set organization that is able to execute on 8-, 16-, 32- or 64-bit-width data as well as selected 128-bit-width data. This gives the processor core the ability to process, in a pipelined fashion, a broad set of data and minimize idle time for any of the processors.

- ▶ *Three-level storage organization* consisting of the register file, Local Store and main storage

This is a radical departure from the conventional microprocessor architecture and programming models. The Cell/B.E. processor uses parallelism separating computations into smaller units executing in parallel. All three of these levels of storage are under the complete control of the programmer allowing them to best optimize this storage use for the task at hand.

Local Store consists of 256 kilobytes of dedicated, on-chip private memory, making memory access extremely fast compared to the main memory access done by conventional microprocessors.

This overall approach to storage organization reduces unnecessary data transfers and miss-fetches found in conventional processors.

- ▶ *A flexible Direct Memory Access (DMA) engine* effectively improving memory bandwidth, latency and utilization of moved data

In essence, the DMA engine is responsible for moving data and instructions between the Local Store of an SPE, other SPE Local Stores, main memory and input/output (I/O) systems. Because it is separate and asynchronous to the SPE, this movement of data happens while the SPE is doing other processing, helping to more effectively use each SPE.

- ▶ *A 128-bit register file stores all data types* such as, integers, single-precision and double-precision floating-point, scalars and vectors

The large size of the file and its unified architecture achieve high performance without the use of expensive hardware techniques. By incorporating a large register file, the number of interactions with memory is reduced, making access to more data locations faster.

Programming for Cell/B.E. processors

In order to take full advantage and reap the rewards of the Cell Broadband Engine architecture some code modifications may be required. To support the programmer who is doing these modifications, IBM has developed and is enhancing a Software Development Kit (SDK) to help them. Systems based on multicore microprocessors are here to stay. Those companies that are willing to modify their applications to take full advantage of the Cell Broadband Engine architecture will benefit from these high performance processors and leap ahead of their competitors.

Applications that are modified to take full advantage of this architecture have experienced dramatic processing performance improvements over conventional processors when performing single-precision, floating-point operations.

The SDK, based on Eclipse and developed to run on Linux®, enables your programmers to develop or modify applications to effectively use Cell Broadband Engine technology. This SDK is also for a robust software development community that engages in the development of applications that are optimized to Cell/B.E. technology. Tools and solutions that are created by IBM Business Partners ease implementation of applications. Currently, the SDK includes:

- ▶ IBM XL C/C++ compiler
- ▶ IBM Full System Simulator
- ▶ Sysroot Image for System Simulator
- ▶ Support for Hybrid Model programming wizard

Summary

Today's IT workloads are increasing and stressing the conventional systems that are now in place. Power, cooling and space constraints are increasingly expensive for companies that are trying to handle these workloads and the mountains of associated data.

Systems based on Cell/B.E. technology may be the answer for companies that want their knowledge workers to make timely decisions based on the dynamic and often large volumes of data. Currently, Cell/B.E. technology is available from IBM on select IBM BladeCenter products and from IBM Business Partners in industry standard form factors. To take full advantage of such systems, applications need to be structured appropriately for parallelism, and tools are available for that task.

Cell/B.E. technology-based systems are helping companies deploy power-efficient, cost-effective, high-performance systems that deliver dramatically improved real-time responses for compute, memory and data-intensive workloads.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.



Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BladeCenter®

IBM®

Power Architecture™

The following terms are trademarks of other companies:

Cell Broadband Engine and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems Inc. in the United States, other countries, or both.

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