

Packaging the Cell Broadband Engine Microprocessor for Supercomputer Applications

P. Harvey¹, R. Mandrekar¹, Y. Zhou¹, J. Zheng¹, J. Maloney¹, S. Cain¹, K. Kawasaki², G. Lafontant¹, H. Noma², K. Imming¹, T. Plachy¹, D. Questad¹

1. IBM Systems & Technology Group, USA . 2. IBM Global Engineering Solutions, Japan
pmharvey@us.ibm.com

Abstract

The Cell Broadband Engine™ (Cell BE) processor initially designed for high-end consumer electronics, has been enhanced by IBM for supercomputer applications. The enhancements to the chip also necessitated the design and development of a new package. The modifications to the chip included replacement of the 3.2 Gb/s XDR interface with a 800Mb/s DDR2 interface of equal bandwidth. This required the addition of several hundred chip-level connections (C4's) and package BGA balls. Incorporating this and other enhancements to the chip resulted in a ~20% larger chip and a larger and more complex package. Additional noise from this large memory interface also drove decoupling requirements that necessitated mounting capacitors on both the top and bottom sides of the package.

This paper describes the design of this new package as well as the analysis and characterization techniques used to address the packaging concerns outlined above. It includes a comprehensive noise analysis as well as a thorough characterization of the DDR2 interface in the final prototypes. The paper also outlines the design and analysis of the power distribution to the various voltage domains on the chip. Along with electrical design and performance, the paper also includes finite element modeling of the mechanical stresses resident in this FCPBGA package. Finally, the concluding portions of the paper will discuss the trade-offs between electrical performance and mechanical stability, reliability and relative cost.

Introduction

The initial Cell BE processor was designed for high end consumer electronics applications and is now shipping in high volume [1,2]. But the Cell architecture was intended for much more than consumer electronics. Shortly after completion of the original Cell BE processor, IBM and Los Alamos National Labs announced plans to use this chip to build a supercomputer with a sustained performance of over 1 petaflop, quadruple the processing performance of today's most powerful supercomputer [3,4,5]. As part of this project, several enhancements to the original Cell BE were planned to augment the processor's capabilities in a supercomputer application environment. The new chip was dubbed the PowerXCell™ 8i (PowerXCell).

The PowerXCell includes an enhanced Synergistic Process Element (SPE) with a five-fold increase in double precision floating point performance [6] to compliment the industry-leading single precision floating point performance. The 3.2 Gb/s XDR memory interface was replaced with an

800 Mb/s DDR2 interface which extended the memory access limit from 2GB to 16GB and enabled the use of industry standard DDR2 DIMM components, an essential building block in the supercomputer system. The DDR2 memory interface required significant area on the chip and was wrapped around one end of the chip. Previously mentioned enhancements to augment the double precision performance were also added to each of the 8 instantiations of the SPE resulting in a new layout for the larger chip that required minimal re-integration. The layout of the enhanced Cell chip is illustrated in Figure 1 below.




Figure 1: PowerXCell floorplan with the DDR2 memory interface and the enhanced double precision (eDP) performance in the SPE's. The chip is shown with the C4 interconnects face up. The chip is flipped over for bonding to the package substrate.

Changes in the processor chip described above as well as the new system environment drove several key changes to the original Cell BE package design. Including test and PLL connections, the DDR2 memory interface required nearly 600 addressable C4's to be wired out on the package. These signals were routed to the available BGA locations on the north, south and east sides of a 47.5mm FPBGA package. A 4-2-4 thin core package construction was necessary to accommodate all the necessary signal wiring. The original Cell BE package maintained a low noise environment for the core elements of the chip with capacitors mounted on the bottom-side of the package directly under the chip [7,8]. It was deemed important to maintain the same low noise environment for these elements in the PowerXCell chip. Consequently, the same bottom-side capacitors were used in the new package; however, the simultaneous switching noise (SSN) from the newly added memory interface drove a need for more decoupling capacitors than could be placed on the bottom-side of the package. As a result, capacitors were placed on both sides of the package laminate.




Figure 2: Cross section drawing of PowerXCell Package (not to scale)

The increase in chip and package size, the laminate construction, and placement of capacitors on both the top and bottom side of the laminate prompted concerns regarding the thermo-mechanical stability and reliability of the fully assembled package. These aspects were studied with finite element mechanical stress modeling and the reliability was confirmed by testing.

DDR2 Memory Interface

The DDR2 interface on the PowerXcell chip consists of two 128-bit channels (144-bit with optional ECC), operating at 800Mb/s using the 1.8V Stub-series terminated logic (SSTL18) technology. Including address and control signals, each channel consists of 192 single ended I/Os and 38 differential I/Os, the DDR2 interface has a total of 540 I/O buffer cells in under 30mm² of chip area. While the DC specification for the supply is +/- 100mV, it was not clear how much AC noise from simultaneous switching (SSN) could actually be tolerated on the buffers of this large interface. A principle objective of the package design process was to determine an acceptable level of SSN and develop a complementary package design that meets this requirement in the context of the existing chip characteristics and proposed system environment.

The SSN on the DDR2 interface was determined by a combination of the current transients on the I/O devices and characteristics of the power distribution network (PDN). To model both of these phenomena accurately, a comprehensive system model that accounted for the signal path from the driver to the receiver, as shown in Figure 3, was used for simulation. Each of the individual components were modeled separately and put together in a SPICE environment for simulation.




Figure 3: Comprehensive system model from driver to receiver used for SSN simulation.

The SSN on the interface was found to be a function of the on-chip decoupling in each I/O cell as well as the package decoupling. Given the space constraints on the bottom side of the package, I/O decoupling was placed on the bottom and the top sides of the package (Fig. 2). The bottom side decoupling capacitors were right underneath the chip shadow of the DDR2 interface. The top side capacitors were placed on the north and south sides of the interface. Table 1 lists the number of capacitors placed at each location along with the relative effectiveness of each capacitor as estimated by the effective loop inductance of capacitors at that location. As expected, the bottom side capacitors which offer the shortest path to the chip provide the best performance. The dependence of SSN on the on-chip decoupling was studied by simulating SSN on the interface for a range of on-chip capacitance values. The results are shown in the graph in Figure 4. It was seen that after about 40pF of decoupling per I/O cell, additional on-chip decoupling yielded diminishing improvement.

Location	No. of decaps	Path Inductance	Loop Inductance (Path Ind + ESL)	Loop Ind/Decap
Top-layer North	10	25.84 pH	31.05 pH	310 pH
Top-layer South	14	24.94 pH	28.66 pH	401 pH
Bottom layer	6	14.85 pH	23.53 pH	141 pH

Table 1: Loop inductance between chip and on-module capacitors in located in 3 principle locations in the package.




Figure 4: Relative simultaneous switching noise as a function of capacitance placed in each I/O cell.

The impedance profile of the DDR2 power supply network was generated by exciting the system model shown in Figure 3 using a impulse switch event. Figure 5 shows the impedance profile obtained from both simulation and measurement of prototype parts. Both measured and simulated impedance profiles show a resonant peak at 230-232 MHz. To simulate worst case SSN, a random bit pattern with sufficient harmonic content in the resonant frequency region was used to excite the interface. Figure 6 shows noise with all data lines switching with the same pattern, the peak-to-peak SSN was observed to be 236mV. This simulation assumed worst case component corners in the package as well as worst case silicon corner on the chip. The eye-degradation observed on the interface because of this SSN was found to be acceptable to meet the overall timing requirements on the interface.




Figure 5: Measured and simulated impedance profile (impedance vs frequency) for Vdd18 power distribution in the PowerXCell package.




Figure 6: Simulated SSN noise (noise vs time) from worst case stimulus and worst case package and chip simulation corners. Maximum noise was found to be 236mV. This level of noise did not cause timing errors in the DDR2 interface.

The SSN simulations performed on the DDR2 interface were validated in hardware through a series of measurements. The impedance profile was obtained by measuring the response of the PDN to a single impulse data switch. Figure 5 shows the impedance profile of the DDR2 interface measured for 6 different prototype parts. The resonant peak seen at 232MHz matches extremely well with the simulated impedance profile. The SSN on the interface was measured by exciting the interface using a number of different data patterns. These included the random data pattern that gave worst case SSN in simulation as well as some repetitive data patterns that excited the resonant region of the impedance profile. The measured and simulated values of SSN were in good agreement as can be seen in Table 2. The lab setup available at the time of the measurements only allowed the simultaneous switching of all data lines within a single channel.

	Peak-to-peak SSN (mV)				
	Patterns				
	Impulse	1010	110	1100	Random
Average	74.83	30.9	67.81	54.88	88.71
Simulated	86	24	65	41	113

Table 2: Measured average SSN noise versus Simulated SSN for a variety of bit patterns used as a stimulus.

Core and Array Power Distribution

The power distribution in the original Cell BE processor has been described in earlier work [9]. In order to ensure a very similar noise environment to the original Cell BE chip and avoid any re-engineering of the core components of the chip, the power distribution strategy for the new package was planned to be very similar to the original Cell BE package. Backside decoupling was used to decouple the core power as implemented previously [7]. Enhancements to the PowerXCell chip increased the AC and DC power

significantly. AC noise simulations are summarized in Figure 7 below. A low frequency noise stimulus, an adjustable mid frequency range stimulus, and a composite stimulus with both mid and low frequency stimulus were used to study the package and board response to noise. Results indicated that while top-side capacitors were not as effective as the initial capacitors placed on the bottom-side of the package, these topside capacitors were necessary to meet established noise specifications with sufficient margin.




Figure 7: Summary of simulated AC noise with various stimuli as a function of package capacitor configuration. Noise is normalized and reported as the percent of the total noise specification (lower is better).

Mechanical Analysis and Reliability

The preceding electrical analysis demonstrates the importance of the on-module capacitors in mitigating both core and I/O switching noise. As previously discussed, capacitors would need to be placed on both sides of the laminate to meet all the decoupling requirements for this package. Therefore, it was very important to carefully consider the mechanical issues associated with this capacitor placement scheme. A finite element analysis of the warpage of the fully assembled organic package highlighted the top side regions immediately adjacent to the chip corners as the locations likely to have the greatest warpage. The general results are illustrated in Figure 8. As a result of this study, capacitors were not placed on the top side of the laminate at these locations. Interestingly, one result of placement of capacitors on the south side of the package was that it increased the rigidity and reduced the warpage.




Figure 8: General warpage of fully assembled package at low temperature. The package lid is removed to clearly illustrate the warpage.

Bottom-side and top-side capacitor placement on organic laminates have been extensively analyzed and tested in

previous packages, but they have never been employed together in the same module at IBM. Stress on the individual capacitors was a concern. A finite model was used to determine the maximum principal stress on each capacitor placed on the package and the results are illustrated in Figure 9.

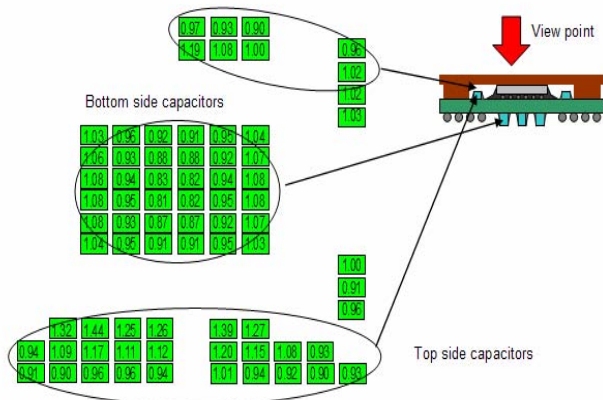


Figure 9: Maximum principal stress in ratios on top side and bottom side capacitors in the package.

Bridging to extensive qualification data from previous evaluations, the stress modeling predicted that capacitors would have no more stress than in previous packaging configurations. Thermal cycling followed by functional verification of the capacitors confirmed the result in this package configuration.

Other mechanical risks for this package were also evaluated against the Cell BE package. As shown in Table 3, the new PowerXCell package has higher mechanical risk than Cell BE package but smaller risk compared to previously qualified IBM packages used in other applications. The overall mechanical risk therefore was considered to be small.

Mechanical risk	Original Cell BE package	Previously qualified IBM package	New PowerXCell package
2nd level interconnect	1.00	1.09X	1.08X
TIM separation/cracking	1.00	1.17X	1.29X
UF corner delamination	1.00	1.63X	1.32X
Laminate resin cracking	1.00	1.16X	1.00X

Table 3: Relative mechanical stress based on finite element modeling for various elements of the package compared to original Cell BE and other previously qualified packages.

In addition to testing the reliability of the capacitors it was also important to verify the reliability of the C4 interconnects between the chip and package and the internal via structure and wiring in the laminate. The PowerXCell chip is currently the largest chip in IBM using the 150um pitch C4s. With a larger DNP (distance from neutral point) than previous 150um pitch C4 chips, the C4s are subjected to increased stress during Thermal Cycling.

Prior to the start of the package qualification some risk mitigation stresses were run with the product laminate and

open (C4s are not connected to each other on the chip) and shorted (all of the C4s are shorted together) chips. Testing with shorted chips facilitates measurement of resistance shifts in the C4 interconnects and in the internal vias and wiring in the laminate during thermal cycle (TC) stress. Testing with the open chip enables leakage measurements on the capacitors, C4 interconnects, and package wiring during temperature and humidity bias (THB) testing. After 2000 cycles of TC (-40C/125C) there were no fails. After 1000 hrs of THB (85C/85%RH/3.6V) there were no failures in the C4's or capacitors. Additional stressing is ongoing and there are no C4 interconnect fails or capacitor solder joint fails.

Conclusion

This paper outlines the design, analysis and validation of an FC-PBGA package for the PowerXCell chip, a second generation Cell BE processor specifically targeted for high performance computing. The large memory interface required a more complex laminate to wire all the addressable signals and switching noise from the memory interface and the core necessitated placing capacitors on both sides of the laminate. An extensive noise analysis confirmed the decoupling requirements. Finally, mechanical analysis guided placement of the capacitors away from high stress locations on the laminate and reliability testing confirmed that all the new design elements in the package were reliable.

References

- Pham, et al, "Overview of the architecture, circuit design, and physical implementation of a first-generation cell processor", *IEEE Journal of Solid State Circuits*. Vol 41, no 1. Jan 2006. pp 179-196.
- Kahle, et al, "Introduction to the Cell Microprocessor", *IBM Journal of Research and Development*. Vol 49 no 4, Sept. 2005. pp 589 – 604.
- "Round Runner: Racing Toward the Home Stretch", Los Alamos National Labs News Bulletin, Dec 10, 2007.
- Los Alamos National Lab High Performance Computing Website, Road Runner Project Overview. Feb 21, 2008 <<http://www.lanl.gov/orgs/hpc/roadrunner/index.shtml>>.
- "IBM to Build First Cell Broadband Engine Based Supercomputer", IBM Press Release, 06 Sept. 2006.
- Flachs, et al, "A streaming processing unit for a CELL processor", *IEEE International Solid-State Circuits Conference*, v 48, 2005. pp. 100-101.
- Goto, et al. "Electrical Design Optimization and Characterization in Cell Broadband Engine Package", *Proceedings of the 56th Electronic Components and Technology Conference*, San Diego, CA. May, 2006. pp 194 – 202.
- Harvey, et al. "Chip/Package Design and Technology Trade-offs in the 65nm Cell Broadband Engine™" *Proceedings of the 57th Electronic Components and Technology Conference*, May, 2007. pp 27-34.
- Zhou, et al. "Distributed On-chip Power Supply Noise Characterization of the Cell Broadband Engine" *Proceedings of Electrical Performance of Electronic Packaging*, Oct. 29-31 Oct. 2007 Page(s):99 – 102.

Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc
PowerXCell is a trademark of the IBM Corporation.