# MRI and the Cell Processor

J. Czechowski[1], S. Zingelewicz[1], B. Moore[1], L. Marinelli[1], C. Hardy[1], and M. El-Demerdash[2]
[1]Global Research, General Electric Company, Niskayuna, NY, United States, [2]Healthcare, General Electric Company, Waukesha, WI, United States

**Introduction:** Computer chip-making capabilities have advanced to the point where multiple, heterogeneous compute engines can be placed onto a single chip. We examine the applicability of one type of these chips, the Sony/Toshiba/IBM Cell processor, to MRI. The Cell processor is a highly parallel, single precision floating point compute engine with a very high-speed on-chip data interconnect. Because of these features, we believe it is well suited to the sort of computational loads characteristic in several types of MR reconstruction algorithms. We describe the acceleration of three of these algorithms using the Cell over conventional computer processors such as the AMD Opteron.

**Methods:** Several relevant MR reconstruction techniques were implemented on a Cell. In particular, this included a Cartesian MR recon using a 2D IFFT, a homodyne recon, and ASSET recon. Our test datasets included 4 slices of 512x512 images from 8 coils. While the input data consisted of 16-bit integers, all calculations were performed using single precision floating point operations.

The algorithms were implemented and tested on an IBM Cell prototype blade consisting of two 2.4GHz Cell processors and 512MB of main memory. Each Cell consists of a general-purpose PowerPC core and eight Synergistic Processing Elements (SPEs) as shown in Figure 1. These SPEs are 128bit SIMD engines, which can process four 32-bit data elements in a single instruction. Each algorithm was implemented and compared against reference implementations on conventional PC processors.

Each SPE has limited memory (128KB) and no cache, so a double buffering scheme was used. The Cell architecture allows DMA operations to occur in parallel with computations, and its high-speed bus (IEIB) is sufficiently fast that (for this size image) the DMA operations are completely covered by the IFFT computations. Figure 2 illustrates the general methodology and partitioning used for these algorithms.

**Results and Discussion:** The timings in Table 1 were attained using only one (of two) Cell chips on an IBM Cell blade (version DD2). In general, the performance gain increases with the complexity of the computations. Sample results are shown in Figure 3.

Exploiting even a fraction of the theoretical potential of multiprocessor systems is non-trivial. As multiprocessor systems proliferate, software tools will evolve to make the programming task more manageable. In this particular case, the Cell implementation of these MR reconstruction algorithms was written in C using IBM provided libraries and compilers. One program was written to run on the PowerPC and one program was written to run on each of the SPEs. The latter was possible because of the symmetry of the processing needed for the MR algorithms. Each SPE simply focuses on processing a different stripe of the data during each phase of the computations.

Additional performance gains are possible and expected. In particular, the code generated in this study can be further optimized, and more Cell chips could be used, in parallel. Also, further Cell chip performance increases are anticipated. In conclusion, we can expect MRI system performance to benefit from multi-processor computer chips. This is likely to be needed as the number of coils expands, the size of images increases, and the usage of 3D datasets continues to grow.

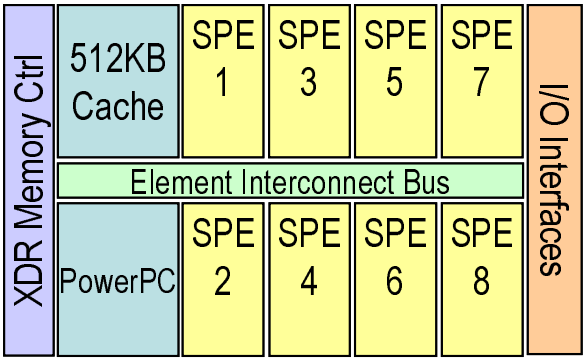


**Figure 1**: Block diagram of Cell processor

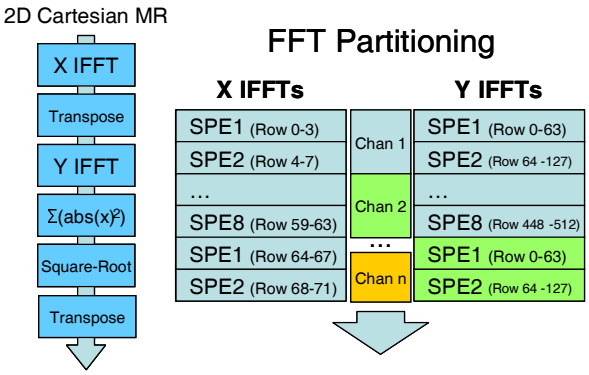


**Figure 2**: Processing steps and iFFT partitioning

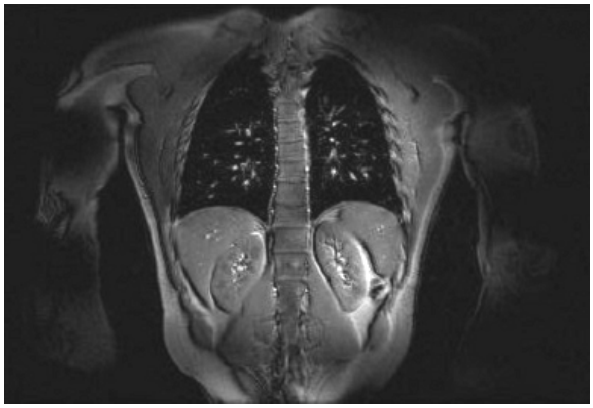

**Figure 3**: Cell reconstructed MR image

| | Opteron | Cell |
|---|---|---|
| Cartesian | **0.09** | **0.05** |
| Homodyne | **0.50** | **0.13** |
| ASSET | **2.74** | **0.18** |

**Table 1**: Performance comparison (seconds to process 4 slices of 512x512 images from 8 coils)