

Platform for high-performance imaging and visualization

by **Olivier Bockenbach**, Mercury Computer Systems
and **Prof. Dr. Marc Kachelriess**, Institute of Medical Physics, Erlangen

This article evaluates a CBE-based architecture in performing spiral CT and C-arm CT reconstructions. Image reconstruction basically consists of data weighting, convolution, and back-projection. The last step is the most computationally demanding, but fortunately all steps can be parallelised quite easily.

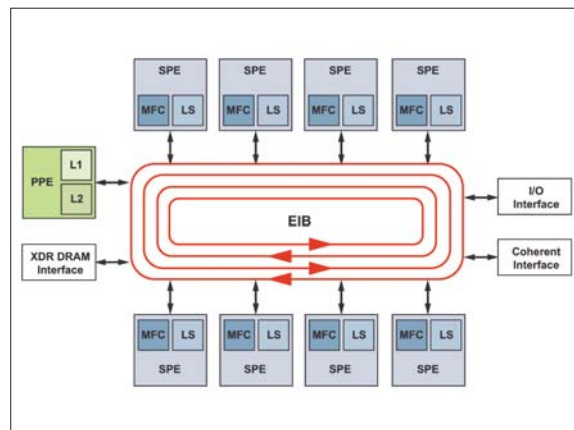


Figure 1. Cell processor

■ Medical imaging is an information processing technique that takes the data samples from medical devices such as MRI or CT scanners and translates those samples into 2D, 3D or even 4D images. The advances in the sensor technology allow for the generation of an increasing number of images per procedure and per patient, posing a tremendous challenge for the efficient, in-time processing and visualization of the resulting images.

CT is used in a growing number of clinical applications, and keeps challenging the scientific community to continuously propose new algorithms with improved image quality while reducing the X-ray dose. CT also challenges manufacturers to design computer systems that enable the object to be reconstructed within a timeframe that is compatible with the workflow in a hospital, while keeping the costs within reasonable bounds. The implied trade-offs have often led to the use of approximate algorithms such as advanced single slice rebinning-based approaches (ASSR-type image reconstruction) for spiral CT or Feldkamp-type algorithms for flat panel detector-based systems such as C-arm CTs or micro-CT. However, keeping the reconstruction times acceptable has also forced the design of special-purpose reconstruction platforms based on FPGAs or ASICs.

Originally developed for the gaming industry

by IBM, Toshiba and Sony, the industry-leading Cell microprocessor technology is going to be deployed for use in medical, defense, and commercial applications, mainly because the Cell technology offers significant performance improvement compared to other available technologies. Clocked at 3 GHz, the Cell processor offers a peak performance of ~200 Gflops and a memory bandwidth of 25 GB/s which is significantly higher than any off-the-shelf solution. To reach this level of performance, the architecture of the Cell articulates itself around a multiprocessor design, where the processing elements are connected together through a high-speed data bus, the EIB. The EIB is capable of peak data rates of ~200 GB/s (see figure 1).

Connected to this bus, the processing power is built from one PowerPC processing element (PPE) and eight synergistic processing elements (SPE). The EIB also connects to the memory controller, the I/O interface and the coherence interface. The memory controller and the I/O interfaces give all processing elements access to main memory and I/O space while the coherence interface allows for building multi-processor systems in symmetric multi processor mode, as depicted in figure 2.

The Cell architecture allows all kinds of distributed applications to run on the processing units, using elaborate data transfer techniques

to design any combination of parallel/pipelined approaches. Therefore, all applications that have the ability to subdivide the main tasks can be subdivided into a finite number of subtasks, and take advantage of the immense processing power implemented in the Cell processing elements, while keeping the data transfer latency very low. Indeed, the Cell Broadband Engine (CBE) processor proposes a new alternative to build a computer system capable of performing modern cone-beam spiral CT reconstruction without hosting dedicated and expensive devices. The level of parallelism along with the vast I/O capabilities permits the Cell processor to efficiently implement complex CT reconstruction algorithms with close to real-time performance. The CBE allows for designing systems where the radiologists can view images obtained from better algorithms, with higher quality much sooner than ever before. Critical decisions can be made more rapidly and more accurately.

This article evaluates the performances of a CBE-based architecture to perform spiral CT as well as C-arm CT reconstructions. Basically, image reconstruction consists of data weighting, convolution, and back-projection. The last step is the most dominating and computationally most demanding step during reconstruction. Fortunately, all image reconstruction steps can be parallelized quite easily. To assess

back-projection performance, an implementation of a parallel back-projection algorithm (2D) and of a perspective cone-beam back-projection algorithm (3D) was ported to the Cell processor. Several constraints had to be followed. The local store (LS) is limited to 256 KB and only small portions of the full problem can be handled by each worker. To accommodate demand, the image and raw data (sinograms) were tiled into sub-images and sub-sinograms. The size of the sub-images and the size of the sub-sinograms or sub-detectors were chosen to allow for double buffering of the sinogram data. Two sub-sinograms plus one sub-image plus code stack must fit into the 256 KB local store. Only those portions of a projection that were needed by a worker's particular sub-image comprise the sub-sinogram and were transferred by direct memory access (DMA) to the worker. While the worker is busy back-projecting the first sub-sinogram, the DMA of the other sub-sinogram was active. In this way, we were able to fully hide the DMA latency behind the back-projection process.

Furthermore, care was taken to make use of the 128 available registers per SPU to fully fill the execution pipelines. Manual loop unrolling and reordering of instructions ensured to achieve a throughput of more than one instruction per clock cycle. Vectorization was achieved by treating each quadruple of neighboring pixels as a 4-vector of floating point values; thereby the number of pixels per row must be a multiple of four. For image sizes that do not fulfil this criterion (e.g. 512²511 pixels) up to three fringe columns of dummy pixels are added prior to being passed to the workers.

With one CBE at 3.2 GHz running a parallel beam 2D CT parallel-beam back-projection for clinical CT, we are able to reconstruct 165 slices per second, from 512 projections with 5122 pixels per slice. Using one CBE for a flat panel detector-based system, such as C-arm CT or micro-CT, we can back-project a volume of 5123 voxels from 512 projections in 14 seconds. Using both CBEs available on a Dual Cell-based server one can double the performance. The time required for convolution and data reorganization prior to backprojection is in the order of 5% of the time required for back-projection and therefore negligible.

The performance achieved for spiral CT with a Cell processor-based architecture makes a major breakthrough with off-the-shelf hardware. Cell-based systems bring the processing power needed for affordable high-performance medical imaging systems, capable of living up to the requirements of modern detectors such as those designed for CT Gantries and of delivering accurate data to the radiologist to make accurate diagnostics. ■

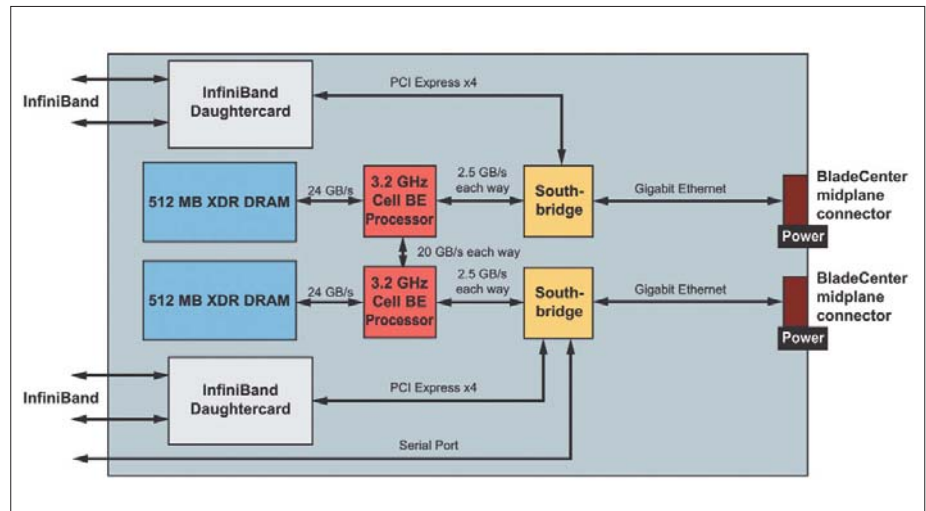


Figure 2. Dual Cell-based blade