

A high performance platform for real-time X-ray imaging.

X-ray tomography has been proven to be a valuable tool for understanding internal, otherwise invisible, mechanisms in biology and other fields. Unfortunately it requires computationally intensive reconstruction algorithms with long post-processing times and was therefore mostly limited to static objects. To enable investigation of technological and biological processes in 3D and with a temporal resolution down to the milliseconds range, a high performance imaging station is currently developed at KIT. The setup includes high precision mechanics, an ultrafast detector system, and a computing platform for a real-time data processing. At a second stage image-based feedback loops are foreseen to control both the technological process under study and the configuration of the imaging station. The key component of the detector system is a new type of high-performance camera developed at our institute. The embedded Xilinx FPGA integrates programmable logic for trigger, compression, and fast control algorithms. The camera prototype is connected to the imaging station using PCI express (gen2, x4) bus and able to continuously stream 2 Mpix image at 300 fps with a dynamic range of 10 bits generating a data stream of approximately 750 MB/s. The final specification aims to achieve rates of up to 50,000 fps with a readout bandwidth of 4GB/sec, a five fold increase compared to the fastest commercially available cameras (see Fig. 1).

To process such amount of data in real-time, we have parallelized image processing software employed at ANKA [1]. Our pipelined architecture utilizes all available resources. For instance, while performing tomographic reconstruction CPUs are pre-processing images; GPUs are reconstructing a 3D volume from 2D projections interleaving memory transfers and computations; results are streamed to the storage subsystem. To fully utilize GPU potential, we have developed multiple reconstruction kernels adapted to recent GPU architectures. The GT200 kernel benefits from the texture engine to perform linear interpolation and to cache data accesses. Performance of the Fermi texture engine has not changed much while computation performance was boosted several times. Our Fermi kernel reduces the amount of texture fetches by the price of computing linear interpolations manually. However, this increased register usage significantly, hence reducing the occupancy. To hide latencies by overlapping of independent instructions, the Fermi kernel processes 4 pixels at once. These changes have increased reconstruction performance by 40% compared to our standard GT200 implementation. After thorough evaluation of existing parallel architectures, we have built a high performance imaging station based on GeForce line of NVIDIA graphic cards. The system is based upon the SuperMicro 7046GT-TRF equipped with 6 GTX580 adapters, 2 Xeon 5650 processors, and 96 GB of memory. Two GPUs are connected directly to the PCIe bus and four are installed into an external expansion box from One Stop Systems sharing a single PCIe x16 connection. By interleaving data transfers to the GPUs that share the bus, it was possible to limit the performance penalty to 3% as compared with GPUs connected directly. Efficiently utilizing resources of the described system, we have reached a processing throughput of about 500 MB/s (only 5 MB/s was achieved with a Xeon server). The storage subsystem consists of 16 SATA hard drives organized as RAID6 in the external Areca enclosure and is connected to the server using SAS interface. Since the sequential bandwidth of magnetic hard drives is significantly degrading with the offset from the beginning of disk, we are using only the first 6 GB of disk space, which is able to sustain 1 GB/s of data flow, for real-time streaming. To avoid file system penalties, these 6 GB are not formatted and data is written into a raw partition. After the end of an experiment, the data is extracted and moved to an ext4 partition for longterm storage. Owing to extensibility of our platform, we can further boost reconstruction performance by replacing internal GPUs with external boxes allowing up to 12 GPUs. The I/O performance can be further increased by attaching more storage boxes as well. However, to handle the expected data flow of 4 GB/s we are now building a small-scale cluster system consisting of several described servers connected with QDR Infiniband interconnect.

The data processing chain at ANKA needs to be very flexible. Different algebraic and analytic techniques are used for tomographic and laminographic reconstruction. Automated post-processing,

like segmentation and optical-flow, is sometimes performed as well. To simplify the implementation of efficient image processing applications for different scenarios using our hardware platform, a parallel processing framework has been developed. This framework operates in a stream-like fashion on data that is either coming from our Unified Camera Abstraction library (streamlining access to a range of high-speed cameras) or pre-recorded sequences. Image processing algorithms are implemented as pipelines of pre-defined filter nodes. Each filter is able to execute on a GPU via OpenCL and can use optimized code for specific hardware platforms if necessary. To simplify development of filters, the framework abstracts some details of OpenCL, e.g. automatic transfer of data between GPUs and host system. The framework incorporates easy to use interfaces for end-users, Python bindings to simplify development, and interfaces to the popular synchrotron control system *Tango*.

The prototype system is currently under test at ANKA. The basic layer of the parallel processing framework has been implemented and a couple of image processing algorithms including tomographic reconstruction are already ported to the framework. New algorithms will be added as plugins for the framework as required. The clustering solution is ordered and software implementation is expected by mid of 2012. Finally, the 4 GB/s version of camera is expected in 2013.

[1] S. Chilingaryan, A. Mirone, A. Hammersley, C. Ferrero, L. Helfen, A. Kopmann, T. dos Santos Rolo, P. Vagovic: A GPU-Based Architecture for Real-Time Data Assessment at Synchrotron Experiments. IEEE TNS, 58,4 (2011) 1447-1455.

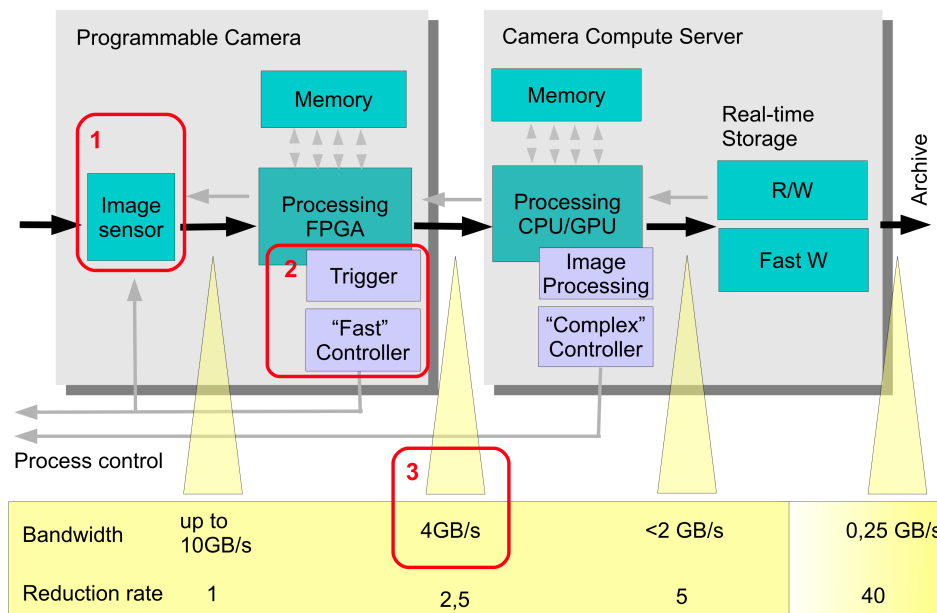


Figure 1: Concept of programmable camera and attached processing server. Highlights of the programmable camera are its modular design with a replaceable image sensor (1), application specific camera-side trigger, compression and control algorithms (2) and the high-speed interface to the compute server (3).