FPGA-Based Acceleration of an Image Registration Algorithm

Jay Brockman, Daniel Rinzler, and Peter Bui
University of Notre Dame, Notre Dame, IN

Frank Iannarilli
Aerodyne Research, Inc., Billerica MA

Image registration is a critical component of many applications, ranging from stabilization and tracking, to medical diagnosis. The image registration problem is defined as the process of finding a set of parameters for an affine transformation from a target image to a transformed image, such that the pixel-by-pixel, squared error is minimized. In this paper/poster, we will describe a methodology for using an FPGA accelerator, coupled to an AMD Opteron via a HyperTransport channel, which can perform 500 registrations per second. We assume an image size of 256x256 pixels, where each pixel is a 16-bit grayscale value.

Our approach is based on the Thevenaz image registration algorithm [1]. The main steps in this algorithm are:

- *Acquire image:* receive a 256x256 pixel frame from the image sensor
- *Pyramid:* perform a hierarchical decomposition of the image using intensity averaging techniques
- *Prefilter and gradient:* perform an initial convolution filter and compute spatial gradient
- *Transform:* given the current guess at a set of affine transformation parameters, transform the image using the inverse transformation method
- *Error and Hessian:* compute the error and its Hessian between the target and transformed images
- *Optimize:* guess at a new set of transform parameters using the Marquardt-Levenberg method.

With the exception of the Optimize step, each of these steps can be parallelized, which a profiling analysis shows to be more than 98% of the workload when executed on an Opteron alone. Analysis shows that we can maximize the speedup of the parallel portion of the algorithm by using a combination of techniques including vector operations, pipelining, and multithreading. To implement this, we are using a novel accelerator architecture that associates a custom-designed, yet programmable "pixel processor" with blocks of a partitioned image. The pixel processor is based on an ultra-lightweight processor-in-memory scheme, which was developed as an outgrowth of the DARPA HPCS Cascade project. Our hardware platform is based the XtremeData XD1000 accelerator module, which consists of a module with an Altera Stratix FPGA that is pin-compatible with an AMD Opteron, and that can communicate with an Opteron using the HyperTransport protocol. The module can be swapped for an Opteron on any multiprocessor Opteron motherboard.

At this point in time, we have completed a simulation of the system and are in the process of implementation. By the time of the HPEC Workshop, we will have initial measured performance results on communication between the Opteron and the FPGA module, as well as between processors within the FPGA to supplement our simulations and to give predictions on scalability.

1. P. Thévenaz, U.E. Ruttimann and M. Unser, "A Pyramid Approach to Subpixel Registration Based on Intensity," IEEE Transactions on Image Processing, vol. 7, no. 1, pp. 27-41, January 1998.