really what we want. There are a number of I/O libraries that can provide this, two popular ones being NetCDF and HDF5. Both of these have parallel versions of the serial libraries and allow architecture-neutral files to be created. However, the parallel versions of these libraries are still at the developmental stage. They are built on top of MPI-IO and it is necessary to understand aspects of MPI-IO in order to use the libraries effectively.

MPI-IO allows the input and output of binary files using all the processes within an MPI group and offers the advantage of output to a single file. File access is performed using MPI derived datatypes, allowing fast I/O using collective operations. Data is packed into a file in a manner consistent with a serial program and so data can be read out on any number of processors and hence files are re-usable on a given machine. There are in fact three different format options for writing data: native, internal and external32. The first of these is the native format of the machine; internal is understood by the whole MPI environment, even if it happens to be heterogeneous; and external32 is a completely portable, machine-independent format. Unfortunately, external32 is not currently available on Altix machines but once it is, it will make MPI-IO an even more attractive option.

Again, the details of using MPI-IO are beyond the scope of this article. However, if your appetite has been whetted, then we'd like to point you in the direction of our "MPI One-Sided Communication and MPI-IO" course. In this one day course, we delve into the mysteries of these important features of the MPI-2 standard and explain how they can be used to improve the performance and scaling of your code. A course is likely to be scheduled in the near future. Please contact jon.gibson@manchester.ac.uk to register your interest now, as places are likely to be limited.

Technical Symposium on Reconfigurable Computing with FPGAs, 21-22 February 2005

Kevin Roy and Mike Pettipher
Research Support Services, Manchester Computing, University of Manchester

In February 2005, the University of Manchester hosted a 2 day symposium on Reconfigurable Computing with FPGAs (Field Programmable Gate Arrays). This meeting was sponsored by Cray and SGI, and supported by the Ohio Supercomputer Center (OSC), who hosted a similar meeting in October 2004. The focus of the symposium was the use of FPGAs for High Performance Computing.

For people that have not come across FPGAs before, they are essentially hardware that can be programmed to do whatever they are tasked with. The millions of logic gates on the chip allow a flow of data or bits; flows can be constructed into algorithms to solve complex problems. The real benefit is that, because of their reconfigurability, an algorithm written for an FPGA allows you to create a processor to solve your particular problem rather than using the main CPU which has a rather rigid structure (set numbers of floating point units, integer units, loads/stores per cycle etc). This benefit is highlighted by one of their main uses in prototyping digital circuit designs.

Two of the major HPC vendors are now actively pursuing this technology - Cray have been marketing a HPC system with (optional) FPGAs, the Cray XD1, and SGI are soon to be offering optional FGPA bricks that can be accommodated in an Altix.

Other vendors from the FPGA market are also targeting the HPC community – the symposium had speakers from XIlinx on the underlying hardware and future chips, Nallatech who spoke on the history and commercial realities of FPGAs, Celoxica who spoke about implementing algorithms in FPGAs, Mitrion who spoke on programming for FPGAs and Star Bridge systems who spoke about development environments for FPGAs.

This is not the new area it seems as was highlighted by many talks from the research community, including the University of Durham, University of Saarland and NASA.

It was clear from the event that FPGAs have a long way to go to achieve a more widespread adoption in the HPC community. The potential benefit for some applications seems impressive but this is less clear for floating point arithmetic, and the development environments seem immature with algorithms needing to be coded in lower

level languages (of particular concern for many users there appears to be no Fortran compiler and none on the horizon either).

The most encouraging thing is there does seem to be a drive both from the vendors of HPC and the user community to see where FPGAs can go, and to see if they can solve some of the future problems facing HPC systems.

The meeting was very successful attracting over 100 delegates from very varied backgrounds and countries – the objective of bringing together people with a common interest in this technology was certainly achieved. The new OpenFPGA forum (see separate article) will also help to drive forward the use of FPGAs in HPC.

Figure 1: Overview of Virtex by Clive Walker of Xilinx.

For further information, please see the following review: http://www.hoise.com/primeur/05/articles/monthly/CL-PR-03-05-1.html, or contact either of the organisers Kevin Roy or Carl Ward at the University of Manchester. A CD of the presentations is available from Carl Ward (carl.ward@manchester.ac.uk).

OpenFPGA Effort Announced at Manchester Reconfigurable Computing Conference

Eric Stahlberg, Senior Systems Manager and Kevin Wohlever, Technical Director, Ohio Supercomputer Center

The Manchester Reconfigurable Computing Conference was the perfect opportunity for the Ohio Supercomputer Center to announce an effort to bring together developers and hardware manufacturers, academic, government and commercial organizations to work together to advance the use of FPGA technology in high level applications. The effort has the mission to promote the use of Field Programmable Gate Arrays in high-level and enterprise applications by collaboratively defining, developing and sharing critical information, technologies and best practices for exploiting FPGA applications.

Following an earlier OSC conference on reconfigurable computing hosted in October 2004, efforts began to determine the interest of these diverse communities in banding together to solve common problems of portability, interoperability and intra-application communication. With international interest confirmed, the task to create the organization commenced in February with the formation of an ad hoc steering group with representatives from multiple application areas, computing centers, government, academic and commercial organizations spanning multiple countries. This group has established objectives to pursue in several areas including characterizing best practices, exploring standardization, improving education and promotion of reconfigurable computing solutions and encouraging broad participation and collaboration.

More information on this effort can be obtained at the organization's new website located at www.openfpga.org where those interested can register and become a part of the effort.