Experiences on 64 and 150 FPGA Systems

Olaf O. Storaasli
Oak Ridge National Laboratory
Olaf@ornl.gov

Dave Strenski
Cray Inc.
Stren@cray.com

Abstract

Four FPGA systems were evaluated: the Cray XD1 system with 6 FPGAs at ORNL and Cray, the Cray XD1 system with 150 FPGAs at NRL and the 64 FPGAs on Edinburgh's "Maxwell". Their hardware and software architectures, programming tools and performance on scientific applications are discussed. FPGA speedup (over a 2.2 GHz Opteron) of 10X was typical for matrix equation solution, molecular dynamics and weather/climate codes and upto 100X for human genome DNA sequencing. Large genome comparisons requiring 12.5 years for an Opteron took less than 24 hours on NRL's Cray XD1 with 150 Virtex FPGAs for a 7,350X speedup.*

Hardware and Software Architectures

The 64 FPGA Edinburgh Maxwell and the Cray XD1 supercomputer with 150 FPGAs at NRL along with the development software and tools used are described, highlighting their strengths and weaknesses.

Equation Solver, MD & Climate Applications

Matrix Equation Solution, Molecular Dynamics and Climate/Weather applications results were typically 10X faster for one FPGA and scalable in proportion to the number of FPGAs. Double precision accuracy was obtained rapidly by adaptively refining single-precision calculations.

Genome Sequencing Application

FASTA [**1**] DNA sequence searches produce alignment scores. FASTA uses the same input as BLAST [**2**], and ssearch34 [**3**] consumes 98.6% of time in a Smith-Waterman FPGA pipeline kernel [**3-4**] (**Fig. 1**) to calculate the maximum alignment score for two sequences.

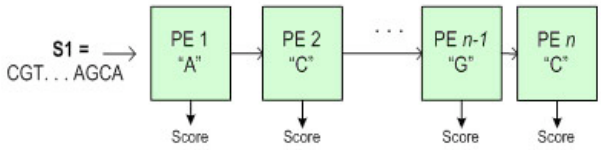

Figure 1. Smith-Waterman FPGA Pipeline

*The Naval Research Laboratory/Edinburgh University kindly provided access to 150/64 FPGA Supercomputing systems.

One query character is preloaded into each processing element to calculate scores in the column of that query character. The database string (S1) is shifted through the pipeline so each query and database character are compared in parallel, resulting in a table of scores.

Genome Sequencing Results:

FPGA timing results (for up to 150 FPGAs) were obtained and compared with up to 150 Opterons for sequences of varying size and complexity (e.g. 4GB openfpga.org human DNA benchmark and 155M human vs. 166M mouse DNA). **1 FPGA: Bacillus_anthracis DNA compare**: Genomes were matched on Virtex2, Virtex4 and Opteron Cray XD1 systems on a large database, AE016879 for 18 DNA query sequences AE017024-AE017041 for detailed and minimal I/O. Comparing each 300K character query sequence with a 5M character database took over 3 days (2.2 GHz Opteron). Ssearch34 (for 16k and 8k query sizes and both output options) on VirtexII/Virtex4 FPGAs at ORNL/Cray achieved 50x/100x speedup whether 8k or 16k query size. **Fig. 2** shows 100X speedup to compare DNA (for minimal output) on one Virtex4 FPGA (over a 2.2 GH Opteron),

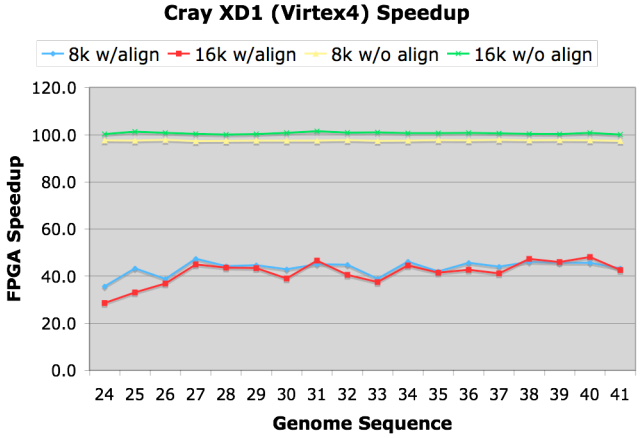

Figure 2. Virtex-4 LX160 speedup

reducing DNA search times from 14 weeks to one day. Detailed output results were less; 40X faster as the Opteron halted the FPGA to perform I/O. The Virtex4 was twice the speed of the Virtex2

Increased Query and Database Sizes: The same query sequence and database was run 30 times splitting the query

sizes: 1k, 2k, 4k, 8k, 16k characters and database: 16k, 32k, 64k, 128k, 256k,512k characters. FPGA speedups increased with query size and were independent of database size (4).

Scalability to 150 FPGAs: The FASTA code was run on 150 FPGA nodes (144 Virtex2 and 6 Virtex4) to compare the 155M humanX and 165M mouseX DNA characters, 155Mx106 x 165Mx106 x 2 = $51x10^{15}$ compares. Sequentially, this would take 138 days (12M seconds), a solution rate of 4.3 TCUPS (Tera Cell Updates/Sec). Since this rate was for the detailed output option, based on identical runs on single single Virtex4 FPGAs cited above, one should achieve an additional 2.5X speedup just by selecting the minimum output option, or 10.75 TCUPS.

All 150 FPGAs were not available simultaneously, but our 150 jobs achieved considerable parallelism by assigning them as FPGAs became available (over a two week period). The horizontal lines in Fig. 3 show the time each job took. The shorter lines (6 Virtex4s) finished in half the time.

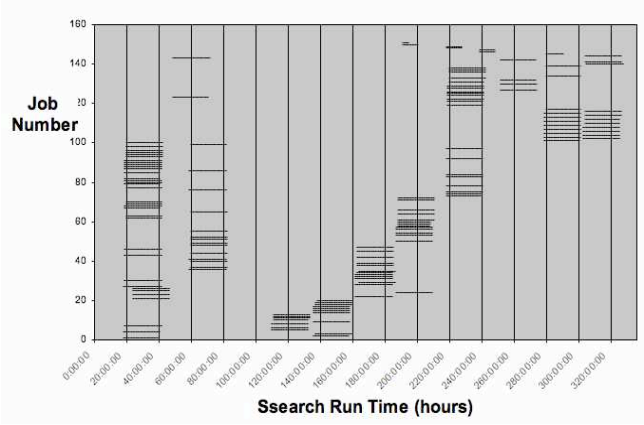

Fig. 3. Cray XD1 speedup (150 FPGAs) with 46 TCUPS

This parallelism, albeit not perfect, reduced wall clock time from 138 days to 12.9 days (1M seconds), a solution rate of 46 TCUPS, perhaps a new world record.

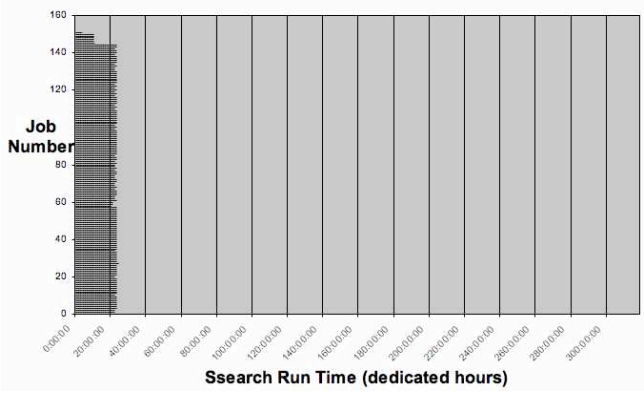

Fig. 4. Ideal Cray XD1 speedup (150 FPGAs) with 0.6 PCUPS

Ideally, Fig. 4. Shows how dedicating all 150 FPGAs to our DNA comparison could reduce the wall clock time to 1 day (85K seconds), a solution rate of 605 TCUPS or 0.6 PCUPS (Peta Cell Updates/Sec). The minimal output option here should result in a solution rate of 1.5 PCUPS.

Observations:

Examining carefully the results shows that unlike the previous Human DNA example exhibiting a 2.5X speedup resulting by simply selecting the minimal output, in this case the speedup may be considerably more. This is indicated by an examining a profile of the total run time of 24 hours and noting that consistently for all 150 jobs, only 2 hours is spent on FPGA time for DNA comparison with 22 hours spent on I/O, for a potential speedup factor of 10X. This would indicate that simply reducing the output option to minimal could result in a solution rate of 43TCUPS for our actual parallel solution, and ideally 6 PCUPS.

A DNA sequencing application that would take one 2.2 GHz Opteron 12.5 years to solve, was found to take 6 weeks for 150 Opterons to solve in parallel. Using 150 FPGAs on NRL's Cray XD1 reduced this time to 24 hours, a total speedup of 7,350X over a single Opteron.

Conclusions

Both single and multiple (64 and 150 FPGA systems) were evaluated. Single FPGA results achieved typically 10X for Matrix Equation Solution, Molecular Dynamics and Climate/Weather applications, while DNA Sequencing (using the FASTA code) achieved 100X speedup.

Scalability was shown by using 150 FPGAs in parallel with an actual solution rate of 46 TCUPS and an ideal rate of 605 TCUPS (with all FPGAS available to begin simultaneously). The 150 FPGAs achieved a 7,350X speedup over a single 2.2 GHz Opteron processor in the solution of a massive HumanX comparison with MouseX DNA. These results indicate similar speedups are likely for acceleration modules (DRC, XtremeData, etc.) that fit in Opteron sockets, to significantly speedup calculations on small embedded systems and Cray XT supercomputers.

References

[1] FASTA Sequence Comparison Code: fasta.bioch.virginia.edu
[2] MitrionC/BLAST: http://www.hpcwire.com/hpc/1274236.html
[3] Margerm, Steve, and Maltby, Jim; Accelerating the Smith-Waterman Algorithm on the Cray XD1, Cray WP-0060406 2006.
[4] Storaasli, Olaf, Yu, Weikuan, Strenski, Dave, and Malby, Jim; Perfomance Evaluation of FPGA-Based Biological Applications, Cray Users Group Proceedings, Seattle WA, May 2007.

Acknowledgment

Research sponsored by the Laboratory Directed Research & Development Program of ORNL managed by UT-Battelle for the U. S. Department of Energy on Contract.DEAC0500OR22725.