

Engineering Applications on NASA's FPGA-based Hypercomputer

Olaf O. Storaasli, Sr. Research Scientist, NASA Langley
Olaf.O.Storaasli@nasa.gov Phone: 757-864-2927

Introduction

This purpose of this presentation is to describe how NASA Langley's reconfigurable Field Programmable Gate Array (FPGA)-based hypercomputer research is addressing the solution of comprehensive engineering and scientific calculations. Two approaches are used to exploit Langley's Star Bridge Hypercomputer Systems for analysis calculations:

1. Develop analyses codes in VIVA for Hypercomputer (fully exploits parallelism)
2. Use Hypercomputer to accelerate time-consuming (bottleneck) calculations

Since NASA C++/FORTRAN legacy codes do not exploit all of the FPGA parallelism possible (hundreds of operations/cycle), the algorithms were entirely written in the VIVA language using the first approach¹⁻². However, the second approach was used for a large legacy code where over 95% of the finite element equation solution computations are concentrated in a two-page FORTRAN kernel. This matrix-factor kernel was replaced by VIVA "gateway" to exploit FPGA parallelism. This VIVA kernel development involves researchers at Alpha-Star Corporation ([GENOA](#) structures code), Starbridge Systems ([VIVA](#) developers), and NASA (developed [GPS Solver](#) used in [GENOA](#)).

NASA FPGA-based research initially focused on rapid structural analysis, but has now been extended to include linear algebra, matrix equation solution and integration (Runge-Kutta for fluid dynamics and Newmark-Beta for finite element structural mechanics). The presentation builds on previous NASA research¹⁻² and describes "lessons learned", via close Starbridge collaboration, to overcome early HAL-15 and VIVA limitations, which has led to rapid and accurate scientific and engineering analysis calculations.

Related Work

High-speed parallel Vector-Matrix 64-bit floating-point (SAXPY) calculations are the key to rapid, accurate structural analysis. Unlike 32-bit (imaging) FPGA applications, Langley focused FPGA research on supercomputer applications. Langley was first to install Starbridge Hypercomputers and collaborates with the National Security Agency, U.S. Air Force and other experts to speed progress. This research has led to the new \$15M NASA Reconfigurable Scalable Computing (RSC) project for space applications.

NASA Langley Hypercomputer

Langley's HC-38m has 5 Xilinx Virtex II 6000 FPGAs with the following characteristics:

Gates	6 million (97x)
Memory (on chip)	2.6 Mb (175x)
18x18 on-chip Multiplies	144 (none)
Comp Logic Blocks-CLB = 4 slices	8,448 (4x)
Memory Speed	5 Tb/s (11x)

Table 1. Xilinx XC2V6000 Specifications (improvement over HAL 15 FPGA)

Figure 1 shows 10 FPGAs, each with 2-8GB DRAM, connected via 225 Gb/sec inter chip communications situated on a dual board seated in a standard PCI-X slot. There are 256 external I/O pins available for use over 20 64-bit parallel memory channels.



Fig 1. Starbridge FPGA board installed in PCI-X slot

Application Algorithms

NASA Langley was first to develop key algorithms to exploit Hypercomputer FPGAs¹: matrix/vector and linear algebra operations including $[A]\{x\} = \{b\}$ matrix equation solvers, greatest common divisor, factorial, transcendental (trig, log, etc.) functions via the Cordic algorithm, differentiation and integration (Runge-Kutta for fluid dynamics & Newmark-Beta for structural mechanics), structural dynamics: $[M]x'' + [C]x' + [K]x = f(t)$ where $[M]$, $[C]$ and $[K]$ are mass, damping and stiffness matrices, $f(t)$ is the force vector, x is the displacement and (') denotes the time derivative. Nonlinear applications using analog computing methods are solved using VIVA with digital accuracy. The presentation will focus on "lessons learned" from these initial applications and show how FPGA-based systems can now be used to solve comprehensive structural analyses (and related applications) where the solution of a large systems of matrix equations dominates the calculations. The tradeoff of completely rewriting legacy application codes in VIVA versus accelerating legacy codes using VIVA kernels to exploit FPGA parallelism will be discussed. The most challenging application involves the Space Shuttle Columbia re-entry analysis which takes in excess of 660 hours without FPGA computation speedup.

References

1. Storaasli, Olaf "Computing Faster without CPUs: Scientific Applications on a Reconfigurable, FPGA-based Hypercomputer," 6th MAPLD Conference, Sept 9-11, 2003.
2. Storaasli, Olaf et al. "Scientific Applications on a NASA Reconfigurable Hypercomputer," 5th MAPLD Conference, Sept 10-12, 2002.