SAND2015-0128C

Preparation of codes for Trinity

Courtenay T. Vaughan, Mahesh Rajan, Dennis Dinge, Clark Dohrmann, Ken Franko, Mike Glass, Kendall Pierson, and Mike Tupek
Sandia National Laboratories[1]
P.O. BOX 5800, Albuquerque, NM 87185

Abstract - Los Alamos National Laboratories (LANL) and Sandia National Laboratories (SNL) have formed a partnership, the New Mexico Alliance for Computing at Extreme Scales (ACES), to acquire and deploy a production capability system for the Department of Energy NNSA ASC program to support the national Stockpile Stewardship Program. Trinity, a Cray XC40, with 9,780 nodes with the Intel Haswell processors and 9984 nodes with the Intel Knights Landing processors is first of the NNSA's new Advanced Technology Systems, ATS-1, anticipated to be installed in 2016. As part of the Trinity Center of Excellence with Cray and Intel, SNL code developers and support staff are working on porting three SIERRA mechanics codes to Trinity. These complex multi-physics applications are: SIERRA/Solid Mechanics (SM), SIERRA/Structural Dynamics (SD) and SIERRA/Aerodynamics. In this paper, we will detail the preliminary work that we have undertaken in preparation for efficient use of these production applications on Trinity. Trinity's architecture challenges these codes to achieve high node level and thread level parallelism, high vectorization efficiency and efficient use of the 576 Burst Buffer I/O nodes with NVRAM. We will present profiles of the codes with tools including CrayPat, vTune, TAU and HPCtoolkit on our current large Cray XE6 called Cielo, and on our large Sandy Bridge/Infiniband cluster called Chama and on our Knights Corner test beds called Compton and Morgan. All the three SIERRA applications (for SIERRA/SM the implicit case) for the target simulations of interest show a large fraction of the simulation time in matrix equation solves, albeit with different characteristics. These performance profiles and scaling studies will focus on leveraging Trillinos and MKL/Pardiso sparse solver kernels. In SIERRA/SM the preconditioning step dominates consuming more than 90% of the solve time, as it is computed at each time step. The iterative linear solve done with the FETI (Finite Element Tearing and Interconnect) algorithm requires a local solve, a coarse solve and a preconditioner solve. In contrast, computations for SIERRA/SD are dominated by forward/backward triangular solves associated with already factored matrices for local and global (coarse) problems. That is, much more time is spent applying the preconditioner than in its initialization. We also will present our efforts with vectorization of key compute kernels using tools like Intel's Vector Advisor. For code kernels for which the current generation of compilers is unable to auto-vectorize, we will discuss data structure layout modifications and the direct use of Intel vector intrinsics to systematically improve vector performance.

Abstract for program - Sandia and Los Alamos National Laboratories are acquiring Trinity, a Cray XC40, with half of the nodes having Haswell processors and the other half having Knight's Landing processors. As part of our Center of Excellence with Cray, we are working on porting three codes, a Solid Mechanics code, a Solid Dynamics code, and an Aerodynamics code, to effectively use this machine. In this paper, we will detail the work that we have done in porting the codes in preparation of getting the machine. We have started by profiling the codes using tools including CrayPat, which showed that a large portion of the time is being spent in the solvers. We will describe the work we are doing on the solvers such as ongoing work on Haswell processors and Knight's Corner machines.

[1]Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.