Code Performance

CIG plans to primarily use four codes on XSEDE for research and further develop these codes to improve their performance and scalability. The scalability and performance of these codes measured on XSEDE resources is presented below.

Calypso

Calypso is being developed for magnetohydrodynamics based geodynamo studies. It uses a pseudo spectral method for toroidal and poloidal components in combination with a finite difference method for radial components. We tested the scalability and performance of Calypso on Stampede up to 16384 cores (see Figure 1). The largest scaling test above corresponds to roughly a 17 million DOFs for one scalar and scales well up to 34304 cores. The scaling result on Stampede2 shows good scalability for the both KNL and Skylake (XEON) nodes. Comparisons between the KNL processor and XEON processor on Stampede2 show the strong scaling on both processors are very similar.

Rayleigh

CIG also supports the development of Rayleigh, a state of the art code for dynamo simulations in collaboration with Dr. Nick Featherstone (University of Colorado Boulder). Rayleigh is based on the solar dynamo code ASH – Anelastic Spherical Harmonic, and scales efficiently to more than 10,000 cores. Rayleigh has been successfully applied to simulate planetary and the geodynamo. Figures 2 show the recently measured scaling of Rayleigh code on TACC Stampede2. The performance of both the KNL and XEON processors are very similar.
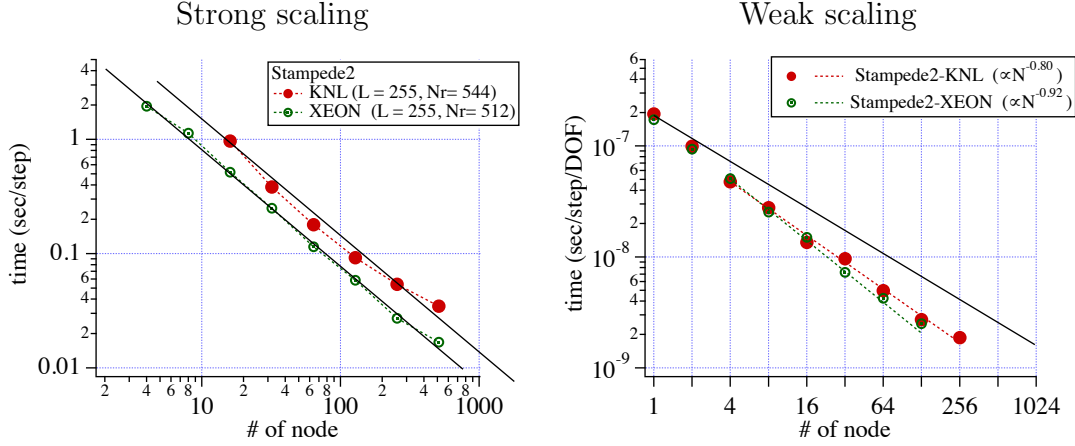
Figure 1: Comparison of Calypso's scaling on the TACC Stampede2. Strong scaling results is shown in the left panel, and weak scaling results are shown in the right panel. Ideal scaling ($O(N_{node}^{-1})$ and $O(N_{node}^{-2/3})$ for the strong and weak scaling, respectively) is plotted by solid lines.
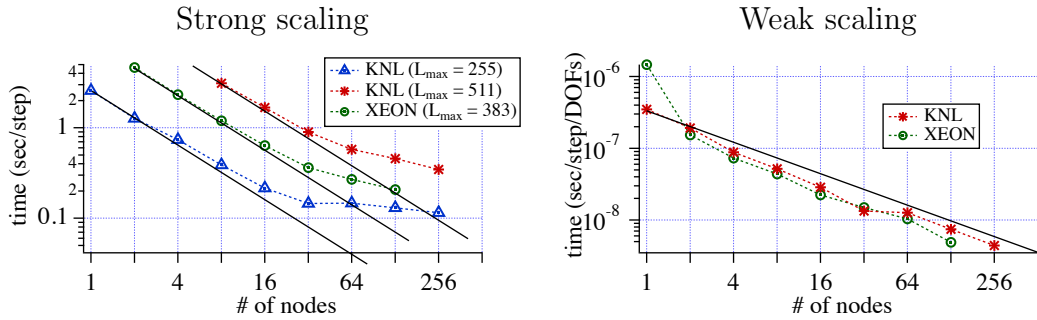

Figure 2: Rayleigh's scaling on the TACC Stampede 2. Strong scaling results is shown in the left, and weak scaling results are shown in the right. Ideal scaling ($O(N_{node}^{-1})$ and $O(N_{node}^{-2/3})$ for the strong and weak scaling, respectively) is plotted by solid lines. Rayleigh only uses MPI parallelization, and 64 of 68 processor cores are used in the each KNL node.

2

SPECFEM3D_GLOBE

In collaboration with Princeton and CNRS (France), CIG offers this software, which simulates global and regional (continental-scale) seismic wave propagation using the spectral-element method (SEM). The SEM is a continuous Galerkin technique, which can easily be made discontinuous; it is then close to a particular case of the discontinuous Galerkin technique, with optimized efficiency because of its tensorized basis functions. Dr. Hiroaki Matsui recently measured scaling of SPECFEM3D_GLOBE on TACC Stampede2 (see Figure 3).
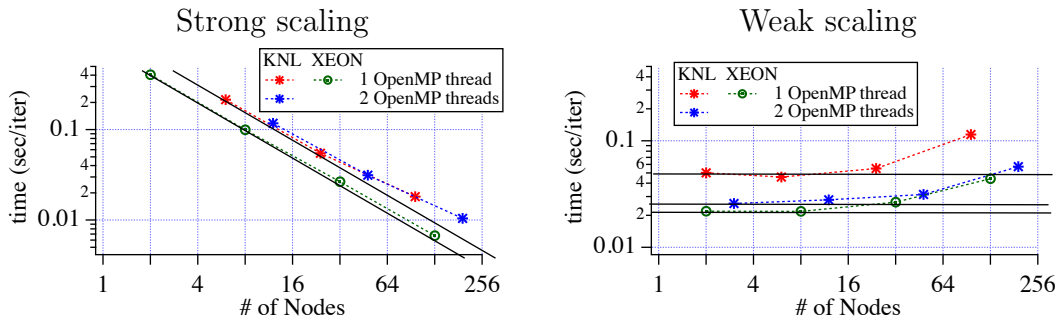

Figure 3: SPECFEM3D_globe scaling on the TACC Stampede2. Strong scaling results is shown in the left, and weak scaling results are shown in the right. Ideal scaling ($O(N_{node}^{-1})$ and constant for the strong and weak scaling, respectively) is plotted by solid lines. In the scaling tests for the KNL nodes, 64 of 68 processor cores are used and 1 or 2 OpenMP threads cases are tested.

ASPECT

ASPECT performs mantle convection and lithospheric deformation simulations using a finite element method. ASPECT utilizes the Trilinos library for linear solvers and their preconditioners (support of the PETSc library is experimental). Scaling capabilities of the stable ASPECT version for large-scale 3D mantle convection simulations on Stampede2 are shown in Figure 4 for a spherical shell model. The result shows that ASPECT scales well on problems up to 1.3×10^8 DOFs and up to over 1000 processes.

Recently, Prof. Timo Heister and Thomas Clevenger have implemented a new experimental Geometrical Multigrid (GMG) solver into ASPECT and performed strong scaling on Skylake nodes of Stampede2. The new solver that will be increasingly utilized during the next phase of XSEDE, keeps good scaling up to 24576 (512 nodes) with 3.4B degree of freedoms problem (see Figure 5). These results are a significant improvement over the solver in the stable release version of ASPECT.
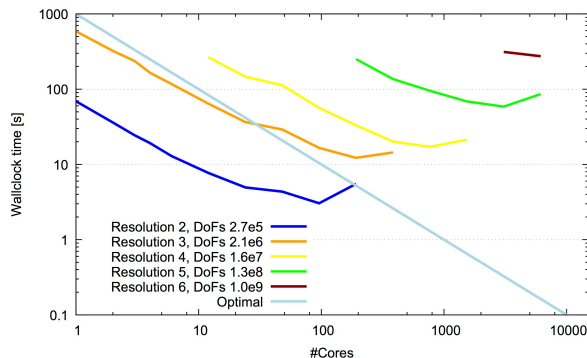

Figure 4: Legend: Strong-scaling results for ASPECT on Skylake nodes on TACC Stampede2. The scaling is derived from the wall-clock time of a thermal-mechanical model run for two time steps.
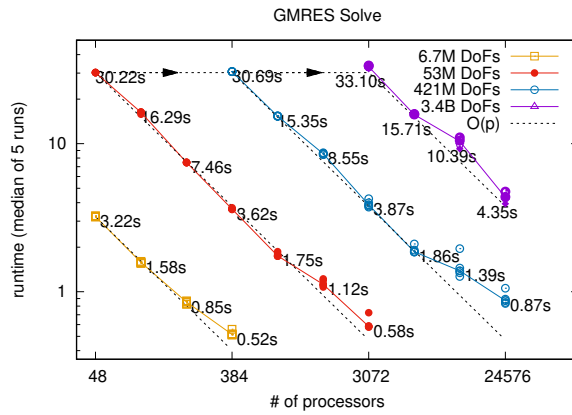
Figure 5: Elapsed time for a globally refined Stokes problem of various refinements and varying number of processors. The solver uses a block Schur complement preconditioner where we estimate the inversion the velocity block of the Stokes system with one GMG V-cycle. We plot 5 separate runs and the solid lines connect the median of the 5 runs. We see good scaling to between 17K-34K DOFs per processor.