

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 a recently developed magnetohydrodynamics (MHD) simulation code for geodynamo studies. It uses a pseudo spectral method for solenoidal and poloidal components in combination with a finite difference method for radial components. We tested the scalability and performance of Calypso on Stampede with $(N_r, N_\theta, N_\phi) = (512, 384, 768)$ up to 16,384 cores (see Figure 1. Calypso keeps 83% of parallel efficiency with 16,384 cores.

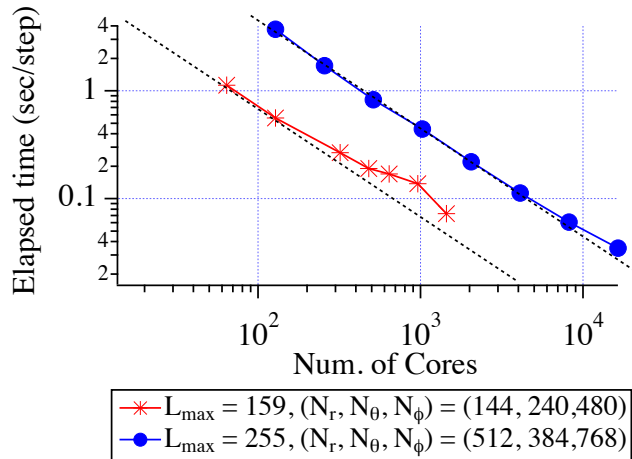


Figure 1: Calypso’s strong scaling on the TACC Stampede system for different sized simulations. Ideal scaling is plotted as a dotted line.

Because Calypso is parallelized by MPI and OpenMP, we also test Calypso by using MIC coprocessor (see Figure 2). The results suggests that we need further optimization to improve the performance on the MIC coprocessor.

Further optimization development for Calypso is being carried out for graphics processing units (GPUs) using CUDA. For the first step, CUDA is

implemented for the Legendre transform, which is the most time consuming part in Calypso. Strong scaling of the Legendre transform on TACC Maverick's GPU is plotted in Figure 3.

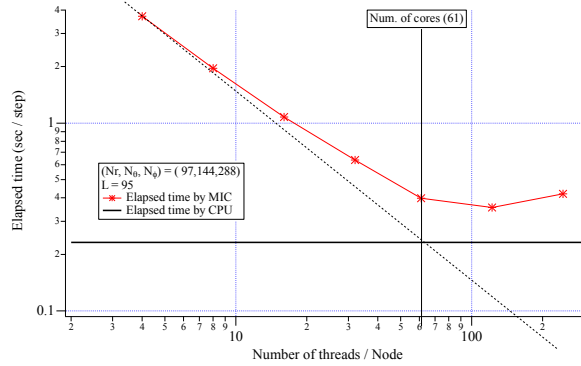


Figure 2: Calypso scaling using MIC processor on TACC Stampede. Ideal scaling is plotted as a dotted line.

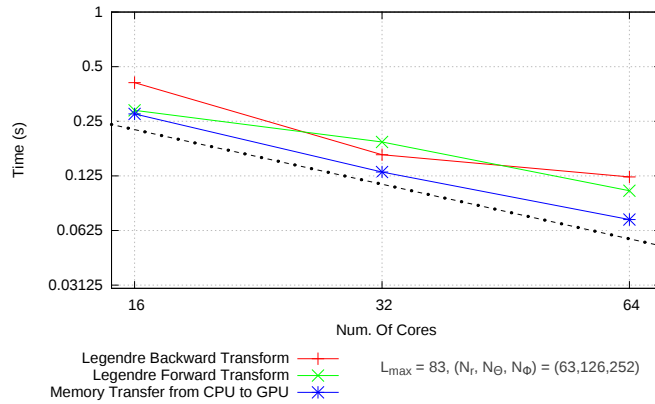


Figure 3: Calypso strong scaling using GPUs on TACC Maverick. Ideal scaling is plotted as a dotted line.

Rayleigh

CIG is also currently developing Rayleigh, a state of the art code for dynamo simulations in collaboration with Dr. Nick Featherstone (JILA, University of Colorado Boulder). Rayleigh development is based on a solar dynamo code (ASH – Anelastic Spherical Harmonic) which scales efficiently to more than 10,000 cores. Rayleigh scales well up to this number on Stampede and above on ALCF Mira. Figure 4 shows scaling of Rayleigh up to 16,384 cores, which is the maximum number of nodes for 'large' job class on Stampede. Figure 5 shows the performance results for the Rayleigh code using up to

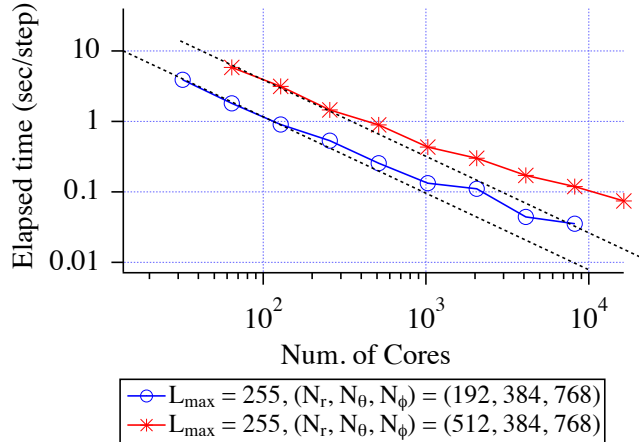


Figure 4: Rayleigh’s strong scaling on the TACC Stampede system for different sized simulations. Ideal scaling is plotted as a dotted line.

131,072 single-threaded cores on ALCF Mira. Figure 5(a) shows the parallel efficiency for the non-magnetic and magnetic runs at the 2048^3 grid points. Efficiency is defined as the ratio of the realized speedup to the ideal speedup. Ideal performance is taken relative to 16,384 cores, the minimum core count for this problem size. Figure 5(b) shows the strong scaling results for the same set of runs. Ideal scaling for each case is indicated by the dotted line.

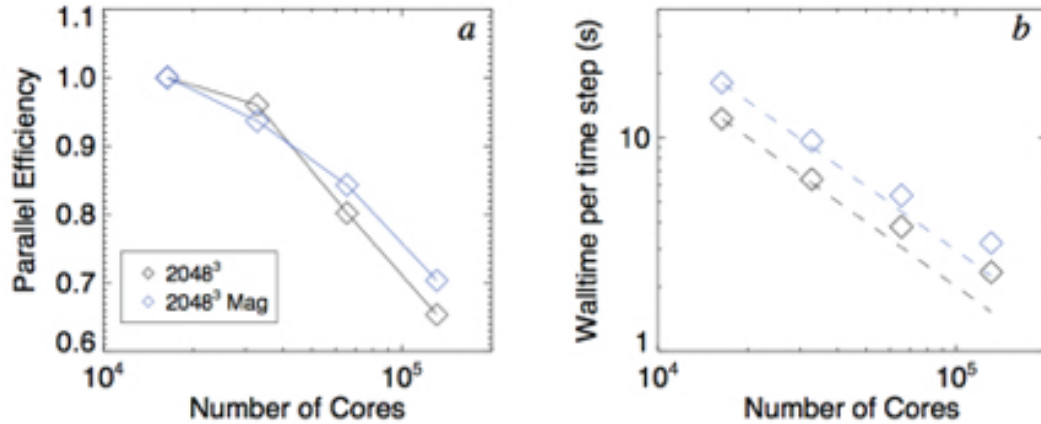


Figure 5: Rayleigh’s strong scaling on ALCF Mira up to 131,072 cores.

ASPECT

ASPECT performs mantle convection simulations using a finite element model and utilizes the Trilinos library for preconditioner and solver support (support for the PETSc library is under development). The scaling capabilities of ASPECT for large-scale 3D mantle convection simulations on Stampede are shown in Figure 6 for a 3D box model and Figure 7 for a spherical shell model. These demonstrate that this next-generation code scales well on problems up to tens of millions of elements and up to over 1000 processors.

The weak and strong scalability of ASPECT was also examined on a local cluster in detail. The results are shown in Figures 8. These plots demonstrate good scalability up to several thousand processors for large problems.

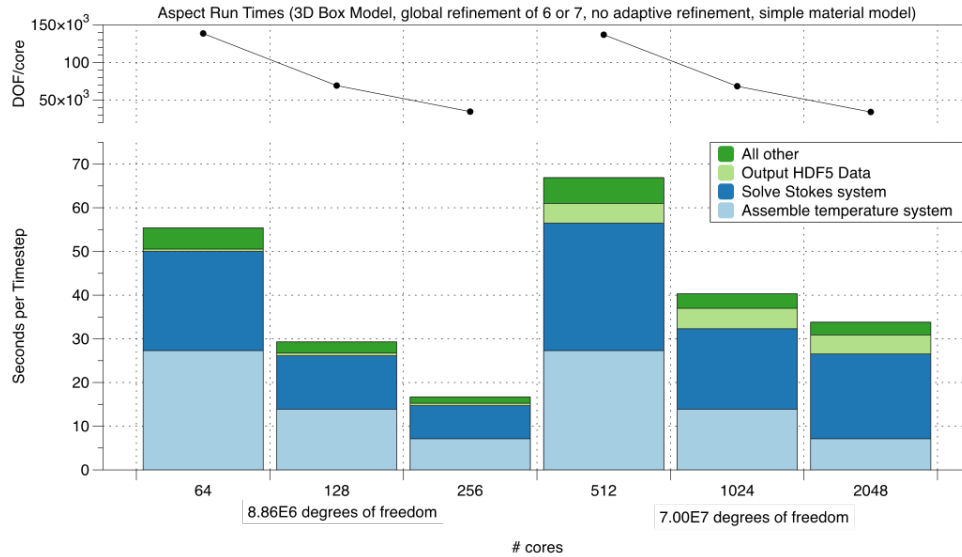


Figure 6: Strong scalability of 3D box model (128^3 or 256^3 elements) on Stampede.

CitcomS

CitcomS can use either conjugate gradient or multigrid solvers and has good scaling properties. On IBM Blue Gene/L, using 64×12 processors on a full spherical model, it has a computational efficiency of 65 percent, compared with the same problem using 12 processors [King et al., 2005]. Tests on TACC Lonestar show good scalability (> 0.6) for regional models with 6.4 million unknowns using up to 256 cores. Details regarding these results are shown in the figures 9.

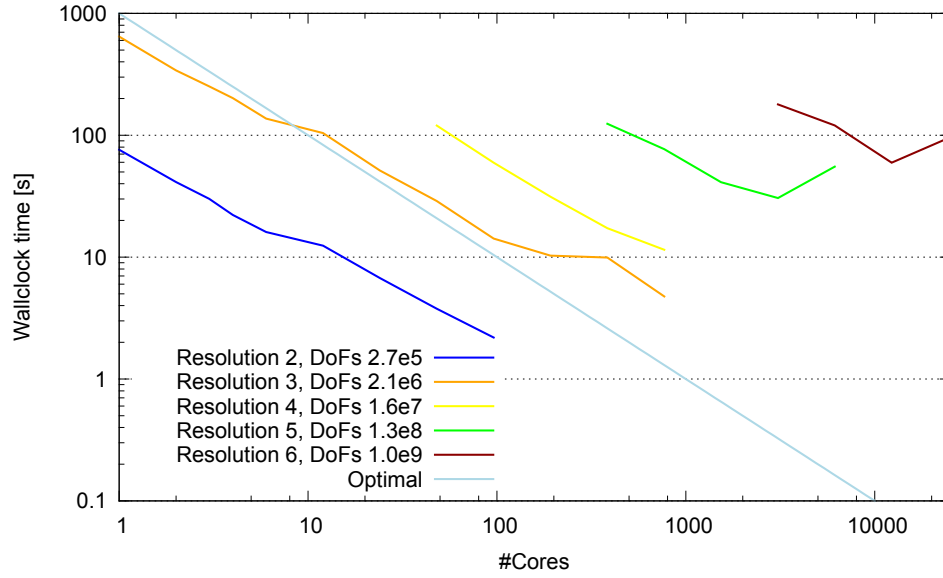


Figure 7: Strong scaling of ASPECT's spherical shell model. The wall clock time for one time step is measured by a mantle convection model including temperature dependent viscosity and composition perturbation in a spherical shell.

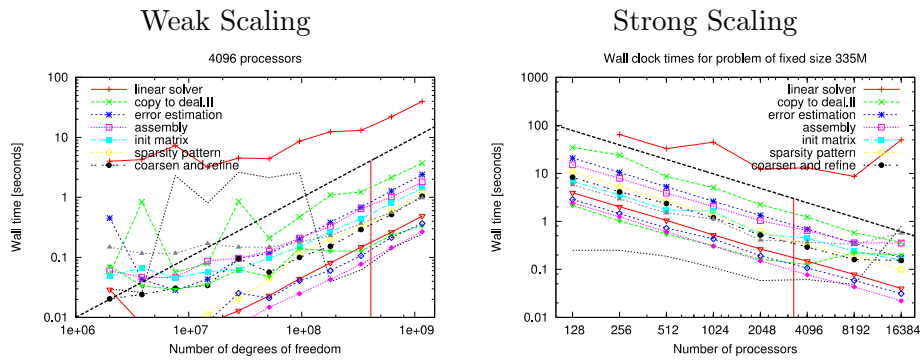


Figure 8: Weak (left) and Strong (right) scaling results for ASPECT.

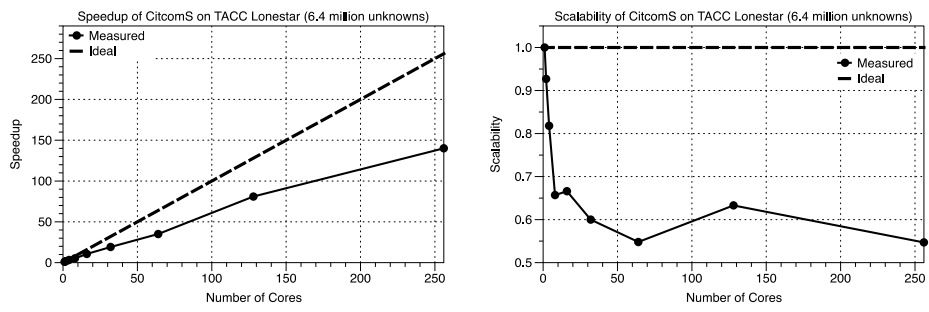


Figure 9: Speedup (left) and scalability (right) of CitcomS on Lonestar.