Scaling and Performance Analysis of Underworld: Towards the One Billion Particle Target

© R. Petersen, D. Stegman, 2010

Institute of Geophysics and Planetary Physics, Scripps Institution of Oceanography, University of California San Diego, La Jolla, USA
rpetersen@ucsd.edu
dstegman@ucsd.edu

We investigate the scaling performance of Underworld, a geodynamic modeling framework, on high performance computing resources. These results inform the configuration and allocation of resources committed to study the complex geodynamic processes in and around subduction zones. We use a 3D thermal convection model that includes a temperature and stress dependent rheology as a proxy for the computational difficulty of the eventual subduction problems. This proxy model is designed with significant variation in viscosity, in conjunction with the non-linear temperature and stress dependent rheology this results in a computational problem that is significantly more challenging than the similar isoviscous convection problem.

”Using an allocation of 30000 service units on Texas Advanced Computing Center (TACC) cluster, Ranger, we ran a suite of models that exercised Underwold by solving problems with element subdomain sizes of 16×16×8 and 16×16×16 per core. For each of subdomain size we ran the models at several global resolutions, from what we call "tiny" models (32×32×32) to what we refer to as "huge" (192×192×192). The global resolutions selected require CPU core allocations from 100’s to 1000’s.

Underworld supports both a basic FEM solution method and Particle In a Cell (PIC). We solve the thermal convection problem using both methods to verify the equivalence of solutions. At the highest resolutions using the PIC method the number of particles approaches 1.5 million. We continue to explore this model with an eye towards systems populated by up to 1 billion particles.

Timing results are measured as the walltime per model timestep. A common steady-state model is first calculated over several thousand timesteps. This steady-state model is then restarted at the appropriate resolution whence the performance data is gathered.