Prospects for Scaling of Seismic Inversion to Petaflops Systems

Omar Ghattas

Institute for Computational Engineering
Sciences, The University of Texas at Austin
omar@ices.utexas.edu

Abstract

The U.S., Japanese, and European governments are all pursuing petaflops computing programs, and the first systems capable of a peak petaflops performance will appear sometime in 2008. Problems in the geosciences have been one of the drivers for the development of such systems. In particular, the seismic inverse problem of determining the distribution of earth properties from surface observations of ground motion in large heterogeneous elastic regions often requires petaflops-class computing. The complexity of solution of this inverse problem stems from the need to solve the forward problem repeatedly for a given regularization parameter, the need to solve the inverse problem repeatedly to estimate the best regularization parameter, the need to estimate the variance in the inverted earth model, and ultimately the need to provide a complete characterization of the model uncertainty. Scaling of inversion methods requires both algorithmic (scaling to large problems sizes) and parallel (scaling to large numbers of processors) scalability. In this talk, we examine both parallel efficiency and algorithmic efficiency of several inverse methods and assess the prospects of scalability of seismic inversion to the petascale.