

Nnodes: a workflow manager for full waveform inversion in large scale HPC clusters

Congyue Cui¹, Jeroen Tromp¹, Lucas Sawade¹

¹ Princeton University

We present a workflow manager that addresses many of the pain points we have experienced when running full-waveform inversion (FWI) on large-scale HPC clusters. FWI can be computationally expensive and complicated in structure, depending on the application. Our workflow manager makes it easy to control the progress and reduces repetitive work. The design of our workflow manager focuses on three key points. First, the workflow definition is flexible to make sure that users can easily make adjustments and experiment with new ideas. Second, the progress management is adaptable so that an interrupted workflow can be stopped and resumed at any point, and it is possible to rewind to a previous state if any parameter does not turn out to perform well; running a partial workflow or merging multiple workflows are also supported. Third, parallel execution of multiple MPI functions is naturally supported. An MPI operation from any part of the workflow will be managed by an MPI executor and users do not need to manually adjust the workflow to make the most of the cluster resources. We have successfully applied the workflow manager to both regional structural inversion and global CMT inversion. Migrating existing tools to this workflow manager was seamless in our case. We hope this is a useful tool for FWI researchers and also the general HPC cluster users.

