Improving the seismic wave propagation performance in heterogeneous parallel environments with adaptive asynchronous work-stealing

João B. Fernandes, UFRN
Italo Augusto Souza de Assis, UFERSA
Tiago Barros, UFRN
Idalmis Milian Sardina Martins, UFRN
Samuel Xavier de Souza, UFRN

Seismic processing consists of a series of steps from the acquisition of seismic data to the generation of subsurface images (2D or 3D). Wave propagation modeling is one important part of seismic processing, as it is included in many of these steps. Modeling involves simulating seismic shots through a 2D or 3D model. The number of shots in a field acquisition designed to cover tens to hundreds of kilometers can easily reach thousands or tens of thousands. Thus, when modeling synthetic data, employing high-performance computers is essential. Over the years, new technologies have been emerging, and supercomputer centers have been adopting them as means to increase performance of modern programs such as seismic applications. As a result, supercomputers are becoming more diversified in the available types of hardware architecture.Different types of CPU, GPU, or FPGA might characterize supercomputers as heterogeneous. The heterogeneity is affected by a type of performance decline: load imbalance. In distributed systems, load imbalance occurs when some program processes finish earlier than others and keep idle until the others finish. To distribute the shots across the processes and avoid unbalancing, the wave propagation algorithms might use scheduling techniques. An alternative solution used in the literature for this type of application is to use dynamic scheduling, but most of the time with a centralized approach, where a main process is responsible for distributing the shots during the execution. However, this method might not be scalable when used in an environment with hundreds or thousands of nodes. In this case, the main process can become a bottleneck in the execution. Because of this, the decentralized dynamic methods, where all processes are co-responsible for managing shots, have shown to be a more suitable form of schedule. Work-stealing is one of the most known methods of decentralized dynamic scheduling. The principle of this scheduler is that faster processes are allowed to steal tasks (in this case, seismic shots) of slower ones. This work focuses on the challenges and solutions associated with executing 3D seismic wave propagation applications in a heterogeneous environment, particularly in addressing load imbalance. For this, we propose a work-stealing method named adaptative asynchronous work-stealing (A2WS), which gathers information about each computer node during the execution and employs this information to improve the distribution. Furthermore, we designed A2WS to avoid bottlenecks and be scalable in large-scale computing environments. We apply A2WS in a 3D seismic wave propagation application and compare it with a centralized dynamic (CD) scheduler, and with another work-stealing method called cycled token work-stealing (CTWS). A2WS performed 16.0% faster in comparison with CD; and 5.88% in comparison with CTWS.

This paper was prepared for presentation during the 18th International Congress of the Brazilian Geophysical Society held in Rio de Janeiro, Brazil, 16-19 October 2023.