Journal of Computational and Applied Mathematics 295 (2016) 103–114

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

Model order reduction with oblique projections for large scale wave propagation

Victor Pereyra

Energy Resources Engineering, Stanford University, CA, United States

ARTICLE INFO

Article history:
Received 5 October 2014
Received in revised form 6 January 2015

Keywords:
Model order reduction
Acoustic wave equation
Slanted projection

ABSTRACT

There are many applications in which it is necessary to solve large scale parametrized wave propagation problems repeatedly. This is still quite a challenging task, even with the largest available computer clusters. In this paper we will discuss the application of Model Order Reduction (MOR) to problems in seismic petroleum exploration, with the aim of diminishing the necessary computing time by a significant factor. We consider POD and some variants. POD is a Model Order Reduction technique that uses snapshots of a few simulations in order to quickly compute related problems with similar accuracy. The method of lines via a Petrov–Galerkin approximation that uses the snapshots as basis functions is the considered approach. The order reduction comes from projecting the wave equation discretized in space to the subspace spanned by the snapshots. This has been shown earlier to work well in two dimensions. The challenge in three dimensions comes from the size of the spatial meshes required and the fact that the method usually requires a number of snapshots that do not fit in fast memory, even for current high end multicore machines. Parallelization is not an option since it is already used for other aspects of this massive problem.

1. Introduction

Seismic exploration is one of many approaches used to obtain information on the underground to locate and extract oil and gas. The seismic method is used inland and off-shore and consists of introducing man-made sources of energy that travels through the earth. The back scatter is recorded on sensors called geophones or hydrophones. It is this time data that is used to produce depth maps of the earth interior through seismic data processing. With the advent of powerful large scale clusters of computers it has become possible to simulate the full wave propagation at various levels of complexity (acoustic, elastic, anisotropic) in order to migrate the data from time to depth to produce three-dimensional maps of the material properties that affect wave propagation: density, wave speeds and anisotropic parameters.

A 3D seismic survey consists of many shots, i.e., activation of energy sources, say explosions, vibrator trucks or air-guns and corresponding receiver arrays. Usually, a simple rectangular geometry with equally spaced sources and receivers is used and only one source is activated at a time with a recording time of a few seconds, that depends on the depth one wants to image and the approximate velocity of the involved rock formations. Because of the complexity and the large scale of the 3D meshes involved the simulation of the time domain wave equation is performed using explicit methods on regular meshes, on a box that contains the region of interest. A typical 3D survey may have thousands of shots and for each shot thousands of receivers that record the back scatter, generating enormous data sets that have to be processed.

Repeated wave equation simulation is used for many different facets of this work, from survey planning and illumination studies, to imaging, migration from time to depth, reverse time migration, seismic tomography and quality control.

E-mail address: vpereyra@yahoo.com.

http://dx.doi.org/10.1016/j.cam.2015.01.029

In recent work we have shown that model order reduction (MOR) is capable of reducing the computational cost of High Fidelity full acoustic wave simulation for shot interpolation and extrapolation in two-dimensions [1–3], when many problems with different source positions (shots) need to be simulated in a seismic survey or for seismic imaging. The requirement there is accuracy (albeit low, say 3 significant figures) and thus we cannot stray too far from the sources that produced the snapshots for MOR. We have shown that extrapolating one shot in each direction from the basis shot or interpolating two interior shots from two end shots works well when the data thus generated is used in a simulation process [3]. Taking into account the overhead, this doubles the speed at which we can perform this task, where many thousands of shots and corresponding simulations are required.

The challenge now is to extend this process to three-dimensions in a competitive manner. To fix our ideas, let us consider a typical 3D simulation that is performed today on a single multicore machine with 128 GB of fast memory. The number of nodes in each direction of the spatial mesh are $nx = ny = 1541$, $nz = 613$, for a total of $n = 1,455,679,453$ mesh points. That is to say, the size of one snapshot in double precision would be approximately 11.6 GB. From our experience in 2D and for the type of source frequencies desired, we expect to need more than 100 snapshots, i.e., just to hold a copy of the whole snapshot matrix in core would require more than 1.2 TB, far more than what is available in a modern single machine. Currently, each High Fidelity simulation of this size takes 3 h in a Sandy bridge box (16 cores, 128 GB memory, 277 GHz). This corresponds to 16″ of simulation with a time step $dt = 0.00191$ (9424 time steps). We could wait for Moore's Law to catch up, but then larger problems will come around: elastic and anisotropic wave propagation, higher frequencies and so on, so that is not the answer.

In this paper we explore a number of avenues to solve this problem competitively using some new variants of MOR. We observe that a full seismic survey simulation requires many thousands of shots and that is where parallelism in a distributed system is already employed. So, parallelism is not the answer for a single shot problem that fits nicely in core for a high fidelity simulation.

For the application of interest, the integration domain is a half space that needs to be artificially limited on five sides, where absorbing boundary conditions should be imposed. Thus, the geometry is very simple (a box), although it would be of interest to have topography, i.e., a non-planar surface as the top boundary. For this type of large wave propagation problems it is now routine to use explicit high order methods on uniform meshes, since they are the most efficient and simple ones available.

The acoustic wave equation in three dimensions with a forcing term and absorbing boundary conditions can be written as:

$$w_{tt} = v^2(x, y, z)\triangle w + \mathbf{b}u(t) - 2\epsilon(x, y, z)w_t - \epsilon^2(x, y, z)w,$$

where v is the velocity of propagation and the function ϵ decays rapidly away from the artificial boundaries. First, this equation is discretized in space on a mesh of size $n = nx \times ny \times nz$, and $k \ll n$ snapshots are collected by running one or several High Fidelity simulations. The snapshots are composed of values of the field variables $w(x, y, x, t)$ at points of the spatial mesh for selected times, ordered in a vector with indices running first in the z direction and then in the y and x ones. They are written in this vector form as columns of an $n \times k$ matrix S. An orthogonal basis U for its column space can be generated either by a truncated SVD process or by an adaptive QR algorithm (see Section 9 of [2]). We then assume that the solution can be approximated by $w = Ua(t)$. Replacing in the wave equation and discretizing the Laplacian in space we obtain the reduced order system:

$$a_{tt} = U^T AUa + U^T\mathbf{b}u(t) - 2U^T D(\epsilon)Ua_t.$$

Here the matrix A contains the discrete Laplacian plus the last term $-\epsilon^2 w$. We use an 8th order discretization of the Laplacian [4], which leads to an $n \times n$ sparse, structured matrix, with only 25 nonzero elements per row that is stored in sparse mode. $\mathbf{b}$ is a vector that describes where the forcing function $u(t)$ is applied. For a point source, $\mathbf{b}$ is all zeros except at the source index, where it is equal to 1. Finally $D(\epsilon)$ is a diagonal matrix. Thus, the main pre-processing task, besides obtaining the snapshots and the orthogonal basis, is to calculate $U^T AU$. Either of these procedures requires all the snapshots to be present in fast memory to be competitive and since for the problem sizes we are interested in this is not feasible we will explore in this paper other alternatives:

(a) The reduction to a lower order system can also be attained without orthogonalization, as we explain in Section 2 on Oblique Projection , where a well conditioned basis is created using a progressive adaptive QR algorithm in reduced row space; (b) we can perform MOR using only a limited number of rows of the snapshot matrix (i.e., rows associated with spatial mesh points selected through a $n_r \times n$ matrix C) and (c) we can get snapshots from selected simulations by only integrating part of the total time.

We also consider the use of randomized algorithms [5–7] to speed up the linear algebra steps that are required to project the high fidelity equations into the reduced ones. Because of the availability of fast least squares solvers for large matrices we will also consider applying them to the oblique projection algorithm.

In order to reduce the computing cost even further we also investigate in Section 5 the use of incoherent encoded sources or super-shots. At this point, the winning combination seems to be MOR with oblique projection (i.e., no orthogonalization) and the use of the compacted snapshots with a limited number of rows or Monte Carlo abbreviated multiplication, to generate the coefficients and reduce the size of the least squares problems that arise. We give some preliminary numerical results for the oblique projection method for large problems in 2D and 3D.