Semi-global multi-parameter FWI using public cloud HPC

Henry A. Debens, David Knodel, Fabio Mancini, and Darryl Harris, Woodside Energy Ltd; Mike Warner and Adrian Umpleby, S-Cube*

Summary

This paper explores the potential for hyperscale public cloud high-performance compute (HPC) to enable efficient deployment of a semi-global approach to multi-parameter full-waveform inversion (FWI) over large areas. We introduce several novel aspects to semi-global FWI that improve convergence and suppress crosstalk, while establishing that the algorithm's embarrassingly parallel nature is well suited for public cloud implementation. We describe how various public cloud services can be taken advantage of to reduce the cost of the inversion and provide a reference architecture for the deployment of semi-global FWI to the Amazon Web Services (AWS) platform. Finally, we apply semi-global FWI to raw data from a large-scale legacy surface seismic dataset acquired offshore Australia as part of a re-processing sequence undertaken recently. Our results demonstrate that semi-global FWI can be effectively parallelized across more than one million logical central processing units (CPUs) and is able to recover an anisotropic velocity model in a few hours and in an automated fashion.

Introduction

Building multi-parameter models of the subsurface for seismic imaging typically requires that the seismic data undergo a series of processing steps prior to analysis, which can take several months to execute and often involve numerous subjective decisions to be made by the processing practitioner. Wave equation-based model building techniques, such as FWI, eliminate the need for any processing of the data prior to model building, as these seek to reproduce the data in their raw form. The use of gradient-based methods like FWI for multi-parameter model building introduces additional challenges, for example, how to deal with crosstalk between parameters and how to scale parameters to which the data have differing sensitivities. One approach that can help in mitigating these challenges is to consider a semi-global approach to FWI (Debens et al., 2015). This approach, however, requires significant computational resources when deployed over large areas.

Our paper is organized as follows. We first introduce semi-global FWI and describe how it can overcome several of the challenges associated with multi-parameter FWI. We then discuss how public cloud HPC can provide the necessary resources to make semi-global FWI practical, while access to spare compute capacity can provide the means to make semi-global FWI affordable. Finally, we demonstrate an application of semi-global FWI in recovering an anisotropic acoustic velocity model from more than 3200 km^2 of surface seismic data from the Bonaparte Basin, offshore Australia. This workload was executed exclusively using public cloud HPC during November 2020 and, at its peak, consumed in parallel more than one million logical CPU cores.

Theory

It is well understood that coupled parameters suffer from crosstalk as well as concerns relating to scaling when being estimated via multi-parameter inversion. These challenges can be lessened by considering curvature information in the form of the inverse Hessian operator, as demonstrated by Métivier et al. (2014). The cost of computing the inverse Hessian matrix for a realistic-size problem is however prohibitive, leading to approximations to the Hessian becoming common, such as the L-BFGS algorithm. Semi-global FWI is an alternative approach that seeks to decouple acoustic velocity from the long-wavelength effects of other parameters that are contained within the wave equation being solved, such as anisotropic terms (Debens et al., 2105), acoustic attenuation (da Silva et al., 2017), or acoustic/shear velocity ratio (da Silva et al., 2019). The approach combines an outer global search for the secondary parameter(s) with nested gradient-based iterations that update the velocity model associated with each member of the global population (Figure 1). The outer global search is feasible in 3D (in terms of computational burden) as the secondary parameter can be defined using only a handful of model parameters, which are downscaled prior to modelling. The velocity model, on the other hand, is defined by a fine mesh as per conventional local FWI. This separation according to scale helps in decoupling the parameters being inverted for, while the use of a global search aids in navigating the highly non-linear misfit space and, as a by-product, precludes the need to choose a multi-parameter optimization strategy, e.g., simultaneous versus sequential parameter updates. The optimization begins from a single estimate for velocity and a population of candidate models for the secondary parameter. This population is created at random, either uniform between some bounds (typically required for stability during modelling) or biased by prior information.

From here on we consider only the inversion of velocity and vertical transverse isotropy parameters. Several new aspects of semi-global FWI are introduced in this work:

- Scaling the initial velocity model for each candidate solution using its associated anisotropy model, which acts to "kick start" the inversion. This scaling is done in the time domain and acts to preserve the velocity at wide angles.

10.1190/image2022-3751932.1

- Use of an objective function that is tolerant of large velocity errors, in this case Warner and Guasch's (2016) AWI misfit, to avoid the effects of cycle skipping.
- Use of prior information, such as regional surfaces and interpretations from nearby well bores, to design the sparsely parametrized anisotropy model. This avoids the use of a sparse regular mesh, which can lead to models that appear non-geological.
- Use of formation tops to penalize depth error during the inversion via a simple 1D time-to-depth conversion of legacy time-domain post-stack data. This assumes that the area of interest has been penetrated previously and that these legacy data are available.

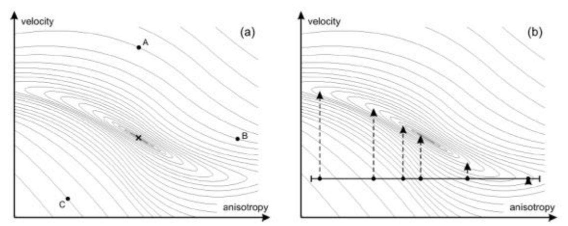

Figure 1: Comparison between (a) local and (b) semi-global FWI for a schematic two-parameter problem (Debens et al., 2015). Points A, B, and C in (a) indicate three different initial models, each with a different optimal optimization strategy.

The assessment of each member of the global population can be considered an independent FWI problem and, as such, can be solved for in parallel as sub-jobs. Semi-global FWI can therefore be considered an embarrassingly parallel problem on both the model and shot levels. This enables efficient parallelization across large amounts of compute. In our implementation (Figure 2), each iteration of the global outer loop is parallelized by (i) distributing the FWI sub-jobs across different data centers within different public cloud regions, (ii) distributing the problem-specific data for each FWI sub-job across worker processes, and (iii) using multi-threading to parallelize wave equation solves across CPU cores. The global outer loops, for which we use quantum particle swarm optimization (Sun et al., 2004), are by themselves computationally inexpensive. Dedicated services offered by public cloud providers, such as Amazon Simple Storage Service (Amazon S3), are taken advantage of to ensure low latency and high throughput. Key to keeping the cost of the inversion comparable with on-premises HPC is the use of spare compute capacity services, such as Amazon EC2 Spot. Services of this kind allow for users to pay a heavily discounted rate to access resources that would otherwise be idle, with the caveat that, should another user wish to pay full price, the resource is reclaimed following a short notice period. As such, it is important to ensure that any algorithm reliant on spare compute capacity is made robust against disruptions. In our case, this is achieved via the use of a bespoke fault-tolerant message-passing library referred to as RIPSSCI.

Field dataset

We applied the approach described above to surface seismic data acquired during a 2001 towed-streamer survey as part of a modern re-processing sequence. The dataset had not previously been processed in depth and, as such, limited information relating to anisotropy was available prior to re-processing. Despite being acquired for development purposes, the dataset is extensive in terms of surface area and therefore poses a significant challenge for any stochastic or semi-stochastic approach to seismic inversion. The overburden is heavily faulted and contains laterally extensive intervals of carbonate-dominated media. Numerous present-day reefs at the seafloor create imaging challenges beneath a shallow-water platform that constitutes roughly half of the survey.

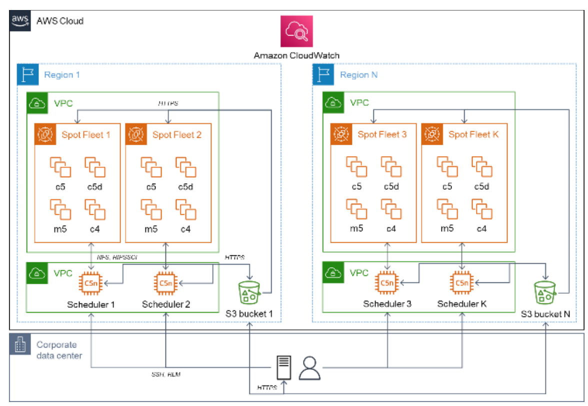

Figure 2: Reference AWS architecture for deployment of semi-global FWI.

Prior to performing semi-global FWI, isotropic local FWI was used to update the long-wavenumber component of the velocity model at the bandwidth that was to be used for semi-global FWI (in this case, up to 5 Hz). Regional interpretations and geological information from well penetrations within the survey area were used to divide the subsurface up both vertically and laterally into regions where the anisotropic properties might be expected to vary, e.g., between clay- and carbonate-rich intervals. The dataset was heavily decimated in the shot, channel, and time dimensions prior to semi-global FWI.

Results and discussion

The epsilon component of the anisotropy model recovered can be seen in Figure 4c. It is smoothly varying by design and is defined by only six parameters, which were optimized for during semi-global FWI. A seventh parameter controlled

10.1190/image2022-3751932.1

the ratio between epsilon and delta, which was treated as uniform across the entire domain. The anisotropy model was free to vary laterally as a function of water depth, as the survey covers a tectonically active region and, as such, it was suggested that the stress regime may change with proximity to a significant nearby trench. In fact, the inversion returned a model with minimal lateral variation. In this example, local multi-parameter FWI for subsequent recovery of higher wavenumber anisotropy was not applied, although Debens et al. (2015) show that this can be of benefit. In such an event, care must be taken over the choice of parameterization; reasonable options include those introduced by da Silva et al. (2014) and Debens et al. (2015).

Deployment of the algorithm involved the parallel execution of 110 FWI sub-jobs across eleven AWS Availability Zones spanning three AWS Regions. Each outer loop of the algorithm took roughly 90 minutes to complete (Figure 3) and, importantly, the algorithm was able to complete its workload despite regular interruptions due to the reallocation of Amazon EC2 Spot resources. The rate of instance turnover during this workload was observed to be approximately 2% per hour. Particularly encouraging was that the total volume of resources requested was provided within roughly ten minutes of the requests being submitted.

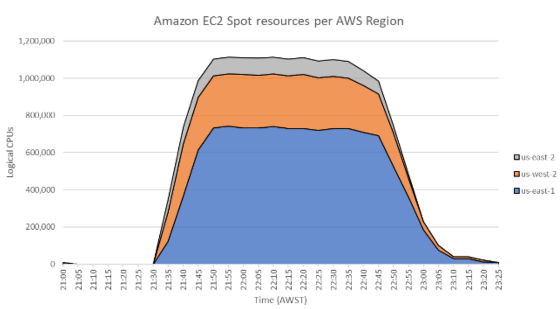

Figure 3: Amazon EC2 Spot resources consumed as a function of time and public cloud region for an outer loop of semi-global FWI.

The anisotropic velocity model recovered was then incorporated into a hybrid workflow containing FWI and reflection travel-time tomography, with the resultant model used during Kirchhoff PreSDM. Each of these subsequent processes were also performed using Amazon EC2 Spot resources in North America. The legacy PreSTM velocity model and image can be seen in Figure 4a, with their re-processed equivalents shown in Figure 4b. The strongly anisotropic interval correlates with a finely layered clay-rich formation, which acts to reduce the depth mis-tie at top reservoir. The introduction of FWI to the model building sequence allowed for the kinematic effects of the seafloor reefs to be captured, although a heuristic approach to water density was required to adequately characterize their reflectivity while using an acoustic wave equation.

The dataset contains a narrow range of offsets, with the longest reaching little more than 4.9 km. As with the travel-time component of local FWI, our approach is driven in principle by turning energy and, as such, question marks exist over the validity of the anisotropy model updates in the deepest parts of the model. Furthermore, only one of the well bores penetrated beyond the target interval, resulting in reduced depth control in the deeper model. Updates in this part of the model will therefore have been reliant on reflection moveout, and FWI using surface seismic data is notoriously insensitive to variations in delta (Alkhalifah and Plessix, 2014). Additional information, such as common-image-point gather flatness, is needed to better constrain the deeper updates during the inversion process.

Whilst the implementation presented here takes advantage of several public cloud services, such as Amazon S3, Amazon Virtual Private Cloud (Amazon VPC), and AWS Spot Fleet, it still closely resembles an infrastructure-as-a-service (IaaS) style of implementation. Improvements in efficiency, cost, and resilience could be achieved by moving towards a platform-as-a-service (PaaS) style of implementation, such as described by Witte et al. (2019). Often in doing so, however, one disadvantage will be a reduction in agnosticism towards cloud provider.

Conclusions

We present here an application of semi-global FWI to raw data acquired during a large-scale marine seismic survey from offshore Australia. Several improvements to the semi-global FWI approach are introduced and a robust approach to deployment using public cloud HPC is presented. We demonstrate this approach as being effective when parallelized across more than one million logical CPUs spread across several regions within the AWS platform.

Further work will seek to exploit information relating to uncertainty that is contained within the population of candidate solutions. Most semi-global FWI projects to date have generated either a single best solution from the population or an average of the best solutions.

Acknowledgements

The authors would like to thank Woodside and Woodside's joint venture partners Timor Gap Greater Sunrise RL, Unipessoal, Lda and Osaka Gas Australia Pty Ltd for their permission to present this work. They also wish to express gratitude towards AWS (in particular, Alex Heggart) and the Fullwave project at Imperial College London for their support of this work, as well as Ka-Hing Cheung for the goofys file system. Any views or opinions expressed in this paper are solely those of the authors and do not necessarily represent those of Woodside or S-Cube.

10.1190/image2022-3751932.1

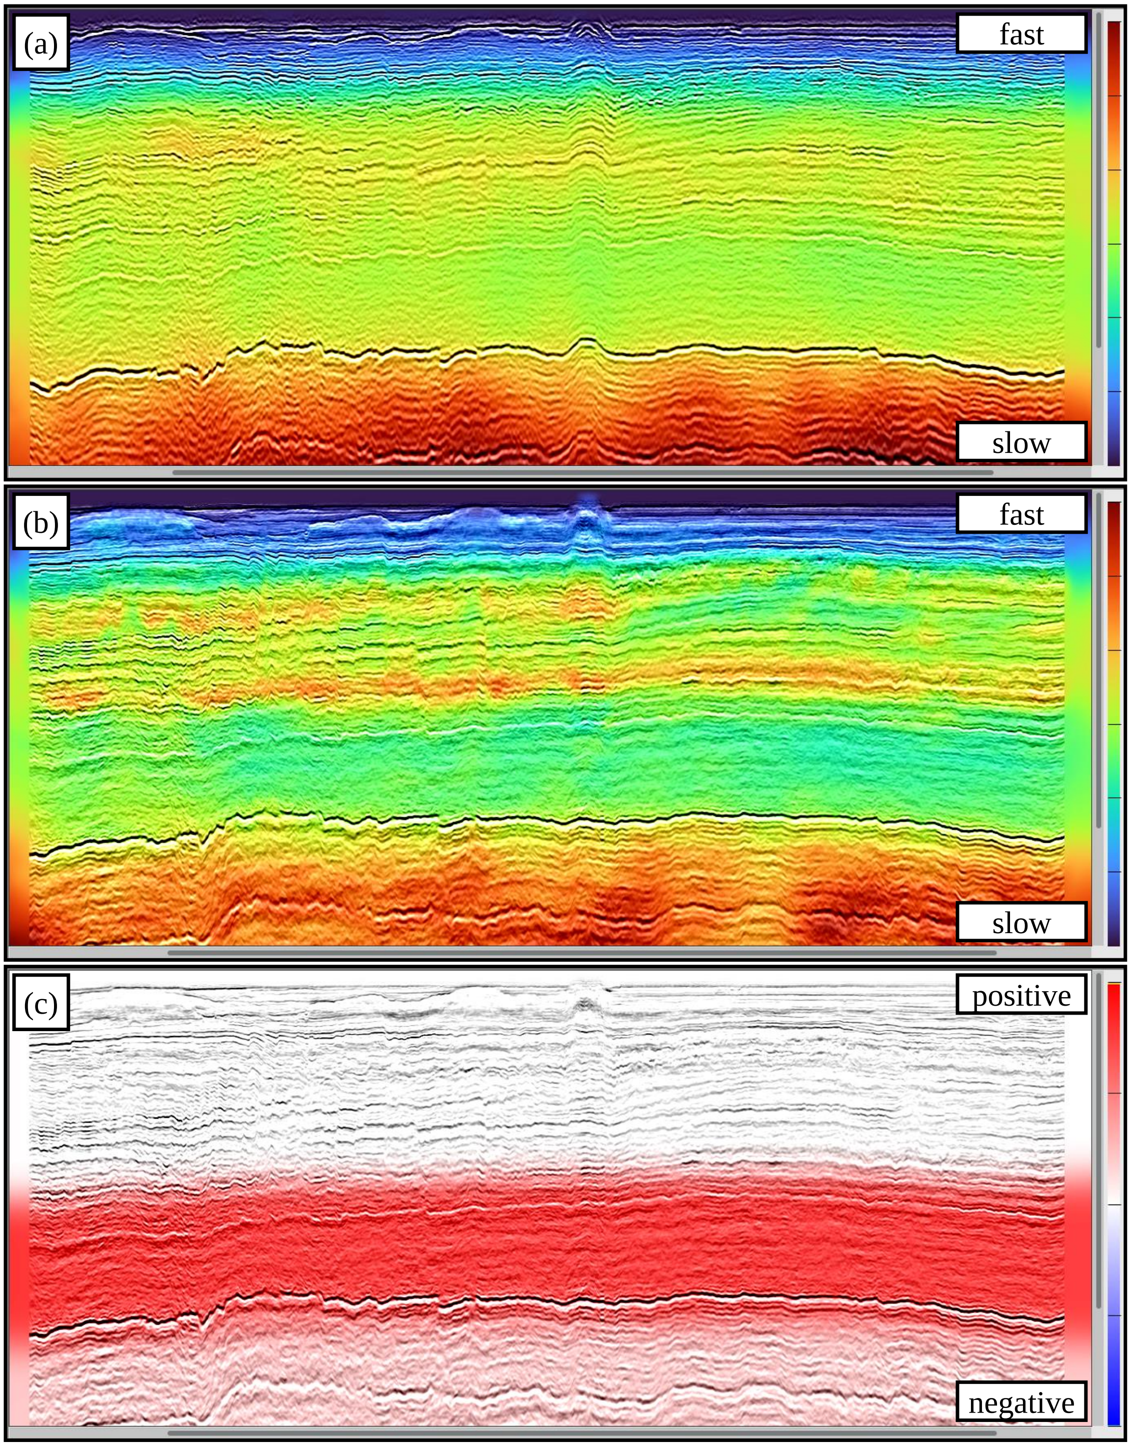

Figure 4: Crossline section through the shallow-water platform. (a) Legacy pre-stack time migration (PreSTM) velocity model and full stack, stretched to depth; (b) re-processed pre-stack depth migration (PreSDM) vertical velocity model and full stack; and (c) semi-global FWI epsilon model and re-processed full stack.

10.1190/image2022-3751932.1

REFERENCES

Alkhalifah, T., and R.-É. Plessix, 2014, A recipe for practical full-waveform inversion in anisotropic media: An analytical parameter resolution study: Geophysics, **79**, no. 3, R91–R101, doi: https://doi.org/10.1190/geo2013-0366.1.

da Silva, N. V., A. Ratcliffe, G. Conroy, V. Vinje, and G. Body, 2014, A new parameterization for anisotropy update in full waveform inversion: 84th Annual International Meeting, SEG, Expanded Abstracts, 1050–1055, doi: https://doi.org/10.1190/segam2014-0664.1.

da Silva, N. V., G. Yao, Ò. Calderón Agudo, G. Stronge, and M. Warner, 2019, 3D elastic semi-global waveform inversion—Estimation of VP to VS ratio: 89th Annual International Meeting, SEG, Expanded Abstracts, 1440–1444, doi: https://doi.org/10.1190/segam2019-3216359.1.

da Silva, N. V., G. Yao, M. Warner, A. Umpleby, and H. Debens, 2017, Global visco-acoustic full waveform inversion: 79th Conference and Exhibition, EAGE, Extended Abstracts, Th A3 02, doi: https://doi.org/10.3997/2214-4609.201700504.

Debens, H. A., M. Warner, A. Umpleby, and N. V. da Silva, 2015, Global anisotropic 3D FWI: 85th Annual International Meeting, SEG, Expanded Abstracts, 1193–1197, doi: https://doi.org/10.1190/segam2015-5921944.1.

Métivier, L., R. Brossier, S. Operto, and J. Virieux, 2014, Multi-parameter FWI—An illustration of the Hessian operator role for mitigating trade-off between parameter classes: 6th Saint Petersburg International Conference and Exhibition, EAGE, Conference Proceedings, Tu BC 01, doi: https://doi.org/10.3997/2214-4609.20140196.

Sun, J., B. Feng, and W. Xu, 2004, Particle swarm optimization with particles having quantum behavior: Proceedings of the 2004 Congress on Evolutionary Computation, 325–331.

Warner, M., and L. Guasch, 2016, Adaptive waveform inversion: Theory: Geophysics, **81**, no. 6, R429–R445, doi: https://doi.org/10.1190/geo2015-0387.1.

Witte, P. A., M. Louboutin, H. Modzelewski, C. Jones, J. Selvage, and F. J. Herrmann, 2019, Event-driven workflows for large-scale seismic imaging in the cloud: 89th Annual International Meeting, SEG, Expanded Abstracts, 3984–3988, doi: https://doi.org/10.1190/segam2019-3215069.1.