The role of High-performance computing and seismic imaging and interpretation

Biondo Biondi (Stanford University), Bob Clapp (Stanford University) and Alejandro Valenciano (Stanford University)

SUMMARY

Progresses in seismic-imaging technology are driven by advancements in data acquisition and high-performance computing. Wide-azimuth acquisition geometries of both marine and land data are dramatically changing the data we image. The commoditization of multi-core processors and the availability of ultra-fast hardware accelerators (FPGAs, GPUs, Cells, …) will change the way that we image and interpret those new data sets. These hardware improvements will enable the application of imaging operators that are more accurate in both the modeling of the underlying physical phenomena (e.g. wave propagation vs. ray-tracing) and the approximation of the actual inversion of the recorded data. The future availability of workstation with multi-core CPUs will enable the exploitation of expensive numerical algorithms to support interpretation. This should lead to dramatic improvements in the structural and stratigraphic interpretation in areas where the complexity of the velocity model requires a tight loop between interpretation and processing

Introduction

Progresses in seismic-imaging technology are driven by advancements in data acquisition and high-performance computing. Wide-azimuth acquisition geometries of both marine and land data are dramatically changing the data we image. The commoditization of multi-core processors and the availability of ultra-fast hardware accelerators (FPGAs, GPUs, Cells, …) will change the way that we image and interpret those new data sets. These hardware improvements will enable the application of imaging operators that are more accurate in both the modeling of the underlying physical phenomena (e.g. one-way vs. two-way wave, acoustic vs. elastic wave propagation) and the approximation of the actual inversion of the recorded data. The future availability of workstation with multi-core CPUs will enable the exploitation of expensive numerical algorithms to support interpretation. This should lead to dramatic improvements in the structural and stratigraphic interpretation in areas where the complexity of the velocity model requires a tight loop between interpretation and processing.

Drivers for computationally-intensive seismic imaging and interpretation algorithms

Historically, the rate of decrease in computational cost has outpaced the rate decrease in data-acquisition cost, as is sketched in Figure 1. The difference in slopes between the green and the yellow lines shown in Figure 1 has created opportunities for the development and application of increasingly more accurate imaging algorithms. Modern imaging methods produce more accurate images than previous methods because they honor more closely the full complexities of wave-propagation in a heterogeneous subsurface. However, they also require substantially more flops to image each recorded sample in the data. The switch from post-stack imaging to pres-stack imaging and the more recent transition from ray-based migration to wavefield-continuation migration are good examples of past important steps forward.

In the near future, it is likely that the widening gaps between the two lines in Figure 1 will lead us to use more accurate wave-propagation methods (two-way wave equation and one-way propagation in non-Cartesian coordinates.) This development in seismic-imaging technology is very well matched by progress in computational hardware. Multi-core CPUs make it economically advantageous to deploy large-memory computers that increase the efficiency of time-domain two-way wave propagation algorithms. These wave-propagation algorithms have regular patterns of both computations and data access and also have high ratio between computations and I/Os. These characteristics well match the requirements of modern specialized hardware accelerators such as FPGAs and GPUs. These hardware accelerators are likely to be widely used for the "routine" imaging by computationally intensive migration methods, such as reverse-time migration.

Another trend that is likely to continue in the near future is the increasing application of "inversion" methods to image the data, at the expense of more conventional "processing" methods. The application of inversion methods can often yield better velocity models and images in difficult areas; however, they are in many ways less robust than conventional processing methods and thus they greatly benefit from being "guided" by *a priori* knowledge of the geology provided by an interpreter. This trend towards more extensive use of interpretative inversion methods will also drive the blurring of boundaries between imaging and interpretation. Interpretation methods will rely more heavily on computationally intensive algorithms and will enable interpreters to drive inversion-based imaging algorithms by testing geologic hypothesis in quasi-real time. The increase in computationally driven interpretation is justified by the difference in slopes between the yellow line and the orange line in Figure 1. Whereas the unit cost of a geoscientist is approximately constant (or slightly increasing) over time, the amount of data to be interpreted continues to increase because of more cost-efficient data acquisition. This widening gap creates the economic incentives to use more computational power to increase the efficiency of interpretation.

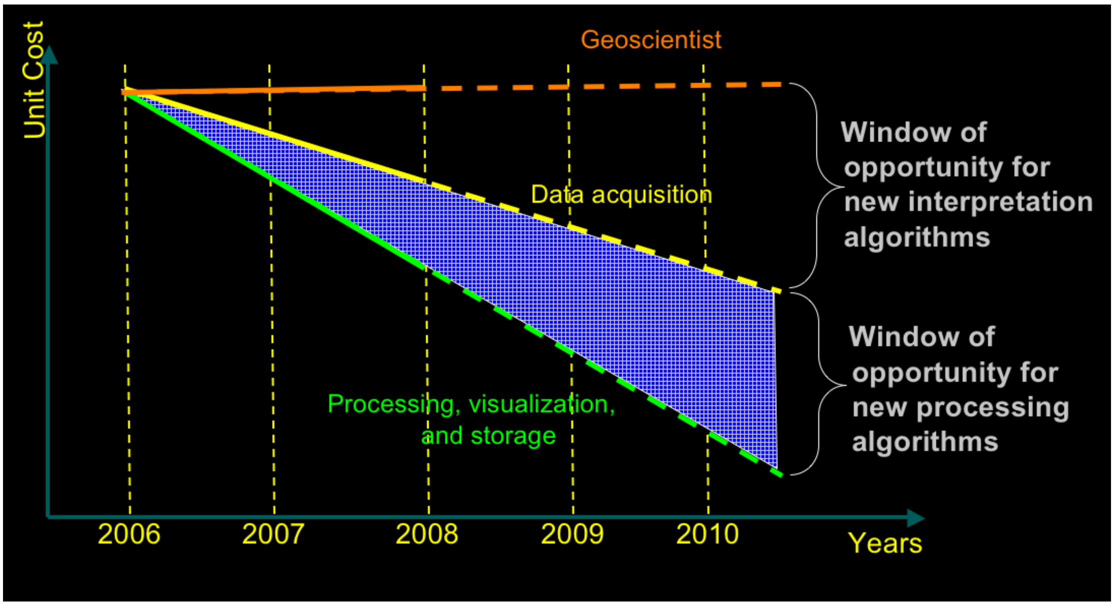

Figure 1 – Conceptual graph of the unit-cost trends for acquiring processing, and interpreting seismic data.

An example of computationally intensive algorithm that can drastically reduce interpretation efforts is the automatic salt-boundaries tracking proposed by Lomask et al. (2007). Examples of real-time driving of inversion algorithms by an interpreter are the tight loop between velocity updating and migration that is enabled by target-oriented migration methods (as proposed by a Ross Hill with Chevron in a recent SEG Recent Advances and the Road Ahead session) and the target-oriented reflectivity inversion presented by Valenciano et al. (2006) and Valenciano (2008), which is described in more detailed in the next section. Both of these interpretative inversions are based on a computational model where large-scale computations are performed ahead of time on a computer cluster and the partial results are organized in a way that enables the quasi real-time interaction with an interpreter.

Target-oriented reflectivity estimation by inversion

Reflectors in the subsurface are commonly imaged by seismic migration. When the overburden is complex and the reflectors are unevenly or insufficiently illuminated, simple migration may not provide an optimal image. In these cases, the numerical inversion of the wave-modeling operator can actually produce better image than migration. However, the inversion problem is poorly constrained because of the band-limited nature of seismic data and the uneven illumination. In order to produce good results, or even to prevent the inversion from diverging, the inversion algorithm needs to be constrained by adding a “regularization” term to the objective function. Best results are obtained when an interpreter chooses the constraints based on the geology of the imaging target. Even the simple choice of the relative weights between the regularization and data-fitting terms would be best done during interpretation by analyzing results obtained with different values of the weights. Fortunately, this kind of interpretative inversion is made possible today by the availability of modern multi-core workstations interconnected by fast networks with large computer clusters.

In the case of target-oriented inversion, an approximation of the Hessian of the objective function can be pre-computed by performing a large number of wave-propagations and correlations on a cluster (Valenciano, 2008). The results are then stored on a disk array, ready for the interpretative step of the procedure. This process requires a large amount of computational resources (hundred of thousand of CPU hours), but once the Hessian is computed the target-oriented inversion can be performed in a matter of minutes on a multi-processor computer, and thus can be interactively driven by an interpreter.

Figures 2-5 illustrate this idea by a simple example. Figure 2 shows the results of target-oriented shot-profile migration computed by employing an accurate one-way propagator. Shadow zones and migration artifacts associated with poor illumination are clearly visible in the image below the salt body; in the image, the salt body was masked and thus corresponds to the uniformly grey area. Some vertical artifacts associated with converted waves area also visible at the right edge of the image. Figures 3-5 show the images obtained by target-oriented regularized inversion with different values for the weight of the regularization term (ε). The regularization is the weakest for Figure 3 and the strongest for Figure 5. The shadow zone is best infilled in Figure 3. From this image the interpretation of many reflectors in the shadow zone is substantially more reliable than from the migrated image (Figure 2). On the other hand, in Figure 3 there are some remnants of truncation smiles in the shadow zone and strong vertical artifacts associated with converted waves. The truncation smiles are almost completely attenuated in Figure 4, and the vertical artifacts are eliminated from Figure 5. If all these three images where available during interpretation, the final results would be much more likely to be accurate than if only one were available. Furthermore, if needed, the interpreter could easily create another inversion results in only few minutes.

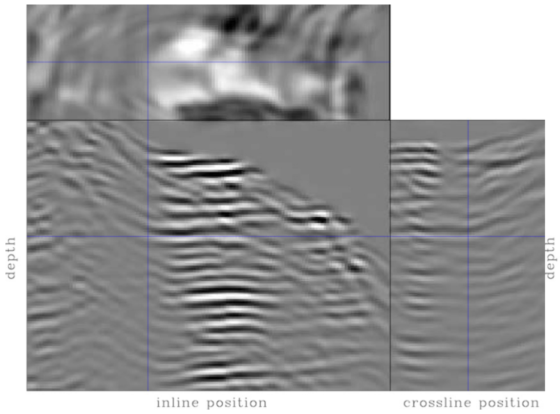

Figure 2 –Target-oriented shot-profile migration. Image from Valenciano (2008).

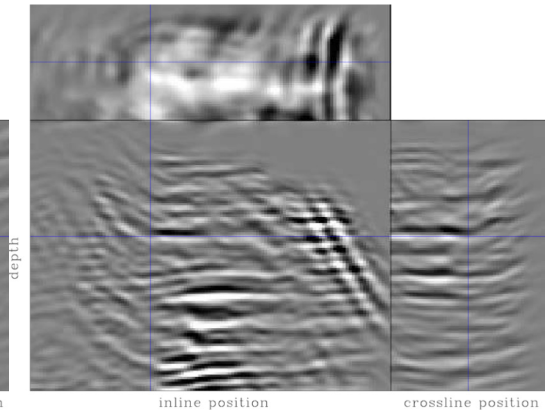

Figure 3 – Target-oriented wave-equation inversion with weak regularization (ε=0.01). Image from Valenciano (2008).

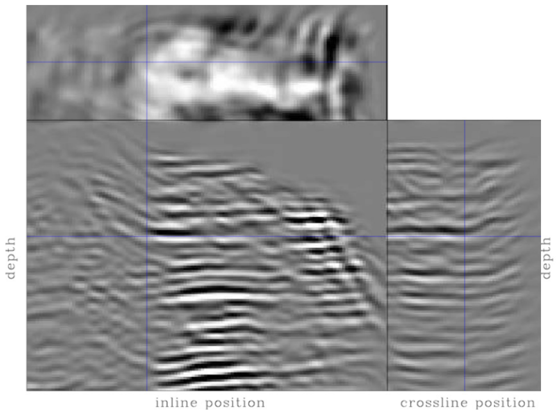

Figure 4 – Target-oriented wave-equation inversion with medium-strength regularization (ε=0.5). Image from Valenciano (2008).

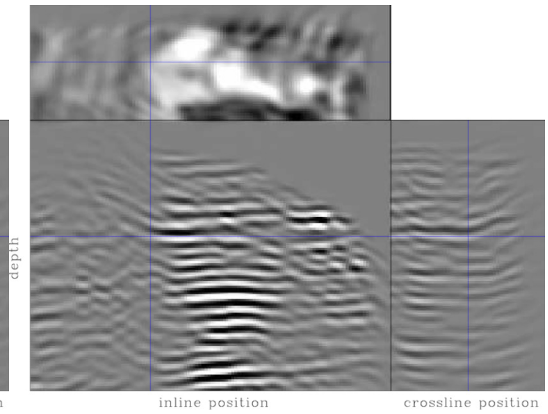

Figure 5 – Target-oriented wave-equation inversion with strong regularization (ε=10). Image from Valenciano (2008).

Acknowledgments

We would like to thank BP and ExxonMobil for providing the Gulf of Mexico data shown in to SEP.

References

Lomask, J., R.G. Clapp, and B. Biondi, 2007. Application of image segmentation to tracking 3D salt boundaries, Geophysics, v. 72, pp. P47-P56.

Valenciano, 2008. Imaging by wave-equation inversion, Ph.D. thesis, Stanford, University.

Valenciano, A., A., B. Biondi, and A. Guitton, 2006, Target-oriented wave-equation inversion: Geophysics, v. 71, A35-A38.