3-D depth migration: parallel processing and migration movies

Dmitri Gavrilov[1], Laurence R. Lines, Henry C. Bland and Tony Kocurko[2]

ABSTRACT

The advent of 3-D depth migration has brought about many challenges and opportunities in the applications of computer science to seismic imaging. We shall examine two fields of computer science, which will impact the 3-D depth migration technology: parallel processing and visualization. Parallel computational applications will make 3-D migration calculations tractable while visualization will allow us to examine the input models, snapshots of computations, and the output images.

In our depth migrations we primarily examine two methods: Kirchhoff migration and reverse-time migration. Zhu and Lines (1998) have examined the computational advantages and disadvantages of these methods. The Kirchhoff algorithm is generally the faster of the two, but is slightly less accurate than reverse-time migration when velocity models are accurately defined since the Kirchhoff is based on ray theory approximations to the wave equation.

PARALLELISM

The advantages of parallel processors are shown by the cartoon in Figure 1 (taken from Kalantzis, 1994) which illustrates scalar, vector, parallel, and vector/parallel processes for putting out a fire with buckets of water. In scalar mode, we have a single processor performing the operation (one bucket carrier). In vector mode, we perform the task by a series of simultaneous operations on a vector or array of elements (the bucket brigade). In parallel, mode we have a number of processors performing similar tasks (several bucket carriers running in parallel). The vector/parallel operation has several lines of bucket brigades operating at once. Obviously, the latter operation would perform the task fastest.

The beauty of the Kirchhoff depth migration lies in the fact that individual traces or shot records can be selectively migrated to depth and summed. The aplanatic surfaces for each trace can be generated by computing seismic wavefronts emanating from the source and summing these wavefronts to those wavefronts generated for each receiver. Seismic amplitudes are distributed over the aplanatic surfaces for each trace. One should also account for obliquity factors in the amplitude, as described by Scales (1995). These depth images are summed together for all traces. Since the same source wavefronts are used in the computation of all aplanatic surfaces in a shot gather, it may be convenient to give each processor a shot gather. The computation of shot migrations can be done on individual processors and then summed together to form a depth migrated image. This computational process has been termed "embarrassingly parallel" by Tony Kocurko since the migration load can be

[1] Multimedia Advanced Computational Infrastructure Project
[2] Memorial University

distributed over an array of processors (Kocurko and Lines, 1998). The speedup of processing is nearly a linear function of the number of processors.

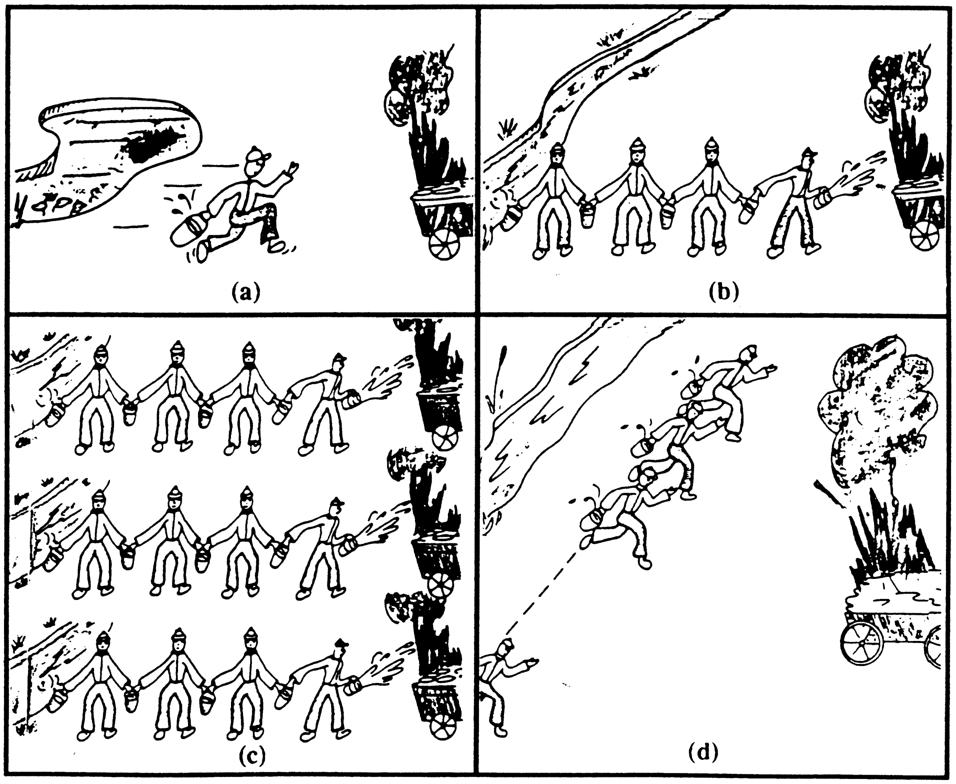

Figure 1. Cartoons illustrating the different processing architectures as they evolved with time: (a) scalar processing; (b) vector processing; (c) vector/parallel processing; and (d) parallel processing (modified after Moriarty, 1987)

Reverse-time depth migration has the advantage of being a method which satisfies the complete wave equation and is therefore able to image arbitrarily steep reflecting beds. Good examples of 3-D reverse-time migration applications have been given by Mufti et al. (1996) for Gulf Coast salt dome imaging and by Kelly (1998) for imaging of faults in Hibernia field. One problem with the Hibernia runs was the lengthy computation time. Runs on seismic workstations often took days to complete the reverse-time migrations. This problem was solved by Dmitri Gavrilov using parallelism.

The reverse-time algorithm uses finite-difference solutions to the wave equation over a grid of points. In our example, the grid size was 150 by 121 by 350 cells with cells being 25m in the x and y directions and 10m in the z direction. The 8910 seismic traces used in the migration contained 1430 samples at 2ms sample interval. (This is a relatively small survey by 3-D standards but served as interesting data set for testing algorithms.)

In reverse-time migration, the seismic traces provide time-varying boundary values as sources for wave propagation back into the subsurface to the point of reflection. Reverse-time wave propagation is computed by using finite-difference algorithms. The trick to invoking parallelism comes from dividing the subsurface grid of subsurface blocks. In each of these blocks, finite-difference calculations can be carried out. Then the task is to allow each block to communicate in finite-difference calculations to each other block at the boundaries. Once this communication is established, the wave equation calculations can be done on N processors for N blocks.

The parallel processing was implemented by Dmitri Gavrilov who created an implementation termed RTM3D (Reverse-Time Migration in 3-D). The RTM3D application was implemented on a cluster of 30 Dec Alphas connected via a Myrinet switch. The Message Passing Interface (MPI) was used to develop a distributed parallel implementation. The use of MPI allows one to port this application to almost any parallel machine (either distributed or shared memory), since MPI implementations are readily available for many parallel computers. The application can be compiled without any MPI calls.

The application uses a 3-D finite-difference scheme to solve the wave equation. Finite-difference schemes are naturally implemented in an MPI environment. The 3-D grid is split into slices, which are distributed among the available processors. After each time step, a processor must exchange the values on the boundaries of their slices since these values are needed for computations on the next time step by neighbouring processors. In such a scheme there is no master process – all processes are working simultaneously on their respective slices of the data array and no external coordination is required. In the application on the provided data set, almost linear speedup can be seen in Figure 2.

In diminishing compute time for days to less than an hour, Gavrilov divided the computations among 13 Alpha processors operating at 500MHz in the alpha cluster at the University of Calgary. Run times were impressively reduced from days to less than 1 hour.

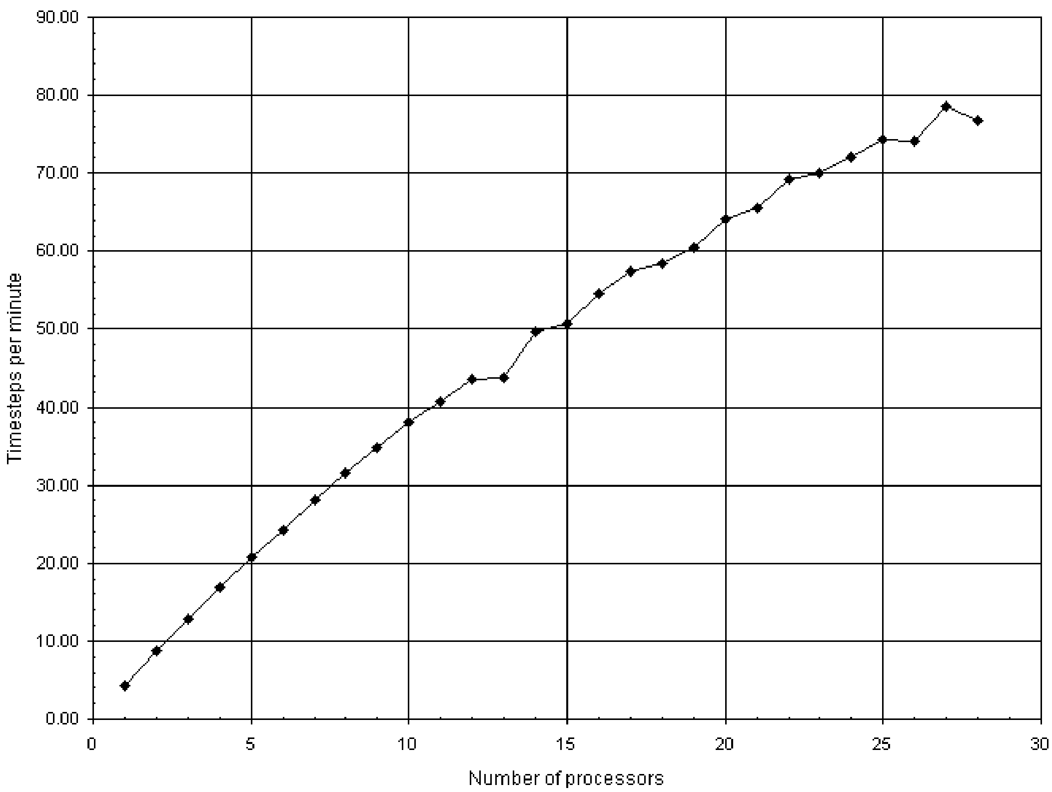

Figure 2. As the number of processors increases, more time-steps can be processed each minute. The almost-linear curve is indicative of optimal parallel processing. A grid of 150 x 121 x 350 elements was processed in each time step.

MIGRATION MOVIES

One of the interesting aspects of wave equation computations is that we can produce a movie of the wave propagation by saving the various snapshots of time steps in the computations. Since reverse-time depth migration uses such calculations, we can develop a “migration movie” which exhibits backward wave propagation over the time duration of the trace to the point of reflection. Bording and Nickerson (1995) showed this type of movie as a means of evaluating migrations.

A movie was created of the 3-D reverse-time migration of the Hibernia data. It illustrates the progression of the migration along a west-east line through the 3-D data volume. A small program was written to convert a series of snap-shots taken at each timestep into digital still image files in PGM (Portable Grey-Map) format. This sequence of images was converted into a movie and stored in the MPEG digital movie format. Spatial and temporal compression are used to greatly reduce the size of the resulting file. It was convenient to use MPEG because of the high availability of MPEG viewing software for personal computers. To create the MPEG file, the “Berkeley MPEG Encoder” program was used (Gong and Rowe, 1994). The encoder reads a series of still frames, and generates an output MPEG file.

Contents

Figure 3 shows an example of a reverse-time migration for a west-east line, as viewed from the north. The Murre fault is a basin-bounding fault which is present in the lower right hand portion of the figure. The reverse-time movie shows the seismic definition of the Murre fault, as waves are backward propagated to the point of reflection in a 3-D volume.

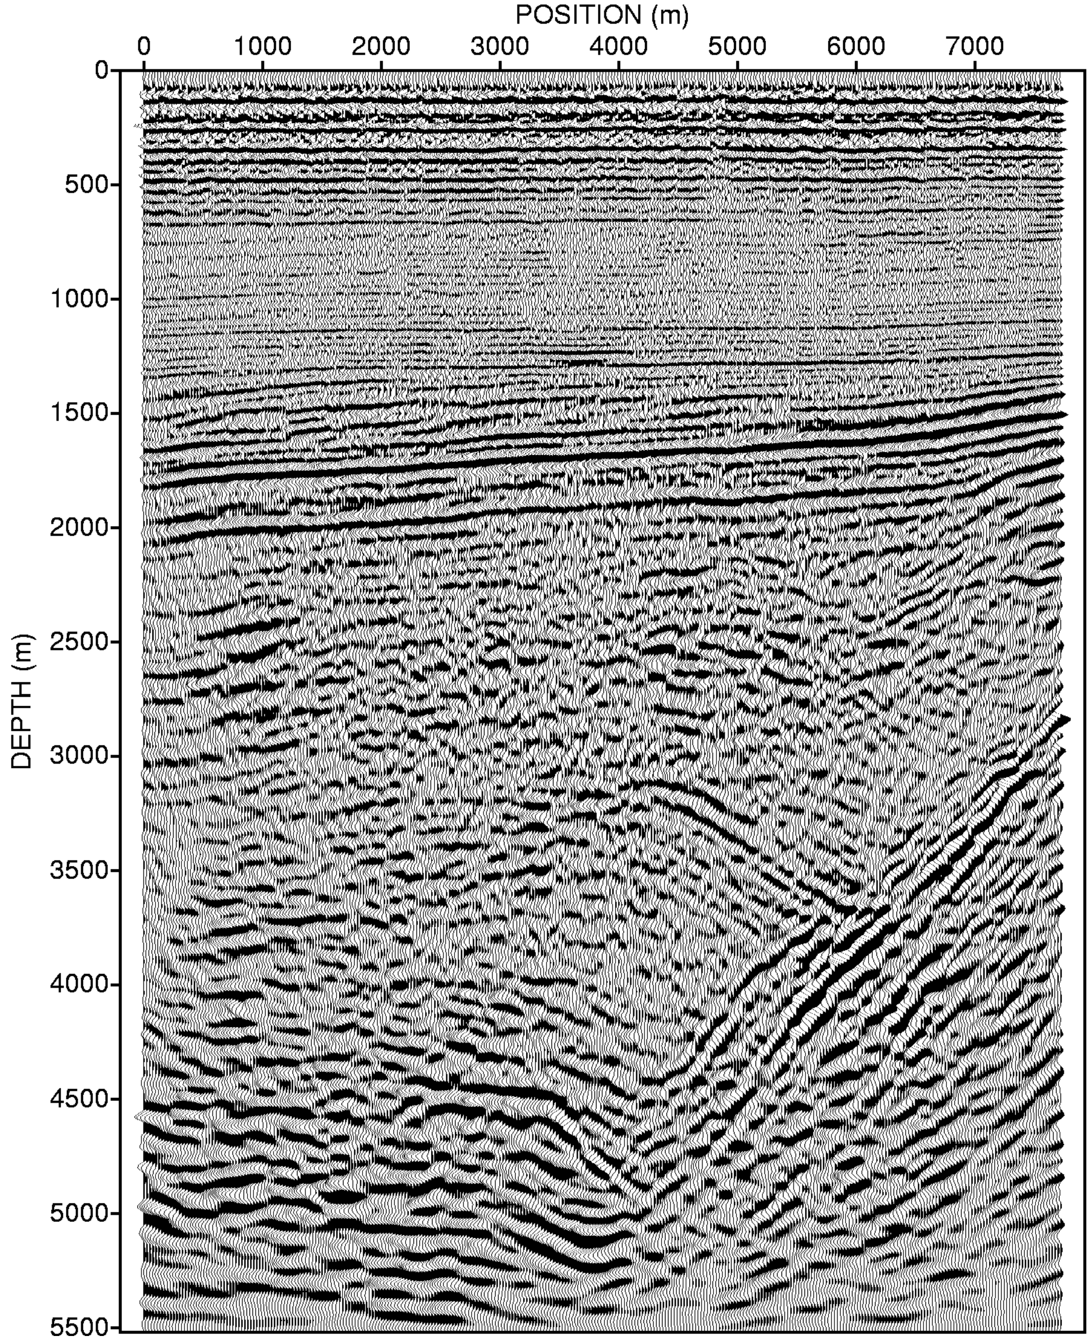

Figure 3. Reverse-time migrated seismic section from a marine survey of the Hibernia field (Newfoundland, Canada). This section is located at crossline 210 of the 3-D survey.

The migration movie includes 1430 frames (one per time step) and lasts about 47 seconds. A sparse sampling of frames from the movie are shown in Figure 4. The movie's final frame represents the same migrated section shown in Figure 3.

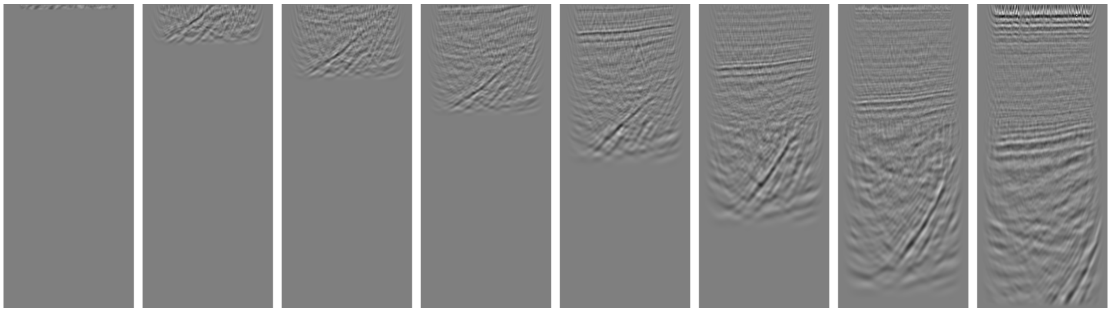

Figure 4. The migration progresses with each successive time step. Shown are frames from the migration movie sampled every 200 time steps (from left to right). The actual movie includes one frame for each time steps.

CONCLUSIONS

Three-dimensional prestack depth migration poses new computational challenges for hardware and software configurations alike. In our research thus far, we have examined parallelism as a computational engine and seismic movies as visualization tools. Parallelism can be invoked through "master-slave" processing for Kirchhoff migration (Kocurko and Lines, 1998) or through Message Passing Interfaces as invoked in the parallelization of reverse-time migration. In either case, significant improvements in run-time can be achieved. Computationally intensive seismic processing algorithms are often good candidates for parallel execution. By harnessing all available processors, either in a single multi-processor sytem, or in a network of several uniprocessor systems, we can obtain significant speed increases. These speed increases enable us to use better and more powerful computation operators when procesing seismic data.

ACKNOWLEDGEMENTS

The researchers thank Gary Taylor and Hibernia Management and Development Corporation for the chance to migrate their data. Also, we thank Irene Kelly for doing the first migrations for her M.Sc. on workstations despite the fact that computer runs took days to complete.

REFERENCES

Bording, R.P., and Nickerson, W., 1995, Process and visualization: The migration movie: The Leading Edge, v.15, 38-42.

Gong, K.L., and Rowe, L.A., 1994, Parallel MPEG-1 Video Encoding, Proceedings of the 1994 Picture Coding Symposium, Sacramento, Calif.

Kalantzis, F., 1994, Imaging of reflection seismic and radar wavefields: Monitoring of steam-heated oil reservoirs and characterization of nuclear waste repositories: Ph.D. thesis, University of Alberta, Edmonton, Alberta.

Kelly, I.G., 1998, Modeling and migration of Hibernia seismic data: M.Sc. thesis, Memorial University of Newfoundland, St. John's, Newfoundland.

Kocurko, T., and Lines, L.R., 1998, An embarrassingly balanced parallel aplanatic migration algorithm: paper presented at the Geo-Triad meeting in Calgary, Alberta, June 16-19, 1998.

Mufti, I.R., Pita, J.A., and Huntley, R.W., 1996, Finite-difference depth migration of exploration scale 3-D seismic data: Geophysics, 61, 776-794.

Scales, J.A., 1995, Theory of Seismic Imaging: Samizdat Press (web publication) and Springer Publications.

Zhu, J., and Lines, L.R., 1998, Comparison of Kirchhoff and reverse-time migration methods with applications to prestack depth imaging of complex structures: Geophysics, 63, 1166-1176.

Contents