

Increasing the parallelism in common azimuth migration with overlap domain decomposition

Sean Crawley, 3DGeo Inc.

Summary

Common azimuth migration (COMAZ) is an efficient algorithm for generating wave equation quality images of narrow azimuth (in particular towed streamer) seismic data. The algorithm is naturally parallel over input frequency. We increase the parallelism of the algorithm using overlap domain decomposition to divide the spatial axes of the data.

Introduction

There are many ways to produce a depth image. Starting from narrow azimuth marine data, common azimuth migration (Biondi and Palacharla, 1996) is a cost-effective way to get a good quality, wave equation result.

Common azimuth migration is naturally parallel over the input frequency, but obvious parallelism ends there. In contrast to many other migration algorithms, the entire spatial extent of the input data migrates at once, usually with some mixed-domain operator. This can lead to a couple of complications. For a simple implementation, it caps the number of cores that can efficiently work on a single input frequency at the number that can share memory, commonly 4 or 8 on a current Linux cluster node. For a large, high frequency input data slice, more cpu power comes in handy.

It also means that (again with a nod to efficiency and simplicity of implementation) memory requirements grow with survey size. Depending on the details of the downward continuation algorithm, storing the required frequency and image slices for a large dataset can overtake the memory on a single cluster node.

Calandra et al (2007) presented an adaptation of common azimuth migration to the IBM Blue Gene supercomputer. They get increased parallelism with a Fourier finite difference (FFD) operator using a distributed 3D FFT to run on an impressive 16k cpus.

An alternative to a distributed mixed-domain operator, is domain decomposition. Domain decomposition methods can provide significant advantages for simulating wave propagation, in making them suitable for parallel computation. Many approaches use non-overlapping subdomains (Carcione, 1991; Faccioli et al 1996), with care taken at the boundaries. Overlap domain decomposition (ODD) methods have been applied to Fourier modeling methods by Liao and McMechan (1993) and Fan et al (1997).

We present an ODD method where the input frequencies are divided into overlapping subdomains, each of which migrates independently between occasional (in depth) domain exchanges. All wavefield computations (in particular the FFTs) are local to a single cluster node. The size of the overlaps and to a lesser extent the frequency of subdomain exchanges control the efficiency and exactness of the method.

Theory

Common azimuth migration (Biondi and Palacharla, 1996) derives from the double square root (DSR) of Claerbout (1985). Briefly, the wavefield recorded at the surface (and preprocessed such that its crossline offset is zero) is downward propagated into the earth according to

$$D_{z+dz}(\omega, k_{mx}, k_{my}, k_{hx}) = D_z(\omega, k_{mx}, k_{my}, k_{hx}) e^{-ik_z dz} \quad (1)$$

where k_z is given by the DSR,

$$k_z = \omega \left[\sqrt{\frac{1}{v(r, z)^2} - \frac{1}{4\omega^2} \left\{ (k_{mx} + k_{hx})^2 + (k_{my} + k_{hy})^2 \right\}} + \sqrt{\frac{1}{v(s, z)^2} - \frac{1}{4\omega^2} \left\{ (k_{mx} - k_{hx})^2 + (k_{my} - k_{hy})^2 \right\}} \right] \quad (2)$$

with the missing k_{hy} provided by the common azimuth approximation:

$$k_{hy} \approx k_{mx} \left[\frac{\sqrt{\frac{1}{v(r, z)^2} - \frac{1}{4\omega^2} (k_{mx} + k_{hx})^2} - \sqrt{\frac{1}{v(s, z)^2} - \frac{1}{4\omega^2} (k_{mx} - k_{hx})^2}}{\sqrt{\frac{1}{v(r, z)^2} - \frac{1}{4\omega^2} (k_{mx} + k_{hx})^2} + \sqrt{\frac{1}{v(s, z)^2} - \frac{1}{4\omega^2} (k_{mx} - k_{hx})^2}} \right] \quad (3)$$

There are a number of ways to do the downward continuation, including split step Fourier (Stoffa et al, 1990), Fourier finite difference (Ristow, 1984), and phase screen (Wu, 1994). Our implementation follows the "phase shift plus interpolation" (PSP) strategy of Gazdag and Sguazzerro (1984).

Domain decomposition and common azimuth migration

A key point is that the downward continuation depends on all the spatial wavenumbers present in the data, which implies that the problem is not easily divided spatially. Fan et al (1997) develop an overlap domain decomposition method for modeling, and apply it to both space-domain finite difference and Fourier pseudo-spectral (Kosloff and Baysal, 1982) modeling, relying on adequate overlap regions and tapering (Liao and McMechan, 1993) to prevent Fourier artifacts. A difference between the modeling cases described by Fan and our downward continuation case is that, propagating at 90 degrees, an impulse will traverse the entire computational domain in a single step. We typically go many steps between domain exchanges, and so do not precisely duplicate the single domain computation with ODD. As a practical matter, we find very little difference between the final migration results obtained with and without, presumably because other well known limitations of downward continuation in complex velocity supersede.

Method

We divide frequencies into subdomains which migrate independently for some depth range (typically chosen to be 100s – 1000m), exchange information between domains, migrate independently again, and repeat.

A frequency slice is divided into overlapping subdomains along the crossline axis. The size of the overlap is adjustable, but is typically chosen to be 0.5 – 1 km. With 20-40m crossline sampling, the overlap corresponds to about 20-50 points. A 20-point taper region is added to either side to prevent Fourier wraparound. All told, the computational footprint of each subdomain is increased by 80-140 lines, or half that at the two ends of the survey. This adds some overhead, but much less than migrating the regions independently. Keeping the memory and survey area relatively small provides further opportunities for meaningful optimization (by affording room in memory for more copies of the wavefield); the benefit gained here may in practice outweigh the benefit of the smaller overlap. In any case, they are complimentary.

We migrate the subdomains independently on different compute nodes, through some range of depths. At some depth interval (typically many times larger than the migration depth step), a portion of the overlap in each subdomain is replaced with the wavefield from the same location in the adjacent subdomain. Then the subdomains go back to migrating independently until the next meeting depth. Referring to Figure 3 below, a pseudo-code representation of the method would look like:

```
For all depths{
  Subdomain1 migrate (A,D)
  Subdomain2 migrate (B,E)
```

```
  If at a meeting depth{
    Subdomain1(C,D) = Subdomain2(C,D)
    Subdomain2(B,C) = Subdomain1(B,C)
  }
}
```

The crossline axis is convenient for common azimuth migration because the lack of an offset axis makes the subdomain exchange simple, but nothing prevents dividing the data further.

The compute nodes working on each subdomain may reach the pre-arranged meeting depth at different times, and the faster nodes will be idle while waiting for the slower nodes. In practice, because the subdomains are chosen to be the same size and the nodes working on each subdomain tend to be identical (usually neighbors in a cluster – a rack of 64 nodes will be divided into e.g. 16 teams of 4 nodes each), the idle time is quite small.

The ODD method effectively increases the maximum survey size that will fit into the migration algorithm, and allows more cpu power to be applied. Potential drawbacks are an inexact propagator due to infrequent domain exchanges and too-small overlap; and/or inefficiency due to frequent domain exchanges and too-large overlap regions. We evaluate these with an impulse response example and the C3 data.


Examples

Figure 1 shows an impulse response migrated as a single domain (a) and with ODD in 2 subdomains (b). The subdomain boundary is in the center. The steepest part of the impulse response shows some elliptical artifacts, though the event itself is intact. The overlap in this case is 600m, with a subdomain exchange every 200m, or 20 depth steps. Figure 1(c) shows a difference section. Aside from the steepest part of the deep impulse, the differences are mostly differing wraparound artifacts.


Figure 2 shows a crossline section from the C3 synthetic computed in 2 subdomains. Figure 2(a) is computed treating the data as a single domain, 2(b) as two subdomains, and 2(c) is the difference panel. The images are essentially identical.

Because data in the overlap zones are migrated twice rather than once, speed does not quite scale with number of compute nodes. The overlaps need to be wide enough to accommodate taper zones in addition to providing for energy exchange between subdomains. Commonly around 20-30 percent of the data migrating in a single compute node is redundant, with an efficiency of around 0.7.


Domain decomposition and common azimuth migration



(a)




(b)




(c)


Figure 1: Impulse response test for ODD method. Panel (a) shows an impulse response computed in a single domain, panel (b) shows the impulse response computed in two subdomains. Panel (c) shows the difference panel.



(a)



(b)



(c)

Figure 2: Crossline section from C3 test for ODD method. Panel (a) shows a section from a single-domain migration. Panel (b) shows a section from a two-domain migration. Panel (c) is the difference panel.

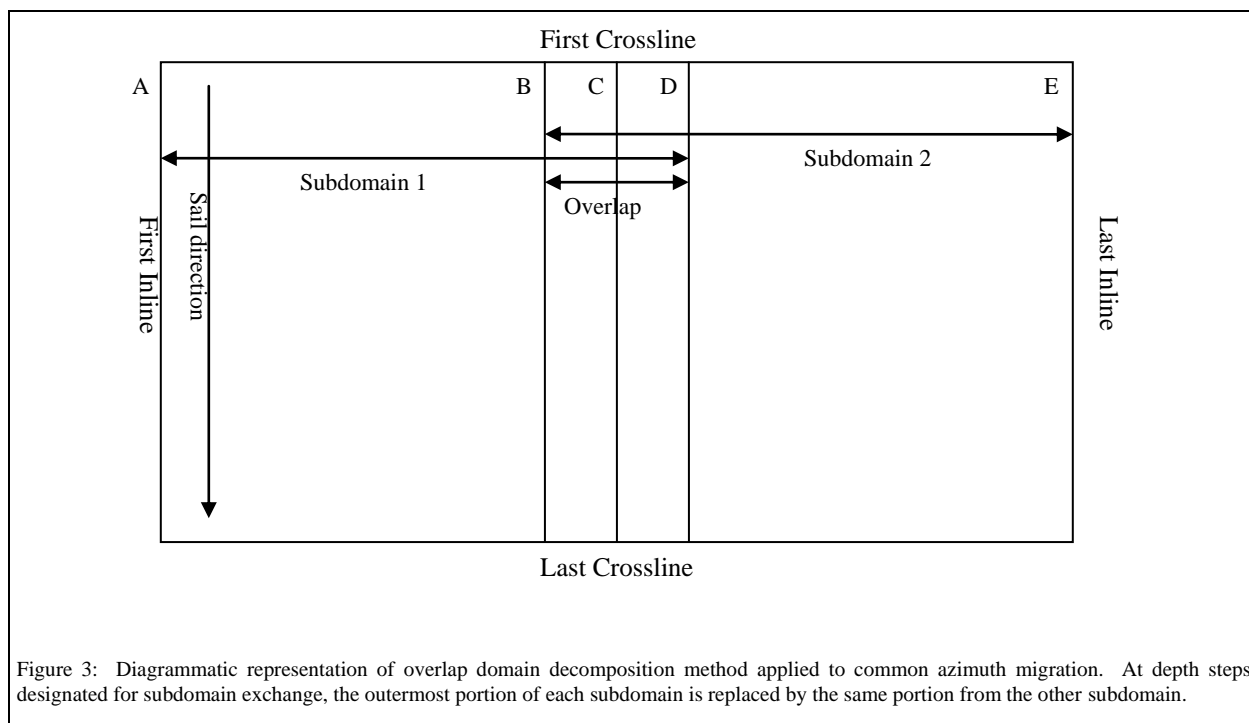
Domain decomposition and common azimuth migration

Conclusions

We present an efficient and straightforward method for parallelizing the spatial axes in a common azimuth migration, using overlap domain decomposition. The method increases the maximum survey size that will fit into the migration program, and speeds turnaround time by increasing the amount of cpu power that can be used for a given number of frequency slices. The resulting operator is an approximation of the single-domain common azimuth migration, but the differences are small.

Acknowledgments

The author would like to acknowledge Joe Higginbotham for helpful discussions.



EDITED REFERENCES

Note: This reference list is a copy-edited version of the reference list submitted by the author. Reference lists for the 2008 SEG Technical Program Expanded Abstracts have been copy edited so that references provided with the online metadata for each paper will achieve a high degree of linking to cited sources that appear on the Web.

REFERENCES

- Biondi, B., and G. Palacharla, 1996, 3D prestack migration of common-azimuth data: *Geophysics*, **61**, 1822.
- Calandra, H., A. Khoury, F. Bothorel, and P. Vezolles, 2007, A massively parallel implementation of the common-azimuth depth migration: 77th Annual International Meeting, SEG, Expanded Abstracts, 2285.
- Carcione, J., 1991, Domain decomposition for wave propagation problems: *Journal of Scientific Computations*, **6**, 453–472.
- Claerbout, J. F., 1985, *Imaging the Earth's interior*: Blackwell.
- Faccioli, E., F. Maggio, A. Quarteroni, and A. Tagliani, 1996, Spectral-domain decomposition methods for the solution of acoustic and elastic wave equations: *Geophysics*, **61**, 1160–1174.
- Fan, J., K. Nihei, L. Myer, N. Cook, and J. Rector, 1997, Overlap domain decomposition method for wave propagation: Annual International Meeting, SEG, Expanded Abstracts, 1485–14.
- Gazdag, J., and P. Sguazzero, 1984, Migration of seismic data by phase shift plus interpolation: *Geophysics*, **49**, 124–131.
- Kosloff, D., and E. Baysal, 1982, Forward modeling by a Fourier method: *Geophysics* **47**, 1402.
- Liao, Q., and G. McMechan, 1993, 2D pseudospectral viscoacoustic modeling in a distributed-memory multiprocessor computer: *Bulletin of the Seismological Society of America*, **83**, 1345–1354.
- Ristow, D., and T. Ruhl, 1994, Fourier finite-difference migration: *Geophysics*, **59**, 1882.
- Stoffa, P. L., J. Fokkema, R. M. de Luna Freire, and W. P. Kessinger, 1990, Split-step Fourier migration: *Geophysics*, **55**, 410–421.
- Wu, R. S., 1994, Wide-angle elastic wave one-way propagation in heterogeneous media and an elastic wave complex-screen method: *Journal of Geophysical Research*, **99**, 751–766.