Introduction

Whilst advanced imaging algorithms have proved to be a crucial tool for many imaging problems, the hindrance for imaging deep and difficult targets is often associated with data coverage. Effective seismic imaging and model building in 3D hinges on several key factors, such as the offset range (determining the rough maximum depth of imaging), the azimuthal coverage, and the density of sampling (for both sources and receivers).

Many acquisition solutions exist that aim to improve offset coverage, azimuthal coverage, and sampling density, however a significant increase in acquisition cost is inherent. These solutions may feature multiple acquisition vessels, multiple source vessels, or varied shooting patterns, but the factor that increases survey cost the most dramatically is time. A typical survey crew costs roughly $500,000 per day, with 3D surveys often lasting 3-6 months. The costs incurred while shooting these data are enormous, and so methods of increasing suvery efficiency are of huge interest. Additionally, these datasets are extremely large, and expand with increase in offset, azimuth and sampling density. Consequently, the imaging approach used must often be adapted to accomodate these huge data quantities.

A typical Narrow Azimuth Towed Streamer (NATS) acquisition vessel (Figure 1, top left) features a typical maximum offset of 10 km in the inline direction, and 1 km in the crossline direction. Typical receiver sampling intervals are 25 m and 50 m in the inline and crossline directions respectively, the source sampling is typically comparable. The rectangular patch indicates the subsurface illumination achieved with this configuration. The sampling density NATS is reasonable, however the offset range limits the depth of imaging to 10 km (for flat geology) and both the azimuthal richness and the crossline offset are very low. Since seismic model building techniques rely on redundant reflection point information in the subsurface, in terms of offset and azimuth, then accurately building these models in the crossline direction becomes difficult.

One method is to reacquire the survey along a different azimuth. The azimuthal range of these data increases linearly with azimuths acquired, but so does survey time, and hence cost. Another method is Wide Azimuth Towed Stream (WATS), also shown in Figure 1. As many as four extra source boats are used, greatly increasing subsurface illumination, by virtue of both offset and azimuth (again, indicated by the rectangle.)

A WATS type approach has been shown to produce excellent images, **?**, but a major bottleneck is the increase in survey time. After each shot point has been fired a sufficient waiting period is elapsed before the next shot point, so that energy has sufficiently dissipated. Typical imaging algorithms assume a single shot point per record, so high-amplitude and coherent overlapping arrivals between sources will be very damaging imaging results. Often 15 seconds are elapsed before the next point is acquired, so for a survey comprising of four source boats, a full minute must elapse before a given source boat acquires it's next location.

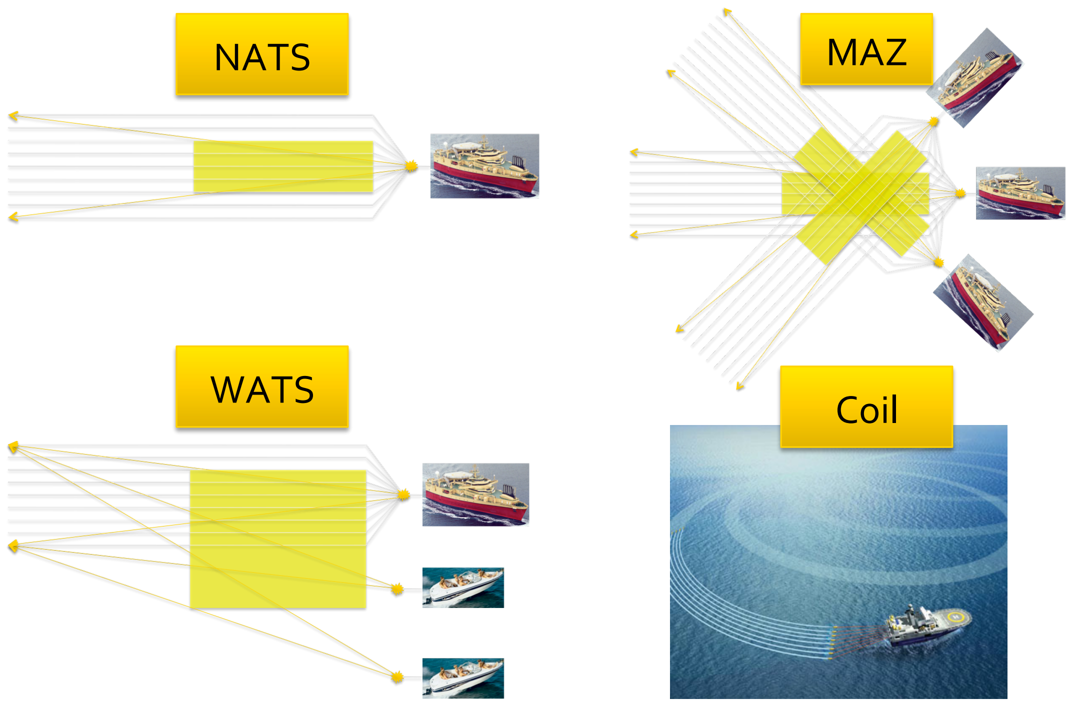

Figure 1: Rough plots of four acquisition techniques - Narrow Azimuth Towed Streamer (NATS), Multiple Azimuth Acquisition (MAZ), Wide Azimuth Towed Stream (WATS) and Coil shooting. [**NR**]

If this waiting period restriction was relaxed, many more source points could be acquired in a given time interval (**?**). This concept has the potential to drastically reduce survey time, and the financial savings can be drastic. Acquiring data in this manner is referred to as simultaneous source data, continuously recorded data, and blended data, amongst others. The recorded data will feature significant, high amplitude, overlapping arrivals between shot records, and this must be taken into account during processing, model building and imaging. Naturally, this will make the data processing more expensive, but this increase in processing cost will be drastically less than the money saved through quicker acquisition.

BLENDED DATA

There are numerous opportunites, and challenges, provided by blended acquisition. A number of studies have looked at extracting subsurface properties directly from the overlapping records, such as **?**, however these approaches implicitly require an accurate velocity model. If blended surveying is to become an option for exploration type surveying, then processing techniques which do not rely on prior subusurface knowledge are essential.

Developing a methodology that can effectively separate these overlapping shot gathers into their conventionally acquired equivalent records would provide the link between efficient data acquisition, and contemporary model-building and imaging work flows. Thus, the focus of this thesis will be on data separation, rather than direct imaging of these data. There are a number of existing separation techniques, and many of them work extrememly well. In particular, the work by **?** is now used widely by BP. **?**, **?** and **?** all have demonstrated useful separation procedures, by exploiting the randomness of source points acquired.

Figure 2 shows a 2D dataset where randomly delayed overlap between sources has been permitted. By studying the receiver gather panel, it is clear that this overlapping energy is randomly scrambled in this domain. A selection of random noise attenuation methodologies can now be employed, and this is how the mentioned authors achieved their separation. However, Figure 3 shows these same data, but with constant delays between sources. There is now no coherency difference between the source and receiver domains; attempting to apply one of these mentioned techniques would now entirely fail.

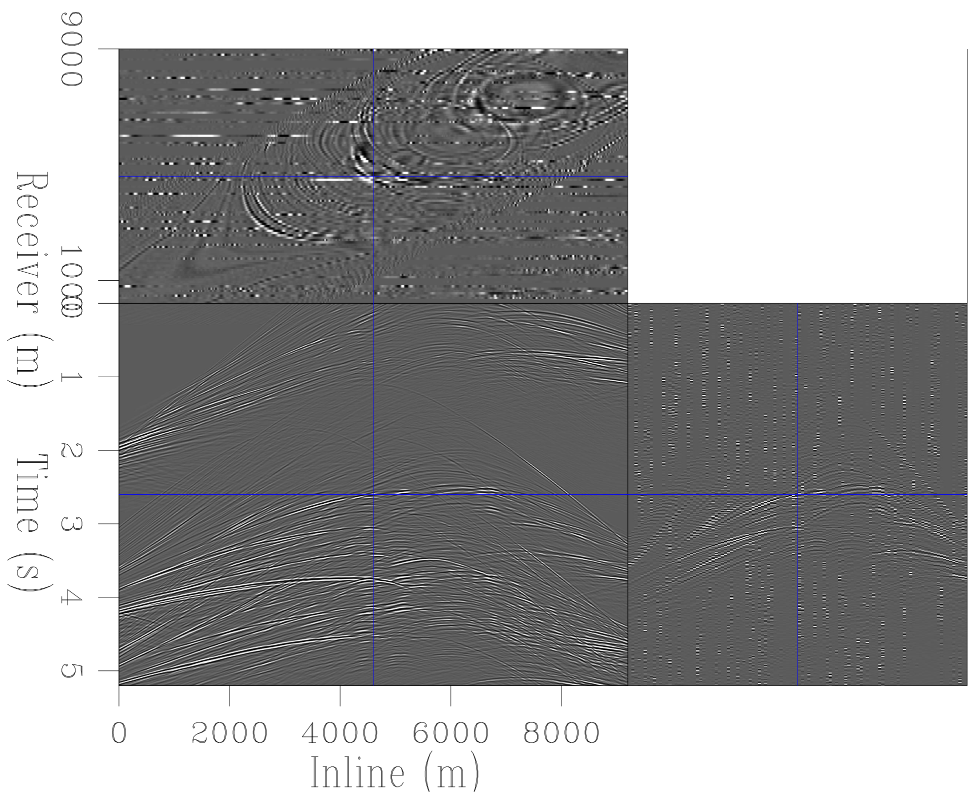

Figure 2: A dataset acquired with random delays between overlapping sources. The left panel is an example shot gather, the right panel, an example receiver gather. [**CR**]

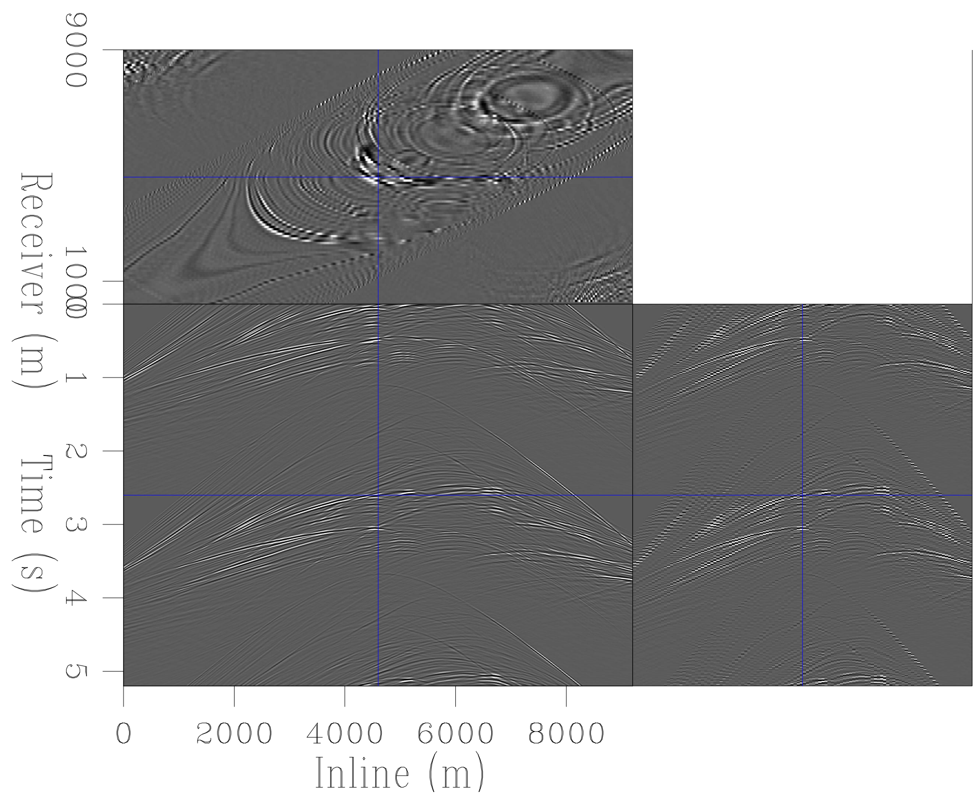

Figure 3: A dataset acquired with constant delays between overlapping sources. [**CR**]

Image space separation

These successful data-space approaches use either Fourier transforms, Radon transforms or wavelet transforms to move to a domain where coherency attributes can be isolated. By a series of domain transforms, or by an inverse process, random noise is iteratively removed and clean receiver gathers formed. Subsequently, by the principle of reciprocity, separated shot gathers have been created.

It is also possible to transform to the extended image domain, rather than the Fourier, Radon or wavelet domain. The image domain is often usefully employed, since it is of much lower dimensionality, and has a high signal-to-noise ratio, when compared to the data-space. This image space transform, known as migration, requires a subsruface velocity model. Instead, by transforming to the extended image domain, the requirement of an accurate velocity model is relinquished. Kinematic and amplitude information contained within these data are all preserved. It is possible to manipulate these extended images, according to a set of criteria, and to then transform back to the data-space. **?** used a similar method for surface-related multiple removal.

This thesis will use the imaging transform as a type of coherency filter, and separate overlapping data using the image domain. This approach is effective for both randomly delayed shots and constantly delayed shots, since even for consistent delays the artifacts are not coherent between separate images. Chapter 4 will formally introduce how an inverse forward modeling procedure can be formulated, for accurate data separation under a variety of blending methodologies.

Artificial blending

As alluded to above, an increase in acquisition complexity leads to an increase in the quantity of data acquired. For data acquired in an unblended fashion, subsequent, artificial blending can be a useful endeavour. This is known as phase encoding.

For seismic imaging and inversion, shots are migrated and modeled invidually, and then combined. It is possible to delay and sum shot gathers together, and then image and model them concurrently (**?**). Such a methodology induces crosstalk artifacts, but these can be reduced according to how shots are scaled and summed together. Furthermore, with stacking these artifacts will be greatly diminished, while real events will actively sum. The difference between phase encoded and field blended data is that the processor has direct control over how the shots are combined.

Chapter 3 will analyse two methods of data space reduction - phase encoding, and randomized sampling, whereby rolling subsets of the data are selected during separate iterations of an inverse process. This is in contrast to the different set of tools needed to process field-blended data, and provides a meaningful discussion and comparison between field blending and artificial blending.

HPC METHODS

The extended imaging, demigration, and inversion approaches applied throughout this thesis are extrememly computationally intensive. Simulating a simplified form of the wave-equation is the engine of all these methodologies, and this simulation requires millions of finite-difference stencil convolutions to apply the spatial derivative.

To account for this demand, a significant portion of the research and work necessary for this thesis was in accelerating wave-equation imaging and modeling. The main focus of this was in the use of Graphical Processing Units (GPUs) to accelerate this stencil convolution. GPUs are an example of Single Instruction Multiple Data (SIMD) computing, where a single set of instructions (a kernel) are applied to a large array, or set of arrays, very efficiently. A GPU, at base level, can be considered as several hundred (to several thousand) mini CPUs, each acting independently, and each with an individual, and small, low latency memory level. Acoustic wave propagation is a close to ideal candidate for GPUs, since the bulk of the computation can be made to fit this SIMD framework.

GPU computing poses challenges, however. Each individual card features a small global memory, so for meaningful computation an array of GPUs must be used in parallel, and communication between these units is paramount. Chapter 6 will discuss GPU computing in detail, and elaborate on many of these details.

It should be noted that the vast majority of results contained in this thesis use a GPU based code library for wave propagation, thus most results are labelled Conditionally Reproducible (CR). This is because an available GPU card with compute capability of CUDA 4.0+ is required for the bulk of the computation, despite run times usually being brief. The codes have been tested and validated on M2090, K10 and K40 cards.

THESIS OVERVIEW

Wave-equation imaging and linearized inversion: Chapter 2 introduces the core concepts of wave-equation imaging, and it's extension to linearized inversion. These techiques are used throughout the entire thesis, and so a formal introduction is necessary. With particular reference to Reverse Time Migration, results from a complex synthetic model are shown for adjoint imaging and for data-space inversion. Quantitive measures of these effectiveness of these methods are introduced and compared, and methods for accelerating the process discussed. The concept of preconditioning is discussed, and the use of a few simple operators are employed, with the goal of increasing inversion efficiency, as a function of cost. This will be contrasted later in the thesis, where regularization methods and HPC methods are used to improve and accelerate linearized inversion.

Phase-encoding and randomized sampling: In Chapter 3 two methods of data-space reduction are introduced - phase-encoding and randomized sampling. The concepts of both static and dynamic phase-encoding are discussed, the algorithms contrasted, and a set of inversion results shown. Residual reductions are plotted as functions of both iteration number and cost, and the improvement of convergence discussed. Phase-encoded is then contrasted with randomized sampling, where a rolling subset of the data is used between iterations. Results are plotted as functions of iteration and cost, again. Results of augmenting these two methodologies are then mentioned.

Simultanous shot separation: Chapter 4 contains much of the core work of the thesis. Existing separation techniques, and their drawbacks, are discussed at length, and then the concept of image space separation is introduced. The domain coherency differences of a variety of blending techniques shown clearly, and the manifestation of these different delays in the extended image space, for a simple and complicated synthetic model, are shown. Isolating events according to the curvature by using a parabolic Radon inversion in the angle domain is demonstrated, and the results show strong event separation. The original concept of extended demigration for shot separation is formally presented, as is it's extension to an inverse forward modeling process. The concept of incorporating a blending operator into an image-space deblending technique is demonstrated. Separation results using this innovative technology are presented, for both accurate and inaccurate velocity models, over a range of three synthetic models of increasing complexity. Quantitiative measures of the effectiveness of this concept are used, and the results show good convergence.

Field data example: In Chapter 5, several of the concepts discussed in Chapter 2 and Chapter 4 are applied to a field dataset. This dataset featured Ocean Bottom Node receivers, and a set of anisotropic models. This required a non-negligible extension of these procedures, and the results of linearized inversion show large improvements. These data are then blended and deblended, using an acoustic deblending engine, with accurate and inaccurate velocity models. The images constucted from these separated data are directly comparable, however these separate receiver gathers show a number of artifacts. Possible sources of these are then discussed, and improvements for the procedure suggested.

High performance computing solutions: Finally, Chapter 6 presents all the research and work that went into building the wave-equation engine that created the results for all previous chapters. Both CPU and GPU based linearized inversion solutions are discussed and contrasted, and the requirement of different source wavefield domain boundaries mentioned. The extension to multiple GPU propagation is then demonstrated in detail, for propagation, RTM and Born modeling. How the necessary extension to anisotropic propagation was done (to construct the results of Chapter 5) is then mentioned, with reference to how GPU memory and storage structures were made use of.