

Accelerating Subsurface Offset Gathers for 3D Seismic Applications

Oliver Pell (Maxeler Technologies) and Robert G. Clapp (Stanford University)

Introduction

Downward continued based migration is often superior to Kirchhoff methods in complex areas. For velocity estimation and lithological determination they require a change in strategy compared to Kirchhoff based approaches. In Kirchhoff based approaches moveout and amplitude information is evaluated as a function of offset. For downward continuation based methods angle gathers are constructed from subsurface offset gathers (de Bruin et al., 1990; Prucha et al., 1999; Biondi and Symes, 2004).

The cost of constructing the subsurface offset gathers is trivial for source-receiver migration methods based on the Double Square Root Equation but can be a dominant cost in shot profile and plane wave methods (Rickett and Sava, 2002). With the increased use of wide azimuth geometries (Michell et al., 2006) and the resulting 3D angle gathers, over 90% of CPU cycles can be spent in constructing the subsurface offset gathers.

Hardware accelerators are emerging as a powerful solution to computationally intensive problems. A standard desktop PC or cluster node can be augmented with additional hardware dedicated to providing substantially increased performance for particular applications. FPGA-based hardware accelerators can offer order-of-magnitude greater performance than conventional CPUs, providing the algorithm to be accelerated performs a large number of operations per data point. Constructing subsurface offset gathers involves a significant number of operations for each data point, making it an ideal candidate for acceleration.

We show that FPGAs can achieve up to 48x speed-up for 2D angle gathers. For 3D gathers, we project that speedups of over 200x will be obtainable with upcoming FPGAs.

Angle Gathers from Shot Profile Migration

Claerbout (1971) noted that you could simulate sources and receivers in the subsurface by applying the DSR equation to the wavefield recorded at the surface. For 3D source-receiver based migration methods a 4D volume (cmpx, cmpy, hx, hy) is downward continued and the zero-time, zero-offset portion of the volume is the image of the subsurface at that depth. In shot profile migration the source and receiver wavefields are downward continued separately by the Single Square Root (SSR) equation. To obtain an image I at a given depth z and location x :

$$I(x, z) = \sum_s \sum_w S(x, z, w, s) \cdot G(x, z, w, s)^* \quad (1)$$

where S is the source wavefield, G is the receiver wavefield, w is the temporal frequency, s is the shot index.

de Bruin et al. (1990), Prucha et al. (1999) and Biondi and Symes (2004) all provided mechanisms to create reflectivity as a function of angle based on the focusing of the energy around zero-offset. In source-receiver migration, the analysis is naturally done by analyzing the focusing in the h plane. In shot profile migration h does naturally occur. Rickett and Sava (2002) explained how to create subsurface offset for shot profile migration by cross-correlating the source and receiver fields by various shifts. As a result subsurface offset volumes can be created by applying:

$$I(h, x, z) = \sum_s \sum_w S(x-h, z, w, s) \cdot G(x+h, z, w, s)^* \quad (2)$$

The cost of constructing these sub-surface offset gathers can be significant. The cost of downward continuing a single shot a single depth step is dominated by the FFT cost which grows by $n \log n$ with the size of the data n . The cost of constructing the subsurface offset gathers is on the order of $nh \times n$, where nh is the number of subsurface offsets. Constructing a 2D subsurface offset gather can be a significant cost in shot profile migration, constructing 3D subsurface gathers can be the dominant cost.

Sub-surface Offset Imaging Condition on FPGA

The shot profile imaging condition has a high arithmetic density, suggesting suitability for FPGA acceleration. The zero time part of the imaging condition (Equation 2), requires a summing over all frequencies meaning that each output requires a summing over nf values. In addition each input point is used in nh cross-correlations. The on-chip memory requirement is $O(nf \times nh)$, well within the capabilities of modern FPGAs for hundreds of frequencies and dozens of subsurface offsets.

To test the applicability of this approach we compared the result of constructing angle gathers for the 2-D Marmousi synthetic dataset. Figure 1 shows the zero-subsurface offset image obtained from implementing the imaging step of shot profile migration on both the processor and the FPGA. The images are indistinguishable. Figure 2 shows an angle gather constructed from the offsets computed with both CPU and FPGA. Note that visually the kinematics are identical.




Figure 1: The zero-subsurface offset image from implementing the shot profile imaging condition on the CPU (left) and the FPGA (right).




Figure 2: The angle gather obtained for $X=5000$ by implementing the same imaging condition on the CPU (left) and the FPGA (right).




Figure 3: Speed-up from implementing 2D offset gatherers for 400 inline CMPs and 200 frequencies on a Xilinx Virtex-4 FX100 FPGA compared to a 2.8Ghz AMD Opteron CPU. Performance projections generated by the Maxeler Parton analysis tool are compared to measured results. Results are shown for (left to right) 25 subsurface offsets, 41 subsurface offsets and 55 subsurface offsets.

To test the speed-up offered by the FPGA implementation we ran a larger 3D problem. Specifically the cost of constructing 25, 41 and 55 subsurface offset gatherers from 400 inline cmps, 200 frequencies and varying number of crossline cmps. Figure 3 shows a comparison of FPGA performance to a 2.8Ghz AMD Opteron PC with 16GB of RAM, showing speedups of up to 48x. The software implementation was written in C and compiled using both *gcc* and *icc* with full optimization.

The FPGA accelerator was implemented on a Maxeler MAX-1 FPGA platform equipped with a single Xilinx Virtex-4 FX100 FPGA. The accelerator circuit consumes 76% of the logic resources of the device running at 125Mhz.

3D Offset Gatherers

3D subsurface offset volumes can be created by applying:

$$I(h_x, h_y, x, z) = \sum_s \sum_w S(x - h_x, y - h_y, z, w, s) \cdot G(x + h_x, y + h_y, z, w, s)^* \quad (4)$$

If the 2D offset gather involved constructing nhx subsurface offsets, the 3D gather to construct nhx inline offsets and nhy crossline offsets has a computational complexity (and associated runtime cost on a CPU) nhy times greater. On FPGA, this greater complexity gives potential for more acceleration since there is a higher arithmetic density per data point transferred to/from the CPU, although a larger number of output points must now be generated, depending on the relative balance of the number of frequencies and subsurface offsets. For 200 frequencies and 41 inline subsurface offsets, there is an increase in acceleration potential of approximately 5x.

3D gatherers also have a substantially increased on-chip memory requirement of $O(nf \times nhx \times nx)$. For 200 frequencies, 41 inline subsurface offsets and 1000 inline cmps the actual memory requirement is approximately 16MB; this is beyond the capabilities of on-chip memory, requiring the use of a small amount of slower off-chip but on-card memory.

We have not yet implemented the 3D gatherers on an FPGA. Figure 4 shows the results of a projection of the expected speedup over a 2.8Ghz AMD Opteron as the number of crossline subsurface offsets is increased. Results for three different FPGA platforms are shown. Initially FPGA performance scales better than linearly, as the extra performance potential is absorbed, after this runtime scales linearly with problem size. Speed-ups of over 200x are achievable with an upcoming FPGA chip.

Conclusion

We implemented the subsurface offset imaging condition for shot profile migration on an FPGA. We show that up to a factor of 50x speedup can be obtained compared to a conventional CPU for 2D gatherers. For 3D gatherers, speedups of over 50x are obtainable at present, 200x will be achievable using an upcoming Virtex-5 SX95T chip.

Acknowledgement This briefing note is a summary of an abstract presented at the 77th SEG Meeting (Pell and Clapp, 2007).

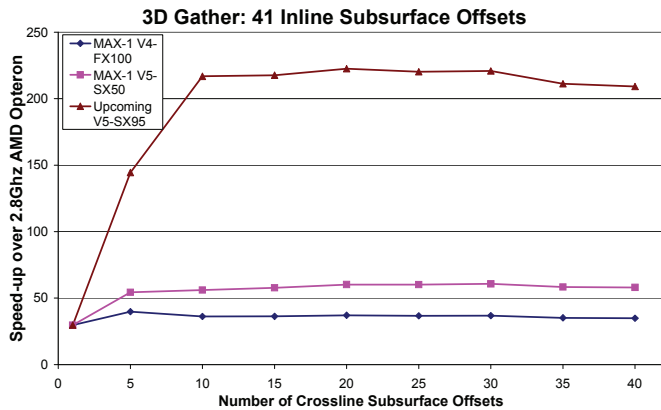


Figure 4: Speedup projections generated by the Maxeler Parton analysis tool for implementation of 3D offset gatherers on an FPGA with 16-bit data compared to a 2.8Ghz AMD Opteron processor. Projections are shown for different Xilinx FPGA platforms: (a) Virtex-4 FX100, as used for the results in Figure 2, (b) a newer Virtex-5 SX50 platform and (c) a Virtex-5 SX95 platform (available Q1 2008).

References

Biondi, B. and W. Symes, 2004, Angle-domain common-image gatherers for migration velocity analysis by wavefield-continuation imaging: *Geophysics*, 69, 1283–1298.

Claerbout, J. F., 1971, Toward a unified theory of reflector mapping: *Geophysics*, 36, 467–481.

de Bruin, C. G. M., C. P. A. Wapenaar, and A. J. Berkhout, 1990, Angle-dependent reflectivity by means of prestack migration: *Geophysics*, 55, 1223–1234.

Michell, S., E. Shoshitaishvili, D. Chergotis, J. Sharp, and J. Etgen, 2006, Wide azimuth streamer imaging of mad dog; have we solved the subsalt imaging problem?: 76th Ann. Internat. Mtg., 2905–2909, Soc. of Expl. Geophys.

Pell, O. and Clapp, R. G, 2007, Accelerating Subsurface Offset Gatherers for 3D Seismic Applications using FPGAs: 77th SEG Mtg., 2383–2387

Prucha, M., B. Biondi, and W. Symes, 1999, Angle-domain common image gatherers by wave-equation migration: 69th Ann. Internat. Mtg., 824–827, Soc. of Expl. Geophys.

Rickett, J. E. and P. C. Sava, 2002, Offset and angle-domain common image-point gatherers for shot-profile migration: *Geophysics*, 67, 883–889.