10.1071/PVv2011n152p29

Massively parallel 3D inversion of gravity and gravity gradiometry data

Glenn Wilson

Glenn Wilson[1,4], Martin Čuma[1,2], and Michael S. Zhdanov[1,3]

[1]TechnoImaging, Salt Lake City, Utah, USA.
[2]Center for High Performance Computing, The University of Utah, Salt Lake City, Utah, USA.
[3]Department of Geology and Geophysics, The University of Utah, Salt Lake City, Utah, USA.
[4]Corresponding author. Email: glenn@technoimaging.com

Reliance on desktop computers limits the scale of 3D inversion of gravity and gravity gradiometry surveys, making it impractical to achieve an appropriate level of resolution and detail for geological interpretation. To begin with, airborne surveys are characterised by very large data volumes. They typically contain hundreds to thousands of line kilometres of data with measurement locations every few metres. Often, surveys cover thousands of square kilometres in area with tens of thousands of line kilometres of data. Regional surveys may be even larger and denser as the result of merging multiple and/or historic surveys. Secondly, 3D modelling of large-scale surveys exceeds the capacity of desktop computing resources. And finally, gravity data are finite and noisy, and their inversion is ill posed. Regularisation must be introduced in order to recover the most geologically plausible solutions from the infinite number of mathematically equivalent solutions. Various strategies for 3D inversion have been previously proposed but few lend themselves to truly large-scale 3D inversion. In this paper, we describe how gravity and gravity gradiometry surveys can be inverted to 3D earth models of unprecedented scale (i.e., hundreds of millions of cells) within hours using cluster computers.

Introduction

Structural interpretations of gravity and gravity gradiometry data are often based on some form of Euler deconvolution, wavelet analysis, or analytic signal method. While such methods may provide information about the sources, it is not immediately obvious how this information can be quantified in terms of the density distribution within a 3D earth model. For this reason, inversion of gravity data to a 3D density distribution is an important step in quantitative interpretation. Generalised inversion methods first discretise the 3D earth models into cells of constant density. Then, regularisation is introduced. Regardless of the inversion methodology used, all geological constraints manifest themselves as regularisation that can be quantified through a choice of data weights, model upper and lower bounds, model weights, an a priori model, and the type of stabilising functional used. The stabilising functional incorporates information about the class of models from which a unique solution is sought, and its choice should be based on the user's geological knowledge and prejudice.

It has been common (if not ubiquitous) practice to use smooth stabilising functionals, which minimise the deviation from an a priori model and/or the gradients of the 3D density distribution (Li and Oldenburg, 1998; Li, 2001). However, smooth density distributions are rare in real geology. Economic geology is typically characterised by sharp boundaries of contrasting density, for example, between an ore deposit and host rock, or across a discontinuity. It follows that the various smooth stabilisers can produce results that bear little relevance to real geology. To overcome this problem, Portniaguine and Zhdanov (1999) introduced focusing regularisation that makes it possible to recover 3D density models with sharp boundaries and contrasts. Below, we use this technique. We refer the reader to Zhdanov (2002, 2009) and Zhdanov *et al.* (2004) for further details on focusing regularisation.

For gravity, computational complexity increases linearly with the size of the problem. There are two major obstacles in large-scale 3D inversion. The first one being that storing the kernels of the forward modelling operators requires a large amount of computer memory. Even a small-sized 3D inversion of thousands of data to 3D earth models with hundreds of thousands of cells can exceed memory available on a desktop computer. The second obstacle is the amount of CPU time required to apply the dense matrix of the forward modelling operator to the data and model vectors. The translational invariance of the kernels has been used to reduce the matrices to Toeplitz block structure and use FFTs for matrix-vector multiplication (Pilkington 1997; Zhdanov *et al.*, 2004). This strategy, and others like it, dramatically reduces memory requirements and CPU time. However, these methods presume that the data lies on a regular grid of a flat surface above the topography. This means FFT-base modelling is applicable only if the data have been upward continued to a flat surface or in other special cases (e.g. marine gravity).

Another strategy for 3D inversion is compression (Portniaguine and Zhdanov, 2002; Li and Oldenburg, 2003). However, for the large-scale 3D inversion of tens of thousands of data to models with millions of cells, the compressed linear operators can still be too large to store and manipulate on a desktop computer. As a result, large surveys are often divided into subsets and each subset is inverted separately. The resulting 3D earth models are stitched together post-inversion (Phillips *et al.*, 2010). Depending on the functionality of the software environment, such work flows can become complicated and time consuming. Our goal is to use massively parallel 3D inversion so as to eliminate the need for stitching and to deliver results within hours. Our inversion methodology is similar to those of Zhdanov *et al.* (2004) in that we use the re-weighted regularised conjugate gradient method for minimising the objective functional. Additionally, we have incorporated a wide variety of regularisation options.

Inversion methodology

The gravity potential, $U(\mathbf{r}')$, is linear with respect to the 3D density distributions, $\rho(\mathbf{r})$:

$$U(\mathbf{r}') = \int \psi(\mathbf{r}', \mathbf{r})\, \rho(\mathbf{r}) d^3\mathbf{r}$$

where the kernel functiony $\psi(\mathbf{r}', \mathbf{r})$ is the Green's function for the gravity potential. All components and gradients of the gravity field can be derived from spatial differentiation of $\psi(\mathbf{r}', \mathbf{r})$. Closed-form solutions for the volume integrals over right rectangular prisms of constant density have been previously derived (Okabe, 1979). While exact, these analytic solutions are inefficient to implement; for example, the gravity response requires evaluation of 16 logarithms and 8 arctangents (Li and Chouteau, 1998, p. 344). However, the volume integrals can be evaluated numerically. Zhdanov (2009) showed how for gravity gradiometry, single-point Gaussian integration with pulse basis functions was as accurate as the analytic solution, provided the depth to the centre of the cell exceeded twice the dimension of the cell. This implies that for an airborne gravity gradiometry survey with 80 m ground clearance, the 3D earth model can be discretised to 40 m cubic cells. The advantage of numerical integration is that it significantly decreases the run time when compared to the corresponding analytic solutions.

The above kernels also represent the sensitivity of the data to the variations of the density due to the linearity of gravity fields. Dransfield (2010a) used the same kernels to investigate instrument sensitivity. He demonstrated that at a limited distance, which we call the footprint, the receiver is no longer sensitive to the 3D earth model. The size of the footprint is often less than the size of an airborne survey. Cox *et al.* (2010) previously introduced the concept of a moving footprint for 3D inversion of airborne electromagnetic (AEM) data. They showed that for a single transmitter–receiver pair, there was no need to calculate the responses or sensitivities beyond the AEM's footprint. The sensitivity matrix for the entire 3D earth model could be constructed as the superposition of footprints from all transmitter–receiver pairs. The framework of this approach can be described as follows: for a given receiver, compute and store the sensitivities for those inversion cells within the footprint. The radius of the footprint is based on the rate of sensitivity attenuation. As an analogue of this 3D AEM inversion strategy, we introduce a moving footprint for 3D potential field inversion.

For example, we can consider an instrument 60 m above a homogeneous earth model. Figure 1 presents the integrated sensitivities for each of the gravity fields and gravity gradients. The figure shows that the gravity gradients have approximately 95% of the sensitivity within a 15 km footprint. It also shows the integrated sensitivity for the total magnetic intensity (TMI). The sensitivity of the TMI with respect to the footprint radius behaves similarly to the gravity gradients with about 95% of the sensitivity being within the 15 km footprint. This behaviour is fully expected since the kernels have similar spatial dependencies. Past a 15 km radius, the sensitivity decays very slowly. Increasing the footprint radius beyond 15 km is not practical. Therefore, we conclude that 15 km is an optimal footprint radius for gravity gradiometry.

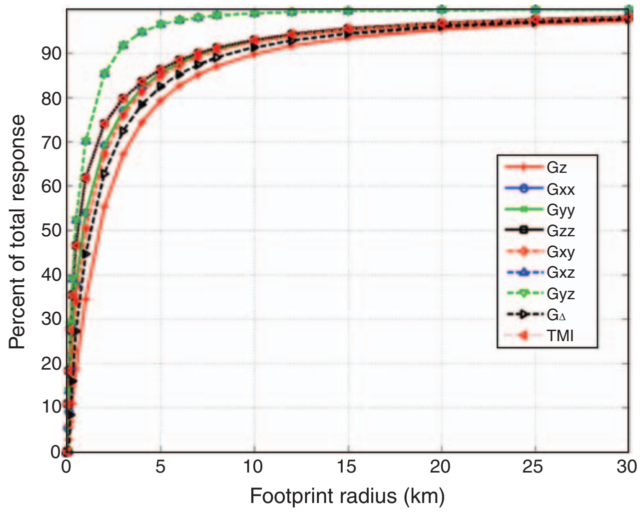

Fig. 1. *Percentage of total response of the integrated sensitivity for the vertical gravity component, various gravity gradients and total magnetic intensity as a function of footprint radius. Note that as the footprint is symmetrical in x and y directions, some gravity gradients overlap.*

Parallel performance

Our 3D inversion algorithm has been implemented as a multilevel parallel application. The 3D inversion domain is divided in a distributed fashion over Message Passing Interface (MPI). On a fine-grained level, loops over the data points and a few other auxiliary loops within each MPI process are further parallelised with a shared memory OpenMP standard. This two-level approach has multiple advantages. It reduces the number of MPI communicating processes, minimising communication stress on the network. It also saves memory, since there are data structures needing to be replicated by each process and most of the data is shared by the OpenMP threads. Finally, it allows for better locality of the processes/threads on the node's boards and sockets, which improves data transfers to/from the main memory. The data locality is critical on modern non-uniform memory architecture (NUMA) computers with a growing number of CPU cores.

In a typical cluster configuration, we run one or two MPI processes per cluster node. Each of these processes launches a number of OpenMP threads – one thread per processor core. The current generation of clusters ship with two hexa-core CPUs (i.e., 12 cores) per node. We have found that it is optimal to run one MPI process per socket (i.e., two per node), with six OpenMP threads per MPI process. The advantage of this is the ability to pin the process to the CPU socket, so that it does not move from one socket to another, which improves the memory performance. We have found that without pinning, the performance can degrade by up to 20%.

Our 3D inversion is relatively light in MPI communication, largely thanks to the linearity of the forward modelling operators. Most MPI communication consists of accumulation of the sensitivities and the regularisation as reduction operations. As a result, the program exhibits excellent parallel scaling. Parallel scaling is usually evaluated with two different metrics. The first one is called *strong scaling*. It measures the performance of a fixed problem size with an increasing number of processors. Another parallel scaling evaluation metric is *weak scaling*. It relates the time to complete one unit of work on one

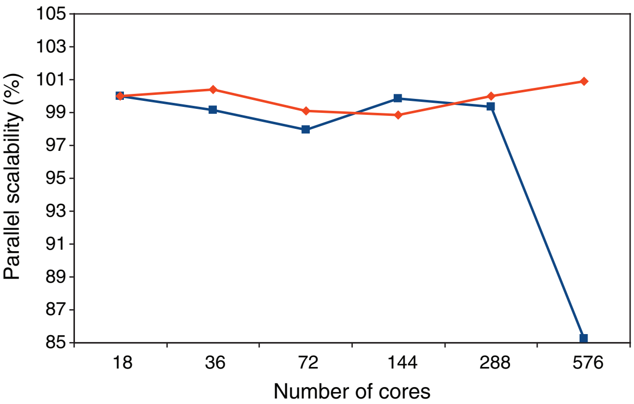

Fig. 2. *Parallel scaling efficiency for 3D inversion of the Vredefort FALCON® data. Strong scaling is shown in blue, and weak scaling is shown in red.*

processing element to the time to perform *N* units of work on *N* processing elements. In both cases, ideal (linear) scaling is 100%. Any scaling below 100% is sublinear, and any scaling above 100% is superlinear. As a side note, it is possible to achieve superlinear scaling due to hardware architectural features that multiprocessor programs can exploit.

We have evaluated the parallel efficiency of our software. All results presented in this paper were run on the University of Utah Center for High Performance Computing's Ember cluster which has 260 nodes, each equipped with two hexa-core (i.e., 12) Intel Xeon CPUs running at 2.8 GHz with 24 GB of RAM and QDR InfiniBand interconnect. Figure 2 shows the parallel scaling efficiency of the subsequent Vredefort case study. In the case of strong scaling, as depicted by the blue line in Figure 2, we chose a 3D model with about 11 million cells and 600 000 data. The scaling efficiency is excellent from 18 to 288 cores. We see a drop at 576 cores. This is due to running 12 rather than 6 cores per process. The memory load is much more uneven for the single MPI process sharing threads on both CPU sockets in the node, which decreases the efficiency by 15%. The weak scaling, depicted by the red line in Figure 2, varied the number of inversion cells from about 11 million cells on 18 cores to about 350 million cells on 576 cores. Again, the scaling is nearly linear with a 1 to 2% difference, which can be attributed to system noise. We draw two conclusions from our scaling analysis. First, our 3D inversion software shows linear scaling and is expected to scale well to thousands of cores. Second, we have identified that process and thread locality is critical in achieving optimal performance, and that one MPI processes should be bound to each socket.

Case study – Vredefort, South Africa

We have applied our massively parallel 3D inversion with a moving footprint to a FALCON® airborne gravity gradiometry (AGG) survey acquired over the Vredefort dome in the Republic of South Africa, approximately 120 km southwest of Johannesburg within the Witwatersrand Basin of the Kaapvaal craton. The Vredefort dome is known as the largest and oldest impact structure on Earth, with a diameter of 250 to 300 km, it is larger than the 200 km Sudbury Basin impact structure in Canada and the 170 km Chicxulub impact structure in Mexico. The impact structure has since been deformed via erosion and tectonic processes, though the centre remains largely unaltered. The centre of the dome is approximately 40 km in diameter and contains an uplifted Archaen basement surrounded by upturned, sub-vertical sediments of the Witswatersrand Supergroup and volcanics of the Ventersdorp Supergroup.

In February 2007, Fugro Airborne Surveys flew a FALCON® AGG survey of 4800 line km over the Vredefort dome area (Dransfield, 2010b). The survey was comprised of two blocks. This study uses 2460 line km of data from the eastern block. The eastern block was flown north-south with a line spacing of 1 km and with 2 east-west tie-lines spaced at 40 km, over an area

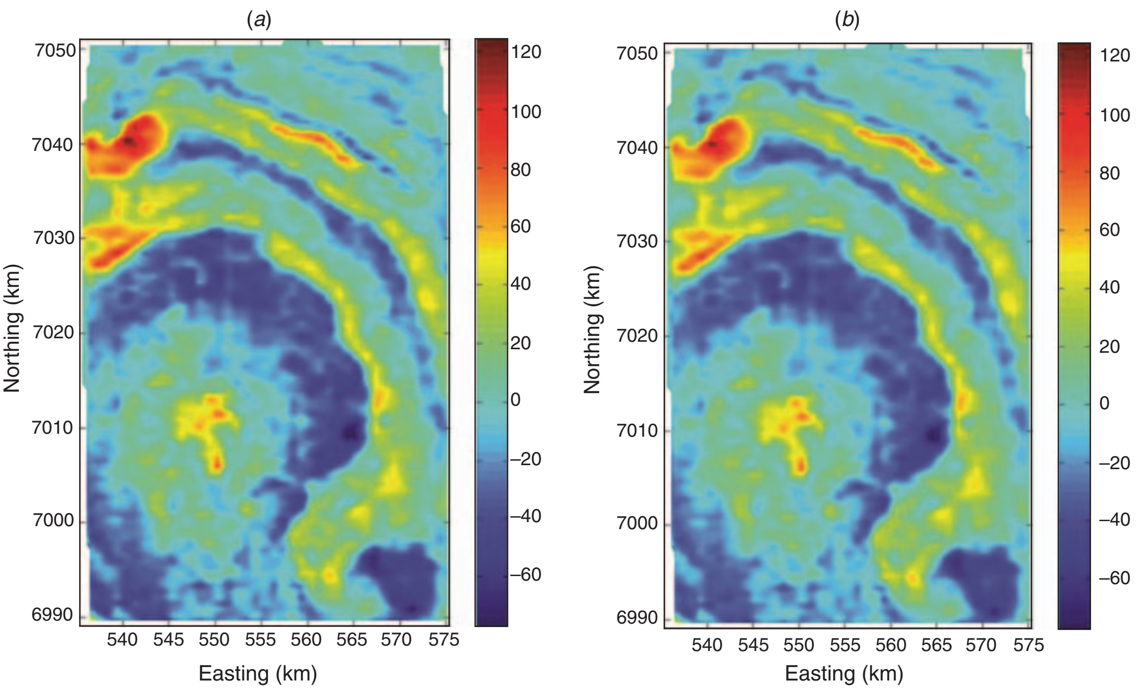

Fig. 3. *(a) Observed and (b) predicted terrain corrected g_{zz} data from joint inversion of all gravity gradient components.*

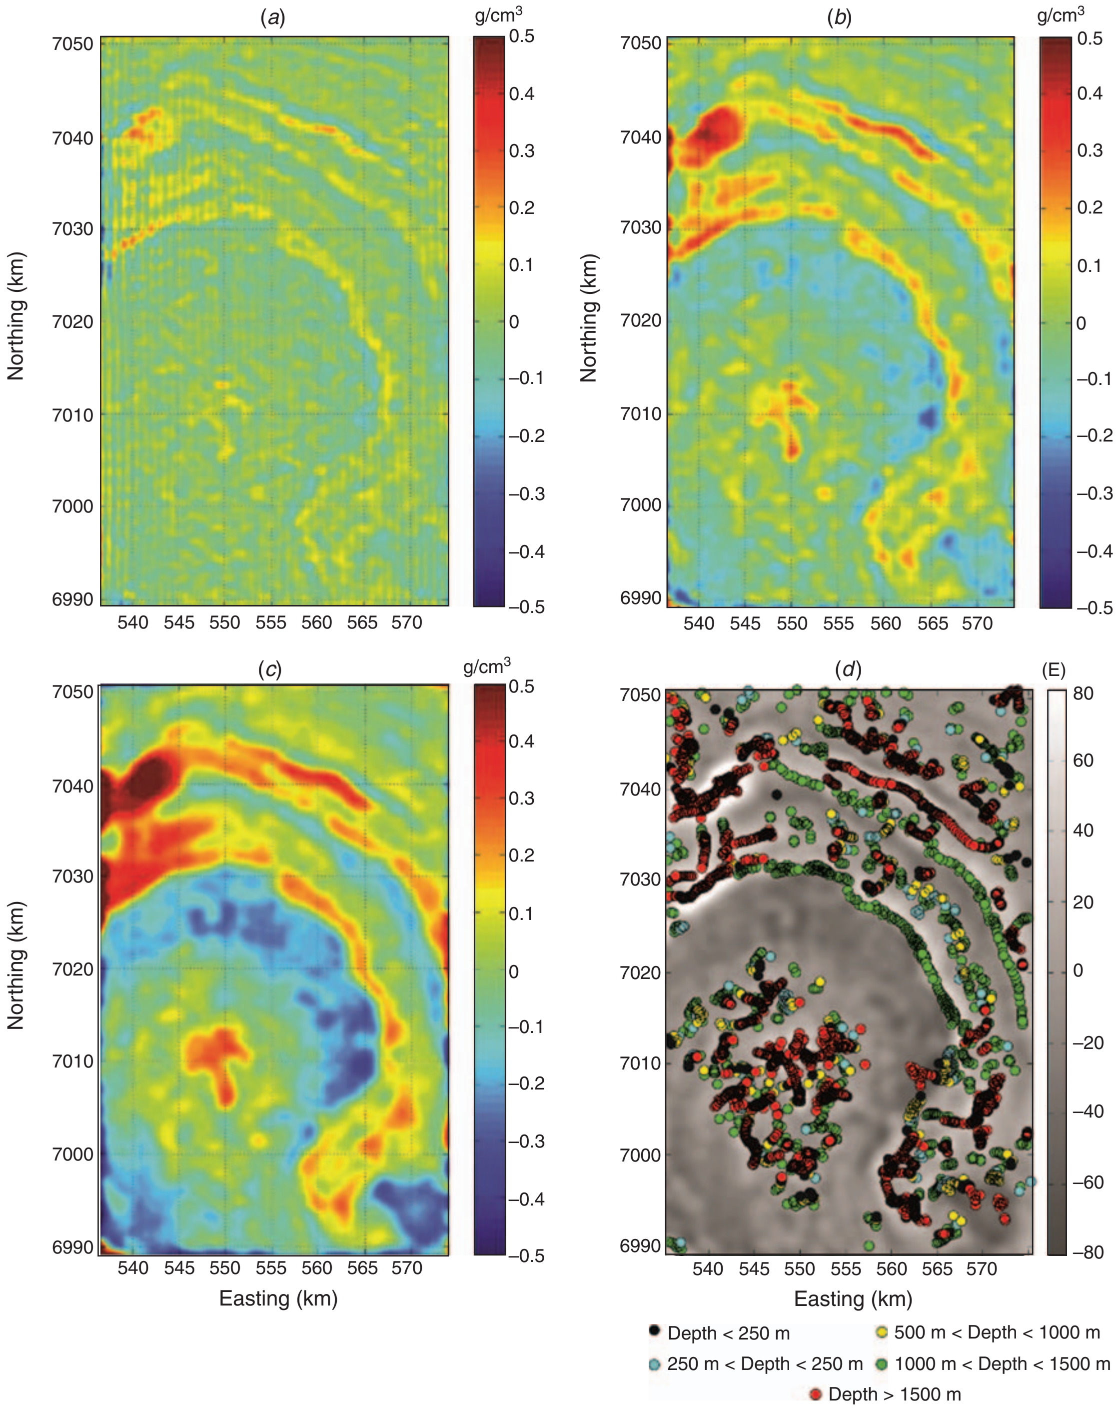

Fig. 4. *Comparison of horizontal cross sections through the 3D density model obtained from joint inversion of all FALCON® gravity gradiometry data at (a) 1137 m ASL, i.e., 500 m depth from topographic peak; (b) 637 m ASL, i.e., 1000 m depth from topographic peak; and (c) 137 m ASL, i.e., 1500 m depth from topographic peak. Panel (d) shows the estimated location of causative bodies based on their depths to the centre of mass from eigenvector analysis (from Beiki and Pedersen, 2010).*

60 km north-south by 40 km east-west, covering most of the Vredefort dome structure. The ground clearance was nominally 80 m flown in a drape over the terrain, corresponding to ellipsoidal heights of between 1430 m and 1740 m. Summer conditions meant moderate to high turbulence at this survey height. The measured gradients were processed by the usual multistep FALCON® processing procedures. After the initial reduction of error due to the residual effects of aircraft motion, the data were demodulated and low-pass filtered with a sixth order Butterworth low-pass filter at a cut-off frequency of 0.18 Hz. The demodulated data were corrected for the self-gradient effects of the aircraft and the tie-lines were levelled. The resulting differential curvature gravity gradient data were further processed to produce terrain-corrected data using a density of 2.67 g/cm^3, and hence the full gravity gradient tensor. In the processing, a low-pass filter with cut-off wavelength of 1000 m was applied to the data.

The 37.8 km × 61.9 km × 2.4 km inversion domain was discretised to over 358 million cubic cells of 25 m dimension. The inversion

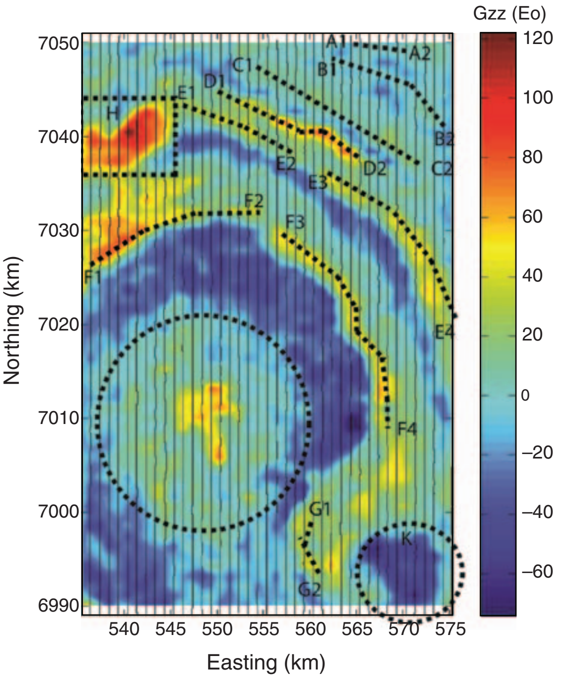

Fig. 5. *Regional geological structure of the Vredefort dome (after Beiki and Pedersen, 2010) superimposed on terrain corrected g_{zz} data from the Vredefort dome FALCON® survey.*

domain conformed to topography and contained no a priori density model. Our previous experience with moving footprint inversion (Cox *et al.*, 2010) has indicated that inversion of redundant data does not aid model recovery. As such, the survey data were decimated by a factor of four, resulting in a data density of one point every 25 m along line. Inversion was run for 85 970 stations, each containing all seven gravity gradients, giving a total of 601 790 data. An example of the observed and predicted data for the g_{zz} component is shown in Figure 3.

Figure 4 shows horizontal cross-sections through the 3D density model at depths of 500 m, 1000 m, and 1500 m below the peak topography, respectively. As we compare our results to the known geological structures (Figure 5), we are able to distinguish the ring structures E and F from the deeper ring structures C and D. The density high in the central part of the dome (J) is related to the deeper structures. Borehole drilling has confirmed that the underlying rock is peridotite, the source of which is open to debate. For example, it is not clear whether these rocks are related to the Bushveld igneous event approximately 2060 Ma (Henkel and Reimold, 1998) or represent mantle material which was uplifted to the surface as a result of the Vredefort impact approximately 2020 Ma (Tredoux *et al.*, 1999). Figure 4 shows a very good agreement between our 3D inversion results and the estimated depths to mass centres obtained from eigenvector analysis by Beiki and Pedersen (2010).

Conclusions

We have developed massively parallel software for the practical large-scale 3D regularised inversion of gravity and gravity gradiometry data to models of unprecedented size. We have also implemented kernels and positivity constraints for 3D magnetic inversion, but that discussion is beyond the scope of this paper. We have achieved linear strong and weak scaling with our parallelisation. Our software can be confidently installed on massively parallel computing architectures. We have introduced a moving footprint, which allows us to represent large, dense linear operators using sparse matrices. The moving footprint approach reduces memory requirements and operation counts for matrix-vector multiplications significantly. Computing the linear operators as needed allows us to handle problems of unlimited size. The effectiveness of our approach has been demonstrated with a case study for 3D inversion of 2460 line km of FALCON® data from Vredefort, South Africa, which included the joint inversion of over 600 000 gravity gradient data to a 3D earth model with over 350 million cells. The computational time for the above inversion totalled about 24 hours using a cluster with 576 CPUs. The results of our inversion agree well with the known geology and independent analyses of the same data.

Acknowledgements

TechnoImaging is acknowledged for support of this research and permission to publish. Čuma and Zhdanov also acknowledge support of the University of Utah's Consortium for Electromagnetic Modeling and Inversion (CEMI), and the Center for High Performance Computing (CHPC). Fugro Airborne Surveys, in particular Dr. Mark Dransfield, are acknowledged for providing the Vredefort FALCON® data, and for permission to publish.

Author Profiles

Glenn Wilson is Chief Geophysicist of TechnoImaging. He holds degrees in physics from Central Queensland and Griffith universities. He previously held positions with OHM Rock Solid Images, BP Exploration and Production, CSIRO Exploration and Mining, and The University of Utah.

Martin Čuma is a Computer Scientist with TechnoImaging and the University of Utah's Center for High Performance Computing. He holds degrees in metallurgy and physical chemistry from the Technical University of Ostrava and Southern University of Illinois at Carbondale.

Michael S. Zhdanov is a Professor in the Department of Geology and Geophysics at The University of Utah, Director of the Consortium for Electromagnetic Modeling and Inversion, and Chief Executive Officer of TechnoImaging. He holds degrees in geophysics, physics and mathematics from the Gubkin State University of Oil and Gas, and Moscow State University. He was previously Director of the Geoelectromagnetic Research Institute of the Russian Academy of Sciences, and Deputy Director of the Institute of Terrestrial and Ionospheric Radio Wave Propagation of the Russian Academy of Sciences. He is an Honorary Gauss Professor of the Göttingen Academy of Sciences, Fellow of the Russian Academy of Natural Sciences, Fellow of the Electromagnetic Academy, and Honorary Professor of the China National Centre for Geological Exploration Technology.

References

Beiki, M., and Pedersen, L. B., 2010, Eigenvector analysis of gravity gradient tensor to locate geologic bodies: *Geophysics*, **75**, I37–I49.

Cox, L. H., Wilson, G. A., and Zhdanov, M. S., 2010, 3D inversion of airborne electromagnetic data using a moving footprint: *Exploration Geophysics*, **41**, 250–259.
Dransfield, M., 2010*a*, Potential field searchlights: Presented at 72nd EAGE Conference and Exhibition, Barcelona.
Dransfield, M., 2010*b*, Conforming FALCON gravity and the global gravity anomaly: *Geophysical Prospecting*, **58**, 469–483.
Henkel, H., and Reimold, W. U., 1998, Integrated geophysical modeling of a giant, complex impact structure – anatomy of the Vredefort structure, South Africa: *Tectonophysics*, **287**, 1–20.
Li, Y., 2001, 3-D inversion of gravity gradiometry data: 71st Annual Meeting, SEG, Expanded Abstracts, 1470–1473.
Li, X., and Chouteau, M., 1998, Three-dimensional gravity modeling in all space: *Surveys in Geophysics*, **19**, 339–368.
Li, Y., and Oldenburg, D. W., 1998, 3-D inversion of gravity data: *Geophysics*, **63**, 109–119.
Li, Y., and Oldenburg, D. W., 2003, Fast inversion of large-scale magnetic data using wavelet transforms and a logarithmic barrier method: *Geophysical Journal International*, **152**, 251–265.
Okabe, M., 1979, Analytical expressions for gravity anomalies due to homogeneous polyhedral bodies and translations into magnetic anomalies: *Geophysics*, **44**, 730–741.
Phillips, N., Nguyen, T. N. H., Thomson, V., Oldenburg, D., and Kowalczyk, P., 2010, 3D inversion modelling, integration, and visualization of airborne gravity, magnetic, and electromagnetic data – the QUEST Project: Presented at International Workshop on Electrical, Gravity and Magnetic Methods, Capri.
Pilkington, M., 1997, 3-D magnetic imaging using conjugate gradients: *Geophysics*, **62**, 1132–1142.
Portniaguine, O., and Zhdanov, M. S., 1999, Focusing geophysical inversion images: *Geophysics*, **64**, 874–887.
Portniaguine, O., and Zhdanov, M. S., 2002, 3-D magnetic inversion with data compression and image focusing: *Geophysics*, **67**, 1532–1541.
Tredoux, M., Hart, R. J., Carlson, R. W., and Shirey, S. B., 1999, Ultramafic rocks at the centre of the Vredefort structure – further evidence for the crust on edge model: *Geology*, **27**, 923–926.
Zhdanov, M. S. 2002, *Geophysical Inverse Theory and Regularization Problems*, Elsevier.
Zhdanov, M. S., 2009, New advances in 3D regularized inversion of gravity and electromagnetic data: *Geophysical Prospecting*, **57**, 463–478.
Zhdanov, M. S., Ellis, R. G., and Mukherjee, S., 2004, Regularized focusing inversion of 3-D gravity tensor data: *Geophysics*, **69**, 925–937.