

Evolving Inversion Methods in Geophysics with Cloud Computing

- a case study of an eScience collaboration

J Craig Mudge, Pinaki Chandrasekhar
Collaborative Cloud Computing Lab
School of Computer Science
University of Adelaide
Adelaide, SA 5005, Australia
craig.mudge@adelaide.edu.au
pinaki.chandrasekhar@adelaide.edu.au

Graham S. Heinson, Stephan Thiel
Discipline of Geology and Geophysics
School of Earth and Environmental Sciences
University of Adelaide
Adelaide, SA 5005, Australia
graham.heinson@adelaide.edu.au
stephan.thiel@adelaide.edu.au

Abstract— Magnetotellurics is a geophysics technique for characterisation of geothermal reservoirs, mineral exploration, and other geoscience endeavours that need to sound deeply into the earth – many kilometres or tens of kilometres. Central to its data processing is an inversion problem which currently takes several weeks on a desktop machine. In our new eScience lab, enabled by cloud computing, we parallelised an existing FORTAN program and embedded the parallel version in a cloud-based web application to improve its usability. A factor-of-five speedup has taken the time for some inversions from weeks down to days and is in use in a pre-fracturing and post-fracturing study of a new geothermal site in South Australia, an area with a high occurrence of hot dry rocks. We report on our experience with Amazon Web Services cloud services and our migration to Microsoft Azure, the collaboration between computer scientists and geophysicists, and the foundation it has laid for future work exploiting cloud data-parallel programming models.

Keywords- *Magnetotelluric Inversions, Cloud Computing, Parallelisation, Geophysical Modelling, Parallel Programming, Geothermal, South Australia*

I. INTRODUCTION

Computing resources needed for geophysical data processing and inversion have grown at a rate faster than can be handled by conventional desktop computers. The issues are twofold: first, memory is not sufficient to store information pertinent to the phenomena under investigation. Because the physical memory limits of 32-bit machines (and lower priced 64-bit computers) do not allow a full analysis, the problems are typically broken down into more manageable fractions. The second issue is one of time. Modelling and inversion of electromagnetic data is highly non-linear in its parameters, and thus most computations require significantly sized and repeated matrix operations. When implemented in a sequential program, these are increasingly taking many weeks or even months to

perform on a desktop computer. In common with many scientists, we find the parallel execution available on High Performance Computing (HPC) infrastructure available to us to be too hard to use without specialist training.

Our Collaborative Cloud Computing Lab (C3L) was set up in mid 2010 to address the computing challenges of electromagnetic geophysics with cloud computing and to grow an eScience capability in earth sciences more broadly. The university housing this collaboration supports an exciting geophysical precinct because the state of South Australia has some of the best hot rocks suitable for electricity generation from enhanced geothermal systems, has a vibrant minerals exploration industry, and is home to a massive new mine, Olympic Dam, owned by BHP Billiton, the world's largest mining company.

Because many problems in earth sciences are solved as inverse problems, we were pleased to focus our first efforts on inversion.

II. CLOUD COMPUTING AND eSCIENCE

Cloud computing presents compute and storage resources in an easy-to-use, flexible manner and as an elastic utility just as we acquire electricity, with a matching business model, namely pay-per-use. Because of the huge scale of the data centres owned by public cloud service providers, the cost of resources and storage is less than that of on-premises infrastructure, up to a factor-of-seven cheaper [1][2]. Finally, the fault tolerance provided by the operations software of the major cloud services providers is of value in our long-running inversion programs.

Parallel programming is receiving much more attention now because of (a) the need to program multi-core chips, as single-chip microprocessors hit the power wall, and (b) the ready availability of inexpensive parallel computing structures in the cloud.

The cloud hardware architecture of a warehouse-sized data centre housing tens of thousands of cheap computers, originally built for search and machine learning, is a parallel structure less familiar to High Performance Computing (HPC) users than shared memory multiprocessors, for example. It is thus useful to look towards programming models appropriate to this massive computing and data resource.

Message Passing Interface (MPI) has become the most popular parallel programming model and many important program libraries provide MPI implementations. However, when an algorithm causes a lot of message passing, programs written in MPI do not map well to the architecture used in the cloud with its low cost interconnect between processors.

Another category of programming languages with a natural expression of parallelism contains declarative languages. There is increasing interest in this category, because it can be naturally matched to the massive resources of a cloud data centre architecture. MapReduce, Hadoop, LINQ, DryadLINQ, and F# are examples [12]. These languages are influenced by functional programming, a core idea of which is avoidance of side effects. Because data-independent computations cannot interfere with each other, data-driven computation can readily be mapped to large-scale cloud resources. However, because most programmers are trained in imperative languages, it is often necessary to find an approach bridging to these new ideas, starting from an existing programming base, and this is the approach reported here. Nevertheless, this is just a bridging step en route to complete reimplementations of inversions in DryadLINQ or MapReduce.

In addition to mapping computationally expensive geophysics problems onto the cloud, we wished to build an eScience test bed developing cost-effective infrastructure anticipating a data deluge as geophysical instruments generate more data and to explore important systems design questions such as effective ways to move computation closer to its data. Moreover, as geophysics moves to discovery from mining electromagnetic, seismic, and gravity soundings, we need a machine learning platform to support our research. Finally, we need to use private, community, and public clouds, and learn to link to a collaborator who has a private cloud

with GPU accelerators without compromising sensitive data.

Public clouds were selected for several reasons -- they were already available, they make the best use of scale-driven effects (costs, fault tolerance, for example), and public cloud services give us the best way to gain from the innovation occurring in data centre design (energy efficiency, security, and cost-performance). We also value the automation of operations which lower our systems administration and data base administration costs. In order to learn how far this outsourcing of information and communication technology infrastructure can take us, we are a "No machines Lab" where the only machines we allow are those desktop, laptop, and tablet machines which access the cloud. Our source code control system SVN and document sharing systems are also in the cloud, not using the university infrastructure.¹

III. ELECTROMAGNETIC GEOPHYSICS

A. Electrical conductivity

The determination of subsurface electrical conductivity σ (or its inverse resistivity ρ) plays an important role in a variety of applications from tectonic evolution, mineral and geothermal exploration and groundwater studies. In contrary to other geophysical parameters, i.e. seismic wave velocities and density, Fig.1 shows that the electrical conductivity varies over more than seven orders of magnitude. The large spread in conductivity ensures sufficient contrast to image geological structures of elevated conductivities in an otherwise electrically resistive host rock (igneous and metamorphic host rocks in Fig. 1).

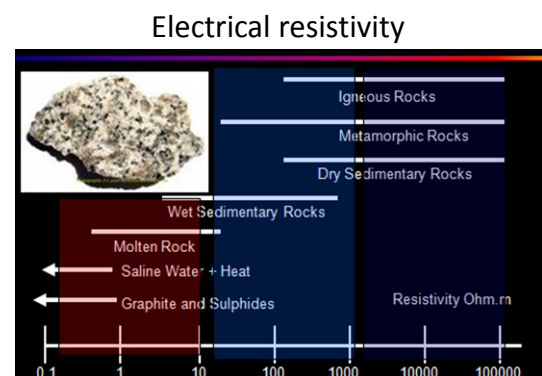


Figure 1. Electrical resistivity as a function of lithology. Dry igneous and metamorphic rocks show several orders of magnitude higher electrical resistivity than rocks with minor conducting agents, such as fluids, graphite and sulfides. In general resistivity also decreases with temperature.

¹ One exception to the non-machines rule will be the head node we are required to have when DryadLINQ is available on Windows Azure.

In general, earth's lithosphere is quite resistive (>1000 Ωm), but frequently zones of higher conductivity are observed. The causes of enhanced conductivities are diverse, but can be attributed in most cases with additional geological and other geophysical constraints [3]. Among the common agents for enhanced conductivities are minor conducting phases like graphite and metallic sulfides formed through metamorphic processes frequently involving shear [4]. Aqueous fluids distribute between rock grains of porous rocks and can significantly enhance conductivities if they are well connected and increase with temperature and salinity of the fluid. Fluids play a major role in the characterization of geothermal areas and also in the understanding of ore genesis.

B. Magnetotelluric method

The subsurface electrical conductivity is measured with natural and control-source electromagnetic (EM) methods (Fig.2). The depth of penetration δ increases with lower frequencies f of the varying EM signal and apparent resistivity ρ_a of the subsurface according to the equation given by

$$\delta[m] = 500 \cdot \sqrt{\rho_a f^{-1}} \quad (1)$$

The magnetotelluric (MT) ([5], [6]) method uses natural sources such as lightning strikes and solar magnetic storm interactions with the ionosphere with frequencies ranging from 10^4 - 10^5 Hz [5]. The frequencies relate to penetration depths of few tens of meters to hundreds of kilometers and the MT method is therefore suitable to a range of application from environmental and geothermal studies in the upper crust to mineral exploration in the upper/middle crust to lithospheric scale studies of crust and mantle interaction.

Fig. 3a illustrates the field setup of an MT station at Paralana, South Australia, which measures the natural variations of the earth's electric (E) and magnetic (B) field in orthogonal directions. Non-polarizing electrodes measure the potential difference in the north-south and east-west directions which yields the electric field strength. Depending on the target depth, the magnetic field is measured using either fluxgate magnetometers (frequencies between 10^{-1} to 10^{-5} Hz) or induction coils (10^4 to 10^{-3} Hz). Outliers are removed from the measured time-series of electric and magnetic fields and subsequently converted by a Fast Fourier Transform [8]. In the frequency domain the electric and magnetic field components are linked by the complex and frequency-dependent impedance tensor Z :

$$\begin{bmatrix} E_x \\ E_y \end{bmatrix} = \begin{bmatrix} Z_{xx} & Z_{xy} \\ Z_{yx} & Z_{yy} \end{bmatrix} \begin{bmatrix} B_x \\ B_y \end{bmatrix} \quad (2)$$

The impedance tensor components are commonly expressed as apparent resistivity $\rho_a(f)$ and phase $\phi(f)$ [9] indicating the bulk change in electrical properties with frequency, which according to (1) is related to depth. As a final step, the frequency-dependent impedance tensor data of all MT stations are used as input into a three-dimensional inversion [9]. Starting from a half-space with usually uniform conductivity, the resistivities of the model cells are changed until the misfit between observed and modelled data is minimized.

C Geothermal monitoring at Paralana, South Australia

Monitoring fluid injection and reservoir development of Enhanced Geothermal Systems (EGS) is one of the crucial steps in maximizing energy extraction in the renewable geothermal energy sector. The MT project at Paralana is aimed at using the MT method to better understand the processes involved during the stimulation of an EGS fluid-system. The initiation and propagation of the underground fluid-reservoir is dependent on several parameters, including the stress regime, geometry of pre-existing faults, and the reservoir lithology.

Currently, measurements of shear-wave splitting (SWS) defines the fracture characteristics of the EGS system, while micro-earthquake (MEQ) arrays are a standard tool to monitor the extent of the reservoir by locating seismic events associated with the fracture front during the injection process. Nevertheless, the MEQ array is not directly sensitive to the injected fluid itself and a spatial relationship between fracture development and fluid reservoir extent remains questionable.

MT stations are deployed across the injection site prior, during and after the injection to take a snapshot of

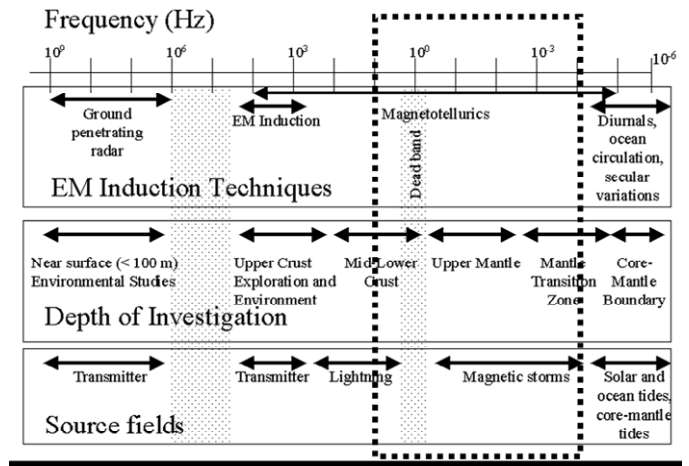


Figure 2. Overview of EM surveying techniques. The frequency (f) range determines the depth of investigation from a few meters to hundreds of kilometres. For high frequencies a transmitter is used, and natural EM sources for $f < 10^3$ Hz.




Figure 3a. Field setup of an MT station at Paralana, South Australia which measures the natural variations in the Earth's Electric and Magnetic fields in orthogonal directions.

the injected fluid reservoir extent over time. A survey prior to the injection provides a reference model to allow identification of even smaller reservoirs. Furthermore, the monitoring array will also provide a 3D resistivity background image that will help in defining fault locations and basement depth.

The pre-injection surveys have been completed with a total of 56 stations positioned along two main north-south and east-west transects across the injection well.

The data have been inverted for three-dimensional resistivity structure using the cloud based 3D inversion code and for speed comparison with the serial version.

D. Processing steps

The processing is in two stages. The first takes the MT time series data which has been logged in multiple stations in the field. After a cursory inspection by hand, the data are run through the BIRRP step [10]. BIRRP is a robust statistical method that removes outliers along with converting the data, via a Fast Fourier Transform, into sets of impedance tensors. Each set of impedance tensors correspond to a particular desired evaluation frequency. The impedance tensors in turn become the input to the second stage, the MT inversion step as shown in Fig. 4.

MT inversion computations are dominated by sensitivity matrix and forward modeling calculations. These calculations are carried out for each evaluation frequency which are independent of each other and are parallelized by our cloud implementation. The inversion processing is shown in Fig. 4, where each of the steps to compute the sensitivity matrix, the forward misfit, and the smoothness misfit are distributed across multiple processors.

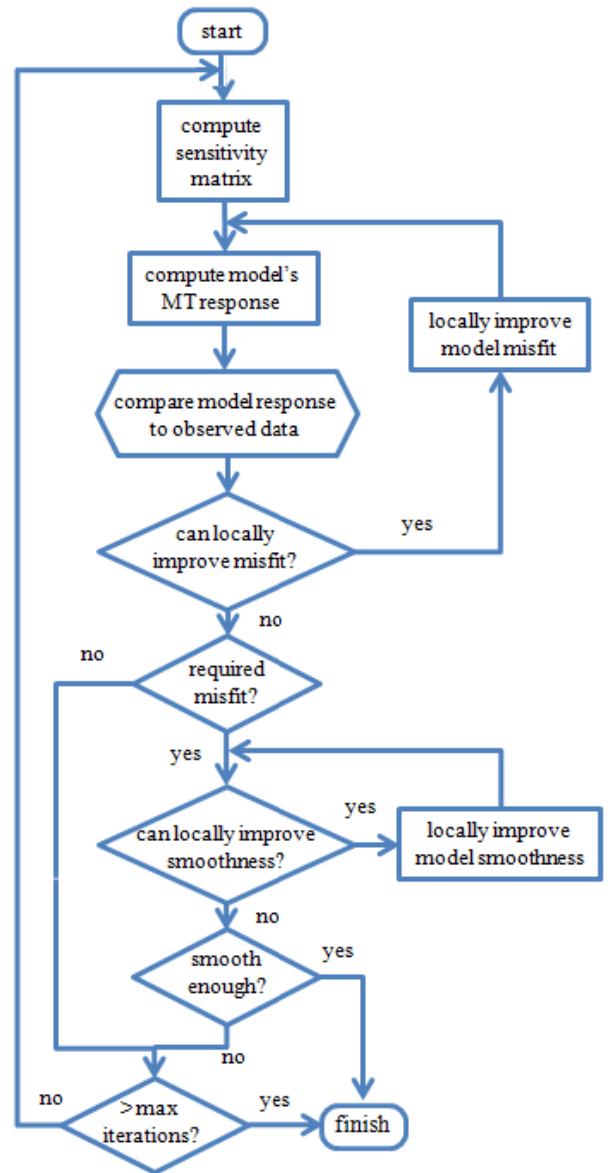


Figure 4 Details on the inversion showing the steps which are distributed across multiple processors

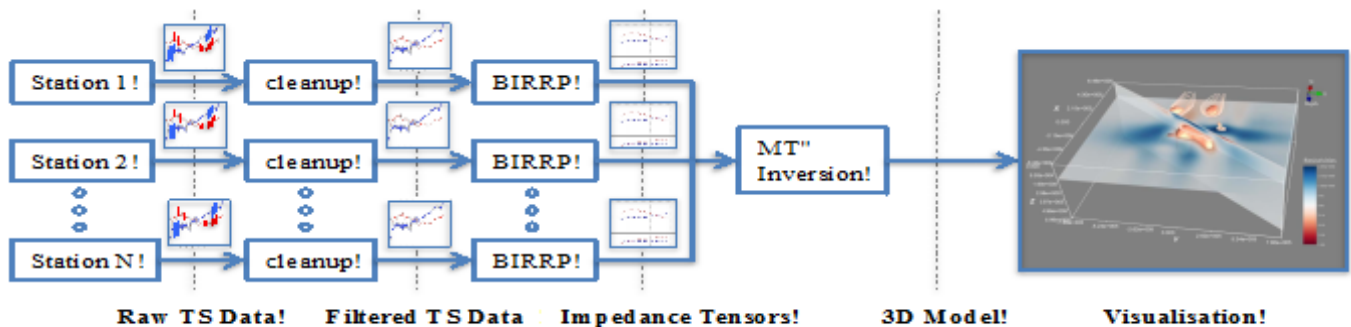


Figure 3b Overview of the MT processing steps

IV. CLOUD ARCHITECTURE

There are two sections of the code: the parallelized FORTRAN code and a web application. The FORTRAN program, 22,000-lines in length, is widely used in geophysics. The parallelized code abstracts the logic which computes the sensitivity and forward modelling calculations per evaluation frequency. The computations for each frequency are run in parallel. To improve the extent of its adoption, and to improve the user experience, we embedded the resulting program in a web application. The web application is implemented using Django, a popular web application framework, written in Python, which follows the model-view-controller architectural pattern.

The remote tasks are executed in remote EC2 instances standing by for this particular inversion.² A complete EC2 instance is dedicated for any remote task. Also, there is no shared memory between processes with the necessary state variables required for the separate tasks persisted on file and transferred between the instances.

A. Amazon EC2 Implementation

The entire distributed process is implemented on the Amazon EC2 cloud using standard available EC2 instances. The whole Inversion process is further packaged as a MT 3D inversion software product which is accessible anywhere through a secure login via a common-or-garden browser. Amazon EC2 machines are automatically acquired on demand only for the duration of the inversion and released soon after. Relevant results and configuration data are stored in Amazon S3 storage.

Our lab is exploring different types of cloud services in order to understand costs and benefits of using cloud resources for large scale eScience applications. Microsoft Azure is a Platform as a Service (PaaS) model whereas Amazon Web Services is an Infrastructure as a Service (IaaS) model. We will also use community cloud services, such as those utilities provided by our federal government for university research. Our migration to Azure is complete and has the architecture shown in Figure 6. Note the symmetry of the Amazon and Azure architectures.

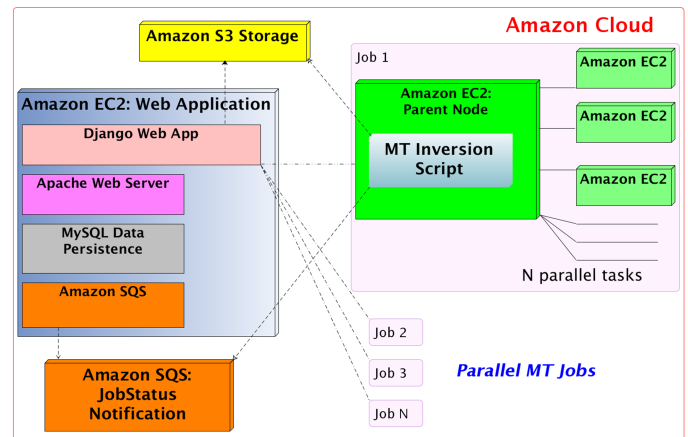


Figure 5. Architecture overview of MT running on Amazon Web Services cloud.

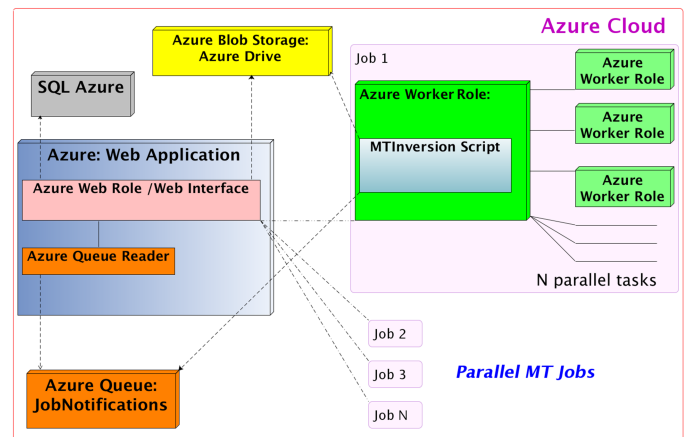


Figure 6. Architecture overview of MT running on Windows Azure cloud.

V. RESULTS

A. Speedup from parallelisation

The computation time versus a range of frequencies used in the inversion are shown in Fig 7. On average, with our cloud-based implementation, the modelling time was reduced by a factor of four reducing the computation time from weeks to days.

B. Taking chunks of FORTRAN is achievable in a timely manner

A straightforward approach to converting legacy code to a parallelized distributive process is to transfer all control logic and optimization functions of the 3D Inversion to the control script/master process. The idea is

² The multicore nature of some processors on the EC2 instances has not been exploited in order to avoid distraction from our goal.

to reduce the FORTRAN code to computational functions and stripped of as much decision logic as possible [11]. In order to do this we segregate optimization, global and local variables in the FORTRAN code to begin with. There are certain optimization logic decisions which if embedded deep inside the FORTRAN makes it very hard to be stripped out. These variables are not removed but are made sure to be consistent and up to date in both the FORTRAN code and the master control script. The idea is to not rework any logic or algorithmic technique in the FORTRAN code which would require it to then go through a rigorous testing cycle.

The control script/master process is implemented using Python. It controls the forward and inverse optimization loops, remote task delegation of parallel & independent tasks, aggregation of remote task outputs, File I/O operations and upload of results, errors & log files onto Cloud storage.

There were 18 FORTRAN routines totalling 22497 lines of code. Four of the 18 routines are unchanged. The Python script has 600 lines. Overall the project took

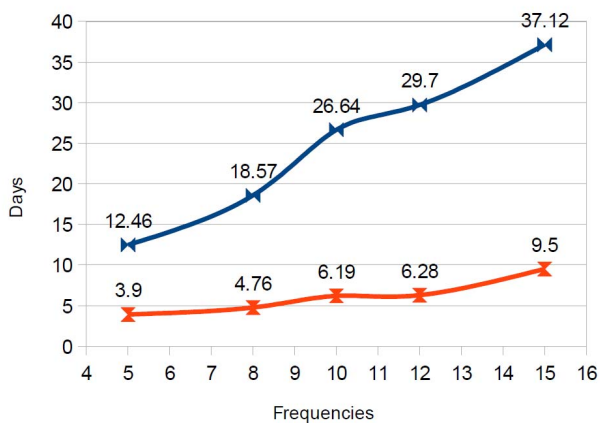


Figure 7. The upper curve shows the execution times for the serial program in days, while the lower curve shows the faster parallel times. The speedup increases with the number of frequencies.

six months. To help those undertaking similar efforts in the future, the time taken (by the one person working on the programming) is given.

After a month of start up actions – learning about MT, studying the sequential program and learning to configure and run Amazon cloud instances, three months was spent constructing the parallel version. Writing the web application and learning the geologists’ user preferences took another two months. The time spent are supported by regular status reports and the dates of documents and emails produced during the project

C. Capability building in our group

The first portion of the processing in Fig. 3b is cleaning data logged in the field and removing outliers using BIRRP (Bounded Influence, Remote Reference Processing) [10] a robust statistical method. Because the BIRRP step is independent for each station, we can process station data in parallel. We planned to achieve this parallel execution by using a workflow language, but not until we had finished the parallelization of the inversion code. However, a team member, Jared Peacock, a geophysics doctoral student was unwilling to wait. Using a newly acquired workstation whose CPU chips had 8 cores and used Python(x,y) 2.6 to write a script for the parallel multi-core execution. This is the type of result we had hoped for in creating a true eScience group, where geophysicists learn about parallel computing and do their own implementation.

A second example of capability building is provided by the geophysicists using the parallel inversion routines. They were able to do so, including making necessary changes to model parameters, without low-level hand-holding by the computer scientists. This resulted from two factors: the user oriented design of the web application running the parallel execution and the effort of the geologists to understand the essence of our research theme.

D. Improved geophysics practices

The primary outcome is that we now can undertake 3D inversion of magnetotelluric (MT) data in a way that is repeatable and testable rather than as a single experiment in which we had little knowledge of model space outcomes.

VI. NEXT STEPS

A. MT inversion

We will optimise the existing deterministic algorithms in which data are mapped onto a single model space. Thus, we will be able to do this quicker, and be able to test several hypotheses and constraints. Secondly, the ability to distribute computation across many thousands of nodes will allow us to use non-deterministic algorithms to produce stochastic models that map confidence limits of model parameters. These include many global sampling

We will explore new approaches to searching the space of matching the observed response to a model response, such as evolutionary algorithms, that require many independent calculations, which would be infeasible on most desktops.

B. Cloud infrastructure and programming model

Increasing the model size (increasing the granularity) by an interesting factor takes the memory size needed

(predominantly used by the sensitivity and beta matrices) from their present 7 Gbytes to 60 Gbytes. We will be exploring solutions to the maximum memory sizes placed on us by public cloud service providers.

Purpose Distributed Data-Parallel Computing Using a High-Level Language, OSDI'08 Eighth Symposium on Operating Systems Design and Implementation, accessed on 22 June 2010 at http://www.usenix.org/events/osdi08/tech/full_papers/yy_y/yy_y.html/

ACKNOWLEDGMENT

Financial support came from a Jim Gray Seed Grant from Microsoft Research. Other funding has been from IMER (Institute for Minerals and Energy Research, University of Adelaide); South Australian Centre for Geothermal Energy Research; and Primary Industries Resources (SA). Wei Wang helped with systems programming expertise on Windows and Windows Azure, Jared Peacock helped us better understand the MT processing steps, discussions with Brad Alexander and Andrew Wendelborn improved the systems aspects, and the eScience guidance and encouragement from Catharine van Ingen was invaluable.

REFERENCES

- [1] Armbrust, M, A Fox, R Griffith, A D. Joseph, R H. Katz, A Konwinski, G Lee, D A. Patterson, A Rabkin, I Stoica and Mi Zaharia, 2009, *Above the Clouds: A Berkeley View of Cloud Computing*, Technical Report No. UCB/EECS-2009-28, accessed on 22 June 2010 at <http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html>
- [2] James Hamilton. Internet Scale Storage. Keynote address, SIGMOD. 2011. Slides available at http://mydirona.com/jrh/TalksAndPapers/JamesHamilton_Sigmod2011Keynote.pdf
- [3] G.S. Heinson, N.G. Direen, and R.M. Gill, "Magnetotelluric evidence for a deep-crustal mineralizing system beneath the Olympic Dam iron oxide copper-gold deposit, southern Australia," *Geology*, 2006, vol. 34, pp. 573-576.
- [4] G. Nover, "Electrical Properties of Crustal and Mantle Rocks: A Review of Laboratory Measurements and their Explanation," *Surveys in Geophysics*, 2005, vol. 26, pp. 593-651.
- [5] L. Cagniard, "Basic theory of the magneto-telluric method of geophysical prospecting," *Geophysics*, SEG, 1953, vol. 18, pp. 605-635.
- [6] A. Tikhonov, "The determination of the electrical properties of deep layers of the Earth's crust," *Dokl. Acad. Nauk. SSR*, 1950, vol. 73, pp. 295-297.
- [7] K. Vozoff, "The magnetotelluric method," in *Electromagnetic methods in applied geophysics – Applications*, M. Nabighian, eds., Society of Exploration Geophysicists, 1991, pp. 641-711
- [8] F. Simpson and K. Bahr, *Practical Magnetotellurics*, Cambridge University Press, 2005.
- [9] W. Siripunvaraporn, G. Egbert, Y. Lenbury, and M. Uyeshima, "Three-dimensional magnetotelluric inversion: data-space method," *Physics of The Earth and Planetary Interiors*, 2005, vol. 150, pp. 3-14.
- [10] Alan D. Chave and David J. Thomson. Bounded influence estimation of magnetotelluric response functions, *Geophysics.J.Int.*, 157, 988-1006, 2004.
- [11] Pinaki Chandrasekhar and J.Craig Mudge. Restructuring a FORTRAN program for parallel execution in the cloud. Collaborative Cloud Computing Lab. University of Adelaide. TN 2011-3
- [12] Yu, Y, M Isard, D Fetterly, M Budiu, Ú Erlingsson, P Kumar Gunda and J Currey, 2008, *DryadLINQ: A System for General-*