J. Geomag. Geoelectr., **45**, 1013–1026, 1993

Two-Dimensional Nonlinear Magnetotelluric Inversion Using a Genetic Algorithm

M. E. EVERETT and A. SCHULTZ

Institute of Theoretical Geophysics,
Department of Earth Sciences,
University of Cambridge, Cambridge CB2 3EQ, U.K.

(Received June 24, 1993; Revised September 13, 1993; Accepted September 14, 1993)

We have developed a nonlinear magnetotelluric inversion based on a standard finite difference TE/TM mode forward solution, including static distortion effects, and a new genetic algorithm for general functional optimisation and hypothesis testing. We have used this to invert a subset of the COPROD2 data in terms of best-fitting 2-D electrical conductivity distributions. Our optimal electrical conductivity model, defined by 66 electrical conductivity parameters and 20 static shift coefficients, attains an rms misfit of 1.48, for standard errors in the data of at least 10% in apparent resistivity and 3° in phase. This may represent the minimum level of misfit given this coarse parameterisation of the earth. The optimal model contains certain features, including the North American Central Plains conductivity anomaly and a surface layer of 1000 S conductance, that are consistent with previous electromagnetic inversions and the local geology. The global optimisation took ~12 days to compute on a ~20–40 Mflop (million floating point operations per second) computer. We have chosen not to seek a smooth model consistent with the data, a task well handled by existing, faster regularized inversions, but instead to use the genetic algorithm for the more demanding task of extracting the global best-fitting conductivity model.

1. Introduction

The North American Central Plains (NACP) conductivity anomaly within the Trans–Hudson orogen is thought to be the electromagnetic expression of a geosuture zone formed during the collision of the Superior and Churchill Provinces in the Early Proterozoic (CAMFIELD and GOUGH, 1977). The NACP anomaly was discovered by GDS techniques (REITZEL *et al.*, 1970) and its location beneath the Phanerozoic Williston Basin has been refined by subsequent analyses of additional electromagnetic data, including the COPROD2 magnetotelluric (MT) data set (JONES and CRAVEN, 1990).

The full COPROD2 data set, 35 MT sites collected over a 407 km traverse (see JONES and CRAVEN, 1990, for location map) with 40 responses in the range 384 Hz to 1820 s, has been distributed by A. G. Jones of the Geological Survey of Canada to workers prior to the 1992 MT Data Interpretation Workshop in Wellington, New Zealand (JONES, 1993). The distribution is intended to form the basis for cross-validation, comparison and development of various 2-D MT inversion codes. As part of this multinational collaborative effort, we have developed a nonlinear 2-D inversion based on a genetic algorithm and a standard finite difference TE/TM mode forward solution, and have inverted a subset of the COPROD2 data for an optimal (*i.e.* minimum misfit) 2-D electrical conductivity distribution.

In particular, we present the results of a search for a 2-D electrical conductivity model whose response minimises the misfit to the COPROD2 observed TE/TM apparent resistivity and phase data, given the fewest *a priori* assumptions about the structure this model. Elsewhere in this issue, the reader will find results of regularised inversion schemes which produce maximally

smooth models. Unlike existing regularised schemes, the approach taken here is fully nonlinear, and the intent is to find the global minimum misfit model.

Although our algorithm can handle general nonlinear functionals of electrical conductivity, such as smoothness, we chose to minimise misfit since our eventual aim is to develop analytical and numerical algorithms capable of determining the best possible fit between responses of multidimensional electrical models of the earth and arbitrary collections of MT observations. We are motivated by the desire to develop a multidimensional analogue of the D^+ 1-D minimum misfit algorithm (PARKER, 1980; PARKER and WHALER, 1981). Such algorithms are important not necessarily as techniques to construct geologically tenable models of the subsurface, but rather to provide an absolute lower bound on the misfit, so that the goodness-of-fit of arbitrary models can be meaningfully evaluated. Furthermore, the genetic algorithm is such that one must place prior bounds on all model parameters. By careful selection of *a priori* information, nonlinear hypothesis testing is possible. One may, for instance, ask whether *any* 2-D model exists that fits the data, given particular bounds placed on one or more model parameters. Finally, fully arbitrary objective functionals (e.g. misfits) may be easily defined, making it possible to direct the search in an extremely flexible manner.

In general, finding best-fitting multidimensional electrical conductivity models is problematic. There are uncertainties inherent in numerically modelling forward responses, and the use of finite dimensional model spaces to represent the earth's electrical structure is clearly inadequate. Additional difficulties arise when global search algorithms are employed to find the minimum of the misfit functional. For example, large amounts of computer time are often required to search model space and the possibility exists for the optimisation routine to become trapped in a local minimum. Global search algorithms based on numerical solutions to the forward problem therefore provide only an approximate value for the true minimum misfit. Furthermore, the result depends strongly on the degrees of freedom permitted by the earth's electrical structure.

2. The Forward Solution

The earth is assumed to possess an electrical conductivity structure which varies with position in the vertical and one horizontal co-ordinate, but is invariant in the orthogonal horizontal (strike) direction. Furthermore, any lateral conductivity variations are contained within a central, inhomogeneous region, outside of which the conductivity depends only on depth. In such cases, Maxwell's equations for an incident, down-going plane wave source reduce to a pair of independent scalar differential equations in the along-strike electric (TE mode) and along-strike magnetic (TM mode) fields. The forward code we have developed to solve the two scalar equations is based on the finite difference method and is similar to that of SMITH and BOOKER (1991), e.g. the discretisation of the governing equations and the boundary conditions are equivalent.

In our formulation, the finite difference mesh is truncated at large lateral and vertical distances from the inhomogeneous region, where the earth is assumed to be locally 1-D, and boundary conditions are applied there. For both TE and TM modes, the fields at the sides of the mesh are set to the known, exact 1-D solutions. The impedance, or ratio of a field to its vertical derivative, at the bottom of the mesh is fixed to that of the half-space below.

The boundary conditions imposed at the top of the mesh are fundamentally different for the two modes. In the TM mode, the surface of the earth coincides with the top of the mesh and a unit magnetic field is specified as the source. However, since the electric field in air is sensitive to the structure inside the earth, for TE mode computations the finite difference mesh must be extended upwards from the surface. At the top of the air layer, set to $\sim 10^4$ km, we fix the vertical gradient of the electric field such that the associated horizontal magnetic field is unity. At the sides of the mesh, we force the electric field to increase linearly with height above the surface of the earth, in a manner that is consistent with the other boundary conditions.

The linear systems of equations in the unknown nodal values of along-strike electric field (TE mode) and magnetic field (TM mode) are solved on the Institute of Theoretical Geophysics Convex C210 (~20–40 Mflop) computer using Veclib complex band matrix routines. Meshes used in the work reported here contain as many as 91×91 nodes for the longest period, leading to linear systems of up to 8281 equations. Apparent resistivities and phases are derived by numerically differentiating the computed field values at nodes near the surface of the earth, and are linearly interpolated onto the measurement sites. The apparent resistivity at each site is then scaled by a frequency independent static shift coefficient to account for the effects of near-surface conductivity heterogeneities too small to be individually modelled.

It is well-known that errors in the finite difference solutions to electromagnetic equations diminish as the node density increases near large variations in electrical conductivity. However, our forward solver must contend with an immense variety of irregular conductivity distributions as the inversion proceeds. Adapting the finite difference mesh for each forward problem is not a recommended strategy in our case, since inconsistent numerical errors between forward problems are likely to result, biasing the misfit in an unpredictable manner. We experimented with an adaptive meshing scheme in which the uniform node spacing depended on both the period and an averaged conductivity of the model, but abandoned this approach for the reason just stated. We prefer consistent solutions, though they may be individually less accurate than those computed on specially tailored meshes. As a result, our meshes are uniform, with the node spacing depending only on the period. In particular, we enforce a mesh density of exactly 5 nodes per skin depth, where the latter is calculated using the geometric mean of the smallest and largest permissible values of electrical conductivity. These values are respectively the infimum of the lower bounds and the supremum of the upper bounds placed on individual model parameters.

3. The Genetic Algorithm

Genetic algorithms (HOLLAND, 1975) have been used recently to solve nonlinear optimisation problems in many disciplines, including 1-D seismic inverse problems (GALLAGHER *et al.*, 1991; SCALES *et al.*, 1992; SAMBRIDGE and DRIJKONINGEN, 1992; JIN and MADARIAGA, 1993), and the 1-D magnetotelluric inverse problem (SCHULTZ, 1992; SCHULTZ *et al.*, 1993). Since our GA will be described in detail for 1-D conductivity studies (*Genetic algorithms for nonlinear magnetotelluric inversion and hypothesis testing*, manuscript in preparation), we merely outline the essential strategy and briefly discuss some of the novel features of our implementation.

The basic idea behind the genetic approach for optimising nonlinear functionals is to mimic the manner by which biological organisms evolve to produce more successful organisms. The algorithm is an iterative one. An initial population of Q models is specified. Each model, or "chromosome", is either randomly generated or read in as a set of N real numbers, which are the model parameters, or "genes". In our case the model parameters are logarithms of electrical conductivities ($\log \sigma$) of the elements of the 2-D electrical model and, optionally, TE and TM static shift coefficients at each measurement site. The values of individual model parameters are constrained to lie within user-specified upper and lower bounds. The forward problem is solved for each of the models, and the misfit to the data of each model response is computed. Based on the misfit information and constrained by probablistic rules, the population of Q models is refined, and the forward problems are solved again. The process is repeated until the misfit can no longer be reduced. The evolution of the population is directed such that models which fit the data well are given a good chance of surviving into the next generation, while models which do not fit rapidly die out. SAMBRIDGE and DRIJKONINGEN (1992) describe the reproduction, crossover and mutation stages involved in a single iterative step. We perform an identical crossover to Sambridge and Drijkoningen, but have developed new reproduction and mutation strategies.

3.1 Reproduction

Natural selection, or survival of the fittest, occurs during the reproductive stage. After the misfits for the current generation of models have been evaluated, a survival tournament is held. Random pairs of models are drawn from the population, but only one model from each pair will survive to the next generation. Normally, the survivor is the better fitting of the two. However, to facilitate exploration of a wide range of model space, the conductivity model with the higher misfit occasionally goes forward to the next generation. The choice of which model survives is determined by the flip of a coin and an input parameter called the tournament probability. If tournament probability is nearly unity, then the better fitting model nearly always survives and the population is inexorably driven towards a few promising neighbourhoods in model space. A random walk through model space results as tournament probability approaches its lower limit of 0.5. Our experience in 2-D MT inversion is that a tournament probability set between 0.8–0.9 leads to adequate global coverage of model space yet provides fine scale probing of the objective function in the vicinity of a local minimum.

3.2 Mutation

Without mutation, crossover (SAMBRIDGE and DRIJKONINGEN, 1992) is the only stage at which new models are created. Crossover is a mechanism by which genetic material (values of individual model parameters) is transferred between pairs of chromosomes. However, crossover alone is not enough to ensure a rich gene pool, as no new genetic material is introduced into the gene pool, and in the absence of mutation the diversity of the models deteriorates within a few generations: a situation we have termed saturation. In cases where the mutation rate is set too low we have observed the genetic algorithm to converge prematurely to an incorrect value, probably not even a local minimum of the objective function. There is simply not enough genetic diversity to further explore model space. As the mutation rate is increased, the model search becomes increasingly a random walk and, if set too high, Monte Carlo behaviour once again governs, inhibiting convergence.

Saturation can be avoided by increasing the size Q of the population, which provides more genetic material for the algorithm to manipulate in its search for the optimal model. However, this approach is unsatisfactory if the forward problem is costly to evaluate since adding new members to the population requires additional solutions the forward problem at each generation.

To stave off saturation and yet maintain a respectable convergence rate, while restricting $Q \sim 100$ to keep the number of forward problems within practical CPU-time limitations, we decided to implement a floating mutation mechanism. Individual mutations are performed in the manner of SAMBRIDGE and DRIJKONINGEN (1992), but the novel concept is that the mutation rate is dynamically controlled between generations by the genetic diversity of the population. For example, if genetic diversity at the end of a generation is high, the mutation rate is automatically lowered. Conversely, the mutation rate is raised whenever the population diversity falls below some critical value. The diversity of a single model parameter is measured by the number of different values that the parameter assumes within a generation, normalised by the total number of models Q. The diversity of the entire population is simply the sum of these individual diversities, divided by the total number of model parameters N. Thus, model diversity is a statistic which ranges from zero to unity. We have obtained good results when the mutation rate is adjusted to keep the model diversity statistic approximately 0.02–0.1.

4. Synthetic Inversion

We have tested the GA on a synthetic data set generated using the same TE/TM mode forward solver that is incorporated in the genetic algorithm. The synthetic study is therefore a true test of the performance of the GA; we are not testing the performance of the forward

code. Foward modelling errors and the finite dimensionality of the synthetic model are essentially irrelevant, while the target misfit is exactly zero. The electrical conductivity model we aim to recover is a uniform, outcropping rectangular conductor (σ=0.1 S/m) 30 km deep and 60 km wide embedded in a uniformly resistive half-space (σ=0.02 S/m). The response to this model is a set of TE and TM mode apparent resistivities and phases.

We define the objective functional Φ to be the rms misfit

$$\Phi = \left(\frac{1}{M} \left\| \frac{\rho_a - \rho_a^{obs}}{\Delta\rho_a^{obs}} \right\|^2 + \frac{1}{M} \left\| \frac{\phi - \phi^{obs}}{\Delta\phi^{obs}} \right\|^2 \right)^{1/2}, \tag{1}$$

where M is the number of degrees of freedom. The right hand side of (1) contains the sum of two Euclidean norms,

$$\|f\| = \left(\sum_{i=1}^{N_x} \sum_{j=1}^{N_T} |f(x_i, T_j)|^2 \right)^{1/2}, \tag{2}$$

one which measures the distance between the computed ρ_a and observed (or synthetic) ρ_a^{obs} apparent resistivities, and one which measures the distance between computed ϕ and observed (synthetic) ϕ^{obs} impedance phases. The quantities $\Delta\rho_a^{obs}$ and $\Delta\phi^{obs}$ are the standard errors in apparent resistivity and phase, respectively. There are N_T periods (T), N_x measurement sites (x) and 2 modes (TE and TM); therefore the number of degrees of freedom M=$4N_T \cdot N_x$, where the additional factor of 2 reflects the fact that apparent resistivity and phase, real quantities derived from the complex surface impedance measurements, are independent degrees of freedom. The objective functional (1) attributes an equal weighting to each datum regardless of its mode (TE or TM), location along the traverse, period, and type (ρ_a or ϕ). The expected value of the rms misfit, denoted E(Φ), is unity. In the special case of synthetic data without noise added, $E(\Phi)$=0.

Table 1. The locations and periods of the TE and TM mode apparent resistivity and phase COPROD2 data used in the inversion. The locations are measured in kilometers east of the town of Macoun, Saskatchewan, which is 18.87 km west of longitude 103°W.

positions	periods
−55.7 km	85.33 s
−45.8 km	170.7 s
−35.0 km	341.3 s
−25.9 km	682.7 s
−14.6 km	
−5.9 km	
4.9 km	
22.8 km	
41.8 km	
54.5 km	

Synthetic TE/TM mode apparent resistivities and phases were generated for the test model, at the 10 locations and 4 periods shown in Table 1, hence the total number of degrees of freedom M=160. The conductor in the test model is centered about position 0.0 km, the longitude of Macoun, Saskatchewan. Standard errors of $\Delta\rho_a^{obs}/\rho_a$=10% and $\Delta\phi$=3° are assumed. These standard errors are similar in magnitude to those typically estimated from actual data. The synthetic data and standard errors were then inverted using the genetic algorithm with population size Q=100, tournament probability 0.85, and crossover probability 0.85. The latter was chosen on

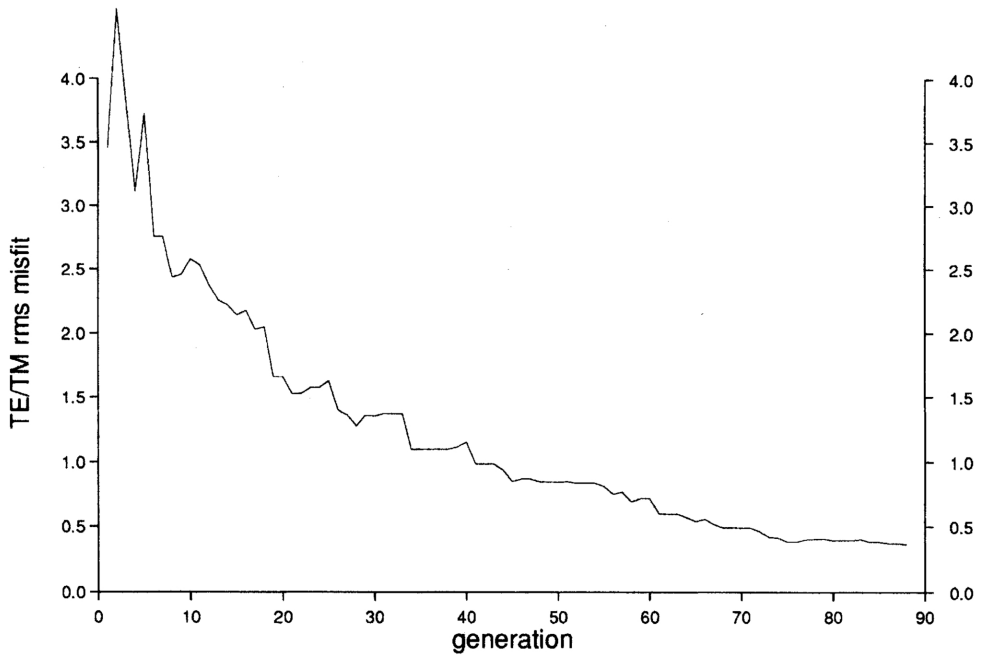

Fig. 1. The joint TE/TM mode rms misfit Φ of the best-fitting model of each generation of the synthetic inversion. The run was terminated after 88 generations with a best misfit of 0.37. Synthetic TE and TM mode apparent resistivities and phases for the test model described in the text were computed at each combination of period and position given in Table 1.

the basis of prior experience inverting 1-D MT data using the genetic algorithm. We maintained a fixed mutation probability 0.002 throughout, *i.e.*, the floating mutation rate mechanism was turned off.

The embedded electrical conductor is parameterised by a fixed, uniform 4×4 grid of constant conductivity rectangular blocks. Two additional model parameters include the conductivity of the terminating half-space beneath the inhomogeneity and the conductivity of the semi-infinite layers to the left and the right of the inhomogeneity. We found it convenient, although not necessary, to keep the conductivities of the left and right semi-infinite layers equal. The static shift option was disabled, *i.e.* all static shift coefficients remained unity throughout the run. Hence, there are N=4×4+2=18 independent model parameters. To reduce the search of model space, and to ensure that the finite difference meshes are of manageable size, individual model conductivities were constrained within the bounds 0.01–1.0 S/m.

The initial population was generated randomly. The best fitting randomly generated model achieved a misfit of ~3.5 compared to the synthetic data, using formula (1). The misfit of the best-fitting model as a function of generation for the entire run is plotted in Fig. 1. The misfit was still decreasing steadily when CPU-time limitations required that the run be prematurely terminated after 88 generations at which point the misfit was 0.37. Over the entire run, TE and TM mode responses were evaluated at four different periods for a total of $88Q$=8.8×10^3 different conductivity models, corresponding to a total of 7.0×10^4 finite difference matrix inversions. At approximately 750 matrix inversions per C210 CPU-hour, the run took 93 CPU-hours to complete.

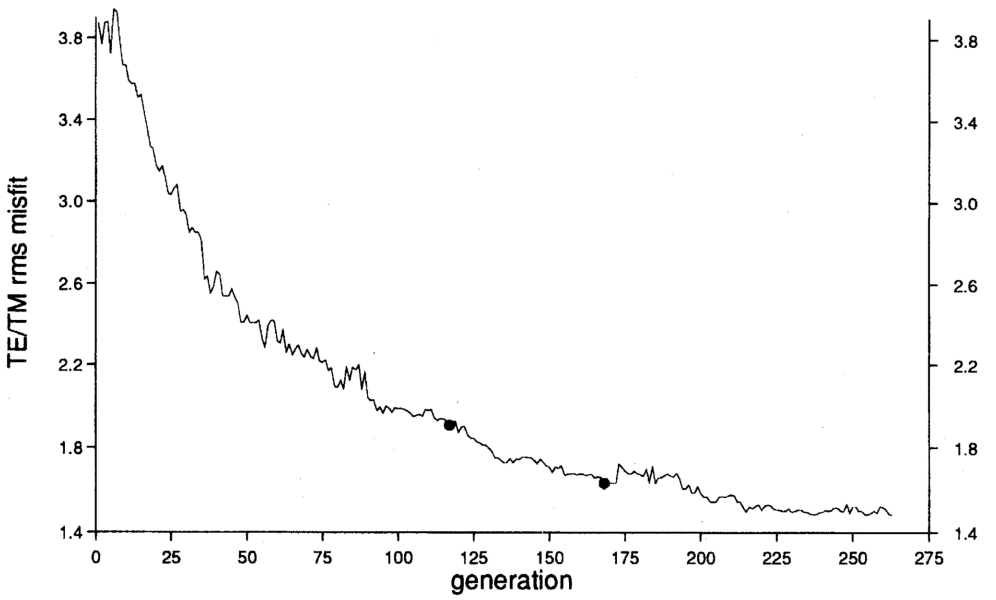

Fig. 2. The joint TE/TM mode rms misfit Φ of the best-fitting model of each generation of the COPROD2 inversion. The run was terminated after 263 generations with a best rms misfit of 1.48. The solid circles at generations 117 and 168 mark the widening of the bounds on permissible electrical conductivities and the onset of static shift coefficients as model parameters, respectively.

5. COPROD2 Inversion

The genetic algorithm was then run for the COPROD2 data set (see Table 1 for the positions and periods inverted). The standard errors on ρ_a were taken to be the greater of 10% or the standard error quoted in the distributed data set. Similarly, the minimum standard error on ϕ was set to 3°. The population size, tournament probability, crossover probability and bounds on the values of individual model conductivities were kept as they were during the synthetic inversion. Initially, the static shift option was turned off in order to keep the global search as simple as possible during the early iterations. The size of the inhomogeneous region was extended to 120 km wide by 60 km deep and was parameterised by a uniform 8×8 grid, hence the total number of model parameters is N=66. The initial population consisted of a 0.1 S/m half-space and $Q-1$ randomly generated models. The dynamic mutation rate was turned on.

The genetic algorithm ran for 263 generations, which took 278 C210 CPU-hours. After 117 generations, the lower bound on individual model conductivities was decreased to 0.002 S/m since it was observed that 50% of the block conductivities were at the minimum value of 0.01 S/m. After 168 generations, the static shift option was switched on, increasing the total number of model parameters to N=86. The rms misfit of the best-fitting model versus generation is shown in Fig. 2. The initial 0.1 S/m half-space achieved a misfit of 3.87. The expected value E(Φ) was not attained in 263 generations, and the misfit had stabilised when the run was terminated. The best rms misfit is 1.48. The additional degrees of freedom gained by switching on static shifts brought the misfit down to 1.48 from 1.63.

The responses of the best-fitting model are compared to the original data set in Fig. 3. The most notable systematic (site-independent) discrepancies are in the short period TM mode phase and the TE mode apparent resistivity for period 170.7 s. These systematic biases may be explained by the lack of a sufficiently thin surface layer in our parameterisation, or by 3-D effects in the data. Alternatively, the biases may be due to the misalignment of the finite difference meshes with conductivity block boundaries (we essentially solve a different conductivity model for each frequency). This effect can be removed by ensuring that, for each period, every conductivity block boundary coincides with either a row or column of nodes in the finite difference mesh. We highly recommend the adoption of this strategy in the future. There may also be numerical problems introduced by inconsistencies in the linear interpolation of the computed apparent resistivities and phases onto measurement sites, since the horizontal locations of finite difference nodes vary with period.

The evolution of the best-fitting model at various stages of the inversion is displayed in Fig. 4. The electrical conductivity of the left and right semi-infinite layers does not change after generation 27, suggesting that the misfit is very sensitive to this model parameter. By contrast, the conductivity of the terminating half-space changes from generation 55 to 117. This suggests that the responses are insensitive to the terminating conductivity. Inside the inhomogeneous region, the electrical conductivity in the first few rows near the surface is more or less fixed by generation 55.

The reduction in misfit afterwards appears to be due to fluctuations in the conductivity of the lower part of the inhomogeneous region and the terminating half-space. There are several intriguing conductivity structures in the upper part of the inhomogeneous region. The most persistent feature is a layer of moderately high conductivity across the top of the model; this has clearly established itself by generation 27. At generation 117, the average conductance across the topmost row of the inhomogeneous region is 1040 ± 100 S, corresponding to a 7.5 km thick layer of conductivity $\sigma \sim 0.15$ S/m. By generation 55, a single, highly conductive (~ 1 S/m) block appears in the upper central part of the inhomogeneous region. This feature persists throughout the remainder of the run. The overall background conductivity of the inhomogeneous region and terminating half-space is quite low.

6. Forward Modelling Error

Following the COPROD2 inversion, we analysed the numerical errors in forward responses from a suite of test structures in an attempt to illuminate the role of forward modelling error in directing the search for the optimal model. In particular, we computed TE and TM responses of the three electrical conductivity structures shown at the right in Fig. 5. The checkerboard (100:1 conductivity contrasts) and the single conductor (10:1 conductivity contrast) models were chosen since they should produce responses with large and small horizontal variations, respectively. We therefore expect the forward solver to require a large mesh to compute accurate checkerboard responses, while relatively small meshes should suffice for computation of the smoother responses of the single conductor model. The third test model we considered is the best-fitting COPROD2 conductivity structure (e.g. bottom right, Fig. 4) obtained by the genetic algorithm.

We used the finite difference forward solver to compute responses of the three test structures, the results are shown at the left in Fig. 5. The responses for the highest frequency used in the inversion were calculated on three uniform meshes of systematically varying node density. The coarsest mesh contained 4 nodes in both horizontal and vertical directions per conductivity block, the medium mesh contained 8 horizontal nodes per block, and the finest mesh 16 nodes. The meshes were specially constructed so that block boundaries coincided with rows and columns of nodes. The results indicate that all three TE responses are relatively stable with respect to mesh size beyond 8 nodes per conductivity block. Similarly the TM mode responses of the COPROD2

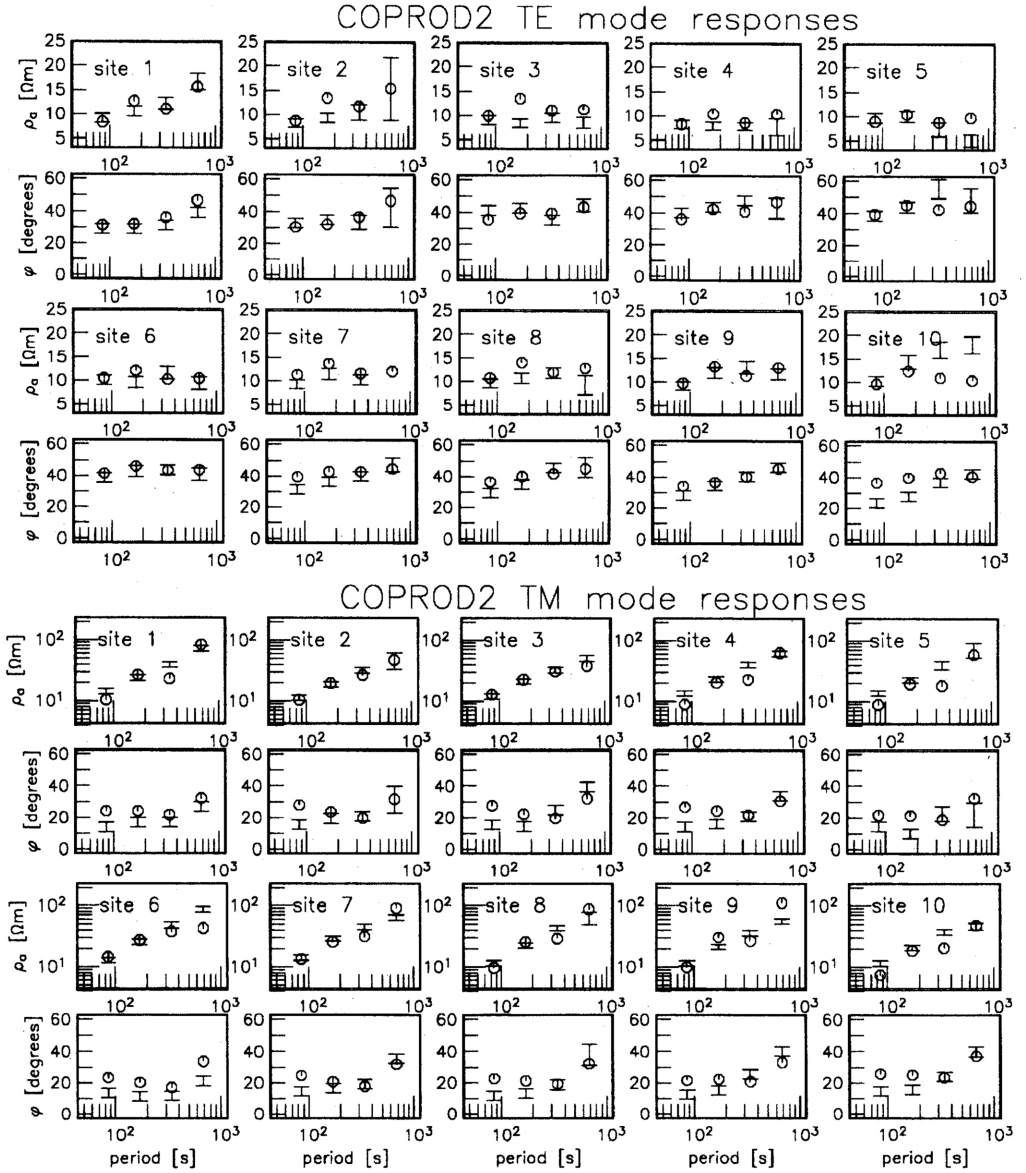

Fig. 3. The responses of the best-fitting electrical conductivity model (open circles) plotted alongside a subset of the COPROD2 data, with standard error, used in the inversion. The site number refers to the position of the measurement, see Table 1, increasing eastward.

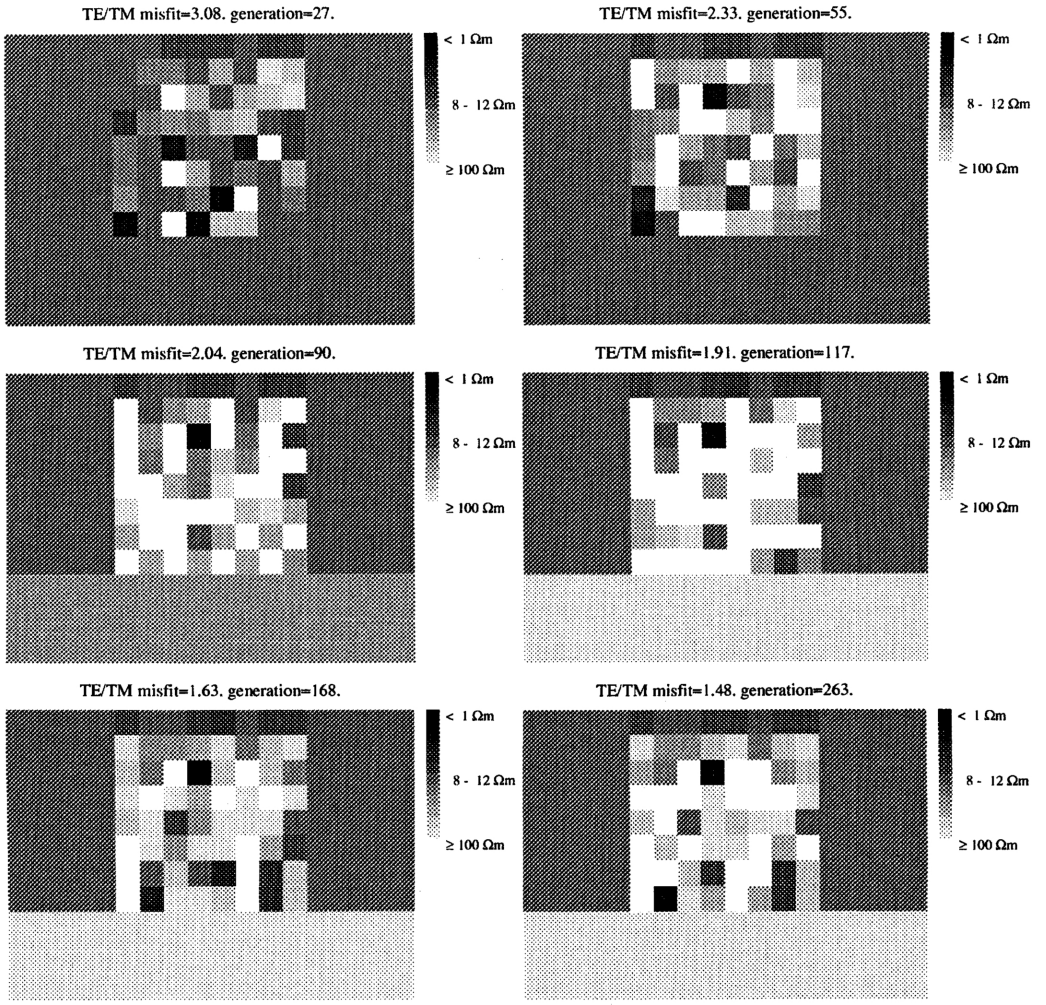

Fig. 4. This sequence of 2-D electrical cross-sections depicts the evolution of the best-fitting model during the COPROD2 TE/TM inversion. The inhomogeneous region is 120 km wide. VE = 2:1.

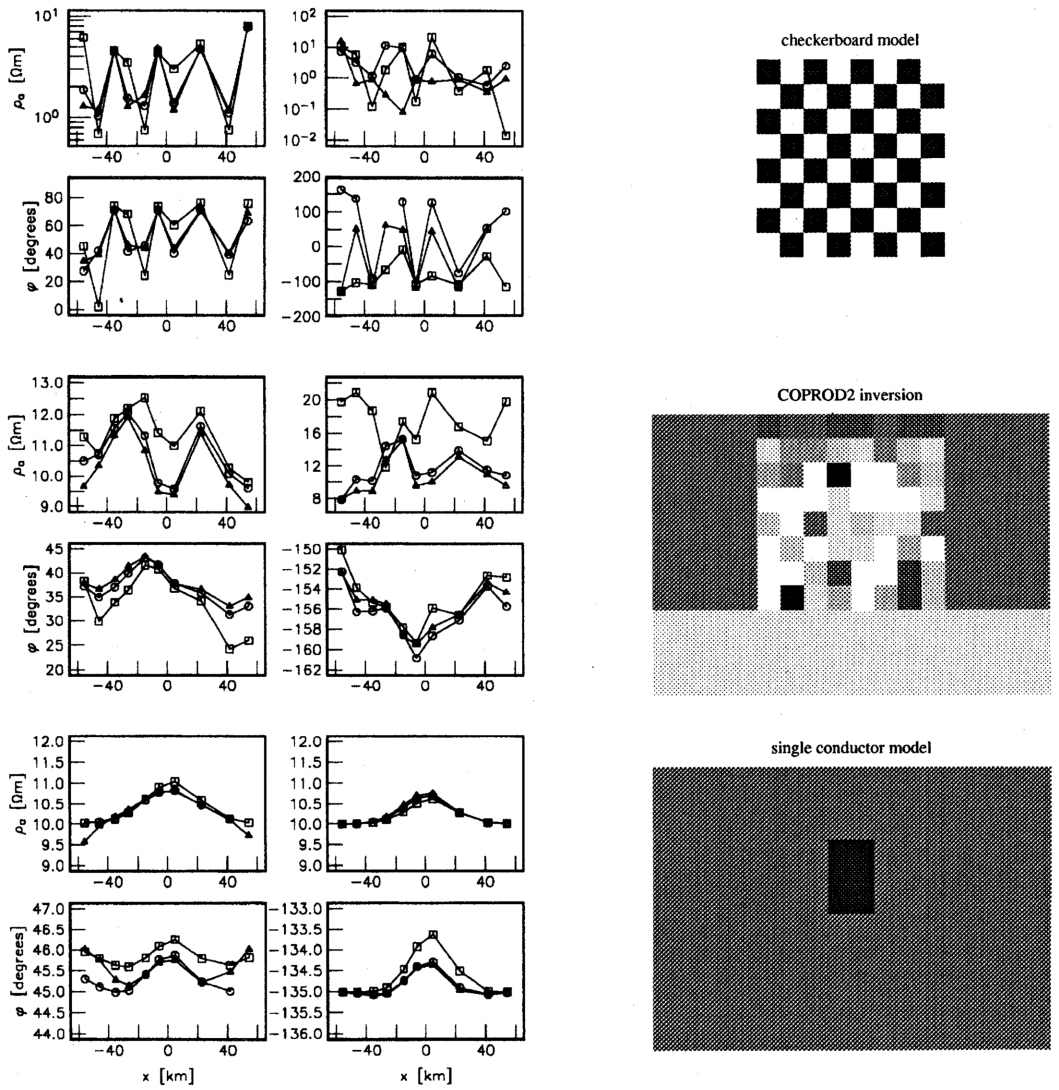

Fig. 5. Three electrical conductivity structures and their TE (left column) and TM (right) mode responses at a period of 85.33 s. The responses are computed on finite difference meshes containing 4 (squares), 8 (circles) and 16 (triangles) nodes per conductivity block.

and single conductor models are stable beyond 8 nodes per block. The TM mode checkerboard responses are very different when computed on the medium and fine meshes, however. Finally, of all the responses, only the apparent resistivity of the single conductor model appears to be stably determined on the coarse mesh.

7. Discussion

The genetic algorithm, validated by an inversion of synthetic data for a simple 2-D structure, has been used to invert a subset of the COPROD2 data in terms of optimal 2-D electrical conductivity models. Our best-fitting model attained an rms misfit of 1.48, for standard errors of at least 10% in apparent resistivity and 3° in phase. The misfit has stabilised when the inversion was terminated. Our coarse parameterisation of the electrical structure is clearly insufficient to breach the expected value of unity. However, we are confident that a finer parameterisation of the earth and/or an increase in the bounds on individual block conductivities should result in a model which does. Similarly, lower misfits would also be attained if the lateral terminations of the earth were a general n-layer conductivity structure, instead of the present layer over a half-space geometry. We would also like to explore the performance of the algorithm when the conductivity blocks are spaced logarithmically with depth, giving rise to relatively more model parameters at the shallow depths where spatial resolution is highest. The parameterisation used in this example of 2-D GA inversion is inadequate to reproduce fine details of the electrical structure in this area. For example, other workers have required a very thin conductive surface layer in order to fit the data. The results do, however, demonstrate the basic features and convergence properties of this nonlinear inversion method.

The best-fitting electrical conductivity model (bottom right, Fig. 4) contains certain features that are surprisingly consistent with the local geology and previous electromagnetic inversions. We interpret the uniform surface layer of conductance ~1000 S to be due to the sediments of the Phanerozoic Williston Basin. The single, highly conducting block in the upper central part of the inhomogeneous region corresponds to the centre of the NACP conductivity anomaly (c.f., JONES and CRAVEN, 1990). The depth to the NACP anomaly is consistent with previous studies. The overall low background conductivity ($\leq$0.01 S/m) of the inhomogeneous region is also in agreement with previous electromagnetic observations (JONES and CRAVEN, 1990).

During the inversion of the COPROD2 data set we fixed the size of the finite difference mesh at 5 nodes per skin depth, and did not enforce block boundaries to coincide with rows and columns of nodes. At the highest frequency, this mesh size corresponds to approximately 5.1 nodes per conductivity block. The results of the error analysis indicate, however, that at least 8 nodes per conductivity block are required for the stable determination of TE and TM mode responses of electrical conductivity structures with moderate conductivity contrasts. Although we were unable to run the inversion at the critical mesh density of 8 nodes per block due to computational limitations, the fact that our final model is consistent with previous inversions and local geology suggests that the forward modelling error, though significant, did not overwhelm the directed search of model space.

Large amounts of computer time are required for global optimisation to find the best-fitting electrical conductivity structure. This makes the use of genetic algorithms for 2-D magnetotelluric inversion practical only on the most powerful computers. Although the genetic algorithm takes much longer to run than other recently published inversions such as the Occam2 algorithm of DEGROOT-HEDLIN and CONSTABLE (1990) or the very fast RRI algorithm of SMITH and BOOKER (1991), the output is fundamentally different. Whereas the rapid inversion codes simply find the smoothest model from amongst the multitude which have a given misfit to the data, the genetic algorithm attempts the more demanding task of extracting the global best-fitting model. As mentioned in the introduction, minimum misfit algorithms are useful not necessarily to construct

geologically tenable models of the subsurface. Rather, they provide an absolute lower bound on misfit, given a particular model parameterisation and set of prior bounds.

An alternative strategy for nonlinear optimisation is the simulated annealing (SA) approach. Rather than choosing a biological analogue, SA mimics a thermodynamic system such as crystallisation of melt. A 1-D MT inversion based on SA is described in DOSSO and OLDENBURG (1991). SCALES *et al.* (1992) report that for large numbers of model parameters, GAs have more freedom to explore model space, and are less prone to converging to purely local minima than SAs. For problems with a small number of model parameters, GAs also contrast with linearised inversions which use, say, a steepest descent approach to converge to a local minimum near some starting model. One consequence of this is that GAs tend to be far more computationally demanding than corresponding regularised inversions, or than extremal nonlinear inversions based on linear or quadratic programming such as D^+. Unlike its extremal counterparts, however, the GA is easily adapted to any forward solution or objective function, and allows one to impose *a priori* information on the inversion.

The GA approach is, on the other hand, far less computationally demanding than conventional Monte Carlo algorithms. The GA rapidly determines the most likely part of model space in which minimum misfit solutions are likely to be found, and directs the search in those subregions of model space. For the inversions described in this paper (e.g. 66 model parameters, each of which is represented by an 8 bit binary word), discretised model space spans 2^{258}, or 8.8×10^{158} possible models. If we consider that a forward solution for each model, for four frequencies and two modes, takes ~38 s to compute on our present system, a complete Monte Carlo search through model space (or an ordered grid search examining every possible discretised model) would take $\sim10^{143}$ times the age of the universe.

We believe that application of the genetic algorithm and other nonlinear optimisation techniques to 2-D inverse problems, currently well within the capabilities of existing computational facilities for 1-D, will become increasingly practical as massively parallel processing architectures evolve. The GA is inherently a parallel algorithm. Work is underway in porting the code to such a machine.

We are indebted to P. J. Savage of PanCanadian Petroleum Limited of Calgary, Alberta for making the raw (uncorrected for static distortion) data available to the Geological Survey of Canada. We would like to acknowledge J. Torquil Smith for valuable discussions on magnetotelluric inversion and genetic algorithms and also for providing the software to plot the electrical conductivity models. A. Schultz thanks Malcolm Sambridge whose many conversations during the early stages of development of the 1-D MT GA influenced the choice of tournament selection. Contribution Number 3362, Department of Earth Sciences and Institute of Theoretical Geophysics, University of Cambridge.

REFERENCES

CAMFIELD, P. A. and D. I. GOUGH, A possible Proterozoic plate boundary in North America, *Can. J. Earth Sci.*, **14**, 1229–1238, 1977.

DEGROOT-HEDLIN, C. and S. CONSTABLE, Occam's inversion to generate smooth, two-dimensional models from magnetotelluric data, *Geophysics*, **55**, 1613–1624, 1990.

DOSSO, S. E. and D. W. OLDENBURG, Magnetotelluric appraisal using simulated annealing, *Geophys. J. Int.*, **106**, 379–385, 1991.

GALLAGHER, K., M. SAMBRIDGE, and G. DRIJKONINGEN, Genetic algorithms: an evolution from Monte Carlo methods for highly-nonlinear geophysical optimization problems, *Geophys. Res. Lett.*, **18**, 2177–2180, 1991.

HOLLAND, J. H., *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbour, 1975.

JIN, S. and B. MADARIAGA, Background velocity inversion with a genetic algorithm, *Geophys. Res. Lett.*, **20**, 93–96, 1993.

JONES, A. G., The COPROD2 dataset: Tectonic setting, recorded MT data and comparison of models, *J. Geomag. Geoelectr.*, this issue, 933–955, 1993.

JONES, A. G. and J. A. CRAVEN, The North American Central Plains conductivity anomaly and its correlation with gravity, magnetic, seismic, and heat flow data in Saskatchewan, Canada, *Phys. Earth Planet. Int.*, **60**,

169–194, 1990.

PARKER, R. L., The inverse problem of electromagnetic induction: existence and construction of solutions based on incomplete data, *J. Geophys. Res.*, **85**, 4421–4428, 1980.

PARKER, R. L. and K. A. WHALER, Numerical methods for establishing solutions to the inverse problem of electromagnetic induction, *J. Geophys. Res.*, **86**, 9574–9584, 1981.

REITZEL, J. S., D. I. GOUGH, H. PORATH, and C. W. ANDERSON, Geomagnetic deep sounding and upper mantle structure in the western United States, *Geophys. J. Roy. Astron. Soc.*, **19**, 213–235, 1970.

SAMBRIDGE, M. and G. DRIJKONINGEN, Genetic algorithms in seismic waveform inversion, *Geophys. J. Int.*, **109**, 323–342, 1992.

SCALES, J. A., M. L. SMITH, and T. L. FISCHER, Global optimization methods for multimodal inverse problems, *J. Comp. Phys.*, **103**, 258–268, 1992.

SCHULTZ, A., A genetic algorithm for nonlinear magnetotelluric inversion: nonlinear hypothesis testing,—or—"Cutting your throat with Occam's Razor" (abstract), *11-th Workshop on Electromagnetic Induction in the Earth*, Wellington, New Zealand, August, 1992.

SCHULTZ, A., R. KURTZ, A. D. CHAVE, and A. G. JONES, Conductivity discontinuities in the upper mantle beneath a stable craton, *Geophys. Res. Lett.*, 1993 (accepted).

SMITH, J. T. and J. R. BOOKER, Rapid inversion of two- and three-dimensional magnetotelluric data, *J. Geophys. Res.*, **96**, 3905–3922, 1991.