Computationally Efficient Regularized Inversion for Highly Parameterized MODFLOW Models

Matthew Tonkin[1], Tom Clemo[2], John Doherty[3]
[1]*S. S. Papadopulos & Associates, Inc., matt@sspa.com, Bethesda, MD, USA*
[2]*Boise State University, tomc@cgiss.boisestate.edu, Boise, ID, USA*
[3]*Watermark Numerical Computing, jdoherty@gil.com.au, Corinda, QLD, Australia*

ABSTRACT

Though popular in the geophysical modeling community, specification of spatially distributed parameters at a scale commensurate with prevailing geological heterogeneity has not been possible in common groundwater modeling practice. The principal reasons for this are (1) the high computational burden of obtaining derivatives necessary for parameter estimation, (2) the memory required to store the derivative and coefficient matrices generated in classical Levenberg-Marquardt methods, and (3) lack of experience within the groundwater modeling community with regularized inversion. The development of adjoint state derivatives calculation within MODFLOW-2000 removed the first of these roadblocks. Efficient compression of the sensitivity matrix (Jacobian) within the inversion code PEST dramatically reduces memory requirements while increasing solution speed. An independent regularization model allows for the specification of arbitrarily complex linear and non-linear regularization schemes. These developments have enabled investigation of the role of various regularization schemes within systems comprising many hundreds to thousands of parameters. Some example regularization schemes are presented for use in different model calibration settings. The methods presented complement zone and pilot-point based parameterization schemes for complex 2- and 3-dimensional groundwater systems, due to the ability to accommodate the estimation of a large number of parameters in a numerically stable and geologically realistic manner.

INTRODUCTION

The inverse problem in groundwater modeling is generally ill-posed and non-unique. The typical approach adopted for resolving non-uniqueness is parameter parsimony – parameterizing the model on the basis of a few zones or pilot points such that the number of observations (*nobs*) is considerably greater than the number of estimated parameters (*npar*). Prior information such as preferred or expected parameter values may be specified to condition the problem. Defining parameter zones enforces the condition that all cells within a zone possess the same parameter value, e.g. conductivity. This is effectively a form of regularization employed to make the inverse problem tractable, and inversion results (parameter estimates) are conditional upon the zonation scheme. Repercussions for model predictions of the chosen zonation scheme are difficult to analyze (Moore and Doherty, 2003).

An alternative approach is to parameterize the model domain using many hundreds or thousands of parameters, and employ a regularization scheme adopted for the problem at hand. The regularization is chosen to stabilize the inverse problem by mitigating parameter correlation, enhance parameter sensitivity, and guarantee convergence to a unique solution (e.g. Engl and others, 1996). This approach is not adopted in groundwater modeling partly due to the computational burden of obtaining the Jacobian and the memory required to store the derivative and coefficient matrices, and partly due to lack of research into appropriate regularization methods. The field of hydraulic tomography represents one implementation of regularized inversion in groundwater problems (e.g. Lui and others, 2002).

The intention of this paper is to describe some recent advances in sensitivity analysis, parameter estimation, and the development of a regularization model that together enable the rapid inversion of highly parameterized models in a stable fashion. These developments will enable the investigation of appropriate regularization schemes for a variety of geologic settings, and may promote analysis of model predictions under varying assumptions of geologic structure that are beyond inference in classical model calibration.

REGULARIZED INVERSION

Modified-Newton parameter estimation techniques aim to minimize an objective function defined as:

$$\Phi_m = (d - M(p))^{-1} Q_1 (d - M(p)) \qquad \text{(eq. 1)}$$

where p is a vector of order *npar* containing the model parameters; and d is a vector of order *nobs* containing measurements (observations); and M is a matrix with dimensions *npar*-columns and *nobs*-rows that describes the action of the model. Matrix Q_1 incorporates weights assigned to the measurement observations. Iterative solution of the inverse problem is designed to minimize Φ_m, termed the measurement objective function. A regularization objective function can be formed as:

$$\Phi_r = (e - R(p))^{-1} Q_2 (e - R(p)) \qquad \text{(eq. 2)}$$

where *R* is a regularization operator; e expresses a preferred condition; and Q2 incorporates weights assigned to the regularization observations. The regularized form of the inverse problem, sometimes referred to as penalized least squares (e.g. Engl and others, 1996), becomes

$$\Phi_g = \Phi_m + \mu\Phi_r \qquad \text{(eq. 3)}$$

where μ is a regularization weight multiplier. Principal features of a good regularization scheme are stability and convergence to a unique solution. Regularization relationships typically range from linear to quadratic, yielding smooth solutions (e.g. Vogel, 1997). Though this provides a stable solution, smoothness is not appropriate for all settings. Alternative schemes developed in the geophysical community for different geologic settings include:

- Focused regularization (e.g. Portniaguine and Zhdanov, 1999)
- Directionally varying and spatially varying regularization (e.g. Pain and others, 2002)
- Iteratively re-weighted regularization (e.g. Portniaguine and Zhdanov, 2002)

OBTAINING DERIVATIVES

Sensitivities necessary for parameter estimation can be obtained in three ways:

1) Finite difference or perturbation sensitivities, obtained by perturbing each parameter in turn and calculating, by finite differences, the sensitivity of each observation to the parameter perturbation. This requires *npar+1* runs to obtain the forward or backward derivatives, and *(npar*2)+1* runs for more accurate central derivatives (Doherty, 2002).
2) Solution by the sensitivity equation method, as employed in MODFLOW-2000 (Hill and others, 2000). This requires *npar+1* of the finite difference equations, but typically provides more accurate derivatives than perturbation methods (Hill and others, 2000).
3) Adjoint state sensitivities, obtained by perturbing each observation in turn and calculating the derivative of the equation of state with respect to each parameter. This requires *nobs+1* runs to obtain the sensitivities (e.g. Sun, 1994).

For problems characterized by *npar>>nobs*, finite difference and sensitivity equation methods are computationally prohibitive, and the adjoint state is required. The adjoint has been described in the groundwater literature (e.g. Townley and Wilson, 1985; Sun, 1994), but has not seen wide use. Until recently, the most accessible application of the adjoint within the groundwater community was implemented in MODFLOWP (Hill, 1992) where it was employed to calculate derivatives with respect to the total objective function. Adjoint sensitivity capabilities have recently been added to the MODFLOW-2000 Observation and Sensitivity Processes (Clemo and others, 2003; Harbaugh and others, 2000; Hill and others, 2001). The new adjoint sensitivity code calculates the sensitivity of each observation with respect to each parameter, allowing construction of the Jacobian matrix. Tests using highly parameterized systems indicate better-than order-of-magnitude speed up times with respect to sensitivities obtained by finite-differences or sensitivity equations.

MEMORY, STORAGE AND SOLUTION OF THE MATRIX EQUATIONS

Regularized inversion typically involves determining many hundreds or thousands of relationships between model parameters. Regularized inversion requires calculation of derivatives of these relationships, comprising rows of matrix M in equation 1. Whereas rows of the derivatives matrix pertaining to the sensitivity of observations with respect to parameters are full, rows pertaining to items of regularization may contain only a very small number (often one) of non-zero entries, since regularization generally relates a single parameter to second parameter or to a preferred condition. Hence, the entire derivatives matrix is very sparse. Generalized matrix compression utilities may not take full advantage of this sparsity, since for different cases it may not conform to generalized sparse-matrix structures. The authors developed a custom compression for file reading/writing based on the premise that the regularization model prepares all entries for derivatives pertaining to regularization items. Only non-zero entries are read/written. The structure for file reading/writing is shown conceptually in figure 1. Custom file retrieval and matrix manipulation routines added to PEST access entries in the compressed Jacobian as necessary for solving the matrix equations. Tests conducted on highly parameterized systems indicate hard-disc and RAM savings of over 98 percent. Modifications of the solver routine for solving large-parameter systems reduced computational time for solving the large normal matrices by about one order of magnitude.

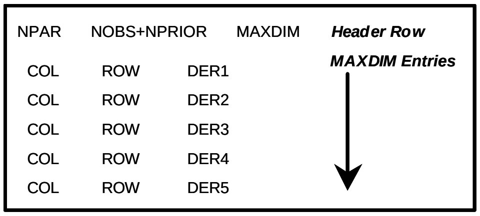

Figure 1: Compressed Matrix Form

REGULARIZATION MODEL

The regularization model accomplishes three principal objectives, namely:

1) Calculates linear and (or) non-linear regularization items of arbitrary complexity, specified by the modeler.
2) Compiles the adjoint-state calculated sensitivities together with analytically determined derivatives of the regularization items into a single, highly compressed derivatives matrix.
3) Provides the interface between the adjoint sensitivity program and PEST.

Combining the three codes into a single program would save file reading and writing. However this represents a very small portion of the total execution time, and this savings would be at the expense of the enormous flexibility offered by the current configuration. Each step of the combined process, shown schematically in figure 2, is (largely) model-independent, but offers rapid solution of its component of the inversion problem.

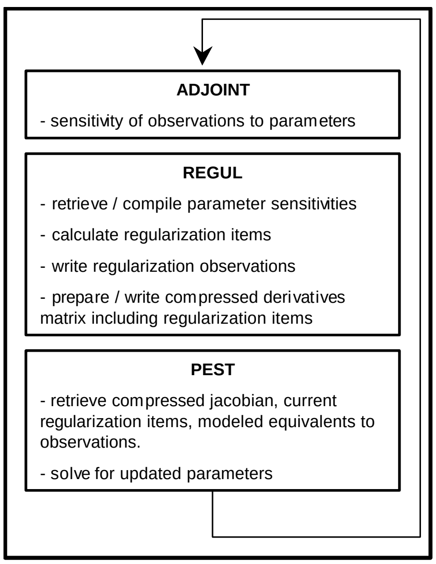

Figure 2: Three Components of the Regularized Inversion

TWO EXAMPLE REGULARIZATION SCHEMES

Anisotropic regularization for a smoothly varying field

This exercise is intended to demonstrate the impact of regularization anisotropy on the correlation of an estimated field with a true (synthetic) field, and on the prediction of advective transport velocities. A synthetic model domain was constructed comprising 1 layer, 20 rows, and 20 columns. Hydraulic conductivity was assigned to the domain using Fieldgen (Doherty, 2002), a geostatistical field generator with an anisotropy of 1.5 in the azimuth 060 degrees. A small resistive structure mimicking a fault was superimposed with an orientation of 045 degrees. Twenty synthetic, randomly located water level

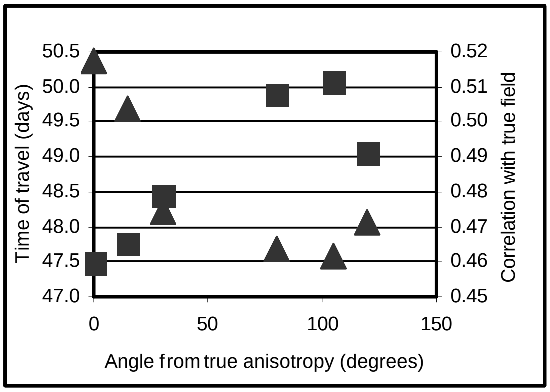

Figure 3: Time of travel (squares) and correlation with true field (triangles)

observation locations were added. The time-of-travel for a particle released within the domain to reach the down gradient boundary was recorded for each simulation. The true time-of-travel is 42 days. All simulations achieved sum-of-squared-residuals under $10e^{-5}$ ft^2. The regularization adopted was a simple 8-node logarithmic differencing scheme, with an arbitrary anisotropy direction and magnitude. The regularized inversion problem comprises 400 parameters, 20 observations, and 3200 regularization equations, resulting in 14,400 non-zero entries in the derivatives matrix. Without compression, this would require approximately 20MB storage (double precision); under compression 350KB are required, or less than 2 percent. Execution time for the full inversion is less than 2 minutes on a 1.7GHz, 512 MB RAM laptop.

Observations and conclusions

This exercise demonstrates the use of appropriate regularization anisotropy in a smoothly varying field. Figure 3 indicates that the correlation of the estimated field with the true field and the estimated exit-time of the particle are improved when the direction of regularization anisotropy approaches the true direction of anisotropy. Exercises such as this may support or comprise a component of more formal approaches to regularization scheme selection such as generalized cross-validation (e.g., Yao and Roberts, 1999).

Identification of a categorical (zone) field using continuous regularization

The object of this exercise is to compare some regularization schemes for identifying areas of equivalent hydraulic properties (zones). A synthetic model comprising 1 layer, 40 rows, and 30 columns is shown in figure 4. A uniform transmissivity of 100 ft^2/d is assigned to approximately 90% of the domain. A small area of transmissivity 2 ft^2/d is assigned in the middle of the domain. Fifty synthetic, randomly located water level observation locations were added. Three regularization schemes were adopted (a) maximal smoothness, (b) focusing regularization (maximal smoothness with a power penalty), and (c) a two-point regularization scheme. The two-point scheme forces parameters to specified fixed bounds using a non-linear power function that tends to zero at preferred values. (b) and (c) are non-linear and could not be described in the form of linear prior-information. All simulations achieved sum-of-squared-residuals under $10e^{-3}$ ft^2. The regularized inversion problem comprises 1200 parameters, 50 observations, and 12,000 regularization equations, resulting in 81,600 non-zero entries in the sensitivity matrix. Without compression, this would require approximately 230MB storage (double precision); under compression 2MB are required, less than 1 percent. Execution time for the full inversion ranged from 30 minutes to 120 minutes on a 1.7GHz, 512 MB RAM laptop.

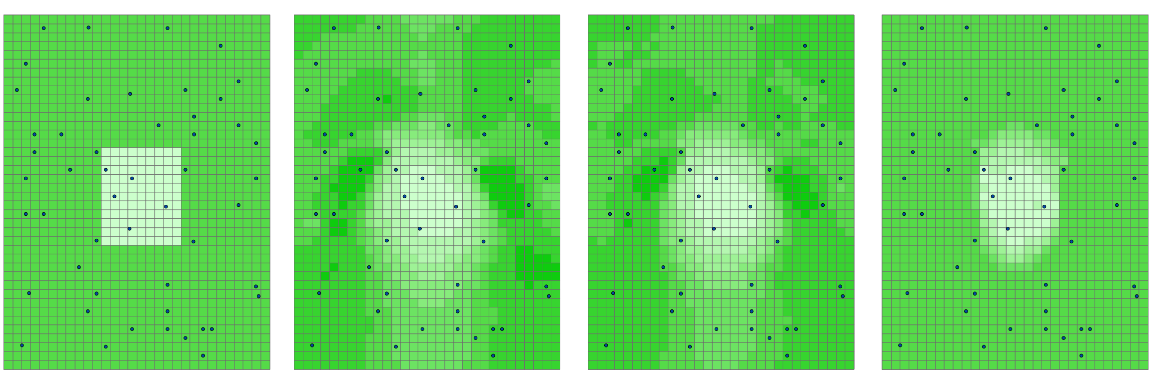

Figure 4: (a) True field, (b) maximally smooth, (c) focused, and (d) two-point simulated fields.

Observations and conclusions

The maximal-smoothness scheme identifies the general area of lower transmissivity, but spreads the anomaly throughout the domain. Oscillatory artifacts of the scheme related to Gibbs phenomenon develop in the vicinity of the sharp interfaces (e.g. Hobro and others, 2003). The power regularization scheme performs better at targeting the location of the anomaly. The two-point scheme shows promise for identifying sharp boundaries. Both non-linear schemes perform better than the linear smoothing.

CONCLUSIONS

The suite of programs described enables rapid regularized inversion of highly parameterized MODFLOW models. This will facilitate an assessment of the most appropriate regularization scheme(s) for a range of settings. Schemes based on smoothness constraints may not be suitable in many geological contexts. Early efforts at identifying discrete areas of equivalent properties (zones) using continuous regularization methods show some promise, but further development of appropriate regularization schemes is warranted. Extending this work to explicitly include derivates with respect to advective or reactive transport will yield valuable additional sensitivity information that was absent from this study.

REFERENCES

Clemo, T., P. Michaels, and R.M. Lehman, 2003. Transmissivity Resolution Obtained from the Inversion of Transient and Pseudo-Steady Drawdown Measurements (this volume).

Doherty, J. E., 2002. PEST Model-independent parameter estimation, users guide.

Engl, H. W., Hanke, M., and Neubauer, A., 1996. Regularization of inverse problems, Klewer Academic Publishing, 251 p.

Harbaugh, A.W., Banta, E.R., Hill, M.C., and McDonald, M.G., 2000. MODFLOW-2000, the U.S. Geological Survey modular ground-water model – user guide to modularization concepts and the ground-water flow process, U.S. Geological Survey Open-File Report 00-92.

Hill, M.C., 1992. A computer program (MODFLOWP) for estimating parameters of a transient, three-dimensional, ground-water flow model using nonlinear regression: U.S. Geological Survey Open-File Report 91-484, 358 p.

Hill, M.C., Banta, E. R., Harbaugh, A. W., and Anderman, E. R., 2000. Modflow-2000, The U.S. Geological Survey Modular Ground-Water Model — User guide to the Observation, Sensitivity, and Parameter-Estimation Processes and three postprocessing programs: U.S. Geological Survey Open-File Report 00-184, 209 p.

Hobro, J. W. D., Singh, S. C., and Minshull, T. A., 2003. Three-dimensional tomographic inversion of combined reflection and refraction seismic traveltime data, Geophys. J. Int., 152, 79-93.

Lui, S., Yeh, T.-C. J, and Gardiner, R., 2002, Effectiveness of hydraulic tomography: Sandbox experiments, Wat. Res. Res., 38 (4), 5-1.

Moore, C., and Doherty, J.E., 2003. Quantifying Model Predictive Uncertainty In Heterogeneous Aquifers - An Hypothesis Testing Method (this volume)

Pain, C. C., Herwanger, J. V., Worthington, M. H., and de Oliveira, C. R. E., 2002. Effective multi-dimensional resistivity inversion using finite-element techniques, Geophys. J. Int., 151, 710-728.

Portniaguine, O., and Zhdanov, S., 1999. Focusing geophysical inversion images, Geophysics, 64 (3), 874-887.

Portniaguine, O., and Zhdanov, S., 2002. 3-D magnetic inversion with data compression and image focusing, Geophysics, 67 (5), 1532-1541.

Sun, N.-Z., 1994. Inverse problems in groundwater modeling, Klewer Academic Publishing, 352 p.

Townley, L.R., and Wilson, J. L., 1985. Computationally efficient algorithms for parameter estimation and uncertainty propagation in numerical models of groundwater flow. Wat. Res. Res., 21 (12), 1851-1860.

Vogel, C.R., 1997. Nonsmooth regularization, *in* Inverse Problems in Geophysical Applications, ed. H.W. Engl, A.K. Louis, and W. Rundell, SIAM, 1-11.

Yao, S. Z., and Roberts, R. G., 1999. A practical regularization for seismic tomography, Geophys. J. Int., 138, 293-299.