

**Use of PEST and Some of its Utilities
in Model Calibration and Predictive
Error Variance Analysis:-
A Roadmap**

prepared by

**John Doherty
Watermark Numerical Computing
May 2007**

Table of Contents

1. Introduction.....	1
1.1 General.....	1
1.2 Utility Software.....	2
1.3 Flow Charts.....	3
1.4 Further Documentation.....	3
1.5 Workshop.....	3
1.6 The JACTEST Utility.....	3
2. Pre-Calibration Analysis.....	4
2.1 The PEST Input Dataset.....	4
2.1.1 Construction of a PEST Input Dataset.....	4
2.1.2 Inclusion of a Prediction.....	4
2.1.3 Initial Parameter Values.....	4
2.1.4 Observation Weights.....	4
2.2 Uncertainty and Error.....	5
2.2.1 The Bayesian Approach.....	5
2.2.2 Model Calibration.....	6
2.3 Linear Predictive Uncertainty Analysis – the PREDUNC Suite.....	6
2.3.1 General.....	6
2.3.2 Theoretical Basis.....	7
2.3.3 Obtaining a JCO File.....	7
2.3.4 Predictive Sensitivity Vector.....	7
2.3.5 Parameter Uncertainty File.....	7
2.3.6 PREDUNC1, PREDUNC4 and PREDUNC5.....	8
2.3.7 Optimisation of Future Data Acquisition.....	8
2.3.8 Limitations of PREDUNC Analysis.....	8
2.4 Linear Predictive Error Variance Analysis – the PREDVAR Suite.....	9
2.4.1 General.....	9
2.4.2 Theoretical Basis.....	9
2.4.3 Parameter Scaling.....	9
2.4.4 Limitations of PREDVAR Analysis.....	9
2.5 Flow Chart.....	10
3. Model Calibration.....	11
3.1 Initial Parameters.....	11
3.2 Adding Regularisation.....	11
3.2.1 General.....	11
3.2.2 The ADDREG1 Utility.....	11
3.2.3 The GENREG Utility.....	12
3.3 Computing Sensitivities.....	12
3.3.1 Computation of Base Parameter Sensitivities.....	12
3.3.2 The JCO2JCO Utility.....	12
3.4 Jacobian Manipulation Utilities.....	13
3.5 Preparing for SVD-Assisted Parameter Estimation.....	13
3.5.1 The SVDAPREP Utility.....	13
3.5.2 Number of Super Parameters.....	13

3.6 SVDA and Internal Parameter Scaling	14
3.7 SVD-Assisted Parameter Estimation	15
3.7.1 The PARLAM Variable	15
3.7.2 The PARREP Utility	15
3.8 Obtaining Super Parameter Definitions	15
3.9 Flow Chart	16
4. Linear Post-Calibration Error Variance Analysis	17
4.1 Predictive Error Variance Analysis	17
4.1.1 General	17
4.1.2 RESPROC and Related Utilities	17
4.1.3 The PREDERR Utility	17
4.1.4 Re-Computation of Parameter Sensitivities	17
4.1.5 Observation Uncertainty File	18
4.1.6 Parameter Error Covariance Matrix	18
4.2 Post-Calibration PREDVAR and PREDUNC Analysis	18
4.3 Flow Chart	18
5. Nonlinear Post-Calibration Error Variance Analysis	19
5.1 Constrained Maximization/Minimization of a Key Model Prediction	19
5.1.1 Solution Space and Null Space Constraints	19
5.1.2 REGPRED and Related Utilities	19
5.1.3 Using SVD-Assist for Nonlinear Predictive Error Variance Analysis	20
5.2 Null Space Monte Carlo	20
5.2.1 Discrete and Continuous Stochastic Fields	20
5.2.2 The RANDPAR and PNULPAR Utilities	20
5.2.3 Model Re-Calibration using Pre-Calculated Sensitivities	21
5.2.4 Steps Required for Null Space Monte Carlo Analysis	21
5.3 Flow Chart	23
6. References	24
Appendix A	25
Appendix B	27
Appendix C	29
Appendix D	30
Appendix E	32

1. Introduction

1.1 General

The purpose of this document is to provide a brief “roadmap” of methodologies available through PEST and its utilities for undertaking model calibration and predictive uncertainty analysis. It is assumed that the model employs many parameters – probably many more parameters than can be estimated uniquely - and that these parameters reflect the fact that natural systems are complex and that certain predictions required of the model are dependent on this complexity. Parameters may be of one or many types, for example pilot point parameters and/or zone parameters describing the spatial distribution of one or more property types throughout the model domain. Parameters can also include distributed or “lumped” parameters employed by spatially variable or stand-alone submodels that provide input and/or boundary conditions to an overarching spatial model.

Modern calibration software allows parameters to be represented (and adjustable) in a model at a level that is commensurate with the inherent complexity of the system which the model purports to represent, at least to the extent that a prediction may be sensitive to such complexity. Thus the pre-emptive (and often heavy-handed) pre-calibration parsimonizing required by previous approaches to model calibration is no longer needed. Simplifications necessary to achieve a well-posed solution to the inverse problem of model calibration are undertaken through the regularised inversion process itself. This approach to model calibration and model-based data analysis brings with it the following advantages.

1. The calibration process is given maximum flexibility in extracting information from a calibration dataset.
2. Thus predictions made by the model will be closer to maximum likelihood (and thus possess minimum potential wrongness) than they otherwise would.
3. The extent of possible predictive wrongness can be quantified.

The last point is particularly important. Most predictive error (especially where a prediction is sensitive to fine system detail that “slips between the cracks of the calibration process”) is dominated by the so-called “null space term” of the predictive error variance equation. This term is ignored in traditional post-calibration error analysis undertaken as an adjunct to over-determined parameter estimation. Such analysis is therefore based on the premise that “if it cannot be estimated it doesn’t exist”. In contrast, highly parameterised inversion explicitly acknowledges the fact that a prediction may be sensitive to system details and properties that cannot be estimated due to lack of information in the calibration dataset. To the extent that aspects of this detail are indeed inestimable, and therefore reside in the calibration null space, the potential error of a prediction which depends on them is not reduced through the calibration process, and is therefore the same as it would have been if the model had not been calibrated at all. This uncertainty is a function of the innate variability of system hydraulic properties – this

being encapsulated in what will be referred to herein as the $C(\mathbf{p})$ covariance matrix (in which the \mathbf{p} vector represents parameters employed by the model). See Moore and Doherty (2005) and Moore and Doherty (2006) for more details.

On the other hand, to the extent that a prediction is dependent on aspects of the system that are estimable through the calibration process, the potential error associated with that component of the prediction arises from the fact that the measurements from which these aspects of the system were estimated are contaminated by noise, the stochastic descriptor of which is provided by another covariance matrix $C(\epsilon)$ (with ϵ representing noise associated with these measurements). This noise has both a structural and measurement component. (Note that the former can often be substantially reduced by the use of many parameters instead of just a few as required by traditional approaches to model calibration.)

Use of the methods described below is based on the assumption that the user “knows” $C(\mathbf{p})$. This is not as worrying an assumption as it may first appear, for $C(\mathbf{p})$ is, in fact, as much an encapsulation of a modeller’s ignorance as it is of his/her knowledge. Hence it can be simply an expression of “what may happen down there that I don’t know about” rather than an exact geostatistical descriptor of hydraulic property spatial variability. It is also assumed that the modeller knows the structure of $C(\epsilon)$, but is informed of its magnitude through the calibration process. In most modelling contexts $C(\epsilon)$ is assumed to be diagonal. Its magnitude is estimated to be that which is commensurate with the level of model-to-measurement misfit achieved through the calibration process. The user often chooses a suitable level of model-to-measurement misfit through trial and error based on repeated calibration attempts. The use of a large number of parameters often allows a very good fit to be achieved between model outputs and field data. However a fit that is “too good” is indicated by estimated parameter values that are unrealistic. Acceptance of a less-than-perfect fit, thereby accepting that measurements are contaminated by structural or observation noise, then becomes an important aspect of the calibration process; in implementing this process, the level of suspected structural plus observation noise is thereby estimated.

1.2 Utility Software

Three suites of utility programs are available to assist with PEST usage, and to undertake various types of analysis based on PEST input datasets. The first set (referred to herein as “PEST Utilities”) are supplied with PEST. Another set is the “Groundwater Data Utilities”. Many members of this latter suite were written to expedite PEST usage in the MODFLOW/MT3D modelling context; however other members of this suite are more general in nature and can be used with other models as well. A third set, known as the “Surface Water Utilities” is also available; these will not be discussed herein.

PEST and all of its utility suites can be downloaded from the following site:-

<http://www.sspa.com/pest>

1.3 Flow Charts

The Appendices to this document contains a number of flow charts which synthesise the information contained in the text. These can be used for easy reference, and the text consulted for more detail.

1.4 Further Documentation

A full description of all software and methodologies discussed in this document can be found in PEST documentation, this comprising the PEST manual and its Addendum. See also the manuals for the Ground and Surface Water Utilities.

1.5 Workshop

A workshop is available in which many of the techniques and methodologies discussed herein are described in more detail. Files and documentation are provided, these allowing a user to implement these methodologies him/herself in the context of a simple MODFLOW model. These are provided to PEST course attendees. They can be provided to certain others on request.

1.6 The JACTEST Utility

PEST-based calibration and uncertainty analysis rests on the premise that model outputs are differentiable with respect to model parameters. Unfortunately, this condition is often violated, sometimes as an outcome of non-convergent iterative solution processes, and sometimes as an outcome of poor model algorithmic bases – see Kavetski et al (2005) for more details. PEST will tolerate the fact that some degree of “numerical granularity” in model outputs can tarnish derivatives computed through finite parameter differences. However if the problem is too severe, many of the methods described herein will not work. If PEST performance is poor, and the user suspects that model derivatives may be the problem, the integrity of these derivatives can be tested with the JACTEST utility (a PEST utility) once a PEST input dataset has been prepared.

2. Pre-Calibration Analysis

2.1 The PEST Input Dataset

2.1.1 Construction of a PEST Input Dataset

A PEST input dataset is readily built either by hand, or with the help of various PEST and commercial utility programs. However, as was discussed above, it is recommended that such an input dataset cite many parameters rather than just a few – as many parameters as key model predictions may be sensitive to, on a scale that allows representation of the level of heterogeneity to which these predictions may also be sensitive. For these same reasons it is likely that few, if any, of these parameters will be tied or fixed.

Let it be assumed, for the moment, that in the PEST control file which forms the centrepiece of the PEST input dataset, PEST is instructed to run in “estimation” mode.

2.1.2 Inclusion of a Prediction

Under some circumstances it may be advisable to consider “carrying” one or more predictions of interest in the PEST input dataset as “observations” with weights of zero. If calculation of these predictions can be undertaken with little or no computational burden beyond that required for calculation of calibration-related model outputs, then these are immediately available for predictive error analysis as sensitivities of these predictions to model parameters are computed as part of normal PEST usage.

2.1.3 Initial Parameter Values

A significant difference between over-determined and under-determined parameter estimation (the latter describing the type of inversion required where there is an abundance of parameters), is that more care should be taken in providing initial values to model parameters in the latter case than in the former case. Initial parameter values should be those which are considered to be of highest likelihood based on geological/scientific evidence. This strategy, together with the regularization devices described below, provides a realistic “fall back position” for parameters and parameter combinations that are inestimable through the calibration process.

2.1.4 Observation Weights

The choice of observation weights is a subjective one. Where uncertainty analysis is to be undertaken using the PREDVAR or PREDUNC utility families discussed below, it is essential that these be proportional to the inverse of the assumed standard deviation of measurement noise (including structural noise as necessary). In most cases it is easiest to adopt a proportionality constant of 1; thus if, for example, a user estimates that the noise associated with all head measurements is characterised by a standard deviation of 2m, the weight assigned to each head measurement in the PEST control file should be 0.5.

On the other hand, if a PEST input dataset has been prepared in order to form the basis of a calibration exercise, a different weighting philosophy may sometimes be appropriate. When calibrating a model myself, I assign observations to a number of different observation groups. Grouping is normally on the basis of observation type; however observations of the same type may be subdivided into further groups by location, layer, assumed relevance to a prediction, or some other basis for similarity or collective relevance to the calibration/prediction process. In many cases I then consider that each such group should neither dominate the objective function, nor be dominated by other components of the objective function, for I may consider that the information contained in each such group is of roughly equal salience to the parameter estimation process. This can be particularly relevant where each such group holds information that is separately informative of different aspects of the system; in such a case it is important that the individual “message” from each such group be “heard” by PEST. Hence while within-group weighting may indeed be such as to reflect relative observation “believability” (the inverse of observation noise), it may be important that the weighting between groups be such as to ensure that each group makes roughly the same contribution to the objective function at the start of the parameter estimation process. This can be achieved through undertaken the following steps.

1. Set NOPTMAX to zero in the PEST control file.
2. Run PEST in order to calculate the objective function, as well as the contribution made to the objective function by each observation group. (PEST will run the model only once when NOPTMAX is set to zero.)
3. Use the PWTADJ1 utility (a PEST utility) to write a new PEST control file in which observation weights are adjusted such that the contribution by each observation group to the overall objective function is the same.

Alternatively the ADJOBS utility (from both the Groundwater and Surface Water Utility suites) can be employed for weights adjustment. This allows weights to be formulated as functions of observation values (with different formulations allowed for each group), and allows subdivision of existing observation groups to form new groups, using observation names as a basis for such subdivision.

2.2 Uncertainty and Error

2.2.1 The Bayesian Approach

To a Bayesian, the notion of “model calibration” is a strange concept. What is the point of assigning unique values to model parameters if uniqueness is achieved through a simplification process that results in spatial parameter value assignments that are a pale reflection of the complexity of real-world systems? How can a model parameterized in this abstract way possibly yield “the right answer” when it is used to make a prediction? Such an undertaking could only be justified if predictions could thereby be guaranteed to be of “maximum likelihood” as a result. But a Bayesian would readily point out that restricting model usage to simply the making of such a prediction would lack integrity,

for knowledge of the width of the predictive likelihood function is as important to a user as knowledge of the point of highest likelihood.

In fact Bayes theorem does not even include the notion of “calibration”. Instead it speaks of a posterior parameter and/or predictive probability distribution - “posterior” in the sense of having included in computation of this distribution the fact that model parameters must be constrained such that model outputs under historical stresses replicate historical measurements of system state to within the level of measurement/structural noise associated with the latter. Once a posterior predictive probability distribution is known, the uncertainty associated with that prediction can be computed, this quantifying the range of values that this prediction may take with reasonable probability. Bayes theorem (and our intuition) readily informs us that predictive uncertainty is dependent on both hydraulic property uncertainty (as expressed by $C(\mathbf{p})$), the sensitivity of the prediction to different properties of the system, the extent to which values taken by these properties are constrained by the calibration process, and the extent to which such constraints are weakened by the presence of measurement noise.

2.2.2 Model Calibration

Because Bayesian analysis is often computationally burdensome, and because the culture of our industry is built on the ill-founded belief that models are the modern-day equivalent of crystal balls, most models are “calibrated” prior to being used for the making of predictions. (I have never found a word to be as misunderstood as “calibration”; many in our industry seem to believe that it implies that a model is invested with supernatural powers.) Once a model has been calibrated the concept of “predictive uncertainty” can no longer be used to characterize the fact that predictions made by that model may actually be wrong. Instead, analysis must focus on determining the error variance of predictions made by the model. This characterizes the probability distribution of “potential wrongness” of model predictions. The outcome of such an analysis is the likelihood that the “right answer” (which we will never know) is within a certain distance (in prediction space) of the prediction as made by the model at a certain level of confidence.

If care is taken in parameterising a model, and in estimating values for model parameters through the calibration process, the probability distribution of predictive error and that of predictive uncertainty will not be too different. However the latter will normally be smaller than the former. “Predictive uncertainty” is a “purer” concept than “predictive error”, being strictly a function of stochastic descriptors such as $C(\mathbf{p})$ and $C(\epsilon)$, rather than the details of how a unique solution was achieved to a fundamentally nonunique inverse problem, to thereby achieve a “calibrated model”.

2.3 Linear Predictive Uncertainty Analysis – the PREDUNC Suite

2.3.1 General

Using the PREDUNC suite (these are PEST utilities) linear uncertainty analysis can be undertaken on either a calibrated or uncalibrated model. In the former case parameter

values employed in the analysis should be those which are considered most likely from a hydrogeological perspective. Note that many more parameter types than are actually estimated during the calibration process can be employed in this analysis. Thus, for example, the values assigned to boundary conditions, system inputs such as recharge, even historical pumping rates (or multipliers thereof) can be included in the analysis; the contributions made by these various model components to the uncertainties of key model predictions can then be readily analysed.

2.3.2 Theoretical Basis

Members of the PREDUNC suite compute predictive uncertainty through conditioning of the pre-calibration prediction/observation covariance matrix in accordance with the observation dataset, and the noise associated therewith. See PEST documentation for full details.

2.3.3 Obtaining a JCO File

Use of any of the PREDUNC utilities is predicated on the assumption that a PEST control file exists (as well as template and instruction files cited therein), together with a Jacobian matrix (JCO) file computed on the basis of this PEST control file. The control file must instruct PEST to run in “estimation” mode. Weights must be assigned as the inverse of observation noise. The JCO file can be produced by setting the PEST NOPTMAX variable to -1 or -2 and then running PEST.

2.3.4 Predictive Sensitivity Vector

Use of the PREDUNC suite also requires the existence of a file holding the predictive sensitivity vector (i.e. the derivative of the prediction of interest with respect to each parameter involved in the analysis). Sensitivities must be stored in “PEST matrix file” format. Such a file is easily extracted from a JCO file using the JROW2VEC utility. The JCO file may be the same as that discussed above; this will be the case if the prediction is “carried” with the calibration dataset as formerly discussed. If not, a JCO file from which it can be extracted can be produced by undertaking an entirely different PEST run specifically for this purpose in which “observations” cited in the PEST control file are in fact predictions whose uncertainty is to be analysed. The weights assigned to these predictions do not matter (unless they are all zero, for PEST will then cease execution before the Jacobian matrix is calculated, saying that the objective function is zero); nor do their “observed values” as provided in the PEST control file.

2.3.5 Parameter Uncertainty File

Before using any members of the PREDUNC suite, the user must prepare a “parameter uncertainty file”. This encapsulates the $C(\mathbf{p})$ matrix – the matrix of innate parameter variability. $C(\mathbf{p})$ submatrices can be provided as individual parameter standard deviations (where these submatrices are diagonal), or as full matrix files (in PEST matrix file format). In the latter case, construction of these files may sometimes be assisted through

use of the PPCOV (Groundwater Data) Utility which builds a $C(\mathbf{p})$ matrix for pilot point parameters based on a user-specified variogram.

2.3.6 *PREDUNC1, PREDUNC4 and PREDUNC5*

PREDUNC1 computes both the pre- and post-calibration uncertainty associated with a specific prediction. PREDUNC4 allows the user to compute the contribution made to the uncertainty of a specified prediction by different parameter groups; both pre- and post-calibration uncertainty contributions are thus computed. PREDUNC5 allows the user to compute the worth of different observation groups (including individual observations) in making a specific prediction. Observation (group) worth can be computed as either the increase in predictive uncertainty accrued through omitting the observation(s) from the calibration dataset, or as the reduction from pre-calibration uncertainty achieved through using only that (those) observation(s) in constraining the posterior prediction probability distribution.

2.3.7 *Optimisation of Future Data Acquisition*

PREDUNC5 can be employed to test the worth of observations that have not even been made. This is possible because linear analysis requires only that observation and predictive sensitivities (and not actual observation and prediction values) be employed. Thus PREDUNC5 can be used as a basis for optimisation of future data acquisition. In this respect it is vastly superior to OPR-PPR statistics developed for this purpose by the USGS. The latter can only be employed after enough pre-calibration parsimonizing has been performed to reduce the dimensionality of the null space to zero. They thus fail to recognise the (mostly significant) contribution made to predictive uncertainty by inestimable parameterisation fine detail. In fact the less data that is presently available, and hence the more urgent is the need for supplementary data, the greater will be the amount of simplification required to compute these statistics, the broader will be the dimensions of the (ignored) null space, and the more misleading will these statistics be (as can be readily verified by studies based on synthetic models) What is the observation that is of most worth? It is normally that which tells a modeller something that he/she does not already know; by definition this information presently lies within the null space. Therefore, by their very construction, OPR-PPR statistics are flawed.

2.3.8 *Limitations of PREDUNC Analysis*

Unfortunately, analysis of the type undertaken by the PREDUNC suite is not possible if the observation dataset is too large, for matrices required in the analysis become too big to invert. In such cases the PREDVAR suite must be used instead. The results may be slightly different however, as PREDVAR analysis is based on a notional calibration exercise undertaken using truncated singular value decomposition on scaled parameters; thus its focus is on predictive error variance rather than predictive uncertainty. However while differences exist as a result of this, they are normally not significant enough to invalidate the analysis.

2.4 Linear Predictive Error Variance Analysis – the PREDVAR Suite

2.4.1 General

Members of the PREDVAR suite (i.e. PREDVAR1, PREDVAR4 and PREDVAR5) perform the same roles as respective members of the PREDUNC suite. These utilities, however, compute predictive error variance rather than predictive uncertainty (and parameter/observation contributions thereto). As stated above, predictive error variance is not quite as “pure” a quantity as predictive uncertainty; however its computation is feasible even where observations number in the tens of thousands.

2.4.2 Theoretical Basis

PREDVAR analysis is based on a notional calibration exercise undertaken using truncated singular value decomposition. The optimal singular value truncation level is computed by PREDVAR1; this can provide a sound basis for selection of an appropriate number of super parameters to use in a subsequent SVD-assisted calibration exercise (see the next section).

Because an explicit singular value decomposition of the (weighted) calibration Jacobian matrix is undertaken, members of the PREDVAR suite are able to compute contributions to predictive error made by the calibration solution and null spaces. This further distinguishes these utilities from those of the PREDUNC suite.

2.4.3 Parameter Scaling

It can be shown that predictive error variance is minimized where parameter estimation (whether notional as in PREDVAR analysis or actual as in real-life model calibration), is undertaken on the basis of parameters which are normalised with respect to their innate variabilities. Using the SCALEPAR utility, a new PEST control file, and its associated JCO file, can be built from an original PEST control file and the JCO file computed from it. This scaled pair should then be used for PREDVAR analysis, as the results of this analysis will then more closely resemble the results of PREDUNC analysis. Note however, that if this is done, predictive sensitivities must also be modified so that they pertain to scaled parameters rather than native parameters. This will happen automatically if predictions are being “carried” with observations in a single PEST control file as discussed above, but will require another SCALEPAR run if predictive sensitivities are computed using a separate PEST run (in which the model is presumably run using inputs and boundary conditions salient to the prediction). Note also that the parameter uncertainty file employed for PREDVAR analysis will also require modification when parameters are scaled. Fortunately it is an easy matter to make this modification, for all parameter variances become, by definition, 1.0.

2.4.4 Limitations of PREDVAR Analysis

If there are many observations, use of the PREDVAR suite becomes problematical if there are more than about 1800 parameters because the time required to undertake

singular value decomposition of the Jacobian matrix becomes excessive. At the time of writing it is planned to overcome this problem using the PROPACK solver (Larsen, 1998) in place of SVD.

2.5 Flow Chart

See Appendix A – “Parameter, Observation and Null Space Contributions to Predictive Uncertainty” - for a summary of procedures discussed in this section.

3. Model Calibration

3.1 Initial Parameters

As has already been stated, when undertaking highly-parameterized regularized inversion initial parameter values should be of highest likelihood from a geological point of view. Sometimes, however, assistance in the assignment of initial parameter values can be gained through undertaking traditional parameter estimation based on a small number of parameters (for example parameters that encompass entire model layers) prior to undertaking SVD-assisted estimation of parameters pertaining to finer system detail. This is especially useful in highly nonlinear contexts such as can occur if the calibration dataset includes contaminant concentration measurements.

3.2 Adding Regularisation

3.2.1 General

Let it be assumed that a PEST input dataset has been constructed. Let it be further assumed that in the PEST control file PEST is asked to run in “estimation” mode. If many parameters are cited, the inverse problem will probably be ill-posed; hence regularisation is required for its solution. PEST employs two broad types of regularisation, viz. Tikhonov and subspace regularisation. These can be employed separately or simultaneously. When undertaking SVD-assisted parameter estimation, the use of super parameters as required by this process constitutes a form of subspace regularisation; Tikhonov regularisation is implemented through the addition of constraints on parameter values (or differences between parameter values) embodied in a series of prior information equations.

3.2.2 The ADDREG1 Utility

The easiest way to add Tikhonov regularisation to a PEST control file is through the ADDREG1 utility (a PEST utility). ADDREG1 provides one prior information equation for each adjustable parameter, these informing PEST that each parameter should be preferentially equal to its initial value. Before running ADDREG1 however, it is advisable to assign parameters of different types (and/or parameters of the same type in different parts of the model domain) to different parameter groups. This allows ADDREG1 to assign prior information equations pertaining to these different parameter groups to different regularisation groups. When the PEST IREGADJ variable (part of the “regularisation” section of the PEST control file) is set to 1, PEST is then able to apply differential weighting to each resulting group of prior information equations, with weighting applied more strongly to those groups for which cited parameters are less informed by the calibration process. The PEST control file produced by ADDREG1 informs PEST that it must run in regularisation mode; it includes a “regularisation” section in the file which it writes, and sets IREGADJ to 1 in that section.

When PEST is run in regularisation mode a target objective function must be set. ADDREG1 assumes that this is 1.0; this almost certainly will require adjustment by the user. ADDREG1 also sets the FRACPHIM regularisation control variable to 0.1. Thus, during each iteration of the regularised inversion process, PEST temporarily sets the target measurement objective function to one tenth of the current value of the measurement objective function (allowing the beneficial effects of Tikhonov regularisation to be thereby felt from the very start of the parameter estimation process). However it is never set lower than the user-specified global target measurement objective function.

3.2.3 The GENREG Utility

If the type of regularisation added to a PEST control file by ADDREG1 is considered to be too simplistic, or to be inappropriate for use in a specific parameter estimation context, regularisation can be added to a PEST control file by other means. The reader's attention is drawn to the GENREG utility supplied with the Groundwater Data Utilities. GENREG facilitates formulation of complex suites of regularisation equations which, once formulated, are automatically added to an existing PEST control file.

3.3 Computing Sensitivities

3.3.1 Computation of Base Parameter Sensitivities

Prior to undertaking SVD-assisted parameter estimation, a Jacobian matrix file (i.e. a JCO file) must be created based on initial parameter values as listed in the PEST control file for the current parameter estimation problem. Ideally, it is better to do this after regularisation prior information equations have been added to this file than before this has occurred. This allows PEST to compute the first set of super parameter sensitivities “for free” during the ensuing SVD-assisted parameter estimation exercise. However, especially if super parameters are few in number, failure to do this does not incur a large computational cost – a cost, in fact, that may be very small compared with that associated with re-calculating an entire base parameter JCO file. (Note that in SVD-assist parlance, “base parameters” are the parameters actually employed by the model; “super parameters” are linear combinations of these base parameters as estimated by PEST through the SVD-assisted parameter estimation process.)

3.3.2 The JCO2JCO Utility

The JCO2JCO utility writes a JCO file corresponding to a new PEST control file based on that corresponding to an old PEST control file. The new PEST control file may differ from the old file in the following ways.

1. Prior information may have been added or subtracted.
2. Parameters and/or observations may have been removed.
3. Observation weights may have been altered.

4. The mode in which PEST is run (i.e. “estimation”, “prediction” or “regularisation”) may have been altered.
5. Parameters may be tied or fixed in the new PEST control file that were not tied or fixed in the old PEST control file.

Thus, for example, if a JCO file was produced on the basis a PEST control file for which PREDVAR or PREDUNC analysis was performed, and prior information was subsequently added to this file using the ADDREG1 utility prior to undertaking SVD-assisted parameter estimation, a JCO file pertaining to the modified PEST control file can be readily produced using JCO2JCO. While this file lacks sensitivities corresponding to the new prior information equations, it will nevertheless be useable by PEST for computation of super parameters upon commencement of the ensuing SVD-assisted parameter estimation process. If it is desired that the JCO file actually contain prior information sensitivities so that super parameter sensitivities can be computed “for free” during the first SVD-assist optimisation iteration, then PEST can be run for one iteration based on the prior-information-augmented base parameter PEST control file with the “/i” switch included in the PEST command line. Use of this switch informs PEST that it must read an existing JCO file to obtain sensitivities to all observations rather than calculate these itself on the basis of finite parameter differences. After having read this base parameter JCO file, it writes another JCO file in which the sensitivities to both observations and prior information equations are recorded.

3.4 Jacobian Manipulation Utilities

A number of utilities have been written to obviate the need to re-compute sensitivities if these sensitivities are available in part or in whole in existing JCO files. If model run times are long, use of these utilities can accrue large savings in computation time. See, for example. JCOPCAT (combines sensitivities with respect to different groups of parameters into a single JCO file), JCOADDZ (adds rows or columns to the Jacobian matrix contained in a JCO file), JCOORDER (re-orders rows/columns in a JCO file), etc.

3.5 Preparing for SVD-Assisted Parameter Estimation

3.5.1 The SVDAPREP Utility

The writing of a PEST control file for SVD-assisted parameter estimation is undertaken using the SVDAPREP utility. (On BPROC machines, however, it is better to use MSVDAPREP instead; this is documented with the MPEST version of PEST.) In the course of its execution, SVDAPREP requests a number of items of information from the user; in most cases default values for these items can be accepted by simply pressing the <Enter> key.

3.5.2 Number of Super Parameters

How many super parameters should be estimated? A prior PREDVAR1 analysis will suggest a suitable answer to this question (That is, the number of super parameters should

be roughly equal to the number of singular values at which predictive error variance is minimized.). However rarely will such an analysis have been undertaken prior to running SVDAPREP (if for no other reason than that a suitable prediction is not available for the purpose of PREDVAR-based predictive error variance analysis). In this case, trial and error may constitute an important part of the answer to this question.

On my first SVD-assisted PEST run, I normally employ a relatively small number of super parameters. Thus the parameter estimation process is completed quickly; alternatively, if there is anything wrong with either the model or PEST input dataset, I will quickly be informed of this through failure of the parameter estimation process. If the parameter estimation process does indeed appear to progress well, I inspect the condition number file; this has the same filename base as the PEST control file, but possesses an extension of “.cnd”. As long as condition numbers recorded in this file are less than about 10^3 , the inverse problem is well enough posed and the number of super parameters being estimated is therefore not unduly high. (Condition numbers even lower than this are often desirable, particularly where derivatives computation is compromised by numerical noise.) Alternatively, if I feel that I could justifiably obtain a lower measurement objective function than PEST has achieved with the current number of super parameters, and if condition numbers are low, I will normally increase the number of super parameters and repeat the SVD-assisted parameter estimation process.

3.6 SVDA and Internal Parameter Scaling

For reasons that are described in PEST documentation, it is best if super parameter definition takes place on the basis of scaled (by their sensitivities) parameters rather than native parameters. PEST provides a number of means by which appropriate internal parameter scaling coefficients can be computed – see a description of the SVDA_SCALADJ variable. In general, the best value to select for this variable (when running SVDAPREP to construct the super parameter PEST control file) is 2. However this value is only appropriate where some parameters are log transformed (for the sensitivities of such parameters are used as a basis for scale adjustment of non-log-transformed parameters). In other cases a value of 4 can be employed. However the user should then exercise caution in the setting of parameter bounds, for the bounded interval of each parameter is then employed as a basis for its scaling. (Note that it is common practice when using the FEHM model to allow FEHM, rather than PEST, to undertake log transformation of parameters such as permeability. In general it is better to allow PEST to undertake the logarithmic transformation, especially when undertaking SVD-assisted parameter estimation, because of the important role that log transformed variables play in determining parameter scaling coefficients. Note also that if log transformation is undertaken by FEHM rather than PEST, then absolute, rather than relative, parameter increments should be employed for finite-difference derivatives computation.)

3.7 SVD-Assisted Parameter Estimation

3.7.1 The PARLAM Variable

SVD-assisted parameter estimation is implemented by running PEST on the basis of a PEST control file produced by SVDAPREP. If using parallel PEST, (or if using PEST on a BPROC machine) set the PARLAM variable in the PEST run management file to -6 or higher (i.e. -5, -4 etc). A setting of -6 forces PEST to limit the number of trial Marquardt lambdas employed on any one parallel run package to 6. This is normally more than sufficient, and eliminates excessive wastage of model runs.

3.7.2 The PARREP Utility

When undertaking SVD-assisted parameter estimation, PEST actually estimates values for super parameters rather than the base parameters employed by the model. As in normal PEST operation, best parameter values are recorded in a PAR file (i.e. a file with the same base name as the PEST control file, but with an extension of “.par”). However the super parameter values recorded in this file have no meaning to anyone but PEST. On the other hand, best base parameter values are recorded in a BPA file. It is important to note that the filename base of this file is the same as that of the base parameter PEST control file and not that of the super parameter PEST control file. (It is thus over-written on subsequent SVD-assisted PEST runs based on the same base parameter set.)

When the SVD-assisted parameter estimation process is complete, a new base PEST control file citing optimised parameter values as initial parameter values can be built using the PARREP utility. NOPTMAX can then be set to 0 in this file and PEST then run. PEST will thus undertake one model run on the basis of optimised base parameter values. All model input files will thereby contain optimised parameter values, and all model output files will contain model outputs computed from these values. (Note that the undertaking of a final model run based on optimised parameters occurs automatically if PEST undertakes non-SVD-assisted parameter estimation. Unfortunately this cannot occur when SVD-assisted parameter estimation is complete as super parameter definitions may alter as the parameter estimation process proceeds, and previous definitions are lost to PEST. If optimal parameter values were computed using a lost definition, recovery of the optimal base parameter set is impossible for PEST.)

3.8 Obtaining Super Parameter Definitions

When PEST undertakes SVD-assisted parameter estimation, the values of base parameters are computed from current super parameter values prior to undertaking each model run by the PARCALC utility. SVDAPREP inserts this utility (together with PICALC for computation of prior information equations) into the model batch/script file which it writes as part of its preparations for an SVD-assisted parameter estimation run. Prior to each model run, PEST writes a file named *parcalc.in* for the use of PARCALC on that model run. PARCALC then undertakes super-to-base parameter translation and writes all model input files.

If, at any stage of the parameter estimation process, a user wishes to inspect the linear combination of base parameters that define any super parameter, the PCLC2MAT utility can be employed. When inspecting the PCLC2MAT output file, the following should be remembered.

1. In definition of any super parameter, the squares of base parameter coefficients sum to 1.
2. The base parameter composition of any super parameter can change as the parameter estimation process progresses; this will occur if any base parameter hits its upper or lower bound in the course of this process.

3.9 Flow Chart

See Appendix B – “SVD-Assisted Model Calibration” - for a summary of procedures discussed in this section.

4. Linear Post-Calibration Error Variance Analysis

4.1 Predictive Error Variance Analysis

4.1.1 General

Once a model has been calibrated, it is an easy matter to compute the error variance of any prediction made by that model using linear analysis. This analysis takes full account of the null space contribution to predictive error variance. Unlike PREDVAR-based predictive error analysis discussed in Section 2, the analysis described below is implemented in the context of a specific calibration exercise, and is based on information forthcoming from that exercise.

4.1.2 RESPROC and Related Utilities

Whenever PEST undertakes any form of regularised inversion, it writes a “regularisation data file”, this having the same filename base as the PEST control file, but possessing an extension of “.rsd”. This is a binary file, and so is not available for direct inspection. Before any post-calibration predictive error analysis is undertaken, this file (together with other files written by PEST) should be processed using the RESPROC utility. RESPROC calculates the **R** and **G** matrices used in linear predictive error variance analysis (see PEST documentation for more details) from the information contained in these files and stores these matrices in binary form (they are often very large). If desired, the RESWRIT utility can be used to re-write these matrices in PEST matrix file format (these are ASCII files). MAT2SRF can then be employed to re-write the resolution matrix (i.e. the **R** matrix) as a SURFER grid file; the diagonal dominance (or otherwise) of this matrix can then be inspected using SURFER.

4.1.3 The PREDERR Utility

The error variance of a model prediction can be computed using the PREDERR utility (or any other member of the PREDERR suite). PREDERR obtains sensitivities of the prediction of interest directly from a user-nominated row of a sensitivity matrix embodied in a JCO file; thus these sensitivities do not need to be extracted from this file using the JROW2VEC utility. This will be the same JCO file as written during the calibration process if the prediction was in fact “carried” through this process. Otherwise, these sensitivities will reside in another JCO file. However that JCO file must cite the same parameters in the same order as the JCO file produced during the calibration process. Ideally, it should also have been produced on the basis of calibrated parameter values.

4.1.4 Re-Computation of Parameter Sensitivities

Computation (by RESPROC) of post-calibration **R** and **G** matrices after completion of SVD-assisted parameter estimation is a complicated matter; see PEST documentation for details. Accuracy of computation of these matrices (and hence of linear predictive error variance analysis) is enhanced if base parameter sensitivities are re-calculated using

optimised parameters after SVD-assisted parameter estimation is complete. This is easily achieved if NOPTMAX is set to -1 rather than 0 in the post-calibration PARREP-produced PEST control file described above in which initial parameter values are in fact optimised parameter values. Note that this PEST control file should have a different name from that employed for computation of base parameter sensitivities prior to SVD-assisted parameter estimation in order to prevent overwriting of the original JCO file used by PEST for super parameter definition; this is necessary because both of these JCO files are needed by RESPROC. (However in the event that one of these JCO files is not present, errors incurred by using the other twice - particularly the post calibration JCO file - are not great.)

4.1.5 Observation Uncertainty File

Because weights employed in the calibration process will not necessarily be proportional to the inverse of observation uncertainties, the latter must be provided to PREDERR in a separate “observation uncertainty file”, this embodying the $C(\epsilon)$ matrix discussed above. See PEST documentation for construction details of this file. A parameter uncertainty file (embodying the $C(p)$ matrix) must also be provided. However no parameter scaling is necessary as is recommended when undertaking PREDVAR analysis; as stated above, predictive error variance is computed by PREDERR on the basis of an actual, rather than a notional, calibration exercise.

4.1.6 Parameter Error Covariance Matrix

If desired, the parameter error covariance matrix can be computed using the PARAMERR utility. However I rarely, if ever, do this.

4.2 Post-Calibration PREDVAR and PREDUNC Analysis

There is no reason why PREDVAR and PREDUNC analysis should not be performed on the basis of a calibrated parameter set rather than on the basis of a pre-calibration parameter set as described in Section 2 of this document. In fact, certain aspects of this analysis (such as parameter and observation contributions to predictive error variance and optimisation of data acquisition) are better undertaken on the basis of a calibrated parameter set than an initial parameter set.

4.3 Flow Chart

See Appendix C – “Post-Calibration Linear Error Variance Analysis” - for a summary of procedures discussed in this section.

5. Nonlinear Post-Calibration Error Variance Analysis

5.1 Constrained Maximization/Minimization of a Key Model Prediction

5.1.1 Solution Space and Null Space Constraints

If run in predictive analysis mode after calibrating a highly parameterised model using regularised inversion, PEST can be asked to maximize/minimize a key model prediction while ensuring that:-

1. The measurement objective function does not change by more than a certain amount from that obtained during the calibration process; and
2. Parameters do not change by more than a certain amount from that obtained during the calibration process.

It should be noted that:-

1. The above problem is formulated in terms of parameter *changes* rather than absolutes.
2. For the objective function “a certain amount” is assessed in terms of $C(\epsilon)$, the stochastic descriptor of measurement noise.
3. For parameters “a certain amount” is assessed in terms of $C(p)$, the stochastic descriptor of innate parameter variability.

Constraints on the objective function imposed by $C(\epsilon)$ translate to constraints on changes allowed to parameter combinations occupying the calibration solution space. Constraints imposed by $C(p)$ are actually enforced on parameter changes projected onto the calibration null space. Thus the two components of the dual-component objective function are independent. This allows formulation of a convenient objective function limit based on normal variates through which the maximisation/minimization process can thereby be used to define a predictive confidence interval.

5.1.2 REGPRED and Related Utilities

Constrained predictive maximization/minimization can only be undertaken as an adjunct to a prior calibration process. After the latter is complete, PARREP must be employed to create a PEST control file based on optimised parameters. OBSREP is then employed to create a PEST control file in which “observations” are in fact model outputs at observation points under calibration conditions. A predictive analysis PEST control file in which both of the $C(\epsilon)$ and $C(p)$ constraints are employed is then built using the REGPRED utility. REGPRED also alters the model batch file such that the VECLOG, MATDIFF and MATPROD utilities are run as part of the model in order to undertake necessary null space parameter projections.

5.1.3 Using SVD-Assist for Nonlinear Predictive Error Variance Analysis

A problem in applying the above methodology in a highly parameterised context is that a large number of parameters are potentially involved in the predictive error variance analysis process; hence the number of model runs per iteration may be very high. Contrast this with SVD-assisted parameter estimation where the use of super parameters reduces the number of model runs per iteration to a potentially very low number. If a super parameter set could be defined to span a “predictive solution subspace” in the same manner that a super parameter set is defined to span the “calibration solution subspace” during SVD-assisted parameter estimation, then the predictive analysis process could also be undertaken with great computational efficiency.

This in fact can be done. PEST can comfortably run in “predictive analysis” mode, with its SVD-assist process activated. Furthermore, super parameter definitions can be supplied by the user rather than computed by PEST on the basis of parameter sensitivities pertaining to a calibration dataset. Tonkin, Doherty and Moore (2007) show how super parameters can actually be defined on the basis of “surrogate predictions” whose computation is very fast. Once an optimal predictive solution subspace is defined on the basis of sensitivities of these surrogate predictions to all parameters, the full model can be run by PEST in an SVD-assisted predictive analysis process in which a true model prediction is maximised/minimised. Only as many model runs per iteration are required as there are super parameters defined. Though no utilities are provided to specifically automate computation of predictive super parameters, the JCO2MAT and MATSVD utilities can be useful in this endeavour.

5.2 Null Space Monte Carlo

5.2.1 Discrete and Continuous Stochastic Fields

Null space Monte Carlo analysis will be described only for the case where model parameterisation is comprised of a large number of zones of piecewise constancy and/or where parameters are lumped by other means. Its use in cases where continuous stochastic parameter distributions are employed to characterise parameter heterogeneity is only slightly more complicated. See documentation of the PPSAMP utility from the Groundwater Data Utility suite for more details.

5.2.2 The RANDPAR and PNULPAR Utilities

In implementing null space Monte Carlo analysis, different realisations of parameter sets are first generated on the basis of an appropriate $C(\mathbf{p})$ matrix using the RANDPAR utility (a PEST utility). The calibrated parameter field is then subtracted from each such set. The resulting parameter differences are then projected onto the calibration null space (or an approximation to it), and the projected differences are then re-added to the calibrated parameter field; these steps are implemented using the PNULPAR utility. If the model were perfectly linear each new parameter set thus obtained would calibrate the model to within confidence limits set by $C(\epsilon)$.

5.2.3 Model Re-Calibration using Pre-Calculated Sensitivities

However because model linearity is a phenomenon only rarely encountered in real-world modelling practice, each newly parameterised model obtained in the above manner must be re-calibrated. In many cases this is easily achieved in only one optimisation iteration of SVD-assisted parameter estimation. Furthermore, the same set of super parameters can be employed for all such re-calibration exercises. In addition to this, because PEST is able to provide the first parameter update “for free” in an SVD-assisted parameter estimation process, finite-difference computation of super parameter sensitivities is not required. Thus only a few model runs need to be undertaken per random parameter set in order to create a suite of entirely different model parameterisations, all of which calibrate the model on the one hand (to within limits set by $C(\epsilon)$) and all of which respect parameter reality constraints on the other hand (to within limits set by $C(\mathbf{p})$). If it transpires that more than one optimisation iteration is needed to restore a particular parameter set to calibration status, then PEST will indeed require that super parameter derivatives be computed for the second optimisation iteration; this will require as many model runs as there are super parameters. Alternatively if, on the basis of just one pre-calibration model run, it is obvious that a stochastic parameter set computed in the manner described above is unlikely to be warped back into calibration on just one iteration because its calibration status is too bad, it can be rejected before re-calibration is attempted.

5.2.4 Steps Required for Null Space Monte Carlo Analysis

Null space Monte-Carlo analysis can be implemented as follows..

1. First calibrate a model using SVDA. The base PEST control file will probably include Tikhonov regularisation (expressed as prior information), though this is not essential.
2. Once the calibration process is complete, build a new base parameter PEST control file in which initial parameters are in fact optimised parameters; use the PARREP utility to achieve this.
3. Remove all prior information from this new base PEST control file; inform PEST that it must run in “estimation” mode.
4. Set NOPTMAX to -1 or -2 in this file and run PEST to compute a new JCO file.
5. Now set NOPTMAX to 1 in this base parameter PEST control file. This will be referred to as the “primary base PEST control file”.
6. Make a copy of this file; this will be referred to as the “reference base PEST control file”.
7. Run SVDAPREP to build a super parameter PEST control file based on the “primary base PEST control file”. Make sure that NOPTMAX is set to 1 in this new super parameter PEST control file. Also make sure that Broyden Jacobian

- upgrading is allowed (see PEST documentation for a discussion of this feature) and that the Marquardt lambda control variable PHIRATSUF is set to 0.1 rather than 0.3.
8. Create a parameter uncertainty file which encapsulates the $C(\mathbf{p})$ matrix.
 9. Run the RANDPAR utility on either the “primary base” or the “reference base” PEST control file to produce a series of “parameter value files” encapsulating realisations of $C(\mathbf{p})$.
 10. Run the PNULPAR utility to create calibration-constrained random parameter sets (each residing in a parameter value file of its own) citing the name of the “reference base PEST control file” in response to the pertinent prompt.
 11. Run the PARREP utility using a command of the type:-

PARREP *random_par reference_base primary_base*

In the above command *random_par* is the name of a parameter value file produced by PNULPAR, *reference_base* is the name of the reference base PEST control file and *primary_base* is the name of the primary base PEST control file. Thus the primary base PEST control file now contains projected parameters added to calibrated parameters as its initial parameter set.

12. Run PEST on the super parameter PEST control file built from the reference base PEST control file.

Steps 11 and 12 should be repeated for each PNULPAR-generated parameter value file. One iteration of SVD-assisted parameter estimation will be undertaken on each random parameter field if the above steps are taken. If desired, alterations to the above procedure may include the following.

1. A “filtering” step may be added whereby a model run is undertaken for each new parameter set before SVD-assisted parameter estimation is attempted. If the objective function is too high, the parameter field is rejected and re-calibration is not undertaken. (This process would be made more efficient if it were undertaken internally by PEST rather than requiring a separate model run which is then repeated by PEST; this issue will be addressed as further PEST development takes place.)
2. NOTMAX could be set to 2 or higher in the super parameter PEST control file to allow more optimisation iterations than one to take place. However, for these latter iterations, PEST will compute super parameter derivatives, this requiring that one model run be undertaken per super parameter. In contrast the first iteration of the SVD-assisted parameter estimation process requires no super parameter derivatives calculation as super parameter derivatives are computed from pre-calculated base parameter derivatives.

5.3 Flow Chart

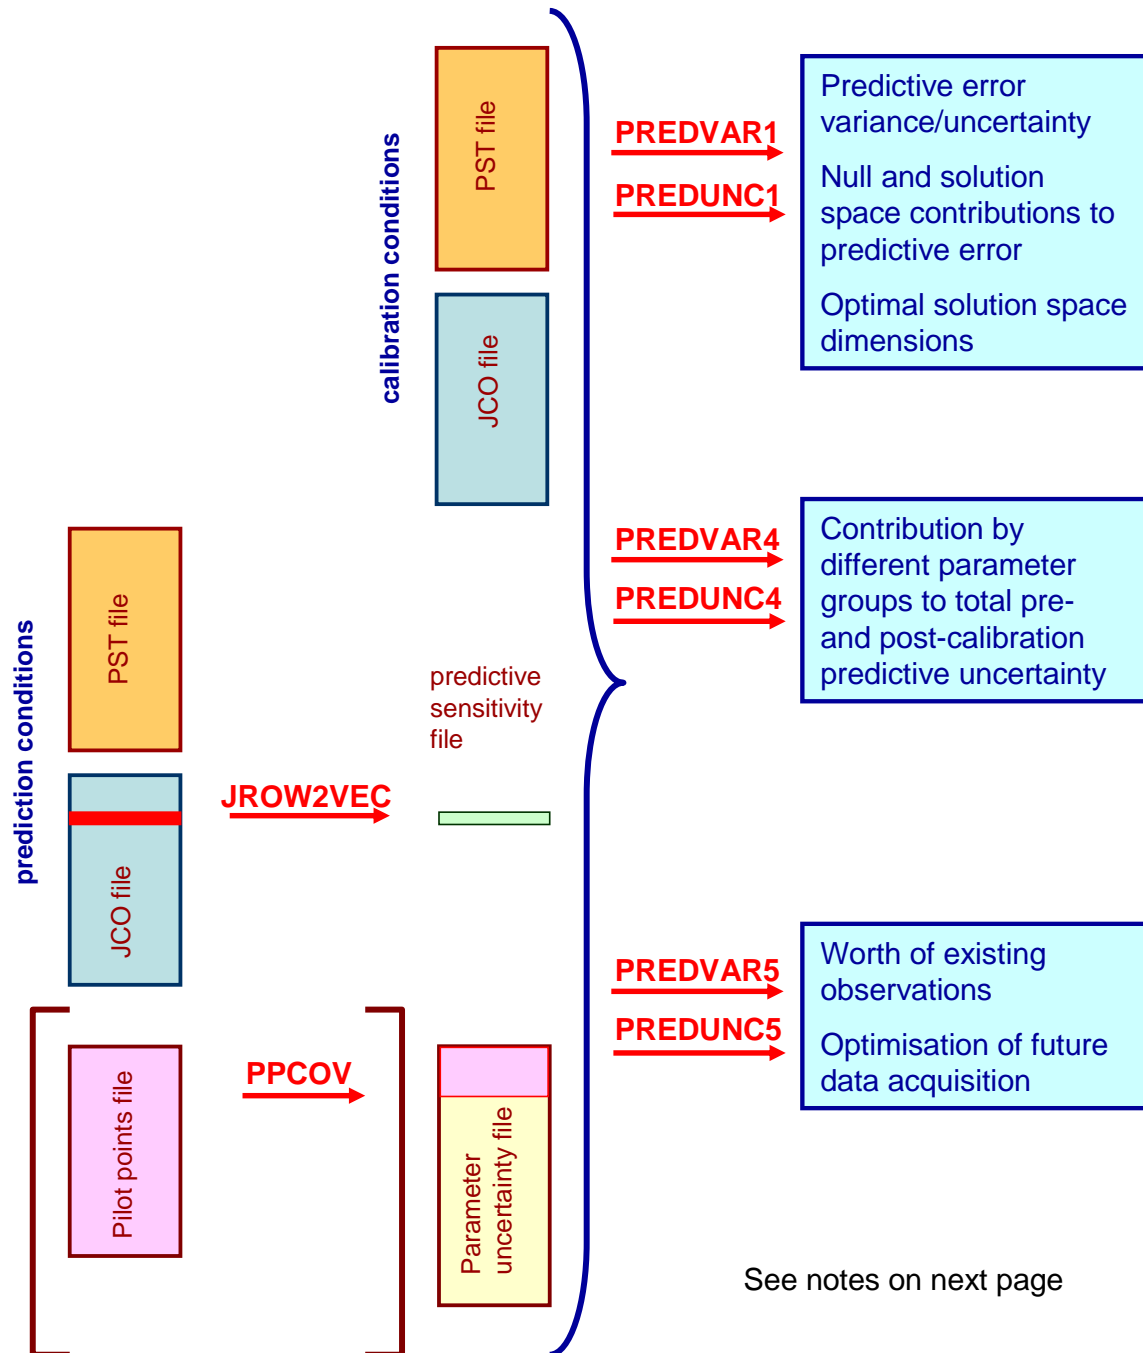
See Appendix D – “Post-Calibration Nonlinear Error Variance Analysis” - for a summary of procedures discussed in this section.

6. References

- Kavetski, D., Kuczera, G., Franks, S.W., 2005. Calibration of conceptual hydrological models revisited: 1. Overcoming numerical artefacts. *Journal of Hydrology*, this issue doi: 10.1016/j.jhydrol.2005.07.012.
- Larsen, R.M., 1998. Lanczos bidiagonalization with partial reorthogonalization, Department of Computer Science, Aarhus University, Technical report, DAIMI PB-357, September 1998. This report describes the initial work that lead to PROPACK.
- Moore, C. and Doherty, J., 2005. The role of the calibration process in reducing model predictive error. *Water Resources Research*. Vol 41, No 5. W05050.
- Moore, C., Doherty, J., 2006. The cost of uniqueness in groundwater model calibration. *Advances in Water Resources*. Volume 29, Issue 4, April, pages 605 – 623.
- Tonkin, M.J., Doherty, J. and Moore, C.R. (in press). Efficient Non-Linear Predictive Error Variance for Highly Parameterized Models. Accepted for Publication, *Water Resources Research*.

Appendix A

Parameter, Observation and Null Space Contributions to Predictive Uncertainty

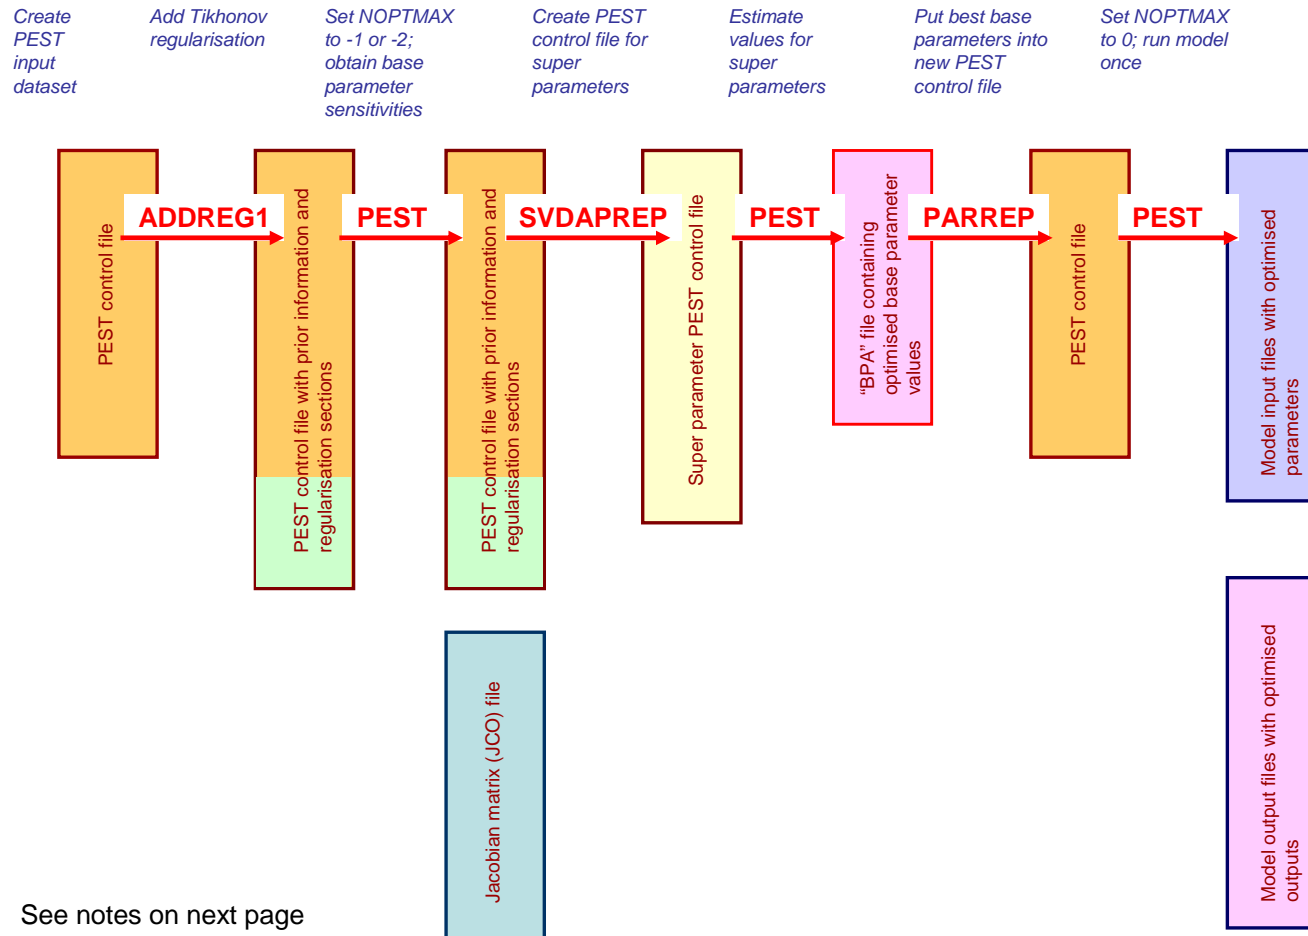


Notes on Previous Figure

1. The model does not need to be calibrated in order to conduct this analysis.
2. Observations employed in the analysis do not need to have been acquired; only sensitivities, and not actual observation values, are used in the analysis.
3. The calibration PEST input dataset should contain no regularisation prior information.
4. All weights within the calibration PEST control file should be proportional to the inverse of observation uncertainty (including that induced by “structural noise”).
5. The parameter uncertainty file is easily prepared with a text editor, especially where parameters are uncorrelated. Where correlation exists, the PPCOV utility can be used to build a covariance matrix file for a parameter subset. PPCOV is a member of the Groundwater Data Utilities.

Appendix B

SVD-Assisted Model Calibration

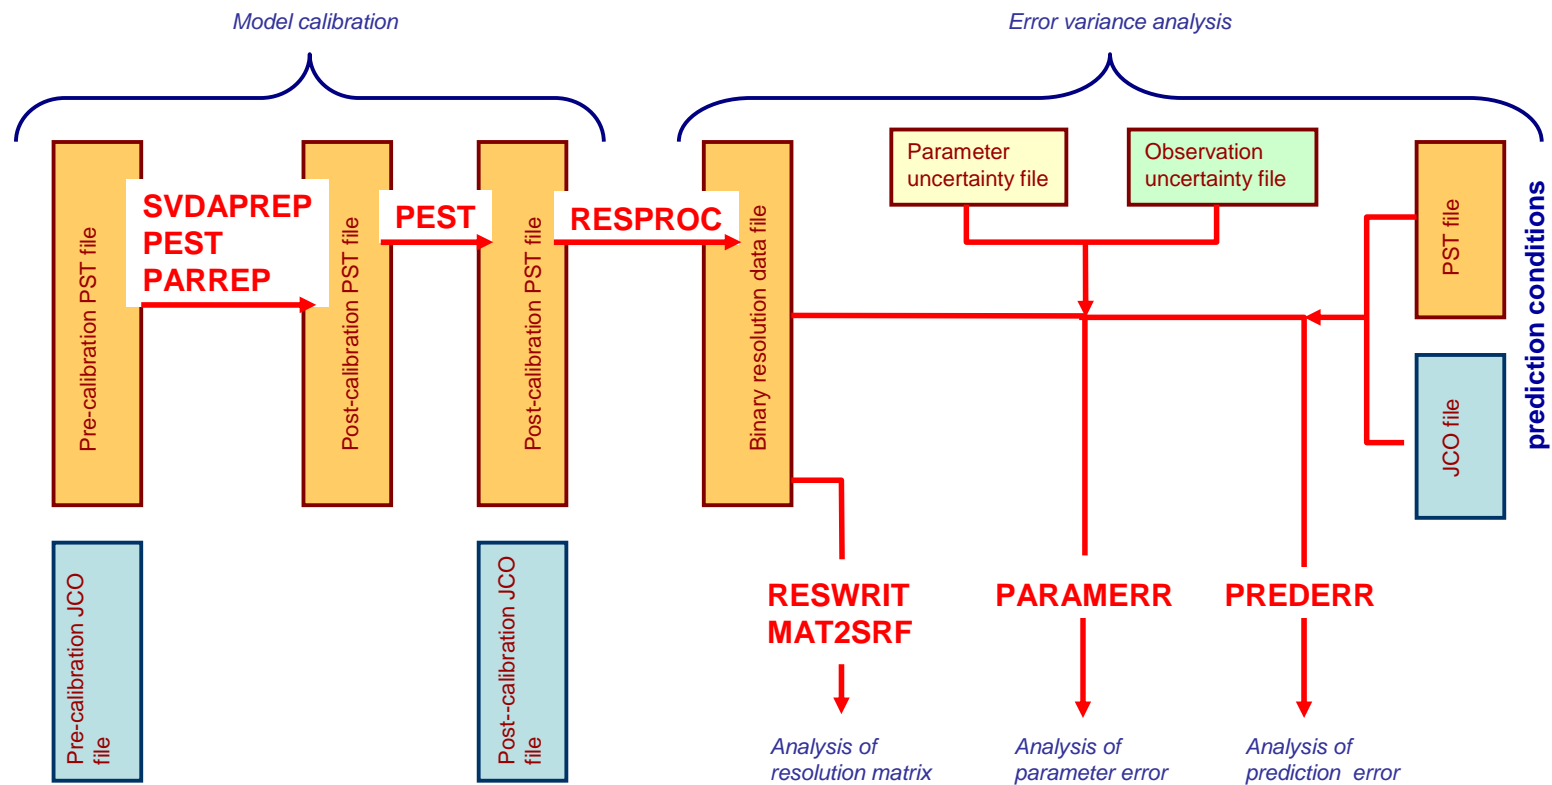


Notes on Previous Figure

1. The integrity of base or super parameter derivatives can be tested using the JACTEST utility.
2. The JCO2JCO utility can be employed to build a new JCO file from an existing one where changes to a PEST control file are restricted to:-
 - a. adding or subtracting prior information,
 - b. fixing, tying, re-transforming or removing parameters,
 - c. re-weighting or removing observations.
3. Use MPEST and MSVDAPREP instead of PEST and SVDAPEST on BPROC machines.

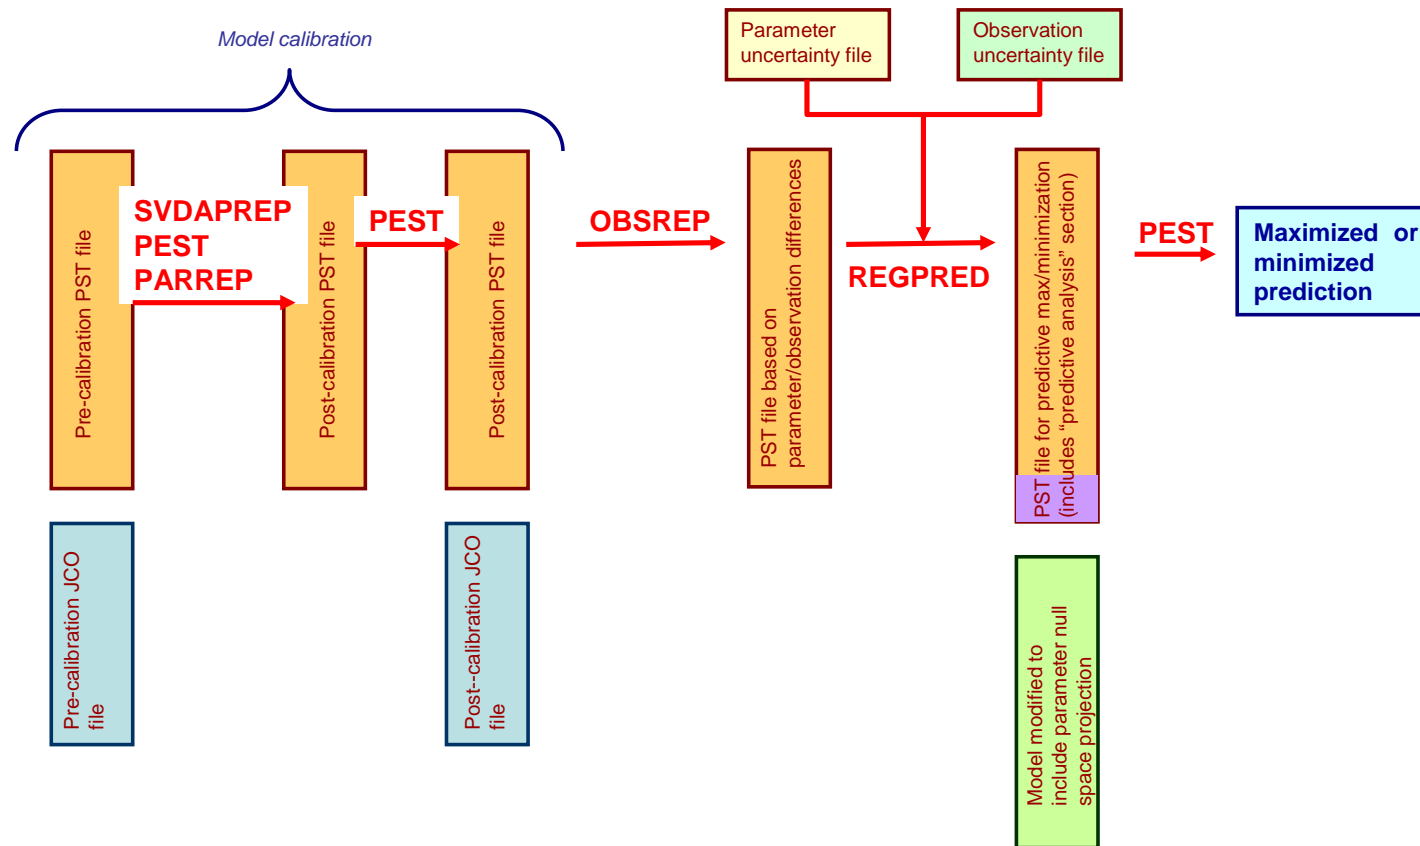
Appendix C

Post-Calibration Linear Error Variance Analysis



Appendix D

Post-Calibration Nonlinear Error Variance Analysis



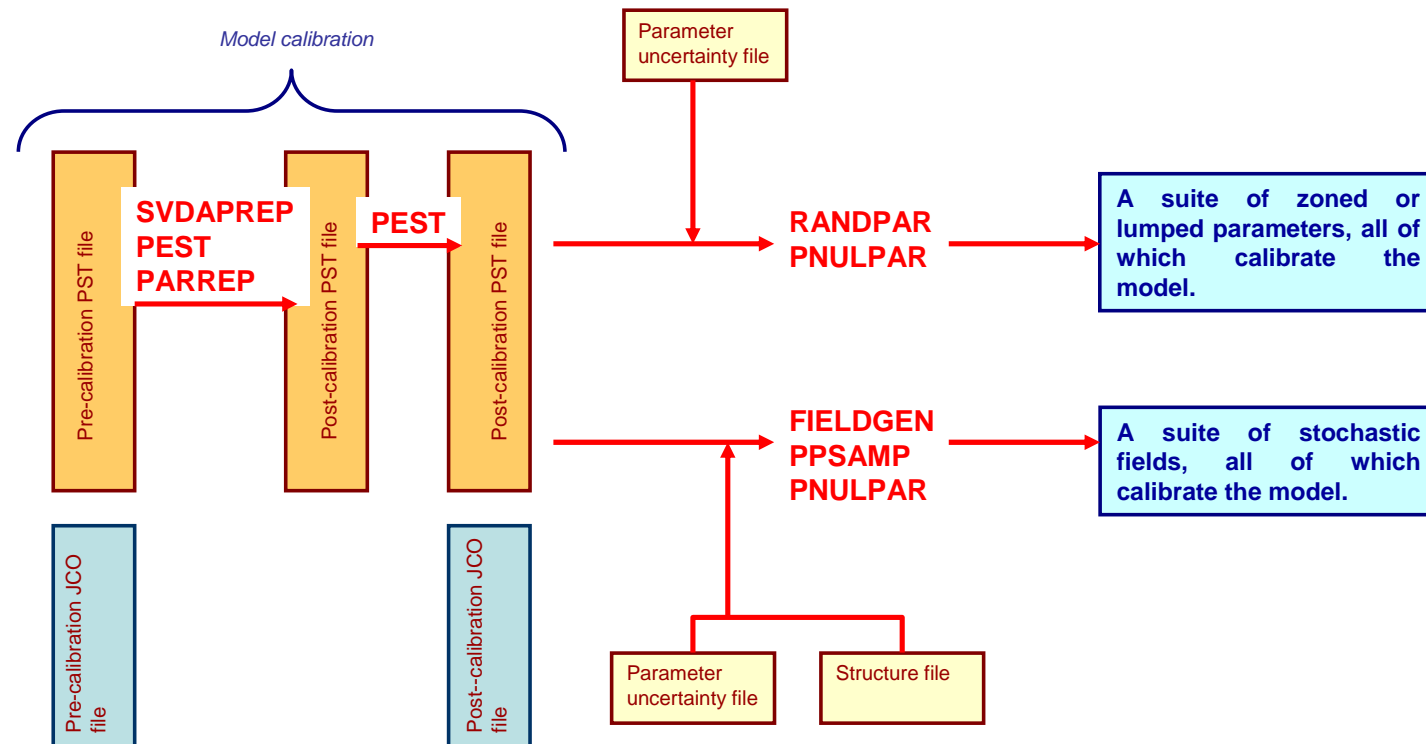
See notes on next page

Notes on Previous Figure

1. The model batch file may need to be modified further than REGPRED's alterations in order to run the model under predictive conditions as well as under calibration conditions.
2. The predictive maximization/minimization process can be made much more efficient if predictive analysis is undertaken on the basis of super-parameters defined to span the "predictive solution space". PEST's ability to undertake SVD-assisted predictive analysis on the basis of user-defined super parameters allows this. See Tonkin, Doherty, Moore (2007)

Appendix E

Null-space Monte Carlo Analysis



See notes on next page

Notes on Previous Figure

1. Null space Monte Carlo Analysis is actually a little more complicated than suggested in the previous figure. It involves the completion of an (SVD-assisted) PEST run for each random parameter field. However as each run may involve only one PEST iteration, and as super-parameter sensitivities for all such runs can be identical (and pre-calculated), the generation of multiple parameter sets or fields which, on the one hand, respect specified parameter stochastic distributions while, on the other hand, calibrate the model, can be numerically quite cheap.