Water supply security

For the township of Biggenden

A GMDSI worked example report

by Mark Gallagher and John Doherty

gmdsi
Groundwater Modelling Decision Support Initiative

BHP
Flinders University
National Centre for Groundwater Research and Training
RioTinto

PUBLISHED BY

The National Centre for Groundwater Research and Training
C/O Flinders University
GPO Box 2100
Adelaide SA 5001
+61 8 8201 2193

DISCLAIMER

CITATION

For bibliographic purposes this report may be cited as: Gallagher, M. and Doherty J. (2020). Water supply security for the township of Biggenden: A GMDSI worked example report. National Centre for Groundwater Research and Training, Flinders University, South Australia.

ISBN: 978-1-925562-39-2

DOI:10.25957/5f4c5815fc0d1. DOI as a weblink https://doi.org/10.25957/5f4c5815fc0d1

COPYRIGHT

PREFACE

The Groundwater Modelling Decision Support Initiative (GMDSI) is an industry-funded and industry-aligned project focused on improving the role that groundwater modelling plays in supporting environmental management and decision-making. Over the life of the project, it will document a number of examples of decision-support groundwater modelling. These documented worked examples will attempt to demonstrate that by following the scientific method, and by employing modern, computer-based approaches to data assimilation, the uncertainties associated with groundwater model predictions can be both quantified and reduced. With realistic confidence intervals associated with predictions of management interest, the risks associated with different courses of management action can be properly assessed before critical decisions are made.

GMDSI worked example reports, one of which you are now reading, are deliberately different from other modelling reports. They do not describe all of the nuances of a particular study site. They do not provide every construction and deployment detail of a particular model. In fact, they are not written for modelling specialists at all. Instead, a GMDSI worked example report is written with a broader audience in mind. Its intention is to convey concepts, rather than to record details of model construction. In doing so, it attempts to raise its readers' awareness of modelling and data-assimilation possibilities that may prove useful in their own groundwater management contexts.

The decision-support challenges that are addressed by GMDSI worked examples include the following:

- assessing the reliability of a public water supply;
- protection of a groundwater resource from contamination;
- design of a mine dewatering system;
- assessing the environmental impacts of mining; and
- management of aquifers threatened by salt water intrusion.

In all cases the approach is the same. Unwanted outcomes of a management strategy are identified. The ways in which modelling is used to explore whether these outcomes are possible, given all information that is available at the time of modelling, are reported. Model design choices are explained in a way that modellers and non-modellers can understand.

The authors of GMDSI worked example reports make no claim that the modelling work which they document cannot be improved. As all modellers know, time and resources available for modelling are always limited. The quality of data on which a model relies is always suspect. Modelling choices are always subjective, and are sometimes made differently with the benefit of hindsight.

What we do claim, however, is that the modelling work which we report has attempted to implement the scientific method to address challenges that are typical of those encountered on a day-to-day basis in groundwater management worldwide.

We thank and acknowledge our collaborators, and GMDSI project funders, for making these reports possible.

Dr John Doherty, GMDSI Project, Flinders University and Watermark Numerical Computing.
Dr Phil Hayes, GMDSI Project, University of Queensland.

Glossary

Anisotropy

A condition whereby the properties of a system (such as hydraulic conductivity) are likely to show greater continuity in one direction than in another. At a smaller scale it describes a medium whose properties depend on direction.

Bayesian analysis

Methods that implement history-matching according to Bayes equation. These methods support calculation of the posterior probability distribution of one or many random variables from their prior probability distributions and a so-called "likelihood function" – a function that increases with goodness of model-to-measurement fit.

Boundary condition

The conditions within, or at the edge of, a model domain that allow water or solutes to enter or leave a simulated system.

Boundary conductance

The constant of proportionality that governs the rate of water movement across a model boundary in response to a head gradient imposed across it.

Time-variant specified head (CHD) package

A Dirichlet (i.e. "fixed head") boundary condition implemented by MODFLOW in which the head can vary with time on a stress-period-by-stress-period basis.

Covariance matrix

A matrix is a two-dimensional array of numbers. A covariance matrix is a matrix that specifies the statistical properties of a collection of random variables - that is, the statistical properties of a random vector. The diagonal elements of a covariance matrix record the variances (i.e. squares of standard deviations) of individual variables. Off-diagonal matrix elements record covariances between pairs of variables. The term "covariance" refers to the degree of statistical inter-relatedness between a pair of random variables.

Ensemble

A collection of realisations of random parameters.

Drain (DRN) package

A one-way Cauchy boundary condition implemented by MODFLOW. Water can flow out of a model domain, but cannot enter a model domain through a DRN boundary condition.

Evapotranspiration (EVT) package

MODFLOW's implementation of water withdrawal from a groundwater system whereby the extraction rate can increase, up to a user-supplied maximum, as the head approaches a user-prescribed level from below. In MODFLOW 6, the rate of increase of extraction with head can be linear, or linear-segmented.

General head boundary (GHB) package

This is MODFLOW parlance for a Cauchy boundary condition. Water flows into or out of a model domain in proportion to the head difference between that ascribed to the boundary and that calculated for neighbouring cells. The rate of water movement through the boundary in response to this head differential is governed by the conductance assigned to the boundary.

Hydraulic conductivity

The greater is the hydraulic conductivity of a porous medium, the greater is the amount of water that can flow through it in response to a head gradient.

Jacobian matrix

A matrix of partial derivatives (i.e. sensitivities) of model outputs (generally those that are matched with field measurements) with respect to model parameters.

MODFLOW

A family of public-domain, finite-difference groundwater models developed by the United States Geological Survey (USGS).

MODFLOW package

An item of simulation functionality that describes one aspect of the operation of a groundwater system, for example recharge or a boundary condition. The word “package” describes the computer code that implements this functionality, as well as its input and output file protocols.

Null space

In the parameter estimation context, this refers to combinations of parameters that have no effect on model outputs that are matched to field observations. These combinations of parameters are thus inestimable through the history-matching process.

Objective function

A measure of model-to-measurement misfit whose value is lowered as the fit between model outputs and field measurements improves. In many parameter estimation contexts the objective function is calculated as the sum of squared weighted residuals.

Parameter

In its most general sense, this is any model input that is adjusted in order to promulgate a better fit between model outputs and corresponding field measurements. Often, but not always, these inputs represent physical or chemical properties of the system that a model simulates. However there is no reason why they cannot also represent water or contaminant source strengths and locations.

Phreatic surface

The water table.

Pilot point

A type of spatial parameterisation device. A modeller, or a model-driver package such as PEST or PEST++, assigns values to a series of points in two- or three-dimensional space. A model pre-processor then undertakes spatial interpolation from these points to cells comprising the model grid or mesh. This allows parameter estimation software to ascribe hydraulic property values to a model on a pilot-point-by-pilot-point basis, while a model can accept these values on a model-cell-by-model-cell basis. The number of pilot points used to parameterize a model is generally far fewer than the number of model cells.

Prior probability

The pre-history-matching probability distribution of random variables (model parameters in the present context). Prior probability distributions are informed by expert knowledge, as well as by data gathered during site characterisation.

Posterior probability

The post-history-matching probability distribution of random variables (model parameters in the present context). These probability distributions are informed by expert knowledge, site characterisation studies, and measurements of the historical behaviour of a system.

Probability density function

A function that describes how likely it is that a random variable adopts different ranges of values.

Probability distribution

This term is often used interchangeably with “probability density function”.

Quadtree mesh refinement

This term refers to a means of creating fine rectilinear model cells from coarse rectilinear model cells by dividing them into four. Each of the subdivided cells can then be further subdivided into another four cells. However it is a design specification of a quadtree refined grid that no cell within the domain of a model be connected to more than two neighbouring cells along any one of its edges.

Realisation

A random set of parameters.

Regularisation

The means through which a unique solution is sought to an ill-posed inverse problem. Regularisation methodologies fall into three broad categories, namely manual, Tikhonov and singular value decomposition.

Residual

The difference between a model output and a corresponding field measurement.

Singular value decomposition (SVD)

A matrix operation that creates orthogonal sets of vectors that span the input and output spaces of a matrix. When undertaken on a Jacobian matrix, SVD can subdivide parameter space into complementary, orthogonal subspaces; these are often referred to as the solution and null subspaces. Each of these subspaces is spanned by a set of orthogonal vectors. The null space of a Jacobian matrix is composed of combinations of parameters that have no effect on matrix outputs, and hence are inestimable.

Solution space

The orthogonal complement of the null space. This is defined by undertaking singular value decomposition on a Jacobian matrix.

Specific storage

The amount of water that is stored elastically in a cubic metre a porous medium when the head of water in which that medium is immersed rises by 1 metre.

Specific yield

The amount of accessible water that is stored in the pores of a porous medium per volume of that medium.

Stochastic

A stochastic variable is a random variable.

Stress

This term generally refers to those aspects of a groundwater model that cause water to move. They generally pertain to boundary conditions. User-specified heads along one side of a model domain, extraction from a well, and pervasive groundwater recharge, are all examples of groundwater stresses.

Stress period

The MODFLOW family of models employs this terminology to describe each member of a series of contiguous time intervals that collectively comprise the simulation time of a model. Within each of these time intervals, all model stresses are assumed to be time-invariant.

Tikhonov regularisation

An ill-posed inverse problem achieves uniqueness by finding the set of parameters that depart least from a user-specified condition, often one of parameter equality and hence spatial homogeneity.

Executive Summary

This GMDSI worked example report describes a groundwater modelling project that is far from ordinary. Construction and calibration of the model rests on a comparatively small amount of data. As such, it addresses the commonly asked question, “how much data does one need in order to justify construction of a groundwater model to support the making of a decision?”. The response is that the premise of the question is incorrect. Lack of data instils high uncertainties in some model predictions. It is the model’s task to quantify these uncertainties. If a decision must be made in the midst of uncertainty, then it is better that this uncertainty be known than not. Furthermore, if even a small amount of data can support reduction of this uncertainty, then it is better that it be reduced through appropriate model-based data-processing than not.

As it turns out, uncertainties ascribed to decision-critical predictions made by the model that is the subject of this report are smaller than were anticipated. This is because these predictions are closely related to observations that comprise the history-matching dataset. Assimilation of these data was an important component of the model construction and deployment process. However this alone did not ensure the model’s integrity in exploring the behaviour of the simulated system under future conditions that are likely to be more extreme than those for which data exist in the past. Model construction, parameterization and history-matching was also tuned to the assimilation of information contained in “soft data” that encapsulates qualitative and semi-quantitative knowledge of system behaviour. Constraints imposed on model parameters that ensure adherence to this behaviour mitigate aberrant model performance when it is asked to simulate future climatic conditions that, on occasion, may become extreme.

The model described herein was built to assess water supply security for the township of Biggenden in south east Queensland, Australia. Biggenden draws its water from a small alluvial system to which it is adjacent. Extraction rates are monitored. Water level measurements are available from one of the two production bores from which water is extracted, and from three neighbouring observation bores. Measurements spanning the period 2003 to 2017 comprise the history-matching dataset.

The model domain spans a 4.4 km reach of the Degilbo Creek alluvium. All of its boundaries are open. The southern and northern boundaries of the model domain receive groundwater from, and deliver groundwater to, upstream and downstream alluvium. Its lateral boundaries receive water from a regional groundwater system which is of generally lower permeability than that of the alluvial system from which Biggenden draws its water. The alluvium itself receives water from diffuse and creek-bed recharge during wet seasons. Water is lost to the creek, and to near-creek evapotranspiration, during dry seasons.

Recharge processes are simulated using a simple, fast-running, lumped-parameter model named LUMPREM. LUMPREM’s calculations are based on a single soil moisture store in which recharge, macropore recharge, evapotranspiration and runoff are nonlinear functions of currently stored water. Five instances of LUMPREM calculate recharge over five different parts of the domain of the Biggenden groundwater model.

Seasonal fluctuations of the water table at model domain boundaries are also simulated using LUMPREM. LUMPREM provides functionality through which fluctuations in stored soil moisture can be transformed into fluctuations of head or drawdown.

Flow of groundwater within the domain of the Biggenden groundwater model is simulated using MODFLOW 6. LUMPREM-calculated recharge and boundary heads are transferred to the groundwater model using time series functionality provided by MODFLOW 6. The MODFLOW 6 grid is quadtree-refined in the vicinity of creeks and bores.

Simulation of water movement within the domain of the Biggenden groundwater model, and of exchange of water between the alluvial and regional groundwater systems, requires an appropriate level of process and parameterization complexity. Parameterization complexity is also required for representation of potential hydraulic property heterogeneity within the domain of the model. Spatially dense arrays of pilot points are used to represent hydraulic conductivity and specific yield - 628 pilot points for hydraulic conductivity and 148 pilot points for specific yield. A further 169 pilot point parameters are used to represent spatial variability of conductance along model boundaries and along the creeks which transect the alluvial system. A total of 81 parameters are associated with instances of LUMPREM, while another 2 parameters describe anisotropy of hydraulic conductivity and specific yield heterogeneity within the Degilbo Creek alluvial system.

A total of 1028 parameters are therefore ascribed to the Biggenden groundwater model. This is many more than can be estimated uniquely. However deployment of a large number of parameters reduces the possibility of history-matching-induced predictive bias. It also removes obstacles to calculation of predictive uncertainty incurred through failure to represent parameters, and combinations of parameters, that cannot be uniquely estimated. The last point is important, as inverse problem nonuniqueness is often the dominant contributor to groundwater model predictive uncertainty. It follows that accommodation of this nonuniqueness through introduction to the model domain of parameters, and combinations of parameters, that cannot be estimated, in addition to those that can, forestalls under-representation of the uncertainties of decision-critical model predictions.

Assessment of water supply security requires that the Biggenden groundwater model be provided with many different stochastic realizations of future weather; these are comprised of daily time series of precipitation and potential evaporation. One hundred such realizations spanning a 125 year predictive simulation period are employed. It also requires that these predictive model runs be repeated with many different realizations of model parameter fields, all of which are reasonable from a hydrogeological perspective, and all of which allow the model to replicate what is known of the historical behaviour of the system. Recall that this knowledge is encapsulated in water use records, borehole water level measurements, and in expert knowledge of system behaviour. Generation of stochastic sequences of daily rainfall and potential evapotranspiration is not described in this report. However this report does describe in some detail the workflow through which 250 realizations of calibration-constrained, hydrogeologically realistic parameter fields are generated.

First the Biggenden groundwater model was calibrated using the PEST_HP parameter estimation package. In describing this aspect of the workflow, the report distinguishes the notion of “calibration” from that of “history-matching”. The former term has a precise meaning; it is history-matching conducted in pursuit of a model parameter field that approaches that of minimized error variance. Uniqueness of this parameter field is achieved through implementation of an appropriate regularisation strategy. For the Biggenden model, Tikhonov regularisation constraints are enforced. These comprise a suite of strategic penalty functions that discourage the introduction of parameter heterogeneity unless it is geologically plausible. At the same time, introduced heterogeneity is limited to the minimum amount that is required for model outputs to attain a good with the calibration dataset.

Once a parameter field of minimized error variance has been attained in this way, the PESTPP-IES ensemble smoother is employed to seek a further 250 parameter fields in which the potential for hydraulic property heterogeneity is allowed full expression, but in which the necessity for a tight fit with hard and soft components of the calibration dataset is maintained. Initial parameter realisations for the ensemble inversion process are sampled from an approximate posterior parameter probability distribution calculated using the Jacobian matrix attained by PEST_HP. The ensemble inversion process benefits from the "head start" that this workflow provides.

Following completion of the above tasks, the Biggenden groundwater model is deployed to assess water supply security for the township of Biggenden. Security is specified using a model-calculated "delivery ratio". The delivery ratio is the amount of water that the aquifer can deliver expressed as a fraction of the amount of water that is requested. A delivery ratio of less than unity denotes a shortfall in supply. A delivery ratio exceedance probability curve is constructed using outputs from the Biggenden model when run into the stochastic future. Using different realizations of future weather and model parameter fields, upper and lower uncertainty bounds are placed on delivery ratio exceedance probabilities.

This report partitions delivery ratio uncertainty between its two contributors – that arising from uncertainty in future weather on the one hand, and that arising from uncertainty in model parameters on the other hand. (The latter express shortfalls in knowledge of system properties and behaviour). It is demonstrated that uncertainties in Biggenden water supply security are dominated by uncertainties in future weather rather than by uncertainties in system properties and behaviour. This is a rather surprising conclusion, given the dearth of data on which construction of the Biggenden groundwater model rests. However it indicates that data of the right type exists at the right time and at the right places to reduce the uncertainties of predictions that matter. Uncertainty reduction is achieved through matching model outputs to these data. The uncertainty that remains is quantified through the modelling workflow.

Modelling that was undertaken in order to assess water supply security for the township of Biggenden was designed to serve the decision-support imperatives of uncertainty quantification and reduction. Although there is room for improvement in all modelling work, we consider that construction, history-matching and deployment of the Biggenden groundwater model has served these imperatives well.

CONTENTS

1. INTRODUCTION

1.1 The Site

Biggenden has a population of about 850. It is situated 300 km NNW of Brisbane, Queensland on the road between Maryborough (a coastal town with a population of about 16,000) and Gayndah (a regional centre with a population of about 2000). The name is derived from “bigindhan”, meaning “place of stringybark” in local Aboriginal (Kabi) dialect. Biggenden was founded in 1889 as a service centre for the short-lived goldrush towns of Paradise and Shamrock, and for coach passengers travelling west from Maryborough. Nowadays, primary production (mainly beef cattle) is the local area’s most significant industry; grain crops are also grown. Biggenden’s location is shown in Figure 1.1.

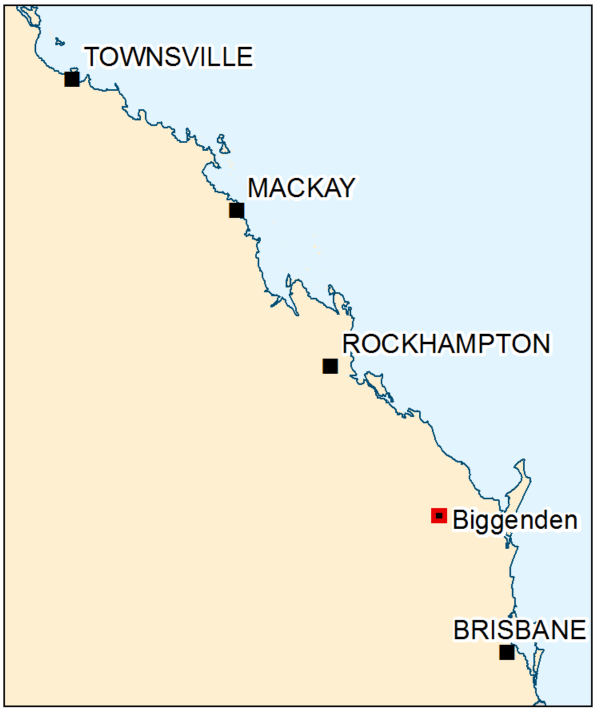

Figure 1.1. Location of Biggenden in relation to some major Queensland towns.

Biggenden’s yearly average rainfall is 882 mm. However this has varied from as low as 334 mm to as high as 1800 mm. Most rain falls in the summer months between October and March.

Degilbo Creek passes about 2 km east of Biggenden. The closest gauging station on this creek is about 25 km downstream. Flow is seasonal. Mungore Creek joins Degilbo Creek just to the southeast of Biggenden.

Since 1965, Biggenden has drawn its water from local alluvium and underlying basalt bordering Degilbo Creek. A small number of stock and domestic bores also draw water from the Degilbo Creek alluvium in the vicinity of Biggenden. Figure 1.2 shows Biggenden township and the two bores from which it draws water, namely RN156052 and RN155332. About 75% of this supply is drawn from the former bore while the remainder is drawn from the latter bore. Three nearby observation wells (RN13600296, RN13600186 and RN13600186) are also shown. A series of five, more distant, observation wells are situated south (and upstream) of the Biggenden bore field along the Maryborough-Biggenden road. The frequency of water level measurements in these observation bores varies from weekly to six-monthly.

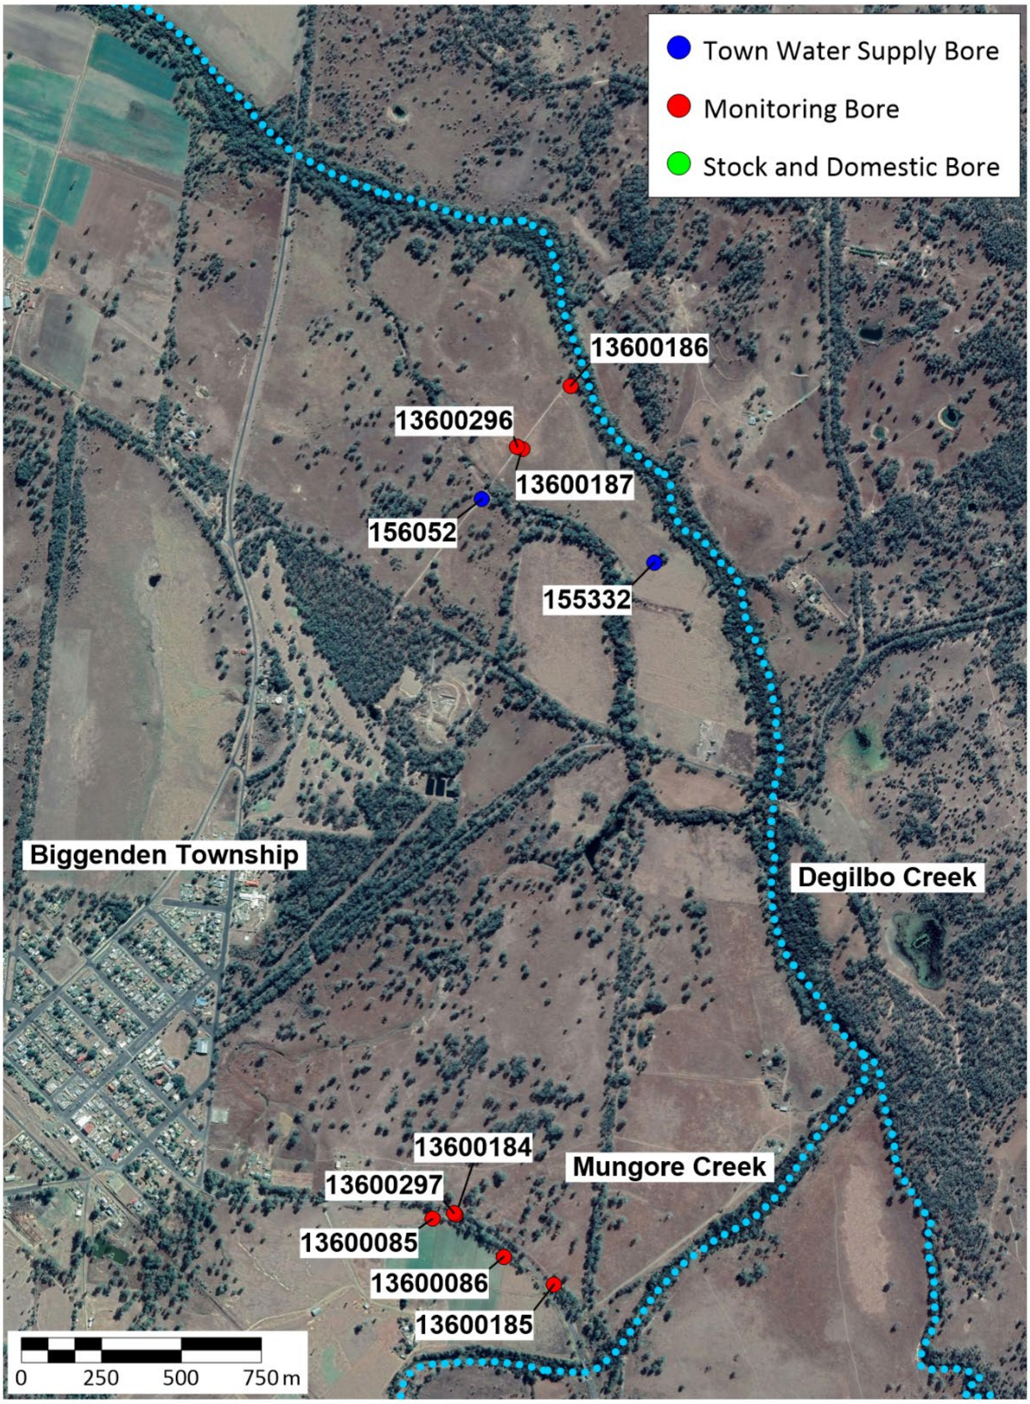

Figure 1.2. Location of pumping bores, monitoring bores and watercourses near Biggenden.

RN156052 and RN155332 were drilled in 1965 and 1980 respectively. Intermittent measurements of groundwater level have been made in the former bore (the main Biggenden supply well). Some of these measurements were made when the pump was running, while other measurements were made when it was switched off. No water level measurements have been made in the other production bore. Water levels measured in RN156052 are graphed in Figure 1.3; observations are linked by dotted lines to improve clarity.

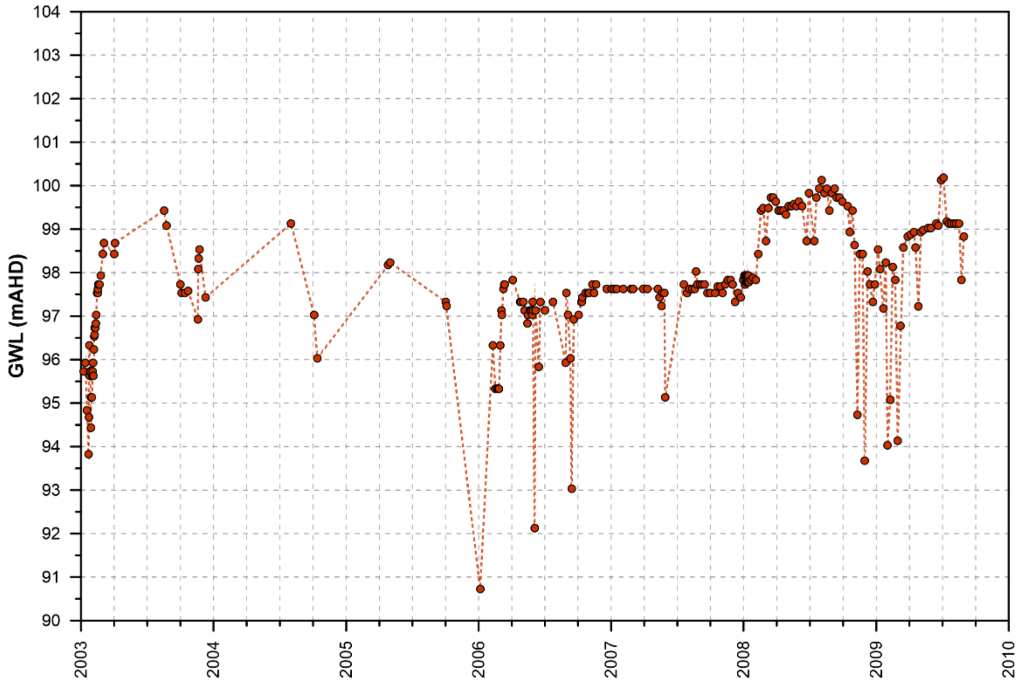

Figure 1.3. Measured water levels in bore RN156052.

Pumping from RN156052 and RN155332 is not metered. However records of water use are available for the Biggenden water treatment facility from September 2008 to the present day.

1.2 The Problem

Queensland Hydrology Unit (QHU) of the Department of Environment and Science (DES) was asked by the Queensland Department of Natural Resources, Minerals and Energy (DNRME) to assess the long-term security of Biggenden's water supply. This was done as part of a larger Regional Water Supply Security Assessment (RWSSA) project conducted by DNRME.

Prior to the Biggenden study, the towns and cities for which water supply security assessments had been undertaken were all reliant on surface water. These assessments were made using existing models of managed river systems. These models are provided with stochastic realisations of daily rainfall and potential evaporation extending up to 1,000 years into the future. Using these realisations of future weather, the models predict volumes of water in the river and creek storages from which these communities draw their water. Water supply failure occurs when available water is insufficient to meet demand. The success (or otherwise) of different management strategies in prolonging water availability can also be tested with this modelling approach. These strategies normally invoke water use restrictions when storage water levels fall below certain thresholds.

1.3 The Challenges

Where water supplies are drawn from man-made storages in rivers and creeks, uncertainties associated with their adequacy arises from uncertainties in future rainfall. The amount of creek and river flow that follows a rain event can be predicted with a high degree of reliability from the historical behaviour of the system, for which long-term records of rainfall and flow are often available. The volumes of river and creek storages are also known. Seepage and evaporative losses from these storages as a function of stored water volume are readily calculable.

In contrast, where water supplies are extracted from a groundwater system, the knowledge base is far smaller. Information inadequacies are exacerbated where measurements of historical groundwater levels have been made in only a few wells, and where these measurements do not extend far back in time. The size of a groundwater storage is only approximately known. Furthermore, its boundaries are pervious; hence unknown amounts of water can flow into and out of a local system from neighbouring systems. Access to stored water by borehole pumps is highly dependent on geological conditions that prevail in the immediate vicinity of boreholes. The extent to which rainfall recharges a groundwater storage is often only vaguely understood; so too is the ability of a groundwater system to gain or lose water from/to a creek or river which transects it.

It follows that uncertainties in predictions of future groundwater availability are likely to be greater than those associated with future surface water availability. Assessment of water supply security must take limited knowledge of system geometry, properties and recharge into account.

1.4 Meeting the Challenges

As stated above, assessment of surface water supply security requires that an appropriate model be run using a suite of sequences of possible future weather. The same applies to assessment of groundwater supply security. However the groundwater model on which this assessment relies should also be run using many stochastic realisations of system hydraulic properties and recharge. Ideally, these realisations should express enough stochastic variability to ensure that they encompass the system's true properties and recharge, hence ensuring that risks to water supply are not under-estimated. At the same time, they should also avoid over-estimation of risk incurred by failure to assimilate all available information. This requires that any realisation that is used to explore the future behaviour of the system must also allow the model to reproduce its historical behaviour when it is run under weather conditions which prevailed in the past. Generation of such a suite of "realistic", calibration-constrained parameter fields is one of the challenges facing water supply security assessment for the township of Biggenden.

The availability of data on which to base construction of a groundwater model for the Degilbo Ck alluvium from which Biggenden draws its water supply is far less than at most sites at which groundwater models are constructed. This does not preclude construction of a model, as this is required for water supply security assessment. However it suggests that the uncertainties associated with model predictions will be high. This makes the need for uncertainty reduction through assimilation of all pertinent information on system behaviour even more pressing. This information is comprised of discrete measurements of historical water levels in a small number of bores. It also includes qualitative knowledge of how the groundwater system behaves. These two sources of information are referred to as "hard data" and "soft data" in this report.

Water levels in one pumped bore, and in a small number of nearby observation bores comprise available hard data. A pumped bore fails when the water level in that bore falls below the pump intake. In anticipation of this condition, water restrictions may be imposed when borehole water levels cross thresholds that are somewhat above this. In a shallow groundwater system such as that from which Biggenden draws its water, centimetres of head in the pumped bore can make the difference between failure and security. While there are no expectations that model predictive uncertainties can be this low, construction and deployment of a water supply security model comes with the responsibility that uncertainties be reduced as much as available data allows. This requires that all stochastic realisations of system hydraulic properties that are used to explore predictive uncertainty allow the model to reproduce historical borehole water levels as well as possible.

“Soft data” is more abundant, but is less accurate. Nevertheless it is important that these data be respected, for they have the potential to constrain model parameters that represent system hydraulic properties to values that deter model predictions from attaining unrealistic extremes. By design, the model will be asked to make predictions of groundwater behaviour under conditions that have not been experienced in the past. It is well known that attainment of a good fit with historical hard data (as is sought for the present model) may compromise its performance under extreme conditions unless “reality checks” are enforced on parameters that are estimated through this fitting procedure. Assimilation of soft data therefore assumes extra importance.

For the Biggenden model, items of “soft data” include (but are not limited to) the following aspects of system behaviour.

- Groundwater flows preferentially within, and parallel to, the alluvial system associated with Degilbo and Mungore Creeks.
- Water that enters the local alluvial system from upstream (i.e. from the south) exits the local system downstream (i.e. to the north) through the alluvial system itself. Water is unlikely to flow upstream through either of these boundaries during either wet or dry seasons.
- Groundwater levels in non-alluvial areas which border the Degilbo and Mungore Creek alluvial systems rise and fall with season. Some water enters these systems from upland areas which border them. Little, if any, water is likely to flow in the opposite direction, especially during dry seasons.
- Except during very wet events, the phreatic surface intersects the land surface only at Degilbo and Mungore Creeks.
- During any wet season, the depth to the phreatic surface is unlikely to be greater than 12m in any part of the study area. This is the greatest wet season depth recorded in any local observation bore.
- Given the nature of alluvial deposition, hydraulic property heterogeneities are more likely to be disposed parallel to the creek system than orthogonal to it.
- Studies of similar systems, for example Wohling et al. (2011) and King et al. (2017), suggest that long term recharge to the Degilbo and Mungore Creek alluvial systems is likely to be between 1% and 5% of average rainfall. However it is also likely to be highly seasonal.

The status of Degilbo and Mungore creeks as gaining or losing streams in the vicinity of Biggenden township is unknown. It is possible that the alluvial groundwater system receives recharge from these streams during high flow events. It is most unlikely, however, that it receives recharge from these creeks during dry periods when they do not flow. Periods of most predictive interest are confined to these dry periods.

1.5 Desired Modelling Outcomes

For the Biggenden model, water supply security is assessed by running it for 125 years into the future using different stochastic realisations of rainfall and potential evaporation. For any one of these realisations, the model must calculate water levels in the two production wells. Meanwhile the model must automatically reduce pumping from these wells as predicted water levels decline. In particular, the model must maintain borehole water levels at or above pump intakes. At the same time, it must be capable of simulating pre-emptive reductions in water extraction at higher water level thresholds in accordance with water restrictions imposed in times of drought, this enabling the model to test their efficacy.

Model predictions of most interest are therefore future borehole extraction rates, particularly during times when these are reduced automatically by the model. Notwithstanding assimilation of hard and soft data in the manner described above, it is anticipated that these predictions will be accompanied by uncertainty. They must therefore be made using multiple realisations of history-match-constrained parameter fields, together with multiple realisations of future weather. Both of these sources of uncertainty can thereby be taken into account when assessing risks to water supply reliability.

Generation of stochastic time series of daily rainfall and evaporation is beyond the scope of this report. It is discussed by Vitkovsky (2008). The present report addresses the design and construction of the Biggenden groundwater model, and its population by random parameter fields that allow the model to replicate hard and soft data on historical system behaviour.

1.6 Remainder of this Report

As discussed in the preface, the document that you are reading is not intended to be a modelling report, despite the fact that it describes a groundwater model. Instead, it attempts to describe a particular approach to decision-support modelling using the Biggenden groundwater model as an example.

Construction and parameterisation of the Biggenden model are briefly described in Chapter 2. In order to maintain reader interest, many details are omitted. Nevertheless, we hope that the description is detailed enough for a reader to understand how the modelling goals described above were pursued.

Chapter 3 comprises a further departure from conventional modelling reporting. In this chapter we justify some of the model design features that are described in Chapter 2. We feel that this discussion is better presented after these features are described, than before.

Chapter 4 describes those aspects of the historical behaviour of the Degilbo Creek groundwater system that must be reproduced by the model. These constrain the values of parameters with which the model must be endowed when calculating its future behaviour.

Chapter 5 discusses how model calibration was achieved. It also discusses the notion of “calibration” itself, and how this should be distinguished from “history-matching”.

Chapter 6 then describes how not just one, but many, history-match-constrained stochastic parameter fields were obtained. The use of these parameter fields in examining Biggenden water supply security is addressed in Chapter 7. Chapter 8 concludes this report.

2. SIMULATION

2.1 General

Simulation of groundwater movement in the Degilbo Creek alluvial system must take place over two time periods. The first of these periods is historical, while the second is predictive. During the first of these periods, model runs are controlled by data assimilation software, namely PEST_HP (Doherty, 2020) and PESTPP-IES (PEST++ Development Team, 2020). This is done in order to derive a suite of stochastic parameter fields which allow the model to replicate historical system behaviour. 250 of these parameter fields are generated. Ideally, these parameter fields sample the posterior (i.e. post-history-matching) probability distribution of model parameters.

The second simulation time period is predictive. Over this period the stochastic parameter fields derived by PEST_HP and PESTPP-IES are employed to make predictions of groundwater behaviour under conditions defined by stochastic sequences of future rainfall and potential evapotranspiration. These predictions are thereby imbued with two levels of stochasticity, namely (a) that associated with uncertainties in future weather patterns and (b) that associated with uncertainties in hydraulic properties of the Degilbo Creek alluvial system.

2.2 Simulators

2.2.1 MODFLOW 6

At the time of writing, MODFLOW 6 (Langevin et al, 2019) is the latest generation of the USGS family of MODFLOW simulators. Of primary interest to the present project is its ability to employ an unstructured grid. This supports grid refinement in the vicinity of Degilbo and Mungore Creeks, and in the vicinity of the two extraction wells.

Another convenient feature of MODFLOW 6 is its ability to read time series from external files. This facilitates its use in conjunction which programs such as LUMPREM (see below) which can compute transient stresses and boundary conditions used by MODFLOW 6. The time base pertaining to these external time series can be independent of that used by MODFLOW 6 itself to calculate system states. They can therefore be set by the programs which generate them.

Yet another convenient feature of MODFLOW 6 is its ability to employ multiple incidences of the same stress package. The Biggenden model employs the MODFLOW 6 EVT package to simulate groundwater losses to evaporation in the vicinity of streams. This same package is also used to automate reduction of well extraction as borehole groundwater levels fall. Reduction of extraction occurs for two reasons:

1. To prevent the water level in a production bore from falling below the pump intake level; and
2. To implement water restrictions imposed during drought conditions.

The EVT package allows implementation of the latter option using a stepped sequence of increasing restrictions with reducing borehole water levels. For brevity, use of the EVT package in implementing these restrictions is not discussed further in this document.

2.2.2 The LUMPREM Recharge Model

The LUMPREM model, together with it manual (Doherty, 2020), can be downloaded from the PEST web pages at http://www.pesthomepage.org. “LUMPREM” stands for “lumped parameter recharge model”. LUMPREM simulates processes that are operative in the unsaturated zone overlying the water table using a single soil moisture store. Water can leave the store as either evaporation or groundwater recharge at rates that are calculated as nonlinear, user-parameterizable functions of stored moisture volume. If the soil moisture store overtops, water is lost as macropore recharge to groundwater and/or as runoff. Macropore recharge suffers a user-specified delay before being made available for groundwater recharge. A schematic of this process is provided in Figure 2.1.

Through adjustment of parameters which govern:

- the size of the soil moisture store;
- the nonlinear relationship between recharge and stored soil moisture;
- the nonlinear relationship between evapotranspiration and stored soil moisture; and
- the delay suffered by macropore recharge;

the quantity and timing of recharge can resemble that calculated by much more complicated models (Watson et al, 2013; Meeks et al, 2017).

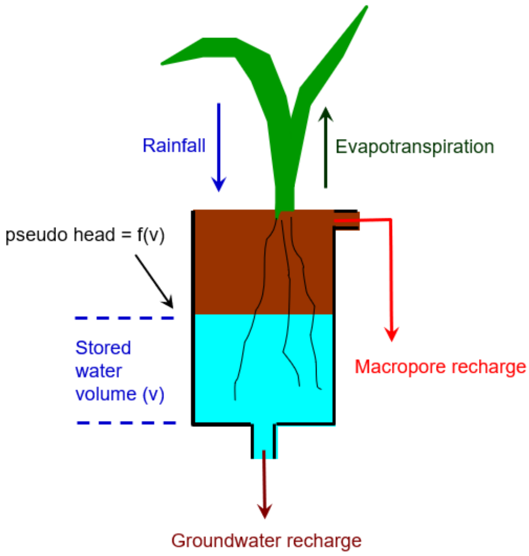

Figure 2.1. Schematic of the LUMPREM recharge model.

Let v' denote the volume of moisture stored in the soil moisture store relative to the total volume of the store. LUMPREM calculates daily evaporation E as function of v' using the equation:

$$E = f\ E_p \frac{1-e^{-\gamma v'}}{1-2e^{-\gamma}+e^{-\gamma v'}} \tag{2.1}$$

where E_p is potential evaporation. f can be considered as a "crop factor" and γ as a "fitting parameter".

LUMPREM calculates daily recharge R from v' using the relationship:

$$R = K_s [v']^l \left[1-\left(1-[v']^{1/m}\right)^m\right]^2 \tag{2.2}$$

where K_s represents soil saturated hydraulic conductivity, l is a pore-connectivity parameter and m is a fitting parameter. See the LUMPREM manual for further details.

LUMPREM operates on a daily time step. It reads daily values of rainfall and potential evaporation from its input files. It records daily values of normal and macropore recharge, runoff, evaporation, and residual evaporation (i.e. potential evaporation minus actual evapotranspiration) on its output file, together with other quantities. A utility program supplied with LUMPREM named LR2SERIES (i.e. "LUMPREM to time series") writes MODFLOW 6 time series files based on the contents of a LUMPREM output file. Thus LUMPREM outputs can serve as inputs for various MODFLOW 6 packages. In particular, recharge calculated by LUMPREM can provide inputs for the MODFLOW 6 RCH (i.e. recharge) package; residual evapotranspiration calculated by LUMPREM can provide the "maximum ET" input for the MODFLOW 6 EVT (i.e. evapotranspiration) package.

LUMPREM also records the daily volume of stored water on its output file. This quantity can rise steeply after rainfall and decay slowly thereafter. As such, its behaviour is not unlike that of groundwater levels in a shallow aquifer. LUMPREM allows a user to calculate a pseudo groundwater head time series from these daily volumes using a nonlinear function whose parameters can be set by the user, and/or adjusted through history matching. This pseudo head time series is also recorded on its output file. Like other time series that are recorded in this file, it can serve as an input time series for a MODFLOW 6 stress package. LUMPREM employs the following equation to transform stored moisture volume to head:

$$h = a + f_1 v' + f_2 (v')^p \qquad (2.3)$$

In this equation v' denotes the volume of stored water relative to the total capacity of the store, h denotes the head that is calculated from it, and a, f_1, f_2 and p are parameters.

As is discussed below, nine instances of the LUMPREM model are used in conjunction with the MODFLOW 6 groundwater simulator in the Biggenden model. Five of these models compute recharge in different parts of the model domain. Four are used for calculation of boundary heads. The following parameters of all of these models are adjusted through history-matching:

- K_s
- m
- f
- γ
- a, f_1 ,f_2 and p (for those instances of LUMPREM which compute boundary heads).

2.2.3 Stress Periods

LUMPREM calculates its outputs on a daily basis. In contrast, time steps and stress periods employed by MODFLOW 6 are user-selectable. For simulation of groundwater flow in the Degilbo Creek alluvium, a 14 day stress period is employed by MODFLOW 6 in both its history-matching and predictive simulations. Pertinent LUMPREM daily outputs are averaged over the time spanning each MODFLOW 6 stress period before being provided to MODFLOW 6.

2.3 The Biggenden Groundwater Model

2.3.1 The Model Domain

Figure 2.2 shows the grid of the Biggenden groundwater model. The domain of the model spans the Degilbo and Mungore Creek alluvial systems. It extends approximately 1.7 km north and 2.7 km south of the Biggenden production wells. The southwestern extremity of the model domain coincides with a line of observation bores beside the Maryborough-Biggenden Road.

Over most of the model domain, cell dimensions are 40m × 40m; they are quadtree-refined to 5m × 5m in the vicinity of pumping wells, and to 10m × 10m near creeks. The model employs just a single, unconfined layer. Its domain covers an area of 6.35 km^2. The top elevation of each cell is set to the topographic elevation. MODFLOW 6 “drains” (implemented using its DRN package), incised 3m into the land surface, are used to represent watercourses. These are thus represented as groundwater sinks; this is their role in the dry season where model predictions are of most interest.

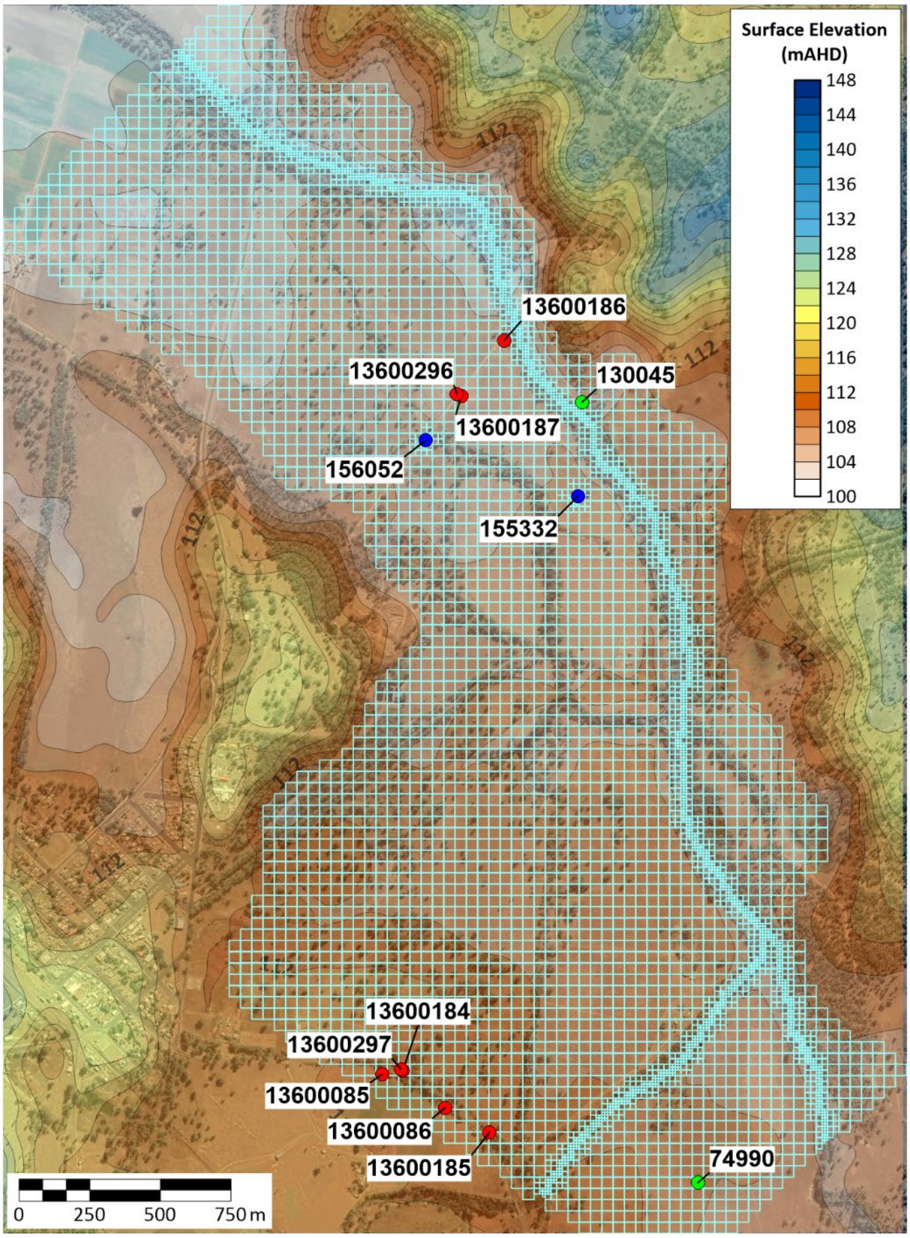

Figure 2.2. The domain of the Biggenden groundwater model; topography is also shown.

2.3.2 Northern and Southern Alluvial Boundaries

Connections to the extended Degilbo and Mungore Creek alluvial systems at the southern (downstream) and northern (upstream) boundaries of the model domain are simulated using fixed head cells; this is implemented using the MODFLOW 6 CHD package. These alluvial heads vary with season. These seasonal variations are represented using the pseudo-head functionality of the LUMPREM recharge model. A single LUMPREM-based model is used to simulate variability of all of these heads. This model is actually comprised of two LUMPREM models in series, such that recharge from one provides “rainfall” input to the other. (Piggybacked LUMPREM models realize a better fit with borehole-observed water levels at model boundaries than a single LUMPREM model.) This dual LUMPREM model was calibrated against water levels measured in bore RN13600185 over the period June 1993 to August 2017. This observation well is situated on the south-western alluvial boundary of the Biggenden groundwater model.

Figure 2.3 shows the fit between water levels measured in this bore, and those calculated by the calibrated, dual LUMPREM model. Figure 2.3a displays LUMPREM-calculated water levels only at borehole measurement times; Figure 2.3b shows LUMPREM-calculated water levels at regular, fortnightly intervals.

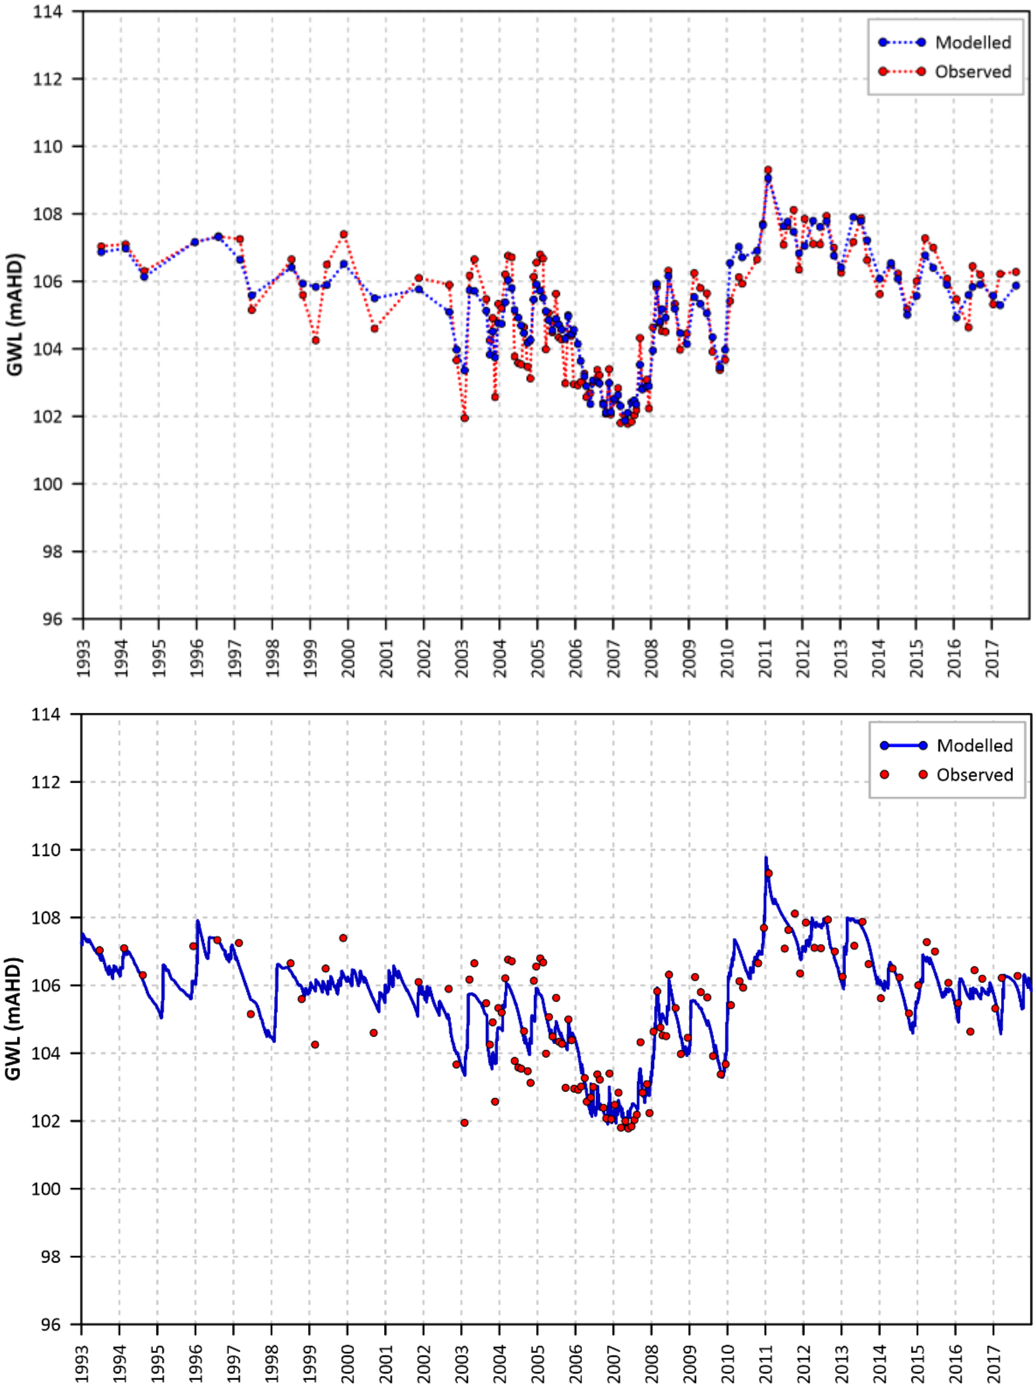

Figure 2.3 Measured and model-calculated water levels in bore RN13600185. The latter are displayed for LUMPREM (a) at times of water level observations (top) and (b) at regular, fortnightly intervals (bottom).

Depths to the phreatic surface at the location of RN13600185 are computed by subtracting LUMPREM-calculated heads from the surface elevation. Heads at all CHD cells along both the northern and southern boundaries of the Biggenden groundwater model are then calculated by subtracting this depth from surface elevations in each cell along these model boundaries.

Table 2.1 lists LUMPREM parameters whose values were estimated by history matching against heads in RN13600185. Tikhonov regularisation and singular value decomposition were used as regularisation devices to ensure numerical stability and achieve parameter uniqueness. Note that, with the exception of hydraulic conductivity anisotropy (see below) these are the only parameters of the Biggenden model for which calibration-constrained stochastic realisations were not generated; the single set of parameter values obtained in the manner described above were therefore used when running the Biggenden model into the future. It was felt that alluvial model boundaries are sufficiently removed from Biggenden pumping wells for their uncertainties to matter little to model predictions of most interest.

Parameter type	**Number of parameters**
K_s	2
m	2
f	2
γ	2
a	1
f_1	1
f_2	1
p	1

Table 2.1 Parameters of the dual LUMPREM model that were adjusted in order to fit heads measured in RN13600185.

2.3.3 Eastern, Western and Southern Boundaries

Along its eastern and western boundaries, and along part of its southern boundary, the Biggenden groundwater model requires connections to a shallow, regional, non-alluvial groundwater system in which hydraulic conductivities are expected to be low. These connections are simulated using MODFLOW 6 general head boundaries (i.e. GHBs). A GHB condition is established in every one of these boundary cells.

Like alluvial boundaries, heads along the eastern, western and southern boundaries of the Biggenden groundwater model are expected to fluctuate with season. These fluctuations are also simulated using the pseudo head functionality of the LUMPREM recharge model. Three LUMPREM models are deployed, one for each boundary. Each of these models was used to calculate water depth along the entire boundary. The head assigned to the GHB ascribed to each cell along these boundaries is calculated by subtracting this depth from the local surface elevation.

The parameters of the LUMPREM models associated with the eastern, western and southern boundaries were estimated through history-matching of the entire Biggenden model, both during calibration of that model and in subsequent history-match-constrained stochastic parameter field generation. Meanwhile GHB conductances are parameterized using pilot points; see below. These are also history-match-constrained.

2.3.4 Recharge and EVT

Recharge throughout the domain of the Biggenden groundwater model is computed using five instances of the LUMPREM model. One of these instances operates in each of five different zones. These zones are depicted in Figure 2.4; they are comprised of four riparian zones and a pervasive open pasture zone. The MODFLOW 6 EVT package also operates within each of the four riparian zones; in each of these zones a time series of maximum EVT rate is provided by the corresponding LUMPREM recharge model. LUMPREM calculates potential evaporation available to the groundwater system by subtracting evapotranspiration calculated using equation 2.1 from potential evaporation for each day.

2.3.5 Summary of Boundary Conditions

Incidences of LUMPREM used to define boundary condition heads on the one hand, and recharge rates on the other hand, are summarized pictorially in Figures 2.4a and 2.4b.

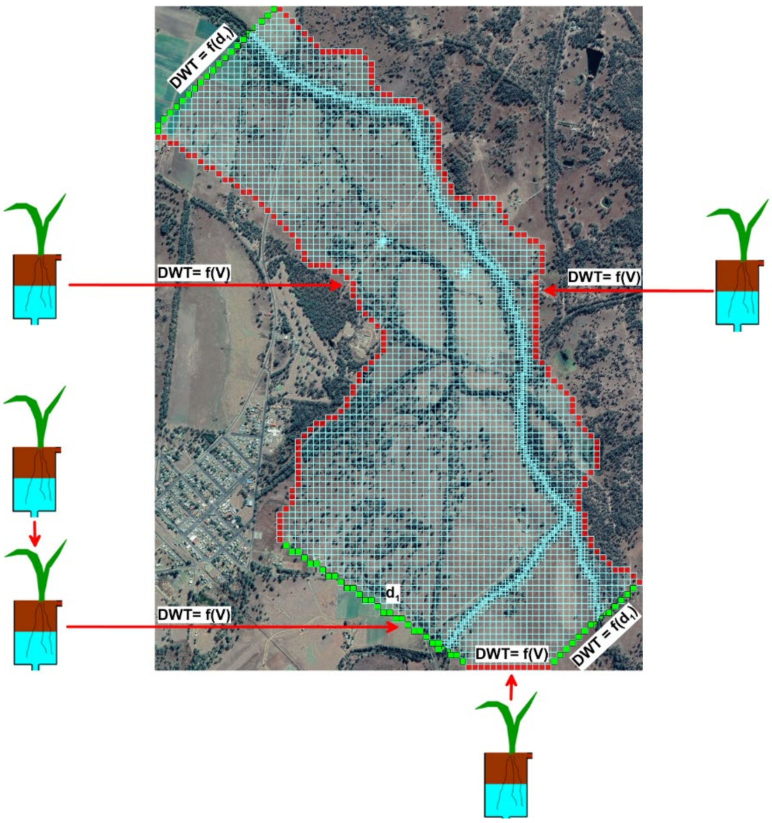

Figure 2.4a. Time-varying heads ascribed to lateral boundary conditions of the Biggenden groundwater model are calculated using four dedicated LUMPREM models for depth-to-water (DWT), one of which is actually a dual LUMPREM model. The boundary conditions on the eastern, western and southern sides of the model are GHBs. Conditions on the northern, south-eastern and south-western alluvial boundaries are represented by CHDs.

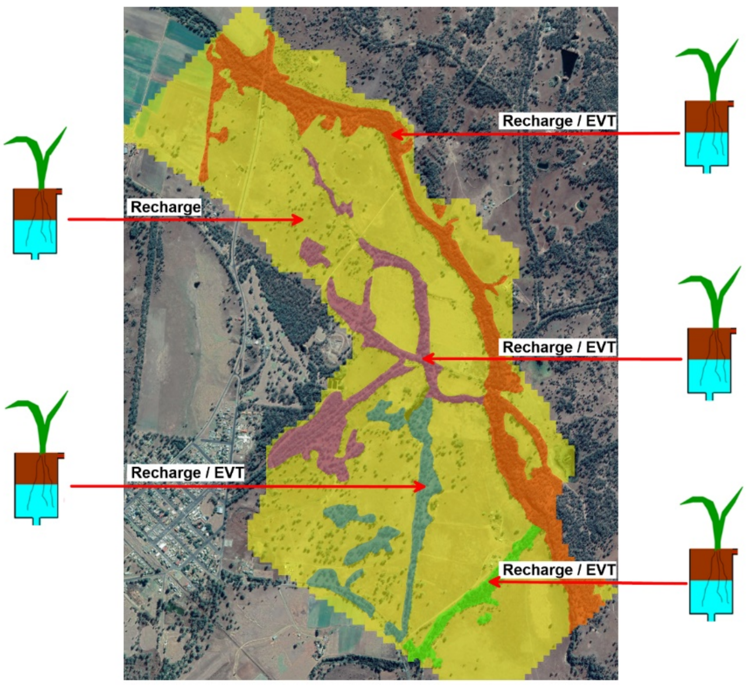

Figure 2.4b Recharge is supplied to the MODFLOW 6 RCH package of the Biggenden groundwater model using five incidences of LUMPREM. Four of these LUMPREM incidences also provide maximum EVT rate to the MODFLOW 6 EVT package.

2.4 Spatial Parameterisation

2.4.1 Hydraulic Conductivity and Specific Yield

MODFLOW 6 requires that values be supplied for hydraulic conductivity (Kh) and specific yield (Sy) throughout its domain. Pilot points are used as a parameterisation device for both of these. Those for Kh are illustrated in Figure 2.5a while those for Sy are illustrated in Figure 2.5b.

As is apparent from Figure 2.5b, Sy pilot points are uniformly distributed throughout the model domain; each is about 200m from its nearest neighbour. Kh pilot points are 100m apart, except in the vicinity of extraction wells where they are 75m apart.

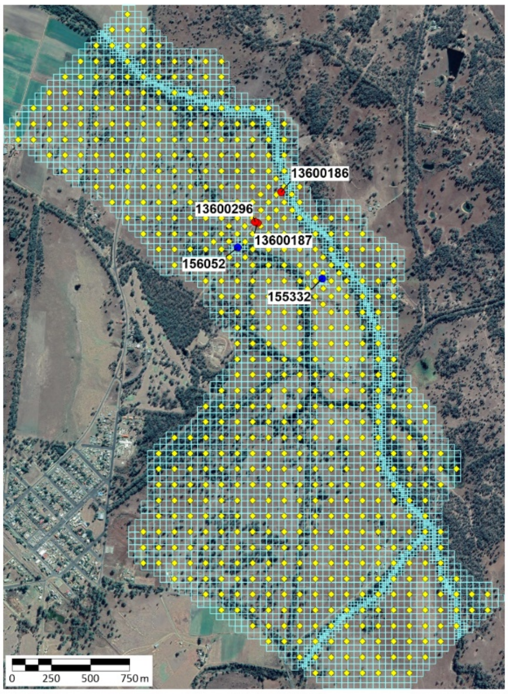

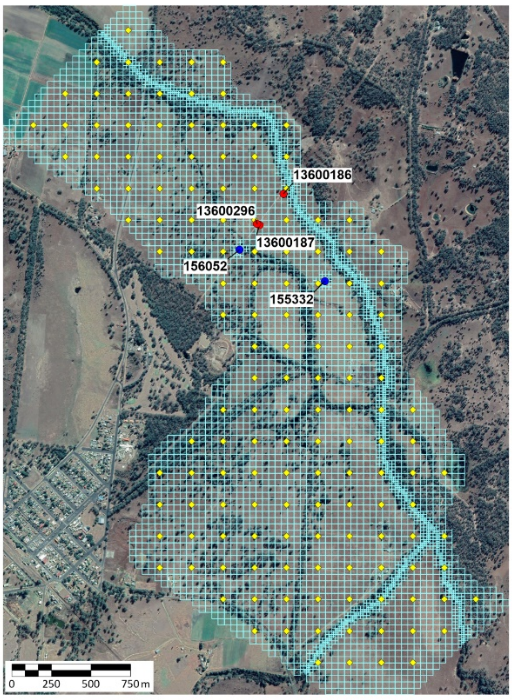

Figure 2.5a. (left) Pilot points used for parameterisation of hydraulic conductivity (Kh).

Figure 2.5b. (right) Pilot points used for parameterisation of specific yield (Sy).

Spatial interpolation from pilot points to the model grid employs the *alluv_boundary_interp()* function of PLPROC. PLPROC (Doherty, 2020) is a "parameter list processor" supplied with the PEST suite. It was written specifically to facilitate flexible pilot points parameterisation of two- and three-dimensional model domains. The *alluv_boundary_interp()* function supports anisotropic interpolation from pilot points to a model grid. The principle axis of anisotropy used in this interpolation varies throughout the model domain in accordance with the orientation of the closest boundary of an alluvial system. A single anisotropy value is employed for interpolation of each of Kh and Sy. These anisotropies were estimated during the model calibration process.

2.4.2 GHB and DRN Conductances

As was discussed above, the Biggenden model employs the MODFLOW 6 GHB package along its eastern and western margins, and along part of its southern margins, to simulate connections with the regional, non-alluvial, groundwater system. At the same time, the MODFLOW 6 DRN (i.e. drain) package allows water to flow from the groundwater system into

Degilbo and Mungore Creeks. Both of these packages require a conductance term that controls the rate of water flow into or out of the groundwater system.

In order for the conductance of these polylinear boundary features to vary along their lengths they, too, are parameterized using pilot points. The locations of these pilot points are shown in Figure 2.6. Interpolation between them and the boundaries that they inform is linear. This is accomplished using a specially written program. It can also be accomplished using so-called "SEGLIST" functionality provided by PLPROC.

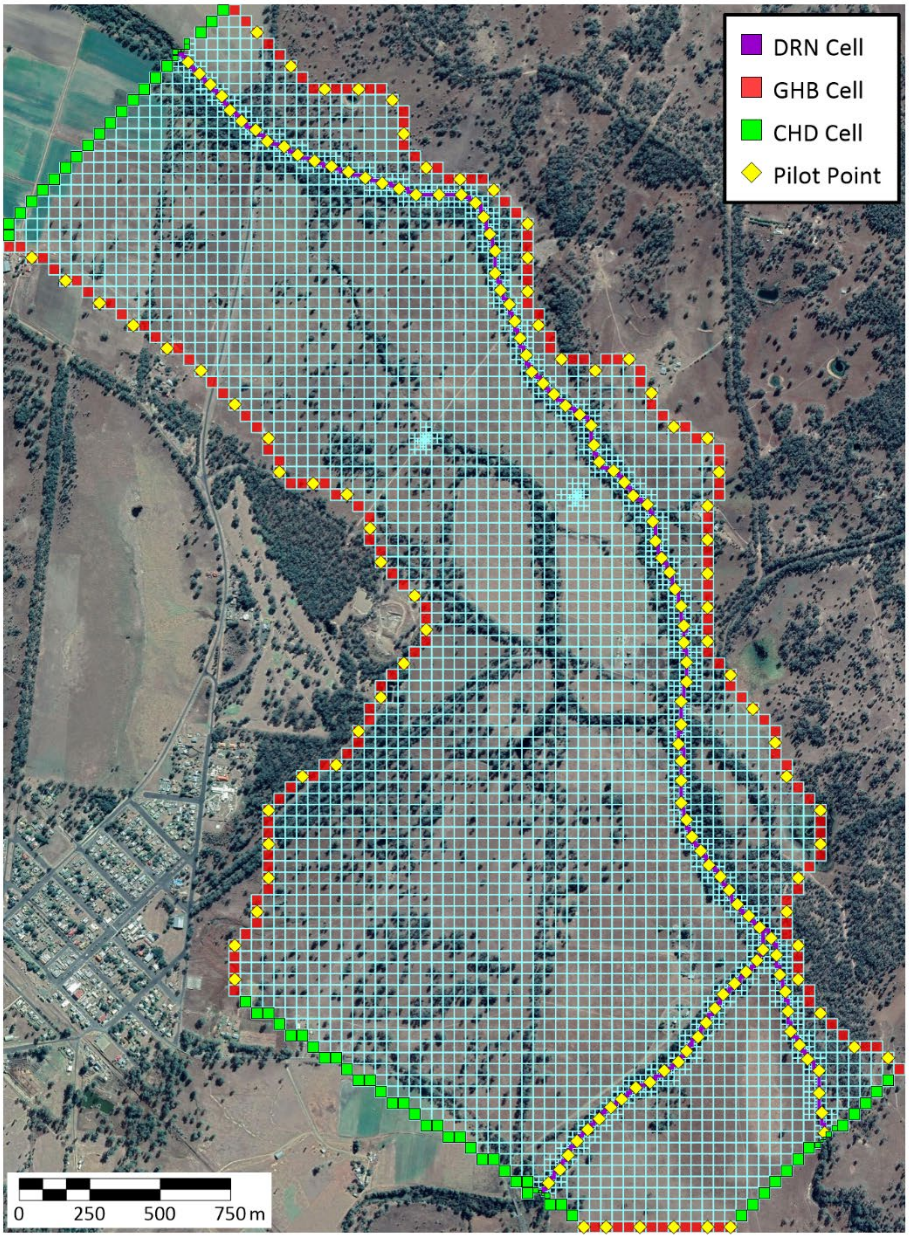

Figure 2.6. Pilot points used for parameterisation of GHB and DRN boundary conductances.

2.5 Parameterisation Overview

2.5.1 Number of Parameters

It is apparent from the above discussion that the Biggenden groundwater model is simple in some respects and complex in others. It is also apparent that it is endowed with many more parameters than can be estimated uniquely. To refresh the reader's memory, parameterisation of the Biggenden model is summarized in Table 2.2.

Parameter type	**Role**	**Number of parameters**
LUMPREM	Calculation of recharge and residual ET in 5 surficial zones	45
LUMPREM	Calculation of head for CHD alluvial boundaries	12
LUMPREM	Calculation of head for GHB lateral boundaries	24
Pilot points	Hydraulic conductivity (Kh)	628
Anisotropy	Interpolation from Kh pilot points to model grid	1
Pilot points	Specific yield (Sy)	148
Anisotropy	Interpolation from Sy pilot points to model grid	1
Pilot points	Creek DRN conductance	99
Pilot points	GHB conductance at lateral boundaries	70
Total		1028

Table 2.2. Parameters used by the Biggenden groundwater model.

As has already been stated, the 12 LUMPREM-related parameters associated with the CHD alluvial boundaries were estimated only once. This was done through a calibration exercise focussed only on water levels in RN13600185. The two parameters that characterize interpolation anisotropy from pilot points to the model grid were also estimated only once, this during calibration of the entire Biggenden model. All other parameters were estimated many times – once during the calibration process, and then repeatedly during calculation of history-match-constrained random parameter fields.

3. SOME NOTES ON MODELLING PHILOSOPHY

3.1 General

At this stage of our report we take the rather unusual step of discussing the rationale for building and parameterizing the Biggenden model in the way that we did. We do this now while its construction and parameterisation details are fresh in the reader's mind.

Aspects of model construction and parameterisation that may appear unusual to some readers include the following:

- the use of so many parameters in a context where there is little data to inform them; and
- implementation of a physically-based model in conjunction with somewhat abstract boundary conditions, the use of which requires parameters whose physical meaning is not clear.

We now address these, and other, potential concerns.

3.2 Parameters

All groundwater systems are complex. All are heterogeneous. Hydraulic properties differ from location to location in unknown ways. Nevertheless, at any particular study site, expert knowledge can generally place limits on the range of prevailing hydraulic property values. Expert knowledge may also be capable of characterizing the nature and disposition of hydraulic property heterogeneity. For example, it may suggest that heterogeneity is more continuous in some directions than it is in others; it may be able to associate heterogeneity correlation lengths with different directions.

The use of many parameters recognizes the heterogeneity of subsurface properties. It is important to recognize, however, that use of a large number of parameters does not prevent attainment of a unique, minimum-error-variance solution to the inverse problem of model calibration. ("Calibration" is precisely defined in the next chapter.) Nor does it propel the history-matching process into over-fitting of field measurements. However, as is described by Doherty (2015), use of a large number of parameters does achieve a number of important outcomes:

1. It prevents inadequacies in representation of potential heterogeneity from compromising the fit attained between field measurements and their model-calculated counterparts. For reasons described in Chapter 1 of this document, attainment of a good fit between model-calculated and borehole-measured heads is of fundamental importance to the decision-support role which the Biggenden model must play.
2. It allows maximum flexibility of parameter response to information contained within a measurement dataset. This flexibility limits the chances of calibration-induced predictive bias.
3. If parameters are endowed with a stochastic characterisation that is reflective of prevailing geology, the use of a large number of parameters reduces the chances that the uncertainties ascribed to decision-critical model predictions are underestimated.

3.3 Boundary Conditions

The boundaries of the Biggenden model domain are not the natural boundaries of the local groundwater system. Hence they cannot be assigned a status of static fixed head nor no-flow. Groundwater levels at these boundaries vary with season. The most important predictions required of the Biggenden model are those associated with droughts. During times of drought, boundary water levels and flows are low, and should be represented as such.

It is impossible to represent groundwater behaviour at these boundaries exactly. Nor is it possible to exactly represent the connection of the local groundwater system with the wider groundwater system through these boundaries. Nevertheless, it is important to represent the fact that heads at these boundaries fluctuate, and that inflow and outflow of water to/from the domain of the Biggenden groundwater model is sensitive to these heads. It is also important that representation of these fluctuations by the model be "realistic", stochastic, and constrained by any hard and/or soft data that has the potential to inform them.

The pseudo-head functionality of the LUMPREM model allows heads ascribed to these boundaries to fluctuate with season. Inclusion of LUMPREM model parameters in the model history-matching process constrains boundary head variability to realistic levels. The use of history-match-adjustable pilot points to characterize boundary conductances limits the extent to which the alluvial groundwater system can draw water from the regional groundwater system.

3.4 Safeguarding Parameter Integrity

As has already been discussed, the Biggenden model is designed to make predictions under conditions that will be somewhat different from those which prevailed during history-matching. Biggenden's water supply has not failed to date. However, when provided with stochastic realisations of future rainfall and potential evaporation, it is likely that there will be times in the stochastic future when conditions will be drier than those which prevailed in the past, and more extreme than those which were experienced over the relatively short calibration period. Model predictions will be uncertain over these periods. In ways that have already been explained, the Biggenden model has been constructed in a way that allows it to quantify these uncertainties. Nevertheless, parameter safeguards must be put in place that circumvent erratic model behaviour under the extreme drought conditions that it has been built to explore.

A model's ability to faithfully replicate the past provides some assurance that its predictions of future system behaviour are reliable when conditions in the future are not dissimilar from those which prevailed in the past. However, it is well known that the attainment of a good fit with a calibration dataset can sometimes achieve exactly the opposite of this under different predictive conditions, particularly those which pertain to climatic extremes. Safeguards were put in place to prevent this when history-matching the Biggenden model. These safeguards are as follows.

- For parameters such as Kh and Sy that have hydrogeological interpretations, regularisation (used during calibration) and stochastic field generation (used during uncertainty analysis) maintained hydraulic property reasonableness (i.e. values for these properties that are in accordance with expectations based on geology).
- For parameters which are not so physically-based (for example, LUMPREM parameters that determine the behaviour of time-varying heads at model boundaries), regularisation and stochastic field generation ensured that these parameters are not endowed with extreme values, and that variability of these

parameters between individual LUMPREM model instances is mitigated. (High parameter variability is a common sign of over-fitting.)

- As has been discussed, “soft data” forms an important component of the dataset used for history-matching of the Biggenden model. The parameter field which is deemed to “calibrate” the model (see the next chapter), and stochastic parameter fields which allow the model to reproduce heads in observation wells, therefore also encouraged the model to behave reasonably (from an expert knowledge point of view) over the history-matching period. Ideally, reasonable behaviour will then continue into the future.

4. HISTORY-MATCHING

4.1 Strategy

Exploration of water supply security requires that the Biggenden model be populated with a suite of random parameter fields, and then run into the future using a suite of random realisations of rainfall and potential evaporation. All of these random parameter fields must be reasonable from an expert knowledge point of view. At the same time, they must promulgate a good fit between model outputs and the "hard data" component of the history-matching dataset, whilst simultaneously respecting its "soft data" component.

Generation of random parameter fields was accomplished in two steps. First the model was "calibrated" (see below). The parameter field that emerged from this process, together with the outcomes of ancillary linear uncertainty analysis, were then used to implement the process of generating random, history-match-constrained parameter fields. The outcome of the entire process was a set of 250 random parameter fields that all satisfy hard and soft data constraints.

4.2 The History-Matching Period

History matching, and concomitant parameter adjustment, was implemented over the period 1st January 2003 to 31st December 2017. The transient simulation which spans this period is preceded by a steady state simulation. The same hydraulic conductivity field is employed for both of these simulations (specific yield is only required for transient simulation).

A steady state simulation does not require direct use of LUMPREM model instances for provision of recharge, EVT rates and boundary heads. Instead, for any set of parameters employed by the transient component of the model, LUMPREM-derived inputs were averaged over the transient period for use in the partnered steady-state simulation.

4.3 The Measurement Objective Function

4.3.1 General

The term "measurement objective function" refers to the sum of squared weighted differences between field measurements and model outputs which correspond to them. The lower is the measurement objective function, the better is model-to-measurement fit.

In contrast, the term "regularisation objective function" refers to discrepancies between parameters and/or relatively simple functions of parameters, and the preferred values (from an expert knowledge point of view) of these parameters and/or functions. When a model undergoes calibration, parameter uniqueness is achieved through minimizing the regularisation objective function subject to the measurement objective function achieving a user-specified value, this value being referred to as the "target measurement objective function". If the measurement objective function cannot be reduced to this target, then it is lowered as far as it can be lowered.

As has already been discussed, the dataset used for history-matching of the Biggenden groundwater model can be partitioned into "hard" and "soft" components. The first is comprised of quantities that were actually measured. The second is comprised of quantities that are considered to be desirable features of simulated groundwater behaviour from an expert knowledge point of view. Attainment of a good fit with hard data is essential to useful deployment of the Biggenden model. Predictions required of the model pertain to the same production well as that in which many historical water level measurements were made. Hence it is likely that attainment of a good fit with historical hard data will do much to reduce the uncertainties of these predictions. (To use PEST jargon, predictions of management interest have high sensitivities to combinations of parameters that occupy the calibration solution space.)

Attaining a good fit with “soft data” serves another purpose. Although conditions in the future will not be radically different from those that prevailed in the past, they will nevertheless be somewhat different. By informing the history-matching process that estimated parameters must support reproduction of not only measured data, but of sensible system historical behaviour, two desirable outcomes of the history-matching process are achieved. The first of these outcomes is that the posterior uncertainties of model parameters are reduced through provision of this extra information; the uncertainties of decision-critical model predictions may thereby also be reduced. The second of these outcomes is a reduced probability of aberrant model behaviour under future simulated drought conditions.

Components of the measurement objective function are now described. This objective function was employed on all occasions on which history matching was undertaken. On the first occasion of history matching, PEST_HP was employed to achieve a parameter field of minimized error variance; this is the calibrated parameter field. On the second occasion of history-matching, PESTPP-IES was employed to attain an ensemble of stochastic parameter fields which respect constraints imposed by both hard and soft data; these constitute samples of the posterior parameter probability distribution.

4.3.2 Hard Data

It is the task of the Biggenden groundwater model to predict future water levels and extraction rates in two pumped bores, namely RN156052 and RN155332. Water level measurements from the former bore are available over the history-matching period. This is the bore from which most water is extracted. Historically, water level measurements were taken both when the pump was operating, and when it was not operating. There is a consistent difference of 5m between these two levels. (This is consistent with cell-to-well head corrections calculated using representative local hydraulic conductivities; see Peaceman, 1978.) In the Biggenden groundwater model, pumping is continous. Over the history-matching period, model-calculated heads in the cell in which the pumping well is situated are matched to “pumps-off” water levels measured in RN156052.

No historical water level measurements are available for RN155332. Nevertheless history-matching constraints arising from soft data were applied to historical pumping rates from this well; see below. (Recall that the Biggenden model is configured to automatically reduce extraction rates if borehole water levels fall below use-specified thresholds.)

Over the history-matching period, model-calculated heads for monitoring bores RN13600186, RN13600187 and RN13600296 were matched to their observed counterparts. These bores are relatively close to the two pumped bores. These are the only observation bores within the model domain, except for those at its south-western boundary. Recall that a LUMPREM model was calibrated against water level variations in one of these boundary wells in order to populate all model alluvial boundaries with time-varying heads over both the calibration and predictive periods.

The “hard data” component of the measurement objective function was supplemented with measurements of temporal head differences. For each of the above bores, the difference between the head measured at any particular time, and the first head measurement available in the bore, was taken. This was done for observations and for model-calculated counterparts to observations. Discrepancies between these observed and model-calculated differences comprised another component of the measurement objective function. Head differences are rich in information on storage and recharge parameters. By explicitly including these differences in the measurement objective function, and by weighing this component of the objective function for visibility, it is ensured that that this information is transferred to model parameters.

4.3.3 Soft Data

The importance of "soft data" in guiding the history-matching process towards rejection of parameters that promote aberrant model behaviour has already been discussed. In a history-matching process, soft data is often embodied in penalty functions. These contribute to the overall measurement objective function only if certain system behaviour thresholds are crossed. Penalty functions steer the history matching process away from hydrogeologically inconsistent parameter sets.

As was previously noted, the Biggenden groundwater model automatically reduces extraction rates from pumped bores as borehole water levels approach pump intake levels. Biggenden model post-processors report the calculated pumping rate and compare it with the desired pumping rate. In order to prevent unrealistic derating of extraction during the history-matching period, a model-extracted pumping rate that is less than 90% of that which is sought over any one-month period is decreed to trigger a residual. This residual is the difference between 90% of the sought pumping rate and that calculated by the model. This residual is weighted and added to the measurement objective function to discourage estimation of parameter sets which prevent the model from matching observed extraction.

Some of the water that enters the Biggenden groundwater model domain as recharge and inflow from lateral boundaries leaves the groundwater system as baseflow to Degilbo and Mungore Creeks. Degilbo Creek is gauged at a station that is installed 25 km downstream of the study area. If, during any month of the history-matching period, the average model-calculated baseflow exceeds the minimum baseflow recorded for that month at the downstream gauging station, an objective function penalty is incurred. This penalty function codifies the assumption that the baseflow contribution to Degilbo Creek flow is likely to be considerably higher at the gauging station than in the immediate vicinity of Biggenden.

MODFLOW 6 CHD boundary conditions line the northern and southern alluvial boundaries of the model domain. To promote an assumed northerly flow direction along the main alluvial trunk, a penalty is incurred if water leaves the domain of the Biggenden groundwater model through the southern CHD alluvial boundary, or if water enters the model domain through the northern CHD alluvial boundary.

When rainfall is low, seepage of groundwater to the surface is low, and is experienced only at the lowest points in the landscape, these being occupied by creeks and gullies. Hence, if the model-calculated depth to the phreatic surface in any other part of the model domain is less than zero during an historical dry season, these negative depths comprised residuals that contribute to an objective function penalty. Conversely, during rainy periods, the depth to groundwater is not expected to exceed its largest recorded wet season value. Consequently, if the model-calculated depth to the phreatic surface was greater than 12m in any part of the model domain during an historical wet season, depths exceeding this threshold comprise residuals that contribute to an objective function penalty. The threshold of 12m corresponds to the deepest wet season water level recorded in any monitoring bore residing in the domain of the Biggenden model.

4.3.4 Weighting

Each of the above-described components of the measurement objective function was assigned to its own "observation group". PEST_HP and PESTPP-IES report the contribution to the total objective function made by these different groups as the history-matching process progresses. Ideally, the history-matching process should commence with each of these contributions visible in the overall objective function; this applies especially to objective function components pertaining to hard data. Rough equality of objective function contributions prevents dominance of the objective function by a single component, and consequential diminishment of the worth of information carried by other components. This strategy was adopted in history matching of the Biggenden model.

This weight-balancing procedure can encounter difficulties when applied to penalty components of an objective function. Penalties can be small, or even zero, for parameter sets chosen by the user to commence a history-matching process. In the present case, weights applied to some penalty function components were initially guessed, with a view to adjusting them if this proved necessary. Little adjustment was required, however.

5 MODEL CALIBRATION

5.1 What is Calibration?

The word “calibration” should not be used interchangeably with “history matching”.

Calibration implies parameter uniqueness. Obviously, there is insufficient data available for unique estimation of parameters assigned to the Biggenden groundwater model. It follows that there is insufficient data available for a unique determination of hydraulic properties within its domain, nor of the parameters which govern seasonal fluctuations of boundary heads, and hence the heads themselves.

The name given to the process through which uniqueness is sought for an ill-posed inverse problem is “regularisation”. This can be implemented in a number of ways. One of these ways is to impose constraints on departures of estimated parameter fields from a default, user-specified condition; this methodology is broadly referred to as “Tikhonov regularisation”. If regularisation constraints posit a preferred state of parameter homogeneity, then the parameter field that emerges from this process is that which exhibits the minimum heterogeneity that is required for model outputs to match field observations.

A parameter field attained through regularised inversion must be seen for what it is. It makes no claim for correctness, only for minimized potential for incorrectness (i.e. minimized error variance). If the information content of a history-matching dataset is low, the potential for calibrated parameter field incorrectness may still be very high. The same applies to some predictions made by the calibrated model.

In PEST’s implementation of Tikhonov regularisation, the inverse problem is reformulated as a constrained minimisation problem. A “regularisation objective function” is minimized. This objective function measures parameter departures from their preferred state. This minimisation problem should be formulated in such a way that departures from this state are more tolerable if they take place in ways that are geologically meaningful. The constraint imposed on the minimisation process is that a “measurement objective function” (a measure of model-to-measurement misfit) achieves a certain (presumably low) value. This value should be set in accordance with the level of measurement noise that is associated with field measurements, or with the level of “structural noise” that characterizes a model’s inability to reproduce field measurements.

The Biggenden model was calibrated according to the above prescription. A good fit between field measurements and corresponding model outputs was attained. Following acquisition of a parameter field of minimized error variance, a further 250 history-match-constrained, heterogeneous parameter fields were obtained. Each of these parameter sets was designed to reflect the level of heterogeneity that may exist within the model domain. Collectively they were used to explore the range of predictive possibilities that are compatible with expert knowledge on the one hand, and the historical behaviour of the system on the other hand.

5.2 So Why Calibrate?

It is legitimate to question the need for procurement of a solution of minimized error variance to the inverse problem of model calibration, when the outcome of a modelling exercise such as this Biggenden example should be quantification of the uncertainties of decision-critical predictions rather than a single prediction of minimized error variance. An alternative to calibration is the use of history-matching methods that attempt to directly sample the posterior parameter probability distribution. Ensemble methods such as that employed by PESTPP-IES can achieve this outcome at a comparatively small numerical cost.

Model calibration can be expensive. It requires calculation of a Jacobian matrix (i.e. a sensitivity matrix). The filling of this Jacobian matrix requires that at least one model run be undertaken for each adjustable parameter. The Jacobian matrix must be re-filled during each iteration of the history-matching process. In contrast, ensemble methods require far fewer model runs per iteration - possibly only a few hundred - regardless of the number of adjustable parameters. (Note, however, that variants of ensemble methods can be used to "calibrate" a model using an approximate, rank-deficient Jacobian matrix. In some situations this may prove to be a viable alternative to conventional calibration. Rank-deficient matrices are those for which at least one row or column is a linear combination of other rows or columns.)

The choice of whether posterior uncertainty analysis should be preceded by a calibration step is personal. However, in making this choice, the following should be taken into account.

- Model construction usually takes place in stages. Concepts on which it rests often undergo considerable refinement during subsequent attempts at history-matching, especially if difficulties are encountered in fitting model outputs to field measurements, and/or if estimated parameter fields are unrealistic. Initial attempts at history-matching therefore comprise a lengthy process wherein hypotheses pertaining to concepts on which the current version of the model rests are tested and refined. It makes sense to test these concepts by seeking a single solution to an inverse problem that, by design, departs minimally from that which is considered to be "most realistic" from an expert knowledge point of view. If a good fit with the calibration dataset cannot be attained with a reasonable parameter field, this becomes readily apparent.
- Experience demonstrates that an inversion process based on a full Jacobian matrix can often promulgate a better fit with a history-matching dataset than an inversion process based on a rank-deficient Jacobian matrix. For reasons already outlined, achievement of a good fit with hard data is an essential component of Biggenden groundwater modelling.
- Experience has also demonstrated that the performance of PESTPP-IES in sampling the posterior parameter probability distribution can sometimes be more numerically efficient if this process begins with random parameter fields that sample a linear approximation to this distribution. Under these circumstances PESTPP-IES may require fewer iterations to achieve a good fit with field measurements than if initiated with samples of the prior parameter probability distribution. A linear approximation to the posterior parameter probability distribution can be calculated using the Jacobian matrix that emerges from calibration of the model. This strategy was adopted in the present example.

5.3 Regularisation

5.3.1 Regularisation in general

Regularisation can be viewed in a number of ways. By definition, it is essential to attainment of parameter uniqueness. It can also be viewed as a set of penalties, mainly (but not always) applied to parameters. The greater the extent to which parameters depart from a preferred condition, the greater is the penalty.

When calibrating a model using PEST, a fundamental difference between a regularisation penalty function and a penalty function which contributes to the measurement objective function is that PEST determines the weights that it applies to the former, but accepts user-specified weights for the latter. PEST ascribes values to regularisation weights that do not compromise its ability to fit hard data, but prevents it from over-fitting this data; the metric for over-fitting is the user-supplied target measurement objective function. Alternatively, if no such threshold is sought, so that the best fit possible with the measurement dataset is pursued, then PEST adjusts regularisation weights in such a way that this fit is attained with the most reasonable parameter set. Regularisation was implemented in the second of these ways when calibrating the Biggenden groundwater model. See PEST documentation for a more detailed description of its regularisation functionality.

5.3.2 Regularisation Specifics

For the five LUMPREM models that simulate recharge in different parts of the model domain, penalties were imposed on differences between values assigned to LUMPREM parameters of the same type ascribed to models operating in different zones. Hence the preferred parameter condition for LUMPREM recharge models was that of homogeneity across all land-use types. In addition to this, a penalty condition was decreed to arise if the average recharge calculated by any one of these five LUMPREM recharge models was less than 1%, or greater than 5%, of average rainfall.

It is a requirement of PEST-based inversion that initial values be provided for all parameters. For pilot point parameters (recall that pilot points are used for parameterisation of Kh and Sy, as well as GHB and DRN conductances), regularisation penalties were incurred to the extent that parameters departed from their initial values. Initial/preferred values for parameters of each type are listed in Table 5.1.

Parameter type	**Initial value**
Hydraulic conductivity	1 m/d
Specific yield	1%
GHB boundary conductance	1 m^2/d
DRN boundary conductance	1 m^2/d

Table 5.1 Initial values of pilot points parameters.

Covariance matrices were used instead of weights for calculating components of the regularisation objective function pertaining to pilot point parameters. This encourages PEST to spatially distribute any heterogeneity that it must introduce to a parameter field in order for the model to fit a calibration dataset, rather than allow this heterogeneity to arise on a point-by-point basis. The covariance matrix employed for regularisation of each of these parameter types was based on an anisotropic variogram whose range in its principle direction is able to vary in space with pilot point density. For all pertinent parameter types, the local range was set to about 4 times the average distance between pilot points. An anisotropy of 2.0 was used throughout the model domain, with the principle axis of anisotropy oriented roughly parallel to the local direction of the alluvial system. Construction of these covariance matrices was enabled using the PPCOV_SVA utility; this is a member of the PEST Groundwater Utility Suite.

5.4 Outcomes of the Calibration Process

5.4.1 Performance of PEST_HP

The Biggenden groundwater model takes around 6 minutes to run on a computer using an Intel Xeon 3.2GHz Octa-core processor (E5-1660 v4) with 128GB RAM. History-matching was undertaken using 36 nodes (1,376 CPU cores) of a Windows-based High Performance Computing (HPC) platform owned by the Queensland Department of Environment and Science (DES).

The PEST_HP version of PEST was used for calibration of the Biggenden groundwater model. PEST_HP was allowed to run for 50 iterations. Initial and final measurement objective function components are listed in Table 5.2. The latter pertain to iteration 24. However objective functions were within 10% of these values after only 12 iterations. The number of model runs required for completion of 12 and 24 iterations was approximately 27,300 and 54,600 respectively. PEST_HP employed three-point finite-difference derivatives to fill the Jacobian matrix; these are more accurate than derivatives calculated using forward differences.

Objective function component	**Initial**	**Final**
Heads in pumping well	30,000	222.7
Heads in observation wells	90,000	204.8
All heads	120,000	427.5
Temporal head differences in observation wells	90,000	197.2
De-rating of extraction	30,000	11.8
Prevention of incorrect flow direction across boundaries	60,000	0.0
Sensibility of wet season heads	30,000	0.0
Sensibility of dry season heads	0.0	0.3
Deterrence of excessive baseflow	30,000	0.5

Table 5.2 Measurement objective function components achieved by PEST_HP.

The measurement objective functions listed in Table 5.2 are somewhat arbitrary. They are an outcome of the “weighting for visibility” strategy that was described above. However they are directly comparable with objective functions attained using the PESTPP-IES ensemble smoother; the latter are listed in Table 6.1.

5.4.2 Heads Calculated by the Calibrated Model

Figure 5.1 shows head measurements and their model-calculated counterparts for bore RN156052 (the pumped bore) and for the three nearby observation wells. (Borehole water levels in RN156052 that were measured while the pump was operating are omitted from Figure 5.1d.) Figure 5.2 shows these same measured heads together with heads calculated by the model at the end of each of its stress periods.

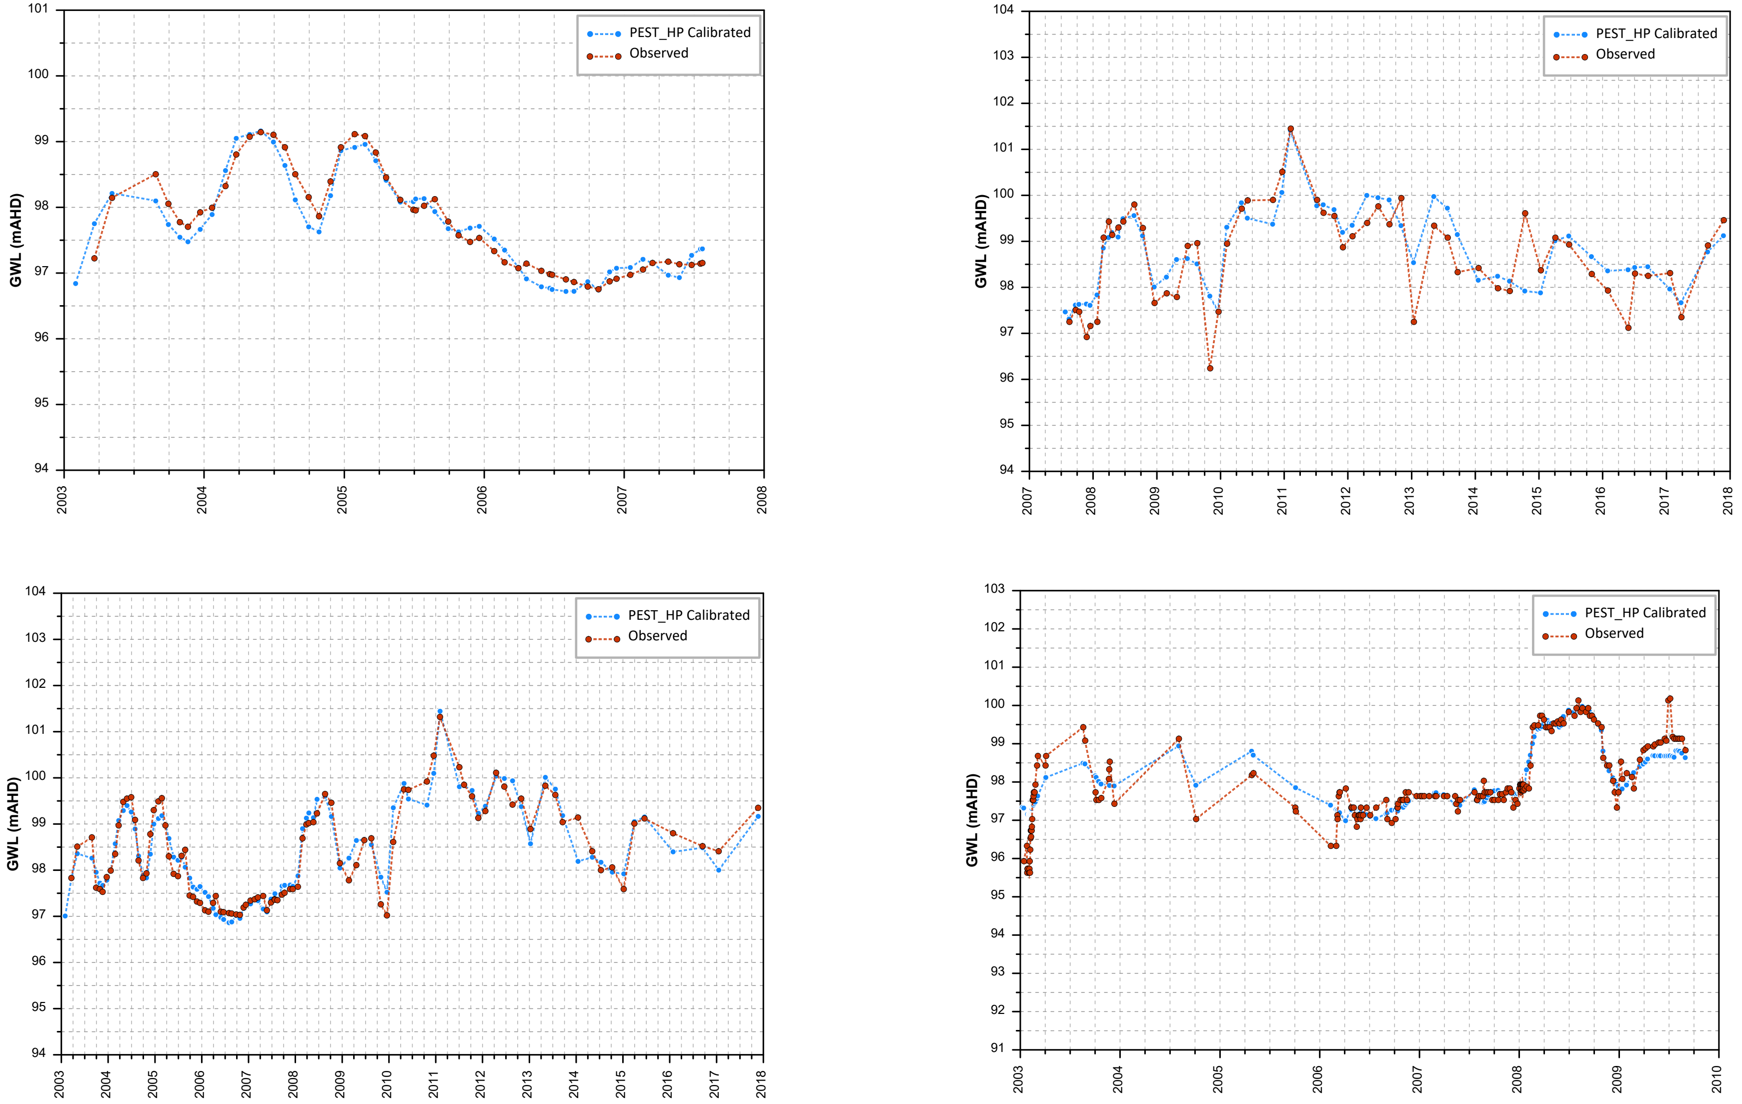

Figure 5.1 Observed and modelled heads for (a) RN13600186 (top left), (b) RN13600187 (top right), (c) RN13600296 (bottom left) and (d) RN156052 (bottom right). Only modelled heads corresponding to field-measured heads are shown in these figures.

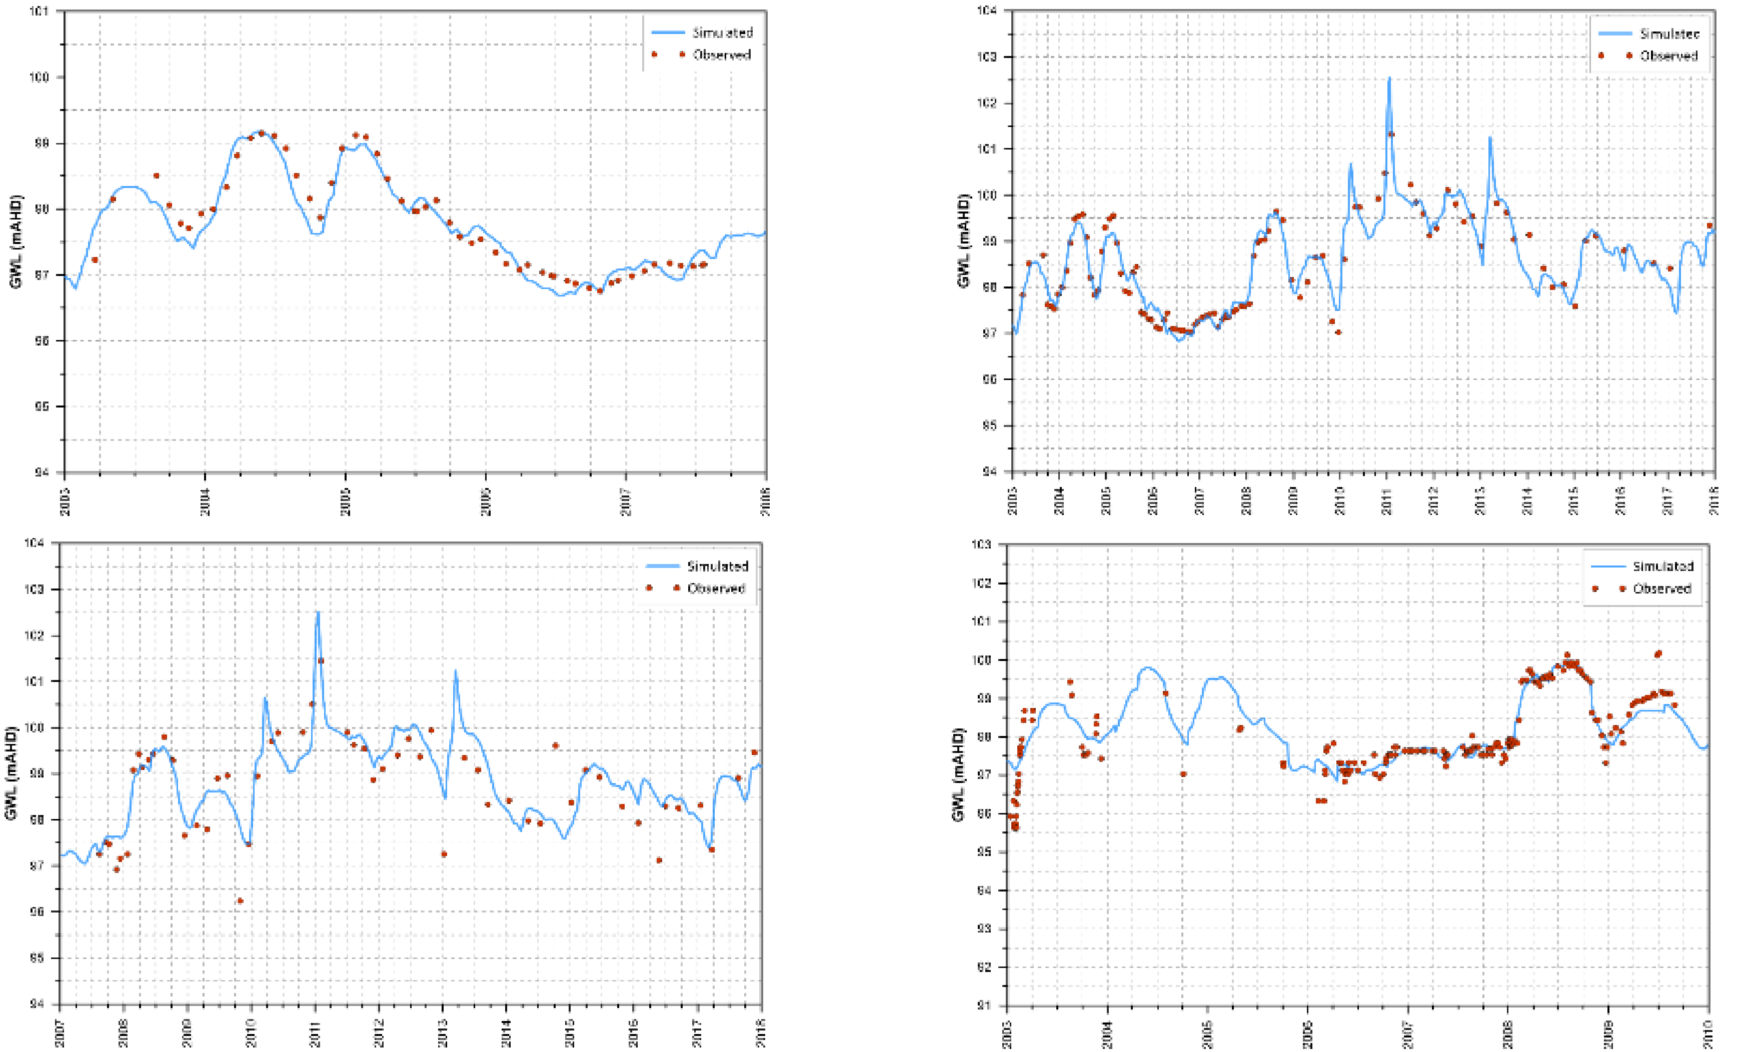

Figure 5.2 Observed and modelled heads for (a) RN13600186 (top left), (b) RN13600187 (top right), (c) RN13600296 (bottom left) and (d) RN156052 (bottom right).

Piezometric surfaces calculated by the model during a particularly wet period (1st February, 2011) and during a particularly dry period (1st November, 2006) are shown in Figure 5.3 and Figure 5.4 respectively.

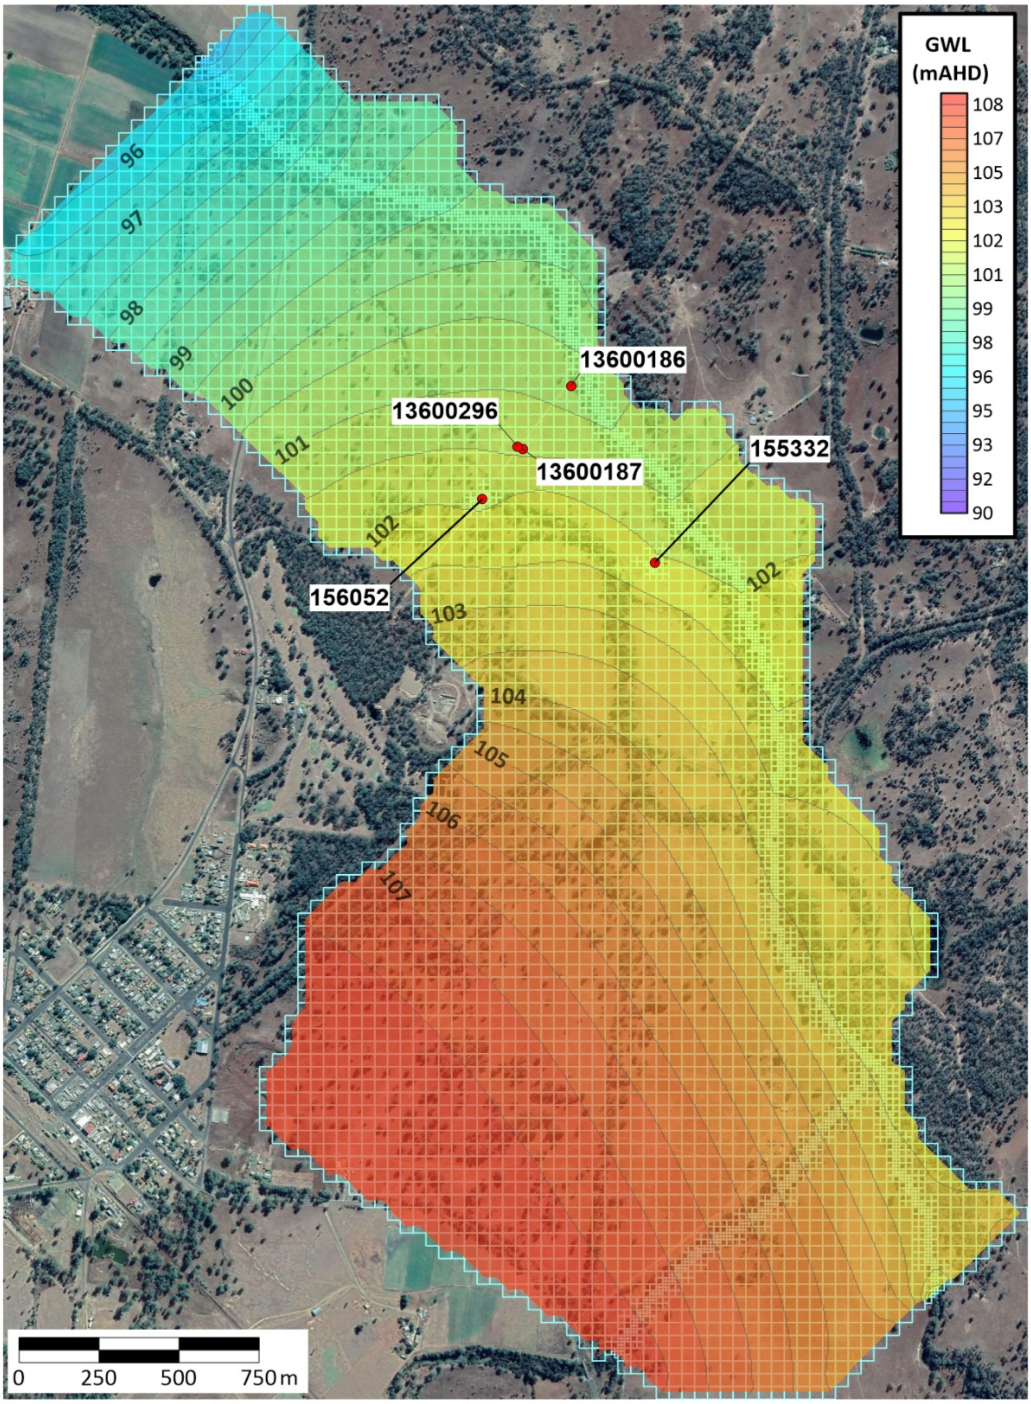

Figure 5.3. Contours of model-calculated groundwater heads for 1st February, 2011.

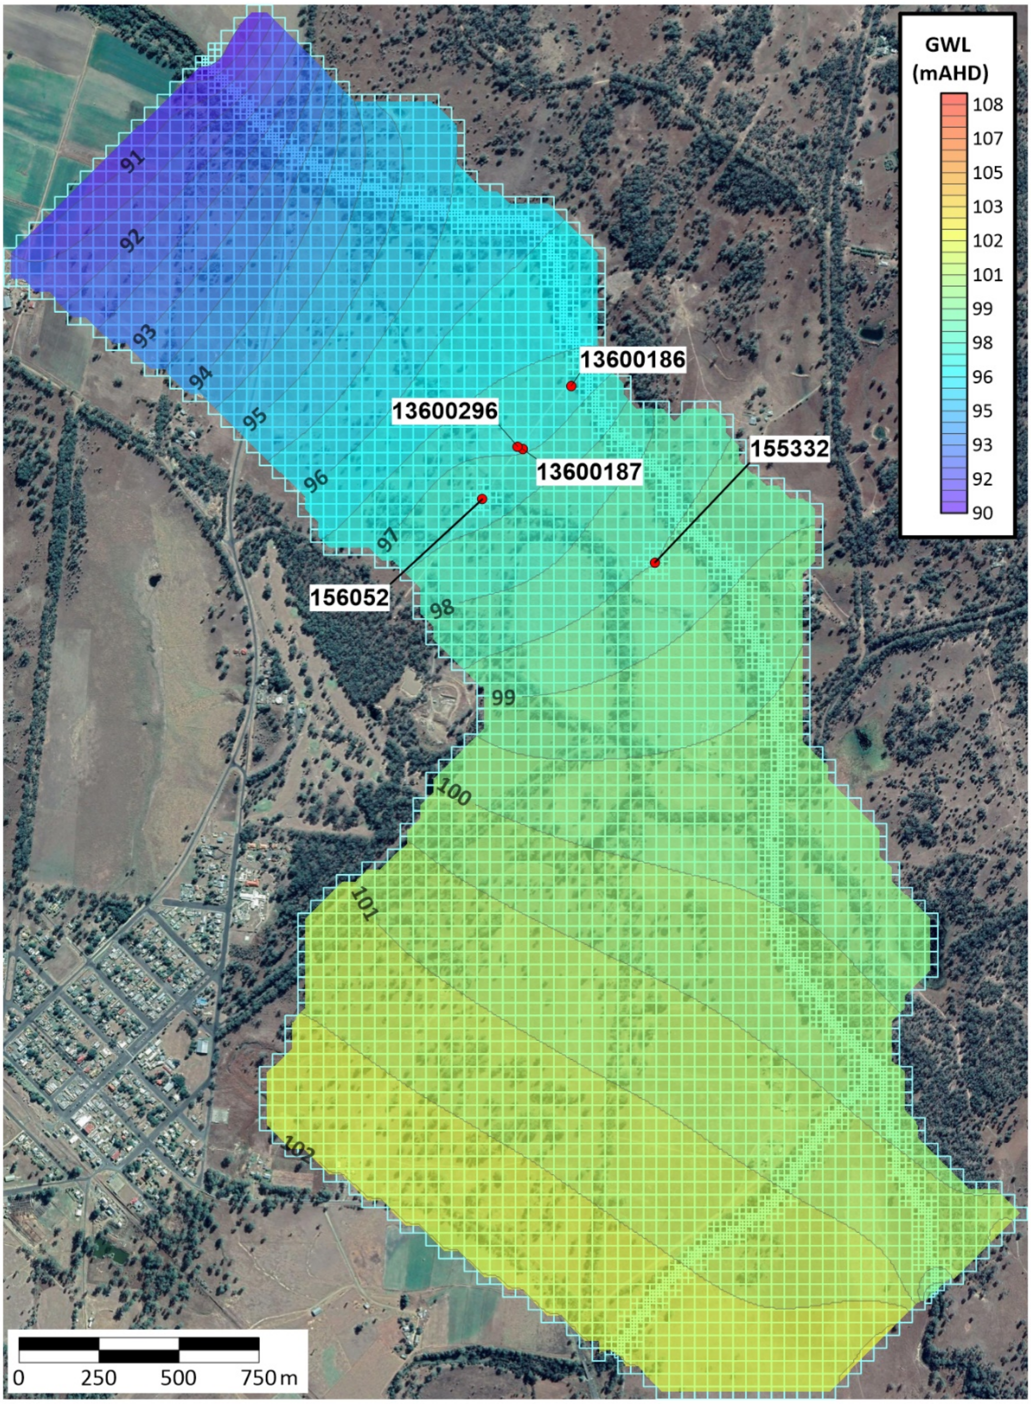

Figure 5.4. Contours of model-calculated groundwater heads for 1st November, 2006.

Figure 5.5 depicts variation of GHB boundary heads with time at points about half way along the eastern and western boundaries of the model domain. Recall that these are calculated using different instances of the LUMPREM model.

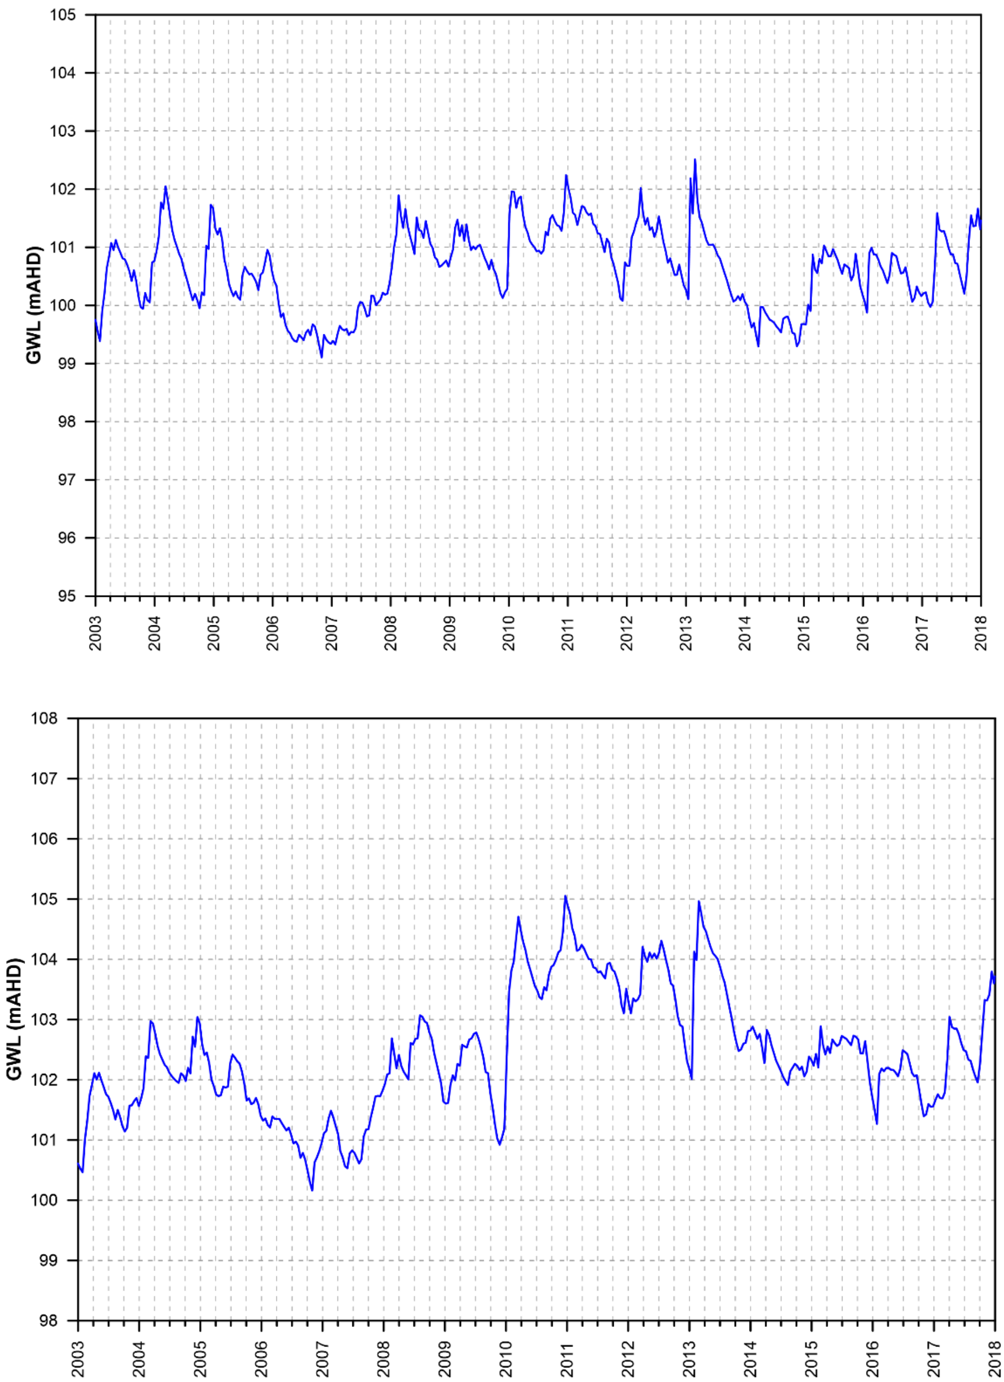

Figure 5.5. Heads at points about half way along the eastern (top) and western (bottom) boundaries of the Biggenden groundwater model domain.

5.4.3 Parameter Fields

The calibrated Kh and Sy parameter fields are shown in Figures 5.6 and 5.7.

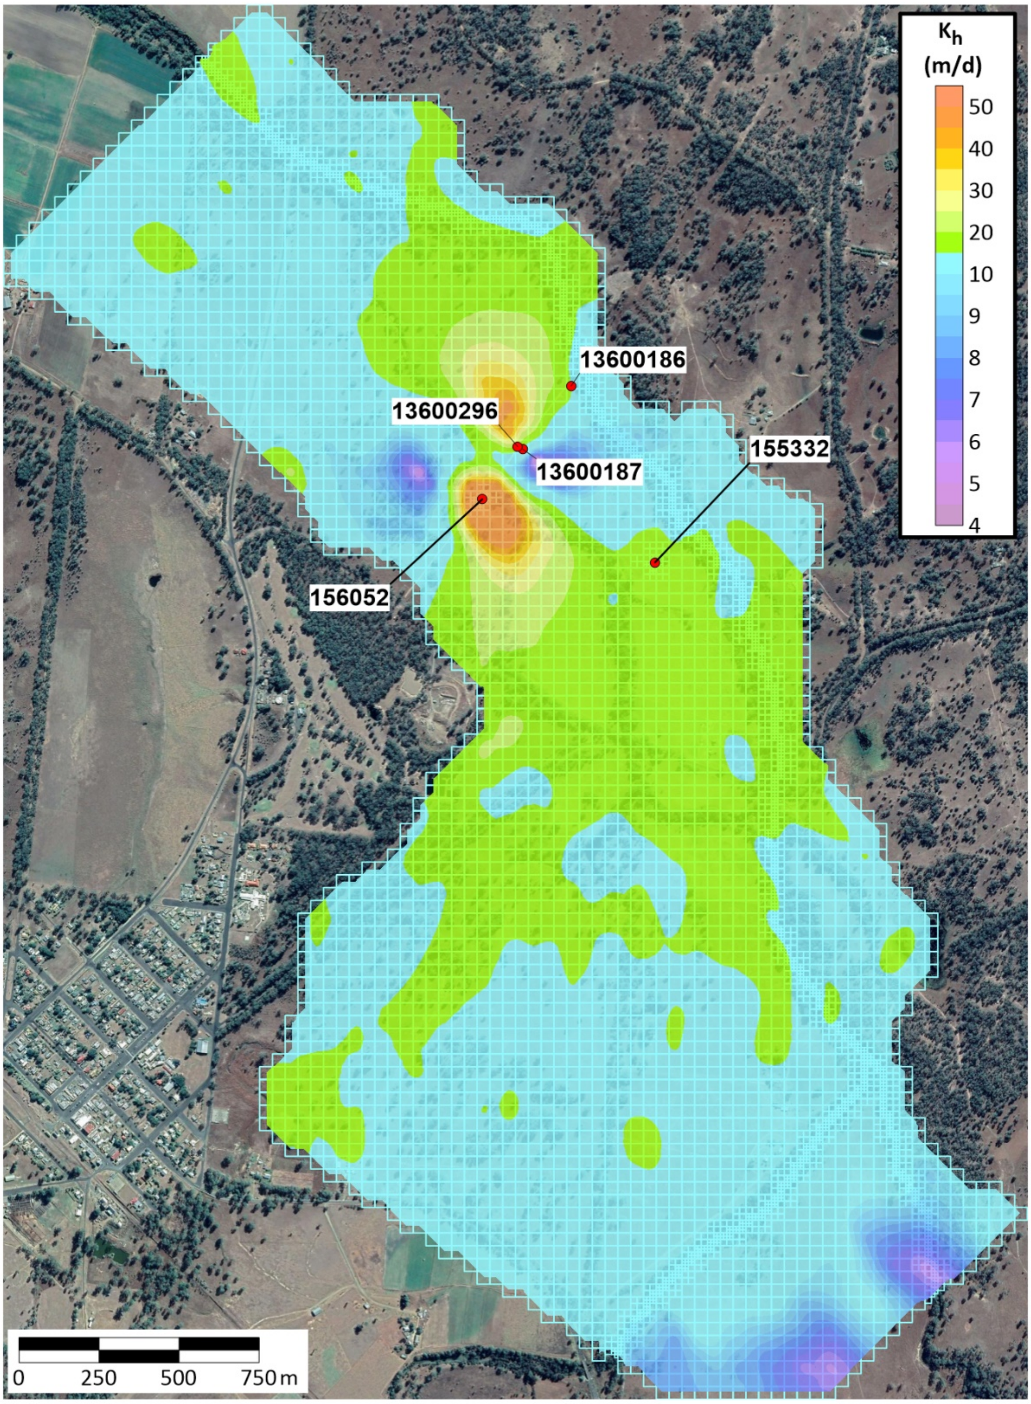

Figure 5.6. Calibrated hydraulic conductivity (i.e. Kh) parameter field.

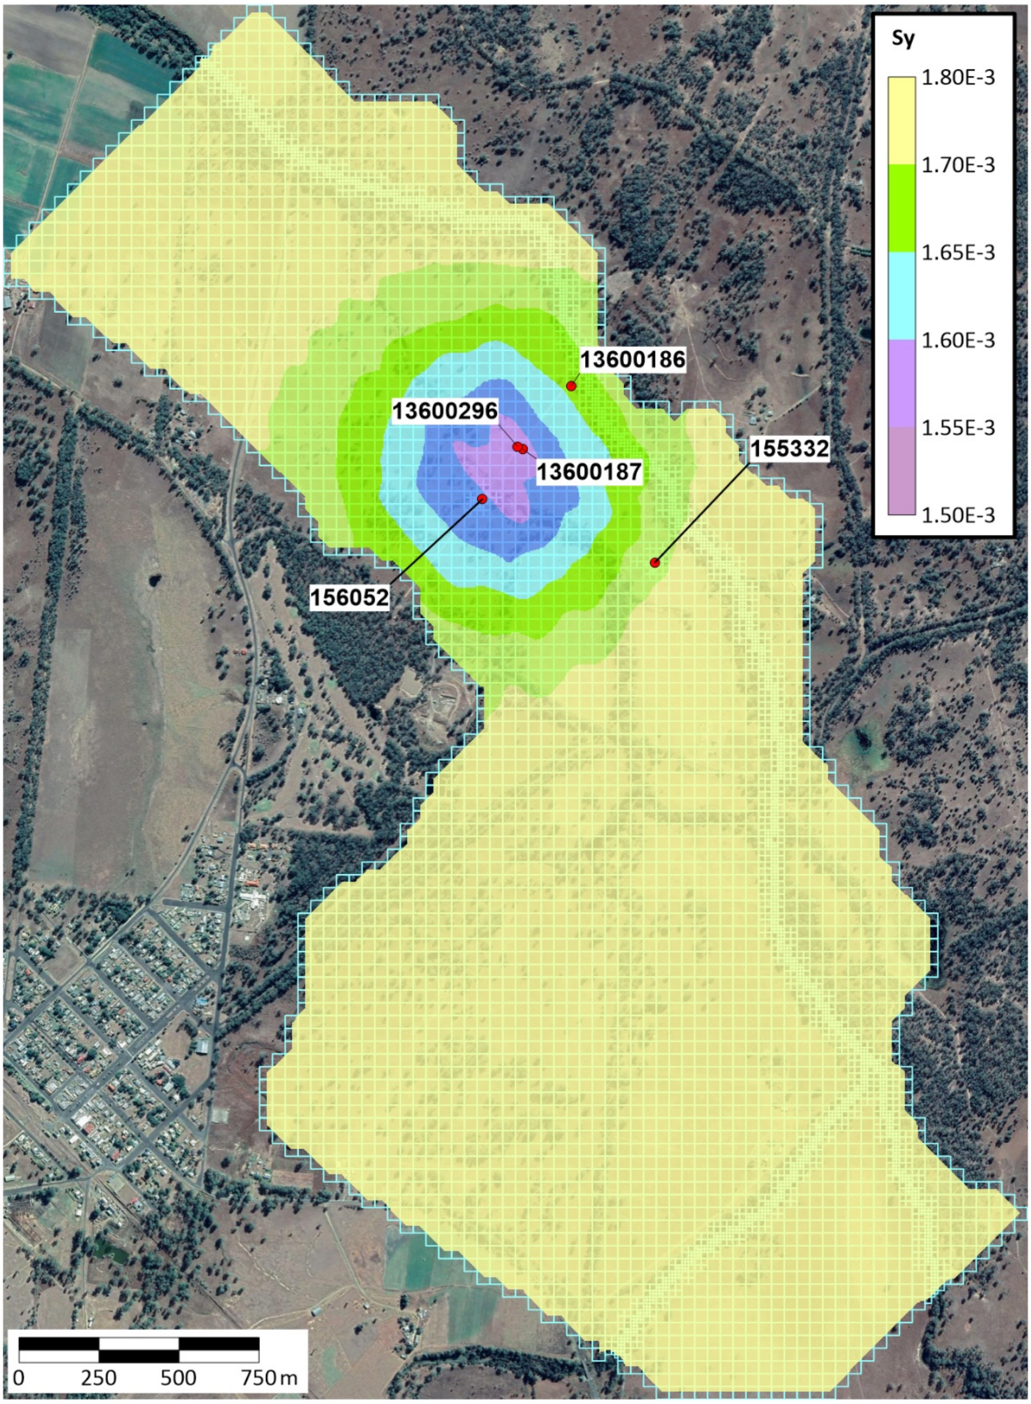

Figure 5.7. Calibrated specific yield (i.e. Sy) parameter field.

It is apparent from Figures 5.6 and 5.7 that, in accordance with regularisation constraints, most heterogeneity is introduced to the model domain in the vincinity of the pumping and observation wells. This is required in order to attain a good fit between model outcomes and the hard component of the calibration dataset. The patterns that appear in Figures 5.6 and 5.7 are likely to be only loosely reflective of the actual patterns of hydraulic properties that prevail in the Degilbo Creek alluvial system. The relationships between parameters that are estimated through regularised solution of an ill-posed inverse problem and the actual hydraulic properties of a system are complex; for linear systems they are revealed by the rows of the so-called resolution matrix (Menke, 2018; Aster et al 2013).

6. SAMPLING THE POSTERIOR

6.1 General

As has already been discussed, the desired outcome of the Biggenden modelling process is a suite of "realistic" parameter fields that all allow the model to replicate historical hard data. Water supply security under any given scenario of future rainfall can then be tested using all of these parameter fields.

The PESTPP-IES iterative ensemble smoother is described by White (2018). The numerical algorithm on which it is based is described by Chen and Oliver (2013). Ideally, an ensemble smoother commences the history-matching process with a suite of parameter fields which comprise samples of the prior parameter probability distribution. Through a succession of iterations of the ensemble inversion process, these parameter fields are adjusted until they allow the model to replicate field observations; in doing so, they thereby constitute samples of the posterior parameter probability distribution. Adjustments to parameter fields are calculated using a rank-deficient Jacobian matrix. This is formed from covariances between individual parameters and individual model outcomes. The number of model runs per iteration that is required to fill the rank-deficient Jacobian matrix is equal to the number of parameter realisations that comprise the ensemble. Ideally, this number should be no lower than the dimensionality of the solution space that characterizes the inverse problem. The ensemble adjustment process gains its efficiency from the fact that the number of runs per iteration need be no higher than this, regardless of the number of parameters requiring adjustment. In history-matching of the Biggenden model, PESTPP-IES employed 500 realisations; those with the lowest 250 measurement objective functions were used to make model predictions.

We undertook a number of PESTPP-IES runs. Based on this experience, we concluded that its performance in attaining a good fit with the Biggenden model history-matching dataset in a reasonably small number of iterations is superior if its initial random parameter fields are centred on a parameter field that already achieves this fit. We found that its performance is further enhanced if initial random parameter fields are sampled from a linear approximation to the posterior parameter probability distribution calculated from the Jacobian matrix computed by PEST_HP during the previous calibration process. We suspect that part of the reason why its performance is so susceptible to improvement in this manner lies in the high degree of nonlinearity introduced to the inverse problem through the use of penalty functions. Penalties are sensitive to parameters over only part of their ranges.

At the time of writing, the use of ensembles is relatively new to the groundwater industry. We assist the reader who is unfamiliar with this technology, but who may be interested in using the methods described herein in his/her own modelling work, by providing a few implementation details.

6.2 Some Implementation Details

6.2.1 Generation of Initial Parameter Fields

Initial random parameter fields were generated using the PEST RANDPAR3 utility. This requires a covariance matrix to specify parameter uncertainty and correlation. This covariance matrix was generated using the PEST PREDUNC7 utility. PREDUNC7 calculates a linear approximation to a posterior parameter covariance matrix based on a user-supplied prior covariance matrix and a matching Jacobian matrix.

The prior covariance matrix employed by PREDUNC7 was calculated using the PEST PPCOV_SVA utility. For Kh, Sy and boundary conductance parameters, this is the same covariance matrix as that which was used for regularisation. Prior standard deviations ascribed to the log (to base 10) of these parameters were all 0.5; the square of this standard deviation (i.e. the variance) comprises the diagonal elements of the covariance matrix. (Note that standard deviations are of no consequence when a matrix is used for regularisation; they are only of consequence when a covariance matrix is used for uncertainty analysis.)

All other parameters employed by the Biggenden model (these predominantly pertaining to instances of the LUMPREM model) were considered to be independent of each other. Prior standard deviations were assigned to them in accordance with the properties that they represent.

6.2.2 Measurement Noise

The same observation weights were employed in the PESTPP-IES history-matching process as those which were employed by PEST_HP. As is discussed above, weighting was based on visibility of objective function components in the overall objective function.

PESTPP-IES allows a user to add realisations of measurement noise to observations comprising the history-matching dataset. These realisations were added only to hard data, namely heads measured in boreholes. The standard deviation chosen for head measurement noise was back-calculated from the fit attained between observed heads and their simulated counterparts by PEST_HP.

6.3 Outcomes of PESTPP-IES History Matching

6.3.1 Performance

Execution of PESTPP-IES continued for 21 iterations, this requiring a total of approximately 14,700 model runs; however parameter fields pertaining to iteration 19 were employed for future model runs, as objective function statistics deteriorated slightly during ensuing iterations. Around 13,400 model runs were required for completion of iteration 19.

6.3.2 Objective Functions

Objective function statistics for iteration 19 are presented in Table 6.1. These are directly comparable with those achieved using PEST_HP; the latter are listed in Table 5.2.

Objective function component	Average	Minimum	Maximum	Standard deviation
Heads in pumping well	210.0	184.3	349.5	16.5
Heads in observation wells	202.3	175.3	397.9	21.9
All heads	412.3	374.0	701.0	35.2
Temporal head differences in observation wells	201.8	174.6	877.3	36.1
De-rating of extraction	27.3	7.4	65.0	9.2
Prevention of incorrect flow direction across boundaries	5.1	0	184.9	13.4
Sensibility of wet season heads	0	0	2.6	0.1
Sensibility of dry season heads	1.7	0	28.0	3.2
Deterrence of excessive baseflow	3.0	0	39.8	5.2

Table 6.1 Measurement objective function components achieved by PESTPP-IES.

Of the 500 realisations that were adjusted by PESTPP-IES, the best 250 were selected for future model use. For these parameter fields the heads component of the objective function varies between 374.0 and 427.6, while the head difference component of the objective function varies between 174.6 and 223.1.

6.3.3 Heads

Figure 6.1 is comparable with Figure 5.1. Figure 6.1 compares measured heads in the pumped bore and three nearby observation bores with their model-calculated counterparts for the best 250 realisations calculated by the PESTPP-IES ensemble. In these figures, measured heads are linked by straight lines. Figure 6.2 shows measured heads in these same wells. However modelled heads are plotted at the end of every simulation stress period. This figure is directly comparable with Figure 5.2.

Wet season (1st February, 2011) and dry season (1st November, 2006) piezometric surfaces throughout the model domain, calculated using four randomly selected realisations from the best 250 realisations achieved by PESTPP-IES are shown in Figures 6.3 and 6.4. These figures are directly comparable with Figures 5.3 and 5.4.

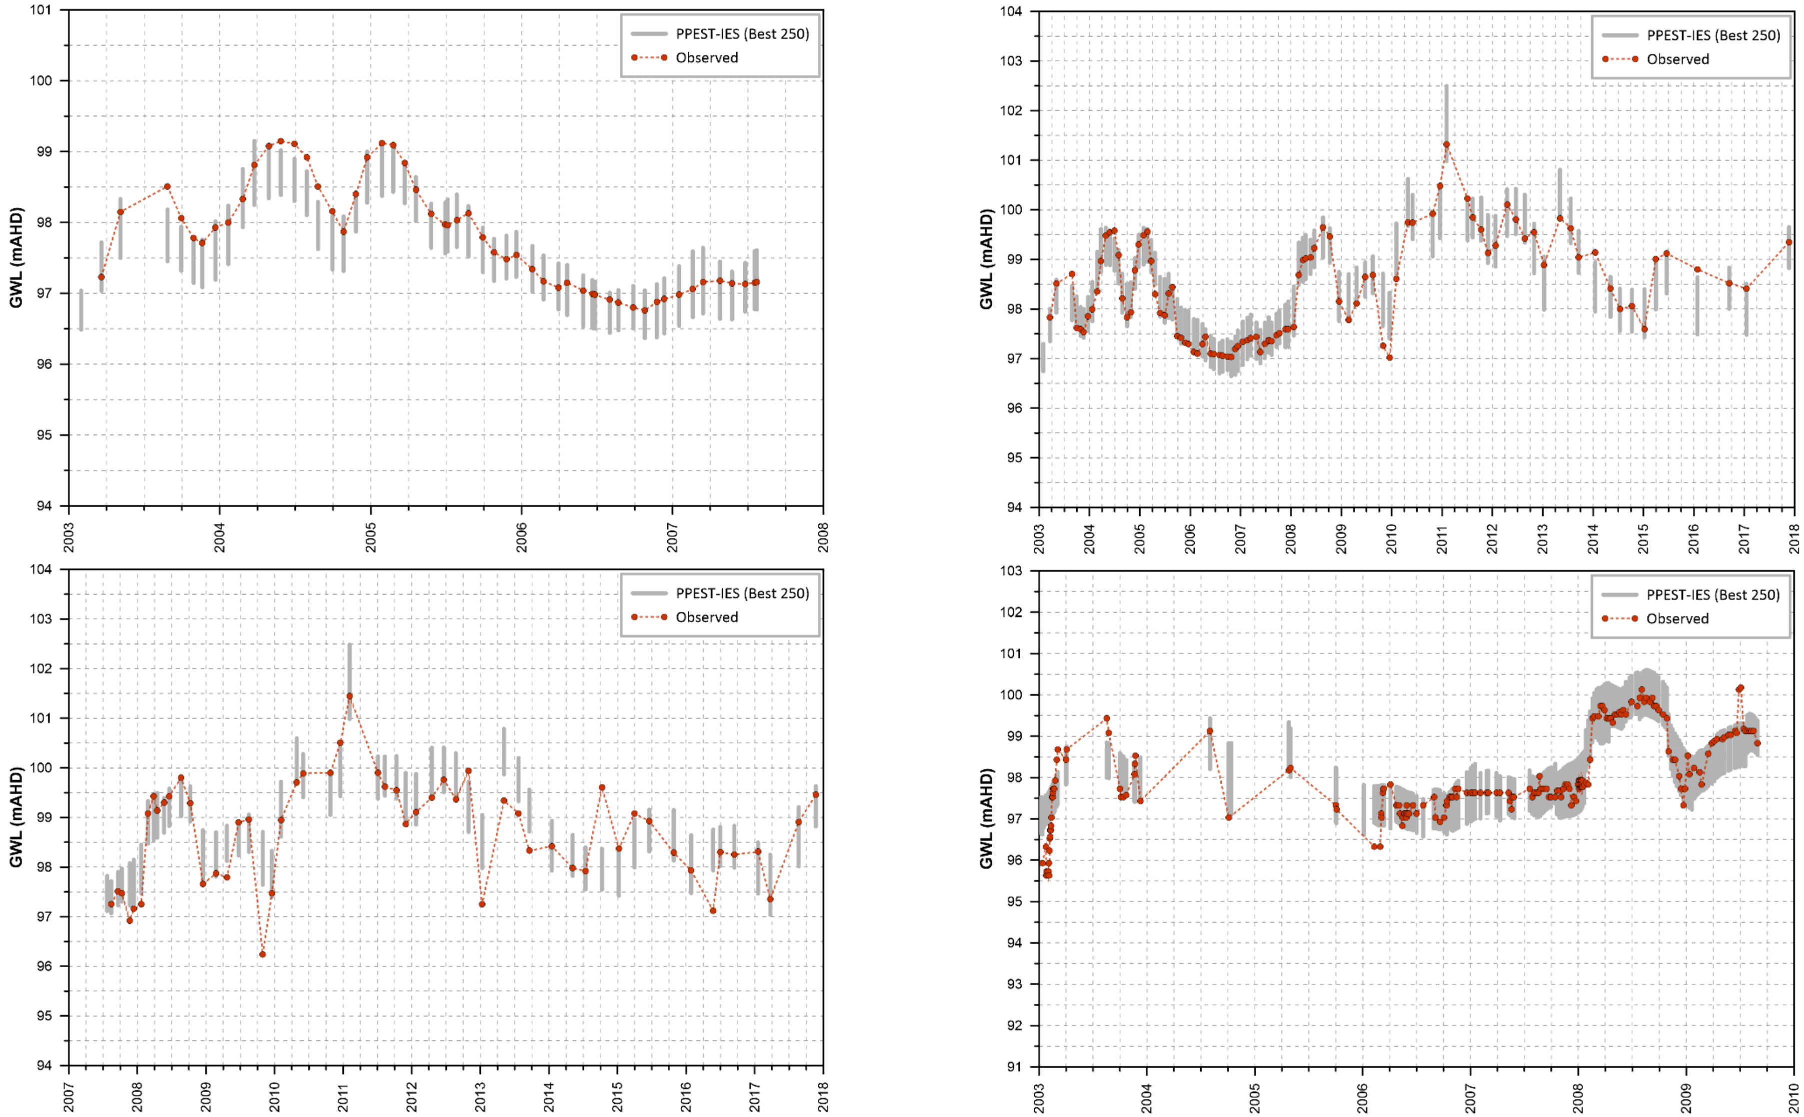

Figure 6.1. Comparison of observed and modelled heads for (a) RN13600186 (top left), (b) RN13600187 (top right), (c) RN13600296 (bottom left) and (d) RN156052 (bottom right) for the best 250 realisations of the ensemble. Only modelled heads corresponding to field-measured heads are shown.

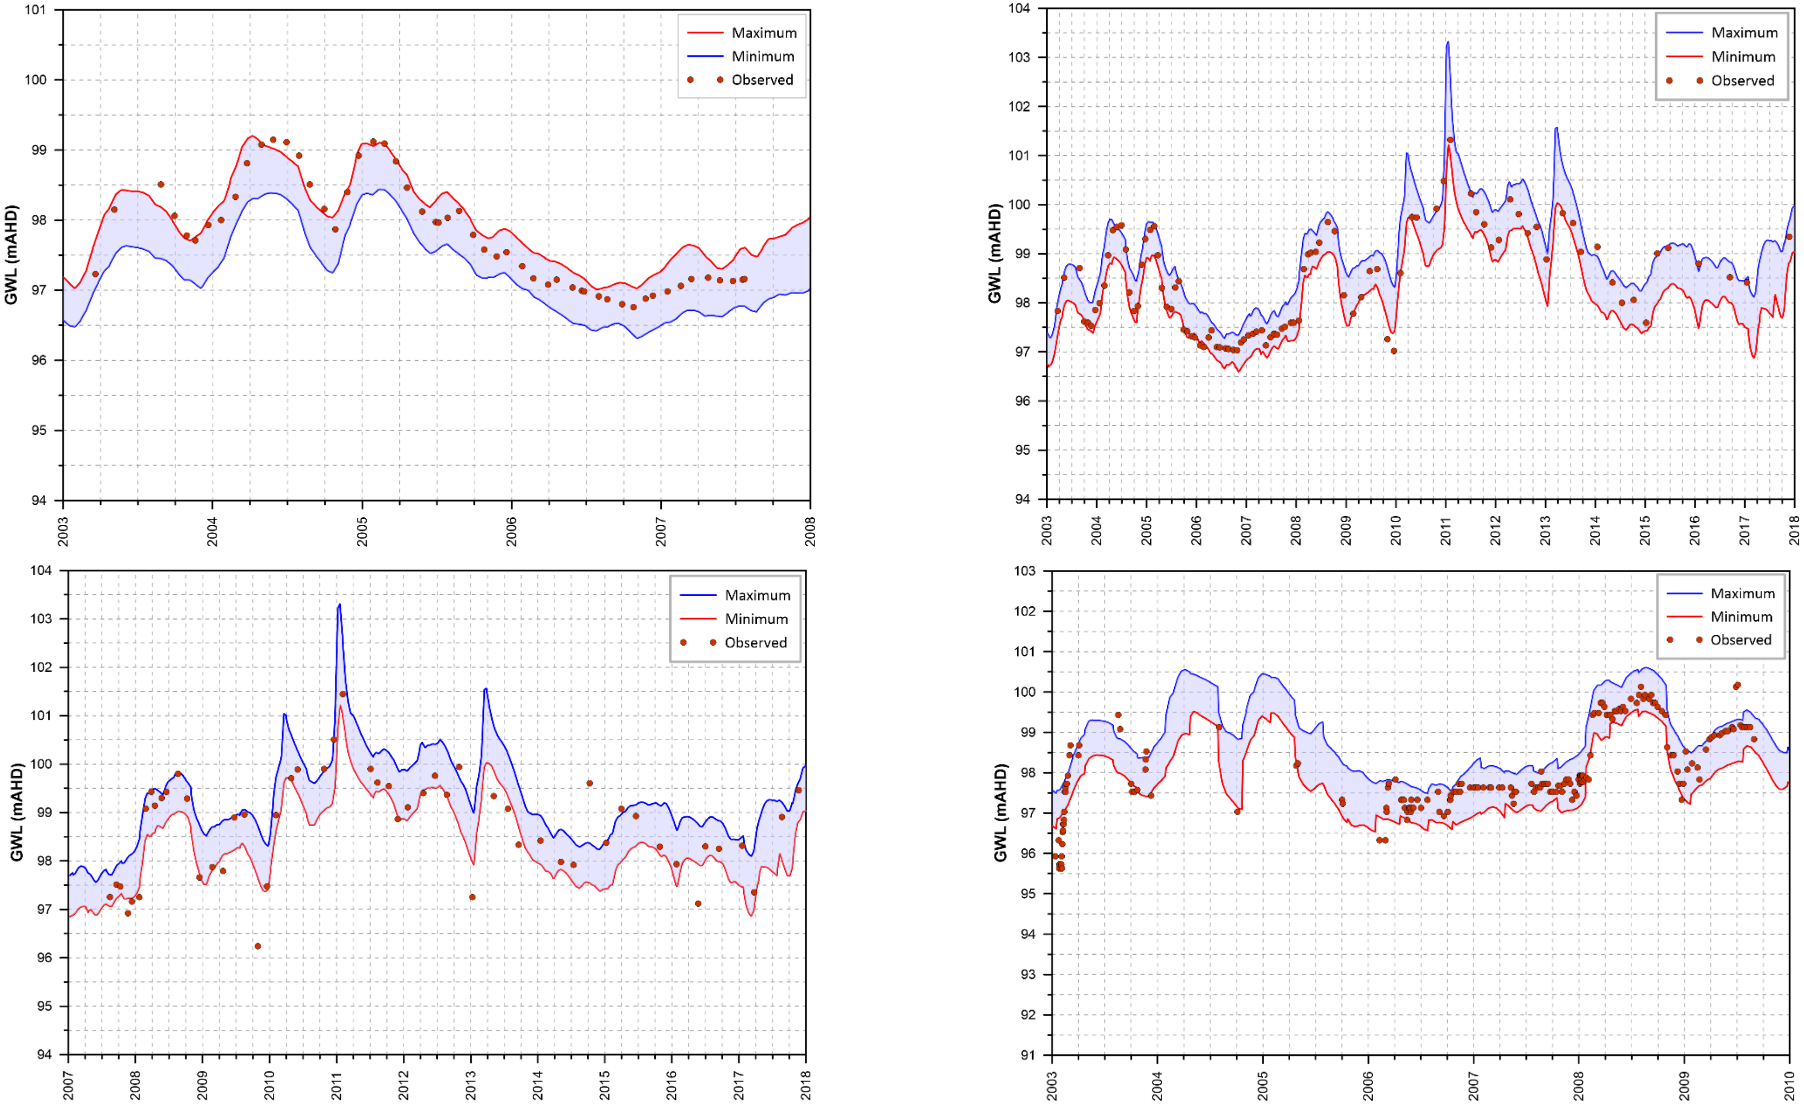

Figure 6.2. Comparison of observed and modelled heads for (a) RN13600186 (top left), (b) RN13600187 (top right), (c) RN13600296 (bottom left) and (d) RN156052 (bottom right) for the best 250 realisations of the ensemble. The envelope of model heads calculated at the end of every MODFLOW 6 stress period is depicted.

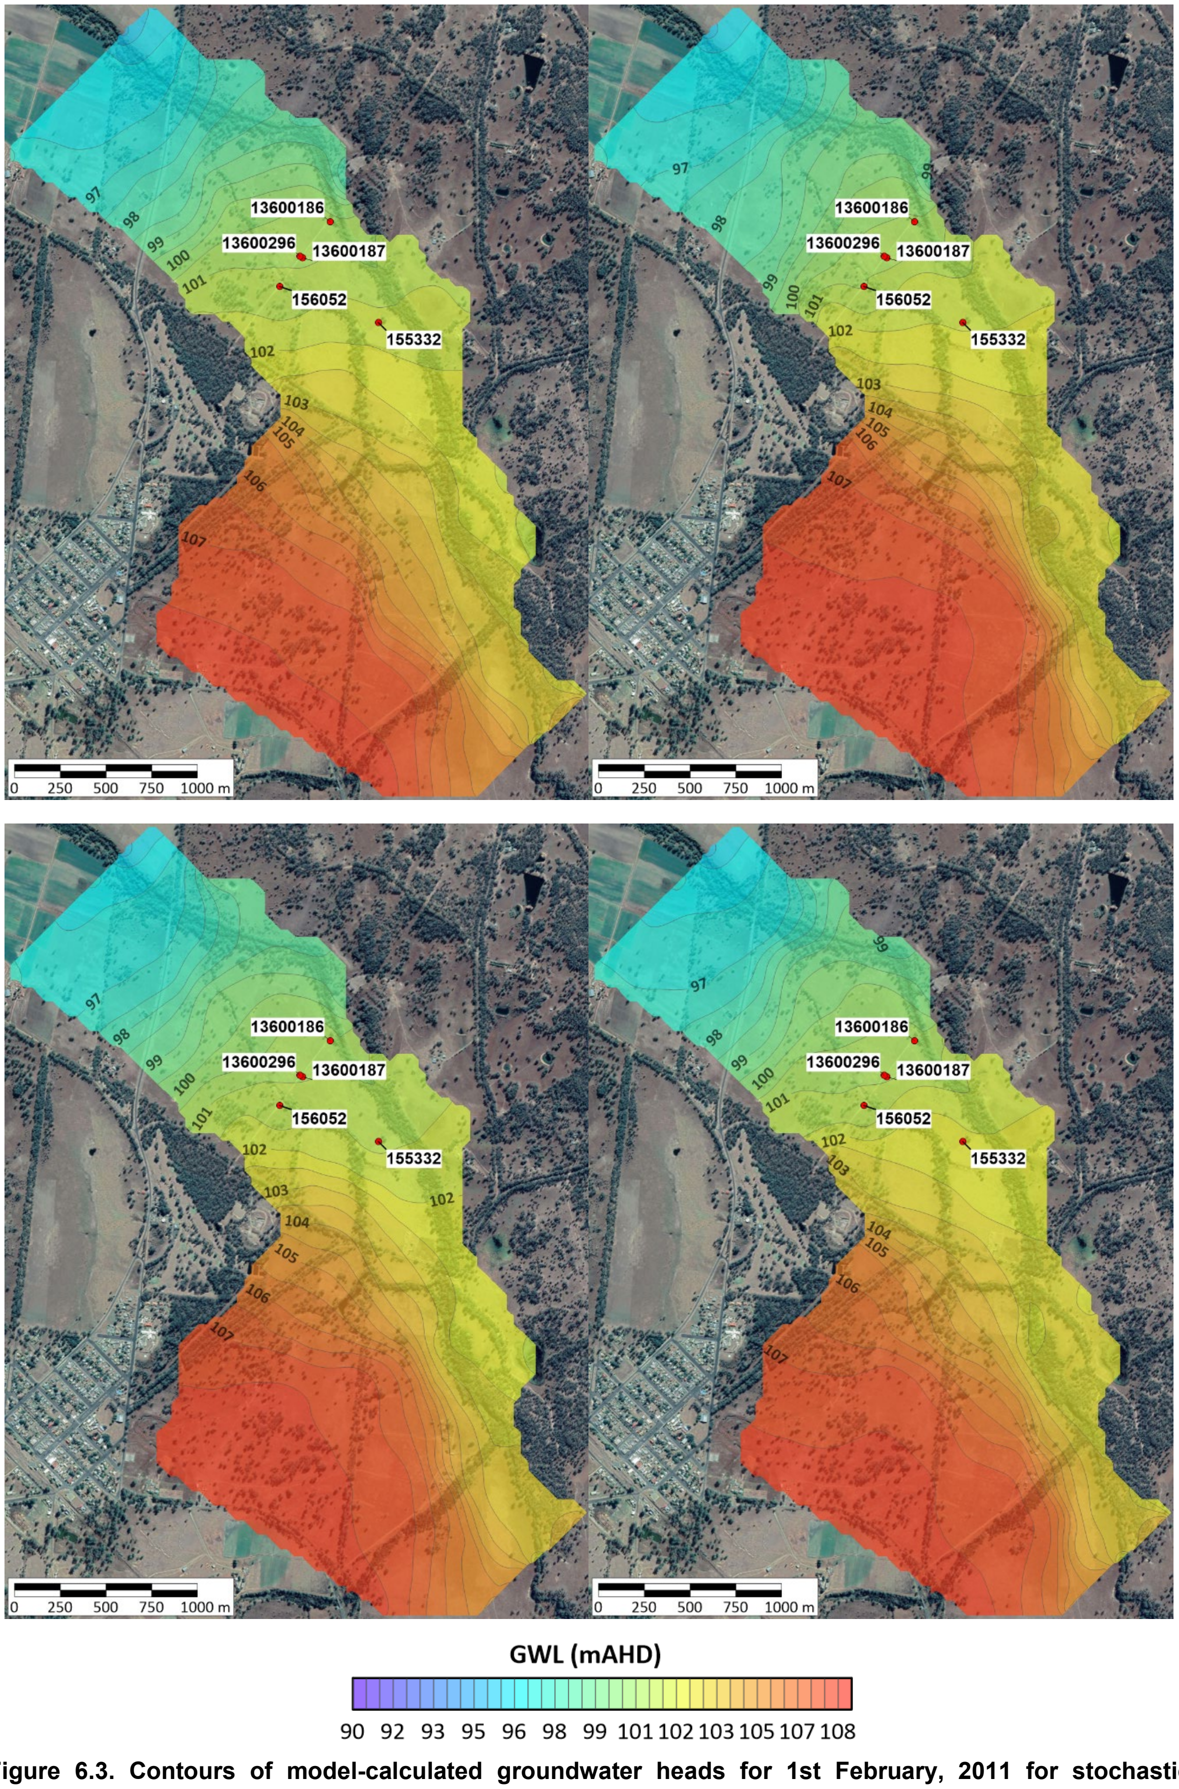

Figure 6.3. Contours of model-calculated groundwater heads for 1st February, 2011 for stochastic realisations 6 (top-left), 173 (top-right), 199 (bottom left) and 213 (bottom right).

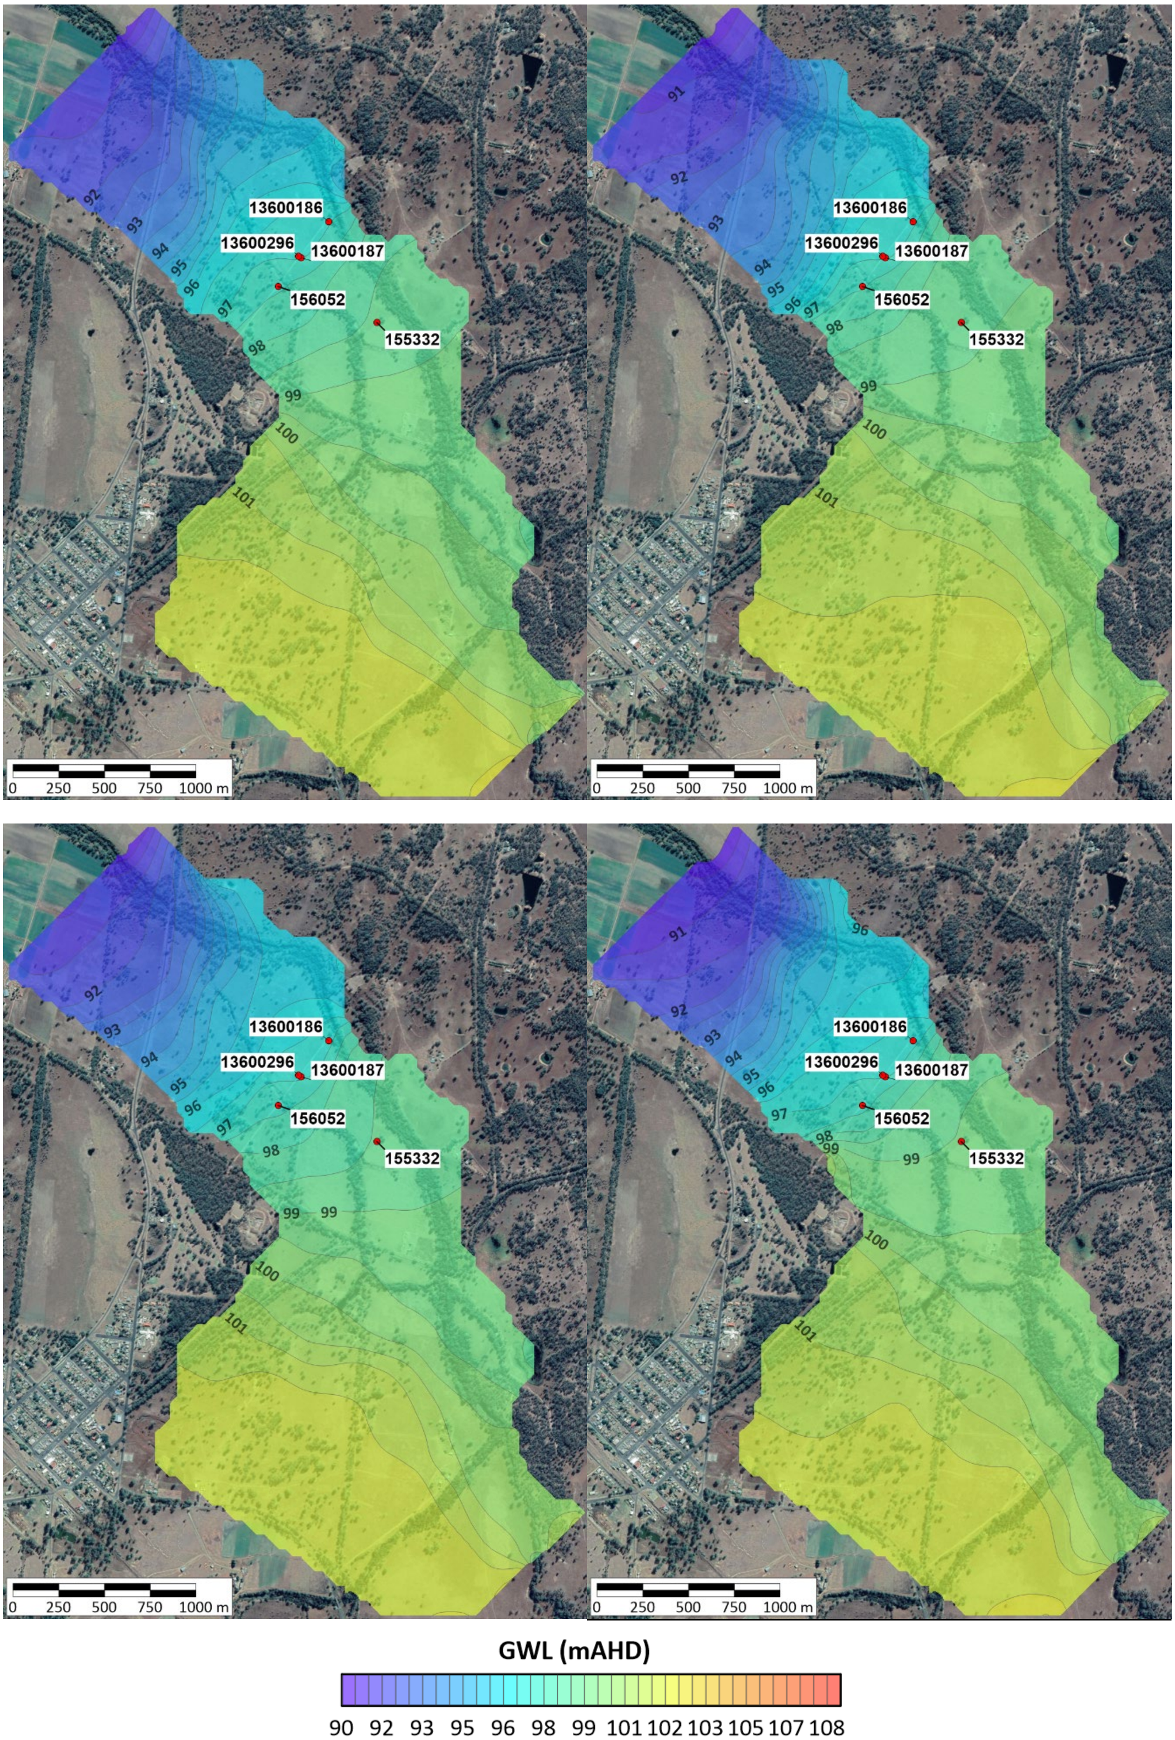

Figure 6.4. Contours of model-calculated groundwater heads for 1st November, 2006 for stochastic realisations 6 (top-left), 173 (top-right), 199 (bottom left) and 213 (bottom right).

6.3.4 Parameter Fields

The four random hydraulic conductivity fields that are used to calculate the piezometric surfaces appearing in Figures 6.3 and 6.4 are shown in Figure 6.5; corresponding specific yield parameter fields are shown in Figure 6.6. These figures are directly comparable with Figures 5.5 and 5.6.

Figure 6.7 shows the spatial distribution of the posterior standard deviation of the best 250 log(Kh) and log(Sy) parameter fields throughout the model domain. It is apparent that the uncertainties of these parameter fields are reduced only in the area where they are constrained by measurements of borehole water levels.

Figures 6.8 and 6.9 depict LUMPREM-calculated boundary heads for the best 250 parameter fields at points about half way along the eastern and western boundaries of the model domain. These are directly comparable with Figures 5.7 and 5.8.

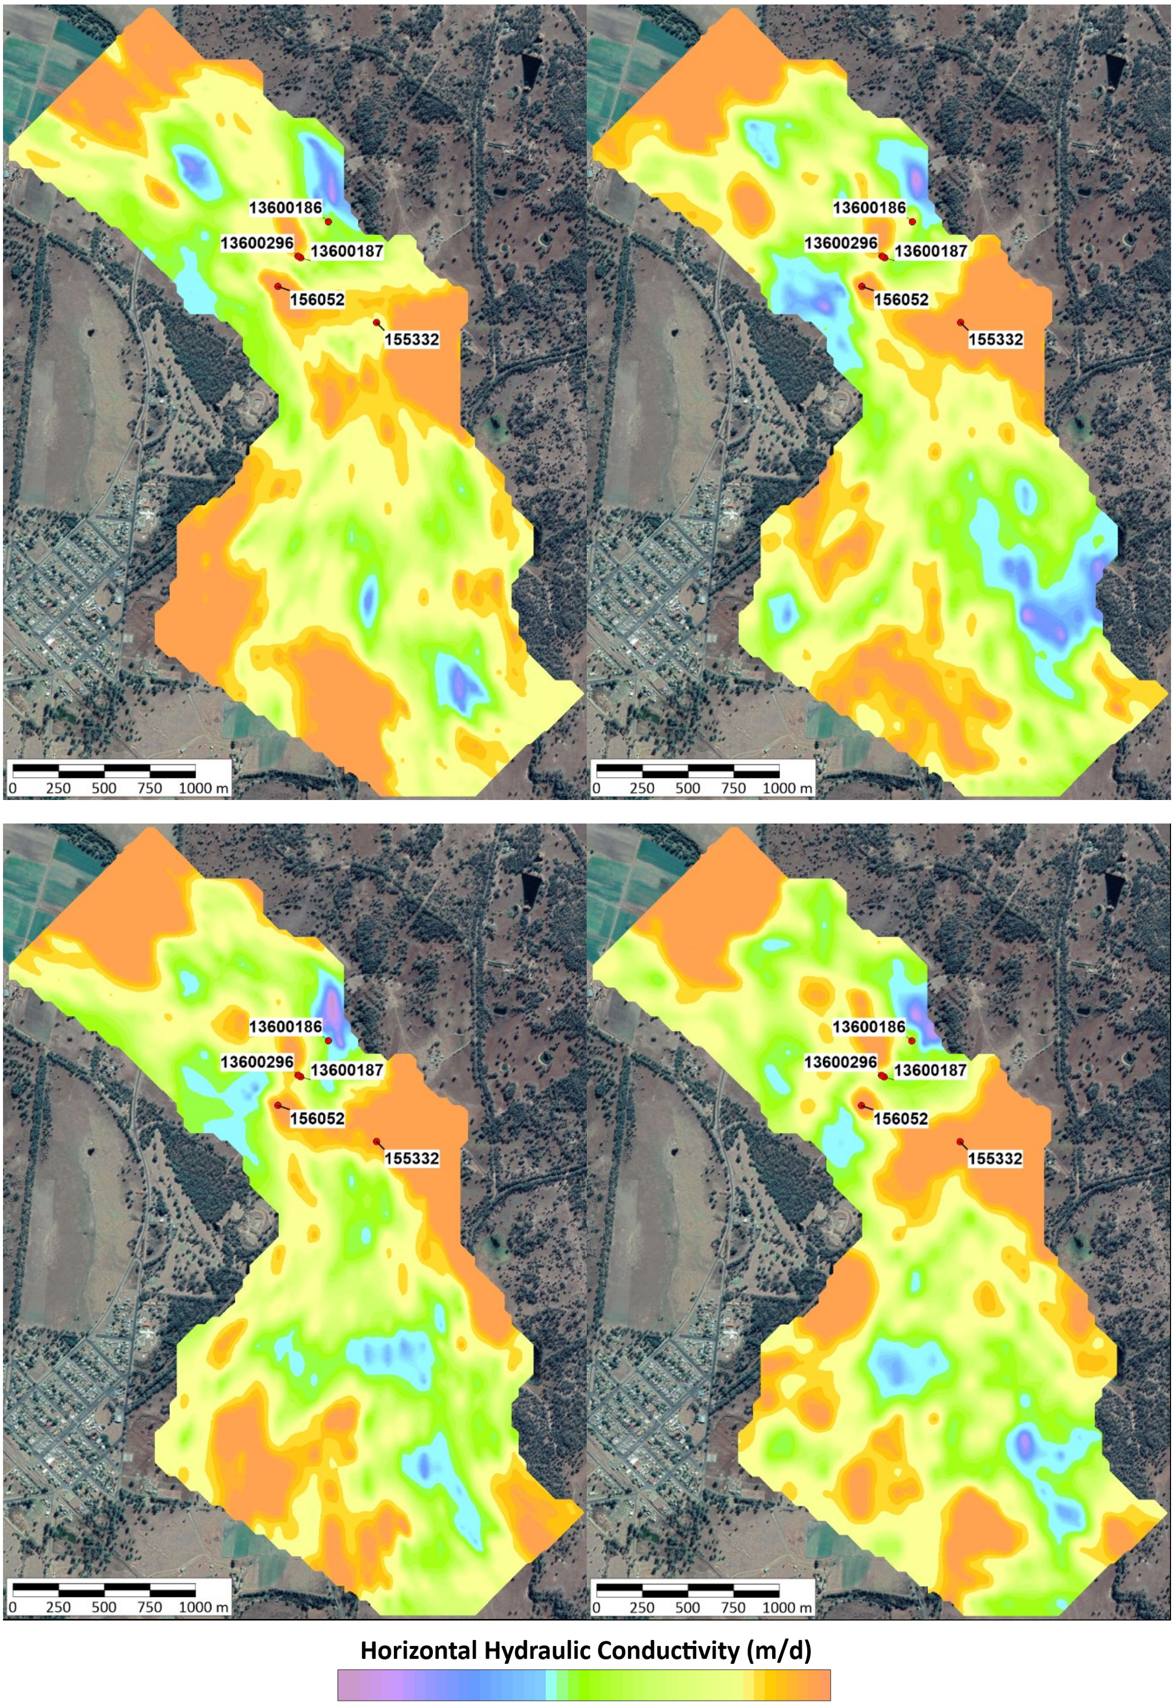

Figure 6.5. Hydraulic conductivity (i.e. Kh) parameter fields for stochastic realisations 6 (top-left), 173 (top-right), 199 (bottom left) and 213 (bottom right).

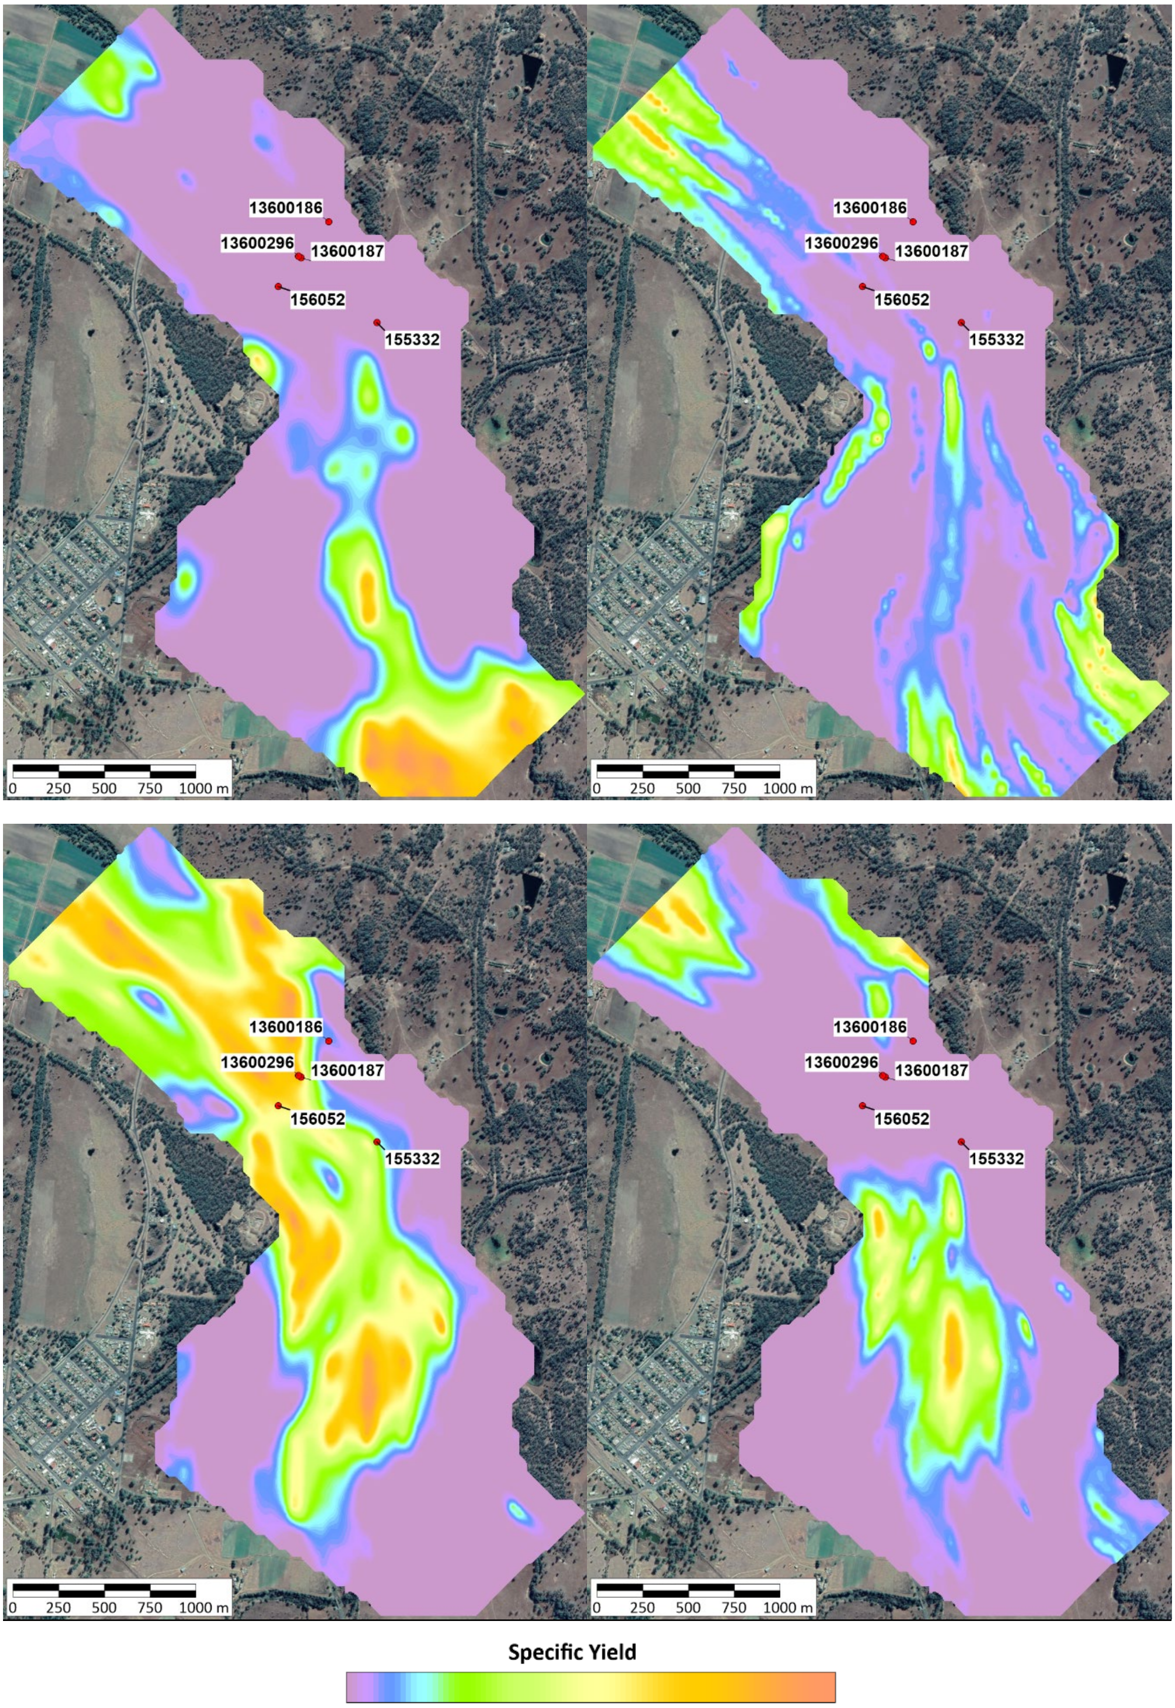

Figure 6.6. Specific yield (i.e. Sy) parameter fields for stochastic realisations 6 (top-left), 173 (top-right), 199 (bottom left) and 213 (bottom right).

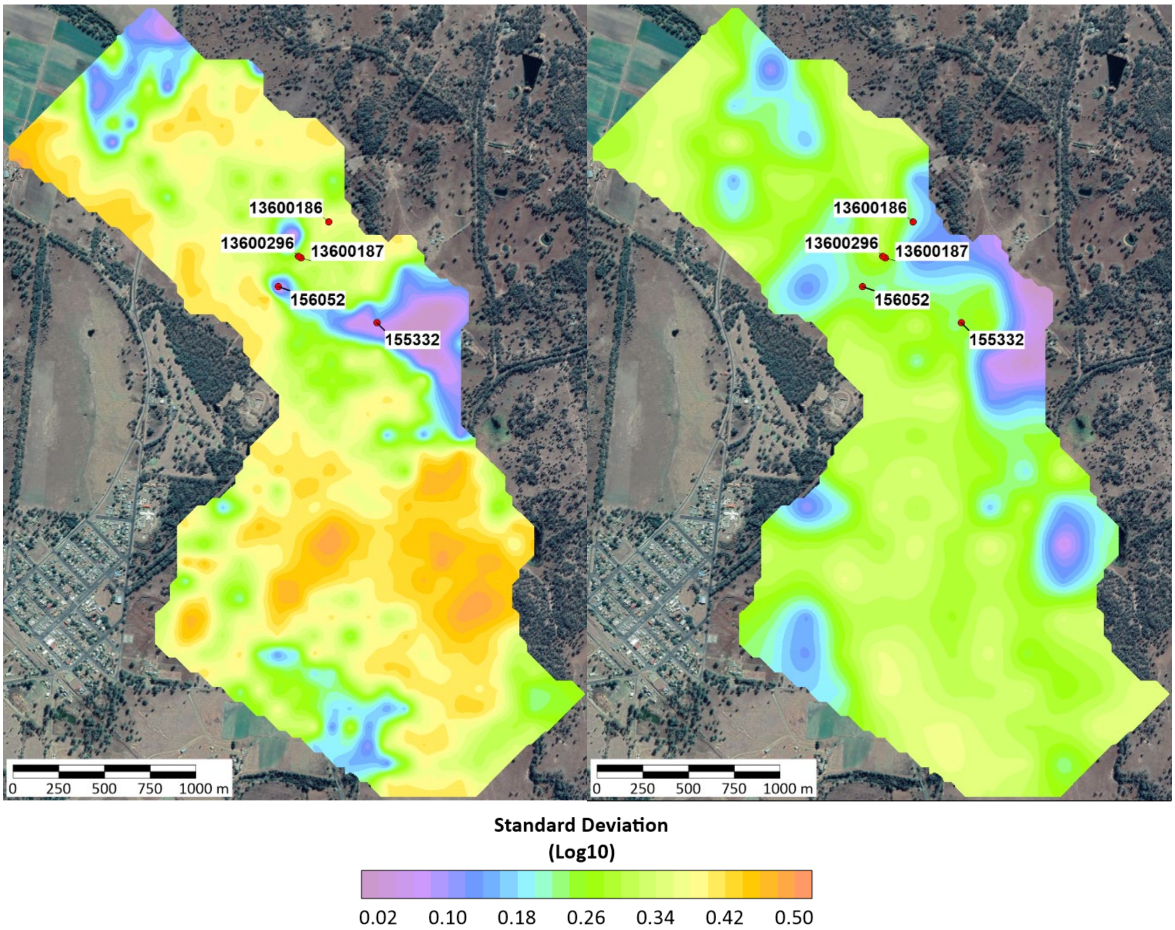

Figure 6.7. Spatial distribution of the posterior standard deviation of the best 250 parameter fields of log(Kh) (left) and log(Sy) (right).

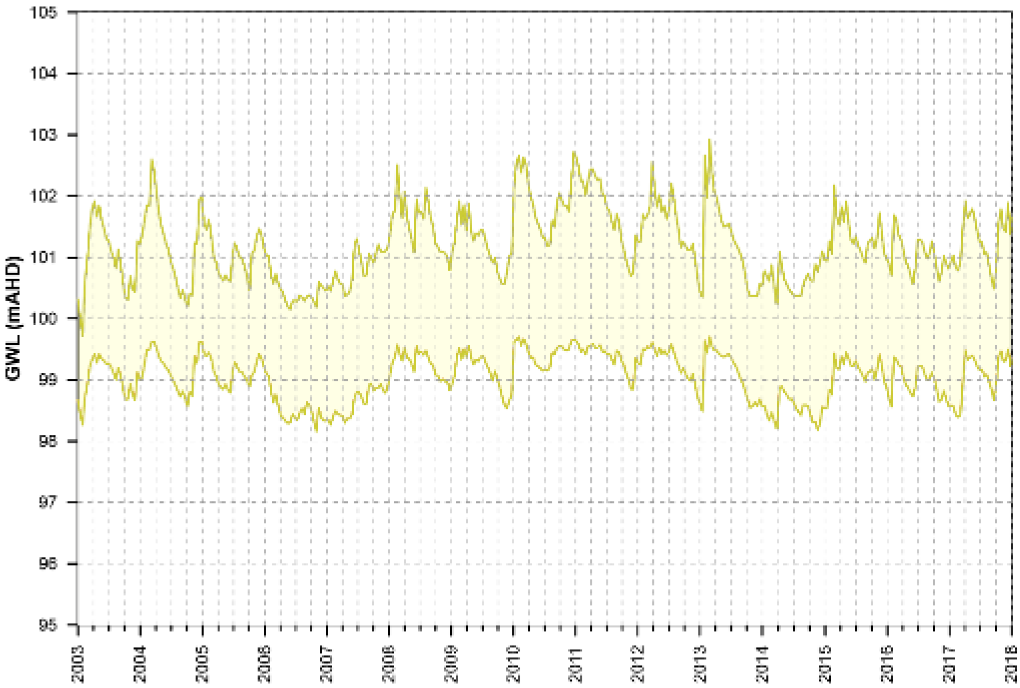

Figures 6.8. LUMPREM-calculated heads for the best 250 parameter fields at a point about half way along the eastern model boundary.

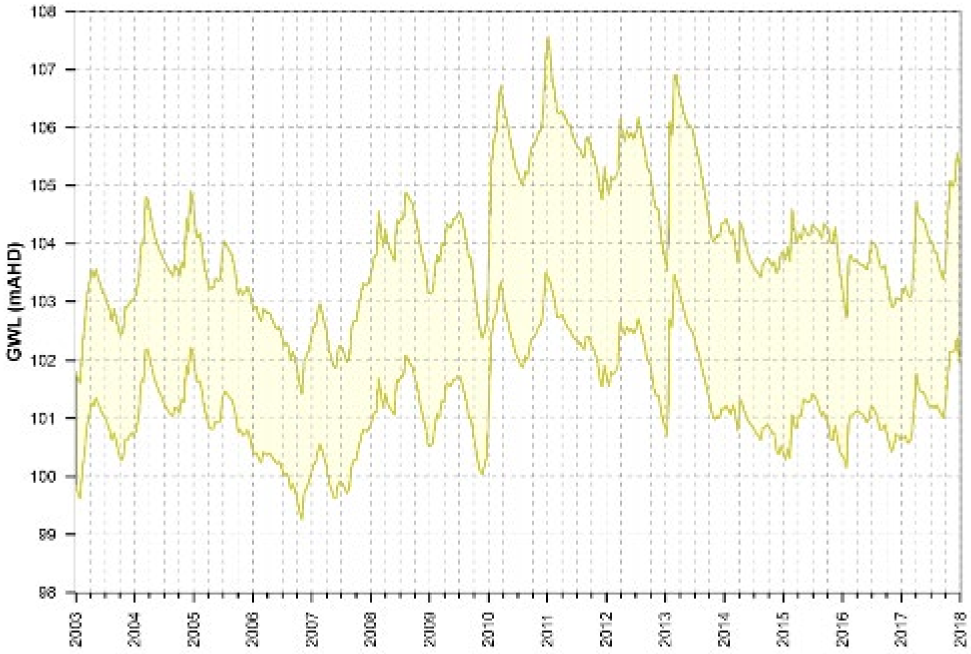

Figures 6.9. LUMPREM-calculated heads for the best 250 parameter fields at a point about half way along the western model boundary.

7. WATER SUPPLY SECURITY ASSESSMENT

In this chapter we illustrate deployment of the model whose construction was described in previous chapters of this report. In doing this, we examine the exceedance probability of a quantity which we refer to as “extraction ratio”. The extraction ratio is 1.0 whenever the water requirements of Biggenden township are met. It drops below 1.0 when extraction rates are reduced in order to prevent water levels in pumping wells from falling below certain thresholds. As was explained in previous chapters, these thresholds can be set in accordance with management-imposed water restrictions. However, regardless of these settings, the model will not allow the water level in a pumped bore to fall below the intake level of the pump. For the deployment example that we describe in the present chapter, it is assumed that no water use restrictions are imposed; hence extraction is reduced only in accordance with the need to safeguard borehole pumps.

In order to calculate exceedance probabilities for various extraction ratios, the model is run for 125 years into the future. As was previously discussed, uncertainties in climate are accommodated through the use of 100 stochastic realisations of future daily rainfall and evaporation. Uncertainties in system properties and behaviour are accommodated through the use of 250 realisations of model parameters calculated by PESTPP-IES.

Three graphs are presented herein.

Figure 7.1 shows extraction ratio exceedance probability as it depends only on parameter uncertainty. This was calculated by running the model using 250 history-match-constrained parameter fields computed by PESTPP-IES, together with a single realisation of future weather.

Four plots are provided in this figure. Two of these pertain to borehole RN155332, while two pertain to both of the bores from which Biggenden extracts its supply, namely RN155332 and RN156052. Recall that 30% of Biggenden’s water needs are provided by the former well while 70% of its needs are provided by the latter well. Extraction ratio exceedance probabilities are presented for two levels of demand, namely 80 Ml/year and 125 Ml/year.

Figure 7.2 shows extraction ratio exceedance probability as it depends only on the uncertainty of future weather. This was computed by running the model 100 times using different realisations of future rainfall and evapotranspiration. The same set of model parameters was used for all of these runs, namely those derived by PEST_HP when it calibrated the model.

Figure 7.3 combines these two sources of uncertainty. It plots the collective outcomes of 250 model runs. Each of these runs is based on a random selection of one of the 100 realisations of future weather together with a random selection of one of the 250 realisations of model parameters.

A comparison of these figures reveals that, for the Biggenden model, uncertainty in future weather contributes more to the uncertainty of water security than uncertainties in properties and behaviour of the groundwater system from which Biggenden draws its water supply. This may have a bearing on how future water supply security assessments are undertaken.

The authors of this report admit to being a little surprised by this conclusion. It indicates that, despite large uncertainties associated with system properties and behaviour, hard and soft data comprising the history-matching dataset contain information that is directly pertinent to predictions of interest. The modelling process was able to effectively assimilate this information.

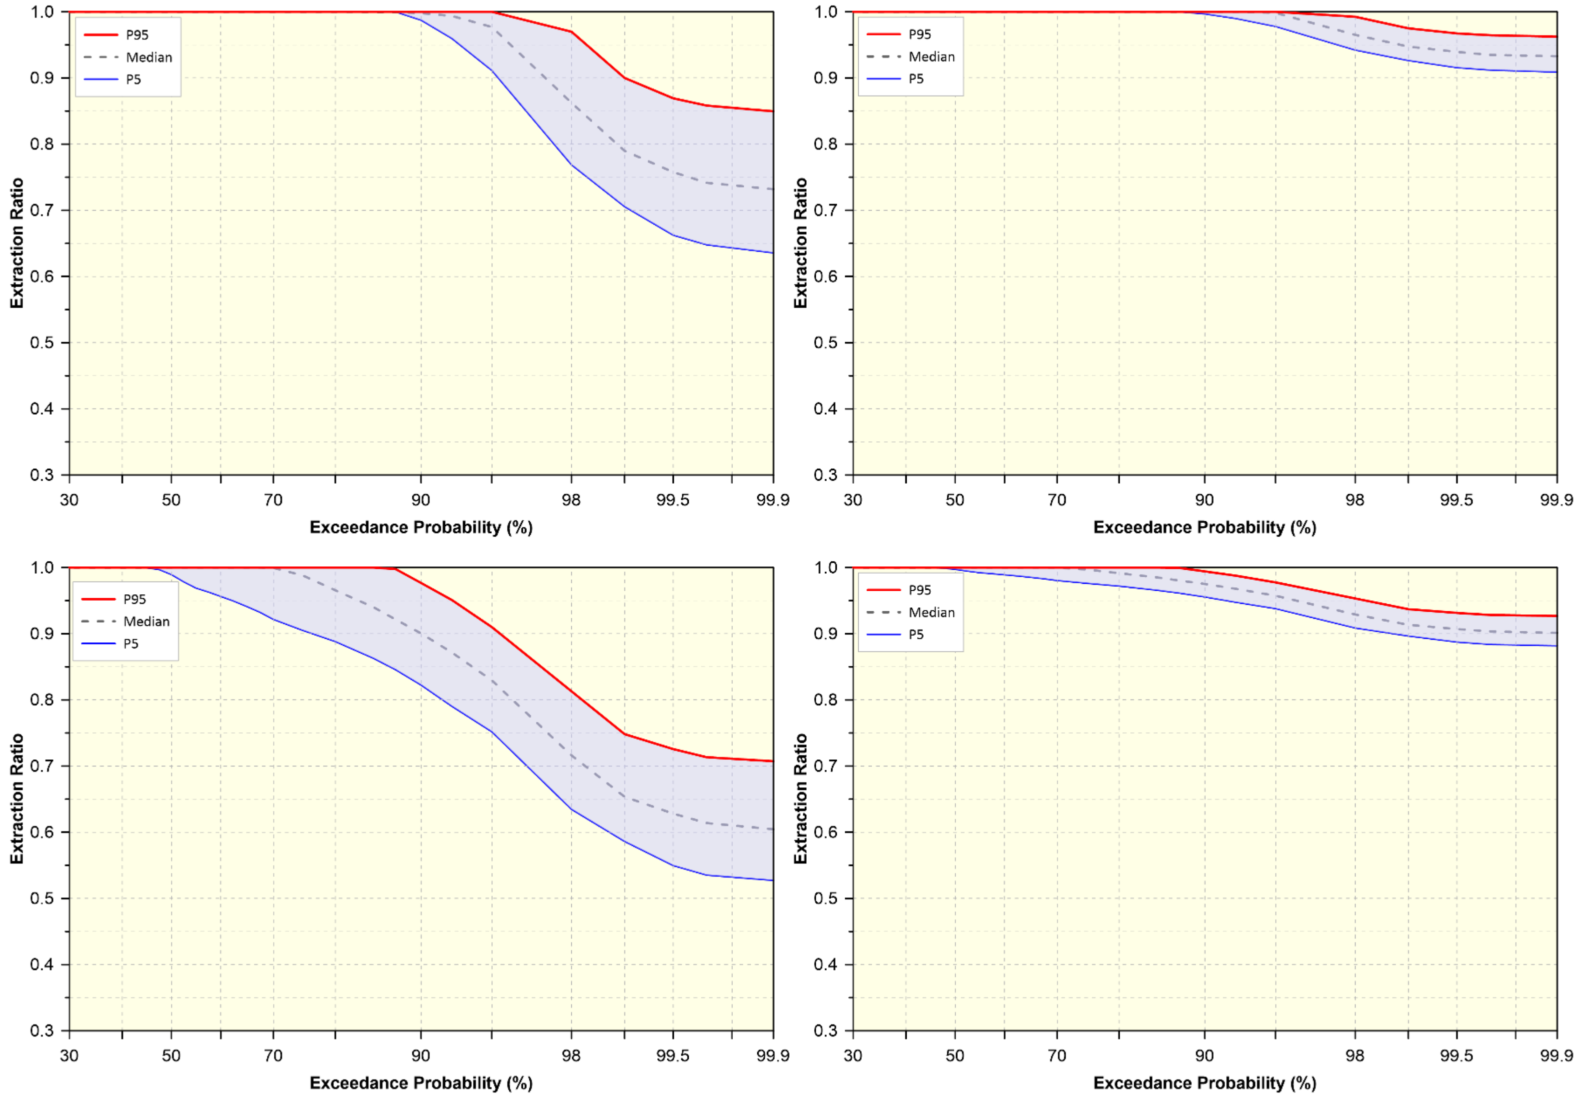

Figure 7.1. Extraction ratio exceedance probability based on parameter uncertainty for RN155332 at 80ML/yr demand (top left), total bore field at 80ML/yr demand (top right), RN155332 at 125ML/yr demand (bottom left) and total bore field at 125ML/yr demand (bottom right).

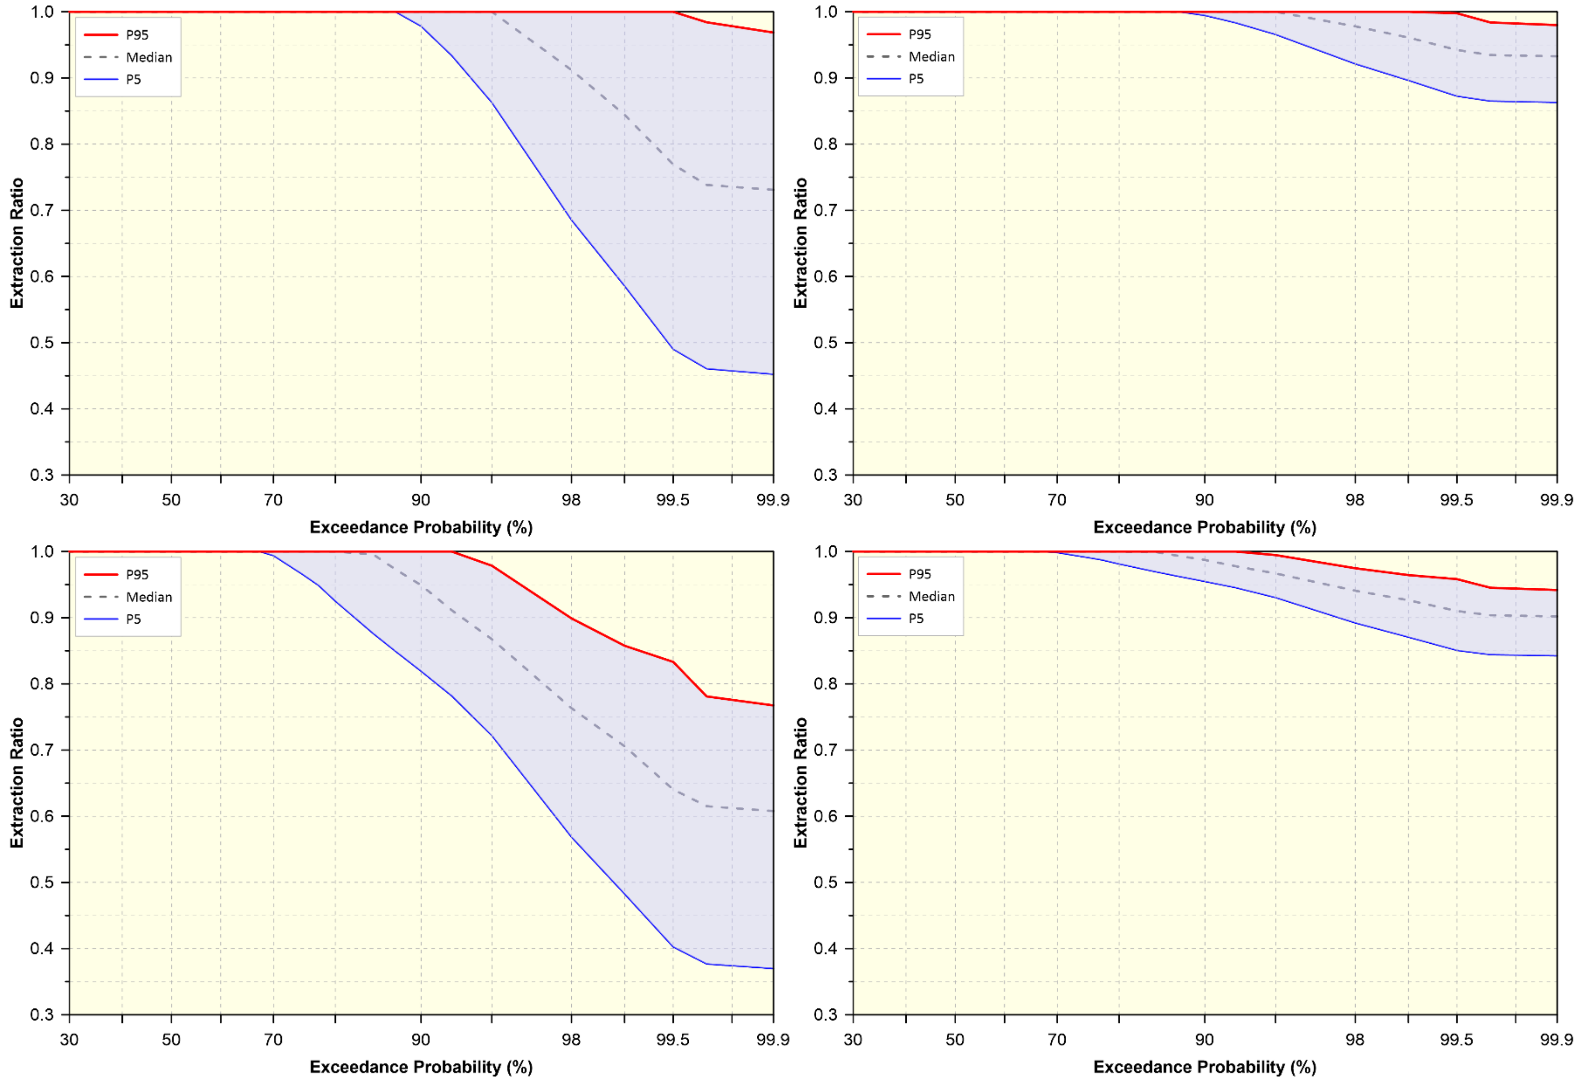

Figure 7.2. Extraction ratio exceedance probability based on the uncertainty of future weather for RN155332 at 80ML/yr demand (top left), total bore field at 80ML/yr demand (top right), RN155332 at 125ML/yr demand (bottom left) and total bore field at 125ML/yr demand (bottom right).

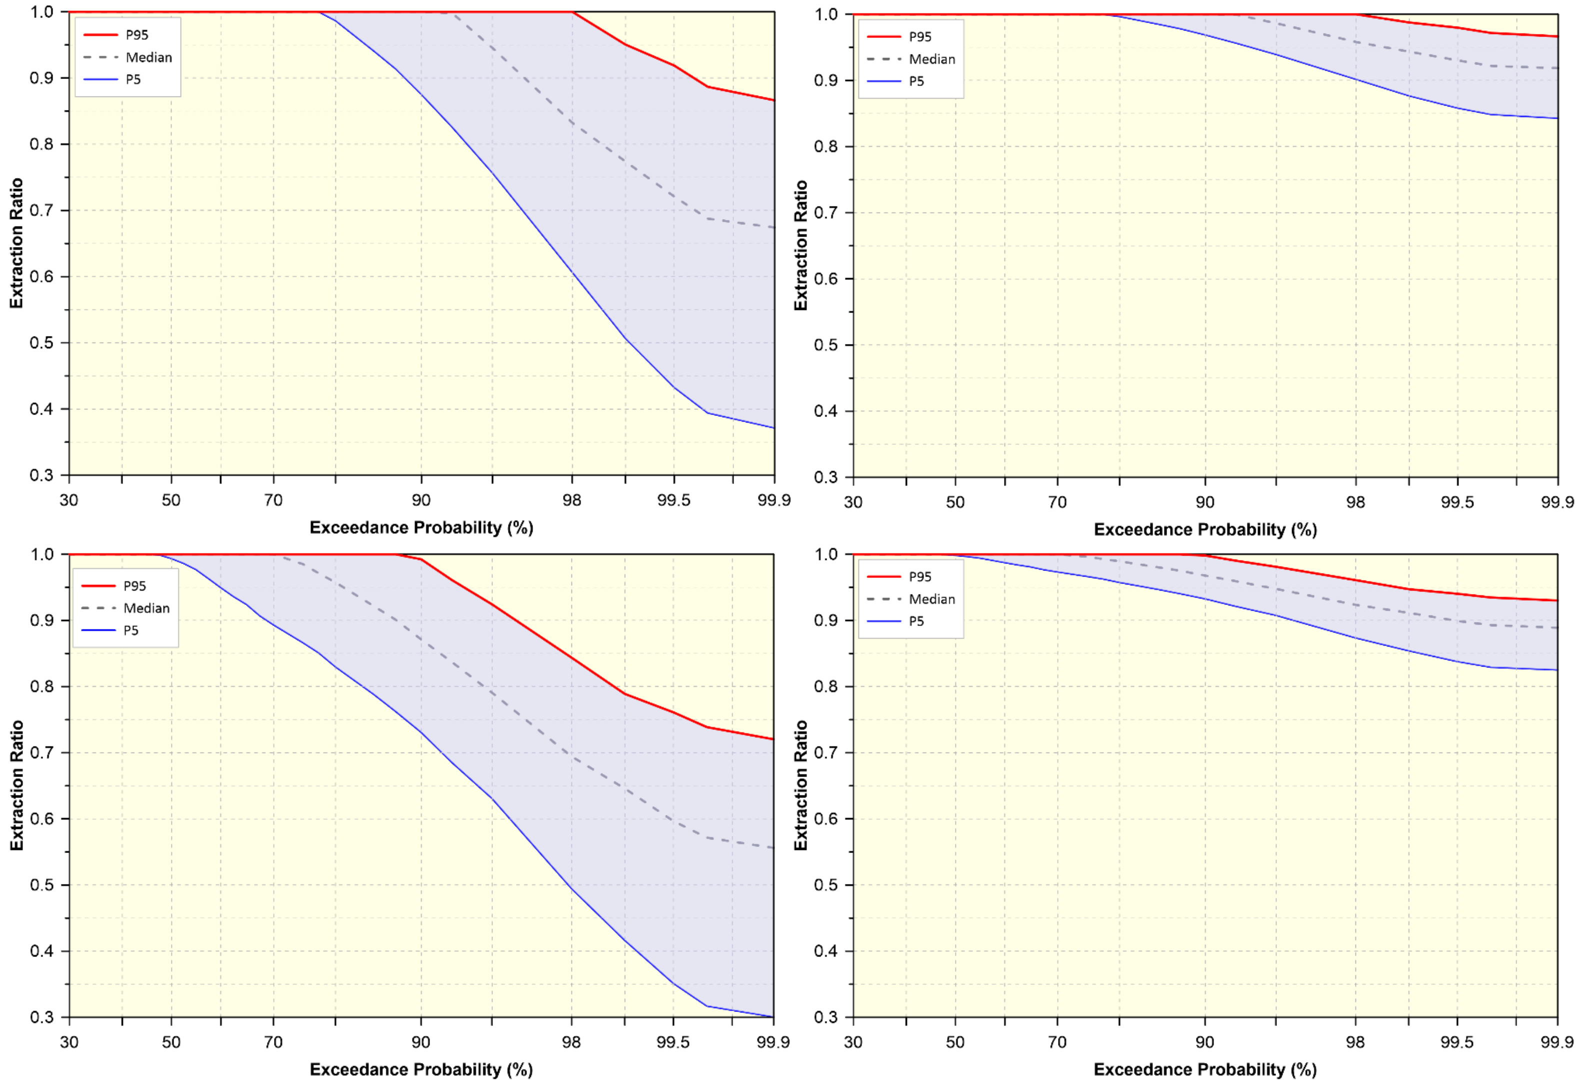

Figure 7.3. Extraction ratio exceedance probability based on both parameter and weather uncertainty for RN155332 at 80ML/yr demand (top left), total bore field at 80ML/yr demand (top right), RN155332 at 125ML/yr demand (bottom left) and total bore field at 125ML/yr demand (bottom right).

8. DISCUSSION

This GMDSI worked example report describes a model that was built to assess water supply security for a small town in south eastern Queensland.

Like any groundwater model, the Biggenden groundwater model has a number of features that make it unique. Data on which its construction is based, and through which its parameters can be inferred, are scarce. Nevertheless, the model respects these data; in doing so, it assimilates the information that they contain. This applies particularly to water level measurements in a small number of bores, including an extraction bore. Maintenance of pumping from extraction bores during dry periods must be assessed by the model.

If the calibration dataset were large, and if it spanned a long enough time interval to ensure that any weather condition that prevails in the future has already been experienced in the past, then assessment of future water supply reliability would not be a difficult undertaking. A simple model, or even a machine learning algorithm, could be trained to calculate borehole water levels from weather sequences. This same algorithm could then be used in conjunction with stochastic sequences of future weather to predict whether borehole water levels are likely to fall below thresholds that restrict water delivery.

In the present study, however, the calibration dataset is relatively small, and spans a relatively short period of time. Predictions required of the model constitute the response of the groundwater system to weather conditions that are not encompassed by this dataset. This requires that the model which makes these predictions has a physical basis. This allows information from other sources to inform its parameters, notwithstanding the fact that this information is qualitative. It includes knowledge of (a) how the Deglibo Creek groundwater system behaves, and (b) materials which host this system. While this information is insufficient to allow definitive estimation of system hydraulic properties, nor to expose the way in which the local alluvial system interacts with the wider groundwater system, it does reduce the range of parameter values that can be used to populate the model. In doing so, it lessens the chances of employing a parameter field that promulgates aberrant behaviour of the model when it encounters possible future weather conditions with which it is unfamiliar.

In assessing water supply security for the township of Biggenden, this worked example has demonstrated how decision-support modelling can express future uncertainties at the same time as it assimilates data whose information content can reduce these uncertainties. A highly-parameterized approach to model construction and history-matching is essential to the achievement of this outcome. It allows free expression of uncertainty at the same time as it provides flexible receptacles for information that resides in hard and soft datasets.

Other facets of model construction, parameterisation and history-matching that have been demonstrated through construction and deployment of the Biggenden groundwater model include the following.

- The LUMPREM recharge model can be used to calculate time series of transient recharge and residual evapotranspiration for use by a groundwater model. It can also be used to calculate time series of boundary heads.
- The inclusion of penalty functions in a history-matching dataset can improve parameter reasonableness, and thereby enhance the performance of a model when making predictions of future system behaviour.
- The PESTPP-IES ensemble smoother provides the means to rapidly sample the posterior probability distribution of model parameters. On some occasions, its performance can be enhanced if its use is preceded by model calibration. Parameter values, and estimates of

posterior uncertainty that are forthcoming from a calibration process, can provide optimal starting realisations for the ensemble inversion process. However the cost (in terms of model runs) of calculating a full rank Jacobian matrix can be high.

9. REFERENCES

Aster, R.C, Borchers, B. and Thurber, C.H., 2013. Parameter Estimation and Inverse Problems. Second edition. Elsevier.

Chen, Y., and Oliver, D.S. 2013. Levenberg–Marquardt forms of the iterative ensemble smoother for efficient history matching and uncertainty quantification. *Computational Geosciences*, 17(4): 689–703.

Doherty, J. 2015. Calibration and uncertainty analysis for complex environmental models. *Watermark Numerical Computing, Brisbane, Australia*.

Doherty, J., 2020. PEST_HP. PEST for highly parallelized computing environments. *Watermark Numerical Computing, Brisbane, Australia*.

Doherty, J., 2020. A simple lumped parameter model for unsaturated zone processes. *Watermark Numerical Computing, Brisbane, Australia*.

King, A.C., Raiber, M., Cox, M.E., and Cendón D.I. 2017. Comparison of groundwater recharge estimation techniques in an alluvial aquifer system with an intermittent/ephemeral stream (Queensland, Australia). *Hydrogeology Journal,* 25(6): 1759–1777.

Langevin, C.D., Hughes, J.D., Banta, E.R., Provost, A.M., Niswonger, R.G., and Panday, Sorab. 2019. *MODFLOW 6 Modular Hydrologic Model version 6.1.0*: U.S. Geological Survey Software Release, 12 December 2019, https://doi.org/10.5066/F76Q1VQV

Meeks, J., Moeck, C., Brunner, P., and Hunkeler, D., 2017. Infiltration under snow cover: modelling approaches and predictive uncertainty. *Journal of Hydrology*, 546, 16-27.

Menke, W., 2018. Geophysical Data Analysis: Discrete Inverse Theory. Third edition. Elsevier.

Peaceman, E.W., 1978. Interpretation of well-block pressures in numerical reservoir simulation. *Soc. Pet. Eng. J.* 183-194; *Trans*., AIME, 253.

PEST++ Development Team (2020). Manual for PEST++. Downloadable from https://github.com/usgs/pestpp

Vitkovsky, J.P. 2008. *NRW_Process Version 1.0*. Water Assessment Group, Department of Natural Resources and Water, Queensland Government, Australia.

Watson, T.A., Doherty. J.E. and Christensen, S., 2013. Parameter and predictive outcomes of model simplification. *Water Resources Research*. 49 3952–3977, doi:10.1002/wrcr.20145.

White, J.T., 2018. A model-independent iterative ensemble smoother for efficient history-matching and uncertainty quantification in very high dimensions. *Environmental Modelling & Software*. 109. 10.1016/j.envsoft.2018.06.009. http://dx.doi.org/10.1016/j.envsoft.2018.06.009.

Wohling, D.L., Leaney, F.W. and Crosbie, R.S., 2012. Deep drainage estimates using multiple linear regression with percent clay content and rainfall. *Hydrol. Earth Syst. Sci*. 16(2): 563-572

Groundwater Modelling Decision Support Initiative

gmdsi.org