UNIVERSITÀ
DEGLI STUDI
DI PADOVA

Sede Amministrativa: Università degli Studi di Padova

DIPARTIMENTO DI MATEMATICA

SCUOLA DI DOTTORATO DI RICERCA IN SCIENZE
DELL'INGEGNERIA CIVILE E AMBIENTALE
CICLO XXV

REDUCED ORDER MODELS AND DATA ASSIMILATION FOR HYDROLOGICAL APPLICATIONS

Direttore della Scuola: Ch.mo Prof. Stefano Lanzoni
Supervisore: Ch.mo Prof. Mario Putti
Co-supervisore: Ch.mo Prof. Alberto Guadagnini
Co-supervisore: Ch.mo Prof. William W-G. Yeh

Dottorando: Damiano Pasetto

Gennaio 2013

Ai miei genitori

Contents

List of Figures

List of Tables

List of Algorithms

Summary

The present thesis work concerns the study of Monte Carlo (MC)-based data assimilation methods applied to the numerical simulation of complex hydrological models with stochastic parameters. The ensemble Kalman filter (EnKF) and the sequential importance resampling (SIR) are implemented in the CATHY model, a solver that couples the subsurface water flow in porous media with the surface water dynamics. A detailed comparison of the results given by the two filters in a synthetic test case highlights the main benefits and drawbacks associated to these techniques. A modification of the SIR update is suggested to improve the performance of the filter in case of small ensemble sizes and small variances of the measurement errors. With this modification, both filters are able to assimilate pressure head and streamflow measurements and correct model errors, such as biased initial and boundary conditions. SIR technique seems to be better suited for the simulations at hand as they do not make use of the Gaussian approximation inherent the EnKF method. Further research is needed, however, to assess the robustness of the particle filters methods in particular to ensure accuracy of the results even when relatively small ensemble sizes are employed. In the second part of the thesis the focus is shifted to reducing the computational burden associated with the construction of the MC realizations (which constitutes the core of the EnKF and SIR). With this goal, we analyze the computational saving associated to the use of reduced order models (RM) for the generation of the ensemble of solutions. The proper orthogonal decomposition (POD) is applied to the linear equations of the groundwater flow in saturated porous media with a randomly distributed recharge and random heterogeneous hydraulic conductivity. Several test cases are used to assess the errors on the ensemble statistics caused by the RM approximation. Particular attention is given to the efficient computation of the principal components that are needed to project the model equations in the reduced space. The greedy algorithm selects the snapshots in the set of the MC realizations in such a way that the final principal components are parameter independent. An innovative residual-based estimation of the error associated to the RM solution is used to assess the precision of

the RM and to stop the iterations of the greedy algorithm. By way of numerical applications in synthetic and real scenarios, we demonstrate that this modified greedy algorithm determines the minimum number of principal components to use in the reduction and, thus, leads to important computational savings.

Sommario

Questo lavoro di tesi riguarda lo studio di tecniche di assimilazione dei dati basate sul metodo di Monte Carlo (MC) per la simulazione numerica di modelli idrologici in presenza di parametri stocastici. I metodi *ensemble Kalman filter* (EnKF) e *sequential importance resampling* (SIR) sono implementati nel modello CATHY, un modello idrologico che accoppia il flusso d'acqua sotterraneo in mezzi porosi con la dinamica del flusso d'acqua superficiale. Il confronto dettagliato dei risultati ottenuti con i due filtri in un caso test sintetico evidenzia i principali vantaggi e inconvenienti associati a queste tecniche. Per migliorare le prestazioni del metodo SIR, in questa tesi è proposta una modifica del passo di update che risulta fondamentale nei casi in cui si usi un ensemble di dimensioni ridotte e la varianza associata all'errore di misura sia piccola. Grazie a questa modifica, entrambi i filtri sono in grado di assimilare misure di carico piezometrico e portata, riducendo la propagazione temporale di errori di modellizzazione dovuti, ad esempio, all'utilizzo di condizioni iniziali o al contorno distorte. La tecnica SIR sembra essere più adeguata dell'EnKF per l'applicazione ai casi test presentati. Si dimostra infatti che l'ipotesi di Gaussianità, che contraddistingue il metodo EnKF, non è soddisfatta in questi casi test, rendendo preferibili metodi più generali come il SIR. Ulteriori approfondimenti sono comunque necessari per stabilire l'affidabilità dei metodi di tipo particle filter, in particolare per garantire l'accuratezza del filtro SIR anche quando viene usato un numero relativamente piccolo di realizzazioni. Siccome il passo di previsione dei metodi SIR ed EnKF è basato sul metodo di MC, la seconda parte della tesi riguarda il problema di ridurre gli onerosi tempi di calcolo associati alla costruzione delle realizzazioni di MC. Con questo obbiettivo, si analizza il risparmio in tempo computazione ottenuto dall'uso di modelli di ordine ridotto (RM) per la generazione dell'ensemble delle soluzioni. La tecnica *proper orthogonal decomposition* (POD) è applicata alle equazioni lineari del flusso d'acqua sotterraneo in mezzi porosi saturi con ricarica stocastica e distribuita spazialmente, oppure con conducibilità idraulica stocastica e descritta per zone. Gli errori di approssimazione introdotti dal modello ridotto sul calcolo delle singole realizzazioni di MC e sulle cor-

rispondenti statistiche sono analizzati in diversi casi test al variare della distribuzione probabilistica dei parametri stocastici. Particolare attenzione è dedicata alla procedura di calcolo delle *principal components* che sono necessarie per la proiezione delle equazioni del modello nello spazio ridotto. Il *greedy algorithm* seleziona gli *snapshots* tra le realizzazioni di MC considerate, facendo in modo che le *principal components* finali siano indipendenti dalla particolare realizzazione dei parametri stocastici. Infine, viene introdotta una stima innovativa della norma dell'errore associato alla soluzione del modello ridotto. Tale stima, basata sul calcolo del residuo, è di fondamentale importanza per stimare la precisione del RM e, quindi, inferire sul numero di *principal components* da usare nella riduzione. Le applicazioni numeriche effettuate su casi test sintetici e reali dimostrano che il *greedy algorithm* così modificato determina un numero minore di *principal components* rispetto al metodo tradizionale, pur mantenendo la medesima accuratezza.

Chapter 1

Introduction

The accurate description of hydrological systems is a crucial topic for modern society, for example to prevent flood events, to assess the transport of contaminants in aquifers, or/and to improve water resources management. In the last decades, the scientific research on the main processes governing the hydrological cycle permitted tremendous improvements in terms of both physical description, governing equation development, and uncertainty assessment. The detailed physical models used nowadays for the simulation of the hydrological processes require an accurate characterization of the spatial and temporal distribution of the model parameters in conjunction with a quantification of system variable uncertainty. The lack of adequate observational techniques of the soil medium properties represents one of the most important issues for the application of these distributed models to real scenarios. It implies the presence of input errors in the initial and boundary conditions, and in the calibration of the model.

An example of these models is the physically-based hydrological model CATHY (Catchment Hydrology) [Camporese et al., 2010]. CATHY is meant to describe the interactions between the surface and subsurface water flow in the simulation of catchment dynamics. Coupling the surface and subsurface flow permits to accurately reproduce several hydrological processes, such as Horton and Dunne runoff generation, ponding, return flow, water reinfiltration, and soil moisture distribution. The application of this model to real hydrological systems requires the knowledge of spatially-distributed soil parameters, such as the saturated hydraulic conductivity, the shape of the retention curves, and the porosity, which are typically unknown on the entire domain.

To take into account these sources of uncertainty, a common approach in forecasting and control theory is to reformulate the problem into a probabilistic framework. In this way, the evolution in time of the probability density function (PDF) associated with the

time-dependent variables becomes the main model unknown. In this probabilistic point of view, data assimilation methods [Jazwinski, 1970] are an important mathematical tool to improve the model forecast. They control the amplification of the uncertainties due to the model nonlinearities and the propagation in time of the PDFs. Data assimilation methods incorporate real system observations (such as pressure head measured at observations wells, soil moisture, streamflow at the outlet) into the hydrological model, in such a way to correct the predicted system state variables toward the measured data, with the possibility of solving inverse problems and quantifying parameter uncertainties (see, e.g., Judd and Stemler [2010] and Margulis et al. [2002]).

The purpose of the first part of this thesis work is to present an experimental comparison between the Ensemble Kalman filter (EnKF) [Evensen, 1994] and sequential importance resampling (SIR) [Gordon et al., 1993] in the context of a detailed process-based hydrological modeling framework. These two Monte Carlo-based data assimilation methods are developed to approximate the solution of the filtering problem in high-dimensional models with nonlinearities and non-Gaussian probability density functions. Camporese et al. [2009a] presented the results of a first application of the EnKF to model CATHY, and demonstrated the great potential of the EnKF to retrieve the correct system state variables when pressure head is assimilated. Nevertheless, an incorrect behavior of the EnKF results when the streamflow is assimilated. This motivated further investigation on suitable data assimilation techniques for the CATHY model.

In Chapter 2, we compare the EnKF and SIR solutions in a synthetic test case. It resembles a rainfall and evaporation events with biased initial and boundary conditions. We give particular attention to the different mathematical assumptions which are at the basis of the two filters, and show the consequences of these assumptions in the numerical results. The Gaussian approximation, which is the core of the EnKF updates, leads to a state estimation that is not consistent with the physics of the model, resulting in a slow-down of the numerical solver. SIR instead duplicates physically-consistent realizations, but, as the state of the realizations is not changed in the updates. SIR can display difficulties in retrieving the correct filtering PDF when the realizations are far from the true state. Moreover, we propose an innovative modification of the SIR algorithm to overcome numerical problems that typically affect the filter performance when the errors on the assimilated measures have small variance. In fact, in this case there is the possibility that the filter misses the update altogether without assimilating the new data. We force an increase of the error variance until the filter can perform the update, and we demonstrated that this modification improves the SIR update [Pasetto et al., 2012a].

The results presented in Chapter 2 show that the main drawback of Monte Carlo (MC) methods is the computational cost associated with the construction of the ensemble. For example, in the CATHY model the propagation in time of each MC realization requires the numerical solution of the highly-nonlinear three-dimensional Richards equation. Moreover, the observational operator, which establishes the relation between the state variables of the dynamical model to the real measurements, in some cases requires the numerical solution of high-dimensional PDEs itself. For example, this is the case of assimilating data obtained by the Electrical resistance tomography (ERT) [Binley and Kemma, 2005]. ERT is a non-invasive technique often employed to detect sub-surface characteristics that can be used to assess the space and time variations of the soil water content. ERT induces an electrical field in the soil and measures the electrical potential differences at the locations of the electrodes. The observational operator for the ERT assimilation consists of two steps: (i) the Archie's law, which relates the water content computed with the hydrological model to the electrical resistivity of the soil; (ii) the solution of a geophysical forward model based on the Maxwell equation to estimate the electric potential at the electrodes. Thus, at each update we must solve an elliptic three-dimensional PDE for each MC realization. As a consequence, application of data assimilation methods to high-dimensional hydrological models is affordable only with small ensemble sizes, thus compromising the accuracy on the forecast and the filter.

To reduce the computational burden associated with MC simulations, in the second part of this thesis work we study the applicability of Galerkin model order reduction techniques (see, e.g., Kunisch and Volkwein [2001]) to groundwater flow problems with stochastic parameters. The idea of model order reduction techniques is to apply a Galerkin projection of the model equations onto a low-dimensional space generated by opportunely chosen basis functions. Then, the fast resolution of the reduced model permits to obtain an approximation of the standard full system model solution. The accuracy of this approximation and the computational efficiency of the reduced model depend on the number and quality of the basis functions selected for the projection, and on the number and properties of the random parameters. For this reason, our main goal is to select parameter-independent basis functions which guarantee the good performance of the method to compute the complete ensemble of MC realizations.

In Chapter 3, we investigate the computational benefits of the reduced model approach in the context of groundwater flow driven by spatially-distributed stochastic recharge. The core of the proposed methodology is the Karhunen-Loève theorem. This theorem identifies the optimal set of basis functions to represent the random field associated to

the solution. The eigenfunction problem used in the Karhunen-Loève theorem is numerically solved with a principal component analysis applied to a set of randomly selected snapshots (full system model solutions). By way of numerical examples, we analyze the dependence between the number of snapshots and principal components used in the reduction with respect to the geostatistical parameters describing the heterogeneity of the distributed recharge, such as the variance and the integral scale. Our results show that the MC solution obtained with the reduced model scheme can decrease significantly the computational burden of a standard MC, while keeping the same degree of accuracy in the mean, covariance and empirical probability density function of the state variables of interest [Pasetto et al., 2011].

A different approach is needed when the random parameter is the hydraulic conductivity. In fact, in this case the nonlinear dependence with the pressure head compromises the efficiency of the reduced model and the computation of the principal components requires a large number of snapshots. As each snapshot corresponds to the solution of a full system model, the introduction of a deterministic procedure to select the optimal snapshots is essential. With this purpose, in Chapter 4 we present a methodology to construct a reduced order model that overcomes the computational burden required solving the transient groundwater flow equation with random hydraulic conductivity. In this methodology, the time- and parameter-independent principal components are computed in an offline algorithm, where the reduced model is iteratively improved until the estimated error on the MC realizations falls below a certain input tolerance. We use a greedy algorithm to select the snapshots in the parameter space [Grepl and Patera, 2005] and to optimally distribute the snapshots in time [Siade et al., 2010]. The main innovative aspect in this procedure is that we introduce a new residual-based estimation of the error associated with the reduced model solution. In particular, we compute the exact relation between error and residual in correspondence of the snapshots and extend this relation to the other realizations by interpolation in the parameter space. This estimation constitutes the main condition for the termination of the greedy algorithm and allows to considerably reduce the number of full system model solutions required to compute the principal components, thus leading to important computational savings. Numerical results on one- and two-dimensional scenarios show that the principal components obtained with our methodology ensure a good approximation of the whole ensemble of realizations. Consequently, the ensemble mean and variance of the head computed with the full and the reduced model solution are qualitatively the same. A necessary condition for the applicability of this methodology to MC-based data assimilation techniques is that the empirical proba-

bility density functions at the observation wells are not changed by the reduced model. For this reason, we compare the empirical probability distributions obtained from the reduced model with the correct distributions. The two-sided Kolmogorov-Smirnov test demonstrates that our reduced model is more accurate in the regions with large drawdown [Pasetto et al., 2012b]. The computational advantages of using the reduced model are evident in the real system application. In fact, the reduced model solution is 1000 times faster than the full system model solution.

Chapter 2

Data assimilation methods for a physically-based coupled surface-subsurface model: ensemble Kalman filter versus particle filters [1]

2.1 Summary

The Ensemble Kalman filter (EnKF) and sequential importance resampling (SIR) are two Monte Carlo-based sequential data assimilation (DA) methods developed to solve the filtering problem in nonlinear systems. Both methods present drawbacks when applied to physically-based nonlinear models: the EnKF update is affected by the inherent Gaussian approximation, while SIR may require a large number of Monte Carlo realizations to ensure consistent updates. In this work we implemented EnKF and SIR into a physically-based coupled surface-subsurface flow model and applied it to a synthetic test case that considers a uniform soil v-shaped catchment subject to rainfall and evaporation events. After a sensitivity analysis on the number of Monte Carlo realizations and the correlation time of the atmospheric forcing, the comparison between the two filters is done on the basis of different simulation scenarios varying observations (outlet streamflow and/or pressure head), assimilation frequency, and type of bias (atmospheric forcing or initial conditions). The results demonstrate that both EnKF and SIR are suitable DA methods for detailed physically-based hydrological modeling using the same, relatively small, en-

---

semble size. We highlight that the Gaussian approximation in the EnKF updates leads
to a state estimation that can be not consistent with the physics of the model, resulting
in a slowdown of the numerical solver. SIR instead duplicates physically-consistent real-
izations, but can display difficulties in updates when the realizations are far from the true
state. We propose and test a modification of the SIR algorithm to overcome this issue
and preserve assimilation efficiency.

2.2   Introduction

It is widely recognized that problems in catchment hydrology and water resource manage-
ment involve strong interactions between surface water and groundwater [Furman, 2008].
A number of physically-based hydrological models that incorporate some representation
of groundwater-surface water interactions have been recently developed [e.g., Camporese
et al., 2010, Kollet and Maxwell, 2006, Morita and Yen, 2002, Panday and Huyakorn, 2004,
Qu and Duffy, 2007, Rigon et al., 2006, Shen and Phanikumar, 2010, Sulis et al., 2010,
VanderKwaak and Sudicky, 2000, Weill et al., 2009]. The objective of these models is to
accurately reproduce a number of hydrological processes, such as rainfall partitioning be-
tween runoff and infiltration, soil moisture redistribution, and groundwater recharge. This
is achieved by solving various equations for the description of surface runoff and variably
saturated flow in porous media. The uncertainties associated to calibrated model param-
eters, initial conditions, and boundary conditions induce model errors that can propagate
in time, yielding final predictions drastically different from reality. For this reason, a
suitable approach to the problem of simulating hydrological processes is to consider them
as realizations of a stochastic process, searching for the most probable one or even for the
full probability distribution of the physical state of the system. Field observations provide
important but noisy information of the true system state. These measurements are nowa-
days often available, at least in the form of discharge at the outlet. In this framework,
data assimilation (DA) methods are techniques that try to combine these measurements
with the simulation model of the system in order to improve accuracy and, most impor-
tantly, to quantify uncertainties (i.e., the probability distribution of the errors) of model
predictions [McLaughlin, 2002, Paniconi et al., 2003].

   Most of the DA techniques are based on the filtering theory [Jazwinski, 1970], in which
the problem unknown is the evolution in time of the probability density function (PDF) of
the system state conditioned to all the previous observations. This PDF is called filtering
PDF and is the solution of a recursive formula that consists of two steps: i) the forecast

step, in which the PDF is propagated in time using the model dynamic equations; ii)
the analysis step, in which the forecast is updated by Bayes theorem using the newly
acquired observations. This approach can be addressed analytically only in few simple
cases. For example, if the model is linear with Gaussian additive noise, filtering PDFs are
Gaussian and the first two moments (mean and covariance) completely characterize the
density function. The widely known Kalman filter Kalman [1960] provides in this case
the full solution of the filtering problem. Several approaches are available for dealing with
nonlinear dynamics or observation relationships

The extended Kalman Filter (EKF) is based on the linearization of the system opera-
tors and on the Gaussian approximation, according to which the filtering probabilities are
again assumed to be Gaussian. The Chapman-Kolmogorov equation is solved linearizing
the system equations along a reference state trajectory based on the previous state esti-
mate. The state covariance matrix is propagated in time extending the KF equations to
the linearized model and it is then used in the analysis step by computing the classical
Kalman gain. Entekhabi et al. [1994] and Hoeben and Troch [2000] demonstrated the
potential of EKF for estimating soil moisture profiles using sequential assimilation of re-
motely sensed surface moisture data in a one-dimensional modeling context, while Walker
et al. [2001] compared direct insertion and the EKF using synthetic data, concluding that
the Kalman filter-based assimilation scheme is superior to the direct insertion method.
The better performance of the KF-based algorithms with respect to simpler assimilation
techniques is a consequence of their ability to adjust the system state over larger portions
of space, as shown also by Camporese et al. [2009a]. There are essentially three drawbacks
in using the EKF. First, the propagation in time of the system state covariance matrix
can be difficult (in hydrological models this has been shown, for instance, by Van Geer
et al. [1991] and Drecourt et al. [2006]) and become impractical for large three-dimensional
models [Evensen, 2007]. The second problem arises in the case of highly-nonlinear equa-
tions, when linearization along the reference state is not sufficient to accurately describe
the complex system dynamics, leading to stability limitations of the filter [Kushner, 1967].
Finally, the third problem is inherent in the use of the Gaussian approximation: if the
true density is not Gaussian, e.g., if it is bimodal, the use of a Gaussian PDF may not
represent an acceptable approximation and the filter can even diverge [Arulampalam and
Ristic, 2000].

To overcome the limitations of EKF, Evensen [1994] introduced the Ensemble Kalman
Filter (EnKF) in the context of ocean modeling. Nowadays EnKF is widely used in
different fields of application, including hydrological modeling, mainly because of its ease

of implementation and computational efficiency. EnKF uses a Monte Carlo approach to evaluate the state mean and covariance matrix and performs an update step based on the Kalman filter. This method has been applied extensively to hydrological models, including, e.g., the one-dimensional Richards equation [Das and Mohanty, 2006], three-dimensional saturated groundwater flow [Chen and Zhang, 2006] and transport [Liu et al., 2008], integral-balance saturated-unsaturated subsurface models [Shu et al., 2005], and conceptual rainfall-runoff models [Aubert et al., 2003, Clark et al., 2008, Pauwels and De Lannoy, 2009, Weerts and El Serafy, 2006].

Although EnKF has given accurate results in many applications, a number of issues remain unresolved. An interesting question to ask is what is the influence of the Gaussian approximation, inherent in any KF based method, when the filtering density function is not Gaussian. Another interesting and partially unanswered question is related to the optimality of the scheme when using KF-based algorithms for nonlinear processes [Margulis et al., 2002, Zhou et al., 2006]. Several modifications to the EnKF and alternative DA techniques have been presented in the literature to handle non-Gaussian distributions. For example, Sun et al. [2009] compares four different deterministic ensemble filters for the estimation of the hydraulic conductivity, while Zhou et al. [2011] and Schöniger et al. [2012] propose two versions of the EnKF with transformed data, in which ensemble realizations are transformed into a Gaussian distribution before the update.

Instead of looking at variations of EnKF that address non-Gaussianity, in this thesis we focus on the sequential importance resampling (SIR) particle filter, a DA method developed specifically to avoid the Gaussian assumption [Gordon et al., 1993, West and Harrison, 1997]. Particle filters (PF) are very flexible and easy to implement, an appealing characteristic that fostered their use in many research areas, such as object recognition, target tracking, signal processing, financial analysis, and robotics. For a comprehensive description of the principal PF methods see, for instance, Arulampalam et al. [2002] and Doucet et al. [2000].

The basic idea of PF is to use directly the Bayesian formula for the computation of the filtering PDF at each analysis step. The resulting update is similar to the static GLUE (generalized likelihood uncertainty estimation [Rojas et al., 2008]) and Bootstrap filters [Smith and Gelfand, 1992]. In SIR the state PDF is approximated via Monte Carlo by associating to each realization (also called particle) appropriate weights, which are updated using the likelihood function of the assimilated observations. A resampling step is used to propagate through the model dynamic equations only the particles that correspond to non-negligible weights. The weights can be thought of as a measure of the

"nearness" of each particle to the true system state. Theoretically, it can be shown that, without any assumption on the relevant PDFs, the empirical probability distribution associated to SIR converges to the filtering density when the ensemble size tends to infinity [Crisan, 2001].

PF methods were used to address parameter uncertainty and state estimation in conceptual hydrological models [Moradkhani et al., 2005, Salamon and Feyen, 2009] and in a coupled hydrogeophysical model applied to an infiltration experiment [Rings et al., 2009]. Noh et al. [2011] introduced a lagged regularized PF for a process-based distributed hydrological model, while Van Delft et al. [2009] combined PF and EnKF to solve a flood forecasting problem. PF had been also used to account for initial condition uncertainties in the framework of ensemble streamflow prediction [DeChant and Moradkhani, 2011a], to assimilate remote sensing derived water levels in a one-dimensional hydraulic model [Giustarini et al., 2011], and to assimilate satellite soil moisture measurements in a one-dimensional mechanistic soil water model [Montzka et al., 2011].

The purpose of this work is to present an experimental comparison between EnKF and SIR in the context of a detailed process-based hydrological modeling framework. Comparison between these two methods are already present in the literature related to hydrological applications, but are mostly restricted to conceptual models. Weerts and El Serafy [2006] compared the EnKF and PF in a conceptual rainfall-runoff model, finding that, for their application, EnKF outperforms SIR for low flows. A comparison on a more complex land surface model was done by Zhou et al. [2006], demonstrating that EnKF can well approximate the optimal filtering PDF also in case of non-normal moisture behavior. DeChant and Moradkhani [2011b] concluded that, using a large number of realizations, SIR is better than EnKF when assimilating microwave radiance data, improving predictions of the snow water equivalent and operational streamflow forecast.

Following the study presented in Camporese et al. [2009b], we want to assess the relative performance of EnKF and SIR as implemented in a coupled model of surface-subsurface water interactions. A further objective of the thesis is to focus on implementation issues related to the assessment of the physical, statistical, and numerical consistency of the state variables within the model dynamics during the entire assimilation process, topics that rarely have been considered in previous data assimilation studies, at least in hydrological models. Finally, we want to discuss the advantages and drawbacks of EnKF and SIR in hydrological applications.

The chapeter is organized as follows. In Section 2.3 we summarize the main aspects of the filtering theory to highlight the theoretical differences between SIR and EnKF.

Then, in Section 2.4 we describe how these methods are combined with our hydrological model, emphasizing the most correct way to initialize and propagate the ensemble of state vectors in order to obtain physically-consistent realizations and to improve the accuracy of the system state estimation. In Section 2.5 we analyze the behavior and performance of the filters in a series of numerical experiments involving a three-dimensional synthetic v-shaped catchment for which a rainfall-drainage simulation is carried out. The two methods are evaluated in terms of their ability to retrieve the correct watershed response and compared for different assimilation frequencies, different scenarios of state variable observations (streamflow only, pressure head only, both), and for biased initial conditions or biased atmospheric forcing. Section 2.6 reports our main conclusions.

2.3   Data assimilation methods

2.3.1   Problem setting

Let $\mathbf{x}_t^{true} \in \mathbb{R}^n$ and $\mathbf{y}_t^{obs} \in \mathbb{R}^m$ represent the vectors of the true system state variables and of the observations at a given assimilation time t, respectively. Given an initial state $\mathbf{x}_0^{true}$, the state $\mathbf{x}_t^{true}$ and the observations $\mathbf{y}_t^{obs}$ satisfy two generic equations of the form:

$$\mathbf{x}_t^{true} = \mathcal{F}_t(\mathbf{x}_{t-1}^{true}, \mathbf{w}_t) \ , \tag{2.1}$$

$$\mathbf{y}_t^{obs} = \mathcal{H}_t(\mathbf{x}_t^{true}, \mathbf{v}_t) \ , \tag{2.2}$$

where $\mathcal{F}_t$ is the discrete transient function arising from the numerical discretization of the physical model and $\mathcal{H}_t$ is the observation function relating the true state to the observations. The variables $\mathbf{w}_t \in \mathbb{R}^n$ and $\mathbf{v}_t \in \mathbb{R}^m$ represent the unknown model and measurement errors and are modeled as independent realizations of two random variables W_t and V_t with given PDF. Filtering methods aim at evaluating the filtering PDF of the system state $p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t}^{obs})$ (where $\mathbf{y}_{1:t}^{obs} = \{\mathbf{y}_1^{obs}, \ldots, \mathbf{y}_t^{obs}\}$), given the knowledge of the forecast distribution, $p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t-1}^{obs})$, and the new measurements available at time t, $\mathbf{y}_t^{obs}$. In sequential methods, this is usually done in two steps.

1. The forward or forecast step, in which the PDF of the system state is propagated in time to obtain the forecast PDF $p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t-1}^{obs})$. Derived from the law of total probability, the Chapman-Kolmogorov equation for the prior PDF characterizes the

forecast step:

$$p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t-1}^{obs}) = \int p(\mathbf{x}_t \,|\, \mathbf{x}_{t-1}) p(\mathbf{x}_{t-1} \,|\, \mathbf{y}_{1:t-1}^{obs}) \, d\mathbf{x}_{t-1} \ . \tag{2.3}$$

2. The update or analysis step, in which Bayes theorem is used to express the filtering PDF as a correction of the forecast PDF:

$$p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t}^{obs}) = C \, p(\mathbf{y}_t^{obs} \,|\, \mathbf{x}_t) p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t-1}^{obs}), \tag{2.4}$$

where C is a normalization constant and $p(\mathbf{y}_t^{obs} \,|\, \mathbf{x}_t)$ is the likelihood function of the observations.

In the case of linear dynamics and measurement model with additive Gaussian noise, the solution of the filtering problem is given by the well-known Kalman filter:

$$\mathbf{x}_t^a = \mathbf{x}_t^f + \mathbf{P}_t^{\mathbf{xy}}(\mathbf{P}_t^{\mathbf{yy}})^{-1}(\mathbf{y}_t^{obs} - \mathbf{H}_t \mathbf{x}_t^f) \ , \tag{2.5}$$

$$\mathbf{P}_t^a = \mathbf{P}_t^{\mathbf{xx}} - \mathbf{P}_t^{\mathbf{xy}} \mathbf{P}_t^{\mathbf{yy}} \mathbf{P}_t^{\mathbf{yx}} \ , \tag{2.6}$$

where $\mathbf{H}_t$ is the observation operator, $\mathbf{P}^{\mathbf{xy}}$ is the cross-covariance matrix between the forward state and the observations, $\mathbf{P}^{\mathbf{yy}}$ is the covariance matrix of the observations, and $\mathbf{P}^{\mathbf{xx}}$ is the prior (i.e., the forecast) covariance matrix.

When the state-space model is nonlinear, we cannot assert that prior and filtering PDFs are Gaussian at all times. Nevertheless, it is possible to demonstrate that the best linear state estimator (in the sense that minimizes the trace of the error covariance matrix) after an assimilation is still given by Eq. (2.5) and Eq. (2.6) [Jazwinski, 1970]. However, this does not ensure that the estimator coincides with the expected value of the filtering PDF. Thus, Eq. (2.5) and Eq. (2.6) employ the so called Gaussian approximation, i.e., the filtering PDF is approximated with the Gaussian PDF $\mathcal{N}(\mathbf{x}_t^a, \mathbf{P}_t^a)$.

Evaluating the expected values and covariance matrices involve the calculation of multidimensional integrals of the form:

$$\int \phi(\mathbf{x}_t) p(\mathbf{x}_t | \mathbf{y}_{0:t}^{obs}) \, d\mathbf{x}_t \ , \qquad \int \phi(\mathbf{x}_t) p(\mathbf{x}_t | \mathbf{y}_{0:t-1}^{obs}) \, d\mathbf{x}_{t-1}, \tag{2.7}$$

where $\phi(\mathbf{x})$ is an integrable function with respect the measure p (e.g., $\phi(\mathbf{x}) = \mathbf{x}$ for the computation of the first moment, $\phi(\mathbf{x}) = \mathbf{x}^2$ for the computation of the second moment). Different DA methods based on the Gaussian approximation use different ways to evaluate

these integrals.

2.3.2   Ensemble Kalman filter

The key idea of EnKF is to use the Monte Carlo method to propagate in time the PDF of the system state and then perform an update on the basis of the Kalman filter. To this aim, an ensemble is constructed from the nonlinear dynamics of the model. Let N_{ens} be the ensemble size and $\{\mathbf{x}_{t-1}^i \in \mathbb{R}^n\}$, $i = 1, \ldots, N_{ens}$, be the collection of realizations of the system state at time $t - 1$. The empirical distribution associated to the ensemble $\{\mathbf{x}_{t-1}^i\}_{i=1,\ldots,N_{ens}}$ approximates $p(\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}^{obs})$, i.e.,

$$p(\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}^{obs}) \approx \frac{1}{N_{ens}} \sum_{i=1}^{N_{ens}} \delta(\mathbf{x}_{t-1} - \mathbf{x}_{t-1}^i),$$

where $\delta(\mathbf{x})$ is the Dirac-delta function. The Chapman-Kolmogorov equation (2.3) is approximated by:

$$\begin{aligned}
p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t-1}^{obs}) &= \int p(\mathbf{x}_t \,|\, \mathbf{x}_{t-1}) p(\mathbf{x}_{t-1} \,|\, \mathbf{y}_{1:t-1}^{obs}) \, d\mathbf{x}_{t-1} \\
&\approx \frac{1}{N_{ens}} \sum_{i=1}^{N_{ens}} p(\mathbf{x}_t \,|\, \mathbf{x}_{t-1}^i) \approx \frac{1}{N_{ens}} \sum_{i=1}^{N_{ens}} \delta(\mathbf{x}_t - \mathbf{x}_t^{i,f})
\end{aligned} \tag{2.8}$$

where $\mathbf{x}_t^{i,f}$ is a sample from $p(\mathbf{x}_t \,|\, \mathbf{x}_{t-1}^i)$ In this way the forecast step is simply obtained by the propagation in time of each Monte Carlo realization using the dynamical model:

$$\mathbf{x}_t^i = \mathcal{F}_t(\mathbf{x}_{t-1}^i, \mathbf{w}_t^i), \qquad i = 1, \ldots, N_{ens}, \tag{2.9}$$

where $\mathbf{w}_t^i$ are independent random samples from the noise W_t. The ensemble average $\mathbf{x}_t^f = \frac{1}{N_{ens}} \sum_{i=1}^{N_{ens}} \mathbf{x}_t^i$ can be considered as the forecast state estimation. Let $\mathbf{X}_t^f$ be the matrix whose columns are $(\mathbf{x}_t^{1,f}, \ldots, \mathbf{x}_t^{N_{ens},f})$. The analysis step is obtained by using Eq. (2.5) to update each realization $\mathbf{x}_t^{i,f}$. To this aim, we draw a sample $\mathbf{y}_t^i \in \mathbb{R}^m$ from $p(\mathbf{y}_t \,|\, \mathbf{x}_t^{i,f})$ for each $i = 1, \ldots, N_{ens}$, and evaluate the sample covariance matrices approximating $\mathbf{P^{yy}}$ and $\mathbf{P^{xy}}$ of Eq. (2.5), yielding [Haug, 2005]:

$$\mathbf{x}_t^{i,a} = \mathbf{x}_t^{i,f} + \mathbf{P}_t^{\mathbf{X}^f\mathbf{Y}}(\mathbf{P}_t^{\mathbf{YY}})^{-1}(\mathbf{y}_t^{obs} - \mathbf{y}_t^i) \tag{2.10}$$

where, indicating with $\mathbf{Y}_t$ the matrix whose columns are $(\mathbf{y}_t^1, \ldots, \mathbf{y}_t^{N_{ens}})$, the sample covariance matrices are:

$$\mathbf{P}_t^{\mathbf{X}^f\mathbf{Y}} = \text{Cov}[\mathbf{X}_t^f, \mathbf{Y}_t],$$

$$\mathbf{P}_t^{\mathbf{YY}} = \text{Cov}[\mathbf{Y}_t].$$

Finally, the filtering PDF $p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t}^{obs})$ is approximated with the empirical measure of $\mathbf{X}_t = (\mathbf{x}_t^1, \ldots, \mathbf{x}_t^{N_{ens}})$, such that the analysis state estimation is given by $\mathbf{x}_t^a = \frac{1}{N_{ens}} \sum_{i=1}^{N_{ens}} \mathbf{x}_t^i$. The filtering PDF $p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t}^{obs})$ is approximated with the empirical distribution associated to $\mathbf{x}_t^{1,a}, \ldots, \mathbf{x}_t^{N_{ens},a}$, and the analysis state estimation is given by their mean. For the implementation details of EnKF see, e.g., [Evensen, 2007].

The main characteristic of Eq. (2.10) is its ability to drive the state of each realization toward the observations, even if the forecast realizations are "far" from the observations. This distinctive feature allows the filter to obtain accurate results in many applications also with small ensemble sizes. The crucial point in the EnKF methodology is the use of the Kalman update also if forecast and measurement error distributions are not guaranteed to be Gaussian. This is usually referred to as Gaussian approximation and can lead to suboptimal properties of the filter. Moreover, the realizations are updated regardless of the physical consistency of the resulting system variables. In other words, the re-initialized realizations $\mathbf{x}_t^{i,a}$ in (2.10) may not be consistent with the dynamics of the model, leading to numerical difficulties and inaccuracies in the simulation of the following time period.

2.3.3 Particle filters and SIR algorithm

PFs are Monte Carlo methods where, differently from EnKF, the Gaussian approximation in the update step is avoided. The key idea is to approximate the filtering PDF by a set of random samples $\mathbf{x}_t^i$ with associated weights, ω_t^i, where samples and weights are calculated directly using the Bayesian filter (2.4). The weighted representation of the filtering PDF derives from the importance sampling technique that approximates the filtering PDF as:

$$p(\mathbf{x}_t \,|\, \mathbf{y}_{1:t}^{obs}) \approx \sum_{i=1}^{N_{ens}} \omega_t^i \delta(\mathbf{x}_t - \mathbf{x}_t^i).$$

The weights and the realizations can be computed in different ways and we refer to [Arulampalam et al., 2002] for the general PF theory. We describe here the SIR procedure, which starts by assigning uniform weights to N_{ens} independent realizations of the initial

PDF $p(\mathbf{x}_0)$, $\{\mathbf{x}_0^i, \omega_0^i\}_{i=1,...,N_{ens}}$, $\omega_0^i = 1/N_{ens}$. As in EnKF, the forecast step is performed using Eq. (2.9) for each realization $\mathbf{x}_t^i$, without varying the weights. The update step given by Eq. (2.4) is simply obtained by computing the new weights recursively through the likelihood function:

$$\tilde{\omega}_t^i = \omega_{t-1}^i p(\mathbf{y}_t^{obs} \,|\, \mathbf{x}_t^i), \tag{2.11}$$

$$\omega_t^i = \frac{\tilde{\omega}_t^i}{\sum_{i=1}^{N_{ens}} \tilde{\omega}_t^i}. \tag{2.12}$$

The main drawback of this algorithm, known as sequential importance sampling, is that after a few steps most of the weights may degenerate to negligible values, leading to an inaccurate state estimation and a waste of computational effort to propagate particles with negligible importance. This deterioration phenomenon, analogous to the filter inbreeding affecting EnKF [Evensen, 2009, Hendricks Franssen and Kinzelbach, 2008], may occur for example when the sample size is small with respect to the space dimension, or when the PDF associated to the observations has a small variance [Snyder et al., 2008]. A measure of the deterioration is given by the effective sample size N_{eff} ($\leq N_{ens}$), [Liu and Chen, 1998]:

$$N_{eff} = \frac{1}{\sum_{i=1}^{N_{ens}} (\omega_t^i)^2}.$$

N_{eff} gives an estimate of how many realizations are effectively contributing to the determination of the empirical PDF. To avoid deterioration that may occur for small values of N_{eff} (typically $N_{eff} < 0.5 N_{ens}$, but this value depends on the ensemble size and the specific application), a resampling step is introduced after the update step: when deterioration occurs, samples with large weights are randomly duplicated and samples with small weights are discarded, maintaining the sample size equal to N_{ens} and associating uniform weights ($1/N_{ens}$) to the newly generated realizations. The resulting algorithm is named sequential importance resampling filter. There are different types of resampling methods; in our application we decided to implement the stratified systematic resampling explained in [Arulampalam et al., 2002].

To minimize the ensemble impoverishment and to improve the filter performance, the resampled particles must differentiate during the evolution in time. This condition is ensured if a large enough variance of the system dynamics noise is imposed.

From a theoretical point of view, the SIR algorithm is more appealing than EnKF, as its convergence can be proved under much weaker hypotheses [Crisan, 2001]. When the sample size tends to infinity, the empirical measure associated to SIR converges P-almost

surely (Eq. (2.13)) and in mean (Eq. (2.14)) to the measure associated to the filtering distribution, i.e.:

$$P\left(\lim_{N_{ens}\to\infty}\sum_{i=1}^{N_{ens}}\omega_t^i\mathbf{z}\left(\mathbf{x}_t^i\right) = \int\mathbf{z}\left(\mathbf{x}_t\right)p\left(\mathbf{x}_t\,|\mathbf{y}_{1,t}^{obs}\right)\,d\mathbf{x}_t\right) = 1, \tag{2.13}$$

$$\lim_{N_{ens}\to\infty}E\left[\sum_{i=1}^{N_{ens}}\omega_t^i\mathbf{z}\left(\mathbf{x}_t^i\right) - \int\mathbf{z}\left(\mathbf{x}_t\right)p\left(\mathbf{x}_t\,|\mathbf{y}_{1,t}^{obs}\right)\,d\mathbf{x}_t\right] = 0\ . \tag{2.14}$$

This holds for every continuous and bounded function $\mathbf{z}\in\mathbb{R}^n$, where $\{\mathbf{x}_t^i,\omega_t^i\}$ are the particles and the weights given by SIR at time t. One of the main drawbacks for the application of SIR to high-dimensional systems is that in the update step the realizations are only discarded or duplicated. If all the realizations were far from the observations, the entire ensemble would be discarded, yielding an update failure and a delayed retrieval of the true state. For this reason, ensemble sizes larger than those employed in EnKF have been reported often in the literature [Zhou et al., 2006].

2.4 Model description

Catchment dynamics are simulated using the model CATHY (CATchment HYdrology). CATHY couples subsurface and surface flow by integrating the three-dimensional Richards equation for variably saturated porous media with a one-dimensional diffusion wave approximation of the Saint Venant equation for surface water routing [Camporese et al., 2010]. The mathematical model can be written as:

$$S_w(\psi)S_s\frac{\partial\psi}{\partial t} + \phi\frac{\partial S_w(\psi)}{\partial t} = \mathrm{div}\left[K_s K_r(S_w(\psi))\left(\nabla\psi + \eta_z\right)\right] + q_{ss}(h_p) \tag{2.15}$$

$$\frac{\partial Q}{\partial t} + c_k\frac{\partial Q}{\partial s} = D_h\frac{\partial^2 Q}{\partial s^2} + c_k q_s(h_p,\psi), \tag{2.16}$$

where S_w is water saturation [-], S_s is the aquifer specific storage coefficient $[L^{-1}]$, ψ is pressure head $[L]$, t is time $[T]$, ϕ is the porosity or saturated moisture content [-], K_s is the saturated hydraulic conductivity tensor $[L/T]$, K_r is the relative hydraulic conductivity function [-], $\eta_z = (0,0,1)^T$, with z the vertical coordinate directed upward $[L]$, and q_{ss} represents both distributed source or sink terms $[L^3/L^3T]$ and the exchange fluxes between the surface and subsurface domains. This exchange term depends on the ponding head h_p $[L]$, i.e., the depth of water on the ground surface. The surface water is routed

using Eq. (2.16) along each single hillslope or channel link defined in a one-dimensional coordinate system s $[L]$ (the drainage network) calculated using state-of-the-art methods of digital elevation model (DEM) analysis [Orlandini et al., 2003]. In this equation, Q is the discharge along the drainage link $[L^3/T]$, c_k is the kinematic wave celerity $[L/T]$, D_h is the hydraulic diffusivity $[L^2/T]$, and q_s is the inflow or outflow exchange rate from the subsurface to the surface $[L^3/LT]$. Values of c_k and D_h are calculated as a function of discharge by means of the Leopold and Maddock relationships Leopold and Maddock Jr. [1953], with different parametrization to distinguish between channel and hillslope links.

A linear tetrahedral finite element method with backward Euler discretization forms the core of the Richards equation solver, together with Newton-like iterations for the nonlinear system [Paniconi and Putti, 1994]. The Muskingum-Cunge or Matched Artificial Diffusion (MAD) scheme solves the surface equation [Orlandini and Rosso, 1996]. Spatial discretization proceeds from a DEM representation of the catchment surface. The DEM cells are triangulated and replicated vertically to form a three-dimensional tetrahedral grid for the underlying soil and aquifer. We denote with *nnod* the number of subsurface grid nodes and with *ncell* the number of surface cells forming the DEM.

Net precipitation fluxes during storm events and potential evaporation during interstorm periods are the main driving forces of the model and are expressed by a boundary condition term denoted here by u_t $[L/T]$. The catchment partitions this atmospheric forcing into surface runoff, infiltration, actual evaporation, and changes in storage via a surface boundary condition switching algorithm [Camporese et al., 2010, Putti and Paniconi, 2004]. Surface saturation or ponding can occur via the infiltration excess or saturation excess mechanisms, and both of these are automatically accounted for by the same switching algorithm. The fluxes q_{ss} and q_s are both functions of the boundary condition term u_t and the ponding head h_p, which is related to the surface water volume $V_{s,k}$ $[L^3]$ at each time t_k.

Considering a DEM cell of surface A $[L^2]$, the variation of water volume over the cell between times t_1 and t_2 is given by:

$$\Delta V_s = \int_{t_1}^{t_2} \left[Q^{in}(t) - Q^{out}(t) \right] \, dt \qquad (2.17)$$

where Q^{in} and Q^{out} are the cell inflow and outflow fluxes calculated by the MC scheme. Ponding head at each time $h_{p,k}$ can be approximated by $h_{p,k} = V_{s,k}/A$, the volume being updated by Eq. (2.17), with the integral numerically evaluated by means of the trapezoidal rule.

At each time step the resulting coupling algorithm can be summarized as follows.

1. Solve the surface module by the explicit MAD:

$$\mathbf{A}(Q_k)Q_{k+1} = \mathbf{B}(Q_k)Q_k + q_s(h_{p,k}, \psi_k),$$

   where $\mathbf{A}$ and $\mathbf{B}$ are the triangular matrices given by MAD and Q_k is the vector whose components are the outflow discharges at each surface cell ($Q_k \in \mathbb{R}^{ncell}$).

2. Calculate new cell volumes $V_{s,k+1} \in \mathbb{R}^{ncell}$ and ponding head $h_{p,k+1}$ on the surface nodes by means of (2.17) and a simple average:

$$V_{s,k+1} = \frac{(Q_{k+1}^{in} - Q_{k+1}^{out}) + (Q_k^{in} - Q_k^{out})}{2}\Delta t + V_{s,k},$$
$$h_{p,k+1} = \frac{V_{s,k+1}}{A}.$$

3. Evaluate the surface/subsurface boundary flux $q_s(h_p)$ and solve the nonlinear system given by the finite element method for the pressure head ψ_{k+1} on the subsurface nodes ($\psi_{k+1} \in \mathbb{R}^{nnod}$):

$$\mathbf{M}(\psi_{k+1})\psi_{k+1} = \mathbf{F}(\psi_k) + q_s(h_{p,k+1}),$$

   where $\mathbf{M}$ and $\mathbf{F}$ are matrices resulting from the finite element spatial discretization and backward Euler scheme.

Note that time stepping sequence, indexed by subscript k, does not coincide with assimilation times, indicated here with subscript t, so that between two consecutive assimilation times $t-1$ and t, several time steps are performed by the numerical model. More details are given in Putti and Paniconi [2004] and Camporese et al. [2010].

2.4.1 Implementation details of DA methods

The model dependent variables are: the pressure head at each subsurface node, collected in vector ψ_t; the inflow and outflow discharge and the water volume of each surface cell, collected in vectors Q_t^{in}, Q_t^{out}, $V_{s,t}$, respectively. The model state vector at time t is given by:

$$\mathbf{x}_t = \{\psi_t, Q_t^{in}, Q_t^{out}, V_{s,t}\},$$

with dimension $n = nnod + 3\,ncell$. The model equations are written at the assimilation time t as:

$$\mathbf{x}_t = \mathcal{F}(\mathbf{x}_{t-1}, \boldsymbol{\alpha}, \mathbf{u}_t, t) \; ; \; t = 1, \ldots, T \; ,$$

where vector $\boldsymbol{\alpha}$ represents the time-invariant soil parameters (saturated hydraulic conductivity, specific storage, porosity, retention curve parameters, and surface routing parameters) and the initial condition is denoted by $\mathbf{x}_0$.

In the following, we assume that the model inputs subject to uncertainties are the initial conditions $\mathbf{x}_0$, the parameters $\boldsymbol{\alpha}$ (except for surface routing parameters), and the time variable potential precipitation/evapotranspiration rates u_t. We postulate an additive normal noise for initial pressure head and a multiplicative log-normal uncertainty distribution for the soil parameters and the atmospheric forcing.

We initialize and propagate the ensemble of random realizations using the same procedure for both EnKF and SIR methods. The initial state of a single realization $\mathbf{x}_0^i$ ($i = 1, \ldots, N_{ens}$) is a spatially-homogeneous perturbation of the nominal initial state. In this way every realization represents a physically-consistent state, at the same time ensuring that the realizations are well differentiated at the beginning of the simulation.

The uncertainty on subsurface parameters is accounted for by associating to each realization a set of random parameters $\boldsymbol{\alpha}^i$, obtained by perturbing the known input nominal values. Each realization is propagated in time perturbing also the atmospheric forcing with a spatially-homogeneous noise time series. The noise on the atmospheric forcing is temporally correlated for each realization as follows [Evensen, 2003]:

$$z_k^i = \gamma z_{k-1}^i + \sqrt{1 - \gamma^2} \zeta_{k-1}^i,$$

where z_k^i is a time correlated sequence of perturbations normally distributed with zero mean and unit variance and ζ_k^i is a white noise sequence. Perturbations z_k^i are then transformed in order to be applied to the atmospheric forcing $\mathbf{u}_k$ as multiplicative log-normal errors with mean equal to 1 and known (imposed) variance. The coefficient γ determines the time decorrelation of the noise on the stochastic forcing and is computed as $1 - \Delta t/\tau$, where Δt is the current time step and τ is the specified time decorrelation length. Thus, the forecast step for both the considered DA methods can be written as:

$$\mathbf{x}_t^i = \mathcal{F}(\mathbf{x}_{t-1}^i, \boldsymbol{\alpha}^i, \mathbf{u}_t^i, t); \qquad t = 1, \ldots, T; \; i = 1, \ldots, N_{ens},$$

where $\mathbf{u}_t^i$ represents the perturbed atmospheric forcing.

We consider assimilation of soil moisture $\mathbf{y}_t^\theta$, pressure head $\mathbf{y}_t^\psi$, and streamflow at the catchment outlet $\mathbf{y}_t^Q$. The pressure head and streamflow measurement operators are simply a linear projection of the state vector components to the measurement nodes/cells. For soil moisture measurements, the van Genuchten retention curve model [Van Genuchten, 1980] is used to relate $\mathbf{y}_t^\theta$ to the pressure head. The densities $p(\mathbf{y}_t \,|\, \mathbf{x}_t)$ are explicitly derived from the formula for the change of variable, to obtain:

$$p(y_{t,j}^\psi \,|\, \mathbf{x}_t) \sim \mathcal{N}(x_{t,i}, \sigma_\psi^2) \ , \tag{2.18}$$

$$p(y_{t,j}^\theta \,|\, \mathbf{x}_t) \sim Log\mathcal{N}(\ln[h_\theta(x_{t,i})], \sigma_\theta^2) \ , \tag{2.19}$$

$$p(y_{t,j}^Q \,|\, \mathbf{x}_t) \sim Log\mathcal{N}(\ln[(x_{t,i})], \sigma_Q^2) \ , \tag{2.20}$$

where j and i are vector components for the measurement and the corresponding system state vectors, respectively, and h_θ represents the functional relation given by the moisture retention curves.

2.5 Numerical results

2.5.1 Model setup

A series of numerical simulations were conducted on a synthetic three-dimensional test case. We considered a tilted 3 m deep v-catchment with a surface area of 1.62 km^2, characterized by a homogeneous and isotropic saturated hydraulic conductivity field, with no-flow conditions at the bottom and at the lateral boundaries. Table 2.1 summarizes the model discretization and parameter values used for the numerical experiments.

We would like to stress that the main objective of this research is to compare the performance of EnKF and SIR in the retrieval of the system state PDFs, and we do not address the problem of parameter estimation. For this reason, each realization maintains the same set of time independent parameters throughout the simulation, without any parameter update. Although the dimension of the parameter space in our numerical experiments is relatively small (uncertainty on 6 soil parameters, initial conditions, and time-variable atmospheric forcings), the dimension of the state space is large (more than 30000). Moreover, the high model nonlinearities related to the unsaturated zone and to the coupling between surface and subsurface flow add to the difficulty of numerically approximating the state PDFs.

The true watershed system state was identified by the numerical solution of our

Grid information	
DEM cell dimensions	20×20 m
Number of cells in the surface grid	$50 \times 81 = 4050$
Soil depth (uniform)	3 m
Vertical discretization (# of layers)	6
Soil layer thickness (top to bottom)	0.27, 0.36, 0.39, 0.66, 0.66, 0.66 m
# of nodes in the 3D grid	$4182 \times 7 = 29274$
# of tetrahedral elements in the 3D grid	145800
Parameters	log-normally distributed
Saturated hydraulic conductivity	$E[K_s] = 1.16 \times 10^{-5}$ m/s, $CV_{K_s} = 100\%$
Aquifer specific storage	$E[S_s] = 5 \times 10^{-4}$ m^{-1}, $CV_{S_s} = 5\%$
Porosity	$E[\phi] = 0.40$, $CV_{\phi} = 5\%$
Van Genuchten curve fitting	$E[\alpha] = 0.47$ m^{-1}, $CV_{\alpha} = 5\%$
parameters	$E[n] = 1.70$, $CV_n = 5\%$
Residual moisture content	$E[\theta_r] = 0.06$, $CV_{\theta_r} = 5\%$
Simulation period	14400 s (4 h)
Initial conditions	hydrostatic profile with normally distributed water-table depth W_{ψ}
True IC	$E[W_{\psi}] = 0.5$ m, $\sigma_{W_{\psi}} = 0.1$ m
Biased IC	$E[W_{\psi}] = 1.0$ m, $\sigma_{W_{\psi}} = 0.2$ m
(scenarios 7-12)	
Atmospheric forcing	log-normally distributed with decorrelation time τ
True atmospheric BC	
from $t = 0$ to $t = 1.5$ h	$E[u_t] = 10.8$ mm/h (rain)
	$CV_{u_t} = 50\%$, $\tau = 1000$ s
from $t = 1.5$ h to $t = 4.0$ h	$E[u_t] = -1.08$ mm/h (evaporation)
	$CV_{u_t} = 50\%$, $\tau = 1000$ s
Biased atmospheric BC (scenarios 1-6)	
from $t = 0$ to $t = 1.5$ h	$E[u_t] = 5.4$ mm/h (rain)
	$CV_{u_t} = 100\%$, $\tau = 1000$ s
from $t = 1.5$ h to $t = 4.0$ h	$E[u_t] = -2.16$ mm/h (evaporation)
	$CV_{u_t} = 100\%$, $\tau = 1000$ s
Measurements	
Measures of pressure head ψ	normally distributed with $\sigma_{\psi} = 0.03$ m
Measures of streamflow Q	log-normally distributed with $CV_Q = 3\%$
Ensemble size N_{ens}	100

Table 2.1: Model discretization and parameter values for the v-catchment test case. CV indicates the coefficient of variation.
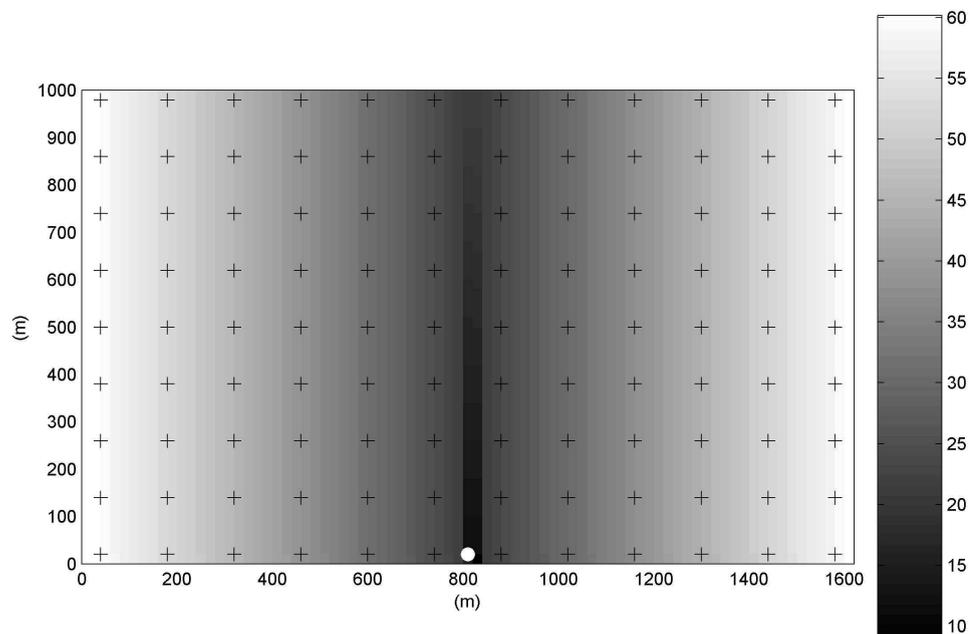
Figure 2.1: Digital elevation model of the tilted v-catchment. Elevations are expressed in meters. Black crosses indicate the location of pressure head observations and the white circle represents the location of the outlet cell, at which streamflow measurements are available.

Scenario	Assimilated Variable (ψ or Q)	Simulation Bias (IC or BC)	Update Interval (min)
1. BC-ψ-H	ψ	BC	15
2. BC-ψ-L	ψ	BC	60
3. BC-ψQ-H	ψ, Q	BC	15
4. BC-ψQ-L	ψ, Q	BC	60
5. BC-Q-H	Q	BC	15
6. BC-Q-L	Q	BC	60
7. IC-ψ-H	ψ	IC	15
8. IC-ψ-L	ψ	IC	60
9. IC-ψQ-H	ψ, Q	IC	15
10. IC-ψQ-L	ψ, Q	IC	60
11. IC-Q-H	Q	IC	15
12. IC-Q-L	Q	IC	60

Table 2.2: Configuration for the V-Catchment Scenarios. Pressure head ψ is measured at the soil surface, and catchment base (216 measures, see Fig. 2.1); streamflow Q is measured at the catchment outlet.

model (2.15), (2.16) using the following: initial conditions (IC) with uniform partially sat-
urated vertical pressure head profile in hydrostatic equilibrium for the whole watershed,
with a water table at 0.5 m below the surface; atmospheric boundary conditions (BC) char-
acterized by a spatially-uniform constant rainfall rate of 3.0×10^{-6} m/s (10.8 mm/h) from
time zero to time 5400 s (1.5 h), followed by a constant evaporation rate of 3.0×10^{-7} m/s
(1.08 mm/h) until the end of the simulation, at 14400 s (4 h). From this true simulation
we collected the observations for the analysis steps of both DA schemes. Streamflow was
measured at the outlet of the basin, while pressure heads were measured in 216 nodes, of
which 108 evenly distributed over the catchment surface (simulating soil moisture mea-
surements) and 108 at the bottom of the corresponding vertical lines at the catchment
base (representing water table measurements) (see Fig. 2.1).

EnKF and SIR were compared in 12 different scenarios, in which the following factors
were varied: simulation bias (BC for scenarios 1-6; IC for scenarios 7-12); assimilation
variable (surface and bottom pressure head for scenarios 1, 2, 7, 8; outlet streamflow for
scenarios 5, 6, 11, 12; both pressure head and streamflow for scenarios 3, 4, 9, 10); as-
similation frequency (15-minute updates for scenarios 1, 3, 5, 7, 9, 11; hourly updates for
scenarios 2, 4, 6, 8, 10, 12). These configurations, together with the input and measure-
ment error statistics, are described in Table 2.2, where they are labeled using an acronym
with the first two letters indicating the biased component (BC or IC) followed by the
assimilated variables (pressure head ψ and/or streamflow Q) and the frequency of update
(H or L for high or low frequency, respectively). For the BC biased scenarios, the rainfall
rate is halved with respect to the true simulation and the evaporation rate is doubled. For
the IC-biased scenarios, the water table is at 1 m depth rather than 0.5 m (see Table 2.1).

Open loop simulations, i.e., Monte Carlo simulations without data assimilation, were
carried out for each scenario, to better evaluate the benefits of using the filters. For
scenarios 3, 4, 9 and 10, where both pressure head and streamflow observations were
assimilated, the heads and the streamflow were normalized with the maximum values
(3.0 m and 4.86 m^3/s, respectively) in order to ensure properly scaled covariance matrices
for the EnKF scheme [Evensen, 2003]. In the SIR algorithm a resampling step is performed
whenever $N_{eff} < 0.6 N_{ens}$, i.e., when the number of effective particles falls below 60% of the
ensemble size. We remark here that, when the resampling step is not performed, different
weights are assigned to the particles and, as a consequence, the empirical ensemble PDF
of the parameters is a varying approximation of the same prior PDF. This fact does not
invalidate the comparison with the EnKF, because each realization maintains the same
set of parameters after an update.
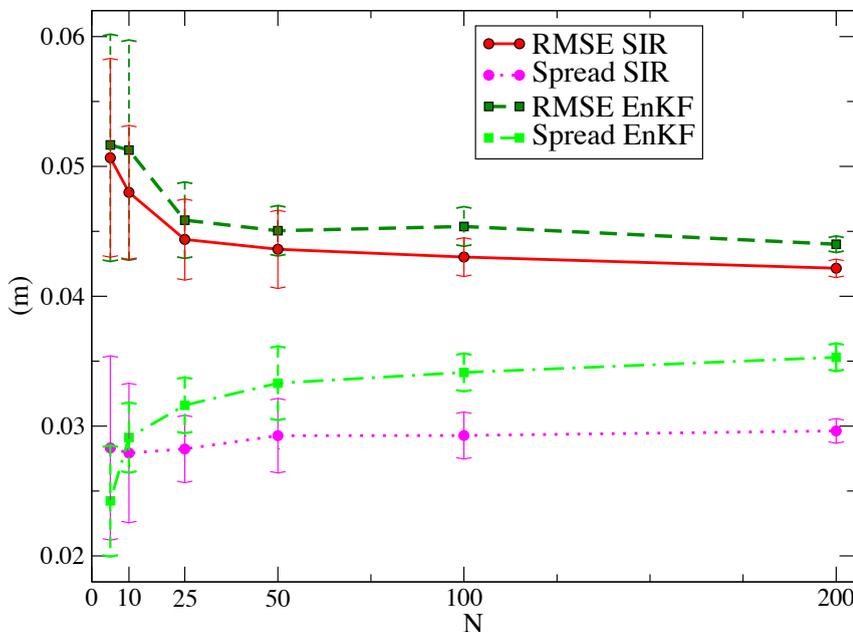
Figure 2.2: Space- and time-averaged RMSE and ensemble spread for the subsurface pressure heads versus the number of realizations of EnKF and SIR. Vertical bars represent ± 1 standard deviation computed over p runs of the filter with different initial realizations. The number of runs corresponding to $N_{ens} = \{5, 10, 25, 50, 100, 200\}$ are $p = \{20, 20, 20, 20, 10, 10\}$, respectively.

2.5.2   Sensitivity analysis to the ensemble size

An analysis on the sensitivity of the two methods to the number of realizations was performed in order to establish the suitability of the ensemble size $N_{ens} = 100$ suggested in Camporese et al. [2009b] for the present application. Obviously, this number is bound to change with problem size (the dimension of parameter and state spaces) or in the case, not addressed in this work, of parameter estimation. EnKF and SIR were compared in an unbiased scenario, with high-frequency assimilation of both pressure head and streamflow measurements. Simulations with an ensemble size of $N_{ens} = \{5, 10, 25, 50, 100, 200\}$ were repeated p times ($p = \{20, 20, 20, 20, 10, 10\}$, respectively) with different initial realizations and perturbations on the model parameters. For each simulation we computed the space and time averaged ensemble spread of the analyzed subsurface pressure head and the corresponding root mean squared error (RMSE).

Fig. 2.2 shows the mean values of the RMSE and of the ensemble spread over the p simulations (the ± 1 standard deviation is indicated by vertical bars) as functions of the ensemble size. As expected, the RMSE for both filters decreases with increasing ensem-

ble sizes, with marginal reduction for N_{ens} larger than 25. Moreover, the ± 1 standard deviation vertical bars shorten for increasing ensemble sizes, indicating that the results of the two filters using a high enough number of realizations have a low dependence on the initial ensemble of realizations.

Contrary to the RMSE, the ensemble spreads increase with the ensemble size and reach stable profiles for N_{ens} greater than 50. The relatively small values of the SIR spread can be attributed to the fact that the SIR analysis frequently consists of only a few particles, thereby leading to very small values of the spread evaluated at the corresponding assimilation times. Fig. 2.2 suggests that the RMSE and spread values of both EnKF and SIR with more than 50 realizations are only weakly dependent on the particular random simulation. Hence we may safely conclude that, in our numerical examples, the use of 100 realizations for both EnKF and SIR is a good compromise between the computational efficiency of the simulation and the independence of the results to the particular random sequence.

Note that our choice of perturbing the initial conditions with a spatially-homogeneous bias is important in ensuring this last conclusion. The drawback of this choice is that the initial volume stored in the system varies with the realizations, leading to a biased ensemble-average of the subsurface water volume. However we consider that physical consistency (w.r.t. the dynamics of Richards equation) of the initial solution is more important and leads to more efficient numerical solutions with respect to the use of a spatially-uncorrelated initial perturbation, as used, e.g., in Camporese et al. [2009a,b]. Furthermore, our approach leads to more differentiated trajectories that better explore the space of variability of the initial state distribution.

2.5.3   Sensitivity analysis to τ

We considered the same unbiased scenario of section 2.5.2 and performed a true simulation using $\tau = 1000$ s (we recall here that τ is the decorrelation time of the noise on the atmospheric forcings). Open loop, EnKF, and SIR were compared using 100 Monte Carlo realizations and $\tau = 500$ s, 1000 s, 2500 s and 5000 s. Fig. 2.3 shows the streamflow sequence obtained from the true run and the average streamflows obtained with the open loop and the EnKF with the different values of τ. For open loop and EnKF simulations the maximum and the minimum values of streamflows in the ensemble are also shown (dashed lines). Results obtained using SIR are similar to results obtained using EnKF and they are not shown for the clarity of the figure. We can observe that increasing the
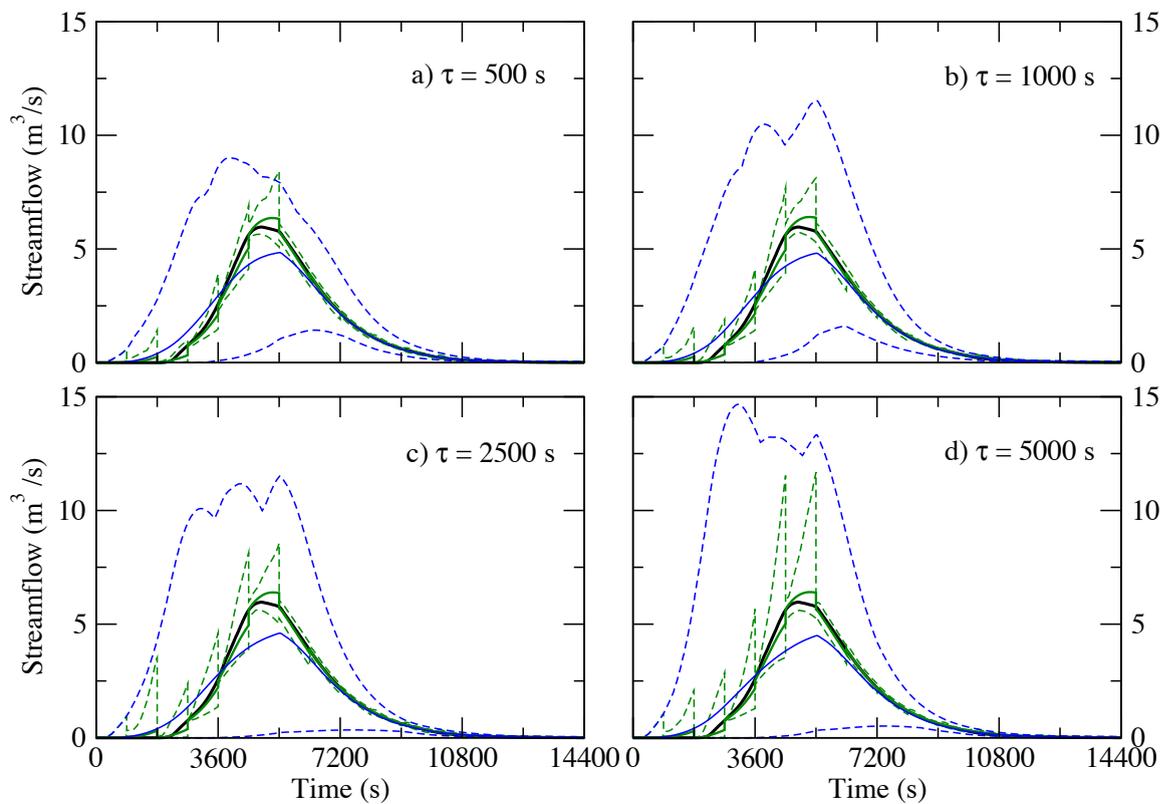
Figure 2.3: Comparison between the streamflow of the true simulation run (black line) and the averaged streamflows of the ensemble of realizations obtained with the open loop (blue continuous line) and EnKF (green continuous line). The blue and green dashed lines represent the maximum and minimum values of streamflows in the open loop and EnKF ensembles, respectively.

decorrelation time length τ does not sensibly change the average value of the realizations, while it results in an increased ensemble spread between the assimilation times, in such a way that the realizations are more differentiated. This effect is due to the differences in the atmospheric boundary conditions associated to each realization. Using a small value of τ, a single perturbed atmospheric forcing sequence displays relatively small but frequent oscillations around the nominal value, while with large τ the BCs remain far from the nominal value with higher probability and for longer times, although the ensemble mean continues to coincide with the nominal value of the atmospheric boundary conditions. This particular dependence between the ensemble and the value of τ can be useful to reduce phenomena such as the filter inbreeding affecting EnKF or the deterioration of the particle ensemble affecting SIR. In practice, the value of τ depends on the characteristic scale of the event and the particular time aggregation of the available rainfall data (e.g., 15-minute, hourly, or daily precipitation rates). In our application, we chose $\tau = 1000$ s, a physically-reasonable value for both precipitation and evapotranspiration events typical of summer storms.

2.5.4   Comparison between EnKF and SIR

Scenarios with biased atmospheric boundary conditions

In Fig. 2.4 EnKF and SIR performances are assessed in terms of their ability to improve both surface and subsurface state predictions (panels "a" and "b", respectively). We distinguish between assimilation of pressure head only (panels 1–2), both pressure head and outlet streamflow (panels 3–4) and outlet streamflow alone (panels 5–6). A further distinction displays the effects of assimilation at high frequency (15 min, left column) and low frequency (1 hour, right column). EnKF and SIR results are shown, as well as the open loop and true solutions. As expected, the biased atmospheric forcing depicts sensibly dryer scenarios with respect to the true solution. The subsurface domain is never completely saturated and the maximum value of the forecast streamflow at the outlet is less than half of the peak value of the true run.

Panels 1 and 2 of Fig. 2.4 show the results associated to scenarios 1 and 2 of Table 2.2, i.e., when only pressure head observations are assimilated, with high and low frequency, respectively. In terms of subsurface state, EnKF and SIR performances are not significantly different. Although the biased atmospheric boundary conditions force the forecast state away from the truth between two consecutive updates, the filters retrieve the true subsurface water storage at each update, both in high- and low-frequency scenarios. As-
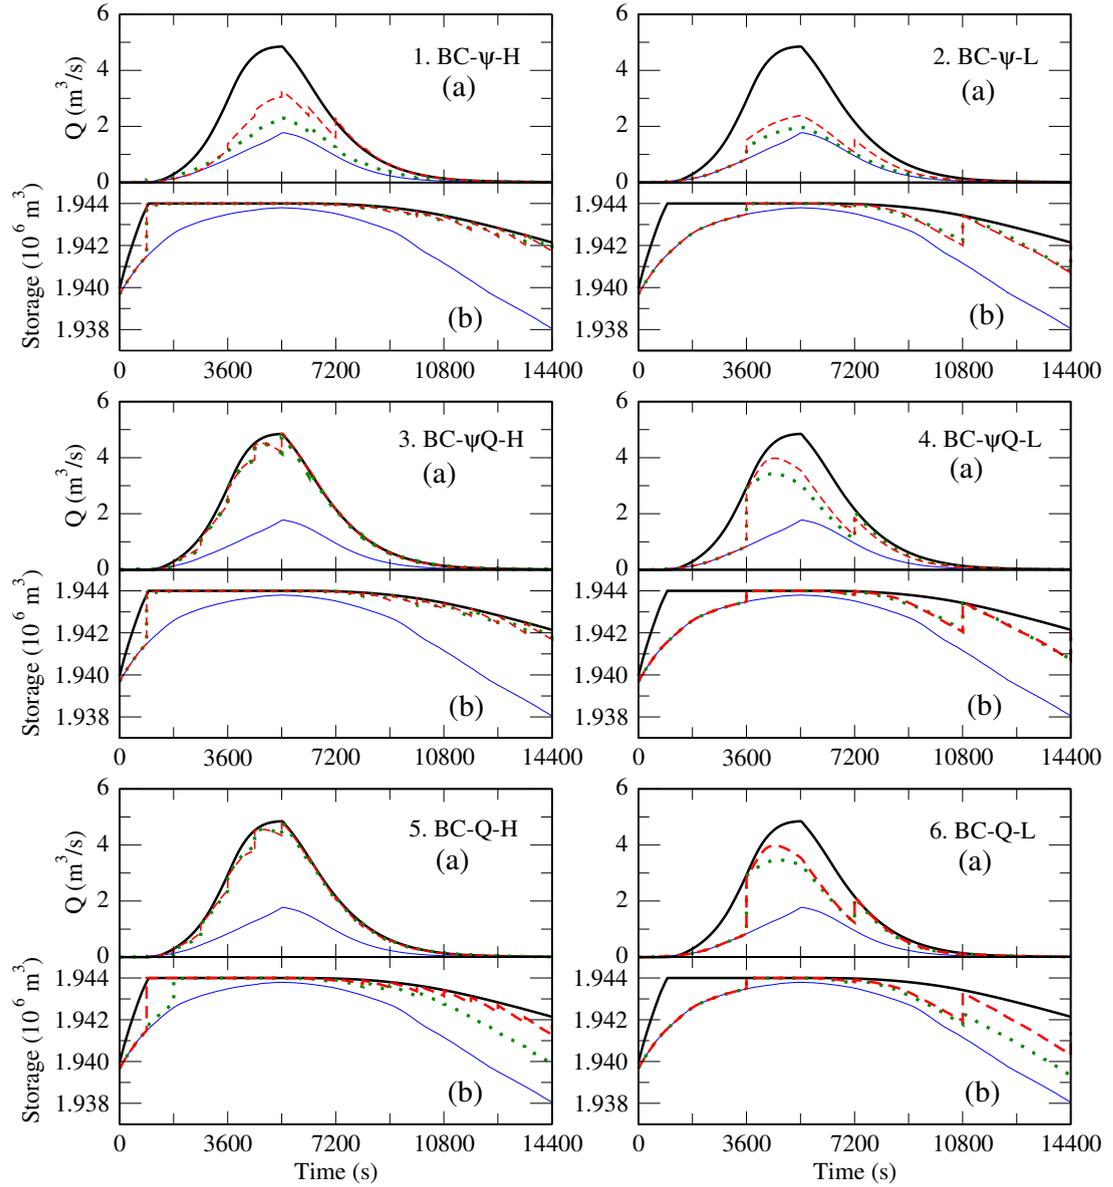
Figure 2.4: Outlet hydrograph (panels a) and subsurface water storage (panels b) for the true solution (black solid line), the ensemble mean of the open loop (blue solid line), and the ensemble mean of the realizations as given by EnKF (green dotted line) and SIR (red dashed line) for the scenarios 1-6 of Table 2.2 (biased atmospheric boundary conditions). The same bias is applied also to the open loop run.

similating pressure head only allows the two filters to slightly correct also the streamflow forecast, with the SIR predictions outperforming EnKF, both in high- and low-frequency scenarios. This improved performance may be attributed to the tighter coupling between surface and subsurface exerted by the SIR update because of the avoidance of the Gaussian approximation.

The retrieval of surface state variables improves when both pressure head and stream-flow are assimilated, as shown in panels 3a and 4a. In the high-frequency scenario (panel 3) the forecast and analysis states of the two filters are not qualitatively distinguishable and at each assimilation time the update retrieves the true values of both subsurface water storage and outlet streamflow. In the low-frequency scenario (panel 4) the time between two assimilations is too large to allow a good prediction of the hydrograph peak. The forecast maximum values of the streamflow given by EnKF and SIR are reached about 15 minutes in advance, and are about 1.42 m^3/s and 0.88 m^3/s smaller than the true peak (4.86 m^3/s), respectively.

Panels 5 and 6 show the results obtained in scenarios 5 and 6 of Table 2.2, i.e., when streamflow at the outlet is the only assimilated variable. The behavior of the two filters in terms of surface state is similar to the one described in scenarios 3 and 4. The SIR scheme continues to retrieve well the subsurface state variables, especially during the rain period. The SIR average value of subsurface volume reaches the true volume already at the first update, and then naturally follows the true value for all the saturated period. During the evaporation period, this volume decreases slightly faster than the true one, due to the bias in the forcing term. The EnKF streamflow assimilations correct the discharge forecasts, but are not sufficient to retrieve accurately the values of the pressure head and hence of the subsurface volumes. This behavior is emphasized in the EnKF high-frequency assimilation results. As panel 5b of Fig. 2.4 shows, during the rain period EnKF requires at least two updates to retrieve the true subsurface water storage and, during evaporation, the accuracy of the subsurface state variables quickly deteriorates. The same issue is manifest in the scenario of low-frequency streamflow assimilation (panel 6b of Fig. 2.4). Looking at the true value of the final subsurface water volume at the end of scenario 5 (1942145 m^3) we observe that SIR results are more accurate with respect to both EnKF and open loop, with differences of -379 m^3, -2248 m^3 and -4105 m^3, respectively.

These results suggest that, when streamflow is the only variable assimilated, both EnKF and SIR improve predictions of both surface and subsurface states with respect to the open loop. However, SIR seems to outperform EnKF in the retrieval of subsurface pressure heads. This result can be attributed to the resampling technique choosing the
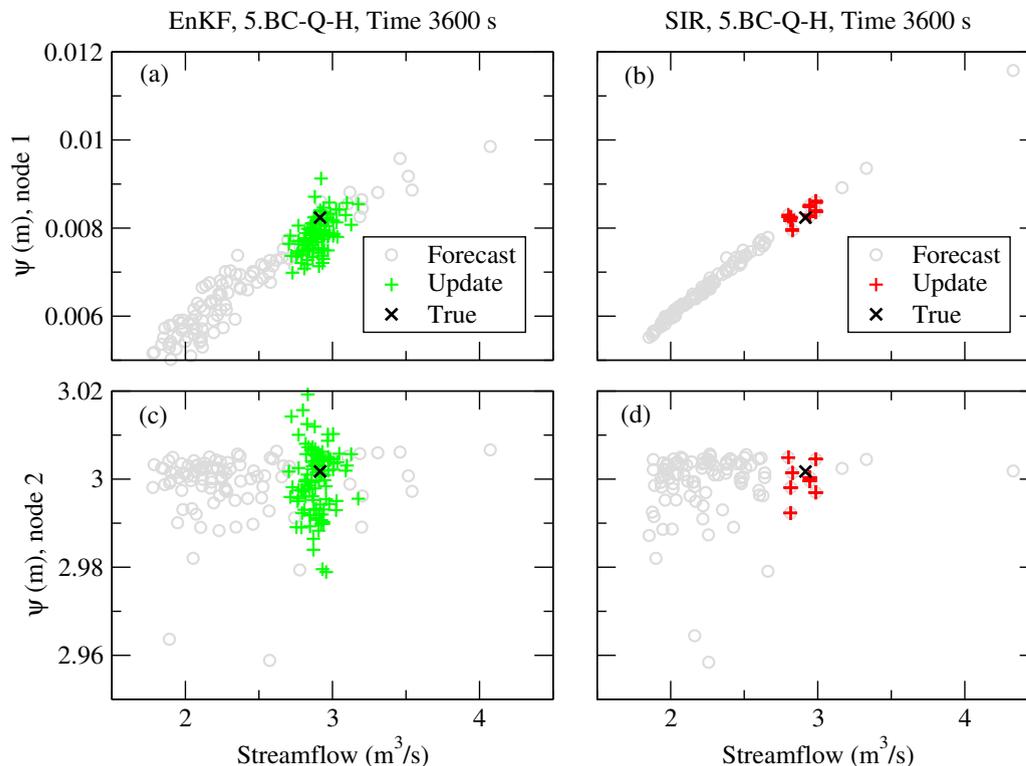
Figure 2.5: Scatter plot of the ensemble pressure head (ψ) with respect to the ensemble streamflow at the outlet. The pressure head is measured at two observation nodes, a surface (node 1) for panels (a) and (b) and the corresponding node at the bottom of the basin (node 2) for panels (c) and (d). The circles and the plus represent the forecast and the update values, respectively, associated to the EnKF ensemble (panels a and c) and the SIR ensemble (panels b and d). The cross indicates the true values.

realizations closest to the streamflow observations. Because of the tight coupling exerted by the hydrological model between surface and subsurface flow, these realizations honor also the subsurface dynamics, giving reason to the better behavior of SIR. On the other hand, in EnKF, the corrections introduced at each assimilation time are related to the spatial and temporal sensitivities of the system state as evaluated by the covariance matrices forming the Kalman gain. The coupling between surface and subsurface states in the EnKF correction is not enforced by the physics of the problem and can become weaker in the analysis vector, leading to a lower accuracy in the retrieval of the subsurface volume with respect to the SIR method. To clarify this point, we report in Fig. 2.5 the scatter plots of the ensemble values of streamflow at the outlet and of pressure head at two reference nodes: a surface node, indicated as node 1 (top panels), and the corresponding node at the bottom of the basin, indicated as node 2 (bottom panels). The results are

	EnKF Forecast	EnKF Update	SIR Forecast	SIR Update
$\rho_{\psi_1,Q}$	0.950	0.530	0.995	0.743
$\rho_{\psi_2,Q}$	0.059	0.012	0.110	0.196

Table 2.3: Cross correlation coefficients between the ensemble values of the streamflow at the outlet (Q) and pressure head at the reference node 1 (ψ_1) and 2 (ψ_1). The values refer to the forecast and the update at time 3600 s of scenario 5 of Table 2.2.

	EnKF Forecast	EnKF Update	SIR Forecast	SIR Update
ψ_1	1	0	1	1
ψ_2	1	0	1	1
Q	1	0	1	1

Table 2.4: Results of the Lilliefors tests on the null hypothesis that the data are a random sample from a normal distribution. The data considered are the EnKF and SIR ensemble values of pressure head at the reference node 1 (ψ_1), reference node 2 (ψ_2), and streamflow at the outlet (Q). The result of the test is 1 if the null hypothesis can be rejected at the 5% significance level, 0 otherwise.

relative to the fourth assimilation step (3600 s) of scenario 5 and include the forecast and analysis states of EnKF (left panels), and SIR (right panels). Table 2.3 reports the cross correlation coefficients (ρ) between streamflow and pressure head and shows that for both EnKF and SIR forecast, the correlation is much higher at the surface node. This reinforces the idea that, assimilating streamflow alone, EnKF has more difficulties in the correction of the subsurface variables. Fig. 2.5 suggests that the ensemble distribution of the forecast is not multivariate Gaussian, as the scatter plots do not show the classical elliptic shape. This observation is confirmed by Lilliefors statistical tests Lilliefors [1967] applied to the ensembles to validate the null hypothesis that pressure head and streamflow are random samples from a normal distribution. The values reported in Table 2.4 confirm that the hypothesis of Gaussianity can be rejected for the ensemble forecasts. The Gaussian approximation enforced in EnKF constrains the update to be Gaussian and generates some realizations with pressure values that are outside the forecast range, while the SIR analysis maintains the non-Gaussianity of the forecast.

Scenarios with biased initial conditions

In Fig. 2.6 the results for scenarios 7-12 with biased initial conditions are presented. Analogously to scenarios 1-6, the biased initial conditions cause scenarios dryer than the true one. Again, in the open loop, the subsurface domain is never completely saturated
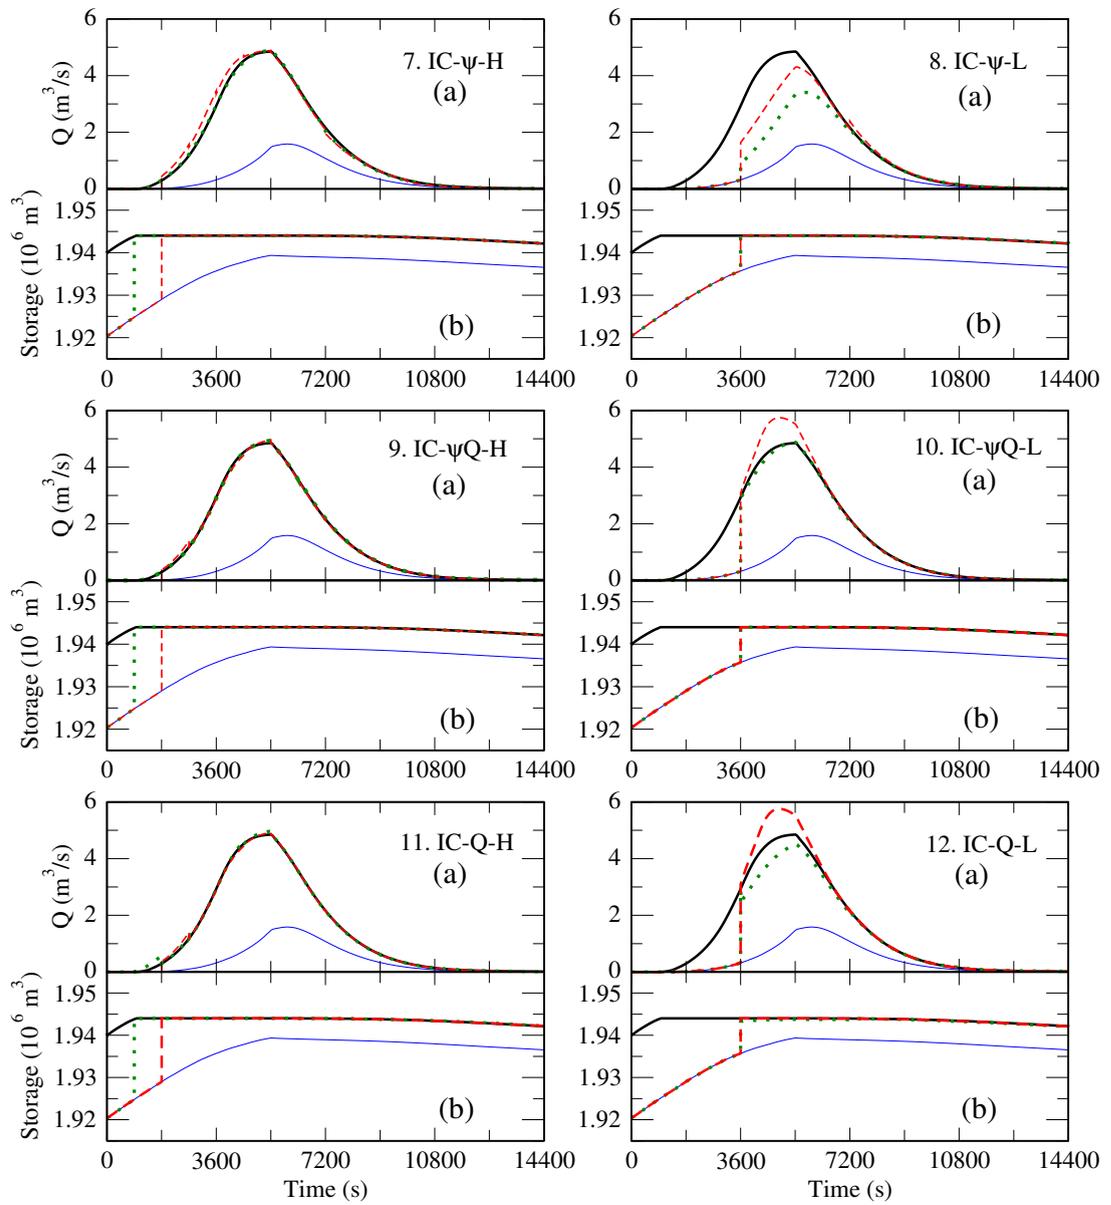
Figure 2.6: Same as Fig. 2.4 but for scenarios 7-12 of Table 2.2 (biased initial conditions).

and only a fraction of the catchment surface contributes to runoff, whose peak is thus 70% smaller than the maximum value of the true run. Considering the high-frequency assimilations (test cases 7, 9, and 11), both EnKF and SIR are able to retrieve the true values of the surface and subsurface state variables with just one update (the first update for the EnKF and the second update for SIR). From that update until the end of the simulation the simulated dynamics correspond in practice to the true run. The particularly good performance of the two filters in these scenarios is due to the fact that the ensemble expected value of the atmospheric forcing corresponds to the atmospheric forcing of the true simulation. Thus, when the biased initial conditions are corrected, the ensemble expected value and the true run practically coincide.

In the low-frequency scenarios (test cases 8, 10, and 12), both EnKF and SIR are able to retrieve the true subsurface water volume after the first update. Looking at the outlet hydrograph, assimilation of pressure head only (Fig. 2.6, panel 8a) causes an underestimation of the hydrograph peak for both EnKF (-1.44 m^3/s with respect to the true run) and SIR (-0.55 m^3/s). The assimilation of both pressure head and streamflow (Fig. 2.6, panel 10a) results in an optimal retrieval of the streamflow for the EnKF simulation and in an overestimation of the hydrograph peak of about 0.90 m^3/s for the SIR simulation. Assimilating only the streamflow at the outlet slightly worsens the EnKF hydrograph (Fig. 2.6, panel 12a) with a peak about 0.34 m^3/s smaller than the true value.

SIR issues

The results of the high-frequency simulations with biased initial conditions (Fig. 2.6, panels 7, 9 and 11) show that EnKF retrieves the true subsurface state at the first update, while SIR skips the first update altogether. This mishap has no further consequences and the algorithm quickly recovers the expected behavior after the first actual assimilation. This behavior is a consequence of an incongruous mismatch between simulated and observed values of all the SIR realizations at the first assimilation time. When a weight associated to a realization is computed using the likelihood function (2.18) or (2.20), the corresponding numerical value can be small if the observations are far from the simulated states, and the weights numerically collapse to zero. If this happens for each realization, the SIR update is undefined and the assimilation cannot be performed. This problem is well known and usually handled by increasing the ensemble size of the SIR simulation or intensifying the uncertainty in the boundary forcing terms to effectively augment the
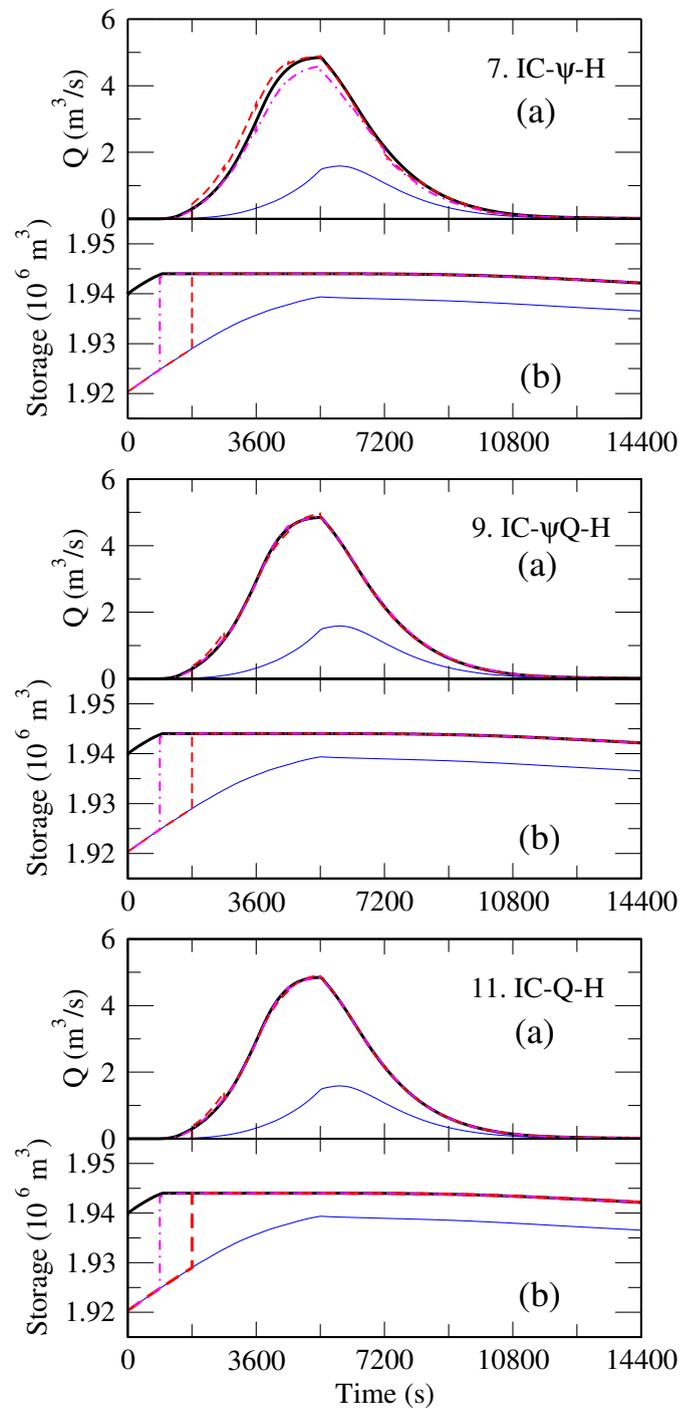
Figure 2.7: Outlet hydrograph (panels a) and subsurface water storage (panels b) for the true solution (black solid line), the ensemble mean of the open loop runs (blue solid line), and the ensemble mean of the standard SIR (red dashed line) and modified SIR (dashed and dotted magenta line) for the scenarios 7, 9, and 11 of Table 2.2 (biased initial conditions), with the same bias applied to the open loop run.
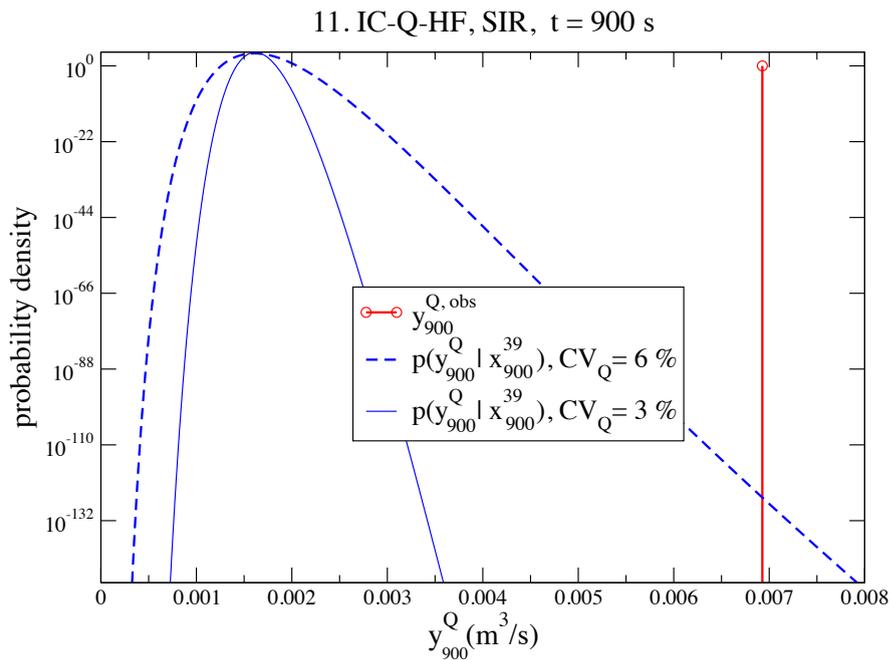
Figure 2.8: Various probability density functions at the first update (time t=900 s) of the SIR method in Scenario 11. The blue lines are the PDFs of streamflow measurements at the outlet given the realization 39 of the ensemble, with i) coefficient of variation CV_Q=3% (solid line) and ii) CV_Q=6% (dotted line). The vertical red line represents the observed value of streamflow. The PDF with CV_Q=3% computed in correspondence of the observation is numerically equal to 0, while the PDF with CV_Q=6% gives a non-zero value.

probability that at least one realization will result close to the observations to yield a large enough likelihood [Simon, 2006].

We tried an alternative approach that, in most of the cases, allowed to overcome this numerical problem in the analysis step, without repeating the forecast step. The idea is to artificially increase the coefficient of variation of the errors associated to the measurements until i) at least one realization has a non-zero weight or ii) the coefficient of variation exceeds a threshold value. In the former case, the approach duplicates the realizations closest to the observations, consistently with the shape of the PDF of the observation errors. In the latter case, the update is not performed and the perturbations to the boundary forcing terms will spread the ensemble in the subsequent forecast step, facilitating the next update (as happened for instance in scenarios 7, 9, and 11 of Fig. 2.6). We denote this method as the modified SIR (MSIR). This procedure is similar to the Kernel estimates for PDFs introduced by Silverman [1986] and to the modification to the Levenberg-Marquardt algorithm proposed by Nowak and Cirpka [2004] for regularization of the Quasi-Linear Geostatistical inversion.

In Fig. 2.7 SIR and MSIR are compared for scenarios 7, 9 and 11. MSIR outperforms the standard SIR at the first update, warranting an optimal retrieval of the true state variables until the end of the simulation, as evidenced by the quicker retrieval of the true subsurface state. To better visualize the differences between standard SIR and MSIR, in Fig. 2.8 we plot $p(y_{900}^{Q} \mid x_{900}^{39})$, i.e., the PDF of streamflow Q at time $t = 900$ s given by the system state corresponding to realization 39. The two curves with solid and dotted lines represent the PDF with coefficient of variation $CV_Q = 3$ % and $CV_Q = 6$ %, respectively. The red line represents the observed streamflow $y_{900}^{Q,obs}$ (the figure refers to scenario 11, where only streamflow is assimilated). In this simulation, the 39^{th} realization is the only realization with a streamflow greater than 0 at $t = 900$ s. We expected that SIR duplicated realization 39, but, as shown in Fig. 2.8, the likelihood function with $CV_Q = 3$ % in correspondence of the observation is numerically 0 (smaller than machine zero). On the other hand, MSIR, using $CV_Q = 6$ %, computes a non-zero weight for realization 39, allowing the system to be updated.

EnKF issues

One of the main drawbacks of EnKF is that the updated state of each realization is computed via a numerical approximation of the Kalman gain, which relies on a hypothesis that is not always satisfied (the Gaussian approximation). Hence the physical consistency

	TIME (s)	BK-ST	NL ITER	N_{ens}	TIME ST
01-EnKF	26146	253	394845	96	193869
01-SIR	22178	6	331618	100	197059
diff	15.1%		16.0%		-1.6%
02-EnKF	20558	1007	307871	83	150829
02-SIR	17544	9	258293	100	157707
diff	14.6%		16.1%		-4.5%
03-EnKF	25255	318	379516	94	186689
03-SIR	22312	4	332227	100	197041
diff	11.7%		12.5%		-5.6%
04-EnKF	20439	195	303701	97	157176
04-SIR	16922	9	246703	100	157707
diff	17.2%		18.8%		-0.3%
05-EnKF	26391	326	398362	95	189690
05-SIR	22413	4	335347	100	197081
diff	15.1%		15.8%		-3.9%
06-EnKF	20767	880	307806	85	150760
06-SIR	16731	9	244471	100	157712
diff	19.4%		20.6%		-4.6 %
07-EnKF	21129	259	339914	96	194903
07-SIR	19347	1	284657	100	197039
diff	8.4%		16.3%		-1.1%
08-EnKF	16246	6	240821	100	157981
08-SIR	14990	6	216122	100	157669
diff	7.7%		10.2%		0.2%
09-EnKF	19093	11	343791	100	197815
09-SIR	16841	1	285818	100	197025
diff	11.8%		16.9%		0.4%
10-EnKF	15630	69	235140	99	157140
10-SIR	14757	7	215230	100	157671
diff	5.6%		8.5%		-0.3%
11-EnKF	18657	384	347906	95	194284
11-SIR	16010	1	283336	100	197044
diff	14.1%		18.6%		-1.4%
12-EnKF	14972	1034	275169	80	149722
12-SIR	13965	7	213785	100	157670
diff	6.7%		22.3%		-5.3%

Table 2.5: Simulation times (TIME) in seconds, number of backsteps (BK-ST), number of nonlinear iterations (NL ITER), final ensemble size (N_{ens}), and number of time steps (TIME ST) for the EnKF and SIR runs in scenarios 1-12.

of the realizations is not always ensured [Chang et al., 2010]. These inconsistencies become critical in the CATHY model especially when the subsurface state of the realizations is not completely saturated. In these cases the strong nonlinearities spatially located in the unsaturated zone do not cope well with the Gaussian approximation. This behavior is reflected by the difficulties encountered by the CATHY nonlinear solver that, in the time steps immediately following an EnKF update, in many cases forces a backstep (a restart of the current time step with a reduced time step size) to ease the convergence of the nonlinear iterations. For some realizations this is not enough and convergence cannot be achieved altogether. In such cases the relevant realization is discarded from the ensemble until the ensemble size drops below a specified value and the simulation halts.

In Table 2.5 we report the simulation times, the number of backsteps, the number of nonlinear iterations, the final ensemble size, and the total number of time steps of the EnKF and SIR simulations for scenarios 1-12. In almost every scenario the number of nonlinear iterations required by EnKF is noticeably larger than that required by SIR, resulting in slower EnKF simulations with respect to SIR. Note that some EnKF realizations do not reach convergence (up to 20 in scenario 12) causing an impoverishment of the ensemble and eventually deteriorating the filter performance. SIR does not display this behavior, since particles that are propagated in time do not loose physical consistency at the update steps as much as EnKF realizations do. In fact, every particle that is replicated by SIR is the numerical solution of the problem, and thus "quasi" physical consistency is guaranteed. In other words, during the SIR resampling step, only the system state variables are duplicated, while each realization preserves its specific set of perturbed time independent soil parameters. This implies that the initial condition for the next time steps is physically consistent not with the actual perturbed dynamics but with a closely related "unperturbed" case. In EnKF, on the other hand, the initial condition is reworked by the Kalman update, without preserving the physical consistency, giving reason to the numerical difficulties described above.

2.6   Conclusions

In this chapter we applied the ensemble Kalman filter (EnKF) and sequential importance resampling (SIR) data assimilation methods to a distributed physically-based hydrological model that couples surface and subsurface flow. Both schemes allow to sequentially incorporate field measurements, such as pressure head, soil moisture, and/or streamflow, into the nonlinear model dynamics. As EnKF and SIR are based on the Monte Carlo

approach, they are suitable for considering different sources of errors in the model, such as uncertainties in soil parameters, atmospheric forcings, and initial conditions.

We analyzed the fundamentals of the two data assimilation schemes in the framework of the filtering probability theory, highlighting the mathematical assumptions at the root of the two filters. The update step in EnKF is performed using the Gaussian approximation, which can result in a suboptimal estimation of the filtering PDF. EnKF algorithm is very attractive for its ease of implementation and its capability to drive the system state towards the observations even for small ensemble sizes. SIR is more robust from a theoretical point of view, as it can be demonstrated that the SIR empirical filtering PDF converges to the true filtering PDF when the ensemble size tends to infinity, regardless of the type of distribution. The principal drawback of SIR with respect to EnKF is that in the update step it duplicates the realizations that are closer to the observations instead of changing the state variable, possibly resulting in low accuracy of the filter if the realizations are far from the observed state. For this reason, EnKF has been so far preferred for applications to distributed models, whereas use of SIR has been limited mostly to conceptual hydrological models.

EnKF and SIR were then compared on the basis of numerical simulations on a three-dimensional synthetic v–catchment with spatially-homogeneous soil parameters. We first performed a sensitivity analysis on the number of realizations to find an appropriate ensemble size ensuring a good compromise between computational effort and estimation accuracy of the two filters. A sensitivity analysis on the decorrelation time length of the atmospheric forcing allowed the determination of a physically-reasonable value of τ sufficient to guarantee a large enough spread of the realizations between the updates. Finally, we compared EnKF and SIR regarding their capability to retrieve the true unknown system state in two series of scenarios characterized by biased forcing terms and biased initial conditions, respectively. Our work leads to the following major conclusions.

1. In all the numerical experiments, using the same relatively small number of realizations (N_{ens}=100), both EnKF and SIR allow to improve significantly the system forecasts with respect to the open loop. With assimilation of streamflow only, EnKF cannot retrieve completely the subsurface pressure heads during the evaporation period, due to the nonlinear dynamics of the unsaturated zone and to the low correlation between the measurements (streamflow at the outlet) and the state variables (distributed pressure heads). In this case, SIR outperforms EnKF as the former duplicates realizations that are closer to the observations and thus naturally

honor also the subsurface system state.

2. Analyzing the computational times and statistics of EnKF simulations, we argue that the Gaussian approximation inherent in the update step of EnKF can render the updated state of the realizations not consistent with the physics of the model. This entails an increase of computational effort for the convergence of the nonlinear numerical solver after each update, with a consequent increase of the total simulation time and sometimes even a loss of some realizations that do not reach convergence. SIR, instead, in the analysis step duplicates solutions of the model equation and thus is not affected by this issue.

3. With small ensemble sizes and biased initial conditions, the SIR updates can be strongly affected by the numerical degradation of the weights associated to each realization, with all the values tending to zero. To overcome this issue, we proposed a modified SIR algorithm (MSIR): during the update, before the normalization of the weights, a control is performed and the update is repeated with an increased standard deviation of the measurement error if all the weights are numerically zero. This modification allows to associate a non-zero weight to the realizations closest to the observations and to perform a correct update. Results from the scenarios with biased initial conditions demonstrate that MSIR is a valid methodology to improve the performance of the SIR algorithm.

Our results show that both EnKF and MSIR can be effectively used as data assimilation algorithms for hydrological simulations with CATHY. Both techniques are ideally suited to take into consideration parameters described by continuous random fields, a topic that will be addressed in future work. We finally stress that a limitation of these methods for applications to parameter identification problems is represented by the large ensemble size required to achieve convergence. For this reason, further efforts will focus toward improvement of the efficiency of the ensemble generation phase, for example by combining DA schemes with model order reduction techniques.

Chapter 3

A reduced order model for the solution of stationary groundwater flow with randomly distributed recharge [1]

3.1 Summary

In this chapter we present a methodology conducive to the application of a Galerkin model order reduction technique, Proper Orthogonal Decomposition (POD), to solve a groundwater flow problem driven by spatially-distributed stochastic forcing terms. Typical applications of POD to reducing time-dependent deterministic partial differential equations (PDEs) involve solving the governing PDE at some observation times (termed snapshots), which are then used in the order reduction of the problem. Here, the application of POD to solve the stochastic flow problem relies on selecting the snapshots in the probability space of the random quantity of interest. This allows casting a standard Monte Carlo (MC) solution of the groundwater flow field into a Reduced Order Monte Carlo (ROMC) framework. We explore the robustness of the ROMC methodology by way of a set of numerical examples involving two-dimensional steady-state groundwater flow taking place within an aquifer of uniform hydraulic properties and subject to a randomly distributed recharge. We analyze the impact of (i) the number of snapshot which are selected from the hydraulic heads probability space, (ii) the associated number of principal

---

[1] The contents of this chapter has been published in Pasetto et al. [2011].

components, and (iii) the key geostatistical parameters describing the heterogeneity of the distributed recharge on the performance of the method. We find that our ROMC scheme can improve significantly the computational efficiency of a standard MC framework while keeping the same degree of accuracy in providing the leading statistical moments (i.e. mean and covariance) as well as the sample probability density of the state variable of interest.

3.2   Introduction

Simulation of groundwater flow in porous media is typically affected by model and parametric uncertainties. In both cases, the groundwater flow problem is conveniently tackled within a probabilistic framework. Parametric uncertainty typically stems from an incomplete knowledge of the distributions of hydraulic parameters, boundary conditions and forcing terms. These are usually modeled as random fields and flow (and possibly transport) processes are governed by stochastic partial differential equations (PDEs) [Dagan and Neuman, 1997].

Solution of a stochastic groundwater flow problem can be achieved in the context of a MC framework. MC methods are based on the idea that it is possible to generate equally likely solutions (or realizations) of the flow problem on the basis of a set of realizations of the distribution of selected system parameters, forcing terms and initial and boundary conditions. Realizations should display a variability reminiscent of that observed in the field and can be conditioned to direct measurements of aquifer properties and other uncertain quantities. Advantages of MC-based approaches include (i) the relatively straightforward implementation of the method by means of standard solvers for deterministic PDEs; (ii) the potential for providing approximations of the probability density function (PDF) of state variables, e.g. hydraulic heads or fluxes, and (iii) the applicability to data assimilation and inverse problems (see, e.g., Camporese et al. [2009a], Hendricks Franssen et al. [2009] and references therein). A key drawback of MC methods is that convergence is sublinear and the number of realizations needed to attain workable accuracy depends on the (space- and eventually time-dependent) moments of interest of the target state variable. Convergence is typically assessed by suitable *a posteriori* local (or global) criteria such as those illustrated by Ballio and Guadagnini [2004]. The large number of MC iterations that are typically needed to ensure stability of the sample moments can lead to prohibitive computational costs in high-dimensional transient systems, especially when the objective is to assess the complete PDF of a groundwater flow state

variable, as required, e.g. in environmental risk assessment practice. A recent comparison amongst three most commonly used MC techniques, i.e., the standard Monte Carlo, Latin Hypercube sampling, and Quasi Monte Carlo has been presented by Zhang et al. [2010]. Application to field scale aquifer systems on the basis of relatively few realizations, as has been the practice to date, may yield plausible representations of reality, which however are random and therefore non-unique and unreliable estimates of the probability distribution (or, alternatively, sample moments) of hydraulic heads and fluxes.

To alleviate the computational burden connected to the evaluation of a large number of independent replicates, the Moment Differential Equations (MDE) approach starts from a stochastic interpretation of the governing flow equations to develop a set of deterministic PDEs that are satisfied by the (ensemble) moments of the state variables. Recursive approximations of these (otherwise non-local) moment equations of steady-state hydraulic heads are presented by Guadagnini and Neuman [1999]. Some features of the approach are reviewed by Winter et al. [2003] and, Zhang [2002]. Limitations of the method include: (a) the system of moment equations is almost never closed so that closure approximations are needed; and (b) only the first two statistical moments, i.e. (ensemble) mean and (co)variance, are typically computed so that a complete description of the PDF of hydraulic heads and/or fluxes is generally not obtained.

Polynomial Chaos expansion (PCE) has been proposed as an efficient non-intrusive technique for the solution of stochastic PDEs in the presence of a limited number of random parameters. The methodology starts from representing the uncertain state variable by a suitable series expansion to determine analytically the statistical moments of a truncated version of the expansion. A general description of the method including a comparison with a standard MC approach can be found in Poles and Lovison [2009]. Examples of application of PCE to groundwater flow problems are provided in Fontaine et al. [2010], and Li and Zhang [2007]. In the case of PDEs with distributed random parameters (e.g. spatially-distributed stochastic recharge or conductivity fields), the large number of random parameters can hamper an efficient implementation of PCE. This problem is usually circumvented by approximating the stochastic parameters via a truncated Karhunen-Loève (KL) representation, so that only a few random variables allow describing the complete random field. A notable problem associated with this technique is that the number of terms to be retained in the KL expansion strongly depends on the degree of spatial coherence of the random field considered (which is governed by parameters such as the correlation scale) so that a large number of terms may be needed in the expansion to provide a proper approximation of the system behavior Zhang and Lu [2004]. Another

difficulty arises form the fact that the distribution of the KL coefficients is analytically known only for particular random fields. For example, the coefficients are Gaussian if the underlying random field is Gaussian [Ghanem and Spanos, 1991]. Assessment of the quality of the approximated moments computed with MDE or PCE is often performed via a detailed comparison against corresponding results obtained by standard MC-based analyses (see, e.g., [Li and Zhang, 2007, Riva et al., 2006] and references therein).

In this chapter we use a Galerkin model order reduction technique to reduce the computational burden associated with the generation of the ensemble of Monte Carlo simulations of a given flow scenario. Model order reduction techniques are traditionally applied to deterministic, typically time-dependent, problems [Kunisch and Volkwein, 2001, Sirovich, 1987]. Our methodology is based on the idea that realizations of hydraulic heads can be efficiently approximated by the solutions of a physically-consistent, low-dimensional system. The key principle of Galerkin model reduction techniques is that the solution of a given PDE can be well approximated by a linear combination of a finite (and small) number of conveniently chosen basis functions. The coefficients of this linear combination can be computed upon solving the reduced system obtained via a Galerkin projection of the PDE. Several different types of Galerkin model order reduction methodologies are available in the literature, depending on the basis functions chosen for the projection. Amongst these, Proper Orthogonal Decomposition (POD) has been widely used in a deterministic setting (see, e.g., Holmes et al. [1996], Kunisch and Volkwein [2002], Li et al. [2011], Luo et al. [2011a,b]). In the POD framework, the basis functions are the eigenfunctions of an integral operator, the kernel of which is rendered by the spatial correlation function of the state variable. In the context of deterministic transient problems the eigenfunctions are often approximated by using the so called snapshot technique which comprises the following steps: (a) the solution of the PDE is computed at some given observation times (snapshots); (b) Principal Component Analysis is applied to the set of the snapshots; and (c) the resulting principal components are then adopted as basis functions for the projection described above. The spatial correlation function is usually estimated via averaging over the calculated temporal snapshots [Henri and Yvon, 2005]. Applications of POD to deterministic groundwater flow problems include, e.g. the works of Vermeulen et al. [2006] where POD is applied in the context of a groundwater flow inverse problem, and Siade et al. [2010] who propose a strategy for a proper selection of the snapshots. McPhee and Yeh [2008] demonstrated that the reduced model is able to reproduce the sensitivities of head with respect to pumping and directly embedded the reduced model in the constraint set of an optimization model for groundwater management. Siade et al. [2012]

used the POD to overcome the computational burden associated with the estimation of the hydraulic conductivity with the quasi-linearization method combined with quadratic programming. Using the concept of POD, Baú [2012] developed a stochastic groundwater flow reduced model which was used for parameter uncertainty analysis in connection with a multiobjective groundwater management problem. POD has also been recently used in the context of stochastic PDEs in conjunction with MC simulations. For example, Wang and Zabaras [2005] use the POD of the heat equation to improve the sampling speed in a heat source reconstruction Bayesian inverse problem. The authors take advantage of the fact that each MC realization entails solving a (time dependent) deterministic equation and adopt the standard snapshot method to reduce the computational cost associated with the solution of each individual ensemble member. The method of snapshots is also employed by Galbally et al. [2010] to construct MC realizations of an advection-diffusion-reaction problem driven by a nonlinear reaction rate upon modeling the two parameters of the rate as uniformly and independent random variables.

As an alternative to POD, the Reduced Basis (RB) method directly uses the snapshots as basis functions for the Galerkin projection (see, e.g., Grepl and Patera [2005], Quarteroni and Rozza [2007]). The ensuing reduced order model is strongly dependent on the choice of the snapshots. As a consequence, the accuracy of RB methods is related to the quality of the snapshots, resulting in a potentially limiting factor when high-dimensional parameter spaces are analyzed. A greedy algorithm has been proposed to circumvent this problem [Bui-Thanh et al., 2007, Grepl and Patera, 2005] and has been applied in Lieberman et al. [2010] to reduce the computational cost of standard Monte Carlo simulations in a groundwater flow inverse problem.

The development of model order reduction techniques in the presence of spatially-distributed random fields is still in its infancy and an exhaustive investigation of their potential is lacking. In the spirit of Galbally et al. [2010], we construct a reduced model on the basis of selected snapshots taken from a few realizations of the governing PDE computed by a standard MC approach. We illustrate the idea and the ensuing ROMC method through a synthetic set-up involving two-dimensional steady state groundwater flow subject to a spatially-distributed random recharge. At the regional scale the aquifer's transmissivity may be regarded as a property associated with slow variations in the horizontal plane and there is field evidence showing that it is characterized by integral scales of the order of kilometers (e.g., Dagan [1989], Dagan et al. [2009], and references therein). Variability of recharge occurs over length scales controlled by local transmissivities and are therefore significantly smaller than those associated with regional transmissivities. A

distinction between local and regional scale transmissivities is presented by, e.g. Copty and Findikakis [2004], Neuman et al. [2007]. It is then appropriate to invoke this scale separation and evaluate the impact of a randomly distributed recharge on the response of a large scale system characterized by a uniform transmissivity. This allows to explore in details the key features of the reduction method we propose in a relatively tractable and still environmentally relevant scenario.

Our key objective is to offer a detailed analysis of:

- the number of snapshots that are extracted from the probability space and adopted in the POD technique;

- the associated number of principal components;

- the effects on the performance of the method of key geostatistical parameters describing the heterogeneity of the distributed recharge.

Results obtained by ROMC are assessed against standard MC iterations, which are taken as a reference, in terms of:

- mean errors in the reproduction of hydraulic head realizations;

- the computational cost;

- the ability of our reduced order model to render the sample estimators of the first two (ensemble) moments and the complete probability distribution of hydraulic heads.

A further and challenging extension of the methodology, which is outside the scope of the present contribution, can comprise the analysis of the effects of a randomly heterogeneous local scale conductivity / transmissivity in the ROMC framework. Our aim is to illustrate how a model reduction technique can lead to relevant savings of the computational time associated with the solution of the flow equation for each MC iteration without resulting in a significant loss of accuracy in the evaluation of the leading (ensemble) moments and the PDF of hydraulic heads.

The chapter is organized as follows. We devote Section 3.3 to a brief analysis of the standard Monte Carlo method in the setting analyzed. Section 3.4 outlines the key ideas underlying a POD-based deterministic model reduction strategy, i.e.:

1. the Karhunen-Loève Representation Theorem;

2. the snapshot technique;

3. the Galerkin projection.

Section 3.4 illustrates how a POD-based model reduction technique can be extended to provide a ROMC solution of the stochastic groundwater flow problem considered. The synthetic flow scenarios and the associated results are presented in Sections 3.5 and 3.6, respectively. Section 3.7 reports our main conclusions.

3.3   Problem setting

We consider steady-state groundwater flow governed by the following stochastic equation,

$$-\nabla \cdot (\mathbf{T}\nabla h(\vec{x})) = f(\vec{x}), \quad \vec{x} \in \Omega \subset \mathbb{R}^2. \tag{3.1}$$

Here, Ω is a two-dimensional compact domain, h represents hydraulic head, $\mathbf{T}$ is a uniform transmissivity, and f is a spatially-distributed (stationary) random recharge. In the standard Monte Carlo (MC) framework one starts by generating an ensemble of N_{ens} realizations of the recharge, $\{f^i(\vec{x})\}_{i=1}^{N_{ens}}$ by means of a suitable method. A number N_{ens} of realizations of hydraulic head, $\{h^i(\vec{x})\}_{i=1}^{N_{ens}}$, are evaluated by solving Eq. (3.1) numerically. In our case we discretize the domain with a regular grid comprising n nodes and t triangular elements, and use the Galerkin Finite Element Method with piecewise linear basis functions. This entails dealing with a high-dimensional linear system for each realization,

$$\mathbf{A}\mathbf{h}^i = \mathbf{b}^i \tag{3.2}$$

where $\mathbf{A}$ is the (sparse) stiffness matrix of dimension $n \times n$, $\mathbf{h}_i$ is the vector of the hydraulic heads calculated at the grid nodes in the i-th realization, and $\mathbf{b}_i$ is the vector representing the i-th realization of the recharge and given boundary conditions.

Let $\mathbf{\Phi}(\vec{x}) = [\phi_1(\vec{x}), \ldots, \phi_n(\vec{x})]^T$ be the vector whose components $\{\phi_j(\vec{x})\}$ are the basis functions of the Galerkin method adopted to solve Eq. (3.1) (superscript T indicates transpose). The value of hydraulic head $h^i(\vec{x})$ at a given location $\vec{x}$ in the i-th MC realization is given by

$$h^i(\vec{x}) = \mathbf{\Phi}^T(\vec{x})\mathbf{h}^i. \tag{3.3}$$

In the following we term Eq. (3.2) as the solution of the Full System Model (FSM).

3.4   The Karhunen-Loève representation theorem

Let $h(\vec{x})$ be a mean square continuous second order random field with continuous covariance function $C_h(\vec{x}, \vec{y})$ between locations $\vec{x}$ and $\vec{y}$,

$$C_h(\vec{x}, \vec{y}) = E[h'(\vec{x})h'(\vec{y})],$$

$E[\cdot]$ representing the expected value and $h'(\vec{x})$ being the field of the (zero-mean) random fluctuations, measuring the deviation of $h(\vec{x})$ from the mean value, i.e., $h'(\mathbf{x}) = h(\vec{x}) - E[h(\vec{x})]$.

Consider the following integral operator, $\mathcal{I}_{C_h}$, with kernel C_h,

$$\mathcal{I}_{C_h}[q](\vec{x}) = \langle C_h(\cdot, \vec{x}), q(\cdot) \rangle_{L_2(\Omega)} = \int_\Omega C_h(\vec{y}, \vec{x}) q(\vec{y}) \, d\vec{y} \ ,$$

where $\langle \cdot, \cdot \rangle_{L_2(\Omega)}$ indicates the inner product in the functional space $L_2(\Omega)$, i.e. the space of square integrable functions over the reference domain, Ω, and $q \in L_2(\Omega)$. $\mathcal{I}_{C_h}$ is a Hilbert-Schmidt operator (see, e.g., [Gohberg and Goldberg, 1981]) and one can show that (a) $\mathcal{I}_{C_h}$ has a countable set of real and positive eigenvalues, $\{\lambda_j\}$, with finite multiplicity, and (b) the sequence of the eigenvalues is nonincreasing and may accumulate only in 0. Let $\{p_j(\vec{x})\}$ be a set of orthonormal eigenfunctions associated with the eigenvalues $\{\lambda_j\}$, i.e.

$$\mathcal{I}_{C_h}[p_j](\vec{x}) = \lambda_j p_j(\vec{x}) \ , \tag{3.4}$$

$$\langle p_i, p_j \rangle_{L_2(\Omega)} = \delta_{i,j} \ , \tag{3.5}$$

$$\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_i \to 0 \ . \tag{3.6}$$

The Karhunen-Loève Representation Theorem states that there is a sequence of zero-mean uncorrelated random variables, $\{a_j\}$ such that the random hydraulic head distribution can be written as

$$h(\vec{x}) = E[h(\vec{x})] + \sum_{j=1}^{\infty} a_j p_j(\vec{x}) \ . \tag{3.7}$$

The random variables a_j have the following properties:

$$E[a_i a_j] = \delta_{ij} \lambda_j \ , \quad E[h'(\vec{x})a_j] = \lambda_j p_j(\vec{x}).$$

The eigenfunctions $\{p_j\}$ in Eq. (3.7) are usually referred to as the Karhunen-Loève modes

or principal components (PC). Let $\hat{h}$ and $\hat{e}$ be defined by:

$$\hat{h} = \sum_{j=1}^{N} \langle h', p_j \rangle_{L_2(\Omega)} p_j \; ; \qquad \hat{e} = h' - \hat{h} \; .$$

The basis formed by the eigenfunctions $\{p_j\}$ is said to be optimal, in the sense that the first N eigenfunctions $\{p_j\}_{j=1}^{N}$ (i.e. the eigenfunctions associated with the largest eigenvalues) minimize the mean square error $E[\|\hat{e}\|_{L_2(\Omega)}^2]$. It is possible to prove that (see, e.g., Ghanem and Spanos [1991]):

$$E[\|\hat{e}\|_{L_2(\Omega)}^2] = \sum_{j=N+1}^{\infty} \lambda_j. \tag{3.8}$$

Eq. (3.8) indicates that our approximation converges to the real solution as $N \to \infty$ because $\lambda_j \to 0$ for $j \to \infty$. The rate with which the eigenvalues converge to zero is not known a priori in our case and requires detailed investigations.

3.4.1 The snapshot technique

We illustrate how the snapshot technique employed in the context of transient deterministic flow scenarios can be adapted to cope with a stochastic problem. We start from $N_{snap} << N_{ens}$ independent realizations of hydraulic head, $\{h^i(\vec{x})\}_{i=1}^{N_{snap}}$. These can be obtained, e.g., by solving the FSM (Eq. (3.2)) on the basis of N_{snap} random realizations of the recharge.

Let $E_{snap}[h(\vec{x})]$ be the expected value of heads calculated with the N_{snap} snapshots, i.e.:

$$E_{snap}[h(\vec{x})] = \frac{1}{N_{snap}} \sum_{i=1}^{N_{snap}} h^i(\vec{x}),$$

and let $\mathbf{z}$ be the vector whose components are $\{h'^i\}$,

$$\mathbf{z}(\vec{x}) = [h'^1(\vec{x}), \dots, h'^{N_{snap}}(\vec{x})]^T.$$

The sample covariance function of hydraulic heads calculated on the basis of the snapshots is given by

$$\tilde{C}(\vec{x}, \vec{y}) = \frac{1}{N_{snap}} \mathbf{z}(\vec{x}) \mathbf{z}^T(\vec{y}).$$

An approximation $\tilde{p}_i(\vec{x})$ of the PCs is obtained by means of Eq. (3.4) upon solving the

eigenvalue problem

$$\left\langle \tilde{C}(\cdot, \vec{x}), \tilde{p}_i(\cdot) \right\rangle_{L_2(\Omega)} = \lambda_i \tilde{p}_i(\vec{x}). \tag{3.9}$$

Eq. (3.9) can be solved numerically by interpolating the snapshots and the PCs from the grid values using piecewise linear functions as in Eq. (3.3), i.e.

$$h^i(\vec{x}) = \mathbf{\Phi}^T(\vec{x})\mathbf{h}^i; \;\; \tilde{p}_i(\vec{x}) = \mathbf{\Phi}^T(\vec{x})\tilde{\mathbf{p}}_i; \;\; E_{snap}[h(\vec{x})] = \mathbf{\Phi}^T(\vec{x})E_{snap}[\mathbf{h}] \,, \tag{3.10}$$

and solving the algebraic eigenvalue problem

$$\frac{1}{N_{snap}}\mathbf{X}\mathbf{X}^T\mathbf{M}\tilde{\mathbf{p}}_i = \lambda_i\tilde{\mathbf{p}}_i \tag{3.11}$$

where $\mathbf{X} = [\mathbf{h}'^1, \ldots, \mathbf{h}'^{N_{snap}}]$, $\mathbf{h}'^j = \mathbf{h}^j - E_{snap}[\mathbf{h}]$, and $\mathbf{M}$ is the lumped mass matrix that accounts for the discretization of the integral operator in Eq. (3.9). The matrix $\frac{1}{N_{snap}}\mathbf{X}\mathbf{X}^T\mathbf{M}$ is a $n \times n$ full (i.e. not sparse) matrix. To reduce the computational effort the eigenvalue problem is solved for the matrix $\frac{1}{N_{snap}}\mathbf{X}^T\mathbf{M}\mathbf{X}$ (dimension $N_{snap} \times N_{snap}$), i.e.

$$\frac{1}{N_{snap}}\mathbf{X}^T\mathbf{M}\mathbf{X}\boldsymbol{\psi}_i = \lambda_i\boldsymbol{\psi}_i \tag{3.12}$$

whose eigenvalues $\{\lambda_i\}$ coincide with those of the matrix $\frac{1}{N_{snap}}\mathbf{X}\mathbf{X}^T\mathbf{M}$ and whose eigenvectors $\boldsymbol{\psi}_i$ are related to the eigenvectors $\tilde{\mathbf{p}}_i$ by:

$$\tilde{\mathbf{p}}_i = \frac{1}{\sqrt{\lambda_i N_{snap}}}\mathbf{X}\boldsymbol{\psi}_i. \tag{3.13}$$

3.4.2   Galerkin model reduction methods

The dimension of Eq. (3.2) can be reduced by the Galerkin projection. The (random) field of hydraulic head is approximated by means of a truncated series in a Karhunen-Loève representation:

$$h(\vec{x}) \approx \tilde{h}(\vec{x}) = E[h(\vec{x})] + \sum_{j=1}^{N_{PC}} a_j \tilde{p}_j(\vec{x}). \tag{3.14}$$

The choice of the number of PCs, N_{PC}, to be used in the approximation relies on the recognition that Eq. (3.8) still holds for the snapshots, i.e.

$$E_{snap}[\|h'^i(\vec{x}) - \sum_{j=1}^{N_{PC}} \left\langle h'^i, \tilde{p}_j \right\rangle_{L_2(\Omega)} p_j(\vec{x})\|^2_{L_2(\Omega)}] = \sum_{j=N_{PC}+1}^{N_{snap}} \lambda_j. \tag{3.15}$$

A typical criterion which is adopted for the choice of N_{PC} is

$$\frac{\sum_{j=1}^{N_{PC}} \lambda_j}{\sum_{j=1}^{N_{snap}} \lambda_j} > 99.99\% \, .$$

The approximated hydraulic head associated with a given realization of the random recharge, $f^i(\vec{x})$, can be computed by replacing $\tilde{h}^i$ into Eq. (3.1) and imposing that the residual

$$r^i(\vec{x}) = \nabla \cdot (\mathbf{T}(\vec{x})\nabla\tilde{h}^i(\vec{x})) + f^i(\vec{x})$$

be orthogonal to all the PCs, $\{\tilde{p}_j\}$. This is tantamount to state that the PCs are chosen as basis functions for the Galerkin Method. The following linear system of order $N_{PC} \times N_{PC}$ can then be written:

$$\sum_{j=1}^{N_{PC}} \langle \tilde{p}_k, -\nabla \cdot (\mathbf{T}\nabla\tilde{p}_j)\rangle_{L_2(\Omega)}\, a_j^i = \left\langle \tilde{p}_k, f'^i \right\rangle_{L_2(\Omega)} \quad , \quad k = 1, \dots, N_{PC} \, , \tag{3.16}$$

where $f'^i = f^i - E[f]$. The realizations a_j^i of the random variables a_j are the unknown terms in Eq. (3.16). Note that here we have taken full advantage of the superposition principle.

Interpolating the PCs on the grid nodes using piecewise linear functions as in Eq. (3.10), allows rewriting Eqs. (3.14) and (3.16) as

$$\mathbf{h}^i \approx \tilde{\mathbf{h}}^i = E[\mathbf{h}] + \sum_{j=1}^{N_{PC}} a_j^i \tilde{\mathbf{p}}_j = E[\mathbf{h}] + \tilde{\mathbf{P}}\,\mathbf{a}^i, \tag{3.17}$$

$$\tilde{\mathbf{P}}^T \mathbf{A}\tilde{\mathbf{P}}\,\mathbf{a}^i = \tilde{\mathbf{P}}^T \mathbf{b}'^i. \tag{3.18}$$

Here, $\tilde{\mathbf{P}}$ is the matrix whose columns are the vectors $\tilde{\mathbf{p}}_j$, i.e. $\tilde{\mathbf{P}} = [\tilde{\mathbf{p}}_1 \dots \tilde{\mathbf{p}}_{N_{PC}}]$; $\mathbf{b}'^i = \mathbf{b}^i - E[\mathbf{f}]$.

3.5 Implementation of the ROMC strategy

The implementation of POD in the Monte Carlo framework comprises two key steps: (i) an initialization step for the computation of the principal components and the reduced system matrix; and (ii) the computation of a desired ensemble of system realizations via the reduced model.

3.5.1   Initialization of the reduced model

Initialization of the reduced model is performed according to the following steps.

1. **Generation of realizations of the random forcing terms**. N_{snap} stochastic realizations of recharge $\{\mathbf{f}^j\}$, are generated by means of a suitable random field generation procedure. The computational cost of this phase is common to the standard MC method and is not considered here.

2. **Computation of the snapshots in probability space**. N_{snap} realizations of hydraulic heads $\{\mathbf{h}^j\}$ are computed by solving Eq. (3.2) on the basis of the generated N_{snap} realizations of recharge. The computational cost using a direct solver is $O(N_{snap} \, n^2)$.

3. **Computation of the principal components**. The collected snapshots are used to compute $E_{snap}[\mathbf{h}]$, the matrix $\mathbf{X}$ and the first N_{PC} eigenvectors $\boldsymbol{\psi}_i$ of the matrix $\frac{1}{N_{snap}}\mathbf{X}^T\mathbf{M}\mathbf{X}$ according to Eq. (3.12). The eigenvectors $\{\tilde{\mathbf{p}}_j\}$ are then computed according to Eq. (3.13). These are collected in the matrix $\tilde{\mathbf{P}}$ and form the set of PCs to be used in the model reduction scheme. The computational cost is $O(n\,N_{snap}^2 + N_{snap}^4)$.

4. **Computation of the reduced model system matrix.** Matrix $\tilde{\mathbf{P}}^T\mathbf{A}\tilde{\mathbf{P}}$ has dimensions $N_{PC} \times N_{PC}$ and can be factored by LU decomposition

$$\tilde{\mathbf{P}}^T\mathbf{A}\tilde{\mathbf{P}} = \mathbf{L}\mathbf{U}.$$

   The computational cost associated with this step is $O(s\,n\,N_{PC} + n\,N_{PC}^2 + N_{PC}^3)$, where s is the average number of non zero elements of $\mathbf{A}$.

3.5.2   Application of the reduced model

We are now in a position to compute a number $N_{res} = N_{ens} - N_{snap}$ of realizations of hydraulic heads via our reduced model. In order to accomplish this, we start by generating N_{res} stochastic realizations of recharge. Each one of these is then employed within the procedure detailed in the following.

1. **Forcing projection**. The term $\tilde{\mathbf{P}}^T\mathbf{b}'$ is computed. The associated computational cost is $O(n\,N_{PC})$.

2. **Solution of the linear system**. We solve the linear system given by Eq. (3.18) to compute $\mathbf{a}^i$:

$$\mathbf{LU}\mathbf{a}^i = \tilde{\mathbf{P}}^T\mathbf{b}'^i.$$

The associated computational cost is $O(N_{PC}^2)$.

3. **Computation of hydraulic head realization**. The (approximated) head realization, $\tilde{\mathbf{h}}^i$, is computed as in Eq. (3.17):

$$\tilde{\mathbf{h}}^i = \tilde{\mathbf{P}}\mathbf{a}^i$$

with a computational cost $O(n\,N_{PC})$.

The total computational cost of the ROMC is then

$$O(n^3 + N_{snap}\,n^2 + n\,N_{snap}^2 + N_{snap}^4 + s\,n\,N_{PC}^3 + N_{res}\,(2\,n\,N_{PC} + N_{PC}^2)),$$

which is to be compared against the corresponding cost associated with a standard MC solution of the FSM (Eq. (3.2)), which is $O(N_{ens}\,n^2 + n^3)$ (without considering the forcing function generation step). For example, if we take $n = 2000$, $N_{ens} = 10000$, $s = 7$, $N_{PC} = 10$ and $N_{snap} = 100$, we obtain $O(8.1 \times 10^{10})$ and $O(1.7 \times 10^{10})$ for standard MC and ROMC, respectively. We note that the difference increases dramatically as n increases or in three-dimensional problems where a direct factorization is not possible.

3.6   Numerical results

We consider steady-state groundwater flow taking place in a two-dimensional domain with uniform isotropic transmissivity, $T = 12 \cdot 10^{-4}$ m/s, and spatially-distributed random recharge. Fig. 3.1 depicts the flow domain and boundary conditions. The latter have been adapted from Hendricks Franssen et al. [2009]. We discretize the domain with a regular grid comprised of 49×50 square cells of 100×100 m. Each cell is divided in two triangles, for a total of $n = 2550$ nodes and $t = 4900$ triangular elements.

A randomly distributed recharge, f, is imposed on the domain and is modeled as a second-order stationary random field, characterized by a log-Normal distribution with uniform mean, $\mu_f = 362.912$ mm/y, and exponential isotropic spatial correlation function,

$$C_f(\mathbf{r}) = \sigma_f^2\,e^{-r/\tau},$$
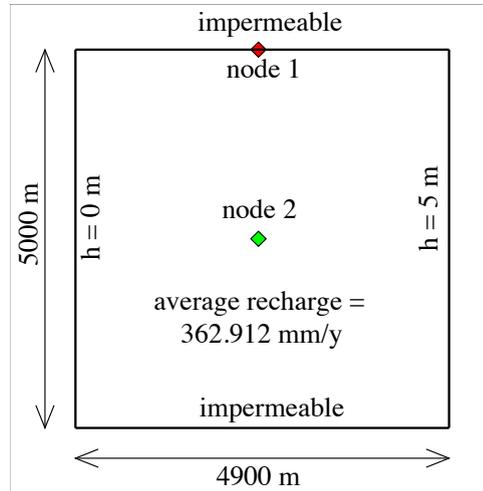
Figure 3.1: Sketch of the two-dimensional domain adopted in the computational examples.

where $r = \|\mathbf{x} - \mathbf{y}\|_2$ and σ_f^2 and τ are the variance and the correlation scale of f, respectively. We consider 4 test cases, characterized by different values of the coefficient of variation (CV) and correlation scales of f:

- TC1: $CV = 100\%$ and $\tau = 500\,m$;

- TC2: $CV = 200\%$ and $\tau = 500\,m$;

- TC3: $CV = 100\%$ and $\tau = 2500\,m$;

- TC4: $CV = 200\%$ and $\tau = 2500\,m$.

An additional test case (TC5) associated with a more refined spatial discretization grid (98×100 cells of 50×50 m), $CV = 100\%$ and $\tau = 500\,m$, is then considered, in order to assess the effect of the spatial resolution of the correlation function, $C_f(\vec{r})$, on the performance of the ROMC approach.

The sensitivity of our ROMC approach to the number of snapshots, N_{snap}, extracted from the head probability space and the associated number of principal components, N_{PC}, is investigated by considering all possible combinations of $N_{snap} = \{50, 100, 200, 500, 1000, 2000\}$ and $N_{PC} = \{5, 10, 20, 30, 40, 50\}$ to construct the reduced order model. Comparison between the standard MC and ROMC approaches is performed on the basis of the results obtained with $N_{ens} = 10000$ realizations. Realizations of f are generated by means of the random field generator HYDRO_GEN [Bellin and Rubin, 1996].
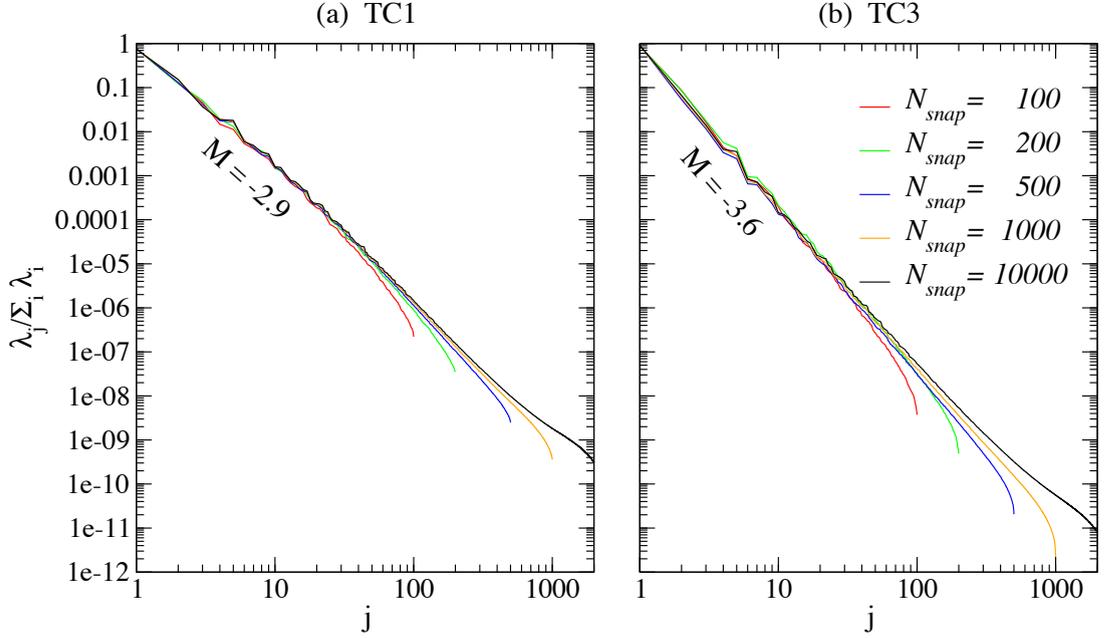
Figure 3.2: Eigenvalues of the operator in Eq. (3.12) normalized by the eigenvalues sum for different values of N_{snap} and for (a) TC1 and (b) TC3.

3.6.1   Eigenvalues

According to Eq. (3.14) the number of principal components required to obtain a good approximation of the second order random field h is strictly related to the eigenvalues λ_i of the operator in Eq. (3.12). Eq. (3.15) indicates that a faster convergence to zero of the eigenvalues implies that a smaller number of terms in the KL representation (Eq. (3.14)) is needed to obtain a given accuracy in the head approximation using the selected snapshots.

Fig. 3.2 depicts the strength of the eigenvalues, λ_j, computed for different values of N_{snap} for TC1 (Fig. 3.2.a) and TC3 (Fig. 3.2.b). Eigenvalues are normalized with respect to the total eigenvalues sum for each N_{snap} selected.

We note that the first 5 and 30 eigenvalues in TC1 account for approximately 90% and more than 99.9% of the total sum of the eigenvalues, respectively. The eigenvalues computed in TC3 display a sharper decay than that observed for TC1. According to Eq. (3.15), this implies that, in principle, a given number of PCs will result in a larger mean square error in TC1 than in TC3. We associate this result with the large correlation scale adopted for f in TC3. Realizations of TC3 display a larger degree of spatial correlation than those of TC1 and this leads to smoother single-realization distributions of f which can be approximated by means of only a few PCs.

In the limiting case when $\tau \to \infty$ the recharge is uniform (but unknown) in each realization, and head fields can be evaluated by simply solving the FSM for one realization and then using the superposition principle, i.e. using only one principal component. The power law decrease displayed in Fig. 3.2 has been observed for all test cases analyzed and imbues us with confidence that using a limited number of eigenvalues can result in a good approximation of the head field, at least for the selected snapshots.

3.6.2   Principal components approximation

There is no general criterion according to which one can (a) provide an estimate of how well the PCs, $\tilde{p}(\vec{x})$, of Eq. (3.14) approximate the actual PCs p_i of the system or (b) select the optimal number of snapshots to consider in the analysis. Fig. 3.3 compares spatial distribution of 5 principal components, $\tilde{p}_i$ $(i = 2, 3, ..., 6)$ obtained for TC1 with $N_{snap} = 100, 200$. The scenario corresponding to $N_{snap} = N_{ens} = 10000$ is reported as a reference against which one can assess the appropriateness of selecting a limited number of snapshots. These results suggest that limiting our sampling in the head probability space to 100 or 200 snapshots insures obtaining sufficiently well approximated distributions of the first PCs. Results of similar quality are obtained for all test cases analyzed.

3.6.3   Analysis of ROMC realizations

A sufficient condition for the ROMC method to reproduce the empirical hydraulic head PDF obtained from a standard MC method is that each head field obtained via ROMC well approximates the corresponding field obtained by standard MC solution of the FSM on the basis of the same random input recharge. Let ρ^i be the relative error between realizations h^i and $\tilde{h}^i$, computed using the FSM and the reduced model, respectively:

$$\rho^i = \frac{\|h^i - \tilde{h}^i\|_{L_2(\Omega)}}{\|h^i\|_{L_2(\Omega)}},$$

and let $E_{res}[\rho]$ be the mean value of ρ^i:

$$E_{res}[\rho] = \frac{1}{N_{res}} \sum_{i=N_{snap}+1}^{N_{ens}} \rho^i.$$

Note that $E_{res}[\rho]$ is computed only over N_{res} realizations, because the snapshots are realizations that coincide for both the standard MC and the ROMC methods. Fig. 3.4
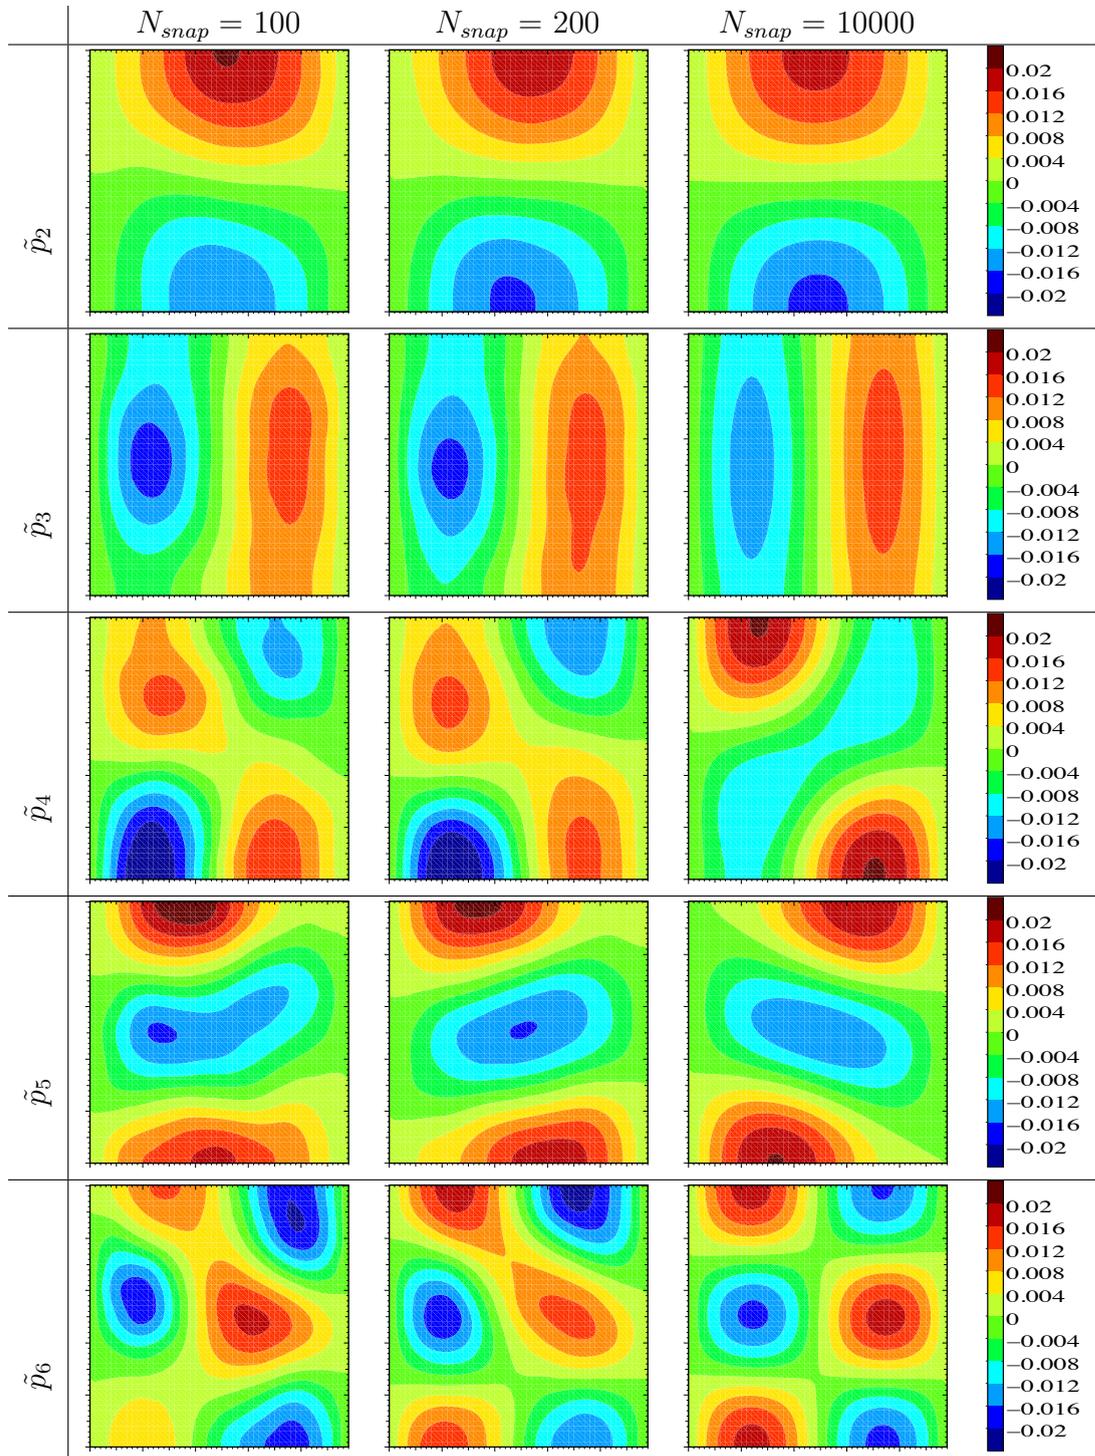
Figure 3.3: Comparison of the spatial distribution within the domain depicted in Fig. 3.1 of 5 principal components, $\tilde{p}_i$ $(i = 2, 3, ..., 6)$, obtained with $N_{snap} = 100$, 200 and 10000. Results associated with $N_{snap} = 10000$ are taken as reference.
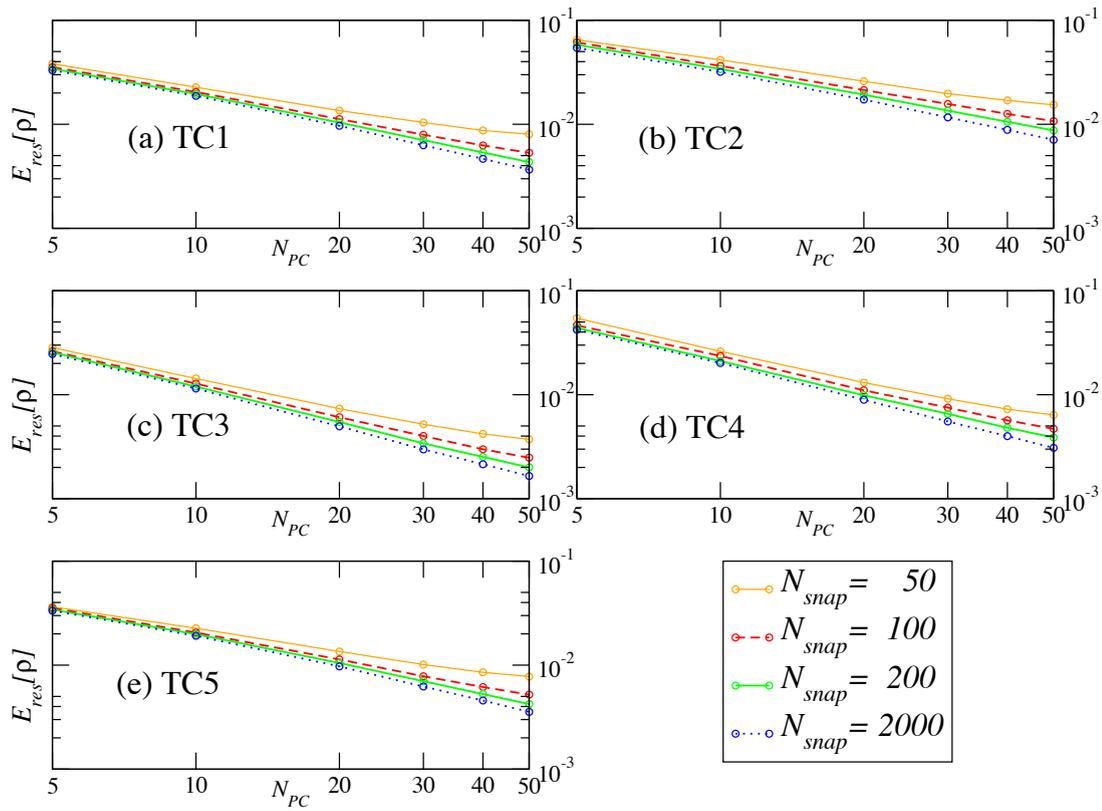
Figure 3.4: Mean value of the relative error, $E_{res}[\rho]$, between realizations computed using the FSM and the reduced model: (a) TC1, (b) TC2, (c) TC3, (d) TC4 and (e) TC5. Model reduction is achieved by $N_{PC} = 5, 10, 20, 30, 40$ or 50 obtained from $N_{snap} = 50$, $100, 200$ or 2000.

FSM	$T_m = 7.118\text{E-}04$		$T_{ens} = 7.1$			
ROMC	$N_{PC} = 5$		$N_{PC} = 10$		$N_{PC} = 20$		$N_{PC} = 30$
T_m	0.428E-04		0.630E-04		1.035E-04		1.451E-04
N_{snap}	T_{in}	T_{ens}	T_{in}	T_{ens}	T_{in}	T_{ens}	T_{in}	T_{ens}
100	0.2	0.7	0.4	1.1	0.7	1.8	1.1	2.6
200	0.3	0.9	0.5	1.2	0.8	2.0	1.2	2.7
500	1.3	2.0	1.4	2.4	1.8	3.1	2.1	3.9
1000	6.9	8.0	7.1	8.4	7.4	9.0	7.8	9.8
2000	52.5	54.2	52.7	54.7	53.1	55.4	53.2	55.8

Table 3.1: Comparison between simulation times (seconds) of the standard MC and ROMC in computing $N_{ens} = 10000$ realizations of head for TC1. Here, T_m is the average time required for the calculation of one realization of head; T_{in} is the time required for the initialization of the reduced model; T_{ens} is the total simulation time, including T_{in}. The FSM (Eq. (3.2)) is solved upon factorizing only once the matrix **A** and using the factorized matrix to compute the remaining head realizations.

compares $E_{res}[\rho]$ for all test cases, when the model is reduced by using $N_{PC} = 5$, 10, 20, 30, 40 or 50 computed on the basis of $N_{snap} = 50$, 100, 200, or 2000. Our results suggest that increasing N_{snap} does not affect significantly $E_{res}[\rho]$, while $E_{res}[\rho]$ displays a power law dependence on N_{PC} with an exponent varying between 0.84 and 1, showing approximately linear convergence. Comparing Figs. 3.4.a and 3.4.b against Figs. 3.4.c and 3.4.d, respectively, one can observe that, consistently with the results of Section 3.6.1, the reduced model based on a given number of PCs yields the highest accuracy for the largest integral scale of the forcing term. Comparison of Figs. 3.4.a and 3.4.c against Figs. 3.4.b and 3.4.d, respectively, suggests that the mean relative error on the realizations slightly increases with the CV of the forcing term, indicating that the performance of our ROMC tends to deteriorate in the presence of very large heterogeneity contrasts in the spatial distribution of f. Finally, Figs. 3.4.a and 3.4.e indicate that the accuracy of the ROMC with respect to the FSM-based MC simulations is not affected by the adopted degrees of descriptive detail of the short-lag spatial correlation of f.

3.6.4 Computational cost

Tables 3.1 and 3.2 report the details of the computational costs associated with the ROMC and the standard MC approaches involved in the solution of the full set of $N_{ens} = 10000$ MC realizations for TC1 and TC5, respectively. Because of the small number of grid nodes involved, we solve the FSM (Eq. (3.2)) by sparse factorization of matrix A, to

FSM	$T_m = 7.63\text{E-}02$		$T_{ens} = 768.6$					
ROMC	$N_{PC} = 5$		$N_{PC} = 10$		$N_{PC} = 20$		$N_{PC} = 30$	
T_m	2.75E-04		2.94E-04		3.42E-04		3.87E-04	
N_{snap}	T_{in}	T_{ens}	T_{in}	T_{ens}	T_{in}	T_{ens}	T_{in}	T_{ens}
100	0.9	11.5	1.6	12.3	2.4	13.6	2.8	14.5
200	1.0	19.6	1.7	20.2	2.4	21.2	3.1	22.5
500	2.2	43.3	5.5	44.0	3.6	45.7	4.3	46.6
1000	7.4	86.8	8.0	87.5	8.7	88.8	9.4	89.8
2000	31.7	187.6	32.4	188.4	33.1	189.4	33.8	190.8

Table 3.2: Comparison between simulation times (seconds) of the standard MC and ROMC in computing $N_{ens} = 10000$ realizations of head for TC5. Here, T_m is the average time required for the calculation of one realization of head; T_{in} is the time required for the initialization of the reduced model; T_{ens} is the total simulation time, including T_{in}. The FSM (Eq. (3.2)) is solved using the preconditioned conjugate gradient method with the incomplete Cholesky preconditioner. The preconditioner is computed only once, at the beginning of the simulation.

minimize the computational cost. This notwithstanding, comparison between the average times, (T_m), needed for the computation of a single head realization in Table 3.1 clearly shows the benefits associated with the ROMC. Adopting our approach over a standard MC procedure results in a time-saving factor of more than 10 or 6 when 5 or 30 PCs are used, respectively.

In the tables, T_{in} indicates the computational time required for the initialization of the reduced model, i.e. the computational time needed to compute the principal components and the reduced system matrix. The results of Table 3.1 show that T_{in} increase quickly when $N_{snap} > 500$, thus hampering the cost-effectiveness of ROMC for large values of N_{snap}.

We solve the FSM for TC5 by means of the conjugate gradient method with the incomplete Cholesky preconditioner. The values of T_m reported in Table 3.2 demonstrates that in this case the reduced model is about 100 times more efficient than the FSM. Note that the most expensive operation in the application of the ROMC to TC5 is not the initialization of the reduced model, but the computation of the snapshots using FSM. This suggests that a tradeoff between the accuracy of the approximated PCs and the number of snapshots is required to balance the accuracy and the computational efficiency of the ROMC.
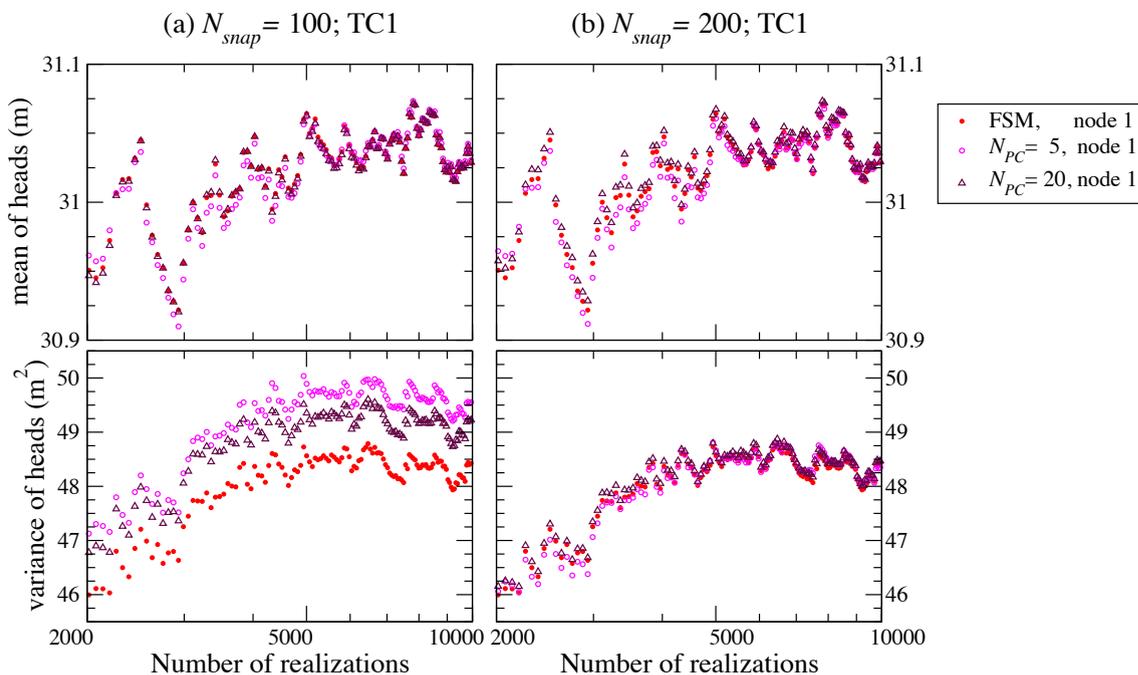
Figure 3.5: Sample mean and variance of the hydraulic head at the reference node 1 of Fig. 3.1, as obtained from the ensemble of realizations calculated with the FSM or reducing the model with $N_{PC} = 5, 20$; results from $N_{snap} = 100$ (a), 200 (b) are reported.

3.6.5 Analysis of leading moments of head PDF

Here, we assess the ability of the ROMC model to reproduce the leading statistical moments (mean and co-variance) as well as the empirical PDF of hydraulic heads by comparing the results of ROMC against those obtained by means of the application of the standard Monte Carlo approach based on the solution of the FSM. By way of illustration, we discuss the (ensemble) results for TC1 at the two reference nodes highlighted in Fig. 3.1. Node 1 in Fig. 3.1 is associated with the largest computed head variance, while Node 2 is located at the domain center, far away from the aquifer boundaries. Figs. 3.5 and 3.6 respectively illustrate the mean and the variance of hydraulic heads at nodes 1 and 2 calculated by the ROMC and FSM a function of the number of realizations. With reference to the ROMC, the figures report results associated with 5 and 20 PCs, obtained from (a) 100 and (b) 200 snapshots. These results demonstrate that the ROMC constructed by 5 PCs and 100 snapshots allows to satisfactorily retrieve the mean head field obtained with the standard MC method. Using 20 PCs and 200 snapshots allows for a satisfactory reproduction of both mean and variance head fields.

The empirical distributions of heads obtained with the standard MC and ROMC
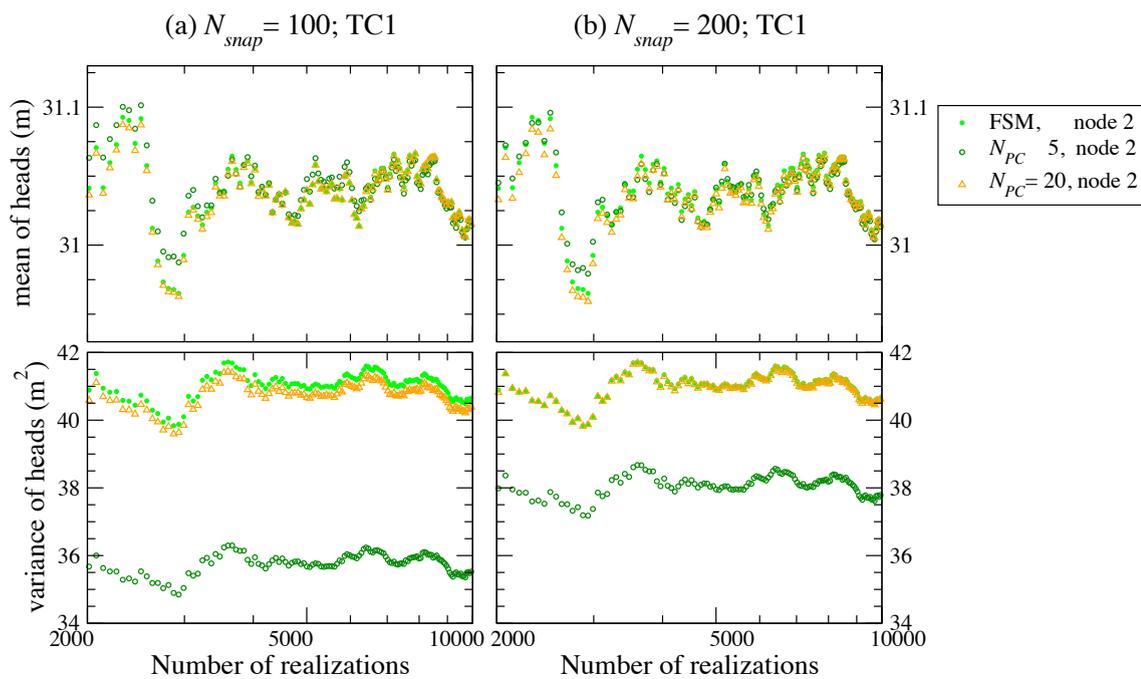
Figure 3.6: Sample mean and variance of the hydraulic head at the reference node 2 of Fig. 3.1, as obtained from the ensemble of realizations calculated with the FSM or reducing the model with $N_{PC} = 5, 20$; results from $N_{snap} = 100$ (a), 200 (b) are reported.

p-values of the U-test								
	$N_{PC} = 5$		$N_{PC} = 10$		$N_{PC} = 20$		$N_{PC} = 30$	
N_{snap}	node 1	node 2	node 1	node 2	node 1	node 2	node 1	node 2
100	0.948	0.229	0.874	0.685	0.875	0.948	0.937	0.961
200	0.875	0.489	0.960	0.865	0.993	0.998	0.987	0.998

Table 3.3: p-values obtained from the U-test, comparing the empirical head distributions at node 1 and node 2 obtained by the standard MC method based on the solution of the FSM and our ROMC with $N_{PC} = 5$, 10, 20 or 30, and $N_{snap} = 100$ or 200.

(a)   (b)   (c)
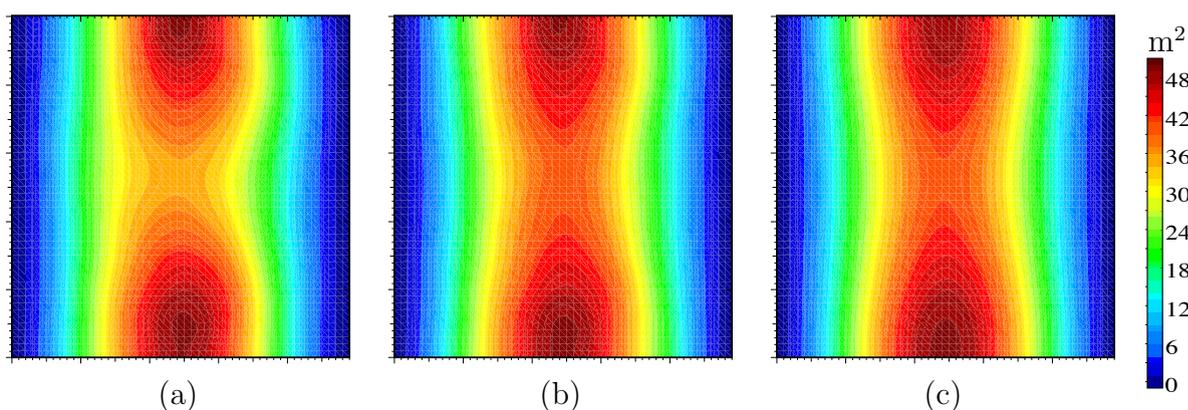
Figure 3.7: Spatial distribution of head variance within the domain depicted in Fig. 3.1 for test case TC1; results for ROMC with $N_{snap} = 100$ and $N_{PC} = 5$ (a), 20 (b); (c) $N_{ens} = 10000$ standard MC simulations of the FSM.

methods at the two nodes are compared by using the Wilcoxon rank-sum test (also called Mann-Whitney U-test [Mann and Whitney, 1947, Wilcoxon, 1945]). The U-test is a non-parametric statistical hypothesis test which allows assessing whether the empirical distributions associated with two independent set of samples are extracted from the same population. Table 3.3 reports the p-values obtained by the application of the U-test. On the basis of these results one can always accept the hypothesis that the empirical distributions provided by the two MC strategies are indistinguishable for both nodes.

Very similar behaviors have been observed for all the remaining cases analyzed.

Fig. 3.7 shows the spatial distributions of the sample variance of heads computed with the ROMC with 5 PCs and 100 snapshots (Fig. 3.7.a) or 20 PCs and 100 snapshots (Fig. 3.7.b), and the standard MC solution of Eq. 3.2 (Fig. 3.7.c). The depiction is complemented by Fig. 3.8, which reports the head covariance referred to the domain center (node 2 in Fig. 3.1) computed by the ROMC with 5 PCs and 100 snapshots (Fig. 3.8.a), or 20 PCs and 100 snapshots (Fig. 3.8.b), and the standard MC solution of Eq. (3.2) (Fig. 3.8.c). These results support our earlier observation that adopting 20 PCs appears
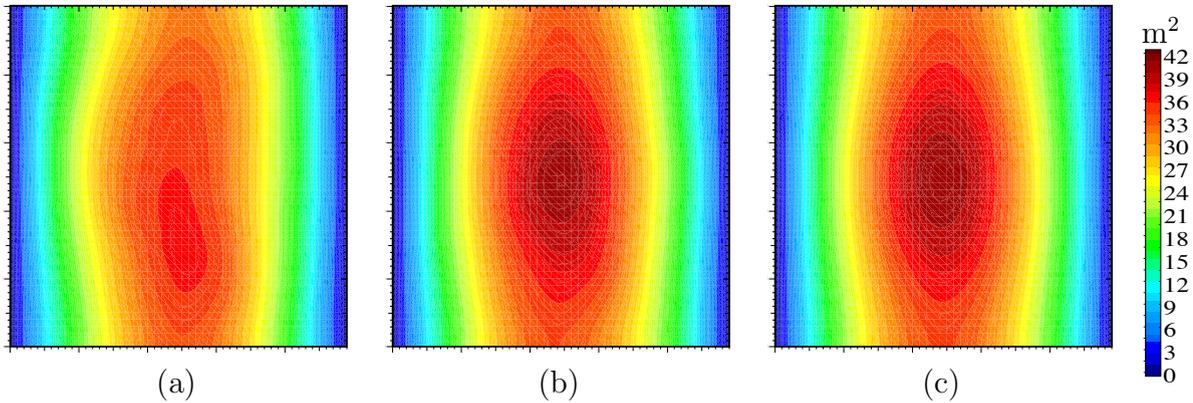
Figure 3.8: Spatial distribution of heads covariance within the domain depicted in Fig. 3.1 (the reference point for covariance calculation is node 2 in TC1); results for ROMC with $N_{snap} = 100$ and $N_{PC} = 5$ (a), 20 (b); (c) $N_{ens} = 10000$ standard MC simulations of the FSM.

to provide an accurate description of the spatial distribution of the second-order moment of the head field in the scenarios investigated. They are also consistent with the p-values rendered by the U-test.

The quantiles of the PDF calculated on the basis of the FSM and ROMC for different values of N_{PC} and N_{snap} provide a valuable quantitative information. We recall that the first and ninth 10-quantiles of the hydraulic heads distribution at location $\vec{x}$, indicated as $q_1(\vec{x})$ and $q_9(\vec{x})$, respectively, are defined as

$$P\left(h\left(\vec{x}\right) < q_1\left(\vec{x}\right)\right) = 10\% \ , \quad P\left(h\left(\vec{x}\right) < q_9\left(\vec{x}\right)\right) = 90\% \ ,$$

P indicating probability. Fig. 3.9 reports $q_1(\vec{x})$ along a cross-section including nodes 1 and 2 of Fig. 3.1 and calculated on the basis of (a) the complete set of N_{ens} realizations of the FSM, (b) the snapshots, and (c) the ROMC constructed by $N_{PC} = 5, 20, 50$ and $N_{snap} = 50, 100, 200, 500, 2000$. Fig. 3.10 depicts the corresponding results associated with $q_9(\vec{x})$. it is apparent that the set of N_{snap} Monte Carlo simulations does not provide an accurate estimate of the spatial distribution of the analyzed quantiles. The N_{res} realizations computed with the reduced model drastically improve the results. The ROMC constructed by 20 PCs and 200 snapshots accurately reproduces both the first and the ninth 10-quantiles obtained using the FSM. Increasing the number of snapshots or the number of principal components does not seem to lead to further significant improvements. Similar results have been obtained for the remaining test cases analyzed.

We analyze the ability of ROMC to retrieve the entire spatial distribution of the variances of the hydraulic heads for every test case by evaluating the relative root mean
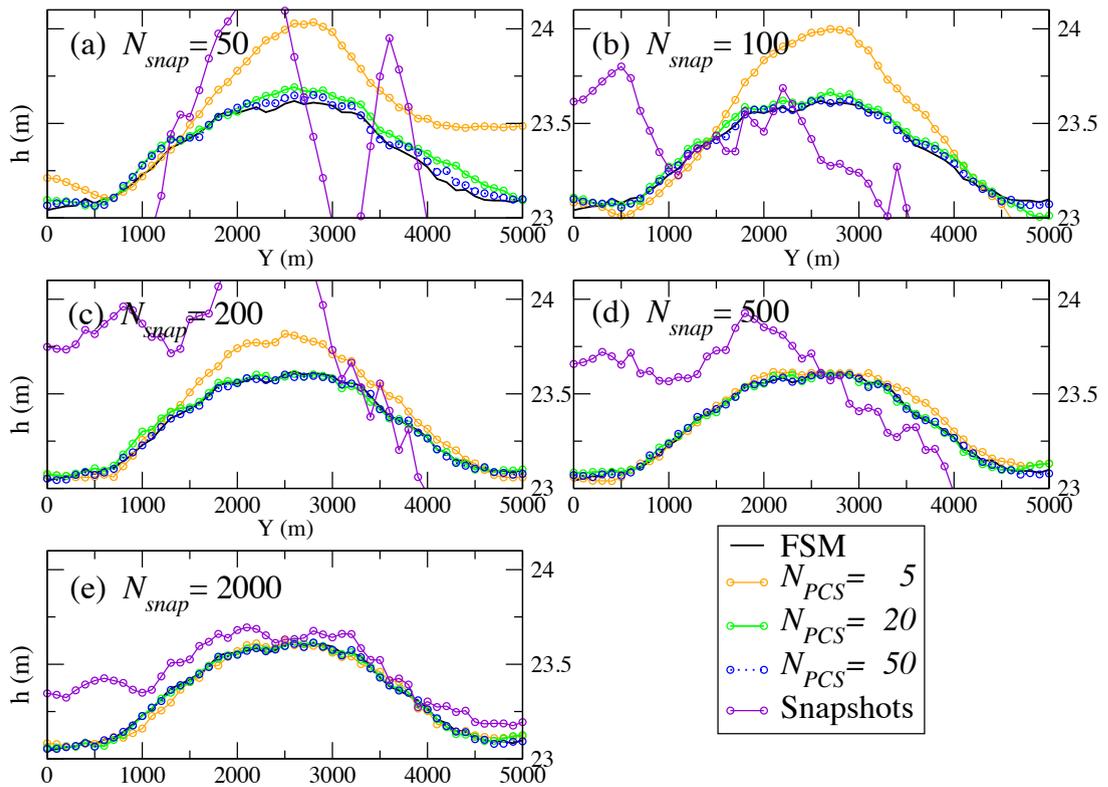
Figure 3.9: First 10-quantile of the pressure head distribution computed along the vertical central transept of the reference domain depicted in Fig. 3.1 for test case TC1. Results for N_{ens} (black) and N_{snap} (violet) standard MC simulations of the FSM and for the ROMC with $N_{PC} = 5$, 20 or 50 and $N_{snap} = 50$ (a), 100 (b), 200 (c), 500 (d) or 2000 (e).
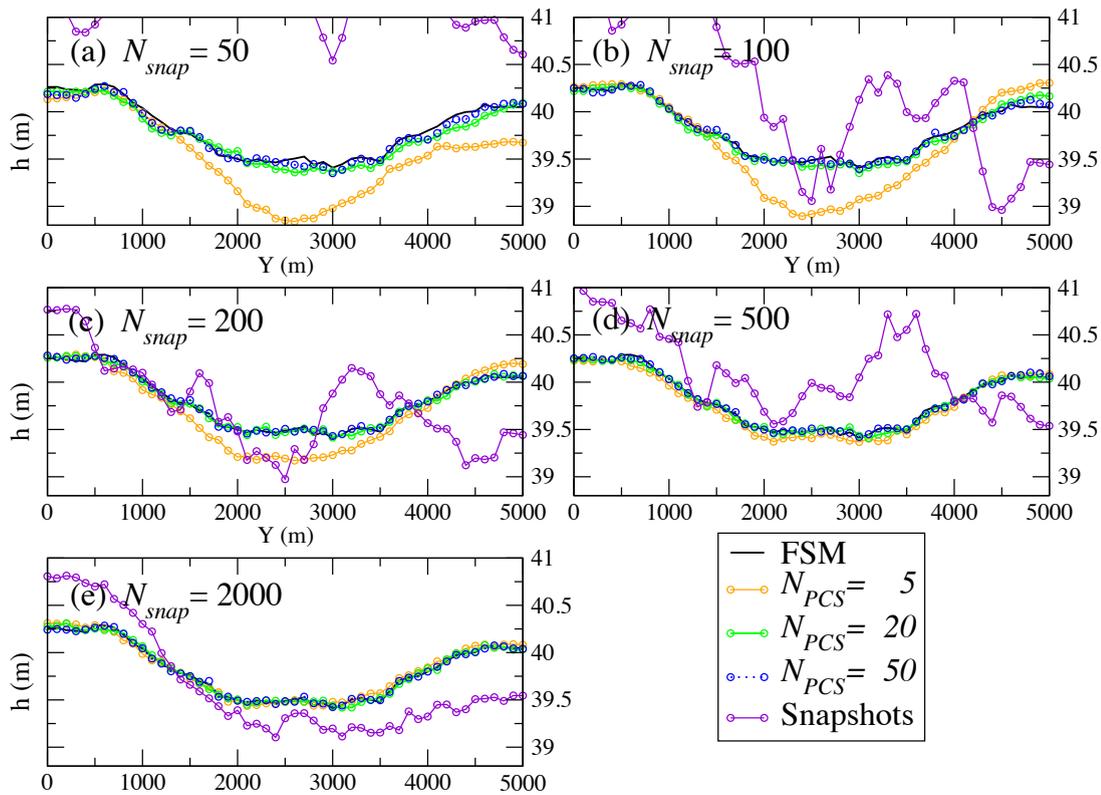
Figure 3.10: Ninth 10-quantile of the pressure head distribution computed along the vertical central transept of the reference domain depicted in Fig. 3.1 for test case TC1. Results for N_{ens} (black) and N_{snap} (violet) standard MC simulations of the FSM and for the ROMC with $N_{PC} = 5$, 20 or 50 and $N_{snap} = 50$ (a), 100 (b), 200 (c), 500 (d) or 2000 (e).
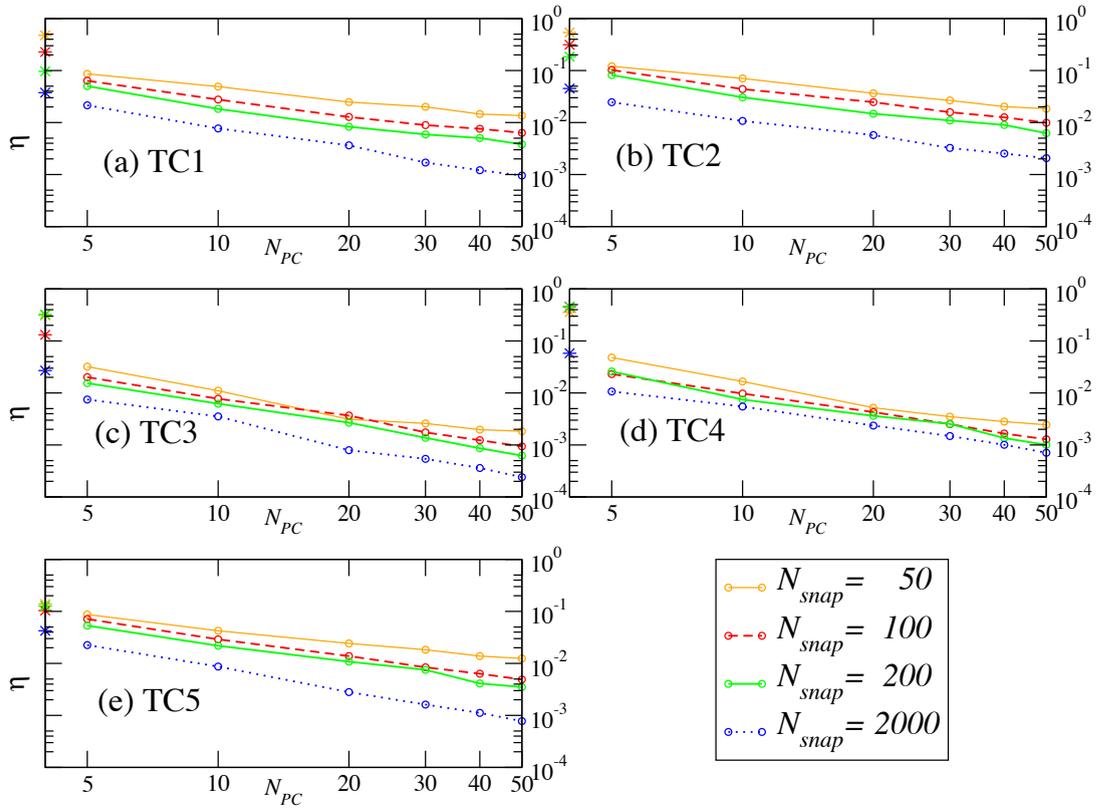
Figure 3.11: Relative RMSE, η, between the spatial variance of the hydraulic heads computed using the standard MC realizations of the FSM and the ROMC: (a) TC1, (b) TC2, (c) TC3, (d) TC4 and (e) TC5. Model reduction is achieved by $N_{PC} = 5$, 10, 20, 30, 40 or 50 obtained from $N_{snap} = 50$, 100, 200 or 2000. The symbols (*) on the vertical axes represent the relative RMSE between the spatial variance of the hydraulic heads computed using N_{ens} and N_{snap} realizations of the FSM.

square error (RMSE) between the spatial distribution of the head variance computed on the N_{ens} realizations obtained by FSM and ROMC:

$$\eta = \frac{\|Var[h] - Var[\tilde{h}]\|_{L_2(\Omega)}}{\|Var[h]\|_{L_2(\Omega)}}.$$

Fig. 3.11 compares the values of η for all test cases, and the same combination of N_{PC} and N_{snap} adopted in Fig. 3.4. Similarly to the previous results, η decreases displaying a power-low dependence on N_{PC}. While the relative error on the realizations is only slightly affected by N_{snap}, η decreases significantly when N_{snap} increases. This result suggests that the ROMC accuracy in retrieving the head variance increases if the reduced model is initialized on the basis of an augmented number of snapshots. This shows that the PCs of the reduced model inherit the statistical properties of the snapshot set, which must be sufficiently rich to guarantee enough accuracy to the ROMC. Increasing the number of snapshots improves the accuracy of ROMC also in the case of a large coefficient of variation for the recharge, while still keeping an affordable computational cost.

We close by noting that the symbols (*) reported in Fig. 3.11 represent the relative RMSE between the variance computed with N_{ens} realizations of the FSM and the variance computed with N_{snap} realizations of the FSM. The difference between these symbols and the corresponding curves highlights the relevance of the N_{res} realizations computed with the reduced model in order to improve the accuracy of the results.

3.7   Conclusions

We have presented and analyzed a methodology conducive to the application of a Galerkin model order reduction technique, Proper Orthogonal Decomposition (POD), to the solution of a synthetically-reconstructed, regionals-scale groundwater flow problem where recharge is modeled as spatially-distributed stochastic forcing term. POD allows reducing the computational cost associated with the solution of a given partial differential equation (PDE) upon identifying a finite (and small) number of conveniently selected basis functions. While POD is mostly adopted in the context of time-dependent deterministic problems and typically relies on solving the governing PDE at some observation times (termed snapshots), we follow the approach of Galbally et al. [2010] and select the snapshots in the probability space of the random state variable of interest (i.e. the hydraulic head). This allows casting a standard Monte Carlo (MC) solution of the selected problem into a Reduced Order Monte Carlo (ROMC) framework.

We start by approximating the random hydraulic head field by a truncated KL representation. The eigenfunctions of the KL representation are calculated on the basis of our reference set of snapshots. We remark again that our snapshots are selected in the probability space, with the idea that a limited number of these can provide a proper description of the first relevant principal components of the statistical behavior of the hydraulic head. The governing PDE is then projected onto the space of the principal components and the distributions of the KL random coefficients are approximated through a set of realizations obtained using the reduced system. Our choice of model reduction by means of principal components (as opposed to, e.g., the RB practice) is motivated by the fact that the resulting reduced model is consistent with the KL representation, in the sense that the optimality criteria associated with the KL representation hold. This implies that the mean error in the reproduction of the realizations using a number N_{PC} of principal components as basis functions is lower than what one could obtain by adopting $N_{snap} = N_{PC}$ snapshots in the RB framework. The drawback associated with our choice is that (in principle) the number of snapshots that are required to properly characterize the principal components might be large. We offer a detailed analysis of (i) the number of snapshots extracted from the probability space and adopted in the POD technique, (ii) the associated number of principal components, and (iii) the effects on the performance of the method of the key geostatistical parameters, i.e. sill and integral scale, describing the heterogeneity of the distributed recharge. We take the standard MC practice as a reference solution. The ROMC results are assessed in terms of (i) mean errors in the reproduction of hydraulic head realization, (ii) computational cost, and (iii) the ability of our reduced order model to render the sample estimators of the first two (ensemble) moments and the complete probability distribution of hydraulic heads, looking also at the tail of the distribution (first and ninth 10-quantiles).

Our work leads to the following major conclusions.

1. The eigenvalues of the operator (3.12) rapidly decay to zero, thus ensuring the effectiveness of our probabilistic model reduction strategy. The rate of decay of the eigenvalues (a) does not show a significant dependence on the coefficient of variation of the recharge, and (b) is influenced by the strength of the spatial correlation of the recharge, displaying the sharpest decay for the largest correlation scale investigated.

2. When compared against standard MC solutions of FSM (Eq. (3.2)), our results suggest that sampling the head probability space by a limited number (100 or 200) of snapshots leads to sufficiently well approximated distributions of the first Principal

Components (PCs). The reduced model based on a given number of PCs attains the highest accuracy for the largest correlation scale of the forcing term, f. The performance of the ROMC tends to deteriorate with increasing values of the coefficient of variation of f, for the range of the investigated parameters. For large values of the coefficient of variation of f it is convenient to increase the number of snapshots for the initialization of the ROMC, in such a way that the moments of the hydraulic head distribution are more accurately represented, while still keeping an efficient reduced model.

3. A key benefit of the method is a major reduction of the computational burden associated with the numerical evaluation of an ensemble of Monte Carlo realization of the flow field, when compared against standard MC modeling practice. Depending on the strength of (a) the spatial persistency and/or (b) the variance of the random forcing term, our ROMC scheme can improve the computational efficiency of a standard MC framework up to a factor of 100, while keeping remarkable degree of accuracy in providing the mean and covariance as well as the sample probability density of hydraulic heads.

Chapter 4

A reduced order model for the solution of transient groundwater flow with random hydraulic conductivity [1]

4.1 Summary

In this chapter we present a model order reduction technique that overcomes the computational burden required for the application of Monte Carlo methods to the solution of the transient groundwater flow equation with random hydraulic conductivity. The construction of the reduced model is based on the Galerkin projection of the high-dimensional model equations onto a subspace, approximated by a small number of optimally chosen basis functions (principal components). To obtain an efficient reduced order model, we develop an offline algorithm for the computation of the parameter-independent principal components. Our algorithm combines a greedy algorithm for the snapshot selection in the parameter space and the optimal distribution of the snapshots in time. Moreover, we introduce a residual-based estimation of the error associated with the reduced model. This estimation allows a considerable reduction of the number of full system model solutions required for the computation of the principal components. We demonstrate the robustness of our methodology by way of numerical examples, comparing the empirical statistics of the ensemble of the numerical solutions obtained using the traditional Monte

---

[1]The content of this chapter has been published in Pasetto et al. [2012b].

Carlo method and our reduced model. The numerical results show that our methodology significantly reduces the computational requirements (CPU time and storage) for the solution of the Monte Carlo simulation, ensuring a good approximation of the mean and variance of the head. The analysis of the empirical probability density functions at the observation wells suggests that our reduced model produces good results, especially at large times, and is most accurate in the regions with large drawdown, i.e., at locations where the change in head is highly influenced by the pumping.

4.2   Introduction

The physical description of groundwater flow in saturated porous media is derived by the solution of a parabolic partial differential equation. The application of this simple mathematical model to real hydrological systems depends on several parameters, such as the initial and boundary conditions, the forcing terms, and the aquifer properties, i.e., the hydraulic conductivity and the specific storage. In large aquifers, the heterogeneity of the soil usually is modeled by subdividing the aquifer into a number of zones, based on the geological formations of the soil. Then, in the classical deterministic approach, the calibration of the numerical model is usually achieved by solving an inverse problem [Oliver and Chen, 2011, Yeh, 1986], where the parameter values in each zone are estimated by minimizing the discrepancies between the model output and observations. Because of the presence of measurement errors as well as the complexity and heterogeneity of the real system, convergence to the correct parameter values is not always guaranteed. The uncertainty of model parameters is taken into account in the probabilistic approach [Dagan, 1982] where the stochastic partial differential equation for the groundwater flow governs the evolution in time of the entire probability density function of the head. This probabilistic approach is used in many applications, such as in a framework for forecasting and risk assessment, since it allows the quantification of the uncertainties in model predictions.

In this chapter we propose an efficient MC method for the solution of the transient groundwater flow equation with stochastic hydraulic conductivity. Following the ideas presented in Chapter 3, we reduce the computational requirements (CPU time and storage) associated with each single solution of the PDE and, consequently, improve the efficiency of the MC method. With this purpose, we apply the POD to the groundwater flow equation. The following main assumptions differ from the case studied in Chapter 3:

1. we consider a transient model, meaning that the reduced order model must take

into account the variation in time of the solution. Consequently we need snapshots not only for different parameter values, but also at different times.

2. The hydraulic conductivity is stochastic, which implies a multiplicative relation between the head and the random parameter.

3. The heterogeneity of the hydraulic conductivity is modeled by a subdivision of the domain in zones.

As described in Chapter 3, the computation of the principal components is the most important step in the construction of the reduced model. On one hand, to avoid large errors in the reduced model solution, we need a sufficient number of principal components to capture the dominating characteristics of the original full system model in the reduced space. On the other hand, the number of principal components determines the dimension of the reduced model and, consequently, the computational time required to solve the reduced model. For this reason, in this chapter we focalize our attention on the offline procedure. We look for an algorithm to establish the appropriate number of snapshots and principal components to use in the reduction.

For what concern the variability of the solution in time, Siade et al. [2010] introduced a methodology for the computation of optimal times for the selection of the snapshots in such a way that the resulting principal components account for the maximal variance of the solution. This methodology is very efficient, as it requires only two runs of the full system model to select the optimal snapshots.

In our application, the randomness of the hydraulic conductivity represent an additional source of variability in the space of the solutions. In theory, to obtain accurate reduced models, we should compute different principal components for each sample of the hydraulic conductivity [Vermeulen et al., 2004]. However, as the computation of the principal components requires the solution of the original full system model, the construction of a different reduced model for each MC realization would defeat the purpose of model reduction. A different approach is based on the idea that, if two realizations of hydraulic conductivity are "close enough", a single set of principal components can be used for the construction of the two RMs. Moreover, we may argue that if the principal components are collected from appropriately chosen hydraulic conductivity values, then we can obtain a unique set of principal components that cover the entire parameter space.

Heuristic methods such as the greedy algorithm [Grepl and Patera, 2005] have been developed to select the snapshots in the parameter space. The goal is to choose a set of

parameter values and, in an offline setting, improve the set of principal components until the solutions of the reduced model satisfy a validation condition. For this purpose, Grepl and Patera [2005] proposed an *a posteriori* error estimation based on the norm of the residual in which the computation of the residual does not involve the solution of the original full system model. We note that the "error" is defined as the difference between the full system model solution and the reduced model solution, while the "residual" is the vector obtained by substituting the reduced model solution into the full system model. Several examples of the use of residual norms to estimate the error norm are presented in the literature [Grepl and Patera, 2005, Haasdonk and Ohlberger, 2011, Hasenauer et al., 2012, Rovas et al., 2006]. Note that, for linear equations, error and residual norms are related by a scaling constant, which is difficult to evaluate [Grepl and Patera, 2005]. Hence, the main drawback in the application of the greedy algorithm with the validation condition based on *a posteriori* error estimation is that the norm of the residual can be much larger than that of the error, causing an overestimation of the error. This results in an inefficient reduced model in which the number of principal components is unnecessarily larger than the one actually needed to obtain the desired accuracy [Hasenauer et al., 2012]. An approach to resolve this problem is proposed by Hinze and Kunkei [2012] who applied the greedy algorithm to a nonlinear model with a one dimensional parameter space and presented a practical way to estimate the reduced model error. They computed a scaling factor between the norm of the error and the norm of the residual for the realizations where the full system model solution is available, and then linearly interpolated these values in the parameter space to compute a scaled residual for all the reduced model solutions.

In this chapter, we present a new methodology for the computation of the principal components that combines the optimal snapshot selection in time of Siade et al. [2010] with the greedy algorithm of Grepl and Patera [2005] for the selection of the snapshots in the parameter space. The efficiency of the proposed approach is improved by the use of the scaled reduced model residual to estimate the error, extending the approach of Hinze and Kunkei [2012] to a parameter space of general dimension and nonlinear interpolation. Hence, the scaling factor needed to relate error and residual norms is calculated exactly for those snapshots for which the full system model solution is available, while for the remaining runs it is interpolated from the "exact" values. The resulting algorithm for the offline process guarantees an efficient computation of the principal components. The unique set of principal components thus obtained is sufficient to cover the variability of the head in both the parameter space and time domain. Additionally, our approach minimizes

the number of full system model runs needed for the computation of the snapshots. The application of the reduced model to the MC method is then straightforward. The reduced model not only dramatically reduces the computational requirements associated to each system solution, but also permits a faster evaluation of the ensemble statistics.

We use two numerical examples to demonstrate the validity and applicability of our proposed methodology: a one-dimensional synthetic test case and a two-dimensional model of the Oristano groundwater aquifer, in Italy. The accuracy of the reduced model results is assessed by way of a comparison with the standard MC solution in terms of ensemble mean and variance on the domain and probability density function of the head at a number of observation wells.

4.3 Problem setting

We consider the governing equation describing a three-dimensional groundwater flow for a confined, isotropic aquifer with pumping [Bear, 1979]:

$$S_s \frac{\partial h(\vec{x}, t)}{\partial t} - \nabla \cdot (\mathbf{K} \nabla h(\vec{x}, t)) = q(\vec{x}, t), \qquad t \in [0, T_F], \qquad \vec{x} \in \Omega \subset \mathbb{R}^3 \qquad (4.1)$$

where Ω is the spatial representation of the aquifer, T_F is the final time [T], $\vec{x}$ is the vector of coordinates (x, y, z), h is the head [L], S_s is the specific storage [L^{-1}], q is the specific volumetric pumping rate [T^{-1}] and $\mathbf{K}$ is the hydraulic conductivity tensor:

$$\mathbf{K} = \begin{bmatrix} K_x & 0 & 0 \\ 0 & K_y & 0 \\ 0 & 0 & K_z \end{bmatrix},$$

with $K_x = K_y = K$ [LT$^-$1]. The initial and boundary conditions are:

$$\begin{cases} h(\vec{x}, 0) = h_0(\vec{x}), & \vec{x} \in \Omega, \\ h(\vec{x}, t) = h_D(\vec{x}, t), & \vec{x} \in \Gamma_D \subset \partial\Omega, \\ -\mathbf{K} \nabla h(\vec{x}, t) \cdot \vec{n}(\vec{x}) = q_N(\vec{x}, t), & \vec{x} \in \Gamma_N \subset \partial\Omega, \end{cases}$$

where Γ_D and Γ_N are the Dirichlet and Neumann boundaries, respectively; h_0, h_D, and q_N are known functions; and $\vec{n}$ is the normal vector at the boundary. In the following, we consider a heterogeneous hydraulic conductivity, $K(\vec{x})$, while the other soil parameters

are homogeneous. We model the heterogeneity of by subdividing the domain Ω into N_z zones, $\Omega_1, \ldots, \Omega_{N_z}$, such that:

$$\bigcup_{i=1}^{N_z} \Omega_i = \Omega, \qquad \text{and} \qquad \Omega_i \cap \Omega_j = \emptyset \quad \text{if} \quad i \neq j,$$

and the hydraulic conductivity $K(\vec{x})$ is constant in each zone, with values $K_1, \ldots, K_{N_z}$ in $\Omega_1, \ldots, \Omega_{N_z}$, respectively. We indicate with $\mathbf{K}$ the vector of the N_z zone values of the conductivity, $\mathbf{k} = \{K_1, \ldots, K_{N_z}\}$.

Without loss of generality, we solve Eq. (4.1) for the drawdown s, defined as the difference between the natural system dynamics (H) and the head (h) resulting from application of the pumping, i.e., $s = H - h$. The initial and boundary conditions for the drawdown are $s_0 = 0$, $s_D = 0$ and $q_N = 0$. We assume that the pumping rate is constant in time, a common practice for pumping tests. The proposed methodology also applies to a time-varying pumping rate. However, if more than one pumping well is present and the pumping rate varies with time, it is necessary to consider the response of each pumping well separately and then apply the principle of superposition.

The solution of Eq. (4.1) is achieved numerically, e.g., via finite elements, finite differences, or finite volumes. A refined spatial discretization of the domain Ω characterized by n nodes results in a high-dimensional linear system of ordinary differential equations (ODEs), written as:

$$\mathbf{B}\frac{d\mathbf{s}(t, \mathbf{k})}{dt} + \mathbf{A}(\mathbf{k})\mathbf{s}(t, \mathbf{k}) = \mathbf{q}, \qquad t \in [0, T_F], \tag{4.2}$$

where $\mathbf{A}$ (stiffness matrix) and $\mathbf{B}$ (mass matrix) are positive definite, symmetric and sparse matrices of dimension $n \times n$, and $\mathbf{s}$ and $\mathbf{q}$ are the vectors of nodal drawdown and source/sinks, respectively. In particular, focusing on the linear finite element method with obvious adjustments for other discretizations, the stiffness matrix can be written as a linear combination of parameter-independent matrices $\mathbf{A}_i$:

$$\mathbf{A}(\mathbf{k}) = \mathbf{A}(K_1, \ldots, K_{N_z}) = \sum_{i=1}^{N_z} K_i \mathbf{A}_i. \tag{4.3}$$

The components of each matrix $\mathbf{A}_i$ are evaluated using unit conductivity over the portion

of the domain encompassing the i-th zone, i.e.:

$$(\mathbf{A}_i)_{rs} = \int_{\Omega_i} \nabla \phi_r \cdot \nabla \phi_s \, d\Omega,$$

where ϕ_r and ϕ_s are the piecewise-linear basis functions used in our finite element approach. An equation similar to Eq. (4.3) also can be derived for the case of finite difference or finite volume schemes, where harmonic means of the conductivity values are used.

The solution in time of Eq. (4.2) is achieved with a backward difference approximation, with variable time step length $\Delta t_l = t_l - t_{l-1}$ leading to the following linear system of algebraic equations:

$$\left(\frac{1}{\Delta t_l} \mathbf{B} + \mathbf{A}(\mathbf{k}) \right) \mathbf{s}(t_l, \mathbf{k}) = \frac{1}{\Delta t_l} \mathbf{B} \, \mathbf{s}(t_{l-1}, \mathbf{k}) + \mathbf{q}, \qquad l = 1, \ldots, l_F. \tag{4.4}$$

We term Eq. (4.4) the full system model (FSM).

Since the hydraulic conductivity values K_i are unknown for real applications, they are modeled as random variables with a given probability distribution [Dagan, 1982]. In this work we consider that K_i's are uniform distribute randon variables:

$$K_i \sim U(K_i^{min}, K_i^{max}),$$

where K_i^{min} and K_i^{max} are the lower and upper bounds of K_i, respectively. Monte Carlo (MC) methods can be used to approximate the temporal evolution of the probability density function (PDF) of the head h, by solving Eq. (4.4) for several independent samples of the hydraulic conductivity $\mathbf{k}$. Let $\mathcal{L} = \left\{ \mathbf{k}^1, \ldots, \mathbf{k}^{N_{ens}} \right\}$ be the realizations of hydraulic conductivity, where N_{ens} is the number of samples. The PDF of the drawdown at time t_l is then approximated by the empirical distribution of the ensemble of the solutions $\mathbf{s}(t_l, \mathbf{k}^1), \ldots, \mathbf{s}(t_l, \mathbf{k}^{N_{ens}})$. MC methods are more accurate when a large number N_{ens} of FSM solutions are used; this procedure is computationally expensive and impractical for high-dimensional models (large n).

4.4 Reduced order methods

Model reduction methods can decrease the computational cost associated with the solution of Eq. (4.4). The idea of the Galerkin model reduction techniques [Kunisch and Volkwein, 2001] is to compute an approximated drawdown, $\tilde{\mathbf{s}}(t_l, \mathbf{k}^j)$, using a suitable lin-

ear combination of a small number of basis functions (also called principal components), $\mathbf{p}_1, \ldots, \mathbf{p}_{N_{PC}}$, where $\mathbf{p}_i$ is a n-dimensional vector and N_{PC} is the number of principal components used in the reduction. Indicating with $\mathbf{P}$ the $n \times N_{PC}$ matrix whose columns are $\mathbf{p}_1, \ldots, \mathbf{p}_{N_{PC}}$, we arrive at

$$\mathbf{s}(t_l, \mathbf{k}^j) \approx \tilde{\mathbf{s}}(t_l, \mathbf{k}^j) = \mathbf{P}\,\mathbf{a}(t_l, \mathbf{k}^j). \tag{4.5}$$

The N_{PC}-dimensional vector $\mathbf{a}$ is the solution of the reduced order equation obtained substituting $\mathbf{s}$ with $\tilde{\mathbf{s}}$ in Eq. (4.4) and applying the Galerkin projection with respect to $\mathbf{P}$:

$$\left(\frac{1}{\Delta t_l}\tilde{\mathbf{B}} + \tilde{\mathbf{A}}(\mathbf{k}^j) \right) \mathbf{a}(t_l, \mathbf{k}^j) = \frac{1}{\Delta t_l}\tilde{\mathbf{B}}\,\mathbf{a}(t_{l-1}, \mathbf{k}^j) + \tilde{\mathbf{q}}, \qquad l = 1, \ldots, l_F, \tag{4.6}$$

where $\tilde{\mathbf{A}}(\mathbf{k}^j) = \mathbf{P}^T\mathbf{A}(\mathbf{k}^j)\mathbf{P}$, $\tilde{\mathbf{B}} = \mathbf{P}^T\mathbf{B}\mathbf{P}$ and $\tilde{\mathbf{q}} = \mathbf{P}^T\mathbf{q}$. We refer to Eq. (4.6) as the reduced model (RM). Note that the matrix $\tilde{\mathbf{B}}$ is parameter independent, so it can be computed once and then stored. Due to the linear dependence of $\mathbf{A}$ on the hydraulic conductivity averages (Eq. (4.3)), the computation of $\tilde{\mathbf{A}}$ also does not depend on the original high-dimension n:

$$\tilde{\mathbf{A}}(\mathbf{k}) = \sum_{i=1}^{N_z} K_i \mathbf{P}^T \mathbf{A}_i \mathbf{P} = \sum_{i=1}^{N_z} K_i \tilde{\mathbf{A}}_i, \tag{4.7}$$

where the matrices $\tilde{\mathbf{A}}_i$ are parameter-independent matrices in the reduced dimension. In this way, the assembly and the solution of the RM (Eq. (4.6)) are performed only in the reduced dimension N_{PC}.

A RM is accurate if the error $\mathbf{e}(t_l, \mathbf{k}^j)$ between the FSM and the RM solutions,

$$\mathbf{e}(t_l, \mathbf{k}^j) = \mathbf{s}(t_l, \mathbf{k}^j) - \tilde{\mathbf{s}}(t_l, \mathbf{k}^j) \tag{4.8}$$

is small in norm. The error associated with the RM is estimated *a posteriori*, using the computation of the residual $\mathbf{r}(t_l, \mathbf{k}^j)$ which is the vector obtained replacing $\mathbf{s}$ with $\tilde{\mathbf{s}}$ in Eq. (4.4) [Grepl and Patera, 2005, Hasenauer et al., 2012, Rovas et al., 2006]:

$$\mathbf{r}(t_l, \mathbf{k}^j) = - \left(\frac{1}{\Delta t_l}\mathbf{B} + \mathbf{A}(\mathbf{k}^j) \right) \mathbf{P}\,\mathbf{a}(t_l, \mathbf{k}^j) + \frac{1}{\Delta t_l}\mathbf{B}\mathbf{P}\,\mathbf{a}(t_{l-1}, \mathbf{k}^j) + \mathbf{q}. \tag{4.9}$$

The error is related to the residual by the following equation:

$$\mathbf{r}(t_l, \mathbf{k}^j) = \left(\frac{1}{\Delta t_l} \mathbf{B} + \mathbf{A}(\mathbf{k}^j) \right) \mathbf{e}(t_l, \mathbf{k}^j) - \frac{1}{\Delta t_l} \mathbf{B} \, \mathbf{e}(t_{l-1}, \mathbf{k}^j). \tag{4.10}$$

Eqs. (4.9) and (4.10) reveal two important properties of the residual: (i) the residual is computed without knowing the corresponding FSM solution and (ii) the residual is zero when the error is zero. Moreover, Haasdonk and Ohlberger [2011] demonstrated the following *a posteriori* error estimate:

$$\|\mathbf{e}(t_l, \mathbf{k}^j)\|_2 \leq C(\mathbf{k}^j) \left(\|\mathbf{r}(t_0, \mathbf{k}^j)\|_2 + \int_0^t \|\mathbf{r}(\tau, \mathbf{k}^j)\|_2 \, d\tau \right) = R(t_l, \mathbf{k}^j), \tag{4.11}$$

where

$$\|\mathbf{e}\|_2 = \frac{\sqrt{\sum_{i=1}^n e_i^2}}{n}$$

and $C(\mathbf{k}^j)$ is a constant that can be approximated by the value 1 for our particular model Haasdonk and Ohlberger [2011]. Also the computation of the time-integrated norm of the residual $R(t_l, \mathbf{k}^i)$ can be performed in the reduced dimension N_{PC}. In fact:

$$
\begin{aligned}
\|\mathbf{r}(t_l, \mathbf{k}^j)\|_2 \;=\; & (\mathbf{a}_l^j)^T \mathbf{P}^T \left(\tfrac{1}{(\Delta t_l)^2} \mathbf{B}^T \mathbf{B} + \tfrac{2}{(\Delta t_l)} \mathbf{B}^T \mathbf{A}^j + (\mathbf{A}^j)^T \mathbf{A}^j \right) \mathbf{P} \, \mathbf{a}_l^j \\
& -2(\mathbf{a}_l^j)^T \mathbf{P}^T \left(+ \tfrac{1}{(\Delta t_l)^2} \mathbf{B}^T \mathbf{B} + \tfrac{1}{\Delta t_l} (\mathbf{A}^j)^T \mathbf{B} \right) \mathbf{P} \, \mathbf{a}_{l-1}^j \\
& -2(\mathbf{a}_l^j)^T \mathbf{P}^T \left(\tfrac{1}{(\Delta t_l)^2} \mathbf{B}^T + (\mathbf{A}^j)^T \right) \mathbf{q} \\
& +(\mathbf{a}_{l-1}^j)^T \mathbf{P}^T \left(\tfrac{1}{(\Delta t_l)^2} \mathbf{B}^T \mathbf{B} \right) \mathbf{P} \, \mathbf{a}_{l-1}^j \\
& +2(\mathbf{a}_{l-1}^j)^T \mathbf{P}^T \left(\tfrac{1}{(\Delta t_l)} \mathbf{B}^T \right) \mathbf{q}+ \\
& +\mathbf{q}^T \mathbf{q},
\end{aligned}
$$

where the matrices $\mathbf{P}^T \mathbf{B}^T \mathbf{B} \mathbf{P}$, $\mathbf{P}^T (\mathbf{A}^j)^T \mathbf{B} \mathbf{P}$, and $\mathbf{P}^T (\mathbf{A}^j)^T \mathbf{A}^j \mathbf{P}$, the vectors $\mathbf{P}^T \mathbf{B}^T \mathbf{q}$, and $\mathbf{P}^T (\mathbf{A}^j)^T \mathbf{q}$, and the scalar $\mathbf{q}^T \mathbf{q}$ can be computed offline (a procedure similar to Eq. (4.7) applies for the parameter dependent quantities). Note that, to simplify the notation, in the previous equation the time and parameter dependences are shown with the indices l and j, respectively.

In the following we present an efficient algorithm to compute the principal components $\mathbf{p}_i$. We search for principal components that are time- and parameter-independent, and that can be applied to the construction of the RM associated with each realization of hydraulic conductivity in $\mathcal{L}$. In this way, the expensive computation of the principal components is performed only once, in an offline setting, i.e., before the application of the

RM to the MC simulation. We select a set of hydraulic conductivity values $\mathcal{K}$,

$$\mathcal{K} = \left\{ \hat{\mathbf{k}}^1, \ldots, \hat{\mathbf{k}}^{N_k} \right\},$$

to validate the RM accuracy (the validation set). The set $\mathcal{K}$ can be different from the set $\mathcal{L}$; thus, we use the symbol '∧' to distinguish between the realizations in $\mathcal{K}$, used for the offline process, and the random realizations in $\mathcal{L}$, used in the MC simulation. We start from a reduced model of dimension one and compute the RM solution for all the realizations in the validation set $\mathcal{K}$. Then, we increase the number of principal components until the error associated with the RM solution is less than a given tolerance for all the realizations in $\mathcal{K}$.

4.4.1   Optimal snapshot selection in time

Algorithm 1 RM initialization

   Compute T_s for $\mathbf{k}^1$
   Compute $\hat{t}_1, \ldots, \hat{t}_{N_{snap}}$
   Solve Eq. (4.4) and store $\mathbf{s}_{\hat{t}_1}, \ldots, \mathbf{s}_{\hat{t}_{N_{snap}}}$
   Compute $\mathbf{p}_1^1, \ldots, \mathbf{p}_{N_{snap}}^1$
   $N_{PC} \leftarrow 1$
   $\mathbf{P} \leftarrow \{\mathbf{p}_1^1\}$
   Solve Eq. (4.6) for $\mathbf{k}^1$
   Compute $\|\mathbf{e}(T_F, \mathbf{k}^1)\|_2$
   while $\|\mathbf{e}(T_F, \mathbf{k}^1)\|_2 \geq \tau_e$ **do**
      $N_{PC} \leftarrow N_{PC} + 1$
      $\mathbf{P} \leftarrow \{\mathbf{p}_1^1, \ldots, \mathbf{p}_{N_{PC}}^1\}$
      Solve Eq. (4.6) for $\mathbf{k}^1$
      Compute $\|\mathbf{e}(T_F, \mathbf{k}^1)\|_2$
   end while

The reduced model is initialized considering the first realization in $\mathcal{K}$, $\hat{\mathbf{k}}^1$. Our choice is to select $\hat{\mathbf{k}}^1$ as the realization with the mean value of the conductivity in each zone. In this way the reduced model reproduces the average response of the system. Then, the principal components are computed based on the transient FSM solution for $\hat{\mathbf{k}}^1$, to capture the dominating characteristics of the solution in time. Given a realization of the hydraulic conductivity $\mathbf{k}$ and given a fixed number N_{PC} of basis functions for the reduced model, it is possible to compute an optimal set of time-independent principal components that minimizes the errors $\mathbf{e}(t_l, \mathbf{k})$, for $l = 1, \ldots, l_F$. The snapshot technique, presented in

Chapter 3 for parameter-dependent functions, here is applied to compute the principal components for time-dependent functions. This requires the storage of the FSM solution at specific times, $\hat{t}_1, \ldots, \hat{t}_{N_{snap}}$. The solution vectors obtained at these times are called snapshots and are indicated by $\mathbf{s}_{\hat{t}_1}, \ldots, \mathbf{s}_{\hat{t}_{N_{snap}}}$. We then obtain the principal components via a principal component analysis (PCA) on the set of the snapshots.

Siade et al. [2010] suggested a general procedure for selecting optimal snapshot times for the groundwater equation (Eq. (4.1)). Let T_s be the time at which the solution $\mathbf{s}(\cdot, \mathbf{k})$ reaches approximately a steady-state. This can be computed by solving the FSM Eq. (4.4) with a termination condition:

$$\frac{\|\mathbf{s}(t_l, \mathbf{k}) - \mathbf{s}(t_{l-1}, \mathbf{k})\|_2}{\|\mathbf{s}(t_l, \mathbf{k})\|_2} \leq \tau_s,$$

with the tolerance τ_s sufficiently small (e.g. $\tau_s = 10^{-3}$) or, when appropriate, with other simple approximations based, e.g., on the Theis equation. The optimal snapshot times $\hat{t}_i$ are then located along an exponential function:

$$t(u) = \frac{T_s}{0.9} \left(\beta \, e^{\alpha u} + \gamma \right), \quad u \geq 0, \tag{4.12}$$

where α, β and γ are computed using the results presented in [Siade et al., 2010]. Finally, as we are interested in the solution of Eq. (4.1) in the time interval $[0, T_f]$, we chose $\hat{t}_1 = \Delta t_1$ and $\hat{t}_{N_{snap}} = T_F$. Let u_1 and $u_{N_{snap}}$ be such that $t(u_1) = \hat{t}_1$ and $t(u_{N_{snap}}) = \hat{t}_{N_{snap}}$ using Eq. (4.12), respectively. The other snapshot times are computed using Eq. (4.12) on equally spaced values of u, with step $\Delta u = (u_{N_{snap}} - u_1)/(N_{snap} - 1)$:

$$\left(u_1, \ldots, \ u_i = u_{i-1} + \Delta u, \ldots, \ u_{N_{snap}} \right).$$

The appealing feature of this method is the possibility of computing N_{snap} quasi-optimal snapshots with only two FSM runs. A PCA is applied on the snapshots and the corresponding principal components $\mathbf{p}_1^j, \ldots, \mathbf{p}_{N_{snap}}^j$ are stored (the index j means that the principal components are computed from the realization $\mathbf{k}^j$). Moreover, as we solve the FSM to get the snapshots, we can store the FSM solution at the output times and then compute the error associated with the RM solution at these times. If the error is over a tolerance value τ_e, then the RM can be improved adding a new principal component. Note that the principal components are ordered in such a way that, if the desired dimension of the RM is N_{PC}, then the RM constructed with the first N_{PC} principal components,

$\mathbf{p}_1^j, \ldots, \mathbf{p}_{N_{PC}}^j$, is the one that minimizes the error. For this reason, the best way to improve the RM is to add the first unused principal component of the set $\mathbf{p}_1^j, \ldots, \mathbf{p}_{N_{snap}}^j$. To summarize, we use Algorithm 1, to initialize the RM.

4.4.2   Snapshot selection in the parameter space

The greedy algorithm is a heuristic method used to determine which parameter values in $\mathcal{K}$ are to be selected for generating the snapshots without applying Algorithm 1, i.e. the FSM, to each realization $\mathbf{k}^i$. The main idea is to compute the new snapshots for the realization in $\mathcal{K}$ where the RM solution gives the worst approximation of the FSM solution. In this way we improve the RM by including the basis functions that were not previously considered. Since the FSM solution is not available for all the realizations in $\mathcal{K}$, we use the estimation of the error based on the residual (Eq. (4.11)) to determine if the RM solution is accurate.

The matrix of the principal components $\mathbf{P}$ is initialized as described in the previous section. Then, we compute the RM solution and the norm of the corresponding residual for each realization in $\mathcal{K}$. If the maximum norm of the residual is above a specified tolerance value τ_e, we compute the FSM corresponding to the realization with the maximum norm of the residual, in agreement with the greedy algorithm. The matrix of the principal components is updated using Algorithm 1. We repeat these operations until the maximum norm of the residual is smaller than a specified tolerance value for all the realizations in $\mathcal{K}$.

We let $\mathcal{K}^* = \left\{ \hat{\mathbf{k}}^{*1}, \ldots, \hat{\mathbf{k}}^{*N_{gr}} \right\}$, $\mathcal{K}^* \subset \mathcal{K}$, be the set of the N_{gr} parameter values selected from the first N_{gr} iterations of the greedy algorithm: in other words, at the N_{gr}-th iteration of the greedy algorithm $\hat{\mathbf{k}}^{*N_{gr}}$ satisfies:

$$\hat{\mathbf{k}}^{*N_{gr}} = \arg\max_{\hat{\mathbf{k}}^i \in \mathcal{K}} R(T_F, \hat{\mathbf{k}}^i). \tag{4.13}$$

Note that we use the symbol '*' to indicate the realizations in $\mathcal{K}$ for which we compute the snapshots. For each realization $\hat{\mathbf{k}}^{*i}$ we select snapshots in time and collect the corresponding principal components $\mathbf{p}^{*i}$ as described in Algorithm 1. To ensure the orthogonality of matrix $\mathbf{P}$, the new principal components are orthonormalized with respect to the columns of $\mathbf{P}$ and possible redundant principal components are discarded. Moreover, at each modification of the principal components, it is convenient to check if the error still satisfies the validation condition for all the realizations in $\mathcal{K}^*$. If the condition is not satisfied, more

principal components are selected from the realization with the maximum error, without any additional FSM run.

Since the norm of the residual $R(T_F, \hat{\mathbf{k}}^i)$ can be much larger than the norm of the error, using Eq. (4.11) in the greedy algorithm may select more principal components than necessary in order to satisfy the desired accuracy [Hasenauer et al., 2012]. Following the suggestion of Hinze and Kunkei [2012], we scale the norm of the residual $R(T_F, \hat{\mathbf{k}}^i)$ to better estimate the error. We note that for all the realizations in $\mathcal{K}^*$ we know both the error and the residual. Then for these realizations we compute an exact scaling factor ρ^{*j} such that

$$\rho^{*j} = \frac{\|\mathbf{e}(T_F, \hat{\mathbf{k}}^{*j})\|_2}{R(T_F, \hat{\mathbf{k}}^{*i})}, \qquad \hat{\mathbf{k}}^{*j} \in \mathcal{K}^*. \tag{4.14}$$

For the realizations in $\mathcal{K}$ which are not in $\mathcal{K}^*$ we approximate the scaling factor ρ^i associated to $\hat{\mathbf{k}}^i$ interpolating the values ρ^{*j} with respect to the hydraulic conductivity. For this purpose we need to introduce a distance in the parameter space that relates the change in drawdown to the change in hydraulic conductivity. We consider the following distance $d^{i,j}$ between $\hat{\mathbf{k}}^i$ and realizations $\hat{\mathbf{k}}^{*j}$ in $\mathcal{K}^*$:

$$d^{i,j} = \sqrt{\sum_{m=1}^{N_z} \left(\frac{1}{\hat{K}_m^i} - \frac{1}{\hat{K}_m^{*j}} \right)^2}, \quad j = 1, \ldots, N_{gr}, \tag{4.15}$$

where the inverse values of K_m^i are used to take into account the inverse relation between the drawdown and the hydraulic conductivity. Let $\hat{\mathbf{k}}^{*i_1}$ and $\hat{\mathbf{k}}^{*i_2}$ be the two conductivities values in $\mathcal{K}^*$ closest to $\hat{\mathbf{k}}^i$ with associated distances d^{i,i_1}, d^{i,i_2} and scaling factors ρ^{*i_1} and ρ^{*i_2}. We propose the following scaling function:

$$\rho^i = \begin{cases} \left| 1 - (1 - \rho^{*i_1}) \exp(-\frac{d^{i,i_1}}{\lambda}) \right| & \text{if } N_{gr} = 1. \\ \\ \left| (1 - (1 - \rho^{*i_1}) \exp(-\frac{d^{i,i_1}}{\lambda}) - (1 - \rho^{*i_2}) \exp(-\frac{d^{i,i_2}}{\lambda}) - \right. \\ \left. (1 - 0.5(\rho^{*i_1} + \rho^{*i_2})) \exp(-\frac{(d^{i,i_1} + d^{i,i_2})}{\lambda}) \right| & \text{if } N_{gr} \geq 2. \end{cases} \tag{4.16}$$

This scaling is defined such that, if the distances d^{i,i_1} and d^{i,i_2} are large, ρ^i is about 1 (no scaling): otherwise ρ^i varies continuously between ρ^{*i_1}, ρ^{*i_2} and $0.5(\rho^{*i_1} + \rho^{*i_2})$. The factor λ controls the shape of the scaling function. When $N_{gr} = 1$, large values of λ result in scaling factors ρ^i closer to ρ^{*i_1}, while small values result in scaling factors closer to 1.

Let $\hat{R}(T_F, \hat{\mathbf{k}}^i)$ be the scaled residual,

$$\hat{R}(T_F, \hat{\mathbf{k}}^i) = \rho^i \, R(T_F, \hat{\mathbf{k}}^i). \tag{4.17}$$

We apply the greedy algorithm with the following modification of Eq. (4.13):

$$\mathbf{k}^{*N_{gr}} = \arg\max_{\hat{\mathbf{k}}^i \in \mathcal{K}} \hat{R}(T_F, \hat{\mathbf{k}}^i). \tag{4.18}$$

The resulting procedure is summarized in Algorithm 2.

Algorithm 2 Offline: modified greedy algorithm

$N_{gr} \leftarrow 0$
$N_{PC} \leftarrow 0$
$\hat{\mathbf{k}}^{*1} \leftarrow \hat{\mathbf{k}}^1$
$\hat{R}(T_F, \hat{\mathbf{k}}^{*1}) \leftarrow 2\tau_e$
while $\hat{R}(T_F, \hat{\mathbf{k}}^{*N_{gr}+1}) \geq \tau_e$ **do**
    $N_{gr} \leftarrow N_{gr} + 1$
    Compute $\mathbf{p}^{*1}, \dots \mathbf{p}^{*N_{snap}}$ from $\hat{\mathbf{k}}^{*N_{gr}}$ (Algorithm 1)
    $i \leftarrow N_{gr}$
    while $\hat{R}(T_F, \hat{\mathbf{k}}^{*i}) \geq \tau_e$ **do**
        $N_{PC} \leftarrow N_{PC} + 1$
        Improve $\mathbf{P}$ with a principal component from $\hat{\mathbf{k}}^{*i}$
        for $j = 1 \to N_{gr}$ (Loop on $\mathcal{K}^*$) **do**
            Solve the RM Eq. (4.6) for $\hat{\mathbf{k}}^{*j}$
            Compute $\|\mathbf{e}(T_F, \hat{\mathbf{k}}^{*j})\|$, $\hat{R}(T_F, \hat{\mathbf{k}}^{*j})$ and ρ^{*j} (Eq. (4.17))
        end for
        $i = \arg\max_{j=1,\dots,N_{gr}} \hat{R}(T_F, \hat{\mathbf{k}}^{*j})$
    end while
    for $j = 1 \to N_k$ (Loop on $\mathcal{K}$) **do**
        Solve the RM Eq. (4.6) for $\hat{\mathbf{k}}^j$
        Compute ρ^j and $\hat{R}(T_F, \hat{\mathbf{k}}^j)$
    end for
    $\hat{\mathbf{k}}^{*N_{gr}+1} = \arg\max_{\hat{\mathbf{k}}^j \in \mathcal{K}} \hat{R}(T_F, \hat{\mathbf{k}}^j)$
end while

A crucial point in the application of Algorithm 2 is the choice of an adequate set $\mathcal{K}$, i.e., the set of realizations of hydraulic conductivity for the validation of the RM over the entire parameter space. A possible choice is to set $\mathcal{K} = \mathcal{L}$, i.e., to validate the RM directly on a subset of the pre-selected set of MC realizations. We prefer to explore another approach, to avoid the dependency of the offline algorithm to the random realizations. We note that (i) the RM is accurate when the parameter values are close to each other,

and (ii) in the offline stage we are interested in the validation of the RM for the worst parameter combinations. Therefore, we propose to select $\mathcal{K}$ as the set of all the possible combinations of upper, lower and average values of the hydraulic conductivity in all the zones. If a different probability distribution describes the hydraulic conductivity, then the same choice is possible for $\mathcal{K}$, considering the tails (e.g., the first and the last 10 -quantiles) and the median of the distribution instead of the upper bound, lower bound and mean. With this approach the size of $\mathcal{K}$ increases rapidly with the number of zones, $N_k = 3^{N_z}$, compromising the computational cost. However, Algorithm 2 can still be efficient because only the RM runs for all the N_k realizations, while the FSM solution is computed only for a few realizations. If the number of combinations 3^{N_z} far exceeds the number of realizations N_{ens}, then we still can apply the offline procedure directly on the set of realizations $\mathcal{L}$ used in the MC method.

4.4.3 Online: Monte Carlo and RM

The online procedure consists of the application of the RM to the Monte Carlo method, i.e., the solution of Eq. (4.6) for all the realizations of hydraulic conductivity $\mathbf{k}^1, \ldots, \mathbf{k}^{N_{ens}}$ and the computation the desired statistics (i.e., spatial mean, spatial variance, probability density function at the output nodes) from the ensemble of solutions $\tilde{\mathbf{s}}(t_i^{out}, \mathbf{k}^1), \ldots,$ $\tilde{\mathbf{s}}(t_i^{out}, \mathbf{k}^{N_{ens}})$ at the output times $t_1^{out}, \ldots, t_F^{out}$. The low dimension of the RM allow us to efficiently evaluate the ensemble statistics. As the matrix $\tilde{\mathbf{B}}$ and the vector $\tilde{\mathbf{q}}$ are computed offline, the RM is assembled simply, using Eq. (4.7) to compute the matrix $\tilde{\mathbf{A}}$ (computational cost $O(N_z (N_{PC})^2)$). Then, Eq. (4.6) requires the solution of a linear system of dimension $N_{PC} \times N_{PC}$ for each time step. For calculating the statistics of the drawdown at the output time t_i^{out}, we let $\boldsymbol{\mu}_{\tilde{\mathbf{s}}}(t_i^{out})$ and $\mathbf{C}_{\tilde{\mathbf{s}}}(t_i^{out})$ be the n-dimensional vector of the mean and the $n \times n$ covariance matrix of the ensemble $\tilde{\mathbf{s}}(t_i^{out}, \mathbf{k}^j)$, respectively. These quantities are efficiently computed using Eq. 4.5 without the solution in the high-dimensional space:

$$\boldsymbol{\mu}_{\tilde{\mathbf{s}}}(t_i^{out}) = \frac{1}{N_{ens}} \sum_{j=1}^{N_{ens}} \tilde{\mathbf{s}}(t_i^{out}, \mathbf{k}^j) = \frac{1}{N_{ens}} \sum_{j=1}^{N_{ens}} \mathbf{P}\, \mathbf{a}(t_i^{out}, \mathbf{k}^j) =$$

$$= \mathbf{P} \frac{1}{N_{ens}} \sum_{j=1}^{N_{ens}} \mathbf{a}(t_i^{out}, \mathbf{k}^j) = \mathbf{P} \boldsymbol{\mu}_{\mathbf{a}}(t_i^{out}),$$

Pumping well, q= 10 m³/d

$s = 0\,m$

| k_1 | k_2 | k_3 | k_4 | k_5 |

$s = 0\,m$

0   20   40   60   80   100

x (m)
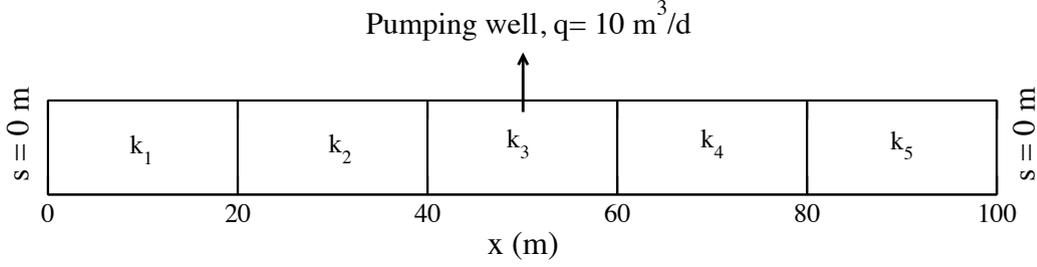
Figure 4.1: Sketch of the one-dimensional aquifer used for TC1.

where $\boldsymbol{\mu_a}$ is the N_{PC}-dimensional vector of the mean values of the coefficients $\mathbf{a}$;

$$
\mathbf{C_{\tilde{s}}}(t_i^{out}) = \frac{1}{N_{ens}-1} \sum_{j=1}^{N_{ens}} \left(\tilde{\mathbf{s}}\left(t_i^{out}, \mathbf{k}^j\right) - \boldsymbol{\mu_{\tilde{s}}}\left(t_i^{out}\right) \right) \left(\tilde{\mathbf{s}}\left(t_i^{out}, \mathbf{k}^j\right) - \boldsymbol{\mu_{\tilde{s}}}\left(t_i^{out}\right) \right)^T =
$$

$$
= \frac{1}{N_{ens}-1} \sum_{j=1}^{N_{ens}} \mathbf{P} \left(\mathbf{a}\left(t_i^{out}, \mathbf{k}^j\right) - \boldsymbol{\mu_a}\left(t_i^{out}\right) \right) \left(\mathbf{a}\left(t_i^{out}, \mathbf{k}^j\right) - \boldsymbol{\mu_a}\left(t_i^{out}\right) \right)^T \mathbf{P}^T =
$$

$$
= \mathbf{P} \mathbf{C_a}\left(t_i^{out}\right) \mathbf{P}^T,
$$

where $\mathbf{C_a}$ is the $N_{PC} \times N_{PC}$ covariance matrix of the coefficients $\mathbf{a}$.

4.5   Numerical results

4.5.1   One-dimensional model

To validate our proposed methodology, we first consider the synthetic test case used by both McPhee and Yeh [2008] and Siade et al. [2012]. Fig. 4.1 depicts a one-dimensional aquifer of length 100 m, with a pumping well located in the center and Dirichlet boundary conditions of 0 m. The thickness of the aquifer (b) is 1 m with a constant specific storage (S_s) of 1 m−1 so that in Eq. (4.1) the elastic storage represents the storage coefficient and the hydraulic conductivity tensor represents scalar transmissivity. The aquifer is subdivided into five zones and the hydraulic conductivity is modeled as a uniformly distributed random variable in each zone, with lower bound $k_i^{min} = 0.1$ m/d and upper bound $k_i^{max} = 20$ m/d. We consider a pumping test with duration of 100 days ($T_F = 100$ d) with a constant pumping rate of 10 m³/d. We name this test case TC1. The numerical simulation is performed with the program Sat2D [Gambolati et al., 1999], a finite element-based software for the simulation of saturated groundwater flow with a preconditioned conjugate gradient solver for the linear system arising from equation Eq. (4.4). The FSM

				One realization (s)	Total time (s)
FSM	$N_{ens}=$	10000	$n=$	303	0.12	1279
RM online	$N_{ens}=$	10000	$N_{PC}=$	30	0.0035	110
RM offline	$N_k=$	243	$N_{gr}=$	12		14
RM total						124

Table 4.1: Comparison between the FSM and RM CPU times for the Monte Carlo simulation for TC1

consists of $n=$ 303 nodes (three lines of 101 nodes each). To take into account the random parameters, we apply the MC method using $N_{ens}=$10,000 independent realizations of hydraulic conductivity. The ensemble-based statistics on the drawdown are computed and stored every five days, for a total of 21 output times. The numerical simulation of this simple scenario requires a CPU time of about 21 minutes, i.e., 0.13 s for each MC realization.

To improve the computational efficiency of the MC simulation we apply our RM, using Algorithm 2 to compute the parameter-independent principal components. The set $\mathcal{K}$ consists of $3^5 =$243 realizations of hydraulic conductivity. The tolerance on the average nodal error τ_e is set to i 10^{-3} m while the value of λ for the computation of the scaling factors ρ^i (Eq. (4.16)) is set to 1000 m^2. For each realization in $\mathcal{K}^*$ we store $N_{snap}=$ 15 snapshots at the optimal times computed by Eq. (4.12) with the following values for the parameters [Siade et al., 2010]:

$$t(0) = 1.11 \times 10^{-7} T_S, \qquad t(1) = T_s, \qquad \gamma = -3.87 \times 10^{-6}.$$

Table 4.1 shows the comparison of CPU times between the FSM and the RM for the MC simulation. It also shows the CPU time required for the offline and online calculations. The offline procedure for TC1 requires the computation of 24 FSM solutions (12 for the computation of the steady-state time T_S and 12 for the selection of the snapshots) and 3295 RM solutions, for a total CPU time of 14 s. The resulting RM has dimension $N_{PC}=$30 (compared with $n=$303), which corresponds to a RM about 40 times faster than the FSM (0.3x10-2 s per realization).

To demonstrate the accuracy and efficiency of Algorithm 2 with the modified greedy algorithm, in Fig. 4.2 we show the norm of the error $\|\mathbf{e}^i\|_2$, the norm of the residual R^i and the scaled norm of the residual $\hat{R}^i$ for all the realizations in $\mathcal{K}$. Panel (a) refers to the first iteration of the algorithm, when the snapshots are taken only from one realization
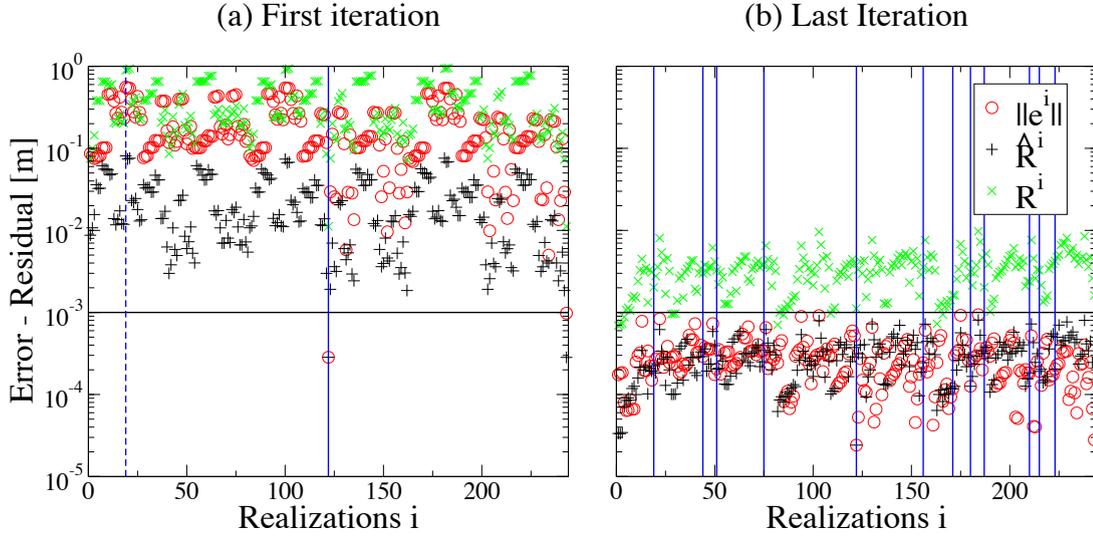
Figure 4.2: Norm of the error $\|e^i\|$, scaled residual norm $\hat{R}^i$, and residual norm R^i associated with all the realizations in $\mathcal{K}$ at the first and last iteration of the greedy algorithm in TC1. The continuous vertical blue lines indicate the realizations selected by the greedy algorithm for the computation of the snapshots. The dashed vertical blue line indicates the realization with the maximum scaled residual after the first iteration of the greedy algorithm.

(solid vertical line) and new snapshots are computed for the realization with maximum $\hat{R}^i$ (dotted vertical line). Panel (b) shows the last iteration. The horizontal line represents the error tolerance $\tau = 10^{-3}$. The continuous vertical lines indicate the realization where the FSM solution is computed. We can see that, in this first iteration, the RM accurately approximates only the realization from which the snapshots are selected and is not parameter independent (because the error is above the tolerance value). Moreover, the scaled norm of the residual underestimates the norm of the error, which is actually closer to R^i. This is due to the fact that we are computing the scaling factors ρ^i, using only one correct value ρ^*, and the large value chosen for λ imposes almost the same value of the scaling factor to all the realizations. Larger values of λ can cause a gross underestimation of the error, with the risk of falsely validating the RM when the real errors may still be very large. Panel (b) shows the results from the last iteration of Algorithm 2. In this case the RM is more accurate for all the parameter values in $\mathcal{K}$, as it consists of 30 principal components obtained from 12 FSM solutions. The scaled norm of the residual, $\hat{R}^i$, provides a good estimate of the norm of the errors, as both these values are below the required tolerance. In contrast, the norm of the residual, R^i, overestimates the error. This shows that if the validation condition is based only on the unscaled residual R^i, the greedy algorithm may proceed with the calculation of additional principal components
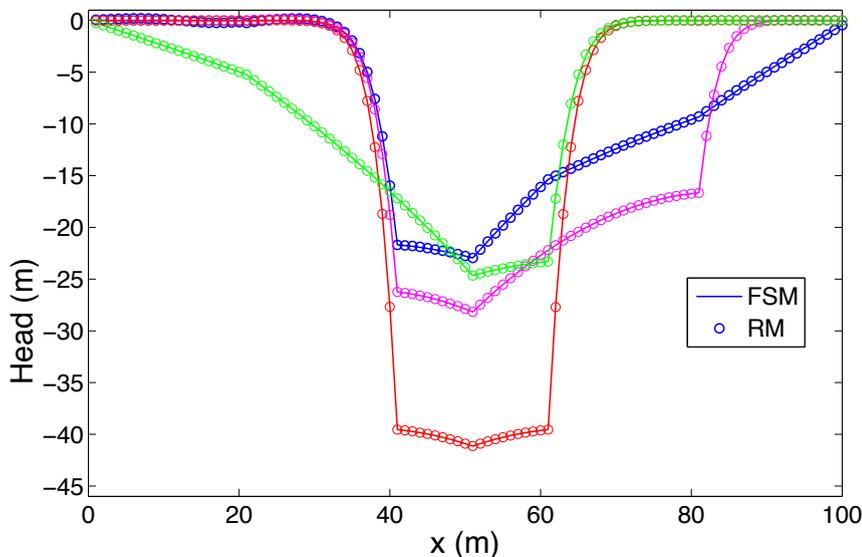
Figure 4.3: Comparison between four FSM solutions and the respective RM solutions for TC1.

that are unnecessary.

Fig. 4.3 compares the FSM and the RM head for four realizations of the hydraulic conductivity in $\mathcal{K}$ at time T_F. Although the four FSM solutions exhibit very different profiles, the solutions obtained from the RM, utilizing the 30 parameter-independent principal components computed in the offline algorithm, are indistinguishable when compared with the FSM solutions.

Finally, to demonstrate the applicability of the RM to MC simulations in TC1, we compare the ensemble statistics produced by the FSM and the RM solutions. For this purpose, we compute the FSM and the RM covariance matrices of the so-called augmented state, the vector of the nodal solutions $\mathbf{s}(t, \mathbf{k})$ augmented with the associated vector of the hydraulic conductivity $\mathbf{k}$. We indicate with $\mathbf{z}(t, \mathbf{k})$ the vector of the augmented state,

$$\mathbf{z}(t, \mathbf{k}) = (z_1, \ldots, z_{n+N_z}) = (s_1(t, \mathbf{k}), \ldots s_n(t, \mathbf{k}), k_1, \ldots, k_{N_z}).$$

Fig. 4.4 shows the matrix of the correlation coefficients associated with the augmented state of the system, computed with the FSM at the final time (T_F). Each element (i, j) of the correlation matrix is computed by the following equation:

$$Corr(z_i, z_j) = \frac{Cov(z_i, z_j)}{\sigma_{z_i} \sigma_{z_j}},$$
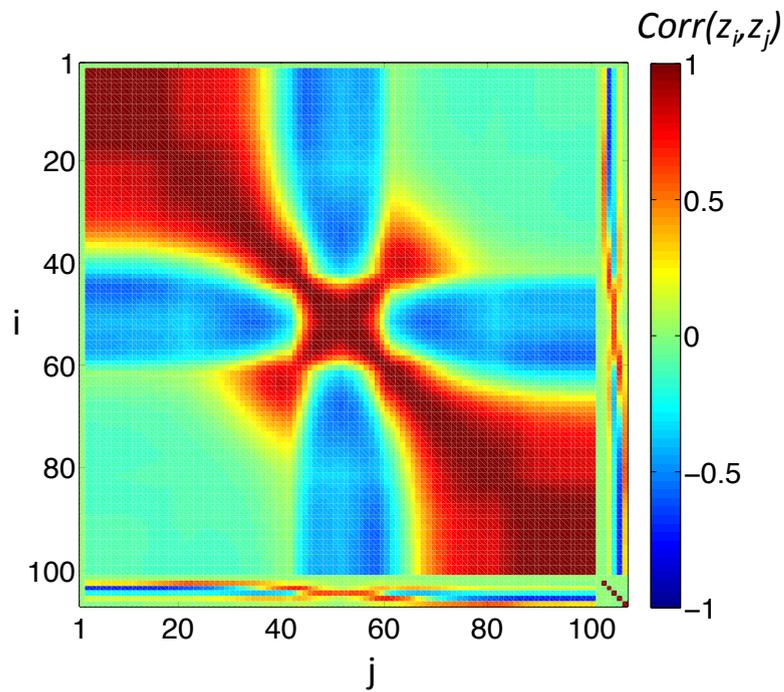
Figure 4.4: Matrix of the correlation coefficients of the extended state associated with the ensemble of FSM solutions for TC1.
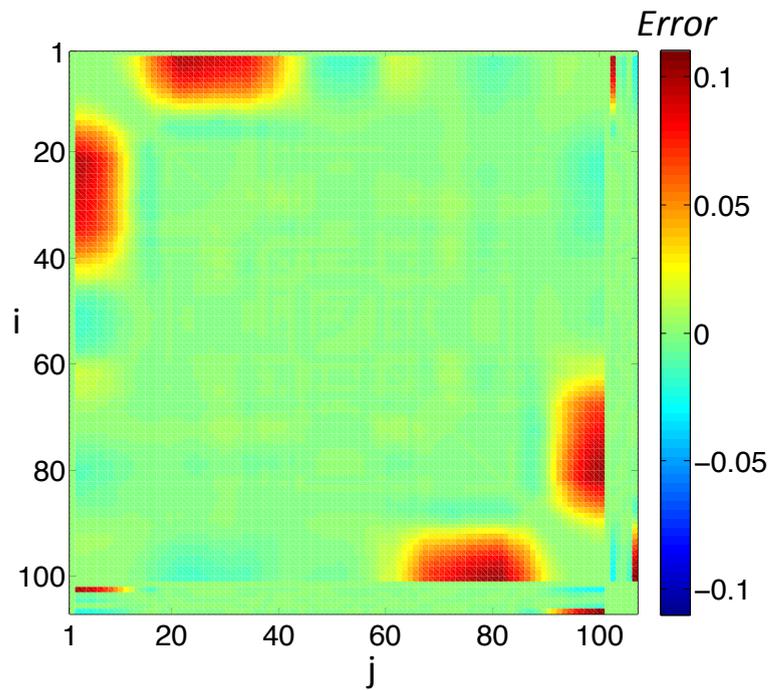
Figure 4.5: Errors between the correlation coefficients computed with the ensemble associated with the FSM and RM solutions for TC1.

where *Corr* is the correlation, *Cov* is the covariance and σ is the standard deviation. Fig. 4.5 shows the errors between the correlation coefficients computed with the FSM and the correlation coefficients computed with the RM. We can see that the statistics obtained by the RM are consistent with the FSM for almost all nodes. Small errors are detected on the Dirichlet boundaries. In fact, the RM statistics slightly overestimate the correlation coefficients between the nodes in the first and second zones and the hydraulic conductivity of the first zone. We observe similar results for zone 5. This behavior is due to the small drawdown values near the boundary and to the oscillatory nature of the principal components. To avoid large errors, the RM coefficients are computed in such a way as to match the FSM solution at nodes with large drawdown, i.e., in the neighborhood of the pumping well. As a consequence, the RM solution can be less accurate where drawdown is small. However, we note that the errors reported in Fig. 4.5 are relatively small with respect to the true values of the correlation coefficients shown in Fig. 4.4. This implies that the RM can be used to perform MC simulation.

4.5.2 Two-dimensional model

In this section we investigate the applicability of the RM to a more complex two-dimensional model, developed to simulate the groundwater flow in a confined aquifer located in the Oristano plain, in west-central Sardinia, Italy [Cau et al., 2002, Siade et al., 2012]. The domain of the Oristano aquifer is discretized using 29197 nodes and 57888 triangular elements. The comparison between the FSM and the RM is performed on the basis of the results obtained in three synthetic test cases (TC2, TC3 and TC4). In these test cases we consider a pumping test with a duration of four days. There are six pumping wells extracting at a constant rate of q=1000 m^3/d each with zero Dirichlet boundary conditions imposed. The aquifer, shown in Fig. 4.6, has a constant thickness (b=110 m) and specific storage (S_s=10^{-5} m^{-1}). Fig. 4.7 shows the zonation patterns used to model the heterogeneous hydraulic conductivity, with three zones in TC2, seven zones in TC3 and 15 zones in TC4. The lower and upper bounds of the hydraulic conductivities are $k_i^{min} = 0.1$ m/d and $k_i^{max} = 20$ m/d in each zone, respectively. The numerical solution of the FSM requires a CPU time of about 45 seconds for each realization of the hydraulic conductivity. As a consequence, the CPU time required by the Monte Carlo method with an ensemble size of N_{ens}=1000 is about 14 h (taking into account the time for the computation of the ensemble statistics). In this situation, the advantage of using the RM for MC simulation is evident. Tables 4.2, 4.3 and 4.4 compare the CPU times
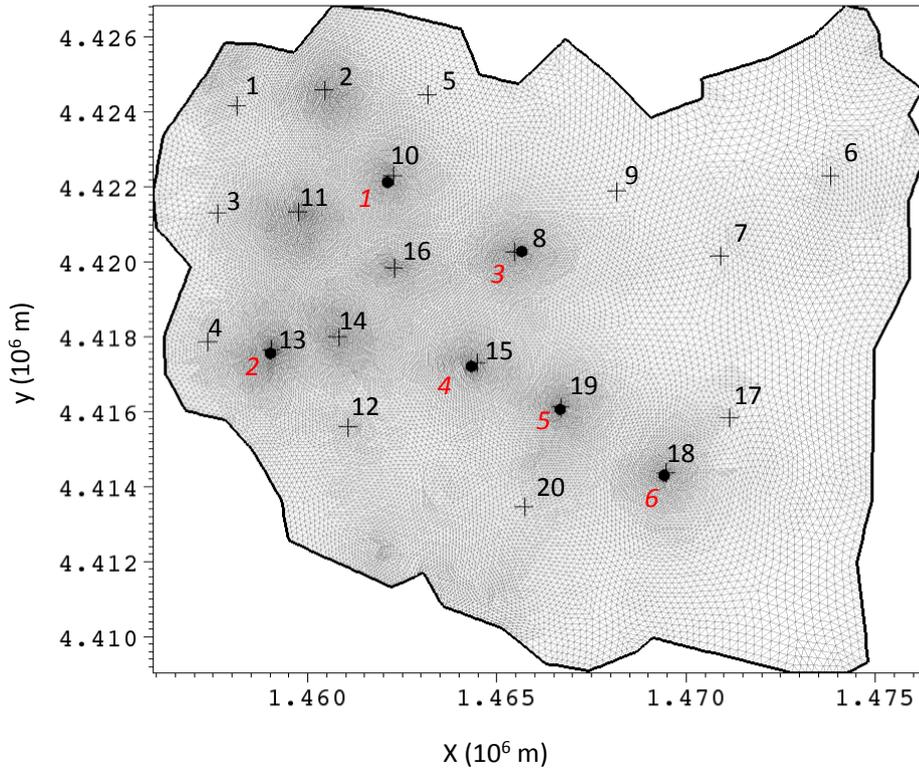
Figure 4.6: Oristano model with the location of the six pumping wells (dots) and the 20 observation wells (crosses).
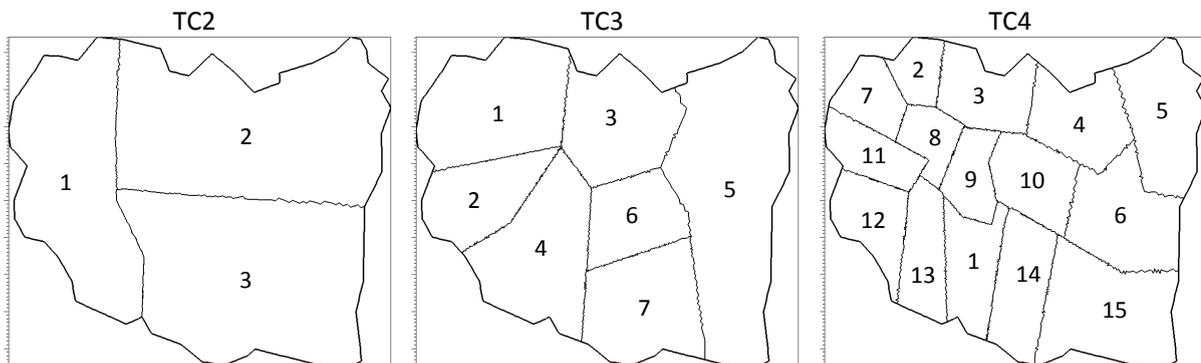
Figure 4.7: Zonation patterns used for TC2 (three zones), TC3 (seven zones) and TC4 (15 zones) for the Oristano model.

					One realization (s)	Total time (s)
FSM	$N_{ens}=$	1000	$n=$	29197	45	50000
RM online	$N_{ens}=$	1000	$N_{PC}=$	28	0.0025	15
RM offline	$N_k=$	27	$N_{gr}=$	13		890
RM total						905

Table 4.2: Comparison between the FSM and RM CPU times for the Monte Carlo simulation for TC2.

for the FSM and the RM for TC2, TC3, and TC4, respectively. The offline process is performed with a tolerance $\tau_e = 10^{-3}$ m. For TC2 the validation set consists of $N_k = 27$ realizations of hydraulic conductivity. Using Algorithm 2, we select snapshots from 13 of these realizations and compute 545 RM solutions, for a total offline CPU time of 890 seconds. The final number of principal components is $N_{PC} = 28$, and the corresponding RM is solved in 0.25×10^{-2} seconds (18000 times faster than the FSM). The MC simulation with the RM requires 15 seconds which, when added to the offline process, yields a total computational time of 905 seconds (55 times faster than the FSM). In TC3, the larger number of random parameters implies a larger variability in the MC solutions. For this reason, it is reasonable to enlarge the set $\mathcal{K}$ for the validation of the RM, with the consequent offline process that is computationally more expensive than in TC2. Using $N_k = 2187$ combinations of hydraulic conductivity, Algorithm 2 requires 31 FSM solutions and 123361 RM runs for a total CPU time of 4115 seconds. The resulting RM consists of 76 principal components and is solved in 0.13×10^{-1} seconds (3400 times faster than the FSM). The final time for the offline and online processes is 4141 seconds. In TC4, the validation of the RM on all the 3^{15} combinations of hydraulic conductivity becomes impractical. In this case we apply the offline process directly to the 1000 hydraulic conductivity realizations used in the MC method. Due to the large number of parameters, the offline process requires more FSM runs ($N_{gr} = 65$) than TC3, with a CPU time of 7132 seconds. The final number of principal components used in the RM is 71 and the total time to apply the MC method is 7158 seconds.

To verify the accuracy of the final RM on the single realizations, we compute the error $\mathbf{e}(T_F, \hat{\mathbf{k}}^j)$ for all the realizations in $\mathcal{K}$. Figs. 4.8, 4.9 and 4.10 show the head contours at the final simulation time (T_F) obtained with the FSM and RM for the realization with the maximum norm of the error (panel (a)) and the maximum nodal error (panel (b)) in TC2, TC3, and TC4, respectively. For example, in TC2 the realization with the maximum erro in norm (Fig. 4.8.a) has hydraulic conductivity values $K_1 = 0.1$ m/d, $K_2 = 10$ m/d,

(a)   (b)
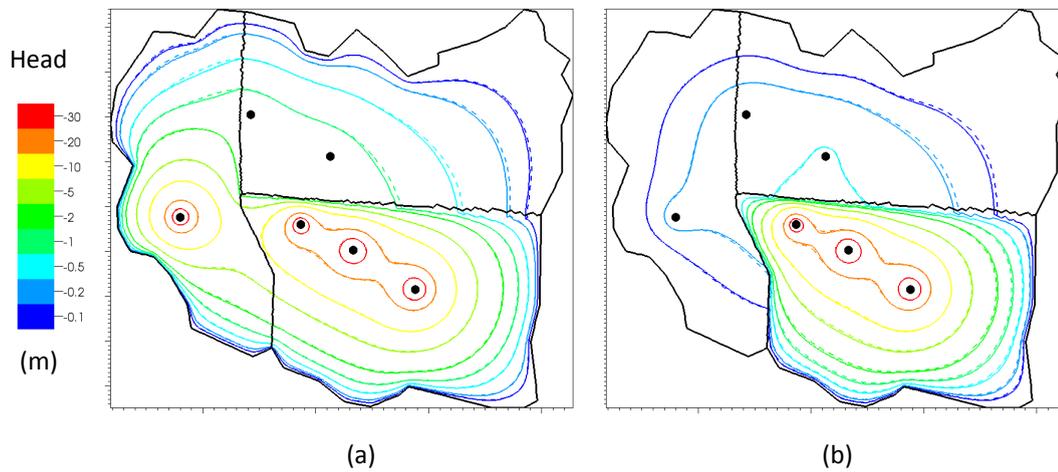
Figure 4.8: Comparison between the head obtained by solving the FSM (continuous line) and RM (dashed line) for two realizations of the hydraulic conductivity for TC2. The conductivity values in panels (a) and (b) are (0.1, 10, 0.1) m/d and (20, 20, 0.1) m/d, respectively.

(a)   (b)
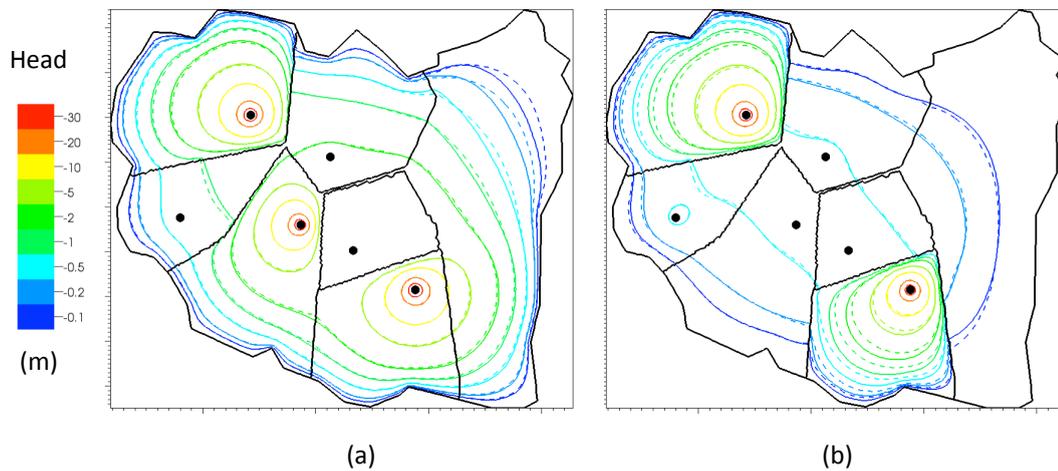
Figure 4.9: Analogous to Fig. 4.8 for TC3.

					One realization (s)	Total time (s)
FSM	$N_{ens}=$	1000	$n=$	29197	45	50000
RM online	$N_{ens}=$	1000	$N_{PC}=$	76	0.013	26
RM offline	$N_k=$	2187	$N_{gr}=$	31		4115
RM total						4141

Table 4.3: Comparison between the FSM and RM CPU times for the Monte Carlo simulation for TC3.
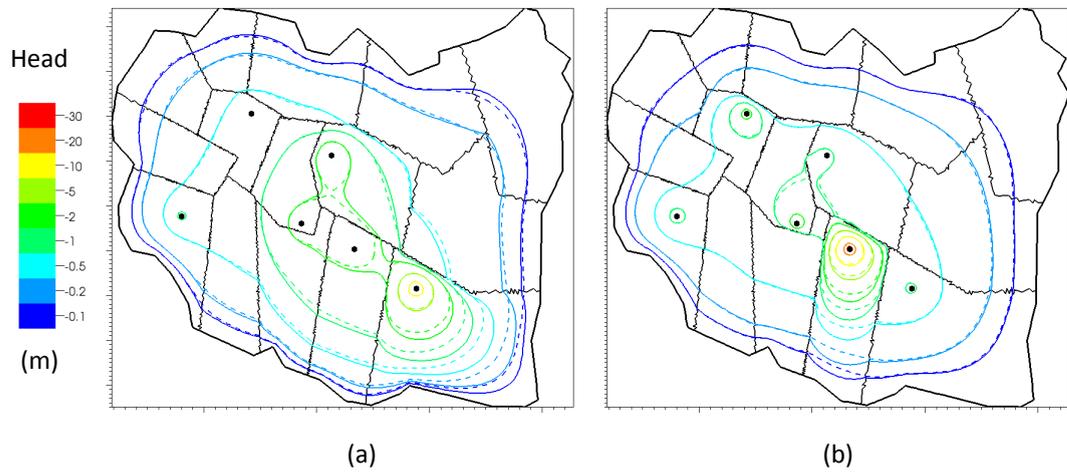
Figure 4.10: Analogous to Fig. 4.8 for TC4.

					One realization (s)	Total time (s)
FSM	$N_{ens}=$	1000	$n=$	29197	45	50000
RM online	$N_{ens}=$	1000	$N_{PC}=$	71	0.01	26
RM offline	$N_k=$	1000	$N_{gr}=$	65		7132
RM total						7158

Table 4.4: Comparison between the FSM and RM CPU times for the Monte Carlo simulation for TC4.
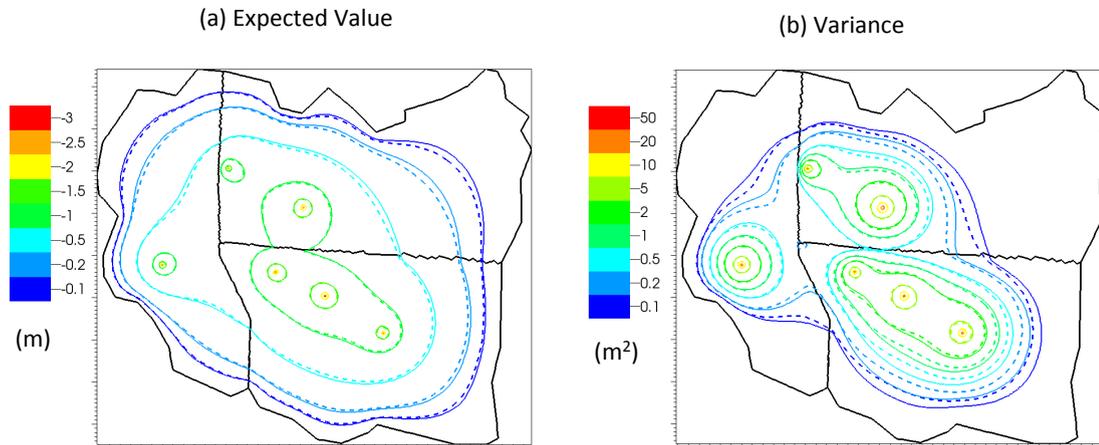
Figure 4.11: Expected value (panel (a)) and variance (panel (b)) of the ensemble of the drawdown obtained with the FSM (continuous lines) and the RM (dashed lines). Results for TC2.

and $K_3 =$ 0.1 m/d, while the realization with the maximum nodal error (Fig. 4.8.b) has hydraulic conductivity values $K_1 =$ 20 m/d, $K_2 =$ 20 m/d, and $K_3 =$ 0.1 m/d. The results show that, even as the head drastically changes for different combinations of the hydraulic conductivity, the RM solutions compare favorably with the FSM solutions. Similar to the results obtained in the one dimensional test case, we note that the RM solution is most accurate in the neighborhood of the pumping wells, i.e., in the regions with larger drawdown, while it becomes less accurate near the boundary, where drawdown is small.

Since the RM solutions do not perfectly match the FSM solutions for each realization, we are now interested in evaluating how these errors affect the leading statistical moments of the head that are usually approximated with the MC methods. With this purpose, in Fig. 4.11, 4.12, and 4.13 we compare the expected value and the variance of the ensemble of the FSM and the RM solutions for TC2, TC3, and TC4, respectively. The results show that the RM satisfactorily produces the mean head field in all the test cases. However, errors in the variance are larger and, as expected from the previous results, there is an underestimation of the variance of the head in areas far from the pumping wells.

Finally, we analyze the empirical distribution of head at 20 observation wells indicated in Fig. 4.6. To compare the data obtained with the FSM and RM, we apply the two-sample Kolmogorov-Smirnov test (KS) with the null hypothesis that the two ensembles are from the same continuous distribution. The test is performed at several output times (1 m, 5 m, 10 m, 30 m, 1 h, 2 h, 6 h, 12 h, 1 d, 2 d, 3d and 4 d), and the results are

(a) Expected Value

(b) Variance
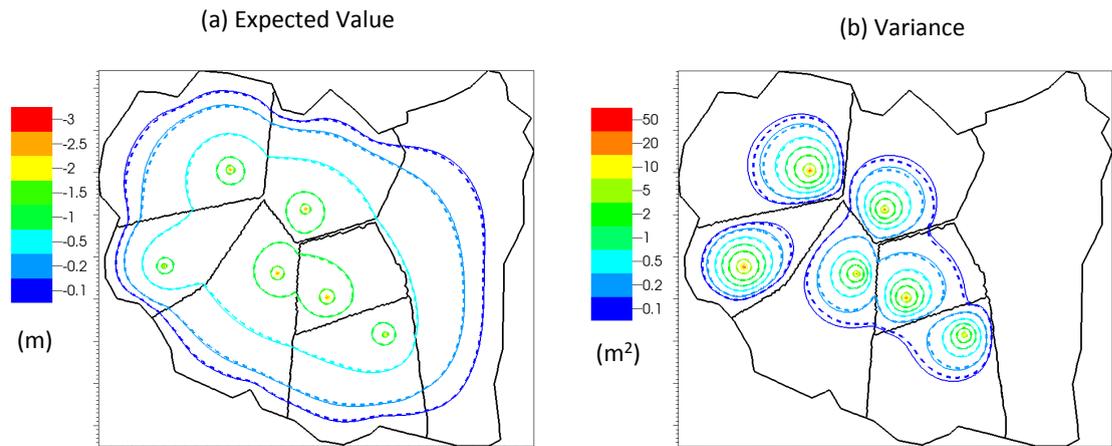
Figure 4.12: Analogous to Fig. 4.11 for TC3.

(a) Expected Value

(b) Variance
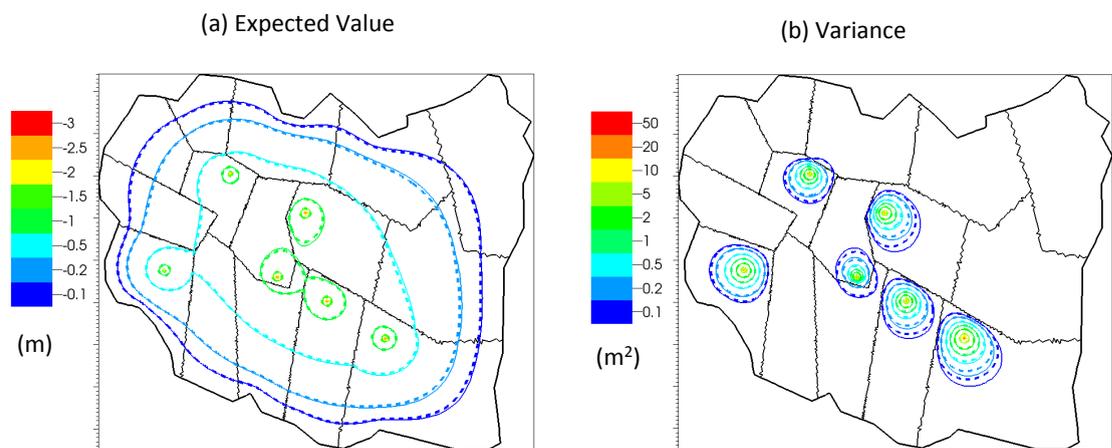
Figure 4.13: Analogous to Fig. 4.11 for TC4.

Time	1 m	5 m	10 m	30 m	1 h	2 h	6 h	12 h	24 h	48 h	72 h	96 h
Well #												
1	1	1	1	1	1	1	0	0	0	0	0	0
2	1	1	1	1	0	0	0	0	0	0	0	0
3	1	1	1	1	1	1	0	0	0	0	0	0
4	1	1	1	1	0	0	0	0	0	0	0	0
5	1	1	1	1	1	1	0	0	0	0	0	0
6	1	1	1	1	1	1	1	0	0	0	0	0
7	1	1	1	1	1	1	0	0	0	0	0	0
8	1	0	0	0	0	0	0	0	0	0	0	0
9	1	1	1	1	1	1	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0	0	0
11	1	1	1	0	0	0	0	0	0	0	0	0
12	1	1	1	1	0	0	0	0	0	0	0	0
13	0	0	0	0	0	0	0	0	0	0	0	0
14	1	1	1	0	0	0	0	0	0	0	0	0
15	1	0	0	0	0	0	0	0	0	0	0	0
16	1	1	1	0	0	0	0	0	0	0	0	0
17	1	1	1	1	1	1	0	0	0	0	0	0
18	0	0	0	0	0	0	0	0	0	0	0	0
19	0	0	0	0	0	0	0	0	0	0	0	0
20	1	1	1	1	1	1	0	0	0	0	0	0

Table 4.5: Results of the two-sample Kolmogorov-Smirnov test to compare the distributions of the ensemble values of drawdown obtained with the FSM and the RM at the observation wells (rows), for different output times (columns). The null hypothesis is that the two ensembles are from the same continuous distribution. The result is 1 if the test rejects the null hypothesis, 0 otherwise. Results refer to TC2.

Time	1 m	5 m	10 m	30 m	1 h	2 h	6 h	12 h	24 h	48 h	72 h	96 h
Well #												
1	1	1	1	1	1	1	1	1	1	1	1	1
2	1	1	1	1	1	1	0	0	0	0	0	0
3	1	1	1	1	1	1	0	0	0	0	0	0
4	1	1	1	0	0	0	0	0	0	0	0	0
5	1	1	1	1	1	0	0	0	0	0	0	0
6	1	1	1	1	1	1	1	0	0	1	1	1
7	1	1	1	1	1	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0	0
9	1	1	1	1	1	1	0	0	0	0	0	0
10	1	0	0	0	0	0	0	0	0	0	0	0
11	1	1	1	0	1	1	0	0	0	0	0	0
12	1	1	1	1	1	1	0	0	0	0	0	0
13	0	0	0	0	0	0	0	0	0	0	0	0
14	1	1	1	1	1	1	0	0	0	0	0	0
15	1	0	0	0	0	0	0	0	0	0	0	0
16	1	1	1	0	0	0	0	0	0	0	0	0
17	1	1	1	1	0	1	0	0	0	0	0	0
18	0	0	0	0	0	0	0	0	0	0	0	0
19	0	0	0	0	0	0	0	0	0	0	0	0
20	1	1	1	1	1	1	0	0	0	0	0	0

Table 4.6: Analogous of Table 4.5. Results refer to TC3.

Time	1 m	5 m	10 m	30 m	1 h	2 h	6 h	12 h	24 h	48 h	72 h	96 h
Well #												
1	1	1	1	1	1	1	0	0	1	1	1	1
2	1	1	1	1	1	1	0	0	0	0	0	0
3	1	1	1	1	1	1	0	1	1	1	1	1
4	1	1	1	1	1	1	0	0	0	0	0	0
5	1	1	1	1	1	1	0	0	0	0	0	0
6	1	1	1	1	1	1	1	1	1	1	1	1
7	1	1	1	1	1	0	0	0	0	0	0	0
8	1	0	0	0	0	0	0	0	0	0	0	0
9	1	1	1	1	1	1	0	0	0	0	0	0
10	1	0	0	0	0	0	0	0	0	0	0	0
11	1	1	1	1	0	0	0	0	0	0	0	0
12	1	1	1	1	1	0	0	0	0	0	0	0
13	1	0	0	0	0	0	0	0	0	0	0	0
14	1	1	1	1	0	0	0	0	0	0	0	0
15	1	0	0	0	0	0	0	0	0	0	0	0
16	1	1	1	1	0	0	0	0	0	0	0	0
17	1	1	1	1	1	0	0	0	0	0	0	0
18	1	0	0	0	0	0	0	0	0	0	0	0
19	0	0	0	0	0	0	0	0	0	0	0	0
20	1	1	1	1	1	1	0	0	0	0	0	0

Table 4.7: Analogous of Table 4.5. Results refer to TC4.

collected in Tables 4.5, 4.6, and 4.7. We can see that the KS-test fails at the initial times when drawdown is small over the entire domain. Consistent with the previous results, the KS-test frequently fails on observations wells that are located closer to the boundary (e.g., wells number 1, 3 and 6). However, most importantly, the null hypothesis is validated at almost all output times for the observation wells that are in the neighborhood of the pumping wells (e.g., well numbers 8, 10, 13, 15, 18, 19).

Offline algorithm: dependency on $\mathcal{K}$

In this section we investigate the dependency of the RM accuracy to the validation set $\mathcal{K}$. We consider Algorithm 2 for TC2 with $\mathcal{K}$ given by: (i) the complete set of combinations of upper bound, lower bound and the mean of hydraulic conductivity ($N_k = 27$), represented by $\mathcal{K}^{up-low}$, and (ii) a set of random realizations, denoted by $\mathcal{K}^{MC}$, of sizes $N_k = 27$, 50, 100, 200, 500, and 1000. To avoid the dependency of the greedy algorithm to the scaled residual, the validation condition is based directly on the norm of the error. The resulting reduced order models are compared on the basis of the maximum error on the solutions associated with the set $\mathcal{L}$ of 1000 MC realizations of hydraulic conductivity (different from the realizations considered in the offline algorithm). Fig. 4.14 shows these maximum errors obtained for the different choices of $\mathcal{K}$ as a function of the dimension
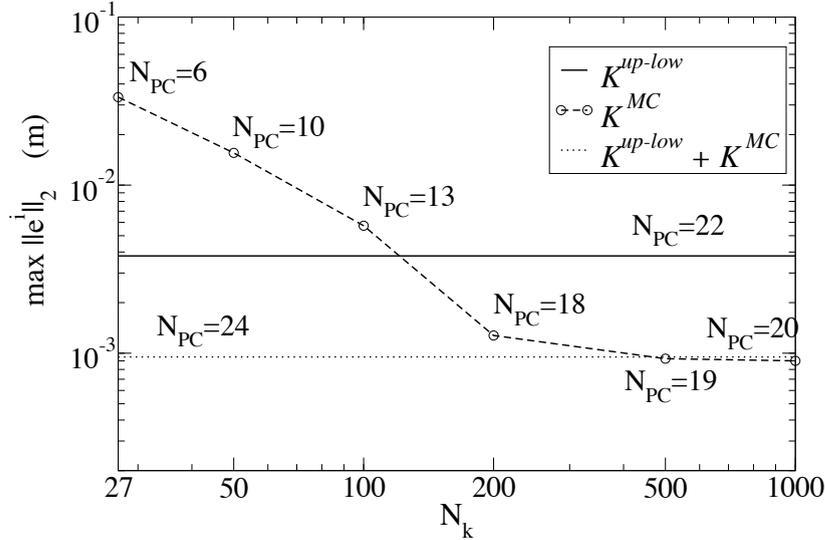
Figure 4.14: Maximum error of the RM on 1000 MC realizations, for different choices of the validation set $\mathcal{K}$ in TC2. The horizontal lines represent the error associated with $\mathcal{K} = \mathcal{K}^{up-low}$ (continuous line) and $\mathcal{K} = \mathcal{K}^{up-low} + \mathcal{K}^{MC}$ (dotted line). The dashed line shows the errors associated with $\mathcal{K} = \mathcal{K}^{MC}$, with N_k =27, 50, 100, 200, 500, and 1000.

N_k. Moreover, in Fig. 4.14 we highlight the number of principal components obtained at the conclusion of the greedy algorithm. The figure shows that our choice of using $\mathcal{K}^{up-low}$ in the offline algorithm provides better results than considering $\mathcal{K}^{MC}$ with 27, 50, and 100 random realizations. However, in TC2 $\mathcal{K}^{up-low}$ has only 27 realizations and is not sufficient to describe the probability space as accurately as $\mathcal{K}^{MC}$ with 200, 500, and 1000 realizations. To see if $\mathcal{K}^{up-low}$ is a reasonable choice for the validation set, we apply the offline algorithm to validate the RM on the set $\mathcal{K} = \mathcal{K}^{up-low} + \mathcal{K}^{MC}$ with 1000 random realizations. In this case the greedy algorithm computes the first 22 principal components from eight realizations belonging to $\mathcal{K}^{up-low}$, achieving a maximum norm of the error of 0.003 m on the entire $\mathcal{K}$. Only the last two principal components are computed from two realizations in $\mathcal{K}^{MC}$. Fig. 4.14 shows that the RM validated on the set $\mathcal{K} = \mathcal{K}^{up-low} + \mathcal{K}^{MC}$ also maintains the desired accuracy (errors smaller than the threshold τ_e) when applied to the independent set $\mathcal{L}$.
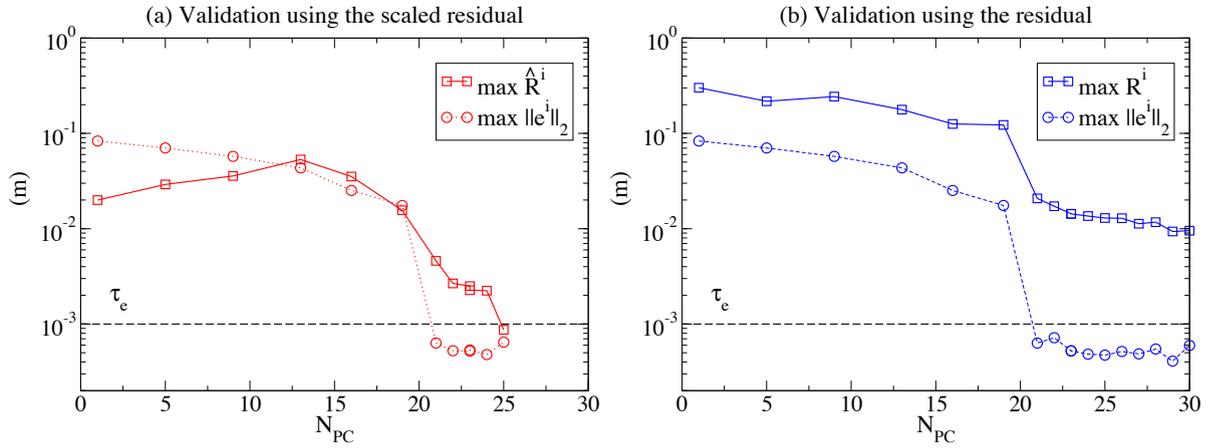
Figure 4.15: Comparison between the estimation of the error at each iteration of the greedy algorithm using the validation condition on the maximum-scaled norm of the residual $\hat{R}^i$ (a) and the maximum norm of the residual R^i (b). The dotted line shows the maximum norm of the error, while the horizontal dashed line is the threshold value τ_e for the termination of the greedy algorithm.

Offline algorithm: scaled residual

In this section we show the practical advantages of using the scaled residual for the estimation of the errors in the greedy algorithm. We consider Algorithm 2 in TC2 with the validation set $\mathcal{K} = \mathcal{K}^{up-low}$. The estimation of the error is performed with (a) the scaled norm of the residual $\hat{R}^i$ and (b) the norm of the residual R^i. Fig. 4.15 shows the main results of the application of these two greedy algorithms, comparing the maximum error with its estimation as a function of the number of the principal components. In both cases, the maximum error falls under the threshold value τ_e after six iterations of the greedy algorithm, with 20 principal components. Using the scaled norm of the residual $\hat{R}^i$ (Fig. 4.15.a) the error is underestimated in the first iterations, due to the lack of information for interpolating the scaling factors. However, in the subsequent iterations we obtain a good estimation of the error and the greedy algorithm stops with only five additional iterations. Instead, estimating the error with the residual R^i (Fig. 4.15.b) shows an overestimation of the error. As a consequence, a larger number of iterations of the greedy algorithm are performed, with unnecessary computation of additional FSM solutions and principal components.

4.6   Conclusions

We have presented a model order reduction technique that overcomes the computational burden associated with the MC simulations of confined groundwater flow models with stochastic hydraulic conductivity. We proposed a new offline algorithm (Algorithm 2) for the computation of parameter-independent principal components, which constitutes the core of the RM. In this offline process we combined a residual-based greedy algorithm for the selection of snapshots in the parameter space with a quasi-optimal method for the selection of snapshots in time. The algorithm starts with an initial set of principal components and improves this set until the RM solution is validated on an appropriately chosen set $\mathcal{K}$ of hydraulic conductivity realizations. The validation condition is provided by an estimation of the error based on the computation of the norm of the residual and a scaling coefficient (Eqs. (4.16) and (4.17)). The greedy algorithm selects the new snapshots corresponding to the hydraulic conductivity value that maximizes the estimation of the RM error (Eq. (4.18)). Then, the FSM solution is used for the selection of the snapshots at the optimal times given by Eq. (4.12). Finally, the principal component analysis performed on the set of snapshots produces the new principal components needed to improve the RM solution. This new methodology allows computing the principal components using only a few essential FSM runs, ensuring the computational efficiency of the algorithm and the accuracy of the RM solution. Once the principal components are computed, the application of the RM to the MC simulation is straightforward. Since the RM is parameter independent, it can be used efficiently to compute the main statistics associated with the ensemble of solutions, such as the expected value and the covariance of the head.

We first verified our methodology on a one-dimensional test case with five zones (TC1) and then applied it to a two-dimensional model of the Oristano aquifer with three, seven and 15 zones (TC2, TC3 and TC4). The RM solution was compared with the classical MC solution in terms of CPU times, expected head values, and head variance in the entire domain as well as the empirical probability distribution function of the head at the observation wells. The numerical results lead to the following conclusions.

1. Our algorithm reduced the 29197 nodes of the Oristano model to less than 100 principal components, with a corresponding RM solution that is at least 1000 times faster than the FSM solution.

2. The computation of the principal components in the offline process is the most expensive part of the procedure, as it requires a certain number of FSM solutions.

The number of FSM runs used by our offline process increased with the number of random parameters, due to the fact that the FSM solutions have more degrees of freedom and, as a consequence, the RM needs more information to cover the space of the solutions. Nevertheless, the numerical results demonstrate that, with 15 zones, our methodology (offline plus online) is more than 10 times faster than the standard MC approach. However, we wish to note that our ultimate goal is to develop a parameter-independent RM that can be used for fast online execution. This goal has been achieved, as demonstrated by an application of the proposed methodology to the Oristano aquifer in Italy, where the FSM was reduced by three orders of magnitude and ran 1000 faster than the FSM.

3. The comparison between the statistics of the ensemble of solutions given by the RM and FSM suggests that our RM accurately produces the expected value of the head over the entire domain, while it slightly underestimates the variance in the regions of small drawdown. In addition, the two-sided Kolmogorov-Smirnov test applied to the empirical distributions of the heads at the observation wells indicates that the RM is most accurate in the neighborhood of the pumping wells, i.e., where the response of the system to the forcing term is higher. We attribute this fact to the methodology that we employed for the generation of the principal components. In fact, with the principal component analysis we keep the principal components corresponding to the larger eigenvalues, which describe the dominating characteristics of the solution, while we discarded the principal components corresponding to small eigenvalues, which are only useful for a detailed description of the solution in regions of low variance and which correspond to the regions with low sensitivity to pumping.

Further research is needed to extend our methodology to compute the principal components for the case of spatially-distributed random hydraulic conductivities, particularly for aquifers with strong local heterogeneity, for which the interpolation scheme employed to evaluate the scaling factors used in Algorithm 2 needs to be completely revised. How to derive reduced models under such situations is a topic of future research.

Chapter 5

Summary and Conclusions

In this thesis work we presented data assimilation methods and model order reduction techniques to improve the numerical simulation of groundwater flows in porous media driven by stochastic parameters. In this chapter we summarize the research topics studied in this dissertation, and highlight the major achievements.

5.1 Data assimilation

Data assimilation techniques allow us to directly incorporate physical measurements into the transient hydrological model. These methods correct the forecast probability density function toward the observed values, taking into account the probability distribution of the measurements errors. The goal of data assimilation technique is to reduce the uncertainty associated to the model forecast and prevent the propagation in time of errors in the modelization, such as biased initial or boundary conditions. In particular, in this thesis we considered the ensemble Kalman filter and the sequential importance resampling, two data assimilation schemes that are based on the Monte Carlo simulations and developed for high-dimensional, nonlinear models. We applied these methods for the assimilation of pressure head and/or streamflow measurements into the model CATHY, which is a distributed physically-based hydrological model that couples surface and subsurface flow. The objective of the research is to better understand the theoretical and practical differences between the two filters. With this goal, we compared their performances on a synthetic test case in presence of different sources of errors in the model, such as uncertainties in soil parameters, atmospheric forcings, and initial conditions.

The major conclusion achieved in this first part of the thesis is that, both EnKF and SIR perform well in the application to high-dimensional hydrological models with a

reasonable small number of Monte Carlo realizations. For the test cases presented, SIR seems to outperform EnKF in the retrieval of the real states and in the computational effort required by the simulation. This is due to the slow convergence of the nonlinear numerical solver after each update of the EnKF. This is a direct consequence of the physical inconsistencies introduced by the Gaussian approximation embedded in the EnKF updates. The Gaussian approximation affects especially the state variables in the unsaturated zone, where the nonlinear Richards equation governs the dynamics of the system. SIR, instead, in the analysis step duplicates solutions of the model equation and thus is not affected by this issue.

We highlighted that, for numerical issues, the SIR method skips the updates where the whole ensemble of measurements are far from the probability distribution of the observed values. We proposed a modification of the SIR update, based on the amplification of the standard deviation of the measurement error. The results from the scenarios with biased initial conditions demonstrate that, with this modification, the filter SIR duplicates the realizations closest to the observations and performs a correct update.

5.2   Reduced Order Model

The large computational time required to generate the ensemble of Monte Carlo realizations is one of the main drawbacks of the data assimilation methods presented in Chapter 2, and compromises their applicability to real hydrological models. For this reason, in the second part of the thesis we explored the possibility of applying model order reduction techniques to speed the Monte Carlo simulation. The idea of model order reduction techniques is to approximate the solution of the full system model with the fast solution of a low-dimensional model. The reduced model equation is achieved by the Galerkin projection of the mathematical equations governing the hydrological model onto the space generated by a low number of basis functions. Since the basis functions are computed using the snapshot technique, the reduced model solution reproduces the main features of the physical process.

In Chapter 3, we implemented a reduced order model for the solution of a synthetically-reconstructed, regional-scale groundwater flow problem where the recharge is modeled as spatially-distributed stochastic forcing term. We showed that the key geostatistical parameters of the distributed recharge play a fundamental role in determining the accuracy of the reduced model solution. In particular, we demonstrated that a random recharge with large coefficients of variation and small integral scales entails using a large number

of snapshots and principal components to reproduce the results of the standard Monte Carlo simulation. As a consequence, in some cases the methodology presented does not provide any reduction of the Monte Carlo computational times. On the contrary, when the recharge is characterized by small coefficients of variation and large integral scales, our Monte Carlo simulation based on the reduced model can improve the computational efficiency of a standard Monte Carlo framework up to a factor of 100. In this case, the reduced order model keeps remarkable degree of accuracy in approximating the mean and covariance of the probability density function of hydraulic heads.

Assessing the number of snapshots and principal components to use in the reduction is a crucial point to construct an efficient reduced order model. Thus, in Chapter 4 we focused our attention on the offline procedure for the automatic selection of the snapshots. We proposed an offline algorithm that combines a residual-based greedy algorithm for the selection of snapshots in the parameter space with a quasi-optimal method for the selection of snapshots in time. In the greedy algorithm an estimate of the error is used to determine the optimal snaphots. The error is generally evaluated by comparing the full and reduced model solutions, and is thus the most expensive phase of the algorithm. To decrease the computational effort, we introduced an estimation of the error based on the computation of the norm of the residual and a scaling coefficient. We applied this algorithm to simulate flow in a confined aquifer with zonal stochastic hydraulic conductivity in a synthetic and a realistic scenario. The reduced model solution was compared with the classical Monte Carlo solution in terms of CPU times, and accuracy on the reconstruction of the expected head values, and head variance in the entire domain as well as the empirical probability distribution function of the head at the observation wells. The numerical results demonstrated that our methodology allows computing the principal components using only a few essential full system model runs, ensuring the computational efficiency of the algorithm and the accuracy of the reduced model solution. We highlighted that the number of independent random parameters affects the computational complexity of the offline procedure. In fact, the number of full system model runs used by our offline process increased with the number of random parameters, due to the fact that the corresponding solutions have more degrees of freedom and, as a consequence, the reduced model needs more information to cover the space of the solutions. However, the application of the reduced model methodology to the Monte Carlo simulation in the real case scenario is 1000 times faster than applying the full system model solution.

5.3   Challenges for future research

In summary, the research work presented in this thesis lays the foundation for a practical implementation of efficient ensemble stochastic simulations applied to complex hydrological models. It achieves this goal by combining data assimilation methods with model order reduction algorithms. A number of issues remain open, however, to exploit the full potential of these methodologies. For example, there is ongoing work to use model order reduction methods to ease the computational burden of Monte Carlo simulations in the case of spatially heterogeneous, randomly distributed hydraulic conductivity in groundwater flow equations. While this issue is easily solved with current approaches for large integral scales, the performance of model order reduction methods quickly deteriorates at smaller spatial scales, i.e., when a low number of principal components cannot completely describe the high variability in space of the solution. Potential improvements could be obtained by combining model order reduction with moment differential equations (see, e.g., Guadagnini and Neuman [1999]) or other means of upscaling. However, the most effective way to achieve this result is still under investigation: an idea we are currently exploring is to use directly the correlation matrix obtained with the moment differential equations for the exact computation of the principal components, avoiding the selection of the snapshots. Moreover, the reduced order model can be applied directly to the moment differential equations for transient problems. Another important topic that must be addressed to overcome the difficulties in the evaluation of the ensemble members is the application of reduced order models to highly nonlinear problems such as the Richards equation. The nonlinearities affect the reduction methodology in two main aspects: on one hand, the nonlinear operator is a function of the solution in the full system state, leading to incomplete reduction of the problem. On the other hand, the errors introduced by the POD approximation may propagate in the nonlinear operator, causing a deterioration of the solution or even divergence phenomena. Recent advances in the construction of efficient nonlinear reduced order models, such as the masked projection of the discrete model equations proposed by Galbally et al. [2010], and the possibility to correct the reduced model forecast using the assimilated measurements, open up an important field of research and motivate additional studies on the development of efficient data assimilation schemes for hydrological models.

Bibliography

Arulampalam, M. S., Maskell, S., Gordon, N., and Clapp, T. (2002). A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE T. Signal Proces.*, 50(2):174–188.

Arulampalam, M. S. and Ristic, B. (2000). Comparison of the particle filter with range-parametrized and modified polar EKFs for angle-only tracking. In *Proc. SPIE*, volume 4048, pages 288–299.

Aubert, D., Loumagne, C., and Oudin, L. (2003). Sequential assimilation of soil moisture and streamflow data in a conceptual rainfall-runoff model. *J. Hydrol.*, 280(1-4):145–161.

Ballio, F. and Guadagnini, A. (2004). Convergence assessment of numerical Monte Carlo simulations in groundwater hydrology. *Water Resour. Res.*, 40:W04603.

Baú, D. A. (2012). Planning of groundwater supply systems subject to uncertainty using stochastic flow reduced models and multi-objective evolutionary optimization. *Water Resour. Manag.*, 26(9):2513–2536.

Bear, J. (1979). *Hydraulics of Groundwater*. McGraw-Hill Inc., New York.

Bellin, A. and Rubin, Y. (1996). HYDRO_GEN: A spatially distributed random field generator for correlated properties. *Stoch. Hydrol. Hydraul.*, 10(4):253–278.

Binley, A. and Kemma, A. (2005). Electrical methods. In Rubin, Y. and Hubbard, S. S., editors, *Hydrogeophysics*, volume 50, pages 129–156. Springer.

Bui-Thanh, T., Willcox, K., and Ghattas, O. (2007). Model reduction for large-scale systems with high-dimensional parametric input space. *SIAM J. Sci. Comput.*, 30(6):3270–3288.

Camporese, M., Paniconi, C., Putti, M., and Orlandini, S. (2010). Surface-subsurface flow modeling with path-based runoff routing, boundary condition-based coupling, and assimilation of multisource observation data. *Water Resour. Res.*, 46:W02512.

Camporese, M., Paniconi, C., Putti, M., and Salandin, P. (2009a). Comparison of data assimilation techniques for a coupled model of surface and subsurface flow. *Vadose Zone J.*, 8(4):837–845.

Camporese, M., Paniconi, C., Putti, M., and Salandin, P. (2009b). Ensemble Kalman filter data assimilation for a process-based catchment scale model of surface and subsurface flow. *Water Resour. Res.*, 45:W10421.

Cau, P. L., Lecca, G., Putti, M., and Paniconi, C. (2002). The influence of a confining layer on saltwater intrusion and surface recharge and groundwater extraction conditions. In Hassanizadeh, S. M., et al., editor, *Computational Methods in Water Resources, Developmetsn in water resources*, volume 1, pages 493–500. Elseveier, Amsterdam, Netherlands.

Chang, H., Chen, Y., and Zhang, D. (2010). Data assimilation of coupled fluid flow and geomechanics using the ensemble Kalman filter. *Spe J.*, 15(2):382–394.

Chen, Y. and Zhang, D. (2006). Data assimilation for transient flow in geologic formations via ensemble Kalman filter. *Adv. Water Resources*, 29(8):1107–1122.

Clark, M. P., Rupp, D. E., Woods, R. A., Zheng, X., Ibbitt, R. P., Slater, A. G., Schmidt, J., and Uddstrom, M. J. (2008). Hydrological data assimilation with the ensemble Kalman filter: Use of streamflow observations to update states in a distributed hydrological model. *Adv. Water Resources*, 31(10):1309–1324.

Copty, N. K. and Findikakis, A. N. (2004). Bayesian identification of the local transmissivity using time-drawdown data from pumping tests. *Water Resour. Res.*, 40:W12408.

Crisan, D. (2001). Particle filters - A theoretical perspective. In Doucet, A., et al., editor, *Sequential Monte Carlo Methods in Practice*, pages 17–41. Springer-Verlag, New York.

Dagan, G. (1982). Stochastic modeling of groundwater flow by uncoditional and conditional probabilities: 1. Conditional simulation and the direct problem. *Water Resour. Res.*, 18(4):813–833.

Dagan, G. (1989). *Flow and Transport in Porous Formations*. Springer, New York.

Dagan, G., Lessoff, S. C., and Fiori, A. (2009). Is transmissivity a meaningful property of natural formations? Conceptual issues and model development. *Water Resour. Res.*, 45:W03425.

Dagan, G. and Neuman, S. P., editors (1997). *Subsurface Flow and Transport: A Stochastic Approach*. Cambridge Univ. Press, New York.

Das, N. N. and Mohanty, B. P. (2006). Root zone soil moisture assessment using remote sensing and vadose zone modeling. *Vadose Zone J.*, 5(1):296–307.

DeChant, C. M. and Moradkhani, H. (2011a). Improving the characterization of initial condition for ensemble streamflow prediction using data assimilation. *Hydrol. Earth Syst. Sc.*, 15:3399–3410.

DeChant, C. M. and Moradkhani, H. (2011b). Radiance data assimilation for operational snow and streamflow forecasting. *Adv. Water Resources*, 34(3):351–364.

Doucet, A., Godsill, S., and Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat. Comput.*, 10(3):197–208.

Drecourt, J. P., Madsen, H., and Rosbjerg, D. (2006). Calibration framework for a Kalman filter applied to a groundwater model. *Adv. Water Resources*, 29(5):719–734.

Entekhabi, D., Nakamura, H., and Njoku, E. G. (1994). Solving the inverse problems for soil-moisture and temperature profiles by sequential assimilation of multifrequency remotely-sensed observations. *IEEE T. Geosci. Remote*, 32(2):438–448.

Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.-Oceans*, 99(C5):10143–10162.

Evensen, G. (2003). The ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dyn.*, 53(4):343–367.

Evensen, G. (2007). *Data assimilation: The ensemble Kalman filter*. Springer, New York.

Evensen, G. (2009). The ensemble Kalman filter for combined state and parameter estimation. *IEEE CONTR. Syst. Mag.*, 29(3):83–104.

Fontaine, V., Mara, T., and Mamode, M. (2010). Probabilistic collocation for efficient uncertainty analysis in groundwater flow. In Carrera, J., editor, *XVIII International Conference on Water Resources, Proceedings*, Barcelona, Spain.

Furman, A. (2008). Modeling coupled surface–subsurface flow processes: a review. *Vadose Zone J.*, 7(2):741–756.

Galbally, D., Fidkowski, K., Willcox, K., and Ghattas, O. (2010). Non-linear model reduction for uncertainty quantification in large-scale inverse problems. *Int. J. Numer. Methods Engrg.*, 81(12):1581–1608.

Gambolati, G., Putti, M., and Paniconi, C. (1999). Three-dimensional model of coupled density-dependent flow and miscible salt transport in groundwater. In Bear, J., et al., editor, *Seawater intrusion in coastal aquifers: concepts, methods and practice*, pages 315–362. Kluver Acad., Dordrecht, Netherlands.

Ghanem, R. . and Spanos, P. D. (1991). *Stochastic Finite Elements: A Spectral Approach*. Springer, New York.

Giustarini, L., Matgen, P., Hostache, R., Montanari, M., Plaza, D., Pauwels, V. R. N., De Lannoy, G. J. M., De Keyser, R., Pfister, L., Hoffmann, L., and Savenije, H. H. G. (2011). Assimilating SAR-derived water level data into a hydraulic model: a case study. *Hydrol. Earth Syst. Sc.*, 15:2349–2365.

Gohberg, I. and Goldberg, S. (1981). *Basic Operator Theory*. Birkhäuser, Boston.

Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc.-F*, 140(2):107–113.

Grepl, M. A. and Patera, A. T. (2005). A posteriori error bounds for reduced-basis approximations of parametrized parabolic partial differential equations. *ESAIM-Math. Model. Num.*, 39(1):157–181.

Guadagnini, A. and Neuman, S. P. (1999). Nonlocal and localized analyses of conditional mean steady state flow in bounded, randomly nonuniform domains: 2. Computational examples. *Water Resour. Res.*, 35(10):3019–39.

Haasdonk, B. and Ohlberger, M. (2011). Efficient reduced models and a-posteriori error estimation for parametrized dynamical systems by offline/online decomposition. *Math. Comp. Model. Dyn.*, 17(2):145–161.

Hasenauer, J., Löhning, M., Khammash, M., and Allgöwer, F. (2012). Dynamical optimization using reduced order models: a method to guarantee performance. *J. Process Contr.*, 22(8):1490–1501.

Haug, J. A. (2005). A tutorial on Bayesian estimation and tracking techniques applicable to nonlinear and non-Gaussian processes. Technical report, MITRE Corporation.

Hendricks Franssen, H.-J., Alcolea, A., Riva, M., Bakr, M., van der Wiel, N., Stauffer, F., and Guadagnini, A. (2009). A comparison of seven methods for the inverse modelling of groundwater flow. Application to the characterisation of well catchments. *Adv. Water Resources*, 32(6):851–872.

Hendricks Franssen, H.-J. and Kinzelbach, W. (2008). Real-time groundwater flow modeling with the ensemble Kalman filter: Joint estimation of states and parameters and the filter inbreeding problem. *Water Resour. Res.*, 44:W09408.

Henri, T. and Yvon, J.-P. (2005). Convergence estimates of POD-Galerkin methods for parabolic problems. In Cagnol, J. and Zolsio, J.-P., editors, *System Modeling and Optimization*, volume 166 of *IFIP International Federation for Information Processing*, pages 295–306. Springer, Boston.

Hinze, M. and Kunkei, M. (2012). Residual based sampling in POD model order reduction of drift-diffusion equations in parameterized electrical networks. *ZAMM-Z. Angew. Math. Me.*, 92(2):91–104.

Hoeben, R. and Troch, P. A. (2000). Assimilation of active microwave observation data for soil moisture profile estimation. *Water Resour. Res.*, 36(10):2805–2819.

Holmes, P., Berkooz, G., and Lumley, J. L. (1996). *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press, Cambridge.

Jazwinski, A. H. (1970). *Stochastic Processes and Filtering Theory*. Academic Press, New York.

Judd, K. and Stemler, T. (2010). Forecasting: it is not about statistics, it is about dynamics. *Philos. T. Roy. Soc. A*, 368(1910):263–271.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *J. Basic Eng.-T. ASME*, 82(Series D):35–45.

Kollet, S. J. and Maxwell, R. M. (2006). Integrated surface-groundwater flow modeling: A free-surface overland flow boundary condition in a parallel groundwater flow model. *Adv. Water Resources*, 29(7):945–958.

Kunisch, K. and Volkwein, S. (2001). Galerkin proper orthogonal decomposition methods for parabolic problems. *Numer. Math.*, 90(1):117–148.

Kunisch, K. and Volkwein, S. (2002). Galerkin proper orthogonal decomposition methods for a general equation in fluid dynamics. *SIAM J. Num. Anal.*, 40:492–515.

Kushner, H. (1967). Approximation to optimal nonlinear filters. *IEEE T. Automat. Contr.*, AC-12:546–556.

Leopold, L. B. and Maddock Jr., T. (1953). The hydraulic geometry of stream channels and some physiographic implications. *Geol. Surv. Prof. Paper*, 252.

Li, H., Luo, Z., and Chen, J. (2011). Numerical simulation based on POD for two-dimensional solute transport problems. *Appl. Math. Model.*, 35(5):2489 – 2498.

Li, H. and Zhang, D. (2007). Probabilistic collocation method for flow in porous media: Comparisons with other stochastic methods. *Water Resour. Res.*, 43:W09409.

Lieberman, C., Willcox, K., and Ghattas, O. (2010). Parameter and state model reduction for large-scale statistical inverse problems. *SIAM J. Sci. Comput.*, 32(5):2523–2542.

Lilliefors, H. W. (1967). On the Kolmogorov–Smirnov test for normality with mean and variance unknown. *J. Am. Stat. Assoc.*, 62(318):399–402.

Liu, G., Chen, Y., and Zhang, D. (2008). Investigation of flow and transport processes at the MADE site using ensemble Kalman filter. *Adv. Water Resources*, 31(7):975–986.

Liu, J. S. and Chen, R. (1998). Sequential Monte Carlo methods for dynamical systems. *J. Am. Stat. Assoc.*, 93:1032–1044.

Luo, Z., Du, J., Xie, Z., and Guo, Y. (2011a). A reduced stabilized mixed finite element formulation based on proper orthogonal decomposition for the non-stationary NavierStokes equations. *Int. J. Numer. Methods Engrg.*, 86:31–46.

Luo, Z., Xie, Z., Shang, Y., and Chen, J. (2011b). A reduced finite volume element formulation and i numerical simulations based on POD for parabolic problems. *J. Comput. Appl. Math.*, 235(8):2098 – 2111.

Mann, H. B. and Whitney, D. R. (1947). On a test of whether one of two random variables is stochastically larger than the other. *Ann. Math. Stat.*, 18(1):50–60.

Margulis, S. A., McLaughlin, D., Entekhabi, D., and Dunne, S. (2002). Land data assimilation and estimation of soil moisture using measurements from the Southern Great Plains 1997 Field Experiment. *Water Resour. Res.*, 38(12):1299.

McLaughlin, D. (2002). An integrated approach to hydrologic data assimilation: interpolation, smoothing, and filtering. *Adv. Water Resources*, 25(8-12):1275–1286.

McPhee, J. and Yeh, W. W. G. (2008). Groundwater management using model reduction via empirical orthogonal functions. *J. Water Res. Pl.-ASCE*, 134(2):161–170.

Montzka, C., Moradkhani, H., Weihermüller, L., Hendricks Franssen, H.-J., Canty, M., and Vereecken, H. (2011). Hydraulic parameter estimation by remotely-sensed top soil moisture observations with the particle filter. *J. Hydrol.*, 399(3–4):410–421.

Moradkhani, H., Hsu, K.-L., Gupta, H., and Sorooshian, S. (2005). Uncertainty assessment of hydrologic model states and parameters: Sequential data assimilation using the particle filter. *Water Resour. Res.*, 41(5):W05012.

Morita, M. and Yen, B. C. (2002). Modeling of conjunctive two-dimensional surface-three-dimensional subsurface flows. *J. Hydraul. Eng.-ASCE*, 128(2):184–200.

Neuman, S. P., Blattstein, A., Riva, M., Tartakovsky, D. M., Guadagnini, A., and Ptak, T. (2007). Is transmissivity a meaningful property of natural formations? Conceptual issues and model development. *Water Resour. Res.*, 45:W03425.

Noh, S. J., Tachikawa, Y., Shiiba, M., and Kim, S. (2011). Applying sequential Monte Carlo methods into a distributed hydrologic model: lagged particle filtering approach with regularization. *Hydrol. Earth Syst. Sc.*, 15(10):3237–3251.

Nowak, W. and Cirpka, O. A. (2004). A modified Levenberg-Marquardt algorithm for quasi-linear geostatistical inversing. *Adv. Water Resources*, 27(7):737–750.

Oliver, D. S. and Chen, Y. (2011). Recent progress on reservoir history matching: a review. *Comput. Geophys.*, 15(1):185–221.

Orlandini, S., Moretti, G., Franchini, M., Aldighieri, B., and Testa, B. (2003). Path-based methods for the determination of nondispersive drainage directions in grid-based digital elevation models. *Water Resour. Res.*, 39(6):1144.

Orlandini, S. and Rosso, R. (1996). Diffusion wave modeling of distributed catchment dynamics. *J. Hydraul. Eng.-ASCE*, 1(3):103–113.

Panday, S. and Huyakorn, P. S. (2004). A fully coupled physically-based spatially-distributed model for evaluating surface/subsurface flow. *Adv. Water Resources*, 27(4):361–382.

Paniconi, C., Marrocu, M., Putti, M., and Verbunt, M. (2003). Newtonian nudging for a Richards equation-based distributed hydrological model. *Adv. Water Resources*, 26(2):161–178.

Paniconi, C. and Putti, M. (1994). A comparison of Picard and Newton iteration in the numerical-solution of multidimensional variably saturated flow problems. *Water Resour. Res.*, 30(12):3357–3374.

Pasetto, D., Camporese, M., and Putti, M. (2012a). Ensemble Kalman filter versus particle filter for a physically-based coupled surface-subsurface model. *Adv. Water Resources*, 47(1):1–13.

Pasetto, D., Guadagnini, A., and Putti, M. (2011). POD-based Monte Carlo approach for the solution of regional scale groundwater flow driven by randomly distributed recharge. *Adv. Water Resources*, 34(11):1450–1463.

Pasetto, D., Putti, M., and Yeh, W. W.-G. (2012b). A reduced order model for groundwater flow equation with random hydraulic conductivity: application to Monte Carlo methods. *Submitted to Water Resour. Res.*

Pauwels, V. R. N. and De Lannoy, G. J. M. (2009). Ensemble-based assimilation of discharge into rainfall-runoff models: A comparison of approaches to mapping observational information to state space. *Water Resour. Res.*, 45:W08428.

Poles, S. and Lovison, A. (2009). A polynomial chaos approach to robust multiobjective optimization. In Deb, K., et al., editor, *Hybrid and Robust Approaches to Multiobjective Optimization*, number 09041 in Dagstuhl Seminar Proceedings. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, Dagstuhl, Germany.

Putti, M. and Paniconi, C. (2004). Time step and stability control for a coupled model of surface and subsurface flow. In *Proceeding of the XV International Conference on Computational Methods in Water Resources (CMWR XV)*, volume 55, part 2, pages 1361–1402. Elsevier.

Qu, Y. Z. and Duffy, C. J. (2007). A semidiscrete finite volume formulation for multiprocess watershed simulation. *Water Resour. Res.*, 43(8):W08419.

Quarteroni, A. and Rozza, G. (2007). Numerical solution of parametrized Navier-Stokes equations by reduced basis methods. *Num. Meth. PDE*, 23(4):923–948.

Rigon, R., Bertoldi, G., and T.M., O. (2006). GEOtop: a distributed hydrological model with coupled water and energy budgets. *J. Hydrometeorol.*, 7(3):371–388.

Rings, J., Huisman, J. A., and Vereecken, H. (2009). Coupled hydrogeophysical parameter estimation using a sequential Bayesian approach. *Hydrol. Earth Syst. Sc.*, 6:6387–6424.

Riva, M., Guadagnini, A., and De Simoni, M. (2006). Assessment of uncertainty associated with the estimation of well catchments by moment equations. *Adv. Water Resources*, 29:676–691.

Rojas, R., Feyel, L., and Dassargues, A. (2008). Conceptual model uncertainty in groundwater modeling: Combining generalized likelihood uncertainty estimation and Bayesian model averaging. *Water Resour. Res.*, 44:W12418.

Rovas, D. V., Machiels, L., and Maday, Y. (2006). Reduced-basis output bound methods for parabolic problems. *IMA J. Num. Anal.*, 26(3):423–445.

Salamon, P. and Feyen, L. (2009). Assessing parameter, precipitation, and predictive uncertainty in a distributed hydrological model using sequential data assimilation with the particle filter. *J. Hydrol.*, 376(3–4):428–442.

Schöniger, A., Nowak, W., and H.-J., H. F. (2012). Parameter estimation by ensemble Kalman filters with transformed data: Approach and application to hydraulic tomography. *Water Resour. Res.*, 48:W04502.

Shen, C. and Phanikumar, M. S. (2010). A process-based, distributed hydrologic model based on a large-scale method for surface–subsurface coupling. *Adv. Water Resources*, 33:1524–1541.

Shu, Q., Kemblowski, M. W., and McKee, M. (2005). An application of ensemble Kalman filter in integral-balance subsurface modeling. *Stoch. Env. Res. Risk. A.*, 19(5):361–374.

Siade, A. J., Putti, M., and Yeh, W. W.-G. (2010). Snapshot selection for groundwater model reduction using proper orthogonal decomposition. *Water Resour. Res.*, 46:W08539.

Siade, A. J., Putti, M., and Yeh, W. W.-G. (2012). Reduced order parameter estimation using quasilinearization and quadratic programming. *Water Resour. Res.*, 48:W06502.

Silverman, B. W. (1986). *Density estimation: for statistics and data analysis.* Monographs on Statisctics and Applied Probability, London: Chapman and Hall.

Simon, D. (2006). *Optimal State Estimation: Kalman, H-infinity, and Nonlinear Approaches.* Wiley & Sons Inc.

Sirovich, L. (1987). Turbulence and the dynamics of coherent structures; Part I: Coherent structures. *Q. Appl. Math.*, 45(3):561–571.

Smith, A. F. M. and Gelfand, A. E. (1992). Bayesian statistics without tears: a sampling–resampling perspective. *Am. Stat.*, 46(2):84–88.

Snyder, C., Bengtsson, T., Bickel, P., and Anderson, J. (2008). Obstacles to high-dimensional particle filtering. *Mon. Weather Rev.*, 136:4629–4640.

Sulis, M., Meyerhoff, S. B., Paniconi, C., Maxwell, R. M., Putti, M., and Kollet, S. J. (2010). A comparison of two physics-based numerical models for simulating surface watergroundwater interactions. *Adv. Water Resources*, 33(4):456–467.

Sun, A. Y., Morris, A., and Mohanty, S. (2009). Comparison of deterministic ensemble Kalman filters for assimilating hydrogeological data. *Adv. Water Resources*, 32(2):280 – 292.

Van Delft, G., El Serafy, G. Y., and Heemink, A. W. (2009). The ensemble particle filter (EnPF) in rainfall-runoff models. *Stoch. Env. Res. Risk. A.*, 23(8):1203–1211.

Van Geer, F. C., Te Stroet, C., and Yangxiao, Z. (1991). Using Kalman filtering to improve and quantify the uncertainty of numerical groundwater simulations: 1. The role of system noise and its calibration. *Water Resour. Res.*, 27:1987–1994.

Van Genuchten, M. T. (1980). A closed-form equation for predicting the hydraulic conductivity of unsaturated soils. *Soil Sci. Soc. Amer. J.*, 44:892–898.

VanderKwaak, J. E. and Sudicky, E. A. (2000). Application of a physically-based numerical model of surface and subsurface water flow and solute transport. In *Calibration and reliability in groundwater modeling: Coping with uncertainty*, ModelCARE 1999 Conf., Zurich, Switzerland, pages 515–523. IAHS Press, Wallingford, UK.

Vermeulen, P. T. M., Heemink, A. W., and Stroet, C. B. M. T. (2004). Reduced models for linear groundwater flow models using empirical orthogonal functions. *Adv. Water Resources*, 27(1):57 – 69.

Vermeulen, P. T. M., Stroet, C. B. M. T., and Heemink, A. W. (2006). Model inversion of transient nonlinear groundwater flow models using model reduction. *Water Resour. Res.*, 42(9):W09417.

Walker, J. P., Willgoose, G. R., and Kalma, J. D. (2001). One-dimensional soil moisture profile retrieval by assimilation of near-surface observations: a comparison of retrieval algorithms. *Adv. Water Resources*, 24(6):631–650.

Wang, J. B. and Zabaras, N. (2005). Using Bayesian statistics in the estimation of heat source in radiation. *Int. J. Heat. Mass. Tran.*, 48(1):15–29.

Weerts, A. H. and El Serafy, G. Y. H. (2006). Particle filtering and ensemble Kalman filtering for state updating with hydrological conceptual rainfall-runoff models. *Water Resour. Res.*, 42:W09403.

Weill, S., Mouche, E., and Patin, J. (2009). A generalized Richards equation for surface/subsurface flow modelling. *J. Hydrol.*, 366(1-4):9–20.

West, M. and Harrison, J. (1997). *Bayesian forecasting and dynamic models*. Springer Verlag Series in Statistic, 2nd edition.

Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics Bul.*, 1(6):80–83.

Winter, C. L., Tartakovsky, D. M., and Guadagnini, A. (2003). Moment differential equations for flow in highly heterogeneous porous media. *Surv. Geophys.*, 24(1):81–106.

Yeh, W. W.-G. (1986). Review of parameter identification procedures in groundwater hydrology: The inverse problem. *Water Resour. Res.*, 22(2):95–108.

Zhang, D. (2002). *Stochastic Methods for Flow in Porous Media: Copying With Uncertainties*. Academic Press, San Diego, CA.

Zhang, D. and Lu, Z. (2004). An efficient, high-order perturbation approach for flow in random porous media via Karhunen-Loève and polynomial expansions. *J. Comp. Phys.*, 194(2):773 – 794.

Zhang, D. X., Shi, L. S., Chang, H. B., and Yang, J. Z. (2010). A comparative study of numerical approaches to risk assessment of contaminant transport. *Stoch. Env. Res. Risk. A.*, 24(7):971–984.

Zhou, H., Jaime Gómez-Hernández, J., Hendricks Franssen, H.-J., and Li, L. (2011). An approach to handling non-Gaussianity of parameters and state variables in ensemble Kalman filtering. *Adv. Water Resources*, 34(7):844–864.

Zhou, Y., McLaughlin, D., and Entekhabi, D. (2006). Assessing the performance of the ensemble Kalman filter for land surface data assimilation. *Mon. Weather Rev.*, 134(8):2128–2142.

Acknowledgments

First and foremost I would like to express my sincere gratitude to my supervisor Prof. Mario Putti for his dedication and enthusiasm in guiding my research. He has been a fantastic source of inspiration and his advice has been invaluable. I am also grateful to my co-supervisors Prof. Alberto Guadagnini and Prof. William W-G. Yeh: their suggestions and comments were very helpful throughout my PhD.

My special thanks goes to Prof. Matteo Camporese, for his help with the complicated code of CATHY, and Prof. Wolfgang Nowak for his detailed review of this thesis.

I thank the entire group of the DMMMSA for the pleasant and stimulating work environment. In particular, I thank Nicola and Gabriele whose friendship made my work experience very enjoyable. Moreover, I am very thankful to Prof. William W-G. Yeh and my fellow PhD students Tim, Scott, Sami, Adam, Manuela, Bruno, Martina and Gonzalo for their hospitality during my stay at the University of California of Los Angeles.

Last but not the least, I am very grateful to my friends, my cousins and, especially, to my family: my parents Paolo and Gina, my sisters Monica and Laura, and my brother Daniele. Thank you for your unconditional support and for always encouraging me to do my best.