Combining Geostatistical Analysis and Flow-and-Transport Models to Improve Groundwater Contaminant Plume Estimation

by

Shahar Shlomi

**A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Environmental Engineering)
in The University of Michigan
2009**

Doctoral Committee:

> **Assistant Professor Anna M. Michalak, Chair
> Professor Peter Adriaens
> Professor Daniel G. Brown
> Professor Roderick J. Little
> Professor Tissa H. Illangasekare, Colorado School of Mines**

To the memory of my mother
עבודה זו מוקדשת לזכרה של אמא

Acknowledgments

I am deeply indebted to many people who have helped me over the last five years at the University of Michigan, in various ways.

First and foremost I would like to thank my advisor Anna Michalak. The amount of time, energy, and resources Anna devoted to me - and to all of her students - is truly remarkable. I have learned a great deal from Anna's insight, comments, and suggestions. She has inspired many of the ideas in this work and helped me overcome many of the obstacles along the way.

I would like to thank my other Committee members, Professors Peter Adriaens, Daniel Brown, Tissa Illangasekare, and Roderick Little. Professor Illangasekare has also been very generous in hosting me at his laboratory and providing me with everything I needed to perform the "wet" experiments described in this work. In addition, Dr Toshi Sakaki and many others at the Colorado School of Mines have contributed to the success of those experiments.

All members of the PUORG Research Group, past and present, have been very helpful all along the way. They patiently heard my presentations over and over again and always

Table of Contents

List of Figures

Abstract

Combining Geostatistical Analysis and Flow-and-Transport Models to Improve Groundwater Contaminant Plume Estimation

by

Shahar Shlomi

Chair: Anna M. Michalak

Groundwater is an important resource, which is often contaminated. In order to ensure a sustainable supply, groundwater has to be monitored, contaminant plumes must be estimated accurately, and remediation operations must be carried out effectively. However, groundwater monitoring networks often do not have enough monitoring wells, and those wells are not always optimally located for the purpose of plume estimation, using existing methods. Moreover, budgetary constraints limit the number of available samples.

Existing methods for plume estimation rely either on the spatial correlation of plume concentrations, or on the underlying physics of groundwater flow and contaminant transport. Often, practitioners who rely on one of these approaches neglect available information which can be used in methods belonging to the other approach. For example, use of kriging, a geostatistical method relying on spatial correlation, often precludes the

use of transport information, which may be readily available. Conversely, flow-and-transport models do not explicitly consider spatial correlation of contaminant concentrations.

In this work, these two approaches are combined, in order to optimally use all available information, to improve the quality of plume estimation. Specifically, two geostatistical methods – Inverse/Forward Modeling and Transport-Enhanced Kriging – are developed that combine transport models with spatial or temporal correlation. These methods are versatile, can apply to a variety of situations, and can work with many kinds of available input data and transport models. A method is also developed to estimate flow and transport parameters simultaneously with the plume concentration, for cases in which this information is unknown or uncertain. Finally, as monitoring network configuration has a dramatic effect on estimation results (but is specific to the plume estimation method used), a method for choosing optimal monitoring sites is presented.

All of the methods were tested in a variety of numerical experiments with synthetic homogeneous and heterogeneous data. In addition, several laboratory experiments were performed in a large sand tank, to assess the performance of the methods. Overall, the new methods yield results that are superior to those obtained by common existing methods such as kriging, with a better reproduction of the true plume shapes and lower uncertainty.

Chapter 1
Introduction

ויפקח אלוהים את-עיניה, ותרא באר מים,

ותלך ותמלא את-החמת, מים, ותשק, את-הנער.

בראשית כא19

And God opened her eyes, and she saw a well of water; and she went,

and filled the bottle with water, and gave the lad drink.

Genesis 21: 19

1. Groundwater Past and Present

Making up more than 98% of fresh water resources on Earth, groundwater has been tapped as an important water resource for human societies for thousands of years [Fetter, 2001]. The earliest wells are known from the Neolithic. In the submerged Pre-Pottery Neolithic B settlement of Atlit Yam in Israel, dated to 8100-7500 BC, a well has been found, which so far is the oldest known [Galili and Nir, 1993].

Today, groundwater is the most widespread and highly used resource, supplied to about 1.5 billion people worldwide [Hetzel et al., 2008]. It is of inestimable value to the inhabitants of dry regions, being the only reliable water resource. The yearly consumption of groundwater world-wide is assumed to be 1000 cubic kilometers, and the global groundwater recharge is estimated at over 12,000 cubic kilometers per year [Döll

and Fielder, 2008]. In many countries, the drinking water is mainly tapped from groundwater, and in some, groundwater is the *only* source for public water supply. The importance of groundwater will grow considerably in the future, as it is a safe and high-quality drinking water resource. If it is used reasonably and sustainably, it can provide an important contribution in solving regional water crises on Earth.

2. Groundwater Contamination

Various sources of contamination affect groundwater, limiting the available supply, and necessitating detection and quantification of the contaminant, and remediation of the groundwater. Contamination of water supplies may be responsible for more human sickness than any other anthropogenic activity [Anderman and Martin, 1986]. Contaminants can come from a variety of sources, which can be generally categorized as either anthropogenic or natural.

In the first category, the most common artificial contaminant is trichloroethylene (TCE), which is an effective solvent for a variety of organic materials [Rivett et al., 2006]. Tetrachloroethylene (PCE), used mainly for dry cleaning, is a related solvent which has often become a groundwater contaminant [Schumacher et al., 2004]. Benzene, toluene, ethylbenzene, and total xylenes (BTEX), which come from gasoline refining, and methyl-tert-butyl-ether (MTBE), which is a fuel additive, are common contaminants in urbanized areas [Rail, 1989], often as the result of leaking underground storage tanks. Groundwater can be contaminated by localized releases from waste disposal sites and landfills. Pesticides, fertilizers, and contaminants from other nonpoint source pollutants are also

major sources of groundwater pollution [Commission on Geosciences, Environment and Resources (CGER), 1993].

Common natural contaminants include iron, magnesium and calcium [Rivett et al., 2006]. Other natural constituents of concern are nitrates and Coliform bacteria, both of which are considered acute contaminants [Palmer, 1996].

3. Physics of Groundwater

In order to monitor contaminants, estimate plumes (i.e. map and/or quantify their spatial distribution in the subsurface), or more generally assess the quality of groundwater, it is important to understand the underlying physical processes that govern its flow, as well as the transport of contaminants.

3.1 Groundwater flow

Groundwater hydrology began as a quantitative science when Darcy (1803 - 1858), a French hydraulic engineer, described the results of an experiment designed to study the flow of water through a porous medium [Darcy, 1856]. The experiment resulted in the formulation of Darcy's Law, which states that the rate of fluid flow through a porous medium is proportional to the potential energy gradient within that fluid. The constant of proportionality is the hydraulic conductivity, which is a property of both the porous medium, and the fluid moving through the porous medium.

Combined with a water-balance equation for steady state flow, Darcy's Law is the basis for the governing groundwater flow equation in a confined aquifer:

$$\nabla \cdot (\mathbf{K}\nabla\varphi) + R - P = 0 \qquad (1)$$

where $\mathbf{K}$ is the hydraulic conductivity (which can in general be a symmetric tensor), $\nabla\varphi$ is the head gradient vector, and R and P represent recharge and pumping rates, respectively [Bear and Verruijt, 1987].

3.2 Contaminant transport

Several processes govern solute transport in groundwater: Advection is chemical movement via groundwater flow due to the groundwater hydraulic (i.e. head) gradient. Dispersion is the longitudinal (forward and backward) and transverse (normal) spreading of the contaminant. Retardation causes the mean chemical velocity to be slower than the groundwater velocity, through sorption processes in which the contaminant is attracted to the solid aquifer matrix [Runnells, 1993]. Contaminants may also react with each other or with the subsurface matrix in various biological and chemical processes, which ultimately affect transport and consequently the concentration of solutes.

The advection – dispersion equation describes the spatial distribution of a contaminant as a function of time, and can be written as follows:

$$\frac{\partial \mathbf{C}}{\partial \mathbf{t}} = \mathbf{D}\nabla^2\mathbf{C} - \mathbf{v}\nabla\mathbf{C} + \frac{Qc}{n} \qquad (2)$$

4

where **C** is the contaminant concentration distribution, **D** is the dispersion tensor, and **v** is the pore velocity (which the water traveling through the pores is experiencing). *n* is the porosity and *Q* represents flow rate from a source with concentration *c*. This too is a simplified version of the governing equation, which does not take into account any processes other than advection and dispersion. Still, even solving only these equations in this form (i.e. assuming no other processes) poses a challenge, because their coefficients, as well as boundary conditions, initial conditions, source/sink terms, etc. are not generally known, and have to be estimated using sparse data.

Throughout this work, both the groundwater flow equation and the contaminant transport equation are solved using two well-established finite-difference numerical codes, MODFLOW-2000 [Hill et al., 2000] and MT3DMS5 [Zheng and Wang, 1999], respectively.

4. Groundwater Management and Monitoring

The high susceptibility of groundwater to contamination can limit the value of the resource to society as a whole. In the United States, hundreds of thousands of sites have contaminated groundwater. The estimated cost of remediating *all* of these sites would range from $480 billion to $1 trillion [Reed et al., 2000]. However, The American Society of Civil Engineers (ASCE) highlights that projected federal expenditures on long term groundwater monitoring for the present decade will only be somewhat more than $5 billion [ASCE, 2003]. Groundwater pollutants pose a threat to the environment and to consumers and therefore must be detected and monitored. Groundwater monitoring is

often defined as the testing of groundwater over an extended time period in order to document groundwater conditions, including the collection of chemical data, such as contaminant concentrations [ASCE, 2003]. In the context of this dissertation, the term groundwater monitoring also includes the estimation of contaminant concentrations throughout the aquifer, not only at data collection sites.

4.1 Legal and regulatory aspects

As the population (in the U.S. and globally) expands, limited water resources are increasingly shared by a larger number of competing consumers, and stakeholders express a growing concern about the quality of source waters and the health risks associated with groundwater contamination. As a result, long-term monitoring of the quantity, quality, and susceptibility of groundwater has been and continues to be an issue of paramount importance. Broad initiatives mandated by The Federal Safe Drinking Water Act [1974, 1986, 1996] require long-term groundwater monitoring records at both regional and local scales. An example is the 1996 amendments to Section 1453 of the Safe Drinking Water Act, which require states to establish and implement a Source Water Assessment Program (SWAP). Source water assessment is intended to provide a strong basis for developing, implementing, and improving a state's source water protection plan. This program requires individual states to delineate protection areas for drinking water intakes, identify and inventory significant contaminants in the protection areas, and determine the susceptibility of public water supply systems to the contaminants released within the protection areas [Bice et al., 2000]. Under CERCLA (Comprehensive

Environmental Response, Compensation, and Liability Act, 1980) and RCRA (Resource Conservation and Recovery Act, 1976), the federal Environmental Protection Agency (EPA) and local state agencies are able to implement maximum contaminant levels in groundwater through enforcement measures, which can include fines of up to $25,000 per day of regulatory noncompliance [Hardisty and Ozdemiroglu, 2005].

4.2 Groundwater contaminant plume estimation

Given that groundwater is the largest store of fresh water within many states, increased investment in long-term groundwater monitoring will be required to better understand the condition of aquifers, and to sustain their long-term capability to provide safe drinking water. The success of future groundwater resource management initiatives will require new tools for optimally balancing monitoring costs and uncertainty. This is particularly important when monitoring groundwater for pollutants such as volatile organic chemicals, pesticides, and other expensive analytes over long time periods.

One of the most basic and important tools in groundwater management is plume interpolation or estimation, which aims to provide the best possible quantitative account of the spatial distribution of a contaminant in an aquifer, along with a measure of the corresponding uncertainty. This tool plays a vital role in a variety of applications (Table 1), the overall aims of which are often to (i) avoid (or at least minimize) pumping of contaminated water into a water distribution system, (ii) contain the contamination plume within a specified area, and/or (iii) treat the contaminated groundwater (remediation) to reduce contaminant levels to within an acceptable threshold [Freshley et al., 2002].

7

Table 1. Various Uses of Plume Estimation

Application	Requirements: why plume estimation	Typical Available Data	Example Reference
Monitoring well selection	Which wells should be sampled? Defined objective: minimize global/local variance/error?	Concentrations at previous times Transport model Source location	Aziz et al. [2000]
Global mass estimation	Estimates of concentrations at nodes	Concentrations at measured locations	Reed et al. [2000], Nyer et al. [1998]
Water distribution systems	Pumping flowrates at wells	Source location (e.g. underground storage tank); Flow model	Weaver et al. [1999]
Remediation	Monitor remediation process	Historical/initial plume, transport model	Boulding and Ginn [2004]
Inverse modeling and parameter estimation	Estimates of parameter(s) at different locations, times	Parameterized model, discrete measurements	Sun [1994]
Network design	Optimal locations for future measurements	Hydrogeological structure, flow regime, source locations	McKinney and Loucks [1992], Meyer et al. [1995]

These applications are often accompanied by monitoring of the groundwater quality over time, which is defined in a variety of ways in the literature [e.g. ASCE 2003; NFESC, 2000], but always involves sampling wells, sometimes followed by estimating concentrations of contaminants elsewhere, and repeating this process after some time

interval, often regularly. Monitoring itself has many specific objectives, such as demonstrating the effectiveness of corrective action, examining the trend of contaminant concentrations [Lindberg, 2006], meeting regulatory requirements, and research [Herzog et al., 1991]. In most cases, some form of spatial estimation of groundwater contamination is an integral part of the monitoring program. Plume estimation is also used for other applications, such as global mass estimation, which often aims to find the total contaminant mass emitted to, or present in, an aquifer, for litigation purposes or as an indicator of natural attenuation [Nyer et al., 1998], respectively.

These and many other groundwater quality management activities require knowledge of the contaminant's spatial distribution. However, as in many other environmental applications, data is typically extremely limited. Even if wells are sampled frequently, they are usually not located close to each other, so the spatial dimension is sparse. Moreover, the spatial configuration of the monitoring wells is rarely ideal for the purposes of sampling, estimation and monitoring. In addition, contaminant concentration fields are highly heterogeneous, anisotropic, and non-stationary phenomena [Reed et al., 2004]. As a result of these factors, plume distributions estimated through interpolation often do not represent the true plume distribution well, and the full spatial distribution of the plume is almost never known.

Interestingly, additional information on the aquifer and its flow regime does often exist – but cannot be integrated in standard interpolation methods (such as geostatistical kriging), which may accept only concentrations as input data. The additional information, often

9

available in the form of a (perhaps incomplete) flow-and-transport model, can potentially constrain the concentration estimates, thus provide a best estimate which is closer to the true concentration, with a lower uncertainty, due to the additional information. The ability to concurrently use information about both (i) the physical behavior as expressed in flow-and-transport models and (ii) the spatial correlation as described in existing geostatistical models, would have a dramatic impact of our ability to represent the shape and extent of groundwater contaminant plumes. The premise of this dissertation is to combine these two types of approaches and to develop methods that would have a significant effect on the management of contaminated groundwater resources, by improving our ability to avoid pumping contaminated water, improving the design of groundwater remediation alternatives, and providing a framework for monitoring remediation progress. The high cost associated with groundwater monitoring and the high risks posed by these contaminants contribute to the importance of developing robust estimation techniques which take into account diverse types of data for plume interpolation.

4.3 National research needs

The two preceding subsections provide the motivation for developing methods that could make better use of existing data, find ways to assimilate different sets of data, and optimize data collections schemes for the purpose of plume estimation. It should be noted that national science organizations have identified similar needs. For example, the National Research Council [CGER, 2000] recommended four research emphases for

DOE's subsurface cleanup mission, which included *Location of subsurface contaminants*, *Conceptual Modeling*, *Containment*, and *Monitoring*. More specifically within the first emphasis, the Committee supported development of improved capabilities for characterizing heterogeneity, measuring contaminants, integrating data collected at different spatial and temporal scales, and integrating such data into conceptual models. Within the second emphasis, the Committee recommended focusing on new approaches for incorporating heterogeneity into conceptual model formulations and integrating process knowledge into model formulations.

One of the United States Geological Survey's major water-quality themes (for the current cycle II) of NAWQA, its National Water-Quality Assessment Program, is *Extrapolation and forecasting* [NRC, 2002]. Within this theme, answers to the following question are sought:

> "How can we best extrapolate (spatial dimension) or forecast (temporal dimension) water-quality conditions for unmeasured geographic areas and future conditions, based on knowledge of contaminant sources, natural characteristics of the land and hydrologic system, and our understanding of governing processes?"

The work presented hereinafter in this dissertation provides a significant contribution to the solution of these problems.

5. Objectives of This Work: Groundwater Quality Monitoring

Broadly, the goal of this work was to develop a set of tools which would **enhance**

groundwater quality monitoring, defined here more specifically as the aggregate of several distinct phases:

1. Designing a network of monitoring wells for an aquifer,

2. Selecting a subset of these wells to sample at a particular point in time,

3. Using the measured concentration, along with other data, to estimate the distribution of contaminants in the aquifer.

Despite this chronological order, most of the work was dedicated to the third element: estimating the contaminant plume by combining transport information in geostatistical analyses. Only Chapter 5 deals with network design and well selection. More specifically, this dissertation addresses the four objectives defined below.

5.1 Objective 1: Integrate transport information into geostatistical analyses

In Chapter 3, a set of geostatistical tools for estimating groundwater contaminant plumes is developed using:

1. Measured concentrations from samples at monitoring wells, and

2. A flow-and-transport model of the aquifer.

The proposed approach proceeds in two steps. First, an initial condition (such as the state of the contaminant plume at an earlier time) or boundary condition (such as the concentration of the contaminant as a function of time at a point source) of the transport model is estimated. This is done using geostatistical inverse modeling, which is a method

used here to infer the initial or boundary condition using concentration measurements in conjunction with a flow-and-transport model of the affected aquifer. Next, the estimated initial or boundary condition is used as input to the transport model (now as a forward-model), in order to estimate the full extent of the plume, i.e. the concentrations throughout the domain. This process is named Inverse/Forward Modeling (IFM).

In order to also take advantage of the spatial autocorrelation of the contaminant plume, one can also integrate the second (forward-modeling) step with a kriging-like constraint. We call this process Transport-Enhanced Kriging (TrEK). Thus, both measured concentrations and a flow and transport model are be used together to infer the plume distribution.

5.2 Objective 2: Introduce transport model uncertainty

The methods developed within Objective 1, IFM and TrEK (as well as geostatistical inverse modeling upon which they are based), do not explicitly consider transport model uncertainty (i.e. they do not separate these errors from measurement errors). However, groundwater modeling relies on many parameters, most of which are usually not known throughout the aquifer. These parameters often have a considerable impact on modeling results. Specifically, uncertainty in the model parameters causes uncertainty in model results, and should therefore be accounted for and treated appropriately, to be reflected accordingly in the estimation results.

The second objective is thus to generalize IFM and TrEK by introducing uncertainty into the transport model. One technique for addressing transport model uncertainty is demonstrated analytically in Chapter 3 and numerically in Chapter 6. It involves generalizing the model-data mismatch error, which is the statistical distribution which the model error is expected to follow. This affects only one parameter of the equation system, without changing the equations (from Chapter 3) themselves.

Chapter 4 presents a method for dealing with an unknown flow/transport model-input field, with an application to an unknown hydraulic conductivity field. It involves simultaneously and iteratively estimating this field along with the initial or boundary condition.

5.3 Objective 3: Optimal sampling well selection for IFM/TrEK

The third objective is to develop an optimal strategy for IFM/TrEK monitoring network design, incorporating measurements of concentrations, the spatial covariance structure of their distribution, and flow-and-transport information. Sampling groundwater quality at different sets of monitoring wells leads to substantially different plume estimates, but the optimal choice of wells depends, among other things, on the specific plume estimation method and on the monitoring objective. For that reason it is important to develop a network design method for the new estimation techniques from Objectives 1 and 2.

In Chapter 5, an optimal measurement site selection method for IFM and TrEK is developed. This method is based on variance-reduction methods (applied in the past to

kriging-based estimation), which seek to minimize either the maximum or the average estimation variance throughout the domain.

5.4 Objective 4: Demonstration and validation of methods in sandbox experiment

The last objective is to demonstrate the methods developed in Objective 1 in a controlled laboratory experiment. Various controlled pollution events were generated in a sandbox with known properties (e.g. hydraulic conductivity), to develop data sets of measurements of simulated plumes. Using this data, the performance of the methods developed in Objective 1 was evaluated.

The ultimate goal of the methods developed throughout this dissertation is to improve groundwater quality monitoring in real aquifers. Laboratory tests serve as an intermediate between computer simulations (with synthetic data sets) and field applications. Computer simulations are very convenient and inexpensive to perform, but may not be realistic, as many of the obstacles of field work are not apparent. Field applications, on the other hand, may be expensive and seldom offer an exhaustive known parameter set, which would enable a model of the true contaminant distribution to compare against.

In Chapter 6, a large sand tank is used to simulate an aquifer, to which a tracer whose concentration we can then measure continuously at several locations. This concentration data, along with a flow-and-transport model (built using the sand tank's known properties), is used to test the performance of IFM.

Chapter 2
Literature Review

This Chapter presents background information of current methods that are used for interpolation of groundwater contaminant plumes, estimation of hydraulic conductivity fields, and design of groundwater quality monitoring networks. More specific information, as it relates to the methods developed in this work, is presented in the following Chapters.

1. Groundwater Contaminant Plume Estimation

Plume estimation can be based either on mathematical or statistical manipulation of measured contaminant concentrations, or on the physics of the transport process that created the plume. When concentration measurements are the only available data, interpolation methods such as geostatistical kriging or inverse-distance weighted interpolation are typically used. When additional hydrological parameters are available, flow-and-transport models are often employed.

In general, kriging can use measurements at sampled locations, incorporate trends, and take advantage of measurements of certain other related variables (e.g. cokriging). Spatial analysis can be performed to identify spatial trends and variogram structures to be used in

the estimation process. Thus, the variogram or covariance function contains additional information, in the form of correlation or variance as a function of distance between points. For example, Cooper and Istok [1988] presented the use of geostatistics to map contaminant concentrations and estimation errors in a groundwater plume, from a set of measured contaminant concentrations. Reed et al. [2004] compared two inverse distance weighting techniques and four kriging techniques for plume interpolation. They ultimately concluded that quantile kriging [Journel and Deutch, 1997] was the most robust method for their specific application. They recommended using deterministic methods only as screening tools. Jones et al. [2003] compared three different three-dimensional plume interpolation techniques (kriging, natural neighbor, and inverse-distance weighted interpolation) and demonstrated that kriging usually results in the lowest error. Geostatistics can also be used to generate stochastic realizations of regionalized variables [Delhomme, 1978]. Boeckenhauer et al. [2000] used non-parametric regression and kriging to produce regional estimates of groundwater contamination by modeling their data as a realization of a log-normal stochastic process. However, many types of supplementary physical data, such as flow and transport information, cannot be directly used in kriging.

Variants and extensions of kriging can be used to introduce additional information into geostatistical analyses. Diggle et al. [1998] defined model-based geostatistics by modeling observations as a generalized linear model. Figueira et al. [2001] measured the concentration of chloride and sodium in plants, and used three additional factors (the distance from the coast, the intensity of rain observed before the sampling date and the

dry period before sampling) in the definition of a space-time trend. Kitanidis and Shen [1996] added one parameter to linear geostatistics to account for the skewness of concentration distributions. Liu [2003] incorporated gradient or sensitivity information into existing kriging techniques for various Multidisciplinary Design Optimization procedures; the approach treated gradients at the sample points as secondary functions. All of these methods analyzed various aspects of the measured distribution, but they did not consider the source of the contamination or the process of transport.

Several studies have also used a combination of kriging/interpolation and groundwater transport models in hydrology. Some have used stochastic [e.g. Wagner and Gorelick, 1987] and geostatistical [e.g. Vyas et al., 2004] methods to estimate parameters for flow-and-transport models, or to interpolate input data to these models [e.g. Feehly et al., 2000]. In these studies, the parameter estimation and the flow-and-transport model constitute distinct sequential phases: the second phase (which incorporates the flow-and-transport model) does not provide information to the first phase (interpolation), and in some cases the total uncertainty is not accounted for properly, because it only considers one phase. Saito and Goovaerts [2001] presented a variant of kriging-with-a-trend for incorporating source location and wind speed in the interpolation of soil contaminant concentrations. While these factors were identified as likely to control pollution spread, none of the actual physics (e.g. in the form of equations) was incorporated into the estimation process. Rather, only the parameters were allowed to affect the trend.

More generally, several researchers have developed various frameworks for the purpose of integrating data from diverse sources into an estimation methodology. For example, Christakos [2005] presented several approaches to geophysical assimilation research, e.g. the ontologic assimilation (OA) approach, which can be used to solve inverse problems. Certain forms of OA may also be used to assess the uncertainty due to model structure error (conceptual uncertainty). He promotes using epistemic cognition assimilation concepts. More specifically, the spatiotemporal Bayesian Maximum Entropy (BME) mapping approach introduced by Christakos [1990, 2000] provides a conceptual framework for processing hard and soft data in the presence of physical laws. Serre et al. [2003] used BME to solve the inverse problem by first assimilating various physical knowledge bases (hydrologic laws, water table elevation data, uncertain hydraulic resistivity measurements, etc.) and then producing robust estimates of the subsurface variables across space. However, this approach has not been used for the specific purpose of recovering the history of contaminants or estimating contaminant plumes. Porter et al. [2000] and Porter [2002] introduce Data Fusion Modeling (DFM), a spatial state estimation and system identification methodology that uses measured data, physical laws, and statistical models for uncertainty in spatial heterogeneities. DFM integrates data, calibrates models, and quantifies uncertainties by solving the causality problem in a generalized Kalman filter. While the presented application is a groundwater flow model, this methodology might also work for contaminant plume estimation. This method may be more suited for real-time updates of models. The multitude of ideas, models, theories,

and approaches, and the large number of discrete steps, make DFM somewhat complex, and this approach, too, has never been applied to the problem of plume estimation.

Many challenges accompany the use of flow-and-transport models for the purpose of estimating the current distribution of existing groundwater contaminant plumes. This type of use would require knowledge of the boundary conditions, including all contaminant sources and their behavior in space and time. Excluding controlled tracer experiments [e.g. Mackay et al., 1986], this information is generally not available. To sum, many investigators have combined elements of both transport models and geostatistics into their analyses; however, it is unknown of any plume estimation method that specifically takes advantage of both a flow-and-transport model (which embodies the laws of physics) and the geostatistical covariance structure of the plume (which embodies the analysis of observations).

Thus, the advantage that transport models have to offer, such as information about the physics controlling advection, dispersion, retardation, and chemical behavior, cannot be fully utilized by methods currently available for estimating the distribution of contaminant plumes.

2. Estimation of Hydraulic Conductivity Within a Geostatistical Framework

The preceding section focused on contaminant plume estimation using mostly mathematical methods, which for the most part interpolate input concentration data

measured at various nodes. In order to take advantage of physical properties of the subsurface, and the transport processes of contaminants, flow-and-transport models are developed. These models typically have many parameters, which must be estimated, generally before using the model. "There is a vast body of literature on the development and application of estimation approaches for these models' parameters, the majority of which has focused on the physical properties of the subsurface, such as, for example, hydraulic conductivity distributions" [Michalak and Shlomi, 2007], which affect groundwater flow models more than anything else [Ricciardi et al., 2007].

As the subject of "estimation approaches for subsurface applications" is too broad to effectively review here, this section focuses primarily on geostatistical methods for estimating the hydraulic conductivity field of an aquifer, which will be carried out in this dissertation. "The geostatistical techniques are attractive because of the simplicity of the procedure and the possibility of easily adapting the assumptions (stationarity, type of covariance, etc.) to the local conditions" [Ahmed and De Marsily, 1987]. Equally important, the plume estimation methodologies from Objective 1 (in the Introduction) were developed *within a geostatistical framework*. In order to have the ability to generalize these methods and integrate the estimation of hydraulic conductivity into the contaminant concentration estimation framework (Objective 2), it is natural to adapt or develop geostatistical methods for hydraulic conductivity estimation, too.

Use of geostatistics in hydrology can be traced back to Delhomme [1976, 1978]. Neuman and Yakowitz [1979], Neuman et al. [1980], and Neuman [1980] all used kriging to

estimate the transmissivity distribution in aquifers. Many early works [e.g. Dagan, 1979; Gutjahr et al., 1978] only used relatively simple approaches such as simple or ordinary kriging. Other variants of kriging were compared by Ahmed and De Marsily [1987], who demonstrated the improvement in estimating transmissivity obtained by integrating more than one type of data (in their case, pump test data and specific capacity). Their comparison included a method that combines kriging with linear regression [Delhomme, 1974, 1976], cokriging [Matheron, 1971], kriging with a trend [Delhomme, 1979], and kriging with a guess field [Delhomme, 1979]. Ahmed and De Marsily [1987] concluded that cokriging and kriging with a trend are generally the more advantageous methods, when additional data is available. Similar methods use either theoretical developments or empirical derivations to find the covariances and cross-covariances for cokriging applications. Ahmed and De Marsily [1993] used a simple relation between transmissivity and head to calculate coherent head covariance and head-transmissivity cross-covariance, based on a chosen transmissivity covariance. They demonstrated the improved estimation (of transmissivities) obtained when piezometric head data was combined with transmissivity data, relative to the use of transmissivity data alone. Many others are still using cokriging approaches to resolve the conductivity field, including for example El Idrysy and De Smedt [2007], who used electrical resistivity data and the slope of the water table as secondary variables. While these methods use parameters which would also be found in transport models, they do not use integrate the knowledge of the actual physical processes of flow and transport into the estimation framework.

To address this issue, in the 1980s, geostatistical tools for inverse modeling began appearing. Kitanidis and Vomvoris [1983] presented a two-step geostatistical solution to the inverse problem, with one-dimensional simulations for recovering permeability. They first used all available information to select geostatistical structure and estimate its parameters; next they found the minimum variance and unbiased linear estimates of the permeability and its estimation variance. Hoeksema and Kitanidis [1984] improved on this work, by using numerical methods for the calculation of the first two joint moments of hydraulic head and the log-transmissivity. They applied their versatile and computationally-efficient method to two-dimensional steady-state flow fields. Dagan [1985] reduced a similar problem to determining the unconditional head-log-transmissivity covariance and head variogram (for a selected log-transmissivity covariance, which depends on a few unknown parameters). He then solved it using a first-order approximation of the flow equations. Dagan et al. [1996] proposed to represent random trajectories of particles in a Lagrangian framework, by relating concentration measurements to trajectories through measurement points and times. While cokriging concentration measurements with flow/transport parameters leads to correct conditioned means of those parameters, the corresponding conditional variance (based on cokriging) depends on the concentration, and the authors concede that it is erroneous.

The geostatistical approach to inverse modeling of Kitanidis and Vomvoris [1983] was generalized by Kitanidis [1995], specifically for problems that may involve nonlinear relations between data and (unknown) parameters. This author noted that the crux of the parameter estimation problem is how to parameterize the distributed parameter system,

and accordingly, the principal conceptual differences among available methods are in the parameterization. He then explained that the random-field parameterization of the geostatistical approach eliminates the need for zoning or other arbitrary and inflexible assumptions. Kitanidis [1996] further emphasized the significance of structural analysis, and compared the geostatistical approach to the inverse problem with a straightforward maximum *a posteriori* probability estimation method [Carrera and Neuman, 1986]. He showed that they use different principles to separate the estimation of covariance parameters from estimation of the spatial variable, and that the maximum *a posteriori* method produces biased estimates of the parameters, as opposed to the geostatistical approach. The quasi-linear inversing approach was chosen for the estimation of hydraulic conductivity in this dissertation, because of its accuracy, effectiveness, relative simplicity, and flexibility, as described in the next paragraph. Another reason for this choice was the similarity to the linear inverse-modeling, used within the methods developed as part of Objective 1.

The quasi-linear geostatistical theory [Kitanidis, 1995] was generalized by Nowak and Cirpka [2004] to handle uncertain trend coefficients (which were usually considered unknown but deterministic), a generalization that could also hold for other kriging and cokriging applications. Nowak and Cirpka [2006a] presented a Bayesian framework for linear cokriging and extended the quasi-linear geostatistical approach [Kitanidis, 1995] to include the generalized cases of uncertain prior knowledge about structural parameters (i.e. covariance and generalized covariance function parameters). Specifically, they incorporated a modified Levenberg-Marquardt algorithm into the framework, for

improving the identification of unknown autocovariance parameters and increasing its stability. Nowak and Cirpka [2006b] presented a geostatistical method for the joint estimation of log-conductivity and log-dispersivities from measurements of hydraulic heads and temporal moments of local breakthrough curves. These or similar extensions may prove useful in generalizing applications based on the quasi-linear inversing approach, and constitute one of the reasons for selecting this approach.

Several other geostatistical inverse methods were described by Zimmerman et al. [1998], who compared their performance in estimating transmissivities, for the ultimate purpose of solute transport predictions, using four test problems. The Fast Fourier Transform (FFT) method and the Fractal Simulation method both used generated multiple realizations of parameter fields using FFT. The Linearized Semianalytical method, Linearized Cokriging method, and the Maximum-Likelihood method used the maximum likelihood theory to estimate structural aquifer parameters, and may thus be biased [Kitanidis, 1995, 1996]. Finally, the Pilot Point method and Sequential Self-Calibration method combined kriging and optimization methods. Four of the methods performed comparably well; however, they seem overly complex (relative to the quasi-linear approach, for example) for the purpose needed here. Moreover, the authors of this study showed that all of the methods had not adequately assessed the prediction uncertainty and asserted that "there is much room for improvement in the inverse methodology." They presented many conclusions regarding the importance of variogram selection, proper parameterization, grid discretization, and the time and experience devoted by the user of the method. Finally, Zimmerman et al. [1998] recommended allowing simultaneous

calibration of transmissivity and boundary conditions while imposing physically-plausible constraints. The approach presented in this dissertation (Chapter 4) takes on this recommendation.

An important source of prediction uncertainty is model error, caused by oversimplified model structure, inexact model parameters, and numerical error [Sun et al., 2006]. Geostatistical applications often lump these together with measurement error as an aggregate model-data mismatch error, or evaluate them through sensitivity studies [Butera and Tanda, 2003], but for the most part do not treat each and every one of these sources of uncertainty. Hernandez et al. [2002, 2003, 2006] formulated a nonlinear inverse problem on the basis of ensemble moment equations on a computational grid. Their approach provided estimates of unknown hydraulic conductivity along with geostatistical (variogram) parameters, and also of second head and flux moments, which constitute measure of predictive uncertainty.

A more comprehensive strategy (which can potentially consider multiple types of data and account for various uncertainties) involves model selection [see review in Burnham and Anderson, 2002] or developing methods for combining several models. For example, Neuman and Wierenga [2003] described a strategy that embodies hydrogeologic conceptualization, model development and predictive uncertainty analysis. Specifically, they used Bayesian Model Averaging to integrate different variogram models for estimating permeability. These approaches, however, are more suitable for larger-scale problems which are beyond the scope of the types of examples presented here. The plume

concentration uncertainty caused by unknown model parameters in this work (Chapter 4) is considered by combining quasi-linear inversing for model parameters with the methods developed in Chapter 3.

3. Groundwater Monitoring Network Design

An important part of virtually any plume estimation method are the solute concentrations measured at specific locations in the aquifer. The locations of these monitoring wells have a significant effect on the estimated distribution of solutes. Thus, one component of improving groundwater quality monitoring is optimizing the selection process of these measurement sites. Moreover, the data collection phase of groundwater monitoring is very expensive. For example, Johnson et al. [1996] estimated that in 1993, some 10,000 samples were taken from more than 1500 groundwater monitoring wells at the Savannah River site at a yearly cost of $10,000,000 in laboratory fees alone. Despite these high costs, systematic sampling methods often give way to inefficient sampling; Johnson et al. [1996] reported that the sampling schedule for a given monitoring well was selected largely by applying rules-of-thumb relating its location to the margins of the contaminant plume. Monitoring costs can grow exponentially as a result of monitoring-dependent operations, such as various remediation schemes. This provides further motivation to optimize the selection of sampling locations for plume estimation.

Many studies have addressed the technical problems of monitoring network design in groundwater hydrology. The design criteria and constraints are generally dictated by the specificities of the addressed problem as well as the objectives of the monitoring

program. Showalter [1985] discussed how the objectives of a groundwater quality monitoring network are developed and how they are implemented in the design process. For example, Hsueh and Rajagopal [1988] developed a model for groundwater quality sampling design with an emphasis on prevention of population exposure to contaminants. The objectives in two examples that Herrera and Pinder [2005] presented were to design a contaminant sampling program to estimate the contaminant concentrations of a moving plume (during a 2-year period), and to estimate the contaminant concentration at the nodes of a mesh for six different dates. The objective of Reed et al. [2000] and Wu et al. [2005] was to minimize long term monitoring costs. Bogaert and Russo [1999] and Chang et al. [2005] presented methodologies for parameter structure identification. The objective of Nunes et al. [2007] was to optimize a monitoring network for plume detection and delineation. Some network design studies have focused on intermediate objectives such as parameter identification, both for geostatistical applications, e.g. variograms [Bogaert and Russo, 1999] and for groundwater modeling [Chang et al., 2005]. However, it is more common to integrate the solution of this kind of a problem in a larger context (such as that of plume estimation), as parameter estimation is not usually the ultimate objective [Müller, 2001]. Other objectives may include background sampling, release detection, flow containment verification, and plume stability (or shrinkage) verification [ASCE, 2003]. In addition to defined objectives, many monitoring programs are subject to legal or regulatory constraints, and must consider these in designing their monitoring plans.

Loaiciga et al. [1992] reported a comprehensive literature review on groundwater quality monitoring network design, in which two main approaches were identified. The first was dubbed the hydrogeological approach, and relied on qualitative and quantitative hydrogeologic information and practitioner judgment, without resorting to advanced statistical and/or probabilistic techniques [Herrera and Pinder, 2005]. In practice, the hydrologic approach has been used extensively in the past. The EPA is supportive of efficient, effective, and innovative tools, but "the regulated community has been reluctant to adopt those approaches readily due to concern over regulatory acceptance. [S]tate and federal regulators are generally risk-averse when approving new characterization, monitoring, and modeling technologies" [Pinder et al., 2003]. The hydrologic approach is not considered further here and not used as a basis for a well-selection methodology, because by nature it does not consider the spatial correlation often observed in concentration measurements. This correlation constitutes meaningful information and forms an important part of the estimation methods at the basis of this dissertation, IFM and TrEK.

The second approach defined by Loaiciga et al. [1992] involves statistical methods, which include statistical comparisons, trend detection, and geostatistical methods. Many optimization techniques under this category either maximize or minimize a given objective function while accounting for some other constraints. Among the proposed methodologies, a number of them consider the minimization of the estimation variance of the monitored parameter estimates. Carrera et al. [1984] applied a method derived from kriging to determine the optimal locations of sampling sites for the estimation of fluoride

concentrations in groundwater. Rouhani [1985] proceeded by minimizing a weighted sum of variance of estimation using a variance-reduction algorithm; sampling sites which provide the maximum reduction in the variance of estimation are added to the monitoring network each at a time. Woldt and Bogardi [1992] used a modified variance-reduction algorithm combined with a multicriteria decision making using composite programming for designing groundwater quality monitoring networks.

These methods are relatively simple and are advantageous in that they are formulated within a geostatistical framework. This makes their modification, to integrate with the plume estimation methods developed in this work, relatively straightforward. The variance-reduction methodology, which forms the basis of many of these works, can preserve its simplicity and advantages while working with the geostatistical plume estimation methods developed in this work, rather than with kriging. However, these methods do not usually consider knowledge of flow and transport processes.

The Task Committee on the State of the Art in Long-Term Groundwater Monitoring Design [ASCE, 2003] built on the categorization of Loaiciga et al. [1992] and added to the second approach probabilistic methods (including Kalman filters, probabilistic simulation, and hierarchical monitoring) and mathematical optimization methods (including integer programming, simulated annealing, and genetic algorithms). Given the multiobjective and uncertain nature of sampling network design problems, various methods within this approach are usually coupled and used together [Zhang et al., 2005]. McKinney and Loucks [1992] presented an optimization algorithm for selecting new

monitoring locations so as to minimize a simulation model's estimation variance. Gorelick [1983], Yeh [1992], Ahlfeld and Heidari [1994], Wagner [1995] and Ahlfeld and Mulligan [2000] provided extensive reviews on coupling simulation models with optimization models. Herrera and Pinder [2005] included some of these within a third approach to groundwater sampling network design, the modeling framework, which "combines the knowledge of physics of groundwater physics and contaminant transport, summarized in a mathematical model, with data." Within this framework, many researchers combined mathematical plume interpolators, such as kriging, with fate and transport models, to estimate contaminant mass. Reed et al. [2000] presented a methodology to reduce an existing monitoring network (without significant loss of accuracy in contaminant mass estimation), using such an approach. The total mass was estimated via three plume interpolation methods: inverse-distance weighting, ordinary kriging, and a hybrid method that combines the two approaches. A genetic algorithm was used to search all potential subsets of wells for the optimal ones for the interpolation. Reed et al. [2001] built on this approach and examined the tradeoffs between sampling cost reduction and local concentration estimation errors. Wu et al. [2005] took a similar approach but added and emphasized the plume moment constraints on the optimization procedure, leading to more accurate sampling designs. Several recent studies have focused on the development and refinement of optimization algorithms for these types of methods designs. Comparisons of such methods have been carried out by Reed and Minsker [2000], Kollat and Reed [2006], and Wu et al. [2006].

Wu et al. [2006] compared two methods that may be more related to the work presented

31

here, as they consider uncertainty in the hydraulic conductivity field. These methods, Monte Carlo simple genetic algorithm and noisy genetic algorithm, couple a genetic algorithm with a numerical flow-and-transport simulator and a global plume interpolator to identify the optimal sampling network for contaminant plume monitoring. Zhang et al. [2005] also considered the uncertainty caused by the lack of knowledge in the hydraulic conductivity field. They solved the optimization problem using a genetic algorithm combined with a Kalman filter. However, most of these methods include separate components for transport modeling simulation and for contaminant plume estimation. Thus, the transport information available to the numerical model does not explicitly assist the plume interpolator, or directly reduce its uncertainty.

The plume estimation methods developed in this dissertation integrate these two components (flow-and-transport model and plume interpolator), and present an estimation variance that accounts for transport information and uncertainty. The well-selection method, like the plume-estimation methods, is developed within a geostatistical framework similar to the variance-reduction methods described above. The biggest difference from these earlier well-selection methods is the underlying plume-estimation method.

Chapter 3
Integration of Transport Information into Geostatistical Analyses[1]

1. Introduction

The goal of groundwater contaminant plume interpolation is to provide an accurate representation of the spatial distribution of a plume given the data limitations associated with sparse monitoring networks with irregular geometries. The extent of a contaminant plume has to be known in order to avoid, or at least minimize, pumping of polluted groundwater into drinking-water wells and water distribution systems. Because samples can typically only be taken at a few discrete points of a plume (i.e. at wells), they have to be interpolated in order to depict the span of the whole plume. Groundwater contaminant plume interpolation is a difficult task, as contaminant concentration fields are highly heterogeneous, anisotropic, and non-stationary phenomena [Reed et al., 2004].

Several forms of information, in addition to concentration measurements, are available in some cases to the practitioner and can potentially improve interpolation. Often, these secondary data (e.g. transmissivity) can be spatially and/or temporally correlated with the measured concentrations. These data can enhance the estimation results if used, for example, in a cokriging setup [e.g. Cassiani and Medina, 1997]. However, many types of data cannot easily be assimilated into a geostatistical analysis. This category may include information about the physical processes that created the contaminant plume, or a groundwater model that quantitatively describes the flow and transport in the aquifer.

In many environmental applications, data are extremely limited. Therefore, the ability to incorporate different forms of data (e.g. physical behavior as expressed in fate-and-transport models and spatial correlation as quantified by geostatistical analyses) would improve the ability to describe the shape and extent of groundwater contaminant plumes. The high cost associated with groundwater monitoring and the high risks posed by these contaminants contribute to the importance of developing a robust plume distribution estimation technique that takes into account diverse types of data for plume interpolation.

2. Objective

As described in Chapter 2, currently available methods for plume estimation cannot fully take advantage of prior knowledge of flow and transport information or the location of a contaminant source. In this Chapter, we present two new methods for estimating the distribution of groundwater contaminant plumes. Unlike existing tools, these methods take into account both the spatial covariance structure of the concentration field, and

available flow and transport information. The concentration data are directly coupled with the transport model to estimate concentrations at unsampled locations. Thus, both data pertaining to the autocorrelation of concentrations, and results of a physical flow-and-transport model are assimilated into a common framework.

The first method is referred to as Inverse/Forward Modeling (IFM). It uses an existing transport model for the aquifer and knowledge of the contaminant source location to estimate the time series of the contaminant release into the aquifer and its associated uncertainty. Together with the transport model, this release history is then used to estimate the current plume distribution and its uncertainty.

The second proposed method, Transport-Enhanced Kriging (TrEK), combines the merits of IFM with geostatistical kriging, in order to also take advantage of the spatial covariance structure of the available concentration measurements. Thus, the resulting predictions are both physically feasible, and exhibit the expected spatial autocorrelation characteristics.

3. Methodology

3.1 Geostatistical inverse modeling

The following section outlines the geostatistical approach to inverse modeling, which we implement in this work to estimate the time series of the contaminant release into the aquifer in both of the proposed methods. No derivations are provided, and the reader is

referred to, for example, Kitanidis and Vomvoris [1983], Snodgrass and Kitanidis [1997] and Michalak and Kitanidis [2004a] for additional details.

Inverse methods use modeling and statistical tools to determine the historical distribution of observed contamination, the location of contaminant sources, or the release history from a known source. The inverse problem is often underdetermined and an infinite number of parameters that are consistent with the data may be obtained [Woodbury and Ulrych, 1996]. Kabala and Skaggs [1998] stress that this nonuniqueness of the solution is caused by the ill-posed nature of the physical problem and is not associated with any particular solution methodology, nor can nonuniqueness be eliminated with any particular procedure. Indeed, there are numerous inverse modeling methods; see, for example, reviews in Michalak and Kitanidis [2004a] and in Atmadja and Bagtzoglou [2001]. Of the available methods, only those that provide a stochastic solution to the problem of source estimation can potentially be applied to improve plume estimation, because the uncertainty of the loading history needs to be quantified in order to, in turn, determine the precision with which the plume distribution can be estimated. Backward tracking [Bagtzoglou et al. 1991; Bagtzoglou et al. 1992] and adjoint-derived source distribution probabilities [Neupauer and Wilson, 2001] are two available stochastic approaches, but these methods are only applicable to identifying the location or release time of a single instantaneous point source.

The geostatistical approach to inverse modeling [Snodgrass and Kitanidis 1997; Michalak and Kitanidis, 2003, 2004a,b, 2005; Butera and Tanda, 2003] and the minimum relative

36

entropy method [Woodbury and Ulrych, 1996, 1998], conversely, are stochastic methods that provide a function estimate to characterize the historical contaminant distribution, source location, or release history. For these methods, the contaminant distribution or source description is not limited to a small set number of fixed parameters but can instead vary in space and/or in time. The applicability of the geostatistical method has already been demonstrated for multidimensional heterogeneous media [Michalak and Kitanidis, 2004a] and is similar in form to geostatistical kriging. The stochastic nature of the geostatistical approach to inverse modeling, which allows the uncertainty of the estimate to be quantified, and its affinity to kriging, which allows for a convenient setup of a familiar equation system, make it an appropriate basis for the proposed plume estimation methods.

The geostatistical approach to inverse modeling is based on the dual criterion of reproducing available concentration data $\mathbf{z}^*$ and preserving spatial or temporal autocorrelation in some unknown function $\mathbf{s}$ $[m \times 1]$, such as the release history of contaminants into the aquifer. The method begins by calculating the sensitivity of available concentration measurements to the unknown function, and assigning this information to a Jacobian matrix $\mathbf{H}^*$ (i.e., $H_{ij}^* = \partial z_i^* / \partial s_j$). This sensitivity information is then used in combination with the available data and their associated uncertainty to estimate a discretized version of the unknown function. This unknown function is assumed to be spatially- and/or temporally-autocorrelated.

This form of inverse-modeling relies on the existence of a groundwater flow and a contaminant transport model. This usually implies knowledge of parameters such as the hydraulic conductivity field, the dispersivity tensor, the boundary conditions, etc. In addition, the transport model must either be linear or meet certain conditions which would enable a quasi-linear approximation [see Kitanidis, 1996, for details]. In this Chapter, as in past applications of the geostatistical approach to inverse modeling for contaminant source identification [e.g. Snodgrass and Kitanidis, 1997; Michalak and Kitanidis, 2002, 2003, 2004a, b], the transport model itself is assumed to be deterministically known. The concentrations of the source and of the measurements, on the other hand, are treated as random functions in a geostatistical framework.

The transport model is sampled at measurement locations and times to yield a sensitivity matrix. If we are estimating the source release history on the basis of concentrations measured in the plume at a single time but at different locations, each element H_{ij}^* of this sensitivity matrix represents the sensitivity of the concentration at location x_i^* $[i = 1..n]$ to the source intensity at time t_j $[j = 1..m]$. The product of this sensitivity matrix and a temporal source function $\mathbf{s}$ reproduces the measured concentrations $\mathbf{z}^*$ at all locations x_i^*, to within a model-data mismatch error ϵ:

$$\mathbf{z}^* = \mathbf{H}^*\mathbf{s} + \epsilon \tag{3}$$

where ε is assumed to be a zero-mean model-data mismatch error with covariance matrix $\mathbf{R}$. Because measurement errors are usually not correlated in space and typically

have a uniform variance σ_R^2, the covariance structure is modeled as $\mathbf{R} = \sigma_R^2 \mathbf{I}$, where $\mathbf{I}$ is an $[n \times n]$ identity matrix. σ_R^2 can either be known a priori, or estimated from available data (see below). Note that the model-data mismatch can also include transport error or microvariability not explained by the model. Thus, even though the model is deterministic, some forms of transport model uncertainties can be taken into account.

The expected value of $\mathbf{s}$ is modeled as $E[\mathbf{s}] = \mathbf{X}_s \boldsymbol{\beta}_s$, where $\mathbf{X}_s$ is a known $\begin{bmatrix} m \times p_s \end{bmatrix}$ matrix of basis functions and $\boldsymbol{\beta}_s$ are p_s unknown drift coefficients. The prior covariance function of $\mathbf{s}$ is a known function $\mathbf{Q}_s(\boldsymbol{\theta}) = E\left[\left(\mathbf{s} - \mathbf{X}_s \boldsymbol{\beta}_s \right) \left(\mathbf{s} - \mathbf{X}_s \boldsymbol{\beta}_s \right)^T \right]$ of unknown parameters $\boldsymbol{\theta}$, which can be estimated, for example, using a Restricted Maximum Likelihood (RML) approach [Kitanidis, 1995]. This approach minimizes the negative log-likelihood of the probability of the measurements with respect to the covariance parameters $\boldsymbol{\theta}$:

$$L_{\boldsymbol{\theta}} = \frac{1}{2}\ln | \boldsymbol{\Sigma} | + \frac{1}{2}\ln | \mathbf{X}_s^T \mathbf{H}^{*T} \boldsymbol{\Sigma}^{-1} \mathbf{H}^* \mathbf{X}_s | + \frac{1}{2}\mathbf{z}^{*T} \boldsymbol{\Xi} \mathbf{z}^* \tag{4}$$

where $\boldsymbol{\Sigma} = \mathbf{H}^* \mathbf{Q}_s \mathbf{H}^{*T} + \mathbf{R}$ and

$$\boldsymbol{\Xi} = \boldsymbol{\Sigma}^{-1} - \boldsymbol{\Sigma}^{-1} \tilde{\mathbf{H}} \mathbf{X}_s \left(\mathbf{X}_s^T \tilde{\mathbf{H}}^T \boldsymbol{\Sigma}^{-1} \tilde{\mathbf{H}} \mathbf{X}_s \right)^{-1} \mathbf{X}_s^T \tilde{\mathbf{H}}^T \boldsymbol{\Sigma}^{-1} \tag{5}$$

Once the covariance parameters $\boldsymbol{\theta}$ have been found, the source is estimated by minimizing:

$$L_{\mathbf{s},\boldsymbol{\beta}_s} = \frac{1}{2}(\tilde{\mathbf{z}} - \tilde{\mathbf{H}}\mathbf{s})^T \mathbf{R}^{-1}(\tilde{\mathbf{z}} - \tilde{\mathbf{H}}\mathbf{s}) + \frac{1}{2}(\mathbf{s} - \mathbf{X_s}\boldsymbol{\beta}_s)^T \mathbf{Q}_s^{-1}(\mathbf{s} - \mathbf{X_s}\boldsymbol{\beta}_s) \qquad (6)$$

which leads to the following system of $n + p_s$ equations

$$\begin{bmatrix} \tilde{\mathbf{H}}\mathbf{Q}_s\tilde{\mathbf{H}}^T + \mathbf{R} & \tilde{\mathbf{H}}\mathbf{X}_s \\ \left(\tilde{\mathbf{H}}\mathbf{X}_s\right)^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\Lambda}_s^T \\ \mathbf{M}_s \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{H}}\mathbf{Q}_s \\ \mathbf{X}_s^T \end{bmatrix} \qquad (7)$$

which can be solved for the $m \times n$ matrix $\boldsymbol{\Lambda}_s$ and the $p_s \times m$ matrix $\mathbf{M}_s$, from which

the best estimate and posterior covariance of the source function can be found:

$$\hat{\mathbf{s}} = \boldsymbol{\Lambda}_s\tilde{\mathbf{z}} \qquad (8)$$

$$\mathbf{V}_{\hat{\mathbf{s}}} = -\mathbf{X}_s\mathbf{M}_s + \mathbf{Q_s} - \mathbf{Q_s}\tilde{\mathbf{H}}^T\boldsymbol{\Lambda}_s^T \qquad (9)$$

3.2 Proposed methods
Inverse/Forward Modeling (IFM)

In the IFM approach, the contaminant release history estimated using geostatistical

inverse modeling is used to obtain an estimate of the distribution of the plume at the time

of sampling, and its associated uncertainty. The contaminant plume is reconstructed by

simulating transport of the estimated release history up to the measurement time(s). This

step involves calculating an expanded sensitivity matrix $\mathbf{H}$ $[N \times m]$ which defines the

sensitivity of N locations in the aquifer at the time when the measurements were taken to

the m discretized times of the contaminant release history. The elements of $\mathbf{H}$ can be

obtained numerically by running the groundwater transport model with unit releases of

contaminant for each time t_j $[j = 1..m]$ and recording the resulting concentration at each of the N locations where the plume concentration is to be estimated. Typically, these N points would be laid out on a regular grid, to allow for easy contouring of the estimated plume distribution. The best estimate of the concentrations is estimated using the resulting linear model

$$\hat{\mathbf{z}} = \mathbf{H}\hat{\mathbf{s}} \qquad (10)$$

where $\hat{\mathbf{z}}$ are the $[N \times 1]$ estimates of the concentrations $\mathbf{z}$. Note that we use $\mathbf{z}^*$ to designate concentration at measurement locations, and $\mathbf{z}$ to designate the concentrations that we are interested in estimating. The same sensitivity matrix $\mathbf{H}$ can also be used to define the uncertainty and covariance of the estimated plume distribution, which is a function of the uncertainty associated with the recovered contaminant release history:

$$\mathbf{V}_{\hat{\mathbf{z}}} = \mathbf{H}\mathbf{V}_{\hat{\mathbf{s}}}\mathbf{H}^T \qquad (11)$$

The diagonal elements of this $[N \times N]$ matrix represent the uncertainty of the estimated concentration spatial distribution.

Transport-enhanced kriging (TrEK)

The second plume estimation method also relies on the estimated release history (eq. (8)) but also takes into account the spatial covariance $\mathbf{Q}_z$ of the plume concentration distribution, where $\mathbf{Q}_z(\boldsymbol{\theta}) = E[(\mathbf{z} - \mathbf{X}_z\boldsymbol{\beta}_z)(\mathbf{z} - \mathbf{X}_z\boldsymbol{\beta}_z)^T]$, and $E[\mathbf{z}] = \mathbf{X}_z\boldsymbol{\beta}_z$. This covariance can also be estimated using the subset of available concentration values $\mathbf{z}^*$ using an RML approach, with the objective function simplifying to:

41

$$L_{\boldsymbol{\theta}} = \frac{1}{2}\ln \mid \mathbf{Q}^*_{z^*} \mid + \frac{1}{2}\ln \mid \mathbf{X}^{*T}_z \mathbf{Q}^{*-1}_{z^*} \mathbf{X}^*_z \mid$$
$$+ \frac{1}{2}\mathbf{z}^{*T} \left(\mathbf{Q}^{*-1}_{z^*} - \mathbf{Q}^{*-1}_{z^*}\mathbf{X}^*_z \left(\mathbf{X}^{*T}_z \mathbf{Q}^{*-1}_{z^*} \mathbf{X}^*_z \right)^{-1} \mathbf{X}^{*T}_z \mathbf{Q}^{*-1}_{z^*} \right)\mathbf{z}^* \tag{12}$$

where the basis functions $\mathbf{X}^*_z$ $[n \times p_z]$ and covariance model $\mathbf{Q}^*_{z^*}$ $[n \times n]$ are the same as $\mathbf{X}_z$ and $\mathbf{Q}_z$, but are only evaluated at the measurement locations. This correlation in deviations of the concentration distribution from its underlying trend $\mathbf{X}_z\boldsymbol{\beta}_z$ may provide information not available through the transport model alone. In assimilating the covariance information $\mathbf{Q}_z$, we simultaneously require that the covariance structure imposed by the physical transport model and that implied by the spatial autocorrelation of the measurements are honored. Explicitly, the residuals between the model predictions $\mathbf{H}\hat{\mathbf{s}}$ and the TrEK estimates $\hat{\mathbf{z}}$ should be consistent with the covariance structure implied by the uncertainty of the release history and by the transport model, i.e. $\mathbf{H}\mathbf{V}_{\hat{\mathbf{s}}}\mathbf{H}^T$. In addition, the residuals from the measurement-space trend $\mathbf{X}_z\boldsymbol{\beta}_z$ should have the spatial structure $\mathbf{Q}_z$. These requirements are assimilated in a second Bayesian step, minimizing:

$$L_{\mathbf{z},\boldsymbol{\beta}_z} = \frac{1}{2}(\mathbf{H}\hat{\mathbf{s}} - \mathbf{z})^T (\mathbf{H}\mathbf{V}_{\hat{\mathbf{s}}}\mathbf{H}^T)^{-1}(\mathbf{H}\hat{\mathbf{s}} - \mathbf{z}) + \frac{1}{2}(\mathbf{z} - \mathbf{X}_z\boldsymbol{\beta}_z)^T \mathbf{Q}^{-1}_z(\mathbf{z} - \mathbf{X}_z\boldsymbol{\beta}_z) \tag{13}$$

with respect to the unknown concentrations $\mathbf{z}$ $[N \times 1]$ and the p_z spatial drift parameters $\boldsymbol{\beta}_z$. The first term in this objective function makes use of the inverse modeling best estimate $\hat{\mathbf{s}}$ and its posterior covariance, whereas the second term requires the estimates to follow the kriging trend $\mathbf{X}_z\boldsymbol{\beta}_z$ and covariance $\mathbf{Q}_z$. Note that the measured concentration values $\mathbf{z}^*$ do not reappear in this second objective function.

The second system of equations to be solved becomes:

$$\begin{bmatrix} \mathbf{Q}_z + \mathbf{H}\mathbf{V}_{\hat{\mathbf{s}}}\mathbf{H}^T & \mathbf{X}_z \\ \mathbf{X}_z^T & \mathbf{0} \end{bmatrix}\begin{bmatrix} \mathbf{\Lambda}_z^T \\ \mathbf{M}_z \end{bmatrix} = \begin{bmatrix} \mathbf{Q}_z \\ \mathbf{X}_z^T \end{bmatrix} \tag{14}$$

Note that these $(N + p_z)$ equations can be set up only after the inverse modeling system has been solved for $\hat{\mathbf{s}}$ and for $\mathbf{V}_{\hat{\mathbf{s}}}$. The solutions of this system, $\mathbf{\Lambda}_z$ $[N \times N]$ and $\mathbf{M}_z$ $[p_z \times N]$, are used to compute the plume concentration distribution and its covariance:

$$\hat{\mathbf{z}} = \mathbf{\Lambda}_z \mathbf{H}\hat{\mathbf{s}} \tag{15}$$

$$\mathbf{V}_{\hat{\mathbf{z}}} = \mathbf{Q}_z - \mathbf{Q}_z\mathbf{\Lambda}_z^T - \mathbf{X}_z\mathbf{M}_z \tag{16}$$

Here $\hat{\mathbf{z}}$ is the TrEK best estimate of the plume distribution, and $\mathbf{V}_{\hat{\mathbf{z}}}$ is the corresponding covariance structure.

Uncertainties modeled in IFM and TrEK

For geostatistical inverse modeling, the matrix $\mathbf{V}_{\hat{\mathbf{s}}}$ represents the covariance of the a posteriori source function residuals $(\mathbf{s} - \hat{\mathbf{s}})$, and its diagonal elements correspond to the uncertainty of the best estimate $\hat{\mathbf{s}}$. For IFM, $\mathbf{V}_{\hat{\mathbf{s}}}$ is used to define the covariance of the full plume concentration distribution, using the model matrix $\mathbf{H}$ to yield $\mathbf{H}\mathbf{V}_{\hat{\mathbf{s}}}\mathbf{H}^T$. Thus, for a given model $\mathbf{H}$, the uncertainty of an estimate at a given point depends on its sensitivity to the source function. For example, the sensitivity will be greater for regions immediately downgradient from the source, while areas further downgradient or away from the principal flow direction are much less sensitive to the release intensity. The

43

result is that the uncertainty will decrease, in general, with growing distance from the contamination source. The uncertainty also depends on $\mathbf{V}_{\hat{s}}$ such that plume areas corresponding to times at which the release was highly uncertain will have higher variance than those corresponding to more certain elements of $\mathbf{s}$.

Similar considerations determine the uncertainty associated with TrEK estimates, because this method also relies on the estimated release history. However, TrEK may be more directly affected by measurement locations, since it also considers the spatial covariance structure of the plume distribution, $\mathbf{Q}_z$. This additional information further constrains the estimate for each point, with the result that the TrEK uncertainty is always smaller than that of IFM.

3.3 Kriging with a trend (KT)

The new methods presented in the previous subsections are compared to kriging in the applications below (Section 4 of this Chapter). This familiar and relatively simple method was chosen in order to keep the focus on IFM and TrEK, without being distracted too much by a particular method's extraneous details, limitations, etc. Therefore, a brief overview of kriging is presented here for reference. In kriging with a trend, we model $\mathbf{z}$ as a random vector with expected value $\mathbf{X}_z \boldsymbol{\beta}_z$, where $\mathbf{X}_z$ is a known $[N \times p_z]$ matrix and $\boldsymbol{\beta}_z$ are p_z unknown drift coefficients, representing the mean and the trend of the unknown concentration distribution. The prior covariance of $\mathbf{z}$ is $\mathbf{Q}_z$, a known function of unknown parameter(s) $\boldsymbol{\theta}$, which can be optimized using RML as presented in

Section 4.2. Once these parameters have been estimated, the best estimate $\hat{\mathbf{z}}$ is obtained

using a linear combination of known measurements $\mathbf{z}^*$. The coefficients $\mathbf{\Lambda}_{\mathbf{z}^*}$ $[N \times n]$

and $\mathbf{M}_{\mathbf{z}^*}$ $[p_z \times N]$ of this linear system are calculated by solving a kriging system of the

form

$$\begin{bmatrix} \mathbf{Q}^*_{z^*} & \mathbf{X}^*_z \\ \mathbf{X}^{*T}_z & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{\Lambda}^T_{z^*} \\ \mathbf{M}_{z^*} \end{bmatrix} = \begin{bmatrix} \mathbf{Q}_{z^*} \\ \mathbf{X}^T_z \end{bmatrix} \tag{17}$$

Note that $\mathbf{Q}^*_{z^*}$ $[n \times n]$ is the covariance matrix of measurements, whereas $\mathbf{Q}_{z^*}$ $[n \times N]$

is the covariance between concentrations at the n measurements locations and the N

estimation locations. Similarly, $\mathbf{X}^*_z$ are the basis functions evaluated at the measurement

locations, and $\mathbf{X}_z$ are those evaluated at the estimation locations. The kriging estimator is

$$\hat{\mathbf{z}} = \mathbf{\Lambda}_{z^*} \mathbf{z}^* \tag{18}$$

and the a posteriori kriging covariance is

$$\mathbf{V}_{\hat{\mathbf{z}}} = \mathbf{Q}_z - \mathbf{Q}^T_{z^*} \mathbf{\Lambda}^T_{z^*} - \mathbf{X}_z \mathbf{M}_{z^*} \tag{19}$$

To account for measurement errors and preserve continuity, continuous part kriging

[Kitanidis, 1997], also referred to as kriging with known measurement error variance

[Wackernagel, 2003], can be applied by adding σ^2_R to the terms on the diagonal of $\mathbf{Q}^*_{z^*}$

in Equation (17).

4. Test Cases

The following section presents sample applications of the two methods developed in this Chapter. These examples involve the estimation of a contaminant plume distribution in a confined aquifer. In all examples, an aquifer is assumed to have been contaminated by a single point source with known location but unknown loading history as a function of time, $\mathbf{s}$. Measurements $\mathbf{z}^*$ taken at time T and knowledge of the hydrogeological conditions, $\mathbf{H}$ and $\mathbf{H}^*$, are used in implementing IFM, TrEK, and KT, to estimate the full plume distributions $\mathbf{z}$. Hypothetical examples were chosen to illustrate and verify the capabilities of the methods in a setup where the true concentration distributions are known.

The two cases considered are a plume in a one-dimensional homogeneous aquifer, and a plume in a two-dimensional heterogeneous aquifer. In addition, sensitivity analyses are performed to evaluate the effects of monitoring network configurations, measurement errors, the time elapsed prior to sampling, and spatially-correlated model-mismatch errors.
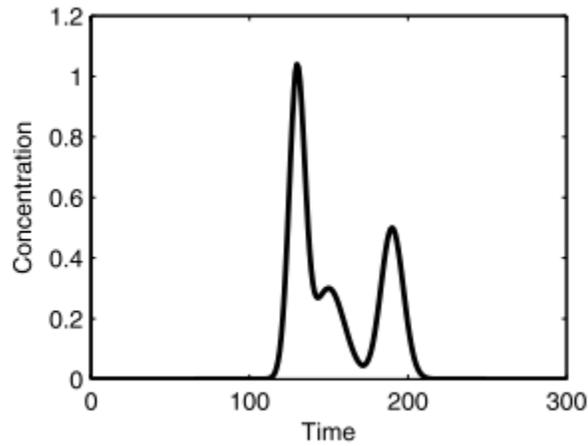
Figure 3.1: Example 1 - Source release history.

4.1 Example 1: One-dimensional homogeneous aquifer

In the first example we use the setup previously implemented by Skaggs and Kabala [1994, 1995] and Snodgrass and Kitanidis [1997], among others. This example involves the release of a conservative solute into a one-dimensional homogeneous aquifer with a steady-state flow field. The solute is released only at the origin ($x_0 = 0$) and the concentration is measured at various points in the aquifer at some later time T. The advective and dispersive transport of this solute can be expressed analytically as:

$$C(x,T) = \int_0^T s(t)g_{1D}(x,T-t)dt \qquad (20)$$

where $C(x,T)$ is the concentration at distance x from the source and time T. The source is a function of time and is expressed by $s(t)$. The one-dimensional transfer function $g_{1D}(x,T-t)$ relates the source concentration at time t to the concentration measured at point x and time T [Skaggs and Kabala, 1994]:

$$g_{1D}(x, T - t) = \frac{x}{2\sqrt{\pi D_L (T - t)^3}} \exp\left[-\frac{(x - v(T - t))^2}{4 D_L (T - t)}\right] \qquad (21)$$

where D_L is the longitudinal dispersion coefficient and v is the seepage velocity.

Note that $z_i = C(x_i, T)$ is the concentration at location x_i.

The synthetic pollution event is obtained through the numerical integration of eq. (20) at $T = 330$ time units, with longitudinal dispersivity $D_L = 1$, flow velocity $v = 1$ and true release history

$$s(t) = \exp\left[-\frac{(t - 130)^2}{50}\right] + 0.3 \exp\left[-\frac{(t - 150)^2}{200}\right] + 0.5 \exp\left[-\frac{(t - 190)^2}{98}\right] \qquad (22)$$

This release history is illustrated in Figure 3.1. Note that this release history is used only for the purposes of simulating the plume, but is then considered unknown in subsequent steps. This one-dimensional plume is sampled at $n = 11$ locations x^*. Negligible measurement error with a variance of 10^{-12} is added to these concentrations to yield the observations $\mathbf{z}^*$. We use these observations $\mathbf{z}^*$ and the transport information (eq. (21)) to estimate the plume concentrations $\mathbf{z}$ at $m = 301$ locations ($x = 0, 1, 2, ..., 300$) using KT, IFM and TrEK.

Kriging with a trend

For kriging with a trend we use a cubic Generalized Covariance Function (GCF):

$$\mathbf{Q}_z(h) = \theta_z h^3 \qquad (23)$$

48

and optimize the parameter θ_z using RML. In order to obtain better structural information, we use more plume samples for this optimization step than for the estimation below. For this example, the estimated parameter is $\theta_z = 5.9 \times 10^{-8}$.

The plume is modeled using a linear trend in x, yielding $\mathbf{X}_z^* = \begin{bmatrix} [\mathbf{1}] & \mathbf{x}^* \end{bmatrix}$, where $[\mathbf{1}]$ is an $[n \times 1]$ vector of ones. The coefficients $\mathbf{\Lambda}_{\mathbf{z}^*}$ and $\mathbf{M}_{\mathbf{z}^*}$, calculated with eq. (17), are used to obtain the best estimate $\hat{\mathbf{z}}$ (eq. (18)) of the plume distribution and the kriging covariance $\mathbf{V}_{\hat{\mathbf{z}}}$ (eq. (19)). The best estimate and associated uncertainty are plotted in Figure 3.2.

Inverse/Forward Modeling

We use the transfer function (eq. (21)) to define the sensitivity matrix $\mathbf{H}^*$ according to [Snodgrass and Kitanidis, 1997]:

$$H_{ij}^* = \frac{x_i^*}{2\sqrt{\pi D_L (T - t_j)^3}} \exp\left[-\frac{(x_i^* - v(T - t_j))^2}{4 D_L (T - t_j)} \right] \Delta t \qquad (24)$$

where $\Delta t = 1$ is the discretization of the source function in time. The covariance structure $\mathbf{Q}_s$ of the source function is also modeled using a cubic GCF. For the presented example, the best estimate for its single parameter (optimized using RML) is $\theta_s = 1.3 \times 10^{-5}$. This information is used to estimate the source release history (eq. (7) - (9)). Its basis functions are $\mathbf{X}_s = \begin{bmatrix} [\mathbf{1}] & \mathbf{t} \end{bmatrix}$, where $[\mathbf{1}]$ is an $[m \times 1]$ vector of ones and the elements of $\mathbf{t} = 0, 1, ..., 300$ correspond to the discretized times of the release history. We then form a second sensitivity matrix $\mathbf{H}$ for the estimation points $\mathbf{x}$, by replacing the measurement locations x_i^* with estimation locations x_i in equation (24). The resulting estimated plume distribution $\hat{\mathbf{z}}$ and its covariance $\mathbf{H}\mathbf{V}_{\hat{s}}\mathbf{H}^T$ are illustrated in Figure 3.2.
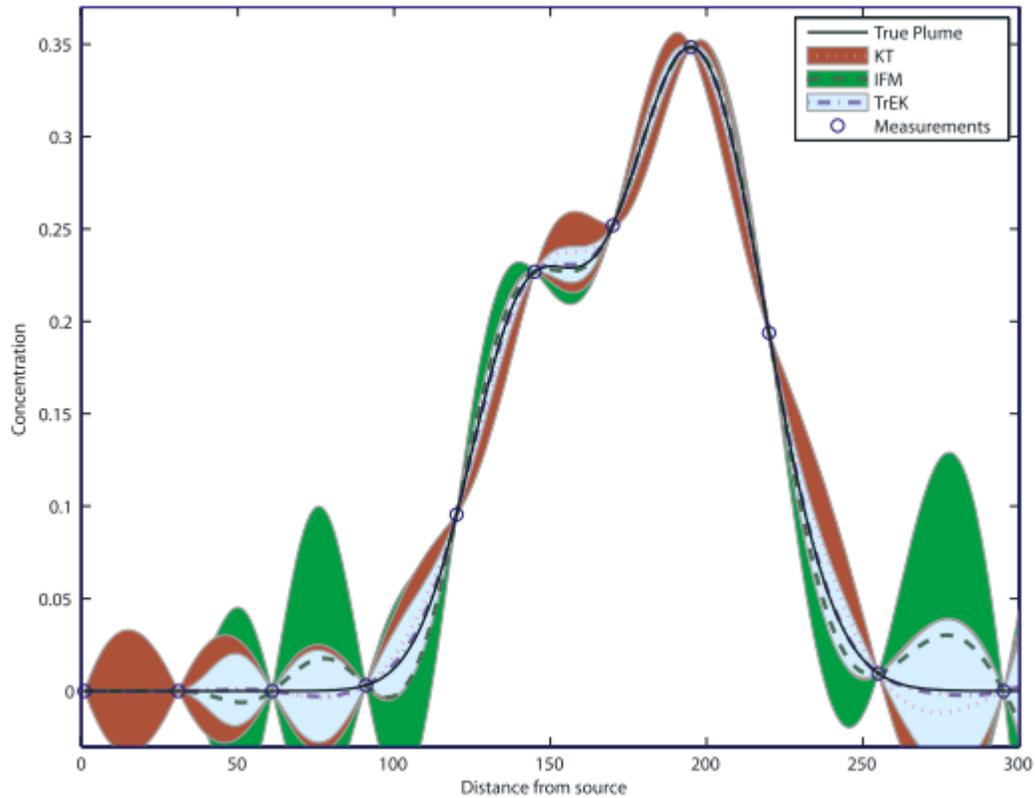
Figure 3.2: Example 1 - Homogeneous one-dimensional plume. Lines represent best estimates; shaded areas represent 95% confidence intervals.

The IFM best estimate fluctuates around the actual plume while reproducing all of the measurements to within the prescribed measurement error. The plume samples were all taken 30 time units after the modeled contamination event ceased; however, the constant flow remained as before. Consequently, uncontaminated water flowed in from the left side of the aquifer. Thus, a significant concentration around $x \approx 0$ is physically nearly impossible according to the transport model. For this reason, IFM's confidence intervals vanish there.

Transport-enhanced kriging

In this third approach, we estimate the release history of the point source as described in Section 4.3. We then use the concentration covariance structure $\mathbf{Q}_z$ together with the inverse modeling results to estimate the plume distribution according to eqs. (14) - (16). Results are presented in Figure 3.2.

With 11 available measurements, and no significant measurement error ($\sigma_R = 10^{-6}$), all three methods do reasonably well in reproducing the true (unknown) concentration distribution throughout the domain. At several points, the KT and IFM curves deviate noticeably from the actual plume, whereas the TrEK estimates remain closest to the actual distribution. The confidence intervals show that uncertainty decreases close to measurement locations for all methods. For any given point, the transport-enhanced kriging (TrEK) has the lowest uncertainty because it combines the information used in KT and IFM. All methods yield reasonable estimates of uncertainty, in the sense that the

real plume falls within the 95% confidence intervals in most areas. However, the TrEK

estimate yields the most accurate representation of the true plume distribution.

4.2 Example 2: Two-dimensional heterogeneous aquifer

In this example, a contaminant is released into a deterministically heterogeneous confined

aquifer (Figure 3.3a). The concentration of the point source located at

$(x_1 = 0,\ y_1 = 260)$ is variable in time and is described by:

$$s(t) = 1.4\exp\left[-\frac{(t-850)^2}{56000}\right] + 1.1\exp\left[-\frac{(t-1700)^2}{13333}\right] \tag{25}$$

The resulting plume is measured at $n = 12$ locations (Figure 3.3b), to yield the

observations $\mathbf{z}^*$. A normally-distributed random error, with a variance of $\sigma_R^2 = 10^{-6}$,

was added to these measurements.
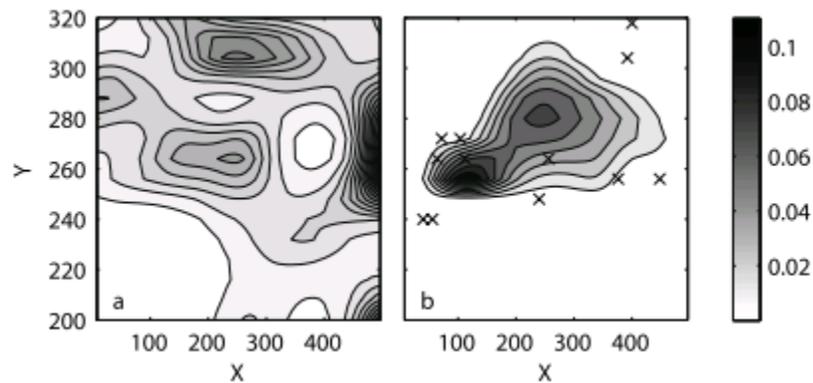
Figure 3.3: Example 2 - Heterogeneous two-dimensional aquifer. (a) Hydraulic conductivity field, (b) Simulated plume and measurement locations. All quantities are dimensionless.

The conductivity field for this aquifer is taken from Michalak and Kitanidis [2004a]. A

constant head difference was imposed between the left and right boundaries of the

aquifer, inducing flow from left to right. On the top and bottom are zero-flux (no flow) boundaries. MODFLOW-2000 [Hill et al., 2000] is used to calculate the corresponding flow field.

The numerical transport model MT3DMS [Zheng and Wang, 1999] is used in implementing IFM and TrEK to determine the effect of a unit concentration pulse on field concentrations. A short pulse is released, and the resultant concentrations are sampled repeatedly on a grid of 16×62 locations, at 188 consecutive times, corresponding to the discretization of the source release history $\mathbf{s}$. Thus, although the transport model is run only once, it effectively yields the required sensitivities for all times of the release history. The concentrations measured at each time step are used to fill one column of the sensitivity matrix $\mathbf{H}$.

We use a typical 10:1 anisotropy [e.g. Delleur, 1998] for the kriging estimation, corresponding to the modeled ratio of longitudinal (0.347) to transverse dispersivity (0.0347). Note that only the ratio of the dispersivity values is needed to define the kriging anisotropy [Chilès and Delfiner, 1999]. Using the measurements $\mathbf{z}^*$, covariance parameters are optimized both for the source function and for the measurement space using cubic GCFs, yielding $\theta_s = 1.3 \times 10^{-8}$ and $\theta_z = 1.4 \times 10^{-9}$, respectively. Note that θ_z is estimated in a transformed coordinate system, in which the y-coordinate was stretched to account for anisotropy.
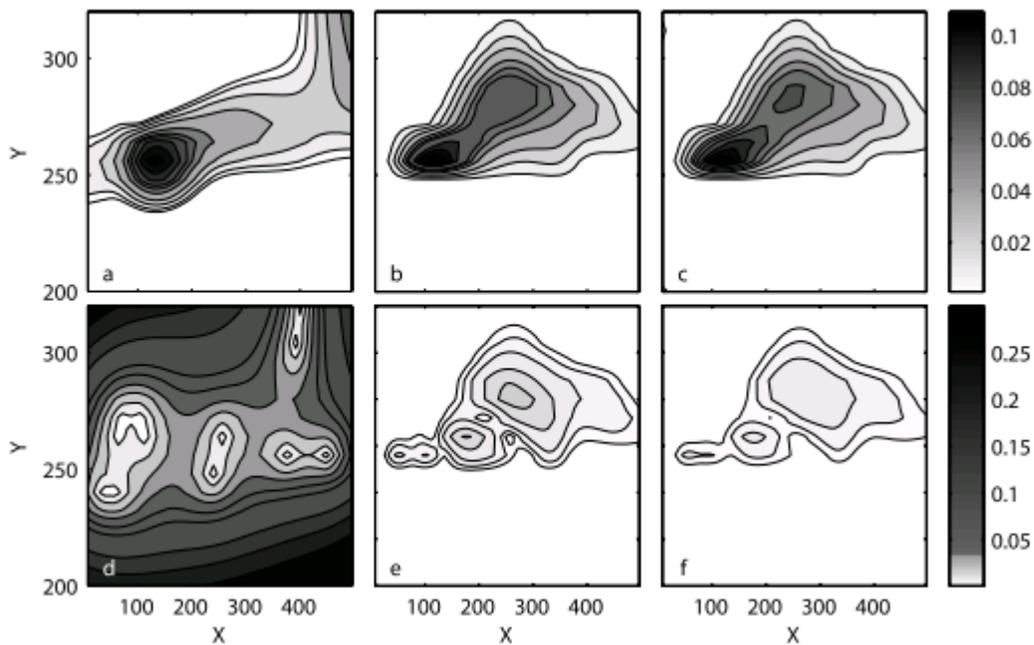
Figure 3.4: Example 2 - Results. (a) Kriging with a trend best estimate, (b) Inverse/forward modeling best estimate, (c) Transport-enhanced kriging best estimate, (d) Kriging with a trend uncertainty, (e) Inverse/forward modeling uncertainty, (f) Transport-enhanced kriging uncertainty. All quantities are dimensionless; uncertainties represented by standard deviation.

As in the previous example, KT, IFM and TrEK are used to estimate the simulated plume. The KT interpolation (Figure 3.4a) reproduces the measurements, but is unable to represent the true shape of the plume, and its uncertainty (Figure 3.4d) grows quickly away from measurement locations. The IFM method reproduces the plume much more precisely (Figure 3.4b), and yields a substantially lower uncertainty (Figure 3.4e). Contrary to KT, the uncertainty decreases with increasing distance from the center of the plume. Finally, the TrEK solution (Figure 3.4c) is very similar to the IFM solution, but slightly better. The uncertainty is always lowest relative to KT and IFM.

54
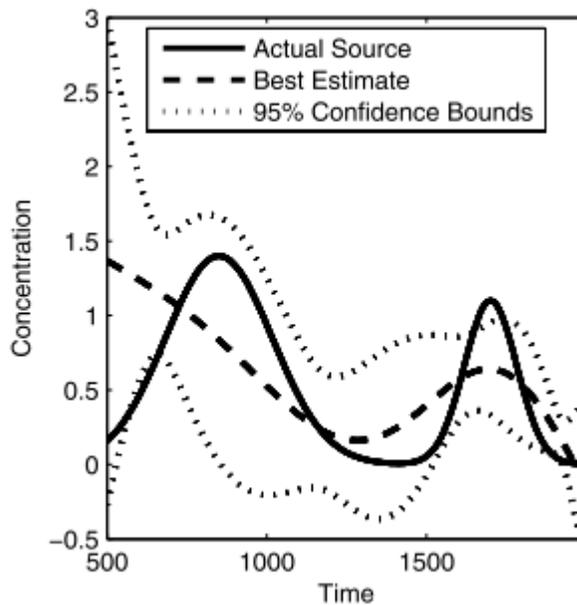
Figure 3.5: Example 2 - Actual and estimated source release history.

Figure 3.5 shows the recovered source function (eq. (8)) for this example and its confidence intervals (eq. (9)). With only 12 concentration measurements, the best estimate of the release history differs from the true loading substantially, and the uncertainty associated with the estimate is large. This demonstrates the strongly underdetermined nature of this inverse problem (12 data points are used to estimate the 188 points of the source function). Note that the estimated source release history is still accurate because the true function lies within the uncertainty bounds of the estimate. Most interestingly, however, even with a poor estimate for the source function, the IFM and TrEK methods reproduce the true plume very well (Figures 3.4b, c). This is due both to the fact that (i) the nature of the flow and transport in the aquifer provides a strong constraint on the possible plume distributions, and that (ii) the forward problem of contaminant transport is a well-posed problem, such that the resulting plume distribution

is unique for a given release history and relatively insensitive to small uncertainties in this release [Sun, 1994]. Therefore, even in cases where the measurement network is not sufficient to constrain the source release history, incorporating information on the flow and transport field leads to significant improvements in the ability to estimate the spatial distribution of the contaminant plume.

4.3 Sensitivity analyses

Sensitivity analyses were performed to investigate the effects of increased measurement errors, correlated model-data mismatch errors, and sparse sampling networks on the performance of the proposed methods. The sparse network example is also repeated for multiple times after the release, to draw conclusions about how the methods behave for plumes that have undergone different degrees of transport and mixing.

Measurement errors

In the first sensitivity test (Figure 3.6), we use a similar setup to that of Example 1, but introduce a measurement error of $\sigma_R = 0.01$. The cubic GCF parameters become $\theta_s = 8.9 \times 10^{-6}$ and $\theta_z = 6.5 \times 10^{-8}$. The best estimates are now conditioned to imperfect measurements, rather than to the actual plume, where this error can represent instrument error and/or uncorrelated errors associated with the transport model. For this case, as expected, the confidence intervals are wider relative to the case with negligible measurement error, especially near the measurement locations. TrEK again provides the most precise plume estimates.
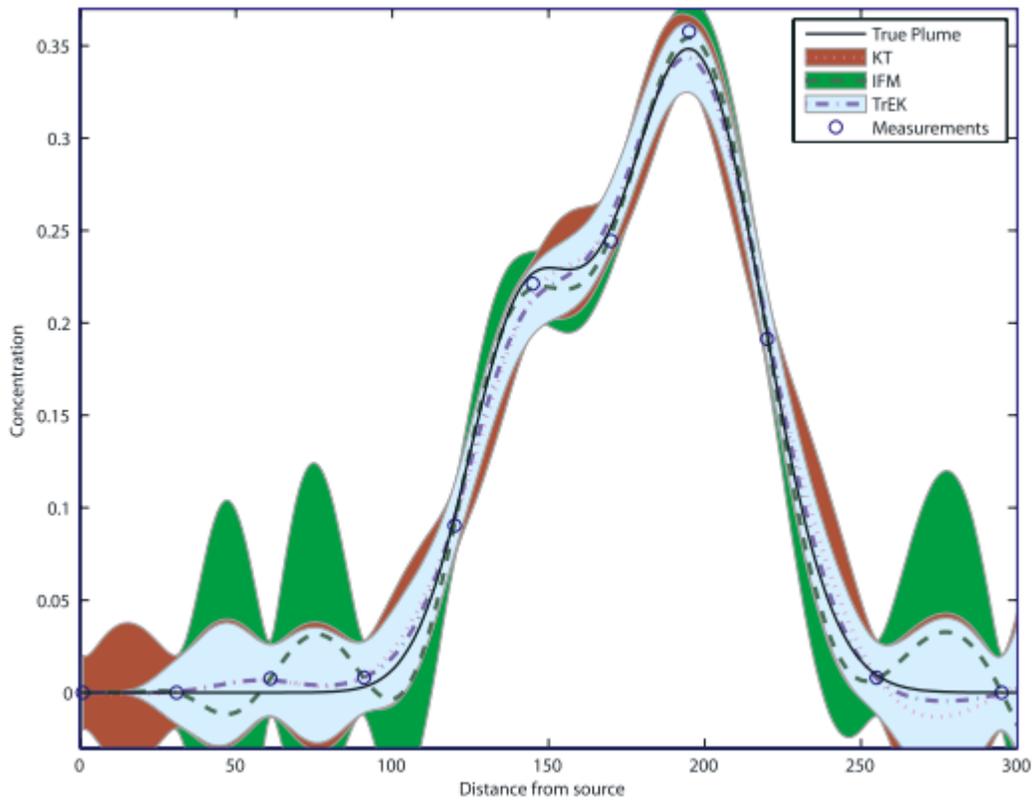
Figure 3.6: First sensitivity analysis - Example 1 with normally distributed measurement errors with standard deviation of 0.01.
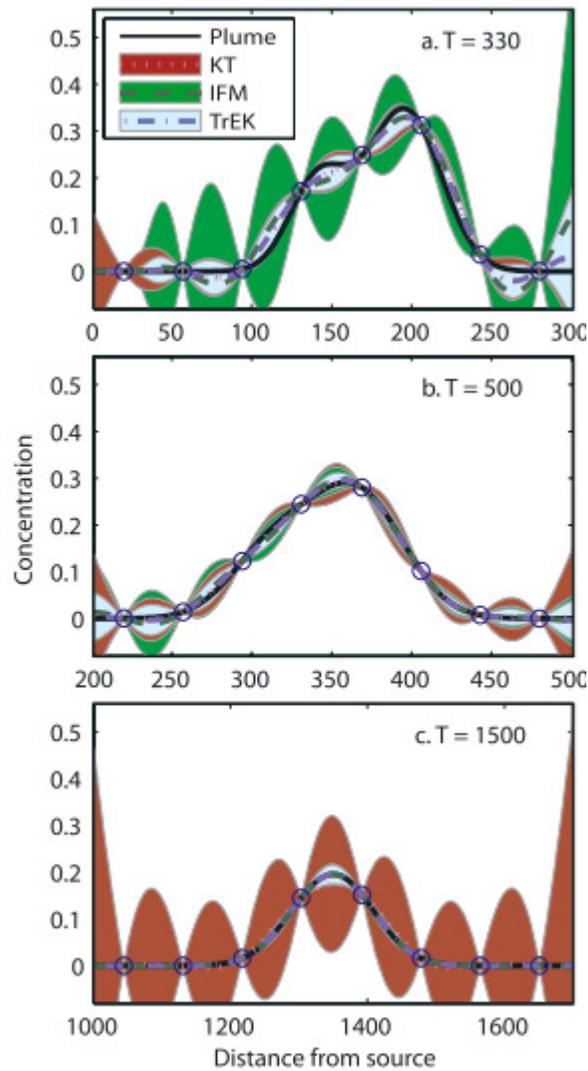
Figure 3.7: Second sensitivity analysis - Example 1 with sparse sampling network of 8 equally-spaced measurements, at three different times. Note that the range of the abscissa is different in different subplots.

Repeated sampling in a sparse sampling network

We again use the setup described in Example 1, but take measurements at only eight locations, to aggravate the ill-posedness of the problem (now only 8 measurements are used to estimate 301 points). We sample this plume as it evolves over time. For $T = 330$ (Figure 3.7a), this sparse network results in poorer best estimates and wider confidence

58

bounds for all methods. Note that the actual plume is still contained within the confidence bounds. At later times, $T = 500$ (Figure 3.7b) and $T = 1500$ (Figure 3.7c), the estimated plumes are closer to the actual plume, because the plume becomes smoother with time. This makes any kind of sparse-network interpolation easier, because these methods tend by their nature to yield smooth estimates. For $T = 500$, the KT confidence intervals are similar to those for $T = 330$, but for IFM and TrEK they decrease considerably. This occurs because the transport information yields more precise estimates as time lapses, because the effects of additional mixing mitigate the effects of the uncertain release history. At $T = 1500$, the plume is longer and the measurement network spans a larger area. Therefore, the distance between adjacent measurement locations is larger, which increases the KT uncertainty. The IFM and TrEK confidence intervals, on the other hand, are narrower owing to the degree of mixing implied by the flow and transport model. After a long time, the actual plume becomes very smooth as a result of prolonged dispersion, and the three best estimates reproduce it very well and are almost indistinguishable (Figure 3.7c). At this large time, the accuracy of the kriging estimate is improved due to the smoothness of the plume, whereas the IFM and TrEK estimates take advantage of this feature as well as the diffusive nature of the transport process.

Correlated model-data mismatch error

If transport errors are present, then measurements cannot be reproduced perfectly, and this error is often parameterized as an additional measurement error. Recall that the first component of the inverse model objective function (eq. (6)) requires the forward model $\mathbf{H}^*\mathbf{s}$ to reproduce the measurements $\mathbf{z}^*$ to within an error with covariance $\mathbf{R}$. Previously,

59

this covariance matrix was modeled as a diagonal matrix, representing only independent measurement errors, with no cross-correlated errors. Non-zero off-diagonal terms in **R** imply that errors at certain points affect the values at other points. These effects can be interpreted as spatially-correlated transport errors; hence we repeat Example 2, replacing the uncorrelated **R** matrix with a cubic GCF with $\theta = 10^{-10}$, to investigate the effect of spatially-correlated transport errors on the performance of IFM and TrEK.

Figure 3.8 shows the results of this sensitivity analysis. As expected, the best estimates are less precise than in Experiment 2, yielding a smoother estimated plume distribution with higher uncertainty. The estimates remain accurate, however, providing a realistic estimate of the uncertainty associated with the plume distribution.
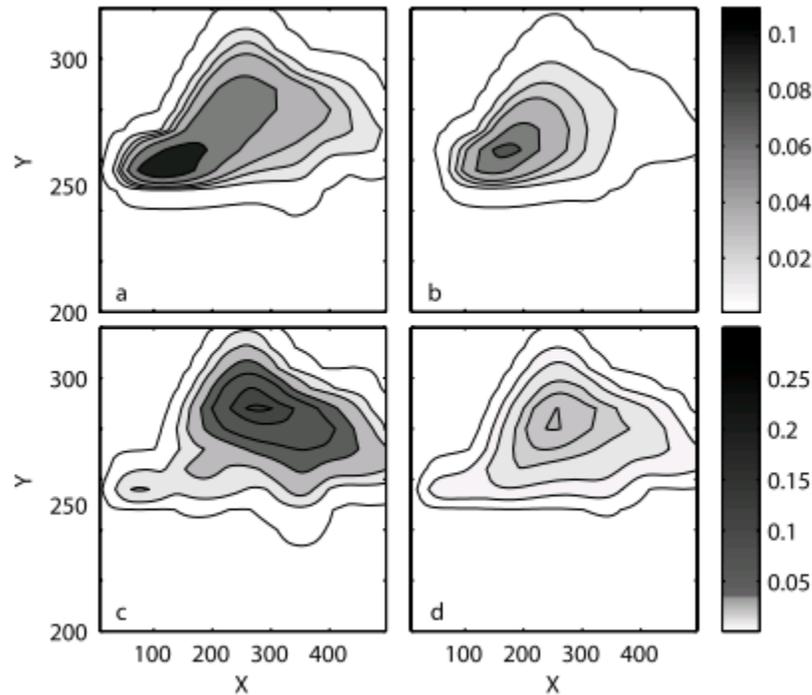
Figure 3.8: Third sensitivity analysis - Example 2 with spatially-correlated model-mismatch error matrix R. (a) Inverse/forward modeling best estimate, (b) Transport-enhanced kriging best estimate, (c) Inverse/forward modeling uncertainty, (d) Transport-enhanced kriging uncertainty. All quantities are dimensionless; uncertainties represented by standard deviation.

5. Discussion

In the presented examples, the IFM and TrEK methods are shown to be robust, providing good estimates under a variety of conditions. Kriging with a trend, in contrast, is strongly dependent on measurement locations. The main reason behind this is that the new methods take advantage of transport information within the aquifer. IFM and TrEK only allow solutions which are feasible in terms of the physical transport process. This is especially advantageous in heterogeneous formations (Example 2), in which the measurement-based covariance function cannot capture small-scale variability. In KT, the estimate reverts to an estimated trend away from measurement locations. It is intuitive

and well-known that for KT and other direct interpolation methods, the uncertainty is lowest at the measurement locations, and increases with distance from them. In contrast, for IFM the highest uncertainties can sometimes be found directly downgradient from the source location (see Example 2), because that is where any uncertainty in the recovered source release history has the most dramatic effect on concentrations. In the transverse direction, the contaminant concentration approaches zero with increasing distance, with little sensitivity to the source magnitude. The consequence is that if monitoring wells were to be drilled around the expected "hot spots" of the contamination, KT and IFM would complement each other: under certain circumstances, the former would have more precise predictions in areas of high concentrations, while the latter would more accurately predict the concentrations farther away. Thus, careful implementation of the TrEK method takes advantage of the best features of the other two approaches.

Interestingly, although a sparse measurement network can lead to high uncertainty in the recovered contaminant release history, the process of forward-modeling often mitigates this problem (Figure 3.5). This indicates that, although the sparse measurements are not sufficient to recover the source in some cases, the information provided by the flow and transport model provides an additional constraint that results in precise estimates of the current contaminant distribution. The flow regime in an aquifer and the contaminant source location provide a strong constraint on the possible distributions of plumes emanating from the source. In addition, groundwater contaminant transport is a dispersive process and consequently there are limits to what can be learned about the history of contamination from measurements of a plume's present spatial distribution

[Skaggs and Kabala, 1998]. However, this dispersion is also the reason that the results of the forward model are relatively insensitive to the details of the release history, i.e. for a given transport model, large variations in the release would cause relatively smaller variations in the plume. Thus, the additional constraint on the plume comes not only from the estimated release history, but also from the transport information itself. As a result, plumes can be estimated accurately even when the measurement network is not sufficient to strongly constrain the contaminant release history.

These examples used deterministic transport models, with no inherent uncertainty. Other sources of uncertainty, such as the sparseness of the observation network, measurement errors, small-scale variations in concentrations, and model-data mismatch contributed to the uncertainty associated with the recovered source release history and plume distribution (Figures 3.2, 3.6, 3.7a). In field situations, incomplete characterization of the subsurface would lead to additional uncertainty in the flow and transport model. This uncertainty would need to be quantified and included in the IFM and TrEK plume distribution estimates, such that the estimated plume reflected the information content of the available model and measurements. In this Chapter, the uncertainty associated with the flow-and-transport model is included in the variance (or covariance in the third sensitivity analysis) specified in the model-data mismatch matrix **R**.

6. Conclusions

The methods presented in this Chapter combine geostatistical kriging and inverse modeling approaches in developing improved tools for estimating the spatial distribution

of contaminant plumes. The methods reconstruct the plume accurately with low uncertainty. In addition, IFM and TrEK can often alleviate the effects of measurement errors by using transport information to eliminate values that are not physically feasible. The TrEK method is most robust because it assimilates information about the transport as well as the spatial covariance structure of the measured concentrations. In cases where such a covariance structure is not available or not credible (e.g. measurements were taken in different hydrogeological zones), IFM would be the preferred choice.

Several assumptions which were made in this Chapter to more easily focus on developing the new plume estimation methods, can be relaxed under certain conditions. First, while the presented sample applications all assumed a single time-varying point source of contamination, the methods can easily be adapted to multiple point sources [e.g. Butera and Tanda, 2003], or to historical plume distributions [e.g. Michalak and Kitanidis, 2004a], as is done in Chapter 4. Second, it was assumed that advection and dispersion are the only active processes in the transport model $\hat{\mathbf{z}} = \mathbf{H}\hat{\mathbf{s}}$, making it linear with respect to the release $\mathbf{s}$. This formulation could also accommodate first-order reactions. For higher-order/nonlinear reactions, a quasi-linear approximation could be used under certain conditions, detailed by Kitanidis [1995].

It is interesting to note that, although we do not enforce nonnegativity in the estimated plume distribution, the IFM and TrEK methods seldom yield estimates with negative concentrations, even though the estimated source function does sometimes exhibit negative concentrations. Nonnegativity-enforcing constraints could be explicitly imposed

64

[e.g. Snodgrass and Kitanidis, 1997; Michalak and Kitanidis, 2003, 2005] to further improve plume interpolation using these methods. Such a modification would be expected to have a modest effect on the best estimates, but could have a more significant impact on the estimated uncertainty.

We presented this work in a geostatistical framework, but the developed principles could potentially also be applied to other inversion methods. For example, the inversion of the release history could theoretically be performed using the Minimum Relative Entropy approach [Woodbury and Ulrych, 1996], with the rest of the method unchanged, as long as the full covariance of the source release history were calculated.

Finally, in the examples presented here, a deterministic transport field was used, and the location of the source was known. In field situations, these parameters may be uncertain, which would result in higher uncertainty in the estimated plume distribution, as demonstrated in the third sensitivity analysis. Moreover, some parameters may be unknown, which would require estimating them before, or while, estimating the plume. The next Chapter focuses on incorporating these additional sources of uncertainty within the IFM framework to further enhance the applicability of this method in field situations.

Chapter 4
Simultaneous Estimation of Hydraulic Conductivity and Contaminant Concentration Using Geostatistical Inverse/Forward Modeling

1. Introduction

In the preceding Chapter, two plume estimation methods were developed that integrate flow and transport information into the process of plume interpolation in a geostatistical framework. While these methods possess many advantages, one limitation mentioned is the requirement to have a full flow-and-transport model, which is not realistic in a field setting, because many parameters cannot be directly measured throughout the aquifer. These parameters must somehow be estimated, a process which introduces uncertainty into the flow-and-transport model.

In this Chapter, Inverse/Forward Modeling, which was developed in Chapter 3, is generalized to simultaneously estimate some of those missing parameters. The objective is thus to explore and develop a method of incorporating transport model uncertainty into IFM. More specifically, the method proposed in this Chapter extends the approaches presented in Chapter 3 to the case where flow and transport information is not perfectly known, which is the case most of interest for field applications. While the forthcoming analysis could be applied to a variety of flow and transport parameters, the examples in

66

this Chapter focus on estimating the hydraulic conductivity distribution $\mathbf{K}$. Regarding the transport model, this Chapter will demonstrate the ability of IFM and TrEK to use a historical plume as an inverse model (instead of a known point source). This capability has an important practical implications, because very often the source(s) of a plume cannot be located, or cannot be modeled because lack of geostatistical structure.

There are many sources of uncertainty in kriging interpolation and in other plume estimation methods. First, the input data (e.g. concentrations) often have measurement error. Second, the sparsity of data often encountered in subsurface hydrology leads to many possible solutions, which are all consistent with the data. Since there is only one true solution among these, this contributes to the uncertainty. Finally, there may be errors in the physical or statistical model used to represent the problem, or in its parameters. In its basic forms, kriging accounts for both the first and second sources of uncertainty through the estimation variance. However, model parameter errors are not routinely considered (they do not form a part of the commonly-used estimation variance).

Similarly, in geostatistical inverse modeling of solute concentrations [e.g. Snodgrass and Kitanidis, 1997], these two sources of uncertainty are considered, but model uncertainties are not always directly taken into account. In fact, in IFM and in TrEK, as well as in other geostatistical inverse-modeling applications, the transport model $\mathbf{h}(\mathbf{s})$ and the associated sensitivity matrix $\mathbf{H}$ are often assumed to be fully known. While this assumption aids in solving theoretical problems, it is not very realistic for field applications. The model-data mismatch covariance, often denoted $\mathbf{R}$, encompasses both

67

the measurement error associated with collecting the data and any random numerical or conceptual inaccuracies associated with the evaluation of the response function $\mathbf{h}(\mathbf{s})$. The zero-bias error term $\boldsymbol{\varepsilon}$ (whose variance-covariance matrix is $\mathbf{R}$) is routinely used to represent the aggregate effect of all errors and uncertainties:

$$\mathbf{z} = \mathbf{h}(\mathbf{s}) + \boldsymbol{\varepsilon} \tag{26}$$

and can represent errors from both $\mathbf{z}$ (i.e. measurement errors) and $\mathbf{h}(\mathbf{s})$ (model errors). The covariance function $\mathbf{R}$ is most commonly modeled as a diagonal matrix, i.e.

$$\mathbf{R} = \sigma_R^2 \mathbf{I} \tag{27}$$

where σ_R^2 is the variance of the data-model mismatch error and $\mathbf{I}$ is an $n \times n$ identity matrix (n is the number of data, or observations). This implies that errors are not spatially correlated, which is not always an optimal assumption for errors or uncertainties in the model $\mathbf{h}(\mathbf{s})$. Alternative methods which more directly address uncertainty associated with model parameters may be more efficient. Kitanidis [1995] mentions two such alternatives, but does not carry out the algebra into full detail. The first (restricted likelihood) is to average the conditional probability of the observations over all possible values of the parameters, as is done with the drift coefficients $\boldsymbol{\beta}$. The second (maximum likelihood) is to minimize the objective function with respect to the parameters, along with the covariance parameters $\boldsymbol{\theta}$. However, Kitanidis [1995] advises the use of only few such parameters ("far fewer than the observations"), so the treatment he suggests for

them is not applicable to a full hydraulic conductivity field distribution, as in the applications presented later in this Chapter.

Butera and Tanda [2003] assess the sensitivity of the geostatistical inverse modeling approach to erroneous flow and transport parameters, and conclude that this uncertainty has a significant impact on the contaminant source estimates. However, they do not suggest a way to incorporate uncertain transport parameters in their analysis. Nowak and Cirpka [2004] generalize the geostatistical inverse modeling approach to handle uncertain trend coefficients ($\boldsymbol{\beta}$ in their notation and in the notation of this work), a generalization that could also hold for other kriging and cokriging applications. Nowak and Cirpka [2006a] extend the approach to include the generalized cases of uncertain prior knowledge about structural parameters (i.e. covariance and generalized covariance function parameters). Sun et al. [2006] present a generic framework for inverse modeling, using a Constrained Robust Least Squares estimator. While their method provides near-optimal solutions to source identification problems under model uncertainty, it is deterministic and thus does not yield estimates of the uncertainty associated with inferred sources. This poses several difficulties, mainly the inability to quantify the precision of the obtained solution, and thus assess its quality. Moreover, such inverse-modeling approaches cannot be used with IFM and TrEK, which require a variance-covariance matrix for the forward-modeling step.

To the best of our knowledge, uncertainty in the transfer function $\mathbf{h}(\mathbf{s})$ has not been previously modeled in geostatistical inverse modeling applications to solute transport.

For the contaminant source identification problem, model errors can be caused by oversimplified model structure, inexact model parameters, and numerical errors (e.g., the errors caused by numerical dispersion or numerical discretization). Sun et al. [2006] observe that in many studies, the effect of model error was only evaluated through ad hoc sensitivity analyses.

The objective of the work reported in this Chapter is to enable the application of IFM and TrEK in cases where the transport model is not fully known, and has parameters that still require estimation. A method is presented that can integrate several sources of data as input, and estimate more than one unknown quantity simultaneously. The method combines quasi-linear geostatistical inversing [Kitanidis, 1995] for recovering a set of parameters (hydraulic conductivity in the examples below), linear inversing [Snodgrass and Kitanidis, 1997] for recovering the history of contamination, and Inverse/Forward Modeling (IFM) [Shlomi and Michalak, 2007] for estimating the spatial distribution of a contaminant plume. Specifically, measurements and expected correlations of concentration, hydraulic head, and hydraulic conductivity are used along with a flow-and-transport model to infer both hydraulic conductivity, $\mathbf{K}$, and the extent of the present-day contaminant plume throughout the study area, $\mathbf{s}$.

2. Methodology

The plume estimation method developed in this Chapter could be seen as a generalized nonlinear variant of IFM, which was described in the previous Chapter. IFM consists of several steps:

1. A transport model is used to construct a Jacobian matrix $\mathbf{H}$ whose elements correspond to the sensitivities of $n \times 1$ measurements $\mathbf{z}$ (typically, concentration measurements at the current or previous times) to an earlier state of the system $\mathbf{s}$ ($m_s \times 1$). The state of the system could be, for example, a temporal release history of a contaminant from a point source [Shlomi and Michalak, 2007], or the spatial distribution of a contaminant plume [Michalak and Shlomi, 2007].

2. Geostatistical inverse-modeling is performed to estimate/infer $\mathbf{s}$, using the Jacobian matrix $\mathbf{H}$ and plume concentration measurements $\mathbf{z}$.

3. A second Jacobian matrix $\tilde{\mathbf{H}}$ is constructed and used to forward-model the earlier state $\mathbf{s}$ to the estimation space $\mathbf{C}_e$, resulting in a plume estimate $\hat{\mathbf{C}}$. The elements of $\tilde{\mathbf{H}}$ correspond to the sensitivities of the estimation space $\mathbf{C}_e$ to the earlier state $\mathbf{s}$. The uncertainty $\mathbf{V}_{\hat{\mathbf{s}}}$ associated with $\hat{\mathbf{s}}$ is also propagated to the plume space, defining a corresponding plume uncertainty.

Shlomi and Michalak [2007] assumed that the transport model was fully known, and only solved for the states of the system, i.e. the historical and present concentration distribution of the contaminant. In this Chapter, the method is generalized to

71

accommodate missing or incomplete transport parameters $\mathbf{r}$ ($m_r \times 1$), which are estimated simultaneously. Here parameters, as opposed to state, pertain to quantities which can be assumed not to change over time as the transport model is used, such as hydraulic conductivity, porosity, or dispersivity. More specifically in the applications below, the hydraulic conductivity $\mathbf{K}$ is estimated along with contaminant plume concentration. The changes in each of the steps of IFM mentioned above, from the previous Chapter to this, can be summarized as follows:

1. The Jacobian matrix $\mathbf{H}$ now contains additional components with elements that correspond to the sensitivities of observations to the missing parameters $\mathbf{r}$. Since $\mathbf{r}$ is estimated simultaneously with $\mathbf{s}$, a new $\mathbf{H}$ (which is a function of $\mathbf{r}$) has to be created at each iteration.

2. Inverse-modeling is carried out iteratively (due to the nonlinearity of the function relating the estimated parameters to the available observations) to estimate both the earlier state $\mathbf{s}$ and the missing parameters $\mathbf{r}$, rather than in one step to estimate only $\mathbf{s}$.

3. The forward-modeling step is performed only after the iterations above have converged. A new Jacobian matrix $\tilde{\mathbf{H}}$ can only be formed at this point (using the final estimate $\hat{\mathbf{r}}$ of the missing parameter from the previous step), prior to forward-modeling of the earlier state $\mathbf{s}$.

2.1 Step 1: Construction of Jacobian matrices

Shlomi and Michalak [2007] and Michalak and Shlomi [2007] were able to construct both of their Jacobian matrices ($\mathbf{H}$ and $\tilde{\mathbf{H}}$) before performing the next steps, because the transport model was fully known. However, this cannot be done when one or more parameters $\mathbf{r}$ are unknown. Because the relationship between flow and transport parameters $\mathbf{r}$ to available observations is typically nonlinear, the system must be solved iteratively, with a new Jacobian being defined at each iteration, using the most recent estimate of $\mathbf{r}$.

Note that the Jacobian matrices can be composed of several parts, corresponding to the various types or times of measurements, e.g.

$$\mathbf{H}_s = \begin{bmatrix} \mathbf{H}_s^{C_1} \\ \mathbf{H}_s^{C_2} \\ \mathbf{H}_s^{K} \\ \mathbf{H}_s^{\Phi} \end{bmatrix} = \begin{bmatrix} \dfrac{\partial \mathbf{C}(t=1)}{\partial \mathbf{s}} \\ \dfrac{\partial \mathbf{C}(t=2)}{\partial \mathbf{s}} \\ \dfrac{\partial \log[\mathbf{K}]}{\partial \mathbf{s}} \\ \dfrac{\partial \boldsymbol{\Phi}}{\partial \mathbf{s}} \end{bmatrix} ; \quad \mathbf{H}_r = \begin{bmatrix} \mathbf{H}_r^{C_1} \\ \mathbf{H}_r^{C_2} \\ \mathbf{H}_r^{K} \\ \mathbf{H}_r^{\Phi} \end{bmatrix} = \begin{bmatrix} \dfrac{\partial \mathbf{C}(t=1)}{\partial \mathbf{r}} \\ \dfrac{\partial \mathbf{C}(t=2)}{\partial \mathbf{r}} \\ \dfrac{\partial \log[\mathbf{K}]}{\partial \mathbf{r}} \\ \dfrac{\partial \boldsymbol{\Phi}}{\partial \mathbf{r}} \end{bmatrix} \tag{28}$$

corresponds to the case in which there are measurements of concentrations C (at two different times t = 1 and t = 2), log-conductivity $\log[\mathbf{K}]$, and head Φ.

2.2 Step 2: Simultaneous quasi-linear inverse modeling

After the Jacobian matrix **H** is constructed at the previous step, inverse modeling is performed to simultaneously estimate the previous state of the system **s** and the missing parameters **r**. As opposed to Shlomi and Michalak [2007] and Michalak and Shlomi [2007], this study uses a quasi-linear variant of geostatistical inverse modeling [Kitanidis, 1995] described below.

Geostatistical inverse modeling is based on Bayes' Theorem, which, for the case examined here, can be written as:

$$p''(\mathbf{s},\mathbf{r}\,|\,\mathbf{z}) = \frac{p(\mathbf{z}\,|\,\mathbf{s},\mathbf{r})\,p'(\mathbf{s},\mathbf{r})}{p(\mathbf{z})} \tag{29}$$

where **z** is a vector of available measurements or observations, such as contaminant concentrations, hydraulic head, and hydraulic conductivity. The right-hand-side has three components:

1. $p(\mathbf{z}|\mathbf{s}, \mathbf{r})$ is the likelihood of the observations for given state and parameter vectors, represented by a multivariate normal distribution with mean **h(s,r)** and covariance matrix **R**:

$$p(\mathbf{z}\,|\,\mathbf{s},\mathbf{r}) \propto \exp\left[-\frac{1}{2}(\mathbf{z}-\mathbf{h}(\mathbf{s},\mathbf{r}))^{T}\mathbf{R}^{-1}(\mathbf{z}-\mathbf{h}(\mathbf{s},\mathbf{r}))\right] \tag{30}$$

The model function **h(s,r)**, or "response" [Kitanidis, 1995] maps the state vector **s** using parameter vector **r** onto the observation space **z**, through the use of a numerical or analytical contaminant transport model:

74

$$\mathbf{z} = \mathbf{h}(\mathbf{s},\mathbf{r}) + \boldsymbol{\varepsilon} \qquad (31)$$

where $\boldsymbol{\varepsilon}$ represents an unbiased model-data mismatch error with covariance $\mathbf{R}$. For parameters such as hydraulic conductivity, which are not generally linearly proportional to observations $\mathbf{z}$, this is an iterative process which requires a first-order approximation about a known point $\mathbf{h}(\mathbf{s}_i,\mathbf{r}_i)$ (namely the previous iteration i):

$$\mathbf{h}(\mathbf{s},\mathbf{r}) \simeq \mathbf{h}(\mathbf{s}_i,\mathbf{r}_i) + \left.\frac{\partial \mathbf{h}}{\partial \mathbf{r}}\right|_{\mathbf{r}_i}(\mathbf{r}-\mathbf{r}_i) + \left.\frac{\partial \mathbf{h}}{\partial \mathbf{s}}\right|_{\mathbf{s}_i}(\mathbf{s}-\mathbf{s}_i) \qquad (32)$$

Defining the blocks of the Jacobian matrix

$$\mathbf{H}_s \triangleq \left.\frac{\partial \mathbf{h}}{\partial \mathbf{s}}\right|_{\mathbf{s}_i} \qquad\qquad \mathbf{H}_r \triangleq \left.\frac{\partial \mathbf{h}}{\partial \mathbf{r}}\right|_{\mathbf{r}_i} \qquad (33)$$

the model-data mismatch $\mathbf{z}-\mathbf{h}(\mathbf{s},\mathbf{r})$ at the i^{th} iteration can be expressed as

$$\tilde{\mathbf{z}} - \mathbf{H}_r\mathbf{r}_i - \mathbf{H}_s\mathbf{s}_i \qquad (34)$$

where $\tilde{\mathbf{z}} \triangleq \mathbf{z} - \mathbf{h}(\mathbf{s}_i,\mathbf{r}_i) + \mathbf{H}_r\mathbf{r}_i + \mathbf{H}_s\mathbf{s}_i$, and the likelihood of the measurements $\mathbf{z}$ can now be written as:

$$p(\mathbf{z}\,|\,\mathbf{s},\mathbf{r}) \propto \exp\left[-\frac{1}{2}\left(\tilde{\mathbf{z}}-\mathbf{H}_s\mathbf{s}-\mathbf{H}_r\mathbf{r}\right)^T \mathbf{R}^{-1}\left(\tilde{\mathbf{z}}-\mathbf{H}_s\mathbf{s}-\mathbf{H}_r\mathbf{r}\right)\right] \qquad (35)$$

2. $p'(\mathbf{s})$ is the prior probability density function of the state $\mathbf{s}$, with $E[\mathbf{s}]=\mathbf{X}_s\boldsymbol{\beta}_s$ as the geostatistical model of the mean, and $\mathbf{Q}_s(\boldsymbol{\theta})$ as its $m_s \times m_s$ geostatistical covariance matrix, with parameters $\boldsymbol{\theta}$. This distribution is also modeled as

multivariate Gaussian:

$$p'(\mathbf{s}) \propto \exp\left[-\frac{1}{2}(\mathbf{s} - \mathbf{X}\beta)^T \mathbf{Q}_s^{-1}(\mathbf{s} - \mathbf{X}\beta) \right] \tag{36}$$

Similarly, $p'(\mathbf{r})$ is the prior probability density function of the parameter vector, with $E[\mathbf{r}] = \mathbf{X}_r\beta_r$ as the geostatistical model of the mean, and $\mathbf{Q}_r$ ($m_r \times m_r$) as the covariance matrix. The full joint distribution $p'(\mathbf{s}, \mathbf{r})$ is thus proportional to:

$$\exp\left[-\frac{1}{2}\left(\begin{bmatrix} \mathbf{s} \\ \mathbf{r} \end{bmatrix} - \begin{bmatrix} \mathbf{X}_s\beta_s & \mathbf{0} \\ \mathbf{0} & \mathbf{X}_r\beta_r \end{bmatrix} \right)^T \begin{bmatrix} \mathbf{Q}_s & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_r \end{bmatrix}^{-1} \left(\begin{bmatrix} \mathbf{s} \\ \mathbf{r} \end{bmatrix} - \begin{bmatrix} \mathbf{X}_s\beta_s & \mathbf{0} \\ \mathbf{0} & \mathbf{X}_r\beta_r \end{bmatrix} \right) \right] \tag{37}$$

Note that this setup assumed no a priori covariance between the state $\mathbf{s}$ and the parameters $\mathbf{r}$.

3. The distribution $p(\mathbf{z}) = \iint\limits_{\mathbf{s},\mathbf{r}} p(\mathbf{z} \mid \mathbf{s}, \mathbf{r}) \, p'(\mathbf{s}, \mathbf{r}) \, d\mathbf{s}d\mathbf{r}$ of the observations is not dependent on the state function $\mathbf{s}$ or parameters $\mathbf{r}$ and therefore can be omitted from the minimization of the objective function in the next step.

Maximizing the joint conditional probability $p''(\mathbf{s}, \mathbf{r} \mid \mathbf{z})$ is equivalent to minimizing its negative log, and the objective function can therefore be expressed as

$$\begin{aligned} L_{\mathbf{s},\mathbf{r},\beta_s,\beta_r} &= \left(\tilde{\mathbf{z}} - \begin{bmatrix} \mathbf{H}_s & \mathbf{H}_r \end{bmatrix} \begin{bmatrix} \mathbf{s} \\ \mathbf{r} \end{bmatrix} \right)^T \mathbf{R}^{-1} \left(\tilde{\mathbf{z}} - \begin{bmatrix} \mathbf{H}_s & \mathbf{H}_r \end{bmatrix} \begin{bmatrix} \mathbf{s} \\ \mathbf{r} \end{bmatrix} \right) \\ &\quad + \left(\begin{bmatrix} \mathbf{s} \\ \mathbf{r} \end{bmatrix} - \begin{bmatrix} \mathbf{X}_s\beta_s & \mathbf{0} \\ \mathbf{0} & \mathbf{X}_r\beta_r \end{bmatrix} \right)^T \begin{bmatrix} \mathbf{Q}_s & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_r \end{bmatrix}^{-1} \left(\begin{bmatrix} \mathbf{s} \\ \mathbf{r} \end{bmatrix} - \begin{bmatrix} \mathbf{X}_s\beta_s & \mathbf{0} \\ \mathbf{0} & \mathbf{X}_r\beta_r \end{bmatrix} \right) \end{aligned} \tag{38}$$

The system of linear equations representing the solution of this minimization is a cokriging system with the state vector, parameter vector, and available measurements all being treated as correlated random fields:

$$\begin{bmatrix} \mathbf{HQH}^T + \mathbf{R} & \mathbf{HX} \\ (\mathbf{HX})^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \Lambda^T \\ \mathbf{M} \end{bmatrix} = \begin{bmatrix} \mathbf{HQ} \\ \mathbf{X}^T \end{bmatrix} \tag{39}$$

where

$$\begin{aligned} \mathbf{H} &= \begin{bmatrix} \mathbf{H}_s & \mathbf{H}_r \end{bmatrix} \\ \mathbf{X} &= \begin{bmatrix} \mathbf{X}_s & \mathbf{0} \\ \mathbf{0} & \mathbf{X}_r \end{bmatrix} \\ \mathbf{Q} &= \begin{bmatrix} \mathbf{Q}_s & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_r \end{bmatrix} \\ \Lambda &= \begin{bmatrix} \Lambda_s \\ \Lambda_r \end{bmatrix} \end{aligned} \tag{40}$$

This system is solved at each iteration for all coefficients Λ and Lagrange multipliers $\mathbf{M}$, yielding the following best estimates, which are used as input to the next iteration:

$$\hat{\mathbf{s}}_{i+1} = \Lambda_s \tilde{\mathbf{z}} \qquad \hat{\mathbf{r}}_{i+1} = \Lambda_r \tilde{\mathbf{z}} \tag{41}$$

The variance-covariance matrix of the estimated errors of $\mathbf{s}$ and $\mathbf{r}$ is

$$\mathbf{V}_{i+1} = \mathbf{Q} - \mathbf{QH}_i^T \Lambda^T - \mathbf{XM} \tag{42}$$

2.3 Step 3: Estimating the plume

When these iterations have converged, the plume at the required time can be estimated

using the last step of IFM, i.e. the forward model, which simply maps the estimate of the state at the last iteration $\hat{\mathbf{s}}_{end}$ forward in time to the plume, using the model matrix

$\tilde{\mathbf{H}}_s^{C_e} = \dfrac{\partial \mathbf{C}_e}{\partial \mathbf{s}}$, which is the sensitivity of the plume at the required estimation time $\mathbf{C}_e$ to the

contaminant history $\mathbf{s}$. The resulting best estimate is

$$\hat{\mathbf{C}} = \tilde{\mathbf{H}}_s^{C_e} \hat{\mathbf{s}}_{end} \tag{43}$$

Note that while the $n_{c_i} \times m_s$ model matrix $\mathbf{H}_s^{C_i}$ maps $\mathbf{s}$ to the n_{c_t} concentration

measurements at time t, the $(N \times m_s)$ matrix $\tilde{\mathbf{H}}_s^{C_e}$ maps $\mathbf{s}$ to the entire domain, i.e. to all of

the N points at which we wish to estimate the concentration of the contaminant.

Similarly, the uncertainty associated with the estimated plume is mapped from the $(m_s + m_r) \times (m_s + m_r)$ inversion uncertainty

$$\mathbf{V}_{end} = \begin{bmatrix} \mathbf{V}_{ss} & \mathbf{V}_{sr} \\ \mathbf{V}_{rs} & \mathbf{V}_{rr} \end{bmatrix} \tag{44}$$

as obtained in the last iteration using Equation (42) to the current plume concentration space. Since we are generally interested in mapping the overall variance of the estimated plume distribution, i.e. the uncertainty originating from the inversion of the concentration, as well as the uncertainty originating from the inversion for the hydraulic conductivity, we use the $N \times (m_s + m_r)$ forward model

$$\tilde{\mathbf{H}}^{C_e} = \begin{bmatrix} \tilde{\mathbf{H}}_s^{C_e} & \tilde{\mathbf{H}}_r^{C_e} \end{bmatrix} \tag{45}$$

78

for projecting the uncertainty forward in time, resulting in the $N \times N$ variance-covariance matrix:

$$\mathbf{V}^{C_e} = \tilde{\mathbf{H}}^{C_e} \mathbf{V}_{end} \tilde{\mathbf{H}}^{C_e^T} \tag{46}$$

To inspect the separate individual contributions to the plume uncertainty of the uncertainty in **s** and that in **r**, we can compare the diagonal elements in

$$\mathbf{V}_s^{C_e} = \tilde{\mathbf{H}}_s^{C_e} \mathbf{V}_{ss} \tilde{\mathbf{H}}_s^{C_e^T} \tag{47}$$

and

$$\mathbf{V}_r^{C_e} = \tilde{\mathbf{H}}_r^{C_e} \mathbf{V}_{rr} \tilde{\mathbf{H}}_r^{C_e^T} \tag{48}$$

respectively. Note that while the diagonal elements of both of these expressions should be smaller than those of $\mathbf{V}^{C_e}$, their sum does not generally add up to $\mathbf{V}^{C_e}$, because of the cross-correlations between the **s-** and **r-**uncertainties. This analysis may not be entirely accurate, as this linearized model does not completely represent the uncertainty. However, we can use it as a qualitative tool to assess the sources of uncertainty.

2.4 Sequential estimation of hydraulic conductivity and plume concentration

One alternative to this new method would be a more conventional three-step process. First, kriging or inverse-modeling could be performed to estimate the hydraulic conductivity [e.g. Feehley et al., 2000]; then, one might treat the results of this first step as given, and proceed as in Shlomi and Michalak [2007] or Michalak and Shlomi [2007],

with linear inverse-modeling for finding the historical plume; finally, the present-day plume could be found using forward modeling. However, such an approach might not account for the uncertainty resulting from estimating the hydraulic conductivity, when calculating and presenting the overall uncertainty, because IFM and TrEK, and more generally geostatistical inverse modeling, do not account for transport model errors. This approach is used in the examples below for comparison to the proposed approach.

3. Applications

Two sample applications of the method to a synthetic aquifer are presented. Both represent the estimation of the hydraulic conductivity along with the spatial distribution of a plume using several types of measurements. In the first application, the geostatistical inverse modeling approach presented above is used to find the state of the plume at an earlier time, along with the aquifer's hydraulic conductivity. In the second application, the same types of measurements are used to characterize the temporal release history of a contaminant from a known point source, along with the conductivity. In both cases, after the simultaneous inverse modeling is performed, the forward model (which includes the estimated hydraulic conductivity) is used to project the plume to the present time. For each case, both the best estimate and the estimated uncertainty are presented.

Two hypothetical examples were chosen to illustrate and verify the capabilities of the methods in a setup where the true conductivity and concentration distributions are known. For both examples, we adopt the synthetic two-dimensional confined aquifer previously designed and used by Michalak and Shlomi [2007]. The conductivity field in

this aquifer (Figure 4.1a) was generated by Michalak and Kitanidis [2004a]. The hydraulic head distribution (Figure 4.1b) was obtained using the numerical package MODFLOW-2000 [Hill et al., 2000] on a 20 × 43 cell grid. All four boundary conditions of the flow model's rectangular domain were defined as constant head, inducing an average net flow from northwest to southeast.
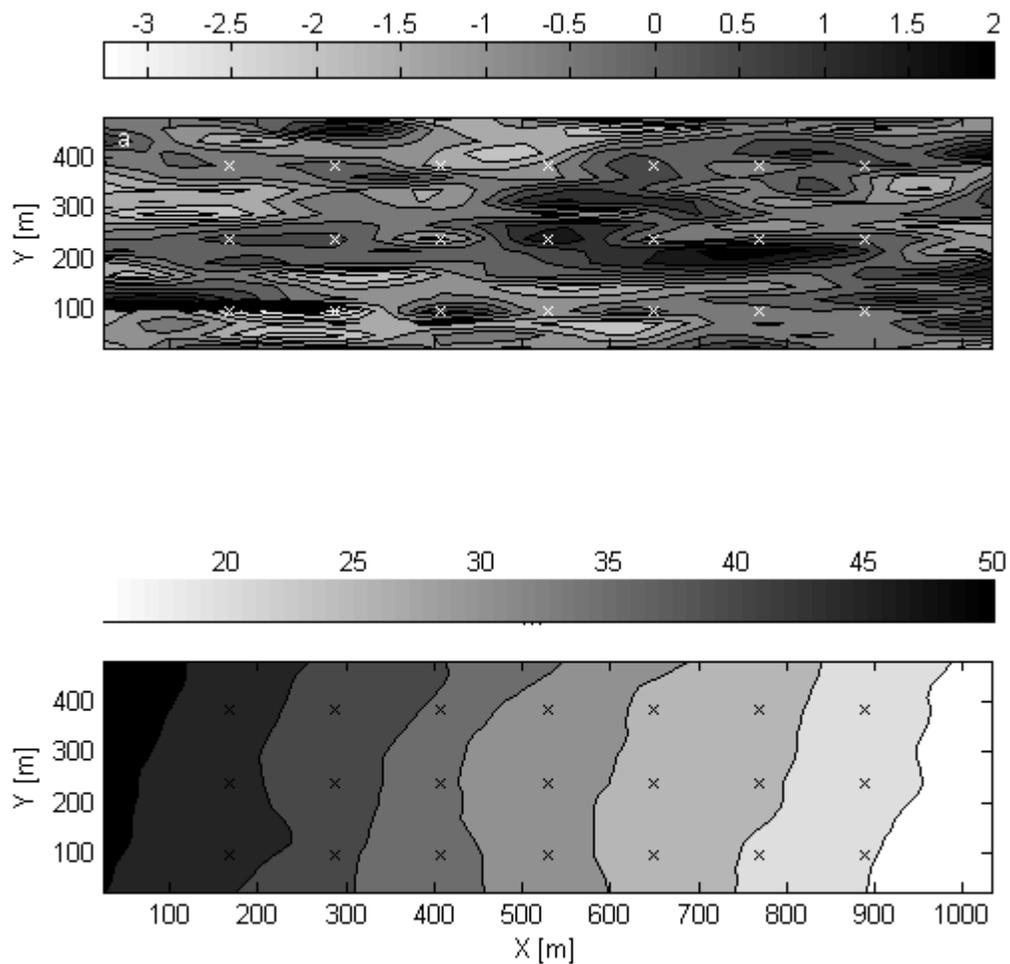
Figure 4.1.   **a) True hydraulic log-conductivity field used for all examples in Chapter 4 [m/s]**
                  b) Resulting hydraulic head throughout the aquifer[m]

The groundwater flow in these examples was assumed steady, and not affected by the contaminant transport. It was assumed that no additional sources of contamination existed, other than the historical plume (in Example 1) and point-source (in Example 2). As opposed to Shlomi and Michalak [2007] and Michalak and Shlomi [2007], the conductivity and head field were assumed to be unknown, except in some cases at several measurement locations. Rather than using the hydraulic conductivity **K** itself, the natural logarithm of the hydraulic conductivity **r** = log[**K**] is generally used for geostatistical estimation because of the empirical evidence that **K** is log-normally distributed [e.g. Freeze, 1975]. **r** was thus treated as a regionalized variable representing the log-hydraulic conductivity distribution in the aquifer. The true log-hydraulic conductivity distribution was numerically generated using a Gaussian covariance model

$$C(h) = \sigma^2 \exp\left(-\left(\frac{h}{l}\right)^2\right)$$
(49)

with sill σ^2 and correlation length l parameters

$$\sigma^2 = 1, \quad l = 100\,m.$$
(50)

with an anisotropy ratio of 5, such that the length parameter along the y-axis was 20 m.

The other unknown for which inverse modeling is performed is **s**, the historical contaminant concentration. More specifically, in the first example, **s** represents the spatial distribution of the plume at an earlier point in time, $T = 0$. This could correspond, for example, to the earliest time at which the existence of the plume was known. In the

second example, **s** represents the concentration of a contaminant release from a known point source as a function of time.

In both applications, we analyze four distinct cases, with measurements of, respectively:

i. plume concentration **C** ($n_c \times 1$), hydraulic conductivity **K** ($n_r \times 1$), and head **Φ** ($n_\phi \times 1$)

ii. plume concentration **C** ($n_c \times 1$) and head **Φ** ($n_\phi \times 1$)

iii. plume concentration **C** ($n_c \times 1$) and hydraulic conductivity **K** ($n_r \times 1$)

iv. **C** ($n_c \times 1$) only, with exhaustive knowledge of the hydraulic conductivity, corresponding to a case where IFM is applied as in Chapter 3, and only the plume concentration is estimated.

3.1 Example 1: Estimation of historical plume and present plume from repeated observations on a fixed grid

The first example involves the estimation of a contaminant plume distribution in a confined aquifer at the most recent of several sampling events, occurring at 600 day intervals (Figure 4.2). At each time, 21 concentration samples were taken on a regular grid of 3 × 7 measurements. Hydraulic head and conductivity (Figure 4.1) were measured at the same locations.
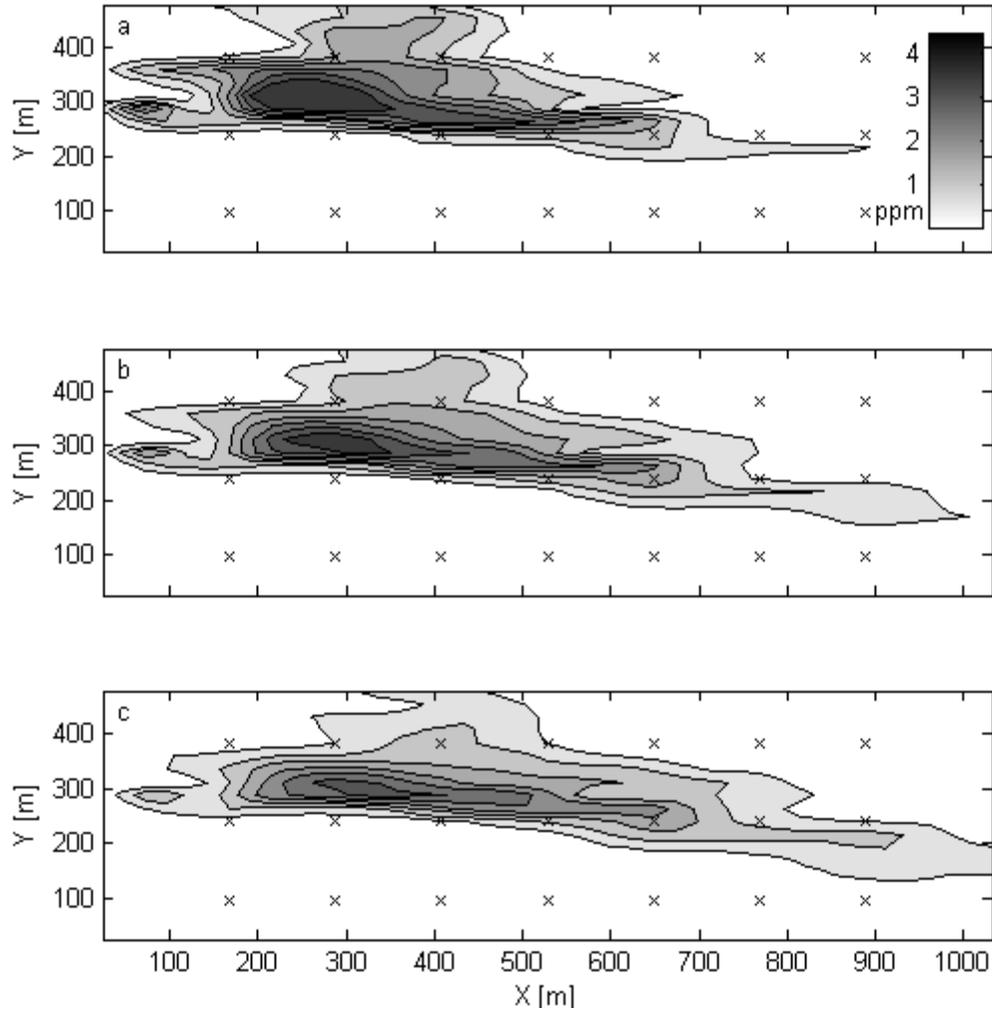
Figure 4.2 Example 1: Spatial and temporal evolution of plume used in sample applications. The black x's indicate sampling locations. Samples were taken at (a) $T = 1200$**, (b)** $T = 1800$**, and (c)** $T = 2400$**. Concentrations are expressed in ppm.**

The actual contaminant plume at $T = 0$ (Figure 4.3), which developed to the plume we are estimating at $T = 2400$, was previously used by Michalak and Shlomi [2007]. Its covariance structure was assumed to follow an exponential covariance function

$$C(h) = \sigma^2 \exp\left(-\frac{h}{l}\right) \qquad (51)$$

with coefficients which were calculated using the full plume distribution and found to be:

84

$$\sigma^2 = 1.49\,ppm^2, \quad l = 84\,m\,. \tag{52}$$

In this example, measurements of all types were assumed to have small independent and identically distributed (i.i.d.) errors drawn from normal distributions with mean zero and variance $10^{-5}\,ppm^2$ for concentration, variance $10^{-2}\,m^2$ for head, and variance 10^{-5} for log-conductivity.

The sensitivity matrix $\mathbf{H}_s$ is generally composed of three distinct components in this example, which are the sensitivities of the various measurement types to the unknown historical concentration $\mathbf{s}$ (throughout the domain), i.e. the plume distribution at $T = 0$. The first of these is

$$\mathbf{H}_s^C = \frac{\partial \mathbf{C}}{\partial \mathbf{s}} \tag{53}$$

the sensitivity of measured concentration to historical concentration. This Jacobian matrix is calculated by repeatedly running the transport model MT3DMS [Zheng and Wang, 1999] and calculating the concentrations resulting from a unit concentration at only one model cell. For $\mathbf{H}_s^C$, concentrations are only calculated at n_c concentration measurement sites; for $\tilde{\mathbf{H}}_s^C$, concentrations are calculated at all N required estimation locations. This step is identical to the implementation of IFM in Chapter 3. The two other components of $\mathbf{H}_s$,

$$\mathbf{H}_s^r = \frac{\partial \mathbf{r}}{\partial \mathbf{s}} \tag{54}$$

and

$$\mathbf{H}_s^\Phi = \frac{\partial \mathbf{\Phi}}{\partial \mathbf{s}} \tag{55}$$

are zero, because the solute was assumed not to affect flow, such that changes in the historical plume **s** would not affect log-conductivity **r** or head **Φ**.

Similarly, the second sensitivity matrix $\mathbf{H}_r$ also has three components. The trivial component pertaining to the measurements of log-conductivity

$$\mathbf{H}_r^r = \frac{\partial \mathbf{r}}{\partial \mathbf{r}} \tag{56}$$

has zeros everywhere except at i^{th} column only of rows i, which correspond to a measurement location. This implies that any change in input log-conductivity in a specific model cell would result in an identical change in the apparent log-conductivity there, and no change at all in other cells.
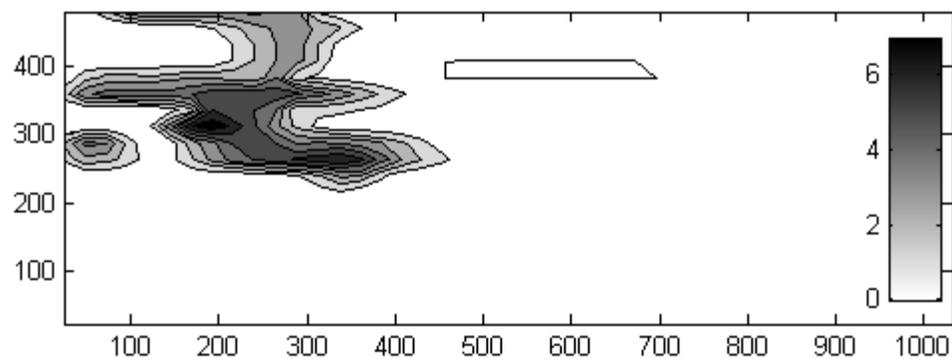
Figure 4.3. Example 1: True contaminant plume at T = 0, as simulated by MT3DMS [ppm].
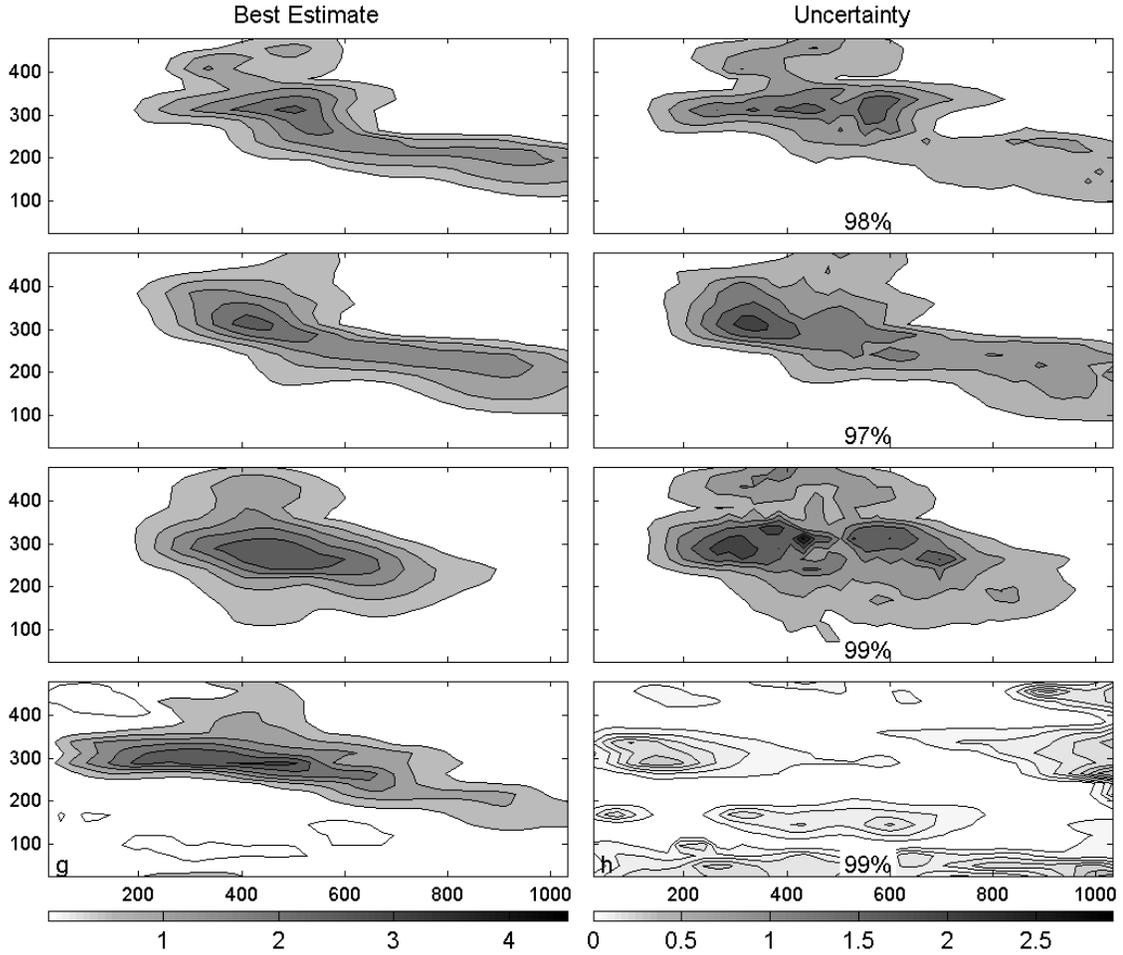
Figure 4.4. Example 1: Recovered plume at T = 2400 days. The left column shows the best estimates, and the right column shows the uncertainties, both in ppm. Concentration measurements taken at three times (T = 1200 days, T = 1800 days, T = 2400 days) were used in all examples. In addition, for the first row (a, b) head and conductivity measurements were taken (case i); for the second (c, d), only head (case ii); for the third (e, f), only conductivity (case iii); for the forth (g, h), the flow field was assumed completely known (inversion was only performed for historical concentration in this case, not conductivity) (case iv). The figures on the right-hand side also display the percentage of estimates which fall within 2σ of the true plume.

The other two components of $\mathbf{H}_r$,

$$\mathbf{H}_r^\Phi = \frac{\partial \mathbf{\Phi}}{\partial \mathbf{r}} \tag{57}$$

and

$$\mathbf{H}_r^C = \frac{\partial \mathbf{C}}{\partial \mathbf{r}}, \tag{58}$$

87

correspond to the sensitivities of the measured head and concentration, respectively, to the log-conductivity. They are calculated numerically, using multiple simulations of the flow and transport models, respectively.

Figure 4.4 shows the results of estimating the contaminant plume for the four cases outlined above. The resulting best estimate resembles the true plume (Figure 4.2c), both in terms of the location of the "hot spot" (where the maximum concentration occurs), and in terms of the plume boundary. The uncertainty tends to be higher where concentrations are higher, as was also described in Shlomi and Michalak [2007] and Chapter 3, due to the relatively high sensitivity of those areas to the historical plume. As opposed to more traditional geostatistical estimation methods, such as the various sorts of kriging, the spatial distribution of the estimation uncertainty does not depend exclusively on the configuration of measurement locations. Equation (46) indicates that the uncertainty of the plume estimate is determined by both the forward model $\tilde{\mathbf{H}}^{C_e}$ and by the uncertainty $\mathbf{V}$ of the inversion (which is itself affected by $\mathbf{H}$). Thus, the primary factor in determining the uncertainty at a location is that specific location's sensitivity to variations in the historical plume $\mathbf{s}$, and in the log-conductivity field $\mathbf{r}$, all of which are inherent in the sensitivity matrix $\mathbf{H}$. The flow in the aquifer was such that the probability of the plume reaching the upper and lower boundaries of the domain was very small. For that reason, the uncertainty was high only in the area of the plume itself – not in the domain boundaries. Conversely, the sensitivities were very high directly downgradient from the plume's previous locations, and variations in the historical plume's estimates have a high impact on the current plume's estimate.

88

The next two panels (Figure 4.4c and 4.4d) were created by applying the method without measurements of hydraulic conductivity (case ii). The best estimate is still close to the true plume, but not as close as for the case where measurements of $\mathbf{K}$ are also available (Figure 4.4a and 4.4b). Since less information is available, the uncertainty is also higher. Several concentration measurements in Figure 4.4b are not reproduced to within the expected model-data mismatch error (expressed by $\mathbf{R}$). A possible explanation is set forth in the discussion of Figure 4.5.

In Figure 4.4e and 4.4f, the plume was estimated using measurements of hydraulic conductivity, but not head (case iii). Here too, the best estimate is not as good as the one using all measurements (Figure 4.4a), but it is still a reasonable estimate of the true plume (Figure 4.2c). The uncertainty seems to be less spread out here, with less high-uncertainty areas on the boundaries than in the previous cases. Since hydraulic head measurements were not available here, only concentration measurements and log-conductivity measurements were used to estimate log-conductivity, with the latter type of measurement being dominant in this process, making the hydraulic conductivity estimate similar to that which would have been obtained using ordinary kriging. To make sure that the best estimate and confidence intervals are realistic, we test how often the best estimate is within 2σ of the true plume concentration, where σ is the standard estimation error (the diagonal of $\mathbf{V}_{end}$). This happens for 97 – 99% of the estimates, which is very close to the ideal 95%.
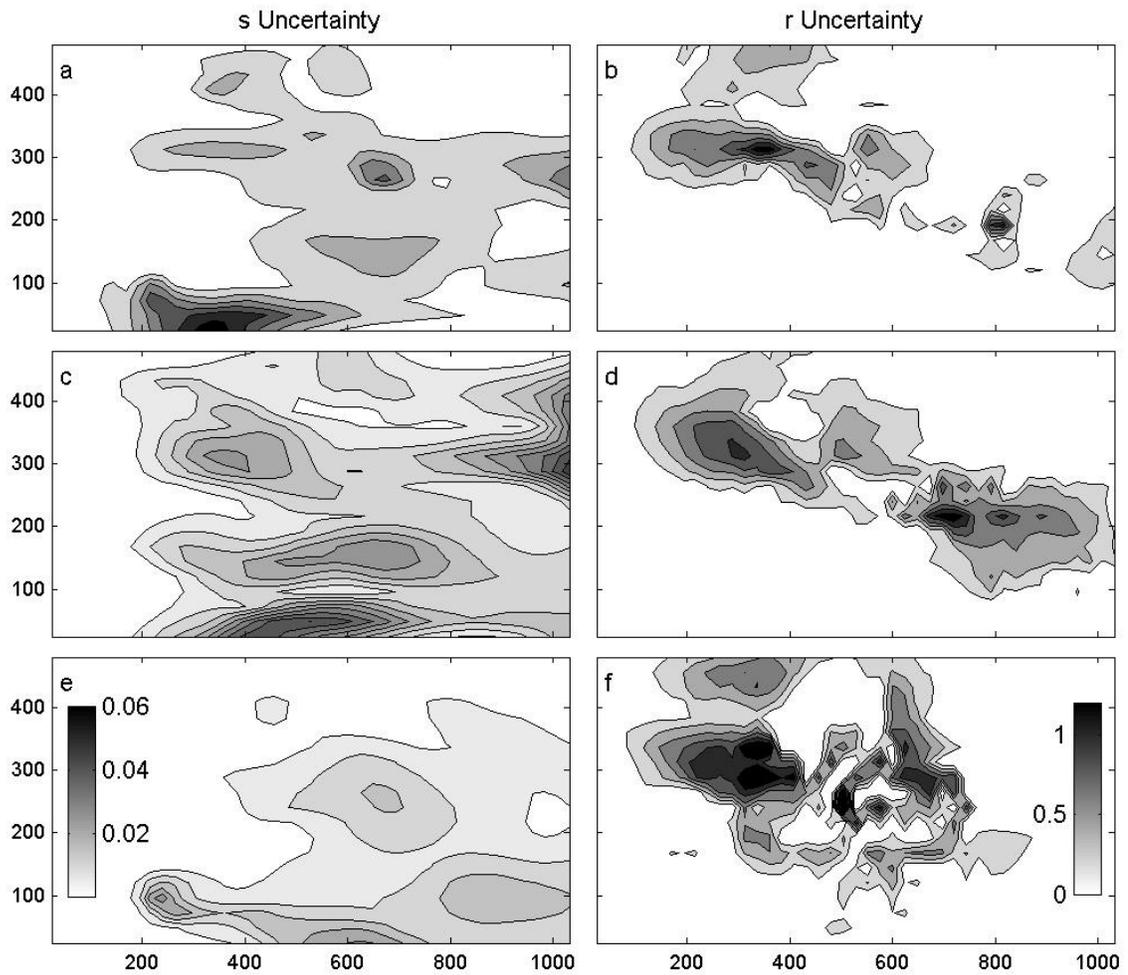
Figure 4.5. Example 1: Individual contributions of uncertainties from the s- and r- components of the inversion to the total plume uncertainties. (a, b) case i, (c, d) case ii, (e, f) case iii. The left column (a, c, e) represents forward projections of the s-component of the inversions and the right column (b, d, f) represents forward projections of the s-component of the inversions. The colorbar in panel e applies also to panels a and c; the colorbar in panel f applies also to panels b and d.

For the sake of comparison to previous methods, we have implemented a linear geostatistical inversion for the concentration only, using the true conductivity field, as in Michalak and Shlomi [2007]. Figures 4.4g and 4.4h (case iv) show the results, which indicate a very good fit of the best estimate, and very low uncertainty. These are due to the abundance of information available in this case: exhaustive conductivity and thus head data are available, and the flow-and-transport models are only missing the initial

90

condition, which is retrieved using the concentration measurements at three points in time. All of this information contributes to constrain the result and reduce the uncertainty. The uncertainty in this case is not only smaller, it is also spatially distributed differently than in the previous cases. When the flow field was unknown, the estimation variance at the plume itself was very high, as was explained before. In this case, while the overall uncertainty is lower, it tends to be higher at the southern and eastern domain boundaries. This shows the large contribution of uncertainty in flow and transport parameters, specifically hydraulic conductivity, to the overall uncertainty.

As mentioned above, we can also analyze the individual contributions of $\mathbf{V}_s$ and $\mathbf{V}_r$ to the total uncertainty (Equations (47) and (48), respectively). Figure 4.5 shows the contribution of uncertainty in $\mathbf{s}$ (first column) and uncertainty in $\mathbf{r}$ (second column) to the overall plume uncertainty. The rows in this figure correspond to cases i - iii. The first general impression from this figure is that the uncertainty associated with historical plume concentration is smoother and more evenly distributed than the uncertainty associated with hydraulic conductivity. The latter appears to be concentrated at the plume, with some peaks of uncertainty occurring occasionally. Elsewhere, the uncertainties from this component seem to be close to zero. The second important observation is that the contributions to the overall uncertainty from the $\mathbf{s}$ component of the inversion (left column) are much lower than the contributions from the $\mathbf{r}$ component of the inversion. This probably has to do with the fact that the flow field serves as a strong constraint on the plume, as seen in Chapter 3. This reinforces the conclusions from Figure 4.4h, which displayed much lower uncertainty as the flow field was completely

91

known. The contribution of the **r** inversion to the plume uncertainty is the source of some large errors at the measurement locations. These errors might be caused by the first-order approximation (Eq. 7) to the model **h**(**s**, **r**). While this approximation is usually a very good one locally (close to the last estimate $\mathbf{r}_i$), it deteriorates as $\|\mathbf{r}-\mathbf{r}_i\|$ increases, which may happen at some iterations. This issue will be further investigated in future work. Finally, note that as these figures represent forward-projections of the uncertainties into the plume space, all contributions to the total uncertainty from the **s** component of the inversion are still lowest at the west boundary, through which it is implicitly assumed that no contaminant originates.

The next set of figures (Figure 4.6) depicts a simpler alternative to this new method, in which the inversions are performed sequentially rather than simultaneously. First, kriging or inverse-modeling is performed to evaluate the hydraulic conductivity; then, the results of this step are used in linear inverse-modeling for finding the historical plume; finally, the present-day plume is found using forward modeling. The elements of the Jacobian matrices used here for the inversion (as two block-column matrices) are the same as those used for the simultaneous estimation (as on matrix with two block-columns), Equations (53) - (58).

The three rows of panels here correspond to cases i – iii, i.e. to the first three rows in Figure 4.4 and to Figure 4.5. These estimates are poorer than the equivalent ones in Figure 4.4, which used simultaneous estimation; while the overall input was identical, in these cases, the log-conductivity was estimated with less data than in the previous cases,

because concentration measurements was not used for the first part, i.e. estimating the log-conductivity. The uncertainties associated with the current plume distribution (column 3) may be lower than those in Figure 4.4, but they represent the projection of the uncertainty in **s** only, and ignore the impact of the uncertainty in the hydraulic conductivity, in effect assuming that the log-conductivity is known. This method does not offer a straightforward way of accounting for the *total* plume uncertainty, from not knowing the log-conductivity and from not knowing the historical concentration.

Thus, although the method used in Figure 4.6 apparently achieved lower estimation variances, they do not realistically represent the estimation variance associated with this method. Moreover, quantitative analysis (detailed results not shown) confirms that the plume best estimates here are inferior to those shown in Figure 4.4, both in terms of the root-mean-square residuals, and in terms of the maximum residual across the domain. This holds true for all three cases i – iii. Overall, this comparison confirms that the proposed simultaneously estimation approach outperforms sequential estimation, which is more akin to existing methods, for this application.
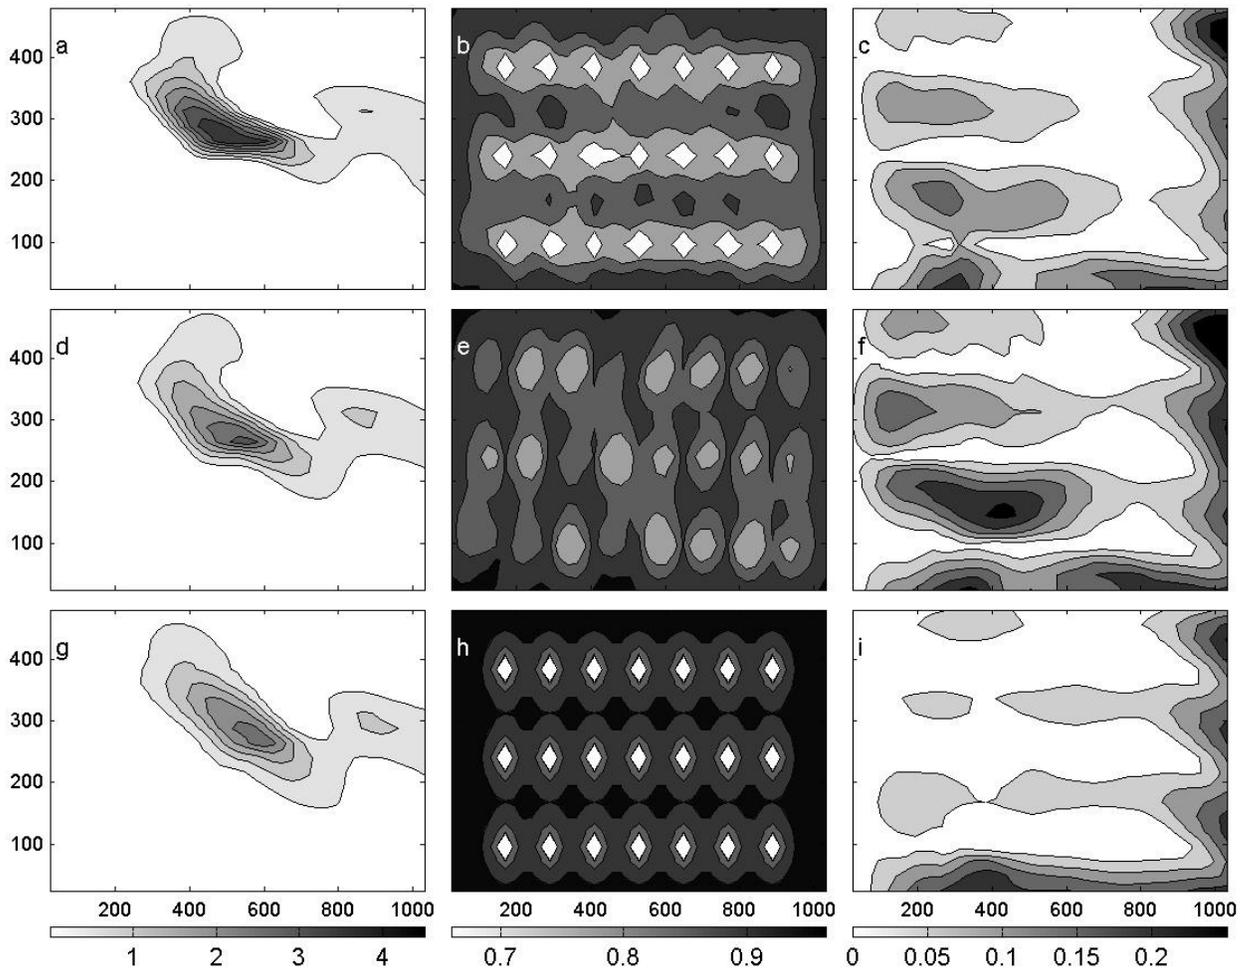
Figure 4.6. Example 1: Sequential estimation. (a - c) case i, (d - f) case ii, (g - i) case iii. The left column (a, d, g) represents the best estimates [ppm], the middle column (b, e, h) represents the log[K] uncertainties [log[m/s]], and the right column (c, f, i) represents the plume concentration uncertainties (not considering log[K] uncertainty) [ppm].

3.2 Example 2: Estimation of temporal release history and present plume from observations on a fixed grid

The Inverse/Forward Modeling approach offers flexibility in terms of the types of measurements that can be used, and state vectors that can be estimated. In Example 1, concentration measurements at three different times were used to recover an earlier

distribution of a plume as part of the estimation process; in Example 2, measurements are only taken at the time at which the plume is to be estimated, and the estimated state represents a temporal release history of contaminant from a single point source with a known location.
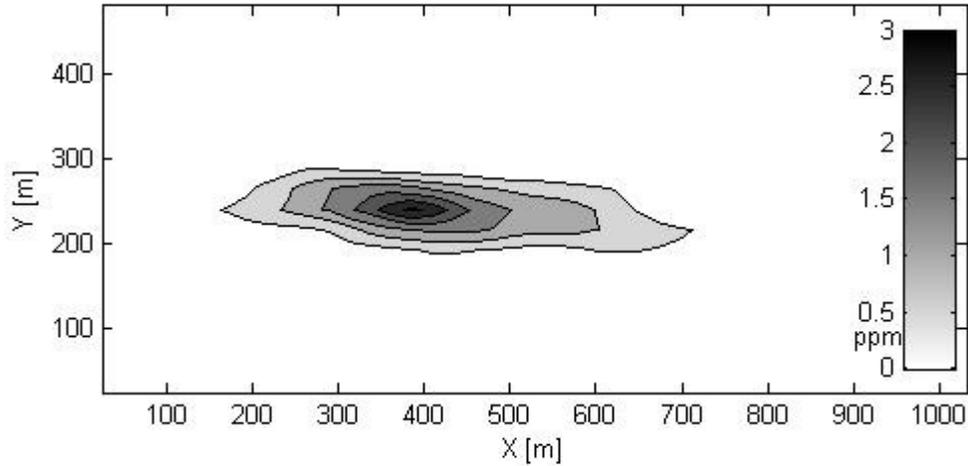
Figure 4.7. Example 2: True contaminant plume distribution, as simulated by MT3DMS.

In this example we use the same aquifer and flow model as above. However, we assume that the plume results from a single point source located at the center of the western boundary, and that this point source was active over some time interval in the past. We use this information to modify the inverse model for **s**, which now represents the *temporal* release history from the point source. We combine this inverse model with the inverse model for log-conductivity.

The actual release history was assumed to be

$$s(t) = 1.4 \exp\left[-\frac{(t-850)^2}{56000} \right] + 1.1 \exp\left[-\frac{(t-1700)^2}{13333} \right] \tag{59}$$

which previously used by Shlomi and Michalak [2007]. We also use the same covariance

structure, a cubic Generalized Covariance Function

$$\mathbf{Q}_s\left(h\right) = \theta h^3 \tag{60}$$

with $\theta = 1.3 \times 10^{-8}$, optimized using the actual release history. Its geostatistical temporal

trend $\mathbf{X}_s \boldsymbol{\beta}_s$ is modeled using a constant component and a linear component.

The equations for this case are identical to those of Example 1. However, the sensitivity

matrix $\mathbf{H}_s$ has a different meaning and is constructed differently. In this case, the Jacobian

matrix $\mathbf{H}_s^C = \dfrac{\partial \mathbf{C}}{\partial \mathbf{s}}$ is calculated by running the transport model (MT3DMS) and observing

the concentrations resulting from a unit release at the point source, at all model cells at

time intervals corresponding to the discretization of the release history. Columns in $\mathbf{H}_s$

correspond to successive points in the temporal release history; rows correspond to

measurement sites (or to estimation locations, in the case of the full $\tilde{\mathbf{H}}_s^{C_e}$).

Since the first step in this example involves recovering the temporal release history,

Figure 4.8 shows the various release histories corresponding to the four examined cases,

and their 95% confidence intervals based on the estimated uncertainty $\mathbf{V}_{ss}$. These cases,

as in Example 1 above, correspond to the use of different sets of measurements (cases i –

iv). The best estimates $\hat{\mathbf{s}}$, for the most part, are qualitatively similar to the true release

history $\mathbf{s}$. However, due to the relatively large measurement errors, they do not replicate

its features (e.g. its two maxima) well. Only the last case, in which the conductivity field

K was fully known, captures the variations in **s**; but even in this case they appear much smoothed out.

Figure 4.8 thus demonstrates the impact of a known conductivity field in a heterogeneous field, and indicates that for inverse modeling, the available measurements may not be sufficient to recover a reasonable representation of the release history of a contaminant for cases where the flow model has high uncertainties. As described below, however, the plume estimates are significantly better than the release estimates.
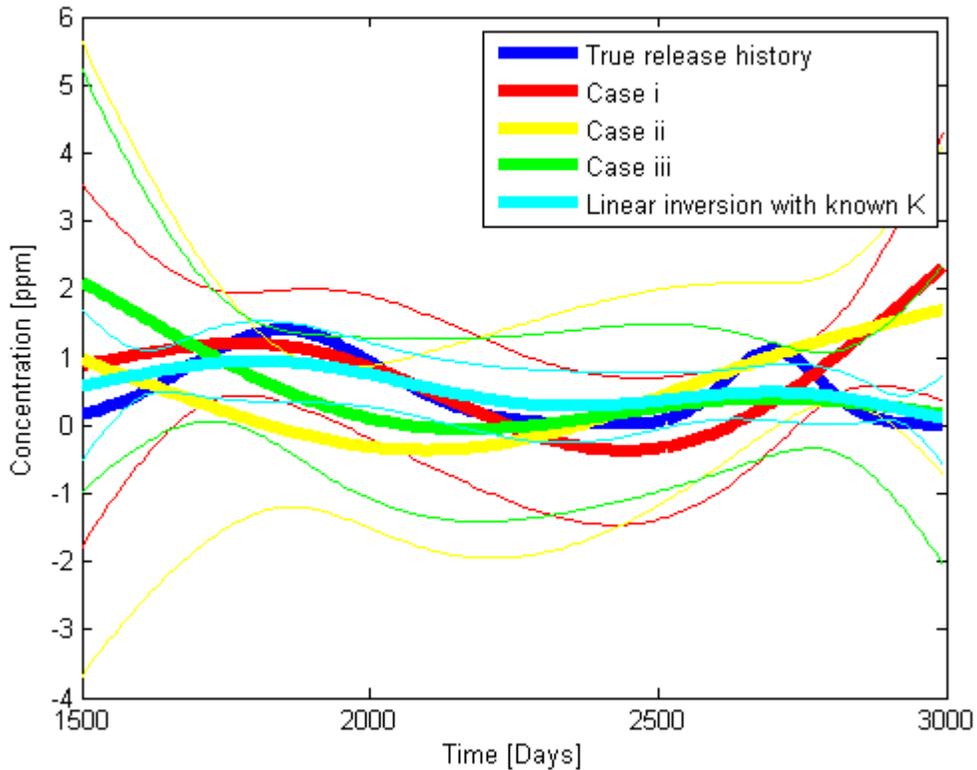
Figure 4.8. Example 2: Temporal release histories, or concentrations at point source, in ppm, as function of time, in days. The thick line of each color represents the best estimate, and the upper and lower thinner lines represent the 95% confidence bounds.

97

It is instructive to see that the width of the confidence intervals for a specific case, which corresponds to the uncertainty, is inversely proportional to the amount of information present. For the linear inversion, where the full conductivity field (and hence flow field) is available, the uncertainty is smallest everywhere. Next is the confidence interval for the case where both $\mathbf{\Phi}$ and $\mathbf{K}$ measurements are available in addition to $\mathbf{C}$. The confidence intervals here at all locations are much wider, because the flow field is only estimated using conductivity and head measurements. Finally, the confidence intervals for the cases in which there is only one type of measurement ($\mathbf{\Phi}$ *or* $\mathbf{K}$) in addition to concentration $\mathbf{C}$, are widest, because there is even less information available for estimating the flow-and-transport model, and consequently the release history.

All confidence intervals are very wide at the beginning of the release history ($T = 1500$), which implies that the available measurements cannot provide adequate information regarding the contaminant release at this early time. Also, the confidence intervals become wider at the end of the modeled release history for a similar reason. Specifically, because there are no measurements very close to the source ($x = 0$), no information can be obtained regarding the release intensity near the end of the estimated period, as this information has not yet reached the measurement locations.

Note that while all of the best estimates of the release history are mostly positive, the lower confidence bounds are mostly below zero. Technically, this means that at each point in the release history, the probability that the release intensity is negative is greater than 2.5%. To relieve this issue, we could constrain the release history to be nonnegative,

for example by using the Gibbs Sampler [e.g. Michalak and Kitanidis, 2003; Michalak, 2008]. However, for the purposes of plume estimation using IFM, there may be no need to constrain nonnegativity in the inverse modeling step: as explained in Chapter 3, this situation is somewhat mitigated after the next step of applying the forward model, as will be shown in the next figures.

Figure 4.9 shows the results of reconstructing the contaminant plume using different combinations of measurements, as in Example 1 (Figure 4.4). Although in this example measurements were taken only at one point in time (the present time), all plume reconstructions are similar to the true plume (Figure 4.7), due to the constraints imposed by the estimated flow and transport models. The most important of these may be the known location of the contaminant release. The pattern of the release history itself had relatively little effect on the plume in this example, as is evident by the variety of release histories (Figure 4.8), corresponding to similar present-day plume (Figure 4.9).

The uncertainties in this example follow a similar pattern to those of the previous example: as more data is available, uncertainties become smaller. Thus, when two types of measurements are available (case i, Figure 4.9b) in additional to concentration measurements, the uncertainty is lower than when only Φ (case ii) or $\mathbf{K}$ (case iii) measurements are available (Figures 4.9d and 4.9f, respectively). When the full transport model is known (case Figure 4.9h), the uncertainty is dramatically lower than in the other cases, as the main source of uncertainty in the other cases – the hydraulic conductivity

field – is completely known in this case. This illustrates again the strong impact that uncertainty in flow and transport parameters has on the overall plume uncertainty.

The estimation here is more conservative than in Example 1, in the sense that 100% of the estimates fall within 2σ of the true plume distribution, as displayed on the right-hand-side column in Figure 4.9.
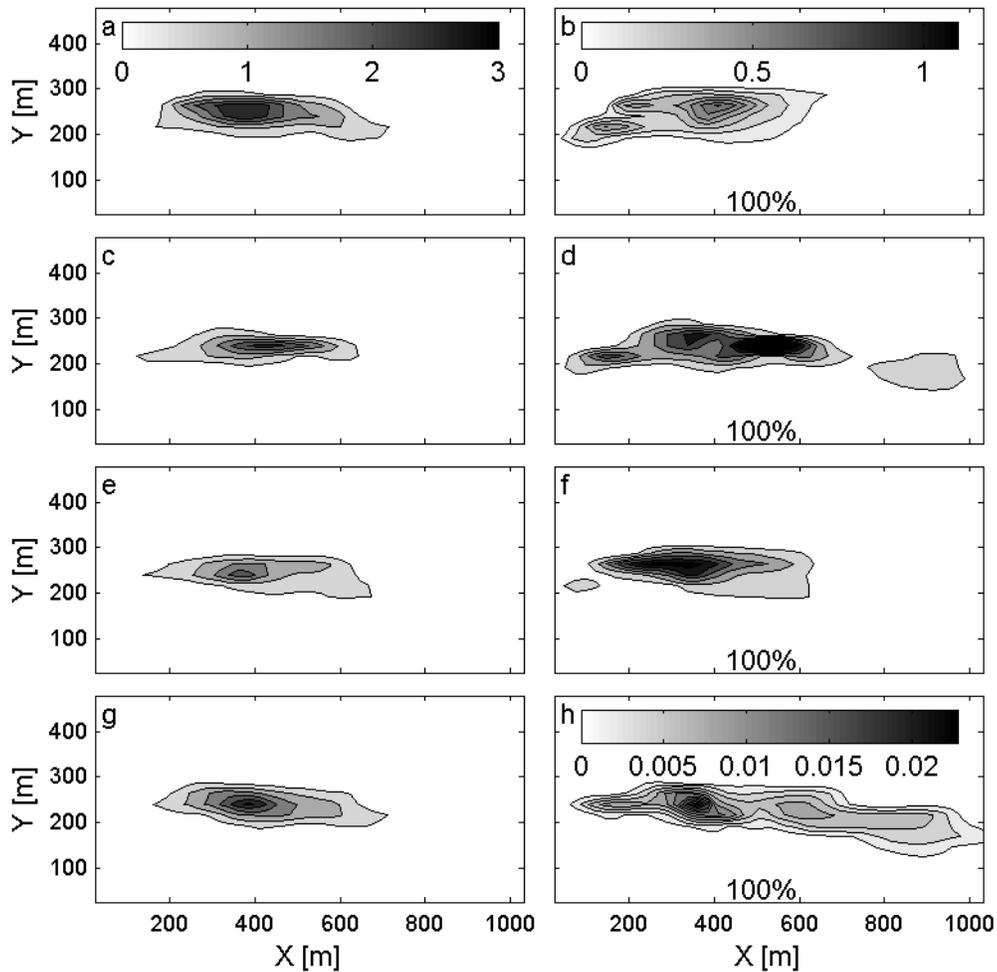
Figure 4.9. Example 2: Recovered contaminant plumes. The left column shows the best estimates, and the right column shows the uncertainties, both in ppm. Rows 1 – 4 correspond to cases i – iv, respectively. The colorbar in panel a also serves panels c, e, and g; the colorbar in panel b also serves panels d and f. Panel h has its own colorbar because the range of uncertainty it represents is much smaller. The figures on the right-hand side also display the percentage of estimates which fall within 2σ of the true plume.

100
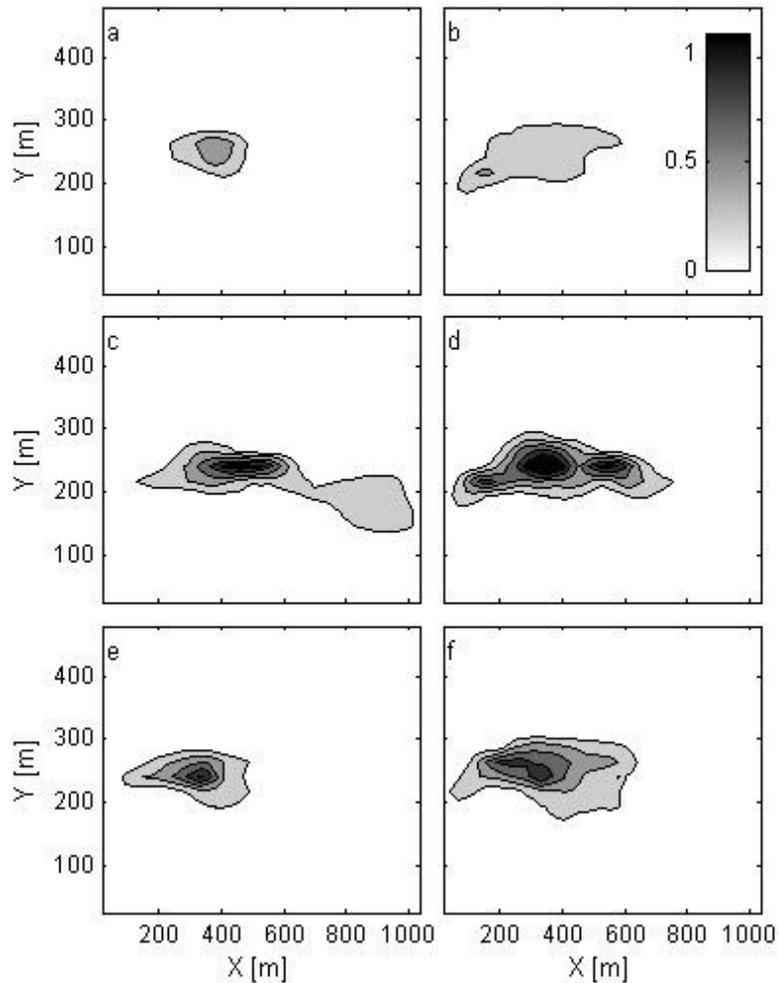
Figure 4.10. Example 2: Individual contributions of uncertainties from the s- and r- components of the inversion to the total plume uncertainties. (a, b) case i, (c, d) case ii, (e, f) case iii. The left column (a, c, e) represents forward projections of the s-component of the inversions and the right column (b, d, f) represents forward projections of the s-component of the inversions. The colorbar in panel b applies to all panels.

While the most important measure may be the total uncertainty, which was displayed in the right column panels of Figure 4.9, it may be instructive to analyze the separate individual contributions of the linear inversion for concentration $\mathbf{C}_0$ and of the nonlinear inversion for log-conductivity log[$\mathbf{K}$] to the total uncertainty. Figure 4.10 represents the

forward projections of these **s** and **r** inversion uncertainties. The left panel corresponds to the uncertainty in the present-day plume caused by uncertainty in the release history **s** (Equation (47)), while the right panel figures correspond to the uncertainty caused by uncertainty in the hydraulic conductivity **r**, calculated by Equation (48). These "partial uncertainties" generally look like the total uncertainties from the right panels of Figure 4.9, with more data translating to lower uncertainties. This is very clear on the left-hand-side panel of Figure 4.10, which corresponds to the uncertainty projected from **s**: the uncertainty in Figure 4.10a is lower than in Figure 4.10c and 4.10e, which only use **Φ** or **K**, respectively (in addition to concentration measurements), but not both. The observation that Figures 4.10c and 4.10d (case ii) display higher uncertainty than Figures 4.10e and 4.10f (case iii), respectively, imply that head observations may provide less information for plume estimation than conductivity observations, in this case. In addition to their high measurement error, a possible explanation may be that as the model used here estimates the hydraulic conductivity, measurements of this field contribute more than measurements of hydraulic head, which is in turn used indirectly to estimate the hydraulic conductivity field.

Another interesting observation from Figure 4.10 is the way in which it is different from its equivalent in Example 1, Figure 4.5. There, the contribution to uncertainty from not knowing **K** was orders of magnitude larger than from not knowing C_0. Here in Example 2, the contributions of these two components seem to be very similar. One reason may be that the previous example had three times as many concentration measurements as this example. However, there seems to be a more dramatic reduction in the contribution of **K**

to uncertainty, than increase in the contribution of **s**. This may be due to the fact that the contaminant release location is known here, reducing the sensitivity of the estimation variance to not knowing the conductivity field. Note again that these figures do not represent the estimation variance of **s** or **r**, but rather the projections of uncertainty (in the inverse models) onto the plume space, or ultimately, their effect on the estimate of the plume.

4. Conclusions

We have generalized Inverse/Forward Modeling, a plume estimation methodology presented in the previous Chapter, and have made it applicable to cases in which the full distribution of flow and transport parameters, such as the hydraulic conductivity field, is unknown. Available data of various types are used to constrain the estimates of both the contaminant plume distribution and the unknown flow and transport parameters. This method allows for the estimation of, for example, concentration and hydraulic conductivity distributions simultaneously, using various types of measurements (e.g. concentration, head, and conductivity) at various times before, during, and also potentially after the time at which the estimate of the plume distribution is required.

In the presented examples, the method developed in this chapter is shown to be robust, providing good estimates using measurements from the concentration, hydraulic head, and/or conductivity fields: the estimation reproduces available measurements, and the estimation variance and root-mean-square residuals are not larger than that of comparable

methods, which estimate hydraulic conductivity and concentration sequentially rather than simultaneously.

An interesting point, previously seen in Shlomi and Michalak [2007], is the high uncertainty associated with the inverse model itself. The problem of reconstructing an earlier plume or a temporal release history, like many other inverse problems, is ill-posed: the solution is not unique, and many alternatives reproduce available data equally well. Hence, the confidence interval on the best estimate is very wide for the inverse problem. In contrast, when that estimate is cast forward, the confidence intervals shrink, because the mixing inherent to the forward problem yields similar plumes for a wide range of historical distributions. Therefore, the uncertainty associated with the estimates of the estimated plume is considerably smaller than that of the historical plume or release history.

Comparing the two examples in this Chapter, it is interesting to note that while the different cases in Example 1 yielded significantly different results, equivalent cases of Example 2 resulted in very similar plumes. Moreover, the uncertainty estimates in Example 2 were also very similar to each other, and even within the projections of the inversion components to the plume uncertainty space (Equations (47) and (48)), the differences were very small (Figure 4.10). This, again, is in contrast to Example 1 (Figure 4.5), in which substantial differences could be seen between the individual contributions of $\mathbf{V}_{ss}$ and $\mathbf{V}_{rr}$ to the total uncertainty. While concentration was measured at three different times in Example 1, and only at one time in Example 2, the latter had the

advantage of having a known release location, which provided a strong constraint on the transport model and the estimate of the plume.

While there are differences within each example as a function of the specific types of data used, all cases confirm that uncertainty in the flow and transport model yields increased uncertainty in the plume estimates relative to the case examined in Chapter 3, as expected. The quality of the best estimates themselves is also lower than in Chapter 3, again reflecting the uncertainty caused by unknown flow and transport parameters. Figure 4.5, which depicts the projection of the uncertainty from the two inversion components into the plume space, quantifies this and indeed shows that the vast majority of the overall uncertainty originates from not knowing the conductivity field. The conclusion from this is that the uncertainty associated with the flow model contributes more to the uncertainty of the final plume estimates relative to the uncertainty caused by the unknown historical plume or release history. This conclusion has implications for groundwater quality monitoring.

That said, the goal of this method is to incorporate available information about both the plume distribution and flow and transport parameters, while accounting for the uncertainty associated with each of these fields. In this way, the approach is a significant improvement over the approach presented in Chapter 3, which required the assumption of a known conductivity field, and other existing interpolation approaches, which ignore available flow and transport information altogether

Future applications may deal with even more unknowns and attempt to estimate them in a similar manner. For example, the distributions of boundary conditions, porosities, or dispersivities could be simultaneously estimated with unknowns such as concentration, to enhance the plume estimate at some point in time. The importance of this process lies in the fact that the many unknowns in groundwater quality monitoring are often difficult to integrate in one consistent estimation framework. Moreover, this type of methodology would offer a degree of flexibility in terms of inclusion or exclusion of various unknown variables.

Finally, the geostatistical framework within which this method was developed could accommodate a large variety of flow-and-transport models, inverse models, and kriging techniques. These individual components should be selected to be appropriate to the application at hand, before merging them into the type of framework presented here. For example, a flow-and-transport model may already exist for an aquifer, or some specific type of kriging may provide good results for a certain type of situation. In many cases, the framework presented in the previous Chapter and in this Chapter can integrate these components to improve on their performance as separate entities.

Chapter 5
Optimal Sampling Well Selection and Monitoring Network Design for IFM/TrEK

1. Introduction

Groundwater quality monitoring consists, in simplistic terms, of

1. designing a monitoring network,

2. periodically selecting wells from this network to be sampled, and finally

3. estimating the concentration of contaminant(s) at each point in the domain.

For any plume estimation technique, measurements of actual contaminant concentrations in the aquifer are a very important component in the method. The location of these measurements is of prime importance, as measurement locations have a dramatic effect on both the best estimates of contaminant concentrations elsewhere in the aquifer, and on the associated estimation variance. Usually, budgetary constraints limit the number of samples that can be taken from the subsurface, further increasing the impact and importance of optimally locating each and every monitoring site.

Over the years, many geostatistical methods have been developed for optimally locating monitoring wells in contaminated aquifers, or in aquifers susceptible to, or suspected of,

contamination. The primary goal of this Chapter is to develop measurement-site selection methods specifically designed for IFM and TrEK. It then serves to demonstrate the ability of IFM and TrEK to easily incorporate existing network design methods that work with established kriging variants.

The importance of this part of the research lies in the fact that well-selection methods are *method-dependent*: that is, for the same aquifer, there is not necessarily only one unique optimal set of measurement locations, which would work equally well for all estimation approaches; rather there is an optimal set for each estimation method. To illustrate this, assume that the selection of wells depends on the estimation variance. The estimation variance is different for different estimation methods, so the chosen wells may also differ. This Chapter presents an adaptation of existing monitoring network design methods to IFM and TrEK, by incorporating IFM or TrEK uncertainty estimates into the design criteria.

2. Methodology

The three components of a network design methodology can be defined as (i) an underlying contaminant plume estimation method (the importance of which was presented above), (ii) design criteria, and (iii) a search technique. *The estimation method* is considered here a component of the overall network design methodology because it is in essence its target, and eventually receives the input of the chosen well selection algorithm: the output of the well-selection algorithm is a set of monitoring locations, where measurements are eventually taken; these measurements are used as input to a

plume estimation method. In this Chapter, the estimation methods we consider are IFM and TrEK, including the variant of IFM presented in Chapter 4, where the hydraulic conductivity field was not known. We also present Kriging with a Trend (KT), again as in Chapter 3, for the sake of comparison.

The design criteria describe what the practitioner seeks in an optimal design, or in other words, what would make such a design better than other candidates. Often, certain constraints are built-in to the design criteria, such as budgetary constraints which may put a maximum on the number of wells from which samples can be taken for analysis. Minimizing the contaminant plume estimation variance is a very common design criterion [Zimmerman, 2006].

Finally, *the search technique* is the method by which the space of all sampling designs $\mathbf{S}$ is systematically examined to find the best design. This must usually be an efficient technique, which is suitable for the design criteria and the nature of the space (e.g. continuous or discrete). As an illustration to the importance of this component, consider for example a domain with 100 potential monitoring wells. In this case, exhaustive searches would have to test 2^{100} designs; even if constrained to select exactly ten wells, $\binom{100}{10} = \dfrac{100!}{10!90!}$ combinations (more than 1.7×10^{13} designs), would have to be tested, a number which is probably not feasible. Despite the speed of today's computers, some sophisticated estimation methods may take a relatively long time to evaluate each design (this may be the case with IFM and TrEK, which work with numerical flow and transport

models). Thus, it is important to choose a proper search algorithm that would reach the optimal network design efficiently.

These three components of the network design methodology are considered in more detail in the following subsections.

2.1 Estimation method and estimation variance

In this work, we focus exclusively on kriging, IFM, and TrEK for interpolation. These geostatistical methods can all produce, in addition to a best estimate, an estimation variance (Eq. 14 and 19 in Chapter 2), which is usually used as a measure of uncertainty. In network design applications, this is a very important quantity, as it is often the critical part of the objective function, which is to be minimized to find the optimal sampling plan in a monitoring network.

In kriging and linear geostatistical inverse-modeling, and therefore also in IFM and TrEK, the formulas for the estimation variance at a certain location do not depend on the value of the measurement there. Thus (for a given covariance model and fixed parameters) the matrices in these equation systems depend only on the distance between measurement points, leading to the result that the kriging, IFM, and TrEK estimation variances will depend solely on the geometry of the sampling network. Consequently, by choosing a given configuration for the sampling sites, the kriging, IFM, and TrEK systems can be solved and therefore the variance of estimation can be calculated using the appropriate equations from Chapter 3. The magnitude of the variance of estimation is

the basis for the criteria to select an optimal monitoring network among a number of possible network layouts.

2.2 Monitoring objectives and design criteria

The optimal design of the groundwater monitoring network often has the objective of finding the distribution of sampling sites which will yield some minimum variance of estimation with minimum monitoring costs. This can be seen as a dual criterion problem because in general, there is a trade-off between these two objectives: minimizing the variance entails taking more measurements, which obviously results in driving up monitoring costs. One way to generate efficient points on the trade-off curve is to minimize one objective subject to a constraint on the maximum value of the other [e.g. Cohon, 2004]. Here, we choose to minimize the variance of estimation subject to an available number of monitoring wells (equivalent to a given budgetary constraint). The variance of estimation is minimized over the whole domain where the estimation is needed. For KT, IFM, and TrEK, the model can therefore be formulated as either

$$\underset{p \in \mathbf{S}}{Minimize}\{\max \mathbf{V}_{\hat{z}}\} \tag{61}$$

or

$$\underset{p \in \mathbf{S}}{Minimize}\{\overline{\mathbf{V}}_{\hat{z}}\} \tag{62}$$

(corresponding to the minimization of maximum or average variance, respectively) where $\mathbf{S}$ denotes the set of all possible sampling plans p, and $\mathbf{V}_{\hat{z}}$ is the estimation variance [cf

Equations 14, 19, and 22 in Chapter 3, corresponding to these methods' variances]. Each sampling plan p is a particular set of monitoring wells. Assuming (for simplicity) that each plan carries the same cost, it is possible to embody this constraint into the set **S**, by only allowing designs with the same number of samples. This assumption is reasonable, as sampling costs often depend on the number of samples, but usually not on which locations are chosen. Zimmerman [2006] notes that these are the two most commonly used design criteria for this type of spatial estimation, i.e. minimizing the *average* kriging variance and minimizing the *maximum* kriging variance. We generalize "kriging variance" to "estimation variance" in the examples below.

The optimization scheme then consists of selecting a given number of monitoring locations from candidate monitoring sites. The total number of monitoring sites to be selected is mainly dictated by the availability of resources. Theoretically, to reach a true optimal configuration of the monitoring network, one has to select the monitoring sites from a two- or three-dimensional *continuous* spatial domain. However, there do not exist analytical techniques for selecting wells in a two-dimensional continuum, even for homogeneous media. Thus, for practical reasons, the monitoring locations are selected from a finite number of prespecified candidate sites. Generally, these candidate monitoring locations are arranged into a regular grid of nodes, if designing a new network, or must be selected from existing monitoring sites, otherwise. As explained above, an important practical consideration in obtaining an optimal monitoring network is the use of the search method for the selection of the optimal sampling sites from the available candidate nodes.

2.3 Search Techniques

For any given configuration of monitoring network layout, the estimation equation system can be solved and the variance of the estimation error can be calculated using the appropriate formulas (given earlier in Chapter 3). To find the best selection of sampling sites, which will yield the lowest values of the estimation variance, a search technique has to be carried out. Two approaches are discussed below: sequential search and simultaneous search.

The sequential search approach consists of dropping or adding a new observation site at each iteration of the optimization process. For example, the optimization process may start by taking into consideration *all* the candidate nodes, and proceed by sequentially eliminating monitoring locations until the allowed number is reached. Several authors [e.g. Loaiciga, 1989; Montasa et al., 2000] have discussed heuristic methods to select wells for elimination. Another way to employ sequential search is to consider only a few nodes initially and then add a new observation node at each stage of the optimization process. Unlike the previous case, as a new node is added to the network, a better estimate of the contaminant distribution will eventually be obtained, hence a smaller estimation variance. Therefore, at a given stage, the new node to be added to the system will be that which causes the largest decrease in the estimation variance. This technique is also referred to more generally as the *variance-reduction approach* [Rouhani, 1985]. It has been used in many studies, and many variants on it have been developed. However, some authors [e.g. Ben-Jemaa et al., 1994] warn that the sequential search method it employs may yield suboptimal solutions. Consider a sequential search method in which

113

we add a new monitoring node at each stage, and suppose that at a given stage we have a set S_0 of n_0 nodes previously selected. In the following stage, the sequential method is supposed to add a new node while keeping the previously selected set of nodes (S_0). The outcome of this process is not necessarily optimal for the simple reason that the new set of $n_0 + 1$ nodes is selected only from the layout of possibilities which include the set S_0. However, there may exist other combinations of $n_0 + 1$ nodes that will yield a lower variance of estimation but do not necessarily include all the nodes of set S_0. Analogously, it can be shown that also by using the reverse process (of starting with all potential measurement nodes and then sequentially deleting nodes), the sequential search method cannot guarantee optimal solutions.

Unlike the sequential search approach or variance-reduction technique, the so-called **simultaneous search approach** selects all of the sampling nodes of the monitoring network simultaneously. A representative of this approach is the branch-and-bound algorithm [Carrera et al. 1984; Carrera and Szidarovszky, 1985; Ben-Jemaa et al., 1994], which consists of searching in a simultaneous way for the optimal monitoring nodes along preconstructed tree branches. The branch-and-bound algorithm sorts through the potential network configurations to find the single combination of measurements that minimizes the estimation variance while meeting the cost constraint. This algorithm is based on the monotonic properties of the estimation variance: the uncertainty cannot increase as data are added to a measurement network and cannot decrease as data are removed from a network. These properties are used to systematically search through the potential network designs and select the optimal sampling plan. If one has to select a

114

monitoring network of n nodes from a total number of W candidate nodes, there will be $\begin{pmatrix} W \\ n \end{pmatrix}$ possible combinations of nodes, or network layouts, and each layout possibility corresponds to one branch of the search tree. The search algorithm consists of a step-by-step optimization in which one node is added at a time while moving along the search tree branches. On the basis of the calculated estimation variance at each stage, the search method proceeds forward along the same branch to add an additional node (if the obtained variance of estimation is the lowest found so far), or leaves the current branch and moves one step back to proceed along another branch (if the obtained estimation variance is larger than the smallest found so far).

As an example, consider the case where there are 12 potential measurements (Figure 5.1). Each of the nodes of the tree represents a potential network design. The first node corresponds to the network with all measurements included. Moving along from one node to another corresponds to removing one measurement. The branch-and-bound algorithm searches the nodes along the tree branches to find the single combination of measurements that minimizes the objective, subject to cost and any other constraints applicable.
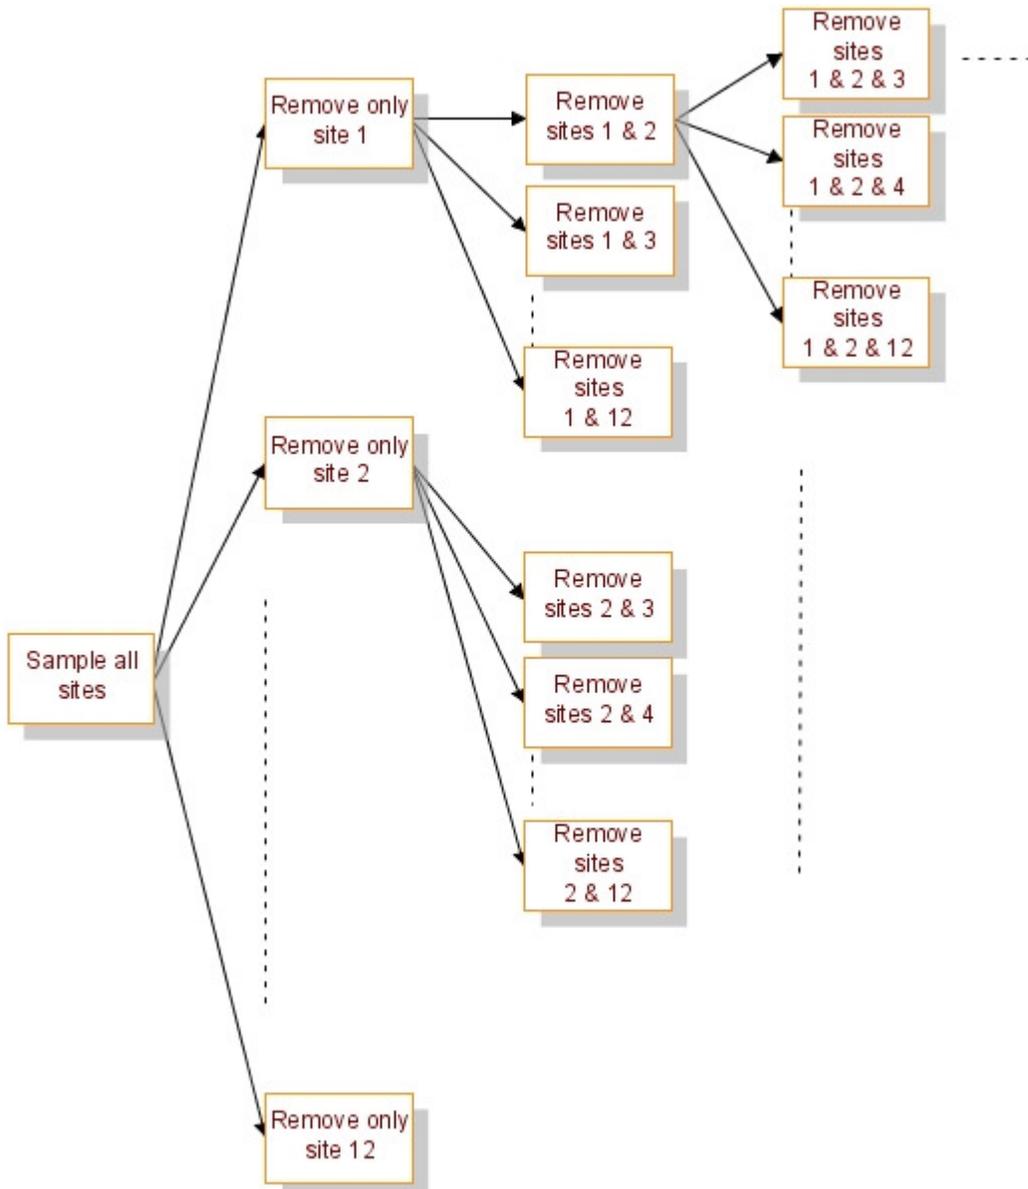
Figure 5.1. Schematic branch-and-bound algorithm showing some of the possible branches for a network of 12 nodes.

The branch-and-bound algorithm begins at the first node of the decision tree, with all measurements in the network. Initially, the first measurement point is removed from the network and this plan's cost is calculated. This forward process of removing a

116

measurement and updating the network cost continues along each branch until we obtain

a feasible network design, i.e., a network whose cost and other constraints comply with

the requirements. At this point, the objective is calculated for this network. The algorithm

then moves back one step, storing this network design, along with its objective and cost,

as the current optimal solution. The algorithm then proceeds forward again by removing a

new measurement and calculating the new objective and network cost. The algorithm

must now decide whether to proceed forward or to move backward and discard the

remaining nodes of this branch. Wagner [1995] recommends taking this decision "on the

basis of the following criteria:

1. If the objective is less than the smallest found thus far and the network cost is
   feasible, this network design is stored as the current optimal solution. The
   algorithm then moves back one node and proceeds forward by removing a new
   measurement.

2. If the objective is not less than the smallest found thus far, the algorithm moves
   backward one node, discarding the remaining nodes of this branch. The algorithm
   chooses to move backward at this point because moving forward means removing
   an additional measurement from the network, which could not improve the
   estimation variance. The search then continues forward by removing a new
   measurement.

3. If a branch contains no feasible solutions, the algorithm moves backward one
   node, discarding the remaining nodes of the infeasible branch. This could
   potentially save evaluating many plans further along that branch.

4. If the final node of a branch has been reached, the algorithm moves back one node
   and either proceeds forward by removing a new measurement or stops because the
   final node of the tree has been reached.

5. If the objective is less than the smallest found thus far but the cost is infeasible,
   the algorithm proceeds forward, removing a new measurement from the network
   and recalculating the prediction uncertainty and cost. The forward movement
   proceeds until one of the above criteria is met."

The branch-and-bound algorithm proceeds along each branch according to the above criteria until all possible network configurations have been considered. It is important to note that the design objective (and constraints, where applicable) are typically calculated for only a fraction of the possible network designs. The remaining designs are removed from consideration according to the criteria listed above.

In the next Section, this search technique is used with both of the design criteria [(1) and (2)] from Subsection 2.2, to find optimal sampling network configurations for the plume estimation methods developed in this work, TrEK and IFM, and for KT as a comparison.

3. Examples, Results and Discussion

To illustrate the methodology presented above, the synthetic aquifer from Example 2 in Chapter 3 is used here again. For the sake of consistency, the contamination event, contaminant plume, and even monitoring well locations (in Examples 1 and 3 here), all remain faithfully identical to Example 2 in Chapter 3. Moreover, we use the exact same estimation methods used there, with the same parameters. However, we now assume that fewer wells are sampled, relative to the 12 we had in Chapter 3.

We select the optimal designs below according to the criteria described above, in Subsection 3.2: for each case, first, the maximum variance is minimized; then the average estimation variance across the domain is minimized. This minimization is achieved via the branch-and-bound algorithm described in Subsection 3.3. Each well-configuration and plume-estimation setup is judged quantitatively by two measures: the root-mean-

square error (RMSE) and the percentage of the estimated domain which is within two standard deviations of the posterior plume distribution. The root-mean-square error can be calculated as follows [Kish, 2000]:

$$RMSE = \sqrt{\frac{\sum (\mathbf{z} - \hat{\mathbf{z}})^2}{N}} \tag{63}$$

where $\mathbf{z}$ are the true concentrations at the $N = 992$ points estimated and $\hat{\mathbf{z}}$ denotes the best estimates. This provides a measure of how good the estimate is, and more importantly allows us to compare estimates obtained by the different methods and/or using different measurement locations. The second quantitative measure is the percentage of the estimates which are within two standard deviations of the posterior plume distribution:

$$\frac{\sum I_{|\hat{\mathbf{z}} - \mathbf{z}| < 2\boldsymbol{\sigma}}}{N} \times 100 \tag{64}$$

where $\boldsymbol{\sigma}$ is the standard error of estimation (i.e. the square root of the kriging-, IFM-, or TrEK-variance) and the indicator function I counts the number of points estimated which fall within the confidence intervals, or in other words fulfill the condition $|\hat{\mathbf{z}} - \mathbf{z}| < 2\boldsymbol{\sigma}$. These two measures provide a quantitative assessment, in addition to the uncertainties (represented by the estimation standard errors), of the performance of the well-selection and plume-estimation methods, which allows to observe potential improvement the methods developed herein may demonstrate.

As in Chapters 3 and 4, it is assumed that the variogram models and their parameters are known. Also, in the first two examples below, the flow and transport models are completely known, except for the release history, which can be seen as a boundary condition at one point in space. Thus, we optimize a network for collecting only concentration data. In the third example, we assume we only have measurements of hydraulic conductivity, rather than its entire spatial distribution (as in case iii in Chapter 4, except that there are no concentration measurements before sampling in this Chapter). The hydraulic conductivity field is estimated using kriging and input to the transport model. Consequently, the estimation variance (upon which the well selection is based) takes into account the uncertainty from this estimation of **K**, in addition to the uncertainty from not knowing the concentrations themselves.

The figures in this section are organized as follows. Each row represents a specific selection of measurement locations corresponding to a specific method, and each column represents estimation using a specific method (KT, IFM, or TrEK). Thus, for example, the top row represents well-selection using KT: from left to right, the panels represent estimates of the KT, IFM, and TrEK methods, respectively, all using measurement locations selected by TrEK. Note that in Examples 1 and 2 below, IFM and TrEK happened to choose the same wells in all cases. Therefore to avoid redundancy, one of the rows was eliminated from the plots, and the row entitled "IFM/TrEK selection" represents the common choice of these two methods. As in Chapter 3, all of the quantities in the Chapter are unitless, as these Examples are based on synthetic data.
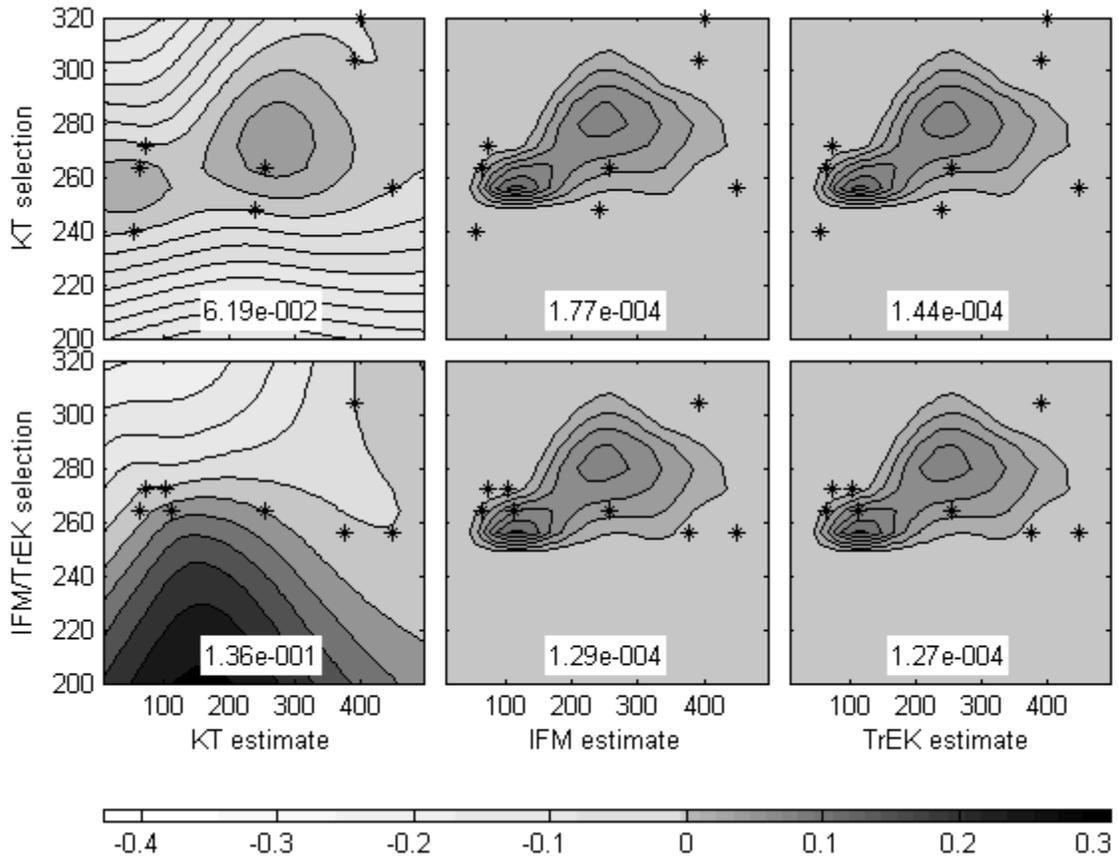
Figure 5.2. Example 1 - Plume best estimate with $\underset{p \in \mathbf{S}}{Minimize}\{\max \mathbf{V}_{\hat{z}}\}$ **criterion. The number in the white rectangle represents the best estimate's root-mean-square-error. All quantities are unitless.**

3.1 Example 1

In the first two Figures (5.2 and 5.3), the design criterion used was minimizing the *maximum* estimation variance in the domain. In this example, TrEK and IFM happened to choose the same monitoring locations. The results of estimating the plume using these measurements in IFM and in TrEK are almost identical, as are the uncertainties. This suggests that the spatial-correlation component in TrEK only had a negligible contribution to the estimation; most of the information used to estimate the contaminant plume actually comes from the flow-and-transport model. More about the similarity of

121

TrEK and IFM's performance can be found in Chapter 3, where these methods are developed.

Five of the measurement locations selected by IFM and TrEK were also chosen by KT (however note that *at least* four wells must have been common, as there were only 12 wells to select from). The most striking feature of this figure is that IFM and TrEK seem to perform very well (all RMSEs are very low, $1\sim2\times10^{-4}$), regardless of the selection of wells, under these conditions. KT, in contrast, does not perform well at all, no matter which method selects the wells. The best result for KT seems to be the case when the well selection is also performed using the KT variance (RMSE = 0.062) – at least some feature of the plume seems to be captured with those measurements, which deliver better results than the TrEK/IFM selection for this method (RMSE = 0.136). This demonstrates the importance of using a well-selection method that is appropriate for the estimation method.

Figure 5.3, which depicts the estimation variance, also shows the importance of using the IFM design criterion or the TrEK design criterion (rather than the KT design criterion) in the monitoring site selection for these methods: using these criteria leads to a much lower estimation variance than using locations selected by KT. In addition, we see as before that the IFM and TrEK uncertainties are generally much lower the KT uncertainties, reflecting the fact that kriging cannot give very useful information with only eight monitoring wells (except at points which are very close to those wells).
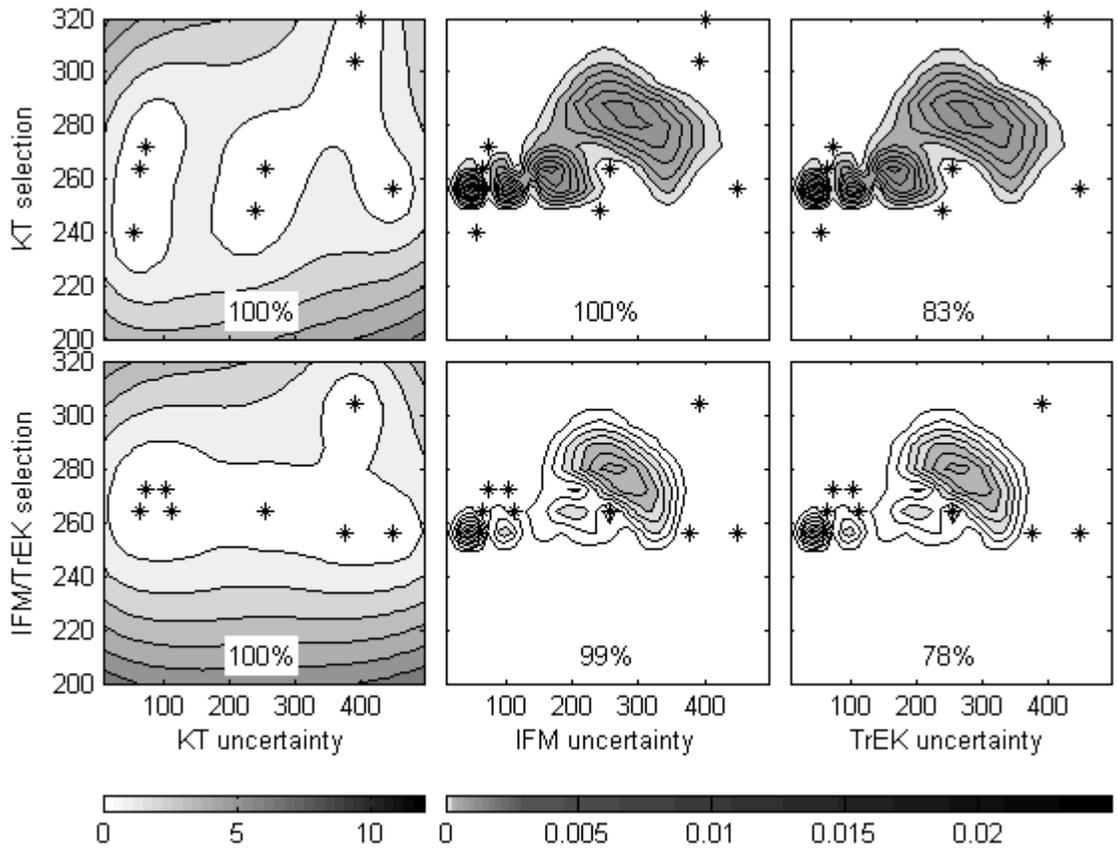
Figure 5.3 Example 1 – Plume concentration uncertainty with $\underset{p \in \mathbf{S}}{Minimize}\{\max \mathbf{V}_{\hat{z}}\}$ **criterion. The number in each graph represents the percentage of points within the 95% confidence intervals.**
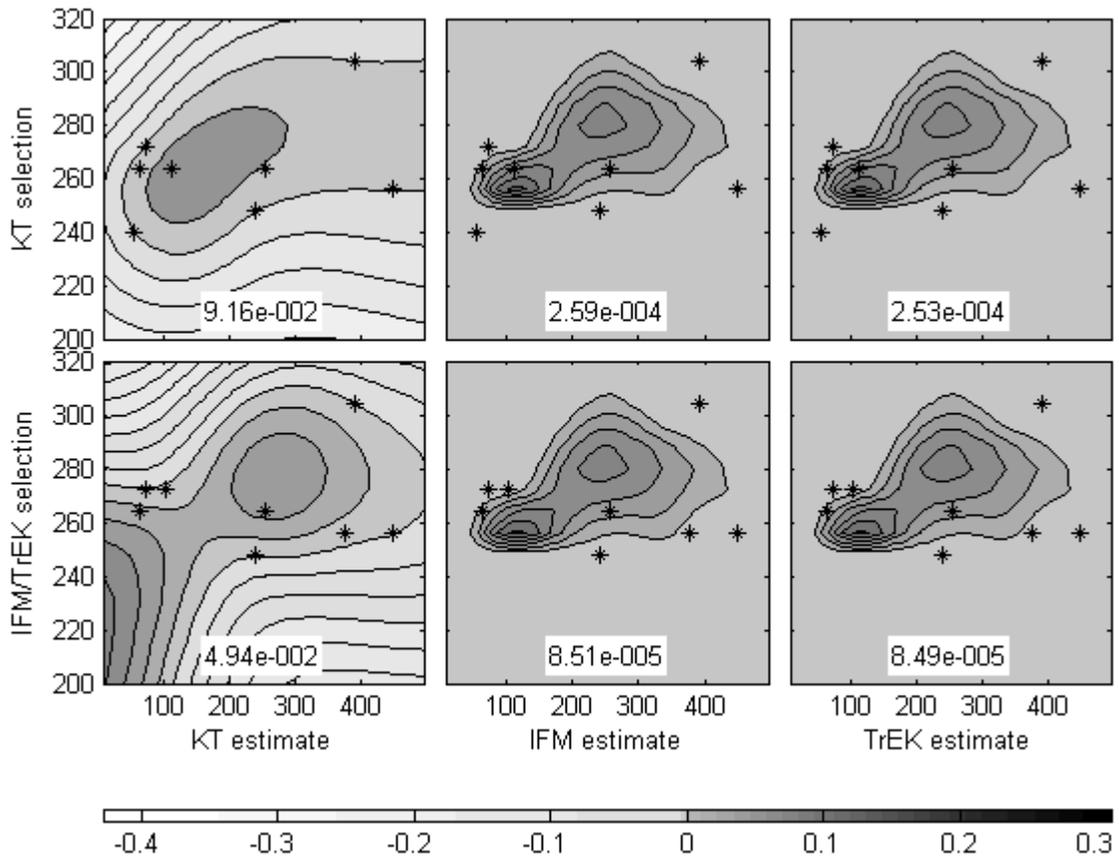
Figure 5.4. Example 1 - Plume best estimate with $\underset{p \in \mathbf{S}}{Minimize}\{\overline{\mathbf{V}}_{\hat{z}}\}$ **criterion. The number in the white rectangle represents the best estimate's root-mean-square-error. All quantities are unitless.**

In the next two Figures (5.4 and 5.5), the design criterion used was minimizing the *average* estimation variance in the domain. Again, TrEK and IFM happened to choose the exact same measurement locations, and demonstrate very similar results. The selection of wells by the kriging-with-a-trend method resulted in six wells in common with the other two methods. The results of using this design criterion seem to improve the performance of kriging, especially when it uses its own selected locations. However, the KT estimates are still poor compared to the IFM and TrEK estimates. The uncertainties seem to paint a similar picture to the one produced by the "Minimize *maximum-*

variance" design criterion: very low uncertainties for IFM and TrEK, and very high uncertainties for KT, throughout the domain. More importantly, the measurement locations selected by IFM and TrEK lead to lower uncertainties in these methods' estimates. In addition, their RMSEs are lower with their own well-selections.

The percentages on the uncertainty figures (5.3 and 5.5) express how often the best estimates lie within two standard deviations (σ) of the true plume. The KT estimates are conservative in this Example, in the sense that the best estimates is always within two standard deviations (σ) of the true plume distribution. The IFM estimates fulfill this condition for 84 – 100% of the points estimated, but the TrEK estimates only for 70 – 83% of the points estimated. The TrEK estimates are very similar to the IFM estimates in this Example, but their uncertainties are always lower, so less of the truth lies within its confidence intervals.

3.2 Example 2

In this next example, the same aquifer, contaminant plume, and methodologies are used again. However, five candidate measurement locations are selected from a regular grid of 5×5. This scenario is more akin to network design cases (as opposed to well selection), where new well locations are often chosen from a larger pool of potential observation sites. It enables us to observe the preferred spatial distribution of measurement sites for each method. Again, both design criteria are tested with all three selection/estimation methods, and the results are presented, as above, in a set of six figure panels.
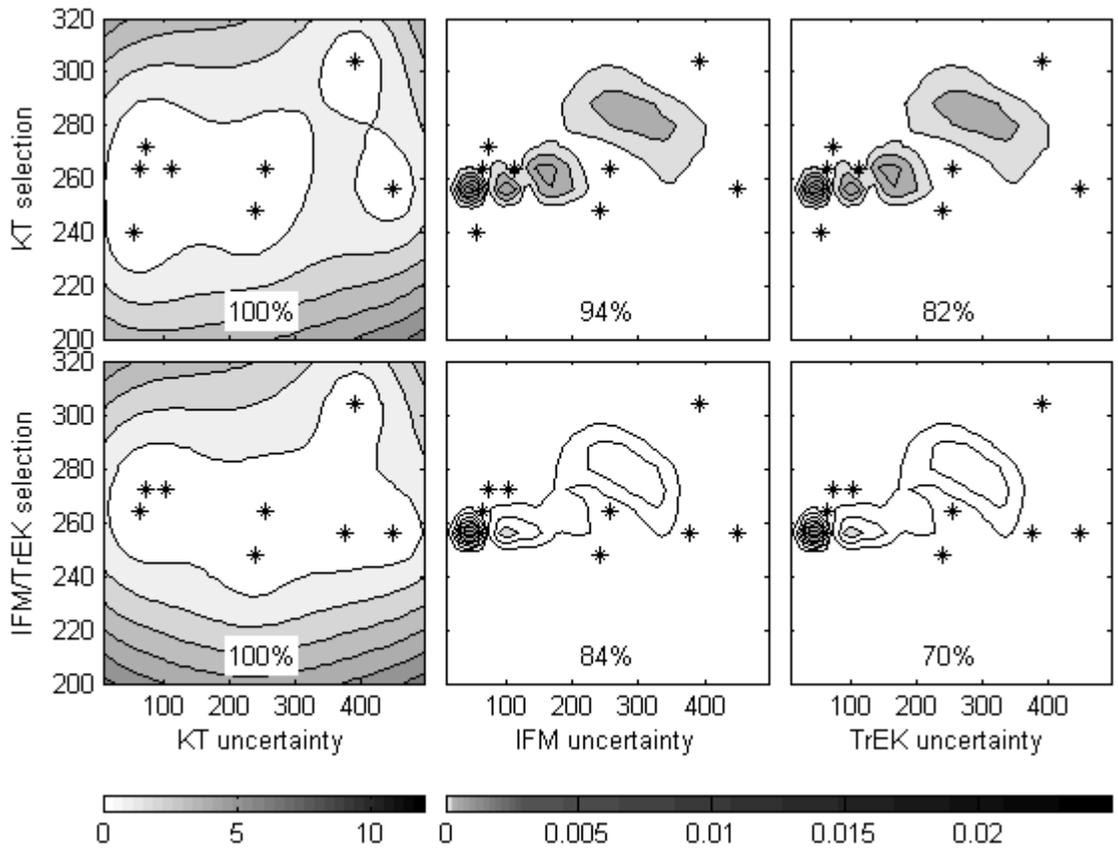
Figure 5.5. Example 1 - Plume concentration uncertainty with $\underset{p \in S}{Minimize}\{\overline{\mathbf{V}}_{\hat{z}}\}$ criterion. The number in each graph represents the percentage of points within the 95% confidence intervals. All quantities are unitless.
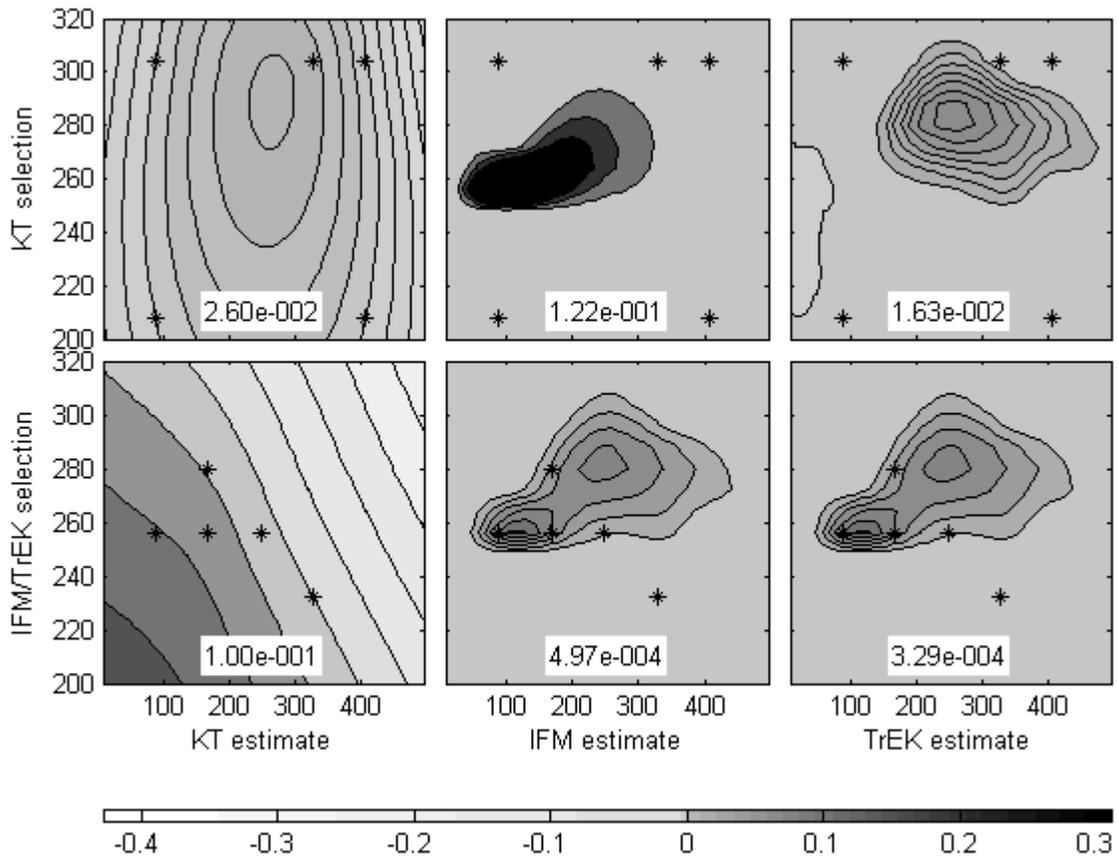
Figure 5.6. Example 2 - Plume best estimate with $\underset{p \in S}{Minimize} \{\max \mathbf{V}_{\hat{z}}\}$ **criterion. The number in the white rectangle represents the best estimate's root-mean-square-error. All quantities are unitless.**

While IFM and TrEK again choose the same observation locations, which seem to favor the plume "hot spots" (the locations with the highest concentrations), the KT selection of wells seems to follow a different pattern, which in general looks to spread out as much as possible. This reinforces the differences in the uncertainties, between the methods, as seen since Chapter 3. This time, the KT choice of wells was nearly identical for both design criteria, differing only in one well. Four of the wells were chosen at the four corners of the rectangular domain, and the fifth well was placed at the center of the domain in order to minimize the average variance, but at a different location to minimize

the maximum variance. Placing that last well in the center of the domain seems more symmetrical or even natural, as it reduces the variance it all directions of the domain almost equally. But a close look at the first (i.e. top-right) panels in Figures 5.7 and 5.9 reveals that the seemingly odd location for the fifth well does indeed reduce the maximum variance (note the high uncertainty on the right-hand side of the first figure, which is smaller in Figure 5.7 than in Figure 5.9).
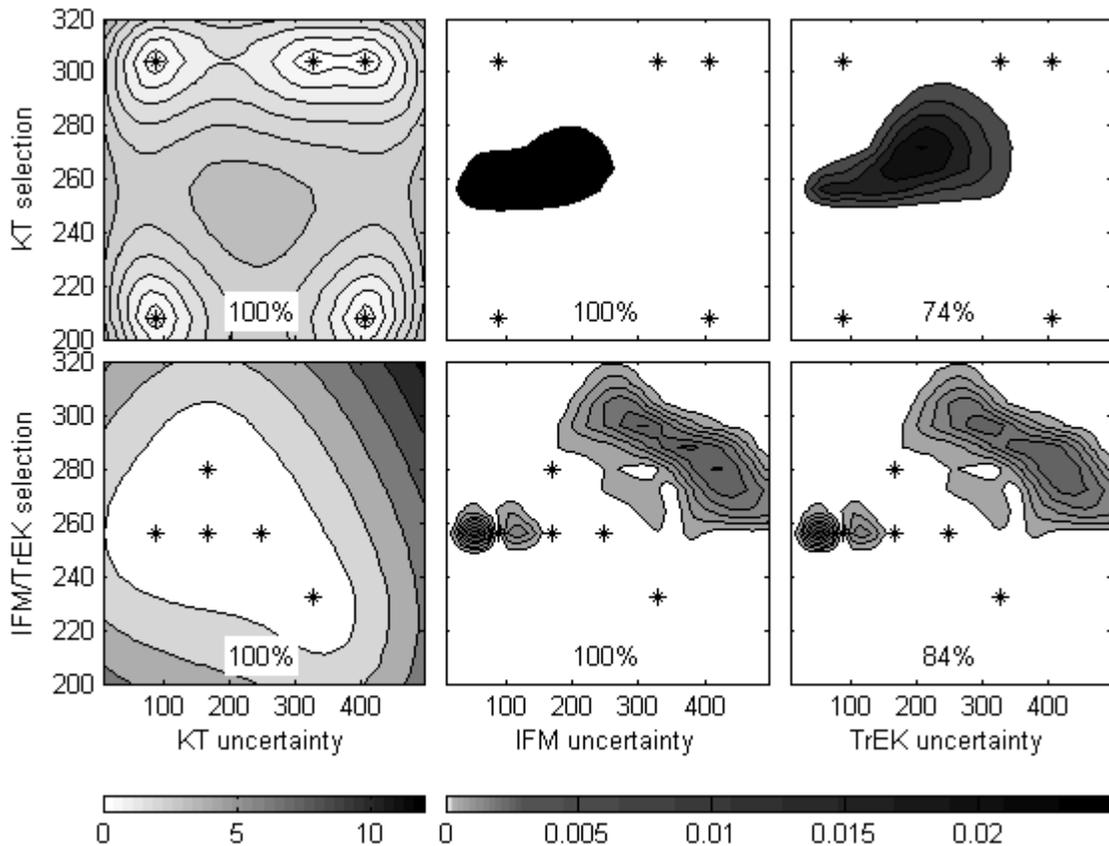
Figure 5.7. Example 2 - Plume concentration uncertainty with $\underset{p \in \mathbf{S}}{Minimize}\{\max \mathbf{V}_{\hat{z}}\}$ **criterion. The number in each graph represents the percentage of points within the 95% confidence intervals. All quantities are unitless.**
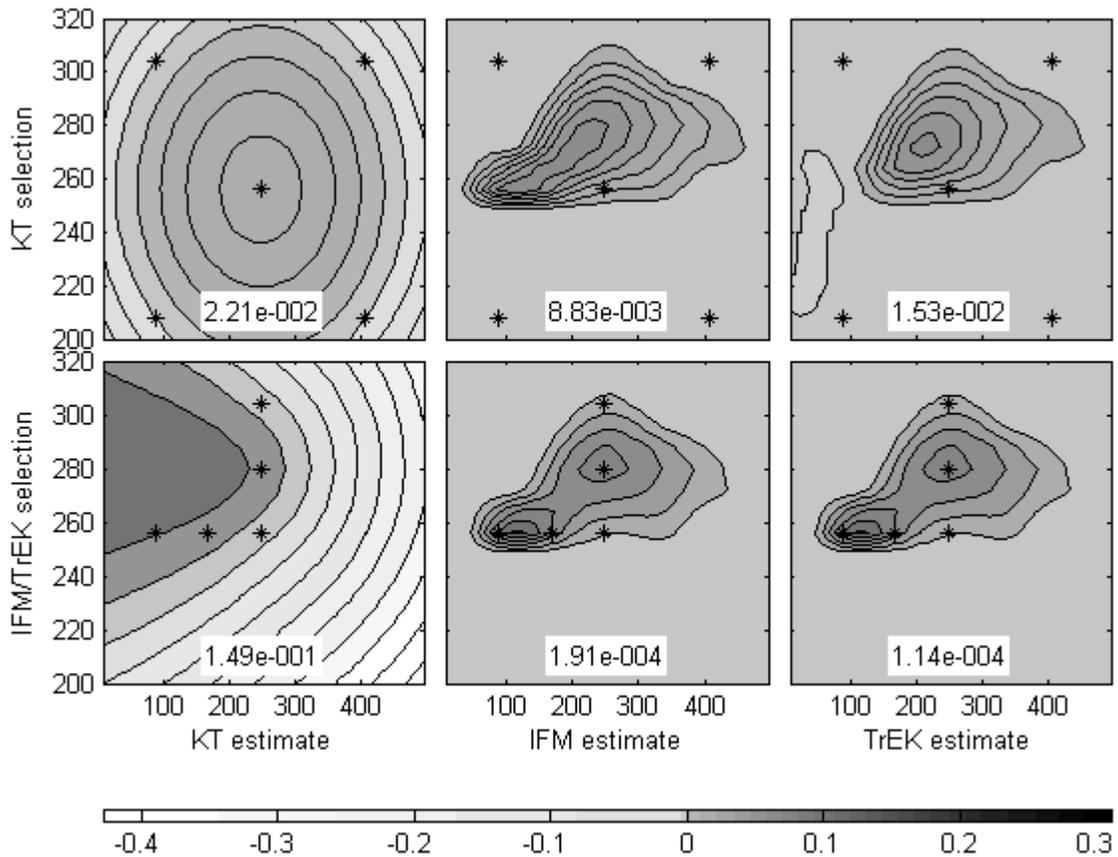
Figure 5.8. Example 2 - Plume best estimate with $\underset{p \in \mathbf{S}}{Minimize}\{\overline{\mathbf{V}}_{\hat{z}}\}$ criterion. The number in the white rectangle represents the best estimate's root-mean-square-error. All quantities are unitless.

The IFM and TrEK methods chose the same wells: again, close to the point of release on the western boundary of the domain, in an area very sensitive to the release. As in Example 1, this choice significantly reduces the estimation variance, compared to the choice KT makes for these methods. As before, TrEK and IFM perform very well, with both design criteria and with all selections. Their best estimates resemble the true plume, and their uncertainties are very low. In contrast, the KT selection does not include any locations on the plume, or around the center of the domain at all. Accordingly, this method demonstrates high uncertainty, due to very little data. Note that in terms of RMSE, for all methods and all design criteria within this Example, the RMSE is much

lower when each method selects the wells for itself, i.e. the kriging variance is used to select wells for KT, IFM variance is used to select wells for IFM, and the TrEK variance is used to select wells for TrEK. These differences may be more pronounced (up to two order of magnitude for TrEK) here in Example 2, because the wells configurations are much more different from case to case than they were in Example 1.

Quantitatively, KT and IFM are again conservative, and 100% of the true plume lies within their 95% confidence intervals. For TrEK, this percentage is as low as 27%, for one case in which the five measurement locations are selected by the KT-variance.
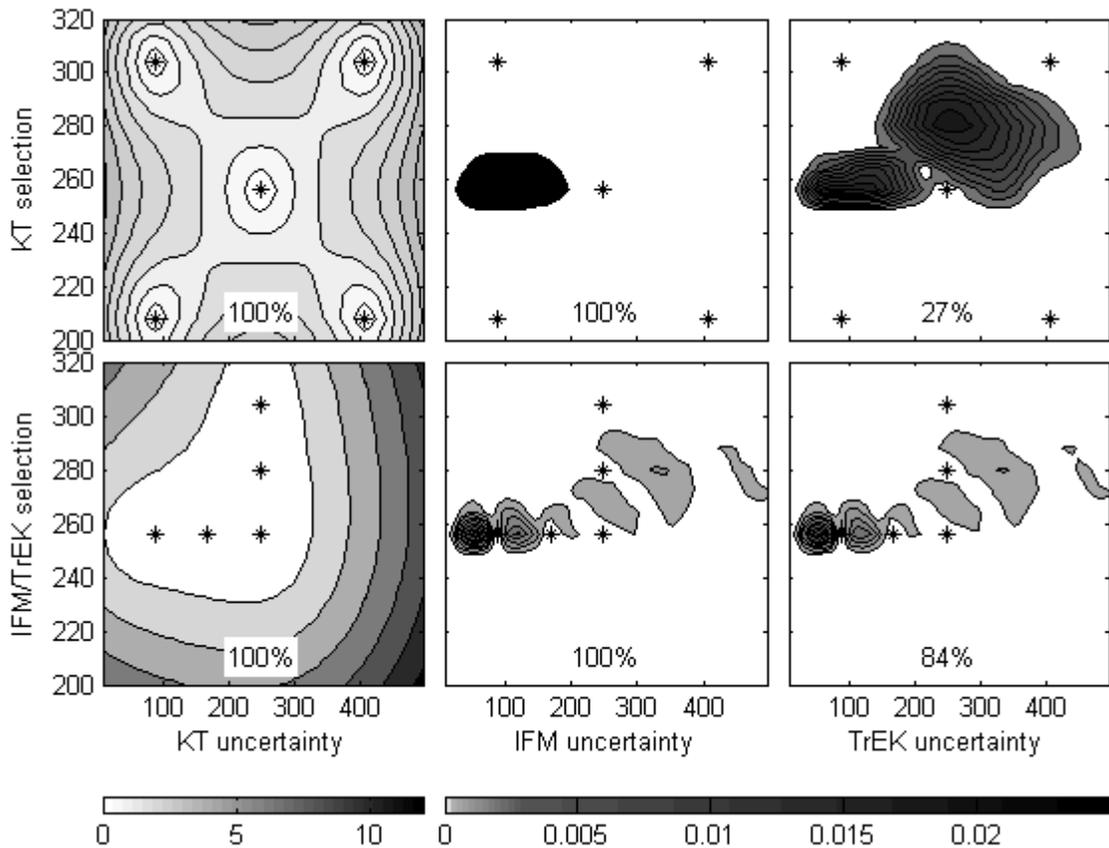
Figure 5.9. Example 2 - Plume concentration uncertainty with $\underset{p \in \mathbf{S}}{Minimize}\{\bar{\mathbf{V}}_{\hat{z}}\}$ **criterion. The number in each graph represents the percentage of points within the 95% confidence intervals. All quantities are unitless.**

3.3 Example 3

In this example, we no longer assume that the numerical flow- and transport-models are perfectly known. Rather, the hydraulic conductivity **K** is only measured at the 12 measurement locations from Chapter 3, and consequently ordinary kriging is performed to estimate the full conductivity field, which is subsequently input to the flow-and-transport model used in IFM and in TrEK. This scenario is more likely closer to typical practical situations, in which there is no way to obtain the true spatial distribution of the entire hydraulic conductivity field, and thus a perfectly-known model is not feasible. While the plume concentration remains the only variable being measured, its uncertainty as represented in this case is more realistic, because it also takes into account the flow-and-transport model uncertainty, in addition to the other sources considered before.

The various Jacobian matrices **H** for IFM and TrEK in this Example, were constructed using models based on the best estimate obtained for **K**. The way in which the plume-concentration estimation variance $\mathbf{V}^{C_e}$ (Eq. 19 in Chapter 4) is calculated in this example, is similar to the way it is calculated in Chapter 4 (Eq. 17): $\widetilde{\mathbf{H}}_s^{C_e}$ and $\widetilde{\mathbf{H}}_r^{C_e}$ are used to project the joint uncertainties both from the unknown contaminant source and from the unknown hydraulic conductivity (Eq. 19 in Chapter 4) onto the plume space. However, there are no iterations here, since the hydraulic conductivity is already estimated, and we are not seeking to refine this estimate using concentration measurements – only to calculate the estimation variance in order to locate the optimal measurement locations. In any case, we cannot refine this estimate because we assume that no concentration measurements have been taken at this point.
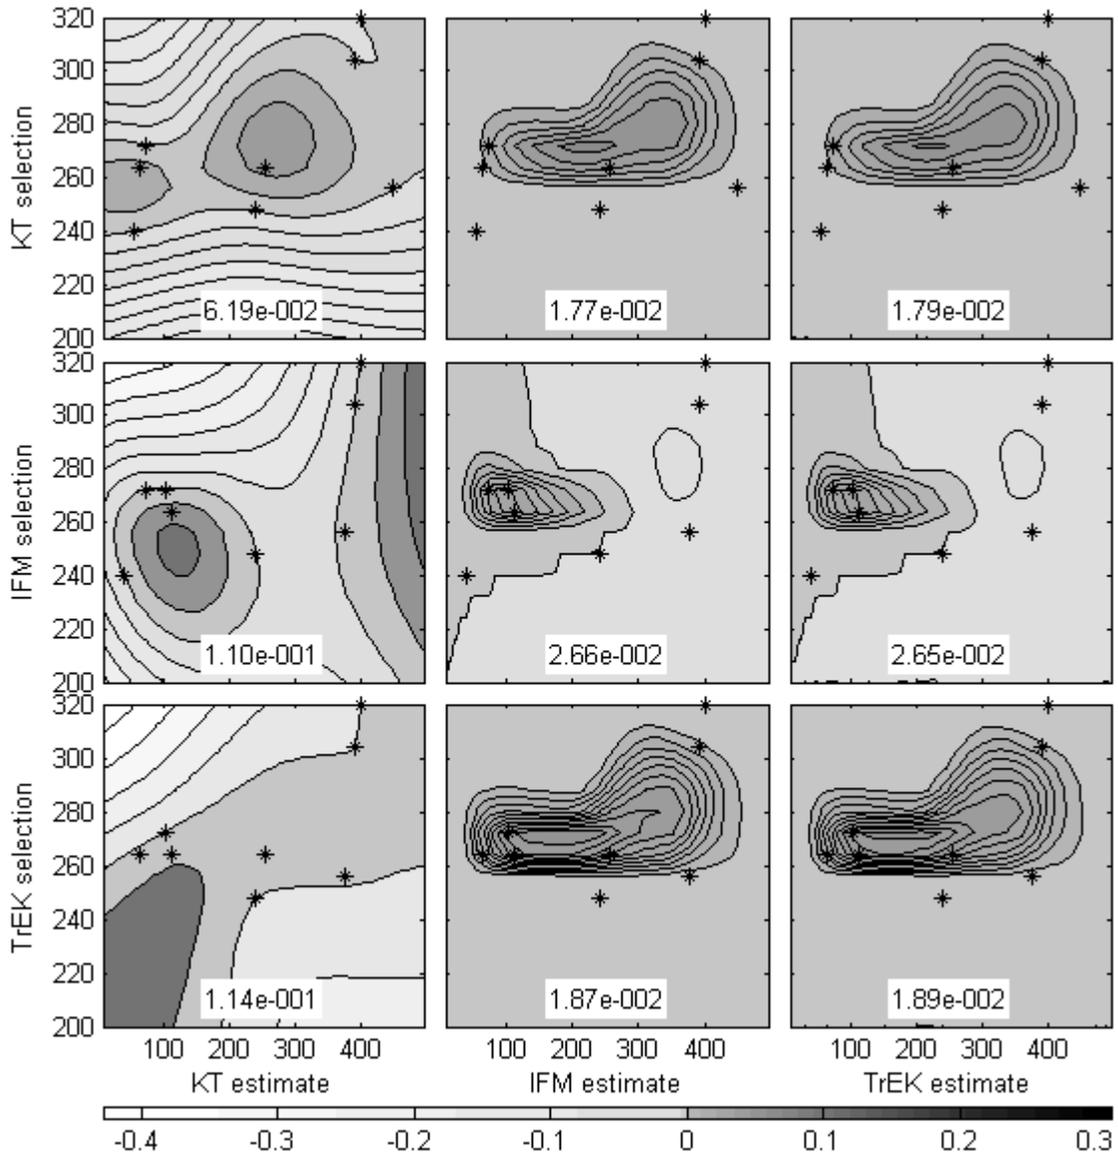
Figure 5.10. Example 3 - Unknown K - Plume best estimate with $\underset{p \in S}{Minimize} \{ \max \mathbf{V}_{\hat{z}} \}$ **criterion. The number in the white rectangle represents the best estimate's root-mean-square-error. All quantities are unitless.**
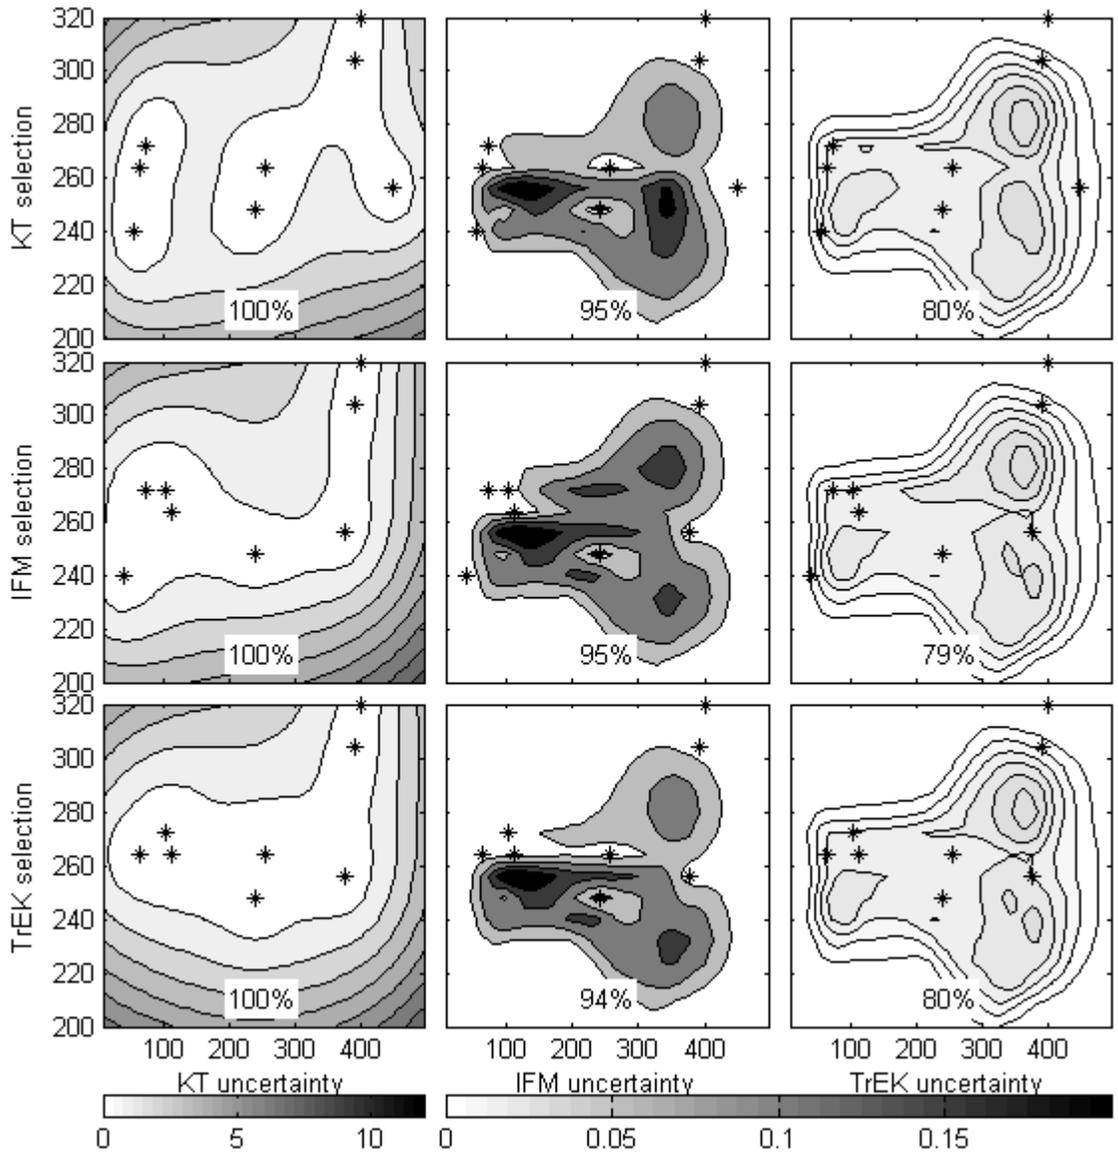
Figure 5.11. Example 3 - Unknown K - Plume concentration uncertainty with $\underset{p \in S}{Minimize}\{\max \mathbf{V}_{\hat{z}}\}$ **criterion. The number in each graph represents the percentage of points within the 95% confidence intervals. All quantities are unitless.**

As we have less data here (than in the first two examples) and the estimation model changes, we would expect the estimates to be further removed from the true plume distribution. Accordingly, the model-data mismatch error is adjusted, its variance increased to $\sigma_R^2 = 10^{-5}$. The estimation of the hydraulic conductivity field with ordinary kriging produced much smoother flow fields, than with those obtained by the numerical

models using the true conductivity. The effect of this smoother flow field on the contaminant plume can be seen in Figures 5.10 and 5.12, which show that the estimated contaminant plumes are also smoother. The estimation variances have become much larger, reflecting the additional uncertainty that derives from not knowing the hydraulic conductivity. Specifically, there is now high uncertainty at the bottom half of the domain (Figures 5.11 and 5.13), because very few hydraulic conductivity measurements were taken there. It is not surprising that the results in this example are inferior to those of Examples 1 and 2 above: The best estimates are further from the true plume than they were in the previous examples, and the uncertainties on those estimates are higher. Moreover, they do not necessarily reliably reflect the true estimation error, as the low percentage points on the uncertainty plots (Figures 5.11 and 5.13) suggest. Likewise, the RMSEs are also much higher, as the uncertain hydraulic conductivity handicaps the flow-and-transport model, relative to the previous examples, resulting in best estimates that are further away from the truth.
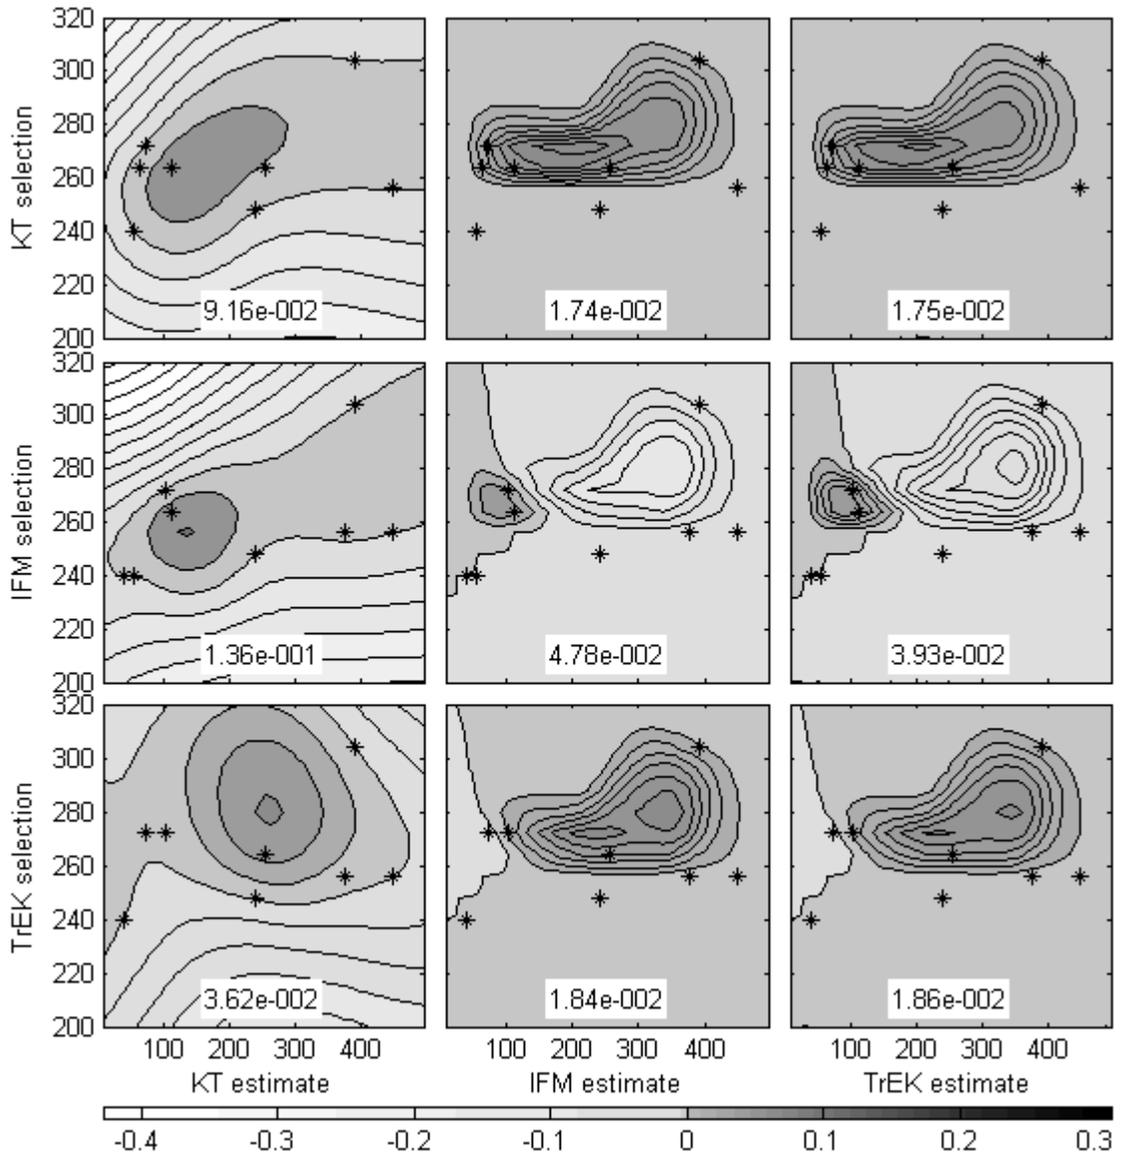
Figure 5.12. Example 3 - Unknown K - Plume best estimate with $\underset{p \in \mathbf{S}}{Minimize}\left\{\overline{\mathbf{V}}_{\hat{z}}\right\}$ **criterion. The number in the white rectangle represents the best estimate's root-mean-square-error. All quantities are unitless.**

For both well-selection criteria, IFM and TrEK seem to prefer a more spread-out network configuration, instead of the more centrally-located measurement locations encountered in Example 1. This is probably due mainly to the smoothing effect of the $\widetilde{\mathbf{H}}^{C_e}$ matrices, which produced a more even distribution of the uncertainty over space. For example, not

knowing the hydraulic conductivity **K** made locations which were previously relatively insensitive to the release history (e.g. the bottom half of the domain), much more uncertain. The contaminant concentration at these locations may change dramatically as **K** changes; moreover, for some **K** distributions, these locations may display a much larger sensitivity to the release history **s**, which also increases their uncertainty.

In terms of the percentages of points that lie within the 95% confidence intervals, the pattern is similar to that of the previous Examples: the KT estimates are most conservative, next are the IFM estimates (94 – 95%), and finally TrEK (79 – 80%). This difference between the TrEK best estimates and the true plume distribution is larger than twice the standard error (2σ) more than the ideal case of 5% (corresponding to 95% confidence intervals). This may be explained by the specific aquifer setup and well configuration used here. The residuals of TrEK may not follow the implicit Gaussian assumption required for the correspondence between the standard deviation σ and the 95% confidence intervals.

4. Conclusions

This Chapter has demonstrated one example of adapting a network design methodology for use with IFM and TrEK. The first conclusion at this point is that **this use of kriging-based methodologies is indeed relatively simple to carry out**. The versatility of IFM and TrEK enables the use of other network-design methods that were originally developed for kriging: the advantage here is that there are already many such methods,

which were developed for a variety of situations. Most of these methods can probably

easily be updated to work with IFM and TrEK.
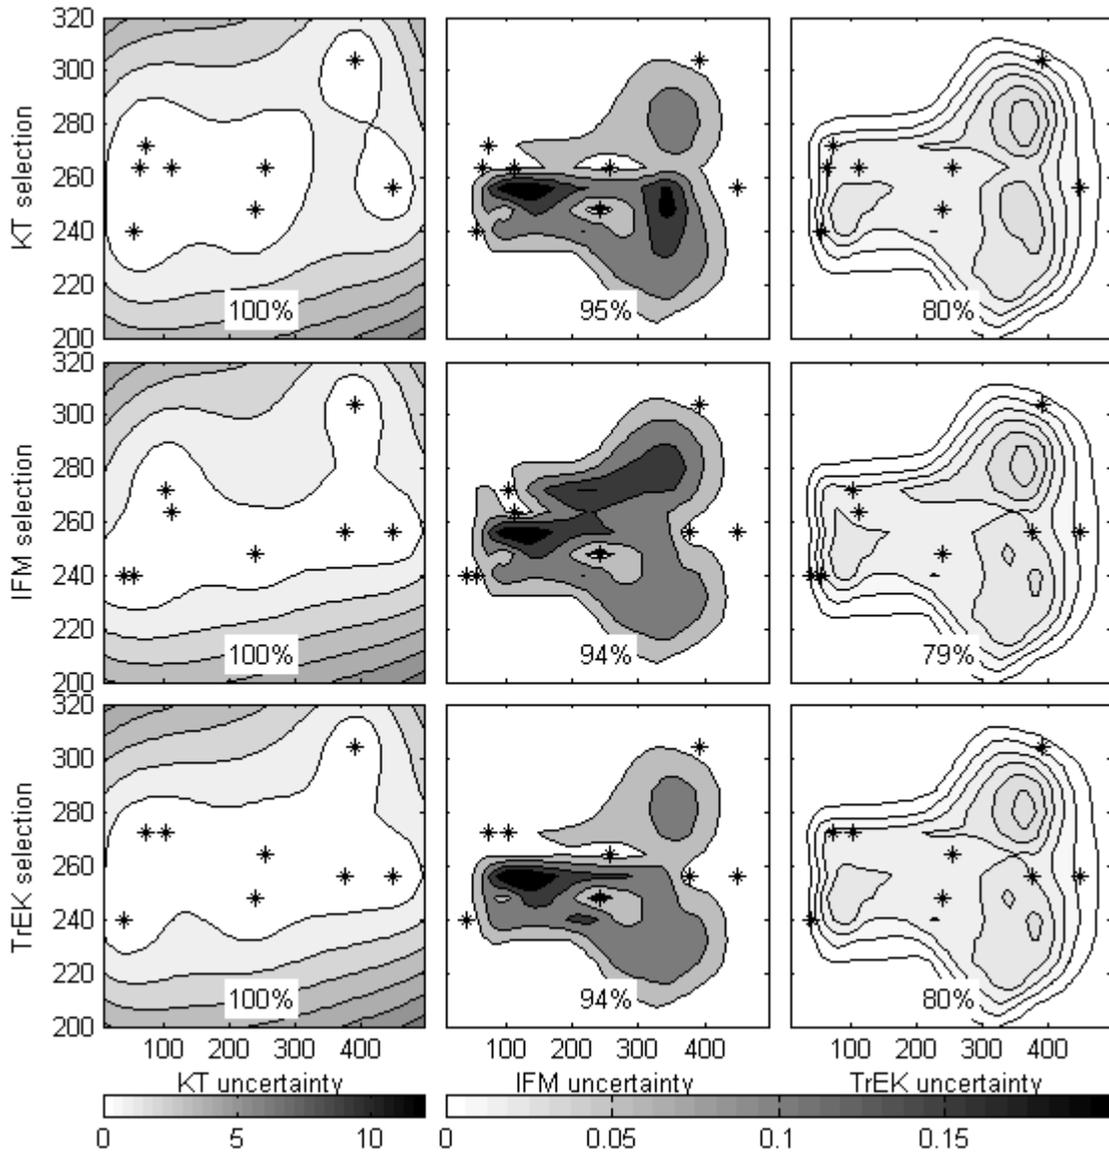
Figure 5.13. Example 3 - Unknown K - Plume concentration uncertainty with $\underset{p \in S}{Minimize}\{\bar{\mathbf{V}}_{\hat{z}}\}$ **criterion. The number in each graph represents the percentage of points within the 95% confidence intervals. All quantities are unitless.**

The second important conclusion is that this use of modified kriging-based

methodologies further improves the results, compared to the original. In all examples, the

results of IFM and TrEK selections were more precise (displayed lower estimation variances), more accurate (displayed lower RMSEs), and not significantly less realistic, in terms of the percentage of the best estimates within the 95% confidence intervals. In Chapter 3, the conclusion was to use IFM and TrEK for estimating groundwater contaminant plumes, whenever possible. Here, we might extend this conclusion to using a TrEK- or IFM- based network design methodology, whenever using these plume estimation methods. The improvement IFM and TrEK have to offer (over KT) – both in the selection of wells and in plume estimation – can be seen on all Figures in this Chapter. As this has to do both with network design and with estimation, it is actually one of the overarching themes of this thesis: **incorporating transport information into this type of analysis improves both the optimization of well selection and the quality of plume-estimation**. Moreover, the combined effect of both is significant.

Furthermore, many of the conclusions of Chapter 3 are valid here as well. First, these results are not unique to the specific model we are using here; other compatible flow and transport models should work as well. Second, although we did not enforce nonnegativity on the release history, we would not expect that to significantly affect the best estimates. The variances may become smaller, and may change the selection of measurement sites. However, many methods of enforcing nonnegativity are iterative, and may not be applicable to the method presented in this Chapter. Using a different type of inversion may also affect the estimation results, but probably would not change the superiority of IFM and TrEK over kriging, also in the selection of wells. While only the Example aquifer from Chapter 3 was used to demonstrate the network design method presented

138

here, there is no reason to believe it would not work in other scenarios where IFM and/or TrEK are applicable.

More specifically, we have again seen the effect of kriging's much larger estimation variance on TrEK: as the information provided by this component of TrEK is relatively low, the performance of TrEK is very similar to that of IFM. This behavior would probably change in cases where the information provided by spatial correlation into TrEK would be comparable to that provided by the inverse/forward transport model. This could turn out to be a very valuable feature in such cases, because choosing a TrEK-based sampling plan may balance between a spatial-correlation-derived tendency to spread out the sampling locations (and thus decrease the variance at the periphery), and an inverse/forward transport model-tendency to concentrate them around the locations which happen to be sensitive to the historical plume or release (and thus decrease the variance around the plume itself).

Even though there are differences in the performance of the methods under different design criteria, these would actually be chosen by the practitioner, so it is not possible to recommend on using one method versus another. Minimizing the maximum variance may be chosen when there exists some regulatory threshold which must not be violated; minimizing the average variance may be better when all of the domain or plume is equally important, for example for mapping or for further [post-] processing of the results, using other applications.

As this Chapter focused on the replacement of kriging (in variance-reduction methods) by IFM and TrEK, the search method did not receive much attention, and may not have been optimal. Wagner [1995] claims that the branch-and-bound algorithm "can require an enormous computational effort, particularly when the number of potential measurements is large (the typical case) and the number of measurements in a feasible solution is small (also typical). The excessive computational requirements occur because of the need to advance farther out on a branch (and perform more objective and cost calculations) before an optimal solution is found or the branch is discarded."

A small disadvantage of using IFM or TrEK instead of kriging, is that the search may take longer to perform. This may be due to two reasons: first, the evaluation of each sampling plan takes longer for IFM or TrEK, as more computations are involved, larger matrices are to be inverted, and most importantly, flow-and-transport models must be repeatedly run. Second, some search methods may take advantage of symmetry that may exist for kriging, but not for IFM or TrEK in heterogeneous domains. For efficient search techniques run on a modern computer, this would probably not be a significant problem.

When a component of the flow and transport model is uncertain, for example hydraulic conductivity, the optimal sampling configuration may differ significantly from the known-model case, as demonstrated by Example 3. Part of the challenge in this case is to determine the effect of the uncertain component on the plume concentration uncertainty, or more importantly the new total plume concentration uncertainty. This was achieved

here by simulating a joint inversion for source release history and hydraulic conductivity, as one iteration of the simultaneous-estimation method (Chapter 4). In this Chapter, this was used only to find the new total plume concentration uncertainty, not to estimate hydraulic conductivity, so iterations were not needed. Note that they could not be performed anyway, since all of this network design is assumed to take place before any measurements of contaminant concentration are taken. In Example 1 and 2, IFM estimates which have used the IFM-selection achieved lower root-mean-square errors than other IFM estimates. TrEK estimates, however, did not achieve a lower RMSE with the TrEK selection, but all of their RMSEs in these examples were the same order-of-magnitude, and lower than the other methods'. Likewise in Example 3, the RMSEs for IFM and TrEK were the same order of magnitude. Note that as opposed to the estimation variance, which is always accessible (regardless of samples), the root-mean-square error is a measure which uses the true underlying distribution, which is not available in practice.

The results of this last example again demonstrate the importance of the hydraulic conductivity field, and the effects of its estimates both on network design and on plume estimation. They also reinforce the need to properly estimate this field, and incorporate the effects of the uncertainty of this process elsewhere, such as in network design and in plume estimation.

Finally, for the sake of simplicity, only the collection of *concentration* data was considered, and the variogram models and parameters were assumed to be known.

Future research may explore more advanced methodologies, which couple parameter-estimation and model-prediction uncertainty analyses with optimization, to identify the mix of hydrogeologic information (e.g., head, concentration, and/or hydraulic conductivity measurement locations) that should minimize a model's estimation variance under a variety of constraints.

Chapter 6
Application of Inverse/Forward Modeling
in a Laboratory Sand Tank[2]

1. Introduction

This Chapter presents work performed at the Colorado School of Mines Center for

Experimental Study of Subsurface Environmental Processes (*CESEP*) laboratory, aimed

at validating the inverse/forward modeling (IFM) method developed in Chapter 3. While

this method has been shown to perform well with synthetic cases [Michalak and Shlomi,

2007; Shlomi and Michalak, 2007], it has not been previously tested using laboratory or

field data.

Although field applications present the ultimate goal for the methods developed in this

dissertation, laboratory experiments are a cost- and time-effective approach for testing

new methods before moving to the field. Moreover, flow and transport model parameters

can be controlled (or at least exhaustively sampled) in a laboratory setting, which is

advantageous in that it enables the development of transport models with true (rather than

estimated) parameter values. These transport models could in turn be used in IFM and in

---

[2] Adapted from Shlomi, S., T. Sakaki, T. Illangasekare, and A.M. Michalak [2007] "Evaluation of Geostatistical Data Assimilation Methodologies for Estimating Groundwater Plume Distributions Using 3D Sand-Tank Tracer-Tests," in *Proceedings of the 37th Mid-Atlantic Industrial & Hazardous Waste Conference*, pp. 86-92, edited by G.A. Sorial and A. Bagtzoglou, Cincinnati, Ohio.

TrEK, and the estimates obtained using these methods could then be compared to laboratory measurements.

In contrast, in field experiments, the true values of parameters such as the hydraulic conductivity are never completely known. Developing or validating methods such as IFM or TrEK directly under these conditions would be problematic, because deviations from expected results may originate either from erroneous aquifer parameter estimates or from a fundamental problem with the estimation method. In field experiments, it is difficult to identify the true cause of such discrepancies. Thus, laboratory experiments present an attractive opportunity for developing and validating estimation methods.

As discussed in Chapter 3, IFM and TrEK are based on geostatistical inverse modeling [e.g. Snodgrass and Kitanidis, 1997], which takes advantage of existing knowledge about groundwater flow and transport at a polluted site. Michalak and Shlomi [2007] extended the applicability of the IFM approach to the estimation of plume distributions based on multiple monitoring events, including measurements taken after the time at which the plume is to be estimated. These studies used synthetic data to demonstrate the benefits of the proposed methods. Chapter 4 generalizes IFM and TrEK by accounting for uncertainty in the flow and transport model, and presents a methodology for simultaneously estimating flow and transport model parameters and the contaminant plume.

While these studies demonstrate the feasibility of IFM and TrEK for synthetic data, these new tools were not validated using data from laboratory or field sites, and the applications in Chapter 3 did not account for inherent transport model errors; rather, it was assumed that the transport model did not contribute to the uncertainty (except in one sensitivity analysis). This Chapter compares the performance of IFM to more traditional ordinary kriging (OK), using tracer tests performed in a carefully-controlled laboratory setting. Even though the magnitudes of the parameter and measurement errors in the laboratory data are low relative to a field setting, the spatial characteristics of the errors encountered in this experiment were more analogous to those that would be expected in the field. These measurement and transport model errors were analyzed and modeled using an empirical error covariance model.

2. Geostatistical Methodology

The main tool used to interpolate tracer measurements in this work is IFM, as described in detail in Chapter 3. In this Chapter, IFM relies on geostatistical inverse modeling to estimate the concentration of a tracer released into a sand tank as a function of time, denoted by $\mathbf{s}$. IFM then proceeds to use the transport model, together with an estimate of the release $\hat{\mathbf{s}}$, to estimate the entire plume distribution at any point in time (for example, the times at which concentration measurements were taken), and quantifies the uncertainties associated with this estimate. Depending on the structure chosen for $\mathbf{s}$, some form of stationarity is implicitly assumed. However, rather than a wide sense stationarity assumption, it can be generally assumed that the mean of $\mathbf{s}$ is either a constant or a sum of

given deterministic functions and that "the assumption of stationary fluctuations from the mean can be generalized to the assumption of stationary increments" [Kitanidis and Vomvoris, 1983].

Empirical error covariance matrix

IFM makes use of a flow-and-transport model for the aquifer to build a sensitivity matrix $\mathbf{H^*}$, the elements of which correspond to the sensitivity of measurements $\mathbf{z}$ to the state vector $\mathbf{s}$. The available measurements are modeled as linear with respect to the state:

$$\mathbf{z} = \mathbf{H^*s} + \boldsymbol{\varepsilon}$$

where $\boldsymbol{\varepsilon}$ is a random error with mean $\mathbf{0}$ and covariance matrix $\mathbf{R}$. In the previous Chapters, as well as in other past studies, $\mathbf{R}$ was modeled as a diagonal matrix of measurement errors, which implies that errors are not correlated. However, transport model inaccuracies are not likely to affect individual *points* in the aquifer independently, but rather whole *areas*; thus, these errors, which may also be represented by $\boldsymbol{\varepsilon}$, are likely to be spatially correlated. By repeatedly measuring concentrations at several points concurrently, and comparing these measurements to the corresponding model predictions $\mathbf{H^*s}$, a covariance matrix $\mathbf{R}$ can be derived from model-data residuals. An error covariance matrix $\mathbf{R}$ derived in this way would represent the observed correlation between errors observed at different locations within the aquifer.

This numerical approach to transport model uncertainty is somewhat similar to the sensitivity analysis presented in Chapter 3. However, this earlier work derived a non-

diagonal **R** *analytically* (from a cubic generalized covariance function), implying that the covariance between two measurement errors is a decreasing function only of their separation distance. In this Chapter, the error covariance is established *empirically*, based on observed mismatch between measured concentrations (resulting from an actual tracer release to the sand tank) and the corresponding results of a numerical flow-and-tranpsort model. This approach allows for a more accurate and complex spatial covariance structure.
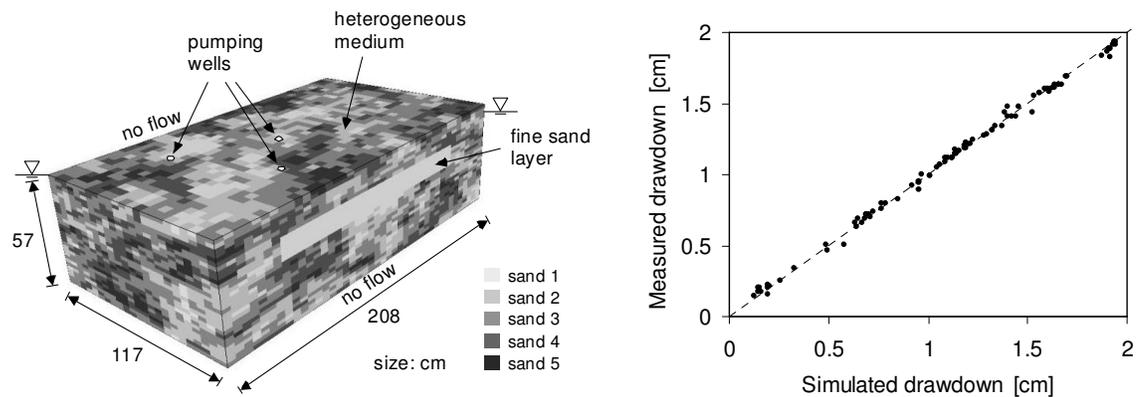
Figure 6.1 (a) 3-D test aquifer, and (b) Measured vs. simulated drawdowns during pumping at the center of the aquifer (after permeability values are calibrated). Source: Sakaki et al. [2006]

3. Experimental Setup

The experimental setup for the tracer test consists of a large [208.3 × 116.8 × 57.1 cm] three-dimensional test tank at the Center for Experimental Study of Subsurface Environmental Processes (www.CESEP.mines.edu). This synthetic aquifer is bounded on two opposite ends by constant head reservoirs, and by no-flow boundaries on the two

other sides, parallel to the flow. The aquifer was built (for earlier, unrelated experiments) by packing many "blocks" of five well-characterized test sands (Figure 6.1a), heterogeneously distributed in a horizontal square grid of 41×23 cells and 30 vertical layers. The hydraulic conductivities of the test sands followed an approximately lognormal distribution, comparable to the hydraulic conductivity distribution at a typical field site [Sakaki et al., 2006].

A series of flow tests performed in the tank prior to the present experiment were used to calibrate a numerical (MODFLOW) flow model, and then compared the measured drawdown distribution in the tank during a pumping test with simulated values (Figure 6.1b). The measured and model simulated drawdown agree closely [Sakaki et al., 2006].

The experiment was initiated by introducing a tracer into the aquifer through an injection well with a screen at layer 22, row 20, column 5. The time-varying injection rate (6.2) was designed to replicate the release history used by Skaggs and Kabala [1994] and others, which was also used in Chapter 3 (Equation 24). The tracer was injected to the aquifer using a computer-controlled syringe pump [Harvard Apparatus PHD2000], enabling an accurate simulation of a time-varying release. During the experiment, tracer breakthrough concentrations were measured using a fiber-optical fluorimeter [Hermes Messtechnik, Germany] which can continuously measure very low (1 ppb) in situ concentrations at 19 predefined locations. The exact plume could only be defined if the tracer concentrations were known at all locations in the tank. Unfortunately, this is not feasible with the direct use of the fluorimeter. 12 of these measurements, which were

148

close enough to the plume, were subsequently used for inverse modeling, IFM and ordinary kriging. The measurements were taken in five different layers (21 – 25), but only layer 22 is reconstructed, because that is where most of the measurements are; there is not enough data to reconstruct other layers, as the vertical correlation length is very short. Due to the direction of flow, the variation of tracer concentration is not isotropic. The anisotropy ratio was assumed to be equal to the ratio of longitudinal to transverse and vertical dispersivities [Chilès and Delfiner, 1999], in this case 10.

In order to characterize the variance-covariance structure $\mathbf{R}$ of the model-data mismatch errors, both calibrated observations and model data were recorded at the measurement locations at ten-second intervals over the course of the experiment. The residuals between these two were treated as a large sample from a multivariate distribution, and were then accordingly analyzed to produce a variance-covariance matrix $\mathbf{R}$, which was consequently used in the plume estimation process.
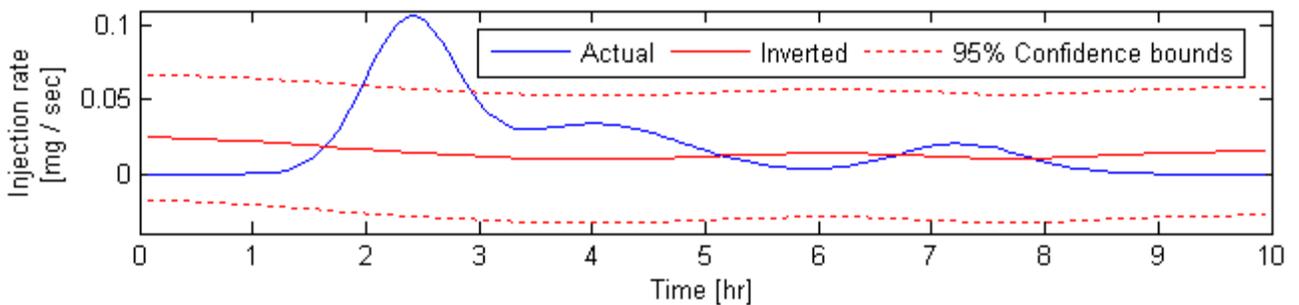
Figure 6.2 Source release histories: the actual release was programmed in the syringe pump which released the tracer; the recovered release and its 95% confidence bounds were calculated using geostatistical inverse modeling.

4. Results

The plume which developed following the tracer injection was modeled numerically using the transport code MT3DMS [Zheng and Wang, 1999], which was calibrated using the measured tracer concentrations, under the assumption that the peak modeled concentrations at each node are indeed the true peak concentrations. A horizontal section through the simulated plume 18 hours after the start of the release is presented in Figure 6.3. Because the flow field was at steady state (at least theoretically) and the aquifer properties were exactly known in the synthetic aquifer, it is reasonable to expect this transport model to be relatively precise. However, as opposed to the previous Chapters, there is no way to depict the actual plume, and Figure 6.3 would be the best representation thereof.
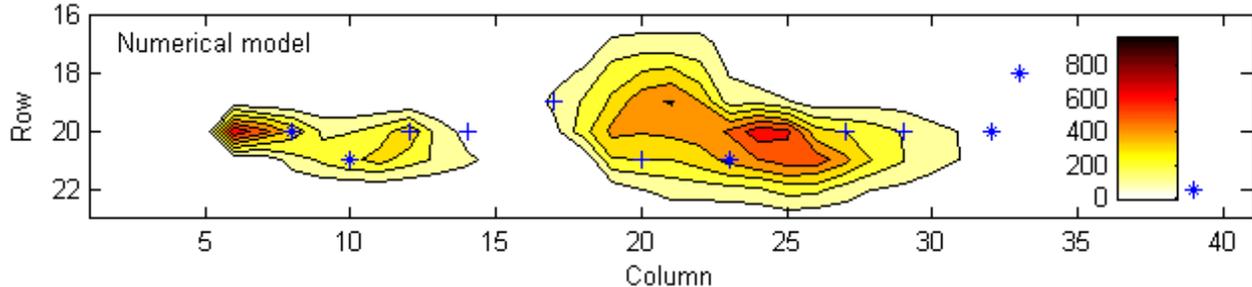
Figure 6.3 Model-simulated plume 18 hours after start of release. Stars denote actual measurement locations in layer 22; Plus signs denote actual measurements in other layers. Concentrations are in ppb.

Figure 6.4 depicts the **R** matrix derived from the residuals of model data vs. observations. Each element (i, j) represents the covariance of two measurement errors, from node i and node j. Note that while there were only 12 measurement locations, the concentrations from three different times were used, hence the dimensions of this variance-covariance are 36×36. Most of its elements are close to 0, implying virtually no cross-correlation

between model-data mismatch at different measurement locations. However, one location (the eighth row/column) has a strong negative correlation with the second and third measurement locations. The diagonal represents measurement errors, most of which were relatively small. Only two measurement locations had very high measurement errors, as represented by the black squares corresponding to the first and tenth measurements.
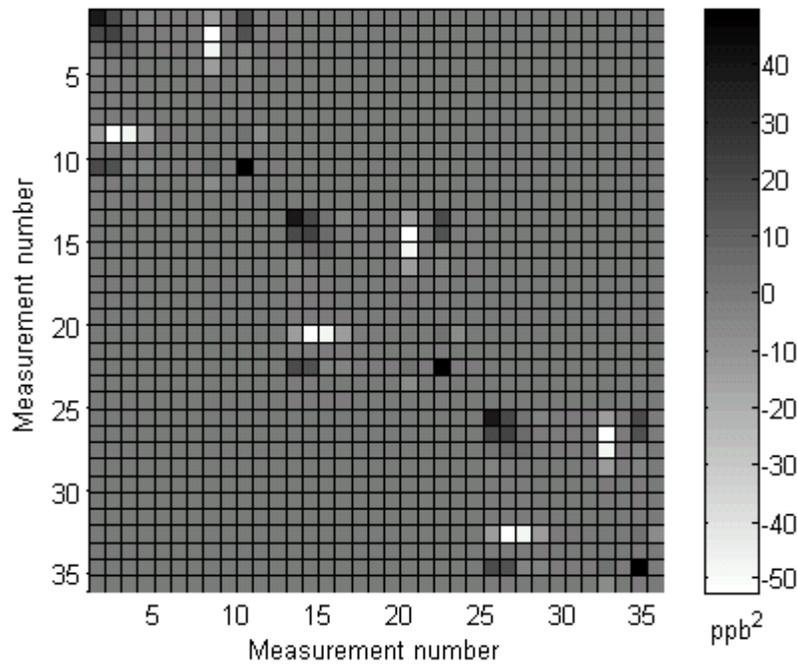
Figure 6.4 Empirical variance-covariance matrix used as model-data mismatch matrix (R) for IFM. All values are in ppb[2].

151
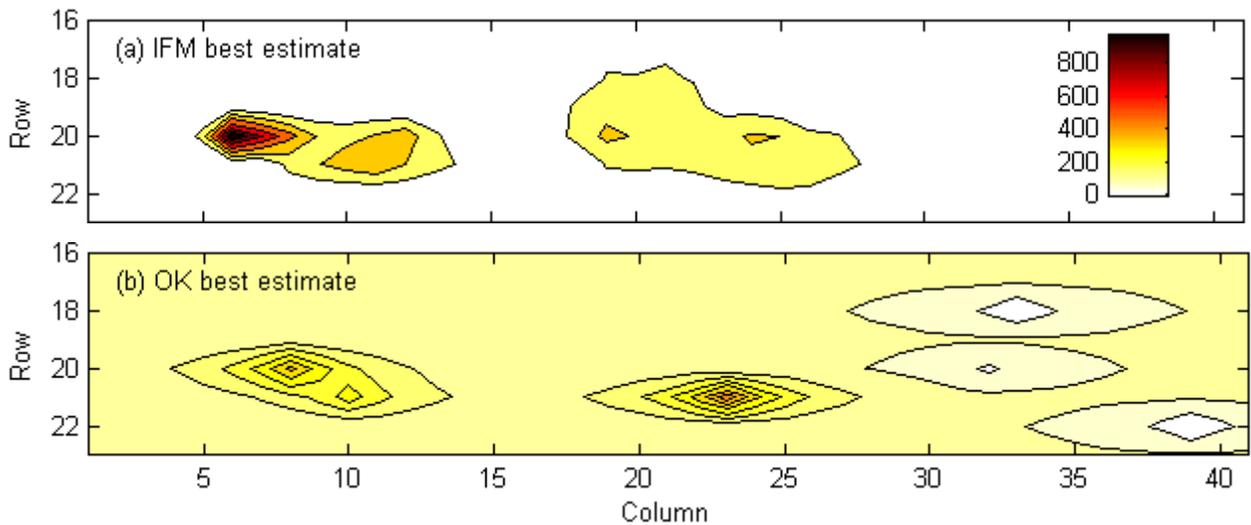
Figure 6.5 Plume best estimates for T = 18 hours. (a) best plume estimate obtained using IFM; (b) best plume estimate obtained using ordinary kriging. All concentrations are in ppb.

Figure 6.5(a) shows the reconstruction of the plume 18 hours after the start of the release using IFM, with measurements taken from the test aquifer 10, 14, and 18 hours after the release started. The release history was assumed to have a Gaussian covariance,

$$C_s(h) = \sigma^2 \exp\left(-\frac{h^2}{l^2}\right)$$

with length parameter 13.2 hours, and variance 0.046 hours^2. All measurements were jointly used to recover the source release history (Figure 6.2), which, in turn, was reintroduced into the model to complete the IFM process. Figure 6.4(b) shows the reconstructed plume using ordinary kriging (implying a weak form of stationarity) with an exponential covariance function

$$C_p(h) = \sigma^2 \exp\left(-\frac{h}{l}\right)$$

and parameters $\sigma^2 = 13.4$ ppb^2 and $l = 7.5$ cm. These parameters were optimized using the full plume distribution, as give by the numerical model (Figure 6.3). Note that only the 12 measurements taken 18 hours after the release were used in kriging, as measurements at other times cannot be used directly in spatial kriging.
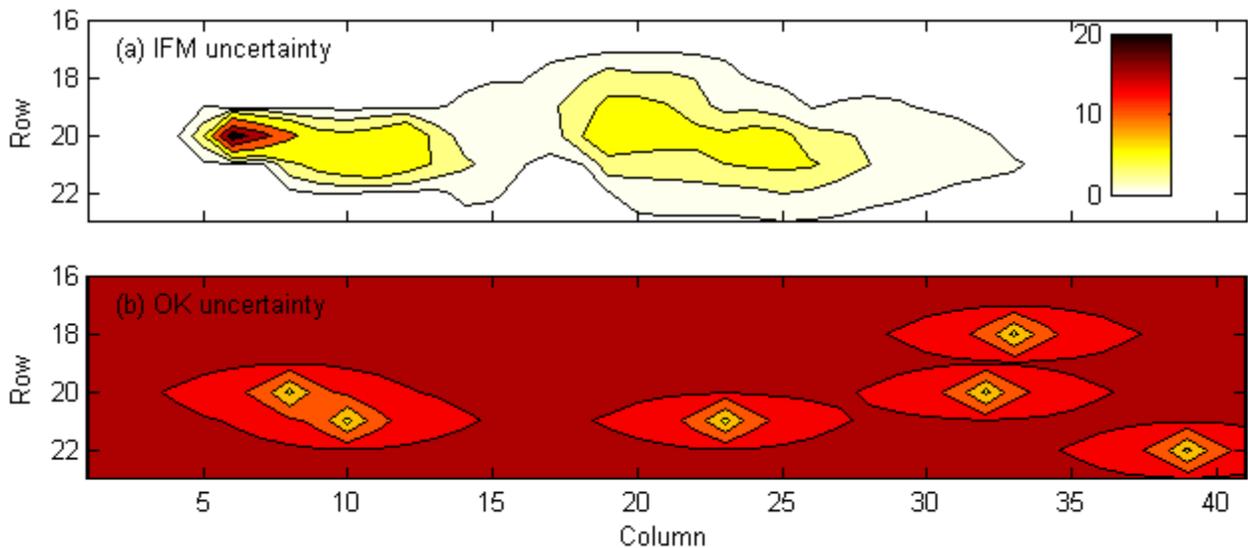
Figure 6.6 Estimation uncertainties. (a) uncertainty using IFM; (b) uncertainty using ordinary kriging (standard deviations in ppb for both).

Figure 6.6 shows the uncertainties associated with IFM and kriging estimates. As was noted in Chapter 3, IFM (Figure 6.6(a)) has high uncertainty in high concentration areas, as these points are typically the most sensitive to the release **s**. Conversely, kriging (Figure 6.6(b)) has low uncertainties at measurement locations, and higher uncertainty at the periphery.

5. Discussion

In the tracer release experiment described above, a numerical model was used to depict the "real" plume, based on the true release history of a tracer, and two geostatistical methods, ordinary kriging and IFM, were used to reconstruct that plume using measurements taken at 12 locations in the test aquifer, 10, 14, and 18 hours after tracer was injected to the aquifer. In the ordinary kriging and IFM applications, the release history was assumed unknown but the hydraulic conductivity distribution of the aquifer was known. The performance of these methods was evaluated by their ability to reproduce the plume, and to provide estimates with low uncertainty.

As can be seen in Figure 6.2, the exact shape of the source release history could not be recovered using the data, because measurements were sparse and were taken a long time after the release. Despite the uncertainty associated with the recovered release history, IFM was able to reconstruct the actual plume more accurately than OK. While the estimate does not reproduce some of the hot spots observed in the simulated plume based on the true release history, the general shape of the plume is similar. The observed differences can be attributed to measurement or calibration errors in the laboratory setup and to the impact of the inaccuracies in the recovered release history. IFM uses measurement data to reconstruct the release history; and subsequently uses the transport model to constrain the simulated plume. However, as was also seen in Chapters 3 and 4, the constraint imposed by the transport model is much less effective for recovering the release history (Figure 6.2) than for projecting it forward in time (Figures 6.5a and 6.6a). As discussed in Chapter 3, groundwater contaminant transport is a dispersive process,

and, consequently, the degree to which the contamination history can be inferred from downgradient measurements is limited [e.g. Skaggs and Kabala, 1994]. Conversely, the results of the forward model are relatively insensitive to the details of the release history, and the release history need not be reproduced exactly for the plume to be recovered.

The **R** covariance matrix used in this experiment (Figure 6.4) jointly represents measurement error, calibration error, and transport model error, but cannot distinguish among them, for lack of an absolute reference. The diagonal **R** elements represent the discrepancy of model predictions and actual observations. These errors exhibited spatial covariance, with adjacent locations often being positively correlated. Conversely, some of the off-diagonal R_{ij} ($i \neq j$) elements were negative, implying that, assuming an unbiased model, a positive residual at location i, would tend to be associated with a negative residual at location j. This could occur, for example, in the case of a model error (e.g. in flow direction) which would cause the predicted contamination to be shifted from a location i to a location j.

The kriging estimation relies only on the 12 measurements taken at the time of estimation, and on the assumed correlation structure of the contaminant field. The states of the plume at earlier times, along with the physics governing groundwater flow and contaminant transport, affect the plume distribution at $T = 18$ hours. However, neither transport information nor measurements at earlier times are explicitly considered in the kriging estimation model. Thus, with significantly less input data, the kriging estimates have two main disadvantages relative to IFM. First, the kriging model does not explicitly

incorporate transport information, relying solely on point measurements. Second, only measurements at the time of estimation (18 hours after tracer release) can be used in the analysis, while in IFM, measurements taken 10 and 14 hours after the start of the release were also used. In addition, the short correlation length (just short of five grid cells in the longitudinal direction; much shorter in the transverse direction) precludes adequate estimates in many areas throughout the domain. As a result of these factors, contaminant concentration estimates in most of the domain are highly uncertain, and could probably only be used for making rough assessments of the plume distribution.

6. Conclusions

This Chapter presents the first application of IFM to laboratory measurements, specifically in assimilating concentration data acquired at several distinct times with available flow and transport information to estimate the distribution of a tracer plume. The IFM approach was shown to have several advantages relative to ordinary kriging, and could have a wide array of applications, such as monitoring network and remediation design.

In past applications of contaminant source identification [e.g. Kitanidis, 1995, Snodgrass and Kitanidis, 1997, and Michalak and Kitanidis, 2004a, b] and IFM [Shlomi and Michalak, 2007; Michalak and Shlomi, 2007], model uncertainty was not taken into account. While a deterministic transport model [MT3DMS] was also applied here, the model-data mismatch $\mathbf{R}$ did not simply contain measurement error variances; rather, covariances between errors were estimated using model-data residuals. Although this

approach still does not fully characterize the errors associated with transport model uncertainty, it begins to shed light on the impact of model errors on the uncertainty associated with the estimated field. Note that the approach taken here for estimating the error covariance structure required knowledge of the release history, and could therefore not be applied in a case where this release is unknown. Assessing the impact of transport model errors on IFM estimates is central to the field-scale applicability of the approach, because models of real aquifers have many unknowns. Nonetheless, once there exists a deterministic model, using a numerically-derived **R** matrix, as was done in this work, may reflect the resulting uncertainty in model predictions. It should be stressed that it was possible to construct such a covariance matrix in this experiment due to the availability of repeated concentration measurements at each location.

Chapter 7
Conclusions and Future Directions

1. Summary and Conclusions

Groundwater contaminant plume interpolation or estimation is essential for numerous functions, such as monitoring, global contaminant mass estimation, network design, design of water distribution systems, remediation, and contaminant source identification and parameter estimation. An accurate representation of groundwater contaminant plume distributions is critical to monitoring the safety of water supplies, estimating the extent of groundwater contamination and designing effective remediation strategies.

This work has contributed to improving plume interpolation capabilities by providing novel tools that integrate inverse modeling with deterministic flow-and-transport models into traditional geostatistical methods. The major methodologies that were developed within this research are Inverse/Forward Modeling and Transport-Enhanced Kriging (**Chapter 3**). These methods take advantage of both geostatistical tools, such as kriging, and groundwater flow-and-transport models, to better estimate contaminant plumes in aquifers.

In **Chapter 4**, these methods were generalized to allow simultaneous estimation of flow or transport model parameters. The importance of this step lies in the fact that parameters such as the hydraulic conductivity are almost never measured exhaustively, such that in practice the usefulness of IFM and TrEK would be limited without the ability to deal with these kinds of uncertainties. While there exist many methods for estimating parameters such as hydraulic conductivity, the simultaneous estimation method presented in Chapter 4 has at least one important advantage, namely the ability to use a larger superset of data, of several types, to estimate the required unknowns, and to directly account for the uncertainty associated with flow and transport parameters in the estimation uncertainty of the contaminant plume distribution.

In **Chapter 5**, a groundwater monitoring network design methodology compatible with IFM and TrEK was presented. The approach taken is based on established methods, but was specifically tailored for IFM and TrEK. This is important because well selection methods must be based on the plume estimation methods to be used at a site. As the shape of an estimated contaminant plume is highly dependent on the specific measurement sites used in the estimation, selecting the optimal wells for this purpose can dramatically improve the quality of interpolation. The presented well-selection tool indicates which wells should be sampled at any point in time, in order to reconstruct the contaminant plume with lower uncertainty relative to other existing methods.

While IFM and TrEK are yet to be tested in a field setting, they have been evaluated using experimental data from a physical sand-tank model. Results from these experiments

are presented in **Chapter 6**, along with an empirical approach for estimating errors in the flow-and-transport model. Although only a few experiments were performed, and many technical challenges have accompanied the calibration and data acquisition phases, the results nevertheless point to a significant advantage of using IFM over kriging, whenever a transport model is available.

Overall, the methodologies developed as part of this dissertation enable better interpolation through integrated use of various types of data and information. They are expected to allow practitioners and stakeholders within the water resource management community to achieve monitoring objectives faster relative to current methods, and using sparser monitoring. As a result, remediation costs and various other related expenses may also be reduced. Finally, as a byproduct, the research carried out within this work may benefit a wider scientific community, by providing generally-applicable estimation tools, such as the capability to account for transport model uncertainties in geostatistical inverse modeling.

2. Recommendations for Future Research

The methods developed within this dissertation provide a firm basis for a new family of robust geostatistical tools that could be further developed, evaluated, and implemented in various applications. Overall, IFM and TrEK combine kriging with transport information. Therefore, obvious extensions could involve altering these two components. Regarding the former, there are numerous existing types of kriging that provide solutions to a wide variety of situations. Some of these could easily be incorporated into the methods

developed here, while others may prove more challenging to integrate into the existing framework. Similarly, while relatively simple flow-and-transport models were used throughout this work, more sophisticated models could potentially be used in IFM and TrEK, if such models exist for the aquifer. The following sections provide several specific suggestions regarding the ways in which the methods developed in this work could be improved, generalized, and validated.

2.1 Optimization of variogram (or covariance function) parameters

Throughout this dissertation, optimization of variogram (or covariance function) parameters was performed using established geostatistical methods, and was not discussed explicitly. While this helped focus attention on the development of the new monitoring methods, this type of parameter estimation is an important part of the overall monitoring process in geostatistical theory, and should be given proper consideration. There exist many proven methods for estimating these parameters, and each carries with it a set of advantages and disadvantages, which often depend on the specifics of the application at hand. These methods could be used as-is, as has been the case in this work, or modified and improved to work with flow-and-transport models, in an IFM-like or TrEK-like framework.

Regarding network design, for example, rather than focusing on designs for minimizing the average or the maximum prediction variance over all prediction points, various studies concentrate on how to construct sampling designs from which the model

parameters can be optimally estimated. Early studies proposed algorithms for finding the optimal combination of sampling points in order to estimate the sample semivariogram, used in geostatistics for analyzing the second moment structure of a spatial stochastic process [Russo, 1984; Russo and Jury, 1988; Warrick and Myers, 1987]. Zimmerman and Homer [1991], Muller and Zimmerman [1999], and Bogaert and Russo [1999] extended these ideas to consider not only estimation of the sample semivariogram but also parametric estimation of the semivariogram model.

One can therefore identify two groups of design criteria, focusing either on spatial prediction or parameter estimation. Lark [2002] illustrated that errors in the model parameters result in errors in the estimated prediction variances. This means that if the primary goal of a monitoring network is to compute spatial predictions, some combination of the two groups of design criteria should be used. This ensures that efficient spatial predictions can be computed while taking proper care of the uncertainties of the model parameters. Hence in this case an efficient design should consist of some sampling points allocated for estimating model parameters and some for computing spatial predictions [Muller, 2001; Martin, 2001].

A way to combine designs for computing spatial predictions with designs for estimating model parameters is to compute spatial predictions and associated prediction variances using what is often referred to as a full Bayesian approach [e.g. Diggle et al., 2003; Handcock and Stein, 1993; Le and Zidek, 1992]. In this approach, the uncertainty of the model parameters, described by the posterior distribution, is automatically incorporated

in the predictions. This means that good parameter estimates, i.e. a narrow posterior distribution, leads to more efficient spatial predictions. Modifying these types of methods to include flow-and-transport information, in a manner analogous to the development of IFM and TrEK, would further improve the quality of interpolation, by potentially displaying better best estimates and lower uncertainties.

2.2 Temporal dimension

Most of the work in this dissertation implicitly assumes that a plume estimate is only needed at one point in time. While historical measurements were integrated into the estimation model, there was no attempt at long-term optimization (or planning) of sampling schedules. Of course, this does not imply that these are less important issues, but the focus in this dissertation was on the specifics of plume estimation and well-selection.

Michalak and Shlomi [2007] first showed the ability of IFM and TrEK to not only seamlessly integrate data from various points in time into the estimation process, but also to estimate a plume at any point in time, just as easily. However, of the two components of TrEK estimation framework, only one – the flow-and-transport model component – related different points on the temporal axis; the second component of TrEK, which takes into account the spatial covariance, did not consider any temporal or spatiotemporal correlation. This latter component only considered the spatial correlation at the time of estimation.

Many methods already exist for optimizing long-term monitoring network design, both in the spatial and in the temporal dimension. Likewise, spatiotemporal geostatistics (which considers the spatial and temporal coordinates of an event, such as a groundwater contamination measurement) has been practiced for years [e.g. Christakos, 2000].

Both the inverse-modeling component and the spatial-correlation component of the methods developed in the preceding Chapters are flexible, and could accommodate various types of measurements, variograms, and basis functions (in the case of universal kriging). It is thus feasible to integrate some form of the existing spatiotemporal theory with IFM and TrEK, wherever this may be appropriate. Depending on the application, inverse modeling, for example, could potentially be performed using a spatiotemporal variogram (if one exists or could be optimized), instead of the spatial **or** temporal covariance functions that were used here. Similarly, what we have considered the spatial-correlation component of TrEK, could also use a spatiotemporal covariance function, whenever such a correlation is available. This could be the case when a contaminant plume is to be estimated at several points in time. This type of modification could potentially enhance the performance of the spatial-correlation component of these methods (which sometimes proved to be the "weak link," in the examples shown throughout this work, as it did not provide a constraint as strong as the transport model). More importantly, it can improve the overall estimation quality of IFM and TrEK, in cases where the conditions are right to implement such modifications, e.g. if more than one type of measurements existed and cross-variograms could be modeled.

2.3 Nonnegativity constraints

Contaminant concentrations, which are the primary unknown sought throughout this research, cannot be negative. However, the methods used here did not impose any limitations on the results, and some conditional realizations, confidence intervals, and even best estimates reached negative values at certain points. While the examples presented in this research did not demonstrate any significant deficiencies originating from lack of nonnegativity constraints, other applications of these types of methods could potentially suffer from the consequences of letting variables assume negative values. This could produce, for example, results which are not physically feasible.

Restricting variables to the nonnegative domain serves two purposes: primarily, it would prevent situations in which the best estimate for a quantity such as concentration is negative. In addition, it could more faithfully represent the entire distribution by properly shifting all of the probability mass to be nonnegative.

In the geostatistical literature, there are several published methods that constrain nonnegativity [e.g. Michalak, 2008; Michalak and Kitanidis, 2003, 2005]. It is hard to say what would be the optimal method to use with IFM or TrEK, because that would depend on the specifics of the application. For example, one can constrain either the inverse model or the forward model, and arguments could be made for both sides.

Finally, other constraints, such as upper limits (which might represent, for example, the solubilities of contaminants in water), could also be imposed on the methods presented

here, for more faithfully representing realistic realizations, and thus improving simulation results.

2.4 Laboratory tests and field demonstrations

IFM and TrEK have thus far primarily been applied to synthetic data. This has enabled to test these methods for several types of data in various conditions. While several sand tank experiments have also been performed, the results obtained are somewhat problematic because of potential calibration issues.

Actual field data has never been used in IFM or TrEK, and should provide an interesting validation of these methods. This could be done either with existing data (from sampling that has already taken place), or in a future field demonstration. Like the synthetic demonstrations in this study, comparisons should be planned against existing methods, such as kriging or other forms of interpolations. Flow-and-transport models might also provide results for comparisons, although one must be careful if using these same models within IFM and TrEK, The dependence of these methods on the flow-and-transport models may limit the validity of this kind of comparison. Thus, a good comparison has to be designed in a way that would differentiate the features of the flow and transport model from those of the geostatistical estimation methods which depend on it, IFM and TrEK.

2.5 Kriging on a flow net

Several examples throughout this study, starting with Chapter 3, indicated that the flow model may be one of the most important components in IFM and TrEK. Using these methods, even with very few measurement data, provided excellent results: the best estimate obtained closely resembled the plume in most examples, and the uncertainties were very low for the most part, orders of magnitude lower in comparison to kriging.

These results may suggest that the transport model may not be completely necessary for performing these kinds of interpolations, and that the flow model may be sufficient to constrain the estimation, especially if a transport model is not available. Shlomi [2007] demonstrated that this is indeed possible, by choosing a more natural coordinate system, in the form of a curvilinear flow net: a grid of stream and equipotential lines, which are everywhere perpendicular to each other. Kriging could be performed on the flow net (e.g. the abscissa could run along stream lines, and the ordinate along equipotential lines), by replacing the actual distances between points by corresponding distances on the flow net.

An interesting application could be formed by incorporating the temporal dimension (see Section 2.2 above) in this type of model. It might actually be easier to relate it to the spatial dimensions on a flow net than it would be on a rectilinear grid (especially under steady state conditions, where a particle would cross equipotential lines at equal intervals), making spatiotemporal cross-variograms relatively easy to model. One unit of space (the distance between equipotential lines along the x-coordinate) would then be equivalent to, or perhaps even interchangeable with, a unit of time (the temporal interval

between crossing these lines). This may allow inclusion of various points in the space/time continuum in one unified correlation structure, making the use of such correlation in further analysis (e.g. in an IFM or a TrEK framework) easier to implement.

2.6 Above ground

All of the research presented in this dissertation pertains directly to, has been developed for, and was demonstrated on, groundwater contamination. However, geostatistical methods are being used for a myriad of applications, including mining, geophysics, atmospheric science, and even finance. While the groundwater flow-and-transport equations (see Chapter 1) are only valid for subsurface contaminants, they belong to more general families of equations that also solve problems in various other fields. Potentially, then, these methods could impact many other realms, and improve estimation for engineers, scientists, and other practitioners from a wide variety of backgrounds.

References

1. (Nrc), N. R. C. (2002), *Opportunities to improve the u.S. Geological survey national water quality assessment program*, NATIONAL ACADEMY PRESS, Washington, DC.
2. Ahlfeld, D. P., and M. Heidari (1994), Applications of optimal hydraulic control to groundwater systems, *Journal of Water Resources Planning and Management-Asce*, *120*(3), 350-365.
3. Ahlfeld, D. P., and A. E. Mulligan (2000), *Optimal management of flow in groundwater systems*, Academic Press.
4. Ahmed, S., and G. De Marsily (1987), Comparison of geostatistical methods for estimating transmissivity using data on transmissivity and specific capacity, *Water Resources Research*, *23*(9), 1717-1737.
5. Ahmed, S., and G. De Marsily (1993), Cokriged estimation of aquifer transmissivity as an indirect solution of the inverse problem: A practical approach, *Water Resources Research*, *29*.
6. Anderman, W. H., and G. Maritim (1986), Effect of public sewers on watershed contamination, *Journal of Environmental Health*, *4*(2), 81-84.
7. Asce (2003), Long-term groundwater monitoring - the state of the art, 116 pp, ASCE-EWRI.
8. Atmadja, J., and A. C. Bagtzoglou (2001), State of the art report on mathematical methods for groundwater pollution source identification, *Environmental Forensics*, *2*(3), 205-214.
9. Aziz, J. J., et al. (2000), Monitoring and remediation optimization system (maros) software user's guide, edited, Air Force Center for Environmental Excellence, Brooks Air Force Base, TX.
10. Bagtzoglou, A. C., et al. (1991), Probabilistic simulation for reliable solute source identification in heterogeneous porous media, paper presented at Water Resources Engineering Risk Assessment, Springer-Verlag.
11. Bagtzoglou, A. C., et al. (1992), Application of particle methods to reliable identification of groundwater pollution sources, *Water Resources Management*, *6*(1), 15-23.
12. Bear, J., and A. Verruijt (1987), *Modeling groundwater flow and pollution*, D. Reidel, Dordrecht, Holland.
13. Benjemaa, F., et al. (1994), Multivariate geostatistical design of groundwater monitoring networks, *Journal of Water Resources Planning and Management-Asce*, *120*(4), 505-522.

14. Bice, L. A., et al. (2000), Source water assessment using geographic information systems, EPA.
15. Boeckenhauer, R. K., et al. (2000), Statistical estimation and visualization of ground-water contamination data, EPA - Office of Research and Development.
16. Bogaert, P., and D. Russo (1999), Optimal spatial sampling design for the estimation of the variogram based on a least squares approach, *Water Resources Research*, *35*(4), 1275-1289.
17. Boulding, J. R., and J. S. Ginn (2004), *Practical handbook of soil, vadose zone, and ground-water contamination: Assessment, prevention, and remediantion*, Second ed., Lewis Publishers.
18. Burnham, K. P., and D. R. Anderson (2002), *Model selection and inference*, Springer.
19. Butera, I., and M. G. Tanda (2003), A geostatistical approach to recover the release history of groundwater pollutants, *Water Resources Research*, *39*(12).
20. Carrera, J., et al. (1984), A method for optimal observation network design for groundwater-management, *Journal of Hydrology*, *73*(1-2), 147-163.
21. Carrera, J., and F. Szidarovszky (1985), Numerical comparison of network design algorithms for regionalized variables, *Appl. Math. Comput.*, *16*(3), 189-202.
22. Carrera, J., and S. P. Neuman (1986), Estimation of aquifer parameters under transient and steady-state conditions .1. Maximum-likelihood method incorporating prior information, *Water Resources Research*, *22*(2), 199-210.
23. Cassiani, G., and M. A. Medina (1997), Incorporating auxiliary geophysical data into ground-water flow parameter estimation, *Ground Water*, *35*(1), 79, doi:10.1111/j.1745-6584.1997.tb00063.x.
24. Chang, L. F., et al. (2005), Optimal observation network design for parameter structure identification in groundwater modeling, *Water Resources Research*, *41*(3).
25. Chiles, J.-P., and P. Delfiner (1999), *Geostatistics: Modeling spatial uncertainty*, John Wiley & Sons, Inc.
26. Christakos, G. (1990), A bayesian maximum-entropy view to the spatial estimation problem, *Mathematical Geology*, *22*(7), 763-777.
27. Christakos, G. (2000), *Modern spatiotemporal geostatistics*, 304 pp., Oxford University Press, New York, NY.
28. Christakos, G. (2005), Methodological developments in geophysical assimilation modeling, *Reviews of Geophysics*, *43*(2).
29. Cohon, J. L. (2004), *Multiobjective programming and planning*, Courier Dover Publications.
30. Commission on Geosciences, E., and Resources (Cger) (2000), *Research needs in subsurface science*, National Academy Press, Washington, DC.
31. Committee on Techniques for Assessing Ground Water Vulnerability, W. S. A. T. B. (Ed.) (1993), *Ground water vulnerability assessment: Predicting relative contamination potential under conditions of uncertainty*, NATIONAL ACADEMY PRESS, Washington, DC.
32. Company, R. M. (1990), *Handbook of ground water development*, Wiley-IEEE.

33. Cooper, R. M., and J. D. Istok (1988), Geostatistics applied to groundwater contamination .2. Application, *Journal of Environmental Engineering-Asce*, *114*(2), 287-299.
34. Dagan, G., et al. (1996), Impact of concentration measurement upon estimation of flow and transport parameters: The lagrangian approach, *Water Resources Research*, *32*(2), 297-306.
35. Darcy, H. (1856), *Les fontaines publiques de la ville de dijon: Exposition et application des principes a suivre et des formules a employer dans les questions de distribution d'eau*, LIBRAIRE DES CORPS IMPÉRIAUX DES PONTS ET CHAUSSÉES ET DES MINES, Paris.
36. Delhomme, J. P. (1978), Kriging in the hydrosciences, *Advances in Water Resources*, *1*(5), 251-266.
37. Delleur, J. W. (1998), *The handbook of groundwater engineering*, CRC Press.
38. Diggle, P. J., et al. (1998), Model-based geostatistics, *Journal of the Royal Statistical Society Series C-Applied Statistics*, *47*, 299-326.
39. Diggle, P. J., et al. (2003), An introduction to model-based geostatistics, in *Spatial statistics and computational methods*, edited by J. Møller, Springer Verlag.
40. Döll, P., and K. Fiedler (2008), Global-scale modeling of groundwater recharge, *Hydrol. Earth Syst. Sci.*, *12*(3), 863-885.
41. El Idrysy, E. H., and F. De Smedt (2007), A comparative study of hydraulic conductivity estimations using geostatistics, *Hydrogeology Journal*, *15*(3), 459-470.
42. Feehley, C. E., et al. (2000), A dual-domain mass transfer approach for modeling solute transport in heterogeneous aquifers: Application to the macrodispersion experiment (made) site, *Water Resources Research*, *36*(9), 2501,.
43. Fetter, C. W. (2001), *Applied hydrogeology*, Fourth ed., Prentice Hall, Upper Saddle River, NJ 07458.
44. Figueira, R., et al. (2001), Use of secondary information in space-time statistics for biomonitoring studies of saline deposition, *Environmetrics*, *12*(3), 203, doi 210.1002/env.1457.
45. Freshley, M. D., et al. (2002), Groundwater protection program science and technology summary description, Size: vp. pp.
46. Galili, E., and Y. Nir (1993), The submerged pre-pottery neolithic water well of atlit-yam, northern israel, and its palaeoenvironmental implications, *The Holocene*, *3*(3), 265-270.
47. Gorelick, S. M., et al. (1983), Identifying sources of groundwater pollution - an optimization approach, *Water Resources Research*, *19*(3), 779-790.
48. Grabovnicov, V., et al. (2003), Laboratory research results of high-density solution migration in fresh groundwaters, *Transport in Porous Media*, *53*(2), 123-143.
49. Graham, W., and D. Mclaughlin (1989), Stochastic-analysis of nonstationary subsurface solute transport .2. Conditional moments, *Water Resources Research*, *25*(11), 2331-2355.

50. Graham, W., and D. Mclaughlin (1989), Stochastic-analysis of nonstationary subsurface solute transport .1. Unconditional moments, *Water Resources Research*, *25*(2), 215-232.

51. Grasso, D., et al. (1994), Closed-loop recovery of site remediation gas-phase contaminants, *Journal of the Air & Waste Management Association*, *44*(8), 983-988.

52. Graybill, F. A. (2000), *Theory and application of the linear model*, Duxbury.

53. Gutjahr, A. L., et al. Stochastic analysis of spatial variability in subsurface flows 2. Evaluation and application, *Water Resources Research*, *14*.

54. Handcock, M. S., and M. L. Stein (1993), A bayesian-analysis of kriging, *Technometrics*, *35*(4), 403-410.

55. Hardisty, P. E., and E. Ozdemiroglu (2005), *The economics of groundwater remediation and protection*, 336 pp., CRC Press.

56. Hernandez, A. F., et al. (2006), Inverse stochastic moment analysis of steady state flow in randomly heterogeneous media, *Water Resources Research*, *42*(5).

57. Herrera, G. S., and G. F. Pinder (2005), Space-time optimization of groundwater quality sampling networks, *Water Resources Research*, *41*(12).

58. Herzog, B., et al. (1991), Ground-water sampling, in *Practical handbook of ground-water monitoring*, edited by D. M. Nielsen, Lewis Publishers, Chelsea, MI.

59. Hetzel, F., et al. (2008), Groundwater and climate change: Challenges and possibilities, edited by G. F. M. F. E. C. A. Development.

60. Hill, M. C., et al. (2000), Modflow-2000, the u.S. Geological survey modular ground-water model -- user guide to modularization concepts and the ground-water flow process, U.S. Geological Survey Open-File, 121 pp, USGS.

61. Hoeksema, R. J., and P. K. Kitanidis (1984), An application of the geostatistical approach to the inverse problem in two-dimensional groundwater modeling, *Water Resources Research*, *20*(7), 1003-1020.

62. Hsueh, Y. W., and R. Rajagopal (1988), Modeling ground-water quality sampling decisions, *Ground Water Monitoring and Remediation*, *8*(4), 121-134.

63. Johnson, V. M., et al. (1996), Reducing the sampling frequency of groundwater monitoring wells, *Environmental Science & Technology*, *30*(1), 355-358.

64. Jones, N. L., et al. (2003), A comparison of three-dimensional interpolation techniques for plume characterization, *Ground Water*, *41*(4), 411, doi:410.1111/j.1745-6584.2003.tb02375.x.

65. Journel, A. G., and C. V. Deutsch (1997), Rank order geostatistics: A proposal for a unique coding and common processing of diverse data., paper presented at 5th International Geostatistics Conference, Kluwer Academic Publishers, Wollongton, Australia.

66. Kabala, Z. J., and T. H. Skaggs (1998), Minimum relative entropy inversion: Theory and application to recovering the release history of a groundwater contaminant, *Water Resources Research*, *34*(8), 2077-2079.

67. Kitanidis, P. K., and E. G. Vomvoris (1983), A geostatistical approach to the inverse problem in groundwater modeling (steady-state) and one-dimensional simulations, *Water Resources Research*, *19*(3), 677-690.
68. Kitanidis, P. K., and R. W. Lane (1985), Maximum-likelihood parameter-estimation of hydrologic spatial processes by the gauss-newton method, *Journal of Hydrology*, *79*(1-2), 53-71.
69. Kitanidis, P. K. (1995), Quasi-linear geostatistical theory for inversing, *Water Resources Research*, *31*(10), 2411-2419.
70. Kitanidis, P. K. (1996), On the geostatistical approach to the inverse problem, *Advances in Water Resources*, *19*(6), 333-342.
71. Kitanidis, P. K., and K. F. Shen (1996), Geostatistical interpolation of chemical concentration, *Advances in Water Resources*, *19*(6), 369, doi:310.1016/0309-1708(1096)00016-00014.
72. Kitanidis, P. K. (1997), *Introduction to geostatistics: Applications to hydrogeology*, Cambridge University Press.
73. Kollat, J. B., and P. M. Reed (2006), Comparing state-of-the-art evolutionary multi-objective algorithms for long-term groundwater monitoring design, *Advances in Water Resources*, *29*(6), 792-807.
74. Lark, R. M. (2002), Optimized spatial sampling of soil for estimation of the variogram by maximum likelihood, *Geoderma*, *105*(1-2), 49-80.
75. Le, N. D., and J. V. Zidek (1992), Interpolation with uncertain spatial covariances - a bayesian alternative to kriging, *J. Multivar. Anal.*, *43*(2), 351-374.
76. Lindberg, J. W. (2006), Results of groundwater monitoring for the 300 area process trenches, Size: PDFN pp.
77. Liu, W. (2003), Development of gradient-enhanced kriging approximations for multidisciplinary design optimization, 177 pp, Notre Dame, Indiana.
78. Loaiciga, H. A. (1989), An optimization approach for groundwater quality monitoring network design, *Water Resources Research*, *25*(8), 1771-1782.
79. Loaiciga, H. A., et al. (1992), Review of groundwater quality monitoring network design, *Journal of Hydraulic Engineering-Asce*, *118*(1), 11-37.
80. Mackay, D. M., et al. (1986), A natural gradient experiment on solute transport in a sand aquifer .1. Approach and overview of plume movement, *Water Resources Research*, *22*(13), 2017-2029.
81. Martin, R. J. (2001), Comparing and contrasting some environmental and experimental design problems, *Environmetrics*, *12*(3), 273-287.
82. Mckinney, D. C., and D. P. Loucks (1992), Network design for predicting groundwater contamination, *Water Resources Research*, *28*(1), 133-147.
83. Meyer, P. D., et al. (Eds.) (1995), *Design of groundwater monitoring networks for landfills*, Cambridge University Press.
84. Michalak, A. M., and P. K. Kitanidis (2003), A method for enforcing parameter nonnegativity in bayesian inverse problems with an application to contaminant source identification, *Water Resources Research*, *39*(2), -.

85. Michalak, A. M., et al. (2004), A geostatistical approach to surface flux estimation of atmospheric trace gases, *JOURNAL OF GEOPHYSICAL RESEARCH*, *109*.

86. Michalak, A. M., and P. K. Kitanidis (2004), Estimation of historical groundwater contaminant distribution using the adjoint state method applied to geostatistical inverse modeling, *Water Resources Research*, *40*(8).

87. Michalak, A. M., and P. K. Kitanidis (2004), Application of geostatistical inverse modeling to contaminant source identification at dover afb, delaware, *Journal of Hydraulic Research*, *42*, 9-18.

88. Michalak, A. M., and P. K. Kitanidis (2004), A method for the interpolation of nonnegative functions with an application to contaminant load estimation, *Stochastic Environmental Research*, *18*, 1-16.

89. Michalak, A. M., and P. K. Kitanidis (2005), A method for the interpolation of nonnegative functions with an application to contaminant load estimation, *Stochastic Environmental Research and Risk Assessment*, *19*(1), 8-23.

90. Michalak, A. M., and S. Shlomi (2007), A geostatistical data assimilation approach for estimating groundwater plume distributions from multiple monitoring events, in *Data integration in subsurface hydrology*, edited by D. W. Hyndman, et al., AGU, Washington, DC.

91. Michalak, A. M. (2008), A gibbs sampler for inequality-constrained geostatistical interpolation and inverse modeling, *Water Resources Research*.

92. Montasa, H. J., et al. (2000), Heuristic space–time design of monitoring wells for contaminant plume characterization in stochastic flow fields, *Journal of Contaminant Hydrology*, *43*(3-4), 271-301.

93. Muller, W. G., and D. L. Zimmerman (1999), Optimal designs for variogram estimation, *Environmetrics*, *10*(1), 23-37.

94. Müller, W. G. (2001), *Collecting spatial data: Optimum design of experiments for random fields*, Physica-Verlag.

95. Neuman, S. P., and S. Yakowitz (1979), Statistical approach to the inverse problem of aquifer hydrology .1. Theory, *Water Resources Research*, *15*(4), 845-860.

96. Neuman, S. P. (1980), A statistical approach to the inverse problem of aquifer hydrology .3. Improved solution method and added perspective, *Water Resources Research*, *16*(2), 331-346.

97. Neuman, S. P., and P. J. Wierenga (2003), A comprehensive strategy of hydrogeologic modeling and uncertainty analysis for nuclear facilities and sites, Office of Nuclear Regulatory Research, Tucson, AZ.

98. Neupauer, R. M., and J. L. Wilson (2001), Adjoint-derived location and travel time probabilities for a multidimensional groundwater system, *Water Resources Research*, *37*(6), 1657-1668.

99. Nfesc (2000), Guide to optimal groundwater monitoring, edited by Navy, p. 112.

100. Nowak, W., and O. A. Cirpka (2004), A modified levenberg-marquardt algorithm for quasi-linear geostatistical inversing, *Advances in Water Resources*, *27*(7), 737-750.

174

101.   Nowak, W., and O. A. Cirpka (2006), Geostatistical inference of hydraulic conductivity and dispersivities from hydraulic heads and tracer data, *Water Resources Research*, *42*(8).

102.   Nyer, E. K., et al. (1998), Beyond the afcee protocol for natural attenuation, *Ground Water Monitoring and Remediation*, *18*(3), 70-77.

103.   Palmer, C. M. (1996), *Principles of contaminant hydrogeology*, CRC Press.

104.   Porter, D. W., et al. (2000), Data fusion modeling for groundwater systems, *Journal of Contaminant Hydrology*, *42*(2-4), 303-335.

105.   Porter, D. W. (2002), Data fusion modeling for groundwater systems using generalized kalman filtering, Birkhauser Boston, Notre Dame, In, Apr 05.

106.   Rail, C. D. (1989), *Groundwater contamination: Sources, control, and preventive measures*, Technomic Pub. Co.

107.   Reed, P., et al. (2000), Cost-effective long-term groundwater monitoring design using a genetic algorithm and global mass interpolation, *Water Resources Research*, *36*(12), 3731-3741.

108.   Reed, P., et al. (2001), A multiobjective approach to cost effective long-term groundwater monitoring using an elitist nondominated sorted genetic algorithm with historical data, *Journal of Hydroinformatics*, *3*(2), 71-90.

109.   Reed, P. M., et al. (2004), Spatial interpolation methods for nonstationary plume data, *Ground Water*, *42*(2), 190, doi:110.1111/j.1745-6584.2004.tb02667.x.

110.   Ricciardi, K. L., et al. (2007), Efficient groundwater remediation system design subject to uncertainty using robust optimization, *Journal of Water Resources Planning and Management-Asce*, *133*(3), 253-263.

111.   Rivett, M., et al. (2006), Chemicals: Health relevance, transport and attenuation, in *Protecting groundwater for health*, edited by O. Schmoll, et al., p. 678, World Health Organization, London.

112.   Rouhani, S. (1985), Variance reduction analysis, *Water Resources Research*, *21*(6), 837-846.

113.   Runnells, D. D. (1993), Inorganic chemical processes and reactions, in *Regional ground-water quality*, edited by W. M. Alley, John Wiley and Sons.

114.   Russo, D., and W. A. Jury (1988), Effect of the sampling network on estimates of the covariance function of stationary fields, *Soil Sci. Soc. Am. J.*, *52*(5), 1228-1234.

115.   Saito, H., and P. Goovaerts (2001), Accounting for source location and transport direction into geostatistical prediction of contaminants, *Environmental Science & Technology*, *35*(24), 4823, doi 4810.1021/es010580f.

116.   Sakaki, T., et al. (2006), Improved subsurface model calibration using soft data on geologic heterogeneity, paper presented at AGU Fall meeting, San Francisco, CA, December 11-15.

117.   Schumacher, J. G., et al. (2004), Assessment of subsurface chlorinated solvent contamination using tree cores at the front street site and a former dry

cleaning facility at the riverfront superfund site, new haven, missouri, 1999—2003, 35 pp, USGS.

118.   Serre, M. L., et al. (2003), A bme solution of the inverse problem for saturated groundwater flow, *Stochastic Environmental Research and Risk Assessment*, *17*(6), 354-369.

119.   Shlomi, S. (2007), Kriging on a flow net, paper presented at 2007 Engineering Graduate Symposium, University of Michigan, Ann Arbor, MI, Friday, November 2, 2007.

120.   Shlomi, S., and A. M. Michalak (2007), A geostatistical framework for incorporating transport information in estimating the distribution of a groundwater contaminant plume, *Water Resources Research*, *43*, 12.

121.   Shlomi, S., et al. (2007), Evaluation of geostatistical data assimilation methodologies for estimating groundwater plume distributions using 3d tracer tests, paper presented at 37th Mid-Atlantic Industrial & Hazardous Waste Conference, University of Cincinnati, Ohio, March 21-23.

122.   Showalter, P. (1985), Developing objectives for the ground-water quality monitoring network of the salinas river drainage-basin, *Ground Water Monitoring and Remediation*, *5*(2), 37-45.

123.   Skaggs, T. H., and Z. J. Kabala (1994), Recovering the release history of a groundwater contaminant, *Water Resources Research*, *30*(1), 71-79.

124.   Skaggs, T. H., and Z. J. Kabala (1995), Recovering the history of a groundwater contaminant plume - method of quasi-reversibility, *Water Resources Research*, *31*(11), 2669-2673.

125.   Skaggs, T. H., and Z. J. Kabala (1998), Limitations in recovering the history of a groundwater contaminant plume, *Journal of Contaminant Hydrology*, *33*(3-4), 347-359.

126.   Snodgrass, M. F., and P. K. Kitanidis (1997), A geostatistical approach to contaminant source identification, *Water Resources Research*, *33*(4), 537-546.

127.   Sun, N.-Z. (1994), *Mathematical modeling of groundwater pollution*, Springer.

128.   Sun, N. Z. (1996), Identification and reduction of model structure for modeling distributed parameter system, in *Parameter identification and inverse problems in hydrology, geology and ecology*, edited by J. Gottlieb and P. Duchateau, Kluwer, Norwell, MA.

129.   Sun, A. Y., et al. (2006), A constrained robust least squares approach for contaminant release history identification, *Water Resources Research*, *42*(4).

130.   Vyas, V. M., et al. (2004), Geostatistical estimation of horizontal hydraulic conductivity for the kirkwood-cohansey aquifer, *Journal of the American Water Resources Association*, *40*(1), 187-195.

131.   Wackernagel, H. (2003), *Multivariate geostatistics: An introduction with applications*, Third ed., 387 pp., Springer, Berlin Heidelberg New York.

132.   Wagner, B. J., and S. M. Gorelick (1987), Optimal groundwater quality management under parameter uncertainty, *Water Resources Research*, *23*(7), 1162-1174.

133.   Wagner, B. J. (1995), Sampling design methods for groundwater modeling under uncertainty, *Water Resources Research*, *31*(10), 2581-2591.

134.   Warrick, A. W., and D. E. Myers (1987), Optimization of sampling locations for variogram calculations, *Water Resources Research*, *23*(3), 496-500.

135.   Weaver, J. W., et al. (1999), Characteristics of gasoline releases in the water table aquifer of long island, paper presented at Petroleum Hydrocarbons and Organic Chemicals in Ground Water, Houston, Texas, November 17-19.

136.   Woldt, W., and I. Bogardi (1992), Ground-water monitoring network design using multiple criteria decision-making and geostatistics, *Water Resources Bulletin*, *28*(1), 45-62.

137.   Woodbury, A. D., and T. J. Ulrych (1996), Minimum relative entropy inversion: Theory and application to recovering the release history of a groundwater contaminant, *Water Resources Research*, *32*(9), 2671-2681.

138.   Woodbury, A. D., and T. J. Ulrych (1998), Minimum relative entropy and probabilistic inversion in groundwater hydrology, *Stochastic Hydrology and Hydraulics*, *12*(5), 317-358.

139.   Wu, J. F., et al. (2005), Cost-effective sampling network design for contaminant plume monitoring under general hydrogeological conditions, *Journal of Contaminant Hydrology*, *77*(1-2), 41-65.

140.   Zhang, Y. Q., et al. (2005), Least cost design of groundwater quality monitoring networks, *Water Resources Research*, *41*(8), doi:10.1029/2005WR003936.

141.   Zheng, C., and P. P. Wang (1999), Mt3dms, a modular three-dimensional multi-species transport model for simulation of advection, dispersion and chemical reactions of contaminants in groundwater systems; documentation and user's guide, 202 pp, U.S. Army Engineer Research and Development Center, Vicksburg, MS.

142.   Zimmerman, D. L., and K. E. Homer (1991), A network design criterion for estimating selected attributes of the semivariogram, *Environmetrics*, *2*(4), 425-441.

143.   Zimmerman, D. A., et al. (1998), A comparison of seven geostatistically based inverse approaches to estimate transmissivities for modeling advective transport by groundwater flow, *Water Resources Research*, *34*(6), 1373-1413.

144.   Zimmerman, D. L. (2006), Optimal network design for spatial prediction, covariance parameter estimation, and empirical prediction, *Environmetrics*, *17*(6), 635-652.