

1 INTRODUCTION

The *National Research Council* (NRC) has estimated that 300,000 to 400,000 sites in the United States have contaminated groundwater (1997). The estimated cost of remediating these sites ranges from \$480 billion to \$1 trillion, or an average cost of \$8,000 per household in the United States (NRC 1993, 1997). Early legislative efforts leading to the Resource Conservation and Recovery Act (RCRA) and the Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA) focused on restoring these resources to their natural states. Cost and technology limitations have since resulted in a shift in the design paradigm for groundwater remediation from resource recovery to long-term risk management.

Risk-based Corrective Action (RBCA) is a direct result of this shifting emphasis [American Society for Testing and Materials (ASTM) 1995]. RBCA uses a tiered approach to remediation, where the cost and level of remediation efforts are controlled by the human health and ecological risks posed by the contaminated resources. The increasing use of RBCA is expected to result in more contaminants being left in place that will require long-term monitoring [NRC, 1999]. Long-term monitoring (LTM) is particularly important for monitored natural attenuation, in which contaminants are mitigated by the natural processes of dilution, dispersion, and degradation. LTM at many sites can require decades of expensive sampling at tens or even hundreds of existing monitoring wells, resulting in hundreds of thousands or millions of dollars for sampling and data management per year.

Developing efficient and effective LTM sampling plans can be difficult when numerous options exist. In any given monitoring period, the number of possible sampling plans is 2^n , where n is the product of the number of wells and the number of possible constituents to be measured. A site with 10 wells where up to 3 constituents can be measured ($n = 30$) requires the decision

makers to identify a sampling plan among the more than one billion that exist while also trying to balance cost and other performance objectives for each sampling scheme.

1.1 Objectives and Scope

In developing long-term monitoring plans, regulators and stakeholders must negotiate monitoring objectives and decision parameters while accounting for technical, social, and regulatory considerations. The primary objective of this research is to develop a highly adaptable multiobjective LTM design methodology that aids this negotiation process by enabling decision makers to discover, understand, and balance tradeoffs among a variety of performance objectives. Three steps are required in the proposed methodology:

- Step 1: Selection and understanding of performance criteria
- Step 2: Selection of a plume interpolation method
- Step 3: Quantify design tradeoffs using multiobjective genetic algorithms (GAs).

The subsequent chapters of this dissertation address the specific challenges practitioners face when completing steps 1 thru 3 of the LTM design methodology.

1.2 Summary of Research

This thesis is the first monitoring research (see Chapter 2) that successfully addresses the two most challenging problems that face monitoring network designers: (1) selecting monitoring objectives and (2) balancing these objectives. These challenges were addressed using the 3-step problem decomposition discussed above in Section 1.1. Problem decomposition guided the selection and development of tools that were used in a multiobjective optimization framework. The optimization framework serves as an interface between the monitoring system being designed and the human decision process. The optimization framework requires effective plume

interpolation for evaluating LTM designs and the efficient use of multiobjective GAs for quantifying LTM objective tradeoffs.

Chapter 5 of this thesis is the first groundwater plume interpolation research to directly illustrate how preferential sampling and the highly skewed nature of groundwater contamination can combine to severely bias performance rankings of interpolation methods. Quantile kriging was the most robust of 6 groundwater plume interpolation methods, showing the least bias from both preferential sampling and the variability of contaminant data. The findings of Chapter 5 warrant further studies into the applicability of quantile kriging to data sets from other fields ranging from mining to life sciences where nonstationary interpolation also plays a vital role. Chapters 3 and 4 develop the first design methodologies for using evolution-based strategies to efficiently solve a new class of high order multiobjective applications (i.e., applications with more than 2 objectives). These methodologies are then used in Chapter 6 to solve the first application of evolutionary multiobjective optimization algorithms to a real-world problem with 4 objectives. The optimization framework developed in this chapter combines quantile kriging with high order multiobjective optimization to select, understand, and balance LTM performance criteria en route to a final negotiated design.

This thesis demonstrates that combining higher order Pareto optimization with visualization can allow designers in any field to assess the mathematical models used to represent their objectives, discover how their objectives are affecting designs, and negotiate a final design that balances their conflicting design preferences. The methods developed in this thesis are powerful tools for enhancing the design of LTM systems. Sections 1.2.1 thru 1.2.5 summarize each of the individual chapters of this thesis in more detail below.

1.2.1 Chapter 2: Literature Review

Chapter 2 summarizes the extensive previous work in groundwater monitoring network design. Previous studies have primarily focused on two problems: (1) the use of geostatistics to augment or design monitoring networks for site characterization (for a review, see *ASCE Task Committee on Geostatistical Techniques* 1990b) and (2) the use of optimization and numerical simulation for screening monitoring plans for plume detection at landfills and hazardous waste sites (for a review, see *Loaiciga et al.* 1992). Recently, a third problem has emerged that seeks to reduce spatial and temporal redundancies in pre-existing well networks for sites undergoing long term monitoring. The LTM design methodology proposed in this dissertation combines elements of the geostatistical characterization approaches with spatial redundancy analysis to balance sampling costs, uncertainty, the quality of plume maps, and the accuracy of contaminant mass estimates (see Chapter 6 for more details).

1.2.2 Chapter 3: Optimization in Pareto Space

Chapter 3 provides practitioners with a design methodology for the Nondominated Sorted Genetic Algorithm (NSGA). This portion of the research represents an extension of the simple GA design methodology presented by *Reed et al.* (2000b) to computationally intensive, multiobjective water resources applications. The NSGA design methodology is demonstrated using an LTM application, in which the tradeoffs between sampling costs and local concentration estimation errors in an existing groundwater monitoring network were quantified. This chapter shows that with proper design and parameterization, the NSGA is able to accurately quantify 2 dimensional tradeoffs.

1.2.3 Chapter 4: Simplifying Optimization in Pareto Space

Chapter 4 extends the design methodology presented in Chapter 3 to the Nondominated Sorted Genetic Algorithm-II (NSGA-II). NSGA-II is a second generation evolutionary multiobjective (EMO) genetic algorithm that significantly improves upon the original NSGA. NSGA-II improves upon the NSGA [see *Deb et al.* 2000] by (1) invoking a more efficient nondomination sorting algorithm, (2) eliminating the sharing parameter, and (3) adding an implicitly elitist selection method that greatly aids in capturing high order Pareto surfaces. Chapter 4 builds on the NSGA design methodology of Chapter 3 and *Lobo* (2000) to introduce a multi-population approach that automates parameter specification for the NSGA-II and significantly reduces the computational costs associated with solving LTM applications. The methodology successfully solved the same LTM application as was solved in Chapter 3 using 80 percent fewer function evaluations (i.e., sampling design evaluations). The combined efficiency of the NSGA-II and design methodology presented in this chapter allows for more challenging higher order Pareto optimization problems (i.e., problems with more than 2 objectives) to be solved [see Chapter 6].

1.2.4 Chapter 5: Spatial Interpolation Methods for Plume Data

Plume interpolation consists of estimating contaminant concentrations at unsampled locations using the available contaminant data surrounding those locations. The goal of groundwater plume interpolation is to maximize the accuracy in estimating the spatial distribution of the contaminant plume given the data limitations associated with sparse monitoring networks with irregular geometries. Beyond data limitations, contaminant plume interpolation is a difficult task because contaminant concentration fields are highly heterogeneous, anisotropic, and nonstationary phenomena. This chapter provides a

comprehensive performance analysis of 6 interpolation methods for scatter-point concentration data, ranging in complexity from intrinsic kriging based on intrinsic random function theory to a traditional implementation of inverse-distance weighting. High resolution simulation data of perchloroethylene (PCE) contamination in a highly heterogeneous alluvial aquifer were used to generate 3 test cases, which show how each interpolation method performs as a function of the amount of available sample data. Overall, the variability of PCE samples and preferential sampling in the source area controlled how well each of the interpolation schemes performed. Quantile kriging was the most robust of the interpolation methods, showing the least bias from both of these factors. Additionally, the method's non-parametric uncertainty estimates successfully predicted zones of high estimation error for each test case. This chapter provides guidance to practitioners balancing opposing theoretical perspectives, ease-of-implementation, and effectiveness when choosing a plume interpolation method.

1.2.5 Chapter 6: Balancing Performance Criteria

This chapter integrates the tools developed in the previous two chapters into a multiobjective optimization framework that can serve as an interface between the physical system being designed and the human decision process. This chapter demonstrates the use of high order Pareto optimization (i.e., optimizing a system for more than 2 objectives) in a highly adaptable optimization methodology. The methodology is implemented on an LTM application that combines quantile kriging [see Chapter 5] and the NSGA-II [see Chapter 4] to successfully balance four objectives: (1) minimizing sampling costs, (2) maximizing the quality of interpolated plume maps, (3) maximizing the relative accuracy of contaminant mass estimates, and (4) minimizing local estimation uncertainty. Optimizing the LTM application with respect to these objectives reduced the decision space of the problem from a total of 500 million designs to

the set of 1156 designs identified on the Pareto surface. Visualization of a total of 8 designs aided in understanding and balancing the application's objectives en route to a single compromise solution. This study shows that high order Pareto optimization holds significant potential as a tool that can be used in the balanced design of water resources systems.

1	INTRODUCTION.....	1
1.1	Objectives and Scope	2
1.2	Summary of Research.....	2
<i>1.2.1</i>	<i>Chapter 2: Literature Review</i>	<i>4</i>
<i>1.2.2</i>	<i>Chapter 3: Optimization in Pareto Space.....</i>	<i>4</i>
<i>1.2.3</i>	<i>Chapter 4: Simplifying Optimization in Pareto Space</i>	<i>5</i>
<i>1.2.4</i>	<i>Chapter 5: Spatial Interpolation Methods for Plume Data.....</i>	<i>5</i>
<i>1.2.5</i>	<i>Chapter 6: Balancing Performance Criteria</i>	<i>6</i>