Designing A New Elitist Nondominated Sorted Genetic Algorithm For A Multiobjective Long Term Groundwater Monitoring Application

Patrick M. Reed
4129 NCEL, MC. 250
205 N. Mathews Ave.
Urbana, IL. 61801
preed@uiuc.edu
217-333-6979

Barbara S. Minsker
3230d NCEL, MC. 250
205 N. Mathews Ave.
Urbana, IL. 61801
minsker@uiuc.edu
217-333-9017

David E. Goldberg
117 Transportation, MC. 238
104 S. Mathews
Urbana, IL. 61801
deg@uiuc.edu
217-333-0897

Abstract

Although usage of genetic algorithms (GAs) has become widespread, the theoretical work from the genetic and evolutionary computation (GEC) field has been largely ignored by practitioners in real-world applications. This paper provides an overview of a three-step method for utilizing GEC theory to ensure robust search and avoid the common pitfalls in GA applications. Additionally, this study presents a niching-based elitist enhancement of the Non-dominated Sorted Genetic Algorithm (NSGA) and tests its performance in identifying the Pareto frontier for a groundwater monitoring application. The Elitist NSGA nearly replicated the true front, finding representative solutions along the entire trade off between cost and estimation error.

1 INTRODUCTION

The major goal of this study is to present a new niching-based elitist enhancement of the NSGA. Additionally, an extension of the genetic algorithm design methodology presented by Reed et al. (2000b) is demonstrated on a multiobjective long-term groundwater monitoring application.

1.1 PREVIOUS WORK

Evolution-based multiobjective optimization (EMO) methods have garnered increased attention since the seminal work presented in Schaffer (1984) (for reviews see Fonseca & Fleming 1995, Coello 1999, Van Veldhuizen 1999). Cieniawski (1993) is one of the earliest studies in the water resources field to utilize EMO methods. The study is an empirical comparison of the performance of VEGA relative to niching-based techniques from Goldberg & Richardson (1987) for identifying a monitoring network to detect potential contaminant leaks from a hazardous waste landfill. Cieniawski (1993) and Cieniawski *et al*. (1995) clearly espouse the efficiency of EMO methods in quantifying tradeoffs between maximizing a groundwater-monitoring network's reliability in detecting contaminants and minimizing the costs associated with remediating the contaminated aquifer at the time of first detection. A common criticism of EMO methods is that the methods often fail to find Pareto optimal solutions along the full extent of the Pareto frontier (Coello 1999, Van Veldhuizen 1999, Zitzler *et al*. 2000). The elitist enhancement developed in this work is intended to address this issue.

This study focuses on the Nondominated Sorted Genetic Algorithm (NSGA) because Zitzler *et al*. (2000) showed that the NSGA performed as well or better than a representative sampling of EMO methods on a suite of test problems with properties similar to our application.

1.2 APPLICATION

The monitoring application uses data drawn from a 38 million-node flow-and-transport simulation performed by Maxwell et al. (2000). The simulation provided realistic historical data for the migration of a plume of perchloroethylene (PCE) in groundwater. PCE is a commonly used industrial solvent that can potentially cause cancer in exposed individuals. Data were provided for a total of 50 hypothetical sampling locations within the 20-well multi-level monitoring network. The data represent a snapshot in time, 10 years after an underground storage tank has continuously released contamination into the site's groundwater. The site is assumed to be undergoing long-term monitoring, in which groundwater samples are used to assess the effectiveness of clean up efforts in reducing the amount of PCE in the subsurface.

During this long-term monitoring phase of a remediation, sampling and laboratory analysis can be a controlling factor in the costs of remediating a site. Quarterly sampling of the entire network described above has a potential cost of over $70,000 annually for PCE testing alone, which could translate into millions of dollars if the

site had a typical life span of 20 to 30 years. The significance of these costs has motivated the development of several approaches for reducing the fiscal burden posed by long term monitoring by identifying redundant wells in groundwater monitoring networks that can be omitted from future sampling periods (Cameron & Hunter 2000, Aziz *et al*. 2000, Reed *et al*. 2000a, Rizzo & Dougherty 2000). These methods define sampling points to be redundant when they minimally affect interpolated concentration estimates (called plume estimates). They employ a variety of single objective optimization techniques ranging from a simple genetic algorithm to trial-and-error heuristics. The objective of these methods is to minimize sampling costs while incorporating performance objectives associated with plume estimates as constraints. The management model presented in this work builds on these previous methods by introducing a sampling design methodology that explicitly and efficiently identifies the tradeoffs encountered when reducing monitoring costs.

2 MANAGEMENT MODEL

To identify which wells are redundant, this study employs a multiobjective approach with the intention of attaining the best-interpolated picture of the PCE plume for the least cost. Equation (1) gives the multiobjective problem formulation for quantifying the tradeoff between sampling costs and maintenance of a high quality interpolated picture of the plume.

$$Minimize\ F(\bar{x}_k) = \left[f_1(\bar{x}_k),\ f_2(\bar{x}_k)\right], \quad \forall k \in \Omega$$

$$f_1(\bar{x}_k) = \sum_{i=1}^{nwell} C_S(i) x_{k\,i}$$

$$f_2(\bar{x}_k) = \sum_{j=1}^{n_{est}} \left(c^*_{all}(\bar{u}_j) - c^k_{est}(\bar{u}_j)\right)^2 \qquad (1)$$

$F(\bar{x}_k)$ is a vector valued fitness function whose components $[f_1(\bar{x}_k),\ f_2(\bar{x}_k)]$ represent the cost and squared relative estimation error (SREE), respectively, for the k th monitoring scheme $\bar{x}_k$ taken from the collection of all possible sampling designs Ω. Equation (2) defines the binary decision variables representing the k th monitoring scheme.

$$x_{k\,i} = \begin{cases} 1, & if\ the\ i^{th}\ well\ is\ sampled \\ 0, & otherwise \end{cases}, \quad \forall k, i \qquad (2)$$

If the ith well is sampled it is assumed that all available locations along the vertical axis of that well will be sampled at a cost of $C_S(i)$. $C_S(i)$ ranged from \$365 to \$1095 for 1 to 3 samples analyzed for PCE solely (Rast 1997). Sampling all available levels within each well reduces the size of Ω from 2^{50} to 2^{20} where 50 and 20 represent the total number of sampling locations and monitoring wells (*nwell*), respectively. Reducing the size of Ω enabled the entire decision space of this application to be enumerated. Enumeration was employed to identify the true Pareto frontier so that the performance of the NSGA could be rigorously tested.

The SREE provides a measure of how the interpolated picture of the plume using data only from wells included in the k th sampling plan compares to the result attained using data from all available sampling locations. The measure is computed by summing the squared deviations between the PCE concentration estimates using data from all available sampling locations, $c^*_{all}(\bar{u}_j)$, and the estimates based on the k th sampling plan $c^k_{est}(\bar{u}_j)$ at each location $\bar{u}_j$ in the interpolation domain. Each $\bar{u}_j$ specifies the coordinates for the jth grid point in the interpolation domain. The interpolation domain consisted of a total of 3300 grid points (n_{est} in equation (1)). The PCE estimates used in the calculation of the SREE for each of the sampling designs were attained using a nonlinear spatial interpolation fitness function.

3 EFFICIENT DESIGN FOR SEARCH & OPTIMIZATION

The NSGA uses nondomination ranking and niching to evolve the Pareto set (for details see Srinivas & Deb 1995). One of the difficulties in applying EMO methods is identifying parameter settings that ensure comprehensive navigation of the decision space and adequate coverage of the Pareto frontier (Van Veldhuizen & Lamont 2000, Cieniawski 1993). Most practitioners use trial-and-error runs to identify the best parameter settings, but this approach is quite time consuming, particularly for applications with computationally intensive fitness functions.

Reed et al. (2000b) present a 3-step methodology for the design of simple genetic algorithms that accounts for population sizing, selection pressure, and the influence of crossover and mutation on real-world computationally intensive applications. The methodology assumes that computationally intensive fitness functions for real-world applications preclude identifying parameter settings for a distribution of initial random number seeds and instead focuses on finding optimal parameter settings for a single random number seed. The following sections summarize an extension of their methodology to the NSGA.

3.1 STEP 1: PRELIMINARY PROBLEM ANALYSIS

The initial step in the methodology consists of preliminary problem analysis to determine a range of potential population sizes and the computational

complexity associated with solving the application as it is currently formulated. In this study, relationships from Mahfoud (1995) were used to attain six population size estimates ranging from 370 to 870. This population size range was used in combination with convergence rate relationships (Thierens *et al*. 1998, Thierens & Goldberg 1994) for stochastic remainder selection as recommended by Srinivas & Deb (1995) to attain estimates of the total number of function evaluations required for this application. Using these relationships, the total required run length was estimated to be approximately 40 generations, yielding a range of potential run times between 10 and 25 minutes. The next step in the method uses relationships from literature and trial runs to set the input parameters for the NSGA.

3.2 STEP 2: PARAMETER SELECTION

Step 1 provided a range of population sizes and an estimate of run length. The remaining input parameters for the NSGA must be specified in the second step. The population sizing relationships presented by Mahfoud (1995) directly account for the potential disruptive effects of crossover and require the specification of the probability of crossover P_c. Comprehensive reviews of the EMO literature (Fonseca & Fleming 1995, Coello 1999, Van Veldhuizen 1999) showed that a majority of applications specify P_c to fall witin the range [0.6, 0.9]. For this application, the lower bound of this range was used to reduce population size estimates. The population sizing relationships assumed that mutation is minimally disruptive. This assumption was enforced by setting the probability of mutation P_m equal to the inverse of the population size, which is the relationship recommended by DeJong (1975).

The NSGA requires additional specification of the parameters controlling fitness sharing in phenotypic space. The relationships presented by Deb & Goldberg (1989) were used to size the niche radius σ_{share}. For this application σ_{share} was set equal to 1.9.

3.3 STEP 3: ELITISM & DRIFT ANALYSIS

Step 3 introduces an elitist enhancement to the NSGA and uses limited trial runs to specify an optimal population size. Zitzler et al. (2000) showed that elitism and population sizing are the primary factors controlling the performance of the NSGA. The importance of elitism and population sizing in the performance of the NSGA relates directly to the selection pressure that nondominated individuals experience as the algorithm evolves the Pareto optimal set. Genetic drift (non-optimal convergence due to crossover and mutation solely) is prevented in this methodology by introducing a niching-based elitist enhancement to the NSGA and using trial runs to identify sufficiently large population sizes.

3.3.1 Seeking the king of the niche

Elitist operators provide a means of ensuring that the best individuals are identified and allowed to pass their traits to latter generations. Unlike elitist sGA applications, EMO methods cannot simply pass a single individual with the current best fitness function value into the next generation. Multiobjective optimization requires that some fraction of the solutions along the current nondominated front be passed on to the next generation. A variety of elitist strategies have been used previously, usually consisting of maintenance of a population of nondominated solutions outside of the normal operators of the given EMO method being employed (for more details see Ishibuchi & Murata 1996, Bäck 1996, Parks & Miller 1998, Zitzler & Thiele 1999). Zitzler et al. (2000) state the primary question practitioners must answer when using elitist strategies as: "When and how are which members of the elite set re-inserted into the population?"

The elitist strategy employed in this study was designed to use previously derived niching relationships to answer this question while maximizing the performance of the NSGA. In an elitist sGA, the best member in the population at generation t, if not present in the new population resulting from selection, crossover, and mutation at generation (t+1), randomly replaces one member of the population. For the NSGA, sharing provides niches that represent stable subpopulations that search for nondominated solutions in subspaces of Ω. Conceptually, the elitist strategy proposed in this study is very similar to the sGA, in that the current best individual in a given niche at generation t, if not present in generation (t+1), is inserted into that subpopulation, ensuring that its traits are available for subsequent search for the Pareto front.

This strategy was implemented by defining σ_{elite} or the elite radius, which is a parameter that allows the user to easily manipulate the amount of elitism. The elite radius defines the distance beyond which members of the current nondominated set are considered independent from one another. Only independent members of the nondominated set are considered for insertion in the next generation. For this application, $\sigma_{elite} \approx \sigma_{share}$ which means that only one representative of each niche in the current nondominated set is considered for elitist reproduction into the next generation. The elitist solutions were selected in the four steps shown below from the nondominated set (or first front) at each generation t.

Step 1: Randomly select an objective f_b for b equal 1 to n_{obj} (the number of objectives)

Step 2: Flip a coin to determine whether to start with either the member in the current nondominated set with

the maximum value of f_b or the member with the minimum value.

Step 3: Identify the next point in the nondominated set that satisfies the following conditions:

(1) Is a distance greater than s_{elite} from the current solution

(2) Is the closest member of the nondominated set to the current position

If none exist, then elitist reproduction is ceased or not performed at all.

Step 4: Repeat Step 3 until elitist reproduction is ceased.

This approach identifies a niched elitist set by systematically stepping through the current nondominated front from one end to the other. After the elitist set of solutions is selected using the above steps, those members who are not represented in generation (t+1) randomly replace individuals. Setting the elite radius equal to the niche radius worked well for this application, but the elite radius parameter allows the practitioner to directly manipulate the elitist selection pressure for other applications if this rule-of-thumb does not work as effectively.

3.3.2 Trial runs to determine population size

The final selection of a population size was completed by performing trial runs for each of the 6 estimates attained in the preliminary problem analysis of step 1. For this study, a total of 12 trial runs were completed to allow the selection of an optimal population size for both the NSGA and the Elitist NSGA. Results from these trial runs identified population sizes equal to 830 and 760 for the NSGA and the Elitist NSGA, respectively. The performance of both forms of the algorithm on the monitoring application are presented and discussed in section 4.

3.3.3 Relative scoring metric (RSM)

Equation (3) defines the deviation between the ηth member of the enumerated Pareto optimal set and the ϕth member of the current nondominated set in generation t to be equal to the absolute difference of their SREE values ($SREE_{true}(\bar{x}_h) - SREE(f(t))$ shown below) if the designs have the same cost. If a cost level present in the Pareto optimal set is not represented in the nondominated set at generation t then equation (3) assumes a maximum deviation of one. The RSM was used to monitor the performance of the NSGA in order to evaluate the effectiveness of the guidelines presented in this section for a realistic application.

$$Deviation(h) = \begin{cases} \left| SREE_{true}(\bar{x}_h) - SREE(f(t)) \right| / SREE_{\max} \\ 1, \qquad otherwise \end{cases}$$

$$RSM = 1 - \left(\sum_{h=1}^{36} Deviation(h) \Big/ 36 \right) \qquad (3)$$

4 RESULTS & DISCUSSION

Figures 1 and 2 show the performance of both the NSGA and the Elitist NSGA. The open circles show the offline performance of each form of the algorithm relative to the true Pareto front attained from enumeration of the problem's decision space. Recall that the population sizes used in these runs are 830 and 760 for the NSGA and the Elitist NSGA, respectively. Also, the offline results shown below were attained after 40 generations, which represents the recommended run time found in step 1.

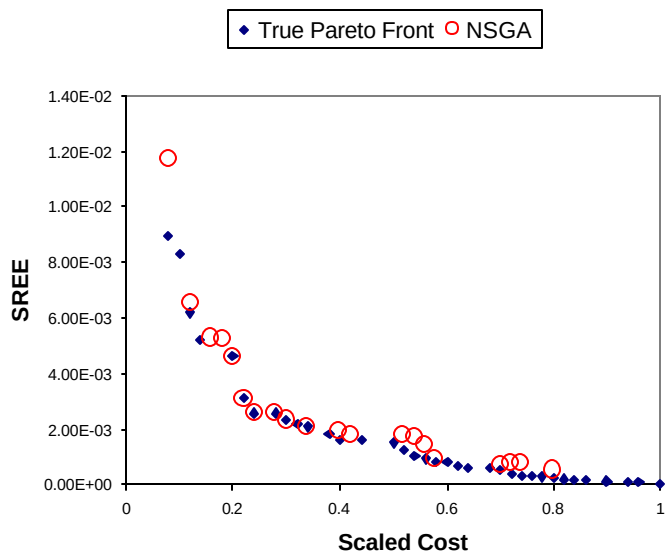

Figure 1: NSGA results without elitism.

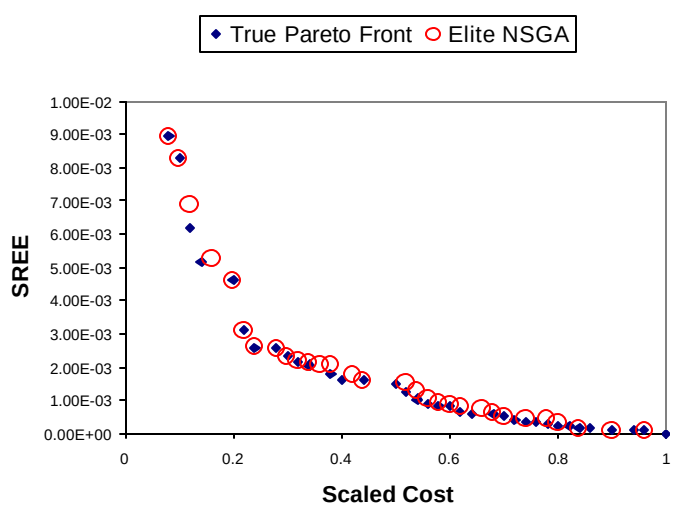

Figure 2: NSGA results with elitism.

Figures 1 and 2 show that the niching-based elitist strategy presented in section 3 significantly improved the NSGA's performance. The Elitist NSGA solution very closely follows the true front, finding representative solutions along the entire trade off between cost and estimation error. Figure 1 shows significant gaps in the extremes of both objectives, which reflect the loss of niches due to deficient selective pressures. The niching-based elitist strategy effectively rescaled the system, increasing selection pressure along the entire extent of the Pareto front and reducing the loss of these niches.

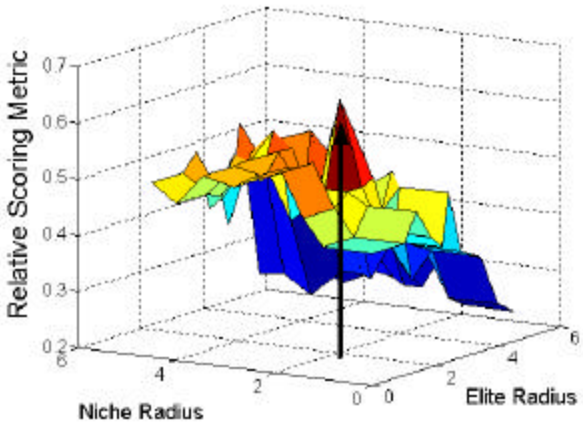

Figure 3: NSGA performance as a function of niching and elitism.

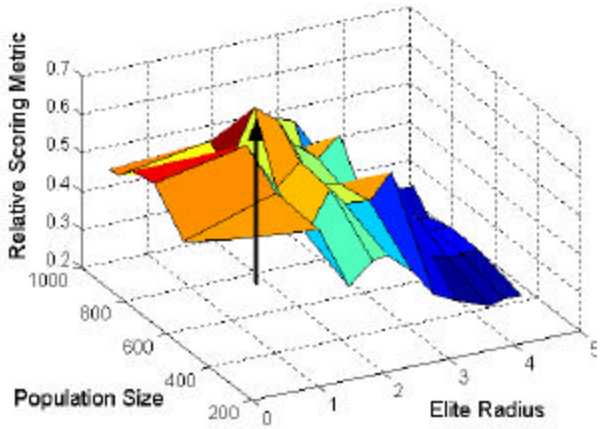

Figure 4: NSGA performance as a function of population size and elitism.

Figures 3 and 4 show the NSGA's performance as a function of elitism, niching, and population sizing. The figures were attained from over 150 trial runs performed for the full range of possible values that the niche radius and elite radius can be assigned as well as the 6 population sizes identified in step 1 of the GA design methodology. Figure 3 shows the NSGA's performance for cases ranging from when the whole population is considered a single species (niche radius = 5) to when each individual is the only member of its niche (niche radius = 0). Additionally, the plot shows the range of performance that results from the absence of elitist reproduction (elite radius = 5) to the case when the entire nondominated front is placed in the elitist set (elite radius = 0). The arrow in the figure highlights a peak in performance that occurs when the niche radius and elite radius are both set equal to 1.9 using the relationships from Deb & Goldberg (1989). Figures 3 and 4 confirm the findings of Zitzler et al. (2000), which showed that population sizing and elitism are the primary factors controlling the performance of EMO methods. Figure 4 shows shows that elitism stabilizes the NSGA's performance for the full range of population sizes. Moreover, the plot confirms that the 3-step GA design method used in this work was able to attain peak or near peak performance in a minimum number of trial runs. The black arrow designates a peak in performance attained when the population size is set equal to 760 (as recommended in section 3.3.2 above) and the elite radius is set equal to 1.9 (as recommended in section 3.3.1 above).

5 CONCLUSIONS

The niching-based elitist enhancement of the NSGA demonstrated in this study captured the true trade off between cost and squared relative estimation error (SREE) for a long-term groundwater monitoring application. Additionally, this paper addresses a primary difficulty and source of criticism for using multiobjective evolutionary methods in real-world engineering applications by extending the design methodology presented by Reed et al. (2000) to the NSGA. The extended methodology enables practitioners to minimize the number of runs required to identify effective parameter settings for the NSGA.

6 ACKNOWLEDGEMENTS

This work was supported by the United States Environmental Protection Agency's Science to Achieve Results (STAR) graduate fellowship program. Dr. Albert Valocchi and Dr. Timothy Ellsworth are gratefully acknowledged for their guidance in the selection and implementation of the nonlinear interpolation method utilized in this study. Additionally, the authors are grateful to Dr. Reed Maxwell for providing the simulation data used in the test case.

7 REFERENCES

Aziz, J. J., Newell, C. J., Rifai, H. S., Ling, M., & Gonzales, J. R. 2000 *Monitoring and Remediation Optimization System (MAROS): Software User's Guide*, Version 1, United States

Air Force Center for Environmental Excellence, Brooks AFB: San Antonio, TX.

Bäck, T. 1996 *Evolutionary Algorithms in Theory and Practice*, Oxford University Press: New York, NY.

Barry, D. A. & Sposito, G. 1990 Three-dimensional statistical moment analysis of the Stanford/Waterloo Borden tracer data. *Water Resources Research*, 26(8), 1735-1747.

Cameron, K. & Hunter, P. 2000 Optimization of LTM Networks: Statistical Approaches to Spatial and Temporal Redundancy. *Proceedings from the American Institute of Chemical Engineers, 2000 Spring National Meeting, Remedial Process Optimization Topical Conference*, Atlanta, GA.

Cieniawski, S. E. 1993 *An Investigation of the Ability of Genetic Algorithms to Generate the Tradeoff Curve of a Multi-Objective Groundwater Monitoring Problem*, Masters Thesis , University of Illinois: Urbana, IL.

Cieniawski, S. E., Eheart, J. W., & Ranjithan, S. 1995 Using genetic algorithms to solve a multiobjective groundwater monitoring problem. *Water Resources Research*, 31(2), 399-409.

Coello, C. A. C. 1996 *An Empirical Study of Evolutionary Techniques for Multiobjective Optimization in Engineering Design*, Doctoral dissertation, Tulane University: New Orleans, LA.

Coello, C. A. C. 1999 A comprehensive survey of evolutionary-based multiobjective optimization techniques. *Knowledge and Information Systems*, 1(3), 269-308.

Deb, K. & Goldberg, D. E. 1989 An investigation of niche and species formation in genetic function optimization. In J. D. Schaffer, ed. *Proceedings of the Third International Conference on Genetic Algorithms*, pp. 42-50, Morgan Kaufmann: San Mateo, CA.

DeJong, K. A. 1975 *An analysis of the behavior of a class of genetic adaptive systems*, Doctoral dissertation, University of Michigan: Ann Arbor, MI.

Fonseca, C. M. and Fleming, P. J. 1995 An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary Computation*, 3(1), 1-16.

Goldberg, D. E. & Richardson, J. 1987 Genetic algorithms with sharing for multimodal function optimization. In J. J. Grefenstette, ed. *Genetic Algorithms and Their Applications: Proceedings of the Second International Conference on Genetic Algorithms*, pp. 41-49, Morgan Kaufmann: San Mateo, CA.

Goldberg, D. E. 1989 *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley: New York, NY.

Goldberg, D.E., Deb, K., & Korb, B. 1989 Messy genetic algorithms: Motivation, analysis, and first results. *Complex Systems*, 3(4), 493-530.

Goldberg, D. E., Deb, K., & Horn, J. 1992 Massive multimodality, deception, and genetic algorithms. In R. Männer & B. Manderick, eds. *Parallel Problem Solving From Nature*, 2, Amsterdam: North-Holland.

Holland, J. H. 1975 *Adaptation in Natural and Artificial Systems*, University of Michigan: Ann Arbor, MI.

Ishibuchi, H. & Murata, T. 1996 Multi-objective genetic local search algorithm. In *Proceedings of 1996 IEEE International Conference on Evolutionary Computation*, pp. 119-124, IEEE: Piscataway, NJ.

Mahfoud, S. 1995 *Niching Methods for Genetic Algorithms*, Doctoral dissertation, University of Illinois: Urbana, IL.

Maxwell, Reed M., Carle, Steven F., and Tompson, F. B. 2000 *Contamination, Risk, and Heterogeneity: On the Effectiveness of Aquifer Remediation*, Lawrence Livermore National Laboratory Report, UCRL-JC-139664, Livermore, CA.

Moré, J. J. 1977 The Levenberg-Marquardt algorithm: Implementation and theory. In G. A. Watson, ed. *Numerical Analysis*, pp. 105-116, Springer-Verlag: New York, 1977.

Parks, G. T. & Miller, I. 1998 Selective breeding in a multiobjective genetic algorithm. In A. E. Eiben, T. Bäck, M. Schoenauer, and H. Schwefel, eds. *Fifth International Conference on Parallel Problem Solving form Nature*, pp. 250-259, Springer: Berlin, Germany.

Rast, R. ed. 1997 *Environmental Restoration: Unit Cost Book*, R. S. Means: Kingston MA.

Reed, P., Minsker B., & Valocchi, A. J. 2000a Cost effective long-term groundwater monitoring design using a genetic algorithm and global mass interpolation. *Water Resources Research*, 36(12), 3731-3741.

Reed, P., Minsker B., & Goldberg, D. E. 2000b Designing a competent simple genetic algorithm for search and optimization. *Water Resources Research*, 36(12), 3757-3761.

Rizzo, D. M., Dougherty, D. E., and Yu M. 2000 An adaptive long-term monitoring and operations system (aLTMOs) for optimization in environmental management, In *ASCE 2000 Joint Conference on Water Resources Engineering and Water Resources Planning and Management*, Minneapolis, MN.

Schaffer, J. D. 1984 *Some experiments in machine learning using vector evaluated genetic algorithms*, Doctoral dissertation, Vanderbilt University: Nashville, TN.

Srinivas, N & Deb, K. 1995 Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary Computation*, 2(3), 221-248.

Thierens, D. & Goldberg, D. E. 1993 Mixing in genetic algorithms. In S. Forrest, ed. *Proceedings of the Fifth International Conference on Genetic Algorithms*, pp. 38-45, Morgan Kaufmann Publishers: San Mateo, CA.

Thierens, D. & Goldberg, D. E. 1994 Convergence Models of Genetic Algorithm Selection Schemes, In Y. Davidor, Hans-Paul Schwefel, and Reinhard Manner, eds. *Proceedings of the Third Conference on Parallel Problem Solving from Nature*, pp. 119-129, Springer-Verlag: New York, NY.

Thierens, D., Goldberg, D. E., & Pereira, A. G. 1998 Domino Convergence, Drift, and the Temporal-Salience Structure of Problems, In *The 1998 IEEE International Conference on Evolutionary Computation Proceedings*, pp. 535-540, IEEE Press: New York, NY.

Van Veldhuizen, D. A. 1999 *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*, Doctoral dissertation, AFIT/DS/ENG/99-01, Air Force Institute of Technology: Wright-Patterson AFB, Ohio.

Van Veldhuizen, D. A. & Lamont, G. B. 2000 Multiobjective evolutionary algorithms: analyzing the state-of-the-art. *Evolutionary Computation*, 8(2), 125-147.

Zitzler, E. & Thiele, L. 1999 Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. *IEEE Transactions on Evolutionary Computation*, 3(4), 257-271.

Zitzler, E., Deb, K., & Thiele, L. 2000 Comparison of multiobjective evolutionary algorithms: empirical results. *Evolutionary Computation*, 8(2), 173-195.