Using Interactive Archives in Evolutionary Multiobjective Optimization: Case Studies for Long-Term Groundwater Monitoring Design

P.M.Reed and V.K. Devireddy

Department of Civil and Environmental Engineering, The Pennsylvania State University,

212 Sackett Building, University Park, PA 16802-1408,

Abstract: Monitoring design is a problem of paramount importance to the environmental engineering field because environmental observation data provide the sole means of assessing if engineered systems are successfully protecting human and ecologic health. The monitoring design problem is extremely challenging because it requires environmental engineers to capture an impacted system's governing processes, elucidate human and ecologic risks, limit monitoring costs, and satisfy the interests of multiple stakeholders (e.g., site owners, regulators, and public advocates). Evolutionary multiobjective optimization (EMO) has tremendous potential to help resolve these issues by providing environmental stakeholders with a direct understanding of their monitoring tradeoffs. This paper demonstrates how **ε**-dominance archiving and automatic parameterization techniques for the NSGA-II can be used to significantly improve the algorithm's ease-of-use and efficiency for computational intensive applications. Results are presented for a 4-objective groundwater monitoring problem in which the archiving and parameterization techniques for the NSGA-II combined to reduce computational demands by more than 90-percent relative to prior published results. The methods of this paper can be easily generalized to other multiobjective applications to minimize computational times as well as trail-and-error parameter analysis.

Keywords: Groundwater; Monitoring; Design; Water resources management; Multiobjective optimization; Genetic algorithms

1. INTRODUCTION

This paper demonstrates the use of evolutionary multiobjective optimization (EMO) to design groundwater monitoring networks for conflicting objectives. Long-term groundwater monitoring (LTM) can be defined as the sampling of groundwater quality over long time-scales to provide "sufficient and appropriate information" to assess if current mitigation or contaminant control measures are performing adequately to be protective of human and ecological health Task Committee on Long-Term Groundwater Monitoring Design [2003]. The LTM problem is ideal for demonstrating how EMO can aid environmental engineers because of the tremendous expense and complexity of characterizing groundwater contamination sites over long time periods.

The 4-objective monitoring design problem presented in this paper is solved using a modified version of the Nondominated Sorted Genetic Algorithm-II (NSGA-II) Deb et al. [2002], which will be termed the **ε**-dominance NSGA-II in this paper using the abbreviated notation, **ε**-NSGA-II. The **ε**-NSGA-II demonstrates how **ε**-dominance archiving (Laumanns et al. [2002], Deb et al. [2003]) can be combined with a parameterization strategy for the NSGA-II Reed et al. [2003] to accomplish the following goals: (1) ensure the algorithm will maintain diverse solutions, (2) eliminate the need for trial-and-error analysis for parameter settings (*i.e.*, population size, crossover and mutation probabilities), and (3) allow users to *sufficiently* capture tradeoffs using a minimum number of design evaluations. A sufficiently quantified trade-off can be defined as a subset of Pareto optimal solutions that provide an adequate representation of the Pareto frontier that can be used to inform decision making.

In this paper, section 2 overviews prior **ε**-dominance archiving and parameterization studies used in the development of the **ε**-NSGA-II. Section 3 discusses the 4-objective groundwater monitoring test case used to demonstrate the **ε**-NSGA-II. Sections 4 and 5 provide a more detailed description of the **ε**-NSGA-II and its performance for the groundwater monitoring test case, respectively.

2. PRIOR WORK

The **ε**-NSGA-II combines the external archiving techniques recommended by Laumanns et al. [2002] and Deb et al. [2003] with automatic parameterization techniques (Reed et al. [2003]) developed to eliminate trial-and-error parameter analysis for the NSGA-II. A primary drawback of using EMO methods for environmental applications lies in the large costs associated with assessing performance (*i.e.*, algorithmic reliability and solution quality). The common practice of assessing performance for a distribution of random seeds employed in the EMO literature is often prohibitively expensive in terms of computational costs and in terms of the time that must be invested by users. The goal of the automated parameterization approaches developed by Reed et al. [2003] is to eliminate the need to assess algorithmic performance for a distribution of initial random number seeds and instead focus on the NSGA-II's reliability and efficiency for a single random seed. Reliability is addressed in the approach by adaptively increasing the size of the population. The method uses multiple runs in which the nondominated solutions are accumulated from the results of the successively doubled population sizes. The runs (and successive doubling of population sizes) continue until either the user-defined maximum run-time is reached or sufficient solution accuracy has been attained.

3. MONITORING TEST CASE PROBLEM

3.1 Test Case Data

The test case developed for this study uses data drawn from a 50 million-node flow-and-transport simulation performed by Maxwell et al. [2000]. The simulation provided realistic historical data for the migration of a hypothetical perchloroethylene (PCE) plume in a highly heterogeneous alluvial aquifer. The hydrogeology of the test case is based on an actual site located at the Lawrence Livermore National Laboratory (LLNL) in Livermore, California. Data were provided for a total of 58 hypothetical sampling locations within a 29-well multi-level monitoring network. If the i^{th} monitoring well was selected for sampling then PCE is sampled at all the possible sampling locations along its vertical axis.

3.2 Problem Formulation

The 4-objective monitoring test case used in this paper combines both the spatial redundancy and geostatistical approaches to monitoring design. The **ε**-NSGA-II and quantile kriging are combined to quantify the tradeoffs among the following four performance criteria: (1) cost, (2) squared relative estimation error (SREE), (3) the relative global mass error (MAE), and (4) local uncertainty as measured by kriging estimation variances. Cost is a linear function of the number of PCE samples that are used in a given monitoring design. SREE measures how the interpolated picture of the plume using data only from wells included in the κ^{th} sampling plan compares to the result attained using data from all available sampling locations. Likewise, the global mass objective error in the total mass of PCE in the subsurface. Lastly, local uncertainty is estimated using the sum of the estimation standard deviations (i.e., the square root of estimation variances) from kriging (for more details see Reed and Minsker [2004]).

3.3 Plume Interpolation using Quantile Kriging

Quantile kriging was selected for plume interpolation in this study based on the findings of Reed et al. [2004], who present 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. Quantile kriging was shown to be the most robust and least biased of the interpolation methods they studied. Additionally, the method's non-parametric uncertainty estimates successfully predicted zones of high estimation error for each test case.

4. OVERVIEW OF THE ε-NSGA-II APPROACH

The proposed algorithm consists of three steps. The first step utilizes the NSGA-II with a starting population of 5 individuals to initiate EMO search. The initial population size is set arbitrarily small (i.e., 5 in this paper) to ensure the algorithm's initial search is done using a minimum number of function evaluations. Subsequent increases in the population size adjust the population size commensurate with problem difficulty. In the second step, the **ε**-NSGA-II uses a fixed sized archive to store the nondominated solutions generated in every generation of the NSGA-II runs. The archive is updated using the concept of **ε**-dominance, which has the benefit of ensuring that the archive maintains a diverse set of solutions. **ε**-dominance requires the user to define the precision with which they want to evaluate each objective by specifying an appropriate **ε** value for each objective.

The third step checks if the user-specified performance and termination criteria are satisfied and the Pareto optimal set has been sufficiently quantified. If the criteria are not satisfied, the population size is doubled and the search is continued. When increasing the population, the initial population of the new run has solutions injected from the archive at the end of the previous run. The algorithm terminates if either a maximum user time is reached or if doubling the population size fails to significantly increase the number of nondominated solutions found across two runs. The following sections discuss the **ε**-NSGA-II in greater detail.

4.1 Searching with the NSGA-II

Development of the **ε**-NSGA-II was motivated by the authors' goal of minimizing the total number of function evaluations required to solve computationally intensive environmental applications, eliminate trial-and-error analysis for setting the NSGA-II's parameters, and avoid the need for random seed analysis. The dynamic population sizing and injection approach applied in the **ε**-NSGA-II exploits computationally inexpensive small populations to expedite search while increasing population size commensurate with problem difficulty to ensure the Pareto optimal set can be reliably approximated.

The initial population size, N_0 is set to some arbitrary small value (*e.g,* 5), as it is expected that the adaptive population sizing scheme will adjust for an undersized population. A randomly selected subset of the solutions obtained using the small population sizes are injected into subsequent larger populations, aiding faster convergence to the Pareto front. This can be viewed as using series of "connected" NSGA-II runs that share results so that the Pareto optimal set can be reliably approximated. Computational savings should be viewed in two contexts: (1) the use of minimal population sizes and (2) elimination of random seed analysis. Note that the number of times the population size needs to be doubled varies with different random seeds, though exploiting search with small populations will on average dramatically reduce computational times. Moreover, our approach eliminates the need to repeatedly solve an application for a distribution of random seeds.

The NSGA-II's remaining parameters are set automatically based on whether an application is being solved using a real or binary coding. The 4-objective problem solved in this paper is solved using binary coded variables, uniform crossover with a probability of 0.5, and a probability of mutation equal to 1/population size [for more details see Reed et al. [2003]].

4.2 Archive Update

The **ε**-dominance archiving approach is particularly attractive for environmental applications because it allows the user to define the precision with which they want to quantify their tradeoffs while bounding the size of the archive and maintaining a diverse set of solutions. The concept of **ε**-dominance requires the user to define the precision they want to use to evaluate each objective. The user specified precision or tolerance vector **ε** defines a grid for a problem's objective space, which biases search towards the portions of a problem's decision space that have the highest precision requirements. The **ε**-dominance archive improves the NSGA-II's ability to maintain a diverse set of nondominated solutions by only allowing 1 archive member per grid cell. In the case when multiple nondominated points reside in a single grid cell, only the point closest to the lower left corner of the cell (assuming minimization) will be added to the on-line archive thereby ensuring convergence to the true Pareto optimal set Laumanns et al. [2002], Deb et al. [2002].

4.3 Injection and Termination

The **ε**-NSGA-II also seeks to speed convergence by pre-conditioning search with larger population runs with the prior search results attained using small populations. In prior efforts, any attempts to inject solutions found using small population into subsequent runs made the NSGA-II prematurely converge to poor representations of the Pareto optimal set, especially for problems with greater than 2 objectives. The **ε**-domination archive's ability to preserve diversity plays a crucial role in overcoming this limitation. As described previously in Section 3.1, the **ε**-NSGA-II begins search with an initial population of 5 individuals from which the **ε**-nondominated solutions identified in this initial run are stored in the archive.

The archive at the end of each run contains **ε**-nondominated solutions that can be used to guide search in future runs and speed up convergence to the Pareto front. This is achieved by injecting **ε**-nondominated solutions from the archive at the end of the run with population size *N* into the initial population of the next run which has a population size *2N*. Two scenarios arise when the **ε**-NSGA-II injects solutions from the archive generated with a population size *N* into the initial generation of a run with a population size *2N.*

In scenario 1, the archive size *A* is smaller than the subsequent population size *2N.* In this case, 100-percent of the **ε**-nondominated archive solutions are injected into the first generation of the subsequent run with *2N* individuals. We have found that the number of injected solutions should be maximized to aid rapid convergence. The **ε**-dominance archive in combination with successive doubling of population size guarantees the **ε**-NSGA-II will maintain sufficient solution diversity. In scenario 2, the archive size *A* is greater than the next population size *2N.* In this case, *2N* **ε**-nondominated archive solutions are selected randomly and injected into the first generation of the next run, again maximizing the impact of injected solutions.

The termination of search across all runs (i.e., across all populations used) compares the rate increase in the archive size at the end of two successive runs of the **ε**-NSGA-II in which the first run uses a population of N sampling designs to evolve a **ε**-nondominated set composed of A individuals, while the second run uses a population of 2N designs to evolve a **ε**-nondominated set of K individuals. The results of these runs are used in equation (1), to define which of the two following courses of action will be taken: (1) population size is again doubled, resulting in 4N individuals to be used in an additional run of the **ε**-NSGA-II or (2) the algorithm stops to allow the user to assess if the **ε**-nondominated set has been quantified to sufficient accuracy.

$$\begin{aligned} & if \quad \Delta < \left(\frac{|K - A|}{A}\right)100 \quad then \ , double \quad N \\ & and \quad continue \quad search \quad else \quad stop \quad search \end{aligned} \qquad (1)$$

The solutions obtained in the archive at the end of the final run represent the Pareto front with the user defined accuracy. In this study, Δ was set equal to 10-percent as recommended by Reed et al. [2003]. Section 4 demonstrates the efficiency of the ε-NSGA-II in solving high order environmental problems.

5. RESULTS AND DISCUSSIONS

The efficiency of the ε-NSGA-II in solving high order multiobjective optimisation problems is demonstrated in this section by solving the 4-objective ground water monitoring problem described in Section 3. The experiments designed in this section are aimed at highlighting the algorithms efficiency in capturing an approximate trade off using a minimum number of functional evaluations.

The problem was solved using various precision settings as described in Table 1. The resolution of the solutions obtained decreases from setting 1 to 5. Setting 1 is of the highest resolution and hence each of the ε's are set to arbitrarily small value values. This would be an ideal setting for a user who would like the entire trade off and does not mind compromising on the computational complexity. To highlight the efficiency of the algorithm in reducing the computational time with a decrease in resolution, four other settings are chosen. In setting 2, the values of ε_{SREE} and ε_{MAE} are set to 0.01 and 10^{-4} respectively based on the range of objective values obtained from setting 1. Setting 3 is obtained by increasing the value of $\varepsilon_{Uncertainty}$ by a factor of 100 while at the same time keeping the other resolutions' constant. Similarly settings 4 and 5 are obtained by increasing ε_{SREE} and ε_{MAE}. ε_{Cost} is not varied in the experiments because it is a discrete integer function of the number of sampling points used.

Table 1. The resolution settings used in solving the ground water problem.

Setting No.	ε_{Cost}	ε_{SREE}	$\varepsilon_{Uncertainty}$	ε_{MAE}
1	1	10^{-5}	0.01	10^{-6}
2	1	0.01	0.01	10^{-4}
3	1	0.01	1	10^{-4}
4	1	0.1	1	10^{-4}
5	1	0.1	1	0.01

The problem was solved using binary coded variables and the parameter settings described in section 4. Table 2 summarizes the number of solutions obtained and the number of function evaluations required for each of the parameter settings. The effects of varying the resolution on the number of solutions found by the ε-NSGA-II as well as the number of function evaluations required are demonstrated in Figure (1). The number of nondominated solutions found by the ε-NSGA-II does not significantly decrease when the resolution is reduced from the values used in setting 1 to those in setting 2. Setting 1 required the highest resolution, which as expected had the highest number of ε-nondominated solutions and reduced the number of function evaluations required to solve the problem by more than 80% over the prior published results Reed and Minsker [2004]. Figure (1) demonstrates that as user-specified resolution requirements decrease, the number of function evaluations reduces by an order of magnitude relative to the 450,000 evaluations utilized by Reed and Minsker [2004].

Table 2. The number of function evaluations and the number of solutions obtained for each of the parameter settings.

Setting No.	No. of Function Evaluations	No. of Solutions
1	82740	880
2	76500	862
3	69900	631
4	51985	515
5	37260	463

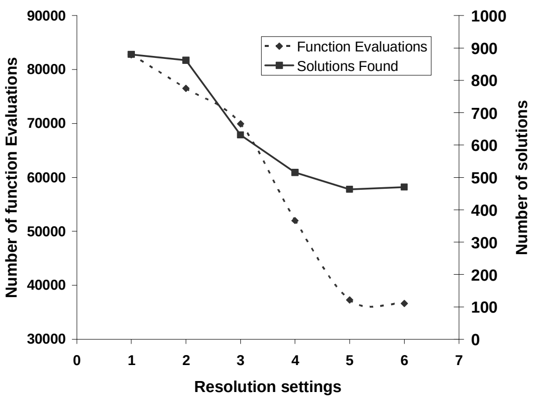

Figure 1. Variation of the number of solutions found and the number of function evaluations required with different resolution settings.

In the LTM application a single constituent is being monitored at 29 monitoring wells, which results in a decision space of more than 500 million possible sampling designs (i.e., 2^{29} sampling designs). Using the ε-NSGA-II to identify the subset of sampling designs that are optimal with respect to the 4 objectives used in this application, reduces the set of designs that must be considered from 500 million to less than a 1000 designs identified on the Pareto surface. Although the 4-dimensional Pareto surface cannot be visualized, the set of solutions designs can inform decision making as follows.

This process begins by analyzing pairs of the objectives that are known to conflict. These 2-dimensional tradeoffs are subsets of the overall nondominated set. These tradeoffs are found by identifying only those solutions that are nondominated in terms of cost and one other objective, independent of the remaining objectives' values [e.g., the Cost—Uncertainty tradeoff in Figure (2)].

Figures (2) shows the Cost—Uncertainty tradeoff generated by Reed and Minsker [2004] designated as the "prior method" as well as the results attained by the ε-NSGA-II using ε settings 1 and 5. The ε-NSGA-II closely approximates the results of the prior method using up to 90-percent fewer function evaluations. These results highlight that a tremendous amount of the computation time originally used to solve this application was spent seeking unnecessarily high-precision results. Non-domination sorting at 6-digits of precision and beyond is computationally expensive and does not significantly improve the representation of the tradeoffs used for decision making.

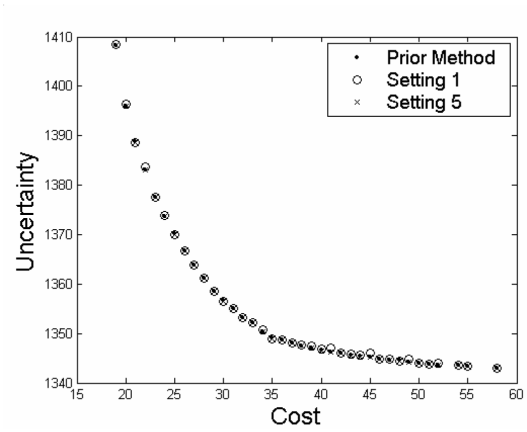

Figure 2. Cost—Uncertainty tradeoff.

6. CONCLUSIONS AND FUTURE WORK

The **ε**-NSGA2 demonstrates how **ε**-dominance archiving can be combined with a parameterization strategy for the NSGA-II to accomplish the following goals: (1) ensure the algorithm will maintain diverse solutions, (2) eliminate the need for trial-and-error analysis for parameter settings (*i.e.*, population size, crossover and mutation probabilities), and (3) allow users to *sufficiently* capture tradeoffs using a minimum number of design evaluations. A sufficiently quantified tradeoff can be defined as a subset of nondominated solutions that provide an adequate representation of the Pareto frontier that can be used to inform decision making. Results are presented for a 4-objective groundwater monitoring case study in which the archiving and parameterization techniques for the NSGA-II combined to reduce computational demands by greater than 90-percent relative to prior published results. The methods of this paper can be easily generalized to other multiobjective applications to minimize computational times as well as trail-and-error parameter analysis.

7. REFERENCES

Deb, K., M. Mohan and S. Mishra, A Fast Multi-objective Evolutionary Algorithm for Finding Well-Spread Pareto-Optimal Solutions. Kanpur, India, Indian Institute of Technology, 2003.

Deb, K., A. Pratap, S. Agarwal and T. Meyarivan, "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II." IEEE Trans. Evol. Computation **6**(2): 182-197, 2002.

Laumanns, M., L. Thiele, K. Deb and E. Zitzler, "Combining Convergence and Diversity in Evolutionary Multiobjective Optimization." Evolutionary Computation **10**(3): 263-282, 2002.

Maxwell, R., F. S. Carle and F. B. Tompson, Contamination, Risk, and Heterogeneity: On the Effectiveness of Aquifer Remediation. UCRL-JC-139664. Livermore, CA, 2000.

Reed, P., T. Ellsworth and B. S. Minsker, "Spatial Interpolation Methods for Nonstationary Plume Data." Ground Water **42**(2): 190-202, 2004.

Reed, P. and B. S. Minsker, "Striking the Balance: Long-Term Groundwater Monitoring Design for Conflicting Objectives." Journal of Water Resources Planning and Management **130**(2): 140-149, 2004.

Reed, P., B. S. Minsker and D. E. Goldberg, "Simplifying Multiobjective Optimization: An Automated Design Methodology for the Nondominated Sorted Genetic Algorithm-II." Water Resources Research **39**(7): 1196, doi:1110.1029/2002WR001483, 2003.

Task Committee on Long-Term Groundwater Monitoring Design, Long-Term Groundwater Monitoring: The State of the Art. Reston, VA, American Society of Civil Engineers, 2003.