

A multiobjective approach to cost effective long-term groundwater monitoring using an elitist nondominated sorted genetic algorithm with historical data

Patrick Reed, Barbara S. Minsker and David E. Goldberg

ABSTRACT

This study presents a methodology for quantifying the tradeoffs between sampling costs and local concentration estimation errors in an existing groundwater monitoring network. The method utilizes historical data at a single snapshot in time to identify potential spatial redundancies within a monitoring network. Spatially redundant points are defined to be monitoring locations that do not appreciably increase local estimation errors if they are not sampled. The study combines nonlinear spatial interpolation with the nondominated sorted genetic algorithm (NSGA) to identify the tradeoff curve (or Pareto frontier) between sampling costs and local concentration estimation errors. Guidelines are given for using theoretical relationships from the field of genetic and evolutionary computation for population sizing and niching to ensure that the NSGA is competently designed to navigate the problem's decision space. Additionally, both a selection pressure analysis and a niching-based elitist enhancement of the NSGA are presented, which were integral to the algorithm's efficiency in quantifying the Pareto frontier for costs and estimation errors. The elitist NSGA identified 34 of 36 members of the Pareto optimal set attained from enumerating the monitoring application's decision space; this represents a substantial improvement over the standard NSGA, which found at most 21 of 36 members.

Key words | groundwater monitoring, multiobjective optimization, genetic algorithm, nondominated sorted genetic algorithm, elitism, guidance

Patrick Reed (corresponding author)
4129 NCEL, MC. 250,
205 N. Mathews,
Urbana,
IL 61801, USA
Tel: +1 217 333 6979
Fax: 217-333-6968
E-mail: preed@uiuc.edu

Barbara S. Minsker
3230d NCEL, MC. 250,
205 N. Mathews,
Urbana,
IL 61801, USA

David E. Goldberg
117 Transportation, MC. 238,
104 S. Mathews,
Urbana,
IL 61801, USA

INTRODUCTION

Evolution-based multiobjective optimization (EMO) methods are pragmatic tools for solving problems with large decision spaces and conflicting objectives. Schaffer (1984) provided the seminal work within the EMO field in which a vector evaluated genetic algorithm (VEGA) was designed to search decision spaces for the optimal tradeoffs among a vector of objectives. Subsequent innovations in EMO have resulted in a rapidly growing field with a variety of solution methods that have been used successfully in a wide range of applications (for reviews see Fonseca & Fleming 1995; Coello 1999; Van Veldhuizen 1999). These solution methods have garnered increased attention over the past decade and have been applied in a variety of contexts within the water resources field.

Cieniawski (1993) is one of the earliest studies in water resources 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. Ritzel *et al.* (1994) compared the performances of VEGA, a domination ranking-based genetic algorithm (Pareto GA), and

mixed-integer chance constrained programming (MICCP) in solving a multiobjective groundwater pollution containment application. The Pareto GA was found to be superior to VEGA in its ability to evolve a better representation of the Pareto frontier. Additionally, the Pareto GA nearly replicated the Pareto frontier found using MICCP while taking less computational effort to solve the problem. Halhal *et al.* (1997) successfully incorporated Pareto domination ranking into the messy genetic algorithm (Goldberg *et al.* 1989) to quantify the tradeoffs in rehabilitating water distribution networks. Gupta *et al.* (1998) combined a downhill simplex method with an evolutionary search strategy implementing Pareto ranking to seek tradeoff solutions when calibrating hydrologic models. These studies show how EMO methods have been adapted to solve a variety of water resource applications. 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.

One of the difficulties in applying EMO methods is identifying parameter settings that ensure competent 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. A major goal of this study is to develop guidelines for using theoretical relationships from the genetic and evolutionary computation literature to ensure that the NSGA competently navigates the problem's decision space. These guidelines are then applied to quantify the tradeoffs implicit in designing sampling strategies for a long-term groundwater-monitoring network. Additionally, a new niching-based elitist enhancement of the NSGA is also presented, which substantially improves coverage of the Pareto frontier.

The test case developed in this study 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 hypothetical perchloroethylene (PCE) plume in a highly




Figure 1 | The 50 potential sampling locations (designated by the x's above) within a 20-well multilevel monitoring network.

heterogeneous alluvial aquifer. The hydrogeology of the test case is based on an actual site located at the Lawrence Livermore National Laboratory in Livermore, California, currently being managed under the United States' Comprehensive Environmental Response, Compensation and Liability Act (CERCLA) program. Data were provided for a total of 50 hypothetical sampling locations within the 20-well multilevel monitoring network shown in Figure 1. The data represent a snapshot in time, 10 years after an underground storage tank has continuously released contamination into the aquifer system. The monitoring wells can sample from 1–3 locations along their vertical axis and have a minimum spacing of 60 m between wells in the horizontal plane.

The site is assumed to be undergoing long-term monitoring, in which groundwater samples are used to assess the effectiveness of current remediation strategies. 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 in Figure 1 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 20–30 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 et al. 2000). These methods define sampling points to be redundant when they minimally affect interpolation-based 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 paper builds on these previous methods by introducing a sampling design methodology that explicitly and efficiently identifies the tradeoffs encountered when reducing monitoring costs.

METHODS

Problem description and formulation

To identify which wells are redundant, this study employs a local concentration 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:

$$\begin{aligned} \text{Minimize } F(\bar{x}_\kappa) &= [f_1(\bar{x}_\kappa), f_2(\bar{x}_\kappa)], \quad \forall \kappa \in \Omega \\ f_1(\bar{x}_\kappa) &= \sum_{i=1}^{n_{well}} C_S(i) x_{\kappa i} \\ f_2(\bar{x}_\kappa) &= \sum_{j=1}^{n_{est}} (c_{all}^*(\bar{u}_j) - c_{est}^\kappa(\bar{u}_j))^2 \end{aligned} \quad (1)$$

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

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

If the i th 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 \$365–1095 for 1–3 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 (n_{well}), 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 under different parameter settings 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 κ 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 estimates using data from all available sampling locations, $c_{all}^*(\bar{u}_j)$, and the estimates based on the κ th sampling plan $c_{est}^\kappa(\bar{u}_j)$ at each location \bar{u}_j in the interpolation domain. Each \bar{u}_j specifies the coordinates for the j th 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 the nonlinear spatial interpolation method described below.

Nonlinear spatial interpolation

The interpolation method used in this study is a variant of the scheme used by Barry & Sposito (1990) in their analysis of tracer plumes at the Borden site located in Ontario, Canada. The Barry & Sposito (1990) interpolation method was selected because it requires minimal modeling assumptions and it has been successfully applied to three-dimensional historical data. This study simplified Barry and Sposito's interpolation scheme to use 3 of the original 7 fitting parameters, resulting in a nonlinear least squares version of inverse distance weighting. Neglecting 4 of the parameters did not appreciably affect the cross-validation residuals and improved the algorithm's stability. Equation (3) shows the interpolation function used in this study,

which estimates PCE concentration $c_{est}^{\kappa}(\bar{u}_j)$ at each unknown location \bar{u}_j by the weighted sum of the $nsamp(\kappa)$ total samples $c(\bar{u}_{\psi})$ taken in the κ th sampling scheme:

$$c_{est}^{\kappa}(\bar{u}_j) = \sum_{\psi=1}^{nsamp(\kappa)} w(\bar{u}_j, \bar{u}_{\psi}) c(\bar{u}_{\psi}) / w_t \quad (3)$$

The weights for each of the samples in the κ th sampling plan are calculated as a function of the distance between the ψ th sample and the j th grid point in the interpolation domain as shown in Equations (4) and (5):

$$w(\bar{u}_j, \bar{u}_{\psi}) = 1 / D_{\kappa}(\bar{u}_j, \bar{u}_{\psi}) \quad (4)$$

$$D_{\kappa}(\bar{u}_j, \bar{u}_{\psi}) = [(u_{j1} - u_{\psi1})^2 / \alpha_1 + (u_{j2} - u_{\psi2})^2 / \alpha_2 + (u_{j3} - u_{\psi3})^2 / \alpha_3]^{0.5} \quad (5)$$

The w_t factor is the sum of the $nsamp(\kappa)$ weights calculated using Equation (4) and serves to scale the system such that the weights sum to one. The parameters α_1 , α_2 and α_3 are fitting parameters that were used to minimize the cross-validation residuals for each of the sampling designs considered in this study. See Barry & Sposito (1990) for details on the cross-validation estimation method used in this study to fit the α parameters shown in Equation (4) to the data provided by each sampling plan.

The basics of the NSGA

The NSGA utilizes the Darwinian process of natural selection to effectively search for solutions that are optimal across a vector of objectives. The algorithm is very similar in form to the simple genetic algorithm (sGA), in that it exploits the operators of selection, crossover and mutation when building a set of optimal solutions. The performance of both algorithms can be described using the building block (BB) theory presented by Holland (1975) and Goldberg (1989). For both the NSGA and sGA, highly fit designs have a higher probability of being selected to mate and pass their traits (or BBs) to succeeding generations. Stochastic remainder selection was used in this study as recommended in Srinivas & Deb (1995). The operators of

crossover and mutation are identical for the two algorithms. Crossover (occurring with probability p_c) exerts an innovative force on the system by allowing favorable traits from parent designs to be juxtaposed in offspring that possess higher relative fitness values (Thierens & Goldberg 1993). Mutation locally refines solutions by randomly changing bit values (with a probability of p_m) from 0 to 1 or vice versa within a design's genotype. The difference between the NSGA and the sGA lies in how fitness is assigned. Unlike the sGA, the NSGA evaluates sampling designs in terms of a vector of objectives. A sampling design cannot be assessed in terms of its performance in any single objective because it may perform poorly with respect to the remaining objectives. The NSGA employs the concepts of *Pareto dominance* and *niching* to assign fitness values to sampling designs in two steps described below (Srinivas & Deb 1995).

The first step in fitness assignment employs the *Pareto dominance* concept defined in Equation (6), using notation adapted from Van Veldhuizen & Lamont (2000):

$$F(\bar{x}) < F(\bar{x}'), \text{ iff } (\forall \beta, f_{\beta}(\bar{x}) \leq f_{\beta}(\bar{x}') \wedge \exists \beta \in (1, n_{obj}): f_{\beta}(\bar{x}) < f_{\beta}(\bar{x}')) \quad (6)$$

Equation (6) states that a design \bar{x} dominates another design \bar{x}' (represented by $F(\bar{x}) < F(\bar{x}')$) if and only if it performs as well as \bar{x}' in all n_{obj} objectives and better in at least one (assuming minimization of all objectives). The NSGA identifies nondominated individual designs within the current population and assigns an arbitrary dummy fitness value to each of them. These individuals are all initially assigned the same fitness value to ensure that they have an equal likelihood of being selected to pass their traits unto latter generations; these designs compose the first nondominated front.

The second step in fitness assignments for the individuals in the first front utilizes the concept of *niching* (for more details see Goldberg & Richardson 1987; Deb & Goldberg 1989; Mahfoud 1995; Horn 1997) to ensure that the NSGA finds a diverse set of solutions defining the entire extent of the tradeoff (or Pareto front) among the objectives. Extending the 'natural selection' analogy to include the phenomena of *niching* and *speciation* helps to satisfy this goal. Horn (1997) defines niching as a 'form

of cooperation' where there is '... the localization of competition around finite, limited resources (niches), resulting in the lack of competition between such areas, and causing the formation of species for each niche'. Niching allows the NSGA to form stable subpopulations of sampling designs (species), each of which are well adapted to search for nondominated solutions specific to subspaces (niches) in Ω . The NSGA elicits niche formation by treating dummy fitness as a limited resource, which is shared by sampling designs using the relationship given in Equation (7) (Goldberg & Richardson 1987; Srinivas & Deb 1995):

$$Sh[d(\kappa, \kappa')] = \begin{cases} 1 - \left(\frac{d(\kappa, \kappa')}{\sigma_{share}} \right)^2, & \text{if } d(\kappa, \kappa') < \sigma_{share} \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

The value of the fitness sharing function $Sh[d(\kappa, \kappa')]$ ranges between 0 and 1 for two sampling designs κ and κ' depending on the ratio of their distance from one another $d(\kappa, \kappa')$ and the niche size defined by the niche radius σ_{share} . The sharing function measures the similarity between designs ranging from completely dissimilar ($Sh = 0$) to being identical ($Sh = 1$). The distance term $d(\kappa, \kappa')$ can measure similarity in several ways that are discussed in the next section.

The third step in assigning the fitness values to the nondominated set identified in step 1 consists of dividing each individual's dummy fitness by the sum of all of its sharing function values (termed the niche count) as shown in Equation (8):

$$niche\ count(\kappa) = \sum_{\kappa'=1}^N Sh[d(\kappa, \kappa')], \quad \forall \kappa \quad (8)$$

When these three steps are completed, the individuals in the first front are removed from consideration and the remaining population members undergo nondomination ranking and sharing to yield the second nondominated front. These individuals are assigned the minimum shared fitness value of the first front decremented by a value ΔD_f . This process continues until the entire population is assigned to the n_f th front and given a shared dummy fitness value. After shared dummy fitness values have been

assigned, the NSGA is operated the same as the sGA, constructing the Pareto front using the traditional operators of selection, crossover and mutation. See Goldberg (1989) for more background on these operators.

Competent design of the NSGA for search and optimization

Construction of the Pareto frontier requires that the NSGA be competently designed to navigate a problem's decision space. Competent design of the NSGA entails careful consideration of the factors controlling the algorithm's performance. Zitzler *et al.* (2000) concluded that population sizing and elitism are the most influential factors influencing the performance of EMO methods. Elitist operators provide a means of ensuring that the best individuals are identified and allowed to pass their traits to latter generations. Unlike sGA applications, EMO methods cannot simply pass a single individual with the current best objective 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. The next four subsections of this paper describe an NSGA design methodology that integrates several relationships from the genetic and evolutionary literature to identify effective population sizes and determine the niching parameters. This methodology is meant to serve as a guide for applying the NSGA to other applications, enabling efficient identification of a competent algorithm. Additionally, a new elitist strategy is presented with the aim of using niching to guide the selection of elite population members and maximize *online* performance (the term 'online performance' means that only the nondominated individuals within a single generation are considered when assessing the NSGA's performance). Online performance was used in this study to avoid the need for external bookkeeping of the nondominated individuals found across all generations, which is required by *offline* performance analysis (Zitzler *et al.* 2000). Performance is finally discussed in the context of selection pressure in the fourth and final section of the NSGA design methodology.

Population sizing: ensuring genetic diversity

The first step in designing a competent NSGA is to perform initial problem analysis to determine a range of population sizes. The goal of this initial step is to provide a means of selecting the best population size for our application with respect to a single random seed. Computationally intensive objective functions that occur in many water resources applications require that effective parameter settings for the NSGA be identified using a minimum number of runs, because a single run can take days or even weeks. The range of population sizes considered in this study were attained using relationships from Mahfoud (1995), who derived population-sizing relationships for genetic algorithms employing niching to solve multimodal problems. The relationships used in this study were designed for applications with multiple optima that have *identical fitness values*. Recall that the NSGA employs niching using members of the same non-dominated set with *identical dummy fitness values*, which makes the population-sizing relationships in Mahfoud (1995) relevant to NSGA applications. Equations (6)–(8) are theoretical equations derived from the schema theorem developed by Holland (1975) to model GA's performance as a function of their primary operators (for more details see Holland 1975; Goldberg 1989; Mahfoud 1995).

Equation (9) gives population size estimates assuming that crossover will not disrupt the traits (or BBs) required to assemble optimal solutions:

$$N \geq q \left[-\ln \left(\frac{1 - \gamma^{1/g}}{q} \right) \right] \quad (9)$$

The relationship was derived as the minimum population size N required to maintain q niches for g generations with reliability γ for the special case when the niches have identical fitness values. Disruption adversely affects the NSGA's performance when crossover between a pair of parent strings destroys traits that are essential to the evolution of nondominated individuals in subsequent generations (for more details see Mahfoud 1995; Zitzler *et al.* 2000; Van Veldhuizen & Lamont 2000). Equation (10) accounts for the potential disruptive

effects of crossover in population size estimates (Mahfoud 1995):

$$N \geq 2 \ln \left(\frac{1 - \gamma^{1/g}}{q} \right) \left[\ln \left[(1 - p_c) \left(1 - \frac{1}{q} \right)^2 + p_c \left(1 - \frac{1}{q^2} \right) - 2p_c(1 - p_d) \frac{1}{q} \left(1 - \frac{1}{q} \right) \right] \right] \quad (10)$$

Equation (10) represents the minimum population size N required to maintain q niches for g generations with reliability γ given a probability of crossover p_c and probability of disruption p_d . The probability of disruption p_d can be conservatively estimated using Equation (11), which assumes the maximally disruptive operator of uniform crossover. In uniform crossover, the binary values at each bit position in two individual strings is swapped with probability p_c :

$$p_d = 1 - \left(\frac{1}{2} \right)^{O(m)} \quad (11)$$

For the binary-coded NSGA, the decision variables (\bar{x}_x in Equation (1)) representing a potential design (termed its *phenotype*) are converted to binary representations and concatenated into a string of binary variables (termed the design's *genotype*) of length l . The traits or (BBs) are actually small subsets of binary digits from each design's genotype. Equation (11) computes the probability that the BBs will be disrupted as a function of their order $O(m)$, which is the smallest number of fixed value digits in a design's genotype that are relevant to the final solutions of the problem.

As was the case in the sGA design methodology presented by Reed *et al.* (2000b), the actual size of BBs in an engineering application, $O(m)$, are unknown, requiring the practitioner to employ conservative assumptions and problem-specific information to calculate a potential range of population sizes from Equations (9) and (10). The p_d was set by assuming that the BB order $O(m)$ may range from 1 to 5 because higher-order BBs will be disrupted by the operators of selection, crossover and mutation (Goldberg 1989), following a similar approach presented by Reed *et al.* (2000b) for designing sGAs.

The remaining parameters in Equations (9) and (10) were set as follows. The population size calculations assumed a reliability of 85% ($\gamma = 0.85$). The number of generations until convergence g can be estimated to fall within the range $[2l, llnl]$ where l is the binary string length of each design (Thierens & Goldberg 1994; Thierens *et al.* 1998). Mahfoud (1995) states that Equations (9) and (10) are not sensitive to the parameter value g and recommends that the user set the value of this parameter to be greater than or equal to the number of generations they expect to run the algorithm. For these reasons, g was set equal to $2l$ (40 generations) for this study. A majority of EMO applications use a value for p_c falling within the range $[0.6, 0.9]$ (Fonseca & Fleming 1995; Horn 1997; Coello 1999; Van Veldhuizen 1999). In this application, p_c was set equal to 0.6 to reduce the potential for disruption and the resulting population size estimates.

The number of niches q was set using information specific to the tradeoff being sought between cost and SREE in this application. The discussion of the spatial interpolation stated that the number of PCE samples $c(\bar{u}_\psi)$ ranged between a minimum value of 4 and a maximum of 50, yielding a tradeoff with a maximum of 46 discrete cost levels. The goal of this application is to find the minimum SREE for each of these discrete cost levels. The number of niches was set equal to 40, which represents a goal of attaining over 85% of the points in the Pareto frontier.

Using these parameters in Equations (9)–(11) yielded six population size estimates, ranging from a minimum value of 370 to a maximum of 870. The lower bound population size estimate of 370 assumes that crossover will not disrupt BBs. The remaining five population size estimates account for disruption of BBs with orders, $O(m)$, ranging from 1 to 5. Given this range of population sizes, the computational complexity of using the NSGA to solve this problem was estimated by multiplying N by the number of generations g to attain the total number of function evaluations. Function evaluations took an average time of 0.04 s on a Dell XPS T800r running Windows NT, yielding estimated total run times between 10 and 25 min.

Finally, it should be noted that both Equations (9) and (10) assume that mutation is minimally disruptive. This

assumption was accounted for by setting p_m to be equal to $1/N$ (DeJong 1975), which limits the number of mutations to l and reduces the operator's influence as N increases. In this application, the optimal population was determined by running the NSGA for each of the six estimates and choosing the run that best defined the full extent of the tradeoff between cost and SREE.

Sharing: sizing of the niche

An important component of ensuring that the NSGA converges to the Pareto frontier is properly setting the size of niches represented by σ_{share} in Equation (7). Deb & Goldberg (1989) provide guidance for setting this parameter in both phenotypic and genotypic space. Recall that a sampling design's phenotype is the floating-point representation of each of its decision variables. Additionally, the design's genotype is its concatenated binary representation (or its chromosome). Equation (7) requires that the distance $d(x, x')$ between two individuals be compared to σ_{share} in order to compute the sharing function values used to assign fitness. If the individuals are within σ_{share} distance of each other, their fitness is penalized to encourage future generations to spread along the entire Pareto frontier.

Two different distance metrics are used to compare the relative magnitudes of $d(x, x')$ and σ_{share} . In genotypic space, $d(x, x')$ is equal to the Hamming distance between the chromosomes representing individuals x and x' . The Hamming distance is simply the total number of positions in the binary strings that have different values. In phenotypic space, the Euclidean distance metric shown in Equation (12) is used to calculate the distance between individuals:

$$d(\kappa, \kappa') = \sqrt{\sum_{i=1}^{npar} (x_{i,\kappa} - x_{i,\kappa'})^2}, \forall \kappa \quad (12)$$

For this application, the total number of decision variables ($npar$) is equal to the number of monitoring wells, $nwell$ in Equation (1).

To determine the appropriate size of σ_{share} for a particular problem, three approaches can be taken. These approaches are described below in order of increasing

domain-specific knowledge required. Deb & Goldberg (1989) derived Equation (13) as a guide for practitioners for sizing niches in genotypic space by deriving σ_{share} to represent ‘... the maximum bits of difference allowed between the strings to make q -subspaces [or niches] in the solution space’:

$$\frac{1}{2^l} \sum_{b=0}^{\sigma_{share}} \binom{l}{b} = \frac{1}{q} \quad (13)$$

Equation (13) assumes that the niches are uniformly spaced, with each niche apportioned $1/q$ of the decision space (Deb & Goldberg 1989). The equation represents a binomial distribution with a probability of 0.5, which can be solved for σ_{share} using the cumulative binomial distribution tables when string lengths l are less than or equal to 25 bits (indexed by b above). In this study, the chromosomes have a length of $l=20$ and the number of niches $q=40$, which resulted in a genotypic σ_{share} equal to 5. For longer string lengths, the binomial distribution can be approximated with the normal distribution. In this case, Equation (14) from Deb & Goldberg (1989) can be used, where the ‘... z^* corresponding to the fraction $1/q$ may be found from a cumulative normal distribution chart’:

$$\sigma_{share} = \frac{1}{2} (l + z^* \sqrt{l}) \quad (14)$$

The second approach for selecting σ_{share} is for phenotypic space. Deb & Goldberg (1989) derive σ_{share} in phenotypic space where ‘... each niche is enclosed in a p -dimensional hypersphere of radius σ_{share} such that each sphere encloses $1/q$ of the volume of the space’. Equation (15) is the resulting expression for the phenotypic niche radius, which was found to have a value of 1.859 in this study:

$$\sigma_{share} = \sqrt{\frac{\sum_{i=1}^{npar} (x_{i,max} - x_{i,min})^2}{npar}} \left| 2q^{1/npar} \right. \quad (15)$$

Phenotypic sharing has been shown to outperform genotypic sharing in most previous applications (Fonseca & Fleming 1995; Mahfoud 1995; Horn 1997; Coello 1999; Van Veldhuizen 1999). A detailed explanation of the

performance difference between the two sharing methods is given in Deb & Goldberg (1989), who found that ‘... a lower bound on the Hamming difference between the strings’ causes the inability of genotypic sharing to form stable subpopulations in applications where non-dominated individuals are not assigned identical fitness values.

Finally, an additional form of phenotypic sharing has been employed successfully in other applications (Fonseca & Fleming 1995; Horn 1997) where the distance metric between individuals is calculated in objective space. Horn (1997) provides guidance to practitioners for sizing niches in objective space using a geometric approach similar to the Deb & Goldberg (1989) phenotypic niche sizing method given in Equation (15). The method requires that the maximum and minimum values of the objectives be known *a priori*. This approach was not used in this study because of the discrete nature of the sampling costs. Moreover, the maximum value of the SREE could not be calculated *a priori*.

The niche sizing relationships given by Deb & Goldberg (1989) and described above have been shown to be robust for a variety of problems (Goldberg *et al.* 1992; Horn 1997) and will be analyzed in greater detail in the results section of this study.

Elitism: seeking the king of the niche

Zitzler *et al.* (2000) showed that elitism is one of the most important factors affecting the performance of EMO methods. 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 minimizing any external book-keeping and maximizing the online 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 non-dominated 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} , the elite radius, which is a parameter that allows the user to easily manipulate the amount of elitism. The elite radius defines the distance (either genotypic or phenotypic) 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_{β} for β equal 1 to n_{obj}

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_{β} 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 σ_{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 within that population. 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 where this rule-of-thumb may not work as effectively.

The importance of elitism in the performance of the NSGA Zitzler *et al.* (2000) relates directly to the concepts of *genetic drift* and *selection pressure*. Genetic drift occurs when promising solutions in the population do not experience sufficient selection pressure and converge to non-optimal values (drift stall) under the influence of crossover and mutation. Elitism increases selection pressure by ensuring that the traits of nondominated individuals remain in the population for use in later generations.

Selection pressure: avoiding drift stall

The importance of elitism in the performance of the NSGA motivated a further analysis of the role that selection pressure has on the algorithm's performance. Stochastic remainder selection selects a particular individual using the ratio of the individual's fitness and the average fitness of the current population in generation t . Recall from the introduction to the NSGA above that selection is based on dummy fitness values, which are decremented by a constant value ΔD_j for each of the n_j successive fronts within the population. For example, analysis of the initial random population of 760 individuals used in the elitist runs discussed in the results section shows that the expected number of copies of members in the nondominated set in the next generation had an average value of 1.7. The expected number of members from the 10th and 20th fronts in the next generation was 1.45 and 1.25, respectively. Although members of dominated fronts are expected to receive fewer copies than the nondominated front, the relative difference between 1.7 and 1.25 is small, which gives rise to an increased potential for drift stall to occur.

This increased potential for drift stall motivated an analysis of the effect of scaling the fitness of successive

fronts on the NSGA's performance. The dummy fitness values for successive fronts are scaled using a scaling coefficient S_c whose value is less than one, such that the minimum fitness in each front is guaranteed to be at least $(1 - S_c)\%$ higher than the maximum fitness in the front that immediately succeeds it. The scaling-based fitness assignments replace the constant decrement ΔD_f assignments used previously in this work. S_c was set to be equal to 0.9, which ensures that the minimum fitness for front $(n_f - 1)$ is at least 10% greater than the maximum fitness in the n_f th front. For the random population of 760 designs discussed above, the expected numbers of individuals in the next generation from the 1st, 10th and 20th fronts will then be 6.2, 2.27 and 0.78, respectively. Note that scaling the system in this manner exponentially decreases the fitness of the members of the fronts succeeding the non-dominated set. Setting S_c requires striking a balance between maintaining a diverse population and ensuring that selection pressure is sufficient to prevent drift stall. Further discussion on this issue is given in the results section below.

Defining a measure of relative performance

To compare performance of the NSGA under different parameter settings, a measure of the algorithm's performance must be defined. In this paper, the performance of the NSGA as a function of its parameters was measured relative to the Pareto frontier for cost and SREE. The frontier is shown in Figure 2, which shows the 36 sampling designs that compose the Pareto optimal set identified using enumeration of the more than 1 million potential designs in Ω . The performance of the NSGA was quantified using a relative scoring metric (RSM) that measures the deviation of the nondominated set in generation t from the true front using Equation (16):

$$\begin{aligned}
 \text{Deviation}(\eta) = & \\
 & \begin{cases} |SREE_{true}(x_\eta) - SREE(\phi(t))| / SREE_{max}, \\ \text{if } f_1(\bar{x}_\eta) = f_1(\phi(t)) \forall \eta, \phi(t) \\ 1, \text{ otherwise} \end{cases} \\
 \text{RSM} = & 1 - \left(\sum_{\eta=1}^{36} \text{Deviation}(\eta) \right) / 36
 \end{aligned} \quad (16)$$

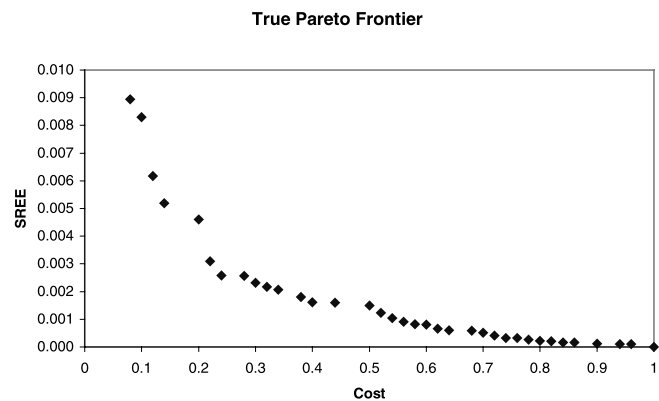


Figure 2 | Actual tradeoff between the squared relative estimation error and cost where each objective has been scaled to fall within the interval [0, 1]. The maximum and minimum SREE values found in Ω were 0 and 13,000,000, respectively. The maximum and minimum cost values found in Ω were \$18,420 and \$1473, respectively.

Although the maximum SREE value ($SREE_{max}$) cannot be calculated *a priori*, its maximum value was recorded while performing the enumeration of the decision space. Equation (16) requires the maximum value because the SREE values have to be scaled such that they fall within the interval [0,1]. Equation (16) defines the deviation between the η th member of the enumerated Pareto optimal set and the ϕ th member of the current nondominated set in generation t to be equal to the absolute difference of their SREE values ($SREE_{true}(\bar{x}_\eta) - SREE(\phi(t))$ shown above) 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 (16) 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.

RESULTS AND DISCUSSION

The guidelines presented in the previous section identify a range of potential population sizes and compute the niching parameters used by the NSGA to search a diverse set of stable subpopulations for the Pareto frontier. Additionally, the niching parameters are integral to properly setting the




Figure 3 | Selection of the proper population size for both the NSGA and the elitist NSGA using the best RSM value attained after 40 generations.

elitist selection pressure. The subsequent sections analyze the effectiveness of these guidelines by presenting the results of over 400 trial runs elucidating the performances of both the NSGA and the elitist NSGA as a function of their parameters for the long-term monitoring design application. The results are presented for a single random number seed because the methodology presented in this study is meant to identify effective parameter settings for a single seed. Such an approach is intended for practitioners solving computationally intensive applications, who cannot perform numerous runs with multiple random number seeds.

Performance of the NSGA and the elitist NSGA

Figure 3 shows the performances of both the NSGA and the elitist NSGA for the range of population sizes attained from Equations (9)–(11) when implementing either genotypic or phenotypic sharing. It is readily apparent from the figure that elitism greatly improves the online performance of the NSGA, regardless of the population size used. Figure 3 was used to select the optimal population size for each of the four forms of the NSGA being considered. The performance trends for the NSGA and the elitist NSGA are nearly identical for both sharing methods. The NSGA's population sizes were set equal to 830 and 660 for genotypic and phenotypic sharing, respectively. These population sizes are consistent with previous studies,

which have found that phenotypic sharing outperforms genotypic sharing. Deb & Goldberg (1989) argue that the reduced performance of genotypic sharing is caused by increased sensitivity of Hamming distance calculations to the assumption that each niche is uniformly apportioned $1/q$ of the decision space, as required by Equation (13). Mahfoud (1995) states that, in addition to the sensitivity of genotypic sharing to the uniformity assumption, genetic drift and population sizing are influential in the performance differences between the two sharing schemes. Larger population sizes are required to ensure that important subpopulations (or niches) receive sufficient selection pressure and are not lost to the 'noisy discrimination of genotypic sharing' (Mahfoud 1995).

Figures 3 and 4 confirm that, for the NSGA to have comparable performance under both sharing schemes, the algorithm requires larger population sizes for genotypic sharing. Figure 4 shows that the NSGA has relatively poor online performance for discerning the true tradeoffs between costs and SREE for the monitoring application presented in this study, regardless of the sharing scheme invoked. It is important to note that genotypic sharing requires the least amount of problem-specific information of those sharing methods discussed in this study, but the inherent tradeoff of the NSGA using this form of sharing lies in the increased computational demands required due to the increased population size necessary for sharing in genotypic space. This tradeoff is not present for the elitist NSGA where Figure 3 clearly shows that the optimal population size is clearly defined by the peak when N is equal to 760, regardless of the sharing method used.

Figure 5 shows that the niching-based elitist strategy presented in this study greatly improved the performance of the NSGA. The algorithm found solutions along the full extent of the Pareto front and was able to identify 17 of 36 members of the Pareto optimal set exactly, regardless of the sharing method used. Figures 3 and 5 show that elitism served to both improve the online performance of the algorithm and help to overcome the 'noisy discrimination of genotypic sharing' relative to sharing in phenotypic space.

Figure 6 shows the performance of the elitist NSGA for the full range of possible values that both σ_{elite} and σ_{share} can be assigned under both sharing schemes (when




Figure 4 | NSGA's performance relative to the true Pareto front using both (a) genotypic and (b) phenotypic sharing. The circles represent the best nondominated front attained in 40 generations.

$N = 760$). Several observations are apparent from analysis of these plots. First, note that decreasing σ_{elite} from its maximum to its minimum value shows the NSGA's behaviour between the extremes of having no elitism (at the maximum value) to the case when all nondominated individuals are selected for elitist reproduction (at the minimum value). Both plots show that the NSGA's performance improves with increasing elitist selection pressure. Figure 6 also confirms the noisy nature of sharing in genotypic versus phenotypic space. Under genotypic sharing, Figure 6(a) shows several small peaks in performance for relatively sporadic combinations of niching and elitist parameter settings. In contrast, Figure 6(b) shows that sharing in phenotypic space yields a smoother surface

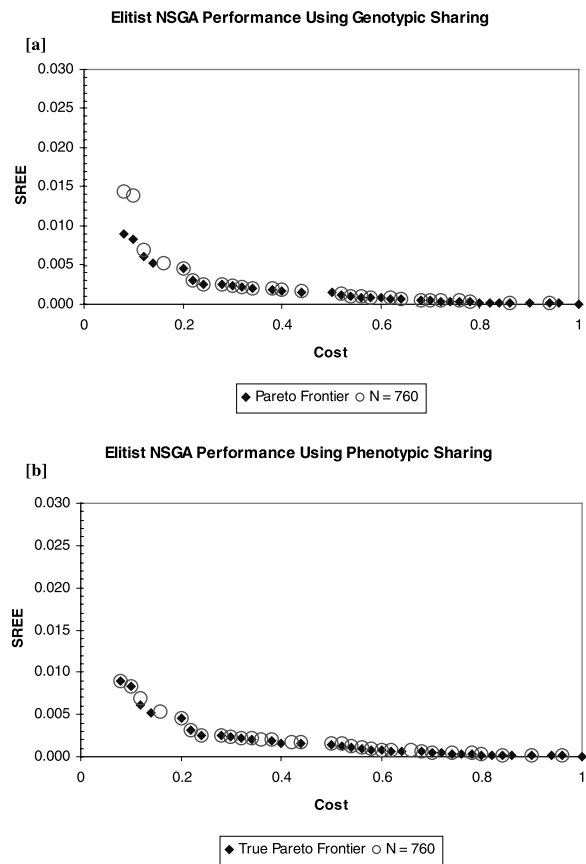


Figure 5 | Elitist NSGA's performance relative to the true Pareto front using both (a) genotypic and (b) phenotypic sharing. The circles represent the best nondominated front attained in 40 generations.

with a well-defined peak relative to the surface in Figure 6(a). Both plots serve to verify the guiding relationships discussed in the methods sections for population sizing, niching and elitist selection. The arrows in the figure designate the algorithm's performance when both σ_{elite} and σ_{share} are set using Equations (12) and (14) taken from Deb & Goldberg (1989).

Figures 4 and 5 show that both the NSGA and the elitist NSGA solutions have gaps in the extreme regions of the Pareto frontier. These gaps occur because the subpopulations or niches representing these portions of the frontier were not able to survive for successive generations and are symptomatic of genetic drift stall. Genetic drift also appears to be prevalent in Figure 4, which shows that the NSGA (in the absence of elitism) is unable to identify

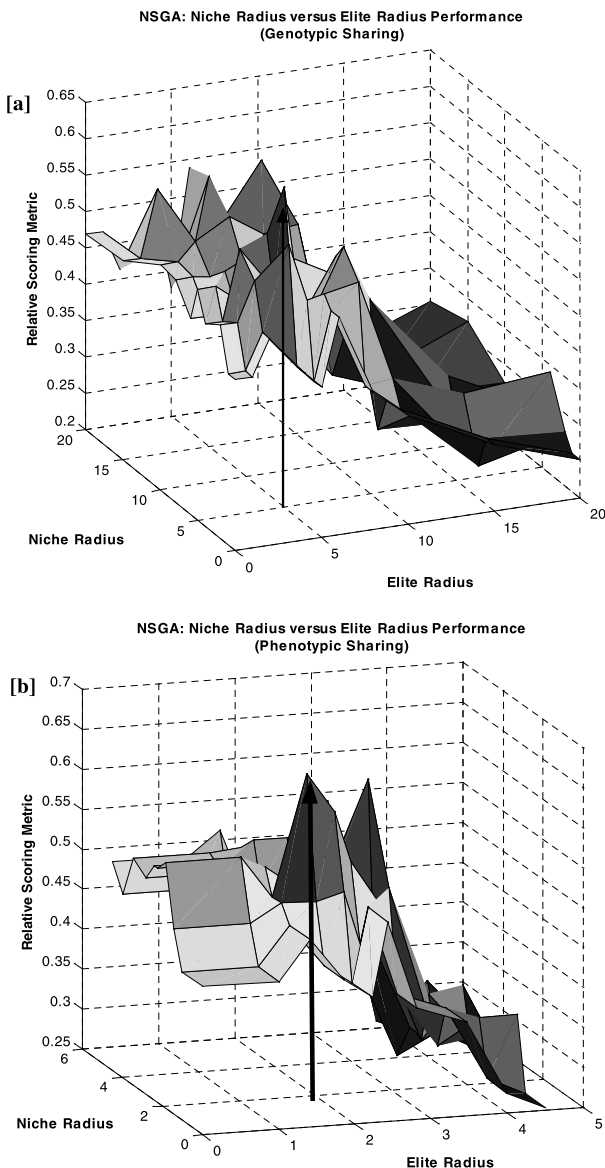


Figure 6 | Analysis of the NSGA's performance for the full range of possible values that the parameters controlling niching and elitism can be assigned when using (a) genotypic and (b) phenotypic sharing. The arrows designate the algorithm's performance when these parameters are set equal to the recommended niche radius attained from the relationships presented by Deb & Goldberg (1989).

more than 3 members of the Pareto optimal set, regardless of the sharing method invoked. Figure 7 shows the performance of all forms of the NSGA as a function of time. The figure shows that elitism allows the NSGA to steadily improve performance over the duration of the run. In the

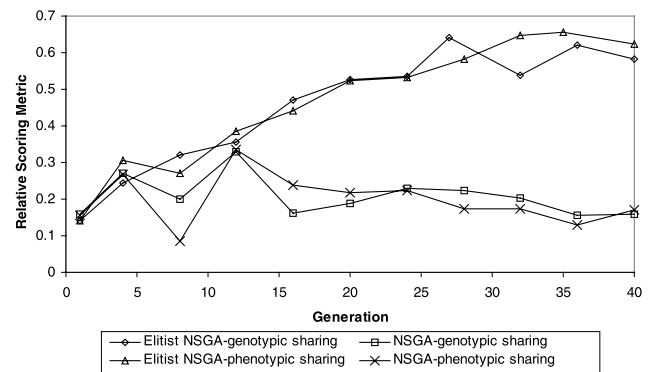


Figure 7 | Performances of both the elitist NSGA and the NSGA as a function of generation.

absence of elitism, the sporadic peaks in performance are steadily degraded over the course of the runs, confirming that drift stall is occurring. Recall that the NSGA decrements the dummy fitness values for successive fronts by a constant value ΔD_f and that stochastic remainder selection selects population members based on the ratio of their dummy fitness relative to the population's average dummy fitness. Drift stall is occurring because the constant decrement ΔD_f does not properly scale the ratio of the non-dominated individuals' fitness values relative to the population's average fitness, making it impossible for the NSGA to distinguish and propagate these individuals into successive generations. These observations motivated further analysis of the influences of selection pressure and drift stall on the algorithm's performance in the following subsection.

Performance under increased selection pressure

The results of the previous subsection showed that both the NSGA and the elitist NSGA converged to non-dominated fronts with gaps in the extreme portions of the Pareto frontier. Additionally, Figure 7 shows that the NSGA's performance only sporadically improves and generally degrades over the duration of the runs, regardless of the sharing method considered. These observations confirm that several of the niches are either being lost or converging to nonoptimal values due to the absence of sufficient selection pressure. Furthermore, it is readily




Figure 8 | Performances of both the elitist NSGA and the NSGA as a function of generation after rescaling the fitness assignments.

apparent that elitism significantly improves the algorithm's performance, confirming the results of Zitzler *et al.* (2000). Elitism serves to increase the selection pressure on the niches and reduce the influence of genetic drift on the system, which in effect is analogous to rescaling the relative fitness values of the population. Modifying the NSGA and elitist NSGA such that successive fronts have exponentially decreasing fitness using the scaling coefficient S_c (as described in the methods section) enabled further analysis of the influence of increased selection pressure on performance. For the monitoring application, analysis of the NSGA's performance for scaling coefficient values ranging from 0.7–0.95 showed that S_c should be set to approximately 0.9. Values for S_c lower than 0.9 greatly reduced the diversity of the populations and caused premature convergence (due to excessive selection pressure). Scaling coefficient values above 0.9 resulted in significant gaps in the extreme regions of the Pareto frontier due to insufficient selection pressure and genetic drift.

Figure 8 shows the performances of the scaled NSGA (sNSGA) and the elitist sNSGA as a function of time when $S_c = 0.9$. The population sizes used in these runs remained the same as those in the previous subsection for all forms of the sNSGA except the non-elitist, genotypic sNSGA, where a population size of 870 was found to be optimal, which represents only a slight increase from the previous section. Rescaling was not able to reduce the sensitivity of genotypic sharing to nonuniformities in the niche spacing,

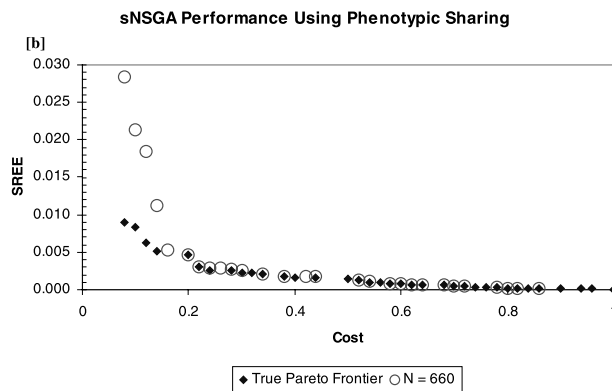



Figure 9 | sNSGA's performance relative to the true Pareto front using both (a) genotypic and (b) phenotypic sharing. The circles represent the best nondominated front attained in 40 generations.

although elitism successfully closed the performance gap between the two sharing methods for the elitist sNSGA. It is immediately obvious from Figure 8 that rescaling the fitness assignments greatly improved the performance of the sNSGA, which no longer shows degradation in performance with time and instead steadily improves towards an upper bound in performance. Note that the combination of rescaling and elitism resulted in the elitist sNSGA exceeding the previous results shown in Figure 7 for the elitist NSGA. The phenotypic elitist NSGA was able to reach a maximum RSM value of 0.66 whereas the genotypic elitist sNSGA attained a maximum RSM value of 0.8.

Figures 9 and 10 show the improved performances of all forms of the NSGA that result from rescaling the fitness assignments. Recall that previously the non-elitist NSGA

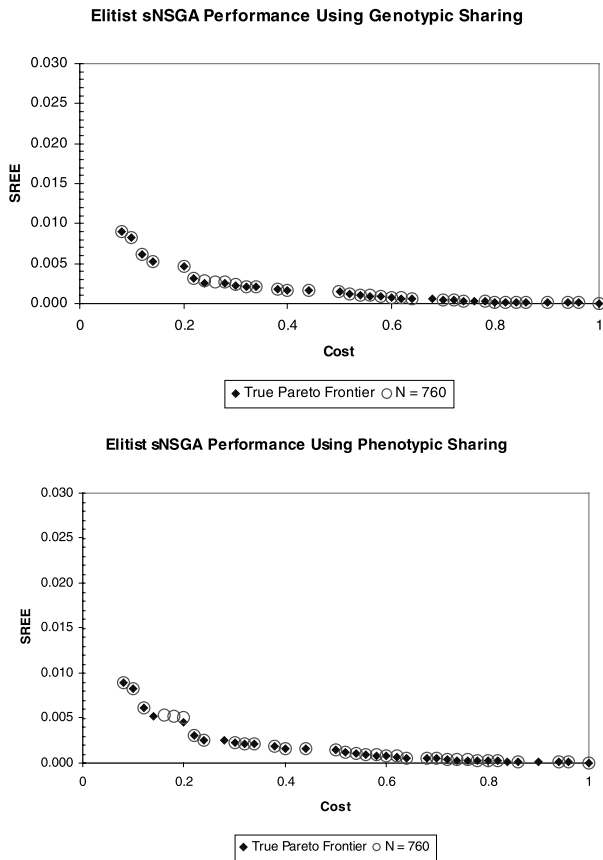


Figure 10 | Elitist sNSGA's performance relative to the true Pareto front using both (a) genotypic and (b) phenotypic sharing. The circles represent the best nondominated front attained in 40 generations.

was able to find a maximum of 3 members of the Pareto optimal set. Rescaling the shared dummy fitness values resulted in finding 21 and 17 members of the Pareto optimal set when implementing genotypic and phenotypic sharing, respectively. The additional use of elitism enabled the sNSGA to almost replicate the Pareto frontier for both sharing methods, as shown in Figure 10. The genotypic elitist sNSGA provided the most complete representation of the frontier by exactly finding 29 members of the Pareto optimal set while also finding close representations of all but 2 of the remaining 36 members. The phenotypic elitist sNSGA found either exact or close representations of 31 of the 36 members of the Pareto optimal set. The slight decrease in performance of the phenotypic elitist sNSGA can be explained with the analysis in Figure 11.

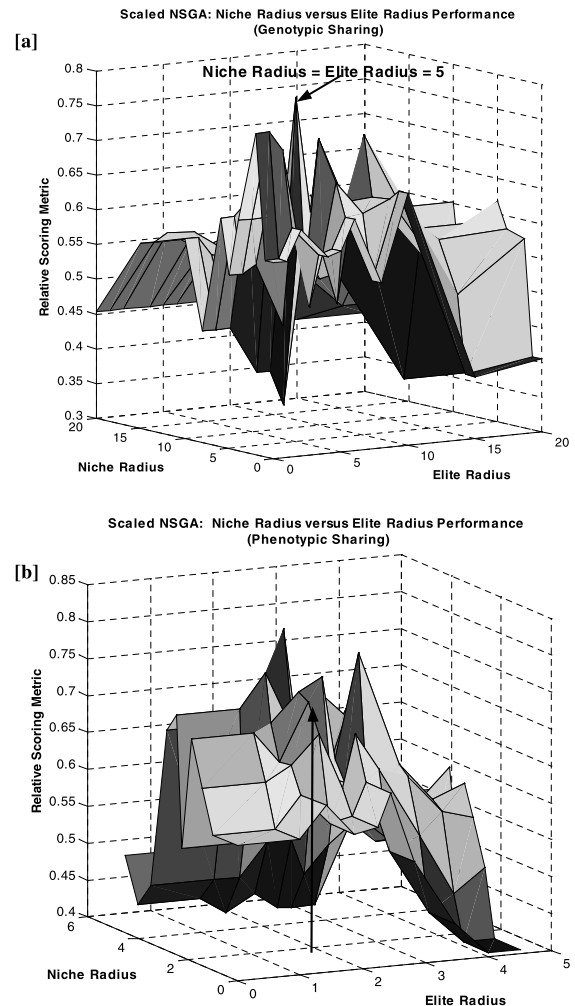


Figure 11 | Analysis of the sNSGA's performance for the full range of possible values that the parameters controlling niching and elitism can be assigned when using (a) genotypic and (b) phenotypic sharing. The arrows designate the algorithm's performance when these parameters are set equal to the recommended niche radius attained from the relationships presented by Deb & Goldberg (1989).

Figure 11 shows the performance of the sNSGA using both sharing schemes for the full range of values σ_{elite} and σ_{share} that can be assigned. The plot again shows that increasing the elitist selection pressure by decreasing σ_{elite} generally improves performance. The slight difference in performance between genotypic and phenotypic sharing was caused by the assumption used in this study that $\sigma_{elite} \approx \sigma_{share}$. Figure 11(a) shows that this assumption results in finding the highest peak in performance for the sNSGA under genotypic sharing, whereas Figure 11(b)

shows that this assumption was not able to exactly find the highest performance peak in phenotypic space. It should be noted from Figures 6 and 11, however, that assuming $\sigma_{elite} \approx \sigma_{share}$ for the monitoring application generally resulted in finding the peak or very near peak performances for the various forms of the NSGA in both sharing spaces, which is probably sufficient for most applications. Of course, the actual performance obtained with these parameter settings will vary somewhat, depending on the random number seed selected.

Figure 11(a) again shows that genotypic space was generally very noisy relative to phenotypic space, which explains why the sNSGA still required larger population sizes to be used in genotypic space in order for the two sharing schemes to attain comparable results in the absence of elitism. Additionally, the increased noise present in genotypic space after rescaling, as shown in Figure 11(a) relative to the previous case in Figure 6(a), explains why N had to be increased from 830 to 870 as discussed above. Overall, the guiding relationships and considerations discussed in the methods section of this paper were able to successfully guide the competent design of the NSGA. Figure 10 shows that the algorithm was able to find 95% of the Pareto optimal set when the influences of population sizing, elitism and genetic drift were carefully considered in its design.

CONCLUSIONS

Evolutionary multiobjective optimization (EMO) methods have been shown to be more efficient than traditional optimization methods such as mixed-integer programming at identifying tradeoffs among multiple objectives (Ritzel *et al.* 1994; Coello 1996). A major difficulty in applying these methods, however, is in identifying appropriate parameter settings to ensure that the problem's decision space is effectively navigated and the entire tradeoff curve is identified. In this paper, we demonstrate how theoretical relationships from the genetic and evolutionary computation (GEC) field can be used to overcome this difficulty. The tradeoff between cost and squared relative estimation error (SREE) for the groundwater monitoring

application presented in this work was accurately quantified by ensuring that the nondominated sorted genetic algorithm (NSGA) was competently designed to navigate the problem's decision space. Preliminary problem analysis identified a range of potential population sizes and the potential computational complexity of solving the problem. Additionally, theoretical relationships from Deb & Goldberg (1989) effectively sized the niches required to maintain stable subpopulations, each of which actively sought different sections of the Pareto frontier. Using the recommended niche radius from Deb & Goldberg (1989) also effectively identified the niched elitist set of solutions for successive generations in each of the runs performed in this study. Elitism greatly improved the efficacy of the NSGA and helped to reduce performance differences between genotypic and phenotypic sharing. Analysis of the algorithm's performance as a function of time showed the adverse effects of genetic drift and motivated further study of the role of selection pressure in the NSGA. Rescaling the shared dummy fitness assignments of the NSGA such that successive fronts had exponentially decreasing fitness values greatly improved the algorithm's performance. Combining elitism with rescaled fitness assignments resulted in identifying all but 2 members of the Pareto optimal set.

This study provides practitioners with a methodology that minimizes the number of runs required to identify the proper parameter settings for the NSGA by: (1) identifying the most appropriate population size, (2) properly sizing niches for fitness-based sharing, (3) correctly setting the elitist selection pressure, and (4) carrying out careful performance analysis to avoid genetic drift stall. Parameter settings for the NSGA attained using a total of 10 runs resulted in peak or near peak performance of the algorithm. Six of the trial runs were necessary to select the proper population size. Elitism greatly improved the NSGA's performance and narrowed any performance differences between sharing in genotypic versus phenotypic space, a clear advantage when problem-specific information required for phenotypic sharing is not available. When such information is available, though, phenotypic sharing is generally preferred because of the increased noise associated with selection in genotypic space in most EMO applications (Fonseca & Fleming 1995; Horn 1997), as was

observed in this application. Finally, the use of the niche radius to set the elitist selection pressure proved to be very effective in attaining peak or near peak performance from the NSGA, but further study is needed to see whether this finding will be appropriate for other applications.

Rescaling the fitness assignments to further increase selection pressure and reduce the influence of genetic drift requires consideration of the potential tradeoff between improving the algorithm's performance and reducing diversity within the population. Additional analysis was necessary to properly address this tradeoff and select the appropriate scaling coefficient. The final selection of a scaling coefficient required four additional trial runs after the six trial runs used to set the population and niching parameters.

The approach used to design the NSGA in this study treats the algorithm as a system with discernable properties and utilizes valuable theoretical work from the field of genetic and evolutionary computation to discern these properties. Theoretical relationships allow engineering practitioners to gain direct insights into the algorithm's performance as a function of its input parameters while also reducing the number of parameter combinations that must be considered relative to the trial-and-error methods that have traditionally been used in the field of water resources. Trial-and-error methods inherently treat the NSGA and other evolution-based methods as black boxes and serve to obfuscate important trends in performance. In this study, recognizing the influence of population sizing and genetic drift on the NSGA's performance facilitated the identification of algorithmic modifications required to best identify the full range of tradeoffs between monitoring costs and SREE. Moreover, the theoretical relationships avoided the black box mystique inherent in a trial-and-error performance analysis, which often leads to unsupported generalizations, both positive and negative, on the effectiveness of evolution-based solution techniques in engineering applications.

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. The review comments of Dr Vlado Babovic and an anonymous reviewer improved the presentation of this study and are appreciated.

NOTATION

$F(\bar{x}_\kappa)$ = Vector valued objective function.

$f_1(\bar{x}_\kappa)$ = First component of the objective function representing sampling costs.

$f_2(\bar{x}_\kappa)$ = Second component of the objective function representing SREE.

β = Index of the objective function's components.

n_{obj} = The total number of objective function's components.

\bar{x}_κ = Vector of binary decision variables representing each sampling plan.

κ = Index of the individual sampling plans being considered.

Ω = The monitoring application's decision space consisting of 2^{20} solutions.

n_{well} = The total number of wells in the monitoring network shown in Figure 1.

$C_s(i)$ = The cost coefficient for the i th well ranging in value from \$365 to \$1095.

i = Index of the groundwater monitoring wells.

n_{est} = The total number of estimates in the interpolation domain.

$c_{all}^*(\bar{u}_j)$ = Estimate of dissolved PCE concentration at location \bar{u}_j in the interpolation domain based on all available sampling locations.

$c_{est}^\kappa(\bar{u}_j)$ = Estimate of dissolved PCE concentration at location u_j in the interpolation domain based on the sampling locations in the κ th sampling plan.

j = Index of the unknown locations in the interpolation domain where dissolved PCE concentrations are estimated.

$nsamp(\kappa)$ = The total number of samples in the κ th sampling plan.

$w(\bar{u}_j, \bar{u}_\psi)$ = The weight of the known PCE sample taken at location \bar{u}_ψ in the estimate at the unknown location \bar{u}_j in the interpolation domain.

\bar{u}_j = Three-dimensional coordinate vector for the j th unknown location in the interpolation domain.

\bar{u}_ψ = Three-dimensional coordinate vector for the ψ th sampling point.

$c(\bar{u}_\psi)$ = Sampled or known dissolved PCE concentration at location u_ψ .

ψ = Index of the available sampling locations.

w_t = Sum of the weights $w(u_j, u_\psi)$ used to calculate the j th dissolved PCE estimate.

$D_\kappa(\bar{u}_j, \bar{u}_\psi)$ = The Euclidean distance between the sampled location \bar{u}_ψ and the unknown location \bar{u}_j .

$\alpha_1, \alpha_2, \alpha_3$ = Fitting parameters used to minimize the cross validation residuals for the κ th sampling plan.

$Sh[d(\kappa, \kappa')]$ = The sharing function value for two sampling plans κ and κ' .

$d(\kappa, \kappa')$ = The distance in parameter space (either genotypic or phenotypic) between two sampling plans κ and κ' .

σ_{share} = The niche radius defining the minimum distance two sampling plans can be from one another and still be considered independent.

N = The population size used by the NSGA.

q = The total number of niches.

g = The total number of generations the NSGA will be run.

γ = The reliability of maintaining q niches for g generations when the population size equals N .

p_c = The probability of crossover.

p_d = The probability that crossover will disrupt building blocks.

$O(m)$ = The order of the building blocks.

l = The string length or total number of binary digits representing each sampling plan.

b = Index of the binary digits in the chromosome representing each sampling.

z^* = Cumulative normal distribution chart value corresponding to a probability equal to $1/q$.

$npar$ = The total number of parameters or decision variables for an application.

σ_{elite} = The elite radius defining the minimum distance required between the members of the population undergoing elitist reproduction.

t = Index of the number generations.

n_f = The total number fronts that members in a given population at generation t are classified into, based on domination.

ΔD_f = The constant decrement used by the NSGA to reduce the fitness values of successive fronts.

S_e = The scaling coefficient used by the NSGA to exponentially decrease the fitness values of successive fronts.

$SREE_{true}(\bar{x}_\eta)$ = The squared relative estimation error for the η th member of Pareto optimal set.

$SREE(\phi(t))$ = The squared relative estimation error for the ϕ th member of the current population at generation t .

$SREE_{max}$ = The maximum squared relative estimation error attained within Ω .

$\phi(t)$ = Index of the members of the current population at generation t .

η = Index of the members of the Pareto optimal set.

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.
- Barry, D. A. & Sposito, G. 1990 Three-dimensional statistical moment analysis of the Stanford/Waterloo Borden tracer data. *Wat. Resources Res.* **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. *Wat. Resources Res.* **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 Inf. Syst.* 1(3), 269–308.
- Deb, K. & Goldberg, D. E. 1989 An investigation of niche and species formation in genetic function optimization. In *Proceedings of the Third International Conference on Genetic Algorithms* (ed. J. D. Schaffer), 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. & Fleming, P. J. 1995 An overview of evolutionary algorithms in multiobjective optimization. *Evol. Comput.* 3(1), 1–16.
- Goldberg, D. E. 1989 *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley: New York.
- Goldberg, D. E., Deb, K. & Korb, B. 1989 Messy genetic algorithms: motivation, analysis, and first results. *Complex Syst.* 3(4), 493–530.
- Goldberg, D. E., Deb, K. & Horn, J. 1992 Massive multimodality, deception, and genetic algorithms. In *Parallel Problem Solving From Nature, 2* (ed. R. Männer & B. Manderick), pp. 37–46. North-Holland, Amsterdam.
- Goldberg, D. E. & Richardson, J. 1987 Genetic algorithms with sharing for multimodal function optimization. In *Genetic Algorithms and Their Applications: Proceedings of the Second International Conference on Genetic Algorithms* (ed. J. J. Grefenstette), pp. 41–49. Morgan Kaufmann: San Mateo, CA.
- Gupta, H. V., Sorooshian, S. & Yapo, P. O. 1998 Toward improved calibration of hydrologic models: multiple and noncommensurable measures of information. *Wat. Resources Res.* 34(4), 751–763.
- Halhal, D., Walters, G. A., Ouazar, D. & Savic, D. A. 1997 Water network rehabilitation with structured messy genetic algorithm. *J. Wat. Resources Plan. Mngmnt* 123(3), 137–146.
- Holland, J. H. 1975 *Adaptation in Natural and Artificial Systems* University of Michigan: Ann Arbor, MI.
- Horn, J. 1997 The nature of niching: genetic algorithms and the evolution of optimal, cooperative populations. *Doctoral dissertation* University of Illinois: Urbana, IL.
- Ishibuchi, H. & Murata, T. 1996 Multi-objective genetic local search algorithm. In *Proceedings of 1996 IEEE International Conference on Evolutionary Computation* IEEE: Piscataway, NJ, pp. 119–124.
- Mahfoud, S. 1995 Niching methods for genetic algorithms. *Doctoral dissertation* University of Illinois: Urbana, IL.
- Maxwell, R. M., Carle, S. F. & Tompson, F. B. 2000 *Contamination, Risk, and Heterogeneity: On the Effectiveness of Aquifer Remediation* Lawrence Livermore National Laboratory Report, UCRL-JC-139664, Livermore, CA.
- Parks, G. T. & Miller, I. 1998 Selective breeding in a multiobjective genetic algorithm. In *Fifth International Conference on Parallel Problem Solving From Nature* (ed. A. E. Eiben, T. Bäck, M. Schoenauer & H. Schwefel), pp. 250–259. Springer-Verlag: Berlin.
- 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. *Wat. Resources Res.* 36(12), 3731–3741.
- Reed, P., Minsker, B. & Goldberg, D. E. 2000b Designing a competent simple genetic algorithm for search and optimization. *Wat. Resources Res.* 36(12), 3757–3761.
- Ritzel, B. J., Eheart, J. E. & Ranjithan, S. 1994 Using genetic algorithms to solve a multiple objective groundwater pollution containment problem. *Wat. Resources Res.* 30(5), 1589–1603.
- Rizzo, D. M., Dougherty, D. E. & 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. *Evol. Comput.* 2(3), 221–248.
- Thierens, D. & Goldberg, D. E. 1993 Mixing in genetic algorithms. In *Proceedings of the Fifth International Conference on Genetic Algorithms* (ed. S. Forrest), pp. 38–45. Morgan Kaufmann Publishers: San Mateo, CA.
- Thierens, D. & Goldberg, D. E. 1994 Convergence models of genetic algorithm selection schemes. In *Proceedings of the Third Conference on Parallel Problem Solving From Nature* (ed. Y. Davidor, H.-P. Schwefel & R. Manner), pp. 119–129. Springer-Verlag: New York.
- 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* IEEE Press: New York, pp. 535–540.
- 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, OH.
- Van Veldhuizen, D. A. & Lamont, G. B. 2000 Multiobjective evolutionary algorithms: analyzing the state-of-the-art. *Evol. Comput.* 8(2), 125–147.
- Zitzler, E., Deb, K. & Thiele, L. 2000 Comparison of multiobjective evolutionary algorithms: empirical results. *Evol. Comput.* 8(2), 173–195.
- Zitzler, E. & Thiele, L. 1999 Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. *IEEE Trans. Evol. Comput.* 3(4), 257–271.