

# Combining Reliability and Pareto Optimality - An Approach Using Stochastic Multi-Objective Genetic Algorithms

Abhishek Singh<sup>1</sup>, Barbara Minsker<sup>2</sup>, and David E. Goldberg<sup>3</sup>

<sup>1</sup>Graduate Research Asst., Dept of Civil & Environmental Engineering, University of Illinois, Urbana-Champaign, 205. N. Mathews Ave, Urbana, IL 61801; [asingh8@uiuc.edu](mailto:asingh8@uiuc.edu)

<sup>2</sup>Associate Professor, Dept of Civil & Environmental Engineering, University of Illinois, Urbana-Champaign, 205. N. Mathews Ave, Urbana, IL 61801; [minsker@uiuc.edu](mailto:minsker@uiuc.edu)

<sup>3</sup>Professor, Dept. of General Engineering, University of Illinois, Urbana-Champaign, 104 S. Mathews Ave., Urbana, IL 61801: [deg@uiuc.edu](mailto:deg@uiuc.edu)

## ***Abstract***

Genetic Algorithms have been successfully applied to numerous water resources problems, including problems with multiple objectives or uncertainty (noise). GAs tackle multi-objective optimization by following three basic principles – advancing the non-dominated frontier; maintaining diversity in the population (through various techniques like sharing, niching, and crowding); and using an elitist. However finding Pareto-optimal solutions becomes complicated when we add uncertainty to the problem. It was found that the solutions obtained using existing multi-objective solvers, although Pareto optimal were not the most robust or reliable solutions. In single-objective problems noise has typically been dealt with using Monte-Carlo-type sampling and some form of aggregate statistics (e.g., the average of the sample fitness). With multiple objectives the noise can interfere in determining non-domination of individuals, diversity preservation, and elitism (the three basic steps in multi-objective optimization). This paper proposes and tests several approaches to tackling some of these problems. These approaches strike a balance between finding the most optimal and the most reliable solution to the problem, thus giving decision makers and designers a practical and robust optimization tool.

## ***Introduction***

Evolutionary and genetic optimization has, over the past few years, ‘evolved’ from being a rare curiosity to a firmly entrenched practice in the circles of engineering and management optimization. These techniques have been applied to many water resources applications, including groundwater remediation design, optimal reservoir system operation, calibrating rainfall-runoff models, remediation policy selection, and solving multiple objective groundwater pollution contaminant problems (e.g., *Ritzel et al*, 1994; *Wang and Zheng*, 1997; *Wardlaw and Sharif*, 1999; *Reed et al.*, 2001). One of the reasons that genetic algorithms have been chosen is that they have been shown to easily handle non-convex, discrete, discontinuous, noisy, and multi-objective problems (*Goldberg*, 1989) that arise frequently in these applications. In addition, evolutionary algorithms are arguably domain independent, which make them excellent candidates for the simulation/optimization methodology commonly used in this field.

Two areas of evolutionary optimization where progress has been in terms of theory and application are multi-objective optimization and noisy or uncertainty based optimization. Evolutionary multi-objective optimization (EMO) methods, which seek to find the most Pareto

optimal set of solutions to a problem, have garnered increased attention. Pioneering work in this area was undertaken by *Fonseca and Fleming* (1993), *Coello* (1999), and *Deb* (2000) among others. *Cieniawski* (1993) and *Ritzel et al* (1994) were among the earlier applications of EMO methods in water resource management. More recently *Reed et al* (2001) developed guidelines for competent Multi-Objective Genetic Algorithm (MOGA) and applied them to a groundwater monitoring problem.

In addition to having multiple objectives, most real world problems have inherent noise. This noise can be due to model approximations, knowledge uncertainty, or inconsistent or sparse data. Noisy GAs seek to find the most reliable and robust optimum in the face of such noisy objective functions. *Miller* (1997) was among the first to analyze and suggest design methodologies for single-objective noisy genetic algorithms. Among the various studies undertaken since, *Gopalakrishnan et al* (2001) demonstrated the applicability of these ideas in water resources management on a simple groundwater remediation problem with uncertain aquifer properties. This study seeks to combine the theoretical and practical work in these two fields to develop multi-objective optimizers capable of handling noise and uncertainty. Some initial work in this area was undertaken by *Hughes* (2001) who used the concept of stochastic dominance, assuming that the noise was from a Gaussian distribution. Although this work provides a good theoretical launching pad for our current study, their simplifying assumption for the noise distribution is either not true or not verifiable in cases such as ours, leading to the present approach for a more general noisy-MOGA.

### ***Multi-Objective Genetic Algorithms***

Most problems in nature have several (possibly conflicting) objectives to be satisfied. Traditionally such problems were handled by converting the multiple objectives to one, by using weighting functions, or using one objective to optimize and the others as constraints. Such an approach has many problems, including the loss of significant tradeoff information and the inability to search the true objective space for a global optimum. The other approach is to consider a set of the locally best alternatives, or the solutions that represent the optimal tradeoffs for the solution. These comprise the non-dominated frontier for the problem and the objective now becomes to find the set of solutions that are globally non-dominated. In other words, the objective is to find solutions to the problem such that there are no feasible solutions that would better one criterion without simultaneously worsening at least one other criterion. This set of non-dominant solutions is often referred to as the Pareto Optimum set, after the famous mathematician Vilfredo Pareto, who generalized the concept of this kind of optimality.

Thus, unlike single objective optimization, multi-objective optimization (MO) has to work to find a set of good solutions. Since genetic algorithms too manipulate populations of candidate solutions, they are a natural approach for this kind of optimization. However loss of diversity (convergence to one solution due to stochastic noise) is one problem with traditional genetic algorithms that make their direct use for MO difficult. That is why research in this area has focused on finding better (more non-dominant) solutions to the problem while still maintaining diversity and preserving the existing front.

Multi-objective optimizers typically use different schemes to give a measure of ‘non-dominance’ to the solutions found. This is done by taking a local measure of non-dominance for each solution. Each individual is compared by all others in the population, and those that are strictly non-dominated are marked as the locally Pareto (or rank 1) solutions. Some methods, like those proposed by *Fonesca and Fleming* (1993), rank each individual as one more than the total number of individuals it dominates. Others, like the non-dominated sorted genetic algorithm (NSGA) (*Srinivas and Deb*, 1994), do layer-wise Pareto ranking. Here rank 1 solutions are those that are not dominated by any other solutions. These are then removed from the list of considered solutions, and rank 2 solutions are defined as those that are strictly non-dominant with respect to the remaining solutions. This process of ranking and removal is followed till all solutions are ranked. In general the objective of the optimizer now becomes to minimize the rank. To maintain good spread over the Pareto front and prevent loss of diversity various methods can be used. These include fitness sharing, mating restrictions, and crowding schemes among others. In addition to preservation of diversity, most multi-objective optimizers have some kind of an elitist scheme, so that the existing Pareto optimal solutions are not lost. As a starting point for our analysis and implementation we used the elitist non-dominant sorting algorithms based on crowding distance (NSGA-II) (*Deb et al*, 2000) which is an efficient algorithm that gives good results over a large range of problem types. This algorithm uses the layered ranking approach and the crowding factor for optimization and diversity preservation. Elitism was implemented by giving first preference in selection to all rank 1 individuals, ensuring that they succeed to the next generation.

### ***Noisy Genetic Algorithms***

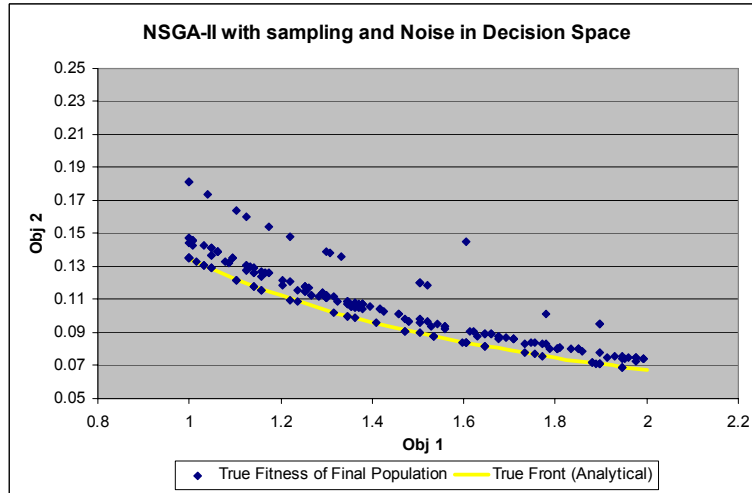
A GA that operates in a noisy environment is referred to as a “Noisy GA”. Much of the theory and practice in this area can be understood if we realize that even in non-noisy environments genetic algorithms are stochastic algorithms, which operate under a certain amount of natural ‘noise’. This is due to the fact that GAs work at the level of building blocks, or sub-solutions to the problem (*Goldberg*, 1989). Combinatorially, GAs are more efficient than random searches, due to the effects of the operations of selection, crossover and mutation, GAs search the decision space to find the optimum building blocks that can then be combined to give the final solution. Since there is no way of explicitly identifying such building blocks, their search and evolution is intrinsically noisy and is handled through adequate population sizing.

In the case of noisy fitness functions, the additional noise adds up to the intrinsic noise and can be dealt with using larger population sizes or adequate sampling to reduce the external noise. In most of these cases sampling is used to find the average value for a given objective, which through the central limit theorem would increasingly tend to the actual mean-average of the distribution for the noise. Since there is an obvious tradeoff between computation time and the increased accuracy through sampling, an optimum sampling size can be computed (*Miller*, 1997). This methodology was tried and tested for a simple one dimensional groundwater remediation problem by *Gopalakrishnan et al* (2001), and gave promising results. Most significantly, highly reliable solutions were found with low amounts of sampling (as few as 5 samples per solution). This was critical in cases such as ours since it was impossible to afford large sample sizes given the expense of fitness evaluations. Another study was undertaken by *Chan Hilton and Culver* (2000) to find robust solutions to groundwater remediation problems

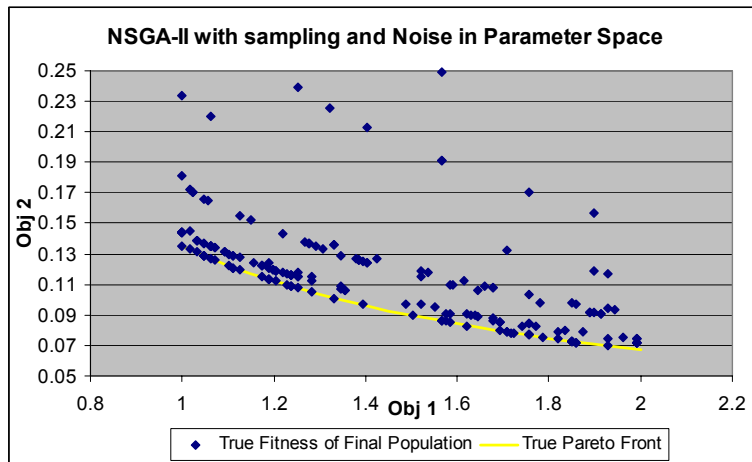
with uncertain aquifer properties, by using the past performance of a particular solution to assess its robustness over uncertainty. It is noteworthy to add that most of the studies undertaken on the theoretical aspects of noisy GAs have considered only noise in objective space. Typically for experimentation a Gaussian signal was added to the objectives. Since for most real world problems noise is seen mostly in the decision variables or the parameters of the objective function, this study also looks at the effect of different kinds of noise on the performance of the multi-objective GA.

### ***First Steps towards a Noisy Multi-Objective GA***

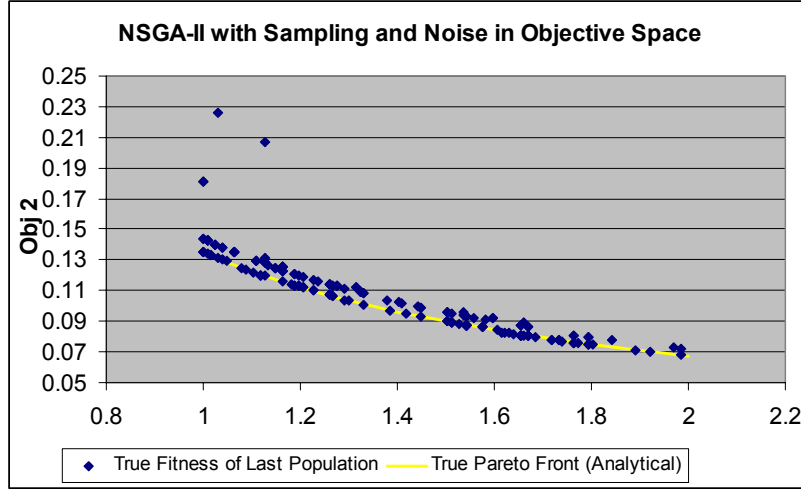
Our first instinct in trying to identify an effective noisy multi-objective GA was the most obvious one. With adequate sampling for each individual and a large enough population size, the noise could be handled in the same way as a single objective case. We thus used the NSGA-II (*Deb et al, 2000*) and introduced sampling at each fitness evaluation averaging fitness for all individuals and taken this average value as the fitness for the individuals.



**Figure 1:** Original NSGA-II with decision variable noise.



**Figure 2:** Original NSGA-II with parameter noise.



**Figure 3:** Original NSGA-II with objective space.

The algorithm was tested on a test bed of multi-objective problems taken from *Deb* (1998) to allow extensive testing on different types of problems before a water resources application was undertaken. For this report we have shown results on a continuous, multi-modal problem which has more than one local Pareto optima. We considered noise in objective, decision, and parameter space. Previous work on population sizing (*Mahfoud*, 1995) was used to come up with conservative population sizes (200) and the GA was run for 60 generations (to ensure convergence, this was almost 4 times the string length used, while normally convergence is expected in two times the string length generations). An artificial noise signal with a standard deviation of 10% of the range was added to the objective, decision, and parameter values. Results with sample size 5 for these three cases are shown in figures 1, 2, and 3. The results show the final population evaluated without noise to indicate how close these were to the original solution. For this objective function, the true Pareto front (obtained analytically) was known *a priori* and is also shown in all Figures. As can be seen, many non-optimal solutions were found in these cases. Moreover, the overall coverage of the front is also not very good.

### ***Problems with the Simple Approach***

On closer analysis of the results and the algorithm, it was clear that noise leads to problems in the major components of the NSGA-II. The first difficulty is in the ranking scheme itself. The ranking scheme compares the absolute fitness values of a particular individual with all the other individuals to decide on a rank. Since the fitness values are themselves uncertain this leads to large fluctuations in the rank of a given individual. Often even for very small variations in the fitness value, the rank can have large variations. In effect the ranking mechanism can potentially *amplify the noise* leading to poor performance. This fact was noticed before by *Hughes* (2001) who demonstrated that different ranking schemes (like the NSGA or the MOGA) have different mean standard deviations of rank values for the same applied noise.

The other potential problem is with the elitist selection mechanism, where all individuals of rank 1, are always selected for the next generation. This sort of string elitism helps non-noisy MOs to maintain the existing good solutions. With noise, however, this approach leads to the retention of outliers, or individuals that at some time got very high fitness due to a favorable random event.

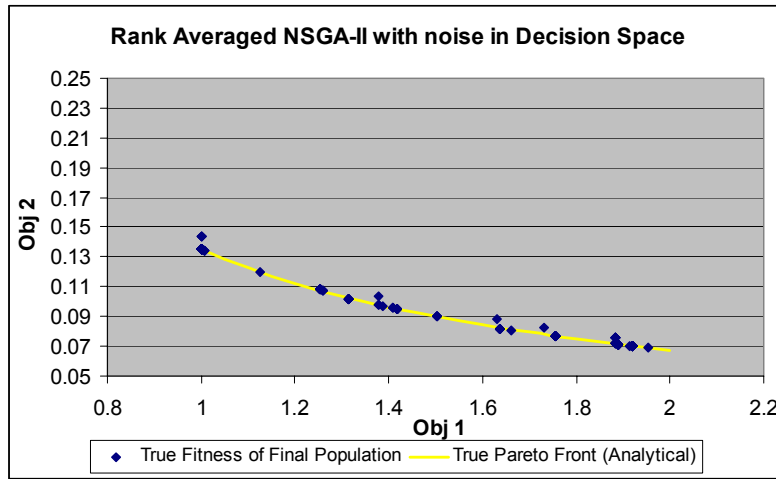
In the existing algorithm there is no way of reevaluating these individuals, so they continue to survive and are part of the final front. Apart from the elitism, another mechanism that is adversely affected by noise is the diversity preservation. The selection procedure of NSGA-II can be viewed as a two step process. The algorithm uses a crowding factor to calculate the distance of each individual from its neighbor on the same Pareto front. This crowding distance is indicative of the relative uniqueness of that particular individual. Small crowding distance means crowding on the Pareto front, while large distances mean the individuals are spread apart. The selection scheme that NSGA-II uses is the modified ' $\mu+\lambda$ ' scheme. This selection mechanism, which is very popular in evolutionary strategies, combines the parent population ( $\mu$ ) with the children population ( $\lambda$ ) and then selects the best  $\mu$  of the combined population. NSGA-II uses two criteria for choosing the best  $\mu$  from the combined population. It first ranks the combined population and then selects the best ranks one by one until the addition of the next ranked solutions would lead to more than  $\mu$  individuals in the selected population. This last rank is then sorted on the basis of the crowding distance factor of the individuals and the individuals with the maximum distance are chosen to fill in the remainder of the population. The first criterion of rank causes the population to converge to the optimum front. After this the second criterion of distance preferentially selects the more unique individuals and causes the solutions to spread over the Pareto front.

With the addition of noise, the problem with this mechanism is two fold. In this case we use the average fitness for distance calculation. First, since the fitness values are all uncertain, the crowding distance of an individual from its neighbor is an unreliable measure of how unique the solution is in the true search space. In fact, multiple copies of the same individual can and often do have different average fitness values. Moreover, with the inclusion of noise, the population almost never converges to one stable Pareto front. Hence, the second selection criterion of distance never becomes very effective and solutions are not spread over the front.

### ***Modifications to Ranking Scheme of Original NSGA-II***

As can be seen the problem of adapting a multi-objective optimizer to a noisy environment is not straightforward. Alterations in the original algorithm may apparently deal with some of the problems discussed above, but they may also hamper the evolutionary process that was the very strength of the original algorithm. A number of approaches for modifying the ranking scheme to address this problem are possible. We first tested a simple modification to the ranking scheme. Future work will test more sophisticated approaches such as that of *Hughes* (2001), who used a stochastic dominance approach, in which the 'probability of dominance' replaces absolute dominance in the ranking scheme. As discussed before the approach assumes a Gaussian distribution for the noise signal and then uses an analytical approach to find the probability of dominance. We intend to combine this with frequent sampling to make it more generally applicable and feasible on various noisy problems. Multi-objective optimization can be viewed as a single-dimensional (though multi-modal) optimization problem with respect to the rank values. Noise in single dimensions is typically handled through sampling and averaging. A comparable approach would be to rank each individual for each fitness sample and then use the average rank as the fitness criterion. When identical individuals within a population were identified the rank was averaged across all members of the population and their samples. All the identical individuals were then given the same common average rank. Under this approach,

fractional average ranks are possible, which means the discrete fronts that are necessary for the selection scheme of the NSGA-II are lost. To address this problem, the layer-wise selection scheme was replaced by selection through tournament selection within the combined population. Two individuals were chosen at random and the one with higher average rank was chosen. If the two had the same ranks then the one with larger crowding distance (still calculated over average fitness) was chosen. This approach was seen to give good results for the case with noise in the decision space (Figure 4). However the performance in the other two cases was not any better than the earlier cases. Moreover, even for the case with noise in decision space, there were many duplicate solutions, and the Pareto front found eventually was sparse. This could be attributed to the problem in the crowding mechanism, which was seen to be unreliable with noise.

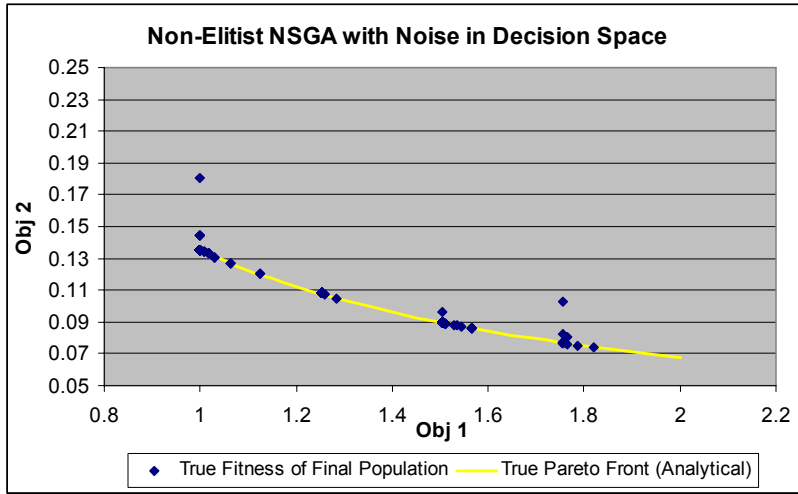


**Figure 4:** NSGA-II with rank based sampling finds optimal (though sparsely spread) solutions for noise in decision space.

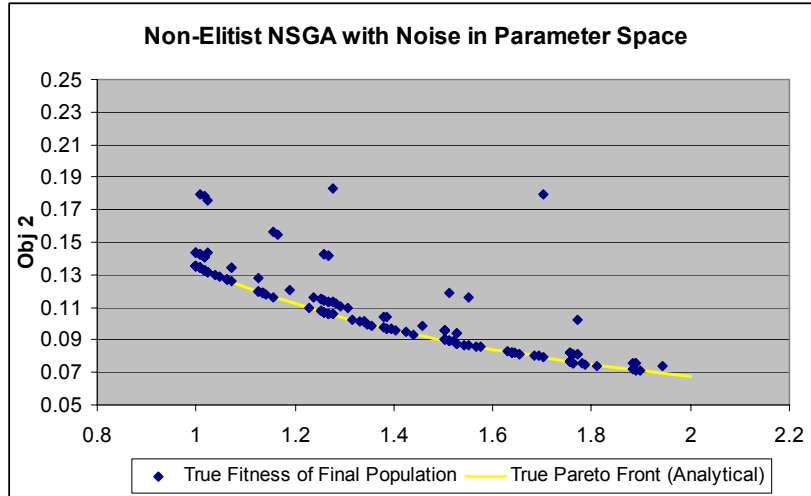
### ***Modification to Selection Scheme***

As discussed earlier, the elitism in the selection scheme causes solutions favored by random events to survive into the final Pareto frontier. On the other hand, elitism can also help stabilize the Pareto front, avoiding the need for the algorithm to have to search again for good solutions. Our first step for solving this problem was to replace the elitist rank based selection scheme, by a simple tournament selection scheme. In this case, once all the individuals have been ranked, two individuals are chosen randomly and the one with lower rank is selected. If two individuals have the same ranks then the one with higher crowding distance is selected. This mechanism can be easily extended for larger tournament sizes, which leads to stronger selection pressure. This approach can be applied to both the average fitness scheme and the average ranking scheme. However initial experiments on the average ranking scheme showed the same problems as before thus all subsequent experiments were performed using averaging on the fitness values only.

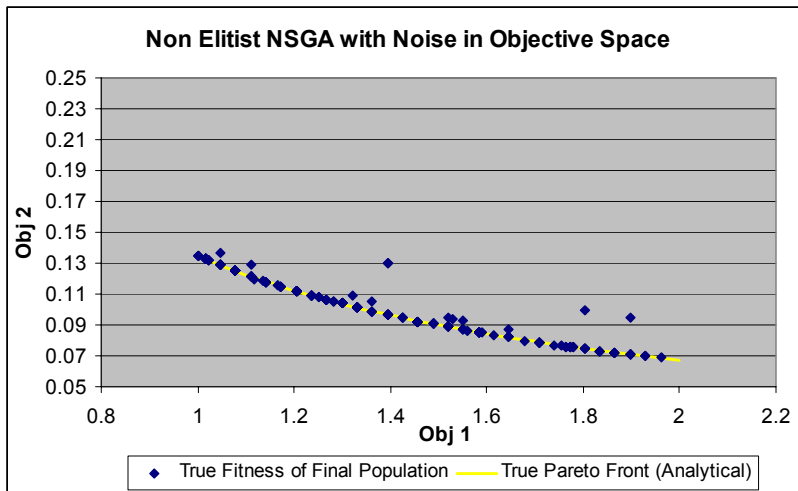
With tournament size 2, and averaging over the fitness values, the results were as shown in figures 5, 6, and 7. Comparing these results to the corresponding results in Figures 1 through 3, it is clear that removing elitism does lead to fewer dominated solutions in the final population. However, the problem of insufficient coverage of the front is still seen, especially in the case with noise in decision space. Since such a dramatic effect is observed by changing the ranking scheme we will study the effect of other selection schemes on the performance of the GA.



**Figure 5:** Improved Performance of Non-Elitist NSGA-II with decision variable noise.



**Figure 6:** Improved Performance of Non-Elitist NSGA-II with parameter noise.



**Figure 7:** Improved Performance of Non-Elitist NSGA-II with objective space noise.

### ***Modifications to the Diversity Preservation Scheme***

Of the various diversity preservation schemes in the GA literature, the attractive feature about the crowding distance used in the NSGA-II is that it does not need any fine tuning or setting of parameters. Other diversity preservation schemes such as niching and sharing are often very sensitive to the parameters, which have to be adjusted through testing and problem knowledge. However as seen, crowding distance is not a very robust scheme in noisy environments. Part of the problem is that due to random sampling, identical individuals can have different average fitnesses, and thus different crowding distances. One way to remedy this is to identify identical individuals within the population and take the *average fitness values* across all of these individuals. All identical individuals are then given the same average fitness value, which is then subsequently used for distance calculation. This is the same as the normalization of rank values for identical individuals, except that now this involves averaging the fitness values instead of the ranks. Unfortunately the results of this approach are not as positive as expected. Normalizing the fitness values leads to improvements only for noise in objective space. For noise in decision and parameter space the performance is seen to actually worsen. The cause of this result needs further study, and we intend to apply different diversity preservation mechanisms like niching and sharing to see if these can improve performance.

### ***Conclusion and Future Work***

This paper identifies the difficulties in designing a robust and efficient noisy multi-objective genetic algorithm using existing techniques. Unlike single objective optimizers, multi-objective GAs need to maintain a population of optimal solutions, which becomes more and more difficult in the presence of noise and uncertainty. It was seen that elitism as used by most multi-objective GAs can be detrimental to performance in noisy environments. Moreover, the ranking scheme used to decide the non-dominant characteristics of the individuals can itself amplify noise and lead to instability. These two problems were handled by trying out two different approaches; that of rank based sampling and non-elitist selection schemes. The performance in both cases was shown to improve, however the GA still suffered from the inability to preserve diversity.

Future work will involve implementing some of the alternatives that have been mentioned in this study, like stochastic dominance, and testing different diversity preservation schemes. The algorithm will also be tested on sample groundwater remediation problems, including a field study of Umatilla Chemical Depot in Oregon, which will serve as a real life case for our GA.

### ***Acknowledgements***

This material is based upon work supported by the U. S. Army Research Office under Grant No. DAAD19-001-1-0025.

### ***References***

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.

- Coello, C. A. C. 1999 An Empirical Study of Evolutionary Techniques for Multiobjective Optimization in Engineering Design. *Doctoral dissertation* Tulane University: New Orleans, L.A.
- Deb, K. (1998). Multi-objective genetic algorithms: Problem difficulties and construction of test problems. Technical Report No: CI-49/98. Dortmund: Department of Computer Science/LS11, University of Dortmund, Germany.
- Deb, K., Agrawal, S., Pratap, A., & Meyarivan, T. (2000) A Fast Elitist Non-dominated Sorting Genetic Algorithm for Multi-objective Optimisation: NSGA-II. *Proceedings of the Parallel Problem Solving from Nature VI Conference*, 849-858.
- Fonseca, C. M. & Fleming, P. J. (1993) Genetic Algorithms for Multiobjective Optimization: Formulation, Discussion and Generalization, In Stephanie Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, 416-423.
- Goldberg, D. E. (1989) *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, New York.
- Gopalakrishnan, G., Minsker, B. S., & Goldberg, D. E. (2001) Optimal Sampling in a Noisy Genetic Algorithm for Risk-Based Remediation Design, "*Bridging the Gap: Meeting the World's Water and Environmental Resources Challenges*", Proceedings of the World Water and Environmental Resources Congress, Orlando, FL.
- Chan Hilton, A. B. and Culver, T. B. (2000) Optimizing groundwater Remediation Design under Uncertainty. *Proceedings of the Joint Conference on Water Resource Engineering and Water Resources Planning and Management*, July 30-August 2, 2000, Minneapolis, MN.
- Hughes, E.J. (2001) Evolutionary Multi-Objective Ranking with Uncertainty and Noise. *First International Conference on Evolutionary Multi-Criterion Optimisation EMO 2001*, Springer LNCS 1993, 329-343, Zurich, Switzerland.
- Mahfoud, S. (1995) Niching Methods for Genetic Algorithms. *Doctoral Dissertation*, University of Illinois: Urbana-Champaign, IL.
- Miller, B.L. (1997) *Noise, Sampling and Efficient Genetic Algorithms*, IlliGAL Rep. 97001.
- Reed, P., B. S. Minsker, and D. Goldberg. Designing a New Elitist Nondominated Sorted Genetic Algorithm for a Multiobjective Long Term Groundwater Monitoring Application, in *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO'2001*.
- Ritzel, B.J., Eheart, J.W. & Ranjithan, S. (1994) Using Genetic Algorithms to Solve a Multiple Objective Groundwater Remediation Problem, *Water Resource Research*, 30(5), 1589-1603.
- Srinivas, N. & Deb, K (1994) Multiobjective Optimization Using Nondominated Sorting in Genetic Algorithms. *Evolutionary Computation*, 2(3), 221-248.
- Wang, M. and Zheng, C. (1997) Optimal Remediation Quality Management under Parameter Uncertainty, *Water Resources Research*, 23(7), 1162-1174.
- Wardlaw R. & Sharif M.(1999), Evaluation of genetic algorithms for optimal reservoir system operation", *Journal of Water Resources Planning & Management*, American Society of Civil Engineers, 125(1), 25-33.