Appendix Contents

Appendix A:

Appendix B: Information Supporting Discussion of the Development of the Recommended Water Quality Management Strategy

Appendix C: Additional Tables

Appendix A: Packaging and Assessment Tool

Overview and Description

With more than $1 billion in funding from the federal government pouring in to the watershed to improve the Lake Tahoe environment, a premium had been placed on developing a suite of state-of-the-science analysis to quantify loads and pollutant sources, assess management activities, and identify and develop innovative management solutions. To facilitate this process, the Lahontan Water Quality Control Board and the Nevada Department of Environmental Protection identified and assembled regional experts into four Source Category Groups (SCGs) to investigate potential Pollutant Control Opportunities for each major source of pollutants entering Lake Tahoe. They were (1) Urban Uplands and Groundwater, (2) Forest Uplands, (3) Atmospheric Deposition, and (4) Stream Channel Erosion SCGs.

Each SCG identified three or four physical settings that represented typical features governing the selection and implementation of specific PCOs, and influencing pollutant fate and transport to the lake. An example of an Urban Upland setting is a concentrated-impervious-steep subbasin, whereas an Atmospheric Deposition setting is a concentric perimeter of uniform distance to the Lake. Each SCG also defined two or three tiers representing different combinations of PCOs, and estimated annualized load reductions from applying each tier to each setting. Table A1 below describes the settings and tiers for each SCG, as established for the IWQMS project.

Table A1. Description of IWQMS Settings and Tiers by Source Category Group

Atmospheric Deposition	
Settings	Four spatially based settings, measured by concentric rings of distance from the lake
Tiers	Four tiers per setting were applied, based on two different treatment levels from two different groups of pollutant sources. The first group was vehicle emisions, and the second group included transportation infrastructure or structural controls
Forest Uplands	
Settings	Three source based settings, including (A) unpaved roads, (B) highly erodible forest and recreational areas, (C) burned, plus harvested, plus relatively undisturbed forest areas
Tiers	Three tiers per setting with increasing degree of treatment: low, medium, and high
Stream Channel	
Settings	Three treatable segments along the top three most sediment-productive streams in the Basin: (1) Blackwood Ck, (2) Upper Truckee, and (3) Ward Creek
Tiers	Three levels of treatment with varying intensities and stabilization activities
Urban Upland	
Settings	Four settings based on the different combinations of slope (moderate or steep) and impervious configuration (concentrated or dispersed).
Tiers	Two tiers of differing intensity and sophistication of treatment activities, plus a third "Pump and Treat" stormwater tier for concentrated impervious areas only

For each tier of controls, associated capital, maintenance, and life-cycle costs were researched, compiled, and normalized for cross-comparison. Cost constraints were not imposed on any tier during the selection of PCOs at this stage in the process. A database of cost and Level of Adoption (LOA) was derived by scaling between the existing baseline loading budget (0 percent LOA) and the maximum load reduction estimated from applying each tier to all areas within applicable settings (100 percent LOA). Different LOA combinations result in different marginal costs for pollutant load reductions.

Faced with such a large search domain of potential alternatives with multiple control objectives and potentially non-linear cost-benefit relationships, a meta-heuristic optimization technique was applied to evaluate the costs-benefits and selection trade-offs among basin-wide pollutant sources. This technique was applied in a Microsoft Excel environment, and was called the Packaging and Assessment Tool (PAT). Output from the various SCG models, methods, and techniques were compiled and applied within PAT to assess and prioritize a wide range of management options, and evaluate potential economic benefits from water quality trading. This methodology was the platform upon which the recommended strategy was derived. The following sections will highlight (1) problem formulation using PAT, (2) describe the underlying computational algorithm, and (3) conclude with a discussion of uncertainty associated with PAT predicted results.

Problem Formulation using PAT

Problem formulation using PAT follows a three step process, as illustrated in Figure A1 below.

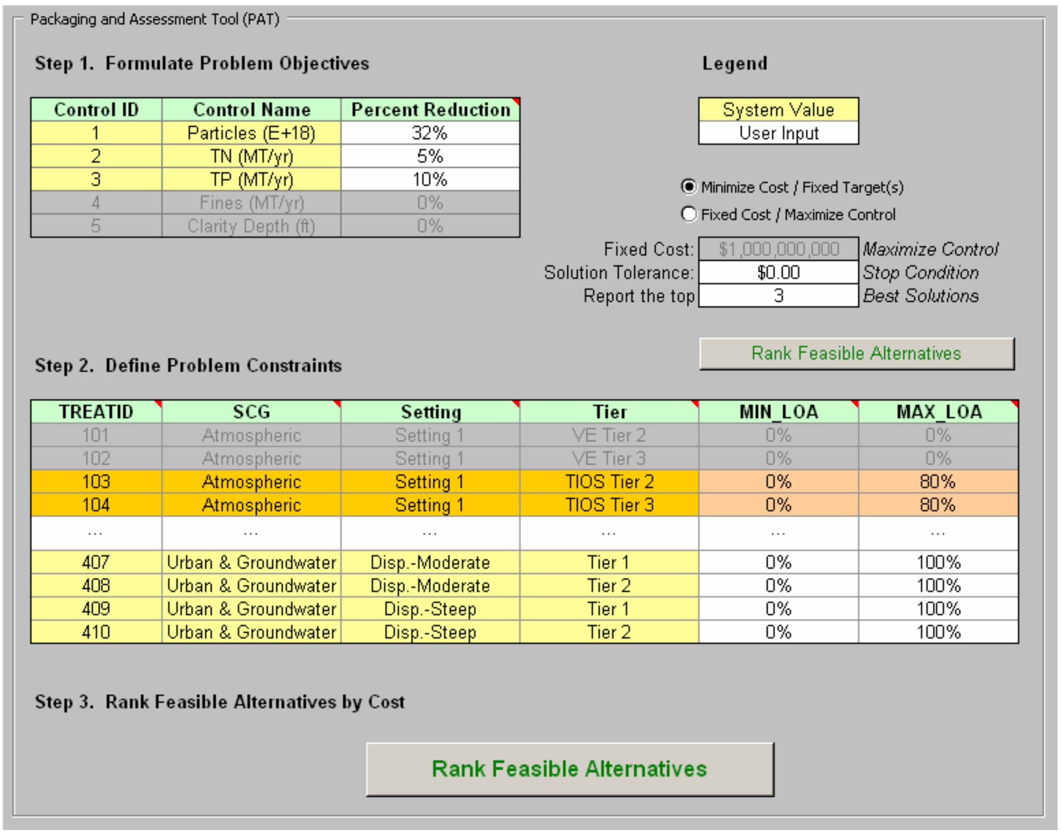

Figure A1. PAT user interface showing the three steps associated with problem formulation

Step 1 – Formulation Problem Objectives

The first step is to formulate the problem objectives. An optimization problem can be formulated to either (1) minimize costs to achieve a set of fixed pollutant targets, or (2) maximize potential reduction of pollutants, given a fixed cost. Since three pollutants (fine sediment, nitrogen, and phosphorus) were all thought to be potentially subjected to management controls, the technique was designed to accommodate multiple pollutant control objectives. For example, the problem could be framed such that an optimum solution must simultaneously meet prescribed fine sediment, nitrogen, and phosphorus targets. Figure A1 shows simultaneous control targets for Particles, TN, and TP (32 percent, 5 percent, and 10 percent, respectively) for the minimize cost for a fixed target option. If the maximize benefit for a fixed cost option is selected, the interpretation of the control targets shifts from percent reduction to control priority (a rank from 0, meaning not controlled, to 10 meaning maximum control at any cost). The other user entries under step 1 include the stop tolerance (which was always set to zero for PAT operation), and the number of best solutions to display. Stop tolerance will be further explained with the description of the underlying algorithm.

Step 2 – Define Problem Constraints

The second step in the process is where problem constraints are introduced into the formulation. Each SCG-Setting-Tier combination represents a controllable entry. Constraints can be placed on each of these entries by adjusting either the minimum or maximum LOA, or both. The number of entries in Figure A1 has been condensed for illustrative purposes. In the Figure A1 example, Atmospheric-Setting 1-VE Tiers 2 and 3 have are not considered in the formulation because the maximum LOA is set to zero. The Atmospheric Transportation Infrastructure or Structural (TIOS) tiers have a maximum LOA of 80 percent, while the Urban and groundwater tiers are not unconstrained (0 percent-100 percent). These constraints are for illustrative purposes only. Rationale and assumptions for some of the actual problem formulations are listed in the Assumptions and Rationale section.

Step 3 – Rank Feasible Alternatives by Cost

Once problem formulation is completed, pressing the “Rank Feasible Alternatives” button launches the optimization process. A progress bar is displayed while the optimizer iterates through the search domain to find a set of optimum solutions. Once the search process is completed, the results are summarized in a series of tables in subsequent tabs, as described below:

BestLoad – In addition to the baseline pollutant load budget, this table presents the pollutant load budgets associated with each of the best solutions in by SCG and setting. Cost information (20-year capital, annual operations and maintenance, and 20-year total) are also summarized and presented with the load budget

BestLOA – Each solution has LOA selections for all SCG-setting-tier combinations. This table presents the selected LOA levels for each of the best solutions

BestTiers – This report presents the load reductions associated with each SCG and tier combination. The results are aggregated for all settings within the SCG and rolled up by tier. In addition to the cost and pollutant reduction information, a weighted aggregate LOA is shown for each SCG and tier combination for each solution, to highlight which tier strategies were preferred during optimization.

SettingsTiers – This report includes the most detailed breakdown of the optimization results. It presents the selected LOA for each SCG-setting-tier combination, the associated load reductions, associated cost information, and a SCG rollup summary.

Handout – Similar to the SettingTiers report, the Handout presents the load reductions as a percentage relative to the original untreated pollutant load budget. Another computed data point included in this report is a cost effectiveness estimate computed as an annual total cost per percent particles removed. Cost effectiveness is a normalized quantity for comparing performance between SCG-setting-tier controls.

Description of the Underlying Algorithm

Scatter search is meta-heuristic search technique that has been explored and used in optimizing complex systems (Glover et al. 2000). Scatter search has some commonalties with traditional Genetic Algorithms (GAs); however, there are also a number of quite distinct features. Both scatter search and GAs are *population-based* approaches which search for best solutions by combining existing elements. An individual in the population is one combination of decision variables (LOA) that all fall within the defined constraints. A population is a group of individuals. The objective of both GA and scatter search is to search among the characteristics (LOA) of individuals in the population to find which ones result in the highest quality. Quality is measured as a computed byproduct of cost and percent reduction. If the objective of the problem is to maximize pollutant reduction and minimize cost, populations and individuals in the population which best achieve this objective are considered of better quality.

GA approaches are predicated on the idea of choosing parents (individuals) randomly to produce offspring (other individuals with slightly different characteristics or LOAs), and further on introducing randomization to determine which characteristics (LOA) of the parents should be combined. By contrast, the scatter search approach does not emphasize randomization because it is indifferent to choices among alternatives. Instead, the approach incorporates strategic responses, both deterministic and probabilistic, that take into account the evaluation history. Scatter search focuses on generating relevant outcomes without losing the ability to produce diverse solutions (Laguna and Marti, 2002). In other words, it initially searches across a wide spectrum of populations with the objective function in mind, and uses the evaluation history to hone in on elements in the population which exhibit the highest quality. The *scatter* process helps to reduce the likelihood of finding a local best solution that is not the true best solution. Because of this feature of scatter search, it can find the near-optimal solution in a more efficient way, and serve as a better optimization engine for the type of goal-seeking required for the underlying PAT dataset. The scatter search procedural adaptations that were applied in PAT were developed by Zhen (2002), and have been tested and applied for solving a variety of optimization formulations (Zhen and Yu, 2002, 2004; Riverson et al., 2004; Lai et al., 2006, Lai et al., 2008). The implementation of scatter search processes generally follows the four step approach outlined below:

1. Generate a starting set of diverse points (individuals in a population)

This is accomplished by initially dividing the range of each decision variable (LOA range) into four sub-ranges of equal size. Then, a solution is constructed in two steps. First, a sub-range is randomly selected, and second, a value is randomly chosen from within the selected sub-range. The starting set of solution points also includes all variables at their lower bound (minimum LOA), all variables at their upper bound (maximum LOA), all variables at their midpoints, and any other solution points or constraints imposed by the user. This represents the widest possible search domain, and sets the stage for further refinement of the search.

2. Choosing a subset of diverse points (individuals) as a reference set (population)

A reference set (reference population) is group of individuals with diverse characteristics (LOA). An individual's characteristics all fall within the user-specified set of constraints; however, the algorithm tries to make the characteristics as diverse as possible. With each successive iteration, a new population is created and compared against the reference set. For PAT, the number of individuals in the population is computed as 4 times the number of decision variables (LOA) ranges. So for the initial configuration, the number of LOA ranges was 43, representing all the different combinations of SCG, settings, and tiers;

therefore, the number of individuals in the unconstrained reference set was 43 times 4 or 172. The higher the number, the more diverse the search domain, the greater the likelihood of finding a near optimum solution; however, the sacrifice is an increase in computational intensity. The PAT underlying database was designed to significantly minimize computation time because (1) LOA increases in 10 percent intervals, and (2) all associated pollutant reductions were pre-computed as rating curve elements in the underlying database; therefore, a set of 172 reference set individual solutions was not computationally prohibitive.

A designed objective of the search approach is to seek individuals within a reference set with diverse characteristics. This is done to reduce the likelihood that the solution technique blindly converges upon a local best solution and ignores other potentially better solutions. Population "diversity" is measured by the Euclidean distance of individuals within the reference set. Euclidean distance is the straight line distance between two points. With respect to the PAT dataset, two points exist in a multi-dimensional space with as many dimensions as there are decision variables (LOA), which would be in 43-dimensional space. An example of Euclidean distance in a two-dimensional plane with point 1 at (x_1, y_1) and point 2 at (x_2, y_2), it is computed as:

$$\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$$

3. Generating new points to update the reference set

While searching for the optimum solution, Scatter Search uses a linear combination method to construct new solution points (individuals) from the reference solution points. This combination is based on the three types of formulations, in which x' and x'' are reference solution points, and x_1 through x_3 are the newly generated solution points:

$$\begin{aligned} x_1 &= x' - d \\ x_2 &= x'' - d \\ x_3 &= x' + d \end{aligned}$$

Where $d = r\dfrac{x'' - x'}{2}$ and r is a random number in the range (0, 1).

The number of solution points created from the linear combination of two reference solution points depends on the quality of the solutions being combined.

In the course of searching for a global optimum, the reference set is continuously updated. The solutions having better quality, while preserving or refining the diversity of the reference set, are chosen as replacements for old points in the reference set. A diversification step is performed to repopulate the reference set. To preserve quality, a small set of the best solutions (or elite set) from the current reference set is used to seed the new reference set. The diversification method is used to repopulate the reference set with solutions that are diverse relative to the elite set. If higher quality elements are found among the new diverse set, they become part of the elite set in successive reference sets. This step continues until no solutions or higher quality are available, or until a user specified stopping criteria is achieved.

4. Stop the search if the stopping criteria are met

The stopping criteria can be defined, as the maximum number of iteration runs, or the minimum improvement between updates of the reference set, or both. If both criteria are specified, then the search will be stopped when either of the criteria is met. In the case of PAT, the minimum improvement was set at zero, so that iterations would stop when absolutely no better solution could be found. After multiple sensitivity runs using the PAT dataset, the maximum number of iterations was empirically computed as:

$$MaxRuns = n^2 \times 4$$

Where n = the number of decision variables.

For the initial PAT configuration, the maximum number of runs was 118,336. While this number seems like a large number, the number of possible solutions for an unconstrained simulation is several orders of magnitude higher. Given 43 decision variables, if all setting-tier and LOA constraints removed, and each decision variable has the option to vary between 0 percent and 100 percent (at 10 percent LOA intervals), there would be 11 possible solutions for each of the 43 variables. The unconstrained number of possible solutions is 11 raised to the 43rd power, or roughly 602,400,691,612,422,000,000,000,000,000,000,000,000,000,000 combinations. The number drops a few orders of magnitude as various combinational constraints are introduced. Nonetheless, it punctuates the potential efficiencies to be gained with the application of meta-heuristic solution techniques like scatter search.

Assumptions and Rationale

There were a few key assumptions that have been made regarding the nature and usage of the underlying database derived from the SCG work products. These assumptions are listed below in bullet form, with associated narrative explanations.

- The SCG products for each setting and tier combination represent maximum application of the associated controls. These controls were quantity based, such that depending on the SCG and setting, controls are expressed in terms of amount of area treated, number of objects controled, or length of segment treated. For this reason, it was both possible and appropriate to scale the LOA for these controls according to the applicable quantity. For each SCG, LOA was scaled linearly from 0 percent (baseline condition with no controls) to 100 percent (maximum application) for a given Tier and setting combination. Pollutant reduction was linearly scaled between baseline loads to reduced loads associated with each setting and tier combination. Associated management costs were scaled from zero cost for the baseline to 100 percent of the cost for full application to a given setting and tier combination.

- The setting is the smallest unit for management, for which there is a fixed manageable quantity. For example, given a specific Urban Upland setting, the fixed manageable quantity is area. Therefore, a 50 percent application level of Tier 1 means that the suite of controls associated with Tier 1 are applied to 50 percent of the total available area. If during the solution search routine, additional controls are found to be required for that specific setting in order to meet the defined objectives, it can be achieved by either (1) increasing the LOA for that particular Tier, (2) applying a different LOA of another Tier (i.e. Tier 2) which has a higher treatment potential, or applying combinations of LOA for more than one Tier (i.e. 50 percent Tier 1, and 20 percent Tier 2, for a total of 70 percent of the total area being treated).

- The maximum LOA for any given setting was assumed to be 80 percent. For practical reasons, it was thought unlikely that any given combination of tiers could be applied so as to treat 100 percent of a given setting. There will always be urban areas which cannot be treated due to restricted access or impracticability, remote forest settings which are naturally erodible and/or are not accessible by conventional means, private property air pollutant sources or vehicle emissions that cannot be managed for various reasons, or stream segments which cannot be easily stabilized and restored.

- There were certain assumptions associated with LOA constraints for the various packages. These include definition of the base package as well as selected minimum/maximum LOA constraints for some of the exploratory packages. These were introduced to limit the selection of some of the more sophisticated, but untested technologies. For example, lets assume that Tier 1 of a given SCG and setting is composed of common conventional practices, while Tier 2 includes some sophisticated and innovative practices. A scenario that focuses on traditional control technologies may restrict the selection of Tier 2 practices, in favor of Tier 1; whereas a scenario that focuses on innovative practices might constrain the selection of Tier 1, and allow more selection of Tier 2 practices. Assumptions associated with scenarios are described in detail in Chapter 3: Development of the Recommended Strategy.

Uncertainty associated with methods and results

All calculations and estimates include uncertainty because of the current limits of scientific understanding. There is uncertainty associated with (1) the data developed in the underlying database, as well as uncertainty associated with (2) the optimization searching processs.

Underlying Database Uncertainty

Since the SCG products all represent the ultimate pollutant delivery to Lake Tahoe associated with the SCG settings and tiers, the predicted confidence of an individual solution can be computed as a weighted average of the confidence associated with the components. A uniform confidence rating convention was developed to help quantify uncertainty associated with the diverse set of SCG products. Each of the SCGs noted uncertainties throughout their analyses and assigned an overall confidence rating to each set of results provided. The assigned confidence values were rated on a one to five scale according to a system designed for the SCG's use. The rating was based on the SCGs own answers to 16 questions about the data sources used, the calculation results and modeling parameters.

SCGs used the following guidance:

A rating of "1" generally indicates:

- Data sources were from a dissimilar area, were unreviewed and not supported by other research
- Calculation results were not similar to other investigations, used mostly professional judgment, had high calculation error, and required unlikely assumptions
- Models were not widely accepted, were poorly calibrated, or were not validated

A rating of "3" generally indicates:

- Data sources were from a similar, cold climate; were reviewed as agency drafts; or were partially supported by other research

- Calculation results were somewhat similar to other investigations, used some professional judgment; had intermediate calculation error or required reasonable assumptions
- Models had been used before, were reasonably calibrated but might not have been well validated

A rating of "5" generally indicates:
- Data sources were from Tahoe, published, and supported by other research
- Calculation results were similar to independent investigations, used little professional judgment, had low calculation error, and were based on conservative assumptions
- Models were widely accepted, well calibrated, and validated on non-calibration data

Overall, ratings of 1 and 2 were used when future values were considered likely to change significantly, and the SCG was not comfortable using them for significant management decisions. Ratings of 3, 4, and 5 were used when future values are not expected to change significantly, and the information is considered appropriate for management decisions. Further detail about confidence ratings for each of the individual SCG components is provided in the Lake Tahoe TMDL PRO Report (Lahontan and NDEP 2007b).

Optimization Uncertainty

Meta-heuristic optimization approaches are based on random number search techniques. Uncertainty increases with the prevalence of local minimums to which the solution technique might become trapped, and miss potentially better solutions within its search vicinity. Specialized approaches like scatter search include considerations for diversity during the searching process, which minimizes the likelihood of falling into a less than optimal local minimum. If the resolution around other potential solutions within the general vicinity of the near-optimum solution is amplified, one can better illustrate and quantify the nature of uncertainty associated with the PAT optimization technique. In optimization space, all points which are said to be optimal will fall along a region called the Pareto optimal frontier.

Pareto optimality is a concept commonly applied in economics and engineering. The underlying premise for PAT optimization involves the allocation of monetary resources to reduce pollutant loads such that the lowest cost alternative is achieved. Additional constraints imposed upon the solution can bend the search space to ensure that the allocation of resources involves all stake holders involved. During the search operation, reallocation of resources is done in such a way that marginally more benefit is achieved with each additional cost incurred. Cost effectiveness is measured by benefit achieved per dollar allocated. To remain Pareto optimal, the solution technique selects the more cost effective solutions first before considering others. A solution is said to be Pareto optimal when no further improvements can be made.

A second optimization formulation was tested against the Scatter Search formulation in PAT to help amplify the resolution of solutions in the general vicinity of the selected near-optimum solution. The primary objective of this formulation is to identify the Pareto optimum frontier. This formulation is a called the Non-dominant Sorting Genetic Algorithm (NSGA-II), developed by Deb et al. (2000). While scatter search refines a scatter pattern around the targeted objectives by replacing members of a reference population, NSGA-II defines a population as individual solutions along a cost-benefit frontier and refines the entire population with better and better solutions until the final solution approaches the true Pareto frontier. Figure A2 is a conceptual representation of the two search routines.

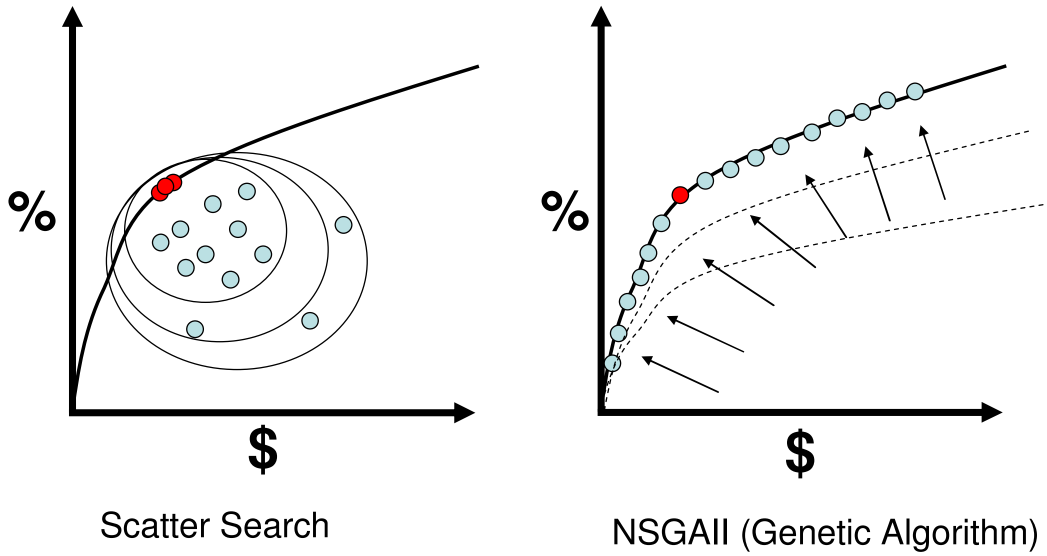

Figure A2. Conceptual representation of the searching techniques for Scatter Search and the Non-dominant Sorting Genetic Algorithm-II

In both graphs, the Pareto Optimum frontier is represented as the solid cost-benefit arc. It conceptually represents the true collection of optimum solutions for every cost ($) and benefit (percent) combination. The concentric circles in the scatter search graph illustrate progressively better (narrower) reference populations, until the final set of best solutions are found clustered around the point along the Pareto frontier. The dashed lines in the NSGA-II graph illustrate progressively improving cost-benefit relationships with each new generation of solutions, until the true Pareto frontier is approximated in the last generation. The red dot on the NSGA-II graph shows how a near-optimum solution for a given objective target might be inferred from the Pareto frontier. The distance between points along the trade-off curve for both methods shows that scatter search clusters solutions around the defined objective on the Pareto frontier, while NSGA-II distributes the solutions along the entire trade-off frontier. Increasing the resolution for NSGA-II means increasing the number of individuals in the population, which exponentially increases the number of generations required to approach the near optimum Pareto frontier.

The scatter search predicted PAT results were compared against results generated by the NSGA-II using the same underlying database and constraints. The NSGA-II population size was set to 1,376, which equals the number of decision variables (43) times 2^5 or 32. The population size must be increased in increments of 2 times the number of decision variables in accordance with requirements associated with parent-offspring Genetic Algorithm formulations. The resulting Pareto frontier is a high resolution rendering based on 1,376 individual solutions along the frontier. Due to the increased complexity of the problem and the high-resolution population size, each NSGA-II simulation required roughly 20 minutes of runtime to complete the simulation to generate the Pareto optimum frontier.

After the simulation, an empirical relationship derived from on sensitivity runs of the Lake Tahoe Clarity Model was used to translate the results from total fines reduction to resultant lake clarity (represented as Secchi Depth). Secchi depth versus 20-year capital cost Pareto Frontiers are presented for some of the exploratory scenarios in Figure A3. These exploratory scenarios are described in greater detail in Chapter 3: Development of the Recommended Strategy.

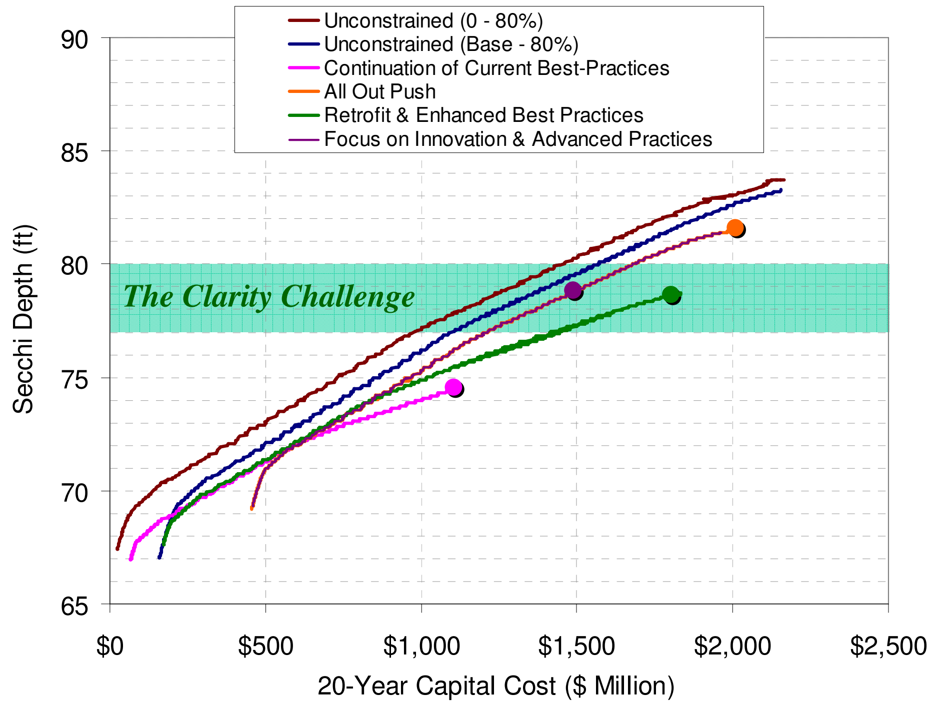

Figure A3. Pareto optimum frontiers for IWQMS exploratory scenarios.

The solid color-coded dots represent the "best" PAT-predicted scatter search solutions plotted along the Pareto Frontier associated with the corresponding scenario. In addition to the individual scenario Pareto frontiers, the unconstrained simulation frontier and a base-package-constrained simulation frontier were also generated for comparison. These illustrate the impact of the LOA constraints on the predicted optimal solution at each cost interval.

Further investigation of the LOA characteristics for the solutions in and around the PAT predicted optimum solutions yields insight into the factors that influence the selected solution at different intervals. Consider the hypothetical unconstrained (0-80 percent) simulation. Each of the 1,376 points along the Pareto optimal frontier includes characteristic information about the selected LOA for every setting and tier combination. If the predicted cost at each percent reduction interval is broken down and summarized by SCG-and-Setting combination, it is possible to see the order in which the optimizer tends to select cost-effective solutions at each cost interval along the Pareto frontier. Figure A4 is a graph of Pareto-optimal cost distribution for SCG-and-settings combinations, normalized at each percent reduction interval along the frontier. Likewise, if cost is broken down and summarized by SCG-and-Tier combination, it is possible to see the order in which the optimizer tends to select cost-effective solutions at each cost interval along the Pareto Frontier. Figure A5 is a graph of Pareto-optimal cost distribution for SCG-and-tier combinations, normalized at each percent reduction interval along the frontier.

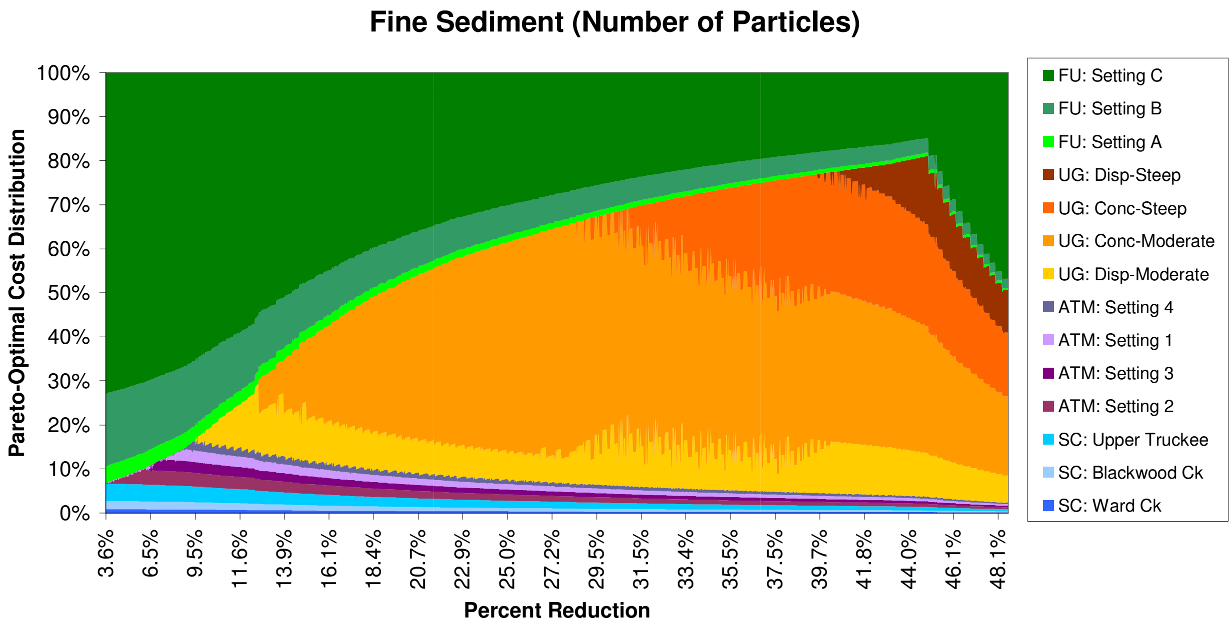

Figure A4. Pareto-optimal cost distribution versus percent fines reduction by setting

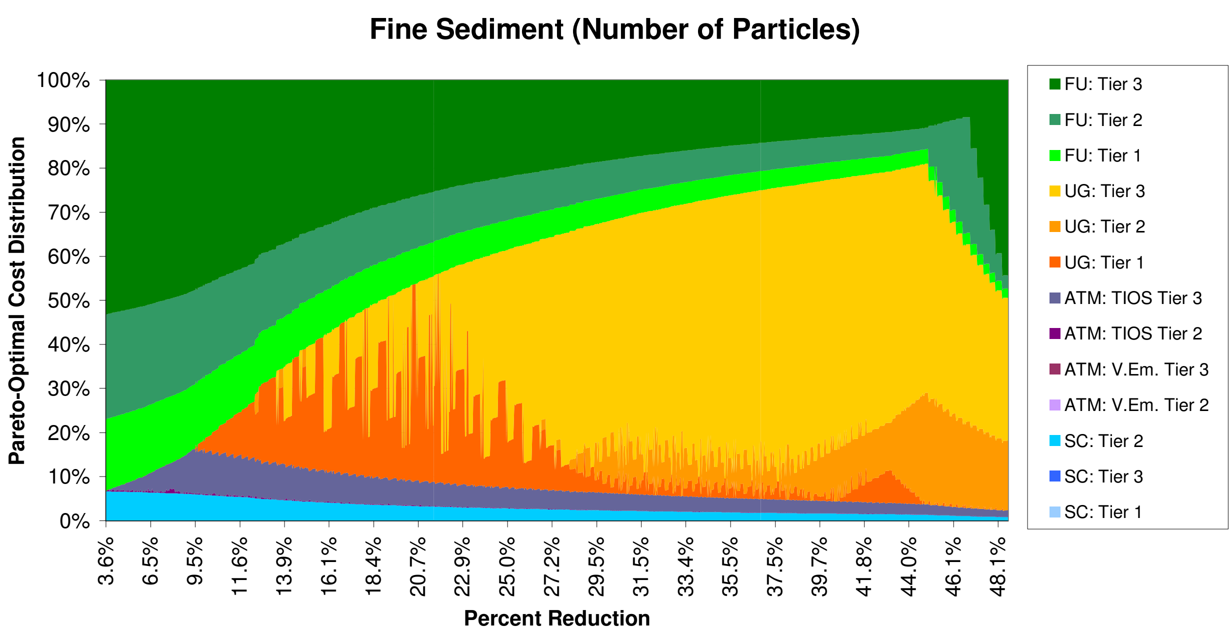

Figure A5. Pareto-optimal cost distribution versus percent fines reduction by tier

At certain intervals, these graphs show a systematic pattern of “noise” in the cost distributions. This occurs when the aggregated cost-effectiveness for different combinations of LOA among settings and tiers are comparable. This pattern is especially prevalent among the urban SCG settings and tiers. There are a number of factors which contribute to this noise in the results. Climate variability around the watershed influences the effectiveness of stormwater controls. Since setting and tier results are aggregated across multiple spatial locations, there are sometimes overlaps in aggregated reduction per dollar invested. When searches target percent reduction criteria that are in relatively stable sections of the

Pareto frontier, scatter-search will tend to find the same solution consistently in consecutive runs. However, when scatter-search explores regions along the Pareto frontier where there is a lot of "noise" generated by the interplay of similar aggregated cost-effectiveness values among individual solutions, it increases the likelihood of finding a local best solution which may actually have slightly better solutions nearby. Figure A6 is a conceptual representation of what a scatter search solution might look like while operating in a noisy region along the Pareto frontier. The dips in the figure are deliberately exaggerated to highlight how a local minimum might be chosen as the solution in lieu of a slightly better one nearby.

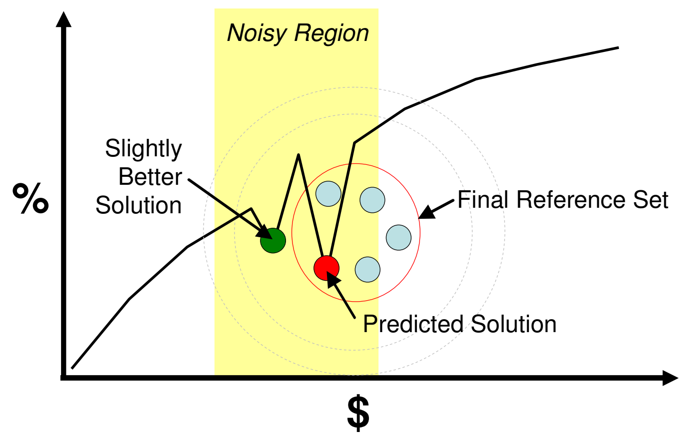

Figure A6. Conceptual representation of scatter search in a noisy region along the Pareto frontier

To validate the predictive performance of scatter search and NSGA-II, Lai et al. (2008) tested both of these search techniques against a known hypothetical linear solution, where no noise existed in the underlying dataset. The objectives were twofold: (1) to evaluate Scatter Search ability to pick a known linear solution for a single BMP given multiple pollutant performance functions and multiple pollutant reduction objective criteria, and (2) to evaluate NSGA-II ability to generate a trade-off curve for a known linear solution for a single BMP. Both the Scatter Search and NSGAII optimization techniques were able to solve a known linear solution with 100 percent accuracy. In addition, the optimization techniques were both able to select an optimum solution given multiple control objectives for controlling sediment and nitrogen simultaneously. The NSGA-II technique was also able to predict a known linear trade-off curve, as shown in Figure A7.

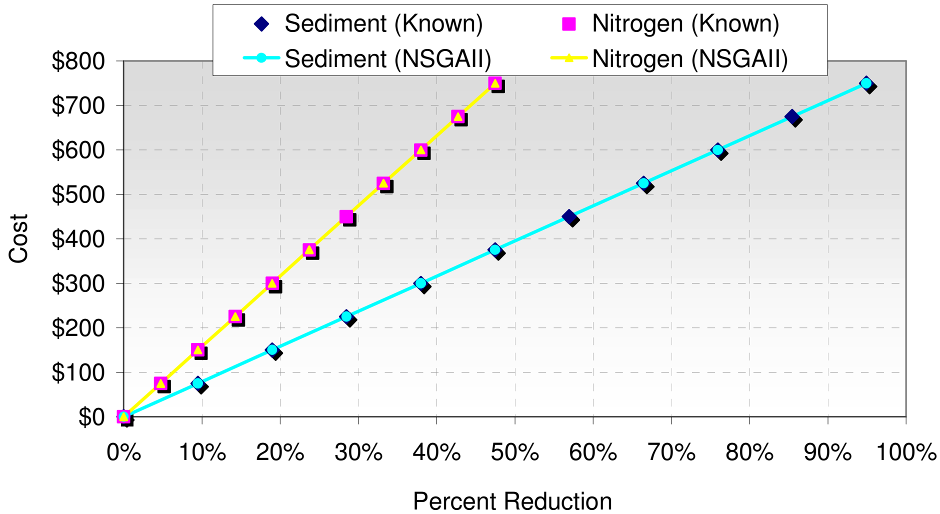

Figure A7. Validation of NSGA-II against a known linear solution (Source: Lai et al., 2008)

In conclusion, it is important to note that the results of meta-heuristic search techniques are usually referred to as near-optimal solutions. Uncertainty of a scatter search solution heavily depends on the amount of "noise" in the region along the Pareto frontier which is being searched. Just as the entire Pareto frontier shifts when LOA constraints are added to given problem formulation, the noise associated with the interplay of cost-effectiveness also varies as constraints are imposed upon the search formulation. In the absence of noise, both scatter search and NSGA-II have been shown to consistently predict results with 100 percent accuracy. The predictive efficiency for scatter search given the nature of the underlying database is probably on the order of 98 percent to 99.9 percent efficient in terms of accuracy. However, further trade-off analysis of this manner is needed to better quantify uncertainty associated with specific scenario results.

Works Cited

Glover, F., M. Laguna, and R. Marti (1999). "Scatter Search to Appear in Theory and Applications of Evolutionary Computation: Recent Trends." A. Ghosh and S. Tsutsui (eds.), Springer-Verlag.

Lai, F., L. Shoemaker, X. J. Zhen, K. Alvi, J. Riverson (2008). "Optimizing BMP Placement at Watershed-Scale Using SUSTAIN." World Water and Environmental Resources Congress 2008 Conference Proceedings, Honolulu, HI.

Laguna, M., and R. Marti (2002). "The OptQuest Callable Library to Appear in Optimization Software Class Libraries." S. Voss and D.L. Woodruff (eds.), Boston: Kluwer Academic Publishers, pp. 193-218.

Riverson, J., J. Zhen, L. Shoemaker, and F. Lai. (2004). "Design of a Decision Support System for Selection and Placement of BMPs in Urban Watersheds." World Water and Environmental Resources Congress 2004 Conference Proceedings, Salt Lake City, UT.

Zhen, X. J., (2002). "Development of Best Management Practice (BMP) Placement Strategies at the Watershed Scale." Ph.D. Dissertation, Department of Civil and Environmental Engineering, Univ. of Virginia, Charlottesville, VA.

Zhen, X. J. and S. L. Yu, (2002). "Optimization of Best Management Practices (BMPs) Placement for Non-Point Source Pollution Control at the Watershed Scale." Conference on Water Resources Planning and Management, ASCE-EWRI and Virginia Tech, Roanoke, VA.

Zhen, X. J. and S. L. Yu. (2004). "Optimal location and sizing of stormwater basins at watershed scale." *Journal of Water Resources Planning and Management* 130(4), July edition.