FORMOSA-P CONSTRAINED MULTIOBJECTIVE SIMULATED METHDOLOGY

Paul M. Keller
Department of Nuclear Engineering
North Carolina State University
P. O. Box 7909
Raleigh, NC 27695-7909
Email: keller@eos.ncsu.edu

Keywords: Optimization methods, nuclear fuel management

ABSTRACT

This paper describes improvements to the multiobjective simulated annealing (MOSA) algorithm for the purpose of exploring the tradeoff surface between multiple objectives for constrained incore fuel management optimization problems. A two step annealing schedule is introduced to control the simulated annealing temperatures for the primary objectives, and an adaptive penalty control algorithm is used to control the simulated annealing temperatures for the constraint objectives. It was found, for the constrained problems of interest, that nondominated archives must be limited to loading patterns without constraint violations. The algorithm's performance is illustrated and compared with separate single objective simulated annealing (SOSA) optimizations. It is shown, for a two objective multiobjective search, that the MOSA algorithm is able to both nearly span the results obtained from SOSA optimizations, and also describe the tradeoff surface between them. For a three objective multiobjective optimization, the method is rather more limited, as although it is able to describe somewhat of the tradeoff surface between SOSA optimizations, it is not able to completely span the results of all three corresponding single objective optimizations.

1. INTRODUCTION

The incore fuel management code FORMOSA-P was developed to determine the family of near-optimum loading patterns (LPs) for pressurized water reactors (PWRs). The FORMOSA-P code utilizes the non-linear stochastic optimization approach of simulated annealing (SA), wherein an objective function is calculated to evaluate the suitability of the candidate LPs. This objective function takes into account a user selected core attribute which is to be minimized or maximized, and a number of constraints on other core attributes. The SA algorithm determines the near-optimum LPs by minimizing the value of the objective function over the course of the optimization. Current versions of FORMOSA-P offer the following user-selected objective functions for minimization or maximization:

- Maximization of end-of-cycle (EOC) reactivity, used as a proxy for fuel cycle length maximization.

- Minimization of radial power peaking.
- Maximization of discharge fuel region average burnup.
- Minimization of feed fuel enrichment (for selected feed fuel assemblies).

These user selected objective functions are referred to in this paper as the primary objectives, as compared with constraint objectives which will be introduced later. Because the core designer's task is to minimize the reactor's fuel cost over a given planning horizon while ensuring that various operational constraints are met, determining the tradeoff surface between two or more of these primary objectives is often of interest (Parks, 1996). In all cases, the tradeoff described between primary objectives must comply with operational constraints, otherwise, the results are of no interest to the core designer. Obtaining such tradeoff information from previous versions of FORMOSA-P, which offered only single objective optimizations, can be difficult and time consuming, both in terms of computational time and, more importantly, the designer's time. This is because single objective SA (SOSA) optimizations tend to provide only narrowly focused results of the best solutions in terms of one of the above objectives, often with very little tradeoff information. While characterization of a multiobjective tradeoff surface using SOSA is possible, it requires performance of multiple SOSA optimizations, each using a different target constraint limit for an objective which is treated as a constraint. For example, to determine the tradeoff surface between radial power peaking and feed fuel enrichment, multiple SOSA feed fuel enrichment optimizations could be performed, each using a different target radial power peaking factor value for the radial power peaking constraint. Because increments as small as 0.01 in target radial power peaking may be required to provide adequate resolution of the tradeoff surface, determining the radial power peaking/ feed fuel enrichment tradeoff surface over a radial power peaking factor span of 0.1 could require as many as ten SOSA optimizations. Obtaining such tradeoff surface information at the cost of fewer LPs evaluated from a single multiobjective optimization would be useful, both from the standpoint of the decreased computational expense, and also due to the decrease in the engineering designer's time required.

Characterization of a tradeoff surface using multiobjective simulated annealing algorithm (MOSA) was first proposed by Engrand (Engrand, 1997), and improved upon by Supppapitnarm and Parks (Suppapitnarm, 1999). This paper describes further improvements to the MOSA algorithm needed to deal with constrained incore fuel management optimization problems.

2. SIMULATED ANNEALING

The SA algorithm, originated by Kirkpatrick (Kirkpatrick, 1983), is a random search technique which not only accepts changes which decrease the objective function f, but also some changes which increase it. The latter are accepted with a probability described by:

$$p = \exp\left(-\frac{\delta f}{T}\right), \tag{1}$$

where δf is the increase in f and T is a control parameter, commonly referred to as the simulated annealing temperature, by analogy with the annealing of crystalline materials. The optimization is started with a simulated annealing temperature sufficiently high that nearly all changes are accepted, and then lowered in accordance with a given control or annealing schedule in order to achieve convergence of the optimization. SA's advantages over other optimization methods are its ability to avoid becoming trapped in local minima, and its lack of need for functional derivative information.

Because of the lack of functional derivative information, all SA algorithms save a collection of solutions for later investigation, referred to as archives, which may be of interest to the designer. In the previous FORMOSA-P single objective SA algorithm, two archives of user specified size, in terms of number of LPs archived, are maintained. Archiving of LPs in these archives is performed after LP acceptance, and only accepted LPs are considered for archiving.

2.1 Multiobjective Simulated Annealing

Engrand's MOSA algorithm (Engrand, 1997) proposed forming a composite objective function G defined as:

$$G = \sum_{i=1}^{N} \ln(f_i) \tag{2}$$

where f_1,..., f_N are the multiple objective functions for which the tradeoff is to be determined. The acceptance probability was then given by:

$$p = \exp\left\{-\frac{[G(\underline{X}_{n+1}) - G(\underline{X}_n)]}{T}\right\} = \exp\left\{-\frac{1}{T}\sum_{i=1}^{N} \ln\left(\frac{f_i(\underline{X}_{n+1})}{f_i(\underline{X}_n)}\right)\right\}. \tag{3}$$

This means that any perturbation which decreases the composite objective function G will always be accepted, while any perturbation which increases G will be accepted with a probability inversely proportional to the increase in G, and with a probability which falls as the simulated annealing temperature T falls. In common with previous FORMOSA-P practice, archiving is performed after the SA acceptance test, and only accepted LPs are considered for archiving.

A major drawback of Engrand's MOSA algorithm, observed by Suppapitnarm and Parks, is the formation of a composite objective function, and its consequent need for careful formulation of the individual objective functions in order to assure that one objective is not favored over another. Suppapitnarm and Parks also fault the Engrand MOSA algorithm for only considering LPs which have been accepted by SA for archiving. Suppapitnarm and Parks observed that this may preclude an LP which moves onto the desired tradeoff surface from being accepted. Instead, Suppapitnarm and Parks proposed submitting candidate LPs for archiving prior to being subjected to the SA acceptance test. If a candidate LP is not archived, then, in the Suppapitnarm and Parks algorithm, each objective function f_i is assigned its own simulated annealing temperature T_i, and each perturbation which was not archived is accepted with a probability described by:

$$p = \prod_{i=1}^{N} \exp\left\{-\frac{[f_i(\underline{X}_{n+1}) - f_i(\underline{X}_n)]}{T_i}\right\}. \tag{4}$$

This eliminates the need for the formation of a composite objective function, and its consequent careful scaling of the individual objective functions, provided a means for automatically determining the initial SA temperatures exist. Suppapitnarm and Parks proposed using White's formula (White, 1984), which is also used in the single objective SA algorithm of FORMOSA-P, which is:

$$T_i = \sigma_i, \tag{5}$$

where σ_ι is the observed standard deviation of objective function i during a temperature initialization survey which precedes the SA optimization in which all feasible solutions are accepted. After SA temperature initialization, the temperatures are periodically lowered in accordance with the formula:

$$T_i^{k+1} = \alpha_i^k T_i^k, \tag{6}$$

where α_i^k is determined using Huang's formula (Huang, 1986):

$$\alpha_i^k = \max\left[0.5, \exp\left\{-\frac{0.7T_i^k}{\sigma_i^k}\right\}\right] \tag{7}$$

Suppapitnarm and Parks further suggest that, in this algorithms, constraints can be handled in several ways. One possibility suggested is to apply constraints as “hard” constraints, that is, any perturbation with a constraint violation can be neither accepted nor archived. Another possibility is that constraints may be handled as additional, nonnegative valued objective functions whose values are to be minimized simultaneously with the primary objectives. In this paper, these additional objectives are referred to as “constraint objectives.” In this work, all calculated constraints are applied as constraint objectives, as these permit the optimization to transition across infeasible design space, and, in nearly all cases, this permits the optimization to obtain superior results in comparison to those results obtained from optimizations in which are applied as “hard,” or true-false constraints.

While the elimination of the composite objective function and introduction of individual SA temperatures for each objective is a very useful innovation, the MOSA algorithm of Suppapitnarm and Parks suffers from a number of drawbacks when applied to constrained incore optimization problems, particularly those with a large number of constraints. First, Eq. (4) suffers from the drawback that it is easily subject to arithmetic underflow and overflow errors when a sampled loading pattern results in a large change, either positive or negative, in one of the objective function values. To correct this, it is possible to observe that the acceptance probability formula may be equivalently written as:

$$p = \exp\left\{-\sum_{i=1}^{N} \frac{[f_i(\underline{X}_{n+1}) - f_i(\underline{X}_n)]}{T_i}\right\}. \tag{8}$$

This avoids arithmetic overflow and underflow errors which may result when an LP perturbation results in a large change in one of the objective function values.

A more serious problem arises with the annealing schedule proposed by Suppapitnarm and Parks. It may be observed that Eq. (7) results in an annealing schedule which not only depends upon the number of active objective functions in the optimization, but also that, as the number of objective functions increases, a more rapid SA cooling schedule results. It was found that this frequently leads to local minima trapping, particularly for problems in which the reference LP is far from the tradeoff surface of near-optimal LPs. Even for cases where it was believed that local minima trapping has not occurred, Eq. (7) did not permit a search length adequate to characterize a tradeoff surface in the presence of constraints, even for practical optimization problems of minimal complexity, involving just two primary objectives and a single active constraint. As a result of these issues, for the primary objectives, an annealing schedule which uses a modification to Huang's formula is proposed:

$$\alpha_i^k = \max\left[0.5^{\frac{1}{N}}, \exp\left\{-\frac{a_l T_i^k}{N\sigma_i^k}\right\}\right], \tag{9}$$

where N is the total number of active objective functions in the optimization, both primary and constraint, and a_l is a constant. For the purpose of Eq. (9), the number of active objective functions is considered to be the total number of objective functions with non-zero standard deviations for the annealing step. For the constraint objectives, the simulated annealing temperatures are controlled by the same adaptive control algorithm used to control the penalty factors which are utilized by the penalty constraints in single objective optimizations (Keller, 1996, and Kropaczek, 1994). Additionally, an objective function dependent multiplier, c_i, is used for the simulated annealing temperature initialization:

$$T_{i = c_i\sigma_i}. \tag{10}$$

Eq. (10) is applied to both the primary objectives and the constraint objectives.

Suppapitnarm and Parks report that their MOSA algorithm will establish the optimization search on the tradeoff surface relatively early in the optimization, with the remainder of the optimization effort filling in and extending the tradeoff surface. While it was found in this work that the MOSA algorithm could be induced to approach the tradeoff surface rather faster than the tradeoff surface could be characterized, nonetheless, a substantial search effort may still be required to arrive at a position close to the tradeoff surface, particularly if the optimization begins with an LP far from the tradeoff surface. A considerable further search is required to adequately characterize the surface and extend the ends of the tradeoff surface. During this period, some further extension of the tradeoff surface is generally observed. As a result of these observations, a two stage annealing schedule for the primary objectives has been introduced. At the start of the SA cooling cycle, the constant a_l in Eq. (9) has a higher value to promote a rapid transition onto the tradeoff surface. After the SA acceptance ratio has fallen below a certain value, a value of 0.15 was used in this work, the constant a_l is substantially reduced to permit an extended SA search along the tradeoff surface at a nearly constant acceptance ratio. For the work presented in this paper, first and second stages values of 0.4 and 0.05, respectively, were used for $a_{l.}$ In coordination with this, the control algorithm which controls the temperatures for the constraint objectives is adjusted such that any constraint violations are eliminated at about the same time that the second stage annealing schedule is initiated. The MOSA optimization search is terminated when the acceptance ratio falls below a certain value, currently a value of 0.02 is used, or the LPs sampled reaches a value described by:

$$N_{max} = 45 N_p N_{ch}, \tag{11}$$

where N_p is the number of primary objectives in the optimization, and N_{ch} is the number of single LP changes available in the design sampling space. Since, for typical design problems of practical interest, N_{ch} will have a value varying from several hundred to as much as a few thousand, this results in multiobjective searches which require that one or two hundred thousand LPs be sampled.

2.2 Archiving

Suppapitnarm and Parks' MOSA algorithm submits candidate LPs as candidates for archiving prior to application of the SA acceptance test. If the solution is archived, then it is automatically accepted. The archive maintained is a nondominated archive. A solution $\underline{X}$ is dominated by solution $\underline{Y}$ if:

$$f_i(\underline{Y}) < f_i(\underline{X}) \forall i = 1, N, \tag{12}$$

that is, if solution $\underline{Y}$ is superior for all objectives. If a candidate solution dominates any solutions in the archives, then those solutions are removed from the archive, and the new solution is added. If the new solution is dominated by any solution in the archive, then the candidate solution is not archived. If the candidate solution neither dominates nor is dominated by any member of the archive, it is added to the archive.

There are several difficulties which arise from such archiving when dealing with constrained multiobjective optimization problems in which the constraints, both primary and constraint objectives, are treated identically as additional objectives functions to be minimized. First, this provides no mechanism for the removal from the archive of solutions with constraint violations in favor of solutions without constraint violations when the solutions without constraint violations have inferior values for one or more of the primary objectives. This negative tradeoff between a constraint and one of the primary objectives, in general, always occurs late in the cooling cycle for constrained optimization problems. It is important to note in this context that the solutions with constraint violations, even if they posses superior values for one or more of the primary objectives, are of no interest to the designer. A second, and more important difficulty, is that, late in the cooling cycle, the negative tradeoff between the constraints and primary objectives results in a situation in which nearly every solution sampled neither dominates nor is dominated by the existing solutions in the archive. If these LPs are added to a nondominated archive, and then accepted by the optimization search algorithm, this results in a convergence failure. As a result of these considerations, if a nondominated archive is to be maintained during a constrained multiobjective optimization, the nondominated archive must be limited to those solutions which have no constraint violations. Such a nondominated archive of LPs meeting all constraints has been added to the FORMOSA-P MOSA methodology.

Concerning the issue of testing LPs for possible addition to the nondominated archive prior testing for acceptance via the SA acceptance formula, it was found that there was only a significant advantage to archiving prior to SA acceptance for three objective multiobjective optimizations. For two objective optimizations, the advantage of nondominated archiving prior to SA acceptance was very small or insignificant, since, typically, only 20 LPs which would have been rejected by SA out of the total 100,000 sampled would be archived by MOSA, and thus accepted. It must be noted in this context that a three objective optimization is of far less practical interest than a two objective optimization.

The two standard FORMOSA-P archives used for single objective SA (Kropaczek, 1991) are also still utilized by the MOSA algorithm, including the existing algorithms for determining archive binning. In the single objective formulation with all adaptive penalty constraints, a candidate LP solution will replace an existing solution if the candidate's augmented objective function, that is, the objective function with all penalties due to constraint violations added, is less than the recalculated augmented objective function of the previously archived solution. The recalculated augmented objective function is the augmented objective function recalculated using the current, adaptively controlled, penalty factors. For multiobjective optimizations, we may define a measure for whether a candidate LP is "superior" to an archived LP in the multiobjective algorithm by observing that the sum in the exponential argument of Eq. (8) may be viewed as a temperatureweighted composite objective function. The candidate LP $\underline{X}_c$ will then replace the archived LP $\underline{X}_a$ if:

$$\sum_{i=1}^{N} \frac{[f_i(\underline{X}_c) - f_i(\underline{X}_a)]}{T_i} < 0 . \tag{13}$$

The FORMOSA-P MOSA algorithm, like that of Engrand and Suppapitnarm and Parks, also periodically performs a "return to base" maneuver in which the optimization search is restarted from a previously archived LP. The FORMOSA-P MOSA return to base maneuver only returns to the current best LPs in terms of each of the primary objectives on an alternating basis. This is in contrast to the Suppapitnarm and Parks methodology, which stochastically returns to either the best LPs or a certain number of previously archived LPs. In the current work, in particular, it was found that returning to archived LPs which were not the best in terms of any objective was highly detrimental to the quality of the final, near-optimal LPs. It is believed that this is the due to the fact that the archived LPs necessarily lag behind the current search progress, and that returning to those LPs only serves to set back the optimization search progress. The use of alternating returns to base in the FORMOSA-P methodology also provides for more consistent results over stochastic sampling for the return to base maneuver. This is attributed to the relatively small number of return to base maneuvers performed during a typically optimization, which, for the two objective optimizations of greatest interest, typically number only a few tens of maneuvers.

3. ALGORITHM PERFORMANCE

Presented in Figure 1 are the radial peaking vs. feed fuel enrichment results for a two objective multiobjective optimization using radial power peaking minimization and feed fuel enrichment minimization. Also illustrated in Figure 1, are the results from two single objective optimizations for each objective separately, and the 22 nondominated LPs remaining in the nondominated archive of the multiobjective optimization. The reactor core modeled in Figure 1 is a quarter core model of a Westinghouse-designed four loop core with a constraint on maximum soluble boron concentration boron which is difficult to meet, with only 4140 LPs out of the total of 107275 LPs sampled during the multiobjective optimization having no constraint violations. For all three optimizations illustrated in Figure 1, only those LPs without constraint violations are shown. The cross section library used for this problem provided a selection of 30 burnable poison designs, a large number typical of current production-quality FORMOSA-P cross section libraries. The single objective enrichment minimization used a target radial power peaking value of 1.40 as its power peaking constraint limit. An input Δk_∞ range for shuffling of 0.15 out of a possible 0.33 was used, resulting in an optimization design problem having 1170 single LP changes available in the design space, and effectively prohibiting changes to the feed fuel pattern during the optimization.

During the course of the multiobjective optimization illustrated in Figure 1, a total of 249 LPs were added to the nondominated archive, and 227 LPs were removed, mainly evenly distributed over the last three-fourths of the optimization. This is indicative of continued optimization progress occurring over the entire optimization. Very few LPs were added to the nondominated archive in the first quarter of the multiobjective optimization because very few LPs without a constraint violation were found during that period.

It may be observed in Figure 1 that the multiobjective optimization minimum enrichment values closely match the values obtained from the single objective enrichment minimization. The minimum radial power peaking values obtained from the multiobjective optimization are less than 0.01 above the minimum power peaking values obtained from the single objective radial power peaking minimization optimization. At the same time, the MOSA minimum radial power peaking LPs have substantially superior feed fuel enrichments. The lack of continuity in the MOSA tradeoff between radial power peaking and feed fuel enrichment is attributed to the discrete nature of the incore fuel management design problem. Between the two apparent groupings of solutions centered around 4.50 and 4.60 w/o enrichment, no feasible solutions appear to exist.

Presented in Figure 2 is the feed fuel enrichment vs. discharge region average burnup (DRABU) tradeoff from a three objective multiobjective optimization using radial power peaking minimization, DRABU maximization and feed fuel enrichment minimization as primary objectives, as illustrated by the nondominated archives remaining at the end of the optimization. Also shown on Figure 2 are the same results from single objective DRABU maximization and feed fuel enrichment minimization cases. Shown in Figure 3 is the radial power peaking(RPP)/DRABU tradeoff for the same three objective

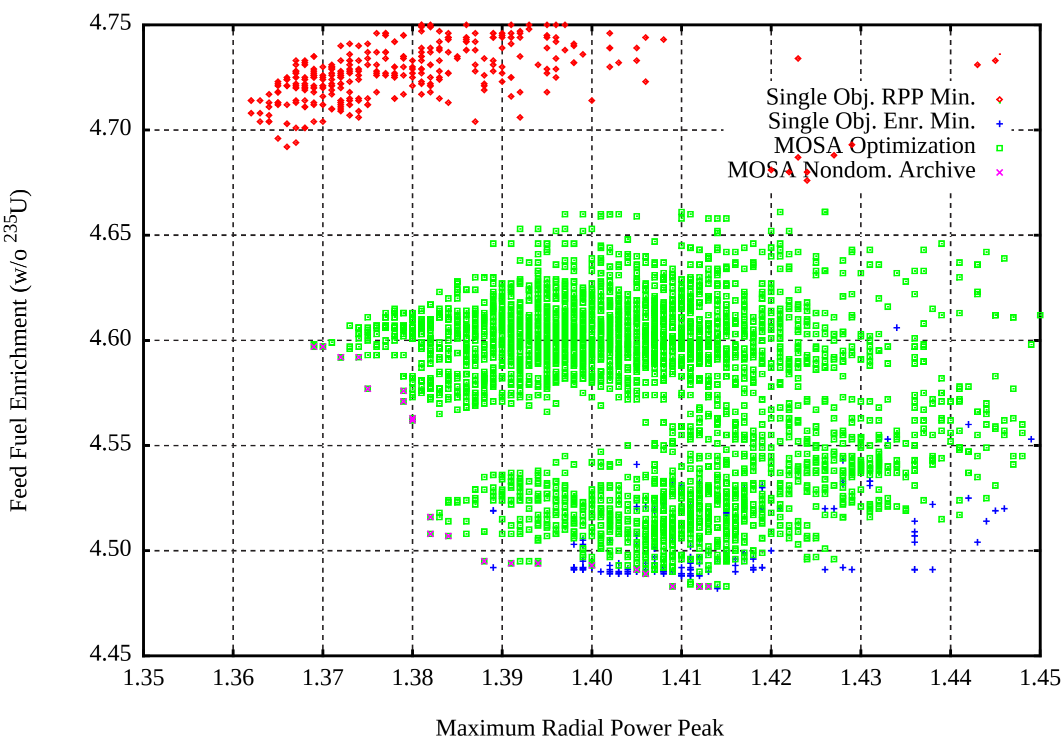

Fig. 1. Comparison Two Objective Multiobjective Results with Single Objective SA.

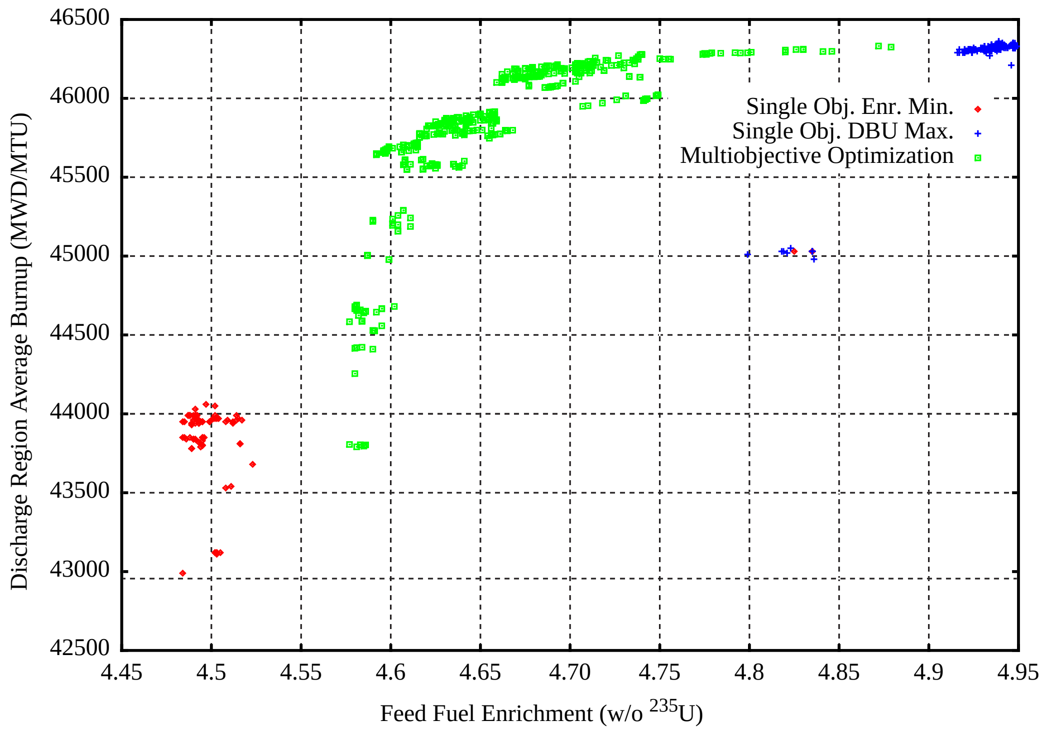

Fig. 2. Feed Fuel Enrichment vs. DRABU Tradeoff for Three Objective Single Objective Optimizations.

MOSA case, and results obtained from single objective RPP minimization and DRABU maximization optimizations. The core design problem used for this case is the same core design, with the same constraints, as was used for the previous result. The MOSA optimization added 2559 LPs to the nondominated archive during the course of the optimization, and removed 2043 LPs, leaving 486 LPs in the final nondominated archive. Once again, the LPs removals were relatively evenly distributed over the final three-fourths of the optimization, indicating that further optimization progress was being made throughout the course of the optimization.

It is important to note that this three objective MOSA case is presented mainly for academic interest and completeness. As discussed earlier, three objective multiobjective optimizations have only very limited practical value, mainly because of the difficulty of characterizing, in a meaningful way, the three way tradeoff between the three objectives. Nonetheless, it may be observed in Figure 2 that the multiobjective optimization was able to duplicate the region average discharge burnup of the single objective DRABU maximization case, and describe one tradeoff surface along lower discharge burnups and lower feed fuel enrichments, and, in Figure 3, another tradeoff surface along lower DRABU and lower radial power peaking factors, even though the multiobjective case was not able to duplicate the minimum feed enrichment of the single objective feed fuel enrichment minimization case.

4. CONCLUSIONS AND RECOMMENDATIONS FOR FUTURE WORK

Improvements have been made to the MOSA algorithm which permit constrained two objective multiobjective optimizations to nearly span the best results obtained from two separate, single objective SA optimizations. The MOSA algorithm is also capable of characterizing, to a limited extent, the three way tradeoff of a three objective multiobjective optimization. A disadvantage of the MOSA methodology is that it is computationally intensive. The two objective multiobjective case presented in this paper required evaluation of approximately 100,000 LPs, while the two single objective optimizations each required evaluation of approximately 29,000 LPs. Nonetheless, the MOSA algorithm likely still possesses a computational advantage over single objective SA in characterizing a tradeoff surface, since four or five single objective optimizations, or more, would probably be required to obtain the same information as provided by the MOSA results.

With the present MOSA time constants of 0.4 and 0.05 for the first and second stage annealing, respectively, the MOSA algorithm must evaluate more LPs than single objective SA optimizations to arrive at a point on the same tradeoff surface when starting from the same reference LP. It does appear that it may be possible to improve the efficiency of the MOSA algorithm by increasing the value of the annealing time constants, and shortening the MOSA search length, and still obtain the same results with fewer LPs sampled. It appears, though, that MOSA is inherently less efficient than single objective SA, and that MOSA will always require the evaluation of somewhat more LPs than single

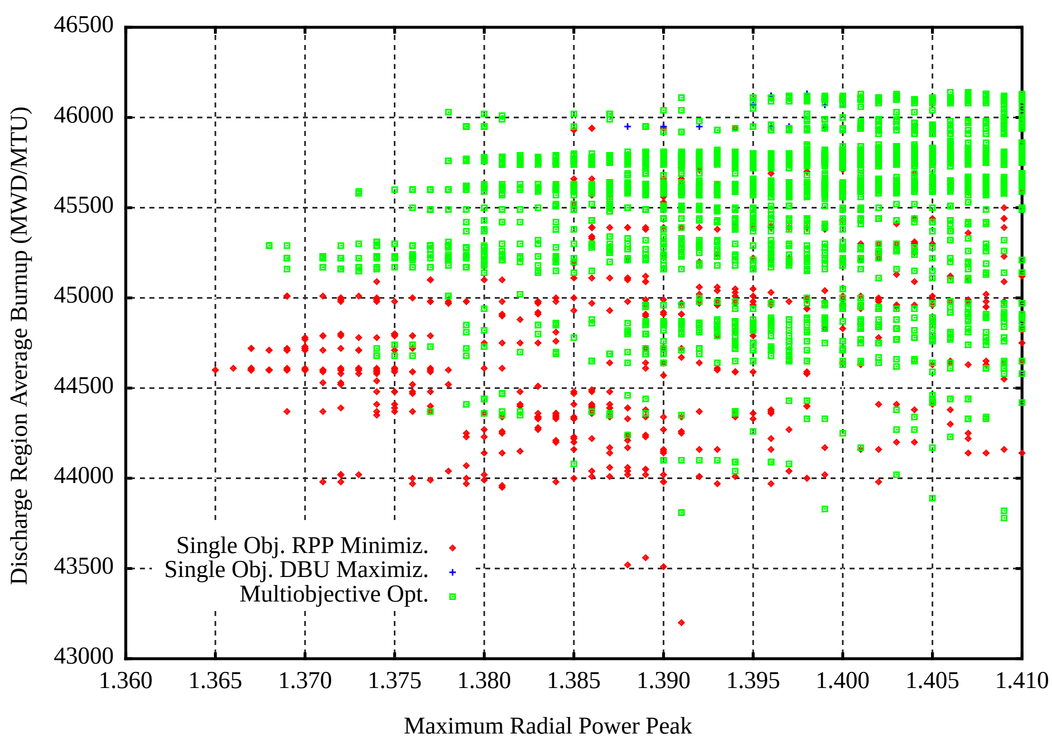

Fig. 3. RPP vs. DRABU Tradeoff for Three Objective and Single Objective Optimizations.

objective SA in order to transition from a reference LP to a point on a tradeoff surface of near-optimal LPs. The reason for this is that it appears that MOSA has a tendency to accept a large number of “lateral” moves--that is, moves which decrease the value of one objective function, while increasing the value of the others, with a resultant overall small improvement in the current solution, and thus acceptance by the SA algorithm. As the number of primary objectives is increased, the search efficiency falls further. As an example, for a two objective MOSA optimization which investigated 106896 LPs, 31663 LPs, (30%) improved both objectives, 39881 LPs (37%) improved one objective, and the remainder, 35352 LPs, (33%) improved neither objective. For a three objective MOSA optimization which investigated 160008 LPs, only 16390 LPs, (10%) improved all three objectives, while 119475 LPs, (75%) improved one or two objectives.

The optimizations presented in this paper used k_∞ ranges did not permit changes to the feed fuel pattern during the course of the optimizations. Although cases which did permit feed fuel pattern changes were performed as a part of this work, it was found that, by the time in the MOSA optimization that the second stage annealing was initiated and all constraint violations eliminated, the feed fuel pattern would become frozen. As a result, any tradeoffs obtained would be about that single feed fuel pattern. This leaves open the possibility that multiobjective genetic algorithms may be able to characterize the tradeoff surface between objectives when the tradeoff surface spans multiple feed fuel LPs.

ACKNOWLEDGMENTS

This work was sponsored by the member organizations of the North Carolina State University Electric Power Research Center (EPRC).

REFERENCES

Engrand, P., 1997. A Multi-Objective *(sic)* Approach Based on Simulated Annealing and its Application to Nuclear Fuel Management, *Proc. 5th Int. Conf. Nucl. Eng.*, Nice, France, ASME.

Huang, M. D., Romeo, F., Sangiovanni-Vincentelli, A., 1986. An Efficient General Cooling Schedule for Simulated Annealing, *Proc. IEEE Int. Conf. Computer Aided Design*, Santa Clara, CA. 381-384.

Keller, P. M., Turinsky, P. J., 1996, FORMOSA-P Multiple Adaptive Penalty Function Methodology, *Trans. Am. Nucl. Soc.*, Washington, DC, **75**, 341.

Kirkpatrick, S., Gerlatt, C. C. Jr., Vecchi, M. P., 1983, Optimization by Simulated Annealing, *Science*, **222**, 671.

Kropaczek, D. J., Turinsky, P. J., 1991, In-core *(sic)* Nuclear Fuel Management Optimization for Pressurized Water Reactors Utilizing Simulated Annealing, *Nucl. Tech.*, **95**, 9-31.

Kropaczek, D. J., Parks, G. T., Maldonado, G. I., Turinsky, P. J., 1994, Application of Simulated Annealing to In-core *(sic)* Nuclear Fuel Management Code FORMOSA-P, *Proc. Int. Conf. on Reactor Phys. and Reactor Comp.*, Tel Aviv, Israel 572-579.

Parks, G. T. 1996, Multiobjective Pressurized Water Reactor Reload Core Design by Nondominated Genetic Algorithm Search, *Nucl. Sci. Eng.*, **124**, 178-187.

Suppapitnarm, A., Parks, G. T., 1999, Multiobjective Optimization of PWR Reload Core Designs Using Simulated Annealing, *Proc. MC'99*, Madrid, Spain.

White, S. R., 1984, Concepts of Scale in Simulated Annealing, *Proc. IEEE Conf. on Computer Design*, Port Chester, IEEE.