APPLICATION OF SIMULATED ANNEALING TO DETERMINING RFSP TARGET EXIT IRRADIATIONS

P.D. Buchan*

Reactor Physics, Pickering NGS, Ontario Power Generation
P.O. Box 160, Pickering, Ontario, CANADA L1V 2R5

C. Stewart

Physics co-op student, University of Waterloo
200 University Avenue West, Waterloo, Ontario, CANADA N2L 3G1

Abstract - A simulated annealing optimization technique was applied to the problem of determining optimal target exit irradiation values for the CANDU† Reactor Fuelling Simulation Program (RFSP-IST). Optimized target exit irradiation values would result in a good match between measured and simulated channel powers, while also achieving a desired core effective multiplication factor. The merits and drawbacks of the method as applied to this problem are discussed.

1.0 Introduction

In order to determine a time-averaged distribution of fuel bundle irradiation and flux in the core, the Reactor Fuelling Simulation Program (RFSP-IST) [1] simulates long term reactor operation with refuelling. This is accomplished by iteratively solving for bundle irradiation, channel axial flux, core flux distribution, and xenon concentration, while conforming to four or eight bundle-shift rules for fuelling. One parameter which affects the calculated axial distribution of bundle irradiations along the length of a channel is the expected bundle exit irradiation for those irradiated bundles removed from the channel at the time of refuelling. This is called the *target exit irradiation* and it must be specified by the RFSP user. The goal in optimizing the target exit irradiations is usually to match simulated channel powers with some desired set of channel powers, and to achieve a preferred core effective multiplication factor ($K_{effective}$). Sometimes other criteria are to be met, such as meeting a specific core-average burnup. At first glance, the problem appears to be quite difficult to solve because changing the target exit irradiation value for one channel affects not only that channel, but channel powers elsewhere in the core, since total reactor core power must be conserved. Furthermore, the value of $K_{effective}$ is dependent upon target exit irradiation distribution.

A variety of methods for finding optimal target exit irradiations are in use, with varying levels of efficiency and effectiveness. A useful method must achieve the necessary acceptance criteria (e.g., match channel powers, achieve desired $K_{effective}$, etc), and use a tolerable amount of computing time in the process. Some methods currently in use include:

* E-mail: david.buchan@opg.com

† CANDU® (CANadian Deuterium Uranium reactor) is a registered trademark of Atomic Energy of Canada Limited.

1) script-based RFSP execution incorporating a root-finding algorithm,
2) manual manipulation of target exit irradiation values with acceptability checking, and
3) generating random target exit irradiations with repeated acceptability checking (hit-and-miss).

The purpose of this study was to investigate the simulated annealing optimization technique as another possible method. At the end, the results with the new technique were compared with those using method 1 above (see Section 7.0).

Note that the acceptability of target exit irradiation values can be evaluated via RFSP execution or by forming a computationally fast, but tolerably accurate, mathematical model to act as a surrogate to actual RFSP execution.

2.0 Target Exit Irradiation Values: An Optimization Problem

The problem at hand involves the adjustment of target exit irradiation values and the monitoring of some acceptance criteria. In this study, the acceptance criteria included achieving desired channel powers and achieving a desired $K_{effective}$. In theory, there should exist target exit irradiation values which result in the best possible values of these parameters (optimal), given the codes and models used. To gauge how optimal a solution might be, the differences between achieved and desired parameter values (in this case, channel powers and $K_{effective}$) are incorporated into a single value – a *metric* of success – which needs to be minimized. The metric can be viewed as a function of the variables to be optimized.

There is a range of traditional optimization methods available, such as Brent's method, the downhill simplex method [4], conjugate gradient methods, Powell's method, etc. In general, these conventional methods suffer from a tendency to occasionally converge on a local minimum in the value of the success metric. Indeed, the initial guess of values can influence the outcome of the optimization process, depending on the proximity to local minima. This problem results directly from:

1) The existence of local minima in the value of the metric function, and
2) An optimization method which always chooses downhill movements.

Here, "downhill movements" are changes in the variables to be optimized that result in a reduced (better) metric value. For complex problems, the shape of the metric function in solution space will not be readily apparent, and the existence of local minima should be assumed.

3.0 The Prototypical Simulated Annealing Method

One stochastic method of overcoming the problems associated with local minima is usually attributed to Metropolis et al. [2]. The problem under their investigation involved finding the probable configuration (equation of state) of a number of interacting particles (molecules) at a given temperature and pressure. The solution they developed allows for finding the configuration which would result in the lowest overall energy state for the system.

In the Metropolis procedure, an initial random configuration of particles is assumed and then a uniform-random displacement (up to some maximum) is applied to each particle and the overall change in energy ΔE of the system is calculated. If $\Delta E < 0$, then the motion was "downhill" and the movement is always accepted. If $\Delta E > 0$, then the movement is accepted with probability $\exp(-\Delta E/kT)$; i.e., a uniform-random number ξ between 0 and 1 is generated, and if $\xi < \exp(-\Delta E/kT)$, then the particle takes the new position. If $\xi > \exp(-\Delta E/kT)$, then the particle returns to its original position. Variable k is Boltzmann's constant, and T is temperature.

The Boltzmann distribution gives the fractional number of particles occupying a set of states which have energy E (as long as the temperature is sufficient to ignore particle interactions at a quantum mechanical level). Thus, the probability of any specific particle in the system having energy E is proportional to $\exp(-E/kT)$, and the particles are said to obey Maxwell-Boltzmann statistics. Applying Metropolis' probabilistic rules to particle movements will result in realistic energy distributions. More importantly, *any* configuration of particles is possible with some probability. The method is thus said to be "ergodic." To find the optimal (lowest potential energy) configuration of particles for the system, the particles are allowed to roam solution space at a series of successively lower temperatures, finally "freezing" at T=0, in a simulated annealing process. It is important to note that the energy of the system must be evaluated repeatedly, and this can become computationally costly depending on the complexity of that calculation.

The crucial innovation to be drawn from the Metropolis procedure is that, like conventional optimization methods, downhill movements are always accepted, but uphill movements are also accepted with some probability determined by the "temperature" of the system. This enhancement can be incorporated into many conventional optimization methods as a way of reducing the likelihood of getting trapped by local minima. Press et al. [3] have applied this enhancement to the traditional downhill simplex method [4] and their method was used for this study.

4.0 Traditional Downhill Simplex Method

The downhill simplex method as envisaged by Nelder and Mead [4] utilizes a geometrical shape called a *simplex* (existing in n-space, where n is the number of variables to be optimized), consisting of n+1 vertices. The position of each vertex is defined by a unique set of coordinates (target exit irradiation values, in our case). The metric can be evaluated at any point in n-space, and the method requires the evaluation of the metric at each vertex following each iteration. The highest (worst) metric value among vertices occurs at vertex denoted P_H. The lowest (best) metric value among vertices occurs at vertex denoted P_L.

Each iteration, sometimes referred to as an *amoeba* iteration because they affect the shape of the simplex, involves taking one of three possible actions: a vertex reflection, a simplex expansion, or a simplex contraction. In each of these three operations, a pseudo-centroid is used (discussed below). A vertex reflection takes place by evaluating the metric at each vertex, and then reflecting the vertex P_H with the highest metric value through the pseudo-centroid. A one-dimensional expansion can be attempted in which the reflected point is moved even further along the same axis as the original reflection in the hope that the downward trend continues. A one-dimensional contraction can be attempted in which the high point P_H is moved toward the pseudo-centroid to a point midway between its current position and the pseudo-centroid. If that new point has a worse metric value, an n-dimensional contraction is performed in which all vertices are brought towards the vertex P_L with the lowest metric value, to midway between their current position and P_L.

The pseudo-centroid position is computed by taking the average of coordinate values over all vertices, *excluding* the vertex P_H which has the highest (worst) value for the metric. Excluding P_H, and the use of an appropriate reflection coefficient, ensures that the simplex volume is conserved upon reflection. The volume of the simplex is conserved for reflections in order to prevent the simplex volume from becoming pathologically long and narrow in some direction of n-space; the simplex becomes degenerate when vertices become coincident.

5.0 Downhill Simplex with Simulated Annealing

Press et al. [3] apply a modified version of the Metropolis technique to the downhill simplex method. In their version, random "temperature" fluctuations are applied to metric values before the acceptability of a new position of the vertices is assessed, as described below.

First, as discussed previously, the metric of success $metric_i$ is evaluated at each vertex P_i of the simplex. Then, *positive* logarithmically distributed random values, proportional to temperature, are *added* to those metric values.

$$metric_i^{\text{with temperature fluctuation}} = metric_i - T \times \log_e(\xi) \tag{1}$$

where, $metric_i$ is the metric value at vertex i,
T is temperature (arbitrary units), and
ξ is a uniform pseudo-random number between 0 and 1.

The maximum metric value $metric_{high}$ of these thermally perturbed values is then found.

Next, a new trial vertex position P_{try} is found via a reflection, expansion, or contraction transformation on the vertex P_{high} associated with $metric_{high}$. The metric value $metric_{try}$ is then calculated for the new trial vertex.

A positive random thermal variation is now *subtracted* from this trial metric value.

$$metric_{try}^{\text{with temperature fluctuation}} = metric_{try} + T \times \log_e(\xi) \tag{2}$$

If $metric_{try}^{\text{with temperature fluctuation}}$ is less than $metric_{high}$, then the new vertex position P_{try} is accepted. In fact, this will be true as long as $metric_{try}$ is less than $metric_{high}$. Thus all downhill movements are accepted. However, if $metric_{try}$ was worse (higher) than $metric_{high}$, the value of $metric_{try}^{\text{with temperature fluctuation}}$ may or may not still be worse than $metric_{high}$. Thus, uphill movements are accepted with probability proportional to $\exp(-\Delta metric/T)$. Note that, unlike the problem solved by Metropolis in his original paper, the metric here does not really have units of energy, and variable T is not a physical temperature, so Boltzmann's constant is not used. As temperature approaches zero, this

method reduces to the pure downhill simplex method and the system effectively "freezes."

An initial temperature is chosen to be high enough to ensure that any local minima can be escaped, and an "annealing schedule" is implemented to control the rate at which temperature is reduced. The annealing schedule allows the simplex to roam about the solution space at a given temperature for a specified number of iterations (by reflection, expansion, or contraction), and then the system is cooled to the next lower temperature and the process repeated. At a given temperature, the simplex will continually reshape itself, effectively sampling the metric value at all points accessible at that temperature. If the temperature is reduced slowly enough, the simplex will likely shrink into the region containing the lowest minimum encountered. Obviously, the choice of initial temperature and annealing schedule can influence the success or failure of the method. In our particular study, each subsequent temperature was 80 % of the previous.

5.1 Problem Simplifications

5.1.1 Grouping Channels into Regions

For Pickering NGS A and B, there are 390 and 380 fuel channels, respectively, and thus 390 and 380 target exit irradiation values to be determined. However, an examination of historical fuelling records indicated that in some areas of the core, de-fuelled bundles from adjacent channels had similar irradiations. Therefore, the core model was divided into regions containing fuel channels with similar exit irradiations. A single value of target exit irradiation was then assigned to each of these regions – five for Pickering B, and six for Pickering A. As a result, the optimization problem was reduced from 380 or 390 variables, to five or six.

Theoretically, at one extreme, each channel could be defined as an individual region, while at the other extreme, all channels could be included in a single region. Mathematically, let us assume there are n regions and thus n target exit irradiation values x_i. Let us define the difference between RFSP-calculated channel power and desired channel power as:

$$y_{channel} = \frac{CP_{RFSP} - CP_{desired}}{CP_{desired}} \times 100\% \tag{3}$$

The average of channel power differences within region i is therefore:

$$\bar{y}_i = \frac{\sum^{\text{over all channels in region i}} y_{channel}}{\#\text{ of channels in region i}} \tag{4}$$

We therefore have n values of $\bar{y}_i$, one for each region.

Let the core effective multiplication factor be $k_{effective}$.

The task here is to simultaneously minimize the values of $|\bar{y}_i|$ by adjusting the x_i values, while achieving a $K_{effective}$ as close to some desired value (typically 1.0) as possible. Note that altering one x_i may affect some or all $\bar{y}_i$ values and $K_{effective}$. The overall cost function to be minimized is a sum of the absolute channel power differences $|\bar{y}_i|$ divided by the number of regions (i.e., an average $\bar{y}_i$), plus the difference between desired and achieved $K_{effective}$. The relative weighting of $K_{effective}$ to average $|\bar{y}_i|$ was usually 1.0, but could be varied.

5.1.2 RFSP Surrogate Model

To alleviate the computational burden associated with repeatedly executing RFSP to update the metric of success for each subsequent set of target exit irradiations, a mathematical model was prepared, which it was hoped, would act as a surrogate to RFSP execution and still give tolerably accurate channel power and $K_{effective}$ values. The strategy adopted here was to execute an initial set of "perturbation" RFSP runs, record the target exit irradiation values used, and extract the resulting channel powers and $K_{effective}$ values. In each of these runs, a single target exit irradiation value was perturbed, while all others were left unchanged. From this data, the relationships between all target exit irradiations and changes in the channel powers of each region as well as changes to $K_{effective}$ could be determined. Several methods were attempted while trying to find a fast, yet acceptably accurate, mathematical fit to this data. By choice, the simplification of assuming that changes in irradiation in one channel only affected that channel's power and $K_{effective}$, was not made, and that option was not explored with annealing in this study. No further RFSP runs were executed during the annealing process, although a final RFSP run was executed after the annealing process in order to confirm the validity of the solution.

5.1.2.1 Surrogate Method 1: Constant Linear Slope

For this surrogate model, a single unique slope was assumed to describe the dependence of channel power in one region or $K_{effective}$, on changes in target exit irradiation in any region.

For example, assuming five regions, $n = 5$, and introducing $y_6 = k_{eff}$, one obtains:

$$\begin{bmatrix} \bar{y}_1 \\ \bar{y}_2 \\ \bar{y}_3 \\ \bar{y}_4 \\ \bar{y}_5 \\ y_6 \end{bmatrix} = \begin{bmatrix} a_{11}x_1 + a_{12}x_2 + a_{13}x_3 + a_{14}x_4 + a_{15}x_5 \\ a_{21}x_1 + a_{22}x_2 + a_{23}x_3 + a_{24}x_4 + a_{25}x_5 \\ a_{31}x_1 + a_{32}x_3 + a_{33}x_3 + a_{34}x_4 + a_{35}x_5 \\ a_{41}x_1 + a_{42}x_2 + a_{43}x_3 + a_{44}x_4 + a_{45}x_5 \\ a_{51}x_1 + a_{52}x_2 + a_{53}x_3 + a_{54}x_4 + a_{55}x_5 \\ a_{61}x_1 + a_{62}x_2 + a_{63}x_3 + a_{64}x_4 + a_{65}x_5 \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \\ b_6 \end{bmatrix} \quad (5)$$

or

$$\begin{bmatrix} \bar{y}_1 \\ \bar{y}_2 \\ \bar{y}_3 \\ \bar{y}_4 \\ \bar{y}_5 \\ y_6 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} \\ a_{31} & a_{32} & a_{33} & a_{34} & a_{35} \\ a_{41} & a_{42} & a_{43} & a_{44} & a_{45} \\ a_{51} & a_{52} & a_{53} & a_{54} & a_{55} \\ a_{61} & a_{62} & a_{63} & a_{64} & a_{65} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \\ b_6 \end{bmatrix} \quad (6)$$

where,

a_{ij} specifies the effect on $\bar{y}_i$ in region i of altering target exit irradiation x_j in region j, and

a_{6j} specifies the effect on y_6 (k_{eff}) of altering target exit irradiation x_j in region j.

This method required the least number of initial RFSP runs. Each target exit irradiation value was perturbed (typically 0.05 n/kb) once from some reference value, so for n regions, only n+1 RFSP runs were required in order to populate the matrix a of equation 6. The reference target exit irradiations were either best guess or previously determined values. This method was expected to be very sensitive to these reference irradiation values. To correct for this, several iterations were performed:

1) Guesses were made for the reference target exit irradiations, slopes computed, and annealing was simulated to find the best irradiation values,

2) The new irradiation values computed in step 1 were used to compute new slopes to be used in another annealing simulation.

Steps 1 and 2 were repeated several times in an effort to converge on the best irradiation values.

5.1.2.2 Surrogate Method 2: Rational Function Fitting

In order to capture channel power and $K_{effective}$ dependencies over the entire range of possible irradiation values, initial RFSP runs were executed for the following irradiations: 0.01, 0.1, 0.15, 0.2, 0.8, 1.4, 2.2, 3.0, 4.5, and 5.9 n/kb, and rational functions were fitted to the channel power and $K_{effective}$ dependency of region *i* on target exit irradiation in any region *j*. Various trials showed that a rational function of the following form was most successful.

$$\bar{y}_i = \frac{p_0 + p_1 x_j + p_2 x_j^2 + p_3 x_j^3 + p_4 x_j^4 + p_5 x_j^5}{1 + q_1 x + q_2 x_j^2 + q_3 x_j^3 + q_4 x_j^4 + q_5 x_j^5 + q_6 x_j^6} \tag{7}$$

where pi and pj are the constants determined during the fit.
Usually rational functions provided an excellent fit to the perturbation RFSP run data, however it was expected to be computationally slower to evaluate during the annealing process.

5.1.2.3 Surrogate Method 3: Piecewise Linear Fitting

Piecewise linear models of region-averaged channel power were then created, based on a set of initial perturbation RFSP runs again for the following irradiations: 0.01, 0.1, 0.15, 0.2, 0.8, 1.4, 2.2, 3.0, 4.5, and 5.9 n/kb.

Specifically, over a small range of Δx, one can express the relationship between change in irradiation Δx and channel power difference $\bar{y}_i$ or $K_{effective}$ as:

$$\bar{y}_i = m_i \Delta x + y_k \tag{8}$$

where, m_i is the slope between point y_{k-1} and y_k,
y_k is the intercept for this particular range, and
Δx is the change in exit irradiation.

This approximation is only valid over small changes in irradiation, so in order to utilize this as a surrogate for repeatedly running RFSP at solution time, the entire range of target exit irradiations was divided into small pieces over which this assumption will be valid. This surrogate model – a set of slopes (m_i) – was developed by performing a set of RFSP perturbation runs in which the target exit irradiations in each region were perturbed ten times while leaving all other target irradiations unaltered, in order to construct an accurate representation of all the interdependencies between adjacent channels. In each perturbation, the resulting region-averaged channel power differences $\bar{y}_i$ and $K_{effective}$, values were recorded. For n regions, this method led to a total of 10*n* initial perturbation RFSP runs to be executed – a significant computational commitment.

6.0 Discussion of Results

This section describes the results of running annealing simulations with each of the three RFSP surrogate models outlined above.

6.1 Surrogate Method 1: Constant Linear Slope

Table 1 reports the results of some constant linear slope annealing simulations.

This surrogate model is computationally fast. However one drawback of the one slope linear surrogate model is that it can't capture the channel power and $K_{effective}$ dependence over the entire possible range of irradiation; typically the slopes a_{ij} are evaluated in the middle of the expected solution range of irradiation, and so behavior at very low irradiation won't be captured.

This method was found to be very sensitive to the assumed values for the reference target exit irradiations used in the initial RFSP runs used to generate the surrogate model. When starting with reference irradiation values about 1 n/kb away from the best solution, repeating the two iterative steps shown in Section 5.1.2.1 usually led to convergence on some unrealistic solution (e.g., some irradiation values were at the chosen maximum of 15 n/b) from which no amount of iteration would lead to escape. If the reference irradiations were close to ideal for five or six regions, the method achieved excellent channel power and $K_{effective}$ results in a short time. The difference in results between the annealing solution and the confirmatory RFSP run were small.

The results reported in Table 1 for the "6regd" case seemed peculiar. This case was expected to produce as good, or better, results than case 6rega, since the only parameter change was a ten-fold increase in the number of amoeba moves at each temperature. But the results for this case, in terms of both channel power differences and $K_{effective}$ were worse. One plausible explanation for this might be that, for case 6regd which had 200,000 amoeba moves per temperature, all amoeba moves made after the initial 20,000 moves would trace out a different path through solution space than that followed in 6rega. Evidently the path followed in 6regd did not pass through the better solution found in 6rega.

When the number of regions was increased from five or six to 380 or 390, the method failed to give reasonable results (very large and/or very small irradiation values, with poor channel power and $K_{effective}$ matches). Some work was done to apply limits to solution space and they were successful in constraining the solution values (irradiations).

6.2 Surrogate Method 2: Rational Function Fitting

More often than not, rational fitting captured the shape of the channel power differences from 0.01 to 5.9 n/kb with great fidelity - something that polynomial fitting was unable to achieve. Ultimately, however, the surrogate model composed of rational functions was unsuccessful because occasionally poles of the rational functions would occur within the irradiation range of interest, leading to computational failures when encountered. When

poles were not present, the rational fits were excellent, and so significant effort was expended attempting to solve this problem, but to no avail.

It appears that rational functions will often give very good fits to a wide variety of data set shapes, but are prone to instability (the presence of poles in the range of interest). Therefore, perhaps in general usage they are best suited to fitting to single sets of data where the results can be graphed and evaluated for goodness of fit. For the purposes here of building a surrogate model to RFSP runs, many rational fits were required and their behavior could not be individually evaluated; given the number of fits, one had to rely on successful *automatic* fitting in every case.

6.3 Surrogate Method 3: Piecewise Linear Fitting

Table 2 reports the results of some piecewise linear annealing simulations.

This surrogate model method was expected to be the most successful at capturing channel power and $K_{effective}$ behavior over a wide range of irradiations while also being computationally fast. However, a comparison of Tables 1 and 2 shows that this surrogate method was not as successful as the constant slope linear method. The channel power differences and $K_{effective}$ were consistently worse when starting from the same reference irradiations as the one slope method, and the computation times were quite a bit longer.

The discrepancy between values reported by the annealing code and the confirmatory RFSP run, are assumed to illustrate the effect of surrogate modelling error. The discrepancies between annealing result and confirmatory RFSP run were consistently worse for the piecewise linear method.

When the number of regions was increased to 380 or 390, the results became unreasonable (very large and/or very small irradiation values, with poor channel power and $K_{effective}$ matches) and they were not improved by increasing the number of amoeba moves to 200,000 per temperature. Our suspicion is that these poor matches are the result of inaccuracies in the surrogate model, leading to a false global minimum in solution space.

7.0 General Observations

By its stochastic nature, the simulated annealing method requires a large number of function evaluations. For this problem, that involved either RFSP executions or RFSP surrogate evaluations. Therefore, there was considerable pressure to find a surrogate that was computationally fast. At the same time however, the surrogate model must be quite accurate or else the annealing method will find a false global minimum; for this problem with large numbers of regions, the difference in depth of local minima and depth of the global minimum may be small.

Presumably, with significantly more computing power, the function evaluations of the surrogate model could be replaced with actual RFSP executions, and this would eliminate

surrogate modelling errors. Reducing the number of amoeba moves per temperature provides significant savings in computation time but increases the risk of missing the global minimum. Similarly, starting from too low an initial temperature can also lead to missing the global minimum. More amoeba moves must be allowed per temperature if more regions are considered (less channels per region). This is because each amoeba move relocates only one of the vertices of the simplex, and ideally, each vertex should be moved a large number of times at each temperature in order to ensure solution space is fully explored.

The assumption that changing irradiation in one channel affects only the power in that channel and $K_{effective}$ (leaving all other channels unaffected) should be explored. However, when channel powers are lumped together in regions (to reduce computations), this simplifying assumption becomes less tenable because changing the target irradiation for a region containing 70 channels is certain to affect other regions of the core.

For the purposes of comparing the efficacy of simulated annealing with other methods, a simple non-stochastic code was written to adjust target exit irradiations in order to match a desired set of channel powers and $K_{effective}$ (method 1 of Section 1.0). This code used Newton's method and repeatedly executed RFSP (i.e., no surrogate model). Each channel's target exit irradiation was adjusted independently. This simple code was able to achieve the desired results (channel power differences less than 1 %, and $K_{effective}$ within 0.1 mk) with about one hundred RFSP executions, demonstrating that complex methods like simulated annealing are not really necessary for the solution of the target exit irradiation problem.

The purpose of applying thermal variations to the pure downhill simplex method was to prevent convergence on local minima. The success of the simple Newton's method code suggests this enhancement is unnecessary for this particular problem. In other words, an optimization method that always takes the downhill movement appears to be sufficient to solve this problem. If this assumption is correct, then using our simulated annealing method with very low temperatures and only one step in the annealing schedule should achieve similar results. Tables 1 and 2 show that starting from too low a temperature causes the results to be much worse, perhaps suggesting that the errors introduced by the surrogate model have introduced sufficient local minima to adversely affect the results.

8.0 Conclusion

The simulated annealing method was shown to be a practical tool for optimizing target exit irradiation values for use in RFSP, but only with a reduced number of variables (target exit irradiation values) to optimize. i.e., with channels lumped into a few regions

The purpose of applying thermal variations to the pure downhill simplex method was to prevent convergence on local minima. The cost of adding this stochastic behavior is the requirement of large numbers of function evaluations. As implemented in the context of optimizing target exit irradiation values for use in the RFSP code, this involved large numbers of RFSP runs or surrogate model evaluations. In the final analysis, the

simulated annealing method proved too computationally demanding for this particular application. Other, much simpler methods are able to achieve the desired results. We expect that other statistical techniques such as genetic methods would similarly demand large numbers of function evaluations.

The success of the simple Newton's method code suggests that the risk of converging on local minima is low for this particular problem. However some tests using nearly pure downhill simplex (very cold simulated annealing) gave poorer results. Our interpretation is that the surrogate model is introducing local minima into solution space.

Regardless, as computing power increases over the years, the simulated annealing method is expected to become more practical and may be an appealing option for a variety of optimization problems, particularly ones that resist simpler methods of solution due to the presence of local minima.

9.0 References

1. B. Rouben, *RFSP-IST, The Industry Standard Tool Computer Program for CANDU Reactor Core Design and Analysis*, Proceedings of the 13th Pacific Basin Nuclear Conference, Shenzhen, China, 2002 October 21-25.

2. N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller, *Equation of State Calculations by Fast Computing Machines*, Journal of Chemical Physics, 1953; 21(6):1087-1092

3. William H. Press, Saul A. Teukolsky, William T. Vetterling, Brian P. Flannery, *Numerical Recipes in Fortran 77*, Second Edition, Cambridge University Press, Cambridge (1992)

4. J.A. Nelder, R. Mead, *A Simplex Method for Function Minimization*, Computer Journal, 1965; 7:308-313

Table 1 Results of Constant Slope Linear Annealing Simulations

Case Name	Metric Value	Maximum CP_{diff} (%)	$K_{effective}$	T_{init}	T_{final}	niter	Lower Limit (n/kb)	Upper Limit (n/kb)	$K_{effective}$ Relative Weight	Number of Function Evaluations	CPU Time
6reg	4.688891958	7.371	0.99959169	100	1.E-04	20,000	0.0001	15	1	1240087	14.4 s
6rega	0.409634775	2.457	1.00000001	1.E+20	1.E-04	20,000	0.0001	15	1	4960446	49.7 s
6regb	4.688891958	7.371	0.99959169	100	1.E-08	20,000	0.0001	15	1	1629715	19.1 s
6regc	0.409657562	2.457	0.99999999	1.E+20	1.E-04	2,000	0.0001	15	1	496451	5.0 s
6regd	5.004876062	16.701	1.00000000	1.E+20	1.E-04	200,000	0.0001	15	1	49600454	8.3 m
6rege	1.266846706	2.293	1.00010724	1.E+10	1.E-04	20,000	0.0001	15	1	2900241	30.1 s
6regf	0.409628683	2.457	1.00000001	1.E+05	1.E-04	20,000	0.0001	15	1	1860133	20.2 s
6regg	10.00867992	21.714	1.00030134	1.E+04	1.E-04	20,000	0.0001	15	1	1660119	16.9 s
390reg	13.14062099	0.094	1.01310783	1.E+20	1.E-04	20,000	0.0001	15	1	4960453	15.6 h
390rega	14.90216705	0.14	1.01485421	1.E+30	1.E-04	20,000	0.0001	15	1	7020666	21.9 h

All runs with same number of regions used slopes generated from same reference values of target exit irradiation. These reference values tended to be quite close to a good solution (usually from previous annealing simulations). Reference values used for Table 1 were the same as for Table 2.

Table 2 Results of Piecewise Linear Annealing Simulations

Case Name	Metric Value	Maximum CP_{diff} (%)	$K_{effective}$	T_{init}	T_{final}	niter	Lower Limit (n/kb)	Upper Limit (n/kb)	$K_{effective}$ Relative Weight	Number of Function Evaluations	CPU Time
6reg	2.51437416	13.974	1.02815333	100	1.E-04	20,000	0.0001	15	1	1240094	31.3 s
6rega	2.52421971	12.949	1.03315014	1.E+20	1.E-04	20,000	0.0001	15	1	4960447	2.0 m
6regb	2.51437416	13.974	1.02815333	100	1.E-08	20,000	0.0001	15	1	1702778	43.7 s
6regc	2.51350049	13.714	1.02985651	1.E+20	1.E-04	2,000	0.0001	15	1	496453	12.1 s
6regd	2.49811116	14.462	1.03322765	1.E+20	1.E-04	200,000	0.0001	15	1	49600454	19.8 m
6rege	2.48449182	14.739	1.02789839	1.E+10	1.E-04	20,000	0.0001	15	1	2900246	72.1 s
6regf	2.52032477	13.15	1.03324008	1.E+05	1.E-04	20,000	0.0001	15	1	1860143	47.1 s
6regg	2.47927494	14.735	1.02336608	1.E+04	1.E-04	20,000	0.0001	15	1	1660122	42.2 s
390reg	4237.09879	10809.9819	3.00173728	1.E+20	1.E-04	20,000	0.0001	15	1	4960466	67.8 h

All runs with same number of regions used slopes generated from same reference values of target exit irradiation. These reference values tended to be quite close to a good solution (usually from previous annealing simulations). Reference values used for Table 2 were the same as for Table 1.