International Conference on Mathematics, Computational Methods & Reactor Physics (M&C 2009)
Saratoga Springs, New York, May 3-7, 2009, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2009)

EVALUATION OF NUMERICAL TECHNIQUES FOR OPTIMIZATION OF ISOCS MODELED DETECTOR MEASUREMENT GEOMETRIES

W. Russ, N. Menaa, D. Nakazawa, A. Bosko, R. Venkataraman, and F. Bronson
Canberra Industries, Inc.
800 Research Parkway, Meriden, CT 06450, USA
wruss@canberra.com

ABSTRACT

Mathematical detector peak efficiency calibrations are often limited in accuracy by uncertain knowledge of source and shielding geometry parameters. Measured spectral peak areas from source emissions of known relative intensity determine the measured relative efficiencies. Absolute calibration accuracy may be improved by optimizing uncertain geometry parameters within specified ranges to best match calculated relative efficiencies to these measured relative efficiencies. The figure of merit (FOM) is a weighted Chi-squared measure of the difference between the measured and calculated relative efficiencies based on Canberra's ISOCS mathematical absolute efficiencies. While this FOM provides a measure of the quality of a given model, a method of efficiently improving the model must also be employed. Several numerical techniques are evaluated in this study for the automated optimization of ISOCS modeled detector measurement geometries. These include an undirected Random Search, Sequential Bracketing, Downhill Simplex, Particle Swarm, Levenberg-Marquardt and Quasi-Newton optimizations. Performance criteria include the speed of convergence and the robustness of the results, with special emphasis on speed given each ISOCS calculation may take several seconds. A variety of geometries, number of free parameters and search ranges were used to test all of the optimization methods to evaluate their performance for this application. The results, in order of average number of FOM calculations to reach a minimum common FOM were: Quasi-Newton (9.5 ± 4.3), Simplex (22.7 ± 11.2), Swarm (43.2 ± 26.9), Sequential (58.1 ± 54.6), and Marquardt (118.5 ± 78.8).

Key Words: ISOCS, modeling, optimization, calibration, efficiency

1. INTRODUCTION

Mathematical efficiency calibrations of radiation measurement systems frequently provide significant advantages over source-based calibrations. It is often impractical to emulate the actual measurement geometry with a physical calibration source, leaving modeling as the only reasonable alternative. Physical calibration sources can also be expensive to obtain, maintain and remove from service. Modeled geometries can be rapidly varied to optimize performance when planning actual measurements.

The quality of results from mathematical efficiency calibrations depends quite strongly on the accuracy of the modeled measurement geometry. This geometry includes an invariant detector model as well as the source and shielding configurations. The detector can be carefully characterized over the full range of energies of interest by optimizing and validating the detector model against precise calibration sources one time. Such a process is performed with Canberra's

In Situ Object Counting System (ISOCS) [1], where the detector is carefully characterized once for later use with a wide set of general geometry templates. Given accurate geometry dimensions, efficiency calibrations can then produce results with uncertainties ranging from 4 to 10 percent depending on energy. However, geometry dimensions are often not well known and the uncertainty budget quickly becomes dominated by the uncertainties in the source and shielding portions of the model. A method of rapidly and automatically optimizing geometry dimensions to yield the most accurate model given the acquired measurement spectra is considered in this study.

2. METHODOLOGY

Some degree of knowledge about measurement geometry is needed to establish an appropriate model template with nominal parameters. In addition, understanding is needed about which geometry parameters are not known to within negligible uncertainty. Those parameters with significant uncertainties must be bounded to reflect the degree of uncertainty. Given such information, it is possible to calculate nominal peak efficiencies with accurately propagated uncertainties and to perform a sensitivity analysis determining which parameter uncertainties have the greatest impact on the result [2]. However, this information can also be used in conjunction with the measured spectrum to determine optimal parameter values that best fit the measured results. When relative emission rates of nuclide energy lines appearing in the spectrum are known, the relative efficiencies are determined according to the corresponding measured peak areas. The most self-consistent peak efficiency calibration yields the same nuclide activity for each peak produced by a given multi-line nuclide. Relative efficiencies calculated using the optimized parameters should provide the best match to the expected relative efficiencies derived from the measured peak areas. For the optimization to be effective, the relative efficiencies must be sensitive to changes in the variable parameters over the energy range represented by the nuclide lines with known relative emission rates. In summary, the prerequisites enabling the optimization of geometry parameters are:

- appropriate geometry model template with nominal parameters
- reasonable bounds specified for uncertain parameters that impact relative efficiencies
- measured peak areas identified as corresponding to known source relative emission rates

A parameter that is not a good candidate for optimization is the source to detector distance. Changing source distance usually does not significantly change differential peak efficiencies.

2.1. Figure of Merit

Since the absolute peak efficiencies are calculated using a non-linear numerical technique (ISOCS), a closed-form analytical solution is not possible for the optimization of the variable geometry parameters. An iterative numerical optimization is required. Trial models created by sampling the bounded variable parameters produce results that must be judged for quality against alternative models. This measure of quality, or figure of merit (FOM), is the metric that drives the automated optimization algorithm. The goal is to minimize the conventional measure of difference between the measured and calculated relative peak efficiencies by measuring the deviation from the ideal of a flat, constant activity across all peaks:

$$\mathrm{FOM} = \frac{\sum_{i}^{peaks} \frac{\left(Activity_i - \overline{Activity}\right)^2}{\sigma^2_{peak\ area,\ i}}}{number\ of\ \ peaks - 1}$$

$$where\ \ Activity_i = \frac{\left(\frac{Peak\ \ Area_i}{Live\ Time \cdot Abundance_i}\right)}{Efficiency_i} \qquad (1)$$

This measure of the consistency of activity as a function of energy, and hence correct relative efficiencies, is summed for each identified nuclide with multiple peaks. In addition to measured spectra, this same FOM calculation can be used to guide an optimization where the goal is defined by a calculated model of known, desired parameters. In such a case, the ideal calculated efficiencies at designated energies replace the measured spectral data ($Peak\ Area_i$ / $Live\ Time$ / $Abundance_i$), yielding a measure of the consistency of ratios of ideal efficiencies to trial efficiencies. This latter case, with known ideal parameters, was only used here to evaluate the optimization methods since such ideal efficiencies are not known in actual applications.

2.2. Candidate Methods

A FOM may indicate the quality of a given model, but it does not directly indicate how to change the free model parameters to optimize the model. A variety of automated numerical optimization techniques are possible. The purpose of this study is to evaluate several promising candidate optimization methods for this particular application. The criteria used to judge the merits of each are speed (number of FOM calculations required) and robustness (consistency and accuracy). The best method will consistently give the lowest FOM in the fewest number of FOM calculations.

The application of optimizing ISOCS geometry parameters includes several constraints. Depending on the complexity and absolute efficiency of geometries, the cost in time to evaluate the FOM at each trial may typically range from a fraction of a second to tens of seconds. Therefore, it is very important to minimize the required number of trial FOM calculations. Given the numerical nature of ISOCS, any method needing derivatives must resort to finite difference calculations, possibly expanding the number of costly FOM calculations. All of the free parameters are explicitly bounded by constant lower and upper limits. Often these bounds are not only reasonable estimates of the likely value ranges, but may also represent physical limitations, such as preventing negative densities. All candidate optimization methods must honor these explicit bounds.

2.2.1. Random

The Random method samples over all of the free model parameters simultaneously using uniform random distributions inside respective bounds (e.g., “throwing darts”). Although simple, the quality of the results depends on the number of trial models calculated, or relative sampling density over the solution space. This undirected approach provides a minimum performance baseline to benchmark the relative performance of the other candidates.

2.2.2. Sequential

The Sequential method [3, 4] treats each free parameter separately, in an iterative sequence of one dimensional grid samples. One free parameter is sampled evenly across the current bounds while the other free parameters are held constant. The best value of the varied parameter is set and then the next free parameter is treated. If the best line sample is not at an edge, parabolic interpolation is used to improve the estimation. This continues until all of the parameters have been optimized. Then the sequence is repeated with bounds that have been reduced and re-centered on the new best values. This reduction in range forcibly sets the rate of convergence, which must be balanced against the risk of falling into a local minimum because of an inadequate search. The Sequential method maintains all sampled points inside the parameter bounds by always maintaining the iteration limits inside the parameter bounds.

2.2.3. Simplex

The Downhill Simplex (Nelder-Mead) method [4] involves continuously improving the FOMs of models represented by points in the solution space at the vertices of a multidimensional form, or simplex. An initial simplex is established with one vertex more than the number of free parameters, and all of these point models are evaluated. The points are sequentially improved by simultaneously adjusting all of the free parameters in the point with the worst FOM. After the worst point is improved and is no longer the worst point, the new worst point is improved. Improvements are performed by reflecting, expanding or contracting the worst point through the centroid of the other points. If none of these three trials improves the worst point to better than the second worst point, all of the points are contracted halfway towards the point with the current best FOM. The Simplex method maintains all sampled points inside the parameter bounds by truncating any parameter values attempting to extend beyond the bounds. The vertices are initialized with one point at the center of each parameter range and the other points randomly located.

2.2.4. Swarm

The Particle Swarm Optimization method [5] involves a group of models or particles simultaneously sampling all of the free parameters within bounds in an iterative manner. Each model remembers the point with the best personal FOM that it has experienced, and all of the models are aware of the global best point any particle has ever experienced. These memories determine where each point should next sample. Each point has a velocity that includes some inertia from the last step as well as stochastically sampled vectors from the current position to the best personal and best global positions. The Swarm method maintains all sampled points inside the parameter bounds by truncating any parameter values attempting to extend beyond the bounds, with a small stochastic offset. The number of particles is set to one more than the number of free dimensions. The particles are initialized with one point at the center of each parameter range and the other points randomly located. Initial velocities are set randomly in proportion to the range of parameter bounds.

2.2.5. Marquardt

The Levenberg-Marquardt method [3, 4, 6] progresses by balancing between steepest descent and inverse Hessian approaches. The trial point starts at the center of each parameter range. The gradient vector and curvature matrix are determined via finite difference (central difference). Each diagonal term of the curvature matrix is multiplied by $1 + \lambda$. This modified curvature matrix is inverted by Gauss-Jordan [3] and is multiplied by the negative gradient to yield the step vector. If the new location improves the FOM, λ is decreased by an order of magnitude and the method continues from the new location, with new finite difference determinations. If the new location fails to improve the FOM, the method stays at the original location and re-calculates a new step vector with λ increased by an order of magnitude. The λ is increased by orders of magnitude and tested until the new location improves the FOM. Decreasing λ approaches no modification of the curvature matrix, approximating purely Newtonian, quadratically convergent steps when in a well-behaved quadratic region sufficiently close to the minimum. Increasing λ approaches steepest descent with smaller and smaller steps in the direction of the negative gradient as a contingency plan when the data indicates the region is not responding to the quadratic expansion approximation, too far from the minimum. Each set of finite difference determinations to calculate the gradient vector and curvature matrix for N free parameters costs $(N + 1)(N + 2)/2$ calculations of the FOM. The Marquardt method maintains all sampled points inside the parameter bounds by truncating any parameter values attempting to extend beyond the bounds.

2.2.6. Quasi-Newton

The Quasi-Newton method [4] is an inverse Hessian approach that progressively determines the inverse curvature matrix using only gradient and location vectors without explicitly determining or inverting the curvature matrix. The trial point starts at the center of each parameter range. The gradient vector is determined via finite difference (forward difference) and the inverse curvature matrix is initialized to the identity matrix. The step vector is determined by multiplying the current inverse curvature matrix by the negative gradient. A scaling factor is applied to the step vector that is either one or a fraction if needed to keep inside all parameter bounds. The FOM is determined at the full step and half step. If neither improves over the original point, the step size is reduced by orders of magnitude until a test point achieves an improved FOM. If the best FOM is a full step, the new step is taken. If the best FOM is a half step, a trial is made at the parabolic interpolation and the best FOM becomes the new location. At the new location, the gradient is determined by finite difference. The old and new location gradients and parameter values are used to directly update the inverse curvature matrix using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) formula [4]. The new step vector is calculated and the iterations progress, with the scaling factor reset to one or the current bounding fraction. Each gradient determination for N free parameters costs $N + 1$ calculations of the FOM. The scaling factor maintains all parameters within the specified limits.

2.3. Test Geometries

Each of the candidate optimization methods was applied to a set of test geometries. The test geometries are intended to represent field measurements where one or more parameters could be significantly uncertain. Using known geometry parameters to calculate ideal target efficiencies

allows the goal parameters to be known *a priori*. However, including a measured spectrum to drive the optimization also provides results for realistic conditions.

2.3.1 Geometry 1

Geometry 1 models an aluminum cylinder viewed by the detector on end, containing uranium oxide with carbon steel and foam shielding, as shown in Figure 1. The optimization methods were tested with parameters varied as shown in Table I, where dashes indicate the parameter was not a free parameter. The tests were conducted both with a measured spectrum, as shown in Figure 2, and with the ideal model used for the FOM calculations.

Table I. Geometry 1 Parameter Settings (minimum / actual / maximum)

Test	U_3O_8 Density (g/cm^3)	U_3O_8 Fill Height (mm)	Al Can Wall Thickness (mm)	Carbon Steel Shield Thickness (mm)
1	0.5 / 2.6 / 3.0	10 / 20.8 / 25	-	-
2	-	-	-	4 / 7 / 15
3	-	10 / 20.8 / 25	-	-
4	0.5 / 2.6 / 3.0	-	-	-
5	0.5 / 2.6 / 3.0	10 / 20.8 / 25	2 / 4.965 / 6	-

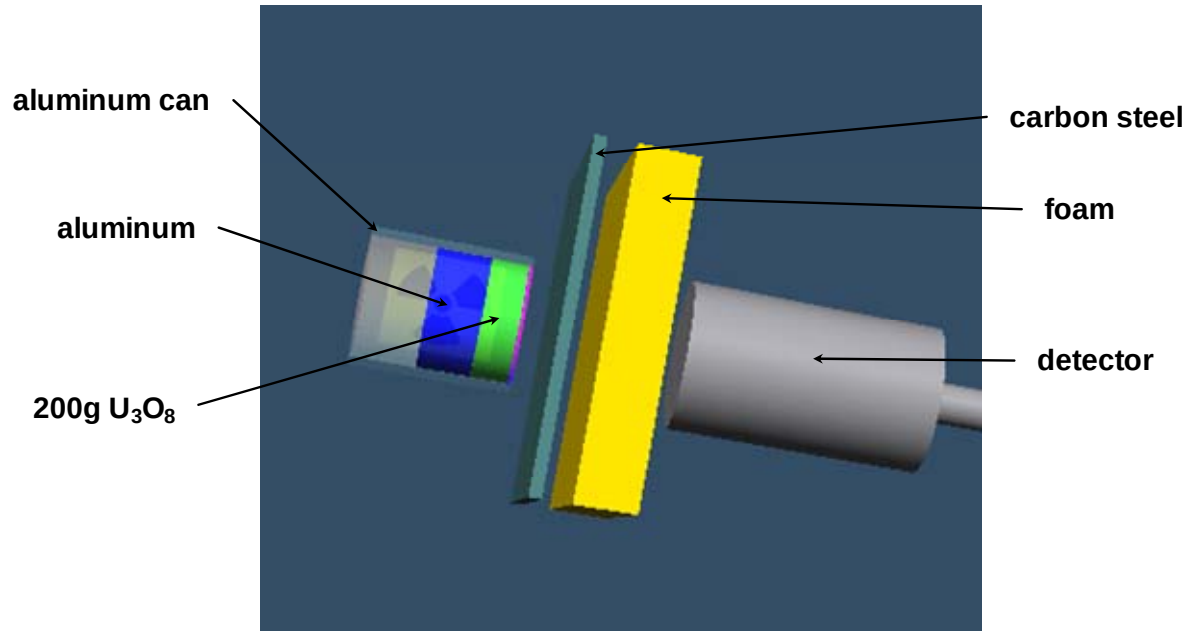

Figure 1. Test geometry 1: aluminum cylinder on end containing uranium oxide with foam and carbon steel shielding. The free parameters include combinations of the source matrix density, source matrix fill height, aluminum container wall thickness, and carbon steel shielding thickness.

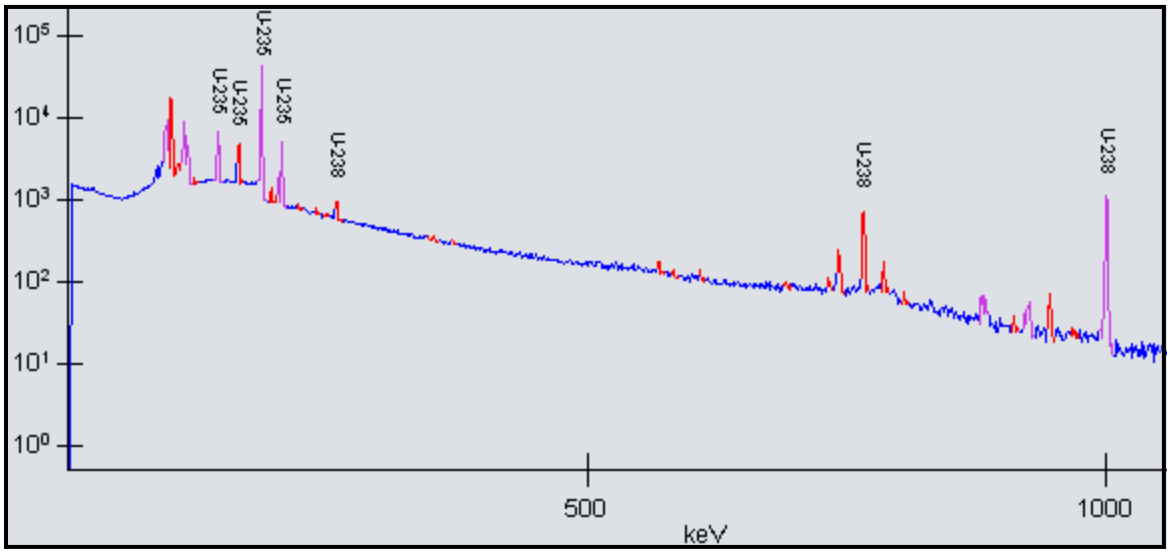

Figure 2. Measured spectrum acquired for geometry 1, with 3 photopeak lines associated with ^{238}U (258, 756, and 1001 keV) and 4 photopeak lines associated with ^{235}U (144, 163, 186, and 205 keV).

2.3.2 Geometry 2

Geometry 2 models a cylindrical polyethylene container of uranium solution in nitric acid being viewed by the detector from the side, as shown in Figure 3. The optimization methods were tested with parameters varied as shown in Table II, where dashes indicate the parameter was not a free parameter. These tests were only conducted with the ideal model used for the FOM calculations. With this geometry, the numbers of parameters were varied and the effects of changing the bounding limits were also explored, as indicated for tests 3, 4 and 5 with changes in the upper limits.

Table II. Geometry 2 Parameter Settings (minimum / actual / maximum)

Test	Uranium Solution Density (g/cm^3)	Uranium Solution Fill Height (mm)	Polyethylene Wall Thickness (mm)	Container Radius (mm)
1	0.5 / 1.192 / 3	80 / 93.4 / 95	0.5 / 2 / 3	50 / 58 /65
2	-	-	0.5 / 2 / 3	-
3	0.5 / 1.192 / 3	-	-	-
4	0.5 / 1.192 / 2	-	-	-
5	0.5 / 1.192 / 10	-	-	-
6	0.5 / 1.192 / 3	80 / 93.4 / 95	-	-
7	0.5 / 1.192 / 3	90 / 93.4 / 96	-	-
8	0.5 / 1.192 / 3	-	0.5 / 2 / 3	-

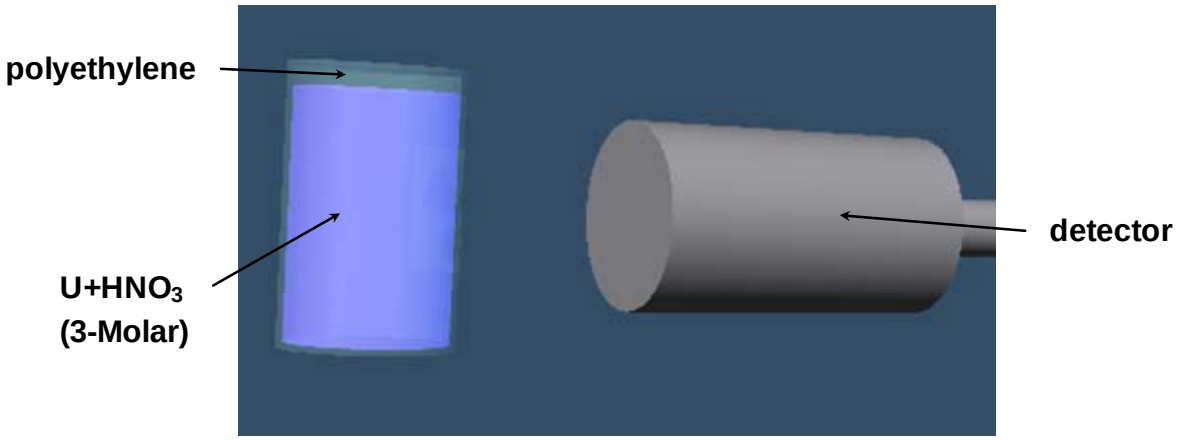

Figure 3. Test geometry 2: polyethylene cylinder from the side filled with a solution of uranium in nitric acid. The free parameters included combinations of source matrix density, source matrix fill height, cylinder radius, and container wall thickness.

3. RESULTS

Examples of optimization progression are shown in Figures 4, 5 and 6, indicating the FOM values associated with each successive FOM calculation. Figure 4 shows the results from modeling geometry 1, test 2 with just the carbon steel shielding thickness as a free parameter. Figure 5 also shows geometry 1, test 2, but for optimization to the measured spectrum data rather than to ideal, modeled data. Figure 6 shows the results from modeling geometry 2, test 1 with four free parameters: uranium solution density, uranium solution fill height, polyethylene container wall thickness, and container radius.

The different methods display very distinct behaviors while optimizing. The Simplex method usually reduces the FOM in very steady, consistently steep slope with evenly spaced steps on a log scale. Swarm also usually progresses consistently but with more spread and shallower slope reflecting the stochastic aspects of the method. The Sequential method also progresses at a constant log slope but with a repeating, saw tooth pattern expected with the converging grid. The Marquardt method shows a more stepped progression, with plateaus indicating times when finite difference steps are being used to develop the curvature matrix or λ is being refined. The Quasi-Newton method also shows smaller plateaus during generation of finite difference gradients and occasional spikes during line search trials.

Table III summarizes results from all of the tests performed. The values indicate the number of FOM calculations required to reach a least common FOM. The least common FOM is the lowest FOM achieved by all methods with acceptable accuracy. Such results are normalized to reflect rates of convergence independent of differences in best FOMs.

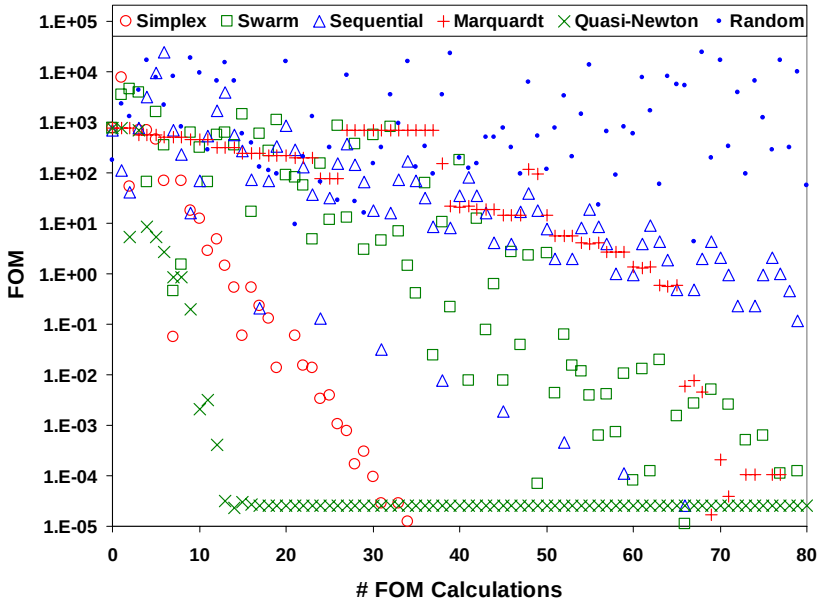

Figure 4. Optimization progression for Geometry 1 Test 2 (modeled), with only one free parameter: steel shielding thickness. The actual thickness was set to 7 mm with the minimum bound at 4 mm and maximum bound at 15 mm.

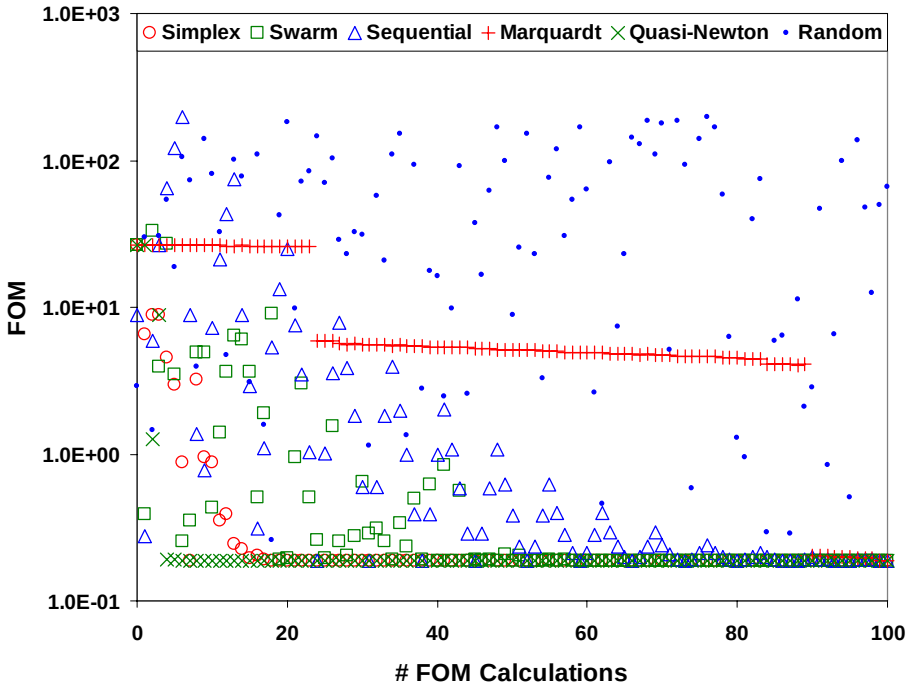

Figure 5. Optimization progression for Geometry 1 Test 2 (measured), with only one free parameter: steel shielding thickness. The actual thickness was set to 7 mm with the minimum bound at 4 mm and maximum bound at 15 mm.

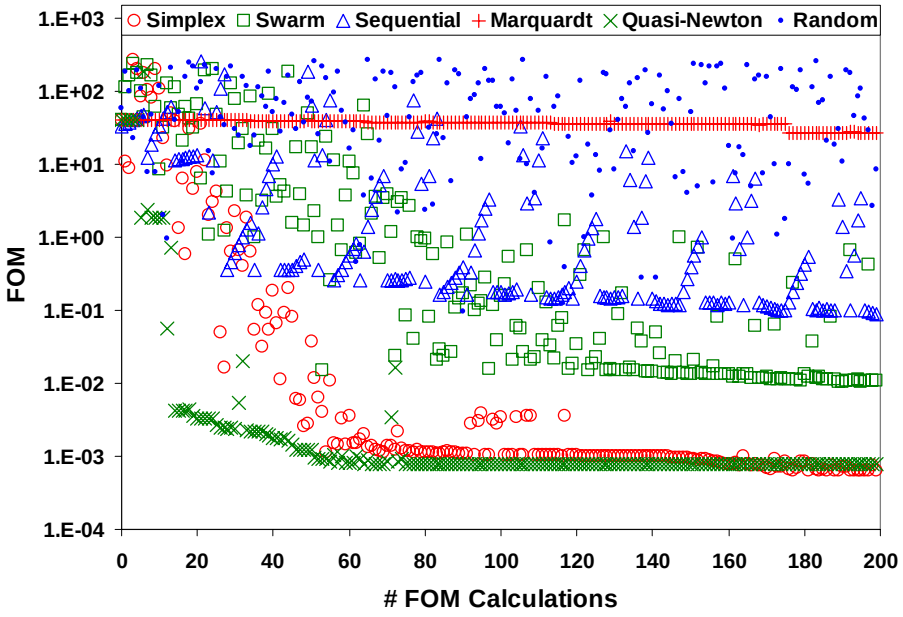

Figure 6. Optimization progression for Geometry 2 Test 1 (modeled), with four free parameters: uranium solution density, uranium solution fill height, polyethylene container wall thickness, and container radius.

Table III. Number of FOM Calculations to Achieve Least Common FOM

Geometry/Test	Sequential	Swarm	Simplex	Marquardt	Quasi-Newton
1/1 (modeled)	34	41	16	77	4
1/1 (measured)	19	3	2	failed (>200)	6
1/2 (modeled)	66	66	31	69	13
1/2 (measured)	23	7	8	91	5
1/3 (modeled)	59	49	29	87	11
1/3 (measured)	5	5	8	146	6
1/4 (modeled)	59	58	30	113	11
1/4 (measured)	12	7	16	223	5
1/5 (modeled)	42	35	27	failed (>200)	5
1/5 (measured)	6	2	7	34	3
2/1	failed (>200)	53	27	failed (>200)	14
2/2	10	76	24	292	14
2/3	179	70	22	95	11
2/4	171	44	23	20	11
2/5	failed (>200)	73	45	failed (>200)	18
2/6	84	70	28	failed (>200)	10
2/7	112	64	33	175	14
2/8	48	55	32	failed (>200)	10
Average ± 1 σ	58.1 ± 54.6	43.2 ± 26.9	22.7 ± 11.2	118.5 ± 78.8	9.5 ± 4.3

4. DISCUSSIONS

The varied mix of optimization methods has yielded a wide range of performance. While it is important to achieve a reasonably accurate solution, the emphasis for this application, with relatively expensive FOM calculations, is primarily on minimizing the number of required calculations. Of course, it is unproductive to rapidly obtain a wrong or inconsistent solution. With a few exceptions, most of the methods were able to reach a minimum level of accuracy for most test trials. Based on the variance of results, the Quasi-Newton, Simplex and Swarm proved to be very consistent, while Marquardt and Sequential were less consistent, sometimes failing to achieve a minimally reasonable solution in an arbitrary number of maximum trials (200).

Interestingly, the candidate method's relative convergence rates varied almost monotonically without much overlap in performance and only occasional swapping of ranks in particular tests. The results suggest that the fastest method is Quasi-Newton, followed in order by Simplex, Swarm, Sequential and Marquardt. The Quasi-Newton method was able to benefit from quadratic convergence with minimal costs in finite difference calculations. Quasi-Newton also used the parameter bounds to more efficiently scale line searches, leveraging additional information about the solution space rather than merely abiding by it. The final convergence of Quasi-Newton may become limited when the optimization step sizes reach an order similar to that of the finite difference step sizes. The finite difference step sizes must be kept above the effective rounding error of ISOCS calculations rather than being limited by machine rounding error. The Simplex method performed very well with a consistently good convergence rate and usually was able to achieve the best final FOM. The Swarm method did not converge as quickly as Simplex because the stochastic aspect of the technique produces a less greedy, more exploratory nature. If the application was more prone to the presence of local minima, the Swarm approach ought to be able to more consistently find the global minimum as it is less dependent on initial conditions.

The poor performance by the Marquardt method was surprising. Marquardt is commonly regarded as the algorithm of choice for optimization problems. However, many applications benefit from the ability to determine analytical derivatives. With four free parameters and several λ optimization trials, the numerical finite difference approach required for this application might need about 20 FOM calculations for what might be considered equivalent to a single "iteration" using analytical derivatives. Ten or more such iterations might be considered acceptable under other circumstances but might require more than 200 expensive FOM calculations for this application.

5. CONCLUSIONS

The results of this evaluation indicate that the Quasi-Newton algorithm should serve as the primary engine when automatically optimizing ISOCS geometry parameters of the type considered. The Simplex and Swarm approaches performed well consistently. The Sequential and Marquardt methods were less robust and were certainly slower.

Further work needed includes refining convergence logic for this application. Although more expensive, a mix of methods might also be evaluated. Combining a limited exploration with Simplex or Swarm to better seed the Quasi-Newton starting point may enhance the robustness of optimization for more complex, general purpose cases than those tested in this study. Finally, future studies should investigate the best ways to handle discrete parameter optimizations. While an exhaustive search of limited discrete parameters (such as a few matrix material choices) might be feasible, it could be more efficient to simultaneously optimize combinations of discrete and continuous parameters explicitly.

REFERENCES

1. "In Situ Gamma Spectroscopy with ISOCS, an In Situ Object Counting System," http://www.canberra.com/pdf/Literature/M2352-InSitu-ISOCS-AN.pdf (2008).
2. F. Bronson, V. Atrashkevich, G. Geurkov and B. Young, "Probabilistic uncertainty estimator for gamma-spectroscopy measurements", *Journal Radioanalytical Nucl. Chemistry*, Vol. 276(3), pp. 589-594 (2008).
3. P. Bevington and D. Robinson, *Data Reduction and Error Analysis for the Physical Sciences*, 2nd ed., pp. 141-147, McGraw-Hill, New York (1992).
4. W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed., Cambridge University Press, (2007).
5. J. Kennedy and R. Eberhart, "Particle swarm optimization," *Proc. IEEE Int'l Conf. on Neural Networks, IV*, Piscataway, NJ, pp. 1942-1948 (1995).
6. K. Bledsoe and J. Favorite, "Using the Marquardt Method for Solutions of Inverse Transport Problems in Two-Dimensional Cylinders," *Trans. Am. Nucl. Soc.*, **98**, pp. 591-593 (2008).