M&C 2017 - International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering, Jeju, Korea, April 16-20, 2017, on USB (2017)

Optimizing HFIR Isotope Production through the Development of a Sensitivity-Informed Target Design Process[1]

Christopher Perfetti*, Susan Hogle, Seth Johnson, Bradley Rearden, and Thomas Evans

Oak Ridge National Laboratory, P.O. Box 2008, Bldg. 5700, Oak Ridge, TN 37831-6170, USA
* perfetticm@ornl.gov

Abstract – *This paper summarizes efforts to improve the efficiency of ^{252}Cf production at Oak Ridge National Laboratory's High Flux Isotope Reactor by using sensitivity analysis to identify potential ^{252}Cf isotope production target design optimizations. The Generalized Perturbation Theory sensitivity coefficient capabilities of the TSUNAMI-3D code within the SCALE Code Package were integrated into the high-performance computing Shift Monte Carlo code in order to obtain sufficiently resolved sensitivity estimates for models containing small concentrations of heavy actinide isotopes. The TSUNAMI-3D sensitivity algorithms were adapted for use in a parallel environment, resulting in a 79% parallel efficiency for simulations using up to 1,000 processors. The potential of several design changes was investigated using the improved TSUNAMI-3D sensitivity analysis tool, including potential changes to the isotope production target density and geometry, and the potential addition of a thin neutron filter material. Several isotope production target design improvements were identified, including a design that featured a lower density target with an indium filter material, resulting in an approximately 1,300% increase in ^{252}Cf production efficiency.*

I. INTRODUCTION

The High Flux Isotope Reactor (HFIR) at Oak Ridge National Laboratory (ORNL) is a valuable national resource for materials irradiation studies and radioisotope production. Scientists designing ^{252}Cf isotope production targets in HFIR facilities must simultaneously consider multiple design objectives, including making efficient use of a limited number of irradiation locations, limiting heat generation in targets, and making efficient use of valuable heavy isotope feedstock. The current heavy curium feedstock that is used for ^{252}Cf production was itself produced at the Sandia River National Laboratory nearly 40 years ago, and about 99% of heavy curium isotopes are lost to fission reactions before they can absorb a sufficient number of neutrons to transmute into ^{252}Cf, as shown in Fig. 1 below.

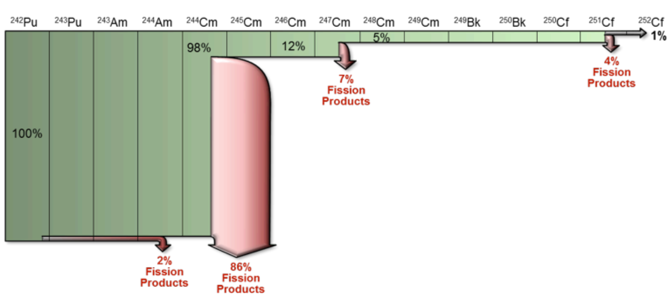

Fig. 1. Heavy actinide loss during ^{252}Cf production. [Ref 1]

There is both a need and the potential to improve the efficiency of ^{252}Cf production at ORNL, and this paper will discuss research and development activities to optimize ^{252}Cf production using sensitivity analysis methods. This paper will begin with an introduction to sensitivity analysis methods, will then discuss their implementation in the massively-parallel Shift Monte Carlo Code, and will then summarize the potential improvements to ^{252}Cf production efficiency that have been identified using the sensitivity methods.

II. GENERALIZED PERTURBATION THEORY (GPT) SENSITIVITY ANALYSIS

The TSUNAMI (Tools for Sensitivity and Uncertainty Analysis Methodology Implementation) capabilities within the SCALE code system make use of sensitivity coefficients for an extensive number of criticality safety applications such as quantifying the data-induced uncertainty in the eigenvalue of critical systems, assessing the neutronic similarity between different systems, quantifying computational biases, and guiding nuclear data adjustment studies [1]. The continuous-energy (CE) TSUNAMI-3D code is a new tool included in SCALE 6.2 that allows for eigenvalue and generalized response sensitivity calculations using high-fidelity CE Monte Carlo methods [2,3]. As shown in Eq. 1, sensitivity coefficients describe the relative change that occurs in a system response, R, due to perturbations or uncertainty in nuclear data parameters (typically cross sections, Σ_x).

$$S_{k,\Sigma_x} = \frac{\delta R / R}{\delta \Sigma_x / \Sigma_x}. \tag{1}$$

CE TSUNAMI-3D contains the Generalized Adjoint Responses in Monte Carlo (GEAR-MC) method, a first-of-its-kind capability for calculating sensitivity coefficients for

[1]This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy.

generalized responses using Generalized Perturbation Theory (GPT) and CE Monte Carlo methods [3]. Rather than computing sensitivity coefficients for the eigenvalue of a system, GEAR-MC calculations compute sensitivity coefficients for the ratio of two reaction rates, R, where

$$R = \frac{\langle \Sigma_1 \phi \rangle}{\langle \Sigma_2 \phi \rangle}. \quad (2)$$

GPT sensitivity analysis has the potential to improve the efficiency of ^{252}Cf production by calculating sensitivity coefficients for ratios of transmutation reaction rates (typically capture-to-fission ratios), which offer insight on what potential design changes can be made to maximize desirable capture reactions and limit heavy actinide destruction through fission reactions.

III. PARALLEL COMPUTING AND GPT SENSITIVITY ANALYSIS

The original CE TSUNAMI-3D GPT sensitivity implementation [3] was shipped with the Beta 4 version of SCALE 6.2 and was completed with the goal of obtaining proof of principle for the new sensitivity capability; the version of this tool that was shipped with the SCALE 6.2 official release includes a number of algorithmic improvements, including significant (typically 60% or more) reductions to the computational memory footprint and simulation runtime, and the ability to compute sensitivity coefficients for multiple reaction rate ratios within a single simulation at the cost of a typically 1–3% increase in memory footprint and runtime per additional response [4].

Sensitivity analysis of systems containing ^{252}Cf production targets may require lengthy simulation runtimes because of the potentially small concentrations of heavy actinides in isotope production targets. In order to obtain sufficiently converged sensitivity tallies in reasonable turnaround times, the sensitivity analysis methods in SCALE 6.2 were extended to parallel Monte Carlo simulations. The parallelization of these sensitivity algorithms was achieved by implementing them in the Shift Monte Carlo code, which has been designed for efficient calculations in a parallel environment [5]. The sensitivity algorithms require tracking a substantial amount of data in order to determine the importance of events that occur during a particle's lifetime, and these algorithms were modified significantly so that they could function efficiently in a parallel environment.

The Iterated Fission Probability methodology used by GPT sensitivity methods requires saving reaction rate information for particles in "chains" of fission events over several generations. This information consists of two pieces of information: 1) a relatively large number of reaction rate tallies (the "Progenitor Tallies"), and 2) a relatively small number of tallies that describe the importance of the Progenitor Tallies (the "Progenitor Importances"). Previously both the Progenitor Tallies and the Progenitor Importances were tied to a particle in a chain of fission events, and were communicated along with the Monte Carlo fission source through several generations of a simulation. The amount of information stored in these tallies often exceeds multiple gigabytes, meaning that parallel simulations (which cannot take advantage of pointers) would require reading, communicating, and writing many gigabytes of information.

These algorithms have been re-written to minimize communication and enable their use in a high-performance computing environment. The large Progenitor Tallies are no longer tied to a given fission chain, but are stored locally on the processor where they originate. Each particle history now only communicates its Progenitor Importances and a unique identifier describing which particle history and node created the corresponding Importances. The creation of Progenitor Importances and Tallies is illustrated in Fig. 2.

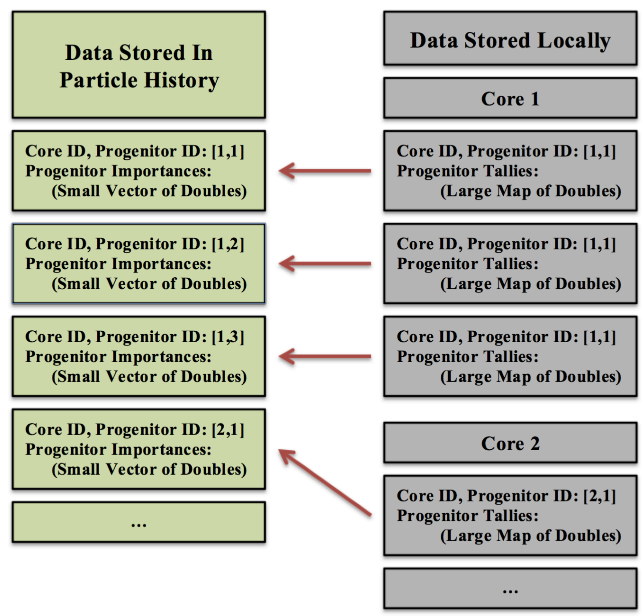

Fig. 2. Progenitor Tallies are stored locally and Progenitor Importance and unique identifier information are communicated to the Master Node.

After several generations the final, "asymptotic" Progenitor Importance is obtained for the Progenitor Tallies in a set of fission chains and is returned to their Progenitor's original processor core by being passed in a "bucket brigade" between neighboring cores. The asymptotic Progenitor Importance is used to weight the stored Progenitor Tallies to produce sensitivity tally estimates.

Fig. 3 shows the results of a weak-scaling study for examining the efficiency of the sensitivity algorithm implementation in Shift. In this study each slave node simulated 500 particles per generation, a value that used the maximum number of particle histories per CPU core given the memory requirements of the GPT sensitivity algorithm's Iterated Fission Probability tallies. The processors on which this simulation took place automatically boost their processors from their ordinary 2.5 GHz speed up to 3.6 GHz when not using all CPUs on a node, resulting in a greater than 100% efficiency for simulations that used less than 32 CPU cores. The parallel efficiency steadily drops for simulations using more than 32 CPU cores, reaching a minimum efficiency of 79% using 1,000 CPU cores. Fig. 4 shows the fraction of compute that was used for various processes during the parallel GPT sensitivity simulations. A large majority of the compute time is spent transporting particle histories and tallying sensitivity coefficient estimates, and a small (but growing) fraction of compute time is used for global sensitivity tally reduction. A very small fraction of time is used for "Response communication," where the Progenitor Importance and unique identifier information is communicated for each particle history, as was described in Figs. 2 and 3.

Although the sensitivity analysis algorithms did not achieve linear scaling, their parallel efficiency was sufficient for the optimization discussed next in this study. Because the batch statistics used by the sensitivity tallies require accumulating first and second moments at each generation, global MPI reductions on large amounts of data are being performed frequently during the simulation. This accounts for the significant and increasing fraction of the compute time as the number of CPU cores increases. Potential exists to improve the parallel scaling of these methods even further, either by performing batch statistics global sums less frequently (perhaps once every 10 generations instead of after every generation), by moving away from batch statistics entirely, or by the "Poor man's parallelism" approach, which involves separating the parallel simulation into some number (30 or more) of repeated simulations, each with a different random seed.

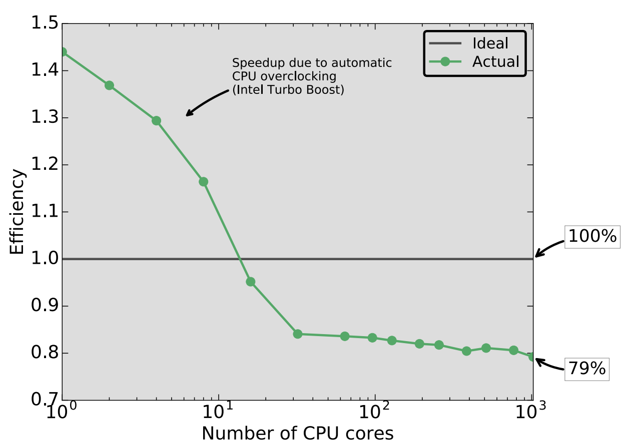

Fig. 3. Efficiency of parallel GPT sensitivity calculations. The greater than 100% efficiency occurs for simulations with less than 32 CPU cores because of automatic CPU overclocking.

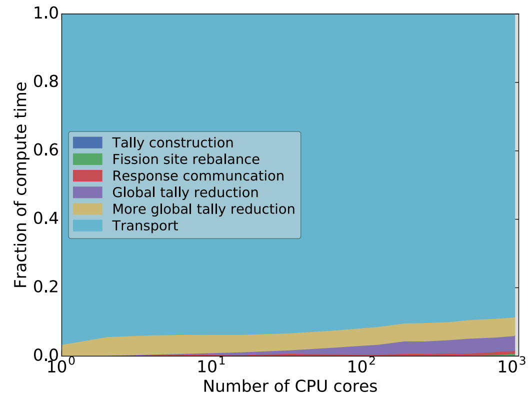

Fig. 4. Compute time used for various processes during parallel GPT sensitivity simulations.

IV. RESULTS OF ^{252}Cf ISOTOPE PRODUCTION TARGET OPTIMIZATION

Having obtained a tool for calculating GPT sensitivity coefficients using parallel computing, the Shift sensitivity analysis tool was used to examine the impact of several potential changes to the design of ^{252}Cf isotope production targets. Each of these simulations used a high-fidelity model of HFIR [6] and required about one full day of runtime. The potential design changes allowed for modifications to the geometry of the ^{252}Cf production targets and the use of a neutron-absorbing filter material for removing neutrons that are likely to cause fission in heavy actinide isotopes.

1. Optimizing the Geometry of ^{252}Cf Production Targets

Although few changes can be made to the HFIR central flux trap in which the ^{252}Cf production targets are placed, it is possible to modify the geometry of the irradiation targets themselves to use either an annular or thin target design. To explore which of these design changes would be optimal, the targets were divided into three equal-area regions (an inner, middle, and outer layer), and sensitivity coefficients were calculated for capture-to-fission ratios in the middle layer with respect to the material density in all three layers. Sensitivity coefficients that are larger (or smaller) for a layer indicate that it is more (or less) important to the transmutation of ^{252}Cf; for example, large, positive sensitivity coefficients in the outer layer would suggest an annular target design.

Table I gives the sensitivity coefficients that were computed for capture-to-fission ratios in the middle layer of the ^{252}Cf targets with respect to the density of the inner, middle, and outer layers. These sensitivity coefficients are unitless and are presented such that a -15% sensitivity implies that 1% increase in the density of a region would cause a 0.15% decrease in the corresponding response. The sensitivity coefficients in Table I are consistently negative, implying that the capture-to-fission ratios can be increased and the efficiency of ^{252}Cf production improved by lowering the heavy actinide number density in any of the three geometry regions. These sensitivity coefficients are larger in magnitude for the outer and middle layers than for the inner layer, which suggests that removing material from the outer and/or middle layers and fabricating a thinner irradiation target would more effectively improve the efficiency of ^{252}Cf isotope production.

Table I. Sensitivity of Heavy Actinide Capture-to-Fission Ratios to the Density of ^{252}Cf Production Targets

Sensitivity of Capture-to-Fission Ratio in the Middle Layer to:			
Isotope	Inner Layer Density Sensitivity	Middle Layer Density Sensitivity	Outer Layer Density Sensitivity
^{244}Cm	-4.21%	-10.79%	-12.42%
^{245}Cm	-0.06%	-0.06%	-0.06%
^{246}Cm	-6.18%	-12.40%	-10.44%
^{247}Cm	-0.18%	-0.24%	-0.19%
^{248}Cm	-7.92%	-12.55%	-10.58%
^{249}Bk	-0.58%	-0.66%	-0.57%
^{250}Cf	-8.44%	-9.63%	-8.53%
^{251}Cf	-0.10%	-0.11%	-0.11%

A possible explanation for the consistently negative sensitivity coefficients is that the heavy actinides in the outermost regions of the targets are over self-shielding the flux at energies corresponding to neutron capture resonances in the targets, and that the neutron flux that causes neutron fissions is not over-shielding (presumably because fission is induced at predominantly faster neutron energies). Lowering the heavy actinide density should decrease the neutron flux depression at the energies corresponding to the location of neutron capture resonances, thereby increasing the capture-to-fission ratios.

Although the calculated sensitivity coefficients suggest that decreasing the amount of heavy actinides in the isotope production targets will increase the heavy actinide capture-to-fission ratios, it should be noted that placing less feedstock material into the isotope production targets is also likely to lower the overall yield of ^{252}Cf from targets (although the transmutation will be more efficient). This effect can be counteracted by placing additional ^{252}Cf production targets within the HFIR flux trap, but these targets will occupy space that might be otherwise used for materials irradiations or other isotope production campaigns. Thus, HFIR scientists may have to decide between having less efficient ^{252}Cf targets, having more efficient targets with less overall ^{252}Cf production, or having more efficient ^{252}Cf production that requires additional irradiation locations in the central flux trap. Thus, any design changes must be weighed by the priorities of the ^{252}Cf production program, which may place more value on producing a certain amount of ^{252}Cf, conserving limited heavy curium feedstock, or using a limited number of irradiation locations in the HFIR flux trap for ^{252}Cf production. Additional flux trap irradiation locations are indeed available to the ^{252}Cf production program, so moving to thin, annular, or lower density irradiation targets is a feasible design change.

2. Selecting an Optimal Neutron Filter

The second potential design change to ^{252}Cf production targets is the placement of a thin filter material around the targets to absorb neutrons that are likely to cause fission in heavy actinides. To explore the viability of different filter materials, an artificial filter was modeled containing a mixture of several potential filter materials and the sensitivity of transmutation reaction rate ratios was determined with respect to the number density of the filter materials. The most promising filter materials would produce positive sensitivity coefficients for desirable reaction rate ratios, which implies that including a full-density filter foil of that material would improve the efficiency of ^{252}Cf production.

Table II gives the sensitivity coefficients that were calculated for several key reaction rate ratios to the presence of several potential filter materials. Rather than simply examining the capture-to-fission ratios (C/F) of all isotopes, this analysis also examined several ratios of capture (cap.) reaction rates that strongly influence the equilibrium concentration of ^{252}Cf. Each of these ratios has a positive impact on ^{252}Cf production and an ideal filter will produce positive sensitivity coefficients for each of these ratios. Identifying an ideal filter material is not simple because a material may (and often does) increase one reaction rate ratio at the expense of another ratio. Therefore, the reaction rate ratio sensitivity coefficients must be weighted by the importance of each ratio to the overall ^{252}Cf production to determine the "Net Sensitivity" of ^{252}Cf production to that material. Fortunately, HFIR scientists have enough experience with ^{252}Cf production to have reasonable estimates for the importance of different reaction rate ratios, as given in Table II.

Table II. Sensitivity of ^{252}Cf transmutation reaction rate ratios to candidate filter materials.

Reaction Rate Ratio	Relative Imp.	Sens. to ^{176}Lu	Sens. to Rh	Sens. to In	Sens. to ^{149}Sm
^{247}Cm C/F	6.76%	-0.11%	-1.06%	-1.16%	-0.23%
^{248}Cm C/F	1.33%	-0.76%	-0.99%	-1.83%	-1.46%
^{251}Cf C/F	9.12%	-0.75%	-0.62%	-1.08%	-5.31%
^{244}Cm cap. / ^{252}Cf cap.	1.70%	2.18%	3.10%	3.45%	2.59%
^{246}Cm cap. / ^{252}Cf cap.	24.49%	3.74%	5.91%	7.51%	5.19%
^{247}Cm cap. / ^{252}Cf cap.	11.27%	0.35%	-6.60%	-7.93%	0.36%
^{248}Cm cap. / ^{252}Cf cap.	29.29%	3.99%	6.00%	7.89%	5.20%
^{251}Cf cap. / ^{252}Cf cap.	14.84%	-6.19%	-11.25%	-16.51%	-26.92%
Net Sensitivity		1.10%	0.61%	0.47%	-1.88%

Of the four potential filter materials, rhodium, indium, and ^{176}Lu produced a positive net sensitivity, indicating that they would likely improve the efficiency of ^{252}Cf production. The ^{149}Sm filter produced a negative net sensitivity coefficient, primarily because of its negative impact on the ^{251}Cf capture-to-fission ratio and the ^{251}Cf / ^{252}Cf capture-to-capture ratio.

3. Effectiveness of the Sensitivity-Informed Design Changes

The effectiveness of the potential design optimizations was evaluated by performing TRITON-3D depletion simulations with the modified ^{252}Cf production targets in the central flux trap of HFIR for three full-power, 30-day irradiation cycles. Each design change was evaluated based on four factors, as shown in Tables VI through VII:

1. The overall yield of ^{252}Cf, (Table IV)
2. The potential ^{252}Cf that was created, (Table V)
3. The potential ^{252}Cf that was destroyed, (Table VI)
4. The efficiency of ^{252}Cf production. (Table VII)

Measuring the "potential" ^{252}Cf created or destroyed gives credit for producing heavy actinides that, although not ^{252}Cf, can be transmuted into ^{252}Cf in future irradiations. Different heavy actinides contribute a different amount of potential ^{252}Cf – for example, ^{251}Cf provides more potential ^{252}Cf than ^{248}Cm. The potential ^{252}Cf present in a sample was determined using conversion factors for each heavy actinide, which describe the fraction of each isotope that would be expected to transmute into ^{252}Cf. These conversion factors have been estimated based on historical yields from previous HFIR ^{252}Cf production campaigns and are given in Table III below.

Table III. Heavy actinide potential ^{252}Cf conversion factors.

Isotope	Potential Californium Factor
^{244}Cm	0.0010
^{245}Cm	0.0033
^{246}Cm	0.0141
^{247}Cm	0.0850
^{248}Cm	0.1800
^{249}Bk	0.3500
^{250}Cf	0.3500
^{251}Cf	0.3500

The efficiency of the ^{252}Cf production was defined as the ratio of the ^{252}Cf yield and the potential ^{252}Cf that was destroyed:

$$^{252}Cf\ Efficiency \equiv \frac{^{252}Cf\ Yield}{Potential\ ^{252}Cf\ Destroyed} \quad (3)$$

The annular and thin target designs each used half the overall mass of heavy actinide feedstock in their targets because of their geometry reductions, and their heavy actinide production results were scaled up by a factor of two for ease of comparison. The lower density design that was investigated used the standard target geometry with 50% of the nominal heavy actinide atom density, and its results were also scaled up by a factor of two.

All of the filtered designs produced lower ^{252}Cf yields (Table IV), but the potential ^{252}Cf produced by these designs (Table V) saw much smaller changes – in most cases, the potential ^{252}Cf increased slightly. These results indicate that they filter materials are slowing down the transmutation of ^{252}Cf because they block some neutrons that would have been captured in the targets. The fact that the potential ^{252}Cf that is destroyed by these designs (Table VI) drops even more significantly than the yields indicates that these filtered designs block more harmful neutrons (i.e. likely to cause fission) than helpful neutrons (i.e. likely to be captured). This observation is reflected in Table VII, where the filtered designs were found to significantly improve the efficiency metric for ^{252}Cf production. The results shown in Table VII should be taken lightly because of their use of approximate potential ^{252}Cf conversion factors in Table III, which means that the potential ^{252}Cf estimates are themselves approximate; furthermore, these efficiency measurement can be deceptive because the efficiency metrics inflate rapidly as the potential ^{252}Cf that is destroyed becomes small (i.e. close to zero). Nonetheless, these results indicate that the potential exists to significantly improve the efficiency of ^{252}Cf production.

Table IV. Yield of ^{252}Cf for potential design changes.

Filter Material	Standard Geometry	Annular Target	Thin Target	Low (50%) Density Target
Unfiltered Target	Baseline	3.8%	13.0%	2.3%
^{176}Lu Filter	-25.2%	-23.7%	-24.4%	-15.3%
Rh Filter	-39.7%	-37.4%	-56.5%	-30.5%
In Filter	-45.8%	-43.5%	-44.3%	-37.4%
^{149}Sm Filter	-58.0%	-56.5%	-58.0%	-51.1%

Table V. Potential yield of ^{252}Cf for potential design changes.

Filter Material	Standard Geometry	Annular Target	Thin Target	Low (50%) Density Target
Unfiltered Target	Baseline	0.0%	0.0%	-0.9%
^{176}Lu Filter	0.7%	0.6%	-0.2%	0.6%
Rh Filter	1.5%	1.5%	1.7%	1.7%
In Filter	1.8%	1.7%	0.9%	1.9%
^{149}Sm Filter	1.7%	1.7%	0.9%	1.8%

Table VI. Potential ^{252}Cf destroyed for potential design changes.

Filter Material	Standard Geometry	Annular Target	Thin Target	Low (50%) Density Target
Unfiltered Target	Baseline	2.1%	0.0%	0.0%
^{176}Lu Filter	-34.0%	-31.9%	-34.0%	-31.9%
Rh Filter	-76.6%	-74.5%	-85.1%	-83.0%
In Filter	-89.4%	-85.1%	-87.2%	-95.7%
^{149}Sm Filter	-87.2%	-85.1%	-87.2%	-93.6%

Table VII. Production efficiency of ^{252}Cf for potential design changes.

Filter Material	Standard Geometry	Annular Target	Thin Target	Low (50%) Density Target
Unfiltered Target	Baseline	0.9%	10.5%	1.3%
^{176}Lu Filter	11.3%	10.5%	11.5%	24.5%
Rh Filter	157.2%	147.4%	181.1%	328.8%
In Filter	375.9%	273.2%	319.4%	1312.5%
^{149}Sm Filter	208.0%	181.1%	229.1%	570.1%

As summarized in Table VII, all of the suggested design changes improved the efficiency of ^{252}Cf production. However, the filter materials that were the most effective for improving ^{252}Cf production efficiency were not the materials that had been predicted to be the most effective in Table II. Furthermore, the addition of ^{149}Sm was expected to reduce the efficiency of ^{252}Cf production, but in resulted in significant efficiency gains in Table VII (although ^{149}Sm did result in the greatest drop in ^{252}Cf yield in Table IV).

Possible explanations for the poorly predicted effects of filter materials include imperfect relative importances in Table II, imperfect potential ^{252}Cf conversion factors in Table III, the infeasibility of using steady-state sensitivity coefficients to optimize a time-dependent design, or the incredible complexity of the ^{252}Cf transmutation chain. The filter material that was predicted to have the greatest positive impact on ^{252}Cf production (^{176}Lu from Table II) resulted in the highest yield of ^{252}Cf among the filter materials in Table IV; likewise, the isotopes with the 2nd, 3rd, and 4th largest sensitivities in Table II also produced the 2nd, 3rd, and 4th highest ^{252}Cf yields in Table IV, respectively. This observation may be coincidental, but it may also suggest that the optimization efforts presented in Table II have optimized the overall ^{252}Cf yield rather than the ^{252}Cf production efficiency.

There are many factors that may be influencing the predictive capabilities of these sensitivity coefficients, and it is difficult to attribute deficient the gap in predictive capabilities to any one factor. For now, these sensitivity methods appear to be more useful for identifying qualitative design changes for optimizing isotope production campaigns. These methods may see improved predictability for isotope production campaigns that are less complicated than the ^{252}Cf campaign, which can require as many as 8 neutron capture events to transmute curium feedstock into ^{252}Cf. The ^{238}Pu production campaign, which requires only one neutron capture, may be a more suitable application for these sensitivity methods.

Overall, the design with half of the nominal actinide number density and an indium filter produced ^{252}Cf most efficiently, resulting in a more than 1300% efficiency increase compared to the standard design. However, this efficiency metric can be deceptive because of the small amount of potential ^{252}Cf that was destroyed. This design may have been more efficient than the standard target, but it also produced a lower overall yield of ^{252}Cf, which highlights a weakness of the efficiency metric and of the reaction rate ratio sensitivity analysis. A design change that decreases both the rate of fission and the rate of capture in the ^{252}Cf targets can produce a positive sensitivity coefficient (and a higher ^{252}Cf efficiency) if it decreases the fission rate more than it decreases the capture rate. If one wishes to avoid the slowed transmutation that occurs in filtered designs, then moving to thin target geometries will most effectively improve the yield of ^{252}Cf and efficiency of ^{252}Cf production. Analysts must prioritize increasing the overall ^{252}Cf yield or conserving limited heavy curium feedstock material.

V. CONCLUSIONS

This paper has documented ongoing research and development activities for using sensitivity analysis to identify potential design optimizations in ^{252}Cf isotope production targets. This sensitivity analysis sought to apply the TSUNAMI-3D GPT reaction rate ratio sensitivity capability to predict how changes to the ^{252}Cf target design would impact ratios of reaction rates (typically capture-to-fission ratios) that are significant to the production of ^{252}Cf. Before performing optimization analysis, the sensitivity methods were implemented in the Shift Monte Carlo code to enable parallel simulations, achieving a parallel efficiency of 79% for a simulation that used 1,000 CPU cores. Next, the sensitivity analysis capability was used to detect the sensitivity of ^{252}Cf production to the geometry of irradiation targets and identified that either an annular, thin, or lower density target would improve the efficiency of ^{252}Cf production. Lastly, the sensitivity capability identified that adding a ^{176}Lu, rhodium, indium, or ^{149}Sm foil filter around the ^{252}Cf production targets would improve production efficiency. When combined, the geometry and filter design changes were found to increase the efficiency of ^{252}Cf production by more than 1300%. Depletion simulations were used to confirm the sensitivity-suggested design changes and it was observed that the reaction rate ratio sensitivity coefficients are more effective at predicting qualitative design improvements rather than quantitative improvements. There is future potential to improve the predictive capability of this sensitivity analysis by calculating sensitivity coefficients for the overall ^{252}Cf yield rather than individual reaction rate ratios.

ACKNOWLEDGEMENTS

This research was sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U. S. Department of Energy. This manuscript has been authored by UT-Battelle LLC under Contract No. DE-AC05-00OR22725 with the US Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

REFERENCES

1. B. T. REARDEN and M. A. JESSEE, Eds., *SCALE Code System*, ORNL/TM-2005/39, Version 6.2, Oak Ridge National Laboratory, Oak Ridge, Tennessee (2016). Available from Radiation Safety Information Computational Center as CCC-834.
2. C. M. PERFETTI and B. T. REARDEN, “SCALE 6.2 Continuous-Energy TSUNAMI-3D Capabilities,” *Proc. ICNC 2015*, Charlotte, North Carolina, USA, September 13-17, 2015.
3. C. M. PERFETTI and B. T. REARDEN, “Continuous-Energy Monte Carlo Methods for Calculating Generalized Response Sensitivities using TSUNAMI-3D,” *Proc. PHYSOR 2014*, Kyoto, Japan, September 28–October 3, 2014, American Nuclear Society (2014).
4. C. M. PERFETTI, B. T. REARDEN, “CE TSUNAMI-3D Algorithm Improvements in SCALE 6.2,” *Trans. Am. Nucl. Soc.*, **114**, 948–951 (2016).
5. T. M. PANDYA, S. R. JOHNSON, G. G. DAVIDSON, T. M. EVANS, S. P. HAMILTON, “Shift: A Massively Parallel Monte Carlo Radiation Transport Package,” *Proc. M&C2015*, Nashville, Tennessee, USA, April 19–23, 2015.
6. R. T. PRIMM, III, N. XOUBI, 2005. Modeling of the High Flux Isotope Reactor Cycle 400, ORNL/TM-2004/251, Oak Ridge National Laboratory.