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

Problem-Dependent ORIGEN Library Compression to Increase Computational Efficiency

S. M. Richards[a], S. E. Skutnik[b]

Department of Nuclear Engineering, University of Tennessee: 1004 Estabrook Road, Knoxville, TN 37996
sricha31@vols.utk.edu[a]
sskutnik@utk.edu[b]

Abstract - *ORIGEN is a highly accurate code to model used fuel isotopic inventory evolution and associated radiation sources. The Oak Ridge Isotope Generation (ORIGEN) code is a depletion model to calculate neutron activation, actinide transmutation, fission product generation, and radiation source terms. This is done through analysis of the full isotopic transition matrix to solve the rate equations that describe the nuclide generation, depletion, and decay processes. This system tracks 2237 isotopes and 54331 transitions, making it the preeminent code for evaluating used nuclear fuel isotopic evolution and production in the world. A next-generation approach to using ORIGEN for accurate source term and depletion models is the direct incorporation of ORIGEN into larger frameworks that rely upon activation and depletion source terms via the modern ORIGEN Application Program Interface (API). However, many problems would benefit most from a reduction in the computational cost of directly performing depletion calculations via ORIGEN. The method to reduce computational cost presented in this paper relies on a lossy compression algorithm to reduce the ORIGEN reactor data library size. This is achieved by a reduction of the nuclides and transitions being tracked by elimination of nuclides and transitions unimportant for the specified problem, whether that is tracking decay heat, dose, or specific isotopic contents for depletion, resulting in a final library size of approximately 300-500 isotopes. These generated problem-specific libraries do not result in a significant sacrifice in accuracy by focusing on the determination of important contributing nuclides in a generalized manner such to enable reliable calculation of source terms without carrying the full transition matrix through the calculation. This method has been shown to potentially reduce run times and memory requirements by an order of magnitude.*

I. INTRODUCTION

ORIGEN (Oak Ridge Isotope GENeration, part of the SCALE code system distributed by Oak Ridge National Laboratory) is a highly accurate code to model and predict used fuel isotopic evolution and associated radiation sources. ORIGEN is a depletion model to calculate neutron activation, actinide transmutation, fission product generation, and radiation source terms. This is done through an analysis of the full isotopic transition matrix that contains the decay and reaction coefficients to solve the rate equations that describe the nuclide generation, depletion, and decay processes. This system tracks 2237 isotopes and 54331 transitions, making it the preeminent code for evaluating used nuclear fuel isotopic evolution and production in the world. However, this system results in a coupled set of linear, homogenous, first-order differential equations with constant coefficients, or in matrix notation

$$\dot{\underset{\sim}{N}} = \underset{\approx}{A}\ \underset{\sim}{N} \qquad (1)$$

where the N vectors represent the nuclide concentrations and the A matrix is the aforementioned transition matrix. This system has a formal solution of a matrix exponential problem that is not directly solvable with matrix inversion due to the transition matrix being a stiff sparse matrix. However, this problem can be solved through a number of methods, such as the newly implemented Chebyshev rational approximation method (CRAM) solver implanted in ORIGEN [1].

Many modern code packages are seeking to directly incorporate an accurate depletion source term calculation directly within the larger analysis workflow, much in the way Next Generation Safeguards Initiative (NGSI) [2], UNF ST&DARDS [3], and CYCLUS nuclear fuel cycle simulator [4] have. Including ORIGEN as part of the core of the process would capture the accurate source term generated into the natural workflow. However, many problems would benefit most from a reduction in the computational cost of incorporating ORIGEN. There are multiple efforts in the development to increase the efficiency of the code. One such strategy, based on reducing the number of tracked nuclides and transitions, has been shown [5] to potentially reduce run times by an order of magnitude, and in the case referenced the run times were reduced by approximately a factor of 20. This is achieved by elimination of nuclides and transitions unimportant for the specified problem whether that is tracking decay heat, dose, or the masses of specific isotopic contents for depletion, thereby reducing the number of nuclides to tracked to approximately 300-500, down from around 2200. However, in this preliminary work was done "by hand", with individual nuclides and transitions being chosen and removed (or reintroduced) based on the personal observations of an individual. Such a labor intensive method is not sustainable do to the significant number of hours

required to reduce each library for any given measure of interest.

In order to make generating these problem-specific libraries quicker, and with a significantly reduced time investment by a user, an automated algorithm was developed. This algorithm can generate problem-specific libraries without a significant sacrifice in accuracy by focusing on the determination of important contributing nuclides in a generalized manner. It does so by considering physical quantities such as decay constants, cross sections, and number of transitions for a given nuclide to determine the most suitable nuclides for removal, following a similar iterative process that the preliminary work employed.

II. THEORY

The method developed relies on restructuring ORIGEN libraries from the existent matrix structure that is used into one physically analogous to the chain of transitions that the burnup of which is being modeled. Nuclides on ORIGEN libraries are designated by an 8-digit identification SIZZZAAA number,

- **S** – Integer ID of sub-library
- **I** – Isomeric state of nuclide
- **ZZZ** – Atomic number
- **AAA** – Mass number

Origen libraries are structured into three sub-libraries based on the designations in Table I. The use of sub-libraries allows for multiple nuclides with the same IZZZAAA to be modeled to distinguish between production pathways. There are many nuclides that would be present in the initial compositions that are also fission product, thus these nuclides are represented in both the "LT" and "FP" sub-libraries.

Table I. Sub-library specification

Integer ID	Short ID	Description	Purpose
1	LT	Activation/light nuclides	Non-actinides that present in materials before irradiation
2	AC	Actinide/heavy nuclides	Heavy nuclides that may fission
3	FP	Fission product nuclides	Nuclides produced by fission

In the physically representative structure, hereafter referred to as the transition system with a basic diagram shown in Figure 1, these sub-libraries are modeled as two separate transition systems. These transition systems are comprised of the "LT" sub-library by itself and the "AC" and "FP" libraries together. Structuring in this manner is possible due to no transitions existing that lead between the "LT" sub-library and the other two; whereas sub-libraries "AC" and "FP" are inherently connected through fission reactions. This structuring also allows for the rebuilt, compressed library to still determine production method of the nuclide despite physically all nuclides with a given IZZZAAA would be indistinguishable at an atomic level.

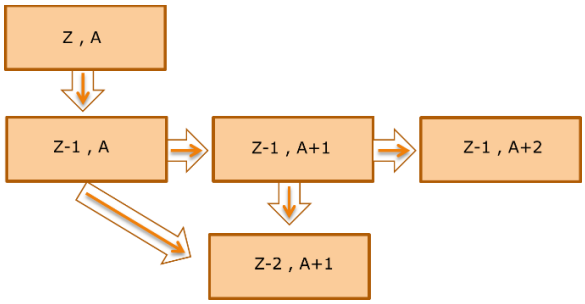

Figure 1. Basic diagram of a transition system for nuclides with atomic numbers given relative to Z and atomic masses given relative to A, with orange arrows representing transitions and the white arrows representing the channels they belong to

The conversion from the ORIGEN library to a transition system involves tracing all sources for all nuclides on the library back to their "parent" nuclides and generating the analogues of these paths in the transition system as nuclides, channels, and transitions. When the nuclide is added to the transition system all "parents" of the nuclide are also added with loss channels that describe the type of transition (i.e. reaction or decay type) and the individual transition that leads to the nuclide being added. In this way, with care being taken to not duplicate nuclides or channels, the entire transition system can be populated. With the library information in this form eliminations can be done by tracing from a removed nuclides parents to its daughters and making appropriate transition changes or removals that lead to and from the nuclide to be removed.

The transition system form explicitly models all transitions, including all fission pathways as well as transitions to and from nuclides. In this way, the system can be reduced to fewer nuclides and transitions while maintaining the systems physicality, as this method places emphasis on a rigorous physical treatment so that no purely artificial transitions or data (decay constants, cross sections, yield, etc.) are created. This is a departure from other methods that perform problem reductions based on numerical correlations in nuclear data [6]. It is because of this physical treatment that the application of weighted metrics can be used to determine appropriate reductions with minimal loss of accuracy in a more generalized and automated method.

1. Metrics

In order to determine the relative importance of a nuclide to a given problem, a problem-specific metric of interest (q) is calculated. The metrics (q) are defined as the sum of the

time integrals of the nuclide-dependent weighting (w_i) multiplied by the number density (n_i) or more explicitly

$$q = \sum_i \int_0^T w_i(t) * n_i(t)\, dt \qquad (2)$$

The weighting functions are then defined based on the quantity of importance for the problem (refer to Table II for specific problem metrics). If the specific contribution to the problem response for nuclide i ($\int_0^T w_i(t) * n_i(t)\, dt$) is below a threshold contribution (typically 1-10pcm depending on the allowable error in the problem), then that nuclide, as well as its creation and removal transitions, is removed from the library. This is referred to as the first stage of creating the simplified library. In the second stage, nuclides that no longer have a means of creation nor any paths for removal (so-called "orphans") are then eliminated from the library as well. In the third and final stage of creating the simplified library, the fission yields are readjusted to preserve total yield. The respected yields are redistributed by giving each daughter of the removed nuclide a yield proportional to the branching ratio of that means of removal; in the case where the daughter is itself a fission product, the yields are treated cumulatively.

Table II. Examples of candidate problem-based weighting factors for the given error to be minimized

Weighting factor	Problem
λ_i	Total Activity
$\frac{\nu\sigma_f}{\sigma_a}$	Criticality
$\int_{E_1}^{E_2} \gamma_i(E) * \lambda_i dE$	Gamma Energies
$\sigma_{a,i}$	Depletion

2. Algorithm

At the beginning of the first stage, ORIGEN is run with the full library of 2237 isotopes and 54331 transitions through the ORIGEN API (full library in figure A.1). This is then used to get the base contribution of each nuclide to the response metric; all individual responses with a contribution below the specified threshold (e.g. 0.01%) are eliminated from the library along with all associated transitions. It is during this stage that the library sees its most significant reduction in size, reducing to fewer than 400 isotopes and 5100 transitions. However, this method of removal has been shown to potentially remove important transition pathways. To correct for this, we employ an iterative approach in which isotopes and their transitions are restored to the library, or potentially further remove them depending on if the desired accuracy or run time are not achieved. This iteration will be explained further later in the paper.

The removal of the orphaned nuclides in the second stage does not contribute as significantly to the simplification of the library as the initial stage, however, it does simplify the library without any further loss of accuracy. To find orphaned nuclides, the algorithm does a full search of the library, checking the number of associated transitions. If the nuclide has no associated transitions, then it is removed from the system. This is due to the nuclide in question no longer having any way of interacting with the system through transitions, meaning that it cannot absorb any particle nor create any new nuclide or particle emission through interaction or decay. Meanwhile, for a nuclide to lack any means of creation results in the nuclide only being present if it was so initially; this is only true for a small number of nuclides (which typically contribute strongly to the problem and thus are unlikely to become orphaned).

The final stage in creating the simplified library is the redistribution of fission yields. In the general 2237-nuclide library, all yields are direct and all decays along mass chains of fission products are modeled explicitly. In the simplified libraries, a combination of cumulative and direct fission yields must be used. The specification for fission yields used in creating the simplified library are direct, cumulative, and "special" fission product yields. For nearly all problems, only the cumulative yield is important and the direct and special yields are for the purpose of weighting certain contributions to the cumulative yield. To redistribute the fission product yield from a removed nuclide, the yield must be redistributed proportionally to the methods of removal of that fission product. In most fission products, removal is dominated by decay; this is even more prominent in the very short-lived nuclides which collectively contribute very little to the problem response metric (and which are first targeted for removal). In this case, the yield is multiplied by the branching ratio for each product and the resulting yield is then added to the cumulative yield for that daughter product.

As stated previously, this is an iterative process. Following the redistribution of the fission product yields, the now simplified library is used in the ORIGEN simulation to obtain a new response metric. This result is compared to that obtained using the full-fidelity library and checked against the given allowed variance (in most cases this variance is limited to 10pcm). In the event that induced error from reducing the library is greater than that which is acceptable, the process begins again from stage one. However, in this and subsequent iterations, the process in the first and second stages are reversed; instead of further removing nuclides and transitions from the library, nuclides with more significant contributions are reintroduced along with all of their transitions. The nuclides are chosen through reducing the

required contribution and searching the original library for nuclides with responses between the original and new contribution to the metric. By reintroducing nuclides, it must be considered that some of the previously-orphaned nuclides are no longer so and thus also need to also be reintroduced. Rather than comparing the original library to the current iteration of the simplified one in search of the orphaned nuclide, the orphaned nuclides are recorded when they are removed. This greatly reduces the time and complexity of this step at minimal cost of storage. Once the nuclides have been restored to the simplified library, the third stage performs the same redistribution as before starting from the original fission product yield in the full library.

III. RESULTS

To test the algorithm outlined in the previous section a simple test problem is used. This test problem is comprised of a single 100-day irradiation followed by a 1-second decay. The library that is being reduced in this problem is one constructed from the ORIGEN default resources. This library contains standard cross section, decay, and yield data for a pressurized water reactor assembly at 0 MWd/MTU burnup. The initial materials for the problem is 1-tonne of 5.0 atom-percent enriched uranium. To reduce the library, the simplest metric, percent of total mass, was chosen. All of the parameters were chosen such that the problem would represent the simplest case that would result in all components of the algorithm being needed for successful compression.

Using the library and the problem defined above, the mass-weighted metric resulted in most of the removals to occur in sub-library 3, fission products. This is due to a majority of fission products being short-lived (i.e. with half-lives less than 0.1 seconds) and having few transitions and low cross-sections. This implies that not only are they unlikely to be present in the final result, their removal will not cause significant propagation of errors. The removals of fission products and their transitions account for approximately 60% of all removals. Another key source of removals are nuclides with low mass inventories that have few loss channels. The fact that these nuclides have low mass inventories but relatively few removal channels indicates that these nuclides are unlikely to be in important mass chains. However, if the low mass inventory nuclides are present in an important mass chain, they are most likely to be located at the end of one of the branches. Due to these nuclides not being found in high transition-density (transitions per nuclides) regions, errors from removals are not propagated to numerous subsequent daughter nuclides. This would also be expected to apply to metastable nuclides because they are predicted to have low transition-density. However, very few metastable nuclides are targeted for removal due to most being present in heavily-weighted mass chains.

To further characterize the tradeoffs in accuracy and runtime, a sensitivity study on removal threshold was done. The sensitivity study related:

- Degree of compression to metric contribution threshold
- Problem runtime to degree of compression
- Compressed library accuracy to degree of compression

The results of this study are presented in Table III. (Note: Aside from the case with the threshold set to 10^{-4}, cases maintain sub-library 1 as it was in the original library, threshold 0 due to no initial masses being present to calculate contribution weights.)

Table III. Results of sensitivity study on threshold for initial removals (***** Case takes advantage of the problem not having any initial material in sub-library 1 to incorporate most nuclides into sub-library 3.)

Threshold	Nuclides/Transitions Remaining	Final Error (% Mass)	Problem Runtime (ms)
0	2237/54331	0	233
10^{-4} *	523/10823	3.83	41
10^{-5}	1083/18785	0.336	80
10^{-6}	1140/20969	0.060	87
10^{-7}	1193/23051	0.057	94

As evident in Table III, significant runtime reduction occurs without drastic reduction in the number of nuclides being tracked while introducing minimal error. Notable results of this study include:

- Despite orders-of-magnitude differences in the metric contribution threshold, difference in reduction were typically under 10% for both nuclide and transitions remaining.
- A factor of two reduction in nuclide inventory and a factor of three reduction in number of transitions results in a factor of four reduction in runtime.
- Despite only reintroducing approximately 5% of the removed nuclides from the 10^{-5} threshold case, the 10^{-6} case saw a significant reduction in error. This indicates that at least one significant mass chain was not reintroduced through the iterative method.

IV. CONCLUSIONS

While it is acknowledged that this current method of reintroducing nuclides based on individual contribution is unsophisticated and can potentially result in the addition of more nuclides than what is strictly necessary when a significant mass chain is omitted, the differences in foreseeable cases are minute and the significant mass chains are expected to be reintroduced after several iterations. Based on previous work performed, as well preliminary

results with the automated method for creating simplified ORIGEN libraries, approximately 10 iterations are needed to meet typical accuracy and efficiency requirements. This additional computational expense at the beginning of a simulation is justified for all large simulations by the order of magnitude reduction in run time by the native ORIGEN API solver. In its current state of development, the automated compression method has been shown to reduce run times by the native ORIGEN CRAM solver from 233ms to 58ms while introducing only 49pcm error to total mass for a single depletion case. This compression was performed with other parameters such as transition coefficients and the number of transitions leading to and from a nuclide are also being considered together with the metric weight. However, in the case of a simple depletion run, a simplified library (created manually by other methods), was shown [5] to reduce run time by the native solver from 237ms to 10ms, a factor of 20 reduction while introducing only 20pcm error in k_{eff}, showing high accuracy can be achieved with even more significant reductions with certain weighting metrics.

Incorporating an accurate source term calculation into the workflow of current and future codes greatly increases model fidelity over "recipe" type models. However, to facilitate ORIGEN being used for these calculations, the memory and run time cost of the code must be reduced. At minimum expense to accuracy of the results, the transition matrix can be reduced in size through iterative removal of nuclides with low contributions to the quantity of interest; resulting in an order of magnitude reduction in nuclides, transitions tracked, and run time.

V. FUTURE WORK

Ongoing development of the algorithm is focused on the identification and reintroduction of important mass chains and optimization of threshold selection. The current proposed solutions to these problems involves consideration of nuclide transition density and total nuclide transition coefficient at the time of nuclide removal and reintroduction selection. It is currently believed that there is a relationship between optimal threshold and total nuclide transition coefficient. Additionally, it is believed that the transition density can be used as an indicator of nuclide importance in a mass chain. These theories will be tested in future editions of the algorithm.

ACKNOWLEDGMENTS

The authors would like to acknowledge the contributions of William Wieselquist and Shane Hart for their assistance and advice in the development of the compression algorithm as well as their work development the prerequisite tools and utilities to make it possible.

This work is supported by the Consortium for Nonproliferation Enabling Capabilities (CNEC) fellowship funded by the National Nuclear Security Administration (NNSA). A portion of this work was also supported by a Nuclear Energy University Programs (NEUP) grant sponsored by the U.S. Department of Energy, Office of Nuclear Energy, award number DE-NE0000737

APPENDIX A: ALGORITHM FLOW DIAGRAM

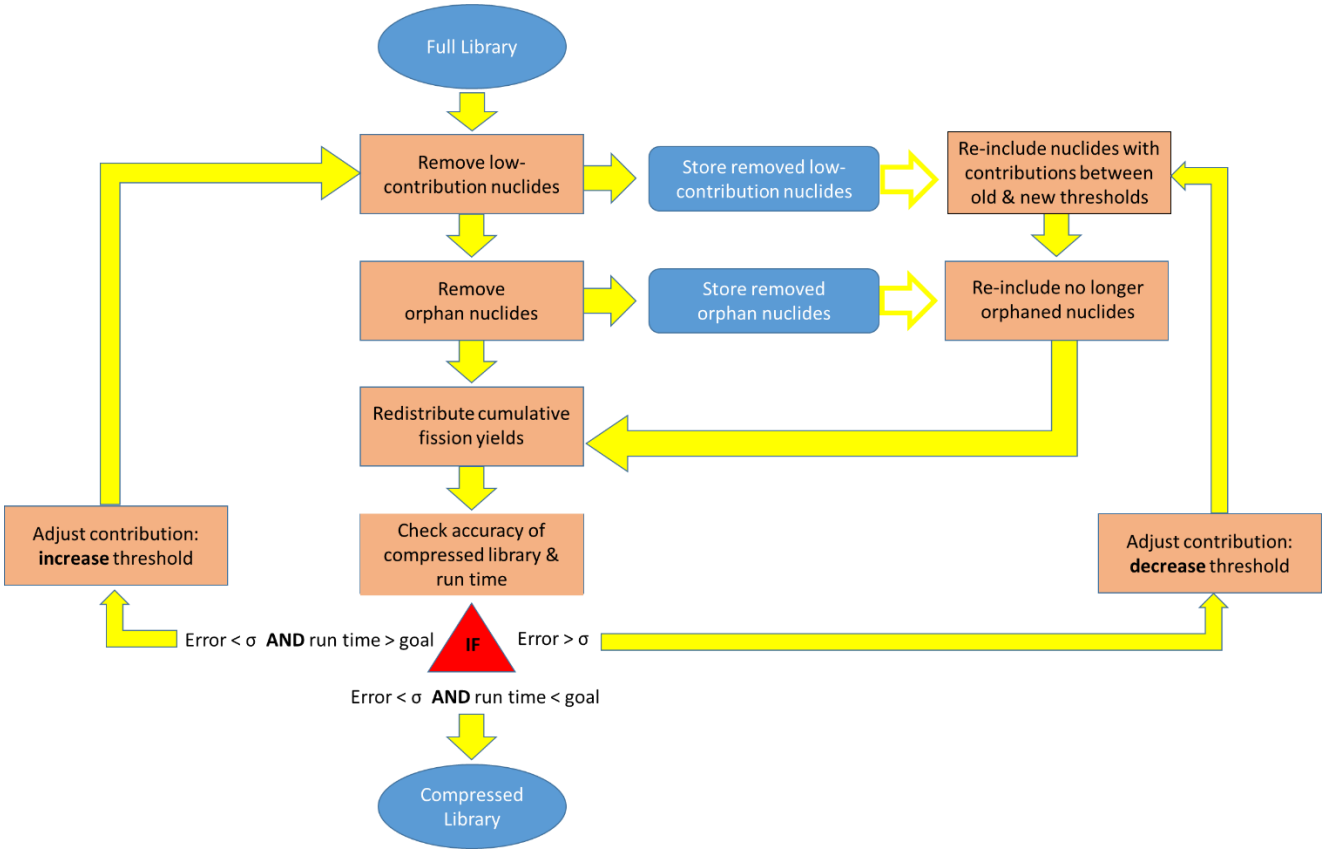

Figure A.2. Flow diagram outlining algorithm steps described in the *Algorithm* section

REFERENCES

1. Isotalo, A.E., Aarnio, P.A., "Comparison of depletion algorithms for large systems of nuclides", Annals of Nuclear Energy, 38 (2-3), pp. 261-268. (2011)

2. S. Skutnik, I.C. Gauld, C.E. Romano, and H.R. Trellue, "Creating NDA 'Working Standards' through High-Fidelity Spent Fuel Modeling," Proceedings of the Institute of Nuclear Materials Management, 53rd Annual Meeting, July (2012)

3. J. Scaglione, K. Banerjee, K. Robb, R. LeFebvre, "The Used Nuclear Fuel Storage, Transportation, and Disposal Analysis Resource and Data System", Proceedings of the Institute of Nuclear Materials Management, 55th Annual Meeting, Atlanta, GA, July (2014)

4. S. E. Skutnik, N. C. Sly, and J. Littell, "An ORIGEN-Based Reactor Analysis Capability for Cyclus," Transactions of the American Nuclear Society, 115, pp. 299-301, Nov. (2016)

5. W. Wieselquist, K.S. Kim, Simplified Burnup Chain Development, CASL PHI Ticket #3432. (2014)

6. S. Zhang, H. Abdel-Khalik, "The Reducibility of Reactor Physics Depletion and Decay Calculation," Reactor Analysis Methods — IV, pp. 1293 – 1296 November (2016)

7. "SCALE: A Modular Code System for Performing Standardized Computer Analyses for Licensing Evaluations", ORNL/TM-2005/39, Version 6, Vols. I–III, January (2009)