Progress in Nuclear Energy xxx (2017) 1–15

Contents lists available at ScienceDirect

Progress in Nuclear Energy

journal homepage: www.elsevier.com/locate/pnucene

Practical techniques for large-scale Monte Carlo reactor depletion calculatons

David P. Griesheimer [a, *], David C. Carpenter [a], Mark H. Stedry [b]

[a] Bettis Atomic Power Laboratory, West Mifflin, PA, USA
[b] Knolls Atomic Power Laboratory, Niskayuna, New York, USA

ARTICLE INFO

Article history:
Received 13 September 2016
Received in revised form
9 January 2017
Accepted 19 May 2017
Available online xxx

Keywords:
Monte Carlo
Reactor
Depletion
Feedback

ABSTRACT

This work presents a brief review of numerical methods for nuclide depletion calculations and a summary of several practical techniques for improving computational speed and reducing memory usage in large-scale Monte Carlo reactor depletion calculations. The techniques covered in the paper include: 1) the use of data hierarchy, 2) separation of absorbing and non-absorbing (precursor) nuclides, 3) optimizations for a backward differentiation formula (BDF) numerical solver, 4) the use of simplified (reduced-order) depletion systems, and 5) the use of a residual fission product absorption correction term to account for the cumulative reactivity effect of nuclides that are not explicitly depleted. In addition, the paper describes several implementation and data management strategies used in the MC21 code, which have proven beneficial for large depletion calculations. A description of these various techniques and strategies are presented along with results from scaling studies and representative reactor depletion calculations that demonstrate the effectiveness of these methods. The results from these studies suggest that large-scale MC depletion calculations including tens- to hundreds-of-millions of depletable material compositions are practical on contemporary mid-range computing clusters.

1. Introduction

Over the past several decades there has been sustained interest in performing large-scale reactor depletion calculations driven by Monte Carlo radiation transport. This interest has motivated significant research and development efforts in this area, including the investigation of improved coupling/time-discretization schemes (Isotalo and Aarnio, 2011a, 2011b; Isotalo and Sahlberg, 2015; Isotalo, 2015; Kotlyar and Shwageraus, 2013), development of more efficient depletion solvers (Isotalo and Sahlberg, 2015; Isotalo, 2015; Kotlyar and Shwageraus, 2013; Pusa, 2011, 2016; Pusa and Leppänen, 2010, 2012; Hykes and Ferrer, 2013), and studies of numerical stability (Dufek et al., 2013; Dufek and Hoogenboom, 2009; Densmore et al., 2013) and propagation of statistical uncertainty (Park et al., 2011; Newell and Sanders, 2015) in multi-timestep depletion calculations. This research, coupled with increases in available computing power, has made full-core Monte Carlo depletion calculations feasible with the current generation of solver codes and computer hardware. Presently, several major Monte Carlo transport codes, such as MC21 (Griesheimer et al., 2015a), TRIPOLI (Brun et al., 2015), MVP (Nagaya et al., 2015), Serpent (Leppänen et al., 2015), and Shift (Davidson et al., 2016), include integrated depletion solvers and have been used for high-fidelity depletion calculations of full-core commercial reactor benchmark problems.

To date, most research in the area of Monte Carlo depletion has focused on the details of solving the Bateman depletion equations. Much of this work, in turn, is built upon decades of research and experience with depletion calculations driven by deterministic diffusion or transport solvers, which has been extended to account for the statistical uncertainty inherent in the MC estimates for local flux and/or reaction rates. This previous work, for both deterministic and stochastic transport solvers, has established that there are many different time-discretization strategies and matrix-exponential solvers that can be used for reactor depletion calculations (Isotalo and Sahlberg, 2015; Isotalo, 2015; Hykes and Ferrer, 2013). However, high-fidelity full-core reactor depletion calculations remain challenging due to sheer problem size, regardless of the solution strategy employed. For example, depletion calculations for commercial light water reactors typically subdivide each fuel

* Corresponding author. P.O. Box 79, West Mifflin, PA 15122, USA.
E-mail addresses: david.griesheimer@unnpp.gov (D.P. Griesheimer), david.carpenter@unnpp.gov (D.C. Carpenter), mark.stedry@unnpp.gov (M.H. Stedry).

http://dx.doi.org/10.1016/j.pnucene.2017.05.018
0149-1970/

element into between 50 and 4000 depletable compositions (assuming axial segments between 1 and 10 cm in height and 1 to 10 radial sub-regions per element). Thus, a large light water reactor (LWR) with 65,000 fuel elements may include over 250 million depletable compositions.

In this paper, we describe a collection of practical techniques for reducing both the computational cost and memory requirements for large-scale MC reactor depletion calculations – while still preserving accuracy in the calculated system reactivities and flux distributions through life. Together, these techniques enable accurate, high-fidelity MC reactor depletion calculations on contemporary computing hardware, while minimizing or eliminating the need for complex data-partitioning methods, such as domain decomposition. In addition, we describe several practical implementation and data management strategies used in the MC21 code, which have proven beneficial for large depletion calculations. It should be noted that, while many of the techniques presented in this paper were originally motivated by the development of the in-line depletion capability in MC21, the techniques are generally applicable to many types of coupled transport-depletion solvers, including systems that use deterministic rather than stochastic transport.

2. Numerical methods for solving depletion equations

The problem of computing the time-dependent change in nuclide concentrations for a single material composition is represented in its most general form as a system of linear ordinary differential equations (ODEs) with variable coefficients,

$$\dot{\mathbf{N}}(t) = \mathbf{A}(t)\mathbf{N}(t), \quad \mathbf{N}(t_0) = \mathbf{N}_0, \tag{1}$$

where $\mathbf{N}(t)$ is a vector of length n containing the concentrations for each nuclide, $\dot{\mathbf{N}}(t)$ is the time derivative (d/dt) of the vector $\mathbf{N}(t)$, $\mathbf{N}_0$ is a vector of initial atom densities, and $\mathbf{A}(t)$ is an $n \times n$ matrix containing coupling coefficients between the nuclides. Typically, systems of depletion equations are stiff, meaning that the coupling coefficients vary over orders of magnitude. The time scales of interest can also vary over many orders of magnitude, typically ranging from hours to years for in-core reactor burnup and shutdown calculations, to millions of years for spent fuel disposal studies.

The rate of change of concentration for a given nuclide as a result of undergoing neutron transmutation and radioactive decay is governed by the differential equation,

$$\frac{dN_i}{dt} = P_i + \sum_j \gamma_j^x N_j(t) - \left(\lambda_i + A_i^a\right) N_i(t), \tag{2}$$

where

P_i = production due to direct fission $= \sum_k y_i^k S^k$, where y_i^k is the yield fraction for nuclide i from the k-th fissile nuclide undergoing fission at a rate S^k,
$\lambda_i N_i$ = loss due to radioactive decay of nuclide i, where λ_i is the decay constant for nuclide i,
$A_i^a N_i$ = loss due to neutron absorption in nuclide i, where the absorption rate is defined as $A_i^a = \int_0^\infty \sigma_a^i(E)\,\phi(E)\,dE$, $\sigma_a^i(E)$ is the total absorption cross section for nuclide i at energy E, and $\phi(E)$ is the scalar neutron flux at that same energy.
$\gamma_j^x N_j$ = production due to transmutation from precursor j, where $\gamma_j^x = \beta_j \lambda_j$ or $\beta_j A_j^x$ (depending on the type of coupling from the precursor), β_j is the branching fraction from the precursor, A_j^x is (analogous to) the absorption term where x denotes a specific type of neutron reaction, including (n,n′), (n,2n), (n,3n), (n,γ), (n,α), (n,p), (n,d) and (n,t).

One key difference among solution techniques for the Bateman equations is the approximation used for the time-dependence coefficient matrix $\mathbf{A}(t)$. A comparison between the various approximations is presented in the following subsections.

2.1. Constant rate solution methods

The most basic approximation for the time dependence of $\mathbf{A}(t)$ is to assume that the coefficient matrix is constant over the timestep. This assumption implies that the microscopic (i.e., density normalized) nuclide reaction rates are constant over the depletion timestep, and is sometimes referred to as a constant reaction rate, constant flux (Carpenter, 2009), or constant extrapolation (CE) assumption (Isotalo and Sahlberg, 2015). Under this assumption, Eq. (1) simplifies to

$$\dot{\mathbf{N}}(t) = \mathbf{A}\mathbf{N}(t), \quad \mathbf{N}(t_0) = \mathbf{N}_0, \tag{3}$$

which has the analytical solution

$$\mathbf{N}(t) = \mathbf{N}(0)e^{\mathbf{A}t}. \tag{4}$$

As noted in a 2003 review paper (Moler and Van Loan, 2003), a variety of different numerical techniques exist for solving the matrix exponential equation shown in Eq. (4). However, many of the general approaches for solving matrix exponential systems are not able to handle stiff systems, such as typically found in nuclide depletion calculations. In the past, many depletion solvers such as ORIGEN2.2 (Croff, 1980) and VESTA (Haeck et al., 2012, 2016) have simplified depletion relationships by removing short-lived nuclides in order to reduce the stiffness of the depletion systems. After the reduced system has been solved, the densities of the omitted short-lived nuclides are determined via an equilibrium calculation. In recent years, the Chebyshev Rational Approximation Method (CRAM) has become an especially popular matrix exponential solver in nuclide depletion solvers (Pusa, 2016, 2011; Pusa and Leppänen, 2012, 2010), due to its ability to handle stiff systems of equations without the need for simplifying approximations.

The primary benefit of the constant coefficient matrix (i.e., constant reaction rate) assumption is simplicity. Matrix exponential solution methods are readily available and relatively straightforward to implement. Furthermore, the change in number densities during a timestep can be computed from a single call to the matrix exponential solver. However, the constant coefficient assumption does not account for the fact that the nuclide reaction rates (and hence the depletion coupling coefficients) are continuously changing due to shifts in the spatial and energy distribution of the neutron flux within each depletion region. This redistribution of flux over time is driven by changes in the relative concentrations of nuclides due to burnup and transmutation, as well as changes in reactor operating conditions (e.g., changes in control rod position, coolant flow rate, power/temperature changes, etc.). As a consequence, single-shot constant-coefficient depletion methods that use only one matrix exponential solution per timestep are typically limited to extremely short timestep sizes in order to ensure stability (Densmore et al., 2013) and achieve reasonable accuracy in calculated end-of-timestep nuclide concentrations. In a 2009 study, Carpenter reported that timestep lengths $\leq$3 h were required to reduce the relative error in depleted ^{155}Gd and ^{157}Gd concentrations at the end of a 700-day reactor operating cycle to 0.1% and 0.3%, respectively, when using a constant coefficient approximation

Please cite this article in press as: Griesheimer, D.P., et al., Practical techniques for large-scale Monte Carlo reactor depletion calculatons, Progress in Nuclear Energy (2017), http://dx.doi.org/10.1016/j.pnucene.2017.05.018