NUCLEAR SCIENCE AND ENGINEERING: **164**, 140–150 (2010)

Computing the Matrix Exponential in Burnup Calculations

Maria Pusa* and Jaakko Leppänen

VTT Technical Research Centre of Finland, P.O. Box 1000, FI-02044 VTT, Finland

Received February 10, 2009
Accepted July 13, 2009

Abstract—*The topic of this paper is the computation of the matrix exponential in the context of burnup equations. The established matrix exponential methods are introduced briefly. The eigenvalues of the burnup matrix are important in choosing the matrix exponential method, and their characterization is considered. Based on the characteristics of the burnup matrix, the Chebyshev rational approximation method (CRAM) and its interpretation as a numeric contour integral are discussed in detail. The introduced matrix exponential methods are applied to two test cases representing an infinite pressurized water reactor pin-cell lattice, and the numerical results are presented. The results suggest that CRAM is capable of providing a robust and accurate solution to the burnup equations with a very short computation time.*

I. INTRODUCTION

The neutronic properties of a reactor fuel are strongly dependent on the isotopic compositions of the fissile materials. The changes in the material compositions must be taken into account in all reactor physics calculations. This is in practice handled by burnup calculation codes. An essential part of a burnup calculation is the solving of the burnup equations that describe the rates by which the concentrations of the various nuclides change. The burnup equations form a system of first-order linear differential equations that can be written

$$\frac{dN_j}{dt} = \sum_{i \neq j} \lambda_{ij} N_i - \lambda_j N_j \ , \quad N_j(0) = N_0 \ , \quad j = 1, \ldots, n \ , \tag{1}$$

where

N_j = concentration of nuclide j

n = total number of nuclides

λ_{ij} = coefficients characterizing the rates of neutron-induced reactions and spontaneous radioactive decay.

*E-mail: Maria.Pusa@vtt.fi

In this paper we consider the burnup system under the assumption that these coefficients are fixed constants. The burnup equations can then be written in matrix notation as

$$\boldsymbol{n}' = \boldsymbol{A}\boldsymbol{n} \ , \quad \boldsymbol{n}(0) = \boldsymbol{n}_0 \ , \tag{2}$$

where

$\boldsymbol{n}(t) \in \mathbb{R}^n$ = nuclide concentration vector

$\boldsymbol{A} \in \mathbb{R}^{n \times n}$ = burnup matrix containing the decay and transmutation coefficients of the nuclides under consideration.

Equation (2) can be formally solved by the matrix exponential method yielding the simple solution

$$\boldsymbol{n}(t) = e^{\boldsymbol{A}t}\boldsymbol{n}_0 \ , \tag{3}$$

where the exponential of the matrix $\boldsymbol{A}t$ is defined as the power series expression

$$e^{\boldsymbol{A}t} = \sum_{k=0}^{\infty} \frac{1}{k!} (\boldsymbol{A}t)^k \ , \tag{4}$$

with the additional definition $\boldsymbol{A}^0 = \boldsymbol{I}$.

There are numerous algorithms for computing the matrix exponential, but many of them are computationally expensive or of dubious numerical quality.[1] Because the decay constants and reaction rates of the nuclides vary extensively, the burnup matrix has a wide spectrum

of eigenvalues, making the approximation of the matrix exponential more difficult. Short-lived nuclides are especially problematic because they can induce eigenvalues of arbitrarily large magnitude. These difficulties have traditionally been solved by using simplified burnup chains or by treating the most short-lived nuclides separately when computing a matrix exponential solution. The selection of a suitable matrix exponential method depends substantially on the characteristics of the problem at hand. For example, the norm and eigenvalue spectrum of the burnup matrix as well as the length of the time step are the key aspects that should be taken into consideration when choosing the matrix exponential method. However, notably little interest and research effort have been shown toward this topic.

The focus of our study was to examine if it is possible to solve a detailed burnup system containing thousands of nuclides by a single matrix exponential method. The motivation for this was the development of the burnup calculation routines in the PSG2/Serpent Monte Carlo reactor physics code.[2] The current burnup calculation implementation in Serpent is based on the TTA method,[3] in which the complicated transmutation chains are resolved into a set of linear subchains that can be solved analytically. The main advantage of this method is that it can handle the extensive variations in the transmutation and decay coefficients and is relatively easy to implement in its basic form using a recursive loop. The most significant problem with the TTA method is that the computation time can easily become excessive if all chains are followed until a stable nuclide is encountered, and cutoffs have to be enforced to terminate insignificant chains. In addition, the current implementation of the method cannot treat chains that form a closed cycle, but the trajectory is terminated if the same nuclide is encountered twice in a single chain.

II. EIGENVALUES OF THE BURNUP MATRIX

In solving the burnup equations with the matrix exponential method, it is beneficial to estimate the character of the matrix eigenvalues, e.g., whether they are real-valued or complex-valued, and, in the latter case, the magnitude of the eigenvalues' imaginary parts.

II.A. Real Parts of Eigenvalues

It is known that the general solution of system (2) is a linear combination of functions of the form

$$t^k e^{\alpha t}\cos(\omega t)\boldsymbol{a}\ ,\quad t^l e^{\alpha t}\sin(\omega t)\boldsymbol{b}\ ,\quad \boldsymbol{a},\boldsymbol{b}\in\mathbb{R}^n\ , \tag{5}$$

where $\lambda = \alpha + i\omega$ runs through all the eigenvalues of $\boldsymbol{A}$ with $\omega \geq 0$ and $k, l \leq m(\lambda) - 1$, where $m(\lambda)$ denotes the algebraic multiplicity of eigenvalue λ (for proof, see Ref. 4). If all eigenvalues of the burnup matrix are real, the concentration of each nuclide is a linear combination of functions of the form $f(t) = t^k e^{\alpha t}$. In this case the eigenvalue determines the rate of exponential growth or decay of the function f. On the other hand, an eigenvalue with a nonzero imaginary part ω indicates that the solution has an oscillating component with period $T = 2\pi/\omega$.

Some understanding of the burnup eigenvalues can be gained by considering the physical constraints related to system (2). For example, it is evident that the concentration of each nuclide must remain bounded at all times. The following theorem (Ref. 4, p. 165) therefore gives a useful characterization of the real parts of the burnup eigenvalues.

Theorem: Every solution $\boldsymbol{n}$ of system (2) remains bounded as $t \to \infty$ if and only if the following hold:

(i) $\mathrm{Re}(\lambda) \leq 0\ \forall\ \lambda \in \Lambda(\boldsymbol{A})$;

(ii) Every $\lambda \in \Lambda(\boldsymbol{A})$ with $\mathrm{Re}(\lambda) = 0$ is a semisimple eigenvalue; i.e., the geometric and algebraic multiplicities agree.

Here, $\Lambda(\boldsymbol{A})$ denotes the set of the eigenvalues of $\boldsymbol{A}$.

The real parts of the eigenvalues of the burnup matrix must therefore all be nonpositive. A purely imaginary eigenvalue would correspond to a nondamped oscillation, which is physically unrealistic in the context of burnup calculation. It can thus be deduced that the real parts of the nonzero eigenvalues of the burnup matrix are always negative.

II.B. Imaginary Parts of Eigenvalues

The characterization of the imaginary parts of the burnup eigenvalues is more difficult. If the burnup chain does not contain any closed cycles—i.e., no paths from any vertex back to itself exist in the burnup matrix—the matrix can be permuted into a triangular form. In this case the eigenvalues are the diagonal elements, and hence, all are real-valued and negative. The nonreal eigenvalues result from closed transition cycles occurring in the burnup chain. However, not all closed transition cycles induce nonreal eigenvalues, and in practice only a fraction of the eigenvalues of the burnup matrix have nonzero imaginary parts.

A suitable mathematical method for establishing a link between the structure and eigenvalues of a matrix is the computation of the strongly connected components of the graph of the matrix.[5] A strongly connected component is defined as a set of vertices such that there exists a path from each vertex to every other vertex. If all of the strongly connected components of a matrix are sorted topologically, the corresponding systems of differential equations can be solved independently in this order. The different cyclic components of a burnup chain can therefore be studied conveniently by calculating its strongly connected components. If the burnup matrix does

not contain any closed cycles, the size of every strongly connected component is one, and the solution of system (2) can be calculated by solving n ordinary linear differential equations. The nonreal eigenvalues can therefore be identified with certain cyclic parts in the burnup transition chain.

We have computed the eigenvalues for a wide range of burnup matrices, and based on our experiments, it seems that they are generally confined to a region near the negative real axis. Based on our observations it appears that a prerequisite for a nonreal eigenvalue is that the majority of the reactions involved in a closed cycle have transition rates that are of the same order. In this scenario the slowest reactions appear to have the most significance for the period of the oscillation. This seems reasonable from a physical standpoint, as well. The cycle that is most likely to induce oscillations appears to consist of an alpha decay followed by successive (n,γ) and β^- reactions. An example of this kind of loop is the transition cycle resulting from the alpha decay of ^{242}Cm. The decay constant of this reaction is of order 10^{-8} 1/s (half-life 162 days), and in a thermal reactor operating at full power, the corresponding cycle typically induces three pairs of complex eigenvalues with imaginary parts of order $\leq 10^{-8}$.

The shortest half-lives encountered in reactor calculations are generally of the order of milliseconds, although there are some even more short-lived nuclides. The half-lives corresponding to neutron-induced reactions are considerably longer. In a thermal reactor operating at full power, most of the transmutation coefficients are of order $\leq 10^{-8}$ 1/s, and they are even smaller in a fast reactor. Therefore, it can be expected that the imaginary parts of the burnup eigenvalues are at most of this order. For every burnup matrix that we have considered, this has also been the case. When the power level is decreased, the transmutation coefficients become smaller. In this case the absolute values of the imaginary parts of the eigenvalues decrease as well. It seems that the oscillations are most likely to occur for reduced power cases where the greatest transmutation coefficients are of order $\leq 10^{-12}$. In general, the eigenvalues of the burnup matrix appear to remain bounded near the negative real axis in all conceivable burnup calculation cases. This observation is exploited in the construction of the matrix exponential method, whose framework is considered in Sec. IV.

III. ESTABLISHED METHODS

III.A. Approximation near Origin

The most obvious approach is to calculate the exponential directly from the definition (4) using a truncated Taylor series. This approximation is naturally most accurate near the origin, so it is ill-suited for burnup calculations, where the matrix norm $\|\boldsymbol{A}t\|$ can become arbitrarily large. In some cases even increasing the number of terms does not improve the approximation because of the accuracy limitations in the computer arithmetics.[1] The accuracy of the series method can be improved by using the method of scaling and squaring, which is based on the identity

$$e^{\boldsymbol{A}t} = (e^{\boldsymbol{A}t/m})^m \quad , \tag{6}$$

where m can be taken as a power of two, $m = 2^k$, so that the norm $\|\boldsymbol{A}/m\|$ becomes sufficiently small. The truncated series is then calculated for the scaled matrix, and the result is squared by repeated multiplications. The accuracy of this technique may be compromised, if the elements of $e^{\boldsymbol{A}t}$ grow before they decay, as t increases. Numerical problems are faced when this so-called "hump" is located between t/m and t (Ref. 1). The series method with scaling and squaring is implemented in the ORIGEN code[6] by excluding short-lived nuclides from the burnup matrix and treating them separately.

The most well-established method for calculating the matrix exponential is probably the rational Padé approximation with scaling and squaring. For example, the matrix exponential function `expm` in MATLAB is based on this approach. Although this method generally outperforms the truncated Taylor series approach, it shares the requirement of $\|\boldsymbol{A}t\|$ remaining relatively small.[1] Accordingly, numerical problems are faced when $\|\boldsymbol{A}\| \gg 1$ and $t \sim 10^6$ s, both of which are plausible values in the context of burnup calculation.

III.B. Krylov Subspace Approach

Various Krylov subspace algorithms are currently very popular, and they have also been recently applied to burnup calculations.[7] In this framework, the original large and sparse matrix $\boldsymbol{A}$ is projected to a lower-dimensional Krylov subspace, and the matrix exponential is then calculated using the series method or the Padé approximation. The projection can be carried out with the well-known Arnoldi iteration, which results in m iteration steps to the partial Hessenberg reduction

$$\boldsymbol{A}\boldsymbol{Q}_m = \boldsymbol{Q}_m\boldsymbol{H}_m + h_{m+1,m}\boldsymbol{q}_{m+1}\boldsymbol{e}_m^T \quad , \tag{7}$$

where $\boldsymbol{Q}_m \in \mathbb{R}^{n\times m}$ is orthogonal, $\boldsymbol{H}_m \in \mathbb{R}^{m\times m}$ is a Hessenberg matrix, and $m < n$. The matrix exponential solution can then be approximated as

$$e^{\boldsymbol{A}t}\boldsymbol{n}_0 \approx \|\boldsymbol{n}_0\|\boldsymbol{Q}_m e^{\boldsymbol{H}_m t}\boldsymbol{e}_1 \quad . \tag{8}$$

This approach appears to be suitable for burnup problems because Krylov subspace methods tend to approximate better the eigenvalues located in the outermost part of the spectrum. These eigenvalues related to the short-lived nuclides are the ones that cause difficulties in most algorithms. Yamamoto, Tatsumi, and Sugimura[7] calculated the matrix exponential using Krylov subspace techniques and

diagonal Padé approximation with EXPOKIT (Ref. 8). They reported promising results in the case where the shortest half-life is ~30 s (^{106}Rh). In this case the burnup matrix norm was approximately $\|\boldsymbol{A}\| \approx 2.3 \times 10^{-2}$.

The m-dimensional Krylov subspace approximation of the matrix exponential is mathematically equivalent to approximating $e^{\boldsymbol{A}t}\boldsymbol{n}_0$ with a polynomial of degree $m-1$ that interpolates the exponential function in the Hermite sense at the eigenvalues of the Hessenberg matrix according to their multiplicities.[9] Consequently, it is evident that the approximation does not work well if the eigenvalues lie far apart from each other, even if the dimension of the subspace is increased. Therefore, the burnup time step must usually be split into smaller substeps in order to keep the eigenvalues located closer to each other.

Selecting the time step is probably the most challenging issue in applying the Krylov subspace approach. Most error estimates are derived from those of truncated Taylor series,[9,10] which leads to highly pessimistic estimates and accordingly impractically short time steps when $\|\boldsymbol{A}\| \gg 1$ and $t \sim 10^6$ s. We found that in such cases error estimation based on the concept of generalized residual[11] gave the most realistic results. However, even when the time step is chosen to be as large as possible, the computation time can easily become prohibitively long. Based on our experiments, it seems that the Krylov subspace approximation alone is not practical for burnup calculations when $\|\boldsymbol{A}\| \gg 1$. However, if the nuclides with the shortest half-lives are excluded from the burnup matrix, this approximation could be a viable replacement for mere Padé approximation with scaling and squaring, for example. Numerical examples are presented in Sec. V.

IV. QUADRATURE FORMULAS AND RATIONAL APPROXIMATION

As mentioned in Sec. II, the eigenvalues of the burnup matrix appear to be generally confined to the vicinity of the negative real axis $\mathbb{R}^-$. This observation is exploited in the matrix exponential method that is described in detail in this section.

IV.A. Relation to Contour Integrals

By the Cauchy integral formula, the solution of system (2) can be represented as a contour integral,

$$\boldsymbol{n}(t) = e^{\boldsymbol{A}t}\boldsymbol{n}_0 = \frac{1}{2\pi i}\int_\Gamma e^z (z\boldsymbol{I} - \boldsymbol{A}t)^{-1}\boldsymbol{n}_0\, dz \ , \qquad (9)$$

where Γ is a closed contour winding once around the spectrum of $\boldsymbol{A}t$. The resolvent of the matrix $\boldsymbol{A}t$ can be written in the form

$$(z\boldsymbol{I} - \boldsymbol{A}t)^{-1} = \frac{\boldsymbol{B}(z)}{\det(z\boldsymbol{I} - \boldsymbol{A}t)} \ , \qquad (10)$$

where

$$\boldsymbol{B}(z) = z^{n-1}\boldsymbol{B}_0 + z^{n-2}\boldsymbol{B}_1 + \cdots + z\boldsymbol{B}_{n-2} + \boldsymbol{B}_{n-1} \qquad (11)$$

with $\boldsymbol{B}_0, \boldsymbol{B}_1, \ldots, \boldsymbol{B}_{n-1}$ matrices with constant elements.[12] It follows that every element of the resolvent is a proper rational function of z with the same denominator polynomial $\det(z\boldsymbol{I} - \boldsymbol{A}t)$. Hence, the poles of these rational functions are the eigenvalues of the matrix $\boldsymbol{A}t$, and calculating $\boldsymbol{n}(t)$ is essentially equivalent to evaluating contour integrals of the form

$$I = \frac{1}{2\pi i}\int_\Gamma e^z f(z)\, dz \ , \qquad (12)$$

where $f = \mathcal{O}(1)$ when $z \to -\infty$, and the singularities of f are the eigenvalues of $\boldsymbol{A}t$.

Integrals of this type are also encountered in the context of Laplace transforms, where they are usually written in the form

$$G(t) = \frac{1}{2\pi i}\int_\Gamma e^{st} g(s)\, ds = \frac{1}{2\pi i}\int_\Gamma e^z g(zt^{-1})t^{-1}\, dz \ . \qquad (13)$$

It should be noted that the solution of system (2) can also be written as an inverse Laplace transform of the form

$$\begin{aligned}\boldsymbol{n}(t) &= \frac{1}{2\pi i}\int_B e^{st}(s\boldsymbol{I} - \boldsymbol{A})^{-1}\boldsymbol{n}_0\, ds \\ &= \frac{1}{2\pi i}\int_B e^z (zt^{-1}\boldsymbol{I} - \boldsymbol{A})^{-1}t^{-1}\boldsymbol{n}_0\, dz \ , \end{aligned} \qquad (14)$$

where B denotes the Bromwich contour running from $-i\infty$ to $+i\infty$.

IV.B. Rational Approximation

When the contour Γ lies in the region of analyticity of f, integral (12) is independent of Γ under mild assumptions. When all of the singularities of function f are confined to a region near the negative real axis, Γ can be widened out to a parabolic or hyperbolic shape in the left complex plane. Because the integrand will decrease exponentially, these contour integrals can be efficiently approximated using numerical methods. These quadrature formulas can be associated with rational functions whose poles are the nodes and residues are the weights of the numerical integration formula. The proof can be found, e.g., in Ref. 13 and is repeated here.

Let $\phi(\theta)$ be an analytic function that maps the real line $\mathbb{R}$ onto the contour Γ that encloses the eigenvalues of matrix $\boldsymbol{A}t$. Integral (12) can then be written

$$I = \frac{1}{2\pi i} \int_{-\infty}^{\infty} e^{\phi(\theta)} f(\phi(\theta)) \phi'(\theta)\, d\theta \ . \tag{15}$$

This integral can be approximated by the trapezoid rule with N points θ_k spaced regularly on the interval $[-\pi, \pi]$ (chosen here for simplicity) yielding the approximation

$$I_N = (iN)^{-1} \sum_{k=1}^{N} e^{z_k} f(z_k) w_k = -\sum_{k=1}^{N} c_k f(z_k) \ , \tag{16}$$

where

$$z_k = \phi(\theta_k) \ ,$$

$$w_k = \phi'(\theta_k) \ ,$$

and

$$c_k = -(iN)^{-1} e^{z_k} w_k = iN^{-1} e^{z_k} w_k \ .$$

By the Cauchy integral formula, this sum can be written

$$I_N = \frac{1}{2\pi i} \int_C r(z) f(z)\, dz \ , \tag{17}$$

where $r(z)$ is a rational function of the form

$$r(z) = \sum_{k=1}^{N} \frac{c_k}{z - z_k} \tag{18}$$

and C is a negatively oriented closed contour that lies in the region of analyticity of f and encloses all the poles z_k.

Let Γ' denote the contour that has the same shape as Γ but lies between the contours C and Γ. Here, $f(z) = \mathcal{O}(1)$ so that $r(z)f(z) = \mathcal{O}(|z|^{-1})$ as $|z| \to \infty$. It follows that the contour C can be deformed to a contour consisting of the union of Γ' and a large circular arc with radius R so that

$$\lim_{R \to \infty} \int_{C_R} r(z) f(z)\, dz = 0 \ . \tag{19}$$

This gives the quadrature rule error estimate

$$I - I_N = \frac{1}{2\pi i} \int_{\Gamma'} (e^z - r(z)) f(z)\, dz \ , \tag{20}$$

which implies that $r(z)$ is a good approximation to e^z near $\mathbb{R}^-$. Therefore, any quadrature formula can be interpreted as a rational approximation. In the same way, every rational approximation can be viewed as a quadrature formula for a contour in the complex plane.

The selection of the contour and quadrature formula has been studied extensively in the context of inverse Laplace transforms. For the case where all of the singularities of f lie on the negative real axis, quite impressive convergence rates have been recently derived.[14] For example, for a cotangent contour originally suggested by Talbot,[15] a convergence rate $\mathcal{O}(3.89^{-N})$ can be achieved by using a trapezoid rule. Of course, it should be kept in mind that the selection of optimal contour and quadrature formula are related to the singularities of the function f.

IV.C. Best Rational Approximation

Another approach to rational approximation is to calculate the best approximation on some subset of the complex plane. This approach was made famous by Cody, Meinardus, and Varga[16] in 1969 in the context of rational approximation of e^{-x} in $[0,\infty)$. Let $\pi_{k,l}$ denote the collection of all real rational functions $r_{k,l}(x)$ of the form

$$r_{k,l} = \frac{p_k(x)}{p_l(x)} \ , \tag{21}$$

where p_j is a polynomial of degree j or less.

It is known from approximation theory that there exists a unique $\hat{r}_{k,l} \in \pi_{k,l}$ such that

$$\sup_{-\infty < x \le 0} |\hat{r}_{k,l}(-x) - e^x| = \inf_{r_{k,l} \in \pi_{k,l}} \left\{ \sup_{-\infty < x \le 0} |r_{k,l}(-x) - e^x| \right\} \ , \quad k \le l \ . \tag{22}$$

Establishing this approximation for given k and l is not easy, but it can be done with the Remes algorithm or the Carathéodory-Fejér method. It has been shown that this Chebyshev rational approximation $\hat{r}_{k,k}$ converges approximately at the rate 9.3^{-k} (Ref. 17). The contour plot of $|e^z - \hat{r}_{14,14}(-z)|$ is shown in Fig. 1, from which it can be seen that this approximation is remarkably accurate in a wide region in the left complex plane. From a computational point of view, it is advantageous that the poles $\{\theta_1, \ldots, \theta_k\}$ of the rational function $\hat{r}_{k,k}$ are distinct, so that it can be computed as a partial fraction expansion[10]

$$\hat{r}_{k,k}(z) = \alpha_0 + \sum_{i=1}^{k} \frac{\alpha_i}{z - \theta_i} \ , \tag{23}$$

where α_0 is the limit of the function at infinity and the scalars α_i are the residues at the poles θ_i. Therefore, the values of α_i and θ_i depend on k. Equation (23) can be derived by noting that $(\hat{r}_{k,k} - \alpha_0) \in \pi_{k-1,k}$ for which the result readily follows from the residue theorem. It should be noted that the poles of $\hat{r}_{k,k}$ come in conjugate pairs, so that for a real-valued variable $x \in \mathbb{R}$, the computational cost can be reduced to half:

$$\hat{r}_{k,k}(x) = \alpha_0 + \mathrm{Re}\left(\sum_{i=1}^{k/2} \frac{\alpha_i}{x - \theta_i} \right) \ . \tag{24}$$

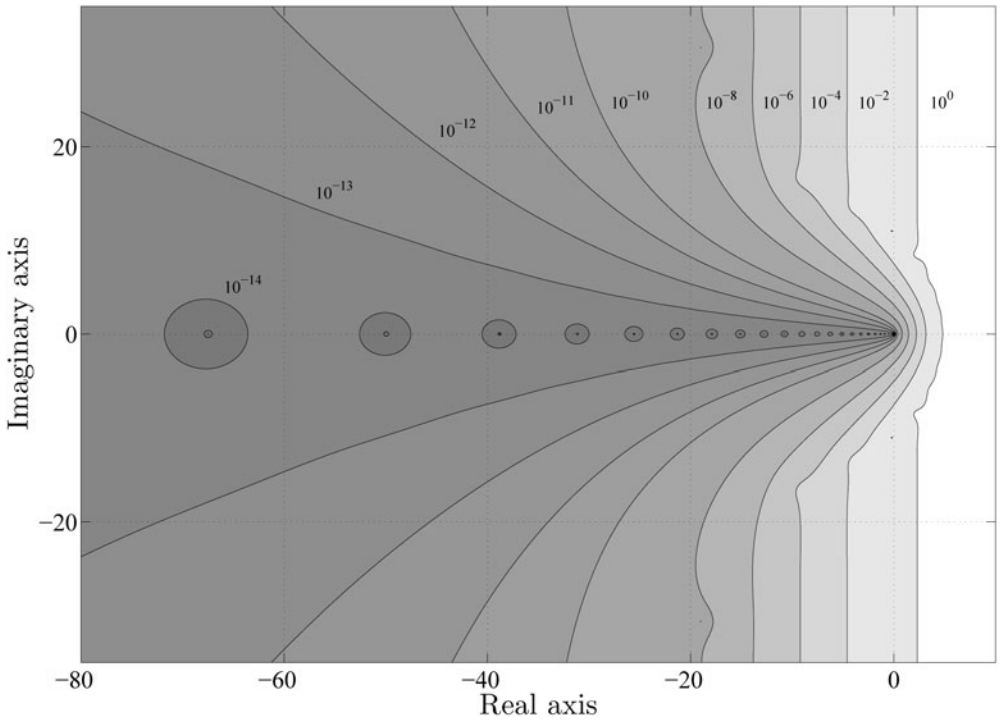

Fig. 1. Contour plot of $|e^z - \hat{r}_{14,14}(-z)|$.

The sets of coefficients for the Chebyshev rational function $\hat{r}_{k,k}$ have been reported for various approximation orders k, so the implementation of this method is relatively straightforward. For example, in Ref. 18 the polynomial coefficients are provided for $k \leq 30$. The partial fraction coefficients α_i and θ_i for each value of k can then be calculated from these polynomial coefficients with the help of a polynomial root finder. The partial fraction coefficients for the cases $k = 10$ and $k = 14$ have been directly given in Ref. 10. They can also be computed with the Carathéodory-Fejér method for $k \leq 14$ with good accuracy, and there is a MATLAB script provided for this purpose in Ref. 13.

Interestingly, the Chebyshev rational approximation $\hat{r}_{k,k}$ can also be interpreted as a quadrature formula for a contour integral of type (12), so the error estimate (20) remains valid.[13] This suggests that the rational approximation could be used for computing the matrix exponential $e^{\boldsymbol{A}t}$ when the eigenvalues of $\boldsymbol{A}t$ are located near the negative real axis. This has also been experimentally verified.[9] From this point of view, the accuracy of the approximation is affected by the magnitudes of the imaginary parts of the eigenvalues of $\boldsymbol{A}$ as long as the eigenvalues remain within the integration contour. However, if the eigenvalues fall outside the contour, Eq. (20) is no longer valid, and this method may yield poor results.

The Chebyshev rational approximation has previously been only occasionally used in scientific applications involving self-adjoint and negative semidefinite matrices.[8,16,19] Equation (20) implies, however, that this approximation is also applicable to non-Hermitian matrices with eigenvalues near $\mathbb{R}^-$. The formal convergence analysis of this special case forms an interesting future research topic.

For the burnup system (2), the matrix exponential solution based on the Chebyshev rational approximation $\hat{r}_{k,k}$ can be computed as

$$\begin{aligned}\boldsymbol{n}(t) &= e^{\boldsymbol{A}t}\boldsymbol{n}_0 \approx \hat{r}_{k,k}(-\boldsymbol{A}t)\boldsymbol{n}_0 \\ &= \alpha_0\boldsymbol{n}_0 - \mathrm{Re}\left(\sum_{i=1}^{k/2}(\theta_i\boldsymbol{I} + \boldsymbol{A}t)^{-1}\alpha_i\boldsymbol{n}_0\right) , \end{aligned} \quad (25)$$

where the last form follows directly from Eq. (24) by replacing x with $-\boldsymbol{A}t$. Using this formula, the concentration vector can be calculated simply by solving $k/2$ sparse linear systems. When the burnup matrix is formed by indexing the nuclides in ascending order with respect to their mass number, these systems can be solved efficiently by first calculating the symbolic lower-upper (LU) factorization of $\boldsymbol{A}$ (Ref. 20) and then performing a Gaussian elimination on this factorization.[21] The structure of a typical large burnup matrix generated in this manner is shown in Fig. 2.

It should be noted that this approximation is ideally suited for decay transmutation calculations, where the absence of closed cycles in the transition chains confines all eigenvalues of the decay matrix to lie strictly on the

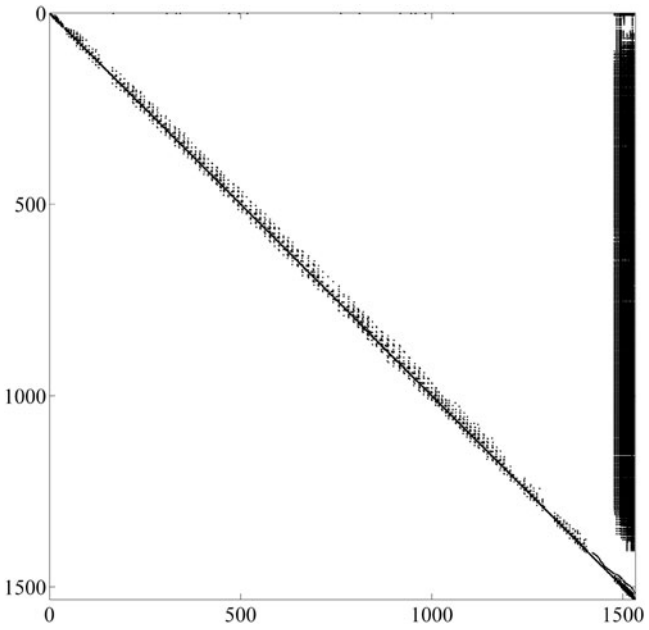

Fig. 2. The structure of the burnup matrix in test case 2 described in Sec. V. The nuclides have been indexed in ascending order with respect to their mass number. The result is that the nonzero elements are concentrated around the diagonal with fission product distributions on the right side.

negative real axis. Another fact worth noticing in Eq. (25) is that the order of the approximation can be adjusted according to needs for accuracy without significant impact on the computational cost as the amount $k/2$ of sparse matrix inversions is directly proportional to the order k of the approximation. However, it should be kept in mind that a rigorous mathematical analysis concerning the convergence properties of this approximation for other than self-adjoint negative semidefinite matrices has not been performed.

V. NUMERICAL RESULTS

The different matrix exponential methods and the TTA method were compared to each other by applying them to several burnup matrices. The two cases presented here can be thought to represent the extreme cases—in terms of matrix size and norm—that were encountered, and they were chosen for evaluating the performance of the different matrix exponential methods.

Both test cases represent an infinite pressurized water reactor (PWR) pin-cell lattice in which the fuel has been irradiated to 25 MWd/kg U burnup. Test case 1 was formed by selecting only the most important actinides and fission products in the calculation, totaling 219 nuclides (41 actinides, 178 fission products and light nuclides). For this case the matrix norm is sufficiently small, $\|\boldsymbol{A}\| \approx 7.3 \times 10^{-4}$, so that $\|\boldsymbol{A}t\| \approx 7.9 \times 10^{3}$. This case is a simplification of test case 2, which contains 1532 nuclides (75 actinides, 1457 fission products and light nuclides). The burnup matrix norm for this case is approximately $\|\boldsymbol{A}\| \approx 2.8 \times 10^{21}$ so that the norm of $\boldsymbol{A}t$ is of order 10^{28}. The time step in both test cases was 125 days corresponding to 5 MWd/kg U burnup.

The TTA results were obtained directly from the Serpent code, and the Chebyshev rational approximation method (CRAM) of order $k = 14$ was implemented as a separate C code that was later added to Serpent. This order for the Chebyshev approximation was chosen because it is generally considered sufficiently accurate[22] and because the partial fraction coefficients for this case are conveniently listed in Ref. 10. The Krylov subspace approximation with adaptive time step and subspace dimension selection was implemented as a MATLAB script. Finally, the standard MATLAB function `expm` was used for the Padé approximation with scaling and squaring.

The numerical results for test case 1 are shown in Fig. 3, from which it can be seen that all results are in good accordance with each other. In particular, the Padé approximation, the Krylov subspace method, and CRAM give almost identical results for this case, as can be seen from Fig. 4, where the absolute values of the relative differences are plotted. For example, the largest relative difference between the Chebyshev and Padé approximation solutions is ~0.00068% for the concentration of ^{252}Cf, for which $N \approx 2.17 \times 10^{-19}$ (b cm)$^{-1}$.

The nuclide concentrations in test case 2 are shown in Fig. 5. The Krylov subspace method could not be applied to this case because the time step selection based on local error estimation became practically impossible. The Padé approximation also faced severe numerical problems producing completely unrealistic results, as can be seen from Fig. 5. On the other hand, the solutions calculated with the TTA method and CRAM are consistent to the same degree as in test case 1. The comparison between these numerical results for the most important nuclides is presented in Table I.

The small differences between the TTA and the CRAM solutions can be attributed to the fact that the closed cycles are terminated in the current implementation of the TTA method. This is supported by the fact that all the concentrations calculated with CRAM are slightly greater, as can be physically expected considering that the feedback transitions are ignored in the TTA calculation. Also, the largest differences occur for nuclides for which the closed transition cycles are significant. In test case 1, for example, the largest relative difference, 0.84%, occurs for the hydrogen isotope ^{3}H, which forms one strongly connected component with the nuclides ^{1}H, ^{2}H, and ^{3}He.

As is pointed out in Sec. IV, the accuracy of CRAM depends on the magnitudes of the imaginary parts of the eigenvalues of $\boldsymbol{A}t$. In this test case the power density in the fuel was 40 kW/kg U, which results in a neutron flux

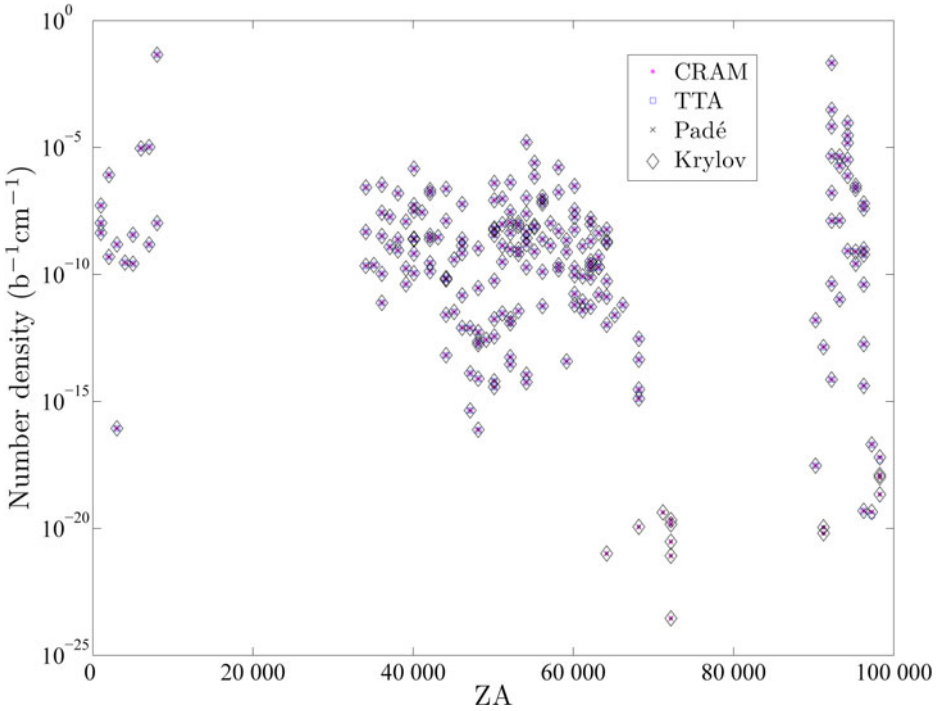

Fig. 3. Nuclide number densities for test case 1. Number densities smaller than 10^{-30} $(\mathrm{b\,cm})^{-1}$ have been omitted. ZA = $1000Z + A$, where Z is the atomic number and A is the mass number of the nuclide.

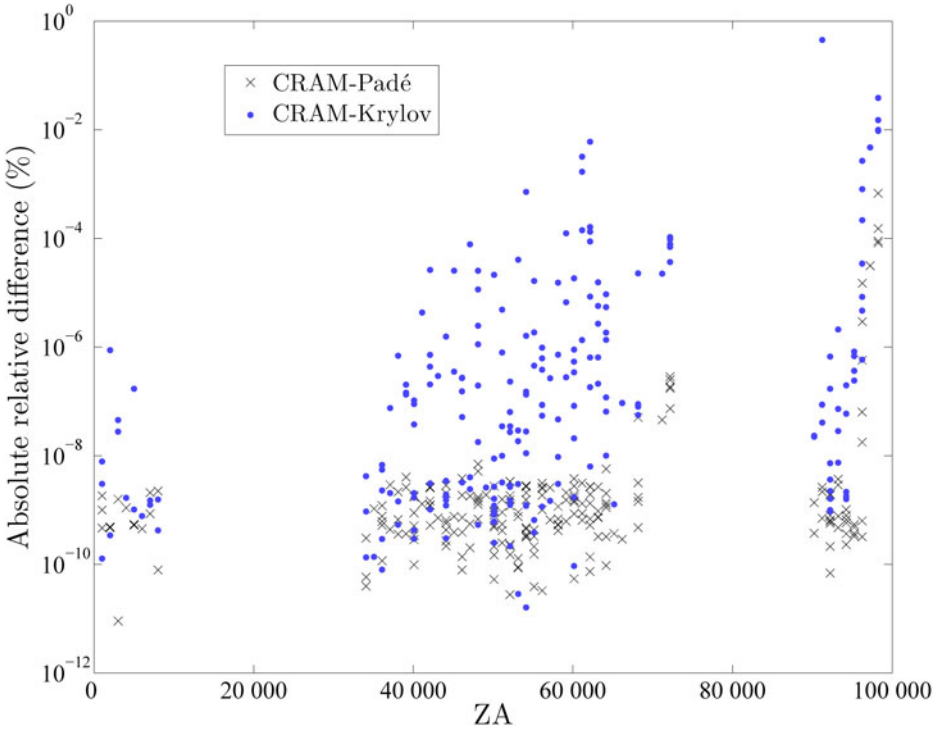

Fig. 4. Absolute values of the relative differences between the results calculated by CRAM, Padé approximation, and the Krylov subspace method. ZA = $1000Z + A$, where Z is the atomic number and A is the mass number of the nuclide.

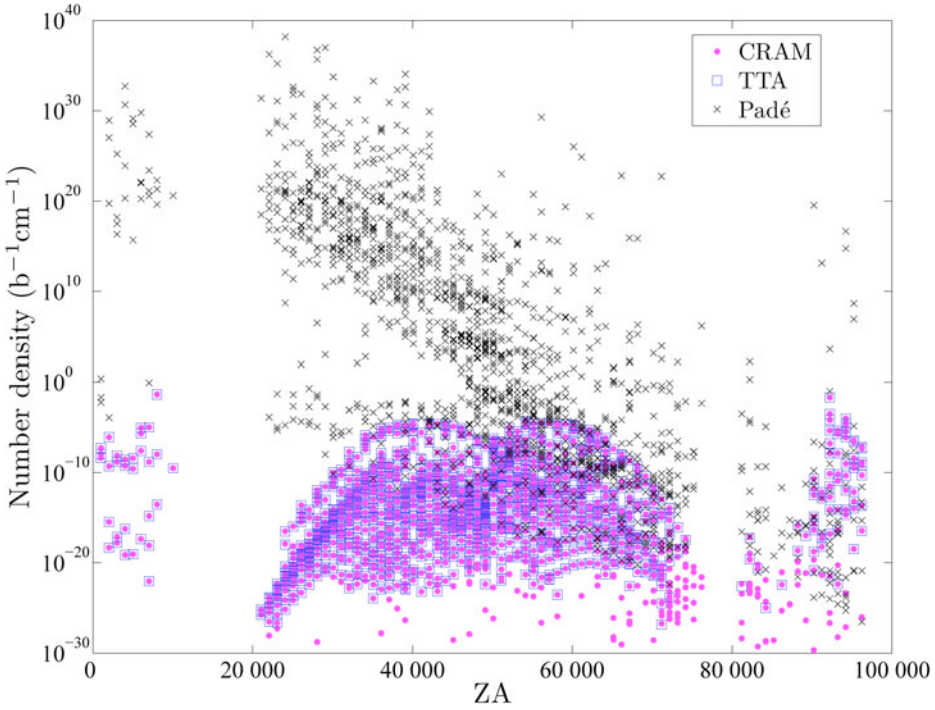

Fig. 5. Nuclide number densities in test case 2. Number densities smaller than 10^{-30} b^{-1} cm^{-1} have been omitted. ZA = $1000Z + A$, where Z is the atomic number and A is the mass number of the nuclide.

of 1.5×10^{14} 1/(cm^2 s) producing transmutation coefficients, most of which are of order $\leq 10^{-8}$ 1/s. The imaginary parts of the eigenvalues were correspondingly of the same order. The time steps in the burnup calculation generally vary from a few days at the beginning of the irradiation cycle to a few hundred days at the end. The time step used here was 125 days, which is of the same order as the practical maximum time step. Time steps greater than this would significantly violate the assumption of constant transmutation coefficients during each step. Therefore, even better convergence can be expected for shorter time steps.

Based on these observations, it seems that the CRAM method is capable of providing robust and accurate solutions regardless of the burnup matrix size or norm. The method is also computationally remarkably effective. The computation time for test case 2 involving a 1532×1532 matrix was only 0.1 s on a 2.6-GHz AMD Opteron CPU. The corresponding computation time for the TTA calculation was ~26 s. To further illustrate the efficiency of the CRAM method, it was compared to the TTA solution method using Serpent in a burnup calculation for a PWR fuel assembly with burnable absorber. The total number of depleted materials was 65, the irradiation history was divided into 42 steps with predictor-corrector calculation, and a total of 3 million neutron histories were run for each Monte Carlo simulation. The overall running time with TTA was 18.5 h, and using CRAM this was reduced to just over 13 h, which can be considered a significant improvement.

When computing the CRAM solution, most of the computation time is spent inverting the sparse matrices of Eq. (25). As pointed out in Sec. IV, the structure of the burnup matrix is crucial to the effectiveness of the solution scheme. In comparison to a random nuclide order, indexing the nuclides according to their mass number led to a computational speedup factor of 40 in test case 2. The corresponding ordered matrix is illustrated in Fig. 2. In this case, the sparse systems can be solved accurately and effectively by first calculating the symbolic LU factorization and then performing a Gaussian elimination on this factorization.

VI. CONCLUSIONS

The magnitude of the transmutation and decay constants of different nuclides vary extensively, which makes calculating the matrix exponential challenging in the context of burnup calculations. Short-lived nuclides are especially problematic because they can increase the matrix norm and induce eigenvalues with absolute values up to order 10^{21}.

We approached this problem by examining the characteristics of the eigenvalues of the burnup matrix. Based

TABLE I

A Comparison of the Numerical Results Computed Using CRAM and the TTA Method for the Most Important Nuclides

Nuclide	Case 1		Case 2	
	Concentration, CRAM ($b^{-1}\,cm^{-1}$)	Relative Difference to TTA (%)	Concentration, CRAM ($b^{-1}\,cm^{-1}$)	Relative Difference to TTA (%)
Actinides				
^{234}U	4.5718×10^{-6}	5.4633×10^{-5}	4.5632×10^{-6}	5.5873×10^{-5}
^{235}U	3.0456×10^{-4}	7.3299×10^{-5}	3.0738×10^{-4}	8.0322×10^{-5}
^{236}U	6.7457×10^{-5}	2.4295×10^{-5}	6.7627×10^{-5}	2.6525×10^{-5}
^{238}U	2.1423×10^{-2}	2.5090×10^{-5}	2.1413×10^{-2}	2.7600×10^{-5}
^{239}U	1.2933×10^{-8}	2.5083×10^{-5}	1.3510×10^{-8}	2.7595×10^{-5}
^{237}Np	4.3857×10^{-6}	3.2838×10^{-5}	4.4366×10^{-6}	4.2891×10^{-5}
^{239}Np	1.8656×10^{-6}	2.4399×10^{-5}	1.9488×10^{-6}	2.6853×10^{-5}
^{238}Pu	7.5350×10^{-7}	2.8311×10^{-5}	7.8289×10^{-7}	4.7964×10^{-5}
^{239}Pu	9.1938×10^{-5}	6.6197×10^{-5}	9.5985×10^{-5}	7.0639×10^{-5}
^{240}Pu	2.9316×10^{-5}	3.9718×10^{-4}	2.9007×10^{-5}	4.6375×10^{-4}
^{241}Pu	1.5002×10^{-5}	3.9580×10^{-4}	1.6400×10^{-5}	4.5646×10^{-4}
^{242}Pu	3.2839×10^{-6}	1.3219×10^{-4}	3.5014×10^{-6}	1.5190×10^{-4}
^{241}Am	2.3450×10^{-7}	1.2989×10^{-4}	2.5499×10^{-7}	1.6080×10^{-4}
^{242}Am	7.9173×10^{-10}	1.3073×10^{-4}	8.6280×10^{-10}	1.6164×10^{-4}
^{243}Am	3.1046×10^{-7}	4.7632×10^{-5}	3.4995×10^{-7}	5.4590×10^{-5}
^{242}Cm	6.2763×10^{-8}	4.7059×10^{-5}	6.8238×10^{-8}	6.1609×10^{-5}
^{244}Cm	3.9601×10^{-8}	9.7363×10^{-5}	4.8298×10^{-8}	1.0704×10^{-4}
Fission products and light nuclides				
^{3}H	5.2271×10^{-8}	8.4335×10^{-1}	5.2535×10^{-8}	8.2818×10^{-1}
^{95}Mo	1.6225×10^{-7}	7.4205×10^{-6}	1.9623×10^{-5}	8.6482×10^{-6}
^{99}Tc	2.9344×10^{-9}	1.0370×10^{-5}	2.5379×10^{-5}	2.1573×10^{-5}
^{103}Ru	6.2905×10^{-11}	5.7469×10^{-5}	2.3941×10^{-6}	8.4875×10^{-5}
^{109}Ag	7.7335×10^{-13}	1.1614×10^{-3}	1.8199×10^{-6}	1.3940×10^{-4}
^{135}Xe	2.7013×10^{-9}	4.9536×10^{-4}	7.2419×10^{-9}	5.6080×10^{-4}
^{133}Cs	8.2702×10^{-9}	3.2331×10^{-5}	2.6399×10^{-5}	6.5866×10^{-5}
^{143}Nd	3.4130×10^{-8}	4.3039×10^{-5}	1.9255×10^{-5}	6.6319×10^{-5}
^{145}Nd	1.7186×10^{-11}	2.3102×10^{-4}	1.6098×10^{-5}	1.2200×10^{-4}
^{147}Sm	1.7783×10^{-10}	1.6831×10^{-5}	1.0618×10^{-6}	1.3204×10^{-5}
^{149}Sm	5.2676×10^{-12}	2.0827×10^{-4}	6.6335×10^{-8}	4.3958×10^{-3}
^{150}Sm	3.2992×10^{-10}	8.6619×10^{-5}	5.3727×10^{-6}	3.4225×10^{-4}
^{151}Sm	2.0391×10^{-9}	4.2694×10^{-4}	3.1044×10^{-7}	5.6033×10^{-3}
^{152}Sm	9.3582×10^{-9}	3.3607×10^{-4}	2.5405×10^{-6}	8.8872×10^{-4}
^{153}Eu	4.3027×10^{-9}	1.2505×10^{-4}	1.8654×10^{-6}	3.8255×10^{-3}
^{155}Gd	1.0272×10^{-12}	1.0841×10^{-1}	6.4533×10^{-10}	1.4526×10^{-2}

on our experiments and physical reasoning, it seems that these eigenvalues are generally confined to a region near the negative real axis. The somewhat obscure fact that the CRAM technique can be interpreted as a numerical contour integral in the left complex plane led us to conduct further experiments with very promising results. We compared this approach with more established matrix exponential methods and the TTA method in solving the burnup equations. Our results imply that the CRAM solution scheme is well-suited for burnup calculation, where it outperformed the more conventional matrix exponential methods in terms of computational accuracy and efficiency. Unlike the previously applied matrix exponential methods, CRAM can readily treat the short-lived nuclides simultaneously with the long-lived nuclides. In addition, the practical maximum time step value can be used in CRAM without compromising the computational accuracy.

For evaluating the matrix exponential methods, we constructed two representative test cases. Test case 1 was designed to be well-behaved in terms of the burnup matrix size and norm ($[\boldsymbol{A}] \approx 200 \times 200$, $\|\boldsymbol{A}\| \sim 10^{-4}$), and test case 2 was designed to be pathologically difficult ($[\boldsymbol{A}] \approx 1500 \times 1500$, $\|\boldsymbol{A}\| \sim 10^{21}$). In test case 1, which

was well within the applicability domain of each tested method, the Padé approximation and CRAM gave virtually identical results. The results obtained with the Krylov subspace matrix exponential method were close, as well. The system was also solved using the TTA method, which gave coherent results. In test case 2, however, all other matrix exponential methods suffered a breakdown, but the results obtained with CRAM remained consistent with those given by the TTA method to the same degree as in the first test case.

Our motivation for the research was the prospect of incorporating a matrix exponential method in the burnup calculation routine of the Serpent code. Based on our positive results, CRAM was added to the code with computational speedup as one of the key improvements.

REFERENCES

1. C. MOLER and C. VAN LOAN, "Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later," *SIAM Rev.*, **45** (2003).

2. J. LEPPÄNEN, "PSG2/Serpent—A Continuous-Energy Monte Carlo Reactor Physics Burnup Calculation Code," VTT Technical Research Centre of Finland; available on the Internet at http://montecarlo.vtt.fi (Nov. 2008).

3. J. CETNAR, "General Solution of Bateman Equations for Nuclear Transmutations," *Ann. Nucl. Energy*, **33**, 640 (2006).

4. H. AMANN, *Ordinary Differential Equations, An Introduction to Nonlinear Analysis*, Walter de Gruyter, Berlin (1990).

5. R. E. TARJAN, "Depth-First Search and Linear Graph Algorithms," *SIAM J. Comput.*, **1**, *2*, 146 (1972).

6. I. C. GAULD, O. W. HERMANN, and R. M. WESTFALL, "Origen-S: Scale System Module to Calculate Fuel Depletion, Actinide Transmutation, Fission Product Buildup and Decay, and Associated Radiation Source Terms," in "SCALE: A Modular Code System for Performing Standardized Computer Analyses for Licensing Evaluation," Vol. II, Sec. F7, Oak Ridge National Laboratory/U.S. Nuclear Regulatory Commission (Nov. 2006).

7. A. YAMAMOTO, M. TATSUMI, and N. SUGIMURA, "Numerical Solution of Stiff Burnup Equations with Short Half Lived Nuclides by the Krylov Subspace Method," *J. Nucl. Sci. Technol.*, **44**, *2*, 147 (2007).

8. R. B. SIDJE, "Expokit: A Software Package for Computing Matrix Exponentials," *ACM Trans. Math. Software*, **24**, *1*, 130 (1998).

9. Y. SAAD, "Analysis of Some Krylov Subspace Approximations to the Matrix Exponential Operator," *SIAM J. Numer. Anal.*, **29**, *1*, 209 (1992).

10. E. GALLOPOULOS and Y. SAAD, "Efficient Solution of Parabolic Equations by Krylov Approximation Methods," *SIAM J. Sci. Stat. Comput.*, **13**, *5*, 1236 (1992).

11. M. HOCHBRUCK, C. LUBICH, and H. SELHOFER, "Exponential Integrators for Large Systems of Differential Equations," *SIAM J. Sci. Comput.*, **19**, *5*, 1552 (1998).

12. L. A. ZADEH and C. A. DESOER, *Linear System Theory: The State Space Approach*, McGraw-Hill Book Company (1963).

13. L. N. TREFETHEN, J. A. C. WEIDEMAN, and T. SCHMELZER, "Talbot Quadratures and Rational Approximations," *BIT Numer. Math.*, **46**, *3*, 653 (2006).

14. J. WEIDEMAN and L. N. TREFETHEN, "Parabolic and Hyperbolic Contours for Computing the Bromwich Integral," *Math. Comp.*, **76**, *259*, 1341 (2007).

15. A. TALBOT, "The Accurate Numerical Inversion of Laplace Transforms," *J. Instrum. Math. Appl.*, **23**, 97 (1979).

16. W. J. CODY, G. MEINARDUS, and R. S. VARGA, "Chebyshev Rational Approximations to e^{-x} in $[0,\infty)$ and Applications to Heat-Conduction Problems," *J. Approx. Theory*, **2**, 50 (1969).

17. A. A. GONCHAR and E. A. RAKHMANOV, "Equilibrium Distributions and Degree of Rational Approximation of Analytic Functions," *Mat. Sb.*, **134**, 306 (1987); see also *Math. USSR Sb.*, **62**, *2*, 305 (1989) (translation).

18. A. J. CARPENTER, A. RUTTAN, and R. S. VARGA, "Extended Numerical Computations on the 1/9 Conjecture in Rational Approximation Theory," in "Rational Approximation and Interpolation," *Lecture Notes in Mathematics*, Vol. 1105, pp. 383–411, P. R. GRAVES-MORRIS, E. B. SAFF, and R. S. VARGA, Eds., Springer-Verlag (1984).

19. T. SCHMELZER and L. N. TREFETHEN, "Evaluating Matrix Functions for Exponential Integrators via Carathéodory-Fejér Approximation and Contour Integrals," *Electron. Trans. Numer. Anal.*, **28**, 1 (2007).

20. D. J. ROSE and R. E. TARJAN, "Algorithmic Aspects of Vertex Elimination on Directed Graphs," *SIAM J. Appl. Math*, **40**, 176 (1978).

21. R. E. TARJAN, "Graph Theory and Gaussian Elimination," CS-TR-75-526, Stanford University, Department of Computer Science (1975).

22. L. LOPEZ and V. SIMONCINI, "Analysis of Projection Methods for Rational Function Approximation to the Matrix Exponential," *SIAM J. Numer. Anal.*, **44**, *2*, 613 (2007).