PUBLICATION II

Rational approximations to the matrix exponential in burnup calculations

In: Nuclear Science and Engineering,
169, 2, pp. 155–167.

NUCLEAR SCIENCE AND ENGINEERING: **169**, 155–167 (2011)

Rational Approximations to the Matrix Exponential in Burnup Calculations

Maria Pusa*

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

Received October 26, 2010
Accepted February 28, 2011

Abstract—*The topic of this paper is solving the burnup equations using dedicated matrix exponential methods that are based on two different types of rational approximation near the negative real axis. The previously introduced Chebyshev Rational Approximation Method (CRAM) is now analyzed in detail for its accuracy and convergence, and correct partial fraction coefficients for approximation orders 14 and 16 are given to facilitate its implementation and improve the accuracy. As a new approach, rational approximation based on quadrature formulas derived from complex contour integrals is proposed, which forms an attractive alternative to CRAM, as its coefficients are easy to compute for any order of approximation. This gives the user the option to routinely choose between computational efficiency and accuracy all the way up to the level permitted by the available arithmetic precision. The presented results for two test cases are validated against reference solutions computed using high-precision arithmetics. The observed behavior of the methods confirms the previous conclusions of CRAM's excellent suitability for burnup calculations and establishes the quadrature-based approximation as a viable and flexible alternative that, like CRAM, has its foundation in the specific eigenvalue properties of burnup matrices.*

I. INTRODUCTION

The topic of this paper is solving the burnup equations using matrix exponential methods based on rational approximation near the negative real axis. Solving the burnup equations is an essential part of the burnup calculations that are necessary to predict the changes in the material compositions in a nuclear reactor.

The burnup equations form a system of first-order linear differential equations that can be written in matrix notation as

$$\boldsymbol{n}' = \boldsymbol{A}\boldsymbol{n} \ , \quad \boldsymbol{n}(0) = \boldsymbol{n}_0 \ , \tag{1}$$

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.

*Email: Maria.Pusa@vtt.fi

The matrix elements A_{ij} characterize the rates of neutron-induced reactions and spontaneous radioactive decay by which nuclide j is transformed to nuclide i. In this paper these coefficients are assumed to be fixed constants.

The burnup equations can be formally solved by the matrix exponential method yielding the simple solution

$$\boldsymbol{n}(t) = e^{\boldsymbol{A}t}\boldsymbol{n}_0 \ , \tag{2}$$

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

$$e^{\boldsymbol{A}t} = \sum_{k=0}^{\infty} \frac{1}{k!} (\boldsymbol{A}t)^k \ , \tag{3}$$

with the additional definition $\boldsymbol{A}^0 = \boldsymbol{I}$. There are numerous algorithms for computing the matrix exponential, but unfortunately, most of them are not well-suited for solving the burnup equations. Because the decay constants of the nuclides vary extensively, the burnup matrix has a wide spectrum of eigenvalues. Short-lived

nuclides are especially problematic since they can induce eigenvalues with absolute values up to an order of 10^{21}, consequently making the differential system extremely stiff. Furthermore, the time steps used in burnup calculations can typically vary from a few days (10^5 s) to several months (10^7 s), and even to thousands of years, if only decay reactions are considered. Most of the established matrix exponential methods, such as the truncated Taylor series approach or rational Padé approximation, are based on approximation near the origin and work well only when the matrix norm $\|\boldsymbol{A}t\|$ is sufficiently small. Consequently, these algorithms are prone to severe numerical problems when applied to the burnup equations, where this norm can be of the order of 10^{27} (Ref. 1).

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. For example, in the ORIGEN code,[2] the matrix exponential is computed with the truncated Taylor series method with scaling and squaring after excluding short-lived nuclides from the burnup matrix to be treated separately. In the AEGIS code, a Krylov subspace method is applied to a simplified burnup chain with 221 nuclides, in which case the burnup matrix norm is of the order of 10^{-2} (Ref. 3). However, it was recently discovered by the author that the eigenvalues of the burnup matrix are generally confined to a region near the negative real axis.[1] This observation led to applying the Chebyshev Rational Approximation Method (CRAM) to solve the burnup equations. This method can be interpreted as the best rational approximation on the negative real axis, and it was shown to give a robust and accurate solution to the burnup equations with a very short computation time. For further information on established matrix exponential methods and their applicability to solving burnup equations, see, e.g., Refs. 1 and 4. For a comparison between CRAM and ORIGEN, see Ref. 5.

The main challenge in using CRAM is determining the coefficients of the Chebyshev rational approximation. The computation of higher-order CRAM coefficients especially can be rather involved. Motivated by these challenges in implementing CRAM, an alternative and easier-to-implement method is presented in Sec. II.B. This method is based on constructing rational approximations from trapezoidal quadrature rules applied to contour integrals in the left complex plane. Although these approximations do not converge as fast as CRAM, they have the advantage that the order of the approximation can easily be adjusted. These approximations are accurate near the negative real axis, so they are well-suited to solving the burnup equations and can be used to obtain extremely high solution accuracy. The convergence and accuracy of the different rational approximations applied to burnup equations are discussed in Sec. III.

II. RATIONAL APPROXIMATION OF THE MATRIX EXPONENTIAL NEAR THE NEGATIVE REAL AXIS

The matrix exponential can be computed based on a rational function $r(z)$ that is known to be a good approximation to the function e^z in some region in the complex plane $\mathbb{C}$. The matrix exponential and the approximating matrix rational function can be defined in various ways. The approach based on the Cauchy integral formula is considered here. Based on this formula, the matrix exponential can be written as a complex contour integral of the form

$$e^{\boldsymbol{A}t} = \frac{1}{2\pi i}\int_{\Gamma} e^z (z\boldsymbol{I} - \boldsymbol{A}t)^{-1}\, dz \ , \tag{4}$$

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

$$(z\boldsymbol{I} - \boldsymbol{A}t)^{-1} = \frac{\boldsymbol{B}(z)}{\det(z\boldsymbol{I} - \boldsymbol{A}t)} \ , \tag{5}$$

where

$$\boldsymbol{B}(z) = z^{n-1}\boldsymbol{B}_0 + z^{n-2}\boldsymbol{B}_1 + \cdots + z\boldsymbol{B}_{n-2} + \boldsymbol{B}_{n-1} \tag{6}$$

with $\boldsymbol{B}_0, \boldsymbol{B}_1, \ldots, \boldsymbol{B}_{n-1}$ matrices with constant elements.[6] 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 $e^{\boldsymbol{A}t}$ is essentially equivalent to evaluating contour integrals of the form

$$(e^{\boldsymbol{A}t})_{kl} = \frac{1}{2\pi i}\int_{\Gamma} e^z R_{kl}(z)\, dz \ , \tag{7}$$

where $\boldsymbol{R} = (z\boldsymbol{I} - \boldsymbol{A}t)^{-1}$, $R_{kl} = \mathcal{O}(1)$ when $z \to -\infty$, and the singularities of R_{kl} are the eigenvalues of $\boldsymbol{A}t$. It follows that when the eigenvalues of $\boldsymbol{A}t$ are confined to a region near the negative real axis, Γ can be extended to a parabolic or hyperbolic shape in the left complex plane. Because the integrand will decrease exponentially, these contour integrals can be approximated efficiently using quadrature formulas. These quadrature formulas can be associated with rational functions, whose poles and residues are the nodes and weights of the numerical integration formula, respectively. In addition, every rational function can be correspondingly interpreted as a quadrature formula applied to a contour integral in the left complex plane (for proof, see Ref. 7).

It is usually advantageous to employ the rational approximation in the partial fraction decomposition (PFD) form. For a rational function $r_{k,k}(z) = p_k(z)/q_k(z)$ with simple poles, and p_k and q_k being polynomials of order k, the decomposition takes the form

$$r_{k,k}(z) = \alpha_0 + \sum_{j=1}^{k} \frac{\alpha_j}{z - \theta_j} \ , \tag{8}$$

where

α_0 = limit of the function $r_{k,k}$ at infinity

α_j = residues at the poles θ_j.

Also rational functions $r_{k-1,k}$ can be written in this form with $\alpha_0 = 0$. The poles of a rational function with real-valued coefficients form conjugate pairs, so the computational cost can be reduced to half for a real variable x:

$$r_{k,k}(x) = \alpha_0 + 2\mathrm{Re}\left(\sum_{j=1}^{k/2} \frac{\alpha_j}{x - \theta_j} \right) \ . \tag{9}$$

The rational approximation to Eq. (2) can then be written

$$\boldsymbol{n} = \alpha_0 \boldsymbol{n}_0 + 2\mathrm{Re}\left(\sum_{j=1}^{k/2} \alpha_j (\boldsymbol{A}t - \theta_j \boldsymbol{I})^{-1} \boldsymbol{n}_0 \right) \ , \tag{10}$$

which requires solving $k/2$ sparse linear systems. It is worth noting that the linear systems in Eq. (10) are independent, so they can be solved in parallel. Notice that Eq. (10) can be used to apply any rational approximation $r_{k,k}$ or $r_{k-1,k}$ to Eq. (2). A MATLAB code implementing this equation is shown in Fig. 1 to further illustrate how these approximations are computed in practice.

II.A. Chebyshev Rational Approximation Method

In CRAM the rational function $r(z)$ is chosen as the best rational approximation of the exponential function on the negative real axis $\mathbb{R}_-$. Let $\pi_{k,k}$ denote the set of rational functions $r_{k,k}(x) = p_k(x)/q_k(x)$, where p_k and q_k are polynomials of order k. The CRAM approximation of order k is defined as the unique rational function $\hat{r}_{k,k} = \hat{p}_k(x)/\hat{q}_k(x)$ satisfying

$$\sup_{x \in \mathbb{R}_-} |\hat{r}_{k,k}(x) - e^x| = \inf_{r_{k,k} \in \pi_{k,k}} \left\{ \sup_{x \in \mathbb{R}_-} |r_{k,k}(x) - e^x| \right\} \ . \tag{11}$$

The asymptotic convergence of this approximation on the negative real axis is remarkably fast. Let us define

$$\delta_k = \sup_{x \in \mathbb{R}_-} |\hat{r}_{k,k}(x) - e^x| \ . \tag{12}$$

It has been proven that

$$\lim_{k \to \infty} \delta_k^{1/k} = \frac{1}{9.289\,025\,49\ldots} = H \ , \tag{13}$$

where H is the Halphen constant that can be represented in closed form using certain elliptic integrals.[8] It follows that for sufficiently large approximation orders k, roughly k correct digits may be expected. Surprisingly, it was recently discovered by Stahl and Schmelzer[9] that this convergence extends to compact subsets on the complex plane and also to Hankel contours in $\mathbb{C} \setminus \mathbb{R}_-$, i.e.,

$$\lim_{k \to \infty} \left(\sup_{z \in K} |\hat{r}_{k,k}(z) - e^z| \right)^{1/k} = \lim_{k \to \infty} \left(\sup_{z \in \Gamma} |\hat{r}_{k,k}(z) - e^z| \right)^{1/k} = H \tag{14}$$

```
function n = rat_apprx(theta, alpha, alpha_0, A, t, n_0)

    % theta   = poles of the rational function r
    % alpha   = residues at these poles
    % alpha_0 = limit of r at infinity

    s = length(theta);
    A = A * t;
    n = 0 * n_0;

    for j = 1 : s
        n = n + (A - theta(j) * eye(size(A))) \ (alpha(j) * n_0);
    end
    n = 2 * real(n);
    n = n + alpha_0 * n_0;

end
```

Fig. 1. MATLAB code illustrating how a rational function $r_{k,k}$ or $r_{k-1,1}$ is applied to Eq. (1) to approximate the matrix exponential solution `n`. The input arguments are the partial fraction coefficients corresponding to the rational function, burnup matrix `A`, time step `t`, and initial composition vector `n_0`.

for any compact $K \subset \mathbb{C}$ and for any Hankel contour $\Gamma \subset \mathbb{C}\backslash\mathbb{R}_-$. However, it is worth noticing that this convergence is related to the asymptotic properties of the sequence $\{\delta_k\}$, and the accuracy of the approximation of order k is still dependent on the choice of the subset and the contour under consideration.

As previously stated, the main difficulty in using CRAM is determining the coefficients of the rational function for a given k. In principle the polynomial coefficients of $\hat{p}_k$ and $\hat{q}_k$ can be computed with Remez-type methods, but this requires delicate algorithms combined with high-precision arithmetics. Fortunately, these coefficients have been computed to a high accuracy by Carpenter, Ruttan, and Varga for approximation orders $k = 0, 1, \ldots, 30$, and they are provided in Ref. 10. In practical applications, however, CRAM approximation is usually needed in the PFD form. Although the PFD coefficients can in principle be computed from the polynomial coefficients, the computation of the polynomial roots is ill-conditioned and requires great care. The only reference providing the PFD coefficients (for approximation orders 10 and 14) is presumably Ref. 11, and the coefficients for approximation of order $k = 14$ have therefore been used in several applications including the EXPOKIT matrix exponential computing package[12] and Serpent reactor physics code.[13] However, it seems that these coefficients suffer from round-off errors and hence do not correspond to the true best approximation. Figure 2 shows the error of order 14 approximation on the negative real axis computed using two different sets of coefficients: the polynomial coefficients from Ref. 10 and the partial fraction coefficients from Ref. 11. According to theory, a necessary and sufficient condition for the best approximation is that the corresponding error function equioscillates; i.e., there exists a set of points where it attains its maximum absolute value with alternating signs. Notice that the approximation computed with the coefficients from Ref. 11 does not exhibit this behavior. In addition, these coefficients result in 10^2 times poorer accuracy than expected by theory.

To provide better accuracy, new sets of partial fraction coefficients for approximation orders $k = 14$ and $k = 16$ were computed from the polynomial coefficients provided in Ref. 10, and these coefficients are listed in Tables I and II. The computations were performed with MATLAB's Symbolic Toolbox using high-precision arithmetics with 200 digits to ensure sufficient accuracy. In Tables I and II the coefficients have been rounded off to

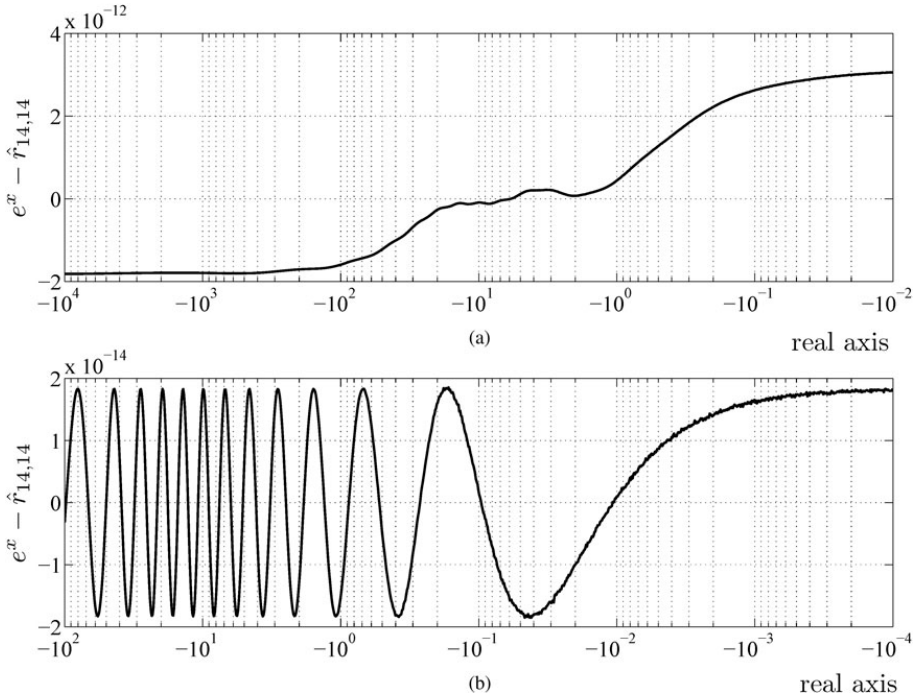

Fig. 2. Plot of $e^x - \hat{r}_{14,14}(x)$ on the negative real axis. In (a) $\hat{r}_{14,14}$ was computed based on the partial fraction coefficients from Ref. 11 and in (b) based on the polynomial coefficients from Ref. 10.

TABLE I

Partial Fraction Decomposition Coefficients for CRAM Approximation of Order 14

Coefficient	Real Part	Imaginary Part
θ_1	$-8.897\,773\,186\,468\,888\,819\,9 \times 10^{0}$	$+1.663\,098\,261\,990\,208\,530\,4 \times 10^{1}$
θ_2	$-3.703\,275\,049\,423\,448\,060\,3 \times 10^{0}$	$+1.365\,637\,187\,148\,326\,817\,1 \times 10^{1}$
θ_3	$-0.208\,758\,638\,250\,130\,125\,1 \times 10^{0}$	$+1.099\,126\,056\,190\,126\,091\,3 \times 10^{1}$
θ_4	$+3.993\,369\,710\,578\,568\,519\,4 \times 10^{0}$	$+6.004\,831\,642\,235\,037\,317\,8 \times 10^{0}$
θ_5	$+5.089\,345\,060\,580\,624\,506\,6 \times 10^{0}$	$+3.588\,824\,029\,027\,006\,510\,2 \times 10^{0}$
θ_6	$+5.623\,142\,572\,745\,977\,124\,8 \times 10^{0}$	$+1.194\,069\,046\,343\,966\,976\,6 \times 10^{0}$
θ_7	$+2.269\,783\,829\,231\,112\,709\,7 \times 10^{0}$	$+8.461\,737\,973\,040\,221\,401\,9 \times 10^{0}$
α_1	$-7.154\,288\,063\,589\,067\,285\,3 \times 10^{-5}$	$+1.436\,104\,334\,954\,130\,011\,1 \times 10^{-4}$
α_2	$+9.439\,025\,310\,736\,168\,877\,9 \times 10^{-3}$	$-1.718\,479\,195\,848\,301\,751\,1 \times 10^{-2}$
α_3	$-3.763\,600\,387\,822\,696\,871\,7 \times 10^{-1}$	$+3.351\,834\,702\,945\,010\,421\,4 \times 10^{-1}$
α_4	$-2.349\,823\,209\,108\,270\,119\,1 \times 10^{1}$	$-5.808\,359\,129\,714\,207\,400\,4 \times 10^{0}$
α_5	$+4.693\,327\,448\,883\,129\,304\,7 \times 10^{1}$	$+4.564\,364\,976\,882\,776\,079\,1 \times 10^{1}$
α_6	$-2.787\,516\,194\,014\,564\,646\,8 \times 10^{1}$	$-1.021\,473\,399\,905\,645\,143\,4 \times 10^{2}$
α_7	$+4.807\,112\,098\,832\,508\,890\,7 \times 10^{0}$	$-1.320\,979\,383\,742\,872\,388\,1 \times 10^{0}$
α_0	$+1.832\,174\,378\,254\,041\,275\,1 \times 10^{-14}$	$+0.000\,000\,000\,000\,000\,000\,0 \times 10^{0}$

TABLE II

Partial Fraction Decomposition Coefficients for CRAM Approximation of Order 16

Coefficient	Real Part	Imaginary Part
θ_1	$-1.084\,391\,707\,869\,698\,802\,6 \times 10^{1}$	$+1.927\,744\,616\,718\,165\,228\,4 \times 10^{1}$
θ_2	$-5.264\,971\,343\,442\,646\,889\,5 \times 10^{0}$	$+1.622\,022\,147\,316\,792\,730\,5 \times 10^{1}$
θ_3	$+5.948\,152\,268\,951\,177\,480\,8 \times 10^{0}$	$+3.587\,457\,362\,018\,322\,282\,9 \times 10^{0}$
θ_4	$+3.509\,103\,608\,414\,918\,097\,4 \times 10^{0}$	$+8.436\,198\,985\,884\,375\,082\,6 \times 10^{0}$
θ_5	$+6.416\,177\,699\,099\,434\,192\,3 \times 10^{0}$	$+1.194\,122\,393\,370\,138\,687\,4 \times 10^{0}$
θ_6	$+1.419\,375\,897\,185\,665\,978\,6 \times 10^{0}$	$+1.092\,536\,348\,449\,672\,258\,5 \times 10^{1}$
θ_7	$+4.993\,174\,737\,717\,996\,399\,1 \times 10^{0}$	$+5.996\,881\,713\,603\,942\,226\,0 \times 10^{0}$
θ_8	$-1.413\,928\,462\,488\,886\,211\,4 \times 10^{0}$	$+1.349\,772\,569\,889\,274\,538\,9 \times 10^{1}$
α_1	$-5.090\,152\,186\,522\,491\,565\,0 \times 10^{-7}$	$-2.422\,001\,765\,285\,228\,797\,0 \times 10^{-5}$
α_2	$+2.115\,174\,218\,246\,603\,090\,7 \times 10^{-4}$	$+4.389\,296\,964\,738\,067\,391\,8 \times 10^{-3}$
α_3	$+1.133\,977\,517\,848\,393\,052\,7 \times 10^{2}$	$+1.019\,472\,170\,421\,585\,645\,0 \times 10^{2}$
α_4	$+1.505\,958\,527\,002\,346\,752\,8 \times 10^{1}$	$-5.751\,405\,277\,642\,181\,997\,9 \times 10^{0}$
α_5	$-6.450\,087\,802\,553\,964\,659\,5 \times 10^{1}$	$-2.245\,944\,076\,265\,209\,605\,6 \times 10^{2}$
α_6	$-1.479\,300\,711\,355\,799\,971\,8 \times 10^{0}$	$+1.768\,658\,832\,378\,293\,790\,6 \times 10^{0}$
α_7	$-6.251\,839\,246\,320\,791\,889\,2 \times 10^{1}$	$-1.119\,039\,109\,428\,322\,848\,0 \times 10^{1}$
α_8	$+4.102\,313\,683\,541\,002\,127\,3 \times 10^{-2}$	$-1.574\,346\,617\,345\,546\,819\,1 \times 10^{-1}$
α_0	$+2.124\,853\,710\,495\,223\,748\,8 \times 10^{-16}$	$+0.000\,000\,000\,000\,000\,000\,0 \times 10^{0}$

20 digits. Some of the newly computed coefficients differ significantly from the ones in Ref. 11. The PFD coefficients for approximation orders $1 \leq k \leq 13$ can be computed with high accuracy by using the approximative Carathéodory–Fejér method. A MATLAB script is provided for this purpose in Ref. 14. With the help of these coefficients, the implementation of the CRAM matrix exponential method for approximation orders $k \leq 16$ is extremely straightforward; as can be seen from Eq. (10), only solving a set of linear equations is required in addition to the coefficients. In order to implement CRAM in MATLAB, only the code from Fig. 1 is needed in addition to these coefficients. CRAM is therefore a very attractive method for solving the burnup equations in reactor physics codes.

However, if higher-order approximations are desired, complications ensue. Based on our experiments, the accuracy of the polynomial coefficients provided in

Ref. 10 is not sufficient for computing the partial fraction coefficients for approximation orders higher than 16. For this reason, an alternative and easier-to-implement method is presented in Sec. II.B. This approach is based on deriving the rational approximation from quadrature formulas applied to a contour integral in the left complex plane. Although these approximations do not converge as fast as CRAM, they have the advantage that the computation of the coefficients can be done on the fly, and therefore, the accuracy of the approximation can be easily adjusted.

II.B. Rational Approximations from Contour Integrals

When the eigenvalues of the matrix $\boldsymbol{At}$ are confined to a region near the negative real axis, the computation of the matrix exponential is effectively equivalent to evaluating contour integrals of the form represented by Eq. (7). Because of the exponential factor in the integrand, the contribution to the integral decays rapidly as $\mathrm{Re}(z) \to -\infty$, and the integral can be approximated efficiently using quadrature rules. These quadrature formulas can furthermore be interpreted as rational approximations that can be used to approximate the matrix exponential.

The idea of constructing rational approximations to the exponential function from quadrature rules was recently resurfaced by Trefethen, Weideman, and Schmelzer[7] and Weideman and Trefethen.[15] In Ref. 15, two types of contours, namely, hyperbolas and parabolas, have been analyzed, and asymptotically optimal parameters for these contours have been derived by balancing the error terms related to the approximation of the contour integrals by quadrature rules. Of these contours the parabola is the simpler one, so it is considered in this paper to illustrate the method. For integrands of type (7) with singularities on the negative real axis, Weideman has proposed the parabola:

$$\phi : \mathbb{R} \to \mathbb{C} \ ,$$

$$\phi(x) = N(0.1309 - 0.1149x^2 + i0.2500x) \ , \qquad (15)$$

which yields the convergence rate $O(2.85^{-N}) = O(e^{-1.05N})$ (Ref. 7). A rational approximation obtained from this contour is applied to solving the burnup equations in Sec. III.B.

The rational approximation of the exponential function based on a contour integral can be simply constructed as

$$r(z) = \sum_{k=1}^{N} \frac{\alpha_k}{z - \theta_k} \ , \qquad (16)$$

where $\theta_k = \phi(x_k)$ are the quadrature points from the contour, and

$$\alpha_k = -\frac{h}{2\pi i} e^{\phi(x_k)} \phi'(x_k) \qquad (17)$$

are the weights of the quadrature rule, where h denotes the interval length used in the quadrature scheme. For a detailed derivation, see Refs. 1 or 7. Figure 3 shows a MATLAB code illustrating how the coefficients corresponding to the contour of Eq. (15) may be computed in practice. After obtaining the coefficients, the code shown in Fig. 1 can be used to apply the approximation to Eq. (1).

The type of convergence analysis applied in Ref. 15 is based on mathematical theorems that characterize the convergence of quadrature rules on the real line $\mathbb{R}$. For integrands that decay sufficiently fast when $x \to \pm\infty$, the rate of convergence is determined by the integrand function's region of analyticity in the vicinity of the real axis. These convergence results can then be extended to contour integrals with the help of conformal functions that map the real axis $\mathbb{R}$ onto the contour Γ under consideration (for details and further information see, e.g., Ref. 16). When the decay of the integrand is exponential, the truncation error can also be assumed to be exponential, and the quadrature scheme can be chosen to give an asymptotic convergence rate with respect to the total error.

```
function [theta, alpha, alpha_0] = quad_coeffs(k)

    x       = pi * (1 : 2 : k - 1) / k;
    theta   = k * (0.1309 - 0.1194 * x .^ 2 + 0.2500 * x * 1i);
    w_j     = k * (-2 * 0.1194 * x + 0.2500 * 1i);
    alpha   = 1i * 1 / k * exp(theta) .* w_j;
    alpha_0 = 0;

end
```

Fig. 3. MATLAB code illustrating how the partial fraction coefficients may be computed for a rational approximation derived from a quadrature formula applied to the contour of Eq. (15). The input argument `k` is the degree of approximation, and it must be an even number.

TABLE III

Accuracy of CRAM Approximation Applied to Test Case 1*

Approximation Order	Mean Error	Maximum Error	Mean Relative Error	Maximum Relative Error
2	3.3901×10^{-7}	3.3110×10^{-4}	8.3015×10^{-2}	1.9561×10^{0}
4	4.0252×10^{-9}	3.8736×10^{-6}	5.6140×10^{-3}	6.3820×10^{-1}
6	4.7339×10^{-11}	4.5163×10^{-8}	2.2452×10^{-4}	3.8184×10^{-2}
8	5.5808×10^{-13}	5.2486×10^{-10}	7.1664×10^{-6}	1.5762×10^{-3}
10	6.5685×10^{-15}	6.0944×10^{-12}	1.9529×10^{-7}	5.1640×10^{-5}
12	7.6474×10^{-17}	6.9690×10^{-14}	4.7280×10^{-9}	1.4323×10^{-6}
14	9.5452×10^{-19}	9.5339×10^{-16}	1.0384×10^{-10}	3.4990×10^{-8}
16	2.0748×10^{-19}	1.6377×10^{-16}	2.1196×10^{-12}	7.7286×10^{-10}

*The errors were computed against a reference solution calculated with high-precision arithmetics.

III. APPLICATION TO BURNUP EQUATIONS

To evaluate the accuracy and study the convergence of the rational approximations, they were applied to two large burnup systems, and the solutions were compared to highly accurate reference solutions that were computed using MATLAB's Symbolic Toolbox and high-precision arithmetics. Both test cases represented an infinite pressurized water reactor pin-cell lattice. The time step used in the computations was 125 days, which is of the same order as the practical maximum time step. In the first test case, the fuel was irradiated to 25 MWd/kg U burnup, and the number of nuclides was 1532.[a] The second test case contained 1290 nuclides, and the burnup system was formed for both fresh fuel and fuel irradiated to 20 MWd/kg U burnup. In the first test case, the burnup matrix norm was of the order of 10^{21}, and in the second test case, it was of the order of 10^{16}.

III.A. Chebyshev Rational Approximation Method

The CRAM approximations of orders $k \leq 12$ were computed using the Carathéodory–Fejér method implemented as a MATLAB script in Ref. 14, and approximations of orders 14 and 16 were formed using the newly computed coefficients provided in Tables I and II. The numerical results for the first test case are shown in Table III and Fig. 4. Note that no nuclides were excluded from these results; i.e., even the nuclides with extremely small number densities have been included in the mean and maximum relative errors. The improvement in accuracy gained by recomputing the CRAM coefficients for approximation order 14 is illustrated in Fig. 5, which shows that the relative error is roughly 10^2 times smaller with the new coefficients provided in Sec. II.A.

[a]This is the same system as test case 1 in Ref. 1.

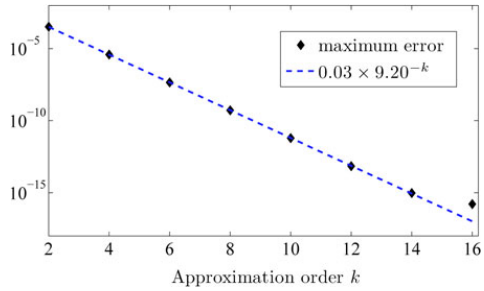

Fig. 4. Convergence of CRAM approximation applied to the first test case with 1532 nuclides.

Overall, it can be seen from the results that even relatively small CRAM approximation orders give remarkably accurate solutions to the first test case. From a practical point of view, orders of approximation as small as $k = 2$ or $k = 4$ could be used for solving the burnup equations for this test case without compromising the accuracy of the entire burnup calculation. For comparison, arguably the most established matrix exponential method, the Padé approximation (with scaling and squaring) suffered a total breakdown when applied to the same test case,[1] which serves well to illustrate the special characteristics of the burnup equations.

To study the convergence rate of CRAM, the error of the nuclide concentrations in the maximum norm was plotted against the approximation order, and the results are shown in Fig. 4. It can be seen that the convergence of the method is clearly geometric. The actual convergence rate can be estimated by performing a least-squares fit to the points, suggesting that the convergence is of the order of $O(9.20^{-N})$ for $k = 2, \ldots, 14$. Notice that this is very close to the theoretical asymptotic convergence rate, giving further proof of this method's

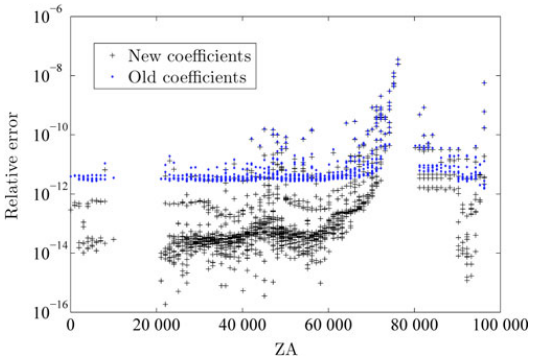

Fig. 5. Absolute values of the errors computed with CRAM approximation of order 14 using the PFD coefficients from Ref. 11 (dots) and using the recently computed coefficients (crosses). $ZA = 1000 \times Z + A$, where Z is the atomic number and A is the mass number of the nuclide.

suitability in the context of burnup equations. For $k = 16$, the round-off error began to affect the result, and hence, the convergence was slightly slower. It should be noted that as an approximation of order k requires $k/2$ linear solves, Fig. 4 also illustrates the computational cost of CRAM versus the maximum error.

It has been observed that the accuracy of CRAM depends relatively little on the fuel composition or neutron spectrum corresponding to the burnup matrix.[1,5] However, it has been noticed that CRAM generally produces less accurate results for fresh fuel than for depleted fuel.[5] When the fuel is fresh, only a few elements of $\boldsymbol{n}_0$ in Eq. (1) are nonzero, and all the nuclides are produced solely from these initial nuclides. For a large part of nuclides, this means both long and complex transmutation chains. It has been suggested that the observed reduction in accuracy is due to this and that the errors might be averaged out in depleted cases, where there are more initial nuclides.[5] To further quantify this effect, CRAM was applied to fresh fuel and fuel irradiated to 20 MWd/kg U burnup. This setup is referred to as test case 2 in this paper, and the numerical results are shown in Table IV. It can be seen from these results that although the absolute errors for both cases are very similar, the relative errors related to the fresh fuel computation are roughly 10^2 times greater than the ones corresponding to the depleted fuel case.

To further investigate this issue, a CRAM approximation of order 16 was explicitly computed for the burnup matrices representing fresh and depleted fuel in test case 2, and the elements of the approximation matrices $\hat{\boldsymbol{E}} = \hat{r}_{16,16}(\boldsymbol{A}t)$ were compared to the reference matrix exponentials $\boldsymbol{E} = e^{\boldsymbol{A}t}$ computed using MATLAB's Symbolic toolbox with high-precision arithmetics. This comparison showed that the accuracy of the approximation is of

TABLE IV

Accuracy of CRAM Approximation Applied to Test Case 2*

Approximation Order	Mean Error	Maximum Error	Mean Relative Error	Maximum Relative Error
Fresh fuel				
2	4.0451×10^{-7}	3.3117×10^{-4}	4.2189×10^{0}	2.3900×10^{3}
4	4.8034×10^{-9}	3.8747×10^{-6}	4.4645×10^{-1}	1.5813×10^{2}
6	5.6347×10^{-11}	4.5181×10^{-8}	4.0548×10^{-2}	1.9953×10^{1}
8	6.5984×10^{-13}	5.2511×10^{-10}	4.5634×10^{-3}	3.9451×10^{0}
10	7.6769×10^{-15}	6.0979×10^{-12}	5.5147×10^{-5}	5.3277×10^{-2}
12	8.9182×10^{-17}	6.9653×10^{-14}	1.6627×10^{-6}	1.7515×10^{-3}
14	1.0448×10^{-18}	8.1185×10^{-16}	4.2531×10^{-8}	4.7451×10^{-5}
16	3.0297×10^{-19}	2.6715×10^{-16}	9.5605×10^{-10}	1.1091×10^{-6}
Depleted fuel				
2	4.0231×10^{-7}	3.3113×10^{-4}	4.7487×10^{-2}	1.7192×10^{0}
4	4.7767×10^{-9}	3.8741×10^{-6}	1.4957×10^{-3}	8.8440×10^{-2}
6	5.6312×10^{-11}	4.5172×10^{-8}	1.3526×10^{-3}	1.6976×10^{0}
8	6.6458×10^{-13}	5.2498×10^{-10}	1.1667×10^{-3}	1.5040×10^{0}
10	7.8380×10^{-15}	6.0962×10^{-12}	1.2991×10^{-6}	1.6558×10^{-3}
12	9.1377×10^{-17}	6.9757×10^{-14}	2.7517×10^{-10}	5.9628×10^{-8}
14	9.7118×10^{-19}	6.7307×10^{-16}	4.6712×10^{-12}	1.1744×10^{-9}
16	3.3790×10^{-19}	3.6082×10^{-16}	2.5916×10^{-13}	2.2667×10^{-10}

*The errors were computed against a reference solution calculated with high-precision arithmetics.

the same order for burnup matrices representing fresh and depleted fuel; i.e., the reduced accuracy is not related to the properties of the burnup matrix. However, a closer examination showed a clear trend between the importance of particular nuclide chains and the relative accuracy of the corresponding matrix elements. That is, the matrix elements corresponding to less important chains are also computed with poorer relative accuracy with CRAM. This is mathematically reasonable since the relative accuracy of CRAM approximation $\hat{r}(x)$ diminishes as $x \to -\infty$.[b]

It follows that the reduced relative accuracy observed in fresh fuel cases is indeed related to the less important nuclide chains. The approximation error for nuclide concentration n_i may be written

$$\varepsilon_i = \left| \sum_{j=1}^{n} (E_{ij} - \hat{E}_{ij}) n_{0j} \right| \quad . \tag{18}$$

When the fuel is fresh, most of the initial concentrations are zero, and hence, the error ε_i is determined by the accuracy of the few matrix elements $\hat{E}_{ij}$ that correspond to the chains originating from these nuclides. As an example, consider the curium isotope ^{246}Cm, which causes the maximum relative error 1.1091×10^{-6} in CRAM approximation of order 16 in test case 2. There are 61 nuclides ranging from ^{232}Th to ^{245}Cm that may contribute to the concentration of this nuclide, and in terms of the burnup matrix alone, nuclides corresponding to simple and short nuclide chains are the most significant ones. The theoretical contribution attributable to such chains is as much as 10^8 to 10^{26} times greater than the contribution originating from uranium isotopes. When the fuel is fresh, however, the uranium isotopes are the only nuclides having nonzero initial concentrations. The reduced relative accuracy is a direct consequence of this. In this example case, the only relevant chain actually originates from ^{238}U, and the relative error corresponding to this chain is of the order of 10^{-6}, which is in accordance with the relative error of the result. For comparison, the relative error for this nuclide is 6.2857×10^{-11} in the depleted fuel case.

As no nuclides were excluded from the results shown in Table IV, the relative maximum errors may reflect individual concentrations corresponding to exotic nuclides or concentrations arbitrarily close to zero. In the depleted case, for example, all relative errors of order 10^0 occur for the unstable dysprosium isotope ^{155}Dy, whose atomic fraction is $\sim 10^{-24}$. The relative errors corresponding to CRAM approximation of order 6 are shown in Fig. 6. It can be seen from Fig. 6 that all other nuclides have relative errors that are several orders of magnitude smaller. The least unlikely nuclide chains for ^{155}Dy originate from fission product nuclides, and hence, the reduced relative accuracy is observed in depleted fuel cases with small approximation orders. Notice that for higher approximation orders the relative errors are also consistently extremely small.

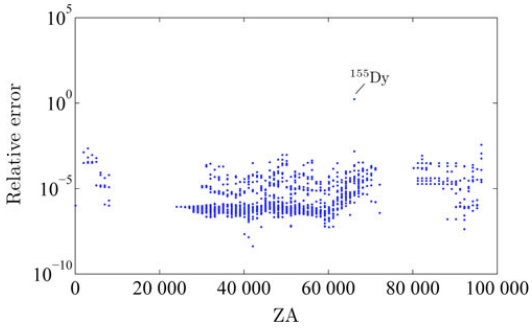

Fig. 6. Relative errors related to CRAM approximation of order 6 applied to the depleted fuel burnup equations in test case 2. The errors were computed against a reference solution calculated with high-precision arithmetics.

From a practical point of view, the results computed with CRAM approximation are remarkably accurate in both test cases. The cross-section uncertainties related to the transmutation reaction rates (i.e., burnup matrix elements) usually range from 10^{-2} to 10^{-1} in typical reactor physics calculations. When Monte Carlo methods are used, the statistical error also reduces the accuracy of the burnup matrix elements. With all this taken into consideration, even approximation orders as low as $k = 6$ could be used without compromising the total accuracy of a burnup calculation.

In addition to accuracy, another advantage of CRAM is its low computational cost. The order of approximation can be easily adjusted to suit needs for accuracy or speed. As can be seen from Eq. (10), the computation time is directly proportional to the approximation order k. If the sparsity pattern of the burnup matrix is properly exploited, the approximations can be computed with remarkably short calculation times. For example, the computation of an approximation of order $k = 14$ for a test case with about 1500 nuclides takes less than a tenth of a second in the Serpent reactor physics code.[1,5]

III.B. Rational Approximation from Quadrature Formula

As explained in Sec. II.B, rational approximations to the matrix exponential can be derived from quadrature rules applied to contour integrals. Here, this technique is illustrated by using the contour of Eq. (15), which is of the form

$$\phi(x) = k(\alpha - \beta x^2 + i\gamma x) \quad , \tag{19}$$

[b]Notice that it is impossible to derive best approximations with respect to the relative error.

where $\alpha, \beta, \gamma > 0$ and $x \in \mathbb{R}$. In the context of burnup equations, the integral under consideration is of the form of Eq. (7) and can be written

$$(e^{At})_{kl} = \frac{1}{2\pi i} \int_{\Gamma} e^{z} R_{kl}(z)\, dz \ ,$$

$$= \frac{1}{2\pi i} \int_{-\infty}^{\infty} e^{\phi(x)} R_{kl}(\phi(x))\phi'(x)\, dx \ . \quad (20)$$

The numerical results for test cases 1 and 2 are shown in Tables V and VI, respectively. By comparing these results to those obtained with CRAM, it can be seen that this approach is well-suited for solving the burnup equations. Notice that although the convergence of this method is much slower than the convergence of CRAM, the results are qualitatively very similar. The results for test case 1 are again more accurate, and the large overall relative errors in the depleted case of test case 2 result from the single dysprosium isotope ^{155}Dy.

The total error in the results consists of the theoretical approximation error and the round-off error related to the finite precision arithmetics used in the computations. When the approximation order is greater than 36, the round-off errors begin to contribute to the accuracy, after which the maximum error does not significantly diminish. However, it can be seen from the results that with sufficiently high quadrature orders, it is possible to obtain at least ten correct digits for all nuclide concentrations, which can be considered quite remarkable. This is a sufficient accuracy for all considerable burnup calculations, and hence, this approach could be suitable for computing reference solutions for other numerical methods.

Another advantage related to rational approximations derived from quadrature formulas is that it is possible to derive conservative convergence estimates for different eigenvalue distributions. Because of the exponential factor in the integral, the convergence of quadrature rules applied to the computation of this integral is determined by the region of analyticity of the integrand function continued to the complex plane. If the integrand function is analytic in the infinite strip

$$D_S = \{z \in \mathbb{C} \,|\, z = x + iy, -d < y < c\} \ , \quad (21)$$

the convergence with respect to the interval length h is of order $O(e^{-2\pi c/h} + e^{-2\pi d/h})$ as $h \to 0$. In addition, the truncation error can be estimated to be of order $O(|g(hN)|)$ as $N \to \infty$ (Ref. 15). The integrand in Eq. (20) can be written

$$g(x + i\omega) = e^{\phi(x+i\omega)} R_{kl}(\phi(x + i\omega))\phi'(x + i\omega) \ . \quad (22)$$

Because the only singularities of the function g are the poles of the rational function R_{kl}, the singularity distribution of g consists of points that are mapped onto the eigenvalues of $\boldsymbol{A}t$. Since the eigenvalues of burnup matrices are all confined to a region near the negative real axis, the convergence rate should be dominated by the factor $O(e^{2\pi c/h})$, where the constant c depends directly on the magnitudes of the imaginary parts of the eigenvalues.

TABLE V

Accuracy of Quadrature Approximation Applied to Test Case 1*

Approximation Order	Mean Error	Maximum Error	Mean Relative Error	Maximum Relative Error
4	1.4726×10^{-6}	1.4526×10^{-3}	2.4786×10^{-1}	1.4951×10^{1}
8	2.2621×10^{-8}	2.1925×10^{-5}	3.5720×10^{-2}	4.7308×10^{0}
12	3.3657×10^{-10}	3.2264×10^{-7}	2.0632×10^{-3}	4.1914×10^{-1}
16	5.0968×10^{-12}	4.8108×10^{-9}	8.8273×10^{-5}	2.2515×10^{-2}
20	7.6246×10^{-14}	7.0702×10^{-11}	3.1568×10^{-6}	9.4284×10^{-4}
24	1.1477×10^{-15}	1.0443×10^{-12}	9.9082×10^{-8}	3.3211×10^{-5}
28	1.7191×10^{-17}	1.5314×10^{-14}	2.8082×10^{-9}	1.0277×10^{-6}
32	3.3333×10^{-19}	2.7756×10^{-16}	7.4740×10^{-11}	2.8729×10^{-8}
36	1.2452×10^{-19}	1.5266×10^{-16}	1.7969×10^{-12}	7.3949×10^{-10}
40	3.2204×10^{-20}	2.4286×10^{-17}	4.4542×10^{-14}	1.7772×10^{-11}
44	2.6421×10^{-19}	3.6082×10^{-16}	5.9574×10^{-15}	4.0187×10^{-13}
48	7.0675×10^{-19}	8.8124×10^{-16}	1.2903×10^{-14}	2.3370×10^{-13}
52	6.5418×10^{-19}	4.9266×10^{-16}	1.6877×10^{-14}	2.4932×10^{-13}
56	2.8871×10^{-18}	3.0947×10^{-15}	5.1656×10^{-14}	1.4797×10^{-13}
60	3.7784×10^{-18}	4.7531×10^{-15}	6.2409×10^{-14}	2.2402×10^{-13}

*The errors were computed against a reference solution calculated with high-precision arithmetics.

TABLE VI

Accuracy of Quadrature Approximation Applied to Test Case 2*

Approximation Order	Mean Error	Maximum Error	Mean Relative Error	Maximum Relative Error
Fresh fuel				
4	1.7504×10^{-6}	1.4528×10^{-3}	3.0494×10^{1}	1.5087×10^{4}
8	2.6904×10^{-8}	2.1931×10^{-5}	3.5070×10^{0}	1.5082×10^{3}
12	3.9894×10^{-10}	3.2276×10^{-7}	4.1829×10^{-1}	3.1038×10^{2}
16	5.9978×10^{-12}	4.8131×10^{-9}	2.7340×10^{-2}	2.2658×10^{1}
20	8.9027×10^{-14}	7.0744×10^{-11}	3.7238×10^{-2}	4.6645×10^{1}
24	1.3364×10^{-15}	1.0451×10^{-12}	9.4550×10^{-4}	1.1684×10^{0}
28	1.9949×10^{-17}	1.5307×10^{-14}	1.2510×10^{-6}	1.4528×10^{-3}
32	3.5814×10^{-19}	2.8449×10^{-16}	3.4718×10^{-8}	4.1451×10^{-5}
36	1.7146×10^{-19}	1.8735×10^{-16}	8.7330×10^{-10}	1.0632×10^{-6}
40	8.0602×10^{-20}	8.3267×10^{-17}	2.0294×10^{-11}	2.5030×10^{-8}
44	2.9928×10^{-19}	3.6776×10^{-16}	5.1227×10^{-13}	5.4906×10^{-10}
48	6.8488×10^{-19}	5.3429×10^{-16}	2.8378×10^{-13}	3.4522×10^{-10}
52	2.0623×10^{-19}	2.3592×10^{-16}	1.2807×10^{-13}	1.5392×10^{-10}
56	4.3947×10^{-18}	3.8441×10^{-15}	4.3086×10^{-13}	5.2005×10^{-10}
60	2.8439×10^{-18}	1.8804×10^{-15}	1.9729×10^{-13}	1.8854×10^{-10}
Depleted fuel				
4	1.7491×10^{-6}	1.4527×10^{-3}	1.5504×10^{-1}	2.2859×10^{0}
8	2.6863×10^{-8}	2.1928×10^{-5}	8.1883×10^{-3}	5.7021×10^{-1}
12	4.0052×10^{-10}	3.2270×10^{-7}	9.2944×10^{-4}	8.4914×10^{-1}
16	6.0748×10^{-12}	4.8119×10^{-9}	1.7091×10^{-4}	2.1048×10^{-1}
20	9.0998×10^{-14}	7.0723×10^{-11}	3.0635×10^{-4}	3.9490×10^{-1}
24	1.3717×10^{-15}	1.0447×10^{-12}	5.3881×10^{-3}	6.9504×10^{0}
28	2.0572×10^{-17}	1.5314×10^{-14}	5.8635×10^{-7}	7.5604×10^{-4}
32	4.2531×10^{-19}	2.8449×10^{-16}	2.1854×10^{-12}	6.4981×10^{-10}
36	1.6180×10^{-19}	1.4572×10^{-16}	5.2742×10^{-14}	1.3987×10^{-11}
40	2.0604×10^{-19}	1.9429×10^{-16}	2.3263×10^{-14}	2.4936×10^{-11}
44	3.9955×10^{-19}	4.1633×10^{-16}	8.4011×10^{-15}	3.8230×10^{-12}
48	9.5099×10^{-19}	7.3552×10^{-16}	6.8619×10^{-14}	6.7985×10^{-11}
52	1.2022×10^{-18}	1.0270×10^{-15}	1.1006×10^{-13}	1.1291×10^{-10}
56	3.7045×10^{-18}	3.0531×10^{-15}	3.8593×10^{-13}	4.0835×10^{-10}
60	4.3363×10^{-18}	3.2127×10^{-15}	1.7439×10^{-13}	1.3113×10^{-10}

*The errors were computed against a reference solution calculated with high-precision arithmetics.

As the value of c is increased from zero, the contour $\phi(x + ic)$ narrows down, and it approaches the negative real axis when c reaches the value $c = \gamma/(2\beta)$. By choosing $h = 2\pi/k$ the convergence of order

$$O(e^{-ck}) = O(e^{-1.0469k}) = O(2.85^{-k})$$

is achieved for the case, where all the eigenvalues lie on the negative real axis. This result can be generalized for burnup matrices having eigenvalues near the negative real axis. It seems to be sufficient to require that $\mathrm{Im}\,\phi(x + ic) > 0.1$ at the origin. From this requirement it is straightforward to derive the following equation for c:

$$k^2\beta^{-1}(\alpha + \beta c^2 - \gamma c)(\gamma - 2\beta)^2 = 0.1 \quad . \qquad (23)$$

In practice, the convergence of the quadrature scheme can be estimated by solving this equation for sufficiently small k with respect to the number of nodes used in the calculations. For quadrature order $k = 4$, Eq. (23) gives $c = 0.93$, and thus, the convergence should be at least of the order of $O(2.53^{-N})$ for this singularity distribution. Note that this estimate corresponds to the worst-case scenario, where practically all the singularities are located on the contour $\phi(x + ic)$. In practice, however, only a fraction of the eigenvalues of $\boldsymbol{A}t$ lie off the real axis, and the majority of them have imaginary parts smaller than 10^{-1}.

The maximum norm error in test case 1 is plotted against the quadrature order in Fig. 7. As in the case of CRAM, the computational cost is directly proportional to the approximation order k, and consequently, Fig. 7 also illustrates the computational cost of the quadrature scheme versus the maximum error. A least-squares fit to the points $k \leq 32$ suggests that the convergence is

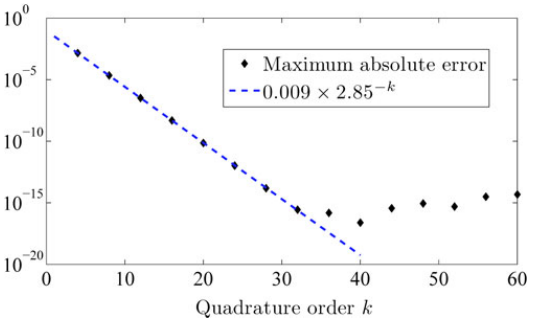

Fig. 7. Convergence of the quadrature approximation applied to the first test case with 1532 nuclides.

actually of the order of 2.85^{-k} until round-off errors begin to dominate the error. Judging by this, the few non-real eigenvalues having small imaginary parts are not relevant when using rational approximations that are accurate near the negative real axis. This gives further credibility to the suitability of this type of approximation in the context of burnup equations.

IV. SUMMARY AND CONCLUSIONS

The computation of the matrix exponential has been considered challenging in the context of burnup equations because the magnitudes of the transmutation and decay constants vary extensively, making the problem extremely stiff. Traditionally, these difficulties have been avoided by using simplified models or by excluding the short-lived nuclides from the computation and treating them separately. However, it was recently discovered by the author that the eigenvalues of burnup matrices are confined to a region near the negative real axis, and this property can be exploited by using rational approximations that are highly accurate there.[1] This allows for simultaneously solving the entire system containing thousands of nuclides in an accurate and efficient manner.

In this paper, two different types of rational approximation were considered for computing the exponential of a burnup matrix—the previously introduced CRAM approximation, which can be characterized as the best rational approximation on the negative real axis, and an approximation method based on quadrature rules applied to a contour integral around the negative real axis. Both methods are very straightforward to implement because only a few lines of program code and a function for solving a set of linear equations are required.

The motivation for introducing the latter method was that although CRAM appears to be the most efficient method for computing the matrix exponential of a burnup matrix, its higher-order partial fraction coefficients are difficult to obtain. In addition, the previous literature values[11] for order-14 coefficients were discovered to contain round-off errors that resulted in relative accuracy two orders of magnitude poorer than expected by theory. To rectify this, new partial fraction coefficients were computed and reported for CRAM approximation orders 14 and 16. Although the accuracy of order-16 CRAM is more than sufficient for most thinkable burnup applications, and it can be readily implemented using the coefficients provided in this paper, the proposed quadrature approach is an attractive alternative. Its benefits include flexibility in terms of balancing efficiency and accuracy for routine burnup calculations with geometric convergence properties as well as the possibility of computing reference results with extremely high accuracy, limited only by the available arithmetic precision. Such reference results offer a distinct benchmark that other matrix exponential methods can be compared against in the future. In this paper, the reference solutions for the two test cases were painstakingly computed using high-precision arithmetics, and they confirmed the accuracy of the high-order quadrature approximations.

Regarding CRAM, this paper is the first to discuss approximation orders other than 14 and to study the convergence properties of the method in the context of burnup equations. In addition, the sources of approximation error were analyzed, and the observed differences in resulting accuracy for fresh and depleted fuel were explained. The new discoveries and observations fully support our previous assessment of CRAM being capable of providing a very accurate and robust solution to the burnup equations with a very low computational cost.

The main motivation and context for the presented work is the Serpent reactor physics code,[13] which was the first of its kind to implement CRAM specifically for solving the burnup equations. As previously mentioned, the literature values for the order-14 CRAM coefficients, on which the implementation of CRAM in the current release version of Serpent is also based, are inaccurate, and future development plans for Serpent include implementing CRAM of orders 2 through 16 using the correct coefficients.

REFERENCES

1. M. PUSA and J. LEPPÄNEN, "Computing the Matrix Exponential in Burnup Calculations," *Nucl. Sci. Eng.*, **164**, 140 (2010).

2. 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 Evaluations," Vol. II, Sec. F7, ORNL/TM-2005/39, Oak Ridge National Laboratory/U.S. Nuclear Regulatory Commission (Nov. 2006).

3. 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).

4. C. MOLER and C. VAN LOAN, "Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later," *SIAM Rev.*, **45**, 3 (2003).

5. A. ISOTALO and P. A. AARNIO, "Comparison of Depletion Algorithms," *Ann. Nucl. Energy*, **38**, *2–3*, 261 (2011).

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

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

8. A. A. GONCHAR and E. A. RAKHMANOV, "Equilibrium Distributions and Degree of Rational Approximation of Analytic Functions," *Math. USSR Sb.*, **62**, 2 (1989).

9. H. STAHL and T. SCHMELZER, "An Extension of the 1/9-Problem," *J. Comput. Appl. Math.*, **233**, 821 (2009).

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

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

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

13. J. LEPPÄNEN, "PSG2/Serpent—a Continuous-Energy Monte Carlo Reactor Physics Burnup Calculation Code," VTT Technical Research Centre of Finland (Nov. 2008); http://montecarlo.vtt.fi (current as of October 26, 2010).

14. T. SCHMELZER, "Carathéodory–Fejér Approximation," Matlab Central (Nov. 2008); http://www.mathworks.com/matlabcentral (current as of October 26, 2010).

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

16. J. LUND and K. L. BOWERS, *Sinc Methods for Quadrature and Differential Equations*, Society for Industrial and Applied Mathematics (1992).