INL/JOU-20-57922-Revision-0

Global Error Analysis of the Chebyshev Rational Approximation Method

January 2021

Olin Calvin, Barry Ganapol, Sebastian Schunert

Changing the World's Energy Future

INL is a U.S. Department of Energy National Laboratory operated by Battelle Energy Alliance, LLC

INL/JOU-20-57922-Revision-0

Global Error Analysis of the Chebyshev Rational Approximation Method

Olin Calvin, Barry Ganapol, Sebastian Schunert

January 2021

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Global Error Analysis of the Chebyshev Rational Approximation Method

Olin Calvin[1], Sebastian Schunert[2], Barry Ganapol[3]

[1]*Department of Nuclear Engineering, University of Idaho*
[2]*Nuclear Science and Technology Directorate, Idaho National Laboratory*
[3]*Department of Aerospace and Mechanical Engineering, The University of Arizona.*

Abstract

The Chebyshev rational approximation method (CRAM) has become a widely adopted method for solving nuclear depletion problems. Therefore, understanding CRAM's accuracy is important for the safe operation of nuclear power plants. This article performs a global error analysis of CRAM and finds that, as the length of the time step approaches zero, the relative error measured between the exact and CRAM solutions at a fixed end time approaches one and infinity for even and odd orders, respectively; for intermediate time step sizes, a minimum in relative error is observed. We show that the reason for CRAM's behavior is that the method is inconsistent. Two best practices for using CRAM, derived from these results, are: (1) use CRAM order 16 or higher, (2) if necessary, increase the CRAM order when multiphysics coupling requires smaller time steps.

Keywords: Chebyshev rational approximation method, Rattlesnake code, doubling method, nuclear depletion

1. Introduction

The Chebyshev rational approximation method (CRAM) is an algorithm for accurately and efficiently evaluating the matrix exponential. In the context of radiological and nuclear engineering, CRAM is widely used for solving the Bateman equations [1]. It was proposed for this purpose in the seminal work of Pusa [2, 3, 4, 5], whose approach rests on the work of Schmelzer and

Trefethen in evaluating matrix functions [6]. In contrast to conventional methods (e.g. implicit ordinary differential equation (ODE) integrators [7] and the transmutation trajectory analysis (TTA) method [8]) CRAM formally solves the Bateman equations by introducing the matrix exponential, thus changing the problem from integrating a stiff system of ODEs to approximating a matrix function.

CRAM is a method for approximating the matrix exponential of the isotopic depletion and decay matrix. The matrix exponential is difficult to approximate for general matrices [9], and most approximations such as a Taylor series or Padé approximations work only if the matrix norm is small [3]. Pusa noticed that the eigenvalues of burnup matrices are confined to the vicinity of the negative real axis [2], making CRAM a suitable candidate for computing the matrix exponential of the Bateman equations. Compared to standard Padé approximations, CRAM is the best rational approximation near the negative real axis [2]. For ensuring this property, CRAM's coefficients must be carefully computed; Pusa uses the Remez algorithm for this purpose [2].

A sequence of papers establishes that CRAM solves decay/transmutation and nuclear depletion problems efficiently and accurately. In Ref. [2], two examples resembling infinite pressurized water reactor lattices are considered. In both cases, fuel is irradiated up to 25 MWd/kg. The first example contains 219 isotopes, while the second contains 1,532 isotopes; 125-day time steps are used. The numerical results indicate that CRAM solves Bateman equations robustly and accurately regardless of problem size. "The method is remarkably efficient" because it solves the larger problem 260 times faster than TTA [10]. Isotalo compares two variants of the CRAM algorithm with three variants of the TTA algorithm and finds all contestants to be sufficiently accurate, though CRAM was always much faster than TTA. However, Isotalo notes that CRAM produces less accurate results for fresh fuel than for burned fuel. In Ref. [3], Pusa addresses this problem by providing highly precise CRAM coefficients for CRAM orders 14 and 16; in addition, she demonstrated an exponential convergence of the CRAM method corresponding with increasing the order. Isotalo

and Wieselquist [11] compare CRAM with the MATREX solver in ORIGEN and find that CRAM is orders of magnitude more accurate while being several times faster. Isotalo explains the difference in CRAM's accuracy for fresh and burned fuel [12]: CRAM is accurate if the change in any isotope density within a single time step is not too large; consequently, errors are dominated by a few fast transitions. Results for burned fuel are more accurate relative to fresh fuel, because a broader range of transitions contributes to each daughter isotope and allows for errors from short transitions to be washed out. Reference [12] shows that dividing time steps into substeps can improve CRAM's accuracy in fresh fuel depletion and decay calculations. Numerical benchmarking for a wide range of problems validates CRAM's substantial advantages over traditional Bateman equation solvers.

Since its introduction to nuclear applications, CRAM has been adopted and implemented by many widely used codes. The accuracy, efficiency, and simplicity of implementing CRAM drives its widespread deployment in the nuclear engineering community. Among others, CRAM is available for solving the Bateman equations in the multiphysics radiation transport code Rattlesnake [13], the Monte-Carlo code Serpent [14, 15], the nuclear engineering toolkit PyNE [16], and ORIGEN [11]. From this viewpoint understanding CRAM's advantages and pitfalls is important for ensuring accurate simulation of nuclear transmutation and decay problems.

In nuclear reactor depletion calculations, the Bateman equations are coupled with the neutron transport or diffusion equation. This coupling makes the Bateman equations nonlinear. In contrast to linear problems, the matrix exponential method does not provide the exact solution to nonlinear problems even if the matrix exponential is evaluated exactly [6]. The solution of nonlinear problems must usually proceed in sufficiently short time steps; the nonlinearity is addressed over each of these time steps by a variety of method (e.g. the exponential Rosenbrock method [17]). In nuclear reactor depletion calculation, predictor-corrector schemes are most commonly used for solving the coupled neutronics-depletion problem with CRAM [18, 19, 20]. In these schemes, the

depletion matrix is assumed to be constant over time steps. For guaranteeing accuracy, these time steps have to be sufficiently small. This raises the question of how CRAM's accuracy is affected when the time step is constrained by the nonlinearity of the coupled physics.

This work reports shortcomings of CRAM that can lead to large relative errors when the time step is short and the CRAM order is too small. Under these conditions conventional numerical methods are expected to give reasonably accurate answers. The root cause for potentially inaccurate CRAM results is that the error in the computed isotopics after a single time step does not asymptote to zero as the length of the time step is reduced to zero; in other words, CRAM is not consistent. This inconsistency causes the global discretization error (i.e. the error at some fixed end time) to increase to one or be unbounded as the number of time steps is increased. In situations where the time step is constrained by coupled physics, this behavior could lead to large errors. Fortunately, the results of this work provide guidance on how to avoid degradation of accuracy. We provide two best practices for using CRAM motivated by the theoretical and numerical results of this work.

The paper is organized as follows: Sec. 2 introduces CRAM, Sec. 3 performs a local truncation error analysis of CRAM (error after a single time step), Sec. 4 performs a global truncation error analysis of CRAM (error at a fixed end time divided into N time steps), Sec. 5 compares CRAM's error with a reference solution obtained using the doubling method for a light water reactor pin-cell problem, and Sec. 6 summarizes the paper and draws conclusions.

2. The Chebyshev Rational Approximation Method

This section briefly describes CRAM. For a thorough discussion, refer to Refs. [2, 3, 4, 5]. The Bateman equations describe the changes in isotope number density caused by decay and transmutation. If the number densities of the M relevant isotopes are collected in vector $\vec{n}$, the Bateman equations can be written

as:

$$\frac{d\vec{n}}{dt} = \boldsymbol{A}\vec{n}(t)$$
$$\vec{n}(0) = \vec{n}_0, \tag{1}$$

where $\boldsymbol{A}$ is the $M \times M$ burnup matrix containing the decay and transmutation coefficients of the isotopes under consideration. The entries of matrix $\boldsymbol{A}$ are described in detail in Ref. [2]. The depletion matrix $\boldsymbol{A}$ is constant if the neutron flux does not change with time (this applies in particular to pure decay problems where the neutron flux is zero).

The solution of the burnup equations can be formally written as:

$$\vec{n}(t) = \exp\left(\boldsymbol{A}t\right)\vec{n}_0, \tag{2}$$

where $\exp\left(\boldsymbol{A}t\right)$ is the matrix exponential. The matrix exponential can be represented using a Taylor series expansion:

$$\exp\left(\boldsymbol{A}t\right) = \sum_{l=0}^{\infty} \frac{1}{l!}(\boldsymbol{A}t)^l. \tag{3}$$

For a numerical evaluation of the matrix exponential, the Taylor series is not suitable because the infinite series may converge slowly [2, 9].

CRAM approximates the matrix exponential via a rational approximation given by:

$$\exp\left(\boldsymbol{A}t\right)\vec{n}_0 \approx \alpha_{0,K}\vec{n}_0 + 2\mathrm{Re}\left[\sum_{j=1}^{K/2} \alpha_{j,K}(\boldsymbol{A}t - \theta_{j,K}\boldsymbol{I})^{-1}\right]\vec{n}_0 = \boldsymbol{C}(t)\vec{n}_0$$
$$\boldsymbol{C}(t) = \left\{\alpha_{0,K}\boldsymbol{I} + 2\mathrm{Re}\left[\sum_{j=1}^{K/2} \alpha_{j,K}(\boldsymbol{A}t - \theta_{j,K}\boldsymbol{I})^{-1}\right]\right\}, \tag{4}$$

where K is the CRAM order, and $\alpha_{j,K}$ and $\theta_{j,K}$ are the complex partial fraction decomposition constants taken from Serpent [15]. The coefficients for orders $K = 1, 2, 3, 4, 5, 8, 15$, and 16 are listed in Tables 3-10. CRAM computes an approximation of $\vec{n}'$ to the exact number densities $\vec{n}$ at time t by:

$$\vec{n}'(t) = \boldsymbol{C}(t)\vec{n}_0. \tag{5}$$

For casting Eq. 4 into a form similar to that of Eq. 3, we use the well-known expansion of the inverse of a sum of matrices:

$$(\boldsymbol{P}+\boldsymbol{Q}\epsilon)^{-1}=\sum_{l=0}^{\infty}(-1)^{l}\,\epsilon^{l}\boldsymbol{P}^{-1}\left(\boldsymbol{Q}\boldsymbol{P}^{-1}\right)^{l}, \tag{6}$$

where $\boldsymbol{P}$ is an invertible $M\times M$ matrix, $\boldsymbol{Q}$ is an $M\times M$ matrix, and ϵ must be sufficiently small [21] (this does not pose a problem as the interest of this work is the limit $\epsilon\to 0$). The CRAM approximation to the matrix exponential is rewritten as:

$$\boldsymbol{C}(t)=\left\{\alpha_{0,K}\boldsymbol{I}+2\mathrm{Re}\left[\sum_{j=1}^{K/2}\left(-\frac{\alpha_{j,K}}{\theta_{j,K}}\right)\left(\boldsymbol{I}+\boldsymbol{A}\left(-\frac{t}{\theta_{j,K}}\right)\right)^{-1}\right]\right\}. \tag{7}$$

Using Eq. 6 we expand the term in parenthesis in Eq. 7 by identifying:

$$\begin{aligned}\boldsymbol{P}&=\boldsymbol{I}\\ \boldsymbol{Q}&=\boldsymbol{A}\\ \epsilon&=-\frac{t}{\theta_{j,K}}.\end{aligned} \tag{8}$$

This leads to an expression for the CRAM approximation of the matrix exponential:

$$\boldsymbol{C}(t)=\sum_{l=0}^{\infty}F_{l,K}\left(\boldsymbol{A}t\right)^{l}, \tag{9}$$

where the coefficient $F_{l,K}$ is given by:

$$F_{l,K}=\delta_{0,l}\alpha_{0,K}-2\mathrm{Re}\left[\sum_{j=1}^{K/2}\frac{\alpha_{j,K}}{\theta_{j,K}^{l+1}}\right]. \tag{10}$$

3. Local Truncation Error Analysis

This section investigates the local truncation error of CRAM. The local truncation error is defined as the error of $\vec{n}'$ after a single time step of length Δt in the limit $\Delta t\to 0$:

$$e=\lim_{\Delta t\to 0}\|\vec{n}'(\Delta t)-\vec{n}(\Delta t)\|=\lim_{\Delta t\to 0}\|\left[\boldsymbol{C}(\Delta t)-\exp\left(\boldsymbol{A}\Delta t\right)\right]\vec{n}_0\|, \tag{11}$$

where $\|\cdot\|$ denotes a vector norm; in this work the two-norm is used exclusively. Using Eqs. 3 and 9, we find:

$$e = \lim_{\Delta t \to 0} \left\| \sum_{l=0}^{\infty} \left[(\boldsymbol{A}\Delta t)^l \left(\frac{1}{l!} - F_{l,K} \right) \right] \vec{n}_0 \right\|. \tag{12}$$

Passing to the limit $\Delta t \to 0$ yields:

$$e = |F_{0,K} - 1| \|\vec{n}_0\|. \tag{13}$$

An approximation is called consistent if the truncation error vanishes as $\Delta t \to 0$. The consistency requirement for CRAM is thus given by $F_{0,K} = 1$. The values of $F_{0,K}$ and $|F_{0,K} - 1|$ are listed in Table 1. CRAM is not consistent because $F_{0,K} \neq 1$ for any finite CRAM order. However, as the order increases, $|F_{0,K} - 1|$ decreases; for orders $K = 16$ it is hardly larger than numerical noise from double precision arithmetic. Therefore, CRAM is effectively consistent if orders of 16 or larger are used. If CRAM orders much lower than 16 are used, the error between exact and CRAM solution in the limit $\Delta t \to 0$ can be as high as 6.7% (i.e. $\sup|F_{0,K} - 1|$ attained for $K = 1$ in Table 1).

The values of the coefficients $F_{0,K}$ are previously listed by Refs. [11] and [6]. The coefficients are restated in Table 1. In this work, the values of $F_{0,K}$ play an important role in the global truncation error analysis to follow.

4. Global Truncation Error Analysis

This section investigates the global truncation error of CRAM. The global truncation error is the difference between the exact and CRAM solution at the end of time interval $[0, t_e]$ that is divided into N time steps of length $\Delta t = t_e/N$ in the limit $\Delta t \to 0$. It follows from the propagation and accumulation of local truncation error over the time interval $[0, t_e]$.

The global truncation error measures the accuracy of the solution at a fixed final time (e.g. a refueling outage) as the number of time steps is increased, or the time step size is reduced. The result of this section strictly applies to the limit of infinitesimally short time steps. However, these asymptotic results

K	$F_{0,K}$	$\lvert F_{0,K} - 1\rvert$
1	1.06683104216185	6.6831e-02
2	0.99264132983042	7.3587e-03
3	1.00079938063634	7.9938e-04
4	0.99991347759305	8.6522e-05
5	1.00000934571316	9.3457e-06
6	0.99999899154562	1.0085e-06
8	0.99999998827735	1.1723e-08
12	0.99999999999840	1.6023e-12
15	1.00000000000044	4.4409e-13
16	1.00000000000000	3.5527e-15

Table 1: The values of $F_{0,K}$ and $|F_{0,K} - 1|$ for various CRAM orders K. $|F_{0,K} - 1|$ indicates how far away from a consistent scheme CRAM of order K is.

are still useful as long as N is sufficiently large and the global truncation error is in the asymptotic regime where the leading order term dominates. Numerical results in Sec. 5 are used to investigate CRAM's behavior outside of the asymptotic regime.

The expectation for consistent numerical schemes is that the global truncation error reduces with some power of Δt and vanishes if $\Delta t = 0$. As CRAM is not consistent, the global truncation error does not vanish. However, the exact behavior in the limit $N \to \infty$ does not immediately follow from the results of Sec. 3.

The isotope number density at t_e can be expressed as:

$$\begin{aligned} \vec{n} &= \exp\left(\boldsymbol{A} t_e\right) \vec{n}_0 = \exp\left(\boldsymbol{A} N \Delta t\right) \vec{n}_0 = \left(\exp\left(\boldsymbol{A} \Delta t\right)\right)^N \vec{n}_0 \\ \vec{n}' &= \left(\boldsymbol{C}(\Delta t)\right)^N \vec{n}_0. \end{aligned} \tag{14}$$

The global truncation error E at t_e is hence given by:

$$E = \lim_{\substack{N \to \infty \\ \Delta t \to 0}} \left\| \left[\left(\boldsymbol{C}(\Delta t)\right)^N - \left(\exp\left(\boldsymbol{A} \Delta t\right)\right)^N \right] \vec{n}_0 \right\|. \tag{15}$$

Substituting Eqs. 3 and 9 for $\boldsymbol{C}(\Delta t)$ and $\exp(\boldsymbol{A}\Delta t)$, respectively, leads to:

$$E = \lim_{\substack{N\to\infty \\ \Delta t\to 0}} \left\| \left[\left(\sum_{l=0}^{\infty} F_{l,K} \left(\boldsymbol{A}\Delta t\right)^l \right)^N - \left(\sum_{l=0}^{\infty} \frac{1}{l!} (\boldsymbol{A}\Delta t)^l \right)^N \right] \vec{n}_0 \right\|. \tag{16}$$

In the limit $\Delta t \to 0, N \to \infty$, terms with $l > 0$ vanish and we obtain the asymptotic error estimate:

$$E = \lim_{\substack{N\to\infty \\ \Delta t\to 0}} E = \lim_{N\to\infty} \left| F_{0,K}^N - 1 \right| \|\vec{n}_0\|. \tag{17}$$

It has already been established in Section 3 that $F_{0,K} \neq 1$, and hence E is nonzero. The value of the relative error $E/\|\vec{n}_0\|$ depends on the magnitude of $F_{0,K}$:

$$\frac{E}{\|\vec{n}_0\|} = \begin{cases} 1 & \text{if } F_{0,K} < 1; \\ \infty & \text{if } F_{0,K} > 1. \end{cases} \tag{18}$$

From Table 1, we observe that even CRAM orders have $F_{0,K} < 1$, while odd CRAM orders have $F_{0,K} > 1$; correspondingly, the relative global truncation error for even CRAM orders is one, while it is unbounded for odd orders.

The result shown in Eq. 18 holds regardless of CRAM order. However, for higher CRAM orders N must be larger for $F_{0,K}^N$ to be significantly different from one. Consequently, we expect higher order CRAM's accuracy to degrade slower than lower order CRAM's. This is corroborated through numerical experiment in Sec. 5.

The conclusions of this section appear to be in contradiction with the findings of Ref. [12] that suggest more "substeps" (referred to as time steps in this manuscript) lead to smaller errors in the computed number densities at a given end time t_e. Reference [12]'s findings apply most prominently to the depletion of fresh fuel, where errors "are dominated by a few fast transitions". However, using "substeps" also improves accuracy for the depletion of "old" fuel. First, it is pointed out that Ref. [12]'s conclusions are based on the average of the isotope-wise relative errors, E_r, computed by:

$$E_r = \frac{1}{M} \sum_{m=1}^{M} \left| 1 - \frac{n'_m}{n_m} \right|, \tag{19}$$

where n'_m and n_m are the m-th components of $\vec{n}'$ and $\vec{n}$, respectively. In contrast, the presented analysis uses a vector norm to measure the error.

We find, through the numerical experiment described in detail in Sect. 5, that average relative errors go through a minimum at intermediate values of N:

- For large time steps, errors are large because number densities change too much within a single time step. In depletion calculations, isotopes with short half-lives (on the order of nanoseconds) rapidly decay and effectively reach a concentration of zero within a time span that is shorter than the time step. However, CRAM predicts a much larger residual concentration for such isotopes if too few time steps are used.

- For small time steps, the consequences of CRAM's inconsistency degrade accuracy; in the limit $N \to \infty$, E_r behaves identical to $E/\|\vec{n}_0\|$.

Therefore, the conclusions in this work are not in contradiction to Ref. [12]'s conclusions; Ref. [12] stopped increasing N before a significant increase in E_r could be observed.

5. Numerical Results

In this section, numerical results obtained with Rattlesnake's CRAM implementation are used to verify the predictions of Sect. 4. The accuracy of CRAM's results is measured against a highly accurate reference solution computed by the doubling method described in Ref. [22]. A brief review of the doubling method is given in Sect. 5.1. The comparison of CRAM and doubling results for a reactor depletion problem using the 297 isotope DRAGON library [23] are presented in Sect. 5.2.

5.1. The Doubling Method

The basic idea of the doubling method is to advance the solution from time step i at time t to time step $i+1$ at time $t+\Delta t$ by solving a sequence of problems indexed by $j = 0, .., J$ for the solution at time $t + 2^j \Delta t/2^J$. The advantage of

this approach is a reduction in the number of operations required to estimate the solution at the end of the time step and the consequent reduction of propagation error. Doubling is a solution strategy that wraps around a differencing algorithm such as backward Euler or Crank-Nicholson.

For simplicity, the Crank-Nicholson (CN) scheme is chosen as the underlying differencing scheme. Applying the CN scheme to the matrix exponential starting at time t to obtain the solution at time $t + \Delta t/2^J$ gives:

$$\begin{aligned} \vec{n}_{i+1,1} &\approx \boldsymbol{P}\vec{n}_i \\ \boldsymbol{P} &= \left[\boldsymbol{I} + \frac{\Delta t}{2^{J+1}}\boldsymbol{A}\right]\left[\boldsymbol{I} - \frac{\Delta t}{2^{J+1}}\boldsymbol{A}\right]^{-1}. \end{aligned} \tag{20}$$

To obtain the solution at time $t + \Delta t$, CN is applied 2^J times:

$$\vec{n}_{i+1} \approx \boldsymbol{P}^{2^J}\vec{n}_i. \tag{21}$$

To drastically reduce the number of required multiplications and enable as large a J as possible, we compute the sequence by doubling:

$$\begin{aligned} \boldsymbol{P}^2 &= \boldsymbol{P}\boldsymbol{P} \\ \boldsymbol{P}^4 &= \left(\boldsymbol{P}^2\right)\left(\boldsymbol{P}^2\right) \\ \boldsymbol{P}^8 &= \left(\boldsymbol{P}^4\right)\left(\boldsymbol{P}^4\right) \\ &\vdots \\ \boldsymbol{P}^{2^J} &= \left(\boldsymbol{P}^{2^{J-1}}\right)\left(\boldsymbol{P}^{2^{J-1}}\right). \end{aligned} \tag{22}$$

Rather than the 2^J matrix multiplications required by the straightforward CN scheme, doubling requires only $J = \log_2\left(2^J\right)$ multiplications. For large values of J, the doubling method allows the computation of highly accurate reference solutions for the Bateman equations.

5.2. Numerical Assessment of CRAM's Accuracy

CRAM's accuracy is investigated by numerical experiment. The selected test problem solves the Bateman equations for a single region using data from

the 297 isotope library of DRAGON5 [23]; this library contains the most relevant isotopes for depletion of nuclear fuel including fissionable isotopes, fission products, and decay products of fissionable isotopes and fission products. This data library is augmented by an additional 17 stable pseudo-isotopes that do not transmute or decay and close the decay chains in the 297 isotope library provided by DRAGON5 (i.e. these 17 isotopes are the stable products of the considered decay chains). One-group transmutation cross-sections are computed using a beginning-of-life light-water reactor spectrum, and the composition at $t = 0$ resembles 4.7% enriched fresh UO_2 fuel. The neutron spectrum is fixed, and consequently, the depletion matrix solved by CRAM does not change between time steps.

The initial number densities are listed in Table 2. One important consideration is the omission of the isotope helium-4 and all isotopes lighter than helium-4, including all isotopes of hydrogen, from the depletion calculation (even though hydrogen is present in the infinite medium mixture used for computing the depletion cross sections). Rattlesnake calculates depletion chains with the buildup of secondary particles, such as alpha particles from alpha decays and hydrogen-1 from (n,p) transmutation reactions, while the reference doubling solution does not calculate secondary particle production originating from these sources. Because the depletion chains in the DRAGON5 data set are isolated for isotopes lighter than helium-4 (i.e. any isotope lighter than helium-4 cannot be transmuted into an isotope heavier than helium-4), this omission does not impact the validity of the comparisons between Rattlesnake and the reference doubling solution for all isotopes heavier than helium-4.

The accuracy of CRAM orders $1, 2, 3, 4, 5, 8, 15$, and 16 is assessed by solving the test problem from $t = 0$ to $t_e = 5.184 \times 10^7$ seconds using a varying number of uniform time steps N ranging from $N = 1$ to $N = 10^5$. The error is computed in the vector norm E and the average of the isotope-wise relative error E_r and plotted versus N in Figures. 1 and 2, respectively.

Increasing the number of time steps increases the vector-norm error. The experimentally observed asymptotic behavior follows the prediction of Eq. 18

Isotope	Number density [atoms / barn cm]
O_{16}	4.66×10^{-2}
O_{17}	1.77×10^{-5}
H_1^a	4.95×10^{-2}
U_{234}	8.45×10^{-6}
U_{235}	1.05×10^{-3}
U_{238}	2.23×10^{-2}

[a] Hydrogen is included in the infinite medium mixture, but it is excluded from the depletion calculation.

Table 2: Initial number densities for the 314 isotope LWR test problem.

and Table 1: numerically computed $E/\|\vec{n}_0\|$ grow without bound for odd orders and approach one for even orders. As the CRAM order is increased, the observed errors decrease dramatically and for order 16 the error is smaller than 10^{-10} for $\Delta t \approx 30$ days. In practice, time steps during reactor depletion calculations are of order $\Delta t \approx 10$ days, although for many reactor systems, initial depletion steps on the order of $\Delta t \approx 20$ minutes may be necessary in order to account for the buildup of reactor poisons such as xenon-135 and samarium-149. The limits of practical time steps are depicted as navy-colored vertical lines in Figures 1 and 2. For CRAM order 16, the relative error is smaller than 10^{-10} for any time step between 20 minutes and 10 days. These results indicate that CRAM order 16 is sufficiently accurate for LWR depletion calculations, but analysis of a wider range of problems should be performed to confirm the generality of this conclusion. Conversely, the presented results suggest that using low order CRAM with short time steps leads to large errors; CRAM orders 4 and 8 may have relative errors larger than 0.1 and 10^{-4}, respectively, within this time step range.

The average of the isotope-wise relative error, E_r, initially decreases with increasing N, assumes a minimum at $N_{\min}$, and then starts increasing; the

asymptotic behavior is identical to the vector norm error. This result is consistent with Isotalo's work [12] that reports a reduction of E_r with increasing the number of time steps. However, Ref. [12] does not increase N sufficiently to observe the upturn in error. It may be desirable for a CRAM user to select N to be close to $N_{\min}$, so the trends of how $N_{\min}$ changes with CRAM order are discussed here.

The value of $N_{\min}$ depends on the CRAM order. For even orders $2, 4, 8$, and 16, $N_{\min}$ is $10, 60, 100$, and 40, respectively, while for odd orders $1, 3, 5$, and 15, $N_{\min}$ is $30, 50, 100$, and 50, respectively. The initial trend is that $N_{\min}$ increases with increasing CRAM order, but then it drops for orders 15 and 16. However, for these two orders the error vs. N curves are essentially flat around $N_{\min}$ making it harder to determine where the minimum actually occurs. In addition, the location of the minimum for orders 15 and 16 is much less important because the error assumed in a wide neighborhood around $N_{\min}$ does not differ much from the minimum value. The observations regarding the minimum of E_r is solely based on numerical results for the presented problem. It is possible these values may be dependent upon the depletion system analyzed, and further research considering different problem statements and different data sets may be warranted.

In Fig. 3 the absolute value of the relative error of four important isotopes (U_{235}, Pu_{239}, Sm_{149}, and Xe_{135}) are plotted versus number of time steps/time step size. We observe that (1) single-isotope relative errors of the four selected isotopes behave very similar to each other, and (2) behave very similar to the vector norm error depicted in Fig. 1. These findings indicate that relative errors of a subset of isotopes mirror the behavior of the vector-norm error. Consequently, conclusions of CRAM's accuracy for the vector-norm error carry over to these isotopes.

It is curious that the single-isotope relative errors do not follow the trend of the average of the isotope-wise relative error E_r at large time steps. We conjecture that for large time steps E_r is dominated by errors in isotopes with very low concentrations (i.e. the uptick in E_r for large time steps is caused by

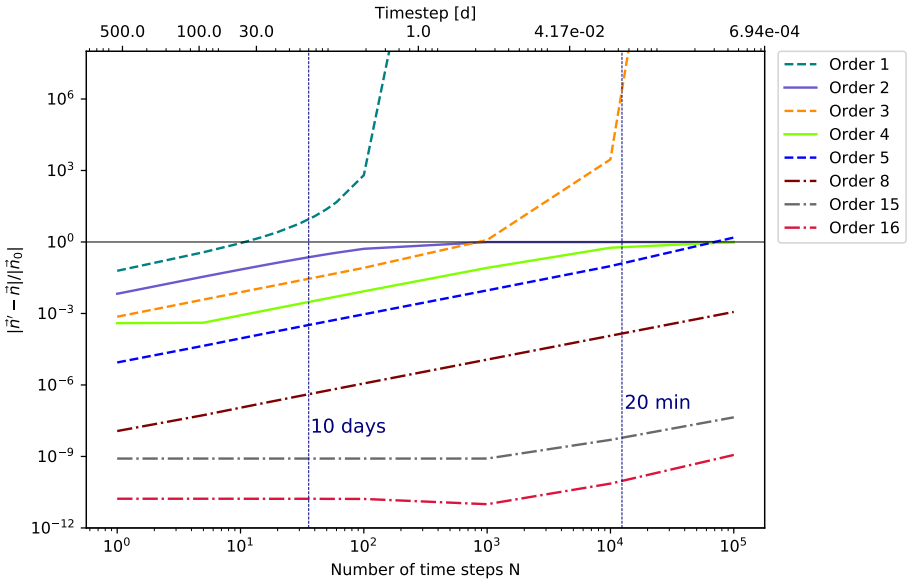

Figure 1: Vector norm error $E/\|\vec{n}_0\|$ plotted versus number of time steps N (bottom axis) and time step Δt (top axis) for the light-water reactor pin-cell test problem.

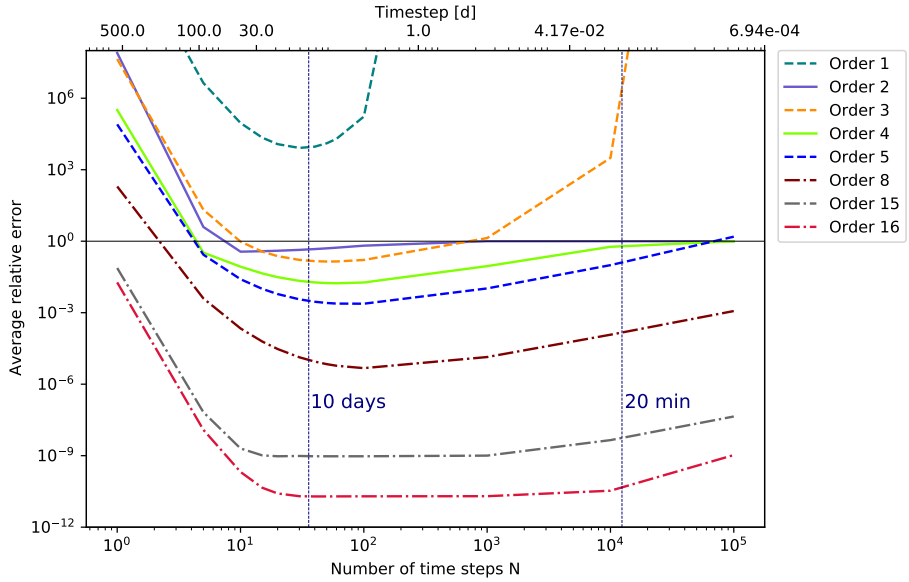

Figure 2: Average of the isotope-wise relative error E_r plotted versus number of time steps N (bottom axis) and time step Δt (top axis) for the light-water reactor pin-cell test problem.

isotopes with very small concentrations). The isotopes we examine in Fig. 3 all have sufficiently large concentrations and their relative errors follow the trend of the vector-norm error where very dilute isotopes have a negligible contribution to the error.

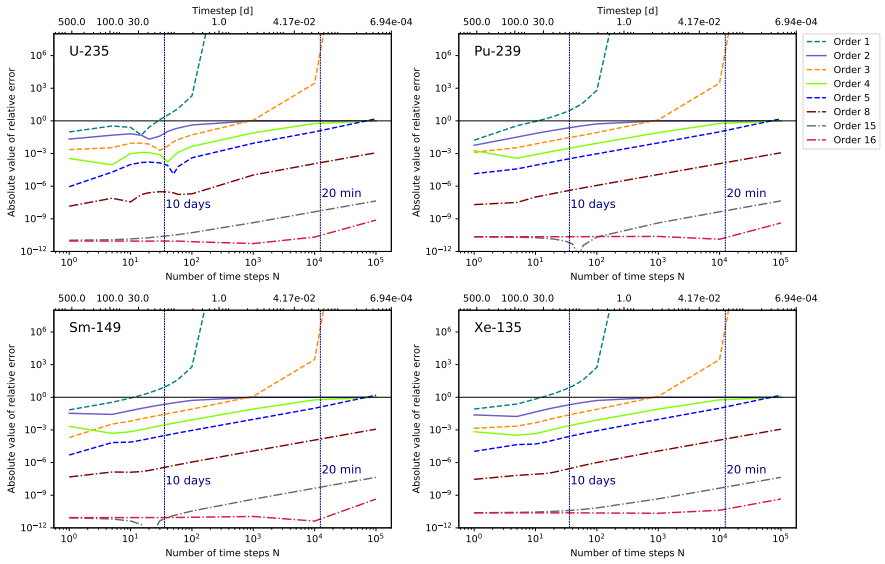

Figure 3: Absolute value of th relative error for isotopes U_{235}, Pu_{239}, Sm_{149}, and Xe_{135} plotted versus number of time steps N (bottom axis) and time step Δt (top axis) for the light-water reactor pin-cell test problem.

6. Conclusions

This article shows that CRAM's relative global truncation error increases to one and infinity for even and odd CRAM orders, respectively, as the time step is decreased. This result is important for users of CRAM because errors may increase when the time step is reduced with the intention to enhance accuracy or by the needs of some coupled physics. We distill our findings into two best practices for using CRAM:

1. Use CRAM order 16 or higher.
2. If multiphysics feedback requires reduction of the time step, increase the CRAM order as needed.

This work demonstrates that CRAM is not consistent; the relative local truncation error does not vanish as $\Delta t \to 0$. However, as the order is increased, CRAM becomes more and more consistent. For CRAM orders 16 and higher, the remaining local truncation error in the limit $\Delta t \to 0$ is smaller or equal to the precision of double precision arithmetic. The inconsistency is the root cause for the unexpected characteristics of CRAM.

We perform a global truncation error analysis and show that the relative global truncation error increases to one and infinity for even and odd CRAM orders, respectively, as the length of the time step is reduced. The global truncation error analysis is augmented by numerical experiments comparing Rattlesnake CRAM results with a highly accurate reference solution obtained through the doubling method. The numerical experiments validate Isotalo's work [12] in its analysis of decreasing relative error in isotope number densities for increasing number of time steps. The improvement in accuracy is driven by limiting the maximum amount of change of each isotope concentration. However, the reduction in error does not continue indefinitely because of the inconsistency in CRAM that leads to a deterioration of accuracy past a critical time step size $\Delta t_{\min}$. The rate at which CRAM's accuracy deteriorates with decreasing time step size decreases with CRAM order. We advise users to increase the CRAM order when the depletion time step size is reduced.

For calculations performed in double precision, CRAM order 16 can be expected to provide results with the anticipated accuracy as documented by Isotalo [12] for any reasonable number of time steps in practical applications. This leads us to advise using CRAM orders of 16 or higher. Since higher-order CRAM calculations require an increasing number of arithmetic operations, runtimes for these higher order solutions also increase, thus a user may be tempted to choose a lower order CRAM solution and compensate for the reduced accuracy by increasing the number of time steps; this is not a recommended practice as a result of the observed deterioration in accuracy.

Acknowledgments

This manuscript has been authored by Battelle Energy Alliance, LLC under Contract No. DE-AC07-05ID14517 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

Appendix

Coefficient	Real Part	Imaginary Part
α_0	$-6.683104216185046 \times 10^{-2}$	$0.000000000000000 \times 10^{0}$
α_1	$+6.563897638922146 \times 10^{-1}$	$0.000000000000000 \times 10^{0}$
θ_1	$-5.789994857980907 \times 10^{-1}$	$0.000000000000000 \times 10^{0}$

Table 3: CRAM coefficients used for Order 1.

Coefficient	Real Part	Imaginary Part
α_0	$+7.358670169580528 \times 10^{-3}$	$0.000000000000000 \times 10^{0}$
α_1	$-1.688587273471294 \times 10^{-1}$	$-8.094501977207793 \times 10^{-1}$
θ_1	$-5.847986326384906 \times 10^{-1}$	$+1.185537781241578 \times 10^{0}$

Table 4: CRAM coefficients used for Order 2.

Coefficient	Real Part	Imaginary Part
α_0	$-7.993806363356878 \times 10^{-4}$	$0.000000000000000 \times 10^{0}$
α_1	$-6.911639229743115 \times 10^{-1}$	$+4.318221807343100 \times 10^{-5}$
α_2	$+1.483848552302971 \times 10^{0}$	$0.000000000000000 \times 10^{0}$
θ_1	$-1.982084787507404 \times 10^{-1}$	$+2.410732102959337 \times 10^{0}$
θ_2	$-1.368849323033854 \times 10^{0}$	$0.000000000000000 \times 10^{0}$

Table 5: CRAM coefficients used for Order 3.

Coefficient	Real Part	Imaginary Part
α_0	$+8.652240695288853 \times 10^{-5}$	$0.000000000000000 \times 10^{0}$
α_1	$-7.339595716394208 \times 10^{-2}$	$+4.499999224740631 \times 10^{-1}$
α_2	$+6.168677956783283 \times 10^{-2}$	$-1.905040979303084 \times 10^{0}$
θ_1	$+3.678453861815378 \times 10^{-1}$	$+3.658121298678667 \times 10^{0}$
θ_2	$-1.548393223297123 \times 10^{0}$	$+1.191822946627425 \times 10^{0}$

Table 6: CRAM coefficients used for Order 4.

Coefficient	Real Part	Imaginary Part
α_0	$-9.345713153026646 \times 10^{-6}$	$0.000000000000000 \times 10^{0}$
α_1	$+2.387886840880759 \times 10^{-1}$	$+1.037386723036788 \times 10^{-1}$
α_2	$-1.872385779133251 \times 10^{0}$	$-3.78206306824073 \times 10^{-1}$
α_3	$+3.271819394452691 \times 10^{0}$	$0.000000000000000 \times 10^{0}$
θ_1	$+1.039505662793796 \times 10^{0}$	$+4.921390733695027 \times 10^{0}$
θ_2	$-1.440595994251471 \times 10^{0}$	$+2.396983525293841 \times 10^{0}$
θ_3	$-2.155415462008556 \times 10^{0}$	$0.000000000000000 \times 10^{0}$

Table 7: CRAM coefficients used for Order 5.

Coefficient	Real Part	Imaginary Part
α_0	$+1.172265211633491 \times 10^{-8}$	$0.000000000000000 \times 10^{0}$
α_1	$-2.812975702329915 \times 10^{-2}$	$+1.157738420892369 \times 10^{-2}$
α_2	$+6.325880513115763 \times 10^{-1}$	$-4.439230969478992 \times 10^{-1}$
α_3	$-2.436240725781788 \times 10^{-2}$	$+3.716755613508318 \times 10^{-1}$
α_4	$+1.831771705771544 \times 10^{0}$	$-9.525608074064298 \times 10^{0}$
θ_1	$+3.408539509663649 \times 10^{0}$	$+8.773034553579034 \times 10^{0}$
θ_2	$-2.694909809127538 \times 10^{-1}$	$+6.082032587493543 \times 10^{0}$
θ_3	$-2.292249142304860 \times 10^{0}$	$+3.600771493483381 \times 10^{0}$
θ_4	$-3.220945239945118 \times 10^{0}$	$+1.193619604620657 \times 10^{0}$

Table 8: CRAM coefficients used for Order 8.

Coefficient	Real Part	Imaginary Part
α_0	$-1.973138996612803 \times 10^{-15}$	$0.000000000000000 \times 10^{0}$
α_1	$+6.122664124472788 \times 10^{-5}$	$-1.367020570275008 \times 10^{-5}$
α_2	$-9.152096770578211 \times 10^{-3}$	$+2.105137145495789 \times 10^{-3}$
α_3	$+2.485673139684196 \times 10^{-1}$	$-1.519186546775811 \times 10^{-1}$
α_4	$-1.874218038231563 \times 10^{0}$	$+2.920080910204411 \times 10^{0}$
α_5	$+1.201982912704081 \times 10^{0}$	$-2.031135193815113 \times 10^{1}$
α_6	$+3.078529179819246 \times 10^{1}$	$+5.938976702938319 \times 10^{1}$
α_7	$-1.135376336279002 \times 10^{2}$	$-6.912824885069587 \times 10^{1}$
α_8	$+1.66370201022818 \times 10^{1}$	$0.000000000000000 \times 10^{0}$
θ_1	$+9.86565548146789 \times 10^{0}$	$+1.795293385233925 \times 10^{1}$
θ_2	$+4.473793268825284 \times 10^{0}$	$+1.493639197784267 \times 10^{1}$
θ_3	$+7.956423718721415 \times 10^{-1}$	$+1.224230235152575 \times 10^{1}$
θ_4	$-1.866307664276102 \times 10^{0}$	$+9.691291108271987 \times 10^{0}$
θ_5	$-3.78002540409208 \times 10^{0}$	$+7.218380899778627 \times 10^{0}$
θ_6	$-5.078618543722945 \times 10^{0}$	$+4.791065618206045 \times 10^{0}$
θ_7	$-5.833760820653117 \times 10^{0}$	$+2.389625482128766 \times 10^{0}$
θ_8	$-6.081751716480002 \times 10^{0}$	$0.000000000000000 \times 10^{0}$

Table 9: CRAM coefficients used for Order 15.

Coefficient	Real Part	Imaginary Part
α_0	$+2.1248537104952237480 \times 10^{-8}$	$0.0000000000000000000 \times 10^{0}$
α_1	$+5.0901521865224915650 \times 10^{-7}$	$-2.4220017652852287970 \times 10^{-5}$
α_2	$-2.1151742182466030907 \times 10^{-4}$	$+4.3892969647380673918 \times 10^{-3}$
α_3	$-4.1023136835410021273 \times 10^{-2}$	$-1.5743466173455468191 \times 10^{-1}$
α_4	$+1.4793007113557999718 \times 10^{0}$	$+1.7686588323782937906 \times 10^{0}$
α_5	$-1.5059585270023467528 \times 10^{1}$	$-5.7514052776421819979 \times 10^{0}$
α_6	$+6.2518392463207918892 \times 10^{1}$	$-1.1190391094283228480 \times 10^{1}$
α_7	$-1.1339775178483930527 \times 10^{2}$	$+1.0194721704215856450 \times 10^{2}$
α_8	$+6.4500878025539646595 \times 10^{1}$	$-2.2459440762652096056 \times 10^{2}$
θ_1	$+1.0843917078696988026 \times 10^{1}$	$+1.9277446167181652284 \times 10^{1}$
θ_2	$+5.2649713434426468895 \times 10^{0}$	$+1.6220221473167927305 \times 10^{1}$
θ_3	$+1.4139284624888862114 \times 10^{0}$	$+1.3497725698892745389 \times 10^{1}$
θ_4	$-1.4193758971856659786 \times 10^{0}$	$+1.0925363484496722585 \times 10^{1}$
θ_5	$-3.5091036084149180974 \times 10^{0}$	$+8.4361989858843750826 \times 10^{0}$
θ_6	$-4.9931747377179963991 \times 10^{0}$	$+5.9968817136039422260 \times 10^{0}$
θ_7	$-5.9481522689511774808 \times 10^{0}$	$+3.5874573620183222829 \times 10^{0}$
θ_8	$-6.4161776990994341923 \times 10^{0}$	$+1.1941223933701386874 \times 10^{0}$

Table 10: CRAM coefficients used for Order 16.

References

[1] H. Bateman, "The solution of a system of differential equations occurring in the theory of radioactive transformations," Proc. Cambridge Philos. Soc., vol. 15, pp. 423–427, 1910.

[2] M. Pusa and J. Leppanen, "Computing the matrix exponential in burnup calculations," Nuclear Science and Engineering, vol. 164, no. 1, pp. 140–150, 2009.

[3] M. Pusa, "Rational approximations to the matrix exponential in burnup calculations," Nuclear Science and Engineering, vol. 169, no. 2, pp. 155–167, 2011.

[4] M. Pusa, "Accuracy considerations for Chebyshev rational approximation method (cram) in burnup calculations," in International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, M&C 2013, pp. 973–984, ANS, 2013.

[5] M. Pusa, "Higher-order Chebyshev rational approximation method and application to burnup equations," Nuclear Science and Engineering, vol. 182, no. 3, pp. 297–318, 2016.

[6] T. Schmelzer and L. N. Trefethen, "Evaluating matrix functions for exponential integrators via Carathéodory-Fejér approximation and contour integrals," Electron. Trans. Numer. Anal., vol. 29, pp. 1–18, 2007-2008.

[7] J. Hykes and R. Ferrer, "Solving the Bateman equations in CASMO5 using implicit ode numerical methods for stiff systems," in International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, M&C 2013, ANS, 2013.

[8] J. Cetnar, "General solution of Bateman equations for nuclear transmutations," Annals of Nuclear Energy, vol. 33, no. 7, pp. 640–645, 2006.

[9] C. Moler and C. V. Loan, "Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later," SIAM Review, vol. 45, 2003.

[10] A. Isotalo and P. Aarnio, "Comparison of depletion algorithms for large systems of nuclides," Annals of Nuclear Energy, vol. 38, no. 2, pp. 261–268, 2011.

[11] A. Isotalo and W. Wieselquist, "A method for including external feed in depletion calculations with cram and implementation into origen," Annals of Nuclear Energy, vol. 85, pp. 68–77, 2015.

[12] A. Isotalo and M. Pusa, "Improving the accuracy of the Chebyshev rational approximation method using substeps," Nuclear Science and Engineering, vol. 183, no. 1, pp. 65–77, 2016.

[13] Y. Wang, S. Schunert, and V. Laboure, Rattlesnake Theory Manual. Idaho National Laboratory, 2017.

[14] J. Leppänen, Development of a New Monte Carlo Reactor Physics Code. PhD thesis, Helsinki University of Technology, 2007.

[15] J. Leppänen, "Serpent - a continuous-energy monte Carlo reactor physics burnup calculation code." `http://montecarlo.vtt.fi/download/Serpent_manual.pdf`, 2015. Accessed on 01/01/2015.

[16] M. I. R. Zweiger, "Implementation of the Chebyshev rational approximation method for burnup calculations into PyNE." `https://pdfs.semanticscholar.org/d54e/08ac78d9f3a2b08b9c274713f39518d00802.pdf`, 2019. Accessed on 03/23/2020.

[17] V. T. Luan and A. Ostermann, "Exponential Rosenbrock methods of order five — construction, analysis and numerical comparisons," Journal of Computational and Applied Mathematics, vol. 255, pp. 417 – 431, 2014.

[18] H. J. Park, D. H. Lee, B. K. Jeon, and H. J. Shim, "Monte Carlo burnup and its uncertainty propagation analyses for vera depletion benchmarks by mccard," Nuclear Engineering and Technology, vol. 50, no. 7, pp. 1043 – 1050, 2018.

[19] A. Isotalo and P. Aarnio, "Higher order methods for burnup calculations with Bateman solutions," Annals of Nuclear Energy, vol. 38, no. 9, pp. 1987 – 1995, 2011.

[20] A. Isotalo and P. Aarnio, "Substep methods for burnup calculations with Bateman solutions," Annals of Nuclear Energy, vol. 38, no. 11, pp. 2509 – 2514, 2011.

[21] G. H. Golub and C. F. V. Loan, Matrix Computations. The Johns Hopkins University Press, 1996.

[22] B. D. Ganapol, "Particle transport in a 3D duct by adding and doubling," Journal of Computational and Theoretical Transport, vol. 46, no. 3, pp. 202–228, 2017.

[23] A. Hebert, Applied Reactor Physics. Presses Internationales Polytechnique, 2016.