A Stable Nuclide Transmutation Procedure Free of Numerical Roundoff

Erwin Müller[1] , Frederik Reitsma[*2] and Paulus P Kruger[2]
[1]*CompuSim AB, SE-723 34 Västerås, Sweden*
[2]*Radiation and Reactor Theory Group, Necsa, P.O. Box 582, Pretoria 0001, South Africa*

Abstract

One of the classical problems faced by developers of nuclear reactor physics codes is that of finding suitable numerical procedures for solving the system of first-order ordinary differential equations of nuclide transmutation. An efficient, accurate and reliable method that places no restriction on the nuclide transmutation paths is generally sought. A crucial requirement is that the method must be stable and free of catastrophic roundoff error. Virtually all methods currently in use are susceptible to serious numerical roundoff error and require some form of approximation to mitigate such problems. The penalty for these approximations is a potential loss of accuracy and more often than not a loss of generality.

Recently, a method that is free from catastrophic roundoff and that fulfills all the requirements for stability, accuracy, reliability and generality was proposed for application in both lattice physics and core simulator codes. The method is based on approximating the matrix exponential by a truncated Taylor series expansion and the application of the uniformization technique to eliminate catastrophic roundoff. By adopting a compact data representation for the nuclide transmutation matrix, numerical operations with zero-valued elements of the matrix can be avoided so as to achieve a high level of efficiency. This nuclide transmutation solution method, as implemented in the MGRAC nodal core simulator, is described and numerical results are presented to illustrate the salient features of the method.

KEYWORDS: nuclide transmutation, depletion method, matrix exponential, lattice physics codes, core simulator codes, MGRAC

1. Introduction

During nuclear reactor operation a variety of materials undergo changes due to radioactive decay and/or nuclear reactions whereby one nuclide may transmute into another nuclide or nuclides. The transmutation processes that are of cardinal importance are fission, neutron capture and radioactive decay. Nuclides that are coupled via their individual transmutation processes (i.e., coupled chains of nuclides) are said to form a *transmutation group* and the number densities of all nuclides that belong to a given transmutation group can be computed at discrete time points by solving a system of first-order ordinary differential equations for that group:

[*] Corresponding author: Tel. +27 12 305 5314, Fax. +27 12 305 5166, E-mail: reitsma@necsa.co.za

$$\frac{d}{dt}\vec{N}(t) = \hat{A}(t)\,\vec{N}(t) \tag{1}$$

Assuming (as we do in this work) the transmutation matrix $\hat{A}(t)$ to be a piecewise constant function of time, the solution for a discrete time step Δt can be written as:

$$\vec{N}(t_0 + \Delta t) = e^{\Delta t \hat{A}(t_0)}\,\vec{N}(t_0) \tag{2}$$

A variety of methods exist for determining the matrix exponential $e^{\Delta t \hat{A}(t_0)}$, none of which are completely satisfactory [1]. From a historical perspective, it is useful to mention three of the most widely used methods for solving the nuclide transmutation equations and to note their deficiencies and hence the motivation for yet another method. These methods are:

1. The analytic method as implemented in the GAUGE core simulator code [2] for fuel depletion.
2. The matrix exponential method as implemented in the ORIGEN-S code [3] for general nuclide transmutation.
3. The Laplace transform method as implemented in the PHOENIX lattice code [4] for fuel depletion.

The analytic method assumes a triangular form for the transmutation matrix in order to facilitate an analytic solution of the transmutation equations. The triangular form excludes feedback reactions from the depletion chains (i.e., cyclic chains) and thereby restricts the transmutation mechanisms that can be described. While this is normally considered to be the main disadvantage of the method, the fact that it is prone to very large roundoff errors must also be considered a serious deficiency. The problem of large roundoff error (catastrophic cancellation) occurs when two (or more) eigenvalues of the transmutation matrix are the same (confluent) or numerically nearly the same [1,5,6].

The matrix exponential method of ORIGEN-S approximates the matrix exponential $e^{\Delta t \hat{A}(t_0)}$ by a standard Taylor series expansion and by that suffers from severe roundoff error (catastrophic cancellation) when the equation system is stiff. In order to contain the roundoff error to some acceptable level the stiffness of the system must be reduced and in the ORIGEN-S code this is achieved by limiting the time step size and by removing short half-life isotopes from the transmutation matrix (their inventories are then separately calculated by means of a Bateman (analytic) method). The effect of removing the short half-life isotopes amounts to an alteration of the transmutation model and is such that significant departure from exact results is possible at short depletion times (i.e., during the early transmutation life of many nuclides) [7]. Thus, the method as implemented in ORIGEN-S not only may be inefficient (very small time step restriction) but it inherently suffers from loss of accuracy.

The Laplace transform method as implemented in the PHOENIX code invokes linearization of the burnup chains and it is thus not readily applicable to general transmutation matrices, a limitation it has in common with the analytic method. The actinides are collected into a transmutation group of their own and therefore the fission-products can be broken up into a number of transmutation groups (the time-step average actinide number densities are employed in the fission yield term of the fission products). The consequence is that the time variation of the actinide and fission product number densities is not strictly synchronized but this is expected to be of little practical consequence if unreasonably large time steps are not used. As pointed out by Moler and Van Loan [1], Laplace transform methods may be seriously affected by roundoff error and in the PHOENIX implementation cognizance is taken of the fact that the contribution of a

distant predecessor nuclide (an indirect parent imposed by the linearization) can become very inaccurate due to catastrophic cancellation. The precaution taken is to neglect any distant contribution when it becomes very small (approaches available computer precision). This makes the method susceptible to loss of accuracy.

The severe round-off error that plagues the Taylor series approach can be removed by means of the uniformization technique [8]. In that way the disadvantages of the Taylor series method as implemented in the ORIGEN-S code can be eliminated and the method becomes a very attractive candidate for application to the nuclide transmutation problem. In a recent study that focused on the nuclide transmutation problem in nuclear reactors, Krüger [9] concluded that the Taylor expansion method combined with the uniformization technique is the most suitable method in such applications. As a result of that work, the method has been implemented in the nodal core simulator MGRAC [10] and it has been evaluated in a number of applications ranging from rather simple transmutation systems (with a small number of burnup chains) to quite large and very stiff transmutation systems.

In this paper numerical results are presented to demonstrate the accuracy and efficiency of the Taylor Series Method with Uniformization in whole-core cycle depletion applications. First, a short description of the method as implemented in MGRAC is given.

2. The Taylor Series Method with Uniformization

The starting point for the method is the regular truncated Taylor series expansion of the matrix exponential, $e^{\Delta t\hat{A}(t_0)}$. This yields

$$\vec{N}(t_0+\Delta t) \approx \sum_{k=0}^{p} \frac{(\Delta t)^k}{k!} (\hat{A})^k \vec{N}(t_0) \tag{3}$$

The diagonal elements of the transmutation matrix $\hat{A}$ are all negative (since they are loss terms) while all the off-diagonal elements are positive (since they are gain terms). Thus the transmutation matrix is *essentially* non-negative.

The main disadvantage of the regular Taylor expansion is that it is subject to severe cancellation error due to terms of alternating sign (since the diagonal of the matrix $\hat{A}$ is negative). It is noted that the larger the stiffness of the matrix $\hat{A}$, the larger the roundoff error becomes since the cancellation error is proportional to the maximum term in the Taylor series expansion with terms of alternating sign. Therefore, by reducing the stiffness as is done in ORIGEN-S the roundoff error is also reduced.

The uniformization technique [8] uses the modified formulation

$$e^{\Delta t\hat{A}}\ \vec{N}(t_0) = e^{-\alpha\Delta t} e^{\Delta t(\hat{A}+\alpha\hat{I})} \vec{N}(t_0) \tag{4}$$

where

$$\alpha = \max_i \left|A_{ii}\right| \tag{5}$$

to obtain:

$$\vec{N}(t_0+\Delta t) \approx e^{-\alpha\Delta t} \sum_{k=0}^{p} \frac{(\Delta t)^k}{k!} (\hat{A}+\alpha\hat{I})^k \vec{N}(t_0) \tag{6}$$

This method is free of the cancellation error caused by terms of alternating sign (the matrix

$\hat{A} + \alpha\hat{I}$ is non-negative) and it becomes a stable method. Because only non-negative terms are added in the series the remaining roundoff error (due to repeated matrix-vector multiplications and vector additions) is independent of the stiffness of the transmutation matrix and is of the same order as the computer precision.

The only shortcoming of uniformization is [8] that p (the number of expansion terms) is likely to be large (in the sense of achieving a certain truncation accuracy) for large $\alpha\Delta t$; for the regular Taylor series approximation p will be smaller. In the MGRAC core simulator the series, which is implemented as a recursive series, is continued as long as each new term contributes more (in the relative sense) to the series than some user-specified tolerance. The truncation tolerance must of course be meaningful with reference to the actual floating point precision that is being used, but it must be emphasized that it is not a measure of the roundoff error that will be tolerated (since that issue is resolved by uniformization), but rather a matter of setting the *accuracy that is desired* from the truncated Taylor series. In MGRAC, floating point computations are performed in double precision and a tolerance of the order of 10^{-12} is, for example, not considered excessively tight (in terms of available precision). Numerical experimentation with the tolerance has been performed and some results are presented later in the paper.

The transmutation matrix is generally a sparse matrix with many zero elements. Since the transmutation matrix is used in repeated matrix-vector multiplications in the Taylor series expansion, the numerical efficiency of the method can be improved (especially if the matrix is large) by eliminating operations with the zero elements of the transmutation matrix and by not storing such elements. The scheme selected in MGRAC for this purpose is a slightly modified version of the Compressed Row Storage (CRS) format [11].

Note that in the rest of this paper any reference to the Taylor series method is taken to imply application of the uniformization technique as well.

3. Description of the Numerical Analysis

The emphasis in this paper is on the numerical experimentation with the Taylor series method (with uniformization) in whole-core cycle depletion applications. For that purpose two test problems were selected, one with a somewhat simplified nuclide transmutation matrix typical of power reactor applications and one with a large transmutation matrix more typical of research reactor applications:

1. The first test problem is a well-known international depletion benchmark problem known as the KWU 2-D PWR burnup benchmark problem [12]. The two-dimensional reactor core contains 193 homogenized fuel assemblies (each with 23 cm side length) and the core has octant symmetry. The fuel transmutation matrix for this problem is lower triangular and represents the linear burnup chains for 7 actinides and 2 fission products (a lumped fission product and Xe^{135}). Xe^{135} is treated as being in equilibrium during the cycle and Xe^{135} is therefore not required to be explicitly included in the fuel transmutation matrix. The nuclear data (burnup chains, initial nuclide inventories, 2-group exposure-independent cross sections) are specified in [12]. This problem is henceforth referred to as the KWU problem. The spatial mesh (nodes) used for this problem consists of 25 nodes per assembly, which is the mesh size used for the reported reference solution [12]. The results presented for this problem consist of the critical

soluble boron concentration (boron let-down) at selected times during a depletion cycle (only the first cycle of this two-cycle problem is modeled).

2. The second test problem is a non-symmetric three-dimensional Materials Test Reactor (MTR) problem with 72 homogenized in-core assemblies (each with 7.71 x 8.1 cm^2 base area and 60 cm height). The core represents an actual loading pattern of the SAFARI-1 reactor and is loaded with 27 fuel assemblies, 6 control assemblies (follower type) and several irradiation rig assemblies. A single depletion cycle is modeled with all the control rods withdrawn (i.e., the fuel follower sections are in the active core) throughout the cycle. The fuel transmutation matrix for this problem is lower triangular and represents the burnup chains for 19 actinides and 19 fission products. Xe^{135} transmutation is explicitly modeled (assuming end of previous cycle xenon for burned fuel and zero xenon for fresh fuel at beginning of current cycle). The nuclear data (burnup chains, initial and exposure-dependent nuclide inventories, 6-group exposure-dependent cross sections) were generated by means of standard lattice depletion calculations for each fuel type in the core. This problem is henceforth referred to as the MTR problem. Near cuboid nodes (of assembly side length) were used for this problem. The results presented for this problem consist of the reactor multiplication factor k_{eff} at selected times during a depletion cycle.

The choice of test problems with lower triangular transmutation matrices was made in order to facilitate comparison of the performance of the Taylor series method against that of the original analytic transmutation method (the GAUGE method) available in a previous version of MGRAC.

The specific calculation methods and options used in MGRAC are:

1. Solution of the multi-group neutron diffusion equations by means of the Multi-Group Analytic Nodal Method [13] using a quadratic transverse leakage approximation and a quadratic intra-node burnup correction. Eigenvalue and nodal power convergence criterion set at 10^{-6}.
2. Depletion calculation by means of a constant-flux explicit time-integration method (one global flux solution per burnup step) or a semi-implicit predictor-corrector method (two flux solutions per burnup step) when using the Taylor series method.
3. Depletion calculation by means of a constant-flux explicit time-integration method (one global flux solution per burnup step) or a three-step flux-renormalization method (one global flux solution per burnup step and three nodal flux renormalizations to preserve constant node powers during a burnup step) when using the analytic transmutation method (older version of MGRAC).
4. Variation of the Taylor series truncation tolerance ranging from 10^{-12} to 10^{-3} to study the sensitivity of the Taylor series method to this parameter.
5. Variation of burnup step sizes to study the sensitivity of the depletion method to step size.

4. Results

4.1 Sensitivity to Truncation Tolerance

The first analysis performed concerns the sensitivity of the Taylor series method to the Taylor series truncation tolerance. The results obtained for the KWU problem using the predictor-corrector method are reported in Table 1. The second column lists the internationally accepted

reference results [12] obtained by KWU using the analytic transmutation method with small burnup steps (4 x 1.5 days, 2 x 12 days, 23 x 15 days, 6.5 days). The same burnup steps were used in MGRAC for this analysis. The last five columns list the MGRAC results. It is seen that the MGRAC results agree quite well (within 1.5 ppm for those cases using tight truncation tolerances) with the reference results. In the table the depletion CPU time (D_{CPU} in seconds) and the percentage of the total CPU time consumed by the depletion procedure (D %) are also given; the depletion procedure incorporates transmutation matrix construction and the transmutation algorithm.

Table 1: Impact of Taylor series truncation tolerance for the KWU problem.

Days since BOC	Boron content (ppm)					
	Reference	$\varepsilon = 10^{-12}$	$\varepsilon = 10^{-6}$	$\varepsilon = 10^{-5}$	$\varepsilon = 10^{-4}$	$\varepsilon = 10^{-3}$
0	1038.6	1038.5	1038.5	1038.5	1038.5	1038.5
30	1081.7	1082.3	1082.3	1082.3	1082.0	1078.0
60	1069.2	1070.3	1070.3	1070.2	1069.6	1060.6
90	1019.0	1020.4	1020.4	1020.3	1019.3	1004.7
120	944.1	945.6	945.6	945.5	944.0	924.0
150	853.1	854.5	854.5	854.3	852.5	827.3
180	751.5	752.9	752.9	752.8	750.6	720.4
210	643.5	644.8	644.8	644.6	642.1	607.5
240	531.7	532.9	532.9	532.6	529.9	491.5
270	418.1	419.2	419.2	419.0	416.0	374.4
300	304.3	305.2	305.2	304.9	301.8	257.8
330	191.0	191.8	191.8	191.5	188.3	142.6
360	79.1	79.8	79.8	79.5	76.2	29.5
381.5	0.0	0.6	0.6	0.3	-2.9	-49.8
D_{CPU}		23.5	20.2	19.7	18.3	18.3
D %		4.7	4.2	4.2	3.9	3.4

It is evident that a truncation tolerance larger than 10^{-4} is inadequate since a significant loss in accuracy (of the order of 50 ppm) is incurred for this problem and that a tighter tolerance of the order of 10^{-6} should be quite sufficient. However, since the truncation tolerance seems to have a moderate impact on the CPU time (D_{CPU}) it is probably prudent to use as tight a tolerance as allowed by the machine precision (say 10 x *machine epsilon*). It is noted that the depletion procedure consumes less than 5% of the overall computation resources for this problem. This is a very acceptable cost and the method can only be considered as highly efficient in this type of application. Auxiliary calculations where the burden of the nodal flux solution was relaxed by changing from a 25 nodes/assembly to 1 node/assembly spatial mesh showed a noticeable increase in the fractional cost of the depletion procedure, but it remained well below 10%. This is still an acceptable cost.

The same conclusion is drawn for the MTR problem, the results of which are given in Table 2 for a case run with 3 day burnup steps (typical step size for routine MTR core analysis). Differences in k_{eff} relative to the case with $\varepsilon = 10^{-12}$ are given in pcm (1 pcm = 10^{-5} Δk_{eff}). It is noted that the depletion procedure consumes only about 1% of the total CPU time. Since the nodal diffusion solution for this problem is far more taxing than that of KWU problem (6 versus

2 neutron energy groups) and the transmutation matrix is much larger (38 x 38 versus 9x9), the results of this case may be taken as an indication that the method should also be quite efficient in lattice physics applications (where the transport solution completely dominates the computational effort even when very large transmutation matrices are used).

Table 2: Impact of Taylor series truncation tolerance for the MTR problem.

Days since BOC	k_{eff}	k_{eff} Differences (pcm)			
	$\varepsilon = 10^{-12}$	$\varepsilon = 10^{-6}$	$\varepsilon = 10^{-5}$	$\varepsilon = 10^{-4}$	$\varepsilon = 10^{-3}$
0	1.04486	0	0	0	0
3	1.03704	0	0	-4	-47
6	1.03229	0	-1	-7	-96
9	1.02754	0	-1	-11	-147
12	1.02277	0	-1	-16	-201
15	1.01795	0	-1	-19	-256
18	1.01308	0	-2	-24	-314
21	1.00816	-1	-3	-29	-375
24	1.00316	0	-2	-33	-437
D_{CPU}	12.4	11.6	11.9	11.5	11.5
D %	0.9	0.8	0.8	0.8	0.8

4.2 Sensitivity to Burnup Step Size

To test the sensitivity of the depletion results to the burnup step size the KWU problem was run with burnup step sizes ranging from half of the reference burnup step sizes to burnup steps of two months. These results are reported in Table 3 for the Taylor series method ($\varepsilon = 10^{-12}$) using both the predictor-corrector and the explicit method. The predictor-corrector calculation with the smallest burnup steps ($\frac{1}{2}\Delta t_{ref}^{predcor}$) is used as reference for the comparison of boron deviations in this table. The following observations can be made.

Halving or doubling the reference time step sizes does not significantly affect the results obtained with the predictor-corrector method (< 1 ppm differences in boron concentration throughout the cycle). The explicit method is slightly more sensitive (up to 3 ppm differences relative to the explicit case with reference time steps) with regard to such time step variations. This is an indication that the reference time steps selected for this benchmark problem are indeed adequate for establishing a reference solution.

The predictor-corrector case with one month burnup steps agrees very closely (within 0.2 ppm) with the case where the reference time step size was doubled since in the latter case the majority of the step sizes are equal to one month steps. In the case of the explicit method there is a slightly larger deviation (of the order of 1 ppm) due to the larger sensitivity to time step size of this method.

The predictor-corrector method with a two month burnup step size yields results that are surprisingly close (within 3 ppm) to the selected reference case ($\frac{1}{2}\Delta t_{ref}^{predcor}$). The explicit method results, on the other hand, deviate significantly in this case (up to 13 ppm). The larger sensitivity to burnup step size of the explicit method is expected since the flux space-energy distribution is assumed to remain invariant during any given burnup step whereas an update of this distribution

is obtained with the predictor-corrector method.

If one were to impose an acceptance criterion of 2 ppm with regard to deviation from the selected reference ($\frac{1}{2}\Delta t_{ref}^{predcor}$) results, then it may be concluded that the predictor-corrector method with a one month burnup step size would be quite adequate for this benchmark problem; it is not uncommon to use one month burnup steps in practical PWR cycle depletion calculations. The explicit method with a one month burnup step does not quite meet this strict 2 ppm requirement. Only when the smallest time steps $\frac{1}{2}\Delta t_{ref}$ are used does the explicit method reach this goal. To achieve roughly the same accuracy the explicit method requires two to four times as many time steps as the predictor-corrector method and in such circumstances it cannot compete with the predictor-corrector method as far as CPU time is concerned (total CPU time T_{CPU} is given in the table). It is noted that for the KWU problem the depletion procedure requires less than 6% of the total CPU time in all cases listed in Table 3.

Table 3: Impact of burnup step size for the KWU problem.

Days since BOC	Boron differences (ppm)									
	Predictor-corrector					Explicit				
	$\frac{1}{2}\Delta t_{ref}$	Δt_{ref}	$2\Delta t_{ref}$	Δt_{1month}	Δt_{2month}	$\frac{1}{2}\Delta t_{ref}$	Δt_{ref}	$2\Delta t_{ref}$	Δt_{1month}	Δt_{2month}
0	1038.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
30	1082.3	0.0	-0.2	-0.4		-0.4	-0.9	-1.9	-2.9	
60	1070.3	0.0	-0.4	-0.5	-2.2	-0.8	-1.8	-3.7	-4.5	-9.9
90	1020.5	-0.1	-0.5	-0.6		-1.2	-2.4	-4.9	-5.6	
120	945.7	-0.1	-0.6	-0.7	-2.7	-1.3	-2.6	-5.5	-6.1	-13.3
150	854.6	-0.1	-0.5	-0.6		-1.3	-2.5	-5.4	-6.1	
180	753.1	-0.2	-0.6	-0.7	-2.8	-1.2	-2.4	-5.1	-5.8	-12.7
210	644.9	-0.1	-0.6	-0.7		-1.0	-2.0	-4.4	-5.1	
240	533.0	-0.1	-0.6	-0.7	-2.7	-0.8	-1.6	-3.6	-4.3	-9.7
270	419.3	-0.1	-0.5	-0.6		-0.6	-1.2	-2.6	-3.3	
300	305.3	-0.1	-0.5	-0.6	-2.5	-0.3	-0.7	-1.7	-2.4	-5.9
330	192.0	-0.2	-0.5	-0.6		-0.2	-0.3	-0.7	-1.4	
360	79.9	-0.1	-0.4	-0.5	-2.2	0.2	0.3	0.3	-0.3	-1.8
381.5	0.7	-0.1	-0.4	-0.5	-2.2	0.3	0.5	0.5	0.2	-1.3
T_{CPU}	757	497	294	226	175	435	307	150	128	97
D %	5.6	4.7	4.9	4.5	4.2	4.2	2.8	3.2	3.3	3.1

The results for the MTR case are shown in Table 4 where the predictor-corrector method using a ½ day burnup step size serves as the reference. Results for the Taylor series method ($\varepsilon = 10^{-12}$) using both the predictor-corrector and the explicit method in combination with various step sizes are included in the table. Differences in k_{eff} relative to the reference k_{eff} are reported in this table. The following observations can be made.

The predictor-corrector method closely reproduces the reference results even when the step size is increased to 6 days, yielding k_{eff} within 4 pcm of the reference value throughout the cycle. With the explicit method the deviations in k_{eff} increase consistently with the burnup step size. For a 1 day step size the effect is 8 pcm at the end of the cycle whereas it is 12 pcm for 1½ day, 26 pcm for 3 day and 54 pcm for 6 day time steps. For this MTR problem the predictor-corrector

method is clearly superior to the explicit method with regard to achieving a given accuracy with the least computation effort. It is noted that for the MTR problem the depletion procedure requires about 1% of the total CPU time in all cases listed in Table 4.

Table 4: Impact of burnup step size for the MTR problem.

Days since BOC	Predictor-corrector					Explicit			
	$\Delta t_{1/2\ day}$	$\Delta t_{1\ day}$	$\Delta t_{1\ 1/2\ days}$	$\Delta t_{3\ days}$	$\Delta t_{6\ days}$	$\Delta t_{1\ day}$	$\Delta t_{1\ 1/2\ days}$	$\Delta t_{3\ days}$	$\Delta t_{6\ days}$
0	1.04486	0	0	0	0	0	0	0	0
3	1.03707	-2	-2	-3		-1	-1	0	
6	1.03232	-2	-3	-3	-4	0	1	4	10
9	1.02757	-2	-3	-3		2	3	8	
12	1.02280	-2	-3	-3	-3	3	4	11	25
15	1.01798	-1	-2	-3		5	7	15	
18	1.01311	-2	-2	-3	-3	6	9	19	40
21	1.00818	-1	-2	-2		7	10	23	
24	1.00319	-2	-2	-3	-3	8	12	26	54
T_{CPU}	6248	3236	2508	1406	831	1854	1395	838	521
D %	0.9	1.0	0.9	0.9	1.1	0.8	0.8	0.8	0.8

4.3 Comparison with the analytic transmutation method

To place the performance of the Taylor series method in perspective a comparison of results obtained with the analytic transmutation method was done, both for the KWU and the MTR problem. The results for the KWU problem as obtained with two month burnup steps are reported in Table 5. The MTR results listed in Table 6 were obtained with 3 day burnup steps.

One of the first observations to make is that the Taylor explicit method and the analytic explicit method are essentially equivalent. This is a gratifying result since it confirms that the Taylor series method is capable of reproducing the analytic solution. The analytic method combined with the three-step flux-renormalization approach (constant power method) is clearly more effective in the case of the MTR problem than it is in the case of the KWU problem. A likely reason for this is that the burnup step size used for the KWU problem is so large that the assumption of invariant flux spectrum during a burnup step is violated. The three-step flux-renormalization and the predictor-corrector methods simply represent two approaches of improving the neutron flux distribution that is used during a depletion step and in that regard the predictor-corrector method is superior in that it actually recalculates this distribution whereas in the former method an ad-hoc renormalization is done. With smaller time-steps the three-step flux-renormalization approach becomes attractive since it clearly improves upon the normal explicit method and it reduces the number of global flux solutions relative to the predictor-corrector method.

Unfortunately, a direct comparison of the efficiency of the Taylor series and the analytic methods is somewhat clouded by the fact that the Taylor series method and the analytic method are implemented in different versions of MGRAC. Since the most significant differences between the two versions of MGRAC are related to the global flux solution procedure the fraction of CPU time consumed by the deletion procedure does not reflect the true performance

of the procedure. Therefore, a direct comparison of the CPU time (in seconds) consumed by the depletion procedure (i.e., D_{CPU}) is more realistic for the purpose of comparing the Taylor series method and the analytic method. Accordingly, it appears that the Taylor series method and the analytic method require essentially the same computation effort per transmutation calculation. From this comparison it is concluded that the Taylor series method is at least as accurate as the analytic method and that it suffers no penalty in computation time relative to the analytic method.

Table 5: Comparison with the analytic transmutation method for the KWU problem.

Days since BOC	Reference $\frac{1}{2}\Delta t_{ref}$	Boron differences (ppm)			
		Taylor: Pred-cor	Taylor: Explicit	Analytic: Explicit	Analytic: Const-power
0	1038.5	0.0	0.0	0.1	0.1
60	1070.3	-2.2	-9.9	-9.9	-8.7
120	945.7	-2.7	-13.3	-13.3	-10.1
180	753.1	-2.8	-12.7	-12.7	-9.3
240	533.0	-2.7	-9.7	-9.7	-7.3
300	305.3	-2.5	-5.9	-5.9	-5.2
360	79.9	-2.2	-1.8	-1.8	-3.2
381.5	0.7	-2.2	-1.3	-1.4	-3.0
D_{CPU}	42.1	7.4	3.0	2.4	7.1
D %	5.6	4.2	3.1	3.8	10.3

Table 6: Comparison with the analytic transmutation method for the MTR problem.

Days since BOC	Reference $\Delta t_{1/2\ day}$	k_{eff} Differences (pcm)			
		Taylor: Pred-cor	Taylor: Explicit	Analytic: Explicit	Analytic: Const-power
0	1.04486	0	0	-1	-1
3	1.03707	-3	0	-1	-4
6	1.03232	-3	4	4	-3
9	1.02757	-3	8	8	-3
12	1.02280	-3	11	11	-3
15	1.01798	-3	15	15	-3
18	1.01311	-3	19	18	-3
21	1.00818	-2	23	22	-3
24	1.00319	-3	26	26	-3
D_{CPU}	55.9	14.6	6.6	7.4	21.7
D %	0.9	0.9	0.8	7.5	19.4

5. Conclusion

A method based on approximating the matrix exponential by a truncated Taylor series expansion combined with the uniformization technique was implemented in the MGRAC core simulator to solve the nuclide transmutation equations. The numerical results presented for two

benchmark problems illustrate the utility of the method and comparison with a traditional analytic transmutation method confirms the accuracy of the method.

The main attributes of the method are:

1. generality - it places no restriction on nuclide transmutation paths
2. reliability - it is free of catastrophic roundoff error and it is stable
3. efficiency - it consumes only a small to modest fraction of the overall computation time of a reactor physics calculation

It is concluded that the Taylor series method combined with the uniformization technique is a viable method for nuclide transmutation calculation and that it can be expected to perform well both in core simulators and lattice physics codes.

References

1) Cleve Moler and Charles van Loan, "Nineteen Dubious Ways to Compute the Exponential of a Matrix," *SIAM Review*, **29**, 801 (1978).
2) M. R. Wagner, "GAUGE: A Two-Dimensional Few Group Neutron Diffusion-Depletion Program for a Uniform Triangular Mesh," AEC Research and Development Report GA-8307, Gulf General Atomic, 1968.
3) 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 Term," NUREG/CR-0200, Rev. 6, Vol. 2, Oak Ridge National Laboratory, 1998.
4) R. J. J. Stamm'ler and M. J. Abbate, *Methods of Steady-State Reactor Physics in Nuclear Design*, Academic Press, 1983.
5) Z. Vukadin, "Estimate of the Error of depletion function expansion series," *KERNTECHNIK*, **60**, 197 (1995).
6) H. Tachihara and H. Sekimoto, "Exact Error Estimation for Solutions of Nuclide Chain Equations," *Journal of Nuclear Science and Technology*, **36**, 1176 (1999).
7) G. F. Thomas and D. H. Barber, "Stiffness in Radioactive Decay Chains," *Annals of Nuclear Energy*, **21**, 309 (1994).
8) R. B. Sidje and W. J. Stewart, "A numerical study of large sparse matrix exponentials arising in Markov chains," *Computational Statistics & Data Analysis*, **29**, 345 (1999).
9) P. P. Krüger, "Numerical Methods to Solve the Fuel Depletion Equations for a Nuclear Reactor," MSc thesis, Potchefstroom University for CHE, Potchefstroom, South Africa, 2004.
10) E. Z. Müller, et. al., "Development of a Core Follow Calculational System for Research Reactors," 9th Pacific Basin Nuclear Conference, Sydney, Australia, 1-6 May 1994.
11) Richard Barrett, et. al., *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, SIAM, 1994. (http://www.netlib.org/templates/Templates.html)
12) "Benchmark Source Situation BSS-19," National Energy Software Center: Benchmark Problem Book, Report ANL-7416, Suppl. 3, pp. 880 & 892, 1985.
13) D. L. Vogel and Z. J. Weiss, "A General Multigroup Formulation of the Analytic Nodal Method," Proc. Int. Topical Meeting on Advances in Reactor Physics, Charleston, South Carolina, USA, 8-11 March 1992.