Annals of Nuclear Energy 120 (2018) 611–624

Contents lists available at ScienceDirect

Annals of Nuclear Energy

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

Essentially nonnegative matrix exponential methods for nuclide transmutation

Erwin Müller

CompuSim AB, Strängnäsgatan 6, SE-72334 Västerås, Sweden

ARTICLE INFO

Article history:
Received 4 February 2018
Received in revised form 4 June 2018
Accepted 7 June 2018
Available online 22 June 2018

Keywords:
Matrix exponential
Essentially nonnegative
Transmutation

ABSTRACT

Two methods originally developed for discrete-time Markov chains are adopted for the solution of the first-order ordinary differential equation of nuclide transmutation. Both methods use Taylor series expansions, which facilitates software implementation. The methods are known, respectively, as the uniformization method and the aggressively truncated Taylor series method. The theory and algorithmic aspects of the two methods, as far as is relevant for software implementation, are presented. A few numerical test problems are employed to compare the two methods and to obtain an impression of their capabilities.

1. Introduction

In the domain of nuclear reactor physics, the temporal tracking of the nuclide inventory of a variety of burnable (depletable) materials in a reactor core plays a key role (Bell and Glasstone, 1985; Henry, 1986). This tracking involves the solution of a first-order ordinary differential equation for various groupings of nuclides which we refer to as transmutation groups. A transmutation group is formed by a collection of nuclides that are linked via their individual nuclear transmutation processes (i.e., coupled chains of nuclides). For any given transmutation group for which a transmutation matrix $\hat{A}(t)$ is defined, the dynamic equation (Bell and Glasstone, 1985)

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{N}(t) = \hat{A}(t)\vec{N}(t); \quad \vec{N}(0) = \vec{N}_0 \tag{1}$$

must be solved to yield the column vector $\vec{N}(t)$ of nuclide concentrations (atom densities) at any desired time $t > 0$. Making the usual assumption (Bell and Glasstone, 1985; Henry, 1986) that the transmutation matrix $\hat{A}(t)$ is a piecewise constant function of time, the formal solution of the dynamic equation can be constructed at discrete time points as (Bell and Glasstone, 1985)

$$\vec{N}(t_{m+1}) = e^{\Delta t_{m+1}\hat{A}_{m+1}}\vec{N}(t_m); \quad m = 0, 1, 2, \ldots \tag{2}$$

with $\Delta t_{m+1} \equiv t_{m+1} - t_m$ and $t_0 = 0$. This scheme implies a sequential computation of the inventory vector: once $\vec{N}(t_m)$ has been determined, $\vec{N}(t_{m+1})$ can be determined. The subscript $m+1$ indicates that the transmutation matrix $\hat{A}_{m+1}$ belongs to time step $m+1$ and that this matrix is time-invariant during the time interval $t \in [t_m, t_{m+1}]$. This time interval is known as the time integration domain for the matrix $\hat{A}_{m+1}$, but we shall also refer to it as a burnup step. Hereafter, we shall consider the solution to the linear dynamic equation for a single burnup step only, thus permitting us to drop the burnup step index. We then write the formal solution as

$$\vec{N}(t_0 + \Delta t) = e^{\Delta t\hat{A}}\vec{N}(t_0) \tag{3}$$

where t_0 now signifies the time at the start of the given burnup step of size Δt. A numerical representation of the matrix exponential $e^{\Delta t\hat{A}}$ is required to complete the calculation of the inventory vector at the end of the burnup step. It is the determination of this matrix exponential, or of its product with a vector, that is the central theme in nuclide transmutation calculations.

A variety of methods exist for determining the matrix exponential, none of which are completely satisfactory in the general sense (Moler and Van Loan, 2003). Consideration of the matrix properties can be a useful guide in selecting an appropriate method. In this regard, the fact that the nuclide transmutation matrix is an essentially nonnegative matrix[1] is of great importance since essentially nonnegative matrices constitute a natural class of matrices in the analysis of matrix exponentials: a matrix $\hat{A}$ is essentially nonnegative if and only if $e^{t\hat{A}}$ is nonnegative for all $t \geqslant 0$ (Varga, 1962). Some

[1] A real square matrix is said to be essentially nonnegative if its off-diagonal entries are nonnegative.

E-mail address: erwin.muller@compusim.se

https://doi.org/10.1016/j.anucene.2018.06.012
0306-4549/

of the available schemes for the computation of matrix exponentials may not be appropriate for this class of matrix while others may be especially tailored for it. For instance, for an essentially nonnegative matrix with a negative diagonal, the naive application of a truncated Taylor series approximation,

$$e^{t\hat{A}} \approx \sum_{k=0}^{K} \frac{t^k}{k!} \hat{A}^k \tag{4}$$

may result in a numerically unstable, inaccurate, and unreliable method as a result of catastrophic cancellation in the summation of terms of alternating sign. On the other hand, numerical methods that have been specially developed for discrete-time Markov chains (Abdallah and Marie, 1993; Sidje and Stewart, 1999; Dulat et al., 2010; Sidje, 2011) may be ideally suited to the nuclide transmutation problem.[2] Two such methods are analysed in this paper. The first one is the well-known (standard) uniformization method (Abdallah and Marie, 1993; Sidje and Stewart, 1999). The other is a more recent development (Shao, 2012; Shao et al., 2014) known as the "aggressively truncated Taylor series method", which is basically an optimised version of the uniformized power method (Abdallah and Marie, 1993; Sidje and Stewart, 1999). The common thread between these methods is the use of the Taylor series method supplemented by the uniformization approach (Sidje and Stewart, 1999) to eliminate catastrophic cancellation. The trademark of these two methods is simplicity, which greatly facilitates implementation in a computer code. The application of the uniformization method to the nuclide transmutation problem has actually been reported previously (Krüger, 2004; Müller et al., 2006) and we include its discussion in this paper primarily because it forms the basis for the development of the aggressively truncated Taylor series method. Another reason is that we want to compare the aggressively truncated Taylor series method, whose application to the nuclide transmutation problem is original, against the uniformization method on a fair basis with both implemented in a common computer code.[3]

The objective in this paper is distinctly modest in that neither a comprehensive review of existing nuclide transmutation methods, nor a numerical comparison against other well known methods is attempted. Instead, we focus entirely on a mutual comparison of our two chosen methods, thereby hoping to stimulate more extensive future research into the practical utility of these methods in the broader nuclear reactor analysis arena. We do, however, include a brief overview of some of the better known (existing) methods in order to gain some perspective.

The numerical analysis presented in this work is restricted to relatively small nuclide transmutation problems because the aggressively truncated Taylor series method is expected to be most suitable to such problems. Since we wish to identify possible applications where the uniformization method is challenged and where the aggressively truncated Taylor series method could be a good alternative, we included in our set of test problems a few very demanding numerical benchmarks developed by Lago and Rahnema (2017).

The remainder of this paper is organised as follows. First, a short overview of the nuclide transmutation methods currently used in the reactor physics field is given. Next, important properties of the nuclide transmutation matrix are noted. The uniformization method and the aggressively truncated Taylor series method are then developed in turn, with the greater part of the text dedicated to a discourse on the algorithmic aspects that are of relevance to the software implementation. This is followed by the numerical analysis. Finally, we draw conclusions and convey our opinion on the prospects of these methods.

2. Overview of nuclide transmutation methods

Nuclide transmutation calculations can be divided into three categories:

1. Reference type calculations involving very large transmutation matrices for thousands of nuclides with extreme variation in effective half-lives (i.e., very stiff problems). These calculations generally simulate reactor fuel depletion histories (power histories) in a highly simplified way since their main purpose is to generate detailed radionuclide inventories for post-irradiation calculations, both of a short-term (using very small burnup steps) and long-term character (using extremely large burnup steps).
2. Lattice physics depletion calculations (Stamm'ler and Abbate, 1983) involving large transmutation matrices for several hundred (fewer than a thousand) nuclides with large variation in effective half-lives (i.e., stiff problems). These calculations employ simplified (reduced) nuclide chains and rather short burnup time steps (ranging from 1 day to a few months). The nuclides included in these calculations are primarily those that have a significant impact on neutron economy. The purpose of these calculations is to prepare data for global reactor core depletion calculations, and sometimes even for reference type calculations.
3. Global reactor core depletion calculations (Henry, 1986) involving small transmutation matrices for fewer than a hundred nuclides with modest to large variation in effective half-lives (i.e., still stiff problems). The nuclide chains in these calculations are even further reduced from those used in the lattice calculations but the burnup time steps are of similar size. These depletion calculations can be, and often are used to prepare fuel depletion histories for reference type calculations.

Very sophisticated nuclide transmutation algorithms are needed in reference type calculations, whereas much simpler methods, such as those developed in this paper, are used in lattice physics and core depletion calculations. One of the best known methods used in reference type calculations is the so-called matrix exponential method embodied in the ORIGEN code (Gauld et al., 2011). This method combines the standard Taylor series approximation of Eq. (4) with a generalised form of the Bateman equations (Bateman, 1910) for linearised nuclide chains. The catastrophic cancellation that plagues the standard Taylor series approach is mitigated by removing the most short-lived nuclide chains from the transmutation matrix and using the Bateman solution for them. The Taylor series method is applied to the reduced matrix only. Because a Horner scheme involving recursive matrix-vector multiplications is used, the matrix exponential (of the reduced matrix) is never actually constructed. Not only is this hybrid scheme complicated, but it also alters the transmutation model such that significant departure from exact results is possible at short depletion times during the early transmutation life of many short-lived nuclides (Thomas et al., 1994).

Recently, another method has gained popularity in this class of application, namely the Chebyshev Rational Approximation Method (CRAM) (Pusa, 2011). This method, which exploits the properties of essentially nonnegative matrices, is based on a rational function approximation of the exponential on the negative real axis. This method is practically insensitive to burnup time step size and the stiffness of the problem (Pusa, 2011; Krüger, 2004). The CRAM has been shown to be a robust and accurate method,

[2] Continuous-time nuclide transmutation chains are actually continuous-time Markov chains (Halász (2018)).

[3] The two methods have been implemented in a nuclide depletion computer code written in the Fortran 2003 programming language.