Department of Applied Physics

Computational Methods for Burnup Calculations with Monte Carlo Neutronics

Aarno Isotalo

Aalto University

DOCTORAL DISSERTATIONS

Aalto University publication series
DOCTORAL DISSERTATIONS 172/2013

Computational Methods for Burnup Calculations with Monte Carlo Neutronics

Aarno Isotalo

A doctoral dissertation completed for the degree of Doctor of Science to be defended, with the permission of the Aalto University School of Science, at a public examination held at the lecture hall K216 of the school on the 4th of December 2013 at 12 noon.

Aalto University
School of Science
Department of Applied Physics

Supervising professor
Rainer Salomaa, Aalto University, Finland

Thesis advisor
Pertti Aarnio, Aalto University, Finland

Preliminary examiners
Gašper Žerovnik, Jožef Stefan Institute, Slovenia
Jussi Kumpula, TVO, Finland

Opponent
Christophe Demazière, Chalmers University of Technology, Sweden

Aalto University publication series
DOCTORAL DISSERTATIONS 172/2013

ISBN 978-952-60-5397-4
ISBN 978-952-60-5398-1 (pdf)
ISSN-L 1799-4934
ISSN 1799-4934 (printed)
ISSN 1799-4942 (pdf)
http://urn.fi/URN:ISBN:978-952-60-5398-1

Unigrafia Oy
Helsinki 2013

Finland

Aalto University

Abstract

Aalto University, P.O. Box 11000, FI-00076 Aalto www.aalto.fi

Author
Aarno Isotalo

Name of the doctoral dissertation
Computational Methods for Burnup Calculations with Monte Carlo Neutronics

Publisher School of Science

Unit Department of Applied Physics

Series Aalto University publication series DOCTORAL DISSERTATIONS 172/2013

Field of research Reactor Physics

Manuscript submitted 21 August 2013 **Date of the defence** 4 December 2013

Permission to publish granted (date) 10 October 2013 **Language** English

☐ **Monograph** ☒ **Article dissertation (summary + original articles)**

Abstract

The work described in this thesis deals with the computational methods and algorithms used in burnup calculations, which model changes in the composition of nuclear fuel under irradiation. While only cases where the neutron transport part of the calculation is handled by the Monte Carlo method are considered, most of the results should also be applicable with deterministic neutronics.

During each step of a Monte Carlo burnup calculation, changes in material compositions are solved by evaluating an explicit solution to the Bateman equations with constant coefficients. Five depletion algorithms capable of doing this while explicitly modeling all of the thousands of nuclides and reactions encountered in burnup calculations were compared. The results are quite conclusive and, together with other studies, show rational approximation based matrix exponential methods to be the best choice for Monte Carlo burnup calculations.

The constant coefficients of the Bateman equations are selected by a coupling scheme that uses one or more steady state neutronics solutions to predict their time development. Because the coefficients must be constant, these predictions are further approximated with their averages. New coupling schemes that use data from the previous step to make higher order predictions are presented. Since the old values are readily available, no additional calculations are required, and the stepwise running time is not affected. The coupling is further improved by dividing the steps into substeps, which are then solved sequentially. Since each substep can use different coefficients for the Bateman equations, this allows piecewise constant, rather than constant, approximation of the predicted behavior. These new methods greatly improve the accuracy obtainable with given step lengths, thus allowing longer steps to be used.

Prior studies have shown that the existing coupling schemes used in Monte Carlo burnup calculations suffer from instabilities caused by spatial xenon oscillations. The new methods are also affected, but it is shown that the simulation models used in these tests actually describe physical xenon oscillations, not a stable state. Thus it is the models, not the methods used to solve them, that are unstable. Regardless, all xenon driven oscillations can be prevented by forcing a mutual equilibrium between the neutron flux and saturated xenon distribution. The equilibrium calculation can be integrated into Monte Carlo neutronics, which provides a simple and lightweight solution that can be used with any of the existing burnup calculation algorithms. However, oscillations driven by nuclides other than xenon may still arise if step lengths are too long.

Keywords nuclear reactor, burnup calculation, computational methods

ISBN (printed) 978-952-60-5397-4 **ISBN (pdf)** 978-952-60-5398-1

ISSN-L 1799-4934 **ISSN (printed)** 1799-4934 **ISSN (pdf)** 1799-4942

Location of publisher Helsinki **Location of printing** Helsinki **Year** 2013

Pages 120 **urn** http://urn.fi/URN:ISBN:978-952-60-5398-1

A! Aalto-yliopisto

Tiivistelmä

Aalto-yliopisto, PL 11000, 00076 Aalto www.aalto.fi

Tekijä
Aarno Isotalo

Väitöskirjan nimi
Computational Methods for Burnup Calculations with Monte Carlo Neutronics

Julkaisija Perustieteiden korkeakoulu

Yksikkö Teknillisen fysiikan laitos

Sarja Aalto University publication series DOCTORAL DISSERTATIONS 172/2013

Tutkimusala Reaktorifysiikka

Käsikirjoituksen pvm 21.08.2013 **Väitöspäivä** 04.12.2013

Julkaisuluvan myöntämispäivä 10.10.2013 **Kieli** Englanti

☐ **Monografia** ☒ **Yhdistelmäväitöskirja (yhteenveto-osa + erillisartikkelit)**

Tiivistelmä

Tässä väitöskirjassa esitelty tutkimus on keskittynyt laskennallisiin menetelmiin ja algoritmeihin, joita käytetään ydinpolttoaineen koostumuksen käytönaikaisia muutoksia mallintavissa palamalaskuissa. Vaikka työssä käsitelläänkin vain tilanteita, joissa laskujen neutroniikkaosio hoidetaan Monte Carlo -menetelmällä, ovat useimmat tulokset sovellettavissa myös deterministisen neutroniikan kanssa.

Monte Carlo -palamalasku muodostuu askelista, joilla muutokset materiaalien koostumuksissa lasketaan evaluoimalla eksplisiittinen ratkaisu vakiokertoimisiin Bateman-yhtälöihin. Työn ensimmäisessä osassa vertailtiin viittä algoritmia, jotka pystyvät tekemään tämän huomioiden eksplisiittisesti kaikki ne tuhannet nuklidit ja reaktiot, joihin palamalaskuissa törmätään. Vertailun tulokset ovat selkeitä ja yhdessä muiden tutkimusten kanssa osoittavat rationaaliapproksimaatioihin pohjautuvat matriisieksponentiaalimenetelmät Monte Carlo -palamalaskuihin parhaiten soveltuviksi.

Bateman-yhtälöiden kerrointen arvot valitaan palama-algoritmilla, joka käyttää yhtä tai useampaa aikariippumatonta neutroniikkaratkaisua ennustaakseen niiden kehityksen. Koska kertoimet halutaan vakioiksi, approksimoidaan näitä ennusteita edelleen niiden keskiarvoilla. Väitöksessä esitetään uusia palama-algoritmeja, jotka käyttävät informaatiota edellisiltä askelilta muodostaakseen korkeamman asteen ennusteita. Koska edellisten askelten arvot ovat valmiiksi saatavilla, tämä ei vaadi ylimääräisiä laskutoimituksia, eikä siten vaikuta askelkohtaiseen ajoaikaan. Tarkkuutta voidaan edelleen parantaa jakamalla askeleet aliaskeleisiin, jotka ratkaistaan peräkkäin. Koska jokainen aliaskel voi käyttää Bateman-yhtälöissä eri kertoimia, tämä mahdollistaa kerrointen ennustetun kehityksen tarkemman seuraamisen. Nämä uudet menetelmät parantavat suuresti saavutettavissa olevaa tarkkuutta mahdollistaen askelpituuksien lisäämisen.

Aikaisemmat tutkimukset ovat osoittaneet vanhojen palama-algoritmien kärsivän spatiaalisista ksenon-oskillaatioista, jotka vaikuttavat myös uusiin menetelmiin. Työssä kuitenkin osoitetaan epästabiileissa testeissä käytettyjen mallien kuvaavan fysikaalisia ksenon-oskillaatioita. Oskillaatioiden taustalla on siis ratkaistavien tapausten, ei mallintamisessa käytettyjen menetelmien, epästabiilisuus. Kaikki ksenon-lähtöiset oskillaatiot voidaan estää pakottamalla saturoitunut ksenon-konsentraatio ja neutronivuo tasapainoon. Tasapainotilan laskeminen voidaan yhdistää Monte Carlo -neutroniikan laskuun, mikä tarjoaa kaikkien palama-algoritmien kanssa toimivan yksinkertaisen ja kevyen ratkaisun. Liian pitkiä askelia käytettäessä voi kuitenkin edelleen esiintyä muiden nuklidien kuin ksenonin ajamia oskillaatioita.

Avainsanat ydinreaktori, palamalasku, laskentamenetelmät

ISBN (painettu) 978-952-60-5397-4 **ISBN (pdf)** 978-952-60-5398-1

ISSN-L 1799-4934 **ISSN (painettu)** 1799-4934 **ISSN (pdf)** 1799-4942

Julkaisupaikka Helsinki **Painopaikka** Helsinki **Vuosi** 2013

Sivumäärä 120 **urn** http://urn.fi/URN:ISBN:978-952-60-5398-1

Preface

The work presented in this thesis was carried out in the Fission and Radiation Physics Group at the Department of Applied Physics in Aalto University and its predecessor, Helsinki University of Technology.

I would like to thank my instructor, Pertti Aarnio, for his invaluable guidance, Prof. Rainer Salomaa for the opportunity to work in the group, and Jarmo Ala-Heikkilä for his continued help and support. I am also grateful to the rest of the research group for the great atmosphere they have maintained and for the stimulating discussions we have shared during meetings, lunches and coffee breaks. Additionally, I would like to express my gratitude towards Jaakko Leppänen and Maria Pusa from VTT, who are the authors of the reactor physics code `Serpent` and the `CRAM` depletion algorithm, respectively. Their innovative work has in no small part enabled my own.

The work has been performed under the NETNUC project of the SusEn programme funded by the Academy of Finland and Fortum Ltd., and the NEPAL project of the SAFIR2014 programme funded by The State Nuclear Waste Management Fund and Finnish nuclear companies.

Helsinki, 18.10.2013,

Aarno Isotalo

Contents

List of Publications

This thesis consists of an overview and of the following publications which are referred to in the text by their Roman numerals.

I A.E. Isotalo, P.A. Aarnio. *Comparison of depletion algorithms for large systems of nuclides*. Annals of Nuclear Energy, 38:2–3, pp. 261–268, February–March 2011.

II A.E. Isotalo, P.A. Aarnio. *Higher order methods for burnup calculations with Bateman solutions*. Annals of Nuclear Energy, 38:9, pp. 1987–1995, September 2011.

III A.E. Isotalo, P.A. Aarnio. *Substep methods for burnup calculations with Bateman solutions*. Annals of Nuclear Energy, 38:11, pp. 2509–2514, November 2011.

IV A.E. Isotalo, J. Leppänen, J. Dufek. *Preventing xenon oscillations in Monte Carlo burnup calculations by forcing equilibrium*. Annals of Nuclear Energy, 60, pp. 78–85, October 2013.

Author's Contribution

Publication I: "Comparison of depletion algorithms for large systems of nuclides"

The author performed the test calculations, analyzed the results and was the main author of the article.

Publication II: "Higher order methods for burnup calculations with Bateman solutions"

The author developed and implemented the new methods, performed the test calculations, analyzed the results and was the main author of the article.

Publication III: "Substep methods for burnup calculations with Bateman solutions"

The author developed and implemented the new methods, performed the test calculations, analyzed the results and was the main author of the article.

Publication IV: "Preventing xenon oscillations in Monte Carlo burnup calculations by forcing equilibrium"

The author performed the test calculations, analyzed the results and was the main author of the article.

1. Introduction

Nuclear fission has been utilized for generating electricity since the fifties, and despite all the slander, it can provide a safe, clean, CO_2 free and economically feasible source of energy, more so now than ever before. Continuous improvements in all aspects of nuclear power have been made possible by development in modeling capabilities, and accurate computational methods remain the foundation of its continued utilization and further development. The importance of computational methods is emphasized in reactor physics as experimental work is often difficult or impossible.

Nuclear reactors and the related physics are too complex for detailed modeling of the entire system at once and must be handled piece by piece. One part of the computational system is formed by burnup calculations, which model long term changes in the composition of nuclear fuel and other materials under irradiation as well as the resulting changes in neutronics properties of the system. Such calculations are typically done at the assembly segment level. The neutronics and material compositions form a complex combined problem that is impossible to solve in truly time-dependent form. Because of this, burnup calculations proceed by sequentially solving the neutronics and changes in material compositions while assuming the other one constant.

Burnup calculations can be divided into 'Monte Carlo' and 'deterministic' based on the type of method used for solving the neutronics. Code systems for Monte Carlo burnup calculation have three distinct parts: The neutronics solver, which calculates steady state neutron flux for given material compositions; a depletion solver, which calculates material changes over a time-step with given microscopic reaction rates; and a coupling scheme, which works as a wrapper algorithm, combining sequential neutronics and depletion solutions to a burnup calculation. The work described in this thesis has focused on the depletion and coupling, which handle all the time dependencies, i.e., the time

integration, of the burnup calculation. Despite most of the existing methods dating back tens of years, this topic has remained relatively unexplored.

Compared to deterministic alternatives, Monte Carlo neutronics are very accurate and problem independent, but slow. Because of this, it pays off to use as accurate and general depletion solvers as possible: The advantages of Monte Carlo are preserved in the whole burnup calculation, and the additional computation cost is still dwarfed by the neutronics. Similarly, the coupling scheme should be as accurate as possible. While any reasonable scheme is exact at the limit of infinitely short steps, better methods allow longer steps to reach given accuracy, meaning that fewer expensive neutronics solutions are required.

Although this thesis is written from the point of view of Monte Carlo burnup calculations, most of the methods and results are also applicable with deterministic neutronics.

1.1 Organization of the thesis

The thesis is organized so that Chapter 2 provides background information and describes general aspects of Monte Carlo burnup calculations to facilitate wider understanding of the topic and putting the author's work into context. The rest of the thesis focuses on different aspects of Monte Carlo burnup calculations, describing the author's work and results.

Chapter 3 presents a detailed description of the different algorithms used for depletion calculations. The focus is on algorithms that can be used to handle a full system of over thousand nuclides without simplifications. Results of Publication I, which evaluates and compares the accuracies and performances of the different algorithms, are presented.

Chapter 4 focuses on the algorithms used for coupling depletion and neutronics solutions to burnup calculations. The existing methods, as well as the new methods proposed by the author in Publications II and III, are presented. Performances of the methods and the underlying error mechanisms are discussed.

Chapter 5 covers the topic of xenon stability in Monte Carlo burnup calculations. The origins and mechanisms of the oscillations are discussed, and the use of integrated equilibrium xenon calculations, suggested in Publication IV for dampening the oscillations, is presented.

Chapter 6 summarizes the work presented herein. Some prospects for future work are also proposed.

2. Monte Carlo burnup calculations

This chapter presents a brief overview of the calculation system and various general aspects of Monte Carlo burnup calculations. Discussion of the author's work, as well as prior research on its topics, is left to chapters 3, 4, and 5.

2.1 Decay and transmutation of nuclides

Some nuclides are naturally unstable, or radioactive, and decay to other nuclides over time. The product might still be unstable, causing it to decay further and so on, creating a chain of decays. Some nuclides can undergo different types of decays leading to different daughter nuclides. Decay chain can thus have branches, which might or might not unite later on.

The different decay modes can be characterized by nuclide and reaction specific partial decay constants $\lambda_{i,k}$, which give the volumetric rate of type k decays of nuclide i as

$$R_{i,k} = \lambda_{i,k} x_i \ , \tag{2.1}$$

where x_i is the atomic density of the decaying nuclide. Decay constants are additive and the total rate of decay is given by the total decay constant $\lambda_i = \sum_k \lambda_{i,k}$. Usually the total decay constant, called simply decay constant, is used instead of the partial ones, and additional constants $b_{i,k} = \lambda_{i,k}/\lambda_i$, known as branching ratios, are defined to specify the relative fractions of different decays:

$$R_{i,k} = b_{i,k} \lambda_i x_i \ . \tag{2.2}$$

Branching ratios can also be generalized between any two nuclides i and j as the fraction of decays of nuclide i that produce nuclide j. Denoting this ratio as $b_{i,j}$, the rate at which nuclide j is produced in the decays of nuclide i is thus $b_{i,j}\lambda_i x_i$. The sum of branching ratios coming from a single nuclide can be greater than one as some reactions, in particular alpha decay, produce two nuclei and spontaneous fission can even produce three.

In addition to natural decay, nuclides can interact with neutrons in various manners causing them to transmute into other nuclides. Interactions with other particles and gamma rays can also cause similar reactions, but they have little importance in burnup calculations of fission reactors and are usually ignored. If needed, such reactions can be handled in the same way as neutron induced ones. While decay always proceeds towards less energetic states, i.e., down to the valley of stability, transmutation reactions can also move towards more energetic states thanks to the binding and kinetic energies of the incident particle. Combining sequential decay and transmutation reactions, any nuclide can, in theory, end up as any other nuclide, although this might require an unrealistic number of unlikely reactions. This results in a web of decay and transmutation reactions, a small part of which is shown in Fig. 2.1.

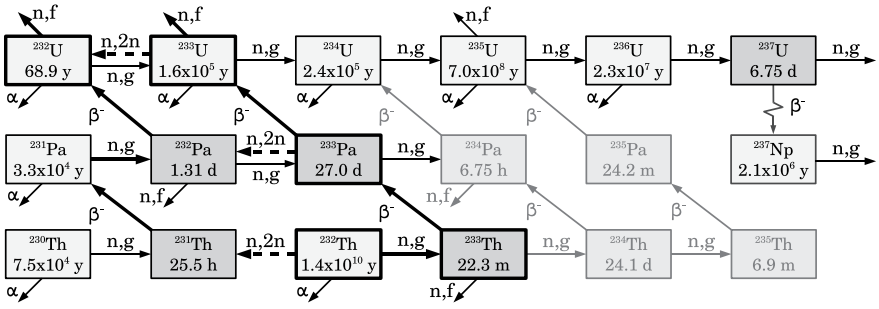

Figure 2.1. Some of the decay and transmutation reactions encountered when thorium bearing nuclear fuel is irradiated. Most of the possible reactions have been omitted for clarity.

The rates of neutron interactions are characterized by microscopic energy dependent cross-sections, $\sigma_{i,k}(E)$, which are nuclide and reaction type specific. The volumetric rate of reaction k of nuclide i with atomic density $x_i(\vec{r})$ caused by incident neutrons of energy E in position $\vec{r}$ is

$$R_{i,k}(\vec{r},E) = x_i(\vec{r})\sigma_{i,k}(E)\phi(\vec{r},E)\ , \tag{2.3}$$

where $\phi(\vec{r},E)$ is the flux of neutrons with energy E in position $\vec{r}$. Total reaction rate is obtained by integrating over all neutron energies. Spatial dependencies need to be discretized for the purpose of numerical solving. This is done by dividing the geometry into suitable volumes where neutron flux is homogenized (i.e. averaged) and atomic densities are assumed spatially constant. The

average macroscopic reaction rate in one such volume is

$$\begin{aligned} R_{i,k} &= \frac{1}{V}\int_V\int_0^\infty R_{i,k}(\vec{r},E)dEdV \\ &= \frac{x_i}{V}\int_V\int_0^\infty \sigma_{i,k}(E)\phi(\vec{r},E)dEdV \\ &= x_i \underbrace{\left(\frac{\int_V\int_0^\infty \sigma_{i,k}(E)\phi(\vec{r},E)dEdV}{\int_V\int_0^\infty \phi(\vec{r},E)dEdV}\right)}_{\sigma_{i,k}} \underbrace{\left(\frac{1}{V}\int_V\int_0^\infty \phi(\vec{r},E)dEdV\right)}_{\phi} \\ &= x_i\sigma_{i,k}\phi = x_i r_{i,k}\ . \end{aligned} \tag{2.4}$$

In the last two forms ϕ is the homogenized one-group flux, and $\sigma_{i,k}$ and $r_{i,k} = \sigma_{i,k}\phi$ are the microscopic homogenized one-group cross-section and reaction rate for the given nuclide, reaction and control volume. These three are the quantities usually used for depletion calculations, and for the remainder of this work ϕ, σ, and r are simply called flux, cross-section and microscopic reaction rate.

It is seen that the macroscopic reaction rate (Eq. 2.4) has the same form as natural decay (Eq. 2.1) if the microscopic reaction rate is interpreted as an effective partial decay constant for the reaction k. This allows decay and transmutation reactions to be treated together in the same way as pure decay systems by defining effective decay constants:

$$\lambda_i^{\text{eff}} = \lambda_i + \phi\sum_k \sigma_{i,k} \tag{2.5}$$

and effective branching ratios between nuclides:

$$b_{i,j}^{\text{eff}} = \frac{b_{i,j}\lambda_i + \sum_k y_{i,j,k}\sigma_{i,k}\phi}{\lambda_i^{\text{eff}}}\ , \tag{2.6}$$

where $y_{i,j,k}$ is the average number of nuclide j produced in reaction k of nuclide i.

2.2 The Bateman equations

Using effective decay constants and branching ratios, the equations governing the decay and transmutation of an arbitrary mixture of N different nuclides in a closed control volume can be written as

$$\frac{dx_i}{dt} = -\lambda_i^{\text{eff}}x_i + \sum_j^N b_{j,i}^{\text{eff}}\lambda_j^{\text{eff}}x_j \qquad \text{for} \qquad i = 1,...,N\ . \tag{2.7}$$

This system of linear first order differential equations is called the Bateman equations after Harry Bateman who first presented a general analytical solution for a linear chain of decay reactions [1], which at that time was the most

general case known. The original equations, which were actually formulated by Ernest Rutherford [2], did not include branching or transmutation, which had not yet been discovered, and the more general system of today is also called decay and transmutation equations or depletion equations.

The difficulties in solving the Bateman equations are twofold. First, the system is large and stiff, i.e., has a large spread of timescales. For example, the JEFF 3.1 library [3] contains decay data for 3851 and cross-sections for 381 nuclides, and the effective half-lives range from microseconds to thousands of years. Second, λ_i^{eff} and $b_{j,i}^{\text{eff}}$ depend on $x_{i=1,\ldots,N}$ through the neutron flux and are thus not actually constants.

The system can be made smaller and less stiff by ignoring unimportant nuclides and by assuming the short-lived ones to decay instantly, thus removing them. It is also possible to lump several low-importance nuclides together and treat them as a single pseudonuclide with averaged properties. The great majority of nuclides can be removed without significantly affecting the neutronics, but which approximations are valid for which nuclides varies from case to case, so the approximations may lead to a degree of problem dependence. The handling and analysis of some results are also complicated as the concentrations of all nuclides are no longer readily available.

The neutron flux reacts to material changes so fast that it is always at an essentially steady state corresponding to the momentary material composition. The second problem can thus be handled by using steady state neutronics solutions with momentary material compositions to predict the time development of the relevant parameters. For a short enough interval even constant extrapolation would be accurate; the predictions simply need to be updated often enough. If all short-lived nuclides have been removed from the system, one can directly integrate dx/dt for $x(t)$ using general purpose numerical integration techniques with every evaluation of the derivatives corresponding to a separate neutronics solution.

Another possibility is to proceed in steps where the Bateman equations are solved assuming the cross-sections and flux, and hence the coefficients of the equations, to remain constant. As changes in the cross-sections and flux are smaller and slower than those in the material compositions, they do not need to be updated as often as the derivatives would. This means that steps can be made longer than with direct integration and fewer neutronics solutions are required. Some of the algorithms for solving the Bateman equations with constant coefficients can also handle the full system of nuclides without any additional approximations, which ensures generality and problem independence.

The tradeoff lies in increased complexity and the computational cost associated with the depletion algorithm. However, even when handling the full system of nuclides, the cost of the depletion calculations is small compared to Monte Carlo neutronics, and practically all Monte Carlo burnup calculations use this approach, although some still make approximations to the system of nuclides.

2.3 Monte Carlo neutronics

What separates Monte Carlo burnup calculations from others is that the neutronics are solved using the Monte Carlo method. Only a brief overview of the method is provided here. More detailed descriptions can be found in various textbooks [e.g. 4, 5] or, for example, in Ref. [6].

In Monte Carlo simulations, individual neutrons are tracked one at a time from emission to eventual removal by capture or leakage. How far the neutron moves before interacting with matter, how it interacts and what results from the interaction are all randomized based on known distributions. These distributions, which are different for each nuclide and depend on the neutron energy, have been the target of a great wealth of experimental and theoretical studies carried out all around the world.

Several organizations maintain evaluated nuclear data libraries, which try to provide a consistent set of data reflecting the best available knowledge on the distributions of reaction probabilities and results. Examples of such libraries include JEFF [3, 7] produced via international collaboration of OECD NEA databank countries, ENDF/B [8] produced by U.S. Cross Section Evaluation Working Group, and JENDL [9] produced by Japan Atomic Energy Agency.

In addition to nuclear data, the neutron transport calculation requires knowledge of the problem geometry, including material compositions, temperatures and densities. Since it is impossible to simulate every neutron in the physical system, an additional normalization condition (e.g., total power or flux density) is also required to determine how many real neutrons one simulated neutron corresponds to.

2.3.1 Progress of the simulation

Depending on the system being simulated, there are two fundamentally different ways to select the initial positions and velocities of the neutrons. Subcritical systems where neutron population is maintained by an external source are simulated using what is called fixed source mode or external source mode. In such

simulations, each neutron is generated from a predefined initial distribution, the source, and followed until absorption or leakage. If the neutron generates secondary neutrons, for example through fission, those secondary particles are simulated after the original. Once all secondary particles have been simulated, a new neutron is sampled from the source distribution and simulated in the same way. This process is repeated until a desired number of neutron histories has been simulated, or some other termination criterion is met.

Critical and supercritical systems, where neutron population is maintained by a fission chain reaction, are simulated using criticality source mode, which is also called fission source mode. The simulation proceeds in source cycles, each consisting of a roughly fixed number of neutron histories. The initial positions and velocities of the neutrons in each cycle are sampled from the fissions caused by the neutrons of the previous cycle. The fission distribution for the first cycle is guessed. Any guess is guaranteed to converge to the correct distribution as the simulation is run, but because the convergence is not instant, a possibly large number of initial unconverged neutron cycles has to be discarded. These discarded cycles are called inactive and the rest active.

Whichever source mode is used, the simulated neutrons form a sample from the neutron population of the system. All physical quantities describing the neutron population, in particular one-group cross-sections and flux, can be estimated by averaging over these individual histories.[1]

2.3.2 Statistics

Since Monte Carlo estimates are stochastic, they will always include some degree of statistical uncertainty. In accordance with the central limit theorem, these uncertainties are proportional to $1/\sqrt{n}$, where n is the number of events that contribute to the estimate. This number depends linearly on the total number of neutron histories, but estimates for averages over small volumes and energy ranges get fewer events per neutron than those for large ones, and more neutrons must be simulated to reach given accuracy. Because of this, Monte Carlo is far less efficient in estimating local than global values.

Typical burnup calculations require estimating the cross-sections and flux for relatively small material regions. This results in relatively large statistical variation in the estimated reaction rates, as well as the material composition calculated using them. However, accuracy of global results, including average

[1] Some non-physical quantities, such as neutron diffusion coefficients, cannot be unambiguously determined, because they do not have a direct counterpart in the real world.

material compositions, is retained because they are still contributed to by the entire geometry, meaning that statistical errors in the compositions and in their effects on the neutronics tend to cancel out.

While the statistical uncertainty behaves as $1/\sqrt{n}$, the computational cost is directly proportional to the number of histories. Halving statistical uncertainty thus requires quadrupling the running time. As the fractions of histories contributing to various estimates and the acceptable levels of statistical uncertainty vary from case to case, running times for different calculations can be very different even when using the same codes in comparable geometries.

So called non-analog Monte Carlo methods, which diverge from the direct simulation of reality, can be used to reduce statistical uncertainty. Some techniques can improve the overall performance, while others provide ways to improve the accuracy of selected estimates at the expense of others. Non-analog methods are important in many applications, but have only a limited significance in burnup calculations.

2.3.3 Advantages and disadvantages of Monte Carlo

Numerous deterministic neutron transport methods [5, 10] provide an alternative to Monte Carlo. These methods are based on a multi-group treatment of energy and various approximations and solution schemes for the spatial and angular dependencies.

Each deterministic method has its own strengths and weaknesses, but overall they are much faster than Monte Carlo. All deterministic methods also avoid the statistical uncertainty associated with the Monte Carlo method, and the difference in performance is especially large when estimating local quantities and multidimensional distributions, for which the statistical efficiency of Monte Carlo is poor. In typical 2D assembly calculations, deterministic methods can be several orders of magnitude faster than Monte Carlo.

The Monte Carlo method, on the other hand, does not require any of the various approximations that are necessary in deterministic calculations. Neutron interaction physics and the whole transport process can be treated to the best available knowledge. The Monte Carlo methods can also handle any geometry consisting of discrete material regions, and while increasing the size and complexity of the geometry slows down the calculation, the relative slowdown is smaller than with deterministic methods. However, if local results are required, deterministic methods may still scale better as the statistical efficiency of Monte Carlo drops.

Rigorous treatment of interaction physics and geometry makes the method not only more accurate, but also problem independent. Deterministic methods, on the other hand, treat neutron interactions and spatial dependencies using approximations and numerical methods that are only valid under certain conditions. Fast and thermal spectra, square and hexagonal lattices and different types of fuel often require different treatments, or at least separate validations. While deterministic codes have been developed for most, if not all, applications, the problem independence of Monte Carlo is still a major advantage: There is no need to use, or learn to use, separate neutronics codes for different problems. This flexibility makes Monte Carlo burnup calculations especially popular in academic research.

The advantages and disadvantages of the neutronics solvers are inherited by burnup calculations using them: Monte Carlo burnup calculations tend to be problem independent and accurate, while deterministic ones can be dramatically faster.

2.4 Error sources and uncertainties

Monte Carlo neutronics can be very accurate, and there are practically no systematic error sources in correctly implemented algorithms. The error mechanisms directly associated with the depletion and coupling calculations are discussed in Chapters 3 and 4, respectively. These parts can also be made very accurate with short enough steps. However, the results of Monte Carlo calculations will always contain a degree of statistical uncertainty, and there are several significant error sources that are 'external' to the calculation system. These errors are associated with nuclear data, geometry models and quantities which the calculation system treats as input.

During the past ten years, there has been wide interest in systematic analysis of propagation of statistical and nuclear data related uncertainties. The magnitude of errors arising from the treatment of the geometry and thermal hydraulics remains relatively unexplored, perhaps because they are very case specific.

2.4.1 Uncertainties from statistics and library data

As discussed in Section 2.3.2, all Monte Carlo estimates contain a level of statistical uncertainty. This affects not only results such as the multiplication factor and power distribution, but also the cross-sections and flux that are used in de-

pletion calculations. As a result, the atomic densities will also contain a degree of statistical uncertainty after the first step. The uncertainty in atomic densities adds to all results of the following neutronics calculations, including those used for the next depletion step, and so on. The uncertainties inherited from earlier steps are called propagated uncertainties.

Another source of uncertainty lies in the nuclear data characterizing the probabilities and results of different neutron interactions. Measuring these probability distributions is difficult, and the available experimental data hence limited in both coverage and precision. Computational models, whose parameters and results are adjusted to fit the experiments, can fill the gaps, but the obtained values are at most as reliable as the experiments they were fitted to. Large parts of the data in evaluated libraries still lack uncertainty estimates, but when such are available, one sigma uncertainties of several percent are the norm rather than the exception.

The most straightforward means of estimating the propagation of uncertainties from statistical variation and library data to final results is the brute force approach: The calculation is repeated a large number of times, each time using a different set of random numbers and nuclear data sampled from the distribution of possible values. These repeats produce independent samples from the distribution of possible results, allowing a straightforward estimation of uncertainties. This method is simple and universally applicable but extremely expensive as it may require the calculation to be repeated hundreds of times. Several schemes have been developed for estimating the uncertainties with only a single set of neutronics solutions [11–14], but they have not yet been widely adapted.

Typical assembly level calculations use sufficient neutron histories in each neutronics solution to keep statistical uncertainties of all relevant quantities well below 1 %, and several studies have shown that unless statistical accuracy is exceptionally low, both direct and propagated statistical uncertainties are small, especially when compared to the uncertainty arising from the library data. [11–13, 15] Estimates of the magnitude of the uncertainties propagated from the library data vary depending on the reactor and fuel types, burnup, accuracy of the data, assumptions made about the accuracy of the data and treatment of correlations between the uncertainties for different quantities.

Estimates for different thermal reactors at 40 – 100 MWd/kgHM burnup include 0.5 – 2 % for k_{eff} [16–19], 1 – 5 % for Doppler coefficients [16, 17] and 1 – 10 % for decay heat [16–20].

Data for high energies is less accurate and the estimated uncertainties for fast reactors correspondingly larger, e.g., 1 – 4 for k_{eff} [16, 17, 20, 21], 3 – 8 % for temperature coefficient [16,17], 10 – 20 % for sodium void coefficient [17,21] and 1 – 20 % for decay heat. Accelerator driven systems used for actinide burning are the worst affected, and uncertainties for various quantities can reach tens of percent [17, 22].

Estimated uncertainties in atomic densities are roughly similar in fast and thermal reactors, depending more on the burnup than the spectrum [16, 18, 20]. The estimates include 1 – 3 % for ^{235}U, 1 – 5 % for ^{239}Pu, 0.1 % for ^{238}U and 0.5 – 6 % for other uranium and plutonium isotopes. The heaviest actinides have largest uncertainties with those for ^{243}Am and $^{243-245}$Cm reaching tens of percent. Uncertainties for fission products are mostly below 1 %.

2.4.2 Statistical bias of the depletion calculations

The Bateman equations are nonlinear with respect to the neutron flux, and since the mean of a nonlinear function of a random variable does not coincide with the value of the function for the mean of the random variable, the calculated compositions will contain bias of a statistical origin [15, 23, 24]. This bias is proportional to $1/n$ and thus small compared to statistical uncertainty, which behaves as $1/\sqrt{n}$, unless statistical accuracy is very poor. For example, the results of Ref. [25] show that the bias in PWR pin cell calculations becomes comparable to the library data uncertainties only when less than a hundred neutron histories are used per step.

In typical assembly level calculations that use millions of neutron histories per step, the bias is negligible, but it might become significant if full core burnup calculations with accurate spatial discretization are ever attempted. Furthermore, if the brute force approach is used to obtain error estimates, statistical uncertainty can, at least in principle, become smaller than the bias. An unbiased estimator for the matrix exponential has been proposed [23, 24] but, as far as the author knows, not used in any burnup code.

2.4.3 Environment effects

Performing a Monte Carlo burnup calculation for the entire core of a power reactor is not feasible without major approximations to the material discretization, and in many applications it is not even possible to define all relevant data about the whole reactor. Thus calculations usually cover only a small part of a reactor, typically an effectively two-dimensional segment of a single fuel assem-

bly, with reflective or periodic boundary conditions to create an infinite lattice. Such calculations ignore the effects of neutron exchange with the surroundings of the solution area. For example, the infinite lattice multiplication factor of a LWR assembly can vary from more than 1.3 to less than 0.8 over its irradiation history. Performing a static k-eigenvalue calculation for such assembly underestimates the contribution from fission spectrum neutrons at low burnups ($k > 1$) and overestimates it at high burnups ($k < 1$).

There are a number of methods that are used to compensate for leakage effects in deterministic burnup calculations, and there have been several studies on applying such methods to the assembly constants generated by Monte Carlo neutronics [e.g. 26–30]. The corrections should also be applied to the one-group cross-sections used in the depletion calculations, but this is still uncommon [31, 32].

Differences between corrected and uncorrected assembly cross-sections from deterministic calculations can be several percent [33, 34], and errors in the cross-sections of some individual nuclides must be at least equally large. This suggests that the propagated errors could be as large as the uncertainties related to library data. Ref. [31] reports tens of percent differences between corrected and uncorrected pin-wise power distribution in a highly heterogeneous MOX assembly, while Ref. [32] shows 300 and 400 pcm reductions in errors for k_{eff} of PWR and pebble bed cores, as well as 1 % reduction in error for assembly-wise power distribution in PWR cores, when using corrected group constants in full core calculations.

The leakage corrections are approximate and cannot account for the true surroundings of the solution geometry, so errors remain even when they are used. Nearby control rods, edges of the core, and neighboring assemblies of different type or burnup can influence the spectrum of neutrons entering the solution area from outside, and hence the spectrum and reaction rates in it. The only way to truly account for these effects is to model the entire reactor, or large enough area around the part of it being studied.

2.4.4 Spatial discretization

The composition and temperature inside a fuel pellet are not spatially constant but change continuously, especially in the radial direction. However, in burnup calculations the geometry must consist of distinct areas with uniform tempera-

tures and atomic densities[2]. Usually each normal fuel pin is treated as a single homogeneous material region, while pins with burnable absorbers are divided into 8 – 10 concentric rings. Most calculations use the same uniform temperature for all fuel regions. Since no spatial discretization can perfectly reproduce the correct continuous distributions, some error is unavoidably generated.

Ref. [36] reports 0.15 % difference in k_{eff} and 0.5 % differences in ^{235}U and ^{239}Pu concentrations between traditional and more detailed spatial discretization for a BWR assembly segment, but very little differences between flat and more detailed temperature distributions. Two recent student projects, on the other hand, found several percent errors in the concentrations of various nuclides from using a flat temperature profile for the fuel in PWR pin cell calculations [37] and up to 0.5 % errors in k_{eff} from using uniform temperature in PWR Gd-pins [38].

2.4.5 Thermal hydraulics and control

The neutronics of a nuclear reactor is coupled not only to fuel depletion, but also to thermal hydraulics and the reactor control system. In most burnup calculations, these feedbacks are ignored and temperatures, thermal expansion, coolant density, control rod positions and coolant boron concentrations treated as input. However, all these parameters depend on the neutronics and cannot be accurately known beforehand. One solution is to treat these dependencies in separate reactor simulation codes. Such codes require homogenized group constants generated by burnup calculations, but the problem does not arise as the group constants are parametrized for the whole range of possible states in terms of variables describing the thermal hydraulics and control variables.

Homogenization is, however, computationally very expensive and reactor simulators add another layer of complexity and include their own error sources, so often burnup calculations are required to produce the desired results directly. In such cases, the required control and thermal hydraulics parameters are usually guessed based on prior experience, and may thus contain significant errors that affect the results. For example, errors in fuel radius, fuel density or coolant density in PWR pin cell calculations cause up to equally large relative differences in reaction rates, atomic densities and group constants [39].

More accurate treatment of thermal hydraulics can be achieved by coupling the neutronics solver with an external thermal hydraulics calculation and solv-

[2]Using so called delta tracking [35] allows arbitrary continuous distributions in transport calculations, but the burnup calculations still require that at least the fuel consists of a finite number of distinct depletion zones.

ing the combined system iteratively [40–44], but this is rare in burnup calculations [45] where the computational cost can become very high. Since such solvers are not exact, some errors related to thermal hydraulics will remain even if they are used. A model for boron concentrations and control rod positions could be coupled in the same way as thermal hydraulics, but since reactivity and control are global phenomena, they cannot be treated accurately at assembly level.

2.5 Applications of Monte Carlo burnup calculations

Burnup calculations or their results are required for practically all analyses involving irradiated nuclear fuel. The fuel composition itself is required as a source term in various design, safety and accident analyses dealing with decay heat or radioactive releases. Because the composition of the fuel affects the neutronics, burnup calculations are also required when analyzing criticality or any other neutronics properties of irradiated fuel or any system containing irradiated fuel. This is also true for full core calculations which are performed with separate reactor simulation codes: While they do not usually track the compositions, they require homogenized few-group constants generated by burnup calculations.

Monte Carlo burnup calculations can be used for all of these applications. Being accurate, flexible and problem independent, they can handle every calculation that deterministic methods can, and can also be used as a reference when testing and validating deterministic codes [46, 47]. However, when applicable deterministic methods are available, the better accuracy of Monte Carlo is often not worth the increased running time. This is especially true for homogenization, i.e., the generation of group constants for full core simulations. Homogenization may require thousands of steps worth of calculations and take months or years of CPU time with Monte Carlo methods. This does not mean that Monte Carlo burnup calculations could not be used, but a decent sized computer cluster is required instead of the tabletop computer that would suffice with fast deterministic methods. Historically, group constant generation with Monte Carlo has been limited to producing some reference results for validating deterministic codes, but recently there has been wider interest in the topic [e.g. 26–30, 48–50], and even large scale homogenization by Monte Carlo has been demonstrated [32, 51].

Outside of full scale homogenization, Monte Carlo burnup calculations can be competitive, but deterministic methods continue to dominate everyday production calculations. These calculations involve specific reactor types, fuels and parameter ranges, for which there are well established, optimized and validated deterministic codes. Even if Monte Carlo burnup codes could be used instead, there is no real reason to do so as long as the errors generated by deterministic neutronics are not a dominant factor. Academic research tends to address more varied topics and ideas as well as innovative and unconventional concepts [e.g. 52–57], where problem independence and flexibility are important. This, and computer time generally being less at premium, makes Monte Carlo burnup calculations more common in academic work, where they appear to be used roughly as often as deterministic ones.

Numerous Monte Carlo burnup calculation codes have been developed by different organizations. Usually these are linked codes that use an external wrapper algorithm to combine independent Monte Carlo transport and depletion calculation programs. MCNP [58], which has been developed at Los Alamos National Laboratory in the USA since the sixties, has become the de facto reference code for a wide range of particle transport applications. It is also the most commonly used neutronics solver in Monte Carlo burnup calculations and is used in numerous linked codes such as Monteburns [59], MOCUP [60], BGCore [61], ALEPH2 [57], MURE [62], BUCAL1 [63], MCODE [64] and MCB [65]. Other linked codes, such as MVP-BURN [66] and the TRITON depletion sequence of the SCALE package [67], use different Monte Carlo neutronics solvers. There are also codes where the entire burnup calculation process has been integrated into a single program. Examples of such codes include MCNPX [68], McCARD [69] and Serpent[3].

2.6 Serpent

Serpent, formerly known as PSG [6], is a continuous energy Monte Carlo reactor physics code developed at VTT Technical Research Centre of Finland since 2004 and publicly distributed since 2009. As of early 2013, Serpent has some 200 users in 88 universities and research organizations in 28 countries [70] and is distributed through OEDC/NEA Data Bank [71]. The transport routines of Serpent have numerous performance improving features. These include the mixed use of ray tracing and delta tracking methods during neutron transport [72], a uniform energy grid for all neutron interaction data [73], pre-

[3]For a complete description of the Serpent code, see http://montecarlo.vtt.fi

calculating energy dependent total cross-sections for various reaction types, and calculation of one-group cross-sections from high resolution spectra rather than tallying them directly.

There are two separate branches of Serpent: Serpent 1 is the official released version and Serpent 2 a complete rewrite still in beta. The main new features of Serpent 2 are a multi-physics interface for efficient coupling with thermal hydraulics codes [74] and a complete redesign of memory management, which lends itself to memory efficient OpenMP-MPI hybrid parallelization and allows the memory hungry optimization features of Serpent 1 to be turned off if needed [75].

All the development and testing in this work, excluding the test calculations of Publication I, was done using Serpent as the platform. Serpent has also been the intended primary application for the results and new methods, but they are by no means Serpent specific and can be applied to any Monte Carlo burnup calculation code.

Serpent 1 includes TTA depletion algorithms implemented by the author. One such algorithm is also found in Serpent 2, but based on the results of Publication I and Ref. [76], CRAM [76] has replaced TTA as the primary depletion algorithm, and TTA is used only for pure decay steps with zero neutron flux. Serpent 2 includes the improved burnup calculation algorithms of Publications II and III which were implemented by the author. The equilibrium xenon method of Publication IV is available in both versions of Serpent.

3. Depletion calculations

There are several methods which can be used for solving the Bateman equations with constant coefficients (Section 2.2). Some of the methods require the removal of unimportant and short-lived nuclides from the system to make it smaller and less stiff, while others are capable of handling the full, unsimplified system with thousands of nuclides and reactions. Due to the slowness of Monte Carlo, there is rarely any reason to use a method that would require simplifying the nuclide system. The accuracies of algorithms for full systems were evaluated and compared in Publication I.

The study presented in Publication I only considered depletion systems of the kind encountered in burnup calculations. Closely related cooling calculations for solving the development of the composition of discharged nuclear fuel at zero flux were not considered. Because the lack of neutron induced reactions makes the system simpler and less interconnected, depletion in zero-flux calculations is generally considered easier then that in burnup calculations. However, while cooling calculations can be, and are, solved with the same algorithm, the system as a whole behaves quite differently and is often observed over much longer periods. This means that in such applications, the relative accuracies and performances of the algorithms might differ from those observed for burnup calculations in Publication I. Performances of the methods when applied to decay-only systems remain a potential topic for future study.

3.1 Transmutation trajectory analysis

Transmutation trajectory analysis (TTA), also known as the linear chains method, is one of the methods for solving the decay and transmutation equations. It is used, for example, in the codes MCB [65] and MCNPX [68] as well as in Serpent 1, where it was implemented by the author. The core of the method is that a complex web of decay and transmutation reactions can be decom-

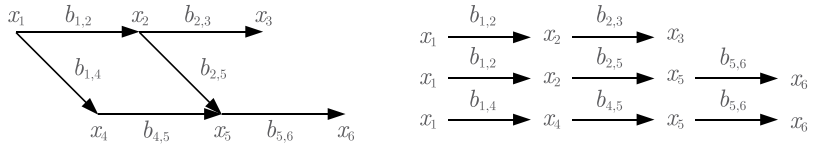

Figure 3.1. Decomposition of a very simple web of decay and transmutation reactions (left) to linear chains (right).

posed into a set of linear chains consisting of all possible routes, or trajectories, through the web as illustrated in Fig. 3.1.

Linear chains are constructed by starting from one nuclide and following all the possible reaction modes. The concentrations of nuclides encountered in each chain are calculated by assuming that only the first nuclide has non-zero initial atomic density. Doing this for each nuclide in the initial composition and superposing the results yields the solution of the original problem.

In a transmutation rich environment, such as a nuclear reactor, the web of decay and transmutation reactions is so complicated that considering all chains is impossible. Fortunately, the great majority of chains are practically meaningless. Only an insignificant number of nuclei pass through them, and the problem is handled by terminating the construction of a chain when its significance is estimated to be small by some criteria. Cyclic chains, e.g., $^{235}\mathrm{U} \xrightarrow{(n,2n)} {}^{234}\mathrm{U} \xrightarrow{(n,\gamma)} {}^{235}\mathrm{U}$..., cannot be linearized even in theory as it would lead to an infinite number of infinitely long chains. Terminating unimportant chains also solves the problem presented by cyclic chains since at most the first few rounds around any loop are significant.

3.1.1 The Bateman solution

A linear chain with n distinct effective decay constants, λ_i^{eff}, and branching ratios, $b_{i,i+1}^{\mathrm{eff}}$, can be solved analytically [1]. Assuming that only the first nuclide has a non-zero initial atomic density, $x_1(0)$, the atomic density of the n:th nuclide after time t is

$$x_n(t) = x_1(0) B_n \sum_{i=1}^{n} \alpha_i^n e^{-\lambda_i^{\mathrm{eff}} t} \,, \tag{3.1}$$

where

$$B_n = \prod_{j=1}^{n-1} b_{j,j+1}^{\mathrm{eff}} \tag{3.2}$$

and

$$\alpha_i^n = \frac{\prod_{j=1}^{n-1} \lambda_j^{\mathrm{eff}}}{\prod_{\substack{j=1 \\ j \neq i}}^{n} (\lambda_j^{\mathrm{eff}} - \lambda_i^{\mathrm{eff}})} \,. \tag{3.3}$$

The Bateman solution fails if the effective decay constants in the chain are not distinct, i.e., if $\lambda_i = \lambda_j$ for some $i \neq j$. This happens in cyclic chains, and might also happen if identical decay constants are input for different nuclides. The simplest possibility is to rely on the weakness of cyclic chains and terminate every chain when a loop is encountered. A more accurate approximation is to introduce small variations to the repeated decay constants to make them distinct, after which Eq. (3.1) can be used. Since the effective decay constants are usually only known to a few digits, these variations are easily kept below the level of uncertainty in initial data.

3.1.2 General analytic solution

A third alternative is provided by a more recent general analytic solution [77], which allows for an arbitrary number of repeated effective decay constants. When this solution is used, the only approximation left is the termination of chains of low importance.

For a chain of n nuclides that has d distinct effective decay constants λ_i^{eff}, each repeated m_i times ($\sum_{i=1}^{d} m_i = n$), the general solution is[1]

$$x_n(t) = x_1(0)\frac{B_n}{\lambda_n^{\text{eff}}} \sum_{i=1}^{d} \lambda_i^{\text{eff}} \alpha_i e^{-\lambda_i^{\text{eff}} t} \sum_{m=0}^{\mu_i} \frac{\lambda_i^{\text{eff}} t}{m!} \Omega_{i,\mu_i - m} \ , \tag{3.4}$$

where

$$\mu_i = m_i - 1 \tag{3.5}$$

is used to simplify notation,

$$B_n = \prod_{j=1}^{n-1} b_{j,j+1}^{\text{eff}} \ , \tag{3.6}$$

$$\alpha_i = \prod_{\substack{j=1 \\ j \neq i}}^{d} \left(\frac{\lambda_j^{\text{eff}}}{\lambda_j^{\text{eff}} - \lambda_i^{\text{eff}}} \right)^{m_j} \tag{3.7}$$

and

$$\begin{aligned} \Omega_{i,j} \quad = \quad & \sum_{h_1=0}^{j} \sum_{h_2=0}^{j} \cdots \sum_{h_{i-1}=0}^{j} \sum_{h_{i+1}=0}^{j} \cdots \sum_{h_d}^{j} \\ & \prod_{\substack{k=1 \\ k \neq i}}^{n} \binom{h_k + \mu_k}{\mu_k} \left(\frac{\lambda_i^{\text{eff}}}{\lambda_i^{\text{eff}} - \lambda_k^{\text{eff}}} \right)^{h_k} \delta \left(j, \sum_{\substack{l=1 \\ l \neq i}}^{d} h_l \right) \ . \end{aligned} \tag{3.8}$$

[1] In the original paper [77] there are two minor notation errors in the final formula for $\Omega_{i,j}$. First, the summation over h_i is not excluded. Second, the summation index in the delta term incorrectly starts from $l = 0$, rather than $l = 1$.

When there are no repeats, i.e., $m_i = 1$ for all i, this reduces to the Bateman solution (Eq. 3.1). In the form presented here, the formula fails if $\lambda_n^{\text{eff}} = 0$, i.e., if the chain ends to a stable nuclide. It would be easy to write the equation in a form where the problematic λ_n^{eff} is canceled. However, it is not necessary as passage can be used to calculate x_n should nuclide n be stable.

3.1.3 Passage and chain termination criteria

Passage, $P_n(t)$, is defined as the fraction of the first nuclide in a chain that has passed through the first n nuclides in the chain at time t [77]:

$$P_n(t) = \sum_{i=n+1}^{m} \frac{x_i(t)}{x_1(0)} = \frac{B_n}{x_1(0)}(1 - \sum_{i=1}^{n} x_i(t)/B_i) \ , \tag{3.9}$$

where m is the number of nuclides in the full chain and is infinite for cyclic chains. The last form can, in theory, be used to calculate the passage, but it is too prone to numerical errors. It is possible to calculate P_n accurately and efficiently at the same time as x_n, for example, by noticing that

$$P_n(t) = \left. \frac{x_{n+1}(t)}{x_1(0)} \right|_{\lambda_{n+1}^{\text{eff}}=0} , \tag{3.10}$$

which can be simplified and rearranged to a form that is very similar to Eq. (3.4).

Unimportant chains can be terminated based on various criteria, but passage provides a particularly robust one: When constructing the chains, daughter j of the n:th nuclide is ignored if

$$b_{n,j}^{\text{eff}} P_n < x_{\text{tot}}(0)/x_1(0) \times \text{cutoff} \ , \tag{3.11}$$

where $x_{\text{tot}}(0)$ is the total initial atomic density and cutoff an input parameter. When using this termination criterion, any chain accounting for less than the fraction cutoff of the total atomic density is ignored. Since a single nuclide might be populated by several chains, this does not guarantee that all absolute errors caused by the termination would be below $x_{\text{tot}}(0) \times \text{cutoff}$, but in practice it does come close to doing so as the production of each nuclide is dominated by at most a few chains. The condition is also efficient in minimizing the number of times Eq. (3.1) or (3.4) needs to be evaluated to bring the absolute errors for all individual nuclides, as well as the sum of absolute errors for all nuclides, below some threshold.

3.1.4 Performance of TTA

The different means of dealing with cyclic chains lead to three variants of the method: basic TTA where cyclic chains are ignored, variation TTA where repeated decay constants are varied to make them distinct and generalized TTA which uses the general analytic solution.[2] Accuracies and performances of the TTA variants were evaluated and compared in Publication I. Fig. 3.2 shows an example of errors generated by each variant after a single step where the initial composition is fresh fuel. Results for old fuel are slightly better. Cutoff of 10^{-15} is sufficient for most, if not all, applications and results in running times of up to a few seconds per depletion solution. With cutoff of 10^{-20}, a single solution can take over 20 seconds and numerical errors start to dominate so that further reducing the cutoff would not improve the results significantly. The general solution is slightly (~10 %) slower than the variation method, which in turn is tens of percent slower than the basic method.

Considering the large uncertainties in library data, the basic version appears accurate enough and might thus be preferable due to its speed. However, since TTA is a slow method, it is only worth using when going for very high accuracy, in which case the variation method should probably be used anyway. Using the general solution rather than varying decay constants had only a minimal effect on results and does not appear to provide any advantage to offset increased running time and complexity.

3.2 Matrix exponential methods

The Bateman equations (2.7) can also be presented in a matrix form

$$\frac{d\vec{x}}{dt} = A\vec{x} \,, \tag{3.12}$$

where

$$A_{i,j} = -\lambda_i^{\text{eff}}\delta_{i,j} + b_{i,j}^{\text{eff}}\lambda_{\text{j}}^{\text{eff}} \tag{3.13}$$

and $\delta_{i,j}$ is the delta function ($\delta_{i=j} = 1$, $\delta_{i\neq j} = 0$). The matrix form has a formal solution

$$\vec{x}(t) = e^{At}\vec{x}(0) \,, \tag{3.14}$$

which employs the matrix exponential notation

$$e^{At} = \sum_{m=0}^{\infty} \frac{1}{m!}(At)^m \,. \tag{3.15}$$

[2] These names are used and suggested by the author, but there are no established names for the different variants.

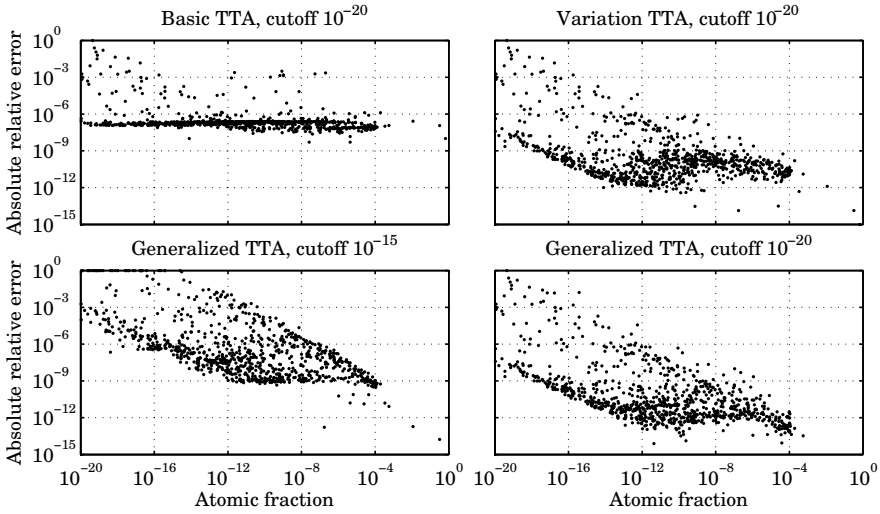

Figure 3.2. Relative errors for the concentrations of individual nuclides in a typical PWR assembly calculation after a single 100 day depletion step starting from fresh fuel. Figure adapted from Publication I.

So-called matrix exponential methods are based on different numerical approximations for evaluating the matrix exponential, which is equivalent to solving the original equations. Matrix exponentials arise in numerous applications and various methods have been developed for evaluating them [78], but only a few are applicable in burnup calculations, and even fewer when solving the full system of nuclides.

3.2.1 ORIGEN

`ORIGEN` [79, 80] is a well known and widely used program for depletion calculations. The method of solution in `ORIGEN` is a power series approximation of the matrix exponential with instant decay and secular equilibrium[3] approximations for handling short-lived nuclides [81]. The same method is used in the modernized `ORIGEN-S` of the `SCALE` package [82]. The use of `ORIGEN` in numerous linked burnup codes, such as `Monteburns` [59], `MOCUP` [60] and `MCODE` [64], makes this probably the most commonly applied depletion algorithm in Monte Carlo burnup calculations.

The solution consists of three phases. First, contributions to the final concentration from short-lived nuclides ($\lambda^{\text{eff}} \leq \ln(0.001)/t$, i.e., $T_{1/2}^{\text{eff}} \lesssim 0.1t$, where t

[3] In secular equilibrium the concentration of a nuclide is at a saturation level defined by the momentary reaction rates and the momentary concentrations of its parents. Such state is reached by nuclides with half-life much shorter than both the observed timescale and the half-lives of their non-short-lived parents.

is the step length) present initially are calculated. This is done by constructing for each nuclide all the populating chains consisting of only short-lived nuclides, and solving these chains the same way as in TTA. The contributions to the final concentrations of short-lived nuclides are saved and the contributions to long-lived nuclides are added to the initial concentrations of those nuclides.

Second, a reduced version of the coefficient matrix A in Eq. (3.14) is constructed by assuming short-lived nuclides to decay instantly[4], thus removing them from the system. The reduced system is then evaluated by truncating the power series for the exponential:

$$e^{At} = \sum_{m=0}^{\infty} \frac{(At)^m}{m!} \approx \sum_{m=0}^{M} \frac{(At)^m}{m!} \; . \tag{3.16}$$

The convergence and numerical stability problems usually encountered when using the power series [78] are avoided because the reduced system does not contain the large matrix elements associated with short-lived nuclides.

Finally, the contributions from long-lived nuclides to short-lived nuclides are solved by assuming these decay and transmutation chains to be in a secular equilibrium at the end of the step:

$$\frac{dx_i}{dt} = \sum_{i=1}^{N} a_{i,j} x_j = 0 \; . \tag{3.17}$$

The resulting greatly reduced form of the system is then solved by iterating

$$x_i^{k+1} = \frac{1}{-a_{i,i}} \sum_{\substack{j=1 \\ j \neq i}}^{N} a_{i,j} x_j^k \; , \tag{3.18}$$

where x_i are the final concentrations obtained from the matrix solution for the long-lived nuclides and unknowns to be solved for the rest. The final concentrations of short-lived nuclides are a superposition of the contributions calculated from other short-lived nuclides and from the long-lived nuclides.

ORIGEN has been verified and validated against experiments, but no prior study on just how accurate it is appears to have been published in the open literature. The accuracy of the algorithm was evaluated in Publication I. Even if the method is far less accurate than TTA, errors for most nuclides are below 1 % (Fig. 3.3) and smaller than the uncertainties in nuclear data. There is, however, no theoretical upper limit for the error caused by multiple instant decay approximations in the populating chain of a nuclide, and it is possible to

[4]The approximation is actually more complex than just instant decay. There is an additional correction that attempts to account for the effects of non-zero half-lives of the removed nuclides by modifying the effective decay constants of their long-lived parents. This feature is not mentioned in the documentations of ORIGEN.

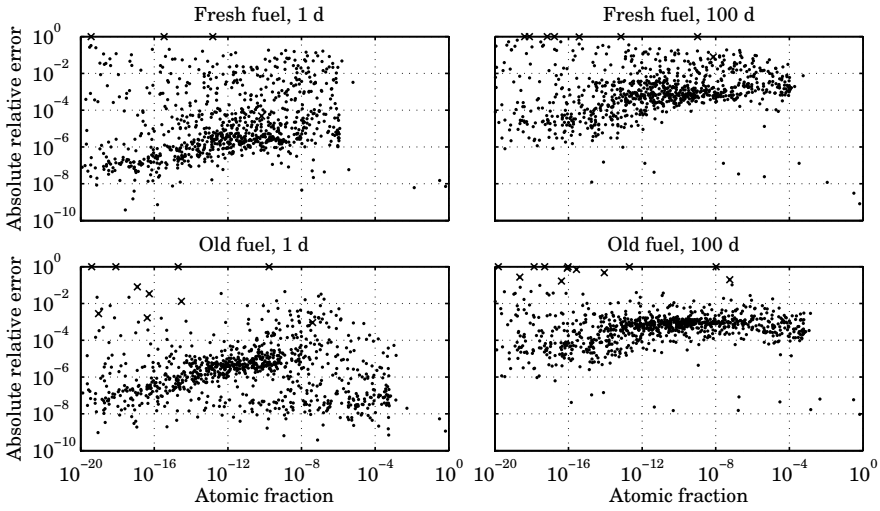

Figure 3.3. Relative errors for the concentrations of individual nuclides in a typical PWR assembly calculation with fresh and old fuel after a single 1 or 100 day depletion step with ORIGEN. Nuclides marked with a cross suffer from errors unrelated to the depletion algorithm. Figure adapted from Publication I.

postulate an artificial chain of decay and transmutation reactions that would results in an arbitrarily large relative error. While such chains do not actually exist in nature, the errors for individual nuclides were observed to reach over 10 %, meaning that the inaccuracy of the algorithm may be a dominant factor for some nuclides in certain situations.

While not the most accurate, the method is extremely fast. For typical depletion steps in burnup calculations, this method is well over hundred times faster than TTA with cutoff of 10^{-15} and over thousand times faster than TTA with cutoff of 10^{-20}.

3.2.2 CRAM

Chebyshev rational approximation method (CRAM) [76] is a new matrix exponential method used by Serpent. It is based on the observation that the eigenvalues of the depletion coefficient matrix, A, are clustered around the negative real axis. This can be exploited by making a Chebyshev rational approximation of the exponential function for the interval $(-\infty, 0]$. The resulting rational function is then decomposed into a pole-residue form (partial fraction decomposition) to avoid numerical instability. When the denominator and numerator orders of the Chebyshev approximation are selected equal and even, the poles form conjugate pairs and the imaginary parts cancel out for a real valued vari-

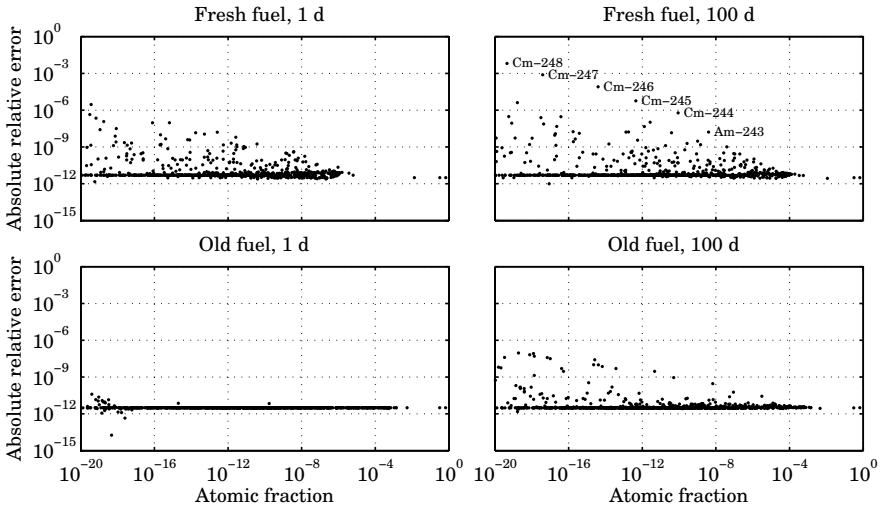

Figure 3.4. Relative errors for the concentrations of individual nuclides in a typical PWR assembly calculation with fresh and old fuel after a single 1 or 100 day depletion step with CRAM of order 14. Figure adapted from Publication I.

able. Thus, an order (k,k) approximation becomes

$$e^z \approx \frac{P_k(z)}{Q_k(z)} = a_0 + \sum_{i=1}^{k} \frac{a_i}{z + \theta_i} = a_0 + 2\,\mathrm{Re}\left[\sum_{i=1}^{k/2} \frac{a_i}{z + \theta_i}\right] . \tag{3.19}$$

where P_k and Q_k are polynomials of order k, whose coefficients have been selected to minimize absolute deviation from exponential function on the negative real axis, a_0 is the limiting value of the approximation at infinity, and a_i and θ_i are the residues and poles. When this approximation is applied to the matrix exponential (Eq. 3.14), it becomes

$$\vec{x}(t) \approx a_0 \vec{x}(0) + 2\,\mathrm{Re}\left[\sum_{i=1}^{k/2} a_i (At + \theta_i I)^{-1}\right] \vec{x}(0) . \tag{3.20}$$

The coefficients a_i and θ_i only depend on the order of the approximation, so they can be pre-calculated and tabulated. A MATLAB script for calculating order $1-13$ coefficient is provided in Ref. [83], and order 14 and 16 coefficients have been published in Ref. [84]. Evaluating the expression thus requires only solving $k/2$ linear systems of the form $(At - \theta_j I)\vec{x}_j = \alpha_j \vec{n}_0$. Due to the special structure of the depletion matrix, this can be done accurately and efficiently by using symbolic LU decomposition [85] and Gaussian elimination [86].

In Publication I, CRAM was found to be far more accurate than either ORIGEN or any form of TTA. An example of the obtained results is shown in Fig. 3.4. These results are in line with the prior theoretical considerations [76] and have been confirmed by latter test calculations [84]. Since there is no theoretical

limit for the imaginary parts of the eigenvalues of the coefficient matrix, neither the theory nor the test calculations provide a strict limit for errors in the general case. However, this seems unlikely to ever be an issue as complex eigenvalues correspond to cyclic chains, which are known to be weak, making the corresponding imaginary parts small. Furthermore, the order of the approximation could be increased to improve accuracy for all nuclides.

In addition to being accurate, CRAM is also remarkably fast. In the tests of Publication I, order 14 approximation for a full system of over thousand nuclides took less than 0.05 s. There is no reason to select any particular order for the approximation as long as it is even. Thus, the order can be used to scale the accuracy versus running time [84]. A range of values might be viable for different applications, but increasing the order past 16 would be very involving as higher order coefficients have not yet been published and are difficult to calculate.

While the TTA methods are easily accurate enough for burnup calculations, CRAM is not only more accurate but also much faster, and there does not seem to be any reason to use TTA over it. The method of ORIGEN is still faster than CRAM and, despite much lower accuracy, might be preferable when the speed of the depletion calculations is a limiting factor. Due to the high computational cost of the neutronics, this could only happen in Monte Carlo burnup calculations if the number of depletion zones is very high compared to the number of source neutrons, implying exceptionally poor statistical accuracy.

3.2.3 Rational approximations from contour integrals

Contour integral based rational approximations [84] are another new method, or a set of methods, similar to CRAM. In these methods the matrix exponential is written as a complex contour integral of the form

$$e^{At} = \frac{1}{2\pi i} \int_\Gamma e^z \left(zI - At\right) dz \,, \tag{3.21}$$

where Γ is a closed contour (a path in the complex plane) around the spectrum of At. This integral is approximated by a quadrature rule, which is then interpreted as a partial fraction decomposition of a rational approximation. This results in exactly the same form as in CRAM (Eq. 3.20) but with a different set of coefficients $(a_\mathrm{i}, \theta_\mathrm{i})$, which depend on the contour and quadrature used as well as the order of the quadrature approximation.

On a practical level, the main difference to CRAM is that the contour integral based coefficients can be calculated on the fly, allowing the order of the approximation to be changed freely. With high enough order, any system should be

approximated as accurately as the arithmetical precision permits, which appears to be 10 or more correct digits for all nuclides [84]. The trade-off for this is that quadrature approximations converge slower than the Chebyshev one, which results in higher running time for given accuracy than with CRAM.

Since CRAM is faster, the contour integral approximations only make sense when aiming for even higher precision. Considering that CRAM of order 14 already appears extremely accurate, it is difficult to see contour based approximations providing a tangible advantage in everyday calculations. However, since the quadrature approximations are still very fast, they do not have any major disadvantages either. It should be noted that higher ($>$16) order CRAM should provide equally high accuracy with shorter running time if the coefficients were calculated and published.

3.3 Other methods

The Krylov subspace method [87] might also be applicable to full systems. The method is based on projecting the matrix exponential on a lower order Krylov subspace, where it is then computed using a Padé approximation. The authors of the method claim that it should be suited for large systems, although they only use it for a system with 221 nuclides and half-lives ranging down to 30 s.

The Bateman equations with constant coefficients can also be solved through numerical integration. This is different from the direct integration mentioned in Section 2.2 as the constant coefficients mean that the neutronics do not need to be recalculated when evaluating the derivatives. Numerical integration has generally been considered unsuited for full systems, but at least the order 5 Radau IIA implicit Runge-Kutta method [88, p. 72] recently implemented to ALEPH2 seems able to handle the full system of nuclides [57]. The results shown in Ref. [57] are too narrow to state anything conclusive about the accuracy or performance of the method, but from what is shown, the method appears accurate and fast enough, although not as fast or accurate as the rational approximation methods presented in Sections 3.2.2 and 3.2.3.

Finally, there are other matrix exponential methods and Runge-Kutta type integration schemes that can be used for depletion calculations but cannot handle the full system of nuclides. Since the primary limiting factor for these methods is the shortest included half-life, any of them could be combined with similar instant decay and secular equilibrium treatment of short-lived nuclides as is used in ORIGEN to handle the full system.

4. Coupling schemes

Sequential neutronics and depletion solutions are combined to a burnup calculation by a wrapper algorithm, which is the outermost layer of the calculation system. In addition to moving data around, the wrapper has a computational function: At each step a limited number of steady state neutronics solutions are used to predict and select representative constant values of cross-sections and flux for the depletion calculations.

The methods used for selecting the constants are called coupling schemes or burnup algorithms. While any reasonable scheme is exact at the limit of infinitely short time steps, better schemes produce more representative constants, allowing longer steps to be used for given accuracy. Publications II and III present improved coupling schemes and demonstrate their advantages over earlier methods.

The notation used for describing different coupling schemes is such that x denotes material compositions in all depletion zones, x_0 the known initial compositions, $\phi(x)$ the spectrum dependent quantities (cross-sections, fluxes, fission yields, etc.) computed with x, $A(\phi)$ the depletion matrix with effective decay constants and branching ratios corresponding to ϕ, $T_0 = 0$ the starting time of the simulation, T_i the ending time of step i and I the total number of steps. The results required from the calculation are x_i and ϕ_i, i.e., the compositions and neutronics parameters at times T_i.

4.1 Traditional coupling schemes

The simplest and most straightforward way to perform burnup calculations is presented in Algorithm 1. At each step, the neutronics is first solved using the beginning of step (BOS) material compositions and changes in material compositions are then solved for the step assuming the microscopic reaction rates to

Algorithm 1

1: **for** $i = 0, \ldots, I-1$ **do** % Loop over steps
2: $\phi_i \leftarrow \phi(x_i)$ % BOS neutronics
3: $x_{i+1} \leftarrow e^{A(\phi_i)(T_{i+1}-T_i)}x_i$ % BOS composition
4: **end for**
5: $\phi_I \leftarrow \phi(x_I)$ % Final neutronics

stay constant at the BOS values. This method is the only option in some codes, e.g., MOCUP [60] and MCB [65], and provided as an alternative by many others.

Algorithm 1 is commonly called the explicit Euler method, or simply the Euler method. However, it is neither the actual Euler method from mathematics nor equivalent to it. The real Euler method would calculate BOS derivatives and assume them constant at each step. Algorithm 1 does this to the microscopic reaction rates, which are not the derivatives of the atomic densities. When the microscopic reaction rates are constant, the derivatives can still change over time due to changing atomic densities. The entire depletion calculation is a result of making the microscopic reaction rates, not the derivatives, constant. Other methods are often referred to in the same way, using the names of 'similar' schemes. The resulting names are sometimes ambiguous.

While simple, using BOS reaction rates for the entire step is a coarse approximation, and various predictor-corrector methods are usually used instead of the simple Euler method. These methods require solving the neutronics twice for each time step, but are usually so much more accurate that step lengths can be more than doubled. There are three predictor-corrector methods in common use:

Algorithm 2 relies on the errors in atomic densities calculated using the BOS and end of step (EOS) reaction rates to cancel out when they are averaged. This method is used, for example, by BUCAL1 [63] and MCODE [64].

Algorithm 3 assumes that the instantaneous reaction rates at the middle of a step would best represent the entire step. This method is used, for example, by Monteburns [59], MCNPX [68] and the TRITON depletion sequence in the SCALE package [67].

Algorithm 4 uses the averages of BOS and predicted EOS cross-sections and fluxes, which can be seen as the average from a linear interpolation between them. This is the default method of Serpent.

Algorithm 2

for $i = 0, \ldots, I-1$ **do** % Loop over steps
% — Predictor —
$\phi_i \leftarrow \phi(x_i)$ % BOS neutronics
$x^{\mathrm{p}}_{i+1} \leftarrow e^{A(\phi_i)(T_{i+1}-T_i)} x_i$ % Predicted EOS composition
% — Corrector —
$\phi^{\mathrm{p}}_{i+1} \leftarrow \phi(x^{\mathrm{p}}_{i+1})$ % EOS neutronics
$x^{\mathrm{c}}_{i+1} \leftarrow e^{A(\phi^{\mathrm{p}}_{i+1})(T_{i+1}-T_i)} x_i$ %
$x_{i+1} \leftarrow (x^{\mathrm{p}}_{i+1} + x^{\mathrm{c}}_{i+1})/2$ % Corrected EOS composition
end for
$\phi_I \leftarrow \phi(x_I)$ % Final neutronics

Algorithm 3

for $i = 0, \ldots, I-1$ **do** % Loop over steps
% — Predictor —
$\phi_i \leftarrow \phi(x_i)$ % BOS neutronics
$x^{\mathrm{p}}_{i+1/2} \leftarrow e^{A(\phi_i)(T_{i+1}-T_i)/2} x_i$ % Predicted midstep composition
% — Corrector —
$\phi^{\mathrm{p}}_{i+1/2} \leftarrow \phi(x^{\mathrm{p}}_{i+1/2})$ % Midstep neutronics
$x_{i+1} \leftarrow e^{A(\phi^{\mathrm{p}}_{i+1/2})(T_{i+1}-T_i)} x_i$ % Corrected EOS composition
end for
$\phi_I \leftarrow \phi(x_I)$ % Final neutronics

Algorithm 4

for $i = 0, \ldots, I-1$ **do** % Loop over steps
% — Predictor —
$\phi_i \leftarrow \phi(x_i)$ % BOS neutronics
$x^{\mathrm{p}}_{i+1} \leftarrow e^{A(\phi_i)(T_{i+1}-T_i)} x_i$ % Predicted EOS composition
% — Corrector —
$\phi^{\mathrm{p}}_{i+1} \leftarrow \phi(x^{\mathrm{p}}_{i+1})$ % EOS neutronics
$\bar{\phi} \leftarrow (\phi_i + \phi^{\mathrm{p}}_{i+1})/2$ % Average neutronics
$x_{i+1} \leftarrow e^{A(\bar{\phi})(T_{i+1}-T_i)} x_i$ % Corrected EOS composition
end for
$\phi_I \leftarrow \phi(x_I)$ % Final neutronics

Algorithm 3 is called the midpoint method or midstep method. There are no established names for Algorithms 2 and 4. There exists also a further variant of Algorithm 4, in which the linear interpolation is done to the microscopic reaction rates, rather than the cross-sections and flux separately. The resulting differences are small and this variant is not considered further.

4.2 New methods

All the above methods can be seen as using some combination of linear interpolations (the corrector of Algorithm 4) and constant extrapolations (everything else) to approximate the behavior of the cross-sections and flux during the step. Since the depletion calculations require the cross-section and flux to be constant at each step, the linear interpolation is further approximated by its average.

Publication II presents a new set of higher order coupling schemes based on using the previous step cross-sections and flux, i.e., the BOS values from the previous step, in addition to the BOS and EOS values of the current step. This allows linear extrapolation and quadratic interpolation to replace the constant and linear ones on the predictor and corrector, respectively. Even if these higher order estimations are further approximated by step averages for the depletion calculations, they should lead to more representative average and better results. Since the previous step values are readily available, no additional neutronics solutions are needed and the stepwise running time is not affected.

Errors that result from having to use step average cross-sections and flux for depletion can be reduced by using substeps as suggested in Publication III. The idea is that while the depletion calculations require microscopic reaction rates to stay constant, nothing requires the whole step to be covered by a single depletion calculation. To take advantage of this, steps are divided into substeps that are solved sequentially. While each substep has to use constant reaction rates, the constants can be selected independently for each of them, allowing piecewise constant, rather than constant, approximation of the behavior predicted for the cross-sections and flux. This is illustrated in Fig. 4.1.

Algorithm 5 presents a general predictor-corrector method combining the higher order estimates and substeps, with '$\mathrm{polyfit}()$' corresponding to the selected predictor and corrector orders. On predictor steps this can be constant extrapolation (`CE`) or linear extrapolation (`LE`) and on corrector steps linear interpolation (`LI`) or quadratic interpolation (`QI`). It is also possible to skip the corrector step and set $x_{i+1} = x^{\mathrm{p}}_{i+1}$. This leads to the six possible combinations

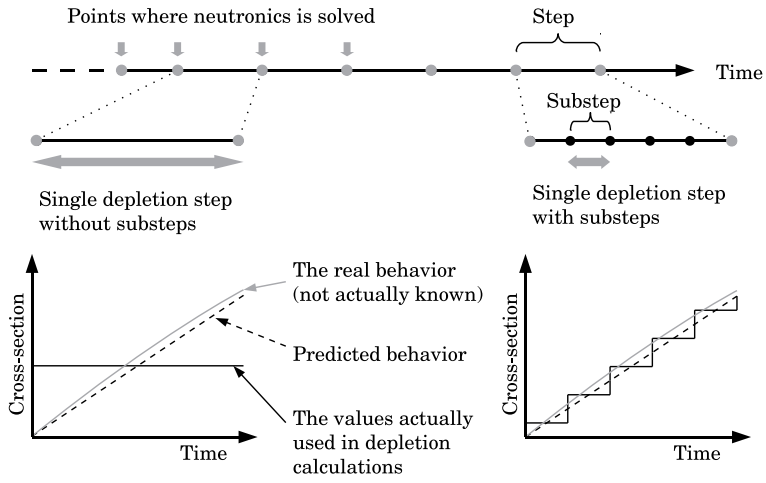

Figure 4.1. Illustration of the substep method. Using substeps allows the predicted behavior of the cross-sections to be followed more closely in the actual depletion calculations.

Algorithm 5

for $i = 0, \ldots, I-1$ **do** % Loop over steps
% — Predictor —
$\phi_i \leftarrow \phi(x_i)$ % BOS neutronics
$\phi(t) = \text{polyfit}(\phi_{i-1}, \phi_i)$ % Extrapolation polynomial
$x_{i+1,0} \leftarrow x_i$
for $s = 0, \ldots, S^{\text{p}} - 1$ **do** % Loop over substeps
$\bar{\phi} = \frac{1}{T_{i+1,s+1} - T_{i,s}} \int_{T_{i+1,s}}^{T_{i+1,s+1}} \phi(t) dt$ % Substep average
$x^{\text{p}}_{i+1,s+1} \leftarrow e^{A(\bar{\phi})(T_{i+1,s+1} - T_{i,s})} x_{i+1,s}$ % Depletion over substep
end for
$x^{\text{p}}_{i+1} = x_{i+1,S^{\text{p}}}$ % Predicted EOS composition
% — Corrector —
$\phi^{\text{p}}_{i+1} \leftarrow \phi(x^{\text{p}}_{i+1})$ % EOS neutronics
$\phi(t) = \text{polyfit}(\phi_{i-1}, \phi_i, \phi^{\text{p}}_{i+1})$ % Interpolation polynomial
$x_{i+1,0} \leftarrow x_i$
for $s = 0, \ldots, S^{\text{c}} - 1$ **do** % Loop over substeps
$\bar{\phi} = \frac{1}{T_{i+1,s+1} - T_{i,s}} \int_{T_{i+1,s}}^{T_{i+1,s+1}} \phi(t) dt$ % Substep average
$x_{i+1,s+1} \leftarrow e^{A(\bar{\phi})(T_{i+1,s+1} - T_{i,s})} x_{i+1,s}$ % Depletion over substep
end for
$x_{i+1} = x_{i+1,S^{\text{c}}}$ % Corrected EOS composition
end for
$\phi_I \leftarrow \phi(x_I)$ % Final neutronics

referred to as CE, LE, CE/LI, CE/QI, LE/LI and LE/QI, using the abbreviations above.

The substep division ($T_{i,s}$) can be selected arbitrarily and could differ from step to step, but only cases where each step is divided into an equal, although possibly different for the predictor and corrector, number of equidistant substeps have been considered thus far. The substep average integrals (lines 7 and 17 in Algorithm 5) do not require significant computational effort as they reduce to weighted sums of the pointwise values to which the polynomial fit was made. The weights are independent of the nuclide and reaction in question, meaning that they only need to be calculated once per substep.

Using only one substep is equal to not using substeps, and of the combinations in the previous paragraph, CE and CE/LI without substeps reduce to Algorithms 1 and 4, which have been in routine use before. The four other combinations can be seen as new methods or higher order versions of Algorithm 4.

At the first step previous step cross-sections and flux are not available, and thus linear extrapolation and quadratic interpolation have to be replaced with constant extrapolation and linear interpolation. Quadratic interpolation is also converted to a linear one at the second step because the step-like changes from the initial formation of short-lived nuclides lead to poor accuracy when the fresh fuel values are used in simple polynomial fits. The same problem affects linear extrapolation, but constant extrapolation is such a bad approximation that it performs worse even at the second step, where linear extrapolation holds poorly. The higher order algorithms must also be restarted after any abrupt changes in the flux or cross-sections. However, apart from the beginning of irradiation, such changes can only occur if the normalization, geometry or materials are changed mid-irradiation by the user.

4.3 Other coupling schemes

There is another new coupling scheme called Projected Predictor Corrector (PPC) [89]. At each step the method performs the same stages as Algorithm 2, but follows this by recalculating the EOS absorption cross-sections and concentrations of ^{155}Gd and ^{157}Gd by assuming linear correlation between their atomic densities and microscopic reaction rates. This requires no transport calculation and has only a very small effect on stepwise running time. Improvement in the results for gadolinium seems comparable to the higher order methods. The au-

thors of the method suggest that the correction could be applied to all nuclides and not just gadolinium, but how well this would work is unclear.

Finally, there are three recently published coupling schemes [90, 91] that have been developed to reduce or remove spatial oscillations encountered in large geometries. It seems unlikely that any of them would be competitive in typical burnup calculations where spatial stability is not a factor and discussion of these methods is left to Chapter 5, which addresses stability.

4.4 Performance of the coupling schemes

Although the traditional coupling schemes of Section 4.1 date back tens of years, their performances have not been comprehensively compared in the open literature. Two studies [92,93] comparing Algorithms 2 and 3 show the first one to be preferable. The studies also show both of them to be preferable over Algorithm 1. As far as the author knows, no comparison of Algorithms 2 and 4 or Algorithms 3 and 4 has been published.

The errors in any coupling scheme falling under Algorithm 5 can be attributed to two sources. First, inaccuracy in the extrapolated or interpolated estimates for the behavior of the cross-sections and flux. Second, having to further approximate these estimates as constant or piecewise constant. In predictor-corrector methods, both the predictor and the corrector have these errors and the errors from the predictor contribute to the inaccuracy of the estimated behavior on the corrector.

Magnitude of the error resulting from the first component is determined by the accuracy of the estimated behavior, i.e., by the predictor and corrector orders, and those from the second by the length of the averaging interval, i.e., by the number of substeps. These error components, and hence the effects of substeps and the order of the method, are largely independent. The total error from coupling is a superposition of the two components.

4.4.1 Effects of the predictor and corrector orders

The predictor and corrector orders affect the results mostly through the step average reaction rates, which dominate the behavior of nuclides with half-life longer than the step-length. Figs. 4.2 and 4.3 show the absolute and relative errors in the atomic densities of ^{155}Gd and ^{235}U calculated with different cou-

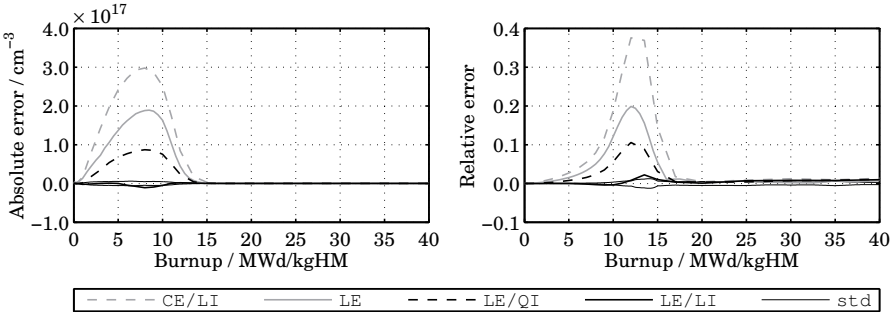

Figure 4.2. Absolute and relative errors for ^{155}Gd atomic density in a typical PWR assembly calculation. The initial atomic density of ^{155}Gd is 1.8×10^{19} cm^{-3}. LE used 43 steps while the rest used 22. Figure adapted from Publication II.

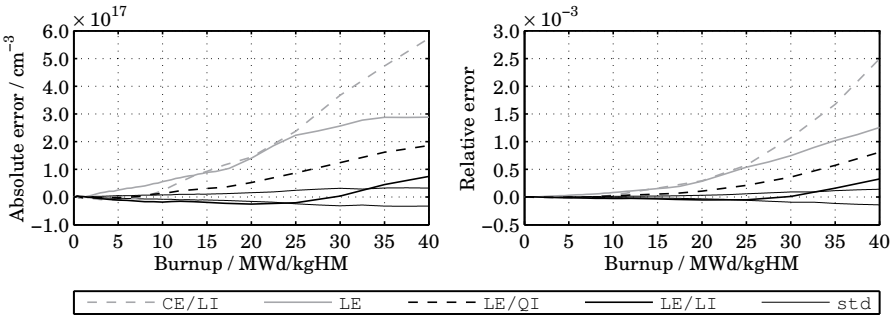

Figure 4.3. Absolute and relative errors for ^{235}U atomic density in a typical PWR assembly calculation. The initial atomic density of ^{235}U is 9.1×10^{20} cm^{-3}. LE used 43 steps while the rest used 22. Figure adapted from Publication II.

pling schemes for a typical PWR assembly segment[1]. Results for CE and CE/QI are not presented as they performed routinely worse than CE/LI, which is an established method. The difference between CE/LI and CE/QI was consistent but not particularly large. CE, on the other hand, produced very poor results, which is not surprising as the method is known to be of low accuracy.

The order of accuracies is mostly the same for other long-lived nuclides, step lengths and test cases. Short-lived nuclides, on the other hand, were barely affected by the orders of the methods. Significant differences in the results for them were observed only when the gadolinium absorber was being depleted, and the resulting rapid changes in the thermal flux led to multiplicative effects on the fission poisons.

Using linear, rather than constant, extrapolation on the predictor steps greatly improves the results. The difference between constant and linear extrapolation is especially dramatic when they are used without corrector steps,

[1]These figures, as well as other results presented in this chapter, only include the error component caused by the time-discretization in the calculations. They have been obtained by comparing the results to reference calculations with very short steps.

and LE might actually provide a viable alternative to the predictor-corrector methods. Due to lower stepwise running time, it would have a major advantage in cases where the required output frequency, rather than the accuracy of the results, is the limiting factor for step lengths.

The effects of using quadratic interpolation on the corrector were mixed and predominantly negative. The reason is that linear corrector results in systematic cancellation of errors, which in most cases improves the results past those from a quadratic one. Due to this reliance on cancellation of errors, LE/LI can actually lose accuracy when step length is decreased. Quadratic interpolation should never be used with constant extrapolation, but the choice between LE/LI and LE/QI is a non-trivial one. While LE/LI was all around more accurate, Figs. 4.2 and 4.3 give an over-optimistic image about its performance. Overall, the difference between accuracies of LE/LI and LE/QI was small, and the slightly worse accuracy of LE/QI might be offset by its more predictable behavior in respect to step lengths.

LE/LI and LE/QI provide a clear advantage over CE/LI, and LE an even larger advantage over CE, but while the higher order estimates do not affect the stepwise running time, they do have one downside: The methods are somewhat sensitive to changes in step lengths. Extrapolating the behavior of a short step to a much longer one amplifies errors and statistical variation. Quadratic interpolation is also affected if the lengths of consecutive steps are very different. As a result, the step lengths of successive steps should not differ by more than a factor of five. Even better if the changes are limited to a factor of two.

4.4.2 Effects of substeps

Short-lived nuclides are in secular equilibrium with their long-lived precursors, and their concentrations are at all times determined by the momentary reaction rates. In particular, the EOS concentrations are determined by the EOS reaction rates. The depletion calculations, however, use average cross-sections and flux which differ from the EOS values, leading to errors for short-lived nuclides. Substeps reduce the length over which cross-sections and flux are averaged, making the difference between the average and momentary values smaller. The differences and the resulting errors are, to a first order approximation, proportional to the substep length, i.e., inversely proportional to the number of equidistant substeps. This is most clearly visible with short-lived fission products, such as ^{99}Mo and ^{131}I in Fig. 4.4, whose production is dominated by neutron induced reactions but removal by decay. For fission poisons

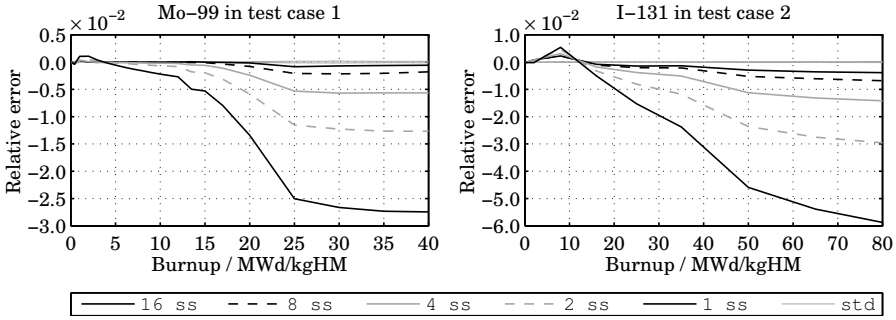

Figure 4.4. Relative errors in atomic densities of ^{99}Mo and ^{131}I in two different PWR test cases with different numbers of substeps on both the predictor and corrector. Figure adapted from Publication III.

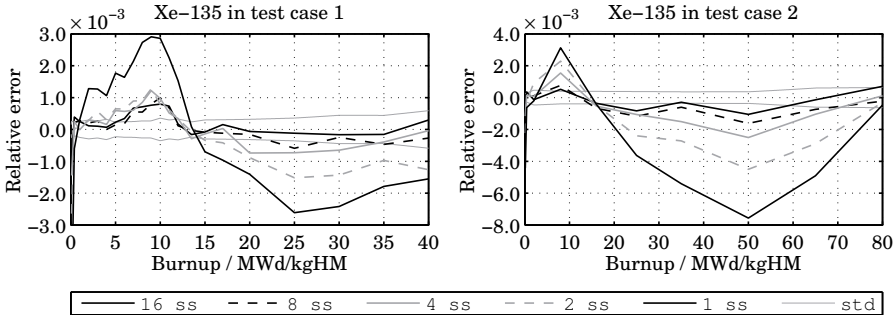

Figure 4.5. Relative errors in atomic density of ^{135}Xe in two different PWR test cases with different numbers of substeps on both the predictor and corrector. Figure adapted from Publication III.

with a large cross-section, e.g., ^{135}Xe in Fig. 4.5, the errors are smaller as both production and removal rates depend on the flux, and the errors in them partly cancel out.

For the momentary, in addition to average, reaction rates to have a significant effect, the effective half-life of the nuclide only needs to be of the same order as the step length. Since the steps in burnup calculations tend to be rather long, this affects many nuclides that are not short-lived in the traditional sense. Examples of such nuclides include all plutonium isotopes (Fig. 4.6) and even ^{235}U to an extent. However, for the nuclides that are not truly short-lived, the average reaction rates, and thus the predictor and corrector orders, also have a significant effect.

The longer the half-life of a nuclide is compared to the step lengths, the less it is affected by momentary reaction rates. However, even if the averages of the cross-sections and flux get correct values, the resulting average microscopic and macroscopic reaction rates are not exact as the product of averages does not generally equal the average of products. Shortening the averaging interval also

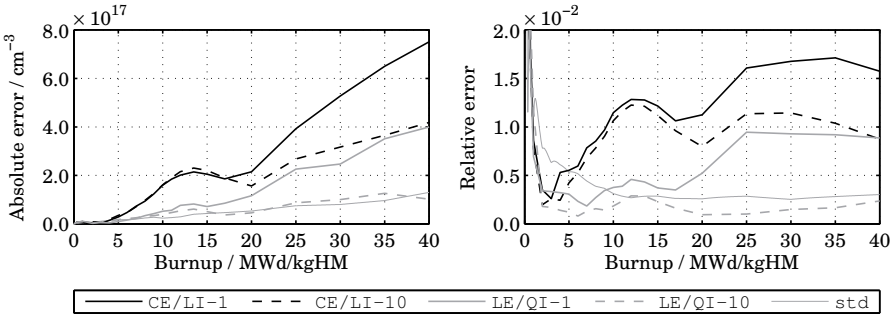

Figure 4.6. Sums of the absolute values of absolute and relative errors for the plutonium isotopes 238 – 242 in a typical PWR assembly calculation with 22 steps and 1 or 10 substeps on both the predictor and corrector. Figure adapted from Publication III.

reduces this error, so substeps can directly affect even truly long-lived nuclides, but this mechanism seems to have only a minor effect.

While Figs. 4.4 and 4.6 use an equal number of substeps on both the predictor and the corrector, the effects of substeps actually arise almost purely from the corrector. Constant extrapolation is not affected by substeps at all, and even on a linear predictor the effect of substeps is only barely detectable. However, when linear extrapolation is not followed by a corrector, i.e., in `LE`, substeps on it have similar effect on the results as those on the corrector steps of a predictor-corrector method.

Using substeps multiplies the number of depletion calculations that have to be performed and thus the amount of time required for them. The relative slowdown depends on numerous factors from the used transport and depletion algorithms to the problem being modeled, but on a very high level, it comes down to the relative speeds of the neutronics and depletion calculations. Because the neutronics dominate running time in typical Monte Carlo burnup calculations, substeps do not usually cause a significant slowdown. However, with many depleted materials and few source neutrons, the slowdown can become noticeable even when using `CRAM` for depletion.

The optimal number of substeps depend on various factors and must be considered on a case by case basis. As a rule of thumb, five substeps should always worth it on the corrector, unless the number of source neutrons is very low compared to the number of depletion zones or the used depletion algorithm very slow. Furthermore, there is usually no pressing need to use more than five of them on the corrector and one or two on the predictor as they already remove 80 % of the errors caused by using constant coefficients, and other parts of the calculation are not exact either.

The effects of using substeps are similar to the linear rate method [92], which works as CE/LI, but instead of using a piecewise constant approximation, it allows truly linear change on the corrector. This is made possible by the use of numerical integration, which does not require the approximation of constant coefficients. Traditionally, Numerical integration has not been usable with a full system of nuclides, so the ability to use any of the specialized depletion algorithms of Chapter 3 gives substeps a major advantage over the linear rate method. This might change as the newly added Radau IIA numerical integration scheme in ALEPH2 actually seems to be able to handle the full system with constant coefficients [57]. If the method retains acceptable performance with non-constant coefficients, it could allow using the linear rate method with a full system and could also be combined with the higher order methods. In any case, only limited improvement over substeps is possible as they already remove nearly all of the error resulting from the use of constant coefficients and cause only minor slowdown.

5. Spatial stability

Burnup calculations have typically focused on assembly segments and other relatively small geometries. As computers and algorithms develop, calculations involving research reactors, entire assemblies and even simplified models of power reactors are becoming increasingly common. While nearly all modeled geometries have been too small or too crudely discretized for spatial oscillations to occur, applications are reaching a point where this in no longer the case.

All coupling schemes commonly used in Monte Carlo burnup calculations have been found to suffer from xenon driven numerical oscillations in large geometries [94,95]. The higher order methods introduced in Section 4.2 are also affected. Publication IV presents efficient means to stabilize the calculations regardless of the coupling scheme used.

5.1 Physical xenon oscillations

^{135}Xe has a very large thermal absorption cross-section and a high cumulative fission yield, giving it a profound effect on neutronics. The combined direct yield of ^{135}Xe ($T_{1/2} \approx 9.2\,\text{h}$) and $^{135\text{m}}$Xe ($T_{1/2} \approx 15\,\text{min}$) from thermal fissions is only 0.2 % while their precursors ^{135}Sb, ^{135}Te and ^{135}I have a combined yield of 6 %. ^{135}Sb and ^{135}Te decay to ^{135}I in seconds, but ^{135}I has a half-life of 6.6 h. Because of the buffer formed by ^{135}I, changes in the flux affect xenon production with a delay, whereas removal, which is dominated by absorption, reacts instantly.

If the flux is tilted, the immediate effect is that in the areas of high flux reactivity starts to increase as xenon is depleted and in the areas of low flux reactivity decreases as xenon builds up. These changes reinforce the flux tilt, which in turn leads to even larger changes in reactivity. Over time, ^{135}I concentrations stabilize and the xenon concentration in high flux areas starts to increase while that in the low flux decreases. Eventually the flux tilts the opposite way and the cycle repeats.

5.2 Numerical xenon oscillations

Burnup calculations aiming to follow long term development use step lengths much longer than the timescale involved in physical xenon oscillations. Due to long steps, ^{135}I and ^{135}Xe concentrations have time to reach saturation levels at each step. This makes the physical xenon oscillation mechanism impossible. Instead, if the flux is tilted, the areas of high flux will get high xenon concentration during the following depletion step and the other way around. This in turn means that in the next neutronics solution the flux will tilt the opposite way, leading to an oscillation.

Since predictor-corrector methods use two neutronics solutions per step, they can exhibit more complex behavior. For example, the following mechanism is possible for the constant-linear predictor-corrector method (Algorithm 4 in Section 4.1). If one end of the geometry has, in the absence of xenon, higher reactivity than the other, it gets higher flux at the beginning of the step (BOS). This results a in high end of step (EOS) xenon concentration and, if the base reactivity difference is not too large, low EOS flux. The other end has opposite fluxes. On the corrector, the average of BOS and EOS fluxes is used, resulting in moderate flux across the entire geometry and thus even xenon concentrations. On the next step, the more reactive area again has higher BOS flux as it has equal xenon as the other end but higher base reactivity, and the cycle continues.

The values that are collected as results from the calculation are the corrected material compositions and the corresponding neutronics, which are calculated at the beginning of the next step. Because they always come from the same phase of the oscillation, all of these may look stable, but the materials were depleted with flux that might not at all represent the ones in the output. The material compositions can also be completely off as the averaged flux used to calculate them is in no way guaranteed to represent the correct one.

Since the results may look stable, comparing the results from successive steps is not sufficient to determine stability. The most direct way to detect the predictor-corrector oscillations is to compare the local predictor and corrector fluxes. Unfortunately, universal or exact threshold for how large a difference implies instability cannot be established as varying differences between the predictor and corrector values are a normal part of predictor-corrector methods.

5.3 Xenon oscillations in Monte Carlo burnup calculations

Earlier studies [94, 95] show that calculations with existing methods exhibit numerical xenon oscillations even with fairly short steps, and conclude that all of the methods are unstable. The new higher order methods also produce similar oscillations. However, in Publication IV it was shown that accurate solutions to simulation models like those used in Refs. [94] and [95] exhibit physical xenon oscillations. An example of the behavior is shown in Fig. 5.1. The models, while geometrically simple, are like those used in typical burnup calculations, only longer in the axial direction.

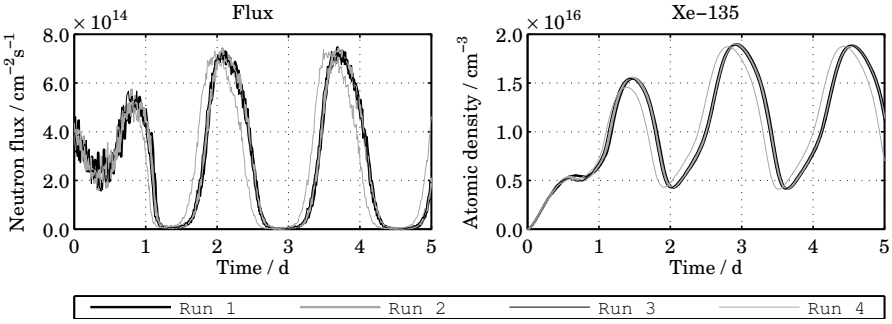

Figure 5.1. Neutron flux and xenon concentration in one end of a 4 m long pin cell with reflective boundary conditions and a 0.1 percent point enrichment difference between the ends. Four identical runs with different random number sequences and 15 min steps. Figure adapted from Publication IV.

Usually burnup calculations use much longer steps, so the physical oscillation mechanism is replaced with a numerical one, but this does not change the fact that the simulation models actually describe physical xenon oscillations. Based on this, it is concluded that the root of the xenon oscillations encountered with long steps is not in the burnup algorithms. The algorithms can handle the computational model for each of the tested cases when used correctly, i.e., with steps that are short enough compared to the timescale of the solution. The real problem is that due to omitting feedback and control mechanisms, the simplified computational models do not even describe the stable state we are trying to simulate.

It should be emphasized that while the models for large and small geometries are fundamentally similar, spatial oscillations are only possible when fluxes in different parts of the geometry are loosely coupled. This requires distances many times larger than the neutron migration length, which is about 6 – 8 cm for light water reactors and 10 – 20 cm for other reactor types.

5.4 Equilibrium xenon calculation

The above does not change the fact that we want to obtain 'stable solutions' with the simplified computation models, but since the problem is caused by the model, the model itself must be changed to make it stable. Perhaps the simplest way to achieve this is to directly require that the flux and saturated xenon concentrations must remain in a mutual equilibrium.

Deterministic full core codes do this via wrapper algorithms that use multiple neutronics solutions to find the equilibrium xenon distribution and the corresponding flux, which is then used for depletion. The same algorithms could also be used with Monte Carlo neutronics and a wrapper algorithm optimized for Monte Carlo has been developed [96]. However, since the equilibrium has to be iterated every time the neutronics is solved, the required number of neutronics solutions is multiplied. For fast deterministic codes this is not necessarily a serious problem, but the use of Monte Carlo burnup calculations is strongly limited by the high computational cost of the neutronics, making such wrapper algorithms prohibitively expensive.

The equilibrium xenon distributions can also be calculated inside a Monte Carlo transport calculation, which is several times faster than the traditional wrapper algorithms [97]. Publication IV explores the use of this option in stabilizing burnup calculations. One algorithm for calculating the equilibrium inside a transport calculation is found in the Monte Carlo transport code `MC21` [97], while another one based on the same principle is available in the reactor physics code `Serpent`. These algorithms allow the equilibrium xenon distribution and the corresponding flux, as well as all related quantities, to be calculated inside a single transport simulation that only takes slightly longer than a normal one. While both algorithms were designed for other purposes [97, 98], they can also be used for removing xenon oscillations in burnup calculations. Simply by applying one of them at each neutronics solution, xenon concentrations and neutron flux can be forced to remain in equilibrium, thus preventing all xenon driven oscillations. Since only the neutronics is affected, this can be done with any coupling scheme.

5.4.1 The equilibrium algorithm of Serpent

The equilibrium calculation in `Serpent` is performed during a criticality source simulation by recalculating the concentrations of ^{135}I and ^{135}Xe after each source cycle using the flux and cross-sections tallied during that cycle. This

is done separately for each fissile material region. The new concentrations are then used during the next source cycle and so on. The result is a continuous iteration between neutronics and the equilibrium concentration of ^{135}I and ^{135}Xe, performed as the transport simulation is run. This means that the concentrations of these two nuclides change through all inactive and active cycles.

The concentrations of ^{135}I and ^{135}Xe are calculated by assuming that ^{135}Xe and its precursors are in a secular equilibrium with fission production, and that the neutron capture rates of the precursors of ^{135}Xe are insignificant compared to radioactive decay. With these approximations, the concentrations become:

$$x_{\mathrm{I}} = \frac{\gamma_{\mathrm{I}} \Sigma_{\mathrm{f}} \phi}{\lambda_{\mathrm{I}}} \tag{5.1}$$

and

$$x_{\mathrm{X}} = \frac{\gamma_{\mathrm{X}} \Sigma_{\mathrm{f}} \phi}{\lambda_{\mathrm{X}} + \sigma_{\mathrm{X}} \phi} \,, \tag{5.2}$$

where x_{I} and x_{X} are the concentrations of ^{135}I and ^{135}Xe, respectively, γ_{I} and γ_{X} (which includes γ_{I}) their cumulative fission yields, λ_{I} and λ_{X} their decay constants, Σ_{f} is the macroscopic total fission cross-section of the material, σ_{X} the microscopic capture cross-section of ^{135}Xe and ϕ the total flux.

All results, including the cross-sections and flux used in depletion calculations, are tallied as before over all active cycles. The concentrations of ^{135}I and ^{135}Xe are collected by averaging over the iterated concentrations from all active cycles. The concentrations of all other nuclides, including the parents and daughters of ^{135}I and ^{135}Xe, still come from depletion calculations.

Updating the xenon concentrations at every cycle means that the cycle-wise concentrations have poor statistics through the entire calculation, but this does not directly concern the above algorithm as it does not rely on the cycle-wise estimate to converge. Instead, it assumes that the statistical errors will cancel out as all results are collected over multiple cycles where the xenon concentrations are continuously recalculated.

While the algorithm has produced good results, its correctness and possible improvements remain topics for a future study. There has been no theoretical analysis on its validity and the estimate of Eq. 5.2 for ^{135}Xe concentrations is known to be biased [97] in the same way as the entire depletion calculation (Section 2.4.2). Because the updates in the algorithm of Serpent use only a single source cycle worth of statistics, the bias might become an issue in some cases despite usually being insignificant.

5.4.2 The equilibrium algorithm of MC21

In the equilibrium xenon algorithm of MC21 [97] the inactive cycles of a normal criticality source transport simulation are divided into N_u xenon update cycles, which consist of N_b source cycles each. The relevant reaction rates are tallied during the transport simulation, and after N_b cycles these estimates are used to recalculate saturated ^{135}Xe and ^{135}I concentrations in each fissile material region in the same way as in the algorithm of Serpent.

During the first N_r update cycles, the reaction rate tallies are reset after xenon concentrations have been updated. This ensures that the effects of the initial poorly converged cycles are not carried on. During the remaining $N_\mathrm{u} - N_\mathrm{r}$ xenon update cycles the tallies are not reset, allowing results to accumulate over multiple xenon updates for higher statistical accuracy. After the last xenon update cycle, the xenon concentrations are fixed and the simulation proceeds to the active cycles, which are performed normally. N_u, N_r and N_b are input parameters. $N_\mathrm{u} \approx 10 - 20$ and $N_\mathrm{r} \approx 4 - 5$ are suggested [97].

This algorithm requires additional inactive cycles to collect sufficient statistics before fixing the xenon distribution, which makes it somewhat slower than the algorithm of Serpent. On the other hand, it does not use the theoretically unjustified variable xenon concentrations during active cycles and reduces any possible bias in estimate (5.2) due to larger sample sizes.

Both this and the previous algorithm can also be applied in the fixed source mode. This only requires dividing the source neutrons into 'batches', and using them in the place of the source cycles of the criticality source mode. In the algorithm of Serpent, the first such batch needs to be treated as inactive, while the algorithm of MC21 should discard as many batches as are required to converge the xenon distribution.

5.4.3 Effects of the equilibrium calculation

The equilibrium xenon method was tested in Publication IV and, unsurprisingly, removed all xenon driven oscillations. The equilibrium algorithm of Serpent caused roughly 10 % increase in the running time per neutron history, but the negative feedback created by constant xenon updates also significantly improved convergence, leading to reduced statistical variation in geometries with very high dominance ratios. The stabilizing effect is independent of the algorithm used for calculating the equilibrium, but slowdown and the effects on convergence may vary.

Even with the xenon and flux forced to equilibrium, strong oscillations were observed when using too long steps. Their mechanisms are similar to the numerical xenon oscillations, with negative reactivity associated with increased local burnup replacing that from the buildup of xenon. While equilibrium xenon does not provide complete stability, it allows calculations to be performed with reasonable step lengths. In the tests of Publication IV 30-day steps were stable while 60-day steps were not. These limits held for all the depletion algorithms in Serpent 2, except for LE, which started to oscillate with 30-day steps but not with 16-day steps. In addition to the choice of burnup algorithm, the step length limit at which oscillations start is likely to be affected by the geometry, material compositions and power density.

While enforcing equilibrium allows stable solutions to be obtained for simplified models, the results are only as accurate as the model they have been calculated for. Forcing equilibrium without modeling the feedback and control systems that would stabilize a real reactor means that their effects on the equilibrium distributions are ignored. Some doubt also remains about the equilibrium xenon algorithm of Serpent, especially in respect to the bias of the xenon concentration at low statistical accuracy. The algorithm of MC21 [97] is not subject to these doubts.

5.5 Stochastic Implicit Euler Method

Another method for stabilizing Monte Carlo burnup calculations, dubbed Stochastic Implicit Euler (SIE) has been proposed [91, 99]. The method is presented in Algorithm 6 using the same notation as in Section 4.1, and c to denote the number of iterations, which Ref. [91] suggests no value for. Another possible implementation of the method is also presented [91], but expected to be inferior to the one in Algorithm 6. This method is fundamentally a wrapper algorithm, but merges the equilibrium and burnup algorithms to ensure that while multiple transport solutions are required, they all contribute to the final estimates. It is too early to conclusively assess the potential of the method, but simple considerations allow a preliminary assessment of the differences between it and the equilibrium xenon treatment of Publication IV.

SIE has one major advantage: It is not xenon specific and should also prevent the depletion related oscillations that arise with the equilibrium xenon method, thus ensuring stability at all step lengths. The largest downside of SIE is that it uses constant backwards extrapolation for predicting the behavior of

Algorithm 6

for $i = 0, 1, ..., I-1$ **do** % Loop over steps
  $x_{i+1}^0 \leftarrow x_i$ % Initial guess
  for $j = 0, 1, ..., c-1$ **do** % — Iteration —
    $\phi_{i+1}^j \leftarrow \phi(x_{i+1}^j)$ % EOS neutronics
    $\bar{\phi}_{i+1}^i \leftarrow \frac{1}{j+1} \sum_{k=0}^{j} \phi_{i+1}^k$ % Average over iterations
    $x_{i+1}^{j+1} \leftarrow e^{A(\bar{\phi})(T_{i+1}-T_i)} x_{\mathrm{i}}$ % Recalculated EOS composition
  end for
  $x_{i+1} \leftarrow x_{i+1}^c$ % Final results for step $i+1$
  $\phi_{i+1} \leftarrow \bar{\phi}_{\mathrm{i+1}}^{\mathrm{c}}$
end for

the cross-sections and flux. This might be no more accurate than the constant forward extrapolation of the explicit Euler method (Algorithm 1). The method might thus require relatively short steps to produce accurate results, which may be a major issue; if short steps are required for accuracy, retaining stability with long steps is useless. Ref. [91] suggests that a higher order version of the method might be possible, which would alleviate this problem. However, while it is easy enough to construct higher order variants, it is not trivially clear whether they would work. In fact, the simplest linear variant has already been found unstable [99].

Another downside is that even when running the same number of active neutron histories, spreading them over multiple transport calculations with different materials means that the overhead associated with initializing the transport simulation, converging source distribution, storing results, depleting materials and calculating averages is multiplied. Depending on how everything is implemented and what is being calculated, this might or might not be significant. The fact that the initial unconverged iterations are used for in the final results might also make the results worse than the number of contributing histories would suggest.

5.6 Other methods

Two more supposedly stable coupling schemes have been published [90]. The first of these, called extended predictor-corrector (EPC) method, divides each step into two parts. The first half is depleted using the average cross-sections from the previous step and flux renormalized to the BOS material compositions in a constant extrapolation akin to Algorithm 1 from Section 4.1. The second half of

the step is then covered by a step of Algorithm 2 from Section 4.1. The second new method, called variable weight Predictor-Corrector (VWPC), is identical to SIE, except that the flux from each iteration gets increasingly small weight in the average. The weights contain two free parameters with no guidelines for selecting them.

While EPC provides an advantage in stability [90], it is difficult to see how it could produce stable results in the general case as neither of the algorithms used for half-steps does according to Ref. [94]. VWPC might have an advantage over SIE with some correctly selected weights, but until such are presented, this remains theoretical.

6. Summary

The work presented in this thesis covers several aspects of the time-integration in Monte Carlo burnup calculations. Apart from the integrated equilibrium xenon algorithms, the new methods and results should also be applicable with deterministic neutronics.

The accuracies and performances of depletion algorithms for full systems of nuclides were evaluated. The results are quite conclusive and, together with other studies, show rational approximation based matrix exponential methods to be the best choice for Monte Carlo burnup calculations. Matrix exponential Taylor series with instant decay and secular equilibrium approximations can be even faster but of limited accuracy. It might still be preferable in special applications where the speed of depletion is critical, but Monte Carlo burnup calculations are not one of them as neutronics dominate the running time.

New coupling schemes, which use data from the previous step to make higher order predictions, were presented. Since the previous step values are readily available, no additional calculations are required, and the stepwise running time is not affected. The coupling is further improved by dividing the steps into substeps, which are solved sequentially. Since each substep can use different coefficients for the Bateman equations, this allows piecewise constant, rather than constant, approximation of the predicted behavior. Using linear extrapolation on the predictor greatly improved results, and when used without a corrector, it even appears to be a viable alternative to the predictor-corrector methods. Increasing the corrector order had predominantly negative effect on the accuracy of the results, but a quadratic corrector might still be preferable due to its more predictable behavior in respect to step lengths. Together, the new methods should allow step lengths to be roughly doubled.

The previously observed spatial oscillations in Monte Carlo burnup calculations were shown to be caused by the true solution of the simulation model being oscillatory, not by instability of the methods. The solutions can be stabilized

by forcing a mutual equilibrium between the flux and xenon concentrations. This can be done efficiently by integrating the equilibrium calculation into the neutronics. However, preventing xenon driven oscillations does not provide absolute stability; if too long steps are used, oscillations driven by other nuclides may emerge.

6.1 Future prospects

Results from the comparison of depletion algorithms are quite conclusive. While the study did not include the contour based rational approximations, which were published after it, the differences between those and CRAM are well covered in Ref. [84]. No comparison of any generality has been performed between the methods for full and simplified nuclide models, but especially with the excellent performance of the rational approximation methods, there seems to be no reason to consider simplified models with Monte Carlo neutronics.

One thing that could be checked is the performance of the depletion algorithms for decay-only calculations, in particular the cooling of spent nuclear fuel. In such calculations the system of nuclides as a whole behaves quite differently from burnup calculations, so there might be differences in the performances of the methods.

Coupling schemes utilizing even higher order or non-polynomial interpolation and extrapolation might, and non-equidistant substeps should, allow for better performance, but the achievable gains are likely to be small. A more pressing need lies in additional testing of the new coupling schemes. Thorough testing that includes the three old predictor-corrector methods presented in Section 4.1 should also be performed. In addition, reverting linear extrapolation to a constant one on the second step should be reconsidered. Using linear extrapolation on the second step leads to better results with properly selected step lengths, but reverting it to a constant one would make the method more robust against traditional step selections, where the second step is made much longer than the first.

The spatial stability of Monte Carlo burnup calculations is still a new and largely unexplored topic, making it the one in need of most work. The equilibrium xenon approach could use far more comprehensive testing, especially with realistic geometries. The recently published stable coupling schemes should also be tested more thoroughly by someone. In addition to stability, these tests should consider the accuracy and performance of each method. Testing the dif-

ferent methods in tandem would allow the results to be compared. However, it might be best to wait for some of the issues and possible improvements identified by the respective authors of each method to be addressed first.

Finally, possibilities for improving the equilibrium algorithm of Serpent should be explored. In particular, the current estimate for updating the xenon concentrations is biased, and the possibility of replacing it with one based on the unbiased estimator of Dumonteil et al. [23, 24] should be explored. The algorithm might also be improved by collecting statistics for the xenon updates over several cycles.

Bibliography

[1] Bateman, H. *Solution of a system of differential equations occurring in the theory of radio-active transformations*. Math. Proc. Cambridge 15, 423–427. (1910)

[2] Rutherford, E. *Radio-Activity*. Second Edition. Cambridge University Press. (1905)

[3] Koning, A. et al. *The JEFF-3.1 Nuclear Data Library*. JEFF Report 21. ISBN 92-64-02314-3. (2006)

[4] Lux, I., Koblinger, L. *Monte Carlo Particle Transport Methods: Neutron and Photon Calculations*. CRC press. ISBN 0-8493-6074-9. (1991)

[5] Lewis, E.E., Miller, W.F. Jr. *Computational Methods of Neutron Transport*. John Wiley & Sons, Inc. ISBN 0-471-09245-2. (1984)

[6] Leppänen, J. *Development of a New Monte Carlo Reactor Physics Code*. VTT Publication 640. ISBN 978-951-38-7018-8 (soft back), 978-951-38-7019-5 (pdf). (2007)

[7] Santamarina, A. et al. *The JEFF-3.1.1 Nuclear Data Library*. JEFF Report 22. ISBN 978-92-64-99074-6. (2009)

[8] Chadwick, M.B. et al. *ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data*. Nucl. Data Sheets 112, 2887–2996. (2011)

[9] Shibata, K. et al. *JENDL-4.0: A New Library for Nuclear Science and Engineering*. J. Nucl. Sci. Technol. 48, 1–30. (2011)

[10] Stacey, W.M. *Nuclear reactor physics*. John Wiley & Sons, Inc. ISBN 0-471-39127-1. (2001)

[11] Takeda, T., Hirokawa, N. Noda, T. *Estimation of Error Propagation in Monte-Carlo Burnup Calculations*. J. Nucl. Sci. Technol. 36, 738–745. (1999)

[12] Tohjoh, M., Endo, T., Watanabe, M., Yamamoto, A. *Effect of error propagation of nuclide number densities on Monte Carlo burn-up calculations*. Ann. nucl. Energy 33, 1424–1436. (2006)

[13] Gracía-Herranz, N., Cabellos, O., Sanz, J., Juan, J., Kuijper, J.C. *Propagation of statistical and nuclear data uncertainties in Monte Carlo burn-up calculations*. Ann. Nucl. Energy 35, 714–730. (2008)

[14] Park, H.J., Shim, H.J., Kim, C.H. *Uncertainty Propagation in Monte Carlo Depletion Analysis*. Nucl. Sci. Eng. 167, 196–208. (2011)

[15] Xu, Z., Rhodes, J., Smith, K. *Statistical Implications in Monte Carlo Depletions*. PHYSOR 2010. Pittsburgh, Pennsylvania, USA, May 9–14, 2010. On CD-ROM, ISBN 978-0-89448-079-9. (2010)

[16] Aliberti, G. et al. *Nuclear Data Sensitivity, Uncertainty And Target Accuracy Assessment for Future Nuclear Systems*. Ann. Nucl. Energy 33, 700–733. (2006)

[17] OECD/NEA. *Uncertainty and Target Accuracy Assessment for Innovative Systems Using Recent Covariance Data Evaluations*. NEA/WPEC-26. ISBN 978-92-64-99053-1. (2008)

[18] Rochman, D., Koning, A.J., da Cruz, D.F. *Propagation of* $^{235,236,238}U$ *and* ^{239}Pu *nuclear data uncertainties for a typical PWR fuel element*. Nucl. Technol. 179, 323–338. (2012)

[19] Wieselquist, W. Zhu, T., Vasiliev, A., Ferroukhi, H. *PSI Methodologies for Nuclear Data Uncertainty Propagation with CASMO-5M and MCNPX: Results for OECD/NEA UAM Benchmark Phase I*. Sci. Technol. Nucl. Install. Article ID 549793, 15 p. (2013)

[20] Stover, T.E. Jr. *Quantification of Back-End Nuclear Fuel Cycle Metrics Uncertainties Due to Cross sections*. Master's Thesis, Idaho National Laboratory. INL/EXT-07-13592. (2007)

[21] Rochman, D., Koning, A.J., Dacruz, D.F., van der Marck, S.C. *Nuclear Data Uncertainty Propagation for a Sodium Fast Reactor*. J. Korean Phys. Soc. 59, 1191–1194. (2011)

[22] Aliberti, G., Palmiotti., G., Salvatores, M., Stenberg, C.G. *Impact of Nuclear Data Uncertainties on Transmutation of Actinides in Accelerator-Driven Assemblies*. Nucl. Sci. Eng. 146, 13–50. (2004)

[23] Dumonteil, E., Diop, M.C. *Unbiased Minimum Variance Estimator of a Matrix Exponential Function. Application to Boltzmann/Bateman Coupled Equations Solving*. M&C 2009. Saratoga Springs, New York, May 3–7, 2009. On CD-ROM, ISBN 978-0-89448-069-0. (2009)

[24] Dumonteil, E., Diop, M.C. *Biases and Statistical Errors in Monte Carlo Burnup Calculations: An Unbiased Stochastic Scheme to Solve Boltzmann/Bateman Coupled Equations*. Nucl. Sci. Eng. 167, 165–170. (2011)

[25] Brun, E., Dumonteil, E., Malavagi, F. *Systematic Uncertainty due to Statistics in Monte Carlo Burnup Codes: Application to a Simple Benchmark with TRIPOLI-4-D*. Prog. Nucl. Sci. Technol. 2, 579–885. (2011)

[26] Shim, H.J., Cho, J.Y., Song, J.S., Kim, C.H. *Generation of few group diffusion theory constants by Monte Carlo code*. Trans. Am. Nucl. Soc. 99, 343–345. (2008)

[27] Fridman, E., Leppänen, J. *On the use of the Serpent Monte Carlo code for few-group cross section generation*. Ann. Nucl. Energy 38, 1399–1405. (2011)

[28] Fridman, E., Leppänen, J. *Revised Methods for Few-Group Cross Sections Generation in the Serpent Monte Carlo Code*. PHYSOR 2012, Knoxville, Tennessee, USA, April 15–20, 2012. On CD-ROM, ISBN: 978-0-89448-085-9. (2012)

[29] Ćalić, D. *Neutron Multigroup Homogenized Cross Section Determination With the Monte Carlo Method*. NENE 2012. Ljubljana, Slovenia, September 5–7. 2012. (2012)

[30] Yamamoto, T. *Monte Carlo method with complex weights for neutron leakage-corrected calculations and anisotropic diffusion coefficient generations*. Ann. Nucl. Energy 50, 141–149. (2012)

[31] Yun, S., Cho, N.Z. *Monte Carlo Depletion Under Leakage-Corrected Critical Spectrum via Albedo Search*. Nucl. Eng. Technol. 42, 271–278. (2010)

[32] Park, H.J., Shim, H.J., Joo, H.G., Kim, C.H. *Generation of Few-Group Diffusion Theory Constants by Monte Carlo Code McCARD*. Nucl. Sci. Eng. 172, 66–77. (2012)

[33] Smith, K.S. *Assembly Homogenization Techniques for Light Water Reactor Analysis*. Prog. Nucl. Energy 17, 303–335. (1986)

[34] Palmtag, S.P. *Advanced Nodal Methods for MOX Fuel Analysis*. Doctor's Thesis, Massachusetts Institute of Technology. (1997)

[35] Woodcock, E.R. et al. *Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry*. In Proc. Conference on the Application of Computing Methods to Reactor Problems. ANL-7050. (1965)

[36] Tohjoh, M., Watanabe, M., Yamamoto, A. *Study of spatial discretization and temperature distribution approximation effects on BWR assembly calculations*. Ann. Nucl. Energy 33, 170–179. (2006)

[37] Rintala, A. *Effect of temperature discretization on the results of neutronics calculations*. Bachelor's thesis, Aalto University. (2012)

[38] Valtavirta, V. *Effect of temperature homogenization on neutronics calculations of Gd-doped fuel*. Special Assignment, Aalto University. (2011)

[39] Rochman, D., Sciolla, C.M. *"Total Monte Carlo" Uncertainty propagation applied to the Phase I-1 burnup calculation*. NRG Report 113696. (2012)

[40] Sanchez, V., Al-Hamry, A. *Development of a Coupling Scheme Between MCNP and COBRA-TF for the Prediction of the Pin Power of a PWR Fuel Assembly*. M&C 2009. Saratoga Springs, New York, May 3–7, 2009. On CD-ROM, ISBN 978-0-89448-069-0. (2009)

[41] Ivanov, A., Sanchez V., Imke, U. *Development of a Coupling Scheme Between MCNP5 and SUBCHANFLOW for the Pin- and Fuel Assembly-Wise Simulation of LWR and Innovative Reactors*. M&C 2011. Rio de Janeiro, RJ, Brazil, May 8–12, 2011. On CD-ROM, ISBN 978-85-63688-00-2. (2011)

[42] Vazquez, M., Tsige-Tamirat, H., Ammirabile, L., Martin-Fuertes F. *Coupled neutronics thermal-hydraulics analysis using Monte Carlo and subchannel codes*. Nucl. Eng. Des. 250, 403–411. (2012)

[43] Guo, Z., Zhou, J., Zhang, D., Chaudri, K.S., Tian, W., Su, G., Qiu, S. *Coupled neutronics/thermal-hydraulics for analysis of molten salt reactor*. Nucl. Eng. Des. 258, 144–156. (2013)

[44] Espel, F.P., Avramova, M.N., Ivanov, K.N., Misu, S. *New developments of the MCNP/CTF/NEM/NJOY code system – Monte Carlo based coupled code for high accuracy modeling*. Ann. Nucl. Energy 51, 18–26. (2013)

[45] Reiss, T., Fehér, S., Czifrus, Sz. *Coupled neutronics and thermohydraulics calculations with burn-up for HPLWRs*. Prog. Nucl. Energy 50, 52–61. (2008)

[46] Noh, J.M., Kim, K-S., Kim, Y., Lee H.C. *Development of a computer code system for the analysis of prism and pebble type VHTR cores*. Ann. Nucl. Energy 35, 1919–1928. (2008)

[47] Rohde, U., Baier, S., Duerigen, S., Fridman, E., Kliem, S., Merk B. *Development and verification of the coupled 3D neutron kinetics/thermal-hydraulics code DYN3D-HTR for the simulation of transients in block-type HTGR*. Nucl. Eng. Des. 251, 412–422. (2012)

[48] Hoogenboom, J.E., Khotylev, V.A., Tholammakkil, J.M. *Generation of Multi-Group Cross Sections and Scattering Matrices With The Monte Carlo Code MCNP5*. M&C + SNA 2007. Monterey, California, April 15–19, 2007. On CD-ROM, ISBN 0-89448-059-6. (2007)

[49] Fridman, E., Shwageraus, E. *Modeling of SFR cores with Serpent–DYN3D codes sequence*. Ann. Nucl. Energy 53, 354–363. (2013)

[50] Rachamin, E., Wemple, C., Fridman, E. *Neutronic analysis of SFR core with HELIOS-2, Serpent, and DYN3D codes*. Ann. Nucl. Energy 55, 194–204. (2013)

[51] Tohjoh, M., Watanabe, M., Yamamoto, A. *Application of continuous-energy Monte Carlo code as a cross-section generator of BWR core calculations*. Ann. Nucl. Energy 32, 857–875. (2005)

[52] Matsunaka, M., Shido, S., Kondo, K., Miyamaru, H., Murata, I. *Burnup calculation of fusion–fission hybrid energy system with thorium cycle*. Fusion Eng. Des. 82, 2779–2785. (2007)

[53] Kramer, K.J., Latkowski, J.F., Abbott, R.P., Boyd, J.K., Powers, J.J., Seifried, J.E. *Neutron Transport and Nuclear Burnup Analysis for the Laser Inertial Confinement Fusion-Fission Energy (LIFE) Engine*. Fusion Sci. Technol. 56, 625–631. (2009)

[54] Mathieu, L. et al. *Possible Configurations for the Thorium Molten Salt Reactor and Advantages of the Fast Nonmoderated Version*. Nucl. Sci. Eng. 161, 78–89. (2009)

[55] Bomboni, E., Cerullo, N., Lomonaco, G. *Simplified models for pebble-bed HTR core burn-up calculations with Monteburns2.0*. Ann. Nucl. Energy 40, 72–83. (2012)

[56] Ochoa, R., Vázquez, M., Álvarez-Velarde, F., Martín-Fuertes, F., Garćia-Herranz, N., Cuervo, D. *A comparative study of Monte Carlo-coupled depletion codes applied to a Sodium Fast Reactor design loaded with minor actinides*. Ann. Nucl. Energy 57, 32–40. (2013)

[57] Stankovskiy, A., Van den Eynde, G. *Advanced Method for Calculations of Core Burn-Up, Activation of Structural Materials, and Spallation Products Accumulation in Accelerator-Driven Systems*. Sci. Technol. Nucl. Install. Article ID 545103, 12 p. (2012)

[58] X-5 Monte Carlo Team. *MCNP – A General Monte Carlo N-Particle Transport Code, Version 5*. LA-UR-03-1987. (2003, Revised 2/1/2008)

[59] Posfon, D.I., Trellue, H.R. *User's Manual, Versio 2.0 for Monteburns version 1.0*. LA-UR-99-4999. (1999)

[60] Moore, R.L., Schnitzler, B.G., Wemple, C.A., Babcock R.S., Wessol, D.E. *MOCUP: MCNP-ORIGEN2 Coupled Utility Program*. INEL-95/0523. (1995)

[61] Kotlyar, D., Shaposhnik, Y., Fridman, E., Shwageraus, E. *Coupled neutronic thermo-hydraulic analysis of full PWR core with Monte-Carlo based BGCore system*. Nucl. Eng. Des. 241, 3777–3786. (2011)

[62] Meplan, O., Nuttin, A., Laulan, O., David, S., Michel-Sendis, F., Wilson, J. *MURE: MCNP Utility For Reactor Evaluation - Description of the methods, first applications and results*. Proceedings of the ENC2005 (European Nuclear Conference). Versailles, France, December 14–18, 2005. (2005)

[63] El Bakkari, B. et al. *Development of an MCNP-tally based burnup code and validation through PWR benchmark exercises*. Ann. Nucl. Energy 36, 626–633. (2009)

[64] Xu, Z., Hejzlar, P. *MCODE, Version 2.2: An MCNP-ORIGEN Depletion Program*. MIT-NFC-TR-104. (2008).

[65] Cetnar, J., Wallenius, J., Gudowski, W. *MCB: a continuous energy Monte Carlo Burnup simulation code*. In Proceedings of the Fifth Information Exchange Meeting on Actinide and Fission Product Partitioning and Transmutation. Mol, Belgium, 25-27 November, 1998. (1998)

[66] Okumura, K., Mori., Takamasa, Nakagawa, M., Kaneko, K. *Validation of a Continuous-Energy Monte Carlo Burn-up Code MVP-BURN and Its Application to Analysis of Post Irradiation Experiment*. J. Nucl. Sci. Technol. 37, 128–138. (2000)

[67] DeHart, M.D. *TRITON: A Two-Dimensional Transport and Depletion Module for Characterization of Spent Nuclear Fuel*. in SCALE: A Modular Code System for Performing Standardized Computer Analyses for Licensing Evaluation. ORNL/TM-2005/39, Version 6 Vol I. (2009)

[68] Pelowitz, D.B. ed. *MCNPX 2.6.0 Manual*. LA-CP-07-1473. (2008)

[69] Shim, H.J., Han, B.S., Jung, J.S., Park, H.J., Kim, C.H. *McCARD: Monte Carlo Code for Advanced Reactor Design and Analysis*. Nucl. Eng. Technol. 44, 161–175. (2012)

[70] Serpent website. Jun 10, 2013. `http://montecarlo.vtt.fi`

[71] OECD/NEA Data Bank Computer Program Services. *SERPENT 1.1.7*. NEA-1840. Jun 10, 2013.
`WWW.oecd-nea.org/tools/abstract/detail/nea-1840`

[72] Leppänen, J. *Performance of Woodcock delta-tracking in lattice physics applications using the Serpent Monte Carlo reactor physics burnup calculations code*. Ann. Nucl. Energy 37, 715–722. (2010)

[73] Leppänen, J. *Two Practical Methods for Unionized Energy Grid Construction in Continuous-Energy Monte Carlo Neutron Transport Calculation*. Ann. Nucl. Energy 36, 878–885. (2009)

[74] Leppänen, J. *Multi-physics Coupling Scheme in the Serpent 2 Monte Carlo Code*. Trans. Am. Nucl. Soc. 107, 1165-1168. (2012)

[75] Leppänen, J., Isotalo, A. *Burnup Calculation Methodology in the Serpent 2 Monte Carlo Code*. PHYSOR 2012, Knoxville, Tennessee, USA, April 15–20, 2012. On CD-ROM, ISBN 978-0-89448-085-9. (2012)

[76] Pusa, M., Leppänen, J. *Computing the Matrix Exponential in Burnup Calculations*. Nucl. Sci. Eng. 164, 140–150. (2010)

[77] Cetnar, J. *General Solution of Bateman Equations for Nuclear Transmutations*. Ann. Nucl. Energy 33, 640–645. (2006)

[78] Mole, C., Van Loan, C. *Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later*. SIAM Rev. 45, 3–49. (2003)

[79] Croff, A.G. *ORIGEN2: A Versatile Computer Code for Calculating the Nuclide Compositions and Characteristics of Nuclear Materials*. Nucl. Technol. 62, 335–352. (1983)

[80] Croff, A.G. *A User's Manual for the ORIGEN2 Computer Code*. ORNL/TM-7175. (1980)

[81] Croff, A.G. *ORIGEN2 - A Revised and Updated Version of the Oak Ridge Isotope Generation and Depletion Code*. ORNL-5621. (1980)

[82] Hermann, O.W., Westfall, R.M. *ORIGEN-S: SCALE System Module to Calculate Fuel Depletion, Actinide Transmutation, Fission Product Buildup and Decay, and Associated Radiation Source Terms*. In SCALE: A Modular Code System for Performing Standardized Computer Analyses for Licensing Evaluation. Rev. 6. ORNL/TM-2005/39. (2009)

[83] Schmelzer, T. *Carathéodory-Fejér approximation*. Matlab Central, File ID: 22055. (2008)

[84] Pusa, M. *Rational Approximations to the Matrix Exponential in Burnup Calculations*. Nucl. Sci. Eng. 169, 155–167. (2011)

[85] Rose, D.J., Tarjan, R.E. *Algorithmic Aspects of Vertex Elimination on Directed Graphs*. SIAM J. Comput. 5, 266–283. (1976)

[86] Pusa, M., Leppänen, J. *Efficient Implementation of CRAM for Solving Burnup Equations*. PHYSOR 2012. Knoxville, Tennessee, USA, April 15–20, 2012. On CD-ROM, ISBN 978-0-89448-085-9. (2012)

[87] Yamamoto, A., Tatsumi, M., Sugimura, N. *Numerical Solution of Stiff Burnup Equation with Short Half Lived Nuclides by the Krylov Subspace Method*. J. Nucl. Sci. Technol. 44, 147–154. (2007)

[88] Hairer, E., Wanner, G. *Solving Ordinary Differential Equations. Stiff and Differential-Algebraic Problems*. 2nd edition. Vol 14 of Springer Series in Computational Mathematics. ISBN 978-3-540-60452-5. (1996)

[89] Yamamoto, A., Tatsumi, M., Sugimura, N. *Projected Predictor-Corrector Method for Lattice Physics Burnup Calculations*. Nucl. Sci. Eng. 163, 144–151. (2009)

[90] Kotlyar, D., Shwageraus, E. *On the use of predictor-corrector methods for coupled Monte Carlo burnup codes*. Ann. Nucl. Energy 58, 228–237. (2013)

[91] Dufek, J., Kotlyar, D., Shwageraus, E. *The Stochastic Implicit Euler Method – A Stable Coupling Scheme for Monte Carlo Burnup Calculations*. Ann. Nucl. Energy 60, 295–300. (2013)

[92] Carpenter, D.C. *A comparion of constant power depletion algorithms*. M&C 2009. Saratoga Springs, New York, May 3–7, 2009. On CD-ROM, ISBN 978-0-89448-069-0. (2009)

[93] Fensin, M.L. et al. *Incorporation of a Predictor-Corrector Methodology and 1-Group Reaction Rate Reporting Scheme for the MCNPX Depletion Capability*. Trans. Am. Nucl. Soc. 95, 317–319. (2006)

[94] Dufek, J., Hoogenboom, J.E. *Numerical Stability of Existing Monte Carlo Burnup Codes in Cycle Calculations of Critical Reactors*. Nucl. Sci. Eng. 162, 307–311. (2009)

[95] Dufek, J., Kotlyar, D., Shwageraus, E., Leppänen, J. *Numerical Stability of the Predictor-Corrector Method in Monte Carlo Burnup Calculations of Critical Reactors*. Ann. Nucl. Energy 56, 34–38. (2013)

[96] Dufek, J., Gudowski, W. *Stochastic Approximation for Monte Carlo Calculation of Steady-State Conditions in Thermal Reactors*. Nucl. Sci. Eng. 152, 274–283. (2006)

[97] Griesheimer D.P. *In-Line Xenon Convergence Algorithm for Monte Carlo Reactor Calculations*. PHYSOR 2010. Pittsburgh, Pennsylvania, USA, May 9-14, 2010. On CD-ROM, ISBN 978-0-89448-079-9. (2010)

[98] Leppänen, J. Personal communication. January 10, 2013.

[99] Dufek. J. *Progress in the HPMC project, WP2. XII Meeting on Reactor Physics Calculations in the Nordic Countries*. Halden, Norway, May 19–20, 2005. (2005)

Accurate computational methods are essential for the continued utilization and further development of nuclear energy. One central part of the computational system is formed by burnup calculations, which model long term changes in the composition of nuclear fuel together with the associated changes in the neutronics properties of the system. In Monte Carlo burnup calculations, the neutronics part is handled by the Monte Carlo method. Such calculations are accurate, flexible and problem independent, but their use is limited by their high computational cost.

The work presented in this thesis covers several aspects of Monte Carlo burnup calculations and presents new methods that improve the performance and stability of the calculations. While the thesis and the underlying research focus on Monte Carlo burnup calculations, many of the results and new methods are also applicable with deterministic neutronics.

Aalto-DD 172/2013

ISBN 978-952-60-5397-4
ISBN 978-952-60-5398-1 (pdf)
ISSN-L 1799-4934
ISSN 1799-4934
ISSN 1799-4942 (pdf)

Aalto University
School of Science
Department of Applied Physics
www.aalto.fi

BUSINESS + ECONOMY

ART + DESIGN + ARCHITECTURE

SCIENCE + TECHNOLOGY

CROSSOVER

DOCTORAL DISSERTATIONS