Annals of Nuclear Energy 60 (2013) 78–85

Contents lists available at SciVerse ScienceDirect

Annals of Nuclear Energy

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

Preventing xenon oscillations in Monte Carlo burnup calculations by enforcing equilibrium xenon distribution

A.E. Isotalo [a,*], J. Leppänen [b], J. Dufek [c]

[a] Aalto University, Department of Applied Physics, P.O.Box 14100, FI-00076 AALTO, Finland
[b] VTT Technical Research Centre of Finland, P.O.Box 1000, FI-02044 VTT, Finland
[c] Royal Institute of Technology, Nuclear Reactor Technology, AlbaNova University Center, SE-10691 Stockholm, Sweden

ARTICLE INFO

Article history:
Received 7 February 2013
Received in revised form 26 April 2013
Accepted 29 April 2013
Available online 24 May 2013

Keywords:
Xenon oscillations
Burnup calculations
Monte Carlo
Spatial stability

ABSTRACT

Existing Monte Carlo burnup codes suffer from instabilities caused by spatial xenon oscillations. These oscillations can be prevented by forcing equilibrium between the neutron flux and saturated xenon distribution. The equilibrium calculation can be integrated to Monte Carlo neutronics, which provides a simple and lightweight solution that can be used with any of the existing burnup calculation algorithms. The stabilizing effect of this approach, as well as its limitations are demonstrated using the reactor physics code Serpent.

1. Introduction

Monte Carlo burnup calculations have typically focused on two dimensional pin cell geometries, assembly segments and other geometries with relatively small dimensions. As computers and algorithms develop, calculations involving research reactors, full 3D assemblies and even simplified models of power reactors are becoming increasingly common. While most modeled geometries have been, and still are, too small or too crudely discretized for spatial oscillations to occur, applications are increasingly approaching the point where this in no longer the case.

Several widely used burnup calculation algorithms have been found to be unstable, at least in long symmetric pin cell geometries (Dufek and Hoogenboom, 2009; Dufek et al., 2013), which is sufficient to show that they cannot handle the general case. These oscillations are driven by xenon, although due to time discretization the mechanisms differ from physical xenon oscillations. Since all existing methods seem to be affected, this effectively prevents expanding Monte Carlo burnup calculations to large and detailed geometries.

Xenon oscillations can also occur in real reactors, or could, if they were not prevented by active control. Due to various approximations oscillations in numerical calculations can be much worse than they would in real reactors, but despite this, explicitly modeling the control system should help. Such solution would, however, be extremely laborious, if at all feasible in the context of Monte Carlo neutronics, and thus a simpler alternative is required.

In deterministic codes, spatial oscillations involving various quantities are dealt with by forcing equilibrium at each time step. This is done via wrapper algorithms that use multiple neutronics solutions to find the equilibrium distributions and the corresponding flux, which is then used for depletion. This approach has also been used in Monte Carlo burnup calculations (Dufek and Gudowski, 2006). However, with Monte Carlo neutronics it is also possible to efficiently calculate equilibrium xenon distributions inside the criticality source simulation (Griesheimer, 2010).

In this paper we suggest utilizing such inline equilibrium xenon calculations for stabilizing Monte Carlo burnup calculations. This provides a lightweight approach that can be used with any burnup calculation algorithm. The inherent instability of computational models used in Monte Carlo burnup calculations and the stabilizing effect of the equilibrium xenon treatment are demonstrated.

2. Theory

2.1. 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 ($T_{1/2} \approx 9.2$ h) and ^{135m}Xe ($T_{1/2} \approx 15$ min) from thermal fissions is only around 0.2%, while its 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

* Corresponding author.
E-mail address: aarno.isotalo@aalto.fi (A.E. Isotalo).

0306-4549/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.anucene.2013.04.031

half-life of 6.6 h. Because of this, changes in the flux affect xenon production rate with a delay, whereas removal rate, which is dominated by absorption, changes 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 in reactivity 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 tilting the flux the opposite way and the cycle repeats.

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 corresponding to the used flux at each step, making the physical xenon oscillation mechanism impossible. Instead, if the flux is tilted, the areas with 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 other way, leading to an unphysical oscillation.

2.2. Equilibrium xenon calculation

All xenon driven oscillations are prevented if the xenon concentrations and neutron flux are forced to remain in equilibrium. Griesheimer (2010) has presented an algorithm that allows the equilibrium to be calculated inside a Monte Carlo criticality source simulation, providing a massive reduction in running time when compared to traditional wrapper algorithms. Another integrated equilibrium calculation algorithm based on the same principle can be found in the reactor physics code Serpent.[1] While both algorithms were designed for other purposes, they can also be used for removing oscillations in burnup calculations simply by applying them to all neutronics solutions. Since only the neutronics is affected, this can be done with any burnup calculation algorithm.

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 the actinides, and that the neutron capture rates of the precursors of ^{135}Xe are insignificant compared to radioactive decay. With these approximations, the concentrations become:

$$n_{\mathrm{I}} = \frac{\gamma_{\mathrm{I}} \Sigma_{\mathrm{f}} \Phi}{\lambda_{\mathrm{I}}} \tag{1}$$

and

$$n_{\mathrm{X}} = \frac{\gamma_{\mathrm{X}} \Sigma_{\mathrm{f}} \Phi}{\lambda_{\mathrm{X}} + \sigma_{\mathrm{X}} \Phi}, \tag{2}$$

where n_{I} and n_{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 cycles. The concentrations of all other nuclides, including the daughters of ^{135}I and ^{135}Xe still come from depletion calculations.

The fission yields used in the equations are fission rate weighted averages of the values for each actinide. The data is typically provided for three incident neutron energies: 0.0253 eV, 400 keV and 14.0 MeV. Even though the energy dependence is taken into account in Serpent burnup calculations, the equilibrium xenon model always uses the data corresponding to the lowest energy.

While the algorithm has produced good results, its correctness and possible improvements remains a topic of future study: There has been no theoretical analysis on its validity, and the estimate of Eq. (2) for ^{135}Xe concentrations is known to be biased (Griesheimer, 2010). Because the updates in the algorithm of Serpent uses only one source cycle worth of statistics, the bias might become an issue in some cases despite being insignificant in the algorithm of Griesheimer (2010).

3. Numerical test calculations

The base case for all tests is a single PWR pin cell identical to the one used by Dufek and Hoogenboom (2009). Fuel pin diameter is 0.82 cm, cladding outer diameter 0.95 cm and there is no gas gap. Lattice pitch is 1.26 cm. The cell is 4 m long and divided into eight 50 cm long axial segments. The segments are numbered 1–8 starting from one end. Reflective boundary conditions are used at all boundaries, including the vertical direction. The cladding is pure zirconium at 600 K and the coolant light water with density of 0.7 g/cm^3 at 600 K. Fuel density is 10 g/cm^3, temperature 900 K and their average enrichment 3.1 wt.%. To break the symmetry, this base case is varied by increasing enrichment in segments 1–4, and lowering it in segments 5–8. The enrichments are selected so that the average remains at 3.1 wt.%. For example, with 0.4 pp (percentage points) difference the enrichments are 3.3 wt.% and 2.9 wt.%. Mean linear power is kept constant at 16 kW/m.

Burnup calculations consisting of 480 steps of 15 min, 40 steps of 3 h, 20 steps of 6 h, 10 steps of 12 h or 10 steps of 1, 2, 4, 8, 16, 30, 60 or 120 d were performed with enrichment differences of 0, 0.1, 0.4, 0.8, and 1.6 pp using the default CE/LI predictor–corrector burnup algorithm (Isotalo and Aarnio, 2011a) of Serpent. The calculations with 15 min steps used version 2.1.9 of Serpent with 1000 inactive and 5000 active cycles of 1000 neutrons, while all other calculations used version 2.1.10 with 1000 inactive and 5000 active cycles of 5000 neutrons.[2] All calculations used the same JEFF 3.1.1 based nuclear data libraries, and were repeated five times with different random number sequences to get an idea of the magnitude of statistical variation.

Additional test calculations, and their results, are described in Sections 3.3 and 3.4. Section 3.5 describes the differences, or lack of, in results obtained with other burnup algorithms.

3.1. Short steps without equilibrium xenon

Step lengths under 1 d are well below those typically used in burnup calculations. These calculations were done to demonstrate that it is impossible, not just unpractical, to avoid the oscillations simply by reducing step lengths and thus to show that a stabilizing scheme is really required. Since reducing step lengths should only

[1] For a complete and up-to-date description of the Serpent code, see http://montecarlo.vtt.fi.

[2] Equilibrium xenon calculation was added to Serpent 2 in version 2.1.10. There are no other changes that affect the results.

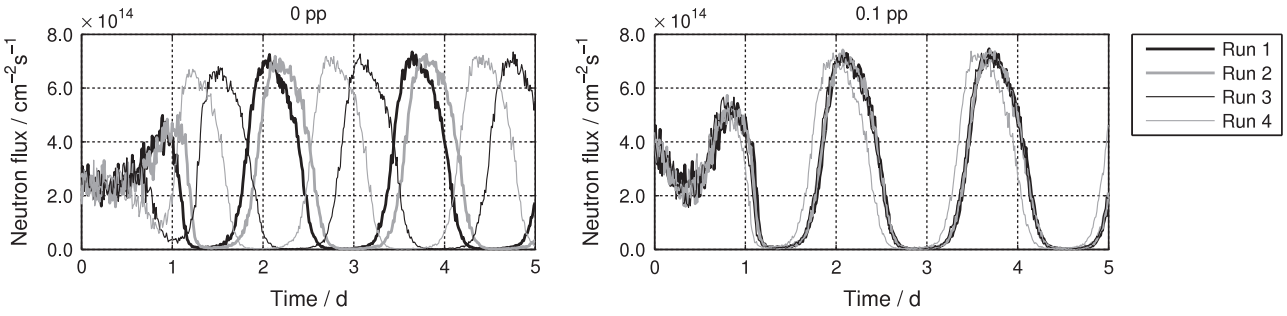

Fig. 1. Flux in segment 1 with 0 and 0.1 pp enrichment differences and 15 min steps without equilibrium xenon.

improve accuracy, the results from these calculations should also be the most accurate.

Fig. 1 shows the flux obtained for segment 1 with 0 and 0.1 pp enrichment differences and 15 min steps without equilibrium xenon in four runs with different random number sequences. In both cases, segments 2–4 show similar oscillations with decreasing amplitude towards the center of the rod, while segments 8–5 have identical results as segments 1–4, but with opposite phases. The only notable differences between these enrichment differences and results with different random number sequences are the phases of the oscillations.

Segment 1 fluxes for enrichment differences of 0.1, 0.4, 0.8 and 1.6 pp without equilibrium xenon are shown in Fig. 2, which also shows the xenon concentrations. With 0.8 and 1.6 pp differences the oscillations are dramatically smaller than with smaller differences, although with 0.8 pp difference the amplitude is growing steadily, which might lead to similar oscillation as with the smaller enrichment differences. Spatial distributions with 0.4 and 0.8 pp differences are shown in Fig. 3. With 0.4 pp and lower enrichment differences the flux oscillates between the ends of the rod, whereas with 0.8 pp and 1.6 pp (not shown) differences the oscillations are much smaller and happen between the higher enriched end and the middle of the rod.

These oscillations follow the physical mechanism, which has a timescale of hours. Because of this, very short steps are required to reproduce this behavior. Results obtained with 3 h steps are in a fairly good agreement with those obtained with 15 min steps (Fig. 1), but with 6 and 12 h steps the crude time discretization already causes large distortions, although the flux still oscillates in similar manner as with shorter steps. 24 h steps are simply too long to support the physical oscillations, and there is a transition to numerical mechanisms.

3.2. Effects of equilibrium xenon

With multi-day steps more typical in burnup calculations the physical xenon oscillation mechanism is replaced with a numerical one as explained in Section 2.1. However, with enrichment difference of 0.8 pp and 1–4 d steps, or enrichment difference of 1.6 pp and 1–8 d steps, the calculations appear to be stable even without equilibrium xenon and the results obtained with and without it are very similar. The only significant exceptions are the initial distributions in each case, where large differences arise due to equilibrium xenon being applied despite zero burnup.

With longer steps or smaller enrichment difference the calculations without equilibrium xenon start to oscillate, but these oscillations happen between subsequent neutronics solutions, i.e., the predictor and corrector, rather than between subsequent steps. Thus the flux is tilted one way during each predictor step, and the other way during each corrector step. As an example, Fig. 4 shows the predictor and corrector fluxes with 0.4 pp enrichment difference after two 16 d steps. These distributions remain very similar through the calculation and with different random number sequences. Oscillations again happen in the same phase in all repeats of the asymmetric cases.

Table 1 lists the average difference between local predictor and corrector fluxes with various enrichment differences and step lengths with and without equilibrium xenon. Varying differences between the predictor and corrector values are a normal part of predictor–corrector methods. While differences as large as the mean flux are a clear sign that the calculation has gone haywire, defining a threshold value for how large a difference indicates instability is difficult, especially as the transition from stable to oscillating behavior is probably continuous.

The oscillation between predictor and corrector leads to results[3] that look relatively stable, but are not even self-consistent. Fig. 5 shows the flux and ^{235}U concentrations calculated with 0.4 pp enrichment difference and 16 d steps. An example of the predictor and corrector fluxes in this case was given in Fig. 4. These results claim that uranium is depleted in segment 8 just as fast as in segment 1 despite segment 1 having over 20 times higher flux and higher enrichment. The discrepancy between segments 5 and 8 is also very clear: both have the same enrichment and segment 5 has 9 times higher flux, yet more uranium is depleted in segment 8. Simply outputting the average flux that is used for depletion instead of the predictor flux would not solve this problem as neither the average flux, nor the compositions calculated with it, are correct. This is apparent, for example, from the fact that more uranium should be depleted in segments 1 and 5 than in segment 8, yet this is not the case.

With step lengths below 60 d, equilibrium xenon removes these oscillations and the resulting discrepancies. As an example, Fig. 6 shows the results for ^{235}U concentration and flux with 0.4 pp enrichment difference and 16 d steps when equilibrium xenon is used. The same case, without equilibrium xenon, was presented in Fig. 5.

However, with 60 and 120 d, oscillation between predictor and corrector happens even with equilibrium xenon as seen from the large difference between their fluxes in Table 1. With 60 d steps the differences are limited to the early part of the calculation and disappear by the fifth step. While this looks like a start-up effect, changing the step lengths of the first four steps to 4, 8, 16 and 30 d only moved the start of the instability back to the first 60 d step. With 120 d steps the difference remains large through the entire calculation.

[3] The results collected from burnup calculations are typically the corrected atomic densities and the flux calculated with these compositions during the following predictor. In the CE/LI algorithm used here, the corrected material densities are calculated using the mean of the predictor and corrector fluxes.

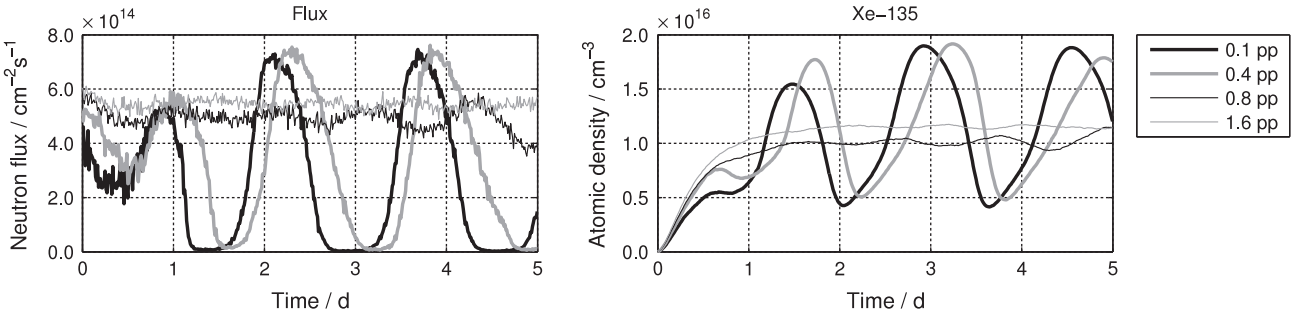

Fig. 2. Flux and ^{135}Xe concentration in segment 1 with various enrichment differences and 15 min steps without equilibrium xenon.

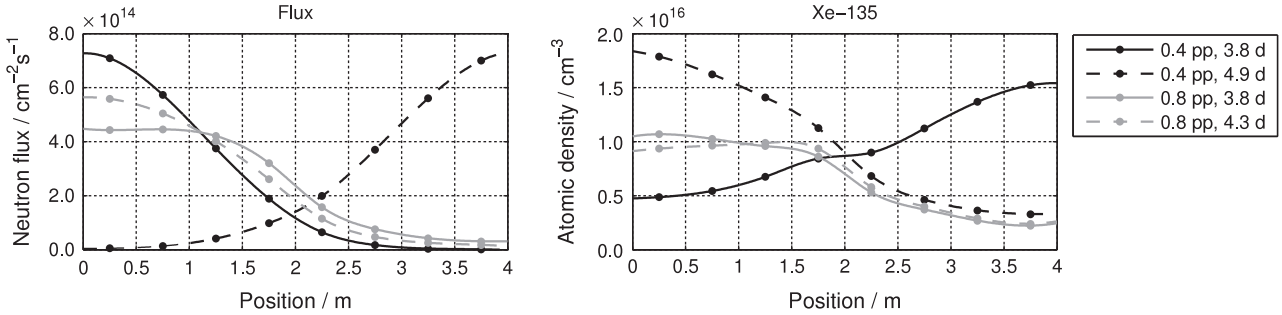

Fig. 3. Flux and ^{135}Xe distributions with 0.4 and 0.8 pp enrichment differences, 15 min steps and no equilibrium xenon at times corresponding to the extremes of the oscillations. The data points are the tallied averages for each segment.

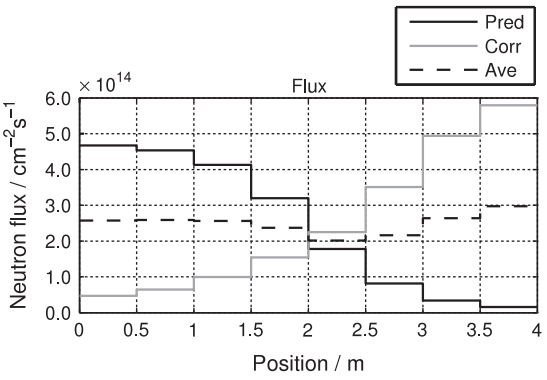

Fig. 4. Predictor and corrector fluxes with 0.4% enrichment after two 16 d step, as well as the average, which is used for depletion on the corrector.

Table 1
Average difference between local predictor and corrector fluxes as a fraction of the global mean flux with various enrichment differences and step lengths.

	1 d	2 d	4 d	8 d	16 d	30 d	60 d	120 d
Without equilibrium xenon								
0.0 pp	0.124	0.493	0.746	0.938	1.194	1.319	1.429	1.529
0.1 pp	0.180	1.000	1.112	1.156	1.239	1.335	1.431	1.529
0.4 pp	0.117	0.812	1.142	1.198	1.237	1.319	1.438	1.540
0.8 pp	0.022	0.032	0.062	0.492	0.990	1.219	1.390	1.524
1.6 pp	0.013	0.014	0.017	0.024	0.416	0.990	1.300	1.419
With equilibrium xenon								
0.0 pp	0.022	0.022	0.021	0.022	0.025	0.025	0.047	0.515
0.1 pp	0.021	0.021	0.021	0.020	0.023	0.029	0.065	0.827
0.4 pp	0.019	0.017	0.019	0.021	0.023	0.041	0.132	0.837
0.8 pp	0.013	0.014	0.014	0.017	0.023	0.048	0.158	0.825
1.6 pp	0.010	0.010	0.010	0.012	0.017	0.043	0.130	0.828

3.3. Accuracy of the equilibrium algorithm of Serpent

The symmetric case with no enrichment difference is special since the symmetry presents an alternate way to stabilize the calculation. This was utilized by performing reference calculations where the rod had only a single 4 m long depletion zone making it essentially two-dimensional, and thus stable without the use of equilibrium xenon. These calculations were otherwise identical to those presented in the previous section.[4] The mean flux and ^{135}Xe concentrations from these 2D calculations were used as reference values and compared to the results obtained with eight segments and equilibrium xenon. To test the role of the bias in the xenon concentration estimate (Eq. 2), part of the equilibrium xenon calculations were also repeated using different numbers of neutrons per source cycle while keeping the total number of neutrons per solution unchanged.

Table 2 shows the mean, μ, and standard deviation, σ, of relative differences between these reference values and corresponding equilibrium xenon calculations with eight depletion zones and 2 d and 30 d step lengths. The values are calculated over all segments, repeats and steps starting after the third one as:

$$\mu = \frac{1}{N_s \times (N_t - 2) \times N_r} \sum_{s=1}^{N_s} \sum_{t=3}^{N_t} \sum_{r=1}^{N_r} \frac{\phi_{s,t,r} - \bar{\phi}_t}{\bar{\phi}_t} \tag{3}$$

and

[4] There was a bug in the library routines of Serpent causing relative error of 4×10^{-3} in xenon concentration when equilibrium xenon is not used, and comparable errors in some other fission products regardless of the equilibrium calculation. The results presented in this section were recalculated after the bug was discovered and fixed, but those in other sections were not as the errors are too small to affect stability.

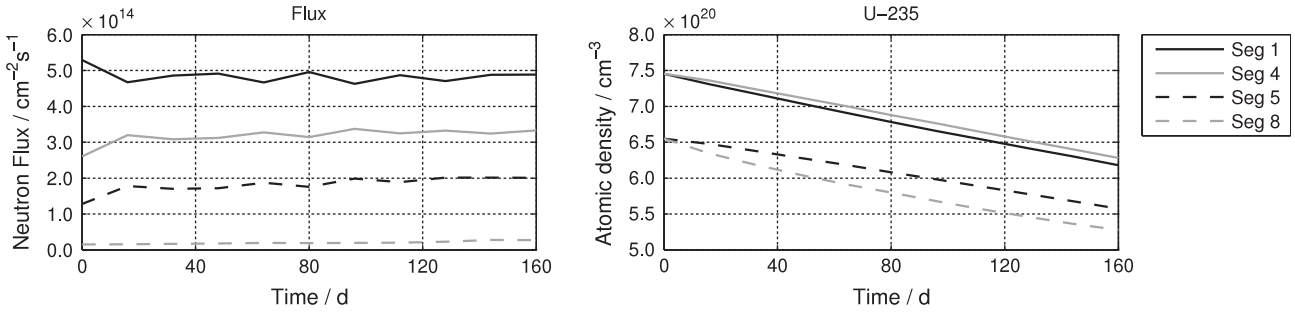

Fig. 5. The output flux and ^{235}U concentrations with 0.4 pp enrichment difference and 16 d steps.

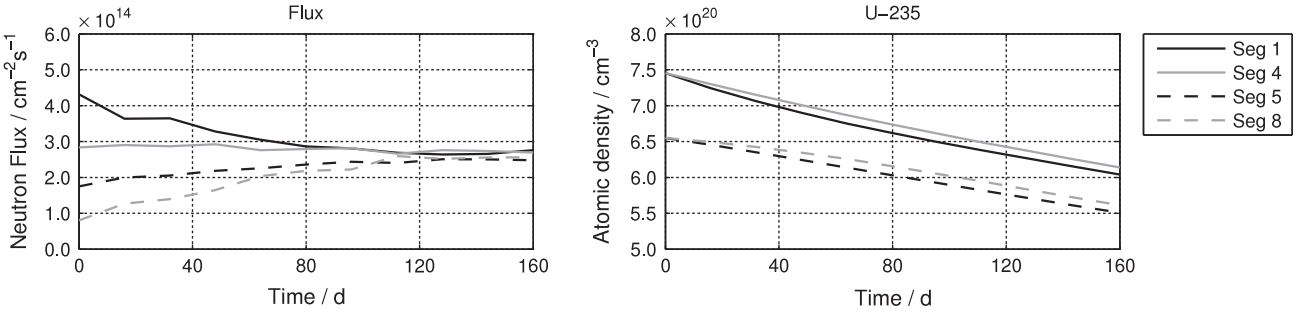

Fig. 6. Flux and ^{235}U concentrations with 0.4 pp enrichment difference and 16 d steps when equilibrium xenon calculation is used.

$$\sigma = \frac{1}{N_s \times (N_t - 2) \times N_r - 1} \times \sqrt{\sum_{s=1}^{N_s}\sum_{t=3}^{N_t}\sum_{r=1}^{N_r}\left(\frac{\phi_{s,t,r} - \bar{\phi}_t}{\bar{\phi}_t} - \mu\right)^2}, \quad (4)$$

where $\phi_{s,t,r}$ is the flux obtained for segment s at time step t in repeat r of the calculation with the eight segments and equilibrium xenon, whereas $\bar{\phi}_t$ is the mean flux obtained in the 2D calculations for step t. $N_s = 8$, $N_t = 10$ and $N_r = 5$ are the numbers of segments, steps and repeats with different random number sequences, respectively. The mean and standard deviation of ^{135}Xe were calculated similarly. The first steps are not included as it takes a few days for the xenon concentration to saturate in the reference calculations. After saturation there were no major time trends in the differences.

The differences in xenon concentration are comparable to, or smaller than, the standard deviation but there is a clear trend with the equilibrium calculation underestimating the concentration with poor statistics and, to a lesser extent, overestimating it with good ones. Differences in the flux reflect those in the xenon concentrations. The relatively large standard deviation in the flux appears to be of a purely statistical origin as it does not depend on the cycle size.

The underestimations of the xenon concentrations at smaller cycle sizes are caused by the bias, while the differences remaining with 50,000 neutrons per cycle result from the use of thermal yields for all fissions and discretization error (Isotalo and Aarnio, 2011b) in the reference calculations. More accurate results may be obtained with intermediate statistics when these errors cancel each other.

3.4. Running times and performance

In the tests of this work, the equilibrium xenon method increased the running times per neutron history by an average of 10%. This value is, however, suggestive at best as the slowdown depends on the geometry and various neutronics parameters as well as the implementations of the neutronics and equilibrium algorithms. Another factor affecting the performance, as measured by the ratio of variance and running time, is that when equilibrium xenon is used, the constant iteration of xenon concentrations gives

Table 2
Mean and standard deviation (in parentheses) of relative differences in neutron flux and xenon concentrations between equilibrium xenon calculations and the 2D reference calculations in the symmetric case when using different number of neutrons per source cycle. The number of source cycles was adjusted to keep the total number of active histories per solution at 25 million, except in the last case, where 250 million histories were used per step.

Cycle	2 day steps		30 day steps	
size	Xe	Flux	Xe	Flux
500	$-3.9 \times 10^{-2}(2.1 \times 10^{-2})$	$2.4 \times 10^{-3}(2.2 \times 10^{-2})$	$-4.9 \times 10^{-2}(2.6 \times 10^{-2})$	$2.6 \times 10^{-3}(2.4 \times 10^{-2})$
1000	$-1.5 \times 10^{-2}(1.0 \times 10^{-2})$	$1.3 \times 10^{-3}(1.9 \times 10^{-2})$	$-1.9 \times 10^{-2}(1.3 \times 10^{-2})$	$1.3 \times 10^{-3}(2.4 \times 10^{-2})$
1500	$-8.4 \times 10^{-3}(7.9 \times 10^{-3})$	$9.7 \times 10^{-4}(1.9 \times 10^{-2})$	$-1.1 \times 10^{-2}(1.0 \times 10^{-2})$	$9.1 \times 10^{-4}(2.3 \times 10^{-2})$
2000	$-5.3 \times 10^{-3}(7.7 \times 10^{-3})$	$8.0 \times 10^{-4}(2.1 \times 10^{-2})$	$-6.5 \times 10^{-3}(8.4 \times 10^{-3})$	$7.0 \times 10^{-4}(2.2 \times 10^{-2})$
3000	$-2.1 \times 10^{-3}(7.0 \times 10^{-3})$	$5.8 \times 10^{-4}(2.1 \times 10^{-2})$	$-2.7 \times 10^{-3}(8.4 \times 10^{-3})$	$5.5 \times 10^{-4}(2.4 \times 10^{-2})$
4000	$-7.5 \times 10^{-4}(6.6 \times 10^{-3})$	$4.9 \times 10^{-4}(2.1 \times 10^{-2})$	$-1.0 \times 10^{-3}(7.8 \times 10^{-3})$	$4.3 \times 10^{-4}(2.3 \times 10^{-2})$
5000	$2.9 \times 10^{-4}(5.3 \times 10^{-3})$	$4.5 \times 10^{-4}(1.7 \times 10^{-2})$	$-1.6 \times 10^{-5}(7.7 \times 10^{-3})$	$9.3 \times 10^{-5}(2.4 \times 10^{-2})$
50,000	$3.4 \times 10^{-3}(1.9 \times 10^{-3})$	$2.7 \times 10^{-4}(6.3 \times 10^{-3})$	$4.0 \times 10^{-3}(2.1 \times 10^{-3})$	$1.4 \times 10^{-4}(6.9 \times 10^{-3})$

Table 3
Average standard deviations of segment-wise fluxes and their 95% confidence intervals (assuming normality) with and without equilibrium xenon for various initial enrichment differences and burnup of 0 and 9.1 MW d per kilogram of initial heavy metal. The values are percentage of mean global flux.

Equilibrium xenon		Fresh fuel		Depleted fuel	
0.0 pp					
	No	5.11	(4.94 5.30)	5.67	(5.48 5.87)
	Yes	1.66	(1.61 1.72)	2.30	(2.22 2.38)
0.1 pp					
	No	3.71	(3.59 3.84)	2.00	(1.93 2.07)
	Yes	1.80	(1.74 1.87)	1.80	(1.74 1.86)
0.4 pp					
	No	2.90	(2.80 3.00)	1.95	(1.89 2.02)
	Yes	1.48	(1.43 1.53)	1.53	(1.48 1.59)
0.8 pp					
	No	1.08	(1.04 1.12)	2.01	(1.94 2.08)
	Yes	1.08	(1.05 1.12)	1.76	(1.70 1.82)
1.6 pp					
	No	0.91	(0.88 0.94)	1.62	(1.57 1.68)
	Yes	0.79	(0.77 0.82)	1.44	(1.39 1.49)

the flux a negative feedback inside the transport calculation. This improves the otherwise poor convergence in geometries with high dominance ratio.

To quantify the effect on statistics, an additional series of stand-alone neutronics calculations without burnup was performed. These calculations used fresh fuel and fuel that has been depleted for 300 d resulting in mean burnup of 9.1 MW d per kilogram of initial heavy metal. Compositions of depleted fuel were taken from the equilibrium xenon burnup calculations. In each case the flux was solved 200 times with and without equilibrium xenon. In the calculations which did not use the equilibrium method, equilibrium amounts of xenon were added to the fixed material compositions. The results from these calculations are shown in Table 3. In the symmetric case statistical variation in the flux was reduced by two thirds, and large reductions are also seen with 0.1 and 0.4 pp differences.

While the above tests disregarded burnup to allow comparable results to be obtained without equilibrium xenon, neutronics solutions in burnup calculations are affected in the same way. Reduced variation compensates for the increased running time per history, meaning that performance can be as good or even better as without equilibrium xenon. The comparison is only strictly applicable to independent neutronics solutions. In the tested geometries, burnup calculations without equilibrium xenon produce incorrect, and completely different, results which might have different variances.

3.5. Other methods

All calculations with step length of 1 d or more, except the 2D ones presented in Section 3.3 were repeated using each of the four other burnup algorithms available in Serpent 2, namely CE, LE, LE/LI and LE/QI (Isotalo and Aarnio, 2011a).[5] Results obtained with these methods were qualitatively similar to those presented in Section 3.2 for CE/LI, and are thus not presented in detail.

Without equilibrium xenon higher order predictor–corrector methods generally experienced slightly stronger oscillations, but behaved otherwise identically. Since CE and LE do not use a corrector, the flux starts to oscillate between steps rather than between the predictor and corrector. This leads to a behavior that is very different from the predictor–corrector methods, but the difference is not particularly significant since both modes of oscillation lead to unusable results.

With equilibrium xenon the only significant difference is that LE starts to oscillate with 30 d steps while all the other methods start at 60 d steps like CE/LI. For LE and CE the oscillations again happen between steps.

The tests were also repeated with CE/LI using equilibrium xenon and nine substeps (Isotalo and Aarnio, 2011b) on the corrector. Stability-wise the results were identical to those obtained without substeps.

4. Discussion

4.1. Instability of the models

The results obtained with 15 min steps for enrichment differences of 0–0.4 pp show a textbook example of physical xenon oscillations with period of about 38 h or 150 step lengths. Such clean oscillations cannot be caused by numerical instability alone. Instead, it appears that when starting from a non-equilibrium state rapidly growing oscillations are the correct solution to the model. With enrichment differences of 0.8 and 1.6 pp oscillations are still present but they are small and happen between the ends of the higher enriched part.

In the tests with non-zero enrichment difference flux and xenon concentrations are out of equilibrium to begin with, so the obtained oscillating solutions are correct for the simulation model. The correct solution with no enrichment difference has to be symmetric because there is no axial dependency in the initial state. However, as statistical variation breaks the symmetry, oscillations are unavoidable when using Monte Carlo neutronics. This difference in the origin of the oscillations explains the differences in their phases seen in Fig. 1. With no enrichment difference the oscillations are initiated by statistical variation and hence have random phases, whereas in the asymmetric case the oscillations have nothing to do with statistics. However, even in the symmetric case one can obtain the correct solution by utilizing the symmetry to perform a 2D calculation instead.

It is important to note that while the step lengths of 15 min (or 3 h) could be used to intentionally simulate physical xenon oscillations, this is not the case here. The model being solved is exactly the same that was used with longer steps; a model which, while simple, is like those used in typical burnup calculations that aim to simulate the development of the steady state. In the asymmetric cases, the model that was supposed to model steady state just turns out to describe physical xenon oscillations instead.

Thus the root of the xenon oscillations encountered with long steps is not in the burnup algorithms: They can handle the computational model for each of the tested cases when used correctly, i.e., with short enough steps. The real problem is that due to omitting feedback and control mechanisms, the simplified computational models do not actually describe the secular equilibrium we want to simulate. When trying to perform a calculation using incorrect model and step lengths longer than the period of oscillations in its solution, it is not surprising that results are unsatisfactory. As with 15 min steps, the oscillation mechanism in asymmetric cases is deterministic, and simply improving the statistics wont't remove the oscillations.

Since the numerical xenon oscillation mechanism encountered with step lengths typical for burnup calculations is different from the physical one, it has different stability properties and hence some burnup algorithm could, in principle, produce stable behavior despite the correct solution being oscillatory. This even seems

[5] CE and LE are single stage methods where cross-sections and flux for each step are set to the average of constant and linear extrapolations, respectively. LE/LI and LE/QI are predictor–corrector methods which follow the LE predictor step with corrector that uses linear or quadratic interpolation.

to be the case in the results obtained with 0.8 pp enrichment difference, where 1–4 d steps produced more stable results than 15 min steps. However, this seems situational at best as none of the methods tested in this or earlier works generates stable results with any generality.

4.2. Burnup calculations with equilibrium xenon

None of the above changes the fact that we still want to obtain 'stable solutions' with simplified computation models. Perhaps the simplest way to achieve this is to directly require that the flux and xenon concentrations must remain in a mutual equilibrium. It should be noted that doing so does not actually produce a stable solution to the original model, but changes the model to include an additional constraint that prevents oscillations. In the tests of this work using equilibrium xenon with step lengths up to 30 d removed all oscillations leaving no clear discrepancies apart from the non-zero xenon concentrations in fresh fuel and the resulting changes in neutronics. The method is also problem independent and does not require any additional input or details in the model.

In those cases where comparable results could be obtained without equilibrium xenon, agreement between them and the equilibrium algorithm of Serpent was good as long as the number of neutrons per source cycle was sufficient. The algorithm also improved the convergence of the neutronics solution, but when the neutron batch size is set too low, results deteriorate due to the bias in the xenon estimate.

It is important to understand that this bias depends on the number of neutrons in a single source cycle, not the total number of neutron histories per step like all other statistical phenomena. Each xenon concentration update in the equilibrium calculation uses statistics from only a single source cycle, and the probability of the flux, and hence xenon concentration, being severely underestimated during given cycle increases rapidly as the number of neutrons per cycle becomes small. Normalization ensures that when the flux in one part of the geometry is underestimated, it must be overestimated equally much somewhere else, but because the dependence of the xenon concentration on the flux is not linear, large overestimations of flux do not increase the xenon concentration as much as equally large underestimations decrease it. Thus the average xenon concentration is always underestimated, resulting in a bias.

Only large statistical variations lead to significant bias as the dependence of xenon on flux, just like any other smooth function, is essentially linear in respect to small variations. Thus the bias can be avoided by making the neutron generation size sufficiently large. The number of source cycles can be made correspondingly smaller so that overall statistical accuracy and the running time required for active cycles remain unaffected. Unfortunately there is no straightforward way to estimate how large cycles are required in a given case without performing test calculations. It should be possible to eliminate the bias by updating the xenon concentrations based on the unbiased estimator of Dumonteil and Diop (2011). This and other possible improvements to the algorithm are a topic for further study.

The stabilizing effect of the equilibrium xenon treatment is not dependent on which algorithm is used for calculating the equilibrium, and for example the algorithm of Griesheimer (2010), which is not susceptible to bias or doubts about correctness, could be used instead of the algorithm of Serpent. Any reasonable integrated algorithm should have comparably low effect on running times, but the effect on convergence may vary greatly. Whichever algorithm is used, it should be kept in mind that 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.

4.3. Other nuclides

Even with equilibrium xenon, strong oscillations were observed when using too long steps. As without equilibrium xenon, the oscillations in different repeats of the asymmetric cases always happen in the same phase indicating that they are triggered by the asymmetry of the geometry, not statistical variation. The mechanisms are similar to the numerical xenon oscillations with negative reactivity associated with increased local burnup replacing that from the buildup of xenon. The step length limit at which these oscillations start is likely to be affected by the geometry, material compositions and power density as well as the choice of burnup algorithm. The equilibrium xenon method only affects xenon, and while it would be possible, in principle, to apply a similar equilibrium treatment for other short-lived nuclides, this is not possible for the long-lived ones.

Since oscillations can occur despite equilibrium xenon, results still need to be checked for signs of instability and step lengths adjusted as needed. With predictor–corrector methods comparing the results from subsequent steps is not sufficient as oscillations can happen between the predictor and corrector. These oscillations can be detected by comparing the local predictor and corrector fluxes, which could easily be automated to generate a warning for differences larger than some predefined threshold. Unfortunately universal or exact threshold cannot be established since varying differences between the predictor and corrector values are a normal part of the methods.

5. Conclusions

Existing Monte Carlo burnup codes suffer from xenon driven spatial oscillations in large geometries. Since the accurate solutions to the simulation models for many of these geometries oscillate, simply reducing step lengths, or solving the model better, would not work in a general case. When using predictor–corrector methods these oscillations may be difficult to detect as they can happen between the predictor and corrector, instead of subsequent steps, which makes the usually collected results look stable.

Forcing xenon and flux to mutual equilibrium offers a simple solution to these oscillations. Since the equilibrium calculation can be integrated to normal Monte Carlo neutronics, it has only a minor effect on running times and can be used with any burnup calculation algorithm.

Oscillations driven by fuel burnup still arise if too long steps are used, but unlike xenon oscillations, these only occur with long steps, allowing calculations to be performed with reasonable step lengths. This is a major improvement, especially in cases where stable solutions simply could not be obtained without the equilibrium method. It is important to remember that while the equilibrium xenon treatment allows stable solutions to be obtained for any model, the equilibrium levels are only as accurate as the model they have been calculated for.

Acknowledgement

Funding from the SAFIR2014 project and computational resources provided by Aalto Science-IT project are acknowledged.

References

Dufek, J., Gudowski, W., 2006. Stochastic approximation for Monte Carlo calculation of steady-state conditions in thermal reactors. Nucl. Sci. Eng. 152, 274–283.

Dufek, J., Hoogenboom, J.E., 2009. Numerical stability of existing Monte Carlo burnup codes in cycle calculations of critical reactors. Nucl. Sci. Eng. 162, 307–311.
Dufek, J., Kotlyar, D., Shwageraus, E., Leppänen, J., 2013. Numerical stability of the predictor–corrector method in Monte Carlo burnup calculations of critical reactors. Ann. Nucl. Energy 56, 34–38.
Dumonteil, E., Diop, M.C., 2011. 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.
Griesheimer D.P., 2010. In-Line Xenon Convergence Algorithm for Monte Carlo Reactor Calculations. PHYSOR2010, Pittsburgh, Pennsylvania, USA, May 9–14, 2010. On CD-ROM, ISBN:978-0-89448-079-9.
Isotalo, A.E., Aarnio, P.A., 2011a. Higher order methods for burnup calculations with Bateman solutions. Ann. Nucl. Energy 38, 1987–1995.
Isotalo, A.E., Aarnio, P.A., 2011b. Substep methods for burnup calculations with Bateman solutions. Ann. Nucl. Energy 38, 2509–2514.