

A simple implicit coupling scheme for Monte Carlo neutronics and isotopic depletion

P. Cosgrove, E. Shwageraus, G.T. Parks

Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, United Kingdom

Abstract

The stochastic implicit Euler scheme and its variants can be used to prevent non-physical behaviour that may emerge when coupling Monte Carlo neutron transport and isotopic depletion solvers for spatially-decoupled reactor problems. However, stochastic implicit methods tend to require many iterations to obtain a stable solution. This paper demonstrates that this is due to using a sub-optimal relaxation scheme: rather than using a variable relaxation factor, a fixed relaxation factor can give a more stable solution in fewer iterations. Furthermore, like stochastic implicit schemes, even though multiple transport solutions are required, using a fixed relaxation factor allows computational effort to be reduced by lowering the number of particles simulated during the corrector step while still providing stable results. This shows that using a fixed relaxation factor to stabilise Monte Carlo burn-up calculations can be more effective than applying the stochastic approximation.

Keywords: Depletion, Neutron transport, Monte Carlo, Stability

1. Introduction

The Monte Carlo method is now routinely applied to neutron transport for reactor analysis problems, often alongside other physics solvers, e.g., for thermal-hydraulics or isotopic depletion. However, for many problems of any complexity, this coupling often produces non-physical results and efforts must be made to stabilise it (Gill et al., 2017; Griesheimer et al., 2017).

When coupling with isotopic depletion in particular, there are two techniques commonly applied to stabilise the solution. The first of these is enforcing xenon equilibrium, wherein the xenon-135 chain is removed from the burn-up matrix and the xenon and iodine concentrations in a given burnable region are instead determined by the value of the flux in that region (Griesheimer, 2010; Isotalo et al., 2013). While this appears to be stable, it has been shown that numerical instabilities may arise in transport-depletion systems even without the presence of xenon (Densmore et al., 2013). The second is the stochastic implicit Euler

Email address: pmc55@cam.ac.uk (P. Cosgrove)

(SIE) scheme and its higher-order variants (Dufek et al., 2013a; Kotlyar and Shwageraus, 2014, 2016). Stochastic implicit (SI) schemes require multiple transport solutions for each burn-up step to perform an averaging over either the calculated reaction rates or nuclide densities. A convenient property of SI schemes is that, with a sufficiently large number of iterations, the nuclide density solution is guaranteed to converge to a fixed vector. The scheme derives from the stochastic approximation (Robbins and Monro, 1951) which has also been applied to coupling Monte Carlo neutronics with thermal-hydraulics (Tuominen, 2015). The stochastic approximation can be shown to be equivalent to either an averaging over multiple iterations or a fixed point iteration with a gradually diminishing relaxation factor (Dufek et al., 2013a).

Recently the phenomenon of neutron clustering has been shown to strongly affect Monte Carlo burn-up simulations (Dumonteil et al., 2014; Zoia et al., 2014; de Mulatier et al., 2015; Nowak et al., 2016; Sutton and Mittal, 2017; Cosgrove et al., 2019). Minimising its effects entails simulating large numbers of neutrons per cycle for fewer cycles (Dumonteil et al., 2014; Sutton and Mittal, 2017). Nevertheless, numerical instabilities in coupling Monte Carlo neutronics and depletion persist, albeit requiring more or longer time-steps to appear than if simulating few neutrons over many generations (Cosgrove et al., 2019).

Anecdotally, in the authors' experience, for the relatively coarsely discretised problems for which burn-up instability is reported, the degree of statistical noise in the Monte Carlo flux solution is quite small, often on the order of 0.1% when simulating a reasonably large number of particles per cycle – say 3,000,000 over tens of cycles. These settings should lessen the effects of both clustering and inter-generational correlation, giving reasonable estimates of statistical noise (Herman, 2014). If statistical noise is relatively minute in the Monte Carlo solution, then the justification for using the stochastic approximation is weakened, given that it was devised primarily to guarantee locating the fixed point of noisy functions (Robbins and Monro, 1951). In fact, the original paper emphasises that '[n]o claim is made that the procedure to be described has any optimum properties (i.e. that it is efficient)' (Robbins and Monro, 1951).

This work was conceived of by supposing that SI schemes have proved successful due to minimising clustering effects, even when simulating relatively small particle populations – this may be due to SI schemes' equivalence to averaging procedures, smoothing out any bias in individual transport solutions contaminated by neutron clustering. However, the coupled transport-depletion system at the end-of-step (EOS) is a fixed point iteration to which SI schemes apply a relaxation which may be necessary for stability, even without clustering effects.

Therefore, following a review of the relevant background, this work numerically investigates the stability behaviour of various depletion coupling schemes while minimising the effects of neutron clustering. As a result, a simpler sibling of the SI scheme is proposed which, rather than varying the relaxation factor with iterations, uses a fixed relaxation factor. This scheme is also investigated numerically with a comparison against the equivalent SI schemes.

2. Theory

This section will discuss the standard transport-depletion coupling schemes, as well as covering some aspects of solving non-linear systems with fixed-point iteration.

2.1. Transport & depletion coupling

Isotopic depletion is now a routine component of many Monte Carlo codes (Leppänen et al., 2015; Griesheimer et al., 2015; Josey, 2017). These codes typically offer a large variety of options for coupling the transport and depletion solvers. Most basically, the Monte Carlo code must have specified some discretisation of the problem into separate burnable regions where reaction rates and fission yields are tallied during the simulation, normalised by a specified operating power or an equivalent. Along with decay constants, this information is often combined into what is known as the burn-up matrix, $\mathbf{A}(\psi)$, with the explicit dependence on the neutron flux solution, ψ , shown here. The burn-up matrix for each burnable region is fed to a depletion solver which must solve the Bateman equations for the evolution of however many nuclides are being followed across the time-step (Bell and Glasstone, 1970):

$$\frac{dN}{dt} = \mathbf{A}(\psi)N \quad (1)$$

Here N is the nuclide density vector in a given region and t is the time variable. Typically on the order of hundreds of isotopes are followed per burnable region (Griesheimer et al., 2017). Often the numerical solution of the Bateman equation involves using the formal matrix exponential solution of Eq. (1) over some time-step of length Δt starting from some initial time t_0 :

$$N(t) = \exp \left[\mathbf{A}(\psi) \cdot \Delta t \right] N(t_0) \quad (2)$$

This matrix exponential can be efficiently and accurately evaluated using the Chebyshev rational approximation method (CRAM) (Pusa, 2013). That said, numerous other methods exist and show comparable speed and accuracy (Griesheimer et al., 2017; Josey et al., 2017).

Given that the neutron flux depends on the nuclide density distribution, the evolution of which likewise depends on the flux, the neutron transport equation and Bateman equation form a non-linear system. The simplest solution method is to neglect this non-linearity and assume that the neutron flux remains fixed over the course of a time-step – this is the explicit Euler scheme, shown in Algorithm 1, with $\psi(N_n)$ denoting the transport solution obtained from the nuclide density distribution at time-step n . Essentially, the explicit Euler scheme can be described as, for a given time-step, evaluating a transport solution, followed by the application of Eq. (2). However, due to error accrued from neglecting the change in spectrum across the time-step, in practice the explicit Euler scheme restricts the user to impractically short steps. Hence, predictor-corrector (PC) methods are often used instead, allowing for longer time-steps due to their higher order but at the cost of additional transport solutions.

One of the more common PC methods is that which Isotalo and Aarnio (2011a) term the CE/LI method, standing for Constant Extrapolation/Linear Interpolation. Initially the scheme is identical to the explicit Euler method: at the beginning-of-step (BOS) a transport

Input: $N_0, n_{\text{steps}}, \Delta t$
for $n = 0, \dots, n_{\text{steps}} - 1$ **do**
 | $\mathbf{A}_n \leftarrow \psi(N_n)$
 | $N_{n+1} \leftarrow \exp[\mathbf{A}_n \Delta t_n] N_n$
end

Algorithm 1: Explicit Euler method.

solution is obtained and the reaction rates are used to burn the problem forward in time to the EOS, assuming the reaction rates are constant. At the EOS another transport solution is obtained and an average of the BOS and EOS reaction rates are obtained and used to burn the BOS nuclide density forward again – in principle with a more time-step representative spectrum. The scheme is shown in Algorithm 2.

Input: $N_0, n_{\text{steps}}, \Delta t$
for $n = 0, \dots, n_{\text{steps}} - 1$ **do**
 | $\mathbf{A}_n \leftarrow \psi(N_n)$
 | $N_{n+1} \leftarrow \exp[\mathbf{A}_n \Delta t_n] N_n$
 | $\mathbf{A}_{n+1} \leftarrow \psi(N_{n+1})$
 | $\mathbf{A} \leftarrow \frac{1}{2}(\mathbf{A}_n + \mathbf{A}_{n+1})$
 | $N_{n+1} \leftarrow \exp[\mathbf{A} \Delta t_n] N_n$
end

Algorithm 2: Predictor-Corrector method.

The accuracy of the CE/LI scheme can be improved upon, however. The first means of doing so is by introducing the notion of ‘substeps’. Substeps in depletion refer to dividing the time interval into a number of sub-intervals across which the reaction rates may vary more smoothly, rather than being fixed (Isotalo and Aarnio, 2011a,b). This is simple to achieve: rather than having a fixed weighting factor for the BOS and EOS reaction rates (0.5 in Algorithm 2), the weights vary across substeps. The only cost to substepping is the additional depletion calculations: rather than having one depletion calculation for a time interval, there are as many as there are substeps. This additional expense is not usually so onerous given that, even for problems with many burnable regions, Monte Carlo solutions are typically many times more computationally expensive than a depletion calculation (Griesheimer et al., 2017).

A second substantial gain in accuracy may also be obtained by using a higher-order extrapolation: rather than assuming reaction rates are constant during the extrapolation, if previous reaction rate data is available one may use it to inform the extrapolation. Isotalo (2013) introduced a successful linear extrapolation which provides a substantial increase in accuracy even if used in isolation, although it is often used as the predictor step in what is termed the LE/LI scheme. The LE/LI scheme is shown in Algorithm 3. In the algorithm the w variables are the weights used to determine the contribution of the BOS, EOS, or the previous BOS during a substep of the LE or LI. In principle, even higher-order interpolation schemes can be used, such as allowing for a quadratic variation during the corrector step – a

QI corrector. However, Isotalo (2013) suggests the benefits of QI over LI are not substantial and notes that the scheme is limited due to not being applicable to cores with discontinuous operating conditions, e.g., when a control rod is withdrawn. Holding an additional set of reaction rates in memory may also be challenging for large burn-up problems (Griesheimer et al., 2017).

```

Input:  $N_0, n_{\text{steps}}, \Delta t, n_{\text{ss}}$ 
for  $n = 0, \dots, n_{\text{steps}} - 1$  do
   $\mathbf{A}_n \leftarrow \psi(N_n)$ 
  if  $n > 0$  then
     $N \leftarrow N_n$ 
    for  $s = 0, \dots, n_{\text{ss}} - 1$  do
       $\mathbf{A} \leftarrow w_{n,s} \mathbf{A}_n + w_{n-1,s} \mathbf{A}_{n-1}$ 
       $N \leftarrow \exp \left[ \mathbf{A} \frac{\Delta t_n}{n_{\text{ss}}} \right] N$ 
    end
     $N_{n+1} \leftarrow N$ 
  else
     $N_{n+1} \leftarrow \exp [\mathbf{A}_n \Delta t_n] N_n$ 
  end
   $\mathbf{A}_{n+1} \leftarrow \psi(N_{n+1})$ 
   $N \leftarrow N_n$ 
  for  $s = 0, \dots, n_{\text{ss}} - 1$  do
     $\mathbf{A} \leftarrow w_{n,s} \mathbf{A}_n + w_{n+1,s} \mathbf{A}_{n+1}$ 
     $N \leftarrow \exp \left[ \mathbf{A} \frac{\Delta t_n}{n_{\text{ss}}} \right] N$ 
  end
   $N_{n+1} \leftarrow N$ 
end

```

Algorithm 3: LE/LI substep method.

While these schemes have allowed for high accuracy even where the flux spectrum may vary strongly over a time-step, large, neutronically-decoupled problems have been noted to pose a challenge to the methods, resulting in non-physical solutions (Dufek and Hoogenboom, 2009; Dufek et al., 2013b; Kotlyar and Shwageraus, 2013; Isotalo et al., 2013; Kępisty and Cetnar, 2015; Josey, 2017). It has since been shown that most of the problems considered will have been strongly and adversely affected by the Monte Carlo settings used, but accounting for this alone will not ensure a physical solution (Cosgrove et al., 2019).

The SIE scheme was proposed as a means of obtaining stable solutions to these large depletion problems (Dufek et al., 2013a). The original scheme would actually induce errors through using the EOS reaction rates in isolation to burn forward in time, and thus its higher-order methods were introduced, in part, to remedy this (Kotlyar and Shwageraus, 2014, 2016). For what is known as the SIMP scheme (or, perhaps more descriptively, the SI-CE/LI scheme) the problem is burned forward using a CE predictor. The LI corrector

step is iterated a number of times, performing both a transport and depletion calculation each time. At the conclusion of iteration k , the nuclide densities of all previous iterations are averaged such that:

$$N^{(k+1)} = \frac{1}{k+1} \sum_{i=0}^k \bar{N}^{(i)} \quad (3)$$

This new nuclide density vector is then used to begin the subsequent iteration. Performing enough iterations, one should obtain a stable, fixed solution according to what has been proved about the Robbins-Monro algorithm (Robbins and Monro, 1951). This procedure may also be applied to reaction rates and similar results have been reported for the two methods. The SI-CE/LI scheme using nuclide density averaging is shown in Algorithm 4. This scheme can also be straightforwardly extended to feature an LE predictor and substepping.

```

Input:  $N_0, n_{\text{steps}}, \Delta t, k_{\text{max}}$ 
for  $n = 0, \dots, n_{\text{steps}} - 1$  do
   $\mathbf{A}_n \leftarrow \psi(N_n)$ 
   $N_{n+1}^{(0)} \leftarrow \exp[\mathbf{A}_n \Delta t_n] N_n$ 
  for  $k = 0, \dots, k_{\text{max}} - 1$  do
     $\mathbf{A}_{n+1} \leftarrow \psi(N_{n+1}^{(k)})$ 
     $\mathbf{A} \leftarrow \frac{1}{2} (\mathbf{A}_n + \mathbf{A}_{n+1})$ 
     $\bar{N}_{n+1}^{(k)} \leftarrow \exp[\mathbf{A} \Delta t_n] N_n$ 
     $N_{n+1}^{(k+1)} \leftarrow \frac{1}{k+1} \sum_{i=0}^k \bar{N}_{n+1}^{(i)}$ 
  end
   $N_{n+1} \leftarrow N_{n+1}^{(k_{\text{max}})}$ 
end

```

Algorithm 4: Stochastic implicit method with nuclide density averaging applied to the CE/LI scheme.

Before describing how SI schemes provide stable solutions, it is worthwhile explaining how non-linear systems of equations are typically solved, as this is what SI schemes aim to do.

2.2. Transport & depletion as a non-linear system

As stated above, the transport-depletion system is non-linear. This can be made more explicit by first noting that the CE/LI scheme can be described succinctly by two equations:

$$N_{n+1} = \exp[\mathbf{A}(\psi_n) \Delta t] N_n \quad (4)$$

$$N_{n+1} = \exp\left[\frac{\mathbf{A}(\psi_n) + \mathbf{A}(\psi_{n+1})}{2} \Delta t\right] N_n \quad (5)$$

The first of these equations is the extrapolation/predictor step and gives the EOS nuclide density for the current time-point, whereas the latter corrects it to give the BOS nuclide

density for the subsequent time-point as carried out by the standard CE/LI method. Here ψ_n and ψ_{n+1} denote, respectively, the BOS and EOS transport solutions. The transport solution can be written as the result from a transport operator, $\tau(\cdot)$, acting on a nuclide density field:

$$\psi = \tau(N) \quad (6)$$

Instead of applying the PC scheme without iteration – which is common in reactor physics – one could iterate on the corrector step. This would give the equations:

$$N_{n+1}^{(0)} = \exp [\mathbf{A}(\psi)\Delta t] N_n \quad (7)$$

$$N_{n+1}^{(k+1)} = \exp \left[\frac{\mathbf{A}(\psi) + \mathbf{A}(\psi^{(k)})}{2} \Delta t \right] N_n \quad (8)$$

Note the subscripts on the transport solutions have been dropped and instead they are differentiated through the BOS solution being unadorned while the EOS solution has an iteration index. Likewise, the nuclide density field at time-point $n + 1$ is now given a corresponding iteration index; there is no longer only a ‘predicted’ and ‘corrected’ EOS nuclide density field, but as many as the number of EOS iterations, plus one. An iterated corrector will be computationally expensive due to the additional transport solutions required. In the literature, excluding the cases where SI schemes are used, the authors are only familiar with an iterated corrector step being used in MC21 ‘to improve convergence of the [EOS nuclide] densities’ without any further elaboration (Griesheimer et al., 2017).

The corrector step given by Eq. (8) may be written to show explicit dependence upon the nuclide density from the previous iteration:

$$N_{n+1}^{(k+1)} = \exp \left[\frac{\mathbf{A}(\psi) + \mathbf{A}(\tau(N_{n+1}^{(k)}))}{2} \Delta t \right] N_n \quad (9)$$

Naturally, there is an equivalent equation when taking n_s substeps during the corrector step (with the first substep as the rightmost matrix multiplication):

$$N_{n+1}^{(k+1)} = \prod_{s=1}^{n_s} \left(\exp \left[\frac{w_{n,s} \mathbf{A}(\psi) + w_{n+1,s} \mathbf{A}(\tau(N_{n+1}^{(k)}))}{n_s} \Delta t \right] \right) N_n \quad (10)$$

As with, say, coupling neutronics and thermal-hydraulics, this iterated corrector step has the form of a non-linear map:

$$N_{n+1}^{(k+1)} = F \left(N_{n+1}^{(k)} \right) \quad (11)$$

The map, F , combines the action of both the transport operator, τ , and the depletion operator, β , like so:

$$N_{n+1}^{(k+1)} = \beta \left(\Delta t, N_n, \psi, \tau \left(N_{n+1}^{(k)} \right) \right) = F \left(N_{n+1}^{(k)} \right) \quad (12)$$

Attaining a stable solution implies finding a fixed point of F , this fixed point being the solution vector N_{n+1}^* which satisfies:

$$N_{n+1}^* = F(N_{n+1}^*) \quad (13)$$

The easiest means of locating this fixed point is through iteration – feeding successive solutions into the map until they no longer differ appreciably. However, depending on the map in question this may not be possible: rather than converging, the vector may oscillate or diverge. In computational physics, convergence to a fixed point is often achieved by using ‘relaxation’. Relaxation implies iterating on a modified version of Eq. (11):

$$N_{n+1}^{(k+1)} = \alpha F(N_{n+1}^{(k)}) + (1 - \alpha)N_{n+1}^{(k)} \quad (14)$$

For ‘underrelaxation’ which is applied when seeking a stable solution, as is the case here, the value of α is chosen under the constraint $0 < \alpha < 1$. The smaller the value of α , the more stable is the iteration but at the cost of converging more slowly towards the fixed point. SI schemes apply a relaxation but of the form:

$$N_{n+1}^{(k+1)} = \alpha_k F(N_{n+1}^{(k)}) + (1 - \alpha_k)N_{n+1}^{(k)} \quad (15)$$

The difference is the iteration-dependent relaxation factor, α_k . The Robbins-Monro algorithm dictates that the sequence of relaxation factors should possess certain properties which happen to be satisfied by (Robbins and Monro, 1951):

$$\alpha_k = \frac{1}{k + 1} \quad (16)$$

As mentioned above, this scheme was introduced to guarantee the convergence of a fixed point iteration subject to a noisy function evaluation with no guarantees of optimality or efficiency (Robbins and Monro, 1951). Therefore it may be the case that a fixed relaxation scheme, or even no relaxation, has a superior performance to the equivalent SI schemes. Consider the extremes of the SI relaxation scheme: on the first iteration, no relaxation is applied as $\alpha_0 = 1$. Therefore, this provides no benefit to stability unless corrector step iteration is already stable. On the other hand, after many iterations, $\lim_{k \rightarrow \infty} \alpha_k \rightarrow 0$ and so there is no additional convergence taking place. Rather, some intermediate value or values of α should presumably be superior on both of these fronts. Hence, the subject of the remainder of this paper is: what is a good value of α for the coupled transport-depletion problem, and for how many iterations?

It should also be highlighted that the current analysis is nearly exactly analogous with methods applied to coupling Monte Carlo neutronics and thermal-hydraulics. Coupled neutronics and thermal-hydraulics system are also described by a non-linear map, often exhibiting instability which is strongly driven by the Doppler feedback between iterations (Gill et al., 2017). By applying relaxation to either the nuclide or temperature field, numerous authors have managed to stabilise this coupled system, some even using the stochastic approximation (Kotlyar and Shwageraus, 2014; Tuominen, 2015). This was thoroughly surveyed by

Gill et al. (2017). The novelty in this work is that, although the stochastic approximation has been applied to attaining a stable neutronics/depletion system before, the fixed relaxation factor equivalent has not. In future, it may be interesting to explore applying to depletion other coupling schemes which have been pioneered in thermal-hydraulics, such as the application of Newton method's to achieve stability (Aufiero and Fratoni, 2017).

3. Numerical investigations and results

The numerical experiments are divided over two subsections: one briefly investigating the effect of iterating on the corrector step without relaxation, the other with relaxation. These are done for the CE/LI scheme without substeps, and the CE/LI and LE/LI schemes with 10 substeps. An LE/LI scheme with substeps and relaxation is shown in Algorithm 5 – all of the schemes investigated here are simplifications of this, except for the SI version which allows α to vary with the corrector iteration.

```

Input:  $N_0, n_{\text{steps}}, \Delta t, n_{\text{ss}}, k_{\text{max}}, \alpha$ 
for  $n = 0, \dots, n_{\text{steps}} - 1$  do
   $\mathbf{A}_n \leftarrow \psi(N_n)$ 
  if  $n > 0$  then
     $N \leftarrow N_n$ 
    for  $s = 0, \dots, n_{\text{ss}} - 1$  do
       $\mathbf{A} \leftarrow w_{n,s} \mathbf{A}_n + w_{n-1,s} \mathbf{A}_{n-1}$ 
       $N \leftarrow \exp \left[ \mathbf{A} \frac{\Delta t_n}{n_{\text{ss}}} \right] N$ 
    end
     $N_{n+1}^{(0)} \leftarrow N$ 
  else
     $N_{n+1}^{(0)} \leftarrow \exp \left[ \mathbf{A}_n \Delta t_n \right] N_n$ 
  end
  for  $k = 0, \dots, k_{\text{max}} - 1$  do
     $\mathbf{A}_{n+1}^{(k)} \leftarrow \psi(N_{n+1}^{(k)})$ 
     $N \leftarrow N_n$ 
    for  $s = 0, \dots, n_{\text{ss}} - 1$  do
       $\mathbf{A} \leftarrow w_{n,s} \mathbf{A}_n + w_{n+1,s} \mathbf{A}_{n+1}^{(k)}$ 
       $N \leftarrow \exp \left[ \mathbf{A} \frac{\Delta t_n}{n_{\text{ss}}} \right] N$ 
    end
     $N_{n+1}^{(k+1)} \leftarrow \alpha N + (1 - \alpha) N_{n+1}^{(k)}$ 
  end
   $N_{n+1} \leftarrow N_{n+1}^{(k_{\text{max}})}$ 
end

```

Algorithm 5: LE/LI substep method with corrector iteration and relaxation.

Table 1. Burn-up schedule.

Time-point	Time (d)	Burn-up (MWd/kg)
1	0	0
2	1	0.034
3	3	0.10
4	8	0.27
5	18	0.61
6	38	1.3
7	78	2.6
8	138	4.7
9	198	6.7
10	258	8.7
11	318	11
12	378	13
13	438	15

Each of these sections will show the results of burn-up calculations on a single 3.66 m tall PWR pin with a pitch of 1.26 cm. The pin is divided into 10 equal-volume burnable regions, each composed of 5% enriched UO_2 with a radius of 0.4095 cm. The pin is clad with natural zirconium with a thickness of 0.0655 cm. The moderator is pure water with a uniform density of 0.7 g/cm^3 . Above and below the pin is a homogenised steel-water reflector material. The boundary conditions are radially reflective and axially vacuum. The pin power is fixed at 60 kW, corresponding to a PWR power density of 104 W/cm^3 , and the depletion schedule followed is the same for all simulations, shown in Table 1. This problem is chosen because its true solution must be symmetric at all times, making it relatively easy to identify whether or not a given burn-up scheme is stable. Such pin geometries with either reflected or vacuum boundaries axially are reasonably representative of unstable reactor problems due to their characteristic length being many times the neutron mean free path. Furthermore, these problems are often considered when investigating Monte Carlo burn-up instabilities and have proven to be a challenging case when ensuring stable solutions are attained (Dufek and Hoogenboom, 2009; Dufek et al., 2013b; Kotlyar and Shwageraus, 2013; Dufek et al., 2013a; Isotalo et al., 2013; Josey, 2017; Cosgrove et al., 2019). In principle, a non-uniform coolant density problem could also be examined; however, in such a case there is no obvious stability diagnostic to exploit to determine whether one scheme or another is superior.

Burn-up calculations are performed using a modified version of Serpent 2.1.31 (Leppänen et al., 2015) coupled with an external MATLAB depletion script. At the conclusion of a transport simulation the modified version of Serpent prints one-group cross-sections, fluxes, and fission yields, as well as decay constants and a list of isotopes being tracked – it then terminates. The MATLAB script reads this information and performs the depletion calculations using a CRAM solver, updating the material compositions in Serpent accordingly. This is necessary given that Serpent does not include iterated or relaxed versions of the CE/LI or LE/LI schemes presently.

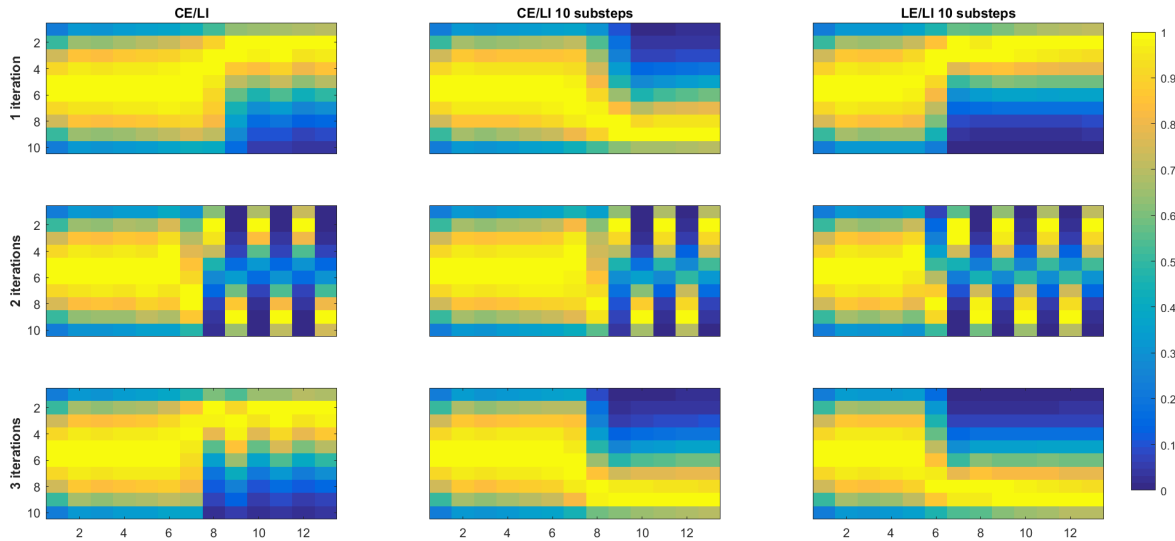


Fig. 1. Flux profile over time for a PWR fuel pin with axially vacuum boundaries using three different burn-up schemes and varying the number of corrector iterations. The horizontal axes on each graph are the burn-up time-points given in Table 1, while the vertical axes are the flux values at a given axial node from 1 to 10, with the fluxes at a given time-point normalised by the maximum value.

Nearly all calculations shown use, for the Monte Carlo settings, 2,000,000 neutrons per cycle, 80 active cycles, and 10 inactive cycles, with source convergence accelerated using Serpent’s response matrix method (Leppänen, 2019). The response matrix solver performed 20 outer iterations with one source batch across a 1D mesh along the height of the geometry, discretised into 60 equal-height nodes. These settings were chosen primarily to minimise the effects of neutron clustering. Where different settings were used, these will be highlighted.

3.1. Iterating on the corrector without relaxation

The first question to be answered is whether the non-linear system describing the corrector step for this problem is stable without any relaxation; it may be the case that, if the fixed point is not reached using a single corrector iteration, it may be closely approached after a number of iterations. This is straightforward to determine: for each of the burn-up schemes considered, the same simulation may be repeated but using differing numbers of corrector steps. These results are shown in Fig. 1, borrowing a presentation style from Gill et al. (2017).

In all cases, regardless of the scheme or the number of corrector iterations, non-physical results are produced, implying that the corrector step without some form of relaxation is not stable, or at least requires many iterations to converge to a stable solution. Using one or three corrector iterations, rather than two, produces superficially quite different results – the flux either tilts to one side and reinforce itself there or will oscillate. However, Isotalo (2013) notes that this difference in instability behaviour can also be observed when comparing the explicit Euler scheme and any PC scheme with a single corrector iteration – flux oscillation




Fig. 2. Flux asymmetry for the CE/LI scheme without substeps while varying the number of corrector iterations.

is still occurring in all cases but is only noticeable when there is an even number of corrector steps.

One may also note that the LE/LI scheme appears marginally more unstable than the others, beginning to tilt more quickly and extremely – the flux has become noticeably non-physical by the sixth step, rather than the seventh or eighth for either version of the CE/LI scheme. It may be the case that the LE predictor more readily excites non-physical solutions than the CE version; this was suggested by Isotalo et al. (2013), who report that, even when enforcing xenon equilibrium, oscillations occur with the LE scheme when taking relatively short time-steps of 30 days.

Although only slight, there is the appearance that additional corrector calculations negatively affect the stability of the simulation, inducing a stronger flux tilt one or two time-steps earlier than they would appear otherwise. This is shown for the CE/LI scheme without substeps in Fig. 2, where the asymmetry of the flux profile at a time-point is given by:

$$\epsilon = \sum_{i=1}^5 \frac{|\phi_i - \phi_{11-i}|}{\max(\phi)} \quad (17)$$

Here ϕ_i denotes the scalar flux calculated in burnable region i at a given time-point.

Given these results, it appears reasonable to conclude that some form of stabilisation is necessary and that it can be achieved by modifying the corrector step.

3.2. Iterating on the corrector with relaxation

To investigate the efficacy of different relaxation schedules, for each of the burn-up schemes specified above, the relaxation factor is varied from 0.3 to 0.6 and the number of corrector iterations is varied from 1 to 3. For comparison, both the corresponding SI




Fig. 3. Flux profile over time for a PWR fuel pin with axially vacuum boundaries using the CE/LI scheme with no substeps and varying the number of corrector iterations and relaxation factor. The horizontal axes on each graph are the burn-up time-points given in Table 1, while the vertical axes are the flux values at a given axial node from 1 to 10, with the fluxes at a given time-point normalised by the maximum value.

($\alpha = \frac{1}{k+1}$) and unrelaxed ($\alpha = 1$) schedules are presented alongside these results. SI schemes using only a single iteration do not perform any relaxation – for these cases, the results from the unrelaxed simulation are presented.

The first of these approaches to be discussed is the CE/LI scheme without substeps, with the flux profiles presented in Fig. 3. While relaxation provides improved results in all cases considered, where the relaxation factor is between 0.3 and 0.4 stable results are achieved, even for a single corrector iteration. Predictably, given that the SI scheme begins with only a mild relaxation, more iterations are necessary to obtain a stable solution. While asymmetry is ultimately damped over iterations when using a relaxation factor of 0.5 and 0.6, the convergence towards a stable solution is relatively poor, especially for 0.6.

It is also worthwhile examining some measure of solution accuracy as, in principle, a small relaxation factor may possess stability at the price of slow convergence to the correct solution. This will be investigated using a difference in eigenvalues compared to a reference solution, as well as the L2 difference of U-235, U-238, Pu-239, Pu-240, Xe-135, and Sm-149 densities across the geometry compared to the chosen reference – each value is normalised to the sum of densities across the reference. Given that the 0.3 relaxation factor with 3 iterations is stable and will have minimised the influence of the predictor solution with multiple iterations, it is chosen as the reference here. For the eigenvalue, two graphs will be presented: one on the effect of increasing the number of iterations when using a relaxation factor of 0.3, the other on the effect of using various relaxation factors alongside 3 corrector iterations. These are shown in Fig. 4. The nuclide density results are likewise presented in two sets, corresponding to fixing the relaxation factor and varying the number of iterations, and vice versa. These are shown in Figs. 5 and 6, respectively.





Fig. 4. Differences in the eigenvalue with burn-up for the relaxed CE/LI scheme without substeps when varying the number of iterations (left) and relaxation factor (right)


The reported standard deviation in the estimate for the eigenvalue in all cases was no greater than 5 pcm, and was often closer to 3.5 pcm. Bearing this error in mind, when maintaining a relaxation factor of 0.3 and varying the number of iterations, the maximum difference from the eigenvalue outside of error was 10 pcm when using 2 iterations, whereas when varying the relaxation factor with 3 iterations the maximum observed bias was 4 pcm from the SI-CE/LI scheme. Hence, in all of these quite stable cases, only a very small bias was observed and only towards the end of the burn-up schedule. Interestingly, even the quite noticeably asymmetric simulation using a relaxation factor of 0.6 for 3 iterations showed relatively good agreement in the eigenvalue – others have made the point that instability is difficult to detect simply through scrutinising the eigenvalue (Dufek et al., 2013b; Kotlyar and Shwageraus, 2013). That being said, without using any form of relaxation or iteration, the maximum error in the eigenvalue was 900 pcm, so this difference can become quite dramatic for blatantly unstable simulations.

Regarding the nuclide density distribution errors, Fig. 5 shows that there is relatively little difference in all nuclide concentrations arising from iterating the corrector step, given that each of the solutions shown is apparently stable. Additional iterations better converge the nuclide densities towards the reference 3 iteration solution. Although the errors in the uranium densities grow over time, they are small in absolute value. Otherwise, for the other isotopes considered, the error after substantial burn-up is dwarfed by the error as fission products and actinides initially accumulate, induced by taking a relatively long initial time-step. The results are similar when varying the relaxation factor and using 3 iterations in Fig. 6: most nuclides' errors are dominated by that on the first step while uranium error – although small – accumulates with burn-up. There is actually little to suggest from these plots of nuclide densities that the scheme using a relaxation factor of 0.6 has manifested a non-physical tilt. In fact, both the Xe-135 and U-235 distributions become asymmetric, but only mildly so.


Before showing additional flux profiles with time, it should be emphasised that, for stable




(a) Uranium-235




(b) Uranium-238




(c) Plutonium-239



(d) Plutonium-240




(e) Xenon-135




(f) Samarium-149


Fig. 5. Normalised L2 nuclide density differences across the PWR pin geometry using the CE/LI scheme without substeps while varying the number of iterations




(a) Uranium-235




(b) Uranium-238




(c) Plutonium-239



(d) Plutonium-240



(e) Xenon-135



(f) Samarium-149

Fig. 6. Normalised L2 nuclide density differences across the PWR pin geometry using the CE/LI scheme without substeps while varying the relaxation factor

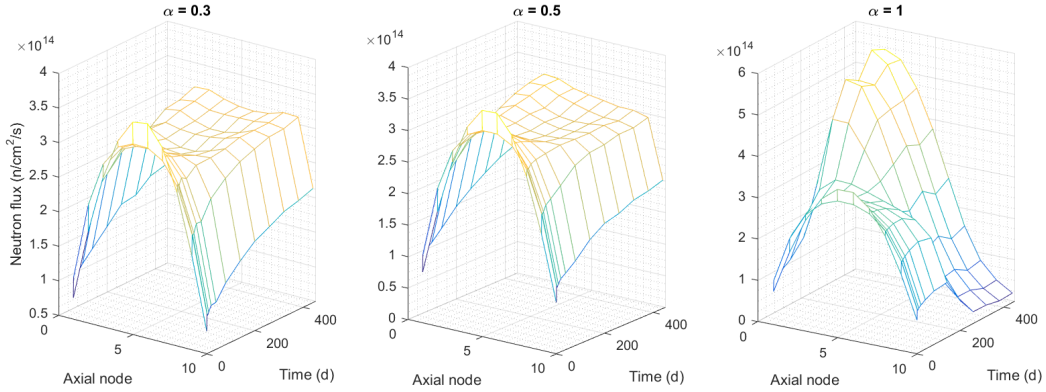


Fig. 7. Flux profiles over time for a PWR fuel pin with axially vacuum boundaries using the CE/LI scheme with no substeps, one corrector iteration and a variable relaxation factor.

simulations, towards the end of the burn-up schedule, the flux tends to become quite flat, albeit with a slight trough in the middle due to the depletion of central fissile material. Given the normalisation, this can exaggerate the effect of the stochastic noise in stable flux solutions. To highlight the differences across these solutions, the time-dependent flux profiles for the above results are plotted without normalisation, in particular, for 1 corrector iteration and with relaxation factors of 0.3, 0.5, and 1. These are shown in Fig. 7. From these results one can see that, in spite of slight noise in the flux profiles in Fig. 3, the $\alpha = 0.3$ solution is quite flat and symmetric, the $\alpha = 0.5$ solution has developed a slight but noticeable tilt, and the $\alpha = 1$ solution is wholly incorrect.

For the CE/LI scheme using 10 substeps, the flux heatmaps are shown in Fig. 8, while the eigenvalue errors relative to using 3 corrector iterations and a relaxation factor of 0.3 are shown in Fig. 9. The nuclide density differences are shown in Figs. 10 and 11. Adding substeps to the corrector appears to have worsened stability, given that a more severe relaxation with additional iterations appears necessary for symmetry – even with a relaxation factor of 0.4, there is still noticeably oscillatory behaviour occurring when applying 2 corrector iterations. That said, a relaxation factor of 0.3 appears adequate to achieve stability, regardless of the number of corrector iterations. Furthermore, even with 3 iterations, the SI scheme still produces a noticeably asymmetric solution.

While varying the number of iterations, the change in the eigenvalue accuracy appears minor – at worst, there was a 3 pcm bias outside of uncertainty when comparing 1 iteration to 3 with a relaxation factor of 0.3. On the other hand, varying the relaxation factor while keeping the number of corrector iterations constant unsurprisingly resulted in much larger errors, given the tilted flux solutions for the less stable relaxation factors. However, even though the SI relaxation schedule results in a flux tilt, it is sufficiently minor as to not significantly affect the eigenvalue, which remained within statistical uncertainty of the reference.

The nuclide densities when fixing the relaxation factor behave much as before, given that each solution with a relaxation factor of 0.3 is apparently stable. When varying the




Fig. 8. Flux profile over time for a PWR fuel pin with axially vacuum boundaries using the CE/LI scheme with 10 substeps and varying the number of corrector iterations and relaxation factor. The horizontal axes on each graph are the burn-up time-points given in Table 1, while the vertical axes are the flux values at a given axial node from 1 to 10, with the fluxes at a given time-point normalised by the maximum value.

relaxation scheme with 3 iterations, however, the three settings for which an asymmetric solution was obtained (0.5, 0.6, and SI) have their nuclide densities diverge substantially from the reference solution. In some cases this difference is more than an order of magnitude greater than when using a relaxation factor of 0.3 and only using a single corrector iteration. Unsurprisingly, given how tightly it couples neutronics and depletion, xenon appears to be the most disrupted isotope, reaching a substantial difference from the reference solution for each of the asymmetric schemes. That being said, as xenon diverges, all uranium and plutonium isotopes simultaneously show a significant divergence (especially noticeable for the $\alpha = 0.6$ case), albeit, of a smaller magnitude. This rapid growth in error, observed here and elsewhere in this paper, is characteristic of numerical instability, i.e., the amplification of a non-physical (and, in this case, asymmetric) solution by the numerical scheme applied, as discussed by Densmore et al. (2013). The changes in the eigenvalue errors in Fig. 9 appear to track the nuclide density errors.

Finally, the flux profiles for the LE/LI scheme with 10 substeps are presented in Fig. 12, while the eigenvalue errors are presented in Fig. 13. The nuclide density differences are shown in Figs. 14 and 15. This scheme shows a marked tendency towards instability, presenting some relatively strange instability behaviour with a single corrector iteration, even when aggressively relaxing the corrector step. In particular, the one-sided tilt is no longer observed, replaced either by a more gentle ‘sloshing’ of the flux over time (for $\alpha = 0.5, 0.6$) or an unpredictable peak flux movement (for $\alpha = 0.3$). More familiar instability behaviour is recovered when using more corrector iterations, though the only stable settings found in these investigations use a relaxation of 0.3 for 2 or 3 iterations, or a relaxation of 0.4 for 3 iterations. Again, the SI relaxation schedule returns noticeably asymmetric results, even





Fig. 9. Differences in the eigenvalue with burn-up for the relaxed CE/LI scheme with 10 substeps when varying the number of iterations (left) and relaxation factor (right)

with 3 corrector iterations.


Fixing the relaxation factor at 0.3 and varying the number of corrector iterations, the single iteration is unsurprisingly erroneous, given an abnormal flux profile. Otherwise, for 2 corrector iterations the maximum bias is quite small. Using 3 corrector iterations with variable relaxation factors also produces predictable results given the extent of instability observed – unless using a relaxation factor of 0.4 or less, all simulations become appreciably biased in the eigenvalue.

Similar results are observed for the nuclide densities: fixing the relaxation factor at 0.3, the single iteration solution deviates substantially and nearly immediately, whereas performing a second iteration results in significantly lowered error in the nuclide densities due to maintaining stability. The rapidity with which the single iteration solution becomes erroneous suggests that perhaps a different instability phenomenon is substantially present: it may be that the small relaxation factor applied with only a single iteration fails to sufficiently dampen non-physical behaviour driven by the LE predictor. When varying the relaxation factor with 3 corrector iterations, the nuclide density behaviour is broadly the same as it was for the CE/LI scheme with 10 substeps: schemes employing a weaker relaxation show a greater and more rapid divergence in their nuclide densities from the reference solution, from which their eigenvalue error in Fig. 13 follows. A notable difference is that the growth in uranium and xenon density errors (and plutonium, for the $\alpha = 0.6$ case) occurs one time-point earlier than when using the CE/LI with 10 substeps: the errors appear to jump following the first 60-day time-step, rather than the second. This may be another indicator of the relative instability of the LE predictor. One point worth noting is the relatively large errors in both the plutonium isotopes at the beginning of the simulation: these both occur momentarily during the third time-point. They almost certainly follow from the first application of the LE predictor taken over a relatively long initial time-step when fission and activation products are rapidly evolving.


Ultimately, in all cases examined, fixing the relaxation factor at a value of about 0.3




(a) Uranium-235




(b) Uranium-238




(c) Plutonium-239



(d) Plutonium-240




(e) Xenon-135




(f) Samarium-149


Fig. 10. Normalised L2 nuclide density differences across the PWR pin geometry using the CE/LI scheme with 10 substeps while varying the number of iterations




(a) Uranium-235




(b) Uranium-238




(c) Plutonium-239



(d) Plutonium-240



(e) Xenon-135



(f) Samarium-149

Fig. 11. Normalised L2 nuclide density differences across the PWR pin geometry using the CE/LI scheme with 10 substeps while varying the relaxation factor




Fig. 12. Flux profile over time for a PWR fuel pin with axially vacuum boundaries using the LE/LI scheme with 10 substeps and varying the number of corrector iterations and relaxation factor. The horizontal axes on each graph are the burn-up time-points given in Table 1, while the vertical axes are the flux values at a given axial node from 1 to 10, with the fluxes at a given time-point normalised by the maximum value.







Fig. 13. Differences in the eigenvalue with burn-up for the relaxed LE/LI scheme with 10 substeps when varying the number of iterations (left) and relaxation factor (right).




(a) Uranium-235




(b) Uranium-238




(c) Plutonium-239



(d) Plutonium-240




(e) Xenon-135




(f) Samarium-149


Fig. 14. Normalised L2 nuclide density differences across the PWR pin geometry using the LE/LI scheme with 10 substeps while varying the number of iterations




(a) Uranium-235




(b) Uranium-238




(c) Plutonium-239



(d) Plutonium-240



(e) Xenon-135



(f) Samarium-149

Fig. 15. Normalised L2 nuclide density differences across the PWR pin geometry using the LE/LI scheme with 10 substeps while varying the relaxation factor

or 0.4 gives equal or superior results to using a relaxation schedule given by the stochastic approximation. This is true with respect to both the stability of the flux profile and the accuracy of the eigenvalue and the nuclide density distribution.

That being said, it is worth performing some further calculations to investigate the sensitivity of this scheme. Serpent allows its implementation of the SIE to use lower statistics during the corrector step calculations, seemingly without significant adverse effect on stability, provided enough iterations are performed. This reduces the computational expense of the scheme significantly as compared to performing many full-statistics transport simulations for each time-step. Before recommending the use of a fixed relaxation factor it would be prudent to verify that it can achieve similar advantages in terms of computational efficiency.

Hence, given the LE/LI scheme with 10 substeps was most challenging in terms of stability, the calculation using a relaxation factor of 0.3 was repeated, as well as the calculation using an SI schedule, each with 3 corrector iterations. However, during the corrector step, each used less than one third of the active particles applied to the predictor, reducing the number of active cycles from 80 to 26. The number of inactive cycles and response matrix acceleration remained the same, as the MATLAB wrapping script does not allow the fission source to be passed from one transport solution to the next, as can be done with Serpent's internal depletion coupling. These results are shown in Fig. 16. Both schemes show nearly unchanged results from their full-statistics counterparts, with the fixed relaxation factor providing a relatively symmetric flux profile while the stochastic approximation has failed to converge sufficiently. This is emphasised by considering the eigenvalue error against the full-statistics version of the $\alpha = 0.3$ simulation, as shown in Fig. 17 – intriguingly, the SI scheme with reduced statistics appears to have a better accuracy than its full-statistics counterpart (Fig. 13), but fails to approach the reference solution. On the other hand, reducing the corrector statistics does not appear to have affected the eigenvalue estimate when using a fixed relaxation factor.

The nuclide density errors against the reference scheme are also shown in Fig. 18: these match the behaviour of the eigenvalue well, showing the sharp error growth in a number of nuclide densities for the SI schedule after 200 days when the eigenvalue error grows outside the range of statistical error. The scheme using a relaxation factor of 0.3 displays only a small, consistent error across all considered nuclides at all times. That said, for both plutonium isotopes and the two fission products, the SI scheme shows the initial, often dominant, error familiar from the previous nuclide density charts, while the error in the fixed relaxation scheme is substantially smaller. It should be emphasised that this is not thought to be any particular advantage of the fixed relaxation scheme or due to numerical instability; it is much more likely a product of stochastic noise, extremely low initial nuclide densities, the LE predictor, and a small difference in weighting between the extrapolated and corrected reaction rates corresponding to each of the schemes applied.

A final test to conduct should investigate the effect of clustering when using a fixed relaxation factor. This work was inspired by suspecting that the success of SI schemes was partly due to their ability to negate clustering and correlation effects, begging the question – will a fixed relaxation factor prove similarly robust? This can be investigated




Fig. 16. Flux profile over time for a PWR fuel pin with axially vacuum boundaries using the LE/LI scheme with 10 substeps, three corrector iterations, and varying the relaxation scheme, while dividing the number of active histories during the predictor equally across corrector iterations. The horizontal axes on each graph are the burn-up time-points given in Table 1, while the vertical axes are the flux values at a given axial node from 1 to 10, with the fluxes at a given time-point normalised by the maximum value.







Fig. 17. Differences in the eigenvalue with burn-up for the relaxed LE/LI scheme with 10 substeps, three iterations and two relaxation schemes which use the same number of active histories across all corrector transport calculations as during the predictor step.




(a) Uranium-235




(b) Uranium-238




(c) Plutonium-239



(d) Plutonium-240



(e) Xenon-135



(f) Samarium-149

Fig. 18. Normalised L2 nuclide density differences across the PWR pin geometry using the LE/LI scheme with 10 substeps, three iterations and two relaxation schemes which use the same number of active histories across all corrector transport calculations as during the predictor step




Fig. 19. Flux profile over time for a PWR fuel pin with axially vacuum boundaries using the LE/LI scheme with 10 substeps, three corrector iterations and two relaxation schemes, when using fewer particles per cycle while maintaining the total number of active histories. The horizontal axes on each graph are the burn-up time-points given in Table 1, while the vertical axes are the flux values at a given axial node from 1 to 10, with the fluxes at a given time-point normalised by the maximum value.

by preserving the total number of active histories applied to a transport solution while lowering the number of particles per cycle – again, the LE/LI scheme was used with 3 corrector iterations, comparing a fixed relaxation factor of 0.3 with the SI scheme. The particle settings were 200,000 particles per cycle over 800 active cycles, and 200 inactive cycles. Unlike previously, the response matrix acceleration scheme was not used, as it was reported in the original paper that lower numbers of particles per cycle could adversely affect the algorithm (Leppänen, 2019). The number of active histories was the same for all transport solutions obtained. The flux profiles are shown in Fig. 19 and their eigenvalues and nuclide densities are compared against the same reference case as above when using the LE/LI scheme, shown, respectively, in Figs. 20 and 21.

From the heat maps, both schemes have been noticeably affected by using fewer particles per cycle, displaying significantly more asymmetry as compared to previous results. That being said, the asymmetry is less pronounced when using a fixed relaxation factor – the stochastic approximation actually does not appear more effective in attaining a stable solution where clustering and correlation is more dominant. This is emphasised further by considering the eigenvalue error: in spite of the asymmetry in the flux profile, there is only a slight bias outside of statistics that can be observed when using the fixed relaxation factor. The SI scheme, on the other hand, obtains a markedly different eigenvalue, with an error that is larger and more sustained than occurred in the preceding simulations where clustering would not have been so dominant. Examining the nuclide density errors, although the magnitudes of the errors have increased for the fixed relaxation scheme, they remain relatively constant with burn-up. On the other hand, the SI scheme displays the previously seen sharp error growth across a number of nuclides after 150 days, indicative of instability. Hence, SI schemes appear to be less robust than simply using a fixed relaxation factor in simulations subject to clustering and correlation.





Fig. 20. Differences in the eigenvalue with burn-up for the relaxed LE/LI scheme with 10 substeps, three iterations and two relaxation schemes, when using fewer particles per cycle while maintaining the total number of active histories.


4. Conclusions

This work has investigated the stability of Monte Carlo burn-up schemes for neutronic large geometries. In particular, a number of the standard PC algorithms were investigated and found to be unstable without modification to the corrector step, even when accounting for neutron clustering. A new, implicit burn-up algorithm was then proposed, closely resembling stochastic implicit schemes, albeit, simplified by keeping the relaxation factor fixed. This fixed relaxation factor scheme was examined on a representative unstable burn-up problem for a variety of PC schemes, relaxation factors, and corrector iterations, with comparison against the stochastic implicit equivalent. In all cases, using a fixed relaxation factor proved to be either equally or more efficient and accurate while achieving a stable solution. These investigations also demonstrate that adding substeps to the corrector step and using a linear extrapolation for the predictor both make the problem more unstable. To demonstrate the scheme’s computational efficiency, it was shown that using reduced statistics on each corrector step, such that the expense of repeated transport solutions was not substantially larger than that during the predictor step, does not degrade the accuracy or stability of the simpler scheme. Finally, it was shown that the scheme has a superior performance to the stochastic approximation in the presence of clustering and correlation effects.


Future work will investigate the selection of an optimal relaxation factor – an effective value was not obvious from the outset, only becoming apparent by examining simulation results. This value will vary with both the time-step size considered and the problem geometry. Shorter time-steps would seemingly permit gentler relaxation and fewer iterations,




(a) Uranium-235




(b) Uranium-238




(c) Plutonium-239



(d) Plutonium-240



(e) Xenon-135



(f) Samarium-149

Fig. 21. Normalised L2 nuclide density differences across the PWR pin geometry using the LE/LI scheme with 10 substeps, three iterations and two relaxation schemes when using fewer particles per cycle while maintaining the total number of active histories

whereas more spatially-decoupled problems would imaginably require more aggressive relaxation. That being said, symmetric PWR pin problems taking 60 day time-steps are some of the more unstable in the authors' experience, and a relaxation factor of 0.3 with 3 corrector iterations, dividing active particles equally between each corrector step seems adequate to obtain a stable and accurate solution. This suggests that, provided a large number of particles are simulated per generation, other challenging problems are not likely to require substantially more severe settings.

Finally, given the success of relaxation methods, it may also be possible to obtain stable solutions to the non-linear system given by the corrector step using Newton's method, potentially by a similar means to that demonstrated by Aufiero and Fratoni (2017). On the one hand, this will require the additional computational expense of computing derivative information – presently this appears relatively computationally inefficient as compared to the currently proposed scheme, which allows the corrector statistics to be reduced while maintaining stability. However, when coupling neutron transport and isotopic depletion with additional physics, there may be some computational advantage to be gained by obviating iteration between transport, depletion, and, say, thermal-hydraulics.

5. Data availability statement

To the best of the authors' knowledge, this paper and references herein contain all the data needed to reproduce and validate the results presented.

Acknowledgements

This work was carried out with the financial support of the UK Engineering and Physical Sciences Research Council (EPSRC) through the Imperial College, University of Cambridge and Open University Centre for Doctoral Training in Nuclear Energy under grant EP/L015900/1.

Declarations

Declarations of interest: none.

References

- Aufiero, M., Fratoni, M., 2017. A new approach to the stabilization and convergence acceleration in coupled Monte Carlo–CFD calculations: The Newton method via Monte Carlo perturbation theory. *Nuclear Engineering and Technology* 49, 1181–1188. doi:10.1016/j.net.2017.08.005.
- Bell, G., Glasstone, S., 1970. *Nuclear Reactor Theory*. US Atomic Energy Commission, Washington, DC.
- Cosgrove, P., Shwageraus, E., Parks, G., 2019. Neutron clustering as a driver of Monte Carlo burn-up instability. *Annals of Nuclear Energy* , 106991doi:10.1016/j.anucene.2019.106991.
- Densmore, J.D., Gill, D.F., Griesheimer, D.P., 2013. Stability analysis of burnup calculations. *Transactions of the American Nuclear Society* 109, 695–698. doi:10.1007/s10967-012-2210-3.2.
- Dufek, J., Hoogenboom, J.E., 2009. Numerical Stability of Existing Monte Carlo Burnup Codes in Cycle Calculations of Critical Reactors. *Nuclear Science and Engineering* 162, 307–311. doi:10.13182/NSE08-69TN.

- Dufek, J., Kotlyar, D., Shwageraus, E., 2013a. The stochastic implicit Euler method: A stable coupling scheme for Monte Carlo burnup calculations. *Annals of Nuclear Energy* 60, 295–300. doi:10.1016/j.anucene.2013.05.015.
- Dufek, J., Kotlyar, D., Shwageraus, E., Leppänen, J., 2013b. Numerical stability of the predictor-corrector method in Monte Carlo burnup calculations of critical reactors. *Annals of Nuclear Energy* 56, 34–38. doi:10.1016/j.anucene.2013.01.018.
- Dumonteil, E., Malvagi, F., Zoia, A., Mazzolo, A., Artusio, D., Dieudonné, C., De Mulatier, C., 2014. Particle clustering in Monte Carlo criticality simulations. *Annals of Nuclear Energy* 63, 612–618. doi:10.1016/j.anucene.2013.09.008.
- Gill, D.F., Griesheimer, D.P., Aumiller, D.L., 2017. Numerical Methods in Coupled Monte Carlo and Thermal-Hydraulic Calculations. *Nuclear Science and Engineering* 185, 194–205. doi:10.13182/NSE16-3.
- Griesheimer, D., Gill, D., Nease, B., Sutton, T., Stedry, M., Dobreff, P., Carpenter, D., Trumbull, T., Caro, E., Joo, H., Millman, D., 2015. MC21 v.6.0 – A continuous-energy Monte Carlo particle transport code with integrated reactor feedback capabilities. *Annals of Nuclear Energy* 82, 29–40. doi:10.1016/j.anucene.2014.08.020.
- Griesheimer, D.P., 2010. In-line xenon convergence algorithm for Monte Carlo reactor calculations, in: Proc. PHYSOR 2010, Pittsburgh, Pennsylvania.
- Griesheimer, D.P., Carpenter, D.C., Stedry, M.H., 2017. Practical techniques for large-scale Monte Carlo reactor depletion calculations[sic]. *Progress in Nuclear Energy* 101, 409–423. doi:10.1016/j.pnucene.2017.05.018.
- Herman, B.R., 2014. Monte Carlo and Thermal Hydraulic Coupling using Low-Order Nonlinear Diffusion Acceleration. Ph.D. thesis. Massachusetts Institute of Technology.
- Isotalo, A., 2013. Computational Methods for Burnup Calculations with Monte Carlo Neutronics. Ph.D. thesis. Aalto University.
- Isotalo, A., Aarnio, P., 2011a. Higher order methods for burnup calculations with Bateman solutions. *Annals of Nuclear Energy* 38, 1987–1995. doi:10.1016/j.anucene.2011.04.022.
- Isotalo, A., Aarnio, P., 2011b. Substep methods for burnup calculations with Bateman solutions. *Annals of Nuclear Energy* 38, 2509–2514. doi:10.1016/j.anucene.2011.07.012.
- Isotalo, A., Leppänen, J., Dufek, J., 2013. Preventing xenon oscillations in Monte Carlo burnup calculations by enforcing equilibrium xenon distribution. *Annals of Nuclear Energy* 60, 78–85. doi:10.1016/j.anucene.2013.04.031.
- Josey, C., 2017. Development and Analysis of High Order Neutron Transport-Depletion Coupling Algorithms. Ph.D. thesis. Massachusetts Institute of Technology.
- Josey, C., Forget, B., Smith, K., 2017. High order methods for the integration of the Bateman equations and other problems of the form of $y' = F(y,t)y$. *Journal of Computational Physics* 350, 296–313. doi:10.1016/j.jcp.2017.08.025.
- Kepisty, G., Cetnar, J., 2015. Instabilities of Monte-Carlo burnup calculations for nuclear reactors—Demonstration and dependence from time step model. *Nuclear Engineering and Design* 286, 49–59. doi:10.1016/j.nucengdes.2015.01.023.
- Kotlyar, D., Shwageraus, E., 2013. On the use of predictor-corrector method for coupled Monte Carlo burnup codes. *Annals of Nuclear Energy* 58, 228–237. doi:10.1016/j.anucene.2013.03.034.
- Kotlyar, D., Shwageraus, E., 2014. Numerically stable Monte Carlo-burnup-thermal hydraulic coupling schemes. *Annals of Nuclear Energy* 63, 371–381. doi:10.1016/j.anucene.2013.08.016.
- Kotlyar, D., Shwageraus, E., 2016. Stochastic semi-implicit substep method for coupled depletion Monte-Carlo codes. *Annals of Nuclear Energy* 92, 52–60. doi:10.1016/j.anucene.2016.01.022.
- Leppänen, J., 2019. Acceleration of fission source convergence in the Serpent 2 Monte Carlo code using a response matrix based solution for the initial source distribution. *Annals of Nuclear Energy* 128, 63–68. doi:10.1016/j.anucene.2018.12.044.
- Leppänen, J., Pusa, M., Viitanen, T., Valtavirta, V., Kältiaisenaho, T., 2015. The Serpent Monte Carlo code: Status, development and applications in 2013. *Annals of Nuclear Energy* 82, 142–150. doi:10.1016/j.anucene.2014.08.024.

- de Mulatier, C., Dumonteil, E., Rosso, A., Zoia, A., 2015. The critical catastrophe revisited. *Journal of Statistical Mechanics: Theory and Experiment* 2015, P08021. doi:10.1088/1742-5468/2015/08/P08021.
- Nowak, M., Miao, J., Dumonteil, E., Forget, B., Onillon, A., Smith, K.S., Zoia, A., 2016. Monte Carlo power iteration: Entropy and spatial correlations. *Annals of Nuclear Energy* 94, 856–868. doi:10.1016/j.anucene.2016.05.002.
- Pusa, M., 2013. Numerical Methods for Nuclear Fuel Burnup Calculations. Ph.D. thesis. VTT Technical Research Centre of Finland.
- Robbins, H., Monro, S., 1951. A Stochastic Approximation Method. *The Annals of Mathematical Statistics* 22, 400–407. doi:10.1214/aoms/1177729586.
- Sutton, T.M., Mittal, A., 2017. Neutron clustering in Monte Carlo iterated-source calculations. *Nuclear Engineering and Technology* 49, 1211–1218. doi:10.1016/j.net.2017.07.008.
- Tuominen, R., 2015. Coupling Serpent and OpenFOAM for neutronics - CFD multi-physics calculations. Master's thesis. Aalto University.
- Zoia, A., Dumonteil, E., Mazzolo, A., de Mulatier, C., Rosso, A., 2014. Clustering of branching Brownian motions in confined geometries. *Physical Review E* 90, 042118. doi:10.1103/PhysRevE.90.042118.