

Burnup calculation methodology in Serpent

Updated March 27, 2012 / Jaakko Leppänen

The purpose of these notes is to:

- Explain with examples the burnup calculation algorithms used in the Serpent code
- Assist code users in the selection of the appropriate calculation scheme for their burnup problems
- Provide support material for lectures on burnup calculation methodology

For questions, contact: Jaakko.Leppanen@vtt.fi

Serpent website: <http://montecarlo.vtt.fi>

Serpent discussion forum: <http://ttuki.vtt.fi/serpent>

Solution of the Bateman depletion equations

- The Serpent code uses built-in burnup calculation routines to simulate the isotopic changes in materials due to neutron-induced reactions and spontaneous radioactive decay. These changes are characterized by the Bateman depletion equations, written as

$$\frac{dN_j}{dt} = \sum_{i \neq j} S_{i \rightarrow j} - \lambda_j N_j - \phi \sigma_j N_j$$

where:

- $S_{i \rightarrow j}$ is the production rate of nuclide j in decay, transmutation and fission reactions
 - $\lambda_j N_j$ is the radioactive decay rate of nuclide j
 - $\phi \sigma_j N_j$ is the transmutation rate of nuclide j , including fission
- The one-group fission and transmutation cross sections are obtained from the Monte Carlo neutron transport simulation, as flux-volume weighted averages of the corresponding continuous-energy microscopic cross sections, and assumed constant over the time interval where the equations are solved. Radioactive decay and fission yield data is read from ENDF format data files.
 - Serpent has two options for solving the depletion equations – the linear chains method, and the CRAM matrix exponential method. The second option is used by default, and it is considered a superior solution to burnup calculation problems involving the full system of nuclides. In Serpent calculations the burnup matrix typically consists of 1200-1700 nuclides, with fission and transmutation cross sections calculated for 200-350 actinides and fission products. The CRAM method was introduced by VTT, specifically for the Serpent code, and a complete description of the theory and practical implementation can be found in three references:
 - [1] M. Pusa and J. Leppänen. "Computing the Matrix Exponential in Burnup Calculations." Nucl. Sci. Eng., **164** (2010) 140-150.
 - [2] M. Pusa. "Rational Approximations to the Matrix Exponential in Burnup Calculations." Nucl. Sci. Eng., **169** (2011) 155-167.
 - [3] M. Pusa and J. Leppänen. "An Efficient Implementation of the Chebyshev Rational Approximation." In. Proc. PHYSOR 2012.
 - The methods used by Serpent 1 and 2 are largely the same.

Burnup algorithm

- The Bateman solution to the depletion equations necessitates that the transmutation cross sections remain constant in time. This is not the case in reality, since changing nuclide compositions affect the level of self-shielding experienced by the materials under irradiation. To overcome the problem, the burnup interval is divided into a number of steps, during which the cross sections are assumed constant. Since the time-consuming transport calculation has to be repeated for each step, finding the optimal step lengths is an important part of the solution.
- The changes in the cross sections within each depletion step can be accounted for to some extent by using constant values that represent the average composition better than the values at the beginning of the step. These values are calculated by assuming some functional behavior for the cross sections, typically using so-called predictor-corrector methods.
- The predictor-corrector methods developed for Serpent are described in the following slides. The methods are based on extrapolated reaction rates for the predictor and interpolated reaction rates for the corrector calculation. The options and the associated abbreviations are:
 - Constant extrapolation for the predictor (CE)
 - Linear extrapolation for the predictor (LE)
 - Linear interpolation for the corrector (LI)
 - Quadratic interpolation for the corrector (QI)

with the following 5 pre-defined combinations: CE, LE, CE/LI, LE/LI, LE/QI.

- In addition to the extrapolated and interpolated reaction rates, Serpent offers the possibility to divide the step-wise solutions into multiple sub-steps, which approximates the continuously changing flux and cross sections closer than a single constant averaged value.
- Complete description of the higher-order predictor-corrector and sub-step methods can be found in references:

[4] A. Isotalo and P. Aarnio. "Higher order methods for burnup calculations with Bateman solutions." Ann. Nucl. Energy, **38** (2011) 1987-1995.

[5] A. Isotalo and P. Aarnio. "Substep methods for burnup calculations with Bateman solutions." Ann. Nucl. Energy, **38** (2011) 2509-2514.

- The depletion routine in Serpent 1 is based on the "conventional" predictor-corrector algorithm (CE/LI). The higher-order approximations and sub-step methods are implemented only in Serpent 2. Some practical considerations with respect to step length and computing time are found in reference:

[6] J. Leppänen and A. Isotalo. "Burnup Calculation Methodology in the Serpent 2 Monte Carlo Code." In. Proc. PHYSOR 2012.

Burnup algorithm

- The selection of step length is left to the user, and it should account for at least the following physical factors:
 - Reactor poison Xe-135 usually saturates within 48 hours from reactor start-up. During this period, flux and reactivity are in continuous change, which calls for shorter steps at the beginning of the irradiation cycle. The same applies to any changes in reactor power level.
 - Gadolinium, used as a burnable absorber in light water reactor fuels, is depleted by 10-15 MWd/kgU burnup. Shorter steps should be used until the absorber is completely gone. Since burnable absorber pins are subject to strong spatial self-shielding, the material should be divided into multiple sub-regions. Other burnable absorbers (boron, europium, etc.) behave similarly, but may be depleted at a different rate.
 - Reactivity changes are caused by the depletion of the primary fissile isotope (U-235) and burnable absorbers, together with the build-up of plutonium and fission products. These changes do not reflect the accumulation of weak absorbers, such as I-131 or Cs-137, which may be important for some other reason (source terms for accident analysis, etc.). The selection of the appropriate step lengths should therefore not be based on the multiplication factor alone.

The same physical mechanisms apply to other reactor and fuel types as well (fast reactors, thorium fuel), but their importance may differ from conventional LWR fuels. In addition to the physical factors, there are considerations related to the computational methods, which will be discussed in the following slides.

- It is important to realize that, even though flux and cross sections are assumed to exhibit a continuous behavior over each burnup interval $[t_i, t_{i+1}]$, the formulation of the depletion equations necessitates the use of constant values. The solution is therefore subject to two approximations:
 1. The behavior is approximated by a continuous analytical function $f(t)$, which is formed using discrete values calculated for one or several burnup steps.
 2. The continuous behavior is further approximated by a constant value, defined by averaging the continuous function over the interval:

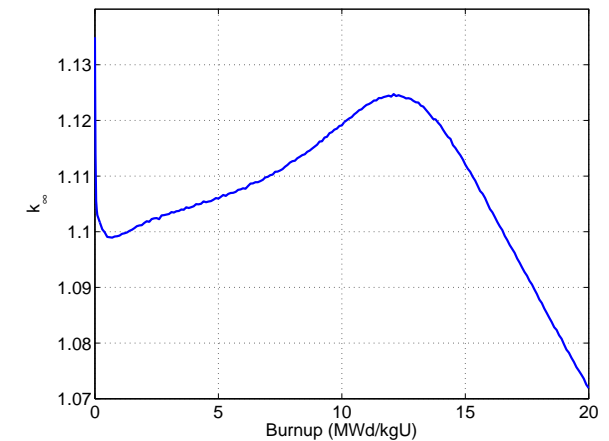

$$\frac{1}{t_{i+1} - t_i} \int_{t_i}^{t_{i+1}} f(t) dt .$$

Sub-step methods divide the integration into multiple parts, after which the depletion equations are solved. The functional dependence, however, remains the same from t_i to t_{i+1} , and the transport calculation is not run multiple times.

Test case

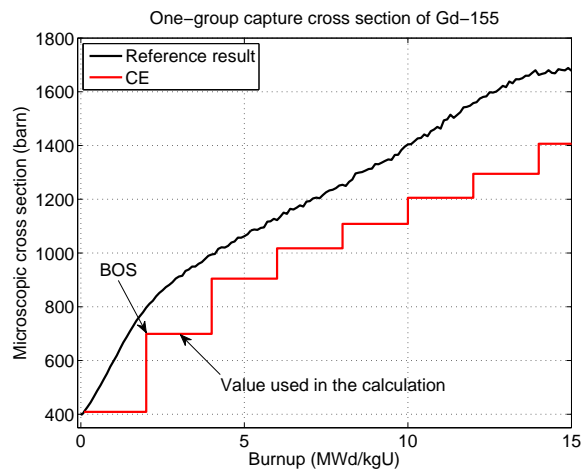

Geometry, materials and depletion history

- The test case used for demonstrating the predictor-corrector and sub-step methods represents a simplified PWR fuel lattice with a single burnable absorber pin surrounded by 15 uranium oxide pins. The uranium pins are treated as a single burnable material, while the absorber is divided into 10 annular regions with equal volume. The fuel is irradiated with a constant 40 kw/kgU power density to 20 MWd/kgU burnup. The figures on the right show the geometry and the infinite multiplication factor as function of burnup.
- The following slides show the depletion of Gd-155 in the outermost ring of the burnable absorber pin, using different burnup algorithms. To emphasize the differences, the step length is set to 2 MWd/kgU, which would normally be considered too long for this type of calculation. The first five steps are 0.01, 0.025, 0.05, 0.075 and 0.1 MWd/kgU, after which the burnup jumps to 2 MWd/kgU. The purpose of taking shorter steps at the beginning is to allow the saturation of short-lived nuclides, in particular Xe-135. The dramatic change in the step length, however, may have a significant impact in some of the algorithms, which is shown in the examples.
- The impact of dividing the burnup solution into multiple sub-steps is demonstrated with I-131, a weak absorber with a continuously-changing production rate and half-life considerably shorter than the step length.
- The reference results are calculated with a step length of 0.1 MWd/kgU throughout the entire cycle, after the shorter steps at the beginning. All calculations were run with 2000 active cycles of 10000 source neutrons. Cross sections, fission yields and radioactive decay data is based on the JEFF-3.1 evaluated nuclear data file. The concentrations of 1561 nuclides were traced for each burnable material, and fission and transmutation cross sections were calculate for 330 nuclides. The code version used in the calculations was Serpent 2.1.4 (beta).



Predictor-corrector methods

Constant extrapolation (CE)

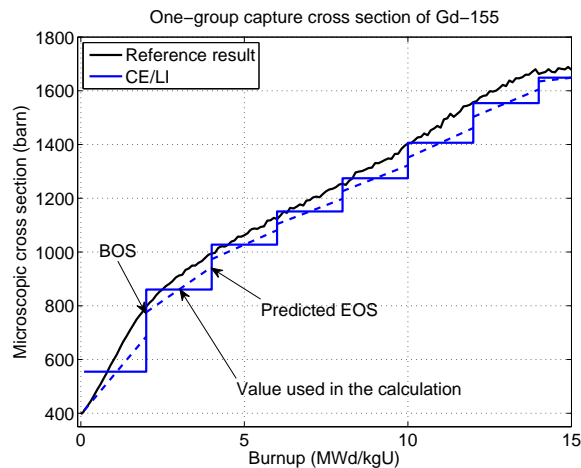



- Constant extrapolation, sometimes called the “Euler’s method”, is the simplest burnup algorithm, that assumes that the flux and cross sections remain at their beginning-of-step (BOS) values throughout the interval. The atomic density of Gd-155 in the example case is plotted in the top figure and the one-group capture cross section of the same isotope in the figure below. The cross section is a step-wise function, as the value is updated at the beginning of each burnup step.
- If the cross section is an increasing function of burnup, which is often the case with burnable absorbers due to reduced self-shielding, the BOS value under-estimates the average cross section over the interval. The result is that the depletion rate is also under-estimated, which, in turn, results in the over-prediction of the nuclide concentration. The same problem occurs in the next interval, and the algorithm simply cannot keep up with the actual depletion rate. The error can be reduced by shortening the step length, but this is often not the optimal solution to the problem.
- The fact that the cross section keeps increasing in the example case even after the concentration is reduced to zero results from self-shielding from the inner depletion regions, in which the absorber depletion doesn’t start until the first ring is sufficiently depleted to let thermal neutrons through.
- This is the default method used by most depletion codes when nothing specific is said about predictor-corrector calculation. The default depletion algorithm in Serpent is predictor-corrector with constant extrapolation on the predictor and linear interpolation on the corrector (CE/LI). This method is described in the next slide.
- The Serpent input option for invoking the CE method is:

```
set pcc 0
```

Serpent 2 also accepts “CE” in place of zero.

Constant extrapolation with linear interpolation (CE/LI)



- Predictor-corrector methods apply multiple transport calculations per depletion step, in order to more accurately represent the continuously changing reaction rate over the interval. The simplest method is based on linear interpolation on the corrector calculation, while the predictor uses constant BOS values. This method is often simply referred to as “predictor-corrector” calculation.

- The solution is divided in two parts:

1. Predictor calculation – flux and cross sections are calculated for the BOS composition.
2. Corrector calculation – the material is depleted over the interval, and new flux and cross sections are calculated for the EOS composition.

The final burnup calculation is carried out from BOS to EOS, using the average of the predictor and corrector flux and cross sections, which corresponds to linear interpolation between the two points.

- The figure shows the linearly changing cross section (dashed line) over each depletion step. The beginning and end points of each line are the predictor and corrector step values, respectively. The solid step-wise function shows the values used in the final calculation. The fact that the values are constant over each step is necessitated by the method used for solving the Bateman depletion equations.


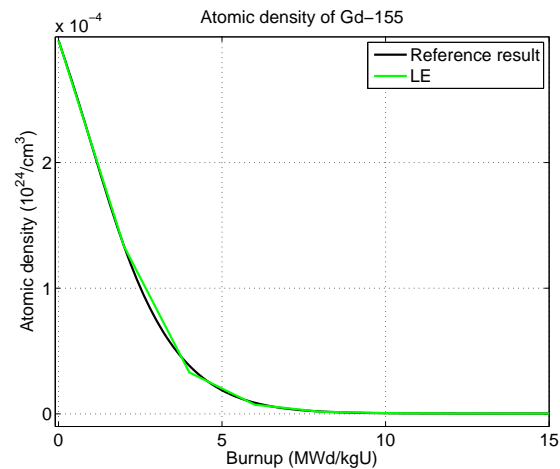
- The CE/LI method usually results in improved accuracy compared to CE, but since two transport solutions are required per depletion step, also the running time is increased. In most cases the method is still superior to CE with step length reduced in half.

- Serpent uses the CE/LI method by default, and it can also be set using input option:

```
set pcc 1
```

Serpent 2 also accepts “CELI” in place of one.

Linear extrapolation (LE)

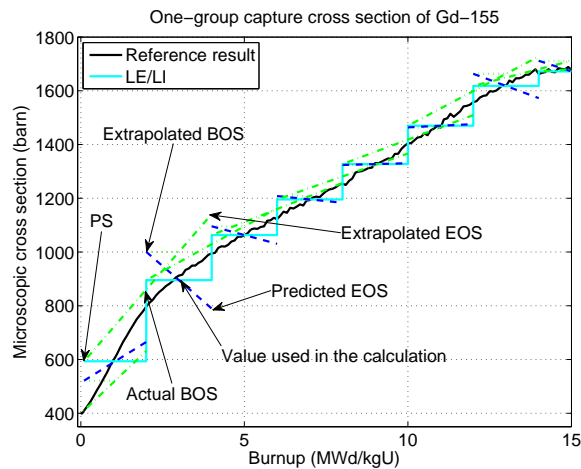



- Linear extrapolation is one of the new burnup algorithms developed for Serpent 2. Instead of performing an additional transport calculation for the EOS composition, the linear behavior of flux and cross sections is extrapolated over the interval using the BOS value and the result of the previous depletion step (PS). The first interval is handled using constant extrapolation, since there is no data from previous steps. Since there is no additional transport calculation involved, the increase in CPU time is negligible compared to constant extrapolation, but additional storage space is needed for keeping the previous values.
- The extrapolation using the two values is plotted with dash-dotted lines. As in linear interpolation, the actual value used in the calculation is constant over the depletion step, and calculated as the average of BOS and the extrapolated EOS value.
- The method has shown good results despite its simplicity, and considering the negligible increase in computing time, it can be considered as the most efficient way to improve the accuracy of burnup calculations. However, since the method is based on extrapolation instead of interpolation, it is typically more sensitive to statistical variation in the results, especially if the burnup interval between PS and BOS is short compared to the interval between BOS and extrapolated EOS.
- The LE method is available only in Serpent 2, and it is set using:

```
set pcc 2
```


The mode number can be replaced by “LE”.

Linear extrapolation with linear interpolation (LE/LI)





- The best result is usually obtained by combining linear extrapolation (LE) with a predictor-corrector method. The simplest solution is to use linear interpolation (LI), similar to the conventional “predictor-corrector calculation”, or CE/LI. The BOS values used for the predictor calculation are then given by the method described in the previous slide. The material is depleted, and a second transport calculation yields the flux and the cross sections corresponding to the predicted EOS composition. The final depletion is carried out using the average of the two.
- The figures on the left show that this method produces a very accurate representation of the true behavior of the cross section in the example case. The slight over-estimation produced by linear extrapolation on the predictor step (“Extrapolated BOS”) is effectively corrected by the second calculation (dashed blue line downwards to “Predicted EOS”), and the final depletion matches almost perfectly the reference result.
- The LE/LI method is available only in Serpent 2, and it is invoked using:

```
set pcc 3
```

 The mode number can be replaced by “LELI”.

Linear extrapolation with quadratic interpolation (LE/QI)



- In addition to linear interpolation, Serpent 2 offers the possibility to use quadratic interpolation for the corrector calculation. The difference to the LE/LI method is that in addition to the BOS and the predicted EOS value, also the previous-step (PS) value is used for the interpolation, in this case performed using a second-order polynomial. Linear interpolation is used for the first interval, since there is no data from previous steps. The PS value is available for the calculation without additional effort, so there is no time penalty compared to the other predictor-corrector methods.
- This method is currently considered the best choice for burnup calculations, although in this particular case the difference to LE/LI seems to be negligible, probably because the linear extrapolation on the predictor already yields a very good approximation.
- The example also demonstrates a problem related to large increases in step length. For the first interval plotted in the figure the step length between PS and BOS is 0.025 MWd/kgU, while the length between BOS and EOS is 1.9 MWd/kgU. Fitting a second order polynomial on the values results in a behavior that is clearly unphysical.
- The LE/QI method is available only in Serpent 2, and it is invoked using:


```
set pcc 4
```

 The mode number can be replaced by "LEQI".

Sub-step methods

Short-lived nuclides in secular equilibrium

- I-131 represents a low-absorbing fission product with a continuously changing production rate. The atomic density of such nuclide is characterized by:

$$\frac{dN}{dt} = S(t) - \lambda N .$$

The time dependence of the production term results from the depletion of U-235 and the build-up of Pu-239, and the fact that the fission yield for I-131 differs by a factor of 17 for the two actinides. Since, however, production rate is approximated by a constant value over each depletion step, the atomic density of I-131 is actually handled using equation:

$$\frac{dN}{dt} = S_{\text{ave}} - \lambda N ,$$

with the solution:

$$N(t) = \frac{S_{\text{ave}}}{\lambda} (1 - e^{-\lambda t})$$

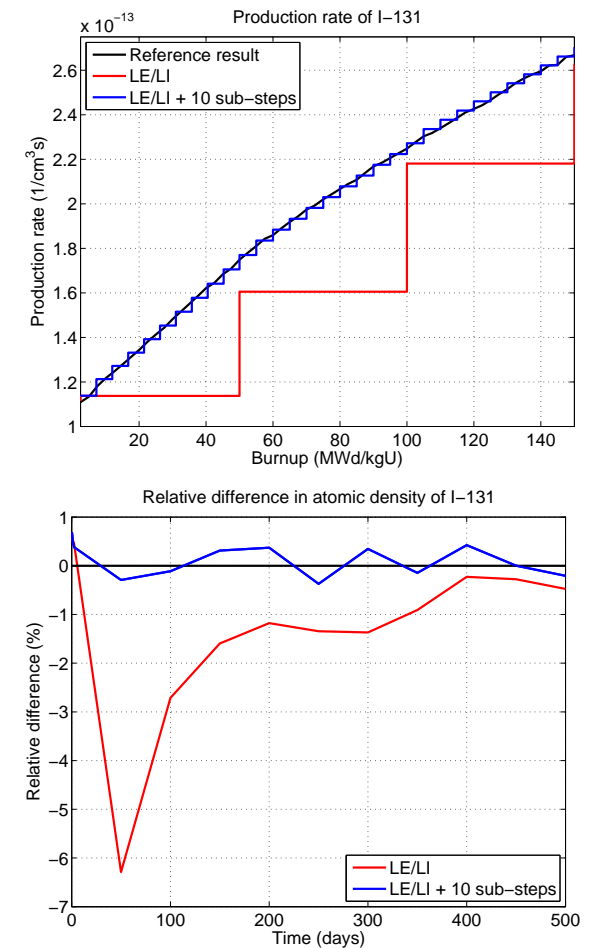
- The time step of 2 MWd/kgU used in the example case corresponds to 50 days, which is considerably longer than the 8-day half-life of I-131. The result is that the concentration of the nuclide reaches almost 99% of its equilibrium concentration (S_{ave} / λ) at the end of the step, and since the average production rate differs significantly from the EOS value, this concentration is grossly under-estimated.
- It is important to realize that the error does not result from the fact that the true behavior of the production rate is poorly represented by the analytical function, but rather that the constant reaction rate approximation breaks down. For this reason the problem cannot be solved simply by using a better fit for $S(t)$.

Sub-step methods

- One solution to the problem would be to use constant production rate that is closer to the EOS value, but defining the appropriate weighting for each nuclide becomes complicated in practice. A more universal solution is to use shorter depletion steps, but since running the transport calculation multiple times does not necessarily improve the approximation of $S(t)$, the same result can be accomplished with less CPU time by dividing only the depletion solution into multiple sub-steps. In Serpent 2, sub-stepping can be used for both predictor and corrector calculation with linear and quadratic extrapolation and interpolation.
- The figures on the right show the results for I-131 in the example case. The top figure shows how the reaction rate is approximated over the first 150 days of the irradiation cycle using linear interpolation, with and without sub-steps. With 10 sub-steps the sub-step-length is twice the step-length of the reference solution (2.5 days). Notice that $S(t)$ is approximated by a linear function in both cases. The second figure shows the relative difference to the reference result.
- It should be noted that the use of sub-steps does not automatically improve the results for all nuclides. The impact is not straightforward, and it depends on how fast the concentrations reach the secular equilibrium (compared to step length), and how closely the constant production rate approximates the time-dependent behavior. For more details, see Ref. [5].
- The sub-step solution is available only in Serpent 2, and currently not used by default. The method is invoked by:

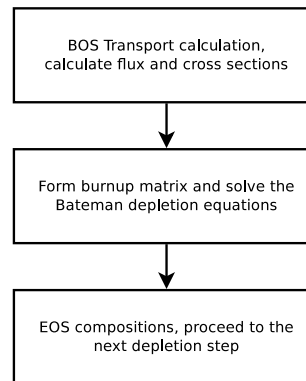
```
set pcc <mode> [ <ssp> <ssc> ]
```

where the two optional parameters are the number of sub-steps for predictor and corrector calculation, respectively.

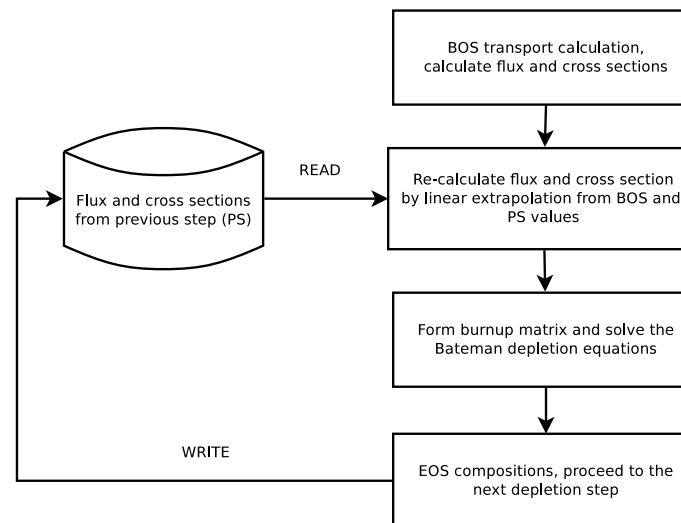


Flowcharts

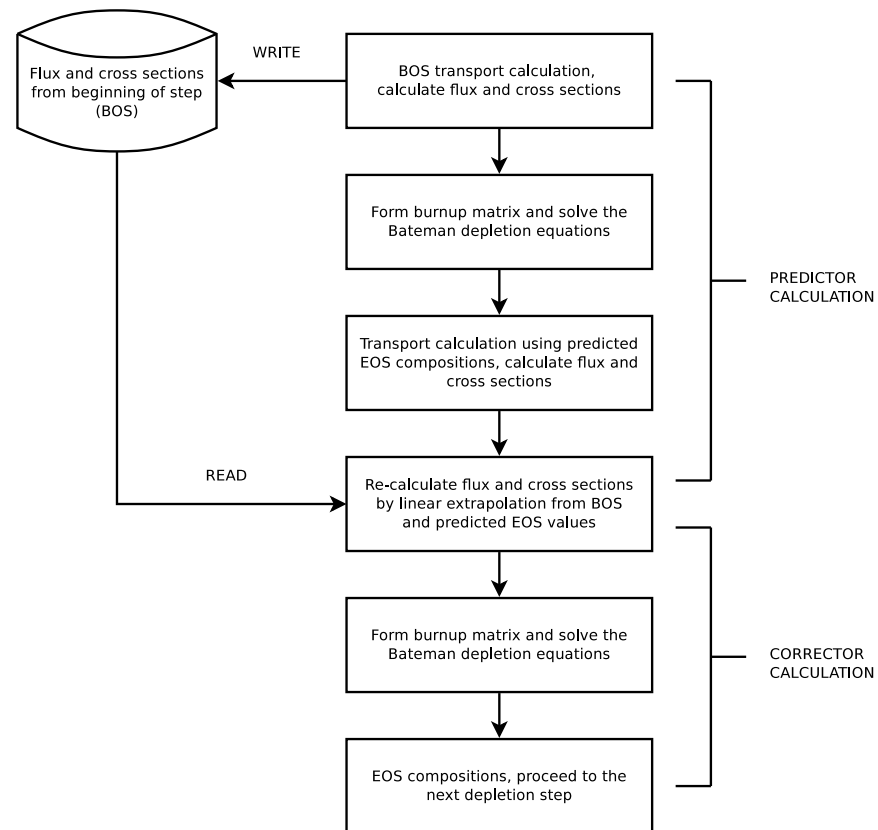
CE (conventional “Euler’s method”)



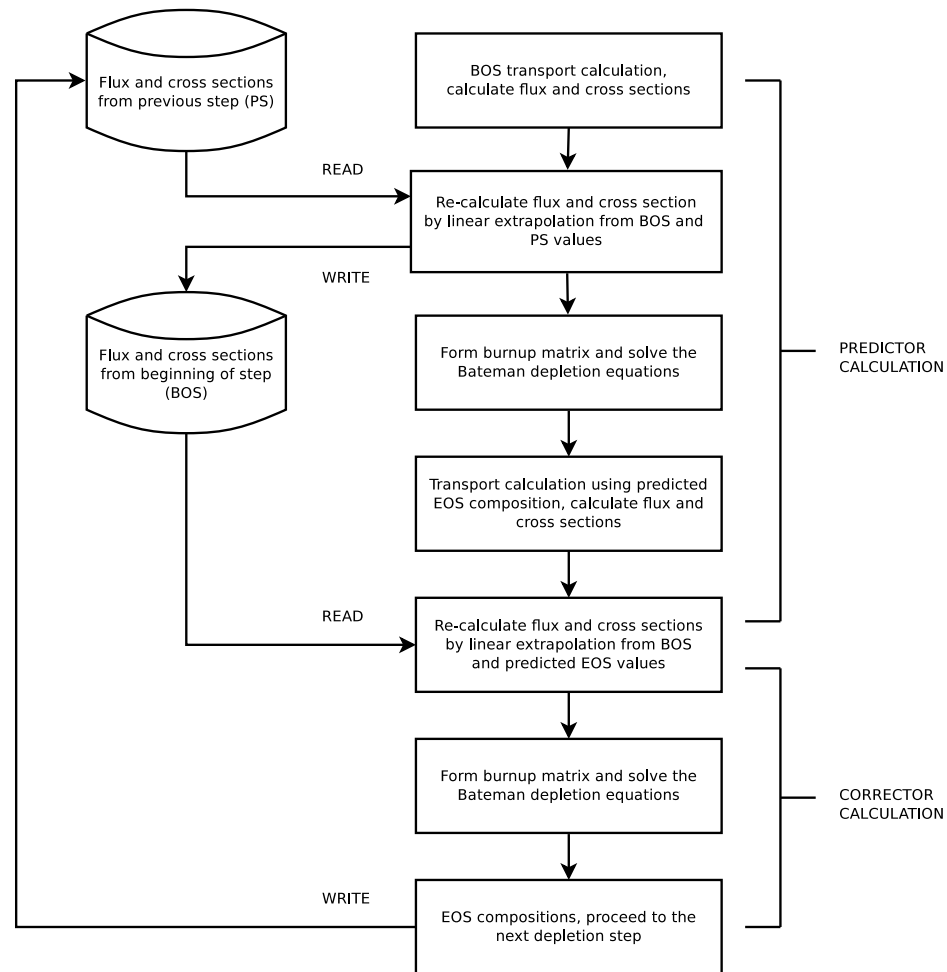
LE



CE/LI (conventional “predictor-corrector”)



LE/LI



LE/QI

