INL/RPT-23-74913
ANL/NSE-23/61

Initial Demonstration of New Griffin Technologies for Simulating the Running-In Phase of Pebble-Bed Reactors

June 17, 2024

Joshua Hanophy[1], Hansol Park[2], Vincent Labouré[1], Yeon Sang Jung[2], Sebastian Schunert[1], Yaqi Wang[1], and Javier Ortensi[1]

[1]*Idaho National Laboratory*
[2]*Argonne National Laboratory*

Idaho National Laboratory

Argonne NATIONAL LABORATORY

INL/RPT-23-74913
ANL/NSE-23/61

Initial Demonstration of New Griffin Technologies for Simulating the Running-In Phase of Pebble-Bed Reactors

Joshua Hanophy[1], Hansol Park[2], Vincent Labouré[1], Yeon Sang Jung[2], Sebastian Schunert[1], Yaqi Wang[1], and Javier Ortensi[1]

[1]**Idaho National Laboratory**
[2]**Argonne National Laboratory**

June 17, 2024

Idaho National Laboratory
Argonne National Laboratory

Prepared for the
U.S. Department of Energy
Office of Nuclear Energy
Under UChicago Argonne, LLC
Contract DE-AC02-06CH11357
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Page intentionally left blank

SUMMARY

Griffin is a reactor multiphysics modeling application based on MOOSE (Multiphysics Object-Oriented Simulation Environment) and specifically targeting transient modeling of advanced reactors. Griffin has been used recently to model pebble-bed reactors for the Nuclear Regulatory Commission (NRC) Office of Nuclear Regulatory Research and the Advanced Reactor Technology program. This modeling work has focused thus far on the direct calculation of equilibrium cores. This report documents an initial demonstration of a new running-in simulation capability. The new running-in capability is verified using the existing direct equilibrium core calculation capability. A simplified pebble-bed reactor model is then used to demonstrate the running-in simulation capability. This demonstration shows that Griffin is able to simulate years of operation during the running-in phase efficiently with each depletion step taking only several seconds. Two new technologies are also presented in this report which have been developed in Griffin that will be essential for improved accuracy both of the direct equilibrium core computation and the new running-in simulation capability. The first technology is an online cross section generation capability specifically targeted for pebble-bed reactors. This will improve the accuracy of the depletion calculation as the cross sections are generated at the exact core status. This also avoids the difficult step of pre-generating a separate standalone multigroup cross section set. Secondly, a newly implemented discretization for discontinuous finite element method (DFEM) S_N transport in cylindrical (RZ) coordinates, which can be solved efficiently using the existing S_N sweep solver, is discussed and some results are shown demonstrating the usefulness of the additional accuracy transport provides over a diffusion approximation.

ACKNOWLEDGEMENTS

This work was supported by the U.S. Department of Energy (DOE), Office of Nuclear Energy, under contract DEAC02-06CH11357 for UChicago Argonne, LLC, and under contract DE-AC07-05ID14517 for DOE Idaho Operations Office and supported by the DOE Nuclear Energy Advanced Modeling and Simulation Program.

This research made use of Idaho National Laboratory's High Performance Computing systems located at the Collaborative Computing Center and supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517.

Page intentionally left blank

CONTENTS

FIGURES

TABLES

ACRONYMS

BCC	body centered cubic
CFEM	continuous finite element method
CFL	Courant–Friedrichs–Lewy
CPM	collision probability method
CRAM	Chebyshev rational approximation method
DFEM	discontinuous finite element method
FCC	face centered cubic
LOFC	loss-of-forced-cooling
MOOSE	Multiphysics Object-Oriented Simulation Environment
PBR	pebble-bed reactor
RIA	reactivity insertion accidents
TRISO	tristructural isotropic

Page intentionally left blank

1. INTRODUCTION

Griffin is a Multiphysics Object-Oriented Simulation Environment (MOOSE) [14] based reactor multiphysics application specifically targeting transient simulations of advanced reactors. Recently, Griffin has been used for analyzing pebble-bed reactors (PBRs) [3, 7, 22, 24, 28, 29]. An overview of PBRs can be found in Reference [10] and general information about such reactors is not reviewed here, but some specific items related to two terms used frequently in this report, equilibrium core and running-in, are reviewed. PBRs can be fueled continuously without interrupting operation. Fuel pebbles flow slowly through the reactor, exiting either the bottom or top of the reactor, depending on the PBR design, and can either be recirculated at the other end or be discharged for spent fuel storage and replaced with fresh fuel pebbles. A PBR equilibrium core is one that does not change while operating; the isotopes fissioned to produce power are made up in the feed of fresh fuel pebbles into the reactor such that the reactivity does not change. Reactor characteristics, like total power, neutron flux shape, and fuel and coolant temperatures, are all constant in time. When operating normally at rated power, PBRs are generally expected to be in an equilibrium core configuration. Analysis of PBRs using Griffin has focused thus far on directly computing the PBR equilibrium core configuration. From this computed core configuration, various types of transient simulations can be initiated, including loss-of-forced-cooling (LOFC) accidents, reactivity insertion accidents (RIA), etc. The direct method implemented in Griffin for computing an equilibrium core is described in Reference [28] and reviewed in Section 1.1.

When a PBR begins operation for the first time, the core changes in time as power is produced. This phase of reactor startup and operation is referred to as the running-in phase. Initially, the reactor is filled entirely with graphite pebbles. A mixture of graphite and fuel pebbles will then be fed into the reactor as graphite pebbles leave the core. A mixture of fuel and graphite pebbles may be used initially instead of only fuel pebbles to prevent these initial fuel pebbles from experiencing unacceptably high power densities. Graphite pebbles are not needed in the equilibrium core because the fuel is burned to some extent and this burned fuel generates a lower power density than the fuel initially loaded into the reactor when it is started because this fuel has not burned at all, it is entirely fresh fuel.

As fuel is added to the core and graphite pebbles removed from the core, there is enough fuel

in the core eventually that it becomes critical. This is the first stage of the burning-in phase. The second stage is the approach to full power operation, where the core power ascension occurs. The graphite pebbles do not produce any power; so while there are graphite pebbles in the core, the reactor power level must generally be limited below the equilibrium core normal operating level to prevent any violation of the maximum power density limits in the fuel pebbles. It can take months to years for all the graphite pebbles in the core to be replaced with fuel pebbles. Even after all graphite pebbles have been removed from the core and the reactor power is raised to the nominal operating level, the average burnup of the fuel will still be lower than it is for the equilibrium core because the fuel has not yet had time to reach nominal burnup levels. Thus, during this phase, less fuel material is fed into the reactor than is burned to increase the average burnup. The running-in phase is complete when the average burnup approaches the equilibrium core level and all other parameters approach their equilibrium core values. The running-in phase generally takes place over years of operation.

This report describes the addition of a new running-in simulation capability to Griffin. The existing direct equilibrium core calculation methodology is reviewed first in Section 1.1 as it provides a foundation for the running-in capability. The model used in this report for investigation is reviewed in Section 2. The running-in methodology is then discussed in Section 3 and demonstrated with a specific example in Section 4. The last two sections cover technologies added to Griffin for supporting PBR modeling that will be important to the running-in simulation but are not fully integrated into the running-in workflow at the time this report was produced. Section 5 discusses a newly implemented discretization for discontinuous finite element method (DFEM)-S_N transport which can be solved efficiently using the existing S_N transport sweep solver in Griffin. Section 6 discusses the new capability to generate cross sections on-the-fly without needing to resort to a pre-generated cross section set. This will improve the accuracy of the depletion calculation as the cross sections are generated at the exact core status, which a pre-generated cross section set may not be able to represent accurately, resulting in more accurate running-in calculation results eventually.

1.1 Equilibrium core modeling in Griffin

The modeling approach used by Griffin for equilibrium core PBR calculations is described by Schunert et. al. [28]. The existing direct equilibrium core capability provided the basis for the new running-in capability and so the purpose of this section it to review key points relevant to this report. Because the simple demonstrative model used in this work involves the pebbles flowing downward under the force of gravity, the methodology presented here is described in terms of pebbles moving downward. Note however that the modelling approach is generally applicable to reactor types where the pebbles flow upward by buoyancy as described in [2]. Pebbles in a PBR core travel mostly downward with little movement radially (see Reference [26] for example), thus this simple flow field can be well approximated with several streamlines. The bottom of the PBR core has a cone to funnel the pebbles from the larger radius of the core to the smaller radius of the discharge chute where the pebble flow does include a radial component, but even here several streamlines describe the entire flow field well. Griffin requires the user to input one or more streamlines to describe the flow field. A flow area is associated with each streamline by dividing the core mesh elements among the specified streamlines based on proximity. The streamlines can be specified independently of the specific way the reactor core region is meshed. The core mesh can be unstructured. A general algorithm is used to transfer relevant quantities between the streamlines and the core mesh in a conservative way. RZ geometry is used for the model investigated in this work, thus the input streamlines represent concentric cylinders. More information about the streamlines and mesh used in this work is covered in Section 2.

The flow paths described by each streamline do not communicate; there is no flow between the streamlines, and thus each streamline specified represents an independent flow problem. The exception to this is that when pebbles leaving the bottom of the reactor are recirculated back to the top, they are assumed to be distributed among all the streamlines, and so the inlet of all streamlines are coupled to all the outlets. Ignoring this slight complication, the simple one-dimensional (1D) flow problems are modeled in Griffin using an Eulerian approach, where each streamline is discretized into a number of control volumes.

As the pebbles move down the core and produce power, their burnup increases. Calculating the burnup of pebbles is desirable for multiple reasons, the burnup of a fuel pebble at the time

it leaves the core is related to deciding whether the fuel pebble is reloaded or discarded. Pebble burnup in Griffin is computed similarly to movement in space, again using an Eulerian approach and the pebbles "advect" through burnup space with power being analogous to flow velocity. Generally, a user knows the maximum burnup expected from the fuel. This burnup space of interest is discretized based on user defined control volumes. The user inputs sequential burnup values which define the edges of control volumes in burnup space. The term bins is used for these control volumes in burnup space. The addition of this flow in burnup space means that a 2D flow problem must be solved solved; a one dimensional spatial flow problem and "flow" in burnup space with a "velocity" given by the power being produced per volume of fuel pebble.

The quantities conserved in each volume are the nuclide densities and pebble volume fractions. Consider $N(t,\tau)_{k,m,q}$ to represent the density of the nuclide q within the m^{th} segment of the k^{th} streamline with burnup τ. To have a discrete equation in burnup space, $N(t)_{k,l,m,q} = \int_{\tau_{l-}}^{\tau_{l+}} N(t,\tau)_{k,m,q}\, d\tau$ is defined where τ_{l+} and τ_{l-} are the bounds of the discrete burnup bins defined by a user for burnup bin l. The q index is dropped and the conserved quantity written as a vector for all nuclides, $\vec{N}(t)_{k,l,m}$. Finally, an index c is added to account for the possibility that different types of fuel pebbles are being used, for example two different fuel types with different enrichments. The nuclide density of each pebble type is tracked separately and given by $\vec{N}(t)_{c,k,l,m}$. In addition to the nuclide densities, the volume fraction occupied by each pebble type within each burnup bin is needed (the second quantity conserved) and is denoted with $n_{c,k,l,m}(t)$, where the indices have the same meaning as discussed previously.

The conservation laws for the nuclide densities and volume fractions are shown in Equations 1 and 2 where $A(\phi(t))$ is the decay and transmutation matrix, which is a function of the scalar flux $\phi(t)$. The indices all have the same meaning as discussed in the previous paragraph and u is the flow velocity and a is the flow area. $\Delta\tau_l$ is the width of burnup bin l and $\bar{\rho}_l$ is the average pebble power density for burnup bin l. A simple upwind scheme is used to derive these discrete equations, as discussed in more detail in Reference [28]. Such a numerical scheme is known to have significant numerical diffusion, and its impact is explored further in Section 3.4. The direct equilibrium equations are generated from Equations 1 and 2 by setting the time derivative to zero. For running-in, the time derivative is treated and time discretization is discussed in Section 3.1.

$$\begin{aligned}\frac{d\vec{N}_{c,k,l,m}(t)}{dt} =& u_{k,m-1}\frac{a_{k,m-1}}{V_{k,m}}\vec{N}_{c,k,l,m-1}(t) + \frac{\bar{\rho}_{c,l-1,m}}{\Delta\tau_{l-1}}\vec{N}_{c,k,l-1,m}(t) \\ &- \left[\left(u_{k,m}\frac{a_{k,m}}{V_{k,m}} + \frac{\bar{\rho}_{c,l,m}}{\Delta\tau_l}\right) I + A(\phi(t))\right]\vec{N}_{c,k,l,m}(t)\end{aligned} \tag{1}$$

$$\begin{aligned}\frac{dn_{c,k,l,m}(t)}{dt} =& u_{k,m-1}\frac{a_{k,m-1}}{V_{k,m}}n_{c,k,l,m-1}(t) + \frac{\bar{\rho}_{c,l-1,m}}{\Delta\tau_{l-1}}n_{c,k,l-1,m}(t) \\ &- \left(u_{k,m}\frac{a_{k,m}}{V_{k,m}} + \frac{\bar{\rho}_{c,l,m}}{\Delta\tau_l}\right) n_{c,k,l,m}(t).\end{aligned} \tag{2}$$

The boundary conditions for Equations 1 and 2 are discussed in detail in Equations 28–36 in Reference [28]. They are summarized and generalized for running-in in Equations 3 and 4 where ∂^- is the inflow boundary for the streamline. K, C, and L are the indices of the last streamline, pebble type, and burnup bin respectively, and M is the index of the outflow segment for the respective streamline. These equations state that the inflow to a streamline is a function of the outflow. Griffin includes logic to determine what fraction of pebbles leaving the core are recirculated to the top, with the remainder being made up with fresh pebble feed, based on user input as described in Reference [28]. In the case of the direct equilibrium core calculation, the functions in Equations 3 and 4 do not change in time, but as discussed in Section 4, time-dependence is accounted for when specifying the dependence of inflow to the core on outflow for running-in simulations.

$$\vec{N}_{c,k,l,\partial^-}(t) = f(\vec{N}_{0\ldots C,0\ldots K,0\ldots L,M}(t)) \tag{3}$$

$$n_{c,k,l,\partial^-}(t) = f(n_{0\ldots C,0\ldots K,0\ldots L,M}(t)) \tag{4}$$

2. MODEL DESCRIPTION

A simplified PBR-like model was used to test and demonstrate the new running-in capabilities described in this work. The dimensions are inspired by the 200 MW_{th} general PBR, which is a generalized gas-cooled PBR first introduced in References [29, 30], but the simplified model consists only of a core and reflector. The shape and dimensions of this model are shown in Figure 1. As described in Reference [28] and reviewed in Section 1.1, Griffin allows the streamlines to be

specified independently of the mesh used for neutronics. However, the simplified geometry can easily be represented with a structured mesh, and so it was natural to have the streamlines exactly correlate to the neutronics mesh. As there are clearly five columns of mesh cells, each with 25 axial cells, there are five streamlines with 25 segments each. Since the streamlines correspond exactly to the core mesh, there is no interpolation of data between them, other than simply moving the values between the core mesh cell and the corresponding streamline cell.

Griffin would typically be coupled to other MOOSE-based applications, which would calculate the fuel, moderator, and reflector temperatures based on the calculated power from Griffin, as well as inputs relevant to these calculations such coolant flow and material thermal properties. For the simple running-in demonstration included in this report, Griffin is not coupled to any other code, and it was deemed sufficient to simply fix the temperatures at 900 K throughout the simulation for the initial demonstration of the running-in capability. For reference, the discretization in burnup space used for all calculations in the report is shown in Table 1. Internally, Griffin uses J/m^3 as the required input format, though we plan on adding the option to input the burnup discretization using other burnup units.

Table 1: Burnup bin boundaries used for the discretization of burnup space

Bin	J/m^3
0	0 - 8.77e13
1	8.77e13 - 1.754e14
2	1.754e14 - 2.631e14
3	2.631e14 - 3.508e14
4	3.508e14 - 4.385e14
5	4.385e14 - 5.262e14
6	5.262e14 - 6.139e14
7	6.139e14 - 7.016e14
8	7.016e14 - 7.893e14
9	7.893e14 - 8.770e14
10	8.770e14 -∞

2.1 Cross Sections

The multigroup cross sections used in this work were generated for an equilibrium core configuration. These cross sections were generated with a similar methodology to that described in Reference [22]. Nine energy groups are utilized with a single cross section set for the entire core

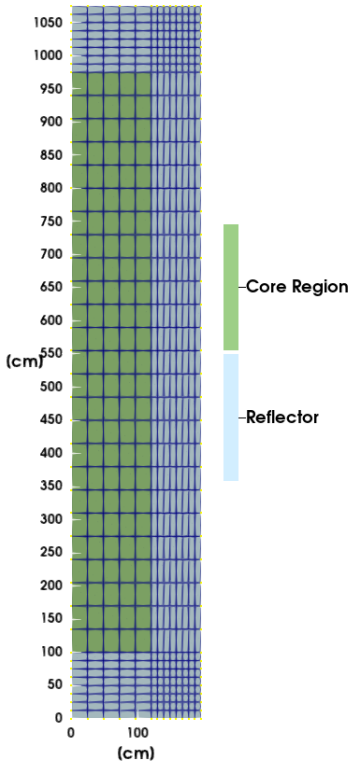

Figure 1: Schematic showing the shape and dimensions of the simplified model.

region and a second set used specifically for the reflector. The cross sections are tabulated with respect to burnup, fuel temperature, and moderator temperature; although, the temperatures for this work were simply held constant at 900 K.

Since the core configuration during the running-in phase is expected to be significantly different from the equilibrium core configuration, these cross sections are not expected to be physical for the running-in simulation. However, they are sufficient for demonstrating the new capability. As such, this report does not include any validation of the running-in simulation. An online cross section capability for Griffin is under development, which is described in Section 6. Validation will be performed once the online cross section capability is fully integrated into the running-in workflow.

3. Running-In Methodology

This section describes the methodologies implemented in Griffin for PBR running-in simulations. The newly implemented numerical methods for handling the time derivatives in

Equation 1 and 2, not present for the previously existing direct equilibrium core calculation, are described in Section 3.1. An initial investigation into the impact of numerical diffusivity on simulation results with the simple upwinding scheme is discussed in Section 3.4. The last section discuss the neutronics methodology used for the running-in simulation although no new neutronics capabilities were implemented specifically for running-in.

3.1 Time Integration

The discrete in time equations for Equation 1 and 2 are shown in Equations 5 and 6. These equations implement forward Euler time-stepping for time integration.

$$\vec{N}^{\ell*}_{c,k,l,m} = \Delta t \begin{pmatrix} \frac{a_{k,m-1}}{V_{k,m}} u_{k,m-1} \vec{N}^{\ell}_{c,k,l,m-1} - \frac{a_{k,m}}{V_{k,m}} u_{k,m} \vec{N}^{\ell}_{c,k,l,m} + \\ \frac{\bar{\rho}_{c,l-1,m}}{\Delta\tau_{l-1}} \vec{N}^{\ell}_{c,k,l-1,m} - \frac{\bar{\rho}_{c,l,m}}{\Delta\tau_{l}} \vec{N}^{\ell}_{c,k,l,m} \end{pmatrix} + \vec{N}^{\ell}_{c,k,l,m}$$
$$\vec{N}^{\ell+1}_{c,k,l,m} = \mathcal{S}_{\Delta t, A(\phi(t^{\ell}))} \vec{N}^{\ell*}_{c,k,l,m} \tag{5}$$

$$n^{\ell+1}_{c,k,l,m} = \Delta t \begin{pmatrix} \frac{a_{k,m-1}}{V_{k,m}} u_{k,m-1} n^{\ell}_{c,k,l,m-1} - \frac{a_{k,m}}{V_{k,m}} u_{k,m} n^{\ell}_{c,k,l,m} + \\ \frac{\bar{\rho}_{c,l-1,m}}{\Delta\tau_{l-1}} n^{\ell}_{c,k,l-1,m} - \frac{\bar{\rho}_{c,l,m}}{\Delta\tau_{l}} n^{\ell}_{c,k,l,m} \end{pmatrix} + n^{\ell}_{c,k,l,m} \tag{6}$$

Operator splitting is used to solve Equation 1 as shown in Equation 5, where Δt is the timestep, ℓ denotes the solution at the current time, $\ell + 1$ the solution at $t^{\ell} + \Delta t$, and $\ell*$ represents an intermediate step. Forward Euler is used to compute an intermediate solution $\vec{N}^{\ell*}_{c,k,l,m}$ from $\vec{N}^{\ell}_{c,k,l,m}$, accounting only for advection in space. The operator $\mathcal{S}_{\Delta t, A(\phi(\ell))}$ represents the Chebyshev rational approximation method (CRAM) [25], which is used to compute $\vec{N}^{\ell+1}_{c,k,l,m}$. Note the subscripts of $\mathcal{S}_{\Delta t, A(\phi(\ell))}$ which denotes a dependence on the scalar flux; this dependence is simply lagged by using the scalar flux from timestep ℓ. Also note the timestep used in $\mathcal{S}_{\Delta t, A(\phi(\ell))}$ is Δt. The time discretized version of Equation 2 is shown in Equation 6. In this case, no operator splitting is needed; the equation simply shows the application of forward Euler.

As is generally known, forward Euler is only conditionally stable for advection problems. However, the limit on the timestep that must be imposed when using forward Euler does not appear to be an issue in practice since the flow velocity of the pebbles is very low. As shown in Figure 6 for the sample problem investigated in this work, the maximum pebble feed rate remains

below 2 pebbles/minute. While the test problem created for this work is highly simplified, it's still expected to be representative of realistic models. In Reference [34], the maximum pebble feed rate during the running-in simulation appears to have been approximately 7 pebbles/minute. This maximum pebble feed rate would lead to a flow velocity of 4.2e-4 cm/s. Each cell of the sample mesh used in this report, shown in Figure 1, is 44.44 cm long, and thus the timestep leading to a Courant number of one is 1.22 days. This is not a prohibitively small timestep since each depletion step is fast, as described in Section 4, and the feed rate is usually much lower than this maximum value, allowing a much larger timestep to be taken if desired.

As noted in Section 1.1, the problem being solved by Griffin is similar to a 2D flow problem with a Courant–Friedrichs–Lewy (CFL) like condition also applying to the "flow" in burnup space. The timestep would be limited by whichever pebble type and burnup bin has the maximum power density considering the entire core. The timestep value that leads to a Courant number of one is plotted in Figure 2. The CFL condition in burnup space appears to be less severe than that for the spatial flow; however, this should be investigated for more models in the future.

Forward Euler time integration exhibits only first-order convergence in time. However, the computational time per time-step is relatively small, and so it is not necessary to take very large time-steps. Taking larger time-steps is of course still generally desirable and may be more necessary for 3D Cartesian models compared to the 2D cylindrical geometry model presented in this work, and so more sophisticated and accurate numerical methods will be implemented and tested in the future for time-stepping.

3.2 Verification Using the Direct Equilibrium Core Calculation

As can be seen in Equations 1 and 2, when the time derivatives are set to zero as they are for the direct equilibrium core calculation, Additional terms are added to the decay and transmutation matrix to account for the downward flow in physical space and burnup space. This matrix is then inverted with a direct linear solver. In the case of obtaining an equilibrium solution using the running-in implementation, considerably more code logic is required for handling the time derivative. As discussed in the previous section, CRAM is being used is used when solving the Bateman equations. CRAM is not a consistent method as was shown in Reference [5], as the truncation error does not go to zero as the timestep goes to zero. In fact, Reference [5] shows

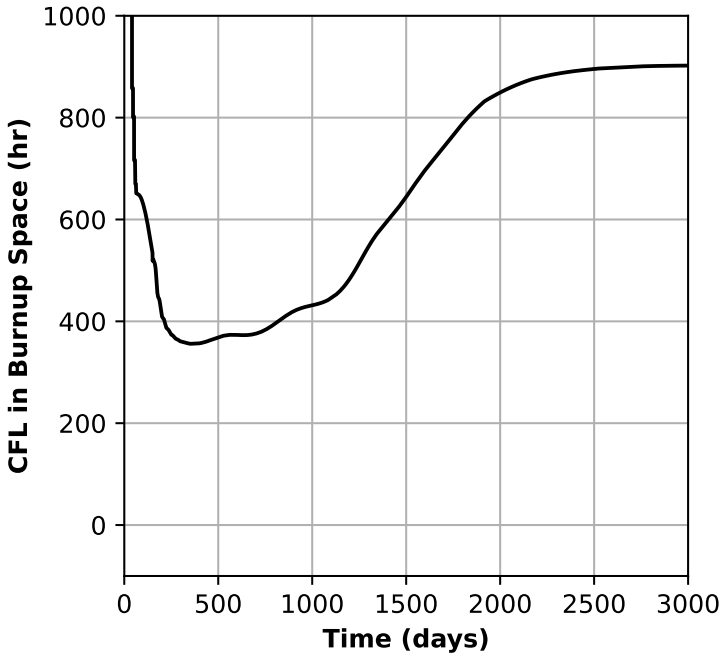

Figure 2: Timestep size that results in a Courant number of one plotted throughout the running-in simulation presented in Section 4.

that, for lower CRAM expansion orders, with timesteps as high as several days, the error grows significantly to orders of magnitude relative error for nuclide compositions. However, for higher CRAM orders it was shown that error remained small for timesteps as low as several minutes. This work uses CRAM order 16, and results are compared between a direct equilibrium calculation and a converged steady state running-in calculation that used a timestep of one hour. At this timestep size, any error from order 16 CRAM should be unnoticeable compared to the error from the forward Euler time integration.

To verify the handling of the time dependent code implemented for running-in simulations, a direct equilibrium calculation was performed using the model described in Section 2 with the parameters shown in Table 2. Note that the burnup limit shown in Table 2 means, in reference to Table 1, that pebbles in burnup bins 9 and 10 are discarded when they leave the bottom of the reactor and this fraction of discarded outflow is made up with fresh 15.5 wt% enriched fuel pebbles while any pebbles in burnup bins 0–8 are recirculated to the top of the core. A running-in simulation was then performed with the same burnup limits and feed rate. The eigenvalue was not controlled in any way during the running-in simulation since only the final steady state

answer mattered. The running-in simulation simply continued stepping through time until the calculation results stopped changing within some specified tolerance. The chosen tolerance was small enough to have a negligible impact on any differences from the direct equilibrium core calculation. The relative difference between the converged running-in calculation and the direct equilibrium core calculation for several integral quantities is shown Table 3. The eigenvalue for the direct equilibrium core calculation was 1.007384 and 1.007486 from the running-in simulation, for a difference of 10 pcm. The converged partial powers are shown in Figure 3 for three burnup bins. The maximum relative difference is only 0.24%.

Table 2: Equilibrium core simulation parameters used for verification of running-in simulation.

Quantity	Value
Pebble Feed Rate	1 pebble/minute
Burnup Limit	7.893e14 J/m^3
Fuel Pebble Enrichment	15.5 wt%
Fuel and Moderator Temperature	900 K

Table 3: Relative difference between converged running-in and direct equilibrium core calculations for several integral nuclide core quantities.

Integral Core Inventory of Nuclide	Absolute Relative Difference
^{239}Pu	1.873E-04
^{149}Sm	1.944E-04
^{235}U	3.120E-04
^{238}U	7.411E-06
^{135}Xe	4.126E-04

3.3 Feed Rate Control

The simple model used in this report to demonstrate the new running-in capability does not have control rods. The reactivity is managed only by varying the pebble feed rate during the simulation. This section describes the logic used to determine the feed rate. A target eigenvalue is input by a user, this is the target eigenvalue for the end of a depletion timestep. At the start of a timestep, the difference between the actual eigenvalue and the target is used to adjust the feed rate with simple proportional control logic. Error bounds above and below the target eigenvalue are user specified, and if, at the end of a timestep, the computed eigenvalue is outside of the error

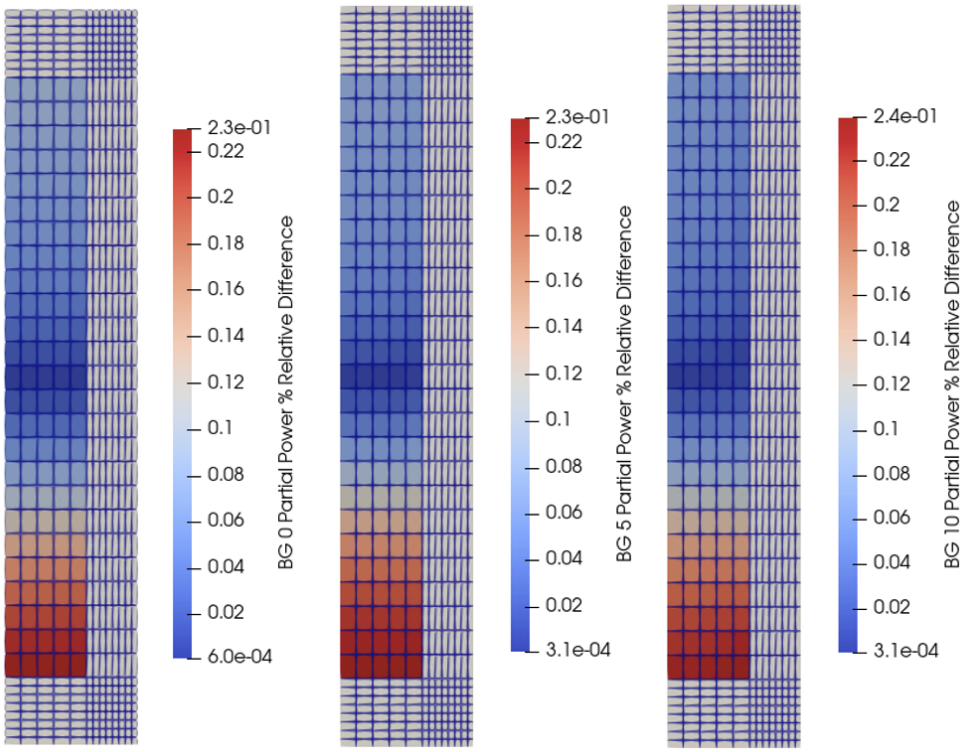

Figure 3: Comparison of the computed partial powers for three different burnup groups (BG) between the direct equilibrium core calculation and converged running-in simulation. Differences shown are absolute.

bounds, the depletion step is recomputed by using a simple weighted bisection method to pick a feed rate that results in an acceptable eigenvalue at the end of the timestep. This bisection wraps the simple proportional control logic since the control logic is not expected to be robust enough to account for step changes in parameters like the pebble type being fed into the reactor.

3.4 Critical-Height Calculation and Numerical Diffusivity

A critical-height calculation capability was initially created for code verification purposes. It is a useful verification calculation because the same code that advects the pebble downward during the running-in simulation is used in the critical-height calculation, but the critical-height simulation is simpler; the power and scalar flux used in the pebble depletion model can be set to zero so the pebbles should not advect in burnup space or undergo transmutation. The critical-height calculation is not intended for validation because appropriate multigroup cross sections for this configuration have not been generated and some potentially important details like the exact

layout of the pebbles at the very top of the reactor core are not accounted for in the Griffin model.

The critical-height calculation was also a useful tool for an initial investigation into the effect that numerical diffusion has on the running-in simulation. The simple spatial discretization currently used in Griffin is known to be diffusive. In actual plant operation, there would be a relatively sharp interface between the graphite pebbles and the initial fuel pebbles loaded into the reactor to bring it to criticality. The numerical diffusion inherent in the currently implemented scheme means this sharp interface cannot be modeled; it becomes a diffuse interface. This can be mitigated by refining the spatial mesh, however, this increases computation time for a timestep and because an explicit time integration scheme is currently implemented, the more refined the mesh become, the smaller a timestep that must be taken. Using a very fine mesh might be used for investigative purposes in the future, but is not a long-term solution for running-in simulations.

The impact of numerical diffusivity is shown qualitatively in Figure 4. Note the porosity used in this model is 0.4, meaning that a pebble volume fraction of 0.6 fills the entire space available to the pebbles. Thus, where the bottom section of the core shown in Figure 4 has a graphite pebble volume fraction of 0.6, the entire space available to pebbles is filled with graphite. The top section of the core is 80% fuel pebbles, with the remaining being graphite pebbles; hence, the volume fraction of graphite pebbles is 20% of 0.6. At time zero, the sample calculation is started from a sharp interface. Two additional plots at $t = 250$ and $t = 500$ hours, are shown for this problem where the feed rate is set to 1 pebble/minute. The sharp interface rapidly diffuses. Some investigation into the the potential importance of the numerical diffusion is discussed next.

Figures 5a and 5b show the result of a critical-height calculation that was started with no fuel in the reactor. The eigenvalue in Figure 5b is plotted with the volume of fuel pebbles (not the volume of tristructural isotropic (TRISO) particles but the volume accounting for the entire pebble) instead of the height, since the height is not necessarily well defined for this diffuse interface. A second calculation was performed that was started from a sharp interface near the critical height found from the simulation shown in Figures 5a and 5b. The eigenvalues for these two calculations are compared in Figure 5c. There is approximately a 4.5% difference in the predicted fuel pebble volume for criticality between the two calculations. This difference is cause by different material configurations predicted between the two calculations. The calculation started very near the critical-height has less numerical diffusion in the solution as it takes some amount of time to fully

develop.

The current work flow used for running-in simulations is to perform a critical-height calculation and then start the running-in simulation from the critical configuration found. As discussed in the previous paragraph, there is a difference between starting the critical-height calculation with no fuel in the reactor or starting it from a sharp interface near the critical height. Ignoring other issues such as multigroup cross section accuracy, the calculation started near the critical height should be more accurate. However, it likely does not matter at this stage of technical development. For example, it would be possible to start the running-in simulation from the computed critical height and have a perfectly sharp interface by simply putting a mesh edge at the critical height, but the sharp interface would rapidly diffuse into a diffusive interface as the running-simulation progressed. More sophisticated numerical methods must be implemented in the future to test the impact of this diffusivity on overall simulation results.

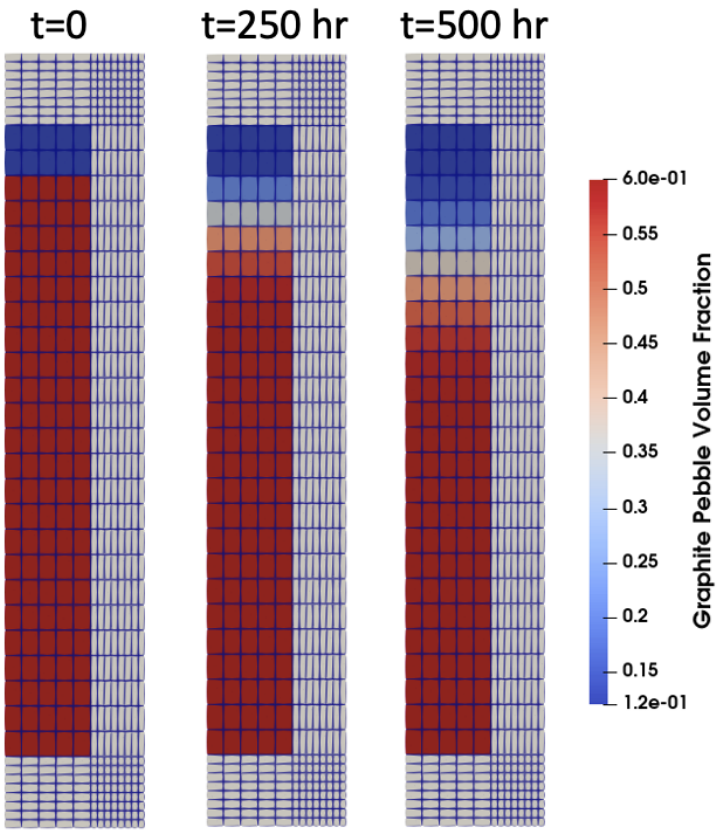

Figure 4: Numerical diffusion causing a diffusive solution in the advection result when started at time zero from a sharp interface configuration.

At this time, it's unclear a more exact determination or specification of the critical height from

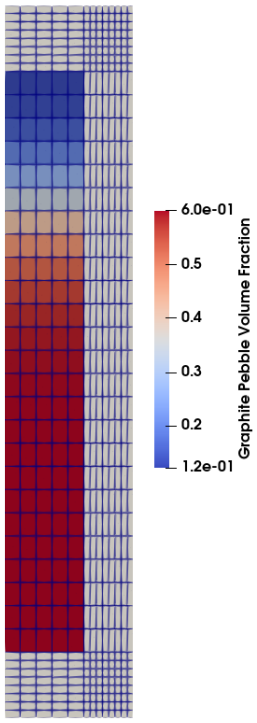

(a)

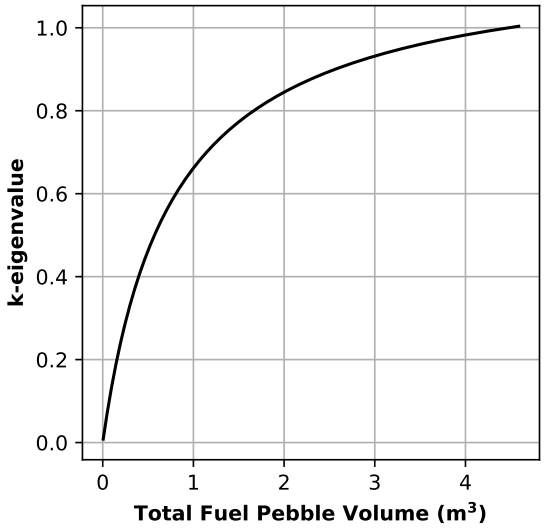

(b)

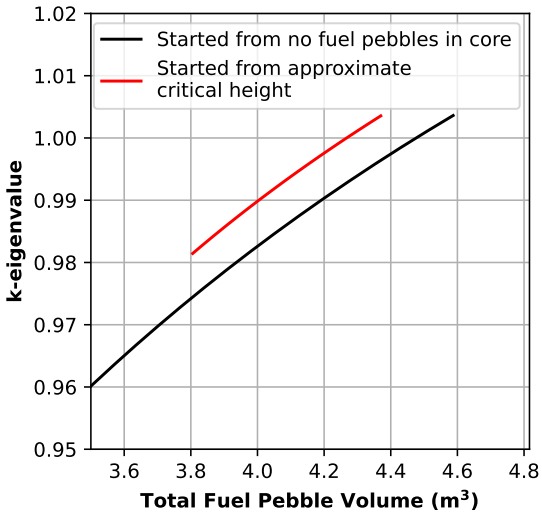

(c)

Figure 5: Results for a critical core configuration computed by a critical height calculation. The graphite pebble volume fraction at the critical height is shown in 5a along with the eigenvalue as a function of total fuel pebble volume in the core shown in 5b. In 5c, the eigenvalue from the calculation shown in Figures 5a and 5b is compared against a result where the calculation is started from a sharp interface close to the critical height.

which to start the running-in simulation is warranted since the solution will experience numerical diffusion as the pebbles begin advecting in space and burnup.

3.5 Neutronics

Multigroup diffusion was used for neutronics in the running-in simulation. The eigenvalue solver in Griffin typically used for diffusion problems is based on SLEPc [33]. At the time that Reference [33] was published, the residual was evaluated on the fly. Since then, a new option has been added for smaller systems to simply assemble and store the entire matrix for the linear system. The matrix can be reassembled for nonlinear iterations if there is some reason for doing this. This new option significantly speeds up the eigenvalue calculation, and the nine energy group problem in RZ geometry investigated in this report is small enough that the matrix can be assembled and stored. Each eigenvalue solve takes a fraction of a second. As described in Section 5, Griffin now includes an efficient sweep solver for DFEM-S_N transport in RZ geometry. That section shows DFEM-S_N transport producing results significantly different from diffusion for a problem with a void region in the top of the core. DFEM-S_N will likely be important for achieving physically correct results in future validation work; however, for the sample running-in simulation included in this work, diffusion should be sufficient. The simplified model used in this work does not have features typically challenging for a diffusion approximation such as void regions.

4. Running-In Simulation Results

An example running-in simulation is documented in this section as a demonstration of the the new capabilities implemented in Griffin. The model used for the sample calculation is described in Section 2. The simulation was started from a critical-height calculation as described in Section 3.4. A relatively small uniform timestep of 5.56 hours was used. A larger timestep can be taken as discussed in Section 3.1, and in particular the timestep could be allowed to vary up to several days. The computational time per timestep on a single workstation processor is approximately four seconds. As discussed in Section 3.1, the currently implemented numerical method is only first-order accurate in time; a predictor-corrector method should be implemented in the future.

Along with this addition, more consideration can be made for what the ideal timestep is.

To demonstrate the running-in capabilities associated with specifying the time-dependent feed of pebbles into the reactor, three different types of pebbles are assumed to be utilized during the running-in phase; a graphite-only pebble, a 5% enriched fuel pebble, and a 15.5% enriched fuel pebble. During the first phase of the simulation, only graphite pebbles leave the bottom of the reactor. Table 4 shows the time-dependent feed of pebbles into the core as a function of the type of pebble discarded from the reactor. For example, in the first 10 days, when graphite pebbles leave the bottom of the core, they are replaced with a mix of 80% fuel pebbles (5% enrichment) and 20% graphite pebbles. From 10 to 30 days, the ratio is changed to 90% fuel. Note that Table 4 shows the makeup feed as a function of discarded pebbles, not just the pebble leaving the bottom of the core. The assumed burnup limit for the 5% enriched fuel is 1.754e14 J/m^3 and 7.893e14 J/m^3 for the 15.5% enriched fuel. Early in the simulation, when a 5% enriched fuel pebble leaves the bottom of the core, it will generally be below this burnup limit and will be recirculated. After the fuel pebble has produced enough power such that, at the exit, the pebble burnup is beyond the user specified limit, it is replaced with a 15.5% enriched fuel pebble as shown in Table 4. The feed regime shown in Table 4 is somewhat arbitrary; it has not been optimized, and the results shown later indicate it could be optimized further. The only condition used for creating this specific feed regime was to ensure the reactor has enough reactivity to stay critical throughout the running-in simulation.

Table 4: Pebble makeup feed specifications as a function of the discarded pebble types.

Fraction of Pebble Type In Makeup Feed	Discarded Pebble Type		
	Graphite Only	5 % Enriched Fuel	15.5 % Enriched Fuel
Time <10 days			
Graphite Only	0.2	0.8	0.0
5 % Enriched Fuel	0.0	0.0	1.0
15.5 % Enriched Fuel	0.0	0.0	1.0
Time <30 days			
Graphite Only	0.1	0.9	0.0
5 % Enriched Fuel	0.0	0.0	1.0
15.5 % Enriched Fuel	0.0	0.0	1.0
Time <150 days			
Graphite Only	0.0	1.0	0.0
5 % Enriched Fuel	0.0	0.0	1.0
15.5 % Enriched Fuel	0.0	0.0	1.0
Time $\geq$ 150 days			
Graphite Only	0.0	0.0	1.0
5 % Enriched Fuel	0.0	0.0	1.0
15.5 % Enriched Fuel	0.0	0.0	1.0

Figure 6 shows several results from a running-in simulation using the makeup feed regime shown in Table 4. The logic used to control the feed rate and manage reactivity is described in Section 3.3. The target k-eigenvalue was 1.005 with a maximum allowable error of $\pm$200 pcm before the weighted bisection method wrapping the proportional controller jumped to an acceptable feed rate. The simple proportional logic used to control the feed rate exhibits some oscillations, particularly at 150 days when, as shown in Table 4, the makeup feed is changed from graphite pebbles being replaced with 5% enriched fuel to 15.5% enriched fuel. The oscillation is small in the feed rate plot but more significant in the k-eigenvalue plot. Such oscillations would likely be damped some if temperature feedback were included as part of a multiphysics simulation, as opposed to simply using the constant temperatures in this calculation. The oscillations quickly dampen, and the reactor slowly approaches an equilibrium configuration, reaching such a state after approximately 2,500 days.

The total reactor power as a function of time is shown in Figure 6. It is specified as a simple function of the volume of fuel pebbles in the core relative to the volume of fuel in the equilibrium core (simply the volume of the core less the porosity). Additionally, this function specifying power as a function of fuel volume is multiplied by another linear function of time, which goes from 0 to 1 over the first 30 days. This second function of time is meant to allow the power to ramp up to some value instead of immediately jumping to the value attained as a function of fuel pebble volume since at time 0 there is already a significant amount of fuel in the reactor.

The maximum power density during the running-in simulation is plotted in Figure 7b. It increases rapidly in the first 300 days to a relatively large value. If this simulation were further optimized, the feed regime could be changed to introduce the 15.5% enriched fuel pebbles more slowly. Logic to control the reactor power based on the maximum power density in the core has also been implemented; although, it was not tested for this report. Critically, the assumption of uniform temperatures of 900 K means that there is no thermal feedback, which would likely limit the maximum power density in a realistic multiphysics simulation. The average discharge burnup is shown in Figure 7a. Note that the average discharge burnup is zero while only graphite pebbles leave the bottom of the reactor and jumps to a non-zero value when fuel pebbles begin to leave the bottom of the reactor.

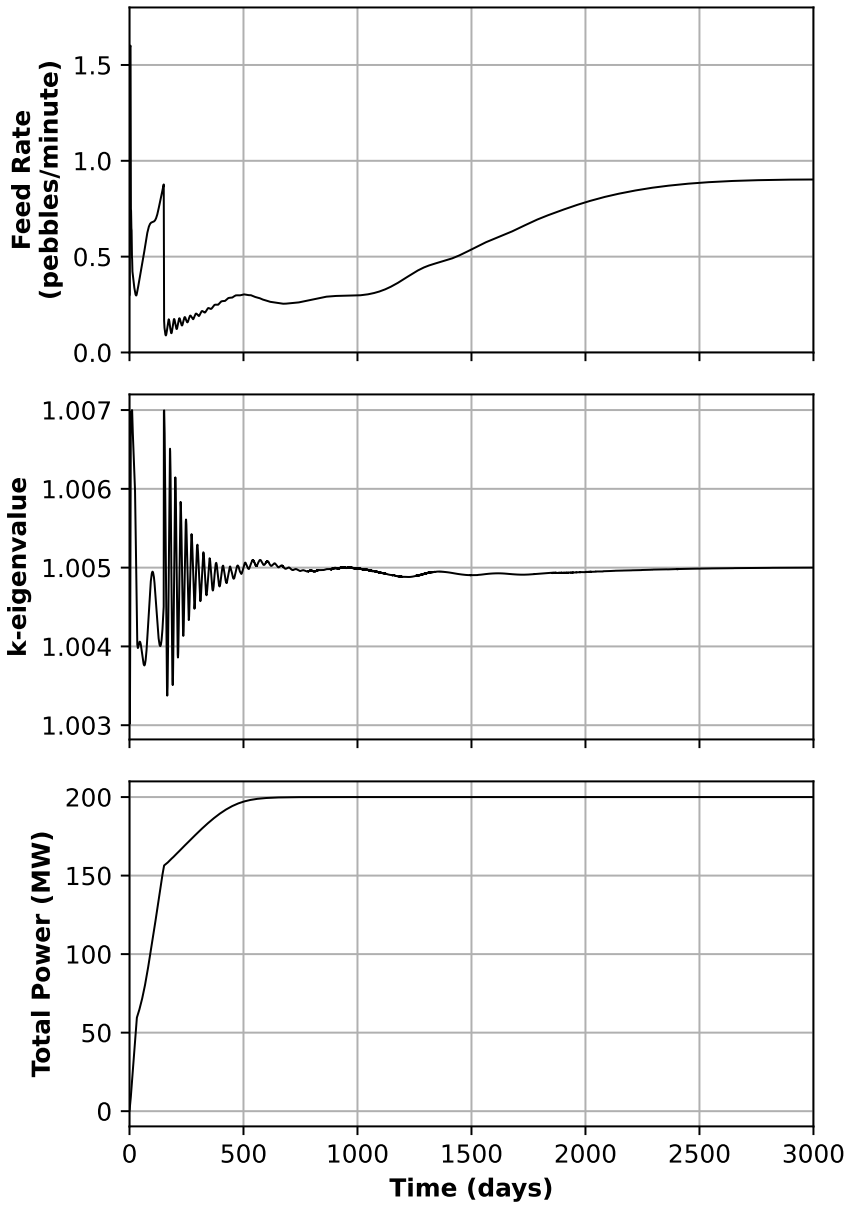

Figure 6: Schematic showing the shape and dimensions of the simplified model.

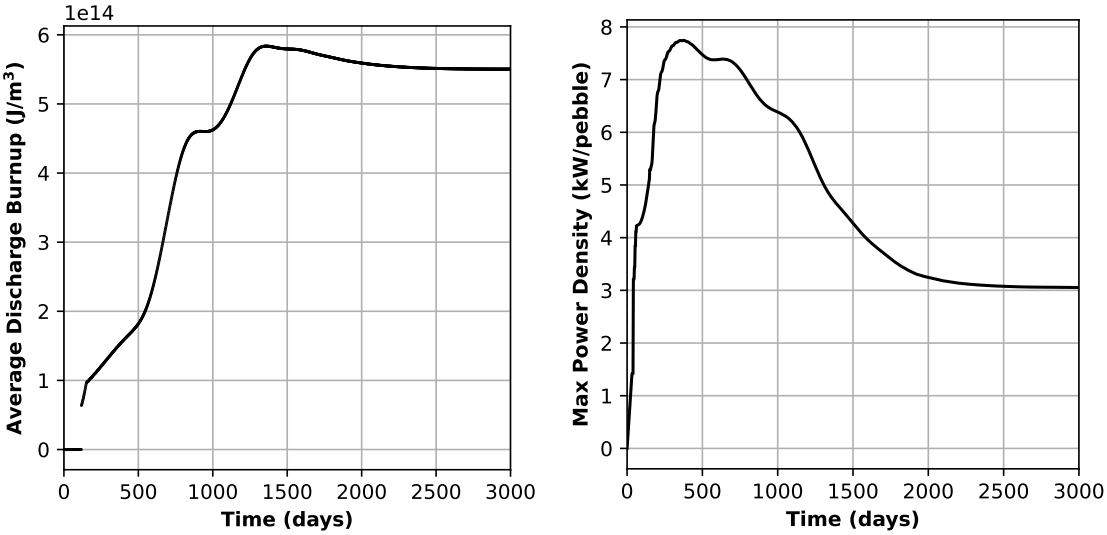

(a) (b)

Figure 7: The average discharge burnup and maximum power density plotted for the duration of the running-in simulation.

5. Applications for DFEM-SN RZ Geometry Sweep Solver

A new discretization was added to Griffin for DFEM-S_N transport in 2D RZ geometry [13] that can be solved with the existing Griffin sweeper, which was previously implemented for efficiently solving DFEM-S_N transport in Cartesian geometry [31]. The new discretization is based on the well known weighted diamond difference method [18]. The implementation was verified by comparing results against a discretization using a P_N expansion [32] to handle the angular derivative, which was already implemented in Griffin but could not be solved efficiently with the transport sweeper and was thus limited to use only for small test problems. More specific details about the DFEM-S_N RZ sweeper in Griffin will be communicated in a future dedicated publication.

The ability to efficiently solve DFEM-S_N RZ geometry will be a key part of Griffin's running-in capabilities in the future. It's not clear the additional accuracy that transport provides over diffusion would be meaningful to the running-in results presented in this work since the online cross section capability was not fully integrated into the running-in workflow. However, the cross sections used in this work are expected to be suitable for an equilibrium core calculation. As such, this section presents results for a reactivity-curve comparing diffusion to transport. This provides a basis for the potential importance of using transport when benchmarking results for running-in, which will be key in addition to improved cross sections and better spatial discretization to address the numerical diffusivity issues reviewed in Section 3.4.

5.1 Reactivity Curve

In order to qualitatively assess the benefits of the DFEM-S_N RZ sweeper compared to a diffusion solver, reactivity curves are compared. For this purpose, the gas-cooled PBR model studied in Reference [12]—featuring movable control rods—is leveraged. To make the comparison more relevant, the same temperature profile (900 K everywhere) and pebble isotopic compositions (the ones obtained via an equilibrium core calculation with the diffusion solver) are imposed. Eigenvalue calculations are then run at various control rod positions with continuous finite element method (CFEM)-diffusion and the DFEM-S_N RZ sweeper (two polar and four azimuthal angles per quadrant), as shown in Figure 8. The main differences observed between the two

solvers occur when the control rods are shallowly or deeply inserted. In other configurations, most of the neutrons reside in highly diffusive regions so the diffusion approximation is adequate. However, when the tip of the rods nears the upper cavity region (which features very low cross sections), diffusion solvers struggle to capture streaming through the cavity. With the diffusion solver, despite using directional diffusion coefficients in the cavity generated using the cumulative migration method [15], a nonphysical increase in the differential worth—from around 500 to 1450 pcm/m when the rod is moved from 7.5 to 9.25 m—is obtained. On the other hand, DFEM-S_N, which can natively handle void or near-void regions, gives a more satisfying control rod worth in that region, compared to the curve obtained with Monte Carlo (not shown here). Interestingly, the two solvers also diverge in their answer for very deeply inserted control rods. This is because, when the rod is nearly fully inserted, the scalar flux peaks towards the top of the core so the treatment of the upper cavity is once again important.

Overall, the benefit of the DFEM-S_N solver for shallowly and deeply inserted control rods is significant because (1) PBRs typically operate at full-power with shallowly inserted rods (so as to limit excess reactivity) and (2) the eigenvalue for fully inserted rods is necessary to accurately predict shutdown margin.

Another observation from Figure 8 is that the use of DFEM-S_N has little to no effect on control rod cusping—which is evidenced by the oscillatory behavior of the differential worth, despite using the decusping treatment in Griffin [27]. Additional improvements to the decusping treatment are planned for next fiscal year in order to ameliorate these oscillations.

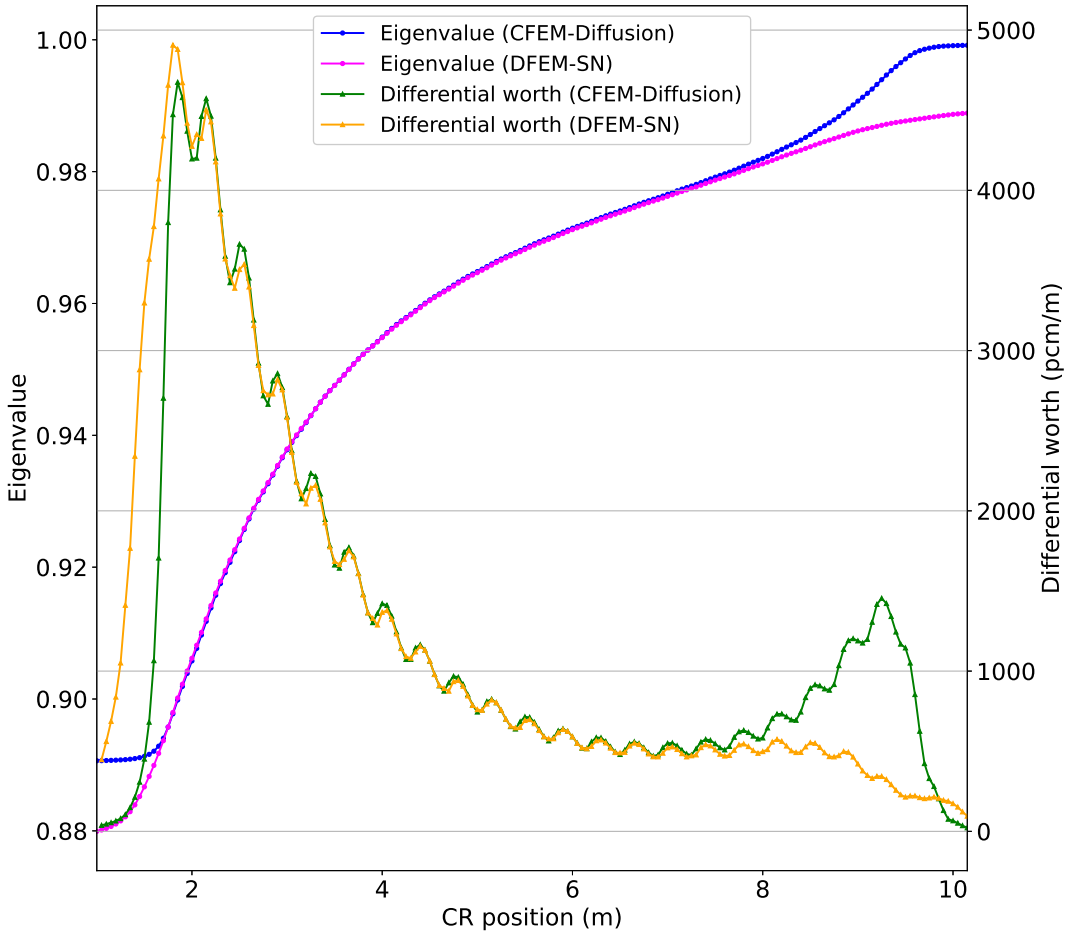

Figure 8: Reactivity curves of the gas-cooled PBR 200 model with equilibrium compositions and a uniform temperature profile (900 K) using CFEM-diffusion and DFEM-S_N (32 angles).

6. Multigroup Cross Section Generation for Pebble-Bed Reactors

In the past, the Griffin application relied on pregenerated cross sections using external codes, like Dragon 5 [17], at predefined reference states. These cross section sets were subsequently linearly interpolated during the multiphysics calculations. Previous work [20] showed that this approach can have significant limitations and that an online cross section generation approach with leakage correction is desirable for gas cooled pebble bed reactors. This section briefly describes the ongoing effort to develop the online cross section generation capability for PBRs. The goal is to generate accurate pebble-type-wise and burnup-group-wise pebble cross sections on the fly for pebble depletion taking into account the local effects (e.g., fuel and moderator temperature) and the long-range effects of spectral interpenetration [21].

The pebble-bed core is divided into several domains within which cross sections of pebbles in the same pebble type and burnup group are assumed to be reasonably similar. Each domain is called a "spectral zone," meaning that the zone is characterized by its representative spectrum. The determination of spectral zone boundaries resorts primarily to engineering judgement, although some studies have attempted to develop a more rigorous methodology [19].

For the cross section generation of a PBR, it is important to consider both the local self-shielding effect and core-wise spectral interpenetration effect owing to its neutronics characteristics. This is the reason why the spectral-zone-wise self-shielding calculation should be performed to capture the local effect and the information from the core transport calculation should be fed into the spectral zone-wise broad-group cross section generation. Taking these into account, the procedure is set up as described in the following subsection, followed by the verification results.

6.1 Online Cross Section Generation Procedure

A generalized procedure to generate cross sections online for the pebble streamline depletion calculation is shown in Figure 9. There are currently two streams of input data required for online cross section generation: the fine group ($\sim$400 groups) base cross section library and the MC^2-3 cross section library. The fine-group base multigroup library contains isotope-wise microscopic cross sections generated using NJOY [16]. The fine-group structure is assumed fine enough ($\sim$400 groups) to use the built-in Maxwellian-1/E-fission within-group spectrum for condensation (in

the GROUPR module). This fine-group library will be distributed as part of the Griffin data; although, in this work, we only generated cross sections for a limited number of isotopes for testing purposes. This number will be expanded next fiscal year to include 423 and 557 isotopes and metastable states for ENDF/B-VII.1 and ENDF/B-VIII.0, respectively. The cross sections in the resolved resonance energy range are replaced with region-dependent self-shielded ones calculated on the fly through the slowing-down calculation. Currently, the MC^2-3 library includes resolved and unresolved resonance parameters for reconstructing point-wise cross sections on the fly used in the slowing down calculation [4].

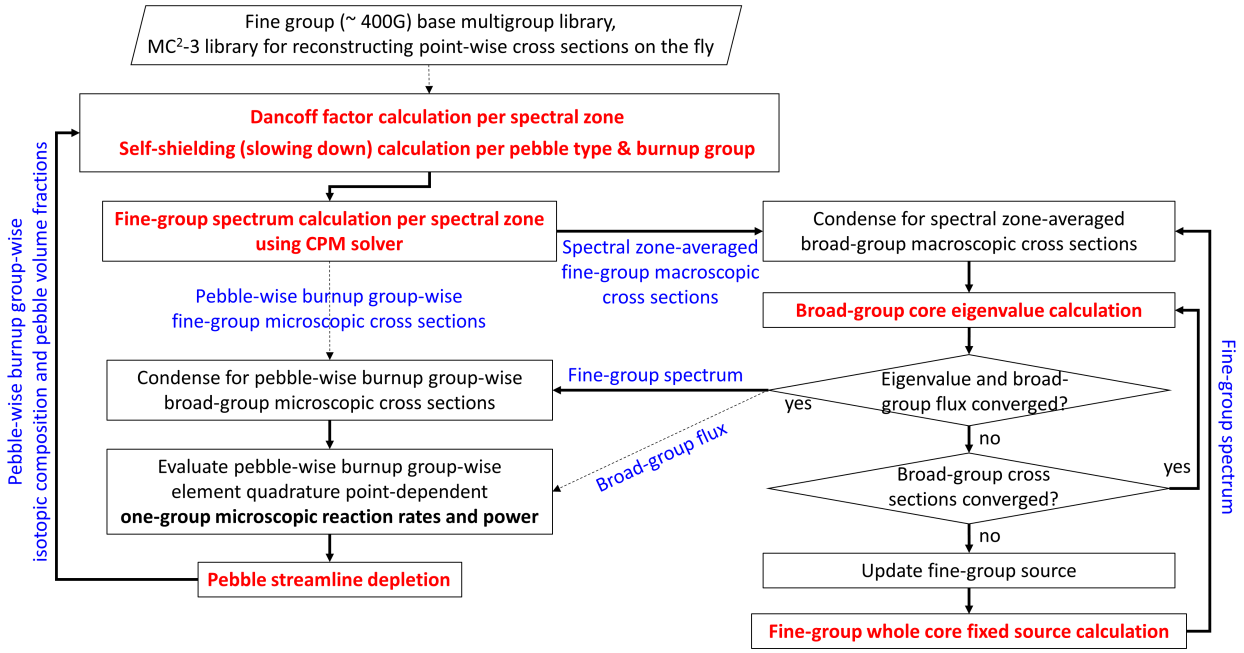

Figure 9: Generalized procedure for generating cross sections on the fly for a PBR, with bold arrows for the procedural sequence and dashed arrows, blue colored texts for required data, and red colored texts for main calculation modules. This procedure will be optimized next year to accelerate convergence for both the equilibrium core and the running-in calculation.

At a depletion timestep or an equilibrium core iteration when isotopic composition changes and pebble volume fractions are determined from the streamline depletion calculation, the self-shielding calculation is performed first to determine intraregion self-shielded fine-group microscopic cross sections within a spectral zone. This calculation consists of two steps: the pebble Dancoff factor calculation and two-step slowing-down calculations per pebble type and burnup group. The pebble Dancoff factor (not TRISO Dancoff factor) is determined using the dual sphere model of Ji's method [9] for each spectral zone. Once the Dancoff factor of each pebble is obtained,

two-step slowing-down calculations are performed for each pebble of fresh pebble type and each burnup group [23]: the first step for the double heterogeneity treatment for TRISO particles and the second step for the pebble, in which the fuel zone with TRISO particles is homogenized using the flux-volume weighting obtained from the first step. Pebble-wise slowing-down calculations are performed in parallel: one processor for one two-step slowing-down calculation.

This is followed by the fine-group spectrum calculation per spectral zone. Within each spectral zone, interactions among pebbles and spectral zone boundaries are taken into account using the collision probability method (CPM) [11], which is deployed to consider problems with multiple pebble types in order to improve results over the traditional approach using an "effective" pebble approach. The probability coefficients for within-pebble, pebble-to-pebble, pebble-to-coolant, pebble-to-boundary, and coolant-to-boundary collisions are obtained using a simple stochastic assumption that any sources from the coolant region, pebbles, or boundaries of a spectral zone are equally treated as isotropic and uniform sources in the coolant region. The number of degrees of freedom for this calculation is small; for example, the number of internal subdomains to reduce the spatial discretization error in a pebble is at maximum 40, the number of boundaries is six, the number of burnup groups $\sim$10, the number of pebble types $\sim$2, and the number of energy groups $\sim$400, resulting in approximate degrees of freedom of roughly $(40 \times (10 \times 2) + 6) \times 400 \sim 3 \times 10^5$ per spectral zone. These spectral-zone-wise spectrum calculations are performed in parallel: one processor per spectral zone. Having more processors than the number of spectral zones will let excessive processors be in idle and the opposite case will assign multiple spectrum calculations to each processor. The CPM calculation with this many degrees of freedom is generally not of much concern, compared to other calculation modules. Technically, this step provides the spatial self-shielding factor (the ratio of region-wise fluxes to the average flux) for homogenization: one for spectral-zone-averaged macroscopic cross sections for the neutron transport calculation and the other for pebble-type-wise and burnup-group-wise pebble-averaged microscopic cross sections for the burnup calculation. Condensation is treated by subsequent procedures.

The next step is to perform the core transport calculation to obtain the neutron flux and thus the reaction rates needed for depletion. This step is an iterative procedure involving a fine-group whole core fixed-source calculation performed at each eigenvalue iteration in the broad-group core transport calculation [6]. The former is to obtain a condensation spectrum for updating broad

group cross sections in spectral zones and reflector regions with consideration of the spectrum interpenetration effect. The fine-group fixed source is calculated using spectral zone-averaged broad-group-to-fine-group production cross sections of the following:

$$\begin{aligned}\overline{\Sigma}^{\ell}_{s,G\to h} &= \frac{\sum_k V_k \sum_{g'\in G} \Sigma^{\ell}_{\mathrm{s},k,g'\to h} f^{\mathrm{SS}}_{k,g'} \overline{\phi}^{\ell,\mathrm{FSP}}_{g'}}{V \sum_{g'\in G} \overline{\phi}^{\ell,\mathrm{FSP}}_{g'}} \text{ for scattering,} \\ \nu\overline{\Sigma}_{\mathrm{f},G\to h} &= \frac{\sum_k V_k \chi^{avg}_{k,h} \sum_{g'\in G} \nu\Sigma_{\mathrm{f},k,g'} f^{\mathrm{SS}}_{k,g'} \overline{\phi}^{0,\mathrm{FSP}}_{g'}}{V \sum_{g'\in G} \overline{\phi}^{0,\mathrm{FSP}}_{g'}} \text{ for fission,}\end{aligned} \tag{7}$$

and broad-group flux solution of the core transport calculation: $\overline{\phi}^{\ell m,\mathrm{Core}}_{G}$ as

$$\overline{S}_h(\vec{\Omega}) = \frac{1}{4\pi}\sum_G \nu\overline{\Sigma}_{f,G\to h}\overline{\phi}^{00,\mathrm{Core}}_{G} + \sum_{\ell,m} Y^{\ell m}(\Omega) \sum_G \overline{\Sigma}^{\ell}_{s,G\to h}\overline{\phi}^{\ell m,\mathrm{Core}}_{G}. \tag{8}$$

where the horizontal line above characters means the spectral zone-averaged quantity, k is a region index within a spectral zone, h and g are fine group indices, G is a broad group index, ℓ is a Legendre order, SS stands for the quantity obtained from the spectral zone-wise spectrum calculation, FSP (fixed source problem) and $Core$ stand for the fine-group fixed source and broad-group core transport calculations, respectively, f_k is the ratio of the flux at region k within a spectral zone to the spectral zone-averaged flux, and Σ_k is the macroscopic self-shielded cross section at region k within a spectral zone. $f^{\mathrm{SS}}_{k,g}$ is obtained from the latest spectral zone-wise spectrum calculation and is fixed during iteration. $\overline{\phi}^{\ell,\mathrm{FSP}}_{g} \equiv \int_{\mathrm{SZ}} \sqrt{\sum_{m=-\ell}^{\ell}\left(\overline{\phi}^{\ell m,\mathrm{FSP}}_{g}(\vec{r})\right)^2}\, d\vec{r} \Big/ V_{\mathrm{SZ}}$ are the latest flux solution during the iteration, which is initialized at the very beginning of the iteration with $\overline{\phi}^{\ell,\mathrm{SS}}_{g}$.

The fine-group fixed-source calculation is performed until the broad-group cross sections are converged. Once converged, the pebble-type-wise and burnup-group-wise microscopic broad-group cross sections are updated as:

$$\tilde{\sigma}^{i}_{G,x,c,l} = \frac{\sum_{g\in G} \tilde{\sigma}^{i,\mathrm{SS}}_{g,x,c,l} \tilde{f}^{\mathrm{SS}}_{g,c,l} \overline{\phi}^{0,\mathrm{FSP}}_{g}}{\sum_{g\in G} \tilde{f}^{\mathrm{SS}}_{g,c,l} \overline{\phi}^{0,\mathrm{FSP}}_{g}}. \tag{9}$$

where the tilde sign above characters means the pebble-averaged quantity, i is an isotope index,

x is a reaction index, c and l are indices for fresh pebble type and burnup group, respectively, $\tilde{\sigma}$ is the pebble-averaged microscopic cross section obtained using number-density-flux-volume weighting, and $\tilde{f}$ is the ratio of the pebble-averaged flux to the spectral-zone-averaged flux obtained from the spectral-zone-wise spectrum calculation.

The core eigenvalue iteration continues until the eigenvalue and flux solutions are converged. Once converged, pebble-type-wise and burnup-group-wise one-group microscopic reaction rates for all isotopes are calculated per quadrature point or element in a neutronics mesh by the Griffin material system using the broad group flux solution as:

$$R^i_{x,c,l}(\vec{r}) = \sum_G \tilde{\sigma}^i_{G,x,c,l} \left(\sum_{g \in G} \tilde{f}^{\text{SS}}_{g,c,l} \overline{w}^{\text{FSP}}_g \right) \phi^{0,\text{Core}}_G(\vec{r}), \tag{10}$$

where $\overline{w}^{\text{FSP}}_{g \in G} = \overline{\varphi}^{0,\text{FSP}}_g / \sum_{g \in G} \overline{\varphi}^{0,\text{FSP}}_g$ is the condensation weighting factor within the spectral zone where the point $\vec{r}$ belongs. Note that the broad-group cross section is computed using the within-fine-group spectrum solution of the fixed-source calculation, and the one-group reaction rate is computed using the broad-group flux of the core-eigenvalue calculation. The reason the fine-group spectrum solution of the fixed-source calculation is not directly used to compute one-group reaction rates from fine-group cross sections is that the fixed-source calculation is only for obtaining the ***within***-fine-group spectrum in an approximated way. It is not guaranteed that the fine-group spectrum in the whole energy range represents the solution of the core-eigenvalue calculation.

These one group reaction rates are transferred to the pebble streamline mesh and are used for depletion. Note that pebble power density does not need to be transferred since it can be calculated in a streamline mesh using the transferred microscopic one-group "kappa fission" reaction rates and isotopic number densities.

6.2 Verification Tests

At the time of writing, the leakage correction part using the fine-group whole-core fixed-source calculation is still being implemented and tested. Thus, this section only includes the results for infinite spectral zone problems and a simple 1D core problem in which the fine group (400 groups) structure was used without condensation. This verification work does not include the pebble

streamline depletion.

6.2.1 Infinite Spectral Zone Problems

The simplest problem was devised such that pebbles are stacked in an infinitely large space. This problem is mainly to verify the CPM solver for pebble-to-pebble interactions. For generating the reference solution, a discrete element method based spherical particle dynamics code, cemfDEM [1], was used to simulate 5,147 pebbles in a 1 m cube. For providing other comparable references, body centered cubic (BCC) and face centered cubic (FCC) regular arrangements were also used for generating Serpent solutions. BCC and FCC configurations were generated to preserve the pebble packing fraction determined from the random model from cemfDEM by adjusting a gap between pebbles.

In a real PBR, fuel pebbles are normally mixed with graphite dummy pebbles. The fuel pebble has a 5-cm-diameter fuel zone in which TRISO particles are randomly bound in a graphite matrix and the fuel zone is surrounded by 5-mm-thick graphite layer. The graphite dummy pebble is a 6-cm-diameter pure graphite sphere. The fuel kernel of TRISO is 17 wt% UO_2 fuel. Three different fuel-to-dummy pebble ratios were tested, 100:0, 75:25, and 50:50, and each ratio has three different cases with varying TRISO packing fractions: 2.4, 11.2, and 52.0%. This is for testing the developed capability in a wide range of spectra.The temperature in all regions is room temperature (293.6 K).

Figure 10 compiles the deviation of eigenvalue results of different Serpent models in pcm with respect to the Griffin results. Serpent results of FCC configurations can vary up to 1000 pcm in K_{eff} ($\sim$500 pcm in reactivity) for high TRISO packing fraction cases mainly due to different extents of the pebble-to-pebble shadowing effect. It was confirmed that the highest reactivity case among FCC results, for example, in 50:50 with the 52.0% packing fraction, came from the configuration where the fuel pebbles are most clustered (the right-most one in Figure 11), and the lowest reactivity case came from the configuration where fuel pebbles are the most uniformly arranged (the middle one in Figure 11). The BCC configuration for the 50:50 case with the 52.0% packing fraction is the perfectly uniform arrangement of fuel pebbles as shown in Figure 11 and its result is very close to the least clustered case of FCC. Pebble clustering makes the resonance absorption at the fuel kernel surface in the inner pebbles occur less than in a regular fuel pebble arrangement due to the shadowing effect of surrounding pebbles.

This is supported by Table 5 and Figure 12 which compare U-238 capture rates among different models for the fuel:dummy=50:50 case with the 52% packing fraction. In Table 5, about 400 pcm difference in reactivity (0.17352 - 0.16947; absolute difference in capture rate per fission neutron is directly translated to reactivity difference) between the least and most clustered cases is solely attributed to U-238 capture rate difference caused by fuel pebble arrangement. The BCC result is closer to that of the minimum K_{eff} case of FCC and the average result of the random configurations is somewhat closer to that of the maximum K_{eff} case of FCC. The Griffin result is in between the average result of FCC and BCC. Figure 12 shows the absolute difference in the number of neutrons captured by U-238 per 10^5 fission neutrons among different models and codes at different incident energies. The reference is the average result of FCC. With respect to the FCC average values, the most clustered configuration gives less capture than the least clustered configuration at resonances, which supports that eigenvalue differences in different configurations mainly come from different extents of resonance capture by U-238 due to the pebble-to-pebble shadowing effect. The BCC and random results are closer to that of the least and most clustered FCC, respectively, which is consistent with Table 5. The Griffin result shows larger changes in the differences, resulting in the net total value somewhere between BCC and FCC average results.

In the eigenvalue results, Griffin results generally agree with the average BCC results within a reactivity difference of 200 pcm even though energy-wise reaction rates do not perfectly follow those of BCC. Nonetheless, these results can be understood from the assumption used in the pebble-bed CPM solver of Griffin that transmission probabilities among the same fuel pebble types are the same, which does not consider the clustering effect. The reactivity deviations among random configurations is generally small in all cases, and their averages are generally positioned in higher reactivity cases for high TRISO packing fraction (i.e., more [fuel pebble] clustered cases) of FCC configurations. This result indicates that the actual pebble configurations randomly stacked in a pebble-bed core inevitably have the effect of fuel pebble clustering. Therefore, the pebble-bed CPM solver needs to be able to consider the fuel pebble clustering effect for a realistic random pebble stack. Work was done in the past to consider this effect in the APPOLO code [8].

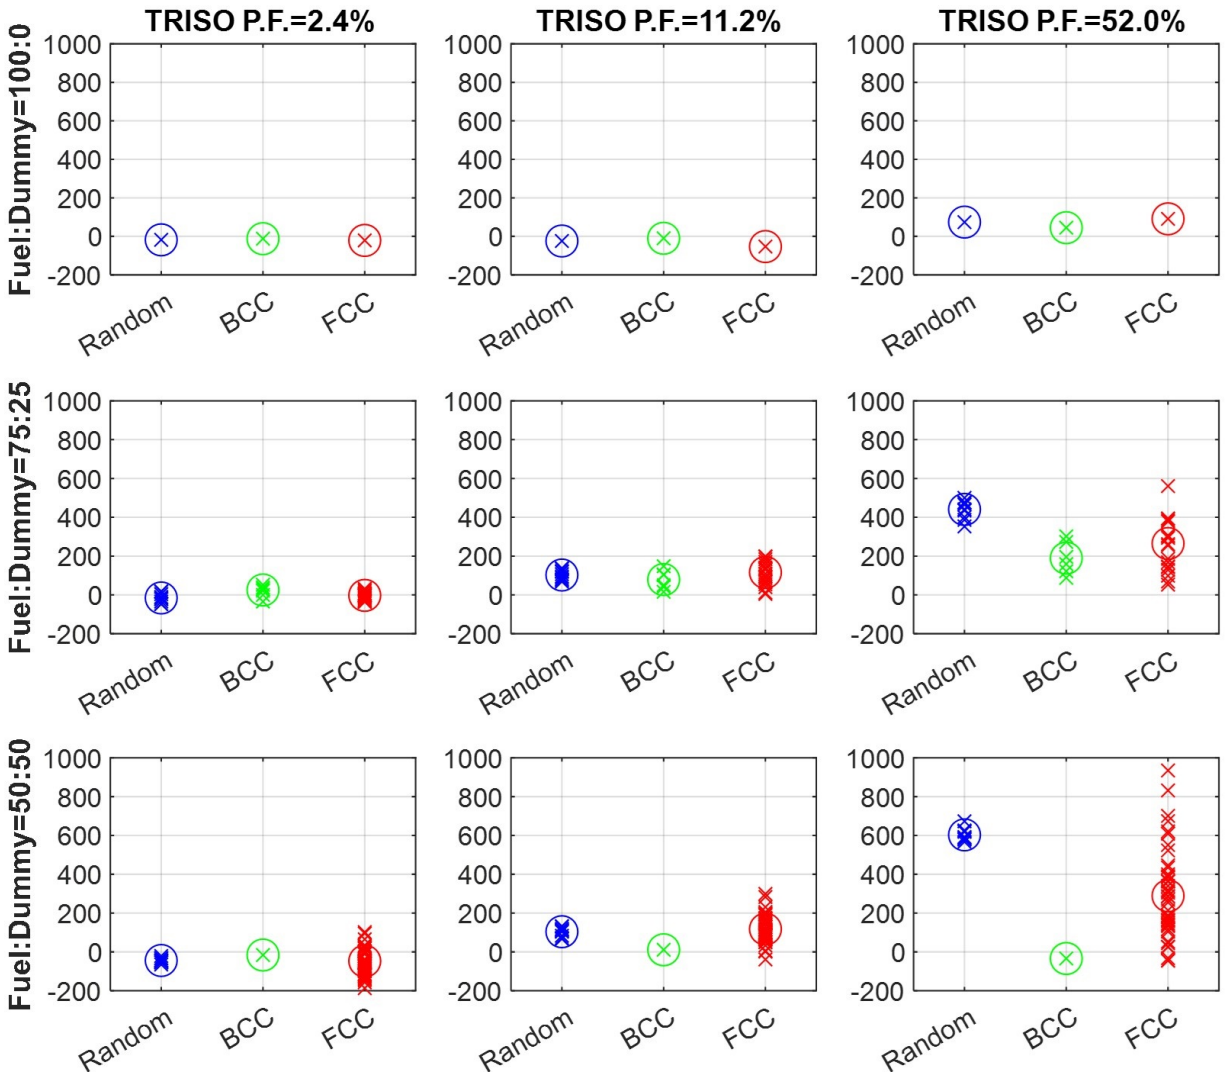

Figure 10: Deviation of K_{eff} (pcm) results of different Serpent models with respect to (w.r.t.) Griffin results, where X marks show individual Serpent results and circle marks are average values.

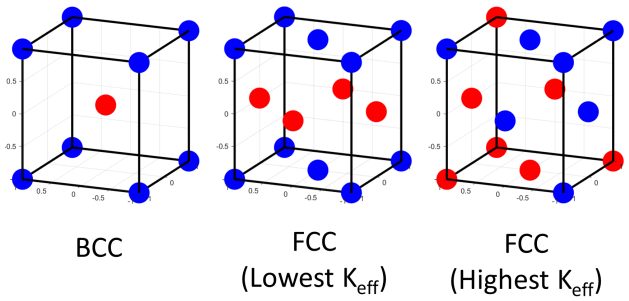

Figure 11: Pebble arrangements for the fuel:dummy=50:50 case. Red: fuel pebble, Blue: graphite pebble. FCC with lowest K_{eff} = least clustered, and FCC with highest K_{eff} = most clustered.

Table 5: Total number of neutrons captured by U-238 per fission neutron for fuel:dummy=50:50 with P.F.=52%.

Serpent					Griffin
BCC	FCC Avg.	FCC Max. K_{eff}	FCC Min K_{eff}	Random Avg.	
0.17341	0.17215	0.16947	0.17352	0.17091	0.17286

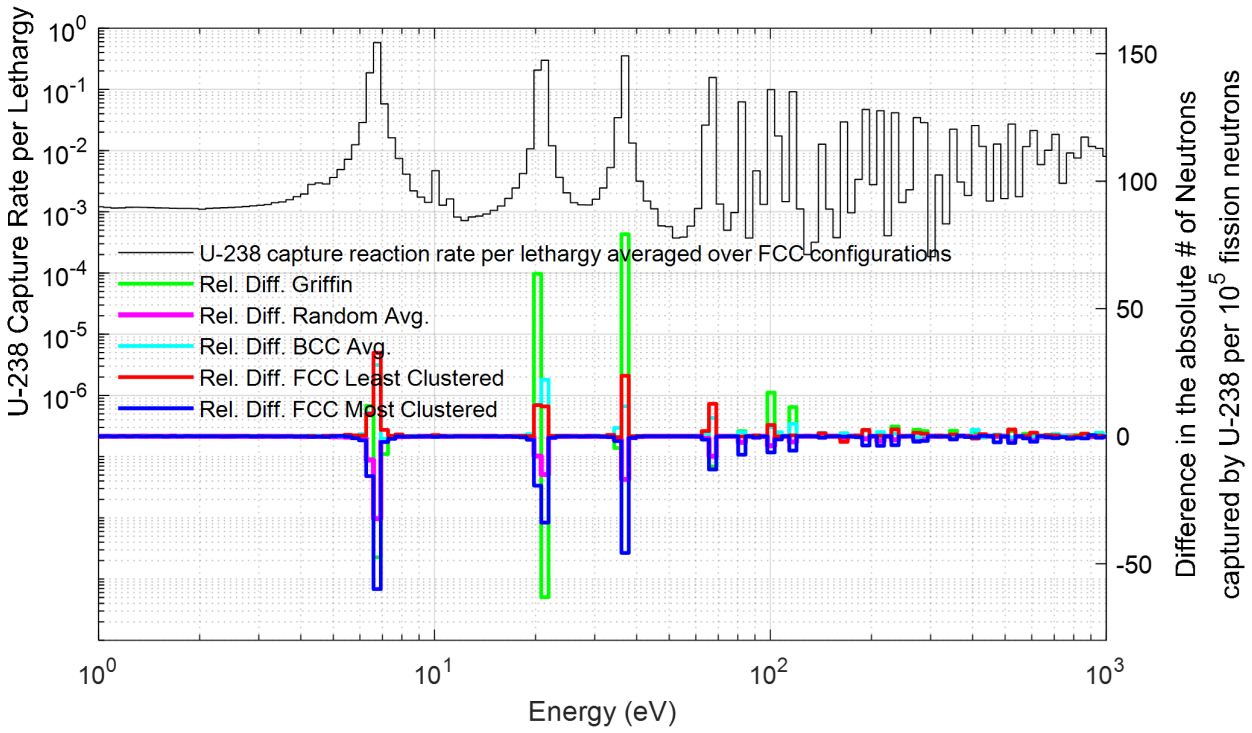

Figure 12: Comparison of U-238 capture rates of Griffin with those of different pebble arrangements in Serpent models for fuel:dummy=50:50 with P.F.=52%. Reference: FCC average.

Additional problems with multiple fuel pebble types were also solved. A 4.25 wt% UO_2 fuel pebble was added to the 50:50 ratio case of the above problem. Packing fractions are assumed to be the same in two fuel pebble types. Two ratios between 17 and 4.25 wt% UO_2 fuel pebbles were tried: 1:1 and 1:3.

The results are shown in Table 6 and are consistent with those shown for the single fuel-pebble-type problem. Griffin results agree the best with the BCC Serpent model and are underestimated against the random Serpent model. Future work will be to solve more challenging but realistic configurations that have burned fuel pebbles mixed with fresh pebbles.

Table 6: Results of multiple fuel-pebble-type problems.

17%:4.25% Ratio	Packing Fraction	Serpent			Griffin			
		K_{eff}			K_{eff}	ΔK_{eff} (pcm)		
		BCC	FCC	Random		w.r.t. BCC	w.r.t. FCC	w.r.t. Random
1:1	2.4%	1.65876	1.66089	1.65913	1.65998	+122	-91	+85
	11.2%	1.66991	1.67010	1.66970	1.66927	-64	-83	-43
	52.0%	1.35539	1.35797	1.36058	1.35518	-21	-279	-540
1:3	2.4%	1.54962	1.54868	1.54883	1.54969	+7	+101	+86
	11.2%	1.62175	1.62167	1.62113	1.62092	-83	-75	-21
	52.0%	1.32221	1.32471	1.32660	1.32178	-43	-293	-482

6.2.2 One-Dimensional Core Problem

Based on the results of the infinite spectral zone problem, a simple 1D slab core problem was devised using the same compositions of 17 wt% UO_2 fuel with 5% TRISO packing fractions. The center 2.95-m-wide pebble-bed region is surrounded by 1-m-wide solid graphite reflectors on both the left and right. The pebble-bed region is composed of 40 repeated BCC cells with a 7.375 cm pitch, resulting in the 2.95 m total length. This configuration leads to graphite pebbles cut in half at the interface of the pebble-bed and reflector regions. The vacuum boundary condition was used on the left and right boundaries.

The Serpent run used 100 inactive and 1,000 active cycles with 100,000 histories per cycle. The Griffin calculation used the 400 group structure directly for the core-eigenvalue calculation. The pebble-bed region has only one cross section region since there is currently no way to consider the core environmental effect on the Dancoff factor and carrying multiple cross section regions won't make a difference in the fine-group cross section sets in different cross section regions. The Serpent and Griffin eigenvalues are 1.69266 and 1.69584, respectively, thus the eigenvalue error of Griffin is +111 pcm. The power distribution of Griffin is -2.5% lower at the periphery of the pebble-bed region and +1.5% higher at the core center, as shown in Figure 13.

The accuracy of the eigenvalue and power distribution results is acceptable; however, we have identified issues in the graphite scattering resonance in the MeV energy range. When using the 400 group cross section library, SSAPI does not apply the self-shielding treatment to the cross sections in such a high energy range. Instead, it relies solely on a fission spectrum as a weighting function, resulting in an infinite dilute condition and thus an overestimation of cross sections. This causes leakage errors for the fast energy range, which is important for core problems. Addressing this

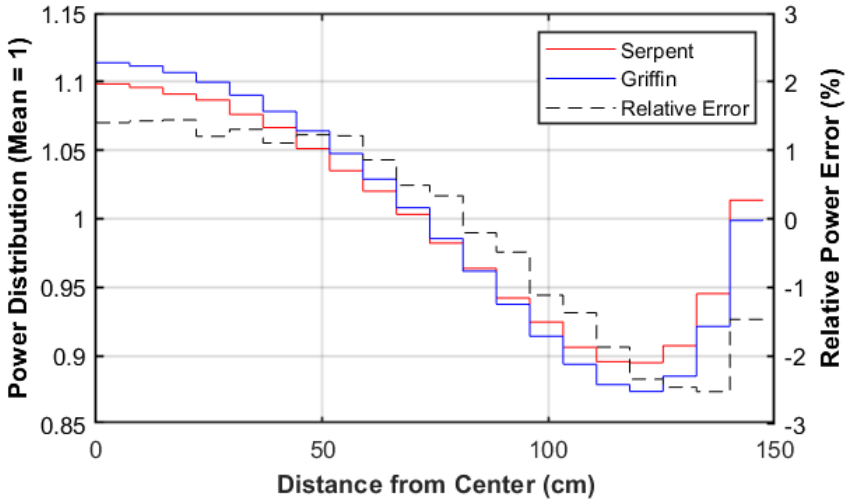

Figure 13: Relative error (%) in power distribution of Griffin w.r.t. that of Serpent.

issue will require more detailed analysis in the future.

7. CONCLUSIONS AND FUTURE WORK

7.1 Conclusions

While Griffin has been used extensively for PBR modeling, previous work has all focused on the direct computation of an equilibrium core. This work is the first documentation of new technologies in Griffin for simulating the running-in phase of PBR operation. While the existing direct equilibrium core capability provided a solid foundation for developing the running-in capability, significant additional work was required as the existing equilibrium core capability had no need to account for the time variability of any quantity in the core. Verification of the new running-in capability is provided in this report as well as a sample running-in calculation. While the running-in simulation presented in this work used a simplified model, it showed that Griffin is able to simulate years of operation during the running-in phase efficiently with each depletion step taking only several seconds. Reactivity is maintained over years of simulation time with different feed regimes for fuel pebbles and a steady state is eventually reached. In the future, more realistic models will be used incorporating things like realistic temperatures for feed back and control rods.

Additionally, a pebble-bed online cross section generation capability has been developed in

Griffin and verified against infinite spectral zone and 1D core problems. The design for integrating the online cross section capability with the streamline depletion capability including equilibrium and running-in simulations is proposed in this work. When this feature is fully integrated into Griffin, it will facilitate accurate calculations avoiding the cumbersome and difficult step of pre-generating a separate standalone multigroup cross section set.

7.2 Future Work

Although Section 3.1 showed that the CFL condition still allows timesteps on the order of days, a relatively small timestep was used in this work for the sample running-in simulation. Larger timesteps will be investigated in the future. Related to this, results with a predictor-corrector method should be compared with the forward Euler results presented in this work. Especially when using larger timesteps, a predictor-corrector scheme may be important for achieving reasonable accuracy. This will become especially important for optimization work where thousands of simulations might be run with varied parameters.

This work did not include results with any coupling between Griffin and other applications for a multiphysics running-in simulation. The sample results were generated from a highly simplified model. The ease of coupling such applications for realistic simulations is a major benefit of the MOOSE framework, and such coupled models are currently being constructed with further analysis to follow. Validating the running-in capability is an ongoing effort related to the ongoing validation of the direct equilibrium core calculation capability. The technologies presented in the last two sections of this report will be fully integrated into the running-in simulation workflow and be important in achieving physically realistic stimulation results.

Section 3.4 included a brief investigation into the importance of numerical diffusion to the running-in calculation. Ultimately, a more sophisticated numerical method with less numerical diffusion will be implemented to better understand the impact numerical diffusion has on results.

Analyses with the infinite spectral zone problems identified the need for considering the fuel pebble clustering effect for reaction rate improvement. Future work will conduct verification against more realistic test problems by integrating streamline depletion and online cross section generation and investigate performance and accuracy improvements of the self-shielding calculation and broad-group core calculation with more sophisticated leakage correction models.

REFERENCES

[1] *Coupled CFD-DEM Modeling: Formulation, Implementation and Application to Multiphase Flows*, John Wiley Sons, Ltd., doi:10.1002/9781119005315, 2016.

[2] Andreades, C., et al., Design summary of the mark-i pebble-bed, fluoride salt–cooled, high-temperature reactor commercial power plant, *Nuclear Technology*, *195*(3), 223–238, doi:10.13182/NT16-2, 2016.

[3] Balestra, Paolo, Schunert, Sebastian, Carlsen, Robert W, Novak, April J, DeHart, Mark D, and Martineau, Richard C, Pbmr-400 benchmark solution of exercise 1 and 2 using the moose based applications: Mammoth, pronghorn, *EPJ Web Conf.*, *247*, 06,020, doi:10.1051/epjconf/202124706020, 2021.

[4] C. H. Lee and W. S. Yang, MC^2-3: Multigroup cross section generation code for fast reactor analysis, *Nucl Sci Eng*, *187*(ANL/NSE-18/12, Revision 0), 268–290, doi:10.1080/00295639.2017.1320893, 2017.

[5] Calvin, O., S. Schunert, and B. Ganapol, Global error analysis of the chebyshev rational approximation method, *Annals of Nuclear Energy*, *150*, 107,828, doi:https://doi.org/10.1016/j.anucene.2020.107828, 2021.

[6] Douglass, S., and F. Rahnema, Subgroup decomposition method, *Annals of Nuclear Energy*, *48*, 84–101, doi:10.1016/j.anucene.2012.04.029, 2012.

[7] Giudicelli, G. L., et al., The virtual test bed (vtb) repository: A library of reference reactor models using neams tools, *Nuclear Science and Engineering*, *197*(8), 2217–2233, doi:10.1080/00295639.2022.2142440, 2023.

[8] Grimod, M., Neutronic modeling of pebble bed reactors in apollo2, Ph.D. thesis, Universite Paris-Saclay, 2010.

[9] Ji, W., C. Liang, and E. N. Pusateri, Analytical dancoff factor evaluation for reactor designs loaded with triso particle fuel, *Annals of Nuclear Energy*, *63*, 665–673, doi:10.1016/j.anucene.2013.09.025, 2014.

[10] Kadak, A. C., A future for nuclear energy: pebble bed reactors, *International Journal of Critical Infrastructures*, *1*(4), 330–345, doi:10.1504/IJCIS.2005.006679, 2005.

[11] Kulik, V., and R. Sanchez, Core homogenization method for pebble bed reactors, in *Mathematics and Computation, Supercomputing, Reactor Physics and Nuclear and Biological Applications*, American Nuclear Society, 2005.

[12] Labouré, V., P. Balestra, J. Ortensi, J. Hanophy, Z. Prince, R. Stewart, and G. Strydom, Multiphysics pebble-bed reactor control rod withdrawal study, *Tech. Rep. INL/RPT-23-74341*, Idaho National Laboratory, 2023.

[13] Lewis, E. E., and W. F. Miller, *Computational Methods of Neutron Transport*, John Wiley & Sons, 1984.

[14] Lindsay, A. D., et al., 2.0 - MOOSE: Enabling massively parallel multiphysics simulation, *SoftwareX*, *20*, 101,202, doi:https://doi.org/10.1016/j.softx.2022.101202, 2022.

[15] Liu, Z., K. Smith, B. Forget, and J. Ortensi, Cumulative migration method for computing rigorous diffusion coefficients and transport cross sections from monte carlo, *Annals of Nuclear Energy*, *112*, 507–516, doi:https://doi.org/10.1016/j.anucene.2017.10.039, 2018.

[16] MacFarlane, R. E., D. W. Muir, R. M. Boicourt, A. C. Kahler, J. L. Conlin, and W. Haeck, The NJOY nuclear data processing system, version 2016, *Tech. rep.*, Los Alamos National Laboratory, 2016.

[17] Marleau, G., A. Hebert, and R. Roy, A user guide for dragon version5, *Tech. Rep. IGE–335*, Ecole Polytechnique de Montreal, 2020.

[18] Morel, J. E., and G. R. Montry, Analysis and elimination of the discrete-ordinates flux dip, *Transport Theory and Statistical Physics*, *13*(5), 615–633, doi:10.1080/00411458408211661, 1984.

[19] Mphahlele, R., A. M. Ougouag, K. N. Ivanov, and H. D. Gougar, Spectral zone selection methodology for pebble bed reactors, *Annals of Nuclear Energy*, *38*(1), 80–87, doi:https://doi.org/10.1016/j.anucene.2010.08.014, 2011.

[20] Ortensi, J., and P. Balestra, Initial study on cross-section generation requirements for a PBR equilibrium core, in *Proceedings of the international conference on physics of reactors-PHYSOR 2022*, pp. 1440–1453, Pittsburgh, PA, 2022.

[21] Ortensi, J., J. J. Cogliati, M. A. Pope, R. M. Ferrer, and A. M. Ougouag, Deterministic modeling of the high temperature test reactor, doi:10.2172/989875.

[22] Ortensi, J., C. M. Mueller, S. Terlizzi, G. Giudicelli, , and S. Schunert, Fluoride-cooled high-temperature pebble-bed reactor reference plant model, *Tech. Rep. INL/RPT-23-72727*, Idaho National Laboratory, 2023.

[23] Park, H., W. S. Yang, and C. Lee, Particulate fuel modeling of mc2-3 using iterative local spatial self-shielding method, *Annals of Nuclear Energy*, *172*, 109,060, doi:10.1016/j.anucene.2022.109060, 2022.

[24] Prince, Z. M., C. Brennan, M. Turkmen, P. Balestra, and G. Strydom, Reduced order models generation for htgrs pebble shuffling procedure optimization studies, *Tech. Rep. INL/RPT-22-68865*, Idaho National Laboratory, 2022.

[25] Pusa, M., and J. Leppänen, Computing the matrix exponential in burnup calculations, *Nuclear Science and Engineering*, *164*(2), 140–150, doi:10.13182/NSE09-14, 2010.

[26] Rycroft, C. H., A. Dehbi, T. Lind, and S. Güntay, Granular flow in pebble-bed nuclear reactors: Scaling, dust generation, and stress, *Nuclear Engineering and Design*, *265*, 69–84, doi:https://doi.org/10.1016/j.nucengdes.2013.07.010, 2013.

[27] Schunert, S., Y. Wang, J. Ortensi, V. M. Labouré, F. N. Gleicher, M. D. DeHart, and R. C. Martineau, Control rod treatment for fem based radiation transport methods, *Annals of Nuclear Energy*, *127*, 293–302, doi:10.1016/j.anucene.2018.11.054, 2019.

[28] Schunert, S., J. Ortensi, Y. Wang, P. Balestra, M. Jaradat, O. Calvin, J. Hanophy, and L. Harbour, An equilibrium core depletion algorithm for pebble-bed reactors in the griffin code, *Annals of Nuclear Energy*, *192*, 109,980, doi:https://doi.org/10.1016/j.anucene.2023.109980, 2023.

[29] Stewart, R., P. Balestra, D. Reger, and E. Merzari, Generation of localized reactor point kinetics parameters using coupled neutronic and thermal fluid models for pebble-bed reactor transient analysis, *Annals of Nuclear Energy*, *174*, 109,143, doi:https://doi.org/10.1016/j.anucene.2022.109143, 2022.

[30] Stewart, R., P. Balestra, D. Reger, and G. Strydom, Novel pebble shuffling capabilities to generate reduced order models training databases, *Tech. Rep. INL/RPT-22-68918*, Idaho National Laboratory, 2022.

[31] Wang, Y., Z. M. Prince, J. T. Hanophy, C. Lee, Y. S. Jung, H. Park, L. H. Harbour, and J. Ortensi, Performance improvements for the griffin transport solvers, doi:10.2172/1822446.

[32] Wang, Y., S. Schunert, V. Laboure, and Z. Prince, Rattlesnake theory manual, *Tech. Rep. INL/EXT-17-42103*, Idaho National Laboratory, 2018.

[33] Wang, Y., et al., Rattlesnake: A moose-based multiphysics multischeme radiation transport application, *Nuclear Technology*, *207*(7), 1047–1072, doi:10.1080/00295450.2020.1843348, 2021.

[34] Zhang, J., F. Li, and Y. Sun, Physical analysis of the initial core and running-in phase for pebble-bed reactor htr-pm, *Science and Technology of Nuclear Installations*, *2017*, 2017.