



# Automated Control for Nuclear Thermal Propulsion Start-Up using MOOSE- based Applications

August 2022

*Changing the World's Energy Future*

Vincent M Laboure, Sebastian Schunert, Stefano Terlizzi, Zachary M Prince,  
Javier Ortensi, Ching-Sheng Lin, Lise Cecile Madeleine Charlot, Mark D  
DeHart



#### **DISCLAIMER**

This information was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness, of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

# **Automated Control for Nuclear Thermal Propulsion Start-Up using MOOSE-based Applications**

**Vincent M Laboure, Sebastian Schunert, Stefano Terlizzi, Zachary M Prince,  
Javier Ortensi, Ching-Sheng Lin, Lise Cecile Madeleine Charlot, Mark D DeHart**

**August 2022**

**Idaho National Laboratory  
Idaho Falls, Idaho 83415**

**<http://www.inl.gov>**

**Prepared for the  
U.S. Department of Energy  
Under DOE Idaho Operations Office  
Contract DE-AC07-05ID14517, DE-AC07-05ID14517**

*Page intentionally left blank*



## SUMMARY

This report presents a Griffin/Bison/RELAP-7 numerical model of a prototypical nuclear thermal propulsion (NTP) system that features fuel assemblies arranged in rings, and which was designed to simulate rapid startup transients. The physics modeled include full-core neutronics, assembly-wise heat conduction, and conjugate heat transfer, with the balance of plant mainly imposed through boundary conditions. In addition, various forms of automated reactivity control were deployed by using the Multiphysics Object-Oriented Simulation Environment (MOOSE) to autonomously drive the model and simulate the reactor transitioning from assumed initial conditions to nominal power in a fraction of a minute.

To generate the cross-sections of the neutronics model, and in an effort to simultaneously account for the tremendous axial temperature gradients in the reactor and to limit the number of state points required for cross-section generation, the average component temperatures and hydrogen densities in the cooling channels were correlated to the average fuel and moderator temperatures, and fixed axial profiles were derived for nominal conditions and then used during the transient. With this approximation, a tractable cross-section library tabulated with fuel/moderator temperatures and control drum (CD) angles was generated using Serpent. The full-core super homogenization (SPH) correction procedure and the CD decussing technology in Griffin, respectively, ensure preservation of the multiplication factor and reaction rates at state points, along with a reasonably accurate reactivity worth between tabulated CD angles, despite using a coarse mesh. Feedback from other physics was calculated by modeling one representative fuel assembly per ring, along with the corresponding fuel and moderator cooling channels.

To limit power overshoots during startup, another layer of multiphysics coupling was added to the model in order to automatically control the drums. Two different technologies presented herein showed outstanding performance in this regard: (1) a novel hybrid proportional integral derivative (PID) controller based on both power and reactivity signals, and (2) a period-generated controller (PGC) that relies on kinetics parameters and reactivity coefficients to predict future behavior and adjust the desired signal accordingly. A challenging benchmark was devised, featuring a power demand curve that exponentially increases by a factor of 500 within 30 seconds, then levels out after that. Both control approaches create a simulated power curve that closely follows the power demand curve and limits power overshoots to 1% or less. While the former approach requires more tuning of the internal parameters, the latter requires additional knowledge of the reactivity feedback coefficients and rates of change of the corresponding variables, including fuel and moderator temperature, which could be difficult to dynamically measure for a real NTP system. Fortunately, some inaccuracy in these quantities will not drastically degrade the PGC performance.

Subsequently, a more realistic startup sequence was considered, in which the mass flow rate and outlet pressures are ramped up to model bootstrap and thrust build-up phases prior to reaching steady-state conditions, demonstrating the ability of the hybrid PID and PGCs to handle such transients, with both types of controllers exhibiting very similar behavior. Nevertheless, a significant chamber temperature overshoot was observed, caused by the demanded power signal and assumed mass flow rate. This issue could be mitigated by deploying a reactor controller that follows the chamber temperature signal and actuates both the control valves and drums (rather than using a power signal based solely on the drums to control reactivity). Enhancement of the hydrogen fluid properties available in MOOSE, as well as a better understanding of prototypical initial conditions, are also needed to further enhance this startup model.

Finally, a study was performed to model decay heat post-shutdown, and to prepare for extending this model to predict shutdown behavior and post-shutdown pulsed cooling requirements.

*Page intentionally left blank*

# CONTENTS

SUMMARY .....	i
1 INTRODUCTION .....	1
2 PROTOTYPICAL DESIGN.....	2
3 FULL-CORE MULTIPHYSICS MODEL .....	6
3.1 Neutronics .....	7
3.1.1 Cross-Section Generation .....	7
3.1.2 SPH Correction .....	11
3.1.3 CD Rotation .....	12
3.1.4 Prompt Power and Decay Heat .....	13
3.2 Fuel Assembly Heat Transfer .....	15
3.2.1 Approach .....	15
3.2.2 Power Density Conversion .....	16
3.3 Thermal-Hydraulics .....	17
3.3.1 Representative Cooling Channels .....	17
3.3.2 System Model.....	18
3.4 Multiphysics Coupling.....	21
3.4.1 Data Transfer.....	21
3.4.2 Tight Coupling .....	23
3.4.3 Workflow .....	23
4 REACTIVITY CONTROL STRATEGIES .....	24
4.1 Exponential Ramp-up Benchmark .....	24
4.2 PID Control of the Drum Rotation .....	25
4.2.1 Power-driven PID Signal .....	25
4.2.2 Reactivity-driven PID Signal .....	27
4.2.3 Hybrid Power-Reactivity PID Signal.....	31
4.3 Period-generated Control of the Drum Rotation .....	34
4.3.1 Methodology.....	34
4.3.2 Reactivity Feedback.....	36
4.3.3 Results .....	38
4.4 Summary .....	40
5 STARTUP RESULTS .....	42
5.1 Assumed Startup Sequence .....	42

5.2	Multiphase Startup Results .....	45
5.3	Computational Cost .....	48
6	PRELIMINARY DECAY HEAT CURVE DETERMINATION FOR SHUTDOWN TRANSIENTS .....	49
6.1	Computation of Decay Heat Curves .....	49
6.2	Simple Model for Decay Heat Curves .....	50
7	CONCLUSIONS AND RECOMMENDATIONS .....	51
7.1	Conclusions .....	51
7.2	Future Work .....	53
8	Acknowledgments .....	55
	References .....	55
	Appendix A Attached File Summary .....	58
A.1	Input Files .....	58
A.2	Mesh Files .....	59
A.3	Cross-Section Files .....	59

*Page intentionally left blank*

## ACRONYMS

<b>CD</b>	control drum
<b>CERMET</b>	CERamic-METallic
<b>DHC</b>	decay heat curve
<b>MARE</b>	mean absolute relative error
<b>MOOSE</b>	Multiphysics Object-Oriented Simulation Environment
<b>NASA</b>	National Aeronautics and Space Administra- tion
<b>NTP</b>	nuclear thermal propulsion
<b>PGC</b>	period-generated controller
<b>PID</b>	proportional integral derivative
<b>PKE</b>	point kinetics equation
<b>RMSRE</b>	root mean square relative error
<b>SPH</b>	super homogenization
<b>SSCV</b>	structural support control valves

*Page intentionally left blank*

# 1 INTRODUCTION

Nuclear thermal propulsion (NTP) systems for interplanetary travel (e.g., manned missions to Mars) rely on nuclear reactors, rather than a more conventional chemical rocket, to heat the propellant and create the required thrust. This enables the use of molecular hydrogen ( $H_2$ ) in lieu of gaseous  $H_2O$ , thereby increasing the specific impulse by a factor of two [1] and reducing the travel time (or, alternatively, increasing the payload) by that same amount.

Compared to a traditional nuclear reactor, an NTP system operates at full power for a significantly lower duration (i.e., on the order of a few hours), with each burn typically lasting 15–60 minutes [2]. Conversely, the reactor must be able to reach full power within one minute [3], particularly to maximize the total impulse. This can be challenging to achieve without significant power and/or temperature overshoots (exceeding their maximum desired values by a large amount, but for a short period of time). Automated reactor control can solve this problem, but standard proportional integral derivative (PID) controllers often require a lot of manual adjustment and are not designed to anticipate sudden changes in power demand. Period-generated controllers (PGCs) [4], initially developed at the Massachusetts Institute of Technology, allow for better anticipation by using current reactor data (e.g., kinetics parameters, reactivity coefficients) to predict future behavior, and can thus reduce power overshoots, as was recently demonstrated using a point kinetics equation (PKE)-based NTP system model [5].

The primary purpose of this report is to investigate various automated control strategies applied to a full-core multiphysics model of a prototypical NTP system. One key finding is that hybrid PID controllers that use both power and reactivity signals can demonstrate a level of performance just as outstanding as that of period-generated control. Nevertheless, both rely on kinetics parameters, including some that may not be easily measured, such as quantities computed via delayed neutron precursor concentrations. The PID controller requires parameter tuning, while the PGC also necessitates providing temperature coefficients and their rates of change, which could be even harder to accurately measure in real time. Fortunately, preliminary results indicate that a certain level of inaccuracy may not drastically reduce PGC performance.

In addition to startup considerations, post-shutdown cooling is needed for NTPs to remove residual decay heat and prevent the system components from reaching excessive temperatures.



To minimize propellant consumption, pulsed cooling is usually relied upon. The secondary goal of this report is to perform a preliminary decay heat study to obtain equations describing the decay power deposited in the core as a function of time. These equations will be incorporated within the Griffin/Bison/RELAP model described in this report in order to simulate the shutdown transient following the startup and operational phases.

The rest of this report is organized as follows: an overview of the design is given in Section 2, and the full-core multiphysics model is presented in detail in Section 3. In Section 4, several reactivity control strategies are investigated, a novel hybrid approach is proposed, and these are compared against a challenging benchmark featuring an exponential ramp-up that increases the initial power by a factor of 500 within 30 s. In Section 5, the most successful strategies (hybrid PID and PGCs) are applied to a more realistic transient sequence in which the mass flow rate and outlet pressures are ramped up to model the bootstrap and thrust build-up phases prior to reaching steady-state conditions. In Section 6, the preliminary decay heat study is detailed, and conclusions are drawn in Section 7.

## 2 PROTOTYPICAL DESIGN

This section gives a brief overview of the reactor. The purpose of this study is to demonstrate methods of studying NTP systems by using MOOSE-based tools. Therefore, the design considered in this work is not meant to be fully realistic. In addition, to avoid export control classification of this document, only limited information is given regarding dimensions and material properties.

The design is based on the BWX Technologies NTP concept shown in Figure 1 and is largely identical to the one further presented in [1]. An overview of the active core design can be found in Figure 2. Specifically, the reactor core is composed of 61 high assay low enrichment uranium (HALEU) CERamic-METallic (CERMET) fuel assemblies arranged in circular rings in a zirconium hydride (ZrH) monolith (in magenta in Figure 2), acting as a neutron moderator. Each fuel assembly contains a fuel matrix (in blue), zirconium carbide (ZrC) insulator (in green), and silicon carbide (SiC) shell (in pink). Cooling channels are placed in the fuel matrix, with the H<sub>2</sub> coolant flowing downward. To simplify cross-section generation, nine representative moderator cooling channels (also in cyan), with H<sub>2</sub> coolant flowing upward, are modeled around each fuel assembly.






Figure 1: BWX Technologies NTP reactor concept (picture courtesy of [6]).

At its radial periphery, the monolith is surrounded by a beryllium reflector (in orange) containing 18 rotating control drums (CD) featuring an absorbing boron carbide ( $B_4C$ ) segment (in green) for controlling reactivity. Axially, the structural materials constituting the axial reflector are homogenized (i.e., represented as uniform materials with volume-averaged compositions).



(a) Radial view.

(b) Axial view.



(c) Zoom-in on a single fuel assembly.

Figure 2: Overview of the Serpent geometry.

Figure 3 gives an overview of how the propellant is assumed to flow through the system. It is pumped from a storage tank (point 0 in Figure 3) and divided into two paths: (1) a portion is preheated by cooling the nozzle and then flowing between the radial reflector and monolith, while (2) the remainder is fed directly into the moderator cooling channels. Both recombine (point 3) before being re-split between a turbine flow path and a bypass flow path. The hydrogen is then directed toward the fuel cooling channels and exits into the chamber (point 5) before passing through the nozzle. A turbo-pump at position 0 (not explicitly modeled) is placed on the same shaft as the turbine, thus allowing bootstrapping of the turbo machinery at reactor startup.

It is assumed that the accumulated operating time of the core will not exceed 4 hours, and that each burn will be under 1 hour, as is typical of an NTP system designed for Mars missions.




Figure 3: Sketch of the NTP system considered.

### 3 FULL-CORE MULTIPHYSICS MODEL

The full-core multiphysics model is based on several single physics models coupled together via the MOOSE framework [7]. The neutronics, heat transfer, and thermal-hydraulics models are presented in Sections 3.1–3.3, respectively. The coupling strategy is then discussed in Section 3.4.

### 3.1 Neutronics


The neutronics calculation relies on a two-step approach in which the cross-sections are first generated under steady-state conditions by the Monte Carlo code Serpent [8], using the 3-D full-core heterogeneous model depicted in Figure 2 for various core conditions, then employed in the multiphysics model by the Griffin neutronics input [9, 10]. To reduce the computational cost of the multiphysics transient simulation, a full-core super homogenization (SPH) procedure is then performed, as detailed in Section 3.1.2, and a cusping treatment technique is applied to model CD rotation, as described in Section 3.1.3.

#### 3.1.1 Cross-Section Generation

The reactor considered in this work has several unique characteristics that affect the process of cross-section generation. First, its very limited burn time removes the need to account for burn-up effects due to fission product buildup, including xenon and samarium poisoning. A fresh core is therefore assumed.

Second, unlike earth-based nuclear reactors, temperatures prior to startup can be well below room temperature. Unfortunately, some cross-section data in this range are missing from most standard libraries (e.g., the thermal scattering law [or  $S(\alpha, \beta)$ ]) for compounds such as ZrH, and should be directly generated (e.g., using NJOY [11]). Currently, the thermal scattering temperature is set to the lowest value available in the ENDF/B-VII.1 library (294 K for most isotopes) when the local temperature goes even below that.

Third, the extreme temperature gradient during nominal operations (exceeding 2500 K/m), combined with the relatively large neutron mean free path (on the order of 1–100 cm), make it debatable whether to only consider local conditions when evaluating the multigroup cross-sections. Indeed, while Doppler feedback effects are strictly local, the multigroup condensation process uses the neutron spectrum to compute cross-sections which depends significantly on conditions elsewhere in the core [1]. A significant approximation in this work—which will have to be addressed in the future—is that the cross-sections can be tabulated by assuming fixed axial profiles for all state variables, as was done in [12]. This is obviously an approximation, particularly because the temperature profile during startup is bound to change from a fairly flat initial distribution to



(a) Correlated to fuel temperature. (b) Correlated to moderator temperature.

Figure 4: Axial profiles assumed for cross-section generation.

a very steep one.

### 3.1.1.1 Axial profiles

A simplified multiphysics model based on a single fuel assembly and corresponding cooling channels was used (with various values of mass flow rate and power) to inform the axial temperature distribution. A total of 110 calculations were performed, with the mass flow rate being varied from 90 to 110% of its nominal value in increments of 2%, and the power being varied from 10 to 100% of its nominal value in increments of 10%. The average axial profiles obtained are given in Figure 4. In particular, Figure 4a suggests that the ZrC temperature, coolant temperature, and coolant density in the fuel cooling channels are strongly correlated to fuel temperature (note that hydrogen density generally decreases as its temperature increases). Similarly, Figure 4b suggests that the SiC temperature, coolant temperature, and coolant density in the moderator cooling channels are mostly correlated to moderator temperature.

### 3.1.1.2 Correlations

Despite fixing the axial profiles of component temperatures and hydrogen densities in the cool-

ing channel, the number of potential tabulation variables remains too large: were the cross-section library to be tabulated with fuel, insulator, shell, and moderator temperatures in addition to the fuel and moderator cooling channel hydrogen temperatures/densities and CD angles, the resulting library would (1) require an intractable number of Serpent calculations and (2) be tremendous in size, and (3) the cross-section interpolation cost in Griffin would be large. The computational expense for each of these issues grows exponentially with the number of tabulation variables—a concept commonly referred to as the curse of dimensionality. The only viable option would then be to have an extremely coarse grid with significantly degraded accuracy.

Instead of using a large number of state variables, the fact that state variables are tightly correlated to each other—as seen in Section 3.1.1.1—can be utilized to reduce the number of state variables without sacrificing much in the way of accuracy. For a  $\bar{T}_{\text{fuel}}$  of 430–2275 K, the average insulator and fuel cooling channel hydrogen temperature and density were found to correlate well to the average fuel temperature:

$$\bar{T}_{\text{ZrC}} = 32.13575831 + 0.578875903 \times \bar{T}_{\text{fuel}}, \quad (1)$$

$$\bar{T}_{\text{H}_2, \text{fuel}} = 95.46106691 + 0.647762551 \times \bar{T}_{\text{fuel}}, \quad (2)$$

$$\bar{\rho}_{\text{H}_2, \text{fuel}} = 10^{-3} \times (1.518264506 + 1461.314462/\bar{T}_{\text{fuel}}). \quad (3)$$

Furthermore, for a  $\bar{T}_{\text{mod}}$  of 135–705 K, the average shell and moderator cooling channel hydrogen temperature and density correlate well to the average moderator temperature:

$$\bar{T}_{\text{SiC}} = 68.36555689 + 0.761900094 \times \bar{T}_{\text{mod}}, \quad (4)$$

$$\bar{T}_{\text{H}_2, \text{mod}} = 37.03174461 + 0.231197231 \times \bar{T}_{\text{mod}}, \quad (5)$$

$$\bar{\rho}_{\text{H}_2, \text{mod}} = 10^{-3} \times (24.97905667 + 3184.311631/\bar{T}_{\text{mod}}). \quad (6)$$

For each of these correlations, the coefficient of determination ( $R^2$ ) exceeded 0.9. This essentially means that the relative average distance between the simulated points and the correlation line was less than 10%, a value deemed satisfactory compared to other modeling errors.

Any of the quantities listed above can then be retrieved at any axial point by multiplying its average by the shape described in Section 3.1.1.1.



### 3.1.1.3 State points

The multigroup cross-section library was then tabulated with respect to  $\bar{T}_{\text{fuel}}$ ,  $\bar{T}_{\text{mod}}$ , and the CD angles:

- $\bar{T}_{\text{fuel}} \in \{475, 775, 1075, 1375, 1675, 1975, 2275\}$
- $\bar{T}_{\text{mod}} \in \{150, 335, 520, 705\}$
- $\theta \in \{0, 60, 120, 180\}$ .

It thus contains a total of 112 state points, each obtained by a Serpent calculation performed with 2,000 cycles of  $5 \times 10^5$  particles after 100 inactive cycles.

### 3.1.1.4 Group structure

Broad energy group boundaries must be carefully selected to enable accurate neutronics calculations. In thermal reactors, special care must be taken to ensure that some energy groups are present between the Maxwellian peak and low-lying resonances in uranium-238, thereby avoiding artificial up-scattering of neutrons into resonance energies [13]. Figure 5a shows the energy spectrum in the bottom fuel layer (where the temperature is set to 2,450 K), for all the rings of fuel assemblies. The central assembly (ring 1) has a much harder spectrum, due to having a much lower moderator-to-fuel-volume ratio in its “neutronic” vicinity. On the other hand, the outermost ring (ring 5) has the softest spectrum, due to its proximity to the radial reflector. Figure 5b further indicates what kind of spectral shift is induced by going from 450 K (top) to 2,450 K (bottom).

Figures 5a and 5b show that the tail of the Maxwellian distribution drops all the way down to the region representing a few eVs. With such considerations in mind, General Atomics’ nine-group structure [14], shown in Table 1, was chosen, noting that several group boundaries (2.38, 3.93, and 17.6 eV) exist in between the thermal and resonance regions. In particular, the 17.6 eV boundary should serve to prevent up-scattered neutrons from reaching the resonance range [13].

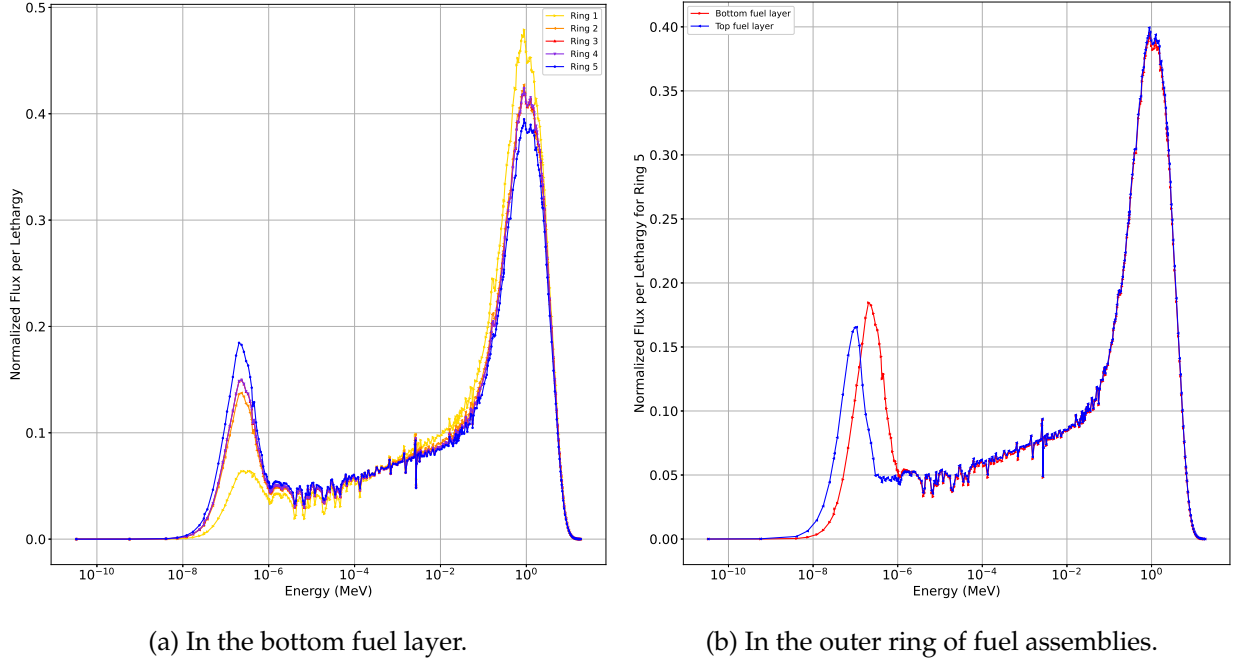


Figure 5: Energy spectrum, using the Tripoli 315-group structure.

Table 1: Energy group (upper) boundaries for the General Atomics nine-group structure [14] used for the Griffin model.

Group	Energy (MeV)	Group	Energy (MeV)
1	1.400E+01	6	2.380E-06
2	1.830E-01	7	1.275E-06
3	9.610E-04	8	8.250E-07
4	1.760E-05	9	1.300E-07
5	3.930E-06		

### 3.1.2 SPH Correction

To reduce the neutronics cost entailed by the multiphysics calculation, a homogeneous neutronics model was used in Griffin, typically necessitating some equivalence techniques to alleviate the otherwise significant homogenization error. Several techniques may be used, such as the SPH method or the use of discontinuity factors—both of which are available in Griffin. The most mature is the preconditioned Jacobian-free Newton-Krylov (PJFNK) SPH method, which can efficiently compute full-core SPH factors via a Newton solve approach [15]. This method was more recently enhanced to encompass problems involving vacuum boundaries and to support spatial restriction [16]—two capabilities that have both proven invaluable for this model.

Spatial restriction is the application of the SPH methodology to a portion of the domain; for the remainder, non-corrected cross-sections are used. This approach is appealing for several reasons: (1) the reference fluxes in the reflector can be quite small, which typically creates convergence problems and increased statistical uncertainty in the Serpent Monte Carlo data; (2) the reference flux tallies in Serpent can be particularly complicated to set up, and would have to be consistent with the mesh; and (3) as long as all the fuel regions are SPH-corrected, the overall eigenvalue and power distribution can still be preserved. Thus, limiting the SPH correction to the fuel assemblies appears highly advantageous. Most of the axial reflector is also excluded for similar reasons.

Furthermore, since the neutronics model uses vacuum boundaries, the simplest equivalence option is to use SPH factors without normalization factors, and to scale the fission terms with the reference eigenvalue (during the SPH procedure only) [16]. The alternative would be to add discontinuity factors, thus requiring the addition of corresponding partial currents to the equivalence libraries in order to achieve limited benefit.

After a neutronics-only SPH calculation is run for each state point, the corrected cross-section library enables reproduction of the reference multiplication factor and the reaction rates at those state points, provided the mesh is unchanged. In particular, this means that a very coarse mesh (e.g., the Griffin mesh [shown in Figure 6] generated using the NEMO software [17]) can be used.

### 3.1.3 CD Rotation

Besides tabulating the cross-sections and SPH factors with respect to CD angle—thereby accounting for changes in energy spectrum—CD rotation modeling relies on the de-cusping technology developed in Griffin to simulate control rod withdrawal and CD rotation with a static mesh [18]. Rather than explicitly meshing and rotating the drum absorber regions, the cross-sections within the rotating part of the drums are assigned based on the position of the various segments of the drum. Whenever an element of the mesh contains both the absorber and non-absorber segments, proper weighting is performed to avoid cusping, which manifests in non-smoothness of the reactivity curve. Figure 7 illustrates a typical reactivity curve obtained for temperatures matching a state point: the procedure ensures the eigenvalue matches the reference eigenvalue at grid points (vertical dotted lines), and the cusping treatment gives a smooth behavior between




Figure 6: Full-core neutronics mesh (picture courtesy of [1]).

grid points, limiting the error to less than 50 pcm—even lower at the region where the reactivity has a fairly linear profile.

### 3.1.4 Prompt Power and Decay Heat

With the exception of Section 6—which specifically covers decay heat—all the power is assumed to be instantaneously deposited. As explained in Section 3.3.2, a portion of the heat generated is considered to be directly used to heat up components outside the active core (radial reflector, nozzle, etc.), and is directly accounted for in the thermal-hydraulics model described in Section 3.3. Therefore, the neutronics model only computes the assembly-homogenized power density, whose integrated value equals  $\mathcal{P}_{\text{tot}}$  (see Equation (11)). In the heat transfer model, the




Figure 7: CD reactivity curve.

split between moderator and fuel heating is performed directly.

## 3.2 Fuel Assembly Heat Transfer

### 3.2.1 Approach

In theory, heat transfer should be modeled at the full-core level, because (among other reasons) a portion of the fission power is directly deposited by photons in the moderator and reflector. The current model, however, assumes that assemblies are perfectly insulated from the reflector block and that their thermal-hydraulics can therefore be modeled independently. This is valid only if the majority of the heat transfer not readily removed by the coolant occurs axially rather than radially. Nevertheless, the validity of this assumption should be further studied and the model enhanced accordingly. In the meantime, a single representative 30-degree fuel assembly was modeled for each ring of assemblies, along with its corresponding fuel and moderator cooling channels, with the mesh shown in Figure 8 extruded over the entire fuel height. Three gap heat transfer models were successively used between the fuel, insulator, shell, and moderator. Note that the moderator cooling channels in the thermal model are not fully consistent with the ones in the Serpent model.




Figure 8: X-Y view of the 30-degree slice thermal mesh (picture courtesy of [1]).

### 3.2.2 Power Density Conversion

To convert the homogeneous power density computed by the neutronics model into the heterogeneous power density to be applied in each representative 30-degree slice, the average power density  $p_{\text{assm},r,i}$  was computed<sup>1</sup> for each assembly ring (of index  $r$ ) over a given axial layer (of index  $i$ , with the total number of layers being set to 10). Next, a volume correction factor was applied to account for volume changes between homogeneous and heterogeneous power densities, and a fixed fraction  $\chi$ —defined by Equation (12)—was assumed to be deposited in the fuel, while the remaining amount was assumed deposited in the moderator. Specifically, the heterogeneous power density deposited in the fuel of ring  $r$  and axial layer  $i$  was given by:

$$p_{\text{fuel},r,i} = \chi \frac{V_{\text{assm}}}{V_{\text{fuel}}} p_{\text{assm},r,i}. \quad (7)$$


The one in the moderator was given by:

$$p_{\text{mod},r,i} = (1 - \chi) \frac{V_{\text{assm}}}{V_{\text{mod}}} p_{\text{assm},r,i}, \quad (8)$$

where  $V_{\text{assm},i}$  is the volume of a homogenized fuel assembly, and  $V_{\text{fuel}}$  and  $V_{\text{mod}}$  are the heterogeneous volumes of fuel and moderator in that assembly (once again, these volumes should be assembly dependent but are currently computed based on the mesh shown in Figure 8). This approach preserves total energy but neglects any radial dependency of the power density in the fuel—which, in reality, should be higher at the periphery of the fuel, due to increased moderation (phenomenon known as spatial self-shielding), as seen in Figure 9a. Either a heterogeneous neutronics model or form functions would be needed to account for this effect. Future work will add an assembly power reconstruction.

---

<sup>1</sup>This was done using the `NearestRadiusLayeredAverage` object in MOOSE, specifically implemented therein for this project.



(a) Power density distribution.

(b) Fission rate and thermal flux.

Figure 9: Typical Serpent results obtained with the CD rotated halfway ( $\theta = 90^\circ$ ).

### 3.3 Thermal-Hydraulics

The thermal-hydraulics model is summarized in Figure 3. It is composed of five moderator and fuel cooling channels—described in Section 3.3.1—with each pair representing all the cooling channels within a given ring of assemblies and coupled together through a simplified system model, as detailed in Section 3.3.2.

#### 3.3.1 Representative Cooling Channels

Each representative moderator cooling channel consists of a single 1-D pipe extending over the height of the fueled regions (thus excluding the axial reflectors), whereas the representative fuel cooling channel also includes the axial reflectors, with the flow area changing between the upper axial reflector, the active core, and the lower axial reflector.

The heat transfer coefficients were computed per:

$$h = \frac{k_{H_2} \text{Nu}}{D_h}, \quad (9)$$

where  $k_{H_2}$  is the fluid thermal conductivity,  $D_h$  is the hydraulic diameter, and the Nusselt number



is given by the Dittus-Boelter correlation:

$$\text{Nu} = 0.023\text{Re}^{0.8}\text{Pr}^{0.4}. \quad (10)$$

In addition, a distributed loss coefficient equal to 0.27 over the length of the channel and lumping the ZrC support structures was assumed in the active core region of the fuel cooling channels.

Due to MOOSE's current inability to accurately model hydrogen properties at high temperatures, ideal gas properties were used instead. The value of the average specific heat capacity was adjusted to match the enthalpy rise from properties provided by National Aeronautics and Space Administration (NASA) over a temperature range of 50–2700 K. This approximation tends to underestimate the temperature of hydrogen prior to 2700 K, and to overestimate it afterward. This should eventually be addressed to get accurate transient results.

### 3.3.2 System Model

#### 3.3.2.1 Energy Balance Assumptions

Currently, the multiphysics model is limited to the active core region shown in Figure 3, while the rest of the system is modeled using boundary conditions and algebraic equations. To this end, assumptions were made regarding the global energy balance of the system. First, the following quantities were defined (the representative values at full power are displayed in parentheses). These values were obtained from a previous study by NASA for a prototypical CERMET design:

- $\mathcal{P}_{\text{fuel}}$ : Total power deposited in the fuel (297.6 MW)
- $\mathcal{P}_{\text{mod}}$ : Total power deposited in the moderator (17.1 MW)
- $\mathcal{P}_{\text{refl}}$ : Total power deposited in the reflector, assumed to heat up the fluid directly (2.63 MW)
- $\mathcal{P}_{\text{RGC}}$ : Total power from regenerative cooling (11.9 MW), heats up the inlet of the fuel cooling channel, cools down the outlet (after the chamber)
- $\mathcal{P}_{\text{turbine}}$ : Total power extracted by the turbine from the coolant (2.53 MW).

The total power deposited in the core is then defined as:

$$\mathcal{P}_{\text{tot}} \equiv \mathcal{P}_{\text{fuel}} + \mathcal{P}_{\text{mod}}. \quad (11)$$

While all these quantities will evolve during startup, a simplifying assumption throughout this work is that the following quantities are independent of time:

$$\chi \equiv \frac{\mathcal{P}_{\text{fuel}}}{\mathcal{P}_{\text{tot}}} = 94.6\% \quad (12)$$

$$\gamma \equiv \frac{\mathcal{P}_{\text{refl}} + \mathcal{P}_{\text{RGC}}}{\mathcal{P}_{\text{tot}}} \quad (13)$$

$$\gamma' \equiv \frac{\mathcal{P}_{\text{turbine}}}{\mathcal{P}_{\text{tot}}}. \quad (14)$$

### 3.3.2.2 Boundary Conditions

A simplified system model—mimicking the system described in Section 2 and illustrated in Figure 3—was then developed, mostly relying on the following assumptions to determine the boundary conditions of the cooling channels modeled inside the active core:

- The temperature of the hydrogen exiting the storage tank is  $T_0 = 50$  K.
- The ratio  $\beta$  of the mass flow rate through the moderator channels to the total mass flow rate is a constant equal to 0.62, a value taken from a preliminary NASA study.
- The total mass flow rate through the system is proportional to the chamber pressure  $p_C$  (which corresponds to the pressure at point 5 in Figure 3) and is given by:

$$\dot{m}(t) = \frac{\dot{m}_{\text{nom}}}{p_{C,\text{nom}}} p_C(t), \quad (15)$$

where the subscript  $_{\text{nom}}$  represents nominal value quantities (i.e., at steady-state).

- The ratio of the chamber pressure to the pressure at the moderator channel exits is a constant

equal to  $\alpha = 1.44$  (taken from the same aforementioned NASA study):

$$p_3(t) = \alpha p_4(t). \quad (16)$$

In practice,  $p_2 = p_3$ . And if the pressure drop is neglected over the active core,  $\alpha$  becomes the ratio of the moderator cooling channel outlet pressure to the chamber pressure (i.e., the fuel cooling channel outlet pressure). A better pressure drop model will be incorporated in the future.

- Assuming that the coolant specific heat capacity  $\bar{C}_p$  is constant over the range of 50–300 K and equal to  $1.6 \times 10^4$  J/kg/K, the temperature at point 1 (see Figure 3) can be expressed as:

$$T_1 = T_0 + \frac{\gamma \mathcal{P}_{\text{tot}}}{(1 - \beta) \dot{m} \bar{C}_p}, \quad (17)$$

where  $\gamma$  is defined by Equation (13), and the temperature of the hydrogen entering the turbine is approximately:

$$T_3 = \beta T_2 + (1 - \beta) T_0 + \frac{\gamma \mathcal{P}_{\text{tot}}}{\dot{m} \bar{C}_p}. \quad (18)$$

The inlet fuel cooling temperature is:

$$T_4 = \beta T_2 + (1 - \beta) T_0 + \frac{(\gamma - \gamma') \mathcal{P}_{\text{tot}}}{\dot{m} \bar{C}_p}, \quad (19)$$

where  $T_2$  is the temperature of the hydrogen exiting the moderator channel and  $\gamma'$  is defined by Equation (14).

- The outlet temperatures (at points 2 and 5 in Figure 3) were computed by assuming that the temperatures of the hydrogen exiting each assembly were fairly similar, thus enabling the difference in heat capacity to be neglected:

$$T_k = \sum_{r=1}^R \frac{n_{\text{assm},r}}{N_{\text{assm}}} T_{k,r}, \quad (20)$$

where  $R = 5$  is the total number of rings of assemblies and  $k$  equals 2 for the moderator channel, or 5 for the fuel channel. For each ring  $r$ ,  $n_{\text{assm},r}$  is the number of assemblies in

that ring, with  $N_{\text{assm}} = \sum_{r=1}^R n_{\text{assm},r} = 61$ .  $T_{k,r}$  is the computed outlet (fuel or moderator, depending on  $k$ ) coolant temperature in the representative assembly of ring  $r$ .

### 3.4 Multiphysics Coupling

#### 3.4.1 Data Transfer

Before detailing the various reactivity control strategies in Section 4, an overview of the multiphysics model in terms of coupling is shown in Figure 10, with a generic controller that determines the angle of the CD, based on a setpoint for a given quantity of interest (e.g., total power) and the data measured<sup>2</sup> or predicted by the full-core model, shown in a black box and consisting of three levels of MOOSE sub-apps. In particular, it reiterates that a single homogenized neutronics model is coupled with five pairs (i.e., one for each ring of fuel assemblies) of representative 30-degree slices: one for the heat transfer model, and one corresponding to the thermal-hydraulics model (with a representative fuel cooling channel and moderator cooling channel), as is also explained in Figure 3.

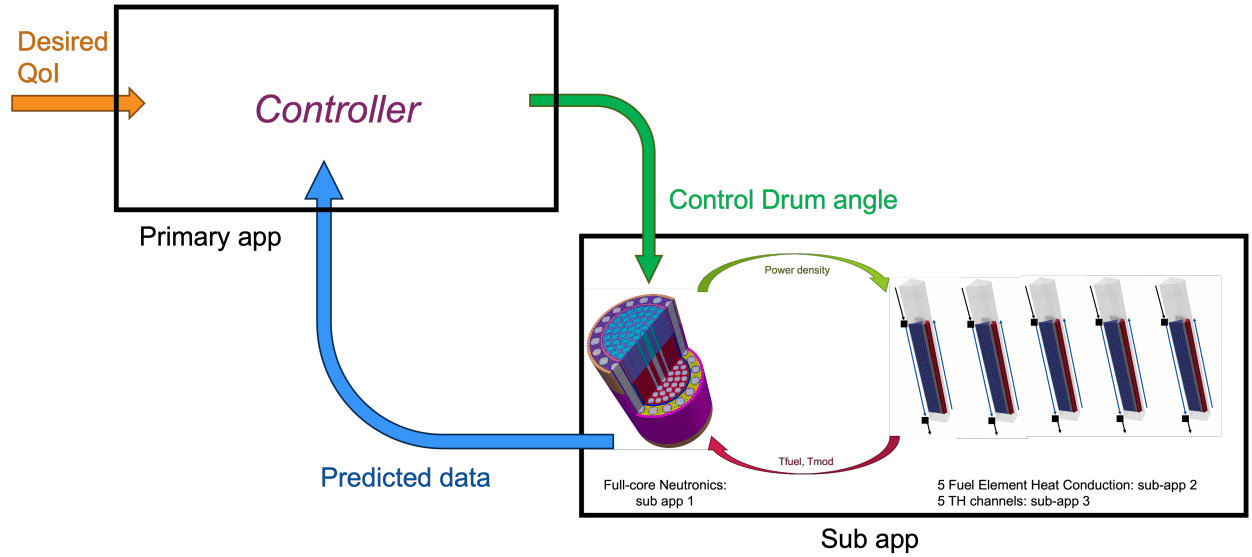


Figure 10: Schematics of the full-core multiphysics model.

A summary of the various pieces of the multiphysics model, along with the data transfers that occur between them, is presented in Table 2. All the controllers considered in this work determine

<sup>2</sup>The reason the terms “measured” and “predicted” are used interchangeably throughout the remainder of this report is that, while this work uses a numerical model to predict reactor behavior, a real controller would require (at least some of) these quantities to be measured from the real system.

Table 2: Data transfer for startup simulation.

From To	Controller	Neutronics	Heat Transfer	Thermal-hydraulics
Controller	X	Total power (+ kinetics and feedback data)	X	X
Neutronics	CD angle	X	Fuel and moderator average temperatures	Outlet moderator cooling channel temperature
Heat transfer	X	Power density	X	Fluid temperatures, heat transfer coefficients
Thermal- hydraulics	X	Inlet fuel cooling channel temperature	Wall temperatures	X

the CD angle to be used by the neutronics model. In contrast, the information required by the controllers varies, depending on the selected control strategy, as discussed in detail in Section 4. Most of them rely on the total power, though some necessitate kinetics data, feedback coefficients, and/or temperature rates of change.

The neutronics also provides an axially-layered-averaged (i.e., averaged over a fixed number of axial layers of equal thickness, set to 10 in this work) power density to each assembly ring's heat transfer model, and receives the fuel and moderator average temperatures (using a formula similar to Equation (20) to average over the entire core). Each heat transfer model transfers axially-layered-averaged (with 40 layers) wall temperatures to the thermal-hydraulics model, and in return obtains axially-layered-averaged (also with 40 layers) fluid temperatures and heat transfer coefficients for both the moderator and fuel cooling channels.

In addition, to compute the inlet temperature of the fuel cooling channels, the average outlet temperature of the moderator cooling channels is needed. To perform the weighted average over the assembly rings, as per Equation (20), the latter is passed to the neutronics application—because there is currently no way in MOOSE to perform averages across sub-apps without bringing the data to the parent application—and the former is computed using Equation (19) before being passed back to the thermal-hydraulics model (in practice, via the heat transfer model).

### 3.4.2 Tight Coupling

To converge the coupling terms and thus solve the multiphysics system at each time step, Picard iterations were performed between the neutronics and the heat transfer models, with the relative and absolute tolerances being set to  $10^{-3}$  and  $10^{-6}$ , respectively. Inner Picard iterations between the heat transfer and the thermal hydraulics models were found unnecessary. No fixed-point iterations were performed at the controller level within a given time step, either. Adding them should be possible—though restrictions on the time step size might ensue, for stability reasons—provided that the delay coming from the instrumentation extracting the measured data needed by the controller is incorporated into the controller model. In the meantime, the time step size acts as this time delay, and fully converging the controller signal with respect to time—through Picard iterations—would amount to assuming the measurements to be instantaneous, which is probably undesirable.

### 3.4.3 Workflow

Any transient presented in this report is preceded by a neutronics steady-state (i.e., eigenvalue) calculation to establish the neutron flux initial condition, and the fission terms of the transient simulation are initially scaled with the multiplication factor to put the reactor in an initial critical state. A fixed temperature distribution is assumed in the fuel and the moderator (potentially with different average values for each). This approach should eventually be improved so as to use more realistic thermal initial conditions. This differs from the traditional approach, in which a coupled multiphysics steady-state calculation is used as the initial condition—as is suitable whenever the initial reactor is in thermal equilibrium, whereas typical NTP startup sequences are too rapid to justify such an assumption.

For most of the control strategies presented in Section 4, kinetics parameters (e.g., dynamic reactivity) are needed inputs to the controller. Therefore, an additional neutronics-only adjoint calculation is run with Griffin prior to the transient simulation, using the same temperature distribution employed to compute the neutron flux initial condition. This adjoint solution is stored and used during the transient simulation to compute time-dependent kinetics parameters. Nevertheless, their computation ignores any change in temperature profile, which is a source of inaccuracy

in that the controller will compute a signal based on potentially inaccurate kinetics parameters.

## 4 REACTIVITY CONTROL STRATEGIES

This section is dedicated to the study of various strategies to automatically control the angle of the drums, thereby adjusting the core reactivity. An automated strategy is necessary, since a priori determination of a time-dependent CD angle would be extremely challenging due to the numerous nonlinear feedback mechanisms that occur inside a nuclear reactor [1], especially considering the very short period of time within which an NTP must reach full power.

### 4.1 Exponential Ramp-up Benchmark

To assess various approaches, a simple yet challenging benchmark was devised with an exponentially increasing power setpoint—or demanded power that the system is to follow—given by:

$$\mathcal{P}_d(t) = \begin{cases} \mathcal{P}_i \left( \frac{\mathcal{P}_f}{\mathcal{P}_i} \right)^{t/t_f} & , \quad 0 \leq t < t_f \\ \mathcal{P}_f & , \quad \text{otherwise} \end{cases} \quad (21)$$

where  $\mathcal{P}_i = 0.5$  MW and  $\mathcal{P}_f = 250$  MW are the initial and final powers, respectively, and  $t_f = 30$  s is the final time at which the ramp-up is to be terminated. In other words, the power should be multiplied by 500 over the span of 30 s, then instantaneously set to a constant value. This is bound to be challenging to achieve without significant power overshoot, which is defined by:

$$O_{\mathcal{P}} = \frac{\max(\mathcal{P}(t))}{\mathcal{P}_f} - 1. \quad (22)$$

For this benchmark, the following further assumptions are made:

- The initial CD angle is set to 120 degrees (see Figure 7).
- The average temperatures in the fuel and moderator regions are initially set to 500 and 200 K, respectively, with the axial shapes from Section 3.1.1.1.
- The total mass flow rate—and therefore chamber pressure, given Equation (15)—is set to its nominal value throughout the transient.

- The inlet fluid temperatures for the moderator and fuel cooling channels are also assumed constant, and are set to 200 and 500 K, respectively. This assumption is mostly made because some of the control strategies will not perform well initially and could therefore drop the component temperatures to very low values, resulting in convergence issues. Use of fixed inlet temperatures enables assessment of such methods' behavior after this initial phase.
- The time step  $\Delta t$  is set to 0.1 s.

## 4.2 PID Control of the Drum Rotation

### 4.2.1 Power-driven PID Signal

Perhaps the most straightforward approach is to rely on a standard PID controller, with the error signal computed directly using the demanded and measured powers ( $\mathcal{P}_d$  and  $\mathcal{P}$ , respectively):

$$\varepsilon(t) = \mathcal{P}_d - \mathcal{P}. \quad (23)$$

Next, the following rotation angle is applied to the drums:

$$\Delta\theta_{\mathcal{P}}(t) = K_p\varepsilon(t) + K_i \int_0^t \varepsilon(t')dt' + K_d \frac{d\varepsilon}{dt}, \quad (24)$$

where  $K_p$ ,  $K_i$ , and  $K_d$  are the proportional, integral, and derivative constants, respectively. The CD angle is then given by integrating the rotation angles over time (i.e., at time step  $n + 1$ ):

$$\theta_{\mathcal{P}}(t_{n+1}) = \theta_i + \sum_{k=1}^n \Delta\theta_{\mathcal{P}}(t_k), \quad (25)$$

where  $\theta_i$  is the initial CD angle. This strategy is further illustrated by Figure 11.

Nevertheless, a disadvantage of PID control is that the constants must often be manually tuned, especially for complex systems whose transfer functions are not determined easily, thus rendering theoretical determination of optimum values arduous if not impossible. Trial and error then becomes the standard approach. Figure 12 shows the results for three values of  $K_p$  while keeping  $K_i = K_d = 0$ :

1.  $K_p = 10^{-8}$  deg/W: The measured power struggles to closely match up with the setpoint,



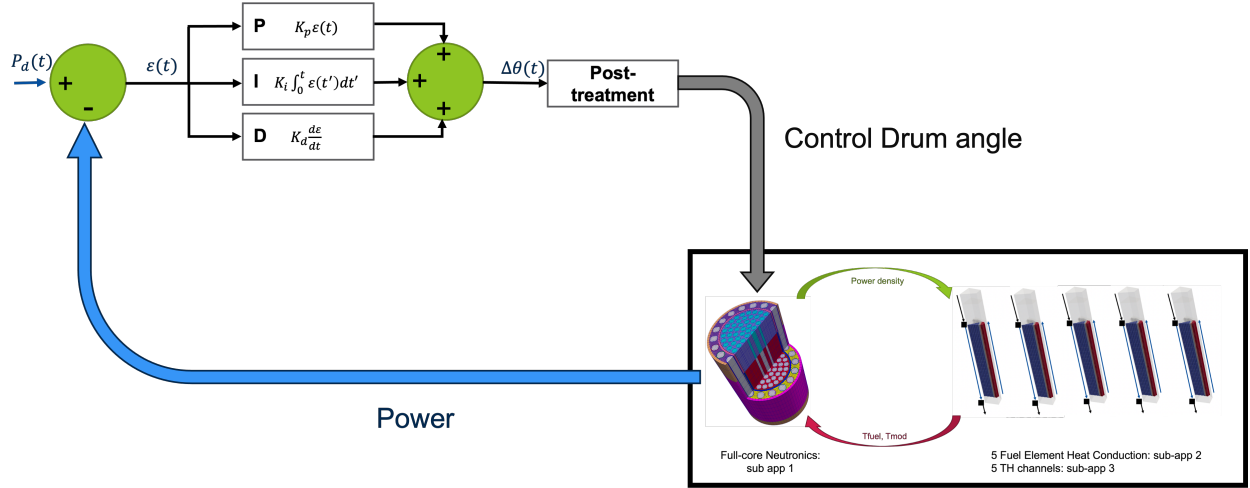


Figure 11: Schematics of power-driven PID control of the full-core multiphysics model.

as it takes nearly 17.2 s to double its initial value, whereas the signal achieves that in less than 3.5 s. This is further evidenced by the CD angle, which remains virtually constant until the power setpoint greatly exceeds its initial value. At 20–25 s, the predicted power catches up, owing to very limited temperature feedback, then nears the signal as the (negative) temperature feedback becomes more and more pronounced. However, it also experiences a significant power overshoot of more than 18%—with the power exceeding 295 MW—further indicating that the controller is not reactive enough and that  $K_p$  should likely be increased.

2.  $K_p = 10^{-7}$  deg/W: The controller behavior is much better, especially in terms of power overshoot, which is reduced to 1.1%, with small and rapidly dampened oscillations observed and the power stabilizing fairly quickly to the desired value. Nevertheless, the behavior early on remains somewhat unsatisfying. The measured power still trails the setpoint significantly for the first 10 s. After that, it is able to catch up and follow the exponential ramp fairly well.
3.  $K_p = 2 \times 10^{-7}$  deg/W: If  $K_p$  is further increased, the time before which the controller is able to closely follow the demand is reduced. Unfortunately, the system appears too reactive, and an onset of CD oscillations appears at 28.8 s, shortly before reaching nominal power. While these oscillations do not diverge, the impact on power is significant, with the measured power rapidly oscillating between 80 and 450 MW.

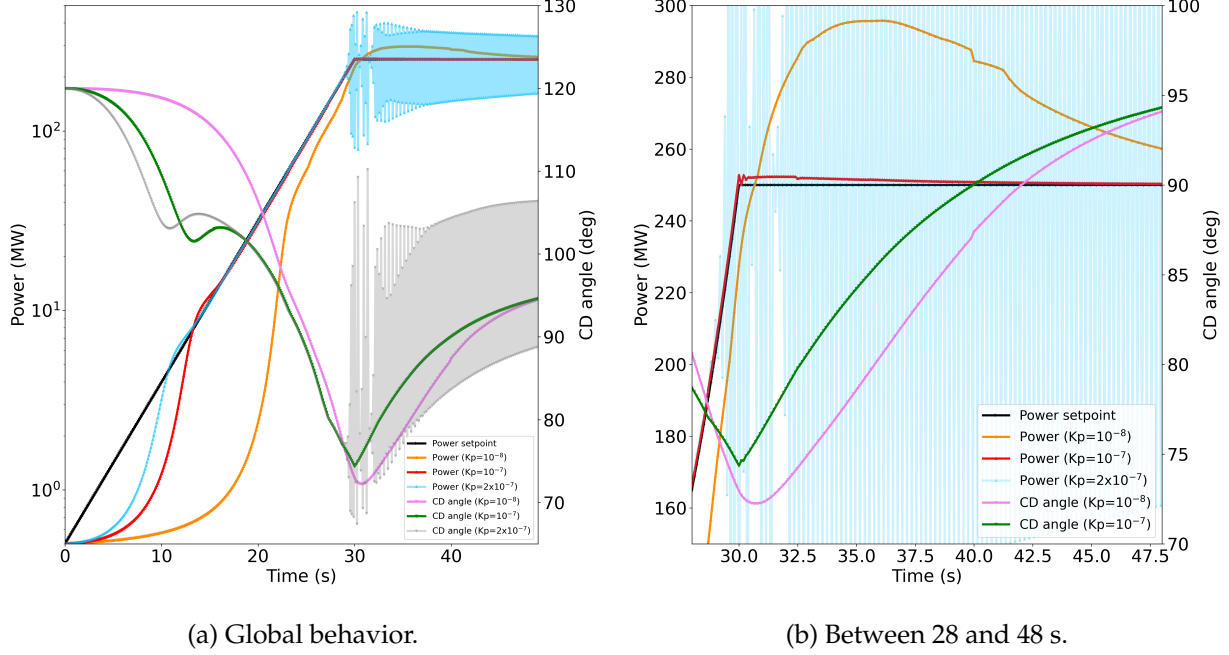


Figure 12: Evolution of the power and CD angles for various values of  $K_p$  (in deg/W) for the power-driven PID controller when applied to the exponential ramp-up benchmark.

All things considered, the power overshoot (on the order of 1%) obtainable via a power-driven PID controller is surprisingly good, given the time delay between the CD rotation and the actual increase in power. However, this delay in response is manifest in the early stages of this benchmark, and though the PID constants could be further optimized by finely tuning  $K_p$ , as well as  $K_i$  and  $K_d$ , it is unlikely that this can be overcome without destabilizing the system. Indeed,  $K_i$  has an integral effect and would thus be unable to correct early discrepancies, whereas  $K_d$  tends to degrade stability and would require careful tuning.

#### 4.2.2 Reactivity-driven PID Signal

Another approach is to use a reactivity signal to drive the controller. This approach is promising because there is virtually no time delay between a rotation of the drums and the corresponding reactivity insertion. The reactivity error signal is then given by:

$$e(t) = \rho_d - \rho. \quad (26)$$

The following rotation angle is then applied to the drums:

$$\Delta\theta_\rho(t) = K'_p e(t) + K'_i \int_0^t e(t') dt' + K'_d \frac{de}{dt}, \quad (27)$$

where  $K'_p$ ,  $K'_i$ , and  $K'_d$  are the proportional, integral, and derivative constants of the reactivity-driven PID controller, respectively. The actual CD angle,  $\theta_\rho$ , is computed using an equation similar to Equation (25). Following the methodology described in [1],  $K'_p$ —which corresponds to the rotation angle (in degree) applied to every dollar of reactivity in the error signal—can be adjusted using the control drum reactivity curve (see Figure 7), by setting  $K'_p$  to the inverse of the reactivity inserted per degree in the operating range of the CD, thus yielding approximately  $K'_p = 25 \text{ deg}/\text{\$}$ . By studying the early behavior of the controller when applied to a simple case,  $K'_i$  was set to 0.25  $\text{deg}/\text{\$-s}$ , while  $K'_d$  was left at 0.

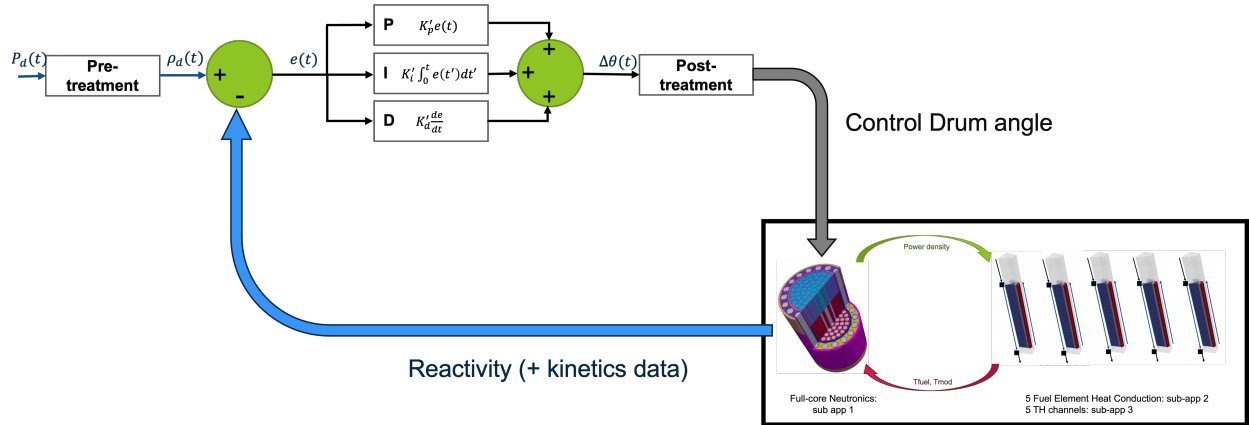


Figure 13: Schematics of reactivity-driven PID control of the full-core multiphysics model.

If the true setpoint is the power (as in the benchmark considered), then this signal must first be converted into a reactivity signal, which can be done by re-arranging the PKE as follows:

$$\rho_d = \frac{\Lambda}{\mathcal{P}_d \beta_{\text{eff}}} \left( \frac{d\mathcal{P}_d}{dt} - \sum_i \lambda_i C_i \right) + 1, \quad (28)$$

where  $\Lambda$  is the neutron mean generation time,  $\beta_{\text{eff}}$  is the effective delayed neutron fraction, and  $\lambda_i$  is the decay constant of the delayed neutron precursor group  $i$  with concentration  $C_i$ . The resulting control strategy is illustrated by Figure 11.

Nonetheless, a clear disadvantage of this method—not considered in [1], since a given reac-

tivity setpoint was assumed—is that having a zero signal does not necessarily guarantee that the power of the system will stabilize to the demanded power. In other words, even if the controller were able to perfectly follow the reactivity signal (which is unlikely to be the case), inaccuracy in the kinetics parameters used to evaluate either the demanded reactivity  $\rho_d$  (see Equation (28)) or the measured reactivity  $\rho$  from the core can result in the final power stabilizing to a value other than the nominal one.




Figure 14: Evolution of the power and CD angle for the reactivity-driven PID controller when applied to the exponential ramp-up benchmark.

In fact, the results displayed in Figure 14 show very good agreement between the demanded and measured powers during the exponential ramp-up phase, suggesting that the reactivity setpoint is being closely matched. Yet, the measured power not only exceeds the final demanded power but stabilizes at above 305 MW—22.4% higher than its expected value. Interestingly, the CD rotation speed (i.e., the time derivative of the CD angle) exhibits some unphysical discontinuity right before  $t = 40$  s. This is actually due to the moderator temperature exceeding its maximum tabulation value of 705 K (see Section 3.1.1.3) because the power is notably higher than its nominal value<sup>3</sup>. Yet, in Griffin, cross-sections are not extrapolated past the last interpolation point but are instead kept at this last value. As such, the (positive) feedback due to moderator temperature ceases at that point, and the drums no longer have to rotate as fast to maintain the power constant.

Figure 15 takes a closer look at the controller performance by directly examining the demanded and measured reactivity as a function of time, as well as the CD signal,  $\Delta\theta_\rho$ , and CD angle,  $\theta_\rho$ . The first observation is that the measured reactivity follows the signal very well for the first 10–15 s, then starts to slightly underestimate it as the temperature feedback becomes more pronounced. The discontinuities observed in the CD signal correspond to the times at which either the moderator or fuel temperature crosses a tabulation point. The exception to this is the discontinuity at  $t = 30$  s, which is due to the sudden change in power and thus reactivity signal. Interestingly, the controller handles these discontinuities without too much difficulty. However, one may wonder why the power ends up significantly exceeding the power setpoint if the measured reactivity is consistently lower than the signal. It is likely this is due to an inaccuracy in the computation of the measured kinetics data used in Equation (28). Indeed, as explained in Section 3.4.3, the computation of these kinetics data in Griffin relies on an adjoint solution computed with the initial temperature profile. Thus, the more the temperature profiles change, the more inaccurate the PKE data, making it very possible that the “true” reactivity seen in the Griffin model is higher than the signal for  $t$  at 15–30 s.

All in all, though it is clear that this strategy stabilizes at a much higher power than desired, and is thus a poor candidate for this type of startup simulation, its performance early on (i.e., before temperature feedback becomes significant) is precisely where the power-driven PID con-

---

<sup>3</sup>Note that in Section 5, the power will reach values higher than this without the fuel and moderator reaching temperatures exceeding their maximum tabulation values, as the coolant will enter the reactor at a much lower temperature.




Figure 15: Evolution of the measured and demanded reactivities, as well as the incremental rotation angle and CD angle for the reactivity-driven PID controller applied to the exponential ramp-up benchmark.

troller from Section 4.2.1 fell short. The following section therefore presents a hybrid strategy combining the two methods.

#### 4.2.3 Hybrid Power-Reactivity PID Signal

In summary, the power-driven PID from Section 4.2.1 performs well, except for early on when the power is low, whereas the reactivity-driven PID from Section 4.2.2 is only satisfactory in the first phase of the benchmark considered, as the power of the system fails to converge to the steady-state power setpoint. A power-reactivity-driven hybrid strategy combining the two approaches

was thus devised, as illustrated in Figure 16.

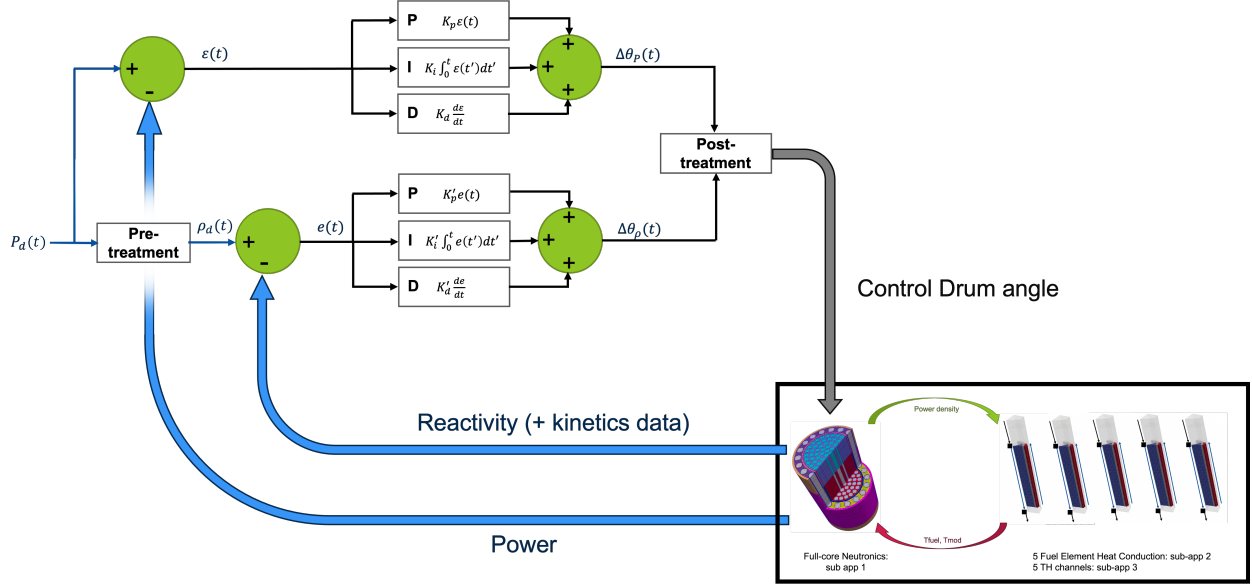


Figure 16: Schematics of hybrid power-reactivity PID control of the full-core multiphysics model.

The CD signal was then computed per the following weighted summation:

$$\Delta\theta = \xi\Delta\theta_P(t) + (1 - \xi)\Delta\theta_\rho(t), \quad (29)$$

where  $\xi$  was chosen to be defined as<sup>4</sup>:

$$\xi \equiv \frac{\ln\left(\frac{\mathcal{P}}{\mathcal{P}_i}\right)}{\ln\left(\frac{\mathcal{P}_f}{\mathcal{P}_i}\right)}, \quad (30)$$

where  $\mathcal{P}_i$  and  $\mathcal{P}_f$  are the initial and final powers, respectively. The PID constants were chosen to be identical to the best performers used in the previous sections<sup>5</sup>:  $K_p = 10^{-7}$  deg/W;  $K_i = K_d = 0$  for the power-based constants; and  $K'_p = 25$  deg/\$,  $K'_i = 0.25$  deg/\$-s, and  $K'_d = 0$  for the reactivity-based constants. The CD angle at time step  $n + 1$  is then given by:

$$\theta(t_{n+1}) = \theta_i + \sum_{k=1}^n \Delta\theta(t_k) = \theta_i + \sum_{k=1}^n (\xi\Delta\theta_P(t_k) + (1 - \xi)\Delta\theta_\rho(t_k)). \quad (31)$$

<sup>4</sup>Note that a linear definition of  $\xi$ , such as  $\xi \equiv \frac{\mathcal{P} - \mathcal{P}_i}{\mathcal{P}_f - \mathcal{P}_i}$ , gave virtually identical results for this benchmark.

<sup>5</sup>Note, however, that they could likely be further enhanced.




Figure 17: Evolution of the measured and demanded powers, as well as the CD signal and CD angle for the hybrid PID controller when applied to the exponential ramp-up benchmark.

Figure 17 shows the results for the hybrid power- and reactivity-driven PID controller. In particular, observe that, as intended, the contribution of the reactivity PID component largely dictates the CD signal early on (i.e.,  $\Delta\theta \approx \Delta\theta_\rho$ ; or equivalently,  $\xi \approx 0$ ), then gradually becomes less and less important as the power increases, enabling the power PID component to take precedence (i.e.,  $\Delta\theta \approx \Delta\theta_p$ ). Once again, discontinuities in the signal occur when the fuel or moderator temperature exceed a state point value—due to the piecewise linear interpolation scheme used for the Griffin cross-sections—but seem to minimally impact the controller response.

The hybrid controller is able to closely follow the power demand during the exponential ramp, and the power quickly stabilizes to the requested value with only a small power overshoot, virtu-



ally identical to the one observed with the power-driven PID: namely, 1.1%.

### 4.3 Period-generated Control of the Drum Rotation

Another control strategy, this one differing from standard PID control by attempting to impose reactor period based on kinetics data and feedback coefficients, was developed in [4] and recently applied to a PKE-based NTP system model [5].

#### 4.3.1 Methodology

The method, summarized in Figure 18, relies on the inverse of the (measured) reactor period:

$$\omega(t) = \frac{\dot{\mathcal{P}}(t)}{\mathcal{P}(t)}, \quad (32)$$

where  $\dot{\mathcal{P}}$  is the time derivative of the power. The error signal is then defined as:

$$e(t) = \ln \left( \frac{\mathcal{P}_d(t + j\Delta t)}{\mathcal{P}(t)} \right), \quad (33)$$

where  $j$  is a positive integer that should be set to “at least 2 to ensure stability against oscillations” [4], and  $\Delta t$  is the time step. Unlike the PID error signal, this error signal does not go to zero if the demanded and predicted powers match exactly, since this controller looks at the future demand rather than the current discrepancy between the demanded and obtained powers.

The demanded inverse period was originally expressed as [4]:

$$\omega_d(t) = \frac{e(t) + \frac{1}{T_i} \int_0^t e(t') dt' + T_d \frac{de}{dt}}{j\Delta t}. \quad (34)$$

However, it appears that the integral and derivative terms do not have the expected effect because the sign of  $e$  does not indicate whether the measured power is trailing or anticipating the demand. Therefore, the choice (in this report) was made to define an instantaneous error signal for use in the derivative and integral terms:

$$\epsilon(t) = \ln \left( \frac{\mathcal{P}_d(t)}{\mathcal{P}(t)} \right), \quad (35)$$

whose sign actually informs whether the measured power is trailing or anticipating the demand.

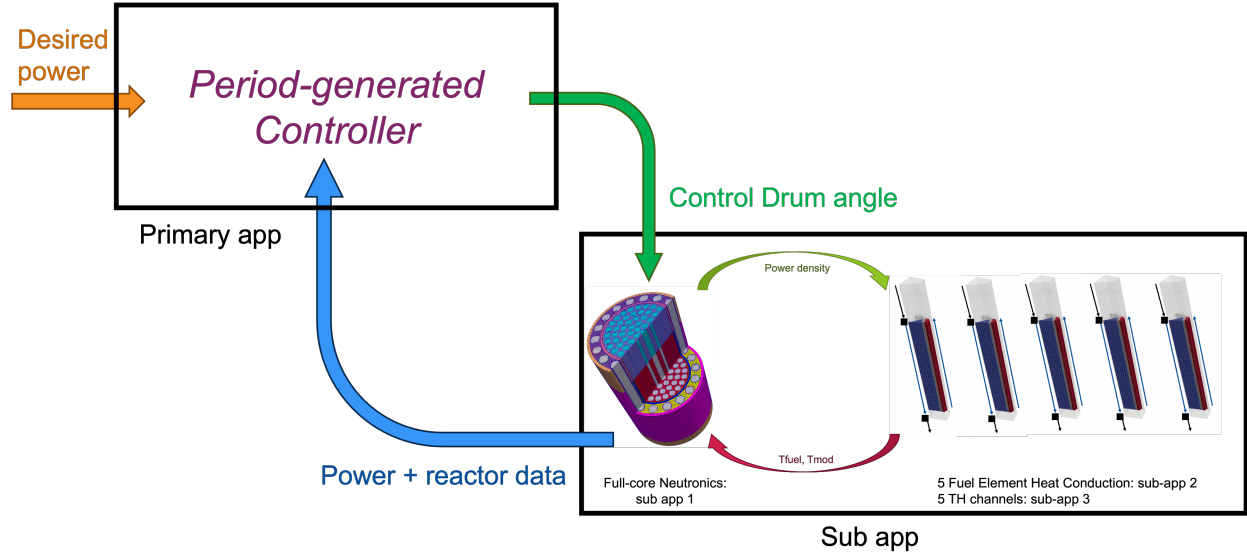


Figure 18: Schematics of period-generated control of the full-core multiphysics model.

In addition, for this work, it was decided to replace  $j\Delta t$  with a single tuning parameter  $\tau$  to avoid artifacts whenever  $\Delta t$  was not constant during the simulation, thus yielding the following modified expression for the demanded inverse period:

$$\omega_d(t) = \frac{e(t) + \frac{1}{T_i} \int_0^t \epsilon(t') dt' + T_d \frac{de}{dt}}{\tau}. \quad (36)$$

In practice, however, this period-generated control works well, at least for the transients considered herein, without the integral and derivative terms (i.e., using  $T_i = \infty$  and  $T_d = 0$ ). Thus, the final expression used in this work is:

$$\omega_d(t) = \frac{\ln \left( \frac{P_d(t+\tau)}{P(t)} \right)}{\tau}, \quad (37)$$

where  $\tau = 0.5$  s.

The required rate of change of reactivity was then introduced [4]:

$$\begin{aligned} \tilde{\rho}(t) = & (\beta_{\text{eff}} - \rho(t)) \omega_d(t) - \lambda'_e \rho(t) - \sum_{i=1}^I \beta_{\text{eff},i} (\lambda_i - \lambda'_e) \\ & - \sum_{u=1}^U \dot{\rho}_u(t) + \frac{\Lambda}{1 - \rho} (\dot{\omega}(t) + \omega_d(t) (\omega_d(t) + \lambda'_e)), \end{aligned} \quad (38)$$

where the  $\sim$  symbol is used to emphasize that this quantity is not measured, but rather is to be imposed on the reactor by the controller.  $\rho$  is the measured reactivity,  $\Lambda$  is the neutron mean generation time,  $\beta_{\text{eff}}$  is the effective delayed neutron fraction, and  $\lambda_i$  is the decay constant of the delayed neutron precursor group  $i$  with concentration  $C_i$  and effective delayed neutron fraction  $\beta_{\text{eff},i}$ , with  $\beta_{\text{eff}} = \sum_i \beta_{\text{eff},i}$ . The alternate effective multigroup decay parameter is expressed as:

$$\lambda'_e = \frac{\sum_{i=1}^I \lambda_i^2 C_i}{\sum_{i=1}^I \lambda_i C_i}. \quad (39)$$

The so-called system acceleration is:

$$\dot{\omega}(t) = \frac{\omega_d(t) - \omega(t)}{k\Delta t}, \quad (40)$$

where the value of  $k$  should be chosen to be small enough but not too small [4] (for all cases presented herein,  $k$  is set to 1).

For each feedback mechanism  $u$ , the reactivity rate of change due to that feedback can be expressed as:

$$\dot{\rho}_u(t) = \frac{\partial \rho}{\partial u} \frac{\partial u}{\partial t}, \quad (41)$$

where  $\frac{\partial \rho}{\partial u}$  is the reactivity feedback coefficient corresponding to  $u$ . For all the transients shown in this report,  $U = 3$  different feedback mechanisms are modeled:  $\bar{T}_{\text{fuel}}$ ,  $\bar{T}_{\text{mod}}$ , and  $\theta$ , as detailed in Section 3.1.1.3.

Finally, the required CD rotation speed to be imposed by the controller can be obtained per:

$$\dot{\theta}(t) = \tilde{\rho}(t) \left( \frac{\partial \rho}{\partial \theta} \right)^{-1}, \quad (42)$$

and the CD angle is given by:

$$\theta(t_{n+1}) = \theta_i + \sum_{k=1}^n \dot{\theta}(t_k) \Delta t_k. \quad (43)$$

#### 4.3.2 Reactivity Feedback

In addition to requiring kinetics data—as with most of the PID controllers considered in Section 4.2—this approach also necessitates knowledge of (1) the reactivity coefficients and (2) the

rate of change of the corresponding variables, which could be difficult to measure in real systems. One interesting consideration involves how inaccuracy in the reactivity coefficients can affect the performance of the period-generated control.

To assess the effect of inaccuracy in reactivity coefficients, and because obtaining very accurate reactivity coefficients for all the values of the variables on which the cross-sections depend would be prohibitively costly, a few approximate options were considered in computing the reactivity coefficients.

1. “Step” feedback: after fixing all other variables (e.g.,  $\bar{T}_{\text{fuel}}$  and  $\theta$ ), the multiplication factors of all computed values for the state variable chosen  $\{u_1, u_2, \dots\}$  (e.g.,  $\{150, 335, 520, 705\}$  for  $\bar{T}_{\text{mod}}$ ) are used to compute piecewise-constant feedback coefficients, defined for a  $u$  of between  $u_k$  and  $u_{k+1}$  as:

$$\frac{\partial \rho}{\partial u} \approx \frac{k_{\text{eff}}(u_{k+1}) - k_{\text{eff}}(u_k)}{(u_{k+1} - u_k)k_{\text{eff}}(u_k)}, \quad (44)$$

with the multiplication factors from Serpent and Griffin guaranteed to match, as detailed in Section 3.1.2.

2. “Griffin” feedback: The same procedure is applied, except that the multiplication factors are computed by Griffin and using a much finer grid (e.g., every 5 K for  $\bar{T}_{\text{mod}}$ ). The results actually differ from the step feedback approach because, although the cross-sections are linearly interpolated between state points,  $k_{\text{eff}}$  does not linearly vary in-between; therefore, the reactivity coefficient seen by Griffin is not constant between state points.
3. “Smooth” feedback: after fixing all other variables, a polynomial regression is performed using all the state points for a single state variable  $\{k_{\text{eff}}(u_1), k_{\text{eff}}(u_2), \dots\}$ , with the maximum polynomial order being determined such that the coefficient of determination (commonly referred to as  $R^2$ ) is greater than 0.99, giving:

$$k_{\text{eff}}(u) \approx \sum_{k=0}^K a_k u^k. \quad (45)$$

In addition, the feedback coefficient is approximated by taking the derivative with respect

to  $u$ :

$$\frac{\partial \rho}{\partial u} \approx \sum_{k=1}^K k a_k u^{k-1}. \quad (46)$$

The main reason for doing so is to obtain a continuous feedback coefficient.

For all approaches, since the value obtained for the feedback coefficients for, say,  $\bar{T}_{\text{mod}}$  could vary for different values of  $\bar{T}_{\text{fuel}}$  and  $\theta$ , an average over all the tabulated values of the non-varying state variables can be performed.

In practice, since the fuel feedback coefficient is found to be almost the same whether step or Griffin feedback is used, the former is used in all cases. As for the CD angle  $\theta$ , the Griffin and smooth approaches are virtually identical (see Figure 7), so the latter is used. However, significant differences are observed among all three feedback approximations in regard to moderator temperature, as illustrated in Figure 19, so they will all be tested on the exponential ramp-up benchmark to estimate the impact of this difference on the results. Given that the determination of reactivity coefficients for all conditions, as well as the real-time measurement of fuel and moderator temperature rates of change, could be challenging for a real system, a more thorough sensitivity analysis should be performed to better assess the impact of inaccuracy on the PGC behavior.

### 4.3.3 Results

Figure 20 presents the results obtained with the PGC on the exponential ramp-up benchmark when using the three different moderator temperature feedback approximations detailed in Section 4.3.2. The first observation is that all three give quite satisfactory results, with outstanding behavior during the ramp-up phase and stabilization to the desired power value. The main difference lies in the behavior immediately following the sudden change in demand. The smooth feedback approach does not actually give the smoothest measured power, because Griffin itself sees a discontinuity in feedback coefficients (see the Griffin curve in Section 4.3.2) but happens to give zero power overshoot (though this is not believed to be generally expected). The step feedback approach initially gives a very low overshoot but manifests a significant discontinuity in derivative at 32.4 s due to the average moderator temperature crossing the 520 K state point and subsequently reaching 0.9% before stabilizing to the expected value. As for the Griffin feedback approach, it shows the overall smoothest behavior, with only a small discontinuity in




Figure 19: Example of moderator feedback coefficients obtained with the smooth, step, and Griffin options described in Section 4.3.2.

derivative (also when  $\bar{T}_{\text{mod}}$  exceeds 520 K) but surprisingly exhibits the largest—though still quite satisfactory—power overshoot (i.e., 1.4%).

Therefore, it appears that the PGC indeed suffers from inaccuracy in the feedback coefficients (and by extension, from inaccuracy in the estimated temperature rates of change), as exhibited by errors on the order of 10%. However, its overall behavior remains robust and comparable to the hybrid PID controller, with specific advantages and disadvantages detailed in the next section.




Figure 20: Evolution of the measured/demanded powers and the CD angle for the PGC under various moderator temperature feedback approximations applied to the exponential ramp-up benchmark.

#### 4.4 Summary

The strengths and weaknesses of each presented control strategy tested on the exponential ramp-up benchmark are summarized in Table 3. Specifically:

- Power-driven PID (see Section 4.2.1): The main advantages are that this controller does not

Table 3: Summary of the strengths and weaknesses of the various control strategies considered in this work. **Green** and **red** are used to indicate an advantage or disadvantage, respectively.

Case	Power overshoot	Accurate initial response?	Need coefficient tuning?	Need kinetics data?	Need feedback coefficients?	Need temperature and CD angle change rate?
Power-driven PID	1.1%	No	Yes	No	No	No
Reactivity-driven PID	22.4%	Yes	Yes	Yes	No	No
Hybrid PID	1.1%	Yes	Yes	Yes	No	No
Period-generated	0-1.4%	Yes	No	Yes	Yes	Yes

require any information from the reactor other than its power, yet is still able to achieve a low power overshoot. A major drawback is its initial delay, resulting in an initial power that significantly trails the power setpoint early on. Moreover, the need to adjust the PID constants makes it more difficult to deploy.

- Reactivity-driven PID (see Section 4.2.2): Compared to the power-driven PID, the initial response of this controller is far superior. However, a critical shortcoming lies in its inability to converge to the final setpoint—with a measured power that stabilizes at more than 22% above the demanded power. This is mainly caused by conversion of the power signal into a reactivity signal (see Equation (28)), for which accurate kinetics data are necessary. Adjustment of the PID constants is also needed but is easier than for the power-driven PID controller, as the main constant ( $K'_p$ ) can be set using the CD reactivity curve.
- Hybrid power-reactivity PID (see Section 4.2.3): This controller combines the main advantages of the power- and reactivity-driven PID controllers—namely, an outstanding initial response and a low power overshoot, with the ability to rapidly stabilize to the final power. However, it also inherits the need for coefficient tuning of the six PID constants. Besides, kinetics data are still needed for the reactivity-based component to convert the power demand into a reactivity demand (see Equation (28)). While  $\beta_{\text{eff}}/\Lambda$  can be measured experimentally, determining the delayed neutron precursor concentrations ( $C_i$ ) could be difficult with a real reactor, and would likely need to be obtained by a model—the results of which may be sufficient, since their accuracy is not as crucial as for the reactivity-driven PID control, given it is only used at the beginning of the ramp-up.
- PGC (see Section 4.3): The performance of this controller is generally very similar to that of



the hybrid PID controller, with an excellent initial response and a low power overshoot. On the one hand, its main advantage over the latter is that most of the tuning parameters seem much more straightforward: no proportional constant appears in the error signal expression, while the integral and derivative terms could be omitted for all the transients considered in this work (see Equation (37)).  $\tau$  and  $k$  in Equations (37) and (40), respectively, remain free parameters, but their values did not need to be adjusted in this work. On the other hand, the period-generated strategy not only requires kinetics data (see Equation (39), among others), but also—unlike the hybrid PID strategy—feedback coefficients and the rate of changes of all the variables contributing to reactivity feedback (e.g., fuel/moderator temperature and CD angle). While a reasonable variation of such quantities did not seem to drastically affect overall performance, it nonetheless remains noticeable, and measurement of such quantities would constitute an additional challenge for real systems.

Based on these considerations, the two control strategies selected for moving forward are the hybrid power-reactivity PID and period-generated controllers.

## 5 STARTUP RESULTS

The ultimate goal of this work is to enable a full operational sequence, from startup to pulsed cooldown, as illustrated in Figure 21. This section is limited to the simulation of a startup sequence that is more realistic (yet still simplified) compared to the exponential ramp-up benchmark studied in Section 4. Particular focus is placed on comparing the two control strategies selected in the previous section. Section 6 presents a preliminary study necessary for adding the shutdown and cooldown phases to this simulation.

### 5.1 Assumed Startup Sequence

A simplified startup sequence—to be enhanced in future work—was considered, and it assumes that temperature conditioning and nuclear startup were completed immediately preceding the beginning of the simulation at time  $t = 0$ . In practice, the initial conditions are assumed to be: (1) CD in a critical position with an initial power of 500 kW; (2) an average temperature of 500 K in the fuel and 200 K in the moderator, insulator, and shell, with the axial shapes from Section 3.1.1.1;

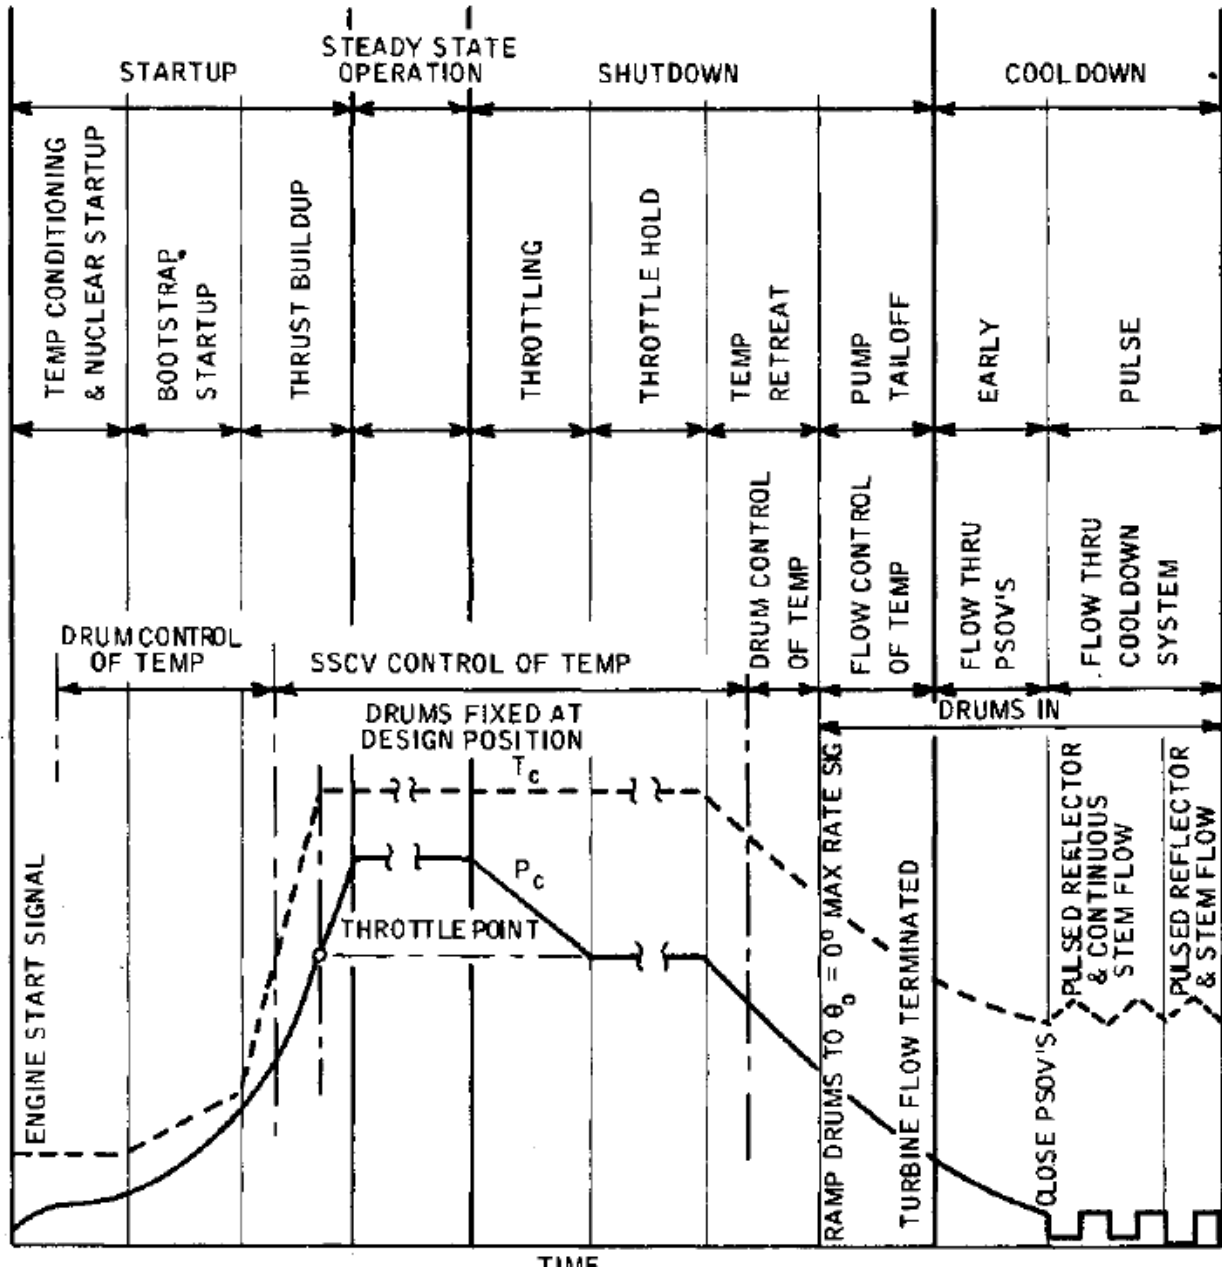


Figure 21: NERVA engine operational phases (picture courtesy of [19]).

and (3) an initial propellant mass flow rate of 2% of the nominal mass flow rate at  $t = 0$  (assuming the cooling effect for  $t < 0$  to be negligible). As discussed in Section 3.3.2, the system model is currently simulated using boundary conditions and is limited to the channels inside the core. In addition, for simplicity, reactivity control of the neutron population was performed solely through CD rotation, rather than also relying on structural support control valves (SSCV) and bypass control valves, and the controller was given a power setpoint rather than both chamber temperature

and pressure setpoints. More specifically, for typical startup sequences, a pressure signal is demanded during bootstrap startup by using both bypass control valves and CD rotation. During thrust buildup, a temperature signal is used up to the throttle point, then bypass control valves are once again relied upon to control the chamber pressure, through turbine power. Temperature signals are achieved by CDs up to a certain temperature and, then by SSCVs thereafter [19].

In this section, the following simplified phases were instead considered:

1. Bootstrap startup for  $0 \leq t \leq 15$  s
2. Thrust build-up (prior to reaching the throttle point) for  $15 \leq t \leq 30$  s
3. Thrust build-up (after reaching the throttle point) for  $30 \leq t \leq 45$  s
4. Steady-state operation for  $t \geq 45$  s.

Specifically, the chamber pressure was chosen to be:

- $p_C(t = 0 \text{ s}) = 0.02 \times p_{C,\text{nom}}$  (47)

- $p_C(t = 15 \text{ s}) = 0.13 \times p_{C,\text{nom}}$  (48)

- $p_C(t = 30 \text{ s}) = 0.65 \times p_{C,\text{nom}}$  (49)

- $p_C(t \geq 45 \text{ s}) = p_{C,\text{nom}}$  (50)

with a linear increase in-between points and the total mass flow rate determined per Equation (15). The exact values were chosen fairly arbitrarily in an attempt to follow the chamber pressure evolution described in [19], including a 65% throttled thrust.

The power setpoint was adjusted as follows:

- $\mathcal{P}_d(t = 0 \text{ s}) = 500 \text{ kW}$  (51)

- $\mathcal{P}_d(t = 13.75 \text{ s}) = 50 \times \mathcal{P}_d(t = 0)$  (52)

- $\mathcal{P}_d(t = 27.5 \text{ s}) = 400 \times \mathcal{P}_d(t = 0)$  (53)

- $\mathcal{P}_d(t \geq 45 \text{ s}) = \mathcal{P}_{\text{tot}} = 315 \text{ MW},$  (54)

with a linear increase in between points. The nominal power is justified in Section 3.3.2, whereas

the other values were chosen to approximately follow the evolution of the chamber temperature, based on values inferred from [19]. It is particularly worth noting that the times at which the changes in power slopes are initiated slightly before the ones for the pressure and mass flow rate (e.g., 13.75 vs. 15 s). This was done to allow the core components to heat up more rapidly (a natural consequence of increasing the power at a rate faster than the coolant mass flow) and then to slow down the heating while approaching the throttle point in order to limit the temperature overshoot. These chamber pressure and power setpoints could be enhanced to further limit this overshoot, but the ideal solution would be to extend the controller strategies so as to directly use the chamber temperature as a setpoint.

The time step  $\Delta t$  was set to 0.05 s and linearly varied to reach 0.1 s at  $t = 15$  s.

## 5.2 Multiphase Startup Results

The results of the simulation using the hybrid PID controller are summarized in Figure 22, showing the evolution of the demanded and measured powers, average fuel and moderator temperatures, outlet coolant temperature (or chamber temperature) and chamber pressure as a function of time. The measured power closely follows the demand, despite having a much larger initial reactivity insertion at the beginning of startup (0.65 vs. 0.16\$ after 1 s, due to having a linear ramp-up instead of an exponential one). It stabilizes to the expected steady-state power, with a power overshoot limited to about 0.3%. Initially, the fuel, coolant, and moderator temperatures slightly decrease before rapidly rising, since the core is not assumed to be in thermal equilibrium. The fuel average temperature reaches almost 2244 K at  $t = 29.5$  s (i.e., shortly before the throttle point is attained) and subsequently stabilizes at around 2165 K. The chamber temperature follows a similar trend, but its maximum of 2861 K is observed at  $t = 30.3$  s, partly because the moderator is still heating up and thus the fuel cooling channel inlet temperature continues to rise accordingly. Note also that it is not unphysical to have an outlet coolant temperature higher than the fuel temperature, because the latter is averaged over the entire core. The chamber temperature stabilizes at 2704 K, but one should keep in mind that adjusted ideal fluid properties are still being used (see




Figure 22: Evolution of power, temperature, and chamber pressure with a hybrid PID controller for the startup transient.

Section 3.3.1). The temperature overshoot, defined as:

$$O_{T_c} = \frac{\max(T_c(t) - T_{C,i})}{T_{C,f} - T_{C,i}}, \quad (55)$$

where  $T_{C,i}$  and  $T_{C,f}$  are the initial and final chamber temperatures, respectively, is then 7.1%. This value should be reduced in the future by either devising an improved power setpoint and/or chamber pressure evolution, or preferably by extending the controllers described in Section 4 to directly accept the chamber temperature as the signal.

The moderator temperature increases much more slowly, partly because it only directly receives 5.4% of the fission power deposited in the core (see Equation (12)), and the rest through

radiation heat transfer in the three gaps successively separating it from the shell, the insulator, and the fuel. It ends up stabilizing at around 685 K. Were the whole of the fission power deposited in the core to be applied to the fuel ( $\chi = 1$ ), the latter would become hotter and the moderator significantly colder. Note however that the current model does not yet preserve the overall mass of moderator so its predicted average value is bound to be inaccurate.




Figure 23: Evolution of power, temperature, and chamber pressure with a PGC for the startup transient.

For comparison purposes, Figure 23 shows the same quantities obtained by using the PGC (with step feedback, as explained in Section 4.3.2). The results look very similar, except that the power and temperature overshoots are slightly smaller (0 and 6.8%, respectively). However, the difference is so small that it cannot be considered a general result. In particular, the slightly lower temperature overshoot is largely due to a slightly lower predicted power at around  $t = 30$  s. The

conclusion in terms of control strategy is that both methods perform virtually identically. The main difference then lies in the disadvantages of each method, as summarized in Table 3.

### 5.3 Computational Cost

A summary of the runtimes required for the simulations described in this report is presented in Table 4, with the exponential ramp-up and multiphase startup referring to the simulations from Section 4 and Section 5, respectively. All the simulations were run on INL’s Sawtooth super-computer, using two nodes with 48 MPI processes each. The number of processors can be further increased, but diminishing returns will be experienced as the number of resources increases (roughly 20 and 35% decreases when using three and four nodes, respectively, as compared to using two), mainly because the convergence tends to degrade with more processors. If less resources are available, the model can also be run with fewer processors, but the runtime would obviously increase (about a 70% increase when using one node instead of two).

Table 4: Runtime in hours using two Sawtooth nodes (96 processors).

Transient	Steady-state runtime (hrs)	Adjoint runtime (hrs)	Number of time steps	Final simulation time (s)	Transient runtime (hrs)
Exponential ramp-up	0.01	0.01	493	49	24
Multiphase startup	0.01	0.01	996	94.6	42

As explained in Section 3.4.3, the steady-state and adjoint calculations are neutronics-only and thus extremely fast. Regarding the actual transients, the runtimes are virtually identical for the hybrid PID and PGCs, which is unsurprising, since the power traces are so similar, so the results are shown only for the latter. For the exponential ramp-up benchmark, shown in Section 4, the computational job was aborted after 24 hours, at which point the simulation had reached  $t = 49$  s, or 493 time steps<sup>6</sup>, which means that about 20.5 time steps were solved on average per hour. For the multiphase startup presented in Section 5, the job wall clock time was limited to 42 hours, at which point the simulation had reached  $t = 94.6$  s, or 996 time steps, which corresponds to an average of 23.7 time steps per hour. This might be slightly higher than for the exponential benchmark, simply because more time steps are performed when the variables are stabilized (and

<sup>6</sup>In a few instances, the Picard iteration failed to converge, and the time step was cut to 0.05 s.

are thus easier to converge): in the latter case, steady-state is reached at around  $t = 30 - 40$  s, with the final simulation time being  $t = 49$  s, as compared to  $t = 45 - 50$  s and around  $t = 95$  s for the former case.

## 6 PRELIMINARY DECAY HEAT CURVE DETERMINATION FOR SHUTDOWN TRANSIENTS

Accurate modeling of decay heat curves is essential for capturing the heat deposition during the shutdown phase following rapid startup and operation transients. Section 6.1 addresses the generation of a reference total decay heat curve (DHC) for the NTP design described in Section 2. A simple model for the DHC is proposed in Section 6.2, based on a generalization of the Kazimi-Todreas equation for decay heat [20]. In future work, the simple model for the DHC, as developed in this section, will be incorporated within the Griffin/Bison/RELAP model described in this report to simulate the shutdown transient following the startup and operational phases.

### 6.1 Computation of Decay Heat Curves

The DHCs were computed with Serpent version 2.1.32 by defining 25 depletion zones (5 radial and 5 axial) in the core (a radial view is reported in Fig. 2). The number of depletion zones was chosen to account for the large axial thermal gradients and the presence of axial and radial reflectors. Three DHCs were computed by simulating nominal power operation for 15, 30, and 60 minutes, followed by 365 days of decay. The durations of the operational phase were chosen to represent typical mission conditions provided by NASA. In computing the DHCs, the following simplifications were made: (1) non-local photon heating was not considered and (2) the Q-value was set to 200 MeV to obtain a conservative value of decay heat. The nominal power was set to 317.4 MW, while the neutron population was set to 500,000 particles for 800 cycles with 200 inactive cycles, leading to 2 pcm uncertainty on the effective multiplication factor.

The Serpent-generated DHCs corresponding to different burning times, denoted in the legend as  $t_b$ , are shown as a function of time from startup in Fig. 24. The decay power slowly builds up during operation to about 1% of the nominal power (recognizable as the quasi-constant portion of



the DHCs) before decaying to 1 kW in less than 10 days (i.e., about 2, 4, and 10 days for the 15, 30, and 60 minute curves, respectively).




Figure 24: Decay heat curve as a function of time from startup.

## 6.2 Simple Model for Decay Heat Curves

Based on [21], Todreas and Kazimi express the approximate decay power as:

$$\frac{P(t)}{P_0} = 0.066 \left( t^{-0.2} - (t + t_b)^{-0.2} \right), \quad (56)$$

where  $P(t)$  and  $P_0$  denote the decay power and the power at shutdown in Watts, while  $t$  and  $t_b$  are the time after shutdown and the burn time at full power in seconds, respectively. As observed from Fig. 25, Equation (56) cannot capture the head and tail of the DHC, with a maximum error exceeding 50%. This is mainly caused by the independence of the exponent (i.e.,  $-0.2$ ) and the pre-exponential factor (i.e., 0.066) from time. Additionally, the values of these coefficients are not tailored to the NTP design under consideration. To overcome these limitations, the following generalized equation was utilized:

$$\frac{P(t)}{P_0} = \alpha_k \left( t^{-\eta_k} + (t + t_b)^{-\eta_k} \right) \quad \forall t \in I_k = [t_{k-1}, t_k], \quad (57)$$

where  $\alpha_k$  and  $\eta_k$  are real numbers obtained from the best fitting of the reference DHC in the time interval  $I_k = [t_{k-1}, t_k)$ . In this work, five time intervals were used. This was sufficient to obtain a root mean square relative error (RMSRE) and mean absolute relative error (MARE) of below 7 and 6%, respectively, for a decay period of 365 days, in regard to all the considered burn times. The values for  $\alpha_k$  and  $\eta_k$  are reported in Table 5 for the different time intervals.

Fig. 25 shows the predictions of Equations (56) and (57) for an operational time of 60 minutes. Equation (57) can more effectively capture the DHC, especially during the initial and final parts of the transient. This observation is quantitatively substantiated by the results shown in Table 6, which reports the maximum relative error and the RMSRE. Equation (57) consistently improves the accuracy of the DHC predictions for all burn times, bringing the RMSRE down from about 30 to 6%, the MARE from 20 to 4%, and the maximum relative error from about 70 to 16% in all cases.




Figure 25: Comparison of the reference Serpent model with the Todreas-Kazimi model (i.e., Equation [56]), and the new best fit (i.e., Equation [57]) for  $t_b = 60$  minutes.

## 7 CONCLUSIONS AND RECOMMENDATIONS

### 7.1 Conclusions

This report presented a Griffin/Bison/RELAP-7 numerical model of a prototypical NTP system designed to simulate rapid startup transients. The physics modeled included full-core neu-

Table 5: Value of parameters in Equation (57).

$k$	$t_{k-1}$	$t_k$	$\alpha_k$	$\eta_k$
1	0 min	10 mins	0.08812231	0.10523162
2	10 min	2.0 days	0.13460035	0.28297623
3	2.0 days	10 days	0.09988645	0.26581696
4	10 days	100 days	0.05553635	0.1813757
5	100 days	365 days	6.6831336	0.55372711

Table 6: MARE, RMSRE, and maximum error for Equations (56) and (57), with respect to the reference Serpent solution, for a shutdown time of 365 days.

Burn time	MARE, %	RMSRE, %	Max error, %
Best fit, Equation (57)			
$t_b = 15$ min	5.21	6.93	18.0
$t_b = 30$ min	3.84	4.98	11.1
$t_b = 60$ min	2.91	3.82	9.88
Todreas-Kazimi, Equation (56)			
$t_b = 15$ min	20.5	25.9	74.1
$t_b = 30$ min	19.6	25.9	73.9
$t_b = 60$ min	19.8	25.5	68.8

tronics, assembly-wise heat conduction, and conjugate heat transfer, with the balance of plant mainly imposed through boundary conditions. In addition, various forms of automated reactivity control were deployed by using MOOSE to autonomously drive the model and simulate the reactor transitioning from assumed initial conditions to nominal power in a fraction of a minute.

To generate the cross-sections of the neutronics model, and in an effort to simultaneously account for the tremendous axial temperature gradients in the reactor and to limit the number of state points required for cross-section generation, the average component temperature and hydrogen densities in the cooling channels were correlated to the average fuel and moderator temperatures, and fixed axial profiles were derived for nominal conditions and then used during the transient. With this approximation, a tractable cross-section library tabulated with fuel/moderator temperatures and CD angles was generated via Serpent. The full-core SPH correction procedure and the CD decussing technology in Griffin, respectively, ensured preservation of the multiplication factor and reaction rates at state points, as well as a reasonably accurate reactivity worth between tabulated CD angles, despite using a coarse mesh. Feedback from other physics was cal-

culated by modeling one representative fuel assembly per ring, along with the corresponding fuel and moderator cooling channels.

Another layer of multiphysics coupling was added to the model to automatically control the drums and attempt to limit power overshoot during startup. Two different technologies presented herein showed outstanding performance in this regard: (1) a novel hybrid PID controller based on both power and reactivity signals, and (2) a PGC that instead relied on kinetics parameters and reactivity coefficients to predict future behavior and adjust the desired signal accordingly. Both approaches were proven successful on a benchmark, exponentially multiplying the power by 500 in 30 seconds with around 1% or less power overshoot. While the former approach required more tuning of internal parameters, the latter additionally required the knowledge of reactivity feedback coefficients and rates of change of the corresponding variables, including fuel and moderator temperature, which could be difficult to dynamically measure for a real NTP system. Fortunately, the performance of the PGC was not drastically degraded as a result of some inaccuracy in these quantities.

Subsequently, a more realistic startup sequence was considered, in which the mass flow rate and outlet pressures were ramped up to model the bootstrap and thrust build-up phases that occur prior to reaching steady-state conditions, thus demonstrating the ability of the hybrid PID and PGCs to handle such transients, with both types of controllers showing very similar behavior. Nevertheless, a significant chamber temperature overshoot (around 7%) was observed, caused by the demanded power signal and assumed mass flow rate and chamber pressure. Alleviating this behavior should be the focus of future work.

Finally, a study was performed to model decay heat post-shutdown, and preparations were made to extend this model to predict shutdown behavior and post-shutdown pulsed cooling requirements.

## **7.2 Future Work**

To enhance the current numerical model, the following non-exhaustive list of tasks to be performed as future work should be considered:

- Enhance the system model to not solely rely on boundary conditions outside the active core.

- Improve the hydrogen fluid properties available in MOOSE, not only to make them more accurate and consistent with NASA properties but also by testing and ensuring their robustness.
- Make the cross-sections depend on the temperature profiles rather than assume a fixed shape that is only valid for steady-state conditions.
- Implement reactor control that follows chamber temperature and pressure signals and is based on both control valves and drums (rather than using a power signal based solely on drums to control reactivity).
- Gain a better understanding of prototypical initial conditions (e.g., after initial temperature conditioning) and adjust the model accordingly.
- Incorporate the shutdown and pulsed cooldown phases into the startup sequence.
- Understand whether radial heat transfer is significant compared to axial heat transfer, and if so, include a full-core heat transfer model.
- Account for the power density spatial distribution within each fuel assembly (through power reconstruction or a heterogeneous transport neutronics model).

With these improvements, the predictions of models based on the methodology presented herein can be made much more reliable, enabling them for use in the design verification and safety analysis of NTP systems.

## 8 Acknowledgments

This manuscript has been authored by Battelle Energy Alliance, LLC under contract no. DE-AC07-05ID14517 with the U.S. Department of Energy. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.

This research made use of the resources of the High Performance Computing Center at Idaho National Laboratory, which is 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.

## REFERENCES

- [1] M. D. DeHart, S. Schunert, and V. M. Labouré, “Nuclear thermal propulsion,” in *Nuclear Reactors* (P. C. L. Pope, ed.), ch. 7, Rijeka: IntechOpen, 2022.
- [2] A. C. Klein, A. Camp, P. McClure, and S. Voss, “Operational Considerations for Fission Reactors Utilized on Nuclear Thermal Propulsion Missions to Mars,” Tech. Rep. NASA/CR20210000387, January 2021.
- [3] Space Nuclear Propulsion Technologies Committee, “Space nuclear propulsion for human mars exploration,” tech. rep., National Academies of Sciences, Engineering, and Medicine, 2021.
- [4] J. A. Bernard and D. D. Lanning, “Considerations in the Design and Implementation of Control Laws for the Digital Operation of Research Reactors,” *Nuclear Science and Engineering*, vol. 110, pp. 425–444, 1992.
- [5] V. Manickam, M. Krecicki, and D. Kotlyar, “Methodology and preliminary results for the transient control of pressure and temperature for nuclear thermal propulsion engines,” in *International Conference on Physics of Reactors 2022 (PHYSOR 2022)*, (Pittsburgh, PA, USA), May 15–20 2022.

- [6] J. Gustafson, "Space nuclear propulsion fuel and moderator development plan conceptual testing reference design," *Nuclear Technology*, vol. 207, no. 6, pp. 882–884, 2021.
- [7] C. J. Permann, D. R. Gaston, D. Andrs, R. W. Carlsen, F. Kong, A. D. Lindsay, J. M. Miller, J. W. Peterson, A. E. Slaughter, R. H. Stogner, and R. C. Martineau, "MOOSE: Enabling massively parallel multiphysics simulation," *SoftwareX*, vol. 11, p. 100430, 2020.
- [8] J. Leppänen, "Serpent– a continuous-energy Monte Carlo reactor physics burnup calculation code.," tech. rep., VTT Technical Research Centre of Finland, 2015.
- [9] Y. Wang, S. Schunert, J. Ortensi, V. Laboure, M. DeHart, Z. Prince, F. Kong, J. Harter, P. Balestra, and F. Gleicher, "Rattlesnake: A MOOSE-based multiphysics multischeme radiation transport application," *Nuclear Technology*, vol. 207, no. 7, pp. 1047–1072, 2021.
- [10] Y. Jung and C. Lee, "PROTEUS-MOC user manual," Tech. Rep. ANL/NE-18/10, Argonne National Laboratory, 2018.
- [11] R. Macfarlane, D. W. Muir, R. M. Boicourt, A. C. Kahler, III, and J. L. Conlin, "The njoy nuclear data processing system, version 2016," 1 2017.
- [12] V. Laboure, J. Ortensi, and A. Hummel, "HTTR 3-D cross-section generation with Serpent and MAMMOTH," Tech. Rep. INL/EXT-18-51317, Idaho National Laboratory, 2018.
- [13] L. Massimo, *Physics of High-Temperature Reactors*. Pergamon Press, first ed., 1976.
- [14] A. Baxter, C. Ellis, M. Fikani, D. Hanson, J. Saurwein, and C. Charman, "Final Report – NGNP Core Performance Analysis, Phase 1," Tech. Rep. 911160, General Atomics, 2009.
- [15] J. Ortensi, Y. Wang, A. Laurier, S. Schunert, A. Hébert, and M. DeHart, "A newton solution for the superhomogenization method: The PJFNK-SPH," *Annals of Nuclear Energy*, vol. 111, pp. 579–594, January 2018.
- [16] V. Labouré, Y. Wang, J. Ortensi, S. Schunert, F. Gleicher, M. DeHart, and R. Martineau, "Hybrid super homogenization and discontinuity factor method for continuous finite element diffusion," *Annals of Nuclear Energy*, vol. 128, pp. 443–454, June 2019.

- [17] J. Ortensi, V. Labouré, S. Schunert, and R. Stewart, “Neutronics enhanced meshing operations (NEMO) users manual version 1.0,” Tech. Rep. INL/MIS-20-00730, Idaho National Laboratory, November 2020.
- [18] S. Schunert, Y. Wang, J. Ortensi, V. Laboure, F. Gleicher, M. DeHart, and R. Martineau, “Control rod treatment for fem based radiation transport methods,” *Annals of Nuclear Energy*, vol. 127, pp. 293–302, May 2019.
- [19] J. H. Altseimer, G. F. Mader, and J. J. Stewart, “Operating characteristics and requirements for the NERVA flight engine,” *Journal of Spacecraft and Rockets*, vol. 8, no. 7, pp. 766–773, 1971.
- [20] N. E. Todreas and M. S. Kazimi, *Nuclear Systems, Volume 1*. CRC Press, 2-nd ed., 1990.
- [21] G. I. Bell and S. Glasstone, *Nuclear Reactor Theory*. 10 1970.



## A Attached File Summary

This appendix describes the various files needed to run the model.

### A.1 Input Files

For each transient, two input files are needed for the steady-state calculation, one for the adjoint calculation, and four for the transient calculation. The transient calculation should be performed last, as it depends on both the steady-state and adjoint outputs (see Section 3.4.3 for more details).

The steady-state input files are:

1. `ntp_fullcore_neut_initial.i`: Eigenvalue neutronics input, with a fixed temperature distribution.
2. `TH_1ch_tr_initial.i`: Initial thermal-hydraulics input to establish the flow with which to begin the transient.

The adjoint input file is:

1. `ntp_fullcore_neut_adj.i`: Adjoint neutronics input, with a fixed temperature distribution.

The transient input files are:

1. `hybrid_power_pid_controller.i` or `mit_controller.i`: Hybrid PID or period-generated controller (as described in Section 4.2.3 and Section 4.3, respectively).
2. `ntp_fullcore_neut_tr.i`: Transient neutronics input, detailed in Section 3.1.
3. `fuel_elem_HC_30_tr.i`: Transient 30-degree-slice fuel assembly heat transfer input, detailed in Section 3.2.
4. `TH_1ch_tr.i`: Transient thermal-hydraulics input, detailed in Section 3.3.

The following position file is needed to position the various models with respect to each other:

1. `ring_positions.txt`: List of center points of each ring's "representative" fuel assembly (to position the five instantiated pairs of sub-apps `{fuel_elem_HC_30_tr.i, TH_1ch_tr.i}`).

## A.2 Mesh Files

The following mesh (Exodus) files are used for the neutronics and full-core thermal models, respectively:

1. `nasa_fuelhom_3D.e`: 3-D full-core, homogenized neutronics mesh shown in Figure 6.
2. `30_degree_3gaps_v1.e`: 2-D 30-degree-slice mesh shown in Figure 8 and extruded to 3-D in `fuelelem_HC_30_tr.i`.

## A.3 Cross-Section Files

The following cross-section files are used for the various neutronics calculations:

1. `nasa_ntp_lib112.xml`: Library containing the tabulated macroscopic cross sections for the 112 state points described in Section 3.1.1.3.
2. `nasa_ntp_equiv_lib112.corrected.xml`: Library containing the tabulated SPH factors for the 112 state points described in Section 3.1.1.3.