



Automated Power-Following Control for Nuclear Thermal Propulsion Startup and Shutdown Using MOOSE-based Applications

June 2023

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 Power-Following Control for Nuclear Thermal Propulsion Startup and Shutdown Using MOOSE-based Applications

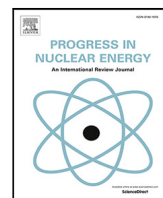
Vincent M Laboure, Sebastian Schunert, Stefano Terlizzi, Zachary M Prince, Javier Ortensi, Ching-Sheng Lin, Lise Cecile Madeleine Charlot, Mark D DeHart

June 2023

**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**



Automated power-following control for nuclear thermal propulsion startup and shutdown using MOOSE-based applications

Vincent M. Labouré^{a,*}, Sebastian Schunert^a, Stefano Terlizzi^a, Zachary M. Prince^a,
Javier Ortensi^a, Ching-Sheng Lin^b, Lise M. Charlot^a, Mark D. DeHart^a

^a Idaho National Laboratory, 1955 N Fremont Ave, Idaho Falls, 83415, ID, USA

^b TerraPower, LLC, 15800 Northup Way, Bellevue, 98005, WA, USA

ARTICLE INFO

Keywords:

Automated control
Nuclear thermal propulsion
Startup
Shutdown
MOOSE
Griffin
Bison
RELAP-7

ABSTRACT

This paper describes an investigation of various automated control strategies applied to a full-core multiphysics Griffin/Bison/RELAP-7 model of a prototypical nuclear thermal propulsion system. In all cases, control is achieved by actuating control drums based on the demanded power and predicted quantities from the numerical model. One key finding is that hybrid proportional integral derivative controllers – a novel type of controller that uses both power and reactivity predicted signals – can demonstrate a level of performance rivaling that of period-generated control. The former requires parameter tuning, while the latter mostly necessitates providing reactivity coefficients and temperature rates of change, which could be very challenging to accurately measure in real-time.

In addition, decay heat plays an important role in determining cooling requirements during cooldown phases. A decay heat model, accounting for burn time and throttling, was derived, and is incorporated within the model to simulate the steady-state and shutdown phases and satisfyingly follow the power demand. However, temperature overshoots and non-constant specific impulse during throttling will necessitate further improvements.

1. Introduction

Nuclear thermal propulsion (NTP) systems for interplanetary travel (e.g., manned missions to Mars) rely on nuclear reactors – instead of more conventional chemical rockets – 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 (DeHart et al., 2022) and reducing the travel time (or, alternatively, increasing the payload) by that same amount.

NTPs were initially introduced in 1947 by the Applied Physics Laboratory (Ruark, 1947) before their development began in 1955 at Los Alamos Scientific Laboratory as part of Project Rover. After the launch of Sputnik by Russia and the subsequent space race, the U.S. National Aeronautics and Space Administration (NASA) was created, in 1958, and became the lead of that effort (Spence, 1968). Later, the Rover project became civilian and focused on providing the Saturn V rocket with an upper stage powered by a nuclear engine. In 1961, the Nuclear Engine for Rocket Vehicle Application (NERVA) program

was created within NASA, with the ambition of designing, producing and testing nuclear rocket engines. Many ground-based reactors were tested as a result, including KIWI and PEWEE. A major milestone was achieved with the NRX-XE reactor which conducted 28 burns for a total operation time exceeding 3.5 h (Stewart, 2015).

Wernher von Braun, who served as the NASA Marshall Space Flight Center director, then promoted the idea of sending astronauts to Mars by August 1982 using NTP-propelled rockets, thinking that “although the undertaking of this mission will be a great national challenge, it represents no greater challenge than the commitment made in 1961 to land a man on the moon” (NASA, 2018). Although NERVA ended after the Cold War, the NASA’s Game Changing Technology for Deep Space Exploration Program initiated in the mid 2010s has been leveraging the experience learned under NERVA. But NTP development continues to face significant challenges which are very different from those encountered in other nuclear applications, including the need to deal with temperatures exceeding 2700 K (National Academies of Sciences, Engineering, and Medicine, 2021).

* Corresponding author.

E-mail addresses: vincent.laboure@inl.gov (V.M. Labouré), sebastian.schunert@inl.gov (S. Schunert), stefano.terlizzi@inl.gov (S. Terlizzi), zachary.prince@inl.gov (Z.M. Prince), javier.ortensi@inl.gov (J. Ortensi), clin@terrapower.com (C.-S. Lin), lise.charlot@inl.gov (L.M. Charlot), mark.dehart@inl.gov (M.D. DeHart).

<https://doi.org/10.1016/j.pnucene.2023.104710>

Received 20 January 2023; Received in revised form 6 April 2023; Accepted 19 April 2023

Available online 5 May 2023

0149-1970/© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

In addition, a NTP system operates at full power for a significantly longer duration (i.e., on the order of a few hours) than traditional nuclear reactors, with each burn typically lasting 15–60 min [Klein et al., January 2021]. Conversely, the reactor must be able to reach full power within one minute (Space Nuclear Propulsion Technologies Committee, 2021), mostly to maximize the total impulse. This can be challenging to achieve without entailing significant power and/or temperature overshoots (exceeding the 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) (Bernard and Lanning, 1992), initially developed at the Massachusetts Institute of Technology, allow for better anticipation by using current reactor data (e.g., kinetics parameters and 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 (Manickam et al., 2022).

The primary purpose of this paper is to investigate various automated control strategies applied to a full-core multiphysics model of a prototypical NTP system. At this point, the scope of the study is limited to control achieved by actuating control drums (CDs) based on the demanded power and predicted quantities from the numerical model—whereas NTP control systems typically rely on both CDs and valves, and are set to follow chamber temperature and pressure set points rather than a power set point (Altseimer et al., 1971). This is why the types of controllers considered herein are classified as *power-following*.

One key finding is that hybrid PID controllers – a novel type of controller that uses both power and reactivity predicted signals – can demonstrate a level of performance just as outstanding as that of period-generated control in terms of following the power demand with little overshoot. This is significant because the former approach is purely reactive (in the sense that it only applies correction based on the current state of the system), whereas the latter is made predictive by attempting to anticipate the time evolution of the system. PID controllers require parameter tuning, while PGCs mostly necessitate providing reactivity coefficients and temperature rates of change, which could be very challenging to accurately measure in real-time. However, preliminary results indicate that a certain level of inaccuracy may not drastically reduce PGC performance. Overall, the accuracies shown by the two approaches during simplified NTP startup phases is quite similar, with very comparable, and satisfying, power overshoots. Nevertheless, temperature overshoots remain a potential issue, largely caused by trying to follow a power demand rather than a temperature one.

For shutdown phases, a simple PID controller relying solely on the power predicted by the numerical model was considered, and was deemed sufficient in following the power demand. However, once again, the desired behavior of the chamber temperature (which is to be maintained at a constant value during throttling in order to maximize the specific impulse) can be difficult to achieve without a temperature-following control strategy.

In addition to control considerations, decay heat plays an important role in determining cooling requirements during shutdown phases. The secondary purpose of this paper is to derive a decay heat model – one that accounts for burn time and throttling – and incorporate it within the controller model to simulate the steady-state and shutdown phases.

The remainder of this paper is organized as follows: an overview of the design is given in Section 2, and the full-core multiphysics model, described in detail in [Labouré et al., 2022], is summarized 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 chamber pressure is

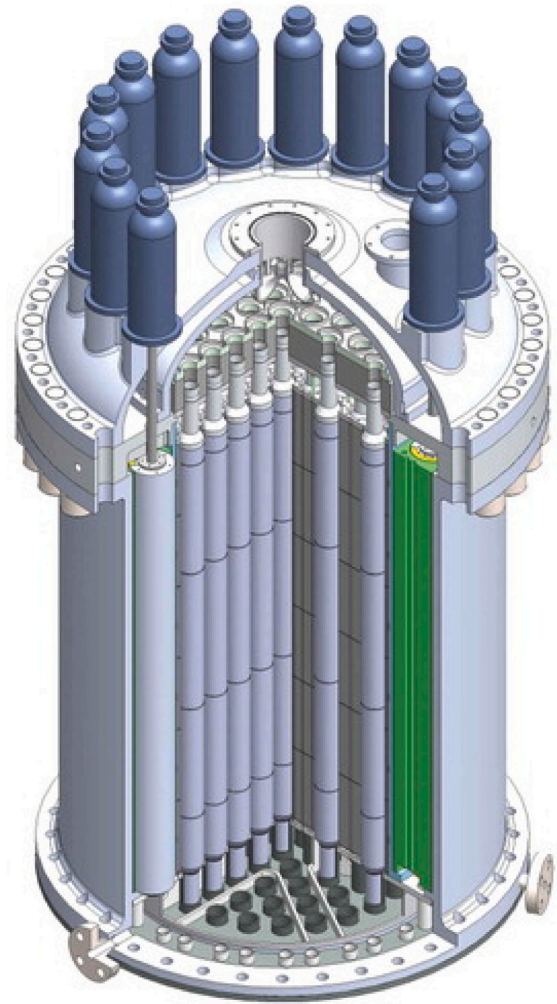


Fig. 1. BWX Technologies NTP reactor concept.
Source: Picture courtesy of Gustafson (2021).

ramped up to model the bootstrap and thrust buildup phases prior to reaching steady-state conditions. In Section 6, the decay heat study is detailed, and shutdown results – including throttling, throttle hold, and temperature retreat – are presented. Conclusions are drawn in Section 7.

2. Prototypal design

This section gives a brief overview of the prototypical reactor considered, with the purpose of examining various methods of studying NTP systems. Therefore, the design considered in this study 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 Fig. 1 and is largely identical to the one further presented in DeHart et al. (2022). An overview of the active core design can be found in Fig. 2. Specifically, the reactor core is composed of 61 high-assay low-enrichment uranium ceramic-metallic (CERMET) fuel assemblies arranged in circular rings within a zirconium hydride (ZrH) monolith (in magenta in Fig. 2) which acts 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_2 coolant flowing

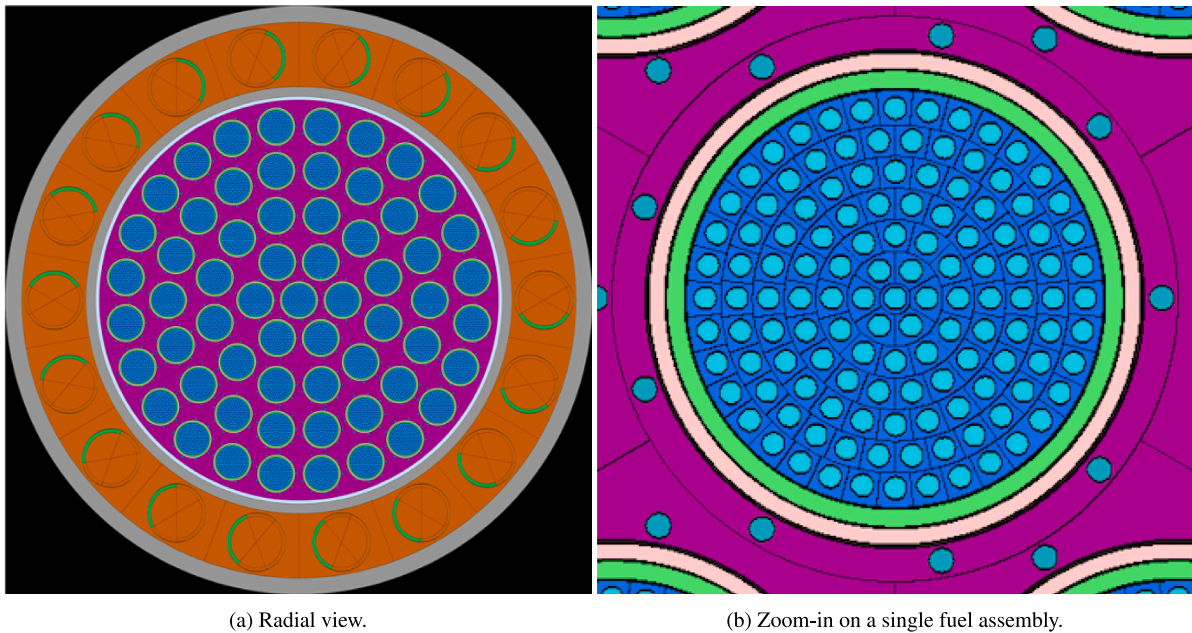


Fig. 2. Overview of the Serpent geometry.

downward. To simplify cross-section generation, nine representative moderator cooling channels (in cyan), with H_2 coolant flowing upward, are modeled around each fuel assembly.

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

Fig. 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 Fig. 3) and divided into two paths: (1) one portion is preheated by cooling the nozzle and then being flowed 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 for 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 in duration, as is typical of a NTP system designed for Mars missions.

3. Full-core multiphysics model

To limit the length of this article and focus on automated control strategies, only a summary of the full-core multiphysics model is presented. A more complete description can be found in [Labouré et al., 2022]. The full-core multiphysics model is based on several single physics models coupled together via the Multiphysics Object-Oriented Simulation Environment (MOOSE) framework [Permman et al., 2020]. The neutronics, heat transfer, and thermal-hydraulics models are presented in Sections 3.1–3.3. 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

Table 1
Energy group (upper) boundaries for the General Atomics nine-group structure [(Baxter et al., 2009)] 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 | | |

the Monte Carlo code Serpent (Leppänen et al., 2015), using the 3-D full-core heterogeneous model depicted in Fig. 2 for various core conditions, then employed in the multiphysics model by the Griffin neutronics input [Wang et al., 2021; Jung and Lee, 2018]. To reduce the computational cost of the multiphysics transient simulation, a full-core super homogenization (SPH) procedure (Ortensi et al., 2018) is then performed to exactly reproduce the multiplication factor, power density profile, and integrated leakage out of the SPH regions at each state point (Labouré et al., 2019), even with a coarse mesh (shown in Fig. 4 and generated using the NEMO software (Ortensi et al., 2020)). A cusping treatment technique is applied to accurately model CD rotation, with a static mesh (Schunert et al., 2019).

The energy group structure studied in [Labouré et al., 2022] and summarized in Table 1 was selected, and the multigroup cross-section library was then tabulated with respect to the average fuel and moderator temperatures in K (\bar{T}_{fuel} and \bar{T}_{mod} , respectively) and the CD angles in degrees (with 0 corresponding to the maximum reactivity):

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

It thus contains a total of 112 state points, each obtained via a Serpent calculation performed with 2,000 cycles of 5×10^5 particles after 100 inactive cycles, yielding a reported standard deviation in multiplication factor of less than 3 pcm. In Griffin, a multilinear interpolation scheme is used to compute cross-sections and SPH factors between temperature state points.

The feedback effects on the cross sections are limited to fuel and moderator temperature, with other components' temperatures and the

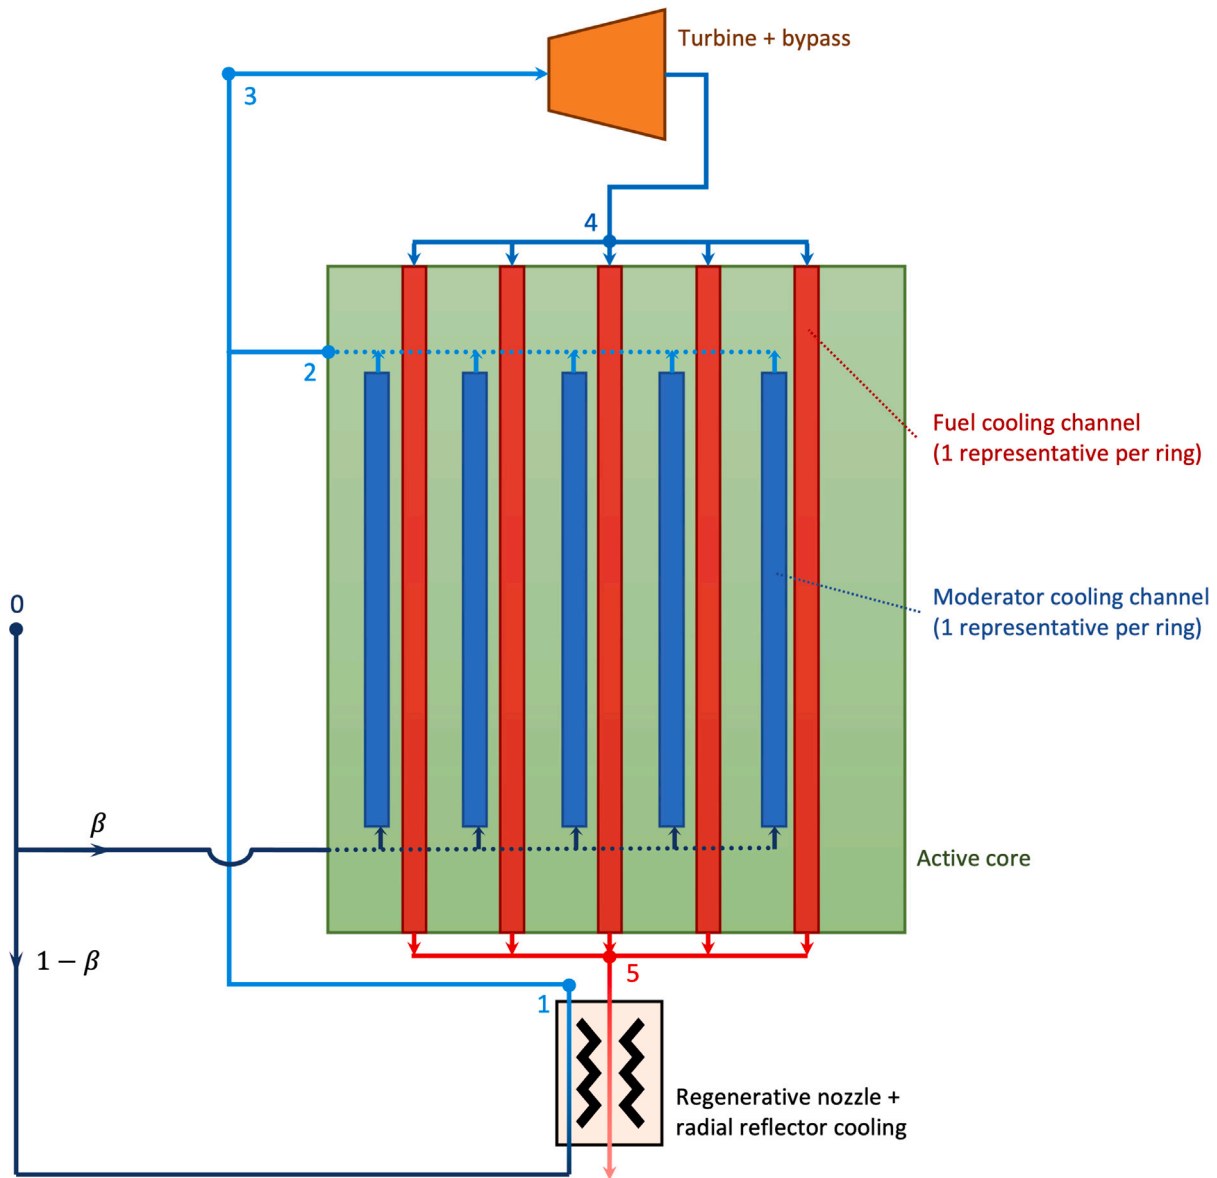


Fig. 3. Sketch of the NTP system considered.

hydrogen density being assigned using correlations and axial profiles. These were determined using a single representative fuel assembly and associated channels with various total mass flow rates and power levels, as detailed in [Labouré et al., 2022]. In particular, the Serpent model is not directly coupled to the multiphysics model and, rather, assumes that these correlations and profiles are solely a function of average fuel and moderator temperatures (which is bound to introduce inaccuracy, since various mass flow rates could correspond to the same average temperatures). The assessment of the accuracy of the current cross-section library between state points – in terms of errors due to (1) interpolation, (2) using assumed correlations and axial profiles, and (3) from the use of SPH factors – will be the focus of future work.

Thermal mechanical feedback is not currently considered, though it was recently shown that it can noticeably impact maximum temperatures and the pressure drop across the cooling channels [Krecicki and Kotlyar, 2022].

With the exception of Section 6 – which specifically covers decay heat during steady-state and shutdown phases – 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 (e.g., radial reflector

and nozzle), 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 P_{tot} (see [Eq. (3)]). In the heat transfer model, the split between moderator and fuel heating is performed directly.

3.2. Fuel assembly heat transfer

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 thus 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 slice of the fuel assembly was modeled for each ring of assemblies, along with its corresponding fuel and moderator cooling channels, with the mesh shown in Fig. 5 being extruded over the entire fuel height. Three gap heat transfer models were successively used to

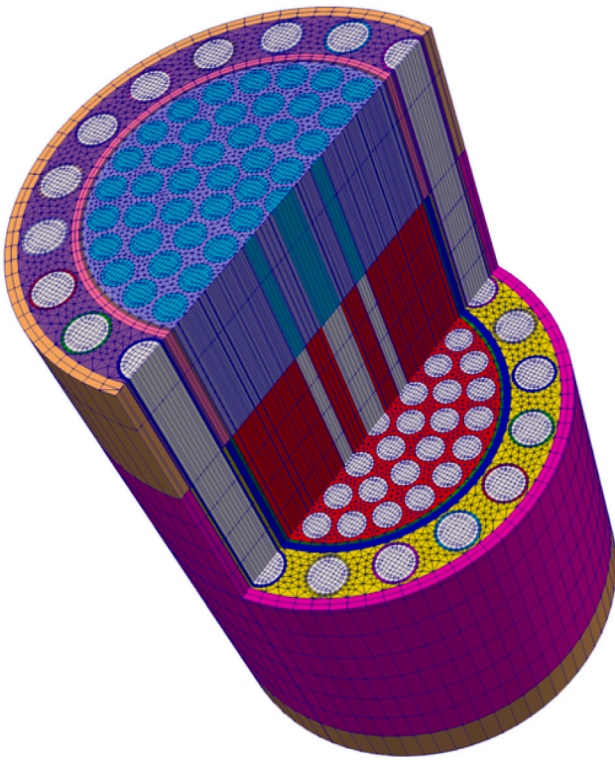


Fig. 4. Full-core neutronics mesh.
Source: Picture courtesy of (DeHart et al., 2022).

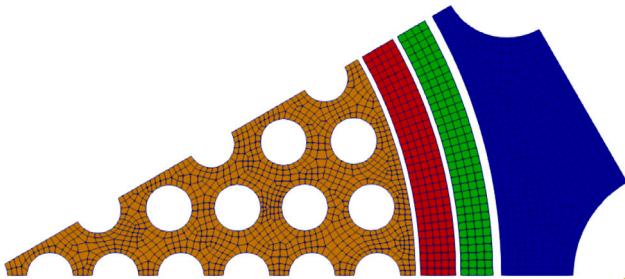


Fig. 5. X-Y view of the 30-degree fuel assembly slice of the thermal mesh.
Source: Picture courtesy of (DeHart et al., 2022).

model heat transfer in the gaps between the fuel, insulator, shell, and moderator.

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 for each assembly ring (of index r) over a given axial layer of index i . Next, a volume correction factor was applied to account for volume changes between homogeneous and heterogeneous power densities. A fixed fraction χ – defined by Eq. (4) – was assumed to be deposited in the fuel, while the remaining amount was assumed to be deposited in the moderator. Specifically, the heterogeneous power density deposited in the fuel of ring r and axial layer i was given by:

$$p_{\text{fuel},r,i} = \chi \frac{V_{\text{assm}}}{V_{\text{fuel}}} p_{\text{assm},r,i}. \quad (1)$$

The one deposited in the moderator was given by:

$$p_{\text{mod},r,i} = (1 - \chi) \frac{V_{\text{assm}}}{V_{\text{mod}}} p_{\text{assm},r,i}. \quad (2)$$

where $V_{\text{assm},i}$ is the volume of a homogenized fuel assembly, and V_{fuel} and V_{mod} are the heterogeneous volumes of fuel and moderator, respectively, in that assembly (once again, these volumes should be assembly-dependent, but are currently computed based on the mesh shown in Fig. 5). This approach preserves the 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 (a phenomenon known as spatial self-shielding), as seen in Fig. 6(a). Either a heterogeneous neutronics model or form functions would be needed to account for this effect, and will be the focus of future work.

3.3. Thermal-hydraulics

The thermal-hydraulics model is summarized in Fig. 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, 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. In addition, the mass flow rate is divided equally across all fuel cooling channels. The same simplifying assumption is made for the moderator cooling channels. The standard Dittus–Boelter correlation is used to compute the heat transfer coefficients. Other correlations will be tested in the future, including the Wolf–McCarthy correlation [McCarthy and Wolf, 1960], which was recently implemented in MOOSE.

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 the NASA over a temperature range of 50–2700 K. This approximation tends to underestimate hydrogen temperatures of less than 2700 K, and to overestimate them from that point on. This should eventually be addressed to obtain accurate transient results.

3.3.2. System model

Currently, the multiphysics model is limited to the active core region shown in Fig. 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). NASA obtained these values from a previous study on a prototypical CERMET design:

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

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

$$P_{\text{tot}} \equiv P_{\text{fuel}} + P_{\text{mod}}. \quad (3)$$

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

$$\chi \equiv \frac{P_{\text{fuel}}}{P_{\text{tot}}} = 94.6\%, \quad (4)$$

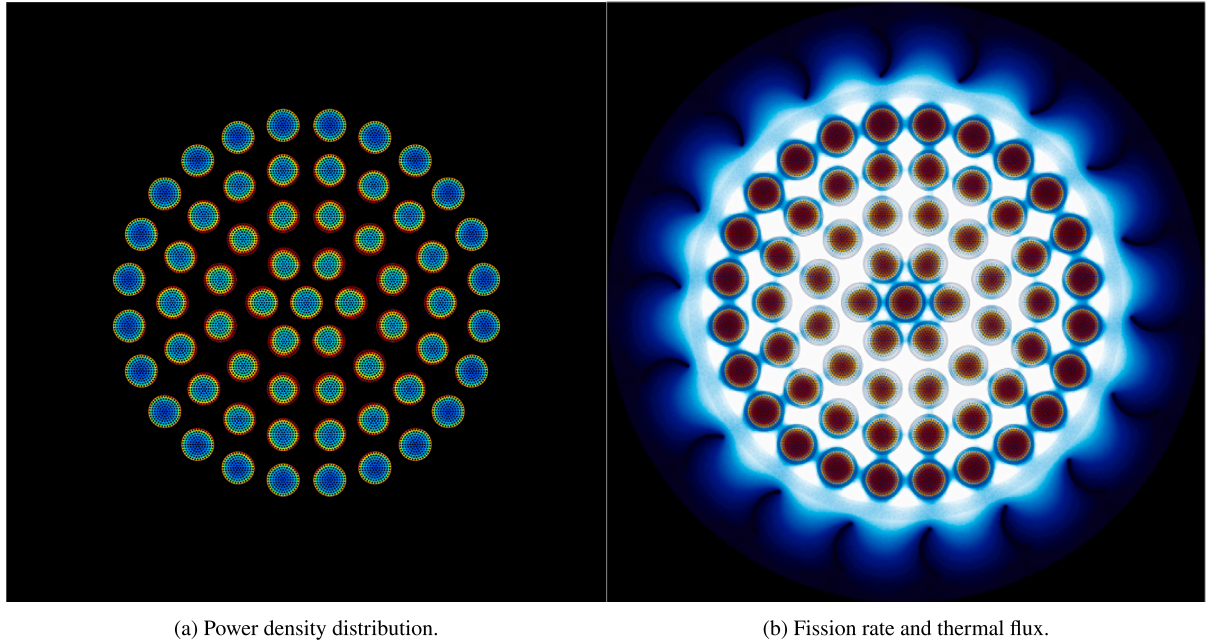


Fig. 6. Typical Serpent results obtained with the CD rotated halfway ($\theta = 90^\circ$).

A simplified system model – mimicking the system described in Section 2 and illustrated in Fig. 3 – was then developed, relying on (among other things) 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 β 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 divided by the square root of the chamber temperature T_C (Manickam et al., 2022) (which correspond to the pressure and temperature at point 5 in Fig. 3, respectively) and is given by:

$$\dot{m}(t) = \frac{\dot{m}_{\text{nom}} \sqrt{T_{C,\text{nom}}}}{p_{C,\text{nom}}} \frac{p_C(t)}{\sqrt{T_C(t)}}, \quad (5)$$

where the subscript nom represents nominal value quantities (i.e., at steady-state), with $\dot{m}_{\text{nom}} = 7.51$ kg/s, $T_{C,\text{nom}} = 2700$ K and $p_{C,\text{nom}} = 6.89$ MPa. In practice, although not necessary, the mass flow rate is only updated once per time step to accelerate convergence of the multiphysics coupling.

- 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 NASA study):

$$p_3(t) = \alpha p_4(t). \quad (6)$$

Many of these assumptions will be relaxed in the future when a more elaborate system model is considered.

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 Fig. 7, with a generic power-following controller that determines the angle of the CD, based on a power set point for a given quantity of

interest and the data measured¹ or predicted by the full-core model, as 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 Fig. 3.

All the controllers considered in this work serve to determine 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 control strategies rely on the total power, though some necessitate kinetics data, feedback coefficients, and/or temperature change rates.

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. Fixed-point iterations were not performed at the controller level within a given time step. Adding them should be possible – though, for stability reasons, restrictions on the time step size may ensue – provided that the time 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 work 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

¹ The reason the terms “measured” and “predicted” are used interchangeably throughout the remainder of this work 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.

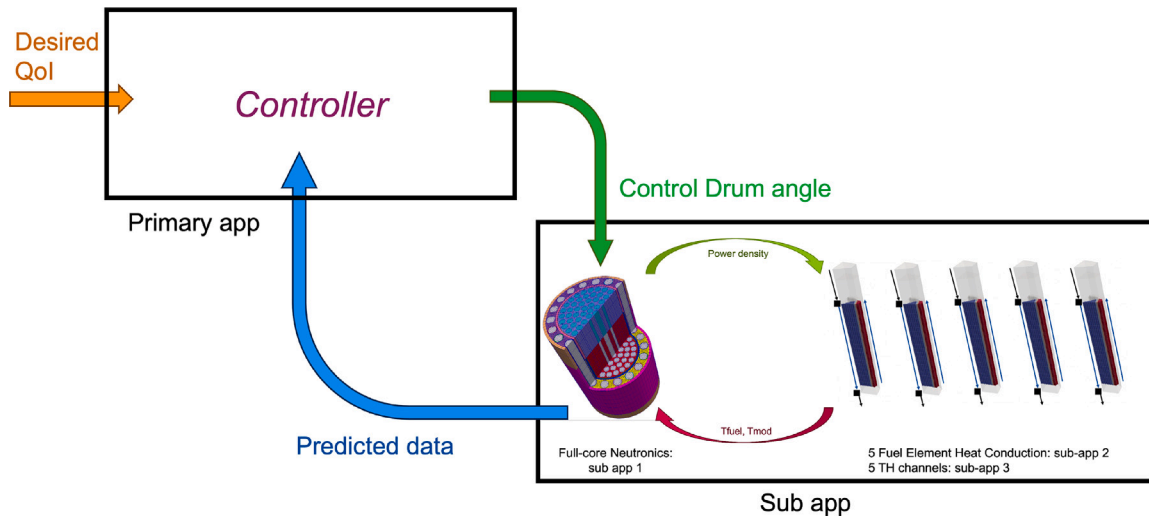


Fig. 7. Schematics of the full-core multiphysics model.

initial critical state. Fixed temperature distributions are assumed in the fuel and moderator, 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 conditions are 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.

A summary of the multiphysics coupling strategy is provided in Algorithm 1. All the solves are performed with an Implicit Euler time discretization.

Algorithm 1 Multiphysics coupling strategy.

```

1: while Final time is not reached do
2:   Compute the next control drum angle.
3:   Update mass flow rate (using Equation (5)).
4:   while Outer Picard Iterations tolerances not met do
5:     Solve neutronics.
6:     Transfer power density to FAHT.
7:     while Inner Picard Iterations tolerances not met do
8:       Solve FAHT.
9:       Transfer wall temperature to thermal-hydraulics.
10:      Solve thermal-hydraulics.
11:      Transfer fluid temperature and HTC to FAHT.
12:    end while
13:    Transfer fuel/moderator temperature to neutronics.
14:  end while
15:  Transfer quantities needed by controller.
16: end while

```

Notations. FAHT: fuel assembly heat transfer; HTC: heat transfer coefficient.

4. Reactivity control strategies

This section focuses on the study of various power-following strategies for automatically controlling 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 exist inside a nuclear reactor (DeHart et al., 2022), especially considering the very short period of time within which a NTP must reach full power. For the purpose of this work, no maximum CD rotation speed was enforced.

4.1. Exponential ramp-up benchmark

To assess various approaches, a simple yet challenging benchmark was devised with an exponentially increasing power set point (i.e., demanded power that the system is to follow) given by:

$$P_d(t) = \begin{cases} P_i \left(\frac{P_f}{P_i} \right)^{t/t_f} & , \quad 0 \leq t < t_f \\ P_f & , \quad \text{otherwise,} \end{cases} \quad (7)$$

where $P_i = 0.5$ MW and $P_f = 250$ 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, as defined by:

$$O_p = \frac{\max(P(t))}{P_f} - 1. \quad (8)$$

For this benchmark, the following further assumptions are made:

- The initial CD angle is set to 120 degrees.
- The average temperatures in the fuel and moderator regions are initially set to 500 and 200 K, respectively, with representative axial shapes (described in [Labouré et al., 2022]).
- The total mass flow rate – and therefore the chamber pressure, given Eq. (5) – is set to its nominal value throughout the transient.
- The inlet fluid temperatures for the moderator and fuel cooling channels are also assumed to be 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

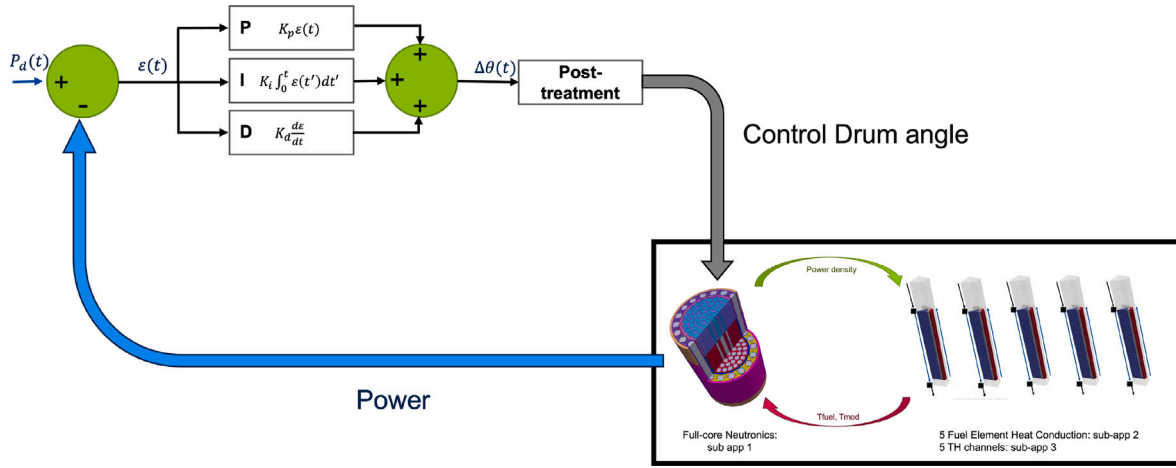


Fig. 8. Schematics of power-driven PID control of the full-core multiphysics model.

low values, resulting in convergence issues. Use of fixed inlet temperatures enables assessment of such methods' behavior after this initial phase.

- The time step Δt 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 being computed directly by using the demanded and measured powers (P_d and P , respectively):

$$\varepsilon(t) = P_d - P. \quad (9)$$

In this work, the measured power is always computed via the 3-D full-core neutronics model. Next, the following rotation angle is applied to the drums:

$$\Delta\theta_P(t) = K_P \varepsilon(t) + K_I \int_0^t \varepsilon(t') dt' + K_D \frac{d\varepsilon}{dt}, \quad (10)$$

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 $n+1$):

$$\theta_P(t_{n+1}) = \theta_i + \sum_{k=1}^n \Delta\theta_P(t_k), \quad (11)$$

where θ_i is the initial CD angle. This strategy is further illustrated by Fig. 8.

Nevertheless, a disadvantage of PID control is that the constants must often be manually tuned, especially for complex systems whose transfer functions are not easily determined, thus rendering theoretical determination of optimum values arduous, if not impossible. Trial and error then becomes the standard approach. Fig. 9 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 set point, 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 set point 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 increasingly 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 the 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 significantly trails the set point 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 controller is sooner able to closely follow the demand. 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.

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 cause degraded stability for large values and would thus require careful tuning.

4.2.2. Reactivity-driven PID signal

Another approach, illustrated in Fig. 10, 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 obtained from the demanded and measured reactivities (ρ_d and ρ , respectively) per:

$$e(t) = \rho_d - \rho. \quad (12)$$

The following rotation angle is then applied to the drums:

$$\Delta\theta_\rho(t) = K'_P e(t) + K'_I \int_0^t e(t') dt' + K'_D \frac{de}{dt}, \quad (13)$$

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, θ_ρ , is computed using an equation similar to Eq. (11). Following the methodology described in DeHart et al. (2022), 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 CD reactivity curve, by setting K'_P to the inverse of the reactivity

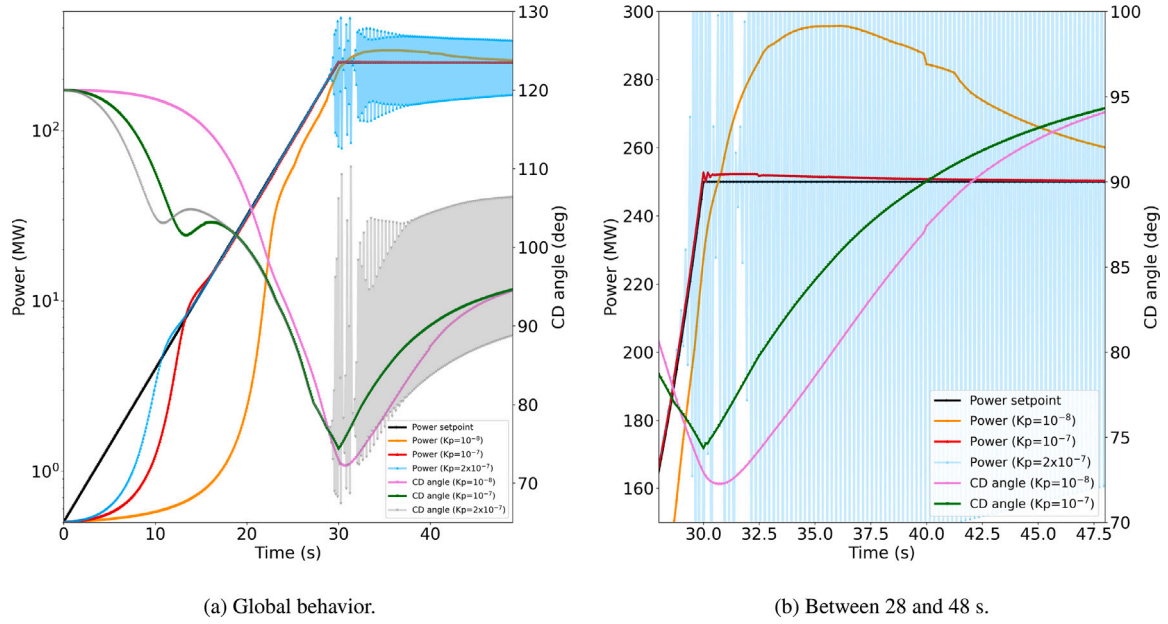


Fig. 9. Evolution of the power and CD angles for various values of K_p (in deg/W) for the power-driven PID controller when applied to the exponential ramp-up benchmark.

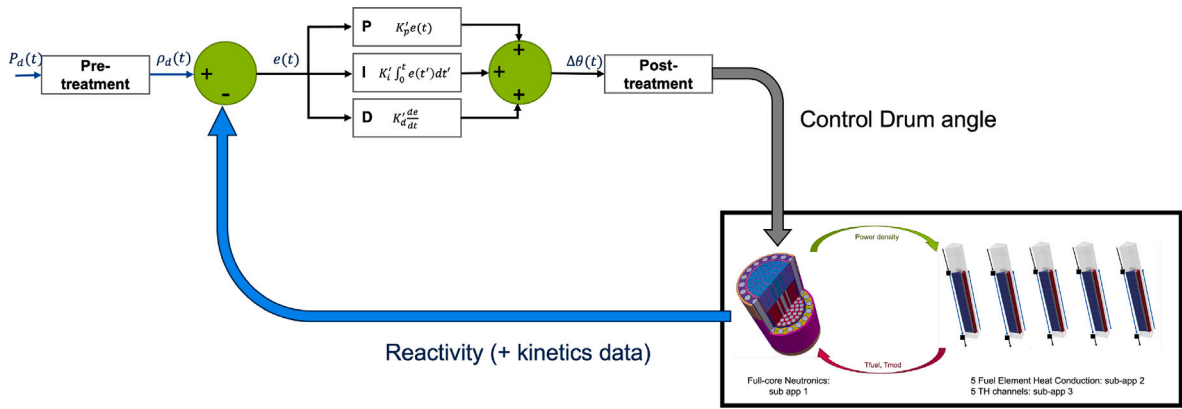


Fig. 10. Schematics of reactivity-driven PID control of the full-core multiphysics model.

inserted per degree in the operating range of the CD, thus yielding approximately $K'_p = 25$ deg/\$. By studying the early behavior of the controller when applied to a simple case, K'_i was set to 0.25 deg/\$-s, while K'_d was left at 0.

If the true set point is the power (as was the case in the benchmark), this signal must first be converted into a reactivity signal, which can be done by rearranging the PKE as follows:

$$\rho_d = \frac{\Lambda}{P_d \beta_{\text{eff}}} \left(\frac{dP_d}{dt} - \sum_i \lambda_i C_i \right) + 1, \quad (14)$$

where Λ is the neutron mean generation time, β_{eff} is the effective delayed neutron fraction, and λ_i is the decay constant of the delayed neutron precursor group i with concentration C_i . The resulting control strategy is illustrated by Fig. 8.

Nonetheless, a clear disadvantage of this method – not considered in DeHart et al. (2022), since a given reactivity set point 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 ρ_d (see Equation [(14)]) or the measured reactivity ρ from the core can result in the final power stabilizing to a value other than the nominal one.

In fact, the results displayed in Fig. 11 show very good agreement between the demanded and measured powers during the exponential ramp-up phase, suggesting that the reactivity set point 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) because the power is notably higher than its nominal value.² Yet, in Griffin, cross sections are not extrapolated past the last interpolation point, but are instead kept at this final 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.

Fig. 12 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, θ_ρ . The first observation is

² 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.

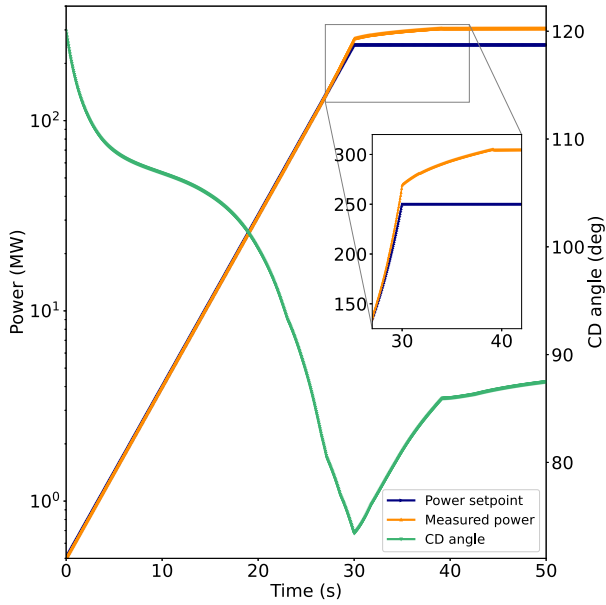


Fig. 11. Evolution of the power and CD angle for the reactivity-driven PID controller when applied to the exponential ramp-up benchmark.

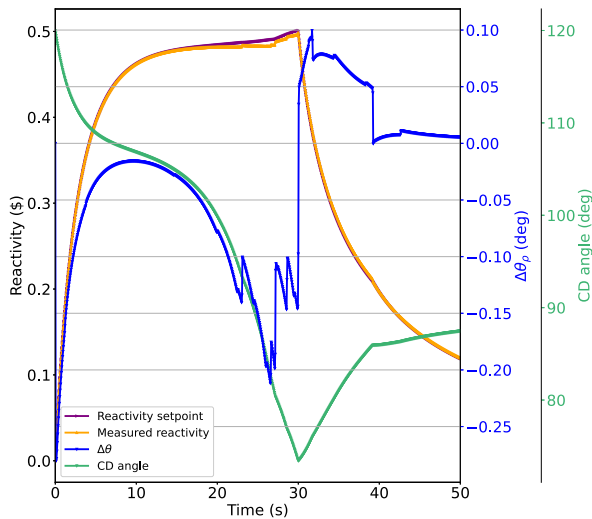


Fig. 12. Evolution of the measured and demanded reactivities, as well as the incremental rotation angle and CD angle for the reactivity-driven PID controller applied to the exponential ramp-up benchmark.

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 set point if the measured reactivity is consistently lower than the signal. This is likely due to an inaccuracy in the computation of the measured kinetics data – especially the mean generation time which can change significantly with temperature – used in Eq. (14). Indeed, as explained in Section 3.4.3, the computation of these kinetics data in Griffin relies on an adjoint solution computed based on 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 controller 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 set point. A power-reactivity-driven hybrid strategy combining the two approaches was thus devised, and is illustrated in Fig. 13.

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

$$\Delta\theta = \xi \Delta\theta_p(t) + (1 - \xi) \Delta\theta_r(t), \quad (15)$$

where ξ was chosen to be defined as³:

$$\xi \equiv \frac{\ln\left(\frac{P}{P_i}\right)}{\ln\left(\frac{P_f}{P_i}\right)}, \quad (16)$$

where P_i and 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 sections⁴: $K_p = 10^{-7}$ 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:

$$\theta(t_{n+1}) = \theta_i + \sum_{k=1}^n \Delta\theta(t_k) = \theta_i + \sum_{k=1}^n (\xi \Delta\theta_p(t_k) + (1 - \xi) \Delta\theta_r(t_k)). \quad (17)$$

Fig. 14 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., $\Delta\theta \approx \Delta\theta_r$; 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 (i.e., 1.1%) that is virtually identical to the one observed when using the power-driven PID.

4.3. Period-generated control of the drum rotation

Another control strategy, this one differing from standard PID control by attempting to impose a reactor period based on kinetics data and feedback coefficients, was developed in Bernard and Lanning (1992) and recently applied to a PKE-based NTP system model (Manickam et al., 2022).

³ Note that a linear definition of ξ , such as $\xi \equiv \frac{P - P_i}{P_f - P_i}$, gave virtually identical results for this benchmark.

⁴ Note, however, that they could likely be further enhanced.

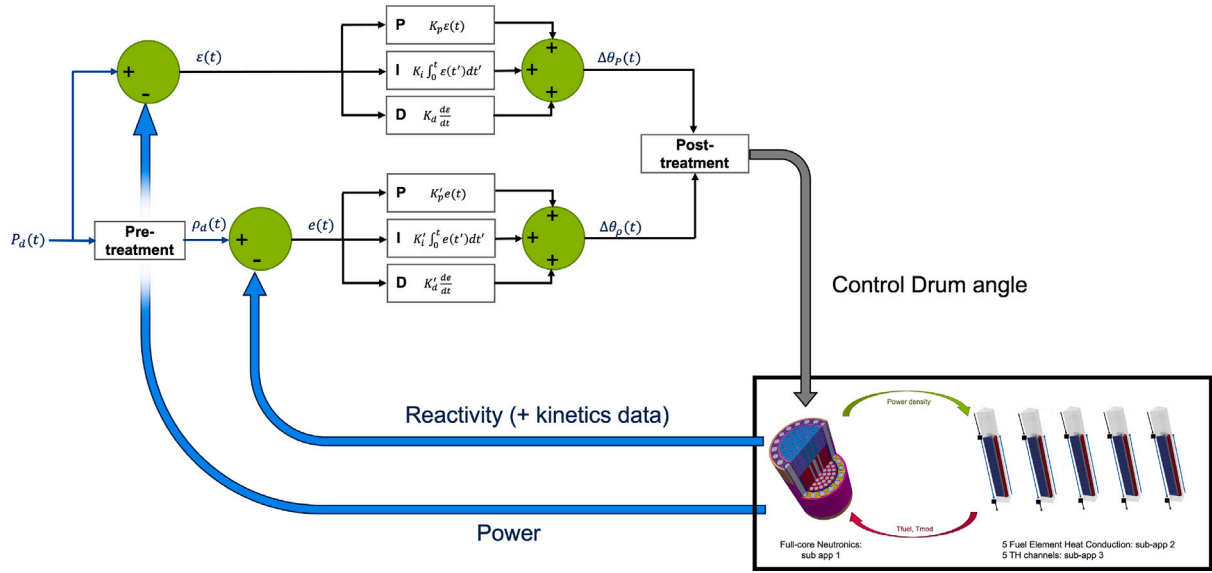


Fig. 13. Schematics of hybrid power-reactivity PID control of the full-core multiphysics model.

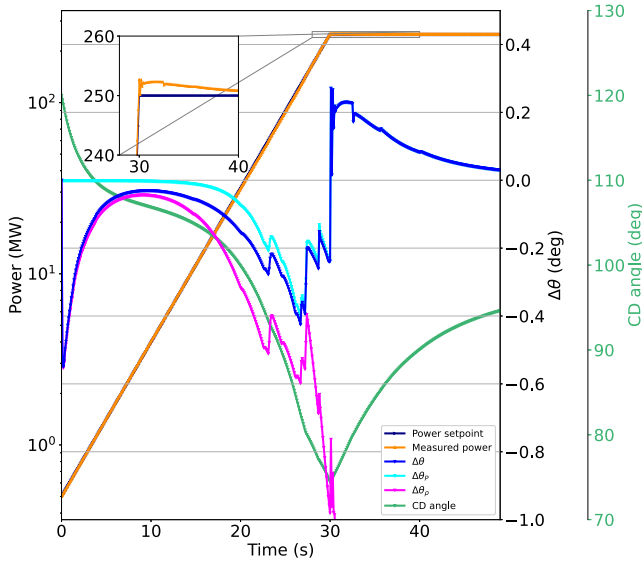


Fig. 14. Evolution of the measured and demanded powers, as well as the CD signal and CD angle for the hybrid PID controller when applied to the exponential ramp-up benchmark.

4.3.1. Methodology

The method, summarized in Fig. 15, relies on the inverse of the (measured) reactor period:

$$\omega(t) = \frac{\dot{P}(t)}{P(t)}, \quad (18)$$

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

$$e(t) = \ln \left(\frac{P_d(t + j\Delta t)}{P(t)} \right), \quad (19)$$

where j is a positive integer that should be set to “at least 2 to ensure stability against oscillations” (Bernard and Lanning, 1992) and Δt is the time step size. Unlike the PID error signal, this error signal does not go to zero if the demanded and predicted powers exactly match, since this controller looks at the future demand rather than the current discrepancy between the demanded and measured powers.

The demanded inverse period was originally expressed as Bernard and Lanning (1992):

$$\omega_d(t) = \frac{e(t) + \frac{1}{T_i} \int_0^t e(t') dt' + T_d \frac{de}{dt}}{j\Delta t}. \quad (20)$$

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

$$e(t) = \ln \left(\frac{P_d(t)}{P(t)} \right). \quad (21)$$

This sign of this error actually informs whether the measured power is trailing or anticipating the demand. In addition, for this work, it was decided to replace $j\Delta t$ with a single tuning parameter τ to avoid artifacts whenever Δt was not constant throughout the simulation, thus yielding the following modified expression for the demanded inverse period:

$$\omega_d(t) = \frac{e(t) + \frac{1}{T_i} \int_0^t e(t') dt' + T_d \frac{de}{dt}}{\tau}. \quad (22)$$

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:

$$\omega_d(t) = \frac{\ln \left(\frac{P_d(t+\tau)}{P(t)} \right)}{\tau}, \quad (23)$$

where $\tau = 0.5$ s.

The required reactivity change rate was then introduced (Bernard and Lanning, 1992):

$$\begin{aligned} \tilde{\rho}(t) = & (\beta_{\text{eff}} - \rho(t)) \omega_d(t) - \lambda'_e \rho(t) - \sum_{i=1}^I \beta_{\text{eff},i} (\lambda_i - \lambda'_e) \\ & - \sum_{u=1}^U \tilde{\rho}_u(t) + \frac{\Lambda}{1 - \rho} (\dot{\omega}(t) + \omega_d(t) (\omega_d(t) + \lambda'_e)), \end{aligned} \quad (24)$$

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. ρ is the measured reactivity, Λ is the neutron mean generation time, β_{eff} is the effective delayed neutron fraction, and λ_i is the decay

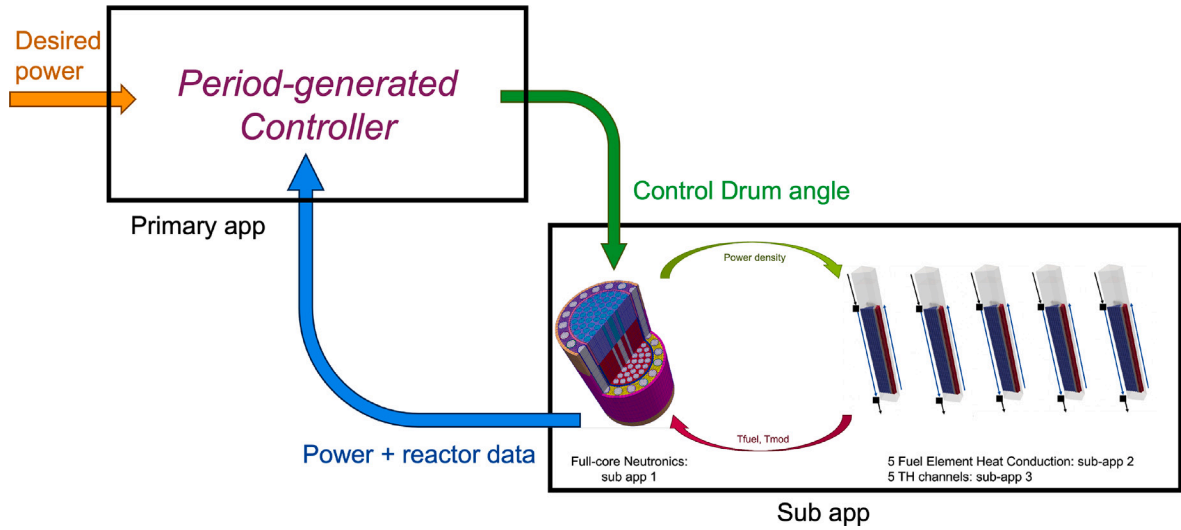


Fig. 15. Schematics of period-generated control of the full-core multiphysics model.

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:

$$\lambda'_e = \frac{\sum_{i=1}^I \lambda_i^2 C_i}{\sum_{i=1}^I \lambda_i C_i}. \quad (25)$$

The so-called system acceleration is:

$$\dot{\omega}(t) = \frac{\omega_d(t) - \omega(t)}{k \Delta t}, \quad (26)$$

where the value of k should be chosen to be small enough but not too small (Bernard and Lanning, 1992) (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:

$$\dot{\rho}_u(t) = \frac{\partial \rho}{\partial u} \frac{\partial u}{\partial t}, \quad (27)$$

where $\frac{\partial \rho}{\partial u}$ is the reactivity feedback coefficient corresponding to u and the rate of change of u is computed at the end of each time step n as:

$$\frac{\partial u}{\partial t} \approx \frac{u^n - u^{n-1}}{\Delta t}. \quad (28)$$

For all the transients shown in this work, $U = 3$ different feedback mechanisms are modeled: \bar{T}_{fuel} , \bar{T}_{mod} , and θ , as detailed in Section 3.1.

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

$$\dot{\theta}(t) = \tilde{\rho}(t) \left(\frac{\partial \rho}{\partial \theta} \right)^{-1}. \quad (29)$$

The CD angle is given by:

$$\theta(t_{n+1}) = \theta_i + \sum_{k=1}^n \dot{\theta}(t_k) \Delta t_k. \quad (30)$$

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 inaccurate 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 approximations were considered in computing the reactivity coefficients.

1. “Step” feedback: After fixing all other variables (e.g., \bar{T}_{fuel} and θ), 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}_{mod}) are used to compute piecewise-constant feedback coefficients, defined for a u of between u_k and u_{k+1} as:

$$\frac{\partial \rho}{\partial u} \approx \frac{k_{\text{eff}}(u_{k+1}) - k_{\text{eff}}(u_k)}{(u_{k+1} - u_k) k_{\text{eff}}(u_k)}, \quad (31)$$

with the multiplication factors from Serpent and Griffin being guaranteed to match, as detailed in Section 3.1.

2. “Griffin” feedback: The same procedure is applied, except that the multiplication factors are computed by Griffin and a much finer grid is used (e.g., every 5 K for \bar{T}_{mod}). The results actually differ from the step feedback approach because, although the cross sections are linearly interpolated between state points, k_{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:

$$k_{\text{eff}}(u) \approx \sum_{k=0}^K a_k u^k. \quad (32)$$

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

$$\frac{\partial \rho}{\partial u} \approx \sum_{k=1}^K k a_k u^{k-1}. \quad (33)$$

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

For all approaches, since the value obtained for the feedback coefficients for, say, \bar{T}_{mod} could vary for different values of \bar{T}_{fuel} and θ , 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 regardless of whether step or Griffin feedback is used, the former is used in all cases. As for the CD angle θ , the Griffin and smooth feedback approaches are virtually identical, so the latter is

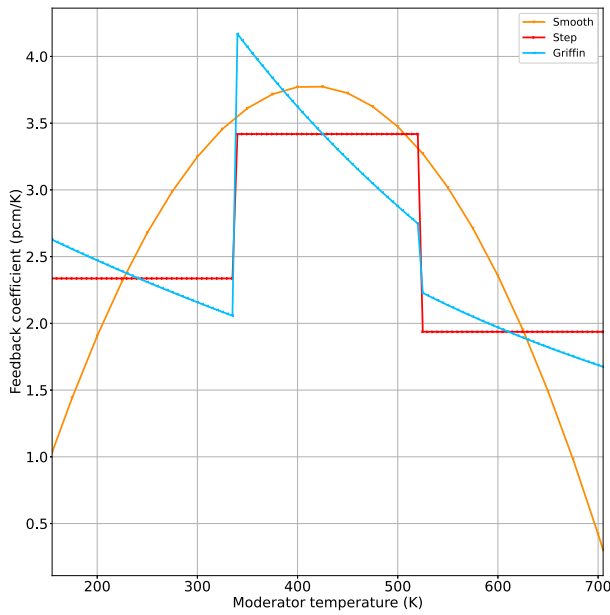


Fig. 16. Example of moderator feedback coefficients obtained via the smooth, step, and Griffin feedback options described in Section 4.3.2.

used. However, because significant differences are observed among all three feedback approximations in regard to moderator temperature (see [Fig. 16]), they will all be tested on the exponential ramp-up benchmark to estimate the impact that this difference has 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 change rates, 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

Fig. 17 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 Fig. 16) 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% overshoot before stabilizing to the expected value. As for the Griffin feedback approach, it shows the overall smoothest behavior, with only a small discontinuity in derivative (also when \bar{T}_{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 change rates), as exhibited by variations on the order of 10%. However, its overall behavior remains robust and comparable to that of the hybrid PID controller, with its specific advantages and disadvantages being detailed in the next section.

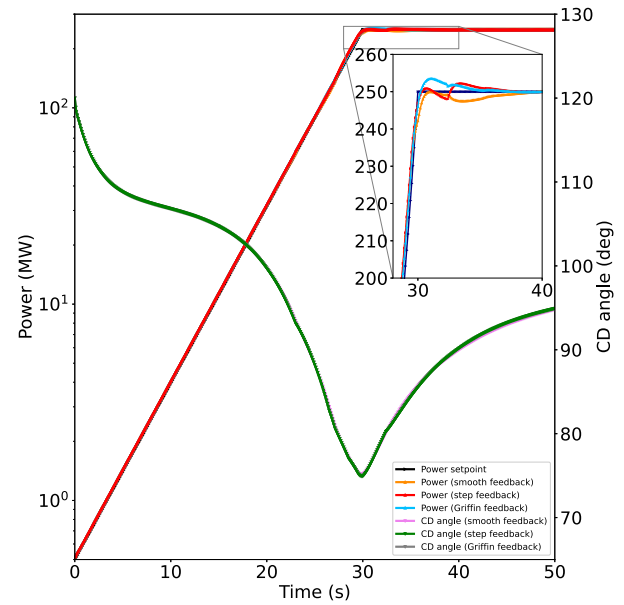


Fig. 17. Evolution of the measured/demanded powers and CD angle for the PGC under the various moderator temperature feedback approximations applied to the exponential ramp-up benchmark.

4.4. Summary

The strengths and weaknesses of each presented control strategy tested on the exponential ramp-up benchmark are summarized in Table 2. Specifically:

- **Power-driven PID** (see Section 4.2.1): The main advantages are that this controller does not require any information from the reactor other than its power, yet can still achieve a low power overshoot. A major drawback is its initial delay, resulting in an initial power that significantly trails the power set point 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 set point—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 [Eq. (14)]), 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 [Eq. (14)]). Both $\beta_{\text{eff}}/\Lambda$ and the delayed neutron precursor concentrations (C_i) can be obtained experimentally [Ciftcioglu and Geckinli, 1980], so this approach could be deployed on a real reactor.
- **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 the tuning of parameters generally seems much more straightforward:

Table 2Summary 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? | Needs coefficient tuning? | Needs kinetics data? | Needs feedback coefficients? | Needs 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 |

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 [Eq. (23)]). τ and k in Eqs. (23) and (26), 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 (e.g., see [Eq. (25)]) but also – unlike the hybrid PID strategy – feedback coefficients and the change rates 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 and would likely need to be achieved via a numerical model.

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

5. Startup phases

The ultimate goal of this work is to enable a full operational sequence, from startup to pulsed cooldown, as illustrated in Fig. 18. This section is limited to the simulation of a startup sequence that is more realistic (yet still simplified) than 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 focuses on the shutdown phases.

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 2.5 MW; (2) an average temperature of 500 K in the fuel and 200 K in the moderator, insulator, and shell, all featuring representative axial shapes (described in [Labouré et al., 2022]); and (3) an initial chamber pressure of 2% of the nominal chamber pressure 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 (SSCVs) and bypass control valves (BCVs), and the controller was given a power set point rather than both chamber temperature and pressure set points. More specifically, for typical startup sequences, a pressure signal is demanded during bootstrap startup by using both BCVs and CD rotation. During thrust buildup, a temperature signal is used up to the throttle point, then BCVs are once again relied upon to control the chamber pressure through turbine power. Temperature signals are achieved by CD rotation at up to a certain temperature, and by adjusting SSCVs thereafter (Altseimer et al., 1971).

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

1. Bootstrap startup for $0 \leq t \leq 15$ s

2. Thrust buildup (prior to reaching the throttle point) for $15 \leq t \leq 30$ s
3. Thrust buildup (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:

$$p_C(t = 0 \text{ s}) = 0.02 \times p_{C,\text{nom}} \quad (34)$$

$$p_C(t = 15 \text{ s}) = 0.13 \times p_{C,\text{nom}} \quad (35)$$

$$p_C(t = 30 \text{ s}) = 0.65 \times p_{C,\text{nom}} \quad (36)$$

$$p_C(t \geq 45 \text{ s}) = p_{C,\text{nom}}, \quad (37)$$

with a linear increase in between points and the total mass flow rate determined per Eq. (5). The exact values were chosen fairly arbitrarily in an attempt to follow the chamber pressure evolution described in Altseimer et al. (1971), including a 65% throttled thrust.

The power set point was adjusted as follows:

$$P_d(t = 0 \text{ s}) = 2.5 \text{ MW} \quad (38)$$

$$P_d(t = 14 \text{ s}) = 10 \times P_d(t = 0) \quad (39)$$

$$P_d(t = 28 \text{ s}) = 80 \times P_d(t = 0) \quad (40)$$

$$P_d(t \geq 45 \text{ s}) = P_{\text{tot}} = 315 \text{ MW}, \quad (41)$$

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 Altseimer et al. (1971). It is particularly worth noting that the changes in power slope occur slightly before the ones for the chamber pressure (e.g., 14 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 limit the temperature overshoot by slowing down the heating while approaching the throttle point. These chamber pressure and power set points 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 set point.

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

5.2. Startup results

The results of the simulation using the hybrid PID controller are summarized in Fig. 19, 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 as compared to the benchmark from Section 4.1 (0.32 vs. 0.16\$ after 1 s, due to the ramp-up being linear instead of exponential). It stabilizes to the expected steady-state power, with a power overshoot limited to about 0.4%. 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 2207 K at $t = 40.3$ s and subsequently stabilizes at around 2177 K. The chamber temperature follows a similar trend, but

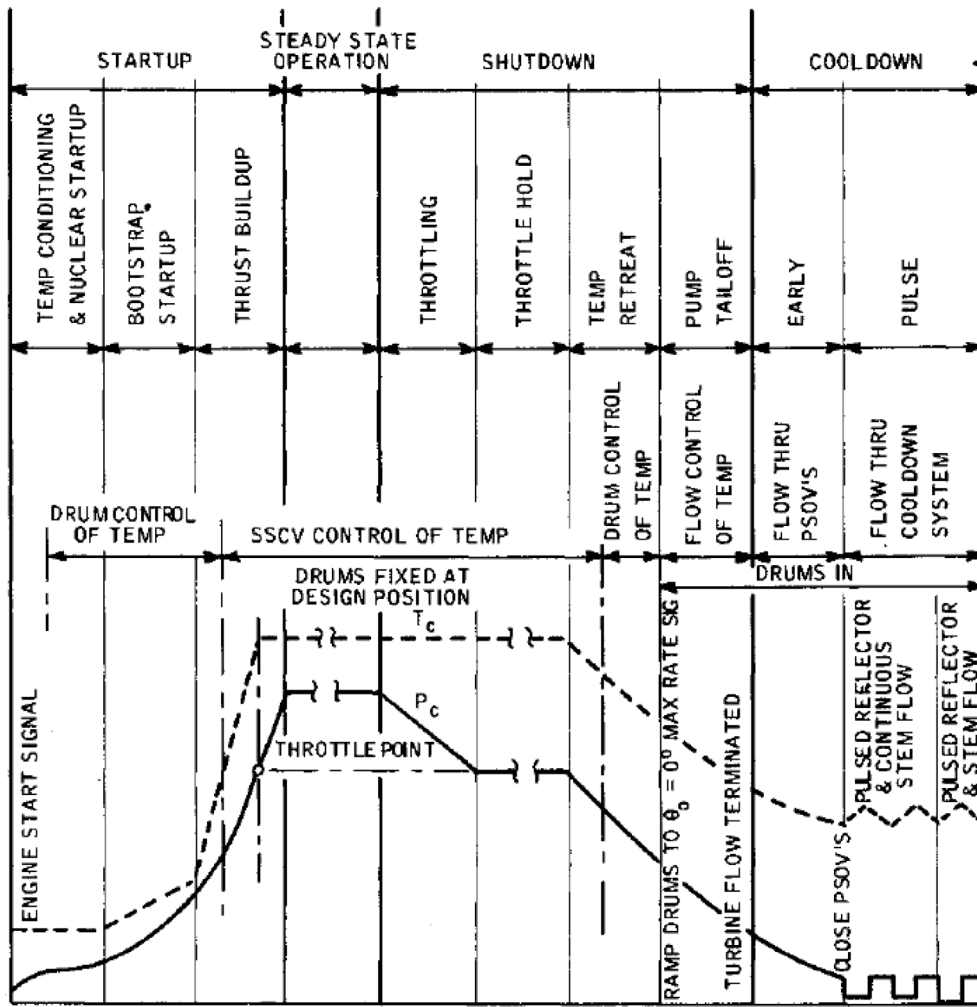


Fig. 18. NERVA engine operational phases.

Source: Picture courtesy of (Altseimer et al., 1971).

its maximum of 2786 K is observed at $t = 38.4$ s. Note also that it is not nonphysical 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 2720 K, but one should keep in mind that adjusted ideal fluid properties are still being used (see Section 3.3.1). The temperature overshoot, defined as:

$$O_{T_C} = \frac{\max(T_C(t) - T_{C,i})}{T_{C,f} - T_{C,i}} - 1, \quad (42)$$

where $T_{C,i}$ and $T_{C,f}$ are the initial and final chamber temperatures, respectively, is then 3.0%. This value should be reduced in the future by either devising an improved power set point 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 [Eq. (4)]), while the rest comes through radiation heat transfer in the three gaps successively separating it from the shell, insulator, and fuel. 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.

For comparison purposes, Fig. 20 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 2.1%, respectively).

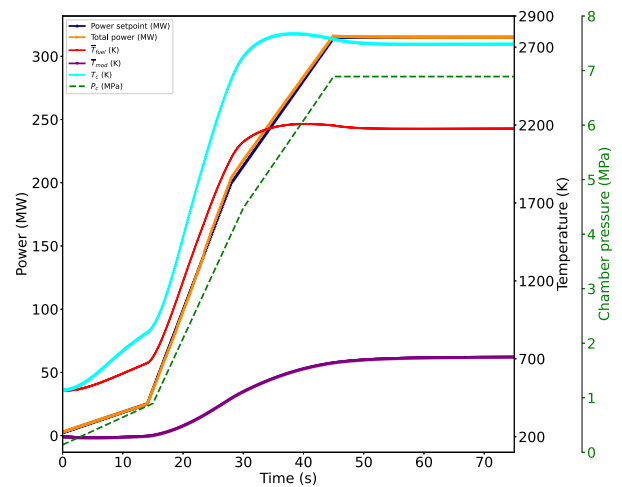


Fig. 19. Evolution of power, temperature, and chamber pressure with a hybrid PID controller for the startup transient.

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

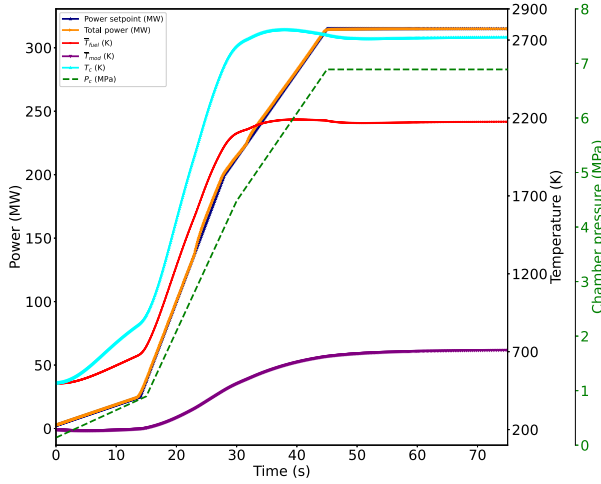


Fig. 20. Evolution of power, temperature, and chamber pressure with a PGC for the startup transient.

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 2.

6. Shutdown phases

In this section, the shutdown phases – namely, throttling, throttle hold, and temperature retreat – are considered. Section 6.1 details a decay heat model that is vital to the accurate modeling of these phases and to eventually adding the cooldown phases to this simulation. The shutdown results are then presented in Section 6.2.

6.1. Decay heat model

Accurate modeling of decay heat curves (DHCs) is essential for capturing the heat deposition during the shutdown and cooldown phases following steady-state operations. Section 6.1.1 addresses the generation of a reference total DHC for the NTP design described in Section 2. A simple model for the DHC is proposed in Section 6.1.2, based on a generalization of the Kazimi–Todreas equation for decay heat (Todreas and Kazimi, 1990).

6.1.1. 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 so as 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 min, 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 decay heat value. The nominal power was set to 317.4 MW, while the neutron population was set to 500,000 particles for 800 cycles in addition to 200 inactive cycles, leading to an uncertainty of 2 pcm on the effective multiplication factor.

In Fig. 21, the Serpent-generated DHCs, corresponding to different burn times and denoted in the legend as t_b , are shown as a function of time from startup. During operation, the decay power slowly builds up to about 5% 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).

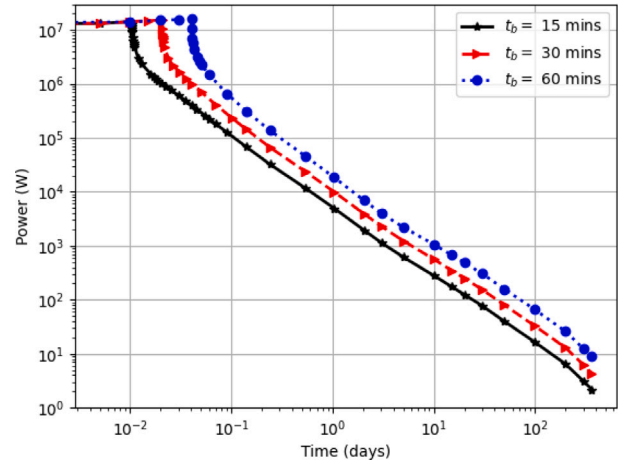


Fig. 21. Decay heat curve as a function of time from startup.

Table 3

Value of parameters in Eq. (44).

| k | t_{k-1} | t_k | α_k | η_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 |

6.1.2. Simple model for decay heat curves

Based on Bell and Glasstone (1970), Todreas and Kazimi express the approximate decay power as:

$$\frac{P(t_s)}{P_0} = 0.066 \left(t_s^{-0.2} - (t_s + t_b)^{-0.2} \right), \quad (43)$$

where $P(t_s)$ and P_0 denote the decay power and the power at shutdown in Watts, while t_s and t_b are the time after shutdown and the burn time at full power in seconds, respectively. As observed from Fig. 22, Eq. (43) 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 being considered. To overcome these limitations, the following generalized equation was utilized:

$$\frac{P(t_s)}{P_0} = \alpha_k \left(t_s^{-\eta_k} - (t_s + t_b)^{-\eta_k} \right) \quad \forall t_s \in I_k = [t_{k-1}, t_k], \quad (44)$$

where α_k and η_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 α_k and η_k are reported in Table 3 for the different time intervals.

Fig. 22 shows the predictions of Eqs. (43) and (44) for an operational time of 60 min. Eq. (44) 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 4, which reports the maximum relative error and the RMSRE. Eq. (44) 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.

6.2. Shutdown results

In this section, the following phases were considered:

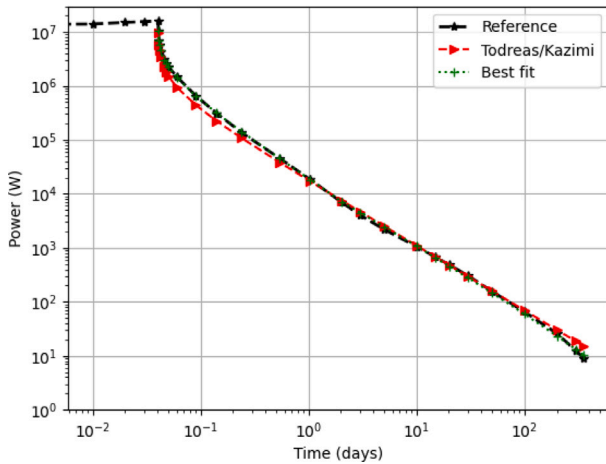


Fig. 22. Comparison of the reference Serpent model with the Todreas–Kazimi model (i.e., [Eq. (43)]), and the new best fit (i.e., [Eq. (44)]) for $t_b = 60$ min.

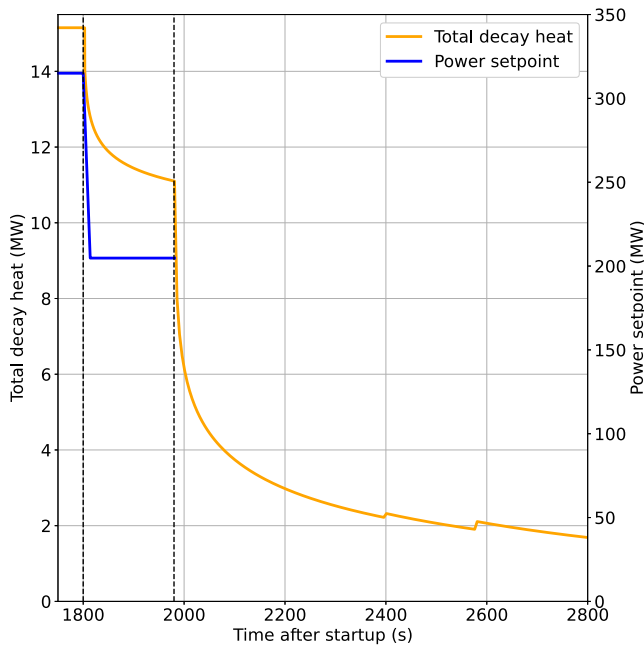


Fig. 23. Decay heat curve, assuming $t_b = 30$ min and a 3 min throttle at 65% thrust.

Table 4

The MARE, RMSRE, and maximum error for Eqs. (43) and (44), with respect to the reference Serpent solution, for a shutdown time of 365 days.

| Burn time | MARE, % | RMSRE, % | Max error, % |
|--------------------------|---------|----------|--------------|
| Best fit, Eq. (44) | | | |
| $t_b = 15$ 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, Eq. (43) | | | |
| $t_b = 15$ 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 |

1. Steady-state operation for $45 \leq t < 1800$ s
2. Throttling for $1800 \leq t < 1815$ s
3. Throttle hold for $1815 \leq t < 1980$ s
4. Temperature retreat $t \geq 1980$ s.

Specifically, the chamber pressure was chosen to be:

- $p_C(1800 \leq t < 1815 \text{ s}) = p_{C,\text{nom}} \left(1 - \frac{0.35(t - 1800)}{15}\right)$
- $p_C(1815 \leq t < 1981 \text{ s}) = 0.65 \times p_{C,\text{nom}}$
- $p_C(t \geq 1981 \text{ s}) = \frac{0.65}{(t - 1980)^{0.5}} \times p_{C,\text{nom}}$

with the total mass flow rate being determined per Eq. (5). The decrease in pressure during the temperature retreat was arbitrarily selected to vanish as $t \rightarrow \infty$, without decreasing too quickly (thus avoiding a significant temperature increase during that phase).

The power set point was adjusted as follows:

- $P_d(1800 \leq t < 1814 \text{ s}) = P_{\text{tot}} \left(1 - \frac{0.35(t - 1800)}{14}\right)$
- $P_d(t = 1814 \text{ s} \leq t \leq 1980 \text{ s}) = 0.65 \times P_{\text{tot}}$
- No set point for $t > 1980$ s.

As for the controller, a simple power-driven PID control, as described in Section 4.2.1, was found sufficient to obtain good results.

6.2.1. Decay heat model with throttling

The DHC was obtained by splitting the total power into two parts: 65 and 35% of the total power, respectively corresponding to the throttled power and the remainder—for which the decay heat precursors stop building up shortly after the beginning of the temperature retreat and throttling, respectively (i.e., around $t = 1980$ s and $t = 1800$ s, respectively). Eq. (44) was applied to each, along with a time delay. Fig. 23 summarizes the total decay heat power assumed as a function of time after startup, with the two vertical lines corresponding to the beginning of throttling and the temperature retreat. At $t = 2400$ s, a small discontinuity is observed that corresponds to the non-throttled part of the DHC as it transitions from $k = 1$ to $k = 2$ in the fit described in Table 3. Another small discontinuity appears when the decay of the throttled power performs that same transition at $t = 2580$ s. It is also noted that for the first few seconds, Fig. 23 constitutes an unphysical extrapolation that gives values larger than the Serpent decay heat at the end of the steady-state operation phase. Therefore, a constant value equal to $P(t_s = 0)/P_0 = 4.81\%$ (obtained from Serpent for $t_b = 30$ min) is used for the first 3 s after $t = 1800$ and 1980 s.

Spatially, the decay heat power density is assumed to have the same distribution as that of the total power density during steady-state operation. The maximum and RMS errors in the spatial decay distribution over the 25 depletion zones are found to be around 11 and 8%, respectively, when compared to the Serpent reference solution.

6.2.2. Steady-state operation

A challenging aspect of simulating steady-state operation lies in the time step selection: if the time steps are kept the same as at the end of the startup simulation ($\Delta t = 0.1$ s), the run time to reach steady-state conditions would be tremendous, as about 10,000 time steps would have to be run. Yet, the controller appears to be unstable if the time step is too large, so a limit of $\Delta t_{\text{max}} = 6$ s was chosen. Another difficulty, stemming solely from the current implementation of using the time step as the controller time delay, is that any change in time step implies a change in time delay. As such, a sudden change in Δt creates a (unphysical) variation of the predicted power if the delayed neutron precursors (DNP) have not yet reached saturation, which takes several minutes at constant power to achieve. This numerical defect could be overcome by incorporating a time-step-independent time delay in the controller numerical model. However, since the solution during the steady-state operation phase is largely known (the power and temperature fields are virtually constant, and only the CD angle adjusts to account for the gradual saturation of the DNPs), the main interest lies in the final solution, at the beginning of the throttling phase. Therefore, the change in controller modeling is left for future work, and the procedure to obtain the final solution is as follows:

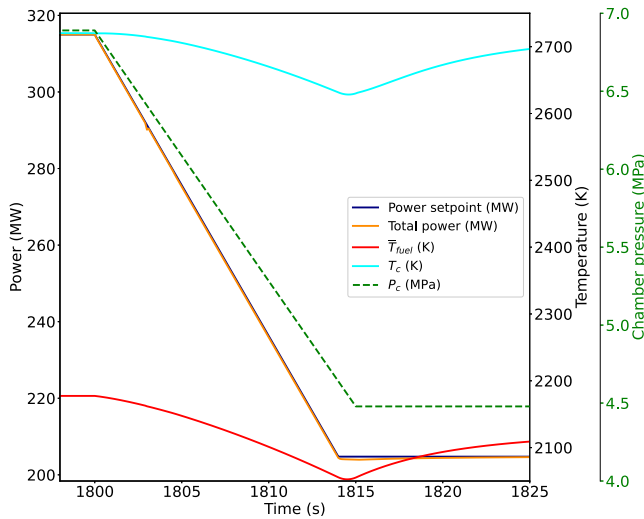


Fig. 24. Evolution of power, temperature, and chamber pressure during the throttling and throttle hold phases.

- Run the startup phase up to $t = 75$ s.
- Instantaneously increase Δt from 0.1 to 6 s and decouple the neutronics from the thermal and thermal-hydraulics model until the unphysical oscillations vanish (to avoid propagating the oscillations to the thermal fields).
- Reactivate the full coupling (in practice, after an arbitrary 80 time steps) and continue the simulation until all fields appear saturated (in practice, until $t = 1075$ s).

6.2.3. Throttling and throttle hold

Throttling is simulated by reducing the mass flow rate and chamber pressure to 65% of their nominal values within 15 s, with $\Delta t = 0.1$ s. The power set point is also reduced by the same amount, but in slightly less time (14 s). The power PID controller can successfully follow the demand, as seen in Fig. 24. The coolant outlet temperature decreases over the first 14.6 s and reaches almost 92 K below its nominal value, indicating a decrease in specific impulse (although throttling should theoretically maintain it). A better behavior could be achieved by adjusting the power and chamber pressure demands. However, this need to manually tune the demands in order to maintain the specific impulse further evidences the limitations of using power-following controllers as opposed to directly controlling the chamber temperature. Once the throttle hold is initiated and both the power and mass flow rates are constant again, the outlet temperature returns to its nominal value.

To speed up the calculation during the throttle hold, which was chosen to last for 165 s, it is tempting to increase the time step. However, the numerical artifact of the controller detailed in Section 6.2.2 once again affects the predicted solution. In practice, it is chosen to start increasing from $\Delta t = 0.1$ s after 10 s of throttle hold ($t = 1825$ s), to $\Delta t = 1$ s, 10 seconds later ($t = 1835$ s).

6.2.4. Temperature retreat

For the temperature retreat, the controller is no longer used. Instead, the drums are rotated at a fixed speed of 5 deg/s until they reach their final position: $\theta = 180$. The large negative reactivity insertion rapidly decreases the total power, as seen in Fig. 25, with the decrease rate being largely driven by the DNP decay constants within the first few minutes, and by decay heat thereafter (see Fig. 22). The sudden change in mass flow rate 1 second after the beginning of the temperature retreat phase also leads to a sharp decrease in fuel, moderator, and chamber temperatures, but cooling remains necessary for some time due to decay heat. Precise cooling requirements and strategies will be the focus of future work.

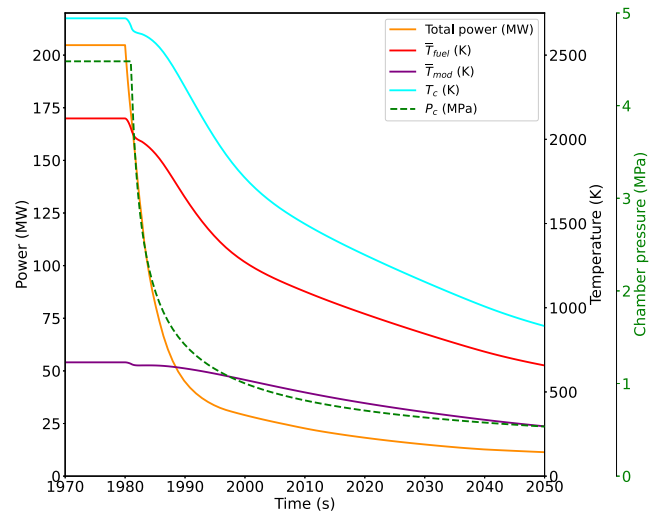


Fig. 25. Evolution of power, temperature, and chamber pressure during the temperature retreat.

6.2.5. Summary

A summary of all the startup and shutdown phases simulated in this work is given in Fig. 26. Overall, the hybrid PID and period-generated controllers achieve great results in following the power demand during startup phases. During shutdown phases, a simple PID using only a power error signal is relied upon, producing similarly good results. Significant chamber temperature overshoot and slightly varying specific impulse during throttling suggest that this work be extended to temperature-following controllers.

6.3. Computational cost

A summary of the runtimes required for the simulations described in this work is presented in Table 5, with the exponential ramp-up, bootstrap and thrust buildup, steady-state operations, throttling and throttle hold, and temperature retreat referring to the simulations from Section 4, Section 5, Section 6.2.2, Section 6.2.3, and Section 6.2.4, respectively. All the simulations were run on the Sawtooth supercomputer at Idaho National Laboratory (Black, 2019), 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 additional 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).

7. Summary and future work

This work presented a multiphysics Griffin/Bison/RELAP-7 numerical model of a prototypical NTP system designed to simulate rapid startup and shutdown transients. The physics modeled included full-core neutronics, assembly-wise heat conduction, and conjugate heat transfer, with the balance of plant mainly being 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 the assumed initial conditions to nominal power within a fraction of a minute, then shutting down after a few minutes of maintaining a 65% throttled thrust.

Another layer of multiphysics coupling was added to the model in order to automatically control the drums by using a power demand and

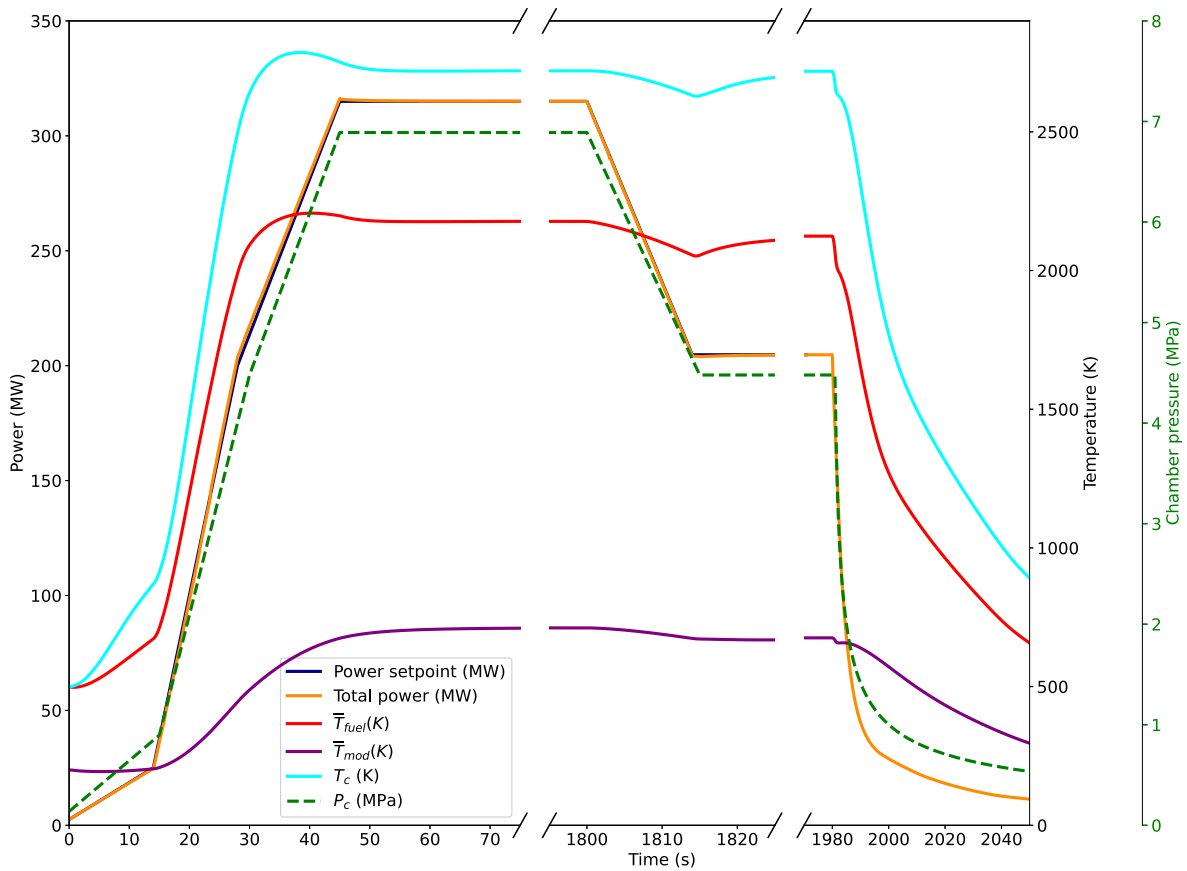


Fig. 26. Summary of the startup and shutdown phases.

Table 5

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

| Transient | Simulation start time (s) | Simulation end time (s) | Number of time steps | Transient runtime (hrs) |
|----------------------------|---------------------------|-------------------------|----------------------|-------------------------|
| Exponential ramp-up | 0 | 49 | 493 | 24 |
| Bootstrap + thrust buildup | 0 | 75 | 809 | 30.4 |
| Steady-state operations | 75 | 1075 | 167 | 4.2 |
| Throttling + throttle hold | 1800 | 1980 | 421 | 16.5 |
| Temperature retreat | 1980 | 2050 | 86 | 5.5 |

to attempt to limit the 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 proved 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 knowledge of reactivity feedback coefficients and change rates 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 degree of 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 buildup phases that occur prior to reaching steady-state conditions, thus demonstrating the ability of the hybrid PID and period-generated controllers to handle such transients, with both types of controllers showing very similar behavior. Nevertheless, a non-negligible chamber temperature overshoot (around 3%) was observed, caused by the demanded power signal and assumed chamber pressure.

Finally, a study was performed to model decay heat during shutdown. Throttling, throttle hold, and temperature retreat phases were simulated. Once again, the power set point was closely followed by the controller, but the chamber temperature, and thus the specific impulse, was not maintained at a perfectly constant value. Both this issue and the temperature overshoot during startup could be alleviated by considering temperature-following controllers instead of power-following ones (as was the focus of this work) in order to avoid manually adjusting the power demand and mass flow rates so as to achieve the desired behavior.

Specifically, the following items are considered to represent future work:

- Implement reactor control that follows chamber temperature and pressure signals and is based on both control valves and drums (rather than using a power-following signal based solely on drums to control reactivity).
- Add the cooldown phases (with continuous and pulsed cooling) to the simulation in order to determine post-shutdown cooling requirements.
- Gain a better understanding of prototypical initial conditions (e.g., after initial temperature conditioning) and adjust the model accordingly.

- Enhance the system model so that it does 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.
- Account for the power density spatial distribution within each fuel assembly (through power reconstruction or a heterogeneous transport neutronics model).
- Include the effect of poisoning from fission products to be able to study the ability of the reactor to perform restart within a few days of shutting down.

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

CRedit authorship contribution statement

Vincent M. Labouré: Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing, Review & editing, Visualization, Supervision. **Sebastian Schunert:** Conceptualization, Methodology, Software, Formal analysis, Review & editing, Supervision, Project administration. **Stefano Terlizzi:** Conceptualization, Methodology, Formal analysis, Investigation, Writing, Review & editing, Visualization. **Zachary M. Prince:** Methodology, Formal analysis, Investigation, Review & editing. **Javier Ortensi:** Review & editing, Supervision. **Ching-Sheng Lin:** Methodology. **Lise M. Charlot:** Software, Review & editing. **Mark D. DeHart:** Formal analysis, Review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors are unable or have chosen not to specify which data has been used.

Acknowledgments

This manuscript was 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.

The authors would like to recognize John Shaver for his excellent and prompt technical review of this manuscript.

References

- Altseimer, J.H., Mader, G.F., Stewart, J.J., 1971. Operating characteristics and requirements for the NERVA flight engine. *J. Spacecr. Rockets* 8 (7), 766–773. <http://dx.doi.org/10.2514/3.59723>.
- Baxter, A., Ellis, C., Fikani, M., Hanson, D., Saurwein, J., Charman, C., 2009. Final Report – NGNP Core Performance Analysis, Phase 1. Technical Report 911160, General Atomics.
- Bell, G.I., Glasstone, S., 1970. Nuclear Reactor Theory. URL <https://www.osti.gov/biblio/4074688>.
- Bernard, J.A., Lanning, D.D., 1992. Considerations in the design and implementation of control laws for the digital operation of research reactors. *Nucl. Sci. Eng.* 110, 425–444.
- Black, D., 2019. Sawtooth supercomputer from HPE comes to idaho national lab. InsideHPC URL <https://insidehpc.com/2019/12/sawtooth-supercomputer-from-hpe-comes-to-idaho-national-lab/>.
- Ciftcioglu, Ö., Geckinli, M., 1980. A CAMAC-based reactivity-meter for nuclear reactors. *Nucl. Instrum. Methods* 177 (2), 321–326. [http://dx.doi.org/10.1016/0029-554X\(80\)90040-3](http://dx.doi.org/10.1016/0029-554X(80)90040-3), URL <https://www.sciencedirect.com/science/article/pii/0029554X80900403>.
- DeHart, M.D., Schunert, S., Labouré, V.M., 2022. Nuclear thermal propulsion. In: Pope, C.L. (Ed.), *Nuclear Reactors*. IntechOpen, Rijeka, <http://dx.doi.org/10.5772/intechopen.103895>, URL <https://doi.org/10.5772/intechopen.103895>.
- Gustafson, J.L., 2021. Space nuclear propulsion fuel and moderator development plan conceptual testing reference design. *Nucl. Technol.* 207 (6), 882–884. <http://dx.doi.org/10.1080/00295450.2021.1890991>.
- Jung, Y., Lee, C., 2018. PROTEUS-MOC User Manual. Technical Report ANL/NE-18/10, Argonne National Laboratory.
- Klein, A.C., Camp, A., McClure, P., Voss, S., January 2021. Operational Considerations for Fission Reactors Utilized on Nuclear Thermal Propulsion Missions to Mars. Technical Report NASA/CR-20210000387.
- Krecicki, M., Kotlyar, D., 2022. Full-core coupled neutronic, thermal-hydraulic, and thermo-mechanical analysis of low-enriched uranium nuclear thermal propulsion reactors. *Energies* 15 (19), <http://dx.doi.org/10.3390/en15197007>, URL <https://www.mdpi.com/1996-1073/15/19/7007>.
- Labouré, V.M., Schunert, S., Terlizzi, S., Prince, Z.M., Ortensi, J., Lin, C.-S., Charlot, L.C.M., DeHart, M.D., 2022. Automated Control for Nuclear Thermal Propulsion Start-Up using MOOSE-based Applications. Technical Report INL/RPT-22-68551, <http://dx.doi.org/10.2172/1893100>, URL <https://www.osti.gov/biblio/1893100>.
- Labouré, V.M., Wang, Y., Ortensi, J., Schunert, S., Gleicher, F.N., DeHart, M.D., Martineau, R.C., 2019. Hybrid super homogenization and discontinuity factor method for continuous finite element diffusion. *Ann. Nucl. Energy* 128 (INL/JOU-18-5106), 443–454. <http://dx.doi.org/10.1016/j.anucene.2019.01.003>.
- Leppänen, J., Pusa, M., Viitanen, T., Valtavirta, V., Kalliaiseno, T., 2015. The serpent Monte Carlo code: Status, development and applications in 2013. *Ann. Nucl. Energy* 82, 142–150.
- Manickam, V., Krecicki, M., Kotlyar, D., 2022. 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.
- McCarthy, J.R., Wolf, H., 1960. The Heat Transfer Characteristics of Gaseous Hydrogen and Helium. Technical Report RR-60-12, Rocketdyne.
- NASA, 2018. Nuclear Thermal Propulsion: Game Changing Technology for Deep Space Exploration. National Aeronautics and Space Administration, URL https://www.nasa.gov/directorates/spacetech/game_changing_development/Nuclear_Thermal_Propulsion_Deep_Space_Exploration.
- National Academies of Sciences, Engineering, and Medicine, 2021. Space Nuclear Propulsion for Human Mars Exploration. The National Academies Press, <http://dx.doi.org/10.17226/25977>.
- Ortensi, J., Labouré, V.M., Schunert, S., Stewart, R.H., 2020. Neutronics Enhanced Meshing Operations (NEMO) Users Manual Version 1.0. Technical Report INL/MIS-20-00730, Idaho National Laboratory.
- Ortensi, J., Wang, Y., Laurier, A., Schunert, S., Hébert, A., DeHart, M.D., 2018. A Newton solution for the superhomogenization method: The PJFNK-SPH. *Ann. Nucl. Energy* 111, 579–594.
- Permann, C.J., Gaston, D.R., Andrs, D., Carlsen, R.W., Kong, F., Lindsay, A.D., Miller, J.M., Peterson, J.W., Slaughter, A.E., Stogner, R.H., Martineau, R.C., 2020. MOOSE: Enabling massively parallel multiphysics simulation. *SoftwareX* 11, 100430.
- Ruark, A.E., 1947. NNuclear Powered Flight. Technical Report APL/JEU-TG-20, Applied Physics Laboratory, Silver Spring, Maryland.
- Schunert, S., Wang, Y., Ortensi, J., Labouré, V.M., Gleicher, F.N., DeHart, M.D., Martineau, R.C., 2019. Control rod treatment for FEM based radiation transport methods. *Ann. Nucl. Energy* 127, 293–302. <http://dx.doi.org/10.1016/j.anucene.2018.11.054>.
- Space Nuclear Propulsion Technologies Committee, 2021. Space Nuclear Propulsion for Human Mars Exploration. Technical Report, National Academies of Sciences, Engineering, and Medicine.

- Spence, R.W., 1968. The rover nuclear rocket program. *Science* 160 (3831), 953–959.
- Stewart, M.E., 2015. A historical review of cermet fuel development and the engine performance implications. In: *Nuclear and Emerging Technologies for Space (NETS) 2015*. American Nuclear Society, Albuquerque, New Mexico, URL <https://ntrs.nasa.gov/api/citations/20150002852/downloads/20150002852.pdf>.
- Todreas, N.E., Kazimi, M.S., 1990. *Nuclear Systems*, Vol. 1, second ed. CRC Press.
- Wang, Y., Schunert, S., Ortensi, J., Labouré, V.M., DeHart, M.D., Prince, Z.M., Kong, F., Harter, J.R., Balestra, P., Gleicher, F.N., 2021. Rattlesnake: A MOOSE-based multiphysics multischeme radiation transport application. *Nucl. Technol.* 207 (7), 1047–1072. <http://dx.doi.org/10.1080/00295450.2020.1843348>.