ORNL/TM-2025/4087

Microreactor Automated Control System: Digital Twin Models and Advanced Control Systems

Pradeep Ramuhalli
Jono McConnell
Wesley Williams
Aiden Meek

August 2025

OAK RIDGE National Laboratory

ORNL IS MANAGED BY UT-BATTELLE LLC FOR THE US DEPARTMENT OF ENERGY

ORNL/TM-2025/4087

Nuclear Energy and Fuel Cycle Division

MICROREACTOR AUTOMATED CONTROL SYSTEM: DIGITAL TWIN MODELS AND ADVANCED CONTROL SYSTEMS

Pradeep Ramuhalli
Jono McConnell
Wesley Williams
Aiden Meek

August 2025

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-BATTELLE LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

LIST OF FIGURES

LIST OF ABBREVIATIONS

BOP	balance of plant
FMU	Functional Mockup Unit
HIL	hardware-in-the-loop
HX	heat exchanger
INL	Idaho National Laboratory
MACS	microreactor automated control system
MAGNET	Microreactor Agile Non-Nuclear Experimental Test Bed
MPC	model predictive control
ORNL	Oak Ridge National Laboratory
PID	proportional-integral-derivative
TRANSFORM	TRANsient Simulation Framework of Reconfigurable Models

ABSTRACT

Automation of control systems is expected to be important in the economic and safe operation of microreactors. Therefore, there is a need to develop and demonstrate automated control for microreactors, along with the development of testbeds for this purpose. This report provides updates on the status of a nonnuclear microreactor automated control system (MACS)—a real-time, hardware-in-the-loop testbed for nonnuclear testing of microreactor control system automation. A real-time hardware-in-the-loop testbed incorporates the realistic dynamics of physical systems into control system development and testing. The collaborative effort between Oak Ridge National Laboratory (ORNL) and Idaho National Laboratory (INL) resulted in the development of a prototypic microreactor plant-level digital twin that includes the reactor and a balance of plant system. Advanced control strategies were incorporated to demonstrate testing of control automation solutions. The gRPC communication protocol, which was implemented in the hardware-in-the-loop testbed by INL, was coupled to a digital twin model developed using the TRANsient Simulation Framework of Reconfigurable Models (TRANSFORM) library in Modelica. This digital twin simulation was tested with the ViBRANT hardware for realistic feedback and visual representation of control action in real time. A modular Python client structure was developed to manage functional mock-up unit–based simulation and real-time gRPC communication. Hardware-in-the-loop testing indicated that the modeled reactor—a natural-convection, molten-salt coolant loop configuration—responds well to control of drum positioning for modulation of reactor core power, as well as system-level control and downstream demand changes. Ongoing research is focused on integrating additional control algorithms that utilize data from newly included sensors within the MACS hardware testbed, as well as demonstrating and assessing the performance of the different automated control algorithms on multiple additional operational scenarios.

1. INTRODUCTION

Microreactors, with their design for mobility and flexible operations [1, 2, 3], are being considered for meeting the energy needs in grid-isolated communities. The lower design power output (typically less than ~100 MWth) leads to a loss of economies of scale. This factor, along with the need for remote operation while maintaining high-capacity factors on demand [4], has led to the need for development and testing of technologies that support the automation of operations [5, 6, 7]. Higher levels of automation in control and operation of microreactors can potentially enable faster decision-making for operational modes that may differ from those used in the current large light-water reactor fleet. Such a capability is especially useful for rapid response to abnormal conditions (such as failure of one or more components) to maintain the plant in a safe condition.

In collaboration with Idaho National Laboratory (INL), Oak Ridge National Laboratory (ORNL) has been supporting the development of a high-fidelity and robust microreactor automated control system (MACS) that can act with minimal need for human-in-the-loop action. This ongoing research effort at ORNL aims to integrate digital twin models with hardware to enable the development and testing of graded automation for microreactors using a hardware-in-the-loop (HIL) testbed. The MACS framework utilizes expected reactor inputs and outputs, such as reactor temperature, control element (i.e., drum or rod) position, coolant temperature and energy transfer to heat sink, and factors, such as reactivity feedback.

Prior work on MACS identified several requirements, showed the potential for automating the control logic using robust plant-level modeling of the microreactor, and integrated the control strategy with the HIL testbed. MACS proposed the separation of the basic reactor control and protection functions from a higher-level coordination layer that can incorporate the automation functions. The resulting control automation function can then be developed as a non–safety-related system, as the separation imposes requirements for

MACS to (1) not perform any safety-related functions, (2) not interfere with the function of any safety-related system, and (3) not override any operator commands. Therefore, MACS can be classified as a Level 2, Level 3, or Level 4 (depending on task) automation system [8, 9], where Level 4 rerpresents a fully autonomous system.

The focus of the present research is on the development and demonstration of MACS when integrated with an HIL simulator. Such integration is expected to result in a testbed capability for developing and evaluating advanced control algorithms by various stakeholders, thereby benefiting the microreactor community and accelerating the deployment of microreactors. The research reported here has focused on updating the previously developed models [6, 10, 11] with a prototypic balance of plant (BOP) model, identifying data or information needs for BOP integration with the MACS digital twins, and integrating a model predictive control (MPC) algorithm with the HIL software platform to demonstrate integrated advanced controls. This work leverages prior research outcomes on automated control systems and autonomous supervisory control systems [12, 13, 14, 15, 16, 17], existing concepts for microreactors [2], and available testbeds [18].

1.1 OBJECTIVES OF THIS REPORT

This report documents the status of research on MACS and describes recent results from the updates to the MACS digital twin models and control algorithms. This work leverages prior research outcomes on automated control systems and autonomous supervisory control systems, existing concepts for microreactors, and available testbeds.

In particular, prior work on MACS identified challenges associated with robust plant-level modeling of the microreactor and the integration of the control strategy with the HIL testbed. The work described in this document addresses many of these challenges, leveraging surrogate models of BOP systems to build a complete plant-level digital twin, and demonstrates robust integration of surrogate models with control automation and the HIL testbed. This document describes the updated digital twins (i.e., surrogate models), control automation formulation, and integration results, as well as ongoing research to enhance the capabilities of the MACS testbed.

1.2 OUTLINE OF THIS REPORT

Section 2 briefly describes previous research and provides an overview of the MACS framework. Section 3 describes the updated plant-level digital twin models and the control automation formulation. Section 4 also reports the results from integrating the control system automation software modules with the digital twin models and those from testing using the HIL platform. A summary of the findings and ongoing research is included in Section 5.

2. BACKGROUND

2.1 PRIOR WORK ON MACS SOFTWARE CLIENTS

The work presented in this document builds upon the success of initial development for a more isolated and simplified problem focus [6, 10, 11]. A microreactor model representative of the MARVEL core and primary coolant loop [2] was developed using ORNL's TRANsient Simulation Framework of Reconfigurable Models (TRANSFORM) library [19, 20]. The model couples two major types of components. The reactor core is modeled using a 6-group point kinetics system of differential equations. Temperature feedback is based on the effective temperature of the NaK coolant in the section of the loop passing through the core. The NaK coolant is represented by multiple pipe sections. These pipes are used to geometrically represent the flow path of NaK in the MARVEL design [2]. The pipe modules provide a multi-physics thermodynamic representation of the viscous flow, considering dynamic calculation of fluid pressure, buoyancy effects in vertical components, and heat transfer to and from the fluid volume. The pipe section closest to the core is simulated with a four-section discretization for mass flow calculations. A distribution of the heat generated by the core is transferred to the nearby coolant pipe, creating a buoyancy difference among the discretized masses and naturally driving convection. Similarly, heat is drawn from the opposing end of the NaK coolant loop to represent the draw of downstream systems. Visualization of the TRANSFORM model is provided in Figure 1.

The truncated model was used to generate a Functional Mockup Unit (FMU) [21], which allows the user to interact with the simulation through programming external to a Modelica compiling environment. This FMU was leveraged to develop a Python-based, modular digital architecture [6, 11] to perform real-time, HIL simulation in concert with the MACS hardware being developed at INL [16, 17, 18]. A test was developed to observe the behavior of the hardware system coupled with the TRANSFORM simulation by moving the physical drums to the demand of a proportional-integral-derivative (PID) controller. Drum position was converted to a corresponding reactivity insertion within the reactor core simulation. The controller was designed to drive the core power to a requested ramp transient. The results of this effort are provided in Figure 2. The data show the performance of the physical drum following of four drums to meet the demanded angle of a PID controller (Fig. 2a) as well as the resulting heat load following of core power in response to controller action (Fig. 2b). Using the HIL setup, the tuned controller showed reasonable power following capability, even in the presence of the realistic communication and hardware reaction delays in the simulation loop. These results show similar acceptable performance at 0.3 s control loop time cycles (0–100 s) as well as 0.5 s time cycles (100–200 s). A similar set of data is displayed for a sinusoidal requested power load (Fig. 3) for the same time step selections. For each time step, amplitudes of 10 kW and 20 kW were explored at frequencies of 1/10 Hz and 1/30 Hz. It was observed that PID control struggled in the high frequency trends, but was able to achieve reasonable power tracking for the 1/30 Hz sinusoidals with some response delay.

2.2 CHALLENGES AND MOTIVATION FOR PRESENT WORK

As discussed above, the prior work focused on the use of simple reactor-only models and testing using baseline control strategies. Tests with the HIL testbed identified communication delays as one of the challenges of real-time automated (or autonomous) controls and operations; additional research is needed in this area. A more complete plant-level model that includes the secondary-side BOP systems was also identified as a key need if the goal is to include control automation and operator decision support through the inclusion of supervisory control systems. Such a model also needs to be fast to allow real-time decision-making for control and operations, and it should ideally include spatial variability of reactor power to enable greater realism for control decision-making. Techniques for evaluating autonomy in the HIL platform are also needed.

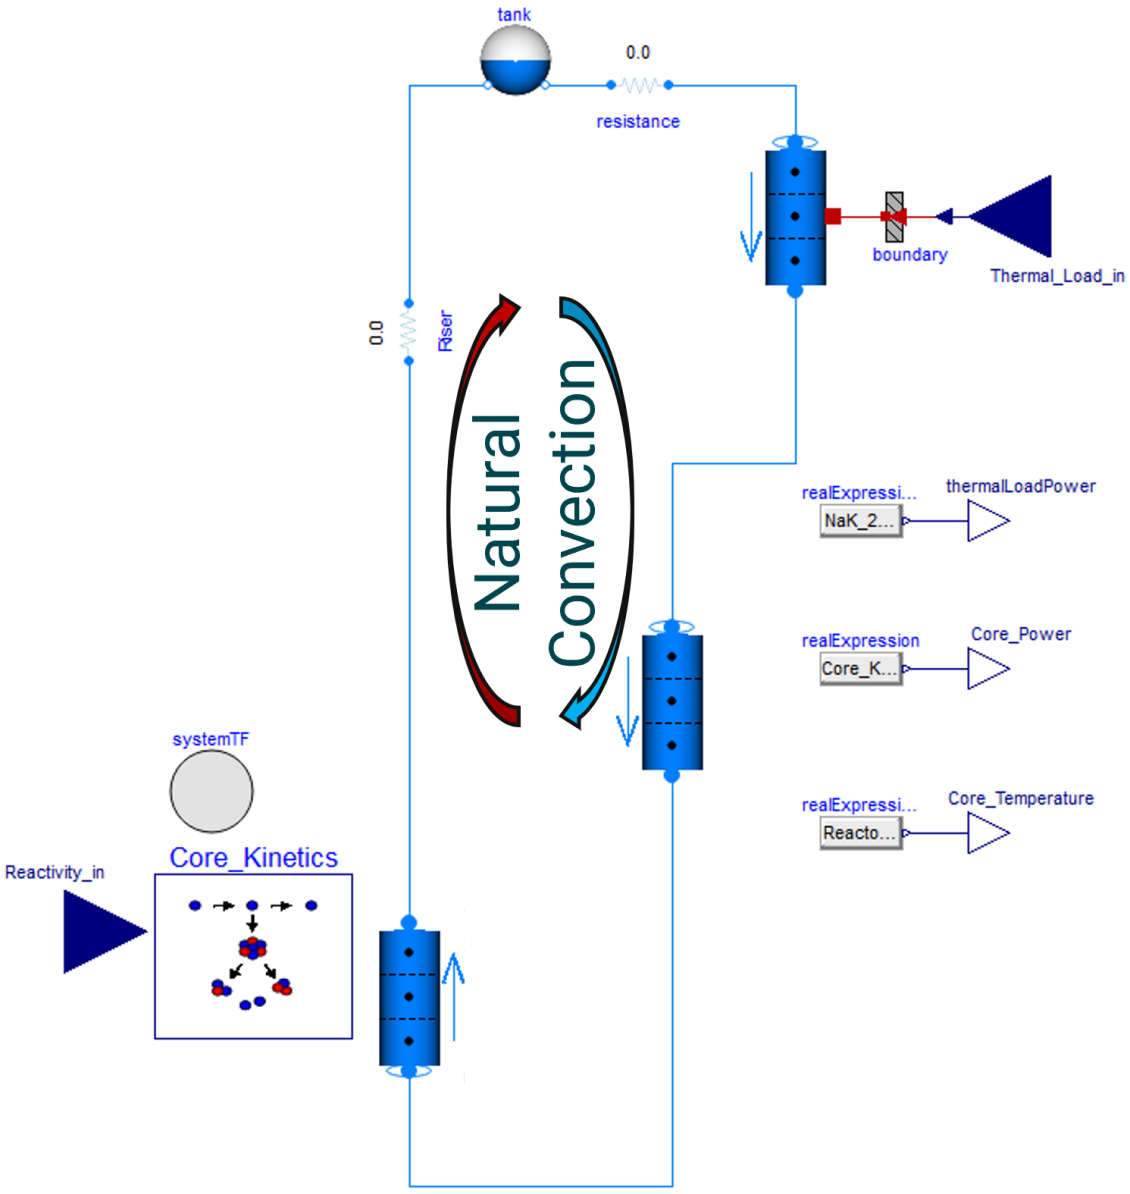

Figure 1. A TRANSFORM model of core behavior coupled with a NaK primary coolant loop under natural convection.

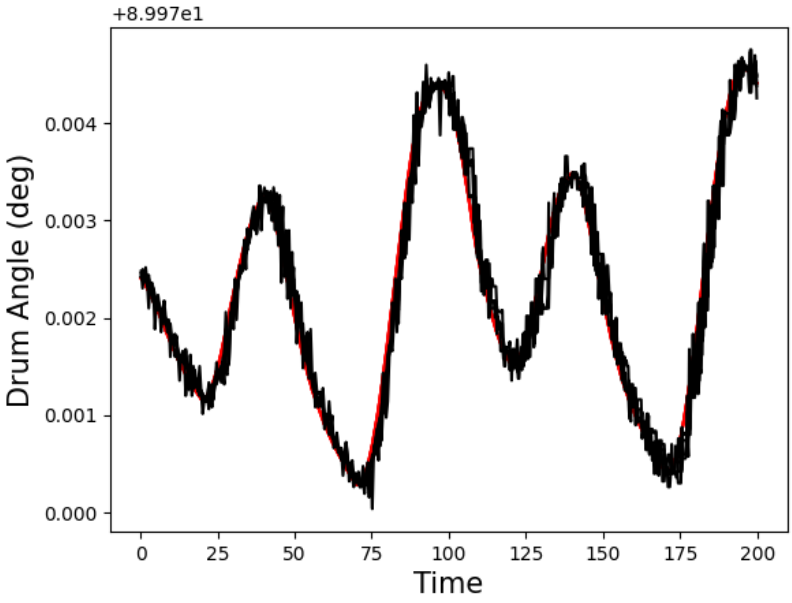

(a)

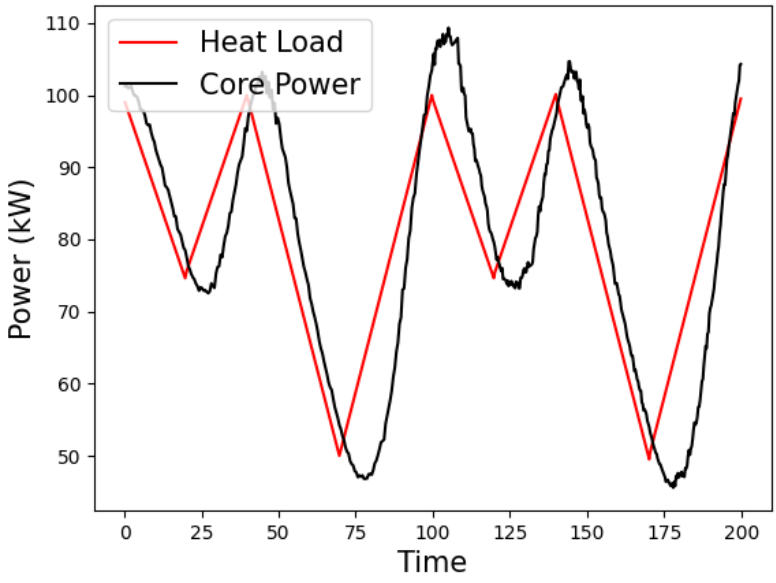

(b)

Figure 2. Hardware-in-the-loop demonstration performance. (a) Drum angle response to control setpoint trend of all four drums superimposed. (b) Power response to ramp changes in desired heat load.

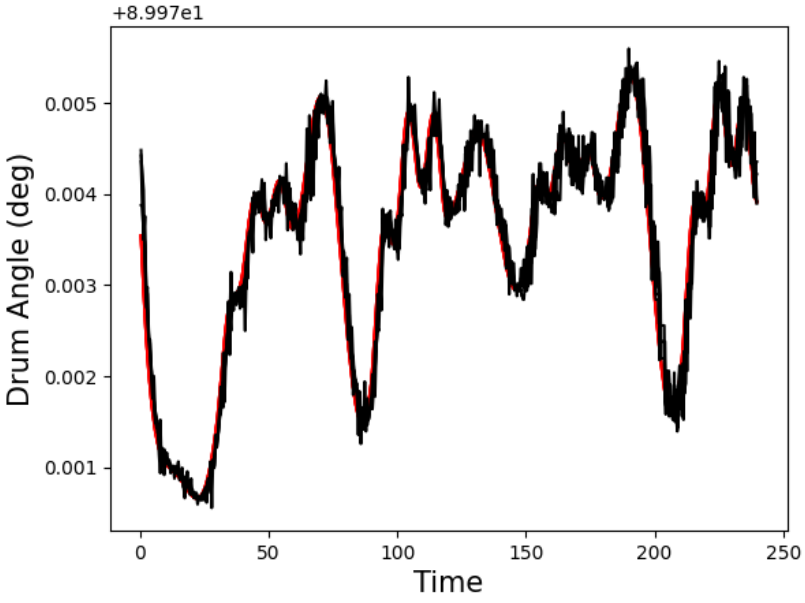

(a)

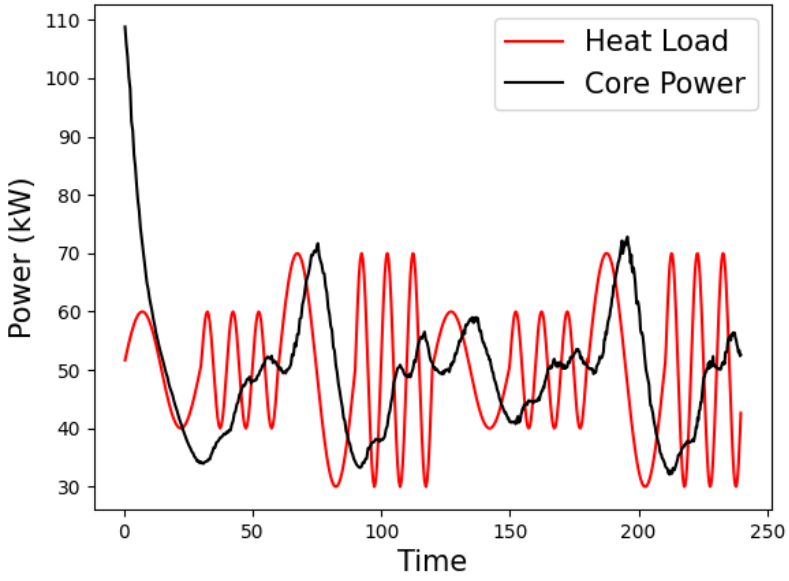

(b)

Figure 3. Hardware-in-the-loop demonstration performance. (a) Drum angle response to control setpoint trend of all four drums superimposed. (b) Power response to sinusoidal changes in desired heat load.

This report describes work conducted to address some of these needs, focusing on system-level models that include spatial variability of reactor power and BOP systems. Work to integrate advanced control algorithms with faster models and integrate these solutions with the HIL software platform is also described.

3. MICROREACTOR AUTOMATED CONTROL SYSTEMS: SURROGATE MODELS

The performance of initial model testing in Figures 2 and 3 provides reasonable proof of concept for real-time, HIL control of a nonnuclear testbed. However, three main areas were identified for further improvement. First, the previous model considers only a single-point kinetics representation of reactor core dynamics. Larger multi-point and distributed models can push the limits of real-time simulation without the integration of high-powered computing. However, it is necessary to approximate the spatial distribution of the core to take advantage of the distributed LED arrays present on the MACS and ViBRANT hardware [16, 17, 18, 22] to improve accuracy of the control solution. Second, expanding the model to represent a more system-level view of nuclear power plant dynamics will provide a more realistic control scenario with multiple inputs, such as reactivity insertion and downstream valves for power, mass flow, and temperature control.

3.1 SPATIAL POWER DENSITY DISTRIBUTION

The task of expanding the model to provide a spatial representation of power output within the core is nontrivial and requires a balance of physical realism with low computational burden to remain a real-time-capable system for HIL testing. For this reason, more accurate but more computationally intensive methods, such as multi-point kinetics or a fully distributed model, such as MCNP, were set aside in favor of a power-shaping augmentation of the current point kinetics dynamic model. This ensures that speed can be maintained; however, to remain flexible for future efforts that explore the effects of fault conditions and reactor tilt on control strategies, it is imperative that the power shaping be asymmetric and account for the individual effects of each control drum position. With this in mind, a power shaping model has been developed in the form of Eqs. (1–5).

$$\rho(r,z,\theta) = \bar{A}P\,\rho_{radial}(r)\,\rho_{axial}(z)\,\rho_{angular}(\theta) \tag{1}$$

$$\rho_{radial}(r) = \cos(\frac{(1-a)r\pi}{2R}) \tag{2}$$

$$\rho_{axial}(z) = \cos(\pi(b+(1-2b)\frac{z}{H}) - \frac{\pi}{2}) \tag{3}$$

$$\rho_{angular}(\theta) = \sum_{i=0}^{N-1} \omega_i(\phi_i)\cos^2(\frac{\theta}{2} + \frac{i\pi}{N}) \tag{4}$$

$$\bar{A} = \frac{\pi(1-a)(1-2b)}{4RH\,sin(\frac{(1-a)\pi}{2R})\;sin(\frac{\pi}{2}-\pi b)\;\sum_{i=0}^{N-1}\omega_i(\phi_i)} \tag{5}$$

In the above equations, ρ represents the volumetric power density in $\frac{W}{m^3}$. The symbols r, z, and θ represent the radial, axial, and angular coordinates within the cylindrical volume of the reactor core, respectively. The symbol P is the total power as defined by point kinetics. The functions ρ_{radial}, ρ_{axial}, and $\rho_{angular}$ are the radial, axial, and angular shaping applied to the total power, respectively. The symbols a and b are the tunable truncation factors associated with radial and axial cosine truncation, respectively, and $\omega_i(\phi_i)$ is the weighting of the individual contribution of each drum position toward tilt. Factors a, b, and ω_i can be fitted

using slower, higher-order simulation to improve the accuracy of the power shaping without an increase in computational burden during real-time simulation. $\bar{A}$ is the correction factor applied to normalize the volumetric power density such that the volumetric integral of the power density is equal to the original total power. Visualization of the functions ρ_{radial}, ρ_{axial}, and $\rho_{angular}$ in isolation as well as a derivation of the correction factor $\bar{A}$ are provided in Appendix A.

A visual depiction of this developed distribution is provided as 2D slices in Figure 4, and a four-angle 3D view is given in Figure 5. As development progresses, this shaping will be fitted to the behavior of a reactor core from higher-order simulation to tune drum position weighting and cosine truncation.

3.2 BALANCE OF PLANT MODELING

The previous iteration of reactor modeling considered power generation from nuclear kinetics, which provided heat to a primary coolant loop that would flow via natural convection. A prescribed heat removal from the primary coolant was implemented to emulate the draw of downstream systems. This modeled system is a truncated representation of the MARVEL design [2] and contains one main control point: reactivity insertion through control drum rotation. To approach a more realistic representation of a reactor power plant, the model must include downstream systems, such as secondary coolant loops with pumps and valves for mass flow control as well as power extraction, to complete the energy balance. Although the MARVEL design includes a secondary coolant loop and sterling engine power extraction, it was decided that a model of the Microreactor Agile Non-Nuclear Experimental Test Bed (MAGNET) facility [23] as the secondary system would provide more value to the modeling effort because it allows for the leveraging of previous model development efforts in the HYBRID library [24]. This selection also creates opportunity for future HIL and digital twin collaboration between the MACS and MAGNET testbeds.

The MAGNET model within the HYBRID library provides a coolant loop of N_2 fluid flow driven by a pump to absorb energy from a heat source. The N_2 is pumped through a preheater heat exchanger (HX) and then passes a heat source, where heat is inserted into the system in a manner similar to that employed for the reactor core–NaK primary coolant interaction. This hot N_2 returns in counter-flow through the preheater before flowing to a second counter-flow HX. This secondary HX transfers heat between the N_2 flow and a cold water flow for system heat extraction.

To augment these models for coupled simulation, a few small changes were made. The prescribed heat extraction from the NaK coolant loop was removed, and the coolant flow was split at the point of heat extraction to create a flow out of the model and a subsequent flow into the model. Similarly, for the MAGNET model, the heat source was removed, and the N_2 flow was split to provide both an outlet and a return flow back into the model. These outflow and inflow ports are connected to a centralized counter-flow HX, which serves to transfer heat dynamically between the two systems. A visual of this model construction is provided in Figure 6.

Initial testing of the coupled system suggested an imbalance between the power levels generated by the reactor core and the heat removal by the MAGNET model tertiary coolant flow of water. As future intentions are to consider more dynamic power extraction systems, such as a Brayton cycle, this tertiary coolant flow was replaced by a dynamic prescribed heat removal, as shown in Figure 7. Initial testing also revealed a delicate balance of energy production and removal that made the selection of steady-state power extraction a nontrivial task. With this in mind, an independent PID controller has been implemented to manage system heat extraction toward maintaining nominal temperatures in the core. Future efforts with MPC will address the temperature control in a more coupled manner.

To observe the performance of the full coupled dynamics system simulation, a ramp power following test was performed within the Dymola/Modelica environment. The system was allowed to settle to steady temperatures at a controlled power of 85 kW. The temperature control PID was used to drive core temperature to 500°C while the rest of the system naturally reached its corresponding steady-state temperatures.

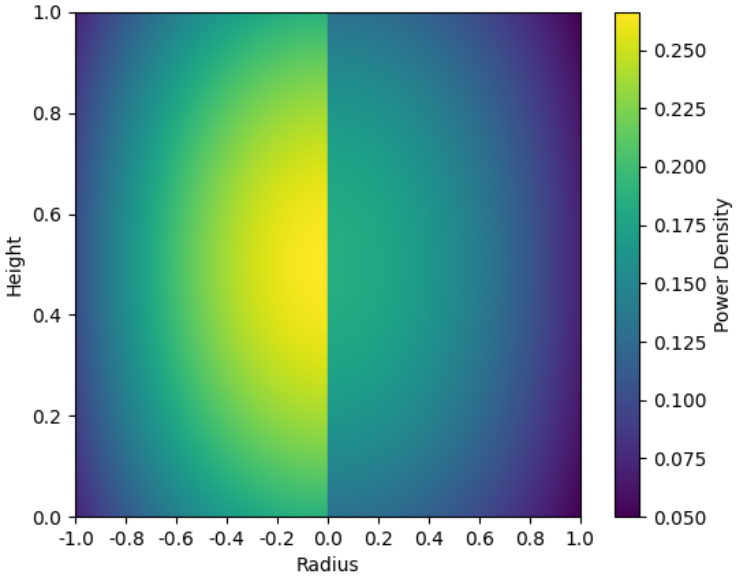

(a)

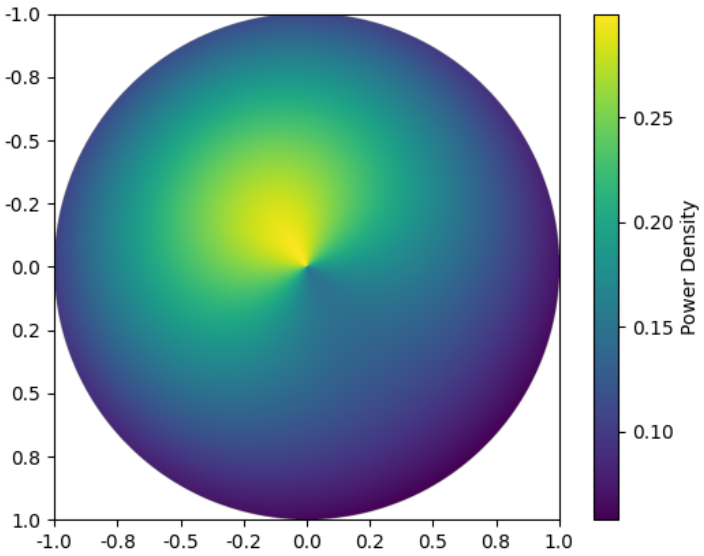

(b)

Figure 4. Slices of 3D power density distribution. Power asymmetry exaggerated for visualization. (a) Side view at angle 0. (b) Top view at height 0.5.

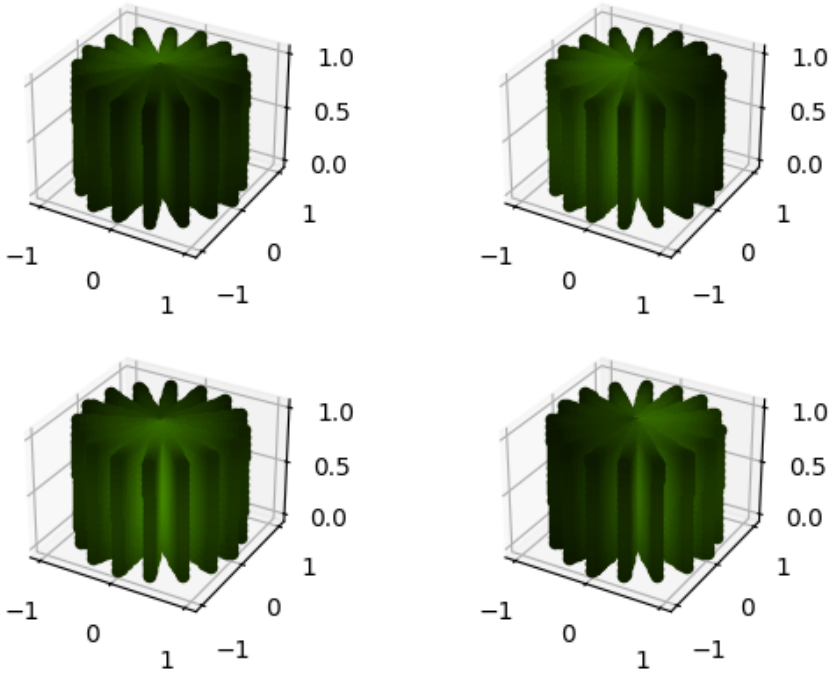

Figure 5. Asymmetric 3D power density distribution shown at four angles, each separated by 90 degrees.

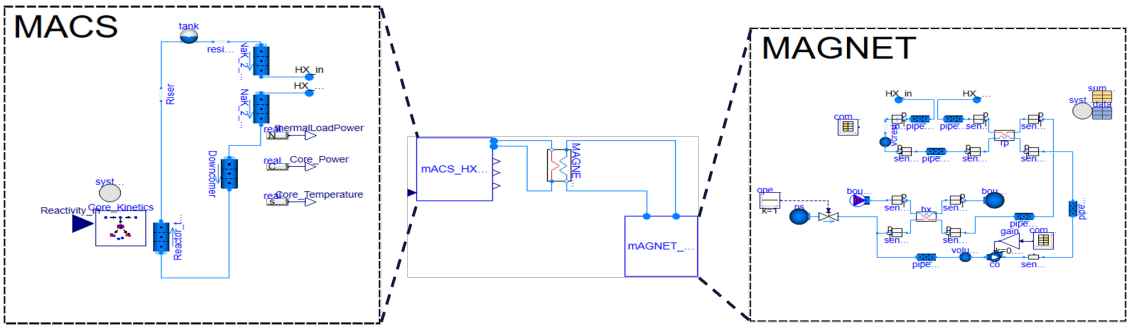

Figure 6. Coupled balance of plant TRANSFORM model of reactor system with NaK primary coolant loop and N_2 secondary system.

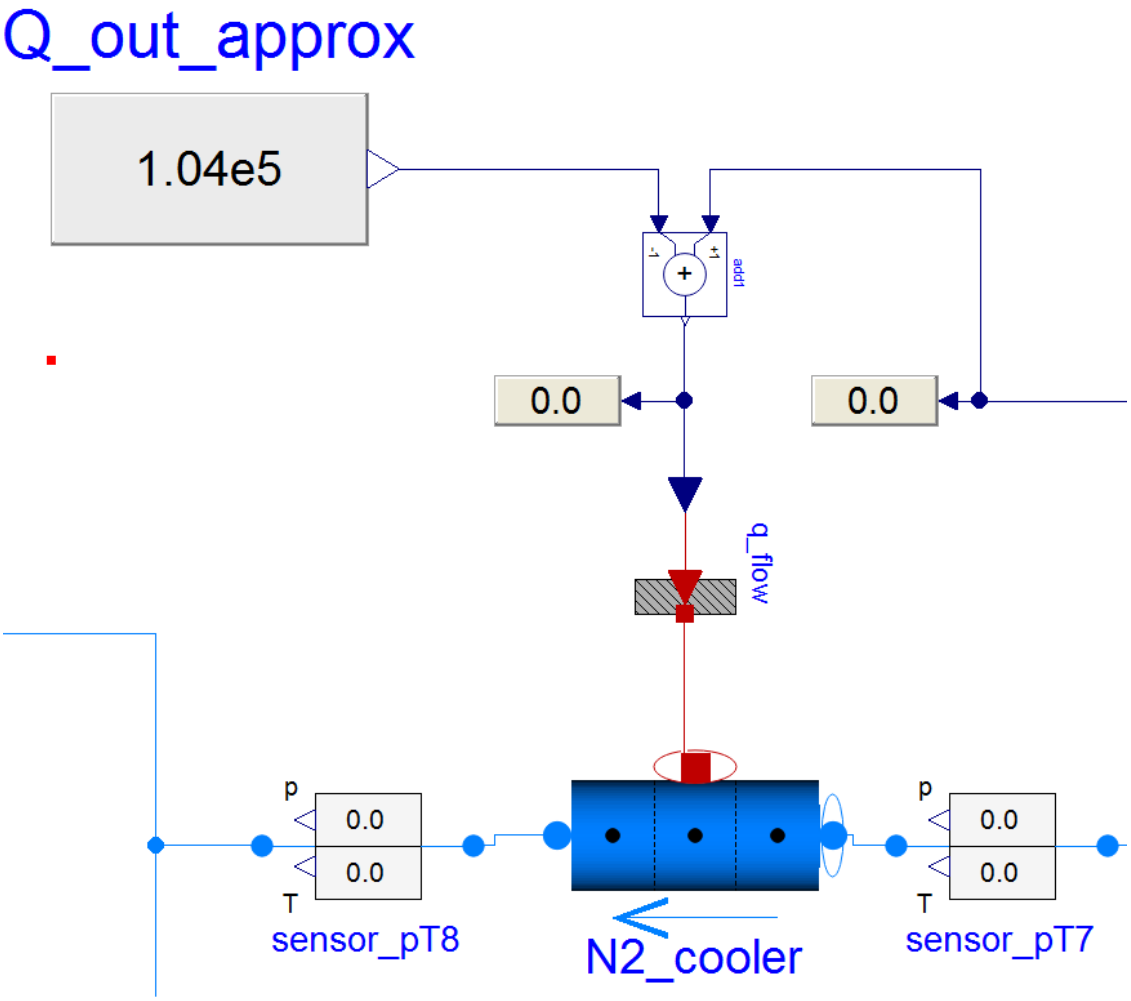

Figure 7. A magnified view of the replacement system for power extraction in MAGNET.

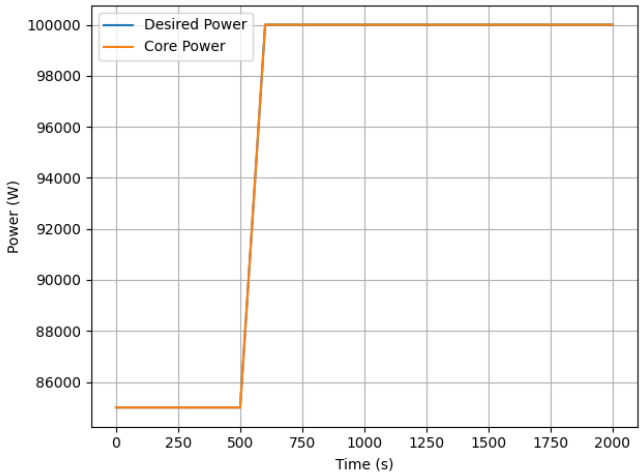

(a)

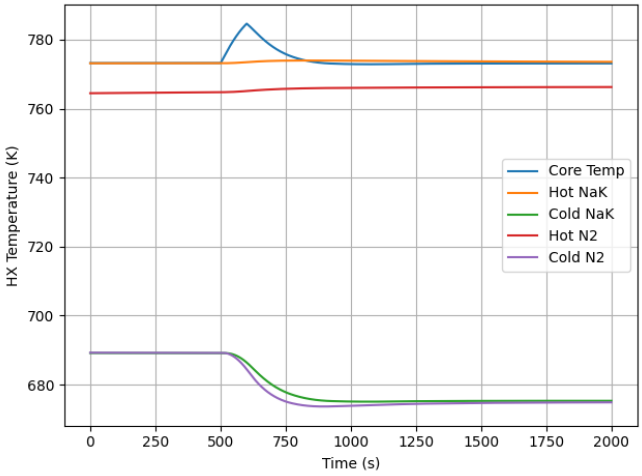

(b)

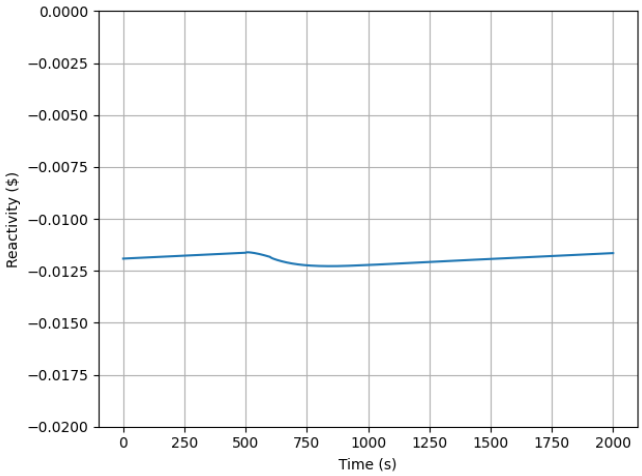

(c)

Figure 8. Balance of plant model performance with core temperature controlled to 500°C and core power controlled from 85 kW to 100 kW over 100 s with independent PID controllers. (a) Power trend following. (b) Temperature response of core and counter-flow inlet and outlet temperatures at the central heat exchanger. (c) Controlled reactivity insertion response to changing power setpoint.

A ramp power setpoint change was then performed from 85 kW to 100 kW over 100 s. The results of this ramp power test are presented in Figure 8. The power setpoint trend following in Figure 8(a) shows a similar well-controlled behavior as that observed for tests performed on the reactor core and primary coolant loop model in isolation [6, 10, 11]. The temperature response in Figure 8(b) shows a swift response from the temperature control PID to remove additional heat from the N_2 flow loop to compensate for the increased power generation upstream. This is actualized as a small spike in core temperature, which is mitigated by a reduction in coolant temperature from additional power extraction. While a more realistic system, such as a Brayton cycle will have efficiency limitations on power draw and the operating temperature of the coolant, the behavior displayed is promising for the implementation of more coupled control strategies, such as MPC.

4. MICROREACTOR AUTOMATED CONTROL SYSTEMS: ADVANCED CONTROLS UPDATES

The introduction of a more complex system model requires the development of a more sophisticated control strategy. MPC is a form of advanced control that originated in the process industries and has become a staple for optimizing system performance under constraints. MPC uses the concept of predicting system behavior over a finite time horizon using a mathematical model of the process or system. At each time step, it solves an optimization problem to determine the optimal control inputs, which are applied in a receding horizon manner (i.e., only the first control action is implemented, and the process repeats at the next time step). MPC is particularly suitable for multivariable systems with constraints. It is a versatile solution that is used across various industrial and engineering domains: application areas include process control; power grid management; autonomous vehicle navigation, path planning, and obstacle avoidance; guidance, navigation, and control of spacecraft; and aircraft control systems. MPC has also been investigated for nuclear power plant control [7, 25, 26] and shown to be capable of effective reactivity control.

In this work, MPC was implemented within the software modules intended for integration with the MACS hardware to demonstrate advanced control systems and autonomous operations in the event of system issues. The incorporation of MPC also provides a pathway for future interface development for advanced control algorithms into the MACS HIL testbed.

4.1 CONTROL PROBLEM FORMULATION

The control problem is formulated as a state-space problem, with state transition and measurement equations. In the continuous case (the case assumed in MACS), this system state-space model is defined as follows.

$$\dot{x}(t) = f(x(t), u(t), p(t), p_{\text{tv}}(t)) \tag{6}$$
$$y(t) = h(x(t), u(t), p(t), p_{\text{tv}}(t)) \tag{7}$$

The states of the system at time t are given by $x(t)$, the control inputs by $u(t)$, (uncertain) parameters by $p(t)$, time-varying (but known) parameters by $p_{\text{tv}}(t)$, and measurements by $y(t)$, respectively.

MPC then uses this setup to solve the optimal control problem over a time horizon N using the cost function and associated constraints as follows.

$$\min_{\mathbf{x}_{0:N+1}, \mathbf{u}_{0:N}, \mathbf{z}_{0:N}} \quad m(x_{N+1}) + \sum_{k=0}^{N} l(x_k, z_k, u_k, p_k, p_{\text{tv},k}) \tag{8}$$
$$\tag{9}$$

subject to the following constraints:

$$x_0 = \hat{x}_0, \tag{10}$$
$$x_{k+1} = f(x_k, u_k, p_k, p_{\text{tv},k}), \qquad \forall k = 0, \ldots, N, \tag{11}$$
$$g(x_k, u_k, p_k, p_{\text{tv},k}) \leq 0 \qquad \forall k = 0, \ldots, N, \tag{12}$$
$$x_{\text{lb}} \leq x_k \leq x_{\text{ub}}, \qquad \forall k = 0, \ldots, N, \tag{13}$$
$$u_{\text{lb}} \leq u_k \leq u_{\text{ub}}, \qquad \forall k = 0, \ldots, N, \tag{14}$$
$$g_{\text{terminal}}(x_{N+1}) \leq 0, \tag{15}$$

In the present work, the cost function $l(\cdot)$ for the MPC optimization used the tracking error between the reactor power setpoint and the actual power [27]. A penalty term was included that penalized large changes in the control inputs (reactivity insertion and coolant flow rate). This term promotes smooth transitions and realistic changes in control which can be actualized in a HIL system.

At each time step, MPC uses the current state estimate $\hat{x}_0$, which is either measured (state-feedback) or estimated based on an incomplete measurement, y_k, to predict the state evolution over the prediction time horizon N and attempts to find the optimal input (control) sequence u_k that minimizes the cost function in 8, subject to the constraints listed above.

The constraints include upper and lower bounds (x_{lb}, x_{ub}) for the states and inputs (u_{lb}, u_{ub}), including power, core temperature, flow, and reactivity insertion (as a surrogate to the control drum position). Terminal constraints can be added using $g_{\text{terminal}}(\cdot)$, and general nonlinear constraints can be defined with $g(\cdot)$, though these constraints were not included in the tests conducted to date. Note that the solution is a sequence of control inputs, and only the first value in that sequence is applied to the system. The entire process then restarts with the new measurements and state estimates.

Given the reactor point kinetics model defined in 4.1.1, we assume that the control variables (inputs) are the reactivity insertion ρ and coolant flow rate $\dot{m}$. Measurements of the reactor power P, coolant temperature T_{cl}, reactivity insertion ρ, and coolant flow rate $\dot{m}$ are assumed to be available, with measurement noise added for each measurement of the system state. For simplicity, the reactivity insertion is assumed to be directly measurable in this setup and the control drum angle estimated from the reactivity; in practice, the control drum position (angle) will be the measured quantity and the reactivity insertion will be estimated from the drum angle.

4.1.1 System Surrogate Model

The MPC algorithm needs a system model that can be used for estimating the system states over the selected prediction horizon. In addition to being relatively accurate, the model needs to be fast in order to support real-time control decisions. In the present study, the reactor point kinetics model coupled with heat transfer from the core to the coolant was used. The six-group point kinetics model [25] is as follows:

$$\dot{P} = \frac{\rho - \beta}{\lambda} P + \sum_{i=1}^{6} \lambda_i C_i + S_0, \tag{16}$$

$$\dot{C}_i = \frac{\beta_i}{\lambda} P - \lambda_i C_i, \quad i = 1, 2, \ldots, 6. \tag{17}$$

Temperature feedback effects were calculated using [28] the following:

$$\dot{T}_f = \frac{P - \left(T_f - Tcl\right) U_f A_{FC}}{C_{fh}}, \tag{18}$$

$$\dot{T}_{cl} = \frac{\left(T_f - Tcl\right) U_f A_{FC} - \left(T_{cl} - T_c\right) U_{cl} A_{CC}}{C_{clh}}, \tag{19}$$

$$\dot{T}_c = \frac{0.5\,(T_{cl} - T_c)\,U_{cl}A_{CC} - \dot{m}C_{ch_s}\,(T_c - T_{in})}{0.5C_{ch}}. \tag{20}$$

A simple HX is assumed to remove heat from the coolant and supply it to a heat sink:

$$\dot{T}_{in} = \frac{1}{mCch_s}\Big(C_{ch_s}\dot{m}T_c - C_{ch_s}\dot{m}T_{in} - H_x k\,(T_{in} - T_{hs})\Big). \tag{21}$$

In these equations ρ, β, and λ represent the reactivity, effective delayed neutron fraction, and prompt neutron lifetime, respectively. β_i and λ_i, for $(i = 1, \ldots, 6)$, denote the fraction of delayed neutrons and the decay constant for each precursor group. S_0 is the neutron source term (assumed to be non-zero only at startup). T_f, T_{cl}, and T_c refer to the fuel, clad, and coolant temperatures. U_f, U_{cl}, A_{FC}, and A_{CC} are heat transfer coefficients and core area–related parameters. C_{fh}, C_{clh}, C_{ch_s}, and C_{ch} denote specific heat capacities (fuel, cladding, coolant). $\dot{m}$ is the coolant mass flow rate. T_{in} and T_{hs} represent the inlet temperature to the heat exchanger and the ultimate heat sink temperature (e.g., a cooling tower or environment). m is the coolant mass, k is the effective heat transfer area, and H_x is the conductive heat transfer coefficient related to the HX's operation and design.

4.1.2 MPC Integration with HIL Platform

The MPC controller routine was modified to integrate with the Python modular gRPC client for future HIL simulation. A LabVIEW gRPC server was utilized to receive and reply to gRPC signals for drum setpoint and current drum position, respectively. This setup follows the data flow displayed in Figure 9. Sensor data (drum positions and reactor power and temperature output) may be pulled from the hardware using gRPC. The drum positions are converted to a single reactivity insertion value to be provided to the FMU for simulation. For initial testing, the TRANSFORM representation of a reactor core and NaK primary coolant loop was employed. The FMU simulation produced updated system variables, such as core power, coolant mass flow, and core temperature. These variable values were used to estimate the state of the system using the simplified system model in 4.1.1; the resulting system state estimate includes internal variables such as precursor power contribution. This state estimate was provided to the MPC algorithm along with the current power setpoint to calculate the next control action. The MPC algorithm produced the required reactivity insertion to achieve the desired power output which was converted to a setpoint for drum positions. The requested drum position was provided to the hardware through gRPC, and the cycle repeated. In practice, the ViBRANT hardware uses PID control for the drum actuators to achieve requested drum position.

An initial test of this data flow was performed by bypassing the hardware with an assumption of instantaneous drum angle actualization while still passing data through the gRPC server. This test, which uses the HIL software platform, mimics that performed in previous work with a simple PID algorithm for reactor power control [6, 10, 11]. The Modelica FMU of the reactor core and NaK coolant loop dynamics replaced the model described in Section 4.1.1 as the simulated reactor to provide a more realistic dynamic response to suggested control inputs from MPC. Internally, the MPC algorithm still employed the simplified model to find the optimal control solution.

4.2 MPC TEST RESULTS

4.2.1 Algorithm Testing

Figures 10 and 12 show examples of the MPC algorithm using the reactor point kinetics model in a software-only setting (i.e., without integration with the HIL software platform). This test, performed to verify functionality of the MPC algorithm, used the simplified reactor point kinetics model described in 4.1.1 for estimating the system state variables and for the control input optimization.

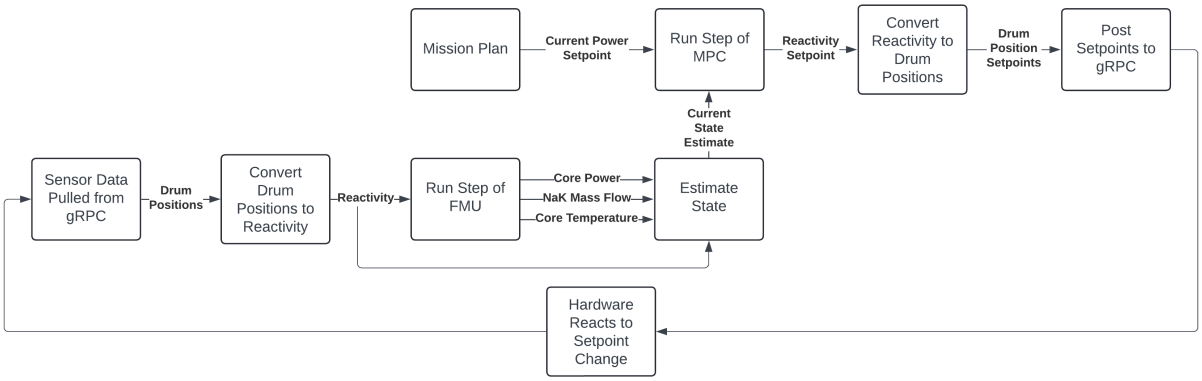

Figure 9. The data flow for controlling the MACS hardware with MPC.

The figures show the power setpoint and the actual reactor power as a result of applying the controller, the core temperature, and the reactivity insertion. The results were obtained using a 0.5 s time-step. The results show that the controller itself is capable of estimating the control inputs correctly to ensure that the reactor meets the desired power output requirements. The results show an early phase where the reactivity, power, and temperature are somewhat unstable before the controller and system settle into a steady state behavior. An examination of the underlying state variables indicates that this early stages of instability may be due to computational issues, arising from slight mismatches in the initial conditions for the various state variables in the simulator or from the simplified representation of the system dynamics leading to numerical instability at higher power levels under the coarse simulation time step. This issue appears to be resolved as the computations converge to consistent values, and the MPC appears to also converge and improve its tracking of the setpoints. Similar results are seen in Figure 11 and Figure 12, with some inconsistent behavior early in the simulation before the simulation appears to converge and the MPC begins to track the setpoints. Once this behavior is reached, the MPC appears to be capable of very accurate tracking in the examples studied.

These results also appear to showcase the capability of MPC, with the controller appearing to stabilize and accurately follow the setpoint for a rapidly varying power demand. However, these are still preliminary results, since the effects of various real-world behaviours such as communication and actuation delays are not included in these simulations. Ongoing work is investigating fine-tuning the model parameters, initial conditions, and constraints to eliminate this mismatch and the apparent initial numerical issues in the calculated control input. The next section discusses the results from the use of the MPC with a more realistic plant-level model to simulate the microreactor and the MPC integration with the HIL platform.

4.2.2 MPC Integration Testing

The integration of the MPC routines with the HIL software platform were evaluated using multiple operational load profiles. The performance of the system in response to a requested sinusoidal heat load change is presented in Figure 13.

The reactor model shows reasonable power-following capabilities for a 15 kW amplitude, 0.01 Hz sinusoidal trend about the nominal power of 85 kW. Most notably, the MPC seems to reduce delay in response time to changes in power setpoint when compared to the performance of a tuned PID (Fig. 3). Initial testing with full data loop performance showed a difficulty for the MPC-integrated client to maintain real-time at 0.3 s time steps, as was achieved with PID control. The results displayed in Figure 13 were collected with a 0.5 s simulation time step. This speed reduction can be somewhat remedied by the implementation

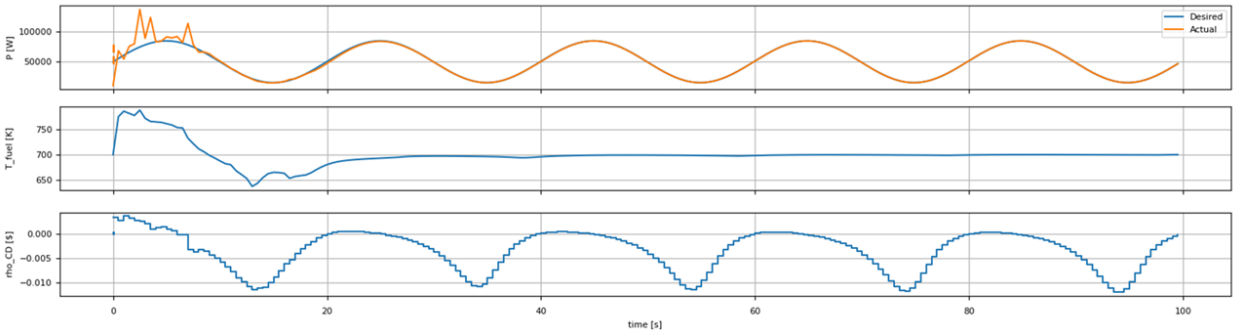

Figure 10. MPC control for 0.05 Hz sinusoidal power demand.

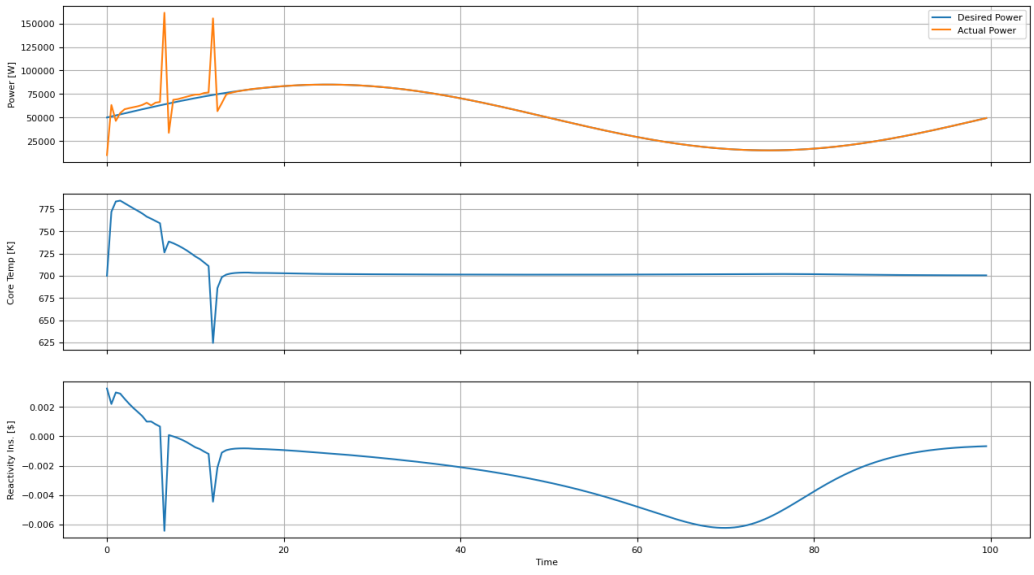

Figure 11. MPC control for 0.01 Hz sinusoidal power demand.

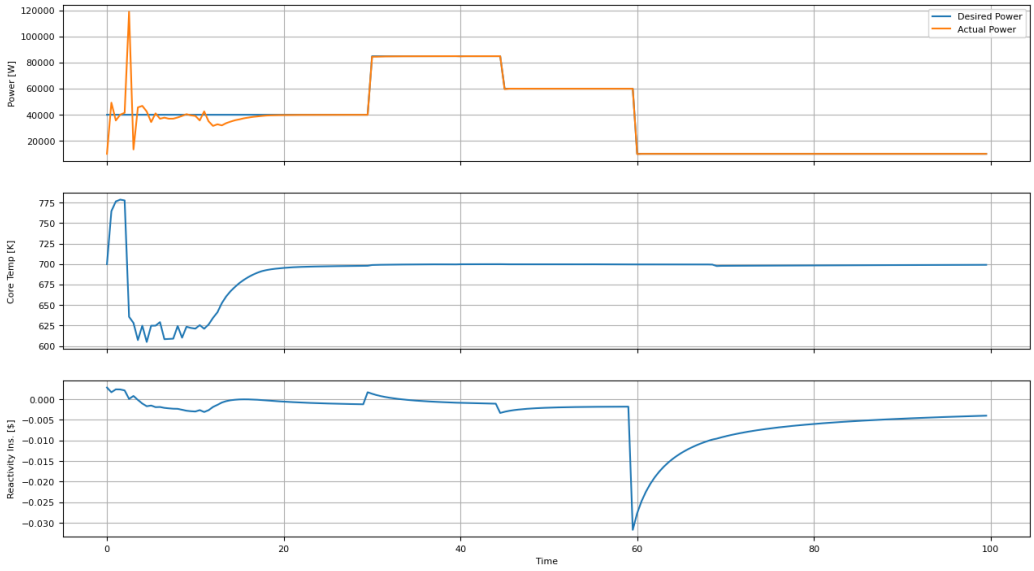

Figure 12. MPC control for step changes in power demand at randomly selected instants.

of machine-learned surrogate models to replace key aspects of the system, such as FMU simulation, state estimation, and MPC optimization. These system speed-ups are considerations for future research. While the resulting core power trend follows the requested system heat load reasonably, the controller requested reactivity shows some inconsistency between low energy and high energy performance. The jagged nature of the control output suggests a few possible diagnoses. The simplified representation of the system dynamics employed by MPC may incur numerical instability at higher power levels under the coarse simulation time step. Additionally, it is possible that differences in FMU performance and MPC prediction may result in a resonance between controller correction and power reaction which is amplified in the small adjustments of the reactivity insertion setpoint. This difference in simulated performance is most likely exacerbated by the slight delay in power load following. Further investigation of the MPC-integrated client performance is being conducted to expand the range of real-time operation.

A similar test is performed in Figure 14 for a step change. Desired power starts at a steady state 40 kW. A step is taken up to 85 kW at 30 s, down to 60 kW at 60 s, and finally down to 0 kW at 90 s. This is intended to represent a sporadic change in power demand. The MPC appears to perform well in power load following, however, system response is noticeably faster for steps down as opposed to steps up. It also appears that stepping to power levels near the nominal full power of 85 kW result in some minor power oscillation, however, this does not appear to have long-term effects on overall performance.

It should be noted, there are key differences between the model used within the MPC algorithm and the TRANSFORM FMU. Specifically, the simpler model considers a lump mass for approximating fuel temperature, whereas the FMU directly applies heat to the relevant section of the NaK coolant. The temperature of the hot-side section of the pipe flow discretization is used as the core reference temperature for temperature feedback. This means that the core temperature value displayed for the FMU-integrated MPC test is more prone to temperature spikes, as local thermal inertia is smaller than that considered for a lumped fuel mass.

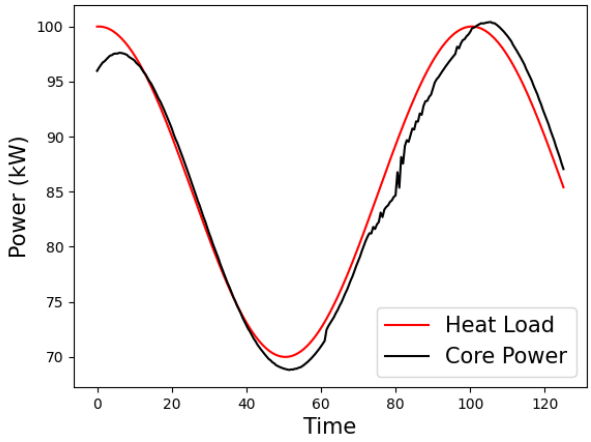

(a)

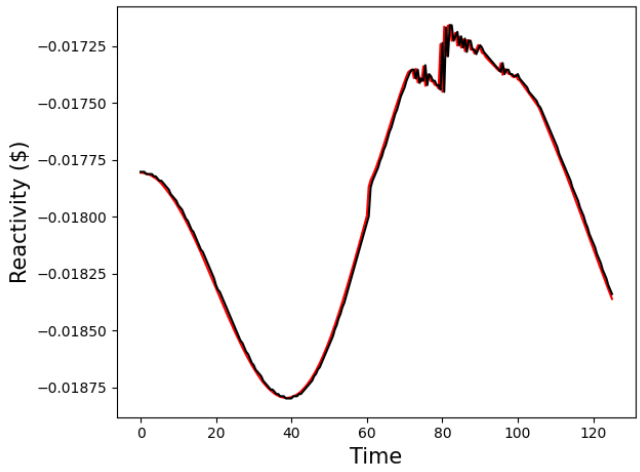

(b)

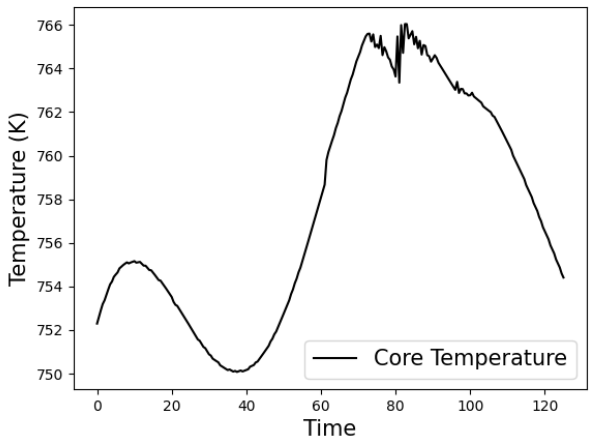

(c)

Figure 13. Performance of MPC on a modular Python client controlling reactivity insertion to a Modelica FMU simulation of reactor core and primary coolant dynamics. (a) Sinusoidal power trend following. (b) Reactivity insertion. (c) Core temperature response.

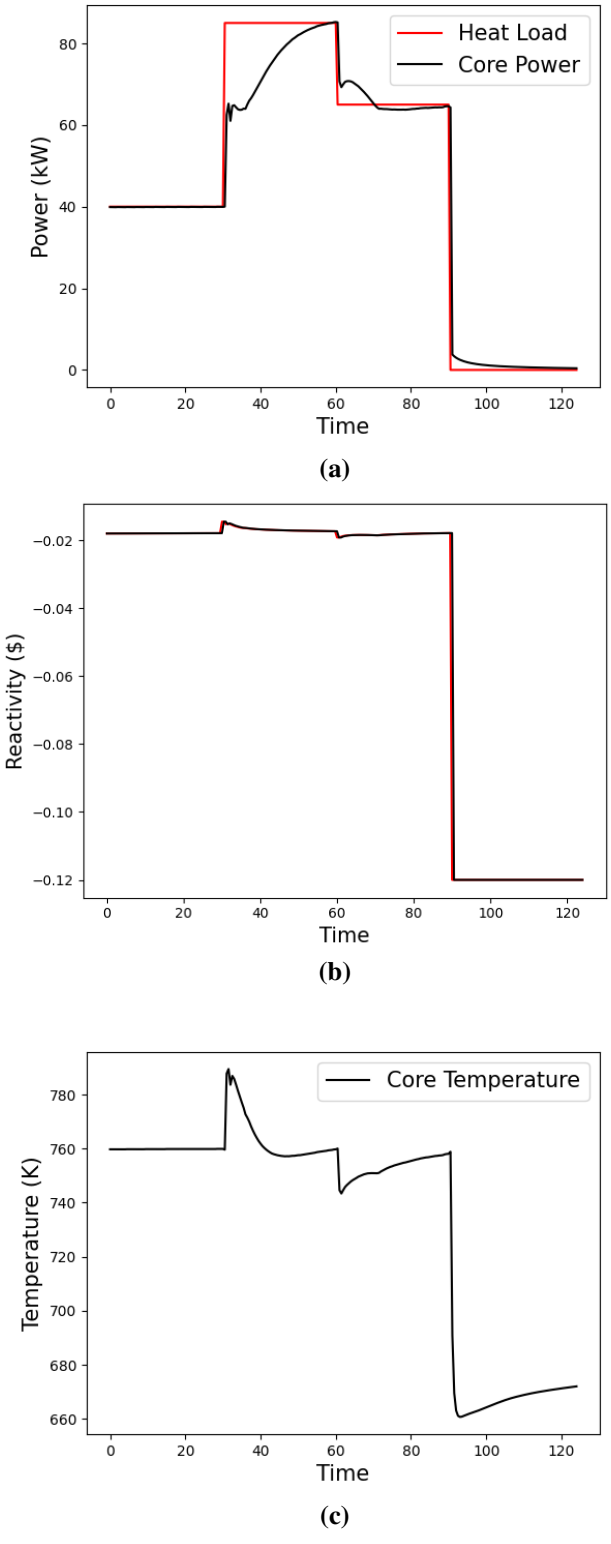

Figure 14. Performance of MPC on a modular Python client controlling reactivity insertion to a Modelica FMU simulation of reactor core and primary coolant dynamics. (a) Step change power trend following. (b) Reactivity insertion. (c) Core temperature response.

5. SUMMARY AND FUTURE WORK

Automation of control systems is expected to be important in the economic and safe operation of microreactors. A MACS testbed that includes HIL simulation was developed in collaboration with INL for the purpose of testing control system automation.

Research to date has resulted in the development of prototypic software for an automated plant-level control system to interact with HIL for automating control under selected scenarios. The software includes a plant-level digital twin that consists of the reactor and a prototypic BOP system. A model-predictive controller was shown to be capable of reading sensor data from the HIL MACS testbed and calculating and sending drum position setpoints to the MACS testbed. Tests to date were performed using the testbed's software platform, with future testing using the MACS hardware planned.

Previous testing at INL demonstrated the ability of the MACS testbed to effectively and directly control the drum positions based on multiple test scenarios using the control setpoints computed by the gRPC client-side modules. The present work expanded on that effort by integrating the plant-level digital twin model with the MPC controller to demonstrate the utility of the digital twin in investigating control solutions in a supervisory control setting. It should be noted that the client-side modules collectively behaved like a supervisory control system in estimating the control setpoints for each scenario, though the client side modules are not, strictly speaking, a supervisory control system. The MPC algorithm also showed the potential for limited evaluation of multiple competing options for control, adding to the capabilities of the testbed toward a future true supervisory system. It is expected that this feature will be incorporated in follow-up work, leveraging work being conducted under other DOE-NE programs.

While results from some transient conditions were simulated in this study, the response of MACS to off-normal microreactor conditions using data from sensors defined by INL remains to be fully evaluated. Modifications to the client-side control algorithms to fully utilize these additional sensor data are ongoing and will be incorporated into the next revision of the software modules. Additional modifications to the software to interface with the COMMAND software system are also planned.

6. ACKNOWLEDGMENTS

The work described in this report would not have been possible without the assistance provided by Dr. Anthony Crawford and Andrew Heim at INL. Andrew Heim's assistance in providing updated software platforms for the MACS-ViBRANT hardware was critical to testing the models and control algorithms. The work presented here was supported by the US Department of Energy (DOE) Office of Nuclear Energy's (NE's) Microreactor Research Program. The authors would like to thank Dr. John Jackson (INL), and Dr. Holly Trellue (LANL) for their support of this research, as well as the technical editors and reviewers for their feedback.

7. REFERENCES

[1] *Roadmap for the Deployment of Micro-Reactors for U.S. Department of Defense Domestic Installations*. Tech. rep. Nuclear Energy Institute, 2018.

[2] D. Gerster and Y. Arafat. *MARVEL 90% Final Design Report*. Tech. rep. INL/RPT-23-74280. INL, 2023.

[3] J. Walksman. *Project Pele Overview: Mobile Nuclear Power For Future DoD Needs*. Tech. rep. 2022.

[4] M. Nichol. *Micro-Reactors*. Presented at GAIN-EPRI-NEI-US NIC Micro-Reactor Workshop. 2019.

[5] P. Ramuhalli and S. Cetiner. *Concepts for autonomous operation of microreactors*. Tech. rep. Oak Ridge National Laboratory, 2019.

[6] J. McConnell and P. Ramuhalli. *An Update on Microreactor Automated Control System (MACS)*. Tech. rep. ORNL/TM-2024/3595. ORNL, 2024.

[7] "Autonomous control for Heat-Pipe microreactor using Data-Driven model predictive control". In: *Annals of Nuclear Energy* 200 (2024), p. 110399. ISSN: 0306-4549.

[8] J. O'Hara et al. *Human Factors Engineering Program Review Model*. Tech. rep. NUREG-0711 Revision 3. U.S.NRC, 2012.

[9] J. O'Hara and S. Fleger. *Human-System Interface Design Review Guidelines*. Tech. rep. NUREG-0700 Revision 3. U.S.NRC, 2020.

[10] J. McConnell, P. Ramuhalli, and W. Williams. "Microreactor Simulation Performance Analysis for Real-Time Hardware-in-the-Loop Testbed". In: *Transactions of the American Nuclear Society*. Vol. 131. ANS. Orlando, FL, Nov. 2024, pp. 318–321.

[11] J. McConnell et al. "Microreactor Automated Control System Test Bed Digital Architecture for Real-Time, Hardware-in-the-Loop Simulation". In: *Proceedings of NPIC&HMIT 2025*. Chicago, IL, June 2025, pp. 402–411.

[12] Michael D. Muhlheim et al. "First-of-a-Kind Risk-Informed Digital Twin for Operational Decision Making". In: *Nuclear Science and Engineering* 0.0 (2025), pp. 1–16. DOI: 10.1080/00295639.2025.2471724. eprint: https://doi.org/10.1080/00295639.2025.2471724. URL: https://doi.org/10.1080/00295639.2025.2471724.

[13] Mustafa Cetiner et al. *Integrated risk-informed decision-making for an ALMR PRISM*. Tech. rep. ORNL/TM-2016/211. Oak Ridge National Laboratory, 2016.

[14] Sacit M Cetiner et al. *Development of an automated decision-making tool for supervisory control system*. Tech. rep. ORNL/TM-2014/363. Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), 2014.

[15] R. Ponciroli et al. *Design and Prototyping of Advanced Control Systems for Advanced Reactors Operating in the Future Electric Grid (Final Report)*. Tech. rep. Argonne National Laboratory (ANL), Argonne, IL (United States), Sept. 2024. DOI: 10.2172/2478767.

[16] A. Crawford. "The ViBRANT Platform, a Surrogate Nuclear Reactor". In: *Proceedings of the Pacific Basin Nuclear Conference*. Idaho Falls, ID, USA: American Nuclear Society, Oct. 2024, p. 389.

[17] A. Crawford and et al. *Complete development of effective communication methods to support the implementation of temperature measurement and light emitters/sensors into the surrogate reactor cell*. Tech. rep. M3AT-24IN0804033. Idaho National Laboratory, 2024.

[18] A. Crawford et al. "ViBRANT Applied as a Closed Loop Flux and Thermal Surrogate". In: *Proceedings of NPICHMIT 2025*. Chicago, IL, June 2025, pp. 824–831.

[19] *TRANSFORM-Library*. URL: https://github.com/ORNL-Modelica/TRANSFORM-Library.

[20] M. S. Greenwood. "TRANSFORM – A Vision for Modern Advanced Reactor System-Level Modeling and Simulation Using Modelica". In: *Transactions of the American Nuclear Society*. Chicago, IL, Nov. 2020, pp. 1097–1100.
[21] Modelica Association. *Functional Mockup Interface*. URL: https://fmi-standard.org/.
[22] J. Farber et al. *Enabling a Physical Twin for Control Methods Evaluation*. Tech. rep. INL/RPT-24-77114. Idaho National Laboratory, 2023.
[23] T. Morton, J. O'Brien, and J. Hartvigsen. *Functional and Operating Requirements for the Microreactor Agile Non-Nuclear Experimental Test Bed*. Tech. rep. INL/EXT-20-58204-Revision-0. Idaho National Laboratory, 2020.
[24] *HYBRID*. URL: https://github.com/idaholab/HYBRID.
[25] Bikash Poudel, Kalpesh Joshi, and Ramakrishna Gokaraju. "A Dynamic Model of Small Modular Reactor Based Nuclear Plant for Power System Studies". In: *IEEE Transactions on Energy Conversion* 35.2 (2020), pp. 977–985.
[26] Linyu Lin et al. "Development and assessment of a model predictive controller enabling anticipatory control strategies for a heat-pipe system". In: *Progress in Nuclear Energy* 156 (Feb. 2023), p. 104527.
[27] James B. Rawlings, David Q. Mayne, and Moritz M. Diehl. *Model Predictive Control: Theory, Computation, and Design*. 2nd Edition. Nob Hill Publishing, 2017.
[28] B. R. Upadhyaya et al. *Autonomous Control of Space Reactor Systems: Final Report*. Tech. rep. DE-FG07-04ID14589/UTNE-06. University of Tennessee Knoxville, 2007.

APPENDIX A. POWER SHAPING VISUALIZATION AND DERIVATION

APPENDIX A. POWER SHAPING VISUALIZATION AND DERIVATION

The form of Eqs. (2–4) were selected to provide a tunable truncation factor of a cosine shaping as well as a tunable weighting for drum position as it relates to reactor tilt. Visuals of the shapes of Eqs. (2–4) are provided in Figures A.1–A.3, respectively.

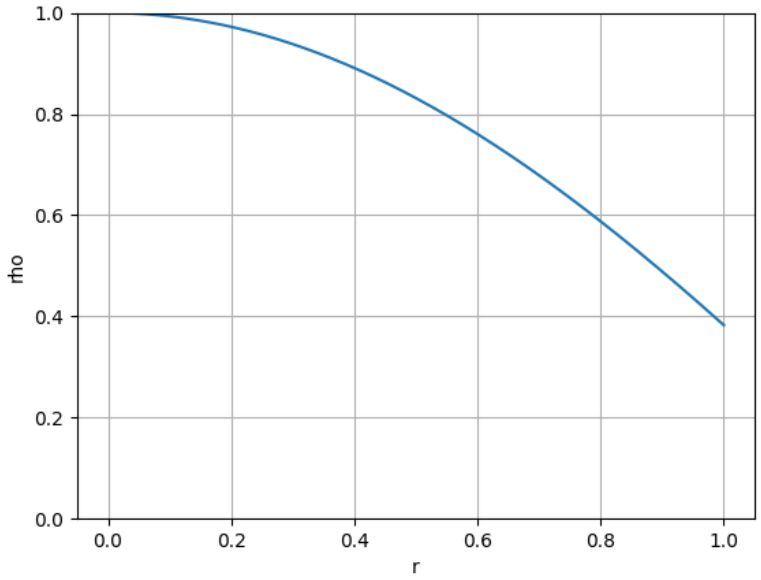

Figure A.1. Isolated radial power shaping with truncation factor a = 0.25.

As stated in Section 3.1, it is necessary to apply a correction factor $\bar{A}$ to the volumetric power distribution to ensure that the volumetric integral of the power density is equal to the original total power defined by point kinetics. For this condition to be true, Eq. (A.1) is as follows.

$$1 = \bar{A} \int_0^{2\pi} \int_0^{H} \int_0^{R} \rho_{radial}(r)\, \rho_{axial}(z)\, \rho_{angular}(\theta)\, drdzd\theta \tag{A.1}$$

Because each of the directional shaping functions depends on only one variable, the integral can be separated as shown in Eq. (A.2).

$$1 = \bar{A} \int_0^{R} \rho_{radial}(r)dr \int_0^{H} \rho_{axial}(z)dz \int_0^{2\pi} \rho_{angular}(\theta)d\theta \tag{A.2}$$

The function ρ_{radial} can be integrated as shown in Eq. (A.3).

$$\int_0^{R} \rho_{radial}(r)dr = \int_0^{R} \cos\left(\frac{(1-a)r\pi}{2R}\right)dr = \frac{2R}{(1-a)\pi} sin\left(\frac{(1-a)\pi}{2R}\right) \tag{A.3}$$

The function ρ_{axial} can be integrated as shown in Eq. (A.4).

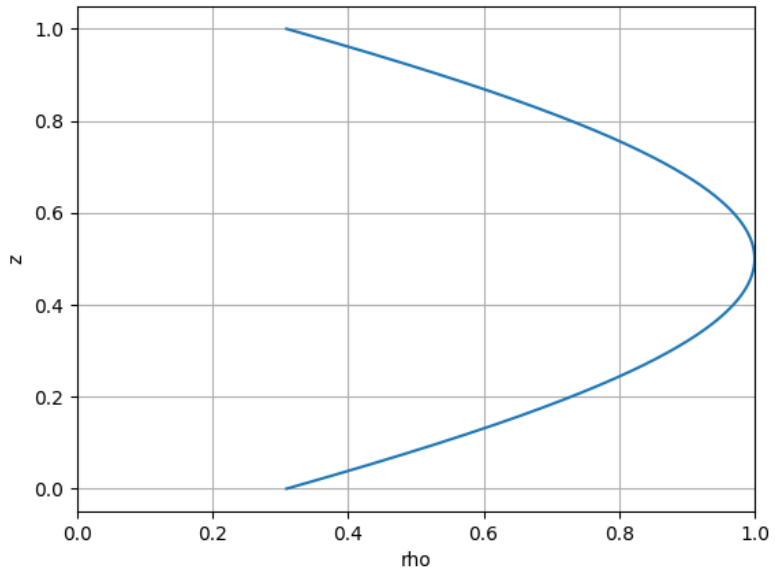

Figure A.2. Isolated axial power shaping with truncation factor b = 0.1.

$$\begin{aligned}\int_0^H \rho_{axial}(z)dz &= \int_0^H \cos\left(\pi\left(b + (1-2b)\frac{z}{H}\right) - \frac{\pi}{2}\right)dz \\ &= \frac{H}{(1-2b)\pi}\left(sin\left(\frac{\pi}{2} - \pi b\right) - sin\left(\pi b - \frac{\pi}{2}\right)\right) \\ &= \frac{2H}{(1-2b)\pi} sin\left(\frac{\pi}{2} - \pi b\right)\end{aligned} \tag{A.4}$$

The function $\rho_{angular}$ can be integrated as shown in Eq. (A.5).

$$\begin{aligned}\int_0^{2\pi} \rho_{angular}(\theta)d\theta &= \int_0^{2\pi} \sum_{i=0}^{N-1} \omega_i(\phi_i)\cos^2\left(\frac{\theta}{2} + \frac{i\pi}{N}\right)d\theta \\ &= \sum_{i=0}^{N-1} \int_0^{2\pi} \omega_i(\phi_i)\cos^2\left(\frac{\theta}{2} + \frac{i\pi}{N}\right)d\theta \\ &= \sum_{i=0}^{N-1} \omega_i(\phi_i)\left(\pi + \frac{i\pi}{N} + \frac{1}{2}sin\left(2\pi - \frac{2i\pi}{N}\right) - \left(\frac{i\pi}{N} + \frac{1}{2}sin\left(\frac{-2i\pi}{N}\right)\right)\right) \\ &= \sum_{i=0}^{N-1} \omega_i(\phi_i)\left(\pi - \frac{1}{2}sin\left(\frac{2i\pi}{N}\right) + \frac{1}{2}sin\left(\frac{2i\pi}{N}\right)\right) \\ &= \pi \sum_{i=0}^{N-1} \omega_i(\phi_i)\end{aligned} \tag{A.5}$$

With each integral calculated, $\bar{A}$ can be solved for in Eq. (A.6).

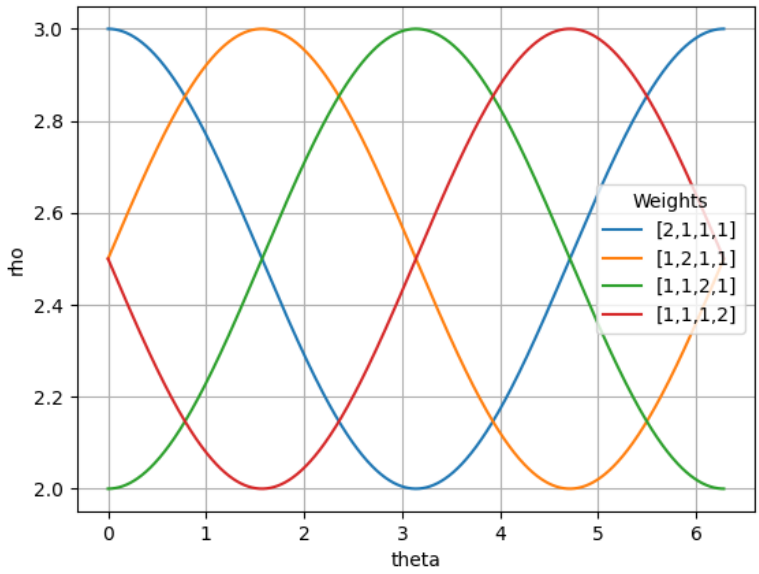

Figure A.3. Isolated angular power shaping with various tilt conditions.

$$\begin{aligned}\bar{A} &= \frac{1}{\left(\frac{2R}{(1-a)\pi} sin\left(\frac{(1-a)\pi}{2R}\right)\right)\left(\frac{2H}{(1-2b)\pi} sin\left(\frac{\pi}{2}-\pi b\right)\right)\left(\pi \sum_{i=0}^{N-1} \omega_i(\phi_i)\right)} \\ &= \frac{\pi(1-a)(1-2b)}{4RH sin\left(\frac{(1-a)\pi}{2R}\right) sin\left(\frac{\pi}{2}-\pi b\right) \sum_{i=0}^{N-1} \omega_i(\phi_i)}\end{aligned} \tag{A.6}$$