Preprint, not peer reviewed

Frequency-Dependent Dynamic Analysis of a Prismatic HTGR Power System for Data Center Applications

Wooseong Park[a,*], W. Reed Kendrick[a], Emilio Baglietto[a], Dean Price[a]

[a]*Nuclear Science and Engineering, Massachusetts Institute of Technology, 60 Vassar St, Cambridge, 02139, MA, US*

Abstract

The rapid growth of hyperscale data centers (DCs) has increased interest in nuclear energy as a reliable low-carbon power source. However, DC demand can exhibit multi-timescale variations, and their impact on nuclear power systems remains largely unexplored. This study investigates the frequency-dependent dynamic characteristics of a prismatic high-temperature gas-cooled reactor coupled to a steam Rankine cycle using an OpenModelica dynamic model. Systematic sinusoidal load perturbations were first applied over a range of amplitudes and periods to identify dynamically sensitive demand characteristics. The resulting characteristics were then examined against a year-long facility-aggregated DC demand profile and a reconstructed workload-resolved profile based on measured AI-workload GPU power traces. The sinusoidal analysis identified an intermediate time-scale region that produced the largest plant deviations, whereas faster fluctuations were attenuated and slower variations were effectively regulated. Consistent with this behavior, the facility-aggregated profile was successfully accommodated despite repeated demand variations between approximately 15% and 100% of nominal power because its dominant variations occur over long time scales. In contrast, retaining faster workload-driven fluctuations in the workload-resolved profile produced substantially larger plant responses. Frequency-band analysis further confirmed that intermediate-frequency components contributed more strongly to these responses than higher-frequency components. For example, under PI control in the energy-matched comparison, the 95th-percentile steam-temperature deviation decreased from 7.2°C to 5.5°C after high-frequency attenuation and to 3.6°C after intermediate-frequency attenuation. These results indicate that the spectral distribution of DC demand is an important consideration when coordinating nuclear generation with flexible resources.

Keywords:

Data center, High-temperature gas-cooled reactor (HTGR), Load following, Frequency-dependent dynamics, OpenModelica

*Corresponding author

Email address: wspark3000@gmail.com (Wooseong Park)

1. Introduction

The rapid expansion of artificial intelligence (AI) and cloud computing is driving unprecedented growth in electricity demand from data centers (DCs). The International Energy Agency (IEA) projects that global DC electricity consumption will more than double from approximately 415 TWh in 2024 to 945 TWh by 2030 [1]. In particular, generative AI has accelerated the deployment of hyperscale DCs with power demands reaching several hundred megawatts. Together with the growing need for reliable and low-carbon electricity, this trend has renewed interest in nuclear energy for hyperscale DCs and has increasingly translated into commercial activities [2, 3, 4]. These developments indicate that nuclear energy is emerging as an important power-supply option for hyperscale DCs.

The operational requirements of DC-connected nuclear systems can differ from those of conventional grid-connected plants. When nuclear generation and DC demand are tightly coupled through a local grid, a substantial portion of the load variation may need to be balanced within the coupled system rather than being fully buffered by external resources or a large interconnected grid, making dynamic flexibility an important design consideration. Flexible nuclear operation includes load following to track variations in electricity demand and frequency control to respond to short-term power imbalances [5]. Although specific requirements depend on reactor design and grid conditions, European Utility Requirements (EUR) and IAEA indicate representative capabilities that include load following from full power to the minimum operating load at approximately 3%/min, as well as frequency-control responses of about ±2% of rated power within 30 s [5, 6]. These capabilities have primarily been established in the context of conventional power-system operations. Whether such capabilities are sufficient for the faster and more irregular demand variations of hyperscale DCs, however, remains unclear.

Previous studies have extensively investigated nuclear load following dynamics with simplified load variations. Rad et al. examined a 5% step load change in a NuScale system [7], while Wu et al., Rivas et al., and Brits et al. evaluated HTGR systems under prescribed linear ramps, typically at 5%/min over broad operating ranges [8, 9, 10]. Other studies considered slower variations, including 24-h periodic loads and representative daily electricity-demand profiles [11, 12]. Related studies have investigated grid-frequency regulation: Wang et al. evaluated AP1000 frequency control following a step load change [13]; Joshi et al. examined SMR flexibility under variable photovoltaic generation [14]; and Masotti et al. analyzed an SMR response to a frequency-deviation transient initiated by a prescribed linear ramp [15]. Collectively, these studies demonstrate substantial nuclear flexibility, but they predominantly characterize responses to simplified step or linear disturbances rather than systematically examining the dependence of plant performance on both disturbance frequency and magnitude.

In contrast, hyperscale DC demand can exhibit irregular and multi-timescale variations driven by computing workloads. Large-scale AI training can synchronize rapid transitions between high- and low-power states across many computing units, producing MW-scale power swings [16]. Job scheduling, workload transitions, and facility-level operations introduce additional variability into aggregate demand [17, 18]. Recent studies have attempted to reconstruct facility-level AI loads using measured workload profiles or stochastic models; however, the resulting degree

of variability differs substantially depending on workload composition, aggregation level, temporal resolution, and modeling assumptions [17, 18, 19]. Despite this uncertainty, these studies suggest that hyperscale DC demand can exhibit substantially more complex temporal characteristics than the monotonic linear ramps commonly used in nuclear load following analysis. Because these variations span multiple time scales, their impact on a nuclear power system cannot be characterized solely by load magnitude or ramp rate. How the spectral characteristics of DC demand interact with the frequency-dependent dynamics of the nuclear power system therefore remains largely unexplored.

To address this gap, this study investigates the frequency-dependent dynamic characteristics of a nuclear power system under hyperscale DC demand. A dynamic model of a prismatic high-temperature gas-cooled reactor (HTGR) coupled to a steam Rankine cycle was developed in OpenModelica, incorporating the major neutronic, thermal-hydraulic, turbomachinery, and control dynamics of the power system. Systematic sinusoidal load perturbations over a range of amplitudes and frequencies were first used to identify dynamically sensitive time scales of the plant. The identified frequency-dependent characteristics were then evaluated against two contrasting facility-level DC demand profiles spanning distinct temporal regimes: 1) a year-long facility-aggregated profile, representing slow variations over a wide operating range, and 2) a workload-resolved reconstructed profile based on measured GPU power traces from AI workloads, retaining faster workload-driven fluctuations. Both proportional-integral (PI) and sliding-mode control (SMC) were considered to assess the dependence of the identified dynamic characteristics on the controller type. Through this analysis, this study establishes a framework for relating the spectral characteristics of DC demand to the dynamic load-following capability of nuclear power systems.

The remainder of this paper is organized as follows. Section 2 describes the reference plant and control system. Section 3 presents the OpenModelica dynamic model and the DC load profiles used in the analysis, including the load reconstruction method. Section 4 presents the dynamic-response results, with particular emphasis on frequency-dependent behavior. Finally, Section 5 summarizes the main conclusions and limitations of the study.

2. Reference Plant and Control System

2.1. Reference design

The reference system consists of a prismatic HTGR coupled to a steam Rankine power conversion cycle. The reactor design is based on the GEMINI+ prismatic HTGR [20]. The reactor has a nominal thermal power of 180 MWth and uses helium as the primary coolant at a nominal pressure of 6 MPa. The helium enters the reactor core at 325°C and exits at 750°C.

Because the balance of plant (BOP) of the original GEMINI+ design is intended for cogeneration, a Rankine power conversion system dedicated to electricity generation was configured for the present study. Hereafter, the term "reference plant" denotes the integrated system consisting of the GEMINI+-based reactor and the BOP configuration developed in this study. Figure 1 and Table 1 show the reference plant configuration and its nominal operating conditions, respectively. The secondary system consists of a once-through steam generator (OTSG) and a regenerative

Rankine cycle with a single stage of feedwater heating. The helical-tube OTSG was sized to transfer the full nominal reactor thermal power while satisfying the specified primary- and secondary-side operating conditions. A portion of the turbine steam is extracted at an intermediate pressure and supplied to the deaerator, which also serves as an open feedwater heater. Moisture separation and steam reheating were omitted to maintain a relatively simple reference configuration, while the high turbine inlet temperature limits the need for additional cycle complexity. The purpose of the present BOP design is not to optimize plant performance, but to establish a physically consistent steady-state reference for the subsequent dynamic analysis.

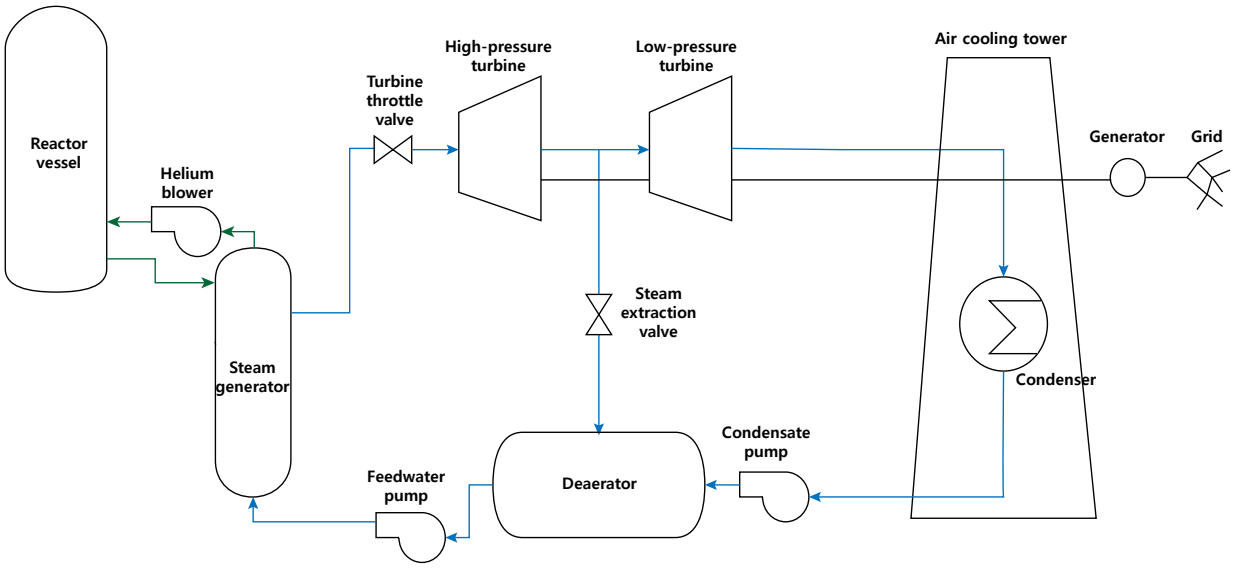

Figure 1: Reference plant configuration. Green and blue lines represent helium and water/steam loops, respectively, while arrows show direction of fluid flow.

Table 1: Nominal operating conditions of the reference plant.

Parameter	Value
Nominal thermal power	180 MWth
Nominal electrical power	68.6 MWe
Cycle efficiency	38.1%
Primary coolant	Helium
Primary coolant pressure	6 MPa
Reactor inlet temperature	325°C
Reactor outlet temperature	750°C
Primary coolant mass flow rate	81.6 kg/s
Main steam pressure	13.8 MPa
Main steam temperature	550°C
Feedwater temperature	198.8°C
Main steam mass flow rate	69.1 kg/s

2.2. Transient operating constraints

Table 2 summarizes the operating constraints applied during the transient simulations. According to the applicable steam-turbine standard, the turbine inlet pressure shall not exceed 105% of its rated value, while the inlet steam temperature shall not exceed the rated temperature by more than 8°C [21]. For the reactor core, the maximum fuel temperature during normal operation was limited to 1250°C to preserve TRISO fuel integrity and suppress fission-product release [22]. The core outlet helium temperature was limited to 850°C to remain within a temperature range compatible with technically mature structural materials such as SA508/533 [23]. The maximum reactivity addition rate was limited to 24 pcm/s based on the HTTR control-rod reactivity insertion limit [24]. This value is used as a reference criterion rather than as a demonstrated control-rod limit for the HTGR system. The allowable grid-frequency deviation depends strongly on the grid configuration and operating requirement. Generator-protection standards permit relatively large excursions, whereas frequency-trigger limits (FTL) used for normal balancing operation are typically on the order of several tens of millihertz [25, 26]. In the present study, a deviation of ±100 mHz was adopted as a representative reference for the calculated frequency response.

The listed thermal-hydraulic and reactivity constraints were used to identify limit exceedances, whereas grid-frequency deviation was used only as a representative screening metric because detailed grid-code compliance is outside the scope of the simplified local-grid model.

Table 2: Transient operating constraints of the reference plant used to evaluate dynamic simulations.

Parameter	Requirement
Steam pressure	≤105% of rated pressure
Steam temperature	≤+8°C of rated temperature
Maximum fuel temperature	≤1250°C
Core outlet helium temperature	≤850°C
Reactivity insertion rate	≤24 pcm/s
Grid frequency deviation	±100 mHz

2.3. Control system

2.3.1. Control strategy

A previous load-following study of an HTGR showed that turbine throttle-valve control provides a rapid means of responding to changes in electrical power demand [10]. Based on this approach, the present study employed a decentralized multi-loop single-input single-output (SISO) control architecture. Figure 2 illustrates the overall control logic. The turbine throttle valve serves as the primary actuator for grid-frequency regulation, while the remaining plant variables are regulated through the helium blower speed, feedwater pump speed, condensate pump speed, and external reactivity insertion.

The control sequence following a decrease in load demand can be summarized as follows:

(i) A decrease in load demand initially produces excess turbine electrical output, causing the grid frequency to increase. The resulting frequency deviation commands the turbine throttle valve to close and reduce the steam flow through the turbine.

(ii) The throttle-valve closure changes the steam pressure and temperature upstream of the valve. The feedwater pump speed is reduced to regulate the steam temperature, while the helium blower speed is also decreased to suppress the pressure rise by lowering the primary-side heat transfer.

(iii) The external negative reactivity is then inserted to regulate the core outlet helium temperature rise induced by the reduction in helium flow rate. Meanwhile, the condensate pump speed is reduced to maintain the condenser water level as the secondary-side mass flow rate decreases.

(iv) These coordinated control actions drive the plant toward a new equilibrium state with reduced reactor thermal power and turbine electrical output.

The opposite control sequence occurs during an increase in load demand. The required external reactivity to regulate the helium temperature is assumed to be ideally realized by the reactor control system, while the detailed dynamics of control-rod motion are not explicitly modeled.

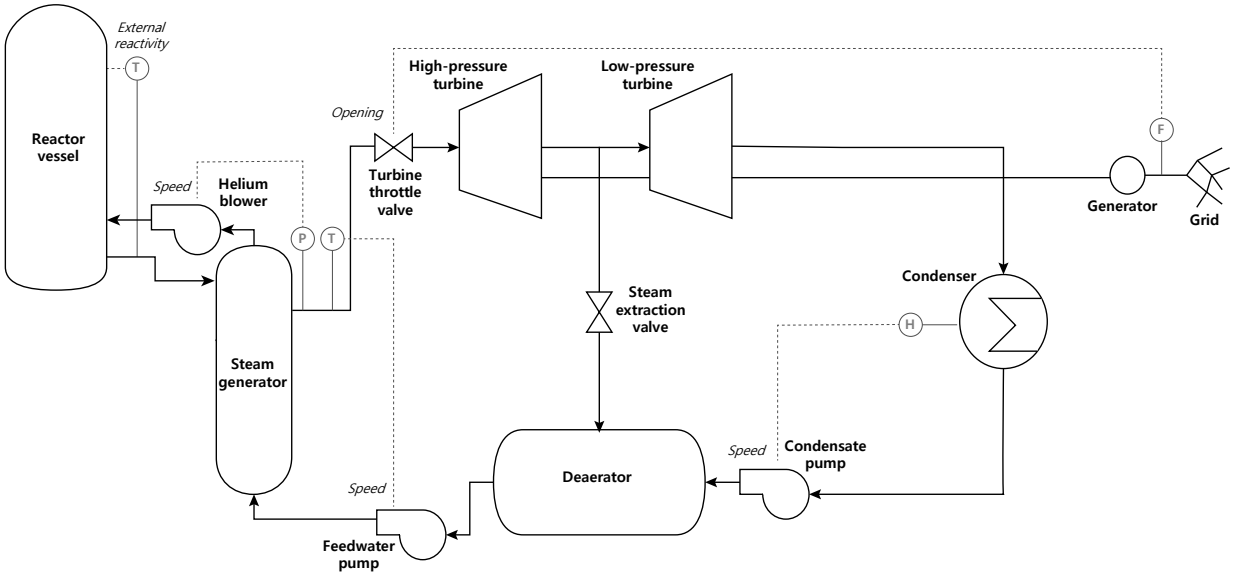

Figure 2: Plant control logic of the reference plant. Gray solid lines indicate measurement signals, where T, P, H, and F denote temperature, pressure, liquid level, and grid frequency, respectively. The corresponding control signals are indicated by gray dashed lines.

2.3.2. PI and SMC controllers

Because the dynamic response of the plant can depend on the control algorithm, both proportional-integral (PI) control and sliding-mode control (SMC) were considered to assess the controller dependence of the identified system characteristics. PI control was selected as a representative linear feedback controller because of its simplicity and widespread use in industrial process control [27]. The PI control law is expressed as:

$$u(t) = K_p\, e(t) + K_i \int_0^t e(\tau)\, d\tau \tag{1}$$

where u is the control action, e is the control error, and K_p and K_i are the proportional and integral gains, respectively.

SMC was additionally employed as a representative nonlinear robust control method. SMC drives the system trajectory toward a prescribed sliding surface and is generally less sensitive to external disturbances and model uncertainties once the sliding condition is established [28]. Conventional first-order SMC employs a discontinuous switching function and may therefore introduce high-frequency chattering. To reduce this behavior, the super-twisting algorithm, a second-order sliding-mode method, was adopted in the present study. In addition, the discontinuous sign function was replaced by a smooth hyperbolic tangent approximation within a finite boundary layer.

The sliding variable was defined as:

$$\sigma(t) = \hat{e}(t) + c\, e(t) \tag{2}$$

where $c > 0$ is the sliding-surface coefficient and $\hat{e}$ denotes the filtered derivative of the tracking error. To suppress high-frequency fluctuations that can be amplified by direct differentiation of the tracking error, a first-order filtered

derivative with a time constant of 5 s was used. The implemented super-twisting control law is:

$$\begin{aligned} u(t) &= \lambda \left|\sigma\right|^{1/2} \tanh\left(\frac{\sigma}{\phi}\right) + u_1 \\ \dot{u}_1 &= \alpha \tanh\left(\frac{\sigma}{\phi}\right) \end{aligned} \tag{3}$$

where λ and α are positive control gains and ϕ is the boundary-layer thickness used to smooth the switching action.

3. Dynamic Modeling and Analysis Methodology

3.1. Dynamic system modeling in OpenModelica

To investigate the dynamic response of the reference plant under DC load variations, a system-level dynamic model was developed in OpenModelica [29]. Version 3.1 of the ThermoPower library was used as the basis for the thermal-hydraulic and turbomachinery components [30], with several additional and modified component models implemented to represent the specific features of the reference plant. OpenModelica system diagrams, including the overall model configuration, are presented in Appendix A.

3.1.1. Reactor core and primary system

The cross-duct connecting the reactor pressure vessel and steam generator was represented by separate hot and cold ducts, while the reactor vessel was divided into the core, annular riser, lower plenum, and upper plenum. The fluid components were primarily constructed using the gas *Header* and *Flow1DFV* models of the ThermoPower library, which provide lumped-volume and one-dimensional finite-volume representations, respectively. The major modeling features are summarized below.

- *Reactor core*: Although the reference core consists of multiple graphite blocks containing fuel and coolant holes, its thermal-hydraulic behavior was represented using two simplified elements corresponding to the average and hot channels. The hexagonal fuel-block geometry surrounding each coolant hole was simplified to an axisymmetric annular representation following the modeling approach of [20]. Lateral and axial bypass flows were neglected in the present model.

  Figure A.2 illustrates the simplified representations of the average and hot channels. The average coolant hole was thermally coupled to a single surrounding fuel-graphite region. For the hot channel, an additional graphite layer and an external convective boundary were introduced to approximately represent lateral heat transfer to the surrounding core. Treating the hot channel as a thermally isolated element resulted in excessive heat accumulation and unrealistically high fuel temperatures during transients. The external heat transfer coefficient was fixed at the minimum value encountered over the simulated operating range (430 W/m^2K), while the boundary temperature was set to the nominal core outlet temperature. These treatments provide an approximate representation of lateral heat removal without explicitly resolving neighboring channels. The corresponding OpenModelica representation of the reactor core is shown in Figure A.3.

Radial heat conduction through the fuel-graphite region was modeled using a newly developed *RadialConduction1D* component based on the *CylinderFourier* model of the ThermoPower library. Radial conduction is solved independently at each axial node, while axial heat conduction in the solid region is neglected.

Fuel heat generation was calculated using a *ReactorPower* model based on the point kinetics equations, including fuel- and moderator-temperature feedback, external reactivity insertion, and Xenon dynamics. The calculated reactor power was distributed to the fuel region as a volumetric heat source.

- *Helium blower*: A fan-type component was used to represent the primary helium blower because the required pressure rise is small relative to the absolute primary-system pressure. The operating point is determined from the prescribed head-flow characteristic curve and the hydraulic resistance of the primary loop. Because a detailed design-specific blower characteristic was unavailable, a quadratic curve was prescribed using the nominal operating point together with assumed shutoff and high-flow conditions.

3.1.2. Steam generator

The helical-tube bundle was represented using equivalent one-dimensional flow channels, with a single *Flow1DFV* component for the water/steam side and another for the helium shell side, thermally coupled through the *MetalTubeFV* model. The steam generator length was divided into three sections corresponding to the economizer, evaporator, and superheater regions. The major modeling features are summarized below.

- *Shell side*: Considering the crossflow of helium over the tube bundle, the Zukauskas correlation was implemented to calculate the shell-side heat transfer coefficient [31]. The shell-side pressure drop was represented using an equivalent hydraulic resistance coefficient calibrated to reproduce the pressure drop obtained from the corresponding design correlation.

- *Tube side*: Heat transfer enhancement due to the helical tube geometry and detailed boiling heat-transfer correlations were not explicitly considered in the present model. This simplification was adopted because the helium-side heat transfer coefficient is approximately one order of magnitude lower than that on the water/steam side. Two-phase pressure-drop effects were represented using an equivalent correction factor applied to the single-phase hydraulic resistance.

3.1.3. Power conversion system

The power conversion system includes the turbines, pumps, condenser, deaerator, and control valves. The components were primarily based on models available in the ThermoPower library, with several modifications to represent off-design and transient behavior. The major modeling features are summarized below.

- *Turbine*: The turbine model was based on the *SteamTurbineStodola* component, which applies Stodola's cone law to relate the turbine mass flow rate to the inlet and outlet pressures [32]. The turbine was divided into high- and

low-pressure sections to accommodate intermediate steam extraction. The Stodola coefficients were determined from the nominal operating conditions, and a nominal isentropic efficiency of 0.85 was assumed for both turbine sections. Variations in isentropic efficiency under off-design conditions were represented by corrections for steam wetness and partial-load operation. The wetness correction was evaluated using Baumann's rule [33]:

$$F_{wet} = 1 - C_B(1 - \chi_m) \tag{4}$$

where χ_m is the average steam quality and C_B is the Baumann coefficient, taken as unity. The partial-load correction was adopted following [15]:

$$F_{part} = 1 - \alpha\left(1 - \frac{\dot{m}}{\dot{m}^n}\right)^{\xi} \tag{5}$$

where $\dot{m}$ and $\dot{m}^n$ are the actual and nominal mass flow rates, respectively, and $\alpha = 0.14$ and $\xi = 1.4$. The transient turbine efficiency was then evaluated as:

$$\eta_{iso} = \eta^n_{iso} F_{wet} F_{part} \tag{6}$$

- *Pump*: The condensate and feedwater pumps were represented using the *Pump* model of the ThermoPower library. Similar to the helium blower, quadratic head-flow characteristics were prescribed using the nominal operating point together with assumed shutoff and high-flow conditions.

- *Condenser and deaerator*: The condenser was based on the ThermoPower *Condenser* model, with the cooling side modified to represent air cooling. The deaerator was modeled as a well-mixed two-phase control volume. Detailed degassing and chemical processes were neglected, and thermodynamic equilibrium between the liquid and vapor phases was assumed. The governing mass and energy balances are as follows:

$$\begin{aligned} \frac{dM}{dt} &= \dot{m}^{steam}_{in} + \dot{m}^{water}_{in} - \dot{m}_{out} \\ \frac{dE}{dt} &= \dot{m}^{steam}_{in}\, h^{steam}_{in} + \dot{m}^{water}_{in}\, h^{water}_{in} - \dot{m}_{out}\, h^{sat}_{out} \end{aligned} \tag{7}$$

  where M and E denote the total mass and energy stored in the deaerator, respectively. The outlet water was assumed to be saturated at the deaerator pressure.

- *Valves*: The turbine throttle and steam-extraction valves were modeled using the *ValveVap* component of the ThermoPower library. The throttle-valve opening was varied by the grid-frequency controller to regulate turbine steam flow, whereas the extraction valve was operated at a fixed opening selected to provide the nominal extraction condition.

3.1.4. Generator and grid-frequency model

The generator and local grid were represented using the *Generator* and *Grid* components of the ThermoPower library. Because the present study focuses on a tightly coupled nuclear-DC system, the external transmission network

was not explicitly modeled. Instead, an equivalent local-grid representation was used to capture the dynamic coupling between generator power, DC load demand, and grid frequency.

The turbine generator was represented using a rotational energy balance. The rotor dynamics and electrical frequency are described as:

$$\begin{aligned} J\omega_m \frac{d\omega_m}{dt} &= P_m - P_d - \frac{P_e}{\eta_g} \\ f &= \frac{\omega_e}{2\pi} \end{aligned} \tag{8}$$

where J is the equivalent rotational inertia, P_m and P_e are the turbine mechanical power and generator electrical power, respectively, η_g is the generator efficiency, and P_d represents equivalent damping of short-term variations. The default parameter values provided by the ThermoPower model were retained in this study: an acceleration time constant of 10 s for the turbine-generator group, a damping coefficient of 0.01, and a generator efficiency of unity.

The local grid was represented using a finite-droop relationship between frequency and net power imbalance:

$$\begin{aligned} f &= f_0 \left(1 + \alpha \frac{\Delta P}{P_{Grid}}\right) \\ \Delta P &= P_e - P_L \end{aligned} \tag{9}$$

where f_0 is the nominal grid frequency, P_L is the DC load demand, P_{Grid} is the reference grid power, and α is the grid droop coefficient. P_{Grid} was set equal to the nominal plant power to represent a tightly coupled local-grid configuration with a droop coefficient of 0.05. Accordingly, an increase in DC demand relative to generator output produces a negative power imbalance and a corresponding decrease in grid frequency. The resulting frequency deviation is supplied to the turbine throttle-valve controller described in Section 2.3.1.

This simplified representation is intended to capture the plant-side electro-mechanical response to load imbalance rather than to evaluate transmission-network dynamics or grid-code compliance.

3.2. Steady-state consistency and controller tuning

3.2.1. Comparison with nominal design values

To confirm the consistency between the reference plant design and the OpenModelica model, the nominal design values were compared with the steady-state simulation results, as summarized in Table 3. The largest deviation is observed in the low-pressure turbine (LPTB) power, with a relative difference of 6.8%. This discrepancy is mainly attributed to the reduction in isentropic efficiency due to steam wetness, which is accounted for in the dynamic turbine model but not in the nominal design calculation. The relative differences of the remaining parameters are within approximately 4%, indicating that the OpenModelica model adequately reproduces the nominal operating conditions of the reference plant.

Table 3: Comparison of nominal design values and OpenModelica steady-state results.

Parameter	Nominal design	OpenModelica	Relative error [%]
Reactor outlet helium pressure [MPa]	6	6.17	2.82
Helium mass flow rate [kg/s]	81.6	84.51	3.57
Steam mass flow rate [kg/s]	69.1	68.99	0.16
HPTB inlet pressure [MPa]	13.8	13.80	0.01
HPTB outlet pressure [MPa]	1.57	1.58	0.35
Deaerator pressure [MPa]	1.43	1.44	0.45
Condenser pressure [kPa]	20	19.85	0.75
HPTB power [MWe]	35.1	34.28	2.34
LPTB power [MWe]	33.5	31.22	6.80
Condenser heat rejection [MWth]	112.2	114.13	1.72
Steam extraction ratio [%]	22	21.61	1.76

3.2.2. Controller tuning

After establishing the steady-state consistency of the dynamic model, a symmetric linear ramp-down and ramp-up transient was used as the reference case for controller tuning. A prescribed ramp was selected to avoid tuning the controllers to a DC-specific load profile. The system was initially maintained at the nominal operating condition for 100 s, after which the load demand was reduced from 100% to 15% of the rated power at a rate of 5%/min. The 15% load level corresponds to the minimum demand encountered in the DC load profile considered in this study. The load was maintained at 15% for 20 min and subsequently increased back to 100% at the same ramp rate.

Although systematic controller-optimization techniques are available, controller optimization itself is outside the scope of the present study. Therefore, the PI and SMC parameters were tuned manually through iterative transient simulations. For both controller types, the parameters were adjusted to achieve stable load tracking while satisfying the transient operating constraints summarized in Table 2.

The resulting controller parameters and system responses to the reference tuning transient are presented in Appendix B. Both PI and SMC provide stable and comparable responses throughout the ramp-down and ramp-up processes. The turbine electrical power closely follows the prescribed load demand, resulting in limited grid-frequency deviations, whereas the reactor thermal power responds more gradually because of the thermal inertia and feedback mechanisms of the primary system. The steam pressure and temperature, core outlet helium temperature, maximum fuel temperature, and reactivity insertion rate remain within the prescribed operating constraints for both controllers.

The steam conditions shown in Figure B.1 correspond to the steam generator outlet; therefore, the nominal pressure is approximately 500 kPa higher than the main steam pressure of 13.8 MPa, which is defined at the high-pressure

turbine inlet after the throttle-valve pressure drop. Overall, the relatively small differences between the PI and SMC responses indicate that both control configurations provide an adequate basis for the subsequent dynamic analysis. The tuned controller parameters were fixed for all subsequent simulations.

3.3. Sinusoidal load perturbation analysis

To systematically characterize the frequency-dependent dynamic responses of the system, sinusoidal variations in electrical load demand were imposed over a range of amplitudes and periods. The system was first stabilized at 50% of the nominal power, after which the load demand was varied according to:

$$P_L(t) = P_r\left(0.5 + A\sin\left(\frac{2\pi t}{T}\right)\right) \tag{10}$$

where P_r is the nominal power, A is the normalized load amplitude, and T is the oscillation period.

The amplitudes considered in the present analysis ranged from 2% to 20%, while the oscillation periods ranged from 0.2 to 300 s. For each amplitude-period combination, 20 complete oscillation cycles were simulated. To minimize the influence of the initial transient, only the final three cycles were used for the subsequent analysis. The resulting responses were compared over the amplitude-period domain to identify dynamically sensitive time scales and their dependence on perturbation magnitude.

3.4. Data center load profiles

Two facility-level DC demand profiles with distinct temporal characteristics were considered. The first is a year-long facility-aggregated demand profile representing slow variations in facility power consumption over extended periods of DC operation (facility-aggregated load profile), while the second is a reconstructed profile that retains faster workload-driven fluctuations suppressed by temporal and workload aggregation (workload-resolved load profile).

3.4.1. Facility-aggregated load profile

The facility-aggregated load profile was adopted directly from the facility-level DC demand reconstructed by Vercellino et al. [18]. In their study, measured GPU power traces of representative AI workloads were combined with job scheduling and utilization distributions derived from operational data from Shanghai AI Lab data center. The resulting year-long facility demand profile has a temporal resolution of 60 s and represents the highly aggregated electrical demand of a DC facility. Because a large number of individual workloads are aggregated in the process, short-term fluctuations are substantially smoothed, whereas longer-term variations associated with DC operation remain in the aggregated load profile. Among the operating scenarios reported by the study [18], the training-workload case with 60% utilization was selected for the present analysis.

Since the scale of the original profile is different from that of the current reference plant considered in this study, the profile was normalized by its observed peak demand and subsequently scaled so that the peak corresponds to the nominal electrical power of the reference plant. In the present study, this profile was used to examine sustained

load-following capability over an extended operating period, including repeated variations between approximately 15% and 100% of the nominal power. The effects of the short-timescale fluctuations suppressed by this aggregation are examined separately using the workload-resolved profile described in the following section.

3.4.2. Reconstruction of the workload-resolved load profile

The aggregated load profile is representative of long-term facility-level demand, but its extensive workload aggregation and 60-s temporal resolution smooth much of the faster variability. The short-timescale characteristics of hyperscale DC demand, however, remain uncertain because they depend strongly on workload composition, synchronization, scheduling, and aggregation level. To examine the plant response under more demanding short-timescale variations, an additional workload-resolved load profile was reconstructed.

The fast component was generated using a semi-Markov model, following the general stochastic load-modeling framework of Chaudhary et al. [34]. A semi-Markov model represents transitions among discrete operating states while explicitly accounting for the residence-time distribution of each state in addition to the transition probabilities between states. This allows nonuniform state durations to be represented without imposing a fixed switching time.

Specifically, the stochastic-model parameters were extracted directly from measured GPU power traces for llama 2 training provided by Vercellino et al. [18]. Twenty-one measured power traces were individually normalized and classified into three statistically distinct power states using a Gaussian mixture model (GMM). These three states can be broadly interpreted as corresponding to the dominant power states of an AI training iteration: Surge, Dip, and Base [34].

For each measured trace, the state-dependent mean and standard deviation, state-transition probabilities, and empirical residence-time distributions were determined. The resulting quantities were then averaged over the 21 measurement cases to construct a representative stochastic model. All three states were retained when extracting these statistical parameters. During signal synthesis, however, the Base state was excluded to represent sustained, high-utilization DC operation, and the transition probabilities among the remaining states were re-normalized accordingly.

Within each state, short-timescale fluctuations around the state mean were represented using an Ornstein-Uhlenbeck (OU) process. Unlike uncorrelated Gaussian noise, the OU process exhibits mean-reverting behavior and therefore introduces temporally correlated fluctuations:

$$d\xi = -\theta\,\xi\,dt + \sigma_s\,dW_t \tag{11}$$

where θ is the mean-reversion coefficient, σ_s is the state-dependent noise intensity, and W_t is a Wiener process. The value of θ was set to 0.8 s^{-1}, and a synthetic workload trace was generated with a time step of 0.2 s.

Before applying the generated workload fluctuations to the facility-level demand profile, the stochastic reconstruction model was first examined independently using the measured llama 2-training signal. Figure 3 compares the measured signal and a synthetic workload trace in the time and frequency domains. A representative 100-s interval is shown in the time domain for clarity. The measured profile exhibits repeated transitions between the Dip and

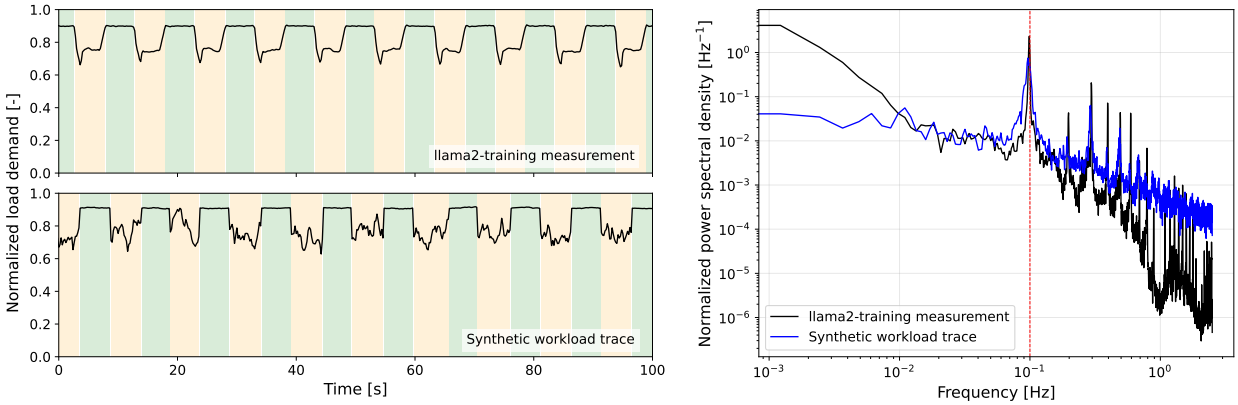

Figure 3: Comparison of the llama2-training measurement and synthetic workload trace in the time (left) and frequency (right) domains. Green and yellow shading denotes the Surge and Dip states, respectively, and the red dashed line indicates a frequency of 0.1 Hz (10 s period).

Surge states with a characteristic timescale of approximately 10 s. This behavior is more clearly identified in the power spectral density (PSD), where a pronounced peak appears near the corresponding frequency. The synthetic workload trace reproduces the alternating state-transition behavior and the dominant spectral peak near 0.1 Hz, although differences remain in the low-frequency content and finer spectral structure. The comparison therefore indicates that the reconstruction process captures the principal short-timescale characteristics of the measured workload relevant to the subsequent analysis.

Based on this workload reconstruction model, N independently generated traces were then averaged to represent workload aggregation at the facility level. The averaged signal was then converted into a zero-mean relative fluctuation. Meanwhile, the 60-s aggregated load profile was linearly interpolated to the 0.2-s reconstruction time step, and a one-hour high-variability interval was selected for the transient analysis. Specifically, the interval exhibiting the largest cumulative change in the load gradient was chosen. The workload-driven fluctuation was subsequently superimposed on the selected interval as:

$$\delta(t) = \frac{P_{\text{fast}}(t) - \overline{P}_{\text{fast}}}{\overline{P}_{\text{fast}}}$$
$$P_{\text{WR}}(t) = P_{\text{agg}}(t)\,[1 + \delta(t)] \tag{12}$$

This reconstruction preserves the slowly varying facility-level demand while introducing fast workload-driven fluctuations whose statistical characteristics are derived from measured GPU power data. Figure 4 summarizes the overall reconstruction procedure.

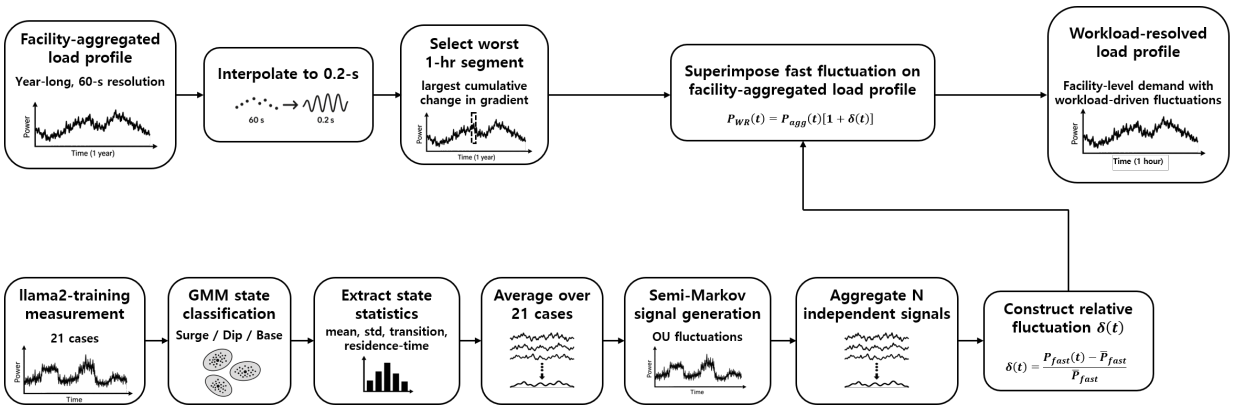

Figure 4: Flow chart of workload-resolved load reconstruction.

4. Results

4.1. Sinusoidal load perturbation analysis

The dynamic response of the system under sinusoidal load perturbations was investigated to characterize its frequency-dependent behavior. Figure 5 compares the turbine mechanical and electrical power responses for selected perturbation periods under PI control. The tracking error is defined as $\frac{P(t)-P_L(t)}{0.5\,P_r} \times 100$, where P denotes either the turbine mechanical or electrical power, P_L is the instantaneous load demand, and P_r is the rated electrical power. The denominator $0.5\,P_r$ corresponds to the initial operating power of the sinusoidal perturbation cases. The mechanical-power tracking error increases as the perturbation period decreases, whereas the electrical-power tracking error reaches a maximum at intermediate periods. This difference arises from the electro-mechanical buffering provided by the turbine-generator inertia. At shorter periods, the thermal-hydraulic plant dynamics cannot respond rapidly enough to the imposed load variation. Consequently, turbine mechanical power changes only slightly over each cycle, resulting in a large tracking error relative to the rapidly varying demand. For the electrical power, in contrast, successive positive and negative power imbalances reverse before substantial changes in rotor kinetic energy can accumulate. The rapid fluctuations are therefore effectively attenuated by the generator-grid dynamics before a significant electrical response develops. As the perturbation period increases into the intermediate range, the disturbance persists long enough for this inertial buffering to become less effective, while still remaining too fast for the control system to fully compensate. Consequently, the electrical-power tracking error reaches its largest values in this region. At still longer periods, the plant controllers have sufficient time to regulate the system, and both tracking errors decrease.

A similar frequency-dependent trend is observed in the other major plant variables, as shown in Figures 6 and 7. These figures summarize the maximum absolute or relative deviations from the nominal operating points for the major plant variables. The turbine mechanical- and electrical-power results are reported using the tracking errors defined above. The grid-frequency response should be interpreted together with the electrical-power tracking error shown in Figure 5. As defined by Equation 9, the simplified grid model relates the grid-frequency deviation directly to the

generator–load power imbalance, exhibiting similar trends between them. Despite differences in the detailed responses, both control configurations exhibit a similar region of elevated sensitivity, indicating that the observed behavior is primarily associated with the characteristic time scales of the modeled system rather than with a specific controller. The magnitude of the deviations generally increases with perturbation amplitude, whereas the location of the sensitive period range remains largely unchanged.

Based on these responses, three distinct response regions can be identified in the frequency domain: (1) a high-frequency region with characteristic periods up to a few seconds, where the plant response is strongly attenuated; (2) an intermediate region with characteristic periods of several tens of seconds, where the largest deviations generally occur; and (3) a low-frequency region with characteristic periods on the order of minutes, where effective regulation is recovered. These results indicate that the ability of the nuclear system to accommodate DC load variations depends not only on the fluctuation magnitude but also strongly on its characteristic time scale.

In the SMC case, the turbine-inlet steam variables exhibit a plateau or a slight increase in deviation in the low-frequency region. Additional simulations showed that this behavior disappeared when the helium-blower controller was replaced by the PI controller. This suggests that the low-frequency feature is associated with the specific SMC implementation and tuning adopted for the helium-flow control loop, rather than with the general frequency-dependent characteristics of the plant. Importantly, the location of the intermediate sensitive region remains broadly consistent between the PI and SMC cases.

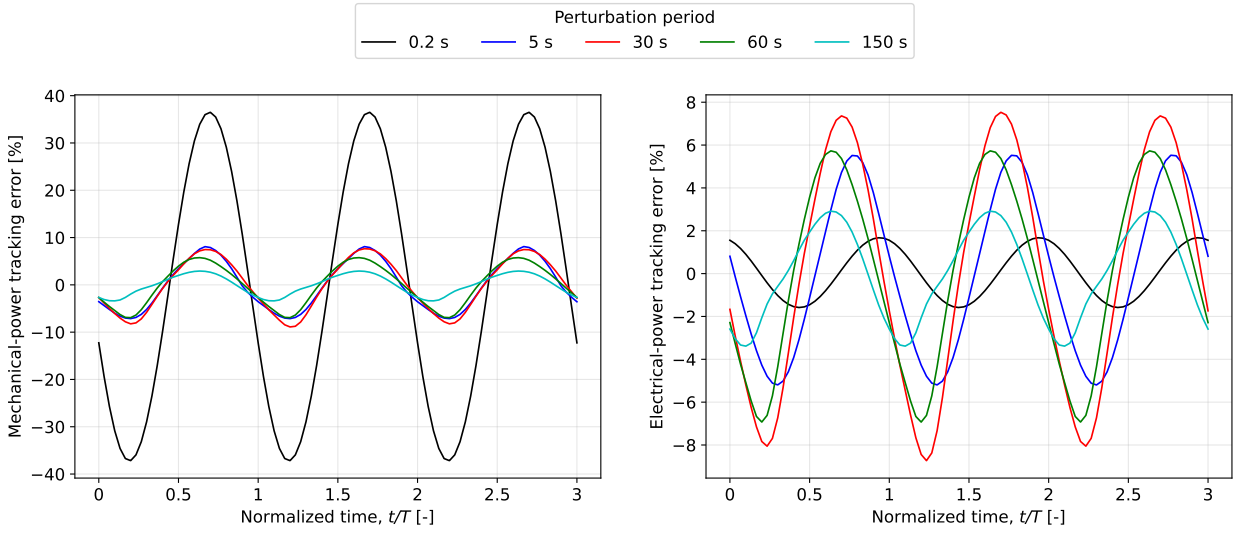

Figure 5: Turbine mechanical (left) and electrical (right) power responses for selected perturbation periods at a 20% load amplitude under PI control. Last three cycles are shown.

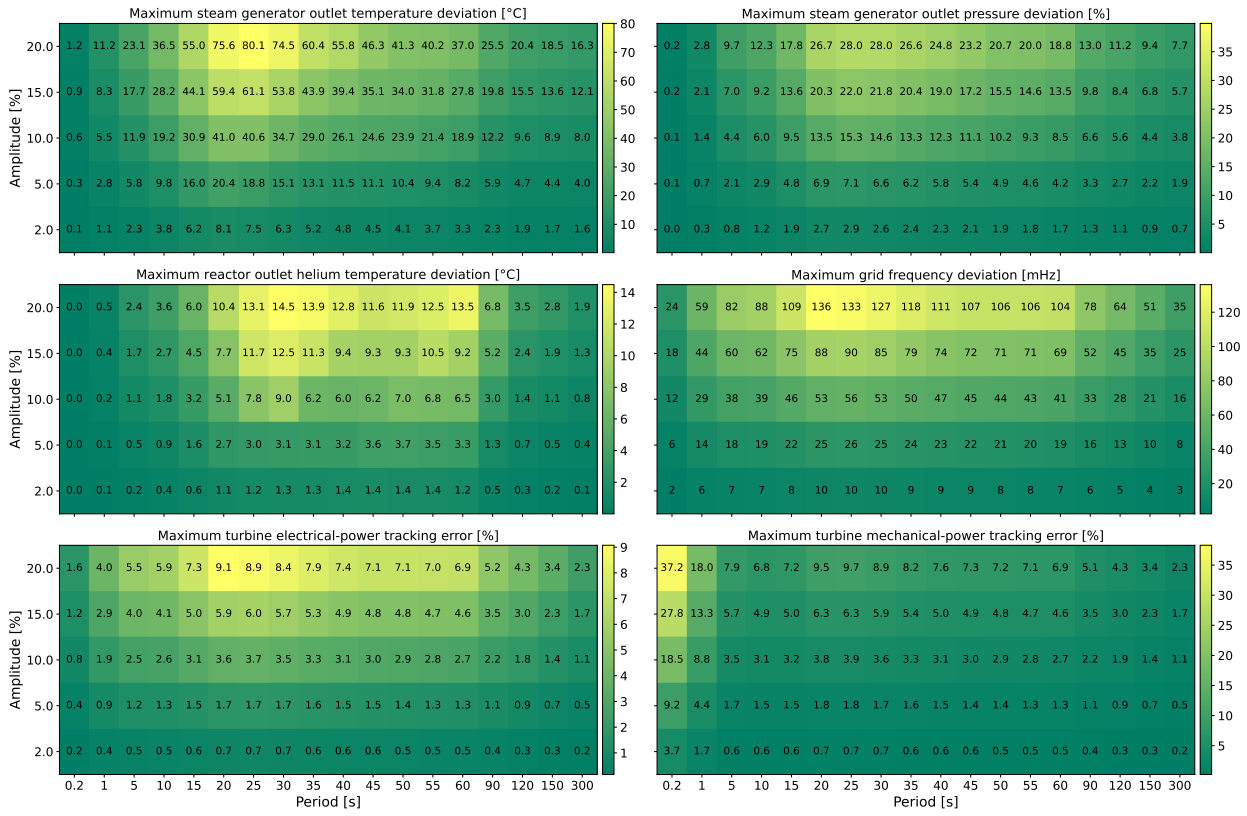

Figure 6: Amplitude-period response map under PI control. Each cell represents the maximum absolute or relative responses. The color scale for each variable is shared between the PI and SMC cases using the combined minimum and maximum values.

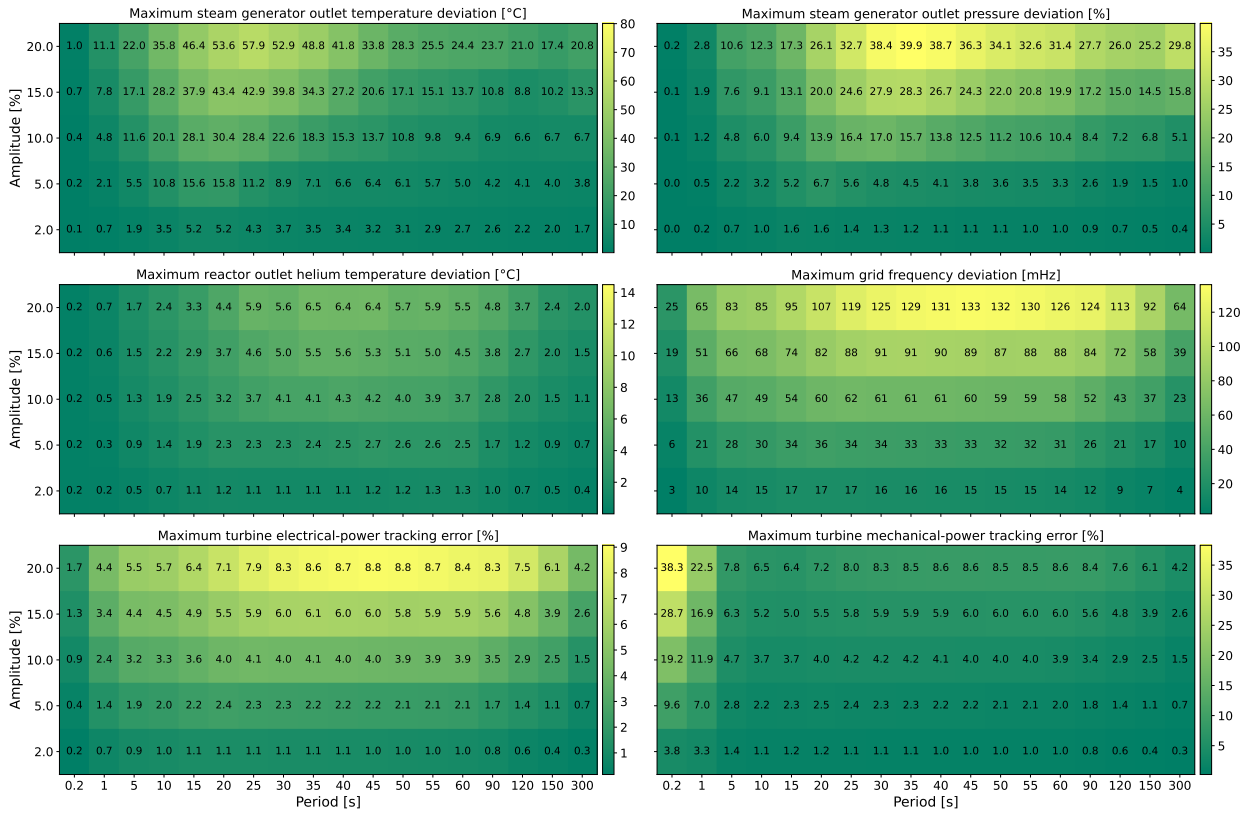

Figure 7: Amplitude-period response map under SMC control. Each cell represents the maximum absolute or relative responses. The color scale for each variable is shared between the PI and SMC cases using the combined minimum and maximum values.

4.2. Dynamics under the facility-aggregated load profile

The plant response under the facility-aggregated load profile was investigated to evaluate sustained load-following capability under long-term and repeated demand variations. Figure 8 shows the aggregated load demand in the time and frequency domains. The demand varies from approximately 15% to 100% of the nominal power, with a maximum ramp rate of around 3%/min. Despite this wide operating range, the variations occur predominantly over relatively long time scales, as also illustrated by the enlarged one-hour profile. Because the original data have a temporal resolution of 60 s, the highest resolvable frequency is the Nyquist frequency of 0.0083 Hz, corresponding to a period of 120 s. The facility-aggregated profile therefore primarily represents the slowly varying component of facility-level DC demand.

Figure 9 summarizes the responses of the major system variables under the aggregated load profile using the monthly maximum, minimum, and average values for both PI and SMC control. Despite repeated operation over a wide power range, the thermal and electrical responses remain relatively limited throughout the year. The monthly average values stay close to their nominal operating points, while the corresponding maximum values satisfy the prescribed transient operating constraints summarized in Table 2. The maximum steam pressure does not exceed 105% of its rated value, and the maximum steam-temperature increase stays within the prescribed +8°C limit. The maximum fuel temperature and reactor outlet helium temperature are also well below their respective prescribed limits. The required reactivity insertion rate is substantially lower than the 24 pcm/s limit, while the condenser water level is also well controlled. Overall, the turbine electrical power tracks the load demand well, resulting in relatively small grid-frequency deviations.

These results are consistent with the sinusoidal perturbation analysis, showing that large-amplitude load variations can be effectively accommodated when their dominant time scales are sufficiently long. Accordingly, the reference nuclear system can sustain long-term facility demand over a wide operating range when short-term workload fluctuations are sufficiently mitigated through workload aggregation and scheduling.

This assessment, however, addresses transient operating performance only; cumulative effects associated with repeated thermal cycling, such as thermal fatigue and material degradation, were not evaluated. Previous studies on HTGR systems have assessed such effects by using transient temperature and pressure histories as boundary conditions for thermo-mechanical analysis [35]. A similar methodology could be applied to the present system to quantify cumulative damage and lifetime impacts under repeated DC load-following operation.

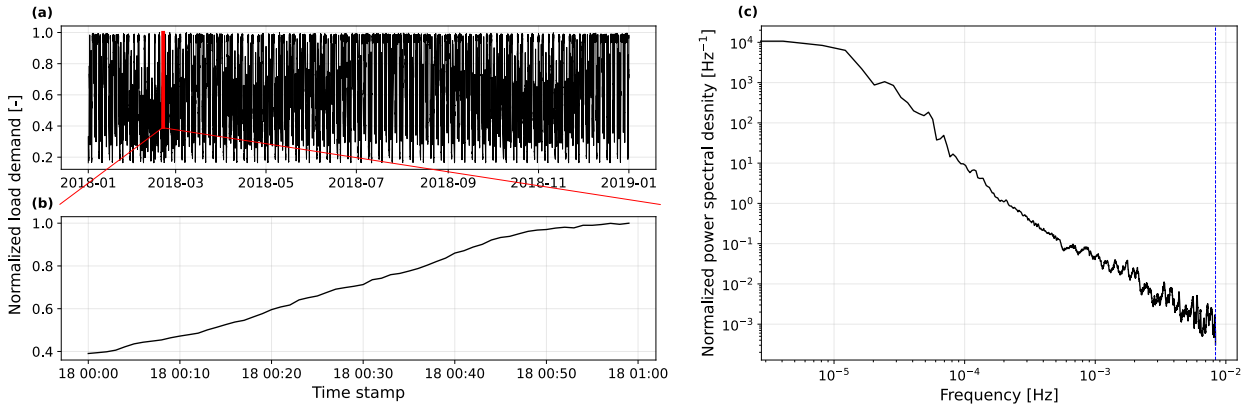

Figure 8: Facility-aggregated load profile in the time and frequency domains. The year-long profile is shown in (a), while a representative one-hour interval on 18 February is shown in (b) for clarity. The blue dashed line indicates the Nyquist frequency of 1/120 Hz.

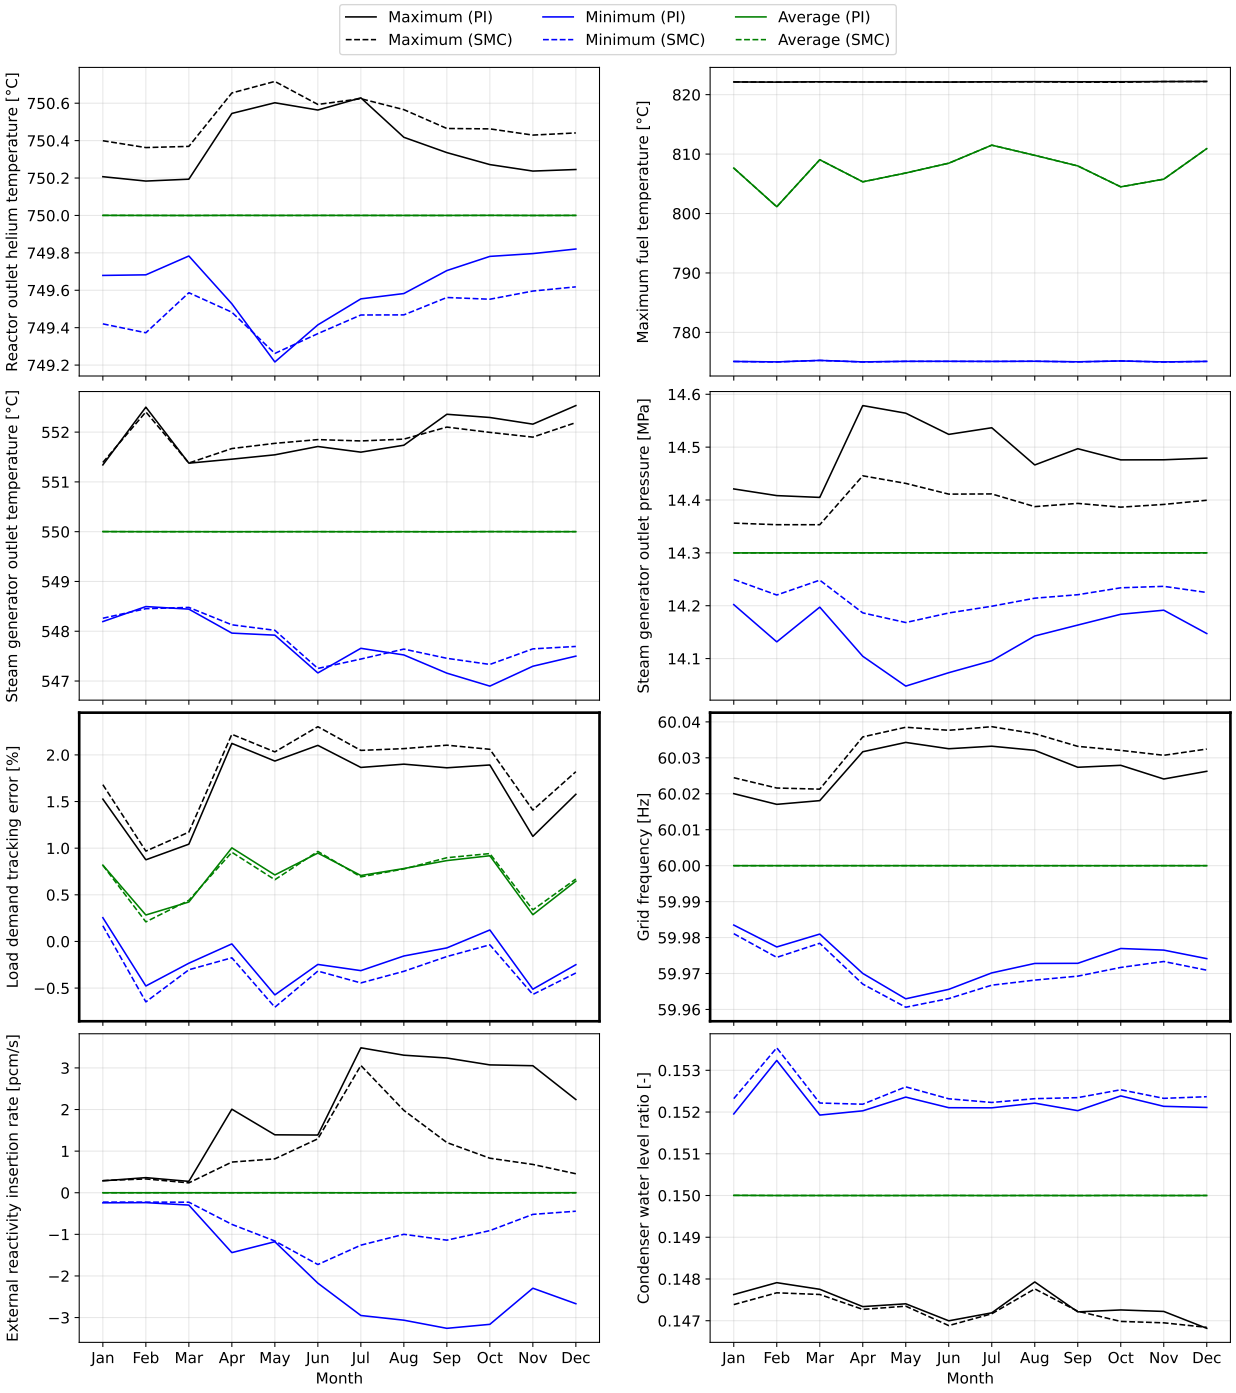

Figure 9: Monthly extrema and averages of the major system variables under the facility-aggregated load profile for PI and SMC control. The thick-bordered panels highlight the load-demand tracking error and the resulting grid-frequency deviation.

4.3. Dynamics under the workload-resolved load profile

4.3.1. Effect of workload aggregation on load profiles

The workload-resolved load profiles were reconstructed using the method described in Section 3.4.2. Figure 10 compares the reconstructed profiles in the time and frequency domains for different aggregation factors, N. For clarity, the time-domain plot shows only a 5-min interval from the 1-h profile used in the transient simulations.

For $N = 1$, the characteristic Dip-Surge transition pattern inherited from the synthetic workload trace is clearly visible, together with the corresponding spectral peak near 0.1 Hz. As N increases, averaging multiple independently generated traces progressively reduces the fluctuation amplitude and weakens the distinct semi-Markov transition pattern. Nevertheless, the characteristic frequency associated with the Dip-Surge transitions remains largely unchanged, indicating that aggregation primarily attenuates the magnitude of the workload-driven fluctuations rather than shifting their dominant frequency. The facility-aggregated load profile, in contrast, contains little spectral energy at these shorter time scales because of its 60-s temporal resolution and extensive workload aggregation.

It should be noted that N should not be interpreted directly as the physical number of workloads or computing jobs operating in the facility. In the present reconstruction, N controls the degree of statistical aggregation among independently generated traces, and therefore simultaneously reduces the fluctuation amplitude and weakens the temporal structure associated with an individual semi-Markov process. Accordingly, N is treated as an effective aggregation parameter rather than a direct representation of the number of active workloads. For independently generated traces, the standard deviation of the averaged fluctuation decreases approximately in proportion to $N^{-1/2}$, consistent with the progressive reduction in fluctuation magnitude observed in Figure 10. For reference, the stochastic load fluctuations can be expressed as an equivalent sinusoidal amplitude based on their standard deviation, $A_{\mathrm{eq}} = \sqrt{2}\,\mathrm{std}(\delta)$. The resulting values are approximately 14.5%, 6.1%, and 3.0% for $N = 1$, 5, and 25, respectively.

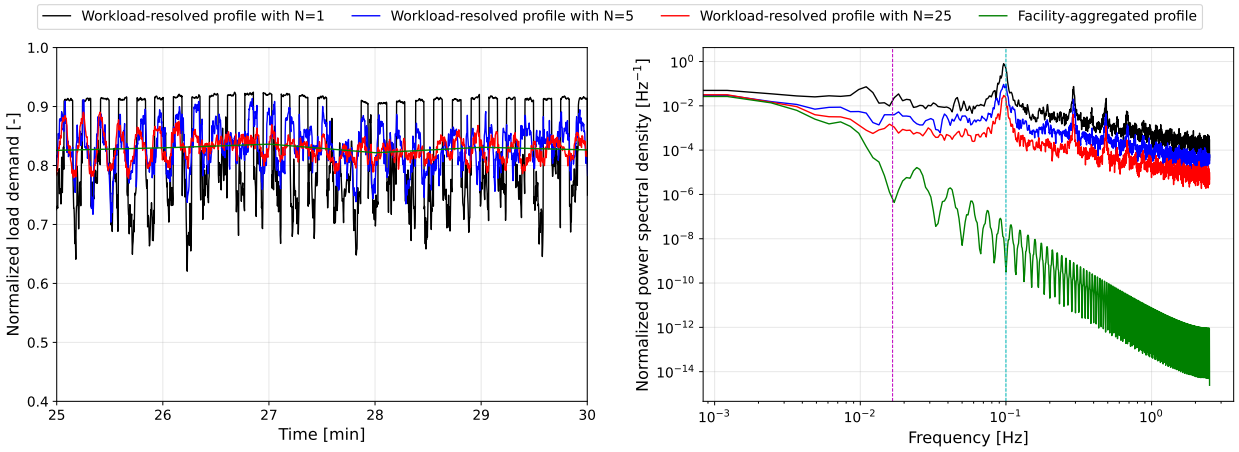

Figure 10: Workload-resolved load profiles in the time (left) and frequency (right) domains with different aggregation factors. Cyan and purple dashed lines indicate the frequencies corresponding to a 10-s dominant workload timescale and the 60-s temporal resolution of the facility-aggregated profile, respectively.

4.3.2. Plant responses to workload-resolved load profiles

Table 4 summarizes the plant responses to the workload-resolved load profiles for different aggregation factors under both PI and SMC control. Responses of the major system variables are characterized using four metrics: maximum, 99th and 95th percentiles, and root-mean-square (RMS) values. Deviations from nominal operating points are reported for steam-generator-outlet temperature and pressure, core outlet helium temperature, load-tracking error, and grid frequency, while absolute values are presented for external reactivity insertion rate and maximum fuel temperature. Steam-temperature and pressure deviations are reported as signed quantities because their operating criteria are upper limits, whereas the remaining deviation-based metrics are evaluated using their absolute magnitudes.

The maximum values present the most severe excursions observed during each simulation, whereas the 99th and 95th percentiles characterize more persistent large responses and the RMS values quantify the overall response magnitude. Across these metrics, the plant responses consistently decrease as the aggregation factor increases.

The operational limits provide a useful basis for interpreting the effect of workload aggregation. At $N = 1$, the large short-timescale fluctuations result in substantial exceedance of the prescribed steam-temperature, steam-pressure, and reactivity-insertion-rate limits, together with large demand-tracking errors and grid-frequency deviations. Increasing N progressively reduces these excursions; however, the steam temperature and reactivity insertion rate remain comparatively restrictive, with their maximum responses still approaching or slightly exceeding the prescribed limits even at $N = 25$. In contrast, the reactor-side thermal variables, including the core outlet helium temperature and maximum fuel temperature, show substantially smaller deviations and retain large margins to their limits, due to the combined effects of the reactor thermal inertia and active reactivity control.

Based on these overall response trends, $N = 5$ represents a transitional case between severe constraint exceedance at low aggregation and well-attenuated behavior at higher aggregation. This case was therefore selected as the representative workload-resolved profile for the subsequent frequency-band analysis.

Table 4: System responses under the workload-resolved profiles with different aggregation factors. Bold numbers indicate cases that exceed the prescribed transient operating constraints.

Parameter	Metric	PI control			SMC control		
		N=1	N=5	N=25	N=1	N=5	N=25
Steam temperature [°C]	max	**26.5**	**17.8**	**8.0**	**27.2**	**18.3**	7.6
	p99	**20.3**	**9.9**	5.6	**20.1**	**9.3**	5.2
	p95	**16.1**	7.2	3.7	**15.3**	6.2	3.1
	rms	**10.1**	4.3	2.3	**9.0**	3.6	1.8
Steam pressure [%]	max	**14.0**	**8.5**	3.4	**21.7**	**9.6**	3.1
	p99	**9.2**	4.2	2.2	**14.4**	4.1	1.7
	p95	**6.9**	3.0	1.5	**8.3**	2.7	1.0
	rms	3.9	1.7	0.9	4.3	1.4	0.6
Helium temperature [°C]	max	4.1	2.4	1.0	3.1	1.9	1.0
	p99	2.8	1.3	0.7	2.1	1.2	0.8
	p95	2.1	1.0	0.5	1.6	1.0	0.6
	rms	1.1	0.5	0.3	0.9	0.5	0.3
Load demand tracking error [%]	max	7.4	3.1	1.6	5.0	3.1	2.0
	p99	3.8	2.1	1.1	3.5	2.2	1.4
	p95	3.2	1.5	0.8	3.1	1.7	1.1
	rms	2.0	0.8	0.4	1.9	0.9	0.5
Grid frequency [mHz]	max	**221.6**	92.7	48.3	**149.7**	92.1	59.3
	p99	**113.7**	62.3	32.6	**105.9**	66.2	40.9
	p95	94.6	46.4	24.2	92.5	52.4	31.7
	rms	60.3	24.3	12.1	55.6	26.7	15.6
Reactivity insertion rate [pcm/s]	max	**98.9**	**50.5**	**24.3**	**147.7**	**69.8**	**29.0**
	p99	**63.0**	**30.9**	16.4	**68.4**	**32.3**	14.9
	p95	**48.1**	**24.0**	12.5	**44.5**	21.4	10.1
	rms	**28.3**	12.4	6.2	23.2	10.3	4.4
Maximum fuel temperature [°C]	max	821.6	819.0	818.6	820.7	818.6	818.1
	p99	819.7	818.3	818.3	818.8	818.1	817.9
	p95	818.7	818.0	818.0	818.1	817.8	817.6
	rms	817.2	817.2	817.4	816.4	817.2	817.3

4.3.3. Frequency-band effects on system dynamics

The sinusoidal perturbation analysis in Section 4.1 identified an intermediate-frequency region where the plant exhibits elevated dynamic sensitivity, whereas substantially faster perturbations produced smaller deviations. To examine whether this frequency dependence is also observed under stochastic DC demand, selected frequency components were removed from the reconstructed workload-resolved profile and the resulting plant responses were compared.

Two frequency ranges were considered. A high-pass filter with a cutoff frequency of 1/10 Hz was used to extract the high-frequency component with periods shorter than 10 s, while a band-pass filter between 1/60 and 1/10 Hz was used to extract the intermediate-frequency component with periods between 10 and 60 s. The filters were implemented using the *scipy.signal* library. Fourth-order Butterworth high-pass and band-pass filters were constructed in second-order-section (SOS) form and applied using zero-phase forward-backward filtering with *sosfiltfilt*, thereby minimizing phase distortion in the reconstructed load signals. This zero-phase filtering was used only for the present offline frequency-band analysis; real-time implementation would require causal filtering and would therefore introduce phase lag. The extracted components were subtracted from the original zero-mean relative fluctuation, $\delta(t)$, to construct the high-frequency-removed (HFM) and intermediate-frequency-removed (IFM) signals:

$$\begin{aligned} \delta_{high} &= HPF(\delta), \qquad \delta_{band} = BPF(\delta) \\ \delta_{HFM} &= \delta - \delta_{high}, \qquad \delta_{IFM} = \delta - \delta_{band} \end{aligned} \tag{13}$$

where HPF and BPF denote the high-pass and band-pass filtering operations, respectively, and δ_{high} and δ_{band} are the extracted high- and intermediate-frequency components, respectively. The modified fluctuations were subsequently recombined with the facility-aggregated profile using the same procedure described in Section 3.4.2. The case with an aggregation factor of $N = 5$ was selected as the original representative workload-resolved profile for the following analysis.

Because the original load contains different amounts of spectral energy in the two frequency ranges, direct removal of the two bands would also produce different reductions in the overall fluctuation magnitude. Therefore, two normalization approaches were considered. In the first approach, the attenuation of each extracted component was adjusted so that the same amount of spectral energy was removed from the high- and intermediate-frequency bands (energy-matched case):

$$\begin{aligned} \delta_{HFM,E} &= \delta - \alpha_{high}\,\delta_{high} \\ \delta_{IFM,E} &= \delta - \alpha_{band}\,\delta_{band} \end{aligned} \tag{14}$$

where α_{high} and α_{band} were determined such that the reduction in the integrated PSD was identical for the two cases. In the second approach, the selected frequency component was fully removed and the remaining fluctuation was rescaled

to recover the standard deviation of the original signal (standard-deviation-matched case):

$$\begin{aligned} \delta'_{HFM} &= \delta - \delta_{high}, \qquad \delta_{HFM,S} = \delta'_{HFM} \frac{\sigma_\delta}{\sigma_{\delta'_{HFM}}} \\ \delta'_{IFM} &= \delta - \delta_{band}, \qquad \delta_{IFM,S} = \delta'_{IFM} \frac{\sigma_\delta}{\sigma_{\delta'_{IFM}}} \end{aligned} \tag{15}$$

where δ'_{HFM} and δ'_{IFM} are the fluctuations after complete removal of the high- and intermediate-frequency components, respectively, and σ denotes the corresponding standard deviation.

Figures 11 and 12 compare the resulting signals in the time and frequency domains. In the energy-matched cases, the target frequency regions are selectively attenuated while the remaining spectral content is only weakly affected, so the total fluctuation variance is reduced relative to the original signal. In the standard-deviation-matched cases, the loss of variability caused by filtering is compensated by amplifying the remaining frequency components. Consequently, these signals retain the same overall fluctuation intensity as the original profile but increase the spectral energy outside the removed band, which can also amplify signal excursions in time domain.

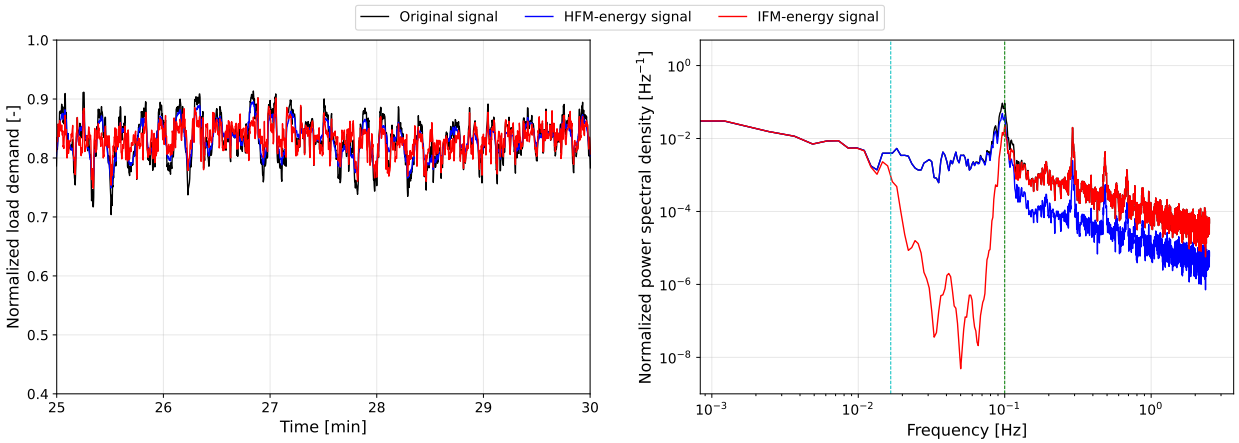

Figure 11: Energy-matched signals in time (left) and frequency (right) domains. Green and cyan dashed lines indicate the frequency-band boundaries at 1/10 and 1/60 Hz, respectively.

Table 5 summarizes the system responses to the different frequency-modified load signals under PI and SMC control. The parameter definitions and response metrics are identical to those used in Table 4. The energy-matched cases compare the mitigation effectiveness of removing different frequency regions while maintaining the same amount of removed fluctuation energy. Across most response metrics, the IFM-energy case produces greater reductions than the HFM-energy case. For example, under PI control, removal of the high-frequency components reduces the 95th-percentile steam-temperature deviation from 7.2 to 5.5°C, whereas removal of the intermediate-frequency components reduces it further to 3.6°C. Similar behavior is observed for the grid-frequency deviation, with the 95th percentile decreasing from 46.4 to 33.6 and 27.5 mHz, respectively. Comparable trends are obtained under SMC control. These results show that, for the same amount of removed fluctuation energy, attenuation of the intermediate-frequency components provides greater mitigation than attenuation of the higher-frequency components. A minor exception is

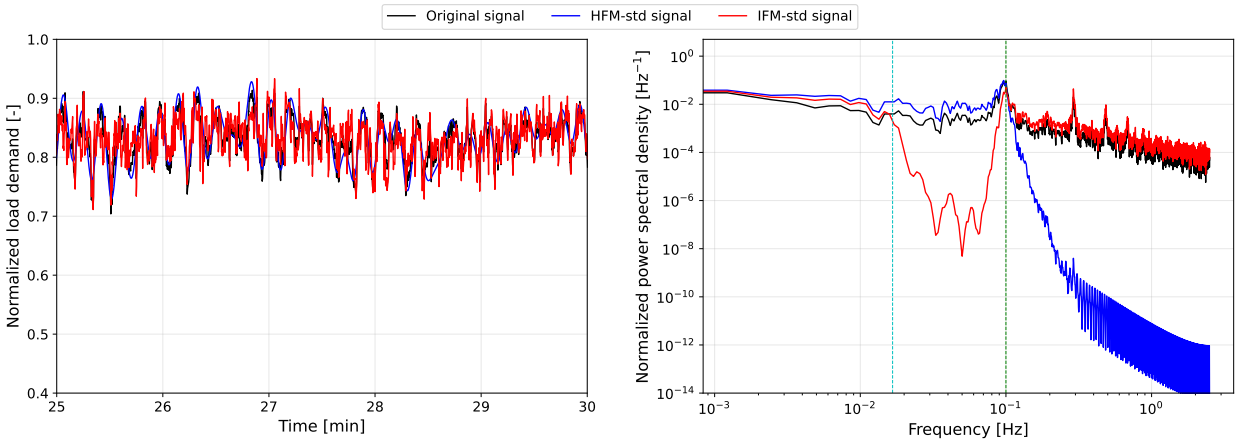

Figure 12: Standard-deviation-matched signals in time (left) and frequency (right) domains. Green and cyan dashed lines indicate the frequency-band boundaries at 1/10 and 1/60 Hz, respectively.

observed in the maximum reactivity insertion rate under PI control, for which the IFM-energy case shows a slightly larger peak than the HFM-energy case. However, the corresponding percentile and RMS metrics retain the overall trend, suggesting that this isolated peak is associated with a controller-specific transient response to the modified load waveform.

The effect of spectral distribution is more clearly demonstrated by the standard-deviation-matched cases, in which the overall fluctuation magnitude is restored after removal of the selected frequency band. Under PI control, HFM-std increases the 95th-percentile steam-temperature deviation from 7.2 to 8.3°C and the grid-frequency deviation from 46.4 to 51.7 mHz, because the remaining intermediate-frequency components are amplified. In contrast, IFM-std reduces these values to 5.2°C and 40.5 mHz, respectively, despite amplification of the remaining higher-frequency components. The same tendency is observed under SMC control. These results confirm that the intermediate-frequency region exerts a greater influence on plant dynamics than the higher-frequency region, demonstrating that the transient burden depends not only on the overall magnitude of load fluctuations but also on their spectral distribution.

These findings suggest that frequency-selective coordination between nuclear power systems and other flexible resources may provide an effective strategy for managing highly dynamic DC demand. In particular, demand components within the dynamically sensitive frequency range could be preferentially handled by faster-response resources, while the nuclear plant follows the remaining component. Such load sharing could be implemented through workload scheduling, energy storage, or support from the external grid.

Table 5: Effects of frequency-band removal on system responses. Values are shown as PI/SMC cases. Bold values highlight the representative results discussed in the main text.

Parameter	Metric	Signal case				
		Original	HFM-energy	IFM-energy	HFM-std	IFM-std
Steam temperature [°C]	max	17.8 / 18.3	15.1 / 14.9	9.6 / 7.9	24.5 / 25.8	14.3 / 12.4
	p99	9.9 / 9.3	7.6 / 6.7	5.1 / 4.3	11.4 / 10.4	7.4 / 6.5
	p95	**7.2** / 6.2	**5.5** / 4.5	**3.6** / 2.9	**8.3** / 6.8	**5.2** / 4.4
	rms	4.3 / 3.6	3.4 / 2.6	2.1 / 1.6	5.2 / 4.1	3.1 / 2.5
Steam pressure [%]	max	8.5 / 9.6	7.4 / 8.3	3.6 / 3.6	12.1 / 14.2	5.2 / 5.5
	p99	4.2 / 4.1	3.4 / 3.0	2.0 / 1.8	5.5 / 5.6	2.9 / 2.9
	p95	3.0 / 2.7	2.5 / 1.9	1.4 / 1.1	3.9 / 3.5	2.0 / 1.7
	rms	1.7 / 1.4	1.4 / 1.1	0.8 / 0.6	2.3 / 1.9	1.2 / 0.9
Helium temperature [°C]	max	2.4 / 1.9	2.1 / 1.8	1.0 / 1.0	3.6 / 2.5	1.5 / 1.4
	p99	1.3 / 1.2	1.1 / 1.1	0.5 / 0.7	1.9 / 1.6	0.7 / 0.9
	p95	1.0 / 1.0	0.9 / 0.9	0.3 / 0.5	1.5 / 1.3	0.5 / 0.7
	rms	0.5 / 0.5	0.5 / 0.4	0.2 / 0.2	0.8 / 0.6	0.3 / 0.3
Load demand tracking error [%]	max	3.1 / 3.1	2.6 / 2.6	2.0 / 2.0	5.1 / 4.6	3.0 / 2.8
	p99	2.1 / 2.2	1.5 / 1.8	1.2 / 1.4	2.5 / 2.6	1.9 / 2.0
	p95	1.5 / 1.7	1.1 / 1.4	0.9 / 1.1	1.7 / 1.9	1.3 / 1.5
	rms	0.8 / 0.9	0.6 / 0.7	0.5 / 0.6	0.9 / 1.0	0.7 / 0.8
Grid frequency [mHz]	max	92.7 / 92.1	76.7 / 79.0	58.5 / 58.8	152.4 / 138.0	91.2 / 84.0
	p99	62.3 / 66.2	44.4 / 52.7	37.5 / 43.1	74.7 / 77.8	55.7 / 58.7
	p95	**46.4** / 52.4	**33.6** / 41.2	**27.5** / 34.2	**51.7** / 57.3	**40.5** / 46.1
	rms	24.3 / 26.7	17.7 / 21.0	14.1 / 17.2	27.5 / 29.6	20.9 / 23.4
Reactivity insertion rate [pcm/s]	max	50.5 / 69.8	**40.2** / 48.8	**40.7** / 27.8	62.8 / 74.4	63.9 / 46.2
	p99	30.9 / 32.3	21.7 / 21.5	19.5 / 17.0	32.0 / 34.6	29.2 / 26.4
	p95	24.0 / 21.4	17.1 / 14.8	14.5 / 11.6	25.1 / 22.7	21.2 / 18.8
	rms	12.4 / 10.3	8.8 / 6.8	7.5 / 5.1	13.1 / 10.7	11.0 / 8.5
Maximum fuel temperature [°C]	max	819.0 / 818.6	819.0 / 818.6	818.0 / 818.0	820.6 / 819.8	818.3 / 818.1
	p99	818.3 / 818.1	818.3 / 818.1	817.7 / 817.9	819.1 / 818.6	817.8 / 817.9
	p95	818.0 / 817.8	818.0 / 817.8	817.5 / 817.7	818.4 / 818.1	817.6 / 817.7
	rms	817.2 / 817.2	817.2 / 817.3	817.2 / 817.2	817.2 / 817.3	817.1 / 817.2

5. Conclusions

This study investigated the frequency-dependent dynamic characteristics of a prismatic HTGR coupled to a steam Rankine cycle under hyperscale DC demand. An OpenModelica-based dynamic model was developed, and the plant response was first characterized through sinusoidal load perturbations. Subsequently, its dynamic responses were evaluated using facility-aggregated and workload-resolved DC demand profiles under both PI and SMC control. The key findings are summarized as follows:

- For the modeled plant-controller-grid configuration, the sinusoidal load perturbation analysis showed that the plant response depends strongly on the time scale of the imposed load variation. An intermediate time-scale region with characteristic periods of several tens of seconds produced the largest plant deviations, whereas faster fluctuations were more strongly attenuated and slower variations were effectively regulated.
- The facility-aggregated DC profile, despite spanning a wide operating range, was successfully accommodated because its dominant variations occurred over relatively long time scales. The major system variables remained well regulated within the prescribed transient operating constraints.
- The workload-resolved profiles produced substantially larger deviations as faster workload-driven fluctuations were retained. Frequency-band analysis further showed that the intermediate-frequency components had a greater influence on the plant response than the higher-frequency components. These results indicate that the spectral distribution of DC demand is an important consideration when coordinating nuclear generation with flexible resources.

Several limitations remain in the present study. First, the electrical grid was represented using a simplified generator and droop model rather than a detailed transmission-network representation. Practical power systems include governor deadbands, additional frequency-control dynamics, and interactions with other generators and grid resources, which may alter both the plant-side transient burden and the resulting grid-frequency response [13, 36]. Accordingly, the quantitative location of the identified sensitive frequency range should be interpreted as specific to the modeled configuration rather than as a universal characteristic of HTGR systems. Second, the present assessment primarily considers transient thermal-hydraulic operating constraints. Cumulative effects associated with repeated cyclic operation, including thermal fatigue, component degradation, and lifetime consumption, were not evaluated and may impose additional restrictions on long-term load-following operation [35]. Finally, the workload-resolved DC demand was statistically reconstructed from a limited set of AI-training measurements and therefore remains dependent on the assumed semi-Markov structure, aggregation level, and available measurement data.

Future work will extend the present framework to a more detailed grid representation and explicitly investigate coordinated operation with flexible resources such as battery energy storage systems. Such analysis will examine how frequency-selective load sharing affects nuclear-system dynamics, grid-frequency response, and the required power and energy capacities of supporting resources.

CRediT Author Statement

Wooseong Park: Conceptualization, Methodology, Formal analysis, Investigation, Data curation, Visualization, Writing – original draft. **W. Reed Kendrick:** Methodology, Writing – review & editing. **Emilio Baglietto:** Funding acquisition, Project administration, Writing – review & editing. **Dean Price:** Conceptualization, Supervision, Writing – review & editing.

Acknowledgment

This work was supported in part by an award from Amazon.
This research made use of Idaho National Laboratory computing resources, which are 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.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used ChatGPT for improving the language and readability of the manuscript. The authors reviewed and edited the output as needed and take full responsibility for the content of the published article.

Data availability

All the results and simulation files to reproduce them are publicly available via git. https://github.com/wspark3000/HTGR-DC-Input-Files

Appendix A. OpenModelica System Diagrams

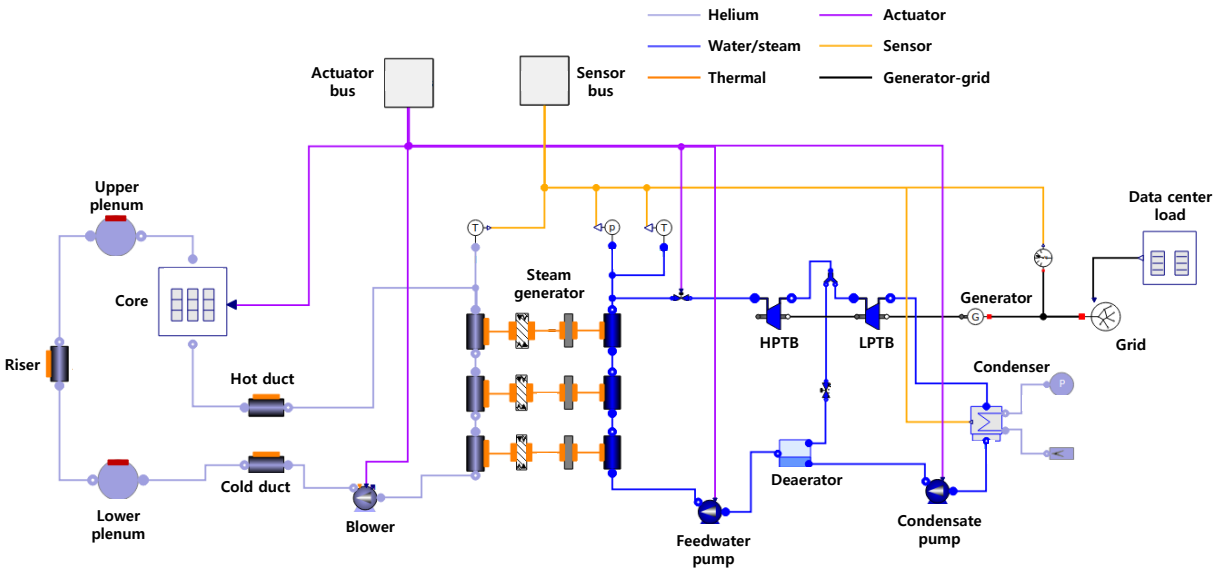

Figure A.1: Dynamic model in OpenModelica.

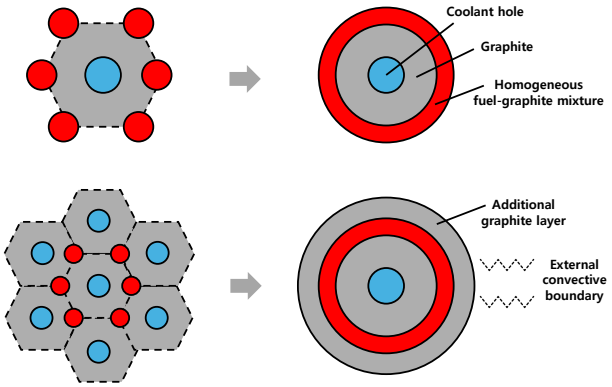

Figure A.2: Simplified representations of average (top) and hot (bottom) channels used in the reactor core model.

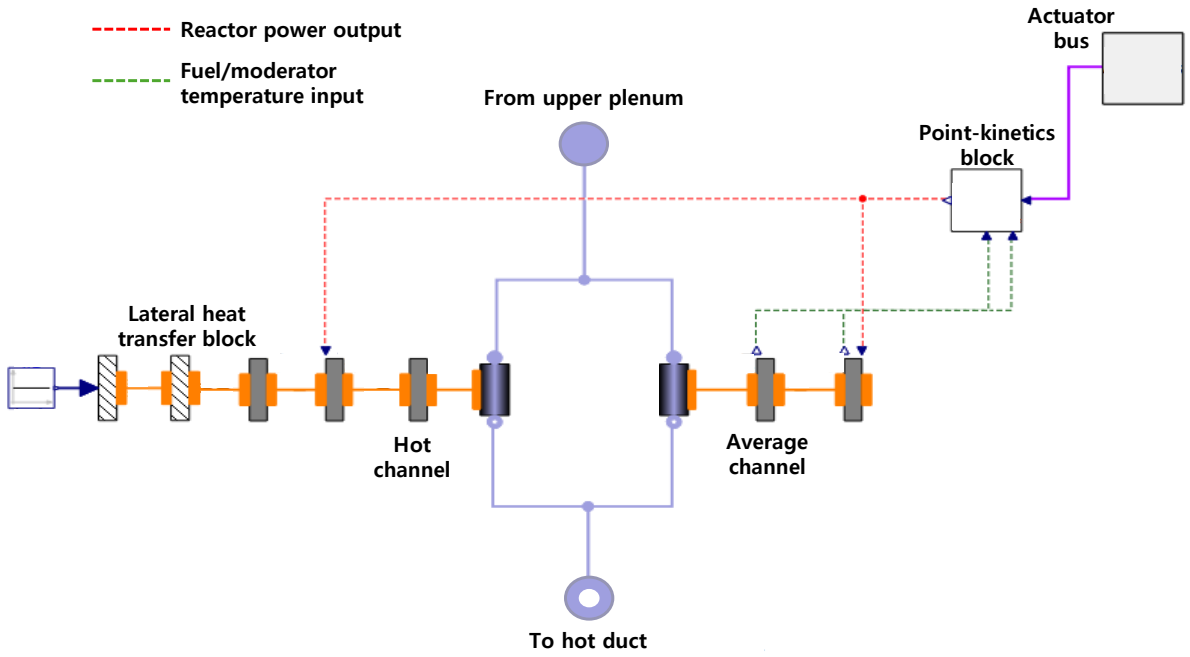

Figure A.3: Reactor core module in OpenModelica.

Appendix B. Tuning Results of OpenModelica Dynamic Model

Table B.1: Tuning results of PI and SMC controllers.

Control type	Controller	Parameter and value
PI	Reactivity	$K_p = 0.5,\ K_i = 25$
	Blower	$K_p = 1.5,\ K_i = 52.5$
	Feedwater pump	$K_p = 1.5,\ K_i = 52.5$
	Condensate pump	$K_p = 3.4,\ K_i = 170$
	Throttle valve	$K_p = 75,\ K_i = 2250$
SMC	Reactivity	$c = 1,\ \lambda = 0.1,\ \alpha = 5 \times 10^{-5},\ \phi = 0.01$
	Blower	$c = 1,\ \lambda = 0.3,\ \alpha = 10^{-3},\ \phi = 0.01$
	Feedwater pump	$c = 1,\ \lambda = 0.25,\ \alpha = 4.5 \times 10^{-4},\ \phi = 0.01$
	Condensate pump	$c = 1,\ \lambda = 0.6,\ \alpha = 6 \times 10^{-4},\ \phi = 0.01$
	Throttle valve	$c = 1,\ \lambda = 8,\ \alpha = 10^{-2},\ \phi = 0.005$

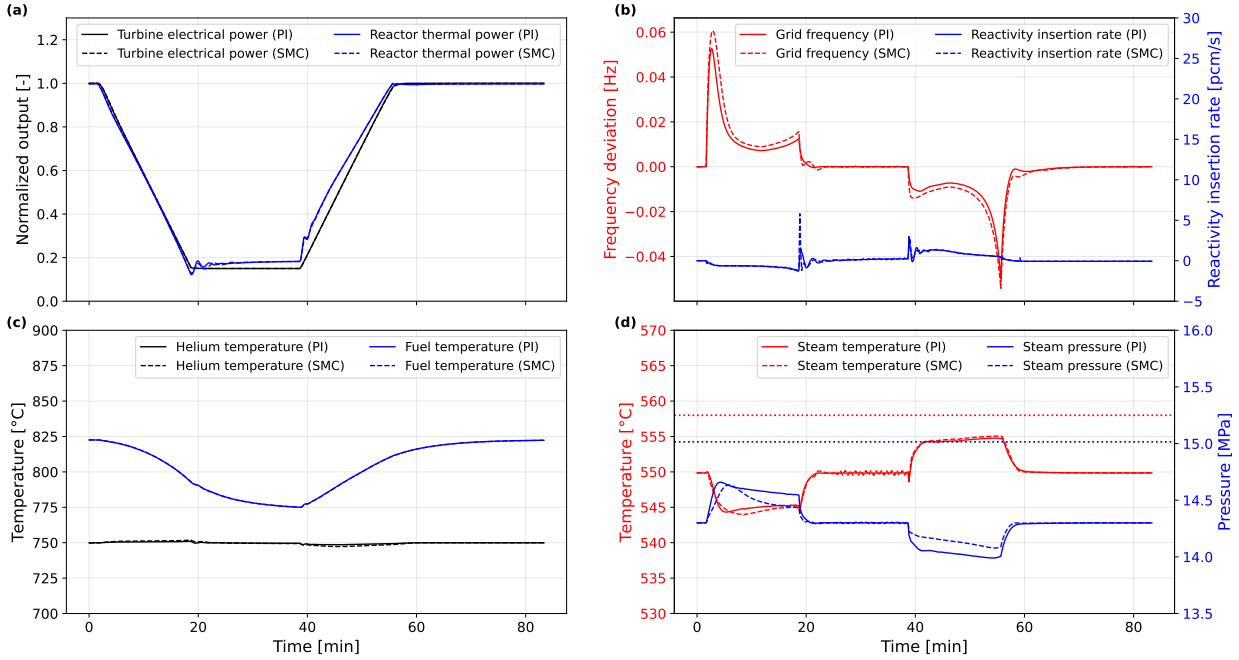

Figure B.1: System responses under the reference load transient used for controller tuning: a) normalized reactor thermal and turbine electrical powers; b) grid-frequency deviation and reactivity insertion rate; c) core outlet helium temperature and fuel maximum temperature; and d) steam-generator outlet temperature and pressure. Horizontal red and blue dotted lines in (d) represent the steam temperature and pressure limits, respectively.

References

[1] International Energy Agency, Energy and AI, Tech. rep., International Energy Agency (2025).

[2] Nuclear Newswire, Amazon buys nuclear-powered data center from Talen, https://www.ans.org/news/article-5842/amazon-buys-nuclearpowered-data-center-from-talen/, American Nuclear Society, Nuclear News (Mar. 2024).

[3] M. Terrell, New nuclear clean energy agreement with Kairos Power, https://blog.google/company-news/outreach-and-initiatives/sustainability/google-kairos-power-nuclear-energy-agreement/, Google (Oct. 2024).

[4] Constellation Energy Corporation, Constellation, Meta sign 20-year deal for clean, reliable nuclear energy in illinois, https://www.constellationenergy.com/news/2025/constellation-meta-sign-20-year-deal-for-clean-reliable-nuclear-energy-in-illinois.html, Constellation Energy Corporation (Jun. 2025).

[5] INTERNATIONAL ATOMIC ENERGY AGENCY, Non-baseload Operation in Nuclear Power Plants: Load Following and Frequency Control Modes of Flexible Operation, no. NP-T-3.23 in IAEA Nuclear Energy Series, INTERNATIONAL ATOMIC ENERGY AGENCY, Vienna, 2018.

[6] OECD Nuclear Energy Agency (NEA), Technical and Economic Aspects of Load Following with Nuclear Power Plants, Nuclear Development, OECD Publishing, Paris, 2021.

[7] A. M. Rad, R. A. Jacob, B. Poudel, M. Mamtimin, J. Zhang, Dynamic modeling and load-following control of small modular reactors with moving-boundary steam generators and thermodynamically coupled rankine cycles (2026). arXiv:2606.19296.

[8] S. Wu, X. Ma, J. Liu, J. Wan, P. Wang, G. Su, A load following control strategy for chinese modular high-temperature gas-cooled reactor HTR-PM, Energy 263 (2023) 125459.

[9] A. Rivas, G. K. Delipei, J. Hou, Modeling a generic pebble bed high-temperature gas-cooled reactor to perform load-following using simulink, Nuclear Engineering and Design 428 (2024) 113506.

[10] Y. Brits, F. Botha, H. van Antwerpen, H.-W. Chi, A control approach investigation of the Xe-100 plant to perform load following within the operational range of 100–25–100%, Nuclear Engineering and Design 329 (2018) 12–19.

[11] H. Zhang, Y. Li, Y. Zhao, J. Wang, S. Du, W. Ma, Modeling and simulation of a micro gas-cooled nuclear reactor using Modelica, Frontiers in Energy Research 11 (2023) 1206755.

[12] K. Frick, S. Bragg-Sitton, Development of the Nuscale power module in the INL Modelica ecosystem, Nuclear Technology 207 (4) (2021) 521–542.

[13] P. Wang, Y. Fu, X. Wei, F. Zhao, Simulation study of frequency control characteristics of a generation III+ nuclear power plant, Annals of Nuclear Energy 115 (2018) 502–522.

[14] K. Joshi, B. Poudel, R. Gokaraju, Investigating small modular reactor's design limits for its flexible operation with photovoltaic generation in microcommunities, Journal of Nuclear Engineering and Radiation Science 7 (3) (2021) 031303.

[15] G. C. Masotti, S. Lorenzi, M. E. Ricotti, Dynamic modelling and control of a small modular reactor in load following by cogeneration mode, Nuclear Engineering and Design 445 (2025) 114518.

[16] E. Choukse, B. Warrier, S. Heath, L. Belmont, A. Zhao, H. A. Khan, B. Harry, M. Kappel, R. J. Hewett, K. Datta, et al., Power stabilization for ai training datacenters, arXiv preprint arXiv:2508.14318 (2025).

[17] M.-S. Ko, H. Zhu, Wide-area power system oscillations from large-scale ai workloads, IEEE Transactions on Power Systems (2026).

[18] R. Vercellino, J. Willard, G. Campos, W. d. S. Pereira, O. Hull, M. Selensky, J. Mueller, Measurement of generative ai workload power profiles for whole-facility data center infrastructure planning, arXiv preprint arXiv:2604.07345 (2026).

[19] C. Chaudhary, A. Abdelkader, Y. Pei, M. Benidris, J. Mitra, Spatial Load Correlation in AI Data-Center-Dominated Power Systems, in: 2026 IEEE Power & Energy Society General Meeting (PESGM), IEEE, 2026, p. 1–5.

[20] E. Joosten, Thermal-hydraulic safety analysis of the httu and gemini+ cores in trace, Master's thesis, KTH Royal Institute of Technology (2022).

[21] International Electrotechnical Commission, IEC 60045-1:1991: Steam turbines - part 1: Specifications (1991).

[22] International Atomic Energy Agency, High temperature gas cooled reactor fuels and materials, Tech. Rep. IAEA-TECDOC-1645, International Atomic Energy Agency, Vienna, Austria (2010).

[23] J. W. Sterbentz, P. D. Bayless, L. O. Nelson, H. D. Gougar, J. C. Kinsey, G. Strydom, A. Kumar, High-temperature gas-cooled test reactor point design, Tech. Rep. INL/EXT-16-38296, Idaho National Laboratory (Apr. 2016).

[24] S. Nakagawa, K. Sawa, K. Ohashi, Safety analysis of abnormal reactivity events in httr, Journal of Nuclear Science and Technology 30 (6) (1993) 579–588.

[25] North American Electric Reliability Corporation, PRC-024-4: Frequency and voltage protection settings for synchronous generators, type 1 and type 2 wind resources, and synchronous condensers, NERC Reliability Standard (2024).

[26] California Independent System Operator, RC0210: Monitoring frequency and balancing authority performance, reliability Coordinator Procedure, Version 1.5, effective February 1, 2025 (2025).

[27] R. Vilanova, A. Visioli, The proportional-integral-derivative (pid) controller, Wiley Encyclopedia of Electrical and Electronics Engineering (1999) 1–15.

[28] Y. Shtessel, C. Edwards, L. Fridman, A. Levant, et al., Sliding mode control and observation, Vol. 10, Springer, 2014.

[29] P. Fritzson, A. Pop, K. Abdelhak, A. Ashgar, B. Bachmann, W. Braun, D. Bouskela, R. Braun, L. Buffoni, F. Casella, et al., The openmodelica integrated environment for modeling, simulation, and model-based development, Modeling Identification and Control 41 (4) (2020) 241–295.

[30] F. Casella, A. Leva, Modelling and simulation of thermal power generation systems with modelica: 20 years of experience with the ThermoPower library, IFAC-PapersOnLine 58 (13) (2024) 227–234.

[31] T. L. Bergman, Fundamentals of heat and mass transfer, John Wiley & Sons, 2011.

[32] D. H. Cooke, On prediction of off-design multistage turbine pressures by stodola's ellipse, Journal of Engineering for Gas Turbines and Power 107 (3) (1985) 596–606.

[33] K. Baumann, Recent developments in steam turbine practice, Journal of the Institution of Electrical Engineers 48 (213) (1912) 768–842.

[34] C. Chaudhary, M. Murillo, M. Ben-Idris, J. Mitra, D. Pandit, A. Bera, Modal analysis of spatial load correlation in ai data center-dominated power systems, arXiv preprint arXiv:2606.13847 (2026).

[35] X. L. Yan, H. Sato, Y. Tachibana, K. Kunitomi, R. Hino, Evaluation of high temperature gas reactor for demanding cogeneration load follow, Journal of nuclear science and technology 49 (1) (2012) 121–131.

[36] G. Wu, P. Ju, X. Song, C. Xie, W. Zhong, Interaction and coordination among nuclear power plants, power grids and their protection systems, Energies 9 (4) (2016) 306.