Design and Comparison of Two Advanced Core Control Systems for Flexible Operation of Pressurized Water Reactors

Guillaume Dupre, Philippe Chevrel, Mohamed Yagoubi, Alain Grossetete

Version of Record: https://www.sciencedirect.com/science/article/pii/S0967066122000685
Manuscript_a474f1c712162412255e5d29f96d574f

Design and Comparison of Two Advanced Core Control Systems for Flexible Operation of Pressurized Water Reactors

Guillaume Dupré[a,b,*], Philippe Chevrel[b], Mohamed Yagoubi[b], Alain Grossetête[a]

[a] *Reactor Process Department, Framatome, 1 pl. Jean Millier, Paris La Défense, 92400, France*
[b] *LS2N (UMR 6004), IMT Atlantique, 4 rue Alfred Kastler, Nantes, 44307, France*

Abstract

This paper focuses on the design of advanced core control systems for future generations of pressurized water reactors. The objective is to improve the flexibility of nuclear power plants to cope with the rapid growth of renewable energies. In practice, this means that the average coolant temperature, the axial power distribution of the reactor core and the position of the control rods have to be properly regulated during power variations. In previous work, conducted by the same authors, two promising approaches were investigated: 1) fixed-structure gain-scheduled control and 2) nonlinear model predictive control. Here, both methods are tested according to industry standards in an attempt to determine the best one for our problem. To achieve this, two different controllers are designed using a new multipoint kinetic model of the reactor core, which provides an accurate representation of the axial power distribution. The advantages and drawbacks of both design methodologies are discussed and then compared on PWRSimu, an intermediate complexity pressurized water reactor simulator developed by Framatome.

Keywords: Nuclear Power Plants, Gain-Scheduling, Model Predictive Control

Nomenclature

ACT	Average Coolant Temperature
AO	Axial Offset
GSC	Gain-Scheduled Control
MS	Multiple-Shooting
NMPC	Nonlinear Model Predictive Control
NP	Nominal Power
PWR	Pressurized Water Reactor
SS	Single-Shooting

1. Introduction

In order to address climate change, many countries are seeking to replace fossil fuel power plants by renewable energy sources [1]. This energy transition poses new challenges in terms of management of the electrical network. Since excess electricity cannot currently be stored on a large scale, generation and consumption have to be continuously balanced to ensure grid stability. For this purpose, a number of power plants must be flexible, i.e., capable of adjusting their power output on demand. However, unlike fossil fuel power plants, which are easily controllable, part of renewable energy sources, such as wind turbines or solar panels, are inherently intermittent. It is therefore essential to enhance the flexibility of conventional power plants to prepare for the upcoming expansion of renewables [2].

A nuclear power plant can be operated either in base-load or in load-following mode [3]. In base-load mode, the power output of the plant remains constant regardless of electricity demand. Most nuclear power plants around the world are still operated in base-load mode because it is simpler and more cost-effective. In load-following mode, the power output of the plant is adjusted depending on the forecasted electricity consumption. Typically, the grid operator sends a daily load profile in advance to the plant operator, who will then set the power target of the turbine accordingly. Whether in base-load or in load-following mode, the nuclear power plant can also provide frequency control to the grid. This is achieved by a dedicated controller that automatically adjusts the power output of the turbine to offset small load imbalances within seconds.

Flexible operation of a nuclear power plant is directly related to the design of the core control system [4]. The main task of the core control system is to maintain the average coolant temperature (ACT) and the axial power distribution of the reactor core, or axial offset (AO), within appropriate limits. These limits are defined upstream to ensure acceptable performance levels and safe operation of the plant. During power variations, ACT regulation is achieved by moving several neutron-absorbing control rods inside the reactor core. However, their movements have a detrimental impact on its axial power distribution. Initially, the movements of the control rods were mitigated by

*Corresponding author
Email addresses: guillaume.dupre@framatome.com (Guillaume Dupré), philippe.chevrel@imt-atlantique.fr (Philippe Chevrel), mohamed.yagoubi@imt-atlantique.fr (Mohamed Yagoubi), alain.grossetete@framatome.com (Alain Grossetête)

adjusting the boron concentration of the primary coolant. This strategy was limiting in terms of flexibility because the boron concentration of the coolant can only change slowly and gradually decreases over the fuel burn-up cycle. For example in mode A, the maximum power variation rate of the plant is typically restricted to 1-2 %NP/min at the beginning of the cycle, and 0.2 %NP/min at the end. To overcome these limitations, advanced core control systems usually comprise two separate groups, or banks, of control rods [4]. In this way, one bank of control rods can be used to perform power variations by, e.g., controlling the ACT, while the other can be used to reduce axial power distribution disturbances. Thus, flexibility is increased with power variation rates of up to 5 %NP/min. After power variations, the xenon distribution of the reactor core will slowly change for several hours until steady-state is reached. During this time, the evolution of xenon concentration has to be counterbalanced to prevent the AO of the reactor core from drifting away. In highly advanced core control systems, such as mode T [5], this can be done either by moving the control rods or by adjusting the boron concentration of the primary coolant. The main motivation for using the control rods is to minimize the volume of effluent. Yet, this can decrease the maneuvering capabilities of the plant, as the control rods may not be properly positioned for a quick return to nominal power. Maximum flexibility is achieved when xenon poisoning is compensated by boron concentration adjustments and when the level of insertion of the control rods is controlled.

In practice, industrial core control systems manage to meet all the aforementioned requirements using single input single output Proportional-Integral-Derivative control along with sophisticated logic rules. Even if this approach gives satisfactory results for now, it may become insufficient against more stringent grid management criteria. Fortunately, a wide variety of core control systems have already been proposed (see [6] and references therein). Popular methods include LQG/LTR control, H_∞ control, sliding-mode control, model predictive control and gain-scheduled control. Nevertheless, these conceptual core control systems do not tackle all the issues that arise from the flexible operation of nuclear power plants. Hence, this paper proposes to bridge the gap between industry and academia by developing two different advanced core control systems using modern design techniques. In view of the strong operational constraints that are placed on the nuclear industry, two promising approaches have been identified: fixed-structure gain-scheduled control [7, 8] and nonlinear model predictive control [9].

To the best of our knowledge, very few papers have tried to apply these techniques to the control of a nuclear reactor core. For instance, it seems that [7, 8] are the first ones that studied fixed-structure gain-scheduled control in this context. Regarding model predictive control, only a handful of examples have been reported in the scientific literature. In [10, 11], a linear model predictive controller is designed to achieve both ACT and AO regulation of a reactor core that is operated in load-following mode. However, the boron concentration of the coolant is adjusted by another module and the nuclear power plant does not provide frequency control to the grid. In another paper [12], the power output of a reactor core is controlled by a robust nonlinear model predictive controller. Yet, this theoretical work is not industrially relevant as it does not even address AO regulation. By contrast, the main contribution of this paper is to conduct a realistic and comprehensive study on the design of nuclear reactor core control systems. The objective is to give practical insights and methodological guidelines that could help industrial practitioners working in nuclear engineering. Two modern control techniques, which were seldom tested in this field, are compared against the technical specifications of Framatome. In order to ensure consistency with industry standards, the actuators and control objectives are identical to those considered in mode T [5].

Following on from the work carried out in [8], the first core control system is designed based on a fixed-structure H_2/H_∞ gain-scheduled control approach. The main improvement is that the controller is tuned using a multipoint kinetic model of the reactor core rather than a point kinetic one. In addition, the mathematical formulation of the H_2/H_∞ synthesis problem has been refined and made more readable. Moreover, the gains are now simultaneously tuned at each operating point and smoothly interpolated as a function of the scheduling variable [13]. This greatly simplifies the design of the controller, as it avoids *a posteriori* interpolation, or post-processing, of the gains. The second core control system is designed based on a nonlinear model predictive control approach [14, 15]. This time, the multipoint kinetic model of the reactor core is embedded into the controller to predict the future response of the plant. Thus, the control inputs can be calculated online by repeatedly solving a constrained finite horizon optimal control problem. To this end, the model of the reactor core is reduced using singular perturbation theory so that it can be efficiently simulated by the controller.

The remainder of this paper is organized as follows: Section II describes the system under study and formalizes the control objectives; Section III presents the design methodology of the fixed-structure H_2/H_∞ gain-scheduled controller and Section IV presents the design methodology of the nonlinear model predictive controller; Then, the advantages and drawbacks of both techniques are discussed in Section V based on a typical load-following MATLAB®/Simulink® simulation scenario; Finally, conclusions and perspectives are drawn in Section VI.

2. Problem Formulation

2.1. System Modeling

To design the core control systems, a new multipoint core kinetic model of a pressurized water reactor (PWR), drawn in Fig. 1, has been developed (a brief overview

of the model is given in [16] Appendix F). Its main advantage, compared to the former point kinetic model presented in [7, 16], is that the behavior of the AO is more accurately represented since the reactor core is now divided into six meshes. However, increasing the number of meshes also increases the size of the model. As illustrated in Fig. 1, every new mesh requires five additional states: two average neutron n and delayed neutron c densities, two iodine I and xenon X_e concentrations, and one inlet temperature T. For each mesh, neutron dynamics are described by point kinetics equations with one group of delayed neutrons. An adequately chosen constant exchange coefficient has been added to take account of the transfer of neutrons between meshes. In addition, the iodine and xenon concentrations are calculated by a multipoint iodine-xenon estimator. The rise in coolant temperature, from the bottom to the top of each mesh, is given by a first order differential equation and is proportional to the amount of power generated by nuclear fission. In each mesh, the growth rate of the fission chain reaction is characterized by a quantity known as reactivity [17, 18]. Criticality is achieved when the fission chain reaction is stable and self-sustaining, i.e., when the reactivity of the reactor core equals zero. Reactivity depends on many factors, namely the fuel and coolant temperatures, the xenon concentration of the reactor core, the position of the control rods and the boron concentration of the primary coolant. The resulting feedbacks and underlying interactions between all involved variables lead to strong plant nonlinearities and coupling.

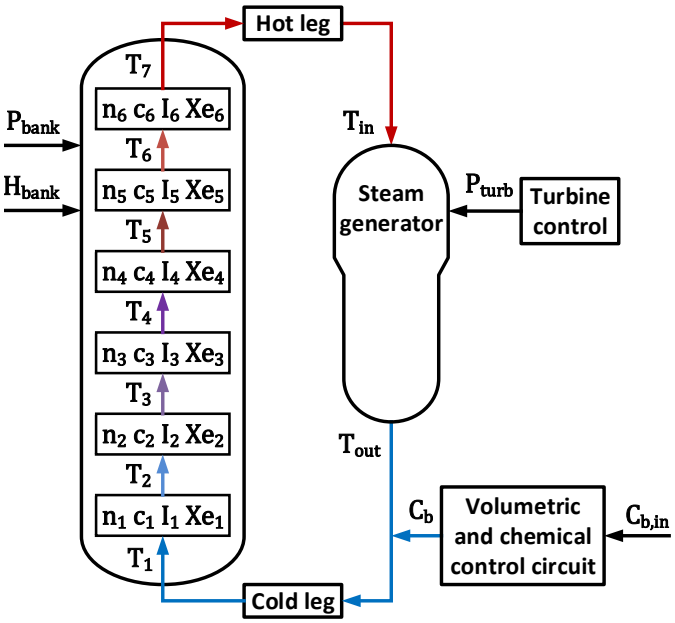

Fig. 1: Simplified diagram of the multipoint PWR core kinetic model

In order to move from reactor core to plant scale, a very simple steam generator is modeled. Once again, it is assumed that the fall in temperature between the inlet and outlet of the steam generator is given by a first order differential equation and is proportional to the power of the turbine. In fact, whenever the power of the turbine is decreased (resp. increased), the average coolant temperature will increase (resp. decrease). As a result, the nuclear power of the reactor core will also decrease (resp. increase) and ultimately reach steady-state because the PWR was designed so that the fission chain reaction can be stabilized by temperature reactivity feedbacks alone (i.e., Doppler effect and moderation) [17]. However, it is very likely that, without control, the average coolant temperature will deviate from the limiting conditions of operation. Furthermore, xenon-induced spatial power oscillations may appear if the axial power distribution of the reactor core is left uncontrolled. Therefore, the role of the core control system is to ensure that both the ACT and the AO stay within their respective limits before, during and after power variations [4]. This can be achieved by slightly altering the reactivity of the reactor core.

In mode T [5], the nuclear reaction is controlled by moving two separate banks of control rods inside the reactor core and by adjusting the boron concentration of the primary coolant. Reactivity can be increased (resp. decreased) either by withdrawing (resp. inserting) the control rods from (resp. in) the reactor core or by decreasing (resp. increasing) the boron concentration of the coolant. The control rods are mainly used to handle abrupt changes in reactivity but adversely impact the power distribution of the reactor core. In fact, reactivity will mostly decrease in the region of the reactor core where the control rods are being inserted. To alleviate this problem, one bank of control rods, denoted by $\mathrm{P_{bank}}$, is dedicated to ACT regulation whereas the other, denoted by $\mathrm{H_{bank}}$, is dedicated to AO regulation. Conversely, the boron concentration of the primary coolant acts uniformly on the reactor core but takes time to become effective. This is due to the fact that the solutions of boric acid and demineralized water, which are injected to modify the boron concentration, have to flow throughout the volumetric and chemical control circuit before reaching the coolant. Hence, boron concentration adjustments are mostly used to counterbalance axial xenon oscillations after power variations. To summarize:

- ACT regulation is achieved by moving $\mathrm{P_{bank}}$
- AO regulation is achieved by moving $\mathrm{H_{bank}}$ during power variations and by adjusting the boron concentration of the primary coolant during xenon oscillations

In addition, the mode T allows the reactor core to quickly return to its nominal power. This is achieved by controlling the level of insertion of $\mathrm{P_{bank}}$ in the reactor core. Last but not least, all three controlled variables can be successfully regulated even when the turbine provides frequency control to the grid.

2.2. Control Objectives

As mentioned before, the goal of an advanced core control system is to keep the ACT, the AO and the position

of $\mathrm{P_{bank}}$ within appropriate limits during power variations and while frequency control is active. These limits were determined beforehand by nuclear reactor engineers to comply with genuine performance and safety requirements. In mode T, their values are given by:

$$\begin{aligned} &|T_a(t) - T_{a,ref}\left(P_{ref}(t)\right)| \leq 1.5\ ^\circ\mathrm{C} \\ &|AO(t) - AO_{ref}| \leq 5\ \% \\ &|P_{bank}(t) - P_{bank,ref}\left(P_{ref}(t)\right)| \leq 30 \text{ steps}, \end{aligned} \tag{1}$$

where T_a is the average temperature of the coolant, AO is the axial offset of the reactor core, P_{bank} is the position of the first bank of control rods, and $T_{a,ref}$, AO_{ref} and $P_{bank,ref}$ are their respective reference signals. Note that the reference signals $T_{a,ref}$ and $P_{bank,ref}$ are both piecewise linear functions of the load profile P_{ref} whereas AO_{ref} is a constant. Shorthand notation for the deviations of ACT, AO and $\mathrm{P_{bank}}$ are:

$$\begin{aligned} \Delta T_a(t) &= T_a(t) - T_{a,ref} \\ \Delta AO(t) &= AO(t) - AO_{ref} \\ \Delta P_{bank}(t) &= P_{bank}(t) - P_{bank,ref}\left(P_{ref}(t)\right). \end{aligned}$$

Actuator saturation should also be considered to make the core control system practically workable. In our case, the reactor core is actuated by changing the speeds of both $\mathrm{P_{bank}}$ and $\mathrm{H_{bank}}$, and by adjusting the boron concentration of the coolant. The speeds of $\mathrm{P_{bank}}$ and $\mathrm{H_{bank}}$ are simply limited by:

$$\begin{aligned} \left|\frac{\mathrm{d}P_{bank}(t)}{\mathrm{d}t}\right| &\leq 75 \text{ steps/min} \\ \left|\frac{\mathrm{d}H_{bank}(t)}{\mathrm{d}t}\right| &\leq 75 \text{ steps/min}, \end{aligned} \tag{2}$$

where H_{bank} is the position of the second bank of control rods. Besides, P_{bank} and H_{bank} are saturated as well since the control rods cannot move beyond their maximum insertion and extraction thresholds:

$$\begin{aligned} 36 \leq P_{bank} \leq 1053 \text{ steps} \\ 9 \leq H_{bank} \leq 411 \text{ steps}. \end{aligned} \tag{3}$$

The upper and lower bounds of the injected boron concentration variation rate depend on the current boron concentration of the coolant. More precisely, the higher the boron concentration of the coolant, the lower the efficiency of a boric acid injection. Similarly, the lower the boron concentration of the coolant, the lower the efficiency of a demineralized water injection. This translates into the following inequalities:

$$-\frac{10}{M_t}C_b(t) \leq \frac{\mathrm{d}C_{b,in}(t)}{\mathrm{d}t} \leq \frac{3}{M_t}\left(C_{b,max} - C_b(t)\right), \tag{4}$$

where $C_{b,max}$ is the maximum boron concentration of the coolant, M_t is the total mass of water in the primary cooling circuit, $C_{b,in}$ is the injected boron concentration and C_b is the boron concentration of the coolant.

In our model, it is assumed that the turbine control system has already been designed and is functioning perfectly. Thus, the power of the turbine can be expressed as $P_{turb} = P_{ref} + \delta P_{ref}$ where P_{ref} is the known load profile and δP_{ref} is an unknown reference signal that is added when frequency control is active. Specifically, P_{ref} appears as a ramp-like signal during power variations and as a step-like signal otherwise. The additional term δP_{ref} is norm-bounded by 2.5 %NP and changes about every second. From the perspective of the reactor core, P_{turb} is seen as a disturbance signal that has to be rejected by the core control system.

3. Fixed-Structure Gain-Scheduled Controller

The idea of gain scheduling is to break down the nonlinear control problem into a finite number of linear subproblems [19, 20]. The main advantage of this divide-and-conquer approach is to rely on well-known linear design tools rather than convoluted nonlinear methods. This is particularly interesting given the recent developments in non-smooth optimization for structured robust control [21].

3.1. Linearization of the Nonlinear Model

The first step in designing the fixed-structure H_2/H_∞ gain-scheduled controller is to linearize the nonlinear model of the reactor core around several operating points. All were determined using Framatome's certified three-dimensional core kinetics computer code, named SMART [22], starting from stationary conditions with equilibrium xenon at 100 %NP. The computation of an operating point begins by setting the nuclear power of the reactor core at a constant level. Then, the positions of $\mathrm{P_{bank}}$ and $\mathrm{H_{bank}}$ are updated so that criticality is achieved with minimal ACT and AO deviations. The boron concentration of the coolant, on the other hand, is fixed according to the fuel burn-up. In total, 21 operating points were computed over an evenly spaced grid of power values between 100 %NP and 50 %NP (i.e., one point every 2.5 %NP). This number of points is sufficient to obtain a fine-grained description of the original system and to ensure a smooth transition of the gains.

Then, the operating points are uniquely defined by a set of scheduling variables that covers the whole operating range of the real plant. Initially, two different scheduling variables were considered: the position P_{bank} and the known load profile P_{ref}. However, using both variables is irrelevant here, as the closed-loop dynamic of P_{bank} should be directly linked to P_{ref}. The question that remains is whether to schedule the controller with respect to the internal variable P_{bank} or the exogenous signal P_{ref}. Theoretically speaking, it would be preferable to choose P_{ref} over P_{bank} since the impact of the hidden coupling terms, which appear in the dynamic of the linearized gain-scheduled controller [20], is less pronounced when an external signal is used for scheduling. Yet, in reality, the

behavior of the reactor core depends heavily on the position of the control rods. Therefore, just like in [8], P_{bank} is selected as the only scheduling variable.

Subsequently, a collection of 21 linear open-loop models is numerically computed with the linearization algorithm of MATLAB®/Simulink®. For all $i \in [\![1, 21]\!]$, the i-th linear open-loop model is obtained by linearizing the nonlinear model of the plant around the operating point defined by $\sigma_i = P^{(i)}_{bank,eq}$. The iodine-xenon estimator has been discarded from the linearization process because its dynamics are much slower than that of the controlled outputs. Finally, the state-space realization of the i-th linearized open-loop model is given by:

$$\begin{aligned} \dot{x}_r(t) &= A_i x_r(t) + Bu(t) + Ed(t) \\ y &= Cx_r(t), \end{aligned}$$

where $x_r \in \mathbb{R}^{25}$ is the truncated state vector (i.e., without iodine and xenon dynamics), $u = \left[\dot{P}_{bank}, \dot{H}_{bank}, \dot{C}_{b,in}\right]^T \in \mathbb{R}^3$ are the control inputs, $d = P_{turb} \in \mathbb{R}$ is the exogenous disturbance, and $y = [T_{in}, T_{out}, AO, P_{bank}, P_{core}]^T \in \mathbb{R}^5$ are the measured outputs. Note that the controlled outputs $z = [T_a, AO, P_{bank}]^T \in \mathbb{R}^3$ are a subset of the measured output y in the sense that:

$$z = Ty \quad \text{with} \quad T = \begin{bmatrix} 1/2 & 1/2 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}.$$

Since the control inputs u are the time derivatives of the state variables P_{bank}, H_{bank} and $C_{b,in}$, the open-loop state matrix A_i and the control matrix B can be decomposed as:

$$A_i = \begin{bmatrix} A_{1i} & A_{2i} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \qquad B = \begin{bmatrix} \mathbf{0} \\ I_3 \end{bmatrix}.$$

Thus, the poles of the i-th linearized open-loop model are the union of the eigenvalues of A_{1i} plus three controllable poles located at $s = 0$. For the chosen set of model parameters, it was observed that every matrix A_{1i} is Hurtwitz stable. This numerical assessment is consistent with reality insofar as the reactor core was designed to be stable under normal conditions. Hence, it can be concluded that the linear open-loop models are all stabilizable.

3.2. Architecture of the Controller

To achieve asymptotic tracking of a class of reference inputs in the presence of another class of disturbances, a model of both exogenous signals should be embedded in the controller. Then, this model should be duplicated as many times as the number of outputs to be controlled [23]. Here, the reference inputs and the disturbance belong to the same class of signals that are constant and piecewise linear. Thus, the controller should comprise at least two integrators per controlled output channel to achieve the desired closed-loop performance. However, an integrator is already included in every input-output transfer function of the linearized open-loop models. Hence, only one extra integrator needs to be added per controlled output channel. This can be realized with the following linear time-invariant output feedback controller:

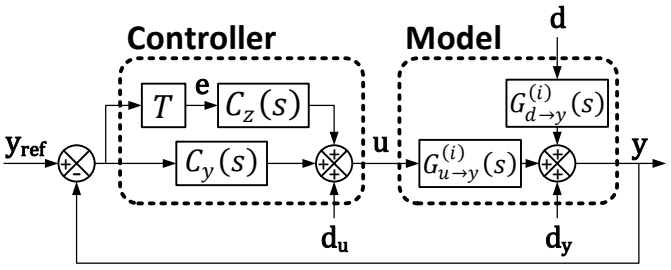

Fig. 2: Block diagram of the linear time-invariant closed-loop system computed at the i-th operating point

$$\begin{aligned} \dot{x}_I(t) &= K_I \left(z_{ref}(t) - z(t)\right) \\ u(t) &= K_P \left(y_{ref}(t) - y(t)\right) + x_I(t), \end{aligned}$$

where $y_{ref} \in \mathbb{R}^5$ are the reference signals of the measured outputs and $z_{ref} \in \mathbb{R}^3$ are the reference signals of the controlled outputs. This controller can be seen as a multiple-input multiple-output version of a Proportional-Integral controller with gain matrices $K_P \in \mathbb{R}^{3\times5}$ and $K_I \in \mathbb{R}^{3\times3}$. The proportional part of the controller is then filtered to provide robustness against high-frequency noise and unmodeled dynamics:

$$\begin{aligned} \dot{x}_I(t) &= K_I(z_{ref}(t) - z(t)) \\ \dot{x}_P(t) &= \operatorname{diag}\left(\frac{1}{a}, \frac{1}{a}, \frac{1}{b}\right)\Big(- x_P(t) + K_P \left(y_{ref}(t) - y(t)\right)\Big) \\ u(t) &= x_P(t) + x_I(t), \end{aligned}$$

where $a > 0$ and $b > 0$ are the time constants of the first-order low-pass filters that are applied to x_P. The time constants of the first and second filters are chosen equal since the dynamic of P_{bank} is expected to be similar to that of H_{bank}. A block diagram of the linear time-invariant closed-loop system that is computed at the i-th operating point, $i \in [\![1, 21]\!]$, is shown in Fig. 2, where:

$$\begin{aligned} G^{(i)}_{d\to y}(s) &= C\left(sI_{25} - A_i\right)^{-1} E \\ G^{(i)}_{u\to y}(s) &= C\left(sI_{25} - A_i\right)^{-1} B \\ C_y(s) &= \operatorname{diag}\left(\frac{1}{as+1}, \frac{1}{as+1}, \frac{1}{bs+1}\right) K_P \\ C_z(s) &= \frac{1}{s} K_I, \end{aligned}$$

and $e = z - z_{ref}$ is the tracking error:

$$e = [\Delta T_a, \Delta AO, \Delta P_{bank}]^T. \tag{5}$$

The signals $d_u \in \mathbb{R}^3$ and $d_y \in \mathbb{R}^5$ are unknown additive disturbances that may appear on the actuators and on the sensors of the plant. The input and the output sensitivity functions:

$$\begin{aligned} S^{(i)}_u(s) &= \left[I_5 + \left(C_y(s) + TC_z(s)\right) G^{(i)}_{u\to y}(s)\right]^{-1} \\ S^{(i)}_y(s) &= \left[I_3 + G^{(i)}_{u\to y}(s)\left(C_y(s) + TC_z(s)\right)\right]^{-1}, \end{aligned}$$

are given by the closed-loop transfer functions from d_u to u and from d_y to y respectively.

3.3. Tuning and Gain-Scheduling

The conventional approach [19, 20] for scheduling the gains of the controller would be to: 1) synthesize a linear time-invariant controller at each operating point using, e.g., a fixed-structure H_2/H_∞ synthesis algorithm and 2) smoothly interpolate the gains of the controller as a function of the scheduling variable P_{bank}. However, independently tuning the controller at each operating point may lead to undesirable jumps in the values of the gains once they have been interpolated. To prevent this, the gains can be tuned against multiple neighboring models rather than just one [8]. Here, the tuning and scheduling stages are blended together using the Gain Surface Tuning method presented in [13]. First, the gain matrices K_P, K_I and the time constants a, b are written as quadratic polynomials of the scheduling variable $\sigma = P_{bank}$:

$$\begin{aligned} K_P(\sigma) &= K_{P0} + K_{P1}\sigma + K_{P2}\sigma^2 \\ K_I(\sigma) &= K_{I0} + K_{I1}\sigma + K_{I2}\sigma^2 \\ a(\sigma) &= a_0 + a_1\sigma + a_2\sigma^2 \\ b(\sigma) &= b_0 + b_1\sigma + b_2\sigma^2, \end{aligned}$$

with coefficients $K_{Pj} \in \mathbb{R}^{3\times5}$, $K_{Ij} \in \mathbb{R}^{3\times3}$, $a_j \in \mathbb{R}$ and $b_j \in \mathbb{R}$ for all $j \in \{1,2,3\}$. Then, these coefficients are tuned to ensure proper closed-loop performance at each operating point. Specifically, let $\mathbf{K} \in \mathbb{R}^{78}$ be the vector of tunable parameters. For each linearized open-loop model, the performance objectives are represented by the following frequency-domain criteria:

$$\min_{\mathbf{K}} \left\| \mathrm{diag}\left(\beta_1,\beta_2,\beta_3\right) H^{(i)}_{z_{ref}\to e}(\mathbf{K},s)\frac{1}{s^2} \right\|_2$$

$$\text{subject to: } \begin{cases} \left\| W_{11} H^{(i)}_{d\to e_1}(\mathbf{K},s) \right\|_\infty \le 1 \\ \left\| W_{12} H^{(i)}_{d\to e_2}(\mathbf{K},s) \right\|_\infty \le 1 \\ \left\| W_{13} H^{(i)}_{d\to e_3}(\mathbf{K},s) \right\|_\infty \le 1 \\ \left\| W_{21} H^{(i)}_{d\to u_1}(\mathbf{K},s) \right\|_\infty \le 1 \\ \left\| W_{22} H^{(i)}_{d\to u_2}(\mathbf{K},s) \right\|_\infty \le 1 \\ \left\| W_{23} H^{(i)}_{d\to u_3}(\mathbf{K},s) \right\|_\infty \le 1 \\ \left\| \gamma H^{(i)}_{d_u\to u}(\mathbf{K},s) \right\|_\infty \le 1 \\ \mathrm{Re}(p^{(i)}(\mathbf{K})) \le -\alpha, \end{cases}$$

where, for all $i \in [\![1,21]\!]$, $H^{(i)}_{u\to y}$ denotes the closed-loop transfer function from input signal u to output signal y for the i-th plant model, $\|\cdot\|_2$ and $\|\cdot\|_\infty$ are the H_2 and H_∞ norms, and $p^{(i)}$ are the closed-loop poles of the i-th plant model. Having two integrators in the loop transfer function means that $H^{(i)}_{z_{ref}\to e}$ has two zeros at $s = 0$. Hence, its H_2 norm can be shaped with a double integrator weighting function which, in this context, is interpreted as a reference signal generator. In fact, minimizing the frequency-weighted H_2 norm of $H^{(i)}_{z_{ref}\to e}$ ensures both asymptotic tracking and disturbance rejection of step-like and ramp-like exogenous signals. The parameters $(\beta_1,\beta_2,\beta_3) \in \mathbb{R}^3_{>0}$ are chosen such that ACT regulation takes precedence over AO regulation which, in turn, takes precedence over $\mathrm{P_{bank}}$ regulation. The scalar weights $(W_{11},W_{12},W_{13}) \in \mathbb{R}^3_{>0}$ and $(W_{21},W_{22},W_{23}) \in \mathbb{R}^3_{>0}$ are used to limit the maximum set-point deviation and the maximum control effort that are induced by a change in the power of the turbine. The reason why the controller is tuned against several single-input, single-output H_∞ constraints, rather than one multiple-input, multiple-output H_∞ constraint, is to make the optimization problem more readable. Finally, the parameters $\gamma \in \mathbb{R}_{>0}$ and $\alpha \in \mathbb{R}_{>0}$ are selected to provide a sufficient level of input-robustness and stability to the closed-loop system.

4. Nonlinear Model Predictive Controller

The principle of model predictive control is to use a model of the plant to predict and optimize its future behavior by repeatedly solving an on-line optimal control problem over a finite horizon [14, 15]. This control technique, also known as receding horizon control, has proven to be very successful in the process industry because it can handle large scale, multi-input multi-output nonlinear systems that are subject to actuator and state constraints [24].

4.1. Model Reduction

The nonlinear state-space representation of the full-order PWR model can be written as:

$$\begin{aligned} \dot{x}(t) &= f_0\left(x(t),u(t),d(t)\right) \\ x(t_0) &= x_0, \end{aligned}$$

where $x \in \mathbb{R}^{37}$ is the full state vector, $x_0 \in \mathbb{R}^{37}$ is the initial state, $u \in \mathbb{R}^3$ are the control inputs, and $d \in \mathbb{R}$ is the exogenous disturbance. In the sequel, it is assumed that the state is perfectly measured. Since the model is intended to be used in a model predictive control framework, special attention is given to its computational complexity. In practice, many numerical simulation runs are performed to solve the optimal control problem at each time step. However, the multiple time scale behavior of the model, ranging from fast neutron to slow xenon dynamics, renders it numerically stiff [25]. Although the full-order model can be simulated using powerful integration solvers, such as [26] or [27], model reduction is still conducted to make the optimal control problem more tractable. State estimation, which is outside the scope of the paper, is another motivation for reducing the order of the model. In the future, a state observer should be designed to implement the NMPC controller [28]. This task would obviously be

easier to achieve with a reduced-order model. Therefore, singular perturbation theory [29] is employed to remove the dynamics that are not essential regarding the control objectives. A well-known approximation in nuclear reactor physics is the prompt-jump approximation [17, 18]. Whenever a fission reaction occurs, two types of neutrons are emitted: prompt neutrons, which are directly emitted from fission, and delayed neutrons, which are emitted a moment later from the radioactive decay of fission products. Under normal operating conditions, the reactor core is predominantly governed by the dynamic of delayed neutrons. Consequently, prompt neutron dynamics can be reasonably approximated by steady-state equations without any significant loss of accuracy [30]. The resulting reduced-order model is given by a set of semi-explicit differential-algebraic equations:

$$\begin{aligned} \dot{x}_d(t) &= f\left(x_d(t), x_a(t), u(t), d(t)\right) \\ 0 &= g\left(x_d(t), x_a(t), u(t), d(t)\right) \\ x_d(0) &= x_{d0}, \end{aligned}$$

where the state is divided into a slow differential part $x_d \in \mathbb{R}^{31}$ and a fast algebraic part $x_a \in \mathbb{R}^6$.

4.2. Formulation of the Optimal Control Problem

At each time step, the NMPC algorithm yields a sequence of control inputs by solving a finite horizon optimal control problem that is initialized with the current state of the plant. Usually, only the first component of the control input sequence is applied open-loop to the plant until a new one is calculated at the next time step. The greatest strength of NMPC is that all the previously defined core control specifications can be directly incorporated into the finite horizon optimal control problem:

$$\min_{\overline{u}(\cdot)} \int_{t}^{t+NT_s} \left(\left\| e\left(\overline{x}_d(\tau), \overline{x}_a(\tau), \overline{d}(\tau)\right) \right\|_Q^2 + \left\| \overline{u}(\tau) \right\|_R^2 \right) \mathrm{d}\tau$$

subject to:

$$\forall \tau \in [t, t+NT_s] \begin{cases} c\left(\overline{x}_d(\tau), \overline{x}_a(\tau), \overline{u}(\tau), \overline{d}(\tau)\right) \leq 0 \\ \dot{\overline{x}}_d(\tau) = f\left(\overline{x}_d(\tau), \overline{x}_a(\tau), \overline{u}(\tau), \overline{d}(\tau)\right) \\ 0 = g\left(\overline{x}_d(\tau), \overline{x}_a(\tau), \overline{u}(\tau), \overline{d}(\tau)\right) \\ \overline{x}_d(t) = x_d(t), \end{cases} \tag{6}$$

where $N \geq 1$ is the length of the prediction horizon, $T_s > 0$ is the sampling time of the controller, and $x_d(t)$ is the differential state of the plant that is measured at time $t \geq 0$. The core control specifications (1) – (4) are represented by the inequality constraints map $c : \mathbb{R}^{31} \times \mathbb{R}^6 \times \mathbb{R}^3 \times \mathbb{R} \to \mathbb{R}^{16}$ and the tracking error (5) is given by the nonlinear map $e : \mathbb{R}^{31} \times \mathbb{R}^6 \times \mathbb{R} \to \mathbb{R}^3$. The weighted 2-norms $\|e\|_Q$ and $\|u\|_R$ are defined by $\|e\|_Q^2 := e^T Q e$ and $\|u\|_R^2 := u^T R u$ where $Q \in \mathbb{R}^{3\times3}$ and $R \in \mathbb{R}^{3\times3}$ are symmetric positive-definite matrices. The bar subscript denotes internal model variables that are used within the controller to predict the behavior of the real plant. In particular, the disturbance signal is written as $d(t) = \overline{d}(t) + \delta d(t)$ where $\overline{d}$ is the known load profile that is sent in advance by the grid dispatcher to the plant and $|\delta d| \leq 2.5$ %NP is the unknown norm-bounded term that appears when frequency control is active.

Transcription of the continuous optimal control problem into a finite-dimensional nonlinear programming problem is commonly achieved using either direct single-shooting or direct multiple-shooting [31, 32]. Here, both options will be implemented and tested thereafter. In a direct method, the continuous-time control inputs are approximated with a finite number of parameters. Most of the time, the control input sequence $\overline{\mathbf{u}} = \{\overline{u}[0], \ldots, \overline{u}[N-1]\}$ is discretized with a piecewise constant parameterization:

$$\forall k \in [\![0, N-1]\!],\ \forall \tau \in [t_k, t_{k+1}],\ \overline{u}(\tau) = \overline{u}[k],$$

where $t_k = t + kT_s$. The key difference between single-shooting and multiple-shooting is the parameterization of the state trajectory.

In single-shooting, the state trajectory is computed in one shot, from t_0 to t_N, outside of the optimization problem. In fact, for a given control input sequence $\overline{\mathbf{u}}$ and a fixed initial state $x_d(t)$, the state trajectory can be entirely determined by simulating the model of the plant:

$$\forall k \in [\![0, N-1]\!],\ \begin{pmatrix} \overline{x}_d(t_{k+1}) \\ \overline{x}_a(t_{k+1}) \end{pmatrix} = \phi_{ss}\left(t_{k+1}; x_d(t), \overline{\mathbf{u}}\right),$$

where $\phi_{ss} : \mathbb{R} \times \mathbb{R}^{31} \times \mathbb{R}^{3\times N} \to \mathbb{R}^{37}$ is a numerical integration operator that returns the solution of the initial value problem:

$$\forall \tau \in [t_0, t_N] \begin{cases} \dot{\overline{x}}_d(\tau) = f\left(\overline{x}_d(\tau), \overline{x}_a(\tau), \overline{u}(\tau), \overline{d}(\tau)\right) \\ 0 = g\left(\overline{x}_d(\tau), \overline{x}_a(\tau), \overline{u}(\tau), \overline{d}(\tau)\right) \\ \overline{x}_d(t_0) = x_d(t), \end{cases}$$

at time t_{k+1}. Therefore, the optimization problem can be reduced to:

$$\min_{\overline{\mathbf{u}}} \sum_{k=0}^{N-1} \left\| e\left(\phi_{ss}(t_{k+1}; x_d(t), \overline{\mathbf{u}}), \overline{d}(t_{k+1})\right) \right\|_Q^2 + \left\| \overline{u}[k] \right\|_R^2$$

subject to:

$$\forall k \in [\![0, N-1]\!],\ c\left(\phi_{ss}(t_{k+1}; x_d(t), \overline{\mathbf{u}}), \overline{u}[k], \overline{d}(t_{k+1})\right) \leq 0.$$

Hence, the advantage of using single-shooting is to remove the equality constraints from the optimization problem. Moreover, the state trajectory is always guaranteed to be consistent with the dynamical equations of the model. However, using single-shooting can increase the computational complexity of the optimization problem because the state equations are repeatedly propagated into both the cost function and the constraints. This is especially true for nonlinear and/or unstable systems.

In multiple-shooting, the state trajectory is cut into N pieces which are independently computed over each

time interval $[t_k, t_{k+1}]$. These pieces are then reassembled within the optimization problem by imposing equality constraints on the differential and algebraic states:

$$\min_{\overline{\mathbf{x}}_\mathbf{d}, \overline{\mathbf{x}}_\mathbf{a}, \overline{\mathbf{u}}} \sum_{k=0}^{N-1} \left\| e\left(\overline{x}_d[k+1], \overline{x}_a[k], \overline{d}(t_{k+1})\right)\right\|_Q^2 + \|\overline{u}[k]\|_R^2$$

subject to:

$$\forall k \in [\![0, N-1]\!],$$

$$\begin{cases} c\left(\overline{x}_d[k+1], \overline{x}_a[k], \overline{u}[k], \overline{d}(t_{k+1})\right) \leq 0 \\ \begin{pmatrix} \overline{x}_d[k+1] \\ \overline{x}_a[k] \end{pmatrix} - \phi_{ms}\left(t_{k+1}; \overline{x}_d[k], \overline{u}[k]\right) = 0 \\ \overline{x}_d[0] - x_d(t) = 0, \end{cases}$$

where $\phi_{ms} : \mathbb{R} \times \mathbb{R}^{31} \times \mathbb{R}^6 \to \mathbb{R}^{37}$ is a numerical integration operator that returns the solution of the initial value problem:

$$\forall \tau \in [t_k, t_{k+1}] \begin{cases} \dot{\overline{x}}_d(\tau) = f\left(\overline{x}_d(\tau), \overline{x}_a(\tau), \overline{u}(\tau), \overline{d}(\tau)\right) \\ 0 = g\left(\overline{x}_d(\tau), \overline{x}_a(\tau), \overline{u}(\tau), \overline{d}(\tau)\right) \\ \overline{x}_d(t_k) = x_d(t_k), \end{cases}$$

at time t_{k+1}. The drawback of multiple-shooting is that the dimension of the optimization problem is substantially increased since the differential and algebraic state sequences $\overline{\mathbf{x}}_\mathbf{d} = \{\overline{x}_d[0], \ldots, \overline{x}_d[N]\}$ and $\overline{\mathbf{x}}_\mathbf{a} = \{\overline{x}_a[0], \ldots, \overline{x}_a[N-1]\}$ become decision variables. However, this is not a real issue because the structure of the resulting nonlinear program is block-sparse, provided that the decision variables are arranged as follows:

$$\left(\overline{x}_d[0]^T, \overline{u}[0]^T, \overline{x}_a[0]^T, \ldots, \overline{x}_d[N-1]^T, \overline{u}[N-1]^T, \overline{x}_a[N-1]^T, \overline{x}_d[N]^T\right)^T.$$

In fact, it has often been reported that efficient sparsity exploitation can actually improve the rate of convergence of the Newton-type algorithm which is used to solve the nonlinear program [33]. Furthermore, having $\overline{\mathbf{x}}_\mathbf{d}$ and $\overline{\mathbf{x}}_\mathbf{a}$ as decision variables is beneficial for initializing the problem since *a priori* information about the state trajectory can be supplied to the algorithm. Last but not least, multiple-shooting is better suited for highly nonlinear and/or unstable plants because integration is broken down into short time intervals, meaning that the nonlinearity is equally distributed over the shooting nodes.

4.3. Practical Implementation Issues

The nonlinear program that is transcribed using single shooting is composed of $3N$ decision variables and $16N$ inequality constraints. On the other hand, the nonlinear program that is transcribed using multiple-shooting is composed of $40N+31$ decision variables, $37N+31$ equality constraints and $16N$ inequality constraints. In both cases, it is expected that the resolution time of the optimization problem becomes greater than the duration T_s in between two sampling instants because the length $N \geq 1$ of the prediction horizon has to be large enough to ensure closed-loop stability. Yet, we still want to recalculate the control inputs at each time step so the plant does not remain in open-loop for too long. Therefore, the optimization routine is deliberately terminated once the resolution time exceeds T_s seconds. Even if the solution is sub-optimal [34], we believe that it should have sufficiently converged to provide an acceptable level of closed-loop performance.

Consequently, a prediction-based delay compensation scheme [35] is incorporated into the NMPC algorithm to account for the non-negligible resolution time $0 < \tau \leq T_s$ of the optimization problem. To illustrate, let $\mu_N(x_d(t))$ be the feedback value given by the (sub-optimal) control input sequence $\overline{\mathbf{u}}^*(x_d(t))$ that (almost) solves the optimal control problem with initial state $x_d(t)$. Because of the computational delay, $\mu_N(x_d(t))$ is applied to the plant at time $t+\tau$ instead of time t. Thus, the greater the mismatch between $x_d(t)$ and $x_d(t+\tau)$, the greater the difference between the expected and the real behavior of the closed-loop system. Knowing the maximum resolution time $\tau_{max} = T_s$ of the optimization problem, the idea of [35] is to solve the optimal control problem in advance, starting from an estimate $\overline{x}_d(t+T_s)$ of the future state in place of the currently measured state $x_d(t)$. The resulting feedback value $\mu_N(\overline{x}_d(t+T_s))$, available at time $t+\tau$, is then applied to the plant accordingly at time $t+T_s$. The next state estimate can be computed straightforwardly by simulating the model of the plant from the newly measured state $x_d(t+T_s)$ with input $\mu_N(\overline{x}_d(t+T_s))$. Note that, since the state is estimated in an open-loop fashion, it is almost certain that the real state $x_d(t+T_s)$ will differ from $\overline{x}_d(t+T_s)$. However, it is still relevant to use the delay-compensated input, as some degree of robustness is expected against estimation errors [36].

5. Simulation Results

5.1. Validation Scenario

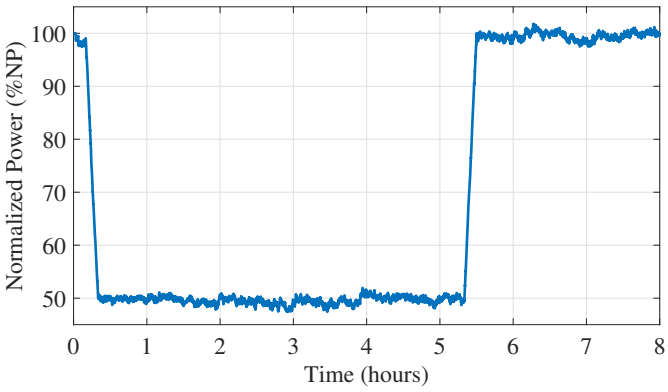

Fig. 3: Normalized power of the turbine with frequency control

To assess and compare the performance of the two core control systems, a typical 8-hour load-following

scenario is conducted on PWRSimu, an intermediate complexity PWR simulator developed by Framatome in MATLAB®/Simulink®. As shown in Fig. 3, this scenario involves two power variations (from 100 %NP to 50 %NP and *vice versa*) that are carried out at a rate of 5 %NP/min with frequency control. The full-order nonlinear PWR model is simulated on Simulink with the stiff adaptive-step integration solver `ode15s` [27].

The fixed-structure gain-scheduled controller has been tuned with the following set of parameters:

$$\begin{array}{lll} \beta_1 = 10/1.5 & \beta_2 = 5/5 & \beta_3 = 2/30 \\ W_{11} = 10/1.5 & W_{12} = 10/5 & W_{13} = 10/30 \\ W_{21} = 10/1.25 & W_{22} = 10/1.25 & W_{23} = 10/0.04 \\ \gamma = 0.5 & \alpha = 0.0001, & \end{array}$$

using `Systune` (MATLAB Control System Toolbox™) [21]. In order to model the whole closed-loop system in Simulink, the quadratic polynomials $K_P(\sigma)$, $K_I(\sigma)$, $a(\sigma)$ and $b(\sigma)$ are approximated by feeding the scheduling variable $\sigma = P_{bank}$ into 1-D lookup table blocks with linear interpolation and clip extrapolation options.

Tuning of the NMPC controller has been achieved in two steps. First, the matrices Q and R of the cost function (6) are selected diagonal with normalized weights:

$$\begin{array}{lll} q_1 = \dfrac{\varepsilon_1}{(e_{1,max})^2} & q_2 = \dfrac{\varepsilon_2}{(e_{2,max})^2} & q_3 = \dfrac{\varepsilon_3}{(e_{3,max})^2} \\ r_1 = \dfrac{\nu_1}{(u_{1,max})^2} & r_2 = \dfrac{\nu_2}{(u_{2,max})^2} & r_3 = \dfrac{\nu_3}{(u_{3,max})^2}, \end{array}$$

where, for all $i \in \{1,2,3\}$, $e_{i,max}$ and $u_{i,max}$ denote the maximum acceptable values of e_i and u_i (see Bryson's rule [37]). Then, the tunable parameters $(\varepsilon_1, \varepsilon_2, \varepsilon_3) \in \mathbb{R}^3_{>0}$ and $(\nu_1, \nu_2, \nu_3) \in \mathbb{R}^3_{>0}$ are used to penalize, in order of importance, the deviations of ACT, AO and P_{bank} without unduly restricting the control effort. In the end, the weighting matrices are given by:

$$Q = \mathrm{diag}\left(\frac{30}{1.5^2}, \frac{20}{5^2}, \frac{15}{30^2}\right), \; R = \mathrm{diag}\left(\frac{0.01}{1.25^2}, \frac{0.01}{1.25^2}, \frac{0.01}{0.04^2}\right).$$

The sampling time $T_s = 60$ s and the prediction horizon $N = 10$ are chosen so as to balance on-line complexity and closed-loop stability. The NMPC algorithm is implemented in MATLAB R2019a with CasADi v3.5.5 [38]. This open-source software provides several building blocks which are especially useful for solving large-scale nonlinear programing problems efficiently. Practically speaking, the key asset of CasADi is its ability to quickly compute Jacobians and exploit their sparsity pattern using algorithmic differentiation and graph coloring techniques. Within this framework, the reduced-order PWR model is simulated with the stiff adaptive-step integration solver `IDAS` [26] and the optimization problem is solved with the primal-dual interior point solver `IPOPT` [39] using an Intel® Core™ i3-6100U processor with 16 GB of RAM. The closed-loop interactions between the controller and the plant are represented by the sequential execution of both MATLAB® and Simulink® files.

5.2. Simulation Results and Discussion

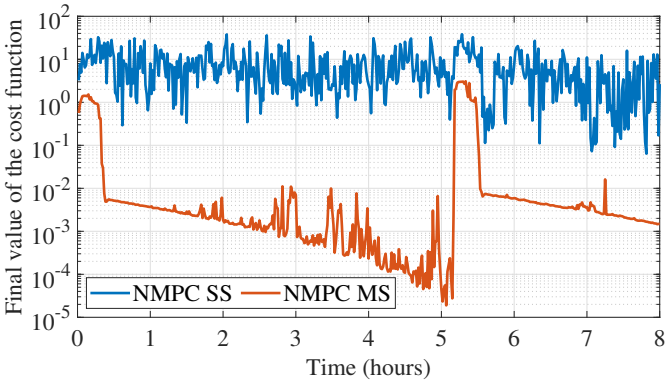

Fig. 4: Semi-log plot of the final value of the cost function (NMPC single-shooting vs NMPC multiple-shooting)

In order to choose between single-shooting and multiple-shooting, both transcription methods were implemented and tested. As can be seen in Fig. A.7 and A.9, the ACT deviation and the position deviation of P_{bank} are almost identical for both methods. However, the AO deviation, shown in Fig. A.8, is significantly smaller in multiple-shooting because the behavior of H_{bank}, shown in Fig. A.11, is less erratic than in single-shooting. The same phenomenon can be observed, but to a lesser extent, for P_{bank} in Fig. A.10 and for the boron concentration in Fig. A.12. The reason behind this, as Fig. 4 demonstrates, is that the optimization problem converges more quickly to the solution when transcribed in multiple-shooting. In fact, it turned out that further increasing the resolution time in multiple-shooting had no effect on the final value of the cost function. This suggests that the optimal solution can always be found within the allocated time, which is why multiple-shooting is selected over single-shooting.

Overall, the performance of the NMPC controller is better than that of the gain-scheduled controller. As can be seen in Fig. A.13, A.14 and A.15, both controllers manage to keep the ACT, the AO and the position of P_{bank} close to their reference signals. Yet, as can be observed in Fig. A.16, A.17 and A.18, the gain-scheduled controller tends to overuse the actuators in comparison with the NMPC controller. This shows that the bandwidth of the controller does not need to be particularly large even when

Table 1: Comparison between the gain-scheduled controller and the nonlinear model predictive controller

GSC	NMPC
Several linear models	One nonlinear model
Frequency domain design	Time domain design
Output-feedback control	State-feedback control
Robust design	Nominal design
Pre-computed gains	Online calculations
Continuous-time control	Discrete-time control
No anticipation	Preventive actions

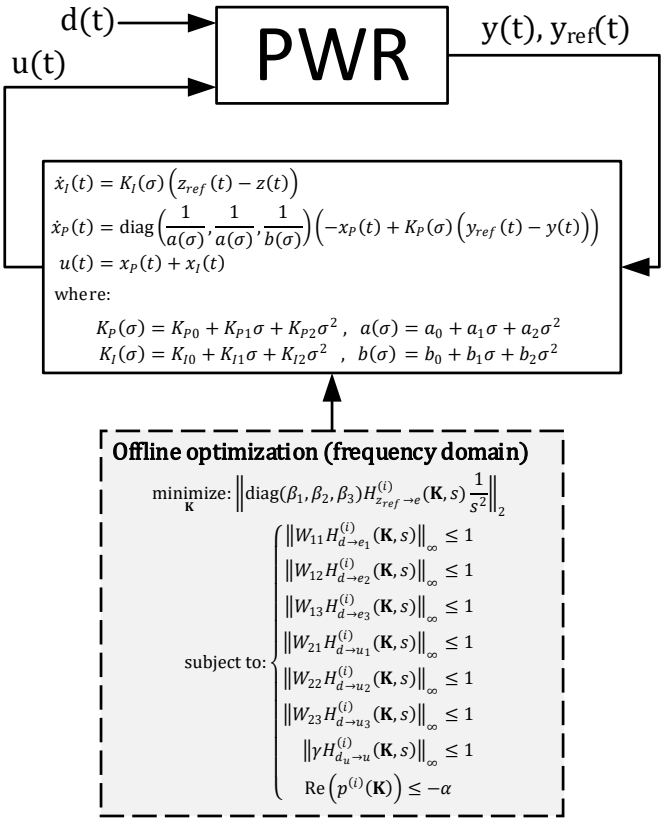

Fig. 5: Fixed-structure gain-scheduled control approach

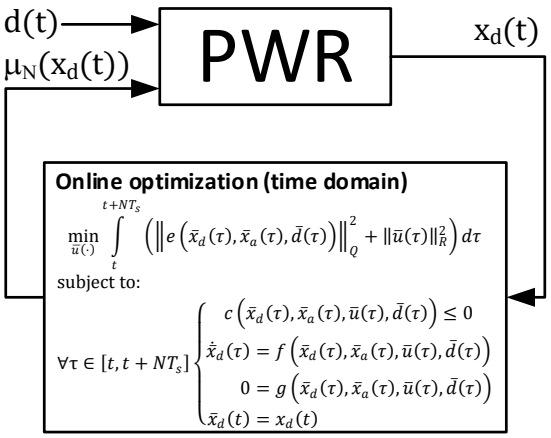

Fig. 6: Nonlinear model predictive control approach

frequency control is active. However, the NMPC controller may struggle to reject other sudden and unexpected disturbances, as the control inputs are only updated every minute. To give a concrete example, note that the principal effect of frequency control is to cause ACT fluctuations of about ±0.8 °C. A closer look at Fig. A.13 reveals that these fluctuations cannot be properly mitigated by the NMPC controller because they are unpredictable and of relatively high frequency. The gain-scheduled controller, on the other hand, gives better results since it provides continuous feedback to the plant and was designed using robust control methods. As a matter of fact, robustness is still an open issue in the field of model predictive control [15]. Therefore, it should be mentioned that the reduced-order PWR model, which is embedded in the NMPC controller, derives from the one that is used for validation. Hence, it is not surprising that the NMPC controller performs better than the gain-scheduled controller, which was designed off-line based on simpler linear models. This is both a strength and a weakness of nonlinear model predictive control. While it is indeed better to use a very detailed and accurate model for prediction, it can also become an issue if the computation time of the NMPC control law gets too long. Worse, using an excessively complex model can even make the optimization problem intractable. Besides, it is also good to recall that the NMPC controller is a state-feedback controller whereas the gain-scheduled controller is an output-feedback one. Unfortunately, some state variables, such as the xenon concentration of the reactor core, cannot be physically measured. Thus, the NMPC controller cannot be practically implemented unless a nonlinear state observer is designed. This not only increases the overall complexity of the core control system but also raises additional concerns about its robustness.

6. Conclusion

In this paper, two advanced core control systems have been designed for flexible operation of future PWRs. The first one is a fixed-structure H_2/H_∞ gain-scheduled controller and the second one is an NMPC controller. The design methodologies of both controllers have been detailed and their performances have been compared on the intermediate complexity pressurized water reactor simulator of Framatome, named PWRSimu. The gain-scheduled controller is well adapted to cope with unexpected disturbances and model uncertainties. Moreover, its architecture is readable and easy to implement. However, some improvements should be made to avoid overusing the actuators (a dead-zone nonlinearity could be taken into account in the design of the controller using, e.g., the circle criterion). By contrast, the NMPC controller can handle a wide range of constraints and can easily deal with disturbances that are known in advance. Its predictive capabilities make it possible to achieve good performance while limiting the control effort. However, it is very difficult to analyze its robustness, especially when used together with a state-observer. In future work, the controllers will be combined in a hierarchical manner to overcome the limitations of both design strategies. The NMPC controller could be located at an upper level to provide a feedforward action to the plant as well as reference output trajectories to the gain-scheduled controller. Then, the gain-scheduled controller could be used at a lower level to reject unknown disturbances and to ensure that the closed-loop system stays close to the reference trajectory despite model uncertainties. This hierarchical scheme is also motivated by

the multiple time scales of the plant. In fact, the upper layer could anticipate future load variations and forecast the evolution of xenon concentration while the lower layer could simply focus on disturbance rejection and setpoint tracking. It would also make sense to leave the calculation of boron concentration adjustments to the upper layer since the evolution of the boron concentration is substantially slower than the movements of the control rods.

Conflict of interest

None declared.

Acknowledgment

The authors would like to thank Framatome for the financial support.

Appendix A. Simulation Results

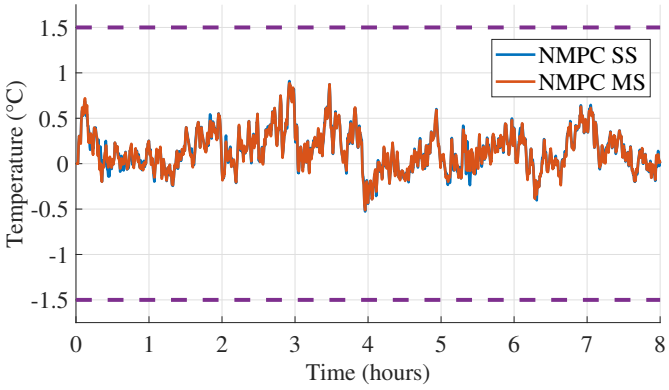

Fig. A.7: Average coolant temperature deviation (NMPC single-shooting vs NMPC multiple-shooting)

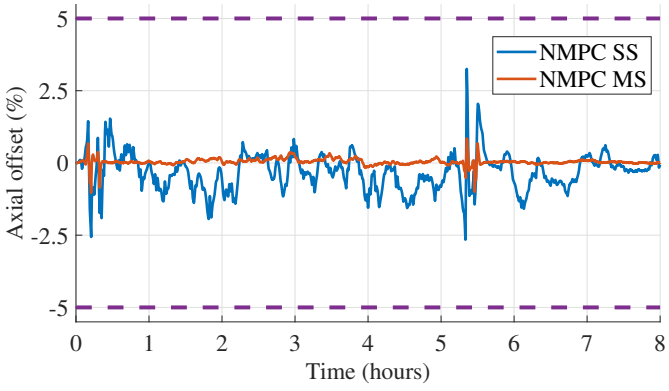

Fig. A.8: Axial offset deviation (NMPC single-shooting vs NMPC multiple-shooting)

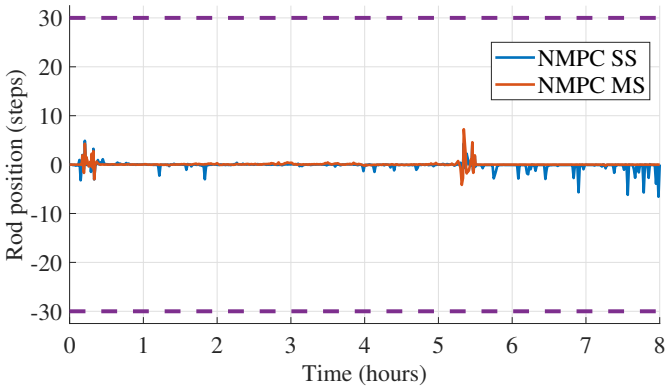

Fig. A.9: Position deviation of P_{bank} (NMPC single-shooting vs NMPC multiple-shooting)

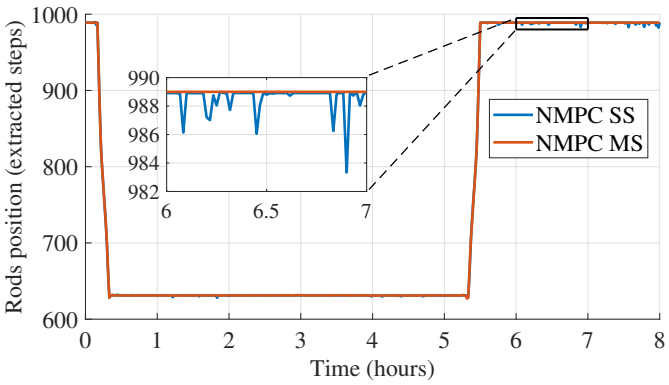

Fig. A.10: Position of P_{bank} (NMPC single-shooting vs NMPC multiple-shooting)

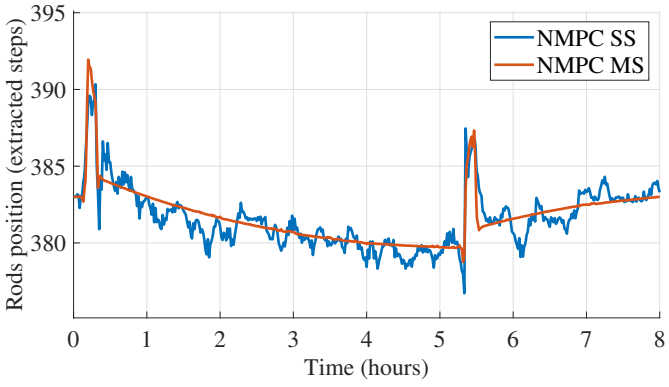

Fig. A.11: Position of H_{bank} (NMPC single-shooting vs NMPC multiple-shooting)

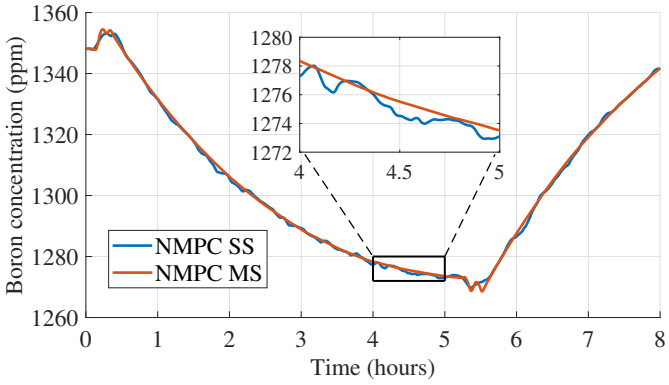

Fig. A.12: Boron concentration of the primary coolant (NMPC single-shooting vs NMPC multiple-shooting)

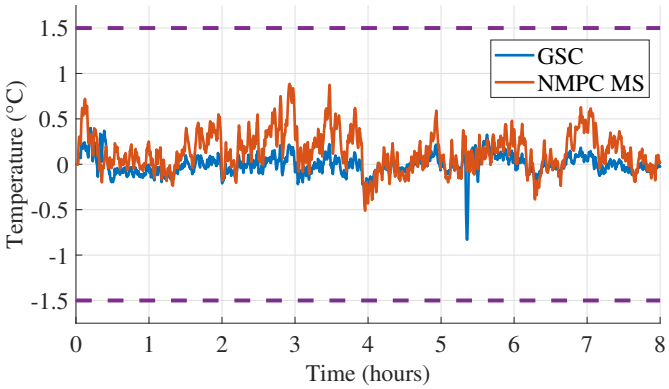

Fig. A.13: Average coolant temperature deviation (Gain-Scheduled Controller vs NMPC multiple-shooting)

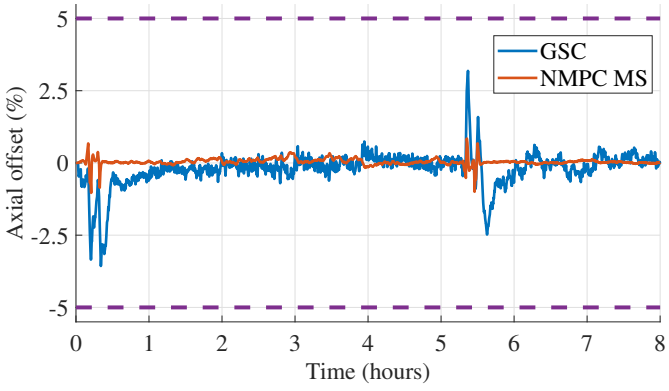

Fig. A.14: Axial offset deviation (Gain-Scheduled Controller vs NMPC multiple-shooting)

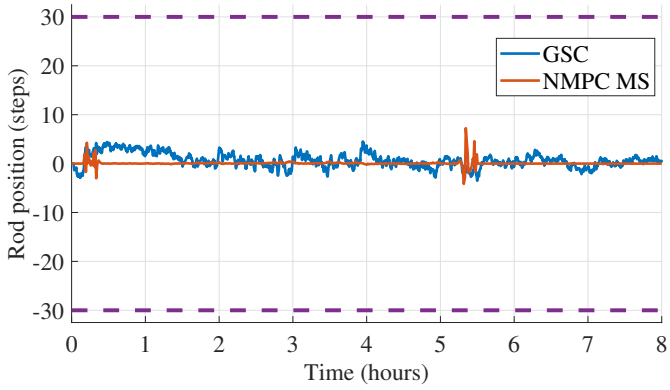

Fig. A.15: Position deviation of P_{bank} (Gain-Scheduled Controller vs NMPC multiple-shooting)

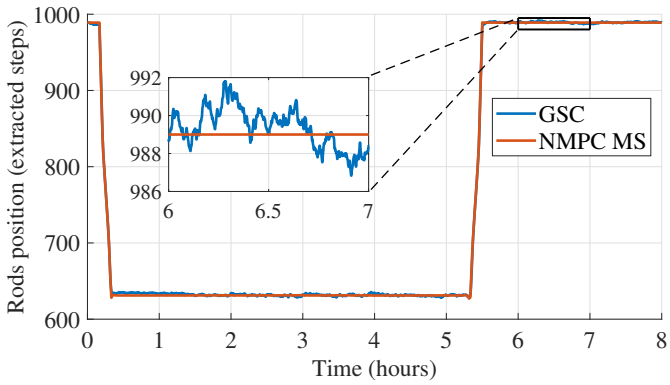

Fig. A.16: Position of P_{bank} (Gain-Scheduled Controller vs NMPC multiple-shooting)

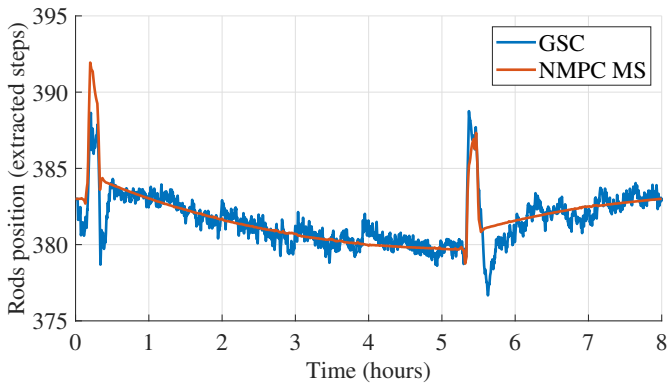

Fig. A.17: Position of H_{bank} (Gain-Scheduled Controller vs NMPC multiple-shooting)

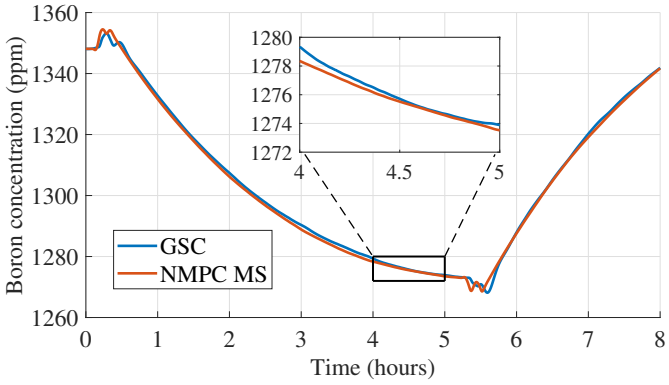

Fig. A.18: Boron concentration of the primary coolant (Gain-Scheduled Controller vs NMPC multiple-shooting)

References

[1] K. Kosowski, F. Diercks, Quo vadis, grid stability? Challenges increase as generation portfolio changes, atw International Journal for Nuclear Power 66 (2) (2021) 16–26.

[2] P. Morilhat, S. Feutry, C. Lemaitre, J. M. Favennec, Nuclear power plant flexibility at EDF, atw International Journal for Nuclear Power 64 (3) (2019) 131–140.

[3] IAEA, Non-baseload Operation in Nuclear Power Plants: Load Following and Frequency Control Modes of Flexible Operation, no. NP-T-3.23 in Nuclear Energy Series, IAEA Viena, 2018.

[4] J.-L. Mourlevat, Évolution des modes de pilotage, Revue générale nucléaire (3) (2007) 29–36.

[5] A. Grossetête, ATMEA1 & EPR™ mode T core control innovative features for high operating flexibility, Transactions of the American Nuclear Society 111 (1) (2014) 1095–1098.

[6] G. Li, X. Wang, B. Liang, X. Li, B. Zhang, Y. Zou, Modeling and control of nuclear reactor cores for electricity generation: A review of advanced technologies, Renewable and Sustainable Energy Reviews 60 (2016) 116–128.

[7] L. Lemazurier, M. Yagoubi, P. Chevrel, A. Grossetête, Multi-objective H_2/H_∞ gain-scheduled nuclear core control design, IFAC-PapersOnLine 50 (1) (2017) 3256–3262.

[8] L. Lemazurier, P. Chevrel, A. Grossetête, M. Yagoubi, An alternative to standard nuclear core control using a multi-objective approach, Control Engineering Practice 83 (2019) 98–107.

[9] G. Dupré, A. Grossetête, P. Chevrel, M. Yagoubi, Enhanced flexibility of PWRs (mode A) using an efficient NMPC-based boration/dilution system, in: 2021 European Control Conference (ECC), IEEE, 2021, pp. 1092–1098.

[10] M. G. Na, D. W. Jung, S. H. Shin, J. W. Jang, K. B. Lee, Y. J. Lee, A model predictive controller for load-following operation of PWR reactors, IEEE Transactions on Nuclear Science 52 (4) (2005) 1009–1020.

[11] J. H. Kim, S. H. Park, M. G. Na, Design of a model predictive load-following controller by discrete optimization of control rod speed for PWRs, Annals of Nuclear Energy 71 (2014) 343–351.

[12] H. Eliasi, M. Menhaj, H. Davilu, Robust nonlinear model predictive control for a PWR nuclear power plant, Progress in Nuclear Energy 54 (1) (2012) 177–185.

[13] P. Gahinet, P. Apkarian, Automated tuning of gain-scheduled control systems, in: 52nd IEEE Conference on Decision and Control, IEEE, 2013, pp. 2740–2745.

[14] D. Q. Mayne, J. B. Rawlings, C. V. Rao, P. O. Scokaert, Constrained model predictive control: Stability and optimality, Automatica 36 (6) (2000) 789–814.

[15] D. Q. Mayne, Model predictive control: Recent developments and future promise, Automatica 50 (12) (2014) 2967–2986.

[16] L. Lemazurier, Conception d'un système avancé de réacteur PWR flexible par les apports conjoints de l'ingénierie système et de l'automatique, Ph.D. thesis, Institut Mines-Télécom Atlantique (2018).

[17] P. Reuss, Précis de neutronique, EDP sciences, 2012.

[18] W. M. Stacey, Nuclear reactor physics, 3rd Edition, John Wiley & Sons, 2018.

[19] D. J. Leith, W. E. Leithead, Survey of gain-scheduling analysis and design, International journal of control 73 (11) (2000) 1001–1025.

[20] W. J. Rugh, J. S. Shamma, Research on gain scheduling, Automatica 36 (10) (2000) 1401–1425.

[21] P. Apkarian, D. Noll, A. Rondepierre, Mixed H_2/H_∞ control via nonsmooth optimization, SIAM Journal on Control and Optimization 47 (3) (2008) 1516–1546.

[22] M. Bouffier, L. Daudin, G. Rio, Science, an advanced 3D fuel management code package, in: TopFuel'97, Thomas Telford Publishing, 1997, pp. 5–30.

[23] B. A. Francis, W. M. Wonham, The internal model principle of control theory, Automatica 12 (5) (1976) 457–465.

[24] S. J. Qin, T. A. Badgwell, A survey of industrial model predictive control technology, Control engineering practice 11 (7) (2003) 733–764.

[25] T. T. Hartley, G. O. Beale, S. P. Chicatelli, Digital simulation of dynamic systems: a control theory approach, Prentice-Hall, Inc., 1994.

[26] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, C. S. Woodward, SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers, ACM Transactions on Mathematical Software (TOMS) 31 (3) (2005) 363–396.

[27] L. F. Shampine, M. W. Reichelt, J. A. Kierzenka, Solving index-1 DAEs in MATLAB and Simulink, SIAM review 41 (3) (1999) 538–552.

[28] R. Findeisen, L. Imsland, F. Allgower, B. A. Foss, State and output feedback nonlinear model predictive control: An overview, European journal of control 9 (2-3) (2003) 190–206.

[29] P. Kokotović, H. K. Khalil, J. O'reilly, Singular Perturbation Methods in Control: Analysis and Design, SIAM, 1999.

[30] X. Chen, A. Ray, On singular perturbation of neutron point kinetics in the dynamic model of a PWR nuclear power plant, Sci 2 (2) (2020) 30.

[31] L. Grüne, J. Pannek, Nonlinear model predictive control: Theory and Algorithms, 2nd Edition, Springer, 2017.

[32] J. B. Rawlings, D. Q. Mayne, M. Diehl, Model predictive control: Theory, Computation, and Design, 2nd Edition, Nob Hill Publishing Madison, WI, 2019.

[33] M. Diehl, H. J. Ferreau, N. Haverbeke, Efficient numerical methods for nonlinear MPC and moving horizon estimation, in: Nonlinear model predictive control, Springer, 2009, pp. 391–417.

[34] L. Grüne, J. Pannek, Analysis of unconstrained NMPC schemes with incomplete optimization, IFAC Proceedings Volumes 43 (14) (2010) 238–243.

[35] R. Findeisen, F. Allgöwer, Computational delay in nonlinear model predictive control, IFAC Proceedings Volumes 37 (1) (2004) 427–432.

[36] R. Findeisen, L. Grüne, J. Pannek, P. Varutti, Robustness of prediction based delay compensation for nonlinear systems, IFAC Proceedings Volumes 44 (1) (2011) 203–208.

[37] A. E. Bryson, Y.-C. Ho, Applied optimal control: optimization, estimation, and control, Routledge, 1975.

[38] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, M. Diehl, CasADi: a software framework for nonlinear optimization and optimal control, Mathematical Programming Computation 11 (1) (2019) 1–36.

[39] A. Wächter, L. T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, Mathematical programming 106 (1) (2006) 25–57.