# Model-predictive kinetic control with data-driven models on EAST

D. Moreau, S. Wang, J.P. Qian, Q. Yuan, Y. Huang, Y. Li, S. Ding, H. Du, X. Gong, M. Li, et al.

OPEN ACCESS
IOP Publishing | International Atomic Energy Agency
Nuclear Fusion
Nucl. Fusion **64** (2024) 126069 (35pp)
https://doi.org/10.1088/1741-4326/ad893b

# Model-predictive kinetic control with data-driven models on EAST

**D. Moreau[1,*], S. Wang[2], J.P. Qian[2], Q. Yuan[2], Y. Huang[2], Y. Li[2], S. Ding[2], H. Du[2], X. Gong[2], M. Li[2], H. Liu[2], Z. Luo[2], L. Zeng[2], E. Olofsson[3], B. Sammuli[3], J.F. Artaud[1], A. Ekedahl[1] and E. Witrant[4]**

[1] CEA, IRFM, 13108 Saint-Paul-lez-Durance, France
[2] IPP, Chinese Academy of Sciences, Hefei 230031, China
[3] General Atomics, San Diego, CA 92186-5608, United States of America
[4] University Grenoble-Alpes, CNRS, GIPSA-Lab, 38000 Grenoble, France

E-mail: didier.moreau@cea.fr and didiermoreau.fr@orange.fr

Received 18 April 2024, revised 27 September 2024
Accepted for publication 21 October 2024
Published 7 November 2024

**Abstract**

In this work, model-predictive control (MPC) was combined for the first time with singular perturbation theory, and an original plasma kinetic control method based on extremely simple data-driven models and a two-time-scale MPC algorithm has been developed. A comprehensive review is presented in this paper. Slow and fast semi-empirical models are identified from data, by considering the fast kinetic plasma dynamics as a singular perturbation of a quasi-static equilibrium, which itself is governed, on the slow time scale, by the flux diffusion equation. This control technique takes advantage of the large ratio between the time scales involved in magnetic and kinetic plasma transport. It is applied here to the simultaneous control of the safety factor profile, $q(x)$, and of several kinetic variables, such as the poloidal beta parameter, $\beta_p$, and the internal inductance parameter, $l_i$, on the EAST tokamak. In the experiments, the available control actuators were lower hybrid current drive (LHCD) and co-current neutral beam injection (NBI) from different sources. Ion cyclotron resonant heating (ICRH) and electron cyclotron resonant heating (ECRH) are used as additional actuators in control simulations. In the controller design, an observer provides, in real time, an estimate of the system states and of the mismatch between measured and predicted outputs, which ensures robustness to model errors and offset-free control. Based on the observer information, the controller predicts the behavior of the system over a given time horizon and computes the optimal actuation by solving a quadratic programming optimization problem that takes the actuator constraints into account. A number of control applications are described in the paper, either in nonlinear simulations with EAST-like parameters or in real experiments on EAST. The simulations were performed with a fast plasma simulator (METIS) using either two control actuators (LHCD and ICRH) in a low density scenario, or up to four actuators at higher density: LHCD, ECRH, and two NBI systems driven in a on/off pulse-width-modulation (PWM) mode, with different injection angles. The control models are identified with the prediction-error method, using datasets obtained from open loop simulations in which the actuators are modulated with pseudo-random

* Author to whom any correspondence should be addressed.

binary sequences. The simulations with two actuators show that various $q(x)$ profiles and $\beta_p$ waveforms can be tracked without offset, within times that are consistent with the resistive and thermal diffusion time scales, respectively. In simulations with four actuators, simultaneous tracking of time-dependent targets is shown for $q(x)$ at two normalized radii, $x = 0$ and $x = 0.4$, and for $\beta_p$. Due to the inherent mismatch between the optimal NBI power request and the delivered PWM power, the kinetic controller performs with reduced accuracy compared with simulations that do not use the NBI/PWM actuators. The first experimental tests using this new control algorithm were performed on EAST when the only available actuator was the LHCD system at 4.6 GHz. The algorithm was thus used in its simplest single-input-single-output version to track time-dependent targets for the central safety factor, $q_0$, or for $\beta_p$. In the closed loop control experiments, the $q_0$ targets were tracked in about one second, consistently with the plasma resistive time constant. Excellent tracking of a piecewise linear $\beta_p$ target waveform was also achieved. When the NBI system became controllable in real time by the EAST plasma control system, new experiments were dedicated to multiple-input-multiple-output MPC control with three actuators: LHCD and two NBI actuators using the PWM algorithm. Given that the minimum time allowed between NBI on/off switching was 0.1 s, i.e. larger than the characteristic time of the fast plasma dynamics, a reduced version of the MPC controller based only on the slow model was used. Various controller configurations were tested during a single experimental session, with up to three controlled variables chosen among $q_0 = q(x = 0)$, $q_1 = q(x = 0.5)$, $\beta_p$ and $l_i$. The main difficulty encountered during this session was the unavailability of the full baseline ICRH and ECRH powers that were used in the reference scenario, and from which the plasma model was identified. This often led to the saturation of one or several actuators, which prevented some targets selected in advance from being accessible. Nevertheless, in cases that were free from actuator saturation, $q_0$ and $q_1$ targets were successfully reached, in a time that is consistent with the resistive diffusion time of the model and with small oscillations that are characteristic of the PWM operation of the neutral beams. During the simultaneous control of $q_0$ and $\beta_p$, the ICRH power was too low and, in addition, the plasma density was much larger than the reference one. The $q_0$ targets were not accessible in this high-density/low-power case, but $\beta_p$ control was successful. Finally, the simultaneous control of $q_0$ and $l_i$ was satisfactory and, during the simultaneous control of, $q_0$, $\beta_p$ and $l_i$, the tracking of $\beta_p$ and $l_i$ was satisfactory but $q_0$ was too large due to the lack of ICRH power and to NBI saturation. In conclusion, the extensive nonlinear simulations described in this paper have demonstrated the relevance of combining MPC, data-driven models and singular perturbation methods for plasma kinetic control. This technique was also assessed experimentally on EAST, although some tests were perturbed by undesired parameter changes with respect to the reference scenario.

Keywords: tokamaks, plasma control, kinetic control, profile control, model-predictive control, two-time-scale control, singular perturbation theory

(Some figures may appear in colour only in the online journal)

## Contents

1. Introduction 3
2. The *ARTAEMIS* semi-empirical state space models 5
3. Identification of *ARTAEMIS* models from nonlinear plasma simulations with EAST-like parameters 7
4. State observer, model errors and two-time-scale MPC controller design 10
   4.1. The augmented two-time-scale model 10
   4.2. The augmented model state observer 11
   4.3. The two-time-scale MPC controller design 13
   4.4. The reduced MPC kinetic controller design on the slow time scale 14
5. Simulations of $q(x)$ and $\beta_p$ control on EAST using LHCD and ICRH 15
   5.1. Control of the safety factor with LHCD 15
      5.1.1. Discrete $q_0$ control 15
      5.1.2. Distributed $q(x)$ control 16
   5.2. Control of $q(x)$ and $\beta_p$ with LHCD, ICRH and the reduced MPC controller 17
   5.3. Faster control of $\beta_p$ using the two-time-scale *ARTAEMIS* model and MPC controller 18
6. Initial SISO control experiments on EAST using the *ARTAEMIS* MPC controller 19
   6.1. *ARTAEMIS* model for SISO control experiments 20
   6.2. Experimental SISO control of $q_0$ with LHCD 20
   6.3. Experimental SISO control of $\beta_p$ with LHCD 20
7. Simulations of two-time-scale MPC control using LHCD, ECRH and NBI pulse width modulation 22
   7.1. Control of the safety factor profile 22

7.2. Simultaneous control of $q(x)$ and $\beta_p$ 23
8. MPC kinetic control experiments on EAST using LHCD and NBI pulse width modulation 24
8.1. Identification of *ARTAEMIS* models from experimental EAST data 24
8.2. Closed loop MIMO control experiments on EAST 25
8.2.1. Discrete $q_0$ control 27
8.2.2. Discrete control of $q_0$ and $q(x = 0.5)$ 28
8.2.3. Simultaneous control of $q_0$ and $\beta_p$ 28
8.2.4. Simultaneous control of $q_0$ and $l_i$ 28
8.2.5. Simultaneous control of $q_0$, $\beta_p$ and $l_i$ 29
9. Summary and conclusion 30
Acknowledgment 34
References 34

## 1. Introduction

The development of advanced tokamak operation scenarios in which an optimization of some plasma parameter profiles, at reduced plasma current, results in a large improvement in energy confinement and MHD stability [1–7] provides the physics basis for the design of a steady state fusion reactor based on the tokamak concept. With the plasma characteristics foreseen in such operation scenarios, a high-gain fusion burn should be achieved at high $\beta$ ($\beta$ is the ratio of thermal to magnetic pressure) while a major fraction of the toroidal current is self-generated by the neoclassical bootstrap effect [8]. Thus, the remaining fraction of the plasma current can also be driven non-inductively by external sources [9], which allows steady state operation or pulse lengths exceeding by far the intrinsic limit of conventional inductive tokamak operation. In present-day experiments, the high performance (high-$\beta$) phase is often limited in duration by the undesired evolution of some plasma parameter profiles and by MHD activity. Advanced control of magnetic and kinetic plasma parameters and radial profiles such as the safety factor profile ($q$-profile), the normalized beta parameter, $\beta_N$, the poloidal beta, $\beta_p$, and the internal inductance parameter, $l_i$, generally referred to as plasma kinetic control, will therefore be essential for the success of advanced steady state operation in ITER [10, 11] and ultimately for the development of nuclear fusion as an attractive source of energy.

Ideally, for integrated profile control involving several magnetic and kinetic parameters and profiles, first-principles nonlinear plasma transport models should be used, as they have, in principle, a universal domain of validity. However, despite active research in this field, available models of increasing complexity still depend on many uncertain or unknown parameters and transport coefficients, and their accuracy cannot be widely assessed, especially in the advanced H-mode operation scenarios and plasma regimes. Simplified versions of such models, together with some ad hoc assumptions or coefficients, have been used for real-time applications and controller designs. Early examples were the reconstruction of the plasma current density profile when real-time measurements are not available [12, 13], and actuator trajectory optimization during plasma current ramp-up to optimally reach a given point in the tokamak operating space [14]. First-principle-driven (FPD) models have also been used for controller synthesis in many plasma control numerical simulations using various integrated transport modeling codes. They were used, for instance, in ITER simulations with some assumptions to allow for real-time update of the profile response models [15]. They were also integrated in a model-based robust feedback control algorithm [16] and in a model-predictive control (MPC) algorithm [17], also for ITER simulations. FPD models have also been combined with Lyapunov-based distributed methods to control the $q$-profile on TORE SUPRA and TCV simulators [18, 19]. Non-linear Lyapunov-based distributed approaches addressed bootstrap current optimization [20] and electron temperature regulation in H-mode plasmas [21] using sum-of-squares polynomials. More recently, robust, nonlinear, model-based control of the current profile and of the plasma energy content or $\beta_N$ was tested in numerical simulations for the EAST tokamak [22] and for DIII-D [23], with control algorithms based on the first-principles magnetic flux diffusion equation combined with 'uncertain' models for the electron temperature, plasma resistivity and non-inductive current drive.

Plasma kinetic controllers based on FPD models have also been implemented and tested experimentally. For instance, a simplified FPD model has been used for the first time for model-based profile control experiments in DIII-D L-mode (low confinement) discharges [24]. Later, simulations and experimental tests of an MPC controller based on a similar simplified FPD model embedding some scaling laws and ad-hoc transport coefficients have also been performed in L-mode plasmas on the TCV tokamak [25]. Experimental results on the $q$-profile and $\beta$ parameter control on TCV were also obtained by designing a Lyapunov-based distributed controller using a two-time-scale approach [26]. More recently, experiments on $q$-profile and $\beta_N$ control were conducted in H-mode plasmas on EAST using feedback control algorithms with a proportional-integral-derivative (PID) structure in which the various gains were optimized using control-oriented FPD models [27].

In the present work, an alternative approach to integrated profile control has been pursued, which we have referred to as the *ARTAEMIS*[5] approach. Model-based controllers are designed starting from the same first-principles model structure (magnetic flux and kinetic diffusion equations), but without specifying complex mathematical expressions for every uncertain or unknown space-dependent coefficient in the equations, and without introducing approximate scaling laws and ad hoc formulas in their various source terms. Instead, this approach is based on the experimental identification [28]

[5] *ARTAEMIS* is an acronym for 'Advanced Real-Time Algorithms based on Empirical Modelling of Integrated Scenarios' that we have used to refer to the two-time-scale profile control algorithms and plasma models originally described in [30, 31] and based on applying the theory of singularly perturbed systems to a set of simplified and linearized plasma transport equations.

of a minimal dynamic plasma model, taking into account the physical structure and couplings of the flux-averaged transport equations (this includes, among many other things, the effect of thermal diffusion on the poloidal magnetic flux and on the bootstrap current), but making no quantitative assumptions on the transport coefficients or on their dependences. The system consists of coupled ordinary differential equations for slow and fast models that govern its evolution on the resistive and the kinetic time scales, respectively. The two-time-scale control-oriented model structure stems from a first order singular perturbation expansion [29] of the linearized plasma transport equations and a projection on appropriate radial basis functions [30]. In other words, in these models and in the associated control algorithms, the fast component of the kinetic plasma dynamics is considered as a singular perturbation of a quasi-static equilibrium, which itself is governed, on the resistive time scale, by the flux diffusion equation. Using singular perturbation methods takes advantage of the small ratio between the thermal and resistive diffusion time scales and provides a natural way to cope with the high dimensionality of the system. The system identification problem is thus made tractable by the partial decoupling of the slow and fast dynamics. For profile control purposes, this technique yields extremely simple data-driven models (hence low control computational cost), at the expense of a restricted applicability to a given device and actuator set, around a given reference plasma equilibrium. Despite this restriction, such models can be used with advanced control algorithms in a relatively broad plasma parameter space around the reference plasma state. The model identification process was originally developed using simulated and experimental data from JET [30] and was improved later with experimental data from JT-60U and DIII-D [31]. It is based on the output prediction-error method (PEM) [28] and specific modulation experiments with random excitation of the various actuators were realized for that purpose, either in plasma simulations or in real tokamak experiments. Similar estimation and identification methods for lumped plasma models of small dimensions were also developed using a combination of subspace and output-error methods [32].

Over the years, the resulting data-driven models were then used to synthesize and test different controllers, going from the simple PID to the more powerful MPC algorithms. The simplest algorithm that combines the two-time-scale linear model structure with linear-quadratic optimal control theory was the so-called near-optimal control algorithm [30]. It consists of a slow proportional-integral (PI) feedback with an anti-windup loop, augmented by a fast proportional feedback that regulates the kinetic variables on the fast time scale. The first experimental tests were performed on JET with the control of the $q$-profile, $q(\mathrm{x})$, at several normalized radii, x, using three heating and current drive (H&CD) actuators: lower hybrid current drive (LHCD), neutral beam injection (NBI) and ion cyclotron resonant heating (ICRH). Control was done either at constant plasma surface loop voltage, or by using the boundary poloidal magnetic flux as a fourth actuator and including the safety factor near the plasma edge in the controlled variables rather than regulating accurately the plasma current through the poloidal field coils. Further experimental tests were performed later on DIII-D high-$\beta_N$ discharges, showing simultaneous control of the internal poloidal flux profile, $\psi(\mathrm{x})$, and of $\beta_N$, as well as effective distributed control of the $q$-profile at normalized radii between $\mathrm{x} = 0$ and $\mathrm{x} = 0.6$ [33]. In these experiments, the profile control actuators were the surface loop voltage and a subset of five other independent H&CD sources: on-axis co-current NBI, off-axis co-current NBI, counter-current NBI, balanced NBI, and electron cyclotron resonant heating (ECRH) and current drive (ECCD). The effectiveness of the *ARTAEMIS* approach for the integrated control of $q(\mathrm{x})$ and $\beta_N$ in steady state fully non-inductive discharges was also investigated in closed loop nonlinear plasma transport simulations using the METIS plasma simulator [34] with DIII-D-like parameters. However, in these simulations [35], the desired steady state $q$-profiles were obtained in about ten seconds at high $\beta_N$, which is too slow compared with the real DIII-D pulse length at high power. With larger gains, the $q(\mathrm{x})$ targets were reached after a large undershoot in the plasma core and a damped oscillation.

Recent work was therefore dedicated to the development of more powerful algorithms incorporating the singular perturbation approach and the *ARTAEMIS* models. $H_\infty$ robust control is an option that was investigated in [36, 37] with METIS simulations and some preliminary experiments on EAST. In this paper, we concentrate on a controller design that combines the simplicity of the *ARTAEMIS* two-time-scale models with the efficiency of MPC techniques [38], which have been successfully used in industry for a long time [39]. In MPC theory, the actuator constraints can be easily embedded in the quadratic programming (QP) optimization process and the model errors can be identified and compensated in real time. The new controllers were designed in view of experimental applications on EAST for the control of the $q$-profile, and of scalar parameters such as $\beta_p$ and $l_i$ in a high-$\beta_p$ fully non-inductive H-mode scenario, at constant plasma current. They were first tested on nonlinear METIS simulations in the aim of tuning various free controller parameters and options. Then, they were implemented on the EAST plasma control system (PCS) and experiments were conducted. The real-time actuators available for the experiments were off-axis LHCD at 4.6 GHz and four co-current NBI sources. At present, the ECRH and ICRH systems can only deliver feedforward power waveforms but they could also be used as control actuators in the future, after the development of real-time actuation from the EAST PCS. They will be used here as additional actuators to extend the controller capabilities in some simulations. The *ARTAEMIS* system identification procedure was applied to experimental data obtained with random actuator modulations that are consistent with the actuator constraints, or to simulated data obtained from the METIS plasma simulator. In both cases, a single model can satisfactorily approximate the coupled response of $\psi(\mathrm{x})$, $q(\mathrm{x})$, $\beta_p$ and $l_i$ to relatively large random variations of the available H&CD actuators compared to their minimum and maximum values, hence with a large diversity of $q$-profile shapes in the

inner plasma. In principle, switching between various models identified around different plasma equilibria would be possible if the plasma regime changes significantly despite the controller regulatory action. This was not necessary for the experiments reported here.

The paper is organized as follows. In section 2, the choice of the relevant state variables and the two-time-scale structure of the *ARTAEMIS* state space models are presented. Then, in section 3, such a model is identified using synthetic data obtained from METIS nonlinear plasma simulations with EAST-like parameters in a typical scenario. In section 4, augmented models including additional disturbance states are introduced, and a state observer is defined for the real-time estimate of the model states and disturbances. The details of the full two-time-scale MPC control algorithm will also be given in this section as well as a reduced version of the kinetic controller that is limited to the slow time scale. In section 5, closed loop simulations performed by combining the METIS plasma simulator with the *ARTAEMIS* MPC controller and using LHCD and ICRH actuators will be described. Examples including discrete or distributed $q(\mathrm{x})$ control, with flat or monotonic $q$-profile targets, and with different $\beta_\mathrm{p}$ target waveforms will be discussed. In section 6, the first experimental tests on EAST, in which LHCD at 4.6 GHz was used to track different targets either for the central safety factor, $q_0$, or for $\beta_\mathrm{p}$, will be described. Section 7 will present the results of closed loop METIS simulations illustrating the simultaneous control of $q_0$, $q(\mathrm{x} = 0.4)$ and $\beta_\mathrm{p}$ with up to four actuators. In particular, the actuators include two co-current NBI systems with different injection angles and driven in an on/off pulse-width-modulation (PWM) mode. Finally, multiple-input-multiple output experiments using LHCD and two NBI/PWM groups were performed with up to three controlled variables among $q_0$, $q(\mathrm{x} = 0.5)$, $\beta_\mathrm{p}$ and $l_i$. They will be discussed in section 8 and general conclusions will be drawn in section 9.

## 2. The *ARTAEMIS* semi-empirical state space models

In a tokamak, the multiple, radially distributed, magnetic and kinetic parameters that define the plasma state are known to be strongly and nonlinearly coupled. Because of this linkage, the parameters and profiles that define a given target plasma state and need to be achieved and regulated in real-time may be reduced to a minimal set of essential ones such as, for example, the safety factor profile, $q(\mathrm{x})$, and the scalar parameter $\beta_\mathrm{N}$ or $\beta_\mathrm{p}$. The specification of $q(\mathrm{x})$ can even be restricted to the inner half of the plasma if the internal inductance plasma parameter, $l_\mathrm{i}$, is also given and regulated. Plasma rotation or temperature profiles may also play a role in the achievement of the desired plasma performance and could be included in the kinetic controlled parameters.

The general structure of the continuous-time *ARTAEMIS* models is postulated from a set of coupled plasma response equations that only depend on a normalized radial coordinate x and on time $t$, and which stem from the linearized flux-averaged plasma transport equations [30]:

$$\partial\psi(x,t)/\partial t = \mathrm{L}_{\psi\psi}\{x\}\cdot\psi(x,t) + \mathrm{L}_{\psi\kappa}\{x\}\cdot\kappa(x,t) + L_{\psi P}(x)\cdot P(t) + V_\mathrm{ext}(t) \quad (1)$$

$$\varepsilon\partial\kappa(x,t)/\partial t = \mathrm{L}_{\kappa\psi}\{x\}\cdot\psi(x,t) + \mathrm{L}_{\kappa\kappa}\{x\}\cdot\kappa(x,t) + L_{\kappa P}(x)\cdot P(t). \quad (2)$$

The system is linearized around a fixed equilibrium state[6] that is called the reference state, and which needs not be known explicitly before the system is identified (see appendix A2 in [30]). Here, the so-called internal poloidal magnetic flux function, $\psi(\mathrm{x},t)$, and a set of kinetic profiles and/or scalar parameters represented by the vector $\kappa(\mathrm{x},t)$, refer to differences with their reference values. They appear as the most natural state variables of the system for a state-space model description. The radial variable, x, is defined as $(\Phi/\Phi_\mathrm{max})^{1/2}$ where $\Phi(\mathrm{x})$ is the toroidal magnetic flux through the poloidal cross-section of a given flux surface, and $\Phi_\mathrm{max}$ its maximum value at the last closed flux surface where $\mathrm{x} = 1$. The internal poloidal flux is obtained by subtracting the plasma boundary flux, $\Psi_\mathrm{b}(t)$, from the total poloidal flux, so that $\psi(1,t) = 0$. It has been introduced in order to eliminate the continuous flux variation that drives the inductive plasma current, and to define system states that reach steady equilibrium values even when the plasma surface loop voltage, $-\mathrm{d}\Psi_\mathrm{b}(t)/\mathrm{d}t$, is not zero. All the unknown differential operators $\mathrm{L}_{\alpha\beta}\{\mathrm{x}\}$ and row vectors $L_{\alpha,P}$ that characterize the linear response of the system depend on x but are independent of time, and the inputs $P(t)$ and $V_\mathrm{ext}(t)$ contain the powers from the various available H&CD systems and the plasma surface loop voltage, respectively, reduced by their reference values.

The small constant parameter $\varepsilon$ ($\varepsilon \ll 1$) represents the typical ratio between the kinetic and the resistive diffusion time scales. It is introduced here to scale the operators $\mathrm{L}_{\psi,\psi}\{\mathrm{x}\}$ and $\mathrm{L}_{\kappa,\kappa}\{\mathrm{x}\}$ so that their largest negative eigenvalues have similar absolute values. As the order of magnitude of $\varepsilon$ is about 0.05 in present-day tokamaks and 0.001 in ITER, the use of singular perturbation methods [29] is quite relevant for model identification and control design. This amounts to expanding each dependent variable in powers of $\varepsilon$, defining an additional independent variable, $\tau = t/\varepsilon$, to describe the fast dynamics while $t$ describes the slow dynamics, and to splitting variables into a sum of a fast and a slow component which depend on $\tau$ and $t$, respectively. A well-posed set of ordinary differential equations is then obtained by grouping terms of equal order in $\varepsilon$ and imposing that, in the asymptotic limit where $\varepsilon$ tends to zero, the initial conditions for the slow dynamics ($t = 0$) must match the quasi-steady-state solution on the fast dynamics ($\tau \to \infty$).

[6] An equilibrium state is defined as a stationary state of the plasma, i.e. a state where the thermodynamic variables are in equilibrium with the particle, momentum and heat sources and sinks. The local plasma parameters are then constant in time as long as these sources and sinks do not change. The poloidal magnetic flux may not be constant but it must vary at a constant rate to induce a constant and homogeneous toroidal electric field in the plasma.

Restricting the singular perturbation expansion to first order in $\varepsilon$, and projecting the partial differential system on a set of radial basis functions with a Galerkin scheme, a lumped-parameter, linear time-invariant model is obtained, in which all distributed variables and unknown operators reduce to finite dimension vectors and matrices [30]. In the following, the functions $a_k(x)$ represent a set of cubic splines[7] and the symbols $\psi$ and $\kappa$ will be used to denote either the functions $\psi(x, t)$ and $\kappa(x, t)$, respectively, or the vectors $\psi(t)$ and $\kappa(t)$ whose elements, $\psi_k(t)$ and $\kappa_k(t)$, are the coefficients of their respective Galerkin expansion:

$$\psi(x,t) = \sum_{k=1}^{N} \psi_k(t)\, a_k(x) + R_{\psi,k}(x,t) \tag{3}$$

$$\kappa(x,t) = \sum_{k=1}^{N} \kappa_k(t)\, a_k(x) + R_{\kappa,k}(x,t). \tag{4}$$

The residuals $R_{\psi,k}(x,t)$ and $R_{\kappa,k}(x,t)$ are defined by:

$$\int_0^1 a_k(x)\, R_{\psi,k}(x,t)\, \mathrm{d}x = 0 \tag{5}$$

$$\int_0^1 a_k(x)\, R_{\kappa,k}(x,t)\, \mathrm{d}x = 0 \tag{6}$$

and will be neglected. Equations (1) and (2) thus reduce to a slow dynamic model,

$$\partial\psi(t)/\partial t = A_S \cdot \psi(t) + B_S \cdot U_S(t) \tag{7}$$

$$\kappa_S(t) = C_S \cdot \psi(t) + D_S \cdot U_S(t) \tag{8}$$

and a fast dynamic model,

$$\partial\kappa_F(t)/\partial t = \varepsilon^{-1}\partial\kappa_F(\tau)/\partial\tau = A_F \cdot \kappa_F(t) + B_F \cdot U_F(t) \tag{9}$$

where $\kappa(t) = \kappa_S(t) + \kappa_F(t)$, and $A_S$ and $A_F$ are regular matrices with negative eigenvalues characteristic of the slow and fast dynamics, respectively. The vector $U(t)$ containing the actuator inputs, $P(t)$ and $V_{ext}(t)$, is also split into a slow and a fast component, $U(t) = U_S(t) + U_F(t)$. The fast inputs will not generate a response from the slow system (7) and (8) and vice versa because $U_F(t)$ does not contain any low frequency component and high frequency components have been removed in $U_S(t)$. It follows that, when a steady state is reached, $U_S(t)$ is constant and $U_F$ vanishes, and therefore $\psi(t)$ and $\kappa_S(t)$ are constant, and $\kappa_F$ vanishes. It is clear from the expansion in powers of $\varepsilon$ of the original system that, apart from anomalous flux redistribution caused by some rapid MHD events (which is not described by the resistive diffusion operator in (1) and should be treated by the controller as an external disturbance), $\psi(t)$ has only a slow evolution. Its fast component can be set identically to zero in the two-time-scale model. This would also be true for any other magnetic variable (current density, safety factor profile, etc). Equation (7) is a slow dynamic state equation obtained from (1) by replacing the kinetic vector $\kappa(t)$ by the quasi-steady solution of equation (2), at zero order in $\varepsilon$. Equation (8) appears as an output equation for the slow model. It includes a term representing a direct feed-through from the input to the output, but only from the low frequency input component, which therefore does not yield any fast response. The fast model, which is obtained from (2) by retaining only the first order terms in $\varepsilon$, is limited to a system of ordinary differential equations (9). Any fast component of $V_{ext}(t)$ can be neglected in equation (9) as $V_{ext}(t)$ does not enter equation (2). More details about the assumptions and approximations leading to equations (1)–(9) can be found in [28].

Linear system identification algorithms may not be efficient enough, especially when using noisy data, to seek models whose order is as large as the dimensions of the $\psi$ and $\kappa$ vectors. It is therefore convenient to reduce the system order, a priori, by retaining a small number of eigenmodes in the models. To do so, we define new state vectors, $X_S(t)$ and $X_F(t)$, that represent the projections of the slow and fast model states in the basis of a reduced set of slow and fast eigenmodes. The system then reads

$$\dot{X}_S(t) = A_S \cdot X_S(t) + B_S \cdot U_S(t) \tag{10}$$

$$\psi(t) = C_\psi \cdot X_S(t) \tag{11}$$

$$\kappa_S(t) = C_S \cdot X_S(t) + D_S \cdot U_S(t) \tag{12}$$

and

$$\dot{X}_F(t) = A_F \cdot X_F(t) + B_F \cdot U_F(t) \tag{13}$$

$$\kappa_F(t) = C_F \cdot X_F(t) \tag{14}$$

so that $\psi$ and $\kappa$ now become output variables and the order of the models to be identified (i.e. the number of columns in the $C_\psi$ and $C_F$ matrices) can be varied until the best fit between the original data and the model prediction is obtained. When some kinetic responses are too fast and the corresponding eigenvalues of the $A_F$ matrix cannot be identified, it was found judicious to add a direct feed-through term, $D_F \cdot U_F(t)$, in the right hand side of equation (14). This term can approximate the fastest response from $U_F(t)$ to $\kappa_F(t)$.

The $q$-profile is controlled through its inverse,

$$\begin{aligned}\bar{\iota}(x,t) &= \iota(x,t)/(2\pi) = 1/q(x,t)\\ &= -[\partial\psi(x,t)/\partial x] \cdot [\partial\Phi(x,t)/\partial x]^{-1}\\ &= -[1/(2\Phi_{max})] \cdot [(1/x) \cdot \partial\psi(x,t)/\partial x]\end{aligned} \tag{15}$$

[7] The set of radial basis functions could be different for each dynamical variable and must be chosen judiciously to provide satisfactory fits of the corresponding parameter profiles.

where $\iota(x, t)$ refers to the rotational transform and is defined, in toroidal plasmas, as the poloidal angle subtended during a single toroidal transit of the field lines on a particular flux surface, expressed in radians. $\Phi_{max}$ is known from the real-time magnetic equilibrium reconstruction that provides the current profile data. At constant vacuum toroidal field and plasma shape, it was checked experimentally that the variations of $q(x,t)$ and $\bar{\iota}(x, t)$ are indeed due mostly to the variations of $\psi(x,t)$ as $\Phi_{max}$ depends weakly on the system inputs in comparison with $\psi(x, t)$. Controlling $\bar{\iota}(x, t)$ rather than $q(x,t)$ is a natural choice due to the inverse dependence of the safety factor with respect to the poloidal flux and current density, and therefore with respect to the H&CD control actuators. The postulated structure of equations (1) and (2) was indeed based upon linear relationships between beam-driven or wave-driven currents and injected powers through current drive efficiencies, and between the ohmic current density and the surface loop voltage through Ohm's law [30].

In order to control $\bar{\iota}(x, t)$, an additional output equation relating $\bar{\iota}(x, t)$ to the model states must be used. Neglecting the residuals, the approximate finite expansion (3) that led to the discrete model equations (7)–(9), readily provides an approximation for $\bar{\iota}(x, t)$:

$$\bar{\iota}(x, t) = -[1/(2\Phi_{max})] \cdot \sum_{N}^{k=1} \psi_k(t)\ \alpha_k(x) \qquad (16)$$

in terms of a finite expansion using the same coefficients, $\psi_k(t)$, as the $\psi(x,t)$ expansion, but on a different set of basis functions, $\alpha_k(x)$. The new basis functions are defined by:

$$\alpha_k(x) = (1/x) \cdot \mathrm{d}a_k(x)/\mathrm{d}x. \qquad (17)$$

The basis functions, $a_k(x)$, used to approximate $\psi(x,t)$ must satisfy a regularity condition at $x = 0$, namely that their first derivative vanishes. This ensures that the new basis functions, $\alpha_k(x)$, are finite at $x = 0$ so that the approximate $\bar{\iota}(x, t)$ in (16) is defined and finite on the magnetic axis, as its physical counterpart is. Once a sufficiently accurate model has been identified for the dynamics of $\psi(x,t)$ through the $\psi_k$ elements of the $\psi$-vector and the $a_k(x)$ basis functions, the dynamics of $\bar{\iota}(x, t)$ can be approximated by equation (16) for controller synthesis. For instance, control of the $\bar{\iota}(x, t)$ profile can be achieved through the control of the $\psi$-vector with an appropriate control objective, based on (16) and involving scalar products of the $\alpha_k(x)$ functions, that forces the $\bar{\iota}(x, t)$ profile to reach the closest least square approximation of a given target profile, $\bar{\iota}_{target}(x)$ [33]. Equation (16) is essential because $q(x,t)$ and $\bar{\iota}(x, t)$ involve the spatial derivative of the poloidal flux, and the experimental $q$-profile data that is obtained from tokamak real-time magnetic equilibrium reconstruction codes is therefore extremely noisy. As a result, using directly the $\bar{\iota}(x, t)$ data for system identification was not found practical, yielding unsatisfactory fits between the data and the model prediction. It was much easier and judicious to identify the model described by equations (10) and (11) for the $\psi(t)$ vector, using the internal poloidal flux data, and to use equation (16) for all the mathematical expressions, scalar products and integrals involving $\bar{\iota}(x, t)$ in the control algorithm.

## 3. Identification of *ARTAEMIS* models from nonlinear plasma simulations with EAST-like parameters

The design of an advanced profile controller using MPC techniques required many tests and simulations before implementing the controller in the EAST PCS and using it for experimental tests. Various control options and tunable parameters were thus optimized from the results of numerical simulations. The METIS plasma simulator [34] was an ideal tool for this because it can produce, with a reasonable CPU time, extensive nonlinear closed loop simulations that allow the controller performance to be evaluated. METIS includes an MHD equilibrium and a current diffusion solver, and combines plasma transport nonlinearity with 0D scaling laws and 1.5D ordinary differential equations. Despite its rapid convergence, METIS integrates basically all the features of real tokamak physics in a simplified but comprehensive and flexible way. Tokamak complexity is restored through the very large number of possible options and sophisticated nonlinear models that the code offers for every elementary physical process. However, the way METIS was constructed would make it very difficult to extract linear versions of the code around a given equilibrium state. Unlike the control-oriented FPD models used in RAPTOR [13] and COTSIM [23] whose equations can be linearized for controller design, METIS is too complex to offer this option. Then, for controller synthesis, linear models must be obtained from system identification using large synthetic data sets obtained from open loop simulations with random excitation of the various actuators. Before running simulations for a particular device in a given scenario, the various models and adjustable parameters in METIS must generally be tuned in order to provide a satisfactory agreement with some experimental data that is relevant to the problem of interest. Such a tuning procedure was previously used for DIII-D and is described in details in [35]. Here, similarly, METIS modeling options and parameters were tuned to approximate two different steady state H-mode scenarios on EAST until a fair agreement with experimental data was obtained. The first scenario is at low plasma density, without NBI, while the second scenario includes NBI, at higher density, and will be used in section 7.

The low-density scenario around which a linear *ARTAEMIS* response model was identified was based on shot #62946, a steady state, fully non-inductive single-null H-mode discharge, at a toroidal magnetic field, $B_T = 2.5$ T, line-averaged density, $n_{el} \approx 2.7 \times 10^{19}$ m$^{-3}$, and plasma current, $I_p = 0.42$ MA. The transition to H-mode occurred at 3.1 s with an H-factor, $H_{98(y,2)} \sim 1.1$. The steady state poloidal $\beta$ and internal inductance parameters were $\beta_p = 1.3$ and $l_i = 1.2$, respectively, and the $q$-profile was characterized by a small negative shear in the plasma core, with a minimum $q$ around 1.5 and $q_0 = q(x = 0) \sim 2$ on axis. The experimental

poloidal flux and safety factor profiles were obtained from EFIT magnetic equilibrium reconstructions, which were available in real-time using magnetic and kinetic measurements, including interfero-polarimetry data from the POINT diagnostic. The reference discharge had been obtained using LHCD (0.6 MW at 2.45 GHz and 2 MW at 4.6 GHz), ICRH at 33 MHz (0.32 MW) and ECRH at 140 GHz (0.3 MW). It must be emphasized that the METIS plasma simulator is a fast code and the simplified model for LHCD in METIS has been developed considering only one frequency. It cannot provide a very accurate simulation of the plasma dynamics in real EAST discharges that use two LHCD systems at different frequencies. But it provides a nonlinear simulation platform that reproduces qualitatively the complex plasma dynamics prevailing in tokamaks, and on which control algorithms and options can be extensively tested and compared. In the METIS simulations, the 2.45 GHz LHCD power was added to the ECRH power, which then amounts to 0.9 MW in the low-density scenario. The ECRH power and current deposition profiles were modified consistently with the assumption that two thirds of the ECRH power used in METIS, i.e. 0.6 MW, is indeed provided by the 2.45 GHz LHCD system. By resorting to these approximations, controller design could be optimized without the limitations encountered with experimental tests, such as the scarcity and short duration of experimental sessions and the intermittent availability of some actuators.

Open loop METIS simulations were run using either steady state or randomly modulated actuator waveforms for the 4.6 GHz LHCD and the ICRH powers, $P_{\rm LH}(t)$ and $P_{\rm IC}(t)$, respectively. Like EAST discharges, the simulations were run in the current control mode, in which $I_{\rm p}(t)$ is accurately regulated by a specific controller through the ohmic poloidal field coil current. $V_{\rm ext}(t)$, which appears as a natural input in equation (1) is therefore reserved for the plasma current controller and it is determined from a strong feedback law that regulates the plasma current. Therefore, it cannot be used as an independent input for system identification [28] and as a profile control actuator, as in [33]. During the plasma current flat-top, $I_{\rm p}(t)$ is strictly constant and the variations of $V_{\rm ext}(t)$ in the linearized system (1) and (2) are related to the variations of the H&CD actuators, $P(t)$, and of $\psi(\mathrm{x}, t)$ and $\kappa(\mathrm{x}, t)$. We have therefore empirically sought models in which the surface loop voltage has been omitted in the right hand side of equation (1). In practice, this assumption did not prevent the system identification process to converge satisfactorily to models yielding good fits to the data. Model errors, including those that could stem from this assumption, will be compensated in real time by the MPC controller.

In the *ARTAEMIS* data-driven model identified from these METIS simulations, the chosen magnetic output variables were the $\psi(t)$ vector elements, and the kinetic output variable was $\kappa(t) = \delta\beta_{\rm p}(t)$, the poloidal beta parameter reduced by its reference value. The dimension of the $\kappa_{\rm S}(t)$ and $\kappa_{\rm F}(t)$ vectors in equations (12) and (14) is thus here equal to 1 as $\kappa$ is independent of x. The input vector $U(t)$, of dimension 2, contains the reduced powers, $\delta P_{\rm LH}(t)$ and $\delta P_{\rm IC}(t)$. The data was processed with the general methodology described in [30] so that the system identification routines only handle datasets with nearly zero mean while searching for the best linear response model. The *ARTAEMIS* identification algorithm is based on the prediction-error method (PEM) [28]. For each output parameter profile in the model, say $Y(\mathrm{x}, t)$, it maximizes a global fit parameter, which is defined as $f = 1-\delta$, with

$$\delta = \sum_{\rm samples} \int_{x_1}^{x_2} (Y(x) - Y_{\rm sim}(x))^2 {\rm d}x \Big/ \sum_{\rm samples} \int_{x_1}^{x_2} \left( Y(x) - \langle Y(x) \rangle_{\rm samples} \right)^2 {\rm d}x \qquad (18)$$

where $Y_{\rm sim}$ represents the data simulated by the current model after reconstruction of the profiles from the basis functions, $[\mathrm{x}_1, \mathrm{x}_2]$ is the radial window on which the model is to apply (here $\mathrm{x}_1 = 0$ and $\mathrm{x}_2 = 1$) and $\langle Y(x) \rangle_{\rm samples}$ stands for the average of the data over the time samples, at a given radius. Equation (18) is the first instance where we can use the approximate mapping (16) between $\psi(\mathrm{x})$ and $\bar{\iota}(x)$. Even though the model outputs involve $\psi$, the fit parameter to be maximized can be defined in terms of $\bar{\iota}$. In the sum over the samples and the integrals corresponding to the magnetic part of the model, $Y(\mathrm{x})$ was replaced by $\bar{\iota}(x)$ expressed in terms of $\psi_k$ and the $\alpha_k(\mathrm{x})$ basis functions, rather than being replaced by $\psi(\mathrm{x})$ in terms of $\psi_k$ and the $a_k(\mathrm{x})$ functions.

Fourteen simulations were used to identify the slow part of the *ARTAEMIS* data-driven model. The simulations were similar to the reference simulation of shot #62946 up to $t = 3.2$ s and variations of the actuators were imposed for $t \geqslant 3.2$ s. Five of them were steady state simulations with various constant values of $P_{\rm LH}$ and $P_{\rm IC}$ from $t = 3.2$ s to $t = 10$ s. They are useful to determine the largest time constant in the model output dynamics. In the remaining simulations, the values of $P_{\rm LH}$ and $P_{\rm IC}$ were modulated from $t = 3.2$ s to $t = 15$ s, following square wave modulations between 1 MW and 2.5 MW for $P_{\rm LH}$ and between 0.5 MW and 1.25 MW for $P_{\rm IC}$. An example of such $P_{\rm LH}$ and $P_{\rm IC}$ waveforms is shown on figure 1. For each actuator, the time between power steps were chosen at random between a minimum of 0.04 s and a maximum of 2 s, with a final steady state period of about 5 s. This was consistent with an expected resistive diffusion time of the order of one second for the chosen reference plasma parameters. The number of magnetic outputs in equation (11) was equal to 10, corresponding to the values of $\psi(\mathrm{x}, t)$ at the cubic splines knots $\mathrm{x}_k = 0, 0.1, \ldots 0.9$ (note that $\psi(1, t) = 0$ by definition). The number of significant eigenmodes in the model was increased at each iteration of the identification process leading to the optimization of the $A_{\rm S}$, $B_{\rm S}$, and $C_\psi$ matrices in equations (10) and (11), until the best fit with the data was obtained from equation (18). Four eigenmodes could thus be found, with eigenvalues equal to $-1.19\ {\rm s}^{-1}$, $-4.24\ {\rm s}^{-1}$, $-11.49\ {\rm s}^{-1}$ and $-13.62\ {\rm s}^{-1}$. The corresponding characteristic times are 0.843 s (defined as the resistive diffusion time, $\tau_S$), 0.236 s, 0.087 s and 0.073 s. A comparison between the reduced $\psi(\mathrm{x})$ data at $\mathrm{x} = 0, 0.1, \ldots$

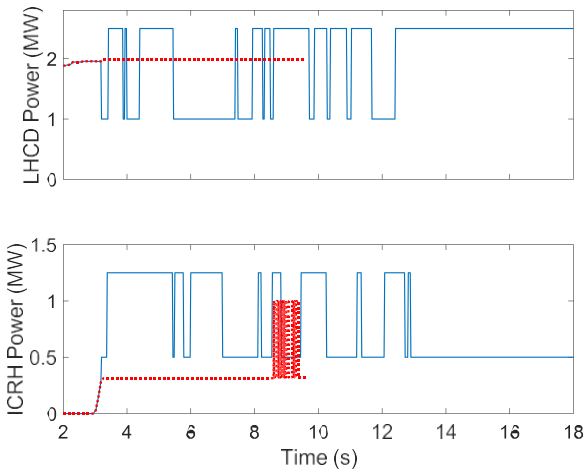


**Figure 1.** LHCD (top) and ICRH (bottom) power waveforms (MW) used in METIS simulations with low frequency modulations (blue) and high frequency modulations (red dotted).

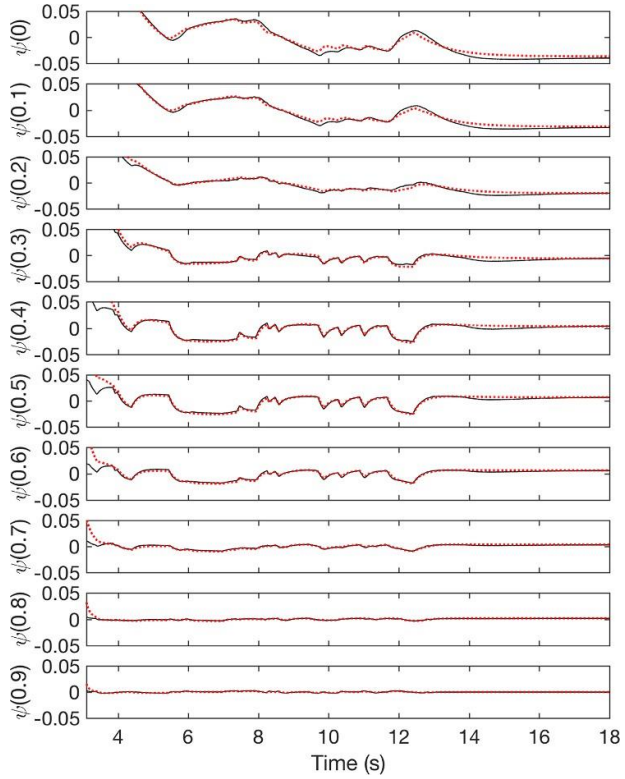


**Figure 2.** Comparison between the reduced $\psi(\mathrm{x})$ data (Wb) at $\mathrm{x} = 0, 0.1, \ldots 0.9$ from METIS (black) and the *ARTAEMIS*-simulated data (red dotted) from the two-time-scale model for the simulation with low frequency modulations whose input data are shown on figure 1. In this simulation, the global fit parameter for $\psi(\mathrm{x})$, integrated from $\mathrm{x}_1 = 0$ to $\mathrm{x}_2 = 1$, is $f = 77\%$.

0.9 from METIS and the *ARTAEMIS*-simulated data from the two-time-scale model is shown on figure 2, with a fit parameter $f = 77\%$. The fourteen values of the fit parameter obtained for the various simulations were found to range between 65% and 86%.

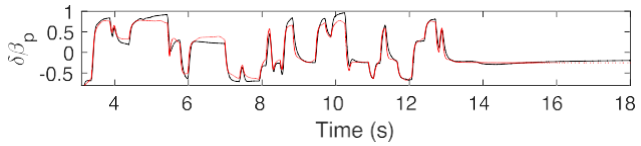


**Figure 3.** Comparison between the reduced $\delta\beta_p$ data from METIS (black) and the *ARTAEMIS*-simulated data (red dotted) for low frequency LHCD and ICRH power modulations. The corresponding input data are shown on figure 1 and the fit parameter for $\delta\beta_p$ in this simulation is $f = 79\%$.

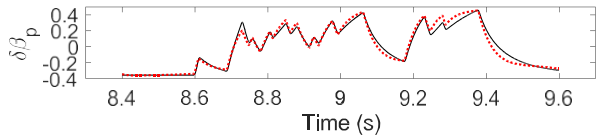


**Figure 4.** Comparison between the reduced $\delta\beta_p$ data from METIS (black) and the *ARTAEMIS*-simulated data (red dotted) for high frequency ICRH power modulations. The corresponding input data are shown on figure 1 and the fit parameter for $\delta\beta_p$ in this simulation is $f = 83\%$.

The slow model $C_S$ and $D_S$ matrices were identified using the nine simulations with modulated actuators among those used for the magnetic part of the model. In order to reject the fast input components and fast $\delta\theta_p$ response, and define their slow components, $U_S(t)$ and $\delta\theta_{p,S}(t)$, respectively, both the input and output data were filtered with a non-causal low-pass filter of order 5, and a cutoff frequency, $f_{filt} = 3.5$ Hz. This was found to provide an optimum separation between the slow and fast data for the two-time-scale model identification. Finally, the fast model (13) and (14) was identified from six new simulations that were also similar to the reference simulation of shot #62946 up to $t = 8.4$ s, but with high frequency square wave modulations of the actuators between $t = 8.4$ s and $t = 9.6$ s. An example of fast ICRH modulations is shown on figure 1. The time between power steps were chosen randomly between a minimum of 0.004 s and a maximum of 0.3 s. The fast input and output components were defined as $U_F(t) = U(t) - U_S(t)$ and $\delta\theta_{p,F} = \delta\theta_p - \delta\theta_{p,S}$, respectively. A model of order 1 was sought and the best fit was obtained with an eigenvalue of $-24.8\ \mathrm{s}^{-1}$ for $A_F$, corresponding to a characteristic thermal diffusion time, $\tau_F = 0.04$ s. Note that the ordering between $\tau_S$ and $\tau_F$ is compatible with the main hypothesis justifying the singular perturbation approximation, namely that $\tau_F \ll \tau_S$. The filter defining the slow and fast components of $U(t)$ and $\delta\theta_p(t)$ was adequately chosen since the inverse of its cutoff frequency, $\tau_{filt} = 1/f_{filt} = 0.29$ s, lies between $\tau_F$ and $\tau_S$. Comparisons between the reduced $\delta\theta_p$ data from METIS and the *ARTAEMIS*-simulated data are shown on figure 3 for a simulation with low frequency modulations and on figure 4 for a simulation with high frequency modulations. The fit parameters obtained for $\delta\theta_p$ from the full two-time-scale model (10)–(14) were found in the range between 71% and 91% for the simulations with either slow or fast power modulations.

## 4. State observer, model errors and two-time-scale MPC controller design

A simple profile control algorithm that combines the two-time-scale linear model structure with linear-quadratic optimal control theory was the so-called near-optimal control algorithm originally described in [30] and tested on JET [30], DIII-D [33] and, later, more thoroughly on METIS simulations [35]. Both experiments and simulations have shown simultaneous magnetic profile control of $\psi(x)$ or $q(x)$, and kinetic control of $\beta_N$, in steady state, fully non-inductive discharges. However, in high-$\beta_N$ long-pulse METIS simulations, the desired steady state $q$-profiles were obtained either much too slowly or after a large $q(x = 0)$ undershoot and a damped oscillation. To overcome these limitations, a controller that combines the simplicity of the *ARTAEMIS* two-time-scale models with the efficiency of model-based predictive control techniques was designed for the first time and is described in this section. The MPC control algorithm is based on the same semi-empirical models, but offers many advantages over the more conventional near-optimal control algorithm. A major advantage has to do with actuator saturation. In the former algorithm, the PI feedback on the slow variables must be supplemented by an anti-windup loop to prevent state winding-up during actuator saturation [30, 35]. This leads to unoptimized control action and delays in the controller response. In MPC theory, the hard physical actuator constraints that must be met during the plasma evolution can be directly embedded in the QP optimization process [38]. Another feature of the new design is related to model uncertainties that were not specifically taken into account in the near-optimal controller. In the MPC algorithm, model errors can be identified through an observer, and compensated in real time to achieve efficient offset-free tracking of assigned control targets [40].

Skogestad [41] nicely compares the advantages and disadvantages of MPC compared to more simple techniques and shows that it also has some drawbacks. Its main disadvantage lies in the complex QP optimization process that needs longer CPU time than PID-like controllers. A second potential difficulty is that MPC often requires a large number of model coefficients to describe the plant response satisfactorily, again increasing the required CPU time. The prediction horizon must be optimally chosen for good control performance, even if the model is accurate. Also, the constant output disturbance assumption, which assumes that the disturbance term is constant in the future (see section 4.1), is often made for simplicity. This may not always yield a good controller performance if there are large input disturbances. These drawbacks did not lead to significant obstacles for the implementation of our MPC controllers and for the applications presented in this paper.

The main concept of MPC is to use the model to predict the future evolution of the system. At each time step, $t$, a sequence of possible future inputs is computed, subject to operating constraints, in order to minimize a given cost function that depends on the predicted evolution of the system during a given time horizon, $\tau_H$. The first value in the optimized actuator sequence is chosen as the control action (controller output) at time $t$, and subsequent values are discarded. At the next time step, a new optimization is solved over a shifted prediction horizon. Thus, with an adequate choice of the tuning parameters such as the horizon and cost function, the controller will achieve its goal if the model steady state gain is exact. However, if the steady state gain matrix of the model is not exact, the prediction cannot be made from the model only, because the system outputs would reach incorrect values. It must somehow include assumptions about disturbances acting on the system and measurement errors. The prediction strategy then becomes another 'tuning parameter'.

### 4.1. The augmented two-time-scale model

In order to adjust the prediction and make the controller robust to disturbances and uncertainties, the identified *ARTAEMIS* model is augmented so that it includes a simple model of the output disturbances. The output disturbances, which represent the mismatch between measured and predicted outputs, will be considered as new states of the augmented system [38, 40]. The estimated disturbances will depend on time and they will therefore vary at each time step. However, we shall make the simple assumption that they remain unchanged during the given time horizon. This assumption will hold only for making the prediction of the optimized actuator sequence (between $t$ and $t + \tau_H$) whose first value provides the control action at time $t$, while subsequent values are discarded. This will be referred to as the constant output disturbance model [38].

In its continuous time version[8], the augmented two-time-scale model for the control of the internal magnetic flux profile, represented by the $\psi(t)$ vector, and of some kinetic parameters or profiles represented by the $\kappa(t)$ vector will then be defined as:

$$\dot{X}_S(t) = A_S \cdot X_S(t) + B_S \cdot U(t) \tag{19}$$

$$\dot{X}_F(t) = A_F \cdot X_F(t) + B_F \cdot U_F(t) \tag{20}$$

$$\dot{d}_\psi(t) = 0 \tag{21}$$

$$\dot{d}_\kappa(t) = 0 \tag{22}$$

$$\psi(t) = C_\psi \cdot X_S(t) + d_\psi(t) \tag{23}$$

$$\kappa(t) = C_S \cdot X_S(t) + D_S \cdot U_S(t) + C_\kappa \cdot X_F(t) + d_\kappa(t) \tag{24}$$

where the vectors $d_\psi(t)$ and $d_\kappa(t)$ are the disturbance states representing the errors on $\psi(t)$ and $\kappa(t)$, respectively. $U_S(t)$ is obtained through a simple low-pass filter with the cutoff frequency that was used for system identification (section 3),

[8] For notation convenience, the systems presented in this paper are continuous-time systems. Their conversion to discrete-time systems is straightforward and was carried out for the controller implementation on the METIS simulator and on the EAST PCS.

and $U_F(t) = U(t) - U_S(t)$. As a minor simplification, we have used the full inputs, $U(t)$, in equation (19) because $U_F(t)$ does not generate a significant response on the slow states. Also, for simplicity and as suggested in [40], we have deliberately chosen the additional states of the augmented model to be equal to the model errors. This is the simplest disturbance model one could choose and it was indeed found appropriate and sufficient to ensure the observability of the augmented system (19)–(24). If necessary, more complex disturbance models could be assumed, e.g. by introducing more general linear dependences of the output variables on the disturbance states on the right hand side of equations (23) and (24), and also more elaborate linear dependences of the original model states on the disturbance states on the right hand side of equations (19) and (20).

Now, for MPC control of the safety factor profile, which will be our main objective here, equations (21) and (23) do not provide useful predictive output equations. The near-optimal control algorithm described in [30] was basically a state control algorithm controlling the $\psi(t)$ vector, for which a good state space model exists, rather than an output control algorithm. For $\bar{\iota}$ control, the specific relationship between $\psi(x,t)$ and $\bar{\iota}(x, t)$ was only introduced in the definition of the quadratic cost function representing the control objective [33, 35]. On the contrary, the MPC algorithm requires an internal model prediction of the $\bar{\iota}$ evolution during the given time horizon, and therefore the model must include an output equation for $\bar{\iota}$ with specific disturbances. To control $\bar{\iota}$ rather than $\psi$, we therefore define a vector, $\bar{\iota}(t)$, whose elements are the values of the inverse safety factor, $\bar{\iota}(x, t)$, at the normalized radii corresponding to the knots of the cubic splines used to define the $\psi(t)$ vector in section 3, i.e. at $x = 0, 0.1, 0.2, \ldots 0.9$. Using equations (11) and (16), the vector $\bar{\iota}(t)$ can be related to the $\psi(t)$ vector and to the states $X_S(t)$ through a linear output matrix equation,

$$\bar{\iota}(t) = \Gamma_\iota \cdot \psi(t) = C_\iota \cdot X_S(t) \tag{25}$$

and the augmented model must also include a disturbance vector $d_\iota(t)$ that represents the mismatch between the predicted and the actual $\bar{\iota}(t)$. The relevant augmented model then becomes:

$$\dot{X}_S(t) = A_S \cdot X_S(t) + B_S \cdot U(t) \tag{26}$$

$$\dot{X}_F(t) = A_F \cdot X_F(t) + B_F \cdot U_F(t) \tag{27}$$

$$\dot{d}_\iota(t) = 0 \tag{28}$$

$$\dot{d}_\kappa(t) = 0 \tag{29}$$

$$\bar{\iota}(t) = C_\iota \cdot X_S(t) + d_\iota(t) \tag{30}$$

$$\kappa(t) = C_S \cdot X_S(t) + D_S \cdot U_S(t) + C_\kappa \cdot X_F(t) + d_\kappa(t). \tag{31}$$

## 4. The augmented model state observer

In order to solve the predictive control problem at time $t$, we need to compute the predicted values of the controlled variables, $\bar{\iota}$ and $\kappa$, over a future time horizon, from our best estimates, at time $t$, of the current states, $X_S(t)$ and $X_F(t)$, and of the disturbances, $d_\iota(t)$ and $d_\kappa(t)$. Therefore, at each time step, an observer must provide a new estimate of the evolving system states and of the model errors, and the controller will use the augmented model (26)–(31) to predict the behavior of the system between $t$ and $t + \tau_H$ and define the appropriate control action at time $t$. The filter states that allow $U_S(t)$ and $U_F(t)$ to be computed in the observer at time $t$ are also transmitted to the controller with the estimated system states and disturbances and with the real-time measurements of the controlled variables, $\bar{\iota}_m(t)$ and $\kappa_m(t)$.

Following [40], we design the state and disturbance estimator as a classical Luenberger observer based on the augmented system above. It is defined as follows:

$$\begin{aligned}\dot{\hat{X}}_S(t) = {} & A_S \cdot \hat{X}_S(t) + B_S \cdot U(t) \\ & + K_{11} \cdot \left( -\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t) \right) \\ & + K_{12} \cdot \left( -\kappa_m(t) + C_S \cdot \hat{X}_S(t) \right) \\ & + D_S \cdot U_S(t) + C_\kappa \cdot \hat{X}_F(t) + \hat{d}_\kappa(t)\end{aligned} \tag{32}$$

$$\begin{aligned}\dot{\hat{X}}_F(t) = {} & A_F \cdot \hat{X}_F(t) + B_F \cdot U_F(t) \\ & + K_{21} \cdot \left( -\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t) \right) \\ & + K_{22} \cdot \left( -\kappa_m(t) + C_S \cdot \hat{X}_S(t) \right) \\ & + D_S \cdot U_S(t) + C_\kappa \cdot \hat{X}_F(t) + \hat{d}_\kappa(t)\end{aligned} \tag{33}$$

$$\begin{aligned}\dot{\hat{d}}_\iota(t) = {} & K_{31} \cdot \left( -\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t) \right) \\ & + K_{32} \cdot \left( -\kappa_m(t) + C_S \cdot \hat{X}_S(t) \right) \\ & + D_S \cdot U_S(t) + C_\kappa \cdot \hat{X}_F(t) + \hat{d}_\kappa(t)\end{aligned} \tag{34}$$

$$\begin{aligned}\dot{\hat{d}}_\kappa(t) = {} & K_{41} \cdot \left( -\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t) \right) \\ & + K_{42} \cdot \left( -\kappa_m(t) + C_S \cdot \hat{X}_S(t) \right) \\ & + D_S \cdot U_S(t) + C_\kappa \cdot \hat{X}_F(t) + \hat{d}_\kappa(t)\end{aligned} \tag{35}$$

where symbols with a hat represent the observer estimates of the system states and disturbances, and the $K_{i,j}$ matrices are chosen so that the state observer is stable and converges rapidly. In particular, the matrix formed by the top row $[K_{31}\ K_{32}]$ and the bottom row $[K_{41}\ K_{42}]$, governing the evolution of the estimated disturbance states, must be non-singular. The idea of a Luenberger observer is to estimate the contribution of non-measurable states in equations (30) and (31) such that the residual errors between the measured outputs and the model outputs (including the observed disturbances) decay to zero steady state values. As seen from equations (34) and (35), the

observer then tracks the measurements without steady state errors, i.e.

$$\bar{\iota}_{m,\infty} = C_\iota \cdot \hat{X}_{S,\infty} + \hat{d}_{\iota,\infty} \tag{36}$$

and

$$\kappa_{m,\infty} = C_S \cdot \hat{X}_{S,\infty} + D_S \cdot U_{S,\infty} + C_\kappa \cdot \hat{X}_{F,\infty} + \hat{d}_{\kappa,\infty} \tag{37}$$

where the $\infty$ symbol indicates steady state values. In addition, from the definition of the fast variables and the discussion following equation (9), $U_F$, $X_F$ and $\kappa_F$ vanish in steady state, so that equation (37) becomes:

$$\kappa_{m,\infty} = C_S \cdot \hat{X}_{S,\infty} + D_S \cdot U_\infty + \hat{d}_{\kappa,\infty}. \tag{38}$$

The observer also satisfies:

$$\hat{X}_{S,\infty} = -A_S^{-1} \cdot B_S \cdot U_\infty \tag{39}$$

and

$$K_S \cdot U_\infty = \begin{bmatrix} \bar{\iota}_{m,\infty} - \hat{d}_{\iota,\infty} \\ \kappa_{m,\infty} - \hat{d}_{\kappa,\infty} \end{bmatrix} \tag{40}$$

with

$$K_S = \begin{bmatrix} -C_\iota \cdot A_S^{-1} \cdot B_S \\ -C_S \cdot A_S^{-1} \cdot B_S + D_S \end{bmatrix}. \tag{41}$$

If the image space of $K_S$ spanned by the actuators has a dimension that is equal to, or larger than the number of controlled variables, offset-free control can be achieved [40]. Then, the control objective is to make $\bar{\iota}_m$ and $\kappa_m$ equal to assigned targets or set points, $\bar{\iota}_{target}$ and $\kappa_{target}$, respectively. The singular decomposition of $K_S$ may indicate that some combinations of actuators corresponding to the smallest singular values of $K_S$ have little or even negligible effect on the system dynamics, even at high power. Such combinations of actuators may therefore lead to unnecessary actuation cost and should be discarded in the computation of the optimal control action. In this aim, the allowed actuator space will be limited to the first $n_{svd}$ singular vectors of $K_S$, i.e. $U = T_{svd} \cdot V$ and $U_\infty = T_{svd} \cdot V_\infty$, where $n_{svd}$ (the dimension of $V$) is a free tuning parameter.

We can now explicitly introduce the filter equations that define the separation of the slow and fast time scales into the augmented model and the observer definition. We shall use here a simple first order filter represented by a square diagonal matrix, $A_v$, whose dimension is equal to the number of inputs in $U(t)$, and whose diagonal elements are all equal to $f_{filt} = -1/\tau_{filt}$. The filter equations then read:

$$\dot{v}(t) = A_v \cdot [v(t) - U(t)] \tag{42}$$

$$U_S(t) = v(t) \tag{43}$$

$$U_F(t) = -v(t) + U(t) \tag{44}$$

where $v(t)$ is the filter state vector, which is updated at the input of the observer at each time step, using the last measured value of $U(t)$. The state space observer equations can thus be written in terms of new inputs, $v(t)$, $U(t)$, $\bar{\iota}_m(t)$ and $\kappa_m(t)$ instead of $U_S(t)$, $U_F(t)$, $\bar{\iota}_m(t)$ and $\kappa_m(t)$:

$$\begin{aligned} \dot{\hat{X}}_S(t) = {} & A_S \cdot \hat{X}_S(t) + B_S \cdot U(t) \\ & + K_{11} \cdot \left( -\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t) \right) \\ & + K_{12} \cdot \left( -\kappa_m(t) + C_S \cdot \hat{X}_S(t) \right) \\ & + D_S \cdot v(t) + C_\kappa \cdot \hat{X}_F(t) + \hat{d}_\kappa(t) \end{aligned} \tag{45}$$

$$\begin{aligned} \dot{\hat{X}}_F(t) = {} & A_F \cdot \hat{X}_F(t) - B_F \cdot v(t) + B_F \cdot U(t) \\ & + K_{21} \cdot \left( -\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t) \right) \\ & + K_{22} \cdot \left( -\kappa_m(t) + C_S \cdot \hat{X}_S(t) \right) \\ & + D_S \cdot v(t) + C_\kappa \cdot \hat{X}_F(t) + \hat{d}_\kappa(t) \end{aligned} \tag{46}$$

$$\begin{aligned} \dot{\hat{d}}_\iota(t) = {} & K_{31} \cdot \left( -\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t) \right) \\ & + K_{32} \cdot \left( -\kappa_m(t) + C_S \cdot \hat{X}_S(t) \right) \\ & + D_S \cdot v(t) + C_\kappa \cdot \hat{X}_F(t) + \hat{d}_\kappa(t) \end{aligned} \tag{47}$$

$$\begin{aligned} \dot{\hat{d}}_\kappa(t) = {} & K_{41} \cdot \left( -\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t) \right) \\ & + K_{42} \cdot \left( -\kappa_m(t) + C_S \cdot \hat{X}_S(t) \right) \\ & + D_S \cdot v(t) + C_\kappa \cdot \hat{X}_F(t) + \hat{d}_\kappa(t) \end{aligned} \tag{48}$$

Now, the various observer gain matrices, $K_{i,j}$, must be chosen judiciously in order for the observer to be stable and to converge reasonably fast. In order to come up with a simple observer design, we can intuitively associate equations (45) and (47) to the observation of $d_\iota(t)$ and equations (46) and (48) to the observation $d_\kappa(t)$. Thus, we have set the off-diagonal elements $K_{1,2}$, $K_{2,1}$, $K_{3,2}$ and $K_{4,1}$ to zero. Then, to ensure fast convergence to steady state, the brackets that multiply $K_{3,1}$ and $K_{4,2}$ must rapidly vanish and therefore the gains $K_{3,1}$ and $K_{4,2}$ must have a large norm compared to $K_{1,1}$ and $K_{2,2}$. A simple way of satisfying this requirement is to introduce pseudo-inverse matrices of $C_\iota$ and $C_\kappa$ in $K_{1,1}$ and $K_{2,2}$, respectively, and a large scalar gain, $G_{obs}$, in $K_{3,1}$ and $K_{4,2}$. Moreover, for normalization purposes, the slow time constant, $\tau_S$, can be associated to the evolution of $X_S(t)$ and $d_\iota(t)$ and can be used to normalize $K_{1,1}$ and $K_{3,1}$. Similarly, the fast time constant, $\tau_F$, can be associated to the evolution of $X_F(t)$ and $d_\kappa(t)$ and can be used to normalize $K_{2,2}$ and $K_{4,2}$. An adequate choice for the observer gain matrices can therefore be the following:

$$K_{11} = -\mathrm{inv}(C_\iota)/\tau_S, \qquad K_{12} = 0 \tag{49}$$

$$K_{21} = 0, \quad K_{22} = -\mathrm{inv}(C_\kappa)/\tau_F \tag{50}$$

$$K_{31} = -G_{obs} \cdot Id_\iota/\tau_S, \qquad K_{32} = 0 \tag{51}$$

$$K_{41} = 0, \quad K_{42} = -G_{obs} \cdot Id_\kappa/\tau_F \tag{52}$$

where the inv function is a pseudo-inverse function since $C_\iota$ and $C_\kappa$ may be rectangular matrices, and $G_{obs}$ is a large positive tuneable gain. $Id_\iota$ and $Id_\kappa$ are identity matrices whose dimensions are equal to the number of controlled magnetic and kinetic variables, respectively. The choice of the observer gain matrices represented by equations (49)–(52) was indeed found to be adequate in all our control simulations and experiments, including those that are not reported in this paper.

When implementing the controller either on a simulation platform or a machine PCS, it is advisable to use normalized input, output and state variables, so that the observer and controller tunings are to a great extent independent of the plasma parameters and machine performance. The transformation from the real dimensional variables to normalized ones is straightforward and was indeed performed in all the implemented controllers used in our simulations and experiments, but they will not be represented in this paper, for the sake of clarity. The choice given in equations (49)–(52) assumes that the variables handled by the observer are normalized. The inclusion of $\tau_S$ and $\tau_F$ in the equations is suggested by comparison with the values of the $A_S$ and $A_F$ eigenvalues in the leading terms of equations (45) and (46). In all cases, except for initial tuning simulations, the observer gain $G_{obs}$ was chosen equal to 100, which resulted in good observer performance.

### 4.3. The two-time-scale MPC controller design

Once the observer has estimated the current system states and disturbances at time $t$, the MPC algorithm then solves a QP problem using the predicted evolution of the augmented system over the future time horizon, $\tau_H$. The choice of $\tau_H$ is obviously important for the controller performance. Intuitively, it should be long enough to characterize the evolution of the slow model on the resistive time scale. On the other hand, a long time horizon would be costly in terms of computation time, and it may not be meaningful given that the prediction is made with the assumption that the estimated error at time $t$ will be constant between $t$ and $t+\tau_H$. So, as a compromise, $\tau_H$ will be chosen of the order of the resistive time, $\tau_S$, or slightly smaller. From now on, we shall assume that the kinetic controlled variables are scalars, independent of x, so that the QP problem to be solved at time $t$ reads as follows:

For $t \leqslant t' \leqslant t+\tau_H$, find $V(t')$ that minimizes

$$\begin{aligned} I_H(t) = & \int_t^{t+\tau_H} dt' \left\{ \int_0^1 dx\, \mu(x)^2 [\bar{\iota}(x,t') - \bar{\iota}_{target}(x,t')]^2 \right\} \\ & + \int_t^{t+\tau_H} dt' \left\{ [\kappa(t') - \kappa_{target}(t')]^+ \cdot \lambda_{kin}^2 \cdot [\kappa(t') - \kappa_{target}(t')] \right\} \\ & + \int_t^{t+\tau_H} dt' \left\{ X_F(t')^+ \cdot \lambda_{fast}^2 \cdot X_F(t') \right\} \\ & + \int_t^{t+\tau_H} dt' \left\{ [U(t') - U_\infty]^+ \cdot R \cdot [U(t') - U_\infty] \right\} \end{aligned} \tag{53}$$

with $U(t') = T_{svd} \cdot V(t')$, subject to the actuator constraints $L(t') \cdot U(t') \leqslant M(t')$, while $X_S(t')$, $X_F(t')$, $\bar{\iota}(t')$, $\kappa(t')$, $d_\iota(t')$ and $d_\kappa(t')$ evolve according to the augmented system (26)–(31), with the initial conditions:

$X_S(t'=t) = \hat{X}_S(t)$, $X_F(t'=t) = \hat{X}_F(t)$, $d_\iota(t'=t) = \hat{d}_\iota(t)$ and $d_\kappa(t'=t) = \hat{d}_\kappa(t)$. In order to keep the computation time small, the elements of $V(t')$ are constrained to be piecewise constant functions of time with only $n_{nodes}$ independent unknowns equi-distributed over the horizon $\tau_H$. In equation (53), a + superscript indicates matrix transposition. The function $\mu(x)$ and the real diagonal matrix $\lambda_{kin}$ are weights given to the control of the magnetic and kinetic variables, respectively. The real diagonal matrix $\lambda_{fast}$ contains additional weights limiting the amplitude of the fast model states, which was found useful to moderate the kinetic control response. Finally, $R$ is a positive semi-definite matrix that can moderate the controller actuation effort, if necessary. Note that the need for a strict positivity of $R$ is alleviated by the linear constraints setting bounds on the actuators, which ensure the positivity of the matrix involved in the Cholesky factorization performed by the QP algorithm [42] at each time step. The radial integrals in the first term of equation (53) can be written as a quadratic expression in matrix form by replacing $\bar{\iota}(x,t')$ and $\bar{\iota}_{target}(x,t')$ by their expansions in terms of the cubic splines $a_k(x)$ defined in section 2, using vectors whose elements are the values of $\bar{\iota}(x,t')$ and $\bar{\iota}_{target}(x,t')$ at the knots, $x_k = 0, 0.1, 0.2, \ldots, 0.9$. The weight function, $\mu(x)$, can either be a continuous function for a distributed control of $\bar{\iota}(x)$, or a sum of Dirac distribution functions involving only a few discrete values of x among $0, 0.1, \ldots 0.9$ where discrete control would be applied. Once the QP problem has been solved, the first sample $U(t'=t)$ is used for the actuator commands at time $t$. The minimized cost function penalizes, with appropriate tuneable weights, the deviations of the predicted controlled outputs from their targets, as well as the actuator powers if $R \neq 0$. Another way of moderating the actuation effort and avoiding overshoots and oscillations of the slow magnetic variables is to reshape the $\bar{\iota}_{target}(t')$ waveforms in $I_H(t)$ so that, for $t'=t$, they start from the measured values of $\bar{\iota}$, i.e. $\bar{\iota}_m(t)$, and, when $t'$ approaches $t+\tau_H$, they approach exponentially the set-points, $\bar{\iota}_{target}(t)$, with a time constant, $\tau_{target}$, of the order of $\tau_S$ or smaller [38]. In equation (53), we thus use:

$$\bar{\iota}_{target}(x,t') = \bar{\iota}_{target}(x,t) - [\bar{\iota}_{target}(x,t) - \bar{\iota}_m(x,t)] \cdot e^{-(t'-t)/\tau_{target}} \tag{54}$$

and

$$\kappa_{target}(t') = \kappa_{target}(t). \tag{55}$$

Note that, apart from the reshaping function, no other change in the target values is made in (54) and (55) during the prediction horizon, even if the assigned set-points change during that period. Just as the disturbance states, the targets at time $t'=t$ will be assumed to hold all over the prediction horizon. Thus, the computation of the control action a time $t$ will not anticipate any possible change of the set-points at a future time.

MPC is implemented in the form of a standard quadratic program with linear constraints so that it becomes a convex optimization problem, which is efficiently solvable (e.g. [43]). The QP software that was used in this work [42] always converged after only a few iterations. No other particular measure was taken to guarantee that the iterative procedure converges under all circumstances, but the issue never appeared in simulations or experiments.

### 4.4 The reduced MPC kinetic controller design on the slow time scale

The two-time-scale *ARTAEMIS* models describe the fast kinetic dynamics of the plasma as a singular perturbation of a quasi-static equilibrium, which is slowly evolving due to the coupling between the kinetic and the magnetic plasma parameters. Local dependences of the plasma transport coefficients on the safety factor profile or on the magnetic shear are well-known examples of the various causes that lead to such coupled dynamics. When attempting to control simultaneously the safety factor profile and some other kinetic plasma parameters (e.g. $\beta_p$ or $\beta_N$), it may be unnecessary or even sometimes undesirable to request changes of such parameters on a time scale that is too short compared to the resistive evolution of the plasma equilibrium. Instead, it may be of interest to preserve a quasi-static equilibrium relationship between magnetic and kinetic plasma parameters during the transient evolution from an initial plasma state to the desired high performance steady state. Restricting the *ARTAEMIS* model to the zero-order equations in the singular perturbation analysis, i.e. to the slow model, will result in a slower kinetic control. For slow kinetic control, the model is thus reduced to equations (26)–(31). However, in equation (31), since the kinetic states, $X_F(t)$, are neglected, the kinetic disturbance, $d_\kappa(t)$, can be directly computed from the measured value of the controlled kinetic variable, $\kappa_m(t)$:

$$d_\kappa(t) = \kappa_m(t) - [C_S \cdot X_S(t) + D_S \cdot v(t)]. \tag{56}$$

Therefore, the only states that need to be evaluated are those linked to the magnetic variables, i.e. $X_S(t)$ and $d_\iota(t)$, and the observer thus reduces to equations (45) and (47) in which the last term involving kinetic variables is omitted:

$$\dot{\hat{X}}_S(t) = A_S \cdot \hat{X}_S(t) + B_S \cdot U(t) + K_{11} \cdot \left(\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t)\right) \tag{57}$$

$$\dot{\hat{d}}_\iota(t) = K_{31} \cdot \left(\bar{\iota}_m(t) + C_\iota \cdot \hat{X}_S(t) + \hat{d}_\iota(t)\right). \tag{58}$$

The kinetic states, $X_F(t)$, are also neglected in the QP problem associated with equation (53), in which $\lambda_{fast} = 0$. Then, during the prediction horizon, $X_S(t')$, $\bar{\iota}(t')$ and $d_\iota(t')$ evolve according to the augmented system (26), (28), (30), with $U(t') = T_{svd} \cdot V(t')$, subject to the actuator constraints $L(t') \cdot U(t') \leqslant M(t')$, and with the initial conditions given by the observed states, $X_S(t' = t) = \hat{X}_S(t)$ and $d_\iota(t' = t) = \hat{d}_\iota(t)$. Substituting $U_S(t')$ by its expression from equation (43), the predicted kinetic variables evolve according to

$$\kappa(t') = C_S \cdot X_S(t') + D_S \cdot v(t') + d_\kappa(t') \tag{59}$$

with constant $d_\kappa(t') = d_\kappa(t)$ between $t' = t$ and $t' = t + \tau_H$, consistently with the constant output disturbance model (see section 4.1), and where $d_\kappa(t)$ is given by equation (56). The filter states, $v(t')$, are governed by equation (42) with $U(t') = T_{svd} \cdot V(t')$ and with the initial condition $v(t' = t) = v(t)$. Now, when solving the QP problem which leads to the optimization of $V(t')$ over the time horizon, $\tau_H$, the kinetic prediction (59), in which we have neglected the fast eigenmodes, is not representative of the actual system kinetic response. Simulations have shown that a more relevant response can be restored by adding a direct feed-through component, $D_F \cdot U_F(t')$, from the fast inputs to the $\kappa(t')$ outputs in the right hand side of equation (59). This additional vector replaces the effect of the fast eigenmodes, i.e. of $X_F$ in equation (31), without the need of solving a set of differential equations to compute their evolution. It was found appropriate to choose $D_F$ proportional to the steady state gain of the fast model[9] (defined as $K_F = -A_F^{-1} \cdot B_F$) so that, in the reduced MPC controller synthesis, and after using equation (44), equation (59) becomes:

$$\begin{aligned}\kappa(t') &= C_S \cdot X_S(t') + D_S \cdot v(t') + D_F \cdot U_F(t') + d_\kappa(t') \\ &= C_S \cdot X_S(t') + (D_S - D_F) \cdot v(t') + D_F \cdot U(t') + d_\kappa(t')\end{aligned} \tag{60}$$

with

$$D_F = -\lambda_F \cdot A_F^{-1} \cdot B_F \tag{61}$$

where $\lambda_F$ is a tuneable positive parameter smaller than 1.

An advantage of using this reduced MPC kinetic controller design, and therefore of neglecting the fast model dynamics altogether, is that it reduces the dimension of the QP problem to be solved at each time step, and therefore it alleviates the real-time computation effort. This could be necessary for more demanding applications such as the simultaneous distributed control of the $q$-profile and of kinetic profiles (e.g. temperature, rotation). This reduced MPC controller is also more appropriate than the full two-time-scale MPC controller in situations where some actuators cannot respond sufficiently fast to high frequency changes of the controller commands. Examples using NBI pulse width modulation will be discussed in sections 7 and 8.

[9] We can speak of a steady state gain for the fast model although $X_F$ always vanishes in steady state. This occurs because, by definition, $U_F$ contains only high frequency components and therefore it vanishes in steady state. Hence $X_F$ also vanishes.

## 5. Simulations of $q(x)$ and $\beta_p$ control on EAST using LHCD and ICRH

To illustrate and validate the *ARTAEMIS* MPC control algorithm presented above, we shall now discuss the results of nonlinear closed loop simulations in a high-$\beta_p$ non-inductive scenario on EAST. For this particular application, the kinetic variable, $\kappa(x, t)$, defined in section 2 reduces to a single scalar parameter, $\beta_p(t)$, independent of x, and therefore the kinetic vector has only one element. The simulations were performed by inserting the METIS plasma simulator at the output of the MPC controller in a MATLAB®/Simulink model, and feeding the appropriate METIS input powers and controlled output variables back into the observer/controller Simulink blocks. Many plasma parameters or profiles such as the plasma shape, the toroidal field, $B_T$, the line-averaged plasma density, $n_{el}$, etc, will be assumed independent of the control actuators and were fixed external inputs to the METIS code, together with all the chosen METIS options for modeling the various physical phenomena. The time evolution of these parameters and profiles was based on actual experimental data from the reference EAST shot #62946, until $t = 3.2$ s when control was switched-on. They were held constant afterwards, during the control phase. Constant LHCD power at 2.45 GHz (0.6 MW) and ECRH power (0.3 MW) were used in the steady state reference shot. So, as in section 3, the 2.45 GHz LHCD power was included in the ECRH power for the METIS simulations. A constant feedforward ECRH power of 0.9 MW was thus used in all the control simulations, with the appropriate power and current deposition profiles to simulate the combined effect of the two systems. For $t \geqslant 3.2$ s, at each time step and with a cycle time of 0.02 s, the 4.6 GHz LHCD and 33 MHz ICRH actuator powers were prescribed by the MPC controller commands and the evolution of all the plasma parameters and profiles that depend on the injected power, e.g. $V_{ext}(t)$, $l_i(t)$, $\beta_p(t)$, $\psi(x, t)$, $q(x, t)$, temperature and pressure profiles, etc., were computed by the METIS simulator.

The controller filter cutoff frequency separating the slow and the fast inputs, $U_S(t)$ and $U_F(t)$, respectively, was chosen as $f_{filt} = 5$ Hz. It is slightly different from the cutoff used for the identification process in section 3 (3.5 Hz), because it was adjusted, a posteriori, with the eigenvalues and characteristic times of the identified model, and with the required ordering: $\tau_F = 0.04\,\text{s} < \tau_{filt} = 0.2\,\text{s} < \tau_S = 0.84\,\text{s}$. The number of nodes during the horizon was chosen as the minimum that can be used, i.e. $n_{nodes} = 2$. These controller parameters were unchanged for all the simulations presented in section 5. The constraint matrices $L(t)$ and $M(t)$ used at each time step when solving the QP problem (53) were constant during all the control phases, and defined as to limit the 4.6 GHz LHCD power to the interval $0 \leqslant P_{LH} \leqslant 3$ MW and the ICRH power to $0 \leqslant P_{IC} \leqslant 1.5$ MW.

Typical examples from three sets of METIS simulations will be reported in this section. They were all performed with the two-time-scale MPC controller designed in section 4 and based on the *ARTAEMIS* model identified in section 3. The first set of tests of the MPC algorithm was dedicated to the control of the $q$-profile using the 4.6 GHz LHCD actuator. It involves only the slow part of the model and of the controller since the safety factor evolves on the slow resistive diffusion time scale. The second set of tests addressed the simultaneous control of $q(x)$ and $\beta_p$ with LHCD and ICRH, using only the reduced kinetic MPC controller (section 4.4). Control of $\beta_p$ was thus achieved on the same time scale as the control of $q_0$, i.e. on the resistive time scale. In the third set, the full two-time-scale MPC controller was tested, using the LHCD and ICRH actuators to control $q(x)$ on the resistive time scale and $\beta_p$ on the kinetic time scale, simultaneously.

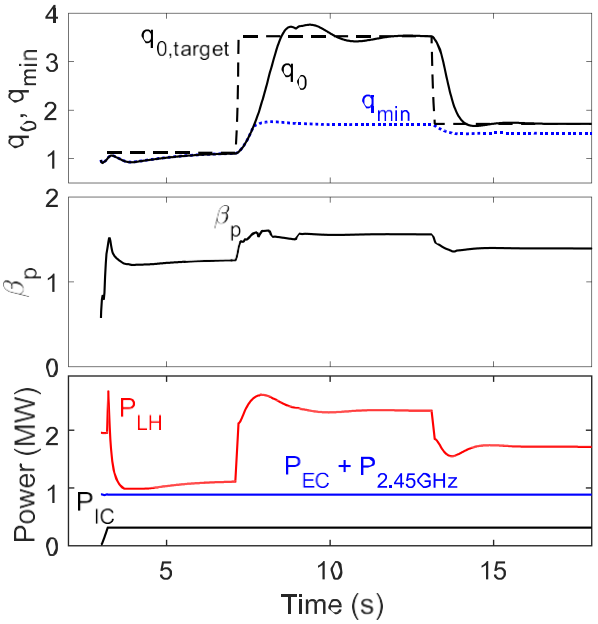


**Figure 5.** Control of $q_0$ with LHCD. Top: $q_0(t)$ (solid black), $q_{0,target}(t)$ (dashed) and $q_{min}(t)$ (dotted blue) vs time. Middle: $\beta_p$ vs time (not controlled). Bottom: 4.6 GHz LHCD actuator power (red) and feedforward powers: ICRH (black) and combined ECRH plus 2.45 GHz LHCD (blue). Control starts at t = 3.2 s.

### 5.1. Control of the safety factor with LHCD

*5.1.1. Discrete $q_0$ control.* The simplest test of the controller consists in tracking a given target value of the safety factor at a given normalized radius, using the 4.6 GHz LHCD actuator only. In this case, offset-free MPC control is possible with the observer and controller synthesis described in section 4. This was proved in [32] when the number of controlled variables is equal to the number of actuators or smaller. Here, there is no need for the kinetic equations (27)–(31) in the augmented model, and the observer reduces to equations (57) and (58). In the QP problem (53), we set $\lambda_{kin} = \lambda_{fast} = 0$ since there is no kinetic control.

At constant plasma current, the area where $q$-profile regulation is most sensitive is the plasma core, and in particular the magnetic axis, because $q(x)$ is almost clamped near the plasma boundary and its shape does not vary much in the outer half of the plasma. An example of such control, starting at $t = 3.2$ s, is displayed on figure 5. In this example, the ICRH power was fixed at 0.32 MW as in the EAST reference

discharge #62946, and the controller was requested to track successively three different target values of $q_0 = q(x = 0)$, with a prediction horizon chosen as $\tau_H = \tau_S = 0.84$ s, equal to the resistive diffusion time. The $q_0$ target waveform is represented by the dashed line on the top frame of figure 5, namely $q_{0,\text{target}} = 1.1$ in the time interval $3.2\text{ s} \leqslant t \leqslant 7.1$ s, $q_{0,\text{target}} = 3.5$ in the time interval $7.2\text{ s} \leqslant t \leqslant 13.1$ s and $q_{0,\text{target}} = 1.7$ in the time interval $13.2\text{ s} \leqslant t \leqslant 18$ s. These $q_0$ values were within the range obtained during the modulation experiments mentioned in section 3 and they were therefore accessible with the available LHCD power. During the prediction horizon, the target reshaping function was computed from equation (54) with $\tau_{\text{target}} = \tau_H = \tau_S = 0.84$ s. To control only the discrete parameter $q_0$, the weight function on $\bar{\iota}(x, t') - \bar{\iota}_{\text{target}}(x, t')$ in equation (53) was replaced by the Dirac distribution, $\mu(x) = \delta(x)$. Now, since we use only one actuator, the input vector, $U(t)$, is one-dimensional and, in equation (41), $K_S$ is a $1 \times 1$ matrix for which we can only retain one SVD component: $n_{\text{svd}} = 1$, $T_{\text{svd}} = 1$ and $V(t) = U(t)$. Finally, the results displayed on figure 5 were obtained with $R = 0$ since there was no need for a specific moderation of the LHCD actuator in solving the QP problem. Figure 5 also shows the time evolution of $q_{\min}$ (the minimum value of $q(x)$ across the plasma cross-section), and of the poloidal beta parameter, $\beta_p$, which is not controlled in this particular example. The tracking of the different $q_0$ targets in the time intervals $3.2\text{ s} \leqslant t \leqslant 7.1$ s, $7.2\text{ s} \leqslant t \leqslant 13.1$ s and $13.2\text{ s} \leqslant t \leqslant 18$ s, respectively, is performed in about 2–3 s, i.e. in a few resistive times, and without steady state offset. The time evolution of the minimum $q$-value shown on the top frame of figure 5 (dotted blue trace) indicates that the second $q_0$ target corresponds to a $q$-profile with a strong negative shear in the plasma core ($q_{\min} = 1.7$). The bottom frame of figure 5 shows the evolution of the 4.6 GHz LHCD power command delivered by the controller. The three successive targets are reached with $P_{LH} = 1.11$ MW, 2.34 MW and 1.72 MW, respectively. The constant feedforward powers from ICRH and from the combined ECRH and 2.45 GHz LHCD systems are also shown.

5.1.2. *Distributed q(**x**) control.* Distributed control of the $q$-profile can also be performed using continuous weight functions, $\mu(x)$, in the cost function $I_H(t)$ of equation (53). In this case, genuine offset-free control cannot be expected over the entire region where $\mu(x)$ is positive. Instead, the controller achieves a least-square minimization of the radially integrated error signals, as can be seen in the definition of $I_H(t)$. When $I_p$ is regulated, there is no need for an additional control of the safety factor in the outer region of the plasma. However, in the plasma interior where the $q$-profile is sensitive to the H&CD actuators it is important to select target profile shapes that are accessible (or nearly accessible) with the available actuators so that the least-square approach is meaningful. In practice such profiles can be obtained offline from open loop simulations using a plasma simulator such as METIS, or more sophisticated models.

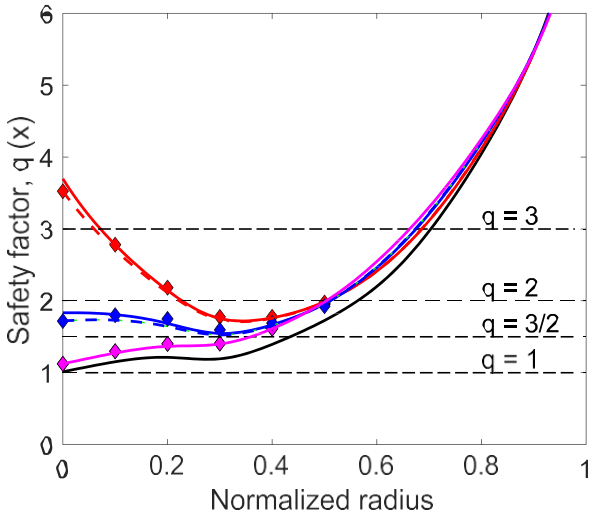


**Figure 6.** Distributed $q(x)$ control with $\mu(x) = 1$ for $0 \leqslant x \leqslant 0.5$ and with LHCD only. Achieved $q(x)$ at $t = 3.2$ s (black), 7.1 s (magenta), 13.1 s (red solid), and 18 s (blue solid). Target $q$-profiles are constant during these intervals (diamond symbols). Dashed lines are profiles achieved with $q_0$ control only (see figure 1).

For comparison with the previous case, a distributed control example is shown on figure 6 with three different $q(x)$ profile targets having the same $q_0$ values as in the previous example shown on figure 5. The other controller parameters, constraints and feedforward powers were also the same except for $\mu(x)$ where we have used piecewise linear functions defined at the radial knots of the basis functions, $x_k = 0$, 0.1, 0.2, …, 0.9, and equal to 1 for $k = 1$–6 (i.e. for $0 \leqslant x_k \leqslant 0.5$) and to 0.01 elsewhere. The target profiles are represented by diamond symbols on figure 6. They were chosen from profiles obtained at three different times in an open loop METIS simulation with modulated LHCD power (see section 3), with the same settings and the same constant ECRH (9 MW) and ICRH (0.32 MW) powers. Note that the target $q$-profiles selected from METIS simulations do not correspond to exact steady state plasma equilibria. They are not either steady state solutions of the identified *ARTAEMIS* model. Now, with only one varying parameter, $P_{LH}$, a steady state value of $q_0$, corresponding to a given LHCD power in the METIS simulations, predetermines the rest of the steady state $q(x)$ profile. Therefore, if the target $q$-profiles were true steady state solutions in METIS, the controller would drive $P_{LH}$ to the same values for the distributed $q(x)$ control and the discrete $q_0$ control since the $q_0$ targets are the same in both cases. The achieved $q_0$ values would also be the same in the two cases.

On figure 6, the black curve represents the $q$-profile at the start of the distributed $q(x)$ control, i.e. at $t = 3.2$ s. Then, the first target profile was a monotonic $q$-profile with $q_0 = q_{\min} = 1.1$ represented by magenta diamonds, and it was tracked for $3.2\text{ s} \leqslant t \leqslant 7.1$ s. The profile represented by the magenta trace is the achieved $q$-profile at $t = 7.1$ s, when the plasma has practically reached a steady state. Note that this first target profile is indeed achieved with no apparent offset. However, $P_{LH} = 1.13$ MW at $t = 7.1$ s versus

1.11 MW in the discrete control case, and a detailed analysis shows that, in both cases, the achieved $q$-profiles are not quite in a true steady state at $t = 7.1$ s. Between $t = 7.1$ s and $t = 7.2$ s the target profile changes to a negative shear safety factor profile with $q_0 = 3.5 > q_{\min} = 1.7$ (red diamonds on figure 6), until $t = 13.1$ s when it changes again to a weak shear $q$-profile with $q_0 = 1.7$ and $q_{\min} = 1.6$ (blue diamonds on figure 6). The solid red and blue traces show the $q$-profiles achieved at $t = 13.1$ s and $t = 18$ s, respectively. The controller behaves very similarly as in figure 5, but with small steady state offsets on axis and a better tracking of the target profile in the region of minimum $q$, between $x = 0.3$ and $x = 0.5$, where $q(x)$ is less sensitive to small power variations. Here, $P_{LH} = 2.38$ MW at $t = 13.1$ s and 1.79 MW at $t = 18$ s versus 2.34 MW and 1.72 MW in the discrete control case. For comparison, the dashed lines on figure 6 represent the $q$-profiles achieved at the same times in the previous case, i.e. with $q_0$ control only. The reasons why the two sets of profiles are not strictly identical are that neither the targets nor the achieved profiles are true steady states, and that METIS is highly nonlinear so that, with the same steady state power but different power history, its results may not always be strictly identical.

### 5.2 *Control of $q(\mathbf{x})$ and $\beta_p$ with LHCD, ICRH and the reduced MPC controller*

An example of the simultaneous control of $q(x)$ and $\beta_p$ on the slow (resistive) time scale, using LHCD and ICRH actuators, is displayed on figures 7 and 8. In this example, four different $\beta_p$ targets were tracked: $\beta_{p,target} = 1.2$ from $t = 4.2$ s to $t = 5.1$ s, and $\beta_{p,target} = 2$, 1.5 and 2.5 for 5.2 s $\leqslant t \leqslant$ 7.1 s, 7.2 s $\leqslant t \leqslant$ 9.1 s and $t \geqslant 9.2$ s, respectively. The $q$-profile is controlled from $t = 3.2$ s and between $x = 0$ and $x = 0.5$ with the same piecewise linear weight function as in the previous example: $\mu(x_k) = 1$ for $0 \leqslant x_k \leqslant 0.5$ and 0.01 for $x_k \geqslant 0.6$. In order to assess the robustness of the $q$-profile control at different plasma pressures (see figure 8), the target $q(x)$ is kept constant while the $\beta_p$ target changes. The $R$ matrix was set to zero and $\lambda_{kin}$ was set equal to 1 because normalized input, output and state variables are used in the controller implementation. As before, we set $\lambda_{fast} = 0$ because the fast model and its kinetic states are not used here, but we set $\lambda_F = 0.3$ in the substituted direct feed-through term of the model kinetic output equations (60) and (61).

Now, the second magnetic eigenmode has a characteristic time of 0.24 s and the $\beta_p$ response in the truncated equation (31) may be somewhat faster than $\tau_S$. Thus, we have reduced the prediction horizon and the target reshaping time to half the resistive diffusion time, $\tau_{target} = \tau_H = \tau_S/2 = 0.42$ s, to see the effect of a faster control on $q(x)$. Finally, in our METIS implementation of the internal controller model, the LHCD and ICRH input powers were normalized to 3 MW and 1.5 MW, respectively. For $t \geqslant 4.2$ s when both $q(x)$ and $\beta_p$ are controlled simultaneously, with distributed $q(x)$ control, the $K_S$ matrix in equation (41) is a 11 × 2 matrix that depends

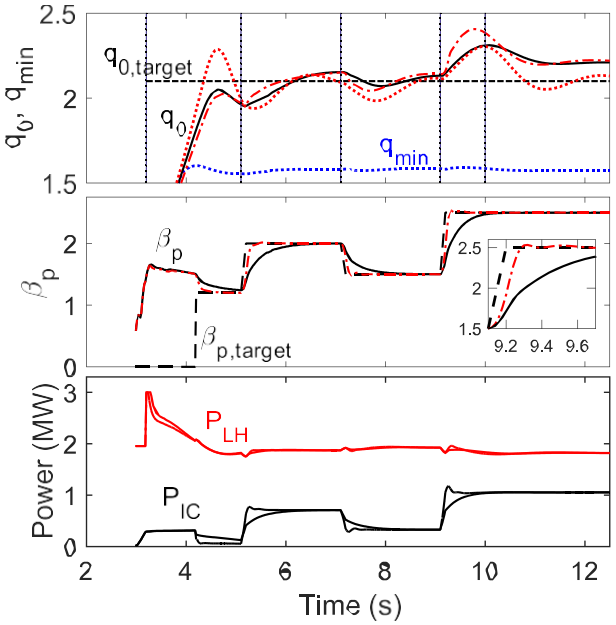


**Figure 7.** Simultaneous distributed $q(x)$ control and $\beta_p$ control with LHCD and ICRH. Control of $q(x)$ starts at 3.2 s while $\beta_p$ control starts at 4.2 s. Top: $q_0(t)$ with slow (solid black) and fast (dash-dot red) $\beta_p$ control, $q_{0,target}(t)$ (dashed black) and $q_{min}(t)$ (dotted blue). For comparison the dotted red trace is from discrete $q_0$ and slow $\beta_p$ control. Middle: $\beta_p(t)$ with slow (solid black) and fast (dash-dot red) control, and $\beta_{p,target}(t)$ (dashed black). A zoom between $t = 9.1$ s and $t = 9.7$ s is displayed in a small insert. Bottom: LHCD (red) and ICRH (black) actuator powers. In all frames, the solid traces are for slow $\beta_p$ control and the dash-dot traces are obtained with the full two-time-scale MPC controller, which yields faster $\beta_p$ control (see insert).

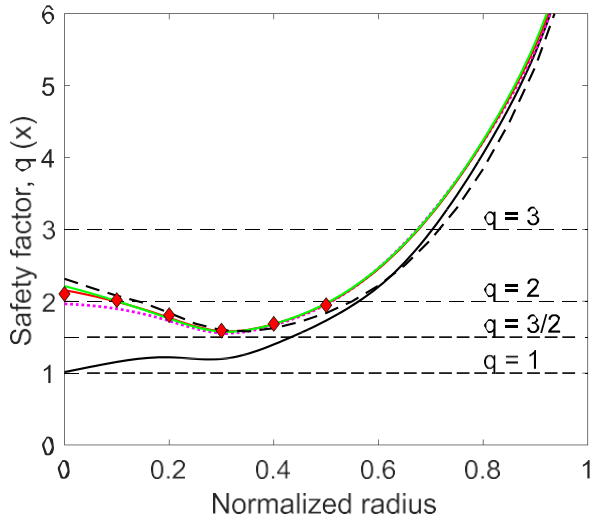


**Figure 8.** Combined distributed $q(x)$ control and slow $\beta_p$ control with LHCD and ICRH. Achieved $q(x)$ at $t = 3.2$ s (solid black), 5.1 s (dotted), 7.1 s (red), 9.1 s (blue), 10 s (dashed black) and 12.5 s (green). These times are depicted on figure 3 by vertical lines. The target $q$-profile is constant (diamond symbols) and $\mu(x) = 1$ for $0 \leqslant x \leqslant 0.5$. $q(x)$ control starts at 3.2 s while $\beta_p$ control starts at 4.2 s.

on the power normalizations, and whose positive singular values are equal, here, to 0.57 and 0.41. Therefore, its first two SVD components are significant and can be retained. We have therefore chosen $n_{svd} = 2$ and $V(t) = U(t)$.

Figure 7 shows the time evolution of $q_0$, $q_{min}$, $\beta_{p,target}$, and of the achieved $\beta_p$ and actuator powers. Figure 8 shows the target $q$-profile and the achieved $q$-profiles at the start of the control phase and at the end of each constant $\beta_{p,target}$ phase. An additional profile is shown at $t = 10$ s, which corresponds to the largest transient $q_0$ offset during the transition to the $\beta_p = 2.5$ plasma equilibrium (see figure 7 and the dashed black line in figure 8).

As mentioned before, small steady state offsets cannot be avoided with only two actuators for distributed $q(x)$ control and simultaneous control of $\beta_p$. The $q$-profile offset is mostly apparent near the magnetic axis where the safety factor is highly sensitive to any perturbation. The steady state $q_0$ offset disappears when only $q_0$ and $\beta_p$ are controlled, as shown by the dotted red trace on figure 7 (top frame) at $t = 7.1$ s, 9.1 s and 12.5 s.

### 5.3 *Faster control of $\beta_p$ using the two-time-scale ARTAEMIS model and MPC controller*

To illustrate the implementation of MPC theory in controllers using singularly perturbed dynamic models, we shall describe now the simultaneous control of $q(x)$ and $\beta_p$ including in the MPC algorithm the first-order perturbation of the identified model, i.e. the fast model and the kinetic states. The effect can already be seen on figure 7 where we have added, for comparison, the result of the fast $\beta_p$ control using the full two-time-scale MPC controller with $\lambda_{fast} = 1.5$ (red dash-dot traces) to the slow control results obtained with the same common controller parameters, and with $\lambda_F = 0.3$ in equations (60) and (61).

Another example is shown on figures 9 and 10, where we have combined a $\beta_p$ ramp request, from $\beta_p = 1.5$ up to $\beta_p = 2.5$ for $3.2\text{ s} \leqslant t \leqslant 5$ s, to the subsequent tracking of three different $q(x)$ profiles at constant $\beta_p$, either with discrete $q_0$ control (solid lines in figure 10) or with distributed $q(x)$ control (dotted lines in figure 10). Both $q(x)$ and $\beta_p$ controls start at $t = 3.2$ s. The first $q$-profile target has $q_0 = 2.1$ and $q_{min} = 1.6$, the second one has $q_0 = q_{min} = 1.4$, and the last one has $q_0 = 3.5$ and $q_{min} = 1.7$ with a large negative magnetic shear over a broad region of the plasma ($x \leqslant 0.32$). It can be noted from figure 9 that the regulation of $\beta_p$ is not perturbed by the changing $q(x)$ targets. Here, we chose the same time horizon as in section 5.1, $\tau_H = \tau_S = 0.84$ s, but we have increased the target reshaping time to $\tau_{target} = 2\tau_S = 1.68$ s in order to exhibit the difference in the approach of $q_0$ to its targets. As expected and confirmed by comparing the top frames in figures 7 and 9, increasing $\tau_{target}$ results in a smoother approach to the various $q_0$ targets. The $K_S$ matrix is a $2 \times 2$ matrix in the discrete $q_0$ and $\beta_p$ control case, with singular values equal to 1.19 and 0.49, and a $11 \times 2$ matrix in the distributed $q(x)$ and $\beta_p$ control case, with $\mu(x_k) = 1$ for $0 \leqslant x_k \leqslant 0.5$ and 0.01 for $x_k \geqslant 0.6$, with singular values equal to 0.57 and 0.41. Again, two SVD components can then be retained in both cases so we chose $n_{svd} = 2$ and $V(t) = U(t)$. Also in both cases, the $R$ matrix was set to zero, and $\lambda_{kin}$ was set equal to 1.

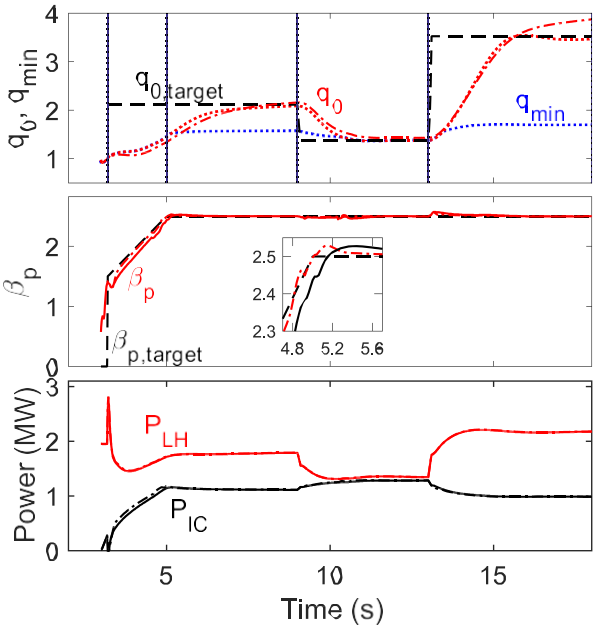


**Figure 9.** Combined $q(x)$ and $\beta_p$ control from $t = 3.2$ s using LHCD and ICRH. Top: $q_0(t)$ for discrete $q_0$ and fast $\beta_p$ control (dotted red), $q_0(t)$ for distributed $q(x)$ and fast $\beta_p$ control (dash-dot red), $q_{0,target}(t)$ (dashed black), and $q_{min}(t)$ (dotted blue). Middle: $\beta_p(t)$ for discrete $q_0$ and fast $\beta_p$ control (dash-dot), $\beta_p(t)$ for discrete $q_0$ and slow $\beta_p$ control (solid black), and $\beta_{p,target}(t)$ (dashed black). Bottom: LHCD (red) and ICRH (black) powers. In all frames, the solid traces are for slow $\beta_p$ control and the dash-dot traces are obtained with the full two-time-scale MPC controller.

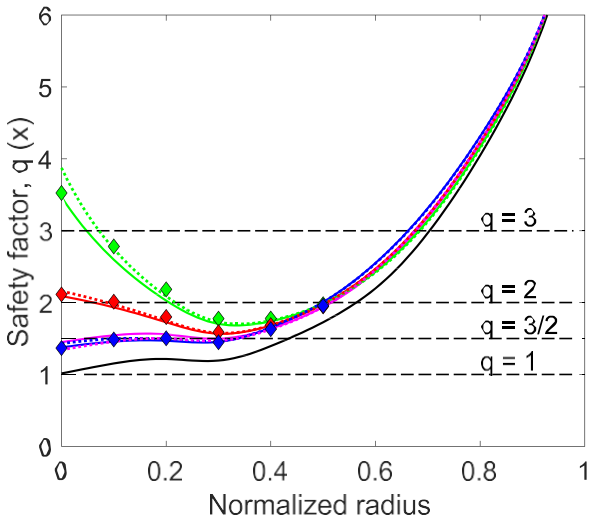


**Figure 10.** Simultaneous $q_0$ and fast $\beta_p$ control with LHCD and ICRH. Achieved profiles at $t = 3.2$ s (black), 5 s (magenta), 9 s (red), 13 s (blue), and 18 s (green). The $q(x)$ targets (diamond symbols) are held constant between these times, which are depicted by vertical lines on figure 9. Both $q(x)$ and $\beta_p$ control start at 3.2 s. The dotted lines are from distributed $q(x)$ control and fast $\beta_p$ control, for comparison.

Figure 9 (top frame) and 10 show the results obtained with the full two-time-scale MPC controller, for both discrete $q_0$ and $\beta_p$ control and distributed $q(x)$ and $\beta_p$ control, with $\lambda_{fast} = 3$. As in figure 7, the $q$-profile offset is mostly apparent for $t > 15$ s near the magnetic axis and it disappears when only $q_0$ and $\beta_p$ are controlled (dotted red trace on the top frame of figure 9). The middle frame in figure 9 shows the difference between the evolution of $\beta_p$ during the slow ($\lambda_F = 0.3$) and fast $\beta_p$ control, with simultaneous distributed $q(x)$ control in both cases. The gain in using the full two-time-scale MPC controller can be assessed by looking at the small insert where the $\beta_p$ traces are magnified.

Finally, on figures 11 and 12, combinations of four different values of $\beta_{p,target} = 1.5$, 2, 2.5 and 3, and three different $q$-profiles are tracked successively using the same distributed $q(x)$ control parameters as before, with $\tau_H = \tau_S = 0.84$ s, $\tau_{target} = 2\tau_S = 1.68$ s, $n_{svd} = 2$ (the singular values of the $11 \times 2$ $K_S$ matrix are 0.57 and 0.41) and $\lambda_{kin} = 1$. One can compare the results obtained with fast $\beta_p$ control ($\lambda_{fast} = 3$, dash-dot lines) and with slow control only ($\lambda_F = 0.3$, solid lines). Again, the benefit of using the full two-time-scale MPC controller can be assessed from the small insert where the $\beta_p$ traces are magnified. The $q$-profile targets are all reached in about 2.5 s (~3 resistive times) and the actuators adjust to reach the various $\beta_p$ targets within about 0.2 s ($\sim 5\tau_F$) with fast control and 0.4–0.5 s ($\geqslant 10\tau_F$ or $0.5\tau_S$) with slow control, while restoring the desired $q$-profile shape after each large $\beta_p$ perturbation.

A last remark can be made concerning the relative effects of the actuators on the evolution of $q(x)$ and $\beta_p$. From the bottom frames of figures 9 and 11, the $q$-profile is clearly more sensitive to $P_{LH}$ while $\beta_p$ is more sensitive to $P_{IC}$. This is, however, peculiar to the chosen LHCD and ICRH systems, and in the configuration in which they were used in the high-$\beta_p$ experiments on EAST. It does not represent the general kinetic control case with a larger number of actuators, for which the present controller has been designed. Even in the simulations described above, $P_{LH}$ and $P_{IC}$ cannot be fully mapped to the individual controlled variables, $q(x)$ and $\beta_p$, so the coupled magnetic and kinetic controller synthesis is relevant.

## 6. Initial SISO control experiments on EAST using the *ARTAEMIS* MPC controller

The reduced *ARTAEMIS* MPC controller defined in section 4.4 was used for the first time on the EAST tokamak at a time when only one actuator (LHCD at 4.6 GHz hereafter referred to as LH2) was available, with enough dynamics and reliability when the PCS commands were bound to the range 1 MW to 2.5 MW. ECRH feedforward power was also available from two gyrotrons delivering 0.9 MW at 140 GHz, and the 2.45 GHz LHCD system (LH1) could also be used for short 0.5 MW pulses. In this configuration, single-input-single-output (SISO) control was tested, either for the central safety factor, $q_0$, or for the poloidal beta parameter, $\beta_p$. In both

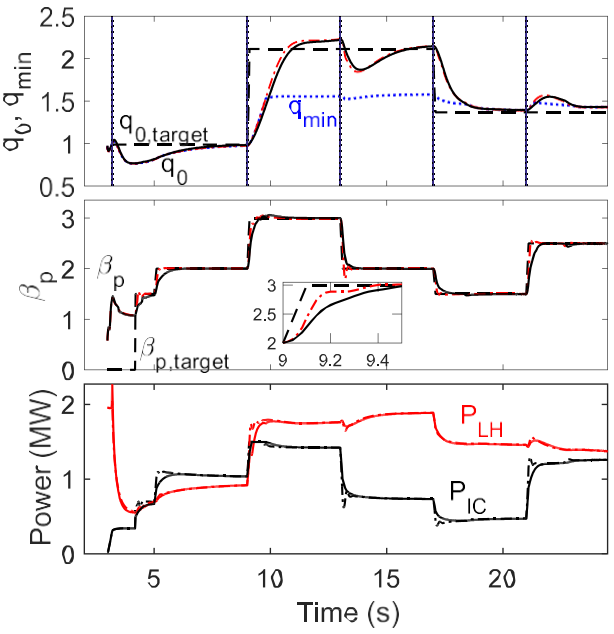


**Figure 11.** Simultaneous control of $q(x)$ and $\beta_p$ with LHCD and ICRH. Control of $q(x)$ starts at 3.2 s while $\beta_p$ control starts at 4.2 s. Top: $q_0(t)$ for distributed $q(x)$ control and slow $\beta_p$ control (solid black), $q_0(t)$ for distributed $q(x)$ control and fast $\beta_p$ control (dash-dot red), $q_{0,target}(t)$ (dashed black) and $q_{min}(t)$ (dotted blue). Middle: $\beta_p(t)$ for distributed $q(x)$ control and slow $\beta_p$ control (solid black), $\beta_p(t)$ for distributed $q(x)$ control and fast $\beta_p$ control (dash-dot red) and $\beta_{p,target}(t)$ (dashed black). Bottom: LHCD (red) and ICRH (black) actuator powers. In all frames, the solid traces are for slow $\beta_p$ control and the dash-dot traces are obtained with the full two-time-scale MPC controller.

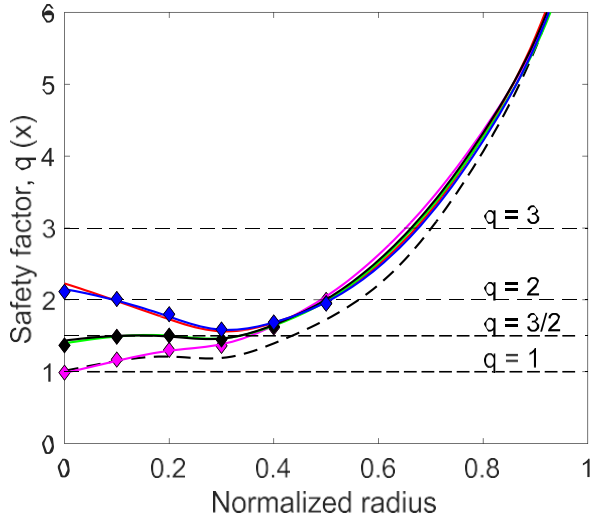


**Figure 12.** Simultaneous distributed $q(x)$ and fast $\beta_p$ control with LHCD and ICRH, and with $\mu(x) = 1$ for $0 \leqslant x \leqslant 0.5$. Control of $q(x)$ starts at 3.2 s while $\beta_p$ control starts at 4.2 s. Achieved $q(x)$ at $t = 3.2$ s (dashed black), 9 s (magenta), 13 s (red), 17 s (blue), 21 s (green), and 25 s (solid black). The $q(x)$ targets (diamond symbols) are held constant between these times, which are depicted by vertical lines on figure 11.

cases, offset-free MPC is possible with the controller synthesis described in section 4.

### 6.1. ARTAEMIS model for SISO control experiments

The *ARTAEMIS* system identification procedure described in section 3 was applied here to real experimental data in a typical high-density single-null H-mode scenario. In this scenario, the toroidal field was $B_T = 2.5$ T, the plasma current was $I_p = 350$ kA and the line-averaged electron density was $n_{el} \sim 4.2 \times 10^{19}$ m$^{-3}$. In all shots, a small additional 0.5 MW LH1 pulse was injected from 0.95 s to 2.25 s in order to reduce the poloidal flux consumption during the plasma current ramp-up, and a constant 0.9 MW ECRH power was injected from the two gyrotrons during the current flattop. The system identification data was obtained from a steady state reference discharge with a constant LH2 power at 1 MW (shot #93295) and from three other discharges, with chirping frequency (shots #93296 and #93297) and pseudo-random binary sequence (PRBS) modulations of the LH2 actuator (shot #93298), with powers between 1 and 2.5 MW.

To cope with the nonlinear response of the LH2 system to the PCS command, a simple additional PI feedback controller was used to track the desired power modulations. The PI gains were chosen following the so-called SIMC tuning rule [44] and the actuator controller was also used in cascade with the MPC controller module to track the MPC power command in the closed loop experiments. The magnetic and kinetic parameters of interest were estimated by the real-time EFIT reconstruction code.

A linear two-time-scale state space model having the *ARTAEMIS* structure defined in sections 2 and 3, with five significant eigenmodes in the slow model and only one eigenmode in the fast model, was identified from the modulation data. Both the input and output data were filtered with a non-causal low-pass filter of order 5, and a cutoff frequency, $f_{filt} = 3.3$ Hz for the separation between the slow and fast data. The fit parameters, $f$, obtained from equation (18) in the four discharges were found as 84%, 51%, 63%, and 67% for $\bar{\iota}(x)$ and as 57%, 49%, 75% and 73% for $\delta\theta_p$. They are somewhat smaller and more broadly distributed than the values reported in section 3, obtained with simulation data. This is due to the noise and uncertainties inherent to the experimental data, and also to undesired variation of some plasma parameters in the experiments. Nevertheless, as in figures 2 and 3, the identified model was found to reproduce satisfactorily the coupled evolution of the internal poloidal flux profile, $\psi(x,t)$, of the inverse of the safety factor profile, $\bar{\iota}(x,t) = 1/q(x,t)$ and of the poloidal beta parameter, $\theta_p(t)$. Another example of the *ARTAEMIS* system identification using experimental data will be given in section 8. Two important parameters characterizing the identified models are the largest (negative) eigenvalues of $A_S$ and $A_F$, obtained here as $-0.95$ s$^{-1}$ and $-60.6$ s$^{-1}$, which correspond to time constants $\tau_S = 1.05$ s and $\tau_F = 0.017$ s for the resistive diffusion and thermal transport, respectively.

### 6.2. Experimental SISO control of $q_0$ with LHCD

The first MPC control experiment aimed at tracking piecewise-constant waveforms for the safety factor on axis, $q_0$. The control cycle time was set to 0.02 s so that, at each time step, we had enough CPU time to reconstruct the plasma magnetic equilibrium from the real-time measurements and solve the QP problem in the MPC algorithm. In the discharge #94832, the target $q_0$ was set at 2.4 from $t = 2.7$ s to $t = 4.5$ s and was raised to $q_{0,target} = 2.8$ at $t = 4.52$ s.

The controller parameters were chosen as $n_{svd} = 1$ (since there is only one actuator), $R = 0$, $\tau_H = 0.52$ s $\sim \tau_S/2$ for the horizon, and $\tau_{target} = 0.26$ s $\sim \tau_S/4$. The target reshaping time is smaller than the value chosen for the simulations in section 5 in the aim of approaching the targets faster. The $q$-profile weight function, $\mu(x)$, was replaced by the Dirac distribution, $\delta(x)$, and the kinetic weights $\lambda_{kin}$ and $\lambda_{fast}$ were set to 0 since there is no kinetic control in this experiment. The time evolutions of $q_0$ (solid black trace) and of the piece-wise constant $q_0$ target (dashed red trace) are shown on figure 13(*a*). The LH2 power commands obtained from the controller (dotted magenta trace) and at the output of the PI actuator control module (dashed red trace) are shown on figure 13(*b*).

The LH2 power effectively coupled to the plasma is also shown on figure 13(*b*) (solid black trace). The coupled power closely matches the original controller command (and therefore the MPC algorithm request), showing the effectiveness and appropriate tuning of the actuator control module. The tracking of the $q_0$ targets is successfully performed without steady state offset, in about one second, which is of the order of the resistive diffusion time. This is relatively shorter than in the simulations of section 5 thanks to the smaller target reshaping time. In these experiments, the inputs to the observer, $U(t)$, $\bar{\iota}_m(t)$ and $\kappa_m(t)$ in equations (45)–(48), were taken from noisy raw measurements and both the power command (figure 13(*a*)) and the controlled safety factor (figure 13(*b*)) were affected by high-frequency noise.

### 6.3. Experimental SISO control of $\beta_p$ with LHCD

In other discharges, $\theta_p$ SISO control was tested. An important feature of plasma transport physics, on which the postulated two-time-scale *ARTAEMIS* model structure is based (section 2), is that the slow evolution of the kinetic variables, $\kappa(t)$, is not only driven by low frequency components of the heating powers, but it is also partly driven by the evolving magnetic equilibrium, i.e. by the poloidal flux profile, $\psi(x,t)$. This is why the dynamic model couples both $\psi(x,t)$ and $\theta_p(t)$ to the slow states of the system through the $C_\psi$ and $C_S$ matrices in equations (11) and (12), in which the kinetic vector, $\kappa$, now stands for $\theta_p$. The slow states of the system, $X_S(t)$, are defined as the eigenmodes of the resistive diffusion equation. They govern the dynamics of $\psi(x,t)$ and $\bar{\iota}(x,t)$but also the slow dynamics of $\theta_p(t)$, and therefore they must be kept in the controller internal model even when $q(x)$ is not controlled. Now, the MPC controller described in section 4 was designed for the control of the safety factor profile, $q(x)$,

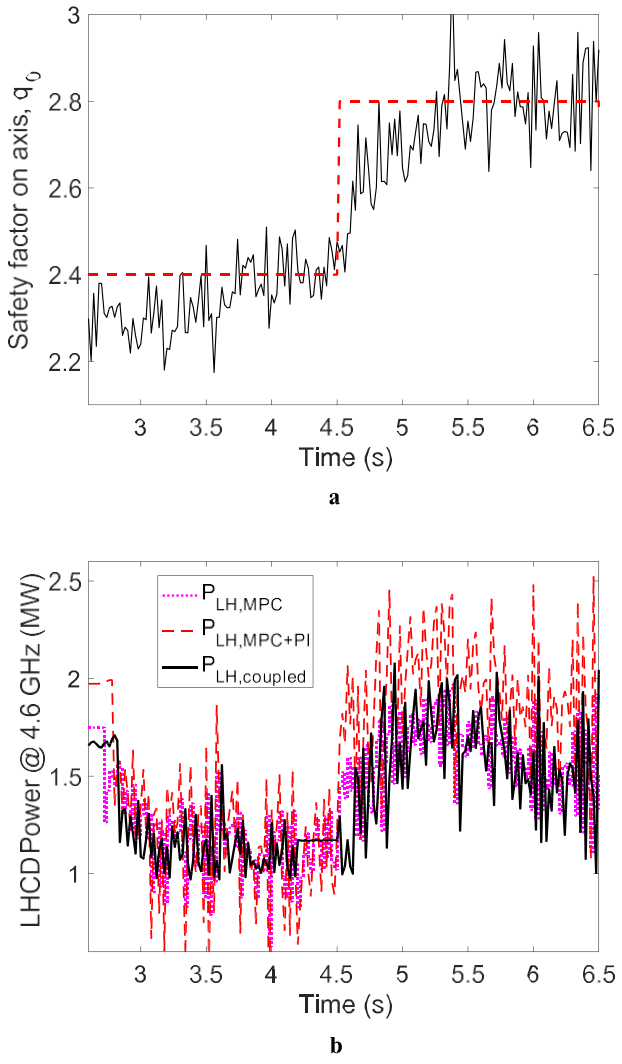


**Figure 13.** (*a*) Safety factor on the magnetic axis, $q_0(t)$, from real-time EFIT magnetic equilibrium reconstruction (solid) and $q_{0,\mathrm{target}}(t)$ (dashed red) in EAST discharge #94832. Control starts at $t = 2.7$ s. (*b*) Coupled LH2 power (solid black) and MPC command before (dotted magenta) and after (dashed red) the PI actuator control in EAST discharge #94832. The coupled power (solid) closely matches the original controller command (dotted magenta). Control starts at $t = 2.7$ s.

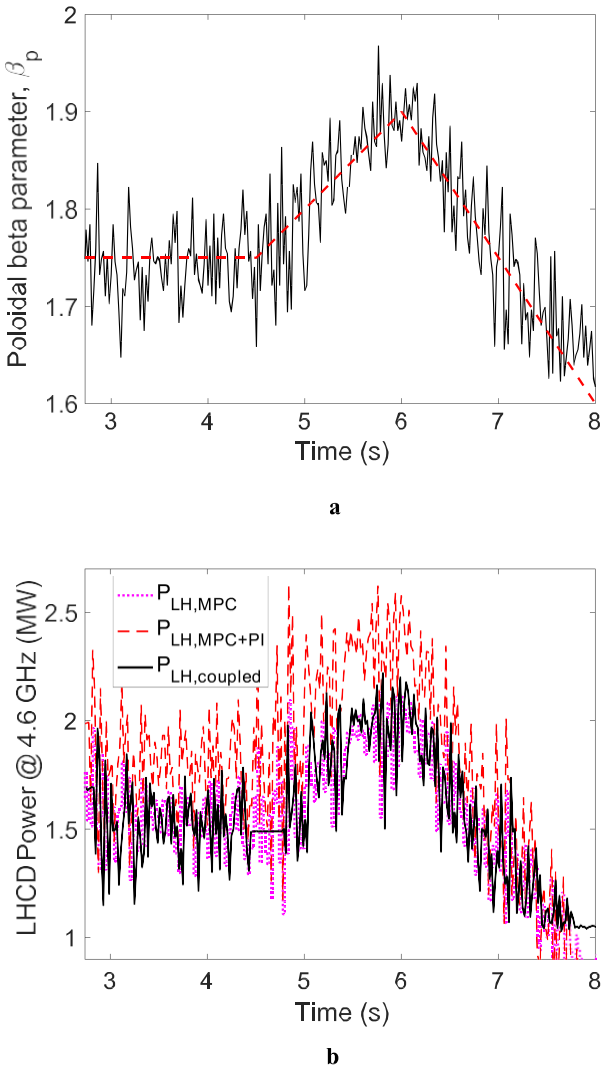


**Figure 14.** (*a*) Poloidal beta parameter, $\beta_p(t)$, from real-time EFIT magnetic equilibrium reconstruction (solid) and $\beta_{p,\mathrm{target}}(t)$ (dashed red) in EAST discharge #94829. Control starts at $t = 2.7$ s. (*b*) Coupled LH2 power (solid black) and MPC command before (dotted magenta) and after (dashed red) the PI actuator control in EAST discharge #94829. The coupled power (solid) closely matches the original controller command (dotted magenta). Control starts at $t = 2.7$ s.

or for the simultaneous control of $q(\mathrm{x})$ and of a kinetic vector, $\kappa$, which is here equal to $\beta_p$. It can also be used for the control of $\psi(\mathrm{x})$ and $\beta_p$ (similarly to what was done in [33] for the control of $\psi(\mathrm{x})$ and $\beta_N$) by keeping the vector $\psi(t)$ in the model equations (21) and (23) and in the following controller/observer equations, instead of using the vector $\bar{\iota}(t)$ and the model equations (28) and (30). The simplest way to achieve $\beta_p$ SISO control with the multi-variable controller described in section 4, is indeed to select $\psi(\mathrm{x})$ and $\beta_p$ as controlled variables and to increase the weight $\lambda_{\mathrm{kin}}$ (and optionally $\lambda_{\mathrm{fast}}$) in the QP cost function to values much larger than the norm of $\mu(\mathrm{x})$.

In the discharge #94829, a piecewise-linear $\beta_p$ target waveform with $1.6 < \beta_{p,\mathrm{target}} < 1.9$ was tracked. The controller parameters were chosen as $n_{\mathrm{svd}} = 1$ (there is only one actuator), $R = 0$, $\tau_{\mathrm{H}} = 0.52\,\mathrm{s} \sim \tau_s/2$, $\tau_{\mathrm{target}} = 0.26\,\mathrm{s} \sim \tau_s/4$. The weight function on the $\psi(\mathrm{x})$ control was the Dirac distribution, $\mu(\mathrm{x}) = \delta(\mathrm{x})$, and the kinetic weight, $\lambda_{\mathrm{kin}}$, was set to 1000 in order to make the contribution of $\psi(\mathrm{x})$ in the controller QP cost function negligible, given that normalized variables are used. The effect of the fast system state on the controller commands was also negligible with $\lambda_{\mathrm{fast}} = 1.5 \ll \lambda_{\mathrm{kin}}$. The controller filter that separates the slow and fast variables was chosen with $\tau_{filt} = 0.1$ s ($f_{filt} = 10$ Hz) and the number of nodes during the horizon was $n_{\mathrm{nodes}} = 2$. The time evolution of $\beta_p$ (solid black trace) and its target waveform (dashed red trace) are shown on figure 14(*a*). The LH2 power commands obtained from the controller (dotted magenta trace) and at the output of the PI actuator control module (dashed red

trace) are shown on figure 14(*b*), together with the LH2 power effectively coupled to the plasma (solid black trace). As in the previous case (section 6.2), the coupled power closely matches the MPC algorithm request. The $\beta_p$ waveform tracking was excellent despite significant noise in the measured observer/-controller inputs.

## 7. Simulations of two-time-scale MPC control using LHCD, ECRH and NBI pulse width modulation

The initial experiments described in section 6 were limited to SISO control. When the real-time command of the neutral beam injectors became available on EAST, the *ARTAEMIS* MPC control algorithm could be tested in the more sophisticated multiple-input-multiple-output (MIMO) control configuration. In the perspective of such experiments, METIS simulations were performed using several actuators including NBI to simultaneously control the $q$-profile at different radii and the poloidal beta parameter, $\beta_p$. For these simulations, up to four actuators were used: LHCD at 4.6 GHz (LH2), two sets of co-current NBI injectors at different angles (NB1 and NB2) and ECRH at 140 GHz. The actuators have been tuned in METIS so that they have nearly the same characteristics and the same effect on the plasma as the H&CD systems used on EAST. At present, real-time ECRH actuation is not yet available from the EAST PCS but, anticipating on future developments, ECRH was used as an actuator in METIS simulations in order to extend the controller capabilities. Also, unlike the LHCD klystrons, which can deliver a continuous power waveform within given lower and upper bounds, NBI injectors can either be running at full power or switched off, and switching from one state to the other can only be done after a minimum time. The MPC algorithm must therefore be complemented with a PWM algorithm that delivers NBI pulses in such a way that, when averaged over a given period of time, the injected power reproduces at best the power that is requested by the controller at the beginning of each averaging period. The NBI power request can therefore be updated only after a number of time samples while the LH2 power request is updated at each time step. The minimum pulse width that can be handled by the EAST NBI system is 0.1 s, which is larger than the characteristic thermal diffusion time and is not best suited for the control of kinetic parameters. So, in order to reduce the time delay in the NBI action, we have assumed that the minimum switching time can be reduced to 0.04 s in the simulation tests presented here. The power averaging time was chosen as 0.4 s so that the NBI power request from the MPC controller is distributed over a maximum pattern of ten on/off periods. Thus the averaged power delivered by each beam during the averaging period can only vary by steps of 10% of the available beam power. With a controller cycle time of 0.02 s, and an averaging period of 0.4 s, the NBI MPC command is updated every twenty cycles. For the other actuators, the commands are updated at each time step, and they are optimally computed by solving the QP problem with the partial knowledge of the NBI power pattern during the time horizon of the model prediction.

We shall now describe the results of closed loop $q(x)$ and $\beta_p$ control simulations in a high-$\beta_p$ non-inductive scenario on EAST with the NBI actuators used in the PWM mode. As in section 5, the simulations were performed by inserting the METIS plasma simulator at the output of the MPC controller in a MATLAB®/Simulink model, and feeding the appropriate METIS outputs back into the controller. Many plasma parameters or profiles are assumed independent of the actuators and are fixed external inputs to the code, together with all the chosen METIS options for modeling the various physical phenomena. The reference scenario around which the linear response model was identified was based on a real high-density discharge (#94429), a steady state single-null H-mode discharge, at a toroidal magnetic field, $B_T = 2.5$ T, line-averaged electron density, $n_{el} = 4.4 \times 10^{19}$ m$^{-3}$, and plasma current, $I_p = 0.4$ MA. This reference discharge uses NBI and has a higher plasma density than that used in sections 3 and 5. It was obtained using a 0.5 MW LH1 prelude from 0.95 s to 2.25 s during the plasma current ramp-up. Then, during the current flattop, 0.9 MW from ECRH, 1.6 MW from LH2 and $2 \times 1.3$ MW from the two sets of co-current injectors at a beam voltage of 65 kV were injected in steady state. The steady state poloidal beta and internal inductance parameters were $\beta_p = 2.3$ and $l_i = 0.75$, respectively, and the $q$-profile exhibited a negative shear in the plasma core, with a minimum $q$ around 3.1 at $x \sim 0.4$ and $q_0 \sim 4$ on axis. In order to identify the various matrices in the *ARTAEMIS* model corresponding to the selected operation scenario, nonlinear open loop simulations with various PRBS power modulation waveforms were performed. Three slow eigenmodes and one fast eigenmode could be identified from the data. The largest (negative) eigenvalues of $A_S$ and $A_F$ were $-1.6$ s$^{-1}$ and $-24.7$ s$^{-1}$, which correspond to time constants $\tau_S = 0.6$ s and $\tau_F = 0.040$ s for the resistive and thermal diffusion characteristic times, respectively.

In all the closed loop control simulations, the time evolution of plasma parameters and profiles was based on the actual experimental data from shot #94429 until $t = 3.2$ s when control was switched-on. For $t \geqslant 3.2$ s, with a cycle time of 0.02 s, the LH2, the ECRH and the two co-current NBI actuator powers were prescribed by the MPC controller commands (using PWM with an amplitude of 1.3 MW for NB1 and NB2), and the evolutions of the plasma parameters and profiles (e.g. $\psi(x, t)$, $q(x, t)$, $\beta_p$, $l_i$, etc) were simulated. The controller filter for the separation of the slow and fast control was chosen with $\tau_{filt} = 1/f_{filt} = 0.22$ s and the number of nodes during the horizon was $n_{nodes} = 2$. These parameters have been unchanged in all the simulations.

### 7.1 Control of the safety factor profile

With three independent actuators (LH2, NB1 and NB2), one could expect to control the safety factor profile at three different radii. However, with the three actuator powers normalized to 1.3 MW, the singular value decomposition of the model steady state response, i.e. of the $K_S$ matrix, limited to three $q$-profile outputs at $x = 0$, 0.3 and 0.6, yields singular values of 0.494, 0.016 and 0.005. For only two outputs at $x = 0$

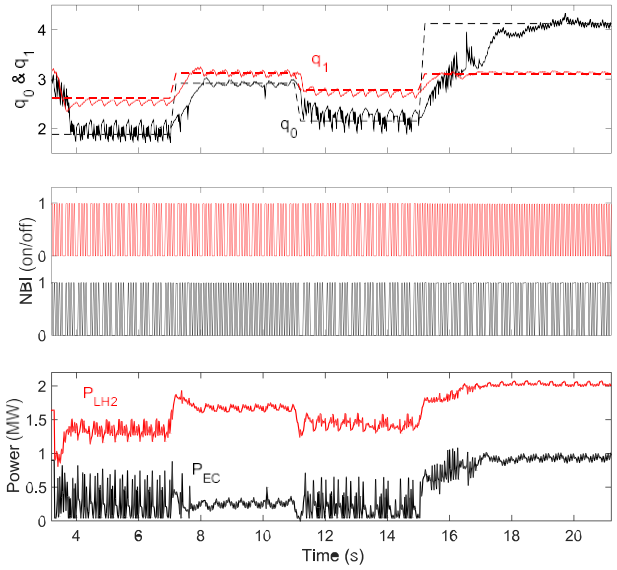


**Figure 15.** Simultaneous control of $q_0$ and $q_1 = q(\mathrm{x} = 0.4)$ with LHCD, ECRH and NBI. Top frame: $q_0$ (solid black trace), $q_{0,\mathrm{target}}$ (black dashed), $q_1$ (solid red trace) and $q_{1,\mathrm{target}}$ (red dashed) vs time. Middle frame: NB1 (bottom, black) and NB2 (top, red) PWM (0 = off, 1 = on). Bottom frame: LH2 (red) and ECRH (black) actuator powers.

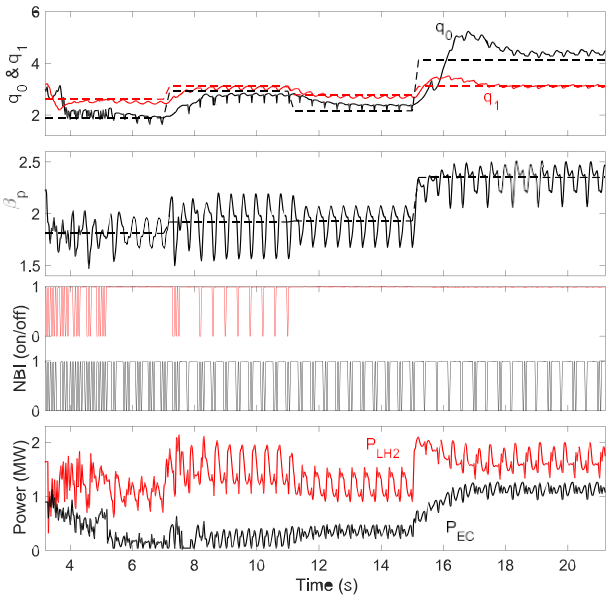


**Figure 16.** Simultaneous control of $q_0$, $q_1 = q(\mathrm{x} = 0.4)$ and $\beta_\mathrm{p}$ with LHCD, ECRH and NBI. Top frame: $q_0$ (solid black trace), $q_{0,\mathrm{target}}$ (black dashed), $q_1$ (solid red trace) and $q_{1,\mathrm{target}}$ (red dashed) vs time. Second frame: $\beta_\mathrm{p}$ (solid black trace), $\beta_{\mathrm{p,target}}$ (black dashed). Third frame: NB1 (bottom, black) and NB2 (top, red) PWM (0 = off, 1 = on). Bottom frame: LH2 (red) and ECRH (black) actuator powers.

and $\mathrm{x} = 0.4$ the singular values are 0.472 and 0.014. A strong drop between the first and second singular values of $K_S$ means that there is significant redundancy in the plasma responses to the different actuators. Adding that the PWM NBI actuation cannot be accurate and tends to produce oscillations at the period of the beam averaging time, it is difficult to control the $q$-profile at more than one radius without an extra degree of freedom from another actuator. We shall here anticipate that, in the future, ECRH can ultimately be controlled by the EAST PCS in a relatively high frequency PWM mode ($\gg$ 50 Hz). So, we shall assume in METIS that an additional ECRH actuator can deliver a continuous power. Then, with this additional actuator also normalized to 1.3 MW, the $K_S$ singular values are 0.578 and 0.109 for $q(\mathrm{x})$ control at $\mathrm{x} = 0$ and $\mathrm{x} = 0.4$, so we can use $n_{\mathrm{svd}} = 2$ in the controller.

A closed loop simulation result is displayed on figure 15 where the controller is tracking four different target pairs for $q_0$ and $q_1 = q(\mathrm{x} = 0.4)$. The targets are $[q_{0,\mathrm{target}}, q_{1,\mathrm{target}}] = [1.9, 2.6]$, $[2.9, 3.1]$, $[2.2, 2.8]$ and $[4.1, 3.1]$ during the time periods $t = 3.2$–7 s, 7–11 s, 11–15 s and 15–21 s, respectively. These targets correspond to monotonic, weak shear and negative shear profiles. The controller parameters were chosen as $n_{\mathrm{svd}} = 2$, $R = 0$, $\tau_{\mathrm{H}} = 0.8$ s and $\tau_{\mathrm{target}} = 0.4$ s. The weight function, $\mu(\mathrm{x})$, was replaced by the sum of Dirac distributions, $\delta(\mathrm{x})+\delta(\mathrm{x}\text{-}0.4)$, and $\lambda_{\mathrm{kin}} = \lambda_{\mathrm{fast}} = 0$ since there is no controlled kinetic variable. The constraint matrices $L$ and $M$ in the minimization of the cost function (53) were defined as to limit the LH2 power to $P_{\mathrm{LH2}} \leqslant 2.5$ MW, the ECRH power to $P_{\mathrm{EC}} \leqslant 2$ MW and the beam power to $P_{\mathrm{NB1}} \leqslant 1.3$ MW and $P_{\mathrm{NB2}} \leqslant 1.3$ MW. The control of $q_0$ and $q_1$ is excellent (see figure 15). The time required for tracking the different targets is consistent with the resistive time and there is no steady state offset, as could be expected with $n_{\mathrm{svd}} = 2$ for two controlled variables. The middle frame shows the evolution of the two NBI PWM waveforms (0 = power off, 1 = power on), and the bottom frame shows the LH2 and the ECRH powers requested by the controller.

### 2 *Simultaneous control of q(**x**) and $\beta_\mathrm{p}$*

For the simultaneous control of $q(\mathrm{x})$ and $\beta_\mathrm{p}$ with LHCD, ECRH and the two sets of co-current NBI injectors, the singular value decomposition of the $K_S$ matrix, limited to three outputs, $q_0$, $q_1 = q(\mathrm{x} = 0.4)$ and $\beta_\mathrm{p}$, with equal normalizations of the four actuators, yields singular values that are equal to 1.13, 0.229 and 0.096. With $n_{\mathrm{svd}} = 3$, simultaneous control of three parameters such as $q_0$, $q_1$ and $\beta_\mathrm{p}$ should therefore be possible without offsets. In the example shown on figure 16, four different sets of $q_0$, $q_1$ and $\beta_\mathrm{p}$ targets were tracked, namely $[q_{0,\mathrm{target}}, q_{1,\mathrm{target}}, \beta_{\mathrm{p,target}}] = [1.9, 2.6, 1.81]$, $[2.9, 3.1, 1.92]$, $[2.2, 2.8, 1.93]$ and $[4.1, 3.1\ 2.35]$ during the time intervals $t = 3.2$–7 s, 7–11 s, 11–15 s and 15–21 s, respectively. As before, the controller parameters were chosen as $R = 0$, $\tau_{\mathrm{H}} = 0.8$ s, $\tau_{\mathrm{target}} = 0.4$ s and the weight function, $\mu(\mathrm{x})$, was replaced by the sum of Dirac distributions, $\delta(\mathrm{x})+\delta(\mathrm{x}\text{-}0.4)$. But here, for the additional control of $\beta_\mathrm{p}$, we chose $\lambda_{\mathrm{kin}} = 1$ and $\lambda_{\mathrm{fast}} = 1.5$. The PWM algorithm introduces a 0.4 s delay and generates low frequency on/off power jumps in the NBI controller actuation. The resulting oscillations of the controlled variables around their targets, at a frequency of 2.5 Hz, are small and acceptable

for the control of $q(x)$, but their amplitude is large for $\beta_p$, showing limited kinetic control accuracy. Increasing the maximum frequency at which NBI can be switched on and off and decreasing the power averaging time would be necessary for better kinetic control.

## 8. MPC kinetic control experiments on EAST using LHCD and NBI pulse width modulation

Experimental tests of the *ARTAEMIS*/MPC kinetic control algorithm were done in the MIMO configuration during a recent experimental campaign on EAST. The neutral beam injectors and the LHCD klystrons could be controlled in real time by PCS commands so that a maximum of three independent actuators could be used for plasma control. The first actuator was the LHCD system at 4.6 GHz (LH2), as in the initial SISO experiments reported in section 6, and two additional actuators were obtained by grouping beams that have the same injection angle with respect to the plasma. A total of four deuterium beams can be injected through ports A and F of the EAST tokamak with two beams in each port, the right beam and the left beam. The geometry of the beams is the same in both ports, but the right beams (R-beams) and the left beams (L-beams) are tilted in the co-current direction at different angles with respect to perpendicular injection. As a result, the injection from the L-beams is more tangential than the injection from the R-beams and we can consider the L-beams and the R-beams as two different actuators named NB-L and NB-R, respectively. Each actuator consists of two equivalent injectors driven at the same beam voltage, one through the A-port and the other one through the F-port. The power delivered by the NB-L actuator is the sum of the powers injected by the two L-beams from ports A (NB1-L) and F (NB2-L), irrespectively of the contribution from each beam. The same applies for the NB-R actuator and the two R-beams (NB1-R and NB2-R).

An important difference between these experimental tests and the simulations discussed in section 7 is that the minimum allowed pulse width from the EAST NBI system is 0.1 s instead of 0.04 s as assumed above for the simulations. The power averaging time was not increased above 0.4 s in order to keep the same delay in the NB-L and NB-R action as in the simulations. However, this choice was made at the expense of a less accurate tracking since the NBI power request from the MPC controller was now distributed over a maximum PWM pattern of only four on/off periods of 0.1 s and, therefore, the averaged power delivered by each beam during the averaging period can only vary by steps of 25% of the available beam power. Given that there are two beams in each NBI actuator, the averaged actuator power in a 0.4 s period can vary by minimum steps of 12.5% of the total actuator power.

### 8.1. *Identification of ARTAEMIS models from experimental EAST data*

New experiments were dedicated to the identification of *ARTAEMIS* models with the LH2, NB-L and NB-R inputs. The reference operation scenario around which the models were obtained was based on a negative toroidal field discharge (#122565) with $B_T = -2.5$ T, $I_p = 0.4$ MA and a line-averaged plasma density, $n_{el} = 4.5 \times 10^{19}$ m$^{-3}$. In order to reduce the poloidal flux consumption during the plasma current ramp-up and extend the current flat-top duration from $t = 2.2$ s until $t = 13.5$ s, a short 0.5 MW pulse was injected from the 2.45 GHz LHCD system (LH1) between $t = 1$ s and $t = 2.5$ s. Then, the following powers were injected during the current flat-top: ECRH (1.8 MW from three gyrotrons at 140 GHz) for $t \geqslant 5.1$ s, ICRH (1.5 MW from two antennas) for $t \geqslant 5.4$ s, and LHCD at 4.6 GHz (LH2) with a constant 1.2 MW power flat-top between $t = 2.6$ s and $t = 5$ s, a power ramp from 1.2 MW to 2 MW between $t = 5$ s and $t = 6$ s and a constant power of 2 MW for $t \geqslant 6$ s. Stable and reproducible stationary H-mode plasmas could be obtained between $t = 6$ s and $t = 13.5$ s using this operation scenario.

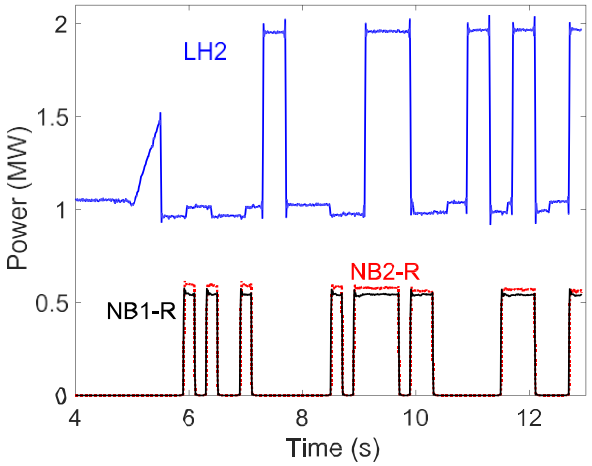


**Figure 17.** Slow LH2, NB1-R and NB2-R power modulations injected from PRBS commands on discharge #126960. The NB-L actuator is off.

Many successful discharges were then run with the same parameters and the same feedforward power sequence from LH1, ECRH and ICRH, but with power modulations from the control actuators. The LH2 power was modulated between 1 MW and 2 MW, and the power delivered by the four beams was modulated between 0 MW and full power, i.e. ~ 0.6 MW, with beam voltages ~ 47.5 kV and with minimum pulse widths of 0.1 s. The power modulations were obtained from an algorithm that can generate many different PRBS waveforms such as those shown on figures 17 and 18. Figure 17 shows an example of low frequency PRBS modulations for the LH2 and NB-R actuators applied on discharge #126960, in which the NB-L actuator was off. Figure 18 shows an example of high frequency PRBS modulations applied to the LH2 actuator on discharge #126963. They were used for the identification of the fast response model. Because of the minimum pulse width of 0.1 s, the NBI powers cannot be modulated at high frequency. In this discharge, the NB-L power was constant ($P_{NB\text{-}L} = 1.1$ MW) and the NB-R actuator was off.

A two-time-scale *ARTAEMIS* model was then identified from the experimental plasma response to the modulated

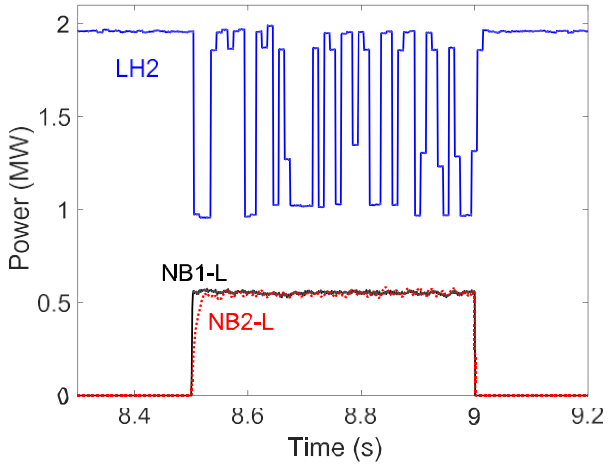


**Figure 18.** Fast LH2 power modulations (blue trace) injected from PRBS commands on discharge #126963. NB1-L (solid black trace) and NB2-L (dotted red trace) powers are also shown. The NB-R actuator is off.

inputs. Both the input and output data were filtered with a non-causal low-pass filter of order 5, and a cutoff frequency, $f_{\mathrm{filt}} = 4$ Hz, for the separation between the slow and fast data. For the experimental control of $q(\mathrm{x})$, $\beta_{\mathrm{p}}$ and $l_{\mathrm{i}}$ using the LH2, NB-L and NB-R actuators, a good model could be identified from three steady state shots with different combinations of actuator powers and seven shots in which PRBS modulations were applied either individually to the three actuators or to different combinations of them. However, no attempt was made to minimize the number of discharges that were used to identify the models, either from simulated METIS data or from EAST data because the focus in these experiments was on testing and validating the MPC controller with the best possible data-driven models. With the future development of more accurate nonlinear simulation codes, large datasets obtained from simulations can be used for system identification if the cost of running dedicated experimental discharges is unacceptable. Another significant improvement of the method in this respect could come from the development of real-time adaptive identification of the *ARTAEMIS* models.

Four significant eigenmodes were found for the slow model and two eigenmodes were found for the model governing the fast dynamics of $\beta_{\mathrm{p}}$ and $l_{\mathrm{i}}$. The largest (negative) eigenvalue of $A_{\mathrm{S}}$ corresponds to a resistive time constant $\tau_{\mathrm{S}} = 1.49$ s. Note that this time constant is significantly larger than in the model used for the SISO control experiments of section 6 ($\tau_{\mathrm{S}} = 1.05$ s) because the reference scenario for the SISO experiments had a much lower plasma temperature. The total reference power (ECRH +LH2) was 1.9 MW while here the total power (ECRH +ICRH +LH2) in the reference discharge amounts to 5.3 MW for about the same plasma density. The $A_{\mathrm{F}}$ matrix was found triangular. Its first eigenvalue corresponds to a time constant of 0.126 s and is mostly related to the dynamics of the internal inductance parameter. The second eigenvalue corresponds to a smaller time constant of 0.030 s, representative of the poloidal beta parameter dynamics. As an example of the fits obtained from the identified two-time-scale *ARTAEMIS* model, a comparison between the reduced $\psi(\mathrm{x})$ experimental data at $\mathrm{x} = 0, 0.1, \ldots 0.9$ and the simulated data obtained from the model is shown on figure 19(*a*) for the discharge #126960. The corresponding integrated fit parameter obtained from equation (18) for $\psi(\mathrm{x})$ is $f = 60\%$. Comparisons between the reduced $\delta\beta_{\mathrm{p}}$ and $\delta l_{\mathrm{i}}$ experimental data and the simulated data obtained from the *ARTAEMIS* model are also shown for the same discharge on figures 19(*b*) and (*c*), respectively. The fit parameters obtained here for $\delta\beta_{\mathrm{p}}$ and $\delta l_{\mathrm{i}}$ from the full two-time-scale model are 60% and 59%, respectively.

### 8. *Closed loop MIMO control experiments on EAST*

The *ARTAEMIS* model found above from the experimental plasma response to PRBS input data was then used to perform the first *ARTAEMIS*/MPC closed loop experiments combining the control of the $q$-profile and kinetic parameters on EAST. The objective was to test the controller ability to simultaneously track several different targets for the safety factor on the magnetic axis, $q_0$, and at half radius, $q_1 = q(\mathrm{x} = 0.5)$, as well as the poloidal beta and internal inductance parameters, $\beta_{\mathrm{p}}$ and $l_{\mathrm{i}}$. In closed loop, the power commands sent to the LH2 actuator and to the four beams, NB1-L, NB1-R, NB2-L and NB2-R, are computed in real time by the *ARTAEMIS*/MPC algorithm combined with the PWM algorithm for the beams. These commands were computed with a controller cycle time of 0.02 s (or 0.025 s in the last example), which was larger than the maximum time for executing the combined MPC/PWM algorithm in the EAST PCS. This maximum occurs at the first time step in each 0.4 s beam averaging period, when the optimal LH2, NB-L and NB-R commands are computed from the MPC algorithm and cost function in equation (53), with a one-second time horizon ($\tau_{\mathrm{H}} = 1$ s). The resulting NBI actuator commands (NB-L and NB-R) are then distributed on the four beams by the PWM algorithm, during the twenty time steps of the averaging period. This is done in such a way that the 0.1 s minimum pulse width constraint is met, and that the average NB-L and NB-R powers during the twenty steps are as close as possible to the original NB-L and NB-R power requests found by solving the QP problem at the first step of the period. During the nineteen following steps, the four beam powers are thus known and the QP solver in the MPC algorithm only has to find the optimal LH2 input sequence during the time horizon, given that the NB-L and NB-R power sequences are known[10]. The constraint matrices $L(t)$ and $M(t)$ used at each time step were constant during all the control phases, and defined as to limit the LH2 command to the interval $1.0\ \mathrm{MW} \leqslant P_{\mathrm{LH2}} \leqslant 2.2\ \mathrm{MW}$ and the two NBI actuators to $0 \leqslant P_{\mathrm{NB\text{-}L}} \leqslant 1.1$ MW and $0 \leqslant P_{\mathrm{NB\text{-}R}} \leqslant 1.1$ MW. The four beam voltages were $\sim 47.5$ kV, as in the identification experiments.

Now, since the 0.4 s NBI power averaging period is longer than the characteristic times of the fast model (0.126 s and

[10] When the prediction horizon for the computation of the LH2 request exceeds the twentieth time step of the beam-averaging period, the NB-L and NB-R patterns are still assumed to be known and to simply repeat themselves beyond the twentieth time step with the same average power.

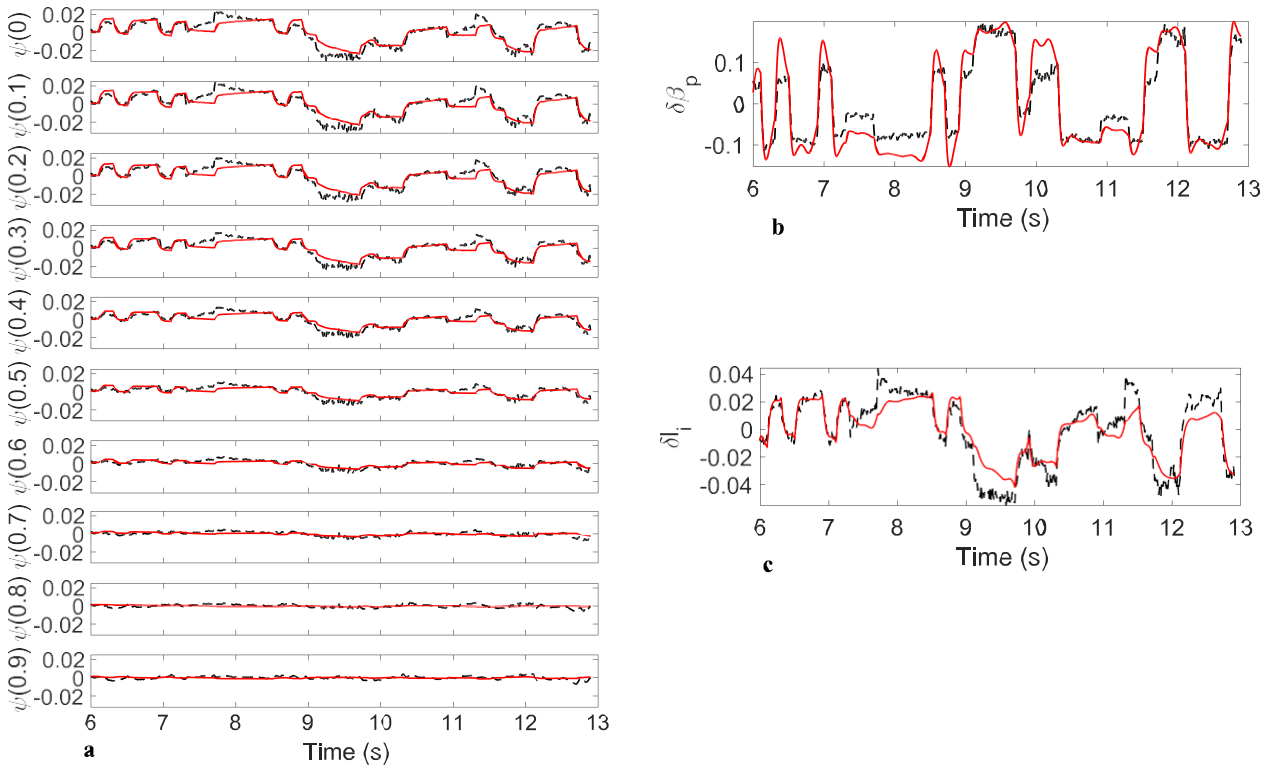


**Figure 19.** (*a*) Reduced poloidal flux response (Wb) at $x = 0, 0.1, \ldots 0.9$ to the PRBS commands in discharge #126960 (see figure 17). Comparison between the experimental data (dashed black traces) and the *ARTAEMIS*-simulated data (solid red traces). (*b*) Reduced poloidal beta response to the PRBS commands in discharge #126960 (see figure 17). Comparison between the experimental data (dashed black trace) and the *ARTAEMIS*-simulated data (solid red trace). (*c*) Reduced internal inductance parameter response to the PRBS commands in discharge #126960 (see figure 17). Comparison between the experimental data (dashed black trace) and the *ARTAEMIS*-simulated data (solid red trace).

0.030 s), efficient control of the kinetic variables on the fast time scale was not expected. Therefore, the MIMO control experiments were all performed with the reduced MPC kinetic controller defined in section 4.4, using the slow model only, with $\tau_{filt} = 0.25$ s in the $A_v$ matrix of equation (42). Thus, we set $\lambda_{fast} = 0$ because the kinetic states are not used, and we set $\lambda_F = 0.3$ in the substituted direct feed-through term of the model kinetic output equations (60) and (61), as in section 5.2. The weight matrix on the actuators in the cost function was set to $R = 0$, the target reshaping time during the horizon was set to $\tau_{target} = 0.25\ \text{s} = \tau_H/4$, and the number of input nodes for the model prediction during the horizon was $n_{nodes} = 2$. In the aim of reducing the noise in the inputs to the observer, an additional low-pass pre-filter with a characteristic time $\tau_{prefilt} = 0.4$ s was applied to the measured $U(t)$, $\bar{\iota}_m(t)$ and $\kappa_m(t)$ data entering equations (45)–(48). The controller parameters defined above have been unchanged for all the experiments described in this section.

Ideally, during closed loop experiments, the main plasma parameters and the operation scenario should be consistent with the controller internal model. In particular, the toroidal field, plasma current and density, and the requested feedforward powers, as well as their time sequence, should be the same as those used in the system identification experiments. As a matter of fact, the selection of accessible and consistent targets was made in advance using data from these open loop experiments, in order to save time during the closed loop session and test many different controller configurations in a limited time. However, the LH1 system at 2.45 GHz was not available during the closed loop session. Also, despite many attempts to improve ICRH power coupling, the maximum power injected from the two antennas did not exceed 0.9 MW, or 1.1 MW in some discharges, instead of 1.5 MW requested. Finally, the ECRH power was sometimes reduced to 0.5 MW delivered by only one gyrotron instead of 1.8 MW requested from three gyrotrons. This lack of power in most discharges had important consequences on the results because the controller internal model and the control targets to be tracked were all obtained with the LH1 prelude during current ramp-up and with the maximum available ICRH and ECRH powers. The limitation of the injected powers in the feedforward mode with respect to the powers foreseen for the selection of the control targets, often led to the saturation of one or more actuators, which, in many cases, prevented the targets from being hit. In this respect, the power limits of the LH2 actuator were another source of control inaccuracy. Although, the LH2 power command at the output of the controller was ranging from 1 MW to 2.2 MW, the output power from the klystrons was restricted

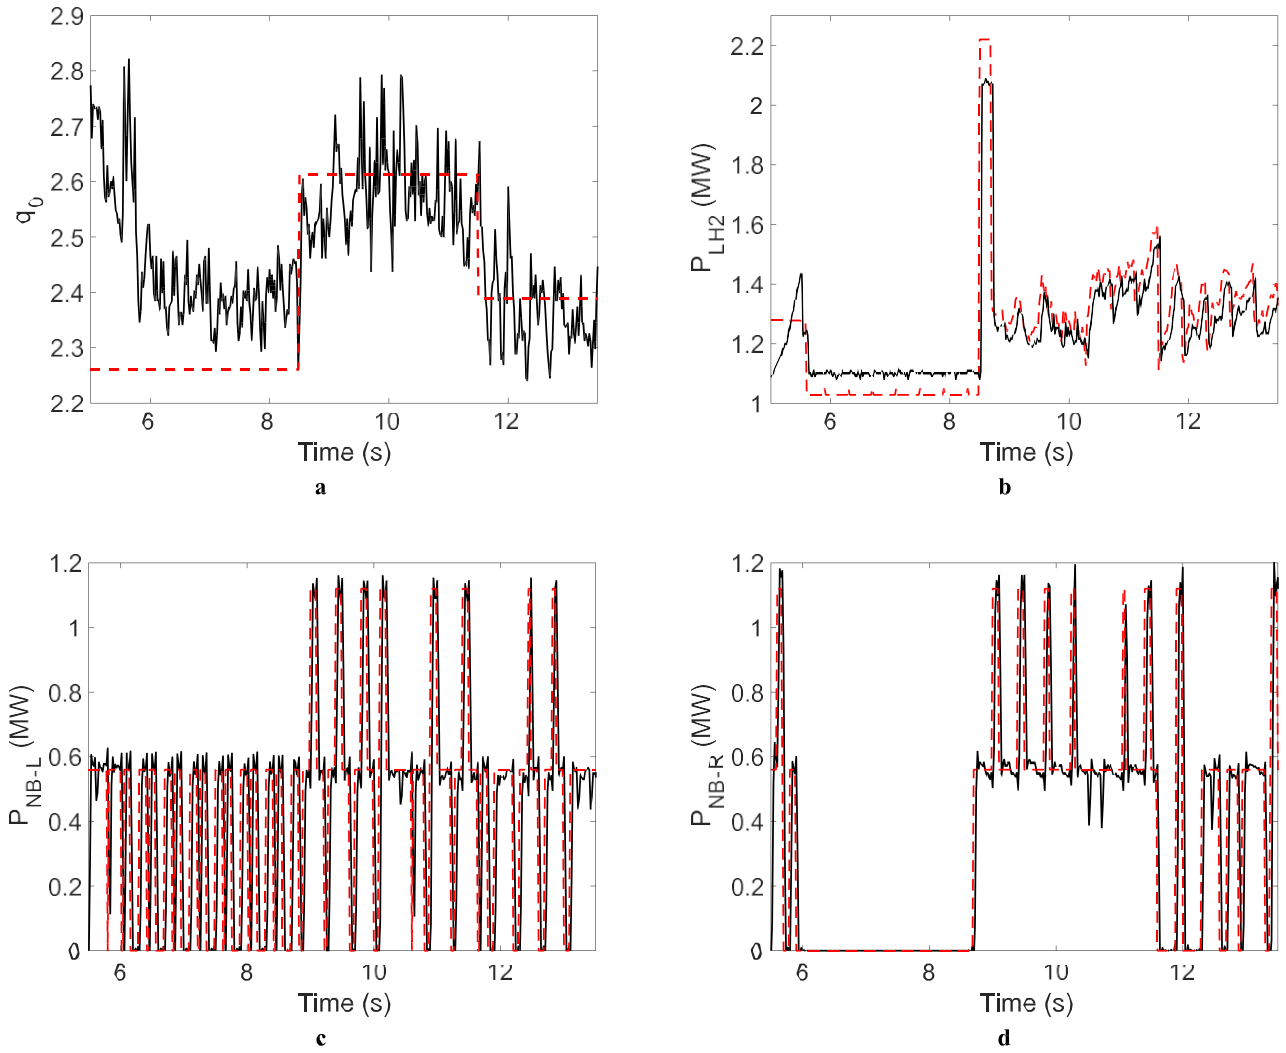


**Figure 20.** (*a*) Control of the safety factor on the magnetic axis with LH2, NB-L and NB-R: $q_0(t)$ from real-time EFIT magnetic equilibrium reconstruction (solid black) and $q_{0,\mathrm{target}}(t)$ (dashed red) in EAST discharge #128569. Control starts at $t = 5.5$ s. The ECRH and ICRH feedforward powers are limited to 0.5 and 0.9 MW, respectively, instead of 1.8 and 1.5 MW requested. (*b*) Coupled LH2 power (solid) and PCS command (red dashed) in EAST discharge #128569. Control starts at $t = 5.5$ s. Actuator saturation at 1.1 MW between 5.6 s and 8.5 s. (*c*) Injected NB-L power (solid) and PCS command (red dashed) in EAST discharge #128569. Control starts at $t = 5.5$ s. (*d*) Injected NB-R power (solid) and PCS command (red dashed) in EAST discharge #128569. Control starts at $t = 5.5$ s. Actuator saturation at 0 MW between 6 s and 8.7 s.

to the interval 1.1 MW to 2 MW in order to ensure optimum real-time operation.

Five controller configurations have been tested during the time allowed for the experimental session, all using the three actuators, NB-L, NB-R and LH2, and typical results are described below for each configuration.

*8.2.1. Discrete $q_0$ control.* The first configuration aimed at the control of the safety factor on the magnetic axis. Figure 20(*a*) shows the evolution of $q_0$ in the discharge #128569 where three successive targets were tracked: $q_{0,\mathrm{target}} = 2.26$ for 5.5 s $\leqslant t \leqslant$ 8.48 s, $q_{0,\mathrm{target}} = 2.61$ for 8.5 s $\leqslant t \leqslant$ 11.48 s and $q_{0,\mathrm{target}} = 2.39$ for 11.5 s $\leqslant t \leqslant$ 13.5 s. Figures 20(*b*)–(*d*) show the powers requested by the controller and the powers really injected. In this configuration, the $q$-profile weight function, $\mu(x)$, was replaced in the cost function (53) by the Dirac distribution, $\delta(x)$, and we set $\lambda_{\mathrm{kin}} = 0$ since there is no kinetic control. With three actuators and only one controlled variable, the $1 \times 3$ $K_S$ matrix has only one SVD component, so $n_{\mathrm{svd}} = 1$. In the example shown, the ICRH and ECRH feedforward powers were only 0.9 MW (instead of 1.5 MW) and 0.5 MW (instead of 1.8 MW), respectively. As a result, the lowest target, $q_{0,\mathrm{target}} = 2.26$, could not be reached and the LH2 actuator was saturated at its minimum value of 1.1 MW (while the controller assumed that a minimum of 1 MW was available and the PCS command did request 1 MW). The NB-R actuator was also saturated at its minimum (power off).

Despite the minimum pulse width constraint in the on/off operation of the NBI actuators, the next two targets were successfully reached in a time that is of the order of the characteristic resistive time of the model, identified as $\tau_S = 1.49$ s for this range of plasma parameters and this operation scenario. A similar result was obtained for the SISO control of $q_0$ with LH2, in section 6.2 and figures 13(*a*) and 20(*a*) are indeed quite comparable, with about the same experimental noise level. Interestingly, a clear signature of the MPC/PWM algorithm can be seen on figure 20(*b*), which exhibits a 0.4 s periodic pattern on the LH2 power. The amplitude of these power oscillations could in principle be reduced, and their frequency increased, by reducing the minimum pulse width of the beams and the beam averaging time. On the technical side, it must be noted from figures 20(*c*) and (*d*) that the power modulation PCS commands generated by the combined MPC/PWM algorithms were perfectly executed by the NBI system.

*8.2.2 Discrete control of $q_0$ and $q(\mathbf{x} = 0.5)$.* The second configuration that was tested corresponds to the simultaneous control of the safety factor at two different normalized radii, $\mathbf{x} = 0$ and $\mathbf{x} = 0.5$. In this configuration, the $q$-profile weight function, $\mu(\mathrm{x})$, was replaced in the cost function of equation (53) by the sum of Dirac distributions, $\delta(\mathrm{x})+\delta(\mathrm{x}\text{-}0.5)$ and $\lambda_{\mathrm{kin}} = 0$ since there is no kinetic control. In equation (41), the model steady state response, i.e. the $K_S$ matrix, was limited to two $q$-profile outputs at $\mathbf{x} = 0$ and $\mathbf{x} = 0.5$ and was therefore a $2 \times 3$ matrix. The NB-L, NB-R and LH2 actuators were normalized to 1.12 MW, 1.12 MW and 1.2 MW, respectively, and with these normalizations, the singular values of the $K_S$ matrix were equal to 0.0927 and 0.0037. The strong drop between the first and second singular values indicates that the combination of actuators corresponding to the second singular vector of $K_S$ has little effect on the system dynamics. It also means that an offset-free control objective would have resulted in large amplitudes of the actuator requests. In other words, the values of the safety factor at $\mathbf{x} = 0$ and $\mathbf{x} = 0.5$ are closely related and it is difficult to decouple them. To reduce the occurrence of actuator saturation, $K_S$ was therefore truncated to the first element of its singular value expansion for the computation of the optimal control action ($n_{\mathrm{svd}} = 1$). The controller then only achieved a least-square minimization of the offsets between the targets and the measured values of $q_0$ and $q_1 = q(\mathrm{x} = 0.5)$. The evolution of $q_0$ and $q_1$ is shown in figure 21(*a*) for the discharge #128550 in which two successive sets of targets were tracked. The targets were $[q_{0,\mathrm{target}}\ q_{1,\mathrm{target}}] = [2.39\ 2.86]$ for 5.5 s $\leqslant t \leqslant$ 8.48 s, and $[q_{0,\mathrm{target}}\ q_{1,\mathrm{target}}] = [2.58\ 3.0]$ for 8.5 s $\leqslant t \leqslant$ 11.5 s.

Figures 21(*b*)–(*d*) show the powers requested by the controller and the powers really injected. It is important to note that, in this discharge, the line-averaged plasma density was limited to $n_{\mathrm{el}} = 4 \times 10^{19}$ m$^{-3}$ instead of the reference $4.5 \times 10^{19}$ m$^{-3}$, and the feedforward ICRH power was limited to 0.7 MW instead of the requested 1.5 MW that was used in the shots from which the different $q(\mathrm{x})$ targets were chosen. This is why the first set of targets could not be reached between $t = 5.5$ s and $t = 8.5$ s. In this time interval, both the NB-R and LH2 actuators were saturated at their lowest possible values, 0 MW and 1.1 MW, respectively, so that $q_0$ and $q_1$ could not decrease further without the full 1.5 MW ICRH power and the requested $4.5 \times 10^{19}$ m$^{-3}$ plasma density. Here also, the LH2 power was technically bounded above 1.1 MW instead of 1 MW as specified in the controller (see figure 21(*b*)).

The second set of targets was successfully reached, consistently with the resistive diffusion time of the model ($\tau_S = 1.49$ s) and with the characteristic oscillations that are due to the on/off operation of the neutral beams and to the MPC/PWM algorithm.

*8.2.3 Simultaneous control of $q_0$ and $\beta_p$.* In the discharge #128557, the simultaneous control of the central safety factor and of the poloidal $\beta$ parameter was tested. As in section 8.2.1, the $q$-profile weight function, $\mu(\mathrm{x})$, was replaced in the controller cost function (53) by the Dirac distribution, $\delta(\mathrm{x})$, and the weight of $\beta_p$ control, $\lambda_{\mathrm{kin}}$, was set to 0.3. The actuator normalizations were the same as in sections 8.2.1 and 8.2.2 and the singular values of the $2 \times 3$ $K_S$ matrix were equal to 0.2898 and 0.0215. Here, the ratio between the two singular values was more favorable than in the previous configuration so the full $K_S$ matrix was retained ($n_{\mathrm{svd}} = 2$) to possibly realize offset-free control of $q_0$ and $\beta_p$. Two successive sets of targets were tracked: $[q_{0,\mathrm{target}}\ \beta_{p,\mathrm{target}}] = [2.26\ 1.71]$ for 5.5 s $\leqslant t \leqslant$ 9.48 s, and $[q_{0,\mathrm{target}}\ \beta_{p,\mathrm{target}}] = [2.61\ 2.01]$ for 9.5 s $\leqslant t \leqslant$ 13.5 s. Unfortunately, in this discharge, plasma density control was poor and the line-averaged density reached $4.8 \times 10^{19}$ m$^{-3}$ between $t = 6$ s and $t = 9.5$ s and even $5 \times 10^{19}$ m$^{-3}$ for $t > 9.7$ s. In addition, the feedforward ICRH power was limited to 0.9 MW instead of the requested 1.5 MW that was used in the shots from which the different $q_0$ and $\beta_p$ targets were chosen. The evolution of $q_0$ and $\beta_p$ is shown on figure 22(*a*). The powers requested by the controller and the powers really injected are shown in figures 22(*b*) and (*c*). The $q_0$ targets could not be reached in these high-density/low-power conditions. Nevertheless, $\beta_p$ control was excellent during the first phase with $\beta_{p,\mathrm{target}} = 1.71$. The $\beta_p = 2.01$ target was also nearly hit despite the saturation at 1.1 MW of the NB-R actuator and also, for $t > 12.5$ s, of NB-L.

*8.2.4 Simultaneous control of $q_0$ and $l_i$.* The fourth configuration tested in this experimental session addressed the simultaneous control of the central safety factor and the internal inductance parameter, $l_i$. The weights in the controller cost function were the Dirac distribution, $\delta(\mathrm{x})$, for the $q$-profile, and $\lambda_{\mathrm{kin}} = 0.3$ for $l_i$. Here, the NB-L, NB-R and LH2 actuators were normalized to 0.5 MW, 0.5 MW and 2.2 MW, respectively, and the singular values of the $2 \times 3$ $K_S$ matrix were equal to 0.1302 and 0.0053. The full $K_S$ matrix was retained ($n_{\mathrm{svd}} = 2$) to possibly realize offset-free control of $q_0$ and $l_i$. In the discharge #128560, the tracked targets were $q_{0,\mathrm{target}} = 2.79$ and $l_{i,\mathrm{target}} = 0.85$ for 5.5 s $\leqslant t \leqslant$ 9.5 s. The feedforward ICRH power was limited to 1.1 MW instead of the requested 1.5 MW that was used in the shots from which the targets were

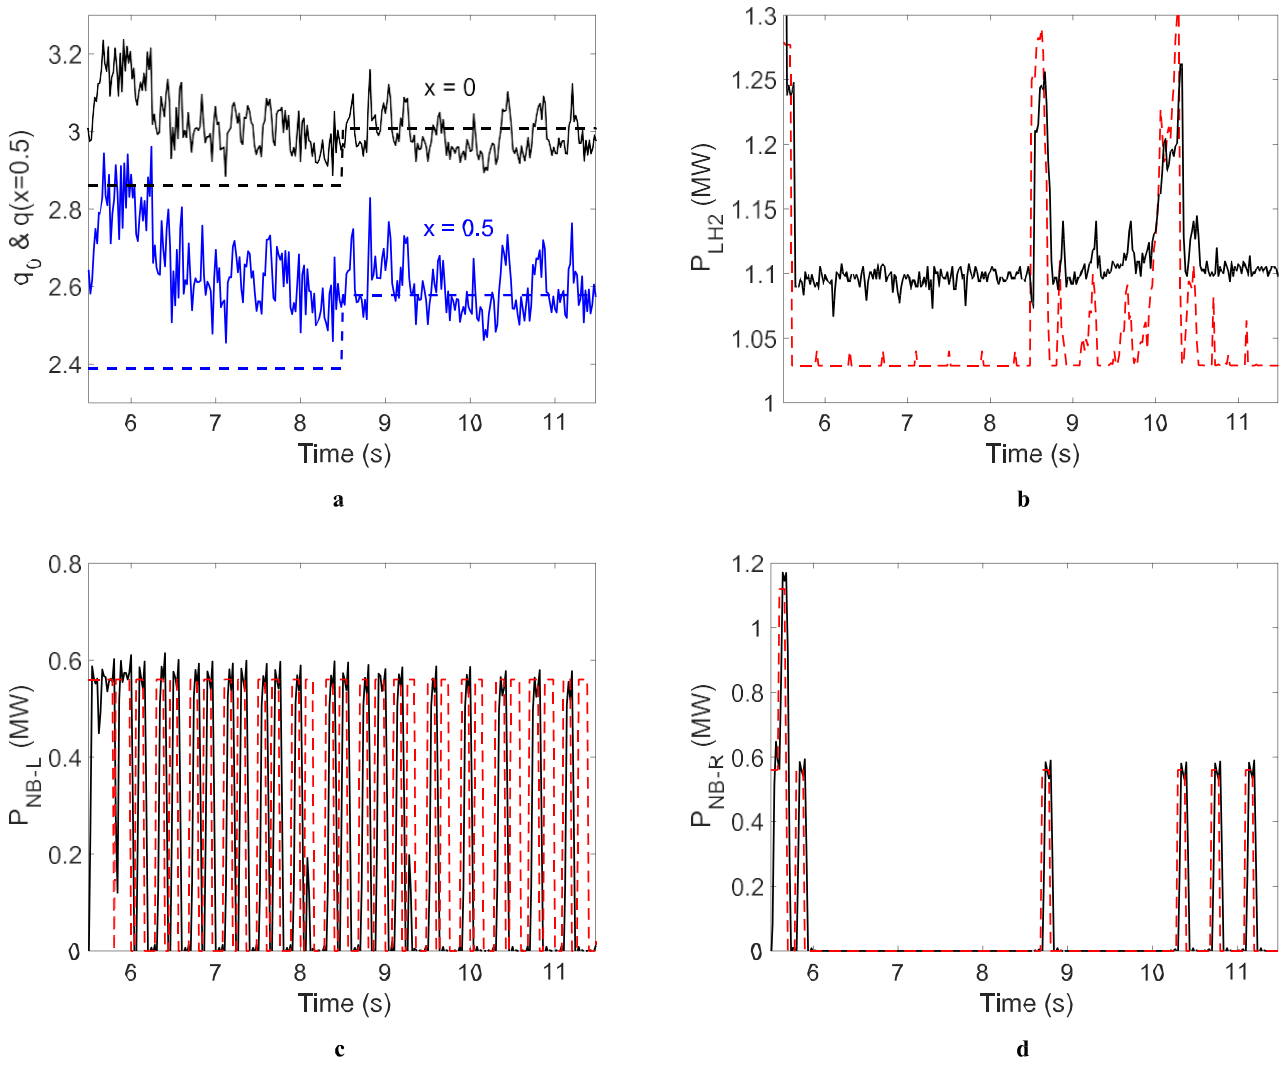


**Figure 21.** (*a*) Simultaneous control of $q_0$ and $q_1 = q(\mathrm{x} = 0.5)$ with the LH2, NB-L and NB-R actuators: $q_0$ (solid black trace), $q_{0,\mathrm{target}}$ (dashed black), $q_1$ (solid blue) and $q_{1,\mathrm{target}}$ (dashed blue) vs time in EAST discharge #128550. Control starts at $t = 5.5$ s. The line-averaged plasma density is limited to $4 \times 10^{19}$ m$^{-3}$ instead of the requested $4.5 \times 10^{19}$ m$^{-3}$. The feedforward ICRH power is limited to 0.7 MW instead of the requested 1.5 MW. (*b*) Coupled LH2 power (solid) and PCS command (red dashed) in EAST discharge #128550. Control starts at $t = 5.5$ s. Actuator saturation at 1.1 MW between 5.6 s and 8.5 s. (*c*) Injected NB-L power (solid) and PCS command (red dashed) in EAST discharge #128550. Control starts at $t = 5.5$ s. (*d*) Injected NB-R power (solid) and PCS command (red dashed) in EAST discharge #128550. Control starts at $t = 5.5$ s. Actuator saturation at 0 MW between 6 s and 8.7 s.

chosen. Nevertheless, there is no actuator saturation except between $t = 5.9$ s and $t = 6.3$ s when LH2 is limited at 2 MW instead of 2.2 MW requested. The evolution of $q_0$ and $l_i$ is shown on figure 23(*a*). The powers requested by the controller and the powers really injected are shown in figures 23(*b*) and (*c*). The values of $q_0$ and $l_i$ are perfectly regulated around the targets.

*Simultaneous control of $q_0$, $\beta_p$ and $l_i$.* The last controller configuration was tested in the discharge #128566 with the same three actuators and with three simultaneously controlled variables: the central safety factor, the poloidal $\beta$ parameter and the internal inductance parameter. As in the previous case, the $q$-profile weight function, $\mu(\mathrm{x})$, was replaced by the Dirac distribution, $\delta(\mathrm{x})$. Here, the kinetic vector $\kappa(t)$ contains two elements, $\delta\beta_p(t)$ and $\delta l_i(t)$, and the diagonal elements of the $2 \times 2$ kinetic weight matrix, $\lambda_{\mathrm{kin}}$, were both set to 0.2, thus giving equal weights on the control of $\beta_p$ and $l_i$. Using the same actuator normalizations as in the previous case, the singular values of the $3 \times 3$ $K_S$ matrix were 0.1088, 0.0226 and 0.0034. Again, the full $K_S$ matrix was retained ($n_{\mathrm{svd}} = 3$) to possibly realize simultaneous offset-free control of three plasma parameters, $q_0$, $\beta_p$ and $l_i$. The tracking of two sets of targets is illustrated on figure 24. These targets are $[q_{0,\mathrm{target}}\ \beta_{p,\mathrm{target}}\ l_{i,\mathrm{target}}] = [2.47\ 1.85\ 0.92]$ for $5.5\ \mathrm{s} \leqslant t \leqslant 8.48$ s and $[q_{0,\mathrm{target}}\ \beta_{p,\mathrm{target}}\ l_{i,\mathrm{target}}] = [2.31\ 1.75\ 0.96]$ for $8.5\ \mathrm{s} \leqslant t \leqslant 11.5$ s. In this controller configuration, it was found that the CPU time required for reconstructing the plasma magnetic equilibrium and solving the QP problem could periodically exceed 0.02 s, every 0.4 s, i.e. at times when the three actuators are unknown. The controller cycle time was therefore raised to

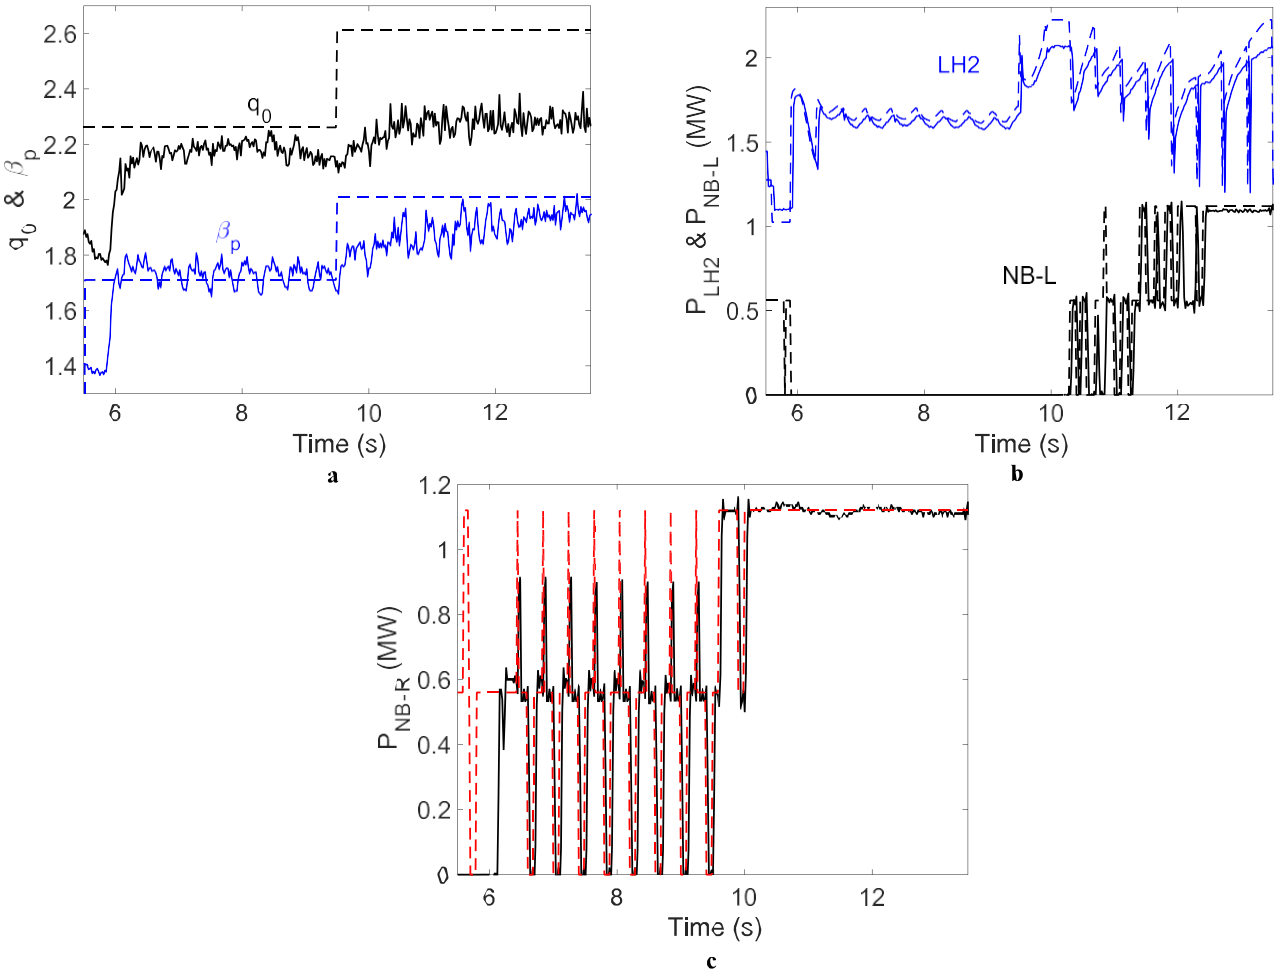


**Figure 22.** (*a*) Simultaneous control of $q_0$ and $\beta_p$ with LH2, NB-L and NB-R: $q_0$ (solid black trace), $q_{0,target}$ (dashed black), $\beta_p$ (solid blue) and $\beta_{p,target}$ (dashed blue) vs time in EAST discharge #128557. Control starts at $t = 5.5$ s. The line-averaged plasma density is equal to $4.8 \times 10^{19}$ m$^{-3}$ between $t = 6$ s and $t = 9.5$ s, and up to $5 \times 10^{19}$ m$^{-3}$ for $t > 9.7$ s, instead of the requested $4.5 \times 10^{19}$ m$^{-3}$. The feedforward ICRH power is limited to 0.7 MW instead of the requested 1.5 MW. (*b*) Coupled LH2 power (solid blue), LH2 PCS command (dashed blue), injected NB-L power (solid black) and NB-L PCS command (dashed black) in EAST discharge #128557. Control starts at $t = 5.5$ s. NB-L actuator saturation between 6 s and 10.3 s, and for $t > 12.5$ s. (*c*) Injected NB-R power (solid black) and NB-R PCS command (dashed red) in EAST discharge #128557. Control starts at $t = 5.5$ s. Actuator saturation at 1.1 MW for $t > 10.1$ s.

0.025 s for these experiments. The evolution of $q_0$, $\beta_p$ and $l_i$ in the discharge #128566 is shown in figure 24(*a*) and the powers requested by the controller and really injected are shown in figures 24(*b*) and (*c*). Here again, the feedforward ICRH power was limited to 1.1 MW instead of the requested 1.5 MW that was used in the shots from which the targets were chosen. For these targets, the NB-R actuator was saturated (power off). The tracking of $\beta_p$ and $l_i$ was relatively successful despite the oscillations generated by the MPC/PWM commands, but the central safety factor was always larger than its target value. This is partly due to the lack of ICRH power, and possibly to the unoptimized relative weights of the various controlled variables in the controller cost function. Figure 24(*a*) suggests that an extension of the control phases to durations larger than three seconds could have allowed the three targets to be reached, but there was not enough time in this experimental session to repeat the tests and optimize the controller parameters in each control configuration.

## 9. Summary and conclusion

A real-time control method that combines the efficiency of MPC with the use of singular perturbation theory has been developed for the first time. This original method is particularly suitable for the control of complex systems with multiple time scales, such as tokamak plasmas. It allows plasma kinetic controllers based on extremely simple data-driven models to be synthesized. The so-called two-time-scale models govern the evolution of the system on the resistive and kinetic time scales, and they can be identified using the classical prediction-error method. They are then augmented to include new disturbance states and outputs so that a real-time estimation of the mismatch between measured and predicted outputs can be made through a state observer. Offset-free control can thus be obtained when the number of actuators is equal to the number of controlled variables or larger, despite disturbances acting on the system and on the measurements.

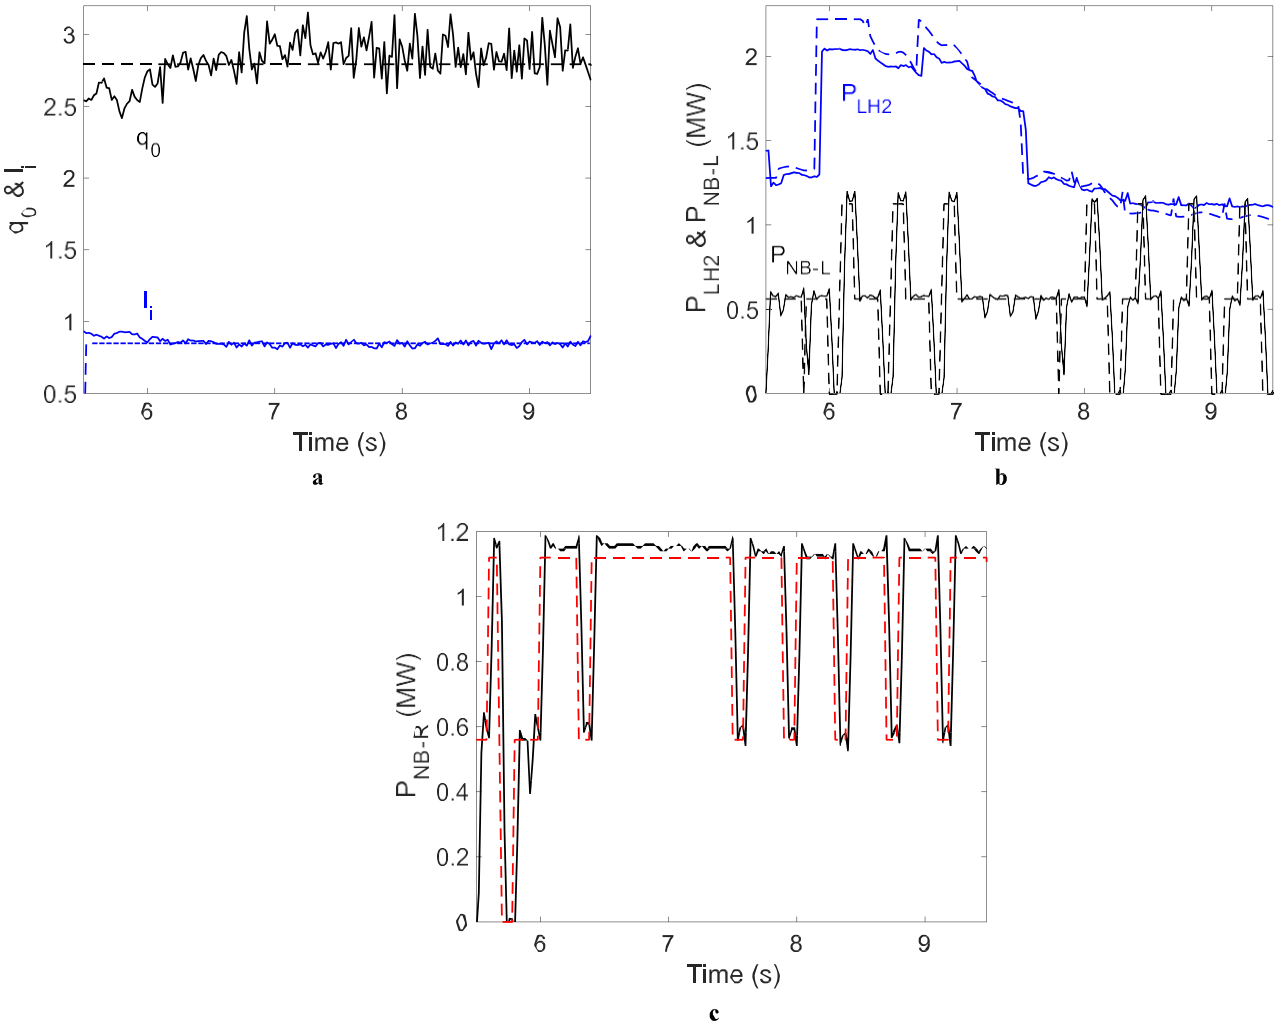


**Figure 23.** (*a*) Simultaneous control of $q_0$ and $l_i$ with LH2, NB-L and NB-R: $q_0$ (solid black trace), $q_{0,\text{target}}$ (dashed black), $l_i$ (solid blue) and $l_{i,\text{target}}$ (dashed blue) vs time in EAST discharge #128560. Control starts at $t = 5.5$ s. The feedforward ICRH power is limited to 1.1 MW instead of the requested 1.5 MW. (*b*) Coupled LH2 power (solid blue), LH2 PCS command (dashed blue), injected NB-L power (solid black) and NB-L PCS command (dashed black) in EAST discharge #128560. Control starts at $t = 5.5$ s. LH2 actuator saturation between 5.9 s and 6.3 s. (*c*) Injected NB-R power (solid black) and NB-R PCS command (dashed red) in EAST discharge #128560. Control starts at $t = 5.5$ s.

The efficiency of the method has been demonstrated through extensive nonlinear closed loop METIS simulations applied to high-$\beta_p$ operation scenarios in the EAST tokamak, with different controlled variables and different sets of actuators. In the first set of simulations, control of the plasma safety factor profile, $q(x)$, and of the poloidal $\beta$ parameter was tested using LHCD and ICRH in a scenario with $B_T = 2.5$ T, $I_p = 0.42$ MA and $n_{el} \approx 2.7 \times 10^{19}$ m$^{-3}$. In a second set of METIS simulations, simultaneous control of $q(x)$ at two radii and of $\beta_p$ was also tested in a scenario with a higher plasma density, $n_{el} = 4.4 \times 10^{19}$ m$^{-3}$, and using up to four actuators: LHCD, ECRH, and also two sets of co-current NBI injectors at different angles driven in a PWM mode. Adequate two-time-scale models were identified in both scenarios from simulation data obtained with randomly modulated actuator waveforms and pseudo-random binary sequences.

In the low-density scenario, three typical control examples were reported. The first tests were dedicated to the control of the $q$-profile using only the 4.6 GHz LHCD actuator. Offset-free tracking of various $q_0$ targets was achieved in a few resistive times in the case of a discrete control of $q_0$. Distributed $q(x)$ control for $0 \leqslant x \leqslant 0.5$ was also shown with targets including monotonic, weak shear and negative shear profiles. The controller behaved very similarly as in the discrete $q_0$ control case, but with small steady state offsets on axis compensated by a better tracking of the target profile between $x = 0.3$ and $x = 0.5$. The second tests addressed the simultaneous control of $q(x)$ and $\beta_p$ with LHCD and ICRH, using the MPC controller reduced to the resistive time scale to track different combinations of $q(x)$ and $\beta_p$ targets. Steady state offsets cannot be avoided with only two actuators for distributed $q(x)$ control and simultaneous control of $\beta_p$. The $q$-profile offsets are mostly apparent near the magnetic axis and they disappear when only $q_0$ and $\beta_p$ are controlled. In the third tests, the full two-time-scale MPC controller was used with the same actuators to simultaneously control $q(x)$ and $\beta_p$. The benefit of

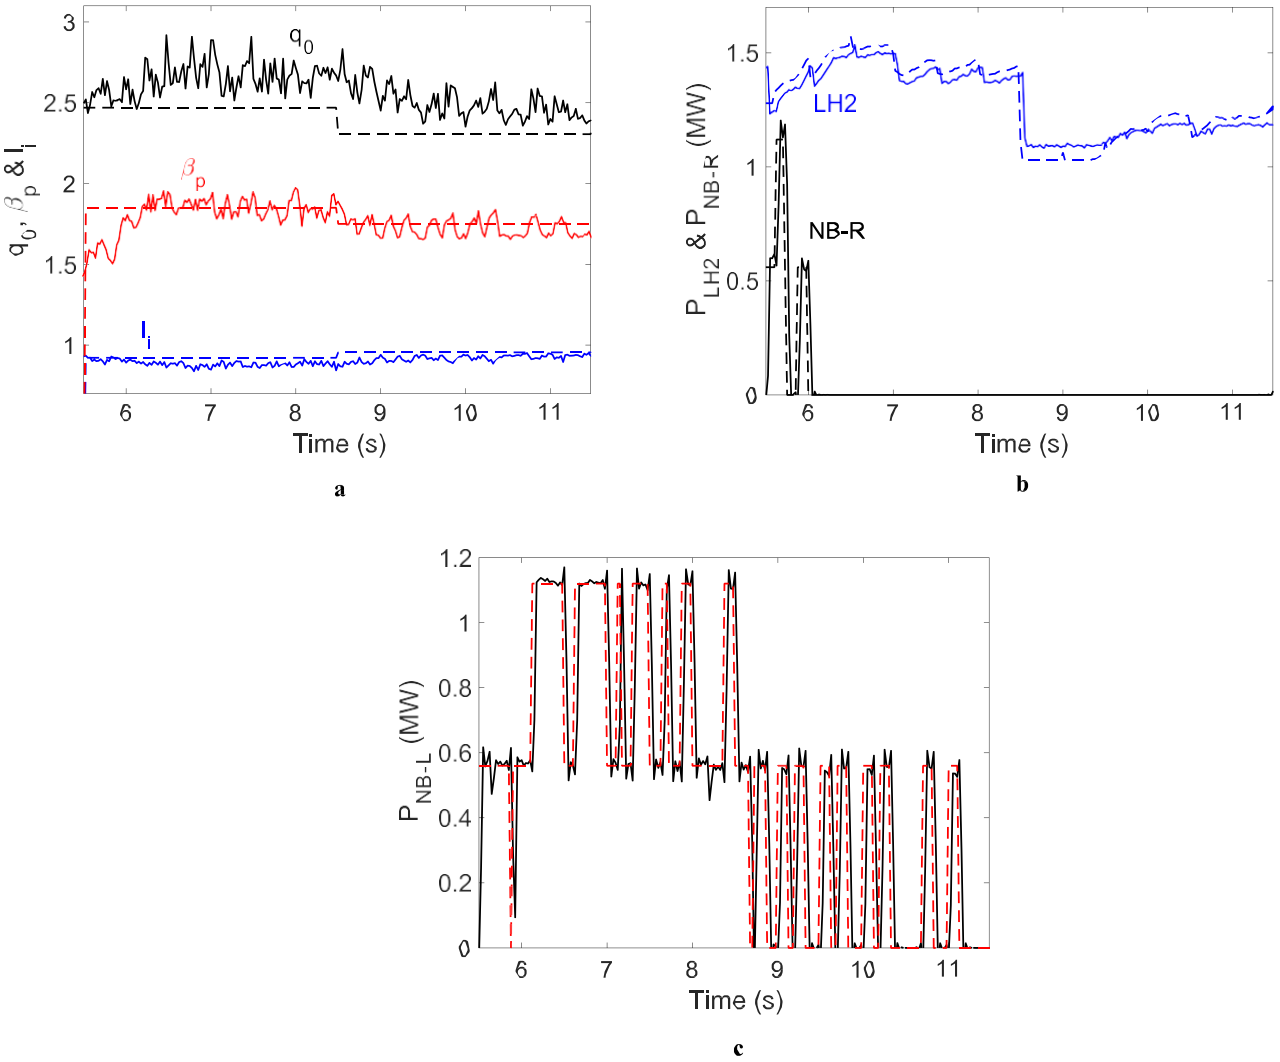


**Figure 24.** (*a*) Simultaneous control of $q_0$, $\beta_p$ and $l_i$ with LH2, NB-L and NB-R in EAST discharge #128566: $q_0$ (solid black trace), $q_{0,target}$ (dashed black), $\beta_p$ (solid red) and $\beta_{p,target}$ (dashed red), $l_i$ (solid blue) and $l_{i,target}$ (dashed blue) vs time. Control starts at $t = 5.5$ s. The feedforward ICRH power is limited to 1.1 MW instead of the requested 1.5 MW. (*b*) Coupled LH2 power (solid blue), LH2 PCS command (dashed blue), injected NB-R power (solid black) and NB-R PCS command (dashed black) in EAST discharge #128566. Control starts at $t = 5.5$ s. LH2 actuator saturation between 8.6 s and 9.3 s and NB-R saturation between 6.1 s and 11.5 s. (*c*) Injected NB-L power (solid black) and NB-L PCS command (dashed red) in EAST discharge #128566. Control starts at $t = 5.5$ s.

using the full controller was demonstrated. The $q$-profile targets were all tracked in about 2.5 s and the various $\beta_p$ targets were reached within about 0.2 s with the full controller and 0.4–0.5 s with the reduced controller.

In the high-density scenario, simultaneous control of $q(x)$ at two radii and of $\beta_p$ was tested in METIS simulations using LHCD, ECRH and NBI actuators. Technically, NBI injectors can either be 'on' or 'off' so they cannot deliver a continuously varying power. They must therefore be operated in a PWM mode so that the controller power commands computed at a given time are applied in an average sense during a number of future time steps corresponding to a fixed power-averaging period. Such a PWM algorithm was therefore included in the MPC controller and validated in the closed loop METIS simulations. In the first simulation tests, the MPC/PWM controller was used to track four discrete target pairs for $q_0$ and $q_1 = q(x = 0.4)$, corresponding to monotonic, weak shear and negative shear $q$-profiles. The tracking of both $q_0$ and $q_1$ was excellent, in a time that is consistent with the resistive time, and with no steady state offset. In the second simulation tests, four different sets of $q_0$, $q_1$ and $\beta_p$ targets were tracked. The PWM algorithm introduces a 0.4 s delay corresponding to the NBI power-averaging period, and generates low frequency oscillations of the controlled variables around their targets. The oscillations are small and acceptable for the control of $q(x)$, but their amplitude is large for $\beta_p$, showing reduced kinetic control accuracy. Increasing the maximum frequency at which NBI can be switched on and off would be necessary for better kinetic control.

The MPC control technique described in the present paper was also tested experimentally on EAST. Two control configurations were considered. In the first scenario, LHCD at 4.6 GHz was the only available actuator and the SISO version of the MPC controller was validated. The toroidal field

was $B_T = 2.5$ T, the plasma current was $I_p = 350$ kA and the plasma density was $n_{el} \sim 4.2 \times 10^{19}$ m$^{-3}$. An experimental *ARTAEMIS* model was identified using data from real discharges with chirping frequency and PRBS modulations of the LHCD actuator. The first control experiments aimed at tracking piecewise-constant waveforms for $q_0$. Tracking of the $q_0$ targets was successfully performed without steady state offset, and in about one second, i.e. one resistive diffusion time. In other discharges, $\beta_p$ control was tested showing excellent tracking of a piecewise-linear $\beta_p$ waveform despite significant noise in the measurements.

More recent experiments were dedicated to MIMO control. The operation scenario was based on a negative toroidal field, $B_T = -2.5$ T, $I_p = 0.4$ MA and a relatively high plasma density, $n_{el} = 4.5 \times 10^{19}$ m$^{-3}$, with constant ECRH and ICRH powers of 1.8 MW and 1.5 MW, respectively, to sustain a good H-mode discharge. A maximum of three independent actuators could be operated from real-time PCS commands and were used for plasma kinetic control: LHCD at 4.6 GHz (LH2) with a minimum power of 1 MW, and two NBI actuators, NB-L and NB-R, with different injection angles. Open loop power modulation experiments with LH2, NB-L and NB-R PRBS inputs were dedicated to the identification of an experimental *ARTAEMIS* model for this scenario, including the constant ECRH and ICRH powers. The minimum allowed pulse width from the EAST NBI system was 0.1 s, instead of 0.04 s assumed for the simulations, but the NBI power-averaging time was 0.4 s, as in the simulations. The objective of the experiments was to simultaneously track different targets for the safety factor on the magnetic axis, $q_0$, and at half radius, $q_1 = q(x = 0.5)$, and for $\beta_p$ and $l_i$. Since the neutral beam power-averaging period was longer than the characteristic times of the fast model, the control experiments were performed with the reduced kinetic controller, i.e. using the slow model only, and with a cycle time of 0.02 s or at most 0.025 s in one case. Five controller configurations have been tested during the time allowed for the experimental session, all with the three actuators, NB-L, NB-R and LH2, with up to three controlled variables chosen among $q_0$, $q_1$, $\beta_p$ and $l_i$. In all the tests, the power modulation commands generated by the combined MPC/PWM algorithm were perfectly executed by the NBI systems. The chosen cycle time was large enough to reconstruct the plasma magnetic equilibrium from the real-time measurements and solve the QP optimization problem at each time step. It was adequate for tracking different steady state targets in a time that is consistent with the plasma resistive time scale and with the required averaging period of 0.4 s for the NBI/PWM actuators. This was the goal of the experiments. It would not be adequate if the discharges had shown rapid transition phases, which are not considered in this paper, and which would require a specific control strategy and hierarchy, with much faster actuation. An example of a faster control using the full two-time-scale model and fast actuators only, but the same 0.02 s cycle time necessary for magnetic reconstruction and QP optimization on EAST, was shown in section 5.3.

The main difficulty encountered during the experiments was the inaccessibility (without hitting actuator saturation limits) of control targets that were selected in advance, in situations where the baseline feedforward powers and, sometimes, the plasma density were far from their desired reference values. Thus, the lack of ECRH and/or ICRH power in most discharges had important consequences on the results because the controller internal model as well as all the chosen control targets had been obtained with reference powers that were available in a previous experimental session. This led to the saturation of one or several actuators, which prevented some selected targets to be reached. Nevertheless, in cases that were free from actuator saturation, $q_0$ and $q_1$ targets were successfully reached, in a time that is consistent with the resistive diffusion time and with the characteristic oscillations that are due to the PWM operation of the neutral beams. During the simultaneous control of $q_0$ and $\beta_p$, the plasma density was too high and the feedforward ICRH power was limited to 0.9 MW instead of the requested 1.5 MW. In these high-density/low-power conditions, the $q_0$ targets were not accessible but $\beta_p$ control was successful. There was almost no actuator saturation during the simultaneous control of $q_0$ and $l_i$, so their values were also successfully regulated. In the last example with three controlled variables, $q_0$, $\beta_p$ and $l_i$, the tracking of $\beta_p$ and $l_i$ was satisfactory but $q_0$ was always too large due to the lack of ICRH power and to NB-R saturation.

As a general conclusion, both nonlinear simulations and real tokamak experiments have demonstrated the relevance of combining MPC, data-driven models and singular perturbation methods for plasma kinetic control. The performance of the controllers is validated a posteriori, based on the successful tracking of the controlled parameters towards their targets, in a time that is consistent with the plasma resistive time scale. In previous designs based on near-optimal control and the use of two-time-scale models, the convergence to the targets was much too slow and often characterized by large oscillations [35]. In both approaches, there were no particular quantitative requirements that the controllers would have to satisfy for validation. Some experimental tests presented in this paper were not fully conclusive because the required baseline power from H&CD systems was not available and the actuators could not compensate for the lack of power from other systems. It would have been interesting to select control targets that were consistent with the available baseline power, on the day of the experiment, or to repeat some tests with different controller parameters, but this was not possible in the allocated time. It would require longer and several experimental sessions.

Finally, as a perspective, this control method could be improved by implementing a real-time adaptive identification of the two-time-scale models. With adequate actuators, it could also be easily extended to the simultaneous control of the $q$-profile and other kinetic parameters or profiles such as $\beta_N$, the plasma rotation, the ion or electron temperature, and/or the fusion reaction rate in burning plasmas. For such further developments, the advantage of using empirical models

obtained by applying singular perturbation methods to system identification is that the resulting two-time-scale state space models are extremely simple and this reduces the computational burden in the control algorithm. From the origin, with near-optimal control, the *ARTAEMIS* framework was based on the separation of the resistive and kinetic time scales [30]. This is indeed a very relevant approximation, particularly for large tokamak devices and fusion grade plasmas. Although this approximation was not generally needed for profile control investigations on several middle-size tokamaks, it provides interesting simplifications for the development of MIMO controllers for large fusion plasmas, and it may be of interest for future research on this subject. Extrapolation of MPC kinetic profile control to burning plasmas may require approximations of this kind. Among other possible developments, the recently proposed hybrid MPC scheme [45], which combines constraints that apply to continuous actuators with the discrete Boolean nature of NBI constraints is an interesting extension of conventional MPC. It is more demanding in terms of computational time but it could be worth trying with the *ARTAEMIS* models and framework.

## Acknowledgment

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under Grant Agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. This work was also supported in part by the China Scholarship Council (CSC) and the Excellence Program of Hefei Science Center CAS 2019HSC-UE013.

## ORCID iDs

D. Moreau https://orcid.org/0000-0001-5754-4103
Q. Yuan https://orcid.org/0000-0003-4292-1302
Y. Huang https://orcid.org/0000-0002-7247-7786
S. Ding https://orcid.org/0000-0002-1930-0439
M. Li https://orcid.org/0000-0002-3658-8243
H. Liu https://orcid.org/0000-0001-6892-358X
Z. Luo https://orcid.org/0000-0002-9560-6720
E. Olofsson https://orcid.org/0000-0001-9019-212X
A. Ekedahl https://orcid.org/0000-0002-0376-5119
E. Witrant https://orcid.org/0000-0002-2026-8915

## References

[1] Ding S. and Garofalo A.M. 2023 *Rev. Mod. Plasma Phys.* **7** 4
[2] Garofalo A.M. *et al* 2018 *Plasma Phys. Control. Fusion* **60** 014043
[3] Holcomb C.T. *et al* 2014 *Nucl. Fusion* **54** 093009
[4] Ferron J.R. *et al* 2013 *Phys. Plasmas* **20** 092504
[5] Luce T.C. 2011 *Phys. Plasmas* **18** 030501
[6] Litaudon X. *et al* 2011 *Nucl. Fusion* **51** 073020
[7] Suzuki T. *et al* 2009 *Nucl. Fusion* **49** 085003
[8] Kikuchi M. 1990 *Nucl. Fusion* **30** 265
[9] Fisch N.J. 1987 *Rev Mod. Phys.* **59** 175
[10] Gormezano C. *et al* 2007 *Nucl. Fusion* **47** S285
[11] Gribov Y., Humphreys D., Kajiwara K., Lazarus E.A., Lister J.B., Ozeki T., Portone A., Shimada M., Sips A.C.C. and Wesley J.C. 2007 *Nucl. Fusion* **47** S385
[12] Witrant E., Joffrin E., Brémond S., Giruzzi G., Mazon D., Barana O. and Moreau P. 2007 *Plasma Phys. Control. Fusion* **49** 1075
[13] Felici F., Sauter O., Coda S., Duval B.P., Goodman T.P., Moret J.-M. and Paley J.I. 2011 *Nucl. Fusion* **51** 083052
[14] Felici F. and Sauter O. 2012 *Plasma Phys. Control. Fusion* **54** 025002
[15] Kim S.H. and Lister J.B. 2012 *Nucl. Fusion* **52** 074002
[16] Barton J.E., Besseghir K., Lister J. and Schuster E. 2015 *Plasma Phys. Control. Fusion* **57** 115003
[17] Maljaars E., Felici F., de Baar M.R., van Dongen J., Hogeweij G.M.D., Geelen P.J.M. and Steinbuch M. 2015 *Nucl. Fusion* **55** 023001
[18] Bribiesca Argomedo F. *et al* 2013 *Nucl. Fusion* **53** 033005
[19] Bribiesca Argomedo F. *et al* 2014 *Safety Factor Profile Control in a Tokamak SpringerBriefs in Electrical and Computer Engineering—Control, Automation and Robotics* (Springer)
[20] Gahlawat A. *et al* 2012 Bootstrap current optimization in tokamaks using sum-of-squares polynomials *Proc. 51st IEEE Conf. on Decision and Control* (*Maui, Hawaii, 10–13 December 2012*) (available at: https://ieeexplore.ieee.org/document/6426638)
[21] Mameche H. *et al* 2019 Nonlinear PDE-Based control of the electron temperature in H-mode tokamak plasmas *58th IEEE Conf. on Decision and Control* (*Nice, France, 11–13 December 2019*) (available at: https://ieeexplore.ieee.org/document/9029527)
[22] Wang H. and Schuster E. 2019 *Fusion Eng. Des.* **146** 688
[23] Pajares A. and Schuster E. 2021 *Nucl. Fusion* **61** 036006
[24] Barton J.E., Boyer M.D., Shi W., Schuster E., Luce T.C., Ferron J.R., Walker M.L., Humphreys D.A., Penaflor B.G. and Johnson R.D. 2012 *Nucl. Fusion* **52** 123018
[25] Maljaars E. *et al* 2017 *Nucl. Fusion* **57** 126063
[26] Mavkov B., Witrant E., Prieur C., Maljaars E., Felici F. and Sauter O. 2018 *Nucl. Fusion* **58** 056011
[27] Schuster E. *et al* 2020 *28th IAEA Fusion Energy Conf.* (*Nice, France, 10–15 May 2021*) (available at: https://nucleus.iaea.org/sites/fusionportal/Shared%20Documents/FEC%202020/fec2020-preprints/preprint1024.pdf)
[28] Ljung L. 1999 *System Identification: Theory for the User* (Prenctice Hall PTR)
[29] Kokotovitch P.V. *et al* 1986 *Singular Perturbation Methods in Control: Analysis and Design* (Academic)
[30] Moreau D. *et al* 2008 *Nucl. Fusion* **48** 106001
[31] Moreau D. *et al* 2011 *Nucl. Fusion* **51** 063009
[32] Mavkov B., Witrant E., Prieur C. and Moreau D. 2017 *Control Eng. Pract.* **60** 28
[33] Moreau D. *et al* 2013 *Nucl. Fusion* **53** 063020
[34] Artaud J.F. *et al* 2018 *Nucl. Fusion* **58** 105001
[35] Moreau D. *et al* 2015 *Nucl. Fusion* **55** 063011
[36] Wang S., Witrant E. and Moreau D. 2021 *Fusion Eng. Des.* **162** 112071
[37] Wang S., Moreau D., Witrant E., Qian J.P., Yuan Q.P., Huang Y. and Zeng L. 2021 *Plasma Phys. Control. Fusion* **63** 125001
[38] Maciejowski J.M. 2002 *Predictive Control with Constraints* (Pearson Education Limited)

[39] Qin S.J. and Badgwell T.A. 2003 *Control Eng. Pract.* **11** 733
[40] Borrelli F. and Morari M. 2007 Offset free model predictive control 46th *IEEE Conf. on Decision and Control (Proc. IEEE Conf.)* (*New Orleans, USA, 2007*) p WePI19.2 (available at: https://ieeexplore.ieee.org/document/4434770)
[41] Skogestad S. 2023 *Annu. Rev. Control* **56** 100903
[42] Olofsson E. private communication
[43] Domahidi A. *et al* 2012 Efficient interior point methods for multistage problems arising in receding horizon control *51st IEEE Conf. on Decision and Control (Proc. IEEE Conf.)* (*Maui, Hawaii, 10–13 December 2012*) pp 668–74 (available at: https://ieeexplore.ieee.org/document/6426855)
[44] Grimholt C. and Skogestad S. 2012 Optimal PI-control and verification of the SIMC tuning rule *IFAC Conf. on Advances in PID Control* (*Brescia (Italy), 28–30 March 2012*) p ThPI.1 (available at: https://www.proceedings.com/18095.html)
[45] Leard B., Wang Z., Paruchuri S.T., Schuster E. and Rafiq T. 2024 *Nucl. Fusion* **64** 086052