Profile control simulations and experiments on TCV: a controller test environment and results using a model-based predictive controller

E. Maljaars^a, F. Felici^a, T.C. Blanken^a, C. Galperti^b, O. Sauter^b, M.R. de Baar^c, F. Carpanese^b, T.P. Goodman^b, D. Kim^d, S.H. Kim^d, M. Kong^b, B. Mavkov^e, A. Merle^b, J.M. Moret^b, R. Nouailletas^g, M. Scheffer^f, A.A. Teplukhina^b, N.M.T. Vu^g, the EUROfusion MST1-team*, and the TCV-team$^{**\,b}$

^{a}*Eindhoven University of Technology, Mechanical Engineering, Control Systems Technology, P.O. Box 513, 5600MB Eindhoven, The Netherlands*
^{b}*SPC-EPFL, CH-1015 Lausanne, Switzerland*
^{c}*FOM Institute DIFFER, PO Box 6336, 5600HH Eindhoven, The Netherlands*
^{d}*ITER Organization, Route de Vinon-sur-Verdon, CS 90 046, 13067 St. Paul Lez Durance Cedex, France*
^{e}*Univ. Grenoble Alpes, CNRS, GIPSA-lab, F-38000 Grenoble, France*
^{f}*Eindhoven University of Technology, Applied Physics, Science and Technology of Nuclear Fusion, P.O. Box 513, 5600MB Eindhoven, The Netherlands*
^{g}*CEA, IRFM, F-13108 Saint-Paul-Lez-Durance, France*

Abstract

The successful performance of a model predictive profile controller is demonstrated in simulations and experiments on the TCV tokamak, employing a profile controller test environment. Stable high-performance tokamak operation in hybrid and advanced plasma scenarios requires control over the safety factor profile (q-profile) and kinetic plasma parameters such as the plasma beta. This demands to establish reliable profile control routines in presently operational tokamaks.

We present a model predictive profile controller that controls the q-profile and plasma beta using power requests to two clusters of gyrotrons and the plasma current request. The performance of the controller is analyzed in both simulation and TCV L-mode discharges where successful tracking of the estimated inverse q-profile as well as plasma beta is demonstrated under uncertain plasma conditions and the presence of disturbances. The controller exploits the knowledge of the time-varying actuator limits in the actuator input calculation itself such that fast transitions between targets are achieved without overshoot.

A software environment is employed to prepare and test this and three other profile controllers in parallel in simulations and experiments on TCV. This set of tools includes the rapid plasma transport simulator RAPTOR and various algorithms to reconstruct the plasma equilibrium and plasma profiles by merging the available measurements with model-based predictions. In this work the estimated q-profile is merely based on RAPTOR model predictions due to the absence of internal current density measurements in TCV. These results encourage to further exploit model predictive profile control in experiments on TCV and other (future) tokamaks.

Keywords: profile control, plasma profiles, tokamak transport, model predictive control, TCV

1. Introduction

Control over the safety factor profile and plasma pressure or stored energy is important for high performance tokamak operation, especially in hybrid and advanced scenarios [1]. Specific safety factor profiles can lead to improved confinement by reducing turbulent transport or to steady state operation by maintaining a significant amount of self-driven bootstrap current at

zero loop voltage. Some safety factor profiles (q-profiles) are unstable to MHD modes and need to be avoided. Therefore, it is important to establish reliable profile control routines in currently operational tokamaks that can be transferred to future tokamaks.

Recently, many model-based profile controllers have been developed using a wide variety of controller models and control methods. The applied methods include adaptive control [2], back-stepping control [3], passivity-based control [4, 5], Lyapunov control [6, 7], linear-quadratic-integral control [8, 9], model predictive control [10, 11] and robust control [12]. Some of these have been implemented in experiments, others only in simulations.

The highest plasma performance in tokamaks is often achieved close to actuator limits (e.g. maximum available heating source power) and in addition close to areas of the plasma parameter space that are prone to disruptions or deleterious MHD behavior (e.g. Neoclassic Tearing Modes (NTMs)). Even if these actuator and plasma physics limits may not be restrictive at the target operating point, these are still limiting during transient phases (e.g. ramp-up, ramp-down or transitions between operating points). Therefore, control methods are required that can effectively deal with these actuator limits and ensure operation within the safe plasma parameter space. Most of the control methods mentioned above impose only limits on the actuator signal after it has been computed by the controller. This approach has the important limitation that the controller is not aware of the limits in actuators and plasma in the actuator input calculation and cannot anticipate for these limits.

Model predictive control (MPC) is a well established control method that can take these time-varying actuator and process parameter limits into account in the optimal actuator input calculation itself [13, 14, 15, 16]. MPC was first applied to profile control in simulations with simple models and only fixed actuator constraints [17, 18]. We presented in [10] a model predictive controller that uses multiple linearized models to control the q-profile while effectively dealing with constraints on actuators and physics limits in ITER simulations. In [19] we extended this work by including other controlled variables, time-varying references and nonlinear constraints, and the estimation of state disturbances. The controller performance was demonstrated in ASDEX-Upgrade H-mode simulations. Even more recently, input-constrained MPC is also successfully applied in profile control experiments at the DIII-D tokamak [11].

In this work we present a model predictive profile controller and its performance in simulations and experiments in TCV L-mode discharges during the flat-top phase with transitions to various targets. Successful tracking of targets for the plasma beta (a measure of the volume-integrated plasma pressure, denoted by β) and inverse safety factor profile in the presence of disturbances is achieved as well as effective handling of time-varying input-constraints. It is important to note that, since measurements of the core current density profile on TCV are presently missing, the q-profile that is used as a basis for feedback control in this work is only a model-based estimate. As such, we present these results as a demonstration of implementation and operation of a profile controller rather than a demonstration of having achieved a given q-profile in TCV plasmas.

Implementation of controllers for e.g. the density, temperature and current distribution requires knowledge of these quantities in real-time. These quantities can be reconstructed using real-time equilibrium reconstruction [20, 21] or dynamic state observers [22] that integrate the available (noisy) diagnostic signals into an estimate of the plasma state.

It is beneficial to prepare controllers well before being tested in experiments in order to minimize testing and commissioning time on the experiment. This requires a controller development and test environment involving interfaces to plasma state reconstruction, a fast simulator and possibly experimental data. One example of such a system has been used on DIII-D [23] and is now also used at EAST, KSTAR, and NSTX-U [24]. A similar software environment is being prepared for ITER, where it is known as PCSSP [25]. These environments allow to simulate controllers in closed-loop (also on the control system hardware).

In this work we present also a profile controller development and implementation software environment that is employed for pre-experiment simulations of four profile controllers as well as testing their performance in TCV experiments. This software environment allows to develop profile controllers in representative simulations (both on a local computer and in the TCV control system) and afterwards test them in experiments without changing the controller code itself. The

software environment is used to prepare and test in parallel, next to the model-based predictive controller discussed here, also an adaptive controller [2], a Lyapunov-based controller [7] and a passivity-based controller (see [5] for initial results of this controller using the framework reported here), confirming the validity of the controller preparation and testing approach.

The remainder of this work is organized as follows. In Section 2 we summarize the plasma transport modelling using RAPTOR, that is used to design the MPC controller as well as for closed-loop testing. The experimental physics scenario is presented in Section 3 where also the control problem is defined. The controller development and testing environment is presented in Section 4. The MPC controller design is summarized in Section 5. The performance of the designed controller is analyzed in simulations (Section 6) and experiments on TCV (Section 7). The real-time model-based profile reconstructions are compared to other (off-line) profile reconstructions in Section 8. Finally, conclusions and suggestions for further research are given in Section 9.

2. Modeling plasma transport using RAPTOR

A rapid, control-oriented plasma transport simulator is used in this work for closed-loop controller testing and plasma state reconstruction (Section 4), as well as for deriving the MPC controller (Section 5). We employ RAPTOR [22, 26] in this work, a control-oriented, physics-based 1D transport code. RAPTOR solves the non-linear coupled transport of the electron temperature T_e and the poloidal magnetic flux ψ as a function of the normalized square-root toroidal flux ρ, represented by partial differential equations (PDEs) [27]:

$$\sigma_\parallel \left(\left.\frac{\partial \psi}{\partial t}\right|_\rho - \frac{\rho \dot{\Phi}_b}{2\Phi_b} \frac{\partial \psi}{\partial \rho} \right) = \frac{F^2}{16\pi^2 \mu_0 \Phi_b^2 \rho} \frac{\partial}{\partial \rho} \left[\frac{g_2 g_3}{\rho} \frac{\partial \psi}{\partial \rho} \right] - \frac{B_0}{2\Phi_b \rho} V_\rho' \left(j_{\text{aux}}(u_k) + j_{\text{bs}} \right), \qquad (1)$$

$$\frac{3}{2}(V_\rho')^{-5/3} \left(\frac{\partial}{\partial t} - \frac{\dot{\Phi}_b}{2\Phi_b} \frac{\partial}{\partial \rho} \rho \right) \left[(V_\rho')^{5/3} n_e T_e \right] + \frac{1}{V_\rho'} \frac{\partial}{\partial \rho} \left(-\frac{g_1}{V_\rho'} n_e \chi_e \frac{\partial T_e}{\partial \rho} \right) = P_e(u_k) \qquad (2)$$

Details of the terms in these equations in RAPTOR depend on the plasma scenario or application. We choose here:

- Bootstrap current profile j_{bs} and neoclassical conductivity $\sigma_\parallel$ are calculated using the Sauter-equations [28, 29].

- Multiple sources and sinks of thermal energy are modeled (P_e) including:

  - Ohmic heating.
  - Power deposition by Electron Cyclotron Heating (ECH). The EC actuator power is given in the actuator inputs u_k.
  - Losses from Bremsstrahlung, line radiation, and electron-ion heat exchange. In these calculations, the $T_e(\rho)/T_i(\rho)$ ratio is prescribed.

- The auxiliary current drive $j_{\text{aux}}(u_k)$ by Electron Cyclotron Current Drive (ECCD) is calculated using a scaling law including a $\frac{T_e}{n_e}$ scaling [22].

- An ad-hoc transport model described in [22] is used to calculate the electron thermal diffusivity χ_e.

- Geometry profiles quantities $V' = \frac{\partial V}{\partial \rho}, F = RB_\phi, g_1 = \langle (\nabla V)^2 \rangle, g_2 = \langle (\nabla V)^2/R^2 \rangle, g_3 = \langle 1/R^2 \rangle$ as well as the scalar Φ_b (toroidal flux enclosed by plasma) can be time-varying and provided by a real-time (RT) equilibrium reconstruction codes such as LIUQE [21].

- The electron density profile n_e is prescribed, but can be time-varying and provided by a real-time density profile reconstruction code such as [30, 31].

3

- A boundary condition for (1) is obtained by prescribing the total plasma current, which is seen as an actuator input in u_k. For (2), the edge electron temperature is prescribed.

Inside RAPTOR, (1) and (2) are written as a nonlinear state space model with the following state update equation and its corresponding output equation at time instant t_k:

$$f(x_{k+1}, x_k, u_k) = 0, \tag{3}$$
$$y_k = h(x_k). \tag{4}$$

The state vector x_k contains the coefficients that are used to parameterize the T_e and ψ profiles and the vector u_k contains the actuator inputs, in this case the total plasma current and EC actuator powers. The state update equation gives the state at the next time instant x_{k+1} based on the present state x_k and actuator inputs u_k. The output vector y_k can contain many quantities that are a function of the state. Further details on the numerical methods applied to evolve (1) and (2) in RAPTOR are given in [22]. The extension of RAPTOR to include time-varying geometry and density profiles is detailed in [32]. The used nonlinear state space format allows RAPTOR to provide linearizations of these state and output equations that can be employed in an observer and controller.

3. Plasma scenario and control problem definition

In this section we will define the plasma scenario for the profile control experiments in TCV presented in this work. From this plasma scenario we will also define a control problem in this section.

3.1. Plasma scenario

The designed plasma scenario for the TCV experiments described in this work will be introduced here in terms of chosen hardware and plasma configuration.

The physics goal for these experiments is to routinely achieve and maintain a broad range of targets for the inverse safety-factor profile $\iota(\rho) = 1/q(\rho)$ and the plasma $\beta = \frac{\langle p \rangle_V}{B_0^2/(2\mu_0)}$ by using the available actuators. We may distinguish three cases where we wish to achieve and maintain targets for:

- Only the plasma β

- Only the inverse safety-factor profile $\iota(\rho)$

- Simultaneously the plasma β and the inverse safety-factor profile $\iota(\rho)$

TCV has the feature that many different stationary safety factor profiles can actually be achieved in a single experiment. This is due to the small current redistribution time w.r.t. the total shot time (about 150ms versus 2.5 seconds in this plasma scenario). Hence, it is possible to investigate control of the q-profile during the flat-top phase. The plasma β evolves on the energy confinement time (here about 5ms), indicating the wide range of time scales of the quantities of interest.

To achieve a broad range of safety-factor profiles as well as plasma pressures in these plasmas, central ECH and central co- and counter-ECCD is employed at low density in L-mode discharges. In addition, the plasma current I_p is feedback controlled via the inductive voltage from external coils, and can be adjusted in real-time to broaden the range of achievable profiles. The plasma heating and current drive setup is illustrated in Figure 1. The plasma boundary shape (last closed flux surface) shows a plasma in limiter configuration.

Gyrotrons in TCV are grouped in so-called clusters, which share a single power supply. So-called cluster A uses one gyrotron, connected to launcher 1 (L1) to drive counter-current in these experiments. Cluster B has two gyrotrons that are used to drive co-current via launchers L4 and
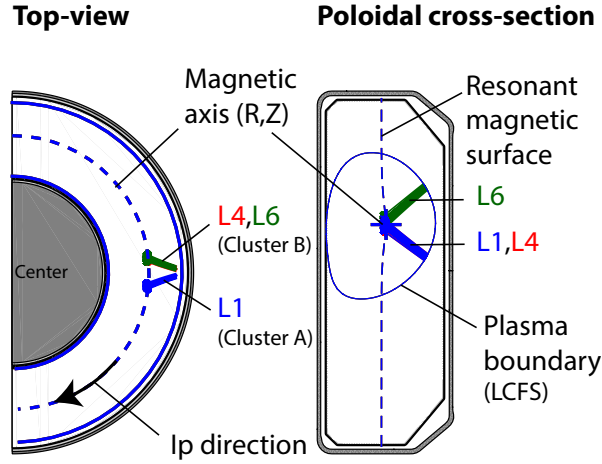
Figure 1: Experimental plasma scenario for TCV plasma profile control experiments using two clusters of gyrotrons/launchers and plasma current I_p as actuators. Plasma is in limiter configuration. Ray-tracing using the TORAY-GA code [33] is used to visualize path EC beams and absorbtion near to the resonant magnetic surface. Cluster A drives counter-current via launcher L1, cluster B drives co-current via launchers L4 and L6.

L6. The magnetic axis is located at the resonant magnetic surface and the aiming of the launchers is such that the EC system effectively provides central co/counter ECCD. The vertical position of the magnetic axis is feedback controlled, whereas only feedforward shape control is used.

The actuator deposition profiles and the achievable plasma profiles are illustrated in Figure 2 for shot #54423 during the flat-top phase. The normalized volume integrated power density profile retrieved from TORAY-GA (a) shows a central and narrow deposition profile. On TCV, an accurate evaluation of the EC deposition width would require a Fokker-Planck code to be used, since radial diffusion of fast electrons broadens the effective width predicted by TORAY-GA [34]. RAPTOR assumes a distribution of EC power prescribed as input to the code. This can be taken from a separate ray-tracing (or Fokker-Planck) code, or manually specified. Here, the EC deposition width in RAPTOR is larger (~ 0.2) than that retrieved from TORAY-GA (~ 0.1). The resulting q-profile will be slightly different in the region where the EC deposition profiles differ. However, the q-profile outside this region depends only on the surface integrated EC current driven, which is in the same order for RAPTOR and TORAY-GA (Figure 2(b)). Since the deposition width is small in any case, it will change the q-profile only on a few radial points. The larger deposition width chosen allows to use a lower spatial resolution for the T_e- and ψ-profiles in RAPTOR, which is essential to achieve real-time capable calculations in TCV. Hereby some freedom is lost in describing highly localized T_e- and ψ-profile features. However, the present main purpose of profile controller performance evaluation does not require many degrees of freedom in describing the q-profile. An offset in the surface integrated driven current can be noticed in Figure 2(b), mainly caused by using different electron temperature and electron density profiles in these calculations (Thomson measurements for TORAY-GA and real-time profile estimates for RAPTOR). The normalized surface integrated current density profile (b) reveals the opposite current drive direction of the EC-clusters.

The reconstructed profiles at this time for T_e and $\iota = 1/q$ are given in panels (c) and (d) respectively. The temperature profile shows that the scenario is in L-mode (no pedestal). The range of achieved profiles in multiple experiments presented in this work is also provided (grey, shaded). Note that some of the reconstructed ι-profiles are nonphysical, significantly higher than 1, which would correspond to $q \ll 1$. In the plasma, the sawtooth instability would prevent such a q-profile, but this effect was not included in the reconstructions shown here (see Section 3.4).
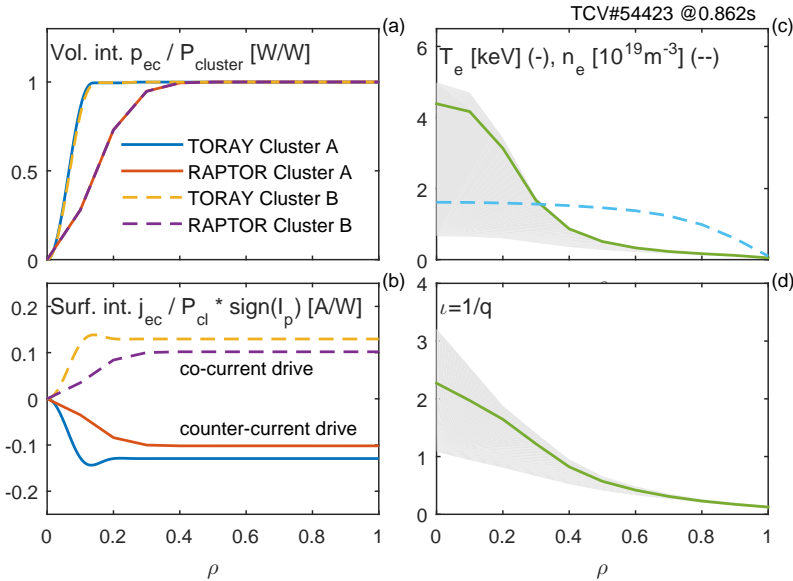
Figure 2: Illustration of actuator deposition profiles and achievable plasma profiles. Normalized volume integrated power and surface integrated current density deposition profiles for EC clusters A and B (a)-(b) obtained from TORAY and RAPTOR; RT reconstructed temperature T_e and density n_e profiles (c), and ι-profile (d). Data from shot #54423 at 0.862s. Note that the EC deposition width used in RAPTOR is larger than the experimental value; this is done to reduce the required spatial resolution to allow the calculations to be real-time capable. The achieved range of T_e and ι-profiles in an number of experiments is also given in (grey, shaded) (color online).

The density (c) is purposely chosen low $(n_e(\rho = 0) \approx 2 \cdot 10^{19} m^{-3})$ in order to maximize the driven current by ECCD, which scales with $\frac{T_e}{n_e}$. The volume-averaged density was controlled to constant value in this plasma scenario.

3.2. Control problem

The defined physics goal in the previous section cannot be achieved using only pre-programmed (feedforward) actuator requests. In that case the ι-profile and plasma β will likely vary due to shot-to-shot varying experimental conditions and the presence of disturbances and MHD-events. Hence active feedback control is required to ensure that targets are achieved and maintained. The physics goal is therefore translated into a control problem with requirements for the profile controller.

The profile controller should be able to:

- Minimize the tracking error for time-varying ι- and/or β-references as given in a performance index (see Section 3.3).

- Be robust against model mismatches and disturbances in the specific class of L-mode discharges as defined in the experiment overview during the current flat-top phase.

- Handle time-varying actuator input constraints (amplitude and ramp-rates).

- Stay within the available 0.7ms computational time per 1ms time step.

We choose to evaluate the controller performance in the flat-top phase, involving transitions between multiple ι- and/or β-targets that require different levels of heating and current. Therefore, the controller is only activated during the flat-top phase. Before controller activation, the plasma shape, density profiles and EC power can evolve towards their nominal values. The controller cycle time $T_s = 1$ms is chosen roughly 5 times smaller than the energy confinement time ($\tau_E \approx 5$ms).
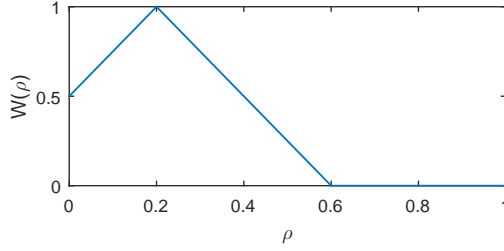
6

Figure 3: Weighting profile $W(\rho)$ defining the region of main interest for ι-profile tracking used in this work.

To achieve the control goal, the controller can act on the plasma current I_p and the power requests to the EC-clusters P_A and P_B such that the actuator input vector u_k at time t_k can be defined as:

$$u_k = \begin{bmatrix} I_\mathrm{p}(t_k) \ [A] \\ P_\mathrm{A}(t_k) \ [W] \\ P_\mathrm{B}(t_k) \ [W] \end{bmatrix} \tag{5}$$

The actuator inputs are subject to (time-varying) amplitude and ramp-rate constraints that can be written as:

$$
\begin{aligned}
u_{\min,k} &\leq u_k &\leq u_{\max,k}, \\
\Delta u_{\min,k} &\leq \frac{u_{k+1}-u_k}{T_s} &\leq \Delta u_{\max,k},
\end{aligned}
\tag{6}
$$

Time-varying power limits are especially important for cluster A. For technical reasons, powers on cluster A in the range between 100kW and 550kW cannot be delivered longer than 470ms. After this time, the requested power must be in the range 550kW-750kW. Cluster B with its two connected gyrotrons does not have this technical limit and can deliver powers between 360kW and 900kW (2x180kW and 2x450kW). No ramp-rate limits are imposed on the power requests to cluster A and B.

Although it is not common in tokamak operation to modify I_p (and hence q_{95}) during the flat-top phase, there is in principle no reason to not change I_p when remaining within appropriate limits. Therefore, besides amplitude limits, ramp-rate limits are used to restrict changes in the plasma current, as large and sudden changes in I_p may lead to plasma performance degradation. The actual values of the imposed limits on I_p vary between simulations and experiments and are given in the result sections 6 and 7.

3.3. Controller performance criterion

The performance of a profile controller is evaluated based on control error indicators for β and the ι-profile. To obtain a scalar error indicator from a ι-profile error, we introduce a weighted norm of the difference between the reference and achieved ι-profile. A mixed norm can also be defined, representing a weighted average of the two norm functions:

$$
\begin{aligned}
J_\beta(t_k) &= \frac{(\beta(t_k) - \beta_\mathrm{ref}(t_k))^2}{(\beta_\mathrm{ref}(t_k))^2} \\
J_\iota(t_k) &= \Sigma_{i=1}^{n_\rho} W(\rho_i) \frac{(\iota(\rho_i,t_k) - \iota_{ref}(\rho_i,t_k))^2}{\iota_{ref}(\rho_i,t_k)^2} \\
J_\mathrm{tot}(t_k) &= \nu_\beta J_\beta(t_k) + \nu_\iota J_\iota(t_k)
\end{aligned}
\tag{7}
$$

The weighting profile $W(\rho)$ allows to set the importance of tracking the ι-profile in a certain region. The chosen weighting profile $W(\rho)$ for this work is given in Figure 3, indicating that we

7

choose to track the ι-profile in the plasma core. Since the outer half of ι-profile has zero weight, the controller is not forced to track the ι-profile near the plasma boundary, i.e. q_{95}, that would fix I_p. Instead, it can use I_p as actuator to ensure faster transitions between ι-profiles in the plasma core. The norm $J_\iota(t_k)$ is computed by evaluating the ι-profile on $n_\rho = 11$ equidistant ρ-grid points.

The weights $\nu_{(\cdot)}$ allow to set the relative importance of β or ι-control and can be used for the three profile control purposes defined in the physics problem:

- β-only control: $\nu_\beta = 1, \nu_\iota = 0$

- ι-only control: $\nu_\beta = 0, \nu_\iota = 1$

- β and ι control: $\nu_\beta = \nu_\iota = \dfrac{1}{2}$

3.4. Note on current density profile estimates

As will be discussed in Section 4.2, the ι-profile and plasma beta (and other plasma profiles and parameters) are computed by a model-based state observer algorithm, which merges model-based predictions with diagnostic measurements. For the temperature profile, a central temperature measurement provides an effective real-time constraint on the central temperature. However, as TCV presently lacks measurements of the internal magnetic field in the plasma region, it is important to note that the q-(or ι-) profile estimates are calculations based exclusively on a poloidal flux diffusion model. As a consequence, the q-profile drops well below 1 in many cases since the effect of sawteeth was not considered, while this is physically unrealistic. In this work we focus, nevertheless, on the performance of the controller in this imperfect situation. While we expect the true q-profile to globally follow the trends reported here, we do not claim that the core q-profile estimate accurately represents the situation in the plasma. A comparison of the real-time reconstructed q-profiles used for control with reconstructed profiles from the equilibrium reconstruction code LIUQE [21] as well as interpretative ASTRA [35] simulations is given in Section 8.

4. Controller development and implementation environment

As mentioned in the introduction, it is important to prepare controllers thoroughly before testing them in experiments. In this section we present the set of controller development and implementation tools that facilitates the efficient preparation of profile controllers in simulations and implementation in experiments. These tools are all prepared in the MATLAB Simulink block-programming language, which greatly facilitates re-using components in different stages of the development and implementation with minimal changes.

First a closed-loop simulation tool is prepared, separate from the TCV control system software environment that allows to largely prepare the profile controller on a local computer. Next we present the controller development and implementation tools inside the TCV control system environment that allows to develop the controller in more comprehensive closed-loop simulations and finally test it in experiments. It is important to note that it is straightforward to transfer these tools to another control testing environment, e.g. PCSSP [25], for re-use in simulations for ITER and other tokamaks.

4.1. Controller development and validation in simulations

A simulation environment has been developed to interface the profile controllers with the RAPTOR simulator for closed-loop simulations. The aim of this simulation environment is to test the profile controller in stand-alone in an ideal situation, without needing to consider interfaces to e.g. state reconstruction codes at the early stage of the controller development. This simulation environment is visualized in Figure 4.

The RAPTOR simulator uses the actuator inputs and prescribed density and geometry information to evolve the plasma profile state and provides the plasma profiles and parameters at the
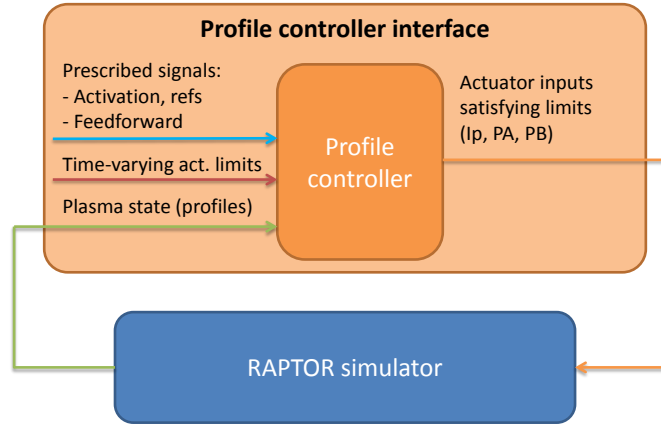
Figure 4: Controller test environment where profile controller is interfaced with the RAPTOR simulator. Profile controller computes actuator inputs using the activation, reference and feedforward signals, the actuator limits and the plasma profiles and parameters provided by RAPTOR simulator. The RAPTOR simulator takes the actuator inputs to compute the plasma profiles for the next time step.

next time step to the profile controller. The profile controller calculates the actuator inputs at each time step based on the provided signals including the RT actuator limits.

To test controllers under circumstances that could be expected during experiments, test cases with different sets of plasma parameters were defined with various targets for the ι-profile and β. These test cases range from different parameters for current drive efficiency and electron heat transport, to simulated events such as the appearance of an Neoclassical Tearing Mode (NTM) resulting in a confinement degradation. Results of closed-loop control simulations using this setup are shown in Section 6.

4.2. Controller development and implementation in TCV control system software environment

Once the controller has been tested and verified in the environment described in Section 4.1, it is ready to be tested in the more complete TCV control system software environment that includes other real-time algorithms. The digital TCV control system is distributed over several computational nodes linked via shared reflective memory and accommodates hundreds of diagnostic inputs and actuator outputs [36]. Details on the recent upgrades to and applications of the TCV control system can be found in [37]. The control system is entirely programmable by MATLAB Simulink, allowing rapid algorithm development and implementation using code generation [38]. The developed profile controller development and implementation environment within the TCV control system allows three crucial applications:

1. Closed-loop testing in simulation using a plant simulator, on a host computer.

2. Hardware-in-the-Loop simulation on the actual real-time hardware, after automated generation of C-code.

3. Experimental testing on tokamak.

More details on these three applications will be given next after which we discuss the profile controller implementation and state reconstruction algorithms.

Closed-loop control simulations using a plant simulator

The profile controller can be further developed and tuned under experimental conditions in closed-loop simulations e.g. by using experimental geometry and density profiles of a specific shot. In addition most dynamics and delays of the closed-loop are included that may affect the controller
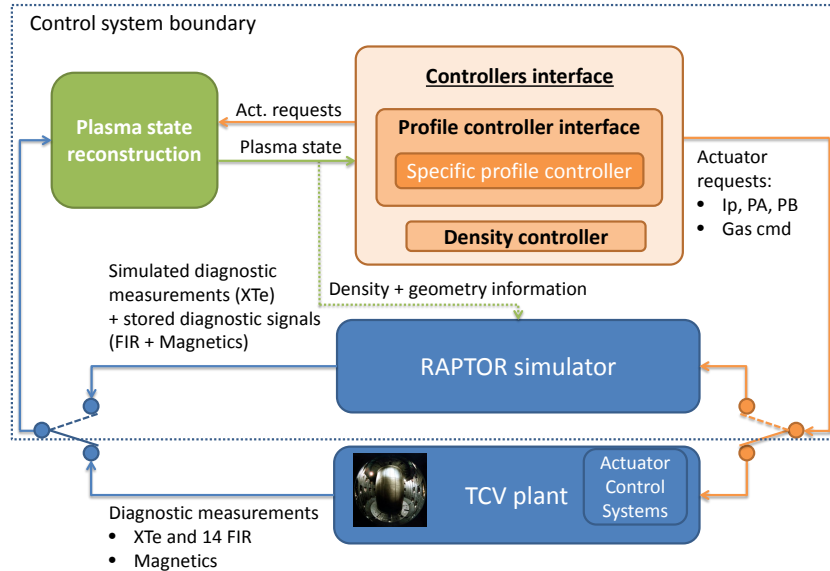
Figure 5: Implementation of profile controller in TCV RT control system allowing to test the controller in experiments on the TCV plant and also in closed-loop simulations. Profile controller and density controller receive plasma state from a set of RT plasma state reconstruction algorithms (see Figure 6 for more details).

performance: e.g. signal routing delays and reconstruction delays. The effect of controller gains can be easily analysed in these closed-loop simulations. In addition, the design of new shot references is eased by checking e.g. the expected proximity of actuator limits.

Hardware-in-the-Loop simulations

In this case, the simulator runs on a dedicated computational node of the control system, using the same interface to the true experimental signals. This Hardware-in-the-Loop test can be seen as the ultimate verification step before deploying the entire suite of reconstruction and control systems in experiments. It is particularly useful for detecting exceeding of computational time limits, or issues related to the integrated use of all control system components.

Experimental testing

The final step is testing the profile controller in closed-loop experiments. A profile controller can be tested in experiments if it has been shown to work properly in closed-loop simulations inside the TCV control system environment, it passes code compilation tests and satisfies the computational time constraint in hardware-in-the-loop simulations.

The codes can also be run in open-loop in background during experiments performed for other purposes (the actuators do not respond to the profile controller, but to e.g. pre-programmed actuator requests). These are known as piggyback experiments and allow to test and improve the codes without the need for dedicated experiments. This has been especially useful in developing and interfacing the state reconstruction algorithms.

Overview of implementation and involved algorithms

An overview of the profile controller implementation is given in Figure 5.

The profile controllers can be tested in experiments in closed-loop on the TCV plant or in closed-loop simulations. The controllers receive the plasma state from a set of reconstruction algorithms that use either the diagnostic measurements or the simulated diagnostics in combination with stored diagnostic signals from a previously performed shot. The set of RT state reconstruction algorithms is explained next. The interface to the profile controller itself is identical to Figure 4.
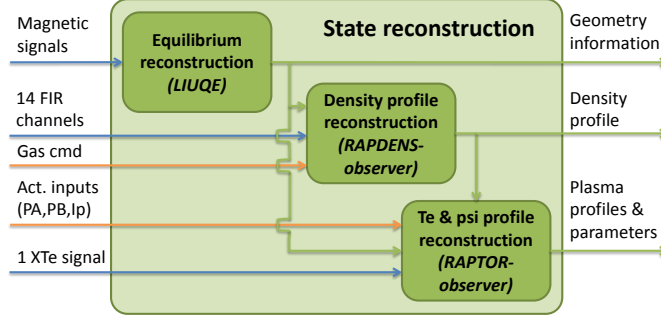
Figure 6: Real-time state reconstruction algorithms and there interfaces for estimation of electron temperature and poloidal magnetic flux, density and equilibrium.

Density control in the experiments is performed using a recently developed robust density controller [30], details of which will be provided in a separate publication. This controller computes the request to a single gas valve to achieve and maintain a specified volume averaged electron density value.

Plasma state reconstruction

Reconstruction of the plasma state from the scarce available measurements is essential for closed-loop control of plasma profiles and parameters. This plasma state involves here knowledge of the 2D plasma equilibrium, density profile and the electron temperature and poloidal magnetic flux profiles. The algorithms used to reconstruct the plasma state and their interfaces are illustrated in Figure 6.

The following real-time reconstruction algorithms provide every 1ms a new plasma state estimate:

RAPTOR-observer The electron temperature and poloidal magnetic flux is reconstructed in an Extended Kalman Filter (EKF) scheme [39]. This algorithm uses the density profile from RAPDENS-observer, the geometry and plasma current estimate from LIUQE, a core electron temperature measurement from soft X-ray (XTe) and the EC power requests from the profile controller to merge diagnostics with RAPTOR model predictions. The estimated central temperature $T_e(\rho = 0)$ will slowly track the XTe value, as the EKF estimates a disturbance heating source/sink that compensates for discrepancies between measurements and the undisturbed model predictions. Due to the lack of RT current density distribution measurements in TCV, the ι-profile is fully determined by the RAPTOR model predictions, constrained only by the experimental total plasma current measurement. The RAPTOR-observer gives many scalar and profile outputs, the most important of which are, for the present purposes:

- Electron temperature profile $T_e(\rho)$
- Poloidal magnetic flux profile $\psi(\rho)$
- Inverse safety factor profile $\iota(\rho)$
- Plasma β

More details on the implementation of RAPTOR-observer in TCV can be found in [40].

RAPDENS-observer Using the same EKF approach as in RAPTOR-observer, RAPDENS-observer reconstructs the density profile from far-infrared (FIR) measurements in combination with model-based predictions from a density evolution model [30, 31]. Automatic detection and correction for diagnostic artefacts like fringe jumps in FIR channels is included.

LIUQE A 2D Grad-Shafranov equilibrium reconstruction code that uses magnetic measurements [21]. It supplies equilibrium geometry information to the RAPTOR-observer, the RAPDENS-observer and the profile controller. At present the RAPTOR q-profile estimate is not fed back to LIUQE.

More details on recent advances of these state reconstructions algorithms will be provided in separate publications.

5. Model Predictive Controller design

We will now design a profile controller using Model Predictive Control (MPC) following the controller requirements defined in Section 3.2. First we discuss MPC after which we summarize the controller design. We conclude with the main controller settings and their effect on the controller performance.

5.1. Brief introduction to MPC

MPC is a well-established advanced control method that has been used for decades to control multiple-input-multiple-output (MIMO) processes in industry that are subjected to input and state constraints [13, 14, 15, 16]. Its principle is visualized in Figure 7.

To compute the optimal actuator requests, an MPC controller relies on a prediction model of the involved process to calculate the future evolution of the controlled variables up to a prediction horizon as a function of the present state and future actuator input sequence. This prediction model is used in an optimization problem to find the future actuator input sequence that minimizes a cost function. If the cost function is set-up to minimize the difference between controlled variables and references up to the prediction horizon, the MPC controller will yield actuator commands to track these references. Only the first time step of the computed actuator sequence is sent to the actuators and at the next time step a new optimization problem is solved based on the estimated present state and possibly updated references. In this optimization problem constraints on actuator inputs and other variables can be taken into account, which may be time-varying due to changing conditions in actuators and plant.

5.2. Outline of controller design procedure

Here the controller design is summarized, which is largely based on previous work in [10, 19], with minor changes to simplify the implementation and meet the set (computational) requirements. Further details on the controller design can be found in [41].

Selected state and controlled variable representation

The controller requires a definition of the process state and of the controlled variables. We will employ a linearized model from RAPTOR in this work (will be explained next), therefore we choose to use the same state vector x_k as in RAPTOR (see Section 2). This state vector contains cubic spline coefficients $\hat{\psi}$ and $\hat{T}_e$ to parameterize the T_e and ψ profiles ($n_b = 12$ coefficients for each profile, see [22] for more details). As controlled variables z_k we choose both the $\iota \propto \frac{1}{\rho}\frac{\partial \psi}{\partial \rho}$-profile (at the locations of the provided 11 point equidistant ρ-grid where $W(\rho) \neq 0$, see Figure 3), and the plasma beta $\beta \propto \int_0^1 V' n_e T_e d\rho$. Both can be calculated as combinations of the state vector for a given plasma geometry and density profile. In vector form these read as:

$$x_k = \begin{bmatrix} \hat{T}_{e1}(t_k) \\ \vdots \\ \hat{T}_{en_b}(t_k) \\ \hat{\psi}_1(t_k) \\ \vdots \\ \hat{\psi}_{n_b}(t_k) \end{bmatrix}, \quad z_k = \begin{bmatrix} \iota_{\rho=0}(t_k) \\ \iota_{\rho=0.1}(t_k) \\ \vdots \\ \iota_{\rho=0.5}(t_k) \\ \beta(t_k) \end{bmatrix}. \tag{8}$$
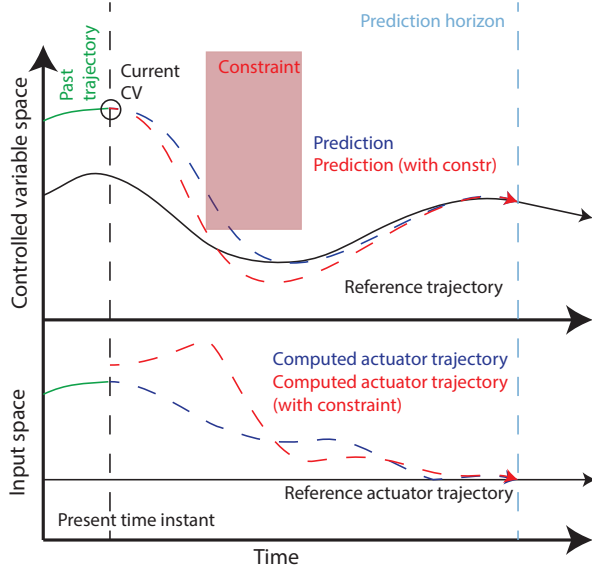
12

Figure 7: Illustration of MPC principle. MPC uses a prediction model to predict the evolution of controlled variables up to a prediction horizon (a) based on the present state and future actuator inputs (b). The MPC controller computes future actuator input trajectories that will minimize the tracking error between the controlled variable and its reference, even in the presence of constraints on actuators or controlled variables.

Prediction model

The prediction model used to predict the future controlled variable evolution is constructed from a single linearized model provided by RAPTOR (see Section 2). As part of the control design procedure, a simulation is run at plasma current and EC-power levels expected in the experiment, until stationary profiles are achieved and then the linearized model is selected at the linearization point (u^0, x^0, z^0). Noting that as the current redistribution time in TCV is approximately 150ms for these plasmas, we choose as prediction horizon 100 time steps of 1ms to capture the main part of the slow evolution of the ι-profile within the predictions.

The controller is expected to be able to compensate for the inevitable mismatches between the linearized model and the actual plasma transport dynamics. The plasma transport dynamics are expected to be only weakly nonlinear in these experiments due to starting control in the current flat-top phase, keeping (almost) the same actuator and plasma configuration, and the absence of confinement mode transitions. This was confirmed in a comparison of the observed nonlinearity of the plasma transport dynamics in the conducted experiments (presented in Section 7) to the linearized controller model [41].

Controller performance objective

The controller performance objective is to minimize the tracking error on the controlled variables and to avoid aggressive control actions, which can be expressed in the cost function J_k at time t_k as follows:

$$J_k = \nu_\iota \underbrace{(R_k - Z_k)^T Q_\iota (R_k - Z_k)}_{\text{Tracking error } \iota} + \nu_\beta \underbrace{(R_k - Z_k)^T Q_\beta (R_k - Z_k)}_{\text{Tracking error } \beta} + w_{\Delta \tilde{U}} \underbrace{\Delta \tilde{U}_k^T R_{\Delta \tilde{U}} \Delta \tilde{U}_k}_{\text{Input changes}}, \qquad (9)$$

The first two terms penalize deviations between the future reference sequence R_k and future controlled variable sequence Z_k from time t_k to the horizon t_{k+N}. The matrices Q_ι and Q_β select the appropriate controlled variables for β or ι respectively and includes the weighting profile $W(\rho)$. The parameter $w_{\Delta \tilde{U}}$ can be used to set the importance of penalizing aggressively changing inputs

13

in the future input sequence, where $\Delta\tilde{U}$ is the sequence of input changes between time steps. More freedom to set this penalty for individual actuators is available in $R_{\Delta\tilde{U}}$. The chosen cost function ensures that if the controller cost $J_k \to 0$, also the provided error norm $J_{\text{tot}} \to 0$ in (7).

Optimization problem

The optimal future actuator inputs are computed by minimizing the quadratic control performance objective (9) subject to the actuator amplitude and ramp-rate limits (6). This can be formulated as a standard Quadratic Programming optimization problem that is solved at each time step using a fast open-source QP-solver (Quadprog++ [42]) that was embedded in Simulink. The actuator requests for the next time step are taken from the QP-problem solution and sent to the actuators.

It should be noted that although constraints on plasma parameters (state constraints) can be readily imposed in the optimization problem, only actuator input constraints are applied in this work. Given the available computational time of 0.7ms per time step, imposing also state constraints was expected to be infeasible at this timescale. Analysis of the used computational time per time step in the presented experiments revealed that it never exceeded 0.3ms [41].

State and disturbance estimation

The controller requires the present plasma profile state x_k as defined in (8). This state can be retrieved from the T_e and ψ profiles and their gradients that are provided by RAPTOR-observer.

In the absence of integral action in a profile controller, both static disturbances acting on the system and model mismatches between the controller model and the actual plasma profile evolution may result in steady state tracking errors. One way to overcome this in an MPC controller is to estimate in a disturbance observer the discrepancies between the controller-model-based expected states and controlled variables and those estimated from measurements. In this work we use such a disturbance observer to estimate the slowly varying component of these discrepancies. The MPC controller uses these disturbance estimates to improve the model-based prediction and ensure that targets are achieved with minimal steady state error. This effectively introduces integral action in the MPC controller [13, 16].

5.3. Main controller settings

The controller behavior can be tuned by changing a relatively small number of controller parameters with a clear impact on the control behaviour:

- The penalty on input changes $w_{\Delta\tilde{U}}$ given in (9). Increasing $w_{\Delta\tilde{U}}$ results in smoother actuator trajectories but a slower convergence of controlled variables to the references.

- The relative penalty on changes in I_p vs. EC powers. This parameter $R_{\Delta\tilde{U}I_\text{p}/P_{EC}}$ is part of the matrix $R_{\Delta\tilde{U}}$ defined in (9). Increasing this value results in more aggressive use of EC rather than I_p.

- The time required to let the disturbance estimates converge towards the observed disturbances, where τ_x sets the convergence time for state disturbance estimates and τ_z the convergence time for disturbances estimates on controlled variables. These parameters set how fast the controller anticipates on changes in the disturbances, high values yield a slow response to changes in disturbances, which translates to a slow convergence towards the references in stationary conditions.

The chosen controller settings in the simulation and the experiments are given in Table 1.

6. Simulation results

Off-line simulations have been performed using the MPC-controller interfaced with the RAPTOR-simulator as described in Section 4.1. These closed-loop simulations are used to tune the controller

Table 1: Controller settings in simulation and experiments. Listed in order as presented.

Sim/Exp	$w_{\Delta \tilde{U}}$	$R_{\Delta \tilde{U} I_\mathrm{p}/P_{EC}}$	$\tau_x = \tau_z$
Sim 1	5	10	10ms
54385	40	20	10ms
54423	30	20	10ms
54402	20	20	10ms
54414	20	80	60ms

parameters and to test the controller's robustness in various test cases. The parameters of the transport model and current drive efficiency in the RAPTOR-simulator are chosen representatively of the physics setup described in Section 3.

We present here a test case involving both control of the ι-profile and plasma β to demonstrate the performance and main properties of the controller. We introduce a significant model mismatch between the physics parameters in the controller and the simulator. This model mismatch is a 30% higher current drive efficiency for both clusters in the simulator parameters, resulting in different required powers to reach the references than those powers used to generate the references. We use the controller settings in Table 1 and use both clusters A (counter-ECCD) and B (co-ECCD) as well as the plasma-current. Actuator amplitude limits are shown next and the ramp-rate limits for I_p are set to $\|\frac{dI_\mathrm{p}}{dt}\| \leq 2\mathrm{MA/s}$.
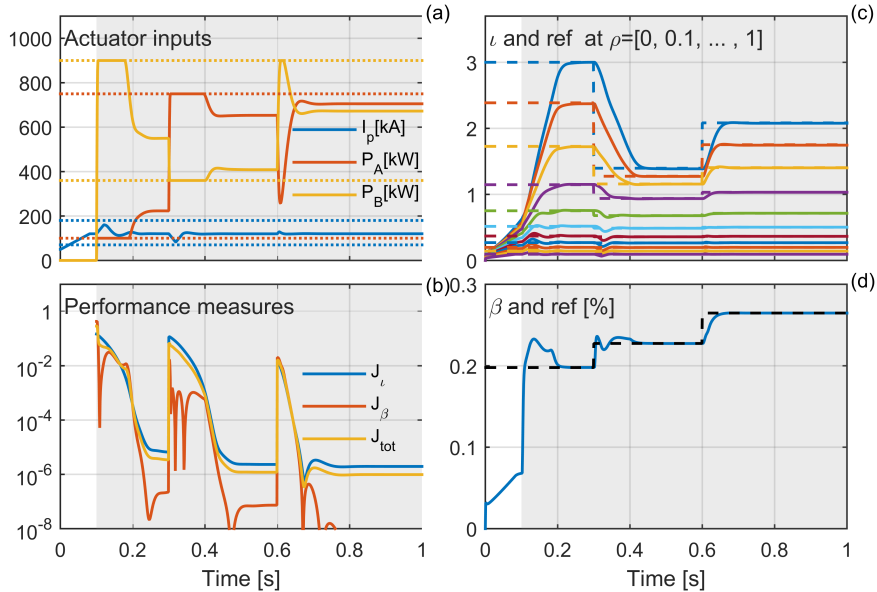

Figure 8: Performance of MPC profile controller in closed-loop simulation using plasma current and gyrotron clusters A (counter-ECCD) and B (co-ECCD) as actuators (a), where various ι (c) and β (d) targets are successfully tracked. This good tracking is also visible in the performance measures (b).

Figure 8 presents the performance of the MPC controller in the closed-loop simulation, where three references for β and ι are successfully tracked. The actuators (panel (a)) are all used to track the ι-profile (c) and β (d), which can also be noticed in the performance measure (b) that reaches small values despite the model mismatch. Actuator limits are given in dashed lines (a).

As can be seen in (a) at t=0.1s, the controller uses the overshoot in the plasma current I_p together with the maximum co-current (cluster B) and minimum counter-current (cluster A) to steer the ι-profile as quickly as possible towards the first target, at the expense of some overshoot in β (b). The same behavior can be noticed during transition to the third target and opposite

behavior can be observed at the start of the second target. When the ι-target is almost reached, both cluster requests are changed so as to minimize the error on β and prevent an overshoot on ι. The transition to the second target is slowed down by the active actuator limits on both clusters. A very small steady state error is achieved for all controlled variables (b) at all targets, thanks to the disturbance estimation that captures the steady state effect of the model mismatch in the disturbance estimates.

This simulation result shows the successful tracking of both ι and β references and effective handling of the actuator limits so as to realize fast transitions between the targets. We conclude from this (and other) simulations that the controller is ready for testing in the TCV control system software environment.

7. Experimental results on TCV

In this section we present the results of using the controller in closed-loop on TCV in four discharges with different control objectives, references and actuators, using the implementation described in Section 4.2. The results are presented in the order of increasing complexity, starting with β-only control using only EC-clusters in feedback, followed by combined β and ι control using both EC-clusters and finally adding I_p as an actuator. The plasma scenario was discussed in Section 3 and the controller settings used are listed in Table 1.

7.1. Plasma β control using both EC clusters

In the first presented shot (#54385), the controller needs to track a β-reference with multiple steps, where the controller has the freedom to use both cluster A and B. The β value at the moment of controller activation is taken as the initial reference value and a staircase reference is added to this value. In addition, a disturbance is purposely applied to the system by adding a ramp-up and ramp-down in the plasma current I_p after 1.4s. In the controller settings a high penalty on aggressive input changes $(w_{\Delta\tilde{U}})$ was chosen, such that smooth actuator input trajectories are expected.

Figure 9 presents the performance of the controller in this β-control shot. We first briefly introduce the various traces shown in the figures here, since these will come back in all experimental results. The actuator inputs are given in (a)-(c), where the shaded area indicates the active feedback control phase for the actuators used by the profile controller. The density control is visualized in (d) where the density reference and estimate as well as the gas command are given. The β and ι estimates and references (dashed, for those controlled) are given in (e) and (f) respectively, and the shaded area indicates the active control phase for the controlled variables. The performance measures (error norms) are given in the panel (g) for the controlled variables. Other information about the temperature (h) and density evolution (i), as well as the estimated NTM normalized likelihood for several mode numbers (j) is also provided. The normalized NTM likelihood is provided by the recent implementation of [43]. If the NTM likelihood marker is close to 1, then this mode is strongly present in the magnetic signals.

The steps in the β-reference are successfully tracked (e) by adjusting the powers of the EC-clusters (b)-(c). As a high penalty on input changes was chosen, the actuator input traces are smooth but the controller cannot react quickly to changes in the density (i), resulting in some overshoot in β.

After about 1.2s, cluster A has spent too much time below 550kW and the lower limit is raised to 550kW, avoiding EC power supply shut-off. The power requested on cluster B drops to the minimum power and the reference cannot be reached.

When the I_p disturbance (a) starts to increase, the density increases even though the gas valve is closed (i). The EC-powers cannot be decreased (b)-(c) due to the minimum power limits, resulting in a significant error in β ((e) and (g)). When the I_p disturbance decreases again (a), the density decreases too (i) as well as β (e), such that the controller quickly raises the EC requests (b)-(c). The plasma disrupts when the plasma current I_p drops below 100kA (a), possibly due to the onset of a strong 2/1 NTM (j).
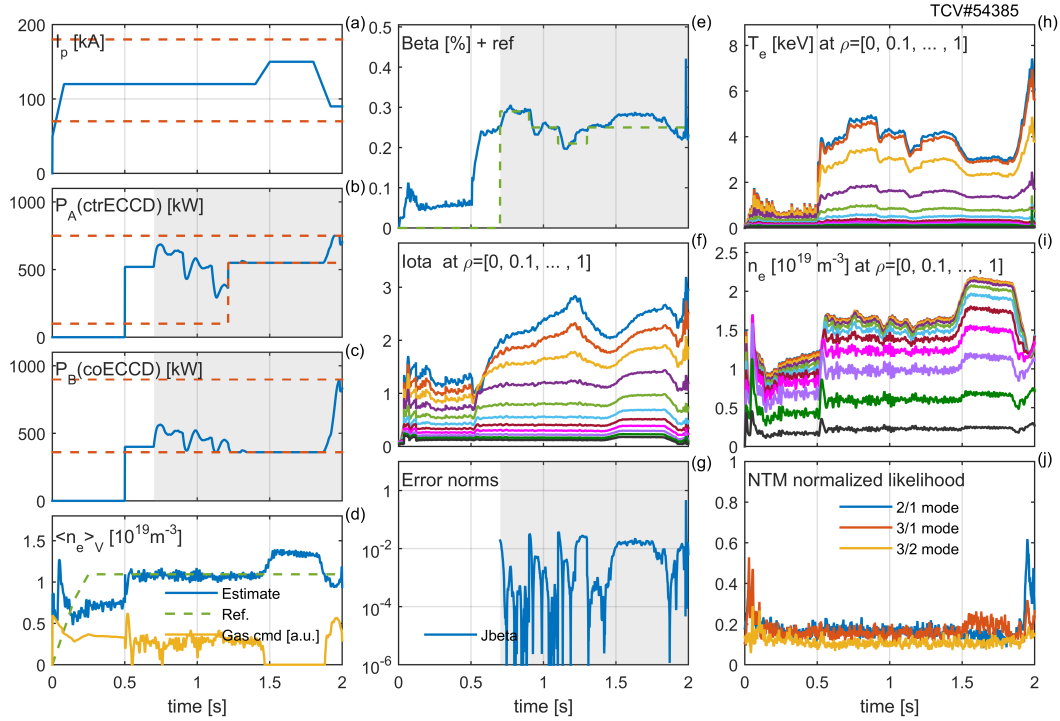
Figure 9: Performance of β-control in shot #54385 using both cluster A and B with a prescribed perturbation from the nominal I_{p} (a). Calculated actuator inputs for cluster A and B (b)-(c) to successfully track the β-reference with effective handling of actuator limit change in cluster A (b). During the applied I_{p} disturbance (a) the density increases (d,i), resulting in a high β. The heating power cannot be reduced as both clusters are on their respective minimum powers (b) and (c). At the moment the plasma current I_{p} drops below 100kA (a), a NTM is triggered (j) and the plasma disrupts just before 2s.

This experiment shows that the controller can successfully control the plasma β as far as it is reachable within the actuator limits. Also the change in the minimum power for cluster A is effectively taken into account. The small overshoot in β could be reduced by choosing a lower penalty on actuator input changes, such that the controller can react rapidly to the changes in the density. This is done in the next presented experiment.

7.2. Combined β and ι-control using both EC-clusters

In the next two experiments both EC-clusters are used to control simultaneously β and ι, where we start in the first experiment with multiple steps in β while keeping the ι-profile constant. Subsequently an experiment is performed where the controller needs to track a sequence of four targets for both β and ι.

Steps in β reference at constant ι

In this experiment (#54423) the controller needs to track a β-reference with multiple steps whereas the ι-target is kept constant.

The results for shot #54423 are presented in Figure 10, where the controller successfully uses both clusters A and B (b)-(c) to keep the ι-profile constant (f), while at the same time quickly reaching the β-targets with hardly any overshoot (e). The clusters A and B move in the same direction during the β-steps to effectively not change the driven current. Immediately after controller activation, the EC-clusters make tiny opposite movements so as to compensate the ohmic current density relaxation while keeping the β constant. Although the volume averaged
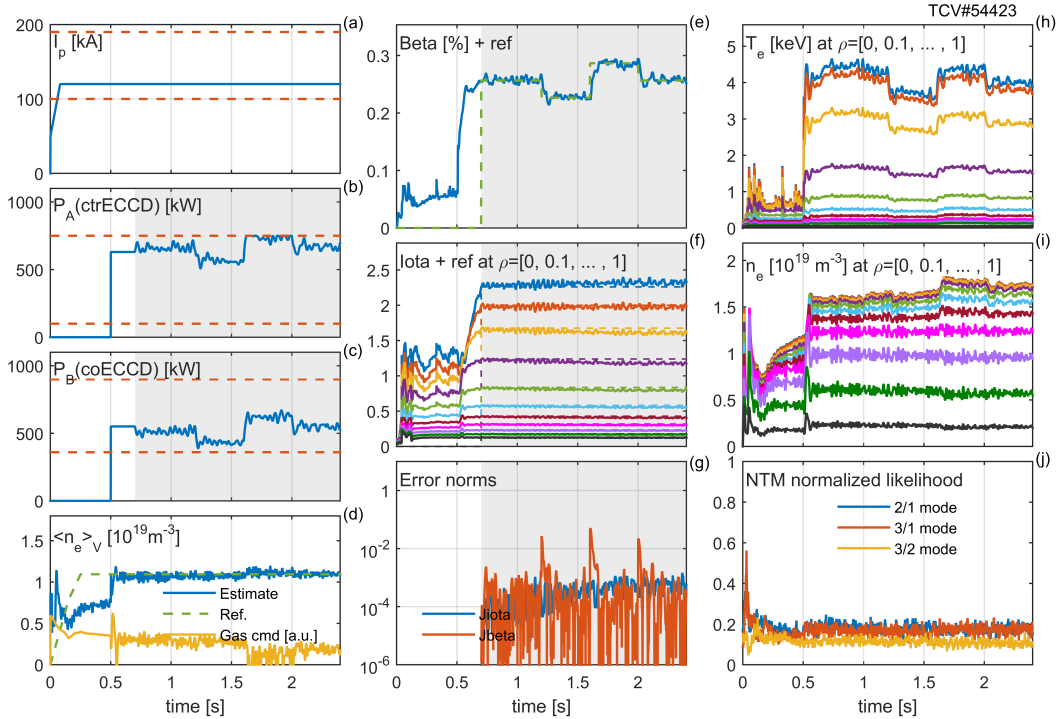
Figure 10: Successful tracking of steps in β (d) at constant ι (f) in shot #54423 using both EC-clusters (b)-(c).

density is constant (d), the density profile distribution peaks slightly during the phase with the highest β-target.

We conclude for this experiment that the controller can tightly control the plasma β while at the same time the ι-profile can be kept constant. Smooth and quick transitions are achieved without overshoot in β or the ι-profile, indicating that the controller effectively takes relevant couplings into account.

Multiple targets for β and ι

In the next shot #54402 the controller needs to track four different targets for both β and ι where it can use both clusters.

The controller performance is presented in Figure 11, where the second (between $t = 700$ms and $t = 1100$ms) and third target (between $t = 1100$ms and $t = 1500$ms) are achieved with a very small remaining error (e)-(g), while the other two targets have some remaining error for various reasons. The transition from one target to another (e.g. from second to third at $t = 1100$ms) is very fast without overshoot, by pushing first both clusters against opposite actuator limits (b)-(c) and then adapting so as to prevent overshoot on ι (f) while keeping β constant (e). A similar result was obtained in simulations, see Figure 8. This can only be achieved by accounting for the actuator limits over the prediction horizon in the control input computation. It demonstrates how a predictive controller can outperform more traditional feedback control methods in terms of constraint handling.

During the short feedforward phase between 0.5 and 0.7s, a transient electron internal transport barrier (eITB) was created by the high counter current drive from cluster A (b) [44], resulting in high temperature T_e (h) and high β (e). The high temperature is estimated from the core electron temperature measurement XTe, while the transport model parameters do not use such scenarios with local confinement improvement at reversed shear profiles (the related model parameters are set to zero). The eITB is lost at the start of the feedback control phase when the cluster A
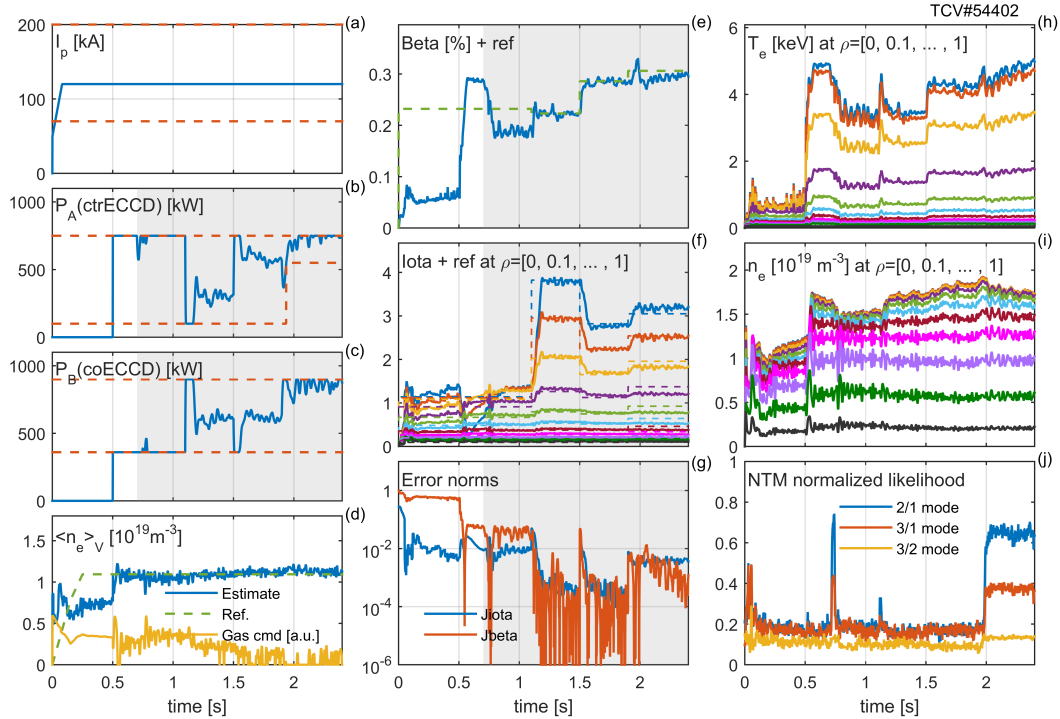
18

Figure 11: Successful tracking of several β (d) and ι (f) targets in shot #54402 using both EC-clusters (b)-(c). Four different targets are requested, where the first and last are not fully reachable with the used actuator settings and density profile. Fast transitions between targets are obtained by exploiting knowledge of actuator limits in the control input computation.

power is lowered to reach the β target and MHD activity is detected for a short time (j). Post-shot analysis of soft X-ray data and electron temperature and density profile evolution revealed a minor disruption that leads to the loss of the eITB, which happens when internal MHD stability is lost [45]. It was shown in [46] that ideal internal modes can trigger minor disruptions due to the high pressure gradient in a low magnetic shear region (infernal modes), even at low normalized beta.

The first target cannot be reached due to the tight actuator limits. The ι-profile in the center could potentially be lowered by using more counter-current drive (cluster A) or less co-current drive (cluster B). However, cluster A and B are already at their upper and lower limit respectively, such that the target ι-profile cannot be achieved. At the same time the first β-reference cannot be reached due to the low density (i) compared to the $n_{e,0} = 2 \cdot 10^{19} m^{-3}$ used to generate the references. The last β-target cannot be reached due to the presence of NTMs (j), while the last ι-target cannot be reached as the last target was generated at a higher plasma current $I_{\mathrm{p}} = 130$kA.

The reference and achieved ι-profiles at 100ms after the target switch are given in Figure 12. Target 2 and target 3 are best achieved as discussed before. Note that the controller uses the weighting profile $W(\rho)$ such that e.g. for target 3 (at $t = 1600$s) the profile error is minimal at the peak of the weighting profile $W(\rho)$.

This experimental result shows the controller's capability to simultaneously control β and the ι-profile, ensuring fast transitions between targets without overshoot. The controller takes advantage of knowing the actuator limits in the control input calculation to realize these fast transitions, such that targets are closely achieved within 100ms, even faster than the current redistribution time scale of about 150ms.
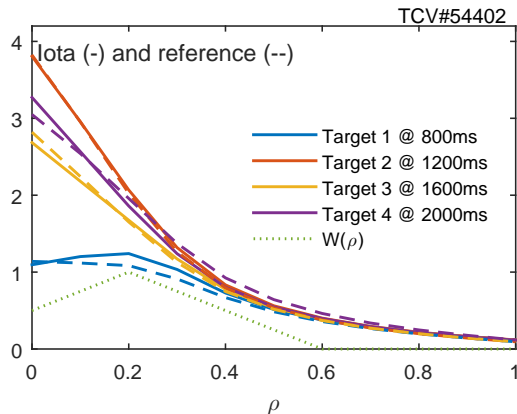
Figure 12: Reference and achieved ι-profiles for the 4 targets in shot #54402 at 100ms after a target switch. Targets are best achieved close to the peak of the weighting profile $W(\rho)$.

7.3. Combined β and ι control including plasma current

In the last presented experiment #54414 we repeated the previous experiment #54402 with the four targets for β and ι while adding the plasma current I_p as actuator. This should provide more control freedom, especially enabling to reach the last target that was generated with $I_p = 130$kA. In this shot the plasma current ramp-rate was restricted to $-0.5\text{MA/s} \leq \frac{dI_p}{dt} \leq 0.5\text{MA/s}$, changes in I_p where significantly more penalized than changes in cluster powers, and a slow convergence of the disturbance estimates to measured disturbances was set (see Table 1).

This last experimental result is presented in Figure 13, where in general the performance of the controller seems worse than without including I_p. The control-loop involving the plasma current I_p is clearly oscillating, while being bounded by the ramp-rate limits (a).

Changes in I_p immediately result in changes in the density (i) due to changing particle confinement and recycling conditions, for which the density controller tries to compensate (d). The fluctuations in I_p lead also to plasma volume variations and in alternating between a limiter and divertor shape (not shown). An eITB was again created during the feedforward phase (e), but just before the active feedback control phase starts ($t = 700$ms) the eITB is lost due to MHD activity (j).

Remarkably, comparing the error norm for ι (g) to the one in Figure 11(g) shows improved ι-tracking in this experiment for the first and last target. To reach the first target, I_p is lowered to its minimum so as to make the ι-profile less peaked. The plasma current I_p is correctly raised in the last target, oscillating around the 135kA, close to the value used to generate the reference. Contrarily, the error in β is clearly larger, mainly due to the induced density oscillations and the inability of the controller to compensate quickly for these using the clusters as the disturbance estimates are converging too slow to the measured disturbances.

The main expected origins of the closed-loop instability in the loop involving I_p are listed here in order of their likely contribution:

- Unmodelled closed-loop delays and actuator dynamics for I_p resulting in a large delay of 7-15 ms before the I_p-values requested by controller were noticed in the measured I_p fed to the RAPTOR-observer. This delay exceeds by far the energy confinement time (5ms).

- Unmodelled effects of changing the plasma current I_p on the density profile, whereas it is known that particle confinement and recycling changes with the plasma current.

- No active plasma shape feedback control is used. The plasma volume changes significantly in shot #54414 and also changes frequently from limiter to divertor shape.
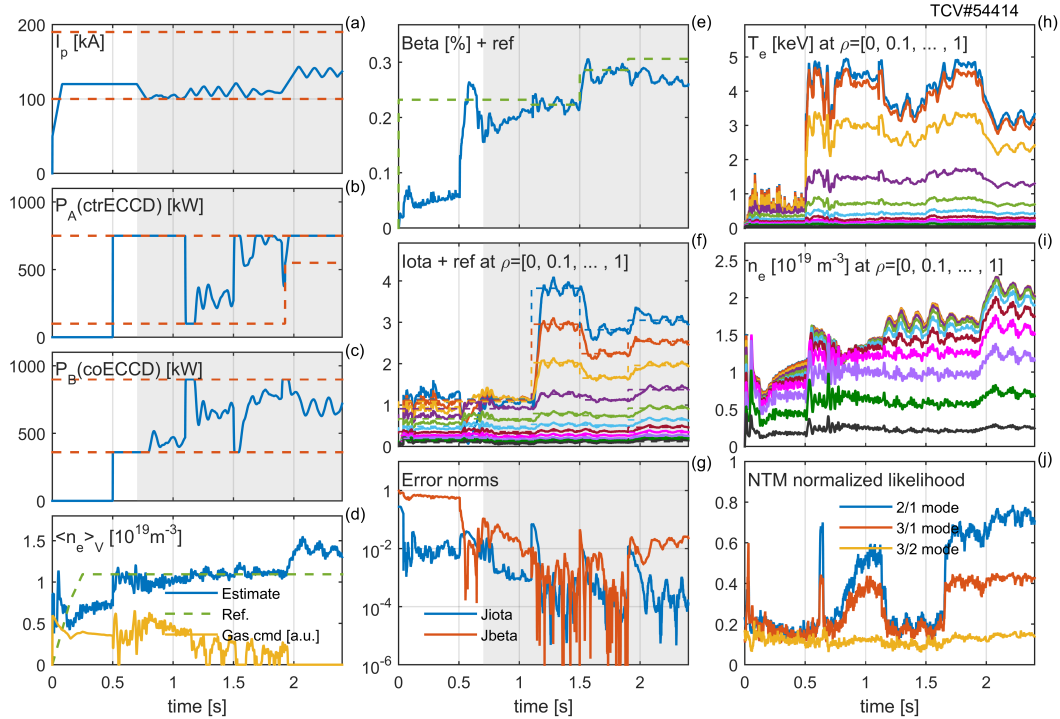
Figure 13: Performance of controller in shot #54414 where the same four targets for β and ι as in #54402 need to be tracked, but with the plasma current I_p as third actuator. Although the ι tracking improves for the first and last target (f)-(g) by lowering or increasing I_p (a), the control-loop involving I_p is clearly unstable (but bounded by the ramp-rate limits). Changes in I_p correspond to changes in density (d,i) and lead to the presence of NTM (j), resulting in a bad β-tracking (e).

We expect that correctly modeling the delays and actuator dynamics will already remove the closed-loop instability. The last two expected origins of the closed-loop instability can be solved by adding an equation for the density in RAPTOR and including shape control in the experiments.

This experiment shows the potential and limitations of using the plasma current I_p in feedback in the present software configuration. It is shown that using the plasma current as feedback actuator allows to achieve a broader range of q-profiles. Applying the described remedies to remove the closed-loop instability should enable the controller to efficiently use the additional control freedom provided by I_p.

8. Comparison of online and off-line profile reconstructions

To investigate the sensitivity and accurateness of the reconstructed profiles by RAPTOR-observer, we compare these to other (off-line) reconstructed profiles. For the shot #54402, we will now compare the real-time reconstructed q-profiles and plasma β from RAPTOR-observer with those obtained in:

- Real-time reconstruction using LIUQE [21]. See Section 4.2 for more details.

- Off-line reconstruction using LIUQE [21]. The off-line version of LIUQE uses different basis functions than the real-time version, allows more iterations of the numerical scheme and includes measurements from the diamagnetic loop (DML).

- Interpretative simulations using the ASTRA transport code [35]. In these simulations of the poloidal magnetic flux evolution, ASTRA uses the EC deposition profiles from TORAY-GA [33], the Thomson electron density and temperature profile measurements, a simplified
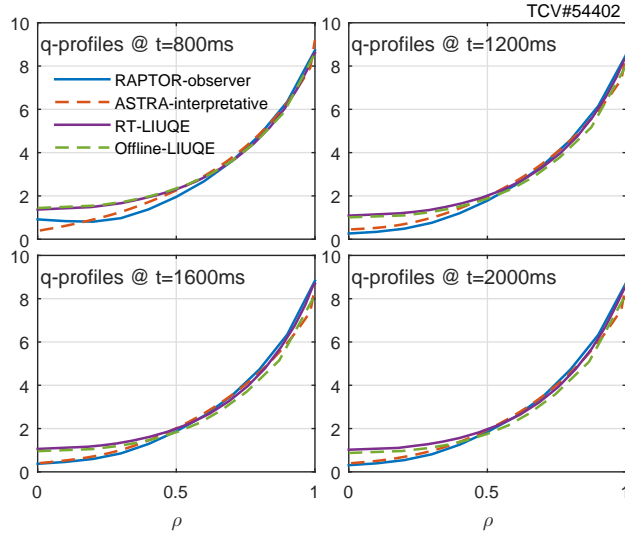
Figure 14: Comparison of reconstructed q-profiles (at the same four time steps as in Figure 12) from the RAPTOR-observer to real-time reconstructions of LIUQE and interpretative ASTRA simulations.

(3-moment) equilibrium based on the plasma boundary reconstructed by LIUQE, and the Z_{eff} value was constrained using the edge loop voltage.

The q-profiles at the time steps given in Figure 12 are compared in Figure 14 to these three alternative reconstructions. Note the agreement in the reconstructions from mid axis to plasma edge. The discrepancy in the center between the LIUQE reconstructions and the RAPTOR reconstruction and ASTRA simulation is due to the fact that while RAPTOR and ASTRA include the simulated effect of the EC current drive on the core q-profile, LIUQE employs a parameterization of the q-profile with only a few degrees of freedom to fit the (magnetics-only) measurements (which is why we do not use the the real-time reconstructed q-profiles from LIUQE in these profile control experiments). However, RAPTOR and ASTRA do not take into account here the effect of sawteeth which usually reduces the radius of the $q = 1$ surface and increases the on-axis value of q.

Figure 15 compares the time evolution of the ι-profile (a), electron temperature T_e (b), electron density n_e (c) in RAPTOR-observer and in ASTRA interpretative simulations, the actuator inputs (d) are also plotted. Electron temperature and density profiles for both reconstructions are given at three time instants in Figure 16.

The ι-profile evolutions from both reconstructions are overall in good agreement, except during the phase between $t = 0.5$s and $t = 0.7$s when an electron-Internal Transport Barrier (eITB) is present, and during transitions between targets. During the presence of the eITB, the Thomson diagnostic measures a twice as high central temperature than the (line-averaged) XTe diagnostic (see Figure 15(b) and Figure 16(a)), resulting in more peaked conductivity and current density profiles, i.e. a higher $\iota(\rho = 0)$ (lower $q(\rho = 0)$). Notice that during transitions between targets, the $\iota(\rho = 0)$ in ASTRA moves initially in the opposite direction as RAPTOR-observer. This is caused by a difference in the temperature and density profiles. For example, when cluster A (driving counter-current) is pushed to full power at $t = 1.5$s, the Thomson diagnostic measures quickly afterwards at $t = 1.53$s a stronger increase in central temperature than the XTe diagnostic (see Figure 15(b) and Figure 16(c)-(e)). Simultaneously, a decrease in the core density can be noticed (see Figure 15(c) and Figure 16(d)-(f)). The redistribution of particles in the plasma core cannot be observed in RAPDENS-observer as it relies on line-integrated FIR measurements. Both differences result in a more peaked conductivity and current density profile (i.e. higher $\iota(\rho = 0)$).

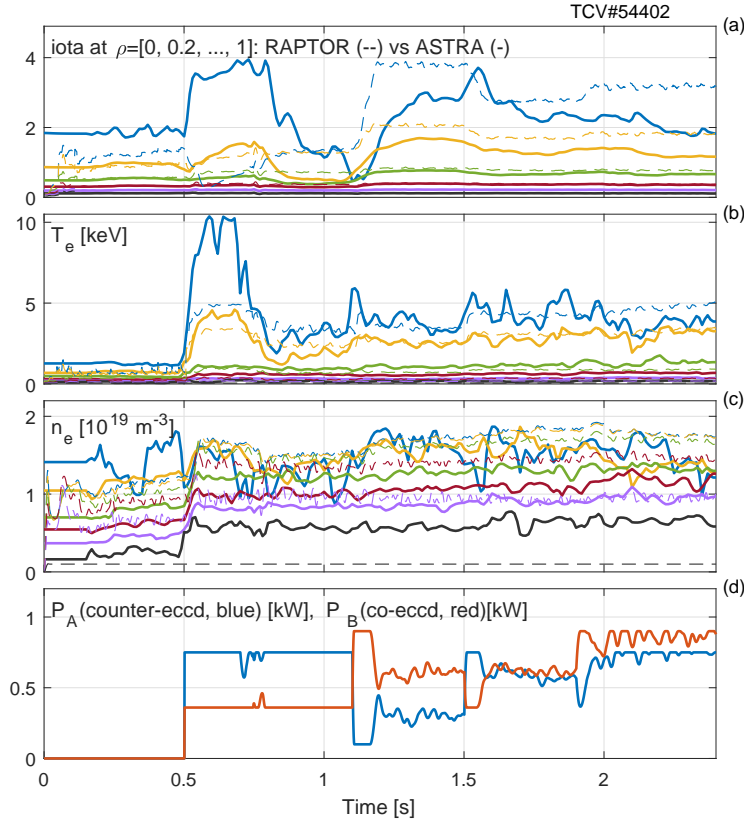The plasma β evolution as reconstructed in the RAPTOR-observer is also compared to off-

Figure 15: Comparison of reconstruction by real-time RAPTOR-observer and off-line interpretative ASTRA. Time evolution of ι-profile (a), electron temperature profile (b), electron density profile (c) and actuator inputs (d).
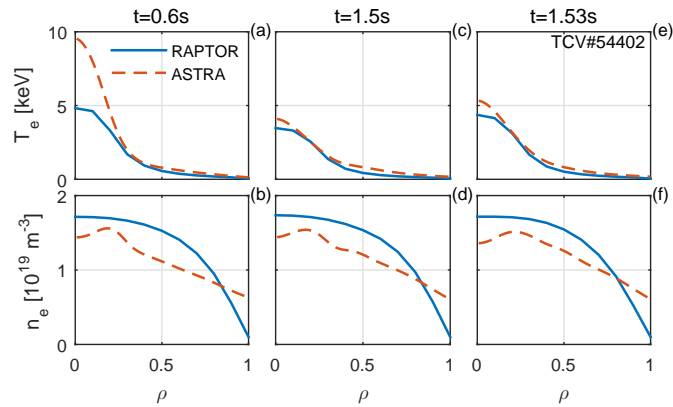

Figure 16: Comparison of reconstructed temperature and density profiles in real-time RAPTOR-observer and fits of Thomson scattering measurements used as input to off-line interpretative ASTRA simulations.
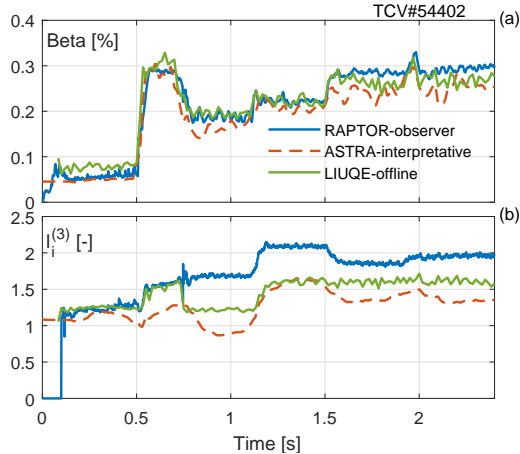
Figure 17: The plasma beta (a) and normalized internal inductance (b) evolution reconstructed from RAPTOR-observer, post-shot reconstruction using LIUQE and interpretative ASTRA simulations.

line reconstructions from LIUQE as well as the interpretative ASTRA simulation in Figure 17(a). Note the agreement between all three reconstructions, indicating that the actual plasma β behaves as reconstructed in the RAPTOR-observer.

In addition, the normalized inductance $l_i^{(3)} = 2V\langle B_p^2 \rangle_V/[(\mu_0 I_p^2)R_0]$ (defined in [47]) for these three reconstructions is given in Figure 17(b). The scalar $l_i^{(3)}$ is a measure of the peakedness of the current density profile and is well constrained in LIUQE by magnetic measurements, contrarily to the q-profile itself. Note that $l_i^{(3)}$ as reconstructed from RAPTOR is very close to LIUQE during the ohmic phase. During the eITB phase (0.5s-0.7s) these $l_i^{(3)}$-reconstructions are also similar, despite the underlying q-profiles being very different (as we have seen in Figure 14). As we have seen also in Figure 15, RAPTOR misses the collapse of the eITB. The current density profile of RAPTOR after 0.7s is too peaked (i.e. a too high $l_i^{(3)}$). This high $l_i^{(3)}$ is mainly caused by the fact that the $q(\rho = 0)$ in RAPTOR-observer becomes very small. This could be avoided in the future by using the sawtooth module implemented recently in RAPTOR [48] and incorporating real-time Thomson measurements instead of the line-averaged FIR and XTe measurements may further improve RAPTOR predictions. The reconstruction by ASTRA interpretative simulations has a similar trend as LIUQE, although it is lower, which is linked to the differences in the current density (or q-) profile.

From the comparisons given in this section, we can conclude that the plasma beta was reasonably reconstructed in RAPTOR-observer, but as its q-profile reconstruction is entirely model-based and relying on indirect measurements of the electron density and temperature, it is very sensitive to the diagnostic chosen for these quantities. Using direct measurements of the internal current density profile (e.g. using the Motional Stark Effect diagnostic or polarimetry) would avoid this strong dependency on indirect measurements and allow to control the true plasma q-profile. At the same time, the model-based predictions in the RAPTOR-observer could be improved by using the sawtooth model [48] and real-time Thomson measurements of the electron density and temperature that can measure these quantities in the plasma core more accurately and with higher spatial resolution than the line-integrated measurements from XTe and FIR.

9. Conclusions and outlook

This work has demonstrated the successful performance of a model predictive profile controller in experiments in the TCV tokamak, employing a profile controller test environment.

We designed a linear MPC controller including disturbance estimation and demonstrated its performance in both simulations and experiments in a TCV L-mode plasma scenario. The results

show successful tracking of the inverse safety factor profile as well as the plasma beta using two clusters of gyrotrons/launchers in the presence of uncertain plasma conditions and disturbances. The controller exploits the knowledge of the time-varying actuator limits in the actuator input calculation such that fast transitions between targets are achieved without overshoot, demonstrating how a predictive controller can outperform other control methods in terms of effective input-constraint handling. Avoiding overshoots can ensure to avoid reaching stability limits when controlling high performance plasmas inherently close to several stability limits.

Secondly, a profile controller development and implementation environment is presented that is used to prepare and test profile controllers both in simulations and experiments on TCV. Next to the MPC-controller, it is also used to prepare and test three other profile controllers in parallel, confirming the efficiency of this framework. It allowed us to first prepare profile controllers interfaced to a simulator on a local computer. Next the controllers were tested more comprehensively inside the TCV real-time control system including plasma state reconstruction codes and experimental data. The employed software was used to automatically generate code for the control system including profile controllers to test the real-time performance in hardware-in-the-loop simulations and finally in experiments.

This work can be extended in several ways to further utilize the advantages of model predictive profile controllers. The controller performance can be further improved by accurately modeling the actuator dynamics and delays. The controller would be able to use the plasma current more successfully as actuator if the plasma current actuator dynamics and delays as well as its coupling with the density profile dynamics would be included in the linearized controller model and shape feedback control would be employed in experiments.

Adding internal current profile diagnostics as input for the RAPTOR-observer would yield more accurate q-profile estimates, ensuring that the tracked q-profile is closer to the true q-profile of the plasma rather than a purely model-based estimate as is presently the case. We noticed that the reconstructed q-profiles are sensitive to the transport model parameters as well as to the electron density and temperature estimates that are fed to the model. We expect that using improved transport models [49, 50], including the effect of sawteeth [48] and using real-time Thomson measurements with sufficient spatial resolution in the plasma core will alleviate these effects.

Exploiting model predictive profile control throughout multiple plasma regimes (e.g. H-mode, internal transport barriers) would be eased if the controller was independent of a predefined linear model. This can be realised by using real-time linearizations provided by the RAPTOR-observer. Fully nonlinear MPC, involving a nonlinear prediction model, is not computationally feasible on currently operational devices, but can be considered for e.g. ITER with its slower characteristic time scales and hence increased available computational time.

Demonstrating handling of plasma parameter limits (e.g. normalized plasma pressure) in experiments on currently operational tokamaks is important for developing the control expertise to ensure reliable high-performance operation close to these limits in (future) large tokamaks. In this work we imposed only input constraints, whereas MPC can also handle state constraints (as shown in [10, 19]). Sufficient computational time is available to add such a limit to the controller in future profile control experiments at TCV.

In (future) plasma control systems, a profile controller will be interfaced with a supervisory controller that may set controller activation, references and parameters as well as available actuators in real-time based on the plasma state and detected events [51]. The MPC controller's ability to deal with time-varying references, activations and actuator limits encourages to use it in the further development of integrated control strategies involving the management of actuators shared between several control tasks.

References

[1] Gormezano, C. *et al.* *Progress in the ITER Physics Basis Chapter 6: Steady State Operation* 2007 *Nuclear Fusion* **47** S285

[2] Kim, S.H. *et al.* 2012 *Nuclear Fusion* **52** 074002

[3] Boyer, M.D. *et al.* 2014 *IEEE Trans. Contr. Sys. Techn.* **22** 1725

[4] Vu, N.M.T. *et al.* 2014 *IFAC Proceedings Volumes* **47** 11398

[5] Vu, N.M.T. *et al.* 2017 *Fusion Engineering and Design*

[6] Bribiesca Argomedo, F. *et al.* 2013 *Nuclear Fusion* **53** 033005

[7] Mavkov, B. 2017 *Control of Coupled Transport in Tokamak Plasmas (Chapter 3)* Theses Communauté Universite Grenoble Alpes

[8] Boyer, M.D. *et al.* 2013 *Plasma Physics and Controlled Fusion* **55** 105007

[9] Moreau, D. *et al.* 2013 *Nuclear Fusion* **53** 063020

[10] Maljaars, E. *et al.* 2015 *Nuclear Fusion* **55** 023001

[11] Wehner, W. *et al.* 2016 in *2016 IEEE Conference on Control Applications* IEEE 629–634

[12] Barton, J.E. *et al.* 2014 *IFAC Proceedings Volumes* **47** 5223

[13] Maciejowski, J. 2002 *Predictive Control with Constraints* Prentice-Hall

[14] Mayne, D.Q. 2014 *Automatica* **50** 2967

[15] Camacho, E. *et al.* 2004 *Model Predictive Control* Springer London ISBN 9781852336943

[16] Rossiter, J. 2013 *Model-Based Predictive Control: A Practical Approach* Taylor & Francis ISBN 9780203503966

[17] Ou, Y. *et al.* 2011 *Control Engineering Practice* **19** 22

[18] Ouarit, H. *et al.* 2011 *Fusion Engineering and Design* **86** 1018

[19] Maljaars, E. *et al.* 2015 *IFAC-PapersOnLine* **48** 314

[20] Ferron, J. *et al.* 1998 *Nuclear Fusion* **38** 1055

[21] Moret, J.M. *et al.* 2015 *Fusion Engineering and Design* **91** 1

[22] Felici, F. *et al.* 2012 *Plasma Physics and Controlled Fusion* **54** 025002

[23] Ferron, J.R. *et al.* 1995 in *Fusion Engineering, 1995. SOFE'95. Seeking a New Energy Era., 16th IEEE/NPSS Symposium* volume 2 IEEE 870–873

[24] Walker, M.L. *et al.* 2015 in *2015 IEEE 26th Symposium on Fusion Engineering (SOFE)* IEEE 1–8

[25] Walker, M.L. *et al.* 2015 *Fusion Engineering and Design* **9697** 716

[26] Felici, F. *et al.* 2011 *Nuclear Fusion* **51** 083052

[27] Hinton, F.L. *et al.* 1976 *Rev. Mod. Phys.* **48** 239

[28] Sauter, O. *et al.* 1999 *Physics of Plasmas (1994-present)* **6** 2834

[29] Sauter, O. *et al.* 2002 *Physics of Plasmas (1994-present)* **9** 5140

[30] Blanken, T.C. *et al.* 2015 in *42nd EPS Conference on Plasma Physics, Lisbon, Portugal*

[31] Blanken, T.C. *et al.* 2017 *Control-oriented modeling of the plasma particle density in tokamaks and application to real-time density profile reconstruction* Accepted for publication in Fusion Engineering and Design

[32] Teplukhina, A.A. *et al.* *Simulation of profile evolution from ramp-up to ramp-down and optimization of tokamak plasma termination with the RAPTOR code* Accepted for publication in Plasma Physics and Controlled Fusion: `https://doi.org/10.1088/1361-6587/aa857e`

[33] Matsuda, K. 1989 *IEEE transactions on plasma science* **17** 6

[34] Nikkola, P. *et al.* 2003 *Nuclear fusion* **43** 1343

[35] Pereverzev, G. *et al.* 2002 *ASTRA. Automated system for transport analysis in a tokamak* Technical report Max-Planck-Institut für Plasmaphysik

[36] Le, H.B. *et al.* 2014 *Fusion Engineering and Design* **89** 155

[37] Anand, H. *et al.* 2017 *Nuclear Fusion* **57** 056005

[38] Felici, F. *et al.* 2014 *Fusion Engineering and Design* **89** 165

[39] Felici, F. *et al.* 2014 in *American Control Conference (ACC), Portland, USA, 2014* 4816–4823

[40] Felici, F. *et al.* 2016 in *Preprint: 2016 IAEA Fusion Energy Conference, Kyoto EX/P8-27*

[41] Maljaars, E. 2017 *Model predictive profile control and actuator management in tokamaks (Chapter 4)* Theses Technische Universiteit Eindhoven

[42] Di Gaspero, L. 2016 *Quadprog++: a C++ library for Quadratic Programming which implements the Goldfarb-Idnani active-set dual method.* Available from: http://github.com/liuq/QuadProgpp

[43] Galperti, C. *et al.* 2014 *Plasma Physics and Controlled Fusion* **56** 114012

[44] Zucca, C. *et al.* 2009 *Plasma Physics and Controlled Fusion* **51** 015002

[45] Turri, G. *et al.* 2008 in *Journal of Physics: Conference Series* volume 123 IOP Publishing 012038

[46] Martynov, A. 2005 *EPFL thesis No. 3218 (Chapter 6.3)* Ph.D. thesis Lausanne

[47] Jackson, G. *et al.* 2008 *Nuclear Fusion* **48** 125002

[48] Piron, C. *et al.* 2015 in *42nd EPS Conference on Plasma Physics, Lisbon, Portugal*

[49] Citrin, J. *et al.* 2015 *Nuclear Fusion* **55** 092001

[50] Kim, D. *et al.* 2016 *Plasma Physics and Controlled Fusion* **58** 055002

[51] Humphreys, D.A. *et al.* 2015 *Physics of Plasmas (1994-present)* **22** 021806