

First-principles-driven model-based current profile control for the DIII-D tokamak via LQI optimal control

Mark D Boyer¹, Justin Barton¹, Eugenio Schuster¹, Tim C Luce²,
John R Ferron², Michael L Walker², David A Humphreys²,
Ben G Penaflor² and Robert D Johnson²

¹ Department of Mechanical Engineering and Mechanics, Lehigh University, Bethlehem, PA 18015, USA

² General Atomics, 3550 General Atomics Court, San Diego, CA 92121, USA

E-mail: m.dan.boyer@lehigh.edu

Received 17 January 2013, in final form 25 June 2013

Published 2 September 2013

Online at stacks.iop.org/PPCF/55/105007

Abstract

In tokamak fusion plasmas, control of the spatial distribution profile of the toroidal plasma current plays an important role in realizing certain advanced operating scenarios. These scenarios, characterized by improved confinement, magnetohydrodynamic stability, and a high fraction of non-inductively driven plasma current, could enable steady-state reactor operation with high fusion gain. Current profile control experiments at the DIII-D tokamak focus on using a combination of feedforward and feedback control to achieve a targeted current profile during the ramp-up and early flat-top phases of the shot and then to actively maintain this profile during the rest of the discharge. The dynamic evolution of the current profile is nonlinearly coupled with several plasma parameters, motivating the design of model-based control algorithms that can exploit knowledge of the system to achieve desired performance. In this work, we use a first-principles-driven, control-oriented model of the current profile evolution in low confinement mode (L-mode) discharges in DIII-D to design a feedback control law for regulating the profile around a desired trajectory. The model combines the magnetic diffusion equations with empirical correlations for the electron temperature, resistivity, and non-inductive current drive. To improve tracking performance of the system, a nonlinear input transformation is combined with a linear-quadratic-integral (LQI) optimal controller designed to minimize a weighted combination of the tracking error and controller effort. The resulting control law utilizes the total plasma current, total external heating power, and line averaged plasma density as actuators. A simulation study was used to test the controller's performance and ensure correct implementation in the DIII-D plasma control system prior to experimental testing. Experimental results are presented that show the first-principles-driven model-based control scheme's successful rejection of input disturbances and perturbed initial conditions, as well as target trajectory tracking.

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

1. Introduction

1.1. Motivation

For nuclear fusion to become a commercially viable means of producing energy, tokamak reactors must be capable of

operating at a high fusion gain (the ratio of power produced to power required to sustain a discharge) for extended periods of time, ideally reaching steady-state operation. ITER, the next experimental step for fusion research, is attempting to show the scientific feasibility of such operation and must address numerous challenges in order to be a success. As

part of this effort, much work has been done to identify and study advanced tokamak (AT) operating scenarios, marked by improved confinement, magnetohydrodynamic stability, and a large fraction of non-inductively driven plasma current arising from the self-generated bootstrap current. Achieving and maintaining such scenarios could lead to steady-state operation and reduced size and cost for future reactors. The approach to reaching certain AT scenarios involves setting up a suitable spatial distribution profile of the toroidal plasma current or the safety factor, q , a related quantity (see definition in section 2), during the ramp-up and early flattop phases of the discharge and actively maintaining it throughout the remainder of the shot. Control of the minimum value of the q profile, $q_{\min}(t)$, is a requirement for mitigating plasma instabilities, while tailoring the entire q profile will be necessary to maximize the fraction of bootstrap current. Developing techniques to actively control the evolution of the entire q profile is an important issue that has recently begun to receive a great deal of attention.

1.2. Prior work

In previous experiments in this area [1, 2], real-time feedback control of scalar parameters characterizing the current profile was demonstrated. The internal inductance parameter, l_i , a measure of the current profile shape, was controlled using LHCD (lower hybrid current drive) on the Tore Supra tokamak [1], while feedback control of either $q(0, t)$, the safety factor at the magnetic axis of the plasma, or $q_{\min}(t)$, the minimum value of the safety factor profile, was achieved on the DIII-D tokamak [2] by modifying ECH (electron cyclotron heating) or NBI (neutral beam injection). In both experiments, simple non-model-based proportional control laws were used to control single scalar parameters describing some characteristic of the current profile. Additional work considering non-model based real-time control of the current profile can be found in [3–5].

For certain advanced operating scenarios, it may be necessary to precisely shape the entire current profile. The high dimensionality of this problem, along with the strong nonlinear coupling between magnetic and kinetic profiles motivates the use of model-based control designs that can accommodate this complexity through embedding the known physics within the design. By capturing the response of the system to the available actuators in a control-oriented model, such control designs can achieve improved performance without the need for extensive trial-and-error tuning. Control-oriented models of the current profile dynamics can be divided into two groups: those created largely from experimental data, and those that depend on limited experimental data and are primarily first-principles driven.

Using data gathered during dedicated experiments, system identification techniques can be used to develop linear dynamic models of the plasma profile response to various actuators. This approach has been carried out at JET [6–8], JT-60U, and DIII-D [9–11]. In these results, data-driven models were used to design controllers for simultaneously regulating magnetic and kinetic plasma profiles around desired references during the flat-top phase of a plasma discharge. Because the data-driven models are linear, they are only valid close to the

reference scenario used for identification and controllers based on them may not perform well when the state of the system moves far from the reference. Additionally, dedicated system identification experiments are required to apply this approach to new devices or, potentially, to different operating scenarios.

To avoid the issues associated with identified models, controllers can be developed using models driven by a first-principles description of the current profile dynamics. Use of a first-principles model allows a control design to incorporate the nonlinear coupling of plasma parameters, enabling improved system performance and potentially allowing for operation over a much wider range of conditions than control schemes based on linear data-driven models. Additionally, first-principles based controllers can potentially be adapted to different scenarios or devices without significant changes to the structure of the control law and while avoiding the need for dedicated model identification experiments. In practice, it is often necessary to make some approximations and simplifications to the first-principles model in order to facilitate control design or to close the equations of the model. If these approximations are made carefully, the performance and flexibility of a first-principles based approach to control design can be mostly retained while the complexity of the control design can be greatly reduced. Recently, first-principles-driven control-oriented models of the evolution of the plasma magnetic flux have been developed [12–14]. In [15, 16], the model developed in [13] was used to determine optimal feedforward actuator trajectories for achieving a desired safety factor profile, utilizing extremum-seeking and nonlinear programming approaches. Some recent work on first-principles-driven feedback control designs have been presented in [17–21]. In these results, robust, optimal, and sliding mode techniques are used to design feedback controllers, which are then tested in simulations.

1.3. Results of this work

The control approach we propose is to combine feedforward actuator trajectories, calculated offline (as in [15, 16]) or based on previous experimental results, with a first-principles-driven model-based feedback control law to adjust the actuator trajectories in real-time in order to track a desired current profile evolution. This combined approach adds robustness and versatility, as it makes it possible to either reproduce profile evolutions achieved in previous experiments despite perturbations in initial conditions or other disturbances, or to track a reference profile evolution for which feedforward actuator trajectories are not well known, e.g. evolutions which have not been previously produced in an experiment. In this work, we focus on the design, implementation, and experimental testing of a possible approach to the feedback portion of the current profile control scheme. The controller is designed to achieve a desired reference trajectory of the poloidal flux gradient profile, which is inversely related to the safety factor profile and is denoted in this work as θ . We base the design on the control-oriented model developed in [13], which combines the magnetic diffusion equation with empirical correlations for electron temperature, resistivity, and

non-inductive current to model the current profile evolution in low confinement mode (L-mode) discharges in DIII-D. The dominant input nonlinearities of the model are retained by choosing a set of virtual actuators, defined by nonlinear transformations of the physical actuators. The nonlinear partial differential equation (PDE) describing the current profile evolution is then discretized in space using a finite difference method, reducing the system to a high-dimensional bilinear time-varying model with nonlinear actuation. An optimal control technique with integral action is used to synthesize a control law for the nonlinear virtual actuators based on a linear-time-invariant approximation of the state dynamics that neglects the bilinear effect. While the progressive simplification of the model used to facilitate feedback controller design could be avoided by implementing more complex control laws, the successful simulation and experimental results presented in this work validate the approximations made to some extent. We emphasize that our approach is one of the first demonstrations of closed-loop current profile control and, in practice, the ideal trade-off of design complexity and performance will depend on many factors, possibly requiring an iterative design process. Using the plasma current, NBI heating power, and plasma density as actuators, the optimal control law minimizes a cost function weighting the error of the profile at several spatial locations and the actuator effort. By manipulating the weight matrices, the controller can be intuitively tuned to respond differently depending on the goals of a particular experiment or operating scenario. As part of this work, a general framework for implementing real-time feedforward control of magnetic and kinetic plasma profiles was implemented in the DIII-D plasma control system (PCS). The framework and controller design was studied in simulations and subsequently tested in a series of experiments to demonstrate the performance of the control scheme. These results are part of the first experimental campaign to demonstrate successful current profile control using a first-principles-driven model-based control design. Results of our other approaches to first-principles driven model-based feedback control design that were tested during the same campaign can be found in [22, 23].

1.4. Organization

The paper is organized as follows. In section 2, a PDE model for the current profile evolution in L-mode discharges is introduced. The model reduction process and controller design are presented in section 3. A discussion of the real-time implementation and the framework used for simulation is provided in section 4. Results of a simulation study of the proposed controller are presented in section 5, and several experimental test cases are discussed in section 6. Conclusions and plans for future work are discussed in section 7.

2. Current profile evolution model

We use ρ to represent a coordinate indexing the magnetic flux surfaces within the tokamak plasma. Any quantity constant on each surface could be chosen as the indexing variable. Here we

choose a form of mean effective radius of the magnetic surface as the variable ρ , i.e. $\pi B_{\phi,0} \rho^2 = \Phi$, where Φ is the toroidal magnetic flux and $B_{\phi,0}$ is the reference magnetic field at the geometric major radius R_0 of the tokamak. We normalize the quantity by the similarly defined mean effective minor radius of the last closed magnetic surface, ρ_b , to obtain $\hat{\rho} = \rho/\rho_b$. The safety factor, a quantity related to the toroidal current density, is given by $q(\rho, t) = -d\Phi/d\Psi(\rho, t)$, where Ψ is the poloidal magnetic flux. This expression can be written as

$$q(\hat{\rho}, t) = -\frac{B_{\phi,0} \rho_b^2 \hat{\rho}}{\partial \Psi / \partial \hat{\rho}}, \quad (1)$$

where ψ is the poloidal stream function ($\Psi = 2\pi\psi$), by noting the constant relationship between ρ and Φ , i.e. $\pi B_{\phi,0} \rho^2 = \Phi$, and the definition of ρ_b . As the safety factor depends inversely on the spatial derivative of the poloidal flux, we define

$$\theta(\hat{\rho}, t) = \frac{\partial \psi}{\partial \hat{\rho}}(\hat{\rho}, t), \quad (2)$$

and take this quantity as the to-be-controlled variable.

To obtain a PDE describing the evolution of $\theta(\hat{\rho}, t)$, we start from the well known magnetic diffusion equation [13, 24, 25], which describes the poloidal magnetic flux evolution. This equation is given by

$$\frac{\partial \psi}{\partial t} = \frac{\eta(T_e)}{\mu_0 \rho_b^2 \hat{F}^2} \frac{1}{\hat{\rho}} \frac{\partial}{\partial \hat{\rho}} \left(\hat{\rho} \hat{F} \hat{G} \hat{H} \frac{\partial \psi}{\partial \hat{\rho}} \right) + R_0 \hat{H} \eta(T_e) \frac{\langle \bar{j}_{\text{NI}} \cdot \bar{B} \rangle}{B_{\phi,0}}, \quad (3)$$

where ψ represents the poloidal stream function, t is time, η is the plasma resistivity, which is dependent on the electron temperature, T_e , μ_0 is the vacuum permeability, \bar{j}_{NI} is the non-inductive current density (from sources such as NBI, ECCD, etc), \bar{B} is the toroidal magnetic field, and $\langle \rangle$ denotes the flux-surface average of a quantity. \hat{F} , \hat{G} and \hat{H} are spatially varying geometric factors of the DIII-D tokamak that are described in [13]. The boundary conditions are given by

$$\frac{\partial \psi}{\partial \hat{\rho}} \Big|_{\hat{\rho}=0} = 0, \quad \frac{\partial \psi}{\partial \hat{\rho}} \Big|_{\hat{\rho}=1} = -\frac{\mu_0}{2\pi} \frac{R_0}{\hat{G} \Big|_{\hat{\rho}=1} \hat{H} \Big|_{\hat{\rho}=1}} I(t), \quad (4)$$

where $I(t)$ is the total plasma current.

Based on experimental observations at DIII-D, simplified scenario-oriented empirical models for the electron temperature, non-inductive current density, and plasma resistivity during L-mode discharges were identified [13]. The model for the electron temperature is given by

$$T_e(\hat{\rho}, t) = k_{T_e} T_e^{\text{profile}}(\hat{\rho}) \frac{I(t) \sqrt{P_{\text{tot}}(t)}}{\bar{n}(t)}, \quad (5)$$

where k_{T_e} is a constant, $T_e^{\text{profile}}(\hat{\rho})$ is a reference profile, $P_{\text{tot}}(t)$ is the total average NBI power and $\bar{n}(t)$ is the line averaged plasma density. The model for the non-inductive toroidal current density is given by

$$\frac{\langle \bar{j}_{\text{NI}} \cdot \bar{B} \rangle}{B_{\phi,0}} = k_{\text{NI}} j_{\text{NI}}^{\text{profile}}(\hat{\rho}) \frac{I(t)^{1/2} P_{\text{tot}}(t)^{5/4}}{\bar{n}(t)^{3/2}}, \quad (6)$$

where k_{NI} is a constant and $j_{\text{NI}}^{\text{profile}}(\hat{\rho})$ is a reference profile for the non-inductive current deposition. Since the plasma current is mainly driven by induction during L-mode discharges, the effects of the bootstrap current on the poloidal magnetic flux are neglected in this model. The plasma resistivity $\eta(T_e)$ is given by

$$\eta(\hat{\rho}, t) = \frac{k_{\text{eff}} Z_{\text{eff}}}{T_e^{3/2}(\hat{\rho}, t)}, \quad (7)$$

where k_{eff} and Z_{eff} are constants.

The models (5), (6), and (7) allow us to write the magnetic diffusion equation (3) as

$$\frac{\partial \psi}{\partial t} = f_1(\hat{\rho}) u_1(t) \frac{1}{\hat{\rho}} \frac{\partial}{\partial \hat{\rho}} \left(\hat{\rho} f_4(\hat{\rho}) \frac{\partial \psi}{\partial \hat{\rho}} \right) + f_2(\hat{\rho}) u_2(t), \quad (8)$$

with boundary conditions given by

$$\frac{\partial \psi}{\partial \hat{\rho}} \Big|_{\hat{\rho}=0} = 0, \quad \frac{\partial \psi}{\partial \hat{\rho}} \Big|_{\hat{\rho}=1} = -k_3 u_3(t), \quad (9)$$

where

$$f_1(\hat{\rho}) = \frac{k_{\text{eff}} Z_{\text{eff}}}{k_{T_e}^{3/2} \mu_0 \rho_b^2 \hat{F}^2(\hat{\rho}) (T_e^{\text{profile}}(\hat{\rho}))^{3/2}}, \quad (10)$$

$$f_2(\hat{\rho}) = \frac{k_{\text{eff}} Z_{\text{eff}} R_0 k_{\text{NI}} \hat{H}(\hat{\rho}) j_{\text{NI}}^{\text{profile}}(\hat{\rho})}{k_{T_e}^{3/2} (T_e^{\text{profile}}(\hat{\rho}))^{3/2}}, \quad (11)$$

$$f_4(\hat{\rho}) = \hat{F}(\hat{\rho}) \hat{G}(\hat{\rho}) \hat{H}(\hat{\rho}), \quad (12)$$

$$k_3 = \frac{\mu_0 R_0}{2\pi \hat{G}|_{\hat{\rho}=1} \hat{H}|_{\hat{\rho}=1}},$$

$$u_1(t) = \left(\frac{\bar{n}(t)}{I(t) \sqrt{P_{\text{tot}}(t)}} \right)^{3/2}, \quad u_2(t) = \frac{\sqrt{P_{\text{tot}}(t)}}{I(t)}, \quad (13)$$

$$u_3(t) = I(t).$$

Equation (8) admits diffusivity, interior, and boundary actuators u_1 , u_2 , and u_3 , respectively, which each represent nonlinear combinations of the physical actuators, $I(t)$, $P_{\text{tot}}(t)$ and $\bar{n}(t)$. Note that the controller proposed in this work will generate waveforms for these physical actuators. These waveforms represent references to be sent to existing dedicated controllers for each of the respective quantities.

We expand (8) with the chain rule to obtain

$$\frac{\partial \psi}{\partial t} = f_1 u_1(t) \frac{1}{\hat{\rho}} \left(\hat{\rho} \frac{\partial \psi}{\partial \hat{\rho}} \frac{\partial f_4}{\partial \hat{\rho}} + f_4 \frac{\partial \psi}{\partial \hat{\rho}} + \hat{\rho} f_4 \frac{\partial^2 \psi}{\partial \hat{\rho}^2} \right) + f_2 u_2(t). \quad (14)$$

We then insert (2) into (14), resulting in

$$\frac{\partial \psi}{\partial t} = f_1 u_1 \frac{1}{\hat{\rho}} (\hat{\rho} \theta f_4' + f_4 \theta + \hat{\rho} f_4 \theta') + f_2 u_2, \quad (15)$$

where $(\cdot)' = \partial/\partial \hat{\rho}$ and the dependencies on time and space have been dropped to simplify the representation. By differentiating (15) with respect to $\hat{\rho}$, the PDE governing the evolution of $\theta(\hat{\rho}, t)$ is found to be

$$\frac{\partial \theta}{\partial t} = h_0 u_1 \theta'' + h_1 u_1 \theta' + h_2 u_1 \theta + h_3 u_2, \quad (16)$$

with boundary conditions:

$$\theta \Big|_{\hat{\rho}=0} = 0, \quad \theta \Big|_{\hat{\rho}=1} = -k_3 u_3, \quad (17)$$

and where

$$h_0 = f_1 f_4, \quad (18)$$

$$h_1 = f_1' f_4 + f_1 f_4 \frac{1}{\hat{\rho}} + 2 f_1 f_4', \quad (19)$$

$$h_2 = f_1' f_4' + f_1' f_4 \frac{1}{\hat{\rho}} + f_1 f_4' \frac{1}{\hat{\rho}}, \quad -f_1 f_4 \frac{1}{\hat{\rho}^2} + f_1 f_4'', \quad (20)$$

$$h_3 = f_2'. \quad (21)$$

3. Model reduction and controller design

3.1. Model reduction via spatial discretization

To reduce the infinite dimensional system (16) describing the profile evolution to a finite dimensional approximation, we discretize the spatial domain $\hat{\rho} = [0, 1]$ into l nodes (equally spaced at $\Delta \hat{\rho} = 1/(l-1)$) using a central difference formula for the $m = l-2$ interior nodes and forward/backward difference formulae for the left/right boundary nodes. The resulting set of ODEs can be expressed as

$$\dot{\alpha}(t) = M \alpha(t) v_1(t) + N v_2(t) + Z v_3(t), \quad (22)$$

where $\alpha = [\theta_2, \dots, \theta_{l-1}]^T \in \mathbb{R}^{m \times 1}$ is the value of $\theta(\hat{\rho}, t)$ at the interior nodes, the control input vectors are given by

$$[v_1(t), v_2(t), v_3(t)]^T = [u_1(t), u_2(t), u_1(t) u_3(t)]^T \in \mathbb{R}^{3 \times 1}, \quad (23)$$

and the system matrices are $M \in \mathbb{R}^{m \times m}$, $N \in \mathbb{R}^{m \times 1}$ and $Z \in \mathbb{R}^{m \times 1}$. The elements of these matrices are given in appendix A.

We consider $v_{\text{ff}}(t) = [v_{1\text{ff}}(t), v_{2\text{ff}}(t), v_{3\text{ff}}(t)]^T$, a set of feedforward control input trajectories associated with a state trajectory $\alpha_{\text{ff}}(t)$ that are computed offline or taken from experimental measurements with a nominal initial state $\alpha_{\text{ff}}(0)$. The feedforward state trajectory satisfies

$$\dot{\alpha}_{\text{ff}}(t) = M \alpha_{\text{ff}}(t) v_{1\text{ff}}(t) + N v_{2\text{ff}}(t) + Z v_{3\text{ff}}(t). \quad (24)$$

Due to disturbances in the inputs, initial conditions, or model parameters, a deviation from the desired feedforward trajectory, $z(t) = \alpha(t) - \alpha_{\text{ff}}(t)$, may occur. We include the terms $v_{\text{fb}}(t) = [v_{1\text{fb}}(t), v_{2\text{fb}}(t), v_{3\text{fb}}(t)]$ as the feedback portion of the control inputs (for which a control law will be designed), allowing us to write

$$\dot{\alpha}_{\text{ff}}(t) + \dot{z} = M \alpha_{\text{ff}}(v_{1\text{ff}} + v_{1\text{fb}}) + M z (v_{1\text{ff}} + v_{1\text{fb}}) + N (v_{2\text{ff}} + v_{2\text{fb}}) + Z (v_{3\text{ff}} + v_{3\text{fb}}). \quad (25)$$

By substituting (24) in (25) we obtain a bilinear model. While we have preserved the dominant input nonlinearities through the use of the nonlinear transformations (13) and (23), here we approximately linearize the state dynamics by neglecting the bilinear terms. This approximation is valid assuming

$v_{1\text{ff}} \gg v_{1\text{fb}}$, which can be ensured through the choice of the control effort weighting matrices in the controller design, or by assuming $\alpha_{\text{ff}} \gg z$. Though it would indeed be possible to account for the bilinear state behavior during the control synthesis (see [18]), simulations and experiments show the closed-loop system to be robust to this approximation, such that we can avoid the extra complexity arising from the bilinear terms. The dynamics of the error can then be expressed as the time-varying system

$$\dot{z} = A(t)z + B(t)v_{\text{fb}}, \quad (26)$$

where $A(t) = Mv_{1\text{ff}}(t) \in \mathbb{R}^{m \times m}$ and $B(t) = [M\alpha_{\text{ff}}(t), N, Z]$. We consider measurement of all of the states z , i.e., the output equation is $y = Cz + Dv_{\text{fb}}$ with $C = I$ and $D = 0$.

3.2. Model reduction via singular-value decomposition

We next decouple system (26) at steady-state through singular-value decomposition of the steady-state transfer function. This reduces the underactuated, coupled system to a set of relevant control channels. We can then choose to keep only the most significant singular values, that is, the control channels through which the most influence on the system is realized, to further reduce the problem and avoid using excessive controller effort to realize small improvements in profile tracking. The details of this step are provided in appendix B.

Using the ‘significant’ bases obtained through the steady-state analysis, we can write the dynamic system (26) as

$$\begin{aligned} \dot{z} &= A_s(t)z + B_s(t)v_{\text{fb}_s}^*, \\ y_s^* &= C_s z + D_s v_{\text{fb}_s}^*, \end{aligned} \quad (27)$$

where $y_s^* = \Sigma_s^{-1}U_s^T Q^{1/2}y$ is the significant output, $A_s(t) = A(t)$, $B_s(t) = B(t)R^{-1/2}V_s$, $C_s = \Sigma_s^{-1}U_s^T Q^{1/2}$, $D_s = 0$, $Q \in \mathbb{R}^{m \times m}$ is a symmetric positive definite weighting matrix for the reference tracking error, and $R \in \mathbb{R}^{3 \times 3}$ is a positive definite weight matrix for the controller effort. Definitions of the partitions Σ_s , U_s and V_s are given in appendix B.

3.3. LQI optimal control design

At this point, we note that we have reduced the dynamic system (16) to a linear time-varying system cascaded with a nonlinear transformation of the physical inputs. A linear-quadratic-regulator (LQR) optimal control problem could be solved for the system (26) in order to minimize a cost function based on a weighted combination of the state z and the control effort v . However, the classical LQR controller is a static state feedback control law without a term proportional to the integral of the error signal, which may cause poor performance in the presence of disturbances or when tracking a constant nonzero reference r . We could improve upon the disturbance rejection and tracking performance by including a term involving the time integral of the output y in the cost function, thereby formulating a linear-quadratic-integral (LQI) control problem. However, since the system (26) has multiple (3) inputs and many (m , the number of spatial discretization nodes) outputs,

such a design would result in a high order control law. Also, since the system is underactuated, the resulting controller could tend to integrate components of the output that are difficult to control, leading to the use of excessive amounts of control effort without achieving a worthwhile improvement in the system performance. By exploiting the previous model reduction step, we can instead design an LQI controller for the reduced system (27), which has s (the number of significant singular values) inputs and outputs. This reduces the order of the controller and avoids the problem of integrating outputs that are very difficult to control.

To solve the LQI optimal control design problem, we define an augmented state vector that includes the integral over time of the significant output y_s^*

$$x = \begin{bmatrix} \int_0^t y_s^* dt' \\ z \end{bmatrix}, \quad (28)$$

and proceed with a classic LQR design, i.e. we state the optimal control design problem as

$$\min_{v_{\text{fb}}} J_+ = \frac{1}{2} \int_{t_0}^{\infty} \left[x^T Q_+ x + (v_{\text{fb}_s}^*)^T R_+ v_{\text{fb}_s}^* \right] dt, \quad (29)$$

where $Q_+ \in \mathbb{R}^{m+s \times m+s}$ is a symmetric positive semi-definite matrix and R_+ is a positive scalar. The optimal control law, detailed in appendix C, can be written in state-space form as

$$\begin{aligned} \dot{x}_c &= A_c x_c + B_c z, \\ v_{\text{fb}} &= C_c(t)x_c + D_c(t)z, \end{aligned} \quad (30)$$

where x_c is the controller state representing the integral of the significant output, $A_c = 0$, $B_c = \Sigma_s^{-1}U_s^T Q^{1/2}$, $C_c(t) = -R^{-1/2}V_s K_I(t)$ and $D_c(t) = -R^{-1/2}V_s K_P(t)$.

The control design is illustrated for clarity in figure 1 with the addition of the optional nonzero reference r . The profile error e and the integral over time of the significant error e_s^* are input to the LQI control law. Note that by updating the feedforward trajectories and reference profiles used to develop the first-principles-driven model, the LQI control law can easily be recalculated for different target profiles and plasma scenarios. In contrast, a control design based on a data-driven modeling approach could require additional dedicated identification experiments to identify an updated model and retune the control law. Once v_{fb} is calculated by the controller, the values are added to the feedforward values v_{ff} . The result, $v = v_{\text{ff}} + v_{\text{fb}}$, is used to determine the required values of the physical actuators through the inverse of the transformations (13) and (23), i.e.

$$I = \frac{v_3}{v_1}, \quad P_{\text{tot}} = \left(\frac{v_2 v_3}{v_1} \right)^2, \quad \bar{n} = \frac{v_2 v_3^2}{v_1^{4/3}}. \quad (31)$$

These nonlinear transformations preserve an important part of the system dynamics that is neglected in linear-data-driven approaches.

Finally, because we have chosen an infinite time horizon for the optimal control problem and the system matrices A_+ and B_+ remain approximately time-invariant after the short

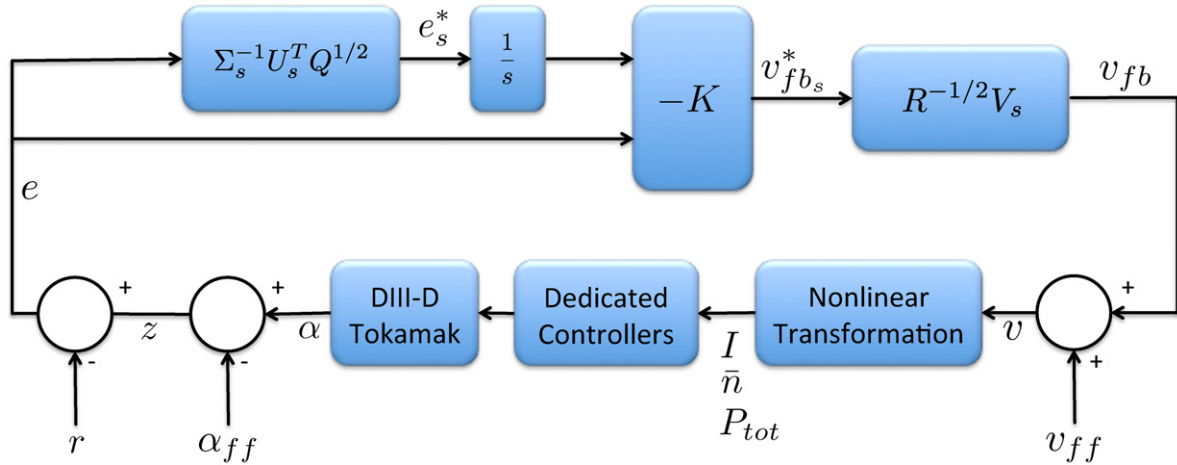


Figure 1. LQI control scheme.

ramp-up phase of the discharge, the controller matrices will be approximately constant for most of the discharge. This approximation, given in appendix C, significantly simplifies the implementation of the controller without significantly impacting performance.

4. Controller implementation in the DIII-D PCS

In this section, we present the real-time control algorithm implemented in the DIII-D PCS. An overview of the framework for real-time magnetic and kinetic profile control is provided in section (4.1). In section (4.2), we describe the simulation framework, referred to as a Simserver, used to test the control design and real-time implementation on the PCS prior to experiments.

4.1. Real-time algorithm

We have implemented a general framework for real-time feedforward+feedback control of both magnetic and kinetic profiles in the DIII-D PCS, allowing us to test this and other profile control designs. The algorithm is designed to be general enough to allow a variety of control laws to be tested within the same framework. Through a simple set of menu selections, an operator can choose which parameters to control, load a particular set of feedforward actuator trajectories, and select a feedback control law. Magnetic profiles that can be controlled include the safety factor, q , the rotational transform, $\iota = 1/q$, the poloidal magnetic flux, ψ , and the poloidal flux gradient, θ . Available kinetic profiles include the electron temperature, the ion temperature, and the toroidal rotation velocity. The framework also enables testing controllers for scalar quantities, including the normalized plasma beta β_N , the minimum q , or the internal inductance of the plasma. The feedback portion of the algorithm is implemented as a discrete time state-space system with a selectable sampling time. A sampling time of 20 ms was used in this work based on the modulation of the MSE (motional Stark effect) beam used to obtain q profile measurements in real-time. In this case the beam was on for 10 ms and off for 10 ms. The real-time magnetic equilibrium reconstruction code rtEFIT [26]

and the real-time charge-exchange recombination code rtCER [27] interface with the algorithm to provide measurements of magnetic and kinetic profiles. Since the controller presented in this work is designed solely for poloidal magnetic flux gradient profile control, only the diagnostics from rtEFIT were needed for feedback. The real-time algorithm performs the necessary coordinate transformation to construct the variable of interest α (recall that α is the vector of measurements of the θ profile) from the data provided by rtEFIT, i.e. the plasma current $I(t)$, the poloidal stream function at the magnetic axis ψ_{axis} and the plasma boundary ψ_{bdry} , and the safety factor q at 64 evenly spaced points in the normalized flux spatial domain $\psi_n = (\psi - \psi_{\text{axis}})/(\psi_{\text{bdry}} - \psi_{\text{axis}})$. Within the algorithm, the discrete time approximation of system (30) produces the output v_{fb} , which is added to the feedforward signal, v_{ff} , before being passed to the nonlinear transformations (31). The resulting output, $I(t)$, $P_{\text{tot}}(t)$ and $\bar{n}(t)$ is then sent to existing controllers for the respective quantities. The control framework is designed to allow us to introduce artificial disturbances through the signal v_d , and to specify feedforward and target trajectories through the signals α_{ff} and r . Because actuator saturation could cause the integral component of the feedback controller to wind up and possibly degrade system performance, an anti-windup compensator is included in the PCS implementation. This anti-windup scheme, which can be designed for both LTI and LTV systems, adds an additional input signal to the feedback portion of the controller. When there is no actuator saturation, the anti-windup compensator leaves the nominal closed-loop system unmodified. See [28] for an example of the anti-windup compensator design employed. A block diagram representing the PCS implementation is given in figure 2.

4.2. Simserver architecture

The simserver architecture is a simulation environment that enables real-time control algorithms implemented in the DIII-D PCS to exchange data with a Matlab/Simulink model created to generate simulated diagnostics. The Simulink model must utilize inputs and outputs that are consistent with the input and output channels of the PCS. This type of simulation framework enables tests of the real-time code

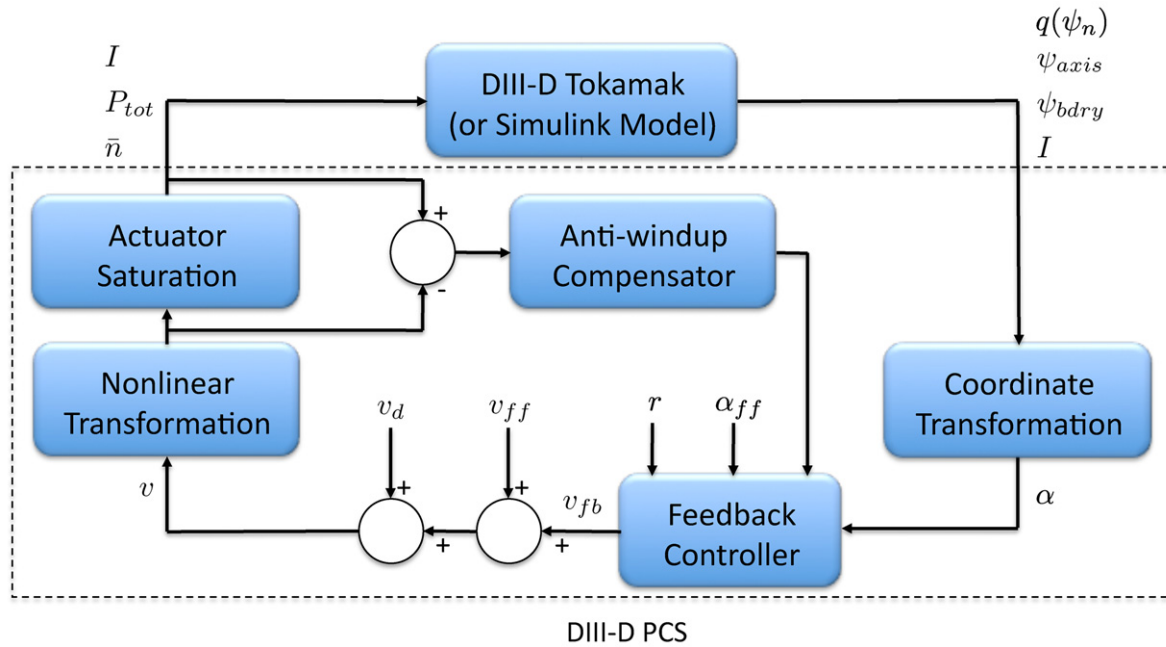


Figure 2. DIII-D PCS implementation of magnetic profile control algorithm. Note that the PCS code can either be connected to the DIII-D tokamak for experiments or, through the simserver architecture, a Simulink model of the magnetic diffusion equation for simulation tests.

and assessment of the effectiveness of control designs prior to conducting experimental tests [29]. Once the simulation phase is completed the PCS code can be used to control the actual device.

To test the proposed control scheme, a Simulink model of the magnetic diffusion equation (3) was developed and integrated into a simserver and the real-time implementation of the controller was programmed into the PCS [30]. In order to construct the Simulink model, the magnetic diffusion equation was discretized in space, resulting in a coupled set of ordinary differential equations that were evolved in time. Since the real-time EFIT (rtEFIT) reconstruction code is used to provide the profile measurements to the controller [26], the Simulink model was created to output the same set of measurements that is provided to the control algorithm by rtEFIT.

5. Simulation results

In this section, we present results of one part of the simulation study used to test and tune the controller design and implementation prior to experimental testing. The simulations presented here demonstrate the ability of the controller to track a desired profile evolution for which the necessary feedforward actuator trajectories are not well known. It is important to note that, due to the underactuated nature of the problem, not all arbitrary profiles can be achieved. To obtain an achievable target evolution, we first simulated the system using a particular set of feedforward inputs and saved the resulting θ profile evolution to be used as the target for subsequent simulated shots. We then programmed a different set of feedforward inputs (i.e. the controller is not provided with the set of feedforward inputs used to produce the desired target profile evolution) and ran two shots: one without feedback (open

loop) and one with feedback (closed-loop). During the closed loop shot, the controller was turned on from 0.5 to 2.5 s, then turned off from 2.5 to 3.5 s to see how the uncontrolled system evolved. From 3.5 s to the end of the simulation, the controller was turned back on to see if it could recover the desired target profile. In this study, we utilized the LTI approximation of the optimal control law (30) and kept one singular value in the model reduction process since, for the particular weight matrices used, $\sigma_1 \gg \sigma_2 > \sigma_3$.

Time traces of the safety factor q , which is related to θ by the relationship (1), are given in figure 3 for several spatial locations. The results of the open-loop shot and the closed-loop shot are compared with the target. The controller successfully drove the closed-loop traces toward the target traces during the first phase of the shot ($t = 0.5$ to $t = 2.5$ s). When the controller was turned off at $t = 2.5$ s, the closed-loop traces moved away from the target and back toward the traces achieved in the open-loop shot. Finally, when the controller was turned back on at $t = 3.5$ s, the closed-loop traces were driven toward the target traces once again. In figure 4, the q and θ profiles achieved in the closed-loop shot and the open-loop shot are compared with the target profiles at several times. Figure 4(a) shows the controller's progress toward achieving the target at the start of the flattop phase, while figure 4(b) shows that the closed-loop system nearly achieved the desired profiles by the time it was turned off at $t = 2.5$ s. Figure 4(c) shows the error in the profile that resulted after the uncontrolled drift phase ($t = 2.5$ s to $t = 3.5$ s). Progress toward recovery of the desired profile after the controller was turned back on is visible in figure 4(d). Finally, the actuator values during the closed-loop and open-loop shots are compared in figure 5. It can be seen how, in order to track the target profile evolution and reject the disturbance in the initial θ profile at the start of each controlled portion of the shot, the feedback

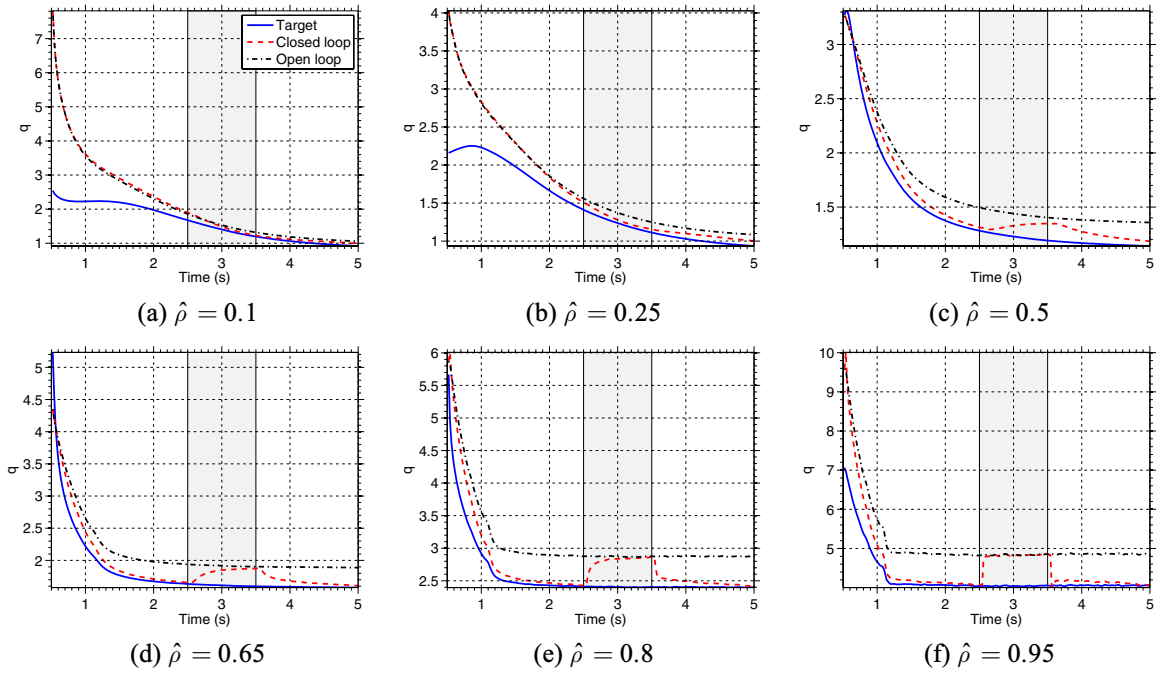


Figure 3. Time trace of q at various points during the simulation test comparing the target simulation (blue solid), the open-loop simulation (black dashed–dotted), and the closed-loop simulation (red dashed). Note the effect of turning off the controller during the closed-loop shot between $t = 2.5$ s and $t = 3.5$ s (shaded regions of the plots).

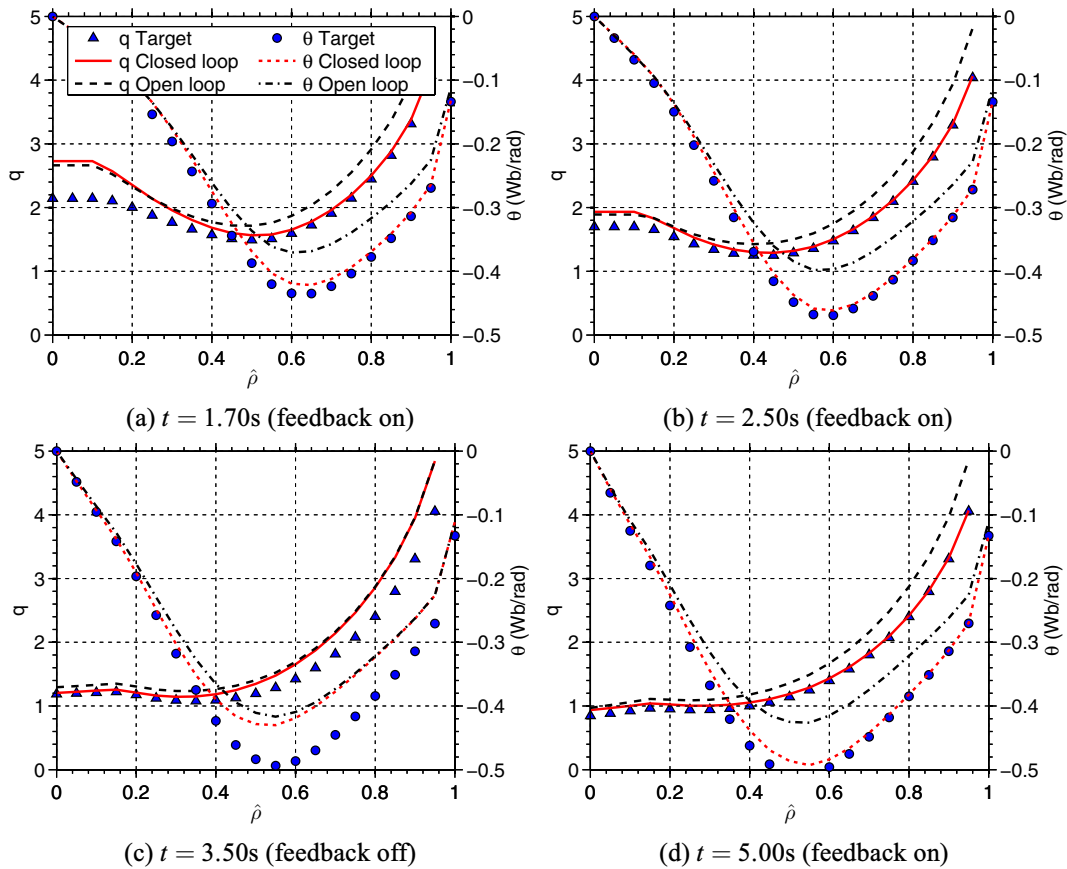


Figure 4. Comparison of q and θ profiles at various times for the target simulation (q : blue triangular markers, θ : blue circular markers), the open-loop simulation (q : black dashed, θ : black dashed–dotted), and the closed-loop simulation (q : red solid, θ : red dotted). Progress toward the target profiles is seen in (a), while the target is nearly achieved in (b). The effect of turning off the controller can be noted in (c). Recovery of the target profiles after the controller is turned back on can be observed in (d).

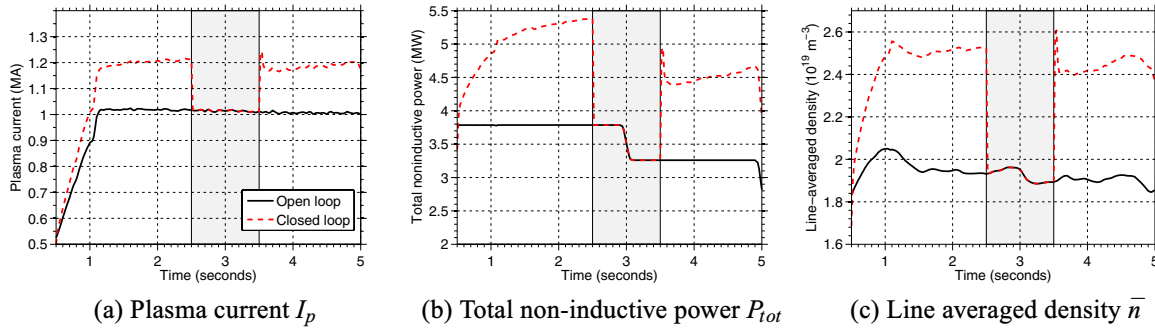


Figure 5. Comparison of actuator values during the open-loop and closed-loop simulations. During the closed-loop simulation, the feedback control was turned off between $t = 2.5$ s and $t = 3.5$ s (shaded regions of the plots).

component of the control scheme modified the feedforward actuator trajectories.

6. Experimental results

In this section, we present experimental results showing the controller's performance on the DIII-D device. As noted in the previous section, due to actuator saturation limits and the limited number of actuators, it is not possible to achieve any arbitrary target profile evolution. This limitation is present regardless of what type of profile control strategy is used. Since the goal of these experiments was to test the design and implementation of a particular feedback control law, we began the experiments by establishing two achievable θ profile evolutions. This was done by performing two open loop (feedforward only) discharges (shot numbers 145477 and 146411) using different actuator trajectories. The resulting θ profile evolutions were used to generate targets for the closed-loop (feedforward + feedback) shots (shot numbers 146413, 146414, and 146460). The first two closed-loop shots utilized the results of the open-loop shots as targets and test the controller's reference tracking and disturbance rejection capabilities, respectively. In the third test, the results of one of the open-loop shots were modified to create a target profile that is not necessarily achievable, representing a much more challenging problem for the control scheme. As we did in the simulation results, we utilized the LTI approximation of the optimal control law (30) and kept one singular value in the model reduction process since, for the particular weight matrices used, $\sigma_1 \gg \sigma_2 > \sigma_3$.

6.1. Reference tracking

As a test of the reference tracking capability of the controller, we used the feedforward inputs for shot 146411 with the θ profile evolution of shot 145477 as the target profile. During the closed-loop shot, the controller was turned on from 0.5 to 3.0 s, then turned off from 3.0 to 3.5 s to see how the uncontrolled system would evolve. From 3.5 s to the end of the shot, the controller was turned back on to see if it could recover the desired target profile despite the large error caused by the uncontrolled phase of the discharge.

Time traces of q at several points along the profile are given in figure 6. The results of the open-loop shot 146411 and the

closed-loop shot 146413 are compared with the target, which was generated in the open-loop shot 145477. The controller successfully drove the closed-loop traces toward the target traces during the first phase of the shot ($t = 0.5$ s to $t = 3.0$ s). When the controller was turned off at $t = 3.0$ s, the measured values of q in shot 146413 moved away from the target and back toward the values achieved in the open-loop shot 146411. When the controller was turned back on at $t = 3.5$ s, the closed-loop measurements were driven toward the desired target traces once again. In figure 7, the q and θ profiles achieved in the closed-loop shot 146413 and the open-loop shot 146411 are compared with the target profiles obtained from shot 145477 at several times. Figure 7(a) shows the controller's progress in achieving the target at the start of the flattop phase, while figure 7(b) shows that the closed-loop system nearly achieved the desired profiles by the time it was turned off at $t = 3.0$ s. Figure 7(b) shows the error in the profiles that resulted after the uncontrolled drift phase ($t = 3.0$ to $t = 3.5$ s). Progress toward recovery of the desired profiles after the controller was turned back on is clear in figure 7(c). Finally, the actuator requests and achieved values are compared in figure 8. It should be noted that while the total plasma current was tightly controlled and the requests were reproduced quite well, the requests for total power and line averaged density were often not very closely followed. Despite this additional disturbance, good tracking results were still achieved.

6.2. Disturbance rejection

In this section, experimental results showing the disturbance rejection capability of the proposed control scheme are presented. We used the θ profile evolution measured during the open-loop shot 145477 as the target for the closed-loop shot 146414. During shot 146414, an artificial input disturbance of -0.15 MA in the plasma current and -0.5 MW in the total non-inductive heating power was added to the reference inputs (taken from 145477) from $t = 0.5$ to $t = 2.5$ s. The feedback controller was turned on from $t = 0.5$ to $t = 2.0$ s to test disturbance rejection and switched off from $t = 2.0$ to $t = 2.5$ s to allow the θ profile to drift away from the desired one under the influence of the input disturbance. At $t = 2.5$ s the disturbance was changed to 0.15 MA in the plasma current and 0.5 MW in the heating power and the controller was turned back on to see if the controller could recover the desired profile

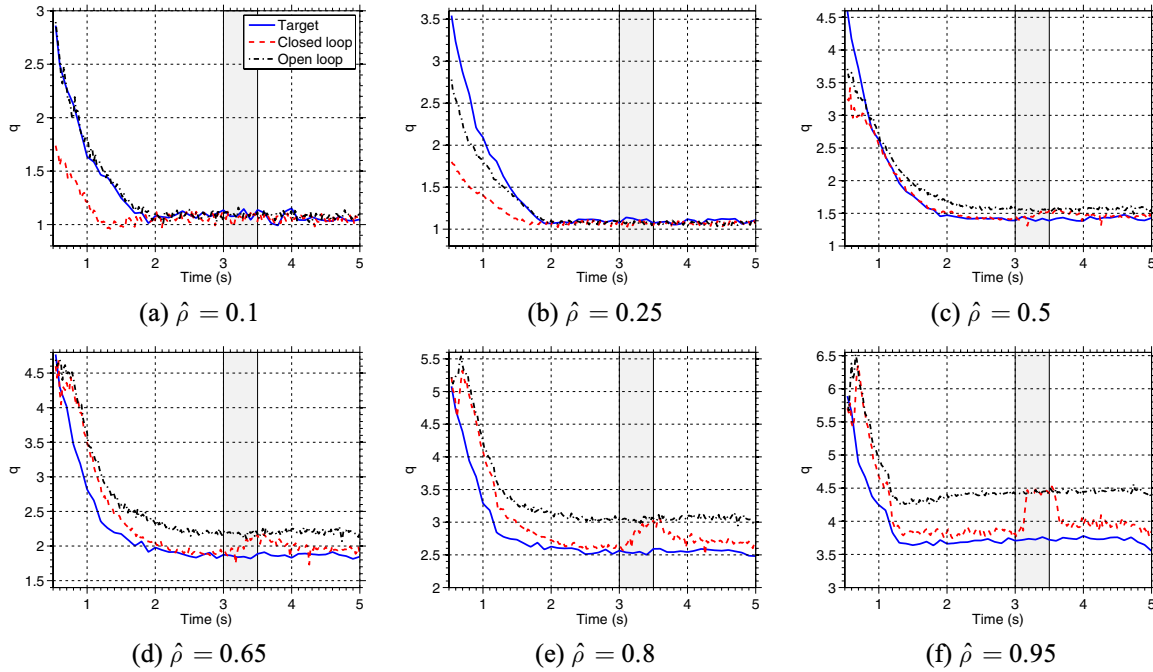


Figure 6. Time trace of q at various spatial locations during the reference tracking experiment. Plots compare the target shot 145477 (blue solid), the open-loop shot 146411 (black dashed–dotted), and the closed-loop shot 146413 (red dashed). Note the effect of turning off the controller during the closed-loop shot between $t = 3$ and $t = 3.5$ s (shaded regions of the plots).

despite the error caused by both the drift and the presence of the new disturbance. At $t = 4.0$ s, the controller was switched off to observe the effect of the uncontrolled disturbance on the θ profile.

Time traces of q at several points along the profile are given in figure 9. The results of the closed-loop shot 146414 are compared with the target generated in the feedforward shot 145477. The controller successfully rejected the disturbance during the first phase ($t = 0.5$ to $t = 2.0$ s) and the effect of the uncontrolled disturbance was apparent during the drift phase ($t = 2.0$ to $t = 2.5$ s). When the disturbance was changed and the controller was turned back on at $t = 2.5$ s, the target values of q were quickly recovered. Finally, the uncontrolled response of the system to the second disturbance can be seen after $t = 4.0$ s. In figure 10, the q and θ profiles achieved in the closed loop, disturbed shot 146414 are compared with the desired reference profiles obtained from shot 145477 at several times. Figure 10(a) shows that the controller mostly rejected the disturbance and recovered the desired profiles shortly before it was turned off at $t = 2.0$ s. Figure 10(b) shows the error that resulted from the disturbance after the uncontrolled drift phase ($t = 2.0$ to $t = 2.5$ s) and the successful recovery of the desired profiles after the controller was turned back on for a short time can be seen in figure 10(c). The error caused by the second disturbance after the controller was switched off is visible in figure 10(d). Finally, the actuator requests and achieved values are compared in figure 11. It should be noted that while the total plasma current and total power were tightly controlled and the requests were reproduced quite well, the request for line averaged density was not followed as closely. This represented an additional input disturbance aside from the one intentionally added to the feedforward input references.

The closed-loop system appeared to be robust to the loose control of the density.

6.3. Frozen target tracking

In the previous sections, the results of an open-loop shot were used as the target for the closed-loop shots. This guaranteed that the target was a physically achievable θ profile evolution. In this section, we again use the results of an open-loop shot as the target, however, for a short time interval, we hold the target profile constant, rather than allowing it to evolve as it did in open loop. By freezing the target, the controller must attempt to force the system to evolve in a manner that was not observed in any open-loop experiments and may not even be achievable given the fact that the system is underactuated and that the actuators have magnitude constraints. For the experiment presented here, the θ profile evolution from the feedforward shot 145477 was used to generate the target. The target was made identical to the measured results of 145477 except that the target profile was frozen at $t = 1.5$ s and held constant until $t = 4.0$ s. After $t = 4.0$ s, the unmodified results of shot 145477 were utilized as the target profile. In the closed-loop shot, the controller was active until $t = 3.0$ s then left off until $t = 4.0$ s, when it was turned back on until the end of the shot.

In appendix B, through analysis of the singular-value decomposition of the steady-state transfer function of the error system (26), it is shown that the reference r can be written as the sum of a component r_t that is trackable at steady-state and a non-trackable component r_{nt} . By adding $r_{t,s}$ (the component of the trackable reference associated with the significant singular values) to the feedforward profile evolution, we can obtain a

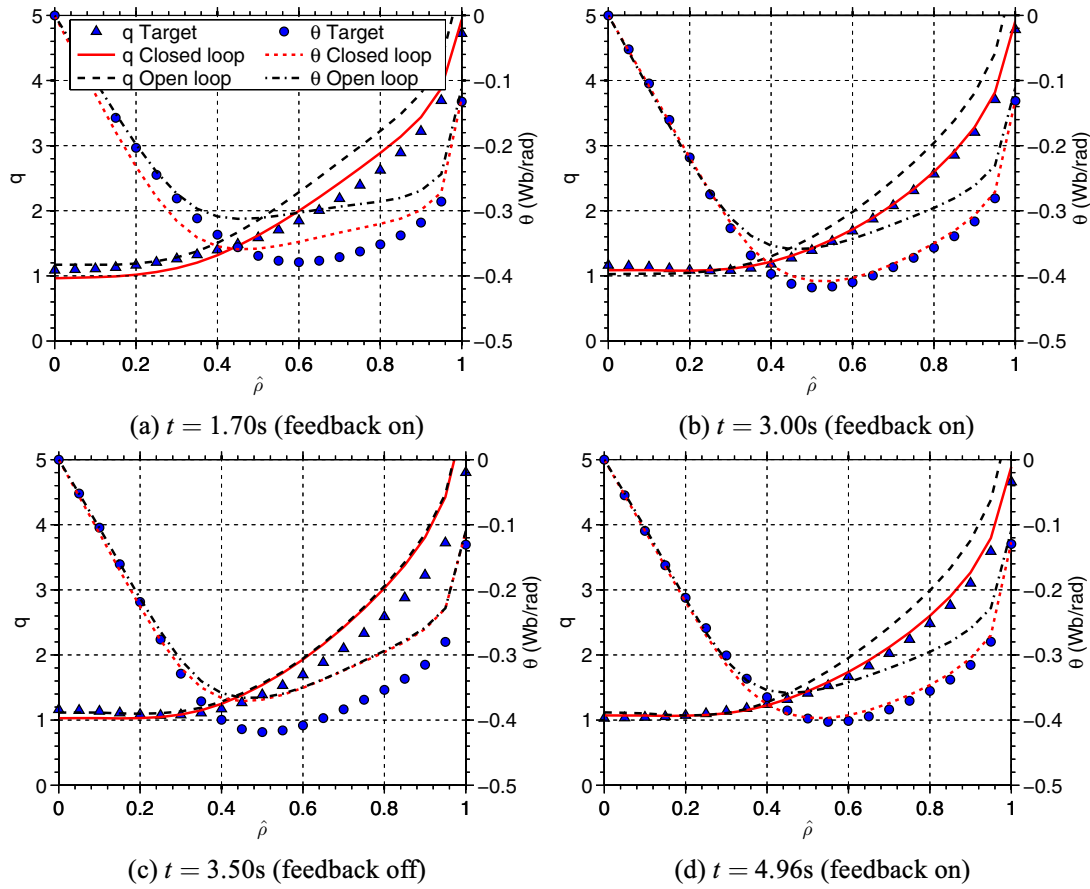


Figure 7. Profiles achieved during the reference tracking experiment. Comparison of q and θ profiles at various times for target shot 145477 (q : blue triangular markers, θ : blue circular markers), the open-loop shot 146411 (q : black dashed, θ : black dashed-dotted), and the closed-loop shot 146413 (q : red solid, θ : red dotted). Progress toward the target profiles is seen in (a), while the target is nearly achieved in (b). The effect of turning off the controller can be noted in (c). Recovery of the target profiles after the controller is turned back on can be observed in (d).

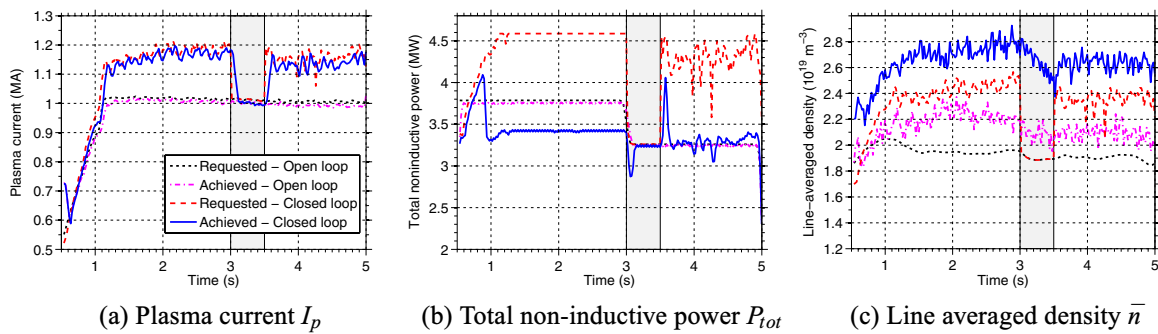


Figure 8. Requested and achieved actuator values during the reference tracking experiment. Comparison of the open-loop shot 145477 and the closed loop, disturbed shot 146413. During the closed-loop shot, the feedback control was turned off between $t = 3.0$ s and $t = 3.5$ s (shaded regions of plots).

representation of the significant, trackable component of the target profile, as predicted by the model, i.e.

$$\alpha_{\text{trackable},s} = r_{t,s} + \alpha_{\text{ff}}. \quad (32)$$

For a target profile not generated directly from experimental or simulation results, like the frozen target used in this scenario, the trackable component of the target may differ greatly from the prescribed target. The trackable target $\alpha_{\text{trackable},s}$

can therefore be very useful in interpreting the results of experiments of this type.

Time traces of q at several points along the profile are given in figure 12. The results of the closed-loop shot 146460 and the open-loop shot 145477 are compared with the prescribed target, as well as the q associated with the trackable target, as defined in (32). During the interval in which the target was frozen, the controller increased the value of q at the edge of the plasma (see figure 12(f)). The effect of this increase

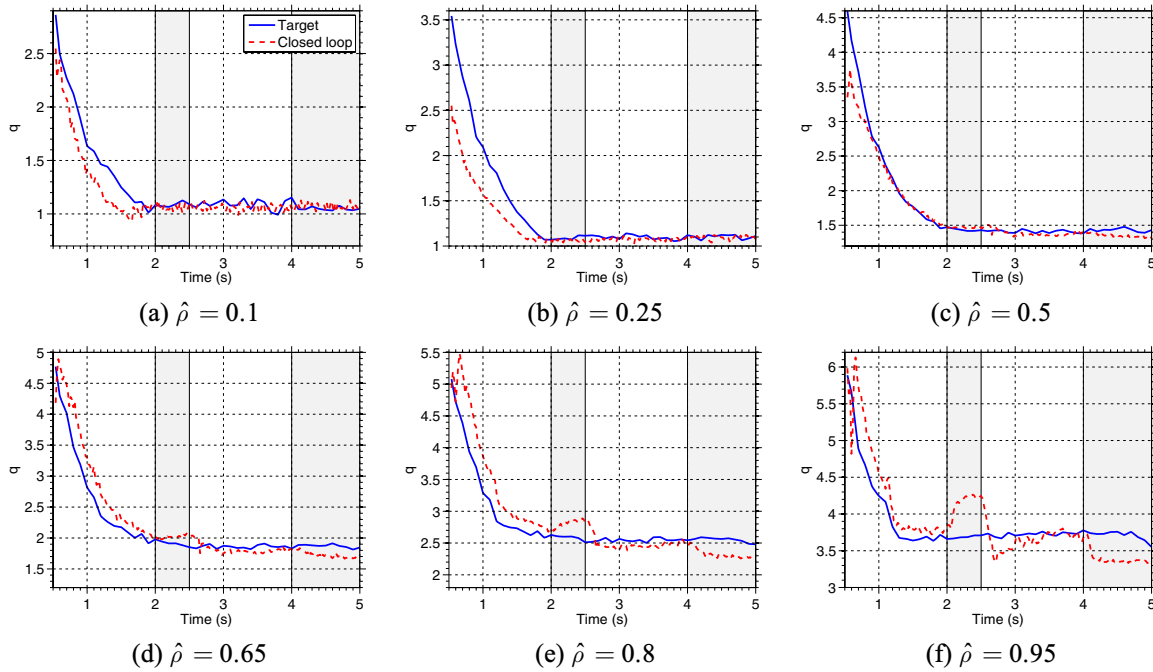


Figure 9. Time trace of q at various spatial locations during the disturbance rejection experiment comparing the reference shot 145477 (blue solid) and the closed loop, disturbed shot 146414 (red dashed). Note the effect of turning off the controller when $2.0 \text{ s} \leq t \leq 2.5 \text{ s}$ and $t \geq 4.0 \text{ s}$ (shaded regions of the plots).

diffused inward from the boundary and brought the interior points of the profile closer to the trackable component of the target at the diffusive time-scale, as seen in figure 12(d) and the controller would have likely achieved most of the trackable reference had the test been continued for a longer duration. When the controller was turned off at $t = 3.0 \text{ s}$, the progress toward achieving the target was lost and the profile relaxed back toward the results of the open-loop shot. After $t = 4.0 \text{ s}$, feedback was turned back on and the controller regulated the profile around the target. In figures 13 and 14, respectively, the q and θ profiles achieved in the open-loop shot 145477 and the closed-loop shot 146414 are compared with both the desired and trackable target profiles at several times. Note that there is a significant difference between the prescribed target and the trackable target, indicating that the prescribed target did not lie within the significant output basis. Figures 13(a)–(c) and 14(a)–(c) show how the controller progressed toward achieving the trackable component of the target profile over time. Figures 13(d) and 14(d) shows that the profiles returned to the open-loop results after the controller was turned off. Finally, the actuator requests and achieved values are compared in figure 15. Note that, rather than using the available actuators excessively in a futile attempt to achieve the component of the target profile lying outside of the significant output basis, the optimal controller took into account the results of the SVD analysis performed in section 3 to track only the significant component of the target profile and minimize a weighted combination of the profile error and actuator effort. In practice, the weight matrices Q and R could be changed by operators in order to tailor the controller performance to achieve the goals of a particular experiment. For example, increased weight could be placed on the interior points of the profile

or on points close to the anticipated location of $q_{\min}(t)$ to improve tracking of a desired $q(0, t)$ or $q_{\min}(t)$ trajectory, at the expense of reduced tracking performance in the other parts of the profile. By including more actuators in the control scheme, the set of achievable profiles could be expanded and the target tracking results in experiments like this one could be improved.

7. Conclusions

We have presented a current profile controller designed using a first-principles-driven dynamic model (with minimal parameters determined from experiment) of the evolution of the poloidal magnetic flux in L-mode discharges in DIII-D. The feedback controller was designed to complement any arbitrary set of feedforward inputs and drive the spatial profile of the variable θ to the desired target profile. Through a nonlinear transformation of the inputs and spatial discretization, a finite dimensional, time-varying model for the profile error was obtained. A singular-value decomposition technique was utilized to reduce the multi-input multi-output coupled system to a set of the most relevant control channels. A linear-quadratic-integral controller was then designed for the reduced order model. The resulting feedback control law outputs, added to the feedforward values and passed through a nonlinear transformation, represent desired reference values for the total plasma current, non-inductive power, and plasma density. These references are then sent to dedicated existing controllers on the DIII-D device. A simulation study was used to test the implementation and tune the control law. Experimental results for three different tests were also presented. These results were part of the

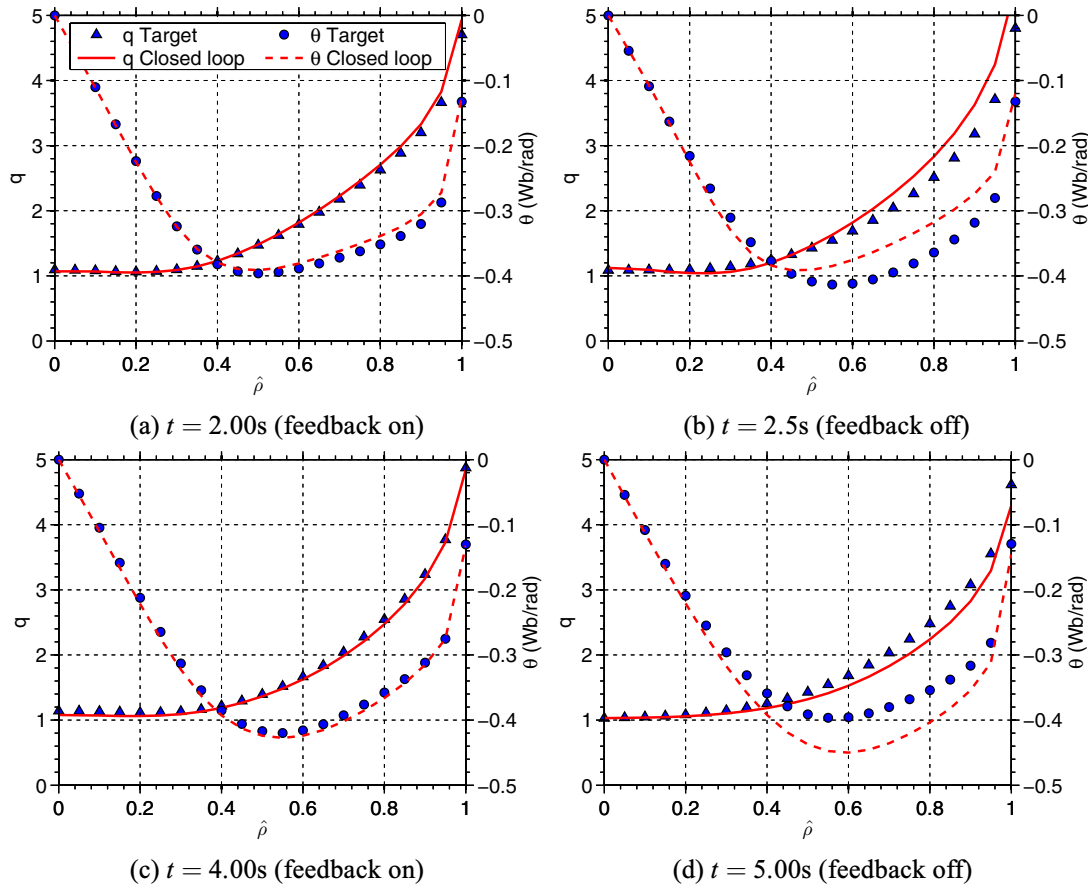


Figure 10. Profiles achieved during the disturbance rejection experiment. Comparison of q profiles at various times for reference shot 145477 (q : blue triangular markers, θ : blue circular markers) and the closed loop, disturbed shot 146414 (q : red solid, θ : red dashed). Successful disturbance rejection is seen in (a), while the effect of the uncontrolled disturbance can be noted in (b). Recovery of the target profiles after the second disturbance is applied and the controller is turned back on can be observed in (c), while the effect of the uncontrolled disturbance is seen in (d).




Figure 11. Requested and achieved actuator values during the disturbance rejection experiment. Comparison of the open-loop shot 145477 and the closed loop, disturbed shot 146414. During the closed-loop shot, the feedback control was turned off between $t = 3.0$ s and $t = 3.5$ s and after $t = 4.0$ s (shaded regions of plots).

first ever experimental campaign to use first-principles-driven model-based controllers for current profile control. In the tracking and disturbance rejection cases, the controller was shown to perform quite well despite the presence of additional disturbances caused by the physical actuators and noisy real-time measurements of the θ profile. In the third test, the controller was tasked with tracking a profile evolution that had not been produced in previous experiments and was not necessarily achievable. By taking into account the steady-

state response of the system to the available actuators, the controller drove the system to the physically achievable profile shape optimizing a weighted sum of the error measurements and actuator effort.

We note that although the model used here best describes the early, inductive phase of the plasma discharge, the controller performs well throughout the flat-top phase of the L-mode discharges studied. This is likely because self-generated non-inductive current sources are typically small in L-mode.

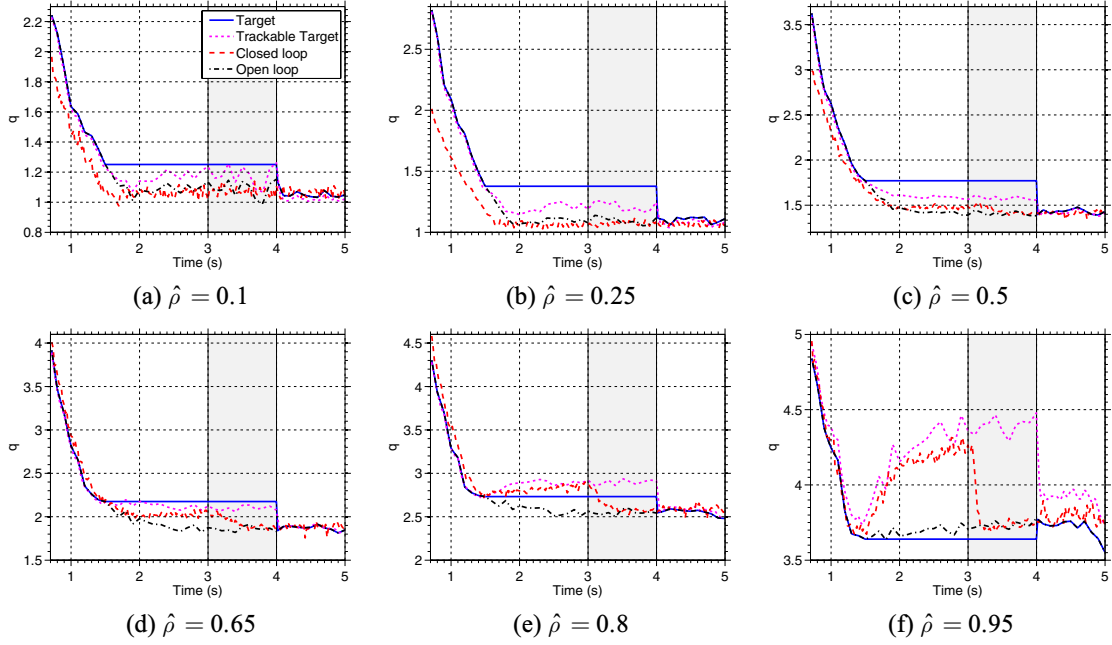


Figure 12. Time trace of q at various spatial locations during the frozen target tracking experiment comparing the open-loop shot 145477 (blue solid), the closed-loop shot 146460 (red dashed), the target (black dashed-dotted), and the steady-state trackable component of the target profile as predicted by the model (magenta dotted). Note the effect of turning off the controller in the closed-loop shot between $t = 3.0$ s and $t = 4.0$ s (shaded regions of the plots). Also note the different scale used in each plot.

Additional work will be needed to extend the model to H-mode discharges, for which the self-generated non-inductive current source neglected by the model used here becomes significant. Work toward developing a nonlinear control-oriented PDE model of the poloidal magnetic flux profile during H-mode discharges will be carried out using a similar approach to the one used in this work, with the inclusion of a bootstrap current model like the one presented in [31]. Additional degrees of freedom will be incorporated by modeling co-injection neutral beams, counter-injection neutral beams, and electron cyclotron heating separately. In doing so, controllability of the profile should be improved, expanding the set of achievable profile evolutions. The long term goal is the development of first-principles-driven model-based feedforward+feedback control strategies for simultaneous control of magnetic and kinetic plasma profiles during H-mode discharges.

Acknowledgments

This work was supported by the NSF CAREER award program (ECCS-0645086) and the US Department of Energy (DEFG02-09ER55064, DE-FG02-92ER54141 and DE-FC02-04ER54698).

Appendix A. Model reduction via spatial discretization

The elements of the system matrices for the interior node $i = 2$ are given by

$$\begin{aligned} M_{1,1} &= h_2(\Delta\hat{\rho}) - \frac{2h_0(\Delta\hat{\rho})}{(\Delta\hat{\rho})^2}, & M_{1,2} &= \frac{h_0(\Delta\hat{\rho})}{(\Delta\hat{\rho})^2} + \frac{h_1(\Delta\hat{\rho})}{2\Delta\hat{\rho}}, \\ N_1 &= h_3(\Delta\hat{\rho}), & Z_1 &= 0. \end{aligned} \quad (\text{A.1})$$

The elements of the system matrices for the interior region, $3 \leq i \leq m$, are given by

$$M_{i-1,i-2} = \frac{h_0(\Delta x)}{(\Delta\hat{\rho})^2} - \frac{h_1(\Delta x)}{2\Delta\hat{\rho}}, \quad (\text{A.2})$$

$$M_{i-1,i-1} = h_2(\Delta x) - \frac{2h_0(\Delta x)}{(\Delta\hat{\rho})^2}, \quad (\text{A.3})$$

$$M_{i-1,i} = \frac{h_0(\Delta x)}{(\Delta\hat{\rho})^2} + \frac{h_1(\Delta x)}{2\Delta\hat{\rho}}, \quad (\text{A.4})$$

$$N_{i-1} = h_3(\Delta x), \quad Z_{i-1} = 0, \quad (\text{A.5})$$

where $\Delta x = (i-1)\Delta\hat{\rho}$. The elements of the system matrices for the interior node $i = l-1$ are given by

$$M_{m,m} = h_2(\Delta x^*) - \frac{2h_0(\Delta x^*)}{(\Delta\hat{\rho})^2}, \quad (\text{A.6})$$

$$M_{m,m-1} = \frac{h_0(\Delta x^*)}{(\Delta\hat{\rho})^2} - \frac{h_1(\Delta x^*)}{2\Delta\hat{\rho}}, \quad (\text{A.7})$$

$$N_m = h_3(\Delta x^*), \quad Z_m = -k_3 \left(\frac{h_0(\Delta x^*)}{(\Delta\hat{\rho})^2} + \frac{h_1(\Delta x^*)}{2\Delta\hat{\rho}} \right), \quad (\text{A.8})$$

where $\Delta x^* = m\Delta\hat{\rho}$. The remaining entries in the M matrix are zero. Note that the values of θ at the boundary nodes $i = 1$ and $i = l$ are known from the boundary conditions (17) and are therefore not included as states of the reduced model (22).

Appendix B. Model reduction via singular-value decomposition

Noting that all of the states are assumed to be measured (i.e., $y = z$), and by assuming closed-loop stability, we can write

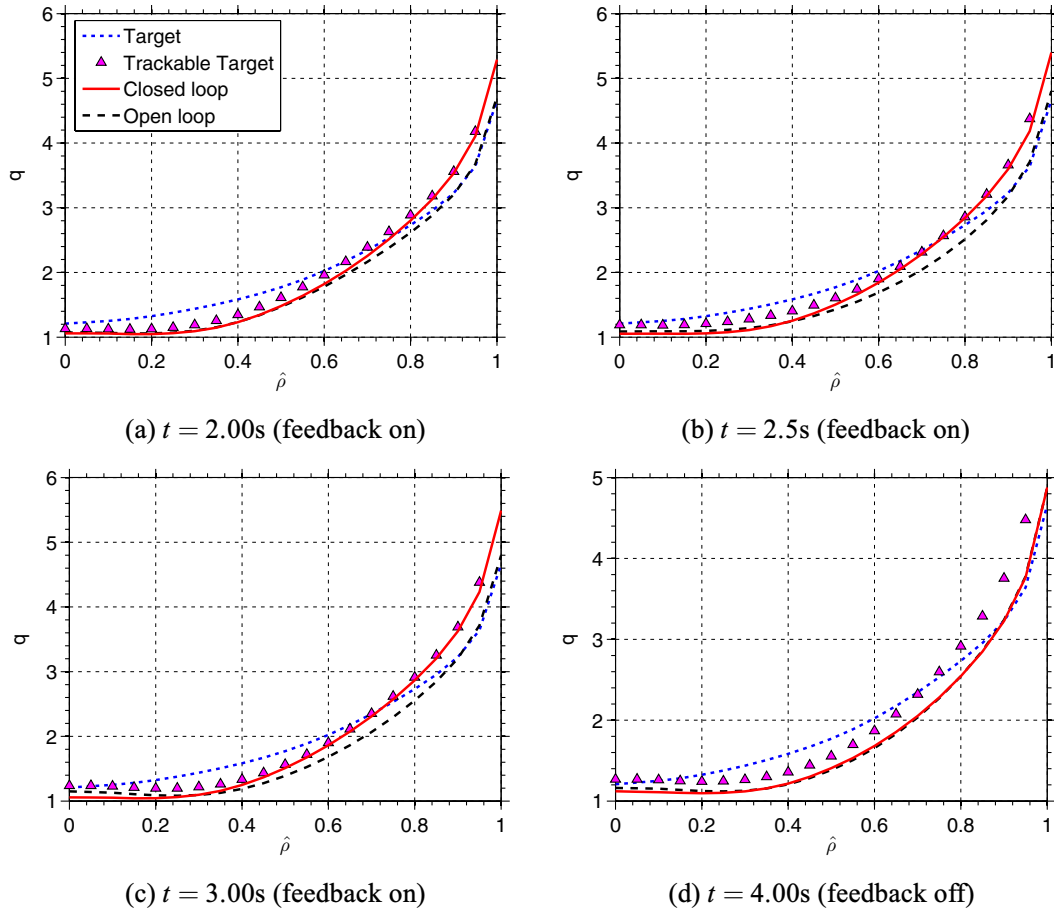


Figure 13. Profiles achieved during the frozen target tracking experiment. Comparison of q profiles at various times for open-loop shot 145477 (black dashed), the closed-loop shot 146460 (red solid), and the target (blue dotted). The steady-state trackable component of the target profile as predicted by the model is also shown (magenta triangular markers). Progress toward achieving the trackable component of the target profile is visible in (a), (b) and (c). After the controller is turned off at $t = 3.00$ s, the profile relaxes back to the profile achieved in open loop, as seen in (d).

the steady-state input-output relationship of (26) as

$$\bar{y} = \bar{z} = -\bar{A}^{-1}\bar{B}\bar{v}_{fb} = \bar{G}\bar{v}_{fb}, \quad (\text{B.1})$$

where steady-state values are denoted with an overbar. We then define the weighted steady-state transfer function

$$\tilde{G} = Q^{1/2}\bar{G}R^{-1/2} \Leftrightarrow \bar{G} = Q^{-1/2}\tilde{G}R^{1/2}, \quad (\text{B.2})$$

where $Q \in \mathbb{R}^{m \times m}$ is a symmetric positive definite weighting matrix for the tracking error and $R \in \mathbb{R}^{3 \times 3}$ is a positive definite weight matrix for the controller effort. The singular-value decomposition of (B.2) is given by $\tilde{G} = U\Sigma V^T$ where $\Sigma = \text{diag}(\sigma_1, \sigma_2, \sigma_3) \in \mathbb{R}^{3 \times 3}$ and $U \in \mathbb{R}^{m \times 3}$, and $V \in \mathbb{R}^{3 \times 3}$ are unitary matrices, i.e.,

$$U^T U = \mathbb{I}, \quad V^T V = V V^T = \mathbb{I}, \quad (\text{B.3})$$

where \mathbb{I} is the identity matrix. Using (B.2) and the singular-value decomposition of \tilde{G} , the steady-state input-output relation (B.1) can then be expressed as

$$\bar{y} = Q^{-1/2}U\Sigma V^T R^{1/2}\bar{v}_{fb}. \quad (\text{B.4})$$

Since the columns of the matrix $Q^{-1/2}U\Sigma$ define a basis for the subspace of obtainable steady-state output values, we can write

$\bar{y} = Q^{-1/2}U\Sigma\bar{y}^*$ where $\bar{y}^* \in \mathbb{R}^{3 \times 1}$. Noting this, there may be some component of the reference r that does not lie in this subspace. We therefore write the reference vector as the sum of a trackable component r_t and a non-trackable component r_{nt} , i.e. $r = r_t + r_{nt}$, where

$$r_t = Q^{-1/2}U\Sigma r^* \Leftrightarrow r^* = \Sigma^{-1}U^T Q^{1/2}r \\ (\Sigma^{-1}U^T Q^{1/2}r_{nt} = 0). \quad (\text{B.5})$$

We can also define $\bar{v}_{fb}^* = V^T R^{1/2}\bar{v}_{fb}$ and use (B.4) to show that there is a one-to-one relationship between the inputs \bar{v}_{fb}^* and the outputs \bar{y}^* , i.e.

$$\bar{y}^* = \Sigma^{-1}U^T Q^{1/2}\bar{y} \\ = \Sigma^{-1}U^T Q^{1/2}Q^{-1/2}U\Sigma V^T R^{1/2}\bar{v}_{fb} = \bar{v}_{fb}^*. \quad (\text{B.6})$$

The bases obtained through singular-value decomposition of the model for particular choices of weight matrices are shown in figure B1. The input singular vectors $V^T R^{1/2}$ show the contribution of each component of \bar{v}_{fb} to the input \bar{v}_{fb}^* , while the output singular vectors $Q^{-1/2}U\Sigma$ show contribution of each component of \bar{y}^* to \bar{y} , weighted by the associated singular value. Essentially, the three profiles in figure B1(b) represent the steady-state output that would result from a step input in each of the three components of \bar{v}_{fb}^* . Clearly, the first

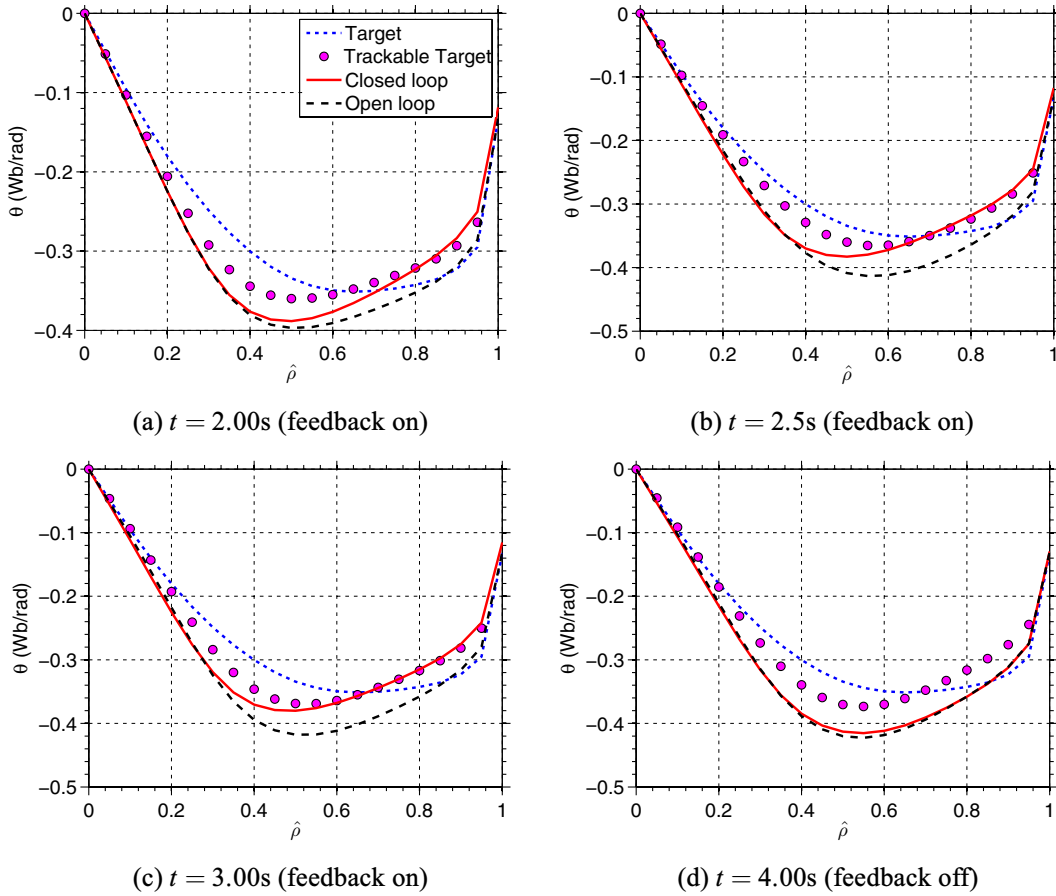


Figure 14. Profiles achieved during the frozen target tracking experiment. Comparison of θ profiles at various times for open-loop shot 145477 (black dashed), the closed-loop shot 146460 (red solid), and the target (blue dotted). The steady-state trackable component of the target profile as predicted by the model is also shown (magenta circular markers). Progress toward achieving the trackable component of the target profile is visible in (a), (b) and (c). After the controller is turned off at $t = 3.00$ s, the profile relaxes back to the profile achieved in open loop, as seen in (d).




Figure 15. Requested and achieved actuator values during the frozen target tracking experiment. Comparison of the open loop shot 145477 and the closed-loop shot 146460. During the closed-loop shot, the feedback control was turned off between $t = 3.0$ and $t = 4.0$ s (shaded regions of plots).

input can produce the largest effect on the profile, while a step input in the third component of \bar{v}_{fb}^* produces little change in the output. The first profile represents the preferred direction of the system and the dominant shape of the trackable reference r_t . The smaller magnitude of the other two profiles indicates that it takes significantly more control effort to track components of the reference in those directions.

For the model reduction process, we consider the task of minimizing the steady-state performance index

$$\bar{J} = \bar{e}^T Q \bar{e} + \bar{v}_{fb}^T R \bar{v}_{fb}, \quad (\text{B.7})$$

where $\bar{e} = \bar{y} - \bar{r} = Q^{-1/2} U \Sigma (\bar{y}^* - r^*) = Q^{-1/2} U \Sigma (\bar{e}^*)$. Since system (B.6) is a square decoupled system, the performance index can be expressed as

$$\bar{J} = (\bar{e}^*)^T \Sigma^2 \bar{e}^* + (\bar{v}_{fb}^*)^T \bar{v}_{fb}^* = \sum_{i=1}^3 \sigma_i^2 (\bar{e}_i^*)^2 + (\bar{v}_{fb_i}^*)^2. \quad (\text{B.8})$$

We note that the i th singular value acts as a weight on the i th component of the error in (B.8). Therefore, if $\sigma_i \gg \sigma_{i+1}$, the tracking error associated with σ_{i+1} will contribute little to the performance index compared to the i th component.

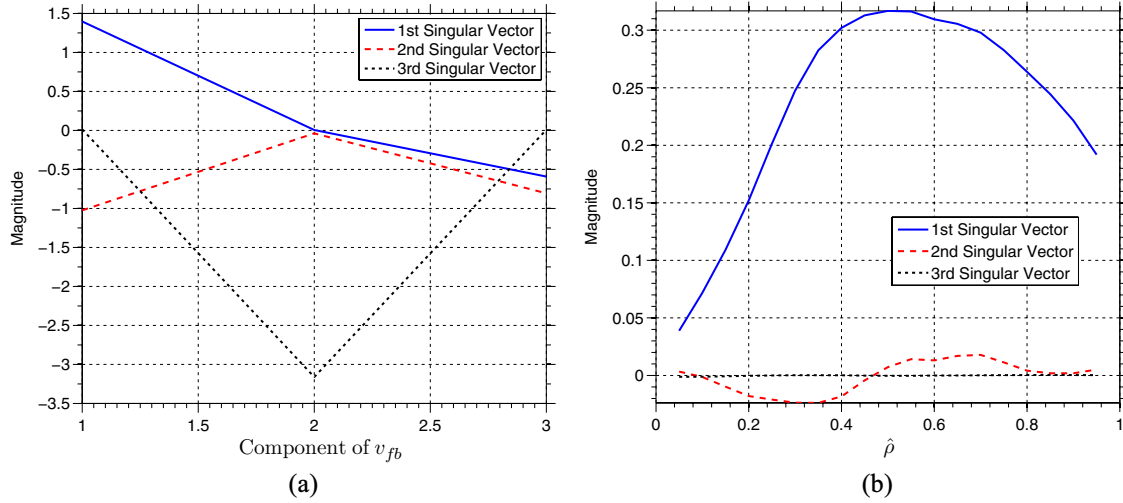


Figure B1. (a) Input singular vectors $V^T R^{1/2}$ and (b) output singular vectors $Q^{-1/2} U \Sigma$ resulting from the SVD analysis for $R = \mathbb{I}$, $Q = \text{diag}[1, 1, 1, 1, 5, 1, 1, 1, 1, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5]$.

To avoid spending excessive control effort (through v_{fb}^*) on realizing insignificant reductions in the performance index, we can separate the singular values into s significant singular values Σ_s and n negligible singular values Σ_n . We then write

$$U = [U_s \ U_n], \quad V = [V_s \ V_n], \quad \Sigma = \begin{bmatrix} \Sigma_s & 0 \\ 0 & \Sigma_n \end{bmatrix}, \quad (\text{B.9})$$

$$\bar{e}^* = [\bar{e}_s^* \ \bar{e}_n^*]^T, \quad \bar{v}_{fb}^* = [\bar{v}_{fb_s}^* \ \bar{v}_{fb_n}^*]^T,$$

where

$$\begin{aligned} \bar{e}_s^* &= \Sigma_s^{-1} U_s^T Q^{1/2} \bar{e}, & \bar{v}_{fb_s}^* &= V_s^T R^{1/2} \bar{v}_{fb} \\ \bar{e}_n^* &= \Sigma_n^{-1} U_n^T Q^{1/2} \bar{e}, & \bar{v}_{fb_n}^* &= V_n^T R^{1/2} \bar{v}_{fb} \\ \Rightarrow \bar{e} &\approx Q^{-1/2} U_s \Sigma_s \bar{e}_s^* & (\text{B.10}) \\ \bar{v}_{fb} &\approx R^{-1/2} V_s \bar{v}_{fb_s}^*. \end{aligned}$$

We can then utilize the ‘significant’ bases obtained through the steady-state analysis to approximate the cost function (B.8) as

$$\bar{J} \approx \bar{J}_s = (\bar{e}_s^*)^T \Sigma_s^2 \bar{e}_s^* + (\bar{v}_{fb_s}^*)^T \bar{v}_{fb_s}^*, \quad (\text{B.11})$$

and the dynamic system (26) as

$$\begin{aligned} \dot{z} &= A_s(t)z + B_s(t)v_{fb_s}^* \\ y_s^* &= C_s z + D_s v_{fb_s}^*, \end{aligned} \quad (\text{B.12})$$

where $A_s(t) = A(t)$, $B_s(t) = B(t)R^{-1/2}V_s$, $C_s = \Sigma_s^{-1}U_s^T Q^{1/2}$, $D_s = 0$.

Appendix C. LQI optimal controller design

Based on the definition of the state x in (28), we can write the augmented system dynamics

$$\dot{x} = A_+(t)x + B_+(t)v_{fb_s}^*, \quad (\text{C.1})$$

where $A_+(t) = \begin{bmatrix} 0^{s \times s} & C_s \\ 0^{m \times s} & A_s(t) \end{bmatrix}$ and $B_+(t) = \begin{bmatrix} D_s \\ B_s(t) \end{bmatrix}$. We then proceed with a classic linear-quadratic-regulator (LQR) design for (C.1), i.e. we state the optimal control design problem as

$$\min_{v_{fb}} J_+ = \frac{1}{2} \int_{t_0}^{\infty} \left[x^T Q_+ x + (v_{fb_s}^*)^T R_+ v_{fb_s}^* \right] dt, \quad (\text{C.2})$$

where $Q_+ \in \mathbb{R}^{m+s \times m+s}$ is a symmetric positive semi-definite matrix and R_+ is a positive scalar. The optimal control law is given by

$$v_{fb_s}^*(t) = -K(t)x(t), \quad (\text{C.3})$$

where $K(t) \in \mathbb{R}^{s \times m+s}$ is given by

$$K(t) = R_+^{-1} B_+^T(t) S_+(t), \quad (\text{C.4})$$

and $S_+(t)$ is the solution to the matrix differential Riccati equation

$$\dot{S}_+ = -S_+ A_+ - A_+^T S_+ + S_+ B_+ R_+^{-1} B_+^T S_+ - Q_+, \quad (\text{C.5})$$

subject to the condition $S_+(\infty) = 0$. The first s elements of $K(t)$ represent the integral gains, $K_I(t)$, corresponding to each of the significant outputs y_s^* , while the remaining elements are state feedback gains $K_P(t)$. We can then express the optimal control law as

$$v_{fb_s}^*(t) = -K_I(t) \int_0^t y_s^* dt' - K_P(t)z. \quad (\text{C.6})$$

We note the definitions of $v_{fb_s}^*$, y_s^* , to write the control law as

$$\begin{aligned} v_{fb} &= -R^{-1/2} V_s K_I(t) \\ &\times \int_0^t (\Sigma_s^{-1} U_s^T Q^{1/2} z) dt' - R^{-1/2} V_s K_P(t)z. \end{aligned} \quad (\text{C.7})$$

Finally, the controller can be put into a state-space form

$$\begin{aligned} \dot{x}_c &= A_c x_c + B_c z, \\ v_{fb} &= C_c(t)x_c + D_c(t)z, \end{aligned} \quad (\text{C.8})$$

where x_c is the controller state representing the integral term in (C.7), $A_c = 0$, $B_c = \Sigma_s^{-1} U_s^T Q^{1/2}$, $C_c(t) = -R^{-1/2} V_s K_I(t)$ and $D_c(t) = -R^{-1/2} V_s K_P(t)$.

Because we have chosen an infinite time horizon for the optimal control problem and the system matrices A_+ and B_+ remain approximately time-invariant after the short ramp-up phase of the discharge, S_+ will be approximately a constant

matrix \bar{S}_+ for most of the discharge. The matrix \bar{S}_+ can be obtained by solving the algebraic Riccati equation

$$0 = -\bar{S}_+ A_+(t_f) - A_+^T \bar{S}_+ + \bar{S}_+ B_+(t_f) R_+^{-1} B_+^T(t_f) \bar{S}_+ - Q_+, \quad (\text{C.9})$$

where t_f is a time during the flat-top phase. Utilizing this approximation, the controller gain could be reduced to a constant $\bar{K} = R_+^{-1} B_+^T(t_f) \bar{S}_+$ and the linear time-varying (LTV) system (30) would become linear and time-invariant (LTI).

References

- [1] Wijnands T, Houtte D V, Martin G, Litaudon X and Froissard P 1997 *Nucl. Fusion* **37** 777
- [2] Ferron J *et al* 2006 *Nucl. Fusion* **46** L13
- [3] Barana O, Mazon D, Laborde L and Turco F 2007 *Plasma Phys. Control. Fusion* **49** 947
- [4] Suzuki T 2005 Recent RF experiments and application of RF waves to real-time control of safety factor profile in JT-60U *AIP Conf. Proc.* **787** 279–86
- [5] Suzuki T *et al* 2008 *Nucl. Fusion* **48** 045002
- [6] Moreau D *et al* 2003 *Nucl. Fusion* **43** 870
- [7] Moreau D *et al* 2008 *Nucl. Fusion* **48** 106001
- [8] Laborde L *et al* 2005 *Plasma Phys. Control. Fusion* **47** 155
- [9] Moreau D *et al* 2011 *Nucl. Fusion* **51** 063009
- [10] Shi W *et al* 2012 Multivariable robust control of the plasma rotational transform profile for advanced tokamak scenarios in DIII-D *Proc. 2012 American Control Conf. (Montreal)*
- [11] Wehner W *et al* 2012 Optimal feedback control of the poloidal magnetic flux profile in the DIII-D tokamak based on identified plasma response models *Proc. 2012 American Control Conf. (Montreal)*
- [12] Felici F *et al* 2011 *Nucl. Fusion* **51** 083052
- [13] Ou Y *et al* 2007 *Fusion Eng. Des.* **82** 1153
- [14] Witrant E *et al* 2007 *Plasma Phys. Control. Fusion* **49** 1075
- [15] Ou Y *et al* 2008 *Plasma Phys. Control. Fusion* **50** 115001
- [16] Xu C *et al* 2010 *IEEE Trans. Plasma Sci.* **38** 163
- [17] Kim S and Lister J 2012 *Nucl. Fusion* **52** 074002
- [18] Ou Y *et al* 2011 *IEEE Trans. Control Syst. Technol.* **19** 432
- [19] Gaye O *et al* 2011 Sliding mode stabilization of the current profile in tokamak plasmas *Proc. 50th IEEE Conf. on Decision and Control and European Control Conf. (CDC-ECC) (Orlando, FL)* pp 2638–43
- [20] Ou Y, Xu C and Schuster E 2010 *IEEE Trans. Plasma Sci.* **38** 375
- [21] Argomedo F, Witrant E, Prieur C, Georges D and Bremond S 2010 Model-based control of the magnetic flux profile in a tokamak plasma *Proc. 49th IEEE Conf. on Decision and Control (CDC)* pp 6926–31
- [22] Barton J *et al* 2012 *Nucl. Fusion* **52** 123018
- [23] Boyer M D *et al* 2012 Backstepping control of the plasma current profile in the DIII-D tokamak *Proc. 2012 American Control Conf.*
- [24] Hinton F L and Hazeltine R D 1976 *Rev. Mod. Phys.* **48** 239
- [25] Blum J 1989 *Numerical Simulation and Optimal Control in Plasma Physics: with Applications to Tokamaks* (Paris: Wiley)
- [26] Ferron J *et al* 1998 *Nucl. Fusion* **38** 1055
- [27] Piglowski D, Ferron J, Gohil P, Johnson R and Penaflor B 2007 *Fusion Eng. Des.* **82** 1058
- [28] Schuster E, Walker M, Humphreys D and Krstic M 2003 Anti-windup scheme for plasma shape control with rate and magnitude actuation constraints in the DIII-D tokamak *42nd IEEE Conf. on Decision and Control*
- [29] Walker M *et al* 2007 *Fusion Eng. Des.* **82** 1051
- [30] Barton J, Ou Y, Xu C, Schuster E and Walker M 2011 *Fusion Eng. Des.* **86** 1116
- [31] Sauter O, Angioni C and Lin-Liu Y 1999 *Phys. Plasmas* **6** 2834