

Predictive Control of the Tokamak q Profile to Facilitate Reproducibility of High- q_{\min} Steady-State Scenarios at DIII-D

William Wehner, Menno Lauret, Eugenio Schuster, John R. Ferron, Chris Holcomb, Tim C. Luce, David A. Humphreys, Michael L. Walker, Ben G. Penaflor and Robert D. Johnson

Abstract—We consider control of the q profile while simultaneously regulating the plasma stored energy for the DIII-D tokamak. The main objective is to improve the shot-to-shot reproducibility and facilitate the accessibility of operating conditions that have steady-state potential, i.e. plasmas with large non-inductive current drive fractions. At DIII-D, non-inductive current sources including electron cyclotron current drive (ECCD) and neutral beam injection (NBI) allow the possibility of shaping the plasma current density distribution, and therefore enabling control of the q profile. A feedback controller is designed in a model predictive control framework to regulate the q profile while simultaneously regulating the plasma stored energy. The effectiveness of the control approach is demonstrated with experiments.

I. INTRODUCTION

In a reactor-grade tokamak, a plasma (hot ionized gas), typically of hydrogen ions, is confined by magnetic fields and heated to temperatures on the order of several keV . At such high temperatures, collisions between ions can overcome the repulsive forces due to their Coulomb fields thereby resulting in fusion reactions. The construction of a tokamak involves wrapping a set of coils poloidally around a toroidal vacuum vessel to produce a toroidal magnetic field. An additional poloidal field component is generated by running a current through the plasma. Combining the component fields, the net magnetic field lines wind helically around the torus. Following any magnetic field line a number of times around the torus maps out a closed flux surface, see Fig. 1.

The plasma current can be driven inductively by a transformer effect, where the plasma acts as secondary and a coil located at the center of the tokamak acts as the primary of the transformer. Tokamak operation typically also involves current driven by non-inductive sources. DIII-D is equipped with electron cyclotron current drive (ECCD), which involves injecting radio-frequency waves into the plasma, which drive current and also resonate with the gyro-kinetic orbit of the electrons, heating the plasma by an effect known as electron cyclotron resonant heating. Additional auxiliary current can be driven by neutral beam injection (NBI), which consists of injecting beams of highly energetic neutral particles into the plasma, driving current and heating the plasma through

collisions. Besides the auxiliary sources of ECCD and NBI, a self-generated source of current known as “bootstrap” current also contributes to the non-inductive current drive. Bootstrap current results from charged particles trapped between surfaces of constant magnetic flux, thus it is an automatically generated current [1].

A key issue for the success of a reactor-grade tokamak is the ability to operate in steady-state, meaning the plasma current is fully sustained with non-inductive sources. Development of the steady-state scenarios is an on-going area of research at the DIII-D tokamak (General Atomics, San Diego, USA) and other tokamaks around the world (see [2] for a review of progress towards steady-state operation). Maximizing the self-generated bootstrap current is naturally an important issue in developing a steady-state scenario. Two important quantities associated with maximizing the bootstrap current fraction are the q profile, which measures the pitch of the magnetic field lines, and the ratio of plasma kinetic pressure to normalized magnetic pressure, β_N . The bootstrap current fraction is proportional to q times β_N ($f_{BS} \propto q\beta_N$). Therefore, one path to realizing a steady-state scenario relies on high q_{\min} , which has the added benefit of raising the β_N limit associated with the onset of deleterious magneto-hydro-dynamics (MHD) instabilities such as tearing modes [3] and resistive wall modes (RWM) [4]. However, in order to maintain high fusion gain, proportional to β_N/q_{95}^2 , the q profile should be optimized by increasing q_{\min} while maintaining low q_{95} (value of q at the plasma edge) [5]. This motivates an interest in active control of the q profile to assist testing of operating conditions that are amenable to steady-state conditions.

In this work, we make use a first-principles-driven, control-oriented model of the current profile evolution suitable for high-confinement (H-mode) discharges in DIII-D. Advances towards developing low-order, control-oriented models for current profile control in various tokamaks are discussed in [6], [7], [8]. The models generally combine the magnetic diffusion equation [9] with empirical correlations for the electron temperature, resistivity, and non-inductive current drive.

Model predictive control of the tokamak q profile has already been considered in simulations [10], [11], [12]. As pointed out in [11], it is typically desired to operate tokamaks near stability limits in order to explore interesting physical phenomenon, thus, the optimal control solution for q profile control often lies at the intersection of various constraints. In particular, constraints associated with actuator limits and the

This work was supported in part by the U.S. Department of Energy (DE-SC0010661, DE-AC05-00OR23100, DE-FC02-04ER54698). W. Wehner (wehner@lehigh.edu), M. Lauret, and E. Schuster are with the Department of Mechanical Engineering and Mechanics, Lehigh University, Bethlehem, PA 18015, USA. C. Holcomb is with Lawrence Livermore National Laboratory, Livermore, CA, 94550, USA. J.R. Ferron, D.A. Humphreys, T.C. Luce, M.L. Walker, B.G. Penaflor, and R.D. Johnson are with General Atomics, San Diego, CA 92121, USA.

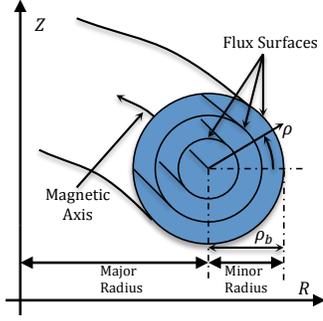


Fig. 1. Magnetic configuration of a tokamak (field lines follow a helical path around the tokamak). Flux surfaces represent points of constant poloidal magnetic flux.

parameter limits necessary to maintain stable plasmas (absent MHD). The most effective control approach is then naturally one that anticipates constraint violations and corrects for them in a systematic way. While we do not yet include state constraints in our control implementation, we have laid the ground work in building a suitable model predictive control framework for q profile control at the DIII-D tokamak that could later be expanded to include state constraints associated with the avoidance of deleterious MHD activity.

Model predictive control has developed significantly over the last few decades [13], while originally only applicable to problems with slow time scales due to the intense computational requirements. Improvements in optimization algorithms that exploit the problem structure have made MPC applicable to medium sized problems requiring fast update times between 1-5 ms [14]. We make use of a simple active set method [15], which combined with *warm-starting* of the optimization problem as described in Section III-B allows for sufficiently fast control computation times on average of 1 ms.

This work is organized as follows. The model structure of the q profile and plasma stored energy is described in Section II, details of the control design approach are given in Section III, experimental evidence of the effectiveness of the controller in reaching the target is presented in Section IV, and, finally, conclusions are made in Section V.

II. MODELING THE CURRENT PROFILE EVOLUTION AND PLASMA STORED ENERGY

The q profile is defined as the ratio of toroidal to poloidal magnetic flux gradients. Assuming a large aspect ratio tokamak, such as DIII-D, we can take the cylindrical approximation,

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

where t represents time, Φ is the toroidal magnetic flux, ψ is the poloidal magnetic flux per radian ($\psi = \Psi/2\pi$), ρ is the spatial coordinate, and $\hat{\rho}$ is the normalized spatial coordinate defined as ρ/ρ_b , where ρ_b is the value of ρ at the last closed magnetic flux surface. The evolution of ψ can be well approximated by the magnetic diffusion equation [16],

$$\frac{\partial\psi}{\partial t} = \frac{\eta(T_e)}{\mu_0\rho_b^2\hat{F}^2} \frac{1}{\hat{\rho}} \frac{\partial}{\partial\hat{\rho}} \left(\hat{\rho}\hat{F}\hat{G}\hat{H} \frac{\partial\psi}{\partial\hat{\rho}} \right) + R_0\hat{H}\eta(T_e)j_{\text{NI}}, \quad (2)$$

along with the boundary conditions $\frac{\partial\psi}{\partial\hat{\rho}}\Big|_{\hat{\rho}=0} = 0$, $\frac{\partial\psi}{\partial\hat{\rho}}\Big|_{\hat{\rho}=1} = -k_{I_p}I_p$, where η is the plasma resistivity, T_e is the electron temperature, μ_0 is the vacuum permeability, \hat{F} , \hat{G} , and \hat{H} are spatially varying geometric factors pertaining to the magnetic configuration of a particular plasma equilibrium [17], I_p is the total plasma current, and k_{I_p} is a constant. Contributions to the non-inductive current drive, j_{NI} , include the self-generated bootstrap current (j_{BS}) and each of the auxiliary sources such as ECCD (j_{EC}) and NBI (j_{NBI}),

$$j_{\text{NI}}(\hat{\rho}, t) = j_{\text{BS}}(\hat{\rho}, t) + \sum_{i=1}^{n_{\text{EC}}} j_{\text{EC},i}(\hat{\rho}, t) + \sum_{i=1}^{n_{\text{NBI}}} j_{\text{NBI},i}(\hat{\rho}, t). \quad (3)$$

At DIII-D there are $n_{\text{EC}} = 6$ ECCD sources and $n_{\text{NBI}} = 8$ NBI sources. The driven current density from each source depends on the plasma temperature, plasma density, and the physical alignment of the source with the plasma. We use time and spatially varying physics-based correlations to describe the plasma resistivity, electron temperature, and current drive efficiency for each auxiliary source [8]. The bootstrap current density, $j_{\text{BS}}(\hat{\rho}, t)$ can be modeled as a function of the temperature, density, and poloidal magnetic flux profiles and their gradients [18], [19].

As the q profile depends inversely on the spatial derivative of the poloidal flux, we define the terms

$$\theta(\hat{\rho}, t) \triangleq \frac{\partial\psi}{\partial\hat{\rho}}(\hat{\rho}, t), \quad \iota(\hat{\rho}, t) \triangleq \frac{1}{q(\hat{\rho}, t)} \triangleq \frac{-\theta}{B_{\phi,0}\rho_b^2\hat{\rho}}, \quad (4)$$

which will be useful for control purposes. We can differentiate (2) in space to obtain an expression for the evolution of θ ,

$$\frac{\partial\theta}{\partial t} = \frac{1}{\mu_0\rho_b^2} \frac{\partial}{\partial\hat{\rho}} \left\{ \frac{\eta}{\hat{F}^2} \frac{1}{\hat{\rho}} \frac{\partial}{\partial\hat{\rho}} \left(\hat{\rho}\hat{F}\hat{G}\hat{H} \frac{\partial\psi}{\partial\hat{\rho}} \right) \right\} + \frac{\partial}{\partial\hat{\rho}} (R_0\hat{H}\eta j_{\text{NI}}). \quad (5)$$

Thus, we can control q indirectly by controlling either θ or ι , and avoid the nonlinearity associated with the inverse of the spatial gradient of poloidal magnetic flux in (1).

The plasma stored energy, i.e. volume averaged energy density over the plasma volume, can be well approximated by the nonlinear first order system,

$$\frac{dE}{dt} = -\frac{E}{\tau_E} + P_{\text{tot}}(t), \quad (6)$$

where τ_E is the global energy confinement time. We use the ITER-98 (IPB98($y, 2$)) [20] scaling law to model the energy confinement time, $\tau_E \propto I_p^{0.93} \bar{n}_e^{0.41} P_{\text{tot}}^{-0.69}$. The total absorbed power, P_{tot} is equal to the auxiliary power injected into the plasma by NBI and ECCD, $P_{\text{aux}} = \sum_{i=1}^{n_{\text{NBI}}} P_{\text{NBI},i} + P_{\text{EC}}$, plus the power from the ohmic coil, P_{ohm} , minus the radiative power, P_{rad} , $P_{\text{tot}} = P_{\text{aux}} + P_{\text{ohm}} - P_{\text{rad}}$. The ohmic and radiative power can be described as a function of the electron density, electron temperature and toroidal current density, which can be determined from the poloidal magnetic flux profile [8].

To transform the model (5) to state-space form, we discretize the system in space using finite difference approximations to the spatial derivatives. The domain of interest,

$\hat{\rho} = [0, 1]$, is truncated to l evenly spaced nodes, separated by $\Delta\hat{\rho} = 1/(l-1)$ to obtain the finite-dimensional system,

$$\dot{\boldsymbol{\theta}} = \mathbf{f}(\boldsymbol{\theta}, \boldsymbol{u}), \quad \boldsymbol{\iota} = \mathbf{C}\boldsymbol{\theta}, \quad (7)$$

where the model state is $\boldsymbol{\theta} = [\theta_2, \theta_3, \dots, \theta_{l-1}]^T$, the output is $\boldsymbol{\iota} = [\iota_2, \iota_3, \dots, \iota_{l-1}]^T$, and the output matrix \mathbf{C} is obtained from (4). The available actuators represented by \boldsymbol{u} , include the total plasma current, the line averaged electron density, and the various gyrotron (ECCD) and neutral beam source powers, $\boldsymbol{u} = [I_p, \bar{n}_e, P_{EC,1}, \dots, P_{EC,n_{EC}}, P_{NBI,1}, \dots, P_{NBI,n_{NBI}}]^T$. Note that in addition to the non-inductive sources, we also have control of the total plasma current, I_p , allowing control of the profile boundary. The total plasma current, represents the sum of non-inductive and inductive currents. A low level dedicated controller exists to regulate the central coil voltage to drive any missing current between the requested total current and the sum of non-inductive sources.

III. CONTROL DESIGN

The control objective is to reach a specified target profile shape at a specified time. To accomplish this, first, an open-loop control problem is formulated as a trajectory optimization problem to find a feasible path from the expected initial condition to the desired target. The problem involves the minimization of a scalar objective over a set of constraints associated with the dynamics of the system (model of the q profile evolution), actuator constraints (physical limits such as maximum NBI power), and bounds on the acceptable current profile shape through the ramp-up phase (see [21] for details). The result of the optimization procedure is an open-loop sequence of NBI powers, ECCD powers, and total plasma current and a corresponding state trajectory that reaches the target, at least according to the model.

Let \boldsymbol{u}_{FF} represent the feedforward control sequence and $\boldsymbol{\theta}_{FF}$ (equivalently $\boldsymbol{\iota}_{FF}$) represent the corresponding feedforward state trajectory. After discretizing the system (7) in time and linearizing around the nominal feedforward trajectory (\boldsymbol{u}_{FF} , $\boldsymbol{\theta}_{FF}$, $\boldsymbol{\iota}_{FF}$), we obtain the linear time-varying system,

$$\tilde{\boldsymbol{\theta}}_{k+1} = \mathbf{A}_k \tilde{\boldsymbol{\theta}}_k + \mathbf{B}_k \tilde{\boldsymbol{u}}_k, \quad \tilde{\boldsymbol{\iota}}_k = \mathbf{C} \tilde{\boldsymbol{\theta}}_k, \quad (8)$$

where $\tilde{\boldsymbol{\theta}}_k$, $\tilde{\boldsymbol{\iota}}_k$, and $\tilde{\boldsymbol{u}}_k$, represent deviations from the feedforward trajectory, i.e. $\tilde{\boldsymbol{\theta}}_k = \boldsymbol{\theta}_k - \boldsymbol{\theta}_{FF,k}$, $\tilde{\boldsymbol{\iota}}_k = \boldsymbol{\iota}_k - \boldsymbol{\iota}_{FF,k}$, and $\tilde{\boldsymbol{u}}_k = \boldsymbol{u}_k - \boldsymbol{u}_{FF,k}$. Given that there is a one-to-one relationship between $\tilde{\boldsymbol{\theta}}$ and $\tilde{\boldsymbol{\iota}}$, we can eliminate $\tilde{\boldsymbol{\theta}}$ for simplicity of control design,

$$\tilde{\boldsymbol{\iota}}_{k+1} = \mathbf{A}_k \tilde{\boldsymbol{\iota}}_k + \mathbf{B}_k \tilde{\boldsymbol{u}}_k, \quad (9)$$

where we have overwritten \mathbf{A}_k and \mathbf{B}_k with $\mathbf{A}_k \leftarrow \mathbf{C}\mathbf{A}_k\mathbf{C}^{-1}$, $\mathbf{B}_k \leftarrow \mathbf{C}\mathbf{B}_k$.

A. q Profile Control

We use the total plasma current to regulate the q profile at the edge (q_{95}) with a classical linear quadratic integral (LQI) controller, the details of which are identical to [22]. This ensures that the controller will hit the target q_{95} value. In the control design approach that follows, including I_p as

a control variable is problematic because the target q profile is expected to be difficult to achieve. If I_p were included as an actuator, the controller might attempt to reduce I_p to undesirably low values to obtain the interior profile target. While it can be advantageous to use I_p to control the profile shape as a boundary actuator, we elect to dedicate I_p towards control of the profile edge to avoid the possibility of undesirably low I_p values. The following design could potentially be modified to include I_p as an actuator with added constraints to prevent I_p from falling undesirably low.

To control the q profile interior, we consider the trajectory tracking problem formulated as a finite-horizon, optimal control problem,

$$\begin{aligned} \text{minimize } J_k &= \tilde{\boldsymbol{\iota}}_{H_p}^T \mathbf{P} \tilde{\boldsymbol{\iota}}_{H_p} + \sum_{t=1}^{H_p-1} \tilde{\boldsymbol{\iota}}_{k+t}^T \mathbf{Q} \tilde{\boldsymbol{\iota}}_{k+t} \\ &\quad \{ \Delta \tilde{\boldsymbol{u}}_{k+t} \}_{t=0}^{H_u} \\ &\quad + \sum_{t=0}^{H_u} (\Delta \tilde{\boldsymbol{u}}_{k+t}^c)^T \mathbf{R} \Delta \tilde{\boldsymbol{u}}_{k+t}^c \end{aligned}$$

$$\begin{aligned} \text{subject to } \tilde{\boldsymbol{\iota}}_{k+t+1} &= \mathbf{A}_{k+t} \tilde{\boldsymbol{\iota}}_{k+t} + \mathbf{B}_{k+t}^c \tilde{\boldsymbol{u}}_{k+t}^c + \mathbf{B}_{k+t}^{nc} \tilde{\boldsymbol{u}}_{k+t}^{nc} \\ \tilde{\boldsymbol{\iota}}_k &= \tilde{\boldsymbol{\iota}}(k) : \text{initial condition} \\ \Delta \tilde{\boldsymbol{u}}_{k+t}^c &= \tilde{\boldsymbol{u}}_{k+t}^c - \tilde{\boldsymbol{u}}_{k+t-1}^c \\ \tilde{\boldsymbol{u}}_{k-1} &= \text{previously applied control} \\ \tilde{\boldsymbol{u}}_{k+t}^{nc} &= \tilde{\boldsymbol{u}}_k^{nc} \text{ for } t = 0, 1, \dots, H_u \\ \tilde{\boldsymbol{u}}_{k+t} &\in \tilde{\mathcal{U}}_{k+t} \text{ for } t = 0, 1, \dots, H_u \end{aligned} \quad (10)$$

The cost function J_k includes an instantaneous cost on deviations of the $\boldsymbol{\iota}$ profile from the desired feedforward trajectory ($\boldsymbol{\iota}_{FF}$) over the prediction horizon, H_p . Also, an instantaneous cost is applied to deviations in the control, $\Delta \tilde{\boldsymbol{u}}_{k+t}^c = \tilde{\boldsymbol{u}}_{k+t}^c - \tilde{\boldsymbol{u}}_{k+t-1}^c$, implying no cost for the control sequence to be away from the value associated with feedforward trajectory, \boldsymbol{u}_{FF} , but there is a cost for fast rate changes. The actuators have been split into controlled $\boldsymbol{u}_k^c = [P_{NB,3}, \dots, P_{NB,n_{NBI}}]$ and uncontrolled $\boldsymbol{u}_k^{nc} = [P_{NB,1}, P_{NB,2}, P_{EC,1}, \dots, P_{EC,n_{EC}}, I_p]$. The first two NBI are dedicated to diagnostics, the total plasma current is controlled via LQI control, and the ECCD are only controlled via feedforward. We allow the prediction horizon associated with the control, H_u , to be less than the prediction horizon associated with the state, to reduce the complexity of the problem. We assume no further update in the control beyond the control horizon, i.e. $\boldsymbol{u}_{k+t}^c = \boldsymbol{u}_{k+t-1}^c$ for $t \geq H_u$ and we replace the future uncontrolled actuators (which are yet unknown) with their current values. The term $\{ \tilde{\boldsymbol{u}}_{k+t} \in \tilde{\mathcal{U}}_{k+t} \}_{t=0}^{H_u}$ describes a set of linear constraints associated with the actuator limits. The solution to this optimization problem, \boldsymbol{z}^* , consists of the optimal predicted state and the optimal control update sequence,

$$\boldsymbol{z}^* = [\Delta \tilde{\boldsymbol{u}}_k^*, \tilde{\boldsymbol{u}}_{k+1}^*, \Delta \tilde{\boldsymbol{u}}_{k+1}^*, \dots, \Delta \tilde{\boldsymbol{u}}_{k+H_u}^*, \dots, \tilde{\boldsymbol{\iota}}_{k+H_p}^*]. \quad (11)$$

Of course, we cannot simply apply the resulting control sequence because the model used to predict the future states is not perfectly accurate. Therefore, the common practice is to apply the first step of the sequence, then sample the state again, and repeat the optimization procedure, which

introduces feedback to the control. Introducing feedback in this manner, i.e. solving the optimization problem, sampling the state, and solving the updated optimization problem again in a repetitive fashion leads to the feedback scheme known as “model predictive control” (MPC).

For control design purposes the stored energy evolution is approximated by its linearized dynamics,

$$\frac{dE}{dt} = -\frac{E}{\tau_{E_{\text{eq}}}} + P_{\text{aux}}(t), \quad (12)$$

where the contributions of ohmic power and radiative power are dropped since they are relatively small compared to the auxiliary power, and take $\tau_{E_{\text{eq}}}$ as the global energy confinement time associated with the operating point of the target q profile. The approximate energy dynamics (12) describe a linear first order system, therefore with a simple proportional-integral (PI) controller we can obtain any desired closed loop performance. With the PI controller,

$$P_{\text{aux}}^{\text{req}}(t) = k_p(E^d(t) - E(t)) + k_i \int_0^t E^d(\tau) - E(\tau) d\tau, \quad (13)$$

we can obtain a request for total auxiliary input power, $P_{\text{aux}}^{\text{req}}(t)$, which can enter into the MPC problem as an equality constraint on the total auxiliary power. The combination of q profile MPC with energy control constraint is depicted in Fig. 2. The constraint on total NBI power can be written as

$$\sum_{i=1}^{n_{\text{NBI}}} \tilde{P}_{\text{NBI},i} = P_{\text{aux}}^{\text{req}}(t) - \sum_{i=1}^{n_{\text{NBI}}} P_{\text{NBI,FF},i}, \quad (14)$$

In this manner we can obtain the desired plasma stored energy, and then allow the MPC controller to find the best combination of on-axis and off-axis NBI, satisfying the total power constraint to reach the desired q profile.

The optimization problem (10) combined with the constraint (14), is composed of the minimization of a quadratic function over a set of linear constraints, a quadratic program (QP).

$$\begin{aligned} & \underset{z}{\text{minimize}} && \frac{1}{2} z^T \mathbf{H} z + \mathbf{f}^T z \\ & \text{subject to} && \mathbf{A}_{\text{eq},k} z = \mathbf{b}_{\text{eq},k} \\ & && \mathbf{A}_{\text{in},k} z \leq \mathbf{b}_{\text{in},k} \end{aligned} \quad (15)$$

where $\mathbf{H} = \text{diag}\{\mathbf{R}, \mathbf{Q}, \mathbf{R}, \dots, \mathbf{R}, \mathbf{P}\}$, and $\mathbf{f} = 0$. The equality constraints consist of the total power constraint (14) and the dynamics constraints of (10), which can be easily converted to matrix form [23]. Note that the equality constraints have a subscript k denoting time dependence, which arises because the linearized dynamics are taken around the feedforward trajectory. The inequality constraints consist of actuator limits, which are also time varying because the limits are considered relative to the time varying feedforward control values. Actuator limits are implemented in the same fashion as [11].

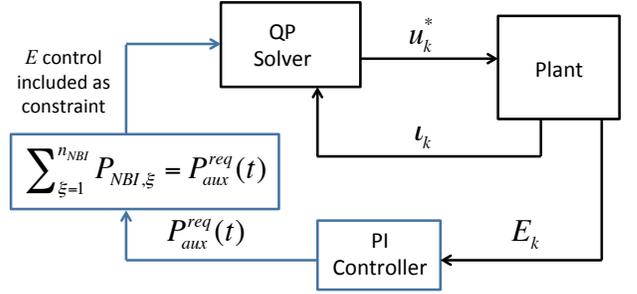


Fig. 2. q profile MPC with total power constraint to satisfy desired plasma stored energy.

B. Optimization

Consider the standard form QP (15)¹, for which we want to find the optimizer z^* . At the optimal solution, some of the inequalities will be *active*, i.e. the solution resides at the boundary of the inequality. If the set of active inequalities at the optimal solution, the *active set*, were known a priori, the remaining inactive inequalities could be dropped from the problem and the optimal solution could be found by solving an equality constrained problem.

Let \mathcal{A} represent the set of active inequalities constraints at the optimal solution, and let \mathbf{A}_{act} and \mathbf{b}_{act} represent each of the active inequalities. Assuming \mathbf{H} is positive definite, i.e. the problem is strictly convex, then the unique optimizer can be found from the single stationary point of the Lagrangian,

$$\begin{aligned} \mathcal{L}(z, \alpha, \lambda) = & \frac{1}{2} z^T \mathbf{H} z + \\ & \mathbf{f}^T z + \alpha^T (\mathbf{A}_{\text{eq}} z - \mathbf{b}_{\text{eq}}) + \lambda^T (\mathbf{A}_{\text{act}}^T z - \mathbf{b}_{\text{act}}), \end{aligned} \quad (16)$$

where α and λ have been introduced as Lagrange multipliers. The stationarity point of (16) can be found by solving the linear system

$$\begin{bmatrix} \mathbf{H} & \mathbf{A}_{\text{eq}}^T & \mathbf{A}_{\text{act}}^T \\ \mathbf{A}_{\text{eq}} & 0 & 0 \\ \mathbf{A}_{\text{act}} & 0 & 0 \end{bmatrix} \begin{bmatrix} z \\ \alpha \\ \lambda \end{bmatrix} = \begin{bmatrix} -\mathbf{f} \\ \mathbf{b}_{\text{eq}} \\ \mathbf{b}_{\text{act}} \end{bmatrix}. \quad (17)$$

The symmetric indefinite system (17) can be solved efficiently with the Schur complement method [24]. Let $\bar{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_{\text{eq}}^T & \mathbf{A}_{\text{act}}^T \end{bmatrix}^T$ and $\bar{\mathbf{b}} = [\mathbf{b}_{\text{eq}}^T \quad \mathbf{b}_{\text{act}}^T]^T$, and for convenience rewrite the system (17) as

$$\begin{bmatrix} \mathbf{H} & \bar{\mathbf{A}}^T \\ \bar{\mathbf{A}} & 0 \end{bmatrix} \begin{bmatrix} z \\ \begin{bmatrix} \alpha \\ \lambda \end{bmatrix} \end{bmatrix} = \begin{bmatrix} -\mathbf{f} \\ \bar{\mathbf{b}} \end{bmatrix}. \quad (18)$$

With the assumption that \mathbf{H} is positive definite, we can multiply the first equation in (18) by $\bar{\mathbf{A}} \mathbf{H}^{-1}$ and then subtract the second equation to obtain a linear system in the vector $[\alpha^T, \lambda^T]$ alone,

$$\left(\bar{\mathbf{A}} \mathbf{H}^{-1} \bar{\mathbf{A}}^T \right) \begin{bmatrix} \alpha \\ \lambda \end{bmatrix} = \left(\bar{\mathbf{A}} \mathbf{H}^{-1} \mathbf{f} - \bar{\mathbf{b}} \right). \quad (19)$$

Solve the symmetric positive definite system for α and λ , then recover z from the first equation of (18),

$$z = -\mathbf{H}^{-1} \left(\mathbf{f} + \bar{\mathbf{A}}^T \begin{bmatrix} \alpha \\ \lambda \end{bmatrix} \right) \quad (20)$$

¹For the remainder of this section we drop the time dependence of the problem, subscript k , for convenience.

This method requires H^{-1} and factorization of matrix $\bar{A}^T H^{-1} \bar{A}$, therefore it is most efficient when the inverse of H^{-1} can be easily computed, i.e. is diagonal, or can be precomputed, which is the case in the MPC approach described in this work.

An active set algorithm is a procedure to determine the active set \mathcal{A} . Numerous possibilities exist for finding \mathcal{A} , we use the simple active set method reported in [15]. For a convex inequality constrained QP, necessary and sufficient conditions for optimality are satisfied by the Karush Kuhn Tucker (KKT) conditions, which can be summarized as follows

$$\left. \begin{aligned} \mathbf{H}\mathbf{z} + \boldsymbol{\alpha}^T \mathbf{A}_{\text{eq}} + \sum_{i \in \mathcal{W}} \lambda_i \mathbf{a}_{\text{in},i}^T &= -\mathbf{f} \\ \mathbf{A}_{\text{eq}}\mathbf{z} &= \mathbf{b}_{\text{eq}} \\ \forall i \in \mathcal{A}: \quad \mathbf{a}_{\text{in},i}^T \mathbf{z} &= b_{\text{in},i} \\ \forall i \in \mathcal{I} \setminus \mathcal{A}: \quad \mathbf{a}_{\text{in},i}^T \mathbf{z} &\leq b_{\text{in},i} \\ \forall i \in \mathcal{A}: \quad \lambda_i &\geq 0 \end{aligned} \right\} \begin{array}{l} \text{KKT(1),} \\ \\ \\ \text{KKT(2),} \end{array} \quad (21)$$

where $\mathbf{a}_{\text{in},i}^T$ denotes the i^{th} row of the inequality constraint matrix \mathbf{A}_{in} , and \mathcal{I} represents the set of all inequality constraints. Let \mathcal{W} represent a guess of the active set \mathcal{A} . Using a guess of the active set, the solution to the problem determined by solving (17) with \mathcal{A} replaced by \mathcal{W} will satisfy KKT(1), but will not necessarily satisfy KKT(2) unless \mathcal{W} is equal to the optimal active set \mathcal{A} . The procedure to determine the active set starts by first taking a guess to the active set, call it the “working set” \mathcal{W} . From the working set, we construct an equality constrained QP and obtain a candidate solution, $\hat{\mathbf{z}}$. With the candidate solution, we can check the conditions KKT(2), any violated inequality constraints are added to the working set and any constraint i for $\lambda_i < 0$ currently in the working set is removed.² The process is repeated until no constraints are added or removed from the working set. The active set algorithm is formally summarized in Algorithm 1.

Algorithm 1: QP by Active Set

```

iter ← 0
repeat
  Compute candidate solution  $\hat{\mathbf{z}}$  via (19)-(20),
  if  $\mathbf{a}_{\text{in},i}^T \hat{\mathbf{z}} > b_{\text{in},i} |_{i \notin \mathcal{W}}$  then
    add  $i$  to  $\mathcal{W}$ 
  end if
  if  $\lambda_i < 0 |_{i \in \mathcal{W}}$  then
    remove  $i$  from  $\mathcal{W}$ 
  end if
  iter ← iter + 1
  if iter > MAXITER then
    return Fail
  end if
until  $\mathbf{z}$  satisfies KKT(2)
 $\mathcal{A} = \mathcal{W}$  and  $\mathbf{z} = \hat{\mathbf{z}}$ 
return  $\mathcal{A}, \mathbf{z}$ 

```

²A negative multiplier implies the optimal solution lies within this inequality constraint.

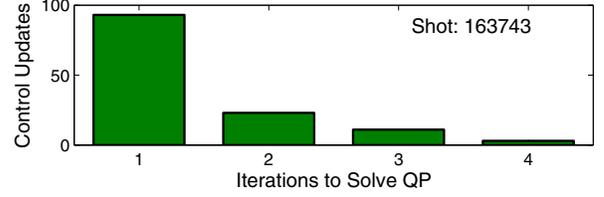


Fig. 3. Number of active set (Algorithm 1) iterations to solve the QP for each control update.

The efficiency of the algorithm is primarily due to the fact that the number of active inequality constraints on sequential control updates does not change dramatically, and thus the active set information from the previous control update can be used to warm start the QP on the next control update. This allows for a significant reduction in the number of iterations, and therefore linear system solves, necessary to solve the QP. For more than 90% of control updates, the QP is solved in fewer than 3 iterations as shown in Fig. 3. However, this particular active set algorithm is not guaranteed to succeed. In the instances in which it fails, we default the optimization procedure of the feedback controller to a basic interior point algorithm [24].

IV. SIMULATIONS AND EXPERIMENTAL RESULTS

In this section we consider two q profile control experiments on the DIII-D tokamak demonstrating the effectiveness of the controller. A new control request is computed every 20 ms, however, we apply a zero-order-hold on updates to the model matrices (9) of 100ms, reducing the set of QP matrices to 24 for the entire tokamak discharge (we control from 0.4-3 s).

The target q profile is a zero shear profile with $q_{\text{min}} = 1.6$ and $q_{95} = 5$. The first experiment involves q profile control only (Fig. 4) and the second involves q profile plus energy control (Fig. 5), where the energy control is incorporated via a constraint to the MPC problem as described in Section III-A. For both experiments, we plot the obtained q profile at the top of the respective figures in comparison to the q profile obtained with feedforward control alone. The failure of the feedforward control action to reach the target is primarily due to slight modeling errors in the resistive diffusion rate. This emphasizes the importance of feedback control, which is able to account for the modeling errors and bring the q profile back on target.

In the middle section of the figures, the plasma stored energy is plotted, comparing the target value with that obtained with feedforward control only and feedforward + feedback control. Comparing Fig. 4 vs Fig. 5, we see that controlling the q profile alone leads to rather aggressive control action and therefore a choppy undesirable response in the plasma stored energy when compared to that obtained with q + energy control. Note, that at the bottom of Fig. 4, the effects of the aggressive control action are observed in the response of β_N , which can be problematic. If β_N is allowed to go too high too early this can potentially seed magnetic islands that can grow and corrupt the confinement of the plasma.

V. SUMMARY AND CONCLUSIONS

We described a MPC trajectory tracking control approach for regulation of the q profile in the DIII-D tokamak that simultaneously maintains the desired plasma stored energy via a constraint on total auxiliary power. The controller was formulated as a strictly convex QP and was solved using an active-set algorithm that exploits the consistency between active constraints in subsequent control steps for fast computation speed. Experimental results successfully demonstrated the capability of the control approach to reach desired q profile targets while regulating plasma stored energy.

REFERENCES

- [1] J. Wesson, *Tokamaks*. Oxford, UK: Clarendon Press, 1984.
- [2] G. Gormezano *et al.*, *Nucl. Fusion*, vol. 47, pp. S285–S336, 2007.
- [3] R. LaHaye, *Physics of Plasmas*, vol. 13, no. 055501, 2006.
- [4] A. Bondeson *et al.*, *Phys. Rev. Lett.*, vol. 72, pp. 2709–2712, 1994.
- [5] J. Ferron *et al.*, “Optimizing the β limit in DIII-D Advanced Tokamak Discharges,” *Phys. of Plasmas*, vol. 12, p. 056126, 2005.
- [6] E. Witrant *et al.*, “A control-oriented model of the current profile in tokamak plasmas,” *Plasma Phys. and Control. Fusion*, vol. 49, pp. 1075–1105, 2007.
- [7] F. Felici *et al.*, “Real-time physics-model-based simulation of the current density profile in tokamak plasmas,” *Nuclear Fusion*, vol. 51, no. 083052, 2011.
- [8] J. Barton *et al.*, “Physics-based Control-oriented Modeling of the Current Density Profile Dynamics in High-performance Tokamak Plasmas,” in *52nd IEEE Conference on Decision and Control*, 2013.
- [9] F. Hinton *et al.*, *Rev. Mod. Phys.*, vol. 48, pp. 239–308, 1976.
- [10] H. Ouarit, “Validation of plasma current profile model predictive control in tokamaks via simulation,” *Fusion Engr. and Design*, vol. 86, pp. 1018–1021, 2011.
- [11] E. Maljaars, F. Felici *et al.*, “Control of the tokamak safety factor profile with time-varying constraints using MPC,” *Nucl. Fusion*, vol. 55, no. 527420, 2015.
- [12] I. Garrido *et al.*, “Low Effort l_i Nuclear Fusion Plasma Control Using Model Predictive Control Laws,” *Mathematical Problems in Engineering*, vol. 2015, p. 023001, 2015.
- [13] J. H. Lee, “Model Predictive Control: Review of the Three Decades of Development,” *Intl Journal of Control, Automation, and Systems*, vol. 9, no. 3, pp. 415–424, 2011.
- [14] Y. Wang and S. Boyd, “Fast Model Predictive Control Using Online Optimization,” *IEEE Trans. on Control Systems Technology*, vol. 18, no. 2, p. 267, 2010.
- [15] S. Kuindersma, F. Permenter, and R. Tedrake, “An Efficiently Solvable Quadratic Program for Stabilizing Dynamic Locomotion,” in *International Conference on Robotics and Automation (ICRA)*, 2014.
- [16] F. Hinton and R. Hazeltine, “Theory of plasma transport in toroidal confinement systems,” *Rev. Mod. Phys.*, vol. 48, pp. 239–308, 1976.
- [17] Y. Ou, T. Luce, E. Schuster *et al.*, “Towards model-based current profile control at DIII-D,” *Fusion Engineering and Design*, vol. 82, pp. 1153–1160, 2007.
- [18] O. Sauter *et al.*, *Physics of Plasmas*, vol. 6, no. 7, p. 2834, 1999.
- [19] —, *Physics of Plasmas*, vol. 6, no. 7, p. 2834, 1999.
- [20] ITER Physics Basis, *Nuclear Fusion*, vol. 39, p. 2137, 1999.
- [21] J. Barton *et al.*, “Physics-model-based nonlinear actuator trajectory optimization and safety factor profile feedback control for advanced scenario development in DIII-D,” *Nuclear Fusion*, vol. 55, no. 9, p. 093005, 2015.
- [22] W. Wehner *et al.*, “Current Profile Control for the Development of Consistent Discharges in DIII-D,” in *54th IEEE Conference on Decision and Control*, 2015.
- [23] F. Borrelli, A. Bemporad, and M. Morari, *Predictive control for linear and hybrid systems*. Cambridge University Press, 2011, In press (version updated: April 22, 2014), 2014.
- [24] J. Nocedal and S. J. Wright, *Numerical Optimization 2nd edn*. Berlin: Springer, 2006.

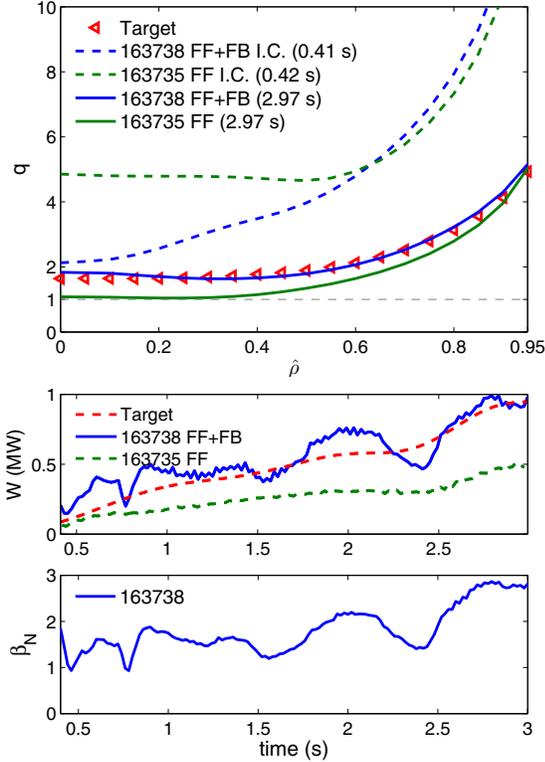


Fig. 4. Shot 163738: q profile control only. Top: Achieved q profile at target time for feedforward control (shot 163735) and feedforward + feedback control (shot 163738) cases. Middle: Plasma stored energy. Bottom: β_N .

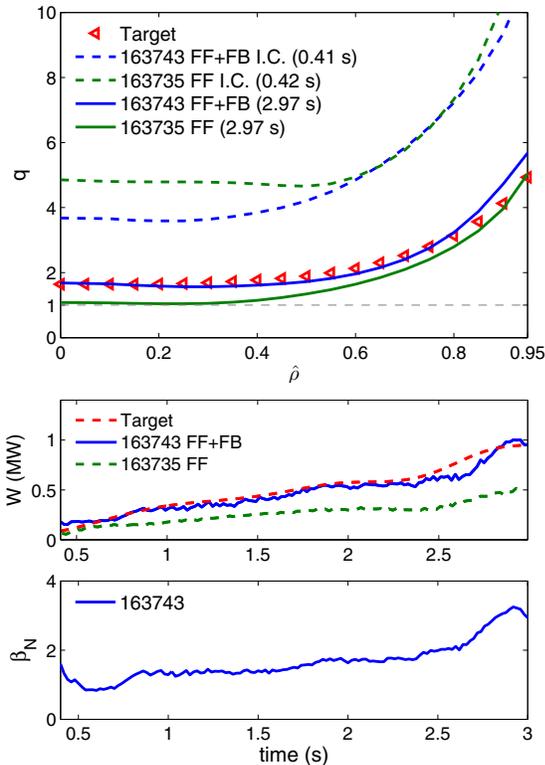


Fig. 5. Shot 163743: q profile + energy control. Top: Achieved q profile at target time for feedforward control (shot 163735) and feedforward + feedback control (shot 163743). Middle: Plasma stored energy. Bottom: β_N .