Physics-model-based Real-time Optimization for the Development of Steady-state Scenarios at DIII-D

W.P. Wehner, A. Pajares, E. Schuster
Lehigh University
Bethlehem, Pennsylvania 18015, USA.
Email: wehner@lehigh.edu (Corresponding author W.P. Wehner)

J.R. Ferron, D.A. Humphreys, R.D. Johnson, B.G. Penaflor, K.E. Thome, M.L. Walker
General Atomics
San Diego, California 92121, USA.

C.T. Holcomb, B.S. Victor
Lawrence Livermore National Laboratory
Livermore, California, 94550, USA.

Abstract

In this work, a novel q profile control approach and recent DIII-D experimental results aimed at reaching stationary plasmas characterized by a flat loop voltage profile are presented. The control approach combines commands computed both offline (feedforward) and online (feedback). Both command components are computed via numerical optimal control techniques. The key advantage of the numerical computation approach is that it allows for the explicit incorporation of state and input constraints to prevent the controller from driving the plasma outside of stability limits and obtain, as closely as possible, stationary conditions characterized by a flat loop voltage profile. Using a suitable control-oriented model, the simulated plasma evolution in response to the actuators is embedded into a nonlinear optimization problem that provides a feedforward control policy (set of actuator waveforms) that under ideal conditions guides the plasma evolution to the desired state. The time trajectory of the plasma current, gyrotron power, and neutral beam power are optimized to guide the plasma to stationary state characterized by a flat loop voltage profile. It is shown in simulations that an overshoot in the plasma current during ramp-up combined with a particular timing of the gyrotron and neutral beam injection can improve the uniformity of the loop voltage profile. The feedback controller computes updates to the feedforward control law to account for variability in plasma conditions; optimizing in real-time the model-predicted plasma evolution in response to the available actuator set.

1 Introduction

A primary goal for the DIII-D research program over the next five years is to develop the physics basis for a high q ($q_{\min} > 1.5 - 2.0$), high β_N steady-state scenario[1] that can serve as the basis for a steady-state ITER scenario at (fusion gain) $Q = 5$. Various approaches are being considered to maximize the bootstrap current contribution, so that fully noninductive ($f_{NI} = 1$) discharges can be obtained for several resistive current diffusion times. It is anticipated that the upgrades to DIII-D including an additional off-axis neutral beam injection (NBI) system in 2019 will provide sufficient auxiliary current drive to maintain fully noninductive plasmas at high β_N. However, much work is necessary to investigate MHD stability, adequate confinement, and early achievement and sustainment of the steady-state condition. Recent work in active control of the q profile accomplished with a combination of offline (feedforward control) and online (feedback control) is playing a crucial role at DIII-D in achieving and sustaining advanced scenarios of interest and in facilitating investigation of their properties [1].

With the use of an appropriate control-oriented model, a feedforward (open-loop) control problem is formulated as a nonlinear optimization problem and is solved using numerical optimization techniques. The result comprises a sequence of feedforward control requests and a corresponding state evolution which evolves from the anticipated initial plasma state to the desired target, which is defined by a flat loop voltage as early as possible in the discharge. Previous works have considered feedforward control synthesis for profile control via nonlinear optimization approaches [2–6] and extremum seeking [7].

Anticipating that the un-augmented feedforward control would be insufficient to reliably reach the target q profile due to mismatch between the model and real system, the feedforward control sequence is combined with a feedback controller, which computes updates to the feedforward control based on measurements of the plasma state in real-time during the discharge. The feedback controller is based on a model-predictive control (MPC) or receding-horizon control (RHC) framework, which continuously simulates the q profile evolution over a short horizon (number of future control time-steps) to help inform the feedback controller of the best control action by

[1]Steady state scenario is characterized by a plasma state that is fully relaxed and the plasma current is composed entirely of intrinsic (bootstrap) and noninductive auxiliary current drives.

solving an optimization problem in real time. The primary motivation for the MPC feedback approach is that it allows for explicit incorporation of constraints associated with the plasma state. This opens the possibility of designing an aggressive profile tracking controller while restricting the controller from driving the plasma outside of stability limits. Additionally, the MPC framework enables the inclusion of constraints associated with the achievement of steady-state conditions defined by a zero loop voltage profile, or, less ambitiously, stationary conditions defined by a flat loop voltage profile. Recent experimental tests have shown successful testing of model predictive control of the tokamak q profile at DIII-D [8] and ASDEX upgrade [9].

This paper is organized as follows. In section 2, the structure of the control-oriented model describing the q profile and plasma stored energy dynamics is briefly outlined. The discussion then proceeds to the model-based control design process starting with feedforward control design via nonlinear optimization in section 3 and following with MPC feedback design in section 4. The combined feedforward + feedback control system is embedded into the DIII-D plasma control system (PCS), and initial tests are presented in Section 5.

2 Control-oriented Modeling of the Poloidal Magnetic Flux Profile and Plasma Stored Energy

A first-principles-driven (FPD) control-oriented model of the poloidal flux evolution developed for DIII-D H-mode plasmas [10], is used as the basis for control design. Assuming an axisymmetric plasma and taking flux surface averages of all quantities, Ampere's law, Faraday's law, and Ohm's law can be combined to form a 1D partial differential equation describing the evolution of the poloidal magnetic flux, referred to as the magnetic diffusion equation (MDE) [11, 12]. To simplify the model to a control-oriented form, the MDE can be combined with physics-based correlations for the electron temperature, plasma resistivity, and the efficiency of each of the current drive sources including neutral beam injection (NBI), electron cyclotron current drive (ECCD), and bootstrap current drive [10, 13]. In summary, the model describing the evolution of the poloidal magnetic flux profile, $\Psi(\hat{\rho},t)$, and the plasma stored energy, $E(t)$, in response to the various heating and current drive sources can be briefly written as

$$\frac{\partial\psi}{\partial t} = \frac{\eta(T_\mathrm{e})}{\mu_0\rho_\mathrm{b}^2\hat{F}^2}\frac{1}{\hat{\rho}}\frac{\partial}{\partial\hat{\rho}}\left(\hat{\rho}\hat{F}\hat{G}\hat{H}\frac{\partial\psi}{\partial\hat{\rho}}\right) + \frac{R_0\hat{H}\eta(T_\mathrm{e})}{B_{\phi,0}}\left(j_\mathrm{aux}+j_\mathrm{bs}\right), \qquad \frac{dE}{dt} = -\frac{1}{\tau_E}E + P_\mathrm{tot}, \tag{1}$$

where $\psi(\hat{\rho},t)$ is the poloidal magnetic flux per radian, i.e. $\Psi(\hat{\rho},t) = 2\pi\psi(\hat{\rho},t)$, and $\hat{\rho}$ is the normalized spatial coordinate defined in terms of the toroidal magnetic flux, Φ, and the toroidal magnetic field strength at the magnetic axis, $B_{\phi,0}$, i.e. $\pi B_{\phi,0}\rho^2 = \Phi$ ($\hat{\rho} = \rho/\rho_b$), and the boundary conditions given by

$$\left.\frac{\partial\psi}{\partial\hat{\rho}}\right|_{\hat{\rho}=0} = 0, \qquad \left.\frac{\partial\psi}{\partial\hat{\rho}}\right|_{\hat{\rho}=1} = -\frac{\mu_0}{2\pi}\frac{R_0}{\hat{G}|_{\hat{\rho}=1}\hat{H}|_{\hat{\rho}=1}}I_\mathrm{p}(t) \tag{2}$$

Of the parameters included in (1)-(2), η is the plasma resistivity, T_e is the electron temperature, μ_0 is the vacuum permeability, R_0 is the plasma major radius, $\hat{F}$, $\hat{G}$, and $\hat{H}$ are spatially varying geometric factors pertaining to the magnetic configuration of a particular plasma equilibrium, I_p is the total plasma current, j_aux is the noninductive current density from all sources, j_bs is the bootstrap current density, and ρ_b is the value of ρ at the boundary. The resistivity is modeled by a simplified Spitzer model, and the electron temperature, which evolves much faster than the poloidal magnetic flux diffusion time, is modeled as a fixed profile shape that scales with line averaged electron density, total plasma current, and total input power [10].

The bootstrap current density, $j_\mathrm{bs}(\hat{\rho},t)$, which is calculated using the Sauter model [14, 15], is a function of the electron temperature, electron density, and poloidal magnetic flux profiles and their gradients. The auxiliary noninductive current drive, j_aux, includes contributions from ECCD, j_EC, and NBI, j_NBI, i.e. $j_\mathrm{aux}(\hat{\rho},t) = \sum_{i=1}^{n_\mathrm{EC}} j_{\mathrm{EC},i}(\hat{\rho},t) + \sum_{\xi=1}^{n_\mathrm{NBI}} j_{\mathrm{NBI},\xi}(\hat{\rho},t)$, where n_EC is the number of ECCD sources and n_NBI is the number of NBI sources.

The parameter τ_E represents the global energy confinement time, and is approximated by the ITER-98 ([16]) scaling law, $\tau_E \propto I_\mathrm{p}^{0.93}\bar{n}_e^{0.41}P_\mathrm{tot}^{-0.69}$, where $\bar{n}_e$ is the line averaged electron density, and the total absorbed power, P_tot, is equal to the total injected power from auxiliary sources plus the power from the ohmic coil, P_ohm, minus the radiative power, P_rad, i.e. $P_\mathrm{tot} = P_\mathrm{aux} + P_\mathrm{ohm} - P_\mathrm{rad}$.

Plasma parameters important for control design include the safety factor profile, q, the normalized plasma β, β_N, and the plasma loop voltage, V_loop, which are given by

$$q(\hat{\rho},t) = \frac{d\Phi}{d\Psi} = -\frac{\partial\Phi/\partial\rho}{2\pi\partial\psi/\partial\rho} = -\frac{B_{\phi,0}\rho_\mathrm{b}^2\hat{\rho}}{\partial\psi/\partial\hat{\rho}}, \qquad \beta_N = \frac{(2/3)(E/V_\mathrm{p})}{B_{\phi,0}^2/2\mu_0}\frac{aB_{\phi,0}}{I_\mathrm{p}[MA]}, \qquad V_\mathrm{loop}(\hat{\rho},t) = -2\pi\frac{\partial\psi}{\partial t}(\hat{\rho},t), \tag{3}$$

where a is the plasma minor radius and V_p is the total plasma volume.

3 Feedforward Control Design via Trajectory Optimization

To find a feedforward control policy, the time trajectories of the plasma current, ECCD power, and NBI power are optimized so as to maximize the stationarity of the plasma in the flattop, i.e. flatten the loop voltage profile as much as possible subject to various constraints. The problem is formulated as a constrained nonlinear optimization problem, and solved with sequential quadratic programming (SQP). Actuator and physics-based constraints are imposed such as limits on plasma current ramp rate, limits on available ECCD and NBI power, shape of the q profile, and β_N. It is generally found in the optimization results that an overshoot in the plasma current during ramp-up combined with a certain EC and NBI timing can produce a nearly flat loop voltage profile.

The system defined by (1) and (2) represents a PDE system for which it is very challenging and usually impossible to design optimal feedback control laws due to the infinite dimensionality of the system that arises from the continuous spatial domain. To transform the system to a finite dimensional state-space form, it can be discretized in space using finite difference approximations to the spatial derivatives. Let the discretized system be represented by

$$\dot{\boldsymbol{\psi}} = \boldsymbol{f}_{\boldsymbol{\psi}}(\boldsymbol{\psi}, \mathbf{u}), \tag{4}$$

where $\boldsymbol{\psi} = [\psi_2, \psi_3, \ldots, \psi_{l-1}]^T$, ψ_i are the values of $\psi(\hat{\rho}, t)$ at evenly spaced internal radial nodes, and the input, $\mathbf{u} = [I_\mathrm{p}, \bar{n}_\mathrm{e}, P_{\mathrm{EC},1}, \ldots, P_{\mathrm{EC},n_\mathrm{EC}}, P_{\mathrm{NBI},1}, \ldots, P_{\mathrm{NBI},n_\mathrm{NBI}}]^T$, represents the set of actuators including, total plasma current, line averaged electron density, individual ECCD sources, and individual NBI sources[2]. A state-space model for the combined evolution of $\boldsymbol{\psi}$ and E can be written as

$$\dot{\mathbf{x}} = \boldsymbol{f}(\mathbf{x}, \mathbf{u}) = \begin{bmatrix} \boldsymbol{f}_{\boldsymbol{\psi}}(\boldsymbol{\psi}, \mathbf{u}) \\ -\frac{E}{\tau_\mathrm{E}} + P_\mathrm{tot}, \end{bmatrix}, \quad \mathbf{y} = \mathbf{h}(\mathbf{x}), \tag{5}$$

where the model state is $\mathbf{x} = [\boldsymbol{\psi}^T, E]^T$, and the system output includes the q profile at various $\hat{\rho}$ locations, $\boldsymbol{q} = [q_{\hat{\rho}=0.05}, q_{\hat{\rho}=0.01}, \ldots, q_{\hat{\rho}=.95}]^T$, β_N, and the loop voltage profile $\mathbf{V}_\mathrm{loop} = [V_{\mathrm{loop},\hat{\rho}=0.05}, V_{\mathrm{loop},\hat{\rho}=0.01}, \ldots, V_{\mathrm{loop},\hat{\rho}=.95}]^T$, i.e. $\mathbf{y} = [\boldsymbol{q}^T, \beta_N, \mathbf{V}_\mathrm{loop}^T]^T$.

The control objective is to design a set of actuator waveforms such that the plasma evolves to a stationary condition characterized by a uniform loop voltage profile in the plasma flattop. A feedforward control policy can be determined from an optimization problem involving the minimization of a scalar objective over a set of constraints associated with the dynamics of the system (model of the ψ profile evolution), actuator constraints (physical limits such as maximum NBI power), and other physics-based constraints,

$$\text{(P1)} \qquad \begin{aligned} &\underset{\mathbf{u}(t)}{\text{minimize}} && J(\mathbf{y}(t)) && && \text{Minimize Cost Function } J \\ &\text{subject to} && \mathbf{x}(t_0) = \mathbf{x}_0, \quad \dot{\mathbf{x}} = \boldsymbol{f}(\mathbf{x}, \mathbf{u}), \quad \mathbf{y} = \mathbf{h}(\mathbf{x}, \mathbf{u}), && && \text{Impose System Dynamics Constraint} \\ &&& \mathbf{g}_\mathrm{in}(\mathbf{y}(t), \mathbf{u}(t)) \leq 0, && \Big\} && \text{Impose Constraints for Allowable} \\ &&& \mathbf{g}_\mathrm{eq}(\mathbf{y}(t), \mathbf{u}(t)) = 0. && && \text{Operating Space} \end{aligned}$$

This is often called a trajectory optimization problem because it involves the search for a state trajectory, $\mathbf{x}(t)$, that starts from the initial state, $\mathbf{x}_0$, and reaches some goal state, which is consistent with the system dynamics described by the equality constraints $\dot{\mathbf{x}} = \boldsymbol{f}(\mathbf{x}, \mathbf{u})$. Distance to the target state is quantified by the scalar cost function J. Additionally, the optimization problem requires the state trajectory to avoid undesirable regions of the state space, which are quantified by the constraints, $\mathbf{g}_\mathrm{in}$ and $\mathbf{g}_\mathrm{eq}$. The undesirable portion of the state space equates to the regions of the tokamak operating space that are associated with MHD instabilities. The most successful approach to solving an optimal control problem like (P1) is to parameterize the problem with a finite set of decision variables, and then to solve it by using numerical optimization methods [17].

The simulated evolution of the plasma response to the actuators is embedded into the problem, and the sensitivities of the cost function and constraints are computed with respect to each optimization variable (actuator values during the discharge). The SQP algorithm begins with an initial guess solution, and iteratively improves on the solution until reaching a feasible local minimum of the cost function [18]. On each iteration, cost and constraint values and sensitivities with respect to the current iterate solution are passed to the SQP algorithm which uses this information to improve on the guess solution. See [19] for more details associated with solving the trajectory optimization problem with SQP as well as parameterizing the problem and computing sensitivities.

[2]The end points, ψ_1 and ψ_l, are not part of the state space description because they are fixed by the boundary conditions.

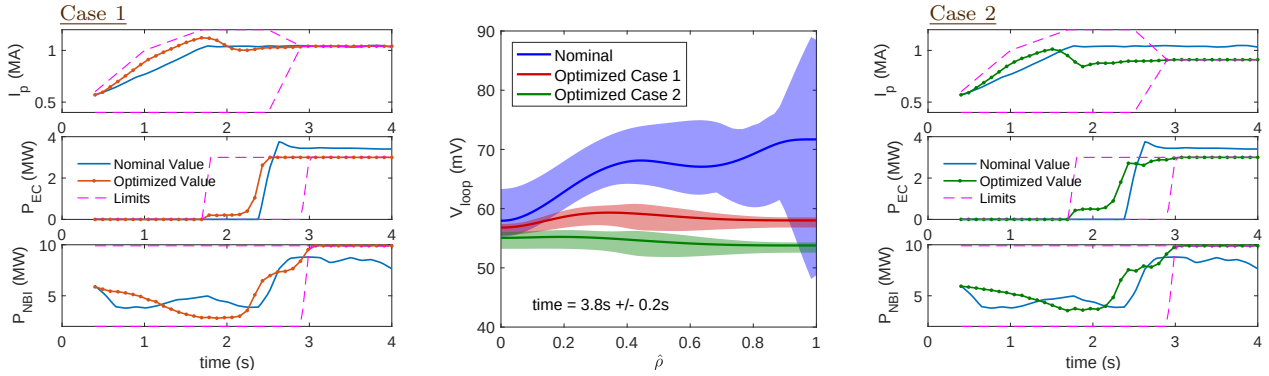

(a) Case 1: flattop $I_\mathrm{p} = 1.04$ MA. Nominal (blue) and optimized (red).

(b) Nominal and optimized loop voltage profiles.

(c) Case 2: flattop $I_\mathrm{p} = 0.91$ MA. Nominal (blue) and optimized (green).

Figure 1: Results of the optimization. (a) Case 1 with flattop I_p fixed to 1.04 MA to align with reference shot 172538 (nominal case). (b) The simulated loop voltage profiles associated with nominal control solution (blue), optimized control solution case 1 (red) and optimized control solution case 2 (green). (c) Case 2 with flattop I_p fixed to 0.91 MA for a noninductive current fraction around 80% (found by predictive TRANSP runs).

3.1 Cost Function and Constraints

The cost function is chosen as

$$J(\mathbf{x}(t)) = \int_0^{t_F} Q(t) \int_0^1 \left(V_\mathrm{loop}(t,\hat{\rho}) - V_\mathrm{loop}(t,1)\right)^2 d\hat{\rho}\, dt, \tag{6}$$

which takes the form of a path cost, i.e. it integrates over the entire trajectory (path of the temporal state evolution). Initial work attempted to find a stationary state at the end of the discharge, by assigning cost only at the final time (t_F), i.e. $J = J(\mathbf{x}(t_F)) = \int_0^1 \left(V_\mathrm{loop}(t_F,\hat{\rho}) - V_\mathrm{loop}(t_F,1)\right)^2 d\hat{\rho}$. While this approach was successful in obtaining a flat loop voltage profile at the target time, t_F, it may only be flat momentarily. With a time-integrated cost function, a profile can be obtained that remains flat, or as flat as possible, through the entire flattop phase. This approach ensures, the optimization does not find a solution that simply passes through a stationary state temporarily.

The integrand penalizes deviations of the loop voltage profile with respect to its edge value, and a time dependent weight, $Q(t)$, is introduced to focus the cost towards the flattop phase of the discharge, i.e. $Q(t)$ is zero before the flattop. In previous work [4], a cost function of the type, $J = \int_0^1 \left(\partial V_\mathrm{loop}/\partial\hat{\rho}\right)^2 d\hat{\rho}$, was used to penalize deviations of the spatial derivative of the loop voltage profile from zero. The idea being that if its spatial derivative is identically zero, then the loop voltage profile will be flat. However, cost functions of this type were found to have local minimums at plasma states corresponding to monotonically increasing loop voltage profiles, and, therefore, not suitable for local optimization approaches such as SQP [18].

The constraints include bounds and rate limits on the actuators as well as physics-based constraints. The reference plasma current is treated as an actuator and it is assumed the inductive coil current is regulated via a dedicated controller to meet the desired target current. A rate limit is imposed on the plasma current ramp of 2 MA per second,

$$\left|\frac{dI_\mathrm{p}}{dt}\right| \leq 2 \text{ MA/s}, \tag{7}$$

and physical limits are imposed on the physical actuators. The counter-current beams are not considered to avoid the possibility of triggering undesirable MHD, such as NTMs or locked modes. The neutral beams are treated as continuously variable actuators, while in reality they function in an ON/OFF manner. During actual experiments, appropriate modulation of the beams is necessary to deliver the desired power on average.

A limit on the q profile evolution is imposed to avoid the onset of sawtooth oscillations and other undesirable MHD, a limit on density associated with the Greenwald density limit, and a lower bound on the total auxiliary input power is imposed in order to prevent back transitions to L-mode, i.e.

$$q(\hat{\rho}) > 1, \qquad \bar{n}_e[10^{20}\mathrm{m}^{-3}] \leq \frac{I_\mathrm{p}[\mathrm{MA}]}{\pi a^2}, \qquad P_\mathrm{aux} \geq 2 \text{ MW}. \tag{8}$$

3.2 Optimization Results

DIII-D shot 172538 is taken as the starting point for the feedforward control optimization. This shot is used as the reference shot for experimental testing (see Section 5). It represents a promising starting point for high

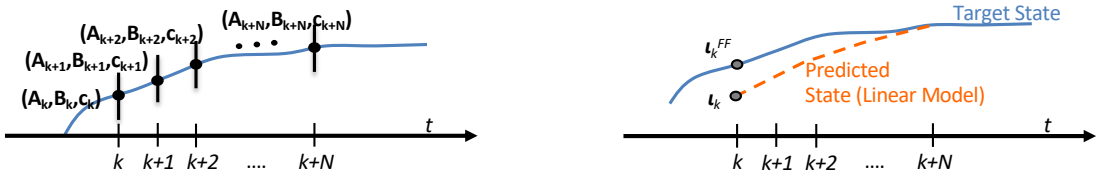

(a) Linearizations around feedforward trajectory.

(b) Finite horizon predictive reference tracking.

Figure 2: Diagrams for linearization around the feedforward trajectory and predictive feedback control as reference tracking around the feedforward trajectory.

performance steady-state scenario development because it has high $q_{\min}$ with high triangularity and elevated energy confinement ($H_{98,y,2} \approx 1.5$), and plasma current of 1.04 MA.

Two optimization cases are presented. In case 1 of Figure 1, the plasma current is constrained to be 1.04 MA from 2.5 seconds until the end of the discharge to ensure the same flattop current as that in the reference discharge 172538. Also, the neutral beam power and ECCD power is constrained to be at maximum starting at 3 seconds until the end of the discharge to force the highest possible noninductive current fraction. The results of the optimization (case 1) are shown in Figures 1(a) and 1(b). As can be seen in Figure 1(a), the optimized control policy is parameterized as a piece-wise linear function with updates every 100 ms, and bounds are applied to the actuators (dashed magenta)[3]. The loop voltage profile is simulated with the model described in Section 2. In Figure 1(b), the nominal (blue) and optimized (red) loop voltage profiles are plotted at the end of the flattop phase. Large strides are made in flattening the loop voltage profile by introducing an overshoot or wiggled current ramp-up and early timing of the ECCD turn-on and the bulk of the NBI power. In case 2 of Figure 1, the plasma current is now constrained to be 0.91 MA during the flattop. Based on predictive TRANSP runs, it was estimated a flattop current of 0.91 MA would provide an inductive current fraction of around 80-85%, which was taken as the baseline for experimental testing. The results of the optimization (case 2) are shown in Figures 1(b) and 1(c). The obtained loop voltage profile is plotted in green in Figure 1(b). Similar result to that of the original optimization case (case 1) was obtained. As shown in Figure 1(c), the primary differences being a slightly later turn-on time of ECCD and later ramp-up of the NBI power. The same characteristic current overshoot in the ramp-up phase can be noted.

4 Real-time Optimization of the q Profile Evolution

In the previous section, the simulated plasma evolution in response to the actuators was embedded into a nonlinear optimization problem that solves for a feedforward control policy (set of actuator waveforms) that under ideal conditions guides the plasma evolution to the desired state. However, mismatch between the model and real plasma dynamics and variability in the plasma startup and tokamak wall conditions amongst other things can lead to poor performance with feedforward control alone. Therefore, to account for variable plasma conditions and mismatch between the model and real system, the feedforward control is combined with feedback control.

The feedback control signal is computed by optimizing in real-time the plasma response to the available actuator set over a finite horizon (number of future control time-steps). Similar to the feedforward control design, the feedback control strategy is formulated as an optimization problem and is solved with numerical optimization techniques. The difference is that the system dynamics are linearized to make the problem strictly convex, the constraints are softened so that the problem is always feasible, and the evolution is only considered over a short horizon (small number of time steps) to reduce the size of the problem so that it is real-time solvable. The real-time optimization problem is solved repeatedly on every control time-step in a scheme known as model predictive control (MPC) or receding horizon control (RHC).

As the q profile depends inversely on the spatial derivative of the poloidal flux, we can introduce the inverse safety factor profile,

$$\iota(\hat{\rho},t) \triangleq \frac{1}{q(\hat{\rho},t)} = \frac{-\frac{\partial\psi}{\partial\hat{\rho}}(\hat{\rho},t)}{B_{\phi,0}\rho_{\mathrm{b}}^2\hat{\rho}}, \tag{9}$$

which is useful for control design purposes. Designing a feedback controller for ι instead of q avoids the nonlinearity associated with the inverse of the spatial gradient of the poloidal magnetic flux in (3).

Let $\mathbf{u}^{\mathrm{FF}}(t)$ represent the feedforward control sequence, and let $\boldsymbol{\iota}^{\mathrm{FF}}(t)$, $\mathbf{V}_{\mathrm{loop}}^{\mathrm{FF}}$, β_N^{FF} represent the corresponding feedforward output trajectory obtained from the feedforward control optimization[4]. To design a predictive feedback controller, a time varying linear model defined by matrices $\mathbf{A}_k, \mathbf{B}_k$, and $\mathbf{c}_k$ is used by taking a first order approximation to the system defined by (1)-(2) (see Figure 2(a)), which represents a good approximation of the system dynamics in a neighborhood of the feedforward trajectory.

[3]To aid convergence, the number control updates is gradually increased starting with updates every 250 ms.

[4]Again the bold variable designates discretization in space, i.e. $\boldsymbol{\iota} = [\iota_{\hat{\rho}=0.05}, \iota_{\hat{\rho}=0.01}, \ldots, \iota_{\hat{\rho}=0.95}]^T$.

To control the q profile, we consider the reference tracking problem formulated as a finite-horizon, optimal control problem. As illustrated in Figure 2(b), the feedback controller predicts the $\boldsymbol{\iota}$ profile evolution over a short horizon and updates the control action to maintain $\boldsymbol{\iota}$ on the desired evolution, which in this work is the feedforward trajectory[5]. At time k, we consider the quadratic optimization problem[6]

$$
\begin{aligned}
&\underset{\{\mathbf{u}_t^{\mathrm{FB}}\}_{t=k}^{k+N}}{\text{minimize}} \quad J_k = \sum_{t=k}^{k+N} \left\|\boldsymbol{\iota}_t - \boldsymbol{\iota}_t^{\mathrm{FF}}\right\|_{\mathbf{Q}}^2 + \left\|\mathbf{u}_t - \mathbf{u}_t^{\mathrm{FF}}\right\|_{\mathbf{R}}^2, && \text{Minimize Deviations from Desired Feedforward Evolution}\\
&\text{subject to} \quad \mathbf{x}_{t+1} = \mathbf{A}_t\mathbf{x}_t + \mathbf{B}_t\mathbf{u}_t + \mathbf{c}_t, \\
&\qquad\qquad \mathbf{y}_t = \begin{bmatrix} \boldsymbol{\iota} \\ \beta_N \\ \mathbf{V}_{\mathrm{loop}} \end{bmatrix}_t = \mathbf{h}(\mathbf{x}_t, \mathbf{u}_t), && \text{Predict State and Output Evolution with Linearized Model}\\
&\qquad\qquad \beta_{N,t} \le \beta_N^{\max}|_t, && \text{Impose MHD Stability Limit}\\
&\qquad\qquad \mathbf{V}_{\mathrm{loop},t} \in \text{ Bounds on } \mathbf{V}_{\mathrm{loop}}, && \text{Impose Flat Loop Voltage Profile}\\
&\qquad\qquad \mathbf{u}_t = \mathbf{u}_t^{\mathrm{FF}} + \mathbf{u}_t^{\mathrm{FB}} \in \mathcal{U}_t. && \text{Impose Actuator Limits - Bounds Defined by } \mathcal{U}_t
\end{aligned}
\tag{P2}
$$

Relative to the feedforward problem (P1), the feedback optimization (P2) problem considers a shorter horizon (number of future time steps), $k, k+1, k+2, \ldots, k+N$, and a linearized form of the system dynamics (poloidal flux evolution model) is used in problem (P2) to make the problem real-time solvable. The feedback control ($\mathbf{u}^{\mathrm{FB}}$) represents an update to the feedforward control ($\mathbf{u}^{\mathrm{FF}}$) with the aim of minimizing deviations from the target $\boldsymbol{\iota}$ profile evolution, $\boldsymbol{\iota}^{\mathrm{FF}}$, which describes a plasma evolution that evolves to a flat loop voltage profile as found in the feedforward control design section. Additionally, an upper limit on β_N is imposed to avoid the onset of magnetohydrodynamic instabilities (MHD), and a constraint on the loop voltage profile is included (see Section 4.1). If deviations from the desired loop voltage profile are predicted, then the feedback controller makes updates, as best as possible, to the feedforward control so as to minimize the constraint violations.

Because of limits on the actuators and other constraints, the optimal solution to problem (P2) cannot be obtained in closed-form, and the problem must therefore be solved numerically. Fortunately, this type of quadratic optimization problem involving a positive definite cost function ($\mathbf{R} > 0$ and $\mathbf{Q} \ge 0$) can be solved efficiently using active set techniques, which take advantage of the fact that the set of active 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 solution on the next control update [8].

4.1 Bounds on Loop Voltage Profile

The bound on V_{loop} of problem (P2) can be imposed in a variety of ways. For instance, a flat loop voltage profile can be characterized by a null gradient $\frac{\partial V_{\mathrm{loop}}}{\partial \hat{\rho}} \equiv 0$. Since the spatial derivative of the loop voltage profile is proportional to the time derivative of the ι profile (from (9)), a flat loop voltage can be imposed by applying the constraint $\partial\iota/\partial t = 0$, or equivalently

$$
\boldsymbol{\iota}_{t+1} - \boldsymbol{\iota}_t = 0 \quad \text{for} \quad t = k, k+1, \ldots, N-1. \tag{10}
$$

However, it certainly does not make sense to apply this constraint throughout the entire control phase. Instead there should be a trigger event or a specific time during the control phase at which this constraint is applied. For this work, we choose to apply the constraint beginning at a specified time, t_{ss}. In order to ensure the optimization problem (P2) remains feasible, the loop voltage constraint can be softened with a forgiveness parameter,

$$
-\epsilon_{ss} \le (\boldsymbol{\iota}_{t+1} - \boldsymbol{\iota}_t) \le \epsilon_{ss}, \quad \epsilon_{ss} \ge 0. \tag{11}
$$

where ϵ_{ss} represents a window on forgiveness of satisfying the loop voltage constraint. The forgiveness parameter is included as an optimization variable in the MPC problem (P2) by replacing the optimization objective with

$$
J_k \leftarrow J_k + W_{ss}\epsilon_{ss}^2, \tag{12}
$$

in which case the MPC problem will minimize violation of the loop voltage constraint, where W_{ss} is introduced as a control weight. Varying the relative size of $\mathbf{Q}$ to W_{ss} can be used to align control effort towards tracking the desired reference profile evolution $\boldsymbol{\iota}^{\mathrm{FF}}$ or satisfying the loop voltage constraint.

[5]Other target profile evolutions are indeed possible.

[6]The function $\|\boldsymbol{z}\|_{\mathbf{W}}^2$ denotes the squared weighted Euclidian norm, i.e. $\|\boldsymbol{z}\|_{\mathbf{W}}^2 = \mathbf{z}^T\mathbf{W}\mathbf{z}$.

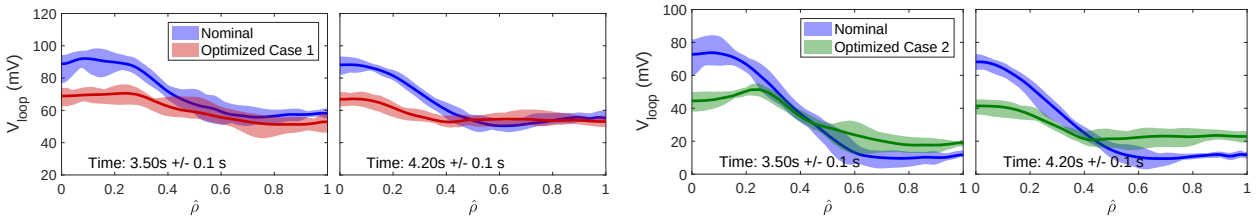

(a) TRANSP test: feedforward case 1 – flattop $I_p = 1.04$ MA. (b) TRANSP test: feedforward case 2 – flattop $I_p = 0.91$ MA.

Figure 3: Predictive TRANSP testing: the loop voltage profile obtained with optimized feedforward control policies of Figure 1 in comparison to the loop voltage obtained with a non-optimized control policy (actuator waveforms of shot 172538 labeled nominal). In (a) the optimized feedforward case 1 (red) and in (b) the optimized feedforward case 2 (green) are plotted. The associated TRANSP run identification numbers are nominal – 725381122, case 1 – 1725381608, and case 2 – 1725381614.

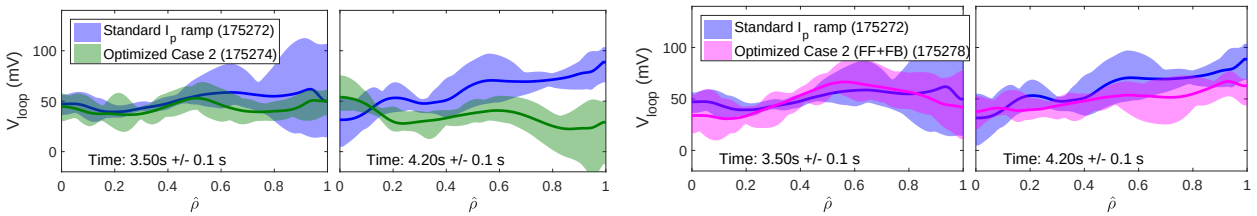

(a) Experimental test: feedforward control alone (case 2). (b) Experimental test: feedforward (case 2) + feedback control.

Figure 4: Experimental testing: the loop voltage profile obtained with optimized feedforward is compared to a standard I_p ramp (shot 175272). Two tests are shown, (a) feedforward control alone shot 175274 (green) and (b) feedforward + feedback shot 175278 (magenta). The loop voltage profile is obtained with interpretive TRANSP runs, and associated run identification numbers are standard I_p ramp – 1752722601, feedforward case 2 – 1752742601, and feedforward + feedback – 1752782601.

5 Experimental Testing

The ultimate objective of this work is to design a sophisticated control strategy that can reliably assist access to high $q_{min} > 1.5 - 2.0$ discharges in a highly noninductive regime ($f_{NI} > 90\%$) with a uniform loop voltage profile obtained as early as possible in the discharge. Shot 172538 is taken as the reference shot; a high q_{min} discharge with high triangularity and elevated energy confinement ($H_{98,y,2} \approx 1.5$), and plasma current of 1.04 MA, toroidal field $B_T = 1.65$, major radius $R = 1.66$ m, and minor radius $a = 0.61$ m. The broad current profile in 172538 is obtained with ECCD applied at $\hat{\rho} = 0.2 - 0.65$ to drive off-axis current and maintain the elevated q_{min}, and the 150° beamline is tilted off-axis by 16° to increase the off-axis current drive. The reference flattop current is scaled down to 0.91 MA for higher a noninductive fraction ($f_{NI} \approx 80 - 85\%$).

Testing begins with predictive TRANSP runs of the optimized feedforward control policies developed in Section 3. In Figure 3, predictive TRANSP runs of the feedforward optimization cases of Figure 1 are shown; case 1 (flattop I_p of 1.04 MA) is shown in Figure 3(a) and case 2 (flattop I_p of 0.91 MA) is shown in Figure 3(b). For comparison, a non-optimized case (labeled nominal) is obtained with predictive TRANSP runs of the actuator waveforms of shot 172538 with flattop I_p of 1.04 MA (same as 172538) and flattop I_p scaled down to 0.91 MA for comparison with optimized case 2. In both cases it can be noted the optimized feedforward control signals have obtained flatter loop voltage profiles.

Experimental testing of the optimized feedforward control policy (case 2) was conducted in DIII-D discharges with feedforward control alone (Figure 4(a)) and feedforward + feedback control (Figure 4(b)). The loop voltage profile obtained with optimized feedforward (shot 175274) is compared to a standard I_p ramp with β_N control (shot 175272) in Figure 4(a). Notable improvements to the flatness of the loop voltage profile are obtained in the optimized case. The loop voltage profile cannot be directly measured; instead it is found by TRANSP analysis runs. In the test demonstrating combined feedforward + feedback control, the feedforward control was once again defined by the actuator waveforms of Figure 1(c) (case 2). Recall that the control strategy involves designing a feedforward control that evolves to the desired plasma state under ideal conditions, and then using feedback control to help follow the feedforward state evolution during actual experiments, i.e. the feedback corrects for disturbances and model mismatch. Stated plainly, the feedforward control produced a set of actuators waveforms and corresponding state evolution (q evolution) that evolves to a flattop loop voltage, which serves as a path to the desired state. The feedback ensures the plasma follows that path. This is why the control strategy for loop voltage control still functions primarily as a q profile controller, only now it follows a q profile evolution that evolves to a stationary state. The V_{loop} constraint is incorporated for added insurance that a flat loop voltage can be obtained. In Figure 4(b), the results of shot 175278 are shown, in which the flat V_{loop} constraint is applied starting at $t_{ss} = 4$ s. The loop voltage profile during the flattop phase of shot 175274 is plotted in comparison to the loop voltage profile obtained with the standard I_p ramp (shot 175272). In this case, only a minor improvement in flattening the loop voltage profile can be seen around 4.2 seconds.

It should be noted that this controller represents the necessary first steps. First, improvements are being made to the model so that accurate predictions to the profile evolution over 1-2 confinement times can be made. Second, the feedback solver efficiency needs to be improved so that a long horizon time (of length 1-2 confinement times) can be used in the control design. With these improvements, the controller can make decisions in the ramp-up while considering how those control actions will effect the profiles in the flattop. Presently, the controller uses a horizon time of 5-10 control time steps (roughly 200ms). This needs to be pushed up to 1-2 seconds, enabling a controller that can impose the V_{loop} constraint through the entire discharge.

6 Conclusions

This work has presented numerical optimal control methods for combined feedforward + feedback control of the plasma q profile evolution to assist access of stationary plasmas characterized by a flat loop voltage profile. It was found, based on numerical optimization studies involving a simplified control-oriented model, that an overshoot in the target flattop current during ramp-up combined with an appropriate timing of the NBI and ECCD turn-ON times is beneficial for the development of stationary plasmas. The optimized feedforward control policy has also shown better results with respect to the non-optimized case in more comprehensive, as compared to the control-oriented model, predictive-TRANSP simulations. While experimental tests were less persuasive, it is important to note that the control approach is largely dependent on a good plasma model. The quality and reliability of both the optimized feedforward control policy and predictive feedback control will improve as more physics knowledge is embedded into the control-oriented model such as a dynamic plasma temperature model to provide better prediction of the plasma resistivity.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Awards DE-FC02-04ER54698 and DE-SC0010661. DIII-D data shown in this paper can be obtained in digital format by following the links at `https://fusion.gat.com/global/D3D_DMP`.

References

[1] Schuster E *et al.* 2017 *Nuclear Fusion* **57** 116026

[2] Xu C *et al.* 2010 *IEEE Transactions on Plasma Science* **38** 163–173

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

[4] Barton J *et al.* 2015 *Nuclear Fusion* **55** 093005

[5] Wehner W *et al.* 2017 *Fusion Engineering and Design* **123** 513 – 517

[6] Dongen J V *et al.* 2014 *Plasma Physics and Controlled Fusion* **56** 125008

[7] Ou Y *et al.* 2008 *Plasma Physics and Controlled Fusion* **50** 115001

[8] Wehner W *et al.* 2016 *2016 IEEE Conference on Control Applications (CCA)* pp 629–634

[9] Maljaars B *et al.* 2015 vol 48 pp 314 – 321 5th IFAC Conference on Nonlinear Model Predictive Control NMPC 2015

[10] Barton J E *et al.* 2013 *52nd IEEE Conference on Decision and Control* pp 4182–4187

[11] Hinton F L *et al.* 1976 *Rev. Mod. Phys.* **48**(2) 239–308

[12] Blum J 1989 *Numerical Simulation and Optimal Control in Plasma Physics: With Applications to Tokamaks* (Wiley, Paris)

[13] Barton J 2015 Phd thesis Lehigh University, USA

[14] Sauter O *et al.* 1999 *Physics of Plasmas* **6** 2834–2839

[15] Sauter O *et al.* 1999 *Physics of Plasmas* **6** 2834–2839

[16] ITER Physics Basis Editors and ITER Physics Expert Group Chairs and Co-Chairs and ITER Joint Central Team and Physics Integration Unit 1999 *Nuclear Fusion* **39** 2137

[17] Betts J T 1998 *Journal of Guidance, Control, and Dynamics* **21** 193–207

[18] Nocedal J *et al.* 2006 *Numerical Optimization 2nd edn* (Berlin: Springer)

[19] Wehner W P 2017 Phd thesis Lehigh University, USA