Princeton Plasma Physics Laboratory

PPPL-5372

Feedback control design for non-inductively sustained scenarios in NSTX-U using TRANSP

M. Boyer, R. G. Andre, D. A. Gates, S. P. Gerhardt,
J. E. Menard, and F. M. Poli

March 2017

Prepared for the U.S. Department of Energy under Contract DE-AC02-09CH11466.

Princeton Plasma Physics Laboratory
Report Disclaimers

Full Legal Disclaimer

Trademark Disclaimer

PPPL Report Availability

Princeton Plasma Physics Laboratory:

http://www.pppl.gov/techreports.cfm

Office of Scientific and Technical Information (OSTI):

http://www.osti.gov/scitech/

Related Links:

U.S. Department of Energy

U.S. Department of Energy Office of Science

U.S. Department of Energy Office of Fusion Energy Sciences

Feedback control design for non-inductively sustained scenarios in NSTX-U using TRANSP

M. D. Boyer[1], R. G. Andre[1], D. A. Gates[1], S. P. Gerhardt[1], J. E. Menard[1], F. M. Poli[1]

[1]Princeton Plasma Physics Laboratory, Princeton, New Jersey, USA

E-mail: mboyer@pppl.gov

December 2016

Abstract. This paper examines a method for real-time control of non-inductively sustained scenarios in NSTX-U by using TRANSP, a time-dependent integrated modeling code for prediction and interpretive analysis of tokamak experimental data, as a simulator. The actuators considered for control in this work are the six neutral beam sources and the plasma boundary shape. To understand the response of the plasma current, stored energy, and central safety factor to these actuators and to enable systematic design of control algorithms, simulations were run in which the actuators were modulated and a linearized dynamic response model was generated. A multi-variable model-based control scheme that accounts for the coupling and slow dynamics of the system while mitigating the effect of actuator limitations was designed and simulated. Simulations show that modest changes in the outer gap and heating power can improve the response time of the system, reject perturbations, and track target values of the controlled values.

1. Introduction

The National Spherical Torus eXperiment Upgrade facility (NSTX-U) [1], which recently completed its first campaign of plasma operation, looks to span the gap between earlier spherical torus devices, like NSTX [2] or the Mega-Ampere Spherical Tokamak (MAST) [3], and potential future facilities intended to study plasma-material interaction [4], nuclear components [5] or production of fusion power [6, 7]. NSTX-U will build upon the results of NSTX [8] to explore several issues for such future devices, including the scaling of electron transport with toroidal magnetic field and plasma current [9, 10, 11], the physics of fast particles [12, 13, 14], and the achievement and sustainment of non-inductive, high-β scenarios [15, 16, 17, 18]. The latter point is especially critical for spherical torus based designs because their compact size combined with the need for tritium breeding blankets and neutron shielding in such facilities leaves little to no room for a central solenoid to induce plasma current. The recently completed and commissioned upgrades to NSTX-U will enable the study of non-inductive scenarios, including start-up, ramp-up, and flattop current sustainment. One of the primary

components of the upgrade project was the replacement of the 'center stack' (which contains the inner-leg of the toroidal field (TF) coils, the Ohmic heating (OH) solenoid, and some divertor coils) to enable fields up to 1.0T and to provide more Ohmic flux for longer inductive discharges. The other major upgrade was the addition of a second neutral beam injector with three neutral beam sources aimed more tangentially, which significantly increases the auxiliary heating power and current drive and adds flexibility in shaping the spatial deposition of these quantities in the plasma.

Advanced plasma control will be an important tool for achieving the research goals of the NSTX-U program. The NSTX-U Control System (NCS) (which includes the real-time hardware, protection systems, and software) includes the flexible Plasma Control System (PCS) software platform provided by General Atomics that allows customized categories of various control algorithms to be developed within a powerful real-time control infrastructure. Many improvements to the NCS hardware have been made as part of the upgrade to NSTX to increase the computational power available for the PCS, reduce latency, and expand the number of diagnostics and actuators under real-time control [19, 20]. A significant amount of software development was completed to create the new control algorithms needed to optimally handle the complex dynamics of the system. Ongoing development (e.g., [21, 22]) aims to enable current and rotation profile control, power and particle exhaust control, and edge transport barrier control, building on the successful advances made during NSTX operations [23, 24, 25]. This paper extends this development by examining potential approaches to active real-time control during non-inductive scenarios using a framework for feedback control simulations in the integrated modeling code TRANSP [21]. Active control schemes will be necessary in such scenarios to tailor the response time of the discharge evolution, to enable reproducible discharges in the presence of disturbances (e.g., changing wall conditions or confinement characteristics), to enable controlled scans of physics and engineering parameters, and for active avoidance of plasma instabilities.

1.1. Organization

The paper is organized as follows: section 2 describes the TRANSP code and the modules and settings used in the simulations. Section 3 describes the actuators and measurements considered in the control design, along with the identification of a control-oriented model, and an overview of the proposed control algorithm. Results of closed-loop TRANSP simulations using the algorithm are described in section 4, followed by conclusions and plans for future work in section 5. Details of the proposed control design are provided in Appendix A.

2. Predictive simulation approach

TRANSP [26, 27] is a time-dependent integrated modeling code for tokamak discharge prediction and interpretive analysis of experimental data. Its predictive mode has been

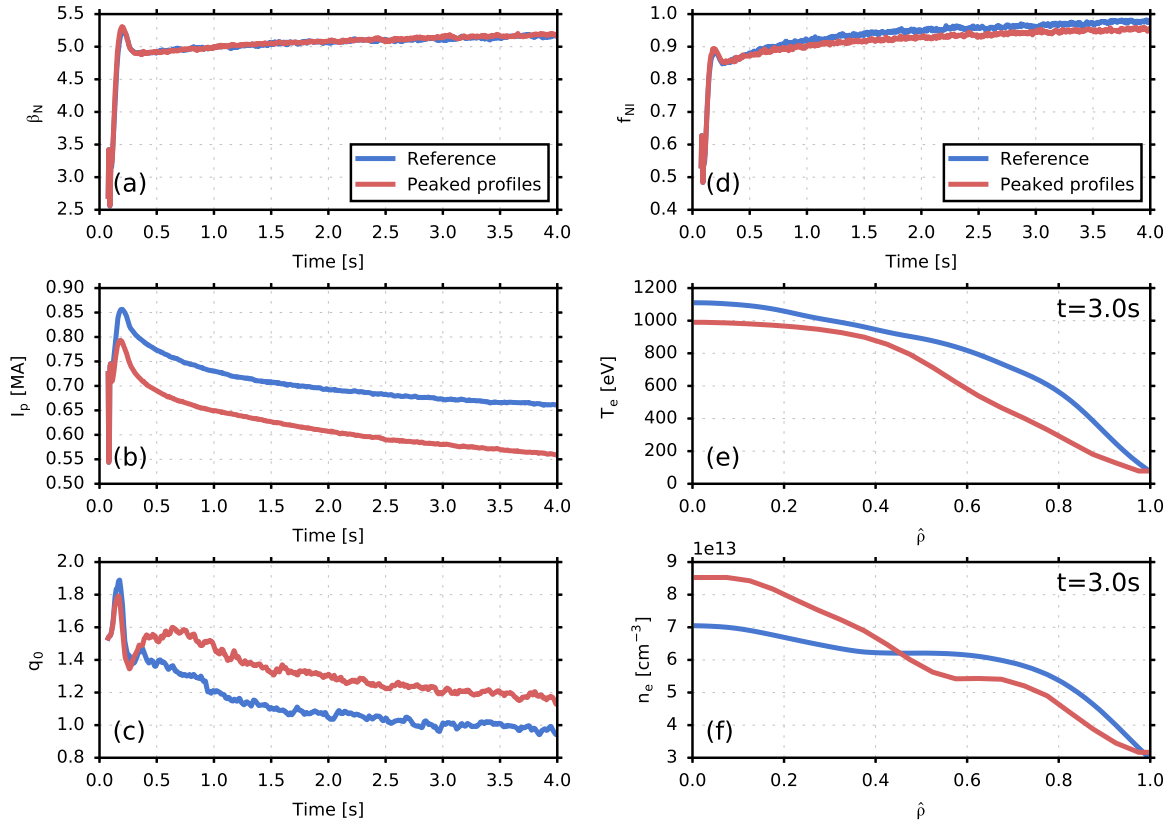

Figure 1: Comparison of (a) β_N, (b) plasma current, (c) central safety factor, (d) non-inductive fraction, (e) electron temperature profile, and (f) electron density profile during the reference simulation and the simulation with more peaked profiles.

used for scenario development on NSTX-U to explore the potential equilibrium operating space of the device, including fully non-inductive scenarios [28], and has been used to explore non-inductive plasma current ramp up [29]. Recently, the ability to include feedback control algorithms in TRANSP simulations has been developed to study control algorithms for stored energy and plasma profiles in inductive scenarios [21, 22]. The framework for feedback simulations in TRANSP uses the NUBEAM [30] module for calculating neutral beam heating and current drive, and the ISOLVER free-boundary equilibrium solver [31, 32] to evolve the discharge shape and current distribution. In this work, ISOLVER is used in a mode that chooses the coil current evolution to match a prescribed target plasma boundary shape in a least-squares sense. The Chang-Hinton neoclassical model is used to predict the ion temperature profile evolution, and the ITER-98y,2 confinement scaling expression [33] is used to constrain the electron temperature based on the TRANSP predicted volume-averaged power balance. The electron temperature profile shape is prescribed ahead of time for each simulation. The electron density is modified throughout the simulation to match a prescribed trajectory for the total particle inventory, with the shape of the density profile prescribed *a priori*.

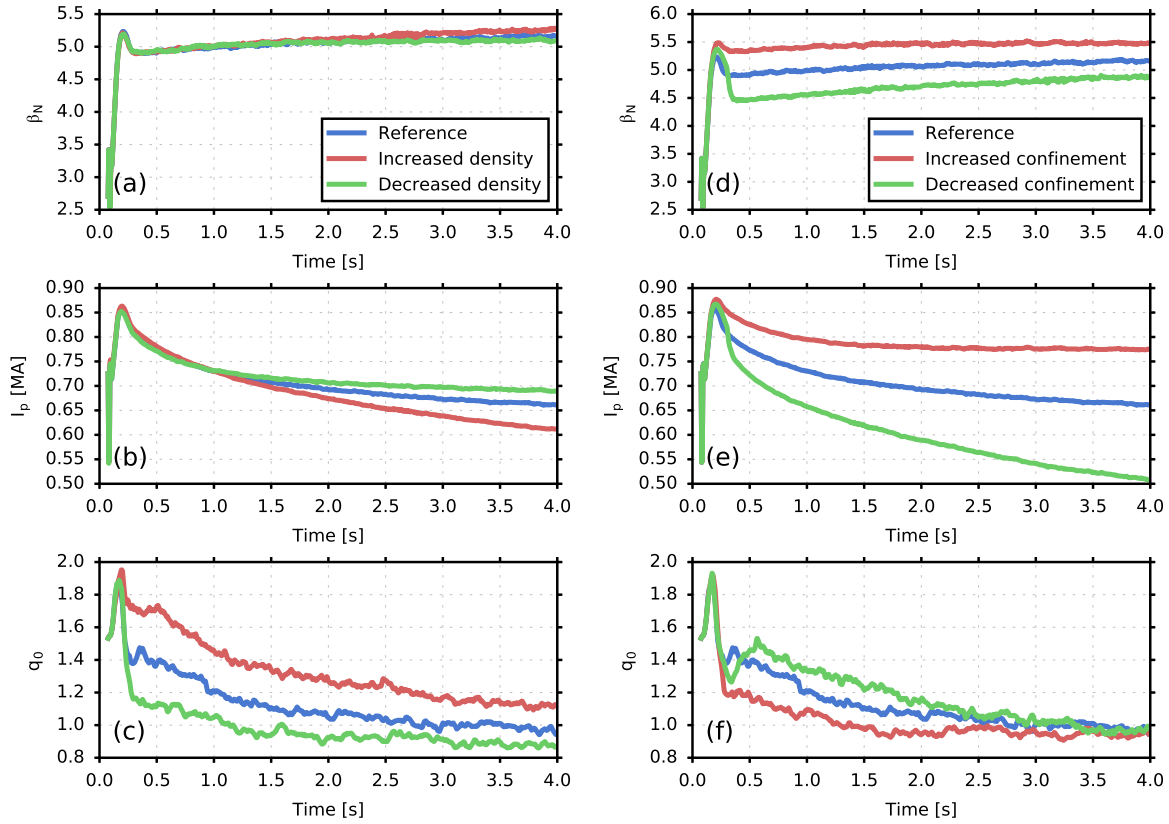

Figure 2: Comparison of β_N, plasma current, and central safety factor during density disturbances (a, b, and c) and during confinement disturbances (d, e, and f).

The ion density is calculated from quasi-neutrality assuming a flat $Z_{eff} = 2$ profile and carbon as the only impurity. While experimental studies of non-inductive start-up and ramp-up are planned, the earliest non-inductive scenario development studies on NSTX-U will likely start with an inductively formed plasma and, at some point during the shot, the Ohmic coil current will be clamped to observe the plasma behavior as it relaxes to a fully non-inductive state. This approach is mimicked in TRANSP by beginning with an inductively formed plasma and fixing the Ohmic coil current throughout the simulation starting at $t = 0.1s$. An open loop (no feedback control) simulation was performed for use as a reference throughout the rest of the study. The profile shapes used throughout the simulation were broad profiles taken from NSTX discharge 142301. Beam sources 1A (R_{tan} = 70cm), 1B (R_{tan} = 60cm), 2A (R_{tan} = 130cm), and 2B (R_{tan} = 120cm) were on throughout the simulation, the boundary shape was held fixed with a mid-plane outer gap of 15cm, and the electron inventory was held fixed at 6.65×10^{20} ($f_{GW} \approx 0.7$). Results of the reference case are shown in figure 1. During the reference simulation, the plasma slowly settled to a steady state with $\beta_N \approx 5.1$ and $I_p \approx 660kA$ (see figures 1a and 1b), taking roughly 4s to fully relax (as indicated by reaching 100% non-inductive fraction). On the same time scale, the safety factor on axis relaxed to

close to 1.0, which could potentially lead to discharge-ending MHD activity. The current redistribution time for these discharges was $\tau_{CR} \approx 1.4\frac{a^2\kappa T_e[keV]^{3/2}}{Z_{eff}} \approx 0.65s$ while the energy confinement time was $\tau_E \approx 0.03$s, indicating that the coupling of kinetic and magnetic profile dynamics results in a slowed plasma response in this scenario. The final beam-driven current fraction for this scenario was $f_{NBI} \approx 0.35$, with the rest of the current sustained by bootstrap current.

To test the sensitivity of the non-inductive scenario to changes in parameters, simulations were run with disturbances, including changes in the shape of the electron temperature and density profiles, the density magnitude, and the confinement quality. Figure 1 shows that more peaked profiles led to reduced plasma current with a slightly slower response time, higher central safety factor, and a nearly identical β_N evolution. The steady-state fractions of beam-driven and bootstrap current were approximately the same as in the reference case. Figures 2a-c show that the final value of β_N varied (slightly) proportionally to the applied density perturbations of $+15\%$ and -10%, while the current was reduced as the density increased. The central safety factor elevated with increased density, but dropped below 1.0 around 2 seconds faster than the reference case with reduced density. The steady-state fraction of beam-driven current was found to increase to $f_{NBI} \approx 0.39$ in the case of decreased density. Figures 2d-f show that the applied confinement increase ($+10\%$) led to increased β_N and I_p and a faster settling time. The final value of q_0 was nearly unaffected, however, it settled much more quickly with increased confinement. Decreased confinement (-10%) resulted in lower β_N and I_p, a slower response time, and no effect on the final value of q_0. Decreasing confinement slightly decreased the steady-state beam-driven current fraction to $f_{NBI} \approx 0.32$. These simulations indicate that, given a desired scenario, disturbances could lead to significant changes in performance or MHD-shortened discharges. This motivates development of feedback control algorithms to reject such disturbances and recover, as closely as possible, the reference evolution. Because the open-loop response time of the discharge is comparable to the discharge limit dictated by coil heating and limits on neutral beam pulse length, the ability of feedback to improve the response time and track requested target scenarios will be important for efficient use of experimental time.

3. Feedback control approach

3.1. Actuators and measurements

The actuators considered for control of non-inductive scenarios in this work are the six neutral beam sources and the plasma boundary shape. The central solenoid current is not considered as an available actuator in this work to mimic future solenoid-free devices. The neutral beam sources, three of which are new for NSTX-U, were designed to enable the current drive deposition and heating to be tailored in real-time. Initial experimental results on NSTX-U have provided empirical evidence that the new sources produce a different plasma response compared to the existing sources, and that they

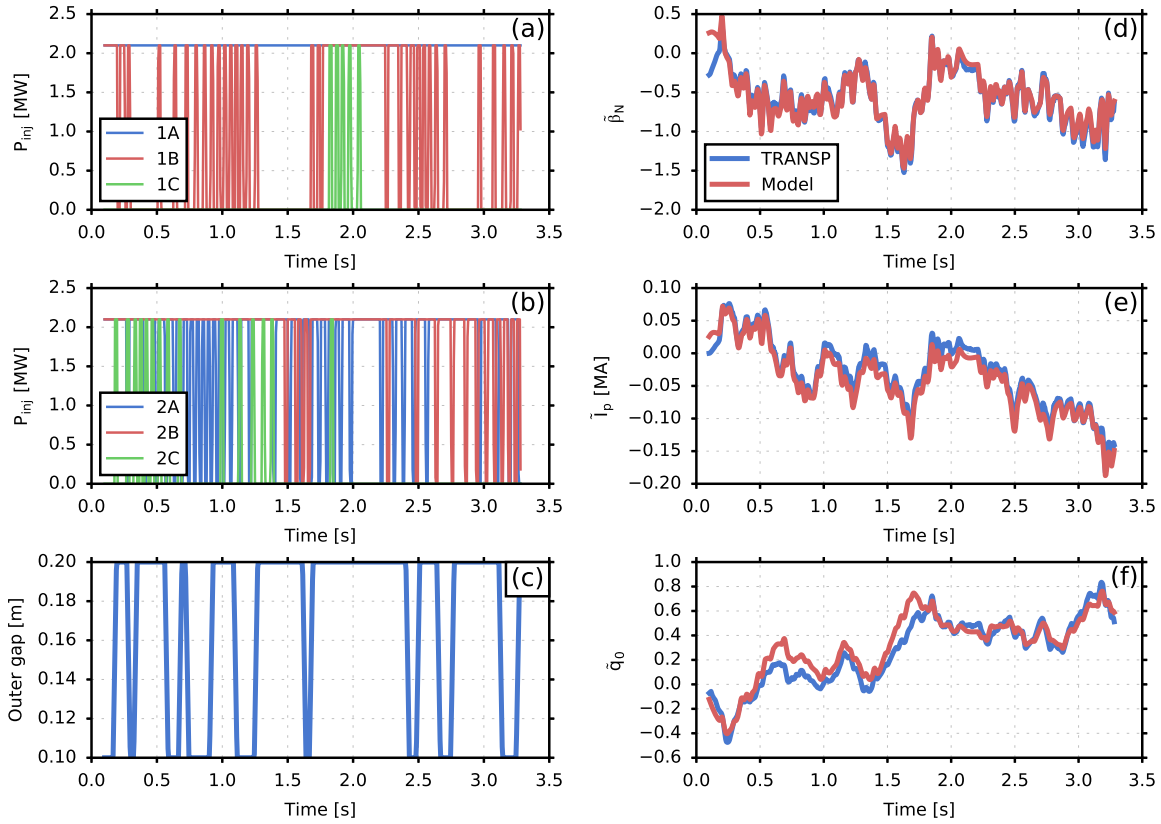

Figure 3: Modulations of (a) beam line 1 source powers, (b) beam line 2 source powers, and (c) the mid-plane outer gap during simulation for validating identified model. Comparison of deviation from the reference values of (d) β_N, (e) plasma current, and (f) central safety factor during TRANSP simulation to the predictions of the identified linear model.

can provide off-axis deposition [34, 35, 36, 37]. The latter was shown by the ability to change the propagation direction of Toroidal Alfven Eigenmodes (an indication of producing a hollow fast-ion beta profile) by injecting source 2A ($R_{tan} = 130$cm) [34, 37]. The primary plasma boundary shape parameter that was considered in this work was the mid-plane outer gap. Two target boundaries, one with a small outer gap and the other with a large outer gap, were chosen as references. Based on the requested outer gap from the feedback controller, the target boundary used by ISOLVER to determine the coil currents was interpolated between the two reference boundaries. Increasing the size of the outer gap changes shaping parameters in such a way that bootstrap current is increased and moves the neutral beam deposition further off-axis, resulting in an increase in the central safety factor. Due to the strong coupling between kinetic and magnetic profile dynamics in non-inductive scenarios, varying any of these actuators during the discharge can alter the plasma current, stored energy, and central safety factor, which were taken as the to-be-controlled variables in this work. These values were chosen due

to their importance to scenario development and their effect on plasma performance and stability. Estimates of these quantities will be available in real-time from an instance of the real-time plasma equilibrium reconstruction code, rtEFIT, incorporating both external magnetic and motional Stark effect measurements.

3.2. Control-oriented model identification

To understand the response of the controlled variables to the actuators and to enable the systematic design of real-time control laws, TRANSP simulations were done in which the actuators were modulated around the values used in the reference simulation, and a linear dynamic response model was fit to the resulting data. The modulation pattern was formed by switching the actuators between their minimum and maximum allowed values at randomized times to create an information-rich dataset for identification. The order of the identified model (the number of states of the system) was chosen by comparing the prediction error of models of different orders on a separate validation simulation (i.e., one not used in the fitting procedure). Figures 3a-c show the beam modulations and outer gap modulations used in one of the validation simulations. Figures 3d-f compare the deviation of the TRANSP outputs during the modulated simulation from those obtained in the reference simulation, and the prediction of these deviations based on the identified linear model with 13 states. Evidently, the simplified model captures the dominant dynamics of the system well enough for use in control design and initial testing of algorithms.

To get a sense of the system coupling and the effect of actuator constraints on the achievable outputs, the steady-state output for a collection of different combinations of the considered actuators (within their constrained ranges) was calculated. The result is shown in figure 4a. While each output has a reasonably large achievable range, it can be seen that coupling and actuator constraints limit the ability to independently control the outputs, e.g., it is not possible to achieve the maximum β_N and maximum q_0. Also, it is evident that for a given value of I_p, the achievable values of β_N and q_0 are approximately constrained along a line. Figure 4b shows the output response generated by applying the maximum and minimum values of each actuator to the system independently, enabling a comparison of the relative magnitudes and the directions of influence. A filled bar represents the influence of the maximum value of the actuator, while an empty bar represents the influence of the minimum value of the actuator. Neutral beam sources were either on at full power or off in the reference shot, so they are essentially unidirectional actuators. Neglecting the small effect they have on q_0, sources 1C ($R_{tan} = 50$cm) and 1B have very symmetric influence on the outputs. Likewise, 2C ($R_{tan} = 110$cm) and 2B are approximately symmetric. 2A affects each of the outputs in the same direction and is the source with the most effect on q_0. For this reason, throughout the control design process, the beam sources are split into three groups: 1) 1C and 1B, 2) 2C and 2B, and 3) 2A, and the power request for each group of sources is taken as the manipulated variable, reducing the total number of manipulated

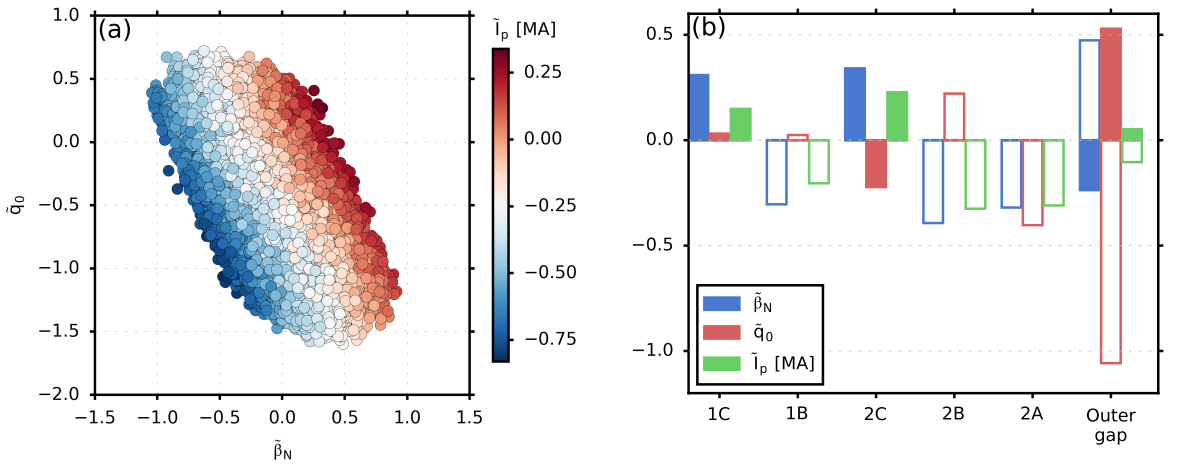

Figure 4: (a) Achievable steady-state output range considering constraints on actuators and (b) influence of applying maximum (filled) and minimum (empty) actuator values on steady-state output.

variables to four. While this reduces the degrees of freedom and the control design approach proposed in this work can handle arbitrarily many actuators, the difference between these degrees is evidently small and grouping the actuators is beneficial from the perspective of beam operation. Experimentally, modification of the injected power from a source is achieved by rapidly modulating the source on and off using a pulse-width-modulation (PWM) algorithm, which can lead to fatigue and failure of beam line components. By grouping sources with similar effects on the plasma together, the requested injected power for a particular group can be achieved by modulating just one of the sources in the group, whereas treating the sources individually may result in simultaneous modulation of multiple sources. Finally, the outer gap is evidently very effective at modifying q_0, about as effective as the individual beam sources at changing β_N, and has only a small influence on I_p.

3.3. Control-design overview

The control design approach proposed in this work is a model-based multi-input-multi-output (MIMO) scheme. It embeds the identified dynamics of the system in the control law to account for the coupling and multiple time scales, while also mitigating the effects of actuator saturation on the performance of the closed-loop system. The proposed scheme includes four main parts:

(i) A dynamic observer to estimate the unmeasured states of the identified model as well as unmodeled disturbances (assumed to be constant for the purpose of design).

(ii) A feedforward compensator to calculate adjustments to a reference actuator trajectory to track the operator-provided target values of the plasma parameters as closely as possible. The limits of the actuators and the disturbances estimated by the observer are taken into consideration (targets are assumed to be constant

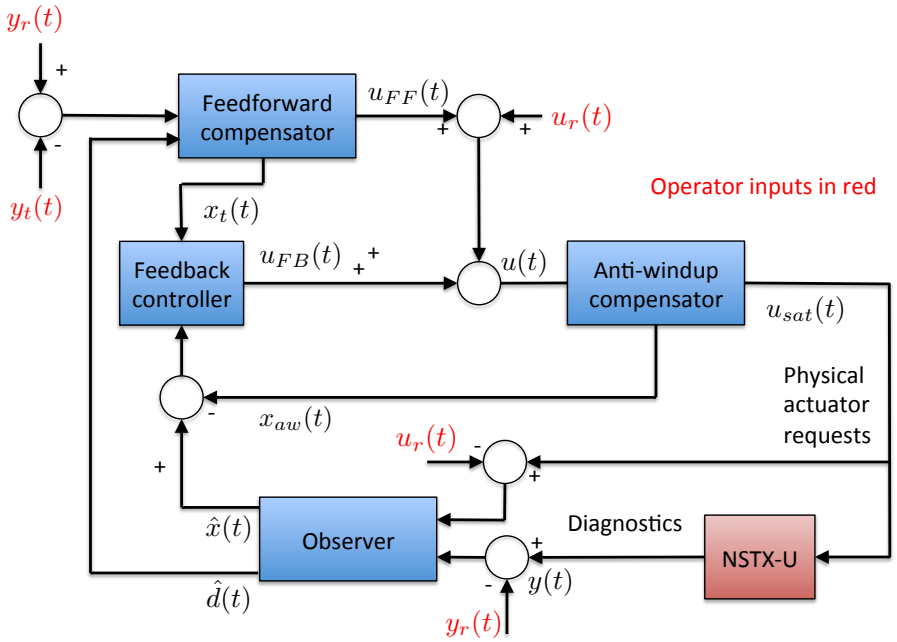

Figure 5: Schematic of multi-input multi-output control algorithm including observer, feedforward, feedback, and anti-windup components.

offsets from the reference trajectory for the purpose of design).

(iii) A state-feedback control law designed using the linear-quadratic-regulator approach to improve the response time of the system.

(iv) An anti-windup scheme to limit the effect of actuator saturation on the feedback portion of the controller.

A schematic of the proposed scheme, which is described in detail in Appendix A, is shown in figure 5. While more complex than an empirical proportional-integral-derivative (PID) control approach, tuning the proposed approach is expected to be more intuitive for operators. They only need to provide a reference shot, target outputs, and physically-meaningful relative weights determining importance of tracking each controlled quantity and penalizing the use of each actuator.

4. Feedback control simulation results

Initial closed-loop (controlled) simulations were performed using the identified state-space plasma response model to test the system response and tune the design parameters of the control algorithm, e.g., the output and actuator weightings. The resulting controller was then tested in fixed-Ohmic-coil-current TRANSP simulations to assess its robustness to the increased complexity of the model and to study the ability of the controller to track targets or to handle disturbances in profile shapes, density, and confinement. Although in actual NSTX-U experiments modification of the power injected by each neutral beam source is achieved through a pulse-width-

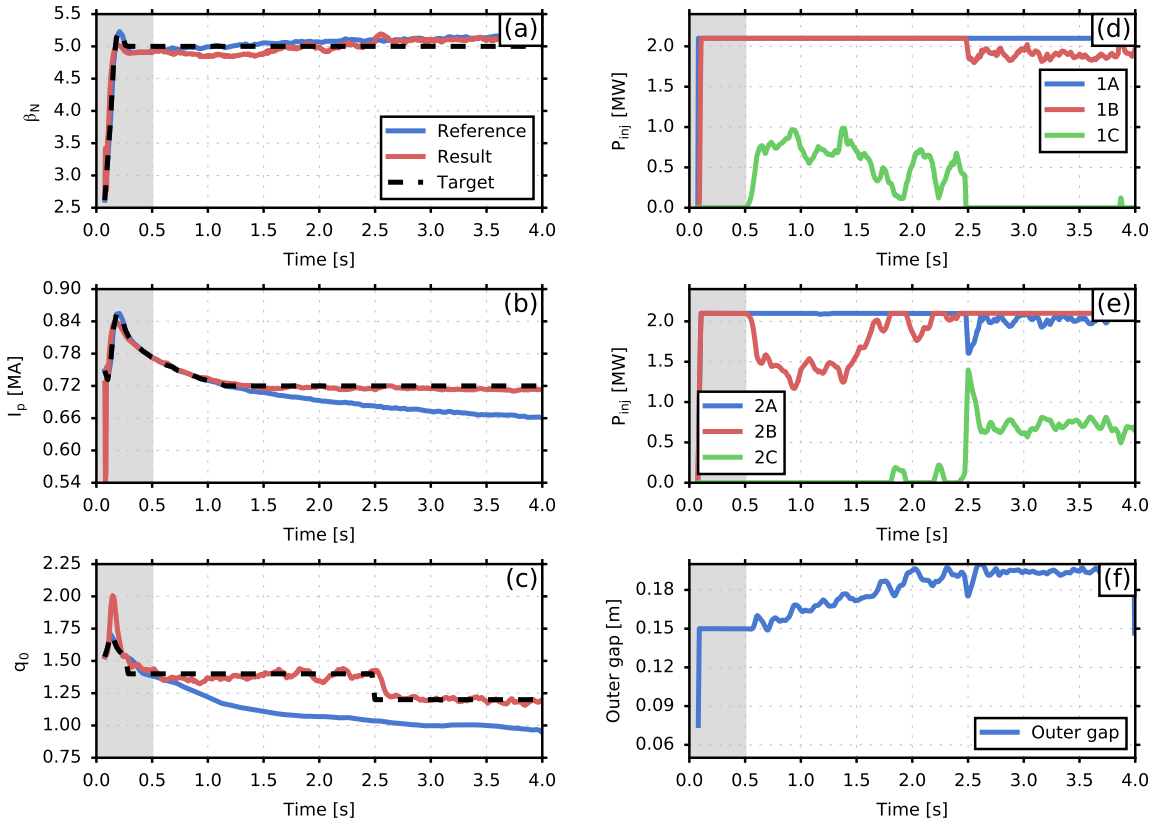

Figure 6: Comparison of reference (blue-solid), programmed target (black-dashed), and resulting evolution with controller active during simulation weighting increased I_p and q_0 targets for (a) β_N, (b) I_p, and c) q_0. Requested source powers for (d) beam line 1 and (e) beam line 2, along with (f) the requested outer gap. Feedback control was activated after 0.5s (unshaded region of plot).

modulation (PWM) scheme, preliminary simulations approximated the source behavior with continuously variable power requests. Later simulations included the PWM algorithm to test the effect of modulations on the performance of the close-loop system. During testing, actuator constraints were found to limit the possible controllable range of plasma parameters when weighting all output quantities roughly equally. However, the optimal control strategy proposed makes it possible for the designer or operator to adjust output weighting to ensure that the quantities that are most critical to a particular experiment are most tightly controlled, even when actuator saturation occurs. In the simulations presented in this section, the outputs q_0 and I_p were weighted more heavily than β_N. In each simulation, the control algorithm was activated after $t = 0.5$s so the effect of uncontrolled disturbances can be seen in the first interval of each simulation (denoted by gray shading in the figures in this section) and the performance of the controller can be seen in the remainder of the simulation.

4.1. Target tracking

Time-dependent results of a closed-loop simulation in which the targets for the output quantities differed from the reference evolution are shown in figure 6. The target for β_N was held fixed at $\beta_N = 5.0$ throughout the controlled interval of the simulation. The I_p target followed the reference shot early in the simulation, but was clamped at $I_p = 0.72$MA at $t \approx 1.2$s. The target for q_0 was held fixed at $q_0 = 1.4$ until $t = 2.5s$ at which point it was stepped down to $q_0 = 1.2$. The results show that the controller, which was tuned to weight β_N less heavily than the other quantities, can improve the response time of the system and quickly transition between different targets. This capability will enable precise experimental scans of parameters during a single discharge. The actuator trajectories show that during the interval $0.5s < t < 2.5s$ the controller responds to the requested targets by increasing the outer gap and replacing power from source 2B with power from source 1C, resulting in an increase of q_0 and I_p at roughly constant β_N. After $t > 2.5$s the 2C source power was increased to decrease q_0 and maintain elevated I_p.

4.2. Disturbance rejection

Results of a simulation in which the profile shapes were replaced with more peaked profiles and the controller targets were set to match the reference evolution are shown in figure 7. During the initial uncontrolled phase (denoted in gray), the system follows the trends shown in figure 1, with increased q_0, decreased I_p and a β_N evolution similar to the reference. The controller, which again weighted q_0 and I_p more heavily than β_N, increased I_p by increasing source 2C power. The β_N and q_0 evolutions were kept close to their targets through decreasing 1B and 2A and increasing the outer gap.

Results of a simulation in which a -10% density disturbance was applied are shown in figure 8. During the initial uncontrolled phase (denoted in gray), the system follows the trends shown in figure 2(a-c), with decreased q_0, and an evolution of I_p and β_N similar to the reference. The controller increased q_0 by decreasing source 2B power, offsetting this change by increasing 2C power to keep β_N and I_p close to their targets. The outer gap was not changed significantly in this case. Note that in the uncontrolled case shown in figure 2(a-c), q_0 drops below $q_0 = 1.0$ shortly after $t = 1.1$s, while the controller is able to maintain the reference evolution of q_0, which is above 1 until after 2.5s.

Results of a simulation in which a +10% confinement disturbance was applied are shown in figure 9. The PWM algorithm used in the actual experiment to vary the inject source power was used during this simulation. This algorithm modulates the sources to achieve the duty cycle needed to achieve the requested injected power, while maintaining a minimum source on/off time of 20ms (recall from section 3.2 that the controller was designed to output power requests for three groups of sources, chosen based on having similar effects on the plasma, rather than individual source powers to reduce the number of modulations from the PWM algorithm.) In the plots of beam power, figures 9d-e, the

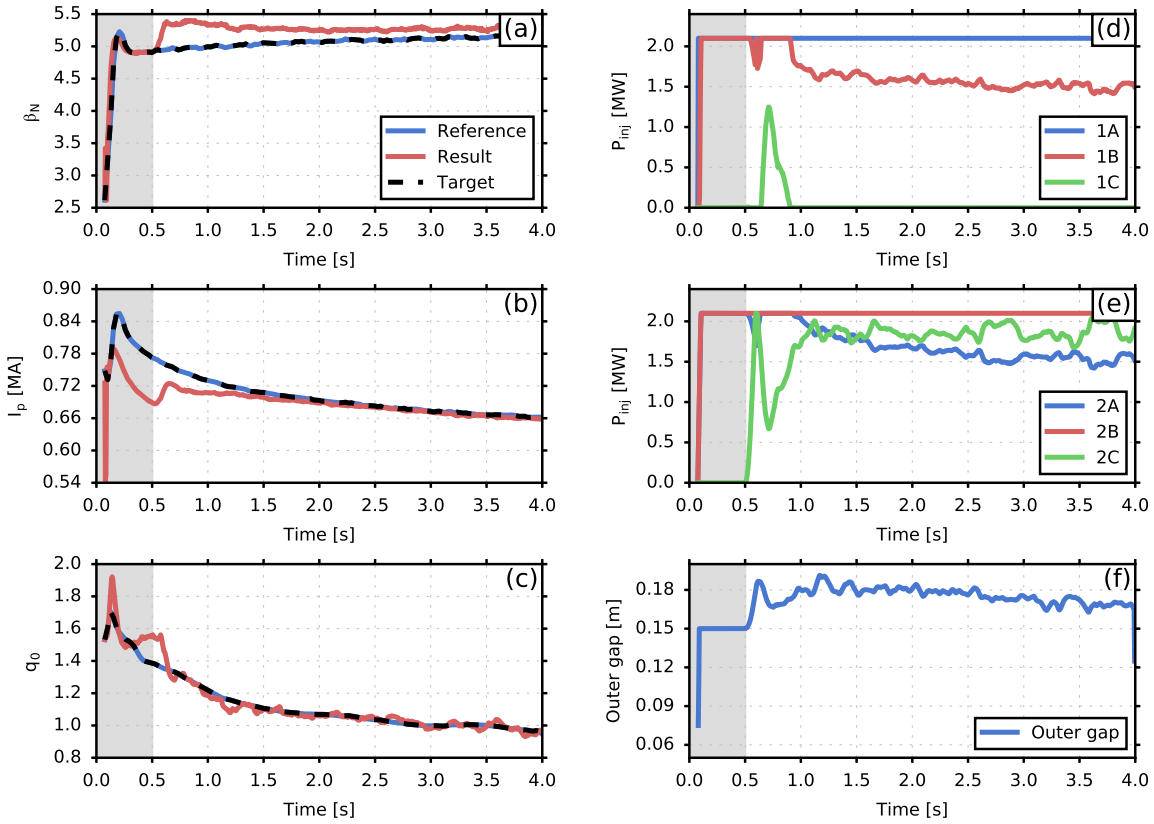

Figure 7: Comparison of reference (blue-solid), programmed target (black-dashed), and resulting evolution with controller active during simulation with a peaked profiles for (a) β_N, (b) I_p, and c) q_0. Requested source powers for (d) beam line 1 and (e) beam line 2, along with (f) the requested outer gap. Feedback control was activated after 0.5s (unshaded region of plot).

thin lines represent the applied power waveform resulting from the PWM algorithm, while the thick lines depict a time-averaged version of the applied power waveform so that trends can be visualized. During the initial uncontrolled phase (denoted in gray), the system follows the trends shown in figure 2(d-f), with increased β_N and I_p and decreased q_0. The controller increased q_0 and decreased β_N and I_p by decreasing injected power from sources 1B, 1C, 2A, and 2B, and decreasing the outer gap. In the uncontrolled case shown in figure 2(d-f), q_0 drops below $q_0 = 1.0$ shortly after $t = 1.5$s, while the controller is able to maintain the reference evolution of q_0, which is above 1 until after 2.5s. The modulation of the beam sources results in small oscillations in the I_p and q_0 evolution. The oscillations in β_N are larger because the minimum on/off time of the sources (20ms) is not much smaller than the energy confinement time ($\approx$ 30ms). The oscillations in the outputs lead the actuator requests to oscillate, occasionally resulting in spurious beam pulses, like the two pulses of 1C during the simulation. Because modulation of the beam sources can contribute to increased fatigue

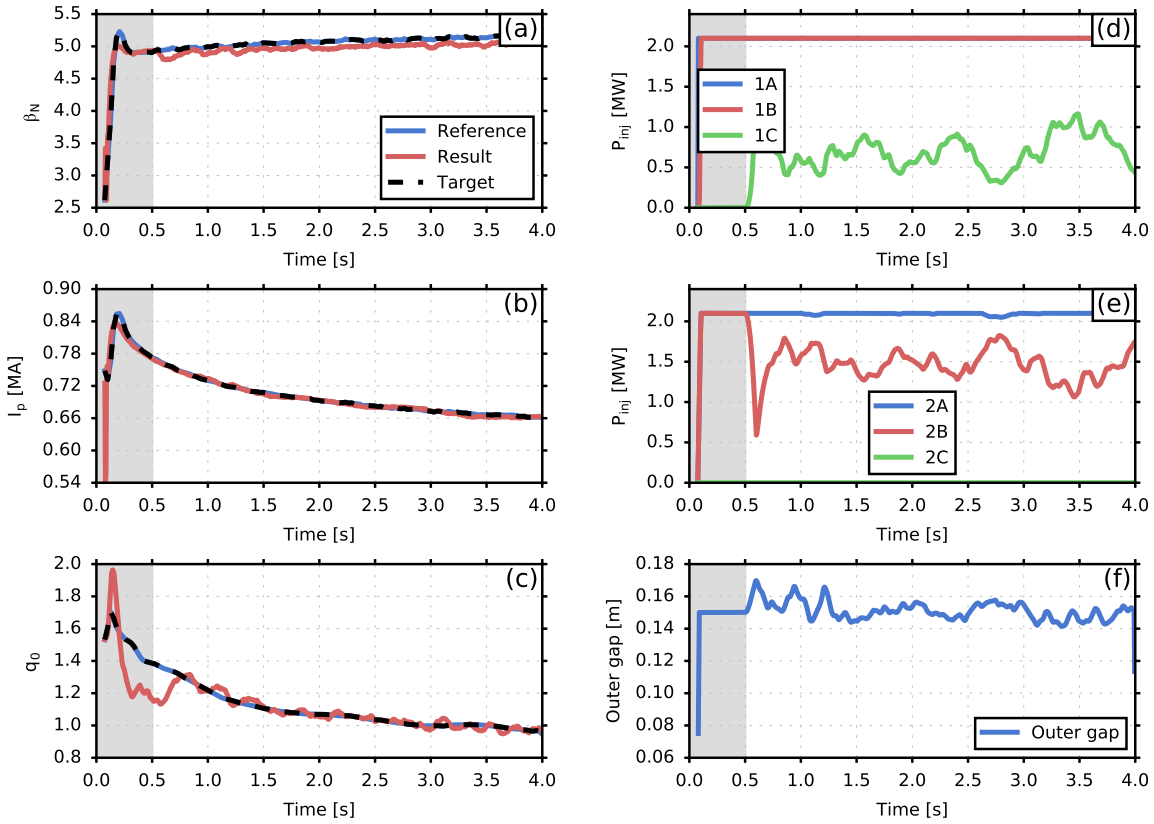

Figure 8: Comparison of reference (blue-solid), programmed target (black-dashed), and resulting evolution with controller active during simulation with a -10% density disturbance for (a) β_N, (b) I_p, and c) q_0. Requested source powers for (d) beam line 1 and (e) beam line 2, along with (f) the requested outer gap. Feedback control was activated after 0.5s (unshaded region of plot).

of beam line components, it is important to explore methods for reducing spurious pulses. One such method is to include a deadzone in the PWM algorithm that forces the source off if the requested power is less the $f_{deadzone} \times P_{source}$ where $f_{deadzone}$ is the deadzone design parameter and P_{source} is the source power, and the source is forced on if the requested power is above $(1 - f_{deadzone}) \times P_{source}$. The result of applying this approach to the previous simulation scenario with $f_{deadzone} = 0.3$ is shown in figure 10. This approach results in a small offset in the output parameters (since the deadzone alters the applied power away from the controller requested values when active) but reduces the total number of times a source is turned off from 120 to 93 (22.5% fewer). Along with the output offset, this approach also leads to somewhat larger, slower oscillations in the outputs since the deadzone causes the applied power to be, in the worst case, $(100 \times f_{deadzone})$% higher/lower than the controller requested power.

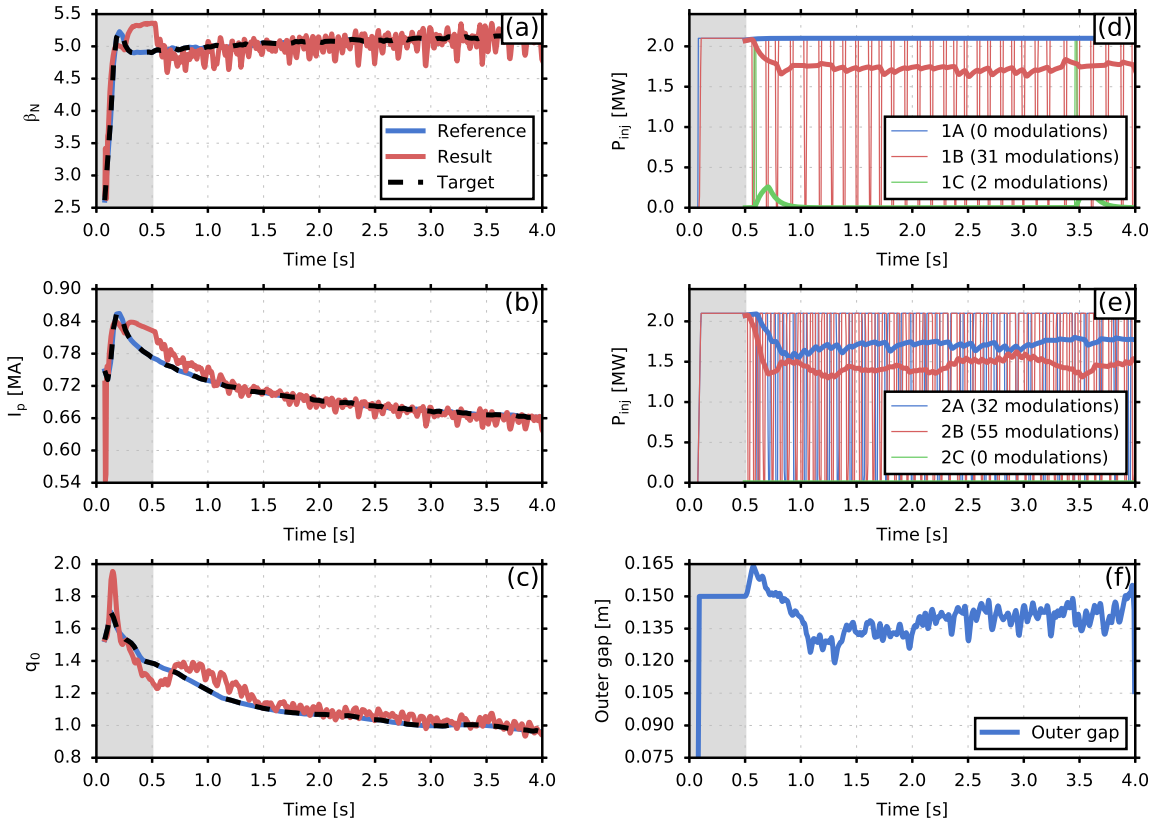

Figure 9: Comparison of reference (blue-solid), programmed target (black-dashed), and resulting evolution with controller active during simulation with a +10% confinement disturbance and beam modulations activated for (a) β_N, (b) I_p, and c) q_0. Applied source power waveforms (thin lines) and time-averaged power (thick lines) are shown for (d) beam line 1 and (e) beam line 2, along with (f) the requested outer gap. Feedback control was activated after 0.5s (unshaded region of plot).

4.3. Discussion

Based on the simulation results presented in this section, the proposed approach to controlling non-inductive scenarios on NSTX-U shows good performance in a variety of situations. The algorithm was able to speed up the response time of the system and quickly achieve steady I_p and q_0 values in the tracking case, and was able to recover the nominal reference evolution of the output quantities during changes in profile shapes, density, and confinement. These results were all obtained with the same choice of design parameters in the control algorithm, demonstrating that the performance of the controller is robust to changes in plasma parameters. It was found in [21] that good closed-loop performance could be obtained using either PID and MIMO control strategies in the case of controlling β_N and q_0 in inductive scenarios with total beam power and outer gap size as actuators. However, the increased coupling, additional actuators considered (individual source powers), and the need to control I_p as part of

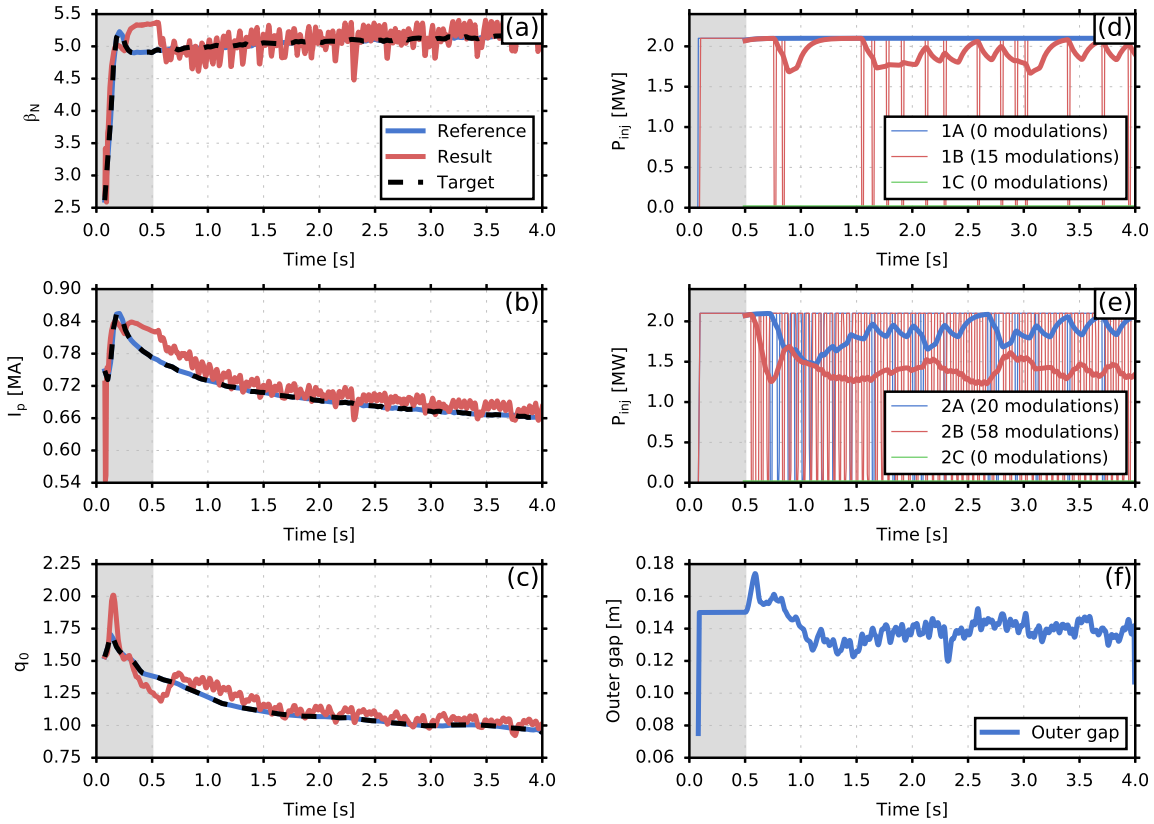

Figure 10: Comparison of reference (blue-solid), programmed target (black-dashed), and resulting evolution with controller active during simulation with a +10% confinement disturbance and beam modulations activated with a deadzone parameter of $f_{deadzone} = 0.3$ for (a) β_N, (b) I_p, and c) q_0. Applied source power waveforms (thin lines) and time-averaged power (thick lines) are shown for (d) beam line 1 and (e) beam line 2, along with (f) the requested outer gap. Feedback control was activated after 0.5s (unshaded region of plot).

the control scheme in this work made the use of MIMO model-based techniques critical. The proposed approach enables systematic design of the control algorithm, considers the coupling and actuator constraints explicitly, and provides the operator with an intuitive way to tune the algorithm through physically-motivated weight parameters (as opposed to a large number of PID gains). The approach is also more amenable to adding more outputs and actuators as the needs and capabilities of the device change. Finally, the approach is shown to be robust to the PWM scheme used to achieve requested source powers, though the modulations cause oscillations in the outputs and the actuator trajectories requested by the controller can lead to a large number of modulations. Along with requesting injected powers for groups of sources with similar effects on the plasma instead of each individual source, a deadzone modification to the PWM algorithm was proposed to reduce the number of times sources are switched on and off, though at

the cost of potential output tracking offsets and slightly larger oscillations. In order to achieve offset free, oscillation free control of the outputs, additional continuously varying actuators would need to be considered, e.g., additional shaping parameters, RF heating, or real-time variation of the beam source voltages.

5. Conclusion

In this paper, a novel approach to controlling β_N, I_p, and q_0 in non-inductive scenarios in NSTX-U, in which the outer-midplane wall gap and individual beam powers serve as manipulated variables, has been proposed. The predictive simulation capabilities of TRANSP were used as both a modeling tool to develop a control-oriented model of the response of the outputs to the actuators, and a test bed for studying the performance of the proposed control algorithm in closed-loop simulations. The proposed control scheme extends the classic linear-quadratic-integral (LQI) control approach, which was applied to related problems in [38, 21] and [22], to explicitly include actuator constraints and avoid the problem of integrator wind-up, without the computational complexity of model-predictive-control (MPC) [39, 40, 41]. While MPC has the advantage of actively avoiding time-varying state constraints in addition to actuator constraints, the increased computational burden of MPC makes implementation on smaller tokamaks with fast time scales a challenge. Like LQI and MPC, the design parameters are physically-motivated weights, making the scheme intuitive for operators to tune. In future work, the approach will be applied to controlling multiple measurements of the safety factor and rotation profiles in both inductive and non-inductive scenarios. The effect of the control scheme on plasma stability and controllability will also be explored, e.g., the effect of modifying the outer gap on coupling to the passive plates and, subsequently, vertical stability will be studied. The proposed control algorithm, along with the approach of using TRANSP as a control-oriented modeling tool and algorithm testbed, is also planned to be applied to other machines, including DIII-D, MAST-U, KSTAR, EAST, and ITER.

Acknowledgements

This work was supported by the US Department of Energy Grant under contract number DE-AC02-09CH11466.

Appendix A. Control algorithm design details

Defining the actuator and measurement trajectories from the reference simulation described in section 2 as u_r and y_r, respectively, the linear model identified in section 3.2 for the deviation of the measured values, $\tilde{y} = y - y_r$ in response to the deviation of the actuators $\tilde{u} = u - u_r$, is given by

$$\dot{x} = Ax + B\tilde{u},$$

$$\tilde{y} = Cx, \tag{A.1}$$

where A, B, and C are the identified matrices and x is the state vector of the identified system. The controller requests the actuator deviation based on the sum of the feedforward, feedback, and anti-windup compensator calculations, i.e.,

$$\tilde{u} = u_{ff} + u_{fb} + u_{aw}. \tag{A.2}$$

Appendix A.1. Observer design

Because the identified model contains a number of unmeasured states, a full-state feedback control approach, like the standard linear-quadratic-regulator used in this work, must be complemented with a dynamic observer that forms an estimate of the state vector, i.e., $\hat{x}$. To ensure that the outputs of the observer $\hat{y}$ converge to the measured values $\tilde{y}$ even in the presence of unmeasured disturbances, a disturbance estimate $\hat{d}$ is augmented to the typical observer design to form a proportional-integral observer [42], i.e.,

$$\begin{aligned}
\dot{\hat{x}} &= A\hat{x} + B\tilde{u} + K_I C^T \hat{d} - L_x\left(\hat{y} - \tilde{y}\right), \\
\dot{\hat{d}} &= -D_f I \hat{d} - L_d\left(\hat{y} - \tilde{y}\right), \\
\hat{y} &= C\hat{x},
\end{aligned} \tag{A.3}$$

where K_I is an integral gain matrix (design parameter), D_f is a fading coefficient (design parameter) that can be used to improve transient response, and L_x and L_d are gain matrices (design parameters) that inject the error between the predicted and measured outputs of the system. The augmented system can be written as

$$\begin{aligned}
\dot{x}_a &= A_a x_a + B_a u_a - L(\hat{y} - \tilde{y}), \\
\hat{y} &= C_a x_a,
\end{aligned} \tag{A.4}$$

where

$$A_a = \begin{bmatrix} A & K_I C^T \\ 0^{n_m \times n_s} & -D_f I^{n_m \times n_m} \end{bmatrix}, \tag{A.5}$$

$$B_a = \begin{bmatrix} B \\ 0^{n_m \times n_a} \end{bmatrix}, \tag{A.6}$$

$$C_a = \begin{bmatrix} C & 0^{n_m \times n_m} \end{bmatrix}. \tag{A.7}$$

The system is assumed to be influenced by a zero-mean Gaussian white process noise $w \in \mathbb{R}^{n_a + n_m \times 1}$ through the matrix

$$B_w = \begin{bmatrix} B & 0^{n_s \times n_a} \\ 0^{n_m \times n_a} & I^{n_m \times n_m} \end{bmatrix}, \tag{A.8}$$

as well as a zero-mean Gaussian measurement noise $v \in \mathbb{R}^{n_m \times 1}$. The covariance of these noise processes W and V, respectively, are design parameters chosen based on expected noise values and/or desired response that are used to adjust the relative weight of model predictions and measurements in the state estimation result. The choice of L that minimizes the steady-state covariance of state estimation error (the Kalman gain) is found by setting $L = PC_a^T V^{-1}$ where P is a positive-definite, symmetric matrix solution to the algebraic Riccati equation $A_a P + P A_a^T - P C_a^T V^{-1} C_a P + B_w W B_w^T = 0$. A comparison of the outputs predicted by the observer and the measured values corrupted by large amplitude synthetic noise is shown in figure A1(a-c) in a simulation of the identified model in which large disturbances (unknown to the observer) were added to the system. The observer estimates of the disturbances converged over time, causing the output predictions track the mean of the measured values, while the measurement noise was rejected from the estimates.

Appendix A.2. Feedforward design

To track targets, y_t, and reject the disturbances estimated by the observer, a feedforward actuator term is designed by finding the actuator values u_{FF} minimizing a cost function that weights the steady-state tracking error of the observer system $e_{ss} = y_r + \hat{y}_{ss} - y_t$, where $\hat{y}_{ss} = -CA^{-1}\left(Bu_{FF} + K_I C^T \hat{d}_{ss}\right)$, as well as the deviation from the reference actuator trajectories, u_{FF} (this weight can be used to preserve linearity as much as possible and to limit the steady-state use of particular actuators):

$$J_{ss} = \frac{1}{2} e_{ss}^T Q_{ss} e_{ss} + \frac{1}{2} u_{FF}^T R_{ss} u_{FF}, \tag{A.9}$$

where Q_{ss} is a positive-definite-symmetric weight on the steady-state error, R_{ss} is a positive-definite-symmetric weight on the steady-state actuator deviations. Under the assumption of approximately constant disturbances, the estimate $\hat{d}$ will converge to a steady-state value $\hat{d}_{ss}$ over time. Since the value of $\hat{d}_{ss}$ is not known *a priori*, the value of $\hat{d}$ at the current time is used as an estimate in the calculation of $\hat{y}_{ss}$ and e_{ss}. This constrained optimization problem is solved with a gradient projection method in this work, though a number of other optimization algorithms could be considered. The output of the feedforward compensator is u_{FF}, along with the associated target states for the system,

$$x_t = -A^{-1}\left(Bu_{FF} + K_I C^T \hat{d}\right). \tag{A.10}$$

A comparison of the programmed target outputs to the optimal achievable steady-state targets calculated by the feedforward compensator ($= Cx_t$), along with the output evolution achieved when the feedforward actuation is applied to the identified response model is shown in figure A1(d-f). The achievable steady-state output tracked the programmed targets during the first 2.0s, however, as higher I_p and q_0 were requested after $t > 2.0s$, the requested target left the achievable range, as can be noted by the difference between the feedforward target and the targets for β_N and, after $t = 3.5s$, for

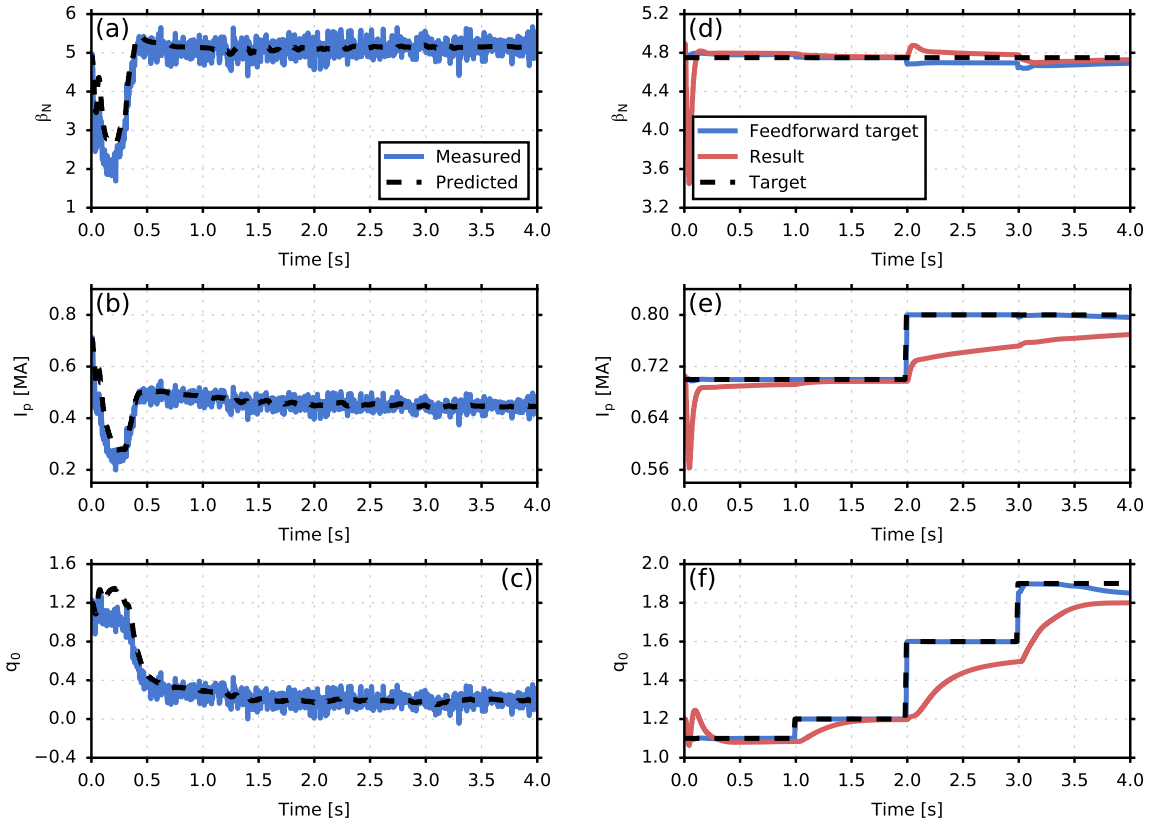

Figure A1: Comparison of simulated noisy measurements (blue-solid) and observer prediction (black-dashed) for (a) β_N, (b) I_p, and (c) q_0 for a simulation with unmeasured input disturbances applied. Comparison of programmed target (black-dashed), optimal achievable steady-state target calculated by the feedforward compensator (blue-solid) and the achieved evolution when the calculated feedforward actuator requests were applied without feedback or anti-windup for (d) β_N, (e) I_p, and (f) q_0.

q_0. The relative amount of deviation in such a situation can be tuned by the selection of the weight matrices Q_{ss}. With the feedforward actuator values applied the system converges toward the achievable targets, however, it does so on the slow timescale of the open-loop system.

Appendix A.3. Feedback design

Under the design assumptions, the outputs of the system would converge to the optimal steady-state values over time (since the system is stable), however, the response time may be slow. To improve response time, a feedback term is designed to drive the estimated state vector toward the state vector corresponding with the feedforward actuator vector and the estimated disturbance, which is given by x_t. Defining $\tilde{x} = \hat{x} - x_t - x_{aw}$, where x_{aw} is a to-be-designed signal from the anti-windup compensator, the tracking error dynamics can be written as

$$\dot{\tilde{x}} = A\tilde{x} + Ax_t + B\tilde{u} + K_I C^T \hat{d} - L_x\left(\hat{y} - \tilde{y}\right) + \dot{x}_t + \dot{x}_{aw}, \tag{A.11}$$

where $\dot{x}_t = -A^{-1}(B\dot{u}_{FF} + K_I C^T \dot{\hat{d}})$. Substituting expression (A.2) for $\tilde{u}$, this is reduced to

$$\dot{\tilde{x}} = A\tilde{x} + Bu_{fb} + Bu_{aw} + L\left(\hat{y} - \tilde{y}\right) + \dot{x}_t + \dot{x}_{aw}, \tag{A.12}$$

The observer and feedforward terms will converge to zero under the design assumptions (and can be designed to do so arbitrarily quickly). The anti-windup input and state adjustments will be designed to ensure that, if no actuator saturation is present, the adjustment will converge to zero. The consideration of actuator limits in the feedforward terms ensures that, under the design assumptions, the steady-state target can be achieved within the actuator limits, i.e., u_{aw} and x_{aw} will converge to zero over time. The feedback law $u_{fb} = -K_{fb}\tilde{x}$ is chosen with K_{fb} taken to be the full-state feedback gain derived using the LQR approach (see standard texts [43, 44]). The state weights Q_{fb} and the actuator weights R_{fb} are design parameters used to adjust response time and actuator use. Under these assumptions, the system (A.12) is input-to-state-stable (ISS) [45] with respect to the anti-windup input adjustment u_{aw} (to be designed), the observer prediction error, and the target state time derivative. Since the latter terms will converge to zero, this implies that $\tilde{x}$ will also converge to zero and the target state x_t will be achieved.

A comparison of the programmed target outputs to the achieved outputs resulting from the combination of feedforward and feedback actuator requests in a simulation of the identified response model is shown in figure A2(a-c). The outputs I_p and q_0, which were weighted more heavily than β_N in both the feedforward and feedback compensators in this simulation, converged much faster toward their steady-state values than in the feedforward only case shown in figure A1(d-f).

Appendix A.4. Control allocation and anti-windup augmentation

Since the response time of the system is relatively slow compared to the length of discharges, large feedback gains may be required to achieve a response time that is significantly faster than the length of discharges. It is therefore expected that even though actuator constraints are considered in the feedforward compensator, the feedback controller may still result in actuator requests that violate the constraints, leading to degraded closed-loop performance. To mitigate the effects of saturation, a model-based anti-windup scheme is added. A model-recovery anti-windup [46] approach is used in which the anti-windup compensator consists of a copy of the plant dynamics driven by the difference between the actuator requests and the saturated actuator requests. The state vector resulting from this system is subtracted from the estimated state vector, effectively hiding the effect of saturation from the feedback controller.

$$\dot{x}_{aw} = Ax_{aw} + B\left[\text{sat}(u_{fb} + u_{aw}) - u_{fb}\right]. \tag{A.13}$$

To make use of potential actuator redundancy to minimize the effect of saturation on the system and to quickly drive the anti-windup modification to zero once the actuator requests are reduced to values that are within the constraints, an actuator request

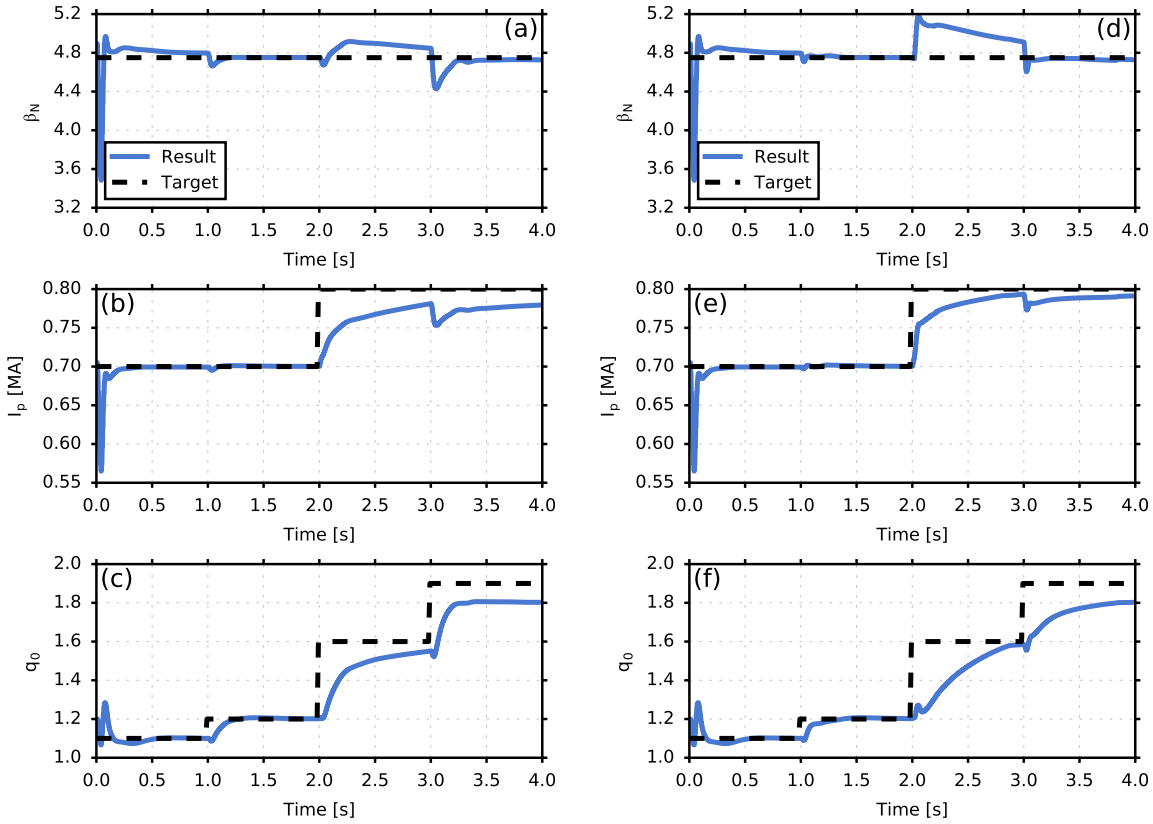

Figure A2: Feedback controlled evolution (blue-solid) compared to programmed target (black-dashed) for (a) β_N, (b) I_p, and (c) q_0 with no anti-windup active. Feedback controlled evolution (blue-solid) compared to programmed target (black-dashed) for (d) β_N, (e) I_p, and (f) q_0 with anti-windup active.

modification u_{aw} is also included, calculated by minimizing

$$J_{aw} = \frac{1}{2}\Delta B Q_{aw} B \Delta + \frac{1}{2} u_{aw}^T R_{aw} u_{aw}, \tag{A.14}$$

where $\Delta = u_{aw} - \rho_{aw} K x_{aw}$, such that the total request $u_{FF} + u_R + u_{FB} + u_{aw}$ is within the actuator limits. A gradient projection method is used to solve this constrained minimization problem at each controller time step. A comparison of the programmed target outputs to the achieved outputs resulting from the combination of feedforward, feedback, and anti-windup actuator requests in a simulation of the identified response model is shown in Figure A2(d-f). The outputs I_p and q_0, which were weighted more heavily than β_N in all three compensators in this simulation, converged faster toward their steady-state values than in the feedforward+feedback case shown in Figure A2(a-c).

[1] MENARD, J. et al., Nuclear Fusion **52** (2012) 083015.
[2] ONO, M. et al., Nuclear Fusion **40** (2000) 557.
[3] SYKES, A. et al., Nuclear Fusion **41** (2001) 1423.
[4] GOLDSTON, R. et al., An Experiment to Tame the Plasma Material Interface FT/P3-12, in *Proc. 22nd Int. Conf. on Fusion Energy*, Geneva, Switzerland, 2008, IAEA.

[5] STAMBAUGH, R. D. et al., Candidates for a Fusion Nuclear Science Facility (FDF and ST-CTF) P2.110, in *37th EPS Conf. on Plasma Physics*, volume 51, Dublin, Ireland, 2010.
[6] MENARD, J. et al., Nuclear Fusion **56** (2016) 106023.
[7] MENARD, J. et al., Nuclear Fusion **51** (2011) 103014.
[8] KAYE, S. et al., Nuclear Fusion **55** (2015) 104002.
[9] KAYE, S. et al., Nuclear Fusion **46** (2006) 848.
[10] KAYE, S. et al., Physical Review Letters **98** (2007) 175002.
[11] VALOVIC, M. et al., Nuclear Fusion **51** (2011) 073045.
[12] FREDRICKSON, E. D. et al., Physics of Plasmas **13** (2006) 056109.
[13] GRYAZNEVICH, M. et al., Nuclear Fusion **48** (2008) 084003.
[14] PODESTA, M. et al., Physics of Plasmas **16** (2009) 056104.
[15] GATES, D. et al., Nuclear Fusion **47** (2007) 1376.
[16] GERHARDT, S. et al., Nuclear Fusion **51** (2011) 073031.
[17] BUTTERY, R. et al., Nuclear Fusion **44** (2004) 1027.
[18] CHAPMAN, I. et al., Nuclear Fusion **51** (2011) 073040.
[19] ERICKSON, K. et al., Fusion Engineering and Design **89** (2014) 853.
[20] PODESTÃ , M. et al., Plasma Physics and Controlled Fusion **58** (2016) 125016.
[21] BOYER, M. et al., Nuclear Fusion **55** (2015) 053033.
[22] GOUMIRI, I. et al., Nuclear Fusion **56** (2016) 036023.
[23] GATES, D. et al., Nuclear Fusion **46** (2006) 17.
[24] GERHARDT, S. et al., Fusion Science and Technology **61** (2010) 11.
[25] KOLEMEN, E. et al., Nuclear Fusion **51** (2011) 113024.
[26] HAWRYLUK, R., An Empirical Approach To Tokamak Transport, in COPPI, B., editor, *Physics of Plasmas Close to Thermonuclear Conditions: Proceedings of the Course Held in Varenna, Italy, 27 August-8 September 1979*, volume 1, pp. 19–46, Varenna, Italy, 1981, Elsevier Ltd.
[27] TRANSP Homepage, 2014.
[28] GERHARDT, S. et al., Nuclear Fusion **52** (2012) 083020.
[29] POLI, F. et al., Nuclear Fusion **55** (2015) 123011.
[30] PANKIN, A. et al., Computer Physics Communications **159** (2004) 157.
[31] ANDRE, R., TRANSP / PTRANSP Isolver Free Boundary Equilibrium Solver, in *American Physical Society, 54th Annual DPP Meeting*, Providence, Rhode Island, 2012, American Physical Society.
[32] HUANG, J. et al., Development of an auto-convergent free-boundary axisymmetric equilibrium solver, in *American Physical Society, 47th Annual DPP Meeting*, Denver, Colorado, 2005, American Physical Society.
[33] ITER Physics Expert Groups, Nuclear Fusion **39** (1999) 2175.
[34] FREDRICKSON, E. et al., Parametric dependence of beam-ion-driven modes in NSTX and NSTX-U, in *International Atomic Energy Agency Fusion Energy Conference*, Kyoto, Japan, 2016.
[35] LIU, D. et al., Initial measurements of beam ion confinement on NSTX-U, in *American Physical Society, 58th Annual DPP Meeting*, Providence, Rhode Island, 2016.
[36] HAO, G. et al., Measurements of fast ion confinement and transport using dual-view Fast Ion D-Alpha Diagnostics on NSTX-U, in *American Physical Society, 58th Annual DPP Meeting*, Providence, Rhode Island, 2016.
[37] MENARD, J. et al., Overview of the first operational and physics results from NSTX Upgrade, in *International Atomic Energy Agency Fusion Energy Conference*, Kyoto, Japan, 2016.
[38] BOYER, M. D. et al., Plasma Physics and Controlled Fusion **55** (2013) 105007.
[39] MALJAARS, E. et al., Nuclear Fusion **023001** 23001.
[40] OU, Y. et al., Control Engineering Practice **19** (2011) 22 .
[41] OUARIT, H. et al., Fusion Engineering and Design **86** (2011) 1018 , Proceedings of the 26th Symposium of Fusion Technology (SOFT-26).

[42] BAS, O. Y. et al., Design of Optimal Gains for the Proportional Integral Kalman Filter with Application to Single Particle Tracking, in *38th IEEE Conference on Decision and Control and European Control Conference*, pp. 4567–4571, Phoenix, AZ, USA, 21999, IEEE.

[43] SKOGESTAD, S. et al., *Multivariable Feedback Control: Analysis and Design*, Wiley, New York, 2005.

[44] ASTROM, K. et al., *Feedback Systems: an Introduction for Scientists and Engineers*, Princeton University Press, Princeton, NJ, 2008.

[45] KHALIL, H. et al., *Nonlinear systems*, Prentice Hall, Englewood Cliffs, NJ, 3rd edition, 2002.

[46] GALEANI, S. et al., European Journal of Control **15** (2009) 418.

PPPL PRINCETON PLASMA PHYSICS LABORATORY

Princeton Plasma Physics Laboratory
Office of Reports and Publications

Managed by
Princeton University

under contract with the
U.S. Department of Energy
(DE-AC02-09CH11466)

P.O. Box 451, Princeton, NJ 08543
Phone: 609-243-2245
Fax: 609-243-2751

E-mail: publications@pppl.gov

Website: http://www.pppl.gov