Plasma Physics and
Controlled Fusion

PAPER • OPEN ACCESS

Experimental demonstration of reactor-relevant adaptive density control on the DIII-D tokamak with coordinated gas and pellet fueling

To cite this article: Hassan Al Khawaldeh *et al* 2026 *Plasma Phys. Control. Fusion* **68** 085001

IOP Publishing Plasma Phys. Control. Fusion 68 (2026) 085001 https://doi.org/10.1088/1361-6587/ae89e3

Plasma Physics and Controlled Fusion

OPEN ACCESS

RECEIVED
17 April 2026

REVISED
6 July 2026

ACCEPTED FOR PUBLICATION
12 July 2026

PUBLISHED
31 July 2026

PAPER

Experimental demonstration of reactor-relevant adaptive density control on the DIII-D tokamak with coordinated gas and pellet fueling

Hassan Al Khawaldeh[1,*], Andres Pajares[2], Sai Tej Paruchuri[1], Vincent Graber[1], Tariq Rafiq[1], Eugenio Schuster[1], June-Woo Juhn[3], Troy Pederson[2], Daisuke Shiraki[4] and Francesca Turco[5]

1 Lehigh University, Bethlehem, PA, United States of America
2 General Atomics, San Diego, CA, United States of America
3 Korea Institute of Fusion Energy, Daejeon, Republic of Korea
4 Oak Ridge National Laboratory, Oak Ridge, TN, United States of America
5 Columbia University, New York, NY, United States of America
* Author to whom any correspondence should be addressed.

E-mail: alkhawaldeh@lehigh.edu

Keywords: density, control, adaptive, indirect, delays, gas, pellet

Abstract

Precise regulation of plasma density is crucial for achieving and maintaining fusion-relevant conditions in reactor-grade tokamaks. Tokamak reactors will utilize both gas puffing and pellet injection as standard fueling mechanisms. However, the coordinated use of gas puffing and pellet injection within feedback control frameworks is largely unexplored, underscoring the necessity to develop and validate dedicated strategies on existing machines. These strategies must address the various challenges associated with both actuators, including actuation delays, unknown fueling efficiencies, and the coexistence of continuous-time and discrete-time dynamics. To overcome these challenges, which surpass the capabilities of traditional empirically tuned proportional-integral-derivative (PID) control, an indirect adaptive control algorithm is proposed in this work for the regulation of the line-averaged electron density through coordinated gas puffing and pellet injection. After initial validation in simulations with a multi-reservoir global particle model, the controller was successfully implemented and tested on the DIII-D tokamak. Experimental results demonstrate robust density tracking under reactor-relevant scenarios, showcasing the controller's ability to seamlessly coordinate actuators while handling disturbances and evolving actuator constraints. This work provides a critical step forward in the control of future fusion reactors.

1. Introduction

To demonstrate the viability of fusion as a large-scale energy source, future reactor-grade tokamaks aim to operate with plasmas sustained over long-pulse durations at high power levels. A key parameter that governs the possibility for achieving such operation is the plasma density, as the generated fusion power is directly proportional to the square of the density. However, maintaining this quantity in reactor-relevant conditions remains a significant challenge due to the complexities of fueling and particle exhaust. Consequently, effective density regulation and fueling are critical for operational stability.

In present tokamaks, plasma is fueled primarily through gas puffing, in which neutral gas is introduced into the vacuum vessel and ionized at the plasma edge. While gas puffing is effective in existing devices, its performance is expected to decrease in future reactor-grade tokamaks. Larger device dimensions and higher plasma densities will increase transit times from gas valves to the plasma edge and limit gas penetration into the core, respectively [1]. Therefore, pellet injection has emerged as a promising solution to be used alongside gas puffing [2–4]. This technique involves rapidly injecting frozen fuel into the plasma, enabling deeper penetration and more efficient core density control. Unlike gas puffing, which provides a continuous and smoothly adjustable particle flow, pellet injection delivers fixed fuel packets separated by idle periods, giving it an inherently discrete-time nature. In addition, pellets must travel through injector guide tubes before reaching the plasma, introducing a non-negligible time delay; the risk of pellet integrity loss also increases at the high extrusion and firing rates required for increased throughput. Thus, the coordinated use of pellet injection and gas puffing for routine density control remains largely unexplored, which motivates the need for control strategies capable of addressing fueling challenges.

The most commonly used feedback control strategy for plasma density regulation in current tokamaks is the proportional-integral-derivative (PID) controller [5, 6]. Although reliable in present devices, standalone PID control faces several limitations in future reactor-grade tokamaks. First, PID controllers rely on fixed parameters that are tuned prior to a discharge, which can lead to degraded performance as plasma conditions evolve in long-pulse operation. Second, trial-and-error tuning in reactor environments is time-consuming and expensive. This tuning complexity increases significantly when moving from single-input single-output (SISO) to multi-input single-output (MISO) architectures required for coordinated fueling. Standard PID frameworks lack intrinsic mechanisms to optimally allocate control authority between multiple actuators, such as gas and pellets. The challenge becomes even greater when moving to multi-input multi-output (MIMO) architectures as needed to regulate the density profile at multiple locations. Third, conventional PID control assumes that fueling actuators are delay-free and continuous. However, both gas puffing and pellet injection exhibit inevitable time delays and a mix of continuous- and discrete-time dynamics. As the associated challenges of pellet injection and gas puffing are not naturally accommodated within traditional PID-based approaches, these approaches may fall short of effectively regulating plasma density without explicit augmentations. Therefore, it is imperative to explore feedback control frameworks that can address the challenges associated with fueling mechanisms.

To overcome the limitations of traditional PID-based density control, recent work has shifted toward model-based approaches. Model-based optimal controllers, such as model predictive control, determine inputs by forecasting the future behavior of the system [7]. They have gained prominence because of their ability to handle actuator constraints and mixed continuous/discrete dynamics in plasma density control [8, 9]. However, their effectiveness ultimately depends on the accuracy of the underlying model. To address modeling uncertainties and time-varying dynamics, model-based adaptive control frameworks have been proposed [10–12]. These controllers adjust their internal parameters in real time to account for evolving plasma dynamics and model mismatches that fixed-model approaches cannot capture [13]. Building on these advances, future progress requires control strategies that explicitly coordinate multiple actuators subject to different time delays. Therefore, developing such mechanisms represents a critical step toward effective and robust density regulation.

A model-based adaptive controller has been developed for active regulation of the line-averaged electron density $\bar{n}_e$. The controller enables coordinated actuation via gas puffing and pellet injection while explicitly accounting for time delays. The control framework leverages the structures of indirect model reference adaptive control (MRAC) [14] as well as the Smith Predictor (SP) [15] to compute the control inputs. The indirect MRAC estimates the system parameters in real time and adjusts the control law accordingly, while the SP provides the delay-free error response of the system based on these estimates. Combining these components yields a control design well-suited to address the challenges considered in this work. Notably, similar frameworks have been explored in other control applications [16, 17]. However, the formulation of the design and the application to plasma density regulation in tokamaks are novel.

The performance of the developed controller has been evaluated using a multi-reservoir particle model that mimics the plasma density dynamics in a tokamak. Simulation studies have been conducted to analyze the controller's performance in response to sudden changes in dynamics, different time delays, and varying actuator efforts. Following successful simulation results, the controller has been experimentally tested on the DIII-D tokamak. The results demonstrate robust tracking capabilities under various plasma conditions while coordinating pellet and gas actuation mechanisms effectively. This experimental validation represents an important step toward the development of advanced fueling control strategies for next-generation tokamaks.

The remainder of this work is organized as follows. Section 2 presents the control design and formulation. Section 3 details the simulation development and control validation. Section 4 discusses the experimental results of the developed controller. Section 5 provides concluding remarks.

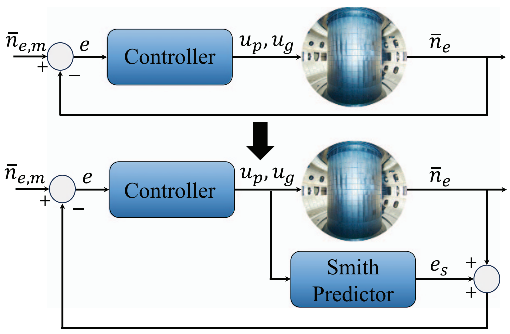

Figure 1. Two block diagrams demonstrating the closed-loop structure of a system without and with the Smith Predictor.

2. Control design and formulation

To facilitate control synthesis, the feedback control technique proposed in this work is based on a simplified model of the electron density. In particular, the response of the zero-dimensional line-averaged electron density $\bar{n}_e$ is approximated as

$$\dot{\bar{n}}_e(t) = A\bar{n}_e(t) + B_g u_g(t - d_g) + B_p u_p(t - d_p) + B_{ff} u_{ff}(t), \tag{1}$$

where $u_g(t)$ is a continuous gas injection particle flow, $u_p(t)$ is a continuous pellet injection particle flow, $u_{ff}(t)$ is a continuous delay-free feedforward source such as neutral beam injection (NBI), $d_g \geqslant 0$ is the time delay of $u_g(t)$, and $d_p \geqslant 0$ is the time delay of $u_p(t)$. Although pellet injection is inherently discrete, $u_p(t)$ is treated as a continuous-time input for the purpose of control design. This assumption implies that $u_p(t)$ represents a continuous 'requested' fueling rate. This signal is subsequently realized by a logic-based modulation algorithm that triggers discrete pellet injection events to match the requested rate on average. For the zero-dimensional regulation of the line-averaged electron density targeted in this work, any realization mismatch is expected to have a negligible impact. Lastly, B_{ff} is known, while A, B_g, and B_p are assumed to be unknown constants, with A negative and B_g and B_p positive. This assumption can be relaxed by treating these parameters as piecewise constant or slowly varying.

As the model parameters A, B_g, and B_p are unknown, MRAC is adopted here as the baseline of the density-control strategy. Thus, the control objective is to design a controller that makes the plant (1) track a reference governed by the following response model

$$\dot{\bar{n}}_{e,m}(t) = A_m \bar{n}_{e,m}(t) + B_m r(t) + B_{ff} u_{ff}(t), \tag{2}$$

where $\bar{n}_{e,m}(t)$ is the reference line-averaged electron density, $r(t)$ denotes the reference inputs, and A_m and B_m are known. This work aims to design a controller that adjusts the gas and pellet fueling rates, $u_g(t)$ and $u_p(t)$, in real time such that $\bar{n}_e(t)$ follows the reference $\bar{n}_{e,m}(t)$ as closely as possible despite delays, uncertainty, or a change in plasma conditions. While traditional MRAC schemes could be applied for control synthesis, the presence of input time delays, d_g and d_p, in (1) violates the assumptions of standard MRAC implementations. To address the time delay challenge, an SP is integrated into the framework. In this scheme, the feedback signal $\bar{n}_e(t)$ is augmented with a correction term known as the Smith error $e_s(t)$. This term quantifies the difference between the predicted delay-free plasma density and the corresponding delayed plasma density. To compute the Smith error, defined as $e_s(t) = \bar{n}_{e,s}(t) - \bar{n}_{e,sd}(t)$, a prediction model is formulated as

$$\dot{\bar{n}}_{e,s}(t) = \hat{A}(t)\bar{n}_{e,s}(t) + \hat{B}_g(t) u_g(t) + \hat{B}_p(t) u_p(t) + B_{ff} u_{ff}(t), \tag{3}$$

$$\dot{\bar{n}}_{e,sd}(t) = \hat{A}(t)\bar{n}_{e,sd}(t) + \hat{B}_g(t) u_g(t - d_g) + \hat{B}_p(t) u_p(t - d_p) + B_{ff} u_{ff}(t), \tag{4}$$

where $\bar{n}_{e,s}(t)$ and $\bar{n}_{e,sd}(t)$ denote the delay-free and delayed model responses, respectively. The parameters $\hat{A}(t)$, $\hat{B}_g(t)$, and $\hat{B}_p(t)$ represent estimates of the unknown coefficients described in (1), generated in real time by the MRAC. Although the true system parameters in (1) are constant, these estimates are time-varying in (3)–(4). This is because the adaptive mechanism continuously adjusts them online as new data becomes available, ensuring the system converges to the desired control goal despite initial uncertainties.

The integration of the SP is illustrated in figure 1. The upper schematic depicts a standard feedback loop where the measured line-averaged electron density $\bar{n}_e(t)$ is compared directly to the reference

Algorithm 1. Pellet injection logic.

Initialize: $p_i \leftarrow 0$, $l_p \leftarrow -\frac{1}{f_{max}}$, $a_{tot} \leftarrow (p_i+1)\,a_i$
for each $t \geqslant 0$ until the end of the discharge **do**
  if $\int_0^t u_p(\tau)\,d\tau V_p \geqslant a_{tot}$ **and** $t \geqslant l_p + \frac{1}{f_{max}}$ **then**
    $u_{fire}(t) \leftarrow 1$
    $p_i \leftarrow p_i + 1$
    $l_p \leftarrow t$
    $a_{tot} \leftarrow (p_i+1)\,a_i$
  else
    $u_{fire}(t) \leftarrow 0$
  end if
end for

$\bar{n}_{e,m}(t)$ to determine the control inputs $u_g(t)$ and $u_p(t)$. The lower schematic displays the modified structure incorporating the SP. Instead of relying solely on the delayed measurement $\bar{n}_e(t)$, the framework computes the Smith error $e_s(t)$ and combines it with the feedback signal. The resulting augmented quantity, $\bar{n}_e(t) + e_s(t)$, is then used to generate the control actions. Accordingly, the objective of this work is to design a controller that ensures

$$e(t) = \bar{n}_{e,m}(t) - (\bar{n}_e(t) + e_s(t)) \to 0 \quad \text{as} \quad t \to \infty. \tag{5}$$

The rationale for this formulation is as follows. The Smith error $e_s(t)$ acts as a correction term that accounts for the fact that the current line-averaged electron density measurement $\bar{n}_e(t)$ has not yet reflected the effect of recent fueling actions due to input delays. By adding $e_s(t)$ to $\bar{n}_e(t)$, the controller is effectively provided with an estimate of how the line-averaged electron density would evolve once the effects of recent fueling are realized, enabling the computation of appropriate control inputs. To that end, the control laws are defined as

$$u_g(t) = \hat{B}_g(t)^{-1}\left[\alpha_1 e(t) + \beta\left(\hat{A}(t)e(t) - \hat{A}(t)\bar{n}_{e,m}(t) + A_m\bar{n}_{e,m}(t) + B_m r(t)\right)\right], \tag{6}$$

$$u_p(t) = \hat{B}_p(t)^{-1}\left[\alpha_2 e(t) + (1-\beta)\left(\hat{A}(t)e(t) - \hat{A}(t)\bar{n}_{e,m}(t) + A_m\bar{n}_{e,m}(t) + B_m r(t)\right)\right], \tag{7}$$

$$\dot{\hat{A}}(t)^{\mathrm{T}} = -\Gamma_A \bar{n}_e(t)\,e(t)^{\mathrm{T}}, \tag{8}$$

$$\dot{\hat{B}}_g(t)^{\mathrm{T}} = -\Gamma_{B_g} u_g(t-d_g)\,e(t)^{\mathrm{T}}, \tag{9}$$

$$\dot{\hat{B}}_p(t)^{\mathrm{T}} = -\Gamma_{B_p} u_p(t-d_p)\,e(t)^{\mathrm{T}}, \tag{10}$$

where $\alpha_1 = \beta\alpha$, $\alpha_2 = (1-\beta)\alpha$, α is a positive scalar, β is constrained to lie between 0 and 1, and Γ_A, Γ_{B_g}, and Γ_{B_p} are positive user-defined scalars. The parameters α and β determine the aggressiveness of the feedback and the distribution of effort between actuators, respectively. A larger value of α produces a faster but less robust corrective action, whereas a smaller value results in a slower but more conservative behavior. Similarly, setting β closer to 1 means that the controller prioritizes gas puffing to correct the tracking error, while values closer to 0 shift the control effort toward pellet injection. Lastly, the learning rates Γ_A, Γ_{B_g}, and Γ_{B_p} determine how rapidly the adaptive mechanism updates the model parameters, where the trade-off between small and large values is similar to that of α. The derivation of the control laws can be found in appendix A.

2.1. Pellet injection logic

The developed controller treats pellet injection as a continuous actuator. Therefore, the feedback component $u_p(t)$ from (7) must be transformed into a binary fire command $u_{\text{fire}}(t)$, where $u_{\text{fire}}(t) = 1$ denotes firing a pellet and $u_{\text{fire}}(t) = 0$ denotes no injection. This conversion is designed to preserve the required fueling implied by $u_p(t)$, ensuring consistency with the control formulation. Thus, let a_i be the number of particles in a pellet, a_{tot} be the total number of injected particles in a discharge, p_i be the number of injected pellets in a discharge, V_p be the plasma volume, l_p be the time at which the last pellet was injected, and $f_{\max}$ be the maximum pellet injection frequency. Then, Algorithm 1 is utilized to transform $u_p(t)$ into $u_{\text{fire}}(t)$, following approaches similar to those reported in [18, 19].

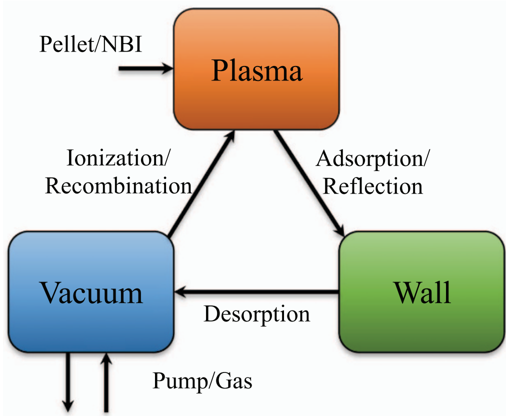

Figure 2. Simplified diagram of the particle interactions considered in the control-oriented simulation model.

3. Simulation-based assessment of the controller

3.1. Model used for assessment

In the following simulation study, the proposed controller is evaluated using a multi-reservoir particle model that captures the simplified dynamics of particle interactions within the plasma. Adopting the approach in [20–22], the model treats the vacuum, wall, and plasma regions of the tokamak as interconnected reservoirs subject to external sources (gas puffing, pellet injection, and NBI) and sinks (pumping). These interactions are illustrated in figure 2. Physically, particles are exhausted from the plasma to the wall, released into the vacuum, and subsequently recycled back into the plasma through ionization. While actual tokamak transport involves more complex phenomena, the interactions considered here capture the fundamental dynamical trends. Crucially, this simulation model exhibits higher-order dynamics and greater structural complexity compared to the simplified first-order design model used for control synthesis. Consequently, it serves as a rigorous testbed for evaluating the controller's performance in the presence of unmodeled dynamics. The complete model is given by

$$\dot{N}_{\mathrm{e}}(t) = -\frac{N_{\mathrm{e}}(t)}{\tau_{\mathrm{e}}} + \frac{N_{\mathrm{v}}(t)}{\tau_{\mathrm{v}}} + S_{\mathrm{p}}(t) + S_{\mathrm{n}}(t), \tag{11}$$

$$\dot{N}_{\mathrm{w}}(t) = -\frac{N_{\mathrm{w}}(t)}{\tau_{\mathrm{w}}} + \frac{N_{\mathrm{e}}(t)}{\tau_{\mathrm{e}}}, \tag{12}$$

$$\dot{N}_{\mathrm{v}}(t) = -\frac{N_{\mathrm{v}}(t)}{\tau_{\mathrm{v}}} + \frac{N_{\mathrm{w}}(t)}{\tau_{\mathrm{w}}} + S_{\mathrm{g}}(t) - S_{\mathrm{pump}}(t), \tag{13}$$

$$\bar{n}_{\mathrm{e}}(t) = \frac{N_{\mathrm{e}}(t)}{V_{\mathrm{p}}}, \tag{14}$$

where $N_{\mathrm{e}}(t)$ denotes the number of electrons in the plasma region, $N_{\mathrm{v}}(t)$ represents the number of neutrals in the vacuum region, $N_{\mathrm{w}}(t)$ denotes the number of retained particles in the wall inventory, $S_{\mathrm{p}}(t)$ is the fueling source due to the pellet injection, $S_{\mathrm{n}}(t)$ is the fueling source due to NBIs, $S_{\mathrm{g}}(t)$ is the fueling source due to the gas puffing, $S_{\mathrm{pump}}(t)$ is the sink due to the pump action, τ_{e} is the electron confinement time, τ_{v} is the vacuum confinement time, and τ_{w} is the wall confinement time. Note that (14) is taken in this work to represent the line-averaged electron density.

3.2. Fueling sources

The fueling sources appearing in (11)–(13) are expressed in units of s^{-1}. To relate these source terms to the actual actuators, appropriate models are introduced. For gas puffing, the input $u_{\mathrm{g}}(t)$ is in units of $\mathrm{s}^{-1}\,\mathrm{cm}^{-3}$; therefore, $u_{\mathrm{g}}(t)$ is mapped to $S_{\mathrm{g}}(t)$ by

$$S_{\mathrm{g}}(t) = V_{\mathrm{p}} u_{\mathrm{g}}(t). \tag{15}$$

The NBI system is actuated by specifying an injection power rather than a direct particle flow rate. Therefore, to accurately model its contribution, the corresponding source term must be derived as a function of the injected power. The term $S_{\mathrm{n}}(t)$ is modeled as

$$S_{\mathrm{n}}(t) = \frac{N^{\mathrm{avg}}}{E_{\mathrm{NBI}}} P_{\mathrm{NBI}}(t), \tag{16}$$

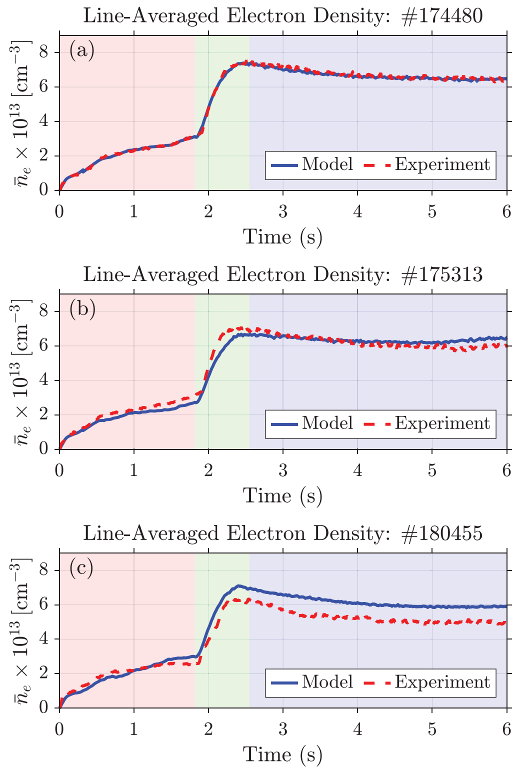

Figure 3. Tailoring the particle model to discharge #174480 on DIII-D and comparing it to discharges #175313 and #180455: (*a*) model's performance in reproducing discharge #174480; (*b*) model's performance in reproducing discharge #175313; (*c*) model's performance in reproducing discharge #180455.

where $P_{\rm NBI}(t)$ is the total NBI power in W, $E_{\rm NBI}$ is the energy of the beam in J, and $N^{\rm avg}$ is the average number of particles in a beam based on its species mix. For instance, the species mix for DIII-D's NBIs at 80 keV is D1 : D2 : D3 = 80 : 14 : 6 (%), which converts to $N^{\rm avg} = 1.26$.

The number of particles extracted by the pumps depends on the local particle density near the pump location. Since the pumps are situated in the vacuum region, this implies that the pumping rate is directly influenced by the vacuum inventory. Therefore, the pumping source term $S_{\rm pump}(t)$ is modeled as

$$S_{\rm pump}(t) = \frac{N_{\rm v}(t)}{\tau_{\rm pump}}, \tag{17}$$

where $\tau_{\rm pump}$ is a pumping time constant.

The pellet injection source is treated as a step-like response rather than an instantaneous event to reduce modeling effort. Specifically, each firing command $u_{\rm fire}(t)$ is assumed to produce a constant particle flow over a time interval $t_{\rm s}$. Accordingly, the pellet injection source $S_{\rm p}$ is modeled as

$$S_{\rm p}(t) = \frac{a_i}{t_{\rm s}} \max_{\tau \in [t-t_{\rm s}, t]} u_{\rm fire}(\tau). \tag{18}$$

Furthermore, the pellet injection system operates under a constraint defined by the maximum frequency $f_{\max}$, which imposes a mandatory waiting period between successive pellet injections. To ensure that the influence of each pellet is fully captured, the value of $t_{\rm s}$ must be smaller than $1/f_{\max}$.

3.3. Model assessment

Figures 3(a)–(c) present the model's performance in reproducing three experimental discharges on DIII-D. The model has been specifically tailored for discharge #174480, as shown in figure 3(a), where it closely reproduces the experimental data. Achieving this level of accuracy required a three-stage fitting process. The stages correspond to the *L*-mode, *L–H* transition, and *H*-mode phases and are represented in the figures by red, green, and blue backgrounds, respectively. Since each stage has its own dynamics, different parameter fits are obtained for each stage by solving an optimization problem that minimizes the difference between the modeled and measured line-averaged electron density. The parameter sets

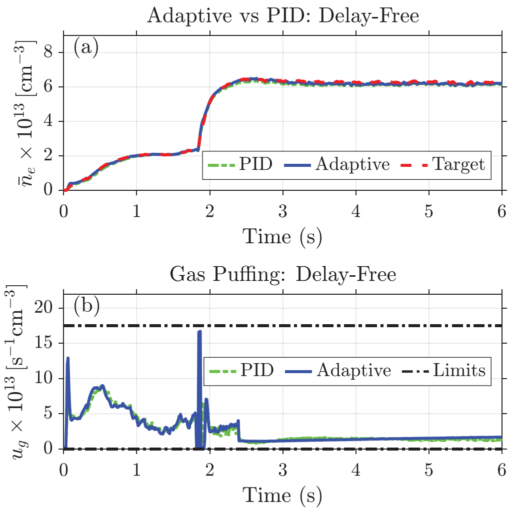

Figure 4. Simulation cases illustrating the responses of adaptive and PID controllers in regulating density for a delay-free plant: (a) closed-loop simulation with no time delay; (b) gas puffing responses with no time delay.

obtained for the three stages are then used to reproduce the discharges. Figure 3(b) and figure 3(c) show the model's predictions for discharges #175313 and #180455, respectively. Notably, discharge #175313 has engineering and operational conditions similar to those of #174480, which explains the model's relatively accurate prediction of the density. By contrast, discharge #180455 differs in its conditions; therefore, the prediction is less accurate but remains reasonable for control-design purposes because it captures the dynamical trends.

In assessing the proposed modeling framework, it is worth noting that while physical plasma density transport exhibits complex and nonlinear characteristics, the dominant dynamics required for real-time feedback regulation are effectively captured by this first-order linear structure. The primary control challenge associated with the density-regulation problem tackled in this work is the highly uncertain nature of the transport properties—particularly during sudden confinement changes such as the L–H transition. The real-time parameter adaptation mechanism of the MRAC is designed specifically to track these variations online, while remaining unmodeled higher-order effects are safely mitigated by the closed-loop feedback action.

3.4. Closed-loop simulations

A series of simulations is conducted to evaluate the performance of the adaptive controller described in section 2. These simulations have been divided into two studies. The first study, described in section 3.4.1, compares the response of the developed adaptive controller to that of a standalone PID controller. The objective of this study is to analyze the performance of each controller in the presence of input time delays on gas puffing. The second study, described in section 3.4.2, examines the tracking capability of the developed adaptive controller using gas puffing and pellet injection simultaneously.

3.4.1. Density regulation via delayed gas puffing

Three simulations are conducted to compare the performance of the PID and adaptive controllers under varying time delays: 0 ms, 100 ms, and 300 ms. For consistency, the controller configuration and tuning are performed only once using the delay-free scenario, ensuring both schemes exhibit similar baseline behavior when no time delays are present. Additionally, as this analysis focuses solely on gas puffing, the weighting factor β for the adaptive controller in (6) is set to 1 throughout the simulations.

Figures 4–6 illustrate the comparative results. In the absence of time delay (figure 4), both controllers achieve accurate tracking with similar control inputs. However, the introduction of delays reveals a significant divergence in performance. As shown in figures 5 and 6 (100 ms and 300 ms, respectively), the PID controller exhibits oscillatory behavior that intensifies as the delay magnitude increases, reflecting its lack of explicit delay compensation. In contrast, the proposed adaptive controller maintains stable and accurate tracking in all scenarios. These results suggest that while PID may be sufficient for negligible delays, its performance degrades significantly as latency increases.

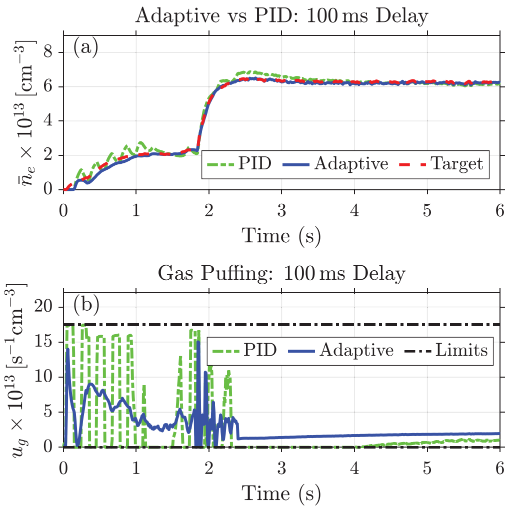

Figure 5. Simulation cases illustrating the responses of adaptive and PID controllers in regulating density for a time-delayed plant: (*a*) closed-loop simulation with 100 ms time delay; (*b*) gas puffing responses with 100 ms time delay.

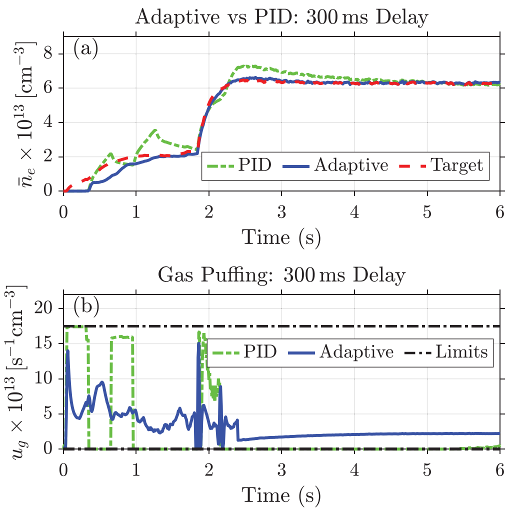

Figure 6. Simulation cases illustrating the responses of adaptive and PID controllers in regulating density for a time-delayed plant: (*a*) closed-loop simulation with 300 ms time delay; (*b*) gas puffing responses with 300 ms time delay.

3.4.2 Coordinated gas puffing and pellet injection in H-mode.

Two simulations are presented to evaluate the controller's ability to coordinate multi-actuator fueling under input delays of 100 ms. Both scenarios track an identical target trajectory designed to exceed the fueling capacity of gas puffing alone, necessitating the use of pellets. The first scenario (figure 7) relies exclusively on gas puffing by holding the weighting factor β at 1. The second scenario (figure 8) enables combined actuation. To ensure a direct comparison of the transition behavior, the second scenario mimics the first by forcing $\beta = 1$ until $t = 3$ s, after which it is modified to 0.5. This constraint forces the adaptive controller to rely solely on gas puffing during the *L*-mode phase (as in the first simulation) before enabling combined actuation for the *H*-mode phase.

As illustrated in figure 7(a), the controller achieves robust tracking of the target in *L*-mode. However, once the plasma transitions to *H*-mode, the density fails to reach the desired values. This limitation is due to the reduced gas puffing efficiency in that regime, as evidenced by the increased actuation effort and negligible density growth. In contrast, figure 8(a)–(c) demonstrates the robustness of the controller in tracking the target by coordinating both actuators. The controller is capable of tracking the desired target despite the mixed continuous-discrete actuator dynamics and the input time delays. Furthermore, figure 8(c) shows that pellet injection occurs at rates adapted to the tracking objective: evenly spaced injections when maintaining constant density and more frequent injections when a rise in density is

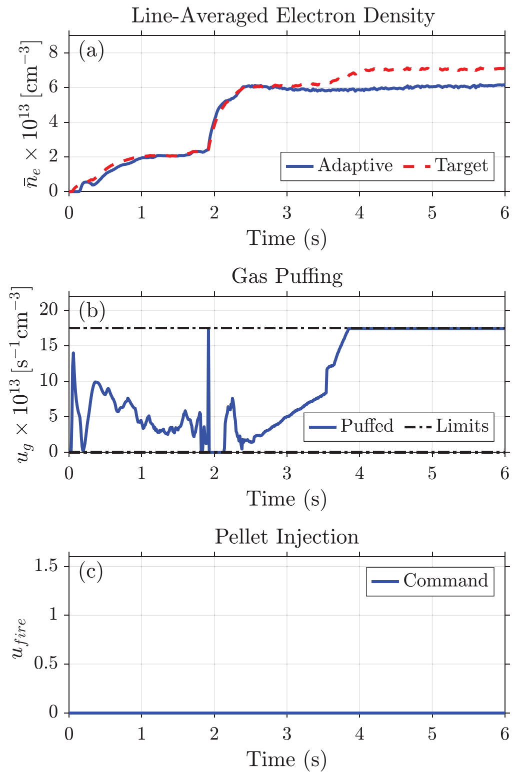

Figure 7. Simulation case illustrating the responses of the adaptive controller in regulating density using only gas puffing: (*a*) evolution of line-averaged density; (*b*) gas puffing response; (*c*) pellet injection response.

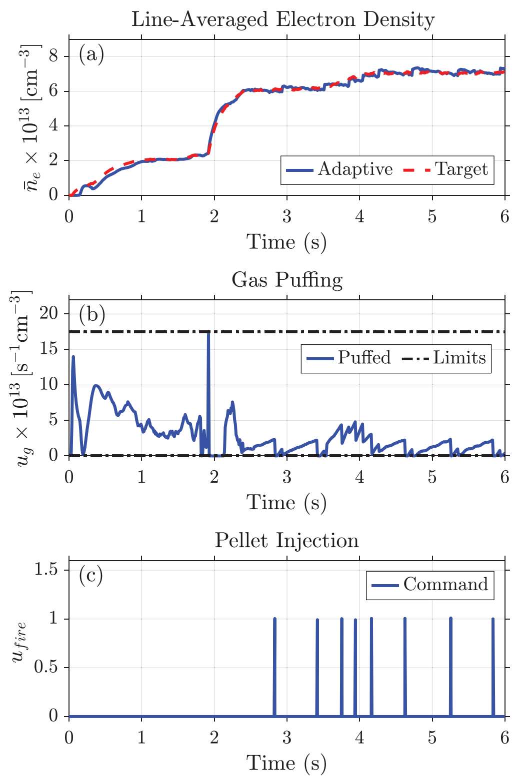

Figure 8. Simulation case illustrating the responses of the adaptive controller in regulating density using gas puffing and pellet injection: (*a*) evolution of line-averaged density; (*b*) gas puffing response; (*c*) pellet injection response.

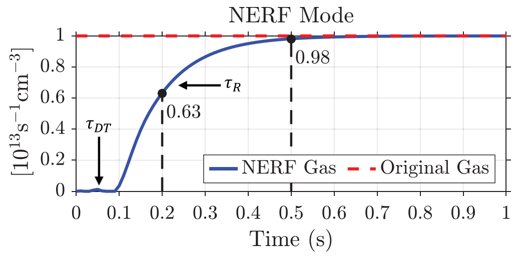

Figure 9. Simulation case demonstrating the gas signal altered by NERF mode compared with the original signal.

required. Overall, these simulations show that the controller maintains robust tracking despite input time delays.

4. Experimental testing on DIII-D

4.1. Integration into the DIII-D plasma control system (PCS)

Present-day devices, such as DIII-D, are significantly smaller than future reactor-grade tokamaks. Consequently, direct experimental testing on these machines does not inherently reproduce the significant fueling transport delays expected in reactor environments. To establish a more representative testing platform, an emulation algorithm known as 'NERF' mode [23] has been implemented in the PCS.

Named after the colloquial term for intentionally reduced capabilities, this operational mode artificially restricts actuator performance to emulate the challenging and constrained fueling environment anticipated in next-generation devices. Specifically, the algorithm intercepts the gas command, modifies it to replicate the delayed response dynamics anticipated in future tokamaks, and subsequently applies the altered signal to the gas valves. The resulting emulated command adheres to a first-order plus dead-time model, characterized by the transfer function

$$G(s) = \frac{\mathrm{e}^{-\tau_{\mathrm{DT}} s}}{\tau_R s + 1}, \tag{19}$$

where τ_{DT} represents the dead-time constant, and τ_R represents the lag-time constant. The dead-time constant defines the time before the signal begins to change. In contrast, the lag-time constant describes how long a signal takes to reach 63% of its final value [24]. By definition, the expected time for the signal to reach 98% of its final value is four times the lag-time constant. For this study, these parameters are set to $\tau_{\mathrm{DT}} = 100\,\mathrm{ms}$ and $\tau_R = 100\,\mathrm{ms}$. An illustration of the NERF mode response is provided in figure 9. When a constant gas flow request of $1 \times 10^{13}\,\mathrm{s}^{-1}\,\mathrm{cm}^{-3}$ is applied, the imposed delay and lag modify the valve actuation significantly. As shown in the figure, the altered flow begins to rise only after 100 ms and gradually reaches the desired value over approximately 400 ms, consistent with the specified τ_{DT} and τ_R parameters. As the NERF mode applies both a dead time and a lag, the altered gas command in this work will be denoted by u_{nerf}.

Figure 10 shows the locations of DIII-D's gas valves and the pellet injectors considered for the adaptive controller in this work. In this initial stage of experimental testing, the use of multiple gas valves and pellet injectors is intentionally avoided to simplify implementation and to enable a clearer performance assessment. Therefore, for the lower single null plasma considered for the experiment, the GASA valve and the high-field side (HFS) midplane tube are chosen as the fueling sources. The GASA valve serves as the standard deuterium gas injector on DIII-D, while the HFS location is used due to its superior fueling efficiency.

The integration of the NERF mode with the controller is illustrated in figure 11. Since the adaptive controller requires both the original and the altered gas signals, its framework is divided into four functional blocks: Plant Estimation, Smith Predictor, Control Law, and Pellet Decision. The Plant Estimation block integrates equations (8)–(10) to give the estimated model parameters. The Smith Predictor block uses (3)–(4) to provide a delay-adjusted measurement. The Control Law block evaluates equations (6)–(7) to provide the required fueling sources. Finally, the Pellet Decision block executes Algorithm 1 (introduced in section 2.1). Overall, the estimated model parameters obtained from the Plant Estimation block are supplied to both the Smith Predictor and Control Law blocks. Then, using the measured

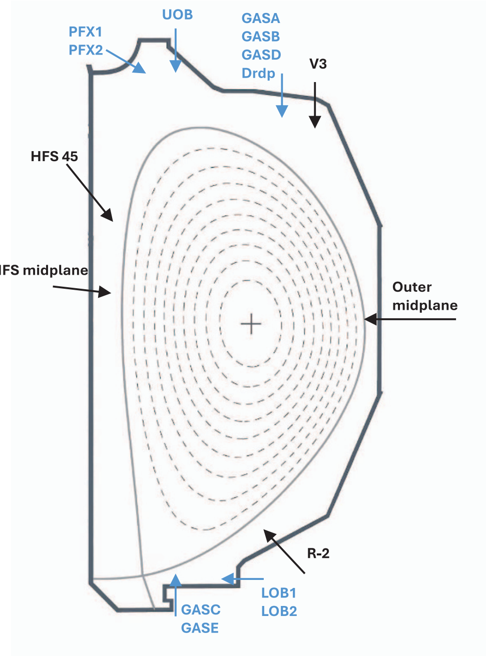

Figure 10. A cross section of the DIII-D tokamak showing the locations of the gas valves and pellet injectors. The blue arrows indicate gas valve positions, while the black arrows denote pellet injection tubes.

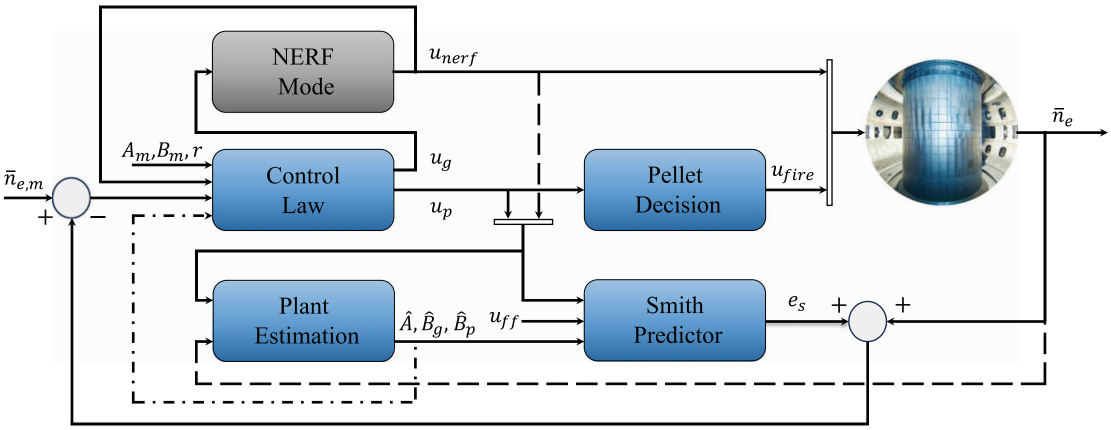

Figure 11. Architecture of the adaptive control scheme. The blue blocks are algorithms executed in the pellet category in the plasma control system (PCS). The gray blocks are algorithms executed in the Gas Category. The overall structure requires constant communication across these two categories.

plasma density together with the prediction made by the Smith Predictor b, the Control Law block computes the needed gas and pellet signals. These signals are first clipped by their saturation limits and then processed by the NERF mode and Pellet Decision blocks to generate the actuation signals.

4.2. *L*-mode testing of adaptive controller

To reliably test the proposed adaptive controller, an appropriate experimental configuration is established. Data from previous discharges are used to select suitable values for the parameters A_m, B_m, and B_{ff} defined in the reference model (2). Based on this model, the reference inputs $r(t)$ are designed to generate a desired target trajectory for the experiment. The controller is then configured with the reference model and its internal parameters are set to $\beta = 0.5$, $\alpha = 1$, $\Gamma_A = 1$, $\Gamma_{B_g} = 1$, and $\Gamma_{B_p} = 1$. Finally, an input time delay of 100 ms is emulated by leveraging the NERF mode algorithm.

Figure 12 presents the experimental results obtained from the adaptive controller. The measured line-averaged electron density and the target trajectory generated by the reference model are shown in figure 12(a). The corresponding gas and pellet commands are displayed in figures 12(b) and (c), respectively, while the NBI and electron cyclotron heating (ECH) powers are plotted in figure 12(d).

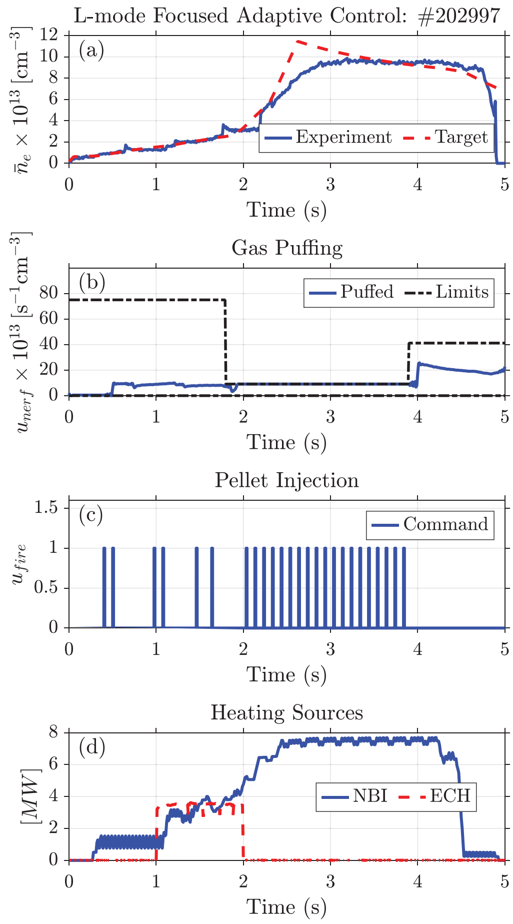

Figure 12. Density regulation in L-mode during discharge #202997: (a) time evolution of measured and target line-averaged density; (b) puffed gas commands and saturation limits; (c) commands sent to the pellet injection system; (d) NBI and ECH heating sources utilized in the discharge.

As shown in figure 12(a), the adaptive controller achieves robust tracking of the desired target for the first few seconds of the discharge. This tracking is achieved through a combination of gas puffing and pellet injection. At approximately $t = 2$ s, an increase in NBI heating triggers an L–H transition. During this phase, the gas actuation is constrained to an upper limit to avoid disrupting the transition. This constraint remains active until $t = 4.5$ s, effectively making pellet injection the only actuation mechanism. Notably, once gas puffing is constrained, the controller increases the pellet commands to compensate, showcasing the adaptive controller's ability to adjust to the current plasma and actuator dynamics. However, although multiple pellet commands are sent, no corresponding pellet-induced density response is observed, suggesting that the commanded pellets were not effectively delivered to the plasma. This loss of effectiveness is primarily attributed to hardware-level limitations inherent to cryogenic pellet injection systems. These include physical constraints on continuous pellet formation and extrusion at high requested repetition rates, as well as mechanical fracturing or severe pellet mass loss and degradation during high-speed transport through the HFS guide tubes [3, 4]. Similar behavior is observed in the subsequent discharges presented in this work, where the same physical interpretation holds.

Figure 13 illustrates the evolution of the parameter estimates in response to the plasma behavior. Prior to the L–H transition, the estimates remain close to their initial values, varying only slightly due to accurate density tracking. This behavior is consistent with the adaptation laws in (9) and (10), which dictate that parameter updates depend on both the tracking error and the computed feedback signal and therefore remain small when either quantity is near zero. After the transition, gas puffing is constrained while commanded pellet injection increases; however, the density tracking error grows despite the changes in feedback effort. In this regime, the combination of a large tracking error and increasing pellet input drives a decrease in the estimated pellet efficiency, $\hat{B}_{\mathrm{p}}$, via (10). Because the physical pellet injections are not delivered, this estimated efficiency continues to decrease (without reaching zero, as discussed in appendix A, Remark 1), which naturally commands a higher pellet request to compensate.

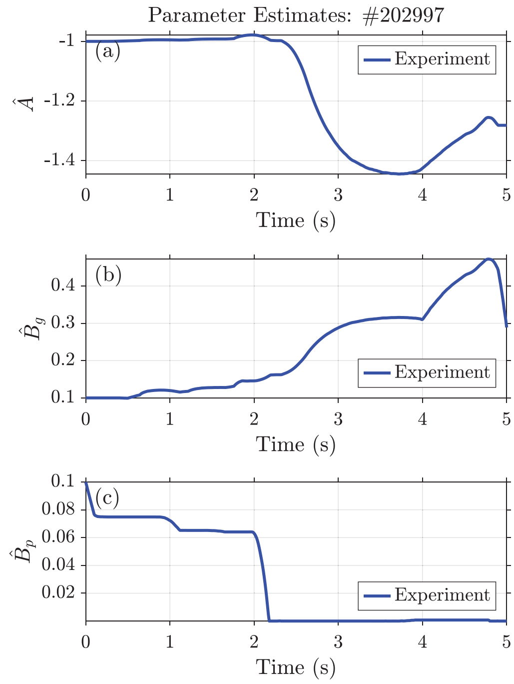

Figure 13. Evolution of parameter estimates during discharge #202997: (*a*) time evolution of density decay estimate; (*b*) time evolution of gas puffing efficiency estimate; (*c*) time evolution of pellet injection efficiency estimate.

Simultaneously, the gas command saturates below its feedforward component, effectively acting as a negative feedback contribution. According to (9), this constrained gas input coupled with the large tracking error results in an increase in the estimate $\hat{B}_\mathrm{g}$. While these parameter variations directly modulate the feedback signals via (6) and (7) to recover performance, it is important to clarify a fundamental characteristic of this framework. Although $\hat{A}$, $\hat{B}_\mathrm{p}$, and $\hat{B}_\mathrm{g}$ map to true physical properties of the plasma (e.g. decay rates and fueling efficiencies), the adaptive control laws are designed to guarantee the minimization of the density tracking error, rather than the convergence of the parameter estimates to their exact physical values. The estimates are dynamically adjusted by the controller strictly to maintain target tracking, as established in section 2 and demonstrated in appendix A. Finally, to maintain conciseness and avoid redundancy, the time traces of the parameter estimates for the subsequent experimental cases are omitted, as similar qualitative adaptation behavior is consistently observed.

Overall, the results of this discharge are significant because they demonstrate the effectiveness of the adaptive controller in tracking the target without an extensive tuning process. Conventional controllers, such as PID controllers, typically require a trial-and-error process in order to obtain the desired performance; however, the adaptive controller eliminates this need. Its 'out-of-the-box' performance allows for robust tracking of the target as long as the system is controllable through effective actuation mechanisms. Therefore, to further study its performance under limited-actuation conditions, such as those encountered in *H*-mode, additional discharges are conducted.

4.3. *H*-mode testing of adaptive controller

To ensure the adaptive controller tracks a feasible target, the reference model and target trajectory are generated following the same procedures explained in section 4.2. In addition, the internal parameters of the adaptive controller are set to $\beta = 0.5$, $\Gamma_A = 1$, $\Gamma_{B_\mathrm{g}} = 1$, and $\Gamma_{B_\mathrm{p}} = 1$, while the aggressiveness parameter α is varied based on the discharge. Finally, it is important to contextualize the fixed 75 ms time delay employed in the SP for these discharges. The operating scenarios presented in this and subsequent sections emulate ITER baseline conditions. Because the expected physical gas-transit delay in ITER is significantly longer than what the spatial scale and discharge duration of the DIII-D tokamak can practically accommodate, a scaled delay of 75 ms is emulated, based on the ratio of the energy confinement times between the two devices. Within this framework, treating the delay as a fixed constant is a physically sound assumption: the true transit time is primarily governed by the fixed geometric distance between the external actuators and the plasma edge, making it predominantly constant. While minor fluctuations can occur in a reactor setting due to evolving edge plasma parameters and neutral pressure gradients, the SP with conservative (not too aggressive) closed-loop performance is robust enough to

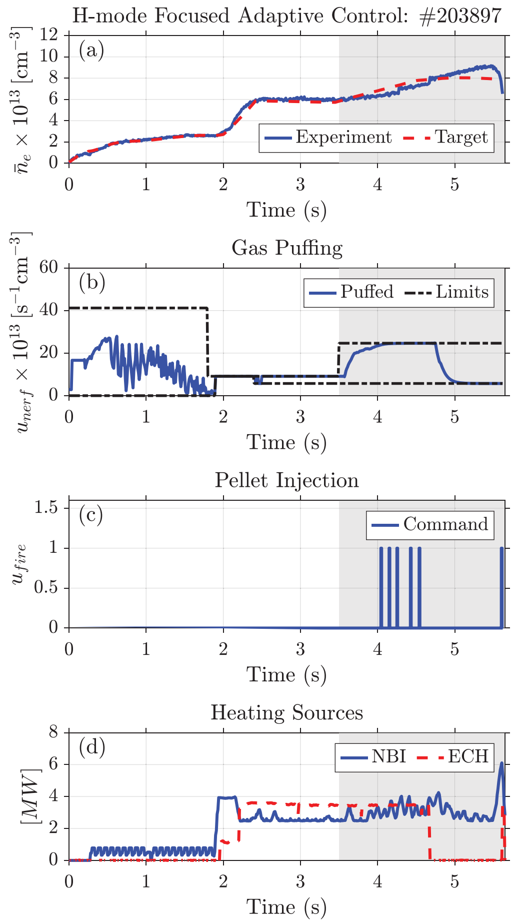

Figure 14. Density regulation in H-mode during discharge #203897: (a) time evolution of measured and target line-averaged density; (b) puffed gas commands and saturation limits; (c) commands sent to the pellet injection system; (d) NBI and ECH heating sources utilized in the discharge.

tolerate these small deviations [25, 26]. It is only when variations from the nominal delay grow substantially that meaningful performance deterioration is observed. Should such large deviations be anticipated in specific future reactor-grade scenarios, the control architecture could be augmented with a real-time delay-estimation algorithm [15, 27]. For the purposes of this study, the fixed-delay configuration provides a rigorous and highly representative environment for evaluating the adaptive controller under the plasma and actuator dynamics expected in future reactor-grade tokamaks.

Several discharges have been conducted to evaluate the performance of the adaptive controller. As the objective is to conduct the evaluation specifically in H-mode, consistent and repeatable L–H transitions are needed to ensure similar H-mode conditions for each discharge. To address this, a suitable L-mode density target is selected. This target, if tracked successfully in L-mode, results in a density value of $6 \times 10^{13}\,\mathrm{cm}^{-3}$ immediately after the L–H transition. While the adaptive controller could be used to track the L-mode target, DIII-D's PID controller is used as it is reliable and consistent in tracking density in L-mode when no input time delays are present. At $t = 3.5\,\mathrm{s}$, the PID controller is turned off and the proposed adaptive controller is activated concurrently with an emulated input time delay of 75 ms. To highlight the active controller during each phase of the discharge, the shaded regions of the figures indicate when the adaptive controller is used with time-delay emulation via NERF mode, while the unshaded regions represent PID-based control without NERF mode emulation. Finally, the procedure of using PID in L-mode and the adaptive controller in H-mode is applied to discharges #203897 and #203899.

Figure 14 illustrates the experimental results obtained from discharge #203897. As shown in figure 14(a), successful tracking of the L-mode density target leads to an L–H transition at the intended density of $6 \times 10^{13}\,\mathrm{cm}^{-3}$. Following this transition, the density target is maintained until $t = 3.5\,\mathrm{s}$, after which it is gradually increased from $6 \times 10^{13}\,\mathrm{cm}^{-3}$ to $8 \times 10^{13}\,\mathrm{cm}^{-3}$. This planned density ramp coincides with the activation of the adaptive controller and the introduction of the emulated input delays, enabling a focused investigation of the controller's performance. During the ramp phase, the adaptive

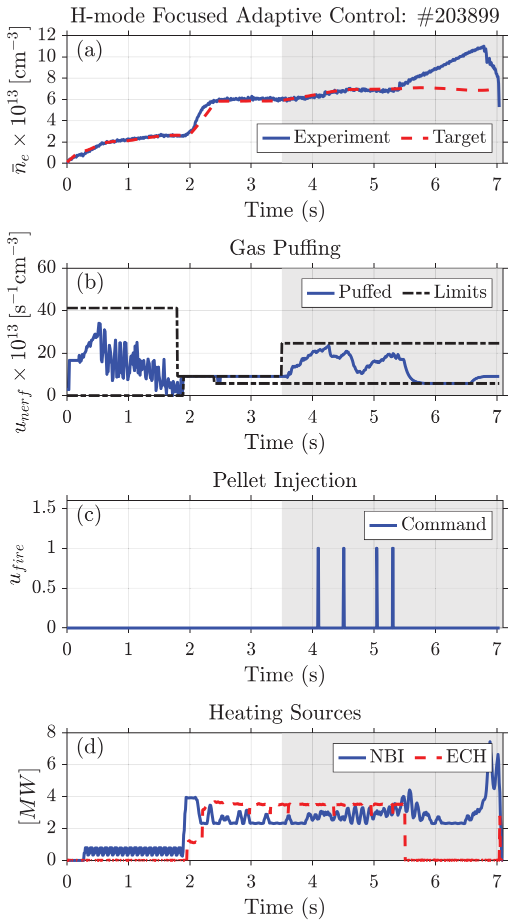

Figure 15. Density regulation in H-mode during discharge #203899: (a) time evolution of measured and target line-averaged density; (b) puffed gas commands and saturation limits; (c) commands sent to the pellet injection system; (d) NBI and ECH heating sources utilized in the discharge.

controller calculates suitable inputs to track the desired trajectory. As observed in figure 14(a), the measured density rises with the reference; however, the response lags behind the desired ramp. This lag is attributed to the gas actuation reaching its upper saturation limit and the pellet injection system failing to deliver all requested pellets (figures 14(b) and (c)). Once the density reaches the target, the controller reduces the gas actuation and stops requesting pellets. Although the desired target is achieved, the high plasma density, combined with changes in particle transport and plasma dynamics induced by pellet injection, triggers the built-in safety protocol of the ECH system. The subsequent shutdown of ECH leads to a sudden rise in plasma density, as shown in figure 14(d). The change in density due to ECH is consistent with the 'density pump-out' effect [28]. Although the adaptive controller attempts to counter this uncommanded rise by reducing actuation to the lower limit, the density continues to rise uncontrollably, ultimately leading to a disruption.

The outcome of this discharge provides valuable insight into the actuation limitations associated with density regulation. First, when the reference density changes more rapidly than the available actuators can respond, the plasma density lags behind the target trajectory. Second, once the density exceeds the target, the absence of a mechanism to reduce density in a controlled manner makes regulation more challenging and dependent solely on the natural confinement of the plasma. It is worth noting that if the density target is chosen too close to the trigger level for ECH machine protection, this inherent inability to actively pull the density down makes tripping the protection system highly likely during a routine overshoot. These observations suggest that selecting feasible targets that align with actuator mechanisms and capabilities is an important aspect to consider for reliable density regulation. To further examine this aspect, an additional discharge with a lower and slower density target is needed.

Figure 15 presents the experimental results for discharge #203899. Similar to discharge #203897, access to H-mode is achieved by tracking a specific target using PID-based control during the L-mode phase. Furthermore, motivated by the ECH-driven shutdown observed in discharge #203897, a reduced H-mode target is selected. Instead of ramping to $8 \times 10^{13}\,\mathrm{cm}^{-3}$, the new target is lowered to

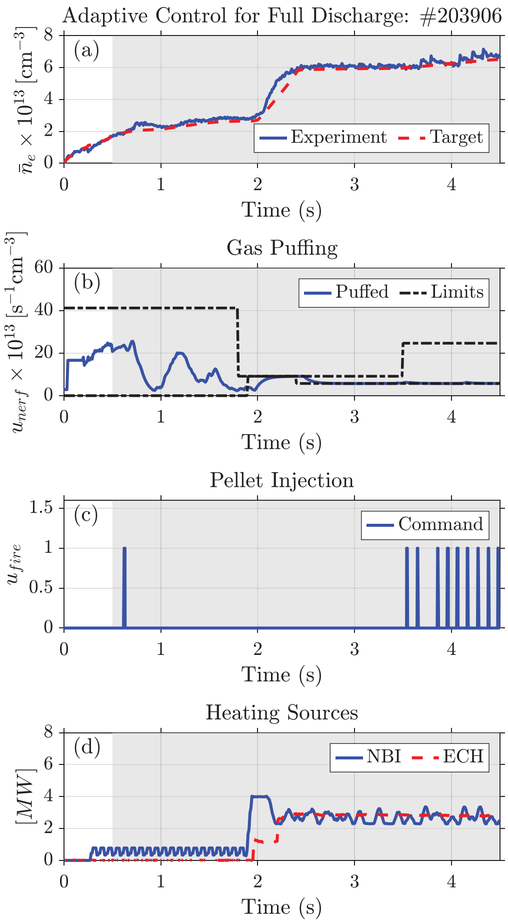

Figure 16. Density regulation during discharge #203906: (a) time evolution of measured and target line-averaged density; (b) puffed gas commands and saturation limits; (c) commands sent to the pellet injection system; (d) NBI and ECH heating sources utilized in the discharge.

$7 \times 10^{13}\,\text{cm}^{-3}$. At $t = 3.5\,\text{s}$, the adaptive controller activates and adjusts the inputs to track the reference density. In contrast to discharge #203897, the measured density closely follows the target trajectory, highlighting the controller's robust and accurate performance. The desired density of $7 \times 10^{13}\,\text{cm}^{-3}$ is successfully achieved and maintained. However, at $t = 5.5\,\text{s}$, a pellet injection results in the density exceeding the target value. In response, the controller immediately adjusts the actuation by reducing gas puffing to its lower saturation limit and ceasing pellet commands. Nevertheless, similar conditions to those observed in discharge #203897 are present, triggering a shutdown of the ECH system and a subsequent disruption.

Overall, the outcome of this discharge demonstrates the robustness of the adaptive controller in tracking the H-mode density target without any extensive tuning process. It is important to note that choosing a slower-rising target eliminated the problem of the measured density lagging behind the target due to actuation limitations. Moreover, reducing the target resulted in successfully maintaining a constant density for an extended period. Unfortunately, at high densities in these scenarios, pellet injection can alter the plasma dynamics in ways that may trigger other mechanisms leading to disruptions. This highlights the need to consider such constraints in control design. Nonetheless, as the controller demonstrates good tracking in both L-mode and H-mode (discharges #202997 and #203899), full-discharge testing is conducted.

4.4. Full discharge testing of adaptive controller

Figure 16 presents the adaptive regulation of the line-averaged electron density during discharge #203906. During plasma formation, small disturbances or oscillations can easily cause instabilities or disruptions. Therefore, the first 500 ms of the discharge is regulated by a PID controller using delay-free gas actuation. At $t = 500\,\text{ms}$, the proposed adaptive controller is activated and a 75 ms gas input delay is emulated through the NERF mode algorithm. The adaptive controller is configured to track the same L-mode target trajectory as in discharge #203899; however, the density target in H-mode is now gradually increased from $6 \times 10^{13}\,\mathrm{cm}^{-3}$ to $7 \times 10^{13}\,\mathrm{cm}^{-3}$ over a longer period of time. Finally, the results shown in figure 16 are limited to $t = 4.5\,\mathrm{s}$. Due to pre-discharge configuration issues, the adaptive controller does not operate as intended after $t = 4.5\,\mathrm{s}$, eventually leading to the same disruption observed in previous cases. Since this behavior was not observed in previous discharges or simulations and is not intrinsic to the controller design, the results after $t = 4.5\,\mathrm{s}$ have been omitted.

Once activated, the adaptive controller successfully tracks the density target across both L-mode and H-mode in the presence of emulated input delays by appropriately adjusting the control inputs. Achieving this level of regulation within a single discharge without extensive prior tuning highlights the controller's ability to accommodate time-delayed actuation and evolving operating conditions. These results underscore the applicability of the proposed approach for density regulation in future reactor-grade tokamaks.

5. Conclusion

An indirect adaptive control strategy has been developed to regulate the line-averaged electron density using coordinated gas puffing and pellet injection. The controller is specifically designed for reactor-grade scenarios, handling unknown plasma dynamics and input time delays by integrating MRAC with a SP. The proposed scheme was first validated in simulation using a multi-reservoir particle model fitted to DIII-D plasma density data. Importantly, this validation model is distinct from and more complex than the simplified model used for control design. The controller was then tested on the DIII-D tokamak under reactor-relevant actuator conditions. Experimental results demonstrate that the adaptive controller achieves robust density tracking despite actuation delays and unknown plasma parameters. The experiments also highlighted critical operational boundaries regarding pellet delivery failures. Because gas puffing will be restricted by neutral penetration to the plasma edge in future reactors, it will not be able to physically substitute for the core deposition of a missed pellet; thus, tracking degradation will be unavoidable without hardware redundancy. Future fault-tolerant control designs must therefore focus on keeping the closed-loop architecture well-behaved under active constraints. Specifically, integrating real-time pellet-detection observers (via microwave or optical diagnostics) would allow the system to instantly recognize delivery failures and freeze internal MRAC parameter adaptation and error-integral states, preventing unhelpful over-adaptation or parameter windup during the physical fault window. Furthermore, extensions of this work may include accounting for time-varying delays and generalizing the architecture to a MIMO framework to regulate multiple spatially distributed plasma quantities (i.e. multi-point control). Because continuous profile control is physically limited by the spatial authority of the actuators, this multi-point approach would naturally allocate gas puffing primarily for edge or pedestal regulation and pellet injection for core density peaking. Overall, this work provides a promising path toward reliable density control in future long-pulse and high-performance tokamak operation.

Acknowledgments

This material is based upon work partly supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences under Awards DE-SC0010661, DE-SC0010537 and DE-FC02-04ER54698, and by the 'KSTAR Experimental Collaboration and Fusion Plasma Research (EN2401-12)" through the Korea Institute of Fusion Energy (KFE).

Data availability statement

The data cannot be made publicly available upon publication because they are owned by a third party and the terms of use prevent public distribution. The data that support the findings of this study are available upon reasonable request from the authors.

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Appendix. A

This appendix derives the control laws presented in section 2. Following the control objective established in (5), the tracking error is defined as

$$\begin{aligned} e(t) =& \bar{n}_{\mathrm{e},m}(t) - (\bar{n}_{\mathrm{e}}(t) + e_{\mathrm{s}}(t)), \\ =& \bar{n}_{\mathrm{e},m}(t) - (\bar{n}_{\mathrm{e}}(t) + (\bar{n}_{\mathrm{e,s}}(t) - \bar{n}_{\mathrm{e,sd}}(t))), \\ =& \bar{n}_{\mathrm{e},m}(t) - \bar{n}_{\mathrm{e}}(t) - \bar{n}_{\mathrm{e,s}}(t) + \bar{n}_{\mathrm{e,sd}}(t). \end{aligned} \tag{A.1}$$

The derivative of the error can be written as

$$\dot{e}(t) = \dot{\bar{n}}_{\mathrm{e},m}(t) - \dot{\bar{n}}_{\mathrm{e}}(t) - \dot{\bar{n}}_{\mathrm{e,s}}(t) + \dot{\bar{n}}_{\mathrm{e,sd}}(t), \tag{A.2}$$

$$\begin{aligned} =& A_m\bar{n}_{\mathrm{e},m}(t) + B_m r(t) \\ & - A\bar{n}_{\mathrm{e}}(t) - B_{\mathrm{g}}u_{\mathrm{g}}(t-d_{\mathrm{g}}) - B_{\mathrm{p}}u_{\mathrm{p}}(t-d_{\mathrm{p}}) \\ & - \hat{A}(t)\bar{n}_{\mathrm{e,s}}(t) - \hat{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t) - \hat{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t) \\ & + \hat{A}(t)\bar{n}_{\mathrm{e,sd}}(t) + \hat{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t-d_{\mathrm{g}}) + \hat{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t-d_{\mathrm{p}}). \end{aligned} \tag{A.3}$$

By defining $\tilde{B}_{\mathrm{g}}(t)$, $\tilde{B}_{\mathrm{p}}(t)$, and $\tilde{A}(t)$ to be the differences between the estimates in (3)–(4) and the unknown parameters in (1), the following is obtained

$$\tilde{B}_{\mathrm{g}}(t) = \hat{B}_{\mathrm{g}}(t) - B_{\mathrm{g}}, \quad \dot{\tilde{B}}_{\mathrm{g}}(t) = \dot{\hat{B}}_{\mathrm{g}}(t), \tag{A.4}$$

$$\tilde{B}_{\mathrm{p}}(t) = \hat{B}_{\mathrm{p}}(t) - B_{\mathrm{p}}, \quad \dot{\tilde{B}}_{\mathrm{p}}(t) = \dot{\hat{B}}_{\mathrm{p}}(t), \tag{A.5}$$

$$\tilde{A}(t) = \hat{A}(t) - A, \quad \dot{\tilde{A}}(t) = \dot{\hat{A}}(t). \tag{A.6}$$

Adding $\hat{A}(t)\bar{n}_{\mathrm{e}}(t) - \hat{A}(t)\bar{n}_{\mathrm{e}}(t)$ to (A.3) and substituting (A.4)–(A.6), the derivative of the error simplifies to

$$\begin{aligned} \dot{e}(t) =& \hat{A}(t)\bar{n}_{\mathrm{e,sd}}(t) - \hat{A}(t)\bar{n}_{\mathrm{e}}(t) - \hat{A}(t)\bar{n}_{\mathrm{e,s}}(t) \\ & + A_m\bar{n}_{\mathrm{e},m}(t) + B_m r(t) - \hat{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t) - \hat{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t) \\ & + \tilde{A}(t)\bar{n}_{\mathrm{e}}(t) + \tilde{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t-d_{\mathrm{g}}) + \tilde{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t-d_{\mathrm{p}}). \end{aligned} \tag{A.7}$$

To develop the needed control laws and ensure stability, a candidate Lyapunov function is constructed to assess the performance of the system under feedback. It takes the following form

$$\begin{aligned} V(t) =& e(t)^{\mathrm{T}}e(t) + \mathrm{tr}\left(\tilde{A}(t)^{\mathrm{T}}\Gamma_A^{-1}\tilde{A}(t)\right) \\ & + \mathrm{tr}\left(\tilde{B}_{\mathrm{g}}(t)^{\mathrm{T}}\Gamma_{B_{\mathrm{g}}}^{-1}\tilde{B}_{\mathrm{g}}(t)\right) + \mathrm{tr}\left(\tilde{B}_{\mathrm{p}}(t)^{\mathrm{T}}\Gamma_{B_{\mathrm{p}}}^{-1}\tilde{B}_{\mathrm{p}}(t)\right). \end{aligned} \tag{A.8}$$

The derivative of this function is

$$\begin{aligned} \dot{V}(t) =& 2e(t)^{\mathrm{T}}\dot{e}(t) + 2\mathrm{tr}\left(\tilde{A}(t)\Gamma_A^{-1}\dot{\tilde{A}}(t)^{\mathrm{T}}\right) \\ & + 2\mathrm{tr}\left(\tilde{B}_{\mathrm{g}}(t)\Gamma_{B_{\mathrm{g}}}^{-1}\dot{\tilde{B}}_{\mathrm{g}}(t)^{\mathrm{T}}\right) + 2\mathrm{tr}\left(\tilde{B}_{\mathrm{p}}(t)\Gamma_{B_{\mathrm{p}}}^{-1}\dot{\tilde{B}}_{\mathrm{p}}(t)^{\mathrm{T}}\right). \end{aligned} \tag{A.9}$$

By substituting the derivative of the error presented in (A.7), the derivative of the Lyapunov function becomes

$$
\begin{aligned}
\dot{V}(t) = 2e(t)^{\mathrm{T}}\Big(&\hat{A}(t)\bar{n}_{\mathrm{e,sd}}(t) - \hat{A}(t)\bar{n}_{\mathrm{e}}(t) - \hat{A}(t)\bar{n}_{\mathrm{e,s}}(t) \\
&+ A_m\bar{n}_{\mathrm{e},m}(t) + B_m r(t) - \hat{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t) - \hat{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t)\Big) \\
&+ 2e(t)^{\mathrm{T}}\tilde{A}(t)\bar{n}_{\mathrm{e}}(t) + 2\mathrm{tr}\big(\tilde{A}(t)\Gamma_A^{-1}\dot{\tilde{A}}(t)^{\mathrm{T}}\big) \\
&+ 2e(t)^{\mathrm{T}}\tilde{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t-d_{\mathrm{g}}) + 2\mathrm{tr}\big(\tilde{B}_{\mathrm{g}}(t)\Gamma_{B_{\mathrm{g}}}^{-1}\dot{\tilde{B}}_{\mathrm{g}}(t)^{\mathrm{T}}\big) \\
&+ 2e(t)^{\mathrm{T}}\tilde{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t-d_{\mathrm{p}}) + 2\mathrm{tr}\big(\tilde{B}_{\mathrm{p}}(t)\Gamma_{B_{\mathrm{p}}}^{-1}\dot{\tilde{B}}_{\mathrm{p}}(t)^{\mathrm{T}}\big).
\end{aligned} \tag{A.10}
$$

To derive the desired control laws, the following trace identities are used

$$
e(t)^{\mathrm{T}}\tilde{A}(t)\bar{n}_{\mathrm{e}}(t) = \mathrm{tr}\left(\tilde{A}(t)\bar{n}_{\mathrm{e}}(t)e(t)^{\mathrm{T}}\right), \tag{A.11}
$$

$$
e(t)^{\mathrm{T}}\tilde{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t-d_{\mathrm{g}}) = \mathrm{tr}\left(\tilde{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t-d_{\mathrm{g}})e(t)^{\mathrm{T}}\right), \tag{A.12}
$$

$$
e(t)^{\mathrm{T}}\tilde{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t-d_{\mathrm{p}}) = \mathrm{tr}\left(\tilde{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t-d_{\mathrm{p}})e(t)^{\mathrm{T}}\right). \tag{A.13}
$$

Substituting equations (A.11)–(A.13) into (A.10) yields the following expression

$$
\begin{aligned}
\dot{V}(t) = 2e(t)^{\mathrm{T}}\Big(&\hat{A}(t)\bar{n}_{\mathrm{e,sd}}(t) - \hat{A}(t)\bar{n}_{\mathrm{e}}(t) - \hat{A}(t)\bar{n}_{\mathrm{e,s}}(t) \\
&+ A_m\bar{n}_{\mathrm{e},m}(t) + B_m r(t) - \hat{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t) - \hat{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t)\Big) \\
&+ 2\mathrm{tr}\left(\tilde{A}(t)\bar{n}_{\mathrm{e}}(t)e(t)^{\mathrm{T}} + \tilde{A}(t)\Gamma_A^{-1}\dot{\tilde{A}}(t)^{\mathrm{T}}\right) \\
&+ 2\mathrm{tr}\left(\tilde{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t-d_{\mathrm{g}})e(t)^{\mathrm{T}} + \tilde{B}_{\mathrm{g}}(t)\Gamma_{B_{\mathrm{g}}}^{-1}\dot{\tilde{B}}_{\mathrm{g}}(t)^{\mathrm{T}}\right) \\
&+ 2\mathrm{tr}\big(\tilde{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t-d_{\mathrm{p}})e(t)^{\mathrm{T}} + \tilde{B}_{\mathrm{p}}(t)\Gamma_{B_{\mathrm{p}}}^{-1}\dot{\tilde{B}}_{\mathrm{p}}(t)^{\mathrm{T}}\big).
\end{aligned} \tag{A.14}
$$

The condition $\dot{V}(t) \leqslant 0$ is satisfied by the following control laws

$$
u_{\mathrm{g}}(t) = \hat{B}_{\mathrm{g}}(t)^{-1}\left[\alpha_1 e(t) + \beta\left(\hat{A}(t)e(t) - \hat{A}(t)\bar{n}_{\mathrm{e},m}(t) + A_m\bar{n}_{\mathrm{e},m}(t) + B_m r(t)\right)\right], \tag{A.15}
$$

$$
\begin{aligned}
u_{\mathrm{p}}(t) = \hat{B}_{\mathrm{p}}(t)^{-1}\Big[&\alpha_2 e(t) + (1-\beta)\Big(\hat{A}(t)e(t) - \hat{A}(t)\bar{n}_{\mathrm{e},m}(t) + \\
&A_m\bar{n}_{\mathrm{e},m}(t) + B_m r(t)\Big)\Big],
\end{aligned} \tag{A.16}
$$

$$
\dot{\tilde{A}}(t)^{\mathrm{T}} = -\Gamma_A\bar{n}_{\mathrm{e}}(t)e(t)^{\mathrm{T}}, \tag{A.17}
$$

$$
\dot{\tilde{B}}_{\mathrm{g}}(t)^{\mathrm{T}} = -\Gamma_{B_{\mathrm{g}}}u_{\mathrm{g}}(t-d_{\mathrm{g}})e(t)^{\mathrm{T}}, \tag{A.18}
$$

$$
\dot{\tilde{B}}_{\mathrm{p}}(t)^{\mathrm{T}} = -\Gamma_{B_{\mathrm{p}}}u_{\mathrm{p}}(t-d_{\mathrm{p}})e(t)^{\mathrm{T}}. \tag{A.19}
$$

Substituting equations (A.15)–(A.19) into (A.14) yields

$$
\dot{V}(t) = -2e(t)^{\mathrm{T}}\alpha e(t). \tag{A.20}
$$

Since $\dot{V}(t) \leqslant 0$, the tracking error $e(t)$ and the parameter estimates $\hat{B}_{\mathrm{g}}(t)$, $\hat{B}_{\mathrm{p}}(t)$, and $\hat{A}(t)$ are bounded. To prove asymptotic stability, $\dot{V}(t)$ needs to be uniformly continuous. This can be proved by examining the boundedness of its derivative. The function $\ddot{V}(t)$ is computed as

$$
\ddot{V}(t) = -4e(t)^{\mathrm{T}}\alpha\dot{e}(t) \tag{A.21}
$$

$$
\begin{aligned}
= -4e(t)^{\mathrm{T}}\alpha\Big[&\hat{A}(t)\bar{n}_{\mathrm{e,sd}}(t) - \hat{A}(t)\bar{n}_{\mathrm{e}}(t) - \hat{A}(t)\bar{n}_{\mathrm{e,s}}(t) \\
&+ A_m\bar{n}_{\mathrm{e},m}(t) + B_m r(t) - \hat{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t) - \hat{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t) \\
&+ \tilde{A}(t)\bar{n}_{\mathrm{e}}(t) + \tilde{B}_{\mathrm{g}}(t)u_{\mathrm{g}}(t-d_{\mathrm{g}}) + \tilde{B}_{\mathrm{p}}(t)u_{\mathrm{p}}(t-d_{\mathrm{p}})\Big].
\end{aligned} \tag{A.22}
$$

By adding $\hat{A}(t)\bar{n}_{e,m}(t) - \hat{A}(t)\bar{n}_{e,m}(t)$ to the derivative, (A.22) can be presented as

$$\begin{aligned}\ddot{V}(t) = -4e(t)^{\mathrm{T}}\alpha\Big[&\hat{A}(t)\bar{n}_{e,sd}(t) - \hat{A}(t)\bar{n}_{e}(t) - \hat{A}(t)\bar{n}_{e,s}(t)\\ &+A_m\bar{n}_{e,m}(t) + B_m r(t) - \hat{B}_g(t)u_g(t) - \hat{B}_p(t)u_p(t)\\ &+\tilde{A}(t)\bar{n}_e(t) + \tilde{B}_g(t)u_g(t-d_g) + \tilde{B}_p(t)u_p(t-d_p)\\ &+\hat{A}(t)\bar{n}_{e,m}(t) - \hat{A}(t)\bar{n}_{e,m}(t)\Big].\end{aligned} \tag{A.23}$$

Taking into account that $e(t) = \bar{n}_{e,m}(t) - \bar{n}_e(t) - \bar{n}_{e,s}(t) + \bar{n}_{e,sd}(t)$, (A.23) is written as

$$\begin{aligned}\ddot{V}(t) = -4e(t)^{\mathrm{T}}\alpha\Big[&\hat{A}(t)e(t) - \hat{A}(t)\bar{n}_{e,m}(t) + A_m\bar{n}_{e,m}(t)\\ &+B_m r(t) - \hat{B}_g(t)u_g(t) - \hat{B}_p(t)u_p(t)\\ &+\tilde{A}(t)\bar{n}_e(t) + \tilde{B}_g(t)u_g(t-d_g) + \tilde{B}_p(t)u_p(t-d_p)\Big].\end{aligned} \tag{A.24}$$

From (A.20) and the formulation of the reference model, all variables inside (A.24) are bounded except for $\bar{n}_e(t)$, $u_g(t)$, and $u_p(t)$. Therefore, their boundedness needs to be proved. Examining equations (A.15) and (A.16), it is clear that the inputs are functions of $e(t)$, $\hat{A}(t)$, $\hat{B}_g(t)^{-1}$, $\hat{B}_p(t)^{-1}$, $\bar{n}_{e,m}(t)$, $r(t)$, A_m, and B_m, which are all bounded. Therefore, the inputs $u_g(t)$ and $u_p(t)$ are bounded as well. Then, it follows that $\bar{n}_e(t)$ is bounded as it depends on $u_g(t)$ and $u_p(t)$. As $\bar{n}_e(t)$ is bounded, $\ddot{V}(t)$ is bounded as well. Thus, $\dot{V}(t)$ is uniformly continuous. It follows from Barbalat's lemma that $\dot{V}(t) \to 0$, hence $e(t) \to 0$ as $t \to \infty$, which proves that the origin is an asymptotically stable equilibrium of the tracking error dynamics [29].

Remark 1. If $\hat{B}_g(t)$ or $\hat{B}_p(t)$ becomes zero, the control laws become singular. To avoid this, the adaptive laws for these estimates are modified so that they never reach zero. This is implemented using standard projection methods [14].

Remark 2. The present MISO architecture regulates a single scalar: the line-averaged electron density. This architecture could be extended to regulate multiple scalar quantities, such as distinct points of the density profile (multi-point control), using a MIMO architecture. Under this formulation, the system parameters would become matrix-valued, and the adaptive laws in (A.17)–(A.19) could be generalized through a Lyapunov-based derivation similar to that presented in appendix A. The matrix inversions in (A.15)–(A.16) would then be replaced by the Moore–Penrose pseudoinverse [14].

ORCID iDs

Hassan Al Khawaldeh 0009-0001-3187-9383
Andres Pajares 0000-0001-9251-9675
Sai Tej Paruchuri 0000-0003-2372-3888
Vincent Graber 0000-0003-2058-2258
Tariq Rafiq 0000-0002-2164-1582
June-Woo Juhn 0000-0002-9435-4349
Troy Pederson 0000-0003-3335-1686
Daisuke Shiraki 0000-0002-1452-6949
Francesca Turco 0000-0002-3723-0438

References

[1] Romanelli M, Parail V, da Silva Aresta Belo P, Corrigan G, Garzotti L, Harting D, Koechl F, Militello-Asp E, Ambrosino R and Cavinato M 2015 *Nucl. Fusion* **55** 093008
[2] Baylor L R, Combs S K, Jernigan T C, Houlberg W A, Owen L W, Rasmussen D A, Maruyama S and Parks P B 2005 *Phys. Plasmas* **12** 056103
[3] Combs S K, Baylor L R, Meitner S J, Caughman J B O, Rasmussen D A and Maruyama S 2012 *Fusion Eng. Des.* **87** 634–40
[4] Geulin E and Pégourié B 2021 Pellet core fueling in tokamaks, stellarators and reversed field pinches *ITC30 - The 30th Int. Toki Conf. on Plasma and Fusion Research* (E-conference)
[5] Kurihara K, Lister J B, Humphreys D A, Ferron J R, Treutterer W, Sartori F, Felton R, Brémond S and Moreau P (JET EFDA Contributors) 2008 *Fusion Eng. Des.* **83** 959–70

[6] Kudlacek O *et al* 2024 *Nucl. Fusion* **64** 056012
[7] Camacho E 2007 *Model Predictive Control* (Springer) published on 2007-05-15
[8] Orrico C A, van Berkel M, Bosman T O S J, Ceelen L, Heemels W P M H, Koechl F and Krishnamoorthy D 2025 *Nucl. Fusion* **65** 076041
[9] Bosman T O S J, Koechl F, Ho A, de Baar M R, Krishnamoorthy D and van Berkel M 2023 *Nucl. Fusion* **63** 126047
[10] Pajares A, Al Khawaldeh H R, Turco F, Shiraki D, Pederson T, Eldon D, Juhn J w, Graber V R, Schuster E and Paruchuri S T 2025 *2025 ITER-Relevant Density Control With Pellets in DIII-D's ITER Baseline Scenario* (*APS Division of Plasma Physics Meeting*)
[11] Graber V and Schuster E 2024 *Nucl. Fusion* **64** 086007
[12] Graber V and Schuster E 2020 Nonlinear adaptive burn control and optimal control allocation of over-actuated two-temperature plasmas *American Control Conf. (ACC)* pp 1411–6
[13] Krstić M, Kanellakopoulos I and Kokotović P V 1995 *Nonlinear and adaptive control design* Adaptive and Cognitive Dynamic Systems: Signal Processing, Learning *Communications and Control* vol 7 (Wiley)
[14] Nguyen N T 2022 *Model-Reference Adaptive Control: A Primer* (Springer)
[15] Bahill A 1983 *IEEE Control Syst. Mag.* **3** 16–22
[16] Yildiz Y, Annaswamy A, Kolmanovsky I V and Yanakiev D 2010 *Automatica* **46** 279–89
[17] Niculescu S I and Annaswamy A M 2003 *Syst. Control Lett.* **49** 347–58
[18] Kudláček O *et al* 2025 *Fusion Eng. Des.* **216** 115071
[19] Ploeckl B Lang P T Kircher M Bock A Gude A Janky F Sieglin B Suttrop W Treutterer W and T Z 2021 Fusion science and technology **77** 199–205
[20] Hirooka Y and Sakamoto M (TRIAM Group T) 2003 *J. Nucl. Mater.* **313–316** 588–594
[21] Maddison G P, Turner A, Fielding S J and You S 2005 *Plasma Phys. Control. Fusion* **48** 71
[22] Ehrenberg J, Andrew P, Horton L, Janeschitz G, de Kock L and Philipps V 1992 *J. Nucl. Mater.* **196-198** 992–6
[23] Graber V 2025 Divertor-safe nonlinear burn control with actuator allocation in ITER plasmas *Ph D Dissertation* (Lehigh University Bethlehem)
[24] Nise N S 2020 *Control Systems Engineering* 8th edn (Wile)
[25] Santacesaria C and Scattolini R 1993 *Automatica* **29** 1595–7
[26] Lee D, Lee M, Sung S and Lee I 1999 *J. Process Control* **9** 79–85
[27] Tan Y and Karim M N 2002 *IFAC Proc.* vol **35** pp 325–30
[28] Wang X, Mordijck S, Doyle E J, Rhodes T L, Zeng L, McKee G R, Austin M E, Meneghini O, Staebler G M and Smith S P 2017 *Nucl. Fusion* **57** 116046
[29] Khalil H 2002 *Nonlinear Systems* (Pearson Education, Prentice Hall)