ADVANCED CONTROL IN DIII-D: SUPERVISORY AND FAIL-SAFE ALGORITHMS FOR FUTURE REACTOR-GRADE TOKAMAKS

A. PAJARES, W. P. WEHNER, A. S. WELANDER, D. A. HUMPHREYS
General Atomics
San Diego, USA
Email: pajaresa@fusion.gat.com

E. SCHUSTER
Lehigh University
Bethlehem, USA

Abstract

Supervisory, fail-safe algorithms will be necessary in ITER and future reactor-grade tokamaks where achieving numerous control objectives will rely on plant supervision and efficient actuator handling. In this work, advanced control algorithms with relevance for the next generation of tokamaks have been successfully tested in DIII-D simulations and experiments. Two algorithms have been developed in this work which have supervisory and actuator-management capabilities. The first algorithm is a reference governor, which supervises the plasma state and accordingly updates the targets in real time to ensure plant safety under feedback. The second algorithm is an actuator manager which optimizes how the tokamak actuators are used to attain the control objectives, even if actuator failures happen. These results illustrate the importance of advanced control techniques to optimize the scientific and economic output desired from reactor-grade devices, while also protecting the high investment that these plants represent.

1. INTRODUCTION

Tokamaks are highly complex dynamical systems which require active control [1] in order to attain the necessary conditions for nuclear fusion to occur. Presently, tokamak operation is customarily enabled by means of control algorithms with different levels of sophistication. These controllers range from manually-tunable proportional-integral-derivative (PID) laws (some examples can be found in [2,3]), to model-based, multi-input multi-output designs (e.g. [4,5]). The main goal of such controllers in present devices is to enable their associated scientific programs by achieving and maintaining the necessary plasma conditions for physics studies. In the next generation of reactor-grade tokamaks (which includes ITER and fusion power-plants), even more complex algorithms will be a must due to the increasing device size, pulse length, and mission complexity of such machines. For example, advanced control algorithms have been recently designed with the ability to coordinate a set of controllers, and with the final goal of scheduling and supervising the plasma discharge (i.e., supervisory algorithms, see for example [6,7]). Another important aspect of an advanced control architecture is the development of actuator managers [8,9], which are algorithms that can select and optimize the use of scarce actuators to fulfill the overall control objectives during a discharge. Other advanced control functions include plant monitoring [10], event detection [11], and exception and fault-response handling [12], among others. The ultimate objective of a control architecture composed of all these elements is to provide a fail-safe control system with the capability of preventing and mitigating potentially dangerous conditions to the plasma and the device. Therefore, the development of advanced control techniques for tokamaks will be vital to protect the large investment that these machines represent, as well as the economic viability of nuclear fusion as a clean-energy source.

In this work, advanced control algorithms with supervisory and actuator-management capabilities have been developed and successfully tested in DIII-D simulations and experiments. Specifically, two algorithms are described in this paper. First, a reference governor has been designed to supervise the plasma state and accordingly update the control targets in real time to avoid future dangerous states and ensure plant safety. The reference governor has been tested in simulations using GSEvolve [13] for different ITER-relevant DIII-D scenarios (see Section 2). Second, an actuator manager has been developed which optimizes the use of the available actuators in the tokamak to attain the overall control objectives, even in the presence of actuator failures. The actuator manager is based on a nonlinear-optimization algorithm and has been experimentally tested in DIII-D for an advanced-tokamak scenario (see Section 3). A conclusion is presented in Section 4.

2. REFERENCE GOVERNOR

Generally speaking, a reference governor is an algorithm with the capability of updating control targets for a feedback controller. Examples include the plasma shape target that is desired during operation and controlled by a shape controller, the total plasma-current target regulated by a current controller using the tokamak ohmic coil, or the normalized beta target for a plasma-energy controller using auxiliary-heating sources. The goal of the reference governor is to maintain the system safe by temporarily changing these targets during transients or unexpected phenomena. System safety can be characterized, for instance, by keeping a minimum plasma-wall clearance, a maximum/minimum plasma current, or a maximum normalized beta that avoids magneto-hydrodynamic (MHD) instabilities. Transients or unexpected phenomena may include impurity influxes, confinement changes, and any other phenomena not envisioned during nominal operation of the tokamak. A general schematic comparing control loops with and without a reference governor are shown in Fig. 1.

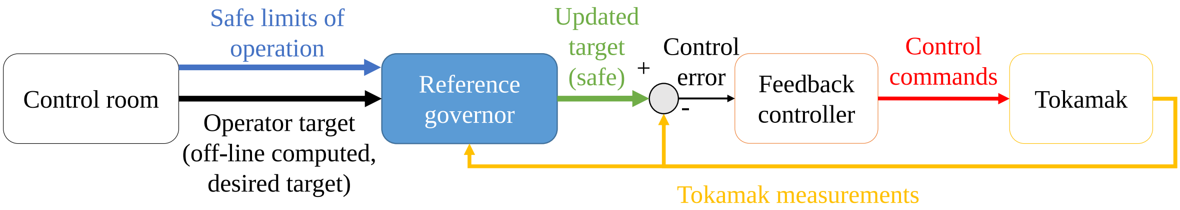

(a) Control loop with a reference governor

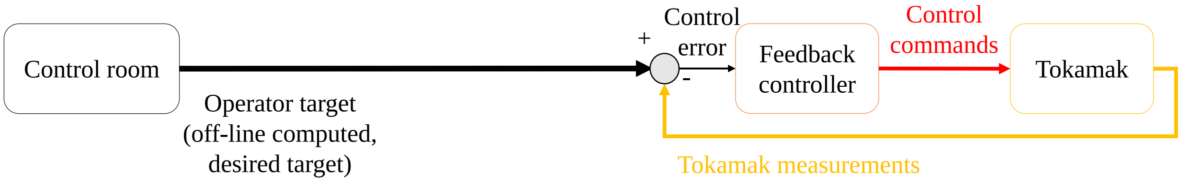

(b) Control loop without a reference governor

FIG. 1. General schematic of a control loop (a) with a reference governor (top) and, (b) without a reference governor (bottom). The target set by the operator in the control room is normally computed off-line and desired to attain a particular operating scenario. In the schematic on the top, such operator target is updated by the reference governor when needed to ensure system safety. On the other hand, in the traditional design at the bottom, the desired target is provided directly to the feedback controller through the control error, and may result in potentially unsafe conditions. It must be taken into account that the control commands are always sent by the feedback controller (which remains unchanged between both designs), and not by the reference governor.

The reference governor designed in this work focuses on **plasma shape and position control**. System safety is specified in terms of the following two sets of variables:

(a) Minimum plasma-wall distances (i.e., minimum gaps), so that plasma-wall contacts are prevented.
(b) Maximum vertical-instability growth rate, so that vertical-displacement events (VDEs) are avoided.

Therefore, the reference governor updates the plasma shape and position targets in real time so that the gaps and vertical-instability growth rate remain within safe limits at all times. Whenever possible, the updated target is set equal to the operator target, i.e., the desired shape for a particular operating scenario. In principle, the updates made by the reference governor are only temporary, unless they are absolutely critical for system safety. In addition, safety is ensured in the presence of disturbances in the total plasma current, plasma poloidal beta, and plasma internal inductance. These disturbances can model changes in plasma transport, MHD stability, or other machine conditions.

The reference governor is a model-based algorithm. It is based on an electromagnetic response model of the plasma surrounded by ohmic and poloidal-field coils as well as the conducting tokamak structures. Such model is provided by the General Atomics TokSys suite of codes [14], and in particular, by GSDesign [13]. The model provided by GSDesign and used within the reference governor is linear and time invariant, so it is particularly well suited for its use with a reference governor such as the one designed in [15]. This type of reference governor is formulated as a linear-optimization problem in real time. While the objective of such optimization is to make the updated shape target as close as possible to the desired shape target, the optimization also imposes hard constraints for system safety (i.e., minimum gaps and maximum vertical-instability growth rate) based on the present and future plasma state estimated by the electromagnetic response model.

Simulation tests of the reference governor have been carried out by means of the equilibrium evolution code GSEvolve [13]. In particular, the DIII-D tokamak is considered in two examples within the next subsections, where comparisons with a traditional control-loop without a reference governor are shown.

2.1. Modification of plasma-elongation target for VDE avoidance

The first simulation case is the prevention of VDEs in the DIII-D ITER-baseline scenario [16]. These types of plasmas have an ITER-similar shape but scaled down to fit within the DIII-D vessel. Parameters of interest for this scenario are given by: plasma current ~ 1 MA, poloidal beta ~ 0.3, toroidal field ~1.75 T, internal inductance ~ 1.3, and q_{95} ~ 4.

The simulation study in this section is an emulation of a previous control experiment in DIII-D reported in [17]. Here, the plasma elongation is chosen beforehand as the to-be-controlled variable. As shown in Fig. 2, the desired target for the plasma elongation is nominally ramped up from about 1.75 until 1.85 (shown as red dashed in the left figure, and denoted as “potentially unsafe” within the legend of the figure). However, the safe target calculated by the reference governor is a capped version of the desired target, limited to a value about 1.8 halfway through the simulation (shown as blue solid in the left figure). This is because the reference governor estimates that the maximum limits for the vertical-instability growth rate (set to 1500 rad/s and shown as magenta dotted in the right figure) may be violated. In fact, without the reference governor, the vertical-instability growth rate gets well beyond the 1500 rad/s limit, as shown in red dashed in the right figure. With the reference governor, the vertical-instability growth rate remains below such limit, as shown by the solid blue line.

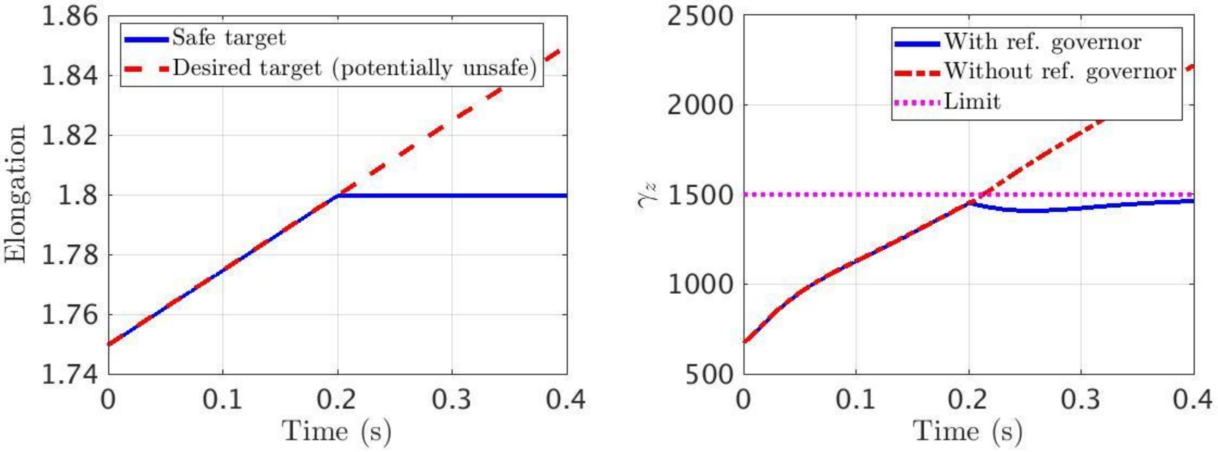

FIG. 2. Simulation of the reference governor to limit the plasma elongation in the ITER-baseline scenario in DIII-D. On the left, the reference governor limits the elongation target (solid blue) to keep the closed-loop system safe. On the right, the vertical-instability growth rate (solid blue) is kept below its safe limit (magenta). Without the reference governor (dashed red), the desired elongation target keeps increasing until it makes the system unsafe when the vertical-instability growth rate goes beyond the safe limit.

As a result of the elongation target updates made by the reference governor, the plasma shape attained is less elongated (see blue shape shown in Fig. 3) than without the reference governor (shown in dashed-red). For reference, the modelled first-wall of the DIII-D tokamak is shown in solid black.

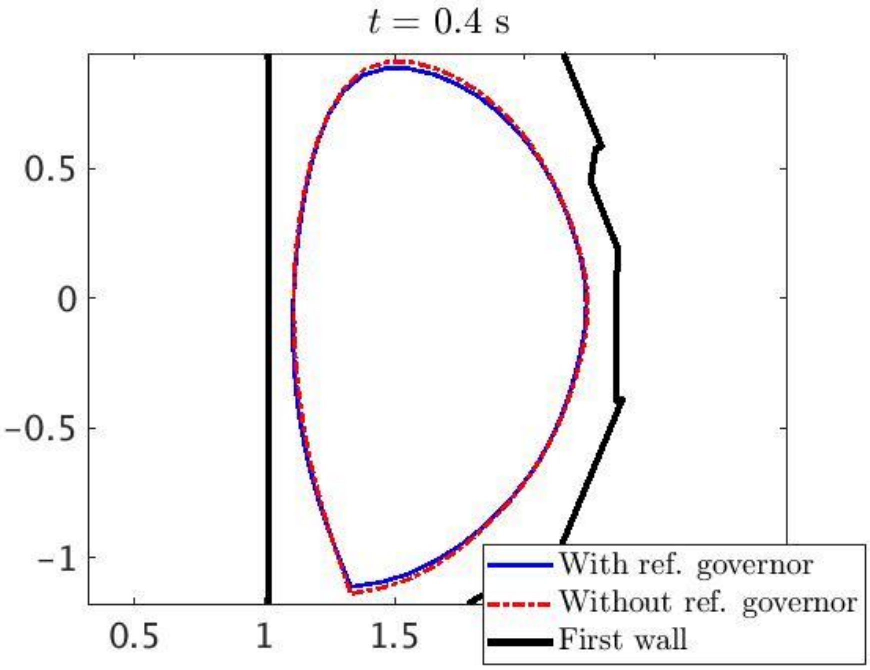

FIG. 3. Plasma shape at the end of the simulation (t = 0.4 s) for the ITER-baseline scenario in DIII-D.

2.2. Inner-gap target control for wall-limiting avoidance

In this second simulation study, the reference governor tries to avoid plasma limiting on the inner wall during an H-L back transition. This capability has been tested in simulations in the DIII-D high-q_{min} scenario, which is a candidate for ITER's *Q=5* mission [18]. Parameters of interest for this scenario are given by: plasma current ~ 1 MA, poloidal beta ~ 1.75, toroidal field ~ 1.7 T, internal inductance ~ 0.8, and q_{95} ~ 5.

The goal of these simulations is to move the plasma close to the wall while ensuring there is no plasma-wall contact. As shown in Fig. 4, the target for the inner gap is nominally set slightly above 2 cm (shown in red dashed). The left figure within Fig. 4 shows the simulation without the reference governor, whereas the figure on the right shows the simulation with the reference governor. Without the reference governor, the inner gap becomes zero slightly before 20 milliseconds, i.e., a plasma-wall contact briefly happens (see black solid line on left figure). On the other hand, as introduced at the beginning of Section 2, the reference governor uses the plasma model and potential range of disturbances to calculate a target that is safe, so it foresees the plasma-wall contact and it temporarily modifies and increases the inner-gap target (shown in blue solid on the right figure) to keep the inner gap evolution (shown in black solid) above the minimum limit (shown in magenta dotted). It can be observed that, after approximately 15 milliseconds, the reference governor estimates that the safe target can be set equal to the desired target. Therefore, the modification of the inner-gap target is only temporary, and the desired inner-gap target is eventually achieved. The temporary modification of the inner-gap target by the reference governor avoids any plasma-wall contact during the entire simulation (see Fig. 5, where the plasma boundary is shown in solid blue together with the first wall of the DIII-D tokamak in the inner-gap region). Meanwhile, without the reference governor (shown in red dashed in Fig. 5), the plasma is always closer to the wall, and a plasma-wall contact briefly happens at around 20 milliseconds, as introduced above. At around 60 milliseconds, both plasma shapes are the same because the inner-gap targets are also the same.

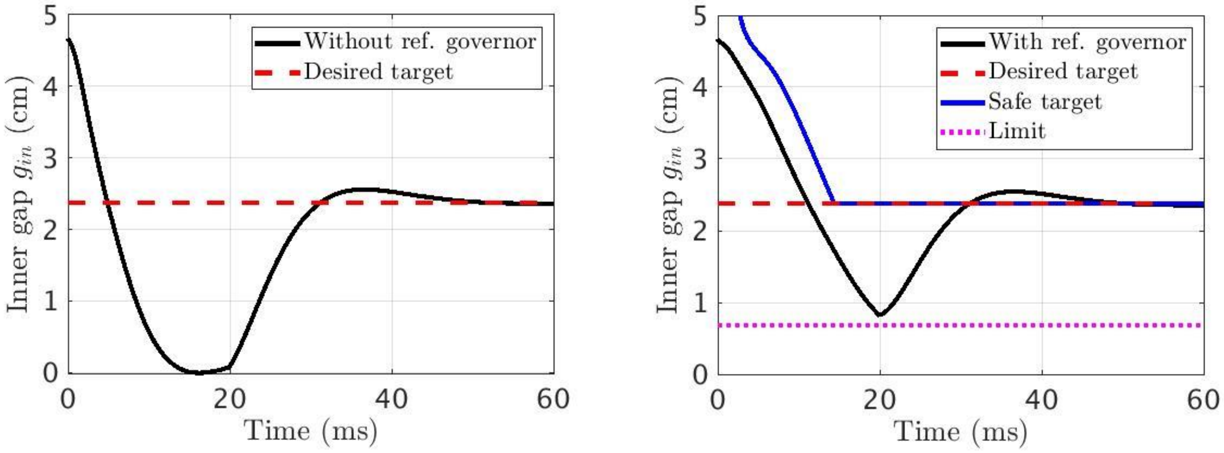

FIG. 4. Simulations without (left) and with (right) the reference governor to regulate the plasma inner-gap in the high-q_{min} scenario in DIII-D. On the left, the lack of a reference governor causes the plasma to hit the wall (inner gap become zero, see solid black line) when trying to achieve the desired but unsafe inner-gap target (shown in red-dashed). On the right, the reference governor increases the safe target (solid blue) temporarily in order to avoid the plasma-wall contact by keeping the inner gap above its limit (shown in magenta dotted).

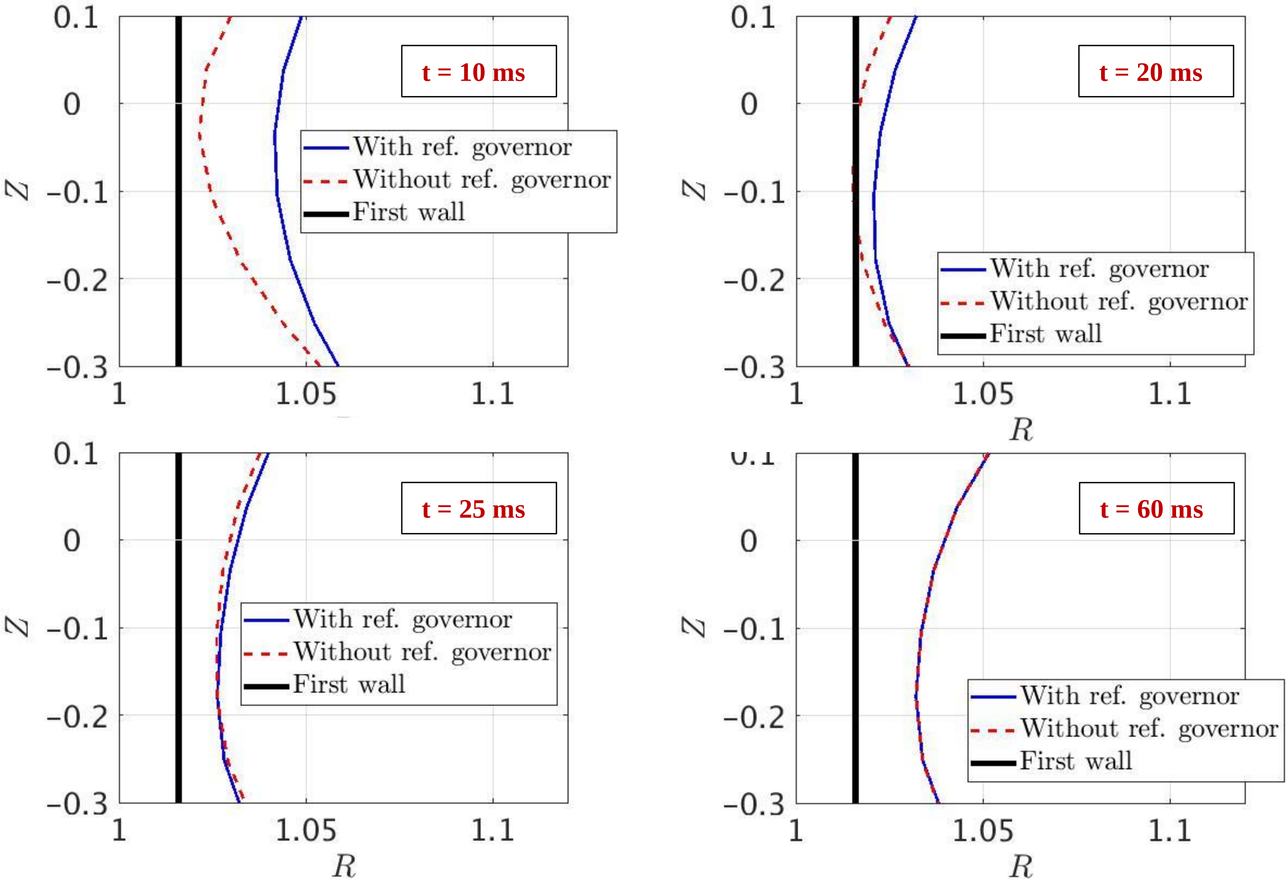

FIG. 5. Plasma boundary in the inner-gap region during simulations at t = 10, 20, 25, and 60 milliseconds with (solid blue) and without (dashed red) the reference governor to regulate the plasma inner-gap in the high-q_{min} scenario in DIII-D. The first wall of the DIII-D tokamak is shown in solid black.

3. ACTUATOR MANAGER

In general, an actuator manager is an algorithm that decides, in real time, how to effectively use the tokamak available actuators in order to attain a set of control objectives. The availability and/or saturation limits of such actuators may change in time, as well as the number and nature of control objectives that need to be attained, making the actuator-management problem highly involved. In addition, an actuator may be shared by more than one controller associated with a control objective. Therefore, in a sense, the actuator manager is also an arbitrator that decides how to use scarce actuators. An example of actuator management is the use of several gyrotrons for electron-cyclotron heating and current drive (ECH&CD) with diverse control objectives such as plasma initiation, current profile control, and MHD stabilization. Such gyrotrons may have different launch ports, frequencies, and available power. In addition, the different control objectives may happen synchronously (such as profile control and MHD stabilization) or asynchronously (e.g., plasma initiation will not happen simultaneously with profile control). Such high variability in the capabilities required for an actuator manager impose a strong need for flexibility in the algorithm, but also a robust and fail-safe design is required.

In this work, an actuator manager has been designed and implemented within the DIII-D plasma control system. The algorithm is based on a nonlinear-optimization scheme which is solved in real time, and whose preliminary version was previously tested in simulations [19]. The overall objective is to fulfill the various control objectives by means of requests that a set of feedback controllers send to the actuator manager. Moreover, constraints on actuator saturation limits and relative priorities of the control objectives are enforced. A schematic of a control loop with and without an actuator manager is shown in Fig. 6.

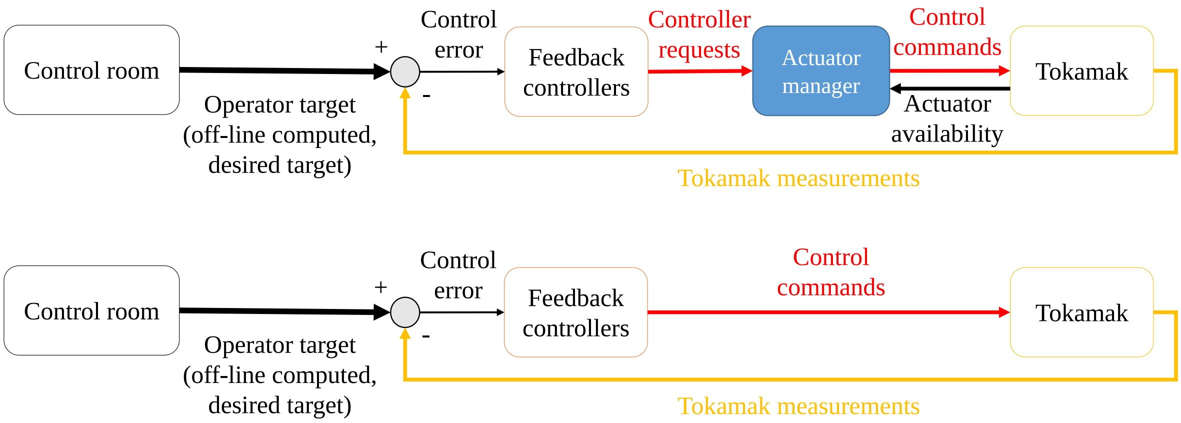

FIG. 6. General schematic of a control loop with an actuator manager (top) and without the actuator manager (bottom). In the schematic on the top, the control commands are calculated by the actuator manager based on the controller requests and actuator availability. On the other hand, in the traditional design at the bottom, the control commands are calculated directly by the feedback controllers. This may result in incompatibility of commands when an actuator is shared by more than one feedback controller.

The actuator manager (AM) has been tested in experiments in the DIII-D tokamak for high-q_{min} discharges [18], as shown in Fig. 7. The control objectives are regulating the plasma thermal energy, W, and the central safety factor, q_0, under feedback (FB). The actuators considered are ECH&CD and neutral-beam injection (NBI). To make the experimental test more challenging, the actuator availability changes in real time and requires actions by the actuator manager. Specifically, ECH&CD is not available before t = 2 s, and a failure in an off-axis NBI is emulated from t = 3 s until t = 3.75 s, as depicted by red and green stars in the top two subfigures of Fig. 7. The time evolution and targets for W and q_0 are shown by the blue and dashed-red lines, respectively, in the bottom subfigures. At t = 2 s, the actuator manager recognizes that EC is available and turns it on, reducing the off-axis NBI power. Similarly, when the off-axis NBI failure is emulated at t = 3 s, the actuator manager replaces the heating and non-inductive current deficits that arise from the loss of the off-axis NBI (see top right subfigure) by increasing the EC power (top left). Within the evolution of W, good tracking performance is observed until a small drop is found at t = 3 s (see bottom left subfigure), when the off-axis NBI fails. This drop is partially recovered by the effect of the FB + AM scheme, although saturation of the EC power (see top left subfigure) did not allow for fully achieving the W target (see bottom left subfigure). On the other hand, control of q_o worsens at t = 2 s (see bottom right subfigure) when the off-axis NBI power is decreased and is substituted by the EC power, which comes on at that time. This suggests a relatively small but noticeable inaccuracy in the actuator manager to carry out this actuator replacement. This is due to the current-drive model for the q_0 dynamics used within the actuator manager [19], which seems to overestimate the capability of ECCD to raise q_0 in the plasma conditions of this scenario.

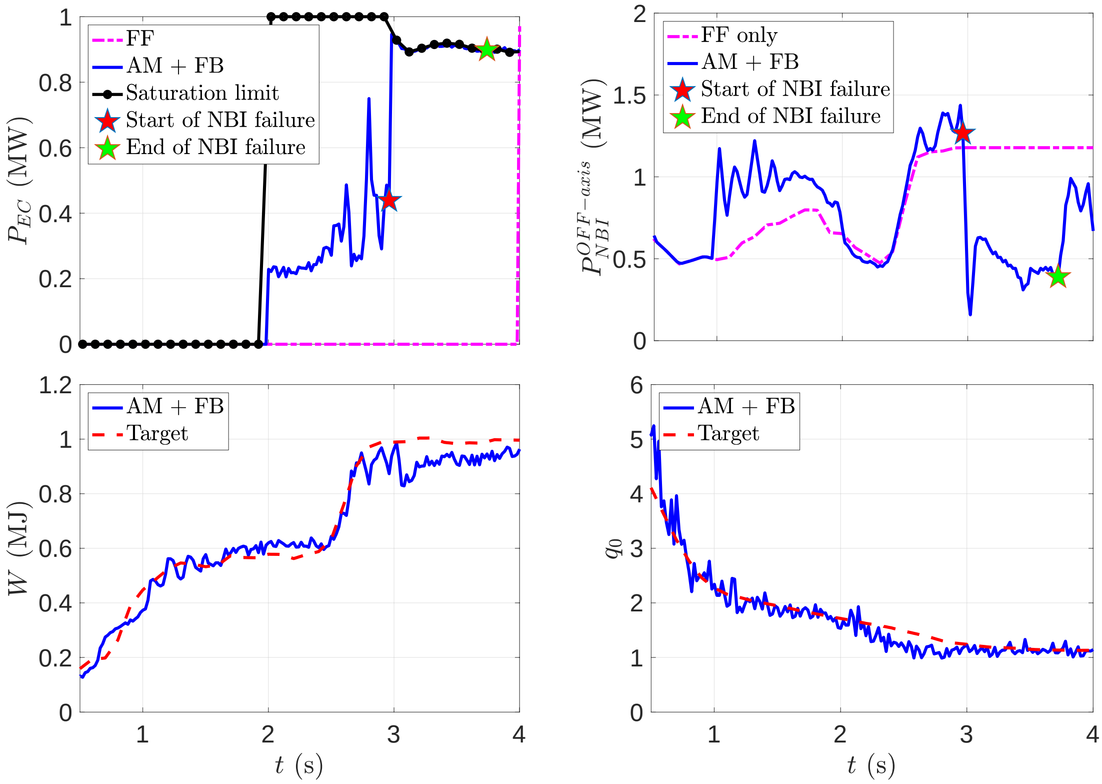

FIG. 7. Experiments on actuator management in DIII-D: EC power (top left), off-axis NBI power (top right), thermal energy evolution (bottom left), and central safety-factor evolution (bottom right). The start and end of a failure of an off-axis NBI is represented with red and green starts, respectively. Also, the FF only evolutions for EC and NBI powers (i.e. without FB control) are shown in magenta dashed-dotted.

4. CONCLUSION

The results reported in this paper illustrate the importance of advanced control techniques to optimize the scientific and economic output desired from reactor-grade devices. Supervisory, fail-safe algorithms will be necessary in ITER and future reactor-grade tokamaks, where achieving numerous control objectives will rely on plant supervision of potential failures and efficient actuator handling. The algorithms presented in this paper perform as expected in simulations and experiments in the DIII-D tokamak to avoid potentially unsafe conditions that may lead to disruptions and/or machine damage. However, further work is required to expand these algorithms for their use in a broader array of control problems. These control problems must not be treated independently in a future, integrated reference-governor + actuator-manager design, but in a combined, simultaneous fashion. Other algorithms with monitoring and exception handling capabilities should also be added to a reference governor and actuator manager in order to provide a full architecture with all the functionalities needed for reactor-grade tokamaks. It must be noted that the model-based nature of these algorithms would allow a relatively easy implementation in different devices. However, the models need to have predictive capability and, therefore, require a minimum level of accuracy and complexity which may increase the computational cost of these control algorithms. Finally, full experimental validation of these algorithms in other scenarios and devices must be sought to assess the robustness of these algorithms to various plasma conditions and tokamak configurations.

ACKNOWLEDGEMENTS

This material is based upon work supported by General Atomics internal funding and the United States (US) Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award(s) DE-FC02-04ER54698 and DE-SC0010661.

REFERENCES

[1] WALKER, M. L. et al, Introduction to Tokamak Plasma Control, 2020 American Control Conference (Denver, CO, USA), pp. 2901-2918 (2020).

[2] HYATT, A. W. et al, Designing, Constructing and Using Plasma Control System Algorithms on DIII-D, IEEE Transactions on Plasma Science **42** (2013) 1-6.

[3] NOUAILLETAS, R. et al, The WEST plasma control system: Integration, commissioning and operation on first experimental campaigns, Fusion Engineering and Design **146** (2019) 6-8.

[4] MELE, A. et al, MIMO shape control at the EAST tokamak: Simulations and Experiments, Fusion Engineering and Design **146** Part A (2019) 1282-1285.

[5] BOYER, M. et al, First-principles-driven model-based current profile control for the DIII-D tokamak via LQI optimal control, Plasma Physics and Controlled Fusion **55** 105007 (2013).

[6] BLANKEN, T. C. et al, Real-time plasma state monitoring and supervisory control on TCV, Nucl. Fusion **59** 026017 (2019).

[7] MALJAARS, E. et al, Control of the tokamak safety factor profile with time-varying constraints using MPC, Nuclear Fusion **55** 023001 (2015).

[8] MALJAARS, E. and FELICI, F., Actuator allocation for integrated control in tokamaks: architectural design and a mixed-integer programming algorithm, Fusion Engineering and Design **122** (2017) 94-112.

[9] KUDLACEK, O. et al, Developments on actuator management, plasma state reconstruction, and control on ASDEX Upgrade, Fusion Engineering and Design **171** 112563 (2021).

[10] MARTIN, V. et al, A Cognitive Vision System for Nuclear Fusion Device Monitoring, International Conference on Computer Vision Systems, Lecture Notes in Computer Science, vol 6962. Springer, Berlin, Heidelberg.

[11] SAMMULI, B. et al, Approaches to tokamak off-normal event detection and response at DIII-D, KSTAR, and EAST, Fusion Engineering and Design **8** (2010) 456-460.

[12] EIDIETIS, N. W. et al, Implementing a finite-state off-normal and fault response system for disruption avoidance in tokamaks, Nuclear Fusion **58** 056023 (2018).

[13] WELANDER, A. S. et al, Closed-loop simulation with Grad-Shafranov equilibrium evolution for plasma control system development, Fusion Engineering and Design **146** Part B (2019) 2361-2365.

[14] HUMPHREYS, D. A. et al, DIII-D Integrated plasma control solutions for ITER and next-generation tokamaks, Fusion Engineering and Design **83** Issues 2-3 (2009) 193-197.

[15] GILBERT, E. G. and KOLMANOVSKY, I., Fast reference governors for systems with state and control constraints and disturbance inputs, International Journal of Robust and Nonlinear Control **9** (1999) 1117-1141.

[16] TURCO, F. et al, The causes of the disruptive tearing instabilities of the ITER Baseline Scenario, Nucl. Fus. **5** 106043 (2018).

[17] BARR, J.L. et al, Development and experimental qualification of novel disruption prevention techniques on DIII-D, Nuclear Fusion **61** 126019 (2021).

[18] BUTTERY R. J. et al, DIII-D Research to prepare for Steady State Advanced Tokamak Power Plants, Journal of Fusion Energy (2019) 38:72-111.

[19] PAJARES A. et al, Actuator Management via Real-time Optimization for Integrated Control in Tokamaks, European Physics Society Annual Conference on Plasma Physics, July 8-12, Milan, Italy, 2019.