PAPER • OPEN ACCESS

Real-time plasma state monitoring and supervisory control on TCV

To cite this article: T.C. Blanken *et al* 2019 *Nucl. Fusion* **59** 026017

View the article online for updates and enhancements.

Real-time plasma state monitoring and supervisory control on TCV

T.C. Blanken[1] , **F. Felici**[2] , **C. Galperti**[2], **N.M.T. Vu**[2] , **M. Kong**[2],
O. Sauter[2], **M.R. de Baar**[1,3], **The EUROfusion MST1 Team**[a]
and The TCV Team[b]

[1] Department of Mechanical Engineering, Control Systems Technology Group, Eindhoven University of Technology, PO Box 513, 5600 MB Eindhoven, Netherlands

[2] École Polytechnique Fédérale de Lausanne (EPFL), Swiss Plasma Center (SPC), CH-1015 Lausanne, Switzerland

[3] DIFFER—Dutch Institute for Fundamental Energy Research, PO Box 6336, 5600 HH Eindhoven, Netherlands

E-mail: federico.felici@epfl.ch

Abstract

In ITER and DEMO, various control objectives related to plasma control must be simultaneously achieved by the plasma control system (PCS), in both normal operation as well as off-normal conditions. The PCS must act on off-normal events and deviations from the target scenario, since certain sequences (chains) of events can precede disruptions. It is important that these decisions are made while maintaining a coherent prioritization between the real-time control tasks to ensure high-performance operation.

In this paper, a generic architecture for task-based integrated plasma control is proposed. The architecture is characterized by the separation of state estimation, event detection, decisions and task execution among different algorithms, with standardized signal interfaces. Central to the architecture are a plasma state monitor and supervisory controller. In the plasma state monitor, discrete events in the continuous-valued plasma state are modeled using finite state machines. This provides a high-level representation of the plasma state. The supervisory controller coordinates the execution of multiple plasma control tasks by assigning task priorities, based on the finite states of the plasma and the pulse schedule.

These algorithms were implemented on the TCV digital control system and integrated with actuator resource management and existing state estimation algorithms and controllers. The plasma state monitor on TCV can track a multitude of plasma events, related to plasma current, rotating and locked neoclassical tearing modes, and position displacements.

In TCV experiments on simultaneous control of plasma pressure, safety factor profile and NTMs using electron cyclotron heating (ECH) and current drive (ECCD), the supervisory controller assigns priorities to the relevant control tasks. The tasks are then executed by feedback controllers and actuator allocation management. This work forms a significant step forward in the ongoing integration of control capabilities in experiments on TCV, in support of tokamak reactor operation.

1. Introduction

In long-pulse scenarios in tokamak reactors, multiple plasma quantities such as current, shape, kinetic profiles, impurities and various magnetohydrodynamic (MHD) activity need to be controlled simultaneously in real time [1–3]. The plasma conditions must stay at the target references despite unknown or unpredictable disturbances such as MHD activity or actuator/diagnostic failure. It is envisioned that multiple feedback controllers will operate concurrently, and that each will control a subset of the relevant plasma quantities [4, 5].

What are the requirements for the plasma control system (PCS) for a tokamak reactor? First, the PCS should execute the pulse schedule during plasma operations. Second, the PCS should decide which plasma quantities must be controlled subject to not only the pulse schedule but also to the state of the plasma and the machine. Third, unforeseen events (such as MHD activity, exceeding operational limits, and deviations from the target scenario, see [3, 6]) must be suppressed, controlled or mitigated by the PCS [4, 7, 8] to maintain plasma performance and stability. If suppression of instabilities or rejection of disturbances is not possible, the PCS must decide on future action, e.g. pursue a lower performance target or initiate an emergency ramp-down.

By definition, a reactor will have limited actuation resources, and it is therefore foreseen that actuators will be shared between the controllers that execute the various control tasks [3]. For example, NBI, ICRH and EC systems provide heating and current drive (H&CD), allowing control of the plasma pressure, the safety factor profile and the MHD modes [9, 10]. Depending on the plasma state and actuator availability, not all control objectives may be achieved simultaneously, so the PCS must prioritize the control tasks and allocate actuation resources to the tasks [3].

In present day tokamak control systems, the integration of various plasma controllers and plasma supervision for the internal plasma quantities is expanding. Although advanced control of e.g. pressure and shape is done routinely, the real-time decision-making is uncommon.

Often, individual control algorithms are developed and used for specific experiments, and are not operated concurrently with other controllers. Significant research has recently been conducted towards real-time control of multiple quantities simultaneously; for instance at TCV [11–14], DIII-D [15–17] and in simulations for ITER [18]. While these represent important steps towards integrated plasma control, they lack dedicated task coordination logic for actuator sharing management among multiple (different) controllers. Furthermore, significant work has been done at ASDEX-Upgrade in switching actuation resources between multiple controllers for simultaneous pressure and NTM control

[19]. While this marks a major step in integrated control, the resource switching is hardcoded for the intended experiment, based on a single trigger. Accommodating additional functionality in a specific implementation often requires substantial rewriting of existing code. Moreover, in present-day practice, signals of (combinations of) individual diagnostics are used to trigger specific control/event handling algorithms. While recent research on real-time off-normal event handling on DIII-D and KSTAR [20] provides the required real-time decision-making, this linked individual diagnostics to specific event handling algorithms and actuators. The links between diagnostics and event handling are done in an ad-hoc fashion, and a modification of the event triggering logic requires manual adjustments of the control code. Also, algorithms are difficult to port from one tokamak to another since they depend on specific diagnostics and actuators.

Instead, it is desirable to have a generic architecture separating

- the estimation of the full generic plasma state from tokamak-specific diagnostics,
- the detection of events in the plasma (or tokamak subsystems) based on the generic state description rather than specific diagnostic signals, and
- the decision logic on how to respond to these events,

as is also elaborated in [3]. This allows for programming arbitrary decisions based on combinations of multiple events.

The PCS of future tokamaks is responsible for high-level discharge supervision, meaning segment scheduling, managing control objectives, but also off-normal event handling. It is important to note that these functionalities must be performed in real-time for long pulse scenarios, as opposed to inter-discharge decision-making common in today's tokamak experiments. The requirements for the supervisory control come from the quantities that should either be controlled, or the quantities that should be monitored since the PCS should respond to them. Also, the set of requirements is different for a small and flexible experimental tokamak like TCV (operated in many different scenarios with different experimental needs on a day), than for a long-pulse experiments like ITER (performing a few scenarios for exploration) or a reactor like DEMO (performing one scenario). Because of the inherent physics coupling between many plasma quantities, information of the plasma state is needed for real-time control decisions. It is desirable to centralize all *high-level* plasma decisions so that decisions may be taken by a single algorithm that has the maximally available information about the plasma and the relevant tokamak subsystems [7, 21]. Here, *high-level* relates to the integrated operation and coordination across all plasma control objectives and algorithms, as opposed to the *low-level* operation of individual control algorithms and hardware control.
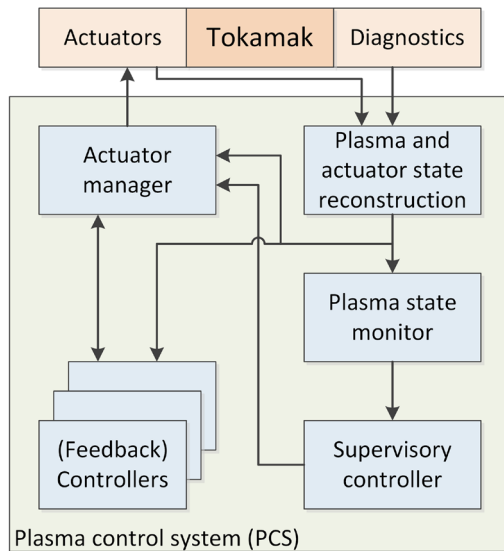
Figure 1. Block diagram of (a selection of) envisioned high-level PCS functionalities. Besides standard feedback control loops, plasma state monitoring and supervision provide high-level decisions regarding multiple control tasks. Moreover, an actuator management algorithm dynamically assigns actuators to controllers, see [22]. These components for pulse supervision and control were also defined in [3, 5]. However, in this paper and in [22], signal interfaces for integrated control are defined explicitly.

The aim of this paper is twofold. First, a conceptual framework for integrated control of tokamak plasmas employing full plasma state monitoring and supervisory control is presented. Second, the results of plasma state monitor and the supervisory controller at TCV applied to experimental simultaneous MHD and kinetic profile control with real-time decisions based on the NTM state are reported.

The conceptual framework has an architecture that is characterized by a strict separation between:

(i) a suite of reconstruction algorithms for the plasma and actuator states,
(ii) a finite-state machine representation for the monitoring of these states,
(iii) a centralized supervisory controller for enabling and prioritizing control tasks, and
(iv) the algorithms responsible for the execution of the control tasks, being (feedback) controllers and actuation resource management.

The ordering in the above list is intentional: each component sequentially feeds information into the next. In figure 1, we show a diagram of several high-level software components in a PCS, corresponding to the above list.

One aim of the real-time state monitoring is to automate the offline analysis of plasma events, such as carried out in [6, 23], for the purpose of plasma control. The result of high-level decisions executed by the supervisory controller consists of the real-time enabling prioritized control tasks, subject to the plasma state and available resources. Each control task represents a specific objective that a (feedback) control

algorithm should accomplish in the plasma (see [22]). Yet, the control task does not contain explicit directives on how a controller must execute this task, such as which actuators and control laws to use. The control task is therefore per definition *tokamak-agnostic*.

The development and application of the architecture at TCV is a collaborative effort. This paper deals primarily with the plasma state monitoring and supervisory controller, as the (compatible) actuator manager and interfaces to feedback controllers and actuators are reported in [22]. One of the difficulties encountered during development, implementation and integration of these complex integrated control algorithms on the TCV control system [24] is that existing input-output signal interfaces were different among various controllers. This proved to be impractical and confusing, particularly while asserting genericity and anticipating future expansions of the control capabilities. Therefore, all signal interfaces between the finite-state machine models, the supervisory controller, the actuator management algorithm and the control algorithms are standardized [22]. This provides ease of development and scalability to additional actuators and control algorithms. The standardization defines the signal dimensions, as well as their meaning in terms of either *tokamak-agnostic* quantities (e.g. kinetic profiles, ECH deposited power and deposition location) or *tokamak-specific* quantities (e.g. ECE data, power supply voltages and mirror angles of EC launchers).

The remainder of this paper is structured as follows: in section 2, the generic software architecture containing the state estimation, state monitor and supervisory control algorithms is discussed. The generic design of the state monitor and the supervisory controller are described in sections 3 and 4, respectively. In section 5, the implementation of the state monitor and supervisory controller on TCV's control system, as well as results of finite state monitoring on TCV discharges are presented. Experimental results with the application of the supervisory control in a TCV discharge are presented. The results of the latter sections are discussed in section 6. Finally, conclusions and an outlook are given in section 7.

2. Generic plasma control system software architecture

In this section, the proposed architecture for integrated plasma control on a PCS is presented and motivated. This architecture facilitates concurrent execution of multiple control tasks using multiple actuators. It is characterized by a strict separation of components based on their functionality as well as their degree of independence to the specifics of the tokamak diagnostics and actuators. First, the concept of control tasks is explained and their usage for handling multiple control tasks is motivated. Second, the proposed control system architecture is discussed. In particular, the functionality and input-output signals of the software components that are described in this paper are highlighted. Last, the design choice of using *tokamak-agnostic* signal interfaces between the algorithms is motivated.

2.1. Task-based supervisory control

In order to facilitate real-time prioritization of multiple control objectives, an abstraction of the representation of these objectives is proposed, called control tasks. As explained in more detail in [22], a control task is defined as a specific objective that should be accomplished in the plasma. Examples of various control tasks are confinement mode control, NTM suppression, NTM preemption, sawtooth pacing, safety factor profile control and impurity removal control. The supervisory control algorithm enables and prioritizes each control task, based on the discharge time, pulse schedule, the state of the plasma as well as the state of the actuators. These priorities represent the (relative) importance of the tasks. Many other research has addressed these tasks, yet with ad-hoc implementations for the specific task and binary prioritizations of tasks, see e.g. [19].

The design choice of abstracting the task execution to facilitate the coordination of integrated plasma control and exception handling has several advantages compared to existing solutions. These are:

- the supervisory decision logic on task priorities may be designed and interpreted more easily, since these decisions do not involve the specifics of the task execution by the responsible controller,
- the tasks are generic and similar among different tokamaks, such that the generic supervisory controller may be used for different devices,
- different feedback controllers may be used interchangeably for a given control task, without changing the supervisory logic,
- since functionalities (e.g. state estimation, task decision logic and task execution) are disentangled and reside in different algorithms with generalized input-output interfaces, the algorithms may be developed separately. This enhances interchangeability and portability of algorithms.

Ultimately, these priorities are taken into account by the actuator management algorithm, which assigns the scarce actuators to the tasks. These latter tasks are executed by controller algorithms given the assigned resources. In the actuator manager, the priorities are used as input to an optimization problem, together with factors such as actuator availability, (possibly conflicting) requirements, and actuation requests from controllers, to assign actuation resources to each task. Examples of such optimization algorithms are presented in [22, 25, 26].

Still, while the supervisory controller determines the priority of tasks, it does not produce explicit commands and directives on how the controllers must accomplish this task, e.g. which specific actuators and control laws to use. This abstraction of the objectives decouples the design of the supervisory control algorithm from the specifics of the task execution, which is handled by the controllers. When enabled, each control task will be executed by one (feedback) controller algorithm.

2.2. Control system software architecture

Different architectures for real-time supervisory control and actuator allocation are possible [25]. Although distributed supervisory control among the controllers is more easy to design for local optimality, they may take decisions that are globally contradictory or counterproductive. For instance, for the suppression of NTMs using ECH it may be needed to temporarily reduce the plasma kinetic pressure, overriding the normal reference for a pressure controller. In contrast, all decision-making can be merged in a single centralized supervisory controller, avoiding contradictory actions by consistent globally optimal decision-making. In this work, the latter centralized supervisor is considered.

In figure 2, a block diagram of the software components and signals is presented, which includes three layers. These layers, the software components, their functions and their input and output signals are introduced next.

2.2.1. Tokamak-agnostic signal definitions between algorithms. The components seen in figure 2 are grouped among three distinct layers, namely a *tokamak-specific layer*, an *interface layer* and a *tokamak-agnostic layer*. Algorithms in the *interface layer* form the boundary between the signals from device-specific hardware (e.g. the diagnostics and actuators in the *tokamak-specific layer*) and signals in the *tokamak-agnostic layer*. The algorithms in the *tokamak-agnostic layer* are characterized by an input-output interface where signals are defined in terms of plasma states (e.g. kinetic profiles, equilibrium and MHD states) and actuation resource signals (e.g. injected heating power and deposition location in terms of normalized radius), rather than device-specific diagnostics signals (e.g. ECE data, line-integrated density) and device-specific actuator signals (e.g. power supply voltages, EC mirror angles). In [22], the architecture and the design choices for these software layers are further motivated.

2.2.2. Plasma and actuator state reconstruction. As introduced in section 1, a single unified representation of the physical state of the plasma and actuators is chosen in order for a single centralized agent to make decisions about all control tasks, as well as providing the plasma state to all controller algorithms for the purpose of feedback control.

In this paper, the state of a system is defined as the (smallest) set of (internal) physical quantities that describe the dynamical evolution in time of the system sufficiently for the purpose of control. The choice of what quantities are represented in the state depend on the level of detail required for the application. For typical control tasks on tokamaks, the state of the plasma and tokamak is roughly represented by the plasma equilibrium, the currents in the conducting structures, all kinetic and magnetic profiles (including fast ions or runaway electrons), the presence, type and amplitude of MHD modes, the availability of actuator systems, and angles of movable mirrors and waveguide switches. These are chosen because they are the typical
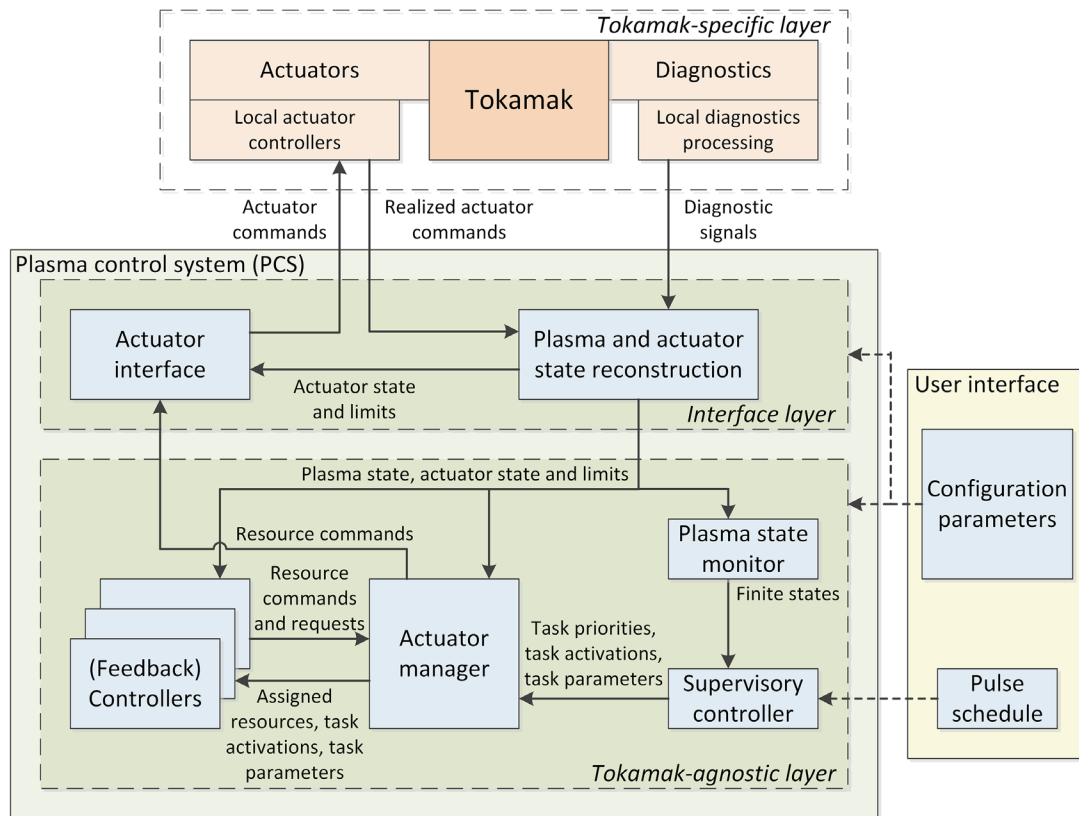
Figure 2. Block diagram of the generic architecture of a plasma control system. This architecture facilitates event detection, integrated control, exception handling and actuator sharing. All signals in the *tokamak-agnostic layer* are defined in terms of generic quantities, while algorithms in the *interface layer* form the boundary between the *tokamak-agnostic* and *tokamak-specific* signal quantities.

quantities that a control system may attempt to manipulate or may need to react to, and are generic for all tokamaks. The representation of this state in the control system software is a parametrization of this physical state.

Note that it may or may not be possible to reconstruct the state of the plasma based on diagnostic measurements, depending on the quality and quantity of diagnostic signals, as well as the modeling understanding of the underlying physics that determine the state evolution. This is discussed in more detail in [27, 28].

All algorithms that estimate the continuous real-valued state of the plasma and actuators based on the specific diagnostic and actuator signals are grouped in the plasma and actuator state reconstruction. The state reconstruction algorithms are typically equilibrium reconstruction (e.g. EFIT [29] and RT-LIUQE [30]), MHD analysis (e.g. [31, 32]), profile estimators (e.g. the RAPTOR-observer [28, 33]), ray-tracing codes (e.g. RT-TORBEAM [34]), and more. The resulting state contains the estimates of kinetic profiles, equilibrium, MHD states, absorbed auxiliary heating power, actuator state, among others. This is the common state fed to all other algorithms.

Depending on the requirements, these state estimates should probably be available at different sample rates. For example for the purpose of control and plasma monitoring, the vertical position must be computed at a higher rate than plasma equilibria. The proposed framework is not restricted to a single sample rate.

2.2.3. Plasma state monitor. Since the decisions about control tasks are taken in the supervisory controller based on discrete assessments of the condition of the plasma, a translation of the continuous representation of the state (e.g. values of physical quantities) to a discrete representation is made. In the plasma state monitor, this discrete state representation is formed using finite state machines (see [35]) representing the state of the plasma and the actuators at every time step during a plasma discharge. The finite state machines represent all the states of the plasma that the PCS should respond to or control. In section 3, the plasma state monitor is discussed in more detail.

2.2.4. Supervisory controller. The supervisory controller is responsible for all real-time decision-making regarding the coordination of all control tasks, as is also mentioned in [5, 7, 21].

The proposed supervisory controller has explicitly defined inputs and output signal interfaces to other algorithms in the PCS. It produces the task priorities, the task activation as well as quantities specific to each task (e.g. references, parameter settings), based on the discrete state of the plasma and the pulse schedule. These task priorities form the basis on which actuators are allocated in real time to the control tasks and their executing control algorithms. In section 4, the proposed supervisory controller is discussed more in detail.

2.2.5. Actuator manager, actuator interface and controller algorithms. An actuator management algorithm (e.g. [22, 25, 26]) assigns the available actuation resources (in terms of heating

power, current drive, fuelling) to the control tasks, based on the task priority, actuator availability and requests from controller algorithms. A set of (feedback) controllers is responsible for the execution of the tasks, using the assigned actuation resources. The focus of this work is on the plasma state estimation and monitoring, and supervisory control. The design of a compatible actuator manager and the interface to controller algorithms and actuator hardware can be found in [22].

## 3. Design of a generic plasma state monitor

In this section, the architecture and design of a plasma state monitor, containing finite-state models, are presented. These models are *tokamak-agnostic*, meaning that they can represent the state of any tokamak plasma. Plasma states and events (such as those in [6, 23]) are formalized using the concept of finite state machines (FSM) [35].

The list of FSMs needed to represent a tokamak plasma for the purpose of real-time control and event handling should be dictated by requirements. The required FSMs are determined by the quantities that need to be controlled or the events that should be monitored since the PCS may need to respond to them (e.g. limit violations, deviations from the target scenario, hardware faults). This list poses requirements on which plasma quantities should be estimated in real time and their reconstruction accuracy. The required quantities will be the confinement mode, detachment state, proximity of quantities (e.g. current, q_{95}, internal inductance, shape parameters, pressure, density, rotation, NTM amplitude and frequency, locked mode amplitude and phase) to their respective physics and operational limits, deviations of the controlled quantities from their references/targets, among others.

In general, each plasma state or set of physically related states that the PCS should respond to is modeled by a single FSM. Discrete quantities such as confinement mode (ohmic, L-mode, ELMy H-mode, quiescent H-mode) and number of X-points may be directly represented by FSMs. In line with modular modeling practice [35, 36], each FSM should represent one quantity (or component) and be as small as possible (atomic), avoiding unnecessary large FSMs which could be divided in sub-FSMs. This prevents redundancy among the FSMs. The combination (or synchronous composition, see [35]) of all individual FSMs represents the full system.

As an example, the finite state machine for the rotational frequency of a 2/1 mode, which may be indicative of the mode locking, is depicted in figure 3 and is further detailed in section 3.2. Its transition conditions are listed in table 1. The meaning of the finite states depends on the requirements for (supervisory) control. For example, the states slow can represent a mode that is predicted to lock within a set amount of time, while the state fast can represent a mode that does not pose immediate threat of locking. Although the threshold values of the transition conditions are specific to the application on a tokamak, they represent physics quantities rather than device-specific measured quantities.

A number of FSMs representing the plasma may be considered for the purpose of integrated plasma control. In
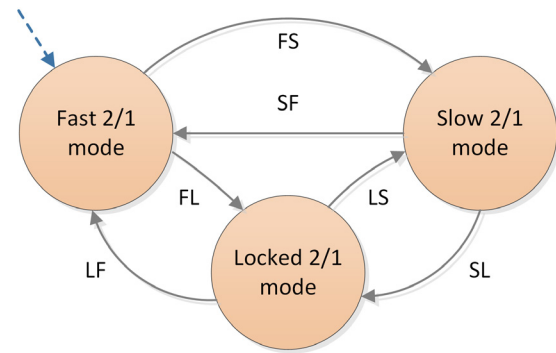
Figure 3. The finite state machine representation of the 2/1 mode frequency. The transitions are triggered on conditional tests, specified in table 1. These state transitions are triggered on the estimated 2/1 mode amplitude and frequency, as well as the estimated $n = 1$ locked mode amplitude. The dashed arrow indicates the initially active state. An non-exhaustive list of required finite state machines, their states and input signals are given in the respective columns of table 2.

Table 1. The conditions for the transitions of the finite state machine in figure 3. These transitions use the plasma state signals $f_{2/1\text{mode}}$, $A_{2/1\text{mode}}$ and $A_{n=1\text{LM}}$, see table 2. Note that small numbers ϵ_A and ϵ_f are added to the signal thresholds on reciprocal transitions, e.g. FS and SF. The resulting hysteresis prevents fast switching between states due to small signal variations. Furthermore, the threshold values are chosen such that $f_{2/1\text{mode}}^{\text{thresh,FS}} > f_{2/1\text{mode}}^{\text{thresh,SL}} + \epsilon_f$. Similar conditions are formulated for all FSMs in table 2.

Transition	Conditional test
FS	$f_{2/1\text{mode}} < f_{2/1\text{mode}}^{\text{thresh,FS}}$ AND $A_{2/1\text{mode}} > A_{2/1\text{mode}}^{\text{thresh,FS}} + \epsilon_A$
SF	$f_{2/1\text{mode}} > f_{2/1\text{mode}}^{\text{thresh,FS}} + \epsilon_f$ OR $A_{2/1\text{mode}} < A_{2/1\text{mode}}^{\text{thresh,FS}}$
FL, SL	$f_{2/1\text{mode}} < f_{2/1\text{mode}}^{\text{thresh,SL}}$ OR $A_{n=1\text{LM}} > A_{n=1\text{LM}}^{\text{thresh}} + \epsilon_A$
LF	$f_{2/1\text{mode}} > f_{2/1\text{mode}}^{\text{thresh,FS}} + \epsilon_f$ AND $A_{n=1\text{LM}} < A_{n=1\text{LM}}^{\text{thresh}}$
LS	$f_{2/1\text{mode}} > f_{2/1\text{mode}}^{\text{thresh,SL}} + \epsilon_f$ AND $f_{2/1\text{mode}} < f_{2/1\text{mode}}^{\text{thresh,FS}} + \epsilon_f$ AND $A_{n=1\text{LM}} < A_{n=1\text{LM}}^{\text{thresh}}$

table 2, an overview of finite state machines F, their states S and their input signals u is given. This table provides an initial set of plasma states which should be used for decision-making in integrated plasma control. Yet, the table is in no way complete or exhaustive. FSMs for additional plasma quantities may be added, for example for (low) plasma rotation, (excessive) tile temperature at strike point locations, MHD mode classification ((neoclassical) tearing mode, resistive wall mode, external kink, etc) and normalized beta limit proximity and violation. In figure 4, the input signals u and output signals Y of the plasma state monitor, which contains the finite state machines, are depicted. These FSMs are briefly introduced next.

Table 2. Non-exhaustive list of finite state machines F that should be considered for representing a tokamak plasma for the purpose of control and event handling, with signals u used for the state transitions. Note that duplicates for different m/n numbers of MHD modes are listed.

Finite state machine F name	States S	Used input signals from u
Plasma current segment	•No plasma, •Ramp-up, •Flat-top, •Ramp-down	Measured plasma current I_P and time derivative $\frac{d}{dt}I_P$
Mode amplitude $m/n = \{2/1, 3/1, 3/2\}$	•Nomode, •Small, •Large	Mode amplitude $A_{(m/n)\mathrm{NTM}}$
Mode frequency $m/n = \{2/1, 3/1, 3/2\}$	•Locked, •Slow, •Fast	Mode frequency $f_{(m/n)\mathrm{NTM}}$, mode amplitude $A_{(m/n)\mathrm{NTM}}$ and locked mode amplitude $A_{(n)\mathrm{LM}}$
Mode acceleration $m/n = \{2/1, 3/1, 3/2\}$	•Locking, •Stationary, •Unlocking	NTM acceleration $\frac{d}{dt}f_{(m/n)\mathrm{NTM}}$ and locked mode amplitude $A_{(n)\mathrm{LM}}$
Locked mode amplitude $n = \{1, 2, 3\}$	•NoLM, •LM	Locked mode amplitude $A_{(n)\mathrm{LM}}$
Position error (slow/static): vertical and radial	•Negative displ., •No displ., •Positive displ.	Vertical and radial position error $Z_{A,e}$, $R_{A,e}$
Position oscillations: vertical and radial	•No oscillations, •Oscillatory	Vertical and radial position oscillation amplitude $A_{Z,\mathrm{osc}}$, $A_{R,\mathrm{osc}}$
Vertical displacement event	•No VDE, •VDE	Vertical displacement indicator $Z_{A,e}\frac{d}{dt}Z_{A,e}$
Greenwald density limit ratio	•Below limit, •Close to limit, •Above limit	Greenwald fraction f_{GW}
Electron density control error	•Below control specification, •Close to zero, •Above control specification	Electron density error $n_{e,\mathrm{ref}} - n_e$
Confinement mode	•Ohmic, •L-mode, •I-mode, •ELMy H-mode •Quiescent H-mode	Auxiliary heating power, H-mode detection, ELM frequency, electron density profile $n_e(\rho)$, electron temperature profile $T_e(\rho)$
Magnetic configuration mode	•Limited, •Lower/upper single null, •Double null	Reconstructed equilibrium
Detachment mode	•No detachment, •Marginally detached, •Detached	Detachment detection, detachment front location
Sawtooth frequency	•No ST, •Slow ST, •Fast ST	Sawtooth frequency f_{ST}
Current density profile class	•Inductive, •Hybrid, •Reverse-shear	Normalized magnetic shear $s = \frac{\rho}{q}\frac{\partial q}{\partial \rho}$
Internal transport barrier location	•No ITB, •ITB collocated with q_{min} location, •ITB non-collocated with q_{min} location	Electron temperature profile gradient $\frac{\partial T_e}{\partial \rho}$, current density profile $q(\rho)$
Thermal radiation	•Close to expected radiation, •Excessive radiation, •Core radiative collapse	Radiated power fraction, electron temperature gradient, tungsten density profile
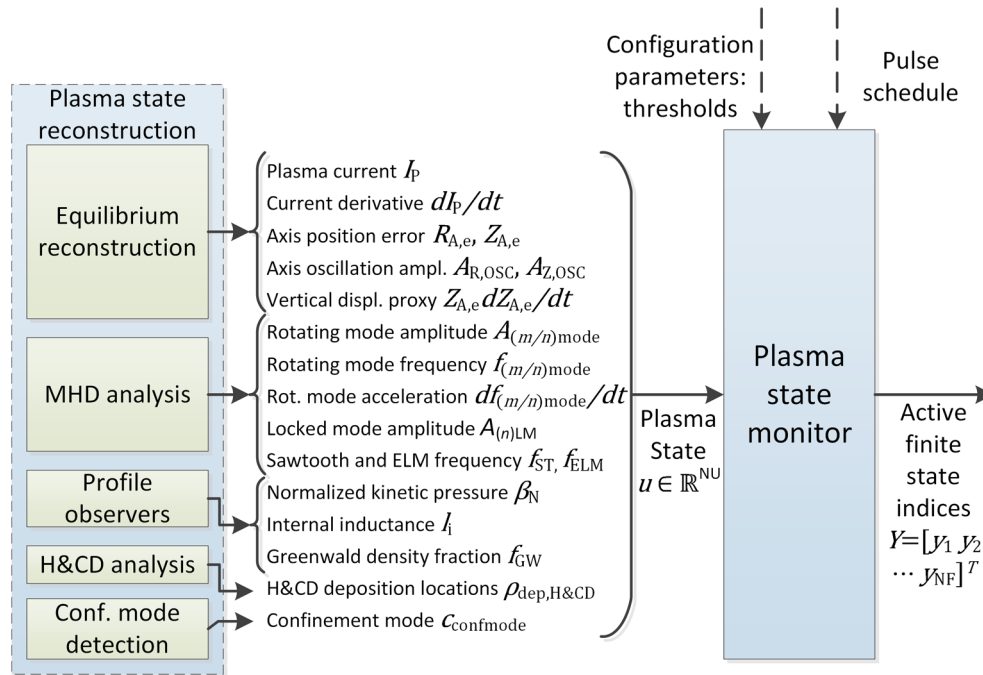
Figure 4. Block diagram of the generic plasma state monitor. Only the signals that can be estimated in real time on TCV at present are listed, which is a subset of those listed in table 2. See section 5 for the set of signals and FSMs that are implemented on TCV.

3.1. Plasma current segment

A simple set of discharge segment labels are considered in this work, namely the ramp-up, flat-top, ramp-down, as well as the absence of plasma. The observation of these discharge phases may be important for enabling control tasks. For instance, fusion power production may only be initiated in a reactor once the plasma is observed to be in flat-top and the current is constant.

In this paper, it is chosen to derive the discharge segment from the measured plasma current and its time derivative. Thresholds tests on the sign and absolute value of the time derivative of the plasma current cue the transitions between the ramp-up, flat-top or ramp-down states. The plasma is considered absent when the absolute value of the measured plasma current is below a given threshold.

3.2. Rotating modes

A number of states of rotating modes are considered, being the magnetic perturbation amplitude, frequency and rotational acceleration of the magnetic perturbation. In this paper, the $m/n = 2/1$, $3/1$ and $3/2$ modes are considered, with separate finite state machines for each mode and each of the aforementioned quantities, see table 2. The mode presence, i.e. an amplitude labeled as 'small' or 'large' may be interpreted as either the 'MHD' or 'NTM' events from [6], can lead to a disruption. The mode acceleration, i.e. the time derivative of its frequency, is considered since it may be indicative (in combination with the frequency) of mode locking and unlocking. In figure 3, the finite state machine for the rotational frequency of a $2/1$ mode is shown. Note that the 'locked' state in the mode frequency FSMs corresponds to the 'ML' label in [6]. The FSMs for the modes require the reconstructed continuous-valued magnetic

perturbation amplitude, frequency and acceleration of each of the m/n modes as input signals.

In order to avoid redundancy among the FSMs, only the mode amplitude FSM indicates the mode presence. Recall that the combination (or synchronous product [35]) of the finite state machines characterizes the full system, thereby eliminating the need for redundancy among atomic FSMs.

3.3. Locked modes

In this work, the presence of toroidally resolved locked modes is considered, specifically the $n = \{1, 2, 3\}$ modes. A mode lock is often a precursor to a disruption [6, 37], and its presence may require an emergency rampdown in a reactor. Thresholds on the magnetic perturbation amplitude determine the transitions between the states listed in table 2. The presence of any locked mode corresponds to the 'ML' label in [6].

3.4. Plasma position displacements and oscillations

Concerning the plasma position, different states and events may be considered, such as position displacements with respect to their references, vertical displacement events in disruptions, and plasma position oscillations. The latter may occur in TCV due to challenges in vertical stabilization control. In a reactor, problems with plasma position control may require a modification of the internal inductance or elongation to avoid a disruption. A vertical displacement event corresponds to the 'VDE' label, while the vertical position displacements and oscillations are related to the 'VSC' (vertical stability control problem) label in [6].

Threshold checks on the reconstructed position errors w.r.t. the references, $Z_{err} = Z_{reconstr} - Z_{ref}$ and $R_{err} = R_{reconstr} - R_{ref}$,

Table 3. Abstract examples of supervisory control decision logic. The task activation conditions in the second column are conditional tests on the discharge time t and the active finite states Y, which were introduced in section 3 and listed in table 2. If a condition holds, the priority is assigned to the corresponding task. Different exclusive conditions may result in different priorities for a task, as seen here for Task 1. Moreover, in the case that a controller can execute one out of several tasks, exclusive conditions among these activation conditions ensure that only one task becomes active at a time. This is shown here for Task 2 and 3.

| Task	Activation condition	Priority
Task 1	$y_1 ==$ state2 OR $y_2 ==$ state1	$p_{Task1} = c_1$
	$y_1 ==$ state1 AND $y_2 \neq$ state1 AND $t_{1,1} < t < t_{1,2}$	$p_{Task1} = c_2$
Task 2	$y_2 ==$ state2 AND $t_{2,1} < t < t_{2,2}$	$p_{Task2} = c_3$
Task 3	$y_2 \neq$ state2 AND $t_{3,1} < t < t_{3,2}$	$p_{Task3} = c_4$

determine the transitions between the position displacement states listed in table 2. Threshold checks on the amplitude of observed oscillations determine the transition between the oscillation states. Moreover, threshold checks on the multiplication $Z_{err} \frac{d}{dt} Z_{err}$ indicates unstable vertical displacement events seen in disruption, as is also done in [23].

3.5. Plasma kinetic quantity deviations and limits

The proximity to and excursions of known operational and physics limits related to kinetic quantities may be used to trigger recovery actions. These limits may include the Greenwald density limit, limits on q_{95}, and normalized kinetic pressure limits. Also, deviations of controlled kinetic quantities w.r.t. their references indicate control problems or indicate that a controller is outside its designed operational space, and may be used to trigger recovery actions [3]. Threshold checks on reconstructed (quantities derived from) kinetic profiles may determine the transitions between operational states.

4. Design of a generic supervisory controller

In this section, the supervisory control algorithm, as mentioned in section 2, is presented. The function of the supervisory controller is to coordinate the execution of various control objectives in the plasma. The supervisory controller prioritizes the various control tasks and activates the low-level controllers responsible for executing the tasks. The priorities are taken into account by the actuator management system; they are an input to the merit function that determines optimal allocations of actuation resources to tasks. This output of the supervisory controller depends on a combination of a preprogrammed response in time, and on states observed in the plasma. For example, the transition to high confinement mode will be programmed by a feedforward waveform for auxiliary heating power; while the observation of an unexpected backtransition should prompt the PCS to allocate additional resources for central heating power or adapt the pressure reference.
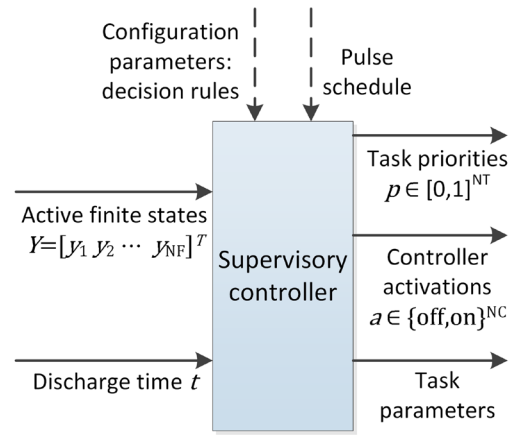
Figure 5. Block diagram of the input and output signals of the supervisory controller. The indices of active finite states originate from the plasma state monitor, discussed in section 3.

4.1. Task-based coordination of controller algorithms

Each control task can be executed by one or more low-level controllers. For example, the task for plasma kinetic pressure control can be executed by either pure feedforward control, or the combination of a feedforward and feedback controller. A single low-level controller may perform multiple control tasks, although not necessarily simultaneously. Rules in the supervisory controller ensure that only viable sets of control tasks are activated at a time. For instance, we consider the stabilization and preemption of one NTM to be mutually exclusive: only one of these tasks can be executed by the NTM controller simultaneously for a given rational q surface. Mutual exclusive conditions for the activation of these tasks impose that only one is enabled at a time for a given rational q surface.

For each control task, the supervisory controller executes a set of conditional logic expressions, which depend on the active plasma states and the discharge time. The outcome of each conditional expression enables the task and assigns a priority to it, or disables the task. As an example, table 3 shows abstract examples of decision rules in the proposed supervisory controller.
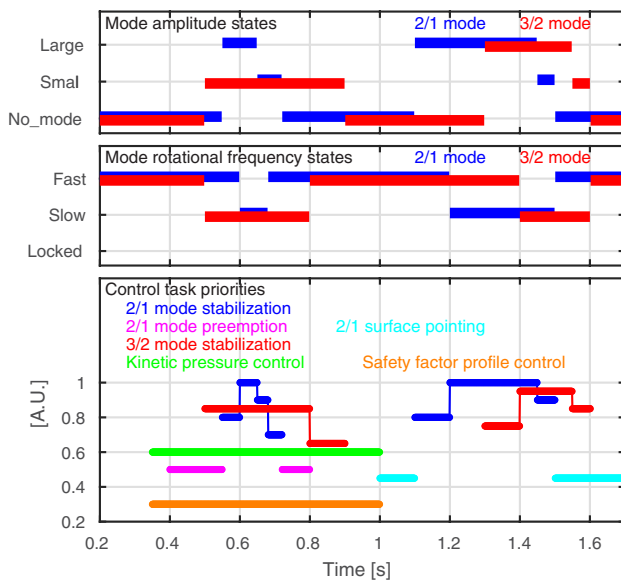
4.2. Supervisory control outputs

The supervisory controller produces multiple outputs, as shown in figure 5. First, it produces the priority for each control task, which is in the range $[0, 1]$. In the proposed framework, these priorities represent the (relative) importance of tasks, and the values are used by the actuator manager to determine an optimal allocation of actuation resources to tasks [22]. When a given task priority is larger than zero, the corresponding task is activated and the task execution by the appropriate feedback control algorithm is enabled.

Second, the supervisory controller activates the controller algorithms if the discharge time is within a pre-set activation time window. This allows a control algorithm to run in the background even if none of its tasks are active and it is not in

Table 4. Decision logic for the simulation in figure 6. Note that all conditions belonging to the same task are mutually exclusive.

Control task	Activation condition	Priority
2/1 mode stabilization	2/1ampl $=$ =large & 2/1freq $=$ =slow	$p_1 = 1$
	2/1ampl $=$ =small & 2/1freq $=$ =slow	$p_1 = 0.9$
	2/1ampl $=$ =large & 2/1freq $=$ =fast	$p_1 = 0.8$
	2/1ampl $=$ =small & 2/1freq $=$ =fast	$p_1 = 0.7$
2/1 mode preemption	$0.4 < t < 0.8$	$p_2 = 0.5$
2/1 surface pointing	$1 < t < 3$	$p_3 = 0.45$
3/2 mode stabilization	3/2ampl $=$ =large & 3/2freq $=$ =slow	$p_4 = 0.95$
	3/2ampl $=$ =small & 3/2freq $=$ =slow	$p_4 = 0.85$
	3/2ampl $=$ =large & 3/2freq $=$ =fast	$p_4 = 0.75$
	3/2ampl $=$ =small & 3/2freq $=$ =fast	$p_4 = 0.65$
Kinetic pressure control	$0.35 < t < 1$	$p_5 = 0.6$
Safety profile control	$0.35 < t < 1$	$p_6 = 0.3$

Figure 6. Simulation of the supervisory controller with artificial plasma state signals, illustrative of the supervisory control possibilities. In the top two panels, synthetic rotating mode state signals are given. In the other panel, the priorities for the various control tasks are shown, computed using the decision rules in table 4. Note that depending on the amplitude and frequency of the 2/1 and 3/2 mode, the supervisory controller prioritizes the stabilization of either mode over the other.

command of any actuation. Last, the supervisory controller sends out a number of signals and parameters that are specific to each control task and are related to its execution. For instance, for a plasma pressure control task, this may include the reference signal for the controlled quantity, while for NTM control tasks this may include the $q = m/n$ target and maximum heating power. Some of these task-specific parameters are discussed in [22].

The supervisory input-output relation can be formally expressed as a task priority function $p = g_P(t, Y)$ and a controller activation function $a = g_A(t)$. Here, $p \in P$ is the vector of control task priorities where, $P = [0, 1]^{N_T}$, N_T is the number of control tasks, and $a \in A$ is the vector of controller activations, where $A = \{\text{off}, \text{on}\}^{N_C}$ and N_C is the number of

(feedback) controllers. Furthermore, t is the discharge time and $Y = \begin{bmatrix} y_1 & y_2 & \dots & y_{N_F} \end{bmatrix}^T$ is the column of all active finite states.

4.3. Supervisory controller simulation

To illustrate the input-output behaviour, capabilities and possibilities of the chosen architecture of the supervisory controller, we present a simulation with synthetic plasma state signals and a predefined set of decision rules.

The input data consists of artificial signals of appearing and disappearing rotating modes, both $m/n = 2/1$ and 3/2 with varying amplitude and rotational frequency. The supervisory control rules are listed in table 4 and represent the preprogrammed response for kinetic pressure, safety factor profile control, mode preemption and idle EC beam pointing as well as the real-time response to observed rotating modes. In this simulation example, the real-world problem of prioritizing the suppression and preemption of modes versus kinetic pressure and q-profile control is displayed. Specifically the case where the rotational frequency and amplitude of the modes determines their relative priority: although a 2/1 mode is more important to suppress than a 3/2 mode, a sufficiently small and fast 2/1 mode is less important to suppress given limited resources than a slow, large 3/2 mode. Although unfavorable safety factor profiles cause rotating modes to appear and therefore control of the safety factor profile is important, the stabilization of modes (when they appear) is valued more important in this case. Therefore the mode stabilization tasks receive a higher priority when a mode is present.

The decision rules internal to the supervisory controller consider the amplitude and frequency of two distinct NTMs, as well as time-based triggers, to prioritize $N_C = 6$ control tasks related to NTM, pressure, and safety factor control. These rules are listed in table 4.

In figure 6, we show the artificial plasma states and resulting control task priorities from the simulation. In this simulation example, the suppression of two appearing and disappearing rotating modes is prioritized over the scheduled tasks of pressure and safety factor profile control.
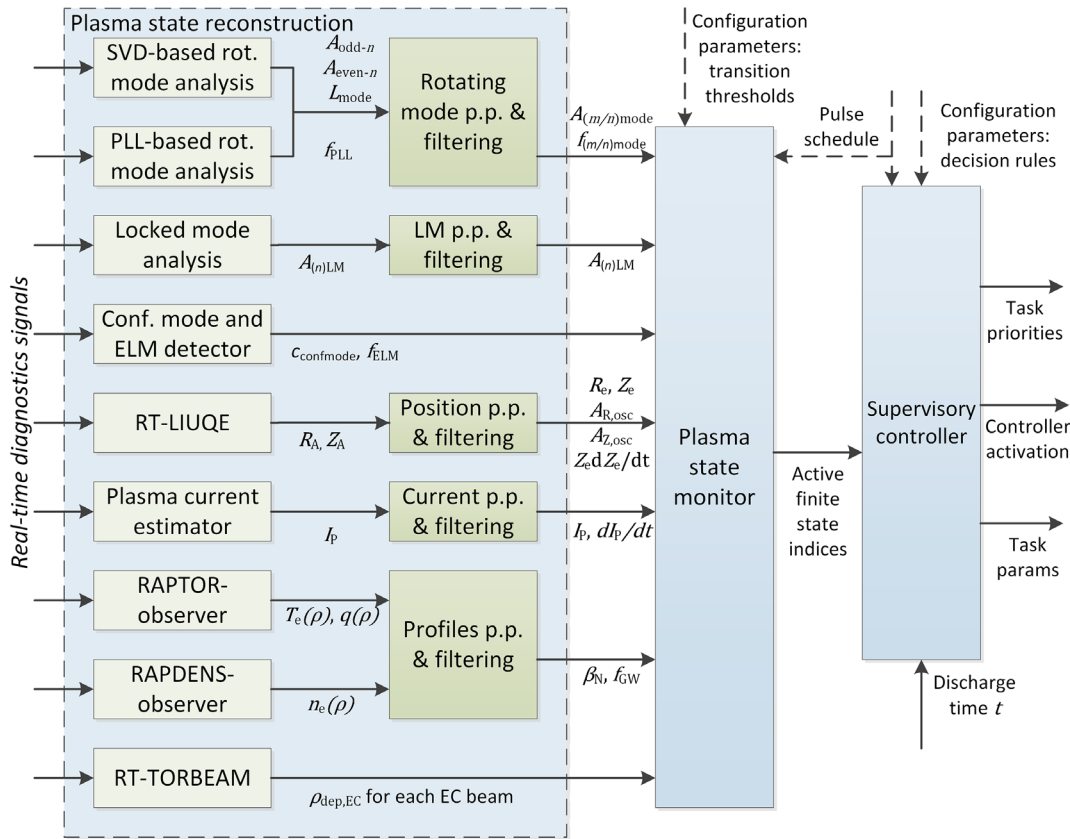
Figure 7. Block diagram of software components related to the plasma state monitor and supervisory control on the TCV digital control system. Various reconstruction algorithm outputs are filtered and processed in order to yield suitable inputs to the finite-state machines. The outputs of the supervisory controller are sent to the actuation resource management and feedback controllers [22], as shown in figure 2.

5. Implementation and results on TCV

The plasma state monitor and the supervisory controller have been implemented on the TCV digital control system [24] and used in experiments. In this section, the implementation and interfacing to other real-time algorithms is discussed, as well as results of experiments on integrated control on TCV.

5.1. Implementation of the plasma state monitor and supervisory controller on the TCV control system

The real-valued plasma state estimates, used in the plasma state monitor described in section 3, are derived from various existing real-time reconstruction algorithms on the TCV control system. These algorithms and their outputs relevant to this work are:

- the rotating mode analysis based on singular value decomposition [31, 38], producing the likelihoods that the two most dominant modes have mode number $m/n = 2/1$, $3/2$ or $3/1$,
- a newly implemented phase-locked loop, providing the rotating mode frequency,
- the standard odd-n and even-n rotating mode amplitudes,
- the real-time mode lock analysis, producing the locked mode amplitude for $n = 1$, $n = 2$ and $n = 3$,
- the real-time event detector [39], providing the confinement mode and ELM frequency,
- the real-time equilibrium reconstruction code LIUQE [30], producing the plasma current centroid position,

- the measurement of the plasma current, based on spatial integration of magnetic probe measurements surrounding the plasma (see [30]),
- the real-time EC raytracing code TORBEAM [34], producing the power deposition locations of injected EC beams,
- the RAPTOR-observer [27, 28], estimating the electron temperature profile, safety factor profile and derived quantities such as kinetic normalized pressure and internal inductance,
- the RAPDENS-observer [40], estimating the electron density profile.

In figure 7, these reconstruction algorithms, as well as newly implemented postprocessing blocks, the plasma state monitor and the supervisory controller are shown, with the relevant input and output signals. Note that in this setup indeed the reconstruction algorithms are *specific* to TCV, while the finite state machines in the plasma state monitor contain no details of TCV hardware and are truly *tokamak-agnostic*. The FSMs in the plasma state monitor are implemented using the MATLAB Stateflow toolbox [41]. The relevant computational node hosting the plasma state monitor, supervisory controller and actuator management runs at a cycle rate of 1 kHz.

On the output side, the supervisory controller is interfaced with the actuator manager [22] and ultimately to individual controllers, being the interchangeable profile controllers [11–14] and the NTM controller [42]. Next, the postprocessing

and interface of the reconstruction algorithms to the plasma state monitor, corresponding to the finite state machines that are used in the results, are discussed.

5.1.1. Rotating mode analysis. A rotating MHD analysis algorithm based on a singular value decomposition (SVD), see [31, 38], is combined with a phase-locked loop to provide information about rotating modes, both using a set of in-vessel poloidal field coils.

The amplitude of odd-*n* and even-*n* rotating modes is computed in the standard way by taking the root mean square value of the difference and sum of two toroidally opposed in-vessel poloidal field coils, respectively. These amplitudes are low-pass filtered with a cut-off frequency of 100 Hz. Then, a poloidal mode number is assigned to these odd-*n* and even-*n* amplitudes yielding m/n numbers, by selecting the maximum over the likelihoods that the two most dominant rotating modes have mode numbers $m/n = 2/1$, $3/1$ or $3/2$. These likelihoods are provided by the SVD-based algorithm.

Note that for the mode amplitude, we consider the amplitude of the magnetic perturbation measured at the coil locations. At present, there is no reconstruction of the magnetic perturbation at the location of the mode.

In this work, the rotating mode frequency is estimated by a phase-locked loop (PLL) [43]. It synchronizes an oscillator to the most dominant periodic component in an appropriate linear combination of toroidally separated poloidal field coil signals. In TCV, rotating modes appear in a wide frequency range up to 10 kHz. The PLL is designed to estimate the mode frequency in the range between 500 Hz and 10 kHz. In the present implementation of the SVD-based analysis, a zero-crossing detector provides an estimate of the rotating mode frequency. Since the PLL does not rely on zero-crossing detection, it is more robust to noise. The resulting mode frequency is low-pass filtered at a cut-off frequency of 100 Hz to smooth over multiple periods of the mode. Similarly as above, a poloidal mode number m is assigned to the estimated frequency by selecting the maximum of the likelihoods. The mode acceleration is computed by numerical derivation of the mode frequency. Additional low-pass filtering at a cut-off frequency of 50 Hz is applied to reject the noise amplification caused by the numerical differentiation. The reconstructed mode amplitude, frequency and acceleration then feed to the FSMs for the rotating modes, as discussed in section 3.2.

5.1.2. Locked mode analysis. A recently implemented real-time mode lock indicator for toroidally resolved mode numbers on TCV's control system is used [24]. The amplitude of locked modes is computed in real time, and is fed to the FSMs for locked modes, see section 3.3. Note that the computation yields the magnetic perturbation amplitude at the coil locations. There is no reconstruction of the physical mode magnitude.

5.1.3. Equilibrium reconstruction and measured plasma current. The magnetic axis position is provided by real-time LIUQE [30] and is subtracted from the reference position and
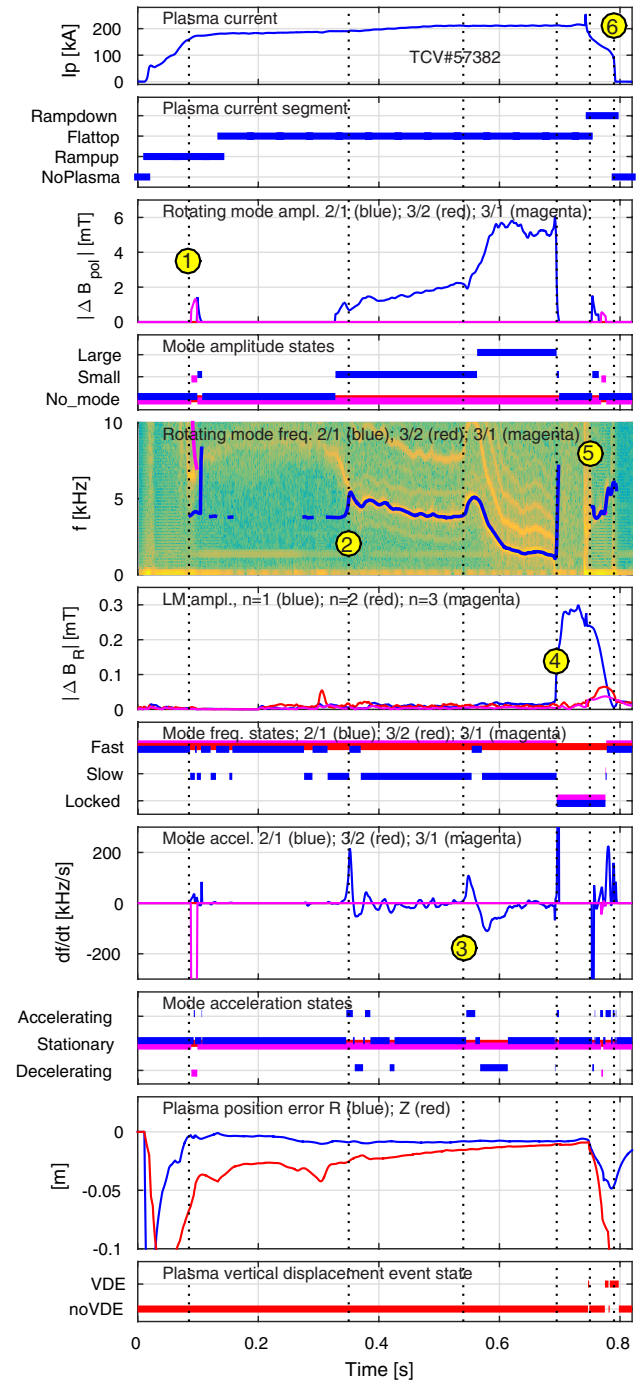
Figure 8. Result of event detection for TCV#57382. At time point ①, a 2/1 mode briefly appears. Around $t = 0.33$ s it appears again and is detected at time point ②. The mode briefly accelerates and decelerates after time point ③, before locking at $t = 0.69$ s at time point ④. The rotating mode reappears at time point ⑤ after which a VDE develops and the plasma disrupts at time point ⑥.

low-pass filtered with a time constant of 100 ms. It is then fed to the position excursion FSMs, see section 3.4. The position oscillation amplitude in both the radial and vertical directions are computed as the square root of the signal power in the frequency range between 10 Hz and 100 Hz, with the latter obtained through band-pass filtering. These feed to the FSMs for position oscillations, introduced in section 3.4. Moreover,
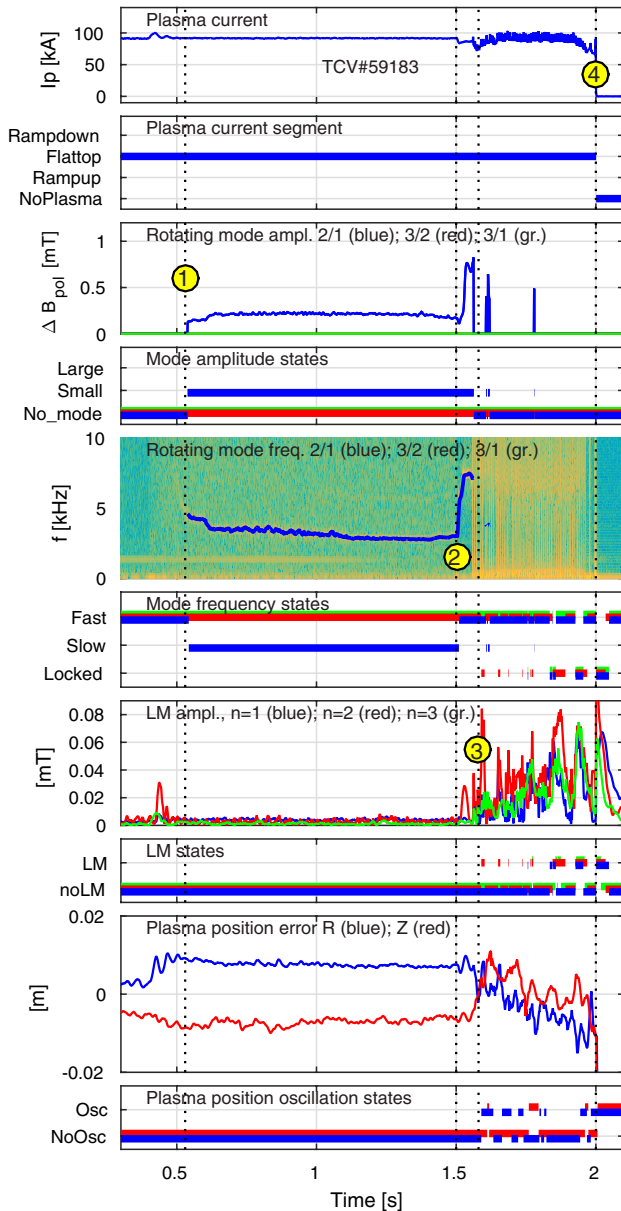
Figure 9. Result of event detection for TCV#59183. A 2/1 NTM appears around 0.55 s (time point ①). After a deliberate massive neon injection at $t = 1.5$ s (time point ②), the NTM accelerates to 6 kHz. At time point ③, the 2/1 NTM disappears while an $n = 2$ locked mode appears. Between 1.6 s and 2 s, the mode consecutively unlocks and locks [44], while the plasma position is oscillating. At time point ④ the plasma disrupts.

the vertical velocity is numerically derived from the vertical position. Similar to [23], the multiplication $Z_{err} dZ_{err}/dt$ is used to indicate unstable vertical movement events, as mentioned in section 3.4.

The plasma current and its time-derivative are used for the plasma current segment states in section 3.1. The plasma current is estimated by the trapeze approximation (see [30]). The time-derivative of the plasma current is numerically computed and is low-pass filtered with a second-order filter at a cut-off frequency of 20 Hz is applied to reject the noise amplification from numerical differentiation.

Table 5. Implemented decision logic for the experimental result in figure 10.

Task	Activation condition	Priority
2/1 NTM stabilization	2/1ampl $==$large OR 2/1ampl $==$small	$p_1 = 1$
2/1 NTM preemption	2/1ampl $==$noNTM AND $0.8 < t < 2$	$p_2 = 0.5$
Beta control	$0.25 < t < 2$	$p_3 = 0.6$
Safety profile control	$0.25 < t < 2$	$p_4 = 0.3$

5.2. Real-time detection of NTMs and locked modes on TCV

In this subsection, the capabilities of the state monitor to detect behaviour of several quantities related to MHD and plasma position control problems is demonstrated.

The results of the plasma state monitoring in TCV discharge #57382 are depicted in figure 8. In this discharge, a 2/1 NTM grows in amplitude, while accelerating and decelerating, before locking. Subsequently, the rotating mode reappears, before a VDE develops and the plasma disrupts. These consecutive events are flagged by the state monitor.

In figure 9, the results of plasma state monitoring results in TCV discharge #59183 are shown. In this discharge, a 2/1 NTM accelerates to twice its initial frequency. Subsequently it disappears while a $n = 2$ locked mode appears at 1.6 s. Although further analysis in [44] reveals that the mode spins up and locks consecutively between 1.6 s and 2 s, these events happen too fast for the real-time MHD analysis to pick up. These fast transients can be seen in the magnetics spectogram. In the last phase of the discharge, radial and vertical position oscillations develop, before disrupting.

Design tradeoffs exist between the detection delay and detection accuracy. In the present implementation, all input signals u are filtered in order to reduce fast switching between finite states due to signal noise. Similarly, a delay is introduced in filtering out the noise introduced by the numerical derivation of the NTM acceleration from its frequency. These delays introduced by the filters propagate to the state monitor, and will cause delayed (re)actions by the supervisory controller and the task-executing algorithms. For the present purposes of control-focused experiments on TCV, 10 ms of delay between the physical event and a supervisory control decision is acceptable.

5.3. Experimental results of integrated kinetic profile and NTM control on TCV

In this subsection, the capabilities of the supervisory controller in a TCV discharge to provide prioritized control task is demonstrated. A low-density limited discharge is considered for simultaneous control of the plasma pressure, safety factor profile and NTMs. The goal is to control the plasma pressure and safety factor profile, while suppressing NTMs. The supervisory controller assigns the priorities to the tasks. Then, the actuation resource management distributes command over

two EC power supplies and the poloidally steerable mirrors on three launchers among the controller algorithms.

In the present implementation on TCV, two controllers are interfaced to the supervisory controller and the actuator management, and a set of control tasks are designated to each controller (see also [22]). The control tasks related to rotating mode control, executed by the NTM controller [42] are

- NTM suppression/stabilization at a specified $q = m/n$ target,
- NTM preemption at a specified $q = m/n$ target,
- flux surface tracking, i.e. aiming an EC launcher at a specified q surface without power injection.

The control tasks related to kinetic profile control, executed by the (interchangeable) profile controllers [11–14] are

- Plasma pressure β control.
- Inverse safety factor profile $\iota = 1/q$ control.

For the present purposes, the supervisory control decisions are programmed by the user. The supervisory decision logic related to the tasks is shown in table 5. Note that the supervisory controller ensures that at most one of the preemption task and stabilization task of the 2/1 mode is active at a time: the activating conditions for these two tasks are mutually exclusive.

In figure 10, the experimental results of TCV discharge #57813 are presented. Although the mode is not stabilized within the time of the discharge. Still, the supervisory controller evolves the task priorities among the relevant control tasks as intended by the user. It first enables the preemption of a 2/1 mode with limited power on launcher #6. At 1.3 s, a 2/1 mode is observed and the supervisor enables the mode suppression task instead of the mode preemption task. This prompts the actuator manager to grant full power availability on launcher #6 to the NTM controller.

Due to technical limitations of the EC system, the gyrotrons of launcher #4 and #6 are on the same power supply such that equal power is delivered by these launchers. Yet, they are not depositing power on the same location. At 1.8 s, the actuator manager assigns a second EC launcher (#4) to the mode suppression task. This prompts the NTM controller to direct launcher #4 to the resonant surface, where launcher #4 was previously aimed at the magnetic axis. This experiment demonstrates the succesful execution of the implemented supervisory control rules, and showcases its capability to manage the execution of multiple control tasks by multiple controllers.

6. Discussion

In this work, we provide a supervisory control framework that deals with several aspects of advanced plasma control. The supervisory controller needs to be programmed by the user according to the intended experimental needs, similar to [19, 20]. However, extending the method to include a complete set of functionalities as required in a reactor is challenging. First, design of a supervisory controller for a highly complex system is difficult due to the large dimension of the

state space [45]. Second, manual design of a supervisory controller for large systems may lead to a supervisor that contains blocking situations, i.e. when the system ends up in a state from which it cannot be driven toward a desired state. A tailored solution from the control engineering community is the requirement-driven design procedure [36, 46], which provides an optimized supervisory controller that is guaranteed to be non-blocking, correct and controllable. This solution may be the topic of future work. Yet, the current proposed method remains generic, promising and flexible to be extended using a requirement-driven design procedure.

In this work, the fact that the signal thresholds that determine the state transitions must be specified by the user implies that the physical meaning of the finite states is arbitrary. In other research, the semantics of events was expressed in terms of a statistical proximity to disruptions [6, 23]. For future applications of supervisory control in tokamaks, the meaning of finite states must be dictated by requirements on the real-time control capability, which includes proximity to disruptions.

In the present implementation on TCV, the set of finite state machines does not cover the complete set of plasma, hardware and control system states that is relevant for real-time control as required in future tokamaks and reactors. As such, this research forms a proof of concept for supervised plasma control. For the purposes of experimental tokamaks, only coverage of states related to the intended experiments is needed, as done in this work and [19, 20]. In [6, 23], larger sets of events are considered, in an offline analysis focussed on disruption causes. Yet, our work can be extended with additional finite state representations for the plasma, and is valuable for the application of event detection to coordination of all relevant control tasks in future reactors.

Moreover, the selection of available diagnostics and reconstruction algorithms determines the physical meaning of the discrete states. In this work, several states are expressed in terms of plasma quantities, such as plasma current, position, density, temperature, safety factor and EC beam deposition. However, the real-time MHD analysis provides the amplitude of NTMs and locked modes as being the perturbation amplitude as observed at the magnetic field coils, rather than the island width or the magnetic field perturbation amplitude local to the mode. This implies that these quantities are not (yet) *tokamak-agnostic*. For the present experimental purposes on TCV, this does not limit the applicability. Yet, in the future a *tokamak-agnostic* state representation is valuable. This allows for developed and proven controller algorithms that take *tokamak-agnostic* state quantities as input to be ported to other tokamak PCS's.

7. Conclusions and outlook

7.1. Conclusions

In next-generation tokamaks and reactors, the plasma control system should deal with the magnetic control as well as a number of control tasks associated with the performance and stability of the plasma. These control tasks will rely on a
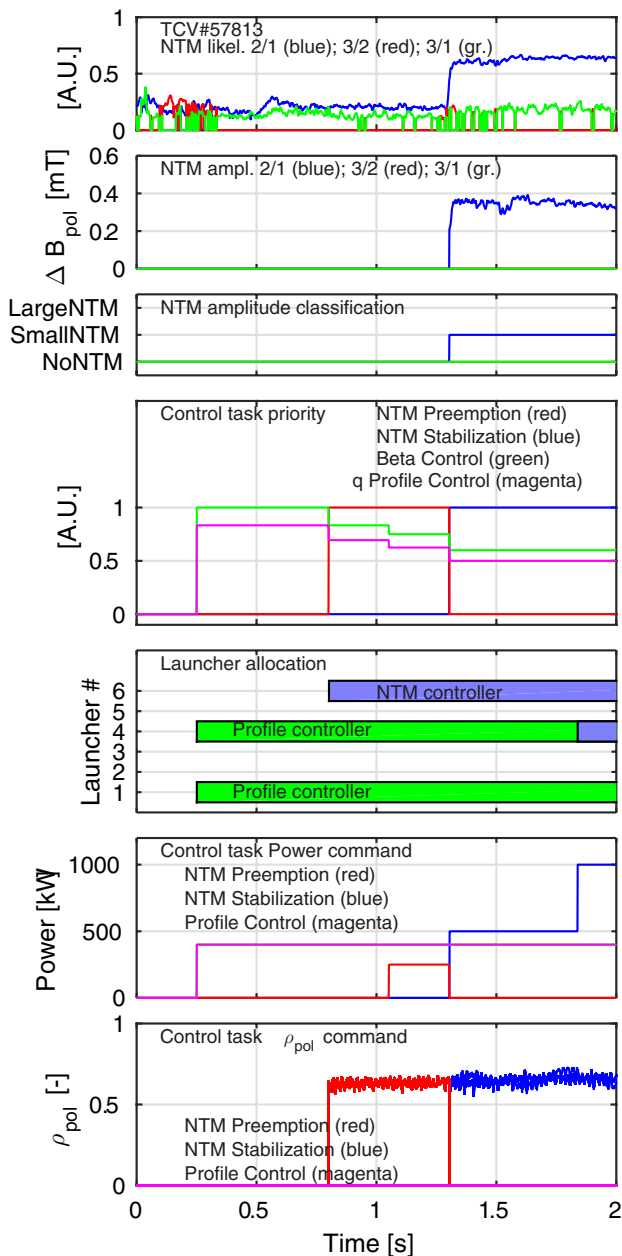
Figure 10. Experimental result of simultaneous pressure and NTM control on TCV #57813. The supervisory controller executes the decision logic shown in table 5. Furthermore, the actuator manager distributes the available EC launchers and power supplies to the NTM controller and the profile controller. The appearance of an NTM prompts the supervisory controller to assign top priority to the stabilization task.

limited shared set of actuators. Supervisory control is a necessary component in future tokamak reactor control systems.

At present, state-of-the-art supervisory control for advanced plasma control on tokamaks is done at ASDEX-Upgrade and DIII-D. In our work, feedback controllers are integrated in a generic architecture for advanced plasma control to allow for the real-time coordination of multiple objectives. This work extends the research at DIII-D by a strict separation between event detection and execution, which allows for a modular and extendable architecture.

This paper shows the design, implementation and experimental demonstration of a generic plasma state monitor and supervisory controller. This supervisory controller sets the relative priority of various control tasks, based on a set of finite states of the plasma. These priorities determine how the available actuators are allocated to the tasks, which is computed by an actuator allocation algorithm. The signal thresholds that determine the boundaries between the finite states must be specified by the user, and need to be adjusted according to the experimental needs of the user. In this way, the supervisory controller ensures that the most important task at any time has highest priority, since it coordinates all tasks. Abstracting the integration of controllers and actuators using tasks and their prioritization yields a less complex control system design in (future) machines, where the number of actuators, controller algorithms and tasks (control functionalities) is much larger than in present-day devices.

In the implementation on TCV, the monitored discrete states are the real-time estimated NTM state, discharge segment, locked mode amplitude and plasma position excursions, oscillations and vertical displacement events.

This work showed that the proposed framework can be used to carry out experiments on TCV with real-time centralized prioritization of control tasks, which can be programmed by the user. In this paper, we have shown real-time event detection and supervisory control for a set of plasma states. More specifically, we have demonstrated examples of rotating mode, mode lock and plasma displacement monitoring, as well as an experimental result of supervisory control of multiple control tasks based on NTM occurence. In the latter, the supervisory controller distributes priority among NTM control tasks and kinetic profile control tasks.

The plasma state monitor and supervisory controller are suited to perform physics experiments on TCV. Moreover, this work supports the integration of control capabilities for ITER and DEMO.

7.2. Outlook

In the future, depending on the real-time available diagnostic signals and plasma state reconstruction algorithms, the number of state machines and discrete states can be extended at will to monitor more complex situations. This is facilitated by the generic architecture of the proposed PCS. Our generic layout can be easily extended with additional discrete states that cover the entire range of physical and technical events. The latter includes both hardware as well as control system events and states.

The semantics of the discrete states may be refined using physics interpretation, statistical analysis [6, 23], as well as hardware and software design specifications [3]. Future applications of the finite-state modeling framework include development of generic algorithms to assess the plasma health and proximity to disruptions based on tokamak-independent representations of the plasma state, from physics interpretation or statistics. These can be tested on existing devices before being ported to larger machines.

In this work, some of the plasma states are not defined in terms of *tokamak-agnostic* quantities. In the future, a real-time reconstruction of the MHD mode width based on the equilibrium and coil geometry can provide a *tokamak-agnostic* interpretation of the MHD states.

The supervisory control framework should be tested on a larger set of discharges, such that its performance can be quantified in a statistical sense.

Note that the control tasks in this work are limited to control of performance and stability quantities. First, additional tasks may be added, e.g. related to control of other MHD and impurity removal [47]. Second, the domain of tasks may be extended to include plasma termination and emergency handling tasks [3, 7, 48].

Acknowledgments

ORCID iDs

T.C. Blanken https://orcid.org/0000-0003-3198-7170
F. Felici https://orcid.org/0000-0001-7585-376X
N.M.T. Vu https://orcid.org/0000-0003-2020-4812

References

[1] Pironti A. and Walker M. 2005 *Control Syst. IEEE* **25** 24–9
[2] Snipes J.A. *et al* 2014 Physics of the conceptual design of the ITER plasma control system *Fusion Engineering and Design* vol 89 (Amsterdam: Elsevier) pp 507–11
[3] Humphreys D. *et al* 2015 *Phys. Plasmas* **22** 021806
[4] Snipes J.A. *et al* 2017 *Nucl. Fusion* **57** 125001
[5] Treutterer W., Humphreys D., Raupp G., Schuster E., Snipes J., De Tommasi G., Walker M. and Winter A. 2014 Architectural concept for the ITER plasma control system *Fusion Engineering and Design* vol 89 (Amsterdam: Elsevier) pp 512–7
[6] De Vries P.C., Johnson M.F., Alper B., Buratti P., Hender T.C., Koslowski H.R. and Riccardo V. 2011 *Nucl. Fusion* **51** 053018
[7] Raupp G. *et al* 2017 *Fusion Eng. Des.* **123** 541–5
[8] Maraschek M. *et al* 2018 *Plasma Phys. Control. Fusion* **60** 014047
[9] Henderson M. *et al* 2015 *Phys. Plasmas* **22** 021808
[10] Snipes J.A. *et al* 2012 *Fusion Eng. Des.* **87** 1900–6
[11] Maljaars E. *et al* 2017 *Nucl. Fusion* **57** 126063
[12] Vu N.M.T., Nouailletas R., Lefèvre L. and Felici F. 2016 *Control Eng. Pract.* **54** 34–45
[13] Kim D., Kim S., Felici F., Maljaars E. and Sauter O. 2016 An active feedback plasma profile control approach applied to TCV plasmas and perspectives toward ITER *43rd European Physical Society Conf. on Plasma Physics (Leuven, Belgium, 4–8 July 2016)* (European Physical Society) (https://research.tue.nl/en/publications/an-active-feedback-plasma-profile-control-approach-applied-to-tcv)
[14] Mavkov B., Witrant E., Prieur C., Maljaars E., Felici F., Sauter O. and The TCV Team 2018 *Nucl. Fusion* **58** 056011
[15] Wehner W., Lauret M., Schuster E., Ferron J.R., Holcomb C., Luce T.C., Humphreys D.A., Walker M.L., Penaflor B.G. and Johnson R.D. 2016 Predictive control of the Tokamak q profile to facilitate reproducibility of high-qmin steady-state scenarios at DIII-D *IEEE Conf. on Control Applications (Buenos Aires, Argentina, 19–22 September 2016)* (IEEE) pp 629–34 (https://doi.org/10.1109/CCA.2016.7587900)
[16] Wehner W., Barton J. and Schuster E. 2017 Combined rotation profile and plasma stored energy control for the DIII-D tokamak via MPC *Proc. American Control Conf. (Seattle, WA, USA, 24–26 May 2017)* (IEEE) pp 4872–7 (https://doi.org/10.23919/ACC.2017.7963709)
[17] Lauret M., Wehner W., Schuster E., Holcomb C.T., Luce T.C., Ferron J.R., Walker M.L., Humphreys D.A., Penaflor B.G. and Johnson R.D. 2017 Current profile and energy control in DIII-D plasmas using discrete-time variable-structure control *25th Mediterranean Conf. on Control and Automation (Valletta, Malta, 3–6 July 2017)* (IEEE) pp 99–104 (https://doi.org/10.1109/MED.2017.7984102)
[18] Barton J.E., Besseghir K., Lister J. and Schuster E. 2015 *Plasma Phys. Control. Fusion* **57** 115003
[19] Rapson C.J. *et al* 2016 *Fusion Eng. Des.* **123** 603–6
[20] Eidietis N., Choi W., Hahn S., Humphreys D., Sammuli B. and Walker M. 2018 *Nucl. Fusion* **58** 056023
[21] Treutterer W. *et al* 2017 *Fusion Eng. Des.* **115** 33–8
[22] Vu N.M.T., Blanken T.C., Kong M., Maljaars E., Felici F., Galperti C., Sauter O. and The TCV team 2018 *Nucl. Fusion* submitted
[23] Gerhardt S.P. *et al* 2013 *Nucl. Fusion* **53** 043020
[24] Anand H., Coda S., Felici F., Galperti C. and Moret J.M. 2017 *Nucl. Fusion* **57** 126026
[25] Maljaars E. and Felici F. 2017 *Fusion Eng. Des.* **122** 94–112
[26] Rapson C.J., Reich M., Stober J., Treutterer W. and the ASDEX Upgrade team 2015 *Fusion Eng. Des.* **96–7** 694–7
[27] Felici F., Sauter O., Coda S., Duval B.P., Goodman T.P., Moret J.M. and Paley J.I. 2011 *Nucl. Fusion* **51** 083052
[28] Felici F. 2016 Real-time model-based plasma state estimation, monitoring and integrated control in TCV, ASDEX-Upgrade and ITER *Preprint: 2016 IAEA Fusion Energy Conf. (Kyoto, Japan, 17–22 October 2016)* pp EX/P8–33 (https://conferences.iaea.org/indico/event/98/session/32/contribution/519)
[29] Ferron J.R., Walker M.L., Lao L.L., St John H.E., Humphreys D.A. and Leuer J.A. 1998 *Nucl. Fusion* **38** 1055
[30] Moret J.M., Duval B.P., Le H.B., Coda S., Felici F. and Reimerdes H. 2015 *Fusion Eng. Des.* **91** 1–15
[31] Galperti C. *et al* 2014 *Plasma Phys. Control. Fusion* **56** 114012
[32] Van Den Brand H., De Baar M., Van Berkel M., Blanken T., Felici F., Westerhof E. and Willensdorfer M. 2016 *Plasma Phys. Control. Fusion* **58** 075002

[33] Felici F., De Baar M. and Steinbuch M. 2014 A dynamic state observer for real-time reconstruction of the tokamak plasma profile state and disturbances *Proc. American Control Conf. (Portland, OR, USA, 4–6 June 2014)* (IEEE) pp 4816–23 (http://ieeexplore.ieee.org/document/6858781/)

[34] Poli E. *et al* 2018 *Comput. Phys. Commun.* **225** 36–46

[35] Cassandras C.G. and Lafortune S. (ed) 2008 *Introduction to Discrete Event Systems* (Boston, MA: Springer) (https://doi.org/10.1007/978-0-387-68612-7)

[36] Baeten J.C.M., van de Mortel-fronczak J.M. and Rooda J.E. 2016 Integration of supervisory control synthesis in model-based systems engineering *Complex Systems. Studies in Systems, Decision and Control* ed G. Dimirovski, vol 55 (Cham: Springer) pp 39–58

[37] Sweeney R., Choi W., La Haye R.J., Mao S., Olofsson K.E. and Volpe F.A. 2017 *Nucl. Fusion* **57** 016019

[38] Galperti C., Coda S., Duval B.P., Llobet X., Milne P., Sauter O., Moret J.M. and Testa D. 2017 *IEEE Trans. Nucl. Sci.* **64** 1446–54

[39] Shousha R. *et al* (the TCV team) 2018 A real-time model-based detector of LH-transitions, HL-transitions and ELMs and its applications to TCV *Nuclear Fusion* in preparation

[40] Blanken T.C., Felici F., Rapson C.J., de Baar M.R. and Heemels W.P. 2018 *Fusion Eng. Des.* **126** 87–103

[41] The Mathworks Inc. 2018 MATLAB Stateflow (https://mathworks.com/products/stateflow)

[42] Kong M. *et al* (The TCV-Team and The EUROfusion MST1 Team) 2017 Real-time control of neoclassical tearing modes and its integration with multiple controllers in the TCV tokamak *44th EPS Conf. on Plasma Physics (Belfast, Northern Ireland (UK), 26–30 June 2017)* p 152

[43] Best R.E. 2007 *Phase-Locked Loops: Design, Simulation, and Applications* 6th edn (London: McGraw-Hill) p 490

[44] Sheikh U.A., Duval B.P., Galperti C., Maraschek M., Sauter O., Sozzi C., Granucci G., Kong M., Labit B. and Merle A. 2018 *Nucl. Fusion* **58** 106026

[45] Wonham W.M., Cai K. and Rudie K. 2017 *IFAC-PapersOnLine* **50** 1791–7

[46] Ramadge P.J. and Wonham W.M. 1987 *SIAM J. Control Optim.* **25** 206–30

[47] Snipes J.A., Gribov Y. and Winter A. 2010 *Fusion Eng. Des.* **85** 461–5

[48] Woo M.H., Hahn S.H., Jeon Y.M., Kim J., Han H., Bak J.G., Oh Y.K., Eidietis N.W., Johnson R.D. and Walker M.L. 2018 *Fusion Eng. Des.* **128** 168–74

[49] Meyer H. *et al* 2017 *Nucl. Fusion* **57** 102014

[50] Coda S. *et al* 2017 *Nucl. Fusion* **57** 102011