

D. Alves, R. Coelho, A.C. Neto, P. Smith, D.F. Valcarcel, P. Card, R. Felton,
P.J. Lomas, P. McCullen and JET EFDA contributors

A Flexible System for the Control of External Magnetic Perturbations in the JET Tokamak

“This document is intended for publication in the open literature. It is made available on the understanding that it may not be further circulated and extracts or references may not be published prior to publication of the original when applicable, or without the consent of the Publications Officer, EFDA, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK.”

“Enquiries about Copyright and reproduction should be addressed to the Publications Officer, EFDA, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK.”

The contents of this preprint and all other JET EFDA Preprints and Conference Papers are available to view online free at www.iop.org/Jet. This site has full search facilities and e-mail alert options. The diagrams contained within the PDFs on this site are hyperlinked from the year 1996 onwards.

A Flexible System for the Control of External Magnetic Perturbations in the JET Tokamak

D. Alves¹, R. Coelho¹, A.C. Neto², P. Smith³, D.F. Valcarcel¹, P. Card³, R. Felton³, P.J. Lomas³, P. McCullen³ and JET EFDA contributors*

JET-EFDA, Culham Science Centre, OX14 3DB, Abingdon, UK

¹*Associação EURATOM/IST, Instituto de Plasmas e Fusão Nuclear, Instituto Superior Técnico, Universidade de Lisboa, 1049-001 Lisboa, Portugal*

²*Fusion for Energy, 08019 Barcelona, Spain*

³*EURATOM-CCFE Fusion Association, Culham Science Centre, OX14 3DB, Abingdon, OXON, UK*

** See annex of F. Romanelli et al, "Overview of JET Results", (24th IAEA Fusion Energy Conference, San Diego, USA (2012)).*

ABSTRACT

External magnetic perturbations are typically utilized in tokamak devices with two operational or experimental purposes: the correction of the intrinsic 3D error fields and the mitigation or suppression of Edge Localized Modes (ELMs). At JET, dedicated coils are used for the *generation* of these toroidally asymmetric perturbations. While error fields exist even in the absence of plasma, in ELM mitigation experiments the external fields are meant to ergodize slightly the magnetic topology in the plasma periphery hence reducing the drive for the destabilization of these instabilities. The control of the magnetic field produced by these coils is achieved by controlling the current flowing in them. The real-time system responsible for this control, recently underwent a number of functional improvements since its original implementation utilizing the present voltage controlled voltage sources.

This work describes the overall system, built-in functionality, control algorithms and presents preliminary experimental results along with performance assessment studies. In particular, the main improvements are: the possibility of automatically reducing the current references in case the plasma amplifies the applied perturbation, real-time limitation of dI/dt to reduce the electromotive force in machine protection diagnostic systems, implementation of a model predictive controller as an alternative to the PID and the possibility of adapting the current references, in real-time, using an external system. The result is a flexible control system which is able to fulfil the operational and experimental requirements of an international and dynamic scientific environment.

1. INTRODUCTION

The use of non-axisymmetric magnetic perturbations is presently being explored in many tokamak devices such as DIII-D, ASDEX-U, MAST and JET. They serve, primarily, two main purposes: error field correction and Edge Localized Mode (ELM) mitigation/suppression [11]. The former tries to address and ameliorate the inevitability of magnetic field asymmetries (due to coil imperfections and the presence of magnetic materials) which can resonantly interact with plasma instabilities causing them to grow and ultimately cause the plasma to disrupt thus imparting significant heat loads onto facing components. The latter relies on the observation that, in certain operational regimes, applying external magnetic perturbations increases the frequency of ELM crashes (non-linear stage of the peeling-ballooning instability and driven by the plasma pedestal current density and pressure gradient [20]) thus reducing their individual intensity and heat loads onto the facing components. Furthermore, low intensity time oscillating magnetic perturbations are also used for Resonant Field Amplification (RFA) studies diagnosing plasma stability near the so-called *no-wall limit* [28].

At JET, four coils placed 90° apart, between the iron limbs and the plasma containing vessel, are dedicated to *generating* external magnetic perturbations, see Fig. 1. Regardless of their present multi-purpose nature, they were historically baptised the Error Field Correction Coils (EFCCs).

Current in these coils is driven by two dedicated Voltage Controlled Voltage Sources (VCVSs) which, in turn, are controlled by the system presented herein.

This work is organized as follows: Section 2 presents a very succinct overview of the EFCC current control system in the JET context and section 3 describes the hardware involved including: the EFCCs, the voltage amplifiers and the control system. Section 4 describes all software with special emphasis on the MARTe framework, synchronisation mechanism and the algorithms implemented in each of the main control modules. Finally, section 5 presents some preliminary results of the operation of the system including performance assessments and in section 6 conclusions are drawn and the anticipated future work is discussed.

2. OVERVIEW OF THE EFCC CURRENT CONTROLLER SYSTEM

The EFCC current controller system at the Joint European Torus (JET) tokamak has been implemented using Versa Module European (VME) technology. Further hardware details are given in section 3 and a basic data flow diagram of the system in JET's context is shown in Fig. 2. The system receives timing information, optically, from JET's Central Timing System (CTS) and uses JET's Asynchronous Transfer Mode (ATM) based Real-Time Data Network (RTDN) [12] for a number of purposes: reporting control errors and operational limits violations to the Real-Time Protection Sequencer (RTPS) [30], receiving current references from the Real-Time Central Controller (RTCC) [21] and measurements from magnetic diagnostics. The system also receives a series of analog measurements: plasma current and mode-lock signal from magnetic diagnostics as well as current and voltage measurements from the VCVSs. The analog outputs of the system are the VCVSs' voltage requests. Finally, the system receives digital stop information from the Plasma Termination Network (PTN) [17].

Typically at JET the person leading the experimental session, also known as the Session Leader (SL), pre-configures the target current waveforms for the EFCCs. In order to do this, a dedicated page is available in JET's pulse schedule editor [35] commonly known as *Level-1* (Fig. 3). Using this time-window based graphical utility, the SL is capable of designing the target current references (either in a parametric, e.g. sinusoidal, or *point-based* form), configure protection related parameters, check the predicted I^2t , choose the control algorithm and decide for which time windows to enable control on references produced, in real-time, by RTCC.

3. HARDWARE

3.1. ERROR FIELD CORRECTION COILS

As previously mentioned, the EFCCs [5] are 4 approximately square shaped coils (side length ≈ 6 m) located 90° apart around the exterior of the plasma containment vessel, see Fig. 1. Each coil has 16 turns, a conductor cross section of 150mm^2 and is allowed to operate up to the I^2t limit of $252 \times 10^6 \text{A}^2\text{s}$. Some electromagnetic parameters of the EFCCs are presented in [24].

3.2. VOLTAGE CONTROLLED VOLTAGE SOURCES

The VCVSs driving the EFCC current consist of two units, VCVS₁₂ and VCVS₃₄. The configuration of these 12-phase thyristor bridge rectifiers has evolved since they were first used at JET for the vertical stabilization of the plasma more than twenty years ago [9]. At that time they were baptised the Poloidal Radial Field Amplifiers (PRFAs). Presently they serve as two quadrant voltage controlled voltage amplifiers rated to provide -3 to $+3$ kV and 0 to 6 kA. Although the original manufacturer's documentation [16] specifies a bandwidth of 75 Hz, close to full voltage swings are limited by the 50 Hz mains supply either on the rising or falling edges depending on the polarity.

3.3. CONTROLLER SYSTEM

The EFCC current controller is a VME based system containing the following cards:

VME Programmable Logic Service (VPLS)

JET specific card providing central timing and digital event information with 1 ms temporal resolution;

Motorola MVME5110

400 MHz PowerPC card with 512 MB of RAM and on-board ethernet for slow monitoring, state transitions, pre-pulse configuration and post-pulse data collection;

Pentland MPV956

16 channel multiplexed Analog-to-Digital Convertor (ADC) with a maximum acquisition rate of 250 kSample/s and 12 bit resolution ($[-10, 10]$ V) and 8 channel Digital-to-Analog Convertor (DAC) with 12 bit resolution ($[0, 10]$ V);

Pentland MPV922

40 digital input channels and 32 digital output channels;

Interphase ATM NIC

ATM interface for real-time I/O.

4. SOFTWARE

The combination of VME, PowerPC and the VxWorks Operating System (OS) has become one of the standards for real-time systems at JET. Therefore, adhering to this standard, the EFCC current controller system runs on top of version 5.5 of the VxWorks real-time OS and was developed using the Multi-threaded Application Real-Time executor (MARTe) framework for real-time systems. This section describes the software components of the application.

4.1. MULTI-THREADED APPLICATION REAL-TIME EXECUTOR

MARTE [27] is a C++ multi-platform framework for the development of modular and data-driven real-time applications and was originally developed at JET. It is built on top of the *BaseLib2* library. At its lowest level, this library implements an abstraction layer of basic functionality such as threads, semaphores or mutexes. This is where all OS-dependent code resides. Building on the previous, a comprehensive set of OS-independent higher level functionality is provided. This includes features such as object management, garbage collection, configuration parsing and management, data-driven object creation, http services and utilities (live introspection), state machines and messaging. The library presently supports the Linux, Linux/RTAI, VxWorks, Solaris, Windows and the MacOSX OSs.

One of the main strengths of the MARTE framework is the ability to develop, test and debug applications in non-real-time environments and subsequently deploy them in real-time targets without code changes due to the OS abstraction feature. Testing and commissioning time is minimized. Furthermore, applications can be developed in user-space [31] (e.g. Linux or MacOSX), where debugging and profiling tools are vastly available, and subsequently deployed in kernel-space [31] (e.g. VxWorks or Linux/RTAI), where it's much harder to diagnose run-time problems. Having been designed with strict modular requirements in order to maximize code reuse (one of the main reasons supporting the choice of the C++ object-oriented language), it also implements a clear separation between physics/engineering algorithms, hardware interaction and interfacing with the outside world. The fact that MARTE is a highly configurable multi-threaded framework means that it is particularly suited for exploring the vastly available and economically accessible modern multi-core architectures such as x86. In addition, it complies with modern real-time programming *paradigms* by, for instance, making use of mutex priority inheritance and real-time schedulers.

Over the last few years, MARTE is becoming increasingly adopted by the magnetic confinement fusion community mainly in Europe. At JET [2, 10], examples of real-time applications using this framework are: the *vertical stabilisation* system (Linux/RTAI) [6], the present and previous versions of the current controller of the EFCCs (VxWorks) [3], the *real-time protection sequencer* (VxWorks) [30], the *vessel thermal map* (Linux) [1], the *wall load limiter system* (Linux) [32] the *advanced predictor of disruptions* (Linux) [22], and the *hard X-ray and gamma-ray profile monitor* (Linux) [13]. In Europe [26], it is actively supported and developed [33] by a dedicated community and, apart from JET, example applications can be found in the ISTTOK [8], COMPASS [34], FTU [7] and RFX-mod [23] tokamaks. Finally, also finding its way through to the tokamak community outside Europe, MARTE is being used in the KSTAR tokamak [36] and is presently under consideration for being included in the plant control system of the ITER tokamak [15].

4.2. SYNCHRONIZATION

The processing time of each control cycle was optimized to run in under 500 μs . Although JET's central timing is made available to the system with a 1 ms resolution interrupting the Central Pro-

cessing Unit (CPU) via a VME mailbox Interrupt ReQuest (IRQ), VxWorks' auxiliary clock timer is able to provide similar functionality, within the CPU itself, at higher rates. Hence, the start of each real-time cycle is triggered by the *release* of a semaphore in the Interrupt Service Routine (ISR) handling the VxWorks timer IRQ configured to expire continuously at a rate of 2 kHz . The determinism of the application's cycle time can be inferred from Fig. 4. In particular, from Fig. 4(b), it can be seen that the occurrence of a control cycle with an absolute jitter larger than 1% of the nominal application's cycle time ($500\mu\text{s}$) is less than 1 every 100 cycles. In order to guarantee consistent time stamping of collected data with respect to JET's central timing, whilst maintaining the 2 kHz control cycle, a scheme has been devised, see Fig. 5, making use of the PowerPC's Time Stamp Counter (TSC), a 64 bit counter incremented every clock cycle. Hence, the sample time t_k attributed to signal samples either acquired or processed during the application's k -th cycle is given by (1), where τ_k is the integer value of the TSC and χ_k is an estimate of the TSC's period given by the exponential moving average $\chi_k = 0.9995\chi_{k-1} + 0.0005\chi_k^m$ and where $\chi_k^m = 10^{-3}/\Delta\tau^{(v)}$ is the latest available measurement of the TSC's period. Also in (1), $t^{(v)}$ and $\tau^{(v)}$ are, respectively, the time given by the VPLS module (with a 1 ms resolution) and the value of the TSC when the VPLS IRQ immediately before the present control cycle occurred.

$$t_k = t^{(v)} + \chi_k (\tau_k - \tau^{(v)}) \quad (1)$$

4.3. CONTROL MODULES

This section describes the various modules performing the global control task. A basic data flow diagram of the control algorithm is depicted in Fig. 6 henceforth used as reference. Detailed descriptions of the different modules and their interactions are presented below in individual sub-sections.

The plant itself essentially comprises the VCVS in an equivalent RL circuit containing the inductors (EFCCs) and a finite resistance. Using the Euler method to discretize a Thèvenin equivalent RL circuit yields (2), where V_k is the output voltage of the VCVS, I_k is the current flowing in the EFCC circuit and R_{eq} , L_{eq} and $F_s = 2\text{ kHz}$ are, respectively, the equivalent resistance, equivalent inductance and the sampling frequency of the system.

$$V_k = R_{eq}I_k + L_{eq}F_s (I_k - I_{k-1}) \quad (2)$$

Mutual inductances between the EFCC circuit and other circuits have been assessed and neglected in (2) to a good approximation. The mutual inductances were assessed by, while changing the currents in each of the other potentially coupled circuits individually, observing the voltage request necessary to control the current in the EFCC circuit to zero. It is worth noting that this was done using standard JET dry-runs and that there was no opportunity to perform systematic tests sweeping the rate of change of the currents in the various circuits thus inducing electromotive

forces of different magnitudes in the EFCC circuit. Because, under these conditions, no mutual inductances have been measured, their effect was neglected in (2).

All results presented herein were obtained with a single VCVS controlling the current in a circuit containing all four EFCCs. Basic analysis of circuit operation in steady-state conditions and of the exponential current decay $\propto \exp[-t/\tau]$, where $\tau = L_{eq}/R_{eq}$, established that $R_{eq} \approx 330 \text{ m}\Omega$ and $L_{eq} \approx 36.7 \text{ mH}$.

4.3.1. Kalman filter current estimation

Using the circuit model (2), and given that experimental measurements of the EFCC current and VCVS output voltage are available, it is possible to use the Kalman Filter (KF) [19] to estimate the EFCC current with a better Signal-to-Noise Ratio (SNR) than the actual measurement. The KF is a linear two-stage recursive estimator which combines theoretical knowledge of the system's temporal evolution (prediction stage) with system measurements (correction stage), in order to minimize the mean-square error of the state estimation. In this particular case, and in the prediction stage, a prior EFCC current estimate I_k^- can therefore be written as (3) where I_{k-1}^+ is the posterior estimate of the previous cycle, V_k is a measurement of the VCVS's output voltage, $\alpha = L_{eq}F_s$ and $\beta = 1/(R_{eq} + \alpha)$.

$$I_k^- = \alpha\beta I_{k-1}^+ + \beta V_k \quad (3)$$

Similarly, in the correction stage and according to the general KF equations, the posterior KF estimate of the EFCC current is given by (4).

$$I_k^+ = I_k^- + K_k \times (I_k^{meas} - I_k^-) \quad (4)$$

It is worth emphasizing that the KF gain K_k needs not be computed in real-time since this is a linear time-invariant KF implementation. In order to pre-calculate the gain K_k one has used $R = 600$ (the measured noise variance in I_k^{meas}) and $Q = 10^{-1} \times R$ (the process noise variance) to obtain $K_k = 0.2673$ after recursively evolving the associated Riccati equation [18] until a steady-state is reached and K_k is considered to be sufficiently stationary. Fig. 7 shows a comparative example between the raw current measurement I_k^{meas} and the KF's posterior current estimate I_k^+ where the SNR improvement is evident. The ratio between the variance of the noise in the KF's posterior current estimate and the variance of the noise in the current's raw measurement has been found to be $\lesssim 1/6$.

In order to further assess the performance of the KF current estimate, the frequency response in terms of amplitude and phase delay, when compared to the raw measurements, has been evaluated. Fig. 8 shows the results of this study for a frequency scan of 10 – 70 Hz in steps of 10 Hz. The DFT is used to calculate the amplitude of both the KF estimate and the raw measurement. Fig. 8(a) shows the difference between the two, normalized to the amplitude of the raw measurement. It can be seen that, up to 70 Hz, the amplitude attenuation of the KF's estimate never exceeds 10%

of the actual amplitude. Fig. 8(b) shows the delay between the KF's estimate and the current measurement, normalized to the period of the oscillation. The results were obtained using standard cross-correlation analysis of both signals and it can be seen that the delay of the KF's estimates never exceeds 7% of the oscillation's period.

Using this scheme and I_k^+ rather than I_k^{meas} for feedback control one reduces the propagation of the measurement noise into the control action itself. The price to pay is some amplitude and phase distortion of the current estimate with respect to the measurement.

4.3.2. Mode-lock compensation

The mode-lock signal is commonly used for protection in tokamak research. It measures the amplitude of the potentially disruptive $n = 1$ MHD activity. At JET this signal is given by (6) and its components are given by (5), where k denotes the discrete time index. Each $S_{x,k}^{Oct_y}$ signal is the magnetic field measurement at time index k obtained by integrating saddle loop x located in octant y , see Fig. 9.

$$m_{n=1,k}^{sin} = [S_{1,k}^{Oct_3} + S_{14,k}^{Oct_3}] - [S_{1,k}^{Oct_7} + S_{14,k}^{Oct_7}] \quad (5)$$

$$m_{n=1,k}^{cos} = [S_{1,k}^{Oct_1} + S_{14,k}^{Oct_1}] - [S_{1,k}^{Oct_5} + S_{14,k}^{Oct_5}]$$

$$m_{n=1,k} = \sqrt{[m_{n=1,k}^{sin}]^2 + [m_{n=1,k}^{cos}]^2} \quad (6)$$

It is often the case that the plasma amplifies these RMPs to the point where the mode eventually locks to the resistive wall and disrupts the plasma imparting large eddy currents and forces on the containment vessel. One way to ameliorate this effect would be to, as soon as the $m_{n=1,k}$ is observed to reach dangerous values, automatically lower the current in the EFCCs to try and reduce the plasma amplification. However, in order to do that, it is assumed that this signal measures exclusively the plasma response and not the contribution of the direct pick-up from the external magnetic perturbation. It is therefore clear that some means to compensate $m_{n=1,k}$, so that it contains exclusively the contribution of the plasma response, is necessary. The form for the compensated mode-lock signal $m'_{n=1,k}$ used herein is shown in (8) and its, individually compensated, components are given by (7).

$$m_{n=1,k}^{sin'} = m_{n=1,k}^{sin} - \left(\sum_{i=0}^{N-1} a_i I_{k-i}^{EFCC_3} - \sum_{i=0}^{N-1} b_i I_{k-i}^{EFCC_7} \right) \quad (7)$$

$$m_{n=1,k}^{cos'} = m_{n=1,k}^{cos} - \left(\sum_{i=0}^{N-1} c_i I_{k-i}^{EFCC_1} - \sum_{i=0}^{N-1} d_i I_{k-i}^{EFCC_5} \right)$$

$$m'_{n=1,k} = \sqrt{[m_{n=1,k}^{sin'}]^2 + [m_{n=1,k}^{cos'}]^2} \quad (8)$$

The N coefficients in (7) were obtained by fitting (linear least-squares), in a discharge without plasma, $m_{n=1,k}^{sin'}$ and $m_{n=1,k}^{cos'}$ to zero for a 4 s window (72 – 76 s as shown in Fig. 10), during which no other circuits were active. The continuous lines represent the uncompensated signals ($m_{n=1,k}^{sin}$ and $m_{n=1,k}^{cos}$) while the dots represent the compensated ones ($m_{n=1,k}^{sin'}$ and $m_{n=1,k}^{cos'}$). For computational reasons the number N of coefficients in (7) should be as small as possible. However, using $N = 1$ implies completely disregarding any eddy currents induced, e.g. in the vessel, due to changes in the EFCC currents. In fact it can also be seen in Fig. 10 that, for the chosen value of $N = 2$, a better compensation is achieved in the region where the currents are stationary, i.e. $m_{n=1,k}^{sin}$ and $m_{n=1,k}^{cos}$ are essentially constant. It is worth stressing that this is not meant to be an exact cancellation mechanism but more of an approximate amelioration scheme.

4.3.3. Reference multiplexing, adaptation and checks

This module is responsible for: (a) checking that the EFCC currents are within pre-determined boundaries; (b) performing the multiplexing¹ of the references pre-configured by the SL and the real-time references from RTCC and (c) adjusting the references according to events in PTN and the compensated mode-lock signal calculated in the previous module.

After multiplexing, the current references I_{k+1}^{ref} are set to zero in the event of a plasma termination request or are adjusted according to (10), where γ_k is given by (9). m_0 and Δm are configurable parameters and their effect on γ_k is exemplified in Fig. 11.

$$\gamma_k = \frac{1}{2} \left[1 + \tanh \left(\frac{4(m_0 - m'_k)}{\Delta m} \right) \right] \quad (9)$$

$$I_{k+1}^{ref'} = \gamma_k \times I_{k+1}^{ref} \quad (10)$$

This mechanism significantly reduces the reference currents while an $n = 1$ mode of considerable amplitude exists in the plasma. As soon as the mode amplitude is reduced, the references automatically return to the desired values.

4.3.4. Reference slope limitation

This module is responsible for implementing a maximum absolute rate of change (slew rate) in the EFCC currents. The reason for this is because the toroidally asymmetric fields are inductively picked-up by the protection system of the toroidal field coils circuit (the direct magnet safety system) leading it to believe that a non-existing problem has occurred. It has been determined experimentally that the pick-up threshold for triggering an alarm is when the EFCC currents are varying at a rate $dI/dt \gtrsim 40$ kA/s. The idea is therefore to calculate the slope obtained through a linear regression of the past N current estimates plus the current reference $I_{k+1}^{ref'}$ desired to be

¹In the present context, the term *multiplexing* is used to describe the selection process by which the current reference used in the control algorithm is chosen between the one pre-configured by the SL and the one made available by RTCC in real-time.

achieved in the next control cycle². Hence a linear regression of $N + 1$ points. If the calculated slope a , given by (11)³, is such that $|a| < |a_{th}|$, where $|a_{th}| = 40$ kA/s, then the new slope limited reference is the same as the adjusted reference calculated in the previous module, i.e. $I_{k+1}^{ref''} = I_{k+1}^{ref'}$. In case this condition is not satisfied, then it is possible to *reverse* (11) and calculate the new reference $I_{k+1}^{ref''}$ that satisfies $|a| = |a_{th}|$, thus optimizing the ramp-down rate to the fastest possible value without violating the slope threshold previously established. In this case $I_{k+1}^{ref''}$ is given by (12) where a_{th} is either ± 40 kA/s depending on the slope violation direction.

$$a = \left\{ N\Delta t \times I_{k+1}^{ref'} + \sum_{i=0}^{N-1} i\Delta t \times I_{k-i}^+ - \frac{1}{N+1} \left(\sum_{i=0}^N i\Delta t \right) \times \left(I_{k+1}^{ref'} + \sum_{i=0}^{N-1} I_{k-i}^+ \right) \right\} \times \quad (11)$$

$$I_{k+1}^{ref''} = \left\{ a_{th} \times \left[\sum_{i=0}^N (i\Delta t)^2 - \frac{1}{N+1} \left(\sum_{i=0}^N i\Delta t \right)^2 \right]^{-1} + \frac{1}{N+1} \left(\sum_{i=0}^N i\Delta t \right) \times \left(I_{k+1}^{ref'} + \sum_{i=0}^{N-1} I_{k-i}^+ \right) - \sum_{i=0}^{N-1} i\Delta t \times I_{k-i}^+ \right\} \times \left\{ N\Delta t - \frac{1}{N+1} \sum_{i=0}^N i\Delta t \right\}^{-1} \quad (12)$$

An example of this mechanism in action can be seen in Fig. 12 for a pulse where there was a termination request at $t \sim 73.55$ s. On every cycle, the last $N = 5$ KF current estimates plus the current reference for the next cycle ($I_{k+1}^{ref'}$) are used in linear regression's slope calculation. It can be seen from Fig. 12 that at $t \sim 73.55$ s the adjusted reference $I_{k+1}^{ref'}$ (in green) is immediately set to zero however, the actual reference used for control $I_{k+1}^{ref''}$ (the slope limited reference shown in red) exhibits a descent slope ~ 40 kA/s as desired. The latter is subsequently used to perform the desired controlled current ramp-down as depicted by the blue curve.

A much simpler mechanism could have been devised to perform the same task however, a significant advantage of the proposed method is that it is suitable for coping with the noisy current estimates. So, instead of using two noisy estimates to calculate a slope, using this method, and

²Given two sets of data points $\{x_i, y_i\}$, with $i = 1, 2, \dots, M$, the slope of the straight line $y = ax + b$ that best fits $\{x_i, y_i\}$ in the least-squares sense, i.e. $\min_{a,b} \sum_{i=1}^M (y_i - ax_i - b)^2$, is given by $a = \text{Cov}[x, y] / \text{Var}[x]$.

³ I_{k-i}^+ is the history of current estimates calculated by the KF module using (3) and (4) and $\Delta t = 1/F_s = 1/2000 = 500 \mu\text{s}$.

assuming that the KF current estimates are sufficiently unbiased (i.e. the noise associated with them is zero-mean), ensures that the slope calculation is obtained using the straight line that fits best the last N estimates plus I_{k+1}^{ref} .

It is also worth noting that, if this mechanism is to be used not only when there has been a pulse termination request, the choice of N may impact on the overall bandwidth of the system. For example, in the case just shown where $N = 5$, means that $\Delta T = N/F_s = 5/2000 = 2.5$ ms which corresponds to a limiting frequency $f_l = 1/\Delta T = 400$ Hz, well above the overall system bandwidth.

4.4. CONTROL MODES

The user can choose between a PID and a MPC based algorithm.

4.4.1. Proportional Integral Derivative control

The PID [14] is, together with the bang-bang⁴ type of control [29], certainly the most widely used control scheme. This mechanism is based on the difference between the desired value and the present value of the controlled parameter. This difference is commonly referred to as the *error* signal. In the case of the present application the error is given by $e_k = I_{k+1}^{ref} - I_k^+$, i.e., the difference between the desired value of the current for the next control cycle and its present value. In its general form, the control action V_k^{req} , i.e. the request to the VCVS, is calculated using (13) where K_p , K_i and K_d are design parameters. The control action is therefore based on a linear combination of the error, its integral and its derivative. It is expected that the effect of the control action steers the system in the direction of reducing $|e_k|$ however, that may not always be the case.

$$V_k^{req} = K_p \times e_k + \frac{K_i}{F_s} \times \sum_{i=1}^k \frac{e_i + e_{i-1}}{2} + K_d F_s \times (e_k - e_{k-1}) \quad (13)$$

Although the PID is a well established and thoroughly studied method, typically credited for its robustness and for its obliviousness of the nature of the underlying process, it is hardly exempt of weaknesses [4]. In particular, the fact that it doesn't take into account information about the underlying process means that, rather than predicting the system trajectory in state-space and anticipating the control action, the PID has to *wait* for the system to deviate from the desired state in order to start taking some action. It is therefore intuitive to realize that, especially under non-stationary conditions, any system output controlled with a PID is always lagging behind its target state. Also, because the control action's calculation is solely based on a linear combination (with fixed coefficients⁵) of linear operators acting on the error signal, means that it cannot be the solution of an optimization problem in the entire operational domain. That is to say that PID control is not optimal.

⁴Bang-bang control, also known as on-off control, is a scheme for which typically the control action is either disabled or at a fixed *intensity* level.

⁵Gain scheduling is certainly possible but assumes some knowledge about the process.

In the present context, the use of external magnetic perturbations to measure the RFA effect on JET plasmas requires operation of the EFCCs at various frequencies within the system’s capabilities, typically in the range 0 – 50 Hz. PID-based control is incapable of providing optimal control across this whole range of frequencies. In the present application one has chosen to, while giving the user the option of using a basic PID algorithm, develop a more sophisticated model-based approach in an attempt to overcome some of the aforementioned PID weaknesses.

4.4.2. Model Predictive Control

The model-based method chosen was MPC. The main principle of MPC [25], also known as receding horizon control, is to solve the open-loop control problem, in real-time and on every control cycle for a finite temporal horizon using the present plant state as an initial condition, in order to determine the sequence of control actions optimizing a given criterion. Although the optimal set of control actions, at a given time instant, is calculated for a finite prediction horizon, only the control action for the present is actually executed. Whenever possible, the most relevant physical limitations of the system, such as actuator bounds and state domain, are taken into account so that the solutions of the optimization process are constrained to a realistic operational space.

Clearly an MPC implementation can be rather demanding in terms of computational burden. Keeping in mind that each optimization process needs to be performed and complete within two consecutive control cycles, the choice of the prediction horizon is tightly coupled with the system’s bandwidth specifications/requirements which, in turn, strongly influence the choice of the control period. One of the greatest strengths of MPC-based control is that it solves the optimal control problem for a finite prediction horizon, in real-time, using the present plant state as an initial condition rather than, as is done in conventional control approaches, pre-computing off-line a control law (even if optimal with respect to some criterion) over all possible system states.

The optimal control problem is expressed as a cost function minimization problem. A typical form for the MPC cost function is given by (14), where N_p is the number of discrete time instants in the prediction horizon and $\|\cdot\|_2$ denotes the 2-norm. \mathbf{u}_k is the set of control variables i.e. the physical quantities used to control the system, and $\Delta\mathbf{u}_k$ denotes their change between two consecutive control cycles. $\mathbf{p}_k^{pred} = \mathbf{p}_k^{pred}(\mathbf{p}_{k-1}^{pred}, \mathbf{u}_k)$ denotes the recursive forecast, within the N_p sample prediction horizon, of the set of process variables being controlled (a function of their starting value \mathbf{p}_{k-1}^{pred} and the inputs \mathbf{u}_k) while \mathbf{r}_k denotes their target values. It is worth noting that, for $i = 1$, $\mathbf{p}_k^{pred} = \mathbf{p}_k^{pred}(\mathbf{p}_{k-1}^{pred}, \mathbf{u}_k)$, i.e., the initial condition is the estimate rather than the prediction of the process variable set in the previous control cycle. Finally, w_p , w_u and $w_{\Delta u}$ are simple weights. The leftmost term on the right-hand side of (14) penalizes deviations of the variables under control from their target values. The middle term penalizes the *intensity* of the control action while the rightmost term penalizes the excursions in consecutive control actions. While minimizing the leftmost term is actually the ultimate goal of the entire optimization process, the other terms provide soft constraints that often help mitigate practical problems related to the control effort. It

is worth noting that, in case J_k is written in the form given by (14) and $\mathbf{p}_k^{pred}(\mathbf{p}_{k-1}^{pred}, \mathbf{u}_{k-1})$ is linear, minimizing J_k is an analytically solvable linear least-square optimization problem.

$$J_k = \sum_{i=1}^{N_p} \{w_p \|\mathbf{p}_{k+i}^{pred}(\mathbf{p}_{k+i-1}^{pred}, \mathbf{u}_{k+i-1}) - \mathbf{r}_{k+i}\|_2^2 + w_u \|\mathbf{u}_{k+i-1}\|_2^2 + w_{\Delta u} \|\Delta \mathbf{u}_{k+i-1}\|_2^2\} \quad (14)$$

Let one assume that the cost function J_k one wishes to minimize in every control cycle k is given by (15). The aim is therefore to calculate the voltage requests, for the present (V_k^{req}) and the subsequent control cycle (V_{k+1}^{req}), that minimizes J_k . Only V_k^{req} , the voltage request performed by the controller at the end of the present control cycle k , is actually performed. In this case, the prediction horizon is $N_p = 2$. Having specified the desired bandwidth of the system to be 100 Hz, it is reasonable to have a control application running at a rate 10 – 20 times faster. A control cycle of 2 kHz was chosen. At this rate, tests revealed that performing the online optimization within the application's cycle time was not possible if $N_p > 2$ with the presently available hardware resources. Therefore, the controller has been implemented with $N_p = 2$.

$$J_k = \left(I_{k+1}^{pred} - I_{k+1}^{ref} \right)^2 + \left(I_{k+2}^{pred} - I_{k+2}^{ref} \right)^2 + \mu \left(V_k^{req} - V_{k-1}^{req} \right)^2 + \varepsilon \left(V_{k+1}^{req} - V_k^{req} \right)^2 \quad (15)$$

In this case there is only one process variable under control which is the estimate of the current (I_k^+) in the EFCC circuit as given by the KF, see (4). In the light of the previous discussion one has that $I_{k+1}^{pred} = I_{k+1}^{pred}(I_k^+, V_k^{req})$ and $I_{k+2}^{pred} = I_{k+2}^{pred}(I_{k+1}^{pred}, V_{k+1}^{req})$. These will therefore need to be obtained and substituted inside (15). It is worth clarifying that I_{k+1}^{pred} is a prediction, performed during control cycle k , of what the current in the circuit will be at the beginning of control cycle $k + 1$, that is, after the voltage request V_k^{req} at the end of control cycle k has been performed.

In order to proceed, let one now consider the behaviour of the VCVS. The VCVS, being a non-perfect voltage amplifier, will have a finite complex transfer function associated with it. This is to say that the VCVS does not instantaneously provide the requested voltage. Let one specify a linear model of the form given by (16) to express the VCVS's output voltage V_k , at time index k , as a linear combination of past output voltage measurements and the VCVS's readback⁶ of past voltage requests V_k^{rbk} . Two assumptions are implicit here: (1) intuition-based assumption that the VCVS's output voltage depends not only on the present and past voltage requests but also on the past output voltages and (2) these are linear dependencies. Note that the two most recent output voltage measurements are assumed to be unknown. The one referring to $i = 0$ (b_0) is a consequence of the present output voltage being on the left-hand side of (16). The absence of the b_1 coefficient is because, as will become clear shortly below, if one uses (16) to calculate V_{k+2}^{pred} , the output voltage

⁶ V_k^{rbk} comes originally from an analog signal, produced by the VCVS, containing the last voltage request acknowledged by the VCVS itself. This signal is subsequently digitized in the controller application. In noise-less environments, the voltage request performed at the end of control cycle k (V_k^{req}) is exactly the same as the readback voltage at the beginning of control cycle $k + 1$ (V_{k+1}^{rbk}).

at time index $k + 1$ (V_{k+1}) will not be known/available yet⁷.

$$V_k = \sum_{i=2}^{N_v} b_i V_{k-i} + \sum_{i=0}^{N_r} c_i V_{k-i}^{rbk} \quad (16)$$

If one now assumes that $V_k^{req} \approx V_{k+1}^{rbk}$ (i.e. the voltage request calculated and performed at the end of control cycle k is approximately the voltage request acknowledged by the VCVS at the beginning of control cycle $k+1$), and therefore that $V_{k+1}^{req} \approx V_{k+2}^{rbk}$ then, using ((16)), yields ((17)). V_{k+1}^{pred} and V_{k+2}^{pred} are the VCVS's output voltage predictions for the future (for time indexes $k+1$ and $k+2$, respectively). It is worth stressing out that these predictions are performed in the present, i.e., at time index k .

$$\begin{cases} V_{k+1}^{pred} = \sum_{i=1}^{N_v-1} b_{i+1} V_{k-i} + c_0 V_k^{req} + \sum_{i=1}^{N_r} c_i V_{k-i}^{rbk} \\ V_{k+2}^{pred} = \sum_{i=0}^{N_v-2} b_{i+2} V_{k-i} + c_0 V_{k+1}^{req} + c_1 V_k^{req} + \sum_{i=2}^{N_r} c_i V_{k-i}^{rbk} \end{cases} \quad (17)$$

System identification techniques are required to determine the optimal set of coefficients b_i and c_i for reproducing the VCVS's behaviour. In order to do this, a frequency scan (0 – 70 Hz in 10 Hz steps) was performed where the VCVS was driven in pre-programmed voltage control, see Fig. 13, and the coefficients were subsequently determined from the acquired data using Matlab's auto-regressive exogenous modelling function. The reported accuracy of fit was 86% and the coefficients obtained are depicted in Fig. 14. It was observed that, for N_v and $N_r > 11$, the improvement in the accuracy of the fit was marginal. It is evident that, in overall terms, the voltage request has a bigger impact, especially in the first four time lags, than the set of past output voltages.

Basic VCVS tests revealed that it was in fact unable to provide an absolute output voltage larger than 1.8 kV. Fig. 15 shows the operational domain of the amplifier in the context of an $N_p = 2$ prediction horizon. In the same figure, the red labels numbered 1 – 8 designate a particular zone in the frontier of the domain and 0 its enclosure. These will be used later on as a reference for the constrained minimization of the cost function.

In order to finalize the derivation of the MPC-based control equations for the problem at hand, let one start from the general RL circuit equation (2) and rewrite it in the form given by (18), very similar to the KF's prior current estimate (3). I_k is the current flowing in the EFCC circuit, V_k is the VCVS's output voltage, $\alpha = L_{eq} F_s$ and $\beta = 1/(R_{eq} + \alpha)$. As previously discussed R_{eq} , L_{eq} and F_s are, respectively, the equivalent resistance, equivalent inductance and sampling frequency. The mutual inductances are assumed to be zero.

$$I_k = \alpha\beta I_{k-1} + \beta V_k \quad (18)$$

⁷Note that this predictive calculation is performed in the present, i.e. in control cycle k .

In the light of the previous, and in order to obtain the current predictions I_{k+1}^{pred} and I_{k+2}^{pred} for the next two control cycles one writes (19). It is worthwhile noting that for the first prediction the starting point is the KF's present current estimate I_k^+ while for the second prediction the starting point is the previous one.

$$\begin{cases} I_{k+1}^{pred} = \alpha\beta I_k^+ + \beta V_{k+1}^{pred} \\ I_{k+2}^{pred} = \alpha\beta I_{k+1}^{pred} + \beta V_{k+2}^{pred} \end{cases} \quad (19)$$

So, substituting (17) in (19) and subsequently the latter in (15) yields $J_k = J_k(V_k^{req}, V_{k+1}^{req})$, where all other quantities are either available measurements/estimates such as I_k^+ , V_{k-i} and V_{k-i}^{rbk} , or known parameters such as R_{eq} , L_{eq} and F_s . It is straightforward to verify that the minimization of this expanded version of (15) expresses a convex problem hence admitting only one minimum which is the necessarily the global.

In order to find V_k^{req} and V_{k+1}^{req} that provides the unconstrained minimum of J_k one has to solve, as usual, $\nabla J(V_k^{req}, V_{k+1}^{req}) = \vec{0}$ yielding (20), where $\psi = \alpha\beta c_0 + c_1$.

$$\begin{pmatrix} \beta^2 (c_0^2 + \psi^2) + \mu + \xi & \beta^2 c_0 \psi - \xi \\ \beta^2 c_0 \psi - \xi & \beta^2 c_0^2 + \xi \end{pmatrix} \begin{pmatrix} V_k^{req} \\ V_{k+1}^{req} \end{pmatrix} = \begin{pmatrix} f(I_k^+, I_{k+1}^{ref}, I_{k+2}^{ref}, V_k^{req}, V_{k-i}, V_{k-i}^{rbk}) \\ g(I_k^+, I_{k+2}^{ref}, V_k^{req}, V_{k-i}, V_{k-i}^{rbk}) \end{pmatrix} \quad (20)$$

The previous is of the form $\mathbf{\Lambda} \times \mathbf{V}^{req} = \mathbf{\Gamma}$. The square matrix $\mathbf{\Lambda}$ admits an inverse if (21) is satisfied. It is straightforward to note that: if R_{eq} , L_{eq} , μ , ξ , c_0 and c_1 are all positive, then (21) is always satisfied.

$$\det(\mathbf{\Lambda}) = \beta^2 (\beta^2 c_0^4 + 2c_0^2 \xi + \psi^2 \xi + \mu c_0^2 + 2c_0 \psi) + \mu \neq 0 \quad (21)$$

Under these assumptions, the solution of the unconstrained minimization problem is always given by $\mathbf{V}^{req} = \mathbf{\Lambda}^{-1} \times \mathbf{\Gamma}$. Although by itself this is, in principle, already an improvement over PID-based control, there is still the VCVS's operational domain to take into account in the minimization process. Strictly speaking it becomes a linear least-squares minimization problem with inequality constraints. Although in principle solvable using Lagrange multipliers or more sophisticated quadratic programming methods, the approach taken herein was far simpler. First the unconstrained minimization problem is solved and, should the solution \mathbf{V}^{req} lie inside the VCVS's operational domain, no further processing is required. Should it lie beyond such domain, then the frontier is analytically scanned to find out the values of V_k^{req} and V_{k+1}^{req} for which J_k is minimal.

In the following section, preliminary results are shown regarding the VCVS modelling and controller performance.

5. PRELIMINARY RESULTS

In this section, preliminary results of the operation of the system are presented. In particular, an example is shown of mixed operation of the controller system interleaving pre-configured current references with the ones provided by RTCC in real-time thus commissioning the reference multiplexing logic.

Also, a basic assessment of the quality of the VCVS modelling and controller performance is done in comparative terms for both the PID and MPC-based control modes. Regarding MPC control, also different configurations of μ and ξ parameters were tested.

5.1. CURRENT REFERENCE MULTIPLEXING LOGIC ASSESSMENT

This section illustrates the commissioning tests of the reference multiplexing logic, see Fig. 16. The blue curve represents the KF posterior estimate I_k^+ of the current flowing in the EFCC circuit, the red curve is the pre-configured SL reference and the black curve is the real-time reference computed at a rate of 500 Hz by RTCC and sent to the controller system over JET's ATM-based RTDN. The dark and light gray areas are pre-configured by the SL, respectively, as the time windows for which control from RTCC isn't and is allowed. Furthermore, even if in control, RTCC can yield it by instructing the controller to revert the current references used for control to the ones pre-configured by the SL. So, analysing Fig. 16 one notes that, for $t < 28$ s, RTCC is requesting control but the controller is following the SL reference. Then, for $28 < t < 29$ s, RTCC refrains its control request. For $29 < t < 31$ s RTCC requests control and is finally able to get it (notice that the current starts to follow the black reference rather than the red one) starting from $t = 30$ s corresponding to the beginning of the allowed window. Further ahead, for $31 < t < 32$ s and although inside the window that permits RTCC control, RTCC yields control and the current goes back to following the SL (red) reference. Finally, for $32 < t < 33.5$ s, RTCC claims back control and the current starts tracking the real-time reference once more. At $t = 33.5$ s, a stop test is performed to evaluate the responsiveness of the system under those circumstances.

This test successfully completes the commissioning of the reference multiplexing feature.

5.2. VCVS MODEL PERFORMANCE

In this section, the quality of the VCVS modelling is succinctly analyzed.

A current reference waveform consisting of a 3 kA baseline with seven 0.5 s time windows of AC operation with a frequency range of 10 – 70 Hz in steps of 10 Hz was pre-configured, see Fig. 17. The amplitude of each AC portion was chosen not to exceed the dI/dt limits so that there was no chance of tripping the toroidal field circuit's protection system, see 4.3.4. Since $dI/dt \propto$

$\omega \times A$, where ω is the AC frequency and A is its amplitude, an increase in ω needs to be compensated by a decrease in A in order to limit dI/dt .

The idea here is basically to assess the accuracy of the output voltage prediction V_{k+1}^{pred} which is based on the VCVS model. The VCVS model, as previously discussed, is a linear model given by (16) and for which the coefficients b_i and c_i were calculated from the data obtained in an experiment performed in pre-programmed voltage control (hence without feedback) as shown in Fig. 13.

Four different controller configurations were used to control the current in the EFCC circuit using the reference depicted in Fig. 18. The data is taken from JET pulses 83737 (blue), 83735 (green), 83736 (red) and 83778 (cyan). The first three controller configurations are MPC-based and characterised by $\mu = \xi = 10^{-4}$, $\mu = \xi = 10^{-3}$ and $\mu = \xi = 10^{-2}$, respectively. The fourth controller configuration is PID-based with $K_p = 5$, $K_i = 20$, $K_d = 0$. A set of relevant time traces obtained (within the 30 Hz time window) are shown in Fig. 18 for the different controller configurations. The first noticeable thing is the saturation of the voltage request in all configurations. The reason for this is that the DC (resistive) component alone consumes almost 1 kV of the total VCVS's available output voltage (1.8 kV) thus essentially leaving ~ 800 V for the AC's positive half-cycles (mainly inductive). The second observation is that, in general and although a more pronounced effect at higher frequencies, there is a lag between the controller request and the VCVS's output voltage thus confirming the need for modelling the VCVS's transfer function. A third remark is that, in general, there is a visibly acceptable agreement between the model's predictions and the true obtained voltage for which a quantification is presented in Fig. ???. By performing the DFT of the output voltage and the predicted output voltage, one is able to plot the difference in their amplitude normalized to the amplitude of the output voltage, i.e. $[|DFT(V_k^{pred})| - |DFT(V_k)|] / |DFT(V_k)|$, as shown in Fig. ??(a) for each of the *scanned* frequencies. The immediate conclusion is that the model tends to overestimate the amplitude of an oscillatory output voltage. The exact cause of this effect is not yet fully understood at this time. By performing the cross-spectrum analysis of the output voltage and its prediction, one is able to estimate their phase delay normalized to the oscillation period and plot it as shown in Fig. ??(b). From this it can be seen that the delay in the estimation of the VCVS's output voltage is never larger than 10% of the period of oscillation.

5.3. CONTROLLER PERFORMANCE

In this section, the controller performance is succinctly assessed.

The same controller configurations and JET pulse data used in the previous section were also used herein and an example of relevant time traces obtained (within the 30 Hz time window) are shown in Fig. ??. It is not surprising to observe that, in Fig. ??(a), for the AC's negative half-cycle both the blue and green curves adequately match the reference whereas for the positive half-cycles that is not the case for any of the curves. The fact that the MPC configurations with $\mu = \xi < 10^{-3}$ are able to provide better control (closer to the target reference) than with the other configurations

can be seen from the voltage requests shown in Fig. ??(b). It is clear that these MPC configurations react faster than the others. This suggests that, for instance, the penalizing factors $\mu = \xi = 10^{-2}$ of the other MPC controller configuration severely limits the output bandwidth of the control action by restricting voltage request excursions in consecutive cycles. Regarding the PID-based control example, one must mention that its settings have not been optimized for the operation at any specific frequency thus its output bandwidth could certainly be improved by increasing K_p and eventually decreasing K_i however, and ultimately, no single set of parameters can ever be optimal for every frequency. The previous discussion is also confirmed by Fig. ??(c) where the value of the cost function J_k is shown. The fact that, for $\mu = \xi = 10^{-2}$, its baseline is at least one order of magnitude above the others and that its normalized excursions are smaller than the others, indicates that the cost function minimization process is dominated by the terms penalizing excursions in consecutive voltage requests. Finally Fig. ??(d) shows, with respect to the VCVS's operational domain (refer to Fig. 15), where the solutions of the constrained minimization of J_k lie. As expected, in the positive AC half-cycles the solutions lie mainly in the frontier of the operational domain whereas in the negative half-cycles they lie inside the operational domain (meaning that the constrained and unconstrained minimization solutions coincide).

In order to further quantify the performance of the different control schemes under scrutiny, the exact same study that was done to assess the VCVS model in the previous section was performed to assess the control performance. By performing the DFT of the target reference current and the estimated current obtained, one is able to plot the difference in their amplitude normalized to the amplitude of the reference current, i.e. $[|DFT(I_k^{est})| - |DFT(I_k^{ref})|]/|DFT(I_k^{ref})|$, as shown in Fig. ??(a) for each of the *scanned* frequencies. Whereas all MPC-based control schemes exhibit similar performance, the PID used herein substantially attenuates the amplitude of the oscillatory output current with respect to the reference. By performing the cross-spectrum analysis of the current reference and the estimated current, one is able to infer their phase delay normalized to the oscillation period and plot it as shown in Fig. ??(b). It can be seen that, with respect to the delay, on average the MPC-based schemes with $\mu = \xi \leq 10^{-3}$ (blue and green curves) behave very similarly. Likewise, both the MPC-based approach with $\mu = \xi = 10^{-2}$ and the PID-based control behave very similarly.

6. CONCLUSIONS AND FUTURE WORK

A real-time application for the control of external magnetic perturbations has been developed, implemented and integrated in the suite of real-time controllers in the JET tokamak. The main features of the system are its ease of use from the SL interface, mode-lock based adaptation and slew rate limitation of the current references, multiplexing of references between pre-programmed configuration and externally provided ones in real-time and the choice between PID-based or model-based control.

Although using hardware technology from the 1990s, which eventually became widely adopted

and a standard at JET, the system is not only capable of meeting its requirements but also of implementing advanced control schemes within a $500\mu\text{s}$ control cycle period. The use of the MARTE framework for real-time applications greatly eased the development and testing process by enabling development work in a non-real-time environment (Linux) and the deployment in the live environment with minimal code changes. The synchronization mechanism devised enables the application to run at a clock speed twice as fast as the central timing signal available in the VME crate while still providing accurate and consistent time stamping. The KF-based current estimation provides a good improvement (1/6 of the raw measurement's noise variance) in terms of SNR of the control variable. This is crucial as all control actions are ultimately driven by it. The mode-lock compensation and slew rate limitation for adjusting in real-time the current references are expected to anticipate and ameliorate events that would, if disregarded, ordinarily stop the entire experiment. In particular, both allow the resume of the originally planned operation conditions in case the offending event is successfully mitigated.

One of the major improvements of this system over its predecessors is the capability of receiving the current references, in real-time, from RTCC. This system allows the implementation of algorithms built upon Simulink-like control/signal processing blocks and used for experimental purposes. This allows the use of arbitrary sets of signals available in RTDN to steer the EFCC currents, in real-time, according to the goals and the evolving state of the plasma pulse. There is already an interest of exploring this feature for the real-time control of the ELM size/frequency crucial to reduce heat loads on to plasma facing components. Finally, the development of an MPC-based control aimed at the optimization of the controller performance especially in RFA experiments is expected to improve the quality and the overall bandwidth of the system. Although the VCVS model, used in the MPC-based control scheme, requires some re-tuning to prevent the over-estimating the VCVS's output voltage, the system has already demonstrated to provide the functionality required to incorporate external magnetic perturbations into plasma experiments at JET.

ACKNOWLEDGEMENTS

This work was supported by EURATOM and carried out within the framework of the European Fusion Development Agreement. IST activities also received financial support from "Fundação para a Ciência e Tecnologia" through project Pest-OE/SADG/LA0010/2011. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

REFERENCES

- [1] D. Alves, R. Felton, S. Jachmich, P. Lomas, P. McCullen, A. Neto, D. F. Valcárcel, G. Arnoux, P. Card, S. Devaux, A. Goodyear, D. Kinna, A. Stephen, and K.-D. Zastrow. Vessel thermal map real-time system for the jet tokamak. *Phys. Rev. ST Accel. Beams*, 15:054701, May 2012.

- [2] D. Alves, A.C. Neto, D.F. Valcarcel, R. Felton, J.M. Lopez, A. Barbalace, L. Boncagni, P. Card, G. De Tommasi, A. Goodyear, S. Jachmich, P.J. Lomas, F. Maviglia, P. McCullen, A. Murari, M. Rainford, C. Reux, F. Rimini, F. Sartori, A.V. Stephen, J. Vega, R. Vitelli, L. Zabeo, and K.-D. Zastrow. A new generation of real-time systems in the jet tokamak. In *Real Time Conference (RT), 2012 18th IEEE-NPSS*, pages 1–9, 2012.
- [3] D. Alves, R. Vitelli, L. Zaccarian, L. Zabeo, A. Neto, F. Sartori, P. McCullen, and P. Card. The new error field correction coil controller system in the joint european torus tokamak. *Fusion Engineering and Design*, 86(68):1034 – 1038, 2011.
- [4] D.P. Atherton and S. Majhi. Limitations of pid controllers. In *American Control Conference, 1999. Proceedings of the 1999*, volume 6, pages 3843–3847 vol.6, 1999.
- [5] I. Barlow, M. Bigi, J. Bird, G. Bonizzoni, R. Buttery, R. Clay, M. De Benedetti, T. Dobbing, T. Gallagher, G. Gervasini, G. Gittini, T.C. Hender, S. Hotchin, J. Last, E. Lazzaro, M. Lennholm, S. Peruzzo, V. Riccardo, L. Rossi, A. Santagiustina, D. Starkey, A. Stevens, M. Tabellini, A. Tanga, A. Terrington, M. Vincent, P. Watkins, M. Way, and M. Williams. The error field correction coils on the jet machine. *Fusion Engineering and Design*, 5859(0):189 – 193, 2001.
- [6] T. Bellizio, G. De Tommasi, P. McCullen, A.C. Neto, F. Piccolo, F. Sartori, R. Vitelli, and L. Zabeo. The software architecture of the new vertical-stabilization system for the jet tokamak. *Plasma Science, IEEE Transactions on*, 38(9):2465 –2473, sept. 2010.
- [7] L. Boncagni, Y. Sadeghi, D. Carnevale, G. Mazzitelli, A. Neto, D. Pucci, F. Sartori, F. Piesco, S. Sinibaldi, V. Vitale, R. Vitelli, L. Zaccarian, S. Monaco, and G. Zamborlini. First steps in the ftu migration towards a modular and distributed real-time control architecture based on marte. *Nuclear Science, IEEE Transactions on*, 58(4):1778 –1783, aug. 2011.
- [8] P. Carvalho, P. Duarte, T. Pereira, R. Coelho, C. Silva, and H. Fernandes. Real-time tomography system at isttok. *Nuclear Science, IEEE Transactions on*, 58(4):1427 –1432, aug. 2011.
- [9] D. Chiron, T. Bonicelli, M. Huart, M. Garribba, P.-L. Mondino, and P. Noll. Power supplies for the stabilisation of plasma vertical position: recent upgrades and future development. In *Fusion Engineering, 1991. Proceedings., 14th IEEE/NPSS Symposium on*, pages 513–516 vol.1, 1991.
- [10] G. De Tommasi, D. Alves, T. Bellizio, R. Felton, A. Neto, F. Sartori, R. Vitelli, L. Zabeo, R. Albanese, G. Ambrosino, and P. Lomas. Real-time systems in tokamak devices. a case study: The jet tokamak. *Nuclear Science, IEEE Transactions on*, 58(4):1420 –1426, aug. 2011.

- [11] T. E. Evans, R. A. Moyer, P. R. Thomas, J. G. Watkins, T. H. Osborne, J. A. Boedo, E. J. Doyle, M. E. Fenstermacher, K. H. Finken, R. J. Groebner, M. Groth, J. H. Harris, R. J. La Haye, C. J. Lasnier, S. Masuzaki, N. Ohyaabu, D. G. Pretty, T. L. Rhodes, H. Reimerdes, D. L. Rudakov, M. J. Schaffer, G. Wang, and L. Zeng. Suppression of large edge-localized modes in high-confinement diii-d plasmas with a stochastic magnetic boundary. *Phys. Rev. Lett.*, 92:235003, Jun 2004.
- [12] R. Felton, K. Blackler, S. Dorling, A. Goodyear, O. Hemming, P. Knight, M. Lennholm, F. Milani, F. Sartori, and I. Young. Real-time plasma control at jet using an atm network. In *Real Time Conference, 1999. Santa Fe 1999. 11th IEEE NPSS*, pages 175–181, 1999.
- [13] A.M. Fernandes, R.C. Pereira, A. Neto, D.F. Valcarcel, D. Alves, J. Sousa, B.B. Carvalho, V. Kiptily, B. Syme, P. Blanchard, A. Murari, C.M.B.A. Correia, and C.A.F. Varandas. Real-time processing system for the jet hard x-ray and gamma-ray profile monitor enhancement. In *Real Time Conference (RT), 2012 18th IEEE-NPSS*, pages 1–7, 2012.
- [14] Gene F Franklin, J David Powell, and Abbas Emami-Naeini. *Feedback control of dynamic systems*, volume 320. Addison-Wesley Reading, MA, 1991.
- [15] B. Goncalves, J. Sousa, B. Carvalho, A.P. Rodrigues, M. Correia, A. Batista, J. Vega, M. Ruiz, J.M. Lopez, R.C. Rojo, A. Wallander, N. Utzel, A. Neto, D. Alves, and D. Valcarcel. Engineering design of iter prototype fast plant system controller. *Nuclear Science, IEEE Transactions on*, 58(4):1439–1446, 2011.
- [16] Holec. *Final Design Report of Poloidal Radial Field Amplifiers*, October 1982.
- [17] J How et al. The JET pulse termination network. *Fusion Technology*, 2:1421–1426, 1980.
- [18] R. E. Kalman and R. S. Bucy. New results in linear filtering and prediction theory. *Transactions of the ASME. Series D, Journal of Basic Engineering*, 83:95–107, 1961.
- [19] Rudolph Emil Kalman. A New Approach to Linear Filtering and Prediction Problems. *Transactions of the ASME—Journal of Basic Engineering*, 82(Series D):35–45, 1960.
- [20] K Kamiya, N Asakura, J Boedo, T Eich, G Federici, M Fenstermacher, K Finken, A Herrmann, J Terry, A kirk, B Koch, A Loarte, R Maingi, R Maqueda, E Nardon, N Oyama, and R Sartori. Edge localized modes: recent experimental findings and related issues. *Plasma Physics and Controlled Fusion*, 49(7):S43, 2007.
- [21] QA King, HEO Brelén, and JET Joint Undertaking. An experimental automatic control facility at JET. *JOINT EUROPEAN TORUS-PUBLICATIONS-JET P*, 1998.
- [22] J.M. Lopez, J. Vega, D. Alves, S. Dormido-Canto, A. Murari, J.M. Ramirez, R. Felton, M. Ruiz, and G. de Arcas. Implementation of the disruption predictor apodis in jet real time

- network using the marte framework. In *Real Time Conference (RT), 2012 18th IEEE-NPSS*, pages 1–4, 2012.
- [23] G. Manduchi, A. Luchetta, A. Soppelsa, and C. Taliercio. The new feedback control system of rfx-mod based on the marte real-time framework. In *Real Time Conference (RT), 2012 18th IEEE-NPSS*, pages 1–5, 2012.
- [24] A. Masiello, M. Bigi, R. Buttery, and V. Riccardo. Finite element electromagnetic analysis of the jet error field correction coils. *Fusion Engineering and Design*, 63-64:455 – 460, 2002.
- [25] Manfred Morari and Jay H. Lee. Model predictive control: past, present and future. *Computers & Chemical Engineering*, 23(45):667 – 682, 1999.
- [26] A.C. Neto, D. Alves, L. Boncagni, P.J. Carvalho, D.F. Valcarcel, A. Barbalace, G. De Tommasi, H. Fernandes, F. Sartori, E. Vitale, R. Vitelli, and L. Zabeo. A survey of recent marte based systems. *Nuclear Science, IEEE Transactions on*, 58(4):1482–1489, 2011.
- [27] A.C. Neto, F. Sartori, F. Piccolo, R. Vitelli, G. De Tommasi, L. Zabeo, A. Barbalace, H. Fernandes, D.F. Valcarcel, and A. J N Batista. Marte: A multiplatform real-time framework. *Nuclear Science, IEEE Transactions on*, 57(2):479–486, 2010.
- [28] VD Pustovitov. Resonant field amplification near the rwm stability boundary in a tokamak. *Journal of Experimental and Theoretical Physics Letters*, 78(5):281–284, 2003.
- [29] LM Sonneborn and FS Van Vleck. The bang-bang principle for linear control systems. *Journal of the Society for Industrial & Applied Mathematics, Series A: Control*, 2(2):151–159, 1964.
- [30] Adam Stephen, G Arnoux, T Budd, P Card, R Felton, A Goodyear, J Harling, D Kinna, P Lomas, P McCullen, and JET EFDA contributors. Centralised coordinated control to protect the jet iter-like wall. October 2011.
- [31] Andrew S. Tanenbaum. *Modern Operating Systems*. Prentice Hall Press, third edition, December 2007.
- [32] D.F. Valcarcel, D. Alves, P. Card, B.B. Carvalho, S. Devaux, R. Felton, A. Goodyear, P.J. Lomas, F. Maviglia, P. McCullen, C. Reux, F. Rimini, A. Stephen, L. Zabeo, and K-D. Zastrow. The JET real-time plasma-wall load monitoring system. *Fusion Engineering and Design*, 2013.
- [33] D.F. Valcarcel, D. Alves, A. Neto, C. Reux, B.B. Carvalho, R. Felton, P.J. Lomas, J. Sousa, and L. Zabeo. Parallel task management library for marte. In *Real Time Conference (RT), 2012 18th IEEE-NPSS*, pages 1–7, 2012.

- [34] D.F. Valcarcel, A. Neto, I.S. Carvalho, B.B. Carvalho, H. Fernandes, J. Sousa, F. Janky, J. Havlicek, R. Beno, J. Horacek, M. Hron, and R. Panek. The compass tokamak plasma control software performance. *Nuclear Science, IEEE Transactions on*, 58(4):1490–1496, aug. 2011.
- [35] H. van der Beken, B.J. Green, C.A. Steed, J. W. Farthing, P. A. McCullen, and J. How. Level 1 software at jet: a global tool for physics operation. In *Fusion Engineering, 1989. Proceedings., IEEE Thirteenth Symposium on*, pages 201–204 vol.1, 1989.
- [36] Sangwon Yun, Woongryol Lee, Taegu Lee, Mikyung Park, Sangil Lee, Andr C. Neto, Anders Wallander, and Young-Kuk Kim. Evaluating performance of {MARTe} as a real-time framework for feed-back control system at tokamak device. *Fusion Engineering and Design*, 88(68):1323 – 1326, 2013. *Proceedings of the 27th Symposium On Fusion Technology (SOFT-27); Lige, Belgium, September 24-28, 2012*.




Figure 1: The error field correction coils at JET.

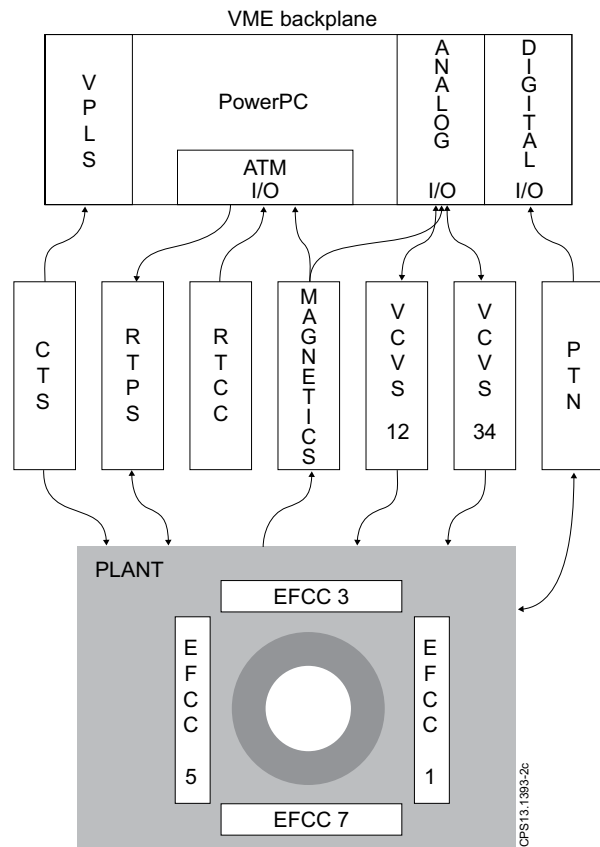


Figure 2: EFCC current controller in the JET context.

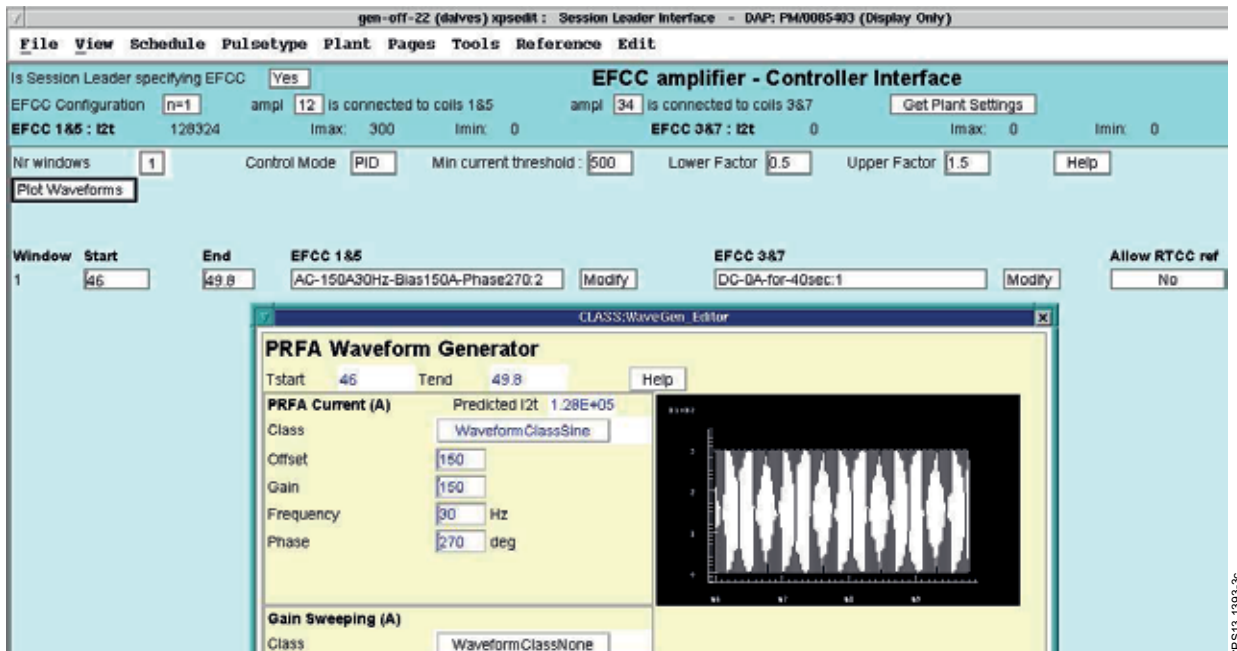


Figure 3: EFCC current controller Level-1 interface.




Figure 4: (a) Histogram of the application's measured cycle time and (b) Probability of having a processing cycle with an absolute jitter larger than a given percentage of the application's nominal cycle time.

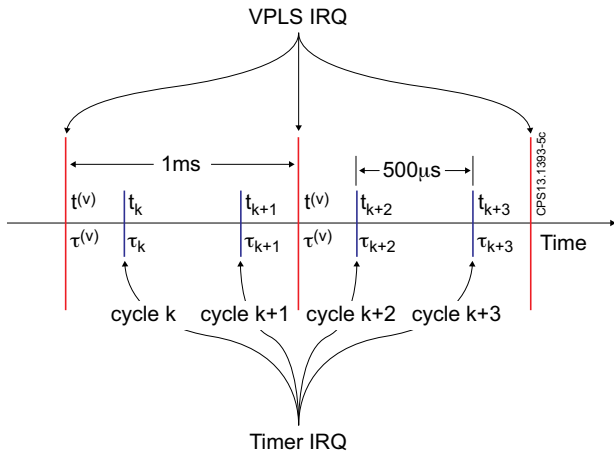


Figure 5: Time stamping and synchronisation mechanism.

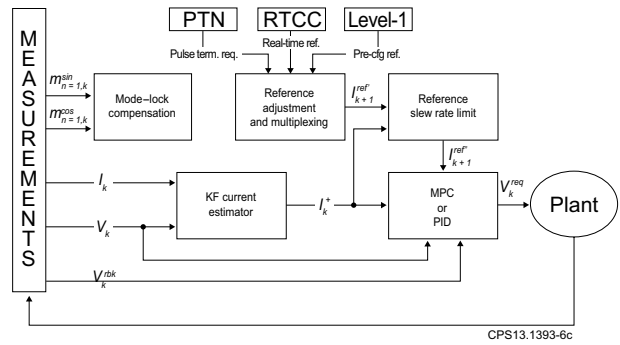


Figure 6: Control algorithm simplified data flow diagram.




Figure 7: Time trace - (blue) Current measurement and (red) KF current estimate.




Figure 8: KF frequency response - (a) Normalized amplitude error and (b) Delay normalized to the period.




Figure 9: Saddle coils used in each of the 4 equally spaced octants to calculate the mode-lock signal.




Figure 10: Mode-lock compensation (color legend in the figure) - continuous line represents the raw signals and dots represents the compensated signals.




Figure 11: Automatic current reference adjustment based on mode-lock compensation.

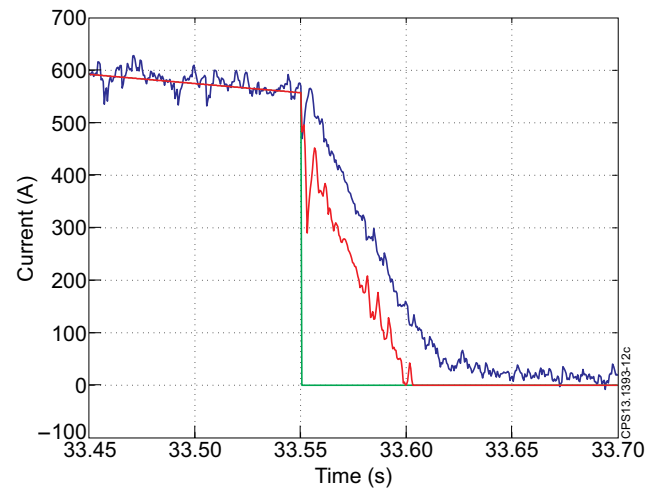


Figure 12: Slew rate - (blue) estimated current; (green) adjusted reference and (red) slope limited reference.




Figure 13: Pre-configured voltage control - (a) Readback of the voltage request acknowledged by the VCVS and (b) Current in the EFCC circuit.

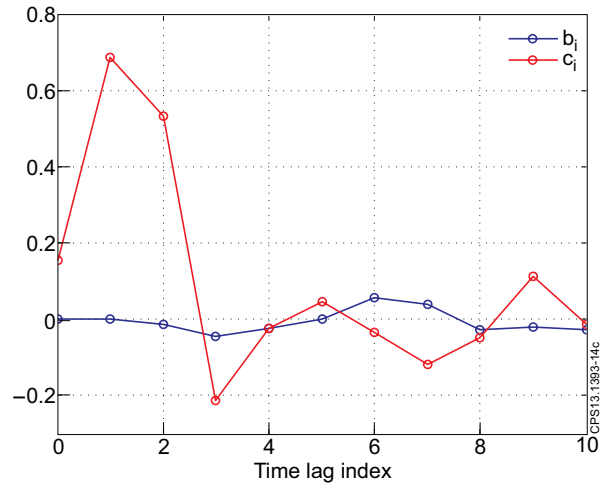


Figure 14: VCVS linear model coefficients.

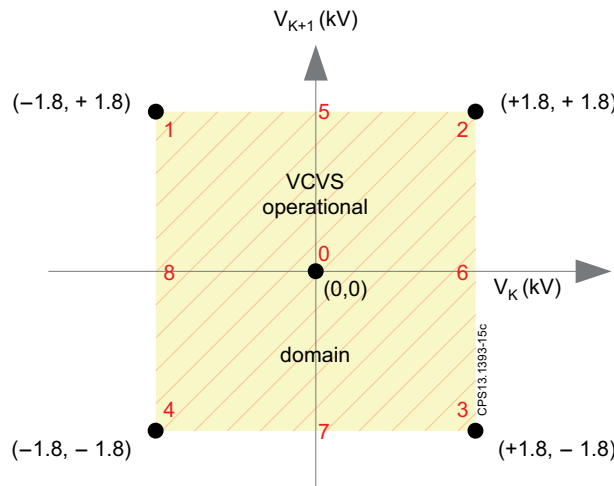


Figure 15: VCVS linear operational domain - red indexes are labels to identify the regions in the frontier of the operational domain (1–8) and inside the operational domain (0).




Figure 17: Pre-configured current reference.

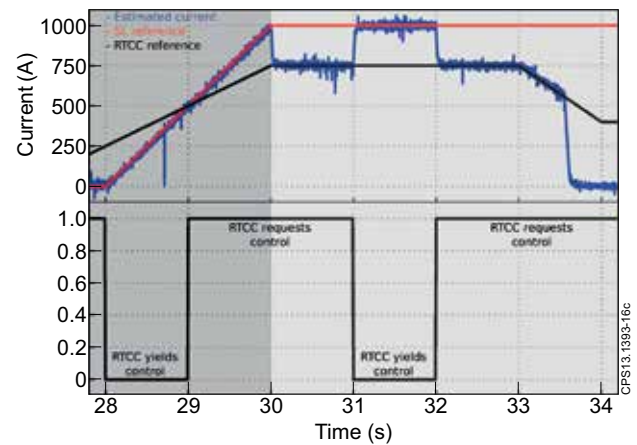


Figure 16: Commissioning of the reference multiplexing logic.




Figure 18: VCVS model assessment - (blue) VCVS output; (green) Controller request; (red) Predicted VCVS output - (a) MPC with $\mu = \xi = 10^{-4}$, (b) MPC with $\mu = \xi = 10^{-3}$, (c) MPC with $\mu = \xi = 10^{-2}$, (d) PID with $K_p = 5$, $K_i = 20$, $K_d = 0$.




Figure 19: VCVS model assessment - (blue) MPC with $\mu = \xi = 10^{-4}$, (green) MPC with $\mu = \xi = 10^{-3}$, (red) MPC with $\mu = \xi = 10^{-2}$, (cyan) PID with $K_p = 5$, $K_i = 20$, $K_d = 0$ - (a) Normalized amplitude error; (b) Delay normalized to period




Figure 20: Control performance - (blue) MPC with $\mu = \xi = 10^{-4}$, (green) MPC with $\mu = \xi = 10^{-3}$, (red) MPC with $\mu = \xi = 10^{-2}$, (cyan) PID with $K_p = 5$, $K_i = 20$, $K_d = 0$ and (black) reference - (a) Estimated current; (b) Requested voltage; (c) Cost function; (d) VCVS operational space index (refer to 15).




Figure 21: Control performance - (blue) MPC with $\mu = \xi = 10^{-4}$, (green) MPC with $\mu = \xi = 10^{-3}$, (red) MPC with $\mu = \xi = 10^{-2}$, (cyan) PID with $K_p = 5$, $K_i = 20$, $K_d = 0$ - (a) Normalized amplitude error; (b) Delay normalized to period.