Implementation of a Novel Real-Time Controller for the Detection and Tracking of Magneto-Hydrodynamic Instabilities on the JET Tokamak

Duccio TESTA, Hervé CARFANTAN[1)], Alex GOODYEAR[2)] and JET-EFDA contributors[*,3)]

Ecole Polytechnique Fédérale de Lausanne (EPFL), Centre de Recherches en Physique des Plasmas (CRPP), Association EURATOM – Confédération Suisse, Lausanne, Switzerland

[1)]*Université de Toulouse, Institut de Recherche en Astrophysique et Planétologie, Centre National pour la Recherche Scientifique, Toulouse, France*

[2)]*Culham Centre for Fusion Energy, Culham Science Centre, Abingdon, United Kingdom*

[3)]*JET-EFDA, Culham Science Centre, OX14 3DB, Abingdon, UK*

(Received 8 October 2012 / Accepted 10 September 2013)

In this work we present the technical implementation of a digital VERSA Module Eurocard (VMEbus) system used to detect and track, in real-time, magneto-hydrodynamic instabilities on the JET tokamak. This VMEbus system runs on a 1 ms clock cycle and performs the unsupervised detection and real-time tracking of the individual components in a multi-harmonic spectrum of coherent electro-magnetic instabilities, actively driven by a set of in-vessel antennas. Its main real-time output signals are the frequency, amplitude, toroidal mode number and damping rate of such modes. Moreover, this controller also provides some of the protection and control tools for the antenna system, such as the reference for the voltage and current control waveforms, and a trip signal related to the shorted-turn protection of the antennas. Current applications of this novel controller focus on the measurement of the damping rate of Alfvén Eigenmodes with different toroidal mode numbers. The successful technical implementation and scientific exploitation of this innovative VMEbus system opens possibilities for the real-time detection and the ensuing control of electro-magnetic instabilities in other present and future fusion devices.

Keywords: MHD modes, real-time detection, real-time control

DOI: 10.1585/pfr.9.1405003

1. Introduction

Analysis of electro-magnetic fluctuations is important for understanding and controlling the magneto-hydrodynamic (MHD) stability of magnetically confined thermonuclear plasmas. Specifically, the problem of unsupervised real-time detection of MHD modes has now become one of the most important aspects for machine protection and control of plasma discharges in thermonuclear fusion experiments. The method routinely used for this analysis involves sampling a (usually rather) small set of input signals, such as measurements of magnetic, density and temperature fluctuations, which in most cases are unevenly sampled in the spatial domain. Appropriate processing of such a set of input data facilitates the detection of the different components in a multi-harmonics spectrum. Furthermore, when the data contains some spatial periodicities, these can be readily used to enhance or eliminate the detection of certain components. A real-time algorithm can then generate a global alarm that is sent to the plant. Under certain specified and pre-determined operational conditions, this may then trigger a feedback control mechanism. For some examples of these activities, the Readers are referred to Chapter 3 and Chapter 7 and references therein in [1] and Chapter 2 and references therein in [2].

One drawback of most of the current MHD detection and control methods is that they can only detect modes when they have become *unstable* (with a growth rate $\gamma_{\mathrm{GROWTH}} > 0$), i.e. when they may have already had some detrimental effect on the actual plasma operation and performance. On the other hand, an alternative and innovative method is in use on the JET tokamak. This diagnostic technique combines the active excitation (via a set of in-vessel antennas) of magnetic field perturbations which have a very small amplitude at the plasma edge (maximum intensity $|\delta B_{\mathrm{DRIVEN}}| < 100\,\mathrm{mG}$, i.e. typically $\sim 10^5$ times smaller than the value of the toroidal magnetic field in JET, $B_{\mathrm{TOR}} \sim 1\,\mathrm{T}$ to $B_{\mathrm{TOR}} \sim 4\,\mathrm{T}$) with synchronous real-time detection of the resonant plasma response to such antenna-driven perturbations. This method then allows detecting

author's e-mail: duccio.testa@epfl.ch

[*] Appendix of F. Romanelli, "*Overview of JET Results*", Nuclear Fusion **51** (2011), 094008 (Proceedings of the 23rd IAEA Fusion Energy Conference 2010, Daejeon, Korea).

MHD modes when they are still *stable* (with a damping rate $\gamma_{\mathrm{DAMP}} > 0$), i.e. before they could have affected the discharge, which is evidently a much more satisfying situation for plasma control and machine protection.

This JET diagnostic system is the so-called Alfvén Eigenmodes (AEs) Active Diagnostic (AEAD) [3], as its original (and still currently predominant) aim is to detect AEs [4–7]. Its real-time controller, the Alfvén Eigenmodes Local Manager (AELM), constitutes one essential and furthermore worldwide unique component of the JET Real Time Data Network (RTDN). The measurement of the mode characteristics, such as the frequency, the amplitude, the toroidal mode number and the damping rate, are obtained in real-time through calculations performed on a sub-millisecond time scale by the AELM software. This data is then passed to the Real Time Signal Server (RTSS) [8], which allows implementing a real-time control and feedback system for the modes detected with the AELM by measuring the distance from the marginal stability limit $\gamma/\omega = 0$, and calling for a reaction of the plant when the plasma is approaching the limit during the discharge.

This paper presents the technical implementation of the AELM hardware and software infrastructure, with a specific focus on the application to the sub-millisecond detection, discrimination and real-time tracking of the individual components in the multi-harmonic, frequency-degenerate spectrum of stable AEs which are excited in the JET tokamak by an array of in-vessel antennas used for MHD diagnostic purposes. This paper is organized as follows. Section 2 gives a short overview of the AEAD system. In Sec. 3 we present a complete technical overview of the AELM system used in JET, with particular attention to its real-time plant control and data analysis hardware and software. Section 4 focuses on the setup of the real-time mode detection and tracking algorithm used in the AELM. Section 5 then shows some illustrative examples of the detection and discrimination between the different toroidal components in a multi-harmonic spectrum of stable AEs. In Sec. 6 we present some statistical analysis of the accuracy of the different algorithms used by the AELM for the real-time mode detection and tracking. Then, in Sec. 7 we summarize our results and give an outlook towards future work. Appendix A reviews the mathematical foundation of the *Sparse Representation* method and of the *SparSpec* code, which is used to extract, both in real-time and post-pulse, the frequency, amplitude, toroidal mode-number and damping rate of the modes actively driven by the AEAD system, its optimization for implementation within the AELM, and an assessment of its real-time numerical performance. Additionally, some applications of the *SparSpec* algorithm other than those related to the JET AELM system are also presented, always focussing on the analysis of magnetic fluctuation in JET tokamak plasmas. Finally, the supplemental information contained in Appendix B provides more technical details on the implementation of a number of functionalities of the AELM hardware and software, so as to facilitate practical understanding of the software and hardware infrastructure required to design, implement and routinely operate such a real-time controller.

2. The Alfvén Eigenmodes Active Diagnostic System in use at JET

The *MHD spectroscopy* technique is a diagnostic tool that uses global waves naturally supported by the plasma with the aim of measuring the parameters that determine their dispersion relation, their absorption and propagation, and the damping and growth of these waves [9, 10]. A simple active method to drive and detect low amplitude modes in the plasma was pioneered and used in many different plasma conditions in the JET tokamak [3]. This is the so-called Alfvén Eigenmodes Active Diagnostic (AEAD) system, as its main aim was (and currently still is) that of driving and detecting plasma Eigenmodes in the Alfvén frequency range. Technical details on the AEAD system used at JET can be found in [3, 11–15] and references therein, and here we only present a brief overview.

As schematically shown in Fig. 1, the two main components of the AEAD system are:

1. the AE exciter, which is built upon a function generator and a single broadband high-power amplifier operating in class-AB mode; the amplifier is then connected to a set of up to eight in-vessel antennas via distribution and isolation transformers and a ∼200 m long transmission line;
2. the AE receiver, which is built upon synchronous detection units and real-time data analysis.

The AEAD exciter operates in the 10 kHz to 500 kHz frequency range, with maximum capabilities of 5 kW/1 kV/15 A-peak delivered into a 50 Ω load. This exciter produces a very small magnetic perturbation at the plasma edge, with maximum intensity of the order of $\max(|\delta B_{\mathrm{DRIVEN}}|) \sim 0.1\,\mathrm{G}$, which is $\sim 10^5$ times smaller than the typical value of the toroidal magnetic field in JET, $B_{\mathrm{TOR}} \sim 1\,\mathrm{T}$ to $B_{\mathrm{TOR}} \sim 4\,\mathrm{T}$. The AEAD receiver collects signals from a set of in-vessel detectors for electro-magnetic fluctuations, such as magnetic pick-up coils, electron cyclotron emission and reflectometry measurements. This receiver is also connected to the AELM to allow for the real-time detection and tracking of the plasma resonant response to the antenna-driven magnetic field perturbation.

The real-time use of the AEAD system is facilitated by the process of *synchronous detection*, which provides the capability of measuring only the plasma response at the frequency corresponding to the antenna excitation, i.e. its *synchronous* component. Figure 2 shows a schematic view of the implementation of this concept. Synchronous detection allows reducing the required bandwidth of the data

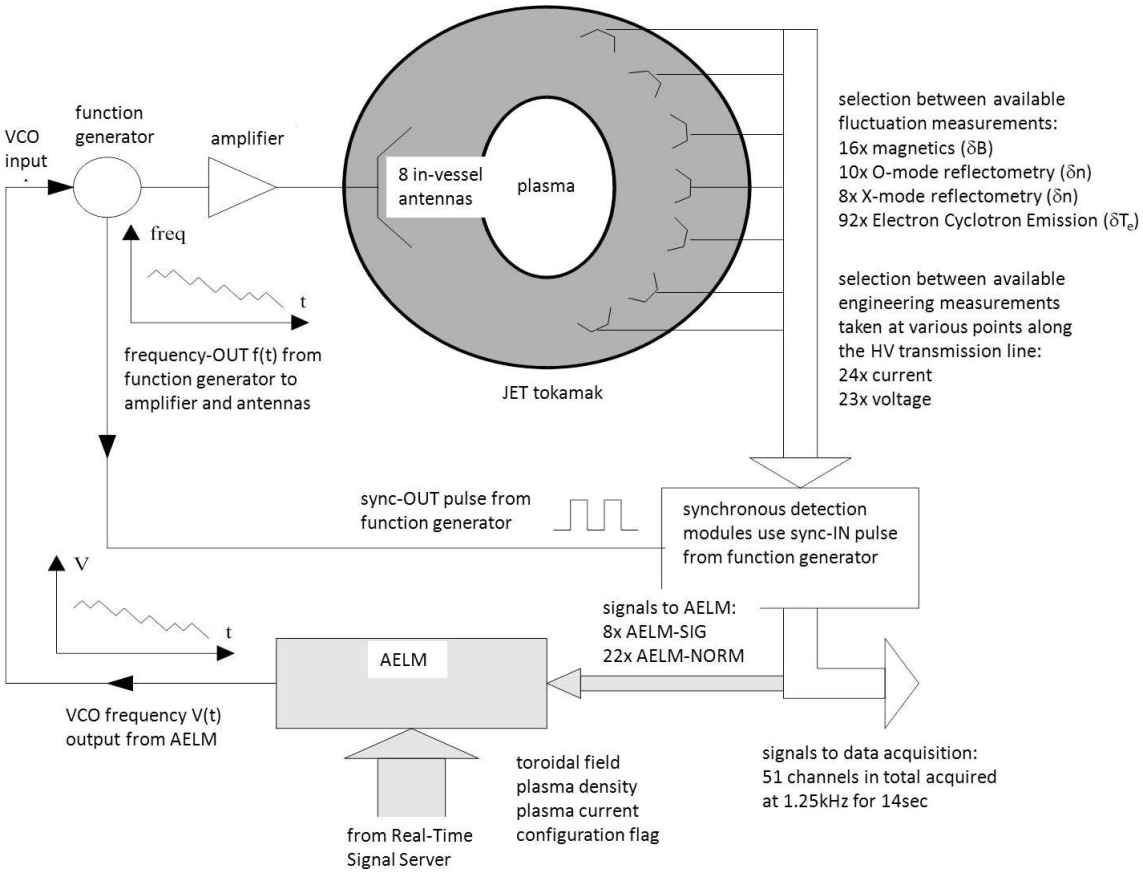

Fig. 1 A schematic overview of the AEAD system in JET. The toroidal field, plasma density and plasma current are retrieved from the RTSS and can be used by the AELM to compute in real-time an initial guess for the antenna excitation frequency freq(t). The AELM then converts this value to a time-dependent voltage $V(t)$ and sends it as a Voltage Controlled Oscillator (VCO) input signal to the function generator which, in turn, converts it back into a frequency freq(t). This signal then drives a 5 kW amplifier connected to up to eight in-vessel antennas via isolation and distribution transformers and a ~200 m long transmission line (not shown). It is also used to provide the pulsed synch-OUT/IN reference for the synchronous detection units. These modules collect a selection of fluctuation and engineering measurements, some of which are also sent back to the AELM for feedback control of the AEAD plant and for mode detection and tracking.

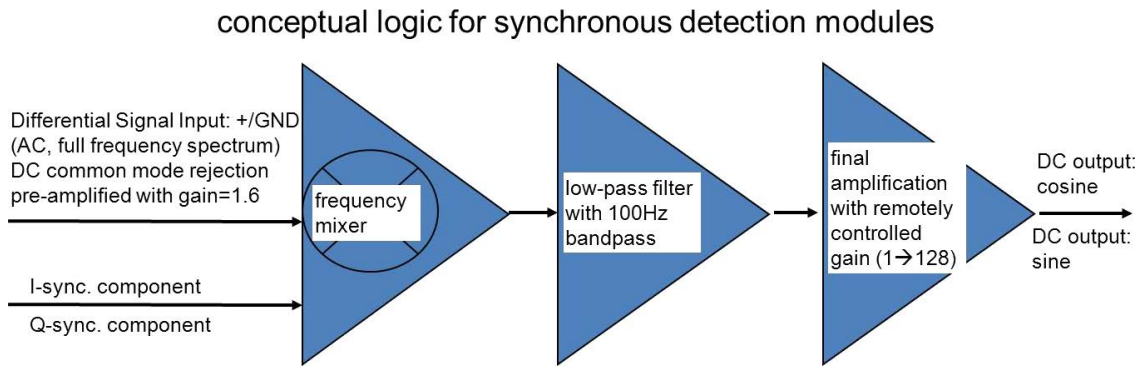

Fig. 2 A schematic conceptual design of the synchronous detection modules. This hardware is conceptually based upon removing the DC common mode from the input AC differential signal, and pre-amplifying this signal with a fixed 1.6 gain. The resulting signal is then separately mixed with the I and the Q synchronous components, and low-pass filtered to extract only the component in the input signal which is at the desired (i.e. synchronous) frequency. Finally, a remotely controlled gain is applied, and the DC output cosine ($\propto$ I-synch) and sine ($\propto$ Q-synch) components are obtained.

acquisition system and removes the need for computationally expensive Fast Fourier Transform (FFT) algorithms to obtain the required frequency component. For the specific case of the AEAD system, we use a 1.25 kHz sampling rate for detecting modes whose frequency can reach 500 kHz, whereas real-time FFT algorithms based on the Nyquist criterion would require a 1 MHz sampling rate, i.e. needing an 800 times larger bandwidth and storage capabilities for the data acquisition. Synchronous detection also allows removing all undesired frequency components, hence dramatically improving the signal-to-noise (S/N) ratio. The AEAD hardware for synchronous detection works conceptually by applying a frequency mixer with the synchronous in-phase (I) and quadrature (Q) components to the incoming differential signal, and then applying a low-pass filter with a < 100 Hz bandwidth to generate the output [cosine (I), sine (Q)] DC components. The final output bandwidth of the AEAD synchronous detection system is then 500 kHz.

3. The Alfvén Eigenmode Local Manager

One essential and worldwide unique component of the AEAD system is the AELM. The AELM is a VERSA Module Eurocard (VMEbus) plant control system, used to monitor and control both the engineering and the scientific aspects of the AEAD plant operation. The AELM operates in real time using a 1 kHz clock derived from the Composite Time and Trigger System (CTTS), which is used to synchronise the JET diagnostic and control systems during an experiment. This 1 kHz clock, i.e. a 1 ms clock cycle, is specified by the requirement for integrating the AELM within the CTTS system and translates into severe data acquisition and processing constraints for the AELM. Synchronous detection is therefore essential for the AELM real-time application, as only a limited amount of processing can be performed within this 1 ms clock limit. Furthermore, as the protocols for input/output data communication with the other elements of the JET plant take around 100 μs to 150 μs to complete, all AELM calculations need to be concluded in less than 850 μs for every 1 ms clock cycle, otherwise an alarm is raised and the AELM will stop processing and return to a safe state. In this respect, the AELM is a *hard* real-time embedded system: the results of missing a deadline are classed as a failure. For JET operation this results in, at best, missing data and/or a wasted experiment, i.e. avoiding the worst case, which would be damage to the AEAD or to the JET machine itself. This is contrary to a *soft* real-time embedded system, which would simply produce a reduced quality of service, such as a Graphical User Interface suffering sluggish mouse response.

The AELM receives the feedback inputs from the AEAD system via analogue-to-digital converter (ADC) modules called VAJ1. The input data consists of 2×8 signals (labelled AELM-SIG channels, with each signal having the real and imaginary component acquired separately) used for mode detection and 2×22 engineering signals (labelled AELM-NORM channels, again each signal having the real and imaginary component acquired separately) used for the AEAD plant control and the normalization of the antenna-driven excitation spectrum. The AELM-SIG channels can be chosen by manually patching to the AELM inputs a suitable selection of measurements from: magnetic pick-up coils mounted on the low- and high-field side vessel walls (16 signals available in total for selection), density fluctuation measurements obtained from X-mode (8 signals available in total for selection) and O-mode reflectometry (10 signals available in total for selection), and electron temperature fluctuation measurements (92 signals available in total for selection) obtained from electron cyclotron emission spectroscopy. The AELM-NORM channels are obtained from current (24 signals available in total for selection) and voltage (23 signals available in total for selection) measurements taken at different points along the transmission line. Two frequency deviation signals from the Ion Cyclotron Resonance Frequency (ICRF) antennas are also supplied as input to the AELM, and can be used for the RF beat-wave control operation [16].

Finally, the RTDN system supplies the AELM with some of the plasma parameters measured in real-time. These are the toroidal magnetic field and the toroidal plasma current, the line-integrated plasma density obtained along a vertical chord passing through the plasma centre, and a flag indicating whether the running plasma configuration is a limiter or an X-point one. These signals can be used to compute in real-time the driving frequency for the in-vessel antennas, for instance the AE frequency. Figure 3 shows a schematic diagram of how the AELM fits into the RTDN.

The AELM produces a number of outputs which are sent back to the AEAD plant through digital-to-analogue converter (DAC) modules, again using the VAJ1 units. The first main output provides the overall timing and data acquisition control sequence for the AEAD system, which includes the processing of alarms from the JET plant via the CTTS (such as a failed breakdown, a disruption or any other signals indicating a soft- or hard-stop of the plasma discharge) and from the AEAD plant itself (the shorted-turn (S/T) trip). The second main output is the antenna driving frequency, called F_{REF}: this is a modulated DC voltage signal $V(t)$, with values in the range $V(t) = 0$ V to $V(t) = 10$ V, connected to the Voltage Controlled Oscillator (VCO) input of a function generator. The VCO frequency output is feedback controlled by the AELM so that tracking of the mode can occur when a resonant plasma response to the antenna-driven perturbation is detected in real-time. The antenna driving frequency is usually selected to match

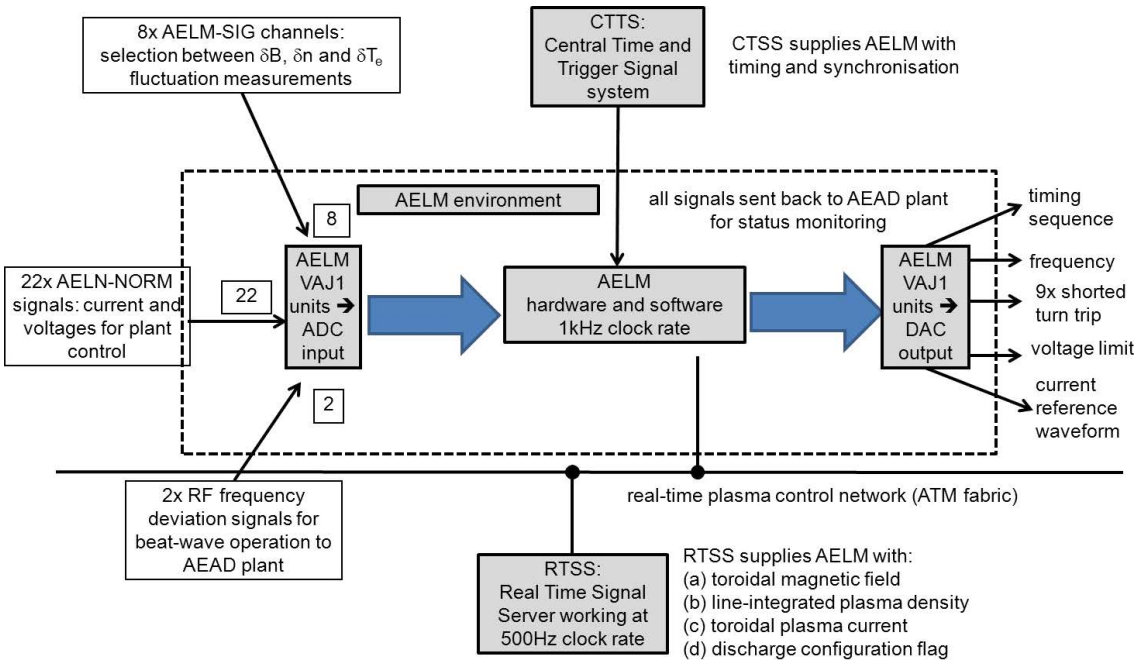

Fig. 3 A schematic diagram of the AELM hardware, illustrating its interfaces with the Real Time Signal Server (RTSS) and the Central Time and Trigger Signal System (CTSS), and the input and output connections to the AEAD plant via the VAJ1 cards and the Ion Cyclotron Resonance Frequency (ICRF) heating system for real-time control of the RF beat-wave operation. The black-dotted line separates the AELM-proper environment, which includes the VAJ1 cards, from the "outside" world.

the AE frequency, which is given by:

$$F_{\mathrm{REF}}\,[\mathrm{kHz}] = (\textit{multiplier}) \times \frac{B_{\mathrm{TOR}}}{4\pi R_{\mathrm{RES}} q_{\mathrm{RES}} \sqrt{m_{\mathrm{p}} \sum_i n_i A_i}}. \quad (1)$$

Here B_{TOR} is the toroidal magnetic field, $\mathrm{n_i}$ and $\mathrm{A_i}$ are the density and atomic mass of all ion species, m_{p} is the proton mass, $q_{\mathrm{RES}}(r_{\mathrm{RES}}) = (2m + 1)/2n$ is the value of the safety factor at the mode resonant position $R_{\mathrm{RES}} = R_0 + r_{\mathrm{RES}}$, where R_0 is the magnetic axis position, r the minor radius coordinate, and m and n are the AE poloidal and toroidal mode numbers, respectively. The quantity *multiplier* defines which class of AEs is being investigated: *multiplier* = 1 is used for Toroidal AEs (TAEs) and *multiplier* = 2 for ellipticity-induced AEs.

The AELM also generates two other outputs without feedback, which are sent to the AEAD control electronics. The first one consists of an amplitude waveform used as the reference current (I_{REF}) in an Active Gain Control (AGC) feedback amplifier loop. This waveform defines the target nominal antenna current that needs to be achieved during the discharge, generally set at a constant value (but with a ramp-up and ramp-down at the start and end of the pulse). The I_{REF} signal is a DC voltage ranging from 0 V to 10 V corresponding to a setting of 0 % to 100 % with respect to the nominal maximum current set in the AEAD hardware. The 100 % I_{REF} setting corresponds approximately to an antenna current I_{ANT} = 15 A-peak and to an amplifier current I_{TOT} = 25 A-peak, both values taken at 200 kHz. The second output, also sent to the AGC, defines the maximum permissible voltage (V_{LIM}) allowed on the antennas, beyond which the AGC gain will be reduced. This is a DC voltage ranging from 0 V to 10 V corresponding to a maximum antenna voltage V_{ANT} = 2 kV-peak at 200 kHz.

Finally, the AELM produces the shorted-turn (S/T) signal for each individual antenna, and a global S/T trip alarm for protection of the plant, which is raised if any of the S/T trips for the individual antennas becomes active. The global S/T trip alarm is generated in real-time via the AELM by computing the change in the antenna impedance (Z_{ANT}), i.e. $\mathrm{d}(V_{\mathrm{ANT}}/I_{\mathrm{ANT}})/\mathrm{d}t = \mathrm{d}Z_{\mathrm{ANT}}/\mathrm{d}t$, V_{ANT} and I_{ANT} being the antenna voltage and current, respectively. If a variation in Z_{ANT} occurs too rapidly in the time/frequency domain and/or is too large in its absolute value in any of the antennas, then it is considered that this variation is not due to the resonant plasma response to the antenna-driven perturbation nor to an electrical resonance in the AEAD system. Therefore a global S/T trip is generated to stop power being delivered to all the in-vessel antennas for at least 20 ms or as long as the alarm exists, because such a fast impedance change is suggestive of a short circuit between one (or more) of the antenna turn(s) and the vessel potential. Note that these S/T trips have only occurred at the time of Edge Localized Modes (ELMs), which are MHD events that often peel off some portion of the edge plasma and dump the associated current onto the wall. These events create a very time- and space- localized short circuit be-

tween the plasma and the vessel potential, which is picked-up by the in-vessel antennas in their earth return circuit, and appears in the synchronous measurements as a very large and rapid variation in Z_{ANT} [17]. By setting a suitable alarm threshold (see Appendix-B4 for further details), the occurrence of these S/T trips has protected the in-vessel antennas without causing problems for the scientific operation of the AEAD system.

The AELM crate contains four main components.

1. A VMEbus Crate Service Module: this is a JET "home-grown" module and acts as the VMEbus system controller, providing timer/trigger synchronization with plant systems; it also monitors the VMEbus crate voltages, temperatures and airflow for operational control (plant failure).
2. A Real-Time Processor (RtProc): this is a commercial off-the-shelf (COTS) Emerson Network Power MVMEbus5500 card with a 1 GHz PowerPC and 512 MB RAM; this card executes software running under the Wind River VxWorks operating system (i.e. the same software used in the NASA's Mars rovers, Spirit and Opportunity), but during a JET pulse the main real-time process is "locked", so that context switching is disabled and most interrupt sources are masked.
3. A Communications Processor (CommsProc): this is a COTS Emerson Network Power MVMEbus5100 card with a 400 MHz PowerPC and 64 MB RAM; it is linked to both the JET real-time Asynchronous Transfer Mode (ATM) and Ethernet networks and isolates the RtProc from asynchronous events that might disrupt its deterministic 1 kHz calculation cycle time, hence guaranteeing deterministic message delivery; the CommsProc also executes software running under the WindRiver VxWorks operating system, and it is used to setup pre-pulse information, synchronize the RtProc with the main time points within the pulse and communicate data recorded during the pulse for archiving.
4. Four COTS Pentland Systems MPV956 analogue and digital input/output cards (VAJ1): these cards are configured to use differential analogue inputs and sample 32 input signals (8 for mode detection, 24 for plant control); these cards are also used to send calculated analogue signals to control the power and frequency of the AEAD plant in real time.

The CommsProc basically coordinates external messages during all phases of the JET experiment (i.e. a JET pulse), such as:

- continuous monitoring of both hardware and software: this ensures that the AEAD is stopped if any asynchronous termination events are generated by the JET protection system;
- continuous monitoring of both hardware and software to provide a coarse, 2 sec AELM system health and data update for the control room computers: this ensures that pulses cannot be started if the system is not ready and "slow" graphical displays in the control room are updated with data throughout the pulse;
- parameter initialisation: the AELM requires approximately 200 setup parameters per pulse (not all change every pulse and they are managed using a separate GUI); any derived parameters are calculated at this moment to avoid unnecessary work during a pulse;
- real-time experiment control communications, operating as a two way street: the AELM provides real-time values for use by other systems to help manipulate their actions during a pulse, and similarly the AELM can receive externally generated instructions for controlling, for instance, the operating frequency for the in-vessel antennas.
- post-pulse data archiving: the AELM records several megabytes of raw ADC and processed data during a pulse, which are archived onto a data warehouse for post-pulse analysis.

The CommsProc can basically be thought of as a synchronization module with four available states:

a) *waiting for the next experiment to start*: the CPU performs basic periodic environment and systems checks to ensure that all hardware and software are running correctly and the networks are able to send and receive "keep alive" data packets;
b) *initialization for the next experiment*: the CPU receives all the parameters for the experiment and configures the hardware and software accordingly;
c) *pulsing/performing the experiment*: the CPU synchronizes the sending and receiving of data over the real-time ATM network and triggers the RtProc based upon a 1 kHz clock interrupt from the Crate Service Module; the CommsProc ensures that the AEAD plant ceases to operate if there is a CPU failure or an external systems failure terminating the experiment prematurely;
d) *data collection*: the CPU halts the RtProc and returns all the data collected to the JET database archiving system for post-pulse analysis and simulation of plant operation.

The RtProc takes care of pre-pulse initialisation by removing all conflicts between deterministic behaviour and the requirement for a flexible system, and of real-time logging, detection and tracking. For the pre-pulse initialisation, the algorithms used to provide flexible signal combinations and alternate methods for mode detection require the system to be dynamic, but this is not compatible with a *hard* real-time system: hence, the worst case scenario for memory requirements is computed and allocated before an experiment starts. The RtProc is only activated during a JET experiment (i.e. a JET pulse), otherwise it remains

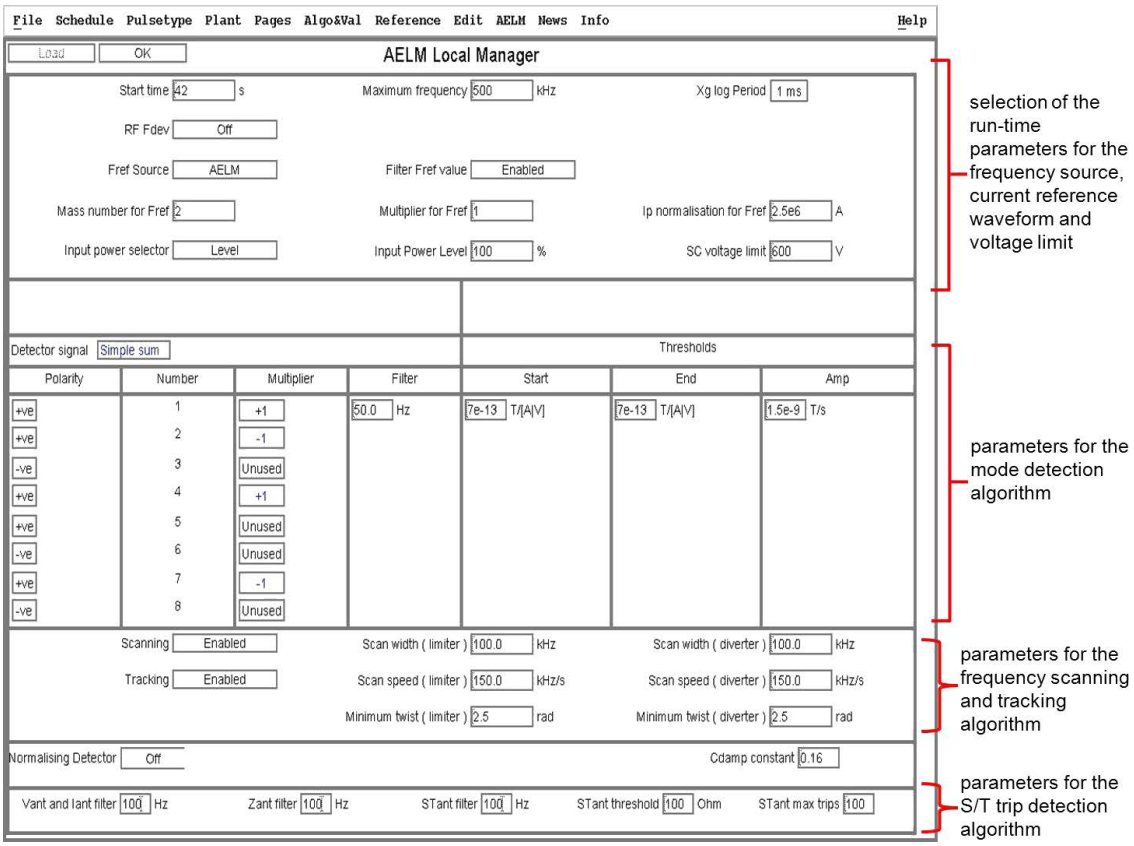

Fig. 4 The AELM configuration panel, shown here for operation with the *SimpleSum* algorithm and using frequency control via the real-time frequency of $n = 1$ Toroidal Alfvén Eigenmodes.

dormant, and operates on the same 1 kHz clock cycle required for compatibility with the CTTS system. The sequence of this cycle is as follows:

a) read and pre-process the analogue signals from the I/O cards (taking usually ~70 μs);
b) perform the calculations required for S/T detection on the 8 antennas, so that if a faster-than-normal change in the antenna impedance is identified, then the AELM will send a trip signal to the amplifier to suspend its output for 20 ms for plant protection (these calculations take usually ~50 μs); the user is able to configure the total number of trips tolerated before one (or more) antenna(s) is (are) permanently excluded from the remainder of the pulse (see Appendix B.4);
c) perform mode detection, discrimination and tracking using one of two user-selectable options, described in Sec. 4 in more detail: the calculations performed in the *SimpleSum* mode take usually ~100 μs, whereas those required for the *SparSpec* mode usually take ~650 μs;
d) finally, the VCO frequency is calculated and sent back to the AEAD plant for the next time step, together with the requested I_{REF} value (taking usually ~50 μs).

The selection of the run-time parameters for the operation of the AEAD system is performed via a configuration panel, and Fig. 4 shows its main operating window for the *SimpleSum* algorithm when the real-time value of the frequency of $n = 1$ Toroidal AEs (TAEs) is used for frequency control.

The real-time experiment control works on the following principles:

- the AELM provides signals to the RTSS systems and can also be controlled by them:
  - the experiment server monitors 35 systems providing > 1000 signals and can be used to control diagnostic, gas, heating and plasma shape control systems; the experiment server uses a Matlab logic block style language to allow signals from a multitude of systems to be combined in algorithms to control gas, pellet and heating injection into the plasma;
  - the protection server monitors 13 systems providing ~450 signals that are used to support machine protection; the protection server uses simple arithmetic and Boolean logic combinations to trigger the machine protection systems.
- the real-time control data is routed over an Asynchronous Transfer Mode (ATM) network: this pro-

Table 1 Physics and engineering constraints which need to be satisfied with the specific operational settings for the AELM.

AELM parameter	AELM value	specification
sampling frequency	1 kHz	determined by the integration of the AELM within the CTTS system, i.e. NOT a user choice.
processing time	850 μs @ 1 GHz	determined by the integration of the AELM within the CTTS system, i.e. NOT a user choice.
number of input channels	mode detection: 8 plasma data: 4 engineering: 22	determined by the RAM and CPU resources available for real-time calculations within the AELM. engineering: compatibility with the JET overall plant control and protection system. mode detection and plasma data: compatibility with the post-pulse data analysis.
number of output channels	VCO frequency: 1 mode amplitude: up to 31 engineering: 12	determined by the RAM and CPU resources available for real-time calculations within the AELM. mode detection and tracking: only the VCO frequency is required for AEAD operation. mode amplitude: needed for all output mode numbers, hence number of output channels depends on which real-time algorithm is used for mode detection and tracking. engineering: compatibility with the JET overall plant control and protection system.
mode frequency	10 kHz to 500 kHz absolute error $\Delta f \leq 100$ Hz	determined by the requirement of compatibility with the post-pulse data analysis. severe error constraint determined by compatibility with synchronous detection.
toroidal mode number	$\lvert n\rvert \leq 12$ relative error $\lvert\Delta n/n\rvert \leq 0.1$	result of simulations: $\lvert n\rvert = 12$ is the maximum mode number that can be evaluated in real-time within the set processing time constraint and the specified error using the available input channels. severe error constraint determined by compatibility with the post-pulse data analysis.
mode amplitude	$\lvert\delta B_{\mathrm{MEAS}}\rvert = 0.1$ mG relative error ≤ 2	determined by the requirement of compatibility with the post-pulse data analysis. very mild error constraint: only time evolution is necessary for real-time analysis.
mode damping rate	$\gamma/\omega = 0.2$ % relative error ≤ 0.3	determined by the requirement of compatibility with the post-pulse data analysis. severe error constraint determined by compatibility with the post-pulse data analysis.

vides guaranteed, deterministic message delivery; ATM networks were originally popular with telecoms companies because of their extremely high reliability (according to switch statistics at JET, none of the systems connected to the RTDN have ever lost a packet of data in 10 years of operation); moreover, using a network infrastructure that differs from Ethernet ensures that time sensitive data is not perturbed by non-urgent, general data.

The physics and engineering constraints that the specific operational settings for the AELM need to satisfy are summarised in Table 1, with more technical details given in Appendix B. In terms of the required accuracy in the data analysis, both real-time and post-pulse, we have set ourselves the high standard of trying to match the nominal tolerances for the ITER measurement requirements for high-frequency instabilities [1, 70, 71], specifically with respect to the accuracy of the toroidal mode number determination. This is very important as, for instance, a nominal 10 % or a ±1 tolerance on the toroidal mode number determination leads to very different requirements for the system depending on which "specific" mode number this tolerance is applied to. In terms of physics interpretation and for real-time plasma protection and control applications, wrongly interpreting the $n = 1$ mode as an $n = 0$ or an $n = 2$ mode (i.e. $n = n \pm 1$) clearly does not have the same implications as wrongly interpreting the $n = 5$ mode as an $n = 6$ or an $n = 4$ mode, i.e. still having a ±1 error on n.

Hence, for post-pulse analysis, we have decided to consider that the acceptable error is ±0 on the toroidal mode number and ±15 % on the measured mode amplitude $\lvert\delta B_{\mathrm{MEAS}}\rvert$ for low-$\lvert n\rvert \leq 5$ modes of importance for plasma protection and control and for real-time measurements. Examples of these modes are the precursor for sawteeth, Edge Localized Modes and disruptions, and radially extended MHD instabilities such as global Alfvén Eigenmodes, Neoclassical Tearing Modes, Resistive Wall Modes

and Alfvén Cascades. Conversely, a mode number measurement error ranging from ±1 to ±3 is deemed to be acceptable for MHD instabilities which are only of "physics" interest, for instance core localized Alfvén Eigenmodes with $|n| = 6 \rightarrow 20$, for which the amplitude $|\delta B_{\mathrm{MEAS}}|$ only needs to be measured within ±30 %. The required post-pulse measurement accuracies on the mode amplitude and toroidal mode number and can then be summarized as follows:

mode number	acceptable error on toroidal mode number calculated post-pulse	acceptable error on mode amplitude calculated post-pulse
$\|n\| \leq 5$	error = 0 ➔ CORRECT detection error ≠ 0 ➔ WRONG detection	error ≤ 15% ➔ CORRECT detection error ≥ 15% ➔ WRONG detection
$\|n\| > 5$	$\|$error$\|$ ≤ min(0.1*n, 3) ➔ CORRECT detection $\|$error$\|$ ≥ max(0.1*n, 3)➔ WRONG detection	error ≤ 30% ➔ CORRECT detection error ≥ 30% ➔ WRONG detection

For real-time analysis, we aim at achieving the same post-pulse accuracy on the mode number for low-n modes up to $n| \leq 12$, i.e. a relative error $|\Delta n/n| = 0$ up to $|n| \leq 5$ and $|\Delta n/n| = 1$ for $6 \leq |n| \leq 12$, whereas we are satisfied with an accuracy of a factor 2 in the measured mode amplitude as the limitations in the computing resources of the AELM prevent setting a higher requirement.

For the measured damping rate, we require a post-pulse and real-time accuracy of better than a factor 2, and this particularly for small values of the damping rate, i.e. we require the capability of accurately distinguishing in real-time a small damping rate $\gamma/\omega = 0.1$ % from a slightly larger value $\gamma/\omega = 0.2$ % as the first one is so close to the marginal stability limit $\gamma/\omega = 0$ that an alarm might have to be triggered so as to activate a feedback system, whereas the second one is sufficiently away from the marginal stability limit that no alarm would need to be raised.

For the mode frequency, we set a required accuracy for its real-time determination of better than 100 Hz, so as to take advantage of the capabilities of the synchronous detection system, which gives us a post-pulse accuracy on the measured mode frequency of 100 Hz.

Finally, it is important to note here that many of the AELM parameters have to be set and their requirements have to be met for compatibility with the overall plant control and protection system of JET, i.e. these are not a user-choice. Conversely, the specific values and the error requirements on the mode frequency, amplitude, toroidal mode number and damping rate are essentially a user-choice which is set for compatibility with the post-pulse data analysis but has to be met within the set processing time and using the RAM and CPU processing speed available for the AELM calculations in real-time: in some cases a trade-off exists between real-time RAM+CPU resources, and real-time vs. post-pulse physics requirements. Hence, a system upgrade (increase of RAM, increase of CPU processing speed, code parallelisation, ...) will allow a more precise determination of these quantities together with the possibility of extending the measurement range and the scope and flexibility of the AELM and AEAD plant.

4. Setting the parameters for the real-time detection and tracking of the resonant plasma response to the antenna-driven perturbations via the AELM

When the *tracking* mode of operation has been selected, the AELM linearly sweeps the antenna frequency around the initial pre-defined F_{REF} guess for the frequency of the modes to be detected and attempts to track in real-time the resonant plasma response to the antenna-driven magnetic field perturbation as the background plasma evolves. Using a selection of the eight input signals which are available in real-time for mode detection, the user can then select either the original *SimpleSum* or the recently developed *SparSpec* algorithm to derive one single amplitude and phase pair that will be used for mode detection and tracking.

Figure 5 shows an example of real-time tracking of TAEs together with a set of schematic diagrams illustrating the basic working ideas of the tracking algorithm as implemented in the AELM. In the full-frequency spectrogram for the calibrated data of one magnetic pick-up coil (acquired at 1 MHz, see Fig. 5a) we find a very faint triangular waveform with amplitude $|\delta B| \approx 5$ mG, performing a very narrow sweep in the frequency range 190 kHz to 230 kHz over the time window 4.50 sec to 12.50 sec. This trace corresponds to the magnetic field perturbation produced by the AEAD system, which was set to scan around the central frequency of the $n = 1$ TAE gap: note that indeed a mode appears out of the background in this frequency range from time = 11 sec onwards. Apart from

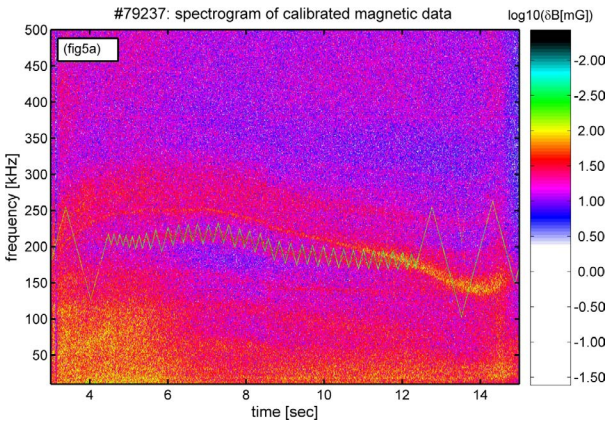

Fig. 5a Spectrogram for the calibrated magnetic data acquired (asynchronously) at 1 MHz with one pick-up coil mounted on the low-field side vessel wall for the JET He4 discharge #79237. Note the faint ($|\delta B| \sim 5$ mG) triangular trace around 200 kHz, which correspond to the waveform for the antenna drive, and the mode appearing out the background turbulence in a similar frequency range.

the antenna-driven signal and this background instability, the magnetic fluctuation spectrum shows no other activity in the entire frequency range up to 500 kHz. Figure 5b shows a tracking representation of the signal synchronously detected at 1.25 kHz for the same magnetic pick-up coil (δB_{TAE}). This tracking representation, where the quantity "time + $|\delta B_{\mathrm{TAE}}|$" is plotted as a function of the antenna driving frequency, immediately reveals the resonant characteristic of the plasma response to the antenna-driven signal, which is highlighted in the insert for clarity. The resonant plasma response to the antenna drive (the so-called antenna/plasma transfer function) is then analysed in real-time using a Lorentzian-type (i.e. and inverted bell-shape) model [3]:

$$\delta B_{\mathrm{TAE}}(\omega, x) = \frac{\alpha(x) + i\beta(x)\left[\omega/(2\pi f_{\mathrm{MEAS}})\right]}{1 - \left[\omega/(2\pi f_{\mathrm{MEAS}})\right]^2\left[1 - 2i\gamma_{\mathrm{DAMP}}/(2\pi f_{\mathrm{MEAS}})\right]}. \tag{2a}$$

$$\delta B_{\mathrm{TAE}}(\omega, x) = \frac{1}{2}\sum_{k=1}^{N/2}\left\{\frac{R_k(x)}{i\omega - P_k} + \frac{R_k^*(x)}{i\omega - P_k^*}\right\}$$

$$= \frac{\sum_{m=0}^{M} b_m(x)(i\omega)^m}{\sum_{n=0}^{N} a_n(i\omega)^n} = \frac{B_M(\omega, x)}{A_N(\omega)} + D(\omega, x). \tag{2b}$$

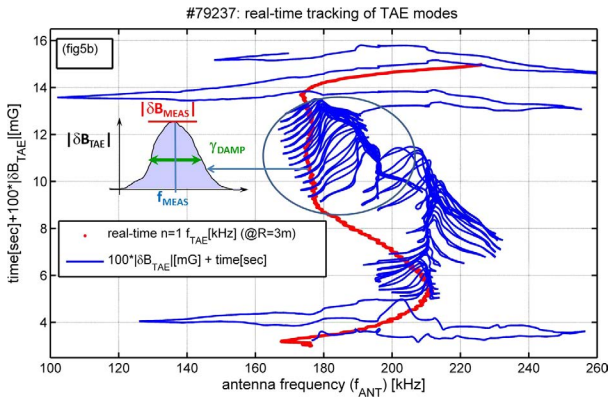

Fig. 5b Tracking representation of the magnetic fluctuation data for the same JET discharge and pick-up coil shown in Fig. 5a. The smaller sweeps of the antenna frequency waveform seen in Fig. 5a around 200 kHz correspond to the time interval when real-time tracking of stable TAEs is obtained, which is indicated by the real-time "tracking" quantity time + $|\delta B_{\mathrm{TAE}}|$ showing an inverted bell-shape feature. Real-time tracking is then obtained in the time interval 4.50 < time [sec] < 12.50, whereas for time < 4.50 sec and time > 12.50 sec the AEAD system is in scanning mode, i.e. it is looking for a resonant plasma response to the antenna drive which fits the Lorentzian-type model given in Eq. (3), and graphically illustrated in Fig. 5c, for the mode characteristics (such as frequency, damping rate and amplitude). The central frequency of the $n = 1$ TAE gap (f_{TAE}) is evaluated on the magnetic axis (i.e. at $R = 3$ m) by the AELM software without the contribution of the plasma toroidal rotation (not directly available in real-time from the RTSS), which produces a Doppler shift of about ~3 kHz in the post-pulse estimate of f_{TAE}.

In Eq. (2a) the label "x" indicates the different sensors that can be used to evaluate δB_{TAE} in real-time, and $\alpha(x)$ and $\beta(x)$ are fitting parameters. As shown in Fig. 5b, the central value of the frequency sweep provides the mode frequency (f_{MEAS}), the normalized half-width at half-maximum provides the quality factor (i.e. the mode damping γ_{DAMP}), and the peak value ($\equiv |\delta B_{\mathrm{TAE}}|$ at $\omega = 2\pi f_{\mathrm{MEAS}}$) provides the mode amplitude ($|\delta B_{\mathrm{MEAS}}|$) in real-time. In the tracking representation of Fig. 5b, note that the antenna frequency (f_{ANT}) does not always closely follow the real-time value of the $n = 1$ TAE frequency ($F_{\mathrm{REF}} = f_{\mathrm{TAE}}$ in this particular case). In fact, the narrow sweeps of f_{ANT} occur only in tracking mode. The much larger sweeps indicate that tracking has been lost as the plasma background evolves and the AELM has now begun scanning its operational frequency range from F_{REF}-*ScanWidth* to F_{REF}+*ScanWidth* to look for new resonances.

Narrow sweeps of the antenna frequency, i.e. mode tracking, occur only when the measured plasma response δB_{TAE} to the antenna driven perturbation is sufficiently close to the pre-set resonant shape corresponding to the model of the resonance given in Eq. (2a), which is shown in graphical form in Fig. 5c and Fig. 5d as a function of the antenna frequency. Remembering now that δB_{TAE} is in fact a complex-valued quantity as it is a synchronously acquired measurement (i.e. it has both the in-phase (I: real) and the quadrature (Q: imaginary) components), we

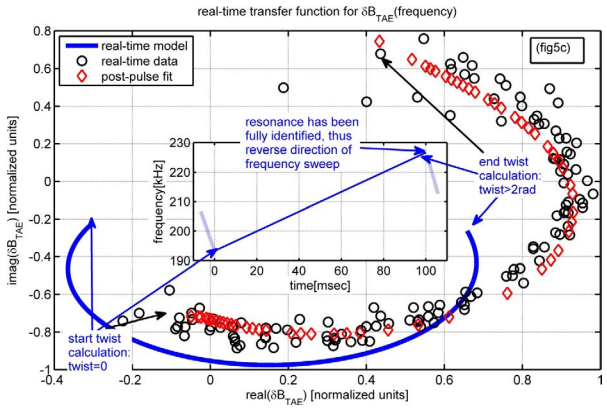

Fig. 5c A diagram illustrating the basic working ideas of the real-time tracking algorithm. When the plasma response δB_{TAE} to the antenna driven perturbation describes a circle in the complex-plane representation imag(δB_{TAE}) vs. real(δB_{TAE}), we determine that a mode resonance has been found. The shift between the real-time model and the actual real-time data originates from a non-resonant coupling term, which is taken into account in the post-pulse analysis.

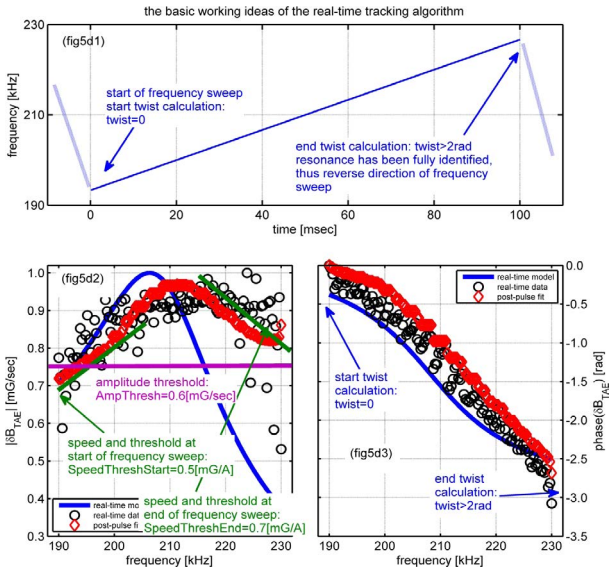

Fig. 5d A second diagram illustrating the basic working ideas of the real-time tracking algorithm. At the beginning of each frequency sweep, the AELM starts integrating the *twist* variable, i.e. the angle between imag($\delta B_{\rm TAE}$) and real($\delta B_{\rm TAE}$) if the signal amplitude $|\delta B_{\rm TAE}|$ and its *speed* $\mathrm{d}(|\delta B_{\rm TAE}/I_{\rm ANT}|)/\mathrm{d}t$ are above the user-set thresholds, AmpThresh = 0.6 mG/sec (shown by the magenta line) and SpeedThreshStart = 0.5 mG/A (shown by the green line). We estimate that the antenna-driven plasma resonance has been fully identified when the integrated *twist* exceeds the user-set threshold = 2 rad and if the *speed* at the supposed end of the frequency sweep is below the user-set threshold SpeedThreshEnd = 0.7 mG/A (again, shown by the green line). Hence, at this moment the direction of the frequency sweep is reversed in an attempt to follow the evolution of the same mode as the background plasma evolves.

easily see that the solution of Eq. (2a) describes a circle in the complex-plane representation imag($\delta B_{\rm TAE}$) vs. real($\delta B_{\rm TAE}$), as shown in Fig. 5c. The model antenna-plasma transfer function shown in Fig. 5c and Fig. 5d was computed (post-pulse but with exactly the same algorithm as that used for the real-time analysis) for a TAE mode with $|\delta B_{\rm MEAS}| = 1$ mG, $f_{\rm MEAS} = 200$ kHz and $\gamma_{\rm DAMP} = 37$ kHz. As seen in Fig. 5c, the actual real-time $\delta B_{\rm TAE}$ data (acquired for #79237 in the time interval 9.980 < time[sec] < 10.020) is shifted with respect to the real-time model: this occurs because of a non-resonant contribution due to a direct coupling term between the antenna and the magnetic sensor(s) that exists already in vacuum, and which is then mediated, thus almost always augmented, by the plasma for an actual plasma discharge.

The post-pulse fit of the complete antenna/plasma transfer function is given in Eq. (2b). This formula then separates the resonant plasma response $B_{\rm M}(\omega, x)/A_{\rm N}(\omega)$ from the non-resonant coupling term $D(\omega, x)$, when $N(= 1)$ in this case is the number of poles P_k (corresponding to the number of modes in the synchronously measured spectrum for any given angular frequency ω), R_k is the residue for each one of these poles, and $M > 2N + 1$ is the degree of the polynomial in Laplace ($i\omega$) which best fits the plasma response to the antenna drive.

The simple real time estimate of the mode amplitude, frequency and damping from the centre and width of the plasma resonant response to the antenna drive (as shown in the insert of Fig. 5b) very closely follows the value obtained from a post-pulse analysis [18], which uses the complex plane representation of $\delta B_{\rm TAE}$ to also extract the non-resonant contribution to the measured fluctuation signals (i.e. the direct antenna/plasma coupling term). Despite this non-resonant term, both real-time and post-pulse analyses give results very close to those of the model [18], as shown here:

	model data	*real-time analysis*	*post-pulse analysis*
mode frequency $f_{\rm MEAS}$	200.00 kHz	198.65 kHz	201.32 kHz
mode damping $\gamma_{\rm DAMP}$	37.00 kHz	37.23 kHz	36.87 kHz
mode amplitude $\lvert\delta B_{\rm MEAS}\rvert$	1.00 mG	0.91 mG	0.97 mG

Comparison between the mode characteristics (frequency, damping, amplitude) as given by a model and by the actual real-time and post-pulse data analyses, for the data shown in Fig.5.

When a plasma resonance has been correctly identified, the antenna frequency is swept back and forth around it, as shown in the insert in Fig. 5c. This approach allows the detection and tracking of hundreds of individual resonances during one single tokamak discharge, and to follow the evolution of the mode characteristics in real time as the background plasma evolves. Moreover, these modes have the same toroidal mode number if the *SparSpec* algorithm has been selected for AELM use.

User input is needed for a number of parameters which the AELM requires for the mode detection algorithm; more technical details are presented in Appendix B.5. As shown in Fig. 4, these parameters are selected via the tab *DetectorSignal*, and Fig. 5d illustrates their use in the basic common ideas behind both the *SimpleSum* and the *SparSpec* real-time tracking algorithms. Note that in Fig. 5d the AELM real-time separate unit system is used, which does not always correspond to the physical units of the post-pulse data. This is due to a frequency normalization that is different in the AELM software when calibrating the TAE-NORM and TAE-SIG channels with respect to the Laplace-space ($s = i\omega = 2\pi i f$) frequency domain where the calibration is performed for the post-pulse data.

At the beginning of each frequency sweep, if a mode has been detected, the AELM software checks that the signal amplitude $|\delta B_{\rm TAE}(\omega)|$ and its *speed* $\mathrm{d}(|\delta B_{\rm TAE}(\omega)/I_{\rm ANT}(\omega)|/\mathrm{d}t$ at the angular frequency $\omega = 2\pi f_{\rm ANT}$ are above the user-set and frequency-independent thresholds *AmpThresh* = 0.6 mG/sec and *SpeedThreshStart* = 0.5 mG/A, respectively, for the case shown in Fig. 5c. Figure 5(d1-d3) provides a graphical illustration of the implementation of this algorithm. When these conditions are met, the AELM software starts

computing the time-integral of the *twist* variable, i.e. the angle between imag($\delta B_{\mathrm{TAE}}(\omega)$) and real($\delta B_{\mathrm{TAE}}(\omega)$) using the complex plane representation of $\delta B_{\mathrm{TAE}}(\omega)$. We estimate that the antenna-driven plasma resonance has been fully identified when the integrated *twist* value exceeds the (user-set, frequency-independent) threshold *MinimumTwist* = 2 rad and the *speed* at the end of the frequency sweep is below the threshold *SpeedThreshEnd* = 0.7 mG/A (frequency-independent, user-set). At this moment the direction of the frequency sweep is reversed in an attempt to follow the evolution of the same mode as the background plasma evolves.

4.1 The *SimpleSum* algorithm for mode detection and tracking

The *SimpleSum* algorithm combines up to eight real and imaginary pairs from the total eight input AELM-SIG channels available to obtain one single output pair of real and imaginary components. This combination is performed using only un-weighted sums and differences, i.e. each I/Q pair has the same importance in determining the single amplitude and phase pair used for mode detection and tracking. The output pair used by the *SimpleSum* algorithm is then constructed as follows:

$$\begin{aligned} output\,(\phi) = & \sum_{k=1\equiv NUMBER}^{8} (Polarity)_k \\ & \times [\mathrm{Real}\,(input\,(\phi_k)) + \mathrm{Imag}\,(input\,(\phi_k))] \\ & \times (Multiplier)_k\,. \end{aligned} \tag{3a}$$

With this algorithm, and assuming that all utilized sensors have the same frequency response over the frequency range selected for operation, a simple discrimination between $|n|$ = odd and $|n|$ = even modes can in principle be obtained when the signals from a single pair of sensors located at toroidal opposite positions are subtracted (*multipliers* $= +1/-1$) or summed (*multipliers* $= +1/+1$), respectively. This can be readily seen by noting that if the sensors "k" and "$k+4$" (for $k = 1 \rightarrow 4$: remember that the AELM can accept up to eight input signals from magnetic pick-up coils) are pairs located at opposite toroidal positions ϕ_k and $\phi_k + \pi$, respectively, then the combination of this two signals can be reformulated using Eq. (3a) as:

$$\begin{aligned} & output\,(\phi) \\ & = A_0 e^{in\phi}\left[1 + \frac{(Polarity)_{k+4}\times(Multiplier)_{k+4}}{(Polarity)_k\times(Multiplier)_k}\,(-1)^{|n|}\right], \end{aligned} \tag{3b}$$

where A_0 is the mode amplitude and n is the toroidal mode number. A more complex mode number discrimination can be obtained when more than one pair is selected, as shown in Eq. (3c):

$$\begin{aligned} & output\,(\phi) \\ & = A_0 e^{in\phi}\sum_{k=1}^{4}\left[1 + \frac{(Polarity)_{k+4}\times(Multiplier)_{k+4}}{(Polarity)_k\times(Multiplier)_k}\,(-1)^{|n|}\right]. \end{aligned} \tag{3c}$$

Utilization of the *SimpleSum* algorithm suffers mainly from three specific limitations. First, JET is equipped with ten sensors equi-spaced along the toroidal angle, but only up to eight input signals from magnetic pick-up coils can be connected to the AELM, hence the intrinsic periodicities of the whole complement of such "toroidal" sensors cannot actually be fully exploited with the AELM. Second, two of these sensors (sensor #8 and #10, hence belonging to different pairs) are now open-circuit (O/C) in-vessel, hence only three pairs can be used with the *SimpleSum* algorithm, further reducing the real-time capabilities for exploiting the intrinsic periodicities of the full complement of toroidal sensors. Third, the sensors in the three surviving pairs are not exactly located at toroidally opposite positions (with a difference of around 0.1 to 0.5 degrees with respect to the optimal π separation, due to installation constraints), and this makes the sum and subtraction less effective in removing and/or selecting specific n-components. Finally, and although neither specifically nor solely related to the *SimpleSum* algorithm, as mainly due to the fact that in real-time the calibration is taken as a fixed value at the single frequency = 200 kHz, the assumption that the frequency response of the individual sensors is the same is not always true, and significant errors can be introduced when performing the *SimpleSum* calculations, particularly when utilizing more than one pair of sensors, and when the mode frequency differs by more than 30 kHz to 50 kHz from the value of 200 kHz at which the calibration is performed in real-time. Due to these limitations, the *SimpleSum* algorithm is commonly used in its most basic form for mode detection and tracking, i.e. selecting just one sensor (simplest form of mode detection and tracking without mode number selection), or a specific pair of sensors (with *multipliers* $= +/-1$, *multipliers* $= 0$ for all other sensors), but this use does not allow for a sufficiently selective real-time discrimination of the toroidal mode numbers of the detected modes. More technical details on the real-time implementation of the *SimpleSum* algorithm are presented in Appendix B.5.

4.2 The *SparSpec* algorithm for mode detection and tracking

The *SparSpec* algorithm can also be used in the real-time AELM software (see Appendix A for a detailed description of the mathematical foundations of this algorithm, with additional details on some applications other than those presented in this work). Basically, the *SparSpec* algorithm performs a least-square minimization of the input data with respect to a sum of complex mode amplitudes selected to belong to a pre-defined mode number base. *SparSpec* then penalizes adding further terms in the

solution using the L1-norm criterion, i.e. the sum of the absolute values of the mode amplitudes, so that the solution invoking the smallest number of modes is retained. This can be expressed as the L1-norm minimization of the quantity J, which can be defined as follows:

$$J_1(x) = \|\boldsymbol{y} - W\boldsymbol{x}\|^2 + \lambda \|\boldsymbol{x}\|_1 = \|\boldsymbol{y} - W\boldsymbol{x}\|^2 + \lambda \sum_{k=-K}^{K} (|x_k|) . \quad (4)$$

In Eq. (4) $\boldsymbol{y}$ is the vector of data, $\boldsymbol{x}$ is the vector of the complex mode amplitudes, $\boldsymbol{W}$ is the spectral window of the measurement system, and λ is the L1-norm penalization parameter, which is related to the noise level in the input data. Note that the value λ must also be set to be consistent with the required accuracy on the measurement.

In its real-time implementation within the AELM, *SparSpec* accepts up to eight real and imaginary input pairs but in principle can detect multiple modes, hence it can produce several amplitude and phase pairs as output signals. Two methods (*highest* and *any*) can be used for selecting the output modes identified with *SparSpec* upon which mode detection and tracking is then performed by the AELM. With the *highest* method the AELM software chooses, at any time point, the detected mode which has the highest amplitude, independently of which modes have been previously selected. The *any* method implements a "running memory" scheme: the AELM chooses the mode whose amplitude is above a given threshold, starting from the minimum mode number selected by the user. If a resonance is then detected, this mode will continue to be selected until tracking is lost when the search for a new mode will re-start. This method has the disadvantage that the highest amplitude solution is not always chosen and then tracked in real-time, but the first one that is found when searching for a mode. A new "running memory" scheme is currently being tested off-line for future real-time implementation, so that at the time point t_j the mode obtained through a time-history-weighted analysis of the modes detected at all previous time points t_{j-k} (with $k = 1, 2, \ldots,$) will be considered as the first guess for mode detection and tracking. This will avoid possible (but not often observed) spurious jumps in the mode tracking algorithm that can occur with the *any* and more commonly with the *highest* tracking methods. More technical details on the real-time implementation of the *SparSpec* algorithm are presented in Appendix B.5.

With respect to the implementation of the *SparSpec* algorithm within a real-time environment such as the AELM, it is important to note that certain mathematical functions (such as square root, sine, cosine and tangent) are very CPU intensive. This problem is solved using two different methods:

- use maths coprocessors, some of which have array and simple matrix functionality – standard for desktop PC systems but not always available within embedded systems using exotic CPUs;
- use simplified lookup tables or approximations – but this can lead to inaccuracies and tables can be a problem in systems with small amounts of memory.

Similarly, algorithm optimisations are most effective at the mathematical formula definition stage. When the formula is interpreted into a computer language, compilers are becoming more adept at translating it into optimised, efficient execution code. Hence, optimisations of the *SparSpec* algorithm used in real-time by the AELM involves:

- pre-computed constants or common calculations that can be stored for later re-use – but compilers are getting better at doing this automatically;
- removing or limiting dynamic calculations, as these are a major problem within a real-time environment; algorithms that require various amounts of memory for arrays or matrices need to be forced into a fixed size – this may mean reserving memory for the worst case or truncating a calculation; memory allocation is a time costly operation and should be avoided within time sensitive systems, and this was indeed one of the major development headaches for the *SparSpec* implementation.
  - memory requirements – arrays or recursion: recursion is a special form of memory allocation and it is used to refer to an algorithm that calls itself; each successive invocation requires more memory so the system can trace where it came from and can unravel itself once the calculation has completed; limits can be placed on the depth of recursion allowed but this may lead to an inaccurate value.
  - loop termination conditions – add a time limit: loop terminations based on achieving a target accuracy or error tolerance are open-ended with regards to time, which was tackled whilst implementing the *SparSpec* algorithm; we used an upper time limit that would force the calculation to terminate at 800 µs even if it hadn't reached the required accuracy; the forced termination is recorded and can be plotted next to the results for post-pulse analysis.

One important point to mention for the real-time implementation of the *SparSpec* algorithm is its capability of determining all the mode characteristics (frequency, toroidal mode number, and mode amplitude and damping rate) with the required accuracy while satisfying the CPU time limit of the AELM. This has been obtained following the off-line optimization of the original post-pulse version of the *SparSpec* code, using real and simulated data. The details of these studies are separately given in the Appendix A.4 for clarity of presentation and briefly summarised in Sec. 7.

Table 2 Comparison between the real-time and post-pulse implementations of the *SimpleSum* and *SparSpec* algorithms.

	***SparSpec*-RT**	***SparSpec*-PP**	***SimpleSum*-RT**	***SimpleSum*-PP**
calibration	fixed value at 200 kHz	full frequency-dependence	fixed value at 200 kHz	full frequency-dependence
input data	up to 8 complex-valued magnetic signals (TAE-SIG) 1 current normalization signal as given by TAE-NORM	up to 16 complex-valued signals corresponding to magnetic, reflectometry and ECE measurements selection between many current normalization signals	up to 8 complex-valued magnetic signals (TAE-SIG) 1 current normalization signal as given by TAE-NORM	up to 10 complex-valued magnetic signals selection between many current normalization signals
output data	many amplitude and phase pairs, one pair for each selected mode number	many amplitude and phase pairs, one pair for each selected mode number	1 single mode amplitude and phase pair	1 single mode amplitude and phase pair
algorithm	choice between *any* and *highest* no least-square renormalization of output amplitude	choice between *any* and *highest* least-square renormalization of output amplitude implemented	default analysis once *polarity* and *multiplier* are selected no additional user-flexibility	default analysis once *polarity* and *multiplier* are selected, no additional user-flexibility
CPU limit	850 μsec @ 1 GHz	un-limited (user choice)	850 μsec @ 1 GHz	un-limited (user choice)
RAM limit	512 MB	un-limited (user choice)	512 MB	un-limited (user choice)
maximum *n*	$\lvert n_{MAX}\rvert \le 15$: mode number determined in real-time with relative error $\lvert\Delta n/n\rvert \le 0.1$	$\lvert n_{MAX}\rvert \le 30$: mode number determined post-pulse with relative error $\lvert\Delta n/n\rvert \le 0.1$	n = odd or n = even selection exact value of mode number cannot be obtained in real-time	$\lvert n_{MAX}\rvert \le 5$ when using all 10 available magnetic sensors, with relative error $\lvert\Delta n/n\rvert \le 0.3$

4.3 Comparison between real-time and post-pulse implementation of the *SimpleSum* and *SparSpec* algorithms

Having described in the two previous sub-sections the main aspects of the *SimpleSum* and *SparSpec* algorithms, it is now important to highlight the differences between their real-time and post-pulse implementations. Table 2 summarizes these differences, where we use the labels "RT" and "PP" to indicate the real-time and post-pulse implementations of both algorithms. It is important to note that the differences between real-time and post-pulse implementation of these algorithms are due to the very demanding constraints of the real-time implementation, namely the need to complete all real-time calculations within a clock time of $< 850\,\mu$sec using a 1 GHz PowerPC with only 512 MB of available RAM. Post-pulse processing does not suffer of these limitations in computing power, hence allows for a much greater flexibility and scope in the analysis.

5. Real-Time Detection of Frequency-Degenerated Toroidal Alfvén Eigenmodes via the AELM

Alfvén Eigenmodes (AEs) are a particularly important example of real-time mode detection and tracking in thermonuclear fusion experiments for two essential reasons. First, these waves are a natural Eigenmode of any magnetically confined plasma: their frequency is simply proportional to the ratio between the magnetic field and the square root of the plasma mass, thus representing the balance between the tension force of the ambient magnetic field lines and the plasma inertia. The analysis of the dispersion relation of AEs can thus provide unique information on the plasma isotopic composition, the safety factor profile and the toroidal rotation frequency via the Doppler shift in the AE mode frequency for different toroidal mode numbers [9, 10, 19]. Second, the fusion-born alpha particles (αs) have a supra-thermal speed at birth that is typically well above the Alfvén speed in the usual thermonuclear tokamak plasma conditions. Resonant interaction with AEs is the first wave-particle interaction encountered by the αs during their thermalization process: hence, this mechanism for phase-space and spatial diffusion needs to be appropriately monitored and controlled to guarantee good confinement of the αs themselves [20, 21].

A particularly challenging task for any real-time mode detection and tracking method is the capability of distinguishing MHD instabilities that have a very similar frequency but a different mode number, i.e. instabilities that belong to a frequency-degenerate spectrum where the half-width at full-maximum of two modes (which, as shown in Fig. 5, is closely related to their damping and growth rate $\gamma/\omega = \mathrm{imag}(\omega)/\omega$) is much smaller than their separation in frequency. Stable AEs (i.e. with a positive damping $\gamma_{DAMP} = \mathrm{imag}(\omega) > 0$) with toroidal mode numbers in the range $|n| \le 15$ are precisely such a class of MHD instabilities, as recently observed in JET [12, 13, 22–25]. The real-time discrimination between stable frequency-degenerate modes is even more challenging when their damping rate is small, for instance below $\gamma_{DAMP}/\omega < 0.5\,\%$, as is the case for the antenna-driven AEs in JET. Moreover, these low-damping instabilities are also the modes that are of most potential interest in future burning plasma experiments such as ITER, as such AEs are those most prone to be driven unstable by resonant fast ions, for instance fusion-born αs. In fact, it is precisely the experimental observation in JET of multiple and stable frequency-degenerate

AEs with very low damping rate and mode numbers in the range of interest for burning plasma experiments that has prompted the development and implementation within the AELM of the more sophisticated algorithm for real-time mode detection and tracking based on the method of the *Sparse Representation of Signals* [26–28] and the *SparSpec* code, which has now almost entirely replaced the previous, and much simpler, *SimpleSum* algorithm.

5.1 Real-time detection of Toroidal Alfvén Eigenmodes using the *SparSpec* algorithm

One example of real-time detection and tracking of multiple, stable, frequency-degenerate AEs in JET is observed during the discharge #79237. For this discharge the AEAD system was configured to drive $|n|$ = odd modes, with maximum amplitude for $|n| < 5$, and the AELM real-time mode detection and tracking algorithm was setup using the *SparSpec-any* mode to look for co-current (positive toroidal mode numbers) and counter-current (negative toroidal mode numbers) propagating AEs with $3 \leq |n| \leq 10$, using seven out of the eight magnetic sensors available in real-time and setting the run-time parameters (see Appendix B.5 for further details on the use of these parameters): *SS-Lambda* = 0.85, *SS-*N_{MAX} = 20, *SS-T'hold* = 5×10^{-10}, *Thresholds/Amp* = 3×10^{-9} [T/s], *Thresholds/Start* = 1×10^{-12} [T/A] = *Thresholds/End*, *MinimumTwist* = 2 [rad].

Figure 6 shows an overview of the main plasma parameters for this discharge #79237. Here $B_{\phi 0}$ is the toroidal magnetic field on the magnetic axis, I_{p} is the plasma current, q is the safety factor profile obtained by combining EFIT [29] reconstruction with Motional Stark Effect and polarimetry measurements when available, s is the magnetic shear profile, κ is the elongation profile, δ is the average top/bottom triangularity profile, T_{e} and T_{i} are the electron and ion temperature profiles, n_{e} is the electron density profile and Z_{EFF} is the plasma effective charge. In Fig. 6, and in the reminder of this work, the suffixes "0" and "95" indicate a value on the magnetic axis and at 95 % of the normalised poloidal flux, and the symbol "$\langle A \rangle$" indicates volume averaging of the quantity "A". The electron density and temperature were measured with a high-resolution Thomson Scattering (when available) or a LIDAR diagnostic system, with typical uncertainties of the order of ±10 % in both cases. The ion temperature and effective charge were measured using charge-exchange (CX) spectroscopy (when available, with typical uncertainties around 15 %) or derived from equilibrium reconstruction using ion-electron energy equi-partition and bremsstrahlung measurements, with typical uncertainties up to 25 % in this latter case. In addition to the background plasma parameters indicated above, the top frame of Fig. 6 also shows: the value of the magnitude (absolutely calibrated) of the antenna-driven radial component of the magnetic field ($|\delta B_{\mathrm{MEAS}}|$) measured

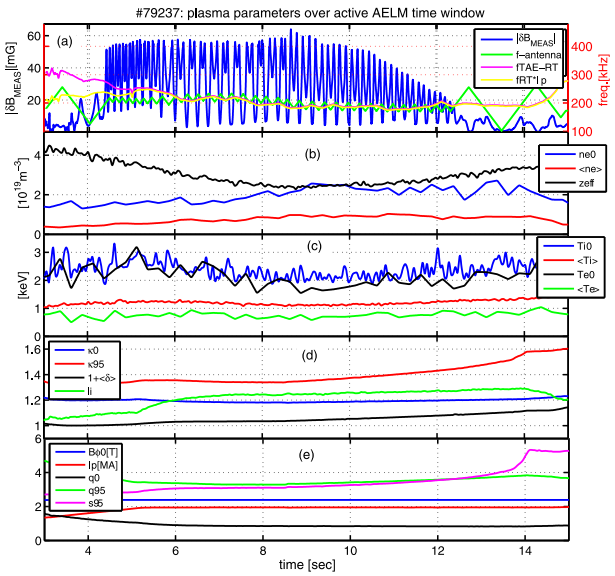

Fig. 6 Overview of the main plasma parameters over the active AELM time window for #79237. Frame-a: antenna frequency waveforms and driven magnetic field amplitude measured at the plasma edge; frame-b: electron density from LIDAR and Z_{EFF} from CX measurements, respectively; frame-c: electron temperature from LIDAR measurements and ion temperature from energy equipartition, respectively; frame-d: elongation, triangularity and internal inductance (l_{i}) from EFIT equilibrium reconstruction; frame-e: magnetic field, plasma current, safety factor and magnetic shear from EFIT equilibrium reconstruction.

with a pick-up coil (BTOR001) mounted on the low-field side vessel wall; the value of the antenna frequency (f_{ANT}); and the value ($f_{\mathrm{RT}} \propto B_{\phi 0\text{-RT}}/R_{\mathrm{GAP}}/q_{\mathrm{GAP}}/\sqrt{A_{\mathrm{EFF}}}/\sqrt{n_{\mathrm{e0\text{-RT}}}}$) of the central frequency of the n = 1 TAE gap computed in real-time using the values R_{GAP} = 3 m and q_{GAP} = 1.5 and a user defined A_{EFF}, without and with normalization ($f_{\mathrm{RT}} I_{\mathrm{p}} = f_{\mathrm{RT}} * I_{\mathrm{p}}(t)/\max(I_{\mathrm{p}})$) with respect to the time evolution of the total plasma current. Finally, note that the discharge #79237 enters into the X-point phase at t = 13.5 sec. This is indicated by the large increase at that time point in the edge magnetic shear s_{95} and edge elongation κ_{95}, which then remain both approximately constant, as shown in the bottom two frames in Fig. 6.

Figure 7 shows the comparison for #79237 between the AELM (real-time) and the post-pulse values for the antenna driving frequency and the reference n = 1 TAE frequency, with and without accounting for the time evolution of the plasma current. There is a ~700 Hz difference in the real-time/post-pulse antenna driving frequency due to electrical pick-up, which causes signal offsets and bit-noise in the post-pulse data acquisition that cannot be exactly compensated for in discharges. This difference is practically constant throughout the entire frequency range of the AEAD measurements, and does not really affect the ensuing data analysis. There are much larger differ-

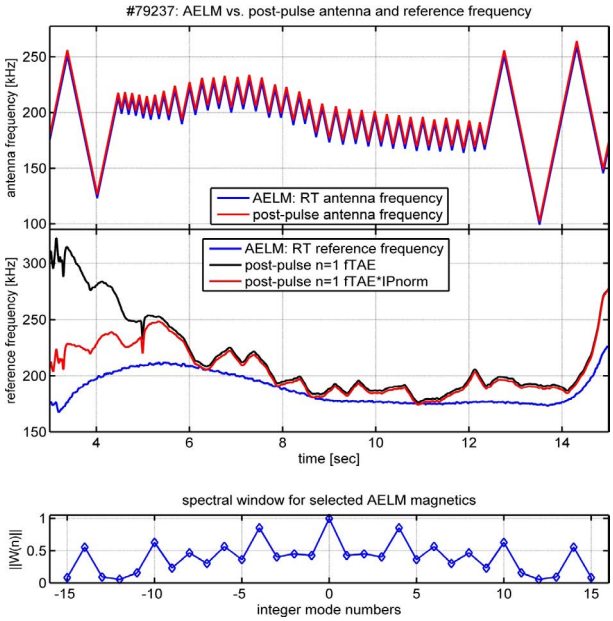

Fig. 7 Comparison between the AELM (real-time) and post-pulse values for the antenna driving frequency and the reference TAE frequency, with and without the normalization to the time evolution of the plasma current. There is an almost constant ~700 Hz difference in the real-time/post-pulse antenna driving frequency, caused by a signal offset in the post-pulse data acquisition that cannot be compensated exactly for all discharges as it is due to bit-noise and electrical pick-up. Additionally, the norm of the spectral window $||W(n)||$ for all the magnetics sensors used in real-time is shown is the bottom frame, to illustrate the toroidal periodicities (here, mostly due to an $n = \pm 4$ component) that we have to account for in the mode number decomposition.

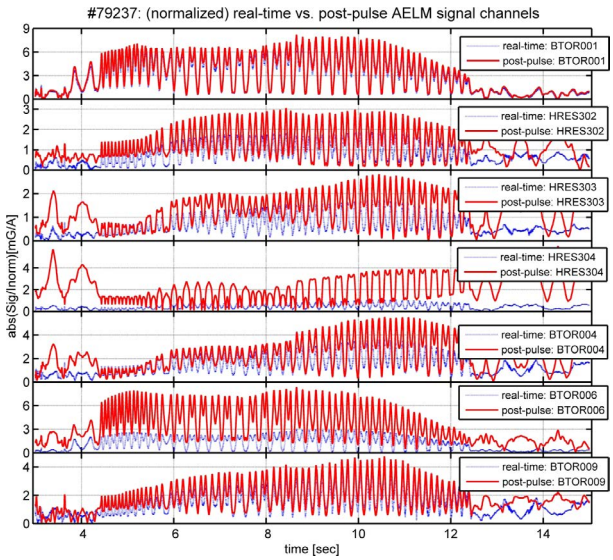

Fig. 8 Magnetic signal channels used in the AELM tracking algorithm for #79237, as processed using the different real-time and post-pulse options. Real-time analysis: using fixed values for the offset and calibration, taken at 200 kHz. Post-pulse analysis: using the full end-to-end frequency-dependent transfer function values for the offset and calibration.

ences between the real-time and post-pulse $n = 1$ reference TAE frequency, due to the different data input and processing options between real-time and post-pulse calculations. Again, our frequency sweep capabilities (up to ±200 kHz with a maximum 400 kHz sweeping speed) are such that these differences are in most instances of no practical consequence for the subsequent data analysis. Finally, the norm of the so-called *spectral window* $W(n)$ for all the magnetics sensors used in real-time by the AELM is shown is the bottom frame of Fig. 7. This illustrates secondary lobes related to the toroidal pseudo-periodicities that we have to account for in the mode number decomposition. For the diagnostic setup chosen for the AELM real-time analysis for #79237, the spectral window has high secondary lobes for $|n_{\mathrm{LOBE}}| = 4$. This in principle implies that if a mode number n is detected, it may actually correspond to a mode $n_0 = n \pm |n_{\mathrm{LOBE}}|$. The probability $P_{\mathrm{FALSE}}(n = n_0)$ for such a *false detection* of the "true" mode number n_0 as a "false" mode number n is $P_{\mathrm{FALSE}}(n_0) \propto \sigma^2(n_0) \times \Sigma(W(n) + W(|n| = n_{\mathrm{LOBE}})/W(n_0)$, where $\sigma(n_0)$ is the uncertainty on the determination of the true n-component (= n_0 in this example) in the absence of pseudo-periodicities in the spectral window, and the sum is intended on all possible combination of n and n_{LOBE} such that $n = n_0 \pm |n_{\mathrm{LOBE}}|$. Then, the noise in data may make it difficult to distinguish between the modes with toroidal mode numbers n and n_0 (see Appendix A for more details), particularly for secondary lobes at low mode numbers $|n_{\mathrm{LOBE}}| \leq 5$ such that $||W(n_{\mathrm{LOBE}})|| > 0.7$, i.e. for low-$|n|$ modes for which the norm of the spectral window is rather large.

Figure 8 shows the comparison between the seven magnetic signal channels used in real-time by the AELM tracking algorithm, and the corresponding channels as processed in the post-pulse analysis. Note that in real-time the magnetic channels are processed and calibrated using fixed values taken at 200 kHz, and are normalized using only the amplitude of the single antenna-current selected as the normalization channel. For the post-pulse analysis we use the full-frequency dependent offset and calibration and normalize the magnetic data with the total current (complex value: amplitude and phase) in all the active antennas. It is therefore clear, as shown in Fig. 8, that these different processing options produce quantitatively different magnetic data for the real-time and post-pulse analysis, although one can see that, qualitatively, the two sets of data points are very similar.

To further demonstrate this similarity, Fig. 9(a,b) shows the decomposition in toroidal mode numbers of the magnetic spectrum obtained for #79237 using the signal channels selected for the AELM tracking algorithm, as processed using the different real-time and post-pulse op-

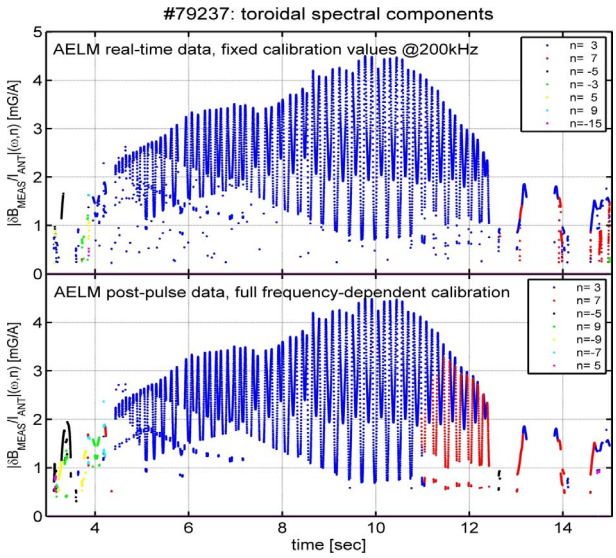

Fig. 9a Decomposition in toroidal mode number of the magnetic spectrum obtained for #79237 using the signal channels selected for the AELM tracking algorithm, as processed using the different real-time and post-pulse options. The seven components with the highest (post-pulse) amplitudes are shown, in decreasing order of amplitude. The *SparSpec-any* method was used for the real-time analysis, hence effectively "forcing" the detection at time t_{j+1} of the component that had already been detected at time t_j, even if the amplitude of that mode was not the highest at the time point time t_{j+1}. Therefore, we have almost continuous detection of an $n = 3$ mode, even if for time > 11 sec an $n = 7$ mode with similar amplitude also appear in the spectrum (note however that as $|\delta B(n = 7)|$ is always slightly below $|\delta B(n = 3)|$, also the *SparSpec-highest* algorithm would have continued locking onto the $n = 3$ mode). Conversely, for the post-pulse analysis all components above the noise level were obtained, included such $n = 7$ mode.

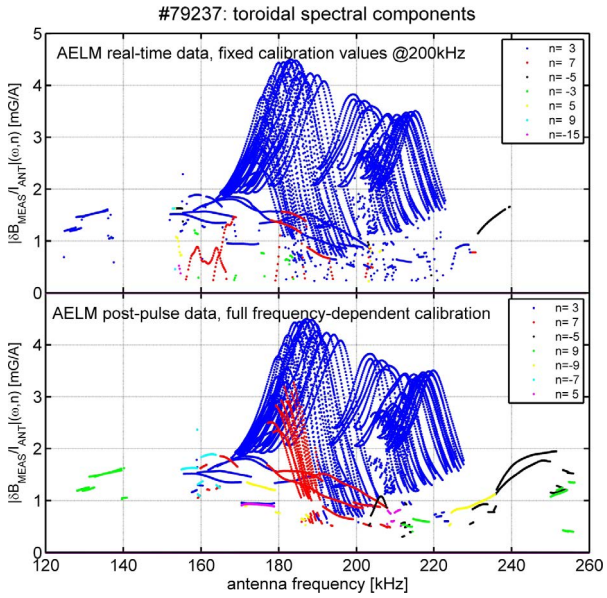

Fig. 9b Decomposition in toroidal mode number of the magnetic spectrum obtained for #79237, plotted as a function of the antenna frequency. With this representation we can see that the different processing options used for offset and calibration in the real-time and post-pulse analysis do not affect the results in the frequency range 150 < freq [kHz] < 240 and for medium toroidal mode number $|n| < 7$, but can cause rather different results for higher mode numbers, particularly at lower frequencies < 130 kHz and at higher frequencies > 250 kHz.

tions. Since the *SparSpec-any* method was used for the real-time analysis, we were effectively "forcing" the detection at time t_{j+1} of the component that had already been detected at time t_j, even if the amplitude of that mode was not the highest at the time point time t_{j+1}. Conversely, for the post-pulse analysis all components above the noise level are obtained. We note that both in real-time and post-pulse the $n = 3$ components is the dominant mode, with very similar amplitudes. The main advantage of the post-pulse analysis resides in the fact that it allows detecting the $n = 7$ mode which appears around the same frequency of the $n = 3$ mode in the time interval 11 < time[sec] < 12. This mode was not detected in real-time, and would have not been even if we had used the *SparSpec-highest* algorithm as its amplitude is slightly below that of the $n = 3$ mode. Despite the $|n| = 4$ toroidal pseudo-periodicity, we can be confident in the correctness of the real-time detection of this $n = 7$ mode for two reasons. Different sensors can be used in the post-pulse analysis: this $|n| = 4$ toroidal pseudo-periodicity is removed, and the $n = 3$ and $n = 7$ modes are again found with almost exactly the same frequency and damping rate (the difference being < 0.1 % for the frequency and around 5 % for the damping rate, well below the uncertainty on the measurements). The noise variance σ^2 is relatively small for the $n = 7$ component, hence the probability of false detection of such mode remains small.

Additionally, post-pulse analysis allows a more precise determination of the n-components when real-time tracking was not successful, for instance for time < 4.3 sec and time > 12.5 sec. Comparing the time (Fig. 9a) and frequency (Fig. 9b) representations of the toroidal mode number spectrum, we can confirm that the different processing options for the offset and calibration used for the real-time and post-pulse analysis do not affect the results in the frequency range 150 < freq[kHz] < 250 and for medium toroidal mode numbers $|n| < 7$, but can cause rather different results for higher mode numbers, particularly at lower frequencies < 130 kHz. Future upgrades of the AELM real-time software will have to take into account a frequency look-up table to use more precise offset and calibration values, so as to improve the real-time analysis in such cases.

Figures 10 and 11 show the real-time and post-pulse measurements of the mode frequency (f_{MEAS}) and damping rate (γ/ω) for #79237, respectively. Three different

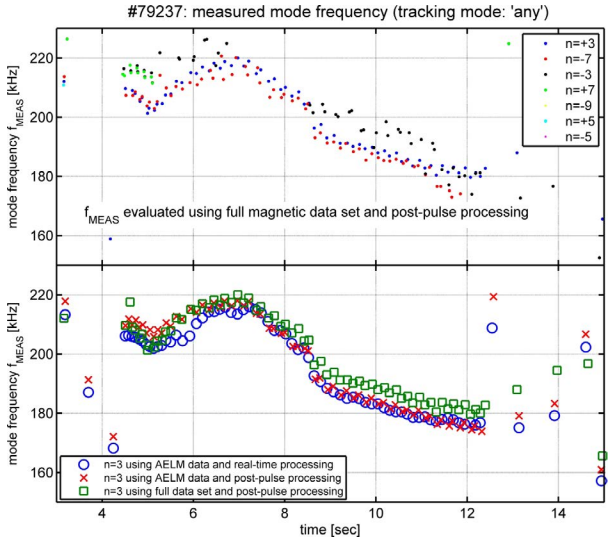

Fig. 10 The real-time and post-pulse measurements of the mode frequency f_{MEAS} for #79237. Three different calculations are compared: the value obtained directly in real-time with the AELM data and the processing options included in the AELM real-time software, the value obtained with the same AELM data but the post-pulse processing options, and finally the values obtained with the full magnetic data set (ten sensors, of which only seven were acquired in real-time) and the post-pulse processing options. As in real-time we can only track one single mode (predominantly the $n = 3$ for #79237), the bottom frame shows the comparison between the f_{MEAS} data for this mode, and the top frame show the f_{MEAS} data for the seven largest-amplitude modes as evaluated post-pulse.

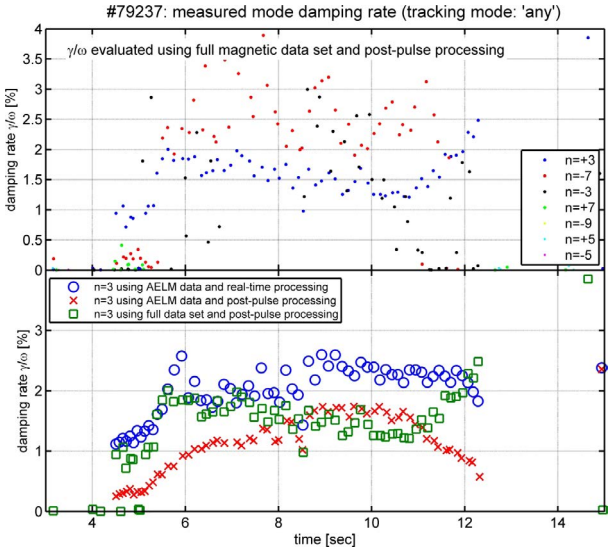

Fig. 11 The real-time and post-pulse measurements of the mode damping rate (γ/ω) for #79237, using the same format and data analysis methods as in Fig. 10.

calculations are compared: the value obtained directly in real-time with the AELM data and the processing options included in the AELM real-time software; the value obtained with the same AELM data but the post-pulse processing options; and finally the values obtained with the full magnetic data set (ten sensors, of which only seven were acquired in real-time) and the post-pulse processing options. It is important to remember that as in real-time we can only track one single mode (predominantly the $n = 3$ for #79237), the bottom frame in both Figs. 10 and 11 shows the comparison between the f_{MEAS} and γ/ω data only for this mode, whereas the top frame show the f_{MEAS} and γ/ω data for the seven largest-amplitude modes as evaluated post-pulse.

From Fig. 10, we note that the measured $n = 3$ mode frequency is indeed very similar when evaluated post-pulse and real-time, independently of the specific algorithm and input dataset used for this calculation. We also note that the frequency of the other modes with a different toroidal mode number are very close to that of the $n = 3$ mode, i.e. we do indeed measure a frequency-degenerate spectrum. This is also partially due to the fact that in this ohmically heated discharge the toroidal rotation frequency of the plasma is rather small, of the order of $f_{\mathrm{TOR}} \sim 1\,\mathrm{kHz}$, hence producing only a minimal frequency Doppler shift between the various n-modes, $\Delta f_{\mathrm{MEAS,n}} = n f_{\mathrm{TOR}}$.

From Fig. 11, we note that the measured $n = 3$ damping rate shows a very similar quantitative trend using both real-time and post-pulse calculations, i.e. it increases from time $= 4\,\mathrm{sec}$ to time $= 8\,\mathrm{sec}$, then it remains relatively constant, then it decreases from time $= 10.5\,\mathrm{sec}$ until time $= 12.5\,\mathrm{sec}$. In terms of absolute values, there is at most a ~30 % difference between the AELM real-time data (blue circles in the bottom frame of Fig. 11) and the post-pulse data obtained using the full magnetic data set and the post-pulse processing options (green squares in the bottom frame of Fig. 11). This difference is due to a non-optimal choice for the factor *Cdamp* that needs to be set pre-pulse in the AELM to convert the real-time calculation of the damping rate as estimated from the frequency width of the detected resonance into physical units. For #79237 a value *Cdamp* = 0.16 was selected, as previously determined for $n = 1$ TAEs [18]. Our recent analysis has determined that the optimal value for *Cdamp* depends on the actual mode number that has been measured in real-time, with a range *Cdamp* = 0.16 for $|n| = 1$ TAEs to *Cdamp* = 0.30 for $|n| > 7$ TAEs. Again, future upgrades of the AELM real-time software may need to take into account a mode number look-up table to use a more precise real-time/post-pulse conversion factor for the mode damping rate.

Figure 12 shows a second example for the real-time *SparSpec* analysis for the discharge #77790, where the *SparSpec-highest* method was used. In this discharge the AEAD system was configured to drive a $|n|$ = odd spectrum peaked towards $|n| = 7$ to $|n| = 11$, with a reduced antenna-drive for lower-$|n|$ and higher-$|n|$ modes. During this discharge an elongation scan was performed to study the dependence of the TAE damping rate on the edge shape (see [12, 22–24] for additional details

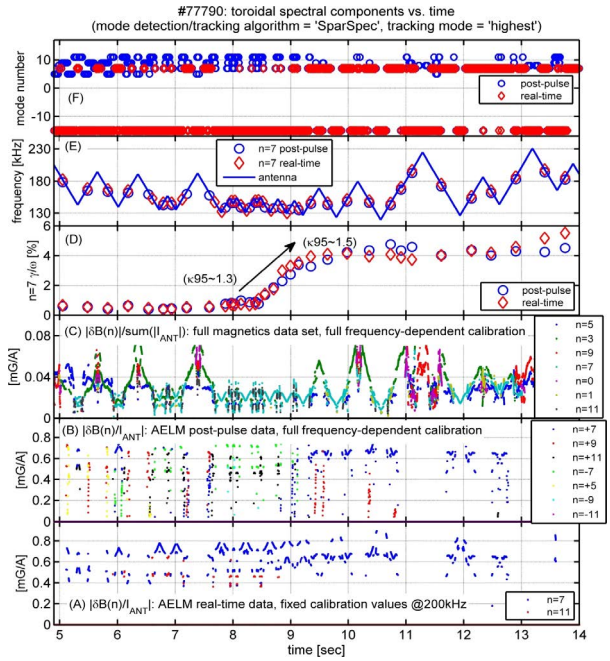

Fig. 12 The toroidal spectral components, and the measured mode frequency and damping rate for the $n = 7$ mode for #77790. Note the sharp increase in γ/ω between time = 8 sec and time = 10 sec, due to a corresponding increase in the edge elongation from $\kappa_{95} \approx 1.3$ to $\kappa_{95} > 1.5$ over the same time window. This feature is observed in both the real-time and post-pulse data, and demonstrates the capabilities of the *SparSpec* algorithm in detecting and following a mode even when its damping rate varies considerably over a relatively short time window.

on these studies). The *SparSpec-highest* algorithm was setup to look for co-current (positive toroidal mode numbers) propagating AEs with $3 \leq n \leq 11$ and counter-current (negative toroidal mode numbers) propagating AEs with $n = [-5, -7, -9]$. Five out of the eight magnetic sensors available in real-time were selected for the analysis, and the following run-time parameters were used (see Appendix B.5 for further details on the use of these parameters): *SS-*N_{MAX} = 20, *SS-Lambda* = 0.85, *SS-T'hold* = 5×10^{-10}, *Thresholds/Amp* = 1×10^{-10} [T/s], *Thresholds/Start* = *Thresholds/End* = 1×10^{-13} [T/A], *MinimumTwist* = 2 [rad].

In Fig. 12 we show the normalized amplitude of the seven largest toroidal spectral component that could have been selected for tracking in real-time (frame-A: only the $n = 7$ and $n = 11$ modes satisfy the detection criteria), and post-pulse applying the full-frequency dependent calibration, using only the five AELM magnetics selected for real-time analysis (frame-B) and the full complement of ten magnetic sensors available for post-pulse analysis (frame-C); frame-D and frame-E shows the damping rate and the mode frequency for the $n = 7$ mode evaluated in real-time and computed post-pulse using the data from frame-A and frame-C, respectively; finally, frame-F shows the toroidal mode number that was actually selected for tracking in real-time and post-pulse.

Comparing the data in frame-A to frame-C, we first note that when using only the magnetic sensors selected for real-time analysis but with the post-pulse processing options, we sometimes mistake the $n = +7$ mode for the $n = -7$ mode, notably in the time window $8 <$ time[sec] < 9. Again, this points to the need for a more accurate real-time implementation of the full frequency-dependent calibration. Moreover, whereas in real-time the $n = 7$ mode has almost always the largest amplitude, using the full complement of magnetic sensors and post-pulse analysis we find that this only occurs for frequencies around 150 kHz, and other modes appear at higher amplitudes in different frequency ranges. This translates into the fact that in real-time the $n = 7$ mode is the one that was almost always selected for tracking, whereas using post-pulse processing other, higher-$|n|$ modes could have been chosen, most notably the $n = 11$ mode during the time interval $7 <$ time[sec] < 10.5. Finally, when comparing the mode frequency and damping rate obtained in real-time and post-pulse, we find a very good agreement between these two sets of data, specifically for the sharp increase in γ/ω between time = 8 sec and time = 10 sec, due to a corresponding increase in the edge elongation from $\kappa_{95} \approx 1.3$ to $\kappa_{95} > 1.5$ over the same time window. This demonstrates the capabilities of the *SparSpec* algorithm in detecting and following a mode even when its damping rate varies considerably over a relatively short time window.

5.2 Real-time detection of Toroidal Alfvén Eigenmodes using the *SimpleSum* algorithm

The *SimpleSum* algorithm provides in most cases some basic capabilities for discriminating between different toroidal mode numbers according to the sensor and sign selection used to construct the signal on which mode detection and tracking is then performed in real-time. Figure 13(a-c) show three illustrative examples of the mode number detection that could in principle be obtained with the *SimpleSum* algorithm for modes whose frequency is exactly 200 kHz, i.e. at the value for which an exact calibration is obtained in real-time, and when all the three surviving pairs of sensors located at toroidally opposite positions are used. In Fig. 13(a-c), the bottom (A), middle (B), and top (C) frames show, respectively, the value of the *multiplier* used in the *SimpleSum* algorithm, the norm of the spectral window ($\|W(n)\|$) for the selected combination of sensors (for comparison with the *SparSpec* algorithm), and the probability of detecting each toroidal component in the antenna-driven radial B-field $B_{\mathrm{RAD}}(n)$, normalized so that $\max(\mathrm{prob}(B_{\mathrm{RAD}}(n))) = 1$. For the data presented in Fig. 13a, the *multiplier* values are $= [1, 1, 0, 1, 0, -1, -1, 0, -1, 0]$ and this in principle favours detection of modes with $|n| = 1$: note that for these modes

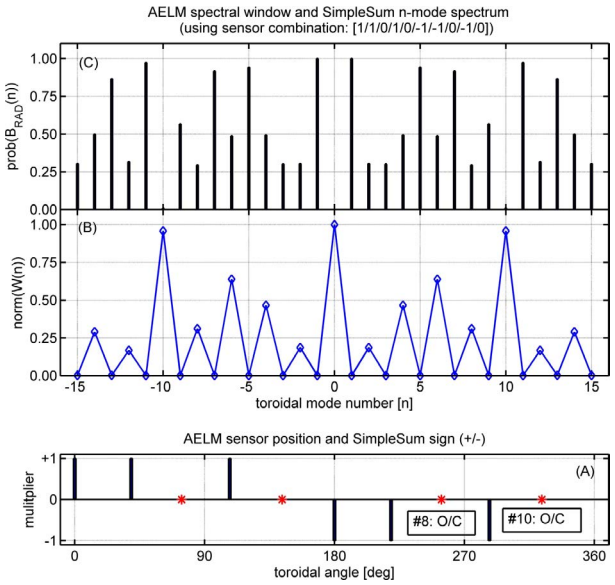

Fig. 13a The toroidal mode number spectrum of the radial component of the antenna-driven magnetic field that can in principle be detected using the *SimpleSum* algorithm to combine the signals from the three available pairs of sensors located at opposite toroidal positions, with *multipliers* = $[1, 1, 0, 1, 0, -1, -1, 0, -1, 0]$. This combination favours detection of $|n|$ = odd modes with low $|n| = 1$. We also show the norm of the spectral window $W(n)$, for comparison with the *SparSpec* algorithm. The high secondary lobe in $\|W(n)\|$ for $|n| = 10$ shows the pseudo-periodicity of the original full complement of 10 toroidal sensors, and additional pseudo-periodicities exist at $|n| = 2$, $|n| = 4$, $|n| = 6$ and $|n| = 8$ due to the specific sensors' selection. The position of the two sensors (#8 and #10) which are open-circuit (O/C) in-vessel is also shown.

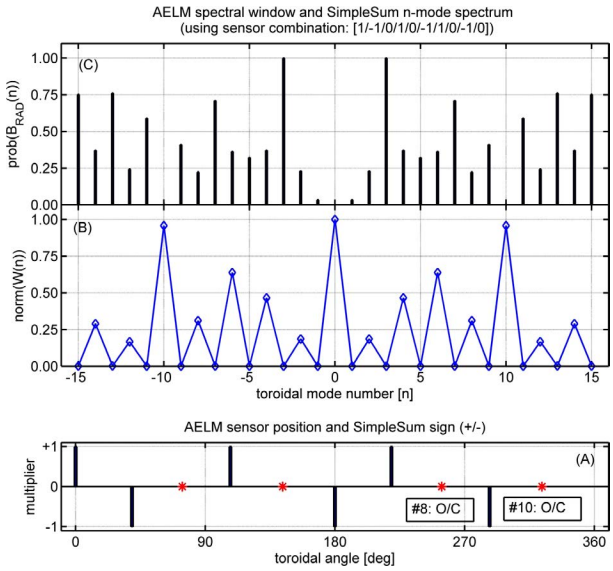

Fig. 13b Plotting the same data as in Fig. 13a, this time using multipliers = $[1, -1, 0, 1, 0, -1, 1, 0, -1, 0]$ and the same subset of sensors located at toroidally opposite positions. This *SimpleSum* combination favours detection of $|n|$ = odd modes with a slightly higher toroidal mode number $|n| = 3$. The pseudo-periodicities at $|n| = 2$, $|n| = 4$, $|n| = 6$ and $|n| = 8$ are the same as in Fig. 13a.

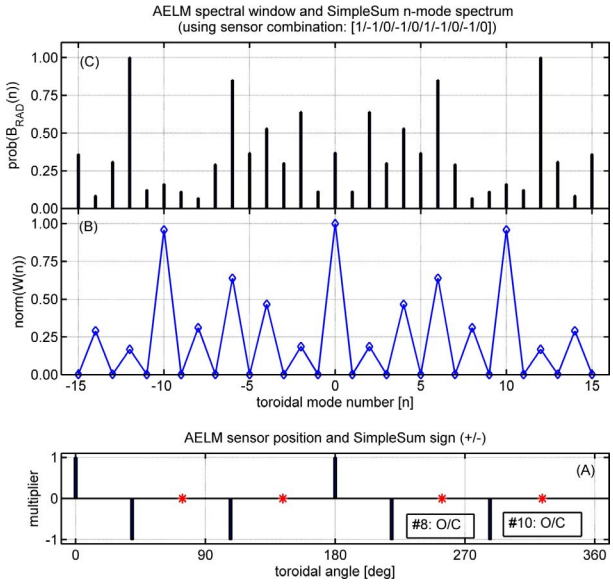

Fig. 13c Plotting the same data as in Fig. 13a, this time using multipliers = $[1, -1, 0, -1, 0, 1, -1, 0, -1, 0]$ and the same subset of sensors located at toroidally opposite positions. This *SimpleSum* combination favours detection of $|n|$ = even modes with a rather high toroidal mode number $|n| = 12$. The pseudo-periodicities at $|n| = 2$, $|n| = 4$, $|n| = 6$ and $|n| = 8$ are the same as in Fig. 13a.

$\|W(n)\| = 0$, i.e. detection is greatly facilitated. The second highest detection probabilities are obtained for $|n| = 5$ and $|n| = 7$, but these values are those of modes that correspond to the sum of $|n| = 1$ (for which detection probability is maximum) with $|n| = 4$ and $|n| = 6$, respectively, i.e. the periodicities given by two of the highest secondary lobes in the spectral window. Similar considerations also apply for Fig. 13b and Fig. 13c, where the *multiplier* values are = $[1, -1, 0, 1, 0, -1, 1, 0, -1, 0]$ and = $[1, -1, 0, -1, 0, 1, -1, 0, -1, 0]$, respectively: these combinations in principle favour detection of toroidal components with $|n| = 3$ and $|n| = 12$, respectively,

One of the most frequent applications of the *SimpleSum* algorithm is using only one pair of sensors at toroidally opposite positions, as the various issues with the signal calibration when the mode frequency is not exactly = 200 kHz are somewhat reduced. Figure 13d shows one example of such analysis, where the *multiplier* values are = $[1, 0, 0, 0, 0, \pm 1, 0, 0, 0, 0]$: note the clear $|n| = 2$ pseudo-periodicities in the detection and the many peaks in prob($B_{\mathrm{RAD}}(n)$), which clearly indicates that actual mode number discrimination is not particularly satisfactory in real-time using this particular detection scheme for the *SimpleSum* algorithm. Finally, Fig. 13e shows some examples of the toroidal mode number spectra that can in

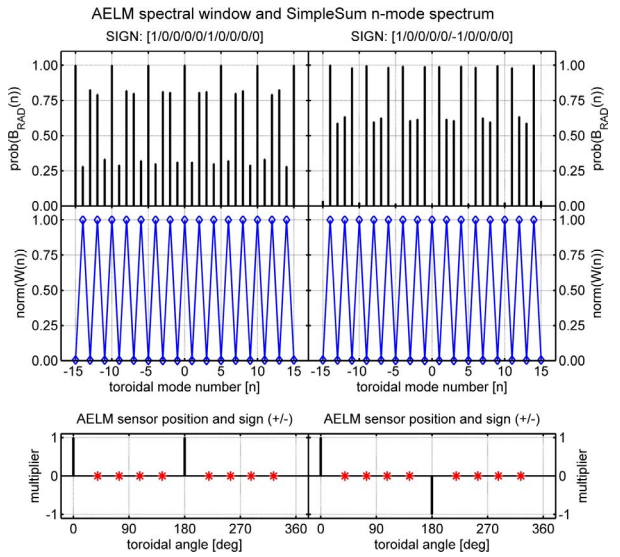

Fig. 13d The toroidal mode number spectra that can be detected using only one pair of sensors at toroidally opposite positions, with ± phasing combinations. Note the many peaks in prob($B_{\mathrm{RAD}}(n)$) and the clear $|n| = 2$ pseudo-periodicities in the detection, illustrating the fact that efficient mode number discrimination cannot actually be obtained in real-time with this scheme.

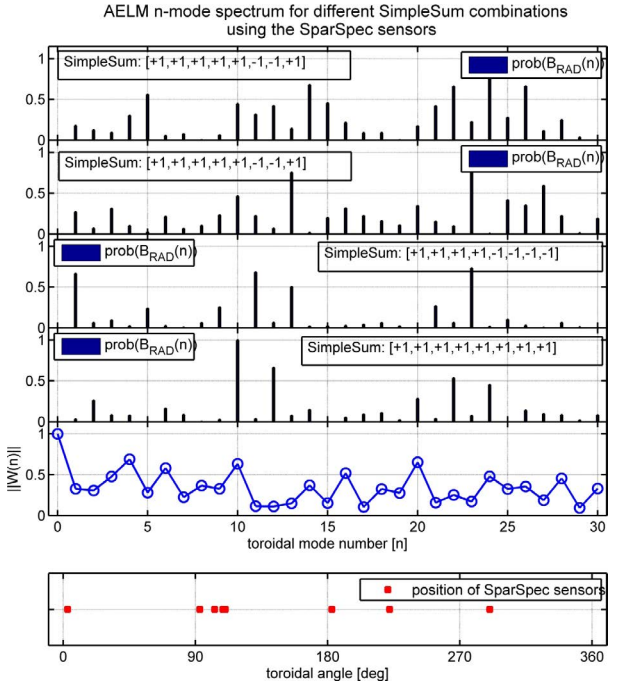

Fig. 13e Some examples of the toroidal mode number spectra that can in principle be detected in real-time using the *SimpleSum* algorithm when all the eight sensors that can be used for the *SparSpec* algorithm are selected, with different phasing ± combinations. Note that as in this case the sensors are neither toroidally equi-spaced nor arranged in pairs at opposite toroidal locations, there is no obvious scheme for selecting/removing specific toroidal components.

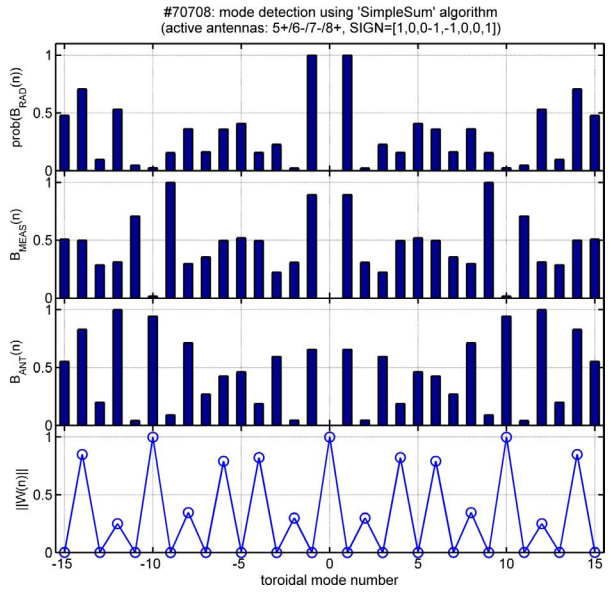

Fig. 14a Mode driving and detection scheme for #70708 using the *SimpleSum* algorithm with two pairs of sensors at opposite toroidal locations, with sign combination = $[1, 0, 0, -1, -1, 0, 0, 1]$, and four active antennas in the same octant, with alternate phasing. Here prob($B_{\mathrm{RAD}}(n)$) is therefore the product of the detection probability using the selected set of magnetic sensors $B_{\mathrm{MEAS}}(n)$ with the antenna-driven radial magnetic field spectrum $B_{\mathrm{ANT}}(n)$, i.e. prob($B_{\mathrm{RAD}}(n)$) = $B_{\mathrm{ANT}}(n) * B_{\mathrm{MEAS}}(n)$.

principle be detected using the *SimpleSum* algorithm when all the eight sensors that can be used in real-time for the *SparSpec* algorithm are selected, with different phasing ± combinations. In this case the sensors are neither toroidally equi-spaced nor all arranged in pairs at opposite toroidal locations, hence there is no obvious scheme for selecting or removing specific toroidal mode numbers. Conversely, the spectral window does not have secondary lobes with very large amplitude $||W(n_{\mathrm{LOBE}})|| > 0.7$ for $|n_{\mathrm{LOBE}}| < 5$, which facilitates the analysis as the low-$|n|$ periodicities are the most difficult to consider appropriately.

One example of real-time detection and tracking of stable AEs using the *SimpleSum* algorithm is observed during the discharge #70708, where four antennas in the same octant were used with alternate phasing [5+/6−/7−/8+] to drive a double-hump spectrum with largest amplitudes around $1 \leq |n| \leq 3$ and $7 \leq |n| \leq 12$. Figure 14a shows the mode driving and detection scheme for #70708: here prob($B_{\mathrm{RAD}}(n)$) = $B_{\mathrm{ANT}}(n) * B_{\mathrm{MEAS}}(n)$ is given as the product of the nominal antenna-driven radial magnetic field spectrum $B_{\mathrm{ANT}}(n)$ (here calculated at the last closed flux surface assuming unitary and equal current in all active antennas) with the detection probability using the selected set of magnetic sensors $B_{\mathrm{MEAS}}(n)$. Figure 14b then shows the tracking results for #70708 using the *SimpleSum* algorithm with two pairs of sensors at opposite toroidal locations, with sign combination = $[1, 0, 0, -1, -1, 0, 0, 1]$, and

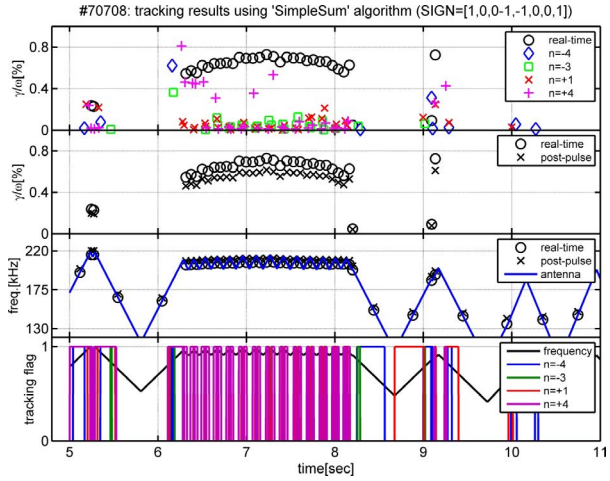

Fig. 14b Tracking results for #70708 using the *SimpleSum* algorithm with two pairs of sensors at opposite toroidal locations, with sign combination = $[1,0,0,-1,-1,0,0,1]$. As no mode number selection can be performed with the *SimpleSum* algorithm, the real-time measurement of the damping rate gives in this case the convolution of the damping rate for all the individual mode numbers detected using post-pulse analysis.

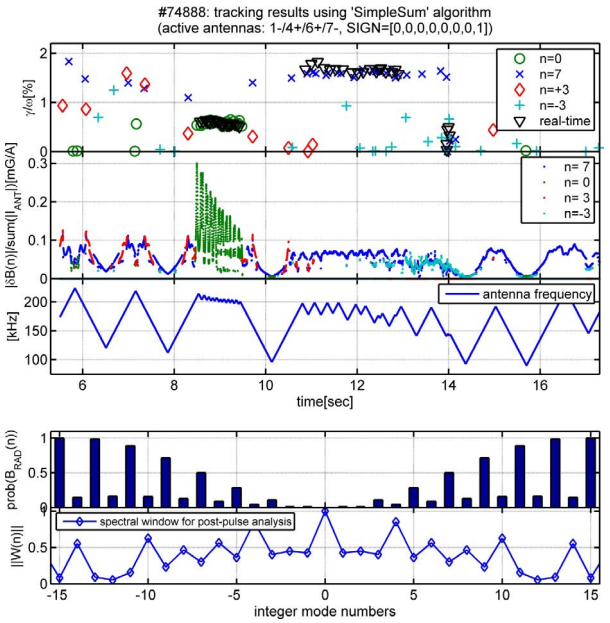

Fig. 15 Tracking results for #74888 using the *SimpleSum* algorithm with just one sensor. Note that in this case when there is a dominant single mode in the measured $|\delta B(n)/I_{\mathrm{ANT}}|$ spectrum, the real-time measurement of the damping rate is in good agreement with the post-pulse evaluation. This occurs, for instance, for the $n = 0$ mode in the time interval $8.2 < \mathrm{time\,[sec]} < 9.7$, and for the $n = 7$ mode in the time interval $10.8 < \mathrm{time\,[sec]} < 14$, although in this case the presence of an $n = 3$ and $n = -3$ frequency-degenerate modes affect the real-time measurements of γ/ω for certain time points.

selecting *MinimumTwist* = 1.5 [rad], *Thresholds/Amp* = 1×10^{-9} [T/s], *Thresholds/Start* = 1×10^{-11} [T/A] = *Thresholds/End*. As no actual mode number selection can be performed in real-time with the *SimpleSum* algorithm, the real-time measurement of the damping rate gives in this case the convolution of the damping rates for all the individual mode numbers which are detected using the post-pulse analysis. Consistently with the value of $\mathrm{prob}(B_{\mathrm{RAD}}(n))$, the mode numbers that are found in the real-time spectrum are $n = -4$, $n = -3$, $n = +1$ and $n = +4$; however, the $n = -1$ mode does not appear in the measured B_{RAD} spectrum, despite having one of the largest detection probabilities. The value of the mode frequency evaluated in real-time is in very good agreement with the post-pulse calculation. Similarly, the damping rate evaluated in real-time when tracking was being obtained, i.e. for $6.2 < \mathrm{time[sec]} < 8.2$, is in good agreement with the convolution of the damping rates for the four individual modes $n = -4, -3, +1, +4$, but is much larger than those for each individual mode number when considered separately. This indicates that the *SimpleSum* algorithm can be used in real-time to find qualitatively the stable plasma resonances corresponding to the antenna-driven modes, but only post-pulse analysis can provide a precise estimation of the damping rate for each individual mode number when the antenna spectrum is frequency-degenerate.

Finally, the tracking results for #74888, shown in Fig. 15, prove that when there is a single dominant mode in the detected $|\delta B(n)/I_{\mathrm{ANT}}|$ spectrum, the real-time measurement of the damping rate obtained with the *SimpleSum* algorithm is in good agreement with the post-pulse evaluation. For this discharge, the antenna excitation spectrum was [1−/4+/6+/7−], one single sensor was selected for real-time analysis, hence $\mathrm{prob}(B_{\mathrm{RAD}}(n)) = B_{\mathrm{ANT}}(n)$ as $B_{\mathrm{MEAS}}(n) = 1$ in real-time for all toroidal mode numbers. Note that despite the very small, almost negligible nominal drive for the $n = 0$ mode when assuming a unitary and equal current in all active antennas, mismatch in the actual time-dependent antenna currents produce a non-negligible amplitude for this mode in certain frequency intervals. Furthermore, this $n = 0$ mode corresponds to a global plasma resonance, called the Global Alfvén Eigenmode [30], which then appears with a rather significant amplitude in the measured spectrum. For this discharge, the real-time and post-pulse estimates of the damping rate are in very good agreement for the $n = 0$ GAE in the time interval $8.2 < \mathrm{time[sec]} < 9.7$, and for the $n = 7$ TAE in the time interval $10.8 < \mathrm{time[sec]} < 14$, respectively, i.e. when these two modes are so largely dominant in the measured magnetic spectrum that it becomes possible to neglect its frequency degeneracy. Finally, note also that in the latter case, the presence of a background $n = 3$ and $n = -3$ frequency-degenerate TAEs affects the real-time measurements of γ/ω for the $n = 7$ TAEs at certain specific time points.

6. Accuracy of the *SimpleSum* and *SparSpec* Algorithms

The measurements of the mode frequency, damping rate, amplitude and toroidal mode number obtained in real-time with the AELM are then passed to the JET RTDN system, and their estimate has previously been shown to be in very good agreement with the result obtained with a more detailed post-pulse analysis [12–15, 18]. This, in principle, allows the implementation of a control system and feedback for the modes, so that when the mode's damping rate reduces excessively and approaches the marginal stability limit $\gamma/\omega = 0$, a control parameter (for instance the edge elongation in the case of low-n and medium-n TAEs [12, 18]) can be changed to bring the plasma back to a situation where it is further away from the marginal stability limit. To confirm the earlier results in a more systematic way, Fig. 16 shows the results of a statistical analysis of the accuracy of the real time measurements of the toroidal mode number, mode frequency, mode damping rate and mode amplitude obtained with the *SparSpec* algorithm. This data was obtained from analysing around 200 different JET discharges where various setups for the *SparSpec* tracking algorithm within the AELM, and different antenna excitation spectra were used, for deuterium, hydrogen and Helium4 plasmas. For this analysis: we consider the expected values for the mode number, frequency and damping rate as the ones given by the post-pulse analysis performed using the full magnetic dataset; use the actual post-pulse estimation on the errors on these quantities, which are typically around the values presented in [12] and define a *confidence level* for the real-time measurements as:

$$\text{confidence level} = \exp\left(-\frac{(RealTimeData - PostPulseData)^2}{(VariancePostPulseData)^2}\right). \quad (5)$$

This *confidence level* is a very stringent criterion to measure whether the scientific requirements for the AELM are satisfied since it has been previously demonstrated that the post-pulse analysis meets its intended accuracy [12–15, 18]. Therefore, a normal distribution of the real-time data centred on an expectation value provided by the post-pulse data, with the variance taken as the variance on such data, i.e. precisely as the one given in Eq. (5), also meets the requirements for post-pulse analysis provided we achieve a confidence level in excess of $e^{-1/4} = 0.7788$. This value for the confidence level is obtained when the absolute difference between real-time and post-pulse data is less than half the variance on the post-pulse data. As the measurement requirements for real-time analysis, given at the end of Sec. 3, and particularly the requirements on their accuracy (which is measured by the variance on the data), are in fact less stringent that those set for the post-pulse analysis, satisfying the criterion given in Eq. (5) is actually a more stringent test for real-time analysis than simply satisfying the measurement requirements.

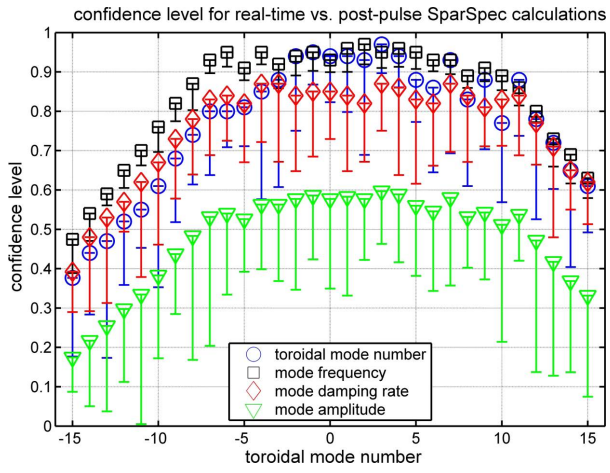

Fig. 16 Confidence level for the real-time evaluation of the mode number, frequency and damping rate when using the *SparSpec* tracking algorithm. We note that the confidence level is quite high for all the real-time data (but for the (less important) absolute mode amplitude) obtained with the *SparSpec* tracking algorithm, at least exceeding 0.8 for all mode numbers in the range $-8 < n < 8$. This confidence level drops below the very reliable value 0.8 for modes with higher toroidal mode numbers $|n| > 10$: this is due to the rather small number of components (SS-N_{MAX} = 30) that can be used in real-time to deconvolve the measured magnetic spectrum using the *SparSpec* real-time mode detection algorithm.

The *confidence level* data obtained with this procedure is shown in Fig. 16 as a function of the toroidal mode number for all the modes in the range $|n| \leq 15$. In Fig. 16 the vertical error bar on the confidence level data indicates the scatter in this quantity across the database. For the damping rate measurement, this scatter includes two separate test-cases: (a) when the real-time damping rate was calculated using the best available estimate for the *Cdamp* factor, and (b) when the (not always optimal) value of *Cdamp* = 0.16 was used.

Apart from the absolute mode amplitude, we note that the confidence level is quite high for all the real-time data obtained with the *SparSpec* tracking algorithm, at least exceeding 0.8 for all mode numbers in the range $-8 < n < 8$. This proves that the *SparSpec* real-time mode detection and tracking algorithm can indeed be used as a valuable diagnostic tool for blind and unsupervised mode discrimination and tracking in a multi-components and frequency-degenerate spectrum. Regarding the real-time measurement of the absolute mode amplitude, we note, first, that the LS re-estimation of the mode amplitude is not performed in real-time (due to CPU and RAM limitations, as already indicated before). Second, obtaining only a factor two accuracy in this quantity is not problematic per se, as what matters most is, in fact, the relative variation during a discharge, which is measured much more accurately.

It is in fact the relative variation in the mode amplitude, i.e. its trend over the discharge evolution, and not the absolute mode amplitude (which may be affected by signal offset, temporal drifts and calibration issues, and shows a very sensitive dependence on the background plasma conditions), which can be used to reliably determine whether a mode becomes a potential danger when combining this real-time estimation of the mode amplitude trend with the real-time damping rate measurements and, possibly, previous knowledge of the mode stability diagram.

The confidence level for the mode frequency, mode damping rate and mode number drops below the very reliable value 0.8 for modes with higher toroidal mode numbers $|n| > 10$. This is actually due to the rather small number of components ($SS\text{-}N_{\mathrm{MAX}} = 30 = 2 \times N_{\mathrm{MAX}}$) that can be used in real-time to deconvolve the measured magnetic spectrum using the *SparSpec* algorithm. Indeed, using post-pulse analysis (see Appendix-A for more details), it can be demonstrated that the optimum value for $SS\text{-}N_{\mathrm{MAX}}$ should be taken as $SS\text{-}N_{\mathrm{MAX}} \geq 3 \times N_{\mathrm{MAX}}$, i.e. $SS\text{-}N_{\mathrm{MAX}} \sim 50$ for real-time analysis, but this typically requires around three times as much CPU and RAM resources than using the standard real-time value $SS\text{-}N_{\mathrm{MAX}} = 2 \times N_{\mathrm{MAX}} = 30$. Future upgrades of the AELM hardware, for instance code parallelization over various CPUs with shared RAM resources, could be considered to take into account a larger number of base functions for the real-time *SparSpec* algorithm if an improved accuracy for the determination of the mode characteristics for $|n| > 10$ AEs is required.

Figure 17 shows the confidence level for the real-time evaluation of the mode number, frequency and damping rate when using the *SimpleSum* tracking algorithm, again combining results from around 100 different JET experiments. As the *SimpleSum* tracking algorithm does not in principle provide any n-number determination, we have considered that the "real-time" value of the toroidal mode number is given by the lower-$|n|$ mode for which prob($B_{\mathrm{RAD}}(n)$) is highest. If this n-mode is then also found in the post-pulse analysis, we take that it has been "correctly" determined in real-time. Two separate cases are shown. The first case corresponds to the situation obtained when the individual mode amplitude $|\delta B(n)|$ is large, $|\delta B(n)| > 0.7 \times \sqrt{\sum_n (|\delta B(n)|^2)}$, so that the detected mode spectrum can be considered to be made up by just one single component. The second case corresponds to the situation obtained when the mode amplitude is small, $|\delta B(n)| < 0.3 \times \sqrt{\sum_n (|\delta B(n)|^2)}$, and in this case resolving the frequency-degeneration of the mode spectrum becomes the most important factor for obtaining accurate real-time estimates. In all cases the mode frequency obtained with the *SimpleSum* algorithm is very accurate, as this quantity is determined to first order by the width of the frequency sweep when in tracking mode, which does not change much when using the *SimpleSum* or *SparSpec* algorithms. We also note that for an almost-pure single mode spectrum

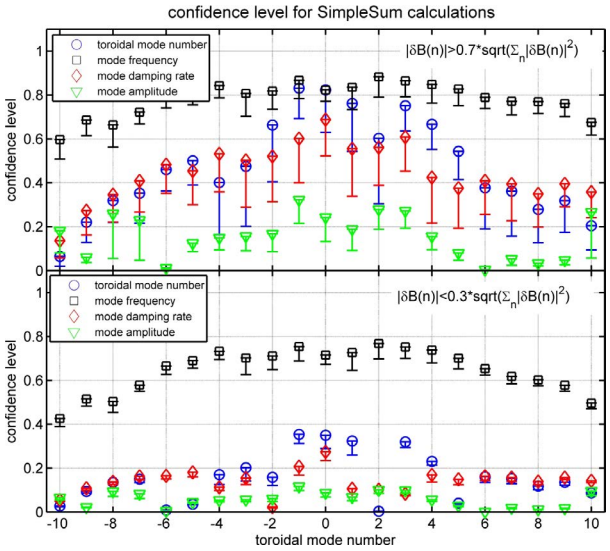

Fig. 17 Confidence level for the real-time evaluation of the mode number, frequency and damping rate when using the *SimpleSum* tracking algorithm. Two cases are shown: when the individual mode amplitude $|\delta B(n)|$ is rather large (top frame), so that the detected spectrum can be considered to be not frequency-degenerate, and when the mode amplitude is small (bottom frame), and in this case resolving the frequency-degeneration of the mode spectrum becomes the most important factor for obtaining accurate real-time estimates.

(top frame in Fig. 17), the *SimpleSum* algorithm is sufficiently accurate in determining at least the mode number and damping rate for lower-$|n|$ modes, up to $|n| = 3$ to $|n| = 5$ typically. Conversely, when the mode spectrum is frequency-degenerate (bottom frame in Fig. 17), in most cases only the mode number for low-$|n| \leq 1$ modes can be determined with sufficient accuracy.

7. Summary and Conclusions

In this work we have reported on the application of a new method for the unsupervised real-time detection and decomposition of a multi-harmonic and degenerate spectrum of high-frequency magnetic instabilities measured on the JET tokamak. This method uses real-time measurements performed on a 1 ms time base, which are then processed by a dedicated VMEbus-based real-time system communicating with the other JET real-time control systems.

The main application of the AELM system has so far been the detection and tracking of Alfvén Eigenmodes with toroidal mode numbers up to $|n| \leq 15$, i.e. those that could be driven unstable by fusion-born alpha particles in future burning plasma experiments such as ITER. In our work we have used two different real-time mode detection, discrimination and tracking algorithms for these AEs. The first one, the so-called *SimpleSum* algorithm, is based on a simple linear combination of the signals from a selected set

of magnetic probes. The second and most successful one, the so-called *SparSpec* algorithm, is based on a novel application of the Sparse Representation of signals, derived from its original applications to astronomical data via the *SparSpec* code.

By appropriate setting of its run-time parameters, the *SimpleSum* algorithm allows in principle a basic discrimination between $|n|$ = odd and $|n|$ = even modes, and also between low-$|n|$ and high-$|n|$ modes. However, a correct mode discrimination relies on all magnetic signals having exactly the same end-to-end transfer function at all frequencies and the magnetic probes being located at exactly periodic spacing, which is not the case for the JET set of magnetic probes. Moreover, the measured AE spectrum is frequency-degenerate, i.e. multiple modes with different toroidal mode number have very close-by frequencies and similar amplitude, and this further complicates the analysis using the *SimpleSum* algorithm. Therefore, whereas the mode frequency is in general correctly determined within the required measurement accuracy, we find that the mode number and the damping rate can typically be correctly determined for low-$|n|$ modes up to $|n| \leq 3$ to $|n| = 5$ only for an almost-pure single mode spectrum. Conversely, when the mode spectrum is frequency-degenerate, i.e. showing multiple components with similar amplitudes, in most cases only the mode number for low-$|n| \leq 1$ modes can be determined with sufficient accuracy.

The real-time (and post-pulse) implementation of the *SparSpec* algorithm at JET has allowed a complete, accurate and numerically efficient analysis of these measurements, which would have not otherwise been possible. Using the rather modest computational resources allocated to the real-time analysis of the AEAD data within the AELM hardware and software (a 1 GHz PowerPC with a 512 MB RAM running on a 1 kHz clock-rate), the multi-components, frequency-degenerate antenna-driven spectrum can be fully resolved within typically ~650 μs for each 1 ms clock-cycle. The results for the mode frequency, damping rate, mode numbers (and scaled mode amplitude) obtained with the real-time *SparSpec* algorithm are in good agreement both statistically and on a shot-by-shot basis with those obtained with the post-pulse implementation of this algorithm. This confirms that the *SparSpec* algorithm as implemented in the AELM software can indeed be valuably used for real-time analysis of MHD instabilities for plasma control purposes.

As shown in Appendix A.4, using the *SparSpec* algorithm we obtain the required confidence level in the real-time measurements of all the mode characteristics (frequency, toroidal mode number, mode amplitude and damping rate) with the required time resolution and within the computational time limits by using the run-time analysis parameters that have been determined from an extensive off-line optimization of the original post-pulse version of the *SparSpec* code. This optimization has been performed using actual data obtained during JET tokamak discharges and simulated data. The three main features of the *SparSpec* algorithm that need to be optimized for real-time analysis are the number of input channels used in the analysis, the size of the dictionary used to model the input dataset, i.e. the parameter *SS-*N_{MAX}, and the penalization value in the L1-norm minimization, i.e. the parameter *SS-Lambda*. We find that the real-time calculations can achieve the required confidence level > 0.8 with the required time resolution of 1 ms and within the computational time limits of $< 850\,\mu$s by using as input the data from between five and seven magnetic sensors, setting *SS-*$N_{\mathrm{MAX}} = 30$ and *SS-Lambda* = 0.85.

To improve on the current capabilities of the AELM system at JET, various options could be considered for a hardware-based evolution of the real-time implementation of the *SparSpec* code, the more intuitive one being parallelization. In this respect, the *SparSpec* algorithm used in the AELM would lend itself rather easily to being parallelised but there would be limits to the performance gains in its current form. It would also be possible to run several versions of the code with different sets of data and/or parameters all running in parallel contributing to a final selection process. Then the only limit would be the number of processors available. Considering that a modern graphics card costing less than 500 USD has more than 500 cores on it, it is immediate to see the attractiveness of such a solution.

For parallelisation of the *SparSpec* code, the following aspects should be considered:

1. compiler optimisations → compilers are getting better at this kind of optimisation but they still need significant hints from an experienced programmer;
2. distinct separate serial code sections → often there are sequences of unrelated code executed serially that can be executed simultaneously on several CPUs;
3. algorithm summations and other code loops → commands repeatedly executed within a loop can be spilt across 2 or more CPUs, with each CPU executing a different section of the loop (for instance: summations are basically loops . . .);
4. dealing with shared resources → most commonly a variable that needs to be written by at least 1 CPU and read by 1 or more other CPUs;
5. communication of initial values and results → most commonly the CPUs will share memory but some systems use fast networks so care must be taken not to move large chunks of data around that will negate any execution speed increases;
6. synchronisation → results may require complex synchronisation algorithms especially when outputs from one calculation are inputs to another or a result is a summation;
7. speculative execution → all possible branches of a decision are executed in advance of the logic that decides which path to choose.

A number of possible upgrades to the real-time *SparSpec* algorithm are also being considered for future work at JET:

1. a frequency look-up table to use more precise offset and calibration values;
2. a mode number look-up table to use a more precise real-time/post-pulse conversion factor for the mode damping rate;
3. a larger number of base functions for the real-time *SparSpec* algorithm if an improved accuracy for the determination of the mode characteristics for $|n| > 10$ AEs are required;
4. implementation of a "memory" in the *SparSpec* algorithm, whereby the solution obtained for the previous time point T_n is taken as the initial guess for the solution at the current time point T_{n+1}.

This particular point is now under test offline, with a view to perform the first tests online during the JET experimental campaigns foreseen for 2013. Two main aspects of such a memory scheme are being investigated, i.e. the correct implementation of error tracking and a relaxation procedure whereby not only the last time point T_n is used, but all previous time points T_j, $j = 1 \rightarrow n$, with a weight related not only to the temporal distance between the "historical" time points T_j and the current time point T_{n+1}, but also to the accuracy of the real-time calculation for any time point T_j.

Considering now fusion devices where operation and, more specifically, advanced real-time MHD control, has recently started or is being envisaged, it would be intriguing to test the capabilities of the *SparSpec* algorithm for stellarators such as the Large Helical Device, where various types of fast-ion driven AEs have been measured and in some instances considered responsible for enhanced transport [31]. The very different magnetic geometry and the different spectral boundary conditions, namely the presence of an additional quantum number for the mode (radial mode number) to the two in tokamaks (toroidal, poloidal), and the absence of a generalized toroidal symmetry, make the efficient application of a Sparse Representation method, such as the *SparSpec* code, a very challenging problem. Newer tokamak devices such as KSTAR [32] and HL-2A [33] could also benefit in their next phase of operations from active suppression of multi-harmonics MHD modes with reliable actuators as simulations may not be able to provide on their own all the required control answers. The experience acquired at JET with the AELM and the *SparSpec* code may then be used to prepare and test offline dedicated control systems.

Finally, for JET, and more generally for future burning plasma experiments such as ITER, and other devices with very energetic ions such as JT60-SA [34], where integrated regimes are being studied to optimize the plasma performance in the presence of significant populations of fast ions, further applications of this new method based on the Sparse Representation of signals open interesting and very useful perspectives for the concurrent real-time detection and control of different MHD instabilities, even when occurring at close-by frequencies, as these can be discriminated very accurately. This approach allows specifically tailored control schemes to be put in place for each individual mode, hence improving the overall control of the plasma operation and fusion performance. This will be particularly important for forthcoming experiments approaching the burning plasma conditions, and in preparation for ITER [35] and DEMO [36] activities, where real-time control of the stability of the fusion born alphas in the background "sea" of MHD modes that are expected to occur in such conditions, represents one of the key ingredients required to achieve a net fusion energy gain.

Acknowledgments

This work, supported by the European Communities under the contract of Association between EURATOM and CRPP-EPFL, was carried out within the framework of the European Fusion Development Agreement. The views and opinions expressed herein do not necessarily reflect those of the European Commission. This work was also partly supported by the Swiss National Science Foundation. The Authors would like to thank the various members, past and present, of the CRPP, MIT and JET staff that have contributed to the design, installation, commissioning and operation of the Alfvén Eigenmode Active Diagnostic system over more than 10 years of experiments at JET. The Authors would also like to thank the Reviewers for their constructive comments.

Appendix A. Sparse Representations and the *SparSpec* Code

The problem of detection and discrimination between the individual components in a multi-harmonic spectrum which is un-evenly sampled in the spatial domain is common to various fields of physics and engineering [37]. Historically, this problem has been addressed using methods essentially based around the Lomb-Scargle periodograms [38–41], and much work has been performed to improve on the limitations of these original methods, essentially in the field of Astronomy and Astrophysics (A&A). This general measurement problem is further complicated in thermonuclear fusion plasmas, and specifically in large-scale tokamak and stellarator devices, by the (often very) low number of measurement points in the spatial domain, which is due to in-vessel engineering and installation constraints, leading to a number of mathematical difficulties. Therefore, analysis method based on the spatial Nyquist criterion cannot in general be used because of the effect of aliasing, particularly if intermediate to large mode numbers need to be resolved.

This has prompted the development and the application of various methods to the analysis of MHD data in

thermonuclear fusion plasmas, such as the Singular Value (SVD) [42, 43] and the wavelet [44] decomposition, the Wigner [45], Choi-Williams [46] and Hilbert [47] Transforms, and a generalization of the Lomb-Scargle periodograms [48]. However, none of these methods can be efficiently used for the decomposition of a frequency-degenerate, stable spectrum of MHD modes with the aim to measure their damping rate, because of their mathematical limitations and computational requirements, particularly when real-time, sub-millisecond calculations are needed.

Conversely, a method based on the Sparse Representation of Signals, as implemented in the *SparSpec* code (freeware available at: http://www.ast.obs-mip.fr/article123.html) [26, 27] has been demonstrated to efficiently and correctly perform the post-pulse [28] and real-time [14, 15] blind and unsupervised signal decomposition of data which are unevenly sampled in the spatial domain using a (very) small number of measurement points.

A.1 Problem statement: astronomy and astrophysics

The A&A problem consists in the analysis of time-series: these can be, for instance, light curves or radial velocity measurements, which are subject to observational constraints, such as day/night alternation and meteorological conditions. The A&A measurements are therefore always obtained through irregular sampling. An example of such a data set is provided in Fig. A1, showing the observation for the radial velocity curve of the Herbig Ae star HD 104237, obtained over five observing nights of high resolution spectroscopy at the South African Astronomical Observatory during April 1999 [49, 50].

In A&A data analysis (as for MHD analysis in thermonuclear fusion plasmas), the main objective is that of looking for periodicities. For the case of variable stars, and multiple star systems, there are several oscillation modes, some of which are related to the stars' orbits and have to be filtered out when oscillations in other quantities are sought. This leads to the estimation of spectral lines from the data. The irregular sampling can be seen as the application of an irregular Dirac comb $w(t)$ to the original signal $y(t)$ and can be well understood in the Fourier domain:

$$y_s(t) = \begin{cases} \sum_{p=1}^{P} y(t_p)\delta(t-t_p) \overset{\text{Fourier Transform}}{\longrightarrow} Y_s(f) \\ = \int_{-\infty}^{\infty} y_s(t) e^{-2j\pi ft} dt = \sum_{p=1}^{P} y(t_p) e^{-2j\pi ft_p} \\ y(t) \times \underbrace{\sum_{p=1}^{P} \delta(t-t_p)}_{w(t)} \overset{\text{Fourier Transform}}{\longrightarrow} Y_s(f) \\ = Y(f) \star \underbrace{\sum_{p=1}^{P} e^{-2j\pi ft_p}}_{W(f)} . \end{cases} \quad \text{(A.1)}$$

In Eq. (A1), $y(t_p)$ are the individual measurements taken at all the time points t_p (for $p = 1, \ldots, P$), from which the irregularly sampled signal $y_s(t)$ is constructed through the Dirac filter $\delta(t-t_p)$, and $Y_s(f)$ is the Fourier Transform (FT) in time of $y_s(t)$. $Y_s(f)$ then corresponds to the convolution of the FT of the original signal $Y(f) = \text{FT}(y(t))$ with the spectral windows $W(f)$, which is the FT of the irregular Dirac comb, $W(f) = \text{FT}(w(t))$.

In the theoretical regular sampling case, the spectral window is a Dirac comb and the Fourier transform of the sampled data corresponds to a periodised version of the original signal's FT. This property leads to the well-known Nyquist-Shannon theorem [51] which is not valid in the irregular sampling case. An example of data FT and spectral window is provided in Fig. A2.

Hence, the analysis problem becomes that of obtaining a deconvolution of the spectral line data $Y(f)$ from the spectral window $W(f)$. The mathematical modelling for this problem is relatively simple: as the original signal is constituted of a sum of pure frequencies, each data point $y(t_p)$ is expressed as a weighted sum of complex sinusoids, the so-called *atoms*:

$$y(t_p) = \sum_{l=1}^{L} c_l e^{2i\pi \upsilon_l t_p} + \varepsilon_p, \quad \text{(A.2a)}$$

where ε_p is the error on the measurement, c_l and υ_l are the complex amplitudes and frequency, and L is the total number of spectral components. The formulation of Eq. (A2a) presents two problems: first, it is non-linear with respect to the frequencies υ_l, and second, L is unknown a-priori.

The problem described by Eq. (A2a), which amounts to fitting multiple complex sinusoids to the input data, is a very general signal processing problem which arises in many fields of physics. Such a spectral analysis problem from irregularly sampled data is very common in A&A, where time series acquisition usually suffer from incomplete temporal coverage, in particular periodic gaps

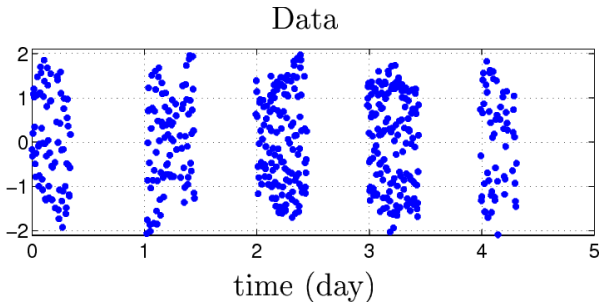

Fig. A1 Observation for the radial velocity curve of the Herbig Ae star HD 104237. These data correspond to five observing nights of high resolution spectroscopy at SAAO (South African Astronomical Observatory) during April 1999. The irregular data sampling due to day/night alternation is very clear.

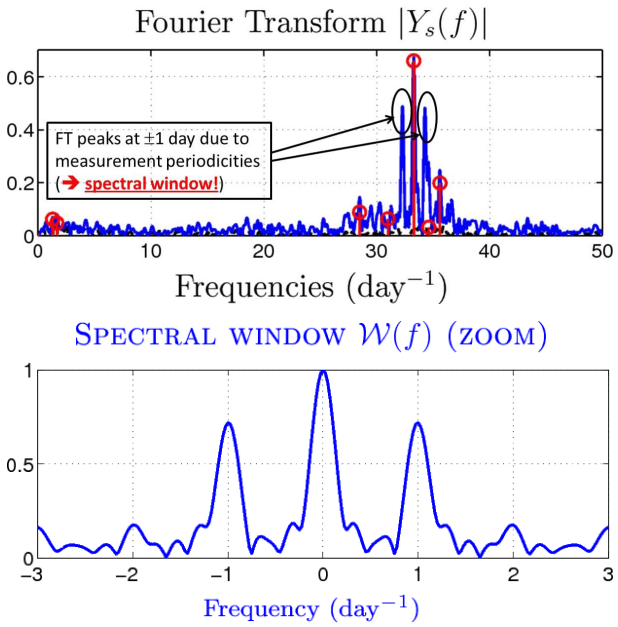

Fig. A2 Top frame: Fourier Transform of the data (blue line) presented in Fig. A1 and *SparSpec* detection results. Various peaks have been detected (indicated by the red vertical lines terminating in a red circle), the lower frequency ones being related to various orbital movements residuals. The black dotted line corresponds to the FT of the estimation residuals. Bottom frame: the (zoomed) spectral window for the measurements presented in Fig. A1: there are very clear ±1 secondary lobes corresponding to the one-day periodicity in the lack of measurements. The sidebands peaks at ±1 day are therefore removed from the FT data in the detection results shown in the top frame.

caused by the Earth's rotation and revolution, and aperiodic interruptions due to the weather. Many methods have been proposed in the fields of A&A to improve the analysis of such irregularly sampled time series, based on generalizations of the Lomb-Scargle periodogram [38, 39] and Data-Compensated Discrete Fourier Transform [52]. These methods involve iterative analysis [41], generally used when dealing with a large number of data points, or fitting periodic signals (Phase Dispersion Minimization [53], string length method [54]) to short data strings. Such methods, however, are inadequate when there are several temporal frequencies and too few measurements.

A major simplification [55] of Eq. (A2a) can be obtained by using a discretization of the frequency axis $f_k = (k/K) * f_{\mathrm{MAX}}$, with $k = [-K, \ldots, K]$, where f_{MAX} is much larger than the largest frequency component that can reasonably be present in the measurements, leading to:

$$y(t_p) = \sum_{k=-K}^{K} x_k e^{2i\pi f_k t_p} + \varepsilon_p. \qquad \text{(A.2b)}$$

The problem is linear with respect to x_k, but we now must deal with an even larger number of unknown amplitudes x_k and frequencies f_k, as we take that $K \gg L$. However, the estimation of the spectral lines $\{x_k, f_k\}$ can then now be greatly simplified imposing the *sparsity* of x_k and f_k, i.e. imposing that the x_k and f_k have only a small number of non-zero components. Such a problem can be tackled through the Sparse Representations principle.

A.2 Basic theory of sparse representations

Formally, *Sparse Representations* [27, 56, 57] are representations that account for all information in the input data $y(t)$ with a linear combination of a small number of elementary signals (for instance: sine waves, Diracs, ...) called *atoms* that belong to a selected family (a *dictionary*) which contains many such elementary signals. The atoms set is a redundant family, i.e. it does not form a basis as the number of atoms ($2K + 1$, see Eq. (A2b)) exceeds the dimension P of the signal space, so any signal can be represented by more than one combination of different atoms. Among all these various possible combinations, the one with the smallest number of atoms is the Sparse Representation of the signal. The sparsity of $\{x_k\}$ can be quantified with the L0 (pseudo-)norm, i.e. the number of non-zero components in $\{x_k\}$: $\|\boldsymbol{x}\|_0 = \#\{k, |x_k| \neq 0\}$. Hence, the Sparse Representations of $\{x_k\}$ is defined as $\hat{x} = \arg(\min_x \|x\|_0)$, subject to $y = W \bullet x$. Here $\boldsymbol{y} = [y_1, y_2, \ldots, y_P]^T$ is the vector of data taken at position t_p; $\boldsymbol{x} = [x_1, x_2, \ldots, x_M]^T$ is the vector of complex amplitudes, and $\boldsymbol{W} = [\boldsymbol{w}_{11}, \ldots, \boldsymbol{w}_{PM}]$ is a matrix where the column vector $\boldsymbol{w}_k$ corresponds to the k-th atom at the time point t_p for $p = \{1, \ldots, P\}$. The *Sparse Approximations* of signals [58–60] is the version of the Sparse Representations adapted to noisy data, i.e. $\hat{x} = \arg(\min_x \|x\|_0)$, subject to $\|y - W \bullet x\|_2^2 < \alpha$, where α is a user-defined threshold related to the noise level. Theoretically, the Sparse Approximation problem can also be written as the minimizer of the criterion:

$$J_0(x) = \|\boldsymbol{y} - W\boldsymbol{x}\|^2 + \gamma \|\boldsymbol{x}\|_0 , \qquad \text{(A.3)}$$

where γ is a penalization parameter related to the noise level. However, to minimize this criterion, one must perform a combinatorial optimization, i.e. sift through all possible combinations of elementary signals, which is intractable for large M. Hence, two kinds of methods have been proposed to get round this problem. The first one, often called a *greedy pursuit* algorithm, iteratively adds atoms to the initial approximation of the signal to improve such approximation [61]. The second one, often called a *convex relaxation* scheme, replaces the L0-norm in Eq. (A3) with another penalization term, generally based on the L1-norm, such that the criterion may be minimized more easily, particularly when considering optimization of the use of CPU time.

Strictly speaking in fact, the sparsest solution minimizes the least-square criterion penalized with the L0-norm, i.e. the number of non-zero components in the solution. However, minimizing such a criterion requires an exploration of all possible combinations of modes in the input

dataset, similarly to the SVD technique proposed in [43], which is very demanding in terms of CPU-time consumption. Such an exploration is in fact avoided in *SparSpec* by considering the L1-norm, i.e. the sum of the absolute values of the mode amplitudes, instead of the L0-norm penalization. Much theoretical work has been performed to determine the conditions of equivalence between the L0-norm and the L1-norm penalization criteria (see for instance [60,62,63]). Hence, the accuracy of our calculations (both real-time and post-pulse) is guaranteed by the comparison between a model input spectrum (with/out background noise) and the output spectrum as calculated by *SparSpec* using the actual geometry of magnetic sensors.

Here we follow this convex relaxation approach, classically using the L1-norm $\|\boldsymbol{x}\|_1 = \Sigma_k |x_k|$, instead of the pseudo-norm L0 in criterion (A3), so that $\hat{x} = \arg(\min_x \|y - W \bullet x\|_2^2) + \lambda \|x\|_1$. Hence we obtain the criterion [27, 56, 57]:

$$J_1(x) = \|\boldsymbol{y} - W\boldsymbol{x}\|^2 + \lambda \|\boldsymbol{x}\|_1 = \|\boldsymbol{y} - W\boldsymbol{x}\|^2 + \lambda \sum_{k=-K}^{K} (|x_k|). \tag{A.4}$$

It can then be easily shown that the criterion of Eq. (A4) is convex, therefore has no local minima, but, as the number of unknowns may be larger than the number of data points, this criterion is not strictly convex, i.e. the solution cannot be a-priori guaranteed to be unique. Moreover, this criterion is not differentiable for $x_k = 0$, which is a necessary (but not sufficient) condition for sparsity [62].

In practice, minimizing this L1-norm penalized Least-Square (LS) fitting criterion is much easier than minimizing the original one based on the L0-norm, and many computationally efficient algorithms have been developed, some of which can be made compatible with a real-time system using a 1 kHz clock-time. However, minimizing Eq. (A4) does not necessarily lead to the same solution as minimizing Eq. (A3), i.e. sufficient conditions for the equivalence between the L0-norm and L1-norm need to be satisfied [60, 62–64].

The choice of the family of atoms is critical in the Sparse Representations (and Approximations) of signals as, with an appropriate choice, these atoms might be well adapted to the signal to be analysed and might lead to a matrix $\boldsymbol{W}$ with good analytical and numerical properties. For example, it can be shown that if the signal can be represented with $\|\boldsymbol{x}\| < (1 + 1/\mu)/2$ components, with $\mu = \max_{k \neq l}(|\boldsymbol{w}_k^{\mathrm{H}} \boldsymbol{w}_l|)$, where $\boldsymbol{W}^{\mathrm{H}}$ is the Hermitian transposition of $\boldsymbol{W}$, then minimizing Eq. (A4) will lead to the selection of the same atoms as the solution minimizing Eq. (A3) [63]. For these reasons, the matrix $\boldsymbol{W}$ is often chosen as a family of relatively uncorrelated atoms, such as wavelets, Diracs, pure sine waves, etc ... Note that for the spectral analysis problem the atoms are driven by the problem and we get $w_k = \exp(2i\pi f_k t_p)$, for $p = \{1, \ldots, P\}$. Moreover, the sparsity of the components x_k, and so the L1 norm, has to be computed on the modulus of the complex amplitudes x_k, while the sparse approximation problem is generally studied for real-value amplitudes. Note that $|x_k| \neq |\mathrm{Re}(x_k)| + |\mathrm{Im}(x_k)|$, so imposing the sparsity on the complex modulus is radically different in terms of the model than sparsity applied separately on the real and imaginary components.

A.3 Relation with the tokamak plasma fusion problem

Considering now the particular application of the Sparse Representation method to magnetically confined thermonuclear fusion plasmas in a tokamak device, the MHD analysis is based on magnetic and turbulence measurements, and typically starts with an initial Fourier decomposition of the data in the time/frequency domain to obtain the individual frequency components $\psi(\omega)$. In a tokamak the plasma column has, to a first approximation, 2D boundary conditions along the longitudinal (the toroidal direction) axis and on the plane perpendicular to it (the poloidal direction). The spatial structure of the MHD instabilities is then determined by further decomposing each frequency component in its toroidal (n) and poloidal (m) harmonics: $\psi(\omega) = \mathrm{e}^{-i\omega t} \Sigma_{n,m} \mathrm{A}_{mn} \mathrm{e}^{in\phi} \mathrm{e}^{im\theta}$. Here ϕ and θ are the toroidal and poloidal angle coordinates, respectively, and we have used the fact that in tokamak geometry one single toroidal component with a given n usually has multiple poloidal components due to toroidicity and various other geometrical effects. The aim of toroidal (poloidal) mode number detection is to determine the mode numbers n (m) of the magnetic instabilities present in the plasma and to estimate their amplitude from data acquired with P detectors unevenly positioned at angles ϕ_p (θ_p), $p = \{1, \ldots, P\}$ being the suffix labelling the individual sensors used for the measurement.

For generality and consistency with the original astrophysics notation, in the tokamak plasma fusion problem the Fourier conjugated variables can still be called [*time*, *frequency*], which can be the real toroidal [ϕ, n] or the poloidal [θ, m] conjugated angle and mode-number variables. This means that our mathematical formulation can be equivalently used for calculating the n- (toroidal) and m- (poloidal) mode numbers by using the relevant sensor geometry. For the determination of the poloidal mode numbers in tokamak geometry, we must remember that we have to consider explicitly the so-called θ_*-*correction* [65–67] to the sensors', so as to run the mode number decomposition analysis using the correct, i.e. equilibrium-dependent, sensor geometry.

Considering now for simplicity of notation the specific case of toroidal mode number analysis, each measurement $y(t_p)$ can be mathematically modelled with a slight variation of Eq. (A2a):

$$y\left(\phi_p\right) = \sum_{l=1}^{L} \alpha_l \mathrm{e}^{in_l \phi_p} + \varepsilon_p, \tag{A.5}$$

where n_l and α_l are the unknown mode numbers and amplitudes, respectively, L is the unknown number of modes and ε_p corresponds to the noise on the data for the given p-th sensor, and periodic boundary conditions in ϕ have been used. Thus, the mode detection problem is strictly equivalent mathematically to the A&A spectral analysis problem.

Evaluating the amplitudes α_l and the mode numbers n_l of multiple modes in a multi-harmonic spectrum is a very difficult problem, even if the number of modes in the input spectrum is actually known a-priori. The usual way to tackle this problem is performing a best LS fitting of the input data. However, this criterion has many local minima for real valued spectral peaks [63, 68, 69], hence in principle requiring a combinatorial exploration for integer-valued mode numbers n_l, and an a-posteriori thresholding scheme to differentiate the "correct" from the "wrong" solutions. This is a very CPU-time intensive process and cannot possibly be adapted for real-time applications on the sub-millisecond time scale required for the analysis of the JET measurements. An alternative solution consists of providing an estimate for the amplitudes of all possible mode numbers in the range $\{-K, \ldots, K\}$ (where $|K|$ is much larger than the maximum mode number that can be conceivably present in the input spectrum), at the same time enforcing that most of these modes actually have a null amplitude, i.e. a utilizing a Sparse Approximation.

The mode detection (i.e. the spectral analysis) problem is particularly difficult in the case of tokamak plasma physics as the data is unevenly sampled and sparse, because of unavoidable installation constraints on the measurement devices. It can be shown that the difficulty of the spectral analysis problem is closely related to properties of the spectral window, such as the height and positions of its secondary lobes. Indeed, the mathematical problem described by Eq. (A5) can be expressed equivalently in the Fourier spatial domain as:

$$\begin{aligned} Y(\nu) &= W(\nu) * \sum_{l=1}^{L} \alpha_l \delta(\nu - n_l) + E(\nu) \\ &= \sum_{l=1}^{L} \alpha_l W(\nu - n_l) + E(\nu), \end{aligned} \tag{A.6}$$

where ν is the spatial frequency, $Y(\nu)$ and $E(\nu)$ corresponds to Fourier transform with respect to the angular position ϕ of the data and the noise, δ is the Dirac delta function, the symbol "$*$" is the convolution operator and $W(\nu)$ is the spectral window of the sampling scheme. Thus, if $W(\nu)$ has high secondary lobe (with an amplitude near to 1) at frequency ν_0, a mode number n will produce in the Fourier transform $Y(\nu)$ a maximum at $\nu = n$ and a secondary maximum at $\nu = n \pm \nu_0$. This means that it will be difficult to distinguish from the actual mode n and aliases modes at $n \pm \nu_0$. If we assume that obtaining the "true" mode number n_0 obeys a normal probability distribution with variance $\sigma^2(n_0)$, i.e. $P_{\mathrm{TRUE}}(n = n_0) \propto \exp(-(n - n_0)^2/\sigma^2(n_0))$, then we also find (using a best fit of the data that led to the analysis presented in [70, 71]) that the probability $P_{\mathrm{FALSE}}(n = n_0)$ for a *false detection* of n as n_0 due to the secondary lobes in the spectral window when using the *SparSpec* algorithm is $P_{\mathrm{FALSE}}(n = n_0) \propto \sigma^2(n_0) \times \Sigma(W(n) + W(|n| = \nu_0)/W(n_0)$ for all possible combinations of n and ν_0 such that $n = n_0 \pm \nu_0$. Then, the noise in data may make it difficult to distinguish between the actual mode n_0 and aliases modes at $n \pm \nu_0$.

In thermonuclear tokamak plasmas these lobes are due to regularities in the sampling (for instance when using a spacing larger than the Nyquist condition) and to the low number of sensors. This situation is further compounded by the failure of sensors over time, a problem that cannot easily be rectified due to restricted in-vessel access. As an example, the spectral window for two families of JET high-frequency magnetic sensors is shown in Fig. A3, comparing the data for the original complete set of 11 sensors that could be used in 1997 for toroidal mode number analysis, and for the 7 sensors in that set which can currently (2012) be used for real-time analysis with the AELM. Note that the original dominant $\nu_0 = \pm 10$ secondary lobe has now been supplemented by an even higher secondary lobe at $\nu_0 = \pm 4$, which is much more difficult to deal with as the most interesting n-number range is actually within $|n| = 1$ and $|n| = 10$.

When applied to thermonuclear plasma physics, the problem described by Eq. (A6) has some additional requirements with respect to the A&A problem described by Eq. (A2b), even if its solution can still be obtained using Eq. (A4). First, the data are complex-valued, implying that the Fourier transform of the data does not satisfy the Hermitian property $\hat{y}(-\nu) = \hat{y}^*(\nu)$ as in the spectral analysis of real-valued data. Obviously, the complex-valued data

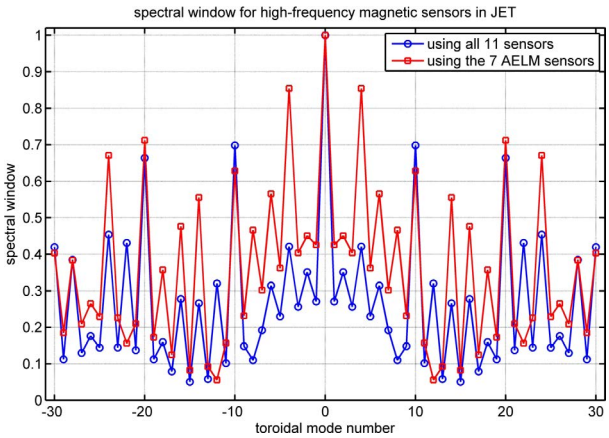

Fig. A3 The spectral windows $W(\nu)$ for the original and complete set of 11 high-frequency magnetic sensors of JET usable for toroidal mode number analysis, and for the seven surviving sensors between them that can currently be acquired in real-time by the AELM. Note that the original secondary lobe at $n = \pm 10$ has now been supplemented by an even higher secondary lobe at $n = \pm 4$, which is much more difficult to deal with.

have to be analysed together, conserving the I/Q phase relation between them, and not independently. Second, the mode numbers n_l can only take positive or negative integer values, while in the general spectral analysis problem frequencies take real values. This is a favourable property as the model (A2b) works on a discretized frequency grid. For A&A problems, a very fine discretization of the frequency grid is required so that real valued frequencies are not too distant from the nearest frequency on the grid. Note however that a posterior estimation of the out-grid estimation of the detected frequencies can be performed, i.e. using a barycentric estimation of the neighbour frequencies (as we are interested in all the $|x_k| \neq 0$ components and not in the approximation of the signal as $y \sim \Sigma_k w_k x_k$). Third, in the real time applications we consider for JET, a set of data is acquired every 1 ms, therefore the spectral analysis must be completed in an unsupervised manner in the short time between each measurement acquisition.

For the analysis presented here, the atoms are imposed by the model setup in Eq. (A5) to be pure complex exponential waves, $\boldsymbol{W} = \{\exp(in_k\phi_p)\}_{p,k}$, for $p = \{1, \ldots, P\}$ and $k = \{1, \ldots, M\}$, with $n_k = k - K + 1$ and $M = 2K + 1$. Due to the irregular sampling, the atoms are strongly correlated. Indeed, it can be shown that $|\boldsymbol{w}_k^{\mathrm{H}}\boldsymbol{w}_l| = W(n_k - n_l)$, so that it corresponds to regular samples of the spectral window. As $W(\nu)$ may take values greater than 1/3 (as shown in Fig. A3), the previous condition guarantees exact detection only if the signal consists of a single mode number. Nevertheless, it has been shown from many simulations and analysis of measurements using comparisons between different numerical methods that such a solution generally gives very satisfactory results in terms of detection, even in the case of multiple modes [28, 68]. Moreover, for irregular sampling, uniqueness of the global minimizer is almost surely guaranteed if it has less than P/2 non-zero components, where P is the data size [26].

In terms of amplitude estimations, it has been shown [26, 27] that minimizing Eq. (A4) leads to an underestimation of the amplitudes of the detected mode numbers due to the L1-norm penalization term. Thus, an a-posteriori LS re-estimation of these amplitudes is usually performed for post-pulse analysis in a second step within the calculations, after the modes have been actually detected. Their amplitudes are computed by minimizing the least square criterion $\|\boldsymbol{y} - \boldsymbol{W}_{\mathrm{DET}}\boldsymbol{x}_{\mathrm{DET}}\|^2$ where only the non-zero amplitudes of the optimization step are preserved in $\boldsymbol{x}_{\mathrm{DET}}$. Note that this a-posteriori amplitude estimation step is not an absolute necessity for the real-time analysis, as its main objective is to detect the actual modes, their mode numbers and frequency width, and not to precisely estimate their absolute amplitudes, a scaled value being sufficient for this purpose.

Many numerical algorithms are available to minimize criteria such as those of Eq. (A4) for Sparse Approximations. While for real-valued unknowns $\boldsymbol{x}_k$ this problem can be written as a classical Quadratic Program, for complex-valued unknowns $\boldsymbol{x}_k$ it corresponds to a Second-Order Cone Program [57]. An algorithm based on an iterative Block Coordinate Descent procedure has been previously proposed [26, 27], and implemented in the *SparSpec* code. This procedure consists of performing successive one-dimensional minimization steps with respect to each complex-valued unknown $\boldsymbol{x}_k$, where each one-dimensional minimization has an explicit solution. This algorithm is very efficient and a correct solution can be typically found in less than 1 ms using the rather modest computational resources available to process real-time JET data [14, 15].

A real-time implementation of the proposed modes detection method requires not only an efficient optimization algorithm to minimize Eq. (A4) but also, even more importantly for a frequency-degenerated spectrum, an efficient unsupervised tuning of the penalization parameter λ. Many numerical algorithms are available to minimize criteria such as those of Eq. (A4) for Sparse Approximations. The penalization parameter λ is related to the noise level [26] and requires an appropriate tuning, since it increases the penalty for those solutions which invoke a larger number of modes. It can be shown that the first order necessary and sufficient optimality condition for convex non-differentiable functions (often knows as the Karush-Kuhn-Tucker optimality conditions [72, 73]), see for instance [74, page 710], provides a physical interpretation for the parameter λ [26]: a) for $\lambda > \lambda_{\mathrm{MAX}} = \max_k(|\boldsymbol{w}_k^{\mathrm{H}}(\boldsymbol{y} - \boldsymbol{W}\boldsymbol{x}_{\mathrm{MIN}})|) = \max_k(|\boldsymbol{Y}(n_k)|)$, the minimizer $\boldsymbol{x}_{\mathrm{MIN}}$ of Eq. (A4) is identically zero, i.e. the unique solution has no detected modes; and b) for a given λ, the minimizer $\boldsymbol{x}_{\mathrm{MIN}}$ of Eq. (A4) satisfies $\max_k(|\boldsymbol{w}_k^{\mathrm{H}}\boldsymbol{r}|) = \max_k(|\boldsymbol{R}(n_k)|) < \lambda$, where $\boldsymbol{r} = \boldsymbol{y} - \boldsymbol{W}\boldsymbol{x}_{\mathrm{MIN}}$ is called the residual (data minus the model corresponding to the estimated modes). Hence λ can be interpreted as the maximum peak amplitude allowed in the FT modulus of the residual, and choosing λ to be a fraction $\lambda_{\mathrm{NORM}} \in [0, 1]$ of the maximum of the FT of the data $\lambda = \lambda_{\mathrm{NORM}} \times \max(|\boldsymbol{W}^{\mathrm{H}}\boldsymbol{y}|)$, ensures the FT of the residual $\boldsymbol{r}$ to be lower up to this fraction relative to the maximum of the data FT. Hence knowledge of the noise level in the measurements helps to determine the optimum value for λ_{NORM} to be used for real-time and post-pulse analysis of MHD fluctuation data.

A.4 Optimisation of the *SparSpec* algorithm for real-time analysis within the AELM

The main issues to be considered for the optimisation of the *SparSpec* algorithm for real-time analysis within the AELM are related to the computational limitations of the AELM itself, namely the need to perform all required data analysis within a hard CPU limit of 850 μsec using an embedded 1 GHz PowerPC processor with 512 MB of RAM. This hard CPU limit is such that if the calculations are not completed within this time limit, the AELM issues an error message, stops processing and returns to a safe operational state, meaning that the remaining part of the discharge is

lost for analysis.

As indicated in Section 3, various tasks need to be completed within this 850 μsec time limit, which leaves at most around 650 μsec for the demanding *SparSpec* calculations once data initialisation has been completed. Therefore, two main features of the *SparSpec* algorithm need to be optimized for real-time analysis, namely the size of the dictionary used to model the input dataset (i.e. the parameter *SS-*N_{MAX} in Table B5), and the penalization value λ_{NORM} (i.e. the parameter *SS-Lambda* in Table B5). A third feature of the *SparSpec* code, namely the LS re-estimation of the mode amplitudes after the calculation has been completed, is available as an option for real-time analysis, but it is in practice always skipped as it has been found that it requires too much CPU time. Finally, increasing the number of input signals up to the maximum value (= eight) has also an effect on the CPU time required to complete the real-time *SparSpec* calculations. However, it turns out that the increase in CPU time is in general relatively small when increasing the number of sensors used in the calculations once at least five (out of the possible eight) are used, which is the actual minimum number that guarantees a correct mode detection. Hence, practically we always use between five and seven sensors for real-time calculations, as three of the eight sensors acquired by the AELM are affected by intermittent connection problems (in-vessel), causing pick-up and drifts in the data that cannot be compensated reliably in real-time.

Optimisation tests for the real-time *SparSpec* calculations are greatly simplified by one feature of the AELM, which allows re-playing offline an already run real plasma discharge while changing some of the parameters used in the *SparSpec* algorithm. Different calculations can then be compared using the actual plasma data and the actual AELM hardware and software, allowing much more realistic CPU time limit tests. These analyses are then complemented with simulations run using Matlab R14 on a 2 GHz laptop with 1024 MB of RAM, where a model input data set is constructed and the results of the real-time and post-pulse *SparSpec* algorithm are compared. For such simulations, the input signal $S_{\mathrm{IN}}(\phi_n)$ at the position $\phi_n \in [0, 2\pi]$ of each magnetic sensor is constructed as an arbitrary superposition of different components at the integer mode numbers n_k, $k \in [-N_{\mathrm{MAX}}, N_{\mathrm{MAX}}]$, where N_{MAX} is the highest mode number in the spectrum:

$$S_{\mathrm{IN}}(\phi_n) = \left[\sum_{k=-N_{\mathrm{MAX}}}^{k=+N_{\mathrm{MAX}}} A_k \exp\left(i n_k \phi_n + i\delta_k\right) + \sigma_{\mathrm{SIG}} \times (r_{1k} + i r_{2k}) \right] + \sigma_{\mathrm{MEAS}}(\phi_n) \times (r_{3n} + i r_{4n}) \,. \quad \text{(A.7)}$$

Each n_k spectrum component can have a fixed or randomized amplitude A_k and relative phase δ_k. The quantities $\sigma_{\mathrm{SIG}} \in [0, 1]$ and $\sigma_{\mathrm{MEAS}}(t_n) \in [0, 1]$ represent the standard deviation in the background noise on each spectral component and on the measurement itself at each sensor, respectively, and are known a priori (i.e. they have a fixed and unique value for each simulation) as they can in principle be measured directly on the system when installed. The quantities $\{r_{1k}, r_{2k}, r_{3n}, r_{4n}\}$ are random numbers chosen from a uniform distribution in the interval $[0.0 \rightarrow 1.0]$; note that the random seed used for $\{r_{1k}, r_{2k}\}$ can be different from the one used for $\{r_{3n}, r_{4n}\}$. With this approach, the noise has independent and un-correlated complex components satisfying the circularity property. In general, σ_{SIG} and σ_{MEAS} can be different and, more importantly, σ_{MEAS} can have different values for different sensors. Intuitively, σ_{SIG} can be associated to background noise from the plasma, for instance due to un-coherent turbulence; conversely, σ_{MEAS} is associated with "engineering" errors, such as tolerances on the position and alignment of the sensors, calibration errors, and various effects such as cross-talk, drifts, offset, signal pick-up and bit-noise in the cabling and electronics.

Figure A4 shows the CPU time required to obtain a confidence level $\tau > 0.7$ for the real-time concurrent evaluation of the mode number, frequency and damping rate as a function of the number of sensors used for the *SparSpec* tracking algorithm. Using the results presented in Fig. 23, only mode numbers in the range $-8 \le n \le 10$

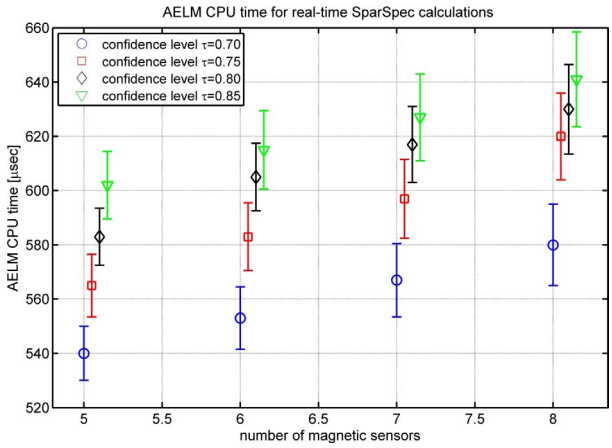

Fig. A4 The CPU time required to obtain a confidence level $\tau > 0.7$ for the real-time concurrent evaluation of the mode number, frequency and damping rate as a function of the number of sensors used for the *SparSpec* tracking algorithm. Only mode numbers in the range $-8 \le n \le 10$ are considered for this analysis, and the data have been obtained by replaying off-line various actual plasma shots by changing only the number of sensors being used from five to eight. The value SS-N_{MAX} = 30 and SS-Lambda = 0.85 were used for these calculations, the tracking mode was set to "highest", and various combinations of the same number of sensors were tested. For graphical purposes the x-axis value is slightly shifted with respect to the integer mode number to avoid overlapping. The vertical error bar indicates the scatter in the CPU time limits for the different discharges and combination of sensors used in this analysis.

are considered for this analysis. The data have been obtained by replaying off-line various actual plasma shots by only changing the number of sensors being used from five to eight. The value *SS-*N_{MAX} = 30 and *SS-Lambda* = 0.85 were used for these calculations, the tracking mode was set to "highest", and various combinations of the same number of sensors were tested. The CPU time needed for the real-time calculations increases effectively linearly with respect to the number of sensors and the confidence level. Considering a hard limit of 650 μsec for these calculations, it is clear that if a confidence level $\tau > 0.8$ is needed for the real-time calculations, only seven sensors at most can be used. We also note that the relative increase in the nominal CPU time required for achieving the same confidence level, for instance $\tau = 0.8$, as a function of the number of sensors once at least five sensors are used is relatively minor: timeCPU = [583, 605, 617, 630]μsec for five, six, seven and eight sensors respectively. This gives a relative increase timeCPU(5:8)/timeCPU(5) ≈ [1.00, 1.04, 1.06, 1.08], i.e. not exceeding 10 % even when eight sensors are used. Therefore, for the *SparSpec* real-time calculations we always use up to seven of the reliable sensors acquired by the AELM.

Figure A5 shows the confidence level in the real-time *SparSpec* calculations as a function of the λ_{NORM} (i.e. the AELM *SS-Lambda* parameter) value when using various combinations of between five and seven sensors, *SS-*N_{MAX} = 30, the "*highest*" and the "*any*" tracking modes. The results shown in Fig. A5 contain simulation and AELM offline re-played data. When we replay off-line actual discharges using the actual AELM hardware and software, we use a CPU time constraint of 650 μsec; when we use simulated data, we define a CPU time constraint of 1.5 sec to take into account the superior computational resources of the laptop used for these tests. Again, only mode numbers in the range $-8 \le n \le 10$ are considered for this analysis. We note that the confidence level increases linearly with λ_{NORM} (as expected from previous simulations [69, 70]), and for $\lambda_{NORM} > 0.8$ the value $\tau = 0.8$ is routinely obtained. Therefore, in real-time the value $\lambda_{NORM} = 0.85 =$ *SS-Lambda* is almost always used.

Finally, Fig. A6 shows the CPU time required to obtain a confidence level $\tau \ge 0.8$ for the real-time concurrent evaluation of the mode number, frequency and damping rate as a function of the size of the dictionary (i.e. the *SS-*N_{MAX} AELM parameter) used for the *SparSpec* tracking algorithm. For this analysis, we use various combinations of between five and seven sensors, *SS-Lambda* = 0.85, the "*highest*" and the "*any*" tracking modes, and again we combine simulation and AELM offline re-played data taking the same CPU time constraints set for the analysis presented in Fig. A5. First, as the minimum value for the highest-|*n*| mode present in the input data is min(|N_{MAX}|) = 10, we start the analysis with *SS-*

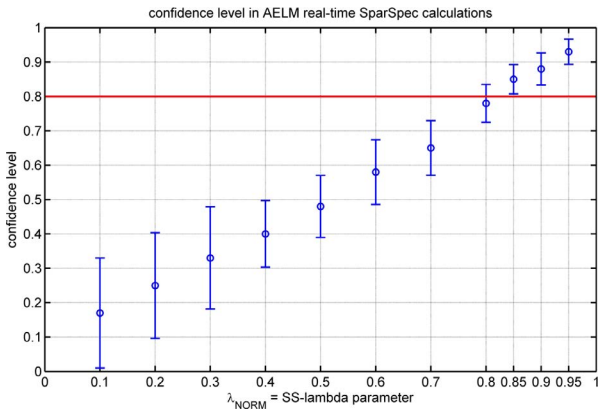

Fig. A5 The confidence level in the real-time *SparSpec* calculations as function of the λ_{NORM} (i.e. the AELM SS-Lambda parameter) value when using between five and seven sensors, with various combinations of them, SS-N_{MAX} = 30, the "highest" and "any" tracking modes, and combining simulations (with a CPU time constraint of 1.5 sec) and actual discharges re-played offline using the actual AELM hardware and software (with a CPU time constraint of 650 μsec). The red horizontal line at the value of the confidence level $\tau = 0.8$ indicates the acceptable value for the real-time calculations. The vertical error bar indicates the scatter in the confidence level for the different computational options used in this analysis.

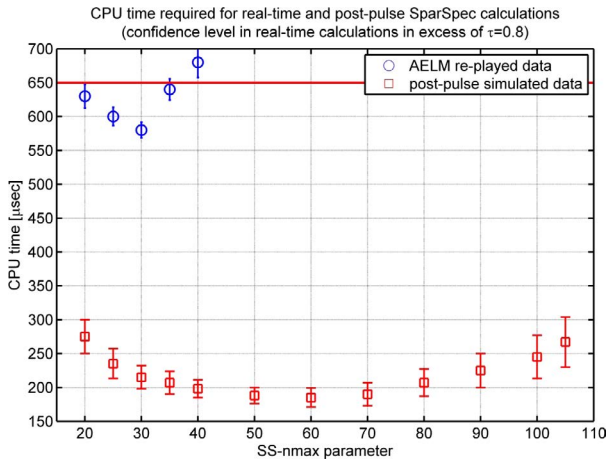

Fig. A6 The CPU time required to obtain a confidence level $\tau \ge 0.8$ for the real-time concurrent evaluation of the mode number, frequency and damping rate as function of the size of the dictionary (i.e. the SS-N_{MAX} AELM parameter) used for the *SparSpec* tracking algorithm. We use between five and seven sensors, with various combinations of them, SS-Lambda = 0.85, the "highest" and the "any" tracking modes, and again we combine simulation and AELM offline re-played data taking the same CPU time constraints set for the analysis presented in Fig. A5. The red horizontal line at the value of the CPU time = 650 μsec indicates the hard AELM time. The vertical error bar indicates the scatter in the confidence level for the different computational options used in this analysis.

$N_{MAX} = 20 = 2 \times \min(N_{MAX})$. Second, as in real-time the maximum value for the highest-$|n|$ mode that can actually be correctly detected is $\max(|N_{MAX}|) = 15$, our last test point becomes limited by the available RAM resources in real-time, and it is always set to $SS\text{-}N_{MAX} = 30$, i.e. $SS\text{-}N_{MAX} = 2\times\max(N_{MAX})$, when replaying actual discharges offline. Conversely, for the simulated data we do not have such strict RAM limitations, and we have added further data points up to $SS\text{-}N_{MAX} = 7 \times \max(N_{MAX}) = 105$. This analysis shows that increasing the size of the dictionary has first a beneficial effect on the computational time required for the *SparSpec* calculation, as the BCD algorithm becomes more efficient for the sparsest data set, but then causes an increase in the required CPU time as the matrices becomes very large to handle and more and more calculations are required. For the post-pulse simulated data, we notice that the optimum value of $SS\text{-}N_{MAX}$ is around three to five times the largest mode number that can reasonably be present in the input dataset. However, already the value $SS\text{-}N_{MAX} = 3 \times |N_{MAX}|$ cannot be used in real-time due to the limitations in the AELM computational resources, which are effectively reached at $SS\text{-}N_{MAX} = 35$, corresponding to $SS\text{-}N_{MAX} = 2.33 \times |N_{MAX}|$ for the usual value $|N_{MAX}| = 15$. Therefore, we use the value $SS\text{-}N_{MAX} = 30$, i.e. $SS\text{-}N_{MAX} = 2\times|N_{MAX}|$ for the usual value $|N_{MAX}| = 15$, in real-time so as to have operational margins with respect to the hard AELM CPU time limit of 650 μsec.

A.5 Additional applications of the *SparSpec* algorithm to the analysis of magnetic fluctuations in JET tokamak plasmas

To complete this overview of the application of the Sparse Representation method and the *SparSpec* code, we present here some results for the analysis of two types of magnetic fluctuations in JET tokamak plasmas: coherent fast-ion driven modes in the Alfvén frequency range, similarly to those detected in real-time by the AELM, and incoherent turbulence spectra.

The JET discharge #55604 is a typical example of a scenario where frequency-sweeping Alfvén Cascades (ACs) [75,76], also known as Reverse Shear Alfvén Eigenmodes (RSAEs) [77–79], are excited by a population of high energy ions, in this case produced by Ion Cyclotron Resonance Heating of a plasma which has a non-monotonic q-profile. One of the main usefulness of these modes is that they can help determining the temporal evolution of the q-profile, particularly the value (q_{MIN}) and the location (R_{MIN}) of its inflexion point through the determination of their poloidal and toroidal mode number. The dispersion relation for these modes is simply given by:

$$f_{AC}(t) \propto \left| n - \frac{m}{q_{MIN}(t)} \right| \frac{\upsilon_A(t)}{2\pi R_{MAG}} + (\text{Doppler+thermal shift}), \quad (A.8)$$

where υ_A is the Alfvén velocity and R_{MAG} the position of the magnetic axis. Hence noting as t_0 the time point where the mode frequency has the minimum value, i.e. $f_{AC}(t_0)$ is just determined from the frequency shift due to the toroidal plasma rotation (Doppler shift) and ion-sound speed (thermal shift), we have that: $q_{MIN}(t_0) = m(t_0)/n(t_0)$. The time evolution of $q_{MIN}(t)$ can then be obtained from the linear dependence $f_{AC}(t)/\mathrm{d}t \propto 1/q_{MIN}(t)$, which give rise to the characteristic upward (i.e. q_{MIN} is decreasing) frequency sweeping of these modes.

Figure A7 shows the value of the toroidal (Fig. A7a) and poloidal (Fig. A7b) mode numbers for these modes as evaluated (post-pulse) using the *SparSpec* code, the latter one having been determined using the θ^*-correction to the probe position. This analysis was performed using 8 magnetic sensors located on the low-field side wall for toroidal mode number analysis, and 12 magnetic sensors located on the low- and high-field side wall in different machine sectors for poloidal mode number analysis. The magnetic data were acquired with two fast digitizer systems sampling the raw data from the pick-up coils at 1 MHz and 250 kHz, respectively, and then combined through their known relative transfer function. The toroidal and poloidal mode number decomposition was performed using the *SparSpec* code on a dictionary base containing 100 atoms with the following run-time parameters: $N_{MAX} = 30$, $\lambda_{NORM} = 0.85$. Figure A7c then shows the comparison between the value of q_{MIN} as determined from the poloidal and toroidal mode number analysis, and that obtained from the EFIT reconstruction of the equilibrium, in this case complemented with Motional Stark Effect (MSE) measurements. Despite a systematic difference of around $q_{MIN}/30$, there is

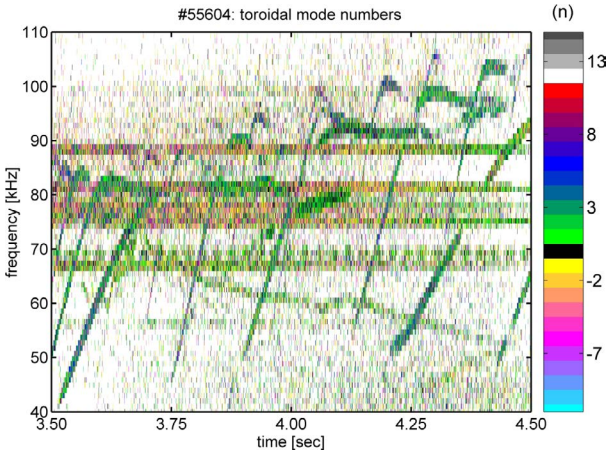

Fig. A7a The toroidal mode numbers for the frequency-sweeping ACs detected for the JET discharge #55604. The analysis was performed using 5 magnetic sensors located on the low-field side wall, which were acquired with a fast digitizer system sampling at 1 MHz (note that the various horizontal lines corresponds to "bad" digitizers bits). The toroidal mode number decomposition was performed using the *SparSpec* code on a dictionary base containing 100 atoms with the following run-time parameters: $N_{MAX} = 30$, $\lambda_{NORM} = 0.85$.

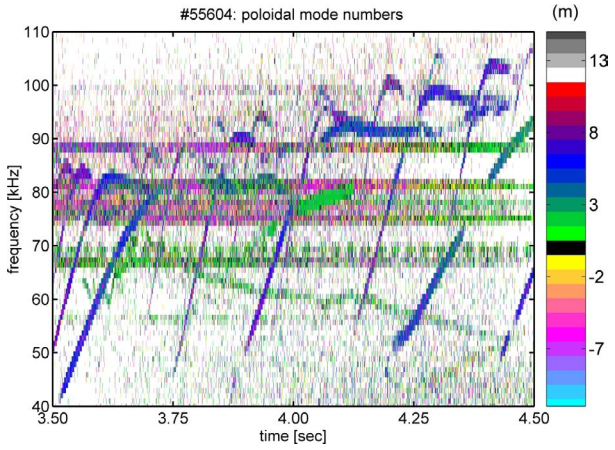

Fig. A7b The poloidal mode numbers for the frequency-sweeping ACs detected for the JET discharge #55604. The analysis was performed using 4 magnetic sensors located on the low-field side wall, which were acquired with a fast digitizer system sampling at 1 MHz (note that the various horizontal lines corresponds to "bad" digitizers bits), and used the appropriate position of the probe as determined with the $\theta*$-correction. The poloidal mode number decomposition was performed using the *SparSpec* code on a dictionary base containing 100 atoms with the following run-time parameters: $N_{\mathrm{MAX}} = 30$, $\lambda_{\mathrm{NORM}} = 0.85$.

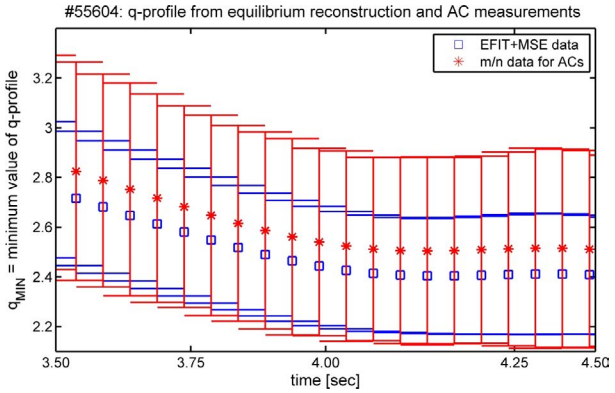

Fig. A7c Comparison between the value of q_{MIN} as determined from the poloidal and toroidal mode number analysis, and that obtained from the EFIT reconstruction of the equilibrium, in this case complemented with Motional Stark Effect (MSE) measurements. The typical error bars on the measurements are also shown, which are around 15 % to 20 % for both datasets. Note that the value of q_{MIN} from AC spectroscopy is slightly larger, by about $0.1 \approx q_{\mathrm{MIN}}/30$, than that obtained from the equilibrium reconstruction, but this discrepancy, although clearly systematic, is entirely within the uncertainty of the measurement, and points more towards an offset in the EFIT + MSE data than to an error in the determination of the toroidal and poloidal mode numbers.

a rather good agreement between these two sets of data point, which overall confirms the applicability of the *SparSpec* algorithm to this analysis.

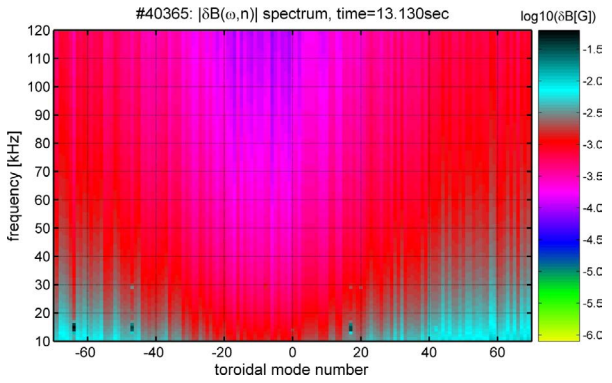

Fig. A8a Spectrogram of $|\delta B_{\mathrm{MEAS}}|$ in the (ω, n) plane, obtained using 9 magnetic sensors located on the low-field side wall, again acquired with two fast digitizer systems sampling the raw data from the pick-up coils at 1 MHz and 250 kHz, respectively, and then combining the data through their known relative transfer function.

The JET discharge #40305 is a typical example of a scenario where incoherent magnetic turbulence was measured using pick-up coils located at the plasma edge [80, 81]. We have used the *SparSpec* code to obtain the turbulence spectra $\delta B_{\mathrm{MEAS}}(\omega, n)$ as function of the toroidal mode number once a standard Fast Fourier Transform in the time/frequency domain has been applied to the raw data from the magnetic pick-up coils. Figure A8(a,b) show some of the results of this analysis at one selected time point: a spectrogram of $|\delta B_{\mathrm{MEAS}}|$ in the (ω, n) plane (Fig. A8a), and the deconvolution of the individual frequency and toroidal mode number components (Fig. A8b), respectively. This analysis was performed using 9 magnetic sensors located on the low-field side wall, again acquired with two fast digitizer systems sampling the raw data from the pick-up coils at 1 MHz and 250 kHz, respectively, then combining the data through their known relative transfer function. The toroidal mode number decomposition was performed using the *SparSpec* code using a dictionary base containing up to 500 atoms with the following run-time parameters: $N_{\mathrm{MAX}} = 100$, $\lambda_{\mathrm{NORM}} = 0.15$ (note that a much smaller value of λ_{NORM} needs to be used as the turbulence spectra are incoherent and have an amplitude which is just above the background noise level). In these figures, the spectral lines with amplitude much larger than the other components correspond to bad digitizer bits at low frequency < 20 kHz. When integrating the turbulence spectra in the ion acoustic frequency range, from 10 kHz to 120 kHz, we find that the amplitude of turbulence associated to Trapped Electron Modes, i.e. with negative toroidal mode numbers, is rather similar to that associated with Ion Temperature Gradient driven modes, i.e. with positive toroidal mode numbers. These measurements can then be used to test the prediction of turbulence codes in this scenario, as reported in [80, 81], and a rather satisfactory agreement is found.

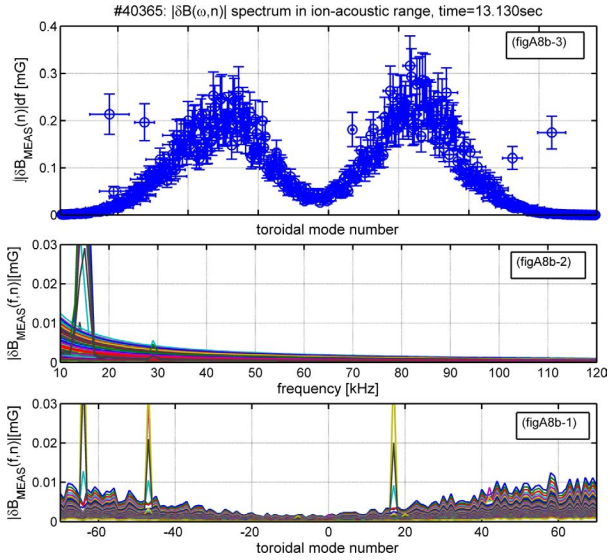

Fig. A8b Deconvolution of the individual frequency and toroidal mode number components of the turbulence spectrogram shown in Fig. A8a. Bottom frame: spectral decomposition as function of the toroidal mode number, superposing the different frequency components (the lines at much larger value correspond to bad digitizer bits at low frequency < 20 kHz). Middle frame: decomposition as function of frequency, superposing the different toroidal mode number components (note the very large components at low frequency < 20 kHz corresponding to the bad digitizer bits). Top frame: frequency integral in the ion acoustic frequency range, from 10 kHz to 120 kHz, of the data shown in the bottom frame, this time removing all "bad components": components with positive and negative toroidal mode numbers correspond to Ion Temperature Gradient driven turbulence and Trapped Electron Modes, respectively.

Appendix B.

B.1 Timing setup for the AEAD system via the AELM

The usual mode of operation of the AEAD plant involves remote control of the timing sequence via the AELM. Figure B1 shows a sketch of the timing waveform and of the AELM sub-panel dedicated to this task. The duration of the AELM time window (AEAD pulse control and data acquisition) is fixed and is limited to 12 seconds. Therefore, the only time point (T1) that needs to be set via the AELM is the start of the AEAD power ON phase, which also corresponds to the starting time for the real-time data acquisition (12 sec @1 kHz). This automatically sets the timing for the HV-ON and HV-OFF commands to the amplifier at the time points T0 = T1-1 sec and T2 = T1 + 12 sec, respectively. At T3 = T1 + 13 sec the power is switched-OFF (using a relay in the internal HV power supply of the amplifier with a 1 sec lag-time), and at T4 = T1 + 14.09 sec the overall AEAD data acquisition for post-pulse analysis (14.09 sec @1.25 kHz) is stopped. The correct setup of the starting time T1 > 40.01 sec from PRE (where PRE is 0.0 s and represents the start of a JET pulse experiment) is verified by the AELM software during the pre-pulse phase. Usually the AEAD active operation (with power to the in-vessel antennas) does not start before T1 = 41 sec to allow for a sufficient plasma density and current to have been established, so that the antenna frequency remains within its operational range (see Appendix B.2 for more details on this).

B.2 Frequency setup for the AEAD system via the AELM

Three possibilities exist for selecting the frequency source $F_{\rm REF}$ in the AELM, and this choice is made by selecting the *FrefSource* tab on the AELM configuration panel (see Fig. 4). The set frequency is then sent to the AEAD plant as the input signal for the AE exciter. The maximum operating frequency ($F_{\rm MAX}$) for the AEAD plant is set by the corresponding tab ($F_{\rm MAX}$ = 500 kHz in the case shown in Fig. 4), to match the available bandwidth of the synchronous detection system.

The first choice for selecting the frequency source $F_{\rm REF}$ corresponds to the *AELM-frequency* mode, as shown in Fig. B2a: this utilises a real-time algorithm based on the measurement of the magnetic field, plasma current and line-integrated density so as to look for AEs, which is the most common operational setup for the AEAD system. In this case, $F_{\rm REF}$ has to match the AE frequency, given by:

$$F_{\rm REF}[\mathrm{kHz}] = (multiplier) \times \frac{B_{\rm TOR}}{4\pi R_{\rm RES} q_{\rm RES} \sqrt{m_{\rm p} \sum_i n_i A_i}}$$
$$= 32.7 \times \frac{(\underline{MultiplierFref}) \times I_{\rm TF}[\mathrm{A}]}{\sqrt{(\underline{MassNumber}) \times \overline{n_{\rm e}}[10^{18}\mathrm{m}^{-2}]}} \times \frac{I_{\rm p}[\mathrm{A}]}{(\underline{IpNorm}[\mathrm{A}])}. \quad \text{(B.1)}$$

In Eq. (B1) $B_{\rm TOR}$ is the toroidal magnetic field, n_i and A_i are the density and atomic mass of all ion species, m_p is the proton mass, $q_{\rm RES}(r_{\rm RES}) = (2m + 1)/2n$ is the value of the safety factor at the mode resonant position $R_{\rm RES} = R_0 + r_{\rm RES}$, where R_0 is the magnetic axis position, r the minor radius coordinate, and m and n are the AE poloidal and toroidal mode numbers, respectively. The quantity *multiplier* defines which class of AEs is being investigated: *multiplier* = 1 is used for Toroidal AEs (TAEs) and *multiplier* = 2 for ellipticity-induced AEs.

The physical form of Eq. (B1), shown in the top line, is then translated into its real-time numerical implementation with the AELM, shown in the bottom line: $I_{\rm TF}$ ($\propto B_{\rm TOR}$) is the real-time value of the current in the toroidal field coils, $\overline{n_{\rm e}}$ is the real-time value of the line-integrated electron density, measured along a vertical chord passing through the plasma centre, *MassNumber* is a user-selected value that accounts for the plasma isotopic composition (so that *MassNumber* $\times \overline{n_{\rm e}} \propto \Sigma_i n_i A_i$), $I_p (\propto 1/q)$ is the real-time value of the toroidal plasma current, *IpNorm* is a user-

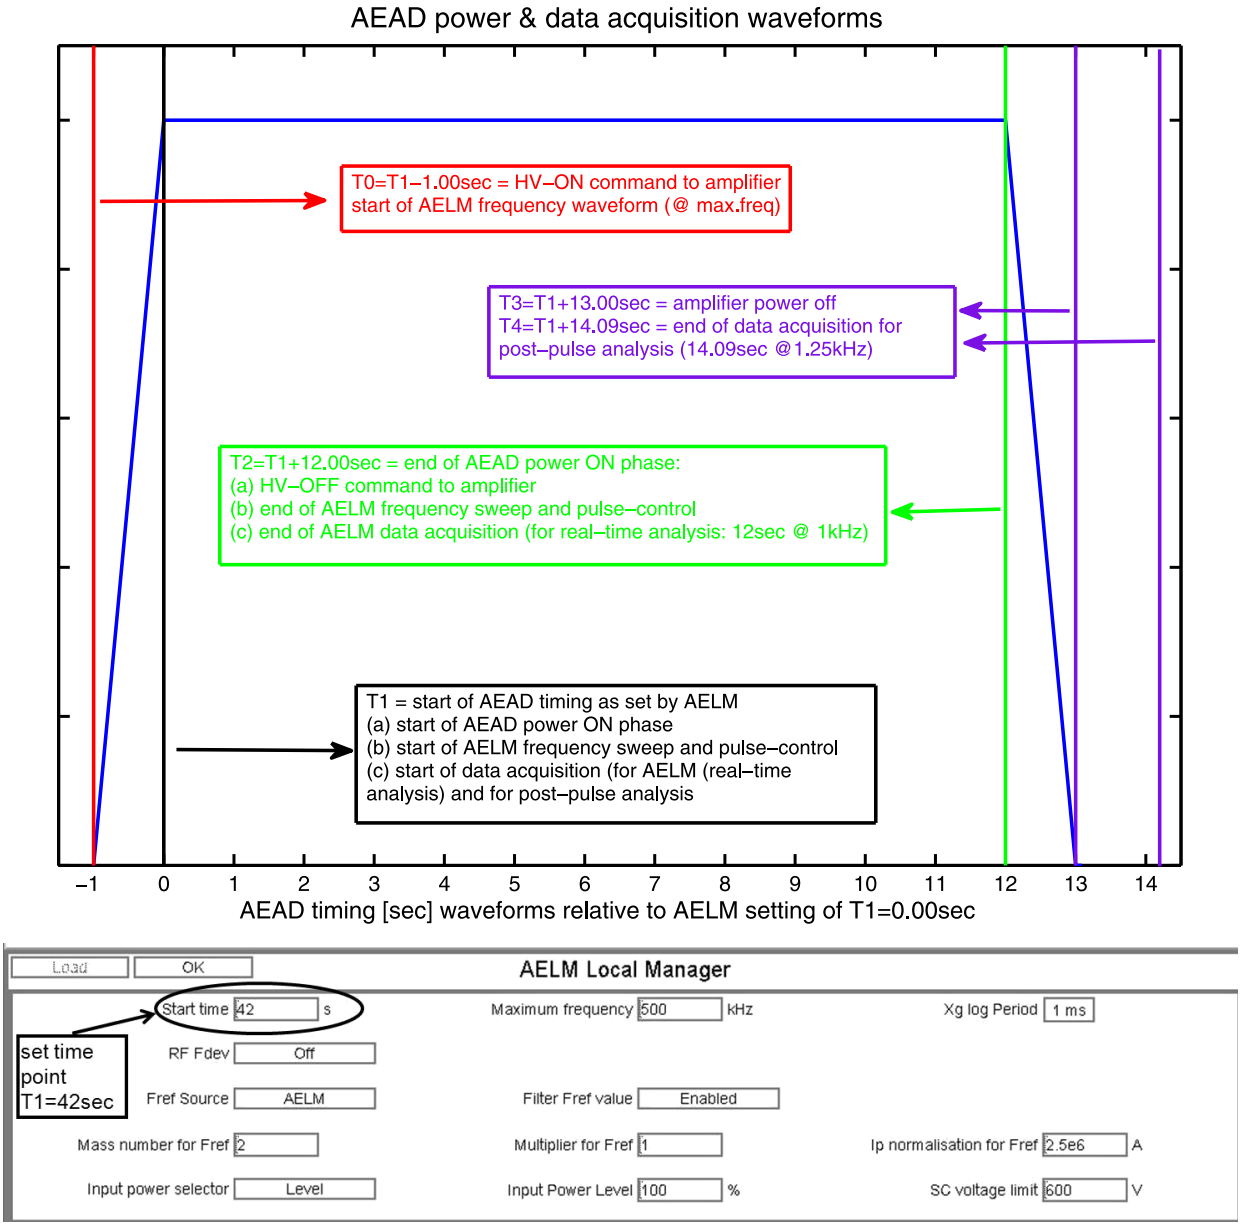

Fig. B1(a,b) The overall pulse timing sequence waveform for the AEAD system (top frame), and the setup of the T1 time point using the AELM (bottom frame).

selected normalization factor, so that $I_p/\underline{IpNorm} \propto q(t = \mathrm{T1})/q(t)$ provides a simple real-time estimate for the evolution of the safety factor profile during the active AELM time window, $\underline{MultiplierFref}$ is a user-selected value that accounts for the selected class of AEs being investigated, and the values $R_{\mathrm{AE}} = 3m$ and $q_{\mathrm{AE}} = 1.5$ are used for all the real-time calculations. Practically, the *AELM-frequency* mode can only be used when operating the AEAD plant on an actual plasma shot, i.e. when $\overline{n_e}$ is measured and has a finite value, as otherwise $F_{\mathrm{REF}} = F_{\mathrm{MAX}}$ throughout the entire length of the pulse, since there would be no density data otherwise, as for instance occurs in the case of a vacuum shot or a dry-run. Finally, note that although effectively only used when investigating AEs, the flexibility of the *AELM-frequency* mode is such that in principle it could be used for any other MHD mode whose frequency can be expressed as a function of the magnetic field, plasma current and plasma density.The parameters required for the *AELM-frequency-mode* are summarised in Table B1.

The second choice for selecting the frequency source F_{REF} is the *WF-frequency* mode, which corresponds to using a pre-defined frequency waveform, as shown in Fig. B2b. This frequency waveform can be selected (and/or edited/modified) using the scroll-down menu appearing on the tab *FrefWaveform*. A frequency waveform can be used on plasma, although this mode is essentially only used for commissioning and testing purposes, for instance during technical discharges such as a vacuum shot or a dry-run. The frequency generated by the waveform editor is given in percentage (0 % → 100 %) of the VCO units,

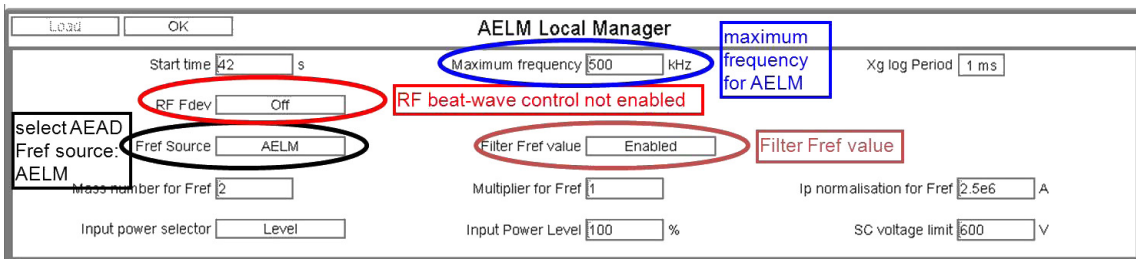

Fig. B2a AEAD frequency control via the AELM software, using the real-time values of the toroidal magnetic field, plasma current and plasma density.

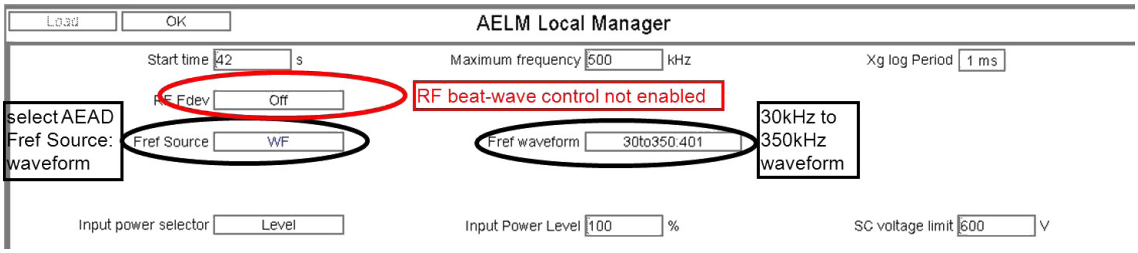

Fig. B2b AEAD frequency control via a pre-defined waveform.

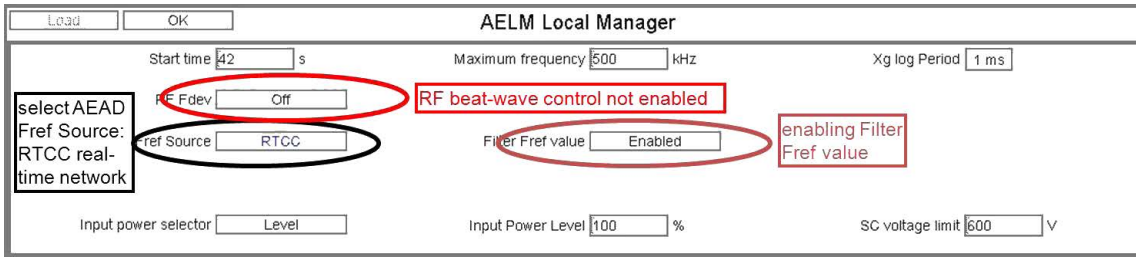

Fig. B2c AEAD frequency control using a real-time RTCC network algorithm.

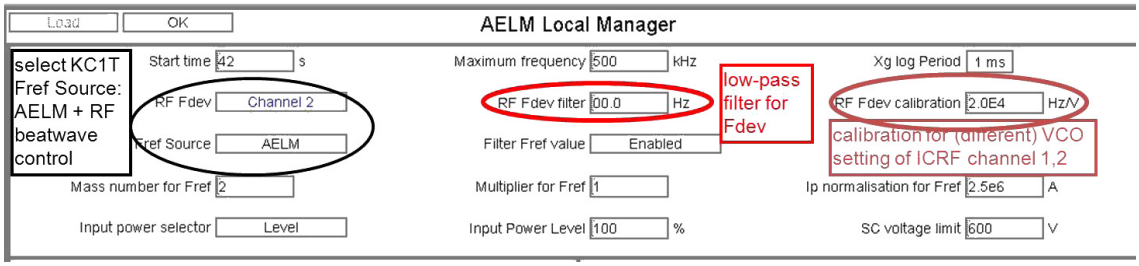

Fig. B2d AEAD frequency control using the real-time frequency provided by the AELM software combined with RF beat-wave control.

i.e. 100 [%] = 10 V = F_{MAX} (as set by the corresponding AELM tab), and must start/end at 100 [%] = 10 V because of the logic of the AELM initialisation algorithm.

The third choice for selecting the frequency source F_{REF} is the so-called *RTCC-frequency* mode, which corresponds to utilising an ad-hoc algorithm running under the Real Time Central Controller (RTCC) network, as shown in Fig. B2c. With this setup, which is actually only very rarely used, a user-prepared RTCC algorithm is utilised to generate the frequency waveform for AEAD operation. This option can only work on a plasma shot, as the RTSS signals are only available on actual plasma discharges, and allows a much more sophisticated real-time frequency control than that provided by the *AELM-frequency* mode, as any of the 1000 + RTSS signals can be utilised in an RTCC algorithm.

If, in addition to any of these choices, the *RF-Fdev-frequency* mode is also selected, as shown in Fig. B2d, then the AELM will take control of two modules of the ICRF plant (both the “master” and the “slave” subsystems) [16],

Table B1 The parameters required for the AELM-frequency-mode of operation, as set by the user. With these conventions, the resulting value of F_{REF} is in units of KHz.

parameter	Source	explanation of use
$\overline{n_e}$	obtained from the RTSS	real-time line-integrated plasma density using the interferometer data (units: [10^{18} m^{-2}])
I_{TF}	obtained from the RTSS	real-time current waveform in the toroidal field coils (units: [A])
I_p	obtained from the RTSS	real-time toroidal plasma current (units: [A])
MassNumber	user-selected value: can take any positive and real (not necessarily integer) value	chosen according to the main plasma gas: • *MassNumber* = 1 for H-plasmas • *MassNumber* = 2 for D- and He4-plasmas
MultiplierFref	user-selected value: can take any positive and real (not necessarily integer) value	chosen to scale the nominal F_{REF} frequency by a constant multiplicative factor
IpNorm	user-selected value: can take any positive value ≥ 0 *IpNorm* = 0 if this normalization is not intended to be used	normalization plasma current (units: [A]) usually chosen as the nominal flat-top value of the plasma current (*IpNorm* = 2.5e6 [= 2.5 MA] in the example shown in Fig.B2a)

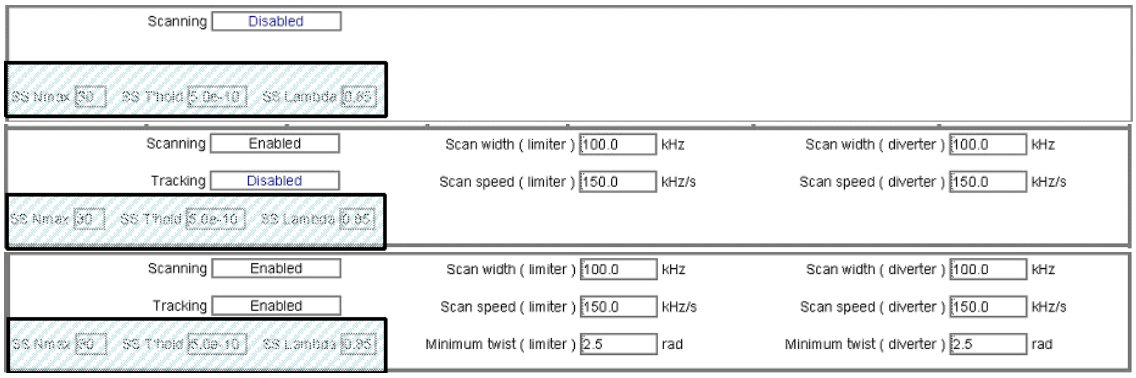

Fig. B3 Selection of scanning and tracking mode for the AELM frequency control; top frame: scanning mode disabled; middle frame: scanning mode enabled without tracking; bottom frame; tracking mode enabled (parameters not relevant for this choice of scanning vs. tracking mode of operation have been shaded in this figure).

so that the difference in frequency F_{DEV} between two of the ICRF antennas (channel-1 = module A, channel-2 = module D) matches the F_{REF} chosen from the *FrefSource* tab, i.e. $F_{DEV} = F_{REF}$. The value of F_{DEV} can also be filtered and a calibration factor is used to match the (possibly different) VCO frequency setting of the two ICRF generators. This mode of operation requires the AEAD plant to take full control of two RF generators, which effectively cannot be used anymore for heating purposes; hence it is only very rarely used. The main advantage of this frequency control is the much larger power delivered to the plasma at the selected beat-wave frequency by the RF generators, which can easily reach 100 kW, i.e. a value at least 20 times larger than that obtained with the AEAD exciter.

Finally, note that if selecting the RTCC or AELM modes for the F_{REF} source, a low-pass filter can then be applied to the F_{REF} value so as to smooth-out any spikes coming from the input plasma data (as shown in Fig. B2(a,c,d): *FilterFrefValue* enabled), which is the routine choice for operation.

Once the choice of the frequency source has been made, the *scanning* or *tracking* mode may be selected. If frequency scanning is disabled (as shown in Fig. B3, top frame), the AELM will output exactly the frequency set using the F_{REF} source, i.e. the antenna frequency is $0 \le F_{ANT} = F_{REF} \le F_{MAX}$.

If frequency scanning is enabled (as shown in Fig. B3, middle frame), a pre-defined symmetric and repetitive

Table B2 The parameters required for the scanning and tracking mode of operation, as set by the user.

parameter	use	range of validity (for both "limiter" and "diverter" configurations)
ScanWidth	both scanning and tracking mode	admissible values: $0 \le$ *ScanWidth* [kHz] ≤ 400 routine values: $100 \le$ *ScanWidth* [kHz] ≤ 200
ScanSpeed	both scanning and tracking mode	admissible values: $0 \le$ *ScanSpeed* [kHz/sec] ≤ 400 routine values: $150 \le$ *ScanSpeed* [kHz/sec] ≤ 250
MinimumTwist	tracking mode only	admissible values: $0 \le$ *MinimumTwist* [rad] ≤ 6.28 ($= 2\pi$) routine values: $2 \le$ *MinimumTwist* [rad] ≤ 2.5

up/down sweep is added to the chosen F_{REF} waveform, with parameters setup via the AELM. The values for the scan half-width (*ScanWidth*, [kHz]) and speed (*ScanSpeed*, [kHz/sec]) can be selected by the user separately for the limiter and X-point (*diverter*) configurations. The permissible values are $0 \le$ *ScanWidth* [kHz] ≤ 400 and $0 \le$ *ScanSpeed* [kHz/sec] ≤ 400 in both limiter and X-point configurations, and the values being typically used for operation are *ScanWidth* = (100 to 250) kHz and *ScanSpeed* = (100 to 250) kHz/sec. When operating in scanning mode, the antenna frequency is swept in the interval $0 \le F_{REF} - ScanWidth \le F_{REF} \le F_{REF} + ScanWidth \le F_{MAX}$.

The *tracking* mode can be enabled only if *scanning* has also been enabled (see Fig. B3, bottom frame). In *tracking* mode, the real-time mode detection algorithm is used to lock the AELM frequency output around the frequency of the detected modes, so as to follow their evolution in time. This is done by reversing the direction of the frequency sweep when the specified resonance conditions are met, or at the extremes of the frequency scan $F_{REF} \pm ScanWidth$ when these conditions are not met. In addition to the speed and width used for the scanning algorithm, the *MinimumTwist* parameter needs to be set in *tracking* mode, again with the possibility of different values for the limiter and X-point magnetic configurations: the usual values are in the range *MinimumTwist* = (2.0 to 2.5)rad. When in tracking mode, the real-time frequency is also always bounded in the interval $0 \le F_{REF} - ScanWidth \le F_{REF} \le F_{REF} + ScanWidth \le F_{MAX}$. The parameters required for the *scanning* and *tracking* mode of operation are given in Table B2.

Finally, note that in *scanning* and *tracking* mode, a check is performed in real-time by the AELM algorithm so that the value of the frequency is always positive and always capped to the maximum set value of F_{MAX} even when the value of the scan width is such that any of these two limits could be exceeded at any particular time point.

B.3 Current and voltage limit setting using the AELM

The usual mode of operation of the AEAD plant involves remote control of the current reference (I_{REF}) waveform and of the maximum allowable voltage ($V_{ANT} \le V_{LIM}$) and current ($I_{ANT} \le \max(I_{REF})$) limit via the AELM. These functionalities are essentially needed for protection of the feedthroughs (current and voltage limit), and also to try achieving an antenna-drive which is as constant as possible, to facilitate the analysis of the data. As the end-to-end transfer function of all the current and voltage measurements has a frequency response which is not constant, the set values for the I_{REF} and V_{LIM} control and trip thresholds only strictly apply to one specific frequency point, which we have chosen to be = 200 kHz (i.e. close to the system self-resonance frequency as seen at the amplifier output, for feedthrough protection purposes).

The voltage limit (V_{LIM} output signal) control is set using the tab *SCvoltageLimit*, as shown in Fig. B4(a,b): the maximum voltage allowed on the antennas is $V_{LIM} = \max(V_{ANT}) = 700$ V-peak, and the usual value chosen for operation is $V_{LIM} = 600$ V-peak. Note that the actual voltage limit on the feedthrough, determined by Paschen breakdown of the vacuum interspace, is $V_{MAX} = 5$ kV-DC, hence the V_{LIM} functionality allows for a significant operational margin on the feedthrough protection.

The current reference (I_{REF} output signal) control is set using the tab *InputPowerSelector*, as shown in Fig. B4(a,b): this choice can be either a fixed-level value or a time-dependent waveform, both given in percentage values in the range 0 % to 100 % of the maximum current setting in the AGC hardware and the usual value chosen for operation corresponds to the maximum available antenna current so as to optimise the antenna-plasma coupling. The power level generated by the waveform editor must also start/end at power = 0 [%] because of the logic of the AELM initialisation algorithm. Note that if the current control request is *power* = 0 for any length of time, this triggers an *I-controlAlarm* in the AGC, which then stops the HV pulse. This functionality can be used to stop the AEAD power pulse for any length of time, and starting again at any time point during the selected AEAD time window. The maximum allowable current $I_{REF} = 100$ % is set directly in the AGC for protection of the feedthroughs and corresponds to a value $\max(I_{REF}) = \max(I_{ANT}) = 15$ A-peak on the antennas and $\max(I_{REF}) = \max(I_{TOT}) = 25$ A-peak on the amplifier. These values are monitored and controlled at the nominal frequency = 200 kHz when

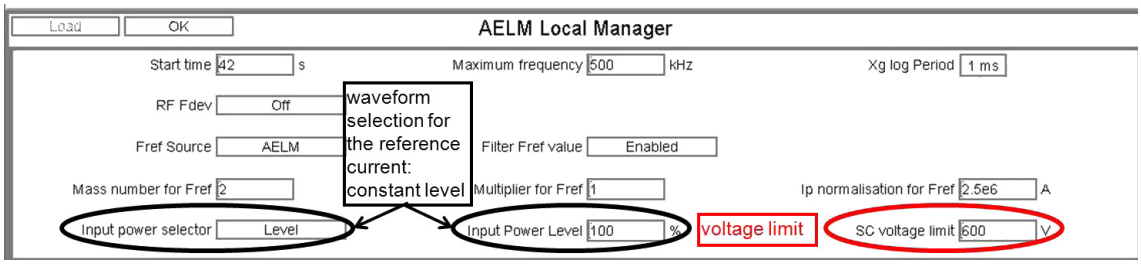

Fig. B4a Fixed-level 100 % power selection for the I_{REF} current reference control algorithm via the AELM, with maximum allowed voltage V_{LIM} = 600 V = max(V_{ANT}) on the antennas.

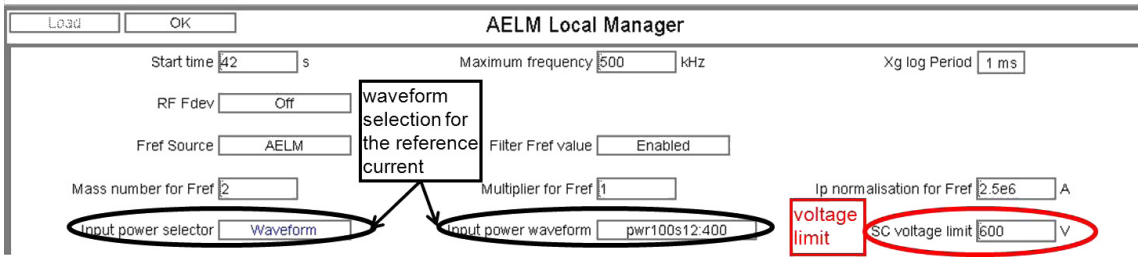

Fig. B4b Waveform selection (12 sec at 100 % level) for the I_{REF} current reference control algorithm via the AELM, again with maximum allowed voltage V_{LIM} = 600 V = max(V_{ANT}) on the antennas.

Fig. B4c AELM normalization channel selection, in this shot channel-7 = I_{ANT7} (out of the 22 possible engineering signals connected to the AELM) was selected for the AGC current control algorithm and used to normalize the measured signal for the AELM real-time mode detection and tracking algorithm (parameters not relevant for this choice of normalization channels have been shaded in this figure).

using any of these currents as the reference value for the AGC control loop. Note that the actual current limit on the feedthrough, determined by arching in the vacuum interspace, is I_{MAX} = 50 A-DC, hence using the max(I_{REF}) functionality also allows for a significant operational margin on the feedthrough protection. When a fixed-value is selected for the I_{REF} control, the user needs to set the *InputPowerLevel*, as shown in Fig. B4a: this value is kept constant for the entire 12 second duration of the AEAD time window. When a waveform is used (see Fig. B4b), then the selection is performed using the menu appearing on the tab *InputPowerWaveform*. The reference (controlled) current to which the power limit is applied is manually selected directly on the AEAD plant, and must be chosen to reflect the expected frequency range of operation and the actual antenna configuration (i.e. either I_{TOT} or one of the active antenna currents I_{ANT-X} should be used).

This choice should also be reflected in the selection of the AELM *NormalisingDetector*, used in *tracking* mode to normalize the measured signal and correctly characterise the antenna-driven resonance, as shown in Fig. B4c. The *NormalisingDetector* can in principle be any of the 22 engineering signals corresponding to a current or voltage measurement, or *OFF* (so that the normalization is not applied). For practical purposes we almost always choose it to be the same current (antenna or amplifier) selected for the AGC control loop. Depending on its definition (i.e. a current or voltage channel), a threshold is also set so that the normalization is skipped when the amplitude of the selected data falls below the specified threshold (hence avoiding spurious peaks in the measured signal that could be mistaken for a plasma resonant response to the antenna-driven perturbation, erroneously triggering the tracking algorithm). The usual values for such thresholds are 0.1 A for a current and 5.0 V for a voltage normalization channel, as shown in Fig. B4c.

B.4 Shorted-turn trip setting using the AELM

The shorted turns (S/T) trip is intended to protect the antennas and the amplifier against an excessively rapid variation in the load impedance, by tripping the HV pulse when the AELM detects a change of the antenna

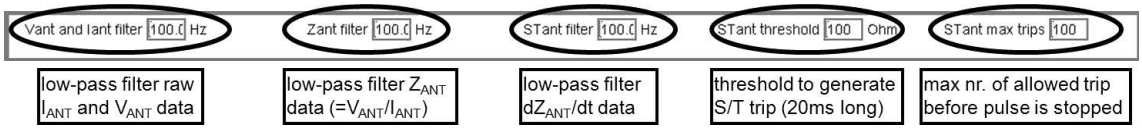

Fig. B5 Setup of the required parameters to generate the S/T trip via the AELM.

impedance, i.e. $\mathrm{d}(V_{\mathrm{ANT}}/I_{\mathrm{ANT}})/\mathrm{d}t = \mathrm{d}Z_{\mathrm{ANT}}/\mathrm{d}t$ which is too fast (in the time/frequency domain) to be due to a system (electrical) or a plasma (mode) resonance, V_{ANT} and I_{ANT} being the antenna voltage and current, respectively. These very fast variations in the antenna impedance have occurred only in concomitance with Edge Localised Modes (ELMs), as one of their main features is to "peel-off" a portion of the edge current and damp it onto the wall [17]. These events create a very time- and space- localized short circuit between the plasma and the vessel potential, which is picked-up by the in-vessel antennas in their earth return circuit, and then appears in the synchronous measurements as a very large and very rapid temporal variation in Z_{ANT}. As the antenna frequency f_{ANT} is typically swept in time during operation, the maximum allowable temporal variation $\max(|\mathrm{d}Z_{\mathrm{ANT}}/\mathrm{d}t|)$ can be determined by comparison with the relative variation $\mathrm{dlog}(Z_{\mathrm{ANT}})/\mathrm{dlog}(f_{\mathrm{ANT}})$, due to an electrical resonance in the system or to a plasma mode, and this value can then be used for operational purposes.

Practically, the S/T trip is generated in-real-time via the AELM by computing $\mathrm{d}Z_{\mathrm{ANT-X}}/\mathrm{d}t$ for all individual antennas using the engineering signals associated with the corresponding $I_{\mathrm{ANT-X}}$ and $V_{\mathrm{ANT-X}}$ measurements. The AEAD operator needs to define various parameters for the algorithm producing this S/T trip, as shown in Fig. B5. First, the input $I_{\mathrm{ANT-X}}$ and $V_{\mathrm{ANT-X}}$ measurements are low-pass filtered to remove spurious noise, which would compromise the calculation of $\mathrm{d}Z_{\mathrm{ANT-X}}/\mathrm{d}t$, and then the derivative itself is also low-pass filtered. The allowable low-pass filter frequency for all the raw $I_{\mathrm{ANT-X}}$ and $V_{\mathrm{ANT-X}}$ data, and for the processed $Z_{\mathrm{ANT-X}} = V_{\mathrm{ANT-X}}/I_{\mathrm{ANT-X}}$ and $\mathrm{STantX} = \mathrm{d}Z_{\mathrm{ANT-X}}/\mathrm{d}t$ data are all in the range 0.1 Hz to 400 Hz, with frequency = 0 Hz meaning that the filter is not applied; the usual values for operation are frequency = 100 Hz for the three filters. The usual value for the S/T trip threshold, defined in the *STantThreshold* tab is $\max(|\mathrm{d}Z_{\mathrm{ANT}}/\mathrm{d}t|) = 100\,\Omega/\mathrm{sec}$, as determined in comparison with the value of $\mathrm{dlog}(\mathrm{Z_{ANT}})/\mathrm{dlog}(\mathrm{f_{ANT}})$ due to an electrical resonance in the system or to a plasma mode. Finally, the usual value for the maximum number of allowed S/T trips during one single discharge, defined in the *STantMaxTrips* tab, is = 100: if this maximum number of trips is exceeded, a permanent alarm is raised and the HV pulse is stopped.

B.5 Setting the parameters for the real-time detection and tracking of the resonant plasma response to the antenna-driven perturbations via the AELM

User input is needed to select a number of parameters which the AELM requires for the mode detection algorithm. These parameters are selected via the tab *DetectorSignal*, and Fig. 5d illustrates their use in the basic common ideas behind both the *SimpleSum* and the *SparSpec* real-time tracking algorithms. Most of these parameters are common to the *SimpleSum* and the *SparSpec* algorithms, and are given in Table B3. Figure B6 shows the graphical implementation of these common user-selectable parameters in the AELM GUI for the *SimpleSum* algorithm, and Fig. B7(a,b) shows the corresponding GUI for the two *SparSpec* algorithms, the "*any*" (Fig. B7a) and "*highest*" (Fig. B7b) tracking modes, respectively.

As shown in these figures, the eight channels AELM-SIG/01-08 (tab: *Number* for the *SimpleSum* algorithm Fig. B6, and tab: *Sensor* for the *SparSpec* algorithms, Fig. B7) can have a positive (+ ve) or negative (− ve) sign (tab: *polarity*). The selected combination of the eight input channels gives a single output, which is then low-pass filtered (tab: *filter*) and processed according to the amplitude threshold (tab: *Thresholds/Amp*), the speed threshold at the beginning (tab: *Thresholds/Start*) and at the end (tab: *Thresholds/End*) of the frequency scan, and the minimum twist (tab: *MinimumTwist*), in limiter and X-point (≡ *diverter*) configurations. The tab *filter* defines a cut-off frequency for a low-pass filter: the default value is *filter* = 50 Hz, and if *filter* = 0 this processing step is not performed. The usual values for the tab *Thresholds/Amp* are in the range $1 \times 10^{-10} \leq$ *Thresholds/Amp* $[\mathrm{T/s}] \leq 5 \times 10^{-8}$ in the AELM own units, with a conversion factor between the AELM [T/s] and the physical [mG] units for the mode amplitude $|\delta B|$ given by $[\mathrm{T/mG}] \times (F_{\mathrm{MAX}}\,[\mathrm{V}]/(2\pi F_{\mathrm{CAL}}\,[\mathrm{MHz}])) \equiv (5/2\pi) \times 10^8$ in the majority of cases, where $F_{\mathrm{CAL}} = 200$ kHz is the frequency point used for the AELM real-time calibration. The values for the tabs *Thresholds/Start* and *Thresholds/End* are typically in the range $1 \times 10^{-13} \leq$ *Thresholds/Start/End* $[\mathrm{T/A(/V)}] \leq 5 \times 10^{-12}$ in the AELM own units, with a conversion factor given by (in the majority of cases) $[\mathrm{T/mG}] \times (F_{\mathrm{MAX}}\,[kHz]) \times (F_{\mathrm{MAX}}\,[kHz]/500) \times (10/F_{\mathrm{MAX}}\,[V]) \equiv 5 \times 10^9$ between the AELM [T/A or T/V] and the physical [mG/A/sec or mG/V/sec] units for the mode speed $\mathrm{d}(|\delta B/I_{\mathrm{ANT}}|)/\mathrm{d}t$ or $\mathrm{d}(|\delta B/V_{\mathrm{ANT}}|)/\mathrm{d}t$. The

Table B3 The common, user-selectable parameters required by both the *SimpleSum* and the *SparSpec* algorithms for real-time mode detection and tracking.

polarity	defines the signal polarity (+ve or −ve) associated to the selected sensor
filter	cut-off frequency for a real-time low-pass filter applied to all the raw data
thresholds	define the amplitude ($\|\delta B\|$) and speed $d(\|\delta B/I_{ANT}\|)/dt$ at the start/end of a frequency scan to recognize if a plasma resonance is indeed associated to a mode of sufficiently high amplitude ($\|\delta B\| > AmpThresh$) and sufficiently close to the marginal stability limit $\gamma/\omega = 0$ ($d(\|\delta B/I_{ANT}\|)/dt > SpeedThresh$) to be of interest for real-time detection and tracking
Minimum Twist	defines the minimum angle threshold ($\text{imag}(\delta B)/\text{real}(\delta B)$) in the complex plane representation to recognize if a mode detected using the amplitude and speed thresholds defined above is antenna-driven and stable ($\gamma/\omega > 0$)
Normalising Detector	used in *tracking* mode to normalize the measured: this can be any of the antenna currents or voltage, or it can be set to "off" (i.e. no normalization is used)
Cdamp Constant	defines the constant conversion factor for the real-time calculation of the damping rate as estimated from the frequency width of the detected resonance [see 18]

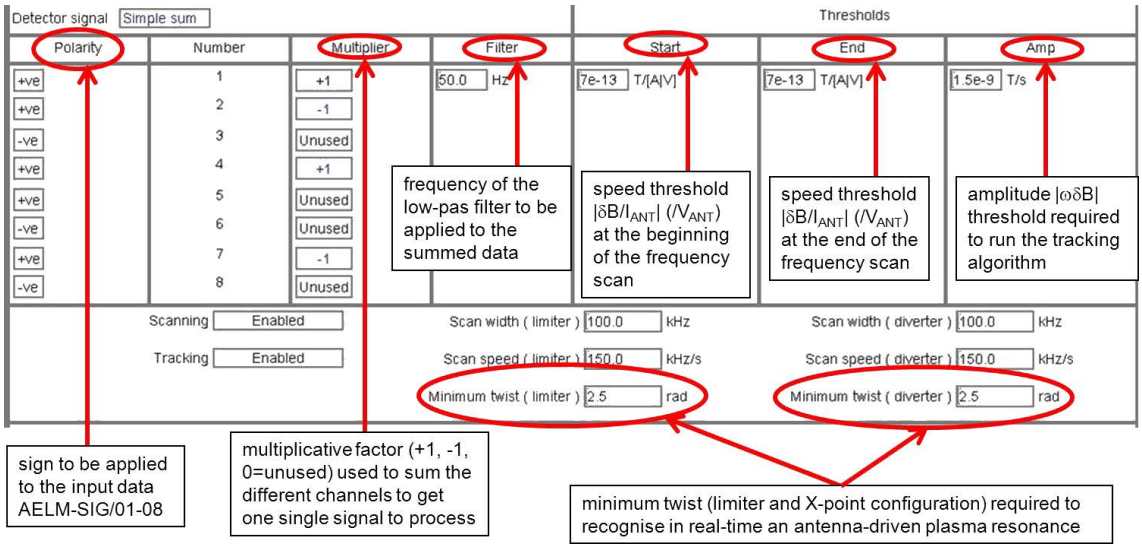

Fig. B6 Setting up the AELM resonance detection parameters using the *SimpleSum* algorithm. Note that there is a conversion factor between the AELM and the physical units for certain quantities. For the typical case considered here of F_{MAX} [V] = 10 corresponding to F_{MAX} [kHz] = 500, we have specifically a conversion factor $\equiv (5/2\pi) \times 10^8$ between the AELM [T/s] and the physical [mG] units for the mode amplitude $|\delta B|$ and $\equiv 5 \times 10^9$ between the AELM [T/A or T/V] and the physical [mG/A/sec or mG/V/sec] units for the mode speed $d(|\delta B/I_{ANT}|)/dt$ or $d(|\delta B/V_{ANT}|)/dt$.

typical values for the *MinimumTwist* tab needed to correctly identify the antenna-driven plasma resonance are in the range $2 \leq MinimumTwist$ [rad] ≤ 3 (in this case the AELM and physical units are the same).

The AELM selection panel for the *SimpleSum* algorithm is shown in Fig. B6, and Table B4 shows the additional two specific parameters other than those already shown in Table B3 (which are common to the *SparSpec* algorithm). These two additional parameters are used to define which sensors are combined (tab: *Number*), and the multiplicative factor (tab: *Multiplier*) used to construct the single output on which the tracking algorithm is then applied.

The AELM selection panel for the *SparSpec* algorithm in the *any* and the *highest* modes of operation are shown in Fig. B7(a,b) in the top and bottom frames, respectively. When the *SparSpec* algorithm is selected, the specific parameters listed in Table B5 then need to be set, in addition to those already shown in Table B3 (which are common to the *SimpleSum* algorithm). If the *SparSpec* algorithm has been selected for tracking via the *DetectorSignal* tab, then the *mode* tab allows selecting between

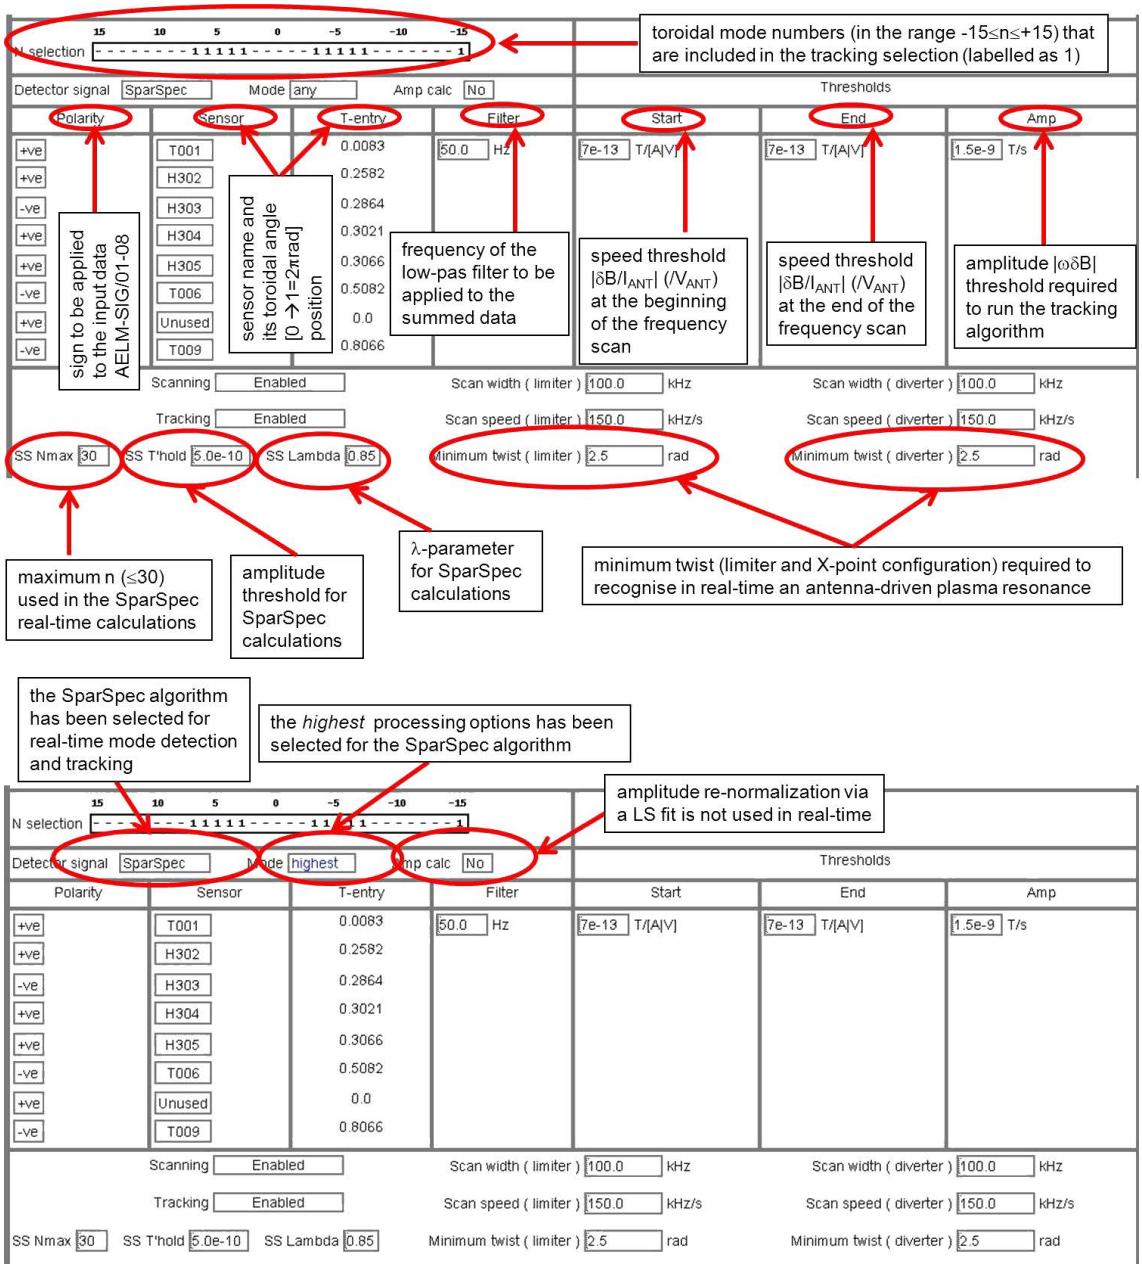

Fig. B7 Setting up the mode detection parameters for the *SparSpec* algorithm: the *any* mode of operation is shown in the top frame (Fig. B7a), and the *highest* in the bottom frame (Fig. B7b).

Table B4 The specific, user-selectable parameters required by the *SimpleSum* algorithm for real-time mode detection and tracking: these parameters are additional to those described in Table B3, which are common to both the *SimpleSum* and the *SparSpec* algorithms.

Number	defines the eight input channels (AELM-SIG/01-08) for this algorithm
Multiplier	multiplicative factors (sum = +1, difference = –1, unused = 0) used to combine the eight input AELM-SIG/01-08 and generate one single output

any and *highest*, and the *AmpCalc* tab if a least-square re-estimation of the computed mode amplitude is to be performed in real time (the usual choice being *AmpCalc* = *NO*, as this re-estimation is very computationally intensive and does not necessarily improve the real-time analysis). The *Sensor* tab then defines which magnetic sensor is connected to the corresponding AELM-SIG/01-08 channels, and the *T-entry* tab then allows inserting their

Table B5 The specific, user-selectable parameters required by the *SparSpec* algorithm for real-time mode detection and tracking, for both the "highest" and "any" processing options: these parameters are additional to those described in Table B3, which are common to both the *SimpleSum* and the *SparSpec* algorithms.

Mode	defines the option used for selecting which n-number mode has to be tracked in real-time (*highest* or *any*)
AmpCalc	defines the option used for computing the absolute amplitude of the mode via a least-square (LS) estimation (*yes* or *no*)
N-selection	defines the n-numbers $\|n\| \le N_{\mathrm{MAX}}$ that need to be detected and discriminated; since the *SparSpec* algorithm defaults to the minimum value of *N-selection* if no mode is found within the allocated CPU time, the value $\|n\| = -N_{\mathrm{MAX}}$ is kept so as to easily see when no mode has been found in real-time, as we set for $-N_{\mathrm{MAX}}$ a "mode" that is not usually associated with the resonant plasma response to the antenna drive
Sensor	defines which sensor is associated with that particular channel (can be "*unused*")
T-entry	defines the angular position ϕ_p (in the normalized range $0 \le$ *T-entry* ≤ 1 corresponding to $0 \le \phi$[rad] $\le 2\pi$) for the selected sensor associated with that particular p-th channel
SS-N_{MAX}	defines the maximum n-number to be included in the calculation of the sparse spectrum; this value must be larger than N_{MAX}, usually *SS-*$N_{\mathrm{MAX}} = 2 \times N_{\mathrm{MAX}}$
SS-T'hold	defines the background noise level threshold in the $\|\delta B_{\mathrm{MEAS}}\|$ spectrum
SS-Lambda	defines the λ_{NORM}-parameter used to penalize the L1-norm criterion for real-time mode selection evaluated by *SparSpec*

actual toroidal angle position using a look-up table. The tab *SS-*N_{MAX} allows selecting the maximum mode number for the *SparSpec* calculation (the default option being *SS-*$N_{\mathrm{MAX}} = 30$, as we take *SS-*$N_{\mathrm{MAX}} = 2 \times N_{\mathrm{MAX}}$ typically). The tab *SS-T'hold* (for *SparSpec* threshold) allows selecting the amplitude threshold below which the input signal is considered as noise, hence the processing is skipped (the default option being *SS-T'hold* $= 5 \times 10^{-10}$ in the AELM own units corresponding to ~ 0.04 mG in the physical $|\delta B_{\mathrm{MEAS}}|$ units). The *SS-Lambda* tab allows selecting the λ_{NORM} parameter ($0 \le \lambda_{\mathrm{NORM}} \le 1$) for the *SparSpec* calculation (the default option being *SS-Lambda* $= 0.85$, with the most useful range $0.3 \le$ *SS-Lambda* ≤ 0.95). Finally, the *N-selection* tab allows choosing which of the output modes produced by the *SparSpec* calculation (with mode number in the range $|n| \le$ *SS-*N_{MAX}) will be considered by the real-time tracking algorithm. An attempt is then made to follow modes that are labelled as "1" as the plasma background evolves, whereas modes labelled as "–" are ignored. Note that the *N-selection* range is limited by default to $-N_{\mathrm{MAX}} = -15 \le NN \le +15 = +N_{\mathrm{MAX}}$ in order to remain within the $\sim 800\,\mu$s CPU time allocated to the real-time calculations. Moreover, the $NN = -15$ must always be selected (even if this is not necessarily a mode of interest) as it is used in the AELM to quickly differentiate in real-time satisfactory ($NN \ne -15$) and unsatisfactory ($NN = -15$) solutions.

[1] K. Ikeda *et al*., "Progress in the ITER Physics Basis", Nucl. Fusion **47** (2007).
[2] E.J. Strait, E.D. Fredrickson, J-M. Moret and M. Takechi, Fusion Sci. Technol. **53**, 304 (2008).
[3] A. Fasoli *et al*., Phys. Rev. Lett. **75**, 645 (1995).
[4] C.Z. Cheng, L. Chen and M.S. Chance, Ann. Phys. **161**, 21 (1985).
[5] R. Betti and J.P. Freidberg, Phys. Fluids B **3**, 1865 (1991).
[6] M.S. Chu *et al*., Phys. Fluids B **4**, 3713 (1992).
[7] A.D. Turnbull *et al*., Phys. Fluids B **5**, 2546 (1993).
[8] Q.A. King and H. Brelen, *An Experimental Control Facility at JET*, JET-P(98)24 (1998) Joint European Torus, Abingdon (UK) (weblink: http://www.iop.org/Jet/fulltext/JETP98024.pdf).
[9] A. Fasoli, D. Testa *et al*., Plasma Phys. Control. Fusion **44**, 159 (2002).
[10] S. Sharapov, D. Testa *et al*., Phys. Lett. A **28**, 127 (2001).
[11] D. Testa *et al*., *The new Alfvén Wave Active Excitation System at JET*, Proceedings 23rd SOFT Conference (2004) (weblink: http://infoscience.epfl.ch/record/143354/files/).
[12] D. Testa *et al*., Nucl. Fusion **50**, 084010 (2010).
[13] T. Panis, D. Testa *et al*., Nucl. Fusion **50**, 084019 (2010).
[14] D. Testa *et al*., Fusion Eng. Des. **86**, 381 (2011).
[15] D. Testa *et al*., Europhysics Letters **92**, 50001 (2010).
[16] A. Fasoli *et al*., Nucl. Fusion **36**, 258 (1996).
[17] D. Testa and M. Bigi, Plasma Phys. Control. Fusion **47**, 733 (2005).
[18] D. Testa and A. Fasoli, Nucl. Fusion **41**, 809 (2001).
[19] A. Fasoli *et al*., Phys. Plasmas **7**, 1816 (2000).
[20] W.W. Heidbrink, Phys. Plasmas **9**, 2113 (2002).
[21] G. Vlad, S. Briguglio, G. Fogaccia and F. Zonca, Nucl. Fu-

sion **46**, 1 (2006).
[22] D. Testa, D. Spong, T. Panis, P. Blanchard and A. Fasoli, Nucl. Fusion **51**, 043009 (2011).
[23] T. Panis, A. Fasoli and D. Testa, Nucl. Fusion **52**, 023013 and 023014 (2012).
[24] D. Testa, N. Mellet, T. Panis *et al*., *Recent JET Experiments on Alfvén Eigenmodes with Intermediate Toroidal Mode Numbers: Measurements and Modelling*, Paper EXW/P7-27, 23rd IAEA Fusion Energy Conference, Daejon, Republic of Korea, 11-16 October 2010, (weblink: http://infoscience.epfl.ch/record/153045).
[25] D. Testa, T. Panis, P. Blanchard and A. Fasoli, Nucl. Fusion **52**, 094006 (2012).
[26] S. Bourguignon, H. Carfantan and T. Böhm, Astron. Astrophys. **462**, 379 (2007).
[27] S. Bourguignon, H. Carfantan and J. Idier, IEEE Journal of Selected Topics in Sig. Proc. **1**, 4 (2007).
[28] A. Klein, H. Carfantan, D. Testa *et al*., Plasma Phys. Control. Fusion **50**, 125005 (2008).
[29] L.L. Lao, H. St John, R.D. Stambaugh, A.G. Kellman and W. Pfeiffer, Nucl. Fusion **25**, 1611 (1985).
[30] L. Villard and J. Vaclavik, Nucl. Fusion **37**, 351 (1997).
[31] H. Yamada for the LHD Experiment Group, Nucl. Fusion **51**, 094021 (2011).
[32] K. Kwon *et al*., Nucl. Fusion **51**, 094006 (2011).
[33] L.W. Yan *et al*., Nucl. Fusion **51**, 094016 (2011).
[34] Y. Kamada *et al.*, Nucl. Fusion **51**, 073011 (2011).
[35] K. Ikeda, Nucl. Fusion **50**, 014002 (2010).
[36] K. Tobita *et al*., Nucl. Fusion **49**, 075029 (2009).
[37] F. Marvasti, *Non-Uniform Sampling: Theory and Practice* (Springer-Verlag, 2001).
[38] N.R. Lomb, Astrophys. Space Sci. **39**, 447 (1976).
[39] J.D. Scargle, Astrophys. J. **263**, 835 (1982).
[40] D. Roberts, J. Lehar and J. Dreher, Astron. J. **93**, 968 (1987).
[41] G. Foster, Astron. J. **109**, 1889 (1995).
[42] J.S. Kim *et al.*, Plasma Phys. Control. Fusion **41**, 1399 (1999).
[43] M. Hole and L. Appel, Plasma Phys. Control. Fusion **49**, 1971 (2007).
[44] F.M. Poli *et al.*, Plasma Phys. Control. Fusion **50**, 095009 (2008).
[45] E.P. Wigner, Phys. Rev. **40**, 749 (1932).
[46] H. Choi and W. Williams, IEEE Trans. Acoust. Speech Signal Process **37**, 86 (1989).
[47] N.E. Huang *et al.*, Proc. Royal Society London **A454**, 903 (1998).
[48] S. Zegenhagen, A. Werner *et al*., Plasma Phys. Control. Fusion **48**, 1333 (2006).
[49] T. Böhm, C. Catala, L. Balona and B. Carter, Astron. Astrophys. **427**, 907 (2004).
[50] A. Fumel and T. Böhm, Astron. Astrophys. **540**, A108 (2012).
[51] C.E. Shannon, *Communication in the presence of noise*, Proc. Institute of Radio Engineers 37 (1949), 10; reprinted in: Proc. IEEE **86**, 447 (1998).
[52] J. Lafler and T.D. Kinman, Astrophys. J. Suppl. **11**, 216 (1965).
[53] M.M. Dworetsky, Mon. Not. R. Astron. Soc. **203**, 917 (1983).
[54] P. Stoica, R.L. Moses, B. Friedlander and T. Söderström, IEEE Trans. ASSP **37**, 378 (1989).
[55] M.D. Sacchi, T.J. Ulrych and C.J. Walker, IEEE Trans. Signal Proc. **46**, 31 (1998).
[56] D.L. Donoho, J. Communications on Pure Applied Math. **59**, 797 (2006).
[57] D.M. Malioutov, *A Sparse Signal Reconstruction Perspective for Source Localization with Sensor Arrays*, Master Thesis, MIT (2003).
[58] S.S. Chen, D.L. Donoho and M.A. Saunders, SIAM J. Sci. Comput. **20**, 33 (1998).
[59] J-J. Fuchs, IEEE Trans. Inf. Theory **50**, 1341 (2004).
[60] J.A. Tropp, IEEE Trans. Inf. Theory **52**, 1030 (2006).
[61] J.A. Tropp, IEEE Trans. Inf. Theory **50**, 2231 (2004).
[62] P. Moulin and J. Liu, IEEE Trans. Inf. Theory **45**, 909 (1999).
[63] J-J. Fuchs, IEEE Trans. Inf. Theory **51**, 3601 (2005).
[64] E. Candès, J. Romber and T. Tao, Comm. Pure Applied Math. **59**, 1207 (2006).
[65] V. Merezhkin, Soviet JETP **4**, 152 (1978).
[66] O. Kluber *et al.*, Nucl. Fusion **31**, 907 (1991).
[67] D. Testa *et al*., Rev. Sci. Instrum. **74**, 1694 (2003).
[68] D.L. Donoho, M. Elad and V.N. Temlyakov, IEEE Trans. Inf. Theory **52**, 6 (2006).
[69] J.A. Tropp, IEEE Trans. Inf. Theory **52**, 1030 (2006).
[70] D. Testa *et al*., Fusion Sci. Technol. **57**, 208 (2010).
[71] D. Testa *et al*., Fusion Sci. Technol. **57**, 238 (2010).
[72] H.W. Kuhn and A.W. Tucker, *Nonlinear programming*, Proceedings of 2nd Berkeley Symposium (Berkeley: University of California Press, 1951) p.481.
[73] W. Karush, *Minima of Functions of Several Variables with Inequalities as Side Constraints*, M.Sc. Dissertation, Dept. of Mathematics, Univ. of Chicago, Chicago, USA (1939).
[74] D.P. Bertsekas, *Nonlinear Programming* (2nd edition) (Athena Scientific, Belmont, USA, 1999).
[75] S. Sharapov and D. Testa, Phys. Lett. A **289**, 127 (2001).
[76] S. Sharapov *et al*., Phys. Plasmas **9**, 2027 (2002).
[77] Y. Kusama *et al*., Nucl. Fusion **38**, 1215 (1998).
[78] H. Kimura *et al*., Nucl. Fusion **38**, 1303 (1998).
[79] H.L. Berk *et al.*, Phys. Rev. Lett. **87**, 185002 (2001).
[80] D. Testa and M. Albergante, Nucl. Fusion **52**, 083010 (2012).
[81] D. Testa and M. Albergante, Europhysics Letters **97**, 35003 (2012).