EFDA–JET–PR(10)16

D. Testa, H. Carfantan, A. Goodyear, P. Blanchard, A. Klein, T. Panis
and JET EFDA contributors

An Algorithm for the Real-Time & Blind Detection, Decomposition & Tracking of the Individual Components in a Degenerate, Multi-Harmonics Spectrum

An Algorithm for the Real-Time & Blind Detection, Decomposition & Tracking of the Individual Components in a Degenerate, Multi-Harmonics Spectrum

D. Testa[1], H. Carfantan[2], A. Goodyear[3], P. Blanchard[1,4], A. Klein[5], T. Panis[1]
and JET EFDA contributors*

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

[1]Ecole Polytechnique Fédérale de Lausanne (EPFL), Centre de Recherches en Physique des Plasmas, Association EURATOM – Confédération Suisse, CH–1015 Lausanne, Switzerland
[2]Laboratoire d'Astrophysique de Toulouse – Tarbes (LATT), Université de Toulouse – CNRS, 31400 Toulouse, France
[3]EURATOM-CCFE Fusion Association, Culham Science Centre, OX14 3DB, Abingdon, OXON, UK
[4]JET-EFDA Close Support Unit, Culham Science Centre, OX14 3DB, Abingdon, OXON, UK
[5]formerly at the Plasma Science and Fusion Centre, Massachusetts Institute
** See annex of F. Romanelli et al, "Overview of JET Results",*
(Proc. 22 [nd] IAEA Fusion Energy Conference, Geneva, Switzerland (2008)).

Preprint of Paper to be submitted for publication in
Europhysics Letters

ABSTRACT.

In this work we report the successful application of an innovative method, based on the Sparse Representation of signals, to perform a real-time, unsupervised and blind detection of the individual components in a frequency degenerate, multi-harmonics spectrum, using a small number of data un-evenly sampled in the spatial domain. This method has been developed from its original applications in astronomy, and is now routinely used in the Joint European Torus thermonuclear fusion experiment [1, p.617] to obtain the decomposition of a spectrum of high-frequency (~10-500kHz range) magnetic instabilities with a faster-than-1ms time resolution, allowing the real-time tracking of its individual components as the plasma background evolves. This work opens a clear path towards developing real-time control tools for electro-magnetic instabilities in future fusion devices aimed at achieving a net energy gain, such as the ITER facility [1, p.711] currently being built in France. More generally, the very high speed and accuracy of this algorithm is recommended for application to instances of physics measurements and control engineering where a real-time, blind and unsupervised decomposition of a degenerate input signal is required from a small number of data.

1. INTRODUCTION

The problem of blind detection of different components in a degenerate multi-harmonics spectrum using un-even sampling is common to various fields of physics and engineering (see for instance the many applications presented in Ref.[1]). This problem is particularly important in the field of Astronomy and Astrophysics (A&A), where the very long measurement series are often interrupted due to different environmental constraints (such as the weather and the Earth's rotation). Much work has been done over the past 20 years to address this mathematical issue and improve on the limitations of the original methods, which were essentially based on the Lomb-Scargle periodograms [2-5].

In the field of plasma physics, analysis of electro-magnetic fluctuations is important for understanding and controlling the Magneto-HydroDynamic (MHD) stability of the plasma. These instabilities can occur when a plasma is trapped and heated in a magnetic configuration and responds by generating its own, usually oscillating, magnetic field, i.e. an MHD instability. Considering now the particular case of magnetically confined thermonuclear fusion plasmas in a tokamak device [6], the MHD analysis is based on magnetic and turbulence measurements, and routinely involves an initial Fourier decomposition of the data in the time/frequency domain to obtain the individual frequency components $\psi(\omega)$. As in a tokamak the plasma column has, to a first approximation, 2D boundary conditions along the longitudinal (toroidal) axis and on the plane perpendicular to it (the poloidal direction), the spatial structure of the MHD instabilities is determined by further decomposing each frequency component in its toroidal (n) and poloidal (m) harmonics: $\psi(\omega)=e^{-i\omega t}\Sigma_{n,m}A_{mn}e^{in\phi}e^{im\theta}$. 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 is

usually made up of multiple poloidal components *m's* due to toroidicity and various other geometrical effects. It is to be noted that the data are actually acquired only at some specific angle positions ϕ_p and θ_p, $p=\{1, \ldots, P\}$, i.e. there is not a continuous measurement coverage of the toroidal and poloidal angle coordinates.

The number of mathematical methods currently used to analyse MHD instabilities in tokamaks is rather high. With evenly spaced sensors, a simple discrete Fourier transform reveals the mode amplitude for each individual (n, m) component up to the Nyquist number [7]. However, not only is this measurement arrangement not usually available because of engineering and installation constraints (essentially related to the large number of sensors required to perform a truly Nyquist-based analysis), but it can also be shown that an even sensor spacing does not necessarily give the best measurement performance because of aliasing effects [1]. Hence, other methods have been developed and applied to fusion plasmas in tokamaks, such as the Singular Value (SVD) [8, 9] and the wavelet [10] decomposition, the Wigner [11], Choi-Williams [12] and Hilbert [13] Transforms, and a generalization of the original Lomb-Scargle periodograms [14]. All these methods work very well for a non-degenerate spectrum of MHD modes, i.e. the case where the different frequency components are at sufficiently well separated frequencies, such that the half-width at full-maximum of two close-by modes (which is closely related to their damping and growth rate $\gamma/\omega=imag(\omega)/\omega$) is much smaller than their separation in frequency. However, multiple instabilities often overlap in frequency space, producing a frequency-degenerate spectrum, and the resulting signal in each individual sensor is the superposition of these modes. In that case the task of a blind spatial decomposition becomes a much more difficult problem, and the methods indicated above are in general unsuitable in such cases because of intrinsic mathematical difficulties, or more simply because of the too long computational time required to perform the necessary calculations. This clearly prevents these methods being applied to the real-time detection of MHD instabilities in tokamaks either for protection of the machine or for optimization of the fusion performance (the so-called burning-plasma scenarios [15]), a functionality very much required for future thermonuclear fusion devices aimed at achieving a net energy gain, such as ITER [6, p.711; 16].

The problem of a blind and unsupervised decomposition of a multi component, frequency-degenerate spectrum, is a typical case where scientific cross-fertilization between two (not-so) similar fields, in our case thermonuclear fusion in tokamak devices and A&A, yields excellent results. Considering the similarities between these two fields, it is clear that an un-evenly distributed measurement time series of A&A data corresponds to an un-evenly spaced measurement array aligned in the toroidal (poloidal) direction in a tokamak. Hence, the canonical time-conjugate in A&A (i.e. a temporal frequency) corresponds to the spatial toroidal (poloidal) mode number in a tokamak. There are, however, minor differences between these two fields: real-valued data and real-valued temporal frequencies in A&A versus complex electro-magnetic fluctuation data and integer-valued (positive and negative) mode numbers in thermonuclear fusion plasmas in tokamaks. Furthermore, the role of the mean value of the data is different in A&A and tokamak plasmas: in the latter this represents the

n=0 (*m=0*) mode, which can be responsible for a major magnetic instability leading to an abrupt vertical displacement of the entire plasma column, which has to be correctly detected to avoid an unwanted plasma termination. In the former, the mean value is generally of no physical interest and even perturbs the analysis as the mean of the data is not generally equal to the theoretical mean of the signal due to the irregular sampling. Finally, the long measurement series of A&A data do not call for any real-time analysis whereas in tokamak plasma, sub-millisecond calculations are routinely required in order to protect the machine from the dangerous effect of MHD instabilities.

To tackle the problem of a blind signal decomposition of irregularly sampled A&A data, a new method has been recently proposed for fitting cissoids to such data [17]. This algorithm is based on the *Sparse Representation of Signals*, as implemented in the *SparSpec* code (freeware available at: HYPERLINK "http://www.ast.obs-mip.fr/Softwares" http://www.ast.obs-mip.fr/Softwares). This method has also been initially adapted for post-pulse data analysis of MHD instabilities in the Joint European Torus (JET) tokamak [6, p.617], where it has been fully benchmarked with simulated and real data [18]. An algorithm based on this method has also recently been used for the analysis and optimization of the measurement performance of the proposed MHD magnetic diagnostic system for ITER [19]. For JET data, the *SparSpec* method has proven to be extremely robust, and is especially useful for resolving the amplitudes and phases of multiple Alfvén Eigenmodes (AEs) (see chapter 7.15 in [6]) in the ~100-300kHz frequency range, which are ringing with the same or nearly the same frequency. The great efficiency with which the *SparSpec* method detects multiple modes in large datasets has then suggested that it may be used in real-time applications, not only for AEs but also for the detection and control of MHD instabilities with much lower frequencies, ~10kHz or even below, such as resistive wall modes (see chapters 6.7 and 12.5 in [6]) and tearing modes (see chapters 6.8 and 7.3 in [6]), among others.

This paper reports on the real-time development of the *Sparse Representation* method, and on its specific application to the sub-millisecond detection, discrimination and tracking of the individual toroidal mode numbers in the multi- components, frequency-degenerate spectrum of stable AEs which are excited in the JET tokamak by an array of external antennas used for MHD diagnostic purposes. This paper is organized as follows. In Section2 we review the mathematical foundation of the *Sparse Representation* method and of the numerical approach used in the *SparSpec* code. Section3 gives a brief overview of the active MHD diagnostic system used in JET, with particular attention to its real-time plant control and data analysis hardware and software. Section4 then shows the first examples of the application of the real-time version of the *SparSpec* code to the detection and discrimination of the different toroidal components in the multi-harmonics spectrum driven by the active MHD diagnostic system. Finally, in Section5 we briefly summarize our results and give an outlook towards future work.

2. THE SPARSE REPRESENTATION METHOD AND THE SPARSPEC PRINCIPLE

In the standard tokamak coordinate system (ϕ, θ), magnetic perturbations can be represented by

functions involving toroidal and poloidal harmonics; considering now the usual case of a perturbation with a specific toroidal mode number n, this can be written as $\psi(\omega,n)=e^{-i\omega t}e^{in\phi}\Sigma_m A_{mn}e^{im\theta}$, where each mode has one single toroidal mode number, but includes several poloidal Fourier harmonics. The aim of toroidal (poloidal) mode numbers detection is to detect the mode numbers n (m) of the magnetic instabilities actually present in the plasma and to estimate their amplitude from data acquired with detectors unevenly positioned at angles ϕ_p (θ_p) in radians, $p=\{1, \ldots, P\}$ being the suffix labelling the individual sensors used for the measurement. Mathematically, each measurement y_p is modelled as:

$$y_p = \sum_{l=1}^{L} \alpha_l e^{in_l\phi_p} + \varepsilon_p \tag{1}$$

where n_l and α_l are the unknown mode numbers and amplitudes, respectively, L the unknown number of modes and ε_p corresponds to the noise on the data for the given p-th sensor. This problem, which amounts to fitting multiple cissoids to the input data, is a very general signal processing problem which arises in many fields of physics. However, it 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. 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 caused by the Earth's rotation and revolution, and a-periodic 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; these are based on generalizations of the Lomb-Scargle periodogram [2, 3] and Date-Compensated Discrete Fourier Transform [20], involving iterative analysis [4, 5], generally used when dealing with a large number of data points, or fitting periodic signals (Phase Dispersion Minimisation [21], string length method [22]) to short data strings. Such methods, however, may all still fail when there are several spectral peaks (i.e. temporal frequencies) and an insufficient number of measurement points to resolve them.

The mathematical problem described by eq.(1) can be expressed equivalently in the Fourier spatial domain (i.e. using a Fourier transform with respect to the angular position ϕ) as:

$$\hat{y}(\nu) = \sum_{p=1}^{P} y_p e^{i\nu\phi_p} = \hat{W}(\nu) * \sum_{l=1}^{L} \alpha_l \delta(\nu - n_l) + \hat{\varepsilon}(\nu) \tag{2}$$

where δ is the Dirac delta function, the "^" symbol indicates the Fourier transform of the data ($\hat{y}(\nu)$) and noise ($\hat{\varepsilon}(\nu)$) samples, the symbol "$*$" is the convolution operator and $\hat{W}(\nu)$ is the *spectral window* of the sampling scheme. 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 of secondary lobes. In A&A, when accounting for long times series, high secondary lobes are due to periodicities, typically daily, of the temporal gaps in the measurement (see fig1, top frame), while 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 (see fig.1, bottom).

When applied to thermonuclear plasma physics, the problem has some additional specificities. 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. Second, the modes numbers n_l can only take positive or negative integer values, while in the general spectral analysis problem frequencies takes real values. Third, in the real time applications we consider for JET, a set of data is acquired every 1ms, therefore the spectral analysis must be completed in an unsupervised manner in the short time between each measurement acquisition. From an estimation viewpoint, evaluating the amplitudes a_l and the modes number n_l is a very difficult problem. For example, consider the best lest-square (LS) fitting: even if the number of modes is known, the LS criterion has many local minima in case of real valued spectral peaks [23] and requires a combinatorial exploration for integer-valued mode numbers n_l. A way to circumvent the problem is to estimate the amplitudes of all mode numbers in the range $\{K, \ldots, K\}$ (where $|K|$ is the maximum mode number), but to enforce the fact that most of these modes have a null amplitude. This amounts to approximating the data with the best linear combination of a small number of elementary known signals, which is called a *Sparse Approximation*.

Formally, *Sparse Representations* of signals [24-26] are representations that account for all information of a signal with a linear combination of a small number of elementary signals called *atoms*. The *atoms* set does not form a basis as the number of atoms exceeds the dimension 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 *Sparse Approximations* of signals is the version of the *Sparse Representations* adapted to noisy data. Theoretically, the *Sparse Approximation* problem consists of minimizing the criterion:

$$J_0(x) = \left|\left|y - Wx\right|\right|^2 + \gamma\left|\left|x\right|\right|_0 \tag{3a}$$

Here $\boldsymbol{y}=[y_1, y_2, \ldots, y_P]^T$ is the vector of data taken at position ϕ_P; $\boldsymbol{x}=[x_1, x_2, \ldots, x_M]^T$ is the vector of complex amplitudes, $\boldsymbol{W}=[\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_M]$ is a matrix where the vector $\boldsymbol{w}_k$ corresponds to the k-th atom, the $L0$-norm of $\boldsymbol{x}$: $\|\boldsymbol{x}\|_0 = \#\{k, |x_k| \neq 0\}$ is the number of non-zero components of $\boldsymbol{x}$ and γ is a penalization parameter. However, to minimize this criterion, one must 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 *greedy pursuit* algorithms, iteratively adds atoms to the approximation of the signal to improve this approximation [27]. The second one, often called *convex relaxation*, replace the $L0$-norm in criterion (3a) with another penalization term such that the criterion may be minimized more easily. In our work, we follow this *convex relaxation* approach, using the classically used $L1$-norm (see for instance [24-26]) as it leads to the convex criterion:

$$J_0(x) = \left\| y - Wx \right\|^2 + \gamma \left\| x \right\|_1 = \left\| y - Wx \right\|^2 + \lambda \sum_{k=-K}^{K} \left(\left| x_k \right| \right) \quad (3b)$$

It can then be easily shown that the criterion of eq.(3b) is convex, therefore has no local minima, but, as the number of unknowns may be larger than the number of data, this criterion is not strictly convex, i.e. the solution cannot be a-priori guaranteed to be unique. In practice, minimizing this *L1*-norm penalized least-square fitting criterion is easiest than minimizing the original *L0* one, and many computationally efficient algorithms can be used, compatible with a real-time system for the problem considered in this paper. However, minimizing eq.(3b) does not lead to the same solution than minimizing eq.(3a). Theoretical conditions guaranteeing the equivalence of both solutions have been established, which are based on properties of the matrix $\boldsymbol{W}$, so depend on the specificities of the problem being considered.

The choice of the family of *atoms* is critical in the *Sparse Representations* of signals as, with an appropriate choice, these *atoms* might be well adapted to the signal to be analyzed 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 $\mu = max_{k \neq l}(|\boldsymbol{w}_k^H \boldsymbol{w}_l|)$, where $\boldsymbol{W}^H$ is the Hermitian transposition of $\boldsymbol{W}$, then minimizing eq.(3b) will lead to select the same atoms than the solution minimizing eq.(3a) [28]. For these reasons, the matrix $\boldsymbol{W}$ is often chosen as a union of incoherent basis, such as wavelets, Diracs, pure sine waves, or other mathematical approaches.

In our case, the *atoms* are imposed by the model setup in eq.(2) 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 $|w_k^H w_l| = \hat{W}(n_l - n_l)$, so that it corresponds to regular samples of the spectral window. As $\hat{W}(v)$ may take values greater than 1/3 (as shown in fig.1), 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 comparison between different numerical methods that such a solution generally gives very satisfactory results in terms of detection, even in the case of multiple modes. 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 [29].

In terms of amplitude estimations, it has been shown that minimizing eq.(3b) leads to an under-estimation of the amplitudes of the detected mode numbers due to the *L1*-norm penalization term of eq.(3b) [17, 29]. Thus, an *a posteriori* least square estimation of these amplitudes is performed in a second step within the calculations, after the modes are actually detected. The amplitudes of the detected modes are computed, minimizing the least square criterion $\|\boldsymbol{y} - \boldsymbol{W}_{DET}\boldsymbol{x}_{DET}\|^2$ where only the non-zero amplitudes of the optimization step are preserved in $\boldsymbol{x}_{DET}$.

A real-time implementation of the proposed modes detection method requires not only an efficient optimization algorithm to minimize eq.(3b) but also, and even more importantly for a frequency-degenerated spectrum, an efficient unsupervised tuning of the penalization parameter λ. Note that

the *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. Many numerical algorithms are available to minimize criteria such as those of eq.(3b) 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* [30]. An algorithm based on an *Iterative Block Coordinate Descent* procedure has been previously proposed [17, 29], 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 1ms using the rather modest computational resources available to process real-time JET data.

Tuning the penalization parameter λ, which is related to the noise level, is still an open question for *Sparse Approximations*. For spectral analysis, this parameter has interesting physical meanings. Indeed, it can be shown that (a) for $\lambda>\lambda_{MAX} = \max(|W^H y| = \max_k(|\hat{y}(n_k)|)$, the minimizer $\boldsymbol{x}_{MIN}$ of (3b) is identically zero, i.e. the unique solution is the zero solution (no detected modes); and (b) for a given λ, the minimizer $\boldsymbol{x}_{MIN}$ of eq.(3b) satisfies $\max(|W^H(y-Wx_{\mathrm{MIN}})|) = \max_k(|\hat{r}(n_k)|) \le \lambda$, where $r=\boldsymbol{y}-\boldsymbol{W}\boldsymbol{x}_{\boldsymbol{MIN}}$ 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 periodogram of the residual. An intuitive description of λ is that, since it increases the penalty for those solutions which invoke a large number of modes, it influences the detection ability of *SparSpe*c (hence its predominant role when analysing frequency-degenerated spectra). Smaller values of λ will call upon a larger number of modes to fit the data, while larger values have the effect of making the method more immune to the effects of noise, however at the expense of disregarding "true" modes whose amplitude is much lower than the dominant one(s). The precise optimum value for λ depends upon the dataset and the level of noise, the number and the spacing between the sensors and, finally, the maximum mode number contained in the data. Hence, knowledge of the noise level (noise variance) in the measurements helps to determine this parameter. An "L-curve" technique described in [17, 29] can also be used to choose the optimum value of λ for a given data set. As $|\boldsymbol{W}^H\boldsymbol{y}|$ corresponds to the periodogram of the data, it can be concluded that choosing λ to be a fraction $\lambda_{NORM}\in[0, 1]$ of the maximum of the periodogram of the data $\lambda=\lambda_{NORM}\times max(|\boldsymbol{W}^H\boldsymbol{y}|)$, ensures the periodogram of the residual $\boldsymbol{r}$ to be lower up to this fraction relatively to the maximum of the data periodogram. The typical values for λ_{NORM} range between $\lambda_{NORM} = 0.1$ in A&A where the number of time samples is very large (typically $>10^3$), to $\lambda_{NORM}=0.95$ when the *SparSpec* algorithm is applied to the real-time detection of toroidal mode numbers in a JET experiment, as described in this paper. In practice, and following very detailed simulations using various models and direct estimates for the noise on each of the magnetic sensors used for real-time MHD analysis at JET, the most suitable value for λ_{NORM} was determined to be $\lambda_{NORM}=0.8$. This value allows for a very rapid

convergence of the optimization algorithm (typically within ~600microseconds) and is sufficient to detect and discriminate multiple modes whose amplitudes are of interest for MHD diagnostic purposes (physics, plant protection and control issues) at JET.

3. THE ACTIVE MHD DIAGNOSTIC SYSTEM IN USE AT JET

A key physics issue for a usable fusion reactor is the understanding and control of the burning plasma regime, a situation in which the energy carried by the fusion produced alpha particles (αs) exceeds that externally injected. One of the main elements of this regime is the interaction of the αs with waves that are naturally excited in the plasma. Such interaction can be resonant, lead to efficient energy and momentum exchange between waves and particles (hence to phase-space diffusion) and drive instabilities, tapping the free energy contained in the as pressure gradient. If a significant spatial redistribution of the αs occurs, then the overall plasma performance can be limited; moreover, if this redistribution goes as far as the machine boundaries, causing net losses of the αs, then damage to the first wall can also occur. Conversely, the knowledge of the mechanisms behind the mode stability, the interaction of the modes with the αs and their redistribution can be turned into tools for the control of their phase-space behaviour and, therefore, the plasma burn itself.

One example of waves that can interact resonantly with the αs is Alfvén Eigenmodes (AEs): these MHD modes are particularly important as they are a natural Eigenmode of any magnetically confined plasma, and also because the fusion-produced αs are born with a supra-thermal speed that is typically super-Alfvénic in the usual thermonuclear tokamak plasma conditions. Therefore, resonant interaction with AEs is the first wave-particle interaction that the αs encounter during their thermalization: hence, this mechanism for phase-space and spatial diffusion needs to be controlled appropriately to guarantee good confinement of the αs themselves. Fortunately, AEs occupy a rather empty portion of the plasma electromagnetic fluctuation spectrum, which is well above the frequencies related to gradient-driven drift instabilities, and well below the gyro-frequencies of all plasma species. Hence, AEs constitute a "unique and very clean way to communicate" with the plasma, which allows them to be used as a powerful diagnostic tool, not only for the αs, but also for the background plasma. This measurement technique using MHD waves that are naturally supported by the plasma is known as *MHD spectroscopy* [31], and 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 [32]. This is the so-called Alfvén Eigenmodes Active Diagnostic (AEAD) system, and fig.2 shows a very simplified schematic overview to illustrate its main features. The AEAD system principally consists of:

1. the AE exciter, built upon a function generator and a (high-power) amplifier connected to a set of in-vessel antennas (up to 8 in our case), whose aim is to send power into the plasma in order to drive a very small magnetic perturbation at the plasma edge, with maximum intensity $max(|\delta B_{DRIVEN}|) \sim 1G$, i.e. 10^4 times smaller than the typical value of the toroidal magnetic field in JET, $B_{TOR} \sim (1\text{-}3)T$;

2. a receiver, built upon synchronous detection units, which is collecting 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 real-time AE Local Manager (AELM) to allow for the detection and tracking of antenna-driven plasma resonances.

The AE exciter is built upon a 5kW class-B power amplifier capable of delivering up to a maximum I_{ANT}~10A-peak and a maximum V_{ANT}~1kV-peak in the frequency range 10kHz→500kHz to each of the 8 in-vessel antennas. The antennas are installed in two groups of four closely-spaced units located at two toroidally opposite positions but at the same poloidal location. This allows the magnetic fluctuations to be driven with a sufficiently high $|\delta B_{DRIVEN}|>10^{-3}G$ at the plasma edge for a spectrum of toroidal mode number extending up to $|n|\sim 30$.

The plasma response to the antenna-driven perturbation is principally measured on a selected subset of signals using synchronous detection [32]. Looking at any electro-magnetic turbulence measurements, such as the signals from the magnetic pick-up coils, we would see that we measure a multitude of mixed frequencies. Conversely, for real-time use of the AEAD diagnostic, we need to measure only the plasma response at the frequency corresponding to the antenna excitation, i.e. the *synchronous* component. The purpose of the synchronous detection system is to select only the portion of the input signal which is at the same frequency as the one chosen for the function generator: this also correspondingly reduces the required bandwidth of the data acquisition system (which has a 1.25kHz sampling rate for 12sec of each JET pulse, compared to the 1MHz sampling rate typical for the other JET fluctuation measurements), and also filters-out all other unwanted frequency components. The measured signal is then directly in relation (i.e. it is synchronous) to the launched antenna signal. The synchronous detection hardware works conceptually by applying a 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 <100Hz bandwidth to generate the output [cosine (I), sine (Q)] DC components. A schematic diagram of the synchronous detection hardware is shown in fig3. Synchronous detection is essential for the real-time applications (with plant protection and control using 24 channels, data analysis using 8 channels) to reduce the bandwidth required for data acquisition (a total of 48 channels in our case) and avoid the need for computationally expensive FFT-type algorithms. Post-pulse analysis can obviously still be performed using the much larger set of raw signals (between 32 and 80 additional channels depending on the operational setup) having the required sampling rate (at least 1MHz) via software-based synchronous detection. Further technical details on the active MHD diagnostic system used at JET can be found in [31-35] and the references provided therein.

The AELM is a digital VME plant control system, used to control the AE excitation in real time, with a 1kHz clock-rate. The AELM crate contains:

1. a VME Crate Service Module: this is a “home grown” module and acts as the VME system controller and provides timer and trigger synchronisation with the plant systems; it also

monitors the VME crate voltages, temperatures and airflow for operational control (plant failure);

2. a Real-Time Processor (RtProc): this is an Emerson Network Power MVME5500 card with a 1GHz PowerPC and 512MB RAM; this card executes software running under the Wind River VxWorks operating system, 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 an Emerson Network Power MVME5100 card with a 400MHz PowerPC and 64MB RAM. It is linked to both the JET real-time ATM and Ethernet networks and isolates the RtProc from asynchronous events that might disrupt its deterministic 1kHz calculation cycle time; the CommsProc also executes software running under VxWorks and is used to setup pre-pulse information, synchronise the RtProc with important time points within the pulse and finally communicates data recorded during the pulse for archiving;
4. four 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 can basically be thought of as a synchronisation module that has 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) *initialisation 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 1kHz clock interrupt from the Crate Service Module; the CommsProc also ensures that the AEAD plant ceases to operate if there is either a CPU failure or an external systems failure that terminates 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 is only activated during an experiment (i.e. a JET pulse), otherwise it remains dormant. In its active state the RtProc operates on a 1kHz cycle, and the sequence of this cycle is as follows:

a) read all the analogue signals from the I/O cards and convert them from complex numbers into polar coordinates (taking usually ~70microseconds);
b) perform "shorted turn" arc detection calculations on the 8 antennas, so that if a faster-than-normal change in the antenna impedance is identified then the AELM trips the amplifier to suspend output for 20ms for plant protection (taking usually ~50microseconds); the user is

able to configure the total number of trips tolerated before an antenna is permanently excluded from the remainder of the pulse;

c) perform mode detection, discrimination and tracking using the selected algorithm; two options are available (a user selection before the beginning of the pulse, which cannot be modified during the pulse), described below in more detail: the "*SimpleSum*" mode takes usually ~100ms, whereas the "*SparSpec*" mode takes usually ~650microseconds;

d) finally, the amplifier frequency is calculated and output back to the AEAD plant for the next time step, together with the amplifier amplitude, which is either a constant or a specified waveform (taking usually ~50microseconds).

A schematic overview diagram for the AELM is shown in fig4.

To detect and track in real-time the individual resonances corresponding to antenna-driven, stable plasma modes, the AELM linearly sweeps the antenna frequency around an initial guess for the AE resonance. The tracking detection algorithm can use one of three references for the scanning frequency: a user supplied waveform, a calculated waveform using real time toroidal field, current and plasma density, corresponding to the AE frequency at the centre of the selected AE frequency gap, or a waveform provided by an external signal. One of two algorithms can be selected to derive a single amplitude and phase pair that will be used for mode detection and tracking: the original (~10 years old) "*SimpleSum*" algorithm combines up to 8 real and imaginary pairs from the total 8 input data available, whereas the newly developed "*SparSpec*" algorithm can also accept up to 8 input real and imaginary pairs but produces many output amplitude and phase pairs. When using the *SparSpec* tracking mode, the AELM has two methods, "*highest*" and "*any*", for selecting the pair to use for tracking: the *highest* method picks the pair with the greatest amplitude, whereas the *any* method looks for a pair where the amplitude is above a given threshold; if a resonance is detected, this pair will continue to be selected until tracking is lost when the search for a new pair will commence. When the AE resonance is met, the exciter frequency is swept back and forth around it. A very simple Lorentzian model fit of the complex antenna/plasma transfer function is then used to obtain the mode frequency and quality factor in real-time, as well as the mode amplitude at the different probe locations. As this simple estimate of the main frequency and damping rate from the centre and width of the driven resonance closely follows the value obtained from a full fit, these quantities can be evaluated in real time.

Figure5 is an illustrative example of the AEAD plant in the tracking mode of operation. In the full-frequency spectrum (shown in the left frame, which was only measured during 4sec, i.e. from *t=13sec* to *t=17sec*) we see a very narrow triangular waveform in an otherwise completely clean portion of the fluctuation spectrum: this is the antenna frequency, which was set to look for resonances around the frequency of an *n=1* Toroidal Alfvén Eigenmode (TAE) as evaluated in real-time for that shot (i.e. around 200kHz). In the right frame, we have the synchronous detected signal ($|\delta B_{MEAS}|$) from one magnetic pick-up coil, showing in red the real-time TAE frequency (f_{TAE}) and in blue the antenna frequency. This signal was measured during 12sec, i.e. from *t=11sec* to *t=23sec*. Narrow

sweeps of the antenna frequency occur when the complex-valued δB_{MEAS} (shown in the two inserts) is sufficiently close to the pre-set resonant shape, i.e. it describes a circle in the complex plan representation with a corresponding bell-shape in the $|\delta B_{MEAS}(\omega)|$ representation. In these instances we are in the *tracking* mode of operation, and the frequency (f_{MEAS}) and damping (γ_{DAMP}) of the detected mode (in this case an $n=1$ TAE) can then also be measured in real-time. Such narrow sweeps occur in fig.5, for instance, between $t=12sec$ and $t=18sec$. After $t=21sec$, a much larger frequency sweep is seen in the right frame, indicating that no resonances close to the pre-set value have been detected in real-time: we are then in the *scanning* mode of operation, and the AELM looks for suitable antenna-driven plasma resonances in a different frequency range.

4. REAL-TIME, BLIND AND UNSUPERVISED DETECTION AND DISCRIMINATION OF DIFFERENT TOROIDAL COMPONENTS IN A FREQUENCY-DEGENERATED SPECTRUM OF MHD INSTABILITIES

One of the first results obtained with the AEAD diagnostic system in the most recent JET experimental campaigns has been that, despite the very low magnetic field driven by the antennas for medium-n AEs (in the plasma core we have $|\delta B_{DRIVEN}| \sim 1\times10^{-3} G$ for $n=5$ compared to $|\delta B_{DRIVEN}| \sim 5\times10^{-2} G$ for the $n=1$ and $n=2$ modes), many modes with $|n| \sim 0\text{-}12$ and very low-damping rate $\gamma/\omega < 0.2\%$ were found to be simultaneously excited in plasmas without populations of resonant fast ions [30]. Correct real-time detection and n-number discrimination of these modes is particularly important as their low intrinsic (i.e. without any drive from resonant fast ions) damping rate makes them very prone to become unstable if resonant fast ions were present in the plasma. It is specifically this experimental observation that has prompted the development of the more sophisticated real-time algorithm for mode-number recognition based on the sparse representation of signals, which has now almost completely replaced the ~10 years old original *SimpleSum* detection method.

The selection between the *SimpleSum* and *SparSpec* detection method is performed by the user via the AELM configuration panel, whose main operating window is shown in fig6. This panel contains many tabs requiring user input for plant configuration and control; those related to the selection of the mode detection algorithm are briefly described below (with a subset further graphically illustrated in fig7a):

DetectorSignal	selects the algorithm for real-time mode detection, either *SimpleSum*, or *SparSpec*
Mode	if the *SparSpec* algorithm has been selected, defines the option used for selecting which n-number mode has to be tracked in real-time (*highest* or *any*)
AmpCalc	if the *SparSpec* algorithm has been selected, defines the option used for computing the absolute amplitude of the mode via a LS estimation (*yes* or *no*)
N-selection	if the *SparSpec* algorithm has been selected, defines the n-numbers that needs to be detected and discriminated, in the range $-15 \le n \le +15$; since the *SparSpec* algorithm defaults to the lowest selected n-number if no mode is found within the

	maximum allocated CPU time, the value $n = -15$ is kept so as to easily notice when no mode has been found in real-time, as the $n = -15$ mode is not usually present in the plasma
polarity	defines the signal polarity (+ or −) associated to the selected sensor
sensor	defines which sensor is associated to that particular channel (can be "unused")
T-entry	defines the angular position of the sensor associated to that particular channel
filter	cut-off frequency for a real-time low-pass filter applied to all the raw data
thresholds	define the amplitude ($d\|\delta B_{MEAS}\|/dt$) and speed ($d(\|\delta B_{MEAS}\|)/dI_{ANT}$) thresholds at the start/end of each frequency scan to recognise if a plasma resonance detected via the *SimpleSum* or *SparSpec* algorithm is indeed associated to a mode of sufficiently high amplitude ($d\|\delta B_{MEAS}\|/dt > AmpThresh$) and sufficiently close to marginal the stability limit $\gamma/\omega = 0$ ($d(\|\delta B_{MEAS}\|)/dI_{ANT} > SpeedThresh$) to be of interest for real-time detection and tracking (see also the illustrative sketch shown in fig.7a)
MinimumTwist	defines the minimum angle threshold $d(imag(\delta B_{MEAS}))/d(real(\delta B_{MEAS}))$ in the complex plane representation to recognise if a mode detected using the amplitude and speed thresholds defined above is antenna-driven and stable ($\gamma/\omega > 0$) (see fig.7(a)
SS-nmax	if the *SparSpec* algorithm has been selected, defines the maximum n-number to be included in the calculation of the sparse spectrum; this value must be at least twice the maximum $\|n\|$ selected for tracking
SS-treshold	if the *SparSpec* algorithm has been selected, defines the background noise level threshold in the $\|\delta B_{MEAS}\|$ spectrum
SS-lambda	if the *SparSpec* algorithm has been selected, defines the λ_{NORM}-parameter
CdampConstant	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 [36])

This approach allows the detection and tracking of hundreds of individual resonances during one single tokamak discharge, which are guaranteed to be the same mode, i.e. to have the same n-number, if the *SparSpec* algorithm has been used imposing one single n for detection.

Figure7(b) shows a real-time time example of the twist and speed calculations for the $n = 11$ TAE mode in the JET Pulse No: 77790. All the digital signals shown at the bottom of fig.7(b) indicate whether detection and tracking of a certain mode has been successful: if the flag is set to high (=1) then the corresponding mode has been correctly detected and tracking is occurring, otherwise the digital signal is set to low (=0). At the start of each real-time scan t, the global "mode" value is initialised to a "bad" tracking value of $n = -15$ and its associated digital validity flag is set to 0 (not shown in fig7(b), where we only shown the digital flags associated to certain specific n's to reduce cluttering). When a mode is successfully detected the "mode" graph is set to the n value

identified and the validity flag set to 1.

Therefore, if no mode is detected the graph will record a "mode" value of $n=-15$ but the validity flag will be 0. By comparing the status of the digital signals for the amplitude (*SigAmp.11*) and speed (*Speed.11*) for the $n=11$ mode, we note that it is only when these two signals are concurrently high that the corresponding digital flag for the differential twist (*TwistRate.11*) is also set to high, hence the integrated twist is being calculated and tracking of the $n=11$ mode can occur (mode signal set to =11) if the integrated twist exceeds the set threshold. Specifically, this mostly happens in the time window $5.0<time(sec)<9.5$ (and an highlight of this is shown in fig.7(c)), and only very spuriously otherwise, indicating a very subtle dependence of the stability of this $n=11$ mode as function of the background plasma parameters. It is also important to note that it is in principle possible to have a valid $n=-15$ mode as long as the validity flag is true (i.e., it is =1): therefore, it is the validity flag that denotes "good" or "bad" tracking for the $n=-15$ mode.

The capability to perform a blind and unsupervised real-time detection and tracking of the individual n-components in the antenna driven spectrum constitutes an invaluable tool, which is unique to the JET tokamak. It provides accurate testing for the code prediction for the damping rate of Alfvén Eigenmodes (see for instance [37]), as it is paramount that the same mode be measured throughout the parameter scan. The full implementation of the *SparSpec* algorithm in the AELM software now allows a detailed quantitative analysis of the recent measurements made with the AEAD system in JET. Mode numbers can be directly separated in real-time and individually tracked to measure the changes in the mode frequency and g/w during the evolution of the plasma background. It is also paramount that the real-time detection matches the results obtained via the more complex and accurate post-pulse analysis method, and this has been discussed and demonstrated in more detail in [34, 35]. Figure8 shows an example of the blind real-time detection and tracking of individual $n=3$, $n=5$ and $n=7$ TAEs for the JET Pulse No: 77417. We notice that the antenna frequency waveform F_{TAE} follows closely the real-time evolution of the reference TAE frequency F_{REF}; the flag "mode" is initialised to a tracking value of $n=10$ ("bad" since it appears in conjunction with a false validity flag =0), so that when it shows in real time a value different from such flag (in our case being "forced" to look for either $n=3$ or $n=5$ or $n=7$), it indicates that a "good" mode has been detected. When this occurs, the AELM locks onto this particular mode (i.e. it discards everything else), allowing the evolution of this mode to be tracked in real-time. The digital flags for the amplitude of the $n=3$ and $n=7$ modes (*SigAmp.3* and *SigAmp.7*, respectively) shown at the bottom of fig8, indicate whether detection and tracking of a certain mode has been successful: if the flag is set to high (=1) then the corresponding mode has been correctly detected and tracking is occurring.

A further example of the blind real-time detection and tracking of the individual n-number components in the antenna-driven spectrum is shown in fig.9(a,b) for the JET Pulse No: 77788, where the antenna were configured to predominantly drive an odd-n spectrum, peaked towards the lower mode numbers $|n|=3$ to $|n|=5$, so as to have a negligible drive for components with $|n|>10$. Here we show separately the detection of all modes with $3\leq n\leq$£8 (fig.9(a)), and the detection of

odd-*n* modes rotating in the co-current (positive *n*'s, the preferential direction for the usual tokamak magnetic equilibrium) and counter-current (negative *n*'s) direction, with $|n|=[3, 5, 7]$ (fig.9(b)).

From fig9a, we note that only very few even-*n* resonances ($n=[4, 6, 8]$) are detected in real-time (and confirmed by post-pulse analysis), compared to the number of $n=3$ resonance; second, the $n=3$ mode dominates the detected spectrum, as this is the one for which the antennas produce the maximum drive. Finally, fig.9(b) demonstrate that not only co-and counter-rotating modes with the same toroidal mode number can be perfectly distinguished in real-time, but also that the plasma preferentially supports co-rotating modes (which are driven by the antennas with the same amplitude as the corresponding counter-rotating modes). This result allows us to discriminate whether the pressure profile of any resonant fast ion population is peaked on axis (reducing the damping rate of co-rotating modes) or off-axis (reducing the damping rate of counter-rotating modes), with important consequences for the control and optimization of the plasma discharge and burn scenario.

SUMMARY AND CONCLUSIONS

In this work we have reported on the application, to tokamak plasmas, of a new algorithm for the real-time, blind and unsupervised detection and decomposition of a degenerate frequency spectrum, where the frequency separation between the various components is less than their full-width at half-maximum. This algorithm is based on the sparse representation of signals, as derived from its original applications to astronomical data via the *SparSpec* code [17]. The development of this method was prompted by the first measurements in the JET tokamak of such a frequency degenerated spectrum of antenna-driven AEs, with toroidal mode numbers $n\sim3\text{-}15$ [34, 35]. The real-time (and post-pulse) implementation of the *SparSpec* algorithm on JET has allowed a complete, accurate and numerically efficient analysis of these measurements, which would have not been possible otherwise. For post-pulse analysis, the antenna-driven spectrum can be decomposed and analysed for one whole frequency scan typically within (2-3)ms of CPU time using the full (more complex and more accurate) implementation of the *SparSpec* algorithm. With an SVD decomposition algorithm similar to the one presented in [9], requiring a combinatorial exploration of all possible solutions and an a-posteriori thresholding scheme to determine the correct ones, or using wavelet-based schemes, such as those presented in [10], this analysis would require a CPU time usually in excess of 150ms for each individual frequency scan. Such computational requirements make it impossible to use these schemes for real-time analysis. Conversely, using the rather modest computational resources allocated to the real-time analysis of the AEAD data within the AELM hardware and software (a 1GHz PowerPC with a 512MB RAM running on a 1kHz clock-rate), the multi-components antenna-driven spectrum can be fully resolved within typically ~650microseconds for each 1ms clock-cycle.

For JET, and more generally fusion plasmas, further application of this new method based on the sparse representation of signals opens interesting and very useful perspectives for the concurrent real-time detection and control of different MHD instabilities, as they can be discriminated very

accurately. This allows specifically tailored control schemes to be put in place for each individual instability, hence improving the overall control of plasma operation. This will be particularly important for future experiments approaching the burning plasma conditions, such as ITER, 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 energy gain.

Furthermore, and while specifically applied for the analysis of astronomical data and mode numbers in thermonuclear fusion plasmas in a tokamak device, the use of sparse approximations methods are ideally suited for applications to all domains of data analysis and control engineering where an efficient decomposition of a multi-harmonics degenerate spectrum is required from irregularly sampled data. Moreover, the computational speed and accuracy of algorithms such as SparSpec makes such a method ideally suited for real-time applications with a small number of data. These domains range from the analysis/optimization of measurement devices (an example of this being the work done for the ITER high-frequency magnetic diagnostic system [19]), to even more sophisticated applications such as the processing of digital images for pattern recognition [1, 38] and the development of the numerous theoretical developments on the Sparse Representations of signals and image processing for practical applications, possibly in real-time (see for instance some examples: "http://nuit-blanche.blogspot.com/" http://nuit-blanche.blogspot.com/)

ACKNOWLEDGEMENTS

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. This work was also partly supported by the Swiss National Science Foundation. The views and opinions expressed herein do not necessarily reflect those of the European Commission. The Authors would like to thank the various members of the CRPP, MIT and JET staff that have contributed to the design, installation, commissioning and routine operation of the new Alfvén Eigenmode Active Diagnostic system, P.Beaumont, C.Bower, S.Dowson and S.G.Sanders (CCFE) for their assistance during the installation and commissioning phase of the new AEAD system, and M.Tsalas (JET-EFDA-CSU) for his preparation work for many experimental sessions during the latest JET experimental campaigns.

REFERENCES

[1]. F. Marvasti, *Non-Uniform Sampling: Theory and Practice*, Springer-Verlag (2001).

[2]. N.R. Lomb, Astrophysics and Space Science **39** (1976), 447.

[3]. J.D. Scargle, The Astroph. Journal **263** (1982), 835.

[4]. D. Roberts, J. Lehar and J. Dreher, The Astron. Journal **93** (1987), 968.

[5]. G. Foster, The Astron. Journal **109** (1995), 1889.

[6]. J. Wesson, *Tokamaks* (3rd edition), Oxford Science Publication, Oxford, 2003.

[7]. P. Zanca et al., Phys. Plasmas **8** (2001), 516.
[8]. J.S. Kim et al., Plasma Phys. Control. Fusion **41** (1999), 1399.
[9]. M. Hole and L.Appel, Plasma Phys. Control. Fusion **49** (2007), 1971.
[10]. F. M.Poli et al., Plasma Phys. Control. Fusion **50** (2008), 095009.
[11]. E.P. Wigner, Phys. Rev. **40** (1932), 749.
[12]. H. Choi, W.Williams, IEEE Trans. Acoust. Speech Signal Process **37** (1989), 86.
[13]. N.E. Huang et al., Proc. Royal Society London **A454** (1998), 903.
[14]. S. Zegenhagen, A. Werner et al., Plasma Phys. Control. Fusion **48** (2006), 1333.
[15]. F. Zonca et al., Plasma Phys. Control. Fusion **48** (2006), B15.
[16]. Various Authors, "Progress in the ITER Physics Basis", Special Issue of Nucl. Fusion **47** (2007).
[17]. S. Bourguignon, H.Carfantan, T.Böhm, Astronomy and Astrophysics **462** (2007), 379.
[18]. A. Klein, H. Carfantan, D. Testa, et al., Plasma Phys. Control. Fusion **50** (2008), 125005.
[19]. D. Testa et al., "*Functional performance analysis and optimization for the high-frequency magnetic diagnostic system in ITER – part1 and part2"* , Fusion Science and Technology, **57**(3) (2010), 208-237 and **57**(3) (2010), 238-273.
[20]. S. Ferraz-Mello, The Astron. Journal 86 (1981), 619.
[21]. J. Lafler and T.D.Kinman, The Astrophys Journal Suppl. **11** (1965), 216.
[22]. M.M. Dworetsky, Mon. Not. R. Astron. Soc. **203** (1983), 917.
[23]. P. Stoica, R.L. Moses, B. Friedlander, and T. Söderström, IEEE Trans. ASSP **37** (1989), 378.
[24]. J.-J. Fuchs, IEEE Trans. Inf. Theory **51** (2005), 3601.
[25]. D.L. Donoho, M. Elad and V.N. Temlyakov, IEEE Trans. Inf. Theory **52** (2006), 6.
[26]. J.A. Tropp, IEEE Trans. Inf. Theory **52** (2006), 1030.
[27]. J.A. Tropp, IEEE Trans. Inf. Theory **50** (2004), 2231-2242.
[28]. D. Donoho, Journal of Communications on Pure & Applied Math. **59** (2006), 797.
[29]. S. Bourguignon, H.Carfantan, J.Idier, IEEE Journal of Selected Topics in Sig. Proc. **1** (2007), 4.
[30]. D.M.Malioutov, *A sparse signal reconstruction perspective for source localization with sensor arrays*, Master Thesis, MIT, July 2003.
[31]. A. Fasoli, D.Testa et al., Plasma Phys. Control. Fusion **44** (2002), 159.
[32]. A. Fasoli et al., Phys. Rev. Lett. **75** (1995), 645.
[33]. D. Testa et al., *The new Alfvén Wave Active Excitation System at JET*, Proceedings 23rd Symposium on Fusion Technology, Venice (Italy), 20-24 September 2004 (weblink: HYPERLINK "http://infoscience.epfl.ch/record/143354/files/" http://infoscience.epfl.ch/record/143354/files/).
[34]. D. Testa et al., *Measurement of the Damping Rate of High-n Toroidal Alfvén Eigenmodes in JET*, Proceedings 11th IAEA Technical Committee Meeting on Energetic Particles, Kyiv (Ukraine), 21-23 September 2009, accepted for publication in Nuclear Fusion, 2010 (weblink: "http://infoscience.epfl.ch/record/143355/files/" http://infoscience.epfl.ch/record/143355/files/).
[35]. T. Panis, D.Testa et al., "*Optimization of the Active MHD Spectroscopy System on JET for the*

Excitation of Individual Intermediate and High-n Alfvén Eigenmodes", Proceedings 11th IAEA Technical Committee Meeting on Energetic Particles, Kyiv (Ukraine), 21-23 September 2009, accepted for publication in Nuclear Fusion, 2010 (weblink: "http://infoscience.epfl.ch/record/142715/files/" http://infoscience.epfl.ch/record/142715/files/).

[36]. D. Testa and A. Fasoli, Nucl. Fusion **41** (2001), 809.

[37]. The ITPA working group, IAEA 2010.

[38]. For different examples of digital image processing methods see: W.Burger and M.J.Burge (2007), *Digital Image Processing: an Algorithmic Approach Using Java*, Springer-Verlag 2008 (ISBN: 978-1-84628-379-6); see also: R.Fisher, KDawson-Howe, A.Fitzgibbon, C.Robertson, E.Trucco, *Dictionary of Computer Vision and Image Processing*. John Wiley 2005, (ISBN: 0-470-01526-8).

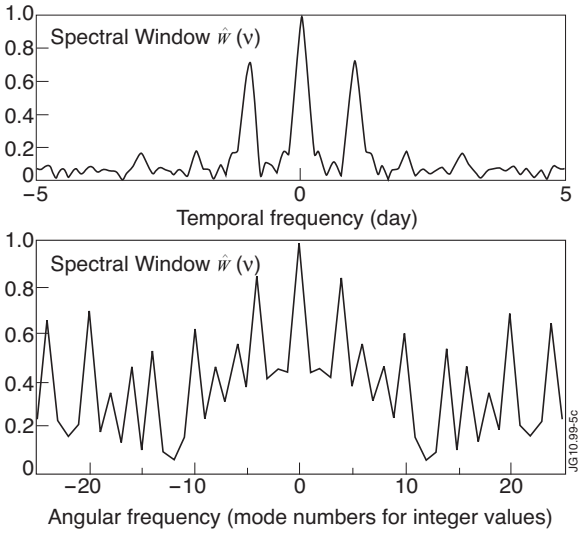

Figure 1: The spectral windows corresponding to an example of astronomical observations during five nights (top frame) and to the actual positions of the only seven JET magnetic sensors that are sufficiently reliable to be used in the real-time data processing algorithm for the analysis of MHD instabilities (bottom frame).

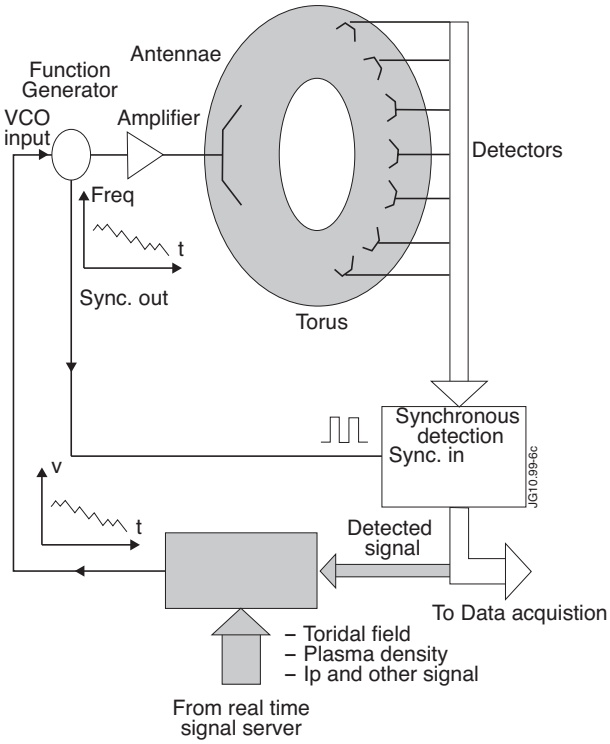

Figure 2: A schematic overview of the active AE diagnostic system in JET. The toroidal field, plasma density and plasma current (I_p) are required to compute in real-time an initial guess for the frequency of Alfvén Eigenmodes. This value is then converted as a time-dependent voltage V(t) and sent as 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 high-power amplifier connected to in-vessel antennas and the synchronous detection units via an optical transducer, so that only the portion of the plasma response which has the same frequency of the antenna drive (i.e. it is synchronous with them) is detected in real-time, which avoids the need for computationally expensive FFT algorithms.

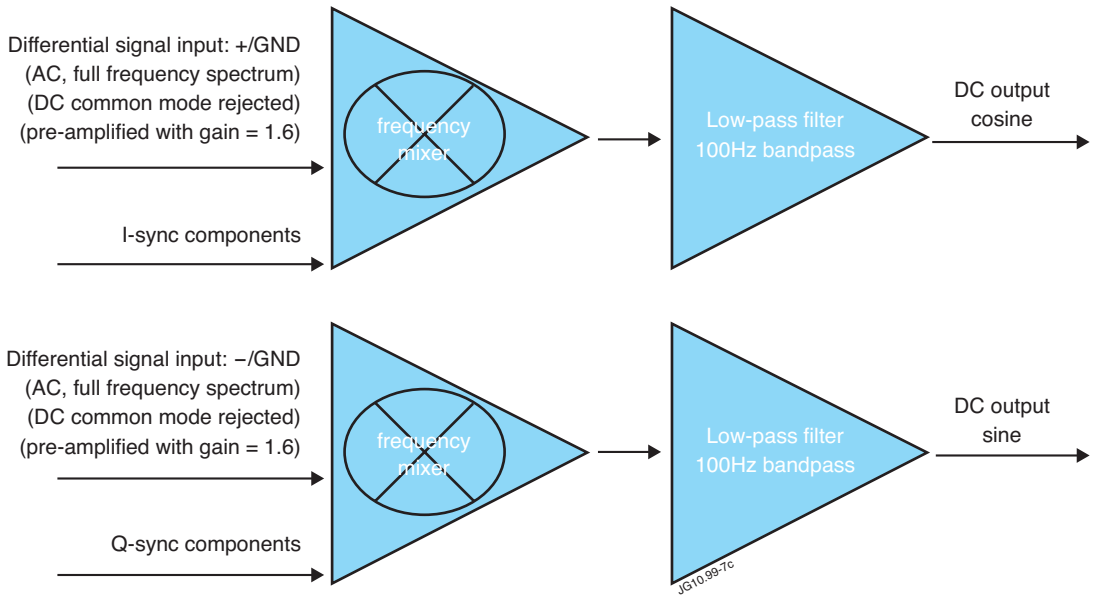

Figure 3: A schematic illustrative diagram of the AEAD synchronous detection hardware. This hardware is conceptually based upon first removing any DC common mode from the input AC differential signal (with a full frequency spectrum), which is then amplified, and then applying a low-pass filter to extract only the component in the input signal which is at the desired (i.e. synchronous) frequency.

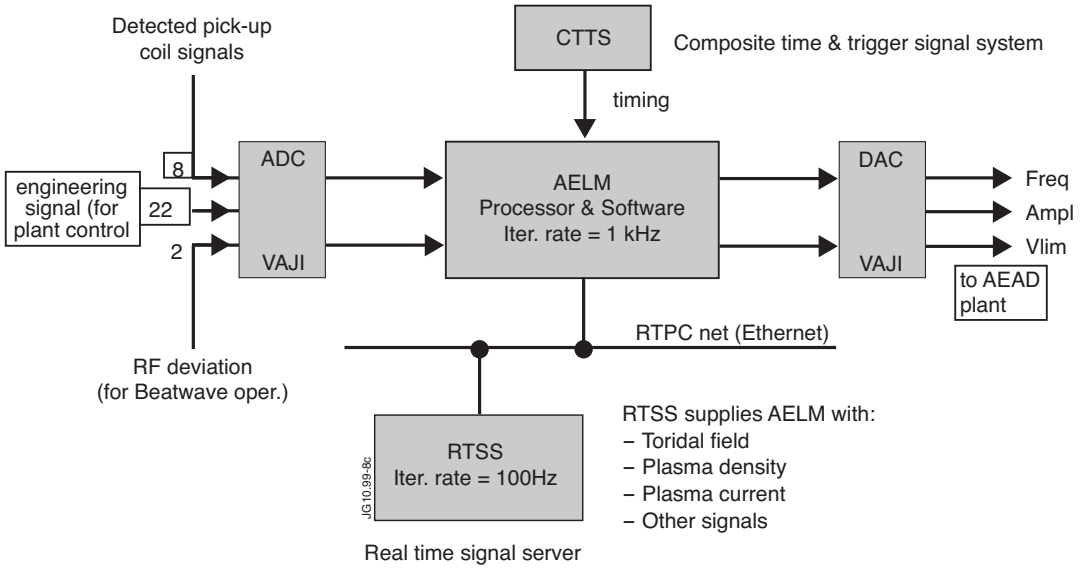

Figure 4: A schematic illustrative diagram of the AELM hardware.

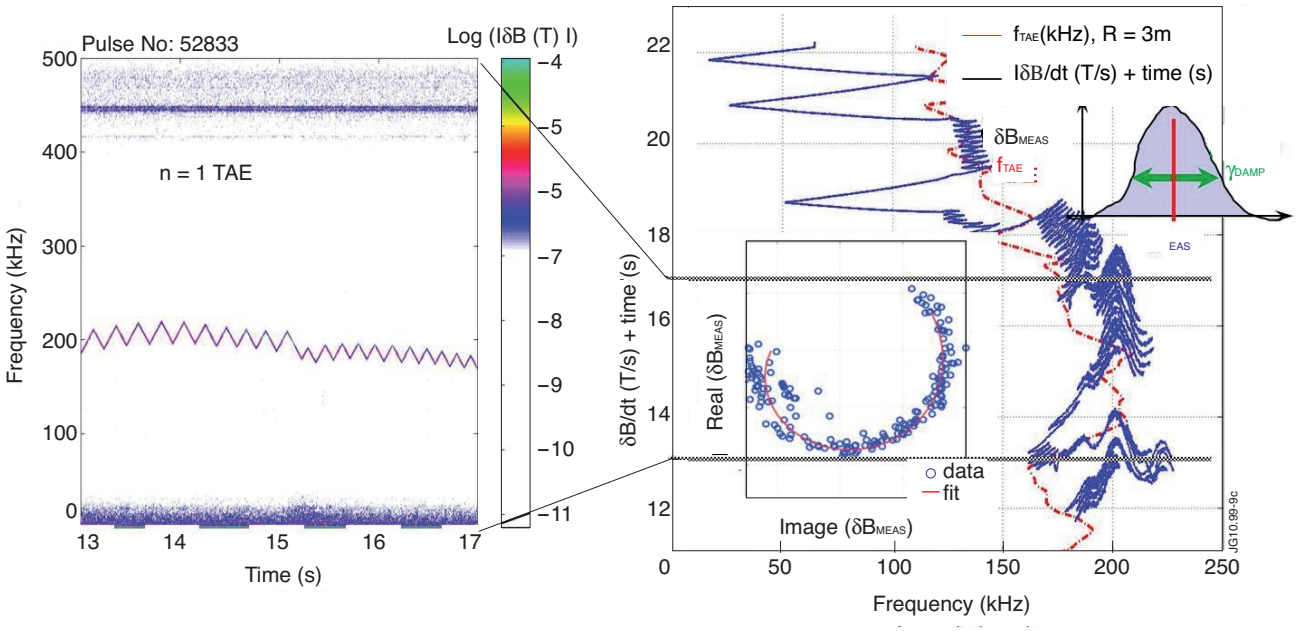

Figure 5: An example of real-time tracking of a resonant Toroidal Alfvén Eigenmode.

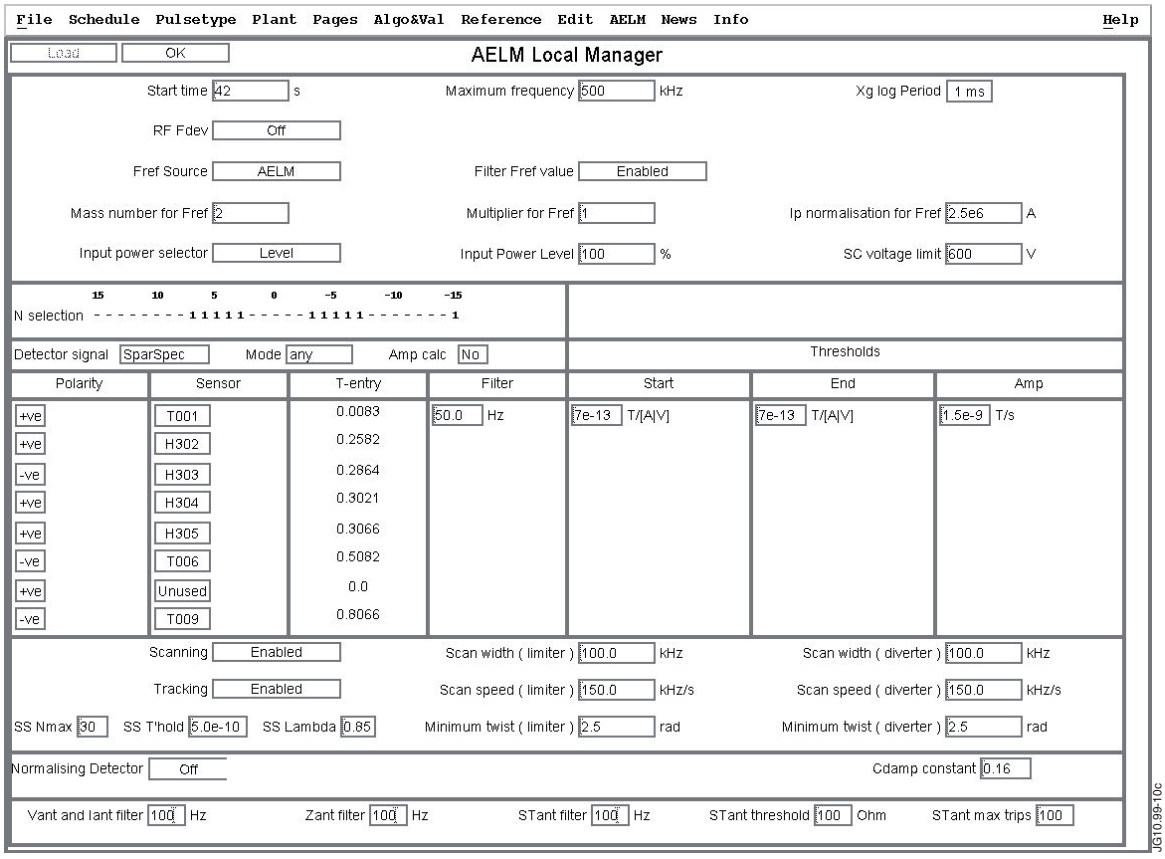

Figure 6: The main AELM configuration panel for selecting the algorithm used for real-time detection, discrimination and tracking of the antenna-driven plasma resonances. All the tabs must be configured by the user and loaded to the plant before the start of the experiment.

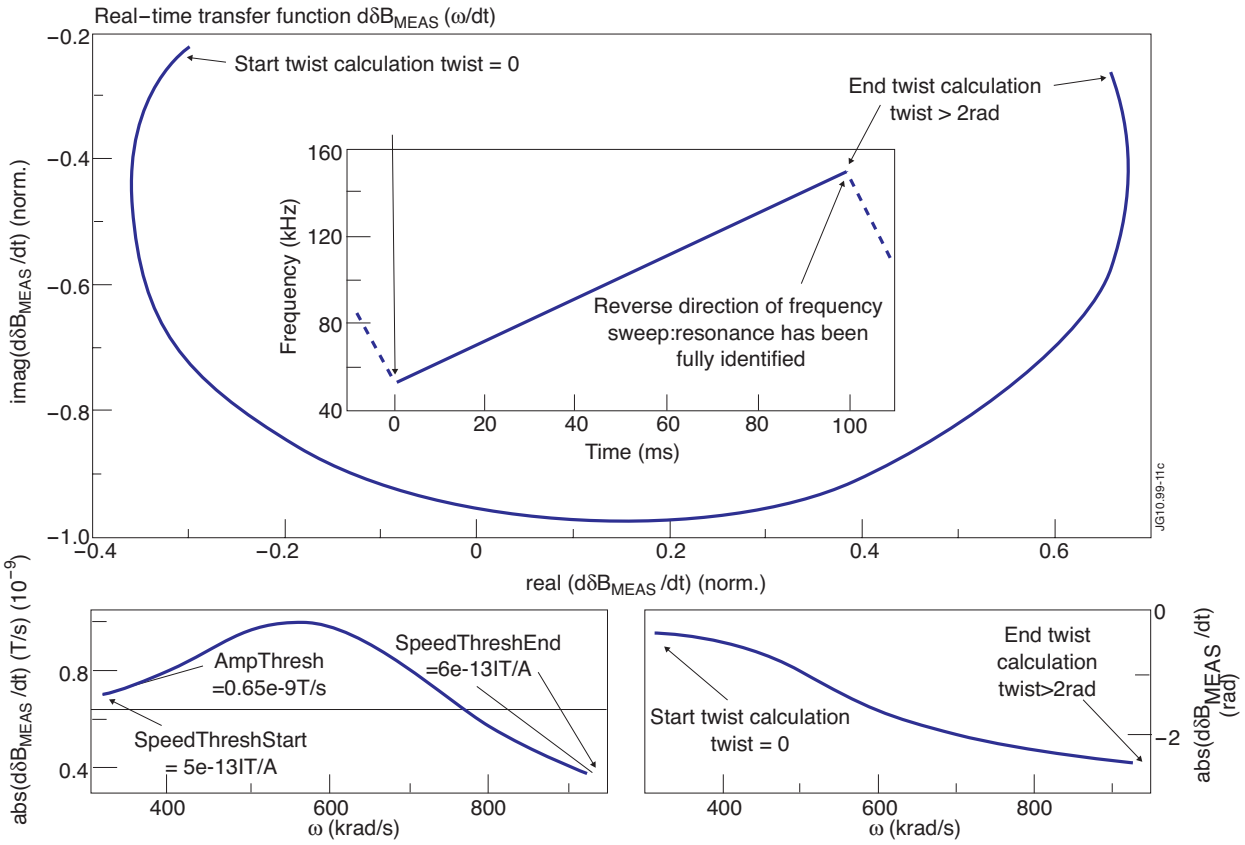

Figure 7(a): A schematic diagram to illustrate the basic ideas of the real-time tracking algorithm. At the beginning of a frequency scan, if a mode has been detected (either via the SimpleSum or SparSpec methods), we check that the amplitude and speed are above the set thresholds (see the bottom left plot: in this case AmpThresh=0.65×10^{-9}[T/s] and SpeedThreshStart = 5×10^{-13}[T/A]); if these conditions are satisfied, we start computing the integrated twist (see the bottom right plot) using the complex plane representation of $\delta B_{MEAS}(w)$ (see the top plot); when the integrated twist exceeds the set threshold =2π and if the speed is below SpeedThreshEnd (in this case SpeedThreshEnd = 6×10^{-13}[T/A]), we estimate that we have fully identified the antenna-driven plasma resonance, hence the direction of the frequency sweep is reversed (as shown in the insert in the top plot) in an attempt to follow the evolution of the same mode as the background plasma conditions change.

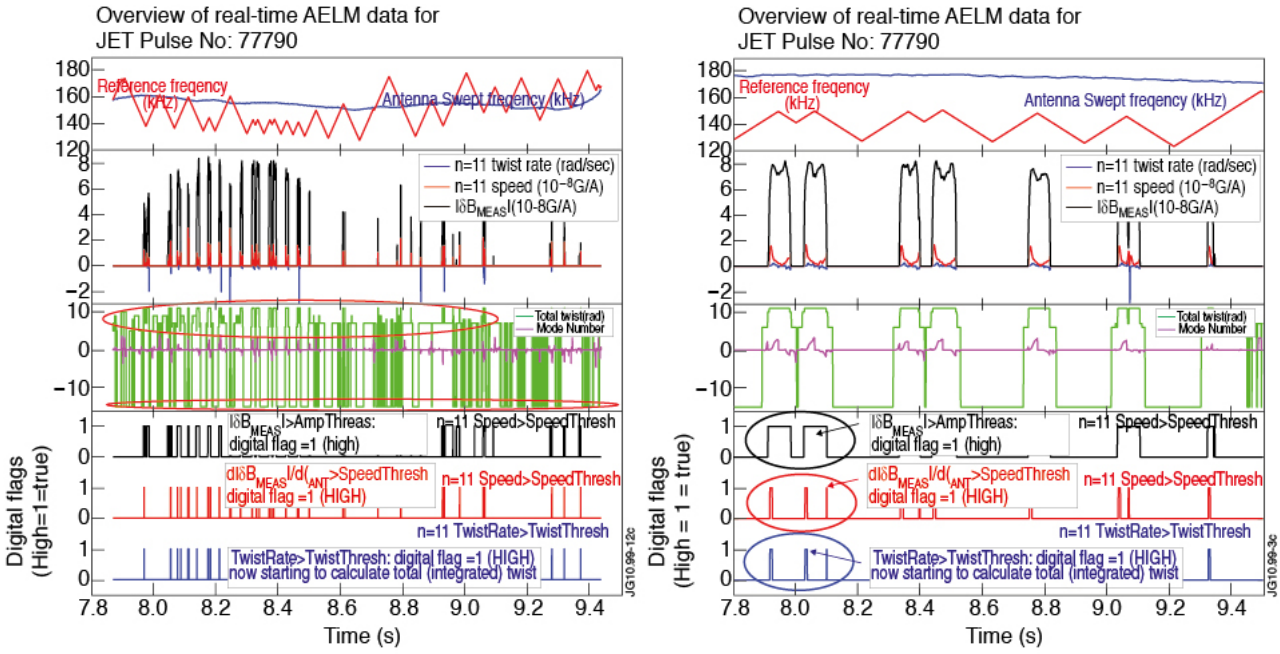

Figure 7(b): (left). A real-time example of the twist and speed calculations for the JET Pulse No: 77790 and the n=11 TAE mode; in green the antenna frequency F_{TAE}, in brown the reference frequency F_{REF}, corresponding in this case to the centre of the n=1 TAE gap, and in yellow the twist signal (only shown on the digital trace); in blue the flag "mode", which indicates the n-number that has been detected. Figure 7(c): (right). A zoom of fig7b to show in more details that the tracking of the n=11 mode only occurs when all the resonance conditions on the mode amplitude, speed and twist are concurrently above the set thresholds.

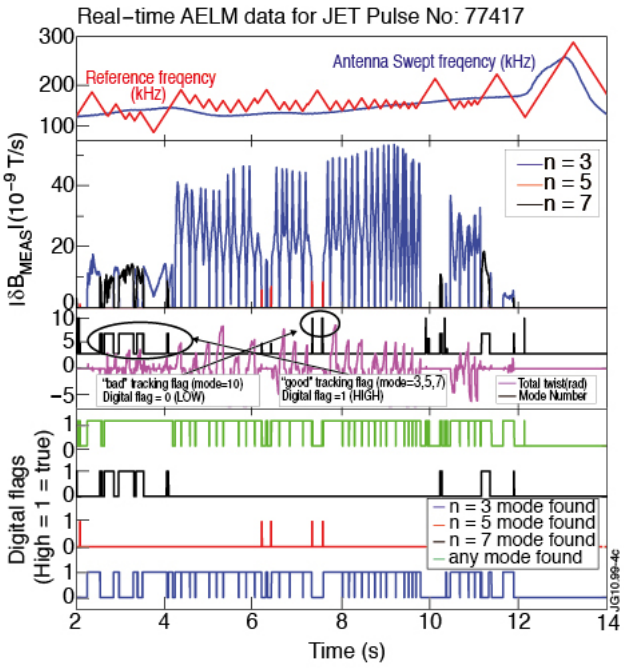

Figure 8: An example of real-time tracking of the individual n=3, n=5 and n=7 TAEs; the digital signal SigAmp.3 (.7) shows the amplitude of the n=3 (n=7) modes as detected in real-time (note that these signals are available also in analogue format for all modes selected via the AELM configuration panel); similar plotting conventions as in fig.7 have been used here for F_{TAE}, F_{REF} and the "mode" flag; the n=10 mode corresponds to "bad" tracking because its digital validity flag is false (=0).

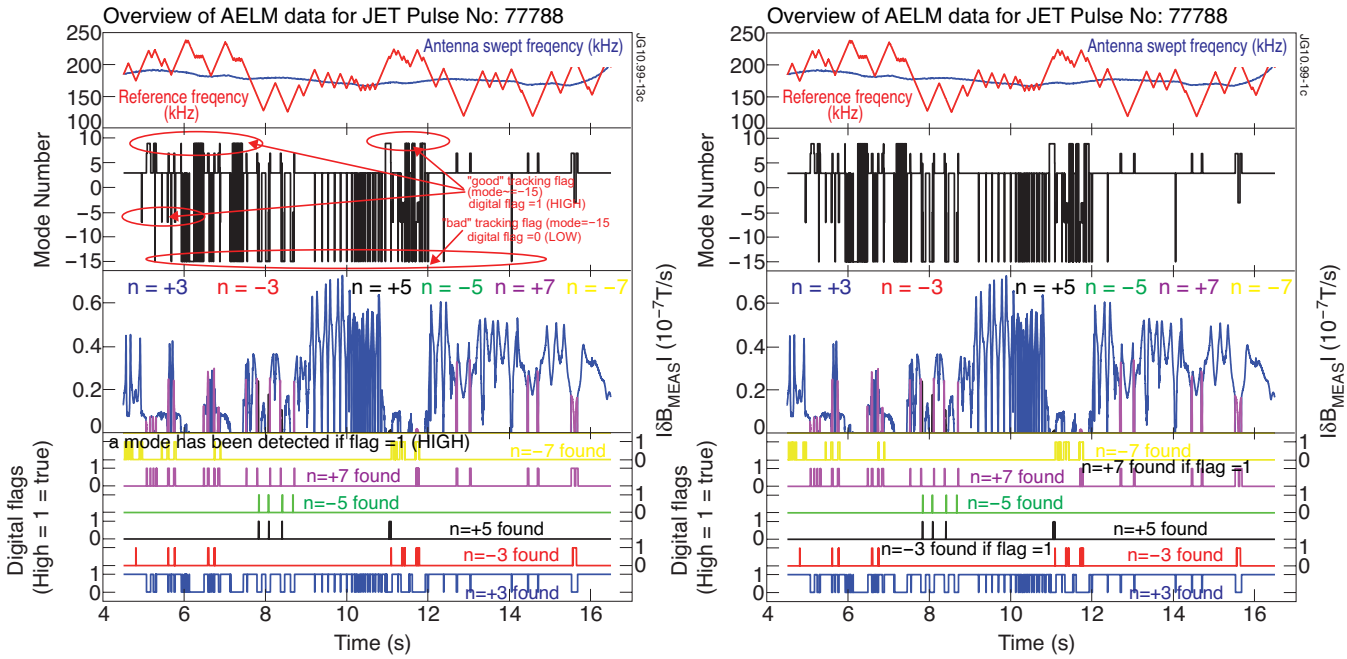

Figure 9(a): (left). An example of real-time tracking of the individual n=3-8 TAEs for the JET Pulse No: 77788, where the antennas were configured to drive predominantly an odd-n spectrum with max($|\delta B_{DRIVEN}|$) for n~3-4: indeed the n=3 mode is the one which is most detected; similar plotting conventions as in fig.7 have been used here for F_{TAE}, F_{REF} and the "mode" flag.