



HAL
open science

A sparsity-based method for the analysis of magnetic fluctuations in unevenly-spaced Mirnov coils

A. Klein, H. Carfantan, D. Testa, A. Fasoli, J. Snipes, Jet Efta Contributors

► **To cite this version:**

A. Klein, H. Carfantan, D. Testa, A. Fasoli, J. Snipes, et al.. A sparsity-based method for the analysis of magnetic fluctuations in unevenly-spaced Mirnov coils. *Plasma Physics and Controlled Fusion*, 2008, 50, pp.5005. 10.1088/0741-3335/50/12/125005 . hal-00392457

HAL Id: hal-00392457

<https://hal.science/hal-00392457v1>

Submitted on 18 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A sparsity-based method for the analysis of magnetic fluctuations in unevenly-spaced Mirnov coils

A Klein¹, H Carfantan², D Testa³, A Fasoli³, J Snipes¹ and JET EFDA Contributors^{1,5}

¹ MIT Plasma Science and Fusion Center, Cambridge, MA 02139, USA

² Laboratoire Astrophysique de Toulouse-Tarbes, Université de Toulouse, CNRS, France

³ CRPP, Association EURATOM – Confédération Suisse, EPFL, Lausanne, Switzerland

⁴ JET-EFDA, Culham Science Centre, OX14 3DB, Abingdon, UK

Received 19 June 2008, in final form 25 September 2008

Published 31 October 2008

Online at stacks.iop.org/PPCF/50/125005

Abstract

A new method for identifying toroidal mode numbers in Mirnov data from toroidal plasmas has been found, and benchmarked with simulated and real data from the JET tokamak. Embodied in the SparSpec code, and originally developed for the analysis of unevenly time-sampled astronomical data, this new method fits signals which are unevenly sampled in the toroidal coordinate to a sum of an arbitrarily large number of toroidal modes with integer mode numbers. The method has proven to be extremely robust, and is especially useful for resolving the amplitudes and phases of multiple Alfvén eigenmodes (AEs) 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 suggests that it may be used in real-time applications such as resistive wall mode or tearing mode control, among others. Examples involving unstable AEs as well as stable AEs, (excited by an array of external antennas) in the JET tokamak are used to illustrate the efficacy of the method. Some additional general considerations regarding optimized spatial sampling strategies are also briefly addressed.

(Some figures in this article are in colour only in the electronic version)

1. Background

Analysis of magnetic fluctuations external to toroidal plasmas is important for understanding the magneto-hydrodynamic (MHD) properties of the plasma, among other things. These properties affect nearly all aspects of behavior of magnetic confinement, and thus are of interest in topics ranging from gross global plasma stability, control and disruption avoidance to the

⁵ See annex 1 of M L Watkins 2006 *21st IAEA Fusion Energy Conf. (Chengdu, China)*.

more subtle areas such as are involved with passive and active MHD spectroscopy, to name a few examples. Mode number analysis is generally accomplished by interpreting signals from a finite number of external Mirnov coils, which typically are unevenly spaced in the toroidal and poloidal coordinates. The toroidal mode number, n , is usually easily determined in tokamaks because different modes generally oscillate at well separated and distinct frequencies, because of toroidal symmetry, and because n is often very low ($n \leq 2$). On occasion, however, multiple modes of various origins may overlap in frequency, and then the signals in the individual sensors are the result of a superposition of these modes. In that case the task of spatial decomposition becomes a difficult problem.

The variety of methods in common use to determine the spatial structure of MHD phenomena in toroidal plasmas from external magnetic measurements is quite small. In the standard tokamak coordinate system, magnetic perturbations at the plasma edge are represented well by functions involving toroidal and poloidal harmonics, written as $\psi_{m,n} = e^{i(n\phi)} \sum_m e^{i(m\theta)}$, where each mode has one single toroidal mode number, but will include several poloidal Fourier harmonics due to toroidicity and other geometric effects. As a result, the poloidal quantum number is not well defined, even in very large aspect ratio plasmas with cylindrical cross-section. In the poloidal dimension, uniform spacing between sensors does not result in any gain in terms of ease of analysis, because the distance between sensor and plasma surface is invariably a function of poloidal angle θ , and because a Shafranov shift and plasma shaping effects will result in effective spacing periods $\Delta\theta$ that are anything but constant. Additionally, because poloidal mode structure can never be represented by a single Fourier term, a technique such as singular value decomposition (SVD) is appropriate, since it makes no assumptions about the nature of the spatial basis functions [1].

For the toroidal aspect of the analysis, however, the eigenfunctions describing perturbations are best represented by simple sinusoids. When sensors are evenly spaced in the toroidal dimension, a simple discrete Fourier transform will reveal the mode amplitudes for each n separately [2], up to the Nyquist mode number given by $n_{\max} = \pi/\Delta\phi$ (here $\Delta\phi$ is the angle between adjacent probes). Unfortunately, evenly spaced sensors are often not available due to engineering constraints. Furthermore, it can be shown that even spacing does not represent an optimum arrangement for magnetic pickup coils, because of the aforementioned Nyquist limitations and aliasing effects [3].

2. Classical methods and SparSpec

A standard approach for determining the n -number of an MHD mode is to fit the phases of data from two or more sensors separated by their toroidal angle to a straight line. The slope of the line then represents the n -number. When there is only one dominant mode with one single n -number, phase fitting is relatively successful and simple (i.e. fast), but when there are multiple modes with multiple n -numbers some other means to resolve the mode numbers must be used.

Other tools employed to estimate sinusoidal contributions in sets of unevenly sampled data have traditionally involved the so-called Lomb periodogram [4], and this method was recently applied to Alfvén eigenmode (AE) identification in the Wendelstein 7-AS stellarator [5], in the single mode case. The maximum of the periodogram corresponds to a least squares (LS) fitting of the data to a single sinusoid, so it is statistically coherent and practically works well if there is only one dominant mode. In the case where multiple modes are present in the data, selecting the most important maxima of the periodogram is no longer statistically coherent in terms of LS fitting. Indeed, the periodogram is perturbed by the spectral window (i.e. the Fourier transform of the sensor positions) and the Lomb periodogram may have many maxima

which do not correspond to actual modes. This is illustrated in the example of figure 1, where it can be seen that, even when the number of actual modes (M) is known ($M = 4$ in this example), selecting the modes corresponding to the M higher values of the periodogram leads to erroneous mode detection. Such aliases have already been noticed in the case of a single mode in [5], but they generally do not lead to wrong detection in this case. Moreover, it is quite difficult to select a value for M from such a representation if it is not known *a priori*.

An optimal LS fitting of multiple sinusoids requires a combinatorial exploration of all possible modes. A recent publication by Hole and Appel [6] describes such an approach which uses SVD methods to fit magnetic data to a predetermined number (M) of Fourier modes, but it was only tested in [6] in the case of a single mode (note that this method is equivalent to the periodogram approach in this case). The principle of such a method is simple: for every combination of M modes it estimates the mode amplitudes by LS fitting (and the use of an SVD technique is a way to compute these amplitudes). The method then selects the combination of modes and corresponding amplitudes which produced the smallest residual (as defined in [6]). The drawbacks of this method are twofold: first, for a chosen M , the number of detected modes will always be M , even if there is only one mode in the data. If the actual number of modes is lower than M , the additional modes will hopefully have a low amplitude, and a posterior thresholding step is necessary to detect the correct number of modes present in the data. However, if the chosen M is smaller than the number of modes in the data, the method will fail. Second, for $M > 1$, it rapidly becomes computationally expensive to run through all possible combinations of M modes, the number of which is given by the binomial coefficient, $(2n_{\max} + 1)!/M!(2n_{\max} + 1 - M)!$ (here n_{\max} is the maximum allowed mode number, and we must consider negative mode numbers as well). For example, for the maximum n -number limited to $n \leq 20$, there are, respectively, 820, 10 660, 101 270 and 749 398 combinations, for $M = 2, 3, 4$ and 5. For $n_{\max} = 30$ and $M = 5$, there are 5949 147 combinations! Although the computation of the Moore–Penrose pseudo-inverse matrices can be done just once and these results stored, the LS solution must be computed for each combination.

The problem of finding periodic waveforms in unevenly sampled data is ubiquitous in the field of astronomy, where much work has been done. It is easily seen that temporal frequencies in astronomical data can correspond to spatial toroidal mode numbers in tokamaks, and that unevenly sampled data in time is the analog of data from unevenly distributed Mirnov sensors in the toroidal coordinate. However, in astronomy, the frequencies sought are obviously allowed to take on any value, while periodic boundaries in toroids ensure that only modes with integer n can exist. Since almost all astronomical data is unevenly sampled (due to weather conditions and the earth rotation), considerable effort has gone into the problem of improving upon the limitations of the Lomb periodogram. For example, the CLEAN [7] and CLEANEST [8] algorithms attempt to remove some of the aliases arising from uneven sampling, to mention only two.

Recently, a new method for fitting sinusoids to irregularly sampled data was proposed [9], based on the principle of sparse representation of signals. It is implemented in the SparSpec computer code (freely available at: <http://www.ast.obs-mip.fr/Softwares>). Applied to the problem of toroidal mode number analysis, the underlying idea is to model the data as a large number (possibly larger than the data size) of pure modes $n = -K \dots K$ (with $K = n_{\max}$). Among the many representations fitting the data, we seek the one with the lowest non-zero amplitude, i.e. a sparse solution. Strictly speaking, the sparsest solution minimizes the LS criterion penalized with the L^0 norm, i.e. the number of non-zero components of the solution. However, minimizing such criterion requires a combinatorial exploration (as in the case of the SVD method but for all M or less than M combination of modes). Such exploration is avoided by considering an L^1 norm instead of a L^0 norm penalization. Such a norm consists of the

sum of the absolute values of the mode amplitudes in the solution. Many theoretical works have been done to determine conditions of equivalence of both settings, one example is found in [10]. Thus, the solution is computed as the minimizer of the penalized LS criterion:

$$J(x) = \frac{1}{2} \|y - Wx\|^2 + \frac{\lambda}{\lambda_{\text{Max}}} \sum_{k=-K}^K |x_k|$$

with

$$\lambda_{\text{Max}} = \text{Max}|W \cdot y|.$$

Here $y = \{y_1, y_2, \dots, y_P\}^T$ is the vector of data taken at position ϕ_p , W is a $P \times (2K + 1)$ matrix with elements $W_{p,k} = \exp(i2\pi\phi_p n_k)$ and $x = \{x_{-K}, \dots, x_K\}^T$ is the vector of complex amplitudes associated with mode numbers n_k , $k = -K \dots K$. The hyper-parameter λ is discussed further below.

This criterion is convex, with no local minima, but as it is not ‘strictly convex’, uniqueness of the global minimizer is not guaranteed. However, it is shown in [11] that the global minimizer is likely to be unique if it has less than $N/2$ non-zero components, where N is the data size. A computationally efficient and convergent optimization strategy has been proposed in [9], based on a block coordinate descent algorithm.

The hyper-parameter λ is unknown and must be fixed to obtain a satisfactory sparse solution. An intuitive description of λ is that, since it increases the penalty for solutions which invoke a large number of modes, it influences the resolution of SparSpec: 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. From a theoretical point of view, λ can be interpreted as the maximum peak amplitude allowed in the periodogram of the residual [9], so that knowledge of the noise level (noise variance) may help to select this parameter. An ‘L-curve’ technique described in [9] can also be used to choose the optimum value of λ for a given data set. However, in our work, we selected this parameter in a more empirical way: the general approach was to use λ values as small as possible while still avoiding erratic or unlikely results in the output. The precise optimum value will depend on the data set and the level of noise, as well as the number of sensors contributing to the input and the number of modes contained in the data.

3. Benchmarking the SparSpec method for toroidal mode number analysis on JET

A batch processing code was developed for tokamak data in the following ways: first, the SparSpec code was adapted to take complex data as input. Then, because of the penalization term in $J(x)$, the minimizing amplitudes are initially underestimated; the original version of SparSpec involves a re-estimation of the detected frequencies using barycentric arguments, which improves the precision when these frequencies can be off the discrete grid, which is the case in astronomy. But in our case, we know that the frequencies are spatial mode numbers, in other words they are integer n , so a simple LS fitting routine was written to re-estimate amplitudes of the detected modes. The code was benchmarked with Monte Carlo simulations using artificial, noisy data. It was found to be very robust, and very fast (2 ms for one time point with 11 complex valued signals on a 2.2 GHz processor). To compare this result with the SVD-based technique described in [6], on a similar processor and for $M = 5$ and $n_{\text{max}} = 20$, it takes about 1 h to compute all required pseudo-inverse matrices, and 1 min to compute every LS solution.

On the JET experiment, there are eleven Mirnov coils at one common poloidal angle, unevenly spaced toroidally. The SparSpec method was applied to simulated data which

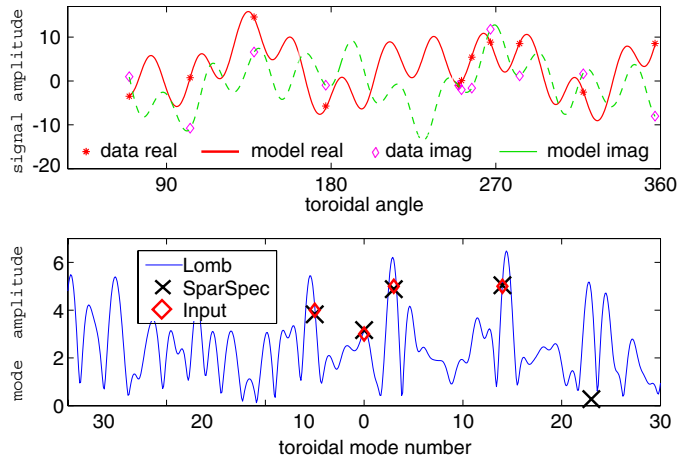


Figure 1. SparSpec calculation using simulated data: top: data (points) and model fit (lines). Bottom: input mode amplitudes versus estimation from SparSpec, along with Lomb periodogram.

mimicked data from JET, using identical sensor coordinates, with varying levels of noise. An example calculation is shown in figure 1, where data were constructed from four random modes (random n , amplitudes, phases), and with 5% random Gaussian noise added to each sensor. (A noise figure of 10% represents the addition of a random number generated with a mean of 0, standard deviation 1, divided by 10, on each of the eleven sensors.) The SparSpec mode number grid was restricted to $-40 < n < 40$. SparSpec correctly identifies the dominant four modes; the small-amplitude solution is due to the added noise. Throughout this paper, the high levels of noise (in both simulated and real data) and the low number of data points (11 sensors) limit the dynamic range of the analysis, and mode amplitudes are presented directly (rather than the logarithm of the amplitudes or power in the modes).

A systematic scan was then performed to benchmark the ability for SparSpec to resolve multiple modes in the Mirnov coil data. In the code, several parameters are set which facilitate better convergence and accuracy of the analysis, depending on the size of the input data array, the number of modes present in the data, the tolerance for precision, etc. The aim of the modeling was to find that set of parameters that would allow for an efficient and automated procedure for toroidal mode number separation. In particular, the results of SparSpec calculations are very sensitive to the hyper-parameter λ . The size of the set mode number grid (range of n to consider) can also play a role in the accuracy of the calculations, and speed. Several levels of random noise were introduced to determine if and under what conditions SparSpec can be trusted. Some of the results of these simulations are shown in the following figures. In general, SparSpec produces very good results when the number of modes in the data is less than five, and when noise levels are reasonable (less than 10% of signal amplitudes). These results are consistent with the theoretical results which guarantee uniqueness of the criterion minimizer only if it has less than five non-zero components for $N = 11$ data [11].

First, in order to qualitatively characterize the ability of SparSpec to determine toroidal mode amplitudes and phases, a set of artificial data was constructed which consisted of a ‘time’ series which included a superposition of 21 toroidal modes ($-10 < n < 10$), with varying amplitude (A_n) and phases (θ). Thus, at each time point the data in the sensor located at ϕ_k consisted of: $\text{data}_k = \sum_n A_n e^{i(n \cdot \phi_k + \theta)}$.




Figure 2. Amplitude and phase of input data for simulated time series. Different colors correspond to different n -numbers, consecutively increasing starting with $n = -10$. The data shown here feeds into SparSpec an average of six modes to identify at any given time.

The data sets were constructed such that each mode had maximum amplitude at a unique time, within a Gaussian envelope. Each mode was designed to experience a full range of phases ($0-2\pi$ radians) within the time in which it had significant amplitude. Many different data sets were constructed, with different widths of Gaussian mode envelopes, so that for different sets the amount of mode overlap was varied. The more mode overlap, the greater the challenge to SparSpec, since it is expected to identify a larger number of simultaneous modes. The generated data contained no noise, i.e. the input to SparSpec represented ideal measurements of amplitude and phase at each sensor (figure 2). One such data set is illustrated in figure 3. SparSpec evaluated each set using three different values of the hyper-parameter, λ . Other parameters in SparSpec were held fixed, for example the mode number grid in which SparSpec searched was limited to $-30 < n < 30$.

The outcome of these calculations illustrates several trends as seen in figure 3: when there is only one mode, nearly any λ value is adequate and SparSpec correctly identifies the amplitude and phase. (For real data with noise, the λ parameter should be set higher in order to preferentially suppress solutions containing multiple modes.) When there are several simultaneous modes, for λ values between 0.2 and 0.9, SparSpec correctly identifies the modes with the strongest amplitudes. This means that SparSpec is fairly robust, errors are expressed in the form of missing information about lower amplitude modes, and SparSpec does not indicate toroidal mode numbers which are not present in the input data. This trend however has limits, and with the addition of noise is only held true by higher values of the λ parameter. Fortunately, the general trend for higher λ s is for SparSpec to ignore lower amplitude modes.

Next, in an attempt to quantify the accuracy of the SparSpec calculations, Monte Carlo type simulations were carried out, using the following scheme: randomly generated data sets with different mode numbers, phases, and relative amplitudes were input into SparSpec. Noise was included in the input data by generating Gaussian noise on each sensor. In each case, the input mode amplitudes were normalized ($\sum |A_q| = 1$, $A_q =$ amplitude in q th mode), and the

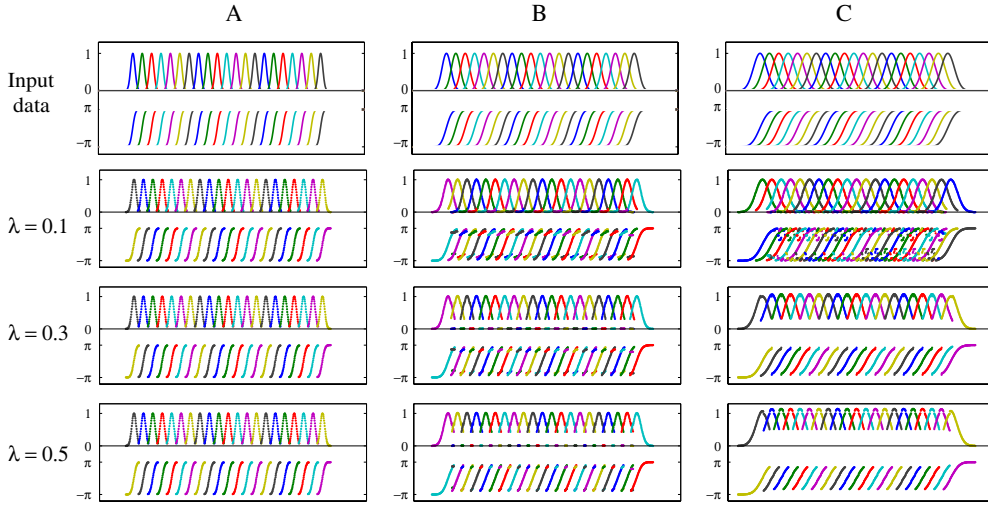


Figure 3. SparSpec estimations for noiseless input data, $-10 < n < 10$, for a few λ parameters, and three levels of ‘difficulty’ (number of simultaneous modes in the data). Good performance is characterized by good reproduction of top row of plots, which depict the input data.

error in the SparSpec output was evaluated by summing the difference between the input mode amplitudes and output amplitudes at all mode numbers, i.e.

$$\sum_n |A_n^{\text{in}}| = 1, \quad \text{err} = \sum_{-n_{\text{max}}}^{n_{\text{max}}} |A_n^{\text{in}} - A_n^{\text{out}}|.$$

If $\text{err} \geq 0.5$, the calculation was considered to be 100% wrong. This is a very conservative way of estimating errors because an error for one mode number can throw the entire result into the 100% error bin, even though there may be several modes which SparSpec has identified correctly. Monte Carlo type simulations then produced results which were histogrammed as shown in figure 4. Evident in the figure is the fact that SparSpec either correctly identifies all the modes and their amplitudes and phases, with little error, or completely misses the mark. This can be understood by recognizing that there are several solutions (with vastly different n s and amplitudes) which come close to reproducing the input data.

1000 calculations were performed for 10 different noise levels, 4 different n -number ranges and 6 different input data types (from having only one n present to finally six distinct modes). The result of one subset of the calculations is shown in figure 5. The full set of results is shown in a set of bar charts in figure 6. The general conclusions are that SparSpec can handle up to five simultaneous modes when the noise level is effectively 0. In the presence of sensor noise, SparSpec calculations are mostly correct when there are four or less modes present in the data even when the mode spectrum contains $n \leq 40$; a greater number of modes in the output should probably not be trusted.

The SparSpec method was also used to look at unstable MHD fluctuations in JET where the toroidal mode numbers were well determined by a phase-fitting technique. The data consisted of ‘pure’ Alfvén and neoclassical tearing modes, in the sense that they contained only one single n -number. These investigations covered the entire available frequency range up to the sampling Nyquist frequency of the coils ($5 \text{ kHz} < f < 500 \text{ kHz}$), in order to ensure that small differences in the transfer function of the coil/cable network between individual Mirnov coils did not lead to any erroneous mode identification with SparSpec. In all cases, SparSpec

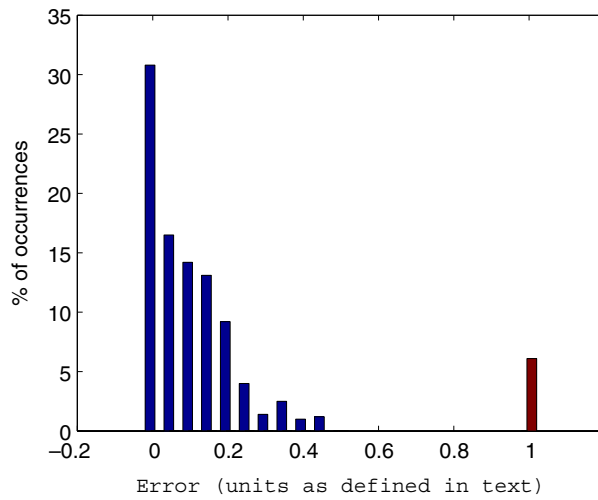


Figure 4. Result of 1000 simulations with three modes included in input, comprised of random n s ($-40 < n < 40$), amplitudes, and phases, with 5% noise, fixed $\lambda = 0.5$. Errors greater than 0.5 are binned into '1', represented by red bar.

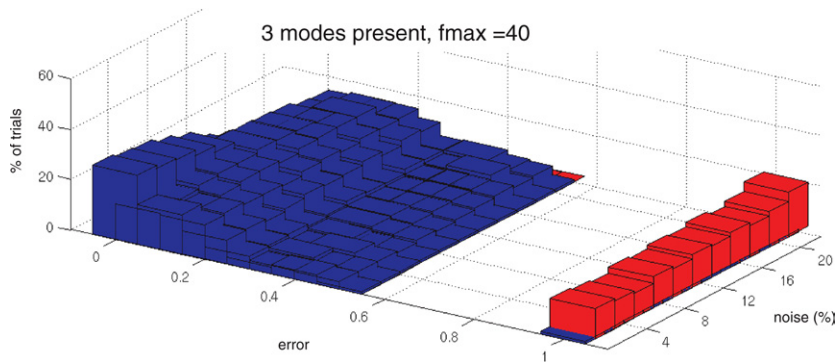


Figure 5. Error statistics for two random modes, $-40 < n < 40$, versus noise level, fixed $\lambda/\lambda_{\max} = 0.5$. SparSpec correctly separates modes 94% of the time, even if there is 20% noise present in the data.

correctly indicated the mode number, and also showed that there were no other small-amplitude mode numbers with other mode numbers in the data. This detailed benchmarking against the phase-fitting technique revealed a number of advantages of SparSpec over the standard phase-fitting method, even when the data contained only one mode number.

First, SparSpec produced a result orders of magnitude faster than finding a best fit to the phases of eight sensors, even when the maximum n -number for SparSpec to consider was set to $n = 30$, while the phase-fitting algorithm was restricted to $0 < n < 7$. Second, unless the phase-fitting routine incorporates some method for disregarding bad fits, this method always produces a result with a definite n -number. But with SparSpec, this is not necessarily the case. Third, the phase-fitting technique does not yield any amplitude information about any of the modes in the data. And finally fourth: when two modes with different mode numbers cross in frequency, phase fitting no longer produces correct results.

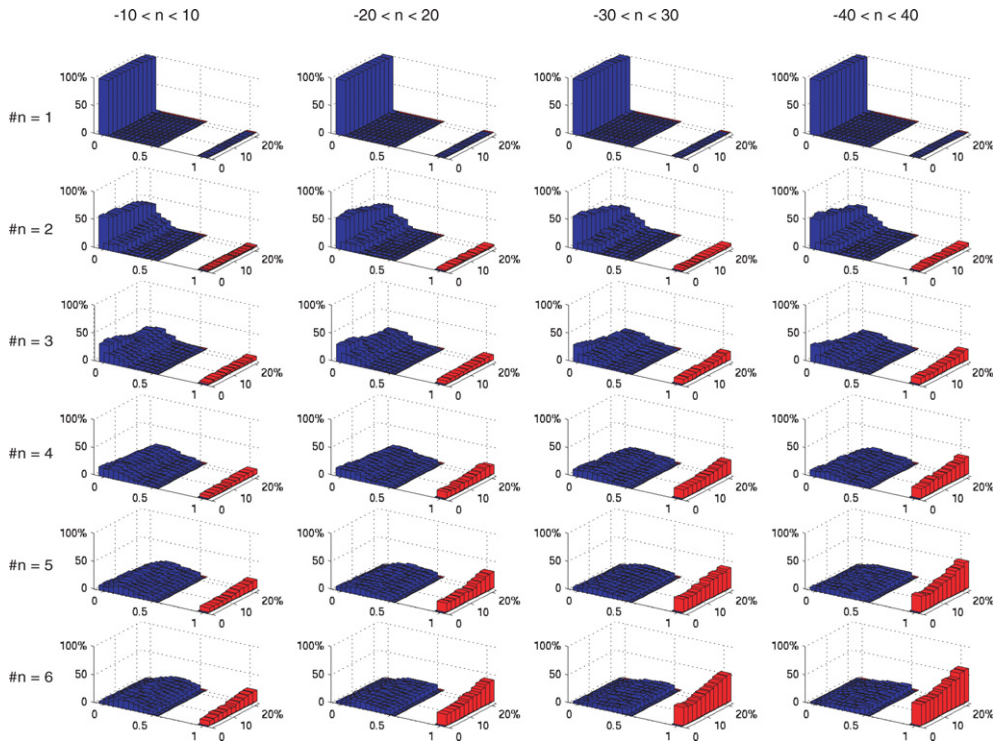


Figure 6. Error results for many noise levels, numbers of modes contained in the data, and maximum mode number included in the data sets (= max n sought), fixed $\lambda/\lambda_{\max} = 0.5$.

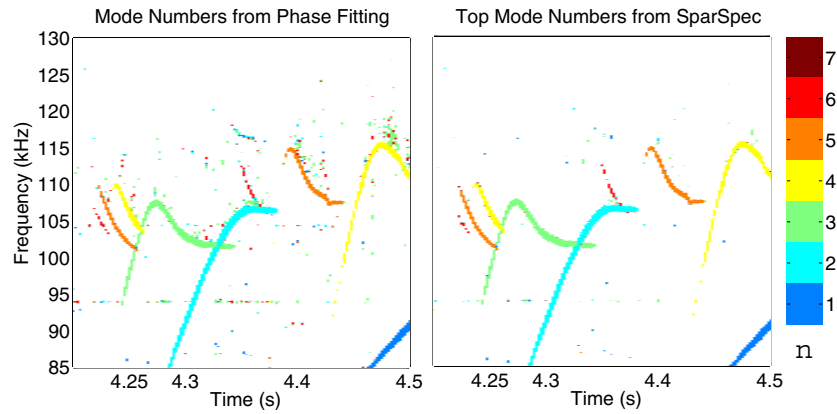


Figure 7. Toroidal mode spectrum analysis for small frequency and time window in JET shot #69436 using eight Mirnov coils. Left: result obtained with LS fit of phases to model of single mode. Right: result obtained with SparSpec.

An example of a SparSpec result versus a phase-fitting result is shown in figure 7, which depicts a period of reversed-shear Alfvén eigenmode (RSAE) activity in a JET discharge. Such modes are common in shear-reversed plasmas when the minimum- q reaches an integer value, characterized by a ‘grand cascade’ of AEs with mode numbers $n = 1, 2, 3, \dots$ RSAEs

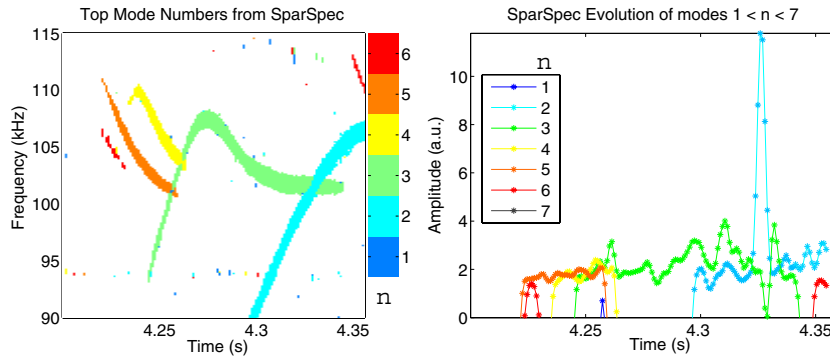


Figure 8. Toroidal mode spectrum analysis for small frequency and time window in JET shot #69436 using eight Mirnov coils. Left: dominant mode numbers obtained with SparSpec, right: detailed mode amplitude evolution during this time/frequency window.

are usually observed to chirp up to the TAE gap frequency [12]. These plots were generated in the following way: a windowed fast-Fourier transform (FFT) was performed on signals from eight Mirnov coils to produce amplitude and phase in frequency and time space for each coil. Only amplitudes in the top 5% were retained for the mode analysis, the rest were set to produce a white space. On the left, the phase-fitting analysis considered only mode numbers with $0 < n < 7$, and took several minutes to complete. The SparSpec figure on the right shows only the most dominant mode found by SparSpec for each pixel. Mode numbers with $-30 < n < 30$ were considered, and the plot was produced within 20 s after the FFTs were complete. Finally, in each plot only modes with toroidal mode numbers $0 < n < 7$ are displayed. Clearly, the two plots are identical, with the exception that SparSpec has eliminated some noise, because it found noisy data to be composed of high n -numbered modes (which are filtered from the plot).

To further illustrate how SparSpec can provide detailed insight into MHD phenomena using magnetic pick up coil data, a smaller sub-region of the same time and frequency window is shown in figure 8. On the left, the RSAEs are seen to sweep in frequency and overlap at some points, depicted in the same manner as in figure 7. On the right are the sum of the modes found by SparSpec ($1 < n < 7$) during this time and frequency window, showing the amplitude evolution of the individual modes. Interestingly, coupling between modes can be observed, so that near $t = 4.33$ the $n = 2$ mode seems to gain energy from the mode with $n = 3$. With SparSpec, such observations can be made routinely.

4. Resolving stable AEs on JET

One area in which the detailed mode number identification of MHD fluctuations is desired is known as MHD spectroscopy [13]. The term refers to a broad topic in which fluctuations are used to deduce information about bulk plasma parameters, for example some properties of the q -profile are revealed with the observation of RSAEs. The TAE antenna project on JET [14, 15] is an example of active MHD spectroscopy, where small magnetic perturbations are actively produced by external antennas, and the plasma response delivers information about toroidal (and other kinds) of AEs. The motivation is that multiple AEs with several toroidal mode numbers are known to be excitable by supra-thermal ion populations, which in some circumstances can lead to rapid loss of fast particle confinement [16, 17]. This would pose

problems for a burning plasma scenario, and therefore measuring the damping rates of AEs in the absence of fast particle drive and in a variety of plasma conditions will aid in understanding some underlying physics relevant to burning fusion plasmas.

The TAE antennas on JET consist of an array of small coils which can excite medium and high- n AEs, depending on the polarity of the individual antennas (one amplifier is the common source). Because of the small spatial extent of the antenna arrays (figure 7, left), the magnetic perturbations to the plasma surface necessarily are composed of a broad spectrum of toroidal mode numbers (figure 7, right). Data from Mirnov coils are obtained via synchronous detection hardware, i.e. a very sharp bandpass filter (± 100 Hz) produces the in-phase and quadrature signal components only at the antenna frequency, and these signals are then used to facilitate real-time resonance detection and tracking [8]. The damping rate of the detected resonant mode is determined post-shot by evaluating the poles and residues in the transfer function between Mirnov signals and antenna currents [14].

On JET, for unstable modes driven by fast ions, multiple AEs with medium toroidal mode numbers ($3 < n < 12$) typically oscillate with nearly the same frequency of a few 100 kHz, but the observed frequencies in the lab frame are separated by a small Doppler shift due to plasma rotation (since the mode rotates with the plasma fluid, $\Delta f = \Delta n \times \Omega$), typically by a few kilohertz. This is enough to determine toroidal mode numbers separately at each frequency using standard techniques. By contrast, for the case of externally excited (stable) AEs, any modes near resonance with the driving frequency will be excited at that same frequency.

In the presence of only one dominant resonant mode, n -number identification is straightforwardly accomplished with straight line phase fitting. The resonances of other, less dominant (more stable or more core-localized) modes with different n -numbers will be missed, however. In addition, when the phase-fit is very bad or jumps from one n to another with incremental changes in frequency, it is a strong indication that there are several modes being detected simultaneously (see figure 10). This latter case can completely prevent n -number identification by traditional methods such as straight line phase fitting. It is however essential to be able to resolve individual toroidal modes, to determine their damping rates accurately and compare these results with theoretical predictions, and it is for this reason that SparSpec was originally adapted for use with toroidal plasma fluctuations.

An example of a SparSpec application is shown in figure 11 for JET shot #70723, which depicts SparSpec output for the six most dominant modes in the time interval $6.4 < t < 7$ s. In this plot, there are two types of regions: near $t = 6.75$ s, region 'A', there is clearly one dominant mode ($n = -1$), which could have been identified using other methods for n -number analysis. In the same region, some other modes (with $n = -2, -3$) are also excited, but at a much lower level. Another region near the beginning of this time interval, region 'B', shows an example of when the antennas are driving several modes with several n -numbers, with competing amplitudes.

The reader should note that negative mode numbers in this context do not relate to the helicity of the mode structure, nor to the toroidal propagation direction of the detected modes; rather they are merely the consequence of calculations which involve complex data (the I and Q from the synchronous detectors). Helicity (and poloidal mode structure) cannot be determined using data sampled at only one poloidal angle, and the mode propagation is determined by the time evolution of the phase of a mode when the antenna frequency is fixed. Pairs of mode numbers with equal amplitudes and opposite sign ($\pm n$) do indicate standing waves, i.e. pairs of counterpropagating modes, which are characterized by the presence of toroidal nodes where the sensor response is always small. This is rarely seen on the JET TAE antenna diagnostic, as any plasma rotation will separate counterpropagating modes due to a Doppler shift in the mode frequency according to $\omega_{\text{lab}} = (\omega_{\text{mode}} + n \times \omega_{\text{plasma}})$.

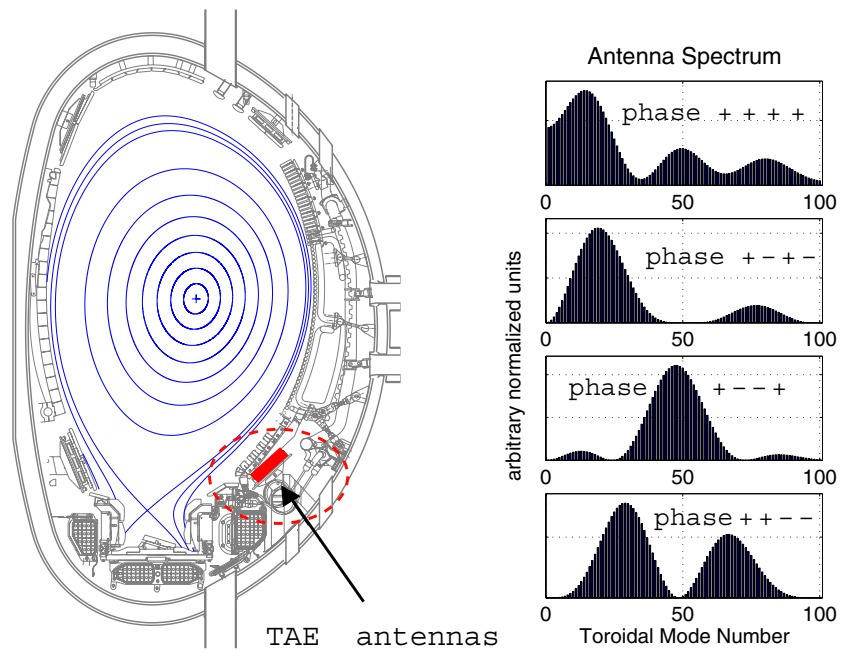


Figure 9. Left: active MHD antenna location in the JET tokamak. Right: calculation of approximate toroidal mode spectra (only positive n shown) being driven in vacuum with four different antenna phase configuration.

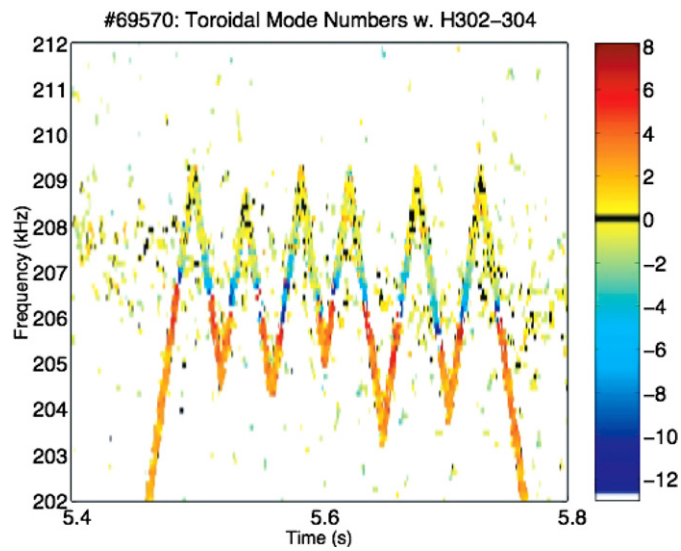


Figure 10. Stable TAEs excited by external antennas in JET shot #69570. Active antennas sweep in frequency to repeatedly scan over a detected resonance. Determination of n -number using simple straight line phase fitting of three closely spaced Mirnov coils shows drastic jumps in n -number for incremental frequency changes and indicates the presence of multiple modes.

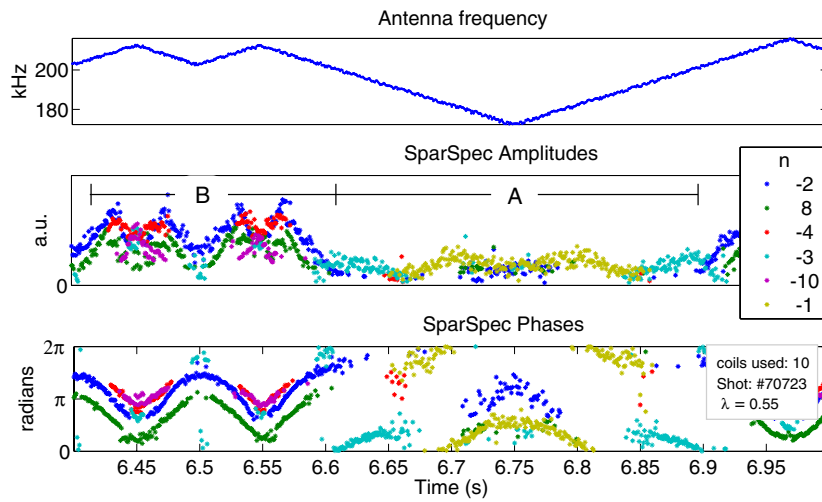


Figure 11. Left: detail JET shot #70723 from $6.4 < t < 7.0$. While antennas sweep in frequency (top plot), several AEs excited as seen in SparSpec output. Region A: one mode with $n = -1$ dominates. Region B: several modes are driven to comparable amplitudes.

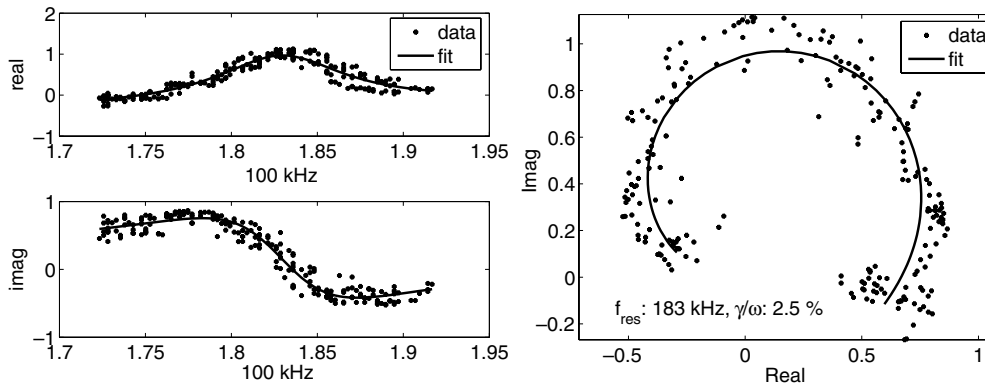


Figure 12. Left: damping calculation using SparSpec output for JET shot #70723 from $6.6 < t < 6.9$ (region A in figure 11). Left: real and imaginary component of $n = -1$ mode versus antenna frequency, right: complex plane representation (solid lines fit according to [14]).

When a very small time interval is analyzed, the changes in amplitude and phase of an individual mode as a function of incremental changes in antenna frequency can be used to calculate the damping rate, γ , of the mode in that time interval [14]. Such calculations are shown in figures 12 and 13, for both region A and B of figure 11. When plotting the in-phase and quadrature (real and imaginary) components in the complex plane, a circle results. SparSpec is able to track amplitude and phase evolution of multiple modes, and calculation of stability parameters for each mode is possible (figure 13).

5. Further considerations

While perhaps obvious, it is nevertheless important to note that in the presence of significant variation across sensors in terms of frequency dependent complex gain (amplitude and phase),




Figure 13. Left: damping calculation using SparSpec output for JET shot #70723 from $6.4 < t < 6.6$ (region B in figure 11). Left: real and imaginary component of modes with $n = -2, -4, 8$ versus antenna frequency, right: complex plane representation (solid lines fit according to [13]).

the SparSpec method no longer produces correct toroidal mode numbers. Such differences may be the result of conducting structures nearby individual probes, subject to different eddy current patterns, or due to differences in probe construction or electrical transmission line characteristics. Thus, for accurate toroidal mode number reconstruction the transfer function of each individual probe must be known, and the data calibrated, at least relatively. Of course, such is also the case for any other mode number identification method. For the magnetic probes on JET, uncertainties in the relative calibration of Mirnov coils at high frequencies lead to systematic errors in the analysis, so for example pure low- n modes will appear to be accompanied by a small-amplitude high- n mode (e.g. $n = 29$) at high frequencies. To thoroughly discuss these errors is beyond the scope of this paper; suffice it to say that such errors are small for JET data below 300 kHz.

Another topic which is of some importance in evaluating the abilities of SparSpec deals with the spacing and positioning of individual Mirnov coils, as well as the number of coils available for use. We have investigated the effect of omitting one or more channels in the analysis, characterizing the efficacy of SparSpec. We merely mention that even with three sensors removed, the SparSpec method generally yields the same results as with all eleven channels, but this depends somewhat on *which* channels are removed. Other tokamaks will have their own sensor geometry, and entirely different applications will naturally have their own sampling strategy. This paper presents a case study of the SparSpec method as applied to JET data from eleven unevenly spaced Mirnov coils and which contains information about small amplitude, medium- n numbered AEs. Of course the probe positions on JET are fixed, but a general investigation might be worthwhile in view of the (not yet completed) design of the magnetic probe diagnostics for ITER. Again it is well beyond the scope of this paper to investigate these topics in any systematic way, because optimized sampling strategy is a topic of some depth, the reader is referred to [3] for a detailed treatment.

Nevertheless, for illustrative purposes, one example of the effect of sensor positioning will now be described. In the modeling of artificial data, it was observed that SparSpec produces very small artifacts near $n = \pm 10$ and ± 20 , i.e. the output contains small peaks at various mode numbers when the input data contains modes $n = 9, 10, 11$. These false peaks, while negligibly small, seem to be result from the toroidal arrangement of the sensors. They can be understood by considering that for regularly spaced coils, the spectral window

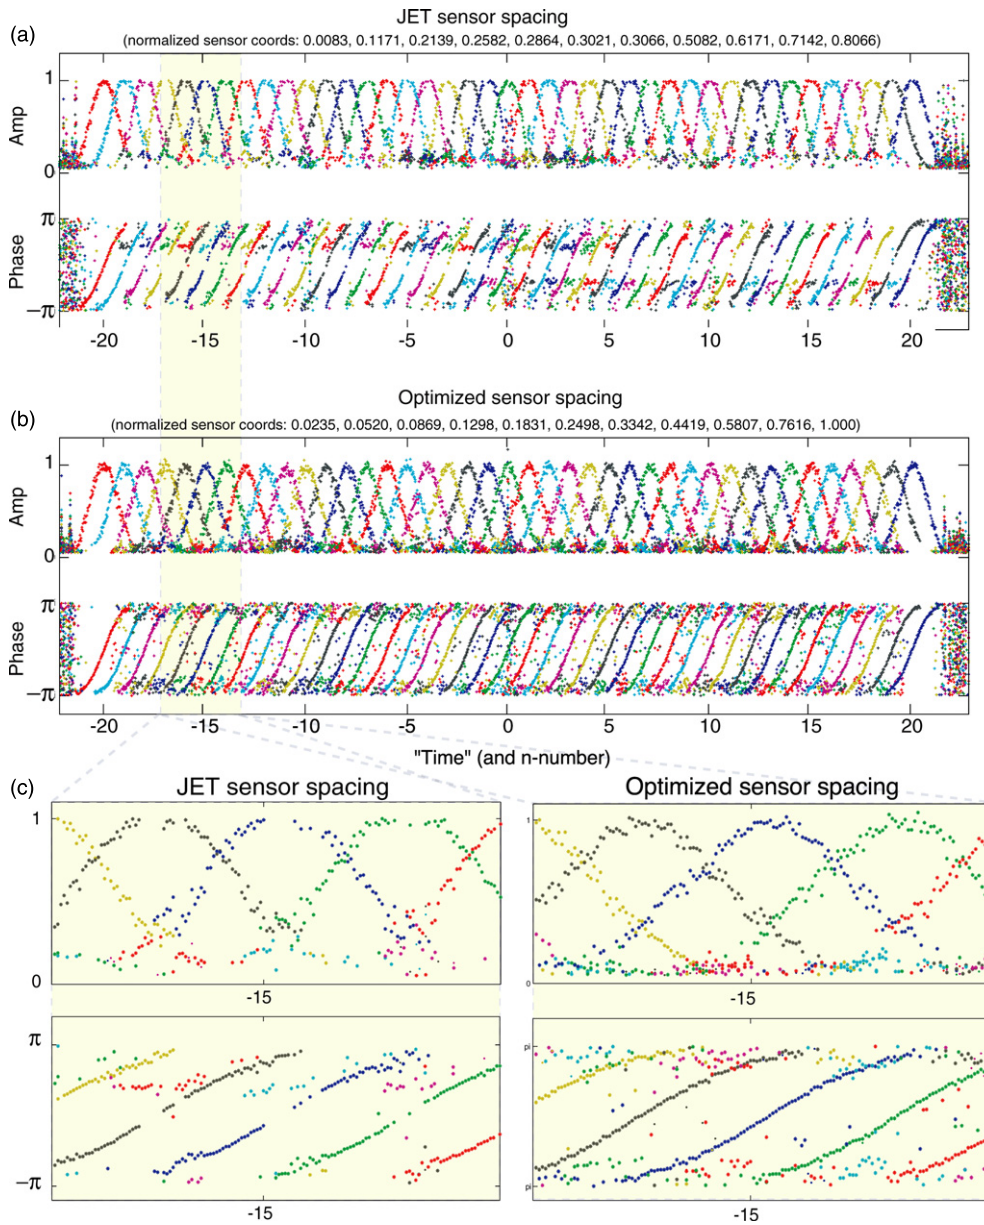


Figure 14. SparSpec calculations estimations for input data with 5% noise, $-20 < n < 20$, λ is fixed at 0.4. Interpretation as in figure 3. (a) Output using JET sensor geometry, (b) output using suggested sensor spacing from [18]. (c) Zoom detail of amplitudes (left) and phases (right) showing minor benefits of optimum spacing.

is a Dirichlet kernel of period N , where N is the number of coils. In the irregular sampling case (such as at JET), the spectral window also can have secondary lobes, due to a hidden regularity in the sampling: 8 of the 11 sensors in use on JET actually form part of an array of 10 equally spaced Mirnov coils. Using simulated data, when the sensor spacing is optimized according to results by Madore and Freedman (table 1 in [18]), SparSpec output

calculations do not exhibit any errors of the kind described. A more noticeable difference in performance emerges when processing data which contains some amount of noise and is discernable in figure 14. This figure depicts SparSpec output data in the same fashion as figure 3 in section 2. Two scenarios were modeled: one in which the sensor geometry is JET-like (a), and one in which the spacing of eleven sensors was ‘optimized’ (b). The JET geometry result has a slightly larger amount of errors, particularly when the phases of modes in the input data take on certain values (figure 14(c)). What appears as missing data points is caused by the failure for SparSpec to identify the particular mode(s). The optimized geometry also allows slightly better resolution into lower amplitudes, as seen by studying the figures closely.

6. Conclusion

A new method for identifying toroidal mode numbers in Mirnov data from toroidal plasmas has been benchmarked. Embodied in the SparSpec code, this new method fits signals which are unevenly sampled in the toroidal coordinate to a sum of an arbitrarily large number of toroidal modes with integer mode numbers. By assigning a larger penalty to solutions that invoke larger numbers of modes, SparSpec determines the best fit with the sparsest spectrum. SparSpec has proven to be extremely robust, and is found to be especially useful for resolving the amplitudes and phases of multiple AEs which are ringing with the same or nearly the same frequency. Examples involving stable AEs in JET, excited by an array of external antennas, were used to illustrate the efficacy of the method. The method also is superior in determining n -numbers when there is only one dominant mode present, as compared with traditional straight line phase-fitting techniques, because both amplitude and phase information is considered, and thus discrimination against noise is improved.

The great efficiency with which SparSpec detects multiple modes in large datasets suggests that it may be readily used in real-time applications such as resistive wall mode (RWM) or tearing mode control, among others. For example, in large tokamaks the growth rates of RWMs are typically orders of magnitude slower than the computation time of SparSpec even when it incorporates dozens of magnetic signals. We plan to use SparSpec in feedback calculations in order to target specific n -numbered modes during active Alfvén spectroscopy, using a subset of eight Mirnov coils, and with loop rates ≤ 1 ms.

Acknowledgment

This work is partially supported by US DOE Grant DEFG02-99ER54563.

© Euratom 2008.

References

- [1] Kim J S *et al* 1999 *Plasma Phys. Control. Fusion* **41** 1399
- [2] Zanca P *et al* 2001 *Phys. Plasmas* **8** 516
- [3] Marvasti F 2001 *Nonuniform Sampling: Theory and Practice* (Berlin: Springer)
- [4] Lomb N R 1976 *Astrophys. Space Sci.* **39** 447
- [5] Zegenhagen S *et al* 2006 *Plasma Phys. Control. Fusion* **48** 1333
- [6] Hole M and Appel L 2007 *Plasma Phys. Control. Fusion* **49** 1971
- [7] Roberts D, Lehar J and Dreher J 1987 *Astron. J.* **93** 968
- [8] Foster G 1995 *Astron. J.* **109** 1889

- [9] Bourguignon S, Carfantan H and Böhm T 2007 *Astron. Astrophys.* **462** 379
- [10] Donoho D 2006 *J. Commun. Pure Appl. Math.* **59** 797
- [11] Bourguignon S, Carfantan H and Idier J 2007 *IEEE J. Sel. Top. Signal Process.* **1** 4
- [12] Sharapov *et al* 2001 *Phys. Lett. A* **289** 127
- [13] Fasoli A *et al* 2002 *Plasma Phys. Control. Fusion* **44** B159
- [14] Fasoli A *et al* 1995 *Phys. Rev. Lett.* **75** 645
- [15] Fasoli A *et al* 2000 *Phys. Plasmas* **7** 1816
- [16] Wong K 1991 *Phys. Rev. Lett.* **14** 1874
- [17] Heidbrink W *et al* 2003 *Plasma Phys. Control. Fusion* **45** 983
- [18] Madore B F and Freedman W L 2005 *Astrophys. J.* **630** 1054