High-definition velocity-space tomography of fast-ion dynamics

M. Salewski[1], B. Geiger[2], A.S. Jacobsen[2], P.C. Hansen[3], W.W. Heidbrink[4], S.B. Korsholm[1], F. Leipold[1], J. Madsen[1], D. Moseev[5], S.K. Nielsen[1], M. Nocente[6,7], T. Odstrčil[2], J. Rasmussen[1], L. Stagner[4], M. Stejner[1], M. Weiland[2] and the ASDEX Upgrade team[2]

[1] Department of Physics, Technical University of Denmark, Kgs. Lyngby, Denmark
[2] Max-Planck-Institut für Plasmaphysik, Garching, Germany
[3] Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kgs. Lyngby, Denmark
[4] University of California Irvine, Irvine, CA, USA
[5] Max-Planck-Institut für Plasmaphysik, Greifswald, Germany
[6] Department of Physics, University of Milano Bicocca, Milano, Italy
[7] Istituto di Fisica del Plasma, Consiglio Nazionale delle Ricerche, Milano, Italy

E-mail: msal@fysik.dtu.dk

CrossMark

Abstract

Velocity-space tomography of the fast-ion distribution function in a fusion plasma is usually a photon-starved tomography method due to limited optical access and signal-to-noise ratio of fast-ion D_α (FIDA) spectroscopy as well as the strive for high-resolution images. In high-definition tomography, prior information makes up for this lack of data. We restrict the target velocity space through the measured absence of FIDA light, impose phase-space densities to be non-negative, and encode the known geometry of neutral beam injection (NBI) sources. We further use a numerical simulation as prior information to reconstruct where in velocity space the measurements and the simulation disagree. This alternative approach is demonstrated for four-view as well as for two-view FIDA measurements. The high-definition tomography tools allow us to study fast ions in sawtoothing plasmas and the formation of NBI peaks at full, half and one-third energy by time-resolved tomographic movies.

S Online supplementary data available from stacks.iop.org/NF/56/106024/mmedia

(Some figures may appear in colour only in the online journal)

1. Introduction

It is often convenient to split ion distribution functions in magnetized fusion plasmas into two parts. One part contains fast ions from fusion reactions or auxiliary heating and the other part thermal ions. The latter is by definition described by a Maxwellian distribution and thus fully determined by a temperature T_i, a density n_i and a drift velocity $\mathbf{v}_i$. As T_i and n_i are nearly constant on a flux surface, bulk-ion measurements often refer to flux surface measurements of these fundamental

parameters [1–4]. Fast-ion velocity distribution functions are much more complex. They are not flux functions due to the large drift excursions of fast ions, and they are in general 6D functions in phase space. Nevertheless, the cyclotron motion in strongly magnetized plasmas implies an approximate rotational symmetry which effectively reduces the dimensionality of velocity space to two. Usually fast-ion measurements do not refer to measurements of the fundamental high-dimensional fast-ion distribution functions but rather to measurements of local 1D projections or other derived quantities such as the

measured spectra in collective Thomson scattering (CTS) [5–8], fast-ion D$_\alpha$ spectroscopy (FIDA) [9–11], neutron emission spectrometry (NES) [12–15] and γ-ray spectrometry (GRS) [16–19]. This is in contrast to bulk-ion measurements where the fundamental parameters T_i, n_i and $\mathbf{v}_i$ are inferred from the spectra and presented as measurements whereas the spectra themselves are of secondary interest.

Velocity-space tomography allows us to infer the fundamental 2D fast-ion velocity distribution functions from measured spectra in analogy to bulk-ion measurements [20–30, 31]. The 2D velocity distribution function is spatially localized in a small measurement volume which we take to be a single point in position space. The fast-ion measurements depend on phase space in a complicated way illustrated by so-called weight functions which have been formulated for FIDA [10, 32], neutral particle analyzers (NPA) [10], CTS [21], fast-ion loss detectors [33], NES [34, 35] and GRS [36, 37]. To exploit the rich information about fast ions contained in the measurements by traditional procedures, we need to consider hundreds of data points, e.g. spectral bins, together with the corresponding weight functions. Further, the spectral measurements also depend on nuisance parameters such as bulk-ion densities or temperatures.

Velocity-space tomography provides a way to process this wealth of information at once by inverting the data. It provides a 2D image that is straightforward to interpret, that is the best useful fit to hundreds of simultaneous measurements from different diagnostics, that shows the fundamental quantity of interest rather than derived quantities, and that accounts for nuisance parameters. The tomography approach also allows an alternative way to compare fast-ion measurements and numerical simulations. Traditionally, this is done by comparing the measurements with simulated measurements in units particular to the diagnostic, e.g. the spectral density of the measured neutron or photon fluxes (D_α, γ, or mm-wave). Velocity-space tomography allows us to use the fundamental 2D velocity distribution function as a meeting ground between theory and observation for any combination of fast-ion diagnostics [23, 29, 31].

Until now velocity-space tomography has relied on standard inversion methods such as truncated singular value decomposition (TSVD), the maximum entropy method, and variants of the Tikhonov regularization [24, 25, 27, 28]. For measurements with high data quality and signal-to-noise ratio, these inversion methods work well. We have previously shown that it is possible to correctly reconstruct the injection energy in plasma heated by neutral beam injection (NBI), and good agreement with TRANSP predictions was found in the absence of strong magnetohydrodynamic (MHD) activity [24, 27]. Fast-ion velocity-space redistribution patterns of sawtooth crashes could also be reconstructed with confidence [26–31]. However, velocity-space tomography is usually a photon-starved enterprise since the signal-to-noise ratio is often low compared with many other tomography applications and the optical access to tokamak plasmas is limited. The FIDA emission is weak compared with other components of the measured spectra, e.g. the beam emission. In addition,

the spectra are often polluted with strong background contributions from bremsstrahlung and impurity line radiation. This hampers the analysis in plasmas with high electron density or high impurity content. In conflict with the limited amount of data, we strive for high-resolution images requiring the inference of many unknowns. Consequently, the inversions based on standard methods have been plagued by artifacts, for example non-zero phase-space densities at energies larger than the NBI energy or negative phase-space densities, even though artifacts can be minimized in optimized discharges usually based on L-mode plasmas with low density and low heating power [24, 26–28]. Artifacts can be attenuated by increasing the number of measurement data points by installation of additional fast-ion diagnostics [22, 27]. However, often economic or technical constraints do not allow this. In high-definition velocity-space tomography we make up for the lack of data through various types of prior information. This not only improves results for the five-view FIDA diagnostic at the tokamak ASDEX Upgrade [38], but it also allows the use of inversion techniques for more common FIDA systems with two or three views. The attenuation of artifacts through the use of prior information further allows an increased resolution of the images.

We discuss prior information for velocity-space tomography in section 2 and the choice of the regularization strength using current inversion techniques in section 3. Substantial improvements brought about by prior information are demonstrated in sections 4 and 5. In section 6 we study fast-ion velocity distributions in NBI heated plasmas and dynamics in sawtoothing plasmas using high-definition movies applying our new techniques. Finally, we draw conclusions in section 7.

2. Prior information for velocity-space tomography

Velocity-space tomography entails the solution of an ill-posed problem in which we seek F^* solving the matrix equation

$$WF^* = S \qquad (1)$$

where W and S are known [22]. F^* is the fast-ion velocity distribution function discretized in n pixels, S holds m fast-ion measurements, and W is an $m \times n$ matrix composed of weight functions. S and W are normalized by the uncertainties of the measurements [23]. Noise in the measurements makes the rows of the matrix equation (1) inconsistent, so that there is no solution irrespective of the choice of n ($n = m$, $n < m$, or $n > m$). One might have hoped that one could instead solve the related least-squares problem

$$F^*_{\text{LS}} = \arg\min_F \lVert WF - S \rVert_2, \qquad (2)$$

where F is an arbitrary fast-ion velocity distribution function and F^*_{LS} is the least-squares solution. However, this is also useless since the matrix W is ill-conditioned and hence the solution F^*_{LS} is not stable. This means that small perturbations in S can lead to large perturbations in F^*_{LS} which is therefore dominated by random jitter. Nevertheless, we can construct a related well-conditioned problem by imposing additional requirements that reflect prior assumptions about the solution

Table 1. Prior information used in this article, their benefits and risks.

Prior	Benefits	Risks
L (1st-order)	Good for smooth functions.	Misses spikes and ridges.
$F > 0$	Always true, improves solutions.	$F < 0$ could point to data errors that are now missed.
$F(E_0, p_0) = 0$	Avoids artifacts in null-measurement region.	Misses marginal densities of fast ions in null-measurement region. Possibly artifacts at the boundary.
$\kappa(E, p)$	Accounts well for NBI peaks.	Misses any peak displacement. Might introduce spurious peaks.
F_{sim}	Good at locating discrepancies in velocity space.	Possible misinterpretation if the absolute scaling is wrong.

Note: Each prior is optionally used when appropriate.

and that provide useful and stable solutions. This is called regularization. A popular regularization method in plasma physics is the Tikhonov regularization [28, 39, 40] in which we solve the minimization problem

$$F^* = \arg \min_F \{ \|WF - S\|_2^2 + \lambda^2 \|LF\|_2^2 \}. \tag{3}$$

In equation (3), F^* minimizes the sum of the residual of the original ill-posed problem (first term) and the norm of the additional requirement on the solution (second term). A common choice is to require that the norm of the gradients in F^* is small, i.e. that the solution is smooth. In this so-called first-order Tikhonov regularization, the penalty matrix L is a matrix operator effecting a finite difference approximation of a gradient [28]. We calculate the gradient in $(v_\parallel, v_\perp)$-space where parallel and perpendicular denote directions with respect to the magnetic field [28]. The free regularization parameter λ balances how well the solution should fit the noisy data and how much it should obey the regularization constraint. Equation (3) shows that for small λ's the residual of the original problem dominates whereas for large λ's the norm of the regularization constraint dominates. The challenge now is to select λ leading to a useful and stable solution as we will discuss in section 3. An equivalent formulation of Tikhonov's minimization problem in equation (3) is

$$F^* = \arg \min_F \left\| \begin{pmatrix} W \\ \lambda L \end{pmatrix} F - \begin{pmatrix} S \\ 0 \end{pmatrix} \right\|_2. \tag{4}$$

Our computations of Tikhonov solutions in the following will be based on the formulation in equation (4) as it is most stable and best suited for numerical computations. Our exposition will focus on Tikhonov regularization, but our methods are applicable to other regularization methods as we will demonstrate for TSVD. The types of prior information presented in the following sections are summarized in table 1 together with their benefits and risks. Each type of prior information can be optionally added when appropriate.

2.1. Null measurements

FIDA measures Doppler-shifted D$_\alpha$-light emitted when a fast deuterium neutral is formed in a charge-exchange reaction from a fast deuterium ion and then decays from the third to the second excited state [9]. In many experiments, there is an upper limit to the observed Doppler shifts. Parts of a spectrum where no FIDA light is observed above the noise floor are referred to as null measurements. The wavelength ranges of null measurements are related to velocity-space regions through weight functions [32]. We refer to such weight functions as null-measurement weight functions. If treated on an equal footing with detections of FIDA light, null measurements already contribute strongly to the reconstruction of the large-scale shape of the velocity distribution function since they tend to decrease the reconstructed phase-space densities in the velocity space covered by the null-measurement weight functions. For that reason the FIDA system at ASDEX Upgrade was upgraded to measure red- and blue-shifted light in all spectra so that the absence of FIDA light could be measured [27]. Still, inversions are plagued by artifacts in velocity-space regions covered by null-measurement weight functions where the phase-space densities should be negligible [24].

Here we remedy such artifacts by analyzing the measurements in two stages. In the first stage we identify regions in velocity space where null measurements suggest phase-space densities below the detection limit of the diagnostic according to

$$\int w_0 F_{\text{true}} \, dE \, dp = S \leqslant \epsilon \tag{5}$$

where w_0 is the null-measurement weight function, F_{true} is the unknown true fast-ion velocity distribution function, S is the measured signal, and ϵ is the noise floor. As is customary, w_0 and F_{true} are given in energy-pitch coordinates (E, p), where E is the energy of the fast ions and p is the pitch defined as $p = \frac{v_\parallel}{v}$. v is the fast-ion speed, and $v_\parallel$ the velocity component anti-parallel to the magnetic field. The non-negativity of w_0 and F_{true} would allow us to conclude from an absolutely certain null-measurement (i.e. $\epsilon = 0$) that F_{true} must be zero in the regions covered by the null-measurement weight function. If noise is present, the strict argument does not hold. Nevertheless, we can still assert that the velocity space covered by the null-measurement weight function contains so few ions that they cannot be detected against the noise floor and thus neglect the small phase-space densities in these regions. This stage does not require the solution of an inverse problem and is thus very reliable if null-measurements can be told apart from measurements of small FIDA intensities, i.e. if the measurement uncertainties can be quantified reliably. In the second stage we solve the tomography problem neglecting the phase-space

densities in the regions covered by null-measurement weight functions:

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda L \end{pmatrix} F - \begin{pmatrix} S \\ 0 \end{pmatrix} \right\|_2 \quad \text{subject to } F^*(E_0, p_0) = 0.$$

(6)

E_0 and p_0 are the energies and pitches covered by the null-measurement weight functions, and $F^*(E_0, p_0)$ are the elements of the vector F^* that represent the null-measurement region. This condition determines an upper energy boundary of the target velocity space that we seek to reconstruct (figure 1). As weight functions are not bounded in energy, the target velocity space has also been judiciously restricted at some energy in previous work. A common choice is 10–50 keV above the injection energy of 60 keV or 90 keV as phase-space densities at larger energies are supposed to be small. Null-measurement weight functions allow us to determine this upper boundary in energies from the measurements in an optimal way that does not allow artifacts in the null-measurement region and that minimizes the number of unknowns in the inversion. This upper boundary is a strong function of pitch (figure 1). Neoclassical simulations by TRANSP [41] agree very well with the shape of the null-measurement velocity space (figure 2). Maxima of the velocity distribution function simulated by TRANSP appear at 60 keV and at 30 keV, corresponding to the full and half injection energy, whereas here no peak appears at 20 keV (one-third energy) as the TRANSP model removes thermalized particles in this region due to the high temperature.

The null-measurement idea is based on the presence of a sharp transition to zero in the fast-ion velocity distribution function. This is fulfilled in NBI discharges at the highest injection energy. In experiments with ion cyclotron resonance heating, the fast-ion distribution function is not expected to have a sharp transition within the energy range accessible to FIDA, such that the null-measurement technique cannot be used here. Nevertheless, we note that null measurements are often found in NES and GRS measurements in plasmas with third harmonic ICRH at JET as these diagnostics can easily detect MeV-range ions [14].

2.2. Non-negativity

One of the advantages of maximum entropy regularization is that it does not allow negative phase-space densities. The TSVD and variants of the Tikhonov regularization (one of which is inspired by Fisher information) do allow negative phase-space densities and in fact inversions based on those methods often contain regions with small negative phase-space densities. Negative phase-space densities have usually been ignored [11, 22–25, 28–31]. In a recent study such negative densities were strongly attenuated by a new approach [27]. It used fictitious measurements with pixel-sized weight functions covering the most negative values in the inversions. These fictitious measurements were given enough weight to iteratively force the phase-space densities


Figure 1. The coloured lines circumscribe the regions covered by null-measurement FIDA weight functions for ASDEX Upgrade discharge #31557 where the five colours represent the five views. These null-measurements suggest negligible fast-ion densities for energies above the black line to the right. The vertical black line at 10 keV is defined as the border between fast and thermal ions. The phase-space densities between the black lines are found by tomographic inversion.


Figure 2. Comparison of the target velocity space identified in figure 1 with a TRANSP simulation of the fast-ion velocity distribution function (10^{16} keV^{-1} m^{-3}). The simulation confirms the form of the null-measurement region with negligible phase-space densities. The colorbar is linear and includes positive phase-space densities as well as negative phase-space densities to allow comparisons with figures 4 and 5.

to negligible, yet still negative, values. Here we simply impose the constraint that the solution be non-negative and solve the minimization problem

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda L \end{pmatrix} F - \begin{pmatrix} S \\ 0 \end{pmatrix} \right\|_2 \quad \text{subject to } F^* \geqslant 0$$

(7)

using a standard non-negative least-squares algorithm [42]. The implementation as constraint has the advantage that no assumptions about the negative regions are needed. We will in the following also impose non-negativity and null-measurement constraints together:
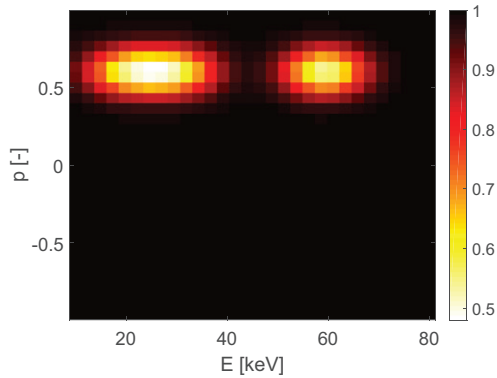
Figure 3. The known positions of the beam injection peaks are encoded by using a 2D regularization given by $\kappa(E, p)$ [-]. The total regularization strength is given by $\lambda\kappa(E, p)$. The half- and one-third energy injection peaks merge due to the low resolution.

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda L \end{pmatrix} F - \begin{pmatrix} S \\ 0 \end{pmatrix} \right\|_2 \quad \text{subject to} \quad \begin{cases} F^*(E_0, p_0) = 0 \\ F^* \geqslant 0 \end{cases}. \tag{8}$$

2.3. Positions of neutral beam injection peaks

Often the fast-ion distribution function has peaks at full-, half and one-third NBI energy which are difficult to reconstruct. We will show a reconstruction of all three NBI peaks based on TSVD exploiting null-measurements in section 4. First-order Tikhonov regularization, which penalizes steep gradients typical for NBI peaks, will tend to reduce the peaks. The positions of the NBI peaks, if present, are actually very well known from the geometry and energy of the beams. Three peaks are expected at $E = (20, 30, 60)$ keV and $p = 0.6$ for NBI Q3 in the plasma center. Here we encode the known geometry of the NBI by using a 2D function $\kappa(E, p)$ varying between 0.5 and 1 and solve the minimization problem

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda\kappa(E, p)L \end{pmatrix} F - \begin{pmatrix} S \\ 0 \end{pmatrix} \right\|_2 \quad \text{subject to} \quad \begin{cases} F^*(E_0, p_0) = 0 \\ F^* \geqslant 0 \end{cases}. \tag{9}$$

This reduces the penalty by a factor of about two at the position of the NBI sources, where we expect large gradients (figure 3), compared with the penalty elsewhere. In this way no prior knowledge of the size of the peak is used, and in fact even a negative peak would equally well be supported by $\kappa(E, p)$.

2.4. Numerical simulations as prior information

In previous studies of velocity-space tomography based on three to five FIDA views or a mix of FIDA and CTS views, it was possible to find useful and stable inversions and hence provide measurements of the fast-ion velocity distribution function and derived quantities such as the fast-ion density [26, 28]. However, if fewer views are available or if the signal-to-noise ratio is unfavourable, this is sometimes out of reach. Here we propose a new goal of the tomography approach for such highly photon-starved situations. Rather than inferring

the full 2D velocity distribution function, we only seek to infer the distribution of any discrepancy between a numerical simulation and the measurement in velocity space. Firstly, this is a simpler task that is possible with fewer views as we will show. Secondly, deviations from a numerical simulation is often what we are actually interested in. For example, the TRANSP code accounts for neoclassical transport and likely provides a good picture of the distribution function in MHD quiescent discharges. In MHD active discharges, we often observe discrepancies from such neoclassical simulations. The discrepancies are attributed to anomalous transport in addition to the neoclassical transport which still serves as a baseline. To locate the distribution of any disagreement in velocity space, we penalize differences from the simulated velocity distribution function F_{sim} and solve the minimization problem

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda\kappa L \end{pmatrix} F - \begin{pmatrix} S \\ \lambda\kappa L F_{\text{sim}} \end{pmatrix} \right\|_2 \quad \text{subject to} \quad \begin{cases} F^* \geqslant 0 \\ F^*(E_0, p_0) = 0 \end{cases}. \tag{10}$$

For very small λ's the inversion is underregularized and is dominated by random jitter as in the previously discussed Tikhonov problems. For very large λ's the solution is overregularized, but here the inversion then approaches the numerical simulation as equation (10) shows. The problem is now again to select λ assigning an appropriate balance between the measurements and the simulation which we will discuss in section 3. The solution F^* represents our estimate of the 2D velocity distribution function considering the simulation and the measurements. As our goal is to locate discrepancies between simulation and measurements in velocity space, we calculate

$$\Delta F^* = F^* - F_{\text{sim}} \tag{11}$$

which shows how the simulation should be modified according to the measurements. Examples will be shown in sections 5 and 6.

3. Standard inversions and the choice of the regularization parameter λ

The number of measurement data points and the number of unknowns in fast-ion velocity-space tomography is fairly small compared with many other tomography problems. Quick automatic inversion of FIDA spectra based on a library of approximate weight functions after each plasma discharge is therefore possible and will be implemented in future work. For that purpose the regularization parameter λ must be computed automatically from the data. In traditional tomography, two popular methods among many are the L-curve method [43, 44] and the generalized cross-validation (GCV) method [45]. The L-curve method has been applied to velocity-space tomography previously [28]. The GCV method seeks to minimize the prediction error by asserting that an arbitrary measurement S_i should be predicted well by the regularized solution based on the other measurements in S and that orthogonal transformations of S should not affect the choice of the

(a)

(b)

Figure 4. Unconstrained first-order Tikhonov tomographic inversions of FIDA measurements in five views in ASDEX Upgrade discharge #31557 in units (10^{16} keV^{-1} m^{-3}). The regularization parameter λ is chosen by two different methods: (*a*) L-curve. (*b*) GCV.

regularization parameter λ [45]. Recall that the Tikhonov solution $F^* = F^*(\lambda)$ is a function of the regularization parameter λ and that, for the unconstrained problem (equation (4)), we can write $F^*(\lambda) = W^{\#}(\lambda) S$ where the matrix $W^{\#}(\lambda)$ defines the solution. Then λ minimizes the GCV function

$$G(\lambda) = \frac{\|WF^*(\lambda) - S\|_2^2}{(\mathrm{trace}(I - WW^{\#}(\lambda)))^2}, \qquad (12)$$

where I is the identity matrix.

Figure 4 shows two unconstrained first-order Tikhonov inversions of FIDA measurements in ASDEX Upgrade discharge #31557. The level of regularization was respectively selected by the L-curve method and the GCV method. The discharge and the five-view FIDA measurements have been discussed previously [28]. The plasma was heated by NBI with the 2.5 MW source Q3 at 60 keV. The five-view FIDA measurements here and throughout the paper were made in the center of the plasma. Figure 2 shows a corresponding TRANSP simulation.

The unconstrained inversion with first-order Tikhonov regularization and L-curve regularization parameter selection in figure 4(*a*) shows the usual characteristics for this method. The inversion is smooth, and the overall shape, including the anisotropy and the location of the merged 20 keV and 30 keV NBI peaks, is close to our expectation. However, there are also limitations of the unconstrained inversion. First, the NBI peak at full injection energy (60 keV) does not appear. It is known that the sharp NBI peaks tend to be attenuated by first-order Tikhonov regularization [28]. Second, there are substantial phase-space densities at energies well in the null-measurement region. Here we plot up to 20 keV above the full injection energy. As already mentioned, TRANSP/NUBEAM predicts almost zero phase-space densities in the null-measurement region (figure 2), which further corroborates the hypothesis that the inferred phase-space densities are artifacts. Third, small patches of unphysical negative fast-ion phase-space densities appear, in this case in the null-measurement region. These three features of the inversions are very likely artifacts since they also appear erroneously in inversions of synthetic data where the true solution is known [22, 24, 28].

The GCV method tends to regularize less than the L-curve method (figure 4(*b*)) and roughly produces a regularization level as was sometimes judiciously selected [11, 24, 27]. The full energy beam injection peak (60 keV) does not appear but there is a clear ridge of large phase-space densities between the expected locations of the beam injection peaks. The ridge does not extend to energies larger than 60 keV as expected. The form of the ridge is consistent with the presence of fast ions that are slowing down due to collisions with electrons, i.e. they loose energy without significant pitch angle scattering. Nevertheless, as the critical energy is about 60 keV, collisions with thermal ions are also important leading to pitch-angle scattering as is apparent in figure 4. The unphysical negative regions are larger than for the L-curve method. A peak appears at $p = -0.5$ which is not expected according to the TRANSP simulation (figure 2) and is likely an artifact.

In velocity-space tomography, the regularization parameter has up to now been set either by judicious choice or by the L-curve method. The choice of regularization parameter or of a method to calculate it remains an open problem as one never knows *a priori* which method works best. Our criteria for good reconstruction of velocity-space distribution functions in NBI heated plasma are that there are no significant negative phase-space densities, that there are no significant phase-space densities far above the beam injection energies, that the distribution function is smooth and that the beam injection peaks are reconstructed well. Neither the L-curve nor the GCV method is consistently superior. Both methods are fairly robust but do not always work. The L-curve method can be applied irrespective of the incorporated prior information. However, the GCV method requires the existence of a regularized inverse matrix mapping the measurement space to the solution space, and so it is not directly applicable to the non-negative least-squares problem formulation or maximum entropy methods. Here we apply either method, depending on what works best for the given problem. For inversions not using a TRANSP simulation as prior information we will choose the regularization parameter λ by the GCV method applied to the unconstrained least-square problem (figure 4(*b*)), and then

Figure 5. First-order Tikhonov tomographic inversions based on five-view FIDA measurements in ASDEX Upgrade discharge #31557. The units are (10^{16} keV^{-1} m^{-3}). Various types of prior information are used. (*a*) Non-negativity: $F^* \geqslant 0$. (*b*) Null-measurement constraints: $F^*(E_0, p_0) = 0$. (*c*) $F^* \geqslant 0$ and $F^*(E_0, p_0) = 0$. (*d*) $F^* \geqslant 0$, $F^*(E_0, p_0) = 0$ and known peak locations: $\kappa(E, p)$.

solve the non-negative least-squares problem with that λ. However, we apply the L-curve method for the inversions using the TRANSP simulations as prior information.

4. High-definition inversions using prior information

In this section we show that prior information based on non-negativity, null measurements and the NBI geometry improves inversions substantially. The effect of prior information based on a numerical simulation is presented in section 5. The inversions in this section represent an estimate of the 2D fast-ion velocity distribution functions based on measurements alone. The inversions in section 5 represent an estimate based on measurements and simulations or, more interestingly, of how a given simulation needs to be modified to match best with the measurements.

Figure 5(*a*) shows a first-order Tikhonov inversion with non-negativity constraint according to equation (7). The non-negativity constraint improves the inversions without any visible disadvantages compared with the unconstrained inversions in figure 4. By virtue of the constraint, there are

no negative regions. Erroneous phase-space densities in the null-measurement region are reduced compared with the unconstrained inversions. The full energy beam injection peak is almost invisible, but there is a clear drop in phase-space densities above the full NBI energy. A possibly spurious peak at $p = -0.5$ is present as also observed in figure 4(*b*).

Figure 5(*b*) shows a first-order Tikhonov inversion with the null-measurement constraint according to equation (6). The inversion shows a clear peak at full injection energy and a clear ridge consistent with ions slowing down due to collisions with electrons. By virtue of the constraint, the phase-space densities in the null-measurement region are zero. Regions of erroneously negative phase-space densities are present at this regularization level. It is unclear if the two peaks appearing at negative pitches are real or if they are artifacts. The peak at $p = -0.5$ also appears as in figures 4(*b*) and 5(*a*).

Figure 5(*c*) shows a first-order Tikhonov inversion with non-negativity and null-measurement constraints according to equation (8). By construction there are no artifacts in the null-measurement region and there are no negative phase-space densities. The white patches of very low fast-ion density circumscribed by larger fast-ion densities seem physically unlikely and would suggest that slightly more regularization

Figure 6. First-order Tikhonov tomographic inversions using a simulation as prior information. The units are $(10^{16}\ \mathrm{keV}^{-1}\ \mathrm{m}^{-3})$. (*a*) F_{true} which is a modified TRANSP simulation with selective ejection of particles. (*b*) $\Delta F = F_{\mathrm{true}} - F_{\mathrm{sim}}$. (*c*) + (*d*) Inversion $\Delta F^* = F^* - F_{\mathrm{sim}}$ based on synthetic measurements of F_{true} with 5% Gaussian noise for a five-view FIDA system (*c*) without and (*d*) with the TRANSP simulation as prior information.

Figure 7. $\Delta F^* = F^* - F_{\mathrm{sim}}$ where F^* is calculated as in figure 6(*d*). In each subfigure the black dotted line shows the true location of the discrepancy. The units are $(10^{16}\ \mathrm{keV}^{-1}\ \mathrm{m}^{-3})$. The colorbar is different for each plot.

**Figure 8.** As figure 7, but using only two FIDA views. The angles of the lines-of-sight to the magnetic field are $\phi = [73°, 153°]$. The colorbar is different for each plot.


Figure 9. Inversions by TSVD with null-measurement constraint after NBI Q3 was turned on in discharge #33138 at $t = 0.4$ s. A movie is provided as supplementary material to this paper (movie 1) (stacks.iop.org/NF/56/106024/mmedia). NBI peaks at full, half, and one-third energy are gradually formed. The units are $(10^{16}\,\text{keV}^{-1}\,\text{m}^{-3})$. (*a*) 0.40 s, (*b*) 0.42 s, (*c*) 0.44 s, (*d*) 0.47 s, (*e*) 0.49 s, (*f*) 0.52 s.

would benefit the solution. The non-negativity constraint also tends to attenuate positive peaks because it reduces high-frequency components of the inversion and hence negative as well as positive extrema. The peak at full injection energy is present but weaker than in figure 5(*b*), and the ridge connecting the beam injection peaks is weaker than in figure 5(*b*). The two small peaks at negative pitches are attenuated compared with figure 5(*b*).

Finally, figure 5(*d*) shows a first-order Tikhonov inversion with non-negativity and null-measurement constraints and known NBI geometry according to equation (9). Two NBI peaks as well as the ridge connecting them appear. No clear artifacts, such as peaks or low-density patches at unexpected locations, appear in this inversion, and by construction there are again no artifacts in the null-measurement region and there are no negative phase-space densities. In the NBI peak

Figure 10. TRANSP simulation ($10^{16}\,\mathrm{keV}^{-1}\,\mathrm{m}^{-3}$) as in figure 9 after NBI Q3 was turned on in discharge #33138 at $t = 0.4$ s. NBI peaks at full, half, and one-third energy appear in the simulation. (*a*) 0.40 s, (*b*) 0.42 s, (*c*) 0.44 s, (*d*) 0.47 s, (*e*) 0.49 s, (*f*) 0.52 s.

region the regularization is comparatively weak so that high-definition features appear. In the regions far from the peaks the solution is smooth due to the stronger regularization. Overall, this inversion resembles the TRANSP simulation from figure 2 the most. This use of prior information will likely make inversions possible at many other tokamaks, where fewer fast-ion data than at ASDEX Upgrade is available.

5. Inversions using numerical simulations as prior information

Inversion techniques can be used to identify differences between theory and observation in velocity space. This could give clues on which physics is not adequately described in the simulation or which systematic error confounds the measurements. For this goal it is advantageous to use the numerical simulation as prior information. Figure 6 illustrates two different ways to locate differences between measurements and simulation in velocity space. Figure 6(*a*) shows a true distribution function F_{true} which we would like to know but which in an experiment is never known. Here we assume a modified TRANSP simulation. We have added the negative Gaussian function shown in figure 6(*b*) to the known TRANSP simulation F_{sim} (figure 2). This selective reduction in phase-space density is a toy model for an assumed anomalous transport phenomenon localized in velocity space which is not modelled in the TRANSP simulation F_{sim}. Hence figure 6(*b*) shows the difference

$$\Delta F_{\mathrm{true}} = F_{\mathrm{true}} - F_{\mathrm{sim}}. \tag{13}$$

The goal is now to reconstruct ΔF_{true}, given synthetic measurements of F_{true} in five FIDA views and F_{sim}. We add 5% Gaussian noise to each synthetic measurement which is


Figure 11. Measured (full lines) and simulated (dashed lines) increases in the fast-ion density for $p \in [0.3, 0.7]$ and five 10 keV wide energy intervals in ASDEX Upgrade discharge #33138 right after the NBI is switched on. The fast-ion densities are obtained by integration of the high-definition tomographic inversion in the specified regions in velocity space. The fast-ion densities in the regions containing NBI peaks grow faster.

a realistic noise level for FIDA measurements at ASDEX Upgrade. The discrepancy between measurement and simulation ΔF^* is calculated according to equation (11) without and with using the TRANSP simulation as prior information, i.e. respectively using equations (9) and (10). Figure 6(*c*) shows ΔF^* for the case without the TRANSP simulation as prior. In the region of the true difference, ΔF^* is negative and has approximately correct amplitudes. However, ΔF^* is dominated by large negative values at the beam injection peaks reflecting the difficulty in reconstructing the NBI peaks. The reconstructed difference ΔF^* therefore does not resemble the true difference ΔF_{true}. Figure 6(*d*) presents ΔF^* for the case where the TRANSP simulation has been used as prior information (equation (10)). In this case the approximate location of the discrepancy between the true distribution and

Figure 12. High-definition first-order Tikhonov tomographic inversions based on four-view FIDA measurements in ASDEX Upgrade discharge #32323. The units are (10^{16} keV^{-1} m^{-3}). 100 frames are supplied in a sawtooth crash movie attached as supplementary material to the paper (movie 2) (stacks.iop.org/NF/56/106024/mmedia). The crash occurs at $t = 2.23$ s. (*a*) 2.1 s. (*b*) 2.13 s. (*c*) 2.16 s. (*d*) 2.19 s. (*e*) 2.22 s. (*f*) 2.25 s.

the TRANSP simulation is found based on experimentally accessible quantities. However, artifacts are present and the amplitudes in the reconstruction are too low by about 30%. The success of this approach does not depend strongly on the position of the discrepancy in velocity space for the five-view FIDA system as we illustrate in figure 7. Each case shows the true location of the discrepancy and the reconstruction using the simulation as prior information. The approximate location of the discrepancy is in each case identified as the region with the lowest amplitudes.

Many FIDA systems have two or three viewing directions rather than the five available at ASDEX Upgrade. DIII-D has three FIDA views [46, 47] and MAST [48, 49], NSTX [50], EAST [51], and LHD have two views [52, 53]. In figure 8 we reconstruct the six cases from figure 7 using just two FIDA views to show that this approach could be useful for other machines. The approximate location of the ΔF_{true} is reconstructed in each case, even though the amplitudes of negative phase-space densities are not reconstructed as well as for the five-view case. The reconstruction of ΔF^* does not require as many measurements as the reconstruction of F^* since the simulation provides cogent prior information about the approximate basic shape. Hence this alternative approach to tomographic reconstruction should be highly useful for FIDA systems with few viewing directions as is common on many machines. Further, the approach should also work for many combinations of fast ion diagnostics installed on many machines [23, 29, 31], for example based on FIDA and CTS at LHD [54, 55], CTS, FIDA, NPA, NES, and GRS at ASDEX Upgrade [5, 7, 11, 56–62], NES, GRS and NPA at JET [14, 15, 63], CTS, GRS, NES, NPA and possibly fast-ion charge-exchange recombination spectroscopy at ITER [64–69]. We will demonstrate this method for the two-view FIDA case using experimental data in the next section.


Figure 13. The measured (full lines) and simulated (dashed lines) total fast-ion densities and the fast-ion densities in selected pitch ranges in ASDEX Upgrade discharge #32323.

6. Studies of neutral beam injection and sawtooth dynamics

Prior information allows us to study the formation and presence of NBI peaks at full, half, and one-third injection energy and the fast-ion dynamics associated with sawteeth. NBI peaks are here reconstructed by TSVD which is well-suited for fine-scale features at the expense of the appearance of some jitter. The target velocity-space is restricted using null-measurements. The beam positions are *not* used as prior information here, and neither is the TRANSP simulation. Negative phase-space densities are present in the TSVD, but they are fairly small and simply ignored as usual for TSVD. The TSVD study confirms the positions of the NBI peaks which we will use as prior information in the Tikhonov inversions of fast-ion dynamics in sawtoothing plasma.

In discharge #33138, FIDA measurements in five views were made just after switching on NBI Q3 to study the

Figure 14. The measured fast-ion densities in selected pitch and energy ranges in ASDEX Upgrade discharge #32323. The red and blue curves also appear in figure 13 and are here split into different energy intervals. (*a*) $0.25 < p < 0.75$. (*b*) $p > 0.75$.


Figure 15. Reconstruction of the velocity-space distribution of the differences between the measurements and the TRANSP simulation $\Delta F^* = F^* - F_{\mathrm{sim}}$ (10^{16} keV^{-1} m^{-3}) according to equation (10). Here the TRANSP simulations was used as prior information in the inversion. The crash occurs at $t = 2.23$ s. A movie is supplied as supplementary material to this paper (movie 3) (stacks.iop.org/NF/56/106024/mmedia). (*a*) 2.1 s. (*b*) 2.13 s. (*c*) 2.16 s. (*d*) 2.19 s. (*e*) 2.22 s. (*f*) 2.25 s.

appearance of the beam injection peaks in the plasma center. The time resolution of the measurements was 2.5 ms. Figure 9 shows six inversions after the NBI was switched on at $t = 0.4$ s. A movie showing the formation of the three NBI peaks in 50 frames is provided as supplementary material to this paper (movie 1) (stacks.iop.org/NF/56/106024/mmedia). The same level of regularization is used for all frames. The NBI peaks at full, half, and one-third injection energy appear reliably at the expected pitch and energies at 20 keV, 30 keV and 60 keV. Even though three peaks do not always appear in inversions, the well-understood locations and the large-scale coherence suggest that the peaks are not artifacts but are supported by the FIDA measurements. Figure 10 shows that a TRANSP simulation of the discharge also predicts the formation of three peaks in agreement with the measured result.

To study the increase of fast-ion density in time at selected positions in velocity space due to switching on the NBI, we integrate the inversion in $p \in [0.3, 0.7]$ and in five 10 keV wide energy intervals (figure 11). The fast-ion density builds up quickly and steadily for the intervals containing the NBI peaks at 20, 30 and 60 keV whereas the intermediate intervals from 32 to 52 keV increase more slowly in agreement with the TRANSP simulation. Simulations and measurements agree very well for the intervals containing the 20 keV and 30 keV NBI peaks whereas the 60 keV interval is predicted to be slightly less populated. The TRANSP simulation shows larger fast-ion densities at 42–52 keV than at 32–42 keV as these intervals are fuelled mostly by slowing down of 60 keV ions and there is some pitch angle scattering. The inversion does not show this behaviour. It should be noted that TSVD

Figure 16. As for figure 15, but using only two FIDA views. The angles of the lines-of-sight to the magnetic field are $\phi = [14°, 73°]$. (*a*) 2.1 s. (*b*) 2.13 s. (*c*) 2.16 s. (*d*) 2.19 s. (*e*) 2.22 s. (*f*) 2.25 s.

showing NBI peaks typically also have undulations elsewhere, which may bias a comparison between the different energy curves. All in all, the TRANSP simulation agrees fairly well with the measurements, and both show the formation of three NBI peaks at the expected locations.

Now we turn to study fast-ion dynamics in the sawtoothing plasma by high-definition first-order Tikhonov inversions using all methods discussed in section 2. First we will use null-measurements, non-negativity, and the NBI peak positions, but *not* the TRANSP simulation, in pure inversions of the data. Then we will use the TRANSP simulation as prior information and locate the distribution of discrepancies between the measurements and the TRANSP simulation. Previous inversions focussed on the redistribution due to the crash [26–31] whereas we here study the dynamics in the entire sawtooth period. In this discharge only four views had a time resolution of 2.5 ms whereas the fifth view had a time resolution of 25 ms as the camera did not allow faster data acquisition in this discharge, and hence we only use four views here. Figure 12 shows first-order Tikhonov inversions with prior information from non-negativity, null-measurements and the position of the NBI peaks but *not* from any simulation. The same level of regularization is used for all frames. Figure 12 covers one sawtooth cycle, and figures 12(*a*) and (*f*) are right after crashes. A movie showing 100 frames covering two sawtooth cycles is supplied as supplementary material to this paper (movie 2) (stacks.iop.org/NF/56/106024/mmedia). The 2D velocity distribution function is strongly depleted at the sawtooth crashes and builds up due to the continuous fuelling by the NBI particle sources at 20 keV, 30 keV and 60 keV between the crashes. The 60 keV peak almost disappears right after the crashes. It gradually builds up as the fast-ion density increases during the sawtooth cycle. The ridge between the beam injection peaks also becomes stronger during the sawtooth cycle.

Figure 13 compares measured total fast-ion densities as well as fast-ion densities in selected pitch intervals with a TRANSP simulation. Here we have added drifting Maxwellian distribution functions to the fast-ion TRANSP simulation as one always measures the complete ion velocity distribution function. The tails of the Maxwellian here make a significant contribution due to the very high ion temperatures of up to 10 keV with sawtoothing time traces. The fast ion densities are obtained by integration of the inversions over all energies and over pitches in the selected intervals. The measured and simulated total fast-ion densities are similar right after the crashes. However, the measured total fast-ion density increases less than the simulated one during the sawtooth cycle. For $p > 0.75$ the measurement and the TRANSP simulation show clear sawteeth in good agreement though the TRANSP simulation is consistently lower than the FIDA measurement. For the intermediate pitch range $0.25 < p < 0.75$ the agreement is also good, but here the TRANSP simulation is consistently higher than the FIDA measurement. For $|p| < 0.25$ the measurements show no evidence of sawteeth, whereas the simulation shows clear sawteeth. We stress that variations in measured signals are much more reliable than their absolute values, and hence we are confident in the presence or absence of sawteeth in the time traces. The previous velocity-space tomography studies of sawtooth crashes have consistently found that fast ions with pitches close to zero are much less affected by the crash than ions with pitches close to one which is consistent with our results [26–31]. Figure 14 subdivides the $p > 0.75$ interval as well as the $0.25 < p < 0.75$ interval into three energy intervals. Sawtoothing time traces are evident at all energies, and we find no strong selection in energy.

In figure 15 we examine the results from figures 13 and 14 by reconstructing the difference $\Delta F^* = F^* - F_{\rm sim}$ between the

13

TRANSP simulation and the high-definition inversions using the TRANSP simulation as prior information. As above we have added a drifting Maxwellian to the TRANSP simulation to account for the thermal ions. We use the available four FIDA views. A movie covering two sawtooth cycles is supplied as supplementary material to this paper (movie 3) (stacks.iop. org/NF/56/106024/mmedia). The results obtained by using and not using the TRANSP simulation as prior information are consistent. The inversions and the simulations agree relatively well right after the crashes. The differences then grow during the sawtooth cycle in the region where the TRANSP simulation underpredicts the measurements ($p > 0.75$) as well as in the region where the TRANSP simulation over-predicts the measurements ($p < 0.75$). Finally, in figure 16 we repeat the tomographic reconstructions of the difference $\Delta F^* = F^* - F_{\text{sim}}$ from figure 15, but only using two FIDA views as available on many other machines. Salient large-scale features such as the regions of overprediction and under-prediction as well as the approximate location of the extrema are well reconstructed using only two views. However, the detailed shape of ΔF^* is somewhat different, and there is a tendency to find larger discrepancies.

7. Conclusions

Velocity-space tomography is usually a highly photon-starved enterprise since the optical access to tokamaks and hence the number of simultaneous measurements is limited and since the signal-to-noise ratio of the measurements is often low compared with many other tomography applications. At the same time, we would like to infer the 2D velocity distribution function in high resolution and hence need to infer as many unknown parameters as the measurement data support. Here we make up for the lack of measurements by using additional prior information. The inversions are substantially improved by using three types of prior information: 1. The non-negativity of phase-space densities. 2. The measured absence of FIDA light restricting the target velocity space. 3. The velocity-space position of NBI peaks.

An inversion based on truncated singular value decomposition using null-measurements but not the other types of prior information reconstructs the three NBI peaks at full, half and one-third energy at the expected locations, and their appearance after switching on an NBI could be studied resolved in time. We could further study the dynamics of the fast-ion velocity distribution function in a sawtoothing plasma resolved in time in unprecedented detail. TRANSP simulations underpredict the measurements for $p > 0.75$ and over-predict the measurements for $p < 0.75$. Measurements show no evidence of sawteeth for $p < 0.25$ in disagreement with the TRANSP simulation. The time-resolved tomographic inversion movies efficiently summarize up to 50 000 data points using all presented prior information.

Lastly, we demonstrate an alternative approach to velocity-space tomography. We infer the 2D fast-ion distribution function considering the measurements as well as a simulation. In this case tomographic inversion uses the simulation as prior information. If the measurements and the simulation are inconsistent, the most likely velocity-space distribution of the discrepancies can then be found by subtracting the simulation from the inversion with the simulation as prior information. This could not be achieved with the pure velocity-space tomography approach, which does not use the simulation, even with five FIDA views. The experimental results obtained with this alternative approach are consistent with results obtained by velocity-space tomography that does not use the simulation as prior information. Experimental results obtained with this approach using only two FIDA views suggest that velocity-space tomography methods can be applied to two-view FIDA systems which are common on many machines.

References

[1] Grierson B.A. *et al* 2012 *Phys. Plasmas* **19** 056107
[2] Grierson B.A. *et al* 2012 *Rev. Sci. Instrum.* **83** 10D529
[3] Stejner M. *et al* 2013 *Plasma Phys. Control. Fusion* **55** 085002
[4] Stejner M. *et al* 2015 *Plasma Phys. Control. Fusion* **57** 062001
[5] Salewski M. *et al* 2010 *Nucl. Fusion* **50** 035012
[6] Moseev D. *et al* 2011 *Plasma Phys. Control. Fusion* **53** 105004
[7] Nielsen S.K. *et al* 2015 *Plasma Phys. Control. Fusion* **57** 035009
[8] Rasmussen J. *et al* 2015 *Plasma Phys. Control. Fusion* **57** 075014
[9] Heidbrink W.W. *et al* 2004 *Plasma Phys. Control. Fusion* **46** 1855–75
[10] Heidbrink W.W. *et al* 2007 *Plasma Phys. Control. Fusion* **49** 1457–75
[11] Geiger B. *et al* 2015 *Plasma Phys. Control. Fusion* **57** 014018
[12] Hellesen C. *et al* 2010 *Nucl. Fusion* **50** 084006
[13] Hellesen C. *et al* 2013 *Nucl. Fusion* **53** 113009
[14] Eriksson J. *et al* 2015 *Nucl. Fusion* **55** 123026
[15] Schneider M. *et al* 2016 Modelling third harmonic ion cyclotron acceleration of deuterium beams for JET fusion product studies experiments *Nucl. Fusion* submitted
[16] Kiptily V.G. *et al* 2010 *Nucl. Fusion* **50** 084001
[17] Tardocchi M. *et al* 2011 *Phys. Rev. Lett.* **107** 205002
[18] Nocente M. *et al* 2012 *Nucl. Fusion* **52** 063009
[19] Nocente M. *et al* 2015 *Nucl. Fusion* **55** 123009
[20] Egedal J. and Bindslev H. 2004 *Phys. Plasmas* **11** 2191
[21] Salewski M. *et al* 2011 *Nucl. Fusion* **51** 083014
[22] Salewski M. *et al* 2012 *Nucl. Fusion* **52** 103008
[23] Salewski M. *et al* 2013 *Nucl. Fusion* **53** 063019
[24] Salewski M. *et al* 2014 *Nucl. Fusion* **54** 023005
[25] Salewski M. *et al* 2015 *Plasma Phys. Control. Fusion* **57** 014021
[26] Geiger B. *et al* 2015 *Nucl. Fusion* **55** 083001

[27] Weiland M. *et al* 2016 *Plasma Phys. Control. Fusion* **58** 025012
[28] Jacobsen A.S. *et al* 2016 *Plasma Phys. Control. Fusion* **58** 045016
[29] Jacobsen A.S. *et al* 2016 *Plasma Phys. Control. Fusion* **58** 042002
[30] Jaulmes F. *et al* 2016 *Nucl. Fusion* **56** 112012
[31] Rasmussen J. *et al* 2016 *Nucl. Fusion* **56** 112014
[32] Salewski M. *et al* 2014 *Plasma Phys. Control. Fusion* **56** 105005
[33] Pace D.C. *et al* 2012 *Rev. Sci. Instrum.* **83** 073501
[34] Jacobsen A.S. *et al* 2014 *Rev. Sci. Instrum.* **85** 11E103
[35] Jacobsen A.S. *et al* 2015 *Nucl. Fusion* **55** 053013
[36] Salewski M. *et al* 2015 *Nucl. Fusion* **55** 093029
[37] Salewski M. *et al* 2016 *Nucl. Fusion* **56** 046009
[38] Stroth U. *et al* 2013 *Nucl. Fusion* **53** 104003
[39] Anton M. *et al* 1996 *Plasma Phys. Control. Fusion* **38** 1849–78
[40] Craciunescu T. *et al* 2009 *Nucl. Instrum. Methods Phys. Res.* A **605** 374–83
[41] Pankin A. *et al* 2004 *Comput. Phys. Commun.* **159** 157–84
[42] Lawson C. and Hanson R. 1974 *Solving Least Squares Problems* (Englewood Cliffs, NJ: Prentice-Hall)
[43] Hansen P.C. 1992 *SIAM Rev.* **34** 561–80
[44] Hansen P.C. and O'Leary D.P. 1993 *SIAM J. Sci. Comput.* **14** 1487–503
[45] Golub G., Heath M. and Wahba G. 1979 *Technometrics* **21** 215–23
[46] Heidbrink W.W. *et al* 2012 *Nucl. Fusion* **52** 094005

[47] Pace D.C. *et al* 2011 *Plasma Phys. Control. Fusion* **53** 062001
[48] Michael C.A. *et al* 2013 *Plasma Phys. Control. Fusion* **55** 095007
[49] Jones O.M. *et al* 2015 *Plasma Phys. Control. Fusion* **57** 125009
[50] Bortolon A. *et al* 2010 *Rev. Sci. Instrum.* **81** 10D728
[51] Huang J. *et al* 2014 *Rev. Sci. Instrum.* **85** 11E407
[52] Ito T. *et al* 2010 *Rev. Sci. Instrum.* **81** 10D327
[53] Ito T. *et al* 2012 *Plasma Fusion Res.* **5** S2099
[54] Kubo S. *et al* 2010 *Rev. Sci. Instrum.* **81** 10D535
[55] Nishiura M. *et al* 2014 *Nucl. Fusion* **54** 023006
[56] Meo F. *et al* 2008 *Rev. Sci. Instrum.* **79** 10E501
[57] Meo F. *et al* 2010 *J. Phys.: Conf. Ser.* **227** 012010
[58] Furtula V. *et al* 2012 *Rev. Sci. Instrum.* **83** 013507
[59] Schneider P.A. *et al* 2015 *Rev. Sci. Instrum.* **86** 073508
[60] Tardini G. *et al* 2012 *J. Instrum.* **7** C03004
[61] Tardini G. *et al* 2013 *Nucl. Fusion* **53** 063027
[62] Nocente M. *et al* 2012 *Nucl. Fusion* **52** 094021
[63] Romanelli F. 2015 *Nucl. Fusion* **55** 104001
[64] Chugunov I.N. *et al* 2011 *Nucl. Fusion* **51** 083010
[65] Donné A.J.H. *et al* 2007 Progress in the ITER Physics Basis Chapter 7: diagnostics *Nucl. Fusion* **47** S337–84
[66] Salewski M. *et al* 2008 *Rev. Sci. Instrum.* **79** 10E729
[67] Salewski M. *et al* 2009 *Plasma Phys. Control. Fusion* **51** 035006
[68] Salewski M. *et al* 2009 *Nucl. Fusion* **49** 025006
[69] Kappatou A. *et al* 2012 *Nucl. Fusion* **52** 043007