Journal of Physics: Conference Series **775** (2016) 012009 doi:10.1088/1742-6596/775/1/012009

Uncertainty estimation and a stopping rule in nonlinear gyrokinetic simulations

Michael Oberparleiter, Hans Nordman

Department of Earth and Space Sciences, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden

E-mail: michael.oberparleiter@chalmers.se

Geert Verdoolaege

Research Unit Nuclear Fusion - Department of Applied Physics, Ghent University, Sint-Pietersnieuwstraat 41, B-9000 Ghent, Belgium

Frank Jenko

Department of Physics and Astronomy, University of California, Los Angeles, CA 90095, USA

Abstract. We present a method to estimate the mean and uncertainty of fluctuating quantities, such as spatially averaged density and temperature fluctuations or radial fluxes, from initial value simulations of the Eulerian gyrokinetic code GENE[1, 2]. Since the time series are autocorrelated in time, the data is grouped into batches based on the autocorrelation time and their means form the sample for further statistical treatment, such as calculating the standard error of the mean.Based on this uncertainty estimate we develop a stopping rule for a nonlinear simulation: First, regression tests ensure that it has reached a stationary (quasi-steady) state and data before this point is discarded. Then the previously described estimate is calculated. If the estimated relative error is below a prescribed threshold, the simulation is stopped.

This scheme is applied to several previously performed GENE simulations ranging from simple benchmarks to modelling of JET and ASDEX discharges. It can be demonstrated that a number of simulations could be around 30% shorter if a maximal statistical relative uncertainty of 5% is desired for all monitored quantities.

1. Introduction

Most modern-day investigations of turbulence in tokamak or stellarator plasmas are performed with massively parallel simulation codes which solve the gyrokinetic Vlasov-Poisson system[3]. This can be done by tracking particles in a Lagrangian framework (particle-in-cell, PIC[4, 5]) or by solving the gyrokinetic equations on a fixed grid (Eulerian approach [2, 6–8]. A semi-Lagrangian approach is also successfully used [9].

Although the availability of computing resources has grown significantly over the last years, nonlinear gyrokinetic simulations of turbulence remain an expensive undertaking and an efficient usage of the available core hours is important. In addition, it is preferable to use the additional computational resources to include more physical phenomena in the simulation model and make

Content from this work may be used under the terms of the Creative Commons Attribution 3.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.
Published under licence by IOP Publishing Ltd

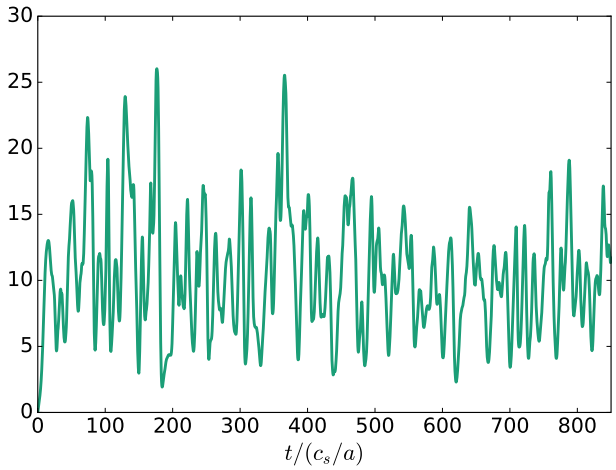

Figure 1: Basic example of a simulation time trace (shown is the flux-surface averaged radial ion heat flux).

it more realistic. Hence, we wish to find a criterion when a simulation can be automatically terminated because it provides results with a sufficient certainty, i.e. with error bars of the size we want.

This immediately raises the question how the (un-)certainty of a simulation result can be estimated. The nonlinear gyrokinetic equation is generally solved numerically as an initial value problem, i.e. simulations provide a time series of the 5D gyrocenter distribution function f and its velocity space moments, which often are physically observable quantities such as density and temperature fluctuations or radial transport fluxes. By taking additional spatial averages they can be reduced to one-dimensional time series which are the most convenient to monitor the state of the simulation.

In Section 2 we discuss one method to find a reliable estimate for the mean and its error of such a time series. Afterwards, in Section 3 we demonstrate how this estimate can be used to establish a stopping rule, i.e. a convergence criterion, for nonlinear simulations.

2. Estimating the mean and uncertainty of fluctuating quantities

In Fig. 1 we show a basic example of a simulation time trace from GENE. Although the figure shows a solid line and the underlying equation systems operate on continuous functions, it is important to remember that the numerical solution and hence our time series are discrete data. In order to apply statistical methods two features are important: First, the time trace contains an initial non-stationary phase ($t \lesssim 50$) which we disregard for the moment and will discuss in Section 3. Secondly, two nearby points in time are not statistically independent, i.e. the data is *autocorrelated*.

We can quantify this by calculating the autocorrelation time t_c, i.e. the time until the normalized discrete autocorrelation function $R(t_i)$ of the time series Q_n has declined to 1/e:

$$R(t_i) = \sum_{j=0}^{N-j} q_i q_{i+j} \quad \text{with } q_i = Q_i - \langle Q \rangle, \tag{1}$$

$$R(t_c) = R(t_0)/\mathrm{e}, \tag{2}$$

where $\langle Q \rangle$ is the average over the entire time trace between the start t_0 and end time t_N. As the samples are not necessarily equidistant in time, we perform averages with the composite

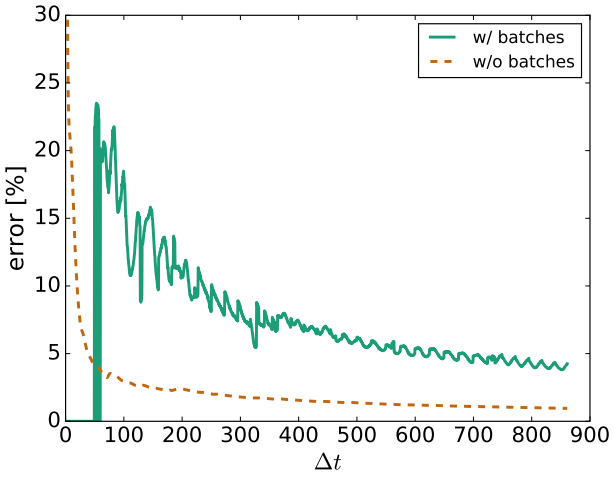

Figure 2: Relative SEM with and without batch mean procedure in dependence of the time interval considered from the time series in Fig. 1.

trapezoidal rule for integration:

$$\langle Q \rangle = \frac{1}{2(t_N - t_0)} \sum_{i=0}^{N-1} (Q_{i+1} + Q_i)(t_{i+1} - t_i)\,. \tag{3}$$

We now use the estimate of the autocorrelation time to construct an uncorrelated set of data: The time interval is separated into N_{win} non-overlapping windows of length $5t_c$ and the data in each window forms a so-called *batch*[10]. This method is also applied, for example, in the field of Markov chain Monte Carlo (MCMC) simulations and has shown a number of good mathematical properties there[11]. The factor 5 ensures that estimation errors of t_c do not lead to remaining autocorrelations. The means $\overline{Q}_n$ of these batches then form the desired set of uncorrelated measurements. The grand mean $\langle Q \rangle$ remains unaffected by this due to its linearity in Q_i and we consider it as the estimate for the average value of the quantity.It is now also straightforward to estimate an uncertainty for said value from the batch means: The standard error of the mean (SEM) from the batch means is defined as:

$$s_{\mathrm{batch}} = \sqrt{\frac{1}{N_{\mathrm{win}}\,(N_{\mathrm{win}} - 1)} \sum_{n=1}^{N_{\mathrm{win}}} \left(\overline{Q}_n - \langle Q \rangle\right)^2} \tag{4}$$

When we apply this procedure to the data of Fig. 1 and look the relative error $s_{\mathrm{batch}}/\langle Q \rangle$ as a function of the length of the time interval considered (Fig. 2), the SEM shows the expected scaling with $N_{\mathrm{win}}^{-1/2}$, i.e. the error becomes smaller the longer the simulation runs. This sets it apart from, for example, the standard deviation of the data which is sometimes used. As an artefact of the batch process additional small oscillations occur but at an amplitude that is not problematic for relative SEM below 10%. Comparison with a naive calculation of the SEM from the raw data on the other hand demonstrates that ignoring the temporal autocorrelation would lead to a massive underestimate of the uncertainty by a factor of 3-4.

Thus, the relative standard error of the mean provides a good quantity to set a threshold value for stopping the simulation. Typically, a threshold of 5% or 10% is a good choice.

3. Regression testing and the algorithm for the stopping rule

With the uncertainty estimator for autocorrelated time series established the remaining condition that should be ensured is *stationarity*. Generally, it can be violated in two ways for gyrokinetic

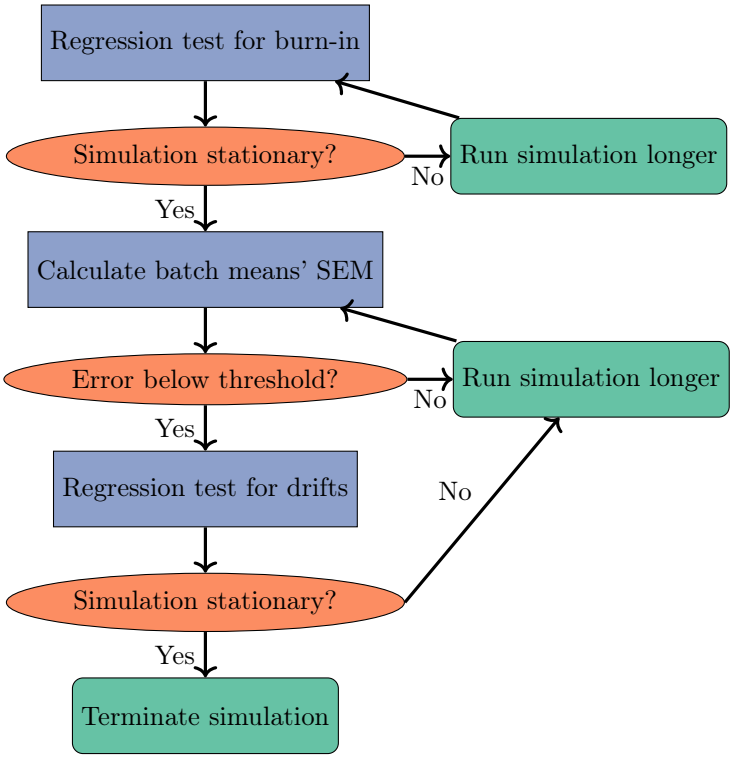

Figure 3: Algorithm for the stopping rule

simulation data: At the start of the simulation the system needs to transition from the initial condition to the quasi-steady state. Typically this occurs as a phase of exponential growth until the nonlinear saturation mechanisms set in. Strictly speaking, it is not necessary for the batch mean procedure to cut off this phase, i.e. allow a "burn-in"[12]. Not doing it, however, significantly increases the required simulation time until the threshold for the SEM is met, as it creates a group of "outliers" in the statistic. The second form of non-stationarity are slow drifts, which can occur due to a number of physical and numerical reasons. In this case, it is preferable to let the code user decide if they consider their simulation run finished.

Both of these issues can be addressed with regression testing. For the burn-in, the last[1] $t_1 - t_0$ time units ($100c_s/a$ in our case) of the simulation are taken and an ordinary least squares linear regression is performed, resulting in the linear equation

$$Q(t) = \alpha(t - t_0) + \beta. \tag{5}$$

If $|\alpha/(t_1 - t_0)| > 0.1\langle Q\rangle$, i.e. the regression line's slope is larger than 0.1 of the data mean in the regression interval or, in case the mean is smaller than the data's standard deviation σ, $|\alpha/(t_1 - t_0)| > 1.5\sigma$, the test fails. As soon as the test succeeds for the first time, the start of the 100 time unit interval is set as the starting point for the procedure of Section 2.

When the threshold for the uncertainty estimate is reached, the regression test is applied again. We, however, perform the regression over the entire time series (without the burn-in) and allow for the slope to be as large as 0.2 of the mean. If this test passes, the simulation can be stopped. The overview over the entire algorithm is shown in Fig. 3.

[1] This is from the perspective of a running simulation. A posteriori, this analysis can be done by taking an initial interval of data and extending it every time the test fails.

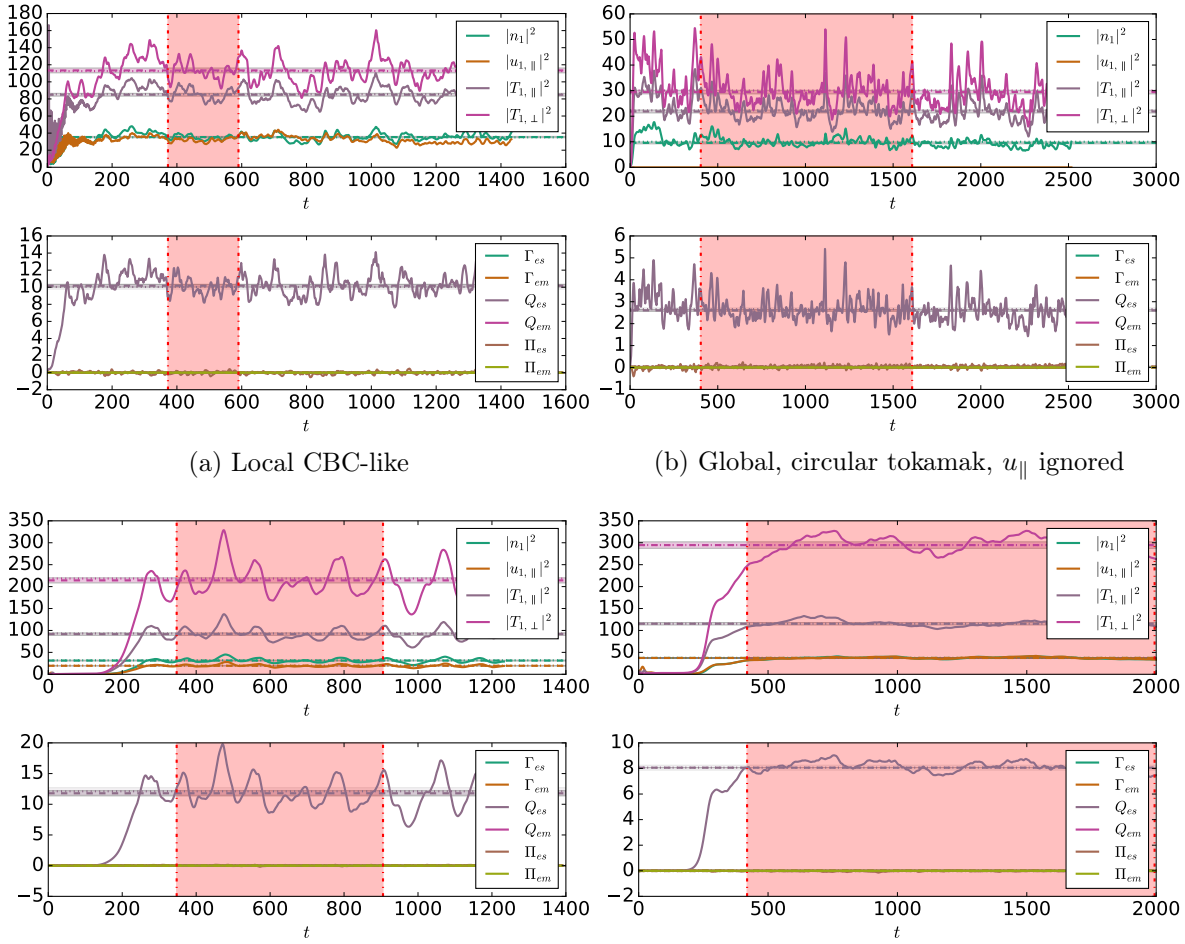

(a) Local CBC-like

(b) Global, circular tokamak, $u_\parallel$ ignored

(c) Full-surface NCSX geometry

(d) Full surface W7X geometry

Figure 4: ITG turbulence simulations with adiabatic electrons

4. Applications to simulation data

The straightforward way to test this algorithm is to a posteriori apply it to a diverse set of already finished simulation runs. In the following we present time traces for the spatially averaged fluctuations of density $|n_1|^2$, parallel flow $\left|u_{1\parallel}\right|^2$, temperature $\left|T_{1\parallel}\right|^2$ and $|T_{1\perp}|^2$, as well as the flux surface averaged flows of particles Γ, energy Q and momentum Π. If a quantity is numerically zero, it is ignored in the analysis. The shaded time intervals show the result of the stopping rule algorithm with an uncertainty threshold of 5% for all quantities. The dashed horizontal lines mark the average of a quantity based on the full time series excluding the burn-in, while the dotted lines show the average based on the stopping rule algorithm.

Fig. 4a shows the most basic type of simulated tokamak turbulence: A local (flux-tube) simulation of a circular tokamak with Cyclone Base Case-like[13] parameters with electrons considered as adiabatic. As can be seen, the stopping rule would terminate the simulation after a surprisingly short time (40% of the total time) with a maximal uncertainty of 3.6% and the mean values agreeing well with their reference from the full time series.

In Fig. 4b the simulation is global, i.e. it takes the radial profiles of density and temperature into account. Otherwise it is similar to the previous case. For the analysis we suppress $u_\parallel$, as it has a long time drift that is deemed irrelevant. The error bar reaches its desired value of 4.9% after 65% of the full time series. Since the cost for a global simulation is about an order

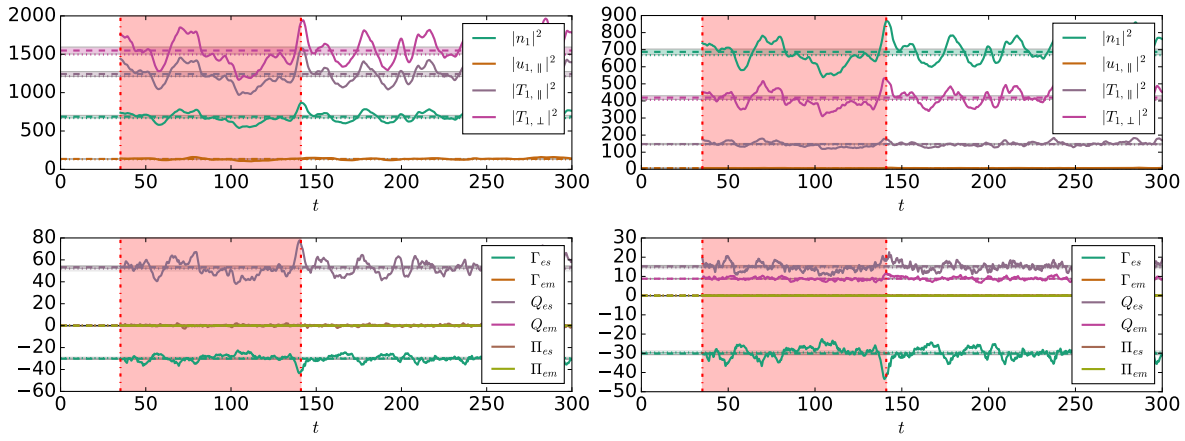

(a) Ions (b) Electrons

Figure 5: Local ITG/TEM turbulence simulation with electromagnetic effects of a JET discharge

of magnitude higher than for a flux-tube, this would imply a greater saving than for the case of Fig. 4a.

Fig. 4c and 4d show time series of full-flux-surface simulations of stellarator geometries. Already from visual inspection it becomes clear that their autocorrelation time is notably longer than previously. For the case of Fig. 4c still an uncertainty of 4.8% is reached after 75% of the simulation time. In the Wendelstein-7X case the final drift regression test fails and the batch means procedure results in an uncertainty of 2.5% for the full time trace.

In all cases presented so far, the regression testing to find the burn-in phase performed completely satisfactory. In the case of a simulation continued from a checkpoint as in Fig. 5, it can, however, cause too much data to be neglected. Hence, we do not apply it in this simulation of a JET discharge. The remaining part of the stopping rule performs well for this physically quite comprehensive modelling and a relative uncertainty of 5.0% is reached after 45% of the used simulation time. It should be noted at this point that for all time series presented so far the average values of the quantities based on the stopping rule agreed with their counterparts based on the full time trace within their uncertainty.

Lastly, we show a simulation of an ASDEX-Upgrade discharge (Fig. 6) where the algorithm does not converge and also does not give a reasonable error estimate. Due to a source of external $\vec{E} \times \vec{B}$ shear switched on at $t = 200$ the time trace is not stationary until fairly late and the remaining data is insufficient for the batch means procedure. In addition, the stationarity of the entire series is not easily to establish. Thus, we deem this one of the scenarios where manual control and evaluation by the code user is the better option.

5. Conclusions

In this article an algorithm was developed to automatically terminate a nonlinear gyrokinetic simulation when it gives results with a predetermined uncertainty, i.e. a previously chosen error bar. For this, the standard error of the mean calculated from (uncorrelated) batch means provides a statistically sound estimate of the uncertainty. It can be combined with linear regression tests for better convergence properties and more reliable average values derived from the simulation. It could be shown in an a posteriori analysis of a number of simulation runs from GENE that significant savings in simulation time are possible with this algorithm. There remain, however, a number of scenarios which require further fine tuning of the method and its parameters.

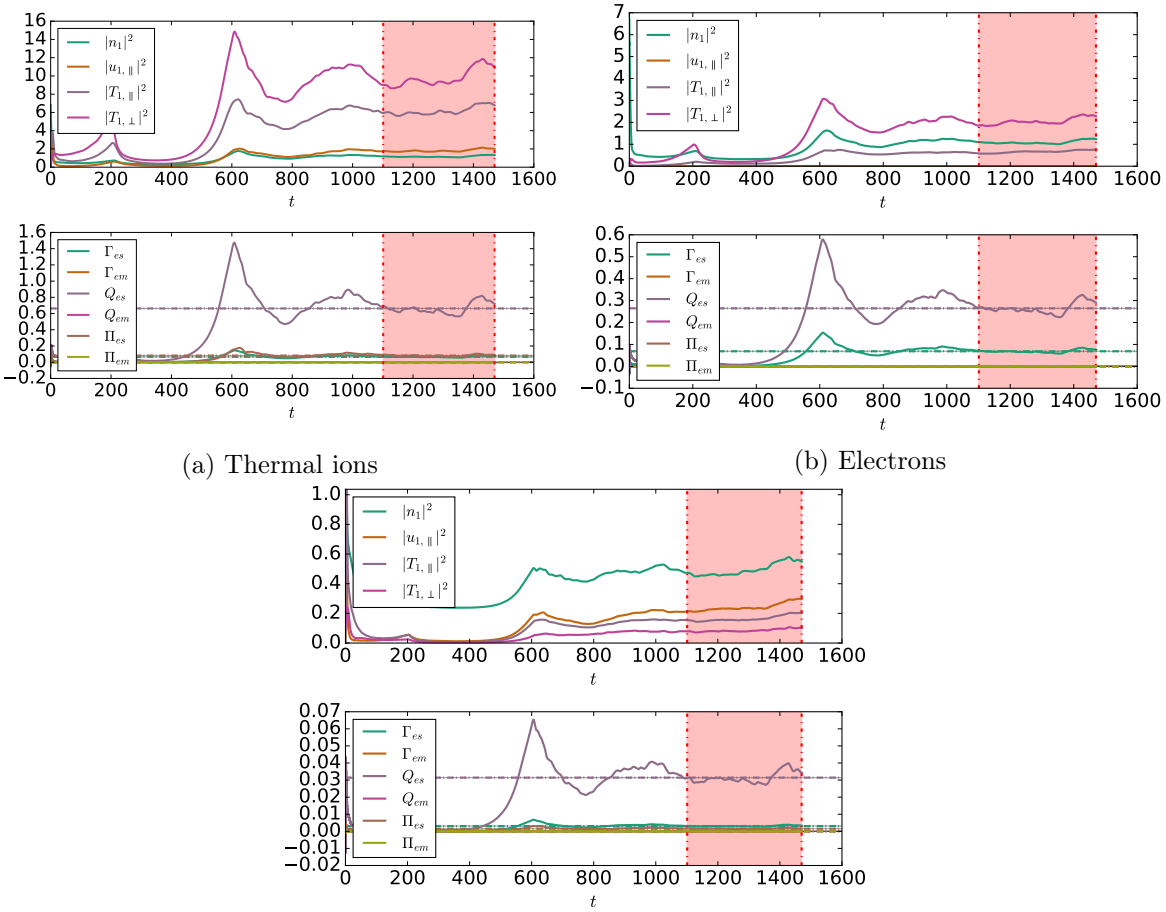

(a) Thermal ions

(b) Electrons

(c) Fast ions

Figure 6: Local ITG/TEM turbulence simulation with a fast particle species and external $\vec{E} \times \vec{B}$ shear for an AUG discharge

6. Acknowledgements

The authors would like to thank A. Bustos, T.Görler, D.Told, M.J.Pueschl, H.Doerk and D.Tegnered for providing significant parts of the simulation data analysed in this work.

References

[1] Jenko F, Dorland W, Kotschenreuther M and Rogers B N 2000 *Phys. Plasmas* **7** 1904–1910

[2] Görler T, Lapillonne X, Brunner S, Dannert T, Jenko F, Merz F and Told D 2011 *J. Comput. Phys.* **230** 7053 – 7071 URL `http://genecode.org`

[3] Brizard A J and Hahm T S 2007 *Reviews of Modern Physics* **79**(2) 421–468

[4] Kornilov V, Kleiber R, Hatzky R, Villard L and Jost G 2004 *Phys Plasmas* **11** 3196

[5] Jolliet S, Bottino A, Angelino P, Hatzky R, Tran T, Mcmillan B, Sauter O, Appert K, Idomura Y and Villard L 2007 *Comput. Phys. Commun.* **177** 409 – 425 ISSN 0010-4655

[6] Jenko F, Dorland W and Hammett G W 2001 *Phys. Plasmas* **8** 4096–4104

[7] Peeters A, Camenen Y, Casson F, Hornsby W, Snodin A, Strintzi D and Szepesi G 2009 *Comput. Phys. Commun.* **180** 2650–2672

[8] Idomura Y, Urano H, Aiba N and Tokuda S 2009 *Nucl. Fusion* **49** 065029

[9] Grandgirard V, Brunetti M, Bertrand P, Besse N, Garbet X, Ghendrih P, Manfredi G, Sarazin Y, Sauter O, Sonnendrücker E, Vaclavik J and Villard L 2006 *Journal of Computational Physics* **217** 395 – 423 ISSN 0021-9991

[10] Law A M and Kelton W D 1991 *Simulation Modeling and Analysis* 2nd ed (McGraw-Hill)

[11] Flegal J M and Jones G L 2010 *Ann. Statist.* **38** 1034–1070

[12] Jones G L, Haran M, Caffo B S and Neath R 2006 *J. Amer. Statist. Assoc.* **101** 1537–1547

[13] Dimits A M, Bateman G, Beer M A, Cohen B I, Dorland W, Hammett G W, Kim C, Kinsey J E, Kotschenreuther M, Kritz A H, Lao L L, Mandrekas J, Nevins W M, Parker S E, Redd A J, Shumaker D E, Sydora R and Weiland J 2000 *Phys. Plasmas* **7** 969–983